Bernoulli percolation on the Random Geometric Graph

Lyuben Lichev, Bas Lodewijks, Dieter Mitsche,[*] and Bruno Schapira

Univ. Jean Monnet, Saint-Etienne, Univ. Claude Bernard, Lyon, and Univ. Aix-Marseille, France

May 23, 2022

Abstract

Given $\lambda > 0$, $p \in [0, 1]$ and a Poisson Point Process $\text{Po}(\lambda)$ in $\mathbb{R}^2$ with intensity λ, we consider the random graph $G = G(\lambda, p)$ with vertex set $\text{Po}(\lambda)$ in which every pair of vertices at distance at most 1 forms an edge with probability p, independently of other pairs. Our study is motivated by the question of *locality* for bond percolation on the standard random geometric graph $G(\lambda, 1)$. We show that for a large class of graphs converging locally to $G(\lambda, 1)$ in a suitable sense, the corresponding critical percolation thresholds converge to the one of $G(\lambda, 1)$. The proof is based on a finite volume criterion. In this direction, we strengthen recent results of Penrose by showing that the size of the largest component rescaled by n converges almost surely to a constant, and providing sharp bounds for the size of the second-largest component.

Keywords: continuum percolation, bond percolation, Schramm's locality conjecture, random geometric graph, giant component, second-largest component
MSC Class: 05C80, 60C05, 60D05, 60K35

1 Introduction

The theory of percolation was introduced by Broadbent and Hammersley [5] more than 60 years ago. In the most classical setting, one is given a subgraph of $\mathbb{Z}^d$ for some $d \geq 2$ where every edge appears with probability $p \in [0, 1]$, independently of all other edges. While the model has a very simple definition, it undergoes a phase transition for the existence of an infinite component which is not completely understood to this day.

Several years after Broadbent and Hammersley, Gilbert [10] proposed a new mathematical model of wireless networks, which gave rise to the field of continuum percolation. His model, known as the *Random Geometric Graph* (RGG), is defined as follows: given $\lambda, R > 0$, the vertices of RGG are given by a Poisson Point Process (PPP) with intensity λ in $\mathbb{R}^2$, and whose edges are given by the pairs of points at distance at most R. In fact, as mentioned in his paper, one of the two parameters λ and R may be put to 1 by a suitable homothety of the plane. Later, Meester and Roy [17] generalized Gilbert's model by connecting randomly and independently pairs of points in the Poisson process with a probability depending on their relative positions, thus introducing the random connection model. In this paper, we aim at studying a particular example of the random connection model obtained by performing standard Bernoulli bond percolation on top of Gilbert's model.

First, we investigate the question of duality between the percolation thresholds associated to the intensity of the Poisson point process on one hand, and the bond percolation parameter on the other

hand. This question is in the spirit of [6] and [11] among others, which were interested in deriving relations between site and bond percolation thresholds on general bounded degree graphs. As we shall see, some interesting and challenging problems appear in our setting, partly due to the continuous setting and the fact that degrees are no longer bounded.

Then, we consider the question of locality for the percolation thresholds, introduced in [3] in the context of vertex-transitive graphs. When suitably rephrased in our continuous setting, we manage to prove some new instance of Schramm's locality conjecture. The proofs rely on a fine understanding of the sizes of the largest connected components in the random geometric graph in finite volume, which is of independent interest. In this direction, we extend and strengthen very recent results of Penrose [19].

1.1 Our results

For $\lambda > 0$, denote by $\mathrm{Po}(\lambda)$ a Poisson Point Process in $\mathbb{R}^2$ of intensity λ. Then, we define $G(\lambda, 1)$ as the random geometric graph with vertex set $\mathrm{Po}(\lambda)$ and edge set the set of pairs of vertices at Euclidean distance at most one. Given $p \in [0, 1]$, we further define $G = G(\lambda, p)$ as the graph obtained from $G(\lambda, 1)$ after Bernoulli bond percolation with probability p.

We shall often consider $\mathrm{Po}(\lambda)$ conditioned on containing the origin, that is, we add artificially the origin to the point process (this construction is known in a more general setup under the name *Palm theory*, see for instance [18]). Then, we shall denote by $\theta(\lambda, p)$ the probability that the connected component of the origin in G is infinite. By classical considerations from ergodic theory (see e.g. [17]) one may deduce the existence of a deterministic threshold $\lambda_0 \in [0, \infty]$ such that:

- for all $\lambda > \lambda_0$, the graph $G(\lambda, 1)$ contains an infinite connected component almost surely, and in particular $\theta(\lambda, 1) > 0$;

- for all $\lambda < \lambda_0$, the graph $G(\lambda, 1)$ contains no infinite component almost surely, and in particular $\theta(\lambda, 1) = 0$.

In fact, a standard coupling argument with site percolation on $\mathbb{Z}^2$ (see [17, 18]) shows that $\lambda_0 \in (0, \infty)$. In a similar way, for every $\lambda \in [\lambda_0, \infty)$ we define

$$p_c(\lambda) = \inf\{p \in [0, 1] : \theta(\lambda, p) > 0\},$$

and also for every $p \in (0, 1]$ we define

$$\lambda_c(p) = \inf\{\lambda \in \mathbb{R} : \theta(\lambda, p) > 0\}.$$

As a first natural question one may ask whether these two functions are strictly monotone, continuous, and inverse of each other. Each of these properties is an easy consequence of the next theorem. For the proof of the first equality, we deduce that p_c is strictly monotone by using a classical result of Hammersley [12], which is then replaced by Theorem 2.3 from [9] in the proof of the second inequality.

Theorem 1.1. *For any $\lambda \in (\lambda_0, \infty)$ and $p \in (0, 1)$, one has $\lambda_c(p_c(\lambda)) = \lambda$ and $p_c(\lambda_c(p)) = p$, respectively. In particular, λ_c and p_c are inverse bijections, and hence continuous and strictly decreasing.*

Our interest in the graph $G(\lambda, p)$ was triggered by the following question, which is related to Schramm's locality conjecture, see [3] and also [8, 14, 16] for some recent progress. Suppose that we are given an integer $k \geq 1$ and a real $\lambda > 0$. Then, remove all vertices of the random geometric graph $G(\lambda, 1)$ which have degree larger than some constant k together with all the edges emanating from them. This defines a subgraph of $G(\lambda, 1)$ in which all vertices have degree bounded by k. Moreover, as $k \to \infty$, these random subgraphs (say rooted at the origin) converge locally in the Benjamini–Schramm sense to $G(\lambda, 1)$ (see e.g. [2, 4] for more on this notion of local convergence). Thus, we were initially aiming to understand whether the bond percolation thresholds on these subgraphs converge to $p_c(\lambda)$, the bond percolation threshold of $G(\lambda, 1)$.

Our first main result confirms a more general version of this statement. We remark that, while Schramm's locality conjecture was stated in terms of vertex-transitive graphs, we assume translation invariance of the distribution of our graphs. First, a random set embedded in $\mathbb{R}^2$ is said to be *locally finite* if for any bounded domain Ω, the restriction of the set to Ω is almost surely finite.

Definition 1.2. Let $K \geq 0$ be given. We say that a random graph with a locally finite vertex set embedded in $\mathbb{R}^2$ is a *K-dependent graph of geometric type* if the following two conditions are satisfied:

- The law of the graph is invariant by any translation of $\mathbb{R}^2$.

- For any two domains Ω_1 and Ω_2 at Euclidean distance at least K, the restrictions of the graph to Ω_1 and Ω_2 are independent.

We also say that a graph is of geometric type if there exists $K \geq 0$ such that it is a K-dependent graph of geometric type.

In particular, each of the graphs $G(\lambda, p)$, or more generally any random connection model with a connection function having bounded support, is of geometric type. The example mentioned above where vertices with degree larger than a constant k are removed is another example of a 2-dependent graph of geometric type. In this setting, a convenient notion of local convergence is the following.

Definition 1.3. A sequence $(\mathcal{G}_k)_{k \geq 1}$ of graphs of geometric type in the sense of Definition 1.2 is said to *converge locally* to another graph $\mathcal{G}$ of geometric type if for any bounded domain $\Omega \subseteq \mathbb{R}^2$, the distribution of the restriction of $\mathcal{G}_k$ to Ω converges to the distribution of the restriction of $\mathcal{G}$ to Ω as a sequence of finite graphs. In other words, for any finite graph H,

$$\lim_{k \to \infty} \mathbb{P}\big((\mathcal{G}_k)_{|\Omega} = H\big) = \mathbb{P}\big(\mathcal{G}_{|\Omega} = H\big).$$

Examples of graph sequences of geometric type which converge to $G(\lambda, p)$ include:

- $G(\lambda_k, p_k)$ for any sequence $(\lambda_k)_{k \geq 1}$ and $(p_k)_{k \geq 1}$ converging respectively to λ and p;

- $\mathcal{G}_k$, obtained after removing all vertices in $G(\lambda, p)$ with degree larger than k;

- $\widehat{\mathcal{G}}_\varepsilon$, obtained after removing all vertices in $G(\lambda, p)$ with another vertex at distance at most ε (in $\mathbb{R}^d$) from them.

A natural guess is that both the bond and the site percolation thresholds are continuous for this notion of convergence in the following sense. Fix $K \geq 0$ and any sequence $(\mathcal{G}_k)_{k \geq 1}$ of K-dependent graphs of geometric type converging locally to some other graph $\mathcal{G}$. Denote by $p_c(\mathcal{G}_k)$ and $s_c(\mathcal{G}_k)$ respectively the bond and the site percolation thresholds of $\mathcal{G}_k$ (and similarly for $p_c(\mathcal{G})$ and $s_c(\mathcal{G})$). Then, one should always have

$$\lim_{k \to \infty} p_c(\mathcal{G}_k) = p_c(\mathcal{G}) \quad \text{and} \quad \lim_{k \to \infty} s_c(\mathcal{G}_k) = s_c(\mathcal{G}).$$

Our first main result proves one direction of this claim for bond percolation when the limiting graph is $G(\lambda, 1)$. We do not address the question of site percolation here, which is a harder problem, see below for some comments on it.

Theorem 1.4. *Fix $\lambda > \lambda_0$, $K \geq 0$, and a sequence of K-dependent graphs $(\mathcal{G}_k)_{k \geq 1}$ of geometric type which converges locally to $G(\lambda, 1)$ in the sense of Definition 1.3. Then, with the previous notation, one has*

$$\limsup_{k \to \infty} p_c(\mathcal{G}_k) \leq p_c(\lambda).$$

If in addition $\mathcal{G}_k$ is a subgraph of $G(\lambda, 1)$ for any $k \geq 1$, then

$$\lim_{k \to \infty} p_c(\mathcal{G}_k) = p_c(\lambda).$$

In particular the theorem applies to the example of the graph $\mathcal{G}_k$ obtained after removing from $G(\lambda, 1)$ vertices with degree larger than k.

Moreover, we remark that the proof works for bond percolation and not for site percolation. This is due to the fact that we use the sprinkling technique at several places in the proof; indeed, revealing parts of the Poisson process as a second step has the potential of forming a dense net in any bounded domain, which is not the case for edge sprinkling since the domain has probability $\Omega(1)$ to be empty.

It may not come as a surprise that the proof of Theorem 1.4 uses a finite-size criterion based on estimates on the sizes of the largest and second-largest connected components in finite volume in the supercritical regime. Such a result has been obtained very recently by Penrose [19] who proved that the size of the largest connected component of $G(\lambda, p)$ restricted to a box of volume n divided by n converges in probability to $\lambda\theta(\lambda, p)$, while the size of the second-largest component divided by n converges in probability to 0. While this would be sufficient for proving Theorem 1.4, we improve Penrose's result in two ways: on the one hand, we improve the convergence concerning the largest component to an almost sure convergence, and on the other hand, we obtain sharp bounds for the size of the second-largest component.

To state our result, we introduce some more notation. For $n \geq 1$, we consider the restriction $G_n = G_n(\lambda, p)$ of G to the box $[0, \sqrt{n}]^2$. Also, we denote by $L_1(G_n)$ the number of vertices in the largest connected component of G_n.

We say that a sequence of events $(\mathcal{E}_n)_{n \geq 0}$ holds *asymptotically almost surely* (which we abbreviate by *a.a.s.*) if $\mathbb{P}(\mathcal{E}_n) \to 1$ as $n \to \infty$. A sequence of random variables $(X_n)_{n \geq 0}$ is said to be a.a.s. of order $\Theta_{\lambda, p}(f_n)$ if there exist positive constants c and C, which might depend on λ and p, such that $\mathbb{P}(cf_n \leq X_n \leq Cf_n) \to 1$ as $n \to \infty$.

Theorem 1.5. *Fix $\lambda > 0$ and $p \in (0, 1]$, such that $\lambda > \lambda_c(p)$. Then $(n^{-1}L_1(G_n))_{n \geq 1}$ converges almost surely towards $\lambda\theta(\lambda, p)$. Moreover, the size of the second-largest component in G_n is a.a.s. of order $\Theta_{\lambda, p}\big((\log n)^2\big)$.*

Theorem 1.5 is the main technical result of this paper. The proof of the almost sure convergence for the largest component is based on estimates for the probability of crossing large rectangles with a fixed length-to-width ratio. The main difficulty here is that two edges that intersect (geometrically) in an interior point may still be in different connected components because of the bond percolation. To solve this problem, we use the classical technique of sprinkling, which consists of revealing the Poisson Point Process in two steps with the idea to locally connect two crossing edges present after the first step with positive probability.

The result on the size of the second-largest component is the most delicate part. Proving the lower bound is the easier part: it consists simply in observing that in the box $[0, \sqrt{n}]^2$, there is a.a.s. a square of side length of order $\log n$ with no point at distance 1 from its boundary and order $(\log n)^2$ points in it. The proof of the upper bound is more elaborate. The main idea is the following: we begin by proving that a.a.s. every point x in the box $[0, \sqrt{n}]^2$ is surrounded by 'many' cycles in the giant component of G_n close to x. Hence, if the connected component of $x \in G_n$ has 'large' Euclidean diameter, it necessarily intersects geometrically each of the above cycles. In this case, the argument of sprinkling that was used for the largest component does not work directly: although adding new points would help to connect x to the giant component, it could also create new components with large diameter. Thus, a refined argument is required. It roughly goes as follows: to begin with, we introduce an auxiliary graph whose vertices are identified with vertical or horizontal rectangles with aspect ratio 2 in the original graph (called *dominos*). We then perform some delicate exploration of the dominos and use sprinkling only in a few carefully chosen dominos which have the potential to connect a path from x at distance at most 1 from them with a cycle in G_n going through them. When doing the sprinkling, we require the formation of a dense net of well-connected points after the second stage, thus making sure that all new points participate in the giant.

A generalization of the model. In fact, Theorem 1.5 holds for the more general random connection model mentioned above. More precisely, denote by $\mathcal{D}$ the unit disc centered at the origin, and fix an even function $g : \mathcal{D} \to (0, 1]$ with bounded support. Then, the random graph $G(\lambda, g)$ may be defined as a subgraph of $G(\lambda, 1)$ in which every edge between two vertices x and y is retained with probability $g(x - y)$, independently of all other edges. In particular, the graph $G(\lambda, p)$ is a special case of this generalized setting with g constant equal to p. Then, by defining $\lambda_c(g)$ and $\theta(\lambda, g)$ as before, Theorem 1.5 holds for this more general setup with almost the same proof.

Stretched exponential decay in the supercritical regime. The proof of Theorem 1.5 easily implies that in the infinite volume setting, the origin is in a component of size at least n without being part of the infinite component with probability $\exp(-\Omega_{\lambda,p}(\sqrt{n}))$. For more details, see Remark 4.9.

Plan of the paper. In Section 2 we introduce notation and several preliminary results. Then, Sections 3, 4 and 5 are dedicated to the proofs of Theorem 1.1, Theorem 1.5 and Theorem 1.4, respectively.

2 Notation and preliminaries

2.1 Notation

Given a graph $\mathcal{G}$, we denote by $|\mathcal{G}|$ its number of vertices, which we also call its *size*, and by $L_1(\mathcal{G})$ the size of its largest connected component. Given an event $\mathcal{A}$, we denote by $\overline{\mathcal{A}}$ the complementary event.

We use classical asymptotic notations, namely $O(\cdot)$, $o(\cdot)$, $\Theta(\cdot)$ and $\Omega(\cdot)$. By default the limit variable is n, and the constants associated to $O(\cdot)$, $\Theta(\cdot)$ and $\Omega(\cdot)$ are independent from the parameters of the problem; if this is not the case, the parameters influencing the constants will be given as lower indices, for example $O_p(\cdot)$ or $\Theta_{\lambda,p}(\cdot)$.

For any $x = (x_1, x_2) \in \mathbb{R}^2$ and $r \geq 0$, denote $\Lambda_x(r) := [x_1 - r, x_1 + r] \times [x_2 - r, x_2 + r]$, and simply write $\Lambda(r)$ when $x = \mathbf{0} := (0, 0)$ (the origin). Furthermore, for all $n \in \mathbb{N}$, we define $\Lambda_n := [0, \sqrt{n}]^2$ and given a domain $\Omega \subseteq \mathbb{R}^2$, we set $\partial\Omega$ to be the boundary of Ω. We also write $\|x\|$ for the Euclidean norm of $x \in \mathbb{R}^2$.

Given two domains $\Omega_1, \Omega_2 \subseteq \mathbb{R}^2$ and a graph $\mathcal{G}$ with vertex set included in $\mathbb{R}^2$, we denote by $\Omega_1 \overset{\mathcal{G}}{\leftrightarrow} \Omega_2$, or simply by $\Omega_1 \leftrightarrow \Omega_2$ when $\mathcal{G}$ is clear from the context, the event that there is a path in $\mathcal{G}$ starting from a vertex in Ω_1 and ending at a vertex in Ω_2. Given a domain $\Omega \subseteq \mathbb{R}^2$, we shall also write $\Omega_1 \overset{\Omega}{\leftrightarrow} \Omega_2$, the same event with the restriction that the path is contained in Ω. Furthermore, when Ω_1 (or Ω_2) is reduced to a singleton $\{x\}$, we simply write x.

Now, we introduce some classical notation from percolation theory. Given a fixed graph $\mathcal{G} = (V, E)$ (which in our case shall mostly be $\mathbb{Z}^2$ or a subgraph of it), a *configuration* ω is an element of $\{0, 1\}^E$. As usual, we often identify ω with the subgraph of $\mathcal{G}$ with vertex set V and edge set $\{e : \omega_e = 1\}$. Edges from this set are called *open*, and the other edges of $\mathcal{G}$ are called *closed*. For $q \in [0, 1]$, Bernoulli bond percolation with parameter q on $\mathcal{G}$ is the product probability measure $\mathbb{P}_q$ for which every edge is open with probability q, independently of the other edges (we omit the reference to the base graph $\mathcal{G}$ in this notation, as it should always be clear from the context to which graph it applies). Note that by definition when the base graph $\mathcal{G}$ is the random geometric graph $G(\lambda, 1)$, the subgraph obtained after Bernoulli bond percolation with parameter p is $G(\lambda, p)$. We shall either denote its distribution by $\mathbb{P}_{\lambda,p}$, when we want to emphasize which parameters we consider, or simply by $\mathbb{P}$, when they are clear from the context.

Given a fixed graph $\mathcal{G} = (V, E)$, a configuration ω is smaller than ω' if for every $e \in E$ one has $\omega_e \leq \omega'_e$. An event $\mathcal{A} \subseteq \{0, 1\}^E$ is increasing if whenever $\omega \in A$ and $\omega \leq \omega'$, then also $\omega' \in A$, and an event $\mathcal{A}$ is decreasing if $\overline{\mathcal{A}}$ is increasing. Then, when considering the random geometric graph $G = G(\lambda, p)$, we say

that an event $\mathcal{A}$ is increasing if whenever ω and ω' are two instances of G such that ω is a subgraph of ω' and ω is in $\mathcal{A}$, then ω' is also in $\mathcal{A}$.

2.2 Preliminaries

The following well-known inequality shows that two increasing events are positively correlated. It applies to a wide variety of random models, and both Bernoulli percolation and the percolated random geometric graph model are among them, see e.g. [7, 19].

Lemma 2.1 (Harris inequality / FKG inequality). *For any $q \in [0, 1]$ and any increasing events $\mathcal{A}$ and $\mathcal{B}$, one has*

$$\mathbb{P}_q(\mathcal{A} \cap \mathcal{B}) \geq \mathbb{P}_q(\mathcal{A})\mathbb{P}_q(\mathcal{B}).$$

Similarly, for any $\lambda > 0$, $p \in [0, 1]$ and any two increasing events $\mathcal{A}$ and $\mathcal{B}$,

$$\mathbb{P}_{\lambda,p}(\mathcal{A} \cap \mathcal{B}) \geq \mathbb{P}_{\lambda,p}(\mathcal{A})\mathbb{P}_{\lambda,p}(\mathcal{B}).$$

A consequence of Lemma 2.1 is the so-called *square-root trick*: given k increasing (or decreasing) events $(\mathcal{A}_i)_{i=1}^k$, and any $q \in [0, 1]$,

$$\max_{1 \leq i \leq k} \mathbb{P}_q(\mathcal{A}_i) \geq 1 - (1 - \mathbb{P}_q(\cup_{i=1}^k \mathcal{A}_i))^{1/k}, \tag{1}$$

and similarly with $\mathbb{P}_{\lambda,p}$ instead of $\mathbb{P}_q$.

Now, we present a couple of results concerning Bernoulli bond percolation on $\mathbb{Z}^2$ which will be used several times in this work. First we need a result on dependent bond percolation. For any $k \in \mathbb{N}$, we say that a bond percolation on a graph is k-*dependent* if the states of any two (families of) edges at graph distance larger than k are independent.

Theorem 2.2 (see Theorem 0.0 in [15]). *For every $k \geq 1$ and $p \in [0, 1)$ there is $q_0 = q_0(k, p) < 1$ such that every k–dependent bond percolation measure on $\mathbb{Z}^2$, satisfying that any edge is open with probability $q > q_0$, dominates Bernoulli bond percolation with parameter p.*

Next, we state a result providing exponential decay of correlations in the supercritical regime.

Theorem 2.3 ([7], Theorem 3.3). *For every $q \in [0, 1/2)$, there exists $c_q > 0$ such that, for all $x \in \mathbb{Z}^2$,*

$$\mathbb{P}_q(\mathbf{0} \leftrightarrow x) \leq \exp(-c_q\|x\|).$$

The last result we shall need on Bernoulli percolation on $\mathbb{Z}^2$ provides some concentration for the size of the largest component in finite volume.

Theorem 2.4 (see [20], Theorem 4). *Fix $n \geq 1$, $H_n = \mathbb{Z}^2 \cap [0, \sqrt{n}]^2$ and $q \in (1/2, 1]$. Denote by $H_{n,q}$ the graph obtained from H_n after Bernoulli bond percolation with parameter q. Then, for every $\varepsilon > 0$ there exists $c = c(q, \varepsilon) > 0$ such that $\mathbb{P}_q(L_1(H_{n,q}) \leq (1 - \varepsilon)\mathbb{E}L_1(H_{n,q})) \leq \exp(-cn)$.*

Since the vertices of our graph are given by a homogeneous Poisson Point Process, the next formula, known under the name *Campbell-Mecke formula*, will be useful in our analysis. We refer to [19] for this version of the theorem, which can simply be deduced from more standard versions by integrating first against the Bernoulli percolation measure, conditionally on the Poisson Point Process.

Lemma 2.5 (Campbell-Mecke formula). *Fix $\lambda > 0$, $p \in [0, 1]$, and let $G = G(\lambda, p)$. For any non-negative measurable function f on the set of pairs $(x, \mathcal{G})$, where $\mathcal{G}$ is a graph whose vertex set is a locally finite subset of $\mathbb{R}^2$ and x is a vertex of $\mathcal{G}$, one has*

$$\mathbb{E}\Big[\sum_{X \in \mathrm{Po}(\lambda)} f(X, G) \Big] = \lambda \int_{\mathbb{R}^2} \mathbb{E}[f(x, G_x)] \, dx,$$

where G_x is obtained by adding x to the vertex set of G, and independently for each other vertex y of G at distance smaller than one from x, an edge is added between x and y with probability p.

6

Finally, we state a standard Chernoff-type inequality for Poisson random variables.

Lemma 2.6. *Let X be a Poisson random variable with mean $\lambda > 0$. Then, for any $\varepsilon > 0$,*

$$\mathbb{P}(|X - \lambda| \geq \varepsilon) \leq 2 \exp\left(-\frac{\varepsilon^2}{2(\lambda + \varepsilon)}\right).$$

3 Proof of Theorem 1.1

First, we fix $\lambda \in (\lambda_0, \infty)$ and show that $\lambda_c(p_c(\lambda)) = \lambda$. We argue by contradiction, assuming that $\lambda_c(p_c(\lambda)) \neq \lambda$. We distinguish two cases:

- if $\lambda_c(p_c(\lambda)) > \lambda$, then there is $\varepsilon > 0$ such that $\theta(\lambda + \varepsilon, p_c(\lambda)) = 0$, and hence $p_c(\lambda + \varepsilon) \geq p_c(\lambda)$. However, p_c is a non-increasing function, so $p_c(\lambda + \varepsilon) = p_c(\lambda)$;

- if $\lambda_c(p_c(\lambda)) < \lambda$, then there is $\varepsilon > 0$ such that $\theta(\lambda - \varepsilon, p_c(\lambda)) > 0$, and hence $p_c(\lambda - \varepsilon) \leq p_c(\lambda)$. However, p_c is a non-increasing function, so $p_c(\lambda - \varepsilon) = p_c(\lambda)$.

In both cases, our assumption leads to the existence of $\lambda_1, \lambda_2 \in (\lambda_0, \infty)$ satisfying $\lambda_1 < \lambda_2$ and $p_c(\lambda_1) = p_c(\lambda_2)$.

Now, set $p = p_c(\lambda_2) + \delta$ for some $\delta > 0$ to be chosen sufficiently small. On the one hand, by a classical result of Hammersley [12] the bond percolation threshold of $G(\lambda_2, p)$, which is $p_c(\lambda_2)/p$, is dominated by its site percolation threshold, which is $\lambda_c(p)/\lambda_2$. Hence, $\frac{\lambda_2}{p} \leq \frac{\lambda_c(p)}{p_c(\lambda_2)}$. On the other hand, $\theta(\lambda_1, p) > 0$ since $p > p_c(\lambda_1)$, so $\lambda_c(p) \leq \lambda_1$. We conclude that $\frac{\lambda_2}{p} \leq \frac{\lambda_1}{p_c(\lambda_2)}$, which rewrites as $\frac{\lambda_2}{\lambda_1} \leq \frac{p}{p_c(\lambda_2)} = 1 + \frac{\delta}{p_c(\lambda_2)}$. Choosing δ sufficiently small (and using $p_c(\lambda_2) > 0$) leads to a contradiction, which proves the equality.

To deduce the equality $p_c(\lambda_c(p)) = p$ for any fixed $p \in (0, 1)$, a similar reasoning provides two distinct $p_1, p_2 \in (0, 1)$ satisfying $\lambda_c(p_1) = \lambda_c(p_2)$, which contradicts the statement of Theorem 2.3 from [9]. □

Remark 3.1. For an infinite connected graph H with site percolation threshold s_c, bond percolation threshold p_c and maximum degree $D \geq 3$, the equality $p_c(s_c(p)) = p$ holds along similar lines without using Theorem 2.3 from [9] as a black box. We provide a sketch of the argument. Suppose for contradiction that $p_c(s_c(p)) \neq p$ for some p as above. Then, as in the proof of Theorem 1.1, there exist $p_1 < p_2$ with $s_c(p_1) = s_c(p_2) \in (0, 1)$. Set $s = s_c(p_2) + \delta$ for some $\delta \in (0, 1 - s_c(p_2))$ to be chosen below and consider the graph $H(s, p_2)$ obtained from H after site percolation with parameter s and bond percolation with parameter p_2. By a result of Grimmett and Stacey [11] we have $s_c \leq 1 - (1 - p_c)^{D-1}$, or equivalently $(1 - p_c)^{D-1} \leq 1 - s_c$ (see also [6] for stronger results in this direction). Moreover, note that $\theta(s, p_1) > 0$ since $s > s_c(p_1)$, and thus $p_c(s) \leq p_1$. Hence, applying the above inequality from [11] for $\widehat{H} = H(s, p_2)$ (which has maximum degree at most D, site percolation threshold $s_c(p_c)/s$ and bond percolation threshold $p_c(s)/p_2$), we infer

$$\left(1 - \frac{p_1}{p_2}\right)^{D-1} \leq \left(1 - \frac{p_c(s)}{p_2}\right)^{D-1} = (1 - b_c(\widehat{H}))^{D-1} \leq 1 - s_c(\widehat{H}) = \frac{s - s_c(p_c)}{s} = \frac{\delta}{s}.$$

Hence, choosing δ sufficiently small (and using $s_c(p_2) > 0$) we have the desired contradiction, and thus $p_c(\lambda_c(p)) = p$.

4 Proof of Theorem 1.5

From now on, $\lambda > 0$ and $p \in (0, 1]$ are fixed parameters such that $\lambda > \lambda_c(p)$. This in particular ensures that $G = G(\lambda, p)$ contains an infinite connected component almost surely.

The proof of Theorem 1.5 is divided into three parts. In Section 4.1 we prove that a.a.s. G_n contains a connected component of size $\Omega_{\lambda,p}(n)$. Despite the fact that this result has already been proved very recently by Penrose [19], its proof serves as a base for the second and third steps. In Section 4.2 we prove that a.a.s. the second-largest component in G_n has size $\Theta_{\lambda,p}((\log n)^2)$, and finally in Section 4.3 we show that $n^{-1}L_1(G_n)$ converges almost surely to $\lambda\theta(\lambda,p)$ as $n \to \infty$.

We introduce some additional notation, which will be used throughout this section. A *horizontal crossing* of a rectangle $[a,b] \times [c,d]$ in G is a path in G with edges, embedded as straight segments in the plane, such that:

- the first edge intersects the segment connecting (a,c) and (a,d);

- all edges in the path but the first and the last ones are included in the rectangle $[a,b] \times [c,d]$;

- the last edge intersects the segment connecting (b,c) and (b,d).

A *vertical crossing* is defined analogously. We denote by $\mathcal{H}([a,b] \times [c,d])$ (respectively $\mathcal{V}([a,b] \times [c,d])$) the event "a path in G crosses the rectangle $[a,b] \times [c,d]$ horizontally (vertically, respectively)". We also let $\mathcal{H}(b,d)$ and $\mathcal{V}(b,d)$ denote the events $\mathcal{H}([0,b] \times [0,d])$ and $\mathcal{V}([0,b] \times [0,d])$, respectively.

4.1 The existence of a giant component

This section is devoted to the proof of the fact that, in the supercritical regime, a.a.s. the graph G_n contains a component of linear size, which we state as a separate theorem.

Theorem 4.1. *Assume that $\lambda > 0$ and $p \in (0,1]$ are such that $\lambda > \lambda_c(p)$. Then a.a.s. one has $L_1(G_n) = \Theta_{\lambda,p}(n)$.*

The proof of this theorem is divided into two main parts. Firstly, we show in Proposition 4.2 that for any $\kappa > 0$ the event $\mathcal{H}(\kappa R, R)$ holds a.a.s. as $R \to \infty$. The proof of this part is inspired by the proof of Corollary 4.2 in [1] in the context of a closely related model but, since we perform bond percolation on top of the random geometric graph, two intersecting edges need not necessarily be part of the same connected component (as is the case for the random geometric graph without percolation). To circumvent this difficulty, we use a carefully designed sprinkling procedure, as was sketched in the introduction.

In the second part, we construct an auxiliary graph that dominates a supercritical Bernoulli bond percolation on $\mathbb{Z}^2$. Then, we rely on known results for supercritical bond percolation on finite boxes of $\mathbb{Z}^2$ and transfer them to the original graph by using that, by construction, any linear-sized subgraph of the auxiliary graph corresponds to a linear-sized subgraph of the original graph.

We start with the first part.

Proposition 4.2. *Assume that $\lambda > 0$ and $p \in (0,1]$ are such that $\lambda > \lambda_c(p)$. Then for every $\kappa > 0$, one has $\mathbb{P}_{\lambda,p}(\mathcal{H}(\kappa R, R)) \to 1$, as $R \to \infty$.*

Proof. The proof has four steps. The first one is to prove the result for $\kappa = 1/3$. The second step uses a result of Penrose [19] saying that with probability bounded away from zero, for some large enough constant K and two squares of sidelength K at distance R from each other, there exist one vertex in the first square and another in the second square that are connected by a path that stays within a box of side-length $4R$ containing its endpoints. In a third step, we prove that with probability bounded away from zero there is a path surrounding the box $\Lambda(R)$ inside $\Lambda(3R)$. Finally, as a last step, we prove that if the result holds for a given $\kappa > 0$, then it also holds for all $\kappa' \in [\kappa, 2\kappa)$.

Step 1. We show that $\mathbb{P}(\mathcal{H}(R/3, R)) \to 1$ as $R \to \infty$. Define the events

$$\mathcal{A}_R = \mathcal{H}([R/6, R/2] \times [-R/2, R/2]), \quad \mathcal{B}_R = \mathcal{H}([-R/2, -R/6] \times [-R/2, R/2]),$$
$$\mathcal{C}_R = \mathcal{V}([-R/2, R/2] \times [R/6, R/2]), \quad \mathcal{D}_R = \mathcal{V}([-R/2, R/2] \times [-R/2, -R/6]).$$

Figure 1: Connecting the box $\Lambda(R/6)$ to the box $\Lambda(R/2)$: in this case, the rectangles associated to the events $\mathcal{C}_R$ and $\mathcal{D}_R$ are separated from the rest of the $R/2 \times R/2$ square by dashed lines, and the path traverses vertically the rectangle on top, so the event $\mathcal{C}_R$ is realized (the rotated rectangles associated to the events $\mathcal{A}_R$ and $\mathcal{B}_R$ are not explicitly separated).

Note that all four events correspond to crossing a rectangle with aspect ratio 3 along its shorter side. Then, connecting the box $\Lambda(R/6)$ to the right (respectively the left, the top, or the bottom) side of $\Lambda(R/2)$ ensures that the event $\mathcal{A}_R$ (respectively $\mathcal{B}_R$, $\mathcal{C}_R$, or $\mathcal{D}_R$) is realized, see Figure 1.

Thus, the square-root trick (1) and the fact that all four events have probability $\mathbb{P}(\mathcal{H}(R/3, R))$ implies that

$$\mathbb{P}(\mathcal{H}(R/3, R)) \geq 1 - (1 - \mathbb{P}(\mathcal{A}_R \cup \mathcal{B}_R \cup \mathcal{C}_R \cup \mathcal{D}_R))^{1/4} \geq 1 - (1 - \mathbb{P}(\Lambda(R/6) \leftrightarrow \mathbb{R}^2 \setminus \Lambda(R/2)))^{1/4}.$$

Since by definition $G(\lambda, p)$ contains an infinite component almost surely, the probability that the box $\Lambda(R/6)$ intersects the infinite component tends to 1 as $R \to \infty$, which finishes the proof of Step 1.

Step 2. For every sufficiently large constant $K > 0$, one has

$$\liminf_{R \to \infty} \mathbb{P}\big(\Lambda(K) \xleftrightarrow{\Lambda(R)} \Lambda_{(R/2,0)}(K)\big) > 0.$$

This result is proved in [19], see the proof of Proposition 2.6 therein, and we refer to this paper for details. Let us stress that, in particular, this proof uses the hypothesis $\lambda > \lambda_c(p)$.

For the third step, we need a couple of new definitions. For $x \in \mathbb{R}^2$ and $r, R > 0$ satisfying $0 < r < R$, we define the annulus
$$A_x(r, R) := \Lambda_x(R) \setminus \Lambda_x(r),$$
and the event
$$\text{Circ}_x(r, R) := \{A_x(r, R) \text{ contains a cycle of } G\}.$$

We also simply write $A(r, R)$ and $\text{Circ}(r, R)$ for $A_{\mathbf{0}}(r, R)$ and $\text{Circ}_{\mathbf{0}}(r, R)$, respectively. In the next step, we prove that large annuli with constant ratio of their radii contain a cycle of G with probability bounded away from zero.

Step 3. We show that $\liminf_{R \to \infty} \mathbb{P}(\text{Circ}(R, 3R)) > 0$. Set $r = R/2$. For $i \in \{1, \ldots, 8\}$, let

$$v_i = (ir - 2R, 2R), \quad v_{i+8} = (2R, 2R - ir), \quad v_{i+16} = (2R - ir, -2R), \quad v_{i+24} = (-2R, -2R + ir).$$

9

Note that the points $(v_i)_{i=1}^{32}$ divide $\partial\Lambda(2R)$ into 32 equal segments of length r. By Step 2 we know that for every sufficiently large K and $R = R(K)$ and for every $i \in \{1, \ldots, 32\}$, we have

$$\mathbb{P}\left(\Lambda_{v_i}(K) \xleftrightarrow{\Lambda_{v_i}(R)} \Lambda_{v_{i+1}}(K)\right) \geq \frac{1}{2}, \tag{2}$$

where $v_{33} = v_1$. Since each of the events in (2) is increasing, we conclude by the FKG inequality (Lemma 2.1) that their intersection, which we call $\mathcal{A}$, has probability at least 2^{-32}. Now, for every $i \in \{1, \ldots, 32\}$, fix a tessellation $\mathcal{T}_i$ of the box $\Lambda_{v_i}(K)$ into squares of side length $1/\sqrt{5}$ (this is possible by choosing K to be an integer multiple of $\sqrt{5}$) and let

$$\mathcal{B} = \{Q \cap \mathrm{Po}(\lambda) \neq \emptyset, \text{ for all } Q \in \mathcal{T}_i \text{ and all } 1 \leq i \leq 32\}.$$

Since $\mathcal{A}$ and $\mathcal{B}$ are both increasing events, by the FKG inequality there is $\rho > 0$ such that $\mathbb{P}(\mathcal{A} \cap \mathcal{B}) > \rho$, for all sufficiently large R. Also, choose $M > 0$ so that the event

$$\mathcal{C} = \{|\Lambda_{v_i}(K) \cap \mathrm{Po}(\lambda)| \leq M, \text{ for all } 1 \leq i \leq 32\},$$

holds with probability at least $1 - \rho/2$. Then, in particular, $\mathbb{P}(\mathcal{A} \cap \mathcal{B} \cap \mathcal{C}) \geq \rho/2$.

Now, observe that all three events $\mathcal{A}$, $\mathcal{B}$ and $\mathcal{C}$ are independent of the state of the edges inside the squares $(\Lambda_{v_i}(K))_{i=1}^{32}$. Indeed, $\mathcal{B}$ and $\mathcal{C}$ depend only on $\mathrm{Po}(\lambda)$, and $\mathcal{A}$ depends on $\mathrm{Po}(\lambda)$ and the state of the edges with at least one endpoint outside these squares. Thus, calling $\mathcal{D}$ the event that, for all $i \in \{1, \ldots, 32\}$, the vertices $\Lambda_{v_i}(K) \cap \mathrm{Po}(\lambda)$ induce a single connected component from $G(\lambda, p)$, we get $\mathbb{P}(\mathcal{D} \mid \mathcal{A} \cap \mathcal{B} \cap \mathcal{C}) \geq p^{32M^2}$. Hence, since $\mathcal{A} \cap \mathcal{B} \cap \mathcal{C} \cap \mathcal{D} \subseteq \mathrm{Circ}(R, 3R)$, we get

$$\mathbb{P}(\mathrm{Circ}(R, 3R)) \geq \mathbb{P}(\mathcal{A} \cap \mathcal{B} \cap \mathcal{C} \cap \mathcal{D}) \geq \mathbb{P}(\mathcal{A} \cap \mathcal{B} \cap \mathcal{C})\mathbb{P}(\mathcal{D} \mid \mathcal{A} \cap \mathcal{B} \cap \mathcal{C}) \geq \frac{\rho p^{32M^2}}{2},$$

which finishes the proof of Step 3.

For the last step, we use a sprinkling argument.

Step 4. If $\mathbb{P}_{\lambda', p}(\mathcal{H}(\kappa R, R)) \to 1$ as $R \to \infty$ for some fixed positive parameters κ, λ' and p, then for any $\lambda > \lambda'$ one also has $\mathbb{P}_{\lambda, p}(\mathcal{H}(2\kappa R, R)) \to 1$.

Suppose that the statement holds for some given κ, λ' and p, and consider the rectangles

$$\Pi = [0, 2\kappa R] \times [0, R], \quad \Pi_1 = [0, \kappa R] \times [0, R], \quad \Pi_2 = [\kappa R, 2\kappa R] \times [0, R].$$

In particular, Π_1 and Π_2 are respectively the left and right halves of Π. Then, divide the segment between the points $(\kappa R, 0)$ and $(\kappa R, R)$ into N segments $I_1, \ldots, I_N$, of equal length, where $N \in \mathbb{N}$ will be chosen appropriately later. Also, for any $i \in \{1, \ldots, N\}$ and $j \in \{1, 2\}$, let $\mathcal{B}_{i,j}$ be the event that Π_j contains a horizontal crossing which intersects I_i. For any fixed N, using our assumption and the square-root trick (1), we have for both $j \in \{1, 2\}$ that

$$\max_{1 \leq i \leq N} \mathbb{P}_{\lambda', p}(\mathcal{B}_{i,j}) \geq 1 - (1 - \mathbb{P}_{\lambda', p}(\mathcal{H}(\kappa R, R)))^{1/N},$$

which tends to 1 as $R \to \infty$. Denote by i_0 the smallest index realizing the maximum above (it is the same for both $j \in \{1, 2\}$ by symmetry of Π_1 and Π_2), and let x_0 be the midpoint of the interval I_{i_0}.

By Step 3, there exist R_0 and $\delta > 0$ such that for every $R \geq R_0$, $\mathbb{P}_{\lambda', p}(\mathrm{Circ}(R, 3R)) \geq \delta$. Assume now that $R/N \geq R_0$ and define

$$K := \left\lfloor \log_4\left(\min(\kappa, 1/2)N\right)\right\rfloor - 1, \quad \text{and} \quad D := R/N.$$

For every $k \in \{0, \ldots, K\}$, define further the annulus $A_k := A_{x_0}(4^k D, 3 \cdot 4^k D)$ and note that, by construction, any horizontal crossing of Π_1 or Π_2 that intersects I_{i_0} also crosses each of these annuli. Denote by $\mathcal{D}_k$

10

Figure 2: In the figure, the central vertex is x_0; around it, there are two cycles, γ_{k_ℓ} and $\gamma_{k_{\ell+1}}$, positioned in the annuli A_{k_ℓ} and $A_{k_{\ell+1}}$. Note that $\gamma'_{k_\ell} = \gamma_{k_\ell}$ but $\gamma'_{k_{\ell+1}} \subsetneq \gamma_{k_{\ell+1}}$ (γ'_{k_ℓ} and $\gamma'_{k_{\ell+1}}$ are thickened while $\gamma_{k_{\ell+1}} \setminus \gamma'_{k_{\ell+1}}$ is transparent). The proportions of the side lengths of the boxes, centered at x_0, are not the real ones in the above schematic representation.

the event that there is a cycle inside the annulus A_k. Recall that by Step 3 and our choice of constants, each of these events has probability at least $\delta > 0$ under $\mathbb{P}_{\lambda',p}$. Denote by $\mathcal{D}$ the event $\{\sum_{k=0}^{K} \mathbb{1}_{\mathcal{D}_k} \geq \delta K/2\}$. Since $(\mathbb{1}_{\mathcal{D}_k})_{k=0}^{K}$ dominates an independent family of $K+1$ Bernoulli random variables with parameter δ, Chernoff's inequality implies that $\mathcal{D}$ holds with probability at least $1 - \exp(-\Omega(\delta K))$. In particular, for every $\varepsilon > 0$ and for every sufficiently large R (allowing for N and K to be sufficiently large as well), the event $\mathcal{B}_{i_0,1} \cap \mathcal{B}_{i_0,2} \cap \mathcal{D}$ holds with probability at least $1 - \varepsilon$.

Now, we first sample the graph $G(\lambda', p)$, and perform a sprinkling to connect the crossings of Π_1 and Π_2 intersecting I_{i_0} to one of the cycles around x_0 on the event $\mathcal{B}_{i_0,1} \cap \mathcal{B}_{i_0,2} \cap \mathcal{D}$. For this, fix some $\lambda > \lambda'$ and recall that by some well-known properties of Poisson Point Processes one may construct the graph $G(\lambda, p)$ by first sampling independent copies of $G(\lambda', p)$ and $G(\lambda - \lambda', p)$, and then adding independently an edge between any pair of vertices $v \in G(\lambda', p)$ and $v' \in G(\lambda - \lambda', p)$ satisfying $\|v - v'\| \leq 1$ with probability p.

So, assume that the event $\mathcal{B}_{i_0,1} \cap \mathcal{B}_{i_0,2} \cap \mathcal{D}$ holds for $G(\lambda', p)$, and let P_1 and P_2 be horizontal crossings of Π_1 and Π_2, respectively. Also, set $L = \lfloor \delta K/2 \rfloor$ and let $(\gamma_{k_\ell})_{\ell=1}^{L}$ be cycles of $G(\lambda', p)$ provided by $\mathcal{D}$ in the annuli $(A_{k_\ell})_{\ell=1}^{L}$, respectively. Note that parts of these paths may be outside Π. Thus, for every $\ell \in \{1, \ldots, L\}$, denote by $\gamma'_{k_\ell} \subseteq \gamma_{k_\ell}$ a path or a cycle whose vertices, except possibly the first and the last one, are in Π and which separates x_0 from the left and the right side of Π, see Figure 2. Then, connecting the paths P_1 and P_2 to the same path or cycle among $(\gamma'_{k_\ell})_{\ell=1}^{L}$ within Π forms a horizontal crossing of Π.

Now, for each $j \in \{1, 2\}$ and each $\ell \in \{1, \ldots, L\}$, there are two edges, one in P_j and one in γ'_{k_ℓ}, that intersect. By the triangle inequality, there is one endvertex of the first edge and one endvertex of the second edge which both lie in Π_j and are at distance at most $3/2$ from each other. Indeed, letting $u_{j,\ell}$ and $v_{j,\ell}$ be the two endvertices of the first edge, $w_{j,\ell}$ be one of the endvertices of the second edge that lies in Π_j, and z be the intersection point of the two edges,

$$\min\{\|u_{j,\ell} - w_{j,\ell}\|, \|w_{j,\ell} - v_{j,\ell}\|\} \leq \frac{\|u_{j,\ell} - w_{j,\ell}\| + \|w_{j,\ell} - v_{j,\ell}\|}{2} \leq \frac{\|u_{j,\ell} - z\| + \|z - v_{j,\ell}\| + 2\|z - w_{j,\ell}\|}{2} \leq \frac{3}{2}.$$

Fix two vertices as above, and note that the area of the region, obtained by intersecting the two unit balls centered around them and Π_j, is at least the area of the intersection of the unit balls around $\mathbf{0}$, around $(3/2, 0)$ and the quarter-plane $\{(x, y) : x, y \geq 0\}$, which we denote by α. We conclude that, for every

11

$\ell \in \{1, \ldots, L\}$, P_1 and P_2 connect to γ_{k_ℓ} after adding $G(\lambda - \lambda', p)$, with probability at least $\mathrm{e}^{-2(\lambda-\lambda')\alpha}p^4$, and these events are all independent of each other since the annuli $(A_k)_{k=0}^K$ are disjoint. Thus, for every sufficiently large R, the probability that $G(\lambda, p)$ contains a horizontal crossing of Π is at least

$$(1 - \varepsilon)(1 - (1 - \mathrm{e}^{-2(\lambda-\lambda')\alpha}p^4)^L) \geq 1 - 2\varepsilon.$$

Since this statement holds for any $\varepsilon \in (0, 1)$, the proof of Step 4 is completed.

To conclude, note that altogether the four steps prove Proposition 4.2. $\qquad\square$

In the proof of Step 4 above we saw how to glue crossings of two parallel copies of the same rectangle to form a longer crossing. In the next remark, we extend this gluing procedure to the setting of two orthogonal copies of the same rectangle, as we shall need this later, e.g. in the proof of Theorem 4.1 below. We omit its proof since it requires only minor modifications to the proof of Step 4 above.

Remark 4.3. Fix $\kappa > 1$ and some $\lambda > \lambda_c(p)$. Then, consider the rectangles $\Pi_3 = [0, \kappa R] \times [0, R]$ and $\Pi_4 = [0, R] \times [0, \kappa R]$. Also, fix any positive integer $N \geq 3$ and denote $\alpha_N = 1/2 - N^{-1}$ and $\beta_N = 1/2 + N^{-1}$. From Proposition 4.2 we know that the rectangle $\Pi_3' := [1, \kappa R - 1] \times [\alpha_N R, \beta_N R] \subseteq \Pi_3$ (respectively $\Pi_4' := [\alpha_N R, \beta_N R] \times [1, \kappa R - 1] \subseteq \Pi_4$) is crossed horizontally (vertically, respectively) in $G(\lambda', p)$ a.a.s. as $R \to \infty$, where $\lambda' = \frac{\lambda + \lambda_c(p)}{2}$. By considering a number of disjoint annuli around the square $[\alpha_N R, \beta_N R]^2$, we deduce as in Step 4 that any horizontal crossing of Π_3' and any vertical crossing of Π_4' connect within $\Pi_3 \cap \Pi_4$ in $G(\lambda, p)$ with probability arbitrarily close to 1 when N is sufficiently large.

We are now in position to conclude the proof of the fact that a.a.s. G_n has a component of linear size.

Proof of Theorem 4.1. Fix some large $R > 0$ so that $\sqrt{n}/R \in \mathbb{N}$. Then, tessellate the square Λ_n into $R \times R$ squares and combine these into horizontal and vertical *dominos*, that is, rectangles of dimensions $2R \times R$ and $R \times 2R$, respectively. By identifying dominos with vertices and declaring that two dominos are neighbors if they are orthogonal and their intersection is an $R \times R$ square, this defines a graph isomorphic to (a finite subgraph of) $\mathbb{Z}^2$. We now construct an auxiliary random subgraph H_{aux} by applying the following bond percolation on $\mathbb{Z}^2$: let $\lambda' := \frac{\lambda + \lambda_c(p)}{2}$, and call a domino *admissible* if a thinner inner rectangle, as described in Remark 4.3, is crossed by an open path in $G(\lambda', p)$. In this case, we associate to this domino an arbitrary one of these paths. Then, we say that the edge between two neighboring admissible dominos Π_1 and Π_2 is open if their corresponding paths are connected by a path in $\Pi_1 \cap \Pi_2 \cap G(\lambda, p)$. By Proposition 4.2 and Remark 4.3, every edge is open with a probability that can be made arbitrarily close to 1 by choosing R sufficiently large (see the proof of Step 4 above). Moreover, the states of two edges at distance larger than 2 are independent. In other words, H_{aux} is obtained via a 2–dependent bond percolation on $\mathbb{Z}^2$. Thus, by applying Theorem 2.2 we deduce that, by choosing R sufficiently large, one can ensure that H_{aux} dominates an independent bond percolation with parameter $q = 3/4$, say. Moreover, it is well known that the critical threshold for independent bond percolation on $\mathbb{Z}^2$ is equal to $1/2$ and also that, in the supercritical regime, by Theorem 2.4, the largest component in the box Λ_n is a.a.s. of size $\Omega_q(n)$ as $n \to \infty$. Finally, notice that, since R is a large but fixed constant, if H_{aux} contains a component of linear size, then G_n does as well, which concludes the proof of Theorem 4.1. $\qquad\square$

4.2 The second-largest component

This section is dedicated to the proof of the following theorem. Recall that we assume throughout that λ and p satisfy $\lambda > \lambda_c(p)$.

Theorem 4.4. *A.a.s. the second-largest component in $G_n(\lambda, p)$ has size $\Theta_{\lambda, p}((\log n)^2)$.*

The proof is divided into two parts. We start with the lower bound.

Proposition 4.5. *A.a.s. the second-largest component in $G_n(\lambda, p)$ has size $\Omega_{\lambda, p}((\log n)^2)$.*

Proof. Tessellate the square Λ_n into subsquares of side length $\frac{\log n}{5\lambda}$. On the one hand, for any given square of the tessellation, the probability that there is no point of $\mathrm{Po}(\lambda)$ in it at distance at most 1 from its boundary is equal to $\exp\left(-\lambda\left(\frac{4\log n}{5\lambda}-4\right)\right) \geq n^{-4/5}$. On the other hand, by Theorem 4.1, the probability that in such a square there is a connected component of size $\Omega_{\lambda,p}((\log n)^2)$, whose vertices are all at distance at least one from its boundary, is equal to $1 - o(1)$. In particular, for any square in the tessellation, the probability that it contains a connected component of size $\Omega_{\lambda,p}((\log n)^2)$ which is disconnected from the rest of the graph G_n is at least $(1 - o(1))n^{-4/5}$. Since there are $n^{1-o(1)}$ squares in the tessellation, by Chernoff's inequality a.a.s. there is a square that satisfies the above condition, concluding the proof of the proposition. $\qquad\square$

We now aim at showing the upper bound. This will be done in several steps. The first point is to show that, roughly speaking, in supercritical Bernoulli bond percolation on $\mathbb{Z}^2$ large boxes are typically crossed by many disjoint paths (more precisely, the number of crossings typically has the same order as the side length of the box), see Corollary 4.8 below for a more precise statement. Then, we use the fact observed in the previous section that, in a sense, the continuous model dominates supercritical Bernoulli bond percolation. We deduce that any small window of side length of order $\log n$ in Λ_n contains many disjoint cycles surrounding its center. Finally, some delicate sprinkling allows us to conclude that a.a.s. for each vertex v in Λ_n, no path from v exits the window around v without connecting to at least one of these cycles.

Theorem 4.4 will then follow immediately by combining Proposition 4.5 with the following upper bound for the second-largest component, which we state as a separate proposition.

Proposition 4.6. *A.a.s. the second-largest component in $G_n(\lambda,p)$ has size $O_{\lambda,p}((\log n)^2)$.*

To prove this proposition, we first need some preliminary results on ordinary Bernoulli bond percolation, as mentioned above. For $k \geq 1$ and $x \in \mathbb{Z}^2$, we define $\mathcal{C}_k(x)$ the *k-th order connected component of* x as the set of vertices $y \in \mathbb{Z}^2$ which are connected to x by a path made of any number of open edges and at most $k - 1$ closed edges. In particular, the first order connected component is the usual connected component of x. Also, for $A \subseteq \mathbb{Z}^2$, define $\mathcal{C}_k^A(x)$ as the k-th order connected component of x in $\mathbb{Z}^2 \setminus A$. We prove the following fact.

Lemma 4.7. *For any $q \in [0, 1/2)$, there are positive constants α and C such that, for any positive integers k and t, one has*
$$\mathbb{P}_q\left(\sup_{x\in\mathcal{C}_k(\mathbf{0})} \|x\| \geq t\right) \leq C^k \exp(-\alpha t).$$

Proof. The proof goes by induction on k. The result for $k = 1$ follows immediately from Theorem 2.3. Now assume that it holds for some k. Then, for any $t \geq 1$,
$$\mathbb{P}_q\left(\sup_{x\in\mathcal{C}_{k+1}(\mathbf{0})} \|x\| \geq t\right) \leq \sum_{\substack{v,w\in\mathbb{Z}^2 \\ \|w-v\|=1}} \mathbb{P}_q\left(\mathbf{0} \leftrightarrow v,\ \mathbf{0} \nleftrightarrow w,\ \sup_{x\in\mathcal{C}_k^{\mathcal{C}_1(\mathbf{0})}(w)} \|x-w\| \geq t-1-\|v\|\right).$$

Then, by conditioning on $\mathcal{C}_1(\mathbf{0})$ and using the induction hypothesis together with the fact that, for any fixed $A \subseteq \mathbb{Z}^2$, $\mathcal{C}_k^A(w)$ is dominated by $\mathcal{C}_k(w)$, we get
$$\mathbb{P}_q\left(\sup_{x\in\mathcal{C}_{k+1}(\mathbf{0})} \|x\| \geq t\right) \leq 4 \sum_{v\in\mathbb{Z}^2} \mathbb{P}_q\left(\mathbf{0} \leftrightarrow v\right) \cdot \mathbb{P}_q\left(\sup_{x\in\mathcal{C}_k(\mathbf{0})} \|x\| \geq t-1-\|v\|\right)$$
$$\leq C^k e^{-\alpha t} \cdot \left(4e^\alpha \sum_{v\in\mathbb{Z}^2} \mathbb{P}_q\left(\mathbf{0} \leftrightarrow v\right)e^{\alpha\|v\|}\right), \tag{3}$$

where the factor of 4 in the first inequality comes from the fact that every vertex in $\mathbb{Z}^2$ has 4 neighbors. Applying again Theorem 2.3 yields the existence of a sufficiently small α so that the sum in (3) is finite, which concludes the proof of the induction step for $C = 4e^\alpha \sum_{v\in\mathbb{Z}^2} \mathbb{P}_q(\mathbf{0} \leftrightarrow v)e^{\alpha\|v\|}$. $\qquad\square$

13

As a corollary, we obtain the following fact.

Corollary 4.8. *Let $q > 1/2$ be given. There exist positive constants C and α, such that for any positive integers k and N,*

$$\mathbb{P}_q\big(\text{There exists } k \text{ disjoint horizontal crossings in } [0, 2N] \times [0, N]\big) \geq 1 - C^k \exp(-\alpha N).$$

Proof. Consider the dual graph of $\mathbb{Z}^2$ (its vertices are the faces delimited by the edges of the graph $\mathbb{Z}^2$, and two faces are adjacent if they share an edge). Recall that one can couple Bernoulli bond percolation with parameter q (and $1 - q$, respectively) on $\mathbb{Z}^2$ (and its dual, respectively) by declaring an edge between two vertices of the dual graph open if the edge separating the two corresponding faces of the original graph is closed, and vice versa. Now it suffices to observe that if there are no k disjoint horizontal crossings of the rectangle $[0, 2N] \times [0, N]$, then the k-th order connected component of one of the vertices, corresponding to the squares with centers $\{(1/2 + i, -1/2)\}_{i=0}^{2N-1}$ in the dual graph, reaches the top of the rectangle. By Lemma 4.7, we know that for any such vertex this happens with probability at most $C^k \exp(-\alpha N)$ for some positive constants α and C. Then, using a union bound over the $2N$ dual vertices on the bottom side and reducing α by factor of 2, say, concludes the proof. $\qquad\square$

We are now ready for proving the upper bound on the size of the second-largest component of G_n. We will write G_n^x for the percolated random geometric graph with vertex set $(\mathrm{Po}(\lambda) \cup \{x\}) \cap \Lambda_n$.

Proof of Proposition 4.6. Fix $\lambda' = \frac{\lambda + \lambda_c(p)}{2}$ and recall the construction of $G(\lambda, p)$ from the independent copies of $G(\lambda', p)$ and $G(\lambda - \lambda', p)$ used in Step 4 of the proof of Proposition 4.2. Consider a sufficiently large R so that $\sqrt{n}/R \in \mathbb{N}$. We construct a version of the auxiliary graph H_{aux} from the proof of Theorem 4.1 as follows. First, reveal $G(\lambda', p)$ and construct a vertex in H_{aux} if its corresponding domino contains a path as described in Remark 4.3. Then, construct an edge between two neighboring vertices of H_{aux} if, firstly, the paths in their corresponding dominos Π_1 and Π_2 intersect at a point surrounded by at least K disjoint cycles in $G(\lambda', p)$ as in Remark 4.3, with K to be fixed, and secondly, they are connected within $\Pi_1 \cap \Pi_2$ in $G(\lambda, p)$. Note that, up to choosing R and K sufficiently large, H_{aux} dominates a Bernoulli bond percolation with parameter $3/4$ in a box of $\mathbb{Z}^2$ with side length roughly $\sqrt{n}/R$. We next tessellate this new box into overlapping horizontal and vertical rectangles with side lengths $C_1 \log n$ and $2C_1 \log n$ called *log-dominos*, where C_1 is some large constant satisfying $\frac{\sqrt{n}}{RC_1 \log n} \in \mathbb{N}$, which will be fixed later. Set

$$W := \left\lfloor \frac{\alpha C_1}{2 \log C} \cdot \log n \right\rfloor,$$

with α and C as in Corollary 4.8. Then, consider the event $\mathcal{A}_n$ that all log-dominos are crossed by at least W disjoint paths in H_{aux} along their longer side. Note that on $\mathcal{A}_n$, since crossings of two orthogonal overlapping log-dominos have a common vertex of H_{aux}, all these paths are connected in H_{aux} and thus correspond to a connected component $\widehat{G}_n$ of G_n of size $\Theta(n)$. Moreover, a union bound and Corollary 4.8 tell us that if C_1 is chosen sufficiently large, then

$$\mathbb{P}(\mathcal{A}_n) \geq 1 - n^2 \exp\left(-\frac{\alpha C_1 \log n}{2}\right) = 1 - o(1). \tag{4}$$

We aim to bound the probability that there is a connected component of G_n whose diameter is at least $4C_1 \log n$ without being connected to this giant component. For this, the strategy will be roughly to show that if such a component exists, say in a square of side length of order $\log n$ centered at a point $X \in \mathrm{Po}(\lambda) \cap \Lambda_n$, it will have to cross many cycles surrounding X in this window, and we estimate the cost of this scenario by sprinkling $G(\lambda - \lambda', p)$ on top of $G(\lambda', p)$. Note that some care is required here; indeed, adding a bit of intensity to the Poisson process could also help creating new connected components with large diameter. In our proof, we first discover the cycles of H_{aux} surrounding X by an exploration from

14

the boundary of a small box around it towards its center, then find possible paths starting from X and coming close to these cycles, and finally perform a sprinkling only in constant-sized regions with the aim to connect the path from X to any of the cycles with probability bounded away from 0.

To be more precise now, for any $x \in \Lambda_n$, consider the box $\Lambda_n^x := \Lambda_x(4C_1 \log n)$ and define the event

$$\mathcal{O}_n^x := \big\{ x \xleftrightarrow{G_n^x} \mathbb{R}^2 \setminus \Lambda_n^x \big\},$$

that is, the connected component of x in G_n^x reaches the complement of Λ_n^x. Let also $\mathcal{U}_n^x$ be the event that in the graph $H_{\mathrm{aux}} \cap \Lambda_n^x$ there are W disjoint cycles surrounding x or, in case x is at distance less than $4C_1 \log n$ from $\partial\Lambda_n$, paths from the boundary of H_{aux} to itself, which all surround the point x when seen as cycles (parts of which may coincide with the boundary of Λ_n). Note the important fact that, by construction,

$$\mathcal{A}_n \ \subseteq \ \bigcap_{x \in \Lambda_n} \mathcal{U}_n^x. \tag{5}$$

On the event $\mathcal{U}_n^x$, we denote by $\mathcal{C}_1$ the outermost cycle of H_{aux} in Λ_n^x that surrounds x (or possibly the outermost path from the boundary of H_{aux} to itself in case x is at distance smaller than $4C_1 \log n$ from $\partial\Lambda_n$). With a slight abuse of notation, we alternatively view $\mathcal{C}_1$ as a cycle in the graph H_{aux}, a set of dominos in Λ_n^x, or a cycle in G_n obtained by connecting the crossings associated to the sequence of admissible dominos forming the cycle in H_{aux}. Let also $\mathcal{C}_1^+$ be the enlarged cycle made of the points in Λ_n^x at distance smaller than one from $\mathcal{C}_1$ (viewed here as a union of dominos).

We next define $\mathcal{C}_2$ as the outermost cycle surrounding x in the restriction of H_{aux} to the region enclosed by $\mathcal{C}_1^+$. Note that on $\mathcal{U}_n^x$, by repeating again this procedure, one can define inductively a sequence of disjoint cycles $\mathcal{C}_3, \ldots, \mathcal{C}_{W/2}$ (from now on, we assume W to be even since decreasing it by 1 does not modify the argument), which all surround x and are at distance at least 1 from each other. Let us define the event $\widetilde{\mathcal{U}}_n^x$ that the cycles $\mathcal{C}_1, \ldots, \mathcal{C}_{W/2}$ as constructed above exist. In particular, $\mathcal{U}_n^x \subseteq \widetilde{\mathcal{U}}_n^x$; it will turn out later that it is more convenient to work with the latter for some reasons related to measurability of these events.

Now, for $i = 1, \ldots, W/2$, denote

$$\mathcal{B}_{n,i}^x := \widetilde{\mathcal{U}}_n^x \cap \{\text{none of the paths from } x \text{ to } \mathbb{R}^2 \setminus \Lambda_n^x \text{ shares a vertex with } \mathcal{C}_i \text{ in } G_n^x\},$$

and

$$\mathcal{B}_n^x := \bigcap_{i=1}^{W/2} \mathcal{B}_{n,i}^x.$$

Recall that on the event $\mathcal{A}_n$, the cycles $\mathcal{C}_1, \ldots, \mathcal{C}_{W/2}$, are all part of the same connected component in G_n. Therefore, if $\mathcal{A}_n$ holds and the connected component of a point $X \in \mathrm{Po}(\lambda) \cap \Lambda_n$ reaches $\partial\Lambda_n^X$ in G_n without being part of the connected component containing the cycles $\mathcal{C}_1, \ldots, \mathcal{C}_{W/2}$, then necessarily $\mathcal{B}_n^X$ must hold. Thus, our aim now is to bound the probability of these events.

We start by observing that a union bound and Lemma 2.5 yield

$$\mathbb{P}\Big(\bigcup_{X \in \mathrm{Po}(\lambda) \cap \Lambda_n} \mathcal{O}_n^X \cap \mathcal{B}_n^X \Big) \leq \mathbb{E}\Big[\sum_{X \in \mathrm{Po}(\lambda) \cap \Lambda_n} \mathbb{1}_{\mathcal{O}_n^X \cap \mathcal{B}_n^X} \Big] = \lambda \int_{\Lambda_n} \mathbb{P}\big(\mathcal{O}_n^x \cap \mathcal{B}_n^x \big) \, dx,$$

and our goal is thus to bound the probabilities on the right-hand side from above. Fix some $x \in \Lambda_n$, and for simplicity assume that it is at distance at least $4C_1 \log n$ from $\partial\Lambda_n$; the proof in the case of x being closer to $\partial\Lambda_n$ is analogous. We first show that the probability that a path in G_n^x between x and $\mathbb{R}^2 \setminus \Lambda_n^x$ does not intersect $\mathcal{C}_1$ is bounded away from 1. For this, note that $\mathcal{C}_1$ can be determined by exploring the restriction of H_{aux} to Λ_n^x by starting from its boundary. More precisely, it suffices to reveal only the connected components of the dual graph that touch the boundary of $H_{\mathrm{aux}} \cap \Lambda_n^x$.

Let U_1 denote the region inside Λ_n^x which is enclosed by $\mathcal{C}_1$ (seen as a set of dominos). Then, conditionally on $\mathcal{C}_1$, the distribution of the point configuration in U_1 is still that of an independent Poisson

Point Process with intensity λ. Moreover, on the event $\mathcal{O}_n^x$, the connected component of x in $G_n^x \cap U_1$ contains at least one vertex at distance smaller than one from $\mathcal{C}_1$. We define the *first one* of these vertices as follows.

First, explore $G(\lambda, p)$ in the entire region $R_1 \subseteq U_1$ enclosed by $\mathcal{C}_1^+$. Then, for $t \geq 0$, let B_t be the set of points at distance at most t from R_1, and let

$$T := \inf\{t \geq 0 : \text{there exists } y \in \partial B_t \text{ such that } x \xleftrightarrow{G_n^x \cap B_t} y\}.$$

Note that T is a stopping time with respect to the filtration $(\mathcal{F}_t)_{t \geq 0}$ defined by

$$\mathcal{F}_t := \sigma(\mathcal{C}_1, G_n^x \cap B_t), \text{ for all } t \geq 0.$$

Moreover, as we already mentioned, one has

$$\mathcal{O}_n^x \subseteq \{T < \infty\}. \tag{6}$$

Let X_T be the (almost surely unique) point on ∂B_T which is connected to x in $G_n^x \cap B_T$. Let Q_T be the (almost surely unique) closest $R \times R$ square from X_T which is part of $\mathcal{C}_1$ (when viewing $\mathcal{C}_1$ as a succession of $R \times R$ squares by dividing every domino in half), and let $S_T := \Lambda_{X_T}(1) \cap (U_1 \setminus B_T)$ be the part of the box $\Lambda_{X_T}(1)$ which is still unexplored when only $\mathcal{C}_1$ and B_T has been revealed. We now define an event $\mathcal{E}_n^x$ on which X_T connects to $\mathcal{C}_1$. Firstly, recall that we set $\lambda' = \frac{\lambda + \lambda_c(p)}{2}$, and we view $\mathrm{Po}(\lambda)$ as the union of two independent Poisson Point Processes with intensities λ' and $\lambda - \lambda'$, respectively. Then, consider a tessellation into squares of side length $1/\sqrt{5}$ of Q_T, and let $\mathcal{E}_n^x$ be the event that the following three conditions are satisfied:

- There exists a set $\mathcal{P} \subseteq \mathrm{Po}(\lambda - \lambda')$ containing exactly one point in every square of the tessellation, and such that every pair of points of $\mathcal{P}$ in adjacent squares is connected by an edge in $G(\lambda - \lambda', p)$ (points in adjacent squares are always at distance smaller than one by construction).

- One of the points in $\mathcal{P}$ is connected by an edge between $G(\lambda - \lambda', p)$ and $G(\lambda', p)$ to the path associated to one of the dominos in $\mathcal{C}_1$ containing Q_T (we recall that by definition these paths are part of $G(\lambda', p)$).

- If X_T is at distance smaller than 0.1 from Q_T, we ask that it is connected by an edge in $G(\lambda - \lambda', p)$ to the closest point of $\mathcal{P}$ (this is possible since $0.1 + \sqrt{2/5} < 1$, where $\sqrt{2/5}$ is the length of a diagonal of a square in the tessellation). If not, we require that in S_T there is a point of $\mathrm{Po}(\lambda - \lambda')$ that is connected by an edge in $G(\lambda - \lambda', p)$ to both X_T and a point of $\mathcal{P}$.

Note that, conditionally on $G(\lambda', p)$, $\mathcal{E}_n^x$ is an increasing event which is measurable with respect to $G(\lambda - \lambda', p)$, and on which, by construction, X_T is connected to $\mathcal{C}_1$ (seen here as a cycle in the original graph G_n), in $G(\lambda, p)$. We claim that there exists $\varepsilon > 0$ (only depending on R, λ and p) such that almost surely

$$\mathbb{P}(\mathcal{E}_n^x \mid \mathcal{F}_T) \cdot \mathbb{1}_{\widetilde{\mathcal{U}}_n^x \cap \{T < \infty\}} \geq \varepsilon \cdot \mathbb{1}_{\widetilde{\mathcal{U}}_n^x \cap \{T < \infty\}}. \tag{7}$$

(Note that, while T is clearly $\mathcal{F}_T$-measurable, $\widetilde{\mathcal{U}}_n^x$ is also $\mathcal{F}_T$-measurable since each of $\mathcal{C}_2, \ldots, \mathcal{C}_{W/2}$ is included in R_1; this explains in particular why it is more convenient to work with $\widetilde{\mathcal{U}}_n^x$ instead of $\mathcal{U}_n^x$.) To see this, notice first that, since T is a stopping time, conditionally on $\mathcal{F}_T$, the distribution of G_n^x in the region $U_1 \setminus B_T$ is that of an independent percolated random geometric graph; in particular, given that the first two conditions are satisfied, the third one is achieved at a constant cost independently of $\mathcal{F}_T$. On the other hand, the first two conditions are more delicate to handle. The first point to observe is that by construction, conditionally on $\mathcal{F}_T$, the only information that we have on Q_T is that it is part of one, two or three admissible dominos whose associated paths get connected in Q_T when sprinkling $\mathrm{Po}(\lambda - \lambda')$. In particular, all dominos disjoint from the interior of U_1 that contain Q_T are admissible and participate in

16

the cycle $\mathcal{C}_1$. This has as an important consequence that, conditionally on $G(\lambda, p)$ in $\Lambda_n^x \setminus (Q_T \cup U_1)$ and on $G(\lambda', p)$ in Q_T, the event that Q_T indeed closes the cycle $\mathcal{C}_1$ when sprinkling $\mathrm{Po}(\lambda - \lambda')$ is increasing. Therefore, by the FKG inequality (see Lemma 2.1) conditionally on this event, the probability that the first two conditions of $\mathcal{E}_n^x$ are satisfied is greater than its corresponding probability without the conditioning, which is at least some positive absolute constant, proving our claim (7).

Now, letting $\widetilde{\mathcal{B}}_n^x := \cap_{i=2}^{W/2} \mathcal{B}_{n,i}^x$ and observing that it is $\mathcal{F}_T$-measurable, we can write

$$\mathbb{P}(\mathcal{B}_n^x \cap \{T < \infty\}) \leq \mathbb{P}(\overline{\mathcal{E}_n^x} \cap \widetilde{\mathcal{B}}_n^x \cap \{T < \infty\}) \leq \mathbb{E}\Big[\mathbb{P}(\overline{\mathcal{E}_n^x} \mid \mathcal{F}_T) \cdot \mathbb{1}_{\widetilde{\mathcal{B}}_n^x \cap \{T < \infty\}}\Big] \leq (1 - \varepsilon)\mathbb{P}(\widetilde{\mathcal{B}}_n^x \cap \{T < \infty\}),$$

where for the last inequality we used that the event $\mathcal{E}_n^x$ has the same distribution conditionally on $\widetilde{\mathcal{B}}_n^x$ and $\widetilde{\mathcal{U}}_n^x$, respectively, and we can use claim (7) with $\mathbb{1}_{\widetilde{\mathcal{B}}_n^x \cap \{T < \infty\}}$ instead of $\mathbb{1}_{\widetilde{\mathcal{U}}_n^x \cap \{T < \infty\}}$; indeed, the fact whether or not there are edges between paths starting from x and some cycle among $\mathcal{C}_2, \ldots, \mathcal{C}_{W/2}$ is independent from the state of the edges between these paths and $\mathcal{C}_1$. Hence, by (6) and an immediate induction we deduce that

$$\mathbb{P}(\mathcal{O}_n^x \cap \mathcal{B}_n^x) \leq \mathbb{P}(\mathcal{B}_n^x \cap \{T < \infty\}) \leq (1 - \varepsilon)^{W/2}.$$

Then, by choosing the constant C_1 (defined in the beginning of the proof) large enough, we can make the previous bound $o(n^{-1})$. As a consequence, we get

$$\mathbb{P}\left(\bigcup_{X \in \mathrm{Po}(\lambda) \cap \Lambda_n} \mathcal{O}_n^X \cap \mathcal{B}_n^X\right) = o(1).$$

Together with (4) this shows that, with probability going to 1 as $n \to \infty$, any connected component which is not connected to $\widehat{G}_n$ is such that it does not exit the box Λ_n^X for any of the vertices X in this component. To conclude, it suffices to observe that, by concentration of Poisson random variables, it is very unlikely that one of these boxes contains much more than $(\log n)^2$ points. Indeed, applying again Lemma 2.5 gives

$$\mathbb{P}\Big(\bigcup_{X \in \mathrm{Po}(\lambda) \cap \Lambda_n} \big\{|G_n \cap \Lambda_n^X| > 2\lambda(8C_1 \log n)^2\big\}\Big) \leq \mathbb{E}\Big[\sum_{X \in \mathrm{Po}(\lambda) \cap \Lambda_n} \mathbb{1}_{|G_n \cap \Lambda_n^X| > 2\lambda(8C_1 \log n)^2}\Big]$$

$$= \lambda \int_{\Lambda_n} \mathbb{P}\big(|G_n^x \cap \Lambda_n^x| > 2\lambda(8C_1 \log n)^2\big)\, dx \leq \lambda n \mathbb{P}\big(|G \cap \Lambda(4C_1 \log n)| > 2\lambda(8C_1 \log n)^2 - 1\big)$$

$$= o(1),$$

where we use Lemma 2.6 for the last equality. This concludes the proof of the proposition. $\qquad \square$

Remark 4.9. Apart from the size of the second-largest component in G_n, with minor modifications the proof also implies a stretched exponential decay of $|\mathcal{C}_1(\mathbf{0})|$ in the infinite-volume limit conditionally on $|\mathcal{C}_1(\mathbf{0})| < \infty$. As above, the key point is to establish an upper bound on the diameter of $\mathcal{C}_1(\mathbf{0})$. By similar arguments one may show that the event $\mathcal{E}_n$ that '$A(n, 2n) = \Lambda(2n) \setminus \Lambda(n)$ contains $\Theta(n)$ cycles in $G(\lambda', p)$ (with the notation from the previous proof) that are pairwise at distance at least 1 from each other which all participate in the infinite component of $G(\lambda', p)$' is satisfied with probability $1 - \exp(-\Omega_{\lambda,p}(n))$. Thus, partitioning $\mathbb{R}^2 \setminus \Lambda(1)$ into the annuli $(A(2^k, 2^{k+1}))_{k \geq 0}$ and applying the argument from the proof of Proposition 4.6 to each of them shows that, conditionally on $\mathcal{E}_{2^k}$, the event $\mathbf{0} \overset{G(\lambda, p)}{\longleftrightarrow} (\mathbb{R}^2 \setminus \Lambda(2^k))$ and simultaneously $|\mathcal{C}_1(\mathbf{0})| < \infty$ holds with probability at most $\exp(-\Omega_{\lambda,p}(2^k))$. Hence, since $\Lambda(2^{\lfloor \log n \rfloor - 2})$ is included in the ball of radius $n/2$ around $\mathbf{0}$, the probability that the Euclidean diameter of $\mathcal{C}_1(\mathbf{0})$ is at least n is at most

$$\mathbb{P}(\mathcal{C}_1(\mathbf{0}) \cap (\mathbb{R}^2 \setminus \Lambda(2^{\lfloor \log n \rfloor - 2})) \neq \emptyset) \leq \mathbb{P}(\mathcal{C}_1(\mathbf{0}) \cap (\mathbb{R}^2 \setminus \Lambda(2^{\lfloor \log n \rfloor - 2})) \neq \emptyset \mid \mathcal{E}_{2^{\lfloor \log n \rfloor - 1}}) + \mathbb{P}(\overline{\mathcal{E}_{2^{\lfloor \log n \rfloor - 1}}})$$

$$= \exp(-\Omega_{\lambda,p}(n)).$$

4.3 The convergence of $(n^{-1}L_1(G_n))_{n \geq 1}$.

First, we recall that the convergence in probability of $(n^{-1}L_1(G_n))_{n \geq 1}$ to $\lambda\theta(\lambda, p)$ is proved in [19]. On the other hand, $L_1(G_n)$ is bounded by the total number of points of $\mathrm{Po}(\lambda)$ in Λ_n (which is distributed as a Poisson random variable with parameter λn), so this sequence is also bounded in L^p for all $p \geq 1$ (e.g. as a consequence of Lemma 2.6). Therefore, the convergence to $\lambda\theta(\lambda, p)$ holds in fact in L^p for all $p \geq 1$.

We now prove the almost sure convergence by using the results of Section 4.2. Fix $\delta < 1/2$, and for any large integer n, let $\mathcal{T}_n$ be a tessellation of Λ_n into $k := \lfloor n^{(1-\delta)/2} \rfloor^2$ squares with volume roughly n^δ. Then, consider the event $\mathcal{A}_n$ that:

- in each square in $\mathcal{T}_n$, the largest component has a size of order at least cn^δ, where $c > 0$ is some sufficiently small constant, and any other component has diameter at most $(\log n)^2$;

- in Λ_n, the second-largest component has size at most $(\log n)^3$.

The proof of Proposition 4.6 shows that, if c is sufficiently small, $\mathcal{A}_n$ holds with probability $1 - o(n^{-2})$.

Moreover, observe that on the event $\mathcal{A}_n$, $L_1(G_n)$ is equal to the sum of k independent terms $(X_i)_{i=1}^k$, all distributed as $L_1(G_{n^\delta})$, up to some error term, which is due to the components lying in the region R of all points in Λ_n at distance at most $(\log n)^2$ from the boundaries of the squares of $\mathcal{T}_n$. For all $i \in \{1, \ldots, k\}$, let us denote $Y_i = \min\{n^{-\delta}X_i, 2\}$, and $Y = Y_1 + \ldots + Y_k$. Then, since $X_i \geq cn^\delta$ on the event $\mathcal{A}_n$, $\mathbb{E}Y_i = \Theta(1)$ for any $i \in \{1, \ldots, k\}$. Therefore, by Hoeffdings's inequality (which is a version of Chernoff's inequality for bounded random variables, see for example [13]), $\mathbb{P}(|Y - \mathbb{E}Y| \geq \frac{k}{\log n}) = \exp(-n^{\Omega(1)})$. At the same time, $Y_i = n^{-\delta}X_i$ on the event that the i-th square in the tessellation contains at most $2n^\delta$ points, which happens with probability at least $1 - \exp(-n^{\Omega(1)})$ (see Lemma 2.6), and in particular $\mathbb{E}Y = n^{-\delta}\mathbb{E}X + o(1)$. As a consequence,

$$\mathbb{P}\left(|X - \mathbb{E}X| \geq \frac{n}{\log n}\right) \leq \mathbb{P}\left(|Y - \mathbb{E}Y| \geq \frac{k}{2\log n}\right) + \sum_{i=1}^k \mathbb{P}(Y_i \neq n^{-\delta}X_i)$$

$$\leq \mathbb{P}\left(|Y - \mathbb{E}Y| \geq \frac{k}{2\log n}\right) + k\exp(-\Omega(n^\delta)) = \exp(-n^{\Omega(1)}).$$

On the other hand, the event $\mathcal{D}_n$ that there are at most $n^{1-\delta/2}(\log n)^3$ points in R (which has area $O(n^{1-\delta/2}(\log n)^2))$ holds with probability $1 - o(n^{-2})$.

Finally, together with the fact that

$$\frac{1}{n}\mathbb{E}[X] \to \lambda\theta(\lambda) \text{ as } n \to \infty,$$

we deduce that,

$$\mathbb{P}\left(|L_1(G_n) - \mathbb{E}L_1(G_n)| \geq \frac{n}{2\log n}\right) \leq \mathbb{P}\left(|X - \mathbb{E}[X]| \geq \frac{n}{\log n}\right) + \mathbb{P}(\overline{\mathcal{A}_n}) + \mathbb{P}(\overline{\mathcal{D}_n}) = o(n^{-2}),$$

and the fact that $\sum_{n \geq 1} n^{-2} < \infty$ together with the Borel-Cantelli lemma concludes the proof of the almost sure convergence.

5 Proof of Theorem 1.4

We first prove that Bernoulli bond percolation preserves local convergence. For completeness, we also include the case of site percolation as it may be of independent interest.

18

Lemma 5.1. *Let $(\mathcal{G}_k)_{k\geq 1}$ be a sequence of graphs of geometric type converging locally to $\mathcal{G}$ in the sense of Definitions 1.2 and 1.3. For $p \in [0,1]$, consider the subgraphs $\mathcal{G}_k^b(p)$ and $\mathcal{G}_k^s(p)$ obtained after performing respectively Bernoulli bond and site percolation on $\mathcal{G}_k$, and similarly for $\mathcal{G}^b(p)$ and $\mathcal{G}^s(p)$. Then, these graphs are of geometric type and moreover $(\mathcal{G}_k^b(p))_{k\geq 1}$ and $(\mathcal{G}_k^s(p))_{k\geq 1}$ converge locally respectively to $\mathcal{G}^b(p)$ and $\mathcal{G}^s(p)$.*

Proof. The fact that a graph of geometric type remains in this class after bond or site percolation is immediate from the definition. Now, assume that $(\mathcal{G}_k)_{k\geq 1}$ converges locally to $\mathcal{G}$, and let $p \in [0,1]$ be fixed. Let us consider first the case of bond percolation, which is slightly easier. Fix Ω to be some bounded domain of $\mathbb{R}^2$, and let H be some finite graph. Note that if H is obtained from another graph H' after bond percolation, then H' belongs to the finite set $\mathcal{A}(H)$ of finite graphs with same vertex set as H and containing H as a subgraph. Therefore, denoting by $E(H)$ the set of edges of a graph H, and taking advantage of the fact that $\mathcal{A}(H)$ is a finite set, we get

$$\lim_{k\to\infty} \mathbb{P}\big(\mathcal{G}_k^b(p)_{|\Omega} = H\big) = \lim_{k\to\infty} \sum_{H'\in\mathcal{A}(H)} \mathbb{P}\big((\mathcal{G}_k)_{|\Omega} = H'\big)p^{|E(H)|}(1-p)^{|E(H')|-|E(H)|}$$

$$= \sum_{H'\in\mathcal{A}(H)} \mathbb{P}\big(\mathcal{G}_{|\Omega} = H'\big)p^{|E(H)|}(1-p)^{|E(H')|-|E(H)|} = \mathbb{P}\big(\mathcal{G}^b(p)_{|\Omega} = H\big).$$

We now consider the case of site percolation. The only additional difficulty is that now the set $\mathcal{B}(H)$ of graphs which can give rise to a fixed graph H after site percolation is infinite. However, denoting by $V(H)$ the set of vertices of a graph H, and using Fatou's lemma, we deduce

$$\liminf_{k\to\infty} \mathbb{P}\big(\mathcal{G}_k^s(p)_{|\Omega} = H\big) \geq \sum_{H'\in\mathcal{B}(H)} \liminf_{k\to\infty} \mathbb{P}\big((\mathcal{G}_k)_{|\Omega} = H'\big)p^{|V(H)|}(1-p)^{|V(H')|-|V(H)|}$$

$$= \sum_{H'\in\mathcal{B}(H)} \mathbb{P}\big(\mathcal{G}_{|\Omega} = H'\big)p^{|V(H)|}(1-p)^{|V(H')|-|V(H)|} = \mathbb{P}\big(\mathcal{G}^s(p)_{|\Omega} = H\big).$$

Since this holds for any finite graph H, by another application of Fatou's lemma, we infer

$$1 = \liminf_{k\to\infty} \sum_{H \text{ finite}} \mathbb{P}\big(\mathcal{G}_k^s(p)_{|\Omega} = H\big) \geq \sum_{H \text{ finite}} \liminf_{k\to\infty} \mathbb{P}\big((\mathcal{G}_k)_{|\Omega} = H\big) \geq \sum_{H \text{ finite}} \mathbb{P}\big(\mathcal{G}^s(p)_{|\Omega} = H\big) = 1.$$

Hence, for any finite graph H, $\lim_{k\to\infty} \mathbb{P}\big(\mathcal{G}_k^s(p)_{|\Omega} = H\big) = \mathbb{P}\big(\mathcal{G}^s(p)_{|\Omega} = H\big)$, which concludes the proof of the lemma. $\qquad\square$

Now, we prove Theorem 1.4. Let $\lambda > \lambda_0$, and $K \geq 0$ be given. Let also $(\mathcal{G}_k)_{k\geq 1}$ be a sequence of K-dependent graphs of geometric type converging to $G(\lambda, 1)$. We need to prove that $\limsup_{k\to\infty} p_c(\mathcal{G}_k) \leq p_c(\lambda)$. To this end, we fix some $p > p_c(\lambda)$ and show that, with the notation from Lemma 5.1, for all k sufficiently large, $\mathcal{G}_k^b(p)$ contains an infinite connected component.

The proof relies on a finite-size criterion. Consider a tessellation $\mathcal{T}$ of $\mathbb{R}^2$ into squares of side length $\sqrt{n}$, where n is a constant to be chosen sufficiently large later. For each square Q in $\mathcal{T}$, consider a partition of the square into 4 smaller squares of side length $\sqrt{n}/2$, say $\{Q_i\}_{1\leq i\leq 4}$. Then, for a random graph with vertex set embedded in $\mathbb{R}^2$, consider the event $\mathcal{A}_Q$ that the following holds:

- The second-largest component in Q has size at most $\frac{\lambda\theta(\lambda,p)}{10}n$.

- The largest components in each of the squares Q_i, for $i \in \{1, 2, 3, 4\}$, have size at least $\frac{\lambda\theta(\lambda,p)}{8}n$.

Also, for two squares Q and Q' of $\mathcal{T}$ sharing a common edge, consider the square Q'' of side length $\sqrt{n}$ in-between Q and Q', which is the union of the two smaller squares of Q touching Q' together with the two smaller squares of Q' touching Q. Then, define

$$\mathcal{A}_{Q,Q'} := \mathcal{A}_Q \cap \mathcal{A}_{Q'} \cap \mathcal{A}_{Q''}.$$

19

On the one hand, on $\mathcal{A}_{Q,Q'}$, the two largest components of Q and Q' are connected by construction. On the other hand, by Theorem 1.5 (or by the results of [19]) we know that for every $\delta > 0$ and every sufficiently large n, $\mathbb{P}_{\lambda,p}(\mathcal{A}_{Q,Q'}) > 1 - \delta$. Now, with a slight abuse of notation, denote by $\mathbb{P}_{k,p}$ the distribution of $\mathcal{G}_k^b(p)$. Then, by definition of the local convergence, Lemma 5.1 and the fact that $\mathcal{A}_{Q,Q'}$ is a local event one has that, for any sufficiently large k and any two squares Q and Q' sharing a common edge, $\mathbb{P}_{k,p}(\mathcal{A}_{Q,Q'}) \geq 1 - 2\delta$. On the other hand, considering the graph isomorphic to $\mathbb{Z}^2$ with vertex set the squares in $\mathcal{T}$, and declaring the edge between two neighboring squares Q and Q' open if the event $\mathcal{A}_{Q,Q'}$ holds, defines a bond percolation process on $\mathbb{Z}^2$. For every $n \geq K$, by hypothesis, this bond percolation is 1-dependent for each of the graphs $\mathcal{G}_k^b(p)$. Therefore, by Theorem 2.2, choosing δ small enough (and then n large enough) allows us to make this graph supercritical for any sufficiently large k, and thus proving that all the graphs $\mathcal{G}_k^b(p)$ have an infinite connected component.

The last part of the theorem is immediate since the critical bond percolation is a decreasing function for the partial order given by the inclusion of graphs. This concludes the proof of Theorem 1.4. □

References

[1] D. Ahlberg, V. Tassion, and A. Teixeira. Sharpness of the phase transition for continuum percolation in $\mathbb{R}^2$. *Probability Theory and Related Fields*, 172(1):525–581, 2018.

[2] D. Aldous and R. Lyons. Processes on unimodular random networks. *Electron. J. Probab.*, 12:1454–1508, 2007.

[3] I. Benjamini, A. Nachmias, and Y. Peres. Is the critical percolation probability local? *Probab. Theory and Related Fields*, 149:261–269, 2011.

[4] I. Benjamini and O. Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6:13pp, 2001.

[5] S. R. Broadbent and J. M. Hammersley. Percolation processes: I. crystals and mazes. In *Mathematical proceedings of the Cambridge philosophical society*, volume 53, pages 629–641. Cambridge University Press, 1957.

[6] L. Chayes and R. H. Schonmann. Mixed percolation as a bridge between site and bond percolation. *The Annals of Applied Probability*, 10(4):1182–1196, 2000.

[7] H. Duminil-Copin. Introduction to Bernoulli percolation. *Lecture notes available on the webpage of the author*, 2018.

[8] H. Duminil-Copin and V. Tassion. A note on Schramm's locality conjecture for random cluster models. *arXiv:1707.07626*, 2017.

[9] M. Franceschetti, M. D. Penrose, and T. Rosoman. Strict inequalities of critical values in continuum percolation. *Journal of Statistical Physics*, 142(3):460–486, 2011.

[10] E. N. Gilbert. Random plane networks. *Journal of the society for industrial and applied mathematics*, 9(4):533–543, 1961.

[11] G. R. Grimmett and A. M. Stacey. Critical probabilities for site and bond percolation models. *Ann. Probab.*, 26(5):1788–1812, 1998.

[12] J. M. Hammersley. Comparison of atom and bond percolation processes. *Journal of Mathematical Physics*, 2(5):728–733, 1961.

[13] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.

[14] T. Hutchcroft. Locality of the critical probability for transitive graphs of exponential growth. *Ann. Probab.*, 48:1352–1371, 2020.

[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] S. Martineau and V. Tassion. Locality of percolation for Abelian Cayley graphs. *Ann. Probab.*, 45:1247–1277, 2017.

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

[18] M. Penrose. *Random geometric graphs*, volume 5. OUP Oxford, 2003.

[19] M. Penrose. Giant component of the soft random geometric graph. *arXiv:2204.10219*, 2022.

[20] M. D. Penrose and A. Pisztora. Large deviations for discrete and continuous percolation. *Advances in Applied Probability*, 28(1):29–52, 1996.