Commun. Math. Phys. 168, 39–55 (1995)

Communications in Mathematical Physics

Simultaneous Uniq^{H}ueness of Infinite Clusters in Stationary Random Labeled Graphs

Kenneth S. Alexander[1]

Department of Mathematics, University of Southern California, Los Angeles, CA 90089-1113, USA. Email: alexandr@mtha.usc.edu

Received: 12 October 1993/in revised form: 19 April 1994

Abstract: In processes such as invasion percolation and certain models of continuum percolation, in which a possibly random label $f(b)$ is attached to each bond b of a possibly random graph, percolation models for various values of a parameter r are naturally coupled: one can define a bond b to be occupied at level r if $f(b) \leqq r$. If the labeled graph is stationary, then under the mild additional assumption of positive finite energy, a result of Gandolfi, Keane, and Newman ensures that, in lattice models, for each fixed r at which percolation occurs, the infinite cluster is unique a.s. Analogous results exist for certain continuum models. A unifying framework is given for such fixed-r results, and it is shown that if the site density is finite and the labeled graph has positive finite energy, then with probability one, uniqueness holds simultaneously for all values of r. An example is given to show that when the site density is infinite, positive finite energy does not ensure uniqueness, even for fixed r. In addition, with finite site density but without positive finite energy, one can have fixed-r uniqueness a.s. for each r, yet not have simultaneous uniqueness.

I. Introduction and Statement of Results

There are various models in which percolation processes are naturally coupled for all values of the order parameter. Typically, a value $f(b)$ is attached to each bond b of an infinite graph $(V, \mathscr{B})$ with site set V and bond set $\mathscr{B}$; the graph and/or the values $f(b)$ may be random. A bond b is said to be *occupied at level* r if $f(b) \leqq r$, and one can consider percolation of occupied bonds at various levels r. Some examples follow.

Example 1.1. In invasion percolation, introduced in the mathematical literature in [6], $(V, \mathscr{B})$ is a (nonrandom) lattice in $\mathbb{R}^d$, and the values $\{f(b)\colon b \in \mathscr{B}\}$ are iid uniform in $[0,1]$. The corresponding percolation model is Bernoulli bond percolation.

[1] Research supported by NSF grant DMS-9206139

Example 1.2. In one example of the "lily pad" model of continuum percolation, the sites are those of a Poisson process of constant intensity in $\mathbb{R}^d$, and connections exist between pairs of sites separated by distance at most some $r > 0$; see [9] Chapter 10. This is equivalent to centering a ball of radius $r/2$ at each site, and considering the connected components of the union of such balls. The clusters formed by the connections are not affected if one considers only pairs of sites which are adjacent in the Voronoi graph, obtained by placing a bond between every pair of sites for which the corresponding polyhedral Voronoi regions $R_x := \{z \in \mathbb{R}^d: d(z,x) = \min\{d(z,y): y \in V\}\}$, $x \in V$, have a face in common (see [17] Lemma 6.2). Thus we can take $(V, \mathscr{B})$ to be either the Voronoi graph or the complete graph on the set of sites. $f(b)$ is the Euclidean length of a bond b; the models for all values of r are then naturally coupled.

Example 1.3. In the "random connection model" of continuum percolation, examined in [16] and [5], the sites are those of a Poisson process of intensity λ in $\mathbb{R}^d$, and there is a nonincreasing right-continuous function $\varphi: (0, \infty) \to [0,1]$ such that for each pair of sites x, y, the bond $\langle x, y\rangle$ is occupied with probability $\varphi(|x-y|)$, independently of other bonds and of the site configuration. It is typically assumed that

$$\int_0^\infty \varphi(s) s^{d-1} ds < \infty ,$$

to ensure that the graph of occupied bonds is locally finite. The intensity λ is the order parameter, or equivalently, instead of varying λ one can fix the intensity, say at 1, choose a value $r > 0$ and occupy each bond $\langle x, y\rangle$ with probability $\varphi(|x-y|/r)$. Let $\varphi^{-1}(t) := \inf\{x \geqq 0; \varphi(x) \leqq t\}$, let $\{Y_{\langle x,y\rangle}: x, y \in V\}$ be iid uniform in $[0,1]$, and let $f(\langle x, y\rangle) := |x-y|/\varphi^{-1}(Y_{\langle x,y\rangle})$. Note $f(\langle x, y\rangle) \leqq r$ if and only if $Y_{\langle x,y\rangle} \leqq \varphi(|x-y|/r)$, which has probability $\varphi(|x-y|/r)$. Thus this process couples all values of r, or equivalently all values of the intensity λ. When r, rather than λ, is the order parameter, the site process can be an arbitrary stationary point process, as in [5].

Example 1.4. In a more general version of the lily pad model of Example 1.2, the ball centered at each site x of a stationary point process has a random radius ρ_x; these radii are iid. For a Poisson process, as in Example 1.3 the order parameter is the intensity λ. Equivalently, and for general stationary point processes, we can fix the point process, choose a value $r > 0$ and for each pair of sites x, y declare the bond $\langle x, y\rangle$ to be occupied if $\rho_x + \rho_y \geqq |x-y|/r$. This means that the label $f(\langle x, y\rangle) = |x-y|/(\rho_x + \rho_y)$.

Example 1.5. Let $(V, \mathscr{B})$ be the d-dimensional hypercubic lattice. Grimmett [10] has shown that for fixed $q > 1$, there exists a stationary process $\{f(b): b \in \mathscr{B}\}$ of $[0,1]$-valued random variables such that for each $r \in [0,1]$, $\{b \in \mathscr{B}: f(b) \leqq r\}$ is the set of occupied bonds of the wired Fortuin–Kastelyn random cluster model [7] with parameters (r, q). The parameters here are such that the probability of a given configuration in a finite box is proportional to the independent-percolation probability at level r, multiplied by q^n, where n is the number of clusters in the configuration; see [7] for the full definition of the process.

Properties of such coupled percolation models are relevant to the study of infinite analogs of the minimal spanning tree; see [2] and [3]. In particular, as implicitly observed in the final section of [2], it is of interest to know whether uniqueness of the infinite cluster a.s. holds simultaneously for all values of r for

which percolation occurs. Here we will answer this question in the affirmative; this result is applied in [3]. A result of Gandolfi, Keane, and Newman [8] ensures that for lattice models with what those authors call "positive finite energy," which include Examples 1.1 and 1.5, for each fixed r, there is a.s. at most one infinite cluster. Their proof does not yield simultaneous uniqueness. Meester and Roy [13] and Burton and Meester [5] proved similar results for the "lily pad" model of Examples 1.2 and 1.4, and for the random connection model of Example 1.3, respectively. Our goals here are twofold: first, find a unifying framework, including in particular an appropriate definition of positive finite energy, which will encompass all of the above examples; second, show that the uniqueness does hold simultaneously.

Let us first describe the sort of process to which our result will apply. The underlying space is either $\mathbb{R}^d$ or a lattice L in $\mathbb{R}^d$, by which we mean a periodic graph as defined in [12]. By a *counting measure* we mean a sum of at most countably many unit point masses at distinct points. Let $\mathbb{R}^d \wedge \mathbb{R}^d$ denote the set of pairs $\langle x, y\rangle$, $x, y \in \mathbb{R}^d$, with $\langle x, y\rangle$ and $\langle y, x\rangle$ identified. A *random labeled graph* X in $\mathbb{R}^d$ can be viewed as a point process in the disjoint union $\mathbb{S}_d := \mathbb{R}^d \cup ((\mathbb{R}^d \wedge \mathbb{R}^d) \times \mathbb{R})$. More precisely, X is a random element of the space $\mathcal{M}(\mathbb{S}_d)$ of Radon measures on $\mathbb{S}_d$, defined on a probability space $(\Omega, \mathcal{F}, \Pr)$, with the following properties a.s.:

(i) X is a counting measure:
(ii) the sets $V = V(X) := \{x \in \mathbb{R}^d : X(\{x\}) = 1\}$ and $\mathcal{B} = \mathcal{B}(X) := \{\langle x, y\rangle \in \mathbb{R}^d \wedge \mathbb{R}^d : X(\{(\langle x, y\rangle, r)\}) = 1 \text{ for some } r \in \mathbb{R}\}$ are such that if $\langle x, y\rangle \in \mathcal{B}$ then $x, y \in V$;
(iii) If $X(\{(\langle x, y\rangle, r)\}) = 1$ and $X(\{(\langle x, y\rangle, s)\}) = 1$ then $r = s$.

Elements of V are sites (or vertices) of X, and elements of $\mathcal{B}$ are bonds (or edges.) We call a fixed Radon measure in $\mathcal{M}(\mathbb{S}_d)$ a *labeled graph* if it has the properties prescribed for X in (i)–(iii). $\mathcal{M}(\mathbb{S}_d)$ is endowed with the topology of weak convergence and the σ-algebra $\mathcal{S}_d$ of corresponding Borel sets. Because of (iii), given X there is a labeling function $f: \mathcal{B} \to \mathbb{R}$ such that $X(\{(\langle x, y\rangle, r)\}) = 1$ if and only if $f(\langle x, y\rangle) = r$. X can then be identified with the triple $(V, \mathcal{B}, f)$.

At times we will consider unlabeled graphs, in which case we will use the same terminology without the word "labeled."

Note that X being a Radon measure means that there are only finitely many sites in each bounded region, though infinitely many bonds may emanate from a single site, and the mean number of sites per unit volume may be infinite. Our definition allows bonds of form $\langle x, x\rangle$, but multiple bonds between the same pair of sites are not possible.

For $t \in \mathbb{R}^d$ the transformation τ_t acts on $\mathbb{S}_d$ by

$$\tau_t x := x - t \quad \text{for } x \in \mathbb{R}^d ;$$

$$\tau_t(\langle x, y\rangle, r) := (\langle x - t, y - t\rangle, r) \quad \text{for } (\langle x, y\rangle, r) \in (\mathbb{R}^d \wedge \mathbb{R}^d) \times \mathbb{R} .$$

This induces a transformation θ_t which acts on $\mathcal{M}(\mathbb{S}_d)$ by

$$(\theta_t m)(B) := m(\tau_t B), \quad B \subset \mathbb{S}_d .$$

We say X is *stationary* if the distribution of X is preserved by θ_t for all $t \in \mathbb{R}^d$. When the underlying space is a lattice L in $\mathbb{R}^d$, these definitions remain the same but with

$\mathbb{R}^d$ replaced by L, and with the added requirement that all bonds in $\mathscr{B}$ be bonds of L; τ_t, θ_t are defined only for sites t of L. If X is stationary, we may assume the following: there is a family H of ergodic stationary probability measures on $\mathscr{M}(\mathbb{S}_d)$, endowed with the topology of weak convergence and with a corresponding class $\mathscr{H}$ of Borel sets, and a probability measure μ on $(H, \mathscr{H})$; X is the first coordinate of a random element $(X, \tilde{P})$ of $\mathscr{M}(\mathbb{S}_d) \times H$ such that $\tilde{P}$ has distribution μ; the probability space is

$$(\Omega, \mathscr{F}, \Pr) = (\mathscr{M}(\mathbb{S}_d) \times H, \mathscr{S}_d \times \mathscr{H}, \Pr) ,$$

where

$$\Pr(C) := \int_H \int_{\mathscr{M}(\mathbb{S}_d)} I_C \, dv \, d\mu(v)$$

and I_C denotes the indicator of an event C, so that

$$\Pr(X \in A | \tilde{P}) = \tilde{P}(A) \, \mu\text{-a.s., for all } A \subset \mathscr{M}(\mathbb{S}_d) .$$

We tacitly here and henceforth restrict to measurable A. We let P denote the distribution of X, so

$$P(A) = \int_H v(A) \, d\mu(v) \quad \text{for all } A \subset \mathscr{M}(\mathbb{S}_d) .$$

A general nonrandom element of H (that is, a possible value of $\tilde{P}$) is denoted v, and a general nonrandom element of $\mathscr{M}(\mathbb{S}_d)$ (that is, a possible value of X) is denoted m.

For each such stationary probability measure P on $\mathscr{M}(\mathbb{S}_d)$ there is a σ-finite *Palm measure* P_0 on $\mathscr{M}(\mathbb{S}_d)$ with the following property. Let λ denote Lebesgue measure and for each $B \subset \mathbb{R}^d$, $m \in \mathscr{M}(\mathbb{S}_d)$ and $A \subset \mathscr{M}(\mathbb{S}_d)$,

$$N_A(m, B) := |\{s \in V(m) \cap B : \theta_s m \in A\}| ;$$

then for bounded B,

$$P_0(A) = \lambda(B)^{-1} \int_{\mathscr{M}(\mathbb{S}_d)} N_A(m, B) \, dP(m) . \tag{1.1}$$

That is, if we shift the origin successively to each graph site in B, and check for each resulting configuration whether A occurs, then $P_0(A)$ is the mean number of occurrences per unit area, which does not depend on B. This is a slight variation on the usual context for defining Palm measures, as we consider only point mass locations in $\mathbb{R}^d$ and not in all of $\mathbb{S}_d$, but the proof of (1.1) in [14] (Theorem II.4) goes through here. For $0 \leqq n \leqq \infty$, let Λ_n denote $[-n, n]^d$; for finite n we can also define random measures $\hat{P}_n(X, \cdot)$ on $\mathscr{M}(\mathbb{S}_d)$ by

$$\hat{P}_n(X, A) := \lambda(\Lambda_n)^{-1} N_A(X, \Lambda_n) .$$

The integer lattice divides $\mathbb{R}^d$ into unit cubes $z + [0,1)^d$, $z \in \mathbb{Z}^d$, and we can set $Y_z := N_A(X, z + [0,1)^d)$ and obtain from the ergodic theorem that

$$\begin{aligned} \hat{P}_n(X, A) &= |\Lambda_n \cap \mathbb{Z}^d|^{-1} \sum_{z \in \Lambda_n \cap \mathbb{Z}^d} Y_z \\ &\to E_{\tilde{P}}(Y_0) = \tilde{P}_0(A) = E_{\tilde{P}} \hat{P}_k(\cdot, A) \text{ as } n \to \infty \text{ a.s. (Pr)} \\ &\quad \text{for each } A \subset \mathscr{M}(\mathbb{S}_d) \text{ and } k \geqq 1; \end{aligned} \tag{1.2}$$

here E_ν denotes expected value with respect to the measure ν. This means that, letting

$$H_k := \{\nu \in H\colon E_\nu \hat{P}_k(\cdot, A) = 0\}\,,$$

we have

$$\begin{aligned}\Pr(\hat{P}_n(X,A) \to 0,\, N_A(X,\mathbb{R}^d) \neq 0) &\leqq \limsup_k \Pr(E_{\hat{P}}\hat{P}_k(\cdot,A) = 0, N_A(X,\Lambda_k) \neq 0)\\ &= \limsup_k \int_{H_k} \nu(N_A(X,\Lambda_k) \neq 0)\, d\mu(\nu)\\ &= 0, \text{ for all } A \subset \mathscr{M}(\mathbb{S}_d)\,. \end{aligned} \tag{1.3}$$

When it exists, we call $\lim_n \hat{P}_n(X,A)$ the *density* of A (in X), or the *density* of the set $\{v \in V(m)\colon \theta_v m \in A\}$. Thus, roughly, events which occur with 0 density actually never occur. It is also important for our purposes that there is a measure, not just a finitely additive function, which for each fixed A a.s. coincides with the density; this means that the density of the limit is a.s. equal to the limit of the densities, for increasing families of events. All of this is of course well-known for stationary point processes in $\mathbb{R}^d$; we have included it here to show that the extra structure (bonds and labels) present here does not alter these facts.

We call $\lambda_Y(X) := \lim_n \hat{P}_n(X, \mathscr{M}(\mathbb{S}_d))$ the *site density* of X.

Instead of shifting the origin to each graph site, we will at times wish to shift it to every possible point, or to each site of some multiple of the integer lattice. Thus define

$$N'_A(X,n,t) := |\{z \in \Lambda_n \cap \mathbb{Z}^d\colon \theta_{2tz}X \in A\}|, \quad 0 \leqq n \leqq \infty\,,$$

$$\hat{P}_{n,t}(X,A) := |\Lambda_n \cap \mathbb{Z}^d|^{-1} N'_A(X,n,t), \quad 0 \leqq n < \infty\,,$$

$$\hat{P}''_t(X,A) := |\Lambda_t|^{-1} \int_{\Lambda_t} I_{[\theta_s X \in A]}\, ds, \quad t > 0\,,$$

$$\hat{P}''_\infty(X,A) := \lim_t \hat{P}''_t(X,A), \quad \text{if this limit exists.}$$

As in (1.2) and (1.3) we have

$$\Pr(\hat{P}'_{n,t}(X,A) \to 0 \text{ as } n \to \infty,\, N'_A(X,\infty,t) \neq 0) = 0 \quad \text{for all } A \subset \mathscr{M}(\mathbb{S}_d) \text{ and } t > 0\,. \tag{1.4}$$

An *r-cluster* is a connected component of the graph denoted $X_{\leqq r}$ which has site set V and bond set $\{b \in \mathscr{B}\colon f(b) \leqq r\}$. We say *percolation occurs* at level r in X if there is an infinite r-cluster, and set

$$r_c(X) := \inf\{r\colon \text{percolation occurs in } X \text{ at level } r\}\,.$$

We say *uniqueness holds* at level r in X if there is at most one infinite r-cluster. The a.s. uniqueness results of [4, 5, 8 and 13] are valid for each fixed level r in certain categories of stationary random labeled graphs. Our first goal is to extend those results to our unified context. A major ingredient is an appropriate definition of

"positive finite energy." Loosely, the idea is that conditioning on the configuration outside some Λ_t cannot rule out the connecting together of all the sites inside Λ_t. But there are more complications than in the nearest-neighbor lattice case of [4], or the long-range lattice case of [8], principally for three reasons: (1) bonds can cross $\partial\Lambda_t$ and connect to arbitrarily distant sites; (2) in models like Example 1.4, connecting sites inside Λ_t together may result in the creation of additional bonds which cross $\partial\Lambda_t$; (3) in continuum models one cannot in general make sense of conditioning on "the entire configuration except one bond."

In particular in our definition we cannot simply condition on what is strictly outside Λ_t, or on what is outside Λ_t together with *all* bonds which cross $\partial\Lambda_t$. Instead, we want to condition on what is outside Λ_t, together with the existence of *at least* certain bonds crossing $\partial\Lambda_t$. This motivates the following preliminaries to our definition. Given a labeled graph $m \in \mathscr{M}(\mathbb{S}_d)$ or the equivalent triple $(V(m), \mathscr{B}(m), f_m)$, we define the *restriction* $m|_B$ of m to B to be the labeled graph with

$$V(m|_B) = V(m) \cap B, \quad \mathscr{B}(m|_B) = \{\langle x, y\rangle \in \mathscr{B}(m)\colon x, y \in B\}\ ,$$

and with label function $f_{m|_B}$ the restriction of f_m to $\mathscr{B}(m|_B)$. We call an infinite component C of a graph a *galaxy* if C includes a finite, nonempty set of sites, called a *core*, such that removing all these sites, and all bonds emanating from them, from C decomposes C into infinitely many disjoint components, all of which are finite. Later we will show galaxies a.s. do not occur, but for now, let $t > 0$ and $q \in \mathbb{Q}$, and consider the event $J := [r_c(X) < q < r_c(X|_{\Lambda_t^c})]$. If the only infinite q-cluster is a galaxy with core in Λ_t, then J can occur. However, it is easily verified that, since each bounded region contains only finitely many sites, there can be at most one $z \in \mathbb{Z}^d$ for which percolation occurs at level q in $\theta_{2tz}X$ but not in $(\theta_{2tz}X)|_{\Lambda_t^c}$, and hence at most one for which $\theta_{2tz}X \in J$. It follows from (1.4) that $P(J) = 0$, and hence that

$$r_c(X) = r_c(X|_{\Lambda_t^c}) \quad \text{P-a.s .}$$

Further, for each $v \in H$ there is a critical point $\rho_c(v)$ such that

$$r_c(X) = \rho_c(\tilde{P}) \quad \text{Pr-a.s .}$$

Applying similar reasoning to the value $\rho_c(\tilde{P})$ in place of q we see that with probability one, for all $r \in \mathbb{R}$ and $t > 0$,

percolation occurs at level r in X if and only if percolation occurs at level r in $X|_{\Lambda_t^c}$. (1.5)

Next we need the concept of a δ-approximate connection event from a bounded set Γ to a set Λ in $\mathbb{R}^d$. Let $\mathscr{D}_\delta := \{\delta x + [0,\delta)^d\colon x \in \mathbb{Z}^d\}$. Suppose $A \subset \mathscr{M}(\mathbb{S}_d)$ is an event of the following form: $\mathscr{C}$ is a finite collection of cubes from $\mathscr{D}_\delta$, each contained in Γ; for each $Q \in \mathscr{C}$, $r_Q = r_Q(X|_{\Lambda^c})$ is a real-valued function of $X|_{\Lambda^c}$; A is the event that

(i) every cube in $\mathscr{C}$ contains exactly one site, and
(ii) for each $Q \in \mathscr{C}$ there is a bond at level r_Q between this site in Q and some site in Λ.

We then call A a *δ-approximate connection event* from Γ to Λ. Note that (a) there are only finitely many such events for fixed Γ, Λ, δ and functions $\{r_Q\colon Q \in \mathscr{D}_\delta\}$; (b)

the δ-approximate connection event does not rule out the existence of other bonds and sites in $\Gamma \cup A$ not specified by the event; and (c) $\mathscr{C}$ can be empty, in which case A is the whole space. The functions $r_Q(\cdot)$ are called the *connection levels* of the event A.

We can now give our definition: we say X (or P) has *positive finite energy* if for every $t > s > 0$, every $r > 0$, every $\delta > 0$ and every δ-approximate connection event A from $\Lambda_t \backslash \Lambda_s$ to Λ_s,

$$P(A \mid X|_{\Lambda_s^c}) > 0 \text{ and percolation at level } r \text{ in } X|_{\Lambda_s^c} \text{ imply}$$

$$P(A;\, X|_{\Lambda_s} \text{ is connected at level } r \mid X|_{\Lambda_s^c}) > 0, \quad P\text{-a.s.} \tag{1.6}$$

Here by a statement of the form "I and J imply K, P-a.s.," we mean $P(I \cap J \cap K^c) = 0$. The restriction to r for which percolation occurs is made because, in models like Example 1.2 or 1.4, if the sites come from a "hard sphere" point process, there may be values of r where no connection is possible in $X|_{\Lambda_s}$, regardless of $X|_{\Lambda_s^c}$. Note that (1.5) makes it easy to verify the percolation condition in (1.6).

It would actually be sufficient for out purposes to require in the definition of positive finite energy only that (1.6) be valid for sufficiently large s, where "sufficiently large" may depend on X, but lacking an example to motivate this added complexity, we will forego it.

The following theorem encompasses results in [4, 5, 8 and 13] in the case of finite site density. The proof we will give in Sect. II is based on that in [8] for lattice models.

Theorem 1.6. *Suppose X is a stationary random labeled graph, in $\mathbb{R}^d$ or in a lattice L, with positive finite energy and finite site density. Then for each $r \in \mathbb{R}$,*

$$P[X \text{ includes at most one infinite } r\text{-cluster}] = 1\,.$$

In Example 1.4 it is clear that the graph has positive finite energy if the support of the radius distribution is not bounded above; everything in Λ_s can then be connected by increasing the associated radii sufficiently for sites in or near Λ_s. Thus our theorem does contain the occupied-cluster result in [13], if the site density is finite; Roy and Meester [13] did not need to assume finite site density in some of their uniqueness results. Similar considerations apply to [5] and Example 1.3. But this is an artifact of the particular models being studied; Theorem 1.6 is not valid for general stationary random labeled graphs without the assumption of finite site density, as the following example shows. It is possible that some other assumption can substitute for finite site density in our results, such as perhaps finite degree of every site, but we have not investigated this.

Example 1.7. Let Y_n, $n \in \mathbb{Z}$, be an iid sequence of nonnegative integer valued r.v.'s with infinite mean, and $\{q_n, n \in \mathbb{Z}\}$ a sequence of nonnegative constants such that

$$\Sigma_n q_n = 1 \quad \text{and} \quad \Sigma_n q_n Y_n = \infty \text{ a.s.} \tag{1.7}$$

This will hold if, for example, $EY_0^{1/2} = \infty$ and $q_n \sim c/n^2$. For each n let V_n be a set of Y_n iid uniform points in $(n, n+1)$. We construct a stationary random graph in $\mathbb{R}$ with site set $V := \mathbb{Z} \cup \bigcup_{n \geq 1} V_n$.

Consider first a long-range bond percolation model on $\mathbb{Z}$ as follows: we choose $p_k \in [0, 1]$ for each $k \geq 0$, and for each $m \neq n \in \mathbb{Z}$ put an (occupied) bond between

m and n with probability $p_{|n-m|}$. Let us choose $\{p_k\}$ so that each p_k is strictly positive, but percolation a.s. does not occur. This is true if $\max_k p_k$ is small and p_k decreases like $k^{-\alpha}$ with $\alpha > 2$; see [1].

For our stationary random graph in $\mathbb{R}$, we will not use a realization of the above model, but only the values p_k with the prescribed properties. As a first stage, for each n and each $v \in V_n$, we select a single site of $\mathbb{Z}$ and connect v to it with an occupied bond, choosing site $n + k$ with probability q_k for all $v \in V_n$ and $k \in \mathbb{Z}$, independently for distinct v. We call the result the *first-stage graph* and denote it X_1. Every component of the first-stage graph clearly contains a single integer site, and we claim that every component is infinite a.s. By stationarity it is sufficient to show this for the component of 0. But

$$\begin{aligned} &\Sigma_n \Pr(\text{there is a bond in } X_1 \text{ from } (n, n+1) \text{ to } 0 \mid \{Y_n, n \in \mathbb{Z}\}) \\ &\quad = \Sigma_n (1 - (1 - q_n)^{Y_n}) \\ &\quad = \infty \text{ a.s.} \end{aligned}$$

by (1.7), so the claim follows from the Borel–Cantelli Lemma.

Next let C_n denote the component of n in X_1. As a second stage, for each m and n, including $m = n$, we order the bonds $\{\langle x, y\rangle : x \in C_m, y \in C_n, x \neq y\}$ by increasing value of $\max[d(x,\{m,n\}), d(y,\{m,n\}))]$, breaking ties using increasing value of $\min[d(x,\{m,n\}), d(y,\{m,n\}))]$; note that with probability one this does break all ties. Here d denotes Euclidean distance. For each m and n we then independently let the k^{th} bond in this ordering be occupied with probability $2^{-k} p_{|n-m|}$. We let X denote the resulting graph consisting of all sites of V and all occupied bonds. (All occupied bonds can be labeled with a fixed constant, say 0, to make a labeled graph.) For $m \neq n$ the probability that there is at least one occupied bond between C_m and C_n is then less than $p_{|n-m|}$, so comparison to the long-range model on $\mathbb{Z}$ shows that, with probability one, only finitely many integer sites are in each component, and there are therefore infinitely many infinite clusters. However, it is easy to see that X has finite energy.

Of course the graph X is invariant only under integer shifts, but we can shift the entire configuration by a random amount U uniform in $[0, 1]$ to obtain stationary random graph in $\mathbb{R}$. Alternatively we could replace $\mathbb{Z}$ with the set of sites of a Poisson process, and adjust V_n accordingly. □

The following improvement of Theorem 1.6 is our main result.

Theorem 1.8. *Suppose X is a stationary random labeled graph, in $\mathbb{R}^d$ or in a lattice L, with positive finite energy and finite site density. Then*

$$\Pr[X \text{ includes at most one infinite } r\text{-cluster for each } r \in \mathbb{R}] = 1 .$$

The following example shows that one can have finite site density and a.s. uniqueness for each fixed r, but yet not have simultaneous uniqueness. The graph does not have positive finite energy.

Example 1.9. Consider a random labeled graph in $\mathbb{R}^2$ with site set $\mathbb{Z}^2$. Vertical bonds are nearest-neighbor; horizontal bonds are long-range. Let us call the subgraph in the vertical line at m, *column* m. Let $\{Y_m : m \in \mathbb{Z}\}$ and $\{Z_{mnk} : m, n, k \in \mathbb{Z}, m \neq n\}$ be independent families of uniform $[0, 1]$ random variables. We label bonds

as follows: all vertical bonds in column m get label Y_m. Each horizontal bond $\langle (m, k), (n, k)\rangle$ gets label $\max(Y_m, Y_n) + Z_{mnk}$. Observe that as the level r increases, all of column m becomes occupied when r reaches Y_m, but no horizontal bond emanating from column m is occupied at level Y_m, so column m is a separate infinite cluster at level Y_m. But as soon as r strictly exceeds the level $\max(Y_m, Y_n)$ where columns m and n are both occupied, there are a.s. many bonds between column m and column n, since $\inf_k Z_{mnk} = 0$ a.s. Therefore, with probability one, at all levels $r \notin \{Y_m\colon m \in \mathbb{Z}\}$ there is at most one infinite cluster.

As in Example 1.7, one can shift the entire graph by an amount U uniform in $[0, 1]^2$ to obtain a stationary random labeled graph in $\mathbb{R}^2$.

Remark 1.10. One could as well define r-clusters using a strict inequality, that is, using bond set $\{b \in \mathscr{B}\colon f(b) < r\}$ instead of $\{b \in \mathscr{B}\colon f(b) \leqq r\}$. Theorems 1.6 and 1.8 remain valid with this change, with no material difference in the proofs.

II. Proofs

We let $\nu(f)$ denote $\int f d\nu$ for functions f and measures ν. We say an event $A \subset \mathscr{M}(\mathbb{S}_d)$ is *local* if there exists t such that $m \in A$ if and only if $m|_{\Lambda_t} \in A$. There exists a countable collection $\mathscr{A}$ of local events which generates the Borel σ-field in $\mathscr{M}(\mathbb{S}_d)$; we may take $\mathscr{A}$ to be a field. Let

$$\mathscr{M}_0 := \{m \in \mathscr{M}(\mathbb{S}_d)\colon \text{there exists } \nu \in H \text{ with } \hat{P}''_\infty(m, A) = \nu(A) \text{ for all } A \in \mathscr{A}\}$$

and for $m \in \mathscr{M}_0$ let ν_m denote the corresponding, necessarily unique, value of $\nu \in H$. From the ergodic theorem,

$$\Pr[X \in \mathscr{M}_0, \nu_X = \tilde{P}] = 1\,. \tag{2.1}$$

Fix some element $\bar{\nu}$ of H and for $m \notin \mathscr{M}_0$, let $\nu_m := \bar{\nu}$, to complete the definition. Since the events in $\mathscr{A}$ are local, letting

$$m^t := m|_{\Lambda_t^c}$$

we have

$$\nu_m = \nu_{m^t} \quad \text{for all } t \text{ and } m\,. \tag{2.2}$$

Let us verify that ν_m is a measurable function of m. It is sufficient to show that $\nu_m(\varphi)$ is measurable for each $\varphi \in C_b(\mathscr{M}(\mathbb{S}_d))$; for this it suffices to show that $\nu_m(A)$ is measurable for each Borel set $A \subset \mathscr{M}(\mathbb{S}_d)$. But this follows from the fact that the class of A for which $\nu_m(A)$ is a measurable function of m is a monotone class containing $\mathscr{A}$.

Our first lemma is an analog of Lemma 1 of [8].

Lemma 2.1. Suppose X is a stationary random labeled graph, in $\mathbb{R}^d$ or in a lattice L. For each $t > 0$ and each event $A \subset \mathscr{M}(\mathbb{S}_d)$, for μ-almost every $\nu \in H$,

$$P(A \mid X^t) = \nu(A \mid X^t) \quad \nu\text{-a.s.}$$

Proof. This proof would be quite straightforward if we knew that a version of $\nu(A \mid X^t)$ could be chosen for each ν so that $\nu(A \mid X^t = m)$ became a measurable function of ν for fixed m; unfortunately this does not seem to be an easy fact.

Fix A and let $A_0 := A \times H$. It is easily checked that

$$P(A \mid X^t) = \Pr(A_0 \mid X^t) \,. \tag{2.3}$$

Let $J_0 \subset \mathscr{M}(\mathbb{S}_d) \times H$ with $J_0 \in \sigma(X^t)$; J_0 is necessarily of the form $J \times H$ for some event $J \subset \mathscr{M}(\mathbb{S}_d)$. Let

$$W := \{v \in H: \int_J \Pr(A_0 \mid X^t)\, dv > v(A \cap J)\} \,,$$

$$W_0 := \mathscr{M}(\mathbb{S}_d) \times W \,.$$

Note that by (2.1) and (2.2), W_0 and $[v_{X^t} \in W]$ differ by a Pr-null set. Therefore if $\Pr(W_0) > 0$, then

$$\begin{aligned}
\Pr(W_0 \cap J_0 \cap A_0) &= \int_H \int_{J \cap [v_{X^t} \in W]} I_{A_0}\, dv\, d\mu(v) \\
&= \int_H \int_{J \cap [v_{X^t} \in W]} \Pr(A_0 \mid X^t)\, dv\, d\mu(v) \\
&= \int_W \int_J \Pr(A_0 \mid X^t)\, dv\, d\mu(v) \\
&> \int_W v(A \cap J)\, d\mu(v) \\
&= \Pr(W_0 \cap J_0 \cap A_0) \,,
\end{aligned}$$

a contradiction. Thus $\Pr(W_0) = 0$, and the same is similarly true if we reverse the inequality in the definition of W. Letting $\mathscr{A}_t := \{[X^t \in C]: C \in \mathscr{A}\}$, where $\mathscr{A}$ is a countable field which generates the σ-algebra in $\mathscr{M}(\mathbb{S}_d)$, we therefore have

$$\mu(\{v \in H: \int_J \Pr(A_0 \mid X^t)\, dv = v(A \cap J) \quad \text{for all } J \in \mathscr{A}_t\}) = 1 \,.$$

But equality for all $J \in \mathscr{A}_t$ implies equality for all $J \in \sigma(X^t)$, so the lemma follows from (2.3). □

Lemma 2.2. *Suppose X is a stationary random labeled graph, in $\mathbb{R}^d$ or in a lattice L, with finite site density. Suppose $A \subset \mathscr{M}(\mathbb{S}_d)$ is a set of labeled graphs in which the origin is a site of an infinite component, and $\{x \in C: \theta_x X \in A\}$ is finite for each infinite component C of X a.s. Then*

$$\mathrm{P}[\theta_x X \in A \text{ for some } x \in V] = 0 \,.$$

In Example 1.7, if we take A to be the set of all graphs in which every integer is a site, then $\theta_x X \in A$ if and only if x is an integer shifted by U, which is true for only finitely many x in each infinite cluster. Thus Lemma 2.2 is false without the assumption of finite site density.

Proof of Lemma 2.2. Let Z be the set of sites $x \in V$, necessarily in infinite components of X, for which $\theta_x X \in A$. Fix $k \geqq 1$. We wish to associate k sites of V to each site of Z, in a translation-invariant way. Given an infinite component C of X, we can order the sites of C in order of increasing distance from the finite set $C \cap Z$, breaking ties by an arbitrary method, say using lexicographic order, and thereby

obtain

$$C = \{v_1, v_2, \dots\}$$

with

$$C \cap Z = \{v_1, \dots, v_n\},$$

where $n = |C \cap Z|$. We can then associate to each site $v = v_i$ of $C \cap Z$ the set $W_v := \{v_{jn+i}, j = 0, \dots, k-1\}$. Doing this for all infinite C, we obtain a collection $\{W_v, v \in Z\}$ of disjoint k-site sets, each containing exactly one site of Z. Let $Z_{t,k} := \{v \in Z: W_v \subset v + \Lambda_t\}$. Then for $s > 0$,

$$|Z_{t,k} \cap \Lambda_s| \leqq k^{-1} |V \cap \Lambda_{s+t}| .$$

Dividing by $|\Lambda_s|$ and letting $s \to \infty$ shows $\tilde{P}_0(0 \in Z_{t,k}) \leqq k^{-1} \lambda_v(X)$ a.s. Since $\tilde{P}_0$ is a measure, letting $t \to \infty$ and then $k \to \infty$ shows $\tilde{P}_0(0 \in Z) = 0$, which by (1.3) proves the lemma. □

For each site x in a graph m, let $C(m, x)$ denote the connected component of x in m. If F is a finite set of sites in a single component, we write $C(m, F)$ for this component. Let V_m denote the vertex set of m. For F a finite set of sites of m, let $m \backslash F$ denote the subgraph of m obtained by deleting F and all bonds emanating from F. If F is contained in an infinite component of m, we define

$$h_m(F) := \{v \in V_m \backslash F: C(m, v) \text{ is infinite, } C(m \backslash F, v) \text{ is finite}\} .$$

We call F a *core* in m if $h_m(F)$ is infinite, i.e. removing F from m creates infinitely many new finite clusters. Note our previous definition of a core of a galaxy is a special case of this. Sites in $h_m(F)$ are called *satellites* of F. A given site may be a satellite of more than one core.

Clearly cores can exist only when there are sites of infinite degree. But finite degree of all sites is not necessary to prohibit cores; the next lemma shows stationarity does so as well.

Lemma 2.3. *Suppose X is a stationary random graph, in $\mathbb{R}^d$ or in a lattice L, with finite site density. Then with probability one, X contains no core, and every site which is in an infinite component of X is the starting point of an infinite self-avoiding path in X.*

In Example 1.7, with probability one there are at most finitely many occupied bonds between C_i and C_j for each i and j. It follows easily that each integer, shifted by U, is a core, and there are a.s. no infinite self-avoiding paths in X. Thus Lemma 2.3 is false without the assumption of finite site density.

Proof of Lemma 2.3. Let us call a core F *minimal* if no proper subset of F is a core. Let Z be the set of all sites which are in minimal cores, and W the set of all sites which are satellites of minimal cores. It is easily seen that every site in Z has infinite degree, and every site of W has finite degree. Let Y be the subgraph of X consisting of site set W and all bonds with both endpoints in W. If F is a minimal core and $q \in h_X(F)$, then since no bond emanating from F is in Y, $C(X \backslash F, w)$ contains $C(Y, w)$. It follows that Y has only finite components. Since each site of Y (i.e. of W)

has finite degree, if $w \in W$ then

$$g_X(w) := \{z \in Z\text{: there is a bond in } X \text{ from } z \text{ to } C(Y, w)\}$$

is finite and nonempty. Hence we can define a mapping φ from W to Z by letting $\varphi(w)$ be the closest site to w in $g_X(w)$, with ties broken as usual by lexicographic order.

Suppose w is a satellite of a minimal core F. We claim that $g_X(w) \subset F$. To see this, suppose $z \in Z \backslash F$ and there is a bond in X from z to $C(Y, w)$. Then z has infinite degree in X, hence also in $X \backslash F$, and $z \in C(X \backslash F, w)$, so $C(X \backslash F, w)$ is infinite, contradicting our assumption that w is a satellite of F and proving the claim. If follows that $\varphi^{-1}(F)$ is infinite for every minimal core F. Let

$$Z_0 := \{z \in Z\text{: } \varphi^{-1}(z) \text{ is infinite}\}\ ;$$

then Z_0 is nonempty if cores exist. For $k \geqq 1$ and $t > 0$ let

$$Z_{t,k} := \{z \in Z_0\text{: } |\varphi^{-1}(z) \cap (z + \Lambda_t)| \geqq k\}\ .$$

Then for $s > 0$,

$$|Z_{t,k} \cap \Lambda_s| \leqq k^{-1} |W \cap \Lambda_{s+t}|\ .$$

As in the proof of Lemma 2.2, it follows that Z_0 has density 0, so Z_0 is empty a.s., so no cores exist a.s.

Suppose C is an infinite component of X and X contains no cores. Then we can start at an arbitrary $v_0 \in C$ and, since $\{v_0\}$ is not a core in C, move to an adjacent v_1 which is in an infinite component C_1 of $C \backslash \{v_0\}$, then delete v_0. Since C contains no cores, neither does C_1, so we can iterate to find v_2, v_3, etc. which form an infinite self-avoiding path in C. □

Proof of Theorem 1.6. This is an adaptation of the proof of Theorem 1 of [8].

By Lemma 2.1, X also has positive finite energy under the law ν in place of P, for μ-almost every $\nu \in H$, so we may assume P is ergodic, and write r_c for $\rho_c(P)$.

Fix $r \geqq r_c$. There exists a constant k such that with probability one there are exactly k infinite r-clusters in X; as in [4, 8, and 15], positive finite energy implies that k must be 0, 1, or ∞, so we need to rule out $k = \infty$.

Thus suppose there are infinitely many infinite r-clusters, a.s. Then for s sufficiently large, there is a positive probability that at least three distinct infinite r-clusters meet Λ_s. Fix such an s. Then since by Lemma 2.3 there are a.s. no galaxies in $X_{\leqq r}$, the following also occurs with positive probability, if $t > s$ is chosen sufficiently large: in $X|_{\Lambda_s^c}$, at least three distinct infinite r-clusters meet $\Lambda_t \backslash \Lambda_s$; for each of these clusters there is a bond in X from one of its sites in $\Lambda_t \backslash \Lambda_s$ to a site in Λ_s; and there are no sites in $\partial \Lambda_s$. When all this occurs, and t is increased slightly if necessary, for sufficiently small $\delta > 0$ there exists a finite collection $\mathscr{C}$ of cubes, each of the form $\delta x + [0, \delta)^d$ for some $x \in \mathbb{Z}^d$ and each contained in $\Lambda_t \backslash \Lambda_s$, such that (i) each cube in $\mathscr{C}$ contains exactly one site; (ii) each of these sites is in a different infinite r-cluster in $X|_{\Lambda_s^c}$, and (iii) each of these sites has a connection at level r in $X|_{\Lambda_t}$ to a site in Λ_s. Because there are only finitely many possible collections $\mathscr{C}$ for fixed δ, there exist δ and $\mathscr{C}$ for which the probability that (i)–(iii) all occur is positive. Let A be the event that (i) and (iii) occur, F the event that (i) and (ii) occur, and G the event that $X|_{\Lambda_s}$ is connected at level r. Then $P(A \mid F) > 0$, and A is

a δ-approximate connection event from $\Lambda_t \backslash \Lambda_s$ to Λ_s, so positive finite energy yields that

$$P(A \cap G \cap F) > 0 \,. \tag{2.4}$$

Let us call Λ_s a *branch node* if (a) there exist three or more infinite r-clusters in $X|_{\Lambda_s^c}$ which are all part of a single larger infinite r-cluster in X, and (b) $X|_{\Lambda_s}$ is connected at level r. Then

$$P(\Lambda_s \text{ is a branch node}) \geqq P(A \cap G \cap F) \,. \tag{2.5}$$

For $z \in \mathbb{Z}^d$ we call $2sz + \Lambda_s$ a *branch node* if Λ_s is a branch node in $\theta_{2sz}X$. Let

$$Z := \{2sz\colon z \in \mathbb{Z}^d,\ 2sz + \Lambda_s \text{ is a branch node}\} \,.$$

For $v \in Z$ let $n_v \geqq 3$ denote the number of infinite r-clusters in the restriction of X to the complement of the corresponding branch node which are all part of a single infinite r-cluster in X, and let $D_v^{(i)}$, $i = 1, \ldots, n_v$, denote the sets of sites of these r-clusters from the complement of the branch node. Define for $q > 0$ and $k \geqq 1$,

$$Z_{q,k} := \{v \in Z\colon |D_v^{(i)} \cap (v + \Lambda_q)| \geqq k \text{ for all } i \leqq n_v|\} \,.$$

Then for $u > 0$, the hypotheses of Lemma 2 of [8] are satisfied for $V \cap \Lambda_{u+q}$ (in place of S), $Z_{q,k} \cap \Lambda_u$ (in place of R), and $D_v^{(i)} \cap \Lambda_{u+q}$ (in place of $C_v^{(i)}$), yielding

$$|Z_{q,k} \cap \Lambda_u| \leqq k^{-1} |V \cap \Lambda_{u+q}| \,.$$

As in the proof of Lemma 2.2, we obtain from this that Z is empty a.s., that is, there are no branch nodes; by (2.4) and (2.5), it cannot be that there are infinitely many infinite r-clusters a.s., and the theorem follows. □

For each site x in a labeled graph m, and $r \in \mathbb{R}$, let $C_r(x)$, or $C_r(m, x)$, denote the r-cluster of x in m, and let $C_\infty(r)$, or $C_\infty(m, r)$, denote the union of all infinite r-clusters in m. To set up the proof of Theorem 1.8 we begin with some observations about the structure X would have to have for uniqueness to hold for fixed r but not simultaneously for all r.

Lemma 2.4. *Suppose X is a random labeled graph, n $\mathbb{R}^d$ or in a lattice L. Suppose that*

$$\textit{for each } r \in \mathbb{R},\ P[X \text{ includes at most one infinite } r\text{-cluster}] = 1. \tag{2.6}$$

Then with probability one the following holds: for each $r > r_c(X)$ there is a unique main infinite r-cluster $C_\infty^m(r)$ which contains $\cup_{s<r} C_\infty(s)$. If T is an ephemeral, i.e. non-main, infinite r-cluster for some $r > r_c(X)$, then for all sites x in T, $C_s(x)$ is finite for all $s < r$, and $C_s(x) = C_\infty^m(s)$ for all $s > r$. Each site of X is in an ephemeral infinite r-cluster for at most one value of r.

An ephemeral infinite cluster, to motivate the name, can thus be thought of as follows: as r is increased, at some level some finite clusters coalesce into an infinite cluster outside the main one, but as r continues to increase, this cluster is immediately absorbed into the main one. Example 1.9 illustrates this; to better illustrate this heuristic, instead of giving every vertical bond in column m the label Y_m, one can let the labels on the vertical bonds in column m be (conditionally) iid uniform in $[0, Y_m)$.

If T is an ephemeral infinite cluster, we call the unique value of r such that T is an ephemeral infinite r-cluster the *index* of T. Lemma 2.4 ensures that, even if they have different indices, distinct ephemeral infinite clusters are disjoint.

Proof of Lemma 2.4. By (2.6), with probability one, there is a unique infinite q-cluster for each rational $q > r_c(X)$; we assume this event, which we denote U, occurs. Note that the graphs $C_\infty(s)$, $s \in \mathbb{R}$, form a nested collection, $C_\infty(q)$ is connected for rational q, and

$$\bigcup_{s<r} C_\infty(s) = \bigcup_{q<r, q\in\mathbb{Q}} C_\infty(q) .$$

It follows that $\bigcup_{s<r} C_\infty(s)$ is connected, and hence it is contained in a unique infinite r-cluster if $r > r_c(X)$.

Suppose U occurs and x is a site of an ephemeral infinite r-cluster T. If $s < r$, then T does not meet $C_\infty(s)$, so $C_s(x)$ is finite. If $s > r$ then there is an infinite s-cluster containing T; this s-cluster meets $C_\infty(r)$ so must be $C_\infty^m(s)$.

The last statement of the lemma follows from the fact that if x is a site in an ephemeral infinite r-cluster, then $r = \inf\{r' : C_{r'}(x) \text{ is infinite}\}$ so there is only one such r. □

In Example 1.1, it follows from Lemma 2.4 that invasion percolation started in an ephemeral infinite r-cluster for some $r > r_c(X)$ will never meet $C_\infty(q)$ for any $r_c(X) < q < r$, since the invasion process never leaves the infinite r-cluster it starts in. But results in [6] and [11] imply that invasion percolation does meet $C_\infty(q)$ for all $q > r_c(X)$ a.s., which shows that ephemeral infinite clusters do not exist, a.s., and thereby proves Theorem 1.8 for Example 1.1. Unfortunately this result on invasion percolation uses the rather deep fact from [11] that the half-space and full-space percolation critical points are the same, a fact which is presumably true but very technical and as yet unproved in models like Examples 1.2 and 1.4 where there is sufficient independence, and which could conceivably be false in some stationary models. Therefore it is desirable to have a proof of Theorem 1.8 based purely on stationarity and positive finite energy.

Proof of Theorem 1.8. Because of Lemma 2.4, the main task is to rule out the existence of ephemeral infinite clusters. The idea, roughly, is to attach an ephemeral r-cluster, if such exists, to the unique infinite q-cluster using positive finite energy, much as three or more infinite clusters were attached to each other in the proof of Theorem 1.6; here $q \in (r_c(X), r)$ is rational. The resulting "attached cluster" is the "arm" we define below. As only finitely many sites are involved in the attachment, one can associate arbitrarily many arm sites to each attachment site. An analog of Lemma 2.2 can then be used to show there are no ephemeral infinite clusters. The complications arise from the fact one needs to be able to identify the attachment sites by looking at the "attached" configuration only. This is problematic, for example, if the ephemeral infinite cluster has an "infinite string of beads" structure, and finite bead strings, attached at one end, form part of the infinite q-cluster; one cannot then tell between which two beads the attachment was made. To solve this problem we will use the notion of "pivot-free" arms.

STEP 1. As in the proof of Theorem 1.6, we may assume P is ergodic, and write r_c for $\rho_c(P)$. With probability one, there is a unique infinite q-cluster in X for each rational $q > r_c$ and at most one for $q = r_c$. Fix a rational $q > r_c$ and suppose that

with positive probability, an ephemeral infinite r-cluster exists for some $r > q$. Then for s sufficiently large, there is a positive probability for the event, denoted A_1, that in X both $C_\infty(X, q)$ and (for some $r > q$) an ephemeral r-cluster meet Λ_s. Fix such an s. Let us call an infinite r-cluster R in $X|_{\Lambda_3^c}$ *pre-ephemeral* (with respect to Λ_s) if for every $r' < r$, R does not contain an infinite r'-cluster in $X|_{\Lambda_3^c}$. Then

$$\text{for every } x \in R,\ r = \inf\{r' \colon x \in C_\infty(X|_{\Lambda_3^c}, r')\}\ . \tag{2.7}$$

Let A_2 be the event that, for some $r > q$,

(a) there is a pre-ephemeral infinite r-cluster R in $X|_{\Lambda_3^c}$;
(b) there is a bond at level r in X from a site of R to a site in Λ_s;
(c) $C_\infty(X, q)$ meets Λ_s.

Note that r and R are not necessarily unique. Since $C_\infty(X, q)$ is a.s. not a galaxy, (c) is a.s. equivalent to

(c′) there is a bond at level q in X from a site of $C_\infty(X|_{\Lambda_3^c}, q)$ to a site in Λ_s.

Let A_3 be the event that (a)–(c) occur for some $r > q$ and

(d) $X|_{\Lambda_s}$ is connected at level q;

when A_3 occurs we say that Λ_s is a *pivot node* or *r-pivot node* in X. We then call the R of (a) and (b) an *arm* or *r-arm* corresponding to Λ_s. By translation we can apply similar terminology to $2sz + \Lambda_s$ for arbitrary $z \in \mathbb{Z}^d$. Roughly, an arm is like an ephemeral infinite cluster, except that it is attached to the infinite q-cluster, with the only such attachment occurring inside the corresponding pivot node. As a shorthand we say an arm *contains* a pivot node if it contains some, and hence every, site in that node. The value r is called the *index* of the arm or pivot node; note this index need not be unique for pivot nodes. We call Λ_s a *shoulder node* or *r-shoulder node* if Λ_s is a pivot node and the arm R as in (a) and (b) can be chosen to contain no pivot node; we call such an arm *pivot-free*. Again similar terminology applies to $2sz + \Lambda_s$.

Recall that, from Lemma 2.4, distinct ephemeral infinite clusters of all indices are disjoint. It follows from Lemma 2.4 applied to X, and from Lemma 2.3 applied to the stationary random graph Y which is the union of all ephemeral infinite clusters of all indices (without labels), that with probability one, no ephemeral infinite cluster in X is a galaxy. Therefore if T is an ephemeral infinite r-cluster in X for some $r > q$, then $T|_{\Lambda_3^c}$ contains an infinite r-cluster R, which is necessarily pre-ephemeral. It follows that if A_1 occurs, then with probability one so does A_2, so $P(A_2) \geqq P(A_1) > 0$. As in the proof of Theorem 1.6, it then follows from positive finite energy that $P(A_3) > 0$. Here we use the fact that, if one specifies the site locations in (b) and (c′) to within a δ-cube, the result is a δ-approximate connection event, because by (2.7) the connection level r of (b) is a function of $X|_{\Lambda_3^c}$.

STEP 2. We claim that for each $r > q$, to each pivot-free r-arm there corresponds a unique shoulder node; thus pivot-free arms avoid the identifiabilty problem discussed preceding Step 1. By definition there is at least one such r-shoulder node, so suppose $2sy + \Lambda_s$ and $2sz + \Lambda_s$ are both shoulder nodes corresponding to the same pivot-free r-arm R, with $y \neq z$. Then R is contained in $((2sy + \Lambda_s) \cup (2sz + \Lambda_s))^c$, but by (b) there is a bond in X at level r from R to $2sz + \Lambda_s$. This bond is part of $X|_{(2sy+\Lambda_s)^c}$, so R is not an r-cluster in $X|_{(2sy+\Lambda_s)^c}$, which is a contradiction.

STEP 3. We have seen that if ephemeral infinite clusters exist then so do arms and pivot nodes, but we need to show that pivot-free arms exist. Let us write $2sy + {}_s \gg_{(R)} 2sz + \Lambda_s$ if $2sy + \Lambda_s$ and $2sz + \Lambda_s$ are pivot nodes, with $2sz + \Lambda_s$ contained in an arm R corresponding to $2sy + \Lambda_s$. Note we do not claim this relation is transitive. We claim (1) there then exists an infinite self-avoiding path at level q in X which starts in $2sz + \Lambda_s$; (2) every infinite self-avoiding path at level q in X which enters R must later enter $2sy + \Lambda_s$; and (3) there is a path from $2sz + \Lambda_s$ to $2sy + \Lambda_s$ at level q, with all sites except its last one in R. Claim (1) is immediate from (c) and Lemma 2.3. For (2), suppose γ is an infinite self-avoiding path at level q in X which enters R but does not later enter $2sy + \Lambda_s$. Removing an initial segment if necessary, we may assume γ starts in R and never enters $2sy + \Lambda_s$. Then since $r > q$ and R is an r-cluster in $X|_{(2sy+\Lambda_s)^c}$, γ is contained in R. But then R is not pre-ephemeral, a contradiction. For (3), the existence of a path from $2sz + \Lambda_s$ to $2sy + \Lambda_s$ at level q is immediate from (1) and (2). Since $r > q$ and R is an r-cluster in $X|_{(2sy+\Lambda_s)^c}$ the portion of this path before its first visit to $2sy + \Lambda_s$ is contained in R. But we can end the path at this first visit, and (3) follows.

Next suppose there is an infinite sequence $2sy_1 + \Lambda_s \gg_{(R_1)} 2sy_2 + \Lambda_s \gg_{(R_2)} \ldots$ of pivot nodes. By Claim (1) there exists an infinite path α at level q in X which starts in $2sy_1 + \Lambda_s$. Let R_0 denote $2sy_1 + \Lambda_s$. If α enters R_i for some $i \geqq 1$ then by Claim (2) α later enters $2sy_i + \Lambda_s \subset R_{i-1}$. Iterating we see that every visit by α to some R_i is followed by a return to $R_0 = 2sy_1 + \Lambda_s$. Since α is self-avoiding and $2sy_1 + \Lambda_s$ contains only finitely many sites, there can be only finitely many such visits and returns. Hence by removing an initial segment from α if necessary, we obtain an infinite self-avoiding path at level q in X which starts in $2sy_1 + \Lambda_s$ and never visits $\bigcup_{i \geqq 1} R_i$. It follows that $2sy_1 + \Lambda_s$ is not contained in, and therefore by (d) does not meet, any R_i. Thus $y_i \neq y_1$ for all $i \neq 1$. Applying this fact to the sequence of pivot nodes starting from $2sy_j + \Lambda_s$ for arbitrary j, we see that all y_i are distinct.

By Claim (3) for each $i \geqq 1$ there is a path γ_i from $2sy_{i-1} + \Lambda_s$ to $2sy_i + \Lambda_s$ in $R_{i-1} \cup R_i$, with only the initial site in R_{i-1}. Let $\tilde{\gamma}_1$ denote γ_1 with the first site and bond removed, so $\tilde{\gamma}_1$ is contained in R_1. By (d), $Q := \tilde{\gamma}_1 \cup (\bigcup_{i \geqq 2} \gamma_i) \cup (\bigcup_{i \geqq 2} X|_{(2sy_i + \Lambda_s)})$ is a graph in $\bigcup_{i \geqq 1} R_i \subset (2sy_1 + \Lambda_s)^c$, connected at level q, which meets R_1. Since the nodes $2sy_i + \Lambda_s$ are distinct and each contains at least one site, Q is infinite. Since R_1 is an r-cluster in $X|_{(2sy_1 + \Lambda_s)^c}$, Q is part of R_1. But this means R_1 is not pre-ephemeral, a contradiction. Thus there can be no infinite sequence $2sy_1 + \Lambda_s \gg_{(R_1)} 2sy_2 + \Lambda_s \gg_{(R_2)} \ldots$ of pivot nodes.

But if $2sy + \Lambda_s$ is a pivot node and there is no arm R and pivot node $2sz + \Lambda_s$ with $2sy + \Lambda_s \gg_{(R)} 2sz + \Lambda_s$, then there is a pivot-free arm corresponding to $2sy + \Lambda_s$. It follows that if pivot nodes exist, then so do pivot-free arms. Thus with probability one, the existence of ephemeral infinite clusters implies the existence of shoulder nodes and pivot-free arms.

STEP 4. Next let us verify that each site of X is in at most one pivot-free arm, and that this arm has a unique index. Suppose w is a site of $R \cap R'$, where $r \leqq r'$, R is a pivot-free r-arm corresponding to some shoulder node $2su + \Lambda_s$, and R' is a pivot-free r'-arm corresponding to a shoulder node $2su' + \Lambda_s$. Since R is pivot-free, R does not meet $2su' + \Lambda_s$. Hence in $X|_{(2su' + \Lambda_s)^c}$, R' is an r'-cluster containing w, and R is connected at level r', so R' contains R. But this means $C_\infty(X, r) \cap R'$ contains R so is not finite. Since R' is pre-ephemeral, we must then have $r = r'$. But then R and R' are non-disjoint connected components of $X_{\leqq r}$, so $R = R'$.

STEP 5. Let W be the set of all sites of all pivot-free arms in X with index greater than q, and

$$Z := \{z \in \mathbb{Z}^d\colon 2sz + \Lambda_s \text{ is a shoulder node}\} .$$

By Steps 2 and 4, there is a mapping φ of W onto Z which takes each pivot-free arm site w to the unique z such that $2sz + \Lambda_s$ is the shoulder node corresponding to the unique pivot-free arm containing w. Since $\varphi^{-1}(z)$ is infinite for all $z \in Z$, we can fix $k \geqq 1$ and define W_z to be the set consisting of the k closest sites to $2sz$ in $\varphi^{-1}(z)$, as usual breaking ties using lexicographic order. Then as in the proof of Lemma 2.2, we obtain that Z has density 0 and is therefore empty a.s. Therefore by Step 3 there are no ephemeral infinite clusters with index greater than q, a.s. Since $q > r_c$ is an arbitrary rational, and by Theorem 1.6 there is a.s. at most one infinite r_c-cluster, the theorem follows. □

The author would like to thank M. Keane for helpful discussions regarding Lemma 2.1 and the "random counting measure" formulation of random graphs.

References

1. Aizenman, M., Newman, C.M.: Discontinuity of the percolation density in one dimensional $1/|x-y|^2$ percolation models. Commun. Math. Phys. **107**, 611–648 (1986)
2. Aldous, D., Steele, J.M.: Asymptotics for euclidean minimal spanning trees on random points. Probab. Theory Rel. Fields **92**, 247–258 (1992)
3. Alexander, K.S.: Percolation and minimal spanning forests in infinite graphs. Ann. Probability, to appear (1994)
4. Burton, R., Keane, M.: Density and uniqueness in percolation. Commun. Math. Phys. **121**, 501–505 (1989)
5. Burton, R., Meester, R.W.I.: Long range percolation in stationary point processes. Rand. Struc. Alg. **4**, 177–190 (1993)
6. Chayes, J.T., Chayes, L., Newman, C.M.: The stochastic geometry of invasion percolation. Commun. Math. Phys. **101**, 383–407 (1985)
7. Fortuin, C., Kastelyn, P.: On the random cluster model. I. Physica **57**, 536–564 (1972)
8. Gandolfi, A., Keane, M., Newman, C.M.: Uniqueness of the infinite component in a random graph with applications to percolation and spin glasses. Probab. Theory Rel. Fields **92**, 511–527 (1992)
9. Grimmett, G.R.: Percolation. Berlin, Heidelberg, New York: Springer, 1989
10. Grimmett, G.R.: The stochastic random-cluster process, and the uniqueness of random-cluster measures. Preprint, 1994
11. Grimmett, G.R., Marstrand, J.M.: The supercritical phase of percolation is well behaved. Proc. R. Soc. Lond. A **430**, 439–457 (1990)
12. Kesten, H.: Percolation theory for mathematicians. Boston: Birkhäuser, 1982
13. Meester, R.W.J., Roy, R.: Uniqueness of unbounded occupied and vacant components in Boolean models. Preprint, 1993
14. Neveu, J.: Processus ponctuels. In: Ecole d'Ete de Probabilities de Saint-Flour VI-1976. Lecture Notes in Mathematics **598**, Berlin, Heidelberg, New York: Springer 1977, pp. 249–447
15. Newman, C.M., Schulman, L.S.: Infinite clusters in percolation models. J. Stat. Phys. **26**, 613–628 (1981)
16. Penrose, M.: On a continuum percolation model. Adv. Appl. Prob. **23**, 536–556 (1991)
17. Preparata, F.P., Shamos, M.I.: Computational Geometry: An Introduction. Berlin, Heidelberg, New York: Springer 1985

Communicated by D. Brydges