

Percolation on transitive graphs

Thesis by
Philip Easo

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

The Caltech logo, consisting of the word "Caltech" in a bold, orange, sans-serif font.

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2025
Defended May 28, 2025

© 2025

Philip Easo
ORCID: 0000-0002-5606-3727

All rights reserved

I am deeply grateful to my advisor Tom Hutchcroft for introducing me to this beautiful area of research, for sharing his ideas and encouraging mine, for bringing me to Caltech, and for much more — beyond mathematics — which he has taught me by way of example.

I am also deeply grateful to my partner Henrietta Coales for her constant support, for keeping me grounded, and for joining me on this adventure in California.¹

¹I also acknowledge that the name “sandcastle” of the object introduced in Chapter 5 was her suggestion.

Percolation on a transitive graph is an idealized mathematical model for a homogeneous system undergoing a phase transition. We will investigate how the geometry of an infinite transitive graph determines whether percolation undergoes a phase transition, and if so, at what critical point. Building on these ideas, we will develop a new theory of percolation on *finite* transitive graphs. This theory unifies the percolation phase transition on infinite transitive graphs with the giant-cluster phase transition in the celebrated Erdős-Rényi model from combinatorics.

INTRODUCTION

Phase transitions appear in our lives in both obvious and subtle ways. The most vivid examples come from physics: as temperature slowly increases past certain critical points, ice suddenly melts and magnets lose their magnetism. The example that received the most news coverage during the recent Covid pandemic concerned the so-called *R*₀-value: whether *R*₀ < 1 or *R*₀ > 1 would predict whether the number of infected individuals would decay or grow exponentially. More hidden phase transitions are present in the formation of traffic jams and the average-case computing time required by an algorithm.

Phase transitions are in fact a very typical occurrence anytime one assembles many tiny identical components (picture: atoms arranged in a lattice) in which each component is only allowed to interact with its neighbouring components and one varies this “local” interaction strength. *Percolation* is the mathematician’s caricature of this setup.

1.1 Percolation

We will model an assembly of tiny components by a *graph* $G = (V, E)$. We will always assume that graphs are connected, undirected, simple, have countably many vertices, and that each vertex has at most finitely many neighbours. To model that all of the tiny components are identical, we will require that G is (*vertex*-)transitive, meaning that for all vertices u and v , there exists an automorphism ϕ of G such that $\phi(u) = v$. This includes, for example, the usual Euclidean lattices \mathbb{Z}^d , regular trees, and more generally, any Cayley graph of a finitely generated group.

Fix a parameter $p \in [0, 1]$, which will be our “local” interaction strength. Build a random spanning subgraph $\omega \subseteq G$ by independently choosing, for each edge $e \in E$, whether to include e (with probability p) or delete e (with probability $1 - p$). (See figure 1.1.) The law of ω is called (*Bernoulli bond*) *percolation*. We denote this law by \mathbb{P}_p . To find the phase transition in this model, track the *clusters* (i.e. connected components) of ω while varying p . For simplicity, let us start by assuming that G is infinite. By Kolmogorov’s 0-1 law and the monotonicity of this model with respect to p , there is always some critical parameter $p_c(G) \in [0, 1]$ such that

$$\mathbb{P}_p(\omega \text{ contains an infinite cluster}) = \begin{cases} 0, & \text{if } p < p_c \\ 1, & \text{if } p > p_c. \end{cases}$$

From a utilitarian perspective, a good reason to study percolation is that this toy model provides a testing ground for new techniques, which often trickle down to more intricate, physically-relevant models some years later. From an aesthetic perspective, percolation provides many simple-to-state yet challenging-to-solve problems that require mathematicians to exercise their creativity. This is best-illustrated by the most notorious open problem in the area: prove that when G is the Euclidean lattice \mathbb{Z}^3 ,

$$\mathbb{P}_{p_c}(\omega \text{ contains an infinite cluster}) = 0.$$

This model was first introduced by Broadbent and Hammersley in 1957 to model the flow of a fluid through a porous medium. In their setup, the graph G is a Euclidean lattice like \mathbb{Z}^d , and indeed, most work on percolation has traditionally taken G to be such a lattice or to be a tree. The scope widened in 1996, when Benjamini and Schramm launched an influential research programme to systematically study percolation on general infinite transitive graphs. This programme is the context in which the present thesis should be understood. None of the results in this thesis are new for Euclidean lattices or trees, nor are we concerned with any other particular transitive graph; our goal has been to understand basic features of the percolation phase transition that hold for all transitive graphs.

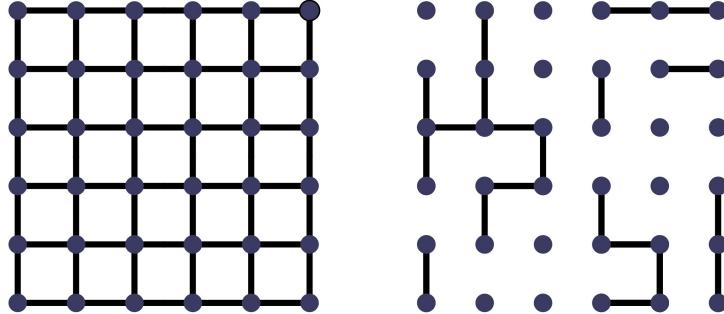


Figure 1.1: Example of a graph G (left) and a spanning subgraph ω (right).

1.2 Transitive graphs

In this section, we will take a tour of the zoo that is the space of transitive graphs. Given a transitive graph $G = (V, E)$, we will always write o to denote an arbitrary vertex. (Statements will always hold independently of the choice of o because G is transitive.) We write $\text{dist} = \text{dist}_G$ to denote the graph metric on G . Given $n \geq 1$ and $u \in V$, we write $B_n(u) = B_n^G(u)$ to denote both the set $\{v \in V : \text{dist}(u, v) \leq n\}$ and the subgraph of G that this set of vertices induces. Hopefully it will

always be clear from context which of these definitions is intended. We also adopt the convention that $B_n := B_n(o)$.

Local and global similarity

It helps to think of the geometry of a given infinite transitive graph G in two parts: the small-scale, *local* structure and the large-scale, *global* structure. Indeed, an important principle in the study of percolation on transitive graphs is that certain key features such as the behaviour of the model around the critical point, or the value of the critical point itself, should be entirely determined by one part or the other. Rather than attempt to intrinsically define what the “local” and “global” structure are of a given infinite transitive graph G , let us instead define what it means for a given pair of graphs to have similar global or local structures.

The *local* metric on the space of transitive graphs is given by, for each pair G, H of transitive graphs,

$$\text{dist}_{\text{loc}}(G, H) := 2^{-\sup\{n \geq 1 : B_n^G \cong B_n^H\}},$$

where $B_n^G \cong B_n^H$ means that the corresponding rooted *subgraphs* (rooted at o) are isomorphic. The topology this induces on the space of all transitive graphs is called the *local* (aka *Benjamini-Schramm*) topology. For example, the sequence of tori $((\mathbb{Z}/n\mathbb{Z})^2 : n \geq 1)$ and the sequence of cylinders $(\mathbb{Z} \times (\mathbb{Z}/n\mathbb{Z}) : n \geq 1)$ both converge *locally* to the square lattice \mathbb{Z}^2 .

On the other hand, *quasi-isometry* provides a way to measure the *global* similarity of graphs. Given metric spaces (V_1, d_1) and (V_2, d_2) , and given constants $C, D > 0$, we say that d_1 is (C, D) -*quasi-isometric* to d_2 if there exists a function $\phi : V_1 \rightarrow V_2$ such that for all $u, v \in V_1$,

$$\frac{1}{C}d_1(u, v) - D \leq d_2(\phi(u), \phi(v)) \leq Cd_1(u, v) + D,$$

and for every vertex $x \in V_2$, there exists some $u \in V_1$ with $d_2(x, \phi(u)) \leq D$. In other words, ϕ roughly preserves distances and is roughly surjective. We simply say that that d_1 and d_2 are quasi-isometric if there exists $C, D > 0$ such that d_1 is (C, D) -quasi-isometric to d_2 , and it can be easily verified that this defines an equivalence relation. We naturally extend all of these definitions to graphs by identifying each graph with its graph metric. For example, the square lattice (graph) \mathbb{Z}_2 is quasi-isometric to the plane \mathbb{R}^2 (with its usual metric), and the cylinder $(\mathbb{Z} \times (\mathbb{Z}/100\mathbb{Z}))$ is quasi-isometric to the line \mathbb{Z} .

When working with infinite transitive graphs, we will not be so interested in the constants of quasi-isometries, and we simply regard two quasi-isometric infinite transitive graphs as having “the same” global structure. However, this notion is not useful when working with finite transitive graphs, since

all finite graphs are quasi-isometric to each other. Instead, the more relevant way to measure the global similarity of two non-trivial compact metric spaces (V_1, d_1) and (V_2, d_2) (such as the graph metrics of two finite graphs) is to first normalise, i.e. consider the metrics $\tilde{d}_1 := \frac{1}{\text{diam } d_1} d_1$ on V_1 and $\tilde{d}_2 := \frac{1}{\text{diam } d_2} d_2$ on d_2 (where diam denotes the diameter of a metric space) then take the Gromov-Hausdorff distance $\text{dist}_{\text{GH}}(\tilde{d}_1, \tilde{d}_2)$. Recall that this Gromov-Hausdorff distance¹ is defined to be the infimum $\varepsilon > 0$ such that there exists a third metric space (V_3, d_3) and isometric embeddings $\phi_1 : (V_1, d_1) \rightarrow (V_3, d_3)$ and $\phi_2 : (V_2, d_2) \rightarrow (V_3, d_3)$ such that the Hausdorff distance with respect to d_3 between the images $\phi_1(V_1)$ and $\phi_2(V_2)$ in V_3 is at most ε . For example, the diameter-rescaled sequence of tori $\left(\left(\frac{2\pi}{n} \mathbb{Z}/n\mathbb{Z} \right)^2 : n \geq 1 \right)$ (where, given a graph G and a constant $c > 0$, we write cG for the graph metric on G where all distances are scaled by c) Gromov-Hausdorff converges to the torus $S^1 \times S^1$ with the L^1 metric.

Growth rates

Let \mathcal{G} be the set of all infinite transitive graphs. An important way to begin classifying the elements of \mathcal{G} is according to their *growth rate*, i.e. the asymptotic behaviour of the function $\text{Gr} : \mathbb{N} \rightarrow \mathbb{N}$ sending $\text{Gr} : n \mapsto |B_n|$, where $|B_n|$ denotes the number of *vertices* in B_n .

In an ideal world, we would have a classification of all possible global geometries of infinite transitive graphs, similar in spirit to the classification of all finite simple groups. For example, we might wish for a family of easy to describe infinite transitive graphs such that every infinite transitive graph is quasi-isometric to one in this list. Unfortunately, this is not the case, and no one expects that such a classification would ever be possible.

However, spectacularly, there is a kind of classification if we restrict to those $G \in \mathcal{G}$ with *polynomial growth*, i.e. those for which there exist $C, D < \infty$ such that $\text{Gr}(n) \leq Cn^D$ for all n . Indeed, thanks to deep, classical work of Gromov and Trofimov, every infinite transitive graph of polynomial growth is quasi-isometric to the Cayley graph of a special kind of group called a *nilpotent* group. While we will not define what it means to be a nilpotent group here, let us simply say that it is a generalisation of being abelian, so for example, the group \mathbb{Z}^d is always nilpotent.

Thanks to this so-called *structure theory* for infinite transitive graphs of polynomial growth, certain important results about probability on general infinite transitive graphs proceed by a *structure vs expansion* dichotomy: either the graph has polynomial growth, in which case we can apply this detailed structure theory, or the graph has *super-polynomial* growth, in which case this fast growth

¹Up to doubling ε , the same notion can be expressed in terms of quasi-isometries with constants, although this latter definition is less commonly used.

can itself be exploited directly.

At the other extreme from polynomial growth, we say that an infinite transitive graph G has *exponential growth* if there exist $c, \delta > 0$ such that $\text{Gr}(n) \geq ce^{\delta n}$ for all n . For example, this includes the d -regular tree for every $d \geq 3$. The hypothesis of exponential growth can itself be helpful when studying percolation. For example, the analogue of the “notorious open problem” about critical percolation on \mathbb{Z}^3 was proved for all graphs of exponential growth. There also exist mysterious transitive graphs that have neither polynomial growth nor exponential growth, and are therefore said to have *intermediate growth* (e.g. the “Grigorchuk group”).

Expansion

One way that having a fast growth rate (“expansion” in the “structure vs expansion” dichotomy) can be exploited is as follows: for transitive graphs, a fast growth rate implies good isoperimetric properties; good isoperimetric properties imply that a random walk on the graph has good escape properties; and these escape properties can be used to prove geometric facts about a graph that serve percolation arguments. This is a common theme running through the heart of both our work on non-triviality and locality (discussed in later sections), for example.

Let us be more precise about what we mean by “good isoperimetric and escape properties”. Let $G = (V, E)$ be an infinite graph with bounded vertex degrees. Given a set of vertices $S \subseteq V$, let ∂S be the edge boundary of S , i.e. the set of all edges having one endpoint in S and the other in $V \setminus S$. Given $d \geq 1$, we say that G *satisfies a d -dimensional isoperimetric inequality* if there exists $c > 0$ such that for every finite set of vertices S ,

$$|\partial S| \geq c |S|^{\frac{d-1}{d}}.$$

Correspondingly, we define the *isoperimetric dimension* of G to be

$$\text{Dim } G := \sup \{d \geq 1 : G \text{ satisfies a } d\text{-dimensional isoperimetric inequality}\}.$$

For example, $\text{Dim } \mathbb{Z}^d = d$ for all $d \geq 1$. Some classical results about random walks on general graphs are that if $\text{Dim } G > 2$, then simple random walk is transient (i.e. has a positive probability to never return to its starting point.), and if $\text{Dim } G > 4$, then the paths of two independent simple random walks have a positive probability to never intersect.

An important consequence of the aforementioned structure theory for polynomial growth is that an infinite transitive graph G has polynomial growth if and only if $\text{Dim } G < \infty$. In particular, if a graph does not have polynomial growth, then $\text{Dim } G = \infty$ and hence, for example, simple random walk on G is transient.

Another notion of good isoperimetry, which is stronger than having $\text{Dim } G = \infty$ and which also plays an important role in percolation, is *nonamenability*. We say that G is nonamenable if there exists $c > 0$ such that for every finite set of vertices S ,

$$|\partial S| \geq c |S|.$$

This is the infinite-graph analogue of being an *expander* graph. A graph that is not nonamenable is said to be amenable. A fundamental result in the study of percolation on general graphs is that if G is amenable, then

$$\mathbb{P}_p (\omega \text{ contains at least two infinite clusters}) = 0 \quad \text{for all } p \in [0, 1],$$

and a major open conjecture is that the converse is true too.

1.3 Non-triviality

The first basic question when embarking on a general study of percolation on arbitrary infinite transitive graphs is to understand for which graphs percolation undergoes a *non-trivial* phase transition, meaning that the critical point p_c is *strictly* between 0 and 1. It is easy to show (exercise!) that $p_c > 0$, and indeed this holds for every infinite (not necessarily transitive) graph whose vertex degrees are bounded above uniformly. So the real question is to understand when $p_c < 1$. This question about whether percolation on a given graph undergoes a non-trivial phase transition is actually equivalent to the analogous questions for many other statistical mechanics models on the same graph, most notably for the Ising model of magnetism.

General graphs

Before restricting to transitive graphs, let us start by considering a more general setup. Let G be *any* infinite (not necessarily transitive) graph².

Does “ $p_c(G) < 1$ ” have a geometric counterpart?

A celebrated argument from 1930s physics gives a geometric condition \mathcal{P} that is *sufficient* (but *a priori* not necessary) for $p_c < 1$, defined in terms of *cutsets*. Given a vertex v , we say that a set of edges Π is a *cutset* from v to ∞ if v belongs to a finite component of the graph $(V, E \setminus \Pi)$. We say that a cutset Π from v to ∞ is *minimal* if Π does not contain a proper subset Π' that is also a minimal cutset from v to ∞ . Consider the exponential growth rate for the number of minimal cutsets of size n :

$$\kappa(G) := \sup_{n \geq 1} \sup_{v \in V} |\{\Pi : \Pi \text{ is a minimal cutset from } v \text{ to } \infty \text{ of size } n\}|^{1/n},$$

²As always, we assume that G is connected, undirected, etc.

and consider the following *uniform* analogue of $p_c(G)$:

$$p_c^*(G) := \inf\{p \in [0, 1] : \inf_{v \in V} \mathbb{P}_p(v \text{ is in an infinite component}) > 0\}.$$

Note that in general, $p_c^*(G) \geq p_c(G)$, and if G happens to be transitive, then $p_c^*(G) = p_c(G)$. Now the so-called *Peierls argument* establishes the following:

Theorem 1.3.1 (Peierls, 1936). *For every infinite graph, if $\kappa < \infty$ then $p_c^* < 1$.*

With Severo and Tassion, we recently established that the converse holds too:

Theorem 1.3.2 (Easo, Severo, Tassion). *For every infinite graph, $\kappa < \infty$ if (and only if) $p_c^* < 1$.*

Transitive graphs

It is easy to see that the line \mathbb{Z} satisfies $p_c = 1$, and indeed, so does every infinite transitive graph that is quasi-isometric to \mathbb{Z} . In their original work introducing percolation on general transitive graphs, Benjamini and Schramm conjectured that conversely, these are in fact the *only* infinite transitive graphs with $p_c = 1$.

Conjecture 1.3.3 (Benjamini and Schramm 1996). *Every infinite transitive graph that is not quasi-isometric to \mathbb{Z} satisfies $p_c < 1$.*

Babson and Benjamini then made the following (a priori) stronger conjecture, which is actually completely deterministic.

Conjecture 1.3.4 (Babson and Benjamini³ 1999). *Every infinite transitive graph that is not quasi-isometric to \mathbb{Z} satisfies $\kappa < \infty$.*

Babson and Benjamini proved their own conjecture in the special case of Cayley graphs of finitely presented groups, and Timar later showed that the property “ $\kappa < \infty$ ” is invariant under quasi-isometries. By the structure theory of polynomial growth, these results imply that for every infinite transitive graph G satisfying $\text{Dim } G < \infty$, if G is not quasi-isometric to \mathbb{Z} , then $\kappa < \infty$ and (hence) $p_c < 1$. In particular, to prove either of the above conjectures, it suffices to work with graphs satisfying $\text{Dim } G = \infty$.

In a breakthrough work, Duminil-Copin, Goswami, Raoufi, Severo, and Yadin resolved the $p_c < 1$ conjecture of Benjamini and Schramm. More precisely, they proved the following theorem, and

³To be historically accurate, they made this conjecture in the case of Cayley graphs, and Benjamini later asked about arbitrary transitive graphs more generally.

by the previous paragraph, this was sufficient to resolve the conjecture. Their proof involved an intricate multi-scale interpolation scheme comparing probabilities of certain events for percolation to those of a percolation-like model arising from the *Gaussian free field*.

Theorem 1.3.5 (Duminil-Copin, Goswami, Raoufi, Severo, and Yadin). *Every infinite graph with $\text{Dim } G > 4$ satisfies $p_c^* < 1$.*

Our proof of Theorem 1.3.2 can be tweaked to also prove the following theorem. This yields a much simpler and shorter proof of Theorem 1.3.5 and resolves the $\kappa < \infty$ conjecture of Babson and Benjamini. It remains open to extend this result to $\text{Dim } G > 1$.

Theorem 1.3.6 (Easo, Severo, Tassion). *Every infinite graph with $\text{Dim } G > 2$ satisfies $\kappa < \infty$.*

1.4 Locality

As explained in the previous section, p_c is strictly between 0 and 1 for every infinite transitive graph that is not quasi-isometric to \mathbb{Z} . A natural follow-up question is: what is the value of p_c *exactly*? Unfortunately, there is no reason in general to expect a simple formula for p_c , even in natural examples like the three-dimensional lattice \mathbb{Z}^3 . (There are a few spectacular exceptions to this rule.)

A weaker question is: which aspect of the geometry of G is responsible for determining the value of p_c ? *Schramm's locality conjecture* asserted that when p_c is strictly between 0 and 1, the exact value of p_c should be entirely determined by the local geometry of G . More precisely, for any $\varepsilon > 0$, there should exist $n \geq 1$ such that for every pair of infinite transitive graphs G and H , neither of which is quasi-isometric to \mathbb{Z} , if B_n^G is isomorphic to B_n^H , then $|p_c(G) - p_c(H)| \leq \varepsilon$. Equivalently, letting \mathcal{H} denote the set of all infinite transitive graphs that are not quasi-isometric to \mathbb{Z} , endowed with the local topology, the function $p_c : \mathcal{H} \rightarrow (0, 1)$ should be continuous. Thanks to fundamental work of Grimmett and Marstrand (which preceded Schramm's conjecture about general transitive graphs), this was known in the special case of Euclidean lattices. In particular, it follows from their work that for every $d \geq 2$, the sequence of graphs $(\mathbb{Z}^d \times (\mathbb{Z}/n\mathbb{Z}) : n \geq 1)$, which converges locally to \mathbb{Z}^{d+1} , satisfies

$$p_c \left(\mathbb{Z}^d \times (\mathbb{Z}/n\mathbb{Z}) \right) \rightarrow p_c \left(\mathbb{Z}^{d+1} \right) \quad \text{as } n \rightarrow \infty.$$

Various authors verified special cases of Schramm's locality conjecture (and some variants to other models). Most notably, by exploiting the structure theory discussed earlier, Contreras, Martineau, and Tassion proved Schramm's locality conjecture for all infinite transitive graphs of polynomial

growth. Building on this work and new ideas, Hutchcroft and I were able to resolve Schramm's locality conjecture in its entirety.

Theorem 1.4.1 (Easo and Hutchcroft). *The function $p_c : \mathcal{H} \rightarrow (0, 1)$ is continuous.*

This theorem, which in this formulation appears quite abstract, can in fact be reduced to a more concrete statement about propagating connection bounds. Very roughly, it suffices to prove that the statement $\mathcal{P}(\varepsilon, n, p)$ that “every pair of vertices at distance $\leq n$ apart are connected under \mathbb{P}_p with probability at least ε ” satisfies a general implication of the form

$$\mathcal{P}(\varepsilon, n, p) \implies \mathcal{P}(\varepsilon', n', p + \delta),$$

where $\delta > 0$, $\varepsilon' < \varepsilon$, $n' > n$, and which is quantitatively strong, i.e. δ is small, n' is large, and ε' is not *too* small. (The actual implication we end up proving is quite a bit more intricate.)

To prove such an implication, we split into various cases depending on the way that the graph *looks* around the current scale n , exploiting a kind of “structure vs expansion” dichotomy as mentioned earlier. In one case, we apply a refinement of the structure theory for transitive graphs of polynomial growth in order to implement the methods of Contreras, Martineau, and Tassion. (In fact, we had to engage with this structure theory more substantially than did previous works on percolation, and in particular, Hutchcroft and I ended up writing a companion paper about groups of polynomial growth.) In another case, we exploit the hypothesis that $|B_n|$ is large in order to implement a new percolation argument that works more efficiently when we have “more vertices packed in a smaller space”. In a third case, we use the trajectories of random walks to probabilistically prove deterministic facts about the geometry of G around scale n . These deterministic facts then play the role that structure theory did in the first case, again allowing us to implement the methods of Contreras, Martineau, and Tassion.

The two key difficulties were that (1) different transitive graphs can have very different geometric properties, e.g. a regular tree and a Euclidean lattice, and therefore require different arguments, and (2) a single transitive graph can have very different geometric properties from one scale to the next, e.g. a graph with (eventually) polynomial growth may look like a tree up to a large finite scale, so we need to thrifitly exploit geometric hypotheses that can only be assumed to hold at a *single scale*.

1.5 Finite graphs

The story of percolation theory on infinite transitive graphs is only one half of the story of this model. In 1960, roughly when Broadbent and Hammersley introduced so-called percolation

theory on Euclidean lattices, Erdős and Rényi introduced percolation on the complete graph K_n with n vertices. This is the well-known *Erdős-Rényi* model (or simply *random graph*) model in combinatorics. In this setting, one studies the threshold for the emergence of a *giant* (as opposed to *infinite*) cluster under percolation. The fundamental result is that percolation on K_n undergoes a phase transition around $p = \frac{1}{n}$ in the sense that for every fixed $\varepsilon > 0$, there exists $\delta > 0$ such that

$$\lim_{n \rightarrow \infty} \mathbb{P}_{\frac{1+\varepsilon}{n}}^{K_n} (\text{the largest cluster contains } \geq \delta n \text{ vertices}) = 1,$$

whereas for every $\delta > 0$,

$$\lim_{n \rightarrow \infty} \mathbb{P}_{\frac{1-\varepsilon}{n}}^{K_n} (\text{the largest cluster contains } \geq \delta n \text{ vertices}) = 0.$$

We say that the sequence $(1/n)_{n \geq 1}$ is the *percolation threshold* for the sequence of complete graphs. (While we call this *the* threshold, note that a percolation threshold is only ever unique up to multiplication by $1 + o(1)$.)

There has since been a tremendous amount of work on this model and on percolation on the sequence of hypercubes $(\{0, 1\}^n)_{n \geq 1}$. Note that hypercubes and complete graphs are both examples of transitive graphs⁴. We have been working to develop a theory of percolation on general *finite* transitive graphs that forms a bridge between these canonical finite graph models and the rich theory of percolation on infinite transitive graphs. As in the Erdős-Rényi model, in this theory, we study the phase transition for the emergence of a *giant* cluster.

Roughly speaking, we can think of the theory of percolation on infinite transitive graphs as the theory of percolation on microscopic (i.e. $O(1)$) scales in bounded-degree finite transitive graphs. In this sense, the finite graph theory generalises the infinite graph theory. (A limitation of this maxim is that not every infinite transitive graph can be locally approximated by finite transitive graphs.) In particular, certain basic questions in the finite graph theory have no natural analogues in the infinite graph theory. For example, the uniqueness/non-uniqueness of giant clusters is not directly related to the uniqueness/non-uniqueness of infinite clusters, which is instead related to the microscopic metric distortion of giant clusters.

Hutchcroft and I established the basic features of the supercritical phase of percolation on finite transitive graphs. Together, we showed that the giant cluster is almost surely unique (except in the trivial cases), resolving a conjecture of Benjamini from 2001. Our argument has already been applied by others to spherical Gaussian ensembles, and with Hutchcroft, we are extending

⁴There has also been a great deal of work on families of finite graphs that are not necessarily transitive, e.g. arbitrary dense graphs.

our arguments to establish results about the Potts model that are new even for tori. In a sequel to this work on uniqueness, with Hutchcroft we established that the density of this unique giant concentrates around some limit and fully characterised those graphs for which this limiting density is *mean-field* (meaning *like on a tree*). Our argument in this sequel relies on a new application of sharp-threshold theory, applied (unusually) to study events that obviously *do not* have sharp thresholds.

This analysis of the supercritical phase did not address the basic question of whether, (in the language of Bollobás, Borgs, Chayes, Riordan for example), percolation on finite transitive graphs *has a phase transition*, meaning that a giant cluster emerges suddenly at some threshold. By combining the supercritical analysis with Vanneuville's new proof of an old result about infinite graphs, we proved that such a phase transition does always occur (except in the trivial cases). Note that this result (and those about the supercritical phase in the previous paragraph) hold for all finite transitive graphs, with possibly diverging vertex degrees. In particular, this result recovers the classical fact that percolation on large complete graphs or hypercubes undergoes a phase transition but via very soft and general arguments.

The theory of percolation on finite transitive graphs is intimately linked with the theory on infinite transitive graphs if we restrict our study to families of finite transitive graphs with uniformly bounded vertex degrees. Indeed, with respect to the local topology, every infinite set \mathcal{G} of finite transitive graphs with bounded degrees is relatively compact, and every graph in the boundary of \mathcal{G} is infinite. Our very recent work combined several ideas from the works discussed above about finite graphs and about Schramm's locality conjecture to establish that in some sense the sudden emergence of a giant cluster for percolation on bounded-degree finite transitive graphs is "*the same*" phase transition as the usual emergence of an infinite cluster for percolation on infinite transitive graphs. More precisely, we considered the following pair of questions, which are actually equivalent:

1. Does percolation on a large bounded-degree finite transitive graph G have a sharp phase transition? This means that in the subcritical phase of percolation, the largest cluster is with high probability not just sublinear but *logarithmic* in the total number of vertices in G .
2. If a finite transitive graph G and an infinite transitive graph H are close in the local sense, does the critical point for the emergence of a giant cluster in G approximately coincide with the critical point for the emergence of an infinite cluster in H ?

Unfortunately, the answer to both of these questions in general is *no*. For example, take the sequence $(\mathbb{Z}_n \times \mathbb{Z}_{f(n)})_{n=1}^\infty$ for any $f : \mathbb{N} \rightarrow \mathbb{N}$ growing fast. This sequence always converges locally to \mathbb{Z}^2 , where the critical point for the emergence of an infinite cluster is $p_c = \frac{1}{2}$. On the other hand, provided that f grows sufficiently fast, the threshold for the emergence of a giant cluster in $\mathbb{Z}_n \times \mathbb{Z}_{f(n)}$ will be as in the sequence of cycles, around $p_c = 1$. Moreover, for percolation of any fixed parameter $p \in (\frac{1}{2}, 1)$ on $\mathbb{Z}_n \times \mathbb{Z}_{f(n)}$, the order of the largest cluster will then typically be much larger than logarithmic but much smaller than linear in the total number of vertices. The problem is that these graphs are long and thin, coarsely resembling long cycles. In particular, after suitably rescaling, their graph metrics (rapidly) converge in the Gromov-Hausdorff metric to the unit circle. We proved that this is in fact the only possible obstacle. This theorem provides a direct way to translate results and conjectures about percolation on general (non-one-dimensional) infinite transitive graphs into statements about percolation on general (non-one-dimensional) finite transitive graphs.

PAPERS INCLUDED

The first paper explores the question: when does percolation on an infinite graph G undergo a non-trivial phase transition?

1. Counting minimal cutsets and $p_c < 1$

Joint with Severo and Tassion

The next two papers resolve Schramm's locality conjecture that the value of the critical point for percolation on a given infinite transitive graph G is typically entirely determined by the *local* (i.e. microscopic) geometry of G . The first paper gives the proof itself, and the second paper establishes a required ingredient about groups of polynomial growth.

2. The critical percolation probability is local

Joint with Hutchcroft

3. Uniform finite presentation for groups of polynomial growth

Joint with Hutchcroft

(Published in Discrete Analysis: <https://doi.org/10.19086/da.127778>)

The next two papers establish the basic features of the giant clusters that form in the supercritical phase of percolation on finite transitive graphs.

4. Supercritical percolation on finite transitive graphs I: Uniqueness of the giant component

Joint with Hutchcroft

(Published in Duke Mathematical Journal: <https://doi.org/10.1215/00127094-2023-0066>)

5. Supercritical percolation on finite transitive graphs II: Concentration, locality, and equicontinuity of the giant's density

Joint with Hutchcroft

The next paper establishes that percolation on a finite transitive graph typically *undergoes a phase transition* in the sense that a giant cluster emerges *suddenly* around a single critical point.

6. Existence of a percolation threshold on finite transitive graphs

(Published in International Mathematics Research Notices: <https://doi.org/10.1093/imrn/rnad222>)

The next paper establishes that this phase transition for percolation on a bounded-degree finite transitive graph G is typically *sharp* in the sense that the subcritical clusters are logarithmically small. Equivalently, if G approximates an infinite transitive graph H , then the critical point for the emergence of a giant cluster in G approximately coincides with the critical point for the emergence of an infinite cluster in H .

7. Sharpness and locality for percolation on finite transitive graphs

COUNTING MINIMAL CUTSETS AND $p_c < 1$

Joint work with Franco Severo and Vincent Tassion

Abstract

We prove two results concerning percolation on general graphs.

- We establish the converse of the classical Peierls argument: if the critical parameter for (uniform) percolation satisfies $p_c < 1$, then the number of minimal cutsets of size n separating a given vertex from infinity is bounded above exponentially in n . This resolves a conjecture of Babson and Benjamini from 1999.
- We prove that $p_c < 1$ for every uniformly transient graph. This solves a problem raised by Duminil-Copin, Goswami, Raoufi, Severo and Yadin, and provides a new proof that $p_c < 1$ for every transitive graph of superlinear growth.

2.1 Main results

Let $G = (V, E)$ be an infinite, connected, locally finite graph. A set of edges $F \subset E$ is called a *cutset* from a vertex v to ∞ if v belongs to a finite connected component of $(V, E \setminus F)$. A cutset is called *minimal* if no proper subset of it is a cutset. Let $Q_n(v)$ be the set of minimal cutsets from v to ∞ of cardinality n and consider the quantity

$$q_n := \sup_{v \in V} |Q_n(v)|. \quad (2.1.1)$$

Here $|\emptyset| := 0$. We emphasize that $q_n = \infty$ is possible, for example for $G = \mathbb{Z}$ and $n = 2$. In this paper, we are interested in cases where the number of cutsets q_n grows at most exponentially with n , and we define

$$\kappa(G) := \sup_{n \geq 1} q_n^{1/n}. \quad (2.1.2)$$

Let P_p denote (Bernoulli bond) percolation of parameter $p \in [0, 1]$ on G , where each edge is open with probability p independently of the other edges. Consider the percolation probabilities $\theta_v(p) := P_p(v \leftrightarrow \infty)$, where $v \leftrightarrow \infty$ denotes the event that v belongs to an infinite open connected component. We define the critical parameter for *uniform* percolation as

$$p_c^*(G) := \inf\{p \in [0, 1] : \theta^*(p) > 0\}, \quad (2.1.3)$$

where $\theta^*(p) := \inf_{v \in V} \theta_v(p)$.

By the classical Peierls argument [Pei36a] if $\kappa(G) < \infty$, then percolation on G has a uniformly percolating phase in the sense that $p_c^*(G) < 1$. Our first theorem establishes the converse.

For every infinite, connected, locally finite graph G we have

$$p_c^*(G) < 1 \iff \kappa(G) < \infty. \quad (2.1.4)$$

Currently, the geometric condition $\kappa(G) < \infty$ is not well understood. Our second result gives a sufficient condition based on the simple random walk. Given a vertex v , let \mathbb{P}_v be the law of a simple random walk $(X_t)_{t=0}^\infty$ on G starting at v . We say that G is uniformly transient if

$$\inf_{v \in V} [d_v \cdot \mathbb{P}_v(\forall t \geq 1 : X_t \neq v)] > 0, \quad (2.1.5)$$

where d_v denotes the degree of v .

Let G be an infinite, connected, locally finite graph. If G is uniformly transient, then $\kappa(G) < \infty$.

2.2 Consequences and comments

In this section, all graphs are assumed to be infinite, connected, and locally finite. Given a set of vertices S in a graph $G = (V, E)$, we define the boundary ∂S to be the set of all edges $\{u, v\} \in E$ such that $u \in S$ but $v \notin S$, and we define the weight $|S|_G := \sum_{u \in S} d_u$. The *isoperimetric dimension* of G is given by

$$\text{Dim}(G) := \sup \left\{ d \geq 1 : \inf_{\substack{S \subseteq V \\ 0 < |S| < \infty}} \frac{|\partial S|}{|S|_G^{\frac{d-1}{d}}} > 0 \right\}.$$

1. We remark that the uniform critical parameter $p_c^*(G)$ slightly differs from the most classical (non-uniform) one given by $p_c(G) := \inf\{p \in [0, 1] : \theta(p) > 0\}$, where $\theta(p) := \sup_{v \in V} \theta_v(p)$. However, these notions often coincide, such as for (quasi-)transitive graphs. See the introduction of [Dum+20a] for a survey of the rich history of the “ $p_c < 1$ ” question and its place in statistical mechanics. Let us just recall that all of the results about percolation here can be translated into analogous statements about many other models, most notably the Ising model.
2. Duminil-Copin, Goswami, Raoufi, Severo, and Yadin proved that every quasi-transitive graph of superlinear growth satisfies $p_c < 1$ [Dum+20a]. This had previously been a long-standing conjecture of Benjamini and Schramm [BS96a]. In fact, the authors of [Dum+20a] established

that $p_c^* < 1$ for every (not necessarily transitive) bounded degree graph G satisfying $\text{Dim}(G) > 4$, and this was known to imply the conjecture about transitive graphs by the classical works of Gromov [Gro81a] and Trofimov [Tro84a].

3. Theorem 2.1 establishes that $p_c^* < 1$ for every graph G satisfying $\text{Dim}(G) > 2$, since such graphs are uniformly transient (see e.g. [LP16a, Theorem 6.41]). We therefore obtain stronger results than [Dum+20a], through a completely new proof. Theorem 2.1 fully realises the idea at the heart of [Dum+20a] to exploit the transience of a simple random walk to prove $p_c < 1$. In particular, we resolve [Dum+20a, Problem 1.4].
4. Our proofs of Theorems 1 and 2 can also be run on *finite* graphs to establish the analogous results about *giant* clusters. (See [HT21e] for background.) In this setting, to define q_n , one should instead count the number of minimal cutsets of cardinality n from a vertex v to another vertex u (and take the supremum over all choices for distinct u and v). The corresponding notion of uniform transience for a given family of finite graphs is that there exists a constant $C < \infty$ such that every graph $G = (V, E)$ in the family satisfies

$$\max_{u, v \in V} \mathcal{R}_G(u, v) \leq C,$$

where $\mathcal{R}_G(u, v)$ denotes the effective resistance from u to v in the graph G .

5. Babson and Benjamini conjectured that $\kappa < \infty$ for every transitive¹ graph of superlinear growth [BB99a]. Notice that this purely geometric conjecture is a priori stronger than the above $p_c < 1$ conjecture of Benjamini and Schramm. Babson and Benjamini verified their conjecture in the special case of Cayley graphs of finitely presented groups by establishing that minimal cutsets in such graphs are coarsely connected. By [Tim07; Gro81a; Tro84a] (see also [CMT24, Lemma 2.1]), this extends to all transitive graphs satisfying $\text{Dim}(G) < \infty$. Given these results, it suffices to show that $\kappa < \infty$ for every transitive graph satisfying $\text{Dim}(G) = \infty$. Theorem 2.1 therefore resolves the $\kappa < \infty$ conjecture of Babson and Benjamini. (Alternatively, taking the results of [Dum+20a] for granted, this conjecture follows from Theorem 2.1.)
6. We establish the existence of a universal constant $\varepsilon > 0$ such that every transitive graph G satisfies $p_c = 1$ or $p_c \leq 1 - \varepsilon$. When G is recurrent, this follows from the proof of [HT21e,

¹In fact, Babson and Benjamini originally made this conjecture in the case of Cayley graphs, and Benjamini later extended this conjecture to allow arbitrary transitive graphs.

Theorem 1.7], and when G is transient, this follows from our proof of Theorem 2.1 because there exists a universal constant $c > 0$ such that a simple random walk in any transient transitive graph has probability at least c never to return to where it started [TT20a, Corollary 1.3]. Previous works had established this result if ε is allowed to depend on the degree of vertices in G [HT21e, Theorem 1.7], or if we instead consider site percolation on a Cayley graph [PS23a; Lyo+23a]. By the proof of Theorem 2.1, we also obtain a universal constant $K < \infty$ such that $\kappa < K$ for every transitive graph of superlinear growth.

7. Much work has been motivated by a desire to find a sharp geometric criterion for a graph G to satisfy $p_c < 1$. Indeed, a well-known open conjecture of Benjamini and Schramm is that every (not necessarily transitive) graph G with $\text{Dim}(G) > 1$ satisfies $p_c < 1$ [BS96a]. We were very surprised to find that the geometric criterion $\kappa < \infty$ (which is arguably simpler and more natural than the isoperimetric criterion) is not just sharp but *exact*. Nevertheless, in light of Theorem 2.1 and this conjecture of Benjamini and Schramm, we encourage the reader to investigate the following: Every graph G with $\text{Dim}(G) > 1$ satisfies $\kappa < \infty$.

The Peierls argument can be used to deduce results that are (a priori) much stronger than $p_c < 1$. To explore these, it helps to consider the *isoperimetric profile* ψ of a graph $G = (V, E)$, given by

$$\psi(n) := \inf_{\substack{S \subseteq V \\ n \leq |\partial S|_G < \infty}} |\partial S|.$$

8. Every graph $G = (V, E)$ satisfying $\kappa < \infty$ admits a *strongly percolating* phase in the sense that for all $p \in (1 - 1/\kappa, 1]$, there is a constant $c > 0$ such that

$$\begin{aligned} \mathbb{P}_p(S \leftrightarrow \infty) &\leq e^{-c\psi(|S|)} && \text{for every finite set } S \subseteq V; \\ \mathbb{P}_p(n \leq |C_v| < \infty) &\leq e^{-c\psi(n)} && \text{for every } n \geq 1 \text{ and } v \in V. \end{aligned} \tag{2.2.1}$$

Thus our work resolves [Dum+20a, Problem 1.6] and implies that percolation on every transitive graph of superlinear growth has a strongly percolating phase. It remains an important open problem to establish that on these graphs, such bounds hold for *all* $p \in (p_c, 1]$. Indeed, this is the “upper bound” half of [HH21b, Conjecture 5.1].

9. Conversely, our proof of Theorem 2.1 (more precisely, Proposition 2.5.1) can be used to show that for every transitive graph $G = (V, E)$ and for every $p > p_c$, there is a constant $c > 0$ such

that

$$\mathbb{P}_p(n \leq |C_v| < \infty) \geq e^{-c\psi(n)} \quad \text{for every } n \geq 1 \text{ and } v \in V.$$

This establishes the “lower bound” half of [HH21b, Conjecture 5.1].

10. A major motivation for studying *anchored* isoperimetric inequalities for graphs and manifolds is the belief that — unlike (*uniform*) isoperimetric inequalities — anchored inequalities should typically be robust under small perturbations of the space [BLS99, Section 6]. We obtain the following concrete statement to this effect by combining Theorem 2.1 with an argument of Pete [Pet08, Theorem 4.1]: for every graph G satisfying $p_c^*(G) < 1$, there exists $\varepsilon > 0$ such that if G satisfies a d -dimensional anchored isoperimetric inequality for any $d \geq 1$ (or f -anchored isoperimetric inequality for any function f) then so does every infinite cluster formed by percolation of parameter $1 - \varepsilon$.
11. By combining the previous item with Theorem 2 and results of Thomassen [Tho92] and Pemantle and Peres [PP96], we deduce that for every graph $G = (V, E)$ with $\text{Dim}(G) > 2$, and for every probability measure μ on $(0, \infty)$, the random weighted network (V, C) with $C = (C(e) : e \in E) \sim \mu^{\otimes E}$ is almost surely transient. (This was previously known if $\text{Dim}(G) > 4$ [Hut23a].)
12. A standard analysis of *Karger’s algorithm* from computer science establishes that every finite graph $G = (V, E)$ with exactly n vertices contains at most $\binom{n}{2}$ *minimum cuts*, i.e. sets of edges F such that $(V, E \setminus F)$ is disconnected but there is no set of edges F' with $|F'| < |F|$ such that $(V, E \setminus F')$ is also disconnected. In the same spirit, in the present paper, we design randomized algorithms to instead count *minimal cutsets*.

Acknowledgements: We are very grateful to Benny Sudakov for telling us about Karger’s algorithm from computer science. This seed is what prompted us to investigate probabilistic approaches to bounding the number of minimal cutsets, ultimately leading to the present work. We thank Itai Benjamini for bringing the $\kappa(G) < \infty$ question to our attention in the first place. PE is grateful for the hospitality provided by ETH Zurich during this project. This project was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No 851565). FS was supported by the ERC grant Vortex (No 101043450).

2.3 Background and notation

In this section, we fix $G = (V, E)$ a locally finite, connected graph.

Paths and connectivity

Let $S \subset V$, $u, v \in S$. A path from u to v in S is a finite sequence $\gamma = (\gamma_0, \gamma_1, \dots, \gamma_\ell)$ of distinct vertices of S such that $\gamma_0 = u$, $\gamma_\ell = v$ and $\{\gamma_{i-1}, \gamma_i\} \in E$ for every $i \in \{1, \dots, \ell\}$. When such a path exists, we say that u is connected to v in S . By extension, a set A is said to be connected to a set B in S if there exists a vertex of A that is connected to a vertex of B in S . A path from u to ∞ in S is an infinite sequence of distinct vertices $\gamma_0, \gamma_1, \dots$ in S such that $\gamma_0 = u$ and $\{\gamma_{i-1}, \gamma_i\} \in E$ for every $i \in \{1, 2, \dots\}$. When such a path exists, we say that u is connected to ∞ in S .

Exposed boundary

Let $S \subset V$ be a finite set. The exposed boundary of S is the set $\partial_\infty S$ of all the edges $\{u, v\}$ such that $u \in S$ and v is connected to ∞ in $V \setminus S$. Notice that the exposed boundary is a subset of the standard boundary defined at the beginning of Section 2.2: for every finite set $S \subset V$, we have $\partial_\infty S \subset \partial S$.

Percolation configurations

An element $\omega \in \{0, 1\}^E$ is called a percolation configuration. Given such a configuration, an edge $e \in E$ is said to be open if $\omega(e) = 1$ and closed if $\omega(e) = 0$. By extension, a path is said to be open if all its edges are open. The cluster of a vertex $u \in V$ is the connected component of u in the graph $(V, \{e \in E : \omega(e) = 1\})$.

Percolation events

A measurable subset $A \subset \{0, 1\}^E$ is called a percolation event. Given $S \subset V$ and $u, v \in S$, we denote by $u \xleftrightarrow{S} v$ the event that there exists an open path from u to v in S , and simply write $u \leftrightarrow v$ when $S = V$. Finally, $u \leftrightarrow \infty$ denotes the event that there exists an open path from u to ∞ in V .

Percolation measures

A percolation measure on G is a probability measure on the product space $\{0, 1\}^E$. For $p \in [0, 1]$, we denote by P_p the standard Bernoulli percolation measure, under which each edge is open with probability p independently of the other edges.

Positive association

A percolation event \mathcal{E} is called increasing if for all percolation configurations ω, ξ satisfying $\omega \leq \xi$ for the standard product (partial) ordering, we have $\omega \in \mathcal{E} \implies \xi \in \mathcal{E}$. Typical examples of

increasing events are the connection events (such as $u \xleftrightarrow{S} v$) introduced above. A percolation measure P is said to be positively associated if

$$P[\mathcal{E} \cap \mathcal{F}] \geq P[\mathcal{E}]P[\mathcal{F}] \quad (2.3.1)$$

for all increasing events \mathcal{E}, \mathcal{F} . This property is often referred to as the FKG inequality. We will use that Bernoulli percolation P_p is positively associated (for every fixed $p \in [0, 1]$) as established by Harris [Har60].

2.4 Exposed boundaries and cutsets

In this section, we fix $G = (V, E)$ an infinite, connected, locally finite graph. In our paper, we will use that minimal cutsets can be obtained by considering the exposed boundary of finite connected sets. In this section, we recall some well-known facts relating the two notions. The first elementary result is that the exposed boundary of a finite connected set is a minimal cutset.

Lemma 2.4.1. Let $S \subset V$ be a finite connected set. For every $u \in S$, $\partial_\infty S$ is a minimal cutset from u to ∞ .

Proof. Any path from u to ∞ in V must traverse an edge in $\partial_\infty S$ (consider the last edge traversed by this path intersecting S). Therefore, $\partial_\infty S$ is a cutset from u to ∞ . To prove that it is minimal, consider an edge $e \in \partial_\infty S$. Since S is connected, there exists a path from u to an endpoint of e in S and by definition of the exposed boundary, there must exist a path from the other endpoint of e to ∞ in $V \setminus S$. The concatenation of these two paths with e connects u to ∞ without using any edges of $\partial_\infty S$ other than e . Hence $\partial_\infty S \setminus \{e\}$ is not a cutset from u to ∞ . \square

The second elementary result identifies the exposed boundary under some simple conditions.

Lemma 2.4.2. Let $u \in V$, let Π be a minimal cutset from u to ∞ . Let A be the connected component of u in $(V, E \setminus \Pi)$ and $B = \{e \cap A, e \in \Pi\}$ be the set of inner vertices of Π . For every set S of vertices, we have

$$(B \subset S \subset A) \implies (\partial_\infty S = \Pi). \quad (2.4.1)$$

Proof. Since A is a maximal connected set in $(V, E \setminus \Pi)$, all the edges at the boundary of A belong to Π , and therefore $\partial_\infty A \subset \partial A \subset \Pi$. By Lemma 2.4.1, $\partial_\infty A$ is a cutset from u to ∞ , hence, by the minimality of Π , the two inclusions above must be equalities:

$$\partial_\infty A = \partial A = \Pi. \quad (2.4.2)$$

Now, let S be a set satisfying $B \subset S \subset A$. Let $e \in \Pi$. Since $e \in \partial_\infty A$, one endpoint of e must belong to B and the other endpoint is connected to ∞ in $V \setminus A$. Therefore, by hypothesis, one endpoint of e belongs to S and the other endpoint is connected to ∞ in $V \setminus S$. This proves the inclusion

$$\Pi \subset \partial_\infty S. \quad (2.4.3)$$

Let $e \in \partial_\infty S$. Let u be the endpoint of e in S , and let v be the endpoint of e connected to ∞ in $V \setminus S$. Then, by hypothesis, $u \in A$ and v is connected to ∞ in $V \setminus B$. Since $\Pi = \partial A$, every edge in Π intersects A and hence intersects B . Therefore, there must exist an infinite path starting at v in the subgraph $(V, E \setminus \Pi)$. In particular, $v \notin A$, and hence $e \in \partial A = \Pi$. This proves that the inclusion above must be an equality.

□

2.5 Full connectivity via positive association

In this section, we consider the following problem: Let B be a finite set in a graph, and P be a percolation measure. What is the probability that all the vertices of B are all connected to each other? Or, in other words, what is the probability that all the vertices of B lie in the same cluster? We prove that this probability is at least exponential in the size of B when the measure is positively associated, and the probability for a point to be connected to B is uniformly lower bounded. This result, formally stated below, will allow us to construct random sets with a prescribed boundary.

Proposition 2.5.1. Let $G = (V, E)$ be a finite, connected graph. Let P be a positively associated percolation measure on G . Let $B \subset V$, let $\theta, p \in (0, 1]$ and suppose that $P(u \leftrightarrow B) \geq \theta$ for every $u \in V$, and $P(e \text{ is open}) \geq p$ for every $e \in E$. Then for every $o \in V$,

$$P\left(\bigcap_{b \in B} \{o \leftrightarrow b\}\right) \geq c^{|B|},$$

where $c := \left(\frac{p\theta}{2}\right)^{3/\theta}$.

Proof. Say that a finite sequence of vertices x_1, \dots, x_k is *chained* if $x_1 = o$ and for all $i \in \{2, \dots, k\}$,

$$\frac{p\theta}{2} \leq P(x_i \leftrightarrow \{x_1, \dots, x_{i-1}\}) \leq \frac{\theta}{2}. \quad (P1)$$

Since there exists at least one chained sequence (take $k = 1$) and V is finite, there must exist a chained sequence x_1, \dots, x_k that is *maximal* in the sense that for every vertex x_{k+1} , the sequence x_1, \dots, x_{k+1} is not chained. Fix a maximal chained sequence x_1, \dots, x_k , and let $X := \{x_1, \dots, x_k\}$.

We claim that, in addition to (P1), this sequence satisfies the following two properties, where $n := |B|$:

$$\forall u \in V \quad P(u \leftrightarrow X) \geq \frac{\theta}{2}, \quad (P2)$$

$$k \leq \frac{2n}{\theta}. \quad (P3)$$

To prove (P2), consider the set of vertices $W \subset V$ that are connected to X with probability at least $\theta/2$ and suppose for contradiction that $W \neq V$. Since W is non empty (because $X \subset V$) and G is connected, we can consider an edge $\{u, v\}$ such that $u \in W$ and $v \notin W$. By positive association,

$$\theta/2 > P(v \leftrightarrow X) \geq P(\{u, v\} \text{ is open}) \cdot P(u \leftrightarrow X) \geq \frac{p\theta}{2}.$$

In particular, x_1, \dots, x_k, v is a chained sequence, contradicting the maximality of x_1, \dots, x_k .

We now prove (P3). To this aim, for each $i \in \{1, \dots, k\}$, let N_i denote the number of clusters that intersect both $\{x_1, \dots, x_i\}$ and B . For every $i \in \{2, \dots, k\}$, the increment $N_i - N_{i-1}$ is equal to 1 if x_i is connected to B but not to the previous points $\{x_1, \dots, x_{i-1}\}$, and it is equal to 0 otherwise. Therefore, for every $i \in \{2, \dots, k\}$, we have the deterministic inequality

$$N_i - N_{i-1} \geq \mathbf{1}_{x_i \leftrightarrow B} - \mathbf{1}_{x_i \leftrightarrow \{x_1, \dots, x_{i-1}\}}. \quad (2.5.1)$$

Taking the expectation, using our hypothesis and (P1), for every $i \in \{2, \dots, k\}$, we get

$$E(N_i) - E(N_{i-1}) \geq \underbrace{P(x_i \leftrightarrow B)}_{\geq \theta} - \underbrace{P(x_i \leftrightarrow \{x_1, \dots, x_{i-1}\})}_{\leq \theta/2} \geq \theta/2. \quad (2.5.2)$$

Summing over $i \in \{2, \dots, k\}$ and using $E(N_1) = P(x_1 \leftrightarrow B) \geq \theta \geq \theta/2$, we get $E[N_k] \geq \frac{\theta}{2}k$. Since N_k is deterministically bounded above by $|B| = n$, this concludes the proof of (P3).

We now explain how the three properties above of the chained sequence imply the desired lower bound in the proposition. First, we estimate the event that all the vertices of X are connected to o : By (P1), (P3) and positive association, we have

$$P\left(\bigcap_{u \in X} \{o \leftrightarrow u\}\right) \geq \prod_{i=2}^k P(x_i \leftrightarrow \{x_1, \dots, x_{i-1}\}) \geq \left(\frac{p\theta}{2}\right)^{k-1} \geq \left(\frac{p\theta}{2}\right)^{\frac{2n}{\theta}}. \quad (2.5.3)$$

Second, we estimate the event that all the vertices of B are connected to X : By (P2) and positive association, we have

$$P\left(\bigcap_{b \in B} \{b \leftrightarrow X\}\right) \geq \left(\frac{\theta}{2}\right)^n.$$

If all the vertices of X are connected to o and all the vertices of B are connected to X , then all the vertices of B are connected to o . Hence, by the two displayed equations above and positive association, we obtain

$$P\left(\bigcap_{b \in B} \{o \leftrightarrow b\}\right) \geq P\left(\bigcap_{u \in X} \{o \leftrightarrow u\}\right) \cdot P\left(\bigcap_{b \in B} \{b \leftrightarrow X\}\right) \geq \left(\frac{p\theta}{2}\right)^{\frac{2n}{\theta}} \cdot \left(\frac{\theta}{2}\right)^n \geq c^n,$$

where $c := \left(\frac{p\theta}{2}\right)^{3/\theta}$. □

2.6 Proof of Theorem 2.1

Let $G = (V, E)$ be an infinite, connected, locally finite graph. In this section, we prove Theorem 2.1, in the following form.

$$(\exists p < 1 \exists \theta > 0 \forall u \in V \quad P_p(u \leftrightarrow \infty) \geq \theta) \iff (\exists K < \infty \forall u \in V \forall n \geq 1 \quad |Q_n(u)| \leq K^n). \quad (2.6.1)$$

The implication \Leftarrow is well-known, and follows from the Peierls argument [Ben13a, Theorem 4.11], which we now recall for completeness. Let $u \in V$. If the cluster of u is finite, then by Lemma 2.4.1, its exposed boundary is a finite minimal cutset from u to ∞ , and all its edges are closed. Hence, by the union bound, for every $p \in [0, 1]$ we have

$$P_p(|C_u| < \infty) \leq \sum_{n \geq 1} q_n (1-p)^n. \quad (2.6.2)$$

If $q_n \leq K^n$ for some constant $K < \infty$, then the right hand side above converges to 0 as p tends to 1. Since the bound is uniform in u , there exists $p < 1$ such that

$$\forall u \in V \quad P_p(u \leftrightarrow \infty) \geq 1/2. \quad (2.6.3)$$

We now prove the implication \Rightarrow . Fix $\theta, p \in (0, 1)$ such that $P_p(u \leftrightarrow \infty) \geq \theta$ for every $u \in V$. Fix $o \in V$ and $n \geq 1$. Writing C for the cluster of o , we show that for every minimal cutset Π from o to ∞ with $|\Pi| = n$,

$$P_p(\partial_\infty C = \Pi) \geq 1/K^n, \quad (2.6.4)$$

where $K = K(p, \theta) \in (0, \infty)$ is a finite constant depending on p and θ only (in particular it does not depend on the chosen vertex o). This concludes the proof since

$$1 \geq \sum_{\Pi \in Q_n(o)} P_p(\partial_\infty C = \Pi) \stackrel{(2.6.4)}{\geq} |Q_n(o)|/K^n. \quad (2.6.5)$$

Let us now prove the lower bound (2.6.4). As in Lemma 2.4.2, let A be the connected component of o in $(V, E \setminus \Pi)$ and B the set of inner vertices of Π . Since any infinite open path from a vertex $u \in A$ must intersect B before exiting A , the hypothesis $P_p(u \leftrightarrow \infty) \geq \theta$ implies

$$\forall u \in A \quad P_p(u \xrightarrow{A} B) \geq \theta. \quad (2.6.6)$$

Let \mathcal{E} be the event that every vertex in B is connected to o by an open path in A . By Proposition 2.5.1 applied to the finite subgraph of G induced by A , we have $P_p(\mathcal{E}) \geq c^n$, where $c = (p\theta/2)^{3/\theta} > 0$. Let \mathcal{F} be the event that all the edges of Π are closed. By independence, we have

$$P_p(\mathcal{E} \cap \mathcal{F}) = P_p(\mathcal{E})P_p(\mathcal{F}) \geq c^n(1-p)^n. \quad (2.6.7)$$

If the event $\mathcal{E} \cap \mathcal{F}$ occurs, then the cluster C of o satisfies $B \subset C \subset A$. Hence, by Lemma 2.4.2 we must have $\partial_\infty C = \Pi$. This concludes that

$$P_p(\partial_\infty C = \Pi) \geq P_p(\mathcal{E} \cap \mathcal{F}) \geq c^n(1-p)^n, \quad (2.6.8)$$

which establishes the desired lower bound (2.6.4) with $K = \frac{1}{c(1-p)} = \frac{1}{(p\theta/2)^{3/\theta}(1-p)}$.

2.7 A covering lemma for Markov chains

In this section, we give conditions under which a killed Markov chain survives long enough to visit every state and then return to its initial state². We will apply this in the next section to prove Theorem 2.1. Here $[n]$ denotes the set $\{1, \dots, n\}$.

Lemma 2.7.1. Let $n \geq 1$. Let $P = (p_{i,j})_{i,j \in [n]}$ be a symmetric matrix of non-negative entries such that³ $\sum_{j \in [n]} p(i, j) \leq 1$ for all $i \in [n]$. Let Γ be the set of all sequences $\gamma = (\gamma_0, \gamma_1, \dots, \gamma_k)$ in $[n]$ (for any $k \geq 1$) with $\gamma_0 = 1$ such that the unique element $i \in [k]$ satisfying both $\gamma_i = 1$ and $\{\gamma_0, \gamma_1, \dots, \gamma_i\} = [n]$ is $i = k$. For every such sequence γ , define

$$p(\gamma) := \prod_{i=1}^k p(\gamma_{i-1}, \gamma_i).$$

For each $\varepsilon > 0$, if every non-empty proper subset I of $[n]$ satisfies

$$\sum_{i \in I} \sum_{j \in [n] \setminus I} p(i, j) \geq \varepsilon, \quad (2.7.1)$$

²In fact, we lower bound the probability that this occurs in $\leq 2n - 2$ steps (which is optimal), where n is the number of states. Contrast this with [BGM13; DK21], both called *Linear cover time is exponential unlikely*; we give conditions under which linear cover time is exponentially *likely*.

³We can think of P as the transition matrix of a Markov chain which is killed at i with probability $1 - \sum_j p(i, j)$.

then $\delta := \frac{\varepsilon^2}{16e^2}$ satisfies

$$\sum_{\gamma \in \Gamma} p(\gamma) \geq \delta^n.$$

Proof. Let $e_1, \dots, e_{2n-2} \in [n]^2 \sqcup \{\emptyset\}$ be an iid sequence of random variables such that for all $u, v \in [n]$,

$$\mathbb{P}(e_1 = (u, v)) = \frac{p(u, v)}{n}.$$

Such random variables exist because these probabilities sum to at most 1. Let H be the *undirected* multigraph with vertex set $[n]$ and edges e_1, \dots, e_{2n-2} . Even though $[n]^2$ consists of *ordered* pairs, we think of each $e_i \in [n]^2$ as encoding an *undirected* edge, loops allowed. (When $e_i = \emptyset$, we simply do not include an edge.)

Consider the iid spanning subgraphs H_1 and H_2 of H that contain only the edges e_1, \dots, e_{n-1} and e_n, \dots, e_{2n-2} , respectively. We will lower bound the probability that each of these graphs is connected. Consider any $k \in [n-1]$. Suppose that we are given all of the connected components C_1, \dots, C_r of the spanning subgraph of H that contains only the edges e_1, \dots, e_{k-1} . If $r \geq 2$, then the conditional probability that e_k connects two of these components is

$$\sum_{z=1}^r \sum_{i \in C_z} \sum_{j \in [n] \setminus C_z} \frac{p(i, j)}{n} \stackrel{(2.7.1)}{\geq} \frac{r\varepsilon}{n}.$$

Therefore by induction on k , and by using the elementary bound $\frac{n^n}{n!} \leq e^n$ in the third inequality,

$$\mathbb{P}(H_1 \text{ is connected}) \geq \prod_{r=2}^n \frac{r\varepsilon}{n} \geq \frac{n! \cdot \varepsilon^n}{n^n} \geq \frac{\varepsilon^n}{e^n}. \quad (2.7.2)$$

Let $\gamma = (\gamma_0, \gamma_1, \dots, \gamma_k)$ be a sequence in Γ . Say that γ is *present* if there exists an injection $\sigma : [k] \rightarrow [2n-2]$ such that for every $i \in [k]$, we have $e_{\sigma(i)} = (\gamma_{i-1}, \gamma_i)$ or (γ_i, γ_{i-1}) . Assume that $k \leq 2n-2$, and note that γ cannot be present otherwise. There are at most $(2n-2)^k$ choices of σ , and given σ , for each i , the probability that $e_{\sigma(i)} = (\gamma_{i-1}, \gamma_i)$ is the same as the probability that $e_{\sigma(i)} = (\gamma_i, \gamma_{i-1})$, both given by $\frac{1}{n}p(\gamma_{i-1}, \gamma_i) = \frac{1}{n}p(\gamma_i, \gamma_{i-1})$. So by a union bound,

$$\mathbb{P}(\gamma \text{ is present}) \leq (2n-2)^k \prod_{i=1}^k \frac{2}{n} p(\gamma_{i-1}, \gamma_i) \leq 4^k p(\gamma) \leq 4^{2n} p(\gamma). \quad (2.7.3)$$

On the other hand, when H_1 is connected and H_2 is connected, then *some* $\gamma \in \Gamma$ must be present in H because every multigraph that contains two edge-disjoint spanning trees must also contain a

spanning subgraph that is connected and Eulerian [Cat92]⁴. Thanks to (2.7.2), this occurs with probability at least ε^{2n}/e^{2n} . So by a union bound,

$$\frac{\varepsilon^{2n}}{e^{2n}} \leq \sum_{\gamma \in \Gamma} \mathbb{P}(\gamma \text{ is present}) \leq 4^{2n} \sum_{\gamma \in \Gamma} p(\gamma). \quad (2.7.4)$$

The conclusion follows by rearranging. \square

2.8 Proof of Theorem 2.1

Let $G = (V, E)$ be an infinite, connected, locally finite graph such that for some constant $\varepsilon > 0$, for every vertex $v \in V$, the simple random walk $(X_t)_{t=0}^\infty$ on G started at v satisfies

$$d_v \mathbb{P}_v(\forall t \geq 1 : X_t \neq v) \geq \varepsilon.$$

Let $G' = (V', E')$ be the graph⁵ obtained from G by replacing each edge by a path of length 2. View V as a subset of V' , and let $m : E \rightarrow V'$ map each edge to its midpoint. Let \mathbb{P}'_u be the law of simple random walk in G' started from a given vertex u , and let $\tau := \sup\{t \geq 0 : X_t = X_0\}$. We claim that for all $z \in V'$,

$$d_z \mathbb{P}'_z(\tau = 0) \geq \varepsilon_1 := \frac{2\varepsilon}{4 + \varepsilon}. \quad (2.8.1)$$

This is trivial when $z \in V$, even with $\varepsilon_1 = \varepsilon/2$, because simple random walk on G' induces lazy simple random walk on G . Otherwise, when $z = m(\{u, v\})$ for some $\{u, v\} \in E$, this follows from the corresponding bounds for u and v by rearranging the following elementary calculation, where $\ell_x := |\{t \geq 0 : X_t = x\}|$:

$$\begin{aligned} \frac{1}{\mathbb{P}'_z(\tau = 0)} &= \sum_{n \geq 0} \mathbb{P}'_z(\tau > 0)^n = \mathbb{E}'_z[\ell_z] = \sum_{t \geq 0} \mathbb{P}'_z(X_t = z) \\ &= 1 + \sum_{t \geq 1} \left[\mathbb{P}'_z(X_{t-1} = u) \cdot \frac{1}{d_u} + \mathbb{P}'_z(X_{t-1} = v) \cdot \frac{1}{d_v} \right] \\ &= 1 + \frac{\mathbb{E}'_z[\ell_u]}{d_u} + \frac{\mathbb{E}'_z[\ell_v]}{d_v} \\ &\leq 1 + \frac{\mathbb{E}'_u[\ell_u]}{d_u} + \frac{\mathbb{E}'_v[\ell_v]}{d_v} = 1 + \frac{1}{d_u \mathbb{P}'_u(\tau = 0)} + \frac{1}{d_v \mathbb{P}'_v(\tau = 0)}. \end{aligned}$$

⁴This general fact can be proved directly as follows: Let E be the set of edges in the multigraph, and let T_1 and T_2 be the two trees. Let o_1, \dots, o_k be the vertices that have odd degree in T_1 . Since the sum of the degrees of all of the vertices in a given graph is always even, we can write $k = 2l$ for some non-negative integer l . For each $i \in [l]$, pick a path P_i in T_2 from o_{2i-1} to o_{2i} . Then, viewing T_1 and each P_i as elements of the $\mathbb{Z}/2\mathbb{Z}$ -vector space $\{0, 1\}^E$, the required subgraph is given by $T_1 + P_1 + \dots + P_l$.

⁵This construction is a technicality that is only necessary if G has unbounded vertex degrees.

Let $C := \{X_t : 0 \leq t \leq \tau\}$ and $\partial := \{e \cap C : e \in \partial_\infty C\}$. Fix $o \in V$, and pick a neighbour $o' \in m(E)$ of o in G' . Fix a finite minimal cutset Π from o to ∞ in G , and set $n := |\Pi|$. We will show that for some finite constant $K = K(\varepsilon) \in (0, \infty)$ depending only on ε ,

$$\mathbb{P}'_{o'}(\partial = m(\Pi)) \geq 1/K^n. \quad (2.8.2)$$

This implies that $\kappa(G) < \infty$ because for all $o \in V$ and $n \geq 1$,

$$1 \geq \sum_{\Pi \in Q_n(o)} \mathbb{P}'_{o'}(\partial = m(\Pi)) \stackrel{(2.8.2)}{\geq} |Q_n(o)|/K^n.$$

Let A be the connected component of o in $(V, E \setminus \Pi)$, let $U := m(\Pi) \cup \{o'\}$, and let $I := A \cup m(\{e \in E : e \subset A\})$. For all $u, v \in U \cup I$, let

$$p(u, v) := \mathbb{P}'_u(\exists t \geq 1 : X_1, \dots, X_{t-1} \in I \setminus \{u\} \text{ and } X_t = v).$$

Extend this to sets of vertices by $p(L, R) := \sum_{u \in L; v \in R} p(u, v)$, and similarly, $p(u, L) := p(\{u\}, L)$ and $p(L, u) := p(L, \{u\})$. We would like to apply Lemma 2.7.1 to the matrix $P := (p(u, v))_{u, v \in U}$. By time-reversing trajectories, we have $p(u, v) = p(v, u)$ whenever $d_u = d_v$, which is for example the case when $u, v \in U$. So P is symmetric, and clearly the entries of P are non-negative and sum to at most 1 along each row. We claim that for every non-trivial partition $U = L \sqcup R$,

$$p(L, R) \geq \varepsilon_2 := \varepsilon_1^2/64. \quad (2.8.3)$$

Indeed, for each $x \in U \cup I$, consider the function (the unit voltage)

$$F(x) := \mathbb{P}'_x(\exists t \geq 0 : X_0, \dots, X_{t-1} \notin L \text{ and } X_t \in R).$$

Given $u \in L$, there exists⁶ $x \in A$ such that $\{u, x\} \in E'$, and if $F(x) \geq 1/2$, then we are done because

$$p(u, R) \geq \mathbb{P}'_u(X_1 = x) \cdot F(x) \geq 1/2 \cdot 1/2 \geq \varepsilon_2.$$

In particular, we may assume that there exists $x \in A$ with $F(x) < 1/2$. By a similar argument, we may assume that there exists $y \in A$ with $F(y) > 1/2$. Since A is connected in G , we can therefore find $\{x', y'\} \in E$ satisfying $F(x') \leq 1/2 \leq F(y')$. Let $z := m(\{x', y'\})$, which has degree 2. Note that

$$F(z) \geq \mathbb{P}'_z(X_1 = y') \cdot F(y') \geq 1/2 \cdot 1/2,$$

⁶If $u = o'$, take $x := o$. If $u = m(e)$ where $e \in \Pi$, take x where $\{x\} = e \cap A$, which exists by Lemma 2.4.2.

and by a union bound,

$$F(z) \leq \sum_{n=0}^{\infty} \mathbb{P}'_z(\tau > 0)^n p(z, R) = \frac{p(z, R)}{\mathbb{P}'_z(\tau = 0)} \stackrel{(2.8.1)}{\leq} \frac{p(z, R)}{\varepsilon_1/2}.$$

So by rearranging, $p(z, R) \geq \varepsilon_1/8$. By a similar argument (i.e. by replacing F by $1 - F$, which switches the roles of L and R , and by recalling that $p(L, z) = p(z, L)$), we deduce that $p(L, z) \geq \varepsilon_1/8$. Now (2.8.3) follows because $p(L, R) \geq p(L, z)p(z, R)$.

Therefore by Lemma 2.7.1, the event \mathcal{E} that the random walk visits every vertex in U then returns to o' before exiting $U \cup I$ satisfies

$$\mathbb{P}'_{o'}(\mathcal{E}) \geq \varepsilon_3^{|U|} \geq \varepsilon_3^{n+1} \quad (2.8.4)$$

for some constant $\varepsilon_3 > 0$ depending only on ε_2 . So by Lemma 2.4.2 and the strong Markov property,

$$\mathbb{P}'_{o'}(\partial = m(\Pi)) \geq \mathbb{P}'_{o'}(\mathcal{E}) \cdot \mathbb{P}'_{o'}(\tau = 0) \stackrel{(2.8.1), (2.8.4)}{\geq} \varepsilon_3^{n+1} \cdot \varepsilon_1/2. \quad (2.8.5)$$

By expanding the definitions of $\varepsilon_1, \varepsilon_2, \varepsilon_3$ we deduce that (2.8.2) holds with $K := 2^{20}/\varepsilon^5$.

2.9 Alternative proof of Theorem 2.1 using the Gaussian free field

Here we sketch an alternative, slightly less elementary proof of Theorem 2.1 along the lines of the proof of Theorem 2.1. Let G be an infinite, connected, locally finite graph that is uniformly transient. Consider the graph $\tilde{G} = (\tilde{V}, \tilde{E})$ obtained by replacing each edge by a path of length 3. Similarly to the proof in Section 2.8, one can prove that \tilde{G} is also uniformly transient. Let $\varphi \in \mathbb{R}^{\tilde{V}}$ with law \mathbb{P} be the (centered) *Gaussian free field (GFF)* on \tilde{G} – see e.g. [BP24] for the required background and definitions. Uniform transience implies that there exists $\varepsilon > 0$ such that $\text{Var}(\varphi(x)) \leq 1/\varepsilon$ for every $x \in \tilde{V}$.

Fix $o \in V$ and let \tilde{C} be the cluster of o in the percolation model induced by the excursion set $\{\varphi \geq 0\} := \{x \in \tilde{V} : \varphi(x) \geq 0\}$. Given every edge e of G , we associate the corresponding mid-edge \tilde{e} in \tilde{G} , with both endpoints of degree 2. For a subset Π of edges in G , we denote by $\tilde{\Pi}$ the associated set of mid-edges in \tilde{G} . We claim that there exists $c = c(\varepsilon) > 0$, depending only on ε , such that for every $\Pi \in Q_n(o)$,

$$\mathbb{P}(\partial_{\infty} \tilde{C} = \tilde{\Pi}) \geq c^n. \quad (2.9.1)$$

Similarly to the previous sections, Theorem 2.1 follows readily from (2.9.1).

We now proceed to prove (2.9.1). Enumerate $\tilde{\Pi}$ by $\tilde{e}_i = \{x_i, y_i\}$, $1 \leq i \leq n$, where x_i and y_i are the inner and outer endpoints, respectively. We first observe that, for some constant $c_1 = c_1(\varepsilon) > 0$,

$$\mathbb{P}(\varphi(y_i) \in [-2, -1] \text{ and } \varphi(x_i) \in [1, 2] \text{ } \forall 0 \leq i \leq n) \geq c_1^n. \quad (2.9.2)$$

Indeed, this follows by successively demanding the desired event at each vertex. Here we use the Markov property of the GFF (see [BP24]) and the fact that the conditional variance of the next vertex given the previous ones is between $1/2$ (since they have degree 2) and $1/\varepsilon$, while the conditional mean remains bounded between -2 and 2 .

Let \mathcal{F} be the event in (2.9.2) and A be the component of o in $(\tilde{V}, \tilde{E} \setminus \tilde{\Pi})$. Notice that

$$\mathbb{P}(\partial_\infty \tilde{C} = \tilde{\Pi}) \geq \mathbb{P}(\mathcal{F}) \mathbb{P}\left(\bigcap_{i=1}^n \{0 \xrightarrow{\{\varphi \geq 0\} \cap A} x_i\} \mid \mathcal{F}\right) \geq c_1^n \mathbb{P}\left(\bigcap_{i=1}^n \{0 \xrightarrow{\{\varphi \geq 0\} \cap A} x_i\} \mid \mathcal{F}\right).$$

By the Markov property, conditionally on \mathcal{F} , the process $\{\varphi \geq 0\} \cap A$ stochastically dominates $\{\varphi_A \geq -1\}$, where φ_A is the centered GFF on A (i.e. associated to the random walk on A killed when reaching $\partial A = \{x_1, \dots, x_n\}$). Therefore, it is enough to prove that, for some constant $c_2 = c_2(\varepsilon) > 0$,

$$\mathbb{P}\left(\bigcap_{i=1}^n \{o \xrightarrow{\{\varphi_A \geq -1\}} x_i\}\right) \geq c_2^n. \quad (2.9.3)$$

Indeed, since the GFF is positively associated (see [BP24]), the desired inequality (2.9.3) follows readily from Proposition 2.5.1 and the following inequality

$$\mathbb{P}(u \xrightarrow{\{\varphi_A \geq -1\}} \partial A) \geq \mathbb{E}(\text{sgn}(\varphi_A(u) + 1)) \geq c_3, \quad (2.9.4)$$

for some constant $c_3 = c_3(\varepsilon) > 0$. The latter follows easily from the Markov property of the GFF. Indeed, let \mathcal{S} be the union of all clusters of $\{\varphi_A \geq -1\}$ intersecting ∂A and note that its closure $\overline{\mathcal{S}}$ (i.e. the union of \mathcal{S} with its neighbours) is a stopping set. Clearly, one has $\text{sgn}(\varphi_A(u) + 1) = 1$ almost surely on the event $\mathcal{G} := \{u \xrightarrow{\{\varphi_A \geq -1\}} \partial A\} = \{u \in \mathcal{S}\}$. On the complementary event \mathcal{G}^c and conditionally on the field on $\overline{\mathcal{S}}$, the Markov property implies that we have a GFF on $A \setminus \mathcal{S}$ with boundary conditions < -1 . In particular, $\text{sgn}(\varphi_A(u) + 1)$ has a negative conditional expectation on \mathcal{G}^c . These observations readily imply the first inequality of (2.9.4). The second inequality follows from the fact that the variance of $\varphi_A(u)$ is at most $1/\varepsilon$.

THE CRITICAL PERCOLATION PROBABILITY IS LOCAL

Joint work with Tom Hutchcroft

Abstract

We prove Schramm’s locality conjecture for Bernoulli bond percolation on transitive graphs: If $(G_n)_{n \geq 1}$ is a sequence of infinite vertex-transitive graphs converging locally to a vertex-transitive graph G and $p_c(G_n) \neq 1$ for every $n \geq 1$ then $\lim_{n \rightarrow \infty} p_c(G_n) = p_c(G)$. Equivalently, the critical probability p_c defines a continuous function on the space \mathcal{G}^* of infinite vertex-transitive graphs that are not one-dimensional. As a corollary of the proof, we obtain a new proof that $p_c(G) < 1$ for every infinite vertex-transitive graph that is not one-dimensional.

3.1 Introduction

In **Bernoulli bond percolation**, the edges of a connected, locally finite graph G are chosen to be either retained (**open**) or deleted (**closed**) independently at random, with probability $p \in [0, 1]$ of retention. The law of the resulting random subgraph is denoted $\mathbb{P}_p = \mathbb{P}_p^G$. Percolation theory is concerned primarily with the geometry of the connected components of this random subgraph, which are known as **clusters**. Much of the interest in the model arises from the fact that it undergoes a *phase transition*: If we define the **critical probability**

$$p_c(G) = \inf \{p \in [0, 1] : \text{there exists an infinite cluster } \mathbb{P}_p\text{-almost surely}\}$$

then “most” interesting graphs have $p_c(G)$ strictly between 0 and 1, so that there is a non-trivial phase where infinite clusters do not exist followed by non-trivial phase where they do exist. We will be primarily interested in percolation on (**vertex-)**transitive graphs, i.e., graphs for which any vertex can be mapped to any other vertex by an automorphism of the graph. We allow our graphs to contain loops and multiple edges, and make the implicit assumption throughout the paper that all transitive graphs are connected and locally finite (i.e. have finite vertex degrees).

Many interesting features of percolation on an infinite transitive graph at and near the critical point are expected to be *universal*, meaning that they depend only on the graph’s *large-scale* geometry and not its microscopic structure. For example, the *critical exponents* governing the power-law behaviour of various interesting quantities at and near criticality are believed to depend only on

the volume-growth dimension of the graph, and should therefore take the same values on e.g. the square and triangular lattice. In contrast, Schramm conjectured around 2008 [BNP11a, Conjecture 1.2] that the *value* of the critical probability $p_c(G)$ should be entirely determined by the *local* (microscopic) geometry of the graph, *subject to the global constraint that* $p_c(G) < 1$. More precisely, he conjectured that if G_n is a sequence of infinite transitive graphs converging to an infinite transitive graph G in the local topology (defined below) and $\limsup_{n \rightarrow \infty} p_c(G_n) < 1$ then $p_c(G_n) \rightarrow p_c(G)$ as $n \rightarrow \infty$. The assumption that $\limsup_{n \rightarrow \infty} p_c(G) < 1$ is needed to rule out degenerate one-dimensional examples such as the cylinder $\mathbb{Z} \times (\mathbb{Z}/n\mathbb{Z})$ (which converges to the square grid \mathbb{Z}^2 but which has $p_c(\mathbb{Z} \times (\mathbb{Z}/n\mathbb{Z})) = 1 \nrightarrow p_c(\mathbb{Z}^2) = 1/2$), and is now known to be equivalent to the graphs G_n having superlinear volume growth for all sufficiently large n [DGRSY20; HT21a]. The fact that p_c is *lower* semi-continuous follows straightforwardly from standard facts about percolation on transitive graphs as observed in [Pet, §14.2] and [DT16a, p.4] and does not require the assumption that the graphs are not one-dimensional; the difficult part of the conjecture is to prove *upper* semi-continuity.

The locality conjecture has inspired a great deal of subsequent work, including both partial progress on the original conjecture [BNP11a; MT17; Hut20a; HH21a; CMT23a], which we review in detail below, and analogous results in other settings including self-avoiding walk [GL17; GL18], the random cluster model [DT19], finite random graphs [BNP11a; KLS20; Sar21a; Hof21; RS22a; ABS22; BZ23; ABS23], and geometric random graphs [HM22; LLMS23].

In this paper we give a complete proof of Schramm's locality conjecture.

Theorem 3.1.1. *Let \mathcal{G}^* be the set of all infinite transitive graphs that are not one-dimensional, endowed with the local topology. Then the function $p_c : \mathcal{G}^* \rightarrow (0, 1)$ is continuous.*

Here, the **local topology** (a.k.a. the **Benjamini-Schramm topology**) on the space of transitive graphs, denoted by \mathcal{G} , is defined so that $(G_n)_{n=1}^\infty$ converges to G if and only if, for each $r \geq 1$, the balls of radius r in G_n and G are isomorphic as rooted graphs for all sufficiently large n . We say that an infinite transitive graph is **one-dimensional** if it has linear volume growth (i.e., if its balls B_n satisfy $|B_n| = O(n)$ as $n \rightarrow \infty$); it follows from (a simple special case of) the structure theory of transitive graphs of polynomial growth that an infinite transitive graph is one-dimensional if and only if it is rough-isometric to \mathbb{Z} [TY16], while the results of [DGRSY20] imply that an infinite transitive graph has $p_c < 1$ if and only if it is not one-dimensional. (In fact the proof of Theorem 8.1.1 also yields a new proof of this theorem as we discuss in detail in Section 3.3.)

Remark 3.1.1. In our forthcoming paper [EH23+a], we prove related results on the locality of the density of the infinite cluster, implying in particular that the percolation probability $\theta(p, G) = \mathbb{P}_p^G(o \leftrightarrow \infty)$ is a continuous function of (G, p) in the supercritical set $\{(G, p) : G \in \mathcal{G}^*, p > p_c(G)\}$. (Theorem 8.1.1 implies that this set is open.) An alternative proof of this result using the methods developed in the present paper is sketched in Section 3.7.

Previous work

In this section we overview previous work on locality, the $p_c < 1$ problem, and the structure theory of transitive graphs of polynomial growth. Our proof will employ many ideas and methods from these earlier papers, including most notably the works [Hut20a; Hut20e; CMT22] and the structure theory developed in [BGT12a; TT21a].

Euclidean lattices. Well before the formulation of Schramm’s conjecture, the first significant work on locality was carried out in the seminal work of Grimmett and Marstrand [GM90a], who proved that the critical probability for percolation on a “slab” $\mathbb{Z}^{d-k} \times \{0, \dots, n\}^k$ converges to $p_c(\mathbb{Z}^d)$ as $n \rightarrow \infty$ provided that $d - k \geq 2$. (In this context \mathbb{Z}^d and $\mathbb{Z}^{d-k} \times \{0, \dots, n\}^k$ refer to the Cayley graphs of these groups *with their standard generating sets*.) This theorem and the methods developed to prove it are of central importance to the study of supercritical percolation in three and more dimensions. (Moreover, one of the motivations for the work of Grimmett and Marstrand was to get closer to proving that the percolation phase transition on \mathbb{Z}^d is continuous, and indeed similar methods were used in [BGN91] to prove continuity for half-spaces.) Although it is not strictly an instance of Schramm’s conjecture since slabs are not transitive, the Grimmett–Marstrand theorem trivially implies that the analogous statement holds for “slabs with periodic boundary conditions” (i.e., $\mathbb{Z}^{d-k} \times (\mathbb{Z}/n\mathbb{Z})^k$ with its standard generating set), which are transitive. The proof of the Grimmett–Marstrand theorem relies heavily on renormalization techniques exploiting the full symmetries of \mathbb{Z}^d and scale-invariance of Euclidean space \mathbb{R}^d , and does not readily generalize to other transitive graphs. Let us also mention that a quantitative version of the Grimmett–Marstrand theorem was proven in the more recent work of Duminil-Copin, Kozma, and Tassion [DKT21] which was very influential in both our work and [CMT22].

Remark 3.1.2. A further classical Euclidean result in the spirit of the locality conjecture was established by Kesten [Kes90], who proved that $p_c(\mathbb{Z}^d) \sim 1/(2d - 1)$ as $d \rightarrow \infty$. See [ABS04a] for a simple proof and [Sla06] for more refined results. While not strictly an instance of the locality conjecture since \mathbb{Z}^d does not converge in the local topology, the intuitive reason for this result to hold is that \mathbb{Z}^d is “locally tree-like” when $d \rightarrow \infty$ in the sense that small cycles have a negligible effect on the behaviour of the percolation model, so that one can define a kind of “local limit” of \mathbb{Z}^d .

as $d \rightarrow \infty$ (in a different technical sense to the one we consider here) in terms of Aldous's *Poisson weighted infinite tree* (PWIT) [AS04].

Progress on locality. Previous works on the locality conjecture can be divided into two strains, with completely different set of methods and domains of application associated to them: The first strain concerns graphs that satisfy various strong, “infinite-dimensional” expansion conditions, while the second concerns “finite-dimensional” graphs (i.e., graphs of polynomial volume growth) where one can hope to develop appropriate generalizations of Grimmett–Marstrand theory. The first strain splits further into two cases according to whether or not the graphs in question are *unimodular*, a technical condition¹ that holds for most familiar examples of transitive graphs including every Cayley graph and every amenable transitive graph [SW88]. Although nonunimodular graphs are often considered to be “pathological” compared to their unimodular cousins, it turns out that *nonunimodularity is actually a very helpful assumption*: in [Hut20e] the second author carried out a very detailed analysis of critical percolation on nonunimodular transitive graphs, which he then used to prove the nonunimodular case of locality in [Hut20a]. Moreover, it was proven in [Hut20a, Corollary 5.5] that the set of nonunimodular transitive graphs is both closed and open in \mathcal{G} , so that to prove Theorem 8.1.1 it now suffices to consider the case that all graphs in question are unimodular.

Let us now discuss previous results for unimodular graphs in the “infinite-dimensional” setting. The first result in this direction was due to Benjamini, Nachmias, and Peres [BNP11a], who proved the conjecture for nonamenable graph sequences satisfying a certain high girth condition (e.g., uniformly nonamenable graph sequences of divergent girth; unimodularity is not required). More recently, the second author [Hut20a] proved the conjecture for graph sequences of uniform exponential growth (meaning that the balls of radius r in the graphs G_n all have volume lower-bounded by e^{cr} for some constant c independent of n and r), and Hermon and the second author [HH21a] proved the conjecture for sequences of graphs satisfying a certain uniform stretched-exponential heat kernel upper bound, a class that includes certain examples of intermediate volume growth (i.e., volume growth that is superpolynomial but subexponential; note however that the spectral condition of [HH21a] is not implied by any growth condition). The works [Hut20a; Hut20e; HH21a] all establish locality for the families of graphs they consider by proving quantitative tail estimates on critical percolation clusters that hold uniformly for all graphs in the family, yielding

¹Here is the definition: A transitive graph $G = (V, E)$ is unimodular if it satisfies the **mass-transport principle**, meaning that $\sum_x F(o, x) = \sum_x F(x, o)$ for every $F : V^2 \rightarrow [0, \infty]$ that is diagonally invariant in the sense that $F(x, y) = F(\gamma x, \gamma y)$ for every automorphism γ of G . We will not directly engage with unimodularity in this paper, but it will appear as a hypothesis in many of our intermediate results since it is needed to apply the two-ghost inequality of [Hut20a].

much more than just locality. (In particular, they also imply that the graphs in question have continuous percolation phase transitions.)

Remark 3.1.3. Although the techniques developed in [Hut20a; HH21a] have yet to be made to work for *nearest-neighbour* percolation models in finite dimension, versions of these arguments have been used in [Hut21a] to analyze certain *long-range* percolation models in finite-dimensional spaces. The results of [Hut21a] can be used to prove versions of the locality conjecture for certain large families of long-range percolation models on unimodular transitive graphs (under the assumption that the long-range edge kernel has a sufficiently heavy tail uniformly throughout the sequence). Further results on locality for long-range percolation can be found in [MS96; Ber02].

Polynomial growth and structure theory. We now discuss the second strain of results, concerning graphs of polynomial volume growth. Let us first briefly review the *structure theory* of transitive graphs of polynomial growth, which plays an important role in these developments. Recall that \mathcal{G} is the space of all infinite transitive graphs and that $G \in \mathcal{G}$ is said to have **polynomial growth** if for some positive reals C and d , the number of vertices contained in a ball of radius n , denoted $\text{Gr}(n) := |B_n(o)|$, satisfies $\text{Gr}(n) \leq Cn^d$ for all $n \geq 1$. The geometry of such graphs is highly constrained: it is a consequence of Gromov's theorem [Gro81b] and Trofimov's theorem [Tro84b] that every $G \in \mathcal{G}$ with polynomial growth is necessarily quasi-isometric to the Cayley graph of a nilpotent group. In particular, for every such graph G , there is a positive real C and a unique positive integer d such that $C^{-1}n^d \leq \text{Gr}(n) \leq Cn^d$ for all $n \geq 1$. The integer d is called the (volume growth) **dimension** of G ; it coincides with the *isoperimetric dimension* and *spectral dimension* of G by a theorem of Coulhon and Saloff-Coste [CS93]. These results are often used in the study of probability on transitive graphs as part of a “structure vs. expansion dichotomy”, wherein each graph either satisfies a high-dimensional isoperimetric inequality (which is often a helpful assumption) or else is quasi-isometric to a nilpotent group of bounded step and rank (which is useful because these graphs are highly explicit and well-behaved); a detailed overview of the structure theory of transitive graphs of polynomial growth and its applications to probability is given in the introduction to [EH23d].

More recently, *finitary* versions of these results have been established, first for groups in the landmark work of Breuillard, Green, and Tao [BGT12a], then for transitive graphs by Tessera and Tointon [TT21a]. These results imply, for instance, that for each constant $K < \infty$ there exists $N < \infty$ such that if we observe that $\text{Gr}(3n) \leq K \text{Gr}(n)$ for some $n \geq N$, then G is $(1, Cn)$ -quasi isometric to the Cayley graph of a virtually nilpotent group where the constant C along with the rank, step, and index of the nilpotent subgroup are all bounded above by some function of K (see Section 3.5 for

further details). This finitary structure theory is extremely useful in applications to problems such as locality in which one wishes to argue in a way that is uniform over some family of graphs. For example, it follows from [TT21a, Corollary 1.5] that for each $d \geq 1$ the set of transitive graphs of polynomial growth with dimension at most d is an open subset of \mathcal{G} , and moreover that if $G_n \rightarrow G$ with G of polynomial growth of dimension d then there exists $n_0 < \infty$ and a constant C such that the ball of radius r in G_n has volume at most Cr^d for every $n \geq n_0$, with constants independent of n . As such, to prove locality in the case that the limit has polynomial growth, it suffices to consider the case that all graphs in the sequence satisfy a *uniform* polynomial upper bound on their growth as well as various other forms of strong uniform control on their geometry.

Besides the original work of Grimmett and Marstrand, the first result on locality for graphs of polynomial growth was due to Martineau and Tassion [MT17], who proved that locality holds for Cayley graphs of abelian groups. Their proof employs a variation on the Grimmett–Marstrand argument, overcoming significant technical difficulties arising due to the loss of rotational and reflection symmetry. This result was greatly extended in the recent work of Contreras, Martineau, and Tassion, who developed a version of Grimmett–Marstrand theory for transitive graphs of polynomial growth in [CMT22] and used this theory together with the finitary structure theory discussed above to deduce the polynomial growth case of the locality conjecture in [CMT23a]. As with the aforementioned works in the infinite-dimensional setting, the works [CMT22; CMT23a] establish not just locality but also many further strong quantitative results about percolation on the classes of graphs they consider. However, while [Hut20a; HH21a] established quantitative estimates on finite clusters in *critical* percolation, [CMT22; CMT23a] instead establish strong results about the geometry of the infinite cluster in *supercritical* percolation. This reflects a fundamental distinction between the two approaches, with the analysis of the polynomial growth case involving estimates on uniqueness of annuli crossings etc. that are simply not true for “big” graphs like the 3-regular tree.

What challenges remain? Given the previous results discussed above, it appears that there are two main cases of the locality conjecture left to consider: arbitrary sequences of superpolynomial growth transitive graphs converging to a superpolynomial growth graph, and the “diagonal” case in which a sequence of polynomial-growth graphs converges to a graph of superpolynomial growth. Moreover, the second case might be split further according to whether the graphs in the sequence have bounded or divergent dimension. (As we will soon explain, our proof will in fact work through a different and less obvious kind of case analysis.) In the first case, a key difficulty is that the geometry of transitive graphs of superpolynomial growth can be highly arbitrary, with

the space of all such graphs being an ineffably complex object in some senses: the challenge is precisely that we need an argument that is robust enough to work for all possible transitive graphs, for which there is nothing like a general classification. Moreover, the best known uniform lower bound on the growth of groups of superpolynomial growth, due to Shalom and Tao [ST10b], is of the form $n^{(\log \log n)^c}$ for a small constant $c > 0$. This growth lower bound (which has not yet been proven for transitive graphs that are not Cayley) is *vastly* weaker² than the assumptions used in the works [BNP11a; Hut20a; HH21a] discussed above. In the diagonal case, one must contend with these same difficulties again together with the total incompatibility of the methods that have thus far been used to handle the high-growth and polynomial growth cases. As such, while the main difficulties in the locality conjecture arise from “unknown enemies” hiding deep within the unknowable expanse of the space of all transitive graphs, there are also explicit examples that seem difficult to handle within existing frameworks. (One such example is the standard Cayley graph of the free step- s nilpotent group on two generators, which converges to a 4-regular tree as $s \rightarrow \infty$ but has finite dimension for each finite s .)

Parallels with $p_c < 1$. Before we begin to describe our proof of the locality conjecture, let us first discuss how its history closely parallels the (older) history of the $p_c < 1$ problem. It follows from the classical work of Peierls [Pei36b] that $p_c(\mathbb{Z}^d) < 1$ for every $d \geq 2$. In their highly influential work [BS96b], Benjamini and Schramm conjectured that $p_c < 1$ for every transitive graph that is not one-dimensional. Benjamini and Schramm also proved in the same paper that $p_c < 1$ for every (not necessarily transitive) nonamenable graph, while earlier results of Lyons [Lyo95] implied that exponential growth suffices in the transitive case. On the other hand, it is a simple consequence of the structure theory that every transitive graph of polynomial volume growth that is not one-dimensional contains a subgraph that is quasi-isometric to \mathbb{Z}^2 , which easily implies that every such graph has $p_c < 1$ (see e.g. [HT21a, Section 3.4] for details). As such, for many years the problem remained open only for groups of intermediate growth. Even in this case the problem was solved for most “known” examples of graphs of intermediate growth, such as the Grigorchuk group [MP01; RY17], with the main remaining difficulty coming from “unknown enemies” as discussed above. See the introduction of [DGRSY20] for a detailed account of this partial progress including several further references.

The $p_c < 1$ problem was eventually solved in full generality by Duminil-Copin, Goswami, Raoufi, Severo, and Yadin [DGRSY20]. More precisely, they established that $p_c < 1$ for any (not necessarily

²While a well-known conjecture of Grigorchuk [Gri14] states that every superpolynomial growth transitive graph has growth at least $\exp[cn^{1/2}]$, this conjecture is completely open, somewhat controversial, and in any case would still require a significant advance on the methods of [Hut20a; HH21a] to be applicable towards the locality conjecture.

transitive) bounded degree graph satisfying a $(4 + \varepsilon)$ -dimensional isoperimetric inequality, with the structure theory and classical results above handling all remaining transitive graphs. Their proof uses a comparison between percolation and the Gaussian free field which works only for values of p very close to 1, making their methods unsuitable for the locality problem. A finitary version of the results of [DGRSY20] was developed by the second author and Tointon in [HT21a], who proved in particular that sequences of *finite* transitive graphs have a non-trivial phase in which a *giant* cluster exists provided that they are “not one dimensional” in an appropriate quantitative sense. This finitary approach also allowed them to prove a *uniform* version of the main result of [DGRSY20], stating that for each $d \geq 1$ there exists $\varepsilon > 0$ such that every infinite transitive graph that has degree at most d and is not one-dimensional has $p_c < 1 - \varepsilon$. (For Cayley graphs it is now known that ε can be taken independently of the degree [PS23b]; see Section 3.7 for some related conjectures.)

Can we do something similar? Continuing to follow the path set by this previous work on the $p_c < 1$ problem, one might hope to prove locality via a similar dichotomy, finding some method that handles all graphs that are “high-dimensional” in some sense, then using the structure theory to separately analyze the remaining “low-dimensional” examples. One technical problem with this approach, which was already a major hurdle in [HT21a], is that one is forced to consider sequences of graphs that may look high-dimensional up to some divergently large scale then switch to looking low-dimensional, meaning that one must find a way to “patch together” the outputs of the two different case analyses at the crossover scale. A more fundamental problem, however, is that to date there have simply been no viable approaches to prove locality under the assumption that the graphs are high dimensional.

As we will see, our proof will instead follow a more subtle approach in which we first dichotomize into two much less obvious cases according to whether or not the graph has *quasi-polynomial growth* on the relevant scale. (Here, a function is said to have *quasi-polynomial growth* if it is bounded by a function of the form $\exp[(\log n)^{O(1)}]$.) In the low growth case, we then employ a second, subordinate dichotomization according to whether the *rate of growth* on the relevant scale is low-dimensional or high-dimensional; it is in this second dichotomy that we can make use of the structure theory. A key technical difficulty when arguing this way is that (as far as we know with the current structure theory), the same graph might oscillate between the quasi-polynomial and super-quasi-polynomial regimes infinitely many times as we go up the scales, so that any “case analysis” we do via this dichotomy must be able to handle this oscillation. Moreover, the delicate nature of our proof leads us to engage with the structure theory literature in a deeper way than had

previously been necessary in applications to probability. Indeed, our proof relies in part on our structure-theoretic companion paper [EH23d] in which we prove a “uniform finite presentation” theorem for groups of polynomial volume growth which we use to make some of the arguments from [CMT22] finitary. Besides this, we must also contend with the fact that the assumption of quasi-polynomial growth is highly non-standard, so that we must spend a significant amount of the paper studying the deterministic geometry of transitive graphs of quasi-polynomial growth (at some scale) with a view to eventually generalizing the methods of [CMT22] from polynomial to quasi-polynomial growth.

Interestingly, our proof of locality yields as an immediate corollary a new proof that $p_c < 1$ for transitive graphs that are not one-dimensional, recovering the main result of [DGRSY20]. This proof works directly with Bernoulli percolation and does not rely on the comparison to the GFF in any way. (On the other hand it is also *much* more complicated than the original proof!)

About the proof

In this section we give an overview of our proof. Let us first establish some relevant notation that will be used throughout the paper. Recall that \mathcal{G} denotes the space of all (vertex-)transitive graphs (which we always take to be connected and locally finite) and that $\mathcal{G}^* \subseteq \mathcal{G}$ is the space of infinite transitive graphs that are not one-dimensional. We also write $\mathcal{U} \subseteq \mathcal{G}$ for the space of all unimodular transitive graphs and write $\mathcal{U}^* = \mathcal{U} \cap \mathcal{G}^*$. Given $d \in \mathbb{N}$, we write \mathcal{G}_d , \mathcal{G}_d^* , and \mathcal{U}_d^* for the subsets of these spaces in which every graph has degree d . Given sets of vertices A and B in a graph, we write $\{A \leftrightarrow B\}$ for the event that there is a path from A to B in the given configuration. We also use the notation $A \leftrightarrow \infty$ to mean that there is an infinite cluster that intersects A . When $A = \{u\}$ and $B = \{v\}$ are singletons, we may simply write $u \leftrightarrow v$ and $u \leftrightarrow \infty$ instead. For each transitive graph G we will write o for an arbitrarily chosen root vertex of G which we will refer to as *the origin*. We write B_n for the graph-distance ball of radius n around o in G and write S_n for the set of vertices at distance exactly n from o .

Uniform estimates and finite-size criteria. Recall that percolation is said to undergo a *continuous phase transition* on an infinite graph G if the function $\theta_G : p \mapsto \mathbb{P}_p^G(o \leftrightarrow \infty)$ is continuous (it is a theorem of Schonmann [Sch99] that θ_G is always continuous at every point other than p_c). One of the best-known conjectures in the study of percolation on general infinite transitive graphs is that percolation should undergo a continuous phase transition on every $G \in \mathcal{G}^*$ [BS96b, Conjecture 4]; this is famously still open when G is the three-dimensional cubic lattice. Although it might not be obvious from their statements, the locality conjecture and the continuity conjecture are very closely related. To see why, let us give two equivalent formulations of Theorem 8.1.1, which will inform

the remainder of our analysis. (For notational convenience we define $\mathbb{P}_p := \mathbb{P}_1$ when $p > 1$.)

Reformulation of the locality conjecture 1. *For each $d \in \mathbb{N}$ and $\varepsilon, \delta > 0$, there exists $r \in \mathbb{N}$ such that*

$$\mathbb{P}_p^G(o \leftrightarrow S_r) \geq \delta \implies \mathbb{P}_{p+\varepsilon}^G(o \leftrightarrow \infty) > 0 \quad (3.1.1)$$

for every $G \in \mathcal{G}_d^*$ and $p \in [0, 1]$.

Reformulation of the locality conjecture 2. *For each $d \in \mathbb{N}$ and $\varepsilon > 0$, there exists a function $h_\varepsilon = h_{d,\varepsilon} : \mathbb{N} \rightarrow (0, \infty)$ with $h_\varepsilon(r) \rightarrow 0$ as $r \rightarrow \infty$ such that*

$$\mathbb{P}_{p_c-\varepsilon}^G(o \leftrightarrow S_r) \leq h_\varepsilon(r) \quad (3.1.2)$$

for every $G \in \mathcal{G}_d^*$ and $r \geq 1$.

These two reformulations are trivially equivalent to one another (the difference amounts to considering either $h_\varepsilon(n)$ or its inverse), but we find the two different viewpoints on the same statement to be illuminating: the first is formulated in terms of a “finite-size criterion for not being very subcritical” while the second is formulated in terms of “uniform estimates on subcritical percolation”. In either formulation, the $\varepsilon = 0$ analogue of the same statement would imply both the locality conjecture and a strengthened, “uniform in G ” version of the conjecture concerning the continuity of the phase transition; this is precisely what the arguments of [Hut20a; HH21a] establish for the classes of graphs they consider. (In fact this uniform version of continuity is implied by locality together with the non-uniform version of continuity by a simple compactness argument.) This suggests a close connection between the two problems, while the freedom to use “sprinkling” (i.e., to increase p by small amounts in an appropriate manner) may make locality significantly more tractable.

Let us now briefly explain why Reformulation 1 is equivalent to the locality conjecture (or more accurately to the upper semi-continuity of p_c , which is the difficult part of locality). In one direction, suppose that $G_n \rightarrow G$ and that $p > p_c(G)$. Since $p > p_c(G)$, the connection probability $\mathbb{P}_p^G(o \leftrightarrow S_r)$ does not decay as $r \rightarrow \infty$. Since for each fixed r we also have that $\mathbb{P}_p^{G_n}(o \leftrightarrow S_r) = \mathbb{P}_p^G(o \leftrightarrow S_r)$ for every sufficiently large n by the definition of local convergence, we may apply Reformulation 1 with $\varepsilon = \mathbb{P}_p^G(o \leftrightarrow \infty) > 0$ and δ an arbitrary positive number to deduce that $p_c(G_n) \leq p + \delta$ for all sufficiently large n . This implies the desired upper semi-continuity since $p > p_c(G)$ and $\delta > 0$ were arbitrary. In the other direction, suppose that Reformulation 1 is *false*, so that there exists $d \in \mathbb{N}$, $\varepsilon > 0$, and $\delta > 0$ such that for each $r \geq 1$ there exists $G_r \in \mathcal{G}_d^*$ and $p_r \in [0, 1]$ with $p_c(G_r) \geq p_r + \varepsilon$ and $\mathbb{P}_{p_r}^{G_r}(o \leftrightarrow S_r) \geq \delta$. Since \mathcal{G}_d^* is compact, there exists a subsequence along which G_r converges locally to some transitive graph $G \in \mathcal{G}_d^*$ and p_r converges to some $p \in [0, 1]$,

which must satisfy $\mathbb{P}_p^G(o \leftrightarrow r) \geq \delta$ for every $r \geq 1$ and hence that $p_c(G) \leq p$. Since the \liminf of $p_c(G_r)$ along this subsequence is at least $p + \varepsilon$ we deduce that p_c is *not* continuous on \mathcal{G}^* , completing the proof of the equivalence.

Remark 3.1.4. The lower semi-continuity of p_c follows by a very similar argument to the deduction of upper semi-continuity from Reformulation 1 that we just gave, but where the relevant finite-size criteria or uniform bounds follow easily from standard facts about the sharpness of the phase transition: The argument of [DT16a] is formulated in terms of finite-size criteria for subcriticality, while that of [Pet] is formulated in terms of uniform bounds on supercritical percolation.

Remark 3.1.5. Although the perspective on the locality problem we have just discussed is highly influential on our approach, we will in fact follow a slightly different approach in order to circumvent some technical problems related to estimating the “burn-in”, i.e., the amount of sprinkling needed to perform the base case of our multi-scale induction scheme. As such, we do not obtain an explicit function h_ε as in Reformulation 2 in this paper. The additional steps required to make our argument completely quantitative and obtain an explicit bound on the function h_ε will be carried out in a forthcoming companion paper [EH23+b].

A non-trivial reformulation. Let us now begin to go into more detail about our methods of proof. As discussed above, the results of [Hut20a; Hut20e] completely resolve the nonunimodular case of the conjecture, and since the space \mathcal{U}^* is both closed and open in \mathcal{G}^* we may restrict from now on to the case that all graphs are unimodular. In this case, the sharpness of the phase transition [Men86; AB87a; DT16a] together with the methods of [Hut20a] allow us to further reformulate Theorem 8.1.1 as follows.

Reformulation of the locality conjecture 3. *For all $d \in \mathbb{N}$ and all $\varepsilon, \delta > 0$, there exists $n \in \mathbb{N}$ such that*

$$\min_{u,v \in B_n} \mathbb{P}_p(u \leftrightarrow v) \geq \varepsilon \implies \lim_{m \rightarrow \infty} \frac{1}{m} \log \mathbb{P}_{p+\delta}(o \leftrightarrow S_m) = 0$$

for every $G \in \mathcal{U}_d^$ and $p \in [0, 1]$.*

The fact that we can replace the statement that $\theta(p + \delta) > 0$ appearing on the right hand side of Reformulation 1 with the statement that the radius has a subexponential tail appearing on the right hand side of Reformulation 3 follows directly from the sharpness of the phase transition [Men86; AB87a; DT16a], which implies that the radius has an exponential tail whenever $p < p_c$. (This does not require unimodularity.) The fact that we can replace the lower bound on the tail of the radius appearing on the left hand side of Reformulation 1 with the lower bound on point-to-point

connection probabilities appearing on the left hand side of Reformulation 3 is a much less obvious³ fact and is the main content of [Hut20a]. (Indeed, in the exponential growth setting studied in [Hut20a], a uniform upper bound on critical point-to-point connection probabilities was already established in [Hut16].) The argument needed to see that Reformulation 3 implies Theorem 8.1.1 is explained in more detail in Sections 3.2 and 3.3.

Sprinkled renormalization of the two-point function. We now explain our unconditional proof⁴ of Reformulation 3, which occupies the bulk of the paper. At a very high level, we will use a “sprinkled multi-scale induction argument”, in which we start with estimates concerning percolation on some scale and deduce similar estimates at a much larger scale after increasing p by some appropriately small amount; if we can do this efficiently enough, so that the total sprinkling is small when we start at a large scale, we can carry the induction up to infinitely many scales and (hopefully) prove that the resulting slightly larger parameter is supercritical (or at least that it is not subcritical).

Since our actual induction hypothesis is rather complicated, let us first illustrate how such an argument might work in principle. Let $d \geq 1$ be fixed and suppose that we were able to prove an implication of the form

$$\min_{u,v \in B_n} \mathbb{P}_p^G(u \leftrightarrow v) \geq \delta(n) \implies \min_{u,v \in B_{\phi(n)}} \mathbb{P}_{p+\varepsilon(n)}^G(u \leftrightarrow v) \geq \eta(n) \quad (3.1.3)$$

held for all $G \in \mathcal{U}_d^*$, $p \in [0, 1]$, and $n \geq 1$, where $\phi : \mathbb{N} \rightarrow \mathbb{N}$ is strictly increasing and $\varepsilon, \delta, \eta : \mathbb{N} \rightarrow (0, 1]$ are decreasing. We claim that such an implication would suffice to prove locality *provided that sufficiently strong quantitative relationships hold between the various functions that appear*. For example, the argument would work provided that $\eta(n) \geq \delta(\phi(n))$, that $\delta(n)$ is subexponentially small as a function of n , and that $\sum_{k=0}^{\infty} \varepsilon(\phi^k(n)) < \infty$, where ϕ^k denotes the k -fold convolution of ϕ . (Consider for example $\phi(n) = (n+1)^2$, $\delta(n) = (\log n)^{-1}$, $\eta(n) = (2 \log n)^{-1}$, and $\varepsilon(n) = (\log \log n)^{-2}$.) To see this, note that if we define the sequence $(n_k)_{k \geq 1}$ by $n_0 = 1$ and $n_{k+1} = \phi(n_k)$ for each $k \geq 1$ then, under this assumption, the implication (3.1.3) implies that

$$\min_{u,v \in B_{n_k}} \mathbb{P}_p^G(u \leftrightarrow v) \geq \delta(n_k) \implies \min_{u,v \in B_{n_{k+1}}} \mathbb{P}_{p+\varepsilon(n_k)}^G(u \leftrightarrow v) \geq \delta(n_{k+1})$$

³Indeed, unlike the tail of the radius, it is possible that point-to-point connection probabilities continue to decay in the supercritical regime, as is the case on the 3-regular tree. This breaks the argument that Reformulation 1 implies Theorem 8.1.1 that we gave above.

⁴As discussed in Remark 3.1.5, we do not quite proceed via Reformulation 3 in order to circumvent certain technical obstacles. Despite these caveats, we still think of our argument to be best understood as “morally” going via Reformulation 3, which in any case is implied by Theorem 8.1.1.

and hence by induction that

$$\begin{aligned} \min_{u,v \in B_{n_k}} \mathbb{P}_p^G(u \leftrightarrow v) &\geq \delta(n_k) \quad \text{for some } k \geq 1 \\ \implies \min_{u,v \in B_{n_i}} \mathbb{P}_{p+\sum_{j=k}^{i-1} \varepsilon(n_j)}^G(u \leftrightarrow v) &\geq \delta(n_{k+1}) \quad \text{for every } i \geq k. \end{aligned}$$

Since the conclusion on the right hand side implies that connection probability are subexponential in the distance at $p + \sum_{j=k}^{\infty} \varepsilon(n_j)$, it would follow from (3.1.3) that if $\min_{u,v \in B_{n_k}} \mathbb{P}_p^G(u \leftrightarrow v) \geq \delta(n_k)$ for some k then $p_c \leq p + \sum_{j=k}^{\infty} \varepsilon(n_j)$. The assumption that $\sum_{j=0}^{\infty} \varepsilon(n_j)$ ensures that the tail sum appearing here is small, verifying Reformulation 3.

Following this approach, we are led to the problem of how to extend point-to-point connection lower bounds from one scale to a much larger scale after sprinkling by a small amount. More specifically, we want to do this *as efficiently as possible*, with the hope of obtaining an inductive statement that is sufficiently strong to imply locality.

Snowballing. In Section 3.4, we develop a new method based on Talagrand’s theory of sharp thresholds [Tal94] and “cluster repulsion” inequalities inspired by the work of Aizenman-Kesten-Newman [AKN87b] that allows us to prove a bootstrapping implication of the form (3.1.3), where the function ϕ depends on the volume of the ball of radius n . While the basic idea of using Talagrand together with Aizenman-Kesten-Newman is already present in [DKT21; CMT22] (both in the polynomial growth case), we find a new way of both implementing and applying⁵ this argument using *ghost fields* that works directly in infinite volume and uses the *two-ghost inequality* of [Hut20a] rather than the classical Aizenman-Kesten-Newman inequality. We call this the **snowballing** method. While the methods of [DKT21; CMT22] needed upper bounds on the growth to work, our method actually becomes *more efficient* as the growth gets larger; if the growth is large enough, the bootstrapping implication we obtain (which is of the form (3.1.3)) is strong enough to prove locality by the argument outlined above. Optimizing the snowballing argument as much as we could (see Remark 3.4.5), we found that this method could prove locality for graph sequences satisfying a uniform growth lower bound of the form $n^{c(\log \log n)^C}$ for some universal constant C and any $c > 0$.

For Cayley graphs, the hypothesis needed for this argument to work is of course frustratingly close to the Shalom-Tao bound $\text{Gr}(n) \geq \exp(c \log n (\log \log n)^c)$ [ST10b], where c is a *small* universal

⁵In particular, the argument we use to efficiently extend two-point estimates to a higher scale given the outputs of this sharp threshold argument is also novel.

constant, which holds for every Cayley graph of superpolynomial growth. Thus, even a modest improvement to Shalom-Tao or the snowballing argument would allow us to prove locality for all sequences of superpolynomial growth via this method. Since we were not able to improve either argument to the required extent (and in any case must also deal with the “diagonal” case of locality), we will instead attack the locality conjecture in a “pincer movement”, where we push the polynomial growth methods of [CMT22] to handle all growths that are too slow to be treated by the snowballing argument. In fact we will push these methods to handle all graphs of *quasi-polynomial* growth $\text{Gr}(n) \leq \exp((\log n)^C)$, so that there is a considerable overlap in the growth regimes handled by the two methods. As mentioned above, a key technical difficulty is that (as far as we know) the same graph may oscillate between the two growth regimes infinitely often, so that the two methods must be harmonized in some way to allow for this.

Chaining via orange peeling. A well-known general approach to efficiently extend connection lower bounds from one scale to another is by a method we will call *chaining*. Consider the event \mathcal{E}_R that $S_R \leftrightarrow S_{10R}$ and the event \mathcal{U}_R that there is at most one cluster intersecting both S_{2R} and S_{5R} . Suppose we knew that for some large R and small $\varepsilon > 0$,

$$\mathbb{P}_p(\mathcal{E}) \geq 1 - \varepsilon \quad \text{and} \quad \mathbb{P}_p(\mathcal{U}) \geq 1 - \varepsilon. \quad (3.1.4)$$

Then, by Harris’ inequality and a union bound, we could deduce that any pair of vertices u and v at distance kR satisfy

$$\mathbb{P}_p(u \leftrightarrow v) \geq \min_{u', v' \in B_{10R}} \mathbb{P}_p(u' \leftrightarrow v')^2 \cdot [1 - k\varepsilon] - k\varepsilon. \quad (3.1.5)$$

If we also had a way to use (3.1.5) to deduce a version of (3.1.4) at scale kR in place of R (possibly after a small increase of p), *and this argument is sufficiently efficient quantitatively*, we might be able to formulate an inductive argument yielding both connection lower bounds and uniqueness of annuli crossings at *all* scales. (Of course such an argument cannot work on e.g. the 3-regular tree.)

This is roughly what is done in [CMT22] (although they consider supercritical percolation, so that crossing probabilities do not decay *a priori*), who establish that \mathcal{U}_R holds with high probability when G has polynomial growth and $R \rightarrow \infty$. Their method, which (following [Gri99]) they refer to as *orange peeling*⁶ and is inspired by the earlier work [BT17], relies on knowing a two-point lower bound within annuli of the form $B_{R+\Delta} \setminus B_R$ for all large R and some appropriate $\Delta(R) \ll R$. This information needs to be proven using information at lower scales, so that the argument is a kind of multi-scale induction or coarse-grained renormalization. The idea is to argue that as two

⁶We do not completely understand the metaphor; perhaps onion peeling would be more appropriate?

clusters cross the thick annulus $B_{5R} \setminus B_{2R}$, they are very likely to become connected to each other because they have to cross many disjoint annuli of the form $B_{R'+\Delta} \setminus B_{R'}$, and in each such annulus they have a good probability to become connected to each other after sprinkling. If there are enough opportunities for any two crossing clusters to merge then all crossing clusters will have merged by the time they reach the outer sphere. The full argument of [CMT22] is highly technical and sophisticated, with the discussion in this paragraph presenting only a simplified cartoon version of some selected parts of their argument.

Working with quasi-polynomial growth. To run the something like the above “orange peeling” method, it is helpful to know that annuli $B_{R_2} \setminus B_{R_1}$ are in some sense well-connected. For example, in the setting of [CMT22], the authors used the fact that for any pair of vertices u, v in the *exposed sphere* S_R^∞ , which is a certain subset of the usual sphere S_R , there is a path from u to v that stays within $B_{R+\Delta} \setminus B_{R-\Delta}$. To prove this they used the structure theory of graphs of polynomial growth (i.e. the fact that such graphs are finitely presented and are one-ended when they are not one-dimensional). As such, this method does not easily generalize to other graphs. (Indeed, any reasonable connectivity-of-annuli statement cannot hold in complete generality — annuli in regular trees are as poorly connected as sets at a given distance can possibly be.) We will prove that a weak connectivity property of annuli does hold if we assume that a quasi-polynomial growth upper bound $\text{Gr}(R) \leq \exp((\log R)^C)$ holds around the relevant scales. This statement, which we call the *polylog-plentiful tubes condition*, says that for any two sets of vertices A and B crossing a thick annulus $B_{3R} \setminus B_R$, we can find many (i.e., at least $(\log R)^{\Omega(1)}$) paths from A to B that are not too long (i.e., have length at most $R(\log R)^{O(1)}$) and that are well-separated from each other (i.e., any two paths in the set have distance at least $(\log R)^{\Omega(1)}$). The proof of this polylog-plentiful tubes condition will exploit a structure vs. randomness dichotomy, where we use the structure theory of [EH23d] to handle the low-dimensional case and handle the high-dimensional case by building the required disjoint tubes using coupled families of random walks. Once this polylog-plentiful tubes condition is verified, we then show it can be used to push the methods of [CMT22] to handle graphs of quasi-polynomial growth. (This requires significant technical changes throughout their entire argument, with the proofs of some intermediate steps being completely different.)

We might hope that the low- and high-growth arguments both imply a common statement similar to (3.1.3). Unfortunately our induction statement is more involved than this. The issue is that our low-growth argument works with point-to-point connections *within finite sets* such as balls and tubes, while the high-growth argument works directly in infinite volume, and our induction hypothesis must be able to handle oscillation between these two cases. Our actual induction statement, which

we explain in detail in Section 3.3, therefore has two parts: a lower bound on the full-space (infinite volume) two-point function, with no restriction on the geometry of the graph at the relevant scale, together with a lower-bound on connection probabilities *inside tubes* that holds only when the *thickness* of the tube happens to belong to a quasi-polynomial growth scale. (The *length* of these tubes are permitted to be quasi-polynomial in their thickness, so that this estimate tells us not just about quasi-polynomial growth scales but also many subsequent larger scales.) Thus, we are able to formulate a multi-scale inductive implication that holds *in all growth regimes* and is strong enough to imply Theorem 8.1.1.

Organization and overview

We now outline the structure of the rest of the paper.

Section 2: In this section we review both the classical Aizenman-Kesten-Newman inequality [AKN87b] and its consequences for the “uniqueness zone” as derived in [CMT22] as well as the *two-ghost inequality* of [Hut20a], which is a kind of infinite-volume version of Aizenman-Kesten-Newman analyzing volumes of clusters rather than their diameters. We also state a useful lemma of [Hut20a] that applies this inequality and that leads to Reformulation 3 above.

Section 3: In this section we formulate the multi-scale induction framework used to prove locality, stating the induction step as Proposition 3.3.1. We then explain how this technical proposition both implies Theorem 8.1.1 and leads to a new proof of $p_c < 1$ for transitive graphs that are not one-dimensional as originally proven in [DGRSY20].

Section 4: In this section we develop a new method of deducing two-point function lower bounds at a large scale from lower bounds at a smaller scale, which we call **snowballing**. This is based in part on an idea already present in [DKT21; CMT22], in which one uses Aizenman-Kesten-Newman bounds as an input to Talagrand’s sharp threshold theorem [Tal94]. Unlike those works, we use the two-ghost formulation of Aizenman-Kesten-Newman from [Hut20a] to work directly with volumes of clusters and prove bounds that hold in arbitrary unimodular transitive graphs; this causes the inequalities we derive to become vastly more efficient as the growth of the graph gets larger. The proof of the main snowballing proposition also involves a novel ghost-based chaining argument to convert this sharp threshold statement into a statement about extending point-to-point connection probabilities. All the arguments in this section work more efficiently as the growth gets larger, but still have content in the polynomial growth case. In Section 2.2, we explain how the snowballing method implies part of the main multi-scale induction step and also easily implies the full locality conjecture for graphs satisfying a mild superpolynomial growth lower bound. (Here the growth

assumption is needed only to ensure that the total amount of sprinkling is small when we start at a large scale and induct up to infinity; no further geometric properties of graphs of high growth are used.) The tools we develop in Section 2.1 also play an indispensable role in our low-growth arguments in Section 3.6.

Section 5: In this section we establish deterministic geometric features of transitive graphs of quasi-polynomial growth (at some scale) that will be used to analyze percolation on these graphs in Section 3.6. The main question addressed is as follows: how can one harness low — but not necessarily polynomial — volume growth to establish that annuli are well-connected? Here, the “well-connectedness” of annuli is made precise via what we call the *polylog-plentiful tubes condition*. The proof that the polylog-plentiful tubes condition holds for graphs of quasi-polynomial growth uses a “structure vs. expansion” dichotomy according to whether the *rate of growth* on the relevant scale is low or high dimensional; in the low-dimensional case we employ the structure theory of [BGT12a; TT21a; EH23d] while in the high-dimensional case we construct large families of disjoint tubes using certain coupled random walks; the quasi-polynomial growth assumption is used to ensure that we can couple two walks started at distance n to coalesce within distance $n(\log n)^{O(1)}$. (The low dimensional case of the analysis is the only place where we directly use the fact that our graphs are not one-dimensional, where it is needed to ensure that “exposed spheres” are well-connected in a certain sense.)

Section 6: Using the polylog-plentiful tubes condition, we run a chaining argument for graphs of quasi-polynomial growth (on some scale) that is inspired by [CMT22]. Many changes to their argument are required to deal with this new geometric setting, and our arguments in this section also make use of the snowballing method developed in Section 3.4. We then use the outputs of this analysis to complete the proof of the induction step and hence of Theorem 8.1.1.

Section 7: In this section we first briefly sketch how the methods of Section 3.4 can be used to prove locality of the *density* of the infinite cluster in the supercritical phase; a significant generalization of this result is proven in full detail (via a different method) in our forthcoming paper [EH23+a]. We finish by discussing some open problems in Section 3.7.

3.2 Cluster repulsion and the *a priori* uniqueness zone

In this section we review various bounds on the probability that two large clusters either meet at a single edge or both intersect a small ball, together with some important consequences of these inequalities. These inequalities originate inexplicitly in the work of Aizenman, Kesten, and Newman [AKN87b] and were brought to the wider attention of the community in the work of Cerf

[Cer15]. These inequalities come in two flavours: finite-volume estimates that work with the *radii* of clusters, and infinite-volume estimates that work with the *volume* of clusters. The first kind of inequality, which is closer to the original vision of [AKN87b; Cer15], works best under a low-growth assumption and plays an important role in Contreras, Martineau, and Tassion's analysis of percolation on transitive graphs of polynomial growth [CMT22; CMT23a], while the second kind of inequality, introduced in [Hut20a] and known as *two-ghost inequalities*, become *more* useful the faster the graph grows. We will also review the argument of [Hut20a] which allows us to deduce locality of p_c from uniform bounds on the two-point function using the two-ghost inequality.

Aizenman-Kesten-Newman and the *a priori* uniqueness zone

We now review the finite-volume Aizenman-Kesten-Newman inequality as presented in [Cer15; CMT22].

Finite-volume two-arm estimates. Given $m, n \in (0, \infty)$ with $m \leq n$, we define $\text{Piv}[m, n]$ to be the event that there exist two distinct clusters in $\omega \cap B_n$ that each intersect both spheres S_n and S_m . (That is, $\text{Piv}[m, n]$ is the event that there is more than one cluster crossing the annulus from m to n .) The following lemma is a minor variation on [CMT22, Proposition 4.1].

Proposition 3.2.1. *For each $0 < \varepsilon < 1/2$, $0 < \eta < 1$, and $d \geq 1$ there exists a constant $C = C(\varepsilon, \eta, d)$ such that if G is a connected, transitive graph of vertex degree d then*

$$\mathbb{P}_p(\text{Piv}[1, n]) \leq C \left[\frac{\log \text{Gr}(n)}{n} \right]^{\frac{1}{2} - \varepsilon},$$

for every $p \in [\eta, 1]$ and $n \geq 1$.

To deduce this proposition from the proof of [CMT22, Proposition 4.1], we will require the following elementary fact about the growth of balls. Here $\partial B_m(o)$ denotes the set of edges that have one endpoint in $B_m(o)$ and the other endpoint in the complement of $B_m(o)$.

Lemma 3.2.2. *For each $d \geq 1$ there exists a constant C_d such that the following holds. Let G be a graph of maximum vertex degree at most d , and let o be a vertex of G . For each integer $n \geq 1$ there exists an integer $n \leq m \leq 2n - 1$ such that*

$$\frac{|\partial B_m(o)|}{|B_m(o)|} \leq \frac{C_d}{n} \log |B_{2n}(o)|.$$

Proof of Lemma 3.2.2. Write $\text{Gr}(m) := |B_m(o)|$ for every $m \in \mathbb{N}$. Since G has vertex degrees bounded above by d , we have that $|\partial B_m(o)| \leq d(\text{Gr}(m + 1) - \text{Gr}(m))$ and hence that

$|\partial B_m(o)| / \text{Gr}(m) \leq d(\text{Gr}(m+1)/\text{Gr}(m) - 1)$. It follows that

$$\sum_{m=n}^{2n-1} \frac{|\partial B_m(o)|}{|B_m(o)|} \leq d \sum_{m=n}^{2n-1} \left(\frac{\text{Gr}(m+1)}{\text{Gr}(m)} - 1 \right).$$

Now, since we also have that $\frac{\text{Gr}(m+1)}{\text{Gr}(m)} \leq d$, there exists a constant $C = C_d$ such that

$$\frac{\text{Gr}(m+1)}{\text{Gr}(m)} - 1 \leq C_d \log \frac{\text{Gr}(m+1)}{\text{Gr}(m)}$$

for every $m \geq 0$. As such, it follows that

$$\sum_{m=n}^{2n-1} \frac{|\partial B_m(o)|}{|B_m(o)|} \leq dC_d \sum_{m=n}^{2n-1} \log \frac{\text{Gr}(m+1)}{\text{Gr}(m)} = dC_d \log \frac{\text{Gr}(2n)}{\text{Gr}(n)}$$

which is easily seen to imply the claim. \square

Proof of Proposition 3.2.1. This follows by exactly the same proof as [CMT22, Proposition 4.1] except that we use our Lemma 3.2.2 instead of their Lemma 4.2 (which is the same estimate specialized to the polynomial growth setting). \square

The *a priori* uniqueness zone. We now discuss how two-arm bounds at a single edge can be used to deduce bounds on the probability of having multiple clusters crossing an annulus. The following lemma, essentially due to Cerf [Cer15], lets us apply Proposition 3.2.1 to bound the probability that there are two distinct crossings of an annulus.

Lemma 3.2.3 ([CMT22, Lemma 6.2]). *Let G be a connected transitive graph. Then*

$$\mathbb{P}_p(\text{Piv}[r, n]) \leq \mathbb{P}_p(\text{Piv}[1, n/2]) \cdot \frac{|S_r|^2 \cdot \text{Gr}(m)}{\min_{a,b \in S_r} \mathbb{P}_p(a \xrightarrow{B_m} b)} \quad (3.2.1)$$

for every $r, m, n \in (1, \infty)$ with $r \leq m \leq n/2$ and every $p \in (0, 1)$.

Remark 3.2.1. The authors of [CMT22] stated this lemma in their context of infinite graphs with polynomial growth. However, their proof works exactly the same for arbitrary connected transitive graphs. (The statement given in [CMT22] has B_{2m} in place of B_m — which is slightly stronger — but this appears to be a typo.)

Since any geodesics from o to two vertices $a, b \in B_m$ are both open with probability at least p^{2m} and the growth satisfies the trivial upper bound $\text{Gr}(m) \leq d^{m+1}$, the quantity multiplying the probability $\mathbb{P}_p(\text{Piv}[1, n/2])$ on the right hand side of (3.2.1) is at most exponential in m when p is bounded away from 0. Proposition 3.2.1 and Lemma 3.2.3 therefore have the following immediate corollary.

Corollary 3.2.4 (Trivial *a priori* uniqueness zone). *Let G be a connected, unimodular transitive graph with vertex degree d and let $\varepsilon, \eta \in (0, 1)$. There exist positive constants $c = c(d, \eta, \varepsilon)$, $C = C(d, \eta, \varepsilon)$, and $n_0 = n_0(d, \eta, \varepsilon)$ such that*

$$\mathbb{P}_p(\text{Piv}[c \log n, n]) \leq C \left[\frac{\log \text{Gr}(n)}{n} \right]^{1/2-\varepsilon}$$

for every $n \geq n_0$ and $p \in [\eta, 1]$.

In other words, if G has subexponential growth then the probability of having two distinct crossings of the annulus from $c \log n$ to n is always small for an appropriately small constant c . (The restriction that p is small could be removed by noting that it is very unlikely for there to be *any* crossings of the annulus when $p \leq 1/d$.)

The two-ghost inequality

We now recall the *two-ghost inequality* of [Hut20a], a form of the Aizenman-Kesten-Newman bound that holds for any unimodular transitive graph, without any growth assumptions. (This version of the bound does *not* imply uniqueness of the infinite cluster since it requires at least one of the clusters to be finite.) Let $G = (V, E)$ be a graph. For each edge e of G and $n \geq 1$ we define $\mathcal{S}_{e,n}$ to be the event that e is closed and that the endpoints of e are in distinct clusters, each of which has volume at least n and at least one of which is finite.

Theorem 3.2.5 (Two-ghost inequality). *Let G be a unimodular transitive graph of degree d . There exists a constant C_d such that*

$$\mathbb{P}_p(\mathcal{S}_{e,n}) \leq C_d \left[\frac{1-p}{pn} \right]^{1/2} \quad (3.2.2)$$

for every $e \in E(G)$, $p \in (0, 1]$ and $n \geq 1$.

Proof. This is an immediate consequence of [Hut20a, Corollary 1.7]: The statement given there concerns the edge volume rather than the vertex volume, but this only makes the statement stronger. \square

Theorem 3.2.5 has the following useful consequence concerning the probability that two large distinct clusters come close to one another; this lets us convert bounds on the two-point function into bounds on the tail of the volume and underlies the fact that Reformulation 3 implies the unimodular case of Theorem 8.1.1.

Lemma 3.2.6. *Let G be an infinite, connected, unimodular transitive graph with vertex degree d . There exists a constant C_d such that*

$$\mathbb{P}_p(|K_u| \geq n \text{ and } |K_v| \geq n \text{ but } u \leftrightarrow v) \leq C_d \cdot d(u, v) p^{-d(u, v)-1} n^{-1/2},$$

for every $u, v \in V(G)$, $n \geq 1$, and $p < p_c(G)$.

Proof. This follows from Theorem 3.2.5 by the same argument used to prove equation (4.2) of [Hut20a]. (Again, the only difference is that we are using vertex volumes rather than edge volumes.) The quantity $d(u, v)p^{-d(u, v)-1}$ arises as a simple upper bound on $((1-p)/p)^{1/2} \sum_{i=1}^{d(u, v)} p^{-i+1}$. \square

3.3 The multi-scale induction step

In this section we state our key technical proposition, Proposition 3.3.1, which encapsulates the multi-scale induction used to prove Theorem 8.1.1. We introduce relevant definitions in Section 3.3, give the statement of the proposition in Section 3.3, and explain how it implies our main theorem in Section 3.3. In Section 3.3 we also explain how Proposition 3.3.1 yields a new proof of the fact that $p_c < 1$ for all infinite, connected, transitive graphs that are not one-dimensional.

Definitions

In this section we establish the notation necessary to state the main multi-scale induction proposition in Section 3.3.

Natural coordinates for sprinkling. We define the **sprinkling function** $\text{Spr} : (0, 1) \times \mathbb{R} \rightarrow (0, 1)$ by

$$\text{Spr}(p; \lambda) = 1 - (1-p)^{e^\lambda} \quad \text{so that} \quad (1 - \text{Spr}(p; \lambda)) = (1-p)^{e^\lambda}.$$

The sprinkling functions $(\text{Spr}(\cdot; \lambda))_{\lambda \in \mathbb{R}}$ form a semigroup in the sense that

$$\text{Spr}(p; \lambda + \mu) = \text{Spr}(\text{Spr}(p; \lambda); \mu)$$

for every $\lambda, \mu \in \mathbb{R}$ and $p \in (0, 1)$. For each $0 < p, q < 1$ we define

$$\delta(p, q) = \log \left[\frac{\log(1 - \max\{p, q\})}{\log(1 - \min\{p, q\})} \right], \quad \text{so that} \quad \max\{p, q\} = \text{Spr}(\min\{p, q\}; \delta(p, q)).$$

Note that if $p \geq 1/d$, as we will assume throughout most of the paper, then for each $D < \infty$ there exists a positive constant $c = c(d, D)$ such that

$$(1 - \text{Spr}(p; \delta)) = (1-p)^{e^\delta-1}(1-p) \leq \left(\frac{2d-1}{2d} \right)^\delta (1-p) \leq (1-c\delta)(1-p) \quad (3.3.1)$$

for every $0 \leq \delta \leq D$, so that Bernoulli bond percolation with parameter $\text{Spr}(p; \delta)$ stochastically dominates the independent union of a Bernoulli- p bond percolation configuration and a Bernoulli- $(c\delta)$ bond percolation configuration. (Note however that $1 - \text{Spr}(p; \delta)$ is much smaller than $(1 - c\delta)(1 - p)$ when p is close to 1.)

The corridor function. Given a path γ , we write $\text{len}(\gamma)$ for its length. Given a path γ and some $r \in (0, \infty)$, we define $B_r(\gamma) := \bigcup_{i=0}^{\text{len}(\gamma)} B_r(\gamma_i)$, and call a set of this form a **tube**. We refer to $\text{len}(\gamma)$ and r as the **length** and **thickness** of the tube respectively. (Note that these parameters depend on the choice of representation of the tube $B_r(\gamma)$, and are not determined by the tube as a set of vertices.) Following [CMT22], we define the **corridor function** by

$$\kappa_p(m, n) := \inf_{\gamma: \text{len}(\gamma) \leq m} \mathbb{P}_p \left(\gamma_0 \xleftrightarrow{B_n(\gamma)} \gamma_{\text{len}(\gamma)} \right)$$

for each $p \in (0, 1)$ and $n, m \geq 1$, so that $\kappa_p(m, n)$ measures the difficulty of connecting points within tubes of thickness n and length at most m . We may also take $n = \infty$ in the definition of the corridor function, where the restriction for connections to lie in a tube disappears and we have simply that

$$\kappa_p(m, \infty) = \kappa_p(m) = \inf \{ \mathbb{P}_p(x \leftrightarrow y) : d(x, y) \leq m \}.$$

Note that the corridor function $\kappa_p(m, n)$ is increasing in p and n and decreasing in m .

Low growth scales. Given a transitive graph G and a parameter $D > 0$, we define the set of **low growth scales** to be

$$\mathcal{L}(G, D) = \left\{ n \geq 1 : \log \text{Gr}(m) \leq (\log m)^D \text{ for all } m \in [n^{1/3}, n] \right\},$$

so that

$$\left\{ n \geq 1 : \log \text{Gr}(n) \leq 3^{-D} (\log n)^D \right\} \subseteq \mathcal{L}(G, D) \subseteq \left\{ n \geq 1 : \log \text{Gr}(n) \leq (\log n)^D \right\}.$$

(We will sometimes call these **quasi-polynomial growth scales** since the function $\exp[(\log x)^C]$ is sometimes known as a quasi-polynomial.) For the purposes of the proof of the main theorem, we will apply this definition only with the (somewhat arbitrary) choice of constant $D = 20$.

The burn-in sprinkle. The first step of our induction will have a different form to the others, which necessitates a possibly larger amount of sprinkling. We now introduce notation describing this initial amount of sprinkling. Given a transitive graph G , $p \in (0, 1)$, and $m \geq 1$ define

$$b(m) = b(m, p) = \max \left\{ b \in \mathbb{N} : 1 \leq b \leq \frac{1}{8} m^{1/3} \text{ and } \mathbb{P}_p(\text{Piv}[4b, m^{1/3}]) \leq (\log m)^{-1} \right\},$$

setting $b(m) = 0$ if the set being maximized over is empty, and define the **burn-in** to be

$$\begin{aligned}\text{Burn}(n, p) &= \text{Burn}(G, n, p) \\ &= \max \left\{ \left(\frac{\log \log m}{\min\{\log m, \log \text{Gr}(b(m))\}} \right)^{1/4} : m \in \mathcal{L}(G, 20) \cap \left[(\log n)^{1/2}, n \right] \right\},\end{aligned}$$

setting $\text{Burn}(n, p) = 0$ if the set being maximized over is empty and setting $\text{Burn}(n, p) = \infty$ if there exists m belonging to the set for which $b(m) \leq 1$. (If $b(m) > 1$ then $\log \text{Gr}(b(m))$ and $\log m$ are both positive.) Note that $\text{Burn}(G, n, p)$ is determined by the ball of radius n in G , so that two graphs whose balls of this radius are isomorphic have the same value of $\text{Burn}(n, p)$ for each $p \in (0, 1)$.

Statement of the induction step

We are now ready to state our main multi-scale induction proposition. We recall that, when applied to logical propositions, the symbol “ \vee ” means “or” while the symbol “ \wedge ” means “and”. The condition $n_0 \geq 16$ appearing in the proposition ensures that $\log \log n_0 \geq 1$.

Proposition 3.3.1 (The main multi-scale induction step). *For each $d \in \mathbb{N}$ there exist constants $K = K(d)$ and $N = N(d) \geq 16$ such that the following holds. Let G be an infinite, connected, unimodular transitive graph with vertex degree d that is not one-dimensional, let $p_0 \in (0, 1)$, and let $n_0 \geq 16$. Let $n_{-1} = (\log n_0)^{1/2}$, let*

$$\delta_0 = \frac{1}{(\log \log n_0)^{1/2}} + K \cdot \text{Burn}(n_0, p_0),$$

define sequences $(n_i)_{i \geq 1}$ and $(\delta_i)_{i \geq 1}$ recursively by

$$n_i := \exp((\log n_{i-1})^9) = \exp^{\circ 3} \left(\log^{\circ 3}(n_0) + i \log 9 \right) \quad \text{and} \quad \delta_i := (\log \log n_i)^{-1/2} = 3^{-i} (\log \log n_0)^{-1/2},$$

and let $(p_i)_{i \geq 1}$ be an increasing sequence of probabilities satisfying $p_{i+1} \geq \text{Spr}(p_i; \delta_i)$ for each $i \geq 0$. For each $i \geq 0$ define the statement

$$\text{FULL-SPACE}(i) = \left(\mathbb{P}_{p_i}(u \leftrightarrow v) \geq \exp \left[-(\log \log n_i)^{1/2} \right] \text{ for all } u, v \in B_{n_i} \right)$$

and for each $i \geq 1$ define the statement

$$\text{CORRIDOR}(i) = \left(\kappa_{p_i}(e^{[\log m]^{10}}, m) \geq \exp \left[-(\log \log n_i)^{1/2} \right] \text{ for every } m \in \mathcal{L}(G, 20) \cap [n_{i-2}, n_{i-1}] \right).$$

If $n_0 \geq N$, $p_0 \geq 1/d$, and $\delta_0 \leq 1$ then the implications

$$\text{FULL-SPACE}(0) \implies [\text{CORRIDOR}(1) \vee (p_1 \geq p_c)], \quad (\text{C}_0)$$

$$[\text{FULL-SPACE}(i) \wedge \bigwedge_{k=1}^i \text{CORRIDOR}(k)] \implies [\text{CORRIDOR}(i+1) \vee (p_{i+1} \geq p_c)], \quad \text{and} \quad (\text{C})$$

$$[\text{FULL-SPACE}(j) \wedge \text{CORRIDOR}(j+1)] \implies [\text{FULL-SPACE}(j+1) \vee (p_{j+1} \geq p_c)] \quad (\text{F})$$

hold for every $i \geq 1$ and $j \geq 0$.

Remark 3.3.1. Note that $\text{CORRIDOR}(1)$ has a significantly different form than $\text{CORRIDOR}(i)$ for $i \geq 2$ since n_{-1} is much smaller than the natural extrapolation of the sequence $(n_i)_{i \geq 0}$ to $i = -1$.

Remark 3.3.2. The condition $p \geq 1/d$ appearing in the hypotheses of Proposition 3.3.1 is in fact redundant: An elementary path counting argument yields that

$$\mathbb{P}_p(u \leftrightarrow v) \leq \mathbb{P}_p(\text{there is a simple open path of length at least } d(u, v) \text{ starting from } u)$$

$$\leq \frac{d}{d-1} \cdot \frac{1}{1-p(d-1)} \cdot (p(d-1))^{d(u,v)}$$

for every $p < 1/(d-1)$ and $u, v \in V(G)$, so that if n_0 is sufficiently large and $\text{FULL-SPACE}(0)$ holds then $p \geq 1/d$. We include this redundant assumption anyway to clarify the structure of the proof.

Remark 3.3.3. In the statement $\text{CORRIDOR}(i)$, the tubes that arise in the relevant corridor function $\kappa_{p_i}(e^{[\log m]^{10}}, m)$ have *thickness* given by the low-growth scale m , but can have length equal to the much larger value $e^{[\log m]^{10}}$. As such, the statement $\text{CORRIDOR}(i)$ gives us strong control of percolation not just at low-growth scales but at a large range of scales above each low-growth scale. In particular, provided n_i is sufficiently large that $e^{[\log(n_i^{1/3})]^{10}} \geq e^{(\log n_i)^9} = n_{i+1}$, the implication (F) holds trivially whenever $n_i \in \mathcal{L}(G, 20)$.

Remark 3.3.4. The “ $\vee(p_{i+1} \geq p_c)$ ” that appears on the right hand side of the implications (C₀), (C), and (F) can be removed if one assumes that G is amenable, or if one works with “wired” connections as discussed in Section 3.7. This would be useful if one wished to use (a modification of) our methods to study the geometry of the infinite cluster in graphs of low growth, extending the results of [CMT22; Hut23b] to this setting.

Most of the paper is dedicated to proving Proposition 3.3.1: The implication (F) is proven in Section 3.4 while the implications (C₀) and (C) are proven in Sections 3.5 and 3.6.

Deduction of the main theorem from the induction step

In this section we deduce our main theorem, Theorem 8.1.1, from Proposition 3.3.1. We write $p_\infty = \lim_{i \rightarrow \infty} p_i$ for the limit of the parameters $(p_i)_{i \geq 0}$ defined recursively by $p_{i+1} = \text{Spr}(p_i; \delta_i)$, so that

$$p_\infty = \text{Spr}\left(p_0; \sum_{i=0}^{\infty} \delta_i\right) \leq \text{Spr}\left(p_0; 2(\log \log n_0)^{-1/2} + K \cdot \text{Burn}(n_0, p_0)\right).$$

We will apply Proposition 3.3.1 via the following corollary, which is a consequence of Proposition 3.3.1 together with the sharpness of the phase transition.

Corollary 3.3.2. *Let $d \geq 1$ and let $K = K(d)$ and $N = N(d)$ be the constants from Proposition 3.3.1. Let G be an infinite, connected, unimodular transitive graph with vertex degree d that is not one-dimensional. Then the implication*

$$\begin{aligned} & \left(\left[\mathbb{P}_p(u \leftrightarrow v) \geq e^{-(\log \log n)^{1/2}} \text{ for every } u, v \in B_n \right] \wedge [\delta_0 \leq 1] \right) \\ & \implies \left[p_c(G) \leq \text{Spr}\left(p; 2(\log \log n)^{-1/2} + K \cdot \text{Burn}(n, p)\right) \right], \end{aligned}$$

holds for every $n \geq N$ and $p \geq 1/d$.

Proof of Corollary 3.3.2 given Proposition 3.3.1. It suffices to prove that $p_c \leq p_\infty$ whenever $n \geq N$ and $p \geq 1/d$ are such that $\mathbb{P}_p(u \leftrightarrow v) \geq e^{-(\log \log n)^{1/2}}$ for every $u, v \in B_n$ and $\delta_0 \leq 1$. Fix one such n and p . Set $n_0 = n$ and define $(n_i)_{i \geq 0}$ as in Proposition 3.3.1. Since **FULL-SPACE(0)** holds by assumption, it follows from Proposition 3.3.1 that either $p_i \geq p_c$ for some $i \geq 1$ or **FULL-SPACE(i)** and **CORRIDOR(i)** hold for every $i \geq 1$. In the former case we may trivially conclude that $p_c \leq p_\infty$, while in the latter case we have that

$$\mathbb{P}_{p_\infty}(u \leftrightarrow v) \geq e^{-(\log \log n_i)^{1/2}} \quad \text{for every } i \geq 0 \text{ and every } u, v \in B_{n_i}. \quad (3.3.2)$$

On the other hand, it follows from the sharpness of the phase transition [Men86; AB87a] that for each $p < p_c$ there exists a positive constant c_p such that $\mathbb{P}_p(u \leftrightarrow v) \leq e^{-c_p d(u, v)}$ for every $u, v \in V(G)$. This is incompatible with the (very) subexponential lower bound (3.3.2), so that $p_\infty \geq p_c$ in this case also. \square

We now apply Corollary 3.3.2 to prove Theorem 8.1.1. The proof we give here will rely on the results of both [Hut20a; Hut20e] (to deal with the nonunimodular case) and [CMT22; CMT23a] (to deal with the case that the limit has polynomial growth). We remark that the quantitative proof of the theorem given in [EH23+b] yields a completely self-contained and “uniform in the

graph” deduction of locality from Proposition 3.3.1 in the unimodular case that does not rely on the results⁷ of [CMT22; CMT23a]. (Doing this requires non-trivial bounds on the burn-in which can be avoided in the case that the limit has superpolynomial growth as we will see below.)

Proof of Theorem 8.1.1 given Proposition 3.3.1. Let $(G_m)_{m \geq 1}$ be a sequence of infinite, connected, transitive graphs converging locally to some infinite transitive graph G , and suppose that the graphs G_m all have superlinear growth. We want to prove that $p_c(G_m) \rightarrow p_c(G)$. If G is nonunimodular the claim follows from the results of [Hut20a; Hut20e] (specifically [Hut20a, Theorem 5.6]), while if G has polynomial growth the result follows from the main result of [CMT23a]. Thus, we may assume that G is unimodular and has superpolynomial growth. Since the set of nonunimodular graphs is both closed and open in \mathcal{G} [Hut20a, Corollary 5.5], we may assume that the graphs $(G_n)_{n \geq 1}$ are all unimodular. We may also assume that the graphs G_n and G all have the same vertex degree d . Let $N_1 = N_1(d)$ and $K = K(d)$ be the constants from Proposition 3.3.1.

Suppose for contradiction that $p_c(G_m)$ does not converge to $p_c(G)$. Since p_c is lower semi-continuous ([Pet, §14.2] and [DT16a, p.4]), we have that $\liminf_{m \rightarrow \infty} p_c(G_m) \geq p_c(G)$. Thus, by taking a subsequence, we may assume that $\inf_{m \geq 1} p_c(G_m) = p_* > p_c(G)$. Let $p_0 = (p_c(G) + p_*)/2$ so that $p_c(G) < p_0 < p_*$. For each $n \geq 1$, let $m(n)$ be minimal such that the balls of radius n are isomorphic in G_m and G for all $m \geq m(n)$. Let $c > 0$ be the constant from Corollary 3.2.4. Since G has superpolynomial growth, we have by [Gro81b; Tro84b] that

$$\lim_{n \rightarrow \infty} \frac{\log \log n}{\log \text{Gr}(c \log n; G)} = 0,$$

where $\text{Gr}(m; G)$ denotes the volume of the ball of radius m in G , and it follows from Corollary 3.2.4 that

$$\lim_{n \rightarrow \infty} \sup_{p \in [1/d, 1]} \sup_{m \geq m(n)} \text{Burn}(G_m, n, p) = 0. \quad (3.3.3)$$

In particular, there exists $N_2 \geq N_1$ (depending on the superpolynomial graph G and the sequence (G_m)) such that if $n_0 \geq N_2$ then $(\log \log n_0)^{-1/2} + K \cdot \text{Burn}(G_m, n_0, p_0) \leq 1$ and

$$\text{Spr}(p_0; 2(\log \log n_0)^{-1/2} + K \cdot \text{Burn}(G_m, n_0, p_0)) < p_*$$

for every $m \geq m(n_0)$. Thus, it follows from Corollary 3.3.2 that for every $n_0 \geq N_2$ and $m \geq m(n_0)$ there exist vertices u and v in the ball of radius n_0 in G_m such that

$$\mathbb{P}_{p_0}^{G_m}(u \leftrightarrow v) < \exp [-(\log \log n_0)^{1/2}].$$

⁷Of course many of the proof techniques remain closely inspired by these works!

Applying Lemma 3.2.6, we deduce that there exists a constant C_1 such that

$$\mathbb{P}_{p_0}^{G_m}(|K| \geq k)^2 \leq C_1 n_0 p_0^{-2n_0} k^{-1/2} + \exp [-(\log \log n_0)^{1/2}]$$

for every $n_0 \geq N_2$, $m \geq m(n_0)$, and $k \geq 1$. Taking $n_0 = \lceil c_2 \log k \rceil$ for an appropriately small constant c_2 (which makes the first term $O(k^{-1/4})$, say, and hence of lower order than the second term), it follows that there exist positive constants C_2 and c_3 such that

$$\mathbb{P}_{p_0}^{G_m}(|K| \geq k) \leq C_2 \exp [-c_3 (\log \log \log k)^{1/2}] \quad (3.3.4)$$

for every $n_0 \geq N_2$ and $m \geq m(n_0)$. Since $p_0 > p_c(G)$, the probability $\mathbb{P}_{p_0}^G(o \leftrightarrow \infty)$ is positive and it follows from (3.3.4) there exist k_0 and m_0 such that

$$\mathbb{P}_{p_0}^{G_m}(|K| \geq k_0) \leq \frac{1}{2} \mathbb{P}_{p_0}^G(o \leftrightarrow \infty) \quad (3.3.5)$$

for every $m \geq m_0$. On the other hand, if m is sufficiently large that balls of radius k_0 are isomorphic in G_m and G then

$$\mathbb{P}_{p_0}^{G_m}(|K| \geq k_0) \geq \mathbb{P}_{p_0}^{G_m}(o \leftrightarrow B_{k_0}^c) = \mathbb{P}_{p_0}^G(o \leftrightarrow B_{k_0}^c) \geq \mathbb{P}_{p_0}^G(o \leftrightarrow \infty),$$

which contradicts the upper bound of (3.3.5). \square

Let us now explain how Proposition 3.3.1 yields a new proof of the $p_c < 1$ theorem as originally established in [DGRSY20]. Recall that \mathcal{G}^* is the space of all infinite, connected, transitive graphs that are not one-dimensional.

Theorem 3.3.3. *Every graph $G \in \mathcal{G}^*$ satisfies $p_c(G) < 1$.*

Proof of Theorem 3.3.3 given Proposition 3.3.1. If G has exponential growth then sharpness of the phase transition easily implies that $p_c \leq (\lim_{r \rightarrow \infty} \text{Gr}(r)^{-1/r}) < 1$ (see also [Hut16; Lyo95]). Since every nonunimodular transitive graph is nonamenable and therefore has exponential growth, it suffices to consider the case that G is unimodular. On the other hand, if G has polynomial growth, then it is well-known that $p_c(G) < 1$ follows from the fact that $p_c(\mathbb{Z}^2) < 1$ and the structure theory of transitive graphs of polynomial growth as explained in detail in [HT21a, Section 3.4].

We now consider the case that G is unimodular and has superpolynomial growth. Let d denote the vertex degree of G and let $N_1 = N_1(d)$ and $K = K(d)$ be the constants from Proposition 3.3.1. It follows by the same reasoning as we gave for (3.3.3) that for every $\eta > 0$ there exists a constant

$M(d, \eta)$ such that $b(m, p) \leq \eta$ for every $m \in \mathcal{L}(G, 20)$ with $m \geq M$ and every $p \geq 1/d$. Since we also have trivially that $\text{Gr}(n) \geq n$ for every $n \geq 1$ since G is infinite, it follows that there exists a constant $N_2 = N_2(d)$ such that

$$(\log \log n)^{-1/2} + K \cdot \text{Burn}(n, p) \leq 1 \quad \text{for every } p \geq 1/d \text{ and } n \geq N_2. \quad (3.3.6)$$

Let $n_0 = n_0(d) = N_1 \vee N_2$. Since $\mathbb{P}_p(u \leftrightarrow v) \geq p^{d(u,v)}$ for every $u, v \in V$ and $p \in [0, 1]$, there exists a constant

$$p_0 = p_0(d) = \frac{1}{d} \vee \exp\left(-\frac{(\log \log n_0)^{1/2}}{n_0}\right)$$

satisfying $1/d \leq p_0 < 1$ such that $\mathbb{P}_{p_0}(u \leftrightarrow v) \geq \exp(-(\log \log n_0)^{1/2})$ for every $u, v \in B_{n_0}$. Since $n_0 = N_1 \vee N_2$, it follows from Corollary 3.3.2 and (3.3.6) that

$$p_c(G) \leq \text{Spr}\left(p_0; 2(\log \log n_0)^{-1/2} + K \text{Burn}(n_0, p_0)\right) \leq \text{Spr}(p_0; 2).$$

The claim follows since the right hand side is strictly less than one. \square

3.4 Making connections via sharp threshold theory

In this section we describe a powerful new way to extend point-to-point connection lower bounds from one scale to another, which we call the “snowballing method”. We develop this method in Section 3.4, then apply it to prove the implication (F) of Proposition 3.3.1 in Section 3.4. In Section 3.4 we will also explain how the method allows us to conclude the proof of locality for unimodular graphs satisfying a mild uniform superpolynomial growth assumption.

Snowballing

We now begin to develop the snowballing method. This method is primarily encapsulated through the following proposition, whose proof is the main goal of this section, but the intermediate lemmas used in its proof can be used to prove results of independent interest as discussed in Section 3.7. Given (not necessarily finite) non-empty sets of vertices A, B , and Λ in a graph $G = (V, E)$ and a parameter $p \in [0, 1]$, we define

$$\tau_p^\Lambda(A, B) := \min_{\substack{a \in A \\ b \in B}} \mathbb{P}_p\left(a \xleftrightarrow{\Lambda} b\right),$$

where we recall that $\{a \xleftrightarrow{\Lambda} b\}$ denotes the event that a is connected to b by an open path all of whose vertices belong to Λ . We will also write $\tau_p^\Lambda(A) := \tau_p^\Lambda(A, A)$ and $\tau_p(A, B) := \tau_p^V(A, B)$. We also use the notion of distance $\delta(p, q)$ between two parameters $p, q \in (0, 1)$ as defined in Section 3.3.

Proposition 3.4.1 (Snowballing). *For each $d \geq 1$ and $D < \infty$ there exist positive constants $c_1 = c_1(D)$ and $h_0 = h_0(d, D)$, and universal positive constants c_2 and c_3 such that the following holds. Let $G = (V, E)$ be a unimodular transitive graph with vertex degree d , let A_1, \dots, A_n be non-empty sets of vertices in G , and suppose that $0 < p_1 < p_2 < 1$ are such that there is at most one infinite cluster \mathbb{P}_p -almost surely for every $p \in [p_1, p_2]$. Let $h \geq (\min_i |A_i|)^{-1}$ and let r be a positive integer with $hr \geq 1$ such that $\mathbb{P}_p(\text{Piv}[1, hr]) < h$ for every $p \in [p_1, p_2]$. If $p_1 \geq 1/d$, $\delta = \delta(p_1, p_2) \leq D$, and $h \leq h_0$ then the implication*

$$\begin{aligned} & \left(h^{c_1\delta^3} \leq c_3 n^{-1} \text{ and } \tau_{p_1}^\Lambda(A_i \cup A_{i+1}) \geq 4h^{c_1\delta^4} \text{ for every } i = 1, \dots, n-1 \right) \\ & \Rightarrow \left(\tau_{p_2}^{B_{2r}(\Lambda)}(A_1, A_n) \geq c_2 \tau_{p_1}^\Lambda(A_1) \tau_{p_1}^\Lambda(A_n) \right) \end{aligned} \quad (3.4.1)$$

holds for every non-empty set of vertices Λ in G . In particular, taking $\Lambda = V$ yields the implication

$$\begin{aligned} & \left(h^{c_1\delta^3} \leq c_3 n^{-1} \text{ and } \tau_{p_1}(A_i \cup A_{i+1}) \geq 4h^{c_1\delta^4} \text{ for every } i = 1, \dots, n-1 \right) \\ & \Rightarrow \left(\tau_{p_2}(A_1, A_n) \geq c_2 \tau_{p_1}(A_1) \tau_{p_1}(A_n) \right) \end{aligned} \quad (3.4.2)$$

whenever $p_1 \geq 1/d$, $\delta = \delta(p_1, p_2) \leq D$, and $h \leq h_0$.

Remark 3.4.1. The fact that we can always find an integer r such that $\mathbb{P}_p(\text{Piv}[1, hr]) < h$ for every $p \in [p_1, p_2]$ can be deduced from the fact that there is at most one infinite cluster \mathbb{P}_p -almost surely for every $p \in [p_1, p_2]$ by an easy compactness argument.

Remark 3.4.2. The $\Lambda = V$ case of this lemma stated in (3.4.2) already allows us to easily deduce that $p_c(G_n) \rightarrow p_c(G)$ when the transitive graphs in the sequence $(G_n)_{n \geq 1}$ all satisfy a uniform superpolynomial growth lower bound of the form $\text{Gr}(r) \geq r^{c(\log \log r)^{10}}$. This is explained in detail in Section 3.4. Working within finite domains as in (3.4.1) will be useful when we apply Proposition 3.4.1 at low-growth scales in Section 3.6.

The basic idea underlying Proposition 3.4.1, which is inspired by earlier works including [CMT22; DKT21], is that one can use the universal two-arm estimates derived from the work of Aizenman, Kesten, and Newman [AKN87b] as reviewed in Section 3.2 to bound the maximum influence of an edge on certain connection events, which can then be used as an input in Talagrand's sharp threshold theorem [Tal94]. Compared to those works, our primary additional insight is that these methods can be made vastly more efficient (especially in the high-growth case) by working with *ghost field connection events* instead of more obvious connection events. Intuitively, these ghost field connection events are “smoother” than ordinary connection events, making it easier to bound the maximum influence of an edge. Moreover, the influence bound we get by using the two-ghost

inequality of [Hut20a] gets better as the size of the relevant sets increases, so that we get extremely strong sharp threshold estimates when the graph has high growth.

Given a set of vertices A in a graph G and a parameter $h \in [0, 1]$, the *ghost field* of intensity h on A is the random subset \mathcal{G}_A of A in which each vertex is included independently at random with probability⁸ h . We denote the law of \mathcal{G}_A by \mathbb{G}_h^A . We record the following reformulation of the two-ghost inequality of [Hut20a].

Lemma 3.4.2. *Let G be a unimodular transitive graph of vertex degree d . There exists a constant $C = C(d)$ such that*

$$\mathbb{G}_h^A \otimes \mathbb{G}_h^B \otimes \mathbb{P}_p(\{x \leftrightarrow \mathcal{G}_A\} \cap \{y \leftrightarrow \mathcal{G}_B\} \cap \{x \leftrightarrow y\} \cap \{x \leftrightarrow \infty \text{ or } y \leftrightarrow \infty\}) \leq C \sqrt{\frac{1-p}{p} h}$$

for every $h \in [0, 1]$, every pair of neighbouring vertices $x, y \in V(G)$, every two sets of vertices $A, B \subseteq V(G)$, and every $p \in (0, 1)$.

Proof. It suffices without loss of generality to consider the case $A = B = V$ since the relevant probability is increasing in A and B . In this case, the probability is the same as if we had a single ghost field instead of two independent ghost fields, since the restrictions of a ghost field to the clusters of x and y are independent when these clusters are disjoint. This version of the lemma then follows easily from Theorem 3.2.5. Alternatively, one can deduce the desired estimate from [Hut20a, Theorem 1.6]. (The only difference is that in that paper the ghost fields are parameterised by $1 - e^{-h}$ and are random sets of *edges* rather than vertices. As with Theorem 3.2.5, this is not a problem since the edge version of the statement is stronger than the vertex version.) \square

The following lemma states roughly that the probability that two low-intensity ghost fields are connected in a region Λ undergoes a sharp threshold with respect to the percolation parameter p . This rough statement has two caveats: as we increase the percolation parameter, we also have to increase the ghost field intensity and thicken the region Λ . Although the argument using ghost field connections is new, the way we adapt it to run inside a given domain Λ (when Λ is not the whole vertex set) is inspired by the analysis of [CMT22, Section 5].

Lemma 3.4.3 (Sharp threshold for ghost connections). *For each $d \geq 1$ and $D < \infty$ there exists a positive constant $c = c(d, D)$ such that the following holds. Let $G = (V, E)$ be a unimodular*

⁸In the literature one often takes this probability to be $1 - e^{-h}$, which makes certain calculations more convenient. The distinction makes little difference since $1 - e^{-h} = h \pm O(h^2)$ as $h \rightarrow 0$.

transitive graph with vertex degree d , let A, B be non-empty sets of vertices in G , and suppose that $0 < p_1 < p_2 < 1$ are such that either

- (i) there is at most one infinite cluster \mathbb{P}_p -almost surely for every $p \in [p_1, p_2]$, or
- (ii) B has finite complement.

If $p_1 \geq 1/d$ and $\delta(p_1, p_2) \leq D$ then the implication

$$\left(\mathbb{G}_h^A \otimes \mathbb{G}_h^B \otimes \mathbb{P}_{p_1}(\mathcal{G}_A \xleftrightarrow{\Lambda} \mathcal{G}_B) \geq h^{c\delta(p_1, p_2)} \right) \Rightarrow \left(\mathbb{G}_{h^c}^A \otimes \mathbb{G}_{h^c}^B \otimes \mathbb{P}_{p_2}(\mathcal{G}_A \xleftrightarrow{B_r(\Lambda)} \mathcal{G}_B) \geq 1 - h^{c\delta(p_1, p_2)} \right)$$

holds for every $h \leq 1/d$ and every set $\Lambda \subseteq V$, where r is the minimum positive integer such that $\mathbb{P}_p(\text{Piv}[1, hr]) < h$ for every $p \in [p_1, p_2]$.

Remark 3.4.3. To prove Proposition 3.4.1 we will apply this lemma only under the hypothesis (i). The version with hypothesis (ii) can be used as part of an alternative of the joint continuity of $\theta(p, G)$ in the supercritical region, as discussed in Section 3.7.

Before proving Lemma 3.4.3 we first recall *Talagrand's inequality* [Tal94], which (in combination with Russo's formula [Rus78]) states that there exists a universal positive constant c such that if $A \subseteq \{0, 1\}^E$ is an increasing event in a finite product space then

$$\frac{d}{dp} \mathbb{P}_p(\mathcal{A}) \geq c \mathbb{P}_p(\mathcal{A})(1 - \mathbb{P}_p(\mathcal{A})) \cdot \left[p(1-p) \log \frac{2}{p(1-p)} \right]^{-1} \log \frac{1}{p(1-p) \max_{e \in E} \mathbb{P}_p(\text{Piv}_e[\mathcal{A}])} \quad (3.4.3)$$

for every $p \in (0, 1)$.

Proof of Lemma 3.4.3. We may assume without loss of generality that A , B , and Λ are finite, exhausting by finite sets and taking a limit otherwise. Since we want to apply Talagrand's inequality to the *inhomogeneous*⁹ product measure $\mathbb{G}_h^A \otimes \mathbb{G}_h^B \otimes \mathbb{P}_p$, we will first encode a random variable with this law as a function of i.i.d. random bits. Define

$$m_E := \left\lfloor \frac{\log(1-p_1)}{\log((d-1)/d)} \right\rfloor, \quad q_1 := 1 - (1-p_1)^{1/m_E}, \quad \text{and} \quad q_2 := 1 - (1-p_2)^{1/m_E}.$$

⁹The fact that Talagrand's inequality for homogeneous product measures implies a version for inhomogeneous measures with parameters close to zero was already observed in [DT22, Appendix A]. In our setting we have some parameters that may be close to zero and others that may be close to 1.

The assumption that $p_1 \geq 1/d$ ensures that $m_E \geq 1$, while it follows from the definitions that $(1 - q_1)^{m_E} = (1 - p_1)$, that $(1 - q_2)^{m_E} = (1 - p_2)$, and that

$$\frac{1}{d} \leq q_1 = 1 - \exp \left(\left(\left\lfloor \frac{\log(1 - p_1)}{\log((d-1)/d)} \right\rfloor \right)^{-1} \log(1 - p_1) \right) \leq \frac{2d-1}{d^2} \leq \frac{2}{d}.$$

(Here we used only that $x/2 \leq \lfloor x \rfloor \leq x$ for $x \geq 1$.) We also define

$$m_G = \left\lfloor \frac{\log h}{\log q_1} \right\rfloor,$$

which satisfies $m_G \geq 1$ since $h \leq 1/d \leq q_1$. For each $q \in (0, 1)$, let $\bar{\mathbb{P}}_q$ be the law of a random variable

$$\text{BITS} = (\text{BITS}_A, \text{BITS}_B, \text{BITS}_\omega) \in \{0, 1\}^{A \times \{1, \dots, m_G\}} \times \{0, 1\}^{B \times \{1, \dots, m_G\}} \times \{0, 1\}^{E \times \{1, \dots, m_E\}} =: \{0, 1\}^\Omega,$$

whose constituent random bits are independent Bernoulli random variables of parameter q . Given $\text{BITS} = (\text{BITS}_A, \text{BITS}_B, \text{BITS}_\omega) \in \{0, 1\}^\Omega$, we define $(\mathcal{G}_A, \mathcal{G}_B, \omega)$ as a function of BITS by

$$\mathcal{G}_A(a) = \prod_{i=1}^{m_G} \text{BITS}_A(a, i), \quad \mathcal{G}_B(b) = \prod_{i=1}^{m_G} \text{BITS}_B(b, i), \quad \text{and} \quad \omega(e) = 1 - \prod_{i=1}^{m_E} (1 - \text{BITS}_\omega(e, i)),$$

so that the triple $(\mathcal{G}_A, \mathcal{G}_B, \omega)$ has law $\mathbb{G}_{q^{m_G}}^A \otimes \mathbb{G}_{q^{m_G}}^B \otimes \mathbb{P}_{1-(1-q)^{m_E}}$. The choice of parameter m_G ensures that $q_1^{m_G} \geq h$ and hence that

$$\bar{\mathbb{P}}_{q_1}(\mathcal{G}_A \xleftrightarrow{\Lambda} \mathcal{G}_B) \geq \mathbb{G}_h^A \otimes \mathbb{G}_h^B \otimes \mathbb{P}_{p_1}(\mathcal{G}_A \xleftrightarrow{\Lambda} \mathcal{G}_B).$$

Note moreover that if $\mathcal{A} \subseteq \{0, 1\}^A \times \{0, 1\}^B \times \{0, 1\}^E$ is an increasing event, $e \in E$ is an edge, and $1 \leq k \leq m_E$ then (e, k) is a closed pivotal for the event $\{\text{BITS} : (\mathcal{G}_A, \mathcal{G}_B, \omega) \in \mathcal{A}\}$ if and only if e is a closed pivotal for the event \mathcal{A} in the configuration ω , so that

$$\begin{aligned} (1 - q) \bar{\mathbb{P}}_q((e, k) \text{ pivotal for } \{\text{BITS} : (\mathcal{G}_A, \mathcal{G}_B, \omega) \in \mathcal{A}\}) \\ = (1 - q)^{m_E} \mathbb{G}_{q^{m_G}}^A \otimes \mathbb{G}_{q^{m_G}}^B \otimes \mathbb{P}_{1-(1-q)^{m_E}}(\text{Piv}_e[\mathcal{A}]). \end{aligned} \quad (3.4.4)$$

On the other hand, an element (x, k) of $A \times \{1, \dots, m_G\}$ can only possibly be an open pivotal if $\text{BITS}_A(x, j) = 1$ for every $j \in \{1, \dots, m_G\}$. Similar considerations also apply with A replaced by B , so that we obtain the coarse bound

$$q \bar{\mathbb{P}}_q((x, k) \text{ pivotal for } \{\text{BITS} : (\mathcal{G}_A, \mathcal{G}_B, \omega) \in \mathcal{A}\}) \leq q^{m_G} \quad (3.4.5)$$

for every $x \in A \sqcup B$ and $1 \leq k \leq m_G$.

Let $\ell := \lfloor h^{-1} \rfloor$ and for each $i = 1, \dots, \ell$ define the event

$$\mathcal{E}_i := \{\mathcal{G}_A \xleftrightarrow{B_{irh}(\Lambda)} \mathcal{G}_B\}.$$

We want to bound the maximum influence of an element of $(E \times \{1, \dots, m_E\}) \sqcup (A \times \{1, \dots, m_G\}) \sqcup (B \times \{1, \dots, m_G\})$ on the event \mathcal{E}_i . For elements of $(A \times \{1, \dots, m_G\}) \sqcup (B \times \{1, \dots, m_G\})$ it will suffice to use the trivial bound of (3.4.5), so that it remains only to bound the pivotality probability $\max_e \mathbb{G}_{h'}^A \otimes \mathbb{G}_{h'}^B \otimes \mathbb{P}_p(\text{Piv}_e[\mathcal{E}_i])$ where $(1-p) = (1-q)^{m_E}$ and $h' = q^{m_G}$. Following [CMT22], we will do this not for every i but instead show that an influence bound of the desired form must hold for an *average* choice of i . (Note that if we are working directly in the case $\Lambda = V$ then this issue does not arise.) More precisely, we claim that for each $q \in [q_1, q_2]$ there exists a set $I(q) \subseteq \{1, \dots, \ell\}$ with $|I| \geq 1/(3h)$ such that

$$\max_{i \in I} \max_{e \in E} q(1-q) \mathbb{G}_{h'}^A \otimes \mathbb{G}_{h'}^B \otimes \mathbb{P}_p(\text{Piv}_e[\mathcal{E}_i]) \leq C_d(h')^{1/2}, \quad (3.4.6)$$

where C_d is a constant depending only on the degree d . There are two separate cases to consider: Edges both of whose endpoints belong to $B_{(i-1)rh}(\Lambda)$ (*bulk edges*) and edges with at least one endpoint not in $B_{(i-1)rh}(\Lambda)$ (*boundary edges*).

Bulk edges. First consider an edge e both of whose endpoints x and y belong to $B_{(i-1)rh}(\Lambda)$. If ω is such that e is pivotal for the event \mathcal{E}_i then at least one of the following two events must occur:

- (i) The endpoints x and y of e belong to distinct ω -clusters, one of which intersects \mathcal{G}_A but not \mathcal{G}_B and the other of which intersects \mathcal{G}_B but not \mathcal{G}_A ; at least one of these clusters must be finite almost surely by the hypotheses of the lemma, which allows us to bound the relevant probability using the two-ghost inequality.
- (ii) The endpoints x and y of e are both ω -connected to the boundary of $B_{irh}(\Lambda)$ but are not ω -connected to each other within $B_{irh}(\Lambda)$, so that $\text{Piv}[1, rh](x)$ and $\text{Piv}[1, rh](y)$ both hold.

Thus, it follows by the two-ghost inequality as stated in Lemma 3.4.2 and the definition of r that

$$q(1-q) \bar{\mathbb{P}}_q(\text{Piv}_e[\mathcal{E}_i]) \leq q(1-q) \left[C_1 \frac{(1-p)^{1/2}}{p^{1/2}} (h')^{1/2} + h \right] \leq C_2 (h')^{\frac{1}{2}}, \quad (3.4.7)$$

for every $q \in [q_1, q_2]$ and $1 \leq i \leq \ell$, where C_1 and C_2 are constants depending only on d and we used that $p \geq 1/d$.

Boundary edges. An edge e not having both its endpoints in $B_{irh}(\Lambda)$ cannot possibly be pivotal for \mathcal{E}_i ; we need only bound $\bar{\mathbb{P}}_q(\text{Piv}_e[\mathcal{E}_i])$ for edges e with both endpoints in $B_{irh}(\Lambda)$ and with

at least one endpoint not in $B_{(i-1)rh}(\Lambda)$. Rather than bound this maximum influence uniformly in i , we will bound it *on average*. Fix $q \in [q_1, q_2]$ and, for each $i \in \{1, \dots, \ell\}$, pick an edge $e_i = e_i(q) \in B_{irh}(\Lambda)$ with at least one endpoint not in $B_{(i-1)rh}(\Lambda)$ that maximises $\bar{\mathbb{P}}_q(\text{Piv}_{e_i}[\mathcal{E}_i])$ over all such edges. Notice that the events $\text{Piv}_{e_i}[\mathcal{E}_i] \cap \{\omega(e_i) = 1\}$ for $i \in \{1, \dots, \ell\}$ are pairwise disjoint. Indeed, if $\text{Piv}_{e_i}[\mathcal{E}_i] \cap \{\omega(e_i) = 1\}$ occurs then e_i must belong to *every* path connecting \mathcal{G}_A to \mathcal{G}_B in $B_{irh}(\Lambda)$, and such a path cannot possibly include e_j if $j > i$. It follows in particular that

$$q \sum_{i=1}^{\ell} \bar{\mathbb{P}}_q(\text{Piv}_{e_i}[\mathcal{E}_i]) = \sum_{i=1}^{\ell} \bar{\mathbb{P}}_q(\text{Piv}_{e_i}[\mathcal{E}_i] \cap \{\omega(e_i) = 1\}) \leq 1.$$

By Markov's inequality, it follows that we can find a subset $I(q) \subseteq \{1, \dots, \ell\}$ with $|I| \geq \frac{1}{3h}$ such that

$$\max_{i \in I} q \bar{\mathbb{P}}_q(\text{Piv}_{e_i}[\mathcal{E}_i]) \leq 6h \leq 6(h')^{1/2}$$

as required. This concludes the proof of (3.4.6).

Now, the assumption that $\delta(p_1, p_2) \leq D$ implies that $(1 - p_2) \geq (1 - p_1)^{e^D}$ and hence that $1 - q_2 \geq (1 - q_1)^{e^D} \geq ((2d - 1)/(2d))^{e^D} > 0$, so that q_1 and q_2 are bounded away from 0 and 1 by constants depending only on d and D . As such, Talagrand's inequality (which is valid to use in our setting since all our events only depend on the status of finitely many bits) together with (3.4.6) yields that

$$\frac{d}{dq} \log \left[\frac{\bar{\mathbb{P}}_q(\mathcal{E}_i)}{1 - \bar{\mathbb{P}}_q(\mathcal{E}_i)} \right] \geq c_1 \log \frac{1}{C_d q^{m_G/2}} \geq c_2 m_G \log \frac{1}{q_1} \geq c_3 \log \frac{1}{h}$$

for every $q_1 \leq q \leq q_2$ and every $i \in I(q)$, where c_1, c_2 , and c_3 are positive constants depending only on d and D . Since $\ell \leq h^{-1}$ and the derivative is non-negative for every i , we can sum over i to obtain that

$$\frac{1}{\ell} \sum_{i=1}^{\ell} \frac{d}{dq} \log \left[\frac{\bar{\mathbb{P}}_q(\mathcal{E}_i)}{1 - \bar{\mathbb{P}}_q(\mathcal{E}_i)} \right] \geq \frac{c_3}{3} \log \frac{1}{h}$$

for every $q_1 \leq q \leq q_2$. Integrating this differential inequality yields that

$$\frac{1}{\ell} \sum_{i=1}^{\ell} \left(\log \left[\frac{\bar{\mathbb{P}}_{q_2}(\mathcal{E}_i)}{1 - \bar{\mathbb{P}}_{q_2}(\mathcal{E}_i)} \right] - \log \left[\frac{\bar{\mathbb{P}}_{q_1}(\mathcal{E}_i)}{1 - \bar{\mathbb{P}}_{q_1}(\mathcal{E}_i)} \right] \right) \geq \frac{c_3}{3} |q_2 - q_1| \log \frac{1}{h}$$

and hence that there exists $i \in \{1, \dots, \ell\}$ such that

$$\begin{aligned} \max \left\{ \log \left[\frac{1}{1 - \bar{\mathbb{P}}_{q_2}(\mathcal{E}_i)} \right], \log \left[\frac{1}{\bar{\mathbb{P}}_{q_1}(\mathcal{E}_i)} \right] \right\} &\geq \max \left\{ \log \left[\frac{\bar{\mathbb{P}}_{q_2}(\mathcal{E}_i)}{1 - \bar{\mathbb{P}}_{q_2}(\mathcal{E}_i)} \right], -\log \left[\frac{\bar{\mathbb{P}}_{q_1}(\mathcal{E}_i)}{1 - \bar{\mathbb{P}}_{q_1}(\mathcal{E}_i)} \right] \right\} \\ &\geq \frac{c_2}{6} |q_2 - q_1| \log \frac{1}{h}. \end{aligned}$$

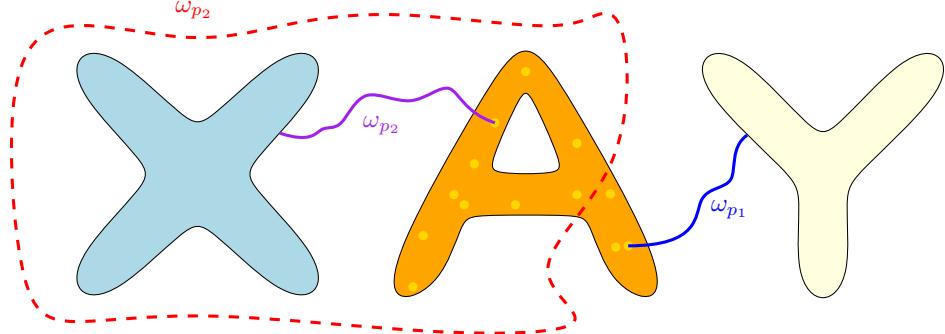


Figure 3.1: Schematic illustration of the event whose probability is estimated in Lemma 3.4.4 when $\Lambda = V$: If any two points in A have a reasonable probability to be connected in ω_{p_1} , then it is unlikely that X is connected to a weak ghost field on A in ω_{p_2} and Y is connected to a weak ghost field on A in ω_{p_1} but that X and Y are not connected in ω_{p_2} .

The claim follows easily from this together with the inequality

$$\begin{aligned} |q_2 - q_1| &= |(1 - p_2)^{1/m_E} - (1 - p_1)^{1/m_E}| \geq \left| (1 - p_2)^{\frac{\log((d-1)/d)}{\log(1-p_1)}} - (1 - p_1)^{\frac{\log((d-1)/d)}{\log(1-p_1)}} \right| \\ &= \frac{d-1}{d} \left[1 - \exp \left(\left(\frac{\log(1-p_2)}{\log(1-p_1)} - 1 \right) \log \frac{d-1}{d} \right) \right] \geq c_4 \delta(p_1, p_2), \end{aligned}$$

which holds by calculus with the constant c_4 depending only on d and D . \square

Our next goal is to run a chaining-like argument but with ghost fields. In this analogy, the previous lemma can be thought of as an *existence* statement whereas the next lemma can be thought of as a *uniqueness* statement. (Note that our proof of the next lemma relies crucially on the previous lemma.) We work with the standard monotone coupling $(\omega_p)_{p \in [0,1]}$ of Bernoulli bond percolation, write $\{A \xleftarrow[p]{\Lambda} B\}$ for the event that A is connected to B by a path that is contained in Λ and open in ω_p , and write $\{A \xleftarrow[p]{\Lambda} B\}$ for the event that A is *not* connected to B by any such path.

Lemma 3.4.4 (Gluing ghost connections). *For each $d \geq 1$ and $D < \infty$ there exist positive constants $c_1 = c_1(d, D)$ and $c_2 = c_2(d, D)$ such that the following holds. Let $G = (V, E)$ be a unimodular transitive graph with vertex degree d , let A , X , and Y be non-empty sets of vertices in G , and suppose that $0 < p_1 < p_2 < 1$ are such that there is at most one infinite cluster \mathbb{P}_p -almost surely for every $p \in [p_1, p_2]$. If $p_1 \geq 1/d$ and $\delta = \delta(p_1, p_2) \leq D$ then the implication*

$$\left[\tau_{p_1}^{\Lambda}(A) \geq h^{c_1 \delta} \right] \Rightarrow \left[\mathbb{G}_h^A \otimes \mathbb{P} \left(X \xleftarrow[p_2]{B_r(\Lambda)} \mathcal{G}_A \text{ and } Y \xleftarrow[p_1]{\Lambda} \mathcal{G}_A \text{ but } X \xleftarrow[p_2]{B_r(\Lambda)} Y \right) \leq 3h^{c_2 \delta^3} \right]$$

holds for every $h \leq 1/d$ and every set $\Lambda \subseteq V$, where r is the minimum positive integer such that $\mathbb{P}_p(\text{Piv}[1, hr]) < h$ for every $p \in [p_1, p_2]$.

An illustration of the event whose probability this lemma estimates is given in Figure 3.1.

Remark 3.4.4. Several of the calculations in the following proof can be simplified significantly if one allows the constants to depend on how small $1 - p_1$ is. Getting the constants to be independent of the choice of p_1 is important when using our methods to deduce that $p_c < 1$ for transitive graphs of superlinear volume growth.

Proof of Lemma 3.4.4. It suffices to prove the claim with $3h^{c_2\delta^4}$ replaced by $\max\{h^{c_1\delta}, 3h^{c_2\delta^4}\}$, as the latter can be bounded by the former after an appropriate decrease of the relevant constant. We write $p_{4/3}$ and $p_{5/3}$ for the parameters defined by $p_{4/3} = \text{Spr}(p_1; \delta/3)$ and $p_{5/3} = \text{Spr}(p_1; 2\delta/3)$, so that $p_1 \leq p_{4/3} \leq p_{5/3} \leq p_2$. To lighten notation we write $\omega_1 = \omega_{p_1}$, $\omega_{4/3} = \omega_{p_{4/3}}$, $\omega_{5/3} = \omega_{p_{5/3}}$, and $\omega_2 = \omega_{p_2}$. We will write \asymp , \preceq , and \succeq for equalities and inequalities holding to within positive multiplicative constants depending only on d and D . We also use the asymptotic big- O and big- Ω notation with all implicit constants depending only on d and D .

Let $c_0 = c_0(d, D)$ be the constant from Lemma 3.4.3. We will prove the claim with $c_1 = (c_0 \wedge 1)/9$. Assume to this end that $\tau_{p_1}^\Lambda(A) \geq h^{c_1\delta}$. Let $K_X(2, r)$ be the set of vertices that are connected to X by an ω_2 -open path in $B_r(\Lambda)$ and let $K_Y(1, 0)$ be the set of vertices that are connected to Y by an ω_1 -open path in Λ . Suppose that there exist two disjoint, connected sets of vertices $C_X \supseteq X$ and $C_Y \supseteq Y$ such that $\mathbb{P}(K_X(2, r) = C_X, K_Y(1, 0) = C_Y) > 0$ and

$$\begin{aligned} & \min\{\mathbb{G}_h^A(\mathcal{G}_A \cap C_X \neq \emptyset), \mathbb{G}_h^A(\mathcal{G}_A \cap C_Y \neq \emptyset)\} \\ & \geq \mathbb{G}_h^A \otimes \mathbb{P}\left(X \xleftrightarrow[p_2]{B_r(\Lambda)} \mathcal{G}_A \xleftrightarrow[p_1]{\Lambda} Y \mid K_X(2, r) = C_X, K_Y(1, 0) = C_Y\right) \geq h^{c_1\delta}; \end{aligned} \quad (3.4.8)$$

If no such pair of sets exists then the conclusion holds trivially (since we are proving a modified version of the claim with $\max\{h^{c_1\delta}, 3h^{c_2\delta^4}\}$ on the right hand side of the implication). For such C_X and C_Y we have that

$$\begin{aligned} \mathbb{G}_h^{C_X} \otimes \mathbb{G}_h^{C_Y} \otimes \mathbb{P}\left(\mathcal{G}_{C_X} \xleftrightarrow[p_1]{\Lambda} \mathcal{G}_{C_Y}\right) & \geq \mathbb{G}_h^{C_X} \otimes \mathbb{G}_h^{C_Y} (\mathcal{G}_{C_X} \cap A \neq \emptyset \text{ and } \mathcal{G}_{C_Y} \cap A \neq \emptyset) \min_{u, v \in A} \mathbb{P}(u \xleftrightarrow[p_1]{\Lambda} v) \\ & = \mathbb{G}_h^A(\mathcal{G}_A \cap C_X \neq \emptyset) \mathbb{G}_h^A(\mathcal{G}_A \cap C_Y \neq \emptyset) \min_{u, v \in A} \mathbb{P}(u \xleftrightarrow[p_1]{\Lambda} v) \\ & \geq h^{c_1\delta} \cdot h^{c_1\delta} \cdot h^{c_1\delta} = h^{3c_1\delta} \geq h^{\frac{c_0}{3}\delta}. \end{aligned}$$

Thus, since $\delta(p_1, p_{4/3}) = \frac{1}{3}\delta(p_1, p_2)$, it follows from Lemma 3.4.3 (applied with p_1 and $p_{4/3}$ rather than p_1 and p_2) that

$$\mathbb{P}\left(C_X \xleftarrow[p_{4/3}]{B_r(\Lambda)} C_Y\right) \leq \mathbb{G}_{h^{c_0}}^{C_X} \otimes \mathbb{G}_{h^{c_0}}^{C_Y} \otimes \mathbb{P}\left(\mathcal{G}_{C_X} \xleftarrow[p_{4/3}]{B_r(\Lambda)} \mathcal{G}_{C_Y}\right) \leq h^{\frac{c_0}{3}\delta}. \quad (3.4.9)$$

Let M denote the maximum cardinality of a set of edge-disjoint paths from C_X to C_Y that are contained in $B_r(\Lambda)$ and are open in $\omega_{5/3}$ with the possible exception of their first and last edges. (We stress that C_X and C_Y are not random, but are fixed sets satisfying (3.4.8).) By Menger's theorem, M is equal to the minimum size of a set of edges separating C_X from C_Y in the subgraph of $B_r(\Lambda)$ spanned by those edges that are either $\omega_{5/3}$ -open or have at least one endpoint in $C_X \cup C_Y$. Observe that M is independent of the event $\{K_X(2, r) = C_X, K_Y(1, 0) = C_Y\}$ since M depends only on edges with neither endpoint in $C_X \cup C_Y$ while the latter event depends only on edges with at least one endpoint in $C_X \cup C_Y$. Conditional on $\omega_{5/3}$, each edge that is open in $\omega_{5/3}$ is closed in $\omega_{4/3}$ with probability

$$\beta := \frac{p_{5/3} - p_{4/3}}{p_{5/3}}.$$

It follows that

$$\mathbb{P}\left(C_X \xleftarrow[p_{4/3}]{B_r(\Lambda)} C_Y \mid M \leq N\right) \geq \beta^N$$

for every $N \geq 0$ and hence by (3.4.9) and the aforementioned independence that

$$\mathbb{P}(M \leq N \mid K_X(2, r) = C_X, K_Y(1, 0) = C_Y) = \mathbb{P}(M \leq N) \leq \beta^{-N} h^{\frac{c_0}{3}\delta} \quad (3.4.10)$$

for every $N \geq 0$. Let $K_Y(5/3, r)$ be the set of vertices that are connected to Y by an $\omega_{5/3}$ -open path in $B_r(\Lambda)$. Conditioned on the event $\{K_X(2, r) = C_X \text{ and } K_Y(1, 0) = C_Y\}$ and the value of M , the size of the boundary $|\partial K_X(2, r) \cap \partial K_Y(5/3, r)|$ stochastically dominates a sum of M Bernoulli random variables of parameter

$$\alpha_1 := \frac{p_{5/3} - p_1}{1 - p_1}.$$

Indeed, if we take some maximal-cardinality set of paths as in the definition of M , then the edge adjacent to C_Y of each path in the set is open in $\omega_{5/3}$ with this probability, and the size of the relevant boundary is at least the number of these edges that are open in $\omega_{5/3}$. Letting Z be a sum of N Bernoulli random variables each of parameter $1 - \alpha_1$, we have that there exist constants c_3 and

c_4 depending only on d and D such that

$$\begin{aligned} \mathbb{P}\left(|\partial K_X(2, r) \cap \partial K_Y(5/3, r)| \leq \frac{\alpha_1}{2}N \mid K_X(2, r) = C_X, K_Y(1, 0) = C_Y, M \geq N\right) \\ \leq \mathbb{P}\left(Z \geq \frac{2 - \alpha_1}{2}N\right) = \mathbb{P}\left(Z \geq \left(1 + \frac{\alpha_1}{2(1 - \alpha_1)}\right)\mathbb{E}Z\right) \\ \leq \exp\left[-\left(\frac{2 - \alpha_1}{2} \log \frac{2 - \alpha_1}{2 - 2\alpha_1} - \frac{\alpha_1}{2}\right)N\right] \quad (3.4.11) \end{aligned}$$

for every $N \geq 1$, where the final inequality follows from the Chernoff bound $\mathbb{P}(Z \geq (1+x)\mathbb{E}Z) \leq (e^{-x}(1+x)^{1+x})^{-\mathbb{E}Z}$, which holds for all sums of independent Bernoulli random variables.

We will now apply (3.4.10) and (3.4.11) with

$$N := \left\lfloor \frac{c_0\delta}{12\log(1/\beta)} \log \frac{1}{h} \right\rfloor, \quad \text{so that} \quad \beta^{-N}h^{\frac{c_0}{3}\delta} \leq h^{\frac{c_0}{6}\delta}$$

to prove the inequality

$$\begin{aligned} \mathbb{P}\left(|\partial K_X(2, r) \cap \partial K_Y(5/3, r)| \leq \frac{c_0\alpha_1\delta}{48\log(1/\beta)} \log \frac{1}{h} \mid K_X(2, r) = C_X, K_Y(1, 0) = C_Y\right) \\ \leq 2h^{\Omega(\delta^4)}. \quad (3.4.12) \end{aligned}$$

To do this, we will make repeated use of the elementary estimates

$$\begin{aligned} p_{5/3}\beta &= (1 - p_{5/3})^{e^{-\delta/3}} - (1 - p_{5/3}) = (1 - p_{5/3})((1 - p_{5/3})^{e^{-\delta/3}-1} - 1) \\ &\geq (1 - p_{5/3})((1 - p_{5/3})^{-\Omega(\delta)} - 1) \geq \delta(1 - p_{5/3}) \log \frac{1}{1 - p_{5/3}} \end{aligned} \quad (3.4.13)$$

and

$$\alpha_1 = 1 - (1 - p_{5/3})^{1-e^{-2\delta/3}} = 1 - (1 - p_{5/3})^{\Theta(\delta)}, \quad (3.4.14)$$

where we recall that the implicit constant appearing here depends only on d and D . First suppose that $N \geq 1$, so that the rounding in the definition of N reduces its size by a factor of at most $1/2$. In this case, (3.4.10) and (3.4.11) yield that there exist positive constants C_1, c_5 and c_6 depending only on d and D such that

$$\begin{aligned} \mathbb{P}\left(|\partial K_X(2, r) \cap \partial K_Y(5/3, r)| \leq \frac{c_0\alpha_1\delta}{48\log(1/\beta)} \log \frac{1}{h} \mid K_X(2, r) = C_X, K_Y(1, 0) = C_Y\right) \\ \leq h^{\frac{c_0}{6}\delta} + \exp\left[-\frac{c_0\delta}{24\log(1/\beta)} \left(\frac{2 - \alpha_1}{2} \log \frac{2 - \alpha_1}{2 - 2\alpha_1} - \frac{\alpha_1}{2}\right) \log \frac{1}{h}\right]. \quad (3.4.15) \end{aligned}$$

We claim that

$$\frac{c_0\delta}{24\log(1/\beta)} \left(\frac{2-\alpha_1}{2} \log \frac{2-\alpha_1}{2-2\alpha_1} - \frac{\alpha_1}{2} \right) \succeq \delta^4. \quad (3.4.16)$$

We prove this by a case analysis according to whether $1 - p_{5/3} \leq e^{-1/\delta}$. If $1 - p_{5/3} \geq e^{-1/\delta}$, it follows from (3.4.13) that $\log 1/\beta = O(\delta)$ (where we stress again that the implicit constants depend only on d and D). On the other hand, since $p_1 \geq 1/d$, we have that $\alpha_1 \succeq \delta$, and together with the elementary inequality

$$\frac{2-\alpha_1}{2} \log \frac{2-\alpha_1}{2-2\alpha_1} - \frac{\alpha_1}{2} \geq \frac{\alpha_1^2}{8}$$

this yields that if $1 - p_{5/3} \geq e^{-1/\delta}$ then

$$\frac{c_0\delta}{24\log(1/\beta)} \left(\frac{2-\alpha_1}{2} \log \frac{2-\alpha_1}{2-2\alpha_1} - \frac{\alpha_1}{2} \right) \succeq \frac{\delta}{\delta} \cdot \delta^2 \succeq \delta^3$$

as claimed. On the other hand, if $1 - p_{5/3} \leq e^{-1/\delta}$ then $1 - \alpha_1 \geq 1 - e^{-\Omega(1)} \succeq 1$, so that

$$\frac{2-\alpha_1}{2} \log \frac{2-\alpha_1}{2-2\alpha_1} - \frac{\alpha_1}{2} \succeq \log \frac{1}{1-\alpha_1} \succeq \delta \log \frac{1}{1-p_{5/3}}.$$

We have under the same assumption that

$$\log(1/\beta) \preceq \log \frac{1}{1-p_{5/3}},$$

and it follows that if $1 - p_{5/3} \leq e^{-1/\delta}$ then

$$\frac{c_0\delta}{24\log(1/\beta)} \left(\frac{2-\alpha_1}{2} \log \frac{2-\alpha_1}{2-2\alpha_1} - \frac{\alpha_1}{2} \right) \succeq \delta^2 \succeq \delta^3.$$

This completes the proof of (3.4.16), which together with (3.4.15) yields the claimed inequality (3.4.12) in the case that $N \geq 1$. Now suppose that $N = 0$, so that

$$\begin{aligned} & \mathbb{P} \left(\left| \partial K_X(2, r) \cap \partial K_Y(5/3, r) \right| \leq \frac{c_0\alpha_1\delta}{48\log(1/\beta)} \log \frac{1}{h} \middle| K_X(2, r) = C_X, K_Y(1, 0) = C_Y \right) \\ &= \mathbb{P} \left(\left| \partial K_X(2, r) \cap \partial K_Y(5/3, r) \right| = 0 \middle| K_X(2, r) = C_X, K_Y(1, 0) = C_Y \right) \leq h^{\frac{c_0}{3}\delta} + (1 - \alpha_1), \end{aligned} \quad (3.4.17)$$

where, as in (3.4.11), the first term on the right hand side of the second line bounds the probability that $M = 0$ and the second bounds the probability of the appropriate event conditional on $M \geq 1$. We will once again prove (3.4.12) via case analysis according to whether or not $1 - p_{5/3} \geq e^{-1/\delta}$. If $1 - p_{5/3} \geq e^{-1/\delta}$ then $\log(1/\beta) = O(\delta)$, and since $N = 0$ it follows that $h = e^{-O(\delta^{-2})}$. As such,

in this case there exists a positive constant c_5 such that $h^{c_5\delta^4} \geq 1/2$, and the desired inequality (3.4.12) follows trivially since probabilities are bounded by 1. Otherwise, $1 - p_{5/3} \leq e^{-1/\delta}$ and $1 - \alpha_1 \leq (1 - p_{5/3})^{\Omega(\delta)}$. Since $N = 0$, we also have that $h \geq (1 - p_{5/3})^{O(\delta^{-1})}$ and hence that $h^{\delta^2} = (1 - \alpha_1)^{\Omega(1)}$. As such, (3.4.17) is stronger than the claimed inequality (3.4.12) when $N = 0$ regardless of the value of $1 - p_{5/3}$. This completes the proof of the inequality (3.4.12).

Since the sets C_X and C_Y were arbitrary sets such that $\mathbb{P}(K_X(2, r) = C_X, K_Y(1, 0) = C_Y) > 0$ and that satisfy (3.4.8), it follows from (3.4.12) that

$$\begin{aligned} \mathbb{G}_h^A \otimes \mathbb{P} \left(X \xleftarrow[p_2]{B_r(\Lambda)} \mathcal{G}_A \xleftarrow[\Lambda]{p_1} Y \text{ but } X \xleftarrow[p_2]{B_r(\Lambda)} Y \right) &\leq h^{c_1\delta} + 2h^{\Omega(\delta^4)} \\ &+ \mathbb{P} \left(|\partial K_X(2, r) \cap \partial K_Y(5/3, r)| \geq \frac{c_0\alpha_1\delta}{48\log(1/\beta)} \log \frac{1}{h} \text{ but } X \xleftarrow[p_2]{B_r(\Lambda)} Y \right), \end{aligned} \quad (3.4.18)$$

where the first term $h^{c_1\delta}$ accounts for the possibility that $K_X(2, r)$ and $K_Y(1, 0)$ are equal to some sets C_X and C_Y that do not satisfy (3.4.8). It remains to bound the final term appearing on the right hand side of (3.4.18). Notice that we can replace the set $K_X(2, r)$ appearing in (3.4.18) by the set \tilde{K}_X of vertices that are connected to X by an ω_2 -open path using only edges that have both endpoints in $B_r(\Lambda)$ and neither endpoint in $K_Y(5/3, r)$, since the two sets $K_X(2, r)$ and \tilde{K}_X are equal when $X \xleftarrow[p_2]{B_r(\Lambda)} Y$. Now let $C_X \supseteq X$ and $C_Y \supseteq Y$ be connected sets of vertices that are disjoint from each other such that $\mathbb{P}(\tilde{K}_X = C_X, K_Y(5/3, r) = C_Y) > 0$ and

$$|\partial C_X \cap \partial C_Y| \geq \frac{c_0\alpha_1\delta}{48\log(1/\beta)} \log \frac{1}{h};$$

If no such sets exist then the last term on the right hand side of (3.4.18) is zero. Each edge that is closed in $\omega_{5/3}$ is open in ω_2 with probability

$$\alpha_2 = \frac{p_2 - p_{5/3}}{1 - p_{5/3}},$$

and we have by independence that

$$\begin{aligned} \mathbb{P} \left(X \xleftarrow[p_2]{B_r(\Lambda)} Y \mid \tilde{K}_X = C_X, K_Y(5/3, r) = C_Y \right) &\leq (1 - \alpha_2)^{|\partial C_X \cap \partial C_Y|} \\ &\leq \exp \left[- \left(\frac{c_0\alpha_1\delta}{48\log(1/\beta)} \log \frac{1}{1 - \alpha_2} \right) \log \frac{1}{h} \right] \leq h^{\Omega(\delta^3)}, \end{aligned}$$

where the final inequality follows by a similar calculation used to prove (3.4.16) above. Thus, summing over all choices of C_X and C_Y , we obtain that

$$\mathbb{P} \left(X \xleftarrow[p_2]{B_r(\Lambda)} Y \mid |\partial \tilde{K}_X \cap \partial K_Y(5/3, r)| \geq \frac{c_0\alpha_1\delta}{48\log(1/\beta)} \log \frac{1}{h} \right) \leq h^{\Omega(\delta^3)}, \quad (3.4.19)$$

and the claim follows from (3.4.18) and (3.4.19). \square

We now apply Lemmas 3.4.3 and 3.4.4 to prove Proposition 3.4.1.

Proof of Proposition 3.4.1. To lighten notation, for each $\rho \in [0, 1]$ define $\mathbb{Q}_\rho := \bigotimes_{i=1}^n \mathbb{G}_\rho^{A_i}$, and for each i write $\mathcal{G}_i := \mathcal{G}_{A_i}$. Let $c_1 = c_1(d, D)$ be the constant from Lemma 3.4.3 and let $c_2 = c_2(d, D)$ and $c_3 = c_3(d, D)$ be the constants from Lemma 3.4.4 (that are called c_1 and c_2 in the statement of that lemma). Define $c = c(d, D) = (2D)^{-5}(c_1 \wedge 1)(c_2 \wedge 1)(c_3 \wedge 1)$. Let $h \geq (\min_i |A_i|)^{-1}$, noting that $\mathbb{G}_h^{A_i}(\mathcal{G}_i \neq \emptyset) \geq 1 - (1 - (\min_i |A_i|)^{-1})^{|A_i|} \geq 1 - 1/e \geq 1/2$ for every i , and let r be the minimum positive integer such that $\mathbb{P}_p(\text{Piv}[1, hr]) < h$ for every $p \in [p_1, p_2]$.

Fix $u \in A_1$ and $v \in A_n$ and suppose that

$$\tau_{p_1}^\Lambda(A_i \cup A_{i+1}) \geq 4h^{c\delta^4} \quad \text{for all } 1 \leq i \leq n-1.$$

We want to bound from below the probability under \mathbb{P}_{p_2} that u and v are connected inside $B_{2r}(\Lambda)$. Let $1 \leq i \leq n-1$ be arbitrary. Note that

$$\mathbb{Q}_h \otimes \mathbb{P} \left(\mathcal{G}_i \xrightarrow[p_1]{\Lambda} \mathcal{G}_{i+1} \right) \geq \mathbb{Q}_h(\mathcal{G}_i \neq \emptyset) \mathbb{Q}_h(\mathcal{G}_{i+1} \neq \emptyset) \tau_{p_1}^\Lambda(A_i, A_{i+1}) \geq 4h^{c\delta^4} \cdot \left(\frac{1}{2}\right)^2 \geq h^{c_1\delta/2}, \quad (3.4.20)$$

and similarly that

$$\mathbb{Q}_h \otimes \mathbb{P} \left(u \xrightarrow[p_1]{\Lambda} \mathcal{G}_1 \right) \geq \frac{1}{2} \tau_{p_1}^\Lambda(A_1) \quad \text{and} \quad \mathbb{Q}_h \otimes \mathbb{P} \left(\mathcal{G}_n \xrightarrow[p_1]{\Lambda} v \right) \geq \frac{1}{2} \tau_{p_1}^\Lambda(A_n). \quad (3.4.21)$$

Let $p_{3/2} = \text{Spr}(p_1; \delta/2)$, so that $\delta(p_1, p_{3/2}) = \frac{1}{2}\delta$. Using (3.4.20), Lemma 3.4.3 implies that

$$\mathbb{Q}_{h^c} \otimes \mathbb{P}(\mathcal{G}_i \xrightarrow[p_{3/2}]{B_r(\Lambda)} \mathcal{G}_{i+1}) \geq 1 - h^{c_1\delta/2} \geq 1 - h^{c\delta}. \quad (3.4.22)$$

Since we also have by choice of c that

$$\tau_{p_{3/2}}^{B_r(\Lambda)}(A_i) \geq \tau_{p_1}^\Lambda(A_i) \geq \tau_{p_1}^\Lambda(A_i \cup A_{i+1}) \geq 4h^{c\delta^4} \geq (h^{c_1})^{c_2\delta/2},$$

we may apply Lemma 3.4.4 (with $[p_1, p_2, h, X, Y, A, \Lambda] := [p_{3/2}, p_2, h^{c_1}, \{u\}, \mathcal{G}_{i+1}, A_i, B_r(\Lambda)]$) to deduce that if $h^{c_1} \leq 1/d$ then

$$\mathbb{Q}_{h^c} \otimes \mathbb{P} \left(u \xrightarrow[p_2]{B_{2r}(\Lambda)} \mathcal{G}_i \xrightarrow[p_{3/2}]{B_r(\Lambda)} \mathcal{G}_{i+1} \text{ but } u \xrightarrow[p_2]{B_{2r}(\Lambda)} \mathcal{G}_{i+1} \right) \leq 4(h^{c_1})^{c_3(\delta/2)^3} = 4h^{c_4\delta^3}, \quad (3.4.23)$$

where $c_4 = c_4(d, D) = (c_1 \cdot c_3)/8$. (Note that using h^{c_1} instead of h changes the hypotheses on r , but this is not a problem since the hypotheses on r associated to h^{c_1} are weaker than

those associated to h .) A similar application of Lemma 3.4.4 (with $[p_1, p_2, h, X, Y, A, \Lambda] := [p_{3/2}, p_2, h^{c_1}, \{u\}, \{v\}, A_n, B_r(\Lambda)]$) yields that if $h^{c_1} \leq 1/d$ then

$$\mathbb{Q}_{h^{c_1}} \otimes \mathbb{P} \left(u \xleftarrow[p_2]{B_{2r}(\Lambda)} \mathcal{G}_n \xleftarrow[p_{3/2}]{B_r(\Lambda)} v \text{ but } u \xleftarrow[p_2]{B_{2r}(\Lambda)} v \right) \leq 4h^{c_4\delta^3}. \quad (3.4.24)$$

Now, we have by a union bound that

$$\begin{aligned} \mathbb{P} \left(u \xleftarrow[p_2]{B_{2r}(\Lambda)} v \right) &\geq \mathbb{Q}_{h^{c_1}} \otimes \mathbb{P} \left(u \xleftarrow[p_{3/2}]{B_r(\Lambda)} \mathcal{G}_1 \xleftarrow[p_{3/2}]{B_r(\Lambda)} \dots \xleftarrow[p_{3/2}]{B_r(\Lambda)} \mathcal{G}_n \xleftarrow[p_{3/2}]{B_r(\Lambda)} v \right) \\ &\quad - \sum_{i=1}^{n-1} \mathbb{Q}_{h^{c_1}} \otimes \mathbb{P} \left(u \xleftarrow[p_2]{B_{2r}(\Lambda)} \mathcal{G}_i \xleftarrow[p_{3/2}]{B_r(\Lambda)} \mathcal{G}_{i+1} \text{ but } u \xleftarrow[p_2]{B_{2r}(\Lambda)} \mathcal{G}_{i+1} \right) \\ &\quad - \mathbb{Q}_{h^{c_1}} \otimes \mathbb{P} \left(u \xleftarrow[p_2]{B_{2r}(\Lambda)} \mathcal{G}_n \xleftarrow[p_{3/2}]{B_r(\Lambda)} v \text{ but } u \xleftarrow[p_2]{B_{2r}(\Lambda)} v \right). \end{aligned} \quad (3.4.25)$$

Using (3.4.21), (3.4.22), and the Harris-FKG inequality to bound the first term and (3.4.23) and (3.4.24) to control the error terms, we obtain that there exists a universal positive constant c_5 and positive constants c_6 and C depending only on d and D such that if $h^c \leq 1/d$ then

$$\begin{aligned} \mathbb{P} \left(u \xleftarrow[p_2]{B_{2r}(\Lambda)} v \right) &\geq \frac{\tau_{p_1}^\Lambda(A_1)}{2} \left(1 - h^{c\delta}\right)^{n-1} \frac{\tau_{p_1}^\Lambda(A_n)}{2} - n \cdot 4h^{c_4\delta^3} \\ &\geq c_5 \left[1 - Cnh^{c_6\delta^3}\right] \tau_p^\Lambda(A_1) \tau_p^\Lambda(A_n), \end{aligned}$$

where we used that $\tau_{p_1}^\Lambda(A_1) \tau_{p_1}^\Lambda(A_n) \geq h^{2c\delta^4} \geq h^{c_1 c_3 \delta^3/2^4}$ to absorb the error term into the prefactor. The proposition follows easily since the vertices $u \in A_1$ and $v \in A_n$ were arbitrary. \square

Graphs of high growth and the implication (F)

We now apply Proposition 3.4.1 to prove the implication (F) of the main induction step Proposition 3.3.1. We state the implication without including the burn-in term in δ_0 ; this will not cause problems since the statement is stronger without this term than with it.

Proposition 3.4.5 (The implication (F)). *For each $d \in \mathbb{N}$ there exists a constant $N = N(d) \geq 16$ such that the following holds. Let G be an infinite, connected, unimodular transitive graph with vertex degree d , let $p_0 \in (0, 1)$, and let $n_0 \geq 16$. Let $\delta_0 = (\log \log n_0)^{-1/2}$, define sequences $(n_i)_{i \geq 1}$ and $(\delta_i)_{i \geq 1}$ recursively by*

$$n_i := \exp^{\circ 3} \left(\log^{\circ 3}(n_0) + i \log 9 \right) \quad \text{and} \quad \delta_i := (\log \log n_i)^{-1/2}$$

and let $(p_i)_{i \geq 1}$ be an increasing sequence of probabilities satisfying $p_{i+1} \geq \text{Spr}(p_i; \delta_i)$ for each $i \geq 0$. Let $n_{-1} := (\log n_0)^{1/2}$. For each $i \geq 0$ define the statement

$$\text{FULL-SPACE}(i) = \left(\mathbb{P}_{p_i}(u \leftrightarrow v) \geq \exp[-(\log \log n_i)^{1/2}] \text{ for all } u, v \in B_{n_i} \right)$$

and for each $i \geq 1$ define the statement

$$\text{CORRIDOR}(i) = \left(\kappa_{p_i}(e^{[\log m]^{10}}, m) \geq \exp[-(\log \log n_i)^{1/2}] \text{ for every } m \in \mathcal{L}(G, 20) \cap [n_{i-2}, n_{i-1}] \right).$$

If $n_0 \geq N$ and $p_0 \geq 1/d$ then the implication

$$\left[\text{FULL-SPACE}(i) \wedge \text{CORRIDOR}(i+1) \right] \implies \left[\text{FULL-SPACE}(i+1) \vee (p_{i+1} \geq p_c) \right] \quad (\text{F})$$

holds for every $i \geq 0$.

Proof of Proposition 3.4.5. Fix $i \geq 0$ and suppose that $\text{FULL-SPACE}(i)$ and $\text{CORRIDOR}(i+1)$ both hold. If $n_i \in \mathcal{L}(G, 20)$, then $e^{(\log n_i)^{10}} \geq 2n_{i+1} = 2e^{(\log n_i)^9}$ whenever N is larger than some universal constant, so that if $u, v \in B_{n_{i+1}}$ then $d(u, v) \leq 2n_{i+1} \leq e^{(\log n_i)^{10}}$ and

$$\mathbb{P}_{p_{i+1}}(u \leftrightarrow v) \geq \kappa_{p_{i+1}}\left(e^{[\log n_i]^{10}}, n_i\right) \geq e^{-(\log \log n_{i+1})^{1/2}} \quad \text{for every } u, v \in B_{n_{i+1}}.$$

That is, $\text{CORRIDOR}(i+1)$ trivially implies $\text{FULL-SPACE}(i+1)$ whenever $n_i \in \mathcal{L}(G, 20)$. Now suppose that $n_i \notin \mathcal{L}(G, 20)$ and suppose that $\text{FULL-SPACE}(i)$ holds, so that

$$\text{Gr}(n_i) \geq \exp((\log(n_i^{1/3}))^{20}) =: h \quad \text{and} \quad \tau_{p_i}(B_n(u_i)) \geq \exp[-(\log \log n_i)^{1/2}].$$

Fix two arbitrary vertices $u, v \in B_{n_{i+1}}$ and let $u = u_1, u_2, \dots, u_k = v$ be the vertices in a geodesic from u to v . It follows from the Harris-FKG inequality that for all j ,

$$\tau_{p_i}(B_n(u_j) \cup B_n(u_{j+1})) \geq \tau_{p_i}(B_{n+1}(u_j)) \geq p_i^2 \tau_{p_i}(B_{n+1}(u_j)) \geq d^{-2} \exp[-(\log \log n_j)^{1/2}].$$

Let c_1, c_2, c_3 and $h_0 = h_0(d)$ be the constants from Proposition 3.4.1 applied with $D = 1$ (so that c_1, c_2 , and c_3 are universal) and let $N_1 = N_1(d)$ be sufficiently large that $\exp(-(\log n)^{20}) \leq h_0$ for every $n \geq N_1$. There exists a constant $N_2 = N_2(d) \geq N_1$ such that if $n_i \geq n_0 \geq N_2$ then

$$h^{c_1 \delta_i^3} = \exp\left[-\frac{c_1}{3^{20}} \frac{(\log n_i)^{20}}{(\log \log n_i)^{3/2}}\right] \leq c_3 n_{i+1}^{-1}$$

and for all j ,

$$\tau_{p_i}(B_n(u_j) \cup B_n(u_{j+1})) \geq d^{-2} \exp[-(\log \log n_i)^{1/2}] \geq 4 \exp\left[-\frac{c_1}{3^{20}} \frac{(\log n_i)^{20}}{(\log \log n_i)^2}\right] = 4h^{c_1 \delta_i^4}.$$

Thus, if $n_0 \geq N_2$ then Proposition 3.4.1 (applied with $D = 1$ and $A_i = B_n(u_i)$) implies that for all j ,

$$\mathbb{P}_{p_{i+1}}(u \leftrightarrow v) \geq \tau_{p_{i+1}}(B_{n_i}(u), B_{n_i}(v)) \geq c_2 \tau_{p_i}(B_{n_i}(u)) \tau_{p_i}(B_{n_i}(v)) \geq c_2 \exp[-2(\log \log n_i)^{1/2}],$$

where the final inequality follows from the assumption that **FULL-SPACE**(i) holds. Since u and v were arbitrary vertices in $B_{n_{i+1}}$, it follows that there exists a constant $N_3 = N_3(d) \geq N_2$ such that if $n_i \geq n_0 \geq N_3$ then **FULL-SPACE**($i+1$) holds as claimed. \square

As mentioned in the introduction, Proposition 3.4.5 already allows us to conclude the proofs of Theorems 3.3.3 and 8.1.1 under a mild uniform superpolynomial growth assumption.

Corollary 3.4.6. *Let G be an infinite, connected, unimodular transitive graph. If $\log \text{Gr}(r) > (\log r)^{20}$ for all sufficiently large r then $p_c(G) < 1$.*

Corollary 3.4.7. *Let $(G_n)_{n \geq 1}$ be a sequence of infinite, connected, unimodular transitive graphs converging to some transitive graph G , and suppose that there exists R such that $\log \text{Gr}(r; G_n) > (\log r)^{20}$ for every $r \geq R$ and $n \geq 1$. Then $p_c(G_n) \rightarrow p_c(G)$*

Proof of Corollaries 3.4.6 and 3.4.7. Observe that if G satisfies $\log \text{Gr}(r; G) > (\log r)^{20}$ for every $r \geq n_0$ then the statement **CORRIDOR**(i) holds vacuously for every i since $\mathcal{L}(G, 20) \cap [n_0, \infty)$ is empty. Thus, these two corollaries follow from Proposition 3.4.5 by the same argument used to deduce Theorems 3.3.3 and 8.1.1 from Proposition 3.3.1. (In fact the proof is slightly simpler since one no longer needs to control the burn-in.) \square

Remark 3.4.5 (Weaker growth conditions and the gap conjecture). The proof of Corollaries 3.4.6 and 3.4.7 extends straightforwardly to (sequences of) graphs satisfying much weaker growth conditions, such as

$$\log \text{Gr}(r) \geq c \log r (\log \log r)^{10}. \quad (3.4.26)$$

It is plausible that this class (and indeed the class treated by Corollaries 3.4.6 and 3.4.7) includes *every* transitive graph of superpolynomial growth, so that the methods of this section would suffice to prove locality and non-triviality of p_c for unimodular transitive graphs of superpolynomial growth. (One would still need the remainder of the paper to handle sequences of graphs of polynomial growth converging to a graph of superpolynomial growth, such as the Cayley graphs of the free step- s nilpotent groups, which converge to trees as $s \rightarrow \infty$.) Indeed, one formulation of Grigorchuk's *gap conjecture* (see [Gri14]) states that there exist universal positive constants c and γ such that if G is a Cayley graph of superpolynomial growth then $\text{Gr}(r) \geq e^{cr^\gamma}$ for every $r \geq 1$, and it seems reasonable to extend this conjecture to transitive graphs. Thus, the lower bound (3.4.26)

required for our arguments to work is much smaller than what might plausibly hold universally for all transitive graphs of superpolynomial growth. On the other hand, the best known bound for the gap conjecture, due to Shalom and Tao [ST10b], states that there exists a universal constant c such that

$$\log \text{Gr}(r) \geq c \log r (\log \log r)^c \quad (3.4.27)$$

for every Cayley graph of superpolynomial growth and every $r \geq 1$. (The authors claim their proof should yield this estimate with the constant $c = 0.01$, but do not carry out the necessary bookkeeping in the paper.) Even if the strong form of the gap conjecture is false, there does not seem to be any reason to believe that the Shalom-Tao bound (3.4.27) is optimal, so that (3.4.26) may very well hold for all transitive graphs of superpolynomial growth.

3.5 Quasi-polynomial growth I: Building disjoint tubes

In this section we prove the geometric facts that we will later use to analyze percolation in the low-growth (a.k.a. quasi-polynomial growth) regime. Let $d \geq 1$, let $G \in \mathcal{G}_d^*$, and let $D \geq 1$ be a fixed parameter. We recall that the set of **low growth scales** $\mathcal{L}(G, D)$ is defined to be

$$\mathcal{L}(G, D) = \left\{ n \geq 1 : \log \text{Gr}(m) \leq (\log m)^D \text{ for all } m \in [n^{1/3}, n] \right\}.$$

It would suffice for all our applications to take e.g. $D = 20$; we keep D as a parameter for now to emphasize that the analysis carried out in this section and Section 3.6 works for arbitrarily large D .

Recall that a **tube** is defined to be a set of the form $B_r(\gamma) = \bigcup_i B_r(\gamma_i)$ where γ is a path and $r \in (0, \infty)$ is a parameter we call the **thickness** of the tube; we call $\text{len}(\gamma)$ the **length** of the tube. We will also sometimes write $B(\gamma, r) = B_r(\gamma)$ to avoid writing large expressions in the subscript. (Strictly speaking the length and thickness of a tube depends on the pair (γ, r) used to represent it, but we will not belabour this point further.) We would like to show that whenever n is in the low growth regime, for all suitable sets (A, B) of vertices at scale n , we can find many disjoint tubes from A to B that are reasonably thick and not unreasonably long.

Definition 3.5.1. Let G be a connected transitive graph. Given $k, r, \ell \geq 1$ (which need not be integers) and $n \geq 1$, we say that G has **(k, r, ℓ) -plentiful radial tubes at scale n** if there exists a family of paths Γ in G with $|\Gamma| \geq k$ satisfying all of the following properties:

- Every path $\gamma \in \Gamma$ starts in the sphere S_n and ends in the sphere S_{4n} ;
- For every two distinct paths $\gamma, \gamma' \in \Gamma$, the tubes $B(\gamma, r)$ and $B(\gamma', r)$ are disjoint.

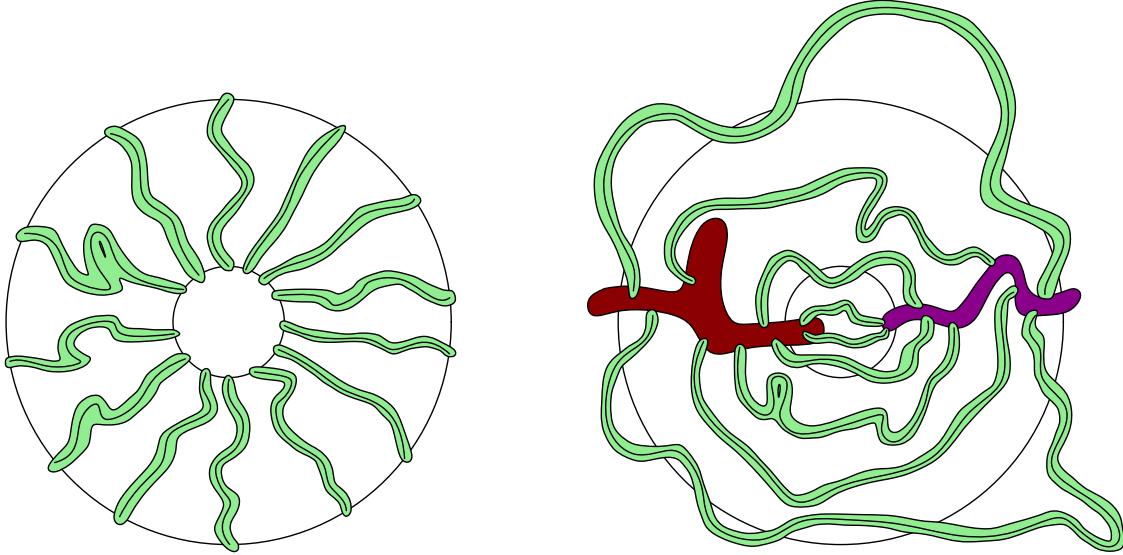


Figure 3.2: Schematic illustration of radial and annular tubes. Left: The (k, r, ℓ) -plentiful *radial* tubes condition means that we can find k disjoint tubes crossing the annulus that all have thickness r and length at most ℓ . Right: the (k, r, ℓ) -plentiful *annular* tubes condition means that for any two crossings of the annulus, we can find k disjoint tubes connecting the two crossings that all have thickness r and length at most ℓ ; these tubes are *not* required to stay inside the annulus.

- Every path $\gamma \in \Gamma$ has length at most ℓ ;

(In particular, the parameter k controls the number of tubes, the parameter r controls the thickness of the tubes, and the parameter ℓ controls the length of the tubes.) Given $m \geq n \geq 1$, we say that a set of vertices $A \subseteq V$ is an (n, m) **crossing** if it contains a path from S_n to S_m . We say that G has (k, r, ℓ) -**plentiful annular tubes at scale n** if for every pair of sets of vertices (A, B) that are both $(n, 3n)$ crossings¹⁰, there exists a family of paths Γ in G with $|\Gamma| \geq k$ satisfying all of the following properties:

- Every path $\gamma \in \Gamma$ starts in A and ends in B ;
- For every two distinct paths $\gamma, \gamma' \in \Gamma$, the tubes $B(\gamma, r)$ and $B(\gamma', r)$ are disjoint;
- Every path $\gamma \in \Gamma$ has length at most ℓ .

We say that G has (k, r, ℓ) -**plentiful tubes at scale n** if G has both (k, r, ℓ) -plentiful annular tubes and (k, r, ℓ) -plentiful radial tubes at scale n . See Figure 3.2 for an illustration. Given parameters

¹⁰The constant 3 that appears here could safely be replaced by any constant strictly larger than 2. We need the constant to be strictly larger than 2 to not cause problems with our application of Proposition 3.5.9.

$c, \lambda > 0$, we say that G has (c, λ) -**polylog-plentiful tubes** at scale n if it has (k, r, ℓ) -plentiful tubes with $k = (\log n)^{c\lambda}$, $r = n(\log n)^{-\lambda/c}$, and $\ell = n(\log n)^{\lambda/c}$. (NB: These tubes are very thick!)

Remark 3.5.1. The property of having (k, r, ℓ) -plentiful tubes at scale n gets stronger as k and r increase and gets weaker as ℓ increases; the property of having (c, λ) -polylog-plentiful tubes at scale n gets stronger as c increases but has no obvious monotonicity in the parameter λ . There is of course a trade-off between the number of tubes and their thicknesses, since e.g. if the $(n, 3n)$ crossings A and B are segments of geodesics from the origin between distances n and $3n$ then we cannot have more than $2n/r$ disjoint tubes of thickness r connecting A and B . In our applications in Section 3.6 we will want to use families of tubes that have thickness $n(\log n)^{-O(1)}$ and length $n(\log n)^{O(1)}$, which is why we phrase Proposition 3.5.2 in terms of polylog-plentiful tubes. As will be clear later in the section, our constructions allow for many other possible trade-offs between k and r which may be useful in future applications.

Ideally, we would like to say that there is some constant $c \in (0, 1)$ such that for every λ , G has (c, λ) -polylog-plentiful tubes at every sufficiently large scale $n \in \mathcal{L}(G, D)$. This is, however, slightly stronger than what we have been able to prove. (Fortunately it is also stronger than we need for our applications!) We instead establish the slightly weaker statement that for every scale $n \in \mathcal{L}(G, D)$, there exists a large interval of scales not much smaller than n on which we have plentiful tubes. We now give the precise statement, the proof of which takes up the rest of this section.

Proposition 3.5.2 (Quasi-polynomial growth yields a large range of consecutive scales with polylog-plentiful tubes). *For each $D \in [1, \infty)$, $\lambda \in [1, \infty)$, and $d \geq 1$ there exist positive constants $c(d, D) \in (0, 1)$ and $n_0 = n_0(d, D, \lambda)$ such that the following holds. Let G be an infinite, connected, unimodular transitive graph with vertex degree d that is not one-dimensional. For each integer $n \geq n_0$ with $n \in \mathcal{L}(G, D)$ there exist integers $n^{1/3} \leq m_1 \leq m_2 \leq n$ with $m_2 \geq m_1^{1+c}$ such that G has (c, λ) -polylog-plentiful tubes at every scale $m_1 \leq m \leq m_2$.*

The constant $c = c(d, D)$ can be thought of as an “exchange rate”, governing the cost to trade between the number of tubes and their thicknesses, while λ is the parameter we vary to make this trade. *It is very important that the constant c does not depend on λ !*

Remark 3.5.2. No non-trivial statements about plentiful annular tubes can be made without some kind of growth upper bound (or other geometric assumption), since the 3-regular tree does not have plentiful (k, r, ℓ) -plentiful annular tubes for any $k, r, \ell > 1$. Similarly, the assumption that the graph is not one-dimensional is needed since the line graph \mathbb{Z} does not have (k, r, ℓ) -plentiful

radial or annular tubes for any $k, r, \ell > 1$. (On the other hand, the assumption of unimodularity is redundant since it is implied by the existence of large scales with subexponential growth [Hut20a, Section 5.1].)

To prove this proposition, we split into two cases according to the *rate of change* of growth in the given interval. More precisely, we will split according to whether or not there exists n in a suitable initial segment of the interval such that G satisfies the small-tripling condition $\text{Gr}(3n) \leq 3^5 \text{Gr}(n)$ at scale n , where the constant 5 could be replaced by any other constant strictly larger than 4. (This is related to the fact that four is the critical dimension for two independent random walks to intersect infinitely often almost surely.) Our proofs in the two cases are completely different from one another.

In the first case, when there does exist such an n , we will build our disjoint tubes by applying the structure theory of transitive graphs of polynomial volume growth [BGT12a; TT21a; EH23d]. Informally, this theory guarantees that, for a large interval of scales, our graph looks approximately like the Cayley graph of a finitely presented group whose relations are generated by cycles of diameter much smaller than the given scale. We will use these techniques to prove the following proposition.

Proposition 3.5.3 (Slow tripling yields plentiful tubes). *For each $d \geq 1$ and $\kappa < \infty$ there exist positive constants $c = c(d, \kappa)$, $C = C(d, \kappa)$, and $n_0 = n_0(d, \kappa)$ such that if G is an infinite, connected, transitive graph of vertex degree d that is not one-dimensional and $n \geq n_0$ is such that $\text{Gr}(3n) \leq 3^\kappa \text{Gr}(n)$, then there exists a set $A \subset [n, \infty)$ with $|A| \leq C$ such that for each $k \geq 1$, G has $(ck, ck^{-1}m, Ck^C m)$ -plentiful tubes at every scale $m \geq Ckn$ such that $m \notin \bigcup_{a \in A} [a, 2ka]$.*

In the second case, when there does *not* exist such a scale n with small tripling, we will prove that the desired plentiful tubes condition holds using random walks. More specifically, we will apply an estimate of [BDKY15], which can be thought of as a “quantitative weak elliptic Harnack inequality” for graphs of subexponential growth. Under our low-growth assumption, this inequality implies that two random walks on G started at distance $n' \in [n^{1/3}, n^{0.9}]$, say, can be coupled to coincide with good probability by the time they reach distance $n'(\log \text{Gr}(n'))^{O(1)}$. On the other hand, the assumption that the growth is rapidly increasing (in an at-least-five-dimensional fashion) lets us prove that two independent pairs of these coupled random walks are unlikely to have their tubes intersect. (This will be proven using fairly standard random walk techniques, most notably the isoperimetric inequality of Coulhon and Saloff-Coste [CS93] and the heat kernel bounds resulting from this inequality together with the work of Morris and Peres [MP05].) Unfortunately these

walks will have length of order $(n')^2(\log \text{Gr}(n'))^{O(1)}$, which is larger than we want by a factor of n' . This can be fixed by a simple coarse-graining argument, using the diffusivity of random walks on low-growth graphs, to replace the random walk by a shorter path with essentially the same tube around it. We will use these techniques to prove the following proposition.

Proposition 3.5.4 (Fast tripling and quasi-polynomial growth yield plentiful tubes). *Let G be an infinite connected unimodular transitive graph with vertex degree d , let $D, \lambda \geq 1$ and let $\varepsilon > 0$. There exist positive constants $c = c(d, D, \varepsilon)$ and $n_0 = n_0(d, D, \lambda, \varepsilon)$ such that if $n \geq n_0$ satisfies*

$$\text{Gr}(m) \leq e^{(\log m)^D} \quad \text{and} \quad \text{Gr}(3m) \geq 3^5 \text{Gr}(3m) \quad \text{for every } n^{1-\varepsilon} \leq m \leq n^{1+\varepsilon}$$

then G has (c, λ) -polylog-plentiful tubes at scale n .

We prove Proposition 3.5.3 in Section 3.5 and Proposition 3.5.4 in Section 3.5. Before doing this, we note that Proposition 3.5.2 follows easily from these two propositions.

Proof of Proposition 3.5.2 given Propositions 3.5.3 and 3.5.4. Let $c_1(d, 5)$, $C(d, 5)$, and $n_1(d, 5)$ be the constants from Proposition 3.5.3 with $\kappa := 5$. Let $c_2(d, D, 0.1)$ and $n_2(d, D, \lambda, 0.1)$ be the constants from Proposition 3.5.4 with $\varepsilon := 0.1$. We may assume that $c_1 \vee c_2 \leq 1/2$ and $C \geq 2$. Suppose that some $n \in \mathcal{L}(G, D)$ is large with respect to d, D, λ , satisfying in particular $n \geq (n_1 \vee n_2)^3$. If there exists $n' \in [n^{1/3}, n^{10/11}]$ such that $\text{Gr}(3n') \leq 3^5 \text{Gr}(n')$, then we may apply Proposition 3.5.3 with $k := (\log n)^\lambda$ to obtain an interval of scales $[m_1, m_2] \subseteq [n^{1/3}, n]$ with $m_2 \geq m_1^{1.1}$ on which G has $(1/(2C), \lambda)$ -polylog-plentiful tubes on every scale. Otherwise, since each $n' \in [n^{0.4}, n^{0.8}]$ satisfies $[(n')^{0.9}, (n')^{1.1}] \subseteq [n^{1/3}, n^{10/11}]$, we can apply Proposition 3.5.4 to each such scale to obtain that G has (c_1, λ) -plentiful tubes at every scale in the interval $[n^{0.4}, n^{0.8}]$. This is easily seen to imply the claim in either case. \square

Using the structure theory of approximate groups

The goal of this subsection is to prove Proposition 3.5.3. Let us first give some relevant context. Given a graph G , a vertex $v \in V(G)$, and a radius $r \geq 0$, we define the **exposed sphere** $S_r^\infty(v)$ to be the set of vertices $u \in S_r(v)$ such that there exists an infinite self-avoiding path started at u that never returns to $B_r(v)$ after its first step. When G is transitive, we set $S_r^\infty := S_r^\infty(o)$. In [CMT22], the authors applied results of Timar [Tim07] to obtain geometric control of exposed spheres in transitive graphs of polynomial growth using the fact that (by Gromov's theorem [Gro81b] and Trofimov's theorem [Tro84b]) these graphs are quasi-isometric to Cayley graphs of nilpotent groups, which are finitely presented.

In this section, we will run a similar argument to build our disjoint tubes under the hypotheses of Proposition 3.5.3. While these hypotheses suffice to guarantee polynomial growth by the results of [BGT12a; TT21a] (discussed in detail below), a key technical difference between our analysis and that of [CMT22] is that we must run all our arguments in a more finitary, quantitative way; we need to build our disjoint tubes at a scale not much larger than the scale where we are assumed to witness relative polynomial growth. This requires us to engage more deeply with the structure theory of approximate groups than was necessary in [CMT22]. Indeed, rather than using Gromov’s theorem and Trofimov’s theorem, we will instead apply finitary versions of these theorems due to Breuillard, Green, and Tao [BGT12a] and Tessera and Tointon [TT21a]. Moreover, we will also apply a *new* structure theoretic result proven in our paper [EH23d], which can be thought of as a “uniform” version of the statement that groups of polynomial growth are finitely presented.

Remark 3.5.3. It will be convenient in this section to let Γ denote a group. This will not conflict with the Γ used to denote a set of disjoint tubes, which we will not use in this section.

Structure theory. We now state the main structure-theoretic results we will use after reviewing some relevant definitions. We begin with Tessera and Tointon’s finitary structure theorem for vertex-transitive graphs of low growth [TT21a]; this theorem builds on Breuillard, Green, and Tao’s structure theorem for approximate groups [BGT12a] as well as Carolino’s extension of this theorem to locally compact approximate groups [Car15]. Recall that a function $\phi : V_1 \rightarrow V_2$ between the vertex sets of two graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ is said to be an (α, β) -**quasi-isometry** (a.k.a. **rough isometry**) if

$$\alpha^{-1}d(x, y) - \beta \leq d(\phi(x), \phi(y)) \leq \alpha d(x, y) + \beta$$

for every $x, y \in V_1$ and every vertex $z \in V_2$ is within distance at most β of $\phi(V_1)$; note that the second property holds automatically if ϕ is surjective. Given a transitive graph G and a subgroup $H \subseteq \text{Aut}(G)$, we write G/H for the associated quotient graph. If H is a normal subgroup of $\text{Aut}(G)$ then the action of $\text{Aut}(G)$ on G descends to a transitive action of $\text{Aut}(G)$ on G/H (see [TT21a, Section 3]), and we write $\text{Aut}(G)_{G/H}$ for the image of $\text{Aut}(G)$ in $\text{Aut}(G/H)$ induced by this action. Note that we view $\text{Aut}(G)$ as a *topological* group where convergence is given by pointwise convergence (see [TT21a, §4]).

Theorem 3.5.5 (Finitary structure theory of transitive graphs of polynomial growth). *For each $K \geq 1$ there exist constants $n_0 = n_0(K)$ and $C = C(K)$ such that the following holds. Let G be a connected (locally finite) (vertex-)transitive graph with a distinguished vertex o , and suppose that*

there exists $n \geq n_0$ such that $\text{Gr}(3n) \leq K \text{Gr}(n)$. Then there exists a compact normal subgroup $H \triangleleft \text{Aut}(G)$ such that:

1. Every fibre of the projection $\pi : G \rightarrow G/H$ has diameter at most Cn .
2. $\text{Aut}(G)_{G/H}$ can be canonically identified with $\text{Aut}(G)/H$.
3. $\text{Aut}(G)/H$ has a nilpotent normal subgroup N of rank, step and index at most C .
4. The set $S = \{g \in \text{Aut}(G)_{G/H} : d_{G/H}(g(Ho), Ho) \leq 1\}$ is a finite symmetric generating set for $\text{Aut}(G)/H$.
5. Every vertex stabiliser of the action of $\text{Aut}(G)/H$ on G/H has cardinality at most C .
6. If for each $v \in G$ we let $g_v \in \text{Aut}(G)/H$ be such that $g_v(\pi(e)) = \pi(v)$ then $v \mapsto g_v$ is a $(1, Cn)$ -quasi-isometry from G to $\text{Cay}(\text{Aut}(G)/H, S)$.
7. If Gr' denotes the growth function of the Cayley graph $\text{Cay}(\text{Aut}(G)/H, S)$ then

$$\frac{\text{Gr}(m_2)}{C \text{Gr}(m_1 + Cn)} \leq \frac{\text{Gr}'(m_2)}{\text{Gr}'(m_1)} \leq \frac{C \text{Gr}(m_2 + Cn)}{\text{Gr}(m_1)}$$

for every $m_1, m_2 \in \mathbb{N}$.

8. The growth bound $\text{Gr}(m_2) \leq C(m_2/m_1)^C \text{Gr}(m_1)$ holds for every $m_2 \geq m_1 \geq m$.

Proof of Theorem 3.5.5. The first five items of the theorem are essentially equivalent to [TT21a, Theorem 2.3] (although that theorem is slightly more general as it allows one to replace $\text{Aut}(G)$ with any other transitive group of automorphisms of G). In their original statement of the theorem, Tessera and Tointon do not explicitly identify the rough isometry from G to $\text{Cay}(\text{Aut}(G)/H, S)$, but the fact we can take it to be of the form above is implicit in their proof. Item 7 is implied by [TT21a, Proposition 9.1], while Item 8 follows from Item 7 together with [BT16, Theorem 1.1]. \square

Remark 3.5.4. The set S has size equal to the union of the stabilizers of the vertices $\{u \in G/H : u \sim v\}$, so that $|S| \leq C(\deg(o) + 1)$.

Remark 3.5.5. For $K = 3^\kappa$ with κ an integer, the growth bound $\text{Gr}(m_2) \leq C(m_2/m_1)^C \text{Gr}(m_1)$ can be improved to the sharp bound $\text{Gr}(m_2) \leq C(m_2/m_1)^\kappa \text{Gr}(m_1)$ under the stronger assumption that the graph satisfies an *absolute* growth bound of the form $\text{Gr}(n) \leq \varepsilon_K n^{\kappa+1}$ at a sufficiently large scale and for a sufficiently small constant $\varepsilon_K > 0$ [TT21a, Corollary 1.5]. This strong bound is not implied by the small-tripling condition (which is the relevant condition for our applications) as

shown in [Tao17a, Example 1.11]. (Indeed, this example suggests that the small-tripling condition $\text{Gr}(3m) \leq 3^\kappa \text{Gr}(m)$ should not imply any bound on the limiting growth dimension stronger than $O(\kappa^2)$. Optimal bounds on growth implied by small tripling will be established in forthcoming work of Tessera and Tointon.)

Uniform finite presentation. We next state our theorem on uniform finite presentation proven in [EH23d]. Given a set of elements A in a group Γ , we define $\langle\langle A \rangle\rangle$ to be the normal subgroup of Γ generated by A and define $\bar{A} = A \cup \{\text{id}\} \cup A^{-1}$. Consider a group Γ with a finite generating set S , let F_S be the free group on S and let $\pi : F_S \rightarrow \Gamma$ be the associated group homomorphism with kernel R . Since $\Gamma \cong F_S/R$, we can think of the sequence of quotients $(\Gamma_r)_{r \geq 1}$ defined by $\Gamma_r := F_S/\langle\langle \bar{S}^r \cap R \rangle\rangle$ as being finitely presented approximations to Γ , since Γ_r admits a finite presentation $\Gamma_r = \langle S \mid R_r \rangle = F_S/\langle\langle R_r \rangle\rangle$ with $R_r = \bar{S}^r \cap R \subseteq \bar{S}^r$. These approximations have the property that the Cayley graphs $\text{Cay}(\Gamma_r, S)$ and $\text{Cay}(\Gamma, S)$ have isomorphic $((r/2) - 1)$ -balls. We record this fact in the following lemma, which is taken from [EH23d, Lemma 5.6]. (Although it is stated there only in the case that r is a power of 2, the same proof works for arbitrary r .)

Lemma 3.5.6. *Let Γ be a group with a finite generating set S . For all $i \geq 1$, the quotient map $\Gamma_r \rightarrow \Gamma$ induces a map of the associated Cayley graphs that restricts to an isomorphism between the balls of radius $\lfloor r/2 \rfloor - 1$.*

The main result of [EH23d] can be stated as follows.

Theorem 3.5.7 ([EH23d], Theorem 1.1). *For each $K, d < \infty$ there exist constants $n_0 = n_0(K)$ and $C = C(K, d)$ such that if Γ is a group and S is a finite generating set for Γ with $|S| \leq d$ whose growth function Gr satisfies $\text{Gr}(3n) \leq K \text{Gr}(n)$ for some integer $n \geq n_0$ then*

$$\#\left\{k \in \mathbb{N} : k \geq \log_2 n \text{ and } \langle\langle R_{2^{k+1}} \rangle\rangle \neq \langle\langle R_{2^k} \rangle\rangle\right\} \leq C.$$

For our purposes, the main output of Theorems 3.5.5 and 3.5.7 is that if the small tripling condition $\text{Gr}(3n) \leq K \text{Gr}(n)$ holds at some sufficiently large n , then at “most” scales $r \geq n$ the graph “looks like” the Cayley graph of a finitely presented group with relations generated by words of length much smaller than r .

Using the structure theory to build disjoint tubes. It remains to apply Theorems 3.5.5 and 3.5.7 to prove Proposition 3.5.3. When working with a group Γ and a finite generating set S of Γ , we will continue to use the notation $(\Gamma_r)_{r \geq 1}$ and $(R_r)_{r \geq 1}$ as defined above.

Let us introduce some more definitions. Let $G = (V, E)$ be a graph, and consider the set $\{0, 1\}^E$ with addition modulo 2 as a vector space over \mathbb{Z}_2 . If $G = \text{Cay}(\Gamma, S)$ for some group $\Gamma = \langle S | R \rangle$ with finite generating set S and (not necessarily finitely generated) relation group $R \triangleleft F_S$, then we can identify every oriented cycle started at the origin with a word in R . Under this identification, we see that if R is generated as a normal subgroup by words of length at most r in F_S , then every cycle in G can be written as a mod-2 sum of cycles of length at most r . Note that this leads to a notion of finite presentation for graphs that are not Cayley graphs, namely that their cycle space is generated as a \mathbb{Z}_2 -vector space by the cycles of some finite length $r < \infty$; see [Tim07] for further details.

We will rely crucially on the following lemma, which is essentially due to Timár [Tim07]. Recall that an infinite graph $G = (V, E)$ is **one-ended** if for every finite set of vertices $W \subseteq V$, the graph $G \setminus W$ has exactly one infinite component; groups of polynomial growth are one-ended if and only if they are not one-dimensional. (Note that when G is the Cayley graph of an infinite finitely-generated group, the property of being one-ended is independent of the choice of finite generating set, so that one can sensibly refer to a *group* as being one-ended without specifying a generating set.)

Lemma 3.5.8. *Let G be an infinite, connected, one-ended transitive graph. If $r \in \mathbb{N}$ has the property that every cycle in G is equal to a sum of cycles of length at most r , then for every $k, n \in \mathbb{N}$ with $r \leq k \leq n$ and every $u, v \in S_n^\infty$ there exists a path from u to v that is contained in $\bigcup_{x \in S_n^\infty} B_{2k}(x)$ and has length at most $3k |B_{3n}| / |B_k|$.*

Proof. This statement is implicit in the proofs of [CMT22, Lemma 2.1 and 2.7], and is an easy consequence of [Tim07, Theorem 5.1]. \square

We will also need two more elementary geometric facts. The first states that if a path travels from a sphere S_r to S_{2r+1} , then it must pass through the exposed sphere S_r^∞ .

Proposition 3.5.9 ([FGO15, Proposition 5]). *Let G be an infinite connected transitive graph and let $r \in \mathbb{N}$. Every path that starts in S_r and ends in S_{2r+1} contains a vertex in S_r^∞ .*

The next lemma lets us pass disjoint tubes through quasi-isometries with all relevant quantities changing in a controlled way.

Lemma 3.5.10. *Let $G = (V, E)$ and $G' = (V', E')$ be two graphs and let $\phi : V \rightarrow V'$ be an (α, β) -quasi isometry for some $\alpha, \beta \geq 1$. Let $u, v \in V$ and suppose that $x, y \in V'$ satisfy $d(x, \phi(u)) \leq \beta$*

and $d(y, \phi(v)) \leq \beta$. For each path γ' from x to y in G' , there exists a path γ from u to v in G such that $\text{len}(\gamma) \leq 10\alpha(\text{len}(\gamma') + \beta)$ and

$$\phi(B_r(\gamma)) \subseteq B_{\alpha r + (5\alpha^2 + 2)\beta}(\gamma')$$

for every $r \geq 0$.

Proof of Lemma 3.5.10. Let $\psi : V' \rightarrow V$ be a function such that $\psi(x) = u$, $\psi(y) = v$, and $d(\phi(\psi(z)), z) \leq \beta$ for every $z \in V$; such a function exists by the definition of an (α, β) -quasi-isometry and our assumptions on u , v , x , and y . Let $\ell = \lceil \text{len}(\gamma')/\lceil \beta \rceil \rceil$, and let the sequence $(n_i)_{i=0}^\ell$ be defined by $n_i = \lceil \beta \rceil i$ for $i < \ell$ and $n_\ell = \text{len}(\gamma')$. For each $0 \leq i \leq \ell$ let $x_i = \gamma'_{n_i}$ and let $u_i = \psi(x_i)$, so that $u_0 = u$ and $u_\ell = v$. For each $0 \leq i < \ell$, the points x_i and x_{i+1} have distance at most $\lceil \beta \rceil \leq 2\beta$ in G' , so that $\phi(u_i)$ and $\phi(u_{i+1})$ have distance at most 4β in G' . Since ϕ is an (α, β) -quasi-isometry, it follows that u_i and u_{i+1} have distance at most $5\alpha\beta$ in G . Let γ be formed by concatenating geodesics between u_i and u_{i+1} for each $0 \leq i < \ell$. The path γ clearly has length at most $5\alpha\beta\lceil \text{len}(\gamma')/\lceil \beta \rceil \rceil \leq 10\alpha(\text{len}(\gamma') + \beta)$. Moreover, given $r \geq 1$, each point in $B_r(\gamma)$ has distance at most $5\alpha\beta + r$ from one of the points u_i , whose image under ϕ has distance at most β from one of the points of γ' . Using the definition of an (α, β) -quasi-isometry again, this implies that every point in $\phi(B_r(\gamma))$ has distance at most $\alpha(5\alpha\beta + r) + 2\beta$ from a point of γ' , which is equivalent to the desired set inclusion. \square

We now have everything we need to prove Proposition 3.5.3.

Proof of Proposition 3.5.3. Fix $\kappa < \infty$ and let $n_0 = n_0(\kappa)$ and $C_1 = C_1(\kappa)$ be the constants from Theorem 3.5.5 applied with $K = 3^\kappa$. Suppose that G is an infinite, connected, transitive graph of vertex degree d that is not one-dimensional and that $n \geq n_0$ is such that $\text{Gr}(3n) \leq 3^\kappa \text{Gr}(n)$. Let H be the normal subgroup of $\text{Aut}(G)$ that is guaranteed to exist by Theorem 3.5.5, let S be the generating set for $\text{Aut}(G)/H$ from Item 4 of Theorem 3.5.5, and write $\Gamma := \text{Aut}(G)/H$. Letting Gr' denote the growth function of the Cayley graph $\text{Cay}(\Gamma, S)$, we have by Items 7 and 8 of Theorem 3.5.5 that

$$\frac{\text{Gr}'(3n)}{\text{Gr}'(n)} \leq \frac{C_1 \text{Gr}(3n + C_1 n)}{\text{Gr}(n)} \leq C_1^2 (3 + C_1)^{C_1}.$$

Moreover, as discussed in Remark 3.5.4, the generating set S has at most $C_1(d+1)$ elements. Thus, if we take $n_1 = n_1(\kappa) \geq n_0$ and $C_2 = C_2(\kappa, d)$ to be the constants from Theorem 3.5.7 applied with $K = C_1^2(3 + C_1)^{C_1}$ and $d = C_1(d+1)$, we have that if $n \geq n_1$ then

$$\#\left\{i \in \mathbb{N} : i \geq \log_2 n \text{ and } \langle\langle R_{2^{i+1}} \rangle\rangle \neq \langle\langle R_{2^i} \rangle\rangle\right\} \leq C_2, \quad (3.5.1)$$

where R_k is the set of relations of Γ that have word length at most k and $\langle\langle R_k \rangle\rangle$ is the smallest normal subgroup of the free group F_S generated by R_k . Let $G' = (V', E')$ be the Cayley graph $\text{Cay}(\Gamma, S)$, let $\phi : V \rightarrow V'$ be the $(1, C_1 n)$ -quasi-isometry that is guaranteed to exist by Theorem 3.5.5 (which we may assume maps ϕ to the identity of Γ , which we denote by id), and let $\psi : V' \rightarrow V$ be such that $d(\phi(\psi(x)), x) \leq C_1 n$ for every $x \in V'$, such a function being guaranteed to exist since ϕ is a $(1, C_1 n)$ -quasi-isometry. This function ψ is easily seen to be a $(1, C_3 n)$ -quasi-isometry for an appropriate choice of constant $C_3 = C_3(\kappa)$. For each $k \geq 1$ let G'_k be the Cayley graph $\text{Cay}(\Gamma_k, S)$ of the group Γ_k defined by $\Gamma_k = \langle S \mid R_k \rangle$.

Fix $n \geq n_1$ and consider the set $A = \{2^i : i \geq \log_2 n \text{ and } \langle\langle R_{2^{i+5}} \rangle\rangle \neq \langle\langle R_{2^i} \rangle\rangle\}$. This set contains at most five times as many elements as the set considered in (3.5.1), so that $|A| \leq 5C_2$. Fix $k \geq n$ and suppose that $m \geq 2kn$ does not belong to $[a, 2ka]$ for any $a \in A$, so that if we define $a(m) = \sup\{a \in A : a \leq m\}$ then $a(m) \leq m/2k$. The definition of A ensures that $\langle\langle R_{2a(m)} \rangle\rangle = \langle\langle R_{16m} \rangle\rangle$ and hence by Lemma 4.5.6 that the Cayley graphs G' and $G'_{a(m)}$ have isomorphic $(8m - 1)$ -balls. For each $r < 4m$ we write $(S_r^\infty)'$ for the exposed spheres in G' and $G'_{a(m)}$, which can be identified by Proposition 3.5.9. We also identify the balls $B'_r(x)$ in G' and $G'_{a(m)}$ for points x of distance at most $4m$ from the identity and all $r \leq m$, where the prime on $B'_r(x)$ reminds us that we are working with G' rather than G .

Now, since $\Gamma_{a(m)}$ has its group of relations generated by its relations of length at most $a(m) \leq m/2k$, it follows from Lemma 3.5.8 that for each $m/2k \leq m_1 \leq m_2 \leq 3m$ and each $u, v \in (S_{m_2}^\infty)'$ there exists a path from u to v in $G'_{a(m)}$ that is contained in $\bigcup_{x \in (S_{m_2}^\infty)'} B'_{2m_1}(x)$ and has length at most $3m_1 \text{Gr}(3m_2)/\text{Gr}(m_1)$, where all sets and growth functions are identical in the two groups $\Gamma_{a(m)}$ and Γ by the restrictions placed on k and m . Since these paths are entirely contained within the ball for which G' and $G'_{a(m)}$ are identical, they exist in G' also. Moreover, since $m_2 \geq m_1 \geq m/2k \geq n$, we have by Item 9 of Theorem 3.5.5 as above that

$$\frac{\text{Gr}'(3m_2)}{\text{Gr}'(m_1)} \leq C_1^2 (3 + C_1)^{C_1} \left(\frac{m_2}{m_1}\right)^{C_1}$$

so that the length of this path is at most $3C_1^2 (3 + C_1)^{C_1} (m_2/m_1)^{C_1} m_1 =: C_4 (m_2/m_1)^{C_1} m_1$. (Everything discussed in this paragraph is still under the assumption that $m \geq 2kn$ does not belong to $[a, 2ka]$ for any $a \in A$.)

We now use the existence of these paths in G' to guarantee the desired plentiful tube conditions in the original graph G . More concretely, we will prove that there exist positive constants $c = c(\kappa, d)$,

$C = C(\kappa, d)$, and $n_2 = n_2(\kappa) \geq n_1$ such that if $n \geq n_2$ then G has $(ck, ck, Ck^C m)$ -plentiful tubes on each scale $m \geq Ckn$ such that m does not belong to $[a, 2ka]$ for any $a \in A$.

We begin by constructing *annular* tubes. It suffices to construct tubes in the case that the two $(m, 3m)$ crossings are both the vertex sets of paths η_1, η_2 from S_m to S_{3m} since any crossing contains the vertex set of such a path. Fix $m \geq 2kn$ and two paths η_1 and η_2 from S_m to S_{3m} in G , and suppose that m does not belong to $[a, 2ka]$ for any $a \in A$. Apply Lemma 3.5.10 with each of these paths and the $(1, C_3 n)$ -quasi-isometry $\psi : V' \rightarrow V$ (taking u and v to be the endpoints of η_i , x to be $\phi(u)$ and y to be $\phi(v)$) to obtain two paths η'_1 and η'_2 in G' . Using the $(1, C_1 n)$ -quasi-isometry property of ϕ , each of these paths starts at distance at most $m + C_1 n$ from $\phi(o) = \text{id}$ and ends at distance at least $3m - C_1 n$ from id . If $m \geq 9C_1 n$ then both paths start at distance at most $\frac{10}{9}m$ from id and end at distance at least $\frac{26}{9}m$ from id . As such, it follows from Proposition 3.5.9 that if $m \geq 9C_1 n$ then the paths η'_1 and η'_2 both intersect the exposed sphere $(S_{m_2}^\infty)'$ in G for each integer $i \in I := \mathbb{Z} \cap [\frac{10}{9}m, \frac{12}{9}m]$. (The only property of these numbers we will need is that $12 > 10$ and $12 < 26/2$.) Fix $m_1 = \lceil m/2k \rceil \leq m$ and for each integer $i \in I$ let x_i be a point of η'_1 belonging to $(S_i^\infty)'$ and let y_i be a point of η'_2 belonging to $(S_i^\infty)'$. If $m \geq \max\{9C_1 n, 2kn\}$ then, since m was assumed not to belong to $[a, 2ka]$ for any $a \in A$, it follows from Lemma 3.5.8 as discussed above that for each such i there exists a path γ'_i from x_i to y_i that is contained in $\bigcup_{(S_i^\infty)'} B'_{2m_1}(x)$ and has length at most $C_4(i/m_1)^{C_1} m_1$. Since $i \leq 3m$, the length of this path can be bounded by $C_5 k^{C_1-1} m$ for an appropriate constant $C_5 = C_5(\kappa)$. Moreover, if $r \geq 1$ and i and j are two integers $i, j \in I$ satisfying $|i - j| > 4m_1 + 2r$ then the tubes of radius r around γ'_i and γ'_j in G' are disjoint since they are contained in disjoint annuli $\bigcup_{(S_i^\infty)'} B'_{2m_1+r}(x)$ and $\bigcup_{(S_j^\infty)'} B'_{2m_1+r}(x)$.

Since Lemma 3.5.10 guaranteed that the paths η'_1 and η'_2 have images under ψ contained in the $4C_1 n$ -neighbourhoods of η_1 and η_2 respectively, we can for each $i \in I$ find points u_i and v_i in η_1 and η_2 respectively such that $d(x_i, \phi(u_i))$ and $d(y_i, \phi(v_i))$ are at most $6C_1 n$. Applying Lemma 3.5.10 to each of the paths γ'_i (with the quasi-isometry ϕ , the points u_i and v_i in G , the points x_i and y_i in G' , and the quasi-isometry constants $\alpha = 1$ and $\beta = 6C_1 n$), it follows that there exist constants $C_6 = C_6(\kappa)$ and $C_7 = C_7(\kappa)$ such that for each $i \in I$ there exists a path γ_i from u_i to v_i of length at most $C_6 k^{C_1-1}$ such that if $i, j \in I$ satisfy $|i - j| \geq C_7 m_1$ then the tubes of radius m_1 around γ_i and γ_j are disjoint. The claim about annular tubes follows easily by taking a $C_7 m_1$ -separated subset of I (i.e., a subset of I in which all distinct pairs of integers have distance at least $C_7 m_1$), since such a set may be taken to have size at least $I/C_7 m_1 \geq c_1 k$ for some positive constant $c_1 = c_1(\kappa)$ whenever n is larger than some constant $n_2 = n_2(\kappa) \geq n_1$. (The freedom to increase n_1 to n_2 lets us make sure that every real number we round down is at least 1, so that rounding cannot reduce any

relevant quantities by more than a factor of $1/2$.)

We now briefly argue that the same construction also yields *radial tubes* crossing a shifted annulus. Run the above construction again with η_1 and η_2 taken to be the two portions crossing from S_m to S_{12m} of a doubly-infinite geodesic passing through o , so that every point in η_1 has distance at least $2m$ from every point in η_2 , but with various constants changed appropriately since we are now working with $(m, 12m)$ crossings instead of $(m, 3m)$ crossings. In particular, the interval I can now be taken to be $\mathbb{Z} \cap [8m, 10m]$, with various other constants changing to reflect this change. Consider the point u in η_1 that has distance $9m$ from o . When we perform the above construction to build paths between η_1 and η_2 , each path γ_i starts at distance at most m from u and ends at distance at least $11m$ from u . Thus, since G is transitive, the family of paths we have constructed verifies the $(ck, ck, Ck^C m)$ -plentiful radial tubes condition holds at the scale m as desired, for some constants $c = c(\kappa, d)$ and $C = C(\kappa, d)$. As before, this works under the assumption that $n \geq n_2$ for some $n_2 = n_2(\kappa)$ and that $m \geq Ckn$ does not belong to $[a, 2ka]$ for any $a \in A$. \square

Using random walk trajectories

In this section we prove Proposition 3.5.4, which verifies the plentiful tubes condition for graphs that have quasi-polynomial absolute growth but a fast rate of relative growth over an appropriate range of scales. As discussed above, we will construct the required collections of disjoint tubes by modifying certain conditioned random walk trajectories. To avoid parity issues, we work with *lazy* random walks throughout the section. We will spend most of the section proving general bounds on the behaviour of random walk on some scale in terms of the growth of the graph at that scale, specializing to the setting of Proposition 3.5.4 only at the very end of the proof. Given a graph G and a vertex u of G , let \mathbf{P}_u denote the law of the *lazy* random walk started from u , which at each step either stays in place with probability $1/2$ or else crosses a uniform random edge emanating from its current position, and let the **heat kernel** $p_t(u, v)$ be defined by $p_t(u, v) = \mathbf{P}_u(X_t = v)$.

We begin by recalling two important facts about random walks on graphs of quasi-polynomial growth that will be used in the proof: the Varopoulos-Carne inequality [Var85; Car85], which implies near-diffusive estimates on the rate of escape, and the total variation inequality of [Yad23, Chapter 7.5], which implies that two walks started from different vertices can be coupled to coalesce by the time they reach a distance that is near-linear in their starting distance.

Diffusive estimates from Varopoulos-Carne. We now state the Varopoulos-Carne inequality, which gives Gaussian-like bounds on the n -step transition probabilities between two specific ver-

tices; see e.g. [LP16c, Chapter 13.2] for a modern treatment. This inequality does not require transitivity, and holds for the random walk on any graph.

Theorem 3.5.11 (Varopoulos-Carne). *Let $G = (V, E)$ be a (locally finite) graph. Then*

$$p_t(u, v) \leq 2 \sqrt{\frac{\deg(v)}{\deg(u)}} \exp \left[-\frac{d(u, v)^2}{2t} \right],$$

for every $t \geq 1$ and every $u, v \in V$.

The Varopoulos-Carne inequality easily implies that the random walk on a graph of quasi-polynomial growth is very unlikely to be at a distance much larger than $t^{1/2}(\log t)^{O(1)}$ from its starting point.

Corollary 3.5.12. *Let $G = (V, E)$ be a (locally finite) transitive graph and let o be a vertex of G . Then*

$$\mathbf{P}_o \left(\max_{0 \leq k \leq t} d(o, X_k) \geq n \right) \leq 2(t+1) \text{Gr}(n) \exp \left[-\frac{n^2}{2t} \right]$$

for every $t, n \geq 1$.

Proof of Corollary 3.5.12. If $\max_{0 \leq k \leq t} d(o, X_k) \geq n$ then there exists $0 \leq k \leq t$ such that X_k has distance exactly n from o . Since the number of points at distance n from o is at most $\text{Gr}(n)$, the claim follows from Theorem 3.5.11 by taking a union bound over the possible values of k and X_k . \square

Remark 3.5.6. For transitive graphs of polynomial growth, Varopoulos-Carne implies a displacement upper bound of the form $\sqrt{t \log t}$ while the true displacement is of order \sqrt{t} with high probability. (This sharp upper bound on the displacement can be proven using a (highly nontrivial) improvement of the Varopoulos-Carne inequality due to Hebisch and Saloff-Coste [HS93].) As such, our reliance on Varopoulos-Carne leads to all of the estimates in this section having polylog terms that are known to be unnecessary for transitive graphs of polynomial growth and are presumably non-optimal for transitive graphs of quasi-polynomial growth also.

Coupling from low growth. We now explain how low growth can be used to couple two walks to coalesce within time not much larger than quadratic in their starting distance; we will eventually concatenate these pairs of coupled random walks to build annular tubes. We first recall some

relevant definitions. Given two probability measures μ and ν on a countable set Ω , the **total variation distance** $\|\mu - \nu\|_{\text{TV}}$ between μ and ν is defined by

$$\|\mu - \nu\|_{\text{TV}} = \sup_{A \subseteq \Omega} |\mu(A) - \nu(A)| = \frac{1}{2} \sum_{\omega \in \Omega} |\mu(\omega) - \nu(\omega)|.$$

The total variation distance is indeed a distance in the sense that it defines a metric on the space of probability measures on Ω . The total variation distance is related to *coupling* (and to the theory of optimal transport) by the variational formula

$$\|\mu - \nu\|_{\text{TV}} = \inf \left\{ \mathbb{P}(X \neq Y) : X, Y \text{ random variables with } X \sim \mu \text{ and } Y \sim \nu \right\}.$$

In particular, given two vertices x and y in a graph, we can couple the lazy random walks started at x and y to coincide at time m with probability $1 - \|\mathbf{P}_x(X_m = \cdot) - \mathbf{P}_y(X_m = \cdot)\|_{\text{TV}}$. Note that if the two walks coincide at time m then we can trivially couple them to remain equal at all subsequent times, so that $\|\mathbf{P}_x(X_m = \cdot) - \mathbf{P}_y(X_m = \cdot)\|_{\text{TV}}$ is a decreasing function of m when x and y are fixed. These couplings will be used when we construct annular tubes using pairs of coupled random walks.

Given a (locally finite) transitive graph $G = (V, E)$, the **Shannon entropy** H_t of the t th step of the lazy random walk is defined to be

$$H_t = -\mathbf{E}_o \left[\log p_t(o, X_t) \right] = - \sum_{x \in V} p_t(o, x) \log p_t(o, x).$$

Since the Shannon entropy of any random variable taking values in a set of size n is at most $\log n$, the quantity H_t satisfies the trivial inequality $H_t \leq \log \text{Gr}(t)$. The following extremely useful inequality¹¹ relates the total variation distance to the increments of the Shannon entropy; versions of this inequality have been rediscovered independently in the works [EK10; BDKY15; Oza18] as discussed in detail in [Yad23, Chapter 7.5].

Theorem 3.5.13. *If G is a (locally finite) transitive graph and o is a vertex of G then*

$$\frac{1}{\deg(o)} \sum_{x \sim o} \|\mathbf{P}_o(X_t = \cdot) - \mathbf{P}_x(X_{t-1} = \cdot)\|_{\text{TV}}^2 \leq H_t - H_{t-1}$$

for every $t \geq 1$.

To apply this inequality, we will need to bound the entropy in terms of the growth. While we always have the trivial bound $H_t \leq \log \text{Gr}(t)$, it is also possible to bound the entropy in terms of

¹¹known in some circles as “the cool inequality”.

$[\log \text{Gr}(t^{1/2})]^2$, which is a significantly better bound when the growth is much larger on scale t than scale $t^{1/2}$. (While this bound is worse than the trivial bound when the growth is subexponential and sufficiently regular, it better fits into our philosophy of understanding the behaviour of the random walk at some scale from the growth of the graph at that scale alone.)

Lemma 3.5.14. *For each $d \geq 1$ there exists a constant $C = C(d)$ such that if G is a (locally finite) transitive graph of degree d then*

$$H_t \leq C \left(\log \text{Gr}(t^{1/2}) \right)^2$$

for every $t \geq 1$.

Proof of Lemma 3.5.14. We may assume that the diameter of G is at least $t^{1/2}$, the claim being trivial otherwise since $H_t \leq \log |V|$. Recall that if X and Y are two random variables defined on the same probability space, the **conditional entropy** $H(X | Y)$ is defined to be the expected entropy of the conditional law of X given Y . Bayes' rule for the conditional entropy states that

$$H(X) = H(Y) + H(X|Y) - H(Y|X) \leq H(Y) + H(X|Y).$$

Applying this inequality with $X = X_t$ and $Y = \mathbb{1}(X_t \in B_r)$, and using that the entropy of a random variable supported on a set of size N is at most $\log N$, we obtain that

$$H_t \leq \log 2 + \log \text{Gr}(r) + [\log \text{Gr}(t)] \mathbf{P}_o(X_t \notin B_r) \leq \log 2 + \log \text{Gr}(r) + 2(t+1) \text{Gr}(r) \exp \left[-\frac{r^2}{2t} \right] \log \text{Gr}(t)$$

for every $r, t \geq 1$. Using the fact that the growth is submultiplicative and that $\text{Gr}(n) \leq d^{n+1}$, we obtain that if $n := \lfloor t^{1/2} \rfloor$ divides r then

$$H_t \leq \log 2 + \frac{r}{n} \log \text{Gr}(n) + 2d(t+1)^2 \exp \left[\frac{r}{n} \log \text{Gr}(n) - \frac{r^2}{2t} \right],$$

and the claim follows by taking r to be a multiple of n closest to $C \log[t \text{Gr}(n)]$ for an appropriately large constant $C = C(d)$. (Note that $\log[t \text{Gr}(n)]$ and $\log \text{Gr}(n)$ are of the same order since the diameter of G is at least $t^{1/2}$ and hence $\text{Gr}(n) \geq n$.) \square

This inequality easily implies the following simple bound on the total variation distance in terms of the growth, yielding in particular that the total variation distance is small whenever t is much larger than $d(x, y)^2 \log \text{Gr}(t^{1/2})^2$.

Corollary 3.5.15. *For each $d \geq 1$ there exists a constant $C = C(d)$ such that if $G = (V, E)$ is a (locally finite) transitive graph of degree d then*

$$\|\mathbf{P}_x(X_t = \cdot) - \mathbf{P}_y(X_t = \cdot)\|_{\text{TV}} \leq \frac{C \log \text{Gr}(t^{1/2})}{t^{1/2}} d(x, y)$$

for every $t \geq 1$ and $x, y \in V$.

Proof of Corollary 3.5.15. It follows by a standard computation that the total variation distance between a $\text{Binomial}(n, 1/2)$ distribution and a $\text{Binomial}(n+1, 1/2)$ distribution is of order $n^{-1/2}$. Indeed, if we let μ be the $\text{Binomial}(n, 1/2)$ distribution and let ν be the $\text{Binomial}(n+1, 1/2)$ distribution then μ is absolutely continuous with respect to ν with density

$$\frac{\mu(k)}{\nu(k)} = \frac{2(n-k+1)}{n+1} \quad \text{for every } 0 \leq k \leq n+1,$$

and we have by an easy computation (using e.g. Jensen's inequality and the linearity of the variance) that

$$\|\mu - \nu\|_{\text{TV}} = \frac{1}{2} \sum_{k=0}^{n+1} \left| \frac{\mu(k)}{\nu(k)} - 1 \right| \nu(k) = \frac{1}{2} \sum_{k=0}^{n+1} \left| \frac{n-2k+1}{n+1} \right| \nu(k) = O(n^{-1/2})$$

as claimed. Since the conditional laws of the lazy random walks X_t and X_{t+1} are the same given that the number of non-lazy steps are the same, it follows that

$$\|\mathbf{P}_x(X_t = \cdot) - \mathbf{P}_x(X_{t-1} = \cdot)\|_{\text{TV}} \leq C_1 t^{-1/2}$$

for every $t \geq 1$ and $x \in V$, where C_1 is a universal constant. Putting this together with Theorem 3.5.13 yields that if x and y are neighbouring vertices on a transitive graph of degree d then

$$\|\mathbf{P}_x(X_t = \cdot) - \mathbf{P}_y(X_t = \cdot)\|_{\text{TV}} \leq \sqrt{4d(H_t - H_{t-1})} + C_1 t^{-1/2}$$

for every $t \geq 1$. Since the left hand side is increasing in t and, by Lemma 3.5.14, there exists $t/2 \leq k \leq t$ with $H_k - H_{k-1} \leq \frac{2}{t} H_t \leq \frac{C_2}{t} \log \text{Gr}(t^{1/2})^2$ for some constant $C_2 = C_2(d)$, it follows that

$$\|\mathbf{P}_x(X_t = \cdot) - \mathbf{P}_y(X_t = \cdot)\|_{\text{TV}} \leq C_1 t^{-1/2} + \sqrt{\frac{4C_2 d}{t} \log \text{Gr}(t^{1/2})^2}$$

for every pair of adjacent vertices x and y and every $t \geq 1$. The analogous bound for arbitrary pairs of vertices follows from this and the triangle inequality for the total variation distance. \square

Hitting probabilities of balls. We now want to argue that tubes around independent random walks started at distant vertices are likely to be disjoint under the assumption that our graph “looks

at least five dimensional” on all relevant scales. In fact we will prove more general versions of these estimates in which “five” is replaced by an arbitrary constant $\kappa > 4$. We begin by noting the following simple analytic consequence of the results of Coulhon and Saloff-Coste [CS93] and Morris and Peres [MP05], which lets us convert growth bounds into bounds on the heat kernel $p_t(u, v)$.

Lemma 3.5.16. *For each integer $d \geq 1$ there exists a positive constant $c = c(d) \in (0, 1]$ such that if $G = (V, E)$ is an infinite unimodular transitive graph with vertex degree d then*

$$p_t(u, v) \leq \frac{1}{\text{Gr} \left(ct^{1/2} [\log \text{Gr} (t^{1/2})]^{-1/2} \right)}$$

for every integer $t \geq 4$ and every pair of vertices u, v in G .

Proof of Lemma 3.5.16. It suffices to prove an inequality of the form

$$p_{2t}(u, v) \leq \frac{1}{c \text{Gr} \left(ct^{1/2} [\log \text{Gr} (t^{1/2})]^{-1/2} \right)} \quad (3.5.2)$$

for every integer $t \geq 4$ and every $u, v \in V$, where $c = c(d)$ is a positive constant depending only on the degree. Indeed, odd values of t can then be handled using the inequality

$$p_{2t+1}(u, v) \leq \frac{1}{d} \sum_{v' \sim v} p_{2t}(u, v') \leq \max_{v' \sim v} p_{2t}(u, v'),$$

while the constant outside of the growth function can be absorbed into the constant inside the growth function using the inequality $\text{Gr}(3nm) \geq n \text{Gr}(m)$, which holds for all positive integers n, m in any infinite transitive graph as an elementary consequence of the triangle inequality.

We now prove an estimate of the form (3.5.2). Let the inverse growth function Gr^{-1} be defined by $\text{Gr}^{-1}(x) := \inf\{n : \text{Gr}(n) \geq x\}$ and recall that the **isoperimetric profile** of G is the function $\Phi : [1, \infty) \rightarrow [0, d]$ defined by

$$\Phi(x) := \inf \left\{ \frac{|\partial W|}{|W|} : W \subseteq V(G) \text{ and } 0 < |W| \leq x \right\}. \quad (3.5.3)$$

For transitive unimodular graphs, the isoperimetric profile and the growth are related by the inequality

$$\Phi(x) \geq \frac{1}{2 \text{Gr}^{-1}(2x)}, \quad (3.5.4)$$

which was proven for Cayley graphs by Coulhon and Saloff-Coste [CS93] and extended to unimodular transitive graphs by Saloff-Coste [Sal95] and Lyons, Morris, and Schramm [LMS08]; we use the statement given in [LP16c, Theorem 10.46]. To make use of this inequality, we will apply the results of Morris and Peres [MP05], which imply that there exists a constant $c_1(d) \in (0, 1)$ such that

$$p_{2t}(u, v) \leq \frac{1}{c_1 \sup \left\{ y : \int_1^y \frac{1}{x\Phi(4x)^2} dx \leq c_1 t \right\}}$$

for every integer $t \geq 1$ and every $u, v \in V$. Using (3.5.4) to estimate the integral that appears here, we have for each $y \in [1, \infty)$ that

$$\int_1^y \frac{1}{x\Phi(4x)^2} dx \leq \int_1^y \frac{4\text{Gr}^{-1}(8x)^2}{x} dx \leq 4\text{Gr}^{-1}(8y)^2 \int_1^y \frac{1}{x} dx = 4\text{Gr}^{-1}(8y)^2 \log(y).$$

Thus, to prove an estimate of the form (3.5.2) it suffices to verify that

$$\text{if } y \geq 1 \text{ satisfies } y \leq \frac{1}{8} \text{Gr} \left(\left[\frac{c_1 t}{4 \log \text{Gr}(t^{1/2})} \right]^{\frac{1}{2}} \right) \quad \text{then} \quad 4\text{Gr}^{-1}(8y)^2 \log(y) \leq c_1 t. \quad (3.5.5)$$

This follows straightforwardly by noting that $8y \leq \text{Gr}(\sqrt{c_1 t / (4 \log y)})$ whenever y satisfies the upper bound on the left hand side of (3.5.5). \square

Lemma 3.5.16 has the following elementary corollary, which we will apply only in situations where $\log \text{Gr}(t^{1/2})$ is much smaller than $n^{-1}t^{1/2}$.

Corollary 3.5.17 (Leaving a ball). *For each integer $d \geq 1$ and real number $\kappa \geq 1$ there exists a constant $C = C(d, \kappa)$ such that if $G = (V, E)$ is an infinite, connected, unimodular transitive graph with vertex degree d and $n, t \geq 1$ are integers such that $\text{Gr}(3m) \geq 3^\kappa \text{Gr}(m)$ for every $n \leq m \leq \frac{1}{2}t^{1/2}$ then*

$$\mathbf{P}_u(X_t \in B_n(v)) \leq C \left[\log \max \left\{ \frac{t}{n^2}, \text{Gr}(n) \right\} \right]^\kappa \left(\frac{n^2}{t} \right)^{\kappa/2}$$

for every pair of vertices u and v in G .

(Note that this corollary holds vacuously when $t \leq n^2$.)

Remark 3.5.7. This estimate is quite similar to that appearing in e.g. [Lyo+20]; the important distinction is that we only assume the tripling condition $\text{Gr}(3m) \geq 3^\kappa \text{Gr}(m)$ for $m = O(t^{1/2})$ rather than for all sufficiently large scales.

Proof of Corollary 3.5.17. Since $\max_u \mathbf{P}_u(X_t \in B_n(v))$ is a decreasing function of t , it suffices to prove the claim under the slightly stronger assumption that $\text{Gr}(3m) \geq 3^\kappa \text{Gr}(m)$ for every $n \leq m \leq t^{1/2}$. Fix $n, t \geq 1$ and let c be the constant from Lemma 3.5.16. We have by submultiplicativity of the growth function that

$$\text{Gr}(t^{1/2}) \leq \text{Gr}\left(ct^{1/2} \left[\log \text{Gr}(t^{1/2})\right]^{-1/2}\right)^{c^{-1}[\log \text{Gr}(t^{1/2})]^{1/2}}$$

and hence that

$$\text{Gr}\left(ct^{1/2} \left[\log \text{Gr}(t^{1/2})\right]^{-1/2}\right) \geq \text{Gr}\left(t^{1/2}\right)^{c[\log \text{Gr}(t^{1/2})]^{-1/2}} = \exp\left[c\sqrt{\log \text{Gr}(t^{1/2})}\right]. \quad (3.5.6)$$

Applying Lemma 3.5.16, it follows that if $r := ct^{1/2} \left[\log \text{Gr}(t^{1/2})\right]^{-1/2} \leq n$ then

$$p_t(u, v) \leq \frac{1}{\text{Gr}(r)} \leq \exp\left[-c^2 \frac{t^{1/2}}{n}\right]$$

for every $u, v \in V$. It follows by a union bound that if $r \leq n$ then

$$\mathbf{P}_u(X_t \in B_n(v)) \leq \exp\left[-c^2 \frac{t^{1/2}}{n}\right] \text{Gr}(n),$$

which is stronger than the desired inequality. Now suppose that $r \geq n$. The assumption $\text{Gr}(3m) \geq 3^\kappa \text{Gr}(m)$ for every $n \leq m \leq t^{1/2}$ guarantees that

$$\text{Gr}(r) \geq \text{Gr}(n) \prod_{i=1}^{\lfloor \log_3(r/n) \rfloor} \frac{\text{Gr}(3^i n)}{\text{Gr}(3^{i-1} n)} \geq \left(\frac{r}{3n}\right)^\kappa \text{Gr}(n).$$

We deduce from Lemma 3.5.16 that there exists a constant C such that

$$\begin{aligned} p_t(u, v) &\leq \frac{1}{\text{Gr}(r)} \leq \min\left\{\left(\frac{3n\sqrt{\log \text{Gr}(t^{1/2})}}{ct^{1/2}}\right)^\kappa \frac{1}{\text{Gr}(n)}, \exp\left[-c\sqrt{\log \text{Gr}(t^{1/2})}\right]\right\} \\ &\leq C \left(\frac{n^2}{t} \log \max\left\{\frac{t}{n^2}, \text{Gr}(n)\right\}\right)^{\kappa/2} \frac{1}{\text{Gr}(n)}, \end{aligned}$$

where the second inequality follows since $\min\{Ax^\kappa, e^{-x}\} \leq A(1 \vee \log(1/A))^\kappa$ for every $A, x > 0$ (as can be checked by case analysis according to whether $x \geq \log(1/A)$). \square

We next analyze the probability of hitting a ball whose radius is much smaller than its distance from the starting point. For transitive graphs of polynomial growth, a similar estimate without the logarithmic term can be proven by a similar calculation as in [Hut20i, Lemma 4.4].

Lemma 3.5.18 (Hitting a distant ball). *For each integer $d \geq 1$ and real number $\kappa > 2$ there exists a constant $C = C(d, \kappa)$ such that if G is an infinite, connected, unimodular transitive graph with vertex degree d and $n, t \geq 1$ are integers such that $t \geq n^2$ and $\text{Gr}(3m) \geq 3^\kappa \text{Gr}(m)$ for every $n \leq m \leq t^{1/2}$ then*

$$\mathbf{P}_u(\text{hit } B_n(v) \text{ before time } t) \leq C \left[\log \max \{d(u, v), \text{Gr}(2n)\} \right]^{(3\kappa+2)/2} \left(\frac{n}{d(u, v)} \right)^{\kappa-2}$$

for every pair of vertices u and v with $d(u, v) \geq 2n$.

Proof of Lemma 3.5.18. For each $1 \leq s \leq t$, let A_s be the event that the random walk hits $B_n(v)$ between times s and t . (We will optimize over the choice of s at the end of the proof.) It follows from Corollary 3.5.17 that there exist constants C_1 and C_2 depending only on d and κ such that

$$\begin{aligned} \mathbf{E}_u \left[\#\{s \leq k \leq 2t : X_k \in B_{2n}(v)\} \right] &\leq C_1 \sum_{k=s}^{2t} \left[\log \max \left\{ \frac{k}{n^2}, \text{Gr}(2n) \right\} \right]^\kappa \left(\frac{n^2}{k} \right)^{\kappa/2} \\ &\leq C_2 \left[\log \max \left\{ \frac{s}{n^2}, \text{Gr}(2n) \right\} \right]^\kappa \frac{n^\kappa}{s^{(\kappa-2)/2}}, \end{aligned} \quad (3.5.7)$$

where the second inequality follows by calculus. On the other hand, it follows from Corollary 3.5.12 and a straightforward calculation that

$$\mathbf{P}_w \left(X_{\tau+k} \in B_{2n} \text{ for every } k \leq \frac{n^2}{8 \log \text{Gr}(n)} \right) \geq \frac{1}{2}$$

for every $w \in B_n(v)$, and since $t + \frac{n^2}{8 \log \text{Gr}(n)} \leq 2t$ it follows by the strong Markov property that

$$\mathbf{E}_u \left[\#\{s \leq k \leq 2t : X_k \in B_{2n}(v)\} \mid A_s \right] \geq \frac{n^2}{16 \log \text{Gr}(n)}. \quad (3.5.8)$$

Putting together the estimates (3.5.7) and (3.5.8) yields that

$$\mathbf{P}_u(A_s) \leq \frac{\mathbf{E}_u \left[\#\{s \leq k \leq 2t : X_k \in B_{2n}(v)\} \right]}{\mathbf{E}_u \left[\#\{s \leq k \leq 2t : X_k \in B_{2n}(v)\} \mid A_s \right]} \leq C_3 \left[\log \max \left\{ \frac{s}{n^2}, \text{Gr}(2n) \right\} \right]^{\kappa+2} \frac{n^{\kappa-2}}{s^{(\kappa-2)/2}},$$

while, since every point in $B_n(v)$ has distance at least $d(u, v)/2$ from u , it follows from the Varopoulos-Carne inequality and a union bound as in the proof of Corollary 3.5.12 that

$$\mathbf{P}_u(\text{hit } B_n(v) \text{ before time } s) \leq 2s \text{Gr}(n) \exp \left[-\frac{d(u, v)^2}{8s} \right].$$

Putting together these estimates yields that

$$\mathbf{P}_u(\text{hit } B_n(v) \text{ before time } t) \leq C_3 \left[\log \max \left\{ \frac{s}{n^2}, \text{Gr}(2n) \right\} \right]^{\kappa+2} \frac{n^{\kappa-2}}{s^{(\kappa-2)/2}} + 2s \text{Gr}(n) \exp \left[-\frac{d(u, v)^2}{8s} \right],$$

and the claimed inequality follows by taking $s = c'd(u, v)^2(\log \text{Gr}(n))^{-1}$ for an appropriately small constant c' . \square

Disjoint tubes from coarse-grained random walks. As noted above, a naive construction of disjoint tubes using random walks is not appropriate for our plentiful tubes condition, since the two walks will couple at a time roughly quadratic in their starting distance rather than roughly linear. To circumvent this issue, we will instead consider tubes around certain coarse-grained versions of the random walk defined through what we call **ironing**, where we replace portions of the random walk with geodesics between their endpoints. This process will also be useful when we analyze intersections between random walk tubes, as the ironing process allows us to circumvent overcounting issues that would arise in a naive first-moment argument.

We now define the ironing procedure formally; see fig. 3.3 for an illustration. Let G be a graph. For every pair of distinct vertices $u, v \in V(G)$, fix a geodesic $\zeta(u, v)$ in G from u to v . (The choice of ζ is irrelevant to our arguments; we need only that it is done deterministically for every pair of vertices before we start running any random walks.) Fix $r > 0$ and let γ be a finite path in G . We define a sequence $(\tau_i)_{i \geq 0}$ recursively as follows: Let $\tau_0 = 0$. For each $i \geq 0$, if $d(\gamma_{\tau_i}, \gamma_k) < r$ for every $\tau_i \leq k \leq \text{len}(\gamma)$ we set $\tau_{i+1} = \text{len}(\gamma)$ and stop. Otherwise, we set τ_{i+1} to be the minimal time k after τ_i that $d(\gamma_{\tau_i}, \gamma_k) \geq r$. We define the **crease number** $\text{cr}(\gamma) = \text{cr}_r(\gamma)$ to be the number of non-zero terms in this sequence, so that $\tau_{\text{cr}(\gamma)} = \text{len}(\gamma)$, call the points $\{\gamma_{\tau_i} : 0 \leq i \leq \text{cr}_r(\gamma)\}$ **crease points**, and define the **ironed path** $\text{iron}_r(\gamma) = \text{iron}_r(\gamma)$ by concatenating geodesics between crease points

$$\text{iron}_r(\gamma) := \zeta(\gamma_{\tau_0}, \gamma_{\tau_1}) \circ \zeta(\gamma_{\tau_1}, \gamma_{\tau_2}) \circ \cdots \circ \zeta(\tau_{\text{cr}(\gamma)-1}, \tau_{\text{cr}(\gamma)}).$$

Thus, the ironed path $\text{iron}_r(\gamma)$ is a finite path in G which has the same start and end points as γ , has length at most $r \cdot \text{cr}(\gamma)$, and satisfies the containment of tubes

$$B_r(\text{iron}_r(\gamma)) \subseteq B_{2r}(\gamma) \quad \text{and} \quad B_r(\gamma) \subseteq B_{2r}(\{\gamma_{\tau_i} : 0 \leq i \leq \text{cr}_r(\gamma)\}) \subseteq B_{2r}(\text{iron}_r(\gamma)).$$

For graphs of quasi-polynomial growth, we can use the Varopoulos-Carne inequality to show that the length of an ironed random walk $\text{iron}_r(X^t)$ is of order at most $r^{-1}t(\log t)^{O(1)}$ with high probability when $r = O(\sqrt{t})$. We write X^t for the path formed by the first t steps of the random walk.

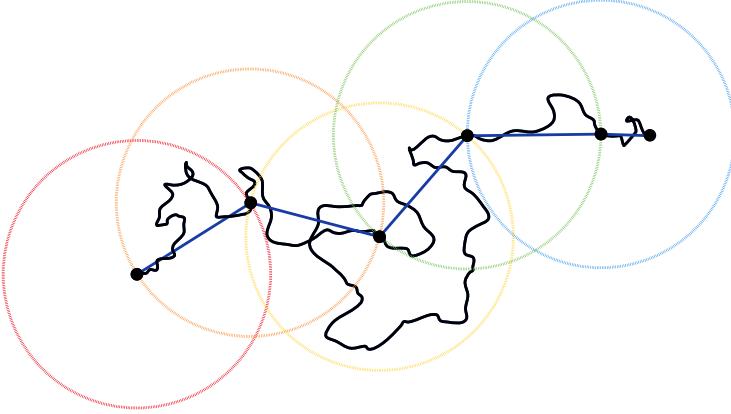


Figure 3.3: Schematic illustration of the ironing procedure applied to a path: The original path is in black, from left to right. The black dots are crease points. A new crease point is formed every time the path leaves a fixed-radius ball centered at the previous crease point. The straight blue line segments are geodesics between consecutive crease points (together with the final point of the walk). The ironed path is formed by concatenating these geodesics.

Lemma 3.5.19. *Let G be a (locally finite) transitive graph and let u be a vertex of G . Then*

$$\mathbf{P}_u \left(\text{cr}_r(X^t) > \frac{t}{m} \right) \leq 2tm \text{Gr}(r) \exp \left[-\frac{r^2}{2m} \right]$$

for every $r, t, \lambda \geq 1$.

Proof. In order for the inequality $\text{cr}_r((X_i)_{i=0}^t) > t/m$ to hold, there must exist $0 \leq i \leq t$ such that $d(X_i, X_{i+m}) \geq r$. As such, the claim follows from Corollary 3.5.12 and a union bound. \square

We now analyze intersections between independent random walk tubes. Given two vertices u and v , we write $\mathbf{P}_u \otimes \mathbf{P}_v$ for the law of a pair of independent lazy random walks X and Y started at u and v respectively.

Lemma 3.5.20 (Intersections of random walk tubes). *For each integer $d \geq 1$ and real number $\kappa > 4$ there exists a constant $C = C(d, \kappa)$ such that if $G = (V, E)$ is an infinite, connected, unimodular transitive graph with vertex degree d and $n, t \geq 1$ are integers such that $t \geq r^2$ and $\text{Gr}(3m) \geq 3^\kappa \text{Gr}(m)$ for every $r \leq m \leq t^{1/2}$ then*

$$\begin{aligned} \mathbf{P}_u \otimes \mathbf{P}_v & \left(\text{there exist } 0 \leq i, j \leq t \text{ such that } d(X_i, Y_j) \leq r \right) \\ & \leq C \frac{t}{d(u, v)^2} \left[\log \max \{d(u, v), \text{Gr}(4r)\} \right]^{(3\kappa+4)/2} \left(\frac{r}{d(u, v)} \right)^{\kappa-4}, \end{aligned}$$

for every $u, v \in V$ with $d(u, v) \geq 4r$.

Remark 3.5.8. Although it will suffice for all our applications, we note that this bound is very wasteful when t is much larger than $d(u, v)^2$. A more careful analysis would use that Y_{τ_k} is typically very far from u when k is large.

Proof of Lemma 3.5.20. Let A be the event that there exists $w \in V$ with $d(u, w) \geq \frac{1}{2}d(u, v)$ and $1 \leq i, j \leq t$ such that $d(X_i, w), d(Y_j, w) \leq r$. It suffices by symmetry to prove that there exists a constant $C = C(d, \kappa)$ such that

$$\mathbf{P}_u \otimes \mathbf{P}_v(A) \leq C \frac{t}{r^2} [\log \max \{d(u, v), \text{Gr}(4r)\}]^{(3\kappa+4)/2} \left(\frac{r}{d(u, v)} \right)^{\kappa-2}.$$

Let $\tau_0, \dots, \tau_{\text{cr}_r(Y^t)}$ be the stopping times used to define the ironed walk $\text{iron}_r(Y^t)$, and observe that if A holds then there must exist $0 \leq k \leq \text{cr}_r(Y^t)$ and $0 \leq i \leq t$ such that

$$d(Y_{\tau_k}, v) \geq \frac{1}{2}d(u, v) - r \geq \frac{1}{4}d(u, v) \quad \text{and} \quad d(X_i, Y_{\tau_k}) \leq 2r.$$

Thus, it follows from Lemma 3.5.18 and a union bound that there exists a constant $C_1 = C_1(d, \kappa)$ such that

$$\begin{aligned} \mathbf{P}_u \otimes \mathbf{P}_v(A) &\leq \mathbf{P}_v\left(\text{cr}_r(Y^t) > \frac{t}{m}\right) + \frac{t}{m} \max \left\{ \mathbf{P}_u\left(\text{hit } B_{2r}(w) \text{ before time } t\right) : w \in V, d(u, w) \geq \frac{1}{4}d(u, v) \right\} \\ &\leq 2tm \text{Gr}(r) \exp\left[-\frac{r^2}{2m}\right] + C_1 \frac{t}{m} [\log \max \{d(u, v), \text{Gr}(4r)\}]^{(3\kappa+2)/2} \left(\frac{r}{d(u, v)} \right)^{\kappa-2} \end{aligned}$$

and the claim follows by taking $m = c'r^2(\log \max\{d(u, v), \text{Gr}(r)\})^{-1}$ for an appropriately small constant $c' = c'(d, \kappa)$. \square

We now have everything we need to prove Proposition 3.5.4.

Proof of Proposition 3.5.4. We will prove the claim concerning annular tubes (which is harder); the changes to the proof needed to establish the claim concerning radial tubes are straightforward and will be explained briefly at the end of the proof.

We will prove a general condition for G to have (k, r, ℓ) -plentiful annular tubes on scale n in terms of the growth of G , which we specialize to give the claim about scales of quasi-polynomial growth at the end of the proof. Fix $n \geq 1$, $k \leq n/2$, $r \leq n$, $t \geq n$ and two $(n, 3n)$ crossings A and B as in

the definition of plentiful annular tubes, so that A and B each contain points at every distance from o between n and $3n$. We will carry out our analysis under the hypotheses that

$$\text{Gr}(3m) \geq 3^\kappa \text{Gr}(m) \quad \text{for every } r \leq m \leq t \quad (3.5.9)$$

and

$$\|\mathbf{P}_x(X_t = \cdot) - \mathbf{P}_y(X_t = \cdot)\|_{\text{TV}} \leq \frac{1}{4} \quad \text{for every } x, y \in B_{3n}. \quad (3.5.10)$$

We will return to the question of when suitable t and r satisfying these hypotheses can be chosen at the end of the proof.

Since A and B each contain at least one point at each distance from o between n and $3n$, and since $n/2k \geq 1$, we may choose for each $1 \leq i \leq k$ points $a_i \in A$ and $b_i \in B$ such that

$$n + (4i-4)\frac{n}{2k} \leq d(o, a_i) \leq n + (4i-3)\frac{n}{2k} \quad \text{and} \quad n + (4i-2)\frac{n}{2k} \leq d(o, b_i) \leq n + (4i-1)\frac{n}{2k}$$

for every $1 \leq i \leq k$, so that the set of points $\{a_i\} \cup \{b_i\}$ is (n/k) -separated (i.e., any two distinct points in the set have distance at least n/k). For each i , let \mathbf{Q}_i be the joint law of a random walk $(X_{i,m})_{m \geq 0}$ started at a_i and a random walk $(Y_{i,m})_{m \geq 0}$ started at b_i , coupled so that $X_t = Y_t$ with probability at least $3/4$ (such a coupling exists by the hypothesis (3.5.10) imposed on the value of t), and let $\mathbf{Q} = \bigotimes \mathbf{Q}_i$ be the law of the collection $\{X_{i,m}, Y_{i,m} : 1 \leq i \leq k, m \geq 0\}$ in which the pairs $((X_{i,m})_{m \geq 0}, (Y_{i,m})_{m \geq 0})$ are sampled independently for each $1 \leq i \leq k$. For each $1 \leq i \leq k$, consider the events

$$\mathcal{A}_i = \{X_{i,t} = Y_{i,t}\} \quad \text{and} \quad \mathcal{B}_i = \left\{ \text{len}(\text{iron}_r(X^t)), \text{len}(\text{iron}_r(Y^t)) \leq \frac{16t}{r} \log t \text{Gr}(r) \right\}.$$

The event \mathcal{A}_i has probability at least $3/4$ for every $1 \leq i \leq k$ by construction. Meanwhile, Lemma 3.5.19 implies that

$$\begin{aligned} \mathbf{Q}(\mathcal{B}_i^c) &\leq 2\mathbf{P}_u \left(\text{cr}_r(X^t) > \frac{16t}{r^2} \log \max\{t, \text{Gr}(r)\} \right) \\ &\leq 4tr^2 [\log \max\{t, \text{Gr}(r)\}] \text{Gr}(r) \exp[-8 \log \max\{t, \text{Gr}(r)\}], \end{aligned}$$

which is less than $1/4$ if t is larger than some universal constant t_0 . Now, for each $1 \leq i, j \leq k$ let

$$\mathcal{I}_{i,j} = \{B_{2r}(X_i^t) \cup B_{2r}(Y_i^t) \text{ has non-empty intersection with } B_{2r}(X_j^t) \cup B_{2r}(Y_j^t)\}.$$

We have by a union bound that if i and j are distinct then

$$\mathbf{Q}(\mathcal{I}_{i,j}) \leq 4 \max \left\{ \mathbf{P}_u \otimes \mathbf{P}_v (B_{2r}(X^t) \cap B_{2r}(Y^t) \neq \emptyset) : d(u, v) \geq \frac{n}{k} \right\}$$

$$\leq C \frac{tk^2}{n^2} [\log \max\{n, \text{Gr}(8r)\}]^{(3\kappa+4)/2} \left(\frac{rk}{n} \right)^{\kappa-4} =: \alpha = \alpha(n, t, k, r).$$

Thus, if for each $1 \leq i \leq k$ we define the event

$$C_i = \{\text{there exist at most } 4\alpha k \text{ values of } 1 \leq j \leq k \text{ with } j \neq i \text{ such that } \mathcal{I}_{i,j} \text{ holds}\}$$

then $\mathbf{Q}(C_i) \geq \frac{3}{4}$ by Markov's inequality. It follows by a union bound that $\mathbf{Q}(\mathcal{A}_i \cap \mathcal{B}_i \cap C_i) \geq 1/4$, and hence that if we define

$$\mathcal{E} = \#\{1 \leq i \leq k : \mathcal{A}_i \cap \mathcal{B}_i \cap C_i \text{ holds}\} \geq \frac{k}{4} \quad \text{then} \quad \mathbf{Q}(\mathcal{E}) > 0.$$

Consider the random set of indices

$$I = \{1 \leq i \leq k : \mathcal{A}_i \cap \mathcal{B}_i \cap C_i \text{ holds but } \mathcal{I}_{i,j} \text{ does not hold for any } j < i \text{ for which } \mathcal{A}_j \cap \mathcal{B}_j \cap C_j \text{ holds}\}.$$

On the event \mathcal{E} , the set I has size at least $\lceil (k/4)/(1 + 4\alpha k) \rceil \geq \frac{1}{20} \min\{k, \alpha^{-1}\}$. Moreover, on this event, the set of paths formed by concatenating $\text{iron}_r(X_i^t)$ and the reversal of $\text{iron}_r(Y_i^t)$ for each $i \in I$ have the property that each such path has length at most $32t/r \log t \text{Gr}(r)$, and the tubes of thickness r around distinct such paths are disjoint. Since this event has positive probability, there must exist a set of paths with this property. Since the sets A and B were arbitrary, it follows that there exists a positive constant $c = c(d, \kappa)$ such that G has

$$\left(c \min \left\{ k, \frac{n^2}{tk^2} [\log \max\{n, \text{Gr}(8r)\}]^{-(3\kappa+4)/2} \left(\frac{n}{rk}\right)^{\kappa-4} \right\}, r, \frac{32t}{r} \log \max\{t, \text{Gr}(r)\} \right) \text{-plentiful}$$

annular tubes on scale n for each $1 \leq r, k \leq n/2$ and $t \geq n^2$ satisfying (3.5.9) and (3.5.10).

We now specialize to the setting of the proposition. Let $D, \lambda \geq 1$ and $0 < \varepsilon < 1$ and suppose that $n \geq 1$ satisfies $\text{Gr}(m) \leq e^{(\log m)^D}$ and $\text{Gr}(3m) \geq 3^5 \text{Gr}(m)$ for every $n^{1-\varepsilon} \leq m \leq n^{1+\varepsilon}$. Let $k = (\log n)^\lambda$, $r = n(\log n)^{-20D\lambda}$, and $t = n^2(\log n)^{2D}$. There exists $n_0 = n_0(d, D, \lambda, \varepsilon)$ such that if $n \geq n_0$ then $n^{1-\varepsilon} \leq r, t^{1/2} \leq n^{1+\varepsilon}$, so that (3.5.9) holds. Moreover, since $\log \text{Gr}(t^{1/2}) \leq (\log t^{1/2})^D \leq (1+\varepsilon)^2(\log n)^D$, it follows from Corollary 3.5.15 that (3.5.10) holds whenever $n \geq n_1$ for some constant $n_1 = n_1(d, D, \lambda, \varepsilon) \geq n_0$. It follows that there exists a constant C such that if $n \geq n_1$ then G has

$$\left(c \min \left\{ (\log n)^\lambda, (\log n)^{(20\lambda - \frac{21}{2})D - \lambda} \right\}, n(\log n)^{-20D\lambda}, C(\log n)^{(3+20\lambda)D} \right) \text{-plentiful}$$

annular tubes on scale n . Since the plentiful annular tubes condition is monotone increasing in the number and thickness parameters and monotone decreasing in the length parameter, it follows that there exists a constant $n_2 = n_2(d, D, \lambda, \varepsilon) \geq n_1$ such that if $n \geq n_2$ then G has $(1/40D, \lambda)$ -polylog-plentiful annular tubes on scale n as claimed (where we used the assumption that n is large to absorb the constant prefactors into the exponents of the logarithms).

For the claim concerning *radial* tubes, one uses random walks started at k equidistant points along a geodesic from o to S_n . Corollary 3.5.17 implies that each of these random walks has a good probability not to belong to B_{3n} at times of order $n^2(\log \text{Gr}(n))^\kappa$, so that we can obtain tubes from S_n to S_{3n} using segments of the resulting ironed paths. The analysis of the number, thickness, and lengths we can take these walks to have with the resulting tubes being disjoint is similar to the annular case above. (Indeed, the analysis is somewhat simpler since we use single walks instead of coupled pairs of walks. This also makes the dependence on the growth better than in the annular case.) \square

3.6 Quasi-polynomial growth II: Analysis of percolation

In this section, we analyze percolation in the low-growth regime, i.e., on scales n where $\text{Gr}(n) \leq e^{[\log n]^D}$ for a constant D . Our goal is to show that a lower bound on (full-space) connection probabilities at some low-growth scale n implies a lower bound on connection probabilities within a tube *at a much larger scale* after sprinkling. This analysis will employ both the polylog-plentiful tubes condition from Proposition 3.5.2 and the outputs of the ghost field technology developed in Proposition 3.4.1.

Recall that \mathcal{U}_d^* is the space of infinite, connected, unimodular transitive graphs of degree d that are not one-dimensional.

Proposition 3.6.1 (Inductive analysis of percolation in the low-growth regime.). *For each $d, D \geq 1$ there exist positive constants $\lambda_0 = \lambda_0(d, D)$ and $c = c(d, D)$ such that for each $\lambda \geq \lambda_0$ there exist constants $K_1 = K_1(d, D, \lambda)$ and $n_0 = n_0(d, D, \lambda)$ such that if $K \geq K_1$ and $n \geq n_0$ then the following holds: If $G \in \mathcal{U}_d^*$ and for each $1 \leq b \leq n$ we define*

$$\delta(b, n) = \delta_K(b, n) = \left(\frac{K \log \log n}{\min\{\log n, \log \text{Gr}(b)\}} \right)^{1/4}$$

then the implication

$$\begin{aligned} & \left(n \in \mathcal{L}(G, D), \kappa_{p_1}(n, \infty) \geq e^{-(\log \log n)^{1/2}}, \mathbb{P}_{p_1}(\text{Piv}[4b, n^{1/3}]) \leq (\log n)^{-1}, \text{ and } \delta(b, n) \leq 1 \right) \\ & \Rightarrow \left(p_2 \geq p_c \text{ or } \kappa_{p_2} \left(e^{(\log n)^{c\lambda}}, n \right) \geq e^{-3(\log \log n)^{1/2}} \right) \quad (3.6.1) \end{aligned}$$

holds for every $n \geq n_0$, $b \leq \frac{1}{8}n^{1/3}$, and $p_2 \geq p_1 \geq 1/d$ with $p_2 \geq \text{Spr}(p_1, \delta(b, n))$.

The parameter λ controls how far a connection probability lower bound is propagated by sprinkling, namely from scale n to scale $e^{(\log n)^{c\lambda}}$, where c is independent of λ . Proposition 3.6.1 morally says

that for any choice of λ , if we sprinkle by a sufficient amount at a sufficiently large scale, we can achieve this propagation from scale n to scale $e^{(\log n)^{c\lambda}}$. (Note that since the conclusion of Proposition 3.6.1 gets stronger as λ increases, the $\lambda \geq \lambda_0$ condition is redundant; we have nevertheless included it because some of our working is simpler if we can assume that λ is large.)

The overall strategy of this section is closely inspired by [CMT22]. A key insight of that paper was that if one is working in the supercritical phase of percolation, then one can sometimes deduce positive information (e.g. a lower bound on set-to-set connection probabilities) from negative information (e.g. an upper bound on set-to-set connection probabilities). Thus, one can analyze the supercritical phase by a case analysis according to whether or not certain point-to-point connection lower bounds hold: both assumptions are useful. This only makes sense because being in the supercritical phase is already a positive hypothesis. For our purposes, we cannot assume that we are in the supercritical phase. However, we can assume that we have a two-point lower bound at some large scale n , which lets us pretend that we are supercritical when working with scales much smaller than n . The positive information that this lets us deduce by working with scales smaller than n is so strong that it implies set-to-set connection lower bounds even at scales much larger than n .

Besides this, there are two main complications we need to address when adapting the methods of [CMT22] in this section. First, our quasi-polynomial growth and plentiful tubes conditions are rather different than the polynomial growth and connectivity of exposed spheres used in [CMT22], and many details of the argument must change to accommodate this. Second, and more seriously, we must use methods that use the growth upper bound *at one scale only*, since that is all we can assume; this is very different from [CMT22] where their graphs are assumed to have polynomial growth at all scales and many of the arguments work by inducting up from scale 1.

We now begin to work towards the proof of Proposition 3.6.1. We begin by recording various notations and important constants that will be used throughout the proof. First, we let $\mathcal{A}(b, n, p_1) = \mathcal{A}_{D,K}(b, n, p_1; G)$ be the statement on the left hand side of the implication (3.6.1):

$$\mathcal{A}(b, n, p_1) = \left(n \in \mathcal{L}(G, D), \kappa_{p_1}(n, \infty) \geq e^{-(\log \log n)^{1/2}}, \mathbb{P}_{p_1}(\text{Piv}[4b, n^{1/3}]) \leq (\log n)^{-1} \text{ and } \delta(b, n) \leq 1 \right).$$

We will always assume that $\delta(p_1, p_2)$ is larger than the quantity $\delta := \delta(b, n) = \delta_K(b, n)$ introduced in Proposition 3.6.1. We regard the choices of $d \in \mathbb{N}$, $D \geq 1$, $K \geq 1$, and $\lambda \geq 20$, as well as the graph $G \in \mathcal{U}_d^*$, as being fixed for the remainder of the section. We write $p_{3/2} := \text{Spr}(p_1; \delta(p_1, p_2)/2)$, so that $p_2 = \text{Spr}(p_{3/2}; \delta(p_1, p_2)/2)$ by the semigroup property of our sprinkling operation, and let $c_1 = c_1(d, D) > 0$ and $N = N(d, D, \lambda)$ be the constants guaranteed to exist by Proposition 3.5.2.

Inspired by [CMT22], we will split the proof of Proposition 3.6.1 into two cases according to how easy it is to connect points at certain well-chosen intermediate scales. For each $1 \leq m \leq n$ we define the **two-point zone**

$$\text{tz}(m) = \text{tz}(m, n) := \sup \left\{ r \geq 0 : \tau_{p_{3/2}}^{B_m}(B_r) \geq (\log n)^{-1} \right\},$$

where we recall that $\tau_p^B(A)$ is the quantity defined by $\tau_p^B(A) := \inf_{x, y \in A} \mathbb{P}_p(x \xrightarrow{B} y)$. We stress that the two-point zone $\text{tz}(m)$ is defined using the intermediate parameter $p_{3/2}$. Knowing the value of $\text{tz}(m)$ for some m provides us with both positive information (on scales smaller than $\text{tz}(m)$) and negative information (on scales larger than $\text{tz}(m)$).

For each $n \in \mathcal{L}(G, D)$ with $n \geq N$, let $n^{1/3} \leq m_1 \leq m_2 \leq n^{1/(1+c_1)}$ be such that $m_2 \geq m_1^{1+c_1}$ and G has (c_1, λ) -polylog-plentiful tubes at every scale $m_1 \leq m \leq m_2$ (such m_1 and m_2 existing by Proposition 3.5.2), and let $\mathcal{S}(n) = \mathcal{S}_{d, D, \lambda}(n) = \{m \in \mathbb{N} : m_1^{1+(c_1/4)} \leq m \leq m_1^{1+(3c_1/4)}\}$. From now on we will mostly work at scales $m \in \mathcal{S}(n)$, so that G has plentiful tubes not just at these scales but at a large range of consecutive scales on either side of every such scale. Let $c_2 = c_2(d)$ be the minimum of the four constants appearing in Proposition 3.4.1 with ‘ D ’ equal to 1 (known in the statement of that proposition as c_1, c_2, c_3 , and h_0 ; the proposition becomes weaker if we replace all four constants by their minimum), define $c_3 = c_3(d) := 2^{-9}c_2$, and consider the statement

$$\begin{aligned} \mathcal{B}(n, p_{3/2}) &= \mathcal{B}_{d, D, \lambda, K}(n, p_{3/2}) \\ &= \left(\text{there exists } m \in \mathcal{S}(n) \text{ such that } \text{tz} \left(\frac{m}{2} \right) [\log n]^{c_3 K} \geq m (\log n)^{3\lambda/c_1} \right). \end{aligned}$$

We think of $\mathcal{B}(n, p_{3/2})$ as our ‘‘positive assumption’’ about percolation on scale n : it means that there is a ‘‘good’’ scale $m \in \mathcal{S}(n)$ such that points in a ball of radius not much smaller than m can be connected with reasonable probability within the ball of radius $m/2$ when $p = p_{3/2}$.

Notational conventions and standing assumptions: **Recall that the choices of $d \in \mathbb{N}$, $D \geq 1$, $K \geq 1$, $\lambda \geq 20$, and $G \in \mathcal{U}_d^*$ are considered to be fixed for the remainder of the section. The constants N, c_1, c_2 , and c_3 used to define $\mathcal{S}(n)$ and $\mathcal{B}(n, p)$ will be used with the same meaning throughout this section. We also define $c_{-1} = c_{-1}(d)$ to be a positive constant such that $\text{Spr}(p; \delta) \geq p + c_{-1}\delta$ for every $p \geq 1/d$ and $0 < \delta \leq 1$, which exists by (3.3.1). Finally, we fix $K_0 = K_0(d, D) = 2^{11}(c_3^{-1} \vee c_{-1}^{-5} \vee 1)$ throughout the section: all subsequent lemmas in this section will include $K \geq K_0$ as an implicit hypothesis. These conventions do not apply outside of this section (Section 3.6).**

Proposition 3.6.1 follows trivially from the following two lemmas.

Lemma 3.6.2 (Concluding from a positive assumption). *There exists a constant $\lambda_0 = \lambda_0(d, D)$ such that if $\lambda \geq \lambda_0$ then there exists a constant $n_0 = n_0(d, D, \lambda) \geq N$ such that the implication*

$$\mathcal{A}_{D,K}(b, n, p_1) \wedge \mathcal{B}_{d,D,\lambda,K}(n, p_{3/2}) \Rightarrow \left(p_2 \geq p_c \text{ or } \kappa_{p_2} \left(e^{(\log n)^{c_1 \lambda/4}}, n \right) \geq e^{-3(\log \log n)^{1/2}} \right) \quad (3.6.2)$$

holds for every $n \geq n_0$, $b \leq \frac{1}{8}n^{1/3}$, and every pair of probabilities $p_2 \geq p_1 \geq 1/d$ with $\delta(p_1, p_2) \geq \delta_K(b, n)$.

Lemma 3.6.3 (Concluding from a negative assumption). *There exists a constant $\lambda_0 = \lambda_0(d, D)$ such that if $\lambda \geq \lambda_0$ then there exist constants $K_1 = K_1(d, D, \lambda) \geq K_0$ and $n_0 = n_0(d, D, \lambda) \geq N$ such that if $K \geq K_1$ and $n \geq n_0$ then the implication*

$$\mathcal{A}_{D,K}(b, n, p_1) \wedge (\neg \mathcal{B}_{d,D,\lambda,K}(n, p_{3/2})) \Rightarrow \left(p_2 \geq p_c \text{ or } \kappa_{p_2} \left(e^{(\log n)^{c_1 \lambda/4}}, n \right) \geq e^{-3(\log \log n)^{1/2}} \right) \quad (3.6.3)$$

holds for every $n \geq n_0$, $b \leq \frac{1}{8}n^{1/3}$, and every pair of probabilities $p_2 \geq p_1 \geq 1/d$ with $\delta(p_1, p_2) \geq \delta_K(b, n)$.

Concluding from a positive assumption

Our next goal is to prove Lemma 3.6.2. We begin with the following lemma, which will later allow us to use the positive assumption $\mathcal{B}(n, p)$ to deduce lower bounds on the corridor function after sprinkling. Note that we are not yet using the polylog-plentiful tubes condition or the assumption $\mathcal{B}(n, p)$, so that the parameter λ does not appear in this lemma. (Recall our standing assumption throughout this section that $K \geq K_0 = K_0(d, D)$.)

We define $p_{7/4} = \text{Spr}(p_1; \frac{3}{4}\delta(p_1, p_2)) = \text{Spr}(p_{3/2}; \frac{1}{4}\delta(p_1, p_2))$, and remind the reader that the two-point zone $\text{tz}(m)$ was defined with respect to the parameter $p_{3/2}$.

Lemma 3.6.4. *There exist a constant $n_0 = n_0(d, D) \geq N$ and a universal constant c_4 such that if $n \geq n_0$ then the implication*

$$\mathcal{A}_{D,K}(b, n, p_1) \Rightarrow \left(\kappa_{p_{7/4}} \left(\text{tz} \left(\frac{m}{2} \right) [\log n]^{c_3 K}, m \right) \geq c_4 e^{-2(\log \log n)^{1/2}} \text{ for every } n^{1/3} \leq m \leq n \right)$$

holds for every $n \geq n_0$, $b \leq \frac{1}{8}n^{1/3}$, and $p_1 \geq 1/d$.

Proof of Lemma 3.6.4. Fix $n^{1/3} \leq m \leq n$ and a path γ starting at some vertex u , ending at some vertex v , and with $\text{len } \gamma \leq \text{tz} \left(\frac{m}{2} \right) (\log n)^{c_3 K}$. Pick a subsequence $u = u_1, u_2, \dots, u_k = v$ of γ with $k \leq 5(\log n)^{c_3 K}$ and $\text{dist}(u_i, u_{i+1}) \leq \frac{1}{4} \text{tz} \left(\frac{m}{2} \right)$ for every $1 \leq i < k$. (There are rounding issues that may prevent such a sequence existing when $\text{tz} \left(\frac{m}{2} \right)$ is small, but this is not a problem when n is

large.) We claim that $\frac{1}{4} \text{tz} \left(\frac{m}{2} \right) \geq b$. Indeed, the hypothesis $b \leq \frac{1}{8} n^{1/3}$ guarantees that $4b \leq m/2$, and we have by a union bound that

$$\begin{aligned} \tau_{p_{3/2}}^{B_{m/2}}(B_{4b}) &\geq \tau_{p_1}^{B_{m/2}}(B_{4b}) \geq \kappa_{p_1}(n, \infty) - \mathbb{P}_{p_1}(\text{Piv}[4b, n^{1/3}]) \\ &\geq e^{-(\log \log n)^{1/2}} - [\log n]^{-1} \geq \frac{1}{2} e^{-[\log \log n]^{1/2}} \end{aligned} \quad (3.6.4)$$

for every n larger than some universal constant. It follows in particular that $\tau_{p_{3/2}}^{B_{m/2}}(B_{4b}) \geq (\log n)^{-1}$ for every n larger than some universal constant n_0 , and hence that $\frac{1}{4} \text{tz} \left(\frac{m}{2} \right) \geq b$ when $n \geq n_0$ by maximality of $\text{tz} \left(\frac{m}{2} \right)$.

Let $h = c_2 \max\{\text{Gr}(b)^{-1}, n^{-1/15}\}$, so that if $n_0 \geq 3$ (to guarantee that $n^{1/15} \geq \log n$) then

$$h^{c_2(\delta/4)^3} \leq h^{c_2(\delta/4)^4} \leq \exp \left(\log h \cdot \frac{c_2}{2^8} \cdot \frac{K \log \log n}{\log h} \right) = (\log n)^{-2c_3 K} \leq (\log n)^{-2} \quad (3.6.5)$$

by the definition of $\delta_K(b, n)$ and the assumption that $K \geq K_0$. We want to apply Proposition 3.4.1 where ‘ n ’ is k , the sets ‘ A_i ’ are the balls $(B_b(u_i))_{i=1}^k$, the superset ‘ Λ ’ is the tube $B_{m/2}(\gamma)$, the ghost field intensity ‘ h ’ is equal to h , ‘ p_1 ’ is $p_{3/2}$, and the sprinkling amount ‘ δ ’ is $\delta/4$. To do this, it suffices to verify that if n is sufficiently large then

$$h^{c_2(\delta/4)^3} \leq c_2 k^{-1} \quad \text{and} \quad \tau_{p_{3/2}}^{B_{m/2}(\gamma)}(B_b(u_i) \cup B_b(u_{i+1})) \geq 4h^{c_2(\delta/4)^4} \quad \text{for every } 1 \leq i \leq k-1.$$

(The assumption that h is sufficiently small holds automatically since we defined c_2 to be the minimum of the constants appearing in Proposition 3.4.1 and set $h \leq c_2$.) The inequality (3.6.5) implies that the first required inequality $h^{c_2(\delta/4)^3} \leq c_2 k^{-1}$ holds whenever n is larger than some constant depending only on d , since $h^{c_2(\delta/4)^3} \leq (\log n)^{-2c_3 K}$, $k \leq 5(\log n)^{c_3 K}$ and $c_3 K \geq 2$. For the second required inequality, note that for every such i and for every $u \in B_b(u_i)$ and $v \in B_b(u_{i+1})$, we have $u, v \in B_{\text{tz}(m/2)-1}(u_i)$ and hence that

$$\tau_{p_{3/2}}^{B_{m/2}(\gamma)}(B_b(u_i) \cup B_b(u_{i+1})) \geq \tau_{p_{3/2}}^{B_{m/2}}(B_{\text{tz}(m/2)-1}) \geq (\log n)^{-1}$$

for all n larger than some universal constant by (3.6.4). Let $n_0 = n_0(d, D) \geq N$ be the maximum of these two constants and N .

Let r be the minimum positive integer such that $\mathbb{P}_q(\text{Piv}[1, rh]) < h$ for every $q \in [p_{3/2}, \text{Spr}(p_{3/2}; \delta/4)]$. The previous paragraph shows that if $n \geq n_0$ then the hypotheses of Proposition 3.4.1 are met. Thus, since $p_{7/4} \geq \text{Spr}(p_{3/2}; \delta/4)$, we obtain from that proposition that for a universal constant $c > 0$,

$$\mathbb{P}_{p_{7/4}} \left(u \xrightarrow{B_{m/2+2r}(\gamma)} v \right) \geq c \tau_{p_{3/2}}^{B_{m/2}(\gamma)}(B_b(u)) \cdot \tau_{p_{3/2}}^{B_{m/2}(\gamma)}(B_b(v)) \geq \frac{c}{4} e^{-2(\log \log n)^{1/2}}.$$

Since γ was arbitrary, this implies that if $n \geq n_0$ then

$$\kappa_{p_{7/4}} \left(\text{tz} \left(\frac{m}{2} \right) [\log n]^{c_3 K}, \frac{m}{2} + 2r \right) \geq \frac{c}{4} e^{-2(\log \log n)^{1/2}}.$$

All that remains is to verify that $2r \leq m/2$ when n is sufficiently large.

Let $q \in [p_{3/2}, \text{Spr}(p_{3/2}; \delta/4)]$ be arbitrary. Since $2/5 < 1/2$, Proposition 3.2.1 yields that there exists a constant C_d such that

$$\mathbb{P}_q (\text{Piv}[1, (m/4)h]) \leq C_d \left[\frac{\log \text{Gr}((m/4)h)}{(m/4)h} \right]^{2/5} \leq C_d \left[\frac{4 \log \text{Gr}(m)}{mh} \right]^{2/5}.$$

Since $h \geq n^{-1/15}$, $n^{1/3} \leq m \leq n$, $\log \text{Gr}(m) \leq (\log m)^D$, and $(2/5) \cdot (4/15) > 1/15$, it follows that there exists a constant $n_2 = n_2(d, D)$ such that if $n \geq n_2$ then

$$\mathbb{P}_q (\text{Piv}[1, (m/4)h]) \leq C_d \left[\frac{4(\log n)^D}{n^{1/3} \cdot n^{-1/15}} \right]^{2/5} \leq n^{-1/15} \leq h,$$

and hence that $r \leq m/4$ as claimed. This completes the proof. \square

We now apply Lemma 3.6.4 together with the polylog-plentiful tubes condition to prove Lemma 3.6.2.

Proof of Lemma 3.6.2. Let $n_0 = n_0(d, D)$ and the universal constant c_4 be as in Lemma 3.6.4. Suppose that $n \geq n_0$, $b \leq \frac{1}{8}n^{1/3}$, and $p_1 \in [1/d, 1]$ are such that $\mathcal{A}_{D,K}(b, n, p_1)$ holds, and write $\delta = \delta(b, n)$. We have by Lemma 3.6.4 that

$$\kappa_{p_{7/4}} \left(\text{tz} \left(\frac{m}{2} \right) [\log n]^{c_3 K}, m \right) \geq c_4 e^{-2(\log \log n)^{1/2}} \text{ for every } n^{1/3} \leq m \leq n$$

and in particular that if $\mathcal{B}_{d,D,\lambda,K}(n, p)$ holds then there exists $m \in \mathcal{S}(n) = \mathcal{S}_{d,D,\lambda}(n)$ such that

$$\kappa_{p_{7/4}}(m(\log n)^{3\lambda/c_1}, m) \geq (\log n)^{-1} \tag{3.6.6}$$

whenever n_0 is larger than a suitable universal constant. (Note that we have not yet put any restrictions on the parameter λ .)

Suppose that $m \in \mathcal{S}(n)$ satisfies (3.6.6) and let $r = m[\log n]^{3\lambda/(2c_1)}$, so that if $n \geq 2$ we have the inclusion of intervals

$$\left[\frac{9}{10}r(\log r)^{-\lambda/c_1}, r(\log r)^{\lambda/c_1} \right] \subseteq \left[m, m(\log n)^{3\lambda/c_1} \right]. \tag{3.6.7}$$

Since $\frac{9}{10}r \in [m_1, m_2]$ (where m_1 and m_2 are as in the definition of $\mathcal{S}(n)$ above), G has (c_1, λ) -polylog-plentiful radial tubes at scale $\frac{9}{10}r$. Let Γ be a family of paths witnessing this fact (which

cross the annulus from $S_{0.9r}$ to $S_{4 \cdot (0.9r)} = S_{3.6r}$) and let $(T_\gamma)_{\gamma \in \Gamma}$ be the associated family of tubes given by $T_\gamma := B(\gamma, \frac{9}{10}r[\log(\frac{9}{10}r)]^{-\lambda/c_1})$. Since each tube T_γ has thickness at least m and length $\text{len } \gamma \leq r(\log r)^{\lambda/c_1} \leq m(\log n)^{3\lambda/c_1}$, it follows from (3.6.6) that

$$\mathbb{P}_{p_{7/4}}(S_{0.9r} \xleftrightarrow{T_\gamma} S_{3.6r}) \geq (\log n)^{-1}$$

for every $\gamma \in \Gamma$ and hence that there exists a constant $n_1 = n_1(d, D, \lambda) \geq n_0$ such that if $c_1\lambda > 4$ and $n \geq n_1$ then

$$\begin{aligned} \mathbb{P}_{p_{7/4}}(S_{0.9r} \leftrightarrow S_{3.6r}) &\geq 1 - \prod_{\gamma \in \Gamma} \mathbb{P}_{p_{7/4}}(S_{0.9r} \xleftrightarrow{T_\gamma} S_{3.6r}) \geq 1 - (1 - (\log n)^{-1})^{(\log(n^{1/3}))^{c_1\lambda}} \\ &\geq 1 - \exp(3^{-c_1\lambda}(\log n)^{c_1\lambda-1}) \geq 1 - \exp(-(\log n)^{3c_1\lambda/4}), \end{aligned} \quad (3.6.8)$$

where we used that $1 - x \leq e^{-x}$ in the second line.

We next use the fact that G also has (c_1, λ) -polylog-plentiful *annular* tubes at scale r . (We will no longer use the paths and tubes from the radial case that we defined in the previous paragraph, so it is not a problem to reuse the same notation for the annular tubes we consider in the rest of the proof.) We will work with the standard monotone coupling $(\omega_p)_{p \in [0,1]}$ of percolation at different parameters. To lighten notation, we write $\omega_1 = \omega_{p_1}$, $\omega_{3/2} = \omega_{p_{3/2}}$, $\omega_{7/4} = \omega_{p_{7/4}}$, and $\omega_2 = \omega_{p_2}$. Let $u, v \in S_r$ and consider the event $A_{u,v} = \{u \leftrightarrow S_{3r} \text{ and } v \leftrightarrow S_{3r} \text{ in the configuration } \omega_{7/4}\}$. Define $C_u := K_u(\omega_{7/4})$ and $C_v := K_v(\omega_{7/4})$, so that the event $A_{u,v}$ is entirely determined by the pair (C_u, C_v) . Whenever $A_{u,v}$ holds, let Γ be a family of paths from C_u to C_v that is guaranteed to exist by the fact that G has (c_1, λ) -polylog-plentiful annular tubes at scale r (choosing these paths as a function of (C_u, C_v)), and let $(T_\gamma)_{\gamma \in \Gamma}$ be the associated family of tubes, noting that $\bigcup_{\gamma \in \Gamma} T_\gamma \subseteq B(2r[\log r]^{\lambda/c_1})$. Define the configurations

$$\alpha := \omega_2 \cap (\partial_E C_u \cup \partial_E C_v) \cap B(2r[\log r]^{\lambda/c_1}) \quad \text{and} \quad \beta := (\omega_{7/4} \setminus \overline{C_u \cup C_v}) \cap B(2r[\log r]^{\lambda/c_1}),$$

where we recall that $\overline{C_u \cup C_v}$ denotes the set of edges with at least one endpoint in $C_u \cup C_v$. In order for C_u to be connected to C_v in the configuration $(\alpha \cup \beta) \cap T_\gamma$, it suffices that in at least one of the tubes T_γ , there is a β -path from an endpoints of an α -open edge in $\partial_E C_u$ to an endpoint of an α -open edge in $\partial_E C_v$. The estimate (3.6.6) together with the interval inclusion (3.6.7) therefore

yields that

$$\begin{aligned}
\mathbb{P}(C_u \xrightarrow{\alpha \cup \beta} C_v \mid C_u, C_v) &\geq \mathbf{1}(A_{u,v}) \cdot \left[1 - \prod_{\gamma \in \Gamma} \mathbb{P}(C_u \xrightarrow{T_i \cap (\alpha \cup \beta)} C_v \mid C_u, C_v) \right] \\
&\geq \mathbf{1}(A_{u,v}) \cdot \left[1 - \left(1 - \left[\frac{c_{-1}\delta}{4} \right]^2 [\log n]^{-1} \right)^{[\log(n^{1/3})]^{c_1\lambda}} \right] \\
&\geq \mathbf{1}(A_{u,v}) \cdot \left[1 - \exp(-3^{-c_1\lambda} (\log n)^{c_1\lambda-2}) \right],
\end{aligned}$$

where we used that $c_{-1}\delta \geq 4(\log n)^{-1/2}$ (which holds by definition of δ and K_0) in the final inequality. As such, there exist constants $\lambda_0 = \lambda_0(d, D)$ and $n_2 = n_2(d, D, \lambda) \geq n_1$ such that if $\lambda \geq \lambda_0$, $n \geq n_2$, and $c_1\lambda \geq 4$ then

$$\mathbb{P}(C_u \xrightarrow{\alpha \cup \beta} C_v \mid C_u, C_v) \geq \mathbf{1}(A_{u,v}) \cdot \left[1 - e^{-[\log n]^{3c_1\lambda/4}} \right].$$

Since $\alpha \cup \beta \subseteq \omega_2 \cap B(2r[\log r]^{\lambda/c_1})$, it follows that

$$\mathbb{P}(C_u \xrightarrow{\omega_2 \cap B(2r[\log r]^{\lambda/c_1})} C_v \mid A_{u,v}) \geq 1 - e^{-[\log n]^{3c_1\lambda/4}}$$

under the same conditions. Letting \mathcal{U} be the event that all $\omega_{7/4}$ -clusters that intersect both S_r and S_{3r} are contained in a single $\omega_2 \cap B(2r[\log r]^{\lambda/c_1})$ -cluster, it follows by a union bound that there exists $\lambda_1 = \lambda_1(d, D) \geq \lambda_0$ such that if $\lambda \geq \lambda_1$ and $n \geq n_2$ then

$$\begin{aligned}
\mathbb{P}(\mathcal{U}) &\geq 1 - \sum_{u,v \in S_r} \mathbb{P}\left(A_{u,v} \cap \{C_u \xrightarrow{\omega_2 \cap B(2r[\log r]^{\lambda/c_1})} C_v\}\right) \\
&\geq 1 - \text{Gr}(r)^2 e^{-[\log n]^{3c_1\lambda/4}} \geq 1 - e^{[\log n]^{2D}} e^{-[\log n]^{3c_1\lambda/4}} \geq 1 - e^{-[\log n]^{2c_1\lambda/3}}. \tag{3.6.9}
\end{aligned}$$

For each vertex x , let \mathcal{E}_x be the event that $S_{0.9r}(x)$ is $\omega_{7/4}$ -connected to $S_{3.6r}(x)$ and let \mathcal{U}_x be the event that all $\omega_{7/4}$ -clusters that intersect both $S_r(x)$ and $S_{3r}(x)$ are contained in a single $\omega_2 \cap B(x, 2r[\log r]^{\lambda/c_1})$ -cluster. Observe that if ζ is a path in G then we have the inclusion of events

$$\{u \xrightarrow{\omega_2 \cap B_n(\zeta)} v\} \supseteq \{u \xrightarrow{\omega_{7/4}} S_n(u)\} \cap \{v \xrightarrow{\omega_{7/4}} S_n(v)\} \cap \bigcap_{t=0}^{\text{len } \zeta} (\mathcal{U}_{\zeta_t} \cap \mathcal{E}_{\zeta_t}).$$

Thus, it follows by the Harris-FKG inequality and a union bound that there exists a constant

$n_3 = n_3(d, D, \lambda) \geq n_2$ such that if $\lambda \geq \lambda_1$, $n \geq n_3$, and the path ζ satisfies $\text{len } \zeta \leq e^{[\log n]^{c_1 \lambda/4}}$ then

$$\begin{aligned} \mathbb{P}_{p_2}(u \xleftrightarrow{B_n(\zeta)} v) &\geq \mathbb{P}_{p_{7/4}}(u \leftrightarrow S_n(u)) \cdot \mathbb{P}_{p_{7/4}}(v \leftrightarrow S_n(v)) \cdot \prod_{t=1}^{\text{len } \zeta} \mathbb{P}(\mathcal{E}_{\zeta_t}) - \sum_{t=1}^{\text{len } \zeta} \mathbb{P}(\mathcal{U}_{\zeta_t}^c) \\ &\geq \left(e^{-(\log \log n)^{1/2}}\right)^2 \left[1 - \text{len}(\zeta) e^{-(\log n)^{3c_1 \lambda/4}}\right] - \text{len}(\zeta) e^{-(\log n)^{3c_1 \lambda/4}} \\ &\geq e^{-3(\log \log n)^{1/2}}. \end{aligned}$$

The claimed lower bound on the corridor function follows since ζ was an arbitrary path of length at most $e^{[\log n]^{c_1 \lambda/4}}$. \square

Concluding from a negative assumption

The next lemma explains how to use the negative information encapsulated in an *upper bound* on the two-point zone $\text{tz}(m)$ to find a set of vertices for which point-to-point connection probabilities within a ball are uniformly small. This lemma plays an analogous role to that of Section 7.2 in [CMT22], but our proof is completely different and relies on the machinery developed in Section 3.4. Note that this lemma does not use the plentiful tubes condition.

Lemma 3.6.5. *There exists a constant $n_0 = n_0(d, D)$ such that if $n \geq n_0$ and $b \leq \frac{1}{8}n^{1/3}$ satisfy $\mathcal{A}_{D,K}(b, n)$ then for every $n^{1/3} \leq m \leq n$, there exists a subset $U \subseteq B_{\text{tz}(m)}$ with $|U| \geq (\log n)^{c_3 K}$ such that*

$$\mathbb{P}_{p_1}(u \xleftrightarrow{B_{m/2}} v) \leq (\log n)^{-c_3 K}$$

for all distinct $u, v \in U$.

(Reminder: The implicit assumption $K \geq K_0$ remains in force.)

Proof of Lemma 3.6.5. Fix n and $b \leq \frac{1}{8}n^{1/3}$ such that $\mathcal{A}_{D,K}(b, n)$ holds. We will assume that the claim is *false* for some particular $n^{1/3} \leq m \leq n$, and show that this implies a contradiction when n is sufficiently large. By definition of the two-point zone $\text{tz}(m) = \text{tz}(m, n)$, there exist vertices $u, v \in B_{\text{tz}(m)}$ such that $\mathbb{P}_{p_{3/2}}(u \xleftrightarrow{B_m} v) < (\log n)^{-1}$. Let γ be a geodesic in $B_{\text{tz}(m)}$ from u to v . Recursively pick a sequence of indices $0 = i_0 < i_1 < \dots < i_k = \text{len } \gamma$ starting with $i_0 := 0$ and for each $j \geq 0$ with $i_j < \text{len } \gamma$ setting

$$i_{j+1} = \max \left\{ \text{len } \gamma, 1 + \max \left\{ i \geq i_j : \mathbb{P}_{p_1}(u_{i_j} \xleftrightarrow{B_{m/2}} u_i) \geq (\log n)^{-c_3 K} \right\} \right\}.$$

To lighten notation, define $v_j := u_{i_j}$ for every $0 \leq j \leq k$. This sequence has the property that $\mathbb{P}_{p_1}(v_j \leftrightarrow v_{j+1}) \geq p_1 (\log n)^{-c_3 K}$ for every $0 \leq j < k$ and $\mathbb{P}_{p_1}(v_j \leftrightarrow v_\ell) \leq (\log n)^{-c_3 K}$ for every

distinct $0 \leq j, \ell < k$. (The last connection probability, from v_{k-1} to $v_k = v$, may be larger.) As such, our assumption guarantees that $k \leq (\log n)^{c_3 K}$.

To conclude the proof, it suffices to show that $\mathbb{P}_{p_{3/2}}(u \leftrightarrow v) \geq (\log n)^{-1}$ when n is sufficiently large. To this end, we would like to apply Proposition 3.4.1 where the sets ‘ A_i ’ are the balls $B_b(v_j)$, the superset ‘ Λ ’ is the bigger ball $B_{m/2}$, the ghost field intensity ‘ h ’ is $h := \min\{\log n, \text{Gr}(b)\}^{-1}$, and the sprinkling amount ‘ δ ’ is $\delta/2$. To do this, we need to verify that

$$h^{c_2(\delta/2)^3} \leq c_2 k^{-1} \quad \text{and} \quad \tau_{p_1}^{B_{m/2}}(B_b(v_j) \cup B_b(v_{j+1})) \geq 4h^{-c_2(\delta/2)^4} \quad (3.6.10)$$

for every $0 \leq j \leq k-1$ when n is sufficiently large. We have by definition of δ (and the assumption $K \geq K_0$) that if n is larger than some universal constant then

$$h^{c_2(\delta/2)^3} \leq 4h^{c_2(\delta/2)^4} \leq 4(\log n)^{-2c_3 K} \leq c_2 [(\log n)^{-c_3 K} + 1]. \quad (3.6.11)$$

Since $k \leq (\log n)^{c_3 K}$, this is easily seen to imply that the first inequality of (3.6.10) holds. Moreover, we have by the same calculation performed in (3.6.4) that

$$\tau_{p_1}^{B_{m/2}}(B_b) \geq \frac{1}{2} e^{-(\log \log n)^{1/2}} \geq (\log n)^{-1} \quad (3.6.12)$$

for all n larger than some universal constant, and it follows by the Harris-FKG inequality that if n is larger than some constant depending only on d then

$$\begin{aligned} \tau_{p_1}^{B_{m/2}}(B_b(v_j) \cup B_b(v_{j+1})) &\geq \tau_{p_1}^{B_{m/2}}(B_b(v_j)) \cdot \mathbb{P}_{p_1}(v_j \xleftrightarrow{B_{m/2}} v_{j+1}) \cdot \tau_{p_1}^{B_{m/2}}(B_b(v_{j+1})) \\ &\geq (\log n)^{-1} \cdot \left[\frac{1}{2d} (\log n)^{-c_3 K} \right] \cdot (\log n)^{-1} \geq (\log n)^{-c_3 K - 3} \end{aligned} \quad (3.6.13)$$

for every $0 \leq j \leq k-1$. The estimates (3.6.11) and (3.6.13) together yield that there exists a constant $n_0 = n_0(d) \geq N$ such that the required estimates (3.6.10) hold whenever $n \geq n_0$. Thus, Proposition 3.4.1 and (3.6.12) yield that there exists a constant $n_1 = n_1(d) \geq n_0$ such that if $K \geq K_0$ and $n \geq n_1$ and we define r to be the minimum positive integer such that $\mathbb{P}_q(\text{Piv}[1, rh]) < h$ for every $q \in [p_1, p_{3/2}]$ then (since $p_{3/2} \geq \text{Spr}(p_1; \delta/2)$)

$$\mathbb{P}_{p_{3/2}} \left(u \xleftrightarrow{B_{m/2+2r}} v \right) \geq c_2 \tau_{p_1}^{B_{m/2}}(B_b(u)) \cdot \tau_{p_1}^{B_{m/2}}(B_b(v)) \geq \frac{c_2}{4} e^{-2[\log \log n]^{1/2}}.$$

Finally, the same argument as in the proof of Lemma 3.6.4 yields that $2r \leq m/2$ when n is larger than some constant depending on d and D , and it follows that there exists a constant $n_2 = n_2(d, D) \geq n_1$ such that if $n \geq n_2$ then $\mathbb{P}_{p_{3/2}}(u \xleftrightarrow{B_m} v) \geq (\log n)^{-1}$ — a contradiction. \square

We now use the existence of this set of poorly connected vertices (negative information) to prove that $S_{\text{tz}(m)}$ is very likely to be connected to the boundary of $S_{m/2}$ (positive information). This only works because we are working under the positive hypothesis of a two-point lower bound at scale n . This step is essentially the same as Section 7 in [CMT22], with our two-point lower bound at scale n playing the role of ‘being in the supercritical regime’ in their setting.

Lemma 3.6.6. *There exists a constant $n_0 = n_0(d, D)$ such that if $n \geq n_0$ and $b \leq \frac{1}{8}n^{1/3}$ satisfy $\mathcal{A}_{D,K}(b, n)$ then*

$$\mathbb{P}_{p_{3/2}}(S_{\text{tz}(m)} \leftrightarrow S_{m/2}) \geq 1 - e^{-(\log n)^{c_3 K^{-1}}}$$

for every $m \in \mathcal{S}(n)$.

Note that we are not actually assuming a negative information assumption (such as $\neg \mathcal{B}(n, p)$) in the hypotheses of this lemma. The lemma holds without any such assumption, but is stronger when $\text{tz}(m)$ is small. The proof of Lemma 3.6.6 will apply the following proposition.

Proposition 3.6.7. *Let G be a finite connected graph, let $p \in [0, 1]$, and let $A, B \subseteq V(G)$. If $\theta \in (0, 1)$ is such that*

$$\min_{x \in A} \mathbb{P}_p(x \leftrightarrow B) \geq \theta \geq 2 |A| \max_{\substack{x, y \in A \\ x \neq y}} \mathbb{P}_p(x \leftrightarrow y),$$

then $\mathbb{P}_q(A \leftrightarrow B) \geq 1 - e^{-\delta(p, q)\theta|A|}$ for every $q \in (p, 1)$.

Proof of Proposition 3.6.7. This proposition is essentially the same as [CMT22, Proposition 7.2], except that it is stated in terms of our sprinkling coordinates introduced in Section 3.3 (which are natural from the perspective of Talagrand’s inequality) and we get a factor δ rather than 2δ in the exponential in the conclusion. Both versions of the proposition are elementary consequences of the differential inequality

$$\frac{d}{dp}(-\log \mathbb{P}_p(A)) \geq \frac{1}{p(1-p)} \mathbb{E}_p[\text{Hamming distance from } \omega \text{ to } A],$$

which holds for every finite graph and every decreasing event A [Gri06, Theorem 2.53]. In our coordinates, this inequality reads

$$\frac{d}{dt}(-\log \mathbb{P}_{\text{Spr}(p;t)}(A)) \geq \frac{\log 1/(1 - \text{Spr}(p;t))}{\text{Spr}(p;t)} \mathbb{E}_{\text{Spr}(p;t)}[\text{Hamming distance from } \omega \text{ to } A].$$

In our case the prefactor $-(\log(1 - \text{Spr}(p;t)))/\text{Spr}(p;t)$ is at least 1 whereas in the original inequality the prefactor $1/(p(1-p))$ is at least 2, leading to the difference between our two conclusions. \square

Proof of Lemma 3.6.6. Let $n_0 = n_0(d, D)$ be as in Lemma 3.6.5, and suppose that $n \geq n_0$ and $b \leq \frac{1}{8}n^{1/3}$ are such that $\mathcal{A}(b, n)$ holds. Let $U \subseteq B_{\text{tz}(m)}$ be the set of vertices guaranteed to exist by Lemma 3.6.5, and let A be a subset of U with $|A| = \lceil [\log n]^{c_3 K - 1/2} \rceil$. Since $\mathcal{A}(b, n)$ holds, we have that

$$\min_{x \in A} \mathbb{P}_{p_1}(x \leftrightarrow S_{m/2}) \geq \kappa_{p_1}(n, \infty) \geq e^{-[\log \log n]^{1/2}},$$

while our choice of A guarantees that

$$2|A| \max_{\substack{x, y \in A \\ x \neq y}} \mathbb{P}_{p_1}(x \xleftrightarrow{B_{m/2}} y) \leq 2\lceil (\log n)^{c_3 K - 1/2} \rceil (\log n)^{-c_3 K} \leq e^{-(\log \log n)^{1/2}}$$

whenever n is larger than some constant $n_1 = n_1(d, D) \geq n_0$. (This constant does not depend on K since $c_3 K \geq c_3 K_0 \geq 2 > 1/2$.) Thus, applying Proposition 3.6.7 with $\theta = e^{-[\log \log n]^{1/2}}$ yields that (since $p_{3/2} \geq \text{Spr}(p_1, \delta/2)$)

$$\mathbb{P}_{p_{3/2}}(B_{\text{tz}(m)} \leftrightarrow S_{m/2}) \geq \mathbb{P}_{p_{3/2}}(A \leftrightarrow S_{m/2}) \geq 1 - \exp\left(-\frac{\delta}{2} e^{-[\log \log n]^{1/2}} |A|\right). \quad (3.6.14)$$

On the other hand, the definition of δ ensures that $\delta \geq [\log n]^{-1/4}$ and hence that there exists a constant $n_2 = n_2(d, D) \geq n_1$ such that

$$\frac{\delta}{2} e^{-[\log \log n]^{1/2}} |A| \geq \frac{1}{2} e^{-[\log \log n]^{1/2}} (\log n)^{c_3 K - 1/2 - 1/4} \geq (\log n)^{c_3 K - 1}$$

whenever $n \geq n_2$, which implies the claim in conjunction with (3.6.14). \square

The next lemma is completely elementary. It tells us that we can find two nearby reals m and m' where $\text{tz}(m)$ is close to $\text{tz}(m')$.

Lemma 3.6.8. *Let $R \geq 1$. There exists a constant $n_0 = n_0(d, D, R)$ such that if $n \geq n_0 \vee N$ then there exists $m \in \mathcal{S}(n)$ such that $m(\log n)^{-R} \in \mathcal{S}(n)$ and*

$$\frac{\text{tz}(m)}{\text{tz}(m(\log n)^{-R})} \leq (\log n)^{8R/c_1}.$$

Proof of Lemma 3.6.8. Let $n \geq N$ so that $\mathcal{S}(n)$ is defined. Let s and t denote the left and right endpoints of $\mathcal{S}(n)$, and define

$$k := \lfloor \log_{[\log n]^R}(t/s) \rfloor = \left\lfloor \frac{c_1 \log(m_1)}{2R \log \log n} \right\rfloor \geq \left\lfloor \frac{c_1 \log n}{6R \log \log n} \right\rfloor$$

where $n^{1/3} \leq m_1 \leq n$ is as in the definition of $\mathcal{S}(n)$. If a suitable $m \in \mathcal{S}(n)$ does *not* exist then, using the trivial inequalities $\text{tz}(t) \leq t \leq n$ and $\text{tz}(s) \geq 1$, we must have that

$$\begin{aligned} n \geq \text{tz}(t) &\geq \text{tz}(s) \prod_{i=1}^k \frac{\text{tz}(s[\log n]^{iR})}{\text{tz}(s[\log n]^{(i-1)R})} \\ &\geq \text{tz}(s) \left([\log n]^{8R/c_1} \right)^k \geq \exp \left(\frac{8R}{c_1} \left\lfloor \frac{c_1 \log n}{6R \log \log n} \right\rfloor \cdot \log \log n \right). \end{aligned}$$

Since $8 > 6$, this yields a contradiction when n is larger than some constant $n_0 = n_0(d, D, R)$ (allowing us to approximately remove the effect of rounding down). \square

We will now combine our lemmas to prove Lemma 3.6.3. The idea here is inspired by the *uniqueness via sprinkling* argument from [CMT22, Section 8], which itself used ideas from [BT17]. Our approach is different because we do not know that exposed spheres are well-connected. Instead, we have the polylog-plentiful tubes condition. This is a much weaker geometric control because the tubes are not constrained to lie within narrow annuli. The main step is to use the strong connectivity bound from Lemma 3.6.6 to deduce that with high probability, every $\omega_{7/4}$ -cluster crossing a thick annulus is contained in a single ω_2 -cluster. (As before we abbreviate $\omega_2 = \omega_{p_2}$ and so on.) In [CMT22], the analogous step was carried out by dividing the thick annulus into thinner annuli before showing that if two clusters cross multiple annuli, then, after sprinkling in those annuli, the clusters will merge with high probability. This works because in every thin annulus, there is some good probability that the clusters will merge after sprinkling in the annulus, thanks to the connectivity of exposed spheres. In our case, we also track how many clusters survive un-merged as they cross through multiple annuli. The difference is that we will have to sprinkle *everywhere* each time we cross a thin annulus. Nevertheless, we will sprinkle so little at each stage that the net effect is to sprinkle by less than $\delta/4$, as required.

Proof of Lemma 3.6.3. Let $K_1 = K_1(d, D, \lambda) = \max\{K_0, 40\lambda(c_1 \vee 1)^{-2}c_3^{-1}, 4c_3^{-1}(D \vee 1)\}$ and suppose that $K \geq K_1$ and $n \geq N$. Fix $R := 5\lambda/c_1$ and let $n_0 = n_0(d, D, R)$ be the constant from Lemma 3.6.8, which by our choice of R depends only on d, D , and λ . We also let $n_1 = n_1(d, D)$ and c_4 be the constants from Lemma 3.6.4.

Suppose that $n \geq n_2 = n_2(d, D, \lambda) = n_0 \vee n_1 \vee N \vee e^{2^R}$ and $b \leq \frac{1}{8}n^{1/3}$ are such that $\mathcal{A}(b, n, p)$ holds, and let $m \in \mathcal{S}(n)$ be the element guaranteed to exist by Lemma 3.6.8 applied with this value of R . Taking $n_2 \geq e^{2^R}$ guarantees that $m/2 \geq m(\log n)^{-R}$ and hence that $m/2 \in \mathcal{S}(n)$ when $n \geq n_2$. Since $n \geq n_1$, we have by Lemma 3.6.4 that

$$\kappa_{p_{7/4}} \left(\text{tz}(m(\log n)^{-R})(\log n)^{c_3 K}, 2m(\log n)^{-R} \right) \geq c_4 e^{-2[\log \log n]^{1/2}}.$$

On the other hand, our choice of R and m ensure that there exists a constant $n_3 = n_3(d, D, \lambda) \geq n_2$ such that if $n \geq n_3$ then we have the inclusion of intervals

$$\left(2m[\log n]^{-R}, \text{tz}(m[\log n]^{-R})[\log n]^{c_3 K}\right) \supseteq \left(\frac{m}{3}[\log m]^{-4\lambda/c_1}, 2\text{tz}(m/2) + 2\right),$$

so that

$$\kappa_{p_{7/4}}\left(2\text{tz}(m/2) + 2, \frac{m}{3}[\log m]^{-4\lambda/c_1}\right) \geq c_4 e^{-2[\log \log n]^{1/2}}. \quad (3.6.15)$$

Note that this estimate holds as a consequence of $\mathcal{A}(b, n, p)$ alone: we have not yet made use of the negative information $\neg\mathcal{B}(n, p)$.

Now suppose that $\neg\mathcal{B}(n, p)$ holds. Since $K \geq K_1 \geq 12c_3^{-1}c_1^{-1}\lambda$, there exists a constant $n_3 = n_3(d, D, \lambda) \geq n_3$ such that if $n \geq n_3$ then

$$\text{tz}(m/2) \leq m[\log n]^{3\lambda/c_1 - c_3 K} \leq m(\log n)^{-3c_3 K/4} \leq \frac{m}{17}(\log m)^{-c_3 K/2} \quad (3.6.16)$$

for every $m \in \mathcal{S}(n)$.

Our next goal is to prove a good upper bound on the probability of the non-uniqueness event $\text{Piv}_{p_{7/4}, p_2}[m/16, m/8]$, where $\text{Piv}_{p, q}[m, n]$ denotes the event that there are at least two distinct ω_p -clusters that each intersect both B_m and S_n but that are not connected to each other by any path in $B_n \cap \omega_q$. We will do this using a variation on the “orange peeling” argument of [CMT22], where we iteratively sprinkle and zoom in closer to $m/16$ over a number of steps.

Let $k := 2\lfloor(\log n)^D\rfloor$, $\varepsilon := (\log n)^{-(D+1)}$, and for each $i \in \{0, \dots, k\}$ set

$$r_i := \frac{m}{8} - \frac{im}{40}[\log n]^{-D} \quad \text{and} \quad q_i := \text{Spr}(p_{7/4}; i\varepsilon).$$

Note that $r_i \in [m/16, m/8]$ and $q_i \in [p_{7/4}, p_2]$ for every $0 \leq i \leq k$. We work with the standard monotone coupling $(\omega_q)_{q \in [0,1]}$, and write $\omega_{(i)} = \omega_{q_i}$ for each $i \in \{0, \dots, k\}$. (Be careful not to confuse this with our previous notational shorthand $\omega_1 = \omega_{p_1}$, $\omega_2 = \omega_{p_2}$.) Given the family of configurations $(\omega_q)_{q \in [0,1]}$, recursively define a set of $B_{m/8} \cap \omega_{(i)}$ -clusters \mathcal{C}_i for each $i \in \{0, \dots, k\}$ as follows:

1. Let \mathcal{C}_0 be the set of all $B_{m/8} \cap \omega_{(0)}$ -clusters that contain a vertex in S_{r_0} .
2. Given \mathcal{C}_i for some $i < k$, let \mathcal{C}_{i+1} be the set of $B_{m/8} \cap \omega_{(i+1)}$ -clusters C such that there exists $C' \in \mathcal{C}_i$ with $C' \subseteq C$ and $C' \cap S_{r_{i+1}} \neq \emptyset$.

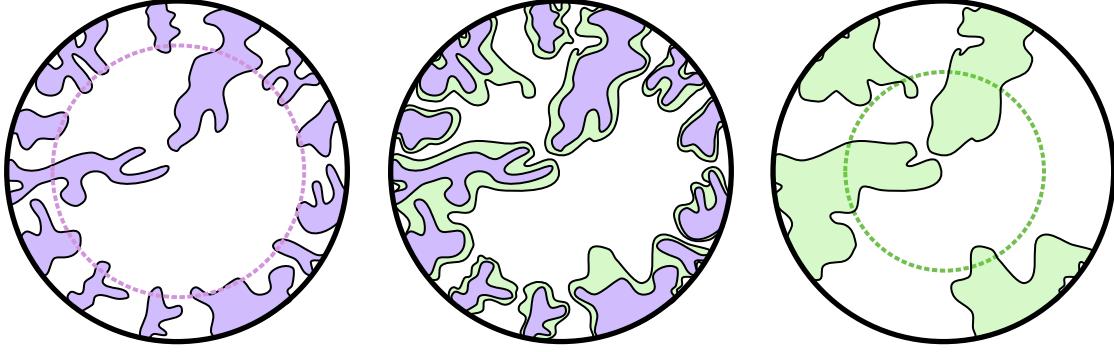


Figure 3.4: Schematic illustration of the construction of \mathcal{C}_{i+1} from \mathcal{C}_i : The black, purple, and green circles represent the spheres S_{r_0} , S_{r_i} , and $S_{r_{i+1}}$ respectively. The purple regions in the left figure represent the clusters making up \mathcal{C}_i . By construction each of these is a $B_{r_0} \cap \omega_{(i)}$ -cluster that connects S_{r_0} to S_{r_i} . The green regions in the middle figure represent the $B_{r_0} \cap \omega_{(i+1)}$ clusters that contain some cluster in \mathcal{C}_i . Finally, the green regions remaining in the third figure represent the subset of these clusters that happen to intersect $S_{r_{i+1}}$; these make up \mathcal{C}_{i+1} .

See fig. 3.4 for an illustration. This definition ensures that the cardinality $|\mathcal{C}_i|$ is a decreasing function of i and that we have the inclusion of events

$$\text{Piv}_{p_{7/4}, p_2}[m/16, m/8] \subseteq \{\mathcal{C}_k \in \{0, 1\}\}. \quad (3.6.17)$$

Roughly speaking, our goal is to show that, for each i , if \mathcal{C}_i is not a singleton then the cardinality $|\mathcal{C}_{i+1}|$ is smaller than $|\mathcal{C}_i|$ by a factor of roughly $1/2$ with high probability under \mathbb{P} .

Ideally, we would like to show that, with high probability on the event $\{|\mathcal{C}_i| > 1\}$, every cluster in \mathcal{C}_i that intersects $S_{r_{i+1}}$ is $B_{m/8} \cap \omega_{(i+1)}$ -connected to a distinct cluster in \mathcal{C}_i . This is what we will prove, except that we will allow one distinguished cluster to go un-merged. For every non-empty set of clusters \mathcal{F} , permanently fix a choice of element $\min(\mathcal{F}) \in \mathcal{F}$ such that $\text{dist}(o, \min(\mathcal{F})) = \text{dist}(o, \bigcup \mathcal{F})$. For each $0 \leq i \leq k-1$, consider the event

$$\mathcal{E}_i = \{\mathcal{C}_i = \emptyset\} \cup \left\{ C \xrightarrow{B_{m/8} \cap \omega_{(i+1)}} \bigcup (\mathcal{C}_i \setminus \{C\}) \text{ for every } C \in \mathcal{C}_i \setminus \{\min(\mathcal{C}_i)\} \text{ with } \text{dist}(o, C) \leq r_{i+1} \right\}.$$

We will prove the following lemma at the end of the section after explaining how it may be used to conclude the proof of Lemma 3.6.3 (and hence of Proposition 3.6.1).

Lemma 3.6.9 (Merging clusters). *There exists a constant $n_{11} = n_{11}(d, D, \lambda) \geq n_4$ such that if $n \geq n_{11}$ then*

$$\mathbb{P}(\mathcal{E}_i) \geq 1 - e^{-[\log n]^{c_1 \lambda/2}}, \quad (3.6.18)$$

for every $0 \leq i \leq k-1$.

(We name this constant n_{11} to leave room for the constants n_6 through n_{10} that will appear in the proof of this claim.)

Let us now conclude the proof of Lemma 3.6.3 given Lemma 3.6.9. On the event \mathcal{E}_i we have that

$$|\mathcal{C}_{i+1}| \leq \left\lfloor \frac{|\mathcal{C}_i| - 1}{2} \right\rfloor + 1.$$

Since $|\mathcal{C}_0| \leq |S_{r_0}| \leq e^{(\log n)^D}$ and $k := 2\lfloor [\log n]^D \rfloor$, it follows that $|\mathcal{C}_k| \in \{0, 1\}$ on the event $\bigcap_{i=0}^{k-1} \mathcal{E}_i$. It follows by a union bound that there exists a constant $n_6 = n_6(d, D, \lambda) \geq n_5$ such that if $n \geq n_6$ then

$$\mathbb{P} \left(\text{Piv}_{p_{7/4}, p_2} [m/16, m/8] \right) \leq \sum_{i=0}^{k-1} \mathbb{P}_p (\mathcal{E}_i^c) \leq 2(\log n)^D e^{-(\log n)^{c_1 \lambda/2}} \leq e^{-(\log n)^{c_1 \lambda/3}}. \quad (3.6.19)$$

On the other hand, (3.6.16) ensures that if $n \geq n_6$ then $\text{tz}(m/2) \leq m/17$, and it follows from Lemma 3.6.6 applied to $m/2$ that

$$\mathbb{P}_{p_{7/4}} (S_{m/17} \leftrightarrow S_{m/4}) \geq 1 - e^{-(\log n)^{c_1 \lambda/3}}$$

if $n \geq n_6$. It follows in particular that there exists $n_7 = n_7(d, D, \lambda) \geq n_6$ such that

$$\mathbb{P}_{p_{7/4}} (S_{m/17} \leftrightarrow S_{m/4})^{e^{[\log n]^{c_1 \lambda/4}}} \geq \frac{1}{2} \quad (3.6.20)$$

if $n \geq n_7$; this will be used to form a chain of connected annuli using the Harris-FKG inequality.

We now apply (3.6.19) and (3.6.20) to bound the corridor function. Let γ be a path of length at most $e^{[\log n]^{c_1 \lambda/4}}$, starting at some vertex u and ending at some vertex v , and observe that we have the inclusion of events

$$\begin{aligned} \{u \xrightarrow{B_n(\gamma) \cap \omega_2} v\} &\subseteq \{u \xrightarrow{\omega_{7/4}} S_n(u)\} \cap \{v \xrightarrow{\omega_{7/4}} S_n(v)\} \cap \\ &\quad \bigcap_{t=1}^{\text{len } \gamma} \left(\{S_{m/17}(\gamma_t) \xrightarrow{\omega_{7/4}} S_{m/4}(\gamma_t)\} \cap \text{Piv}_{p_{7/4}, p_2} [m/16, m/8](\gamma_t) \right). \end{aligned}$$

Applying the Harris-FKG inequality and a union bound, we deduce that there exists a constant $n_8 = n_8(d, D, \lambda)$ such that if $n \geq n_8$ then

$$\begin{aligned} \mathbb{P}_{p_2} (u \xrightarrow{B_n(\gamma)} v) &\geq \mathbb{P}_{p_{7/4}} (o \leftrightarrow S_n)^2 \cdot \mathbb{P}_{p_{7/4}} (S_{m/17} \leftrightarrow S_{m/4})^{\text{len}(\gamma)} \\ &\quad - \text{len}(\gamma) \cdot (1 - \mathbb{P} (\text{Piv}_{p_{7/4}, p_2} [m/16, m/8])) \\ &\geq \frac{1}{2} e^{-2(\log \log n)^{1/2}} - e^{(\log n)^{c_1 \lambda/4}} e^{-(\log n)^{c_1 \lambda/3}} \\ &\geq e^{-3[\log \log n]^{1/2}}, \end{aligned}$$

where we used the assumption that $\mathcal{A}(b, n, p)$ holds to bound $\mathbb{P}_{p_{\gamma/4}}(o \leftrightarrow S_n) \geq \kappa_{p_1}(n, \infty) \geq e^{-(\log \log n)^{1/2}}$. The claimed lower bound on the corridor function follows since γ was an arbitrary path of length at most $e^{[\log n]^{c_1 \lambda/4}}$. \square

It remains only to prove Lemma 3.6.9.

Proof of Lemma 3.6.9. We continue to use the notation from the proof of Lemma 3.6.3, and in particular will use the constants $K_1 = K_1(d, D, \lambda)$ and $n_4 = n_4(d, D, \lambda)$ defined in that proof. Consider the event Ω defined by

$$\Omega := \bigcap_{u \in B_{m/8}} \{S_{\text{tz}(m/2)}(u) \xrightarrow{\omega_{3/2}} S_{m/8}\}.$$

It follows by Lemma 3.6.6 (applied to $m/2$) and a union bound that there exists a constant $n_5 = n_5(d, D, \lambda) \geq n_4$ such that if $n \geq n_5$ then

$$\begin{aligned} \mathbb{P}(\Omega) &\geq 1 - \text{Gr}(m) \sup_{u \in B_{m/8}} \mathbb{P}_{p_{3/2}}(S_{\text{tz}(m/2)}(u) \leftrightarrow S_{m/8}) \\ &\geq 1 - \text{Gr}(m) \sup_{u \in B_{m/8}} \mathbb{P}_{p_{3/2}}(S_{\text{tz}(m/2)}(u) \leftrightarrow S_{m/4}(u)) \\ &\geq 1 - e^{(\log m)^D} e^{-(\log n)^{c_3 K - 1}} \geq 1 - e^{-(\log n)^{c_3 K/2}}, \end{aligned} \quad (3.6.21)$$

where we used that $K \geq K_1 \geq 4c_3^{-1}(D \vee 1)$ in the final inequality.

For each $0 \leq i \leq k-1$, let \mathbb{F}_i be the set $\mathbb{F}_i = \{\mathcal{F} : \mathbb{P}(\mathcal{C}_i = \mathcal{F} \mid \Omega) > 0\}$. (Note that \mathbb{F}_i is a set of sets of sets of vertices.) It follows from the definitions that

$$d\left(u, \bigcup_{C \in \mathcal{F}} C\right) \leq \text{tz}(m/2) \text{ for every } u \in B_{m/8} \text{ and } \mathcal{F} \in \mathbb{F}_i. \quad (3.6.22)$$

By (3.6.21) and a union bound, it suffices to prove that there exists $n_{11} = n_{11}(d, D, \lambda) \geq n_5$ such that

$$\mathbb{P}(\mathcal{E}_i \mid \mathcal{C}_i = \mathcal{F}) \geq 1 - \frac{1}{2} e^{-(\log n)^{c_1 \lambda/2}} \quad (3.6.23)$$

for every $0 \leq i \leq k-1$ and every $\mathcal{F} \in \mathbb{F}_i$.

Before doing this, we will need to prove a purely geometric preliminary claim. Let $0 \leq i \leq k-1$, let $\mathcal{F} \in \mathbb{F}_i$, and let $C \in \mathcal{F} \setminus \{\min(\mathcal{F})\}$. We claim that there exists a constant $n_8 = n_8(d, D, \lambda) \geq n_5$ such that if $n \geq n_8$ and $\text{dist}(o, C) \leq r_{i+1}$ then there exists a set of vertices $U \subseteq B(r_{i+1} + m(\log m)^{-\lambda/c_1})$ with $|U| \geq \frac{1}{2}(\log m)^{c_1 \lambda}$ such that the following hold:

(i) U is $m(\log m)^{-4\lambda/c_1}$ -separated. That is, pairwise distances between distinct points in U are at least $m(\log m)^{-4\lambda/c_1}$.

(ii) For each $u \in U$, the ball $B_{\text{tz}(m/2)+1}(u)$ intersects both C and $\bigcup(\mathcal{F} \setminus \{C\})$.

(As before, we name this constant n_8 to leave room for the constants n_6 and n_7 that will appear in the proof of this claim.) We let $r := m(\log m)^{-2\lambda/c_1}$ and split the proof of this claim into two cases according to whether $\text{dist}(o, C)$ is smaller or larger than r .

Case 1: ($\text{dist}(o, C) \leq r$.) Since $C \neq \min(\mathcal{F})$, we have that $\text{dist}(o, \bigcup(\mathcal{F} \setminus \{C\})) \leq \text{dist}(o, C) \leq r$ also. Let Γ be the family of paths from $B_{3r} \cap C$ to $B_{3r} \cap \bigcup(\mathcal{F} \setminus \{C\})$ that is guaranteed to exist by the fact that G has (c_1, λ) -polylog-plentiful annular tubes at scale r . We now observe that for each $\gamma \in \Gamma$, there exists a vertex u_γ on the path γ satisfying the ball-intersection condition (ii):

- If $\max_t \text{dist}(C, \gamma_t) \leq \text{tz}(m/2)$ then we may take u_γ to be the final vertex of γ .
- Otherwise, if $\max_t \text{dist}(C, \gamma_t) > \text{tz}(m/2)$, we may take $u_\gamma = \gamma_{t_\gamma}$ where t_γ is the maximum index such that $\text{dist}(C, \gamma_t) \leq \text{tz}(m/2)$. To see that this choice of u_γ satisfies (ii), note that $\text{dist}(\gamma_{t_\gamma+1}, C) > \text{tz}(m/2)$ and hence by (3.6.22) that $\text{dist}(\gamma_{t_\gamma+1}, \bigcup(\mathcal{F} \setminus \{C\})) \leq \text{tz}(m/2)$.

Now define $U := \{u_\gamma : \gamma \in \Gamma\}$. Since the family Γ is $2r(\log r)^{-\lambda/c_1}$ -separated (and r was defined to be $r = m(\log m)^{-2\lambda/c_1}$), it follows that there exists a constant $n_6 = n_6(d, D, \lambda) \geq n_5$ such that if $n \geq n_6$ then U is $m(\log m)^{-4\lambda/c_1}$ -separated and

$$|U| = |\Gamma| \geq (\log r)^{c_1 \lambda} \geq \frac{1}{2}(\log m)^{c_1 \lambda}.$$

Moreover, since every path in Γ was contained in $B(3r + r[\log r]^{\lambda/c_1})$, there exists a constant $n_7 = n_7(d, D, \lambda) \geq n_6$ such that if $n \geq n_7$ then $3r + r[\log r]^{\lambda/c_1} \leq m/16 \leq r_{i+1} + m(\log m)^{-\lambda/c_1}$ and hence $U \subseteq B(r_{i+1} + m(\log m)^{-\lambda/c_1})$.

Case 2: ($\text{dist}(o, C) > r$.) Let $v \in C$ be such that $\text{dist}(o, v) = \text{dist}(o, C)$, let γ^a be a path in C from v to $S_r(v)$, and let γ^b be the portion of a geodesic from v to o starting at a neighbour u of v with $\text{dist}(o, u) < \text{dist}(o, v)$ and ending at the first intersection with $S_r(v)$. These path are both finite, start in $S_1(v)$ and end in $S_r(v)$, and are contained in C and disjoint from C respectively. Let Γ be the family of paths from γ^a to γ^b that is guaranteed to exist by the fact that G has (c_1, λ) -polylog-plentiful annular tubes at scale $r/3$. We can construct the desired set U by picking a vertex u_γ in γ satisfying the condition (ii) for each $\gamma \in \Gamma$; the fact that such a vertex exists for each

γ follows by the same argument used in case 1 above. Moreover, it follows by the same argument used in the first case that there exists a constant $n_8 = n_8(d, D, \lambda) \geq n_7$ such that if $n \geq n_8$ then the required bounds on the cardinality and separation of the set $U = \{u_\gamma : \gamma \in \Gamma\}$ hold, as well as the containment $U \subseteq B(r_{i+1} + m(\log m)^{-\lambda/c_1})$. This concludes the proof of the geometric claim.

We now use this geometric claim to establish the estimate (3.6.23), which will complete the proof of the lemma. Let $i \in \{0, \dots, k-1\}$ and let $\mathcal{F} \in \mathbb{F}_i$ be arbitrary. Consider also an arbitrary $C \in \mathcal{F} \setminus \{\min(\mathcal{F})\}$ with $\text{dist}(o, C) \leq r_{i+1}$, and let U be the corresponding set of vertices guaranteed to exist by the geometric claim above. For each $u \in U$, let $\beta_u := B(u; \frac{m}{2}(\log m)^{-4\lambda/c_1}) \cap \omega_{(i+1)}$. Consider an arbitrary vertex $u \in U$. By construction of U , there exists a path γ that starts in C , ends in $\bigcup(\mathcal{F} \setminus \{C\})$, is contained in $B(u; \text{tz}(m/2) + 1)$, and has length at most $2 \text{tz}(m/2) + 2$. The estimate eq. (3.6.16) yields the existence of a constant $n_9 = n_9(d, D, \lambda) \geq n_8$ such that if $n \geq n_9$ then

$$[\text{tz}(m/2) + 1] + \frac{m}{3}(\log m)^{-4\lambda/c_1} \leq \frac{m}{2}(\log m)^{-4\lambda/c_1},$$

so that the tube $B(\gamma; \frac{m}{3}(\log m)^{-4\lambda/c_1})$ associated to this path γ is contained in the ball $B(u; \frac{m}{2}(\log m)^{-4\lambda/c_1})$.

Thus, if $n \geq n_9$, the estimate (3.6.15) yields that

$$\mathbb{P}\left(C \xleftrightarrow{\beta_u} \bigcup(\mathcal{F} \setminus \{C\})\right) \geq c_4 e^{-2[\log \log n]^{1/2}}.$$

We stress that the C and \mathcal{F} appearing in this inequality are deterministic, and do not depend on the configuration β_u . Under \mathbb{P} , the conditional law of β_u given that $\mathcal{C}_i = \mathcal{F}$ is simply (inhomogeneous) bond percolation on $B(u; \frac{m}{2}(\log m)^{-4\lambda/c_1})$ where every edge has probability at least $c_{-1}\varepsilon$ of being open, and every edge that does not touch $\bigcup \mathcal{F}$ has probability q_{i+1} of being open. In particular, recalling that $\varepsilon = \delta(q_{i+1}, q_i) = (\log n)^{-(D+1)}$, it follows that there exists a constant $n_{10} = n_{10}(d, D, \lambda) \geq n_9$ such that if $n \geq n_{10}$ then

$$\mathbb{P}\left(C \xleftrightarrow{\beta_u} \bigcup(\mathcal{F} \setminus \{C\}) \mid \mathcal{C}_i = \mathcal{F}\right) \geq c_{-1}^2 \varepsilon^2 c_4 e^{-2[\log \log n]^{1/2}} \geq (\log m)^{-2D-3}.$$

Notice that under \mathbb{P} , the configurations $(\beta_u)_{u \in U}$ are independent. Moreover, this still holds after conditioning on the event that $\mathcal{C}_i = \mathcal{F}$. So, by independence, there exist constants $\lambda_0 = \lambda_0(d, D) = 8c_1^{-1}(2D+3)$ and $n_{11} = n_{11}(d, D, \lambda) \geq n_{10}$ such that if $\lambda \geq \lambda_0$ and $n \geq n_{11}$ then

$$\begin{aligned} \mathbb{P}\left(C \xleftrightarrow{B_{m/8} \cap \omega_{(i+1)}} \bigcup(\mathcal{F} \setminus \{C\}) \mid \mathcal{C}_i = \mathcal{F}\right) &\geq 1 - \prod_{u \in U} \mathbb{P}\left(C \xleftrightarrow{\beta_u} \bigcup(\mathcal{F} \setminus \{C\}) \mid \mathcal{C}_i = \mathcal{F}\right) \\ &\geq 1 - (1 - (\log m)^{-2D-3})^{\frac{1}{2}(\log m)^{c_1\lambda}} \\ &\geq 1 - e^{-2(\log n)^{c_1\lambda/2}}. \end{aligned}$$

Finally, since $|\mathcal{F}| \leq |S_{r_0}| \leq e^{(\log n)^D}$, we have by a union bound that if $\lambda \geq \lambda_0$ and $n \geq n_{11}$ then

$$\mathbb{P}(\mathcal{E}_i \mid \mathcal{C}_i = \mathcal{F}) \geq 1 - e^{(\log n)^D} e^{-2(\log n)^{c_1 \lambda/2}} \geq 1 - \frac{1}{2} e^{-(\log n)^{c_1 \lambda/2}}.$$

Since \mathcal{F} was arbitrary, this implies the claimed bound (3.6.23), completing the proof. \square

Completing the proof of the main theorem: The implications (C_0) and (C)

In this section we apply Proposition 3.6.1 to complete the proof of the Proposition 3.3.1 and hence of Theorem 8.1.1. Given Proposition 3.4.5, what remains is to verify the implications (C_0) and (C) .

Proof of Proposition 3.3.1. Let $D := 20$, and accordingly let $\lambda_0(d, D)$ and $c = c(d, D)$ be the constants from Proposition 3.6.1 with this value of D . Let $\lambda := \max\{\lambda_0, 10/c\}$, and let $K(d, D, \lambda)$ and $M(d, D, \lambda)$ be the corresponding constants from Proposition 3.6.1 that are there called K_1 and n_0 . (We want to avoid reusing the label n_0 .) We claim that if we define δ_0 using this value of K , then the implications (C_0) and (C) (for all $i \geq 1$) hold whenever $p_0 \geq 1/d$, $\delta_0 \leq 1$, and n_0 is sufficiently large with respect to d , which in particular guarantees that $n_0 \geq \max\{16, M\}$. The implication (C_0) is immediate (i.e., is a direct consequence of Proposition 3.6.1 after unpacking the definitions), so we will just explain how to prove the implication (C) .

Fix $i \geq 1$ and assume that $\text{FULL-SPACE}(i)$ holds and that $\text{CORRIDOR}(k)$ holds for all $1 \leq k \leq i$. Our goal is to establish that $\text{CORRIDOR}(i+1)$ holds provided that n_0 is sufficiently large with respect to d . This follows immediately from Proposition 3.6.1 if we can show that for every $n \in \mathcal{L}(G, 20) \cap [n_{i-1}, n_i]$ there exists some $b \leq \frac{1}{8}n^{1/3}$ such that

$$\left(\frac{K \log \log n}{\min\{\log n, \log \text{Gr}(b)\}} \right)^{1/4} \leq (\log \log n_i)^{-1/2} \quad \text{and} \quad \mathbb{P}_{p_i} \left(\text{Piv}[4b, n^{1/3}] \right) \leq (\log n)^{-1}. \quad (3.6.24)$$

Consider an arbitrary $n \in \mathcal{L}(G, 20) \cap [n_{i-1}, n_i]$ (assuming one exists; the claim is vacuous if not). We split into two cases according to whether $(\log n)^{2/3} \in \mathcal{L}(G, 20)$. First suppose that $(\log n)^{2/3} \notin \mathcal{L}(G, 20)$, so in particular

$$\text{Gr}((\log n)^{2/3}) \geq e^{((\log n)^{2/9})^{20}}.$$

By Corollary 3.2.4, provided n_0 is sufficiently large with respect to d , we know that

$$\mathbb{P}_{p_i} \left(\text{Piv}[4(\log n)^{2/3}, n^{1/3}] \right) \leq (\log n)^{-1}.$$

So in this case both conditions in (3.6.24) are satisfied for $b = (\log n)^{2/3} \leq \frac{1}{8}n^{1/3}$ provided that n_0 is sufficiently large with respect to d .

Now instead suppose that $(\log n)^{2/3} \in \mathcal{L}(G, 20)$. Note that we can always find some $k \in \{1, \dots, i\}$ such that $(\log n)^{2/3} \in [n_{k-2}, n_{k-1}]$. (This is why we took $n_{-1} = (\log n_0)^{1/2}$ as small as we did in the statement of the proposition.) So by our hypothesis that $\text{CORRIDOR}(k)$ holds for this particular value of k , we have that

$$\kappa_{p_k} \left(e^{[\log([\log n]^{2/3})]^{10}}, [\log n]^{2/3} \right) \geq e^{-(\log \log n_k)^{1/2}}. \quad (3.6.25)$$

We now claim that $b := \frac{1}{5} \min\{e^{(\log \log n)^9}, \text{Gr}^{-1}(e^{(\log n)^{1/10}})\}$ satisfies both conditions from eq. (3.6.24) provided that n_0 is sufficiently large with respect to d . The inequality $b \leq \frac{1}{8}n^{1/2}$ again holds trivially when n_0 is large. The definition of b ensures that $\log \text{Gr}(b) \geq \frac{1}{5} \log \text{Gr}(5b) \geq \frac{1}{5}(\log \log n)^9$, so that the first condition also trivially holds when n_0 is large. To see that the second condition holds, we apply Lemma 3.2.3 and Proposition 3.2.1 to obtain that there exists a constant C such that

$$\begin{aligned} \mathbb{P}_{p_i}(\text{Piv}[4b, n^{1/3}]) &\leq \mathbb{P}_{p_i}(\text{Piv}[1, n^{1/3}/2]) \cdot \frac{|S_{4b}|^2 \text{Gr}(5b)}{\min_{a,b \in S_{4b}}(a \xleftrightarrow{B_{5b}} b)} \\ &\leq C \left(\frac{(\log n)^{20}}{n} \right)^{1/4} e^{3(\log n)^{1/10} + (\log \log n_k)^{1/2}} \leq (\log n)^{-1} \end{aligned}$$

whenever n_0 is sufficiently large with respect to d , where we used the estimate (3.6.25) and the fact that $4b \leq e^{[\log([\log n]^{2/3})]^{10}}$ and $5b \geq (\log n)^{2/3}$ (when n_0 is sufficiently large) to bound the term $\min_{a,b \in S_{4b}}(a \xleftrightarrow{B_{5b}} b)$. This completes the proof. \square

3.7 Closing discussion and open problems

Joint continuity of the supercritical infinite cluster density

Recall that \mathcal{G}^* is the space of all infinite, connected, transitive graphs that are not one-dimensional, which we endow with the local topology, and recall that for all $p \in (0, 1)$ and $G \in \mathcal{G}^*$, the **infinite cluster density** is defined to be $\theta(G, p) := \mathbb{P}_p^G(o \leftrightarrow \infty)$. Consider a sequence $(G_n)_{n \geq 1}$ in \mathcal{G}^* converging to some $G \in \mathcal{G}^*$. The main result of the present paper Theorem 8.1.1 states that $p_c(G_n) \rightarrow p_c(G)$. One could ask the following more refined question: does $\theta^{G_n} \rightarrow \theta^G$ pointwise? (One can observe from the mean-field lower bound, say, that a positive answer to this question would imply our result that $p_c(G_n) \rightarrow p_c(G)$.) For $p < p_c(G)$, it follows immediately from the lower semi-continuity of p_c that $\theta(G_n, p) \rightarrow \theta(G, p) = 0$, so the only non-trivial cases are when $p = p_c(G)$ and $p > p_c(G)$. The case $p = p_c(G)$ appears to be hard. Indeed, if one could prove this result in the case of toroidal slabs $(\mathbb{Z}^2 \times \mathbb{Z}/n\mathbb{Z})_{n \geq 1}$ converging to the cubic grid \mathbb{Z}^3 , then it would follow from the main result of [DST16] that $\theta(\mathbb{Z}^3, p_c(\mathbb{Z}^3)) = 0$, a notorious open question. We have nothing interesting to say about this case. However, in our other work [EH23+a]

we prove the following theorem, which together with Theorem 8.1.1 completely resolves the case when $p > p_c(G)$.

Theorem 3.7.1 ([EH23+a]). *Let $(G_n)_{n \geq 1}$ be a sequence in \mathcal{G}^* that converges in the local topology to some $G \in \mathcal{G}^*$. Then $\theta(G_n, p) \rightarrow \theta(G, p)$ as $n \rightarrow \infty$ for every $p > \limsup_{n \rightarrow \infty} p_c(G_n)$.*

Note that the main theorem of [EH23+a] is much more general than this and also establishes a form of locality for the density of the *giant cluster* on *finite* transitive graphs that may have divergent degree. Together with Theorem 8.1.1 this theorem yields the following elementary corollary.

Corollary 3.7.2. *$\theta(G, p)$ is continuous on the open set $\{(G, p) : G \in \mathcal{G}^*, p > p_c(G)\}$. Moreover, if $(G_n)_{n \geq 1}$ is a sequence in \mathcal{G}^* that converges in the local topology to some $G \in \mathcal{G}^*$ then $\theta(p, G_n) \rightarrow \theta(p, G)$ as $n \rightarrow \infty$ for each $p \in [0, 1] \setminus \{p_c(G)\}$.*

Let us roughly indicate how the tools built in Section 3.4 could be used to give an alternative proof of Theorem 3.7.1. This alternative proof is less general and (arguably) more involved than the one given in [EH23+a], but the result is quantitatively stronger: it can be used to prove that $\theta(G, p)$ is not just continuous but even locally Hölder continuous on the supercritical set (with the power in the definition of Hölder continuity possibly degenerating near the boundary of the set).

Let \mathcal{G}_d^* denote the set of infinite transitive graphs with vertex degree exactly d that are not one-dimensional. As explained in detail in [EH23+a], it suffices to prove a tail estimate on the size of finite clusters in supercritical percolation that is uniform over \mathcal{G}_d^* for each d , i.e. it suffices to prove that

$$\lim_{m \rightarrow \infty} \sup_{G \in \mathcal{G}_d^*} \sup_{p \geq p_c(G) + \varepsilon} \mathbb{P}_p^G(m \leq |K_o| < \infty) = 0 \quad \text{for all } \varepsilon > 0 \text{ and } d \geq 1, \quad (3.7.1)$$

where we recall that K_o denotes the cluster of the root vertex o . We will focus on proving an estimate of this form for the set of unimodular graphs \mathcal{U}_d^* instead of \mathcal{G}_d^* ; in the nonunimodular case much stronger results (with optimal dependence on $p - p_c$ and m) can be proven by invoking the results of [Hut20e; Hut22] as explained in detail in [EH23+a]. Our proof will yield quantitatively that for every $\varepsilon > 0$ and $d \geq 1$ there exist constants C and c such that

$$\sup_{G \in \mathcal{U}_d^*} \sup_{p \geq p_c(G) + \varepsilon} \mathbb{P}_p^G(m \leq |K_o| < \infty) \leq Cm^{-c};$$

running the proof of continuity with this quantitative estimate yields the aforementioned local Hölder continuity of $\theta(G, p)$. This bound is quantitatively much better than the bound coming from the proof in [EH23+a]. On the other hand, it is also much worse than the conjectured

optimal bounds, which are stretched exponential in m (see Section 5.3 of [HH21c]). Having any *superpolynomial* tail bound would imply that $\theta(G, p)$ is a smooth function of $p \in (p_c(G), 1]$ for each fixed G ; for \mathbb{Z}^d and for nonamenable graphs it is known that the density is not just smooth but real analytic on this set [HH21c; GP23].

Let $\varepsilon > 0$, $d \geq 1$, $G \in \mathcal{U}_d^*$, $p \geq p_c(G) + \varepsilon$, $m \geq 1$, and $\eta > 0$ be arbitrary, and suppose that $\mathbb{P}_p^G(m \leq |K_o| < \infty) \geq \eta$. It suffices to prove that m is necessarily bounded above by some constant $M(\varepsilon, d, \eta) < \infty$. Let \mathbb{P} denote the canonical monotone coupling $(\omega_q)_{q \in [0,1]}$ of the percolation measures $(\mathbb{P}_q^G)_{q \in [0,1]}$. By the mean-field lower bound and transitivity, one can find vertices $u, v \in V(G)$ such that

$$\mathbb{P}\left(|K_u(\omega_{p-\varepsilon/2})| \geq m \text{ and } |K_v(\omega_p)| \geq m \text{ but } u \leftrightarrow v\right) \geq \frac{\eta\varepsilon}{2},$$

where $K_u(\omega_{p-\varepsilon/2})$ denotes the cluster of u in $\omega_{p-\varepsilon/2}$. In particular, writing $\mathbb{G}_{1/m}$ for the law of a ghost-field \mathcal{G} of intensity $1/m$ on the whole vertex set $V(G)$,

$$\mathbb{G}_{\frac{1}{m}} \otimes \mathbb{P}(u \xrightarrow{\omega_{p-\varepsilon/2}} \mathcal{G} \xrightarrow{\omega_p} v \text{ but } u \xleftrightarrow{\omega_p} v) \geq \left(1 - \frac{1}{e}\right)^2 \frac{\eta\varepsilon}{2} \geq \frac{\eta\varepsilon}{8}. \quad (3.7.2)$$

Assume for now that G is amenable so that there is at most one infinite cluster \mathbb{P}_q^G -almost surely for every $q \in [p - \varepsilon/2, p]$. Then by Lemma 3.4.4 with $(X, A, Y) := (\{u\}, V(G), \{v\})$, one can deduce that for some constants $c_3(\varepsilon, d) > 0$ and $C(\varepsilon, d) < \infty$,

$$\mathbb{G}_{\frac{1}{m}} \otimes \mathbb{P}(u \xrightarrow{\omega_{p-\varepsilon/2}} \mathcal{G} \xrightarrow{\omega_p} v \text{ but } u \xleftrightarrow{\omega_p} v) \leq Cm^{-c_3}. \quad (3.7.3)$$

By combining (3.7.2) and (3.7.3), we deduce that $m \leq M(\varepsilon, d, \eta) := (8C/(\eta\varepsilon))^{1/c_3}$, as required.

Finally, to handle the case when G is nonamenable (but still unimodular), one can still run essentially the same argument as above but with some technical modifications to handle the possible existence of multiple infinite clusters. The key difference is that instead of tracking connections $u \leftrightarrow v$ as usual, we instead track *wired* connections $u \xleftrightarrow{\text{wired}} v$ where

$$\{u \xleftrightarrow{\text{wired}} v\} := \{u \leftrightarrow v\} \cup \{u \leftrightarrow \infty \text{ and } v \leftrightarrow \infty\}.$$

One can verify that the proof of Lemma 3.4.4 works just as well with this alternative notion of connectivity, without requiring the hypothesis about the uniqueness of the infinite cluster. The rest of the argument explained above can then be adapted to work with this wired notion of connectivity also.

As explained in [EH23+a], the estimate (6.1.2) has the following nice interpretation. Consider the space $(0, 1) \times \mathcal{G}^* \rightarrow (0, 1)$ with the product topology, and consider the function $\theta : (0, 1) \times \mathcal{G}^* \rightarrow (0, 1)$ mapping $(p, G) \mapsto \theta^G(p)$. It is natural to ask whether θ is continuous as a function of two variables. One can show that a priori, θ is continuous if and only if the conclusion of Theorem 8.1.1 holds (locality of p_c), the estimate (6.1.2) holds (i.e., there is a uniform tail bound on supercritical finite clusters), and $\theta(G, p_c(G)) = 0$ for every $G \in \mathcal{G}^*$ (continuity of the phase transition), the last statement being one of the most important open conjectures in the general study of percolation in \mathcal{G}^* . (In this decomposition, (6.1.2) handles the interior of the supercritical region $\mathcal{S} := \{(p, G) : p > p_c(G)\}$, Theorem 8.1.1 implies that this region is open, and the continuity conjecture handles the boundary values.)

Finite graphs. As mentioned above, in [EH23+a] we prove versions of Theorem 3.7.1 and eq. (6.1.2) that also apply to families of bounded-degree *finite* transitive graphs. The above sketches work just as well in this context too; we have stated things in terms of infinite graphs purely for simplicity. Moreover, the above sketch can be used to give an alternative proof that for supercritical percolation on bounded-degree finite transitive graphs, the giant cluster is unique and has concentrated density, recovering the results of our two papers [EH21a; EH23+a] in this case. Note however that all of the tools from Section 3.4 break down rather badly when working with families of finite graphs that have large vertex degrees (e.g. vertex degrees that grow at least as a power of the total number of vertices), partly because for such graphs the emergence of a giant cluster can occur around values of p close to 0. To handle this more general setting of arbitrary finite transitive graphs, we know of no alternative proofs of the uniqueness or concentration of the supercritical giant cluster to those we give in [EH21a; EH23+a].

The p_c gap and its witnesses

Since \mathcal{G}_d^* is compact for each $d \geq 1$, it is a consequence of Theorem 8.1.1 that p_c attains its maximum on \mathcal{G}_d^* for each $d \geq 1$. In [PS23b], Panagiotis and Severo improved upon the results of [DGRSY20; HT21a] to establish that there exists a *universal* $\varepsilon > 0$ (independent of the degree) such that every Cayley graph with $p_c < 1$ has $p_c \leq 1 - \varepsilon$; Lyons, Mann, Tessera, and Tointon [LMTT23] give the explicit bound $\varepsilon \geq \exp(-\exp(17\exp(100 \cdot 8^{100})))$. Presumably a similar result holds for transitive graphs that are not Cayley. The following natural conjecture would strengthen this result, and would also imply by Theorem 8.1.1 that p_c attains its global maximum on \mathcal{G}^* .

Conjecture 3.7.3. *There exists a universal constant C such that if G is an infinite, connected, transitive, simple graph of vertex degree d that is not one-dimensional, then $p_c(G) \leq C/d$.*

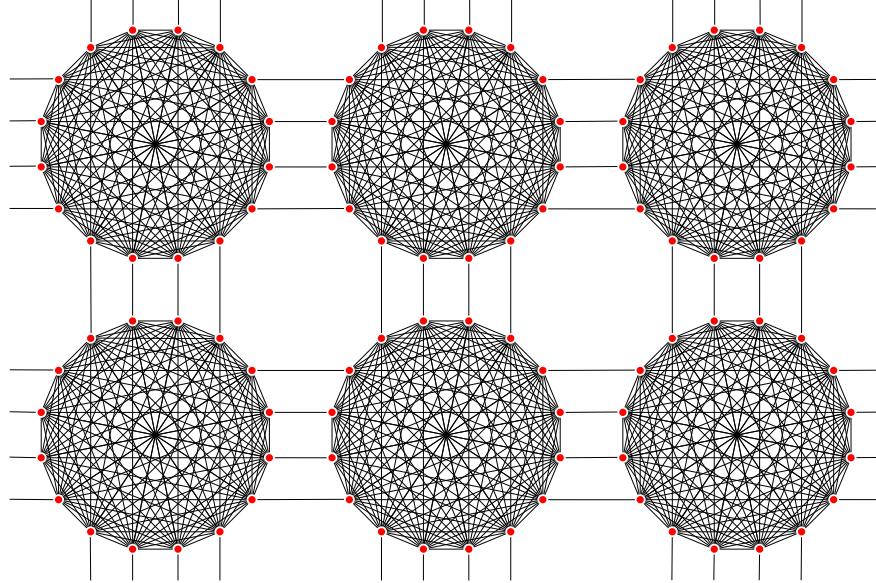


Figure 3.5: The transitive graph formed by laying out copies of K_{4n} in an infinite square grid as above has $p_c \geq (4 \log 2 - o(1))/\deg$ as $n \rightarrow \infty$ as can be seen by coupling with bond percolation on \mathbb{Z}^2 and using a Poisson approximation. Since $4 \log 2 \approx 1.2 > 1$, this shows that the asymptotic estimate $p_c \sim 1/\deg$ can fail for high-degree vertex-transitive graphs even when these graphs are not one-dimensional in any sense. An exact asymptotic estimate $p_c \sim C/\deg$ can be proven with a little further work. (Indeed, the constant $C \approx 3.095$ is the unique solution to the equation $C(1 + C^{-1}W[-e^{-C}C])^2 = 4 \log 2$ where W is the Lambert W function.)

For many natural families of high-degree graphs we have the stronger statement that $p_c \sim 1/\deg$ as the degree diverges. For example, this holds for \mathbb{Z}^d as $d \rightarrow \infty$ by a theorem of Kesten [Kes90] (see also [ABS04a]). Moreover, this is not just a high-dimensional phenomenon: Penrose [Pen93] proved that a similar-estimate holds for the “spread-out” d -dimensional lattice, in which $x, y \in \mathbb{Z}^d$ are connected by an edge whenever $\|x - y\| \leq R$, when d is fixed and $R \rightarrow \infty$. On the other hand, the example illustrated in fig. 3.5 shows that high-degree transitive graphs do not always have $p_c \sim 1/\deg$ even when they are not one-dimensional in any sense.

Once one knows that p_c attains a maximum (either globally on \mathcal{G}^* or on \mathcal{G}_d^*), it becomes interesting to understand which graphs attain this maximum. Martineau and Severo [MS19] proved that p_c is strictly increasing under quotients, so that any maximal graph must have no non-trivial quotients in \mathcal{G}^* . It seems reasonable to believe that the maximal graph would be a lattice of low degree and in low dimension. Consulting tables of numerical values of p_c for these lattices (as can be found on https://en.wikipedia.org/wiki/Percolation_threshold) leads to the highly speculative conjecture that p_c is maximized by the so-called super-kagome lattice (a.k.a. 3-12

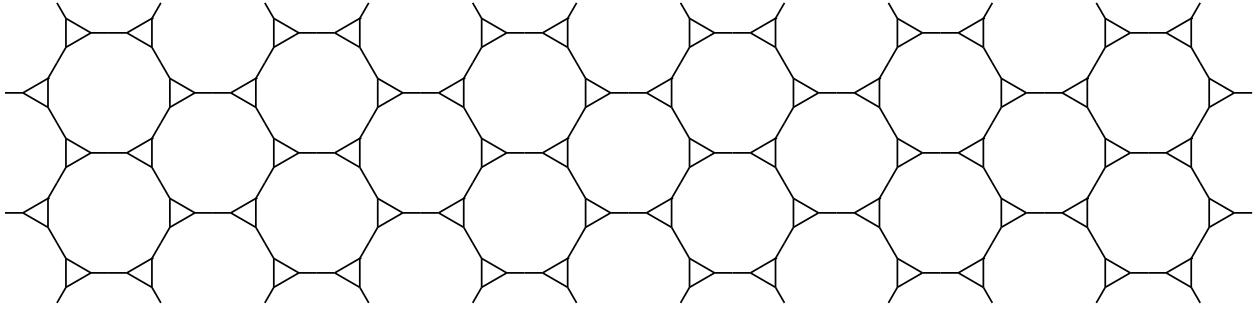


Figure 3.6: The 3-12 (a.k.a. super-kagome) lattice is the current best candidate for the transitive graph with the highest non-trivial value of p_c for bond percolation. Its critical value $p_c \approx 0.7404\dots$ has been estimated to great precision numerically in [SJ20]. The transitive graph with the next highest value of p_c to have been investigated numerically is the truncated trihexagonal lattice, which has $p_c \approx 0.6937$.

lattice); see fig. 3.6 for an illustration.

Problem 3.7.4. *Investigate the transitive graphs in \mathcal{G}_d^* that maximize p_c for each degree d , as well as the global maximum in \mathcal{G}^* if this maximum exists. Are these maxima uniquely attained? Does $p_c : \mathcal{G}^* \rightarrow [0, 1]$ attain its maximum uniquely at the 3-12 lattice (a.k.a. super-kagome lattice), which has $p_c \approx 0.7404207$? When restricted to edge-transitive graphs, does p_c attain its unique maximum at the hexagonal lattice, which has $p_c = 0.65270\dots = 1 - 2 \sin(\pi/18)$?*

One may wish to restrict attention to simple graphs. Similar questions have been investigated for self-avoiding walk by Grimmett and Li [GL20].

Acknowledgments

This work was supported by NSF grant DMS-2246494. TH thanks Ariel Yadin for discussions on the history of the “cool inequality” and thanks Vincent Tassion for discussions on the plausibility of Conjecture 3.7.3. Both authors thank Geoffrey Grimmett, Russ Lyons, and Sébastien Martineau for helpful comments on an earlier version of the manuscript.

UNIFORM FINITE PRESENTATION FOR GROUPS OF POLYNOMIAL GROWTH

Joint with Tom Hutchcroft

Abstract

We prove a quantitative refinement of the statement that groups of polynomial growth are finitely presented. Let G be a group with finite generating set S and let $\text{Gr}(r)$ be the volume of the ball of radius r in the associated Cayley graph. For each $k \geq 0$, let R_k be the set of words of length at most 2^k in the free group F_S that are equal to the identity in G , and let $\langle\langle R_k \rangle\rangle$ be the normal subgroup of F_S generated by R_k , so that the quotient map $F_S/\langle\langle R_k \rangle\rangle \rightarrow G$ induces a covering map of the associated Cayley graphs that has injectivity radius at least $2^{k-1} - 1$. Given a non-negative integer k , we say that (G, S) has a **new relation on scale k** if $\langle\langle R_{k+1} \rangle\rangle \neq \langle\langle R_k \rangle\rangle$. We prove that for each $K < \infty$ there exist constants n_0 and C depending only on K and $|S|$ such that if $\text{Gr}(3n) \leq K \text{Gr}(n)$ for some $n \geq n_0$, then there exist at most C scales $k \geq \log_2(n)$ on which G has a new relation. We apply this result in another paper as part of our proof of Schramm's locality conjecture in percolation theory.

4.1 Introduction

It is a seminal theorem of Gromov [Gro81c] (see also [Kle10; Oza18]) that a finitely generated group has polynomial volume growth if and only if it is virtually nilpotent. This theorem and its extension to transitive graphs due to Trofimov [Tro85] are of foundational importance in the study of geometry and probability on transitive graphs, implying in particular that every transitive graph of polynomial growth has a well-defined volume growth dimension and that this dimension is an integer. In probability, these theorems are often used together with the isoperimetric inequality of Coulhon and Saloff-Coste [CS93] to prove results for general transitive graphs via a “structure vs. expansion” dichotomy: that is, proceeding by a case analysis according to whether the graph is virtually nilpotent or satisfies a d -dimensional isoperimetric inequality for every $d < \infty$. Important results in probability employing the structure theory of transitive graphs in this way include Varopoulos's theorem [Var86] that an infinite transitive graph is recurrent for simple random walks if and only if it has linear or quadratic volume growth, and Duminil-Copin, Goswami, Severo, Raoufi, and

Yadin's proof that transitive graphs admit a percolation phase transition if and only if they have superlinear growth [Dum+20c].

Over the last twenty years, an extensive literature in *approximate group theory* has been developed establishing *finitary* versions of Gromov's theorem and Trofimov's theorem, highlights of which include [BGT12b; ST10a; Hru12; TT21b; BGT11]. See [Bre14] for a detailed overview, [Toi20a] for a textbook introduction, and [Toi20b] for a concise survey. For groups, this theory culminated in the celebrated work of Breuillard, Green, and Tao [BGT12b], a special case of whose results can be stated¹ as follows. Given a group G and a finite generating set S , we write $\text{Gr}(r) = \text{Gr}_{G,S}(r)$ for the cardinality of the ball of radius r in the Cayley graph $\text{Cay}(G, S)$.

Theorem 4.1.1 (Breuillard, Green, and Tao 2012). *For each $K \geq 1$ there exist constants $r_0 = r_0(K)$ and $C = C(K)$ such that the following holds. Let G be a group with finite generating set S , and suppose that there exists $r \geq r_0$ such that $\text{Gr}(3r) \leq K \text{Gr}(r)$. Then $\text{Gr}(mr) \leq m^C \text{Gr}(r)$ for every $m \geq 3$ and there exists a finite normal subgroup $Q \triangleleft G$ such that:*

1. *Every fibre of the projection $\pi : G \rightarrow G/Q$ has diameter at most Cr .*
2. *G/Q has a nilpotent normal subgroup N of rank, step and index at most C .*
3. *The projection $x \mapsto \pi(x)$ is a $(1, Cr)$ -quasi-isometry from $\text{Cay}(G, S)$ to $\text{Cay}(G/Q, \pi(S))$.*

Here, we recall that a function $\phi : V_1 \rightarrow V_2$ between the vertex sets of two graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ is said to be an **(α, β) -quasi-isometry** (a.k.a. **rough isometry**) if

$$\alpha^{-1}d(x, y) - \beta \leq d(\phi(x), \phi(y)) \leq \alpha d(x, y) + \beta$$

for every $x, y \in V_1$, and every vertex $z \in V_2$ is within distance at most β of $\phi(V_1)$. (The second property holds automatically if ϕ is surjective.)

Informally, the Breuillard-Green-Tao theorem states that polynomial growth at one sufficiently large scale forces the group to have polynomial growth at every subsequent scale, and moreover to be metrically “well-modelled” by a nilpotent group at all larger scales. Similar theorems for

¹We state their theorem in a ‘metric’ form that is convenient for our applications, and which is adapted from Tessera and Tointon’s structure theorem for *vertex-transitive graphs* of polynomial growth [TT21b, Theorem 2.3]. Indeed, the statement given below is equivalent to the special case of their theorem in which the graph Γ is the Cayley graph of G , together with the growth bound of [BGT12b, Corollary 11.9].

vertex-transitive graphs that are not necessarily Cayley graphs have recently been established in the work of Tessera and Tointon [TT21b; TT18].

These results have recently found many probabilistic applications, particularly for problems concerning *families* of transitive graphs (such as sequences of finite transitive graphs converging to an infinite graph); such problems often require estimates that are “uniform in the graph”, so that structure theoretic results invoked in their solutions must typically be finitary. Results proven using finitary structure theory include a finite-graph version of Varopoulos’s theorem [TT20b], universality theorems for cover time fluctuations [BHT22], locality of the critical probability for graphs of polynomial growth [CMT23b], non-triviality of the supercritical phase for percolation on finite transitive graphs [HT21b], and “gap at 1” theorems for the critical probability on infinite vertex transitive graphs [HT21b; Lyo+23b; PS23b]. Several of these works exploit finitary versions of the “structure vs. expansion” dichotomy provided by the finitary structure theory of [BGT12b; TT21b], with key technical difficulties arising from the fact that the same graph may exhibit different sides of this dichotomy at different scales.

Uniform finite presentation

Since virtually nilpotent groups are finitely presented, it is a consequence of Gromov’s theorem that every group of polynomial volume growth is finitely presented. The purpose of this paper is to prove a *uniform* version of this fact, stating roughly that every group of polynomial growth has a bounded number of scales witnessing a new relation after the first scale that polynomial growth is witnessed. This result is used in our work [EH23c] as part of our proof of Schramm’s locality conjecture for Bernoulli bond percolation [BNP11a], where it plays an important part in our “uniformization” of the methods of Contreras, Martineau, and Tassion [CMT21]. A comparison of our results with the previous literature is given at the end of this section.

Let us now state our result formally. Let G be a group with finite generating set S , so that $G \cong F_S/R$ for some normal subgroup R of F_S . For each $n \geq 0$, let R_n be the set of words of length at most 2^n in the free group F_S that are equal to the identity in G , and let $\langle\langle R_n \rangle\rangle$ be the normal subgroup of F_S generated by R_n , so that the quotient map $F_S/\langle\langle R_n \rangle\rangle \rightarrow G$ induces a covering map of the associated Cayley graphs that has injectivity radius at least $2^{n-1} - 1$ (see Lemma 4.5.6). We say that (G, S) has a **new relation on scale n** if $\langle\langle R_{n+1} \rangle\rangle \neq \langle\langle R_n \rangle\rangle$. A finitely generated group G is finitely presented if and only if it has a new relation on at most finitely many scales, so that the following theorem can indeed be thought of as stating that groups of polynomial growth are “uniformly finitely presented”.

Theorem 4.1.2. *For each $K, k < \infty$ there exist constants $r_0 = r_0(K)$ and $C = C(K, k)$ such that if G is a group and S is a finite generating set for G with $|S| \leq k$ whose growth function Gr satisfies $\text{Gr}(3r) \leq K \text{Gr}(r)$ for some integer $r \geq r_0$ then*

$$\#\left\{n \in \mathbb{N} : n \geq \log_2(r) \text{ and } (G, S) \text{ has a new relation on scale } n\right\} \leq C.$$

Remark 4.1.1. Considering the abelian group $\prod_{i=1}^k (\mathbb{Z}/n_i\mathbb{Z})$ with its standard generating set, where n_1, \dots, n_k are arbitrary, we see that it is not possible to control on *which* scales we find a new relation; we only claim that the total *number* of scales on which we find a new relation is bounded.

We will prove the following theorem about the number of times we find an “unexpected element” during a breadth-first exploration of a (not necessarily normal) subgroup of a group of polynomial growth; we will see in Section 4.5 that this theorem easily implies Theorem 8.1.1.

Theorem 4.1.3 (Breadth-first exploration of subgroups). *For each K and k there exist constants $r_0 = r_0(K)$ and $C = C(K, k)$ such that the following holds. Let G be a group with finite generating set S satisfying $|S| \leq k$, let H be a subgroup of G , and for each $n \geq 1$ let H_n be the subgroup of H generated by elements that have word length at most 2^n in (G, S) . If $r \geq r_0$ is such that $\text{Gr}(3r) \leq K \text{Gr}(r)$ then*

$$\#\{n \geq \log_2 r : H_{n+1} \neq H_n\} \leq C.$$

Remark 4.1.2. We believe that it should be possible to take the constants in Theorems 4.1.3 and 8.1.1 to be independent of the size of the generating set. We do not pursue this here.

Other previous results. A more classical way to quantify the sense in which a presentation is finite is through *Dehn functions*, *filling length functions*, and so-called *isoperimetric functions* (which do not refer to the same kind of isoperimetry mentioned in our above discussion of the structure vs. expansion dichotomy); see [Bri02] for an overview and [GHR03] for results on nilpotent groups. As far as we can tell, however, the literature on these notions focusses on asymptotic properties of a fixed group and is not suitable for the kind of uniform-in-the-group results we wish to prove. Besides this, the notion of uniform finite presentation we consider is also rather different from these notions in terms of Dehn functions etc.

There is a striking resemblance between our theorem and the following theorem of Tao [Tao17b] (see also [TT17, Appendix A]), which also relies on the structure theory of Breuillard, Green, and Tao.

Theorem 4.1.4 [Tao17b], Theorem 1.9). For each non-negative integer d there exist constants m_0 and C depending only on d such that if G is a group, S is a finite, symmetric generating set for G containing the identity and satisfying $|S^{nm}| \leq m^d |S^n|$ for some integers $n \geq 1$ and $m \geq m_0$ then there exists a continuous, piecewise-linear, non-decreasing function $f : [0, \infty) \rightarrow [0, \infty)$ with $f(0) = 0$ that has at most C pieces, each of which has slope equal to an integer bounded by C , such that

$$\left| \log \frac{|S^{knm}|}{|S^{nm}|} - f(\log k) \right| \leq C$$

for every integer $k \geq 1$.

Informally, this theorem states that, once we witness polynomial growth on a sufficiently large scale, the log-log plot of the growth function is well-approximated by a continuous, piecewise-linear function with bounded, integer valued slopes and a bounded number of “kinks” connecting the different pieces.

Naively, one might hope that our bounded number of scales on which a new relation occurs are in correspondence with Tao’s bounded number of scales on which the growth function has a “kink” in its log-log plot. Unfortunately this is not the case, at least when one allows generating sets of unbounded size: one can have a new relation without having a kink, and can have a kink without having a new relation. Indeed, as explained in [Tao17b, Example 1.11], taking

$$G = \begin{pmatrix} 1 & \mathbb{Z} & \mathbb{Z} \\ 0 & 1 & \mathbb{Z} \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad S = \begin{pmatrix} 1 & [-N, N] & [-N^3, N^3] \\ 0 & 1 & [-N, N] \\ 0 & 0 & 1 \end{pmatrix}$$

for a large integer N yields

$$\log \frac{|S^n|}{|S|} = \begin{cases} 3 \log n \pm O(1) & 1 \leq n \leq N \\ 4 \log n - \log N \pm O(1) & n > N, \end{cases}$$

so that this example’s growth function has a kink at scale $\log_2 N$. On the other hand the pair (G, S) does *not* have a new relation at any $k \geq 3$, and in particular does not have a new relation on the scale where it has a kink when N is large. Indeed, the relations of (G, S) are generated by the usual relations for the Heisenberg group

$$\left[\begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix} \right] = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

$$\begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \text{and} \quad \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}$$

together with the following three sets of relations relating the extra generators in S to the standard generators:

$$\begin{aligned} \left\{ \begin{pmatrix} 1 & a \pm 1 & c \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & a & c \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & \pm 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} : a, b \in [-N, N], c \in [-N^3, N^3] \right\}, \\ \left\{ \begin{pmatrix} 1 & a & c \\ 0 & 1 & b \pm 1 \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & \pm 1 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & a & c \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix} : a, b \in [-N, N], c \in [-N^3, N^3] \right\}, \\ \left\{ \begin{pmatrix} 1 & a & c \pm 1 \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & a & c \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & \pm 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} : a, b \in [-N, N], c \in [-N^3, N^3] \right\}. \end{aligned}$$

These relations all have word length at most five in (G, S) , so that the example has the desired properties. Conversely, taking the direct product $G \times \mathbb{Z}/N$ with generating set $S \times \{-1, 0, 1\}$ yields an example where there is a new relation at scale $\log_2 N$ but where the growth function does not have any kinks.

About the proof. It is natural to describe the proof of Theorems 4.1.3 and 8.1.1 “backwards”, as a sequence of reductions, although we have written it “forwards” as a sequence of extensions and generalizations. In this backwards description, the “first” step (which is the last part of the paper) is to reduce from groups of polynomial growth to nilpotent groups of bounded step using the Breuillard-Green-Tao theorem. Next, this statement about nilpotent groups is in turn reduced to an analogous statement about breadth-first exploration of discrete subgroups in a *Carnot group*, a simply connected nilpotent Lie group carrying the additional structure of a *stratification* and *homogeneous left-invariant metric*. Nilpotent groups are related to Carnot groups for example by *Pansu’s theorem* [Pan83; BL13], which states roughly that the large-scale geometry of a finitely generated nilpotent group is well-modelled by an appropriate Carnot group equipped with a left-invariant homogeneous metric. Finally, this statement about Carnot groups is reduced to a statement about *vector spaces* using the close connection between the discrete subgroups of a simply connected nilpotent Lie group and the additive bracket-closed subgroups of its associated Lie algebra. This step of the reduction is the most involved part of the paper, with the connection between additive and multiplicative lattices being developed at length in Section 4.4. This ends the chain of reductions, and leaves us with a problem we must actually solve directly: Bounding the number of times we find an “unexpected element” of a discrete subgroup of \mathbb{R}^d as we explore the subgroup with an increasing family of convex, symmetric sets. This is done in Section 4.3 as an application of Minkowski’s second theorem, a classical result in the geometry of numbers.

Let us stress again that we have described the argument here in the opposite order to the way we carry it out, so that the result about subgroups of \mathbb{R}^d is the first thing we prove.

Disclaimer: Neither author is an expert in approximate groups, Lie theory, or the geometry of numbers. As such, it is likely that we have included a larger amount of detail in the proofs than would be considered necessary by experts, or have re-derived known results from scratch. While we have attempted to provide appropriate attribution to the intermediate results of the paper as much as possible, we would be happy to receive comments and corrections from experts.

Remark 4.1.3. Since the present paper first appeared, two relevant papers by Tessera and Tointon have since appeared. The first paper [TT24] establishes a stronger version of Theorem 8.1.1 as a corollary, with optimal bounds on the number of scales at which new relations appear that do not depend on the degree. The proof is very different than that given here. The second paper [TT23] allows us to replace the “small tripling” hypothesis $\text{Gr}(3r) \leq K \text{Gr}(r)$ in Theorems 4.1.3 and 8.1.1 with the more natural hypothesis of “small doubling” $\text{Gr}(2r) \leq K \text{Gr}(r)$ because, as the title of the paper says: *Small doubling implies small tripling on large scales*.

4.2 Background on nilpotent groups and Lie groups

In this section we review the relevant background material and establish some notational conventions. We have included a rather thorough account of the basic theory with the hope that our paper can be easily understood by probabilists.

Given a group G , the **commutator** of two elements $x, y \in G$ is defined by $[x, y] = xyx^{-1}y^{-1}$. The **lower central series** of G is defined recursively by $G_1 = G$ and $G_{i+1} = [G_i, G]$ for each $i \geq 1$, where we write $[A, B] := \{[a, b] : a \in A, b \in B\}$ for subsets A and B of G . The group G is said to be **nilpotent** if $G_{i+1} = \{\text{id}\}$ for all sufficiently large i , with the minimal such i denoted by s and known as the **step** of G . A group is said to be **virtually nilpotent** if it has a nilpotent subgroup of finite index. Given $s \geq 1$ and a set S , the **free step s nilpotent group** $N_{s,S}$ is defined to be the quotient of the free group F_S by the step- s nilpotency relations, which state that all iterated commutators of length at least $s+1$ are equal to the identity. The free step s nilpotent group $N_{s,S}$ can also be defined up to unique S -preserving isomorphism by the universal property that it is nilpotent of step at most s , contains S , and every function from S to a nilpotent group of step at most s can be uniquely extended to a homomorphism from $N_{s,S}$ to that group.

(Nilpotent) Lie groups and the Baker-Campbell-Hausdorff formula

Recall that a (real) **Lie group** is a group that is also a finite-dimensional real smooth manifold, in such a way that the group operations of multiplication and inversion are smooth maps $G \times G \rightarrow G$ and $G \rightarrow G$. By Gleason, Montgomery, and Zippin’s solution to Hilbert’s fifth problem [MZ52; Gle52], one can equivalently define a Lie group as a group that is also a finite-dimensional

topological manifold with continuous multiplication and inversion operations; such a group carries a unique smooth structure compatible with its algebraic structure. More generally, Yamabe [Yam50] proved that every locally compact, connected topological group is a projective limit of Lie groups (possibly of divergent dimension). These facts underlie the ubiquity of Lie groups in the scaling limit theory of discrete groups, and in particular are used directly in Gromov’s original proof of his polynomial growth theorem [Gro81c]. Further background on these topics can be found in [Tao14]. For our purposes, Lie groups become relevant primarily via a theorem of Pansu, which allows us to approximate the balls in the Cayley graph of a nilpotent group in terms of the balls in a certain *left-invariant homogeneous metric on a Carnot group*; this is explained in Section 4.2.

Lie algebras. A **Lie algebra** \mathfrak{g} is a vector space equipped with a binary operation, the Lie bracket $[\cdot, \cdot] : \mathfrak{g} \times \mathfrak{g} \rightarrow \mathfrak{g}$, that is bilinear, antisymmetric ($[X, Y] = -[Y, X]$ for all $X, Y \in \mathfrak{g}$), and satisfies the Jacobi identity ($[X, [Y, Z]] + [Y, [Z, X]] + [Z, [X, Y]] = 0$ for all $X, Y, Z \in \mathfrak{g}$). A subset A of a Lie algebra is said to be **bracket-closed** if $[X, Y] \in A$ for every $X, Y \in A$. For ease of reading, we will loosely follow the convention that points in a Lie algebra are denoted using upper-case letters, while points in a Lie group are denoted using lower-case letters. We will also assume without further comment that all Lie algebras are finite-dimensional.

To each Lie group G , we can associate a Lie algebra \mathfrak{g} arising from the tangent space at the identity; the details of this construction are not important to us and can be found in any textbook on the subject. In the concrete case that G is a Lie subgroup of a general linear group GL_n for some $n \geq 1$, the affine space $I + \mathfrak{g}$ is precisely the tangent space at the identity to G in the space of all $n \times n$ matrices, so that \mathfrak{g} is a Lie subalgebra (i.e. a bracket-closed linear subspace) of the Lie algebra \mathfrak{gl}_n of all $n \times n$ matrices with Lie bracket defined by the commutator $[X, Y] = XY - YX$. (In particular, \mathfrak{gl}_n is the Lie algebra associated to the Lie group GL_n .) In fact this case is not particularly special: Ado’s theorem states that every Lie algebra is isomorphic to a Lie subalgebra of \mathfrak{gl}_n for some $n \geq 1$ [Tao14, Chapter 2.3]. (Ado’s theorem does *not* imply that every Lie *group* is isomorphic to a Lie subgroup of a general linear group, although it does imply a “local” version of the same claim.)

The fundamental theorems of Lie (see e.g. [Tao14, Chapter 2.5.1]) state in particular that there is a one-to-one correspondence between (isomorphism classes of) Lie algebras and *simply connected* Lie groups. On the other hand, Lie groups that are connected but not simply connected have universal covers which are simply connected Lie groups with the same Lie algebra.

The **lower central series** of the Lie algebra \mathfrak{g} is defined recursively by $\mathfrak{g}_1 = \mathfrak{g}$ and $\mathfrak{g}_{i+1} = [\mathfrak{g}_i, \mathfrak{g}]$ for each $i \geq 1$, where if A and B are two subsets of \mathfrak{g} then we write $[A, B] = \{[a, b] : a \in A, b \in B\}$.

It is a simple consequence of the Jacobi identity that $[\mathfrak{g}_i, \mathfrak{g}_j] \subseteq \mathfrak{g}_{i+j}$ for every $i, j \geq 1$ [Wan23, Lemma 1.4.3]. The Lie algebra \mathfrak{g} is said to be **nilpotent** if $\mathfrak{g}_{i+1} = \{0\}$ for all sufficiently large i ; the minimal such i is called the **step** of \mathfrak{g} and is usually denoted s . Nilpotence of a Lie *group* is defined as for any other group (meaning that the lower central series terminates at the identity subgroup); a connected Lie group is nilpotent if and only if its corresponding Lie algebra is nilpotent.

Lie polynomials. A **Lie monomial** of degree d in the terms $X_1, X_2, \dots, X_n \in \mathfrak{g}$ is an expression obtained by taking iterated Lie brackets of these terms in some way, so that the sum over i of the total number of times X_i appears is d . In other words, a Lie monomial of degree 1 in the terms X_1, \dots, X_n is an expression of the form $L(X_1, \dots, X_n) = X_i$ for some $1 \leq i \leq n$, while each Lie monomial of degree d can be written $L(X_1, \dots, X_n) = [L_1(X_1, \dots, X_n), L_2(X_1, \dots, X_n)]$ for some Lie monomials L_1 and L_2 whose degrees sum to d . A **Lie polynomial** $P(X_1, X_2, \dots, X_n)$ in elements $X_1, X_2, \dots, X_n \in \mathfrak{g}$ is a linear combination of Lie monomials; it is said to be **homogeneous of degree d** if every Lie monomial in the linear combination has degree d . Thus, homogeneous Lie polynomials of degree d obey the scaling transformation $P(\lambda X_1, \dots, \lambda X_n) = \lambda^d P(X_1, \dots, X_n)$.

The exponential map. Given a Lie group G and associated Lie algebra \mathfrak{g} , there is a canonically defined **exponential map** $\exp : \mathfrak{g} \rightarrow G$, which for Lie subgroups of GL_n coincides with ordinary matrix exponentiation. (We will omit the general definition of the exponential map; everything we need to know about it will be captured by the Baker-Campbell-Hausdorff formula.) The exponential map is smooth, and is a diffeomorphism in a neighbourhood of the identity, but might not be injective or surjective. The **Baker-Campbell-Hausdorff (BCH) formula** [Tao14, Chapter 1.2.5] states that if G is a Lie group with Lie algebra \mathfrak{g} then there exists an open neighbourhood U of the origin such that

$$\log [\exp X \exp Y] = X + Y + \frac{1}{2}[X, Y] + \frac{1}{12}[X, [X, Y]] - \frac{1}{12}[Y, [X, Y]] + \dots = \sum_{i=1}^{\infty} L_i(X, Y)$$

for all $X, Y \in U$, where each L_i is a homogeneous Lie polynomial of degree i with rational coefficients and where we write \log for the inverse of the exponential map on $\exp(U)$. (The Lie polynomials L_i appearing in the BCH formula are universal and do not depend on the choice of Lie group G .) When G is a simply connected nilpotent Lie group, all terms with $i > s$ are identically zero, the exponential function $\exp : \mathfrak{g} \rightarrow G$ is defined globally, and the BCH formula holds for *all* $X, Y \in \mathfrak{g}$. This lets us define the **BCH product** on the Lie algebra \mathfrak{g} as

$$X \diamond Y := \log [e^X e^Y] = X + Y + \frac{1}{2}[X, Y] + \frac{1}{12}[X, [X, Y]] - \frac{1}{12}[Y, [X, Y]] + \dots,$$

so that (\mathfrak{g}, \diamond) is isomorphic to G as a Lie group. Note that $(aX) \diamond (bX) = (a+b)X$ for every $X \in \mathfrak{g}$ and $a, b \in \mathbb{R}$, so that 0 is both the additive and BCH identity and that $(-X)$ is both the additive and BCH inverse of X .

Example 4.2.1. The Heisenberg group H and its Lie algebra \mathfrak{h} are given by

$$H = \begin{pmatrix} 1 & \mathbb{R} & \mathbb{R} \\ 0 & 1 & \mathbb{R} \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad \mathfrak{h} = \begin{pmatrix} 0 & \mathbb{R} & \mathbb{R} \\ 0 & 0 & \mathbb{R} \\ 0 & 0 & 0 \end{pmatrix},$$

with exponential map and logarithm

$$\exp \begin{pmatrix} 0 & a & c \\ 0 & 0 & b \\ 0 & 0 & 0 \end{pmatrix} = \begin{pmatrix} 1 & a & c + \frac{ab}{2} \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad \log \begin{pmatrix} 1 & a & c \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 0 & a & c - \frac{ab}{2} \\ 0 & 0 & b \\ 0 & 0 & 0 \end{pmatrix}.$$

Note that the exponential map here is just the usual matrix exponential, which for $X \in \mathfrak{h}$ satisfies $e^X = I + X + \frac{1}{2}X^2$. The Lie bracket on \mathfrak{h} is given by

$$\left[\begin{pmatrix} 0 & a & c \\ 0 & 0 & b \\ 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & x & z \\ 0 & 0 & y \\ 0 & 0 & 0 \end{pmatrix} \right] = \begin{pmatrix} 0 & 0 & ay - xb \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

Since H is step-2 nilpotent, the BCH multiplication on \mathfrak{h} is given by

$$\begin{aligned} \begin{pmatrix} 0 & a & c \\ 0 & 0 & b \\ 0 & 0 & 0 \end{pmatrix} \diamond \begin{pmatrix} 0 & x & z \\ 0 & 0 & y \\ 0 & 0 & 0 \end{pmatrix} &= \begin{pmatrix} 0 & a & c \\ 0 & 0 & b \\ 0 & 0 & 0 \end{pmatrix} + \begin{pmatrix} 0 & x & z \\ 0 & 0 & y \\ 0 & 0 & 0 \end{pmatrix} + \frac{1}{2} \left[\begin{pmatrix} 0 & a & c \\ 0 & 0 & b \\ 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & x & z \\ 0 & 0 & y \\ 0 & 0 & 0 \end{pmatrix} \right] \\ &= \begin{pmatrix} 0 & a+x & c+z + \frac{ay-bx}{2} \\ 0 & 0 & b+y \\ 0 & 0 & 0 \end{pmatrix}. \end{aligned}$$

We will see in Example 4.4.1 that the non-integer rational coefficient 1/2 appearing in this expression leads to complications when comparing lattices in H and \mathfrak{h} .

Remark 4.2.1. It follows from the BCH formula that

$$\log[e^X e^Y e^{-X} e^{-Y}] = [X, Y] + \frac{1}{2}[X, [X, Y]] + \frac{1}{2}[Y, [X, Y]] + \dots$$

for X and Y in a neighbourhood of the origin in \mathfrak{g} (all of \mathfrak{g} when G is nilpotent and simply connected). As such, the BCH commutator and the Lie bracket agree to first order as $X, Y \rightarrow 0$, but are not exactly equal unless G is nilpotent of step at most 2.

We will also make use of the **Zassenhaus formula** [CMN12], a dual form of the BCH formula which states in particular that if G is s -step nilpotent then

$$\begin{aligned} e^{X+Y} &= e^X e^Y e^{-\frac{1}{2}[X,Y]} e^{\frac{1}{6}(2[Y,[X,Y]]+[X,[X,Y]])} e^{\frac{-1}{24}([[[X,Y],X],X]+3[[[X,Y],X],Y]+3[[[X,Y],Y],Y])} \dots \\ &= e^X e^Y e^{\tilde{L}_2(X,Y)} e^{\tilde{L}_3(X,Y)} \dots e^{\tilde{L}_s(X,Y)} \end{aligned} \quad (4.2.1)$$

for every $X, Y \in \mathfrak{g}$, where \tilde{L}_i is a homogeneous Lie polynomial of degree i with rational coefficients for each $i \geq 2$.

Subgroups and lattices. Let G be a simply connected nilpotent Lie group with Lie algebra \mathfrak{g} . For each set $A \subseteq \mathfrak{g}$, we define $\mathcal{L}(A)$ to be the smallest Lie subalgebra of \mathfrak{g} containing A . The exponential map identifies closed connected subgroups of G with Lie subalgebras of \mathfrak{g} , so that every closed connected subgroup of G is itself a simply connected nilpotent Lie group. (For general Lie groups, the image under the exponential map of a Lie subalgebra might not be closed, but for simply connected nilpotent groups it is always closed since the exponential map is a diffeomorphism.) As such, for each subset A of G , the intersection of all closed connected subgroups of G containing A is a closed connected subgroup of G that is equal to $\exp(\mathcal{L}(\log A))$. We write $\mathcal{C}(A)$ for this minimal closed connected subgroup of G containing A .

Theorem 4.2.2 (Mal'cev). *If G is a simply connected nilpotent Lie group and H is a closed subgroup of G then G/H is compact if and only if H is not contained in any proper closed connected subgroup of G . In particular, $\mathcal{C}(H)/H$ is compact.*

(The quotients G/H and $\mathcal{C}(H)/H$ appearing here are topological spaces, and do not carry group structures in general.) We call a subgroup Γ of a simply connected nilpotent Lie group G a **lattice** in G if it is discrete with compact quotient G/Γ . (Note that discrete subgroups of Hausdorff topological groups are automatically closed.)

Remark 4.2.2. For a general Lie group, a lattice is defined to be a discrete subgroup for which the quotient admits a finite left-invariant measure (a.k.a. Haar measure); for simply connected nilpotent Lie groups this is equivalent to G/Γ being compact by Mal'cev's theorem. This theorem also gives several further characterisations of a discrete subgroup being a lattice that we omit since we do not use them.

A further theorem of Mal'cev [Rag72, Theorem 2.12] states that a simply connected nilpotent Lie group G admits a lattice if and only if its Lie algebra \mathfrak{g} admits a basis e_1, \dots, e_d for which the **structure constants** $(T_{i,j}^k)_{i,j,k}$, defined by $[e_i, e_j] = \sum_k T_{i,j}^k e_k$, are rational. It is a consequence of

this theorem [Wan23, Proposition 2.3.7] that if Γ is a lattice in a simply connected nilpotent Lie group then there exist (additive) lattices Λ^- and Λ^+ in \mathfrak{g} such that $\Lambda^- \subseteq \log \Gamma \subseteq \Lambda^+$. We will prove significantly stronger versions of this fact in Section 4.4.

Carnot groups and Pansu's theorem

A **Carnot group** is a simply connected nilpotent Lie group G of some step s whose Lie algebra \mathfrak{g} is equipped with a decomposition

$$\mathfrak{g} = V_1 \oplus V_2 \cdots \oplus V_s$$

for some non-trivial linear subspaces V_1, \dots, V_s such that $[V_1, V_i] = V_{i+1}$ for every $1 \leq i < s$ and $[V_1, V_s] = \{0\}$. It can be shown that for such a decomposition we have moreover that $[V_i, V_j] \subseteq V_{i+j}$ for every $i, j \geq 1$, where $V_{i+j} := \{0\}$ for $i+j > s$. A decomposition $\mathfrak{g} = V_1 \oplus \cdots \oplus V_s$ satisfying these conditions is known as a **stratification** of \mathfrak{g} ; the subspace V_1 , which generates \mathfrak{g} as a Lie algebra, is known as the **horizontal subspace**. Not every nilpotent lie algebra admits a stratification. While a Carnot group consists of both a nilpotent Lie group and a choice of stratification of its Lie algebra, we will nevertheless write e.g. “let G be a Carnot group” when this does not cause confusion.

Given a Carnot group G and a real number $\lambda > 0$ the **dilation maps** $\delta_\lambda : \mathfrak{g} \rightarrow \mathfrak{g}$ and $D_\lambda : G \rightarrow G$ are defined by

$$\delta_\lambda(x_1 + x_2 + \cdots + x_s) = \lambda x_1 + \lambda^2 x_2 + \cdots + \lambda^s x_s \quad \text{and} \quad D_\lambda(x) = \exp(\delta_\lambda(\log(x))),$$

where we write $x = x_1 + x_2 + \cdots + x_s$ for the decomposition $x \in \mathfrak{g}$ associated to the stratification $\mathfrak{g} = V_1 \oplus V_2 \oplus \cdots \oplus V_s$. These dilation maps satisfy the semigroup properties $\delta_{\lambda\mu} = \delta_\lambda \delta_\mu = \delta_\mu \delta_\lambda$ and $D_{\lambda\mu} = D_\lambda D_\mu = D_\mu D_\lambda$ for every $\lambda, \mu > 0$. It is a consequence of the Baker-Campbell-Hausdorff formula and the fact that $[V_i, V_j] \subseteq V_{i+j}$ that D_λ is a Lie group automorphism of G for every $\lambda > 0$. A metric $d : G \times G \rightarrow [0, \infty)$ on a Carnot group is said to be **left-invariant and homogeneous** if

$$d(zx, zy) = d(x, y) \quad \text{and} \quad d(D_\lambda x, D_\lambda y) = \lambda d(x, y)$$

for every $x, y, z \in G$ and $\lambda > 0$. Note that the abelian group \mathbb{R}^d is a Carnot group with $V_1 = \mathbb{R}^d$, and the left-invariant homogeneous metrics on \mathbb{R}^d seen as a Carnot group are equivalent to norms on \mathbb{R}^d . In general, left-invariant homogeneous metrics are closely analogous to norms, and also have the property that they are determined by the unit ball B around the origin:

$$d(x, y) = \inf\{\lambda : x^{-1}y \in D_\lambda(B)\}.$$

In particular, the ball $\{x : d(0, x) \leq \lambda\}$ is equal to $D_\lambda(B)$.

We now introduce the free step- s nilpotent Lie algebra and the free step- s Carnot group, referring the reader to [BLU07, Chapter 14.1] for proofs the objects we discuss here are well-defined. Let S be a finite set. The **free step- s nilpotent Lie algebra** $\mathfrak{f}_{s,S}$ is defined to be the unique-up-to- S -preserving-isomorphism nilpotent Lie algebra of step s that is generated by S and is such that if \mathfrak{g} is any nilpotent Lie algebra of step at most s and $\phi : S \rightarrow \mathfrak{g}$ is any function, then there exists a unique Lie algebra homomorphism $\mathfrak{f}_{s,S} \rightarrow \mathfrak{g}$ extending ϕ . The Lie algebra $\mathfrak{f}_{s,S}$ may be equipped with a canonical stratification defined in terms of *Hall bases*, making its associated BCH-multiplication Lie group $G_{s,S} = (\mathfrak{f}_{s,S}, \diamond)$ into a Carnot group known as the **free step- s Carnot group** or **free step- s nilpotent Lie group** over S ; see [BLU07, Chapters 14.1 and 14.2] for details. (It is convenient to define the free step- s nilpotent Lie group over S via BCH multiplication so that it contains the set S .) Moreover, this stratification $\mathfrak{f}_{s,S} = V_1 \oplus \cdots \oplus V_s$ has the property that V_1 is equal to the linear span of S . Finally, if we define $\Gamma_{s,S}$ to be the subgroup of $G_{s,S}$ generated by S , then $\Gamma_{s,S}$ is a lattice in $G_{s,S}$ that is isomorphic (via an S -preserving isomorphism) to the discrete free step- s nilpotent group $N_{s,S}$. Indeed, the fact that $\Gamma_{s,S}$ is discrete in $G_{s,S}$ can be proven using Mal'cev's theorem on rational structure constants, since the structure constants in the Hall basis are all equal to 1, while Γ is a lattice since S generates $G_{s,S}$ as a Lie group. (The fact that $\Gamma_{s,S}$ is discrete can also be proven using the techniques of Section 4.4.) Finally, the fact that $\Gamma_{s,S}$ is isomorphic to $N_{s,S}$ can be deduced straightforwardly from the relevant universal properties since every torsion-free finitely generated nilpotent group can be embedded as a lattice in a nilpotent Lie group, sometimes known as the **Mal'cev completion** of the group [Rag72, Theorem 2.18].

In his thesis [Pan83], Pansu proved that Cayley graphs of finitely generated nilpotent groups converge under rescaling to Carnot groups equipped with certain left-invariant homogeneous metrics known as *sub-Finsler metrics*. (We will not need the definition of sub-Finsler metrics in this paper.) Note that the Carnot group arising in this limit might *not* be isomorphic to the Mal'cev completion of the relevant nilpotent group, and indeed the Mal'cev completion might not admit a stratification. However, if N is a torsion-free nilpotent group with finite generating set S and the lie algebra \mathfrak{g} of the Mal'cev completion G happens to be simply connected and admit a stratification $\mathfrak{g} = V_1 \oplus \cdots \oplus V_s$ with $\log S \subseteq V_1$, then there exists a left-invariant homogeneous metric d_G on G such that

$$d_S(x, y) = (1 \pm o(1))d_G(x, y) \quad \text{as } d_S(x, y) \rightarrow \infty, \quad (4.2.2)$$

where d_S denotes the word metric on N . It follows in particular that $(N, \frac{1}{n}d_S)$ converges to (G, d_G) in the Gromov-Hausdorff sense as $n \rightarrow \infty$. See [BL13; Tas22] for details and quantitative refinements of this theorem.

4.3 Exploring abelian lattices with convex sets

In this section we prove the following theorem, which will eventually be used to prove our main theorems by reduction to the abelian case.

Theorem 4.3.1. *Let $d \geq 1$, let Λ be a discrete subgroup of \mathbb{R}^d , and let $K_1 \subseteq K_2 \subseteq \dots$ be an increasing sequence of non-empty, symmetric, convex sets in \mathbb{R}^d . For each $n \geq 1$ let $\Lambda_n = \text{span}_{\mathbb{Z}}(\Lambda \cap K_n)$ so that $\Lambda_1 \subseteq \Lambda_2 \subseteq \dots$ is an increasing sequence of subgroups of Λ . Then*

$$\#\{n : \Lambda_{n+1} \neq \Lambda_n\} \leq d + 1 + \sum_{\ell=1}^d \lfloor \log_2 \ell! \rfloor.$$

Remark 4.3.1. By Stirling's formula, the upper bound appearing here is asymptotic to $\frac{1}{2}d^2 \log_2 d$ as $d \rightarrow \infty$. We have not investigated the optimality of this bound.

We will deduce this theorem as a consequence of Minkowski's second theorem [Cas97, p. 203], which states that if Λ is a lattice in \mathbb{R}^d with $d \geq 1$ and K is a non-empty, symmetric convex subset of \mathbb{R}^d then

$$1 \leq \frac{\text{vol}(\mathbb{R}^d/\Lambda)}{2^d \text{vol}(K) \prod_{i=1}^d \lambda_i(\Lambda, K)} \leq d!, \quad (4.3.1)$$

where

$$\lambda_i(\Lambda, K) = \inf \{ \lambda > 0 : \lambda K \cap \Lambda \text{ contains at least } i \text{ linearly independent vectors} \}$$

for each $1 \leq i \leq d$. For our purposes, the most important feature of (4.3.1) is that the expression $2^d \text{vol}(K) \prod_{i=1}^d \lambda_i(\Lambda, K)$ is determined by K and $\Lambda \cap K$ whenever $\Lambda \cap K$ has real span equal to \mathbb{R}^d . Indeed, if Λ is a lattice in \mathbb{R}^d then every fundamental domain for Λ has volume $\text{vol}(\mathbb{R}^d/\Lambda)$ and if $\Lambda_1 \subseteq \Lambda_2$ are two lattices in \mathbb{R}^d then Λ_1 is a finite-index subgroup of Λ_2 with index

$$[\Lambda_2 : \Lambda_1] = \frac{\text{vol}(\mathbb{R}^d/\Lambda_1)}{\text{vol}(\mathbb{R}^d/\Lambda_2)}.$$

Thus, if $\Lambda_1 \subseteq \Lambda_2$ are two lattices and K is a symmetric convex set such that $\Lambda_1 \cap K$ and $\Lambda_2 \cap K$ are equal and both have real linear span equal to \mathbb{R}^d then Minkowski's second theorem implies that $[\Lambda_2 : \Lambda_1] \leq d!$.

Proof of Theorem 4.3.1. It suffices to prove that if $K_1 \subseteq K_2 \subseteq \dots \subseteq K_n$ is an increasing sequence of non-empty, symmetric convex subsets of \mathbb{R}^d and Λ is a discrete subgroup of \mathbb{R}^d such that the subgroups $\Lambda_i := \text{span}_{\mathbb{Z}}(\Lambda \cap K_i)$ satisfy $\Lambda_{i+1} \neq \Lambda_i$ for every $1 \leq i < n$ then $n \leq d + 1 + \sum_{\ell=0}^d \lfloor \log_2 \ell! \rfloor$. We may also assume without loss of generality that $\Lambda = \Lambda_n$ is a lattice in \mathbb{R}^d , replacing Λ with Λ_n

and \mathbb{R}^d with the subspace spanned by Λ_n otherwise. Fix such a pair Λ and (K_1, \dots, K_n) and for each $0 \leq \ell \leq d$ let i_ℓ be minimal such that the real span of Λ_{i_ℓ} has dimension at least ℓ , setting $i_{d+1} = n+1$ for notational convenience. Since $\Lambda_2 \neq \Lambda_1$, we must have that $i_0 = 1$ and $i_1 \in \{1, 2\}$. For each $1 \leq \ell \leq d$ define V_ℓ to be the real span of Λ_{i_ℓ} , so that Λ_j is a lattice in V_ℓ for every $1 \leq \ell \leq d$ and $i_\ell \leq j < i_{\ell+1}$. (Note that V_ℓ might have dimension strictly larger than ℓ , in which case $i_{\ell+1} = i_\ell$.) Suppose that $0 \leq \ell \leq d$ is such that $i_{\ell+1} \geq i_\ell + 2$. In this case, the subgroups Λ_{i_ℓ} and $\Lambda_{i_{\ell+1}-1}$ are both lattices in V_ℓ with $\Lambda_{i_\ell} \cap K_{i_\ell} = \Lambda_{i_{\ell+1}-1} \cap K_{i_\ell}$ and with $\Lambda_{i_\ell} \cap K_{i_\ell}$ having real span equal to V_ℓ . As such, it follows by Minkowski's second theorem that

$$[\Lambda_{i_{\ell+1}-1} : \Lambda_{i_\ell}] = \frac{\text{vol}(V_\ell / \Lambda_{i_\ell})}{\text{vol}(V_\ell / \Lambda_{i_{\ell+1}-1})} \leq \ell!$$

for each $1 \leq \ell \leq d$ such that $i_{\ell+1} \geq i_\ell + 2$. Now, using that

$$[\Lambda_{i_{\ell+1}-1} : \Lambda_{i_\ell}] = \prod_{j=i_\ell}^{i_{\ell+1}-2} [\Lambda_{j+1} : \Lambda_j] \geq 2^{i_{\ell+1}-1-i_\ell}$$

it follows that $i_{\ell+1} - i_\ell \leq 1 + \lfloor \log_2 \ell! \rfloor$ for every $1 \leq \ell \leq d$ such that $i_{\ell+1} \geq i_\ell + 2$. Since the same inequality also holds trivially when $i_{\ell+1} < i_\ell + 2$, it follows that

$$n = i_{d+1} - i_0 = \sum_{\ell=0}^d (i_{\ell+1} - i_\ell) \leq d + 1 + \sum_{\ell=1}^d \lfloor \log_2 \ell! \rfloor$$

as claimed. \square

4.4 Additive and multiplicative subgroups of nilpotent Lie algebras

Our goal in this section is to clarify the relationship between lattices in a simply connected nilpotent Lie group and its associated Lie algebra. As mentioned above, closed, connected subgroups of a simply connected nilpotent Lie group G are in bijection with Lie subalgebras of the Lie algebra \mathfrak{g} , which are precisely the closed, connected, bracket-closed additive subgroups of \mathfrak{g} . Without the assumption of connectivity, the exponential map need not interact this nicely with the subgroup structure of G : it is possible to have subgroups H of G for which $\log H$ is not an additive subgroup of \mathfrak{g} and to have bracket-closed additive subgroups of \mathfrak{g} whose image under the exponential is not a subgroup of G .

Example 4.4.1. Let the Heisenberg group H and its Lie algebra \mathfrak{h} be as in Example 4.2.1. The set of all elements of H whose matrix entries are integers is a lattice in H whose logarithm is given by

$$\log \begin{pmatrix} 1 & \mathbb{Z} & \mathbb{Z} \\ 0 & 1 & \mathbb{Z} \\ 0 & 0 & 1 \end{pmatrix} = \left\{ \begin{pmatrix} 0 & a & c \\ 0 & 0 & b \\ 0 & 0 & 0 \end{pmatrix} : a, b \in \mathbb{Z}, c \in \frac{1}{2}\mathbb{Z}, 2c = ab \pmod{2} \right\}.$$

This is *not* an additive subgroup of \mathfrak{h} since

$$\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} + \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix} \notin \log \begin{pmatrix} 1 & \mathbb{Z} & \mathbb{Z} \\ 0 & 1 & \mathbb{Z} \\ 0 & 0 & 1 \end{pmatrix}.$$

Similarly, the set of all elements of \mathfrak{h} whose matrix entries are integers is a bracket-closed additive subgroup of \mathfrak{h} whose exponential is not a subgroup of H .

We call a subgroup H of G **harmonious** if $\log H$ is an additive subgroup of \mathfrak{g} that is bracket-closed in the sense that $[\log x, \log y] \in \log H$ for every $x, y \in H$. (This terminology is not standard.) Thus, harmonious subgroups of G are those that most closely mimic the behaviour of the connected subgroups of G under the exponential map. We call a lattice in G that is also a harmonious subgroup of G a **harmonious lattice** in G , noting that if Γ is a harmonious lattice in G then $\Lambda = \log \Gamma$ is a bracket-closed lattice in \mathfrak{g} .

The remainder of this section is devoted to proving the following theorem, which states intuitively that subgroups of G and bracket-closed additive subgroups of \mathfrak{g} are, in some sense, “equivalent up to bounded index”. This result allows us to deduce various statements about lattices in simply connected nilpotent Lie groups from analogous statements for vector spaces (i.e., statements in the geometry of numbers), which are classical. Note that the theorem does not require Γ to be discrete. Recall that if $A \subseteq \mathfrak{g}$ and $\lambda \in \mathbb{R}$ then we define $\lambda \cdot A = \{\lambda a : a \in A\}$, so that if Λ is an additive subgroup of \mathfrak{g} then $(m\lambda) \cdot \Lambda \subseteq \lambda \cdot \Lambda$ for every $\lambda \in \mathbb{R}$ and $m \in \mathbb{Z}$. We also write $\mathcal{B}(A)$ for the smallest bracket-closed set containing A .

Theorem 4.4.2. *Let G be a simply connected nilpotent Lie group of step s with Lie algebra \mathfrak{g} . There exist positive integers C_1 and C_2 depending only on s such that if Γ is a subgroup of G then the sets*

$$\mathcal{H}_-(\Gamma) := \exp\left(C_1 \cdot \text{span}_{\mathbb{Z}}(\log \Gamma)\right) \quad \text{and} \quad \mathcal{H}_+(\Gamma) := \exp\left(C_1 \cdot \mathcal{B}\left(\frac{1}{C_1} \cdot \text{span}_{\mathbb{Z}}(\log \Gamma)\right)\right)$$

are harmonious subgroups of G such that

$$C_1 \cdot \log \Gamma \subseteq \log \mathcal{H}_-(\Gamma) \subseteq \log \Gamma \subseteq \log \mathcal{H}_+(\Gamma) \subseteq \frac{1}{C_2} \cdot \log \Gamma.$$

In corollary 4.4.13 we show moreover that if Γ is discrete then the harmonious subgroups $\mathcal{H}_-(\Gamma)$ and $\mathcal{H}_+(\Gamma)$ have index bounded by a constant depending only on the step and dimension of G .

Example 4.4.3. We continue to analyze the Heisenberg group as studied in Examples 4.2.1 and 4.4.1. Although the subgroup Γ of H consisting of those elements of H with integer matrix entries is not harmonious, we can write

$$\begin{pmatrix} 1 & 2\mathbb{Z} & \mathbb{Z} \\ 0 & 1 & 2\mathbb{Z} \\ 0 & 0 & 1 \end{pmatrix} \subseteq \Gamma = \begin{pmatrix} 1 & \mathbb{Z} & \mathbb{Z} \\ 0 & 1 & \mathbb{Z} \\ 0 & 0 & 1 \end{pmatrix} \subseteq \begin{pmatrix} 1 & \mathbb{Z} & \frac{1}{2}\mathbb{Z} \\ 0 & 1 & \mathbb{Z} \\ 0 & 0 & 1 \end{pmatrix}$$

with the two outer two lattices being harmonious subgroups of H . Moreover, these subgroups arise naturally from the original subgroup Γ as

$$\begin{pmatrix} 1 & 2\mathbb{Z} & \mathbb{Z} \\ 0 & 1 & 2\mathbb{Z} \\ 0 & 0 & 1 \end{pmatrix} = \exp(\text{span}_{\mathbb{Z}}(2 \cdot \log \Gamma)) \text{ and } \begin{pmatrix} 1 & \mathbb{Z} & \frac{1}{2}\mathbb{Z} \\ 0 & 1 & \mathbb{Z} \\ 0 & 0 & 1 \end{pmatrix} = \exp(\text{span}_{\mathbb{Z}}(\log \Gamma)).$$

Similar remarks apply to the bracket-closed additive lattice in \mathfrak{h} consisting of those elements of \mathfrak{h} with integer matrix entries.

The methods used to prove Theorem 4.4.2 are based on those of Breuillard and Green [BG11]. In particular, we will make use of the following lemma of Tessera and Tointon [TT18] that is also proved using the methods of [BG11].

Lemma 4.4.4. Let G be a simply connected nilpotent Lie group of step s with Lie algebra \mathfrak{g} . There exists an integer constant $C = C(s)$ such that if Λ is a bracket-closed additive subgroup of the Lie algebra \mathfrak{g} then $\exp(C \cdot \Lambda)$ is a harmonious subgroup of G .

Proof. This is essentially [TT18, Lemma 4.3]. It is not stated that $C \cdot \Lambda$ is bracket-closed, but this is obvious since $[C \cdot \Lambda, C \cdot \Lambda] = C^2 \cdot [\Lambda, \Lambda] \subseteq C \cdot \Lambda$. \square

Remark 4.4.1. The constant C_1 appearing in Theorem 4.4.2 will be taken to be a multiple of the constant C appearing in lemma 4.4.4. This will be important in the proof of proposition 4.5.1.

We begin by stating the following lemma of Lazard [Laz54] as presented in [BG11, Lemmas 5.2 and 5.3]. This lemma was first applied to the structure theory of approximate groups in the work of Fisher, Katz, and Peng [FKP09]. Given a simply connected nilpotent Lie group G , we define $x^\alpha = \exp(\alpha \log x)$ for every $x \in G$ and $\alpha \in \mathbb{R}$.

Lemma 4.4.5 (Lazard). Let G be a simply connected nilpotent Lie group of step s with Lie algebra \mathfrak{g} . There exists $\ell \geq 1$ and sequences of rational numbers $\alpha_1, \dots, \alpha_\ell, \beta_1, \dots, \beta_\ell, \gamma_1, \dots, \gamma_\ell$, and

$\delta_1, \dots, \delta_\ell$ depending only on s such that

$$\exp(\log x + \log y) = x^{\alpha_1} y^{\beta_2} \cdots x^{\alpha_\ell} y^{\beta_\ell}$$

and

$$\exp([\log x, \log y]) = x^{\gamma_1} y^{\delta_1} \cdots x^{\gamma_\ell} y^{\delta_\ell}$$

for every $x, y \in G$.

Corollary 4.4.6 (Expansion of sums). Let G be a simply connected nilpotent Lie group of step s with Lie algebra \mathfrak{g} . For each $n \geq 2$ there exists $\ell \geq 1$, a sequence of rational numbers $\alpha_1, \dots, \alpha_\ell$, and a sequence of indices $i_1, \dots, i_\ell \in \{1, \dots, n\}$, all depending only on s and n , such that

$$\exp(\log x_1 + \log x_2 + \cdots + \log x_n) = x_{i_1}^{\alpha_1} x_{i_2}^{\alpha_2} \cdots x_{i_\ell}^{\alpha_\ell}$$

for every $x_1, \dots, x_n \in G$.

Corollary 4.4.7. Let G be a simply connected nilpotent Lie group of step s with Lie algebra \mathfrak{g} , let Γ be a subgroup of G and let $\Lambda = \log \Gamma$. For each $n \geq 1$, there exists a natural number $C = C(s, n)$ such that if $X_1, \dots, X_n \in C \cdot \Lambda$ then $X_1 + \cdots + X_n \in \Lambda$.

Proof. Let $C = C(s, n)$ be the least common multiple of the denominators of the numbers $\alpha_1, \dots, \alpha_\ell$ appearing in corollary 4.4.6 when written in reduced form. \square

Lemma 4.4.8. Let G be a simply connected nilpotent Lie group of step s with Lie algebra \mathfrak{g} , let Γ be a subgroup of G and let $\Lambda = \log \Gamma$. There exists a natural number $C = C(s)$ such that

$$\left\{ L(X_1, \dots, X_n) : X_1, \dots, X_n \in C^{n-1} \cdot \Lambda \right\} \subseteq \Lambda,$$

for every multilinear Lie monomial L .

(A Lie monomial $L(X_1, \dots, X_n)$ is multilinear when each variable appears at most once.)

Proof. Let $C = C(s)$ be the least common multiple of the denominators of the numbers $\gamma_1, \dots, \gamma_\ell$ and $\delta_1, \dots, \delta_\ell$ appearing in Lazard's lemma when written in reduced form. We will prove the claim by induction on n , the case $n = 1$ being vacuous. If $n > 1$ then

$$L(X_1, \dots, X_n) = [L_1(X_{\pi(1)}, \dots, X_{\pi(j)}), L_2(X_{\pi(j+1)}, \dots, X_{\pi(n)})]$$

for some $1 \leq j < n$, some multilinear Lie monomials L_1, L_2 and some permutation $\pi : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$. We may assume without loss of generality that π is the identity permutation. By the induction hypothesis, if $X_1, \dots, X_n \in C^{n-2} \cdot \Lambda$ then $L_1(X_1, \dots, X_j), L_2(X_{j+1}, \dots, X_n) \in \Lambda$. As such, if $X_1, \dots, X_n \in C^{n-1} \cdot \Lambda$ then we have by multilineararity that

$$L_1(X_1, \dots, X_j) = C^j L_1(C^{-1} X_1, \dots, C^{-1} X_j) \in C^j \cdot \Lambda \subseteq C \cdot \Lambda$$

and

$$L_2(X_{j+1}, \dots, X_n) = C^{n-j} L_2(C^{-1} X_{j+1}, \dots, C^{-1} X_n) \in C^{n-j} \cdot \Lambda \subseteq C \cdot \Lambda.$$

On the other hand, if $X, Y \in C \cdot \Lambda$ then $e^{\gamma_j X}, e^{\delta_j Y} \in \Gamma$ for every $1 \leq j \leq \ell$, so that

$$\exp([X, Y]) = e^{\gamma_1 X} e^{\delta_1 Y} \cdots e^{\gamma_\ell X} e^{\delta_\ell Y} \in \Gamma$$

and hence that $[X, Y] \in \Lambda$. It follows that if $X_1, \dots, X_n \in C^{n-1} \cdot \Lambda$ then $L(X_1, \dots, X_n) \in \Lambda$ as claimed. \square

Corollary 4.4.9. Let G be a simply connected nilpotent Lie group of step s with Lie algebra \mathfrak{g} , let Γ be a subgroup of G and let $\Lambda = \log \Gamma$. There exists a natural number $C = C(s)$ such that

$$\left\{ L(X_1, \dots, X_n) : X_1, \dots, X_n \in \Lambda, \quad X_i \in C^{d(d-1)} \cdot \Lambda \text{ for some } 1 \leq i \leq n \right\} \subseteq \Lambda,$$

for every Lie monomial L of degree d that depends on every variable.

Proof. It suffices without loss of generality to consider the case that L is multilinear. Let C be the constant from lemma 4.4.8. Let $X_1, \dots, X_n \in \Lambda$, and let $L(X_1, \dots, X_n)$ be a multilinear Lie monomial depending on every variable. Such a Lie monomial necessarily has degree $d = n$. For each $1 \leq i \leq n$ we can write

$$L(X_1, \dots, X_n) = L(C^{(n-1)} X_1, \dots, C^{(n-1)} X_{i-1}, C^{-(n-1)^2} X_i, C^{(n-1)} X_{i+1}, \dots, C^{(n-1)} X_n).$$

If $X_i \in C^{n(n-1)} \cdot \Lambda = C^{(n-1)^2 + (n-1)} \cdot \Lambda$ then $C^{-(n-1)^2} X_i \in C^{(n-1)} \cdot \Lambda$, so that the claim follows from lemma 4.4.8. \square

Lemma 4.4.10. Let G be a simply connected nilpotent Lie group of step s with Lie algebra \mathfrak{g} . There exists a constant C depending only on s such that if Γ is a subgroup of G and we write $\Lambda = \log \Gamma$ then $C \cdot \Lambda$ is bracket-closed and $X + Y \in \Lambda$ for every $X \in C \cdot \Lambda$ and $Y \in \Lambda$.

Proof. We recall the Zassenhaus formula, which states for simply connected nilpotent Lie groups that

$$\exp(X + Y) = e^X e^Y e^{L_2(X, Y)} e^{L_3(X, Y)} \dots e^{L_s(X, Y)}$$

for every $X, Y \in \mathfrak{g}$, where, for each $i \geq 1$, L_i is a homogeneous Lie polynomial of degree i of the form

$$L_i(X, Y) = \sum_{j=1}^{r_i} a_{i,j} L_{i,j}(X, Y),$$

where $a_{i,j}$ are rational numbers and $L_{i,j}$ are Lie monomials depending on both variables. Let C_1 be the least common multiple of the denominators of the rational numbers $a_{i,j}$ and let C_2 be the least common multiple of the constants $C(s, 2), C(s, 3), \dots, C(2, \max_i r_i)$ appearing in corollary 4.4.7. By corollary 4.4.9, there exists a constant $C_3 = C_3(s)$ such that if $X \in C_3 \cdot \Lambda$ and $Y \in \Lambda$ then $L_{i,j}(X, Y) \in \Lambda$ for every $1 \leq i \leq s$ and $1 \leq j \leq r_i$. Let $C = C_1 C_2 C_3$. Since $L_{i,j}$ depends on X , it follows that if $X \in C \cdot \Lambda$ then $L_{i,j}(X, Y) = (C_1 C_2)^{d_{i,j}} L_{i,j}((C_1 C_2)^{-1} X, Y) \in (C_1 C_2) \cdot \Lambda$ for every $y \in \Lambda$, where $d_{i,j}$ is the degree of X in $L_{i,j}$. It follows in particular that if $X \in C \cdot \Lambda$ then $a_{i,j} L_{i,j}(X, Y) \in C_2 \cdot \Lambda$ for every $1 \leq i \leq s$ and $1 \leq j \leq r_i$ and hence by corollary 4.4.7 that $L_i(X, Y) \in \Lambda$ for every $X \in C \cdot \Lambda$ and $Y \in \Lambda$. Since Γ is a subgroup of G , it follows by the Zassenhaus formula that $X + Y \in \Lambda$ for every $X \in C \cdot \Lambda$ and $Y \in \Lambda$ as claimed. Moreover, if $X, Y \in C \cdot \Lambda$ then $[X, Y] = C[X, C^{-1}Y] \in C \cdot \Lambda$ since $[X', Y'] \in \Lambda$ for every $X' \in C \cdot \Lambda$ and $Y' \in \Lambda$, so that $C \cdot \Lambda$ is bracket-closed as claimed. \square

We are now ready to prove Theorem 4.4.2.

Proof of Theorem 4.4.2. We begin by proving the claim concerning $\mathcal{H}_-(\Gamma)$. Let $C_{-1} = C_{-1}(s)$ be the constant from lemma 4.4.10, let $C_0 = C_0(s)$ be the constant from lemma 4.4.4, and let $C_1 = C_{-1} C_0$. Let G be a simply connected nilpotent Lie group of step s , let \mathfrak{g} be the Lie algebra of G , let Γ be a subgroup of G and let $\Lambda = \log \Gamma$. lemma 4.4.10 implies that $C_1 \cdot \Lambda$ is bracket-closed and that $C_{-1} \cdot \Lambda + \Lambda \subseteq \Lambda$, and it follows by induction on the number of terms in a linear combination that $\text{span}_{\mathbb{Z}}(C_{-1} \cdot \Lambda) = C_{-1} \cdot \text{span}_{\mathbb{Z}}(\Lambda)$ is contained in Λ . Since the \mathbb{Z} -span of a bracket-closed set is bracket-closed, it follows that $\text{span}_{\mathbb{Z}}(C_{-1} \cdot \Lambda)$ is a bracket-closed additive subgroup of \mathfrak{g} and hence by lemma 4.4.4 that $\text{span}_{\mathbb{Z}}(C_1 \cdot \Lambda) = C_0 \cdot \text{span}_{\mathbb{Z}}(C_{-1} \cdot \Lambda)$ is a bracket-closed additive subgroup of \mathfrak{g} whose exponential $\mathcal{H}_-(\Gamma)$ is a harmonious subgroup of G satisfying the required set inclusion $C_1 \cdot \log \Gamma \subseteq \log \mathcal{H}_-(\Gamma) \subseteq \log \Gamma$.

Now consider the set $\mathcal{H}_+(\Gamma)$ defined by

$$\mathcal{H}_+(\Gamma) := \exp \left(C_1 \cdot \mathcal{B} \left(\frac{1}{C_1} \cdot \text{span}_{\mathbb{Z}}(\Lambda) \right) \right) = \exp \left(C_1 \cdot \text{span}_{\mathbb{Z}} \mathcal{B} \left(\frac{1}{C_1} \cdot \Lambda \right) \right).$$

Since $\text{span}_{\mathbb{Z}} \mathcal{B} \left(\frac{1}{C_1} \cdot \Lambda \right)$ is a bracket-closed additive subgroup of \mathfrak{g} and C_0 divides C_1 , lemma 4.4.4 implies that $\mathcal{H}_+(\Gamma)$ is a harmonious subgroup of G . Moreover, $\log \mathcal{H}_+(\Gamma)$ trivially contains $\text{span}_{\mathbb{Z}} \Lambda$. As such, it remains only to prove that there exists a constant $C_2 = C_2(s)$ such that $\log \mathcal{H}_+(\Gamma) \subseteq \frac{1}{C_2} \cdot \Lambda$. Since $C_1 \cdot \text{span}_{\mathbb{Z}}(\Lambda)$ is bracket-closed and $[\mathfrak{g}_i, \mathfrak{g}_j] \subseteq \mathfrak{g}_{i+j}$ we have that

$$\left[\left(\frac{1}{n} \cdot \text{span}_{\mathbb{Z}}(\Lambda) \right) \cap \mathfrak{g}_i, \left(\frac{1}{m} \cdot \text{span}_{\mathbb{Z}}(\Lambda) \right) \cap \mathfrak{g}_j \right] \subseteq \left(\frac{1}{C_1^2 nm} \Lambda \right) \cap \mathfrak{g}_{i+j}$$

for every pair of integers $n, m \geq 0$ and $1 \leq i, j \leq s$. Thus, it follows by induction on i that

$$\mathcal{B} \left(\frac{1}{C_1} \cdot \text{span}_{\mathbb{Z}}(\Lambda) \right) \cap \mathfrak{g}_i \subseteq \frac{1}{C_1^{3i-2}} \cdot (\Lambda \cap \mathfrak{g}_i)$$

for every $i \geq 1$ and hence that

$$\mathcal{B} \left(\frac{1}{C_1} \cdot \text{span}_{\mathbb{Z}}(\Lambda) \right) \subseteq \frac{1}{C_1^{3s-2}} \cdot \Lambda.$$

This implies that the claim holds with $C_2 = C_1^{3(s-1)}$. □

Comparing additive and multiplicative indices

In this section we prove bounds on the index $[\mathcal{H}_+(\Gamma) : \mathcal{H}_-(\Gamma)]$. We will deduce these bounds from the following proposition, which lets us compare additive and multiplicative indices in the Lie algebra of a simply connected nilpotent Lie group.

Proposition 4.4.11 (Index sandwich). Let $\Gamma_1 \subseteq \Gamma_2$ be lattices in a simply connected nilpotent Lie group G with Lie algebra \mathfrak{g} , and suppose that $\Lambda_1 \subseteq \Lambda_2$ are additive lattices in \mathfrak{g} .

1. If $\Lambda_1 \subseteq \log \Gamma_1 \subseteq \log \Gamma_2 \subseteq \Lambda_2$ then $[\Gamma_2 : \Gamma_1] \leq [\Lambda_2 : \Lambda_1]$.
2. If $\log \Gamma_1 \subseteq \Lambda_1 \subseteq \Lambda_2 \subseteq \log \Gamma_2$ then $[\Gamma_2 : \Gamma_1] \geq [\Lambda_2 : \Lambda_1]$.

In particular, if Γ_1 and Γ_2 are harmonious in G then $[\Gamma_2 : \Gamma_1] = [\log \Gamma_2 : \log \Gamma_1]$.

The proof of this proposition will require the following classical fact.

Proposition 4.4.12 (Compatibility of Haar measures). Let G be a simply connected nilpotent Lie group and let \mathfrak{g} be the Lie algebra of G . If μ is a translation-invariant, locally finite measure on \mathfrak{g} (i.e., a Lebesgue measure) then the pushforward of μ by the exponential map is a locally finite measure on G that is both left and right invariant (i.e., a bi-invariant Haar measure).

Proof of proposition 4.4.12. It suffices to prove that for every $X \in \mathfrak{g}$, the maps $\ell_X : Y \mapsto X \diamond Y$ and $r_X : Y \mapsto Y \diamond X$ preserve the Lebesgue measure on \mathfrak{g} . For this, it suffices to prove that the total derivatives $D\ell_X$ and Dr_X have $|\det(D\ell_X)| = |\det(Dr_X)| = 1$ at every $Y \in \mathfrak{g}$. To prove this, it suffices to prove that both $D\ell_X$ and Dr_X can be expressed as the sum of the identity and a nilpotent linear transformation, which can be done via an explicit computation with the BCH formula. For details see e.g. [Wan23, Proposition 2.1.1]. \square

Proof of proposition 4.4.11. We start by constructing large sets in \mathfrak{g} that have “small boundary-to-volume ratio” in both the additive and multiplicative senses. (That is, the sets we construct will yield a Følner sequence for both addition and BCH multiplication on \mathfrak{g} .) If G were assumed to be a Carnot group we could use the logarithms of balls in a homogeneous left-invariant metric; we will perform a similar construction for a general Lie group. Let $(\mathfrak{g}_i)_{i \geq 0}$ be the lower central series of \mathfrak{g} , and for each $i \geq 0$ let V_i be such that $\mathfrak{g}_i = V_i \oplus \mathfrak{g}_{i+1}$, so that we can write $\mathfrak{g} = V_1 \oplus V_2 \oplus \cdots \oplus V_s$. For notational convenience we also write $V_i = \{0\}$ for $i > s$. In contrast to the Carnot case, it is not necessarily the case that V_1 generates \mathfrak{g} as a Lie algebra or that $[V_i, V_j] \subseteq V_{i+j}$, but we do have that $[V_i, V_j] \subseteq [\mathfrak{g}_i, \mathfrak{g}_j] \subseteq \mathfrak{g}_{i+j} = \bigoplus_{k=i+j}^s V_k$ for every $i, j \geq 1$. Fix an isomorphism of vector spaces $\mathfrak{g} \cong \mathbb{R}^d$ for some $d \geq 1$ and let $\|\cdot\|_\infty$ be the associated ∞ -norm on \mathfrak{g} . For each $\lambda > 0$ let

$$F_\lambda := \{X \in \mathfrak{g} : \max_i \lambda^{-i} \|X_i\|_\infty \leq 1\} = \{X \in \mathfrak{g} : \max_i \|X_i\|_\infty^{1/i} \leq \lambda\},$$

where we write $X = \sum_i X_i$ for the decomposition of X induced by the decomposition $\mathfrak{g} = V_1 \oplus V_2 \oplus \cdots \oplus V_s$. The volume of F_λ satisfies

$$\text{vol}(F_\lambda) = \lambda^q$$

for an integer $q \geq 1$, which can be expressed as $q = \sum_{i=1}^s i \dim(V_i)$ (this is equal to the *homogeneous dimension* of G). For each $X \in \mathfrak{g}$ we define

$$((X)) = \inf\{\lambda > 0 : X \in F_\lambda\} = \max_i \|X_i\|_\infty^{1/i},$$

so that $((X)) \leq \|X\|_\infty^{1/i}$ for every $1 \leq i \leq s$ and $X \in \mathfrak{g}_i$ with equality if $X \in V_i$. Note that this is *not* a norm on \mathfrak{g} since it does not satisfy $((\lambda X)) = \lambda((X))$ for all $X \in \mathfrak{g}$ and $\lambda > 0$. (Rather, it scales under a certain graded dilation map as in the Carnot case.) Moreover, unlike the metrics we considered on Carnot groups, $((\cdot))$ will not be left-invariant in general. Nevertheless, it does trivially satisfy the triangle inequality in the form

$$\begin{aligned} ((X + Y)) &= \max_i \|X_i + Y_i\|_\infty^{1/i} \leq \max_i (\|X_i\|_\infty + \|Y_i\|_\infty)^{1/i} \\ &\leq \max_i (\|X_i\|_\infty^{1/i} + \|Y_i\|_\infty^{1/i}) \leq ((X)) + ((Y)). \end{aligned} \quad (4.4.1)$$

As with norms, writing $((X)) = (((X + Y) - Y))$ yields the reverse inequality

$$((X + Y)) \geq ((X)) - ((Y)), \quad (4.4.2)$$

so that $F_{\lambda-((Y))} \subseteq F_\lambda + Y \subseteq F_{\lambda+((Y))}$ for every $Y \in \mathfrak{g}$ and $\lambda \geq ((Y))$.

We will need similar inequalities for the BCH product $((X \diamond Y))$. Since the Lie bracket is bilinear and \mathfrak{g} is finite-dimensional, there exists a constant C_1 such that $\|[X, Y]\|_\infty \leq C_1 \|X\|_\infty \|Y\|_\infty$ for every $X, Y \in \mathfrak{g}$. This implies that there exists a constant C_2 such that

$$(([X_i, Y_j])) \leq \|[X_i, Y_j]\|_\infty^{1/(i+j)} \leq C_1^{1/(i+j)} \|X_i\|_\infty^{1/(i+j)} \|Y_j\|_\infty^{1/(i+j)} \leq C_2 ((X_i))^{i/(i+j)} ((Y_j))^{j/(i+j)}$$

for every $X_i \in V_i$ and $Y_j \in V_j$. Together with (4.4.1) this implies that there exists a constant C_3 such that

$$(([X, Y])) \leq C_3 \max \left\{ ((X))^{1-\theta} ((Y))^\theta : \frac{1}{s} \leq \theta \leq \frac{s-1}{s} \right\}$$

for every $X, Y \in \mathfrak{g}$. It follows by induction on $k \geq 2$ that if $L(X, Y)$ is any Lie monomial of degree $k \geq 2$ then, writing $L(X, Y) = [L_1(X, Y), L_2(X, Y)]$ for two Lie monomials of degree $a, k-a < k$,

$$\begin{aligned} ((L(X, Y))) &\leq C_3 \max \left\{ ((L_1(X, Y)))^{1-\theta} ((L_2(X, Y)))^\theta : \frac{1}{s} \leq \theta \leq \frac{s-1}{s} \right\} \\ &\leq C_3^{k-1} \max \left\{ ((X))^{1-\theta} ((Y))^\theta : \frac{1}{s^k} \leq \theta \leq \frac{s^k-1}{s^k} \right\} \end{aligned}$$

for every $X, Y \in \mathfrak{g}$. (The case that one of the monomials L_1 or L_2 has degree *one*, and so is equal to X or Y , must be checked separately.) We deduce from this together with (4.4.1) and the definition of BCH multiplication that there exists a constant C_4 such that

$$((X \diamond Y)) \leq ((X)) + ((Y)) + C_4 \max \left\{ ((X))^{1-\theta} ((Y))^\theta : \frac{1}{s^s} \leq \theta \leq \frac{s^s-1}{s^s} \right\} \quad (4.4.3)$$

for every $X, Y \in \mathfrak{g}$. Since $-Y$ is both the additive and BCH inverse of Y , we can write $X = X \diamond Y \diamond (-Y)$ to obtain that there exists a constant C_5 such that the complementary inequality

$$\begin{aligned} ((X)) &\leq ((X \diamond Y)) + ((Y)) + C_4 \max \left\{ ((X \diamond Y))^{1-\theta} ((Y))^\theta : \frac{1}{s^s} \leq \theta \leq \frac{s^s-1}{s^s} \right\} \\ &\leq ((X \diamond Y)) + C_5 ((Y)) + C_5 \max \left\{ ((X))^{1-\theta} ((Y))^\theta : \frac{1}{s^{2s}} \leq \theta \leq \frac{s^{2s}-1}{s^{2s}} \right\} \end{aligned} \quad (4.4.4)$$

holds for every $X, Y \in \mathfrak{g}$.

We now use the sets F_λ to prove the claim about indices. Suppose that Λ is an additive lattice in \mathfrak{g} and that Γ is a lattice in G . Let K_Λ be a fundamental domain for Λ in \mathfrak{g} and let K_Γ be a fundamental

domain for Γ in G . Since $K_\Lambda \cup \log K_\Gamma$ is compact, $\max\{((Y)) : Y \in K_\Lambda \cup \log K_\Gamma\}$ is finite. As such, it follows from (4.4.1) and (4.4.2) that there exist positive constants C and ε (depending on K_Λ and K_Γ) such that

$$((X)) - C \leq ((X + Y)) \leq ((X)) + C$$

and

$$((X)) - C((X))^{1-\varepsilon} - C \leq ((X \diamond Y)) \leq ((X)) + C((X))^{1-\varepsilon} + C$$

for every $x \in \mathfrak{g}$ and $y \in K_\Lambda \cup \log K_\Gamma$. Since the sets $(X + K_\Lambda : X \in \Lambda)$ and $(X \diamond \log K_\Gamma : X \in \log \Gamma)$ each cover \mathfrak{g} (with distinct sets having measure-zero intersection), it follows that

$$F_{\lambda-C} \subseteq \bigcup_{X \in F_\lambda \cap \Lambda} (X + K_\Lambda) \subseteq F_{\lambda+C}$$

and

$$F_{\lambda-C\lambda^{1-\varepsilon}-C} \subseteq \bigcup_{X \in F_\lambda \cap \log \Gamma} (X \diamond \log K_\Gamma) \subseteq F_{\lambda+C\lambda^{1-\varepsilon}+C}$$

for every $\lambda \geq 1$ such that $\lambda - C\lambda^{1-\varepsilon} - C > 0$. Taking volumes and using that addition and BCH multiplication are both measure-preserving, we deduce that

$$(\lambda - C)^q \leq |F_\lambda \cap \Lambda| \cdot \text{vol}(K_\Lambda) = \text{vol} \left(\bigcup_{X \in F_\lambda \cap \Lambda} (X + K_\Lambda) \right) \leq (\lambda + C)^q$$

and

$$(\lambda - C\lambda^{1-\varepsilon} - C)^q \leq |F_\lambda \cap \log \Gamma| \cdot \text{vol}(\log K_\Gamma) = \text{vol} \left(\bigcup_{X \in F_\lambda \cap \log \Gamma} (X \diamond \log K_\Gamma) \right) \leq (\lambda + C\lambda^{1-\varepsilon} + C)^q$$

for every $\lambda \geq 1$ such that $\lambda - C\lambda^{1-\varepsilon} - C > 0$, so that

$$\text{vol}(K_\Lambda) = \lim_{\lambda \rightarrow \infty} \frac{\lambda^q}{|F_\lambda \cap \Lambda|} \quad \text{and} \quad \text{vol}(\log K_\Gamma) = \lim_{\lambda \rightarrow \infty} \frac{\lambda^q}{|F_\lambda \cap \log \Gamma|}. \quad (4.4.5)$$

This is easily seen to imply the claim. \square

Corollary 4.4.13. Let G be a simply connected nilpotent Lie group of step s and dimension d , and let the constants $C_1 = C_1(s)$ and $C_2 = C_2(s)$ be as in Theorem 4.4.2. If Γ is a discrete subgroup of G then the harmonious subgroups $\mathcal{H}_+(\Gamma)$ and $\mathcal{H}_-(\Gamma)$ satisfy the index bounds

$$[\mathcal{H}_+(\Gamma) : \Gamma][\Gamma : \mathcal{H}_-(\Gamma)] = [\mathcal{H}_+(\Gamma) : \mathcal{H}_-(\Gamma)] \leq (C_2 C_1)^d.$$

Proof. We may assume without loss of generality that Γ is a lattice, replacing G by $\mathcal{C}(\Gamma)$ otherwise. Since $\log \mathcal{H}_+(\Gamma) \subseteq C_1 C_2 \cdot \log \mathcal{H}_-(\Gamma)$ we have that $[\log \mathcal{H}_+(\Gamma) : \log \mathcal{H}_-(\Gamma)] \leq (C_1 C_2)^d$, and the claim follows from proposition 4.4.11. \square

4.5 Proof of the main theorems

In this section we prove our main theorems, Theorems 4.1.3 and 8.1.1. We begin by proving the following proposition, which is a direct analogue of Theorem 4.3.1 for Carnot groups.

Proposition 4.5.1. Let G be a Carnot group with Lie algebra \mathfrak{g} , let d_G be a left-invariant homogeneous metric on G , and for each $r > 0$ let B_r be the ball of radius r around id in (G, d_G) . Then

$$\sup \{ \# \{ \langle H \cap B_{2^k} \rangle : k \in \mathbb{Z} \} : H \text{ a discrete subgroup of } G \} < \infty.$$

It will suffice for our applications that all relevant constants depend on the pair (G, d_G) in an arbitrary fashion.

Proof of proposition 4.5.1. Since d_G is consistent with the usual topology of G , the (closed) ball B_r is a compact subset of G containing a neighbourhood of the identity for each $r > 0$. In particular, there exists a convex, symmetric subset K of \mathfrak{g} and a constant $C_0 \geq 1$ such that $K \subseteq \log B_1 \subseteq C_0 K \subseteq \delta_{C_0}(K)$, where $(\delta_\lambda)_{\lambda > 0}$ is the dilation semigroup on the stratified Lie algebra \mathfrak{g} . Thus, the balls B_r are sandwiched between the exponentials of the dilates of K :

$$\delta_r(K) \subseteq \log B_r \subseteq \delta_{C_0 r}(K) \tag{4.5.1}$$

for every $r > 0$. The sets $\delta_\lambda(K)$ are all convex and symmetric since they are linear images of the convex symmetric set K ; this will allow us to apply Theorem 4.3.1 to an appropriately chosen additive subgroup of \mathfrak{g} .

Let $\Lambda = \log H$, and for each $k \in \mathbb{Z}$ let $H_k = \langle H \cap B_{2^k} \rangle$ and $\Lambda_k = \log H_k$. By Theorem 4.4.2, there exists an integer constant C_1 such that $\tilde{\Lambda} := \text{span}_{\mathbb{Z}}(C_1 \cdot \Lambda)$ is an additive, bracket-closed lattice in \mathfrak{g} whose exponential $\mathcal{H}_-(H)$ is a harmonious subgroup of G that is contained in H . It follows by a direct application of Theorem 4.3.1 that there exists a constant C_2 such that

$$\#\{\text{span}_{\mathbb{Z}}(\tilde{\Lambda} \cap \delta_\lambda(K)) : \lambda > 0\} \leq C_2.$$

For each $k \geq 0$ we have trivially that $H_k \cap B_{2^k} = H \cap B_{2^k}$ and hence that $\Lambda_k \cap \delta_{2^k}(K) = \Lambda \cap \delta_{2^k}(K)$. In particular,

$$\tilde{\Lambda} \cap \delta_{2^k}(K) \subseteq \tilde{\Lambda} \cap B_{2^k} \subseteq \Lambda \cap B_{2^k} = \Lambda_k \cap B_{2^k}.$$

On the other hand, letting $m = \lceil \log_2 C_0 \rceil$, we also have by (4.5.1) that

$$\tilde{\Lambda} \cap \delta_{2^{k+m}}(K) \supseteq \tilde{\Lambda} \cap B_{2^k} \supseteq (C_1 \cdot \Lambda) \cap B_{2^k} \supseteq C_1 \cdot (\Lambda \cap B_{2^k}) = C_1 \cdot (\Lambda_k \cap B_{2^k}).$$

Thus, if we define $\tilde{\Lambda}_k = \text{span}_{\mathbb{Z}}(\tilde{\Lambda} \cap \delta_{2^k}(K))$ for each $k > 0$ then

$$C_1 \cdot \text{span}_{\mathbb{Z}}(\Lambda_{k-m} \cap B_{2^{k-m}}) \subseteq \tilde{\Lambda}_k \subseteq \text{span}_{\mathbb{Z}}(\Lambda_k \cap B_{2^k}) \quad (4.5.2)$$

for every $k \in \mathbb{Z}$.

Consider an interval $[a, b] \cap \mathbb{Z}$ such that $\tilde{\Lambda}_k$ does not change as k varies over $[a, b] \cap \mathbb{Z}$. Then we have by (4.5.2) that

$$C_1 \cdot \text{span}_{\mathbb{Z}}(\Lambda_{b-m} \cap B_{2^{b-m}}) \subseteq \text{span}_{\mathbb{Z}}(\Lambda_a \cap B_{2^a})$$

and hence that

$$[\text{span}_{\mathbb{Z}}(\Lambda_{b-m} \cap B_{2^{b-m}}) : \text{span}_{\mathbb{Z}}(\Lambda_a \cap B_{2^a})] \leq C_1^d.$$

Since strict sublattices have an index of at least 2, this in turn implies that

$$\#\{\text{span}_{\mathbb{Z}}(\Lambda_k \cap B_{2^k}) : k \in [a, b]\} \leq m + \#\{\text{span}_{\mathbb{Z}}(\Lambda_k \cap B_{2^k}) : k \in [a, b-m]\} \leq m + d \log_2 C_1.$$

Since $[a, b]$ was an arbitrary interval over which $\tilde{\Lambda}_k$ remained constant, it follows that there exist constants C_3 and C_4 such that

$$\#\{\text{span}_{\mathbb{Z}}(\Lambda_k \cap B_{2^k}) : k \in \mathbb{Z}\} \leq C_3 \#\{\tilde{\Lambda}_k : k \in \mathbb{Z}\} \leq C_4.$$

Now, for each k , (since the constant C_1 divides the constant appearing in lemma 4.4.4) the set $C_1 \cdot \text{span}_{\mathbb{Z}}(\mathcal{B}(\frac{1}{C_1}(\Lambda_k \cap B_{2^k})) \subseteq \log \mathcal{H}_+(H_k)$ is an additive, bracket-closed subgroup of \mathfrak{g} whose exponential is a subgroup of G that contains a generating set for H_k , so that

$$\Lambda_k \subseteq C_1 \cdot \text{span}_{\mathbb{Z}}\left(\mathcal{B}\left(\frac{1}{C_1}(\Lambda_k \cap B_{2^k})\right)\right) \subseteq \log \mathcal{H}_+(H_k)$$

for every k . As such, if $[a, b] \cap \mathbb{Z}$ is an interval such that $\text{span}_{\mathbb{Z}}(\Lambda_k \cap B_{2^k})$ does not change as k varies over $[a, b] \cap \mathbb{Z}$ then we have that

$$\Lambda_a \subseteq \Lambda_b \subseteq C_1 \cdot \text{span}_{\mathbb{Z}}\left(\mathcal{B}\left(\frac{1}{C_1}(\Lambda_a \cap B_{2^a})\right)\right) \subseteq \log \mathcal{H}_+(\Lambda_a)$$

and hence by corollary 4.4.13 that $[H_b : H_a] \leq C_5$ for some constant C_5 . Arguing as above, this implies that there exist constants C_6 and C_7 such that

$$\#\{H_k : k \in \mathbb{Z}\} \leq C_6 \#\{\text{span}_{\mathbb{Z}}(\Lambda_k \cap B_{2^k}) : k \in \mathbb{Z}\} \leq C_7,$$

completing the proof. \square

Our next goal is to use proposition 4.5.1 to prove the special case of Theorem 4.1.3 in which the group is nilpotent of bounded step.

Proposition 4.5.2 (Exploring subgroups of nilpotent groups). *For each $s, k \geq 1$ there exists a constant $C(s, k)$ such that the following holds. Let N be a nilpotent group of step s generated by some set S with $|S| \leq k$, let H be a subgroup of N , and for each $n \geq 1$ let H_n be the subgroup of H generated by elements that have word length at most 2^n in (G, S) . Then*

$$\#\{n : H_{n+1} \neq H_n\} \leq C.$$

Proof of Proposition 4.5.2. We first argue that it suffices to consider the case that N is equal to the free step- s nilpotent group $N_{s,S}$. Let N , S , and $(H_n)_{n \geq 0}$ be as in the statement of the theorem. Let $N_{s,S}$ be the free step- s nilpotent group over S and let $G_{s,S}$ be the free step- s nilpotent Lie group over S , so that $N_{s,S}$ can be identified with the subgroup of $G_{s,S}$ generated by S . By the universal property of $N_{s,S}$, there exists a homomorphism $\pi : N_{s,S} \rightarrow N$ satisfying $\pi(x) = x$ for every $x \in S$, which is necessarily unique and surjective since S generates N . Thus, π maps the word metric r -ball in $(N_{s,S}, S)$ to the word metric r -ball in (N, S) for every $r \geq 0$. For each $n \geq 0$, let \tilde{H}_n be the subgroup of $N_{s,S}$ generated by the elements of $\pi^{-1}(H)$ that have word length at most 2^n in $(N_{s,S}, S)$. Letting K denote the kernel of π , we observe that the subgroup $K\tilde{H}_n = \{kh : k \in K, h \in \tilde{H}_n\}$ of $N_{s,S}$ is equal to the preimage $\pi^{-1}(H_n)$: On the one hand, since K is normal in $K\tilde{H}_n$, $\pi(K\tilde{H}_n)$ is a subgroup of N that contains the set of words in H that have word length at most 2^n , and therefore contains H_n . On the other hand, if $x = kh$ is an element of $\pi^{-1}(H_n)$ then we can write $h = h_1 h_2 \cdots h_\ell$ as a product of elements of $\pi^{-1}(H)$ of word length at most 2^n in $(N_{s,S}, S)$, so that $\pi(h_i)$ is an element of H of word length at most 2^n in (N, S) and $\pi(x)$ belongs to $\pi^{-1}(H_n)$ as required. Since π is surjective, we have the chain of implications

$$(\tilde{H}_{n+1} = \tilde{H}_n) \Rightarrow (K\tilde{H}_{n+1} = K\tilde{H}_n) \Rightarrow (H_{n+1} = H_n).$$

Thus, it suffices to prove that the theorem holds with N and H replaced by $N_{s,S}$ and $\pi^{-1}(H)$.

From now on we assume that $N = N_{s,S}$ and let $G = G_{s,S}$. Since N , G , and the embedding $N \rightarrow G$ are determined up to isomorphism by s and $|S|$, we are now free to use constants that depend on this data in an arbitrary way (but must still be independent of the choice of subgroup $H \subseteq N$). By Pansu's theorem as formulated in (4.2.2), there exists a left-invariant Carnot metric d_G on G such that

$$\frac{d_G(\text{id}, x)}{d_S(\text{id}, x)} \rightarrow 1 \quad \text{as } x \rightarrow \infty \text{ in } N.$$

(For the argument to work we need only that the embedding of (N, d_S) into (G, d_G) is a quasi-isometry.) In particular, there exist positive constants c and C such that

$$\log N \cap B_{cn} \subseteq \log(\bar{S}^n) \subseteq \log N \cap B_{Cn}$$

for every $n \geq 0$, where we write $\bar{S} = S \cup \{\text{id}\} \cup S^{-1}$. This implies that if H is a subgroup of N then

$$\langle H \cap B_{c2^k} \rangle \subseteq H_k \subseteq \langle H \cap B_{C2^k} \rangle$$

for every $k \geq 0$, and the claim follows easily from this together with proposition 4.5.1. \square

We next deduce Theorem 4.1.3 from Proposition 4.5.2 and the Breuillard-Green-Tao theorem. The proof will use the following elementary lemmas, the first of which is proven in [BT16, Lemma 4.2]. Recall that we write $\bar{S} = S \cup \{\text{id}\} \cup S^{-1}$.

Lemma 4.5.3. Let G be a group, and let H be a subgroup of G with index at most n . If S is a finite generating set for G , then $(\bar{S})^{2n-1} \cap H$ is a generating set for H .

Lemma 4.5.4. Let G be a group, and let $H_1 \subseteq H_2 \subseteq \dots \subseteq H_n$ and $H'_1 \subseteq H'_2 \subseteq \dots \subseteq H'_n$ be two chains of subgroups such that H'_i is a subgroup of H_i for each $1 \leq i \leq n$. Then

$$\#\{H_i : 1 \leq i \leq n\} \leq (1 + \lfloor \log_2(\max_i [H_i : H'_i]) \rfloor) \cdot \#\{H'_i : 1 \leq i \leq n\}.$$

Proof of lemma 4.5.4. By taking a subsequence if necessary, we may assume that $H_{i+1} \neq H_i$ for every $1 \leq i < n$. Let $\ell = \#\{H'_i : 1 \leq i \leq n\}$ and for each $0 \leq k < \ell$ let i_k be the k th time H'_i changes, so that $i_0 = 1$ and $i_k = \min\{i > i_{k-1} : H'_i \neq H'_{i_{k-1}}\}$ for each $1 \leq k < \ell$. We also set $i_\ell = n+1$ for notational convenience. Since $H'_{i_{k-1}} = H'_{i_k}$ is a subgroup of $H_{i_{k-1}}$ for each $1 \leq k \leq \ell$, we have that

$$[H_{i_k} : H_{i_{k-1}}] \leq [H_{i_k} : H'_{i_{k-1}}] \leq \max_i [H_i : H'_i]$$

for every $1 \leq k \leq \ell$. On the other hand, we also have that $[H_{i_k} : H_{i_{k-1}}] = \prod_{i=i_{k-1}}^{i_k-1} [H_{i+1} : H_i] \geq 2^{i_k-1-i_{k-1}}$ for every $1 \leq k \leq \ell$, and hence that

$$n = \sum_{k=1}^{\ell} (i_k - i_{k-1}) \leq \ell \cdot (1 + \lfloor \log_2(\max_i [H_i : H'_i]) \rfloor)$$

as claimed. \square

Corollary 4.5.5. Let G be a group, let G' be a finite-index subgroup of G , and let $H_1 \subseteq H_2 \subseteq \dots \subseteq H_n$ be an increasing sequence of subgroups of G . Then

$$\#\{H_i : 1 \leq i \leq n\} \leq (1 + \lfloor \log_2[G : G'] \rfloor) \cdot \#\{H_i \cap G' : 1 \leq i \leq n\}.$$

Proof. Apply lemma 4.5.4 with $H'_i = H_i \cap G'$ and use that $[H_i : H_i \cap G'] \leq [G : G']$. \square

We now have everything we need to prove Theorem 4.1.3.

Proof of Theorem 4.1.3. Let $K \geq 1$ and let $r_0 = r_0(K)$ and $C_1 = C_1(K)$ be as in Theorem 4.1.1. Suppose that $r \geq r_0$ is such that $\text{Gr}(3r) \leq K \text{Gr}(r)$, and let $Q \triangleleft G$ and $N \triangleleft G/Q$ be as in Theorem 4.1.1. Since N has index at most C_1 in G/Q , lemma 4.5.3 implies that the set $S' := \pi((\bar{S})^{2C_1-1}) \cap N$ is a generating set for N , and the word metric associated to the pair (N, S') is bi-Lipschitz equivalent to the restriction of the word metric on $(G/Q, \pi(S))$ to N , with constants depending only on K .

Let $H' = (QH)/Q$, so that H' is a subgroup of G/Q and $\tilde{H} := H' \cap N$ is a subgroup of N . For each $n \geq 0$ let W'_n and \tilde{W}_n be the set of the elements of H' and \tilde{H} , respectively, that have word length at most 2^n with respect to $(G/Q, \pi(S))$, and define $H'_n = \langle W'_n \rangle$ and $\tilde{H}_n = \langle \tilde{W}_n \rangle$ for every $n \geq 0$. Since $[H'_n : H'_n \cap N] \leq [G/Q : N] \leq C_1$ and W'_n is a finite symmetric generating set for H'_n , the set $(W'_n)^{2C_1-1} \cap N$ is a generating set for $H'_n \cap N$, so that if we define $m = \lceil \log_2(2C_1 - 1) \rceil$ then

$$H'_n \cap N \subseteq \tilde{H}_{n+m} \subseteq H'_{n+m} \cap N \quad (4.5.3)$$

for every $n \geq 0$.

Using the above mentioned bi-Lipschitz equivalence between the two different word metrics on N and the fact that the step of N and the size of the generating set S' are bounded by constants depending only on K and k , it follows from Proposition 4.5.2 that there exists a constant $C_2 = C_2(K, k)$ such that $\#\{\tilde{H}_n : n \geq 0\} \leq C_2$. It follows from this and (4.5.3) that there exists a constant $C_3 = C_3(K, k)$ such that $\#\{N \cap H'_n : n \geq 0\} \leq C_3$, and hence by corollary 4.5.5 that there exists a constant $C_4 = C_4(K, k)$ such that $\#\{H'_n : n \geq 0\} \leq C_4$.

Now, as in the proof of Proposition 4.5.2, we have that $QH_n = \pi^{-1}(H'_n)$ for every $n \geq 0$. Observe that if $QH_{n+1} = QH_n$ but $H_{n+1} \neq H_n$ then there exist $q_n, q_{n+1} \in Q$, $h_n \in H_n$, and $h_{n+1} \in H_{n+1} \setminus H_n$ such that $q_n h_n = q_{n+1} h_{n+1}$, and hence that $h_{n+1} h_n^{-1} = q_{n+1}^{-1} q_n \in Q$. Since Q has diameter at most $C_1 r$, this implies that $h_{n+1} h_n^{-1}$ has word length at most $C_1 r$, contradicting the assumption that $h_{n+1} \notin H_n$ if $2^n \geq C_1 r$. It follows that there exists a constant $C_5 = C_5(K)$ such that

$$\#\{n \geq C_5 + \log_2 r : H_{n+1} \neq H_n\} \leq \#\{n \geq C_5 + \log_2 r : QH_{n+1} \neq QH_n\} \leq \#\{H'_n : n \geq 0\} \leq C_4,$$

which easily implies the claim. \square

It remains only to deduce Theorem 8.1.1 from Theorem 4.1.3. We will need the following lemma about the injectivity radius of the quotient $F_S/\langle\langle R_n \rangle\rangle \rightarrow G$. For each $n \geq 0$ let $G_n = F_S/\langle\langle R_n \rangle\rangle$.

Lemma 4.5.6. *Let G be a group with a finite generating set S . The quotient map $G_n \rightarrow G$ induces a map between Cayley graphs that restricts to an isomorphism between the balls of radius $2^{n-1} - 1$ around the identity.*

Proof of Lemma 4.5.6. It suffices to prove that the quotient map $\pi : G_n \rightarrow G$ is injective on $(\bar{S})^{2^{n-1}}$. (This implies that $\pi(xs) = \pi(y)$ if and only if $xs = y$ for every $s \in \bar{S}$ and x, y in the ball of radius $2^{n-1} - 1$ and hence that the balls of radius $2^{n-1} - 1$ are isomorphic.) Suppose for contradiction that this is false. Then there exist $u, v \in (\bar{S})^{2^{n-1}} \subseteq F_S$ such that $u^{-1}v \in R \setminus \langle\langle R_n \rangle\rangle$. Since u and v both belong to $(\bar{S})^{2^{n-1}}$, the product $u^{-1}v$ belongs to $(\bar{S})^{2^n}$, and since it also belongs to R it must belong to R_n by definition of R_n . This contradicts the assumption that $u^{-1}v \notin \langle\langle R_n \rangle\rangle$. \square

Proof of Theorem 8.1.1. Let $r_0 = r_0(K)$ and $C = C(K, k)$ be the constants from Theorem 4.1.3. Let $n_0 = \lceil 4 + \log_2 r \rceil$ and let $G' = F_S/\langle\langle R_{n_0} \rangle\rangle$. The projection $G' \rightarrow G$ induces a surjective graph homomorphism between the Cayley graphs $\text{Cay}(G', S)$ and $\text{Cay}(G, S)$ that restricts to an isomorphism between the balls of radius $2^{n_0-1} - 1 \geq 4r$. Let H be the subgroup of G' generated by $\bigcup_{n \geq n_0} (R_n/\langle\langle R_{n_0} \rangle\rangle)$ and, for each $n \geq n_0$, let H_n be the subgroup of H generated by $R_n/\langle\langle R_0 \rangle\rangle$. If $r \geq r_0$, we may apply Theorem 4.1.3 to G' and H to obtain that $\#\{H_n : n \geq n_0\} \leq C$, and it follows that

$$\#\{\langle\langle R_n \rangle\rangle : n \geq n_0\} = \#\{\langle\langle H_n \rangle\rangle : n \geq n_0\} \leq \#\{H_n : n \geq n_0\} \leq C$$

as claimed. \square

Acknowledgments

TH thanks Christian Gorski and Mikolaj Fraczyk for helpful conversations and thanks Matthew Tointon and Seung-Yeon Ryoo for comments on a draft. This work was supported by NSF grant DMS-2246494. Both authors thank the anonymous referees for their comments.

SUPERCRITICAL PERCOLATION ON FINITE TRANSITIVE GRAPHS I: UNIQUENESS OF THE GIANT COMPONENT

Joint with Tom Hutchcroft

Abstract

Let $(G_n)_{n \geq 1} = ((V_n, E_n))_{n \geq 1}$ be a sequence of finite, connected, vertex-transitive graphs with volume tending to infinity. We say that a sequence of parameters $(p_n)_{n \geq 1}$ in $[0, 1]$ is *supercritical* with respect to Bernoulli bond percolation \mathbb{P}_p^G if there exists $\varepsilon > 0$ and $N < \infty$ such that

$$\mathbb{P}_{(1-\varepsilon)p_n}^{G_n} (\text{the largest cluster contains at least } \varepsilon|V_n| \text{ vertices}) \geq \varepsilon$$

for every $n \geq N$ with $p_n < 1$. We prove that if $(G_n)_{n \geq 1}$ is sparse, meaning that the degrees are sublinear in the number of vertices, then the supercritical giant cluster is unique with high probability in the sense that if $(p_n)_{n \geq 1}$ is supercritical then

$$\lim_{n \rightarrow \infty} \mathbb{P}_{p_n}^{G_n} (\text{the second largest cluster contains at least } c|V_n| \text{ vertices}) = 0$$

for every $c > 0$. This result is new even under the stronger hypothesis that $(G_n)_{n \geq 1}$ has uniformly bounded vertex degrees, in which case it verifies a conjecture of Benjamini (2001). Previous work of many authors had established the same theorem for complete graphs, tori, hypercubes, and bounded degree expander graphs, each using methods that are highly specific to the examples they treated. We also give a complete solution to the problem of supercritical uniqueness for *dense* vertex-transitive graphs, establishing a simple necessary and sufficient isoperimetric condition for uniqueness to hold.

5.1 Introduction

Let $G = (V, E)$ be a countable graph that is connected and vertex-transitive, meaning that for all vertices $u, v \in V$ there is a graph automorphism $\phi \in \text{Aut } G$ with $\phi(u) = v$. Given $p \in [0, 1]$, *Bernoulli bond percolation* \mathbb{P}_p^G (abbreviated \mathbb{P}_p when the choice of G is clear from context) is the distribution of a random spanning subgraph ω formed by independently including each edge with probability p . We identify ω with an element of $\{0, 1\}^E$ where $\omega(e) = 1$ means that the edge e is present in ω . The edges in ω are called *open* and the rest are called *closed*. We are interested

primarily in the geometry of the connected components of ω , which we refer to as *clusters*. Much of the interest in the model stems from the existence of a *phase transition*: For infinite graphs, there is typically¹ a *critical probability* $0 < p_c < 1$ such that every cluster is finite almost surely when $p < p_c$, while at least one infinite cluster exists almost surely when $p > p_c$. For large finite graphs, one typically observes a similar phase transition in which a *giant* component, containing a positive proportion of all vertices, emerges as p is varied through a small interval. The regime in which an infinite/giant cluster exists is known as the *supercritical phase*. It is now known that the supercritical phase is always non-degenerate for bounded degree transitive graphs that are strictly more than one-dimensional in an appropriate coarse-geometric sense: this was proven for infinite graphs by Duminil-Copin, Goswami, Raoufi, Severo, and Yadin [Dum+20b] and for finite graphs by the second author and Tointon [HT21c].

Once one knows that the supercritical phase is non-degenerate, so that infinite/giant clusters exist for sufficiently large values of p , a central problem is to understand the *number* of these clusters. For *infinite* transitive graphs, this is the subject of a famous conjecture of Benjamini and Schramm [BS96c] stating that the infinite cluster is unique for every $p > p_c$ if and only if the graph is *amenable*. The ‘if’ direction of this conjecture follows from the classical work of Aizenman, Kesten, and Newman [AKN87a] and Burton and Keane [BK89], while the ‘only if’ direction remains open in general; see e.g. [MR3352259; Hut20f; Hut19a; LP16c; PS00] for an overview of what is known.

In contrast, for *finite* transitive graphs, Benjamini [Ben01a, Conjecture 1.2] conjectured in 2001 that the giant cluster should *always* be unique in the supercritical regime, irrespective of the geometry of the graph. Several works, some of which are very classical, have established versions of this conjecture in special cases including for complete graphs [ER61; Bol84], hypercubes [AKS82a; BKŁ92], Euclidean tori of fixed dimension [HR06], and bounded degree expanders [ABS04b]. Each of these works uses methods that are very specific to the example it treats, with the analysis of the tori $(\mathbb{Z}^d/n\mathbb{Z}^d)_{n \geq 1}$ in dimension $d \geq 3$ relying in particular on the important and technically challenging work of Grimmett and Marstrand [GM90b]. A related conjecture giving mild geometric conditions under which it should be impossible to have multiple giant clusters *both above and at criticality* was subsequently stated in the influential work of Alon, Benjamini, and Stacey [ABS04b, Conjecture 1.1].

A central difficulty in the study of this conjecture, and in the study of percolation on finite graphs

¹We keep the meaning of ‘typically’ intentionally vague. One very general conjecture [BS96c, Question 2] is that $p_c < 1$ for all (not necessarily transitive) infinite graphs with isoperimetric dimensional strictly greater than 1.

more generally, is that many of the most important qualitative tools used to study the infinite case, such as the ergodic theorem, break down completely in the finite case. For example, adapting the Burton–Keane uniqueness proof to the case of finite graphs merely shows that each vertex is unlikely to have three distinct large clusters in a small vicinity around it – a statement that need not in general be in tension with the existence of multiple giant components in a finite graph. Indeed, while the Burton–Keane proof applies to arbitrary *insertion-tolerant* automorphism-invariant percolation processes, it is possible to construct insertion-tolerant automorphism-invariant percolation processes on the torus $(\mathbb{Z}/n\mathbb{Z})^2$ that have multiple giant components with high probability. In fact, for the highly asymmetric torus $(\mathbb{Z}/n\mathbb{Z}) \times (\mathbb{Z}/2^n\mathbb{Z})$ one even has multiple giant clusters with good probability for Bernoulli percolation at appropriate values of p , although this arises as a feature of a discontinuous phase transition rather than of the supercritical phase *per se*. See Section 5.5 for further discussion of both examples. In light of these difficulties, any treatment of the uniqueness problem for finite transitive graphs must involve new techniques and use finer properties of supercritical percolation than in the infinite case.

In this paper we resolve Benjamini’s conjecture and hence also the supercritical case of the Alon–Benjamini–Stacey conjecture, giving a complete solution to the problem of supercritical uniqueness on large, finite, vertex-transitive graphs. In the forthcoming work in this series, we will prove moreover that the density of the giant component is concentrated, local, and equicontinuous in the supercritical regime and prove analogous theorems for the Fortuin–Kasteleyn random cluster model, Ising model, and Potts model.

Definition 5.1.1. We will assume all graphs to be locally finite and to contain at least one vertex. Let \mathcal{F} be the set of all isomorphism classes of finite, connected, simple (i.e., not containing loops or multiple edges), vertex-transitive graphs. (We will usually suppress the distinction between graphs and their isomorphism classes as much as possible when this does not cause any confusion.) Given an infinite set $\mathcal{H} \subseteq \mathcal{F}$, a function $\phi : \mathcal{H} \rightarrow \mathbb{R}$, and $\alpha \in \mathbb{R}$ we write $\lim_{G \in \mathcal{H}} \phi(G) = \alpha$ or “ $\phi(G) \rightarrow \alpha$ as $G \rightarrow \infty$ in \mathcal{H} ” to mean that for each $\varepsilon > 0$ there exists N such that $|\phi(G) - \alpha| \leq \varepsilon$ for every $G \in \mathcal{H}$ with at least N vertices, or equivalently that $\phi(G_n) \rightarrow \alpha$ for some (and hence every) enumeration $\mathcal{H} = \{G_1, G_2, \dots\}$ of \mathcal{H} . Similar conventions apply to the definition of $\limsup_{G \in \mathcal{H}}$, $\liminf_{G \in \mathcal{H}}$, and limits that may be equal to $+\infty$ or $-\infty$.

Let $G = (V, E)$ be a countable graph and consider a percolation configuration $\omega \in \{0, 1\}^E$. The connected components of ω are called *clusters*. We write K_u to denote the cluster containing the vertex u and write $u \leftrightarrow v$ for the event that $K_u = K_v$. Given a subset W of V , the *volume* of W is the number of vertices in W , denoted $|W|$, while if G is finite, the *density* of W is defined to be the

ratio $\|W\| := |W|/|V|$. We write K_1, K_2, \dots for the clusters of ω in decreasing order of volume. (Note the slight abuse of notation: here K_1 does not mean the cluster of a vertex labelled ‘1’.)

Given an infinite set $\mathcal{H} \subseteq \mathcal{F}$, we say that an assignment of parameters $p_c : \mathcal{H} \rightarrow [0, 1]$ is a **percolation threshold** if

1. $\lim_{G \in \mathcal{H}} \mathbb{P}_{(1-\varepsilon)p_c}^G (\|K_1\| \geq c) = 0$ for every $\varepsilon, c > 0$, and

2. For every $\varepsilon > 0$ there exists $\alpha > 0$ such that

$$\lim_{G \in \mathcal{H}} \mathbb{P}_{(1+\varepsilon)p_c}^G (\|K_1\| \geq \alpha) = 1,$$

where we set $\mathbb{P}_p^G = \mathbb{P}_1^G$ for $p \geq 1$.

Note that critical thresholds are *not* unique (when they exist), but any two percolation thresholds $p_c, \tilde{p}_c : \mathcal{H} \rightarrow [0, 1]$ must satisfy $p_c(G) \sim \tilde{p}_c(G)$ as $G \rightarrow \infty$ in \mathcal{H} . When a percolation threshold $p_c : \mathcal{H} \rightarrow [0, 1]$ exists, we say that $p : \mathcal{H} \rightarrow [0, 1]$ is *supercritical* if $\mathcal{H}' := \{G \in \mathcal{H} : p(G) < 1\}$ is finite or if \mathcal{H}' is infinite and satisfies

$$\liminf_{G \in \mathcal{H}'} \frac{p(G)}{p_c(G)} > 1.$$

We generalise this definition to include the case that p_c does not exist by saying that p is **super-critical** if $\mathcal{H}' := \{G \in \mathcal{H} : p(G) < 1\}$ is finite or if \mathcal{H}' is infinite and there exists $\varepsilon > 0$ such that

$$\liminf_{G \in \mathcal{H}'} \mathbb{P}_{(1-\varepsilon)p}^G (\|K_1\| \geq \varepsilon) \geq \varepsilon.$$

Note that these two definitions of supercriticality coincide when \mathcal{H} admits a threshold function, and in particular that the definition of supercriticality does not depend on the choice of threshold function. (Without the $(1 - \varepsilon)$ factor in $\mathbb{P}_{(1-\varepsilon)p}^G$, these definitions would not always coincide, for example for the highly asymmetric torus discussed in Example 5.1, which has giant clusters with good probability at a percolation threshold bounded away from 1.) The reason for introducing the set \mathcal{H}' is to ensure that every family has a supercritical sequence of parameters, namely the constant assignment $p(G) := 1$ for all G . It will also be helpful to have the finitary version of this definition: Given a *single* finite, connected, simple, vertex-transitive graph $G = (V, E)$ and given any $\varepsilon > 0$, we say that a parameter $p \in [0, 1]$ is **ε -supercritical** for G if $\mathbb{P}_{(1-\varepsilon)p}^G (\|K_1\| \geq \varepsilon) \geq \varepsilon$ and $|V| \geq 2\varepsilon^{-3}$. (There is some flexibility in how to choose this latter, technical condition that $|V|$ is not too small.)

We say that \mathcal{H} has the **supercritical uniqueness property** if

$$\lim_{G \in \mathcal{H}} \mathbb{P}_p^G (\|K_2\| \geq \varepsilon) = 0$$

for every supercritical $p : \mathcal{H} \rightarrow [0, 1]$ and every constant $\varepsilon > 0$.

We begin by stating our main result in the simplest-to-state and most interesting case when we have an infinite set $\mathcal{H} \subseteq \mathcal{F}$ that is *sparse*, meaning that the average vertex degree $d(G) := 2|E(G)|/|V(G)|$ of $G \in \mathcal{H}$ (which is the exact degree of every vertex because G is regular) satisfies $d(G) = o(|V(G)|)$ as $G \rightarrow \infty$ in \mathcal{H} . Note that in particular, if \mathcal{H} has uniformly bounded vertex degrees (i.e. $\sup_{G \in \mathcal{H}} d(G) < \infty$) then \mathcal{H} is sparse.

Theorem 5.1.2. *Let $\mathcal{H} \subseteq \mathcal{F}$ be an infinite set. If \mathcal{H} is sparse, then \mathcal{H} has the supercritical uniqueness property.*

Remark 5.1.1. The restriction to simple graphs is not very important and could be replaced by e.g. the assumption that there are a bounded number of parallel edges between any two vertices.

Remark 5.1.2. The Alon–Benjamini–Stacey conjecture [ABS04b, Conjecture 1.1] would follow immediately from Theorem 5.1.2 together with the plausible claim that the percolation phase transition is always continuous for bounded degree graph families satisfying the $\text{diam}(G) = o(|V(G)|/\log |V(G)|)$ condition they consider. More formally, such a claim would state that if $\mathcal{H} \subseteq \mathcal{F}$ is an infinite set with uniformly bounded vertex degrees that satisfies this condition, and $p : \mathcal{H} \rightarrow [0, 1]$ is any assignment of parameters such that $\liminf_{G \in \mathcal{H}} \mathbb{P}_{p(G)}^G (\|K_1\| \geq c) > 0$ for some $c > 0$, then p is supercritical in our sense. Unfortunately such a claim seems to be completely beyond the scope of present techniques and is a major open problem even for, e.g. the three-dimensional torus $(\mathbb{Z}/n\mathbb{Z})^3$. Indeed, it appears to be an open problem to prove that there are not multiple giant components at criticality in this example.

All the proofs in our paper are effective in the sense that they can in principle be used to produce explicit bounds on, say, the expected density of the second largest cluster. While we have not kept track of what these bounds are in all cases, we make note of the following simple explicit estimate implying Theorem 5.1.2 in the case that the graphs in question have bounded or subalgebraic vertex degrees.

Theorem 5.1.3. *There exists a universal constant C such that if $G = (V, E)$ is a finite, simple, connected, vertex-transitive graph with vertex degrees bounded by d , and $p \in [0, 1]$ is ε -supercritical*

for G for some $\varepsilon > 0$, then

$$\mathbb{P}_p^G \left(\|K_2\| \geq \lambda e^{C\varepsilon^{-18}} \sqrt{\frac{\log d}{\log |V|}} \right) \leq \frac{1}{\lambda}$$

for every $\lambda \geq 1$.

Note that this bound is only useful under the subalgebraic degree condition $\log d \ll \log |V|$. The constant given by our proof is fairly large, of order around 10^6 .

This bound has not been optimized and is known to be very far from optimal in classical examples. Indeed, the second largest cluster in supercritical percolation is known to be of order $\Theta(\log |V|)$ with high probability on both the complete graph [ER61; Bol84] and the hypercube [BKŁ92], while for a Euclidean torus of fixed dimension d , it is of order $\Theta((\log |V|)^{d/(d-1)})$ [HR06]. Similar results are established for a large class of *dense* graphs in [Bol+10a].

Remark 5.1.3. Let us now explain the relationship between our theorem and the Benjamini–Schramm $p_c < p_u$ conjecture [BS96c]. Suppose $(G_n)_{n \geq 1}$ is a bounded degree expander sequence converging locally to some infinite nonamenable transitive graph G . One may deduce either from our results or those of [ABS04b] (see also [Sar21b]) that there is always a unique giant component with high probability for supercritical percolation on G_n , a result that seems to be in tension with the conjectured existence of a non-uniqueness phase for percolation on the limit graph G . Naively, one might think that our definition of supercriticality for finite graphs should therefore be thought of more properly as an analogue of the *uniqueness phase* ($p > p_u$) for infinite graphs.

This is misleading. Indeed, it was proven in [BNP11a] that if $(G_n)_{n \geq 1}$ is a sequence of transitive, bounded degree expanders converging to an infinite, transitive, nonamenable graph G , then a sequence $(p_n)_{n \geq 1}$ is supercritical if and only if $\liminf p_n > p_c(G)$. The uniqueness/non-uniqueness transition on the limit graph G does manifest itself in the approximating finite graphs G_n , but as a transition in the *metric distortion* of the giant component rather than its uniqueness: the length of the path connecting two neighbouring vertices of G_n given that both vertices belong to the giant is tight as $n \rightarrow \infty$ when $p > p_u(G)$ and is not tight when $p_c(G) < p < p_u(G)$. In the second case, the open path connecting two such vertices in G_n has good probability to be very long, thus the two vertices become disconnected with positive probability in the limit. See [AB07] for related results and open problems for hypercube percolation.

The dense case. We now discuss how our results extend to *dense* graphs, where vertices have degree proportional to the number of vertices. In contrast to Theorem 5.1.2, it is not true in general

that dense graph families have the supercritical uniqueness property. Suppose for example that G_n is the Cartesian product of the complete graphs K_2 and K_n and that $p_n = 2/n$ for every $n \geq 1$. Then we have by the classical theory of Erdős–Rényi random graphs that each copy of K_n contains a giant component with high probability, while the number of ‘horizontal’ edges connecting the two copies of K_n converges in distribution to a Poisson(2) random variable and is therefore equal to zero with probability bounded away from zero as $n \rightarrow \infty$. Thus, the number of giant clusters in this example is unconcentrated and can be equal to either one or two, each with good probability.

Our next main result shows that examples of roughly this form are the only transitive counterexamples to the supercritical uniqueness property.

We will in fact characterise the failure of the supercritical uniqueness property for dense vertex-transitive graphs in two equivalent ways. We say that an infinite set $\mathcal{H} \subseteq \mathcal{F}$ is *dense* if $\liminf_{G \in \mathcal{H}} |E(G)|/|V(G)|^2 > 0$; this is equivalent to the vertex degree $d(G)$ growing linearly in the number of vertices in the sense that $\liminf_{n \rightarrow \infty} d(G)/|V(G)| > 0$.

Definition 5.1.4. Let $\mathcal{H} \subseteq \mathcal{F}$ be a set. Given $m \in \{2, 3, \dots\}$, we say that \mathcal{H} is *m-molecular* if it is infinite and dense and there exists a constant $C < \infty$ such that for each $G \in \mathcal{H}$ there exists a set of edges $F \subseteq E(G)$ satisfying the following conditions:

1. $G \setminus F$ has m connected components;
2. F is invariant under the action of $\text{Aut } G$;
3. $|F| \leq C |V(G)|$.

These conditions imply that the m connected components of $G \setminus F$ are dense, vertex-transitive, and isomorphic to each other. For example, the family of Cartesian products $\{K_n \square K_2 : n \geq 1\}$ discussed above is 2-molecular. We say \mathcal{H} is *molecular* if it is m -molecular for some $m \geq 2$.

Definition 5.1.5. Let $G = (V, E)$ be a finite graph. For each set $A \subseteq V$, we write $\partial_E A$ for the set of edges that have one endpoint in A and the other in $V \setminus A$. For each $\theta \in (0, 1/2]$, the quantity $\text{SEPARATOR}(G, \theta)$ is defined to be

$$\text{SEPARATOR}(G, \theta) :=$$

$$\min \left\{ |\partial_E A| : \theta \sum_{v \in V} \deg(v) \leq \sum_{v \in A} \deg(v) \leq (1 - \theta) \sum_{v \in V} \deg(v) \right\}.$$

In other words, $\text{SEPARATOR}(G, \theta)$ is the minimal number of edges needed to cut G into two pieces of roughly equal size. We say that a set \mathcal{H} of isomorphism classes of finite, simple graphs has *linear θ -separators* if \mathcal{H} is infinite and $\limsup_{G \in \mathcal{H}} \text{SEPARATOR}(G, \theta)/|V(G)| < \infty$.

The following theorem provides a complete solution to the problem of supercritical uniqueness for Bernoulli bond percolation on finite vertex-transitive graphs and implies Theorem 5.1.2 as a special case.

Theorem 5.1.6. *For every infinite set $\mathcal{H} \subseteq \mathcal{F}$, the following are equivalent:*

- (i) \mathcal{H} does not have the supercritical uniqueness property;
- (ii) \mathcal{H} contains a subset that is molecular;
- (iii) \mathcal{H} contains a subset with linear $1/3$ -separators;
- (iv) \mathcal{H} contains a dense subset with linear θ -separators for some $\theta \in (0, 1/2]$.

The dense case of this result sharpens the transitive case of a theorem of Bollobás, Borgs, Chayes, and Riordan [Bol+10a], who proved supercritical uniqueness for any (not necessarily transitive) dense graph sequence converging to an irreducible graphon. In our language, their result states that a dense graph family has the supercritical uniqueness property whenever it does not have any subquadratic separators, i.e., whenever

$$\liminf_{G \in \mathcal{H}} \text{SEPARATOR}(G, \theta)/|V(G)|^2 > 0$$

for every $\theta \in (0, 1/2]$ (see [Bol+10a, Lemma 7]). In fact, since they also prove that the giant cluster density is zero at the percolation threshold, their results imply uniqueness of the giant cluster for *all* (not necessarily supercritical) assignments of parameters. The same authors also established a formula for the limiting critical probability of dense graph sequences that we will use to prove the implication (iv) \Rightarrow (i) of Theorem 8.1.1. Further comparison of our results with those of [Bol+10a] is given in remark 5.4.2.

Remark 5.1.4. In [Eas22] the first author has built on the results of the present paper to characterise which infinite subsets of \mathcal{F} admit a percolation threshold. The obstacle to having a percolation threshold turns out to be the presence of molecular subsets for *infinitely many* values of $m \in \{2, 3, \dots\}$. In particular, every infinite subset of \mathcal{F} that is sparse admits a percolation threshold. So a posteriori, for Theorem 5.1.2 it suffices to work with the original (more natural) definition

of supercritical assignments of parameters, which refers to the percolation threshold, rather than the more general definition. Note the surprising logical order here: we first proved uniqueness of the supercritical giant cluster (in the present paper), then this was used to prove that there exists a percolation threshold for the emergence of a giant cluster (in [Eas22]).

About the proof and organization

We now briefly overview the structure of the paper and outline the proofs of the main steps.

Section 2: Lower bounds on point-to-point connection probabilities. In this section we prove that for ε -supercritical percolation on any finite, simple, connected, vertex-transitive graph, we always have a uniform lower bound $\mathbb{P}_p(x \leftrightarrow y) \geq \delta(\varepsilon) > 0$ on the probability that any two given vertices are connected. This was previously known only in the bounded degree case, with constants depending on the degree. Our argument starts by partitioning the vertex set into classes within which we have such a lower bound then recursively merges these classes until a single class contains the entire graph. The merging step makes use of a new high-degree version of insertion-tolerance, which allows us to open a single edge in a sufficiently large random set of edges.

Section 3: Uniqueness under the sharp density property. In this section we prove that the supercritical uniqueness property holds for any infinite set $\mathcal{H} \subseteq \mathcal{F}$ satisfying the *sharp density property*, meaning that $\mathbb{P}_p(\|K_1\| \geq \alpha)$ has a sharp threshold for every $\alpha \in (0, 1]$ in an appropriately uniform sense. This section is at the heart of the paper and contains the most significant new arguments.

The proof has two parts. First, given any particular supercritical parameter p , we use the sharp density property to (non-constructively) deduce the existence of a smaller parameter $q \leq p$ such that under \mathbb{P}_q there is a giant cluster whose density is *concentrated*, i.e., lies in a small interval with high probability. The point-to-point connection probability lower bound easily implies that this giant cluster is the *unique* giant cluster under \mathbb{P}_q with high probability.

The remainder of the proof consists in showing that non-uniqueness of the giant cluster under \mathbb{P}_p would imply non-concentration of the density of the giant cluster under \mathbb{P}_q , establishing a contradiction. To this end we introduce a new object called a *sandcastle*, which is a large subgraph that is not resilient to q/p -bond percolation. We observe that under the hypothesis that there are at least two giant clusters under \mathbb{P}_p with good probability, at least one of these clusters must be a sandcastle with good probability. As we pass from \mathbb{P}_p to \mathbb{P}_q in the standard monotone coupling of these measures, this sandcastle-cluster disintegrates into small clusters with good probability

by definition. On this event, the giant cluster under \mathbb{P}_q is constrained to the complement of this sandcastle-cluster from \mathbb{P}_p , where edges are distributed (conditionally) independently. We argue that if this is the case then there must be a subset of V with density significantly less than 1 that has good probability under \mathbb{P}_q to contain a cluster of approximately the same size as the typical global size of the largest cluster in V . Finally, we use the existence of this subset together with Harris' inequality and the point-to-point lower bound to deduce that $\|K_1\|$ is abnormally large with good probability under \mathbb{P}_q , contradicting the previously established concentration property.

In the final subsection of this section, we verify that subalgebraic degree graphs have the sharp density property, completing the proof of the main theorems in this case and establishing the quantitative estimate Theorem 5.1.3.

Section 4: Non-molecular graphs have the sharp density property. In this section we prove that the only way for an infinite subset of \mathcal{F} to *fail* to have the sharp density property is for it to contain a molecular subset, completing the proof of the main theorem.

Our argument uses a theorem of Bourgain [Fri99a] formalising the heuristic that increasing events without sharp thresholds are heavily influenced by the state of a bounded number of edges. In our case the event is the existence of a giant cluster of a given density. We apply a delicate sprinkling argument to iteratively reduce the size of this bounded-size set of edges until it contains a single edge. A novel trick in this induction is that during each iteration we use the second author's universal tightness theorem [Hut21b] and the high-degree analogue of insertion-tolerance from Section 2 to stick large (but not necessarily giant) clusters to both endpoints of an edge in the current set, allowing the small number of sprinkled edges to have a disproportionately large effect. This is the most technical part of the paper. Once the set of edges reaches a singleton, we apply Russo's formula to derive a contrasting *lower bound* on the sharpness of the threshold for our event. For this lower bound to not contradict our original upper bound, the graph in question must be dense, completing the proof of Theorem 5.1.2; the proof of the implication (i) \Rightarrow (ii) of Theorem 8.1.1 in the dense case relies on a second, rather subtle application of the sprinkling technology we develop to prove that the graph must in fact be molecular. Finally we show in Section 5.4 that the remaining non-trivial implication (iv) \Rightarrow (i) of Theorem 8.1.1 follows easily from the results of [Bol+10a].

We end the paper with some further discussion and closing remarks in Section 5.5.

5.2 Lower bounds on point-to-point connection probabilities

The goal of this section is to prove that point-to-point connection probabilities are uniformly bounded away from zero in ε -supercritical percolation on finite vertex-transitive graphs, where all relevant constants depend only on the parameter $\varepsilon > 0$. Given a finite graph G and constants $0 < \alpha, \delta \leq 1$, we define

$$p_c(\alpha, \delta) = p_c^G(\alpha, \delta) = \inf\{p \in [0, 1] : \mathbb{P}_p(\|K_1\| \geq \alpha) \geq \delta\}, \quad (5.2.1)$$

so that p is ε -supercritical if and only if $(1 - \varepsilon)p \geq p_c(\varepsilon, \varepsilon)$ and $|V(G)| \geq 2\varepsilon^{-3}$. By continuity and strict monotonicity, $p_c(\alpha, \delta)$ is equivalently the unique parameter satisfying $\mathbb{P}_{p_c(\alpha, \delta)}(\|K_1\| \geq \alpha) = \delta$.

Theorem 5.2.1. *Let $G = (V, E)$ be a finite, connected, simple, vertex-transitive graph and let $\varepsilon > 0$. If $|V| \geq 2\varepsilon^{-3}$ and $p \geq p_c(\varepsilon, \varepsilon)$ then*

$$\mathbb{P}_p(u \leftrightarrow v) \geq \tau(\varepsilon) := \exp[-10^5 \cdot \varepsilon^{-18}]$$

for every $u, v \in V$.

The exact value of this bound is not important for our purposes, and we have not attempted to optimize the relevant constants. For *bounded degree* graphs, a similar estimate follows from an argument essentially due to Schramm, which is recorded in [Ben01a] and in more detail in [HT21c, Lemma 2.1]. This argument yields in particular that if $G = (V, E)$ is a finite vertex-transitive graph and $p \geq p_c(\varepsilon, \varepsilon)$ then

$$\mathbb{P}_p(u \leftrightarrow v) \geq \exp\left[-3\left(\frac{2}{\varepsilon^2} \vee \frac{1}{p}\right)\log\left(\frac{2}{\varepsilon^2} \vee \frac{1}{p}\right)\right] \quad (5.2.2)$$

for every $u, v \in V$. This bound is adequate for our purposes in the bounded degree case, in which the condition $p \geq p_c(\varepsilon, \varepsilon)$ bounds p away from zero when $|V| \geq 2\varepsilon^{-3}$ by lemma 5.2.6 below. As such, readers who are already familiar with (5.2.2) and are only interested in the bounded degree case of our results may safely skip the remainder of this section. The estimate (5.2.2) does *not* yield a uniform lower bound on the two-point function in the high-degree case however, making a more refined analysis necessary at this level of generality.

Remark 5.2.1. The assumption that G is simple is not really needed for this theorem to hold: the proof works whenever G has degree at most $|V|$, and yields a similar statement (with different constants) under the assumption that the degrees are bounded by $C|V|$ for some constant C . No such uniform two-point lower bound holds without this assumption, as can be seen by taking the product $K_n \square K_2$ and replacing each edge of K_n by a large number of parallel edges.

Remark 5.2.2. For *infinite* transitive graphs, it was proven by Lyons and Schramm [LS99] (in the unimodular case) and Tang [Tan19] (in the nonunimodular case) that there is a unique infinite cluster at p if and only if $\inf_{x,y} \mathbb{P}_p(x \leftrightarrow y) > 0$. As such, one might naively expect that we could deduce our main theorems on uniqueness directly from Theorem 5.2.1 via a similar argument. This does not appear to be the case: firstly, we note that there *do* exist large finite transitive graphs and values of ε such that there are multiple giant components with good probability at certain values of $p \geq p_c(\varepsilon, \varepsilon)$, despite there always being a uniform lower bound on the two-point function at such values of p by Theorem 5.2.1. Indeed, the product $K_n \square K_2$ and the elongated torus $(\mathbb{Z}/n\mathbb{Z}) \times (\mathbb{Z}/2^n\mathbb{Z})$ both have this property. Secondly, the proofs of [LS99; Tan19] both rely essentially on *indistinguishability theorems* that are of an ergodic-theoretic nature and do not generalize to the finite-volume setting.

We will deduce Theorem 5.2.1 as an analytic consequence of the following inductive lemma. Recall that we write $\|A\| = |A|/|V|$ for the density of a set of vertices in a finite graph.

Lemma 5.2.2 (Two-point induction step). *Let $G = (V, E)$ be a finite, connected, vertex-transitive graph and let $p \geq 1/(2|V|)$. If $\tau > 0$ and $k \geq 1$ are such that $\|\{v \in V : \mathbb{P}_p(u \leftrightarrow v) \geq \tau\}\| \geq 2^{-k}$ for every $u \in V$ then there exists $\ell \in \{0, \dots, k-1\}$ such that*

$$\left\| \left\{ v \in V : \mathbb{P}_p(u \leftrightarrow v) \geq \tau^{3^3 2^{k-\ell}} 2^{-2\ell-10} \right\} \right\| \geq 2^{-\ell}$$

for every $u \in V$.

Before proving this lemma we first state and prove some general facts that will be used in the proof. The first is a standard bound on the diameter of dense graphs whose proof we include for completeness.

Lemma 5.2.3 (Dense graphs have bounded diameter). *Let $G = (V, E)$ be a finite, simple, connected graph. If every vertex of G has degree at least $a|V|$ for some $a > 0$ then $\text{diam } G \leq (3-a)/a$.*

Proof of Lemma 5.2.3. Let $a > 0$, and let G be a finite, connected graph with minimum vertex degree at least $a|V|$. Let u and v be two vertices of G and let $u = u_0, u_1, \dots, u_k = v$ be a minimal length path from u to v . Writing $k = 3m + r$ where m is a positive integer and $r \in \{0, 1, 2\}$, it suffices to prove that $m + 1 \leq 1/a$. Suppose for contradiction that this is not the case. Then

$$\sum_{i=0}^m \deg u_{3i} \geq (m+1) \cdot \min_{0 \leq i \leq m} \deg u_{3i} > |V|, \quad (5.2.3)$$

so, by the pigeonhole principle, we can find $i, j \in \{0, \dots, m\}$ with $i < j$ such that u_{3i} and u_{3j} have a common neighbour w . It follows that $u_0, u_1, \dots, u_{3i}, w, u_{3j}, \dots, u_{k-1}, u_k$ is a shorter path from u to v than u_0, \dots, u_k , a contradiction. \square

The second ingredient we will require is a quantitative form of insertion-tolerance. In bounded degree contexts, insertion-tolerance usually refers to the fact that the conditional probability of an edge being included in the configuration given the status of every other edge is bounded away from zero. Such a statement need not be valid in regimes of interest in the high-degree case, where p may be very small. Intuitively, the following proposition instead gives conditions in which we can insert exactly one edge into some sufficiently large *set* of edges with good probability. This proposition will be used again in the proof of Lemma 5.4.9.

Proposition 5.2.4 (Quantitative insertion tolerance). *Let $G = (V, E)$ be a finite graph, let $p \in (0, 1)$, and let $F \subseteq E$ be a collection of edges. Let $A \subseteq \{0, 1\}^E$ be an event, let $\eta > 0$ and suppose that for each configuration $\omega \in A$ there is a distinguished subset $F[\omega]$ with $F[\omega] \subseteq F \setminus \omega$ and $|F[\omega]| \geq \eta |F|$. If we define $A^+ := \{\omega \cup \{e\} : \omega \in A \text{ and } e \in F[\omega]\}$ then*

$$\mathbb{P}_p(A^+) \geq \frac{\eta^2}{1-p} \cdot \frac{p|F|}{p|F|+1} \cdot \mathbb{P}_p(A)^2.$$

Note that the hypotheses of this proposition force there to be at most $(1-\eta)|F|$ open edges in $F[\omega]$ whenever the event A holds. The lower bound appearing here has not been optimized, and a more careful implementation of our argument would give a $\mathbb{P}_p(A)/(-\log \mathbb{P}_p(A))$ term in place of the $\mathbb{P}_p(A)^2$ term above. The only important conclusion of this proposition for our purposes will be that if $p|F|$, $\mathbb{P}_p(A)$, and η are all bounded below by some constant $c > 0$ then there exists $\delta = \delta(c) > 0$ such that $\mathbb{P}_p^G(A^+) \geq \delta$.

Proof of Proposition 6.6.4. We will abbreviate $\mathbb{P} = \mathbb{P}_p^G$ and $\mathbb{E} = \mathbb{E}_p^G$. Given a set of edges H , we write $\omega|_H$ for the configuration of open and closed edges in H , which by a standard abuse of notation we will think of both as a subset of H and a function $H \rightarrow \{0, 1\}$. We can sample a configuration ω with law \mathbb{P} using the following procedure:

1. Sample the restriction $\omega|_{E \setminus F}$ of ω to $E \setminus F$.
2. Sample a uniformly random permutation π of the edges of F independently of $\omega|_{E \setminus F}$.
3. Sample a binomial random variable $N \sim \text{Binomial}(|F|, p)$ independently of $\omega|_{E \setminus F}$ and π .

4. Set $\omega|_F := \pi(\{1, \dots, N\})$.

Let $\tilde{\mathbb{P}}$ denote the joint measure of ω , π , and N sampled as in this procedure.

Let A and $F[\omega]$ be as in the statement of the proposition. The assumption that $|F[\omega]| \geq \eta|F| > 0$ and $F[\omega] \subseteq F \setminus \omega$ for every $\omega \in A$ guarantees that $N \leq (1 - \eta)|F| < |F|$ on the event that $\omega \in A$. By construction we have that $\omega = \omega|_{E \setminus F} \cup \pi(\{1, \dots, N\})$ so that for each $n \in \{1, \dots, |F|\}$ we can rewrite

$$\begin{aligned} \{\omega \in A^+ \} \cap \{N = n\} &= \{\omega|_{E \setminus F} \cup \pi(\{1, \dots, N\}) \in A^+ \} \cap \{N = n\} \\ &= \{\omega|_{E \setminus F} \cup \pi(\{1, \dots, n\}) \in A^+ \} \cap \{N = n\}. \end{aligned} \quad (5.2.4)$$

Note that the two events on the second line are independent. One way for the union $\omega|_{E \setminus F} \cup \pi(\{1, \dots, n\})$ to belong to A^+ is for $\omega|_{E \setminus F} \cup \pi(\{1, \dots, n-1\})$ to belong to A and for $\pi(n)$ to belong to the set $F[\omega|_{E \setminus F} \cup \pi(\{1, \dots, n-1\})]$. Since $|F[\nu]| \geq \eta|F|$ for every configuration $\nu \in A$, $\pi(n)$ belongs to this set with probability at least η conditional on $\omega|_{E \setminus F}$ and $\pi(\{1, \dots, n-1\})$ and we deduce that

$$\begin{aligned} \tilde{\mathbb{P}}(\omega \in A^+ \text{ and } N = n) &\geq \eta \tilde{\mathbb{P}}(\omega|_{E \setminus F} \cup \pi(\{1, \dots, n-1\}) \in A) \tilde{\mathbb{P}}(N = n) \\ &= \eta \tilde{\mathbb{P}}(\omega \in A \text{ and } N = n-1) \cdot \frac{\tilde{\mathbb{P}}(N = n)}{\tilde{\mathbb{P}}(N = n-1)}. \end{aligned} \quad (5.2.5)$$

The ratio of probabilities appearing here is given by

$$\frac{\tilde{\mathbb{P}}(N = n)}{\tilde{\mathbb{P}}(N = n-1)} = \frac{\binom{|F|}{n} p^n (1-p)^{|F|-n}}{\binom{|F|}{n-1} p^{n-1} (1-p)^{|F|-n+1}} = \frac{p(|F|-n+1)}{(1-p)n} \quad (5.2.6)$$

and we deduce that

$$\begin{aligned} \tilde{\mathbb{P}}(\omega \in A^+) &\geq \eta \sum_{n=1}^{|F|} \frac{p(|F|-n+1)}{(1-p)n} \tilde{\mathbb{P}}(\omega \in A \text{ and } N = n-1) \\ &= \frac{p\eta}{1-p} \tilde{\mathbb{E}} \left[\frac{|F|-N}{N+1} \mid \omega \in A \right] \tilde{\mathbb{P}}(\omega \in A), \end{aligned} \quad (5.2.7)$$

where we used that $N \leq (1-\eta)|F| < |F|$ whenever $\omega \in A$ in the second line. Since $(|F|-x)/(x+1)$ is convex we may apply Jensen's inequality to deduce that

$$\tilde{\mathbb{P}}(\omega \in A^+) \geq \frac{p\eta}{1-p} \frac{|F| - \tilde{\mathbb{E}}[N \mid \omega \in A]}{\tilde{\mathbb{E}}[N \mid \omega \in A] + 1} \tilde{\mathbb{P}}(\omega \in A). \quad (5.2.8)$$

Using again that $N \leq (1 - \eta)|F|$ whenever $\omega \in A$ and using the bound $\tilde{\mathbb{E}}[N \mid \omega \in A] \leq \tilde{\mathbb{E}}[N]/\tilde{\mathbb{P}}(\omega \in A) = p|F|/\tilde{\mathbb{P}}(\omega \in A)$ we deduce that

$$\tilde{\mathbb{P}}(\omega \in A^+) \geq \frac{\eta^2}{1-p} \cdot \frac{p|F|}{p|F|+1} \cdot \tilde{\mathbb{P}}(\omega \in A)^2, \quad (5.2.9)$$

concluding the proof. \square

We are now ready to prove lemma 5.2.2.

Proof of lemma 5.2.2. Let G^* be the graph with vertex set V in which two vertices u and v are connected by an edge if and only if $\mathbb{P}_p(u \leftrightarrow v) \geq \tau$. The graph G^* is simple, vertex-transitive, and is dense in the sense that its vertex degrees are all at least $2^{-k}|V|$. Let C be a connected component of G^* , noting that $2^{-k} \leq \|C\| \leq 1$. Applying Lemma 5.2.3 to the subgraph of G^* induced by C implies that $\text{diam}(C) \leq (3 - 2^{-k}/\|C\|)/(2^{-k}/\|C\|) = (3\|C\| - 2^{-k})/2^{-k} \leq 3 \cdot 2^k\|C\|$. Thus, if u and v belong to the same connected component of G^* then there exists a sequence of vertices $u = u_0, u_1, \dots, u_n = v$ with $n \leq 3 \cdot 2^k\|C\|$ such that u_i is adjacent to u_{i-1} in G^* for every $1 \leq i \leq n$. It follows by the Harris-FKG inequality that

$$\mathbb{P}_p(u \leftrightarrow v) \geq \prod_{i=1}^n \mathbb{P}_p(u_{i-1} \leftrightarrow u_i) \geq \tau^{3 \cdot 2^k\|C\|} =: \tau_1 \quad (5.2.10)$$

for every u, v in the same connected component of G^* . If G^* is connected then $\|C\| = 1$ and the claim follows with $\ell = 0$, so we may assume that G^* is disconnected and that $2^{-\ell-1} \leq \|C\| < 2^{-\ell}$ for some $\ell \in \{1, \dots, k-1\}$.

Since G is connected there must exist at least one edge of G connecting C to $V \setminus C$. Since the connected-component equivalence relation on G^* is invariant under the automorphisms of G , it follows by vertex-transitivity of G that *every* vertex in C belongs to an edge from C to $V \setminus C$. Letting $\partial_E C$ be the set of edges of G with one endpoint in C and the other in $V \setminus C$, it follows that $|\partial_E C| \geq |C|$. Since there are $|V|/|C|$ connected components of G^* , it follows by the pigeonhole principle that there exists a connected component $C' \neq C$ of G^* such that at least $|C|^2/|V|$ edges of $|\partial_E C|$ have their other endpoint in C' . Let I be the set of *oriented* edges e of G with tail $e^- \in C$ and head $e^+ \in C'$, so that $|I| \geq |C|^2/|V|$.

Fix vertices $u \in C$ and $v \in C'$ that are the endpoints of some edge in I . We claim that

$$\mathbb{P}_p(u \leftrightarrow v) \geq \frac{\tau_1^8}{2^8} \|C\|^2. \quad (5.2.11)$$

Let L be the random set of oriented edges $e \in I$ such that $u \leftrightarrow e^-$ and $v \leftrightarrow e^+$. For each $e \in I$, the Harris-FKG inequality and eq. (5.2.10) imply that $\mathbb{P}_p(u \leftrightarrow e^- \text{ and } v \leftrightarrow e^+) \geq \tau_1^2$ and hence that $\mathbb{E}_p|L| \geq \tau_1^2|I|$. Applying Markov's inequality to $|I \setminus L|$ we obtain that

$$\begin{aligned} \mathbb{P}_p\left(|L| \geq \frac{\tau_1^2}{2}|I|\right) &= 1 - \mathbb{P}_p\left(|I \setminus L| > \frac{1}{2}(2 - \tau_1^2)|I|\right) \\ &\geq 1 - \frac{2 - 2\tau_1^2}{2 - \tau_1^2} = \frac{\tau_1^2}{2 - \tau_1^2} \geq \frac{1}{2}\tau_1^2. \end{aligned} \tag{5.2.12}$$

Let A be the event that $|L| \geq \frac{\tau_1^2}{2}|I|$ and that every edge of L is closed, so that $\{u \leftrightarrow v\} \supseteq \{|L| \geq \tau_1^2|I|/2\} \setminus A$. If $\mathbb{P}_p(A) \leq \tau_1^2/4$ then

$$\mathbb{P}_p(u \leftrightarrow v) \geq \mathbb{P}_p\left(|L| \geq \frac{\tau_1^2}{2}|I|\right) - \mathbb{P}_p(A) \geq \frac{1}{4}\tau_1^2, \tag{5.2.13}$$

which is stronger than the claimed inequality, so we may assume that $\mathbb{P}_p(A) \geq \tau_1^2/4$. In this case, applying Proposition 6.6.4 with $F = I$, $F[\omega] = L$, and $\eta = \tau_1^2/2$ yields that

$$\mathbb{P}_p(u \leftrightarrow v) \geq \mathbb{P}_p(A^+) \geq \frac{\tau_1^8}{64} \cdot \frac{p|I|}{p|I| + 1} \geq \frac{\tau_1^8}{64} \cdot \frac{|I|}{|I| + 2|V|} \geq \frac{\tau_1^8}{64} \cdot \frac{|C|^2}{|C|^2 + 2|V|^2}, \tag{5.2.14}$$

where we used the assumption $p \geq \frac{1}{2|V|}$ in the second inequality and the inequality $|I| \geq |C|^2/|V|$ in the third. Bounding $|C|^2 + 2|V|^2$ by $4|V|^2$ completes the proof of (5.2.11). It follows from this inequality, (5.2.10), and a further application of Harris-FKG that

$$\mathbb{P}_p(u \leftrightarrow w) \geq \mathbb{P}_p(u \leftrightarrow v)\mathbb{P}_p(v \leftrightarrow w) \geq \frac{\tau_1^9}{2^8}\|C\|^2 \tag{5.2.15}$$

for every $w \in C'$. The same inequality also holds for every $w \in C$ by (5.2.10). Thus, recalling that $1 \leq \ell < k$ is such that $2^{-\ell-1} \leq \|C\| < 2^{-\ell}$ and using that $\tau_1 = \tau^{3 \cdot 2^k \|C\|}$, we deduce that

$$\left\| \left\{ w \in V : \mathbb{P}_p(u \leftrightarrow w) \geq \tau^{3 \cdot 2^{k-\ell}} 2^{-2\ell-10} \right\} \right\| \geq \|C \cup C'\| \geq 2^{-\ell}, \tag{5.2.16}$$

completing the proof. (The reason why we have worked so hard to get an extra factor of 2 via $\|C \cup C'\| \geq 2\|C\|$ in the final inequality above will become clear shortly.) \square

lemma 5.2.2 implies the following general inequality by induction, from which we will deduce Theorem 5.2.1 as a special case.

Lemma 5.2.5. Let $G = (V, E)$ be a finite, connected, vertex-transitive graph and let $p \geq \frac{1}{2|V|}$. If $\tau > 0$ and $k \geq 0$ are such that $\|\{v \in V : \mathbb{P}_p(u \leftrightarrow v) \geq \tau\}\| \geq 2^{-k}$ for every $u \in V$ then

$$\mathbb{P}_p(u \leftrightarrow v) \geq \tau^{(54)^k} 2^{-k(54)^k}$$

for every $u, v \in V$.

Proof of lemma 5.2.5. Fix $G = (V, E)$ and $p \geq \frac{1}{2|V|}$. Applying lemma 5.2.2 recursively implies that there exists a decreasing sequence of non-negative integers $k = k_0 > k_1 > \dots > k_m = 0$ with $m \leq k$ such that if we define the sequence of positive real numbers τ_0, \dots, τ_m recursively by $\tau_0 = \tau$ and

$$\tau_{i+1} = \tau_i^{3^3 2^{k_i - k_{i+1}}} 2^{-2k_{i+1} - 10} \quad (5.2.17)$$

for each $0 \leq i \leq m-1$ then

$$\|\{v \in V : \mathbb{P}_p(u \leftrightarrow v) \geq \tau_i\}\| \geq 2^{-k_i} \quad (5.2.18)$$

for each $1 \leq i \leq m$. It follows by induction on i that

$$\tau_i = \tau^{3^3 2^{k_i - k_i}} \prod_{j=1}^i \left(2^{-2k_j - 10}\right)^{3^3(i-j)2^{k_j - k_i}} \quad (5.2.19)$$

for every $0 \leq i \leq m$ and hence that

$$\begin{aligned} \tau_m &= \tau^{3^3 m 2^k} \prod_{j=1}^m \left(2^{-2k_j - 10}\right)^{3^3(m-j)2^{k_j}} \geq \tau^{3^3 m 2^k} \prod_{j=1}^m \left(2^{-2k_j - 10}\right)^{3^3(m-j)2^{k_j}} \\ &= \tau^{3^3 m 2^k} 2^{-(2k+10)3^3 m 2^k \sum_{j=1}^m 3^{-3j}} \end{aligned} \quad (5.2.20)$$

$$\geq \tau^{3^3 m 2^k} 2^{-3^3 m k 2^k}, \quad (5.2.21)$$

where we used the inequality $(2k+10) \sum_{j=1}^m 3^{-3j} \leq 12k \cdot (1/26) \leq k$ for $k \geq 1$ to simplify the final expression. The claim follows since $m \leq k$ and $3^3 \cdot 2 = 54$. \square

To deduce Theorem 5.2.1 from lemma 5.2.5 we will need the following elementary but useful lower bound on the critical probability.

Lemma 5.2.6. If G is a finite graph with maximum degree d and $\varepsilon > 0$ is such that $|V| \geq 2\varepsilon^{-3}$ then $p_c(\varepsilon, \varepsilon) \geq 1/2d$. In particular, if p is ε -supercritical then $p \geq 1/2d$.

Proof. Fix a vertex $v \in V$. For each r , the expected number of open simple paths of length r starting at v is at most $d(d-1)^{r-1} \leq d^r$ and it follows that if $p < 1/2d$ then $\mathbb{E}|K_v| \leq \sum_{r=0}^{\infty} p^r d^r < 2$. On the other hand, if $p \geq p_c(\varepsilon, \varepsilon)$ then

$$\sum_{v \in V} \mathbb{E}_p |K_v| \geq \mathbb{E}_p |K_1|^2 \geq \varepsilon^3 |V|^2,$$

so if the inequalities $p < 1/2d$ and $p \geq p_c(\varepsilon, \varepsilon)$ both hold then $|V| < 2\varepsilon^{-3}$. \square

We are now ready to conclude the proof of Theorem 5.2.1.

Proof of Theorem 5.2.1. Since $p \geq p_c(\varepsilon, \varepsilon)$ we have that $\mathbb{P}_p(\|K_1\| \geq \varepsilon) \geq \varepsilon$ and hence that $\mathbb{P}_p(\|K_u\| \geq \varepsilon) \geq \varepsilon \mathbb{P}_p(\|K_1\| \geq \varepsilon) \geq \varepsilon^2$ for every $u \in V$ by vertex-transitivity. It follows in particular that $\sum_{v \in V} \mathbb{P}_p(u \leftrightarrow v) = \mathbb{E}_p |K_u| \geq \varepsilon^3 |V|$ and hence by Markov's inequality that

$$\left\| \left\{ v \in V : \mathbb{P}_p(u \leftrightarrow v) \geq \frac{\varepsilon^3}{2} \right\} \right\| \geq \frac{\varepsilon^3}{2}. \quad (5.2.22)$$

Moreover, since $|V| \geq 2\varepsilon^{-3}$ and G is simple it follows from lemma 5.2.6 that $p \geq 1/2d \geq 1/2|V|$. Thus, applying lemma 5.2.5 with $\tau = \varepsilon^3/2$ and $k = \lceil \log_2(2/\varepsilon^3) \rceil$ we deduce by elementary calculations that

$$\begin{aligned} \mathbb{P}_p(u \leftrightarrow v) &\geq \exp \left[-(54)^{\lceil \log_2(2/\varepsilon^3) \rceil} \log \frac{2}{\varepsilon^3} - \lceil \log_2(2/\varepsilon^3) \rceil (54)^{\lceil \log_2(2/\varepsilon^3) \rceil} \right] \\ &\geq \exp \left[-54 \cdot (54)^{\log_2(2/\varepsilon^3)} \log \frac{2}{\varepsilon^3} - 54 \cdot (54)^{\log_2(2/\varepsilon^3)} \log_2 \frac{2}{\varepsilon^3} \right. \\ &\quad \left. - 54 \cdot (54)^{\log_2(2/\varepsilon^3)} \right] \\ &\geq \exp \left[-162 \cdot (54)^{\log_2(2/\varepsilon^3)} \log_2 \frac{2}{\varepsilon^3} \right] \end{aligned} \quad (5.2.23)$$

for every $u, v \in V$, where we used the inequality $\log(2/\varepsilon^3) \leq \log_2(2/\varepsilon^3)$ in the final inequality. We have by calculus that $x54^x \leq 64^x/(e \log(64/54))$ for every $x \geq 0$, and using that $(64 \cdot 162)/(e \log(64/54)) = 22449.65 \dots \leq 10^5$ we deduce that

$$\mathbb{P}_p(u \leftrightarrow v) \geq \exp[-10^5 \cdot \varepsilon^{-18}] \quad (5.2.24)$$

for every $u, v \in V$ as claimed. \square

5.3 Uniqueness under the sharp density property

In this section we prove supercritical uniqueness under the assumption that our infinite set $\mathcal{H} \subseteq \mathcal{F}$ satisfies the *sharp density property*, which we now introduce. This section is at the heart of the paper and contains the most significant new techniques.

Definition 5.3.1. Let $G = (V, E)$ be a finite graph and let $\Delta : (0, 1) \rightarrow (0, 1/2]$ be decreasing. Recall from Section 5.2 that for each $0 < \alpha, \delta < 1$ we define $p_c(\alpha, \delta) = p_c^G(\alpha, \delta) := \inf\{p \in [0, 1] : \mathbb{P}_p(\|K_1\| \geq \alpha) \geq \delta\}$. We say G has the Δ -sharp density property if

$$\frac{p_c(\alpha, 1 - \delta)}{p_c(\alpha, \delta)} \leq e^\delta \quad \text{for every } 0 < \alpha < 1 \text{ and } \Delta(\alpha) \leq \delta \leq 1/2.$$

Let $\mathcal{H} \subseteq \mathcal{F}$ be an infinite set. Given a \mathcal{H} -indexed family $(\Delta_G)_{G \in \mathcal{H}}$ of decreasing (i.e. non-increasing) functions $\Delta_G : (0, 1) \rightarrow (0, 1/2]$ such that $\Delta_G \rightarrow 0$ pointwise as $G \rightarrow \infty$ in \mathcal{H} , we say that \mathcal{H} has the $(\Delta_G)_{G \in \mathcal{H}}$ -sharp density property if G has the Δ_G -sharp density property for every $G \in \mathcal{H}$. We say that \mathcal{H} has the sharp density property if there is some \mathcal{H} -indexed family of decreasing functions $(\Delta_G)_{G \in \mathcal{H}}$ with $\Delta_G : (0, 1) \rightarrow (0, 1/2]$ and $\Delta_G \rightarrow 0$ pointwise as $G \rightarrow \infty$ in \mathcal{H} such that \mathcal{H} has the $(\Delta_G)_{G \in \mathcal{H}}$ -sharp density property. Equivalently, \mathcal{H} has the sharp density property if and only if

$$\lim_{G \in \mathcal{H}} \sup_{\beta \in [\alpha, 1]} \frac{p_c^G(\beta, 1 - \delta)}{p_c^G(\beta, \delta)} = 1$$

for every $0 < \alpha < 1$ and $0 < \delta \leq 1/2$.

Graphs with subalgebraic vertex degrees can straightforwardly be shown to satisfy the sharp density property using standard sharp threshold theorems [FK96; Bou+92; Tal94], all of which are proven via Fourier analysis on the hypercube. Indeed, applying these theorems in our setting leads to the following proposition, which will be used in the proof of Theorem 5.1.3 and whose proof is deferred to Section 5.3.

Proposition 5.3.2. *There exists a universal constant C such that the following holds. Let $G = (V, E)$ be a finite, simple, connected vertex-transitive graph with vertex degree d . Then G has the Δ -sharp density property with*

$$\Delta(\alpha) = \begin{cases} \frac{1}{2} \wedge C \sqrt{\frac{\log d}{\log |V|}} & \text{if } \alpha \geq (2/|V|)^{1/3} \\ \frac{1}{2} & \text{otherwise.} \end{cases}$$

In particular, if $(G_n)_{n \geq 1} = ((V_n, E_n))_{n \geq 1}$ is a sequence of finite, vertex-transitive graphs with $|V_n| \rightarrow \infty$ and with subalgebraic degrees $d_n = |V_n|^{o(1)}$ then $(G_n)_{n \geq 1}$ has the sharp density property.

A more general proposition stating that an infinite subset of \mathcal{F} has the sharp density property if and only if it does not have a molecular subsequence is proven in Section 5.4. For now we will focus on the *consequences* of the sharp density property, leaving the verification of this property to later sections.

We now state the main quantitative result of this section. When applying this theorem we will think of $\varepsilon > 0$ as a fixed constant, the number of vertices $|V|$ as being large, and $\Delta(\varepsilon)$ as being small. We have not attempted to optimise the universal constants appearing in this theorem.

Theorem 5.3.3. *Let $G = (V, E)$ be a finite, simple, connected, vertex-transitive graph, let $\varepsilon \in (0, 1]$, and let $\tau(\varepsilon) > 0$ be as in Theorem 5.2.1. If G has the Δ -sharp density property for some $\Delta : (0, 1) \rightarrow (0, 1/2]$ then*

$$\mathbb{P}_p \left(\|K_2\| \geq \lambda \left(\frac{200\Delta(\varepsilon)}{\varepsilon^3\tau(\varepsilon)} + \frac{25}{\varepsilon^2\tau(\varepsilon)|V|} \right) \right) \leq \frac{\varepsilon}{\lambda}$$

for every ε -supercritical parameter p and every $\lambda \geq 1$.

Corollary 5.3.4. *Let $\mathcal{H} \subseteq \mathcal{F}$ be an infinite set. If \mathcal{H} has the sharp density property, then \mathcal{H} has the supercritical uniqueness property.*

In this proof, we think of an event as holding *with high probability* if the probability of its complement is controlled by $|V|^{-1}$ and $\Delta(x)$ for some constant x . Similarly, we think of a real-valued random variable as being *concentrated* if the random variable lies in an interval of width controlled by $|V|^{-1}$ and $\Delta(x)$ for some constant x with high probability. Fix a parameter p that is ε -supercritical with respect to G . Our plan is as follows. First, in Section 5.3, we show that it is possible to find a parameter q with $p_c(\varepsilon, \varepsilon) \leq q \leq p$ such that $\|K_1\|$ is concentrated in a small window under \mathbb{P}_q . Then, in Section 5.3, we deduce that $\|K_2\|$ is small with high probability under \mathbb{P}_q . Finally, in Section 8.4 we introduce the notion of *sandcastles* and use this notion to prove that non-uniqueness of the giant cluster at the fixed parameter p would contradict the established properties of percolation at the well-chosen lower parameter q .

Concentration at a lower parameter

Our first step is to use the sharp density property to find another parameter q with $p_c(\varepsilon, \varepsilon) \leq q \leq p$ such that the largest cluster under \mathbb{P}_q is a giant whose density is concentrated in a small interval.

Lemma 5.3.5. *Let $G = (V, E)$ be a finite graph with the Δ -sharp density property for some Δ . Then for every $\varepsilon \in (0, 1]$ and every ε -supercritical parameter p , there exists a parameter*

$q \in (p_c(\varepsilon, \varepsilon), p)$ and a density $\alpha \geq \varepsilon$ such that

$$\mathbb{P}_q \left(\left| \|K_1\| - \alpha \right| \geq \frac{4\Delta(\varepsilon)}{\varepsilon} + \frac{1}{|V|} \right) \leq 2\Delta(\varepsilon). \quad (5.3.1)$$

Roughly speaking, the idea behind the following proof is that if we pick q such that the *median* — or any other particular quantile — of the density of the largest cluster increases slowly across a small neighbourhood of q then the sharp density property implies that the density of the giant at q must be concentrated; such a q can always be found since a bounded increasing function cannot increase rapidly everywhere.

Proof of Lemma 5.3.5. We may assume $4\Delta(\varepsilon) \leq \varepsilon$, the lemma being trivial otherwise. Consider the increasing sequence of reals q_0, q_1, \dots given by $q_j := e^{j\Delta(\varepsilon)} p_c(\varepsilon, \varepsilon)$ for each $j \geq 0$, and let k be the maximum integer such that $q_{2k} \leq p$. We start by finding a simple lower bound for k . Since p is ε -supercritical, we know $(1 - \varepsilon)^{-1} \cdot p_c(\varepsilon, \varepsilon) \leq p$ and hence that $k \geq r$ for any integer r satisfying $e^{2r\Delta(\varepsilon)} \leq (1 - \varepsilon)^{-1}$. It follows in particular that

$$k \geq \left\lfloor \frac{1}{2} \cdot \frac{\log 1/(1 - \varepsilon)}{\Delta(\varepsilon)} \right\rfloor \geq \left\lfloor \frac{1}{2} \cdot \frac{\log(1 + \varepsilon)}{\Delta(\varepsilon)} \right\rfloor \geq \left\lfloor \frac{\varepsilon}{4\Delta(\varepsilon)} \right\rfloor \geq \frac{\varepsilon}{8\Delta(\varepsilon)}, \quad (5.3.2)$$

where we used the inequality $1/(1 - x) \geq 1 + x$ for $0 \leq x < 1$ in the first inequality, the inequality $\log(1 + x) \geq x/2$ for $0 \leq x \leq 1$ in the second inequality, and the assumption $4\Delta(\varepsilon) \leq \varepsilon$ in the final inequality.

Now, for each $i \geq 0$ we define the density λ_i of K_1 under \mathbb{P}_{q_i} by

$$\lambda_i := \max\{\beta \in [0, 1] : \mathbb{P}_{q_i}(\|K_1\| \geq \beta) \geq \varepsilon\},$$

so that $\lambda_i \geq \lambda_0 = \varepsilon$ for every $i \geq 0$. Since λ_i is increasing in i we have that

$$\sum_{i=1}^k |\lambda_{2i} - \lambda_{2(i-1)}| = \sum_{i=1}^k \lambda_{2i} - \lambda_{2(i-1)} = \lambda_{2k} - \lambda_0 \leq 1,$$

and hence by the pigeonhole principle that there is some $j \in \{1, \dots, k\}$ such that

$$|\lambda_{2j} - \lambda_{2(j-1)}| = \lambda_{2j} - \lambda_{2(j-1)} \leq \frac{1}{k} \leq \frac{8\Delta(\varepsilon)}{\varepsilon},$$

where the final inequality follows from (5.3.2).

We will argue that the values

$$q = q_{2j-1} \quad \text{and} \quad \alpha = \frac{\lambda_{2j} + \lambda_{2(j-1)}}{2}$$

satisfy the conclusions of the lemma. Indeed, by definition of λ , we have that

$$\mathbb{P}_{q_{2(j-1)}} (\|K_1\| \geq \lambda_{2(j-1)}) \geq \varepsilon \quad \text{but} \quad \mathbb{P}_{q_{2j}} (\|K_1\| \geq \lambda_{2j} + |V|^{-1}) < \varepsilon. \quad (5.3.3)$$

We also have by assumption that $4\Delta(\varepsilon) \leq \varepsilon$ and $\varepsilon \leq 1/2$, so that $\Delta(\varepsilon) \leq \varepsilon \leq 1 - \Delta(\varepsilon)$ and hence by the definition of the Δ -sharp density property that

$$\max \left\{ \frac{p_c(\beta, 1 - \Delta(\varepsilon))}{p_c(\beta, \varepsilon)}, \frac{p_c(\beta, \varepsilon)}{p_c(\beta, \Delta(\varepsilon))} \right\} \leq \frac{p_c(\beta, 1 - \Delta(\varepsilon))}{p_c(\beta, \Delta(\varepsilon))} \leq e^{\Delta(\varepsilon)}$$

for every $\beta \in [\varepsilon, 1]$. Applying this inequality with the values $\beta = \lambda_{2(j-1)}$ and $\beta = (\lambda_{2j} + |V|^{-1}) \wedge 1$ and using that $e^{-\Delta(\varepsilon)} q_{2j} = q = e^{\Delta(\varepsilon)} q_{2(j-1)}$, we deduce from (5.3.3) that

$$\mathbb{P}_q (\|K_1\| \geq \lambda_{2(j-1)}) \geq 1 - \Delta(\varepsilon) \quad \text{and} \quad \mathbb{P}_q (\|K_1\| \geq \lambda_{2j} + |V|^{-1}) \leq \Delta(\varepsilon).$$

Since $|\alpha - \lambda_{2(j-1)}|$ and $|\alpha - \lambda_{2j}|$ are both bounded by $4\Delta(\varepsilon)/\varepsilon$, it follows that

$$\mathbb{P}_q \left(|\|K_1\| - \alpha| \geq \frac{4\Delta(\varepsilon)}{\varepsilon} + \frac{1}{|V|} \right) \leq 2\Delta(\varepsilon)$$

as claimed. \square

Concentration implies uniqueness

By applying Lemma 5.3.5 with our fixed parameter p , we obtain a parameter $q \in (p_c(\varepsilon, \varepsilon), 1)$ and a density $\alpha \geq \varepsilon$ that satisfy (5.3.1). We next argue that concentration of $\|K_1\|$ under \mathbb{P}_q implies uniqueness of the giant cluster under \mathbb{P}_q .

Lemma 5.3.6. *Let $G = (V, E)$ be a finite graph, let $q \in (0, 1]$, and let $\tau := \min_{u, v \in V} \mathbb{P}_q(u \leftrightarrow v)$.*

The estimate

$$\mathbb{P}_q (\|K_2\| \geq 2\delta) \leq \left(1 + \frac{1}{4\delta^2\tau}\right) \mathbb{P}_q (|\|K_1\| - \alpha| \geq \delta)$$

holds for every $\alpha, \delta > 0$.

The idea is that by the two-point connection property (and positive association), on any increasing event, we can connect K_1 and K_2 to form a new largest cluster K_1 with good probability. Thus, given that $\|K_1\|$ is concentrated, $\|K_1 \sqcup K_2\|$ must be close to $\|K_1\|$ with high probability. Equivalently, $\|K_2\| = \|K_1 \sqcup K_2\| - \|K_1\|$ must be close to zero with high probability.

Proof of Lemma 5.3.6. By the union bound,

$$\mathbb{P}_q (\|K_2\| \geq 2\delta \text{ and } \|K_1\| \geq \alpha - \delta) \geq \mathbb{P}_q (\|K_2\| \geq 2\delta) - \mathbb{P}_q (|\|K_1\| - \alpha| \geq \delta).$$

On the event that $\|K_2\| \geq 2\delta$ and $\|K_1\| \geq \alpha - \delta$, every pair of vertices u, v with $u \in K_2$ and $v \in K_1$ has $\|K_u \cup K_v\| \geq \alpha + \delta$, and there are at least $4\delta^2 |V|^2$ such pairs. So, by linearity of expectation,

$$\begin{aligned} \max_{u,v \in V} \mathbb{P}_q (\|K_u \cup K_v\| \geq \alpha + \delta) &\geq \frac{1}{|V|^2} \sum_{u,v \in V} \mathbb{P}_q (\|K_u \cup K_v\| \geq \alpha + \delta) \\ &\geq 4\delta^2 [\mathbb{P}_q (\|K_2\| \geq 2\delta) - \mathbb{P}_q (|\|K_1\| - \alpha| \geq \delta)]. \end{aligned} \quad (5.3.4)$$

When $\|K_u \cup K_v\| \geq \alpha + \delta$ and $u \leftrightarrow v$, we are guaranteed to have $\|K_1\| \geq \alpha + \delta$ and hence that $|\|K_1\| - \alpha| \geq \delta$. It follows by Harris's inequality that

$$\begin{aligned} \mathbb{P}_q (|\|K_1\| - \alpha| \geq \delta) &\geq \max_{u,v \in V} \mathbb{P}_q (\|K_u \cup K_v\| \geq \alpha + \delta) \cdot \mathbb{P}_q (u \leftrightarrow v) \\ &\geq 4\delta^2 \tau [\mathbb{P}_q (\|K_2\| \geq 2\delta) - \mathbb{P}_q (|\|K_1\| - \alpha| \geq \delta)], \end{aligned} \quad (5.3.5)$$

and the claim follows by rearranging. \square

Proof of Theorem 5.3.3 via sandcastles

So far we have obtained good control over $\|K_1\|$ and $\|K_2\|$ under \mathbb{P}_q , where q is a well-chosen parameter $p_c(\varepsilon, \varepsilon) \leq q \leq p$. We now need to convert this into an upper bound on the probability that the second largest cluster is large under \mathbb{P}_p . We do this by introducing an object we call a *sandcastle*. This is defined in terms of the canonical monotone coupling (ω_q, ω_p) of the percolation measures \mathbb{P}_q and \mathbb{P}_p with $q \leq p$ on any given graph, where each closed edge of ω_p is also closed in ω_q and each open edge of ω_p is open in ω_q with probability q/p . We write $\mathbb{P}_{q,p}$ for the joint law of this coupling. (Recall that in this coupling, when we condition on ω_p , the states of the edges in ω_q are still independent of each other.) Informally, a sandcastle is a large connected subgraph of G with the property that even knowing that the subgraph is entirely open in ω_p , there remains a good conditional probability that it contains no large cluster for ω_q . We fix this ‘good probability’ to be $1/2$ in the following definition, but we could have used any other universal constant in $(0, 1)$.

Definition 5.3.7. Let $G = (V, E)$ be a finite graph. Let $0 \leq q \leq p \leq 1$ and let $0 \leq \alpha, \beta \leq 1$. A $[(p, \beta) \rightarrow (q, \alpha)]$ -sandcastle is a connected subgraph $S \subseteq G$ such that $\|S\| \geq \beta$ and

$$\mathbb{P}_{q,p} (\|K_1(\omega_q \cap S)\| < \alpha \mid S \subseteq \omega_p) \geq \frac{1}{2}.$$

We now show that non-uniqueness of the giant cluster under \mathbb{P}_p and uniqueness of the giant cluster under \mathbb{P}_q together imply that some cluster must be a sandcastle with good probability under the measure \mathbb{P}_p .

Lemma 5.3.8. *Let $G = (V, E)$ be a finite graph. For each $0 \leq q \leq p \leq 1$, and $0 < \alpha, \beta < 1$ there exists a vertex u such that*

$$\begin{aligned} \mathbb{P}_p(K_u \text{ is a } [(p, \beta) \rightarrow (q, \alpha)]\text{-sandcastle}) \\ \geq \beta [\mathbb{P}_p(\|K_2\| \geq \beta) - 4\mathbb{P}_q(\|K_2\| \geq \alpha)]. \end{aligned}$$

Proof of Lemma 5.3.8. Consider a configuration $\nu \in \{0, 1\}^E$ in which $\|K_2\| \geq \beta$ but no clusters are $[(p, \beta) \rightarrow (q, \alpha)]$ -sandcastles. Let $A := K_1(\nu)$ and $B := K_2(\nu)$. Since $\|A\| \geq \beta$ but A is not a $[(p, \beta) \rightarrow (q, \alpha)]$ -sandcastle, we know by the definition of sandcastles that

$$\mathbb{P}_{q,p}(\|K_1(\omega_q \cap A)\| \geq \alpha \mid \omega_p = \nu) = \mathbb{P}_{q,p}(\|K_1(\omega_q \cap A)\| \geq \alpha \mid A \subseteq \omega_p) \geq \frac{1}{2}.$$

The same result holds for B . Since A and B are disjoint, the restrictions of ω_q to A and B are conditionally independent given ω_p , and hence

$$\mathbb{P}_{q,p}(\|K_1(\omega_q \cap A)\| \geq \alpha \text{ and } \|K_1(\omega_q \cap B)\| \geq \alpha \mid \omega_p = \nu) \geq \frac{1}{4}.$$

The edges in the boundary of A are all closed in ν , disconnecting A from B . Since $\omega_q \leq \omega_p$, these edges are also closed in ω_q when $\omega_p = \nu$. In particular, given that $\omega_p = \nu$, the subgraphs $K_1(\omega_q \cap A)$ and $K_1(\omega_q \cap B)$ are not connected to each other in ω_q , and hence

$$\begin{aligned} \mathbb{P}_{q,p}(\|K_2(\omega_q)\| \geq \alpha \mid \omega_p = \nu) \\ \geq \mathbb{P}_{q,p}(\|K_1(\omega_q \cap A)\| \geq \alpha \text{ and } \|K_1(\omega_q \cap B)\| \geq \alpha \mid \omega_p = \nu) \geq \frac{1}{4}. \end{aligned}$$

Letting \mathcal{E} be the event that $\|K_2(\omega_p)\| \geq \beta$ but no cluster in ω_p is a $[(p, \beta) \rightarrow (q, \alpha)]$ -sandcastle, it follows since $\nu \in \mathcal{E}$ was arbitrary that

$$\mathbb{P}_{q,p}(\|K_2(\omega_q)\| \geq \alpha \mid \omega_p \in \mathcal{E}) \geq \frac{1}{4}.$$

It follows from this and a union bound that

$$\begin{aligned} \mathbb{P}_q(\|K_2\| \geq \alpha) &\geq \mathbb{P}_{q,p}(\|K_2(\omega_q)\| \geq \alpha \mid \omega_p \in \mathcal{E}) \cdot \mathbb{P}_p(\mathcal{E}) \\ &\geq \frac{1}{4} \left(\mathbb{P}_p(\|K_2\| \geq \beta) \right. \\ &\quad \left. - \mathbb{P}_p(\text{some cluster is a } [(p, \beta) \rightarrow (q, \alpha)]\text{-sandcastle}) \right), \end{aligned}$$

which rearranges to give that

$$\begin{aligned} \mathbb{P}_p(\text{some cluster is a } [(p, \beta) \rightarrow (q, \alpha)]\text{-sandcastle}) \\ \geq \mathbb{P}_p(\|K_2\| \geq \beta) - 4\mathbb{P}_q(\|K_2\| \geq \alpha). \quad (5.3.6) \end{aligned}$$

Every cluster that is a $[(p, \beta) \rightarrow (q, \alpha)]$ -sandcastle contains at least $\beta |V|$ vertices by definition, and we deduce by linearity of expectation that

$$\begin{aligned} \max_{u \in V} \mathbb{P}_p (K_u \text{ is a } [(p, \beta) \rightarrow (q, \alpha)]\text{-sandcastle}) \\ \geq \frac{1}{|V|} \sum_{u \in V} \mathbb{P}_p (K_u \text{ is a } [(p, \beta) \rightarrow (q, \alpha)]\text{-sandcastle}) \\ \geq \beta \mathbb{P}_p (\text{some cluster is a } [(p, \beta) \rightarrow (q, \alpha)]\text{-sandcastle}) . \end{aligned}$$

The claimed inequality follows from this and (5.3.6). \square

We want to use the fact that K_u is a sandcastle with good probability to contradict the concentration of $\|K_1\|$ under \mathbb{P}_q . The rough idea is as follows: as we pass from ω_p to ω_q , with good probability this sandcastle disintegrates into only small clusters, none of which are equal to the giant cluster $K_1(\omega_q)$. Since the status of any edge that does not touch the cluster of the vertex u in ω_p remains conditionally distributed as Bernoulli percolation, this implies that there exists a large set of vertices whose complement contains, with good probability, an ω_q cluster whose density is close to the typical density of the largest cluster in the whole graph. Using Harris' inequality and uniqueness of the giant cluster in ω_q , we deduce that $\|K_1(\omega_q)\|$ is abnormally high with good probability, contradicting the concentration of the giant cluster's density under \mathbb{P}_q .

We now begin to make this argument precise. For each subgraph H of G , let \overline{H} denote the set of all edges that have at least one endpoint in the vertex set of H .

Lemma 5.3.9. Let $G = (V, E)$ be a finite, vertex-transitive graph and let $q \in (0, 1]$. The estimate

$$\begin{aligned} \mathbb{P}_q \left(\|K_1(\omega \setminus \overline{H})\| \geq \beta \right) \cdot \frac{2\beta \mathbb{P}_q (\|K_1\| \geq \beta) - \beta^2}{2 - \beta^2} \\ \leq \mathbb{P}_q \left(\|K_1\| \notin \left(\beta, \beta + \frac{\beta^2}{2} \|H\| \right) \right) + \mathbb{P}_q (\|K_2\| \geq \beta) \quad (5.3.7) \end{aligned}$$

holds for every subgraph H of G and every $0 < \beta < 1$.

Proof of lemma 5.3.9. Let X be the set of vertices that are contained in clusters with density at least β , noting that

$$\mathbb{P}_q (X \neq K_1) \leq \mathbb{P}_q (\|K_1\| \leq \beta) + \mathbb{P}_q (\|K_2\| \geq \beta) \quad (5.3.8)$$

and hence that

$$\begin{aligned} \mathbb{P}_q \left(\|X\| \geq \beta + \frac{\beta^2}{2} \|H\| \right) \leq \mathbb{P}_q \left(\|K_1\| \geq \beta + \frac{\beta^2}{2} \|H\| \right) \\ + \mathbb{P}_q (\|K_1\| \leq \beta) + \mathbb{P}_q (\|K_2\| \geq \beta) . \quad (5.3.9) \end{aligned}$$

We have by Markov's inequality applied to $\|H \setminus X\|$ that

$$\begin{aligned}\mathbb{P}_q \left(\|X \cap H\| \geq \frac{\beta^2}{2} \|H\| \right) &= 1 - \mathbb{P}_q \left(\|H \setminus X\| > \left(1 - \frac{\beta^2}{2}\right) \|H\| \right) \\ &\geq 1 - \left(1 - \frac{\beta^2}{2}\right)^{-1} \|H\|^{-1} \mathbb{E}_q \|H \setminus X\| \\ &= 1 - \left(1 - \frac{\beta^2}{2}\right)^{-1} \mathbb{P}_q(\|K_u\| < \beta),\end{aligned}\tag{5.3.10}$$

where u is an arbitrary vertex and we used vertex-transitivity in the last line. Bounding $\mathbb{P}_q(\|K_u\| \geq \beta) \geq \beta \mathbb{P}_q(\|K_1\| \geq \beta)$ we deduce that

$$\mathbb{P}_q \left(\|X \cap H\| \geq \frac{\beta^2}{2} \|H\| \right) \geq 1 - \frac{2 - 2\beta \mathbb{P}_q(\|K_1\| \geq \beta)}{2 - \beta^2}\tag{5.3.11}$$

$$= \frac{2\beta \mathbb{P}_q(\|K_1\| \geq \beta) - \beta^2}{2 - \beta^2}\tag{5.3.12}$$

and hence by Harris' inequality that

$$\begin{aligned}\mathbb{P}_q \left(\|X\| \geq \beta + \frac{\beta^2}{2} \|H\| \right) &\geq \mathbb{P}_q(\|X \setminus H\| \geq \beta) \cdot \mathbb{P}_q \left(\|X \cap H\| \geq \frac{\beta^2}{2} \|H\| \right) \\ &\geq \mathbb{P}_q(\|K_1(\omega \setminus \bar{H})\| \geq \beta) \cdot \frac{2\beta \mathbb{P}_q(\|K_1\| \geq \beta) - \beta^2}{2 - \beta^2}.\end{aligned}\tag{5.3.13}$$

The claim follows by combining (5.3.9) and (5.3.13). \square

Proof of Theorem 5.3.3. Write $\tau = \tau(\varepsilon)$ and $\Delta = \Delta(\varepsilon)$, fix an ε -supercritical parameter p , and let q and α satisfying $\alpha \geq \varepsilon$ be as in Lemma 5.3.5. Define

$$\delta = \frac{4\Delta}{\varepsilon} + \frac{1}{|V|} \quad \text{and} \quad \beta_0 = \frac{25\delta}{\varepsilon^2\tau} = \frac{200\Delta}{\varepsilon^3\tau} + \frac{25}{\varepsilon^2\tau|V|},$$

and fix some $\beta \geq \beta_0$. This value of δ is chosen so that

$$\mathbb{P}_q \left(|\|K_1\| - \alpha| \geq \delta \right) \leq 2\Delta\tag{5.3.14}$$

by Lemma 5.3.5. We will refer to $[(p, \beta) \rightarrow (q, \varepsilon/2)]$ -sandcastles simply as sandcastles for the remainder of the proof. It suffices to prove that

$$\mathbb{P}_p(\|K_2\| \geq \beta) < \frac{200\Delta}{\varepsilon^2\tau\beta} \leq \varepsilon \cdot \frac{\beta_0}{\beta},$$

so we will suppose for contradiction that the reverse inequality

$$\mathbb{P}_p(\|K_2\| \geq \beta) \geq \frac{200\Delta}{\varepsilon^2\tau\beta} \quad (5.3.15)$$

holds. Since $\|K_2\| \leq 1$, in this case we must have that $\beta_0 \leq \beta \leq 1$ and hence that $\delta \leq \varepsilon^2/25 \leq \varepsilon/2$ and $\Delta \leq 1/200$.

Since $q \geq p_c(\varepsilon, \varepsilon)$, we can apply Theorem 5.2.1 to bound the minimal connection probability $\min_{u,v} \mathbb{P}_q(u \leftrightarrow v) \geq \tau = \tau(\varepsilon)$. Thus, applying Lemma 5.3.6 yields that

$$\begin{aligned} \mathbb{P}_q\left(\|K_2\| \geq \frac{\varepsilon}{2}\right) &\leq \left(1 + \frac{4}{\varepsilon^2\tau}\right) \mathbb{P}_q\left(|\|K_1\| - \alpha| \geq \frac{\varepsilon}{2}\right) \\ &\leq \left(1 + \frac{4}{\varepsilon^2\tau}\right) \cdot 2\Delta \leq \frac{10\Delta}{\varepsilon^2\tau}, \end{aligned} \quad (5.3.16)$$

where we used the assumption $\varepsilon/2 \geq \delta$ and eq. (5.3.14) in the second inequality. Applying Lemma 5.3.8, we deduce that there exists a vertex u such that

$$\mathbb{P}_p(K_u \text{ is a sandcastle}) \geq \beta \mathbb{P}_p(\|K_2\| \geq \beta) - \frac{40\beta\Delta}{\varepsilon^2\tau} \geq \frac{\beta}{2} \mathbb{P}_p(\|K_2\| \geq \beta), \quad (5.3.17)$$

where we used the assumption (5.3.15) in the final inequality. By vertex-transitivity, this holds for *every* vertex $u \in V$. Fix a vertex $u \in V$ and let \mathcal{S}_u be the event that $K_u(\omega_p)$ is a sandcastle. Since $\omega_q \leq \omega_p$, no vertex in $K_u(\omega_p)$ is connected to a vertex of $V \setminus K_u(\omega_p)$ in ω_q , and using the fact that $\alpha - \delta \geq \varepsilon/2$ (because $\alpha \geq \varepsilon$ and $\delta \leq \varepsilon/2$), we have the inclusion of events

$$\begin{aligned} \left\{ \|K_1(\omega_q \cap K_u(\omega_p))\| < \varepsilon/2 \right\} \cap \left\{ \|K_1(\omega_q)\| \geq \alpha - \delta \right\} \\ \subseteq \left\{ \left\| K_1\left(\omega_q \setminus \overline{K_u(\omega_p)}\right) \right\| \geq \alpha - \delta \right\}. \end{aligned}$$

Taking probabilities and using the definition of sandcastles, we deduce that

$$\begin{aligned} &\mathbb{P}_{q,p}\left(\left\| K_1\left(\omega_q \setminus \overline{K_u(\omega_p)}\right) \right\| \geq \alpha - \delta \mid \mathcal{S}_u\right) \\ &\geq \mathbb{P}_{q,p}\left(\|K_1(\omega_q \cap K_u(\omega_p))\| < \frac{\varepsilon}{2} \mid \mathcal{S}_u\right) - \mathbb{P}_{q,p}\left(\|K_1(\omega_q)\| \leq \alpha - \delta \mid \mathcal{S}_u\right) \\ &\geq \frac{1}{2} - \mathbb{P}_{q,p}\left(\|K_1(\omega_q)\| \leq \alpha - \delta \mid \mathcal{S}_u\right). \end{aligned} \quad (5.3.18)$$

Using (5.3.1) and (5.3.17), we can bound the error term

$$\mathbb{P}_{q,p}\left(\|K_1(\omega_q)\| \leq \alpha - \delta \mid \mathcal{S}_u\right) \leq \frac{\mathbb{P}_q(|\|K_1\| - \alpha| \geq \delta)}{\mathbb{P}_p(K_u \text{ is a sandcastle})} \quad (5.3.19)$$

$$\leq \frac{4\Delta}{\beta \mathbb{P}_p(\|K_2\| \geq \beta)} \leq \frac{1}{4} \quad (5.3.20)$$

by the assumption that $\mathbb{P}_p(\|K_2\| \geq \beta) \geq 200\Delta/(\varepsilon^2\tau\beta) \geq 16\Delta/\beta$, so that

$$\mathbb{P}_{q,p} \left(\left\| K_1(\omega_q \setminus \overline{K_u(\omega_p)}) \right\| \geq \alpha - \delta \mid \mathcal{S}_u \right) \geq \frac{1}{4}. \quad (5.3.21)$$

Since the left hand side of (5.3.21) can be written as a weighted sum of conditional probabilities given that $K_u(\omega_p)$ is equal to a *specific* sandcastle, there must exist a sandcastle S such that

$$\begin{aligned} \mathbb{P}_{q,p} \left(\left\| K_1(\omega_q \setminus \overline{S}) \right\| \geq \alpha - \delta \mid K_u(\omega_p) = S \right) \\ = \mathbb{P}_{q,p} \left(\left\| K_1(\omega_q \setminus \overline{K_u(\omega_p)}) \right\| \geq \alpha - \delta \mid K_u(\omega_p) = S \right) \geq \frac{1}{4}. \end{aligned} \quad (5.3.22)$$

Since the event $\{K_u(\omega_p) = S\}$ depends only on the status of edges in \overline{S} , it is independent of the restriction of ω_q to $E \setminus \overline{S}$, and we deduce that

$$\mathbb{P}_q \left(\left\| K_1(\omega \setminus \overline{S}) \right\| \geq \alpha - \delta \right) = \mathbb{P}_{q,p} \left(\left\| K_1(\omega_q \setminus \overline{S}) \right\| \geq \alpha - \delta \mid K_u(\omega_p) = S \right) \geq \frac{1}{4}. \quad (5.3.23)$$

On the other hand, using that $\Delta \leq 1/200$ and hence that

$$2\mathbb{P}_q(\|K_1\| \geq (\alpha - \delta)) \geq 2(1 - 2\Delta) \geq \alpha - \delta,$$

lemma 5.3.9 implies that

$$\begin{aligned} \mathbb{P}_q \left(\left\| K_1(\omega \setminus \overline{S}) \right\| \geq \alpha - \delta \right) &\leq \frac{2 - (\alpha - \delta)^2}{2(\alpha - \delta)\mathbb{P}_q(\|K_1\| \geq \alpha - \delta) - (\alpha - \delta)^2} \\ &\cdot \left[\mathbb{P}_q \left(\|K_1\| \notin \left(\alpha - \delta, \alpha - \delta + \frac{(\alpha - \delta)^2}{2}\beta \right) \right) + \mathbb{P}_q(\|K_2\| \geq \alpha - \delta) \right], \end{aligned} \quad (5.3.24)$$

and since $\alpha - \delta \geq \varepsilon/2 \leq 1/2$ and $\beta \geq 16\varepsilon^{-2}\delta$, and

$$\mathbb{P}_q(\|K_1\| \geq \alpha - \delta) \geq 1 - 2\Delta \geq 99/100,$$

it follows that

$$\begin{aligned} \mathbb{P}_q \left(\left\| K_1(\omega \setminus \overline{S}) \right\| \geq \alpha - \delta \right) &\leq \frac{8}{4\varepsilon\mathbb{P}_q(\|K_1\| \geq \alpha - \delta) - \varepsilon^2} \\ &\cdot \left[\mathbb{P}_q(\|K_1\| - \alpha \geq \delta) + \mathbb{P}_q(\|K_2\| \geq \varepsilon/2) \right]. \end{aligned} \quad (5.3.25)$$

Applying (5.3.14) and (5.3.16) to control the two probabilities appearing here we obtain that

$$\mathbb{P}_q \left(\left\| K_1(\omega \setminus \overline{S}) \right\| \geq \alpha - \delta \right) \leq \frac{8}{4\varepsilon(1 - 2\Delta - \varepsilon/4)} \cdot \left[2\Delta + \frac{10\Delta}{\varepsilon^2\tau} \right] \leq \frac{48\Delta}{\varepsilon^3\tau}, \quad (5.3.26)$$

where we used that $\Delta \leq 1/200 \leq 1/8$ in the final inequality. The two estimates (5.3.23) and (5.3.26) contradict each other since $200\Delta/\varepsilon^3\tau \leq \beta_0 \leq 1$. \square

Subalgebraic degree graphs have the sharp density property

In this section we prove Proposition 5.3.2. As stated above, this proposition is a straightforward application of standard sharp-threshold theorems. The details of the implementation of the proof are somewhat technical but do not contain any significant new ideas. We will apply the following straightforward consequence of the results of Talagrand [Tal94]. We refer the reader to [Gri18, Chapter 4] and [ODo14] for general background on sharp threshold theorems.

Theorem 5.3.10. *Let E be a finite set and let $A \subseteq \{0, 1\}^E$ be an increasing event. Let Γ be a group acting on E and for each $e \in E$ let $\Gamma e = \{\gamma e : \gamma \in \Gamma\}$ be the orbit of e under Γ . There exists a universal constant $c > 0$ such that if A is invariant under the action of Γ on $\{0, 1\}^E$ then*

$$\frac{d}{dp} \mathbb{P}_p(A) \geq c \left[p(1-p) \log \frac{2}{p(1-p)} \right]^{-1} \mathbb{P}_p(A)(1 - \mathbb{P}_p(A)) \log \left(2 \min_{e \in E} |\Gamma e| \right)$$

for every $p \in (0, 1)$.

Proof of Theorem 5.3.10. The **influence** of an edge e with respect to A under \mathbb{P}_p is defined to be

$$I_p(A, e) := \mathbb{P}_p(\omega \cup \{e\} \in A, \omega \setminus \{e\} \notin A).$$

Russo's formula states that if A is an increasing event then

$$\frac{d}{dp} \mathbb{P}_p(A) = \sum_{e \in E} I_p(A, e) \tag{5.3.27}$$

for every $p \in [0, 1]$. It is a theorem of Talagrand [Tal94] that there exists a universal constant $0 < c \leq 1$ such that if A is increasing then

$$p(1-p) \log \left(\frac{2}{p(1-p)} \right) \sum_{e \in E} \frac{I_p(A, e)}{\log \frac{1}{p(1-p)I_p(A, e)}} \geq c \cdot \mathbb{P}_p(A)(1 - \mathbb{P}_p(A)) \tag{5.3.28}$$

and hence that

$$\begin{aligned} \sum_{e \in E} I_p(A, e) &\geq c \cdot \mathbb{P}_p(A)(1 - \mathbb{P}_p(A)) \\ &\cdot \left[p(1-p) \log \frac{2}{p(1-p)} \right]^{-1} \log \frac{1}{p(1-p) \max_e I_p(A, e)}. \end{aligned} \tag{5.3.29}$$

(Note that Talagrand states his inequality in terms of *open* pivots, so that his expression differs from ours by some factors of $1/p$.) Intuitively, this inequality implies that any event that does not depend too strongly on the status of any particular edge must have a sharp threshold, i.e., must have

probability changing rapidly from near 0 to near 1 over a short interval. Letting e maximize the influence, we have by (5.3.27) and (5.3.29) that

$$\begin{aligned} \frac{d}{dp} \mathbb{P}_p(A) &\geq \frac{1}{p(1-p)} \cdot \max \left\{ |\Gamma e| p(1-p) I_p(A, e), \right. \\ &\quad \left. c \cdot \mathbb{P}_p(A)(1 - \mathbb{P}_p(A)) \cdot \left[\log \frac{2}{p(1-p)} \right]^{-1} \log \frac{1}{p(1-p) I_p(A, e)} \right\}. \end{aligned} \quad (5.3.30)$$

Since the function $f(x) = \max\{ax, b \log 1/x\}$ attains its minimum when $\frac{1}{x} \log \frac{1}{x} = \frac{a}{b}$, it follows that

$$\begin{aligned} \frac{d}{dp} \mathbb{P}_p(A) &\geq \frac{c}{p(1-p)} \cdot \mathbb{P}_p(A)(1 - \mathbb{P}_p(A)) \\ &\quad \cdot \left[\log \frac{2}{p(1-p)} \right]^{-1} W \left(\frac{|\Gamma e| \log \frac{2}{p(1-p)}}{c \cdot \mathbb{P}_p(A)(1 - \mathbb{P}_p(A))} \right) \end{aligned} \quad (5.3.31)$$

where W is the Lambert W-function (i.e., the inverse function of xe^x). The claim follows since $\mathbb{P}_p(A)(1 - \mathbb{P}_p(A)) \leq 1$ and W is increasing and satisfies $W(x) \geq \frac{1}{2} \log x$ for every $x \geq 1$. \square

We now apply Theorem 5.3.10 to prove Proposition 5.3.2.

Proof of Proposition 5.3.2. Let $G = (V, E)$ be a finite vertex-transitive graph of degree $d \geq 2$. It follows from lemma 5.2.6 that $p_c(\alpha, \delta) \geq 1/2d$ for every $\alpha, \delta \geq \alpha_0 := (2/|V|)^{1/3}$. It suffices to show that there exists a universal constant $C \geq 1$ such that

$$\text{If } \alpha, \delta \geq \alpha_0 \quad \text{and} \quad \frac{p_c(\alpha, 1 - \delta)}{p_c(\alpha, \delta)} \geq e^\delta \quad \text{then} \quad \delta \leq \sqrt{\frac{C \log d}{\log |V|}}. \quad (5.3.32)$$

Fix $\alpha_0 \leq \alpha \leq 1$ and $\alpha_0 \leq \delta \leq 1/2$ and write $p_0 = p_c(\alpha, \delta)$ and $p_1 = p_c(\alpha, 1 - \delta)$. If $p_0 \leq p \leq p_1$ then $\mathbb{P}_p(\|K_1\| \geq \alpha)(1 - \mathbb{P}_p(\|K_1\| \geq \alpha)) \geq \delta(1 - \delta) \geq \frac{1}{2}\delta$ and it follows from Theorem 5.3.10 that there exists a universal constant $c > 0$ such that

$$1 \geq 1 - 2\delta = \int_{p_0}^{p_1} \frac{d}{dp} \mathbb{P}_p(\|K_1\| \geq \alpha) dp \quad (5.3.33)$$

$$\geq \frac{c\delta}{2} \log |V| \int_{p_0}^{p_1} \left[p(1-p) \log \frac{2}{p(1-p)} \right]^{-1} dp, \quad (5.3.34)$$

where we used that every edge has at least $|V|/2$ edges in its $\text{Aut}(G)$ orbit on any vertex-transitive graph. To estimate this integral we first use the substitution $p = \phi(x) := e^x/(e^x + 1)$, which satisfies

$dp/dx = e^x/(e^x + 1)^2 = p(1 - p)$, to write

$$\begin{aligned} \int_{p_0}^{p_1} \left[p(1-p) \log \frac{2}{p(1-p)} \right]^{-1} dp &= \int_{x_0}^{x_1} \left[\log \frac{2(e^x + 1)^2}{e^x} \right]^{-1} dx \\ &\geq \int_{x_0}^{x_1} \frac{1}{|x| + \log 8} dx, \end{aligned} \tag{5.3.35}$$

where we write $x_i = \phi^{-1}(p_i) = \log p_i/(1 - p_i)$ and use the elementary bound $(e^x + 1)^2/e^x = e^x + 2 + e^{-x} \leq 4e^{|x|}$ in the final inequality. The logarithmic derivative of $\phi(x)$ is $1/(e^x + 1)$ so that if $p_1 \geq e^\delta p_0$ then we have that

$$\int_{x_0}^{x_1} \frac{1}{e^x + 1} dx \geq \delta \quad \text{and hence that} \quad x_1 - x_0 \geq (e^{x_0} + 1)\delta.$$

It follows that if $p_1 \geq e^\delta p_0$ then

$$\int_{x_0}^{x_0+(e^{x_0}+1)\delta} \frac{1}{|x| + \log 8} dx \leq \frac{2}{c\delta \log |V|},$$

from which the claim may easily be proven via case analysis according to whether $x_0 \leq 0$ or $x_0 > 0$, noting that $x_0 \geq \phi^{-1}(1/2d) \geq -\log 2d$ since $p_0 \geq 1/2d$. In the first case we use that $|x| + \log 8 \leq |x_0| + \log 8 + \delta \leq |x_0| + 3$ for every $x_0 \leq x \leq x_0 + \delta$ to deduce that

$$\frac{\delta}{3 + \log 2d} \leq \int_{x_0}^{x_0+\delta} \frac{dx}{|x_0| + 3} \leq \frac{2}{c\delta \log |V|} \quad \text{when } x_0 \leq 0, \tag{5.3.36}$$

while in the case $x_0 > 0$ we lower bound the integral by the minimum of the integrand times the length of the interval to obtain that

$$\frac{2}{2 + \log 8} \delta \leq \frac{(e^{x_0} + 1)\delta}{x_0 + (e^{x_0} + 1)\delta + \log 8} \leq \frac{2}{c\delta \log |V|} \quad \text{when } x_0 > 0. \tag{5.3.37}$$

Putting together (5.3.36) and (5.3.37) completes the proof. \square

Proof of Theorem 5.1.3. The claim follows immediately from Theorems 5.2.1 and 5.3.3 and Proposition 5.3.2. \square

5.4 Non-molecular graphs have the sharp density property

In this section we complete the proofs of our main theorems, Theorem 5.1.2 and Theorem 8.1.1. The most important remaining step is to deduce Theorem 5.1.2 and the implication (i) \Rightarrow (ii) of Theorem 8.1.1 from Corollary 5.3.4 by proving the following proposition.

Proposition 5.4.1. *Let $\mathcal{H} \subseteq \mathcal{F}$ be an infinite set. If \mathcal{H} does not have the sharp density property, then \mathcal{H} contains a molecular subset.*

Remark 5.4.1. It follows from Theorem 8.1.1 and Corollary 5.3.4 that the converse of this proposition also holds, that is, that molecular sequences do not have the sharp density property.

Recall from Section 5.3 that a sequence of graphs is said to have the sharp density property if the emergence of a giant cluster of a given density always has a sharp threshold. As we saw in Section 5.3, it is an immediate consequence of standard sharp threshold results [FK96; Bou+92; Tal94] that this property holds whenever the graphs in question have bounded or subalgebraic vertex degrees. Indeed, these results imply that *any* increasing event depending in a sufficiently symmetric way on m i.i.d. Bernoulli- p random variables has a sharp threshold *provided that this threshold occurs around a value of p that is subalgebraically small in m* ; lemma 5.2.6 implies that the latter condition is satisfied for the event $\{\|K_1\| \geq \alpha\}$ whenever G_n has subalgebraic degrees. Unfortunately it is not true in general that every symmetric increasing event has a sharp threshold without this condition on the location of the threshold. For example, the event that the Erdős–Rényi graph contains a triangle has a coarse threshold on the scale $p = \Theta(n^{-1})$ and the event that the Erdős–Rényi graph contains a tetrahedron has a coarse threshold on the scale $p = \Theta(n^{-2/3})$ [AS16, Chapter 10.1]. Thus, to prove Proposition 5.4.1 we will need to use specific properties of the event $\{\|K_1\| \geq \alpha\}$ on non-molecular graphs.

Our proof will apply a theorem first established by Bourgain [Fri99a] and sharpened by Hatami [Hat12a], which, roughly speaking, states that any event that does *not* have a sharp threshold must be heavily influenced by the status of a small number of edges. Throughout this section, the prime in e.g. \mathbb{P}'_p will always refer to a p -derivative. We say that an event \mathcal{A} is *non-trivial* if $\mathcal{A} \neq \emptyset$ and $\mathcal{A}^c \neq \emptyset$. (The following theorem actually only requires that $\mathcal{A} \neq \emptyset$.)

Theorem 5.4.2 (Hatami 2012, Corollary 2.10). *Let $G = (V, E)$ be a finite graph, and let $\mathcal{A} \subseteq \{0, 1\}^E$ be a non-trivial increasing event. For every $p \in (0, 1/2]$ and $\epsilon > 0$, there is a set of edges $F \subseteq E$ such that $\mathbb{P}_p(\mathcal{A} \mid F \subseteq \omega) \geq 1 - \epsilon$ and*

$$|F| \leq \exp\left(10^{13} \left[p \cdot \mathbb{P}'_p(\mathcal{A})\right]^2 \mathbb{P}_p(\mathcal{A})^{-2} \epsilon^{-2}\right).$$

In particular, for each $\epsilon > 0$ there exists a constant $C(\epsilon) < \infty$ such that if $\mathbb{P}_p(\mathcal{A}) \geq \epsilon$ and $p \cdot \mathbb{P}'_p(\mathcal{A}) \leq \epsilon^{-1}$ then there is a set of edges $F \subseteq E$ such that $\mathbb{P}_p(\mathcal{A} \mid F \subseteq \omega) \geq 1 - \epsilon$ and $|F| \leq C(\epsilon)$.

The relevance of this theorem to sharp-threshold phenomena is made clear by the following elementary lemma, which shows that when $\mathbb{P}_p(\mathcal{A})$ does not have a sharp threshold there must be a good supply of parameters where $\mathbb{P}_p(A)$ is not close to 0 or 1 and $p \cdot \mathbb{P}'_p(\mathcal{A})$ is not large. Note that if $\mathcal{A} \subseteq \{0, 1\}^E$ is a non-trivial increasing event then its probability $\mathbb{P}_p(\mathcal{A})$ is a strictly increasing function of $p \in [0, 1]$ by [Gri99, Theorem 2.38] and hence defines an invertible increasing homeomorphism $[0, 1] \rightarrow [0, 1]$.

Lemma 5.4.3. *Let $G = (V, E)$ be a finite graph, and let $\mathcal{A} \subseteq \{0, 1\}^E$ be a non-trivial increasing event. Define $f : [0, 1] \rightarrow [0, 1]$ by $f(p) := \mathbb{P}_p(\mathcal{A})$. If $f^{-1}(1 - \delta) \geq (1 + \varepsilon)f^{-1}(\delta)$ for some $\varepsilon \in (0, 1]$ and $0 < \delta \leq 1/2$ then*

$$\mathcal{L}\left(\left\{p \in f^{-1}[\delta, 1 - \delta] : pf'(p) \leq \frac{4}{\varepsilon}\right\}\right) \geq \frac{1}{2}\mathcal{L}\left(f^{-1}[\delta, 1 - \delta]\right),$$

where \mathcal{L} denotes the Lebesgue measure on $[0, 1]$.

Proof of Lemma 5.4.3. Let $I := f^{-1}[\delta, 1 - \delta]$. The function f is differentiable, as it is a polynomial, and satisfies

$$\int_I pf'(p) \, dp \leq f^{-1}(1 - \delta) \int_I f'(p) \, dp = (1 - 2\delta)f^{-1}(1 - \delta) \leq f^{-1}(1 - \delta).$$

On the other hand, rearranging our hypothesis $f^{-1}(1 - \delta) \geq (1 + \varepsilon)f^{-1}(\delta)$ gives $f^{-1}(1 - \delta) \leq \left(1 + \frac{1}{\varepsilon}\right)[f^{-1}(1 - \delta) - f^{-1}(\delta)] \leq \frac{2}{\varepsilon}\mathcal{L}(I)$ and hence that

$$\frac{1}{\mathcal{L}(I)} \int_I pf'(p) \, dp \leq \frac{2}{\varepsilon}.$$

The result follows by applying Markov's inequality to the normalised Lebesgue measure on I . \square

When the edges in the set F given by Theorem 5.4.2 are open, the event \mathcal{A} occurs with conditional probability at least $1 - \varepsilon$. Since $\mathbb{P}_p(\mathcal{A} \mid F \subseteq \omega) \geq \mathbb{P}_p(\mathcal{A})$ by Harris's inequality, this result is only interesting when $\mathbb{P}_p(\mathcal{A})$ is significantly smaller than $1 - \varepsilon$. In this case, the state of F plays a decisive role in determining whether \mathcal{A} occurs in the sense that the event $\{\omega \notin \mathcal{A} \text{ but } \omega \cup F \in \mathcal{A}\}$ occurs with good probability. This motivates our definition of a subgraph H that *activates* an event \mathcal{A} , a generalisation of being a closed pivotal edge.

Definition 5.4.4. Let $G = (V, E)$ be a finite graph, let $H \subseteq G$ be a subgraph, and let $\mathcal{A} \subseteq \{0, 1\}^E$ be an increasing event. We say H activates the event \mathcal{A} in a configuration ω if $\omega \notin \mathcal{A}$ but $\omega \cup H \in \mathcal{A}$. For every density α , we simply say H activates α to mean H activates the event $\{\|K_1\| \geq \alpha\}$, and we label this event $\text{Act}_\alpha(H)$.

Corollary 5.4.5. *For every $\delta > 0$ there exists $\varepsilon > 0$ such that if $G = (V, E)$ is a finite, simple, vertex-transitive graph and $\alpha, p \in (0, 1]$ are such that $\mathbb{P}_p(\|K_1\| \geq \alpha) \in [\delta, 1 - \delta]$ and $p \cdot \mathbb{P}'_p(\|K_1\| \geq \alpha) \leq \delta^{-1}$ then there exists a subgraph H of G such that $|E(H)| \leq \varepsilon^{-1}$ and $\mathbb{P}_p(\text{Act}_\alpha(H)) \geq \varepsilon$.*

Proof of Corollary 5.4.5. The case $p \leq 1/2$ follows trivially from Theorem 5.4.2. Now assume $p \geq 1/2$. Since $\{\|K_1\| \geq \alpha\}$ is invariant under automorphisms of G and G is vertex-transitive we can apply Theorem 5.3.10 to obtain that there exists a universal constant $c > 0$ such that

$$\mathbb{P}'_p(\|K_1\| \geq \alpha) \geq c \cdot \mathbb{P}_p(\|K_1\| \geq \alpha) \mathbb{P}_p(\|K_1\| < \alpha) \log |V|.$$

Plugging in our assumed bounds on p , $p \cdot \mathbb{P}'_p(\|K_1\| \geq \alpha)$, and $\mathbb{P}_p(\|K_1\| \geq \alpha)$ gives $|V| \leq e^{2/c\delta^3}$, in which case $|E| \leq e^{4/c\delta^3}$ and the result holds trivially with $\varepsilon = \delta \wedge e^{-4/c\delta^3}$. \square

When the event $\{\|K_1\| \geq \alpha\}$ has a *coarse* (i.e., not sharp) threshold on $G = (V, E)$, this lemma gives us a subgraph $H \subseteq G$ that has a good probability of activating α . When H has only a single edge e , we can use this in the other direction to establish a *lower bound* on the sharpness of the phase transition. Indeed, e activates α if and only if e is closed and pivotal for the event $\{\|K_1\| \geq \alpha\}$. So, by Russo's formula and lemma 5.2.6 we have that

$$\begin{aligned} p \cdot \mathbb{P}'_p(\|K_1\| \geq \alpha) p \sum_{f \in E} \mathbb{P}_p(f \text{ is pivotal for } \{\|K_1\| \geq \alpha\}) \\ \geq \frac{|\text{Orb}(e)|}{\deg G} \mathbb{P}_p(\text{Act}_\alpha(e)), \end{aligned}$$

where $\text{Orb}(e)$ denotes the orbit of the edge e under the action of the automorphism group $\text{Aut } G$. Contrasting this with the assumed *upper bound* on the sharpness of the phase transition with which we started, we can extract information about the underlying graph G . For example, we immediately deduce that G is dense, and we are only one step away from concluding that G is part of a molecular sequence. Our main challenge is to reduce to this case, that is to say, to show that we can take H to be a subgraph with a single edge.

During the proof we will want to apply a *sprinkling* argument, where by slightly increasing the parameter p we can make strict subsets of an activator H become activators. One difficulty is that as we increase the percolation parameter, we may form a giant cluster with density at least β , in which case *nothing* can be an activator. As such, we must carefully choose the amount that we sprinkle by at each step. To this end, we will work only with a special sequence of such values constructed by the following lemma.

For the remainder of this section, given a finite, connected, vertex-transitive graph $G = (V, E)$, we write $d = \deg(G)$ and $\alpha_0 = (2/|V|)^{1/3}$. Given $\beta > 0$ and $\delta \in (0, 1/2]$ we write $I = I(G, \beta, \delta) = [p_c(\beta, \delta), p_c(\beta, 1 - \delta)]$ and $Q = Q(G, \beta, \delta) = \{p \in I : p \cdot \mathbb{P}'_p(\|K_1\| \geq \beta) \leq \frac{4}{\delta}\}$.

Lemma 5.4.6 (A good sequence for sprinkling). *Let $G = (V, E)$ be a finite, connected, vertex-transitive graph, let $\delta > 0$, and let $0 < \beta \leq 1$. There exists a sequence $(p_n)_{n \geq 1}$ in Q such that*

$$p_{n+1} - p_n \geq 3^{-(n+1)} \mathcal{L}(Q) \quad \text{and} \quad \mathbb{P}_{p_{n+1}}(\|K_1\| \geq \beta) - \mathbb{P}_{p_n}(\|K_1\| \geq \beta) \leq 2^{-n}$$

for every $n \geq 0$.

Proof of Lemma 5.4.6. We will prove more generally that if $X \subseteq [0, 1]$ is a non-empty closed set and $f : X \rightarrow [0, 1]$ is an increasing (but not necessarily continuous) function then there exists a sequence $(x_n)_{n \geq 1}$ in X such that $x_{n+1} - x_n \geq 3^{-(n+1)} \mathcal{L}(X)$ and $f(x_{n+1}) - f(x_n) \leq 2^{-n}$ for every $n \geq 0$. The claim is trivial if $\mathcal{L}(X) = 0$ since we can take x_n constant in this case, so we may assume that $\mathcal{L}(X) > 0$. First consider the case $X = [0, 1]$. Let $x_0 := 0$ and define $(x_n)_{n \geq 1}$ recursively as follows. Assume we have defined x_n for some $n \geq 0$. Set $x_{n,i} := x_n + i3^{-(n+1)}$ for each $i \in \{1, 2, 3\}$ and define

$$x_{n+1} := \begin{cases} x_{n,1} & \text{if } f(x_{n,2}) - f(x_{n,1}) \leq f(x_{n,3}) - f(x_{n,2}), \\ x_{n,2} & \text{otherwise.} \end{cases}$$

For each $n \geq 0$ write $y_n := x_{n,3}$ so that $x_n \leq y_m$ for every $n \geq m \geq 0$. It follows by induction that $x_n \in [0, 1]$, $x_{n+1} - x_n \geq 3^{-(n+1)}$, and $f(y_n) - f(x_n) \leq 2^{-n}$ for every $n \geq 0$, and the claim follows since $f(x_{n+1}) \leq f(y_n)$ for every $n \geq 0$.

Now let $X \subseteq [0, 1]$ be an arbitrary closed set with $\mathcal{L}(X) > 0$. Since X is closed, we can define an increasing function $\phi : [0, 1] \rightarrow X$ such that

$$\mathcal{L}([0, \phi(x)] \cap X) = x\mathcal{L}(X)$$

for all $x \in [0, 1]$. Construct a sequence $(x_n)_{n \geq 1}$ by the above procedure but for the function $f \circ \phi$ instead of f . Then the sequence $(\phi(x_n))_{n \geq 1}$ has the properties that $\phi(x_n) \in X$, $\phi(x_{n+1}) - \phi(x_n) \geq (x_{n+1} - x_n)\mathcal{L}(X) \geq 3^{-(n+1)}\mathcal{L}(X)$, and $f(\phi(x_{n+1})) - f(\phi(x_n)) \leq 2^{-n}$ for every $n \geq 0$ as required. \square

We now state our key technical sprinkling proposition.

Proposition 5.4.7 (Reducing to a single edge by sprinkling). *Let $G = (V, E)$ be a finite, simple, connected, vertex-transitive graph, let $\alpha_0 \leq \delta \leq 1/2$ and $\alpha_0 \leq \alpha \leq \beta \leq 1$, and suppose that*

$p_c(\beta, 1 - \delta) > e^\delta p_c(\beta, \delta)$. Let $(p_n)_{n \geq 1}$ be as in Lemma 5.4.6. For each $\varepsilon > 0$ there exists $N = N(\alpha, \delta, \varepsilon)$ such that for each $n \geq N$ there exists $\eta_n = \eta_n(\alpha, \delta, \varepsilon) > 0$ such that if H is a subgraph of G with $|E(H)| \leq \varepsilon^{-1}$ and $\mathbb{P}_{p_n}(\text{Act}_\beta(H)) \geq \varepsilon$ then there exists $e \in E(H)$ such that $\mathbb{P}_{p_m}(\text{Act}_\beta(e)) \geq \eta_n$ for every $m \geq n + N$.

Note that the choice of edge $e \in E(H)$ may depend on the choice of $n \geq N$ and that the constants N and η_n are independent of G and β .

We now show how Proposition 5.4.1 follows from Corollary 5.4.5 and Proposition 5.4.7, deferring the proof of Proposition 5.4.7 to Section 5.4.

Proof of Proposition 5.4.1 given Proposition 5.4.7. Let $G = (V, E)$ be a finite, simple, connected, vertex-transitive graph of degree d and let $\alpha_0 = (2/|V|)^{1/3}$. Let $\alpha_0 \leq \alpha \leq \beta \leq 1$, let $\alpha_0 \leq \delta \leq 1/2$, and suppose that

$$p_c(\beta, 1 - \delta) > e^\delta p_c(\beta, \delta).$$

It suffices to prove that there exist positive constants $c = c(\alpha, \delta)$ and $C = C(\alpha, \delta)$ such that if $|V| \geq C$ then $\deg G \geq c|V|$ and there exists an automorphism-invariant set of edges F with $|F| \leq C|V|$ such $G \setminus F$ has at most C connected components.

Let $(p_n)_{n \geq 1}$ be as in Lemma 5.4.6. It follows from Corollary 5.4.5 that there exists $\eta_1 = \eta_1(\delta)$ such that for each $p \in Q$ there exists a subgraph H of G such that $|E(H)| \leq \eta_1^{-1}$ and $\mathbb{P}_p(\text{Act}_\beta(H)) \geq \eta_1$. Applying this fact with $p = p_n$ for an appropriately large constant n , it follows from Proposition 5.4.7 that there exist positive constants $\eta_2 = \eta_2(\alpha, \delta)$ and $N_1 = N_1(\alpha, \delta)$ and an edge $e_0 \in E$ such that

$$\mathbb{P}_{p_n}(\text{Act}_\beta(e_0)) \geq \eta_2$$

for every $n \geq N_1$. Since the edge e_0 activates β if and only if e_0 is closed and pivotal for the event $\{\|K_1\| \geq \beta\}$, we have by Russo's formula that

$$\begin{aligned} \mathbb{P}'_{p_n}(\|K_1\| \geq \beta) &= \sum_{e \in E} \mathbb{P}_{p_n}(e \text{ is pivotal for } \{\|K_1\| \geq \beta\}) \\ &\geq |\text{Orb}(e_0)| \mathbb{P}_{p_n}(\text{Act}_\beta(e_0)) \geq \eta_2 |\text{Orb}(e_0)| \end{aligned}$$

for every $n \geq N_1$. Since $p_{N_1} \geq p_c(\beta, \delta)$ and $\alpha, \delta \geq \alpha_0$ it follows from lemma 5.2.6 that $p_{N_1} \geq 1/2d$ and hence that

$$\eta_2 \frac{|\text{Orb}(e_0)|}{2d} \leq p_{N_1} \cdot \mathbb{P}'_{p_{N_1}}(\|K_1\| \geq \beta) \leq \frac{4}{\delta}$$

for every $n \geq N_1$, where the upper bound follows since $p_n \in Q$ for every $n \geq 1$. Since G is vertex-transitive $|\text{Orb}(e_0)| \geq \frac{1}{2}|V|$ and it follows that the constant $c_2 = c_2(\alpha, \delta) = \delta\eta_2/16$ satisfies

$$d \geq c_2|V|.$$

This establishes the desired density of G .

All that remains is to find an $\text{Aut } G$ -invariant linear-sized set of edges $F \subseteq E$ that disconnects G into a bounded number of components. Let $C_1 = 6/c_2$, let $\eta_3 = \eta_2 \wedge C_1^{-1}$, and let $N_2 = N(\alpha, \delta, \eta_3)$ and $\eta_4 = \eta_{N_1 \vee N_2}(\alpha, \delta, \eta_3)$ be as in Proposition 5.4.7. We claim that if we set $k = N_1 \vee N_2 + N_2$ then there exists a constant $C_2 = C_2(\alpha, \delta)$ such that the set

$$F := \{e \in E : \mathbb{P}_{p_k}(\text{Act}_\beta(e)) \geq \eta_4\}$$

has the desired properties when $|V| \geq C_2$. This set F is clearly $\text{Aut } G$ -invariant. Moreover, by Russo's formula, since $p_k \cdot \mathbb{P}'_{p_k}(\|K_1\| \geq \beta) \leq \frac{4}{\delta}$ and $p_k \geq 1/2d \geq 1/2|V|$, we have that

$$\begin{aligned} |F| &\leq \frac{1}{\eta_4} \sum_{e \in E} \mathbb{P}_{p_k}(e \text{ is pivotal for } \{\|K_1\| \geq \beta\}) \\ &= \frac{1}{\eta_4} \mathbb{P}'_{p_k}(\|K_1\| \geq \beta) \leq \frac{8}{\delta\eta_4}|V|, \end{aligned}$$

so that $|F|$ is at most linear in $|V|$. Since $|F| \leq \frac{8}{\delta\eta_4}|V|$ and $d = \deg G \geq c_2|V|$, we have that

$$\deg(G \setminus F) \geq c_2|V| - \frac{16}{\delta\eta_4}.$$

Thus, there exists a constant $C_2 = C_2(\alpha, \delta)$ such that if $|V| \geq C_2$ then

$$\deg(G \setminus F) \geq \frac{c_2}{2}|V|.$$

It now suffices to prove that $G \setminus F$ is not connected when $|V| \geq C_2$. Indeed, once this is shown it follows automatically that $G \setminus F$ has a bounded number of components since if $G \setminus F$ has m components then $\frac{c_2}{4}|V|^2 \leq |E| \leq m(|V|/m)^2 + |F|$, and hence there exists a constant $C_4 = C_4(\alpha, \delta)$ such that $m \leq 1 + 4/c_2$ when $|V| \geq C_4$.

Suppose for contradiction that $|V| \geq C_2$ and that $G \setminus F$ is connected. Since $G \setminus F$ is connected and $\deg(G \setminus F) \geq \frac{c_2}{2}|V|$, it follows from Lemma 5.2.3 that $\text{diam}(G \setminus F) \leq 6/c_2 = C_1$. Let P be a path of length at most C_1 connecting the endpoints of e_0 in $G \setminus F$. If e_0 activates β then so does the set P , and it follows that

$$|P| \leq C_1 \leq \eta_3^{-1} \quad \text{and} \quad \mathbb{P}_{p_{N_1 \vee N_2}}(\text{Act}_\beta(P)) \geq \mathbb{P}_{p_{N_1 \vee N_2}}(\text{Act}_\beta(e_0)) \geq \eta_2 \geq \eta_3.$$

As such, it follows from the definitions of the quantities $N_2 = N(\alpha, \delta, \eta_3)$ and $\eta_4 = \eta_{N_1 \vee N_2}(\alpha, \delta, \eta_3)$ as in Proposition 5.4.7 that there exists an edge $e_1 \in P$ such that $\mathbb{P}_{p_k}(\text{Act}_\beta(e_1)) \geq \eta_4$. This implies that $e_1 \in F$, a contradiction. \square

Proof of Proposition 5.4.7

In this section we complete the proof of Proposition 5.4.1 by proving Proposition 5.4.7. The proof will proceed inductively, showing that — by changing to a different value of $p \in Q$ if necessary — we can reduce the number of edges in the subgraph H given by Corollary 5.4.5 while keeping $\mathbb{P}_p(\text{Act}_\alpha(H))$ bounded away from zero. More precisely, we will deduce Proposition 5.4.7 as an inductive consequence of the following lemma.

Lemma 5.4.8 (Removing one edge by sprinkling). *Let $G = (V, E)$ be a finite, simple, connected, vertex-transitive graph, let $\alpha_0 \leq \delta \leq 1/2$ and $\alpha_0 \leq \alpha \leq \beta \leq 1$, and suppose that $p_c(\beta, 1 - \delta) > e^\delta p_c(\beta, \delta)$. Let $(p_n)_{n \geq 1}$ be as in Lemma 5.4.6. For each $\varepsilon > 0$ there exists $N = N(\alpha, \delta, \varepsilon)$ such that if $n \geq N$ then there exists $\eta_n = \eta_n(\alpha, \delta, \varepsilon) > 0$ such that if H is a subgraph of G with $|E(H)| \leq \varepsilon^{-1}$ and $\mathbb{P}_{p_n}(\text{Act}_\beta(H)) \geq \varepsilon$ then there exists a subgraph H' of H with $|E(H')| \leq \max\{1, |E(H)| - 1\}$ such that $\mathbb{P}_{p_m}(\text{Act}_\beta(H')) \geq \eta_n$ for every $m > n$.*

Note that the choice of subgraph H' of H may depend on the choice of $n \geq N$ and that the constants N and η_n are independent of G and β . Also, note that N and η_n here are not the same N and η_n as in the statement of Proposition 5.4.7.

Proof of Proposition 5.4.7 given Lemma 5.4.8. Applying Lemma 5.4.8 iteratively $\lfloor 1/\varepsilon \rfloor$ times yields the claim. \square

For the remainder of this section we fix a finite, simple, connected, vertex-transitive graph $G = (V, E)$ of degree d , fix $\alpha_0 \leq \delta \leq 1/2$ and $\alpha_0 \leq \alpha \leq \beta \leq 1$ such that $p_c(\beta, 1 - \delta) > e^\delta p_c(\beta, \delta)$, and let $(p_n)_{n \geq 1}$ be as in Lemma 5.4.6. We also continue to write $I = I(G, \beta, \delta) = [p_c(\beta, \delta), p_c(\beta, 1 - \delta)]$ and $Q = Q(G, \beta, \delta) = \{p \in I : p \cdot \mathbb{P}'_p(\|K_1\| \geq \beta) \leq \frac{4}{\delta}\}$.

When G has bounded vertex degrees, and hence critical parameters are bounded away from zero, Lemma 5.4.8 could be proven easily by (classical) insertion-tolerance. The problem is rather more delicate in general. The idea is to use vertex-transitivity of G to find many copies of H that each activate β simultaneously, then *sprinkle*, i.e. open a small number of additional edges by slightly increasing the percolation parameter, and argue that, after sprinkling, many of the copies of H have strict subgraphs that are activators. To ensure that sprinkling reduces the number of edges

necessary to activate β in a positive proportion of the copies of H , we need these copies to be well-connected to each other in the open subgraph. We guarantee this by sticking a large cluster to each endpoint of an edge in H , using the next lemma.

Lemma 5.4.9. *For every $\varepsilon > 0$ there exists $\eta = \eta(\alpha, \delta, \varepsilon) > 0$ such that if $p \in Q$ and H is a subgraph of G with $|E(H)| \leq \frac{1}{\varepsilon}$ and $\mathbb{P}_p(\text{Act}_\beta(H)) \geq \varepsilon$ then there is an edge $e \in E(H)$ with endpoints u and v such that*

$$\mathbb{P}_p(\text{Act}_\beta(H) \cap \{\|K_u\| \geq \eta\} \cap \{|K_v| \geq \eta d\}) \geq \eta.$$

The proof of this lemma uses the quantitative insertion-tolerance estimate of Proposition 6.6.4 together with the following theorem of the second author [Hut21b, Theorem 2.2], which guarantees that the size of the largest intersection of a cluster with a fixed set of vertices is always of the same order as its mean with high probability. We state a special case of the theorem that is adequate for our purposes.

Theorem 5.4.10 (Universal Tightness). *There exist universal constants $C, c > 0$ such that the following holds. Let $G = (V, E)$ be a countable, locally finite graph, let $\Lambda \subseteq V$ be a finite non-empty set of vertices, and let $p \in [0, 1]$ be a parameter. Set $|M| := \max\{|K_v \cap \Lambda| : v \in V\}$. Then*

$$\mathbb{P}_p(|M| \geq \alpha \mathbb{E}_p |M|) \leq C e^{-c\alpha} \quad \text{and} \quad \mathbb{P}_p(|M| \leq \varepsilon \mathbb{E}_p |M|) \leq C \varepsilon$$

for every $\alpha \geq 1$ and $0 < \varepsilon \leq 1$.

Proof of Lemma 5.4.9. We may assume that $\varepsilon \leq \delta$. We may also assume that H has no isolated points, so that we have the bound $|V(H)| \leq 2|E(H)| \leq \frac{2}{\varepsilon}$. When H activates β we must have that $\|\bigcup_{u \in V(H)} K_u\| \geq \beta$ and hence by the pigeonhole principle that there exists $u \in V(H)$ such that

$$\|K_u\| \geq \frac{\beta}{|V(H)|} \geq \frac{\alpha \varepsilon}{2}.$$

It follows that there exists a *fixed* vertex $u \in V(H)$ such that

$$\mathbb{P}_p\left(\text{Act}_\beta(H) \cap \left\{\|K_u\| \geq \frac{\alpha \varepsilon}{2}\right\}\right) \geq \frac{1}{|V(H)|} \mathbb{P}_p(\text{Act}_\beta(H)) \geq \frac{\varepsilon}{|V(H)|} \geq \frac{\varepsilon^2}{2}. \quad (5.4.1)$$

Since H has no isolated points, u is the endpoint of some edge $e \in E(H)$. Let v be the other endpoint of e , let N be the set of neighbours of v in G , and let X be an ω -connected subset of N

(i.e. a subset of N that is contained in a single ω -cluster) of maximum size. Since $p \geq p_c(\alpha, \delta)$ and G is vertex-transitive,

$$\mathbb{E}_p |X| \geq \mathbb{E}_p [|K_1 \cap N|] = |N| \mathbb{E}_p \|K_1\| \geq \alpha \delta d. \quad (5.4.2)$$

Applying Theorem 5.4.10 it follows that there exists a positive constant $c_1 = c_1(\alpha, \delta, \varepsilon)$ such that $\mathbb{P}_p (|X| \leq c_1 d) \leq \frac{\varepsilon^2}{4}$ and hence by a union bound that

$$\mathbb{P}_p \left(\text{Act}_\beta(H) \cap \left\{ \|K_u\| \geq \frac{\alpha \varepsilon}{2} \right\} \cap \{|X| \geq c_1 d\} \right) \geq \frac{\varepsilon^2}{4}. \quad (5.4.3)$$

To obtain a similar bound with $|K_v|$ in place of $|X|$, we use insertion tolerance to open an edge connecting v to X , which forces $|K_v| \geq |X|$. Unfortunately we cannot argue this way directly since opening this edge may produce a cluster with density at least β , in which case $\text{Act}_\beta(H)$ would no longer hold. To avoid this issue, we first claim that there exists a constant $C_1 = C_1(\alpha, \delta, \varepsilon)$ such that if $|V| \geq C_1$ then

$$\mathbb{P}_p \left(\text{Act}_\beta(H) \cap \left\{ \|K_u\| \geq \frac{\alpha \varepsilon}{2} \right\} \cap \{|X| \geq c_1 d\} \cap \text{Act}_\beta(vX)^c \right) \geq \frac{\varepsilon^2}{8}, \quad (5.4.4)$$

where vX denotes the set of edges with one endpoint equal to v and the other in X . (Note that since X is ω -connected, each edge in vX individually activates β if and only if the entire set vX activates β .) Indeed, suppose that (5.4.4) does not hold. We have by (5.4.3) and a union bound that

$$\mathbb{P}_p (|X| \geq c_1 d) \cap \text{Act}_\beta(vX) \geq \frac{\varepsilon^2}{8}$$

and hence that

$$\sum_{e \in E: e \ni v} \mathbb{P}_p (\text{Act}_\beta(e)) \geq \mathbb{E}_p [|X| \mathbb{1}_{\text{Act}_\beta(vX)}] \geq \frac{c_1 \varepsilon^2}{8} d.$$

Applying Russo's formula and using that G is vertex-transitive, it follows that

$$\mathbb{P}'_p (\|K_1\| \geq \beta) \geq \frac{c_1 \varepsilon^2}{16} d |V|.$$

Since $p \geq p_c(\alpha_0, \alpha_0)$ we also have that $p \geq 1/2d$ by lemma 5.2.6 and hence that

$$p \cdot \mathbb{P}'_p (\|K_1\| \geq \beta) \geq \frac{c_1 \varepsilon^2}{32} |V|.$$

Since $\varepsilon \leq \delta$, this contradicts the hypothesis that $p \in Q$ whenever $|V| \geq C_1 := 128/(c_1 \varepsilon^3)$, completing the proof of the claim.

There are finitely many graphs with $|V| \leq C_1$, hence the lemma holds trivially in this case and we may assume without loss of generality that (5.4.4) holds. Since we have that $p \geq 1/2d$ as above, we can use insertion tolerance Proposition 6.6.4 to open an edge in vX in the event appearing on the left hand side of (5.4.4), giving that

$$\mathbb{P}_p \left(\text{Act}_\beta(H) \cap \left\{ \|K_u\| \geq \frac{\alpha\epsilon}{2} \right\} \cap \{|K_v| \geq c_1 d\} \right) \geq \eta,$$

for some $\eta = \eta(\alpha, \delta, \epsilon) > 0$ as claimed. \square

Once we have many copies of H that activate β and have large clusters stuck to the copies of u and v , we use the following easy fact about equivalence relations to deduce that the copies of e are well-connected to each other. More precisely, we will use this lemma to show that we can find many large disjoint sets of copies of the edge e in which any two copies $\gamma_1(e)$ and $\gamma_2(e)$ of e are ω -connected by paths $\gamma_1(u) \leftrightarrow \gamma_2(u)$ and $\gamma_1(v) \leftrightarrow \gamma_2(v)$.

Lemma 5.4.11. *Let X be a non-empty finite set, let \sim be an equivalence relation on X , and let $Y \subset X$ be a non-empty subset of X . For each $x \in X$ write $[x]$ for the equivalence class of x under \sim . If $\min_{y \in Y} |[y]| \geq \frac{2|X|}{|Y|}$ then there exists a collection $(Z_i)_{i \in I}$ of disjoint subsets of Y such that each Z_i is contained in an equivalence class of \sim and that satisfies the inequalities*

$$|Z_i| \geq \frac{|Y|}{2|X|} \min_{y \in Y} |[y]| \quad \text{and} \quad |I| \geq \frac{|X|}{4 \min_{y \in Y} |[y]|}.$$

Proof of Lemma 5.4.11. Write $m = \min_{y \in Y} |[y]|$ and define

$$Y_- := \left\{ y \in Y : |[y] \cap Y| \leq \frac{|Y|}{2|X|} |[y]| \right\} \quad \text{and} \quad Y_+ := Y \setminus Y_-.$$

Observe that if \mathcal{C} denotes the set of equivalence classes of \sim then

$$\frac{2|X|}{|Y|} |Y_-| \leq \sum_{y \in Y} \frac{|[y]|}{|[y] \cap Y|} = \sum_{C \in \mathcal{C}} |C| \mathbb{1}(C \cap Y \neq \emptyset) \leq |X|,$$

so that $|Y_-| \leq \frac{1}{2}|Y|$ and $|Y_+| \geq \frac{1}{2}|Y|$. Let $(Z_i)_{i \in I}$ be a maximal collection of disjoint subsets of Y_+ such that each Z_i is contained in an equivalence class of \sim and $|Z_i| = \left\lceil \frac{|Y|}{2|X|} m \right\rceil$ for each $i \in I$. Every element $y \in Y_+$ has

$$|[y] \cap Y| \geq \left\lceil \frac{|Y|}{2|X|} |[y]| \right\rceil \geq \left\lceil \frac{|Y|}{2|X|} m \right\rceil.$$

So, by maximality, the union $\bigcup_{i \in I} Z_i$ contains at least half the elements in $[y] \cap Y$ for every $y \in Y_+$. So the union $\bigcup_{i \in I} Z_i$ contains at least half the elements in Y_+ , and we deduce that

$$|I| \geq \frac{|Y|/4}{\left\lceil \frac{|Y|}{2|X|} m \right\rceil} \geq \frac{|X|}{4m},$$

where we used the hypothesis $m \geq \frac{2|X|}{|Y|}$ in the final inequality. \square

We are now ready to complete the proof of Lemma 5.4.8. The final step to reduce the number of edges in H is to open a small number of these well-connected copies of e by slightly increasing the percolation parameter. When H activates β and $u \leftrightarrow v$, $H \setminus \{e\}$ also activates β . Since the copies of e are well-connected, opening this small number of edges is actually sufficient to ensure a positive proportion of the copies of $H \setminus \{e\}$ activate β . By linearity of expectation, we conclude that at this higher parameter the set $H \setminus \{e\}$ activates β with good probability, as required.

Proof of Lemma 5.4.8. Since $\alpha, \delta \geq \alpha_0$ we have by Lemma 5.4.3 and Lemma 5.2.6 that

$$\mathcal{L}(Q) \geq \frac{1}{2} \mathcal{L}(I) = \frac{1}{2} (p_c(\beta, 1 - \delta) - p_c(\beta, \delta)) \geq \frac{1}{2} (e^\delta - 1) p_c(\beta, \delta) \geq \frac{\delta}{4d},$$

where we used that $e^\delta - 1 \geq \delta$ in the final inequality. Fix $\varepsilon > 0$ and $n \geq 1$, and suppose that H is a finite subgraph of G with $|E(H)| \leq \varepsilon^{-1}$ such that $\mathbb{P}_{p_n}(\text{Act}_\beta(H)) \geq \varepsilon$. By Lemma 5.4.9, we can find $\varepsilon_1 = \varepsilon_1(\alpha, \delta, \varepsilon)$ and an edge $e \in E(H)$ with endpoints u and v such that the event

$$\mathcal{A} := \text{Act}_\beta(H) \cap \{\|K_u\| \geq \varepsilon_1\} \cap \{|K_v| \geq \varepsilon_1 d\} \quad \text{satisfies } \mathbb{P}_{p_n}(\mathcal{A}) \geq \varepsilon_1.$$

For each $x \in V$, pick an automorphism $\phi_x \in \text{Aut } G$ such that $\phi_x(v) = x$, so that $\phi_x(u)$ is a neighbour of x for each $x \in V$. These maps exist because G is vertex-transitive. Define $X := \{x \in V : \phi_x^{-1}(\omega) \in \mathcal{A}\}$. We have by linearity of expectation that $\mathbb{E}_{p_n} \|X\| = \mathbb{P}_{p_n}(\mathcal{A}) \geq \varepsilon_1$ and hence by Markov's inequality that $\mathbb{P}_{p_n}(\|X\| \geq \frac{1}{2}\varepsilon_1) \geq \frac{1}{2}\varepsilon_1$. We will now prove that we can take $N = \lceil \log_2(8/\varepsilon_1) \rceil$, so assume from now on that $n \geq \log_2(8/\varepsilon_1)$.

Write $\mathcal{B} := \{\|X\| \geq \frac{1}{2}\varepsilon_1\}$ and consider an arbitrary configuration $\omega \in \mathcal{B}$. Every $\phi_x(u)$ with $x \in X$ has $\|K_{\phi_x(u)}\| \geq \varepsilon_1$. So, by the pigeonhole principle, we can find a subset $Y \subseteq X$ with $\|Y\| \geq \frac{1}{2}\varepsilon_1^2$ such that $\{\phi_x(u) : x \in Y\}$ is contained in a single cluster of ω . By definition of X and \mathcal{A} , every vertex $y \in Y$ has $|K_y| \geq \varepsilon_1 d$.

We next claim that there exists a constant $C_1 = C_1(\alpha, \delta, \varepsilon)$ such that if $|V| \geq C_1$ then $\varepsilon_1 d |Y| \geq 4|V|$. Indeed, since $\|Y\| \geq \frac{1}{2}\varepsilon_1^2$, if this inequality does not hold then we must have that $d \leq 8\varepsilon_1^{-3}$. In this

case, since $p_n \in Q$ and $p_n \geq 1/2d$, and since every edge has $\text{Aut } G$ -orbit of size at least $|V|/2$ by transitivity, it follows from Theorem 5.3.10 that there exists a constant $c_1 = c_1(\alpha, \delta, \varepsilon)$ such that

$$\frac{4}{\delta} \geq p_n \cdot \mathbb{P}'_{p_n}(\|K_1\| \geq \beta) \geq c_1 \log |V|,$$

which rearranges to give the claim. Since graphs with $|V| < C_1$ can be handled trivially, we may assume throughout the rest of the proof that $|V| \geq C_1$ and hence that $\varepsilon_1 d |Y| \geq 4|V|$.

By construction, every vertex $y \in Y$ has $|K_y| \geq \varepsilon_1 d$. So by splitting clusters, we can find an equivalence relation that is a refinement of $\overset{\omega}{\leftrightarrow}$ in which the equivalence class of each vertex $y \in Y$ has between $\varepsilon_1 d/2$ and $\varepsilon_1 d$ total vertices. We now apply Lemma 5.4.11 to this equivalence relation with the sets V and Y in place of X and Y . (The hypothesis of Lemma 5.4.11 is met because $\varepsilon_1 d |Y| \geq 4|V|$.) This yields a collection of disjoint ω -connected subsets $(Z_r)_{r \in R}$ of Y such that for every r ,

$$|Z_r| \geq \frac{\varepsilon_1 d |Y|}{4|V|} \geq \frac{\varepsilon_1^3 d}{8} \quad \text{and} \quad |R| \geq \frac{|V|}{4\varepsilon_1 d} \geq \frac{|V|}{4d}.$$

Whenever H activates β and $u \leftrightarrow v$, $H \setminus \{e\}$ also activates β . On the event \mathcal{B} we must have that $x \leftrightarrow y$ and $\phi_x(u) \leftrightarrow \phi_y(u)$ whenever x, y both belong to Z_r for some $r \in R$. Thus, on the event \mathcal{B} , if there exist $r \in R$ and $x \in Z_r$ such that $x \leftrightarrow \phi_x(u)$ then $\phi_{x'}(H \setminus \{uv\})$ activates β for every $x' \in Z_r$.

In view of this fact, our next step will be to increase the percolation parameter to open an edge $\phi_x(e)$ with $x \in Z_r$ for a positive proportion of the indices $r \in R$, thus making a positive proportion of the copies $\phi_x(H \setminus \{e\})$ with $x \in \bigcup_{r \in R} Z_r$ activate β .

Let $m > n \geq \lceil \log_2(8/\varepsilon_1) \rceil$, and let \mathbb{P} be the joint law of the standard monotone coupling (ω, ω') of percolation with the two parameters $p_n \leq p_m$. It suffices to prove that there exists a constant $\eta_n = \eta_n(\alpha, \delta, \varepsilon) > 0$ such that

$$\mathbb{P}_{p_m}(\text{Act}_\beta(H \setminus \{e\})) \geq \eta_n. \quad (5.4.5)$$

Recall that the increasing sequence $(p_n)_{n \geq 1}$ was defined so that

$$p_m - p_n \geq 3^{-(n+1)} \mathcal{L}(Q) \geq 3^{-(n+1)} \frac{\delta}{4d}$$

and

$$\mathbb{P}_{p_m}(\|K_1\| \geq \beta) - \mathbb{P}_{p_n}(\|K_1\| \geq \beta) \leq 2^{-n+1}$$

for every $m > n$. The assumption that $m > n \geq \log_2(8/\varepsilon_1)$ implies that $\mathbb{P}_{p_m}(\|K_1\| \geq \beta) - \mathbb{P}_{p_n}(\|K_1\| \geq \beta) \leq \frac{\varepsilon_1}{4}$, and since $\mathbb{P}_{p_n}(\mathcal{B}) \geq \frac{\varepsilon_1}{2}$ a union bound gives that

$$\begin{aligned} \mathbb{P}(\omega \in \mathcal{B} \text{ and } \|K_1(\omega')\| < \beta) &\geq \mathbb{P}_{p_n}(\mathcal{B}) - \mathbb{P}(\|K_1(\omega)\| < \beta \leq \|K_1(\omega')\|) \\ &= \mathbb{P}_{p_n}(\mathcal{B}) - \end{aligned} \tag{5.4.6}$$

$$\begin{aligned} &\left(\mathbb{P}_{p_m}(\|K_1(\omega)\| \geq \beta) - \mathbb{P}_{p_n}(\|K_1(\omega)\| \geq \beta) \right) \\ &\geq \frac{\varepsilon_1}{2} - \frac{\varepsilon_1}{4} = \frac{\varepsilon_1}{4}, \end{aligned} \tag{5.4.7}$$

and hence in particular that

$$\mathbb{P}(\|K_1(\omega')\| < \beta \mid \omega \in \mathcal{B}) \geq \frac{\varepsilon_1}{4}. \tag{5.4.8}$$

Consider an arbitrary configuration $\xi \in \mathcal{B}$ and let $(Z_r)_{r \in R}$ be a collection of sets as defined via Lemma 5.4.11 above, which we take to be a function of ξ . (Note in particular that the index set R depends on ξ .) For each $r \in R$ and $x \in Z_r$ we have that

$$\mathbb{P}(\phi_x(e) \in \omega' \mid \omega = \xi) \geq \frac{p_m - p_n}{1 - p_n} \geq 3^{-(n+1)} \frac{\delta}{4d},$$

and since $|\{\phi_x(e) : e \in Z_r\}| \geq \frac{1}{2} |Z_r| \geq \frac{\varepsilon_1^3}{16} d$ (where the factor of 1/2 accounts for the distinction between oriented and unoriented edges, with ϕ_x acting bijectively on the former), we deduce that there exists $\varepsilon_2 = \varepsilon_2(\alpha, \delta, \varepsilon) > 0$ such that

$$\begin{aligned} \mathbb{P}(\exists x \in Z_r : \phi_x(e) \in \omega' \mid \omega = \xi) &\geq 1 - \left(1 - 3^{-(n+1)} \frac{\delta}{4d}\right)^{\varepsilon_1^3 d / 16} \\ &\geq 1 - \exp\left[-3^{-(n+1)} \frac{\varepsilon_1^3 \delta}{64}\right] \geq \varepsilon_2 3^{-n}, \end{aligned} \tag{5.4.9}$$

where we used the inequality $1 - x \leq e^{-x}$ in the second line. Condition on the event $\omega = \xi$ and consider the random set $J := \{r \in R : Z_r \xleftrightarrow{\omega'} Y\}$. Note that for all $r_1, r_2 \in R$ with $Z_{r_1} \not\xleftrightarrow{\omega} Y$ and $Z_{r_2} \not\xleftrightarrow{\omega} Y$, we have $\{\phi_x(e) : x \in Z_{r_1}\} \cap \{\phi_x(e) : x \in Z_{r_2}\} = \emptyset$. So by independence, it follows that there exists a constant $C_2 = C_2(\alpha, \delta, \varepsilon)$ such that if $|R| \geq C_2 3^n$ then

$$\mathbb{P}\left(|J| \geq \frac{\varepsilon_2}{2} 3^{-n} |R| \mid \omega = \xi\right) \geq 1 - \frac{\varepsilon_1}{8}. \tag{5.4.10}$$

We will now proceed by case analysis according to whether $|R(\xi)| \geq C_2 3^n$ for every $\xi \in \mathcal{B}$ or $|R(\xi)| < C_2 3^n$ for some $\xi \in \mathcal{B}$. First assume that $|R(\xi)| \geq C_2 3^n$ for every $\xi \in \mathcal{B}$. On the event

$\omega \in \mathcal{B}$, we write $(Z_r)_{r \in R}$ and J for the corresponding sets defined *with respect to* ω . We have by (5.4.8) and (5.4.10) that

$$\mathbb{P}\left(|J| \geq \frac{\varepsilon_2}{2} 3^{-n} |R| \text{ and } \|K_1(\omega')\| < \beta \mid \omega \in \mathcal{B}\right) \geq \frac{\varepsilon_1}{4} - \frac{\varepsilon_1}{8} = \frac{\varepsilon_1}{8}, \quad (5.4.11)$$

and hence that

$$\mathbb{P}\left(\omega \in \mathcal{B}, |J| \geq \frac{\varepsilon_2}{2} 3^{-n} |R|, \text{ and } \|K_1(\omega')\| < \beta\right) \geq \frac{\varepsilon_1^2}{16}. \quad (5.4.12)$$

Consider a pair of configurations (ω, ω') satisfying the event whose probability is estimated in (5.4.12) and pick $r \in J$ and $y \in Z_r$. Since $\phi_y(H)$ activates β in ω and $\|K_1(\omega')\| < \beta$, we also have that $\phi_y(H)$ activates β in ω' . By definition of J , we can find $x \in Z_r$ such that $\phi_x(e)$ is open in ω' . Since $x, y \in Z_r$ we have by definition that x and y are connected in ω , and since $x, y \in Y$ we have by definition of Y that $\phi_x(u)$ and $\phi_y(u)$ are also connected in ω . Since $\phi_x(e)$ is open in ω' and $\phi_x(v) = x$, we deduce that $x, y, \phi_x(u)$, and $\phi_y(u)$ all belong to the same cluster of ω' . Since $\phi_y(H)$ activates β in ω' and y is connected to $\phi_y(u)$ in ω' , we deduce that $\phi_y(H \setminus \{e\})$ activates β in ω' also. Since this holds for every $r \in J$ and $y \in Z_r$, we have

$$|W| := |\{x \in X : \omega' \in \text{Act}_\beta(\phi_x(H \setminus \{e\}))\}| \geq |J| \min_{r \in J} |Z_r|.$$

Since $|J| \geq \frac{\varepsilon_2}{2} 3^{-n} |R| \geq 3^{-n} \frac{\varepsilon_2}{8} \frac{|V|}{d}$ and every Z_r has $|Z_r| \geq \frac{\varepsilon_1^3}{8} d$, we have that $\|W\| \geq \frac{\varepsilon_1^3 \varepsilon_2}{64} 3^{-n}$ and hence that

$$\mathbb{E} \|W\| \geq \frac{\varepsilon_1^5 \varepsilon_2}{2^{10}} 3^{-n}.$$

Thus, it follows by vertex-transitivity that $\mathbb{P}_{p_m}(\text{Act}_\beta(H \setminus \{e\})) \geq \frac{\varepsilon_1^5 \varepsilon_2}{2^{10}} 3^{-n}$. This bound has the required form, completing the proof of (5.4.5) in this case.

It remains to consider the case that $|R(\xi)| < C_2 3^n$ for some $\xi \in \mathcal{B}$. Since we always have $|R| \geq |V|/4d$, we must have in this case that

$$|V| \leq 4C_2 3^n d.$$

Observe that if $\omega \in \mathcal{B}$ and $\|K_1(\omega')\| < \beta$ then $\omega' \in \mathcal{B}$ also. Thus, by (5.4.8) and the fact that $\mathbb{P}_{p_m}(\mathcal{B}) \geq \frac{\varepsilon_1}{2}$, we have that

$$\mathbb{P}_{p_m}(\mathcal{B}) \geq \frac{\varepsilon_1^2}{8}.$$

On the event \mathcal{B} , pick one of the sets Z_r that are guaranteed to exist by our earlier argument, call it Z , and let $U := \{x \in Z : \omega \in \text{Act}_\beta(\phi_x(e))\}$. If

$$\mathbb{P}_{p_m}\left(\mathcal{B} \cap \left\{\|U\| \geq \frac{1}{2} \|Z\|\right\}\right) \geq \frac{1}{2} \mathbb{P}_{p_m}(\mathcal{B})$$

then we have by vertex-transitivity and the bound $|Z| \geq \varepsilon_1^3 d/8$ that

$$\mathbb{P}_{p_m}(\text{Act}_\beta(e)) \geq \mathbb{E}_{p_m}[\|U\| \mathbb{1}(\mathcal{B})] \geq \frac{\varepsilon_1^3 d}{16|V|} \cdot \frac{\varepsilon_1^2}{16} \geq \frac{\varepsilon_1^5}{2^{10}C_2} \cdot 3^{-n}$$

as required. Conversely, if

$$\mathbb{P}_{p_m}\left(\mathcal{B} \cap \left\{\|U\| \geq \frac{1}{2}\|Z\|\right\}\right) \leq \frac{1}{2}\mathbb{P}_{p_m}(\mathcal{B}),$$

then we use insertion-tolerance (Proposition 6.6.4 - using the fact that $\|Z\| \geq \frac{\varepsilon_1^3}{64C_23^n}$) to open a single edge in $\{\phi_x(e) : x \in Z \setminus U\}$ on the event $\mathcal{B} \cap \{\|U\| \leq \frac{1}{2}\|Z\|\}$. This does not create a cluster with density at least β , since none of the edges $\phi_x(e)$ with $x \in Z \setminus U$ activates β . Arguing as before, after opening this edge, *every* set $\phi_x(H \setminus \{e\})$ with $x \in Z$ activates β . Thus, it follows by the same vertex-transitivity argument as above that there exists a positive constant $\varepsilon_3(n) = \varepsilon_3(\alpha, \delta, \varepsilon, n)$ such that

$$\mathbb{P}_{p_m}(\text{Act}_\beta(H \setminus \{e\})) \geq \varepsilon_3.$$

This completes the proof. \square

Completing the proof of the main theorems

In this section we complete the proof of Theorem 8.1.1 and hence of Theorem 5.1.2.

Proof of Theorem 8.1.1. The implication (i) \Rightarrow (ii) follows immediately from Corollary 5.3.4 and Proposition 5.4.1, while the implication (iii) \Rightarrow (iv) is trivial. As such, it remains only to prove (ii) \Rightarrow (iii) and (iv) \Rightarrow (i).

We begin with the implication (ii) \Rightarrow (iii), i.e., the claim that molecular sets admit linear 1/3-separators. To see this, note that if \mathcal{H} is m -molecular for some $m \geq 2$ then there exists a constant C such that for each $G \in \mathcal{H}$ there is an automorphism-invariant set of edges F such that $|F| \leq C|V(G)|$ and F disconnects G into $m \geq 2$ connected components each of size $|V(G)|/m$. If we take A to be the union of $\lceil m/2 \rceil$ of these components then $|V(G)|/3 \leq |A| = \lceil m/2 \rceil |V(G)|/m \leq 2|V(G)|/3$ and $|\partial_E A| \leq |F| \leq C|V(G)|$ so that \mathcal{H} has linear 1/3-separators as claimed.

We next prove the implication (iv) \Rightarrow (i). We will use the following theorem [Bol+10a] identifying the location of the critical threshold for (not necessarily transitive) dense graph sequences.

Theorem 5.4.12 (Bollobás, Borgs, Chayes, Riordan 2010). *Let (G_n) be a dense sequence of finite, simple graphs with $|V(G_n)| \rightarrow \infty$, and for each $n \geq 1$ let λ_n be the largest eigenvalue of the*

adjacency matrix of G_n . For each $c > 0$

$$\lim_{n \rightarrow \infty} \mathbb{P}_{\lambda_n^{-1}}^{G_n}(\|K_1\| \geq c) = 0,$$

and for each $\varepsilon > 0$ there exists $\delta > 0$ such that

$$\lim_{n \rightarrow \infty} \mathbb{P}_{(1+\varepsilon)\lambda_n^{-1}}^{G_n}(\|K_1\| \geq \delta) = 1.$$

Note that if A is the adjacency matrix of a finite graph then A is self-adjoint and the largest eigenvalue (a.k.a. the Perron–Frobenius eigenvalue) of A coincides with the L^2 operator norm of A . This norm satisfies $\|A\| \geq \langle A\mathbb{1}, \mathbb{1} \rangle / \langle \mathbb{1}, \mathbb{1} \rangle = 2|E|/|V|$ where $\mathbb{1}$ is the constant-one function, and hence if $(G_n)_{n \geq 1}$ is dense then the sequence of largest eigenvalues λ_n satisfies

$$\liminf |V(G_n)|^{-1} \lambda_n > 0.$$

Suppose that \mathcal{H} is dense and admits linear θ -separators for some $\theta \in (0, 1/2]$, so that there exists a constant C_1 such that for each $G \in \mathcal{H}$ there exists a set $A(G) \subseteq V(G)$ with $\theta|V(G)| \leq |A(G)| \leq (1 - \theta)|V(G)|$ and $|\partial_E A(G)| \leq C_1|V(G)|$. For each $G \in \mathcal{H}$ let $H(G)$ and $H(G)^c$ denote the subgraphs of G induced by $A(G)$ and $A(G)^c$ respectively. Since every vertex of G has degree $2|E(G)|/|V(G)|$ we have that

$$\min\{|E(H(G))|, |E(H(G)^c)|\} \geq \theta|E(G)| - |\partial_E A(G)|.$$

For large n we have that $\theta|E(G)| \gg |\partial_E A(G)|$ and hence that $(H(G))_{G \in \mathcal{H}}$ and $(H(G)^c)_{G \in \mathcal{H}}$ are both dense. Thus, it follows from Theorem 5.4.12 that there exists a constant C_2 such that if we set $p(G) = C_2/|V(G)|$ for each $G \in \mathcal{H}$ then both $H(G)$ and $H(G)^c$ contain a giant component with high probability under percolation with parameter $p(G)$. The same also holds at $q(G) = 2C_2/|V(G)|$, which is ε -supercritical for \mathcal{H} for an appropriate choice of $\varepsilon > 0$. But at this same parameter $q(G)$ the expected number of edges connecting $A(G)$ and $A(G)^c$ is bounded by $2C_1C_2$, so that by Poisson approximation the probability that there are no such edges is bounded away from zero uniformly over $G \in \mathcal{H}$. Thus the probability that $q(G)$ -percolation on G contains at least two giant components is bounded away from zero uniformly over $G \in \mathcal{H}$, and hence \mathcal{H} does not have the supercritical uniqueness property. It follows that the supercritical uniqueness property fails for any set \mathcal{H}' with $\mathcal{H} \subseteq \mathcal{H}' \subseteq \mathcal{F}$, by extending q from \mathcal{H} to \mathcal{H}' by setting $q(G) := 1$ for all $G \in \mathcal{H}' \setminus \mathcal{H}$. \square

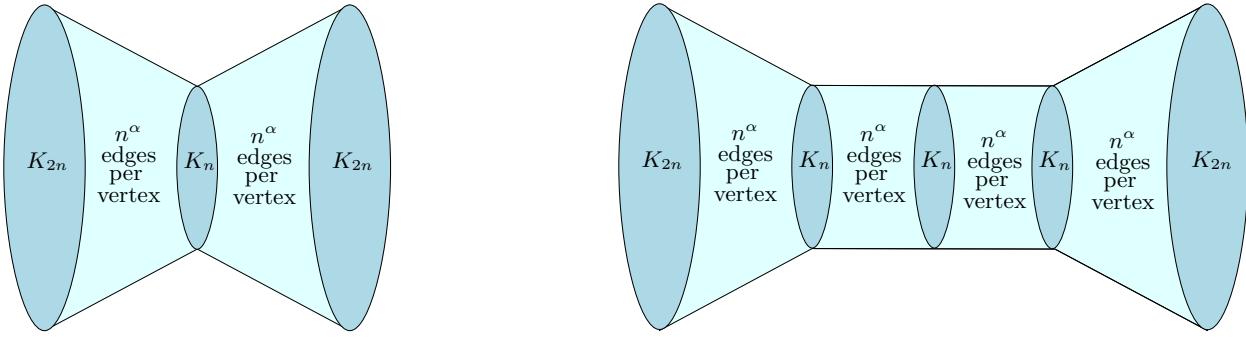


Figure 5.1: Schematic illustration of the graphs discussed in remark 5.4.2

Remark 5.4.2. The proof of (iv) \Rightarrow (i) above shows more generally that any sequence of finite, simple graphs with linear minimal degree and linear θ -separators for some $\theta \in (0, 1/2]$ fails to have the supercritical uniqueness property. On the other hand, the results of [Bol+10a] imply that dense graphs without *subquadratic* separators have the supercritical uniqueness property. It is natural to wonder in light of Theorem 8.1.1 whether the failure of supercritical uniqueness in graphs of linear minimal degree is always characterized by the existence of linear separators, without the assumption of vertex-transitivity.

This is not the case. Indeed, let $0 < \alpha < 1/2$ and suppose that we take two copies of K_{2n} and one copy of K_n arranged in a line with the two copies of K_{2n} at the end and the copy of K_n in the middle. We may glue these copies together in such a way that each vertex is connected to each of the complete graphs adjacent to its own complete graph by between n^α and $3n^\alpha$ edges. It is easily verified that the smallest separators in the resulting graph sequence are of order $n^{1+\alpha}$ and that there will exist two distinct giant clusters with high probability when $p = 3/4n$. We focus on the second claim, which is more involved. For such p the two copies of K_{2n} are supercritical and each contains a giant cluster, while the copy of K_n is subcritical and has largest cluster of order $\log n$ with high probability. Thus, when we add in the edges between the various complete graphs, the probability that there exists a cluster in the copy of K_n that has an edge connecting it to both of its neighbouring copies of K_{2n} is small: a K_n -cluster of size $m = O(\log n)$ has both such edges adjacent to it with probability of order $m^2 n^{2\alpha-2}$, and since there are at most $2n$ such clusters the total conditional probability is $O((\log n)^2 n^{2\alpha-1}) = o(1)$ with high probability, yielding the claim. This gives an example of a linear minimal-degree graph sequence that does not have linear separators but does not have the supercritical uniqueness property either. By considering longer chains of copies of K_n connecting the two copies of K_{2n} as in fig. 5.1, one can obtain similar examples where the minimal size of a separator scales like an arbitrary power of n between n and n^2 .

5.5 Closing remarks

Counterexamples. We now discuss examples demonstrating that Theorem 8.1.1 does not extend to arbitrary insertion-tolerant percolation models on the torus or to critical percolation.

Example 5.5.1 (Multiple giants at criticality). The cycle $\mathbb{Z}/n\mathbb{Z}$ (with its standard generating set) has multiple giant components with good probability when $p = 1 - \lambda/n$, and the set of closed edges converges to a Poisson process on the circle as $n \rightarrow \infty$. As observed by Alon, Benjamini, and Stacey [ABS04b], by considering the highly asymmetric torus $T_n := (\mathbb{Z}/2^n\mathbb{Z}) \times (\mathbb{Z}/n\mathbb{Z})$ (with its standard generating set) one can obtain similar behaviour at values of p that are bounded away from 1.

Since these authors did not include a proof, let us now very briefly indicate how the analysis of this example works. We assume for notational simplicity that n is a power of 2 and hence is a factor of 2^n . Let X be the set of cylinders in $(\mathbb{Z}/2^n\mathbb{Z}) \times (\mathbb{Z}/n\mathbb{Z})$ of the form $[kn, (k+1)n] \times (\mathbb{Z}/n\mathbb{Z})$ whose vertices are incident to some simple cycle of dual edges that wrap around the torus. When $p > 1/2$, it follows by sharpness of the phase transition on \mathbb{Z}^2 that there exists a constant $c_p > 0$ with $c_p \rightarrow \infty$ as $p \rightarrow 1$ such that each particular cylinder has probability at most $e^{-c_p n}$ to belong to X , this probability being bounded by the probability that a box of size n in \mathbb{Z}^2 intersects a dual cluster of diameter at least n . On the other hand, since the correlation length on \mathbb{Z}^2 diverges as $p \uparrow 1/2$, there exists a fixed $p > 1/2$ such that $\mathbb{E}_p |X| \rightarrow \infty$ as $n \rightarrow \infty$. Using this one can prove that if we define p_n to be the unique value such that $\mathbb{E}_{p_n} |X| = 1$ then $\liminf p_n > 1/2$ and $\limsup p_n < 1$. It is fairly straightforward to prove that there are multiple giant clusters with positive probability at p_n . Indeed, using the same exponential decay estimates one can prove that non-neighbouring cylinders are highly de-correlated, so that $\mathbb{E}_{p_n} |X|^2 = O(1)$ and there is a good probability for X to contain at least two elements that are well-spaced around the torus. On this event there must be at least two giant components with high probability: since p_n is bounded away from $p_c(\mathbb{Z}^2) = 1/2$, each vertex of the torus has good probability to belong to a cluster that wraps around the torus, and any two such vertices can be disconnected only if there are two closed dual cut-cycles separating them. The events that two distant vertices belong to such wrapping clusters are highly de-correlated, so that the number of vertices belonging to wrapping clusters is linear with high probability and the claim follows.

Example 5.5.2 (Insertion-tolerance on the torus is not enough). Consider the symmetric torus $(\mathbb{Z}/10n\mathbb{Z})^2$ with its standard generating set. Consider the model defined as follows: First, select a pair of vertical strips of the form $[k, k+n] \times (\mathbb{Z}/10n\mathbb{Z})$ and $[k+5n, k+6n] \times (\mathbb{Z}/10n\mathbb{Z})$ uniformly from among the n available possibilities. Declare each edge belonging to one of these strips open with probability $1/4$ and each edge not belonging to one of these strips open with probability $3/4$.

Using standard properties of subcritical and supercritical percolation on \mathbb{Z}^2 , one easily obtains that this model contains exactly two giant clusters with high probability: the two large high-density strips will each contain a giant cluster with high probability, while there will be no clusters crossing either of the two thin low-density strips with high probability. By also applying a random rotation in $\{0, \pi/2, \pi, 3\pi/4\}$ one obtains an automorphism invariant, uniformly insertion-tolerant, percolation-in-random-environment model on $(\mathbb{Z}/10n\mathbb{Z})^2$ with the same properties. As such, one should not expect any uniqueness-of-the-giant-component results to hold on finite graphs at anywhere near the same generality as found in the Burton–Keane theorem [BK89], even when restricting to symmetric tori of fixed dimension. By taking the relative width of the low-density strips to go to zero in a well-chosen manner as $n \rightarrow \infty$, one can construct a similar example in which the number of giant clusters is either one or two each with good probability and any two vertices are connected with good probability.

The supercritical existence property. Theorem 8.1.1 can be thought of a geometric characterisation of the infinite sets $\mathcal{H} \subseteq \mathcal{F}$ for which supercritical percolation has *at most* one giant cluster with high probability. We now briefly address the complementary problem of whether there is *at least* one giant cluster with high probability in supercritical percolation, noting that the definitions only ensure that such a cluster exists with *good* probability (i.e., with probability bounded away from zero).

Let $\mathcal{H} \subseteq \mathcal{F}$ be an infinite set. We say that \mathcal{H} has the *supercritical existence property* if for every supercritical assignment $p : \mathcal{H} \rightarrow [0, 1]$ there exists a constant $\alpha > 0$ such that

$$\liminf_{G \in \mathcal{H}} \mathbb{P}_{p(G)}^G (\text{the largest cluster contains at least } \alpha |V(G)| \text{ vertices}) = 1.$$

Notice that the sharp density property immediately implies the supercritical existence property. However the converse is false because (as we will show below) molecular graphs also have the supercritical existence property. This might lead one to suspect that the supercritical existence property always holds. In fact, this is not the case, and the counterexamples can once again be exactly characterised in terms of molecular graphs. Let us note that the weaker supercritical existence property

$$\lim_{\alpha \downarrow 0} \liminf_{G \in \mathcal{H}} \mathbb{P}_{p(G)}^G (\text{the largest cluster contains at least } \alpha |V(G)| \text{ vertices}) = 1$$

always holds (even without transitivity) as an immediate consequence of the universal tightness theorem [Hut21b]. The following is a fairly straightforward consequence of our results and those of [Bol+10a].

Corollary 5.5.3. *An infinite set $\mathcal{H} \subseteq \mathcal{F}$ has the supercritical existence property if and only if it is not the case that there exist arbitrarily large integers m for which \mathcal{H} contains an m -molecular subset.*

Sketch of proof. Applying the results of [Bol+10a] as in the proof of Theorem 8.1.1 easily yields that if \mathcal{H} is m -molecular then

$$\liminf_{G \in \mathcal{H}} \mathbb{P}_{(1-\varepsilon)^{-1} p_c^G(\alpha, \varepsilon)}^G (\|K_1\| \leq 1/m) > 0$$

for every $0 < \alpha, \varepsilon < 1$, yielding the forward implication of the claim. We now suppose \mathcal{H} does not have the supercritical existence property and argue that it contains m -molecular subsequences for arbitrarily large values of m . It is an immediate consequence of Proposition 5.4.1 that \mathcal{H} has at least one molecular subsequence. Suppose for contradiction that the supremal value of m such that \mathcal{H} has an m -molecular subsequence is finite, and denote this supremum by M . Since \mathcal{H} does not have the supercritical existence property, there exists an infinite subset $\mathcal{S} \subseteq \mathcal{H}$ such that

$$\limsup_{G \in \mathcal{S}} \mathbb{P}_{(1-\varepsilon)^{-1} p_c^G(\varepsilon, \varepsilon)}^G \left(\|K_1\| \geq \frac{\varepsilon^2}{4M} \right) < 1,$$

and applying Proposition 5.4.1 as before we may assume that this subset is m -molecular for some $2 \leq m \leq M$. By taking a further infinite subset and changing m if necessary, we may assume that this subset has density $\deg(G)/|V(G)|$ converging to some constant $c > 0$ and does not have a further subset that is k -molecular for any $k > m$.

By definition there exists a constant C such that for each $G \in \mathcal{S}$ there exists an automorphism-invariant set $F_G \subseteq E(G)$ with $|F_G| \leq C|V_G|$ such that $G \setminus F_G$ has m connected components. For each $G \in \mathcal{S}$, let H_G be a graph isomorphic to each of the m connected components of $G \setminus F_G$. Since \mathcal{S} does not admit a subset that is k -molecular for any $k > m$, $\mathcal{A} := \{H_G : G \in \mathcal{S}\}$ cannot itself contain a molecular subset. Thus, it follows from Proposition 5.4.1 and Theorem 8.1.1 that \mathcal{A} has the sharp density property and the supercritical uniqueness property. On the other hand, the vertex degrees of H_G with $G \in \mathcal{S}$ are asymptotically equal to those of G in the sense that the ratio tends to 1, and Theorem 5.4.12 allows us to compute the location of the percolation thresholds for these transitive dense graph sequences in terms of their vertex degrees. (Recall that the largest eigenvalue for the adjacency matrix of a regular graph is equal to its vertex degree.) So any assignment $\tilde{p} : \mathcal{A} \rightarrow [0, 1]$ built from the assignment $p : \mathcal{S} \rightarrow [0, 1]$ with $p(G) := (1 - \varepsilon)^{-1/2} \cdot p_c^G(\varepsilon, \varepsilon)$ by arbitrarily picking $\tilde{p}(H) \in \{p(G) : H_G = H\}$ for each $H \in \mathcal{A}$ is itself supercritical for \mathcal{A} . We deduce from Theorem 8.1.1 that

$$\lim_{G \in \mathcal{S}} \mathbb{E}_{p(G)}^{H_G} \|K_2\| = 0.$$

In order for a vertex to belong to the largest cluster in some $G \in \mathcal{S}$, either it must belong to the largest cluster in its copy of H_G , or this is not the case and there is an open edge of F_G incident to its cluster. By vertex transitivity every vertex of G is incident to at most C edges of F_G . It follows that, writing K_i for the i^{th} largest cluster,

$$\begin{aligned}\mathbb{E}_{p(G)}^G \|K_1\| &\leq \mathbb{E}_{p(G)}^{H_G} \left[\|K_1\| + \sum_{i \geq 2} C p(G) \|K_i\| \cdot |K_i| \right] \\ &\leq \mathbb{E}_{p(G)}^{H_G} \|K_1\| + C p(G) |V(G)| \mathbb{E}_{p(G)}^{H(G)} \|K_2\|.\end{aligned}$$

Since $p(G)$ is of order $|V(G)|^{-1}$ the second term tends to zero as $G \rightarrow \infty$ with $G \in \mathcal{S}$, and we deduce that $\mathbb{E}_{p(G)}^{H_G} \|K_1\| \geq \frac{1}{2} \mathbb{E}_{p(G)}^G \|K_1\| \geq \frac{\varepsilon^2}{2}$ for all but finitely many $G \in \mathcal{S}$. Define $p' : \mathcal{S} \rightarrow [0, 1]$ by

$$p'(G) := (1 - \varepsilon)^{-1/2} \cdot p(G) = (1 - \varepsilon)^{-1} \cdot p_c^G(\varepsilon, \varepsilon).$$

Since \mathcal{A} does not have any molecular subsets, it follows by Markov's inequality and Proposition 5.4.1 that

$$\liminf_{G \in \mathcal{S}} \mathbb{P}_{p'(G)}^G \left(\|K_1\| \geq \frac{\varepsilon^2}{4M} \right) \geq \liminf_{G \in \mathcal{S}} \mathbb{P}_{p'(G)}^{H_G} \left(\|K_1\| \geq \frac{\varepsilon^2}{4} \right) = 1,$$

a contradiction. \square

SUPERCRITICAL PERCOLATION ON FINITE TRANSITIVE GRAPHS II: CONCENTRATION, LOCALITY, AND EQUICONTINUITY OF THE GIANT'S DENSITY

Joint with Tom Hutchcroft

Abstract

In the previous paper of this series, we showed that supercritical percolation on a large finite transitive graph typically has exactly one giant cluster. In the present paper, we simultaneously establish that the density of this unique giant cluster is concentrated around its mean and that this mean is equicontinuous with respect to the (suitably scaled) percolation parameter and is determined by the “local geometry” (suitably interpreted) of G . For example, consider the torus

$$\mathbb{T}_n^d := (\mathbb{Z}/n\mathbb{Z}_n)^d.$$

Our general arguments recover the well-known fact that the supercritical giant cluster density for $(\mathbb{T}_n^d)_{n \geq 1}$ ($d \geq 2$ fixed) converges to the infinite cluster density on \mathbb{Z}^d , whereas the supercritical giant cluster density for $(\mathbb{T}_n^d)_{d \geq 1}$ ($n \geq 2$ fixed) converges to the survival probability of a Poisson branching process.

Our proof relies on a new perspective on how to use sharp threshold theory in percolation: to exploit the fact that certain events of interest (such as the event $o \leftrightarrow \infty$) do *not* undergo sharp thresholds. These arguments apply equally well to infinite transitive graphs, yielding analogous, new results in this context too. For example, we show that for all $\varepsilon > 0$ and $d \in \mathbb{N}$, the function given by the restriction of

$$\theta_G(p) := \mathbb{P}_p^G(o \leftrightarrow \infty) \quad \text{to} \quad p \in [p_c(G) + \varepsilon, 1]$$

is uniformly equicontinuous as G varies over all infinite transitive graphs with vertex degree d .

6.1 Introduction

This paper is the second in a series investigating the supercritical phase of Bernoulli bond percolation on finite, connected, vertex-transitive graphs. The overarching goal of both papers is to obtain results that hold for all such graphs, in contrast to earlier works that have focused on particular geometric

settings (such as complete graphs, tori, or expanders) and used methods specific to those examples. In the first paper of this series, we answered the most basic question about the geometry of clusters in this phase by showing that there is typically a *unique* giant cluster with high probability. In this paper, we study the *density* of this giant cluster. We will show that as the volume of the graph tends to infinity, the *density* of this giant cluster concentrates around a *deterministic* value, unless the graph belongs to an explicit family of counterexamples known as *molecular graphs* which we introduced in the first paper. We then investigate the continuity properties of this limiting density, determining the senses in which it is determined by the local geometry of G and continuous in p . As in [EH21a], the theory we develop applies without any constraints on the degree, but is also new in the bounded degree case.

Setting the scene

Before stating our results, we first briefly overview the results of the first paper in the series [EH21a] and the first author's companion paper [Eas22]. In addition to providing important context for our new results, this will also give us an opportunity to introduce relevant definitions and notation.

Let \mathcal{G} be the set of all isomorphism classes of connected, simple (i.e., not containing loops or multiple edges), locally finite, vertex-transitive graphs, and let $\mathcal{F} = \{G \in \mathcal{G} : G \text{ finite}\}$. (We work with isomorphism classes partly to make sure these really are sets; \mathcal{F} is countably infinite while \mathcal{G} has the cardinality of the continuum. We will usually suppress the distinction between graphs and their isomorphism classes as much as possible when this does not cause any confusion.) Given an infinite set $\mathcal{H} \subseteq \mathcal{F}$, a function $\phi : \mathcal{H} \rightarrow \mathbb{R}$, and $\alpha \in \mathbb{R}$ we write $\lim_{G \in \mathcal{H}} \phi(G) = \alpha$ or “ $\phi(G) \rightarrow \alpha$ as $G \rightarrow \infty$ in \mathcal{H} ” to mean that for each $\varepsilon > 0$ there exists N such that $|\phi(G) - \alpha| \leq \varepsilon$ for every $G \in \mathcal{H}$ with at least N vertices, or equivalently that $\phi(G_n) \rightarrow \alpha$ for some (and hence every) enumeration $\mathcal{H} = \{G_1, G_2, \dots\}$ of \mathcal{H} . Similar conventions apply to define $\limsup_{G \in \mathcal{H}}$, $\liminf_{G \in \mathcal{H}}$, and limits that may be equal to $+\infty$ or $-\infty$. Given two positive functions f and g on \mathcal{H} , we write “ $f(G) \sim g(G)$ as $G \rightarrow \infty$ in \mathcal{H} ” to mean that $\lim_{G \in \mathcal{H}} f(G)/g(G) = 1$.

Given a countable graph $G = (V, E)$ and $p \in [0, 1]$, we write $\mathbb{P}_p = \mathbb{P}_p^G$ for the law of Bernoulli- p bond percolation on G , i.e., the random subgraph of G in which each edge is included independently at random with inclusion probability p . Given a percolation configuration $\omega \in \{0, 1\}^E$ on G , the connected components of ω are called **clusters**. We write K_u to denote the cluster containing the vertex u and write $u \leftrightarrow v$ for the event that $K_u = K_v$. Given a subset W of V , the **volume** of W is the number of vertices in W , denoted $|W|$, while if G is finite, the **density** of W is defined to be the ratio $\|W\| := |W|/|V|$. We write K_1, K_2, \dots for the clusters of ω in decreasing order of volume

(breaking ties arbitrarily).

Given an infinite set $\mathcal{H} \subseteq \mathcal{F}$, we say that an assignment of parameters $p_c : \mathcal{H} \rightarrow [0, 1]$ is a **percolation threshold** if

1. $\lim_{H \in \mathcal{H}} \mathbb{P}_{(1-\varepsilon)p_c} (\|K_1\| \geq c) = 0$ for every $\varepsilon, c > 0$, and
2. For every $\varepsilon > 0$ there exists $\alpha > 0$ such that $\lim_{H \in \mathcal{H}} \mathbb{P}_{(1+\varepsilon)p_c} (\|K_1\| \geq \alpha) = 1$, where we set $\mathbb{P}_p = \mathbb{P}_1$ for $p \geq 1$.

Note that critical thresholds are *not* unique when they exist, but any two percolation thresholds $p_c, \tilde{p}_c : \mathcal{H} \rightarrow [0, 1]$ must satisfy $p_c(G) \sim \tilde{p}_c(G)$ as $G \rightarrow \infty$ in \mathcal{H} . When a percolation threshold $p_c : \mathcal{H} \rightarrow [0, 1]$ exists, we say that $p : \mathcal{H} \rightarrow [0, 1]$ is **supercritical** if the set $\mathcal{H}' = \{G \in \mathcal{H} : p(G) < 1\}$ is infinite and satisfies

$$\liminf_{G \in \mathcal{H}'} \frac{p(G)}{p_c(G)} > 1$$

or \mathcal{H}' is finite. We generalise this definition to include the case that a threshold p_c does not exist by saying that p is **supercritical** if \mathcal{H}' is infinite and there exists $\varepsilon > 0$ such that

$$\liminf_{G \in \mathcal{H}'} \mathbb{P}_{(1-\varepsilon)p} (\|K_1\| \geq \varepsilon) \geq \varepsilon,$$

or \mathcal{H}' is finite. Note that these two definitions of supercriticality coincide when \mathcal{H} admits a threshold function, and in particular that the definition of supercriticality does not depend on the choice of threshold function. This definition also guarantees that every infinite set $\mathcal{H} \subseteq \mathcal{F}$ admits a supercritical assignment $p : \mathcal{H} \rightarrow [0, 1]$, namely the trivial supercritical assignment $p(G) \equiv 1$. We say that \mathcal{H} has the **supercritical uniqueness property** if

$$\lim_{G \in \mathcal{H}} \mathbb{P}_p (\|K_2\| \geq \varepsilon) = 0$$

for every supercritical $p : \mathcal{H} \rightarrow [0, 1]$ and every constant $\varepsilon > 0$.

The first paper in this series [EH21a] together with the related work of the first author [Eas22] give simple characterizations of those transitive graph families that have the supercritical uniqueness property and that admit percolation thresholds respectively.

Theorem 6.1.1 ([EH21a, Theorem 1.2]). *An infinite set $\mathcal{H} \subseteq \mathcal{F}$ has the supercritical uniqueness property if and only if it does not contain an infinite molecular subset.*

Theorem 6.1.2 ([Eas22, Theorem 2]). *An infinite set $\mathcal{H} \subseteq \mathcal{F}$ admits a percolation threshold if and only if it does not contain an infinite m -molecular subset for infinitely many values of m .*

Here, given an integer $m \geq 2$, we say that an infinite set $\mathcal{H} \subseteq \mathcal{F}$ is **m -molecular** if it has the following properties, where $E(G)$ and $V(G)$ denote the set of edges and the set of vertices of a given graph G :

1. \mathcal{H} is **dense**, meaning that $\liminf_{G \in \mathcal{H}} \frac{|E(G)|}{|V(G)|^2} > 0$.
2. There exists a constant $C < \infty$ such that for each $G = (V, E) \in \mathcal{H}$ there exists a set of edges $F \subseteq E$ satisfying the following conditions:
 - a) $G \setminus F$ has exactly m connected components;
 - b) F is invariant under $\text{Aut}(G)$;
 - c) $|F| \leq C |V|$.

The stated conditions on G and F imply that the m connected components of $G \setminus F$ are dense, vertex-transitive, and isomorphic to each other. We say that an infinite set $\mathcal{H} \subseteq \mathcal{F}$ is **molecular** if it is m -molecular for some $m \geq 2$. For example, the set of Cartesian products of complete graphs $\{K_n \square K_m : n \geq 1\}$ is m -molecular for each $m \geq 2$. Applying the analysis of percolation on dense graphs carried out in [Bol+10b], it is fairly easy to see that the supercritical uniqueness property does not hold for molecular sets of graphs. For example, if we take $p : \{K_n \square K_m : n \geq 1\} \rightarrow [0, 1]$ defined by $p(K_n \square K_m) = 2/n$ then p is supercritical since each copy of K_n will contain a giant cluster by the classical theory of Erdős-Rényi random graphs, but the total number of edges between distinct copies of K_n converges to a Poisson random variable and hence is zero with positive probability. The main result of [EH21a] is that constructions of this form are the *only* way to get a non-unique giant in supercritical percolation on a transitive graph.

These theorems imply in particular that if $\mathcal{H} \subseteq \mathcal{F}$ is an infinite set that is **sparse**, meaning that $\lim_{G \in \mathcal{H}} \frac{|E(G)|}{|V(G)|^2} = 0$, then \mathcal{H} has the supercritical uniqueness property and admits a percolation threshold function. An important example of such an infinite set \mathcal{H} is \mathcal{F}_d (for each $d \geq 2$), the set of all (isomorphism classes of) finite, connected, simple, vertex-transitive graphs with degrees bounded by d .

Concentration: The density is well-defined

An important quantity associated to percolation on an *infinite* transitive graph at a given parameter p is the *infinite cluster density*, the proportion of vertices contained in infinite clusters. In this

context, if we define the expectation $\theta(p, G) := \mathbb{P}_p^G(o \leftrightarrow \infty)$, where o denotes an arbitrary vertex of G , then it is easily shown that this quantity accurately captures the density of infinite clusters in the sense that if A is any finite set of vertices then $|\{x \in A : x \leftrightarrow \infty\}| = (\theta \pm o(1))|A|$ with high probability when A is large. (Indeed, the variance of this random variable is easily seen to be $o(|A|^2)$. Getting sharp quantitative bounds on the fluctuations of this random variable in general geometry is a very interesting problem closely related to those of [HH19, Section 5.3].) Moreover, for percolation on \mathbb{Z}^d , it is an immediate consequence of the ergodic theorem that

$$\frac{1}{(2n+1)^d} \sum_{x \in [-n, n]^d} \mathbb{1}(x \leftrightarrow \infty) \rightarrow \theta$$

almost surely as $n \rightarrow \infty$. As such, for infinite graphs, the relationship between the almost-sure and in-expectation density of infinite clusters is trivial, and one instead focuses on questions concerning e.g. the dependence of the density on p , with continuity at p_c being a famous open problem for three-dimensional lattices.

For *finite* transitive graphs, the most natural analogue of the infinite cluster density is the *giant cluster density*. In this setting, the ergodic theorem no longer applies and the analogous question on the relation between expected and almost-sure densities become much more subtle. Of course, the picture we would naively expect is that as p increases across some threshold value $p_c(G)$, a unique macroscopic cluster should emerge, and the density of this macroscopic cluster should be concentrated around its mean $\theta(p, G) := \mathbb{P}_p^G(o \in K_1)$. Unfortunately this is not always the case: In the product of an n -vertex complete graph with an edge with $p = \lambda/n$ where $\lambda > 1$, it can be shown the largest cluster either has density close to either $\text{mf}(\lambda)$ or $\frac{1}{2}\text{mf}(\lambda)$ with high probability, where $\text{mf}(\lambda)$ is the limiting density of the giant cluster in the $p = \lambda/n$ Erdős-Rényi graph, with probability approximately $1 - e^{-\lambda \text{mf}(\lambda)^2}$ to have density close to $\text{mf}(\lambda)$; the resulting density $\theta \sim (1 - \frac{1}{2}e^{-\lambda \text{mf}(\lambda)^2}) \text{mf}(\lambda)$ does not adequately capture the bimodal nature of this limiting distribution. Even for bounded degree graphs, it is possible for similar behaviour to hold *at the critical point* as we see on the long torus $(\mathbb{Z}/n\mathbb{Z}) \times (\mathbb{Z}/2^n\mathbb{Z})$ [EH21a, Example 5.1]. Still it seems reasonable to conjecture that the density should be concentrated if we restrict to the supercritical case and impose some mild geometric conditions on the graph at hand.

In order to address this question, we first introduce a relevant definition: We say that an infinite set $\mathcal{H} \subseteq \mathcal{F}$ has the **supercritical concentration property** if

$$\lim_{G \in \mathcal{H}} \mathbb{P}_p \left(\left| \mathbb{E}_p \|K_1\| - \mathbb{E}_p \|K_1\| \right| \geq \varepsilon \right) = 0$$

for every supercritical $p : \mathcal{H} \rightarrow [0, 1]$ and every constant $\varepsilon > 0$. (Note that if concentration occurs at all then it must be around the mean by the bounded convergence theorem.) As with the uniqueness problem, it is reasonably easy to see that this property does *not* hold for molecular sets of graphs, for similar reasons to the concrete example of the product of a complete graph with an edge as discussed above. Our first main theorem states that, as with the supercritical uniqueness property, molecular sets are the *only* obstruction to the supercritical concentration property.

Theorem 6.1.3. *Let \mathcal{H} be an infinite set of (isomorphism classes of) finite, simple, connected, vertex-transitive graphs. The following are equivalent:*

1. \mathcal{H} has the supercritical concentration property.
2. \mathcal{H} has the supercritical uniqueness property.
3. \mathcal{H} does not contain an infinite molecular subset.

Note that the equivalence $2. \Leftrightarrow 3.$ and the implication $1. \Rightarrow 2.$ were established in [EH21a, Theorem 1.2 and Lemma 3.6]; the new content of the theorem is the implication $2 \& 3. \Rightarrow 1.$

Corollary 6.1.4. *Let \mathcal{H} be an infinite set of (isomorphism classes of) finite, connected, vertex-transitive graphs. If \mathcal{H} is sparse, then it has the supercritical concentration property.*

This corollary is new even for $\mathcal{H} = \mathcal{F}_d$, the set of all finite, simple, connected, vertex-transitive graphs with degrees bounded by d . (Again, we stress that for bounded degree graphs the giant can exist and fail to be unique or have concentrated density *at the critical point*, as part of a discontinuous phase transition [EH21a, Example 5.1].)

Locality of the density

Now that we know that $\theta(p, G) = \mathbb{P}_p^G(o \in K_1)$ accurately captures the density of the supercritical giant for large non-molecular finite transitive graphs with high-probability, we would like to understand the nature of the dependency of this quantity on p and G . In particular, we would like to understand whether we can compute the density of the giant in terms of some appropriate local limit object as the size of our graph diverges. Before providing our complete answer to this question, let us introduce the analogous statements in the simpler setting of infinite transitive graphs.

Infinite graphs The most important open problem in the theory of percolation on infinite transitive graphs is to establish that the map $\theta(\cdot, G)$ (i.e. $p \mapsto \theta(p, G)$) is continuous for every fixed $G \in \mathcal{G}^*$, where \mathcal{G}^* is the set of all infinite transitive graphs that are not quasi-isometric to \mathbb{Z} . This map is always continuous on $[0, 1] \setminus \{p_c(G)\}$ [Sch99], but continuity at $p_c(G)$ remains open for many graphs, including $G = \mathbb{Z}^3$. In this paper, we are interested in the orthogonal question: Endow \mathcal{G}^* with the local (Benjamini-Schramm) topology. Now is the map $\theta(p, \cdot)$ continuous for every fixed p ? This variant of continuity arises naturally in the study of percolation on large finite transitive graphs (see below). More generally, is the map $\theta : [0, 1] \times \mathcal{G}^* \rightarrow [0, 1]$ jointly continuous, as a function of two variables? In fact, since $\theta(\cdot, G)$ is monotone, this joint continuity is nothing more than continuity in each argument individually.

Underlying both of these versions of continuity is the problem of showing that $\theta(p, G)$ is *uniformly* well-approximated by

$$\theta_n(p, G) := \mathbb{P}_p^G(|K_o| \geq n, o \leftrightarrow \infty).$$

Indeed, let $d \in \mathbb{N}$, let \mathcal{G}_d^* be the (compact) subset of \mathcal{G}^* of graphs with vertex degree d , and consider some $p \in [0, 1]$ and $G \in \mathcal{G}_d^*$. Now by Dini's theorem, the functions $\theta(\cdot, G)$, $\theta(p, \cdot)$, and $\theta(\cdot, \cdot)$ are continuous if and only if $\theta_n \xrightarrow{n \rightarrow \infty} \theta$ uniformly on $[0, 1] \times \{G\}$, $\{p\} \times \mathcal{G}_d^*$, and $[0, 1] \times \mathcal{G}_d^*$, respectively.

An obstacle to establishing the required uniform convergence on $\{p\} \times \mathcal{G}_d^*$ for fixed p is that for some choices of G , the parameter p might be supercritical (i.e. $p > p_c(G)$), whereas for others, p might be subcritical or critical. It is often easier to build arguments that are tailored to studying percolation in just one of these three phases at a time. In this paper, we directly establish that $\theta_n \rightarrow \theta$ uniformly on the supercritical region $\{(p, G) : p \geq p_c(G) + \varepsilon\}$ for every fixed $\varepsilon > 0$. This immediately yields the following statement. (There is no hypothesis that the graphs are not one-dimensional because this condition does not appear in our argument and because the result holds trivially when the graphs are one-dimensional.)

Theorem 6.1.5. *Let $(G_n)_{n \geq 1}$ be a sequence of infinite transitive graphs converging locally to some infinite transitive graph G . Then*

$$\lim_{n \rightarrow \infty} \theta(p, G_n) = \theta(p, G) \quad \text{for all } p > \sup_{n \geq 1} p_c(G_n).$$

This theorem in particular recovers the fact that $p_c(G) \leq \liminf_{n \rightarrow \infty} p_c(G_n)$, and hence that the map $p_c : \mathcal{G}^* \rightarrow [0, 1]$ is lower semi-continuous. This had previously been established as a consequence of the sharpness of the phase transition [Pet14, §14.2].

In [EH23a], we established that $\theta_n \rightarrow \theta$ uniformly on the subcritical region $\{(p, G) : p \leq p_c(G) - \varepsilon\}$ for every fixed $\varepsilon > 0$. By the same reasoning as above, this implies that the map $p_c : \mathcal{G}^* \rightarrow [0, 1]$ is *upper* semi-continuous. Since, as mentioned, p_c was known to be lower semi-continuous, our work established that p_c is continuous. The continuity of p_c had been known as *Schramm's locality conjecture*.

In stark contrast to our proof of Schramm's locality conjecture, our proof of Theorem 6.1.5 is short and handles all (unimodular) transitive graphs via single argument. By combining these two results, we obtain the following concerning the continuity of $\theta(p, \cdot)$. This corollary was previously known in special cases such as when all of the graphs in question have polynomial growth [CMT22]. We were not able to remove the hypothesis that $p \neq p_c(G)$; if we could, then by approximating \mathbb{Z}^3 by toroidal slabs, we could also prove that $\theta(\cdot, \mathbb{Z}^3)$ is continuous by applying the results of [DST14].

Corollary 6.1.6. *Let $(G_n)_{n \geq 1}$ be a sequence of non-one-dimensional infinite transitive graphs converging locally to some infinite transitive graph G . Then*

$$\lim_{n \rightarrow \infty} \theta(p, G_n) = \theta(p, G) \quad \text{for all } p \in [0, 1] \setminus \{p_c(G)\}.$$

Bounded-degree finite graphs Continuity-in- G questions arise naturally in the study of percolation on large, bounded-degree, finite transitive graphs. Indeed, since every infinite set \mathcal{H} of finite transitive graphs with bounded degrees is relatively compact in the local topology, we can for many purposes assume without loss of generality that \mathcal{H} converges locally to some infinite transitive graph G . It is then natural to ask whether $\lim_{H \in \mathcal{H}} \theta(p, H) = \theta(p, G)$ for a fixed p . All of our earlier discussion relating continuity properties of θ to uniform convergence $\theta_n \rightarrow \theta$ can be adapted to finite graphs mutatis mutandis, where we define $\theta_n(p, G) := \mathbb{P}_p^G(|K_o| \geq n, o \notin K_1)$ when G is a finite transitive graph. In particular, our proof of Theorem 6.1.5 yields the following for supercritical percolation on finite transitive graphs.

Theorem 6.1.7. *Let $(G_n)_{n \geq 1}$ be a sequence of finite transitive graphs converging locally to some infinite transitive graph G . Let p_c be a percolation threshold function for $\{G_n : n \geq 1\}$. Then*

$$\lim_{n \rightarrow \infty} \theta(p, G_n) = \theta(p, G) \quad \text{for all } p > \limsup_{n \rightarrow \infty} p_c(G_n).$$

In [Eas24], the first author combined [EH23a; Eas22; EH21a] to prove an analogue of Schramm's locality conjecture for finite graphs. The trouble here is that the non-one-dimensionality condition becomes more subtle. We need to exclude finite graphs are long and thin, like a circle. Formally, let $\text{dist}_{\text{GH}}(S^1, \frac{\pi}{\text{diam } G} G)$ denote the Gromov-Hausdorff distance between the unit circle S^1 and the

graph metric on G after it has been rescaled by $\frac{\pi}{\text{diam } G}$, where $\text{diam } G$ is the diameter of G . Now we would like to exclude sequences of graphs for this distance tends to zero. (In fact, for the following theorem, we may allow this distance to tend to zero, so long as it tends to zero not too quickly (i.e. faster than $\frac{e^{(\log \text{diam } G)^{1/9}}}{\text{diam } G}$.) By combining this finite graph locality result with Theorem 6.1.5, we obtain the following result describing the asymptotic behaviour of θ on large bounded-degree finite transitive graphs.

Corollary 6.1.8. *Let $(G_n)_{n \geq 1}$ be a sequence of finite transitive graphs converging locally to some infinite transitive graph G . Suppose that*

$$\inf_{n \geq 1} \text{dist}_{\text{GH}} \left(S^1, \frac{\pi}{\text{diam } G_n} G_n \right) > 0.$$

Then

$$\lim_{n \rightarrow \infty} \theta(p, G_n) = \theta(p, G) \quad \text{for all } p \in [0, 1] \setminus \{p_c(G)\}.$$

High degree finite graphs Let us now consider an infinite set \mathcal{H} of finite transitive graphs with $\lim_{G \in \mathcal{H}} \deg G = \infty$, such as the sequence of complete graphs or the sequence of hypercubes. We would like to again relates the asymptotic behaviour of θ to some kind of local limit object for \mathcal{H} . The problem is that since \mathcal{H} has diverging vertex degrees, \mathcal{H} cannot converge in the local topology. So the continuity-in- G questions above do not readily extend. On the other hand, the (equivalent) questions about the uniform convergence of $\theta_n(p, G) \rightarrow \theta(p, G)$ do! For example, the analogue of continuity-in- G is that for a given sequence of parameters $p : \mathcal{H} \rightarrow [0, 1]$,

$$\lim_{n \rightarrow \infty} \sup_{G \in \mathcal{H}} |\theta(p, G) - \theta_n(p, G)| = 0. \quad (6.1.1)$$

By considering a step-by-step exploration of the cluster at o , it is easy to see that for each fixed n , as $G \rightarrow \infty$ in \mathcal{H} , the probability $\theta_n(p, G)$ tends to the probability that a $\text{Poi}(p \cdot \deg G)$ branching process contains n vertices. In particular, writing $\text{mf}(\lambda)$ for the survival probability of a $\text{Poi}(\lambda)$ branching process, (6.1.1) is equivalent to having the following *mean-field approximation*:

$$\lim_{G \in \mathcal{H}} |\theta(p, G) - \text{mf}(p \deg G)| = 1.$$

Thus the analogue of the conclusion of Theorem 6.1.7 is that this mean-field approximation holds whenever p is a supercritical sequence of parameter. In this paper, we will characterise for which families of finite transitive graphs this approximation holds. Since a $\text{Poi}(\lambda)$ -branching process is critical at $\lambda = 1$, the analogue of the locality of the critical point is that $G \mapsto \frac{1}{\deg G}$ is the percolation threshold for \mathcal{H} , under some non-one-dimensionality hypothesis on G . We have nothing to say here about this question.

Unfortunately, the mean-field approximation does not always hold for supercritical percolation on high-degree finite transitive graphs. For example, consider the product $K_n \times (\mathbb{Z}/n\mathbb{Z})^2$ of a complete graph with a two-dimensional torus. It is easy to see that the failure of the mean-field approximation in this example extends to any family of graphs that can be automorphism-invariantly decomposed into a collection of dense graphs by deleting $O(|V|)$ edges. We call such graphs *macromolecular*. Here is the precise definition. Note that every molecular graph is macromolecular.

Definition 6.1.9. We say that a finite transitive graph $G = (V, E)$ is ε -*macromolecular*, where $\varepsilon > 0$, if there exists an $\text{Aut } G$ -invariant set of edges $F \subseteq E$ with $\varepsilon |F| \leq |V|$ such that $(V, E \setminus F)$ is not connected, and the connected component (V', E') of o in $(V, E \setminus F)$ satisfies $|E'| \geq \varepsilon |V|^2$.

We say that an infinite set \mathcal{H} of finite transitive graphs is *macromolecular* $\lim_{\mathcal{H}} \deg G = \infty$ and there exists a constant $\varepsilon > 0$ such that all but finitely many of the graphs in \mathcal{H} are ε -macromolecular.

We will show that macromolecular graphs are in fact the *only* obstacles to the mean-field approximation for high degree finite transitive graphs.

Theorem 6.1.10. *Let \mathcal{H} be a set of (isomorphism classes of) finite, connected, vertex-transitive graphs, and suppose that $\deg(G) \rightarrow \infty$ as $G \rightarrow \infty$ in \mathcal{H} . Then the mean-field approximation*

$$\lim_{G \in \mathcal{H}} |\theta(p, G) - \text{mf}(p \deg G)| = 1,$$

holds for every supercritical assignment $p : \mathcal{H} \rightarrow [0, 1]$ if \mathcal{H} does not have any infinite subsets that are macromolecular. The converse also holds, provided that \mathcal{H} admits least one supercritical assignment $p : \mathcal{H} \rightarrow [0, 1]$ that is non-trivial in the sense that $\liminf_{\mathcal{H}} \theta(p, G) < 1$.

What governs the asymptotic density θ for supercritical percolation on a high-degree macromolecular graph? We can think of percolation on the product of a complete graph with a torus as roughly “simulating percolation” on the torus, by first revealing the states of the edges in the complete graphs — producing a unique giant cluster in each complete graph — then revealing the remaining edges. This perspective can be used to show that

$$\theta(\lambda/n, K_n \times (\mathbb{Z}/n\mathbb{Z})^2) \sim \theta(\lambda^*, (\mathbb{Z}/n\mathbb{Z})^2) \quad \text{as } n \rightarrow \infty,$$

where $\lambda^* = (1 - e^{-\lambda}) \text{mf}(\lambda)^2$. In Section 6.9, we will show that in fact for all macromolecular graphs, the supercritical density θ is determined “locally” by the graph obtained by contracting each dense graph to a vertex. Note that this contraction can produce a torus (as above), and similarly any bounded-degree graph, but also large complete graphs and hypercubes. In particular,

the macromolecular case is at least as hard as the general bounded-degree case and the complete graph and hypercube cases, combined. However, crucially, this contraction cannot produce a macromolecular sequence of graphs.

Equicontinuity

Our results also yield the following general result about the equicontinuity of the supercritical giant cluster density on any finite transitive graph. Recall that \mathcal{F} is the set of all finite transitive graphs.

Theorem 6.1.11. *Let $p : \mathcal{F} \rightarrow [0, 1]$ be supercritical. For each $G \in \mathcal{F}$, consider the function $f_G : [p, 1] \rightarrow [0, 1]$ given by $f_G(q) := \theta(q/p, G)$. Then the family of functions $\{f_G : G \in \mathcal{F}\}$ is uniformly equicontinuous.*

By specialising to families of bounded-degree infinite transitive graphs, we also obtain the following. (This time, we can use q in place of q/p , since p is bounded away from zero.)

Theorem 6.1.12. *Let \mathcal{H} be an infinite set of infinite transitive graphs with uniformly bounded vertex degrees. Let $\varepsilon > 0$, and for each $G \in \mathcal{H}$, consider the function $f_p : [p, 1] \rightarrow [0, 1]$ given by $f_G(q) := \theta(q, G)$. Then the family of functions $\{f_G : G \in \mathcal{H}\}$ is uniformly equicontinuous.*

Proof sketch

Consider supercritical percolation \mathbb{P}_p on a large finite transitive graph G that is not molecular, and say that we are tasked with proving that the giant cluster density is concentrated. The natural subgoal is to prove the existence of a uniform tail on the distribution of non-giant clusters, i.e.

$$\mathbb{P}_p \left(\underbrace{|K_o| \geq n \text{ but } o \notin \text{giant}}_{\heartsuit} \right) \xrightarrow{n \rightarrow \infty} 0 \quad \text{uniformly in } G. \quad (6.1.2)$$

Indeed, all of the following implications are rather easy to prove:

$$\text{Uniform tail} \iff \text{Locality} \implies \text{Equicontinuity} \implies \text{Concentration}.$$

See Section 6.11 for details. Unfortunately, both of the missing reverse implications in this chain are false. So our strategy to prove concentration by first proving eq. (6.1.2) is doomed to fail in general. Let us anyway see how far this plan will take us. We split into cases according to whether the vertex degrees of G are bounded above or tending to infinity.

Bounded degrees

Our starting point is the elementary observation that the required uniform tail on finite clusters holds for typical choices of p when some neighbour u of o happens to belong to the giant, i.e.

$$\mathbb{P}_p \left(\underbrace{|K_o| \geq n \text{ but } o \notin \text{giant} \text{ and } u \in \text{giant}}_{\diamond} \right) \xrightarrow{n \rightarrow \infty} 0 \quad \text{uniformly in } G,$$

because a simple mass-transport argument shows that

$$\mathbb{P}_p (\diamond) \leq \frac{\theta'(p, G)}{n},$$

where $\theta'(p, G)$ denotes derivative with respect to p , and this derivative must trivially be bounded above for most choices of p because $p \mapsto \theta(p, G)$ is an increasing function taking values in $[0, 1]$. In principle, this approach could fail because the parameter p that we were given was carefully chosen to belong a small set of parameters where this derivative is large. This turns out not to be a serious obstacle: if we could prove that eq. (6.1.2) holds whenever this derivative is bounded, we could quite easily deduce that eq. (6.1.2) holds in general. The real challenge is therefore to prove that for all $\varepsilon > 0$ there exists $\delta > 0$ such that

$$\mathbb{P}_p (\heartsuit) \geq \varepsilon \implies \mathbb{P}_p (\diamond) \geq \delta. \quad (6.1.3)$$

Here is an initial, naive approach: by an easy exploration argument (in which we first reveal the edges incident to K_o , then reveal the remaining edges), with probability at least

$$\mathbb{P}_p (\heartsuit) \cdot \theta(p, G)$$

there is an edge uv in the boundary of K_o whose outer endpoint v belongs to the giant and whose inner endpoint u satisfies $|K_u| \geq n$ but does not belong to the giant. The problem is that we do not have any control over the location of this edge uv . If only we knew that uv could be found inside of some bounded-size, deterministic set of edges F ...

An important theme in boolean analysis is that an increasing boolean function typically undergoes a sharp threshold¹ *if and only if* it does not depend too much on a small number of bits. In percolation theory, the ‘only if’ direction is typically applied to show that some connectivity event of interest does undergo a sharp threshold. Our new idea is to apply the ‘if’ direction to the event

$$\{o \in \text{giant}\}$$

¹See Section 6.3 for full background, including the definition of sharp thresholds.

which — in the cases of interest — trivially does *not* undergo a sharp threshold. This guarantees the existence of a bounded-size, deterministic set of edges F such that if F is entirely open, then o is highly likely to belong to the giant cluster. Using this and insertion tolerance, we are able to fix the naive argument above to show a suitable edge uv can indeed be found within F with good probability, as required.

Unbounded degrees

There are two problems with the above approach when it comes to graphs with large degrees. The first is the use of insertion tolerance: typically, we will be working with parameters that on the order of $1/\deg G$, and in particular, close to 0. Therefore, the cost to open all but one of the edges in F is no longer constant. This turns out to be solvable by applying an adaptation of the sprinkling and surgery argument we introduced in the previous paper of the series to prove Theorem 6.2.1. This argument, which relies heavily on transitivity, lets us reduce the size of F to a singleton, as required, by paying only a constant cost. A slightly technical point is that for this argument to work, which involves repeatedly sprinkling an unboundedly large number of times, it is helpful to work with an event that is *increasing*. For this reason, in section Section 6.5, we will show that bounds on the distribution of non-giant clusters (“germs”) — which may not be monotone in p — can be converted into bounds on the probability that the giant cluster does not intersect given large deterministic sets (“holes”) — which clearly is monotone in p , and vice versa.

The second problem is that the argument naively yields an upper bound on

$$\clubsuit_1 := \mathbb{P}_p (|K_o| \geq n \deg G, o \notin \text{giant})$$

for a large constant n , rather than on

$$\clubsuit_2 := \mathbb{P}_p (|K_o| \geq n, o \notin \text{giant}).$$

This is a more substantial problem because for macromolecular graphs, we can indeed have that $|K_o|$ is on the order of $\deg G$ with good probability, without belonging to the giant cluster. We will define a suitable automorphism-invariant relation on the vertices on G , called “being friends”, where, roughly speaking, u and v are friends if it is unlikely that the giant cluster contains many of the vertices in the neighbourhood of u but not of v , or vice versa. Then, we show that if we *cannot* convert our control on \clubsuit_1 into a control on \clubsuit_2 , then the equivalence relation given by the transitive closure of this “friend” relation must actually induce a macromolecular decomposition.

Thus, we can split into two cases: either we have the required control on \clubsuit_2 , or our graph is macromolecular. In the former case we are done as in the bounded-degree case. In the

macromolecular case, we relate percolation on the graph G to percolation on the quotient graph obtained by contracting each equivalence class of G to a point. This essentially reduces the problem back to the bounded-degree, since clusters of size much larger than constant in this quotient graph correspond to clusters of size much larger than $\deg G$ in the original graph G , which we are able to control by \clubsuit_1 . More precisely, we use this coupling to show that the density $\theta(p, G)$ on G can be related to the microscopic cluster distribution in the quotient graph, and thus that this density is uniformly equicontinuous. From this equicontinuity, we then obtain concentration via Theorem 6.2.1.

Note that a key challenge and novelty arising in our work, in contrast to previous works investigating percolation on general graphs under some isoperimetric conditions, is that we do not obtain a general bound on the size of microscopic clusters. Indeed, at our level of generality, such a bound does not exist. (Relatedly, we cannot give a non-trivial upper bound on the variance of $\|\text{giant}\|$ when this density is concentrated around its mean, since by approximating molecular graphs, the rate of convergence could be arbitrarily small.) This leads us to use a expansion-vs-structure dichotomy — either we have good isoperimetric properties and therefore have sufficiently strong control of small clusters, or we have some non-trivial rigid structure (being macromolecular) that we can exploit. This is similar in spirit to the strategy often used in the study of probability on general infinite transitive graphs, for example in our proof of Schramm’s locality conjecture, which uses an expansion-vs-structure dichotomy with graphs with rapid volume growth having good expansion, and graphs with slow volume growth having approximately the structure of the Cayley graph of a nilpotent group (thanks to Gromov’s theorem about groups of polynomial growth and related results).

Remark 6.1.1. In our first draft of this paper, we were only able to prove our results for families of graphs uniformly bounded vertex degrees. Our proof was quite different and less elementary, combining the second author’s two-ghost inequality with a standard application sharp threshold theory - namely, using Talagrand’s inequality to prove that certain connectivity events do have sharp thresholds. Realising that these methods could be relevant to Schramm’s locality conjecture, we paused on this project to complete our proof of Schramm’s conjecture in [EH23a]. These original methods make up [EH23a, Section 2], and we sketched in [EH23a, Section 7.2] how these methods can be used to establish uniqueness, concentration, locality, and equicontinuity for bounded-degree finite transitive graphs.

However, these methods break down completely when trying to analyse the supercritical giant cluster density for finite transitive graphs with large vertex degrees. More precisely, these methods

fail to prove locality and equicontinuity as soon as graphs have even slowly diverging vertex degrees, and they fail to prove uniqueness and concentration once the vertex degrees grow algebraically with respect to the total number of vertices. Indeed, it is clear that these methods must fail in general because they do not detect whether a graph is macromolecular.

6.2 Notation

Graphs Let G be a *graph*. For us, this means that G is an isomorphism class of connected, simple², locally finite, countable graphs. When the choice of graph (or multigraph) G is clear from context, we write V and E for the vertex and edge sets of G ; when it is not clear we instead write $V(G)$ and $E(G)$ to be explicit. We will adopt the convention that the *volume* of G is the number of vertices in G , denoted $|G| := |V(G)|$. If G is (vertex-)transitive, then we write o to denote some fixed vertex of G , which we refer to as the *origin*. We write $\text{dist}(\cdot, \cdot)$ for the graph metric on G and define $\text{diam } G := \max_{u,v} \text{dist}(u, v)$. Given $u \in V$ and $n \geq 0$, we define $B_n(u) := \{v \in V : \text{dist}(u, v) \leq n\}$, and when G is transitive, we set $B_n := B_n(o)$.

Classes of graphs We write \mathcal{G} and \mathcal{F} to denote the set of all infinite and finite transitive graphs, respectively. Later, we will also write \mathcal{U} and \mathcal{N} for set of all unimodular and nonunimodular transitive graphs, respectively. The definition of unimodularity is recalled in Section 6.4. We say that an infinite transitive graph $G \in \mathcal{G}$ is *one-dimensional* if G is roughly-isometric to \mathbb{Z} , and we define $\mathcal{G}^* := \{G \in \mathcal{G} : G \text{ is not one-dimensional}\}$. For each $d \geq 1$, we write \mathcal{G}_d to denote the set of all graphs in \mathcal{G} having vertex degree exactly d , and we analogously define $\mathcal{F}_d, \mathcal{U}_d, \mathcal{N}_d$ and \mathcal{G}_d^* . We endow the whole of $\mathcal{G} \cup \mathcal{F}$ with the *local* (aka *Benjamini-Schramm*) topology. This obviously makes \mathcal{G}_d compact and \mathcal{F}_d relatively compact (for each $d \geq 1$), but in fact, the spaces \mathcal{G}_d^* , \mathcal{U}_d , and \mathcal{N}_d are also compact.

Dense graphs Let $G = (V, E)$ be a finite graph. Let $\varepsilon > 0$. We say that G is *ε -dense* if $|E| \geq \varepsilon |V|^2$. If G is transitive, this is equivalent to $\deg G \geq \frac{\varepsilon}{2} |V|$, where $\deg G$ denotes the common vertex degree of G . We say that an infinite set of finite graphs \mathcal{H} is *dense* if there exists $\varepsilon > 0$ such that all but finitely many graphs in \mathcal{H} are ε -dense, and we say that \mathcal{H} is *sparse* if instead $\lim_{G \in \mathcal{H}} \frac{|E(G)|}{|V(G)|^2} = 0$. Note that every infinite set of finite graphs must contain an infinite subset that is sparse or dense.

²This assumption can be replaced throughout the paper with the assumption that there are a bounded number of edges between any two vertices. In particular, all our results about *bounded degree* graphs do not really require the assumption of simplicity.

Macromolecular graphs Let $G = (V, E)$ be a finite transitive graph. We say that a pair (A, B) is a *macromolecular decomposition* for G if there exists a non-trivial equivalence relation³ on V that is invariant under the action of $\text{Aut } G$ such that A is the subgraph of G induced by the equivalence class containing o , and $B = G/\sim$ is the multigraph obtained from G by contracting each equivalence class to a (distinct) vertex. Given $\varepsilon > 0$, an ε -*macromolecular decomposition* for G is a macromolecular decomposition (A, B) such that A is ε -dense and $\frac{|E(B)|}{|V(G)|} \leq \frac{1}{\varepsilon}$. Say that G is ε -*macromolecular* if G admits an ε -macromolecular decomposition. We say that an infinite set $\mathcal{H} \subseteq \mathcal{F}$ is *macromolecular* if $\lim_{G \in \mathcal{H}} \deg G = \infty$ and there exists $\varepsilon > 0$ such that all but finitely many graphs in \mathcal{H} are ε -macromolecular. The next lemma says that after passing to a subsequence, there is an essentially unique best way to choose the macromolecular decompositions for these graphs. Note that an infinite set of finite transitive graphs is *molecular* (defined in the introduction) if and only if it is both dense and macromolecular. Given $\varepsilon > 0$, say that G is ε -*molecular* if G is ε -macromolecular and ε -dense, so that an infinite set \mathcal{H} of finite transitive graphs is molecular if and only if there exists $\varepsilon > 0$ such that all but finitely many graphs $G \in \mathcal{H}$ are ε -molecular.

Subgraphs Let $G = (V, E)$ be a graph. Given $u, v \in V$, we write uv for the unordered pair $\{u, v\}$, write $u \sim v$ to mean that $uv \in E$, and define $\text{neigh}(u) := \{v \in V : u \sim v\}$. Given a set of vertices X and a single vertex u , we write $uX := \{uv \in E : v \in X\}$. Given X that is a subgraph of G or a subset of V , we write ∂H or $\partial_E H$ for the **edge boundary** $\partial X := \{xy \in E : x \in X \text{ and } y \notin X\}$. Given a set of vertex $W \subseteq V$, we write $G[W]$ for the subgraph of G induced by W and $E[W]$ for the edge set $E(G[W])$. Given two finite subsets W_1 and W_2 of V or a finite subsets of V , we define the **density**⁴ of W_1 in W_2 to be

$$\|W_1\|_{W_2} := \frac{|W_1 \cap W_2|}{|W_2|}.$$

When $W_2 = V$, then we simply write $\|W_1\|$ to mean $\|W_1\|_G$. We will also apply the same notation to *subgraphs* of G , in which case the density of one subgraph in another is defined to be the density of the vertex set of one subgraph in the vertex set of the other.

Configurations and clusters Let G be a graph and let $\omega : E \rightarrow \{0, 1\}$ be a configuration. We say that e is open or closed according to whether $\omega(e) = 1$ or $\omega(e) = 0$. We think of ω as encoding the spanning subgraph of G with edge set $\omega^{-1}(1)$. So $\omega \cup X$ and $\omega \cap X$ are defined by the above conventions for $X \subseteq G$, $X \subseteq V$, or $X \subseteq E$. Given $u, v \in V$, we write $u \leftrightarrow v$ or $u \xleftrightarrow{\omega} v$ to mean that

³‘Non-trivial’ means that there is more than one equivalence class.

⁴This definition is unrelated to the notion of *dense graphs* from earlier.

u and v are ω -connected, i.e. there is a path in the graph encoded by ω from u to v . We also write $v \leftrightarrow \infty$ to mean that there are infinitely many vertices u satisfying $u \leftrightarrow v$.

Given $v \in V$, we define K_v to be the subgraph of ω spanned by $\{u \in V : u \leftrightarrow v\}$ and call this the *cluster* at v . We enumerate the clusters of ω by K_1, K_2, \dots in such a way that $|K_1| \geq |K_2| \geq \dots$. This enumeration is not well-defined when there are multiple clusters of the same volume. In this case, we are happy to break ties arbitrarily. So to avoid this (unimportant) technicality, let us now once and for all fix such a choice of enumeration K_1, K_2, \dots with $|K_1| \geq |K_2| \geq \dots$ of the clusters in every (countable, locally finite, simple) graph.

Now suppose that G is finite. In our setting, K_1 will typically be the unique macroscopic cluster. So we will sometimes use the more suggestive notation $\text{giant} := K_1$. Since $\omega \mapsto \text{giant}(\omega)$ is not an increasing map (with respect to inclusion), for technical reasons we will typically work with the following proxies: for each $\varepsilon > 0$, define giant_ε to be the subgraph of ω induced by the set of vertices v satisfying $\|K_v\| \geq \varepsilon$. Note that the equation ‘ $\text{giant} = \text{giant}_\varepsilon$ ’ is a convenient way to say that there exists a unique cluster with density at least ε .

Percolation \mathbb{P}_p with varying p Let G be a graph and let $p \in [0, 1]$. We write $\mathbb{P}_p = \mathbb{P}_p^G$ for the law of a random configuration $\omega : E \rightarrow \{0, 1\}$ where every $\omega(e)$ is iid $\text{Bernoulli}(p)$. We will write $\mathbb{P} = \mathbb{P}^G$ for the law of the canonical monotone coupling $(\omega_q)_{q \in [0,1]}$ of $(\mathbb{P}_q)_{q \in [0,1]}$. Let \mathcal{E} be an event and suppose that $p \in (0, 1)$. Then we define $\mathbb{P}'_p(\mathcal{E})$ to be the derivative (if it exists) of the map $q \mapsto \mathbb{P}_q(\mathcal{E})$ evaluated at $q = p$.

Densities and supercriticality Let G be a transitive graph and let $p \in [0, 1]$. If G is infinite, then we set $\theta(p) = \theta(p, G) := \mathbb{P}_p(o \leftrightarrow \infty)$. Now suppose that G is finite. We define $\theta(p) = \theta(p, G) := \mathbb{P}_p(o \in \text{giant})$. Moreover, for each $\varepsilon > 0$, we define $\theta_\varepsilon(p) = \theta_\varepsilon(p, G) := \mathbb{P}_p(o \in \text{giant}_\varepsilon)$, and we say that p is ε -supercritical if $\theta_\varepsilon((1 - \varepsilon)p) \geq \varepsilon$. So for an infinite set \mathcal{H} of finite transitive graphs and an assignment of parameters $p : \mathcal{H} \rightarrow [0, 1]$, the assignment p is supercritical (as defined in the introduction) if and only if there exists $\varepsilon > 0$ such that for all but finitely many graphs $G \in \mathcal{H}$, the parameter $p(G)$ is ε -supercritical for G or $p = 1$.

In addition to Theorem 6.1.1, we use the following technical result from our previous work [EH21a, Proposition 4.1 and Remark 4.2] characterising the so-called *sharp density property*, which we now recall. For all $G \in \mathcal{F}$ and $\beta, \varepsilon \in (0, 1]$, we define $p_c^G(\beta, \varepsilon)$ to be the unique parameter satisfying $\mathbb{P}_{p_c^G(\beta, \varepsilon)}(\|K_1\| \geq \beta) = \varepsilon$, which is well-defined by continuity and strict monotonicity of percolation

on G . Now we say that an infinite set $\mathcal{H} \subseteq \mathcal{F}$ has the *sharp-density property* if

$$\lim_{G \in \mathcal{H}} \sup_{\beta \in [\alpha, 1]} \frac{p_c^G(\beta, 1 - \delta)}{p_c^G(\beta, \delta)} = 1 \quad \text{for all } \alpha \in (0, 1) \text{ and } \delta \in (0, 1/2].$$

This property in particular implies that the event $\{\|K_1\| \geq \beta\}$ undergoes a sharp-threshold for every fixed $\beta > 0$ when considering percolation on graphs in \mathcal{H} .

Theorem 6.2.1. *Let \mathcal{H} be an infinite set of finite transitive graphs. Then \mathcal{H} has the sharp-density property if and only if \mathcal{H} does not contain an infinite molecular subset.*

Percolation thresholds We will use the notation p_c both for the usual infinite-cluster percolation threshold $p_c(G) = \inf\{p : \theta(p) > 0\}$ when G is an infinite transitive graph *and* for a percolation threshold $p_c : \mathcal{H} \rightarrow [0, 1]$ when \mathcal{H} is an infinite set of finite transitive graphs (when \mathcal{H} admits such a threshold), as defined in [Eas22]. It is well-known that $p_c(G) \geq \frac{1}{\deg G - 1}$ for every infinite transitive graph G . Let us record here the analogous result for finite transitive graphs, which can be proven by the same argument as [EH21a, Lemma 2.8] or [Eas22, Proposition 5].

Lemma 6.2.2. *Let \mathcal{H} be an infinite set of finite transitive graphs. If $p : \mathcal{H} \rightarrow [0, 1]$ is a supercritical sequence of parameters, then $\liminf_{G \in \mathcal{H}} (\deg G - 1)p(G) \geq 1$. In particular, if $p_c : \mathcal{H} \rightarrow [0, 1]$ is a percolation threshold, then $\liminf_{G \in \mathcal{H}} (\deg G - 1)p_c(G) \geq 1$.*

It is also well-known that every infinite transitive graph G satisfies the mean-field lower bound $\theta((1 + \varepsilon)p_c(G)) \geq \frac{\varepsilon}{1 + \varepsilon}$ for all $\varepsilon > 0$. This was first proven by Menshikov [Men86] and Aizenman and Barsky [AB87b]; various alternative simplified proofs are now available [DT16b; Hut20c; DRT19; Van22a]. For future use, we state a version of this bound applying also to sets of *finite* graphs without infinite molecular subsets. This was proven in [Eas22, Lemma 10] using an adaptation of the methods of [Van22a] together with the uniqueness theorem of the first paper in this series [EH21a].

Theorem 6.2.3. *Let \mathcal{H} be an infinite set of finite transitive graphs that does not contain an infinite subset that is molecular, and let $p_c : \mathcal{H} \rightarrow [0, 1]$ be a percolation threshold function (which exists by [Eas22]). For all $\varepsilon \in (0, \infty)$ and all $\delta \in (0, \frac{\varepsilon}{1 + \varepsilon})$,*

$$\lim_{G \in \mathcal{H}} \mathbb{P}_{(1 + \varepsilon)p_c} (\|K_1\| \geq \delta) = 1.$$

6.3 Applications of Hatami's theorem to percolation

Suppose for each $n \geq 1$ that A_n is an increasing event depending on some finite number N_n of independent random bits, each of which is 1 with probability p and 0 with probability $1 - p$. Recall that the sequence of events $(A_n)_{n \geq 1}$ is said to have a **sharp threshold** if there exists a sequence p_1, p_2, \dots of numbers in $(0, 1)$ such that

$$\lim_n \mathbb{P}_{(1+\varepsilon)p_n}(A_n) = 1 \quad \text{and} \quad \lim_n \mathbb{P}_{(1-\varepsilon)p_n}(A_n) = 0$$

for every $\varepsilon > 0$ (in which case this holds with $p_n = \min\{p \in [0, 1] : \mathbb{P}_p(A_n) = 1/2\}$). In this section, we outline a new application of sharp threshold theory to percolation using theorems due to Friedgut [Fri98], Bourgain [Fri99b], and Hatami [Hat12b], which have previously been overlooked by the percolation community.

Let us first give some further background and definitions. *Russo's formula* [Rus82] states that if A is an increasing event depending on finitely many bits then

$$\frac{d}{dp} \mathbb{P}_p(A) = \sum_{e \in E} \mathbb{P}_p(e \text{ is pivotal for } A),$$

where an edge e is said to be **pivotal** for the (increasing) event A if $\omega \setminus \{e\} \notin A$ and $\omega \cup \{e\} \in A$. The probability that e is pivotal for A is known as the **influence** of e and is denoted by $I_p(A, e)$. Talagrand's inequality [Tal94] states that

$$\frac{d}{dp} \mathbb{P}_p(A) \geq c \left[p(1-p) \log \frac{2}{p(1-p)} \right]^{-1} \mathbb{P}_p(A)(1 - \mathbb{P}_p(A)) \log \frac{1}{p(1-p) \max_e I_p(A, e)},$$

so that events with small maximal influence must have large total influence. As explained in [EH21a, Section 3.4], Talagrand's inequality implies that increasing events that depend on n bits in a sufficiently symmetric way automatically have sharp thresholds *provided that the threshold occurs at a value of p that is subalgebraically small in n* . This condition cannot be removed in general, since the existence of a triangle in the Erdős-Rényi random graph $G(n, p)$ has a coarse threshold at p of order λ/n . (Other powers of n can be obtained from the event that the Erdős-Rényi graph contains a complete subgraph on k vertices.) While this application of Talagrand's inequality is a standard part of the theory of percolation on bounded-degree and infinite graphs, we would like to bring attention to a related but less well-known result due to Friedgut [Fri98]. Friedgut showed that if p is bounded away from 0 and 1, and $\frac{d}{dp} \mathbb{P}_p(A)$ is bounded above, then in fact A can be approximated *arbitrarily well* by events determined by a bounded number of edges (where the bound on the number of edges depends on the accuracy of the approximation). Friedgut's theorem

is the key to proving our results about bounded-degree graphs, where parameters of interest are always bounded away from 0 and 1.

When p is algebraically small in n , the absence of edges of large influence is no longer sufficient to ensure a sharp threshold, and Friedgut's theorem does not apply, but one can instead look to Bourgain [Fri99b] for a result of a similar flavour that is applicable. Bourgain showed that the absence of a sharp threshold must imply the existence of sets of edges of bounded size $(F_n)_{n \geq 1}$ for which $\omega \cup F_n$ is more likely to belong to A_n than ω by at least a constant (for infinitely many n and appropriate choices of the parameter p). When the threshold occurs at a subalgebraically small value of p , the conclusions of Bourgain's theorem are strictly weaker than those of Talagrand, so that Bourgain's theorem is not relevant to percolation on bounded degree graphs. In this paper, we will make use not of Bourgain's theorem but rather a powerful strengthening of this theorem due to Hatami [Hat12b]. Roughly speaking, this theorem allows us to replace the conclusion that $\omega \cup F_n$ is more likely to belong to A_n than ω by a constant with the conclusion that $\omega \cup F_n \in A_n$ with *arbitrarily high probability*. While this distinction may seem minor at first glance, it is in fact very significant. When p is bounded away from 0 and 1, Hatami's theorem's is equivalent to Friedgut's theorem from the point of view of our applications. However, to allow for smaller values of p , which will be necessary in later sections, we will always refer to Hatami's theorem.

We now state Hatami's theorem. For each positive integer n , we write $[n]$ for the set $\{1, \dots, n\}$, and for each parameter $p \in (0, 1)$ write μ_p for the law of a random variable $x = (x_i)_{i \in [n]}$ where the x_i 's are i.i.d. $\text{Bernoulli}(p)$.

Theorem 6.3.1 (Hatami 2012). *Let $n \in \mathbb{N}$, let $f : \{0, 1\}^{[n]} \rightarrow \{0, 1\}$ be non-constant and increasing, and let $p \in (0, 1/2]$. For every $\varepsilon > 0$, there exists a set $S \subseteq [n]$ such that*

$$\mu_p(f(x) \mid x_i = 1 \ \forall i \in S) \geq 1 - \varepsilon \quad \text{and} \quad |S| \leq \exp\left(10^{12} \frac{\lceil J_f(p) \rceil^2}{\varepsilon^2 \mu_p(f(x))^2}\right),$$

where $J_f(p) := 2p(1-p) \frac{d\mu_p(f)}{dp}$.

We will apply this theorem via the following corollary. We write $\mathbb{P}'_p(\mathcal{E})$ to denote $\frac{d\mathbb{P}_p(\mathcal{E})}{dp}$.

Corollary 6.3.2. *Let G be a finite graph, let $p \in (0, 1)$, and let $\varepsilon_1, \varepsilon_2, \varepsilon_3 \in (0, 1)$. If $\mathcal{E} \subseteq \{0, 1\}^E$ is a non-trivial increasing event such that*

$$\mathbb{P}_p(\mathcal{E}) \geq \varepsilon_1 \quad \text{and} \quad p\mathbb{P}'_p(\mathcal{E}) \leq \frac{1}{\varepsilon_2},$$

then there is a set of edges $F \subseteq E$ such that

$$\mathbb{P}_p(\omega \cup F \in \mathcal{E}) \geq 1 - \varepsilon_3 \quad \text{and} \quad |F| \leq \exp\left(\frac{2^{50}}{(\varepsilon_1 \varepsilon_2 \varepsilon_3)^2}\right).$$

Proof of Corollary 6.3.2. The corollary is essentially immediate when $p \leq 1/2$. To allow for the case that $p > 1/2$, we simulate percolation of parameter p on G by percolation of a parameter $\phi(p) \leq 1/2$ on a multigraph with more edges. This argument does not contain any surprises and can safely be skipped on a first reading.

Fix $p \in (1/2, 1)$ and let $H = (V, F)$ be the multigraph formed from G by replacing each edge e by $\ell := \lceil \log_2 \frac{1}{1-p} \rceil$ parallel edges e_1, \dots, e_ℓ with the same endpoints as e . For each configuration ω on H , let $\Phi(\omega)$ be the configuration on G in which an edge e is open if and only if there is some $1 \leq i \leq \ell$ for which e_i is open in ω . For each $q \in (0, 1)$, define

$$g(q) := \mathbb{P}_q^G(\mathcal{A}), \quad h(q) := \mathbb{P}_q^H(\Phi^{-1}(\mathcal{A})), \quad \text{and} \quad \phi(q) := 1 - (1 - q)^{1/\ell},$$

so that $g = h \circ \phi$ and the pushforward measure $\Phi_* \mathbb{P}_{\phi(q)}^H$ satisfies $\Phi_* \mathbb{P}_{\phi(q)}^H = \mathbb{P}_q^G$. By our choice of ℓ , we know that $\phi(p) \in (0, 1/2]$, allowing us to apply Theorem 6.3.1. This guarantees that there is a set of edges $T \subseteq F$ such that

$$\mathbb{P}_{\phi(p)}^H(\omega \cup T \in \Phi^{-1}(\mathcal{A})) \geq 1 - \varepsilon_3 \quad \text{and} \quad |T| \leq \exp\left(10^{12} \frac{\lceil I \rceil^2}{\varepsilon_3^2 (h \circ \phi(p))^2}\right), \quad (6.3.1)$$

where

$$I = 2\phi(p)(1 - \phi(p))(h' \circ \phi(p)).$$

Since $\Phi_* \mathbb{P}_{\phi(p)}^H = \mathbb{P}_p^G$, we know that $\mathbb{P}_p^G(\omega \cup \Phi(T) \in \mathcal{A}) \geq 1 - \varepsilon_3$. Since $g = h \circ \phi$, we have by the chain rule that $h' \circ \phi = g'/\phi'$. We also have by direct calculation that $(\phi'(p))^{-1} = \ell(1 - p)^{(\ell-1)/\ell}$, and since $\log_2 \frac{1}{1-p} \leq \ell \leq 2 \log_2 \frac{1}{1-p}$ we deduce that

$$\frac{1}{\phi'(p)} \leq 4(1 - p) \log_2 \frac{1}{1 - p} \leq 4.$$

Together with the simple observations that $\phi(p) \leq p$ and $1 - \phi(p) \leq 1$, this lets us simplify our bound on $|T|$ from eq. (6.3.1) to

$$|T| \leq \exp\left(10^{12} \frac{\lceil 4 \cdot 2 \cdot p \cdot g'(p) \rceil^2}{\varepsilon_3^2 \cdot g(p)^2}\right).$$

The result now follows since $\lceil x \rceil \leq 2x$ for all $x \geq 1$, $|\Phi(T)| \leq |T|$ and, by hypothesis, $g(p) \geq \varepsilon_1$ and $p \cdot g'(p) \leq \frac{1}{\varepsilon_2}$. \square

The key advantage of this result over the usual sharp-threshold results used in percolation theory is that we can guarantee that $\omega \cup F \in \mathcal{E}$ occurs with with *arbitrarily high probability*. In particular, we can force this event to have a good-probability intersection with any other event \mathcal{B} such that $\mathcal{B} \setminus \mathcal{E}$ has good probability; usually we will apply this in the case $\mathcal{B} \subseteq \mathcal{E}^c$. (On the other hand, this method will tend to give estimates that are very poor quantitatively.)

Given an increasing event \mathcal{E} , a set of edges F , and a configuration $\omega \in \{0, 1\}^E$, we say that F is an **activator** for E if $\omega \notin \mathcal{E}$ but $\omega \cup F \in \mathcal{E}$, and define

$$\text{Act}_{\mathcal{E}}[F] := \{\omega \cup F \in \mathcal{E}\} \cap \{\omega \notin \mathcal{E}\}$$

to be the event that F is an activator for \mathcal{E} . We also write $\text{Act}_{\mathcal{E}}[e] := \text{Act}_{\mathcal{E}}[\{e\}]$ when $e \in E$ is a single edge, so that $\text{Act}_{\mathcal{E}}[e]$ is the event that e is closed and pivotal for \mathcal{E} .

Corollary 6.3.3. *For all $\varepsilon > 0$, there exists $N < \infty$ such that the following holds. Let G be a finite graph. Let \mathcal{E} be an increasing event. Let $p \in (0, 1)$. Suppose that*

$$\varepsilon \leq \mathbb{P}_p(\mathcal{E}) \leq 1 - \varepsilon \quad \text{and} \quad p\mathbb{P}'_p(\mathcal{E}) \leq \frac{1}{\varepsilon}.$$

Let \mathcal{B} be any event with $\mathbb{P}_p(\mathcal{B} \setminus \mathcal{E}) \geq \varepsilon$. Then there is a set of edges $F \subseteq E$ with $|F| \leq N$ such that

$$\mathbb{P}_p(\mathcal{B} \cap \text{Act}_{\mathcal{E}}[F]) \geq \frac{\varepsilon}{2}.$$

Proof. This follows from Corollary 6.3.2 applied with $\varepsilon_1 = \varepsilon_2 = \varepsilon$ and $\varepsilon_3 = \varepsilon/2$. Indeed, letting F be as in the statement of that corollary, we have that

$$\mathbb{P}_p(\mathcal{B} \cap \text{Act}_{\mathcal{E}}[F]) \geq \mathbb{P}_p((\mathcal{B} \setminus \mathcal{E}) \cap (\omega \cup F \in \mathcal{E})) \geq \varepsilon - \frac{\varepsilon}{2} \geq \frac{\varepsilon}{2}. \quad \square$$

A strategy to show that events have low probability. This corollary leads to the following general strategy for showing that an event \mathcal{B} is unlikely under \mathbb{P}_p :

1. Find a suitable increasing event \mathcal{E} .
2. Justify that for one (or many) parameters $q \leq p$, the expression $q\mathbb{P}'_q(\mathcal{E})$ is bounded above.
3. Show that at such q , it is unlikely that *both* $\text{Act}_{\mathcal{E}}[F]$ and \mathcal{B} hold simultaneously, where F is a deterministic bounded-size set of edges.
4. Use Corollary 6.3.3 to deduce that at such q , the event \mathcal{B} is unlikely.

5. Finally, show that because \mathcal{B} is unlikely at q , it is therefore also unlikely at p .

This is to be interpreted as a loose strategy rather than a detailed recipe: each step will involve arguments that are specific to the event at hand. We implement a strategy in the same spirit in Sections 6.4 and 6.7.

We conclude this section by making note of the following elementary lemma, which states that we always have many good parameters q where $q\mathbb{P}'_q$ is bounded above. Here \mathcal{L} denotes the Lebesgue measure.

Lemma 6.3.4. *Let $G = (V, E)$ be a graph. Fix $p, \delta \in (0, 1)$ such that the interval $I := [p, (1+\delta)p]$ is contained in $(0, 1]$. If $\mathcal{E} \subseteq \{0, 1\}^E$ is an increasing event that is determined by the state of finitely many edges then*

$$\mathcal{L}\left(\left\{p \in I : p\mathbb{P}'_p(\mathcal{E}) \leq \frac{1}{\varepsilon}\right\}\right) \geq \left(1 - \frac{2\varepsilon}{\delta}\right) \mathcal{L}(I)$$

for every $\varepsilon > 0$.

Proof. Let $J = \{p \in I : p\mathbb{P}'_p(\mathcal{E}) \leq \varepsilon^{-1}\}$. Since \mathcal{E} is increasing and determined by the state of finitely many edges, the map $q \mapsto \mathbb{P}_q(\mathcal{E})$ is a polynomial that defines an increasing function $[0, 1] \rightarrow [0, 1]$. It follows that

$$\frac{1}{\varepsilon(1+\delta)p} \mathcal{L}(I \setminus J) \leq \int_J \mathbb{P}'_q(\mathcal{E}) dq \leq \int_I \mathbb{P}'_q(\mathcal{E}) dq = \mathbb{P}_{(1+\delta)p}(\mathcal{E}) - \mathbb{P}_p(\mathcal{E}) \leq 1.$$

We deduce the claim by rearranging, using that $\mathcal{L}(I) = \delta p$ and that $\delta^{-1} + 1 \leq 2\delta^{-1}$. \square

6.4 The unimodular bounded-degree case

In this section we prove our main theorems in the case that all graphs in our family are unimodular and have uniformly bounded degrees. (Finite transitive graphs are always unimodular, so this restriction is only relevant for infinite graphs.) All the theorems in this section are special cases of those established in later sections; we present this case separately since it is much less technical and will be the primary case of interest to many readers. It will also serve as a warm-up to the general case, with several of the ideas used in this section appearing again in a more technical context later in the paper. This section also contains all our results about *infinite* graphs.

For every $d \geq 2$, let us fix a consistent choice of percolation threshold function $p_c : \mathcal{F}_d \rightarrow [0, 1]$, which exists by Theorem 6.1.2. In the following proposition, we write $p_c(G)$ to denote the value

of this percolation threshold function at G when G is finite and to denote the usual infinite cluster percolation threshold for G when G is infinite. Moreover,

$$\theta(p, G) := \begin{cases} \mathbb{E}_p \|\text{giant}\| & \text{if } G \text{ is finite} \\ \mathbb{P}_p(o \leftrightarrow \infty) & \text{if } G \text{ is infinite.} \end{cases} \quad (6.4.1)$$

Proposition 6.4.1. *Let \mathcal{H} be an infinite set of (finite or infinite) transitive graphs converging locally to some infinite transitive graph G_∞ . For every constant $p > \limsup_{\mathcal{H}} p_c(G)$,*

$$\lim_{\mathcal{H}} \theta(p, G) = \theta(p, G_\infty).$$

We now give a brief overview of (non)unimodularity, the mass-transport principle and the modular function, referring the reader to [LP16b; Hut20g] for further background. Let $G = (V, E)$ be a connected, locally finite, transitive graph, and let $\text{Aut}(G)$ be the automorphism group of G . The **modular function** $\Delta = \Delta_G : V^2 \rightarrow (0, \infty)$ is defined by

$$\Delta(u, v) = \frac{|\text{Stab}_v u|}{|\text{Stab}_u v|},$$

where $\text{Stab}_v = \{\gamma \in \text{Aut}(G) : \gamma v = v\}$ is the stabilizer of v and $\text{Stab}_v u = \{\gamma u : \gamma \in \text{Stab}_v\}$ is the orbit of u under Stab_v . We say that G is **unimodular** if $\Delta \equiv 1$ and that G is **nonunimodular** otherwise. Every finite transitive graph is unimodular, as is every Cayley graph and every amenable transitive graph [SW90].

The mass-transport principle. Let G be a connected, locally finite, unimodular transitive graph. The **mass-transport principle** states that

$$\sum_{x \in V} F(o, x) = \sum_{x \in V} F(x, o) \quad (6.4.2)$$

for every $F : V^2 \rightarrow [0, \infty]$ that is invariant under the diagonal action of $\text{Aut}(G)$ on V^2 , meaning that $F(\gamma x, \gamma y) = F(x, y)$ for every $x, y \in V$ and $\gamma \in \text{Aut}(G)$. Note that this identity holds trivially when G is finite.

Proof of Proposition 6.4.1

In this subsection we use Hatami's theorem to prove Proposition 6.4.1. We will deduce this proposition from a chain of simple lemmas, several of which do not require the bounded degree assumption. We begin with the following immediate consequence of Corollary 6.3.3.

Lemma 6.4.2. *For every $\varepsilon > 0$ there exists $N < \infty$ such that the following holds. Let G be a graph, let o be a vertex of G , let $p \in (0, 1)$, and let m, n be integers with $m < n$. If*

$$\mathbb{P}_p(m \leq |K_o| < n) \geq \varepsilon, \quad \mathbb{P}_p(|K_o| \geq n) \geq \varepsilon, \quad \text{and} \quad p\mathbb{P}'_p(|K_o| \geq n) \leq \frac{1}{\varepsilon}$$

then there exists a set of edges $F \subseteq E$ with $|F| \leq N$ such that

$$\mathbb{P}_p(\{|K_o| \geq m\} \cap \text{act}_{|K_o| \geq n}[F]) \geq \frac{\varepsilon}{2}.$$

Proof. This follows from Corollary 6.3.3 with $\mathcal{E} := \{|K_o| \geq n\}$ and $\mathcal{B} := \{m \leq |K_o| < n\} \subseteq \mathcal{E}^c$. \square

We next turn this estimate concerning activators into one concerning *pivots* at a cost of $p^{|F|}/2|F|$. This estimate is much less wasteful in the bounded degree case, where the relevant values of p will not be small, than it is in the high degree case.

Lemma 6.4.3. *Let G be a graph, let o be a vertex of G , let $p \in (0, 1)$, and let m, n be integers with $m < n$. For each non-empty finite set of edges $F \subseteq E$ there exists an edge $uv \in F$ such that*

$$\mathbb{P}_p(\{|K_u| \geq m\} \cap \text{act}_{|K_u| \geq n}[uv]) \geq \frac{p^{|F|}}{2|F|} \mathbb{P}_p(\{|K_o| \geq m\} \cap \text{act}_{|K_o| \geq n}[F]).$$

Note that the left hand side of the inequality concerns the cluster of u while the right hand side concerns the cluster of o .

Proof. Let $e_1, \dots, e_{|F|}$ be an enumeration of F . Consider a configuration $\omega \in \{|K_o| \geq m\} \cap \text{act}_{|K_o| \geq n}[F]$ and let $i < |F|$ be the maximum index such that $|K_o(\omega \cup \{e_1, \dots, e_i\})| < n$. We can write $e_{i+1} = uv$ in such a way that

$$\omega \cup \{e_1, \dots, e_i\} \in \{|K_u| \geq m\} \cap \text{act}_{|K_u| \geq n}[uv].$$

Thus, by the pigeonhole principle, there exists a non-random index i with an endpoint labelling $e_{i+1} = uv$ such that

$$\mathbb{P}_p(\omega \cup \{e_1, \dots, e_i\} \in \{|K_u| \geq m\} \cap \text{act}_{|K_u| \geq n}[uv]) \geq \frac{1}{2|F|} \mathbb{P}_p(\{|K_o| \geq m\} \cap \text{act}_{|K_o| \geq n}[F]).$$

Since the event $\{\omega \cup \{e_1, \dots, e_i\} \in \{|K_u| \geq m\} \cap \text{act}_{|K_u| \geq n}[uv]\}$ is independent of the restriction of ω to $\{e_1, \dots, e_i\}$ it follows that

$$\begin{aligned} \mathbb{P}_p \left(\{|K_u| \geq m\} \cap \text{act}_{|K_u| \geq n}[uv] \right) \\ \geq \mathbb{P}_p \left(\omega \cup \{e_1, \dots, e_i\} \in \{|K_u| \geq m\} \cap \text{act}_{|K_u| \geq n}[uv] \right) \mathbb{P}_p(\{e_1, \dots, e_i\} \in \omega) \\ \geq \frac{p^{|F|}}{2^{|F|}} \mathbb{P}_p \left(\{|K_o| \geq m\} \cap \text{act}_{|K_o| \geq n}[F] \right) \end{aligned}$$

as required. \square

Lemma 6.4.4. *Let G be a unimodular transitive graph, and let $m < n$ be natural numbers. Then*

$$\mathbb{P}'_p (|K_o| \geq n) \geq m \cdot \mathbb{P}_p \left(\{|K_u| \geq m\} \cap \text{act}_{|K_u| \geq n}[uv] \right)$$

for every $p \in (0, 1)$ and every edge $uv \in E$.

Proof. The event that $|K_o| \geq n$ is fully determined by the state of finitely many edges. As such, the map $q \mapsto \mathbb{P}_q (|K_o| \geq n)$ is differentiable at p , and, by Russo's formula, has derivative given by

$$\begin{aligned} \mathbb{P}'_p (|K_o| \geq n) &= \sum_{e \in E} \mathbb{P}_p \left(\{ |K_o(\omega \setminus e)| < n \} \cap \{ |K_o(\omega \cup e)| \geq n \} \right) \\ &= \frac{1}{1-p} \sum_{e \in E} \mathbb{P}_p \left(\{ \omega(e) = 0 \} \cap \{ |K_o(\omega)| < n \} \cap \{ |K_o(\omega \cup e)| \geq n \} \right) \\ &= \frac{1}{1-p} \sum_{a \in V} \sum_{b \in \text{neigh}(a)} \mathbb{P}_p \left(\{ o \leftrightarrow a \} \cap \text{act}_{|K_a| \geq n}[ab] \right), \end{aligned}$$

where $\text{neigh}(a)$ denotes the set of neighbours of a . Applying the mass-transport principle to exchange the roles of o and a in the last line yields that

$$\begin{aligned} \mathbb{P}'_p (|K_o| \geq n) &= \frac{1}{1-p} \sum_{b \in \text{neigh}(o)} \mathbb{E}_p \left[|K_o| \mathbf{1}(\text{act}_{|K_o| \geq n}[ob]) \right] \\ &\geq \max_{b \in \text{neigh}(o)} \mathbb{E}_p \left[|K_o| \mathbf{1}(\text{act}_{|K_o| \geq n}[ob]) \right] \geq \max_{b \in \text{neigh}(o)} m \cdot \mathbb{P}_p(|K_o| \geq m, \text{act}_{|K_o| \geq n}[ob]), \end{aligned}$$

and the claim follows by transitivity. \square

Lemma 6.4.5. *For all $d \in \mathbb{N}$ and $\varepsilon \in (0, 1)$, there exists $m = m(d, \varepsilon) < \infty$ such that if $G \in \mathcal{U}_d$, $p \in (0, 1)$ and $n \geq m$ are such that $\mathbb{P}_p (|K_o| \geq n) \geq \varepsilon$ then*

$$\mathcal{L} \left\{ q \in [p, 1] : \mathbb{P}_q (m \leq |K_o| < n) > \varepsilon \right\} \leq \varepsilon.$$

Recall that \mathcal{U}_d denotes the set of all (finite or infinite) unimodular transitive graphs with vertex degree at most d .

Proof. It suffices to consider the case $n \geq 2$, the case $n = 1$ being vacuous. Fix $G \in \mathcal{U}_d$, $p \in (0, 1)$, and integers $n \geq m$ with $n \geq 2$ and suppose that $\mathbb{P}_p(|K_o| \geq n) \geq \varepsilon$. We will prove the contrapositive of the claim: There exists a constant $M = M(d, \varepsilon) \geq 2$ such that

$$(\mathcal{L}\{q \in [p, 1] : \mathbb{P}_q(m \leq |K_o| < n) > \varepsilon\} > \varepsilon) \Rightarrow (m \leq M).$$

Since $n \geq 2$, we know by a union bound that

$$\varepsilon \leq \mathbb{P}_p(|K_o| \geq n) \leq \sum_{u \sim o} \mathbb{P}_p(ou \text{ is open}) \leq pd,$$

and hence $p \geq \frac{\varepsilon}{d}$. We may also assume that $p \leq 1 - \varepsilon$, the claim being trivial otherwise. Take $J := \{q \in [p, 1] : q\mathbb{P}'_q(|K_o| \geq n) \leq 2\varepsilon^{-2}\}$. Since $\mathcal{L}([p, 1] \setminus J) \leq \varepsilon$ by Lemma 6.3.4, it suffices to prove that there exists $M = M(d, \varepsilon) < \infty$ such that

$$(\exists q \in J \text{ such that } \mathbb{P}_q(m \leq |K_o| < n) \geq \varepsilon) \Rightarrow (m \leq M). \quad (6.4.3)$$

To this end, suppose that $q \in J$ is such that $\mathbb{P}_q(m \leq |K_o| < n) \geq \varepsilon$. Applying Lemma 6.4.2 with $\mathcal{E} = \{|K_o| \geq n\}$ and $\mathcal{B} = \{m \leq |K_o| < n\} \subseteq \mathcal{E}^c$, we deduce that there exists $N = N(\varepsilon) < \infty$ and a set of edges $F \subseteq E$ with $|F| \leq N$ such that

$$\mathbb{P}_q(\{|K_o| \geq m\} \cap \text{act}_{|K_o| \geq n}[F]) \geq \frac{\varepsilon}{2}.$$

Using Lemma 6.4.3, it follows that there exists an edge $uv \in F$ such that

$$\mathbb{P}_q(\{|K_u| \geq m\} \cap \text{act}_{|K_u| \geq n}[uv]) \geq \frac{q^{|F|}}{2|F|} \cdot \frac{\varepsilon}{2} \geq \frac{\varepsilon^{1+N}}{4Nd^N} =: c$$

where $c = c(d, \varepsilon)$ is a positive constant. Since G is unimodular, we may apply Lemma 6.4.4 to deduce that $\mathbb{P}'_q(|K_o| \geq n) \geq mc$. Contrasting this with the hypothesis that $q \in J$ yields that

$$m \leq \frac{2}{cp\varepsilon^2} \leq \frac{2d}{c\varepsilon^3}.$$

We deduce that the implication (6.4.3) holds with $M = M(d, \varepsilon) = 2d/(c\varepsilon^3)$, completing the proof. \square

In the next lemma, we write \mathcal{L} for the Lebesgue measure. In the context of finite graphs, we used ‘giant’ to denote an arbitrary choice of largest cluster. Let us extend this definition as follows: if there exists at least one infinite cluster, let ‘giant’ be the union of *all* infinite clusters. The reader may notice that ‘giant’ is undefined when there are arbitrarily large finite clusters but no infinite clusters; we will never use the notation ‘giant’ in such a situation.

Lemma 6.4.6. *Let $d \geq 2$ and $\varepsilon > 0$, there exists $m(\varepsilon, \delta) < \infty$ such that for all $G \in \mathcal{U}_d$,*

$$\mathcal{L}\left\{p \in [p_c(G), 1] : \mathbb{P}_p^G(|K_o| \geq m \text{ but } o \notin \text{giant}) \geq \varepsilon\right\} \leq \varepsilon.$$

Proof of Lemma 6.4.6. We start by dealing with the finite graphs in \mathcal{U}_d . By the definition of a percolation threshold, there exists $\delta > 0$ such that $\mathbb{P}_{p_c+\varepsilon/2}(|K_o| \geq \delta |V|) \geq \delta$ for all but finitely many of the finite graphs $G \in \mathcal{U}_d$. Thus, by Lemma 6.4.5, there exists $m \in \mathbb{N}$ such that

$$\mathcal{L}\left\{p \in [p_c(G) + \frac{\varepsilon}{2}, 1] : \mathbb{P}_p(m \leq |K_o| < \delta |V|) \leq \frac{\varepsilon}{2}\right\} \geq 1 - p_c(G) - \varepsilon \quad (6.4.4)$$

for all but finitely many (isomorphism classes of) finite graphs $G \in \mathcal{U}_d$. By increasing m if necessary, we may take this estimate to hold for *every* finite graph $G \in \mathcal{U}_d$. For each finite graph $G \in \mathcal{U}_d$, let $J(G)$ be the set whose Lebesgue measure is bounded in (6.4.4). By Theorems 6.1.1 and 6.2.1, all but finitely many finite graphs $G \in \mathcal{H}$ have the property that

$$\mathbb{P}_p(|K_o| \geq \delta |V| \text{ but } o \notin \text{giant}) \leq \frac{\varepsilon}{2}$$

for every $p \geq p_c(G) + \varepsilon/2$, and hence by a union bound that

$$\mathbb{P}_p(|K_o| \geq m \text{ but } o \notin \text{giant}) \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon$$

for every $p \in J(G)$. This proves the claim for the finite graphs in \mathcal{U}_d since $\mathcal{L}([p_c(G), 1] \setminus J(G)) \leq \varepsilon$ for every $G \in \mathcal{U}_d$.

We now turn to the *infinite* graphs in \mathcal{U}_d . By Theorem 6.2.3, there exists $\delta > 0$ such that

$$\mathbb{P}_{p_c+\frac{\varepsilon}{2}}(|K_o| \geq n) \geq \mathbb{P}_{p_c+\frac{\varepsilon}{2}}(o \leftrightarrow \infty) \geq \delta$$

for every infinite $G \in \mathcal{H}$ and every $n \in \mathbb{N}$. So, by Lemma 6.4.5, there exists $m = m(d, \varepsilon)$ such that the set of parameters

$$J_n = J_n(G) = \left\{p \in \left[p_c + \frac{\varepsilon}{2}, 1\right] : \mathbb{P}_q(m \leq |K_o| < n) \leq \varepsilon\right\}$$

satisfies $\mathcal{L}([p_c + \frac{\varepsilon}{2}, 1] \setminus J_n) \leq \frac{\varepsilon}{2}$ for every infinite $G \in \mathcal{H}$ and every $n \in \mathbb{N}$. Noting that $J_1 \supseteq J_2 \supseteq \dots$ for every infinite $G \in \mathcal{H}$, we deduce that the intersection $J = J(G) = \bigcap_{n \geq 1} J_n$ satisfies

$$\mathcal{L}([p_c, 1] \setminus J) = \frac{\varepsilon}{2} + \lim_{n \rightarrow \infty} \mathcal{L}\left([p_c + \frac{\varepsilon}{2}, 1] \setminus J_n\right) \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon$$

for every infinite $G \in \mathcal{U}_d$. This implies the claim since

$$\mathbb{P}_p(|K_o| \geq m \text{ but } o \nleftrightarrow \infty) = \lim_{n \rightarrow \infty} \mathbb{P}_p(m \leq |K_o| < n) \leq \varepsilon$$

for every infinite $G \in \mathcal{U}_d$ and every $p \in J(G)$. \square

Lemma 6.4.7. *Let \mathcal{H} be an infinite set of (infinite or finite) transitive graphs converging locally to some infinite transitive graph G_∞ . For all $p : \mathcal{H} \rightarrow [0, 1]$ satisfying $\liminf_{\mathcal{H}} p/p_c > 1$, if*

$$\lim_{n \rightarrow \infty} \limsup_{\mathcal{H}} \mathbb{P}_p(|K_o| \geq n \text{ but } o \notin \text{giant}) = 0, \quad (6.4.5)$$

then

$$\lim_{\mathcal{H}} |\theta(p, G) - \theta(p, G_\infty)| = 0.$$

Proof. For all $G \in \mathcal{H}$ and $n \geq 1$,

$$\begin{aligned} \heartsuit &:= |\theta(p, G) - \theta(p, G_\infty)| \\ &\leq \underbrace{|\theta(p, G) - \mathbb{P}_p^G(|K_o| \geq n)|}_{\heartsuit_n^1} + \underbrace{|\mathbb{P}_p^G(|K_o| \geq n) - \mathbb{P}_p^{G_\infty}(|K_o| \geq n)|}_{\heartsuit_n^2} + \underbrace{|\mathbb{P}_p^{G_\infty}(|K_o| \geq n) - \theta(p, G_\infty)|}_{\heartsuit_n^3}. \end{aligned}$$

Since $\liminf_{\mathcal{H}} p/p_c > 1$, for all n ,

$$\lim_{\mathcal{H}} \mathbb{P}_p(|\text{giant}| \geq n) = 1,$$

and hence, by Theorems 6.1.1 and 6.2.1,

$$\limsup_{\mathcal{H}} \heartsuit_n^1 = \limsup_{\mathcal{H}} \mathbb{P}_p^G(|K_o| \geq n \text{ but } o \notin \text{giant}), \quad (6.4.6)$$

which tends to zero as n tends to infinity if eq. (6.4.5) holds. In particular, if eq. (6.4.5) holds then

$$\begin{aligned} \limsup_{\mathcal{H}'} \heartsuit &\leq \limsup_{\mathcal{H}'} \inf_{n \geq 1} (\heartsuit_n^1 + \heartsuit_n^2 + \heartsuit_n^3) \\ &\leq \limsup_{n \rightarrow \infty} \limsup_{\mathcal{H}'} (\heartsuit_n^1 + \heartsuit_n^2 + \heartsuit_n^3) \\ &\leq \underbrace{\limsup_{n \rightarrow \infty} \limsup_{\mathcal{H}'} \heartsuit_n^1}_{=0 \text{ by eq. (6.4.5)}} + \underbrace{\sup_{n \geq 1} \limsup_{\mathcal{H}'} \heartsuit_n^2}_{=0} + \underbrace{\limsup_{n \rightarrow \infty} \heartsuit_n^3}_{=0} = 0. \end{aligned}$$

\square

We are now ready to complete the proof of Proposition 6.4.1. To make the proof slightly shorter, we will invoke the following well-known result of Schonmann [Sch99]: for every infinite transitive graph G , the density $\theta(\cdot, G)$ is continuous on $(p_c(G), 1]$. This is not strictly necessary: we could bypass this step by using Lemma 6.5.1 (adapting its proof to all unimodular transitive graphs) and Corollary 6.10.2 (implying that that G_∞ is unimodular.)

Proof of Proposition 6.4.1. Pick $\varepsilon > 0$ such that $p > \limsup_{\mathcal{H}} p_c(G) + \varepsilon$. By Lemma 6.4.6, we can find sequences $p^1, p^2, p^3 : \mathcal{H} \rightarrow [0, 1]$ such that $p^1 \rightarrow p - \varepsilon$, $p^2 \uparrow p$, and $p^3 \downarrow p$ as $G \rightarrow \infty$ with $G \in \mathcal{H}$, and such that for all $i \in \{0, 1, 2\}$,

$$\lim_{n \rightarrow \infty} \limsup_{\mathcal{H}} \mathbb{P}_{p^i} (|K_o| \geq n \text{ but } o \notin \text{giant}) = 0.$$

By Lemma 6.4.7, for all i ,

$$\lim_{\mathcal{H}} |\theta(p^i, G) - \theta(p^i, G_\infty)| = 0. \quad (6.4.7)$$

By Theorem 6.2.3, we have $\liminf_{\mathcal{H}} \theta(p^1, G) > 0$. So by eq. (6.4.7), $\lim_{\mathcal{H}} p^1 \geq p_c(G_\infty)$ and hence $p > p_c(G)$. In particular, $\theta(\cdot, G_\infty)$ is continuous at p by [Sch99]. So for both $i \in \{1, 2\}$,

$$\lim_{\mathcal{H}} \theta(p^i, G_\infty) = \theta(p, G_\infty). \quad (6.4.8)$$

By combining eqs. (6.4.7) and (6.4.8), for both $i \in \{1, 2\}$,

$$\lim_{\mathcal{H}} \theta(p^i, G) = \theta(p, G_\infty),$$

which yields the desired claim by monotonicity because p is sandwiched between p^1 and p^2 . \square

6.5 Germs and Holes

In this section, let $\mathcal{H} \subseteq \mathcal{F}$ be an infinite set that does not contain any molecular subsequences, and consider some assignment of positive integers $M : \mathcal{H} \rightarrow \mathbb{N}$. Given an assignment of parameters $p : \mathcal{H} \rightarrow [0, 1]$, let $\text{Germ}(p) = \text{Germ}(p, M, \mathcal{H})$ be the statement that

$$\lim_{n \rightarrow \infty} \limsup_{\mathcal{H}} \mathbb{P}_p (|K_o| \geq nM \text{ but } o \notin \text{giant}) = 0,$$

and let $\text{Hole}(p) = \text{Hole}(p, M, \mathcal{H})$ be the statement that

$$\lim_{n \rightarrow \infty} \liminf_{\mathcal{H}} \inf_{\substack{A \subseteq V(G) \\ |A| \geq nM}} \mathbb{P}_p (\|\text{giant}\|_A \geq 1/n) = 1.$$

In both equations, we use the convention that $\inf \emptyset := 1$.

Lemma 6.5.1. *For every supercritical assignment p , if $\text{Germ}(p)$ holds then $\text{Hole}(p)$ holds.*

Lemma 6.5.2. *For every supercritical assignment p and every constant $\varepsilon > 0$, if $\text{Hole}(p)$ holds then $\text{Germ}((1 + \varepsilon)p)$ holds.*

Lemma 6.5.3. *Let G be a graph, let u, v be vertices, and let $p, q \in [0, 1]$. For all $n \in \mathbb{N}$,*

$$\mathbb{P}\left(o \xrightarrow{\omega_q} u \text{ but } |\partial K_u(\omega_p) \cap \partial K_v(\omega_q)| \geq n\right) \leq e^{-n|p-q|}.$$

We next state the universal tightness theorem [Hut21c, Theorem 2.2], which guarantees that the size of the largest intersection of a cluster with a fixed set of vertices is always of the same order as its mean with high probability. This theorem also holds for percolation on *weighted* graphs; we state a special case that is adequate for our purposes.

Theorem 6.5.4 (Universal Tightness). *There exist universal constants $C, c > 0$ such that the following holds. Let $G = (V, E)$ be a (countable, locally finite) graph, let $\Lambda \subseteq V$ be a finite non-empty set of vertices, and let $p \in [0, 1]$ be a parameter. Set $|M| := \max\{|K_v \cap \Lambda| : v \in V\}$. Then*

$$\mathbb{P}_p(|M| \geq \alpha \mathbb{E}_p |M|) \leq C e^{-c\alpha} \quad \text{and} \quad \mathbb{P}_p(|M| \leq \varepsilon \mathbb{E}_p |M|) \leq C\varepsilon$$

for every $\alpha \geq 1$ and $0 < \varepsilon \leq 1$.

Proof of Lemma 6.5.1. For all $\delta > 0$, $G \in \mathcal{H}$, and $A \subseteq V(G)$, by a union bound and by linearity of expectation,

$$\begin{aligned} \mathbb{P}_p(\|\text{giant}\|_A \geq \delta) &\geq \mathbb{P}_p\left(\max_{u \in A} \|K_u\|_A \geq \delta\right) - \mathbb{P}_p\left(\max_{u \in A \setminus \text{giant}} \|K_u\|_A \geq \delta\right) \\ &\geq \mathbb{P}_p\left(\max_{u \in A} \|K_u\|_A \geq \delta\right) - \frac{1}{\delta} \max_{u \in A} \mathbb{P}_p(|K_u| \geq \delta |A| \text{ but } u \notin \text{giant}). \end{aligned}$$

So for all $m, n \in \mathbb{N}$,

$$\begin{aligned} \heartsuit_{m,n} &:= \liminf_{\mathcal{H}} \inf_{\substack{A \subseteq V(G) \\ |A| \geq nM}} \mathbb{P}_p(\|\text{giant}\|_A \geq 1/m) \\ &\geq \underbrace{\liminf_{\mathcal{H}} \inf_{A \subseteq V(G)} \mathbb{P}_p\left(\max_{u \in A} \|K_u\|_A \geq 1/m\right)}_{\heartsuit_m^1} - \underbrace{\limsup_{\mathcal{H}} \max_{u \in V(G)} m \mathbb{P}_p(|K_u| \geq nM/m \text{ but } u \notin \text{giant})}_{\heartsuit_{m,n}^2}. \end{aligned}$$

By Theorems 6.1.1 and 6.2.1, since p is supercritical,

$$\liminf_{\mathcal{H}} \inf_{A \subseteq V(G)} \mathbb{E}_p \left[\max_{u \in A} \|K_u\|_A \right] > 0,$$

and hence by Theorem 6.5.4, $\lim_m \heartsuit_m^1 = 1$. By Theorems 6.1.1 and 6.2.1, for all $m, n \geq 1$,

$$\heartsuit_{m,n}^2 = \limsup_{\mathcal{H}} m \mathbb{P}_p (|K_o| \geq n \text{ but } o \notin \text{giant}).$$

So if $\text{Germ}(p)$ holds then $\lim_n \heartsuit_{m,n}^2 = 0$ for all m . Therefore, if $\text{Germ}(p)$ holds then by monotonicity of $\heartsuit_{m,n}$,

$$\lim_n \heartsuit_{n,n} = \lim_m \lim_n \heartsuit_{m,n} \geq \lim_m \heartsuit_m^1 - \lim_m \lim_n \heartsuit_{m,n}^2 = 1,$$

which implies that $\text{Hole}(p)$ holds. \square

Proof of Lemma 6.5.2. Let $q := (1 + \varepsilon)p$. By Theorems 6.1.1 and 6.2.1, there is a constant $\varepsilon > 0$ such that

$$\lim_{\mathcal{H}} \mathbb{P}_p(\mathcal{N}) = \lim_{\mathcal{H}} \mathbb{P}_q(\mathcal{N}) = 0 \quad (6.5.1)$$

where \mathcal{N} is the complement of the event that there exists a unique cluster K satisfying $\|K\| \geq \varepsilon$. Suppose for contradiction that $\text{Hole}(p)$ holds but $\text{Germ}(q)$ does not. Then we can find an infinite subset $\mathcal{H}' \subseteq \mathcal{H}$, an assignment $N : \mathcal{H}' \rightarrow \mathbb{N}$ with $\lim_{\mathcal{H}'} N/M = \infty$, and a constant $\eta > 0$ such that for all $G \in \mathcal{H}'$,

$$\mathbb{P}_q(N \leq |K_o| < \varepsilon |V|) \geq \eta. \quad (6.5.2)$$

Consider some $G \in \mathcal{H}'$. Trivially,

$$\nu(G) := \mathbb{P}_q(\mathcal{N}) = \mathbb{E} [\mathbb{P}(\omega_q \in \mathcal{N} \mid K_o(\omega_q))],$$

and hence by Markov's inequality,

$$\mathbb{P}(\mathbb{P}(\omega_q \in \mathcal{N} \mid K_o(\omega_q)) \geq 2\nu/\eta) \leq \eta/2.$$

So by eq. (6.5.2) and a union bound, there is a deterministic set Π of possible outcomes for K_o such that

$$\mathbb{P}_q(K_o \in \Pi) \geq \eta/2, \quad (6.5.3)$$

and every $A \in \Pi$ satisfies $N \leq |A| < \varepsilon |V|$ and

$$\mathbb{P}(\omega_q \in \mathcal{N} \mid K_o(\omega_q) = A) \leq 2\nu/\eta. \quad (6.5.4)$$

Consider some $A \in \Pi$. Let \bar{A} be the set of all edges having at least one endpoint in A . On the event that $K_o(\omega_q) = A$, since $|A| < \varepsilon |V|$ and every edge in ∂A is ω_p -closed, we have $\omega_p \in \mathcal{N}$ if and

only if $\omega_p \setminus \overline{A} \in \mathcal{N}$. This is helpful because $\omega_p \setminus \overline{A}$ is independent of the event that $K_o(\omega_q) = A$. So we can rewrite eq. (6.5.4) as

$$\mathbb{P}(\omega_p \setminus \overline{A} \in \mathcal{N}) \leq 2\nu/\eta.$$

In particular, since A and G were arbitrary, by eq. (6.5.1), Hole(p), and a union bound,

$$\limsup_{\mathcal{H}' \atop B \in \Pi} \underbrace{\mathbb{P}_p(\{\omega \setminus \overline{B} \in \mathcal{N}\} \cup \mathcal{N} \cup \{\text{giant} \cap B = \emptyset\})}_{\mathcal{E}(B)} = 0. \quad (6.5.5)$$

Now let $b(A) = b(A, \omega)$ count the number of edges uv with $u \in A$ and $v \in V \setminus A$ such that

$$\left| K_v(\omega \setminus \overline{A}) \right| \geq \varepsilon |V|.$$

Notice that on the event $\mathcal{E}(A)$, at least one such edge must be ω_p -open. So by independence, for all $n \geq 1$,

$$\mathbb{P}_p(\mathcal{E}(A) \mid b(A) \leq n/p) \geq (1-p)^{n/p} \geq e^{-n} > 0. \quad (6.5.6)$$

By contrasting eqs. (6.5.5) and (6.5.6), we must have for every constant n ,

$$\liminf_{\mathcal{H}' \atop A \in \Pi} \mathbb{P}_p(b(A) > n/p) = 1. \quad (6.5.7)$$

For all G and n , by eq. (6.5.3) and by independence,

$$\begin{aligned} \mathbb{P}(b(K_o(\omega_q), \omega_p) > n/p) &\geq \mathbb{P}_p(K_o \in \Pi) \inf_{A \in \Pi} \mathbb{P}(b(A, \omega_p) > n/p \mid K_o(\omega_q) = A) \\ &\geq \frac{\eta}{2} \cdot \inf_{A \in \Pi} \mathbb{P}_p(b(A) > n/p). \end{aligned}$$

So by eq. (6.5.7), for all n ,

$$\liminf_{\mathcal{H}'} \mathbb{P}(b(K_o(\omega_q), \omega_p) > n/p) \geq \eta/2. \quad (6.5.8)$$

For all G and n , whenever $b(K_o(\omega_q), \omega_p) > n/p$, there must be at least $\varepsilon |V|$ vertices u such that the event

$$\mathcal{T}(u) := \left\{ o \xrightarrow{\omega_q} u \right\} \cap \left\{ |\partial K_o(\omega_q) \cap \partial K_u(\omega_p)| \geq n/p \right\} \quad (6.5.9)$$

holds. So by linearity of expectation, for all G and n ,

$$\diamond_{G,n} := \sup_{u \in V} \mathbb{P}(\mathcal{T}(u)) \geq \frac{1}{|V|} \sum_{u \in V} \mathbb{P}(\mathcal{T}(u)) \geq \varepsilon \mathbb{P}(b(K_o(\omega_q), \omega_p) > n/p),$$

and hence by eq. (6.5.8),

$$\inf_{n \geq 1} \liminf_{\mathcal{H}'} \diamond_{G,n} \geq \frac{\varepsilon \eta}{2}. \quad (6.5.10)$$

On the other hand, by Lemma 6.5.3, for all G and n ,

$$\diamond_{G,n} \leq e^{-(q-p)\cdot n/p} = e^{-\varepsilon n},$$

and hence,

$$\limsup_{n \rightarrow \infty} \limsup_{\mathcal{H}'} \diamond_{G,n} \leq \limsup_{n \rightarrow \infty} e^{-\varepsilon n} = 0,$$

contradicting eq. (6.5.10). \square

6.6 Characterizing discrete phase transitions

We now turn our attention to finite transitive graphs of divergent degree. (In fact most of what we do will also apply in the bounded degree case, but has additional technicalities compared to that case.) Our first goal, which we carry out in this section, is to characterise those graphs that have a particularly degenerate kind of percolation phase transition we call a *discrete* phase transition.

Let \mathcal{H} be an infinite set of finite connected transitive graphs. We say that \mathcal{H} has a **discrete** percolation phase transition if $\lim_{\mathcal{H}} \theta(p, G) = 1$ for every supercritical assignment $p : \mathcal{H} \rightarrow [0, 1]$. (Here, the word ‘discrete’ is used since this kind of phase transition, where the density jumps from 0 to 1 over a window of negligible size, is the extreme opposite of a *continuous* phase transition.) When this occurs, all of our claims about concentration and continuity in the supercritical phase are trivial. In this section we prove the following proposition, which gives a necessary and sufficient condition for the supercritical phase to be discrete. This proposition will play an important role in the remainder of our analysis, allowing us to focus our attention on the case that $p_c(G)$ is of order $1/\deg(G)$.

Proposition 6.6.1. *An infinite set $\mathcal{H} \subseteq \mathcal{F}$ has a discrete percolation phase transition if and only if \mathcal{H} admits a percolation threshold function $p_c : \mathcal{H} \rightarrow [0, 1]$ satisfying*

$$\lim_{\mathcal{H}} \frac{p_c(G)}{1 - p_c(G)} \deg G = \infty. \quad (6.6.1)$$

The condition (6.6.1) neatly encapsulates that a family of finite vertex-transitive graphs has a discrete phase transition if it has a percolation threshold function that is always either very close to 1 or much larger than the reciprocal of the degree for all large elements of the family; it includes the case that $p_c(G) = 1$ for all but finitely many $G \in \mathcal{H}$. For an example where $p_c \rightarrow 1$, take the sequence of cycles (\mathbb{Z}_n) , and for an example where p_c remains bounded away from 1 but $p_c/\deg \rightarrow \infty$, take cartesian products of complete graphs and cycles $(K_n \times \mathbb{Z}_{f(n)})$ where $f : \mathbb{N} \rightarrow \mathbb{N}$ with $\lim f = \infty$ is any sufficiently slowly growing sequence. Since every set of finite transitive graphs with a discrete

percolation phase transition admits a threshold function, the fact that the phase transition is discrete can be phrased in terms of the convergence of θ to a step function when rescaled by $p_c(G)$: An infinite set $\mathcal{H} \subseteq \mathcal{F}$ has a discrete percolation phase transition if and only if it admits a percolation threshold function $p_c : \mathcal{H} \rightarrow [0, 1]$ satisfying

$$\lim_{G \in \mathcal{H}} \theta(\lambda p_c(G), G) = \mathbb{1}(\lambda \geq 1)$$

for each constant $\lambda \neq 1$. (The $\lambda = 1$ limit of $\theta(\lambda p_c(G), G)$ is sensitive to the precise choice of threshold function $p_c(G)$.)

Corollary 6.6.2. *For every infinite set $\mathcal{H} \subseteq \mathcal{F}$ and for every supercritical assignment $p : \mathcal{H} \rightarrow [0, 1]$,*

$$\lim_{\mathcal{H}} p \deg G = \infty \implies \lim_{\mathcal{H}} \theta(p, G) = 1.$$

Proof. Assume for contradiction that $\lim_{\mathcal{H}} p \deg G = \infty$ but that for some infinite subset \mathcal{H}' , we have $s := \sup_{\mathcal{H}'} \theta(p, G) < 1$. By passing to an infinite subset if necessary, we may assume without loss of generality that \mathcal{H}' is dense or sparse. In either case, as explained in [Eas22], it follows from Theorem 6.1.2 that \mathcal{H}' admits a percolation threshold $p_c : \mathcal{H}' \rightarrow [0, 1]$. By passing to a further infinite subset if necessary, we may assume without loss of generality that this percolation threshold satisfies either $\lim_{\mathcal{H}'} p_c \deg G = \infty$ or $\sup_{\mathcal{H}'} p_c \deg G < \infty$. In the first case, Proposition 6.6.1 tells us that \mathcal{H}' has a discrete percolation phase transition. Since p is supercritical, it follows that

$$\lim_{\mathcal{H}'} \theta(p, G) = 1,$$

contradicting $s < 1$. In the second case, by Theorem 6.2.3, for every positive constant x ,

$$\liminf_{\mathcal{H}'} P_p \left(\|K_1\| \geq \frac{x}{1+x} \right) = 1,$$

and hence

$$\liminf_{\mathcal{H}'} \theta(p, G) \geq \frac{x}{1+x},$$

which contradicts $s < 1$ when $x > \frac{s}{1-s}$. \square

We start with the easier ‘only if’ direction of Proposition 6.6.1.

Lemma 6.6.3. *Let $\mathcal{H} \subseteq \mathcal{F}$. If \mathcal{H} does not have a percolation threshold, or if it has a percolation threshold $p_c : \mathcal{H} \rightarrow [0, 1]$ satisfying $\liminf_{\mathcal{H}} \frac{p_c(G)}{1-p_c(G)} \deg G < \infty$, then the supercritical phase of \mathcal{H} is not discrete.*

Proof. First suppose that \mathcal{H} does not have a percolation threshold. Then by the main theorem of [Eas22], \mathcal{H} contains an infinite molecular subset \mathcal{H}' . (In fact it contains an infinite m -molecular subset for infinitely many m , but we will not need to use this stronger fact.) Since molecular sets are dense, it follows by the results of [Bol+10b] that there exists $\alpha < \infty$ (in fact $\alpha = 1$) such that $G \mapsto \frac{\alpha}{\deg G}$ is a percolation threshold for \mathcal{H}' . Then the assignment $p : G \mapsto \frac{2\alpha}{\deg G}$ is supercritical for \mathcal{H}' and satisfies

$$\begin{aligned} \limsup_{\mathcal{H}'} \theta(p(G), G) &\leq \limsup_{\mathcal{H}'} \mathbb{P}_p(|K_o| > 1) \\ &\leq \limsup_{\mathcal{H}'} \left[1 - \left(1 - \frac{2\alpha}{\deg G} \right)^{\deg G} \right] < 1. \end{aligned} \tag{6.6.2}$$

Next suppose that \mathcal{H} *does* have a percolation threshold $p_c : \mathcal{H} \rightarrow [0, 1]$ but

$$\liminf_{\mathcal{H}} \frac{p_c(G)}{1 - p_c(G)} \deg G < \infty.$$

Then there exists an infinite subset \mathcal{H}' with

$$\limsup_{\mathcal{H}'} \frac{p_c(G)}{1 - p_c(G)} \deg G < \infty,$$

and hence $\limsup_{\mathcal{H}'} p_c(G) < 1$ and $\limsup_{\mathcal{H}'} p_c(G) \deg G < \infty$. Without loss of generality (i.e. by passing to a further infinite subset if necessary), we may assume that there exist constants $\varepsilon > 0$ and $\alpha < \infty$ such that $\sup_{\mathcal{H}'} p_c(G) < 1 - \varepsilon$ and $\lim_{\mathcal{H}'} p_c(G) \deg G = \alpha$, and either $\lim_{\mathcal{H}'} \deg G = \infty$ or $\sup_{\mathcal{H}'} \deg G < \infty$. If $\lim_{\mathcal{H}'} \deg G = \infty$, then the assignment $p : G \mapsto \frac{2\alpha}{\deg G}$ is supercritical for \mathcal{H}' and satisfies eq. (6.6.2). If $\sup_{\mathcal{H}'} \deg G < \infty$, then the assignment $p : G \mapsto 1 - \varepsilon$ is supercritical for \mathcal{H}' and satisfies

$$\begin{aligned} \limsup_{\mathcal{H}'} \theta(p(G), G) &\leq \limsup_{\mathcal{H}'} \mathbb{P}_p(|K_o| > 1) \\ &\leq \limsup_{\mathcal{H}'} \left[1 - \varepsilon^{\deg G} \right] < 1. \end{aligned}$$

We have shown that in all cases, we can find an infinite subset $\mathcal{H}' \subseteq \mathcal{H}$ and a supercritical assignment $p : \mathcal{H}' \rightarrow [0, 1]$ satisfying $\limsup_{\mathcal{H}'} \theta(p(G), G) < 1$. Then the assignment $\mathcal{H} \rightarrow [0, 1]$ given by

$$G \mapsto \begin{cases} p(G) & \text{if } G \in \mathcal{H}' \\ 1 & \text{if } G \in \mathcal{H} \setminus \mathcal{H}' \end{cases}$$

is supercritical for \mathcal{H} and satisfies $\liminf_{\mathcal{H}} \theta(p(G), G) < 1$. Therefore the phase transition on \mathcal{H} is not discrete. \square

The proof of Proposition 6.6.4 will apply both the *universal tightness theorem* Theorem 6.5.4 and the following *quantitative insertion tolerance* estimate of [EH21a, Proposition 2.6]. Both results will be used again several times later in the paper.

Note that for percolation on *infinite* graphs, “insertion tolerance” usually refers to the fact that $\omega \cup \{e\}$ has law absolutely continuous with respect to that of ω , or equivalently that $\mathbb{P}(\omega(e) = 1 | \omega|_{E \setminus e}) > 0$ almost surely. This statement is much less useful for finite graphs, particularly when p is very small as is typical in the high-degree case. The following proposition gives conditions under which we can force an edge to appear within a given (possibly random) set without changing the probability of a given event too badly.

Proposition 6.6.4 (Quantitative insertion tolerance). *Let $G = (V, E)$ be a finite graph, let $p \in (0, 1)$, and let $F \subseteq E$ be a collection of edges. Let $A \subseteq \{0, 1\}^E$ be an event, let $\eta > 0$ and suppose that for each configuration $\omega \in A$ there is a distinguished subset $F[\omega]$ with $F[\omega] \subseteq F \setminus \omega$ and $|F[\omega]| \geq \eta |F|$. If we define $A^+ := \{\omega \cup \{e\} : \omega \in A \text{ and } e \in F[\omega]\}$ then*

$$\mathbb{P}_p(A^+) \geq \frac{\eta^2}{1-p} \cdot \frac{p|F|}{p|F|+1} \cdot \mathbb{P}_p(A)^2.$$

We now begin to work towards the proof of the converse of Lemma 6.6.3 in earnest. We begin with the following lemma, recalling that $\theta_\varepsilon(p) = \theta_\varepsilon(p, G) := \mathbb{P}_p(\|K_o\| \geq \varepsilon)$.

Lemma 6.6.5. *For each $\varepsilon > 0$ and $\lambda < \infty$, there exists $\delta > 0$, and $C < \infty$ such that the following holds for every $G = (V, E) \in \mathcal{F}$ satisfying $|V| \geq C$ and $\deg G \leq \delta |V|$. If $p \in (0, 1)$ is ε -supercritical and satisfies $p\theta'_\varepsilon(p) \leq \lambda$ then either $p \leq \frac{C}{\deg G}$ or $\theta_\varepsilon(p) > 1 - \varepsilon$.*

Proof of Lemma 6.6.5. Let $G \in \mathcal{F}$. The claim is equivalent to the statement that for each $\varepsilon > 0$ and $\lambda < \infty$, there exists $C = C(\varepsilon, \lambda)$ and $\delta = \delta(\varepsilon, \lambda) > 0$ such that if $p \in (0, 1)$ is ε -supercritical and satisfies $\theta_\varepsilon(p) \leq 1 - \varepsilon$ and $p\theta'_\varepsilon(p) \leq \lambda$ then either $|V| < C$, $\deg G > \delta |V|$, or $p \leq C/\deg G$.

Fix $p \in (0, 1)$ that is ε -supercritical and satisfies $\theta_\varepsilon(p) \leq 1 - \varepsilon$ and $p\theta'_\varepsilon(p) \leq \lambda$. By Corollary 6.3.3, there exists $C_1 = C_1(\varepsilon, \lambda) \geq 1$ and a set of edges $F \subseteq E$ with $|F| \leq C_1$ such that

$$\mathbb{P}_p(\text{act}_{\|K_o\| \geq \varepsilon}[F]) \geq \frac{\varepsilon}{2}.$$

Applying the pigeonhole principle twice, we deduce that there exists $uv \in F$ such that

$$\mathbb{P}_p\left(\{u \leftrightarrow v\} \cap \left\{\|K_v\| \geq \frac{\varepsilon}{2C_1}\right\}\right) \geq \frac{\varepsilon}{4C_1}. \quad (6.6.3)$$

By Theorems 6.1.1 and 6.2.1, since a sparse family of graphs can never contain an infinite macro-molecular subset, there exists $C_2 = C_2(\varepsilon, \lambda) < \infty$ and $\delta = \delta(\varepsilon, \lambda) > 0$ such that if $|V| \geq C_2$ and $\deg G \leq \delta |V|$ then

$$\mathbb{P}_p \left(\frac{\varepsilon}{2C_1} \leq \|K_v\| < \varepsilon \right) \leq \mathbb{P}_p \left(\|K_1\| < \varepsilon \text{ or } \|K_2\| \geq \frac{\varepsilon}{2C_1} \right) \leq \frac{\varepsilon}{16C_1}. \quad (6.6.4)$$

Let N be an ω -connected subset of $\text{neigh}(u)$ of maximal volume (breaking ties using some total order on the set of subsets of $\text{neigh}(u)$ that is chosen in advance). We have that

$$\mathbb{E}_p |N| \geq \mathbb{E}_p |K_1 \cap N| \geq \varepsilon \deg(G),$$

and hence by the universal tightness theorem (Theorem 6.5.4) that there exists a constant $\delta_1 = \delta_1(\varepsilon, \lambda) > 0$ such that

$$\mathbb{P}_p (|N| \leq \delta_1 \deg) \leq \frac{\varepsilon}{16C_1}.$$

It follows from this, by a union bound, that the event

$$\mathcal{E} := \{u \leftrightarrow v\} \cap \{\|K_v\| \geq \varepsilon\} \cap \{|N| \geq \delta_1 \deg\} \cap \left\{ \|K_2\| < \frac{\varepsilon}{2C_1} \right\}$$

satisfies $\mathbb{P}_p(\mathcal{E}) \geq \varepsilon/(8C_1)$ whenever $|V| \geq C_2$ and $\deg G \leq \delta$.

To conclude, it suffices to prove that there exists a constant $C_3 = C_3(\varepsilon, \lambda) < \infty$ such that if $|V| \geq C_2$ and $\deg G \leq \delta |V|$ then $p \leq C_3/\deg G$. Let

$$(*) := \mathbb{P}_p (\text{act}_{\|K_u\| \geq \varepsilon} [uN] \mid \mathcal{E}),$$

where uN denotes the set of edges with one endpoint equal to u and the other in N . We split into two cases according to whether $(*) \geq \frac{1}{2}$ (Case 1) or $(*) < \frac{1}{2}$ (Case 2).

Case 1 Since all the vertices of N are connected in ω , if $\text{act}_{\|K_u\| \geq \varepsilon} [uN]$ occurs then *every* edge in uN is a closed pivotal for the event $\{\|K_u\| \geq \varepsilon\}$. As such, by the proof of Lemma 6.4.4 (i.e. a simple mass-transport argument),

$$\begin{aligned} \theta'_\varepsilon(p) &= \frac{1}{1-p} \sum_{b \sim o} \mathbb{E}_p [|K_o| \mathbb{1}(\text{act}_{\|K_o\| \geq \varepsilon} [ob])] \\ &\geq \frac{\delta_1}{2} \deg G \mathbb{P}_p(\mathcal{E}, \text{act}_{\|K_u\| \geq \varepsilon} [uN]) \\ &\geq \frac{\varepsilon \delta_1}{16C_1} \deg G, \end{aligned}$$

and, since $p\theta'_\varepsilon(p) \leq \lambda$, it follows that

$$p \leq \frac{16C_1 \lambda}{\varepsilon \delta_1 \deg G}.$$

Case 2 Notice that if we start with a configuration belonging to the event $\mathcal{E} \setminus \text{act}_{\|K_u\| \geq \varepsilon}[N]$ and then open an edge in uN , we obtain a configuration in which

$$\delta_1 \deg G \leq |K_u| < \varepsilon |V|.$$

So, by the quantitative insertion-tolerance estimate of Proposition 6.6.4, there exists a constant $\delta_2 = \delta_2(\varepsilon, \lambda) > 0$ such that

$$\mathbb{P}_p(\{\delta_1 \deg G \leq |K_u| < \varepsilon |V|\} \cap \{\|K_v\| \geq \varepsilon\}) \geq \delta_2.$$

It follows from Lemma 6.4.4 that $\theta'_\varepsilon(p) \geq \delta_1 \delta_2 \deg G$, and hence that

$$p \leq \frac{\lambda}{\delta_1 \delta_2 \deg G}. \quad \square$$

Proof of Proposition 6.6.1. The ‘only if’ direction is Lemma 6.6.3. Suppose for contradiction that the ‘if’ direction is false, so we can find an infinite set \mathcal{H} of finite connected transitive graphs with a percolation threshold p_c satisfying

$$\lim_{\mathcal{H}} \frac{p_c(G)}{1 - p_c(G)} \deg G = \infty \quad (6.6.5)$$

and a supercritical sequence p satisfying

$$\liminf_{\mathcal{H}} \theta(p(G), G) < 1.$$

Let \mathcal{H}' be an infinite subset of \mathcal{H} such that $\sup_{\mathcal{H}'} \theta(p(G), G) < 1$. If $\lim_{\mathcal{H}'} p_c(G) = 1$, then the only supercritical sequence for \mathcal{H}' (up to changing finitely many terms in the sequence) would be the constant assignment $G \mapsto 1$, which trivially satisfies $\lim_{\mathcal{H}'} \theta(1, G) = 1$. So by passing to a further infinite subset if necessary, we may assume without loss of generality that $\sup_{\mathcal{H}'} p_c(G) < 1$. Since eq. (6.6.5) holds but $p_c(G)$ is bounded away from 1 on \mathcal{H}' , we know that

$$\lim_{\mathcal{H}'} p_c(G) \deg G = \infty. \quad (6.6.6)$$

By Lemma 6.3.4, we can find a constant $\varepsilon > 0$ and an ε -supercritical sequence q for \mathcal{H}' such that for all but finitely many $G \in \mathcal{H}'$, we have $\theta_\varepsilon(q) \leq 1 - \varepsilon$ and $q\theta'_\varepsilon(q) \leq \frac{1}{\varepsilon}$. By eq. (6.6.6), we know that \mathcal{H}' does not contain an infinite subset of graphs with edge densities $\frac{|E|}{|V|^2}$ uniformly bounded away from zero [Bol+10b] or vertex degrees uniformly bounded above (trivially, since $p_c \leq 1$). Lemma 6.6.5 therefore implies that

$$\lim_{\mathcal{H}'} q\theta'_\varepsilon(q) = \infty,$$

which yields the required contradiction. \square

6.7 Negligibility of large holes

The goal of this section is to prove the following proposition, which states that sets of size significantly larger than the degree have large intersection with the giant with high probability. Recall that when W and A are two finite sets of vertices we write $\|W\|_A = |W \cap A|/|A|$ for the density of W in A , and that $\text{giant}_\varepsilon = \{v : \|K_v\| \geq \varepsilon\}$ denotes the set of vertices whose clusters have density at least ε . Recall that \mathcal{F} denotes the set of all finite transitive graphs.

Proposition 6.7.1. *For every supercritical assignment $p : \mathcal{F} \rightarrow [0, 1]$,*

$$\lim_{n \rightarrow \infty} \liminf_{\mathcal{F}} \inf_{\substack{A \subseteq V(G) \\ |A| \geq n \deg G}} \mathbb{P}_p (\|\text{giant}\|_A \geq 1/n) = 1,$$

where $\inf \emptyset := 1$.

At the core of our proof is an induction argument that is similar to our proof of [EH21a, Proposition 4.1] in the previous paper of this series. We will start by proving Proposition 6.7.1 conditionally on a technical lemma whose proof is deferred to Section 6.7. Before giving the proof of Proposition 6.7.1, let us give an informal overview of the strategy. Notice that the original event $\{\|\text{giant}_\varepsilon\|_A \geq \delta\}$ is increasing. So, if we suppose for contradiction that the proposition fails at some ε -supercritical p , then the conclusion of the proposition also fails to hold for all $q \leq p$. By Lemma 6.5.2, this implies that at every $(1 - \varepsilon)p \leq q \leq p$, the cluster at o is mesoscopic with good probability. Thus, we must have a whole interval I of parameters such that we reach a contradiction if we can show that there is *some* $q^* \in I$ where K_o is unlikely to be a mesoscopic under \mathbb{P}_{q^*} .

To find such a $q^* \in I$, we will use the following lemma to pick a “good” sequence $(p_n)_{n \geq 1}$ within this interval I . We will then sprinkle repeatedly, moving from $\mathbb{P}_{p_{n_1}}$ to $\mathbb{P}_{p_{n_2}}$ to $\mathbb{P}_{p_{n_3}}$ and so forth, where $(p_{n_i})_{i \geq 1}$ is a well-chosen subsequence of $(p_n)_{n \geq 1}$. As we do so, we will deduce a sequence of increasingly strong estimates about measure \mathbb{P}_{p_n} , eventually finding an n where we can prove that $q^* = p_n$ has the desired properties needed for us to obtain a contradiction.

Lemma 6.7.2. *For every $0 < \varepsilon < 1$ there exists $\delta(\varepsilon) > 0$ such that the following holds. Let $G = (V, E) \in \mathcal{F}$ satisfy $|V| \geq \delta^{-1}$. For every ε -supercritical parameter $p \in (0, 1)$, there exists an increasing sequence of parameters $(p_n)_{n \geq 1}$ in $((1 - \frac{\varepsilon}{2})p, p)$ such that the following inequalities hold for every $n \geq 1$:*

1. $p_n \theta'_\varepsilon(p_n) \leq 8\varepsilon^{-1}$;

$$2. \ p_{n+1} - p_n \geq \varepsilon 3^{-n-3} (\deg G)^{-1};$$

$$3. \ \theta_\varepsilon(p_{n+1}) - \theta_\varepsilon(p_n) \leq 2^{-n}.$$

Proof of Lemma 6.7.2. Fix $G = (V, E) \in \mathcal{F}$ and an ε -supercritical parameter $p \in (0, 1)$. By Lemma 6.3.4, the set $J = \{q \in ((1 - \varepsilon/2)p, p) : q\theta'_\varepsilon(q) \leq 8\varepsilon^{-1}\}$ has Lebesgue measure $\mathcal{L}(J) \geq \frac{1}{4}\varepsilon p$. By Lemma 6.2.2, there exists $\delta > 0$ such that if $|V| \geq 1/\delta$ then $p \geq \frac{1}{2\deg}$, so that $\mathcal{L}(J) \geq \frac{\varepsilon}{8\deg}$. Applying the same argument as in the proof of [EH21a, Lemma 4.7] but with ‘ $\theta_\varepsilon(\cdot)$ ’ used in place of ‘ $\mathbb{P}(\|K_1\| \geq \beta)$ ’ yields a sequence $(p_n)_{n \geq 1}$ in J such that $p_{n+1} - p_n \geq 3^{-n-1} \mathcal{L}(J) \geq \varepsilon 3^{-n-3} (\deg G)^{-1}$ and $\theta_\varepsilon(p_{n+1}) - \theta_\varepsilon(p_n) \leq 2^{-n}$ for every $n \geq 1$ as required. (Indeed, this argument works for any non-decreasing function taking values in $[0, 1]$ as explained in the proof of [EH21a, Lemma 4.7].) \square

Let us now introduce some notation that will be useful in the rest of the section. Let $G \in \mathcal{F}$ and let $\varepsilon, \delta > 0$. For each vertex v of G , define the event

$$\text{meso}_v^{\delta, \varepsilon} := \left\{ \frac{\deg}{\delta} \leq |K_v| < \varepsilon |V| \right\}.$$

To lighten notation, we will also adopt the shorthand $\text{act}_x^\varepsilon[F] := \text{act}_{\|K_x\| \geq \varepsilon}[F]$ for the event that a set of edges F is an activator for the event $\{\|K_x\| \geq \varepsilon\}$. The next lemma is the inductive step for our repeated sprinkling argument. It tells us that by sprinkling along our good sequence $(p_n)_{n \geq 1}$, we can iteratively shrink the size of a certain set of edges F while maintaining a good probability that $\text{meso}_o^{\delta, \varepsilon} \cap \text{act}_o^\varepsilon[F]$ occurs. Once F becomes a singleton, we will have found our parameter $q^* = p_n$ with which we can obtain a contradiction.

Lemma 6.7.3 (Shrinking by sprinkling). *Let $p, \varepsilon, \delta \in (0, 1)$ and $G = (V, E) \in \mathcal{F}$. Suppose that p is ε -supercritical and that $(p_n)_{n \geq 1}$ is a sequence in $((1 - \frac{\varepsilon}{2})p, p)$ satisfying conditions 2 and 3 from Lemma 6.7.2. For each (n, k, α) , let $A_n(k, \alpha)$ be the statement that there exists a set of edges F with $|F| = k$ satisfying*

$$\mathbb{P}_{p_n} \left(\text{meso}_o^{\delta, \varepsilon} \cap \text{act}_o^\varepsilon[F] \right) \geq \alpha.$$

Given any $k \geq 2$ and $\alpha > 0$, we can find $N(\varepsilon, k, \alpha) < \infty$ such that following holds if $|V| \geq N \deg G$: For each $n \geq N$, there exists $\beta(\varepsilon, k, \alpha, n) > 0$ such that for all $m > n$,

$$A_n(k, \alpha) \implies A_m(k-1, \beta)$$

We stress that the constants β and N appearing in Lemma 6.7.3 are independent of the choices of G , p , $(p_n)_{n \geq 1}$, and δ . The proof of Lemma 6.7.3 is highly technical, and is deferred to the next subsection. By repeatedly applying Lemma 6.7.3, we obtain the following.

Lemma 6.7.4. *Let $\mathcal{H} \subseteq \mathcal{F}$ be an infinite family of graphs that is sparse. Let $\varepsilon > 0$ be a constant, and suppose that $p : \mathcal{H} \rightarrow (0, 1)$ is an assignment such that for every G , the parameter $p(G)$ is ε -supercritical for G . Then*

$$\lim_{m \rightarrow \infty} \limsup_{\mathcal{H}} \inf_{q \in [(1-\varepsilon)p, p]} \mathbb{P}_q \left(\text{meso}_o^{1/m, \varepsilon} \right) = 0.$$

Proof. Suppose for contradiction that this is not the case. Then there exists a constant $\eta > 0$ such that for all $m \geq 1$, there is an infinite subset $\mathcal{H}_m \subseteq \mathcal{H}$ such that for all $G \in \mathcal{H}_m$, for all $q \in [(1 - \varepsilon)p, p]$,

$$\mathbb{P}_q \left(\text{meso}_o^{1/m, \varepsilon} \right) \geq \eta.$$

Since \mathcal{H} is sparse, we can moreover assume that for all $m \geq 1$ and for all $G \in \mathcal{H}_m$,

$$|G| \geq m \deg G.$$

Consider some $m \geq 1$ and some $G = (V, E) \in \mathcal{H}_m$. We will show that if m is sufficiently large with respect to η and ε , then we can force a contradiction. Let $\delta_1(\varepsilon) > 0$ be the constant from Lemma 6.7.2. Suppose that $m \geq 1/\delta_1$, and let $(p_n)_{n \geq 1}$ be the sequence in $((1 - \varepsilon/2)p, p)$ that is thereby guaranteed to exist. By Corollary 6.3.3, there is a constant $k(\varepsilon, \eta) < \infty$ such that for every $m \geq 1$, there is a set of edges $F \subseteq E$ (which may depend on m) such that $|F| \leq k$ and

$$\mathbb{P}_{p_n} \left(\text{meso}_o^{1/m, \varepsilon} \cap \text{act}_o^\varepsilon[F] \right) \geq \beta_0 := \frac{\min(\varepsilon^2, \eta, \varepsilon/8)}{2}.$$

Equivalently, in the language of Lemma 6.7.3 (where the constant “ δ ” is $1/m$), $A_n(k, \beta_0)$ holds for all n . Let N_0 be the constant “ $N(\varepsilon, k, \beta_0)$ ” from Lemma 6.7.3. Recursively define the sequence $(N_0, \beta_0), (N_1, \beta_1), \dots, (N_{k-1}, \beta_{k-1})$, starting with (N_0, β_0) as already defined, as follows: Suppose that we have defined N_{i-1} and β_{i-1} for some $i \in \{1, \dots, k-1\}$. Then set β_i to be the constant “ $\beta(\varepsilon, k-i+1, \beta_{i-1}, N_{i-1})$ ” from Lemma 6.7.3 and set $N_i := N \wedge N_{i-1} + 1$ where N is the constant “ $N(\varepsilon, k-i, \beta_i)$ ” from Lemma 6.7.3. Note that $M := N_{k-1}$ and $\delta_2 := \beta_{k-1}$ are constants that are entirely determined by ε and η . In particular, we may assume that $m \geq M$, so that for all $i \in \{1, \dots, k-1\}$,

$$|G| \geq N_i \deg G. \tag{6.7.1}$$

Claim 6.7.5. $A_{p_M}(1, \delta_2)$ holds.

Proof of claim. We will use induction on i to prove more generally that $A_{N_i}(k-i, \beta_i)$ holds for every $i \in \{0, \dots, k-1\}$, which yields the claim as the case $i = k-1$. The base case $i = 0$ holds

because $A_n(k, \beta_0)$ holds for every $n \geq 1$, and therefore in particular for $n = N_0$. For the inductive step, suppose that $A_{N_i}(k-i, \beta_i)$ holds for some $i \in \{0, \dots, k-2\}$. Since $N_i \geq N$ where N is the constant “ $N(\varepsilon, k-i, \beta_i)$ ” from Lemma 6.7.3, and since eq. (6.7.1) holds, Lemma 6.7.3 implies that $A_n(k-(i+1), \beta_{i+1})$ holds for every $n > N_i$, and in particular, for $n = N_{i+1}$, as required. \square

Pick an edge $uv \in E$ witnessing the fact that $A_{p_M}(1, \delta_2)$ holds, i.e. such that

$$\mathbb{P}_{p_M} \left(\text{meso}_o^{1/m, \varepsilon} \cap \text{act}_o^\varepsilon[uv] \right) \geq \delta_2.$$

By a union bound, we may assume without loss of generality that the endpoint $u \in uv$ satisfies

$$\mathbb{P}_{p_M} \left(\text{meso}_u^{1/m, \varepsilon} \cap \text{act}_u^\varepsilon[uv] \right) \geq \frac{\delta_2}{2}.$$

Then by Lemma 6.4.4,

$$\theta'_\varepsilon(p_M) \geq \frac{\delta_2 m \deg G}{2}.$$

By Lemma 6.2.2, when m is large, we must have $p_M \geq \frac{1}{2 \deg G}$. Therefore,

$$p_M \theta'_\varepsilon(p_M) \geq \frac{\delta_2 m}{4},$$

which contradicts the first condition enumerated in Lemma 6.7.2 when m is sufficiently large with respect to ε and η . \square

We will now deduce Proposition 6.7.1 from Lemma 6.7.4. Note that this proof remains conditional on Lemma 6.7.3 (through our use of Lemma 6.7.4), which we will prove in the next subsection.

Proof of Proposition 6.7.1. Suppose for contradiction the claimed equation does not hold. For every fixed $G = (V, E) \in \mathcal{F}$,

$$\diamondsuit_{n,G} := \inf_{\substack{A \subseteq V \\ |A| \geq n \deg G}} \mathbb{P}_p (\|\text{giant}\|_A \geq 1/n)$$

is trivially non-decreasing with respect to n . So there must exist a constant $\gamma > 0$ such that for all $n \geq 1$ there is an infinite set $\mathcal{H}_n \subseteq \mathcal{F}$ such that $\diamondsuit_{n,G} \leq 1 - \gamma$ for all $G \in \mathcal{H}_n$. By picking a distinct element from each \mathcal{H}_n , we can build an infinite set $\mathcal{H} \subseteq \mathcal{F}$ such that for all $n \geq 1$,

$$\limsup_{\mathcal{H}} \diamondsuit_{n,G} \leq 1 - \gamma.$$

In particular, in the language of Section 6.5, the set \mathcal{H} does not have property $\text{Hole}(p, d)$ where $d(G) := \deg G$. By passing to an infinite subset of \mathcal{H} , we may assume without loss of generality

that there is a constant $\varepsilon > 0$ such that for every $G \in \mathcal{H}$, $p(G) \in (0, 1)$ and $p(G)$ is 2ε -supercritical for G .

For each $G \in \mathcal{H}$, pick a parameter $q(G) \in [(1 - \varepsilon)p, p]$ that minimises $\mathbb{P}_q(\text{meso}_o^{1/m, \varepsilon})$, which exists by continuity. Note that $q : \mathcal{H} \rightarrow [0, 1]$ is supercritical because $(1 - \varepsilon)p$ is supercritical because $p(G)$ is always 2ε -supercritical for G . By Lemma 6.7.4,

$$\lim_{n \rightarrow \infty} \limsup_{\mathcal{H}} \mathbb{P}_q(\text{meso}_o^{1/m, \varepsilon}) = 0.$$

In particular, by Theorems 6.1.1 and 6.2.1 and since sparse families of graphs cannot contain molecular subsequences,

$$\lim_{n \rightarrow \infty} \limsup_{\mathcal{H}} \mathbb{P}_p(|K_o| \geq md \text{ but } o \notin \text{giant}) = 0,$$

i.e. \mathcal{H} has the Germ(q, d) property. By Lemma 6.5.1, it follows that \mathcal{H} has the Hole(q, d) property. So by monotonicity, \mathcal{H} also has the Hole(p, d) property, a contradiction. \square

Proof of Lemma 6.7.3

Let $G = (V, E) \in \mathcal{F}$, and write d for the vertex degree of G . Let $p, \varepsilon, \delta \in (0, 1)$, and suppose that p is ε -supercritical for G . Suppose that $(p_n)_{n \geq 1}$ is a sequence in $((1 - \frac{\varepsilon}{2})p, p)$ satisfying conditions 2 and 3 from Lemma 6.7.2. Fix $k \geq 2$ and $\alpha > 0$. Let $n < m$ be arbitrary positive integers. Assume that $A_n(k, \alpha)$ holds, and let F be a set of edges witnessing this. Our goal is to find $N(\varepsilon, k, \alpha) < \infty$ and $\beta(\varepsilon, k, \alpha, n) > 0$ such that if $|V| \geq Nd$ and $n \geq N$, then $A_m(k - 1, \beta)$ must hold. Notice that $A_m(r, \beta) \implies A_m(k - 1, \beta)$ for all $r \in \{1, \dots, k - 1\}$, since we can always extend a set of edges witnessing $A_m(r, \beta)$ to a set of edges witnessing $A_m(k - 1, \beta)$ by adding $(k - 1) - r$ many arbitrary edges.

In the following claims, we will write $c_1, c_2, \dots \in (0, 1)$ for small positive constants that are determined by the triple (ε, k, α) . When a claim involves a new constant c_i that has not previously appeared, we are asserting the existence of a constant $c_i(\varepsilon, k, \alpha) \in (0, 1)$ with $c_i \leq \min(c_1, \dots, c_{i-1})$ that would make the claim true. The constant c_1 is introduced in Claim 1, the constant c_2 in Claim 2, and so forth.

Claim 6.7.6. *If $|V| \geq c_1^{-1}d$, then we can find vertices x and y with $xy \in F$ such that $\mathbb{P}_{p_n}(\mathcal{E}_1) \geq c_1$ where*

$$\mathcal{E}_1 := \text{meso}_o^{\delta, \varepsilon} \cap \text{act}_o^\varepsilon[F] \cap \{y \in \text{giant}_\varepsilon\}.$$

Proof. Similarly to in the proof of Lemma 6.6.5, by applying the pigeonhole principle twice, there are vertices x and y with $xy \in F$ such that

$$\mathbb{P}_{p_n} \left(\text{meso}_o^{\delta, \varepsilon} \cap \text{act}_o^\varepsilon[F] \cap \{y \in \text{giant}_{\frac{\varepsilon}{2k}}\} \right) \geq \frac{\alpha}{2k}.$$

By Theorem 6.1.1 and Theorem 6.2.1, there exists $N(\varepsilon, \alpha, k) < \infty$ such that if $|V| \geq Nd$ then

$$\mathbb{P}_{p_n} \left(\|K_1\| \geq \varepsilon \text{ and } \|K_2\| < \frac{\varepsilon}{2k} \right) \geq 1 - \frac{\alpha}{4k}.$$

The conclusion follows by a union bound, with $c_1 := N^{-1} \wedge \frac{\alpha}{4k}$. \square

We will use x and y to denote the vertices whose existence is guaranteed by this claim throughout the rest of the proof. Given a configuration ω and a vertex u , let $S(u) = S(u, \omega)$ be the largest ω -connected subset of $\text{neigh}(u)$, breaking any ties according to an arbitrary deterministic rule.

Claim 6.7.7. *If $|V| \geq c_1^{-1}d$, then $\mathbb{P}_{p_n}(\mathcal{E}_2) \geq c_2$ where*

$$\mathcal{E}_2 := \mathcal{E}_1 \cap \left\{ \|S(x)\|_{\text{neigh}(x)} \geq c_2 \right\}.$$

Proof. Consider any vertex u . Since $p_n \geq (1-\varepsilon)p$ and p is ε -supercritical, $\mathbb{E}_{p_n} \|\text{giant}_\varepsilon\|_{\text{neigh}(u)} \geq \varepsilon$. So by Markov's inequality,

$$\mathbb{P}_{p_n} \left(\|\text{giant}_\varepsilon\|_{\text{neigh}(u)} \geq \frac{\varepsilon}{2} \right) \geq \frac{\varepsilon}{2}.$$

There can never be more than $1/\varepsilon$ clusters that each contains at least $\varepsilon|V|$ vertices. Therefore

$$\mathbb{E}_{p_n} \|S(u)\|_{\text{neigh}(u)} \geq \frac{\varepsilon^3}{4}.$$

So by Theorem 6.5.4, there exists a universal constant $C \in (1, \infty)$ such that

$$\mathbb{P}_{p_n} \left(\|S(u)\|_{\text{neigh}(u)} \leq \frac{\varepsilon^3}{4} \cdot \frac{c_1}{2C} \right) \leq \frac{c_1}{2}.$$

The conclusion follows by a union bound, with $c_2 := \frac{\varepsilon^3}{4} \cdot \frac{c_1}{2C}$. \square

Claim 6.7.8. *If $|V| \geq c_1^{-1}d$ and $n \geq c_3^{-1}$, then either $A_m(1, c_3)$ holds or $\mathbb{P}_{p_n}(\mathcal{E}_3) \geq c_3$ where*

$$\mathcal{E}_3 := \mathcal{E}_2 \setminus \text{Act}_o^\varepsilon[xS(x)].$$

Proof. Suppose that $|V| \geq c_1^{-1}d$ and $\mathbb{P}_{p_n}(\mathcal{E}_3) < c := \frac{c_2^2}{4}$. Our goal is to show that $A_m(1, c)$ holds if we assume that $n \geq N$ for a sufficiently large choice of $N(\varepsilon, \alpha, k)$. Indeed, this would establish the claim with $c_3 := N^{-1} \wedge c$. By a union bound,

$$\mathbb{P}_{p_n}(\mathcal{E}_2 \setminus \mathcal{E}_3) \geq c_2 - \frac{c_2^2}{4} \geq \frac{c_2}{2}.$$

On the event $\mathcal{E}_2 \setminus \mathcal{E}_3$, we always have

$$\|\{\mathbf{u} \in \text{neigh}(x) : \text{act}_o^\varepsilon[x\mathbf{u}]\}\|_{\text{neigh}(x)} \geq \|S(x)\|_{\text{neigh}(x)} \geq c_2.$$

So there exists $\mathbf{u} \in \text{neigh}(x)$ such that

$$\mathbb{P}_{p_n}(\text{meso}_o^{\delta, \varepsilon} \cap \text{act}_o^\varepsilon[x\mathbf{u}]) \geq \frac{c_2}{2} \cdot c_2 = 2c.$$

Let N be the smallest positive integer such that $\frac{1}{2^{N-1}} \leq c$, and assume that $n \geq N$. By the third condition enumerated in Lemma 6.7.2,

$$\theta_\varepsilon(p_m) - \theta_\varepsilon(p_n) \leq \sum_{k=n}^{m-1} \frac{1}{2^k} \leq \frac{1}{2^{N-1}} \leq c.$$

Therefore, by a union bound and the monotone coupling of \mathbb{P}_{p_n} and \mathbb{P}_{p_m} ,

$$\mathbb{P}_{p_m}(\text{meso}_o^{\delta, \varepsilon} \cap \text{act}_o^\varepsilon[x\mathbf{u}]) \geq \mathbb{P}_{p_n}(\text{meso}_o^{\delta, \varepsilon} \cap \text{act}_o^\varepsilon[x\mathbf{u}]) - (\theta_\varepsilon(p_m) - \theta_\varepsilon(p_n)) \geq 2c - c = c,$$

and hence $A_m(1, c)$ holds. \square

Claim 6.7.9. *If $|V| \geq c_4^{-1}d$ and $\mathbb{P}_{p_n}(\mathcal{E}_3) \geq c_3$, then $\mathbb{P}_{p_n}(\mathcal{E}) \geq c_4$ where*

$$\mathcal{E} := \mathcal{E}_1 \cap \{|K_x| \geq c_2 d\}.$$

Proof. By Lemma 6.2.2, there exists $N(\varepsilon, k, \alpha)$ with $N \geq c_1^{-1}$ such that if $|V| \geq N$ then $p_n \geq \frac{1}{2d}$. Assume that $|V| \geq Nd$, which implies that $|V| \geq N$ and $|V| \geq c_1^{-1}d$, and assume that $\mathbb{P}_{p_n}(\mathcal{E}_3) \geq c_3$. Our goal is to find $c(\varepsilon, k, \alpha) \in (0, c_3]$ such that $\mathbb{P}_{p_n}(\mathcal{E}) \geq c$. Indeed, then the claim holds with $c_4 := N^{-1} \wedge c$. On the event \mathcal{E}_3 , we know that $|S(x)| \geq c_2 d$. So we are trivially done, with $c := \frac{c_3}{2}$, if $\mathbb{P}_{p_n}(\mathcal{E}_3 \cap \{\omega \cap xS(x) \neq \emptyset\}) \geq \frac{c_3}{2}$. So by a union bound, we may instead assume without loss of generality that

$$\mathbb{P}_{p_n}(\mathcal{E}_3 \cap \{\omega \cap xS(x) = \emptyset\}) \geq \frac{c_3}{2}.$$

Notice that if we start with a configuration in $\mathcal{E}_3 \cap \{\omega \cap xS(x) = \emptyset\}$, then open any edge in $xS(x)$, we obtain a configuration in \mathcal{E} . Therefore by Proposition 6.6.4 with $F := x \text{neigh}(x)$, $A := \mathcal{E}_3 \cap \{\omega \cap xS(x) = \emptyset\}$, $\eta := c_2$, and $F[\omega] := xS(x)$,

$$\mathbb{P}_{p_n}(\mathcal{E}) \geq \frac{c_2^2}{1-p_n} \cdot \frac{p_n d}{p_n d + 1} \cdot \left(\frac{c_3}{2}\right)^2 \geq c_2^2 \cdot \frac{1/2}{1/2 + 1} \cdot \left(\frac{c_3}{2}\right)^2 =: c.$$

□

For each vertex z , pick a map $\phi_z \in \text{Aut } G$ with $\phi_z(x) = z$, which exists by transitivity, and define $\mathcal{E}_z := \{\omega \circ \phi_z \in \mathcal{E}\}$.

Claim 6.7.10. *If $|V| \geq c_4^{-1}d$ and $\mathbb{P}_{p_n}(\mathcal{E}_3) \geq c_3$, then $\mathbb{P}_{p_n}(\mathcal{F}_1) \geq c_5$ where*

$$\mathcal{F}_1 := \{\|\{z \in V : \mathcal{E}_z \text{ holds}\}\| \geq c_5\}$$

Proof. Assume that $|V| \geq c_4^{-1}d$ and $\mathbb{P}_{p_n}(\mathcal{E}_3) \geq c_3$. By the previous claim and since every ϕ_z is an automorphism,

$$\mathbb{E}_{p_n} \|\{z \in V : \mathcal{E}_z \text{ holds}\}\| = \mathbb{P}_{p_n}(\mathcal{E}) \geq c_4.$$

Therefore the claim follows from Markov's inequality, with $c_5 := \frac{c_4}{2}$. □

For the remaining claims, we will (explicitly) assume that d is sufficiently large in order to help with certain rounding errors. When d is not this large, we will anyway be able to easily conclude directly.

Claim 6.7.11. *If $|V| \geq c_4^{-1}d$ and $d \geq c_6^{-1}$, then for every $\omega \in \mathcal{F}_1$, there exists a collection of disjoint sets of vertices $(\mathcal{B}_i : i \in I)$ such that all of the following hold, where $\mathcal{B} := \cup_{i \in I} \mathcal{B}_i$:*

1. $\omega \in \bigcup_{z \in \mathcal{B}} \mathcal{E}_z$.
2. For each $i \in I$, the set \mathcal{B}_i is ω -connected.
3. $\{\phi_z(y) : z \in \mathcal{B}\}$ is ω -connected (where y is the vertex from Claim 6.7.6).
4. $c_6 \leq \frac{|V|}{|I|d} \leq 4c_2$.
5. $c_6 \leq \frac{|\mathcal{B}_i|}{d} \leq c_2$ for each $i \in I$.

Proof. Define $c_6 := \frac{c_2 c_5 \varepsilon}{6}$. Assume that $|V| \geq c_4^{-1} d$ and $d \geq c_6^{-1}$, and consider $\omega \in \mathcal{F}_1$. For every $z \in V$, if $\omega \in \mathcal{E}_z$, then $\phi_z(y) \in \text{giant}_\varepsilon(\omega)$. There cannot be more than $1/\varepsilon$ clusters that each contains at least $\varepsilon |V|$ vertices. So there exists a set of vertices Y with $\|Y\| \geq c_5 \varepsilon$ such that $\omega \in \bigcup_{z \in Y} \mathcal{E}_z$ and $\{\phi_z(y) : z \in Y\}$ is ω -connected.

Consider a cluster K of ω with $|K| \geq c_2 d$. We claim that K can be partitioned into sets of vertices $A_1 \sqcup \dots \sqcup A_k$ such that $\frac{c_2 d}{3} \leq |A_i| \leq c_2 d$ for every i . Indeed, it suffices to establish this when K is the set of integers $\{1, \dots, N\}$ for some $N \geq c_2 d$. Let $q := \lfloor c_2 d/2 \rfloor$, and write $N = bq + r$ where $b \in \mathbb{N}$ and $r \in \{1, \dots, q-1\}$. Now take the partition into the consecutive intervals $\{1, \dots, q\}, \dots, \{(b-2)q+1, \dots, (b-1)q\}$ and $\{(b-1)q+1, \dots, N\}$. Each interval has size q or $q+r$. Notice that thanks to our choice of c_6 , and since $d \geq c_6^{-1}$, we have $c_2 d \geq 6$, and hence $\lfloor c_2 d/2 \rfloor \geq \frac{c_2 d}{3}$. Therefore,

$$\frac{c_2 d}{3} \leq q \leq q+r \leq 2q \leq c_2 d.$$

For every $z \in Y$, since $\omega \in \mathcal{E}_z$, we know that $|K_z(\omega)| \geq c_2 d$. Thanks to the previous paragraph, by splitting large clusters, there therefore exists an equivalence relation \sim on V that is a refinement of $\overset{\omega}{\leftrightarrow}$ such that for every $z \in Y$, the equivalence class of z under \sim , say $[z]$, satisfies

$$\frac{c_2 d}{3} \leq |[z]| \leq c_2 d.$$

In particular, using the bound $d \geq c_6^{-1}$ for the second inequality, and using the bound $\|Y\| \geq c_5 \varepsilon$ for the third inequality,

$$\min_{z \in Y} |[z]| \geq \frac{c_2 d}{3} \geq \frac{2}{c_5 \varepsilon} \geq \frac{2|V|}{|Y|}.$$

This allows us apply [EH21a, Lemma 4.12] where $(X, Y, \sim) := (V, Y, \sim)$. This implies the existence of a collection of disjoint subsets $(B_i : i \in I)$ of Y such that

$$|I| \geq \frac{|V|}{4 \min_{z \in Y} |[z]|} \geq \frac{|V|}{4c_2 d},$$

and for every $i \in I$, the set B_i is entirely contained in some equivalence class of \sim (which may depend on i) and satisfies

$$|B_i| \geq \frac{|Y|}{2|V|} \min_{z \in Y} |[z]| \geq \frac{c_5 \varepsilon}{2} \cdot \frac{c_2 d}{3} = c_6 d.$$

Since every B_i is entirely contained in some equivalence class of \sim , we also know that $|B_i| \leq c_2 d$, and therefore

$$c_6 \leq \frac{|B_i|}{d} \leq c_2.$$

Since $B_i \cap B_j = \emptyset$ for all $i \neq j$, and every B_i satisfies $|B_i| \geq c_6 d$, we know that $|I| \cdot c_6 d \leq |V|$, and therefore

$$c_6 \leq \frac{|V|}{|I| d} \leq 4c_2.$$

□

Claim 6.7.12. *There exists $\beta(\varepsilon, k, \alpha, n) > 0$ such that if $|V| \geq c_4^{-1}d$, $n \geq c_7^{-1}$, $d \geq c_6^{-1}$, and $\mathbb{P}_{p_n}(\mathcal{F}_1) \geq c_5$, then $A_m(k-1, \beta)$ holds.*

Proof. Assume that $|V| \geq c_4^{-1}d$, $d \geq c_6^{-1}$, and $\mathbb{P}_{p_n}(\mathcal{F}_1) \geq c_5$. Recall that \mathbb{P} denotes the standard monotone coupling of percolation measures. Our goal is to find $N(\varepsilon, k, \alpha) < \infty$ and $\beta(\varepsilon, k, \alpha, n) > 0$ such that if $n \geq N$ then $A_m(k-1, \beta)$ holds. Indeed, then the claim follows with $c_7 := N^{-1} \wedge c_6$. On the event \mathcal{F}_1 ,

$$\|\mathcal{B}\| \geq |I| \cdot \min_{i \in I} \|\mathcal{B}_i\| \geq |I| \cdot \frac{c_6 d}{|V|} \geq \frac{c_6}{4c_2} \geq \frac{c_6}{4}.$$

So there exists a fixed vertex z such that the event $\mathcal{E}_4 := \mathcal{F}_1 \cap \{z \in \mathcal{B}\}$ satisfies $\mathbb{P}_{p_n}(\mathcal{E}_4) \geq \frac{c_5 c_6}{4}$. Let $N(\varepsilon, k, \alpha)$ be the smallest positive integer such that $\frac{1}{2^{N-1}} \leq \frac{c_5 c_6}{8}$, and assume that $n \geq N$. Then as in the proof of Claim 6.7.8, by the third condition enumerated in Lemma 6.7.2, $\theta_\varepsilon(p_m) - \theta_\varepsilon(p_n) \leq \frac{c_5 c_6}{8}$. So by a union bound, since $\omega_{p_n} \in \mathcal{E}_4$ implies that $\phi_z(o) \notin \text{giant}_\varepsilon(\omega_{p_n})$,

$$\mathbb{P}(\{\omega_{p_n} \in \mathcal{E}_4\} \cap \{\phi_z(o) \notin \text{giant}_\varepsilon(\omega_{p_m})\}) \geq \frac{c_5 c_6}{8}. \quad (6.7.2)$$

When $\omega_{p_n} \in \mathcal{E}_4$ and $\phi_z(o) \notin \text{giant}_\varepsilon(\omega_{p_m})$, then by monotonicity of the coupling, $\omega_{p_m} \in \text{meso}_{\phi_z(o)}^{\delta, \varepsilon}$. Therefore, using the fact that ϕ_z is an automorphism for the first equality,

$$\begin{aligned} \mathbb{P}_{p_m}(\text{meso}_o^{\delta, \varepsilon} \cap \text{act}_o^\varepsilon[xy]) &= \mathbb{P}_{p_m}(\text{meso}_{\phi_z(o)}^{\delta, \varepsilon} \cap \text{act}_{\phi_z(o)}^\varepsilon[\phi_z(xy)]) \\ &\geq \mathbb{P}(\{\omega_{p_n} \in \mathcal{E}_4\} \cap \{\phi_z(o) \notin \text{giant}_\varepsilon(\omega_{p_m})\} \cap \{\omega_{p_m} \in \text{act}_{\phi_z(o)}^\varepsilon[\phi_z(xy)]\}). \end{aligned}$$

In particular, if

$$\mathbb{P}(\{\omega_{p_n} \in \mathcal{E}_4\} \cap \{\phi_z(o) \notin \text{giant}_\varepsilon(\omega_{p_m})\} \cap \{\omega_{p_m} \in \text{act}_{\phi_z(o)}^\varepsilon[\phi_z(xy)]\}) \geq \frac{c_5 c_6}{16}, \quad (6.7.3)$$

then we are done because $A_m(1, \beta)$ holds, and hence $A_m(k-1, \beta)$ holds, with $\beta := \frac{c_5 c_6}{16}$. So we may assume that eq. (6.7.3) is false, and therefore by eq. (6.7.2) and a union bound, that the event

$$\mathcal{A}_1 := \{\phi_z(o) \notin \text{giant}_\varepsilon(\omega_{p_m})\} \cap \{\omega_{p_m} \notin \text{act}_{\phi_z(o)}^\varepsilon[\phi_z(xy)]\}$$

satisfies

$$\mathbb{P}(\{\omega_{p_n} \in \mathcal{E}_4\} \cap \mathcal{A}_1) \geq \frac{c_5 c_6}{16}.$$

On the event that $\omega_{p_n} \in \mathcal{E}_4$, let \mathcal{B}_* denote the class \mathcal{B}_i from the collection $(\mathcal{B}_i : i \in I)$ (defined with respect to ω_{p_n}) that contains z , and define the event

$$\mathcal{A}_2 := \{\omega_{p_m} \cap \{\phi_u(xy) : u \in \mathcal{B}_*\} \neq \emptyset\}.$$

Suppose we condition on ω_{p_n} and find that $\omega_{p_n} \in \mathcal{E}_4$, then condition on the restriction of ω_{p_m} to $E \setminus \{\phi_u(xy) : u \in \mathcal{B}_*\}$. Then the events \mathcal{A}_1 and \mathcal{A}_2 are both increasing events of the unrevealed edges, i.e. the restriction of ω_{p_m} to $\{\phi_u(xy) : u \in \mathcal{B}_*\}$. The conditional law of this restriction of ω_{p_m} is still a product measure, and in particular almost surely satisfies the FKG inequality. Since there are at least $\frac{|\mathcal{B}_*|}{2} \geq \frac{c_6 d}{2}$ unrevealed edges, the conditional probability that \mathcal{A}_2 will occur is almost surely at least

$$1 - \left(1 - \frac{p_m - p_n}{1 - p_n}\right)^{\frac{|\mathcal{B}_*|}{2}} \geq 1 - \left(1 - \frac{\varepsilon}{3^{n+3} d}\right)^{\frac{c_6 d}{2}} \geq 1 - e^{-\frac{c_6 \varepsilon}{2 \cdot 3^{n+3}}} =: \beta_1,$$

where the first inequality used the second condition enumerated in Lemma 6.7.2. Notice that if \mathcal{A}_1 and \mathcal{A}_2 both occur after revealing the remaining edges, then (using that $\omega_{p_n} \in \mathcal{E}_4$)

$$\omega_{p_m} \in \text{meso}_{\phi_z(o)}^{\delta, \varepsilon} \cap \text{act}_{\phi_z(o)}^{\varepsilon} [\phi_z(F \setminus \{xy\})].$$

Therefore, using that ϕ_z is an automorphism in the first line and applying the conditional FKG inequality to obtain the final inequality,

$$\begin{aligned} \mathbb{P}_{p_m} \left(\text{meso}_o^{\delta, \varepsilon} \cap \text{act}_o^{\varepsilon} [F \setminus \{xy\}] \right) &= \mathbb{P}_{p_m} \left(\text{meso}_{\phi_z(o)}^{\delta, \varepsilon} \cap \text{act}_{\phi_z(o)}^{\varepsilon} [\phi(F \setminus \{xy\})] \right) \\ &\geq \mathbb{P}(\{\omega_{p_n} \in \mathcal{E}_4\} \cap \mathcal{A}_1 \cap \mathcal{A}_2) \\ &\geq \beta_1 \cdot \mathbb{P}(\{\omega_{p_n} \in \mathcal{E}_4\} \cap \mathcal{A}_1) \geq \frac{\beta_1 c_5 c_6}{16} =: \beta, \end{aligned}$$

and hence $A_m(k-1, \beta)$ holds. \square

We are now ready to complete the proof of the lemma. Recall that our goal is to find $N(\varepsilon, k, \alpha) < \infty$ and $\beta(\varepsilon, k, \alpha, n) > 0$ such that if $|V| \geq Nd$ and $n \geq N$, then $A_m(k-1, \beta)$ must hold. We will split our argument into two cases according to whether or not $d \geq c_6^{-1}$. (The ultimate choices of N and β are obtained by taking the maximum/minimum of the constants in each case.)

First suppose that $d < c_6^{-1}$. By Lemma 6.2.2, there exists $N_1(\varepsilon, k, \alpha) < \infty$ such that if $|V| \geq N_1$ then $p_n \geq \frac{1}{2d} \geq \frac{c_6}{2}$. By Lemma 6.4.3, if $p_n \geq \frac{c_6}{2}$ then there exists an edge $uv \in F$ such that

$$\begin{aligned} \mathbb{P}_{p_n} \left(\text{meso}_u^{\delta, \varepsilon} \cap \text{act}_u^{\varepsilon}[uv] \right) &\geq \frac{p_n^{|F|}}{2|F|} \mathbb{P}_{p_n} \left(\text{meso}_o^{\delta, \varepsilon} \cap \text{act}_o^{\varepsilon}[F] \right) \\ &\geq \frac{\left(\frac{c_6}{2}\right)^k}{2k} \cdot \alpha =: c_8(\varepsilon, k, \alpha). \end{aligned}$$

As argued in the proof of Claim 6.7.8, by the third condition of Lemma 6.7.2, there exists $N_2(\varepsilon, k, \alpha) < \infty$ such that if $n \geq N_2$ then $\theta_\varepsilon(p_m) - \theta_\varepsilon(p_n) \leq \frac{c_8}{2}$. Set $N := N_1 \vee N_2$, and assume that $|V| \geq Nd$ (hence $|V| \geq N$) and $n \geq N$. Then by a union bound,

$$\mathbb{P}_{p_m} \left(\text{meso}_u^{\delta, \varepsilon} \cap \text{act}_u^{\varepsilon}[uv] \right) \geq \mathbb{P}_{p_n} \left(\text{meso}_u^{\delta, \varepsilon} \cap \text{act}_u^{\varepsilon}[uv] \right) - (\theta_\varepsilon(p_m) - \theta_\varepsilon(p_n)) \geq c_8 - \frac{c_8}{2} = \frac{c_8}{2},$$

and hence, by transitivity, $A_m(1, \beta)$ holds (and hence $A_m(k-1, \beta)$ holds) with $\beta := \frac{c_8}{2}$.

Now suppose that $d \geq c_6^{-1}$. Let $\beta(\varepsilon, k, \alpha, n) > 0$ be the constant from Claim 6.7.12, which we may assume satisfies $\beta \leq c_7$. Set $N := c_7^{-1}$, and assume that $|V| \geq Nd$ and $n \geq N$. By Claim 6.7.8, either $A_m(1, c_3)$ holds or $\mathbb{P}_{p_n}(\mathcal{E}_3) \geq c_3$. If $A_m(1, c_3)$ holds then $A_m(1, \beta)$ holds (and hence $A_m(k-1, \beta)$ holds), so we may instead assume that $\mathbb{P}_{p_n}(\mathcal{E}_3) \geq c_3$. So by Claim 6.7.10, $\mathbb{P}_{p_n}(\mathcal{F}_1) \geq c_5$. Therefore by Claim 6.7.12, $A_m(k-1, \beta)$ holds, as required.

6.8 The non-macromolecular case

The goal of this section is to prove the following proposition, which says that for supercritical percolation on a high-degree non-macromolecular finite transitive graph, o is unlikely to belong to a mesoscopic cluster.

Proposition 6.8.1. *Let $\mathcal{H} \subseteq \mathcal{F}$ be an infinite set of (isomorphism classes of) finite, connected, vertex transitive graphs with $\lim_{G \in \mathcal{H}} \deg G = \infty$ that contains no infinite macromolecular subsets. Then for every supercritical $p : \mathcal{H} \rightarrow [0, 1]$,*

$$\lim_{n \rightarrow \infty} \limsup_{G \in \mathcal{H}} \mathbb{P}_p (|K_o| \geq n, o \notin \text{giant}) = 0. \quad (6.8.1)$$

Equivalently, we must show that if the convergence claimed in eq. (6.8.1) fails, then we can construct suitable macromolecular decompositions for infinitely many of the graphs in \mathcal{H} . In Section 6.8, we show that we can construct these macromolecular decompositions if the giant cluster does not behave as it should in the neighbourhood of the origin. In Section 6.8, we will then show that if the giant cluster does behave as it should in the neighbourhood of the origin, then we

can deduce eq. (6.8.1). Together, these yield Proposition 6.8.1. In Section 6.11, we will show that Proposition 6.8.1 implies our main theorems in the case of high-degree non-macromolecular graphs.

Macromolecular graphs and the neighbourhood of the origin

The goal of this subsection is to prove the following proposition. Informally, this proposition states that if we consider supercritical percolation on a high-degree finite transitive graph that is *not* macromolecular, then the giant cluster includes a positive proportion of the neighbours of the origin with high probability. (This is not the case for macromolecular graphs, since the “local giants” we see near the origin may fail to be included in the global giant.)

Proposition 6.8.2. *Let $\mathcal{H} \subseteq \mathcal{F}$ be an infinite set of (isomorphism classes of) finite, connected, vertex transitive graphs with $\lim_{G \in \mathcal{H}} \deg G = \infty$ that contains no infinite macromolecular subsets. Then*

$$\lim_{\alpha \downarrow 0} \liminf_{\mathcal{H}} \mathbb{P}_p \left(\|\text{giant}\|_{\text{neigh}(o)} \geq \alpha \right) = 1$$

for every supercritical $p : \mathcal{H} \rightarrow [0, 1]$.

When the graphs in \mathcal{H} are dense, then the result follows easily from supercritical uniqueness Theorem 6.1.1. So for most of the proof of this proposition, we may assume that the graphs are sparse. The proof of this proposition will rely on the following definition. Let $G = (V, E) \in \mathcal{F}$, let x be a vertex of G , and let $\alpha, \beta \in (0, 1)$. Given a configuration $\omega \in \{0, 1\}^E$, we say that x is **(α, β) -happy** if there is an ω -connected subset $X \subseteq \text{neigh}(x)$ with $\|X\|_{\text{neigh}(x)} \geq \alpha$ such that $X \subseteq \text{giant}_\beta$, and say that x is **(α, β) -sad** if there is an ω -connected subset $X \subseteq \text{neigh}(x)$ with $\|X\|_{\text{neigh}(x)} \geq \alpha$ such that $X \not\subseteq \text{giant}_\beta$. (A vertex can be both happy and sad, or neither happy nor sad.) Given $x, y \in V$ and parameters $\alpha, \beta, \gamma, p \in (0, 1)$, we say that x and y are **$(\alpha, \beta, \gamma, p)$ -friends** if

$$\mathbb{P}_p \left([\{x \text{ is } (\alpha, \beta)\text{-happy}\} \cap \{y \text{ is } (\alpha, \beta)\text{-sad}\}] \cup [\{y \text{ is } (\alpha, \beta)\text{-happy}\} \cap \{x \text{ is } (\alpha, \beta)\text{-sad}\}] \right) \leq \gamma.$$

That is, x and y are friends if it is unlikely that one is happy while the other is sad. We will eventually see that the subgraph spanned by edges between friends can be identified, in an appropriate sense, with the macromolecular structure of G . Our first lemma gives conditions under which a vertex has at most $O(1)$ neighbours that it is not friends with.

Lemma 6.8.3. *For all $\alpha, \beta, \gamma, \delta \in (0, 1)$, there exists $\varepsilon > 0$ and $N < \infty$ such that the following holds. Let G be a finite connected transitive graph with $\deg \geq \varepsilon^{-1}$ that is not ε -molecular. Suppose*

that $p \in (0, 1)$ satisfies $\theta_\beta((1 - \delta)p) \geq \delta$ and $p\theta'_\beta(p) \leq \frac{1}{\delta}$. Then

$$|\{x \in \text{neigh}(o) : o \text{ and } x \text{ are not } (\alpha, \beta, \gamma, p)\text{-friends}\}| \leq N.$$

Proof. Let $G \in \mathcal{F}$. Fix $\alpha, \beta, \gamma, \delta, p \in (0, 1)$ such that $\theta_\beta((1 - \delta)p) \geq \delta$ and $p\theta'_\beta(p) \leq 1/\delta$, and suppose $\varepsilon > 0$ is such that $\deg \geq \varepsilon^{-1}$ and that G is not ε -molecular. By Lemma 6.2.2, there exists $\varepsilon_0 = \varepsilon_0(\delta, \beta) > 0$ such that if $\varepsilon \leq \varepsilon_0$ then $p \geq \frac{1}{2\deg}$. We will write that two vertices are friends to mean that they are $(\alpha, \beta, \gamma, p)$ -friends.

Fix an edge $xy \in E$ and assume that x and y are not friends. By a union bound, we may assume without loss of generality (swapping x and y if necessary) that the event \mathcal{E} that x is (α, β) -happy but y is (α, β) -sad satisfies $\mathbb{P}_p(\mathcal{E}) \geq \gamma/2$. On this event, let X be an ω -connected subset of $\text{neigh}(x)$ with $\|X\|_{\text{neigh}(x)} \geq \alpha$ that is contained in giant_β , and let Y be an ω -connected subset of $\text{neigh}(y)$ with $\|Y\|_{\text{neigh}(y)} \geq \alpha$ that is not contained in giant_β . By Theorem 6.1.1, there exists $\varepsilon_1 = \varepsilon_1(\beta, \gamma, \delta) \in (0, \varepsilon_0)$ such that if $\varepsilon \leq \varepsilon_1$ then $\mathbb{P}_p(\|K_2\| \geq \beta/2) \leq \gamma/4$, so that if $\varepsilon \leq \varepsilon_1$ then by a union bound

$$\mathbb{P}_p(\mathcal{E} \cap \{\|K_2\| < \beta/2\}) \geq \gamma/4.$$

If \mathcal{E} occurs and $y \in \text{giant}_\beta$, then opening any of the edges between y and Y causes giant_β to increase in size by at least $\alpha \deg$. Thus, using transitivity, Russo's formula implies that

$$\theta'_\beta(p) \geq \frac{1}{2} \mathbb{P}_p(\mathcal{E} \cap \{y \in \text{giant}_\beta\}) (\alpha \deg)^2,$$

and since $p \geq \frac{1}{2\deg}$ and $p\theta'_\beta(p) \leq \delta^{-1}$ it follows that

$$\mathbb{P}_p(\mathcal{E} \cap \{y \in \text{giant}_\beta\}) \leq \frac{4}{\alpha^2 \delta \deg} \leq \frac{4\varepsilon}{\alpha^2 \delta}.$$

Thus, if $\varepsilon \leq \varepsilon_2 = \min\{\varepsilon_1, \alpha^2 \gamma \delta / 32\}$ then $\mathbb{P}_p(\mathcal{E} \cap \{y \in \text{giant}_\beta\}) \leq \gamma/8$ and hence by a union bound,

$$\mathbb{P}_p(\mathcal{E} \cap \{y \notin \text{giant}_\beta\} \cap \{\|K_2\| < \beta/2\}) \geq \gamma/8.$$

By applying quantitative insertion-tolerance (Proposition 6.6.4) to open an edge in yY , we deduce that there is a constant $c_1 = c_1(\alpha, \gamma) > 0$ such that if $\varepsilon \leq \varepsilon_2$ then the event

$$\mathcal{E}_1 := \{x \text{ is } (\alpha, \beta)\text{-happy}\} \cap \{y \notin \text{giant}_\beta\} \cap \{\|K_y\|_{\text{neigh}(y)} \geq \alpha\}$$

satisfies $\mathbb{P}_p(\mathcal{E}_1) \geq c_1$. If the event $\mathcal{E}_1 \cap \{x \leftrightarrow y\}$ holds, then x is not connected to X , and opening any edge between x and X causes the size of giant_β to increase by at least $\alpha \deg$. Thus, as above, Russo's formula together with transitivity imply that

$$p\theta'_\beta(p) \geq \frac{\alpha^2 \deg}{4} \mathbb{P}_p(\mathcal{E}_1 \cap \{x \leftrightarrow y\}), \quad \text{so that} \quad \mathbb{P}_p(\mathcal{E}_1 \cap \{x \leftrightarrow y\}) \leq \frac{4\varepsilon}{\alpha^2 \delta}.$$

Thus, if we define $\varepsilon_3 = \min\{\varepsilon_2, c_1\alpha^2\delta/8\}$ then by a union bound

$$\mathbb{P}_p(\mathcal{E}_1 \cap \{x \leftrightarrow y\}) \geq \frac{c_1}{2}$$

whenever $\varepsilon \leq \varepsilon_3$. Applying quantitative insertion-tolerance to open an edge between x and X , we deduce that there is a constant $c_2 = c_2(\alpha, \gamma) > 0$ such that the event

$$\mathcal{E}_2 := \{\|K_x\|_{\text{neigh}(x)} \geq \alpha\} \cap \{x \in \text{giant}_\beta\} \cap \{\|K_y\|_{\text{neigh}(y)} \geq \alpha\} \cap \{y \notin \text{giant}_\beta\}$$

satisfies $\mathbb{P}_p(\mathcal{E}_2) \geq c_2$ whenever $\varepsilon \leq \varepsilon_3$. In particular, the unordered pair $xy = \{x, y\}$ satisfies

$$M(xy) := \mathbb{E}_p \left[|K_y| \mathbb{1}_{\text{act}_{\|K_y\| \geq \beta}[xy]} \right] + \mathbb{E}_p \left[|K_x| \mathbb{1}_{\text{act}_{\|K_x\| \geq \beta}[xy]} \right] \geq c_2 \alpha \deg$$

whenever $\varepsilon \leq \varepsilon_3$. If we define $\mathcal{E}_x = \{y \in \text{neigh}(x) : x \text{ and } y \text{ are not friends}\}$ then this estimate holds for every $x \in V$ and $y \in \mathcal{E}_x$ provided that $\varepsilon \leq \varepsilon_3$. Using transitivity and summing over this estimate, we obtain that if $\varepsilon \leq \varepsilon_3$ then

$$\theta'_\beta(p) \geq \frac{1}{|V|} \sum_{x \in V} \sum_{y \in \mathcal{E}_x} \mathbb{E}_p \left[|K_x| \mathbb{1}_{\text{act}_{\|K_x\| \geq \beta}[xy]} \right] = \frac{1}{2|V|} \sum_{x \in V} \sum_{y \in \mathcal{E}_x} M(xy) \geq \frac{c_2 \alpha}{2} \deg |\mathcal{E}_o|.$$

Since $p \geq 1/(2 \deg)$ and $p\theta'_\beta(p) \leq \delta^{-1}$, it follows that if $\varepsilon \leq \varepsilon_3$ then $|\mathcal{E}_o| \leq 4(c_2 \alpha \delta)^{-1}$. This completes the proof. \square

We now investigate the subgraph spanned by edges between friends. We say that x and y are $(\alpha, \beta, \gamma, p)$ -**linked** if they are connected in the subgraph of G spanned by those edges of G whose endpoints are $(\alpha, \beta, \gamma, p)$ -friends, and define $\text{pop}(\alpha, \beta, \gamma, p)$ by

$$\text{pop}(\alpha, \beta, \gamma, p) := \frac{1}{\deg} |\{x \in V : o \text{ and } x \text{ are } (\alpha, \beta, \gamma, p)\text{-linked}\}|.$$

Note that this quantity can be larger than 1.

Lemma 6.8.4. *For all $\alpha, \beta, \gamma, \delta \in (0, 1)$, there exists $\varepsilon > 0$ such that the following holds. Let G be a finite connected transitive graph with $\deg G \geq \varepsilon^{-1}$ that is not ε -dense. Suppose that $p \in (0, 1)$ satisfies $\theta_\beta((1-\delta)p) \geq \delta$, $p\theta'_\beta(p) \leq \delta^{-1}$, and $\text{pop}(\alpha, \beta, \gamma, p) \leq \delta^{-1}$. Then G is ε -macromolecular.*

Proof. Let G and p satisfy the given hypotheses, for some constant $\varepsilon > 0$ to be determined. Write pop for $\text{pop}(\alpha, \beta, \gamma, p)$. Say that vertices x and y are linked to mean that they are $(\alpha, \beta, \gamma, \delta)$ -linked, and note that this induces an $\text{Aut } G$ -invariant equivalence relation on $V(G)$. First suppose that this

equivalence relation is trivial, i.e. every vertex is linked to every other vertex. By definition of pop , every equivalence class has size $\text{pop} \cdot \deg G$. Therefore,

$$\frac{|E(G)|}{|G|^2} = \frac{\deg G}{2|G|} = \frac{\deg G}{2\text{pop} \cdot \deg G} \geq \frac{\delta}{2},$$

i.e. G is $\delta/2$ -dense. If $\varepsilon < \varepsilon_0 := \delta/2$, then this is impossible (since G is not ε -dense), and hence the equivalence relation of being linked must be non-trivial. Let (A, B) be the corresponding macromolecular decomposition.

By Lemma 6.8.3, there exists $\varepsilon_1 = \varepsilon_1(\alpha, \beta, \gamma, \delta) \in (0, \varepsilon_0)$ and $N(\alpha, \beta, \gamma, \delta) < \infty$ such that if $\varepsilon < \varepsilon_1$ then (since G being not ε -dense implies that G is not ε -molecular)

$$|\{x \in \text{neigh}(o) : o \text{ and } x \text{ are not } (\alpha, \beta, \gamma, p)\text{-friends}\}| \leq N.$$

In particular, by transitivity, since every pair of friends is trivially linked,

$$\frac{|E(B)|}{|G|} \leq \frac{N}{2}.$$

Similarly, $\deg A \geq \deg G - N$. Hence if $\varepsilon \leq \frac{1}{2N}$, which implies that $\deg G \geq 2N$, then

$$\frac{|E(A)|}{|A|^2} = \frac{\deg A}{2\text{pop} \cdot \deg G} \geq \frac{\deg G - N}{(2/\delta) \deg G} = \frac{\delta}{2} \left(1 - \frac{N}{\deg G}\right) \geq \frac{\delta}{4}.$$

Therefore, if $\varepsilon < \varepsilon_2 := \min\left(\varepsilon_1, \frac{2}{N}, \frac{1}{2N}, \frac{\delta}{4}\right)$, then (A, B) is an ε -macromolecular decomposition. \square

Lemma 6.8.5. *Let $\alpha, \beta, \gamma, p \in (0, 1)$ and let $m \in \mathbb{N}$. Let G be a finite connected graph and suppose that v_1, \dots, v_m is a sequence of vertices such that v_i and v_{i+1} are $(\alpha, \beta, \gamma, p)$ -friends for all $i \in \{1, \dots, m-1\}$. Then*

$$\mathbb{P}_p \left(\{v_1 \text{ is } (\alpha, \beta)\text{-sad}\} \cap \bigcup_{i=2}^m \{v_i \text{ is } (\alpha, \beta)\text{-happy}\} \right) \leq m\gamma.$$

Proof. We have by a union bound that

$$\mathbb{P}_p \left(\{v_1 \text{ is } (\alpha, \beta)\text{-sad}\} \cap \bigcup_{i=2}^m \{v_i \text{ is } (\alpha, \beta)\text{-happy}\} \right) \leq \sum_{i=1}^{m-1} \mathbb{P}_p(v_i \text{ is } (\alpha, \beta)\text{-sad}, v_{i+1} \text{ is } (\alpha, \beta)\text{-happy}) \leq m\gamma$$

as claimed, where the second inequality follows from the definition of $(\alpha, \beta, \gamma, p)$ -friends. \square

Lemma 6.8.6. *For all $\beta, \delta, \eta \in (0, 1)$, there exists $\alpha > 0$ such that for every finite connected transitive graph G , for every parameter $p \in (0, 1)$, if $\theta_\beta(p) \geq \delta$ then*

$$\mathbb{P}_p(o \text{ is neither } (\alpha, \beta)\text{-happy nor } (\alpha, \beta)\text{-sad}) \leq \eta.$$

Proof. Suppose that $\theta_\beta(p) \geq \delta$. Then by transitivity, $\mathbb{E}_p \|\text{giant}_\beta\|_{\text{neigh}(o)} \geq \delta$, and hence by Markov's inequality, $\mathbb{P}_p \left(\|\text{giant}_\beta\|_{\text{neigh}(o)} \geq \frac{\delta}{2} \right) \geq \frac{\delta}{2}$. Now there can never be more than $(1/\beta)$ -many clusters each having density at least β . Thus,

$$\lambda := \max_{u \in V(G)} \|K_u\|_{\text{neigh}(o)}$$

satisfies $\mathbb{P}_p \left(\lambda \geq \frac{\beta\delta}{2} \right) \geq \frac{\delta}{2}$, and in particular, $\mathbb{E}_p [\lambda] \geq \frac{\beta\delta^2}{4}$. So by the universal tightness theorem (Theorem 6.5.4), there is a universal constant $C < \infty$ such that for all $\alpha > 0$,

$$\mathbb{P}_p (\lambda \leq \alpha) \leq C \cdot \frac{4\alpha}{\beta\delta^2}.$$

Now the conclusion follows by choosing $\alpha := \frac{\eta\beta\delta^2}{4C}$, noting that o is neither (α, β) -happy nor (α, β) -sad if and only if $\lambda < \alpha$. \square

We next prove the following deterministic graph theory lemma, via a probabilistic argument.

Lemma 6.8.7. *For all $k \in \mathbb{N}$ there exist $\varepsilon > 0$ and $m \in \mathbb{N}$ such that if G is a finite, connected, regular graph that is not ε -dense and has at least one edge, then G contains a (possibly self-intersecting) path of vertices v_1, \dots, v_m satisfying*

$$\left| \bigcup_{i=1}^m \text{neigh}(v_i) \right| \geq k \deg.$$

Proof. We claim that the result holds with $m := 6k^2$ and $\varepsilon := (2m)^{-1}$. Let $G = (V, E)$ be a finite connected regular graph that is not ε -dense and has at least one edge. We will prove the claim by case analysis according to whether $\text{diam } G \geq 3k$.

First suppose that $\text{diam } G \geq 3k$. Let v_1, \dots, v_{3k} be the first $3k$ vertices in a geodesic path between two points of maximal distance in G . Extend the path v_1, \dots, v_{3k} in an arbitrary way to a path v_1, \dots, v_m , which has the correct length. (For example, repeatedly cross the edge $v_{3k-1}v_{3k}$.) Since v_1, \dots, v_{3k} is a geodesic, we must have that $\text{neigh}(v_i) \cap \text{neigh}(v_j) = \emptyset$ for all $1 \leq i, j \leq 3k$ with $|i - j| \geq 3$, so

$$\left| \bigcup_{i=1}^m \text{neigh}(v_i) \right| \geq \sum_{j=1}^k |\text{neigh}(v_{3j})| = k \deg$$

as required.

Now suppose that $\text{diam } G \leq 3k$. Let \mathbb{P} be the law of an independent sequence u_1, \dots, u_{2k} of vertices in V each chosen uniformly at random, so that

$$\mathbb{P}(v \in \text{neigh}(u_i)) = \frac{1}{|V|} \sum_{u \in V} \mathbb{1}_{v \in \text{neigh}(u)} = \frac{\deg}{|V|} = \frac{2|E|}{|V|^2} \leq 2\epsilon = \frac{1}{m}$$

for each $v \in V$ and $1 \leq i \leq 2k$. It follows by independence and linearity of expectation that

$$\mathbb{E}|\text{neigh}(u_i) \cap \text{neigh}(u_j)| = \sum_{v \in V} \mathbb{P}(v \in \text{neigh}(u_i)) \mathbb{P}(v \in \text{neigh}(u_j)) = |V| \left(\frac{\deg}{|V|} \right)^2 \leq \frac{\deg}{m}$$

for every $i \neq j$ and hence by inclusion-exclusion that

$$\mathbb{E} \left| \bigcup_{i=1}^{2k} \text{neigh}(u_i) \right| \geq 2k \deg - \sum_{i \neq j} \mathbb{E}|\text{neigh}(u_i) \cap \text{neigh}(u_j)| \geq 2k \deg - (2k)^2 \frac{\deg}{m} \geq k \deg.$$

In particular, there is a *deterministic* sequence of vertices $\tilde{u}_1, \dots, \tilde{u}_{2k}$ such that $|\bigcup_{i=1}^{2k} \text{neigh}(\tilde{u}_i)| \geq k \deg$. By picking a geodesic (which necessarily has length $\leq 3k$) from \tilde{u}_i to \tilde{u}_{i+1} for each i , then extending arbitrarily to obtain the correct length, we can find a path v_1, \dots, v_m that contains $\tilde{u}_1, \dots, \tilde{u}_{2k}$ as a subsequence and hence satisfies

$$\left| \bigcup_{i=1}^m \text{neigh}(v_i) \right| \geq \left| \bigcup_{i=1}^{2k} \text{neigh}(\tilde{u}_i) \right| \geq k \deg$$

as desired. \square

We are now ready to complete the proof of Proposition 6.8.2. In addition to the lemmas from the present subsection, our proof will also apply Proposition 6.7.1 from the previous section, stating that sets of larger than degree order have large intersection with the giant with high probability.

Proof of Proposition 6.8.2. Let $q : \mathcal{H} \rightarrow [0, 1]$ be a supercritical assignment of parameters. Our goal is to prove that

$$\lim_{\alpha \downarrow 0} \liminf_{\mathcal{H}} \mathbb{P}_q \left(\|\text{giant}\|_{\text{neigh}(o)} \geq \alpha \right) = 1. \quad (6.8.2)$$

We may assume without loss of generality that $q(G) < 1$ for all $G \in \mathcal{H}$. In particular, we may assume that there exists a constant $\delta_0 > 0$ such that for every $G \in \mathcal{H}$, $q(G)$ is δ_0 -supercritical for G . By Lemma 6.3.4 applied to the interval $I = I(G) := [(1 - \delta_0/2)q, (1 + \delta_0/2)(1 - \delta_0/2)q]$,

$$\mathcal{L} \left(\left\{ p \in I : p \theta'_{\delta_0/8}(p) \leq 1/\epsilon \right\} \right) \geq \left(1 - \frac{2\epsilon}{\delta_0/2} \right) \mathcal{L}(I) \quad \text{for all } G \in \mathcal{H} \text{ and all } \epsilon > 0,$$

where \mathcal{L} denotes the Lebesgue measure. In particular, for all $G \in \mathcal{H}$ there exists $p = p(G) \in I$ such that $p\theta'_{\delta_0/8}(p) \leq 1/\delta$ where $\delta := \delta_0/8$. Note that p is always δ -supercritical because q is δ_0 -supercritical and $p \geq (1 - \delta_0/2)q$. By Theorem 6.1.1 and Theorem 6.2.1,

$$\lim_{\mathcal{H}} \mathbb{P}_q (\text{giant} = \text{giant}_\delta) = 1.$$

So by monotonicity, since $p(G) \leq q(G)$ for all $G \in \mathcal{H}$, it suffices to establish that

$$\lim_{\alpha \downarrow 0} \liminf_{\mathcal{H}} \mathbb{P}_p \left(\|\text{giant}_\delta\|_{\text{neigh}(o)} \geq \alpha \right) = 1.$$

To this end, fix an arbitrary constant $\eta > 0$, and note that the event that $\|\text{giant}_\delta\|_{\text{neigh}(o)} \geq \alpha$ is the same as the event that o is (α, δ) -happy, for all $\alpha > 0$. We will find a constant $\alpha > 0$ such that for all but finitely many $G \in \mathcal{H}$,

$$\mathbb{P}_p (o \text{ is } (\alpha, \delta)\text{-happy}) \geq 1 - \eta. \quad (6.8.3)$$

By Proposition 6.7.1, there exists $\varepsilon_0 > 0$ such that for all $G \in \mathcal{H}$, for every set of vertices $A \subseteq V(G)$ with $|A| \geq \varepsilon_0^{-1} \deg G$,

$$\mathbb{P}_p \left(\|\text{giant}_\delta\|_A \geq \varepsilon_0 \right) \geq 1 - \frac{\eta}{3}.$$

By Lemma 6.8.7, there exists $\varepsilon_1 > 0$ and $m \in \mathbb{N}$ such that every finite, connected, regular graph G that is not ε_1 -dense and has at least one edge must contain a path of vertices v_1, \dots, v_m satisfying

$$\left| \bigcup_{i=1}^m \text{neigh}(v_i) \right| \geq \frac{2}{\varepsilon_0} \deg G.$$

By Lemma 6.8.6, there exists $\alpha_0 > 0$ such that for all $G \in \mathcal{H}$,

$$\mathbb{P}_p (o \text{ is } (\alpha_0, \delta)\text{-happy or } (\alpha_0, \delta)\text{-sad}) \geq 1 - \frac{\eta}{3}, \quad (6.8.4)$$

and by monotonicity, this inequality also holds with any $\alpha \in (0, \alpha_0)$ in place of α_0 . Define $\alpha := \min(\alpha_0, 1/m)$ and $\gamma := \frac{\eta}{3m}$. Now by Lemma 6.8.4, there exists $\varepsilon_2 > 0$ such that for all but finitely many $G \in \mathcal{H}$ (namely, every $G \in \mathcal{H}$ with $\deg G \geq 1/\varepsilon_2$ that is not ε_2 -macromolecular), **(Case A)** G is ε_2 -dense or **(Case B)** G satisfies $\text{pop}(\alpha, \delta, \gamma, p) \geq 1/\varepsilon_1$. We will establish eq. (6.8.3) in each of these two cases in turn.

Case A Note that when o is (α, δ) -sad, there exists a cluster $K \not\subseteq \text{giant}_\delta$ satisfying

$$\|K\| \geq \frac{\alpha \deg G}{|G|} = \frac{2\alpha |E(G)|}{|G|^2}.$$

In particular, if G is ε_2 -dense and o is (α, δ) -sad then $\text{giant}_\delta \neq \text{giant}_{2\alpha\varepsilon_2}$. By Theorem 6.1.1 and Theorem 6.2.1, for all but finitely many $G \in \mathcal{H}$,

$$\mathbb{P}_p \left(\text{giant}_{2\alpha\varepsilon_2} \neq \text{giant}_\delta \right) \leq \frac{2\eta}{3}.$$

So for all but finitely many $G \in \mathcal{H}$, if G is ε_2 -dense then

$$\begin{aligned} \mathbb{P}_p (o \text{ is } (\alpha, \delta)\text{-happy}) &\geq \mathbb{P}_p (o \text{ is } (\alpha, \delta)\text{-happy or } (\alpha, \delta)\text{-sad}) - \mathbb{P}_p (o \text{ is } (\alpha, \delta)\text{-sad}) \\ &\geq 1 - \frac{\eta}{3} - \frac{2\eta}{3} = 1 - \eta. \end{aligned}$$

Case B By Lemma 6.8.3, there exists $N < \infty$ such that for all but finitely many $G \in \mathcal{H}$ (namely, for a suitable constant $\varepsilon_3 > 0$, every $G \in \mathcal{H}$ with $\deg G \geq 1/\varepsilon_3$ that is not ε_3 -macromolecular),

$$|\{x \in \text{neigh}(o) : o \text{ and } x \text{ are not } (\alpha, \delta, \gamma, p)\text{-friends}\}| \leq N.$$

In particular, for all but finitely many $G \in \mathcal{H}$ (those graphs that additionally satisfy $\deg G \geq 2N$),

$$\frac{1}{\deg G} |\{x \in \text{neigh}(o) : o \text{ and } x \text{ are } (\alpha, \delta, \gamma, p)\text{-friends}\}| \geq \frac{1}{2}. \quad (6.8.5)$$

Now suppose that some given $G \in \mathcal{H}$ satisfies $\text{pop}(\alpha, \delta, \gamma, p) \geq 1/\varepsilon_1$, eq. (6.8.5), and (a trivially harmless hypothesis) $\deg G > 0$. Consider the spanning subgraph of G containing only the edges $uv \in E(G)$ such that u and v are $(\alpha, \delta, \gamma, p)$ -friends, and let A denote the connected component of this graph containing o . Note that

$$\frac{|E(A)|}{|A|^2} = \frac{\deg A}{2|A|} \leq \frac{\deg G}{2|A|} = \frac{1}{2\text{pop}(\alpha, \delta, \gamma, p)} \leq \frac{\varepsilon_1}{2} < \varepsilon_1,$$

and the graph A contains at least one edge because $\deg A \geq \frac{1}{2} \deg G > 0$. So by definition of ε_1 , the graph A must contain a path of vertices v_1, \dots, v_m satisfying

$$|T| \geq \frac{2}{\varepsilon_0} \deg A \geq \frac{1}{\varepsilon_0} \deg G \quad \text{where } T := \bigcup_{i=1}^m \text{neigh}(v_i),$$

and without loss of generality, since A is transitive, we may choose this path such that $v_1 = o$. Therefore, by definition of ε_0 ,

$$\mathbb{P}_p \left(\|\text{giant}_\delta\|_T \geq \varepsilon_0 \right) \geq 1 - \eta/3.$$

We always have

$$\begin{aligned}
\sum_{i=1}^m \|\text{giant}_\delta\|_{\text{neigh}(v_i)} &= \frac{1}{\deg G} \sum_{i=1}^m |\text{giant}_\delta \cap \text{neigh}(v_i)| \\
&\geq \frac{1}{\deg G} |\text{giant}_\delta \cap T| \\
&= \frac{|T|}{\deg G} \|\text{giant}_\delta\|_T \geq \frac{1}{\varepsilon_0} \|\text{giant}_\delta\|_T.
\end{aligned}$$

So if $\|\text{giant}_\delta\|_T \geq \varepsilon_0$, then some v_i with $i \in \{1, \dots, m\}$ must be $(1/m, \delta)$ -happy and hence (α, δ) -happy. In particular, we deduce that

$$\mathbb{P}_p \left(\bigcup_{i=1}^m \{v_i \text{ is } (\alpha, \delta)\text{-happy}\} \right) \geq 1 - \frac{\eta}{3}.$$

Now by combining this estimate with eq. (6.8.4) and the estimate appearing in Lemma 6.8.5 via a union bound,

$$\begin{aligned}
\mathbb{P}_p(o \text{ is } (\alpha, \beta)\text{-happy}) &\geq 1 - \mathbb{P}_p(o \text{ is neither } (\alpha, \delta)\text{-happy nor } (\alpha, \delta)\text{-sad}) \\
&\quad - \mathbb{P}_p \left(\bigcap_{i=1}^m \{v_i \text{ is not } (\alpha, \delta)\text{-happy}\} \right) \\
&\quad - \mathbb{P}_p \left(\{v_1 \text{ is } (\alpha, \delta)\text{-sad}\} \cap \bigcup_{i=2}^m \{v_i \text{ is } (\alpha, \delta)\text{-happy}\} \right) \\
&\geq 1 - \eta/3 - \eta/3 - m\gamma = 1 - \eta. \quad \square
\end{aligned}$$

Negligibility of mesoscopic clusters

In this subsection, we complete the proof of Proposition 6.8.1. Given our work in the previous subsection, it suffices to show that for supercritical percolation on a high-degree finite transitive graph, if the giant cluster includes a positive proportion of the neighbourhood of the origin with high probability, then eq. (6.8.1) holds.

Proof of Proposition 6.8.1. Let $\mathcal{H} \subseteq \mathcal{F}$ be an infinite set with $\lim_{G \in \mathcal{H}} \deg G = \infty$ that does not contain any infinite macromolecular subsets. Let $p : \mathcal{H} \rightarrow [0, 1]$ be supercritical. We may assume without loss of generality that $p(G) < 1$ for all $G \in \mathcal{H}$. In particular, we may assume that there exists a constant $\delta > 0$ such that for every $G \in \mathcal{H}$, the parameter $p(G)$ is δ -supercritical for G . Fix an arbitrary constant $\eta > 0$. Our goal is to find a constant $k < \infty$ such that for all but finitely many graphs $G \in \mathcal{H}$,

$$\mathbb{P}_p(|K_o| \geq k, K_o \neq \text{giant}) \leq \eta. \quad (6.8.6)$$

By Theorems 6.1.1 and 6.2.1, for all but finitely many $G \in \mathcal{H}$,

$$\mathbb{P}_p(\text{giant}_\delta \neq \text{giant}) \leq \eta/2.$$

So it suffices to find $k < \infty$ such that for all but finitely many $G \in \mathcal{H}$,

$$\mathbb{P}_p(o \in \mathcal{M}) \leq \eta/2 \quad \text{where } \mathcal{M} := \{u \in V(G) : k \leq |K_u| < \delta |G|\}. \quad (6.8.7)$$

Define $q(G) := (1 - \delta/2)p(G)$ for all $G \in \mathcal{H}$. By Proposition 6.8.2, there exists a constant $\alpha > 0$ such that for all but finitely many $G \in \mathcal{H}$,

$$\mathbb{P}_q\left(\|\text{giant}\|_{\text{neigh}(o)} \geq \alpha\right) \geq 1 - \frac{\eta^2}{200}. \quad (6.8.8)$$

By Theorems 6.1.1 and 6.2.1 again, for all but finitely many $G \in \mathcal{H}$,

$$\mathbb{P}_q(\text{giant}_\delta \neq \text{giant}) \leq \frac{\eta^2}{300}. \quad (6.8.9)$$

By Lemma 6.2.2, for all but finitely many $G \in \mathcal{H}$,

$$p(G) \geq \frac{1}{2 \deg G}. \quad (6.8.10)$$

We will show that if a graph in \mathcal{H} satisfies eqs. (6.8.8) to (6.8.10) then it must also satisfy eq. (6.8.7) with $k := 1000\alpha^{-1}\eta^{-2}\delta^{-1} \log(300\eta^{-2}\delta^{-1})$. Indeed, suppose for contradiction that some particular graph $G = (V, E) \in \mathcal{F}$ satisfies eqs. (6.8.8) to (6.8.10) but not eq. (6.8.7). By the negation of eq. (6.8.7) and Markov's inequality, (and by transitivity and linearity of expectation)

$$\mathbb{P}_p\left(\|\mathcal{M}\| \geq \frac{\eta}{4}\right) \geq \frac{\eta}{4}.$$

By eqs. (6.8.8) and (6.8.9), the set $\mathcal{S} := \{u \in V : \|\text{giant}_\delta\|_{\text{neigh}(u)} \geq \alpha\}$ satisfies

$$\mathbb{P}_q(o \in \mathcal{S}) \geq 1 - \frac{\eta^2}{200} - \frac{\eta^2}{300} \geq 1 - \frac{\eta^2}{64}.$$

So by Markov's inequality,

$$\mathbb{P}_q\left(\|\mathcal{S}\| \geq 1 - \frac{\eta}{8}\right) \geq 1 - \frac{\eta}{8}.$$

Recall that $(\omega_t : t \in [0, 1]) \sim \mathbb{P}$ denotes the standard monotone coupling of percolation measures on G . By a union bound,

$$\mathbb{P}\left(\|\mathcal{M}(\omega_p)\| \geq \frac{\eta}{4} \text{ and } \|\mathcal{S}(\omega_q)\| \geq 1 - \frac{\eta}{8}\right) \geq \frac{\eta}{4} - \frac{\eta}{8} = \frac{\eta}{8}. \quad (6.8.11)$$

Note that on the event being estimated in eq. (6.8.11),

$$\|\mathcal{M}(\omega_p) \cap \mathcal{S}(\omega_q)\| \geq \|\mathcal{S}(\omega_q)\| + \|\mathcal{M}(\omega_p)\| - 1 \geq \frac{\eta}{8}.$$

Moreover, we can rewrite this intersection density as

$$\begin{aligned} \|\mathcal{M}(\omega_p) \cap \mathcal{S}(\omega_q)\| &= \frac{1}{|V|} \sum_{u \in \mathcal{M}(\omega_p)} \mathbf{1}_{\mathcal{S}(\omega_q)}(u) \\ &= \frac{1}{|V|} \sum_{u \in \mathcal{M}(\omega_p)} \mathbf{1}_{\mathcal{S}(\omega_q)}(u) \sum_{v \in K_u(\omega_p)} \frac{1}{|K_u(\omega_p)|} \\ &= \frac{1}{|V|} \sum_{v \in \mathcal{M}(\omega_p)} \frac{1}{|K_v(\omega_p)|} \sum_{u \in K_v(\omega_p)} \mathbf{1}_{\mathcal{S}(\omega_q)}(u) \\ &= \frac{1}{|V|} \sum_{v \in \mathcal{M}(\omega_p)} \|\mathcal{S}(\omega_q)\|_{K_v(\omega_p)}. \end{aligned}$$

Therefore, by transitivity and linearity of expectation,

$$\mathbb{E} \left[\|\mathcal{S}(\omega_q)\|_{K_o(\omega_p)} \mathbf{1}_{\mathcal{M}(\omega_p)}(o) \right] = \mathbb{E} \left[\frac{1}{|V|} \sum_{v \in \mathcal{M}(\omega_p)} \|\mathcal{S}(\omega_q)\|_{K_v(\omega_p)} \right] \geq \frac{\eta}{8} \cdot \frac{\eta}{8} = \frac{\eta^2}{64},$$

and thus by Markov's inequality,

$$\mathbb{P} \left(\{o \in \mathcal{M}(\omega_p)\} \cap \left\{ \|\mathcal{S}(\omega_q)\|_{K_o(\omega_p)} \geq \frac{\eta^2}{128} \right\} \right) \geq \frac{\eta^2}{128}.$$

So by eq. (6.8.9) and a union bound,

$$\mathbb{P} \left(\{o \in \mathcal{M}(\omega_p)\} \cap \left\{ \|\mathcal{S}(\omega_q)\|_{K_o(\omega_p)} \geq \frac{\eta^2}{128} \right\} \cap \{\text{giant}(\omega_q) = \text{giant}_\delta(\omega_q)\} \right) \geq \frac{\eta^2}{128} - \frac{\eta^2}{300} \geq \frac{\eta^2}{300}. \quad (6.8.12)$$

On the event being estimated in eq. (6.8.12), every vertex v in the unique ω_q -cluster of density $\geq \delta$ satisfies

$$\begin{aligned} |\partial K_o(\omega_p) \cap \partial K_v(\omega_q)| &\geq |\mathcal{S}(\omega_q) \cap K_o(\omega_p)| \cdot \alpha \deg G \\ &= |K_o(\omega_p)| \cdot \|\mathcal{S}(\omega_q)\|_{K_o(\omega_p)} \cdot \alpha \deg G \\ &\geq k \cdot \frac{\eta^2}{128} \cdot \alpha \deg G. \end{aligned}$$

Therefore, there must exist a deterministic vertex v satisfying

$$(*) := \mathbb{P} \left(|\partial K_o(\omega_p) \cap \partial K_v(\omega_q)| \geq \frac{\alpha \eta^2 k}{128} \deg G \text{ and } o \not\leftrightarrow v \right) \geq \frac{\eta^2 \delta}{300}.$$

On the other hand, by Lemma 6.5.3,

$$(*) \leq \exp\left(-(p-q)\frac{\alpha\eta^2 k}{128} \deg G\right) \leq \exp\left(-\frac{\delta}{2} \cdot \frac{1}{2\deg G} \cdot \frac{\alpha\eta^2 k}{128} \deg G\right) = \exp\left(-\frac{\alpha\eta^2 \delta k}{512}\right).$$

We now have a contradiction because our choice of k was large enough to ensure that

$$\exp\left(-\frac{\alpha\eta^2 \delta k}{512}\right) < \frac{\eta^2 \delta}{300}.$$

□

6.9 The macromolecular case

In this section, we describe the asymptotic behaviour of the density of the giant cluster for supercritical percolation on high-degree finite transitive graphs that are macromolecular. Recall that $\mathcal{H} \subseteq \mathcal{F}$ is said to be macromolecular if $\lim_{G \in \mathcal{H}} \deg G = \infty$ and there exists a constant $\varepsilon > 0$ such that all but finitely many $G \in \mathcal{H}$ admit an ε -macromolecular decomposition $(A(G), B(G))$. We will often abbreviate $(A, B) := ((A(G), B(G)) : G \in \mathcal{H})$, noting that at most finitely many of these pairs $(A(G), B(G))$ might be undefined. Let us say that \mathcal{H} is *irreducibly* macromolecular if additionally, these ε -macromolecular decompositions $(A(G), B(G))$ can be chosen in such a way that there does not exist an infinite subset $\mathcal{H}' \subseteq \mathcal{H}$ such that $\{A(G) : G \in \mathcal{H}'\}$ is macromolecular. In this case, let us also call (A, B) an *irreducibly* macromolecular decompositon. In this section, we will focus on irreducibly macromolecular families. We will later show that every macromolecular family contains an infinite family that is irreducibly macromolecular (and in fact, the corresponding macromolecular decompositions are essentially uniquely determined). As such, general macromolecular families will inherit the results of this section.

Our main goal is to prove the following proposition, which establishes that this asymptotic density is determined by the local geometry of the quotient graphs B .

Proposition 6.9.1. *Let $\mathcal{H} \subseteq \mathcal{F}$ be an infinite set of (isomorphism classes of) finite, connected, vertex transitive graphs with $\lim_{G \in \mathcal{H}} \deg G = \infty$. Suppose that $p : \mathcal{H} \rightarrow [0, 1]$ is a supercritical assignment satisfying $\sup_{\mathcal{H}} p \deg G < \infty$. For each $G \in \mathcal{H}$, define*

$$\psi(G) := \text{mf}(p(G) \deg G) \quad \text{and} \quad p^*(G) := p(G)\psi(G)^2.$$

If \mathcal{H} is sparse and admits an irreducible macromolecular decomposition (A, B) , then

$$\lim_{n \rightarrow \infty} \limsup_{\mathcal{H}} \left| \theta(p, G) - \psi \cdot \mathbb{P}_{p^*}^B (|K_o| \geq n) \right| = 0. \quad (6.9.1)$$

In the first subsection, we will use the result of Section 6.7 to show that in the setting of this proposition, under \mathbb{P}_p^G , o is unlikely to belong to a mesoscopic cluster whose volume is much larger than $\deg G$. In the second subsection, we prove Proposition 6.9.1. In the third subsection, we will explain how Proposition 6.9.1 implies our main theorems concerning concentration (Theorem 6.1.3) and equicontinuity (Theorem 6.1.11) for graphs that admit irreducible macromolecular decompositions.

Warning: In this section, we will apply our main results about *non*-macromolecular graphs in order to analyse the family of graphs A coming from an irreducibly macromolecular decomposition (A, B) . More precisely, we will apply both Proposition 6.9.1, establishing a uniform tail on the distribution of non-giant clusters, as well as the consequences of this proposition that are derived in Section 6.11, namely the results in column 4 of the table in that section, establishing that the supercritical giant cluster density is concentrated and local. The proofs of these results about non-macromolecular graphs do not rely on the analysis of macromolecular graphs appearing in this section, so there is no danger of circular reasoning. In fact, the analysis in this section only relies on these results for non-macromolecular families of graphs that are also *dense*, since the family of graphs A coming from a macromolecular decomposition is always dense by definition. For this special case of dense non-macromolecular graphs, all of our results could be deduced much more easily and directly from Theorem 6.1.1 and a refinement of Lemma 6.11.3, without having to go through the work of Section 6.8.

Proof of Proposition 6.9.1

Our goal is to establish eq. (6.9.1). It suffices to establish that every infinite subset of \mathcal{H} contains a further infinite subset for which these equations hold. As such, throughout this proof we may without loss of generality replace \mathcal{H} by an arbitrary infinite subset of \mathcal{H} .

We would like to apply our results from Section 6.8 concerning high-degree non-macromolecular graphs to $\mathcal{A} := \{A(G) : G \in \mathcal{H}\}$. It follows easily from the definition of irreducible macromolecular decomposition that $\lim_{\mathcal{H}} \deg A = \infty$ (because $\deg G - \deg A$ is uniformly bounded above) and \mathcal{A} does not contain an infinite subset that is macromolecular. What is less obvious is that p is supercritical for \mathcal{A} . Strictly speaking, it does not make sense to ask whether p is supercritical for \mathcal{A} because graphs in \mathcal{A} do not lie in the domain of p . To avoid this technicality, by passing to an infinite subset of \mathcal{H} if necessary, let us assume that the map $\mathcal{H} \rightarrow \mathcal{A}$ sending $G \mapsto A(G)$ is injective and hence a bijection. This lets us identify our given assignment of parameters $p : \mathcal{H} \rightarrow [0, 1]$ with an assignment $p : \mathcal{A} \rightarrow [0, 1]$.

Claim 6.9.2. *The assignment $p : \mathcal{A} \rightarrow [0, 1]$ is supercritical for \mathcal{A} .*

Proof. Since \mathcal{A} is a family of dense graphs, the assignment $p_c := 1/\deg A$ defines a percolation threshold for \mathcal{A} [Bol+10b]. (See discussion in Section 6.1). So our goal is to show that

$$\liminf_{\mathcal{H}} p \deg A > 1. \quad (6.9.2)$$

By Lemma 6.2.2, since p is supercritical for \mathcal{H} ,

$$\liminf_{\mathcal{H}} p \deg G > 1. \quad (6.9.3)$$

Since $\deg G - \deg A$ is uniformly bounded above and $\lim_{\mathcal{H}} \deg G = \infty$,

$$\lim_{\mathcal{H}} \frac{\deg A}{\deg G} = 1. \quad (6.9.4)$$

Combining eqs. (6.9.3) and (6.9.4) yields eq. (6.9.2). \square

We are now able to invoke Proposition 6.8.2 from Section 6.7 to analyse \mathbb{P}_p^A . We will also invoke the main result of Section 6.8, concerning a uniform tail for non-macromolecular graphs, and its consequences explained in Section 6.11, specifically, the validity of the mean-field approximation and the concentration of the giant cluster cluster density for non-macromolecular graphs. Our next step is to leverage this control of percolation on A to build a suitable coupling of \mathbb{P}_p^A and $\mathbb{P}_{p^*}^{B^*}$. Informally, we want this coupling to confirm the picture that mesoscopic clusters in \mathbb{P}_p^G with volume of order $|A|$ behave like microscopic clusters with volume of order 1 in $\mathbb{P}_{p^*}^{B^*}$. The idea is that these mesoscopic clusters in \mathbb{P}_p^G must be essentially built from joining the “local” giant clusters from a bounded number of the copies of A in G , and these connections between “local” giant clusters behave like $\mathbb{P}_{p^*}^{B^*}$. Let us now introduce the notation required to make this precise.

For each $G \in \mathcal{H}$, let $u \mapsto [u]$ denote the class function for an $\text{Aut } G$ -invariant equivalence relation on V that induces the macromolecular decomposition (A, B) . We naturally identify the set of classes $\{[u] : u \in V\}$ with the vertex set $V(B)$ of B . Let F be the set of all edges $xy \in E$ such that $[x] \neq [y]$. For each $[u] \in V(B)$, make the following definitions: Let $E[u]$ be the set of all edges $xy \in E$ such that $x, y \in [u]$. Let $\omega[u] := \omega|_{E[u]}$, and for each $v \in [u]$, let $K_v[u] := K_v(\omega[u])$. Let \vec{F} be the set of all *ordered* pairs (u, v) such that the corresponding unordered pairs $uv = \{u, v\}$ belongs to F . Let $\mathbb{Q} = \mathbb{Q}^G$ be the law of a random vector $\beta \in \{0, 1\}^{\vec{F}}$ whose entries are iid Bernoulli(ψ).

Thanks to Claim 6.9.2, after passing to an infinite subset of \mathcal{H} if necessary, we may assume that there exists a constant $\varepsilon > 0$ such that for all $A \in \mathcal{A}$, the parameter $p(A)$ is ε -supercritical for A .

Fix some choice of ε for the rest of this proof. Define $m := |A| = |[o]|$. For each $[u] \in V(B)$, define

$$r[u] := \{v \in [u] : |K_v[u]| \geq \varepsilon m\},$$

and given vectors $\omega \in \{0, 1\}^E$ and $\beta \in \{0, 1\}^{\vec{F}}$, let $\mathcal{E}[u]$ be the assertion that for every pair $(x, y) \in \vec{F}$ that satisfies $x \in [u]$ and $\omega(x, y) = 1$,

$$\beta(x, y) = 1 \quad \text{if and only if} \quad x \in r[u].$$

Claim 6.9.3. *For each $G \in \mathcal{H}$, there exists a choice of coupling $\mathbf{P} = \mathbf{P}^G$ of $\omega \sim \mathbb{P}_p^G$ and $\beta \sim \mathbb{Q}^G$ in which $\omega|_F$ and β are independent of each other such that*

$$\lim_{\mathcal{H}} \min_{[u] \in V(B)} \mathbf{P}(\mathcal{E}[u]) = 1. \quad (6.9.5)$$

Proof. Fix $\eta > 0$ and let $G \in \mathcal{H}$. We will show that for all but finitely many choices for G , we can build suitable coupling that satisfies

$$\mathbf{P}(\mathcal{E}[u]) \geq 1 - \eta$$

for every $[u]$. Let n be a large positive integer to be determined. Independently sample the following families of random variables, and write \mathbf{P} for their joint law:

1. Sample $\omega|_F \sim \bigotimes_F \text{Bern}(p)$.
2. Independently for each $(u, v) \in \vec{F}$, sample

$$\nu_{(u,v)} \sim \bigotimes_{E[o]} \text{Bern}(p) \quad \text{and} \quad \beta(u, v) \sim \text{Bern}(\psi)$$

that are coupled so that

$$\Delta := |\mathbb{P}_p(|K_o[o]| \geq n) - \psi| = \mathbf{P} \left(\underbrace{\{|K_o(\nu_{(u,v)})| \geq n\} \triangle \{\beta(u, v) = 1\}}_{\mathcal{E}(u,v)} \right).$$

3. Independently for each $(u, v) \in \vec{F}$, sample

$$\phi_{(u,v)} \sim \text{Unif}(\Gamma_u)$$

where Γ_u is the set of all graph automorphisms of G that map $o \mapsto u$.

4. Sample $v \sim \bigotimes_E \text{Bern}(p)$.

For each $(u, v) \in \vec{F}$, let

$$\hat{\nu}_{(u,v)} : E[o] \rightharpoonup \{0, 1\}$$

be the *partial function* encoding the edges revealed in an exploration of the cluster at o from inside (with respect to an arbitrary deterministic ordering of $E[o]$) that is halted as soon as the event that $|K_o(\nu_{(u,v)})| \geq n$ is determined by the states of the revealed edges. (See [Eas24, Section 2.1] for more discussion of these partial functions and for a warm-up to the argument below.) Given partial functions f and g , we write $f \sqcup g$ for the *override*, i.e. the partial function with domain $\text{dom}(f \sqcup g) = \text{dom}(f) \cup \text{dom}(g)$ that coincides with f on $\text{dom}(f) \setminus \text{dom}(g)$ and coincides with g on $\text{dom}(g) \setminus \text{dom}(f)$. Let x_1, \dots, x_r be an enumeration of the set of all pairs $(u, v) \in \vec{F}$ satisfying $\omega(uv) = 1$, listed according to an arbitrary but deterministic total order on \vec{F} that is fixed ahead of time. Now set

$$\omega := \omega_F \sqcup (\hat{\nu}_{x_1} \circ \phi_{x_1}) \sqcup \dots \sqcup (\hat{\nu}_{x_r} \circ \phi_{x_r}) \sqcup v.$$

Observe that ω and β have the required marginals and independence properties, so what remains is to establish the required control of $\mathcal{E}[u]$ for every $[u]$. We will focus on $\mathcal{E}[o]$, but the same arguments work for all $[u]$. Let \mathcal{F}_1 be the event that $\mathcal{E}(u, v)$ holds for some $(u, v) \in \vec{F}$ with $v \in [o]$, and let \mathcal{F}_2 be the event that there exist $(u, v), (u', v') \in \vec{F}$ such that $v, v' \in [o]$ and $\omega(uv) = \omega(u'v') = 1$ but $\hat{\nu}_{(u,v)} \circ \phi_{(u,v)}$ and $\hat{\nu}_{(u',v')} \circ \phi_{(u',v')}$ are *incompatible* as partial functions, i.e. they disagree on some portion of the intersection of their domains. Note that

$$\mathcal{E}[o] \subseteq \mathcal{F}_1 \cup \mathcal{F}_2.$$

So it suffices to show that $\mathbf{P}(\mathcal{F}_1), \mathbf{P}(\mathcal{F}_2) \leq \frac{\eta}{2}$.

Control of \mathcal{F}_1 Let z denote the number of edges in $\partial[o]$ that are open in ω . By a union bound, almost surely,

$$\mathbf{P}(\mathcal{F}_1 \mid z) \leq z\Delta.$$

So by another union bound and by Markov's inequality, for all $\lambda > 0$,

$$\mathbf{P}(\mathcal{F}_1) \leq \frac{\mathbf{E}[z]}{\lambda} + \lambda\Delta.$$

Note that for some constant $C < \infty$ independent of G , we have $\mathbf{E}[z] = p \deg B < \infty$ because $\sup_{\mathcal{H}} p \deg G < \infty$ by hypothesis, and $\sup_{\mathcal{H}} \frac{\deg B}{\deg G} < \infty$ by definition of macromolecular decomposition. Fix $\lambda := 4C/\eta$ so that $\frac{\mathbf{E}[z]}{\lambda} \leq \frac{\eta}{4}$. By continuity of mf and by eq. (6.9.4),

$$\lim_{\mathcal{H}} |\text{mf}(p \deg G) - \text{mf}(p \deg A)| = 0.$$

So by Theorems 6.1.1 and 6.2.1 and Proposition 6.8.1 and our results about non-macromolecular graphs (specifically, $(T, 4)$ and $(L, 4)$ in Section 6.11),

$$\lim_{l \rightarrow \infty} \limsup_{\mathcal{H}} \left| \mathbb{P}_p (|K_o[o]| \geq l) - \psi \right| = 0.$$

It follows that by picking our constant n to be sufficiently large, we can guarantee that $\Delta \leq \frac{\eta}{4\lambda}$ for all but finitely many G . Fix such a choice of n for the rest of the proof. By combining our bounds, we now have

$$\mathbf{P}(\mathcal{F}_1) \leq \frac{\eta}{4} + \frac{\eta}{4} = \frac{\eta}{2}.$$

Control of \mathcal{F}_2 For each $(u, v) \in \vec{F}$, consider the cluster

$$C_{(u,v)} := K_v (\hat{\nu}_{(u,v)} \circ \phi_{(u,v)}),$$

where any edges with undefined state are treated as closed. Notice that for all $x, y \in \vec{F}$, the clusters C_x and C_y are independent of $\omega|_F$, and if $\hat{\nu}_x \circ \phi_x$ and $\hat{\nu}_y \circ \phi_y$ are incompatible then $C_x \cap C_y \neq \emptyset$. So by a union bound and independence,

$$\mathbf{P}(\mathcal{F}_2) \leq \sum_{x,y} p^2 \mathbf{P}(C_x \cap C_y \neq \emptyset)$$

where the sum is over all $x = (u, v) \in \vec{F}$ and $y = (u', v') \in \vec{F}$ such that $v, v' \in [o]$. Consider some $x = (u, v)$ and $y = (u', v')$. Let \mathcal{K} be the set of all possible outcomes for C_x for any (equivalently every) $x \in \vec{F}$. By independence, we can expand

$$\mathbf{P}(C_x \cap C_y \neq \emptyset) = \sum_{X,Y \in \mathcal{K}} \mathbf{P}(C_x = X) \mathbf{P}(C_y = Y) \cdot \frac{|\{(f, g) \in \Gamma_u \times \Gamma_{u'} : f(X) \cap g(Y) \neq \emptyset\}|}{|\Gamma_u| |\Gamma_{u'}|}.$$

Now by summing over all choices for x and y and exchanging the order of summation,

$$\mathbf{P}(\mathcal{F}_2) \leq (p |A|)^2 \left(\frac{\deg B}{|A|} \right)^2 \min_{X,Y \in \mathcal{K}} \mathcal{P}(f(X) \cap g(X) \neq \emptyset)$$

where \mathcal{P} is the law of two independent uniformly random graph automorphisms f and g of $G[o]$. Recall that $p \deg B \leq C$. Consider some $X, Y \in \mathcal{K}$. By a union bound and independence,

$$\mathcal{P}(f(X) \cap g(X) \neq \emptyset) \leq \sum_{u \in [o]} \mathcal{P}(u \in f(X)) \mathcal{P}(u \in g(X)).$$

Notice that for all $u \in [o]$, we have $u \in f(X)$ if and only if $f^{-1}(u) \in X$, the law of f^{-1} is the same as the law of f , and the law of $f(u)$ is uniform over $[o]$. In particular, for all $u \in [o]$,

$$\mathcal{P}(u \in f(X)) = \frac{|X|}{|A|} \leq \frac{n}{|A|},$$

and the same is true of $g(Y)$. So

$$\mathcal{P}(f(X) \cap g(X) \neq \emptyset) \leq \frac{n^2}{|A|},$$

and since X and Y were arbitrary,

$$\mathbf{P}(\mathcal{F}_2) \leq \frac{C^2 n^2}{|A|},$$

which is smaller than $\frac{\eta}{2}$ for all but finitely many G because $\lim_{\mathcal{H}} \deg A = \infty$.

□

Fix such a family of couplings $(\mathbf{P}^G : G \in \mathcal{H})$ for the rest of the proof. Given some $G \in \mathcal{H}$, we can naturally identify F with the edge set of B . Now our coupling \mathbf{P} induces the following pair of percolation configurations $\tilde{\omega}, \hat{\omega} \in \{0, 1\}^F$ on B : for all $uv \in F$,

$$\hat{\omega}(uv) := \omega(uv) \mathbf{1}_{r[u]}(u) \mathbf{1}_{r[v]}(v) \quad \text{and} \quad \tilde{\omega}(uv) := \omega(uv) \beta(u, v) \beta(v, u).$$

Note that $\hat{\omega}$ only depends on ω , and by construction, the law of $\tilde{\omega}(u, v)$ is exactly $\mathbb{P}_{p^*}^B$. For each $[u] \in V(B)$, let $\Pi[u]$ be the set of all edges $xy \in F$ such that $x \in [u]$ and $\omega(xy) = 1$, let $s[u]$ be the set of all vertices $v \in V \setminus [u]$ such that $v \notin r[v]$ and there exists some $xy \in F$ with $x \in [u]$ and $y \in K_v[v]$ satisfying $\omega(xy) = 1$, and let $t[u]$ be the set of all vertices $v \in s[u]$ such that there exists $x \in V \setminus ([v] \cup r[u])$ with $vx \in E$ satisfying $\omega(vx) = 1$.

Claim 6.9.4. *We have*

$$\lim_{i \rightarrow \infty} \limsup_{\mathcal{H}} \mathbb{P}_p (|\Pi[o]| > i) = 0;$$

$$\lim_{i \rightarrow \infty} \limsup_{\mathcal{H}} \mathbb{P}_p (|s[o]| > i) = 0;$$

$$\lim_{\mathcal{H}} \mathbb{P}_p (t[o] \neq \emptyset) = 0.$$

Proof. We will prove each of the three equations in turn.

First equation For all G , we have $\mathbb{E}_p |\Pi[o]| = p \deg B$. As noted in the proof of the previous claim, it follows easily from the definition of macromolecular decomposition and from $\sup_{\mathcal{H}} p \deg G < \infty$ that

$$\sup_{\mathcal{H}} p \deg B < \infty.$$

So by Markov's inequality,

$$\limsup_{i \rightarrow \infty} \limsup_{\mathcal{H}} \mathbb{P}_p (|\Pi[o]| > i) \leq \limsup_{i \rightarrow \infty} \frac{1}{i} \sup_{\mathcal{H}} p \deg G = 0. \quad (6.9.6)$$

Second equation Let $G \in \mathcal{H}$ and $i \geq 1$. Let X be the set of all vertices $y \in V \setminus [o]$ such that there exists $x \in [o]$ satisfying $xy \in F$ and $\omega(xy) = 1$. The event that $|s[o]| > i$ is contained in the union

$$\mathcal{F}_1(i) \cup \mathcal{F}_2(i)$$

where $\mathcal{F}_1(i)$ is the event that $|X| \geq \sqrt{i}$, and $\mathcal{F}_2(i)$ is the event that there exists $xy \in \Pi[o]$ with $x \in [o]$ and $y \notin r[y]$ such that $|K_y[y]| \geq \sqrt{i}$. Since $|X| \leq |\Pi[o]|$ always holds, it follows from our analysis above of the first equation that

$$\sup_{\mathcal{H}} \mathbb{E}_p |X| < \infty \quad (6.9.7)$$

and

$$\lim_{j \rightarrow \infty} \limsup_{\mathcal{H}} \mathbb{P}_p (\mathcal{F}_1(j)) = 0.$$

By a union bound, independence, and transitivity,

$$\begin{aligned} \mathbb{P}_p (\mathcal{F}_2(i)) &\leq \mathbb{E}_p \left[\sum_{x \in V} \mathbf{1}_{x \in X} \mathbf{1}_{|K_x[x]| \geq \sqrt{i}} \mathbf{1}_{x \notin r[x]} \right] \\ &\leq \mathbb{E}_p |X| \cdot \mathbb{P}_p (|K_o[o]| \geq \sqrt{i} \text{ and } o \notin r[o]). \end{aligned}$$

By Theorems 6.1.1 and 6.2.1 and Proposition 6.8.1,

$$\lim_{j \rightarrow \infty} \mathbb{P}_p (|K_o[o]| \geq \sqrt{i} \text{ and } o \notin r[o]) = 0,$$

and hence by eq. (6.9.7),

$$\lim_{j \rightarrow \infty} \mathbb{P}_p (\mathcal{F}_1(j)) = 0.$$

The conclusion now follows by a union bound.

Third equation Let I be the set of all edges $xy \in E$ such that $[x] = [y]$ or $\{x, y\} \cap r[o] \neq \emptyset$. Let Y be the set of all vertices $y \in V \setminus [o]$ such that there exists $x \in V \setminus ([y] \cup r[o])$ satisfying $\omega(xy) = 1$. Note that

$$t[o] \neq \emptyset \quad \text{if and only if} \quad Y \cap s[o] \neq \emptyset.$$

For all $y \in V$, by a union bound we almost surely have

$$\mathbb{P}_p (y \in Y \mid \omega|_I) \leq p (\deg G - \deg A) =: \Lambda$$

because after having revealed the states of all edges in I , in order for y to ultimately belong to Y , at least one of the at most $(\deg G - \deg A)$ -many unrevealed edges $e \in F$ with $y \in e$ must turn out to be open. So by a union bound and since $s[o]$ is determined by $\omega|_I$, we almost surely have

$$\begin{aligned}\mathbb{P}_p(t[o] \neq \emptyset \mid \omega|_I) &\leq \mathbb{E}_p \left[\sum_{y \in V} \mathbf{1}_{y \in Y} \mathbf{1}_{y \in s[o]} \mid \omega|_I \right] \\ &= \sum_{y \in s[o]} \mathbb{P}_p(y \in Y \mid \omega|_I) \leq \Lambda |s[o]|.\end{aligned}$$

So by another union bound,

$$\mathbb{P}_p(t[o] \neq \emptyset) \leq \mathbb{P}_p \left(|s[o]| \geq \frac{1}{\sqrt{\Lambda}} \right) + \sqrt{\Lambda}. \quad (6.9.8)$$

By definition of macromolecular decomposition, $\deg G - \deg A$ is bounded above uniformly in G . Moreover, $\lim_{\mathcal{H}} p = 0$ because $\lim_{\mathcal{H}} \deg G = \infty$ and $\sup_{\mathcal{H}} p \deg G < \infty$. Therefore,

$$\lim_{\mathcal{H}} \Lambda = 0. \quad (6.9.9)$$

In particular, thanks to our analysis above of the second equation,

$$\lim_{\mathcal{H}} \mathbb{P}_p \left(|s[o]| \geq \frac{1}{\sqrt{\Lambda}} \right) = 0. \quad (6.9.10)$$

The conclusion now follows by combining eqs. (6.9.8) to (6.9.10). \square

We will now use these simple properties and a mass-transport argument to relate the distribution of microscopic clusters in $\mathbb{P}_{p^*}^B$ to the distribution of mesoscopic clusters in \mathbb{P}_p^G . Let

$$k := \frac{1}{m\psi} \left| \bigcup_{v \in r[o]} K_v(\omega) \right|,$$

and let $\text{Round}(k)$ be the integer closest to k , rounding up in case of a tie.

Claim 6.9.5. *For every positive integer n ,*

$$\lim_{\mathcal{H}} \left| \mathbb{P}_p^G(\text{Round}(k) = n) - \mathbb{P}_{p^*}^B(|K_{[o]}| = n) \right| = 0.$$

Proof. Let n be a positive integer, let $\eta > 0$, and let $G \in \mathcal{H}$. Note that

$$\Delta := \left| \mathbb{P}_p^G(\text{Round}(k) = n) - \mathbb{P}_{p^*}^B(|K_{[o]}| = n) \right| \leq \mathbf{P} \left(\underbrace{\{\text{Round}(k) = n\} \Delta \{\tilde{k} = n\}}_{\mathcal{F}} \right) \quad (6.9.11)$$

where \tilde{k} denotes the size of the cluster of $[o]$ in $\tilde{\omega}$, which we view as a configuration on B . Our goal is to show that $\Delta \leq \eta$ for all but finitely many choices of G . Given $i \geq 1$, say that $[u]$ is i -good if all of the following events hold:

- $\omega[u]$ has a unique cluster K satisfying $|K| \geq \varepsilon m$.
- $|r[u] - \psi m| \leq \frac{\psi}{16n}$.
- $|\Pi[u]| \leq i$, $|s[u]| \leq i$, and $t[u] = \emptyset$.

By Theorems 6.1.1 and 6.2.1, our main results about non-macromolecular graphs (row 4 in Section 6.11), and Claim 6.9.4, there exists a constant $i \in \mathbb{N}$ such that

$$\mathbb{P}_p ([o] \text{ is } i\text{-good}) \leq \frac{\eta}{2n} \quad (6.9.12)$$

for all but finitely many choices for G . Fix such a choice of i for the rest of the proof, and simply say that if $[u]$ is *good* to mean that u is i -good. Let X be the cluster of $[o]$ in $\tilde{\omega} \cap \hat{\omega}$, and let Y be the set of all edges $xy \in F$ such that $[x] \in X$ or $[y] \in X$.

Say that an edge $xy \in F$ is *intact* if $\mathcal{E}[x]$ holds and $\mathcal{E}[y]$ holds, and note that this implies that $\hat{\omega}(xy) = \tilde{\omega}(xy)$. In particular, if every edge in Y is intact, then

$$\hat{k} = \tilde{k}, \quad (6.9.13)$$

where \hat{k} denotes the size of the cluster of $[o]$ in $\hat{\omega}$. We claim that if additionally every element of X is good then \mathcal{F} cannot hold. Indeed, if every element of X is good, then $km\psi$, which by definition of k equals the number of vertices ω -connected to $[o]$, satisfies

$$\sum_{[v] \in K_{[o]}(\hat{\omega})} |r[v]| \leq km\psi \leq \sum_{[v] \in K_{[o]}(\hat{\omega})} (|r[v]| + |s[v]|),$$

and moreover,

$$\left(1 - \frac{1}{16n}\right)\psi m \hat{k} \leq km\psi \leq \left(1 + \frac{1}{16n}\right)\psi m \hat{k} + i\hat{k}.$$

So by eq. (6.9.13) and our hypothesis that $\lim_{\mathcal{H}} m = \infty$, for all but finitely many choices of G ,

$$|k - \tilde{k}| \leq \tilde{k} \left(\frac{1}{16n} + \frac{i}{m\psi} \right) \leq \frac{\tilde{k}}{8n}. \quad (6.9.14)$$

By manipulating eq. (6.9.14), we have $\tilde{k} \leq 2k$. So by applying eq. (6.9.14) again,

$$|k - \tilde{k}| \leq \frac{k}{4n}. \quad (6.9.15)$$

Notice that if $\tilde{k} = n$, then by eq. (6.9.14), we have $|k - \tilde{k}| < \frac{1}{2}$, and hence $k = \tilde{k}$. Similarly, if $\text{Round}(k) = n$, then $k \leq n + 1$ and hence $k = \tilde{k}$ by eq. (6.9.15). So \mathcal{F} cannot hold, as claimed.

Let \mathcal{F}_1 be the event that \mathcal{F} holds but not every element of X is good, and let $\mathcal{F}_2 := \mathcal{F} \setminus \mathcal{F}_1$. By our work above, on the event \mathcal{F}_2 , not every edge in Y is intact. For each $[u] \in V(B)$, let $\mathcal{F}[u]$, $\mathcal{F}_1[u]$, $\mathcal{F}_2[u]$, $X[u]$, and $Y[u]$ be the analogues of \mathcal{F} , \mathcal{F}_1 , etc. defined with $[u]$ in place of $[o]$. Given classes $[u]$, $[v]$, let $f_1([u], [v]) \in \{0, 1\}$ be the indicator for the event that $\mathcal{F}_1[u]$ holds, $[v] \in X[u]$, and $[v]$ is not good. Given a class $[u]$ and an edge $e \in F$, let $f_2([u], e) \in \{0, 1\}$ be the indicator for the event that $\mathcal{F}_2[u]$ holds, $e \in Y[u]$, and e is not intact. We know that for all $[u] \in V(B)$, if $\mathcal{F}[u]$ holds then $f_i([u], x) = 1$ for some x and i . So by eq. (6.9.11) and transitivity,

$$\underbrace{\frac{1}{|B|} \mathbf{E} \left[\sum_{[u], [v]} f_1([u], [v]) \right]}_{\Gamma_1} + \underbrace{\frac{1}{|B|} \mathbf{E} \left[\sum_{[u], e} f_2([u], e) \right]}_{\Gamma_2} \geq \Delta, \quad (6.9.16)$$

where the first sum is over all $[u], [v] \in V(B)$ and the second sum is over all $[u] \in V(B)$ and $e \in F$. (Recall that $|B|$ is the number of *vertices* in B .) So it suffices to show that $\Gamma_1, \Gamma_2 \leq \frac{\eta}{2}$.

For every $[u]$, on the event $\mathcal{F}[u]$, we have $|X[u]| \leq n$. So for every class $[v]$, there are never more than n classes $[u]$ satisfying $f_1([u], [v]) = 1$, and of course there are *no* such classes $[u]$ when $[v]$ is good. So for all $[u]$,

$$\mathbf{E} \left[\sum_{[u]} f_1([u], [v]) \right] \leq n \mathbf{P}([u] \text{ is not good}) \leq \frac{\eta}{2}, \quad (6.9.17)$$

where the second inequality comes from eq. (6.9.12). In particular, $\Gamma_1 \leq \frac{\eta}{2}$. Similarly, for every $[u]$, since on the event $\mathcal{F}_2[u]$ we have

$$|Y[u]| \leq \sum_{[v] \in X[u]} |\Pi[v]| \leq ni$$

because every element of $X[u]$ is good, it follows that for all $e \in F$,

$$\mathbf{E} \left[\sum_{[u]} f_2([u], e) \right] \leq 2ni, \quad (6.9.18)$$

the factor of ‘2’ arising from the fact that the classes of the endpoints of the edge $e \in F$ could belong to two distinct $\hat{\omega} \cap \tilde{\omega}$ -clusters. In particular, $\Gamma_2 \leq \frac{\eta}{2}$, as required. \square

We are now ready to conclude the proof of eq. (6.9.1). Fix $\eta > 0$. By Proposition 6.7.1 and Lemma 6.5.2 (as well as Theorems 6.1.1 and 6.2.1), there exists a sufficiently large positive integer N_0 such that for every integer $N \geq N_0$, for all but finitely many $\mathcal{G} \in \mathcal{H}$,

$$|\theta(p, G) - \mathbb{P}_p^G (|K_o| \geq Lm)| \leq \frac{\eta}{4},$$

where $L = L(N, G) := \psi \left(N + \frac{1}{2} \right)$. By Theorems 6.1.1 and 6.2.1 and Proposition 6.8.1, for all but finitely many G ,

$$|\mathbb{P}_p^G (|K_o| \geq L) - \mathbb{P}_p^G (|K_o| \geq L \text{ and } o \in r[o])| \leq \frac{\eta}{4}.$$

Note that $|K_o| \geq L$ and $o \in r[o]$ if and only if $k \geq N$ and $o \in r[o]$. By transitivity,

$$\mathbb{P}_p^G (k \geq N \text{ and } o \in r[o]) = \frac{1}{m} \sum_{u \in [o]} \mathbb{P}_p^G (k \geq N \text{ and } u \in r[o]) = \mathbb{E}_p^G \left[\frac{|r[o]|}{m} \mathbf{1}_{k \geq N} \right].$$

By our main results for non-macromolecular graphs (row 4 in Section 6.11), for all but finitely many G ,

$$\left| \mathbb{E}_p^G \left[\frac{|r[o]|}{m} \mathbf{1}_{k \geq N} \right] - \psi \cdot \mathbb{P}_p^G (k \geq N) \right| \leq \mathbb{E}_p^G \left| \frac{|r[o]|}{m} - \psi \right| \leq \frac{\eta}{4}.$$

By Claim 6.9.5 and the fact that ψ is always bounded above by 1, for all but finitely many G ,

$$\left| \psi \cdot \mathbb{P}_p^G (k \geq N) - \psi \cdot \mathbb{P}_{p^*}^B (|K_o| \geq N) \right| \leq \frac{\eta}{4}.$$

By combining our bounds, we find that for all but finitely many G ,

$$|\theta(p, G) - \mathbb{P}_{p^*}^B (|K_o| \geq N)| \leq \eta,$$

as required.

6.10 The nonunimodular case

In this section we complete the proof of the bounded-degree case of our results by treating the *nonunimodular* case. Perhaps surprisingly, following [Hut20g], percolation is much better understood on nonunimodular transitive graphs than on unimodular transitive graphs. More specifically, we will briefly explain how *very strong* quantitative forms of all our main results can be deduced in the nonunimodular case from the results of [Hut20h], which establish sharp, quantitative tail bounds on the distribution of finite clusters under the L^2 *boundedness condition*, and [Hut20g], which imply that this condition holds in the nonunimodular case; our main task in this section will be to outline how uniform versions of these results can be deduced by a compactness argument. Nonunimodular

transitive graphs cannot arise as limits of finite transitive graphs by Corollary 6.10.2, and we will keep this section brief since it is tangential to the main focus of the paper.

For nonunimodular G , the **tilted mass-transport principle** states that

$$\sum_{x \in V} F(o, x) = \sum_{x \in V} F(x, o) \Delta(o, x) \quad (6.10.1)$$

for every $F : V^2 \rightarrow [0, \infty]$ that is invariant under the diagonal action of $\text{Aut}(G)$ on V^2 .

The following proposition, established in [Hut20b, Corollary 5.4], tells us that the modular function is determined by the local geometry of the graph. This proposition is proven using a probabilistic interpretation of the modular function as a Radon-Nikodym cocycle (see [Hut20g; BC12]) and is not at all obvious from the algebraic definition given above!

Proposition 6.10.1. *Let $(G_n)_{n \geq 1}$ be a sequence of connected, locally finite transitive graphs converging locally to some locally finite transitive graph G , and let $(o_n)_{n \geq 1}$ and o be vertices of $(G_n)_{n \geq 1}$ and G respectively. For each $r \geq 1$ there exists $N < \infty$ such that for every $n \geq N$ there exists an isomorphism ϕ_n from the ball of radius r around o_n in G_n to the ball of radius r around o in G satisfying*

$$\Delta_{G_n}(u, v) = \Delta_G(\phi(u), \phi(v))$$

for every u, v in the ball of radius r around o_n in G_n .

Together with the **cocycle identity** [Hut20g, Lemma 2.3], which states that $\Delta(x, z) = \Delta(x, y)\Delta(y, z)$ for every $x, y, z \in V$ (and hence that G is unimodular if and only if $\Delta(o, x) = 1$ for every neighbour of the origin), this proposition has the following immediate corollary.

Corollary 6.10.2. *For each $d \geq 1$, let \mathcal{G}_d be the space of all isomorphism classes of connected, transitive graphs of degree at most d , and let \mathcal{U}_d and $\mathcal{N}_d = \mathcal{G}_d \setminus \mathcal{U}_d$ be the sets of unimodular and nonunimodular elements of \mathcal{G}_d respectively. The sets \mathcal{U}_d and \mathcal{N}_d are both closed (and hence both open) in \mathcal{G}_d with respect to the local topology.*

Note that the family \mathcal{U}_d contains both finite and infinite graphs.

The following theorem is a “uniform in G ” version of a theorem whose non-uniform version follows from the results of [Hut20g; Hut20h]. The bound on the volume tail of finite clusters it yields is sharp for p not too close to 0 or 1 as explained in detail in [Hut20h].

Theorem 6.10.3. *Let $d \geq 1$ and let \mathcal{N}_d be the set of (isomorphism classes of) nonunimodular transitive graphs of degree at most d . There exist positive constants $c = c(d) > 0$ and $C = C(d) < \infty$ such that*

$$\mathbb{P}_p^G(n \leq |K_o| < \infty) \leq Cn^{-1/2} \exp\left(-c|p - p_c|^2 n\right)$$

for every $n \geq 1$, $G \in \mathcal{N}_d$, and $p \in [0, 1]$.

After sketching the proof of this theorem we will use it to deduce the following corollary, which is much stronger than the equicontinuity statements proven in the other parts of the paper.

Corollary 6.10.4. *Let $d \geq 1$ and let \mathcal{N}_d be the set of (isomorphism classes of) nonunimodular transitive graphs of degree at most d . For each $G \in \mathcal{N}_d$, $\theta(p, G)$ is an analytic function of p on $(p_c(G), 1]$ and there exists a constant $C = C(d)$ such that*

$$\frac{d}{dp}\theta(p, G) \leq \frac{C}{1-p} \quad \text{and} \quad \theta(p, G) \geq 1 - C|1-p|^C$$

for every $G \in \mathcal{N}_d$ and $p \in (p_c, 1]$.

We now introduce some relevant machinery from [Hut20g]. Let $G = (V, E)$ be a connected, locally finite, nonunimodular (vertex-)transitive graph, and let $\Delta = \Delta_G : V^2 \rightarrow (0, \infty)$ be the modular function of G . For each $\lambda \in \mathbb{R}$, the **tilted susceptibility** is defined to be

$$\chi_{p,\lambda} = \chi_{p,\lambda}^G = \sum_{x \in V} \mathbb{P}_p(o \leftrightarrow x) \Delta^\lambda(o, x),$$

so that $\chi_{p,0} = \mathbb{E}_p|K_o|$ is the ordinary susceptibility. It is a consequence of the tilted mass-transport principle that $\chi_{p,\lambda} = \chi_{p,1-\lambda}$, and since $\chi_{p,\lambda}$ is also a convex function of $\lambda \in \mathbb{R}$ this leads to a special role for the critically tilted susceptibility $\chi_{p,1/2} = \min_\lambda \chi_{p,\lambda}$. The **tiltability threshold** is defined to be

$$p_t(G) = \sup\{p \in [0, 1] : \chi_{p,1/2} < \infty\}.$$

The following is a special case of the main theorem of [Hut20g].

Theorem 6.10.5. *Let $G = (V, E)$ be a connected, locally finite, nonunimodular (vertex-)transitive graph. Then $p_c(G) < p_t(G)$.*

In order to prove Theorem 6.10.3, we will argue that Theorem 6.10.5 implies the following “uniform in G ” version of the same theorem.

Corollary 6.10.6. *For each $d \geq 1$ there exist positive constants $c = c(d)$ and $C = C(d)$ such that*

$$\chi_{p_c,1/2}^G \leq C \quad \text{and} \quad p_t(G) - p_c(G) \geq c$$

for every $G \in \mathcal{N}_d$.

To deduce Corollary 6.10.6 from Theorem 6.10.5, we will also apply the following theorem of [Hut20b], which is a consequence of the results of [Hut20g].

Theorem 6.10.7 ([Hut20b], Theorem 5.6). *The critical probability p_c defines a continuous function on \mathcal{N}_d . That is, if G_n is a sequence in \mathcal{N}_d converging to some $G \in \mathcal{N}_d$ then $p_c(G_n) \rightarrow p_c(G)$.*

We will also require the following lemma.

Lemma 6.10.8. *For each fixed $p \in [0, 1]$ and $\lambda \in \mathbb{R}$, the tilted susceptibility $\chi_{p,\lambda}^G$ defines a continuous function $\mathcal{N}_d \rightarrow [0, \infty]$. That is, if G_n is a sequence in \mathcal{N}_d converging to some $G \in \mathcal{N}_d$ then $\chi_{p,\lambda}^{G_n}$ converges to $\chi_{p,\lambda}^G$ as $n \rightarrow \infty$.*

We will prove Lemma 6.10.8 using a tilted version of the “ $\phi_p(S)$ argument” of Duminil-Copin and Tassion [DT16b]. Let G and Δ be as above. For each $\lambda \in \mathbb{R}$, $p \in [0, 1]$, and each finite set of vertices $S \ni o$ we consider the quantity $\phi_{p,\lambda}(S)$ defined by

$$\phi_{p,\lambda}(S) := p \sum_{x \in S} \mathbb{P}_p(o \xrightarrow{S} x) \sum_{y \sim x} \mathbb{1}(y \notin S) \Delta^\lambda(o, y) = p \sum_{x \in S} \mathbb{P}_p(o \xrightarrow{S} x) \Delta^\lambda(o, x) \sum_{y \sim x} \mathbb{1}(y \notin S) \Delta^\lambda(x, y),$$

where $\{o \xrightarrow{S} x\}$ denotes the event that o and x are connected by an open path only using vertices of S and the equality between these two expressions follows from the cocycle identity $\Delta(o, y) = \Delta(o, x)\Delta(x, y)$.

Lemma 6.10.9. *Let $G = (V, E)$ be a connected, locally finite, nonunimodular (vertex-)transitive graph, let o be a vertex of G and let $S \ni o$ be a finite set of vertices. Then*

$$\sum_{x \in S} \mathbb{P}_p(o \xrightarrow{S} x) \Delta^\lambda(o, x) \leq \chi_{p,\lambda} \leq \frac{1}{1 - \phi_{p,\lambda}(S)} \sum_{x \in S} \mathbb{P}_p(o \xrightarrow{S} x) \Delta^\lambda(o, x)$$

for every $p \in [0, 1]$ and $\lambda \in \mathbb{R}$ such that $\phi_{p,\lambda}(S) < 1$. In particular, if $\phi_{p,\lambda}(S) < 1$ for some finite set of vertices $S \ni o$ then $\chi_{p,\lambda} < \infty$.

Proof of Lemma 6.10.9. The first inequality is trivial; we focus on the second. If o is connected to some vertex z by an open path in G but not in S , this path must exit S for the first time using some edge $e = xy$ with $x \in S$ and $y \notin S$. On this event the three events $\{o \xleftrightarrow{S} x\}$, $\{e \text{ open}\}$, and $\{y \leftrightarrow \Lambda z\}$ occur disjointly, and we deduce from a union bound and the BK inequality that

$$\mathbb{P}_p(o \leftrightarrow z) \leq \mathbb{P}_p(o \xleftrightarrow{S} z) + p \sum_{x \in S} \sum_{y \notin S} \mathbb{P}_p(o \xleftrightarrow{S} x) \mathbb{P}_p(y \leftrightarrow z).$$

Let $\Lambda \subseteq V$ be a (finite or infinite) set of vertices containing S . Multiplying both sides by $\Delta^\lambda(o, z)$, summing over $z \in \Lambda$, and using the cocycle identity yields that

$$\begin{aligned} \sum_{z \in \Lambda} \mathbb{P}_p(o \leftrightarrow z) \Delta^\lambda(o, z) &\leq \sum_{x \in S} \mathbb{P}_p(o \xleftrightarrow{S} x) \Delta^\lambda(o, x) + p \sum_{x \in S} \sum_{y \notin S} \mathbb{P}_p(o \xleftrightarrow{S} x) \sum_{z \in \Lambda} \mathbb{P}_p(y \leftrightarrow z) \Delta^\lambda(o, z) \\ &\leq \sum_{x \in S} \mathbb{P}_p(o \xleftrightarrow{S} x) \Delta^\lambda(o, x) + \phi_{p, \lambda}(S) \sup_{y \in \Lambda} \sum_{z \in \Lambda} \mathbb{P}_p(y \leftrightarrow z) \Delta^\lambda(y, z). \end{aligned} \quad (6.10.2)$$

If we knew that $\chi_{p, \lambda} < \infty$ we could conclude by rearranging; a little care will be needed to conclude in a non-circular manner without this assumption.

Suppose that there exists a finite set $S \ni o$ with $\phi_{p, \lambda}(S) < 1$. For each finite set $\Lambda \subseteq V$ there exists a vertex w with $\sum_{z \in \Lambda} \mathbb{P}_p(w \leftrightarrow z) \Delta^\lambda(w, z) = \sup_{y \in \Lambda} \sum_{z \in \Lambda} \mathbb{P}_p(y \leftrightarrow z) \Delta^\lambda(y, z)$ and an automorphism γ of G sending o to w . Using that

$$\sum_{z \in \gamma^{-1}\Lambda} \mathbb{P}_p(x \leftrightarrow z) \Delta^\lambda(x, z) = \sum_{z \in \Lambda} \mathbb{P}_p(x \leftrightarrow \gamma^{-1}z) \Delta^\lambda(x, \gamma^{-1}z) = \sum_{z \in \Lambda} \mathbb{P}_p(\gamma x \leftrightarrow z) \Delta^\lambda(\gamma x, z)$$

for every $x \in V$, we have that

$$\sum_{z \in \gamma^{-1}\Lambda} \mathbb{P}_p(o \leftrightarrow z) \Delta^\lambda(o, z) = \sup_{y \in \gamma^{-1}\Lambda} \sum_{z \in \gamma^{-1}\Lambda} \mathbb{P}_p(y \leftrightarrow z) \Delta^\lambda(y, z).$$

Thus, (6.10.2) implies the inequality

$$\sum_{z \in \gamma^{-1}\Lambda} \mathbb{P}_p(o \leftrightarrow z) \Delta^\lambda(o, z) \leq \sum_{x \in S} \mathbb{P}_p(o \xleftrightarrow{S} x) \Delta^\lambda(o, x) + \phi_{p, \lambda}(S) \sum_{z \in \gamma^{-1}\Lambda} \mathbb{P}_p(o \leftrightarrow z) \Delta^\lambda(o, z),$$

which, since Λ is finite, can be rearranged to yield that

$$\sum_{z \in \gamma^{-1}\Lambda} \mathbb{P}_p(o \leftrightarrow z) \Delta^\lambda(o, z) \leq \frac{1}{1 - \phi_{p, \lambda}(S)} \sum_{x \in S} \mathbb{P}_p(o \xleftrightarrow{S} x) \Delta^\lambda(o, x)$$

when $\phi_{p,\lambda}(S) < 1$. On the other hand, the choice of γ ensures that

$$\sum_{z \in \Lambda} \mathbb{P}_p(o \leftrightarrow z) \Delta^\lambda(o, z) \leq \sum_{z \in \gamma^{-1}\Lambda} \mathbb{P}_p(o \leftrightarrow z) \Delta^\lambda(o, z),$$

and the claim follows by taking Λ to exhaust the entire vertex set V . \square

We now deduce Lemma 6.10.8 from Lemma 6.10.9 and Proposition 6.10.1.

Proof of Lemma 6.10.8. Fix $p \in [0, 1]$ and $\lambda \in \mathbb{R}$ and let $(G_n)_{n \geq 1}$ be a sequence in \mathcal{N}_d converging to some $G \in \mathcal{N}_d$. We need to prove that $\chi_{p,\lambda}^{G_n}$ converges to $\chi_{p,\lambda}^G$ as $n \rightarrow \infty$. Let $B_G(o, r)$ denote the ball of radius r around o in G . The *lower semicontinuity* statement $\liminf \chi_{p,\lambda}^{G_n} \geq \chi_{p,\lambda}^G$ is immediate from the fact that

$$\chi_{p,\lambda}^G = \sup_r \sum_{x \in B_G(o, r)} \mathbb{P}_p^G(o \xleftrightarrow{B_G(o, r)} x) \Delta_G^\lambda(o, x),$$

since each function on the right hand side is continuous in G by Proposition 6.10.1 and any supremum of continuous functions is lower semicontinuous. To conclude the proof it suffices to prove that $\limsup \chi_{p,\lambda}^{G_n} \leq \chi_{p,\lambda}^G$. The claim is trivial when $\chi_{p,\lambda}^G = \infty$ so we may assume that $\chi_{p,\lambda}^G < \infty$. Since the internal vertex boundaries of the balls $B_G(o, r)$ are disjoint for different choices of r , we have that

$$\begin{aligned} \chi_{p,\lambda}^G &\geq \frac{1}{\sum_{y \sim o} \Delta^\lambda(o, y)} \sum_{r=0}^{\infty} \sum_{x \in B_G(o, r)} \mathbb{P}_p(o \xleftrightarrow{B_G(o, r)} x) \Delta^\lambda(o, x) \sum_{y \sim x} \Delta^\lambda(x, y) \mathbb{1}(y \notin B_G(o, r)) \\ &= \frac{1}{p \sum_{y \sim o} \Delta^\lambda(o, y)} \sum_{r=0}^{\infty} \phi_{p,\lambda}(B_G(o, r)). \end{aligned}$$

Since $\chi_{p,\lambda}^G < \infty$ it follows that $\phi_{p,\lambda}(B_G(o, r)) \rightarrow 0$ as $r \rightarrow \infty$. On the other hand, Proposition 6.10.1 and Lemma 6.10.9 imply that

$$\limsup_{n \rightarrow \infty} \chi_{p,\lambda}^{G_n} \leq \frac{1}{1 - \phi_{p,\lambda}(B_G(o, r))} \sum_{x \in B_G(o, r)} \mathbb{P}_p^G(o \xleftrightarrow{B_G(o, r)} x) \Delta_G^\lambda(o, x)$$

for every r such that $\phi_{p,\lambda}(B_G(o, r)) < 1$, and the claim follows since the sum on the right hand side is bounded above by $\chi_{p,\lambda}^G$ for every $r \geq 1$. \square

Proof of Corollary 6.10.6. It suffices to prove the claim about $\chi_{p_c, 1/2}$, the claim about $p_t - p_c$ following from this estimate. Suppose for contradiction that the claim does not hold, so that there exists a sequence $(G_n)_{n \geq 1}$ in \mathcal{N}_d with $\chi_{p_c(G_n), 1/2}^{G_n} \rightarrow \infty$ as $n \rightarrow \infty$. By taking a subsequence

if necessary, we may assume that G_n converges to some $G \in \mathcal{N}_d$ as $n \rightarrow \infty$. It follows from Theorem 6.10.5 and Theorem 6.10.7 that $\lim_{n \rightarrow \infty} p_c(G_n) = p_c(G) < p_t(G)$ and hence that there exists $q < p_t(G)$ such that $p_c(G_n) \leq q$ for all sufficiently large n . Lemma 6.10.8 then implies that $\chi_{q,\lambda}^{G_n} \rightarrow \chi_{q,\lambda}^G < \infty$, which contradicts the assumption that $\chi_{p_c(G_n),1/2}^{G_n} \rightarrow \infty$ since $\chi_{p_c(G_n),1/2}^{G_n} \leq \chi_{q,\lambda}^{G_n}$ for all sufficiently large n . \square

Sketch of proof of Theorem 6.10.3. Let $G = (V, E)$ be a connected, locally finite, transitive graph, and for each $p \in [0, 1]$ define the **two-point matrix** $T_p \in \mathbb{R}^{V \times V}$ by $T_p(u, v) = \mathbb{P}_p(u \leftrightarrow v)$. We define $\|T_p\|_{2 \rightarrow 2}$ to be the operator norm of T_p on $L^2(V)$, which is infinite if T_p does not define a bounded operator on $L^2(V)$. The main result of [Hut20g] states that if G satisfies the **L^2 boundedness condition**, meaning that $\|T_{p_c}\|_{2 \rightarrow 2}$ is finite, then there exists $\delta > 0$

$$\mathbb{P}_p(n \leq |K_o| < \infty) \asymp n^{-1/2} \exp \left[-\Theta(|p - p_c|^2 n) \right] \quad (6.10.3)$$

for every $p \in (p_c - \delta, p_c + \delta)$ and $n \geq 1$. The following observations allow us to deduce Theorem 6.10.3 from this and Corollary 6.10.6:

1. The inequality $\|T_p\|_{2 \rightarrow 2} \leq \chi_{p,\lambda}$ holds for every $p \in [0, 1]$ and $\lambda \in \mathbb{R}$. This is proven in [Hut19b, Theorem 2.9]. Thus, Corollary 6.10.6 implies that the L^2 boundedness condition holds *uniformly* for nonunimodular transitive graphs, with constants depending only on the degree.
2. All the implicit constants in the results of [Hut20h] can be taken to depend only on the degree and on $\|T_{p_c}\|_{2 \rightarrow 2}$. (It may seem that they also depend on p_c , but this is redundant since p_c is bounded below by the reciprocal of the degree and bounded away from 1 by a constant depending only on $\|T_p\|_{2 \rightarrow 2}$.) Although such a claim is not made explicit in that paper, it can be verified by going through the proof and using that all estimates derived from the triangle condition (e.g. on the percolation probability θ and the intrinsic one-arm) can be taken to depend only on the degree and the value of $\|T_{p_c}\|_{2 \rightarrow 2}$. This fact can in turn be seen from the derivation of mean-field critical behaviour from the triangle condition presented in [Hut]. Indeed, [Hut, Eq. 2.1 and Lemma 2.1] yield that the susceptibility satisfies the differential inequality

$$\frac{d}{dp} \chi_p \geq \frac{1}{(1-p)(\log(1-p))^2} \mathbb{E}_p \left[\frac{(\chi_p |K| - \sum_{u,v \in K} T_p(u, v))^2}{\chi_p \sum_{u,v,w \in K} T_p(u, v) T_p(v, w)} \right],$$

and we can bound the two sums appearing here by

$$\sum_{u,v \in K} T_p(u, v) = \langle T_p \mathbf{1}, \mathbf{1} \rangle_K \leq \|T_p\|_{2 \rightarrow 2} \|\mathbf{1}\|_K^2 = \|T_p\|_{2 \rightarrow 2} |K|$$

and

$$\sum_{u,v,w \in K} T_p(u,v) T_p(u,w) = \langle T_p \mathbf{1}, T_p \mathbf{1} \rangle_K \leq \|T_p\|_{2 \rightarrow 2}^2 \|\mathbf{1}\|_K^2 = \|T_p\|_{2 \rightarrow 2}^2 |K|$$

to obtain that

$$\frac{d}{dp} \chi_p \geq \frac{(\chi_p - \|T_p\|_{2 \rightarrow 2})^2}{(1-p)(\log(1-p))^2 \|T_p\|_{2 \rightarrow 2}^2}.$$

This implies that $\varepsilon \chi_{p_c - \varepsilon}$ is bounded above and below by positive constants depending only on the degree and the value of $\|T_{p_c}\|_{2 \rightarrow 2}$; the fact that all other implicit constants appearing in the consequences of the triangle condition can be taken to depend only on the degree and $\|T_{p_c}\|_{2 \rightarrow 2}$ follows from this as can be seen from [Hut; Hut20h].

3. The restriction that $p \in (p_c - \delta, p_c + \delta)$ in (6.10.3) is only really needed for the *lower bound*, since the stated estimate is not sharp for p very close to 0 or 1. Indeed, the claimed upper bound extends to all $0 \leq p \leq p_c$ by monotonicity, and an upper bound of the same form follows for all $p_c + \delta \leq p \leq 1$ by [Hut20h, Proposition 3.1] and the methods of either [HH19] or [Hut23b].

We omit further details. \square

Remark 6.10.1. Most the analysis of [Hut20h] can be skipped if one only wishes to establish our main results in the nonunimodular case, rather than the sharp uniform volume-tail estimate of Theorem 6.10.3. Indeed, the proof of [Hut20h, Proposition 3.1] yields in the transitive case that

$$\Psi_p(S) := \frac{1}{1-p} \sum_{u \in S} \sum_{v \sim u} \mathbb{1}(v \notin S) \mathbb{P}_p(v \leftrightarrow \infty \text{ off } S) \geq \frac{\theta(p)|S|}{p(1-p)^2 \|T_p\|_{2 \rightarrow 2}^2}. \quad (6.10.4)$$

Since $\Psi_p(S)$ is monotone in p , it follows from this together with Corollary 6.10.6, [Hut19b, Lemma 2.4] (which bounds the effect of changing p on $\|T_p\|_{2 \rightarrow 2}$), and the mean-field lower bound that

$$\Psi_p(S) \geq c(p - p_c)|S|.$$

for every connected, locally finite, nonunimodular transitive graph $G = (V, E)$ and every finite set $S \subseteq V$, where c is a positive constant depending only on the degree of G . The fact that this implies a uniform upper bound on the volume tail of the form $\mathbb{P}_p(n \leq |K_o| < \infty) \leq \exp[-c(d, \varepsilon)n]$ for every $p \geq p_c + \varepsilon$ and $n \geq 1$ can easily be deduced from the methods of either [HH19] or [Hut23b].

Proof of Corollary 6.10.4. The analyticity of θ on $(p_c, 1]$ is already established in [HH19]. Since percolation with parameter $1 - (1-p)^n$ can be thought of as the union of n independent copies of

percolation with parameter p , we have that

$$\theta(1 - (1 - p)^n, G) \geq (1 - \theta(p, G))^n$$

for every $p \in (0, 1)$ and $n \geq 1$, which implies the claimed bound on $\theta(p, G)$ for p close to 1 in conjunction with the mean-field lower bound and the fact that p_c is bounded away from zero for elements of \mathcal{N}_d . It remains to prove the uniform upper bound on the derivative. Given a connected, locally finite graph G and $p \in [0, 1]$, the **triangle diagram** is defined by $\nabla_p := \sum_{x,y} \mathbb{P}_p(o \leftrightarrow x) \mathbb{P}_p(x \leftrightarrow y) \mathbb{P}_p(y \leftrightarrow o)$, and satisfies the inequality $\nabla_p \leq \|T_p\|_{2 \rightarrow 2}^3 \leq \chi_{p,1/2}^3$. The convergence of the triangle diagram at p_c is a well-known sufficient condition for mean-field percolation critical behaviour [AN84; BA91; Hut], and in particular, it follows from the results of [Hut] that if $\nabla_{p_c} < \infty$ then there exist positive constants c and C depending only on ∇_{p_c} and the degree of G such that

$$c(p - p_c) \leq \theta(p) \leq C(p - p_c) \tag{6.10.5}$$

for every $p \geq p_c$. (Again, it may seem that these constants should also depend on p_c , but p_c is bounded below the reciprocal of the degree and bounded away from 1 by a quantity depending only on ∇_{p_c} .) Russo's formula yields in our setting that

$$\frac{d}{dp} \theta(p) = \frac{1}{1 - p} \mathbb{E}_p [\Psi_p(K_o) \mathbb{1}(|K_o| < \infty)] \leq \frac{\theta(p)d}{1 - p} \mathbb{E}_p [|K_o| \mathbb{1}(|K_o| < \infty)]$$

for every $p > p_c$, where Ψ_p is defined in (6.10.4). (Note that Russo only applies directly for events depending on at most finitely many edges, but in our setting we may easily take a limit since, by Theorem 6.10.3, the contribution from distant edges is very small. The argument needed to do this is standard and is omitted.) The claim follows from this together with (6.10.5) and Theorem 6.10.3, which implies that $\mathbb{E}_p [|K_o| \mathbb{1}(|K_o| < \infty)] = \sum_{n=1}^{\infty} \mathbb{P}_p(n \leq |K_o| < \infty) \leq C|p - p_c|^{-1}$ for some constant C depending only on the degree. \square

6.11 Proofs of the main theorems

To streamline our proofs, let us introduce some language that will only be used in this subsection. Let \mathcal{H} be a countably infinite set of (finite or infinite) transitive graphs. Say that an assignment $p : \mathcal{H} \rightarrow [0, 1]$ is *supercritical* if there is a constant $\varepsilon > 0$ such that for $\theta((1 - \varepsilon)p, G) \geq \varepsilon$ for all but finitely many $G \in \mathcal{H}$, where $\theta(p, G)$ is as defined in eq. (6.4.1). We say that \mathcal{H} has a *discrete percolation phase transition* if $\lim_{\mathcal{H}} \theta(p, G) = 1$ for every supercritical assignment p . This generalises the definition of discrete percolation phase transition for families of finite transitive graphs given in Section 6.6. Recall the definition of *irreducibly macromolecular* from Section 6.9. Let us moreover say that $\mathcal{H} \subset \mathcal{F}$ is *irreducibly molecular* if \mathcal{H} is dense and irreducibly macromolecular.

Now say that \mathcal{H} is of **type**...

- 1 if every graph in \mathcal{H} is infinite and nonunimodular, and the vertex degrees in \mathcal{H} are bounded;
- 2 if every graph in \mathcal{H} is infinite and unimodular, and the vertex degrees in \mathcal{H} are bounded;
- 3 if every graph in \mathcal{H} is finite, and the vertex degrees in \mathcal{H} are bounded;
- 4 if every graph in \mathcal{H} is finite, $\lim_{\mathcal{H}} \deg G = \infty$, and \mathcal{H} does not contain an infinite subset that is either macromolecular or has a discrete percolation phase transition;
- 5 if every graph in \mathcal{H} is finite, $\lim_{\mathcal{H}} \deg G = \infty$, \mathcal{H} is sparse and irreducibly macromolecular, and \mathcal{H} does not contain an infinite subset that has a discrete percolation phase transition;
- 6 if every graph in \mathcal{H} is finite, \mathcal{H} does not have a discrete percolation phase transition, and \mathcal{H} is irreducibly molecular;

...and say that \mathcal{H} satisfies **property**...

T if for every supercritical p ,

$$\lim_{n \rightarrow \infty} \limsup_{\mathcal{H}} |\theta(p, G) - \mathbb{P}_p (|K_o| \geq n)| = 0;$$

L if for every supercritical p and for every infinite subset $\mathcal{H}' \subseteq \mathcal{H}$:

- if \mathcal{H}' converges locally to some infinite transitive graph G_∞ , then

$$\lim_{\mathcal{H}'} |\theta(p, G) - \theta(p, G_\infty)| = 0;$$

- if $\lim_{\mathcal{H}'} \deg G = \infty$, then

$$\lim_{\mathcal{H}'} |\theta(p, G) - \text{mf}(p \deg G)| = 0;$$

C if for every supercritical p and for every infinite subset $\mathcal{H}' \subseteq \mathcal{H} \cap \mathcal{F}$,

$$\lim_{\mathcal{H}} \mathbb{P}_p (|\|\text{giant}\| - \theta(p, G)| > \varepsilon) = 0;$$

E if for every supercritical p ,

$$\lim_{\delta \downarrow 0} \sup_{\mathcal{H}} \sup_{q \in [p, 1]} |\theta(q, G) - \theta((1 - \delta)q, G)| = 0.$$

We will establish the validity of the following table. Formally, this means that if \mathcal{H} has some type i , then for every property P , if the (P, i) th entry of the following table is a tick (\checkmark), then \mathcal{H} has property P , whereas if the entry is a cross (\times), then \mathcal{H} does not have property P .

	1	2	3	4	5	6
T	\checkmark	\checkmark	\checkmark	\checkmark	\times	\times
L	\checkmark	\checkmark	\checkmark	\checkmark	\times	\times
E	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
C	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\times

Starting the table

In this subsection, we continue to let \mathcal{H} be a countably infinite set of (finite or infinite) transitive graphs.

Lemma 6.11.1. *The following portion of our table is correct.*

	1	2	3	4	5	6
T	\checkmark			\checkmark		
L		\checkmark	\checkmark			
E					\checkmark	
C						

Proof.

(T,1) This follows immediately from Theorem 6.10.3.

(L,2) and (L,3) Suppose that \mathcal{H} is of type 2 or 3, and consider an infinite subset $\mathcal{H}' \subseteq \mathcal{H}$ that converges locally to some infinite transitive graph G_∞ . By Proposition 6.4.1,

$$\lim_{\mathcal{H}} \theta(p, G) = \theta(p, G_\infty)$$

for every constant $p > p_* := \limsup_{\mathcal{H}} p_c(G)$. In particular, $\theta(p, G_\infty) > 0$ for all $p > p_*$. So $(p_*, 1] \subseteq (p_c(G_\infty), 1]$, and thus by [Sch99], $\theta(\cdot, G_\infty)$ is continuous on $(p_*, 1]$. Now consider an arbitrary (possibly non-constant) supercritical assignment p . Since p is supercritical, there exists a constant $\varepsilon > 0$ such that $p \in [p_* + \varepsilon, 1]$ for all but finitely many $G \in \mathcal{H}$. Recall that if a sequence of monotone functions defined on a closed and bounded interval converges pointwise to a continuous function, then the sequence actually converges uniformly. This readily implies that as $G \rightarrow \infty$

with $G \in \mathcal{H}$, the density $\theta(\cdot, G)$ converges uniformly to $\theta(\cdot, G_\infty)$ on the interval $[p_* + \varepsilon, 1]$. In particular,

$$\lim_{\mathcal{H}} |\theta(p, G) - \theta(p, G_\infty)| = 0.$$

(T,4) This follows immediately from Proposition 6.8.1 and Theorems 6.1.1 and 6.2.1.

(E,5) Suppose that \mathcal{H} has type 5. Let $\mathcal{H}' \subseteq \mathcal{H}$ be an infinite subset, and let $q_1, q_2 : \mathcal{H}' \rightarrow (0, 1)$ be a pair of assignments satisfying $\lim_{\mathcal{H}'} q_2/q_1 = 1$ and $q_1, q_2 \in [p, 1]$ for all $G \in \mathcal{H}'$. Since $\theta(\cdot, G)$ is trivially continuous for each $G \in \mathcal{H}$, to establish E, it suffices to show that

$$\lim_{\mathcal{H}'} \underbrace{|\theta(q_2, G) - \theta(q_1, G)|}_{\clubsuit} = 0.$$

By Lemma 6.2.2 and Proposition 6.6.1, this holds if $\lim_{\mathcal{H}'} q_2 \deg G \in \{0, \infty\}$. So it suffices for us to hanfdle the case when there is a constant $C < \infty$ such that $q_2 \deg G \in [1/C, C]$ for all $G \in \mathcal{H}'$. Fix an irreducibly macromolecular decomposition (A, B) for \mathcal{H}' , and similarly to as in the statement of Proposition 6.9.1, for each $G \in \mathcal{H}'$ and for both $i \in \{1, 2\}$, let

$$\psi_i(G) := \text{mf}(q_i(G) \deg G) \quad \text{and} \quad q_i^*(G) := q_i(G)\psi_i(G)^2.$$

For all $G \in \mathcal{H}'$ and $n \geq 1$,

$$\clubsuit \leq \heartsuit_n^1 + \heartsuit_n^2 + \diamondsuit_n,$$

where for both $i \in \{1, 2\}$,

$$\heartsuit_n^i := |\theta(q_i, G) - \psi_i \cdot \mathbb{P}_{q_i}^B (|K_o| \geq n)|,$$

and

$$\diamondsuit_n := |\psi_2 \cdot \mathbb{P}_{q_2}^B (|K_o| \geq n) - \psi_1 \cdot \mathbb{P}_{q_1}^B (|K_o| \geq n)|.$$

By Proposition 6.9.1, $\lim_{\mathcal{H}'} \heartsuit_n^i = 0$ for both $i \in \{1, 2\}$. So it suffices to prove that $\lim_{\mathcal{H}'} \diamondsuit_n = 0$.

Using that ψ and q_2 are trivially bounded above by 1, for all G and n ,

$$\diamondsuit_n \leq \underbrace{|\mathbb{P}_{q_2}^B (|K_o| \geq n) - \mathbb{P}_{q_1}^B (|K_o| \geq n)|}_{\diamondsuit_n^1} + \underbrace{|\psi_2 - \psi_1|}_{\diamondsuit_n^2}.$$

There can never be more than $n \deg G$ closed pivotal edges for the event $\{|K_o| \geq n\}$. So by Russo's formula, for all G and n , the derivative of the function

$$f_{G,n}(p) := \mathbb{P}_p^B (|K_o| \geq n)$$

satisfies

$$f'_{G,n}(p) \leq \frac{n \deg G}{1-p}$$

for all $p \in (0, 1)$. In particular, by the mean-value theorem, for all G and n ,

$$\diamond_n^1 \leq (q_2 - q_1) \cdot \frac{n \deg G}{1 - C/\deg G}. \quad (6.11.1)$$

Since $\lim_{\mathcal{H}} \deg G = \infty$, we have $\frac{1}{1-C/\deg G} \leq 2$ for all but finitely many $G \in \mathcal{H}'$. By plugging this and $q_1 \leq \frac{C}{\deg G}$ into eq. (6.11.1), we obtain

$$\diamond_n^1 \leq \left(\frac{q_2}{q_1} - 1 \right) \cdot 2Cn,$$

and hence $\lim_{\mathcal{H}'} \diamond_n^1 = 0$ for all n . Note that the function $f : [-\infty, +\infty] \rightarrow [0, 1]$ given by $f(x) := \text{mf}(e^x)$ for all $x \in \mathbb{R}$, $f(-\infty) := 0$, and $f(+\infty) := 1$, is uniformly continuous. Since

$$\lim_{\mathcal{H}'} |\log(q_2 \deg G) - \log(q_1 \deg G)| = \lim_{\mathcal{H}'} \left| \log \left(\frac{q_2}{q_1} \right) \right| = 0,$$

it follows that $\lim_{\mathcal{H}'} \diamond^2 = 0$. By combining this with our control of \diamond_n^1 , we have $\lim_{\mathcal{H}'} \diamond_n = 0$, as required. \square

Lemma 6.11.2. *The following portion of our table is correct.*

	1	2	3	4	5	6
T					\times	\times
L						
E						
C	✓	✓				\times

Proof.

(C,1) and (C,2) These hold vacuously because if \mathcal{H} is of type 1 or 2, then none of the graphs in \mathcal{H} are finite.

(T,5) Suppose that \mathcal{H} is of type 5. Since \mathcal{H} does not have a discrete percolation phase transition, by Corollary 6.6.2, there exists a supercritical assignment p satisfying

$$\liminf_{\mathcal{H}} p \deg G < \infty.$$

It now follows easily from Claim 6.9.2 and the definition of macromolecular decompositions that there exists a constant $\varepsilon > 0$ such that

$$\limsup_{\mathcal{H}} \mathbb{P}_p \left(\varepsilon \deg G \leq |K_o| < \frac{1}{\varepsilon} \deg G \right) \geq \varepsilon. \quad (6.11.2)$$

(With probability bounded away from zero, for infinitely many choices of G , every edge in the boundary of the equivalence class of o is closed, but the cluster of o contains a positive proportion of the vertices in the equivalence class of o .)

For all $n \geq 1$ and $G \in \mathcal{H}$, by transitivity,

$$\begin{aligned} \heartsuit_{n,G} &:= \theta(p, G) - \mathbb{P}_p (|K_o| \geq n) \\ &= \frac{1}{|V|} \sum_{v \in V} (\mathbb{P}_p (v \in \text{giant}) - \mathbb{P}_p (|K_v| \geq n)) \\ &\geq \frac{1}{|V|} \sum_v \left(\mathbb{P}_p \left(\varepsilon \deg G \leq |K_v| < \frac{1}{\varepsilon} \deg G \right) - \mathbb{P}_p \left(|\text{giant}| < \frac{1}{\varepsilon} \deg G \right) \right). \end{aligned}$$

So by eq. (6.11.2), Theorem 6.2.1, and the fact that \mathcal{H} is sparse, for all $n \geq 1$,

$$\limsup_{\mathcal{H}} \heartsuit_{n,G} \geq \varepsilon.$$

Since ε is independent of n , this shows that \mathcal{H} does not have property T.

(T,6) Suppose that \mathcal{H} is of type 6. By Theorem 6.1.1, \mathcal{H} does not have the supercritical uniqueness property. So there exists a supercritical assignment p and a constant $\varepsilon > 0$ such that

$$\limsup_{\mathcal{H}} \frac{1}{|V|} \sum_{v \in V} \mathbb{P}_p (|K_v| \geq \varepsilon |G| \text{ but } v \notin \text{giant}) \geq \varepsilon.$$

By arguing along the same lines as in case (T,5), it follows that for all $n \geq 1$,

$$\limsup_{\mathcal{H}} \heartsuit_{n,G} \geq \varepsilon,$$

and hence that \mathcal{H} does not have property T.

(C,6) This follows immediately from [EH21a, Lemma 3.6]. \square

Finishing the table

In this subsection, we continue to let \mathcal{H} be a countably infinite set of (finite or infinite) transitive graphs. We will extend the portion of our table completed in the last subsection to complete the whole table. We will use “ (P, i) ” as shorthand for the statement “if \mathcal{H} is of type i , then \mathcal{H} has property P ”.

Recall that \mathcal{F}_d is the subset of \mathcal{F} of graphs with vertex degree exactly d , and recall that $\text{mf}(\lambda)$ is the probability that a Poisson(λ)-branching process survives forever. The following result is easy to prove and known to experts. It is classical for percolation on complete graphs and hypercubes, and the standard arguments extend without change to the general case. There are many elementary ways to prove this result, so we will omit the details. (The arguments even work more generally for high-degree regular graphs that are not necessarily transitive.)

Lemma 6.11.3. *For every positive constant C ,*

$$\lim_{n \rightarrow \infty} \limsup_{d \rightarrow \infty} \sup_{G \in \mathcal{F}_d} \sup_{p \in [0, C/d]} \left| \mathbb{P}_p^G (|K_o| \geq n) - \text{mf}(pd) \right| = 0.$$

Proof sketch. Let \mathbf{P}_λ denote the law of a Poisson(λ)-branching process, and let N be total number of offspring. Fix $C > 0$. We have (e.g. by Dini’s theorem) that

$$\mathbf{P}_\lambda(N \geq k) \rightarrow \text{mf}(\lambda) \quad \text{as } k \rightarrow \infty,$$

uniformly over all $\lambda \in [0, C]$. Moreover, for every fixed k , we have (e.g. by comparing an exploration of K_o to a Binomial branching process, by coupling to percolation on a regular tree, or by counting the number of trees of a given size) that

$$\left| \mathbb{P}_{\lambda/\deg G} (|K_o| \leq k) - \mathbf{P}_\lambda (|K_o| \leq k) \right| \rightarrow 0 \quad \text{as } G \rightarrow \infty \text{ with } G \in \mathcal{F},$$

uniformly over all $\lambda \in [0, C]$. So given any $\eta > 0$, if we first pick n sufficiently large, then pick d sufficiently large (depending on n), we can ensure that

$$\left| \mathbb{P}_{\lambda/d}^G (|K_o| \geq n) - \text{mf}(\lambda) \right| \leq \eta$$

for all $G \in \mathcal{F}_d$ and all $\lambda \in [0, C]$. □

Lemma 6.11.4. *For every $i \in \{1, \dots, 6\}$,*

$$(T, i) \iff (L, i).$$

Proof.

$(T, i) \implies (L, i)$ Let $\mathcal{H}' \subseteq \mathcal{H}$ be an arbitrary infinite subset that either (1) converges locally to some infinite transitive graph G_∞ , or (2) satisfies $\lim_{\mathcal{H}'} \deg G = \infty$.

1. For all $G \in \mathcal{H}$ and $n \geq 1$,

$$\begin{aligned} \heartsuit &:= |\theta(p, G) - \theta(p, G_\infty)| \\ &\leq \underbrace{|\theta(p, G) - \mathbb{P}_p^G (|K_o| \geq n)|}_{\heartsuit_n^1} + \underbrace{|\mathbb{P}_p^G (|K_o| \geq n) - \mathbb{P}_p^{G_\infty} (|K_o| \geq n)|}_{\heartsuit_n^2} + \underbrace{|\mathbb{P}_p^{G_\infty} (|K_o| \geq n) - \theta(p, G_\infty)|}_{\heartsuit_n^3}. \end{aligned}$$

Since p is supercritical, for all n ,

$$\lim_{\mathcal{H}} \mathbb{P}_p (|\text{giant}| \geq n) = 1,$$

and hence

$$\limsup_{\mathcal{H}} \heartsuit_n^1 = \limsup_{\mathcal{H}} \mathbb{P}_p^G (|K_o| \geq n \text{ but } o \notin \text{giant}), \quad (6.11.3)$$

which tends to zero as n tends to infinity if **T** holds. In particular, if **T** holds then

$$\begin{aligned} \limsup_{\mathcal{H}'} \heartsuit &\leq \limsup_{\mathcal{H}'} \inf_{n \geq 1} (\heartsuit_n^1 + \heartsuit_n^2 + \heartsuit_n^3) \\ &\leq \limsup_{n \rightarrow \infty} \limsup_{\mathcal{H}'} (\heartsuit_n^1 + \heartsuit_n^2 + \heartsuit_n^3) \\ &\leq \underbrace{\limsup_{n \rightarrow \infty} \limsup_{\mathcal{H}'} \heartsuit_n^1}_{=0 \text{ by } \mathbf{T}} + \underbrace{\sup_{n \geq 1} \limsup_{\mathcal{H}'} \heartsuit_n^2}_{=0} + \underbrace{\limsup_{n \rightarrow \infty} \heartsuit_n^3}_{=0} = 0. \end{aligned}$$

2. Thanks to Corollary 6.6.2, the claim is trivial if

$$\lim_{\mathcal{H}'} p \deg G = \infty.$$

So by passing to a further infinite subset if necessary, let us assume without loss of generality that

$$\sup_{\mathcal{H}'} p \deg G < \infty,$$

which will allow us to invoke Lemma 6.11.3. For all $G \in \mathcal{H}$ and $n \geq 1$,

$$\begin{aligned} \diamond &:= |\theta(p, G) - \text{mf}(p \deg G)| \\ &\leq \heartsuit_n^1 + \underbrace{|\mathbb{P}_p^G (|K_o| \geq n) - \text{mf}(p \deg G)|}_{\diamond_n^2}. \end{aligned}$$

In particular, if **T** holds then

$$\begin{aligned}
\limsup_{\mathcal{H}'} \diamond &\leq \limsup_{\mathcal{H}'} \inf_{n \geq 1} \left(\diamond_n^1 + \diamond_n^2 \right) \\
&\leq \limsup_{n \rightarrow \infty} \limsup_{\mathcal{H}'} \left(\diamond_n^1 + \diamond_n^2 \right) \\
&\leq \underbrace{\limsup_{n \rightarrow \infty} \limsup_{\mathcal{H}'} \diamond_n^1}_{=0 \text{ by } \mathbf{T}} + \underbrace{\limsup_{n \rightarrow \infty} \limsup_{\mathcal{H}'} \diamond_n^2}_{=0 \text{ by Lemma 6.11.3}} = 0.
\end{aligned}$$

$(T, i) \iff (L, i)$ Thanks to eq. (6.11.3), our goal is to show that

$$\lim_{n \rightarrow \infty} \lim_{\mathcal{H}} \diamond_n^1 = 0,$$

or equivalently, to show that this holds for every infinite subset $\mathcal{H}' \subseteq \mathcal{H}$ that either (1) converges locally to some infinite transitive graph G_∞ , or (2) satisfies $\lim_{\mathcal{H}} \deg G = \infty$.

1. By a similar argument as above, for all $G \in \mathcal{H}$ and $n \geq 1$,

$$\diamond_n^1 \leq \diamond + \diamond_n^2 + \diamond_n^3.$$

So if **L** holds then

$$\limsup_{n \rightarrow \infty} \limsup_{\mathcal{H}'} \diamond_n^1 \leq \underbrace{\limsup_{\mathcal{H}'} \diamond}_{=0 \text{ by } \mathbf{L}} + \underbrace{\sup_{n \geq 1} \limsup_{\mathcal{H}'} \diamond_n^2}_{=0} + \underbrace{\limsup_{n \rightarrow \infty} \limsup_{\mathcal{H}'} \diamond_n^3}_{=0} = 0.$$

2. By a similar argument as above, for all $G \in \mathcal{H}$ and $n \geq 1$,

$$\diamond_n^1 \leq \diamond + \diamond_n^2.$$

So if **L** holds then

$$\limsup_{n \rightarrow \infty} \limsup_{\mathcal{H}'} \diamond_n \leq \underbrace{\limsup_{\mathcal{H}'} \diamond}_{=0 \text{ by } \mathbf{L}} + \underbrace{\limsup_{n \rightarrow \infty} \limsup_{\mathcal{H}'} \diamond_n^2}_{=0 \text{ by Lemma 6.11.3}} = 0.$$

□

Lemma 6.11.5. *For all $i \in \{1, \dots, 6\}$,*

$$(L, i) \implies (E, i).$$

Proof. Let $\mathcal{H}' \subseteq \mathcal{H}$ be an infinite subset, and let $q_1, q_2 : \mathcal{H}' \rightarrow (0, 1)$ be a pair of assignments satisfying $\lim_{\mathcal{H}'} q_2/q_1 = 1$ and $q_1, q_2 \in [p, 1]$ for all $G \in \mathcal{H}'$. To establish **E**, it suffices to show that

$$\lim_{\mathcal{H}'} \underbrace{|\theta(q_2, G) - \theta((q_1, G)|}_{\heartsuit} = 0.$$

Without loss of generality, we may further assume that either (1) \mathcal{H}' converges locally to some infinite transitive graph G_∞ , or (2) $\lim_{\mathcal{H}'} \deg G = \infty$.

1. For every $G \in \mathcal{H}'$,

$$\heartsuit \leq |\theta(q_2, G) - \theta(q_2, G_\infty)| + \underbrace{|\theta(q_2, G_\infty) - \theta(q_1, G_\infty)|}_{\heartsuit_1} + |\theta(q_1, G_\infty) - \theta(q_1, G)|.$$

So if **L** holds then

$$\limsup_{\mathcal{H}'} \heartsuit \leq \limsup_{\mathcal{H}'} \heartsuit_1,$$

and since q_1 is supercritical,

$$b := \liminf_{\mathcal{H}'} q_1 > p_c(G_\infty).$$

So by [Sch99], the function $\theta(\cdot, G_\infty)$ is uniformly continuous on $[b, 1]$. Since

$$\limsup_{\mathcal{H}'} |q_2 - q_1| \leq \limsup_{\mathcal{H}'} \left| \frac{q_2}{q_1} - 1 \right| = 0,$$

it follows that $\lim_{\mathcal{H}'} \heartsuit_1 = 0$, as required.

2. Arguing as in the case above, if **L** holds then

$$\limsup_{\mathcal{H}'} \heartsuit \leq \limsup_{\mathcal{H}'} \underbrace{|\text{mf}(q_2 \deg G) - \text{mf}(q_1 \deg G)|}_{\heartsuit_2}.$$

Note that the function $f : [-\infty, +\infty] \rightarrow [0, 1]$ given by $f(x) := \text{mf}(e^x)$ for all $x \in \mathbb{R}$, $f(-\infty) := 0$, and $f(+\infty) := 1$, is uniformly continuous. Since

$$\lim_{\mathcal{H}'} |\log(q_2 \deg G) - \log(q_1 \deg G)| = \lim_{\mathcal{H}'} \left| \log \left(\frac{q_2}{q_1} \right) \right| = 0,$$

it follows that $\limsup_{\mathcal{H}'} \heartsuit_2 = 0$, as required.

□

Lemma 6.11.6. *For all $i \in \{3, 4, 5\}$,*

$$(E, i) \implies (C, i).$$

Proof. Assume that \mathcal{H} is of type 3,4, or 5. So every graph in \mathcal{H} is finite, and \mathcal{H} does not contain an infinite subset that is molecular. Fix a constant $\varepsilon > 0$. If

$$\limsup_{\mathcal{H}} \mathbb{P}_p (\|\text{giant}\| \geq \theta(p, G) + \varepsilon) > 0, \quad (6.11.4)$$

then by Theorem 6.2.1, there is a sequence $\delta : \mathcal{H} \rightarrow [0, 1]$ with $\lim_{\mathcal{H}} \delta = 0$ such that

$$\limsup_{\mathcal{H}} \mathbb{P}_{(1+\delta)} (\|\text{giant}\| \geq \theta(p, G) + \varepsilon) = 1,$$

and hence

$$\limsup_{\mathcal{H}} \theta((1 + \delta)p, G) \geq \theta(p, G) + \varepsilon,$$

which implies that E does not hold. The same reasoning shows that E does not hold if

$$\limsup_{\mathcal{H}} \mathbb{P}_p (\|K_o\| \leq \theta(p, G) - \varepsilon) > 0. \quad (6.11.5)$$

If \mathcal{H} does not have property C, then eq. (6.11.4) or eq. (6.11.5) must hold for some choice of $\varepsilon > 0$. So, putting everything together, we have shown that if \mathcal{H} does not have property C, then \mathcal{H} does not have property E either. □

Lemma 6.11.7. $(C, 5) \wedge (L, 5) \implies (E, 6)$.

Proof. Suppose that $(C, 5)$ and $(L, 5)$ hold. Suppose that \mathcal{H} is of type 6, and let (A, B) be irreducibly macromolecular decomposition for \mathcal{H} . For each $G \in \mathcal{H}$, let \sim be an equivalence class inducing the macromolecular decomposition (A, B) . Now the fact that \mathcal{H} has property E follows easily from the fact that the subgraphs induced by equivalence classes of \sim all have properties E and C. Although this can be shown directly, for completeness, let us note that this follows for example from Claim 6.9.5. (Although this claim is proven in the setting that \mathcal{H} is sparse and irreducibly macromolecular, the hypothesis that \mathcal{H} is sparse is not used for this claim.) □

Proofs of the main theorems

By applying the results of Section 6.11, we can extend the entries we found in Section 6.11 to establish that all 4×6 entries of our table are correct. To deduce our main theorems from this table, the only non-trivial fact to check is that every infinite set of finite transitive graphs contains an infinite subset that is of some type $i \in \{3, 4, 5, 6\}$. Even more precisely, it suffices to show that if \mathcal{H} is macromolecular then \mathcal{H} has an infinite subset that is irreducibly macromolecular. This is the content of the following lemma.

Lemma 6.11.8. *Every infinite set of finite transitive graphs that is macromolecular contains an infinite subset that is irreducibly macromolecular.*

Proof. Let \mathcal{H} be an infinite set of finite transitive graphs that is macromolecular. Let $\delta_* \in [0, 1]$ be the largest constant such that for all $\delta \in [0, \delta_*]$ there exists a constant $\varepsilon > 0$ such that infinitely many graphs in \mathcal{H} admit a macromolecular decomposition (A, B) where A is δ -dense and

$$\frac{|E(B)|}{|G|} \leq \frac{1}{\varepsilon}.$$

Since \mathcal{H} is itself macromolecular, we know that $\delta_* > 0$. Now pick a constant $\varepsilon > 0$, an infinite set $\mathcal{I} \subseteq \mathcal{H}$, and a collection $((A(G), B(G)) : G \in \mathcal{I})$ such that $A(G)$ is $\frac{9}{10}\delta_*$ -dense and $\frac{|E(B)|}{|G|} \leq \frac{1}{\varepsilon}$ for all $G \in \mathcal{I}$. It suffices to show that $(A(G) : G \in \mathcal{I})$ does not contain an infinite subset that is macromolecular.

Suppose for contradiction that there exists an infinite subset $\mathcal{J} \subseteq \mathcal{I}$ and a constant $\eta > 0$ such that for each $G \in \mathcal{J}$, we can find an η -macromolecular decomposition $(A'(G), B'(G))$ for $A(G)$. Note that every $G \in \mathcal{J}$ now trivially admits a macromolecular decomposition (A', B'') where

$$\frac{|E(B'')|}{|G|} = \frac{|E(B)|}{|G|} + \frac{|E(B')|}{|A|} \leq \frac{2}{\varepsilon} + \frac{2}{\eta}.$$

We will show that for infinitely many choices of $G \in \mathcal{I}$, the graph A' is $\frac{4}{3}\delta_*$ -dense, contradicting the maximality of δ_* . Note that for all $G \in \mathcal{J}$,

$$\deg A = \deg G - \frac{2|E(B)|}{|G|} \geq \deg G - \frac{2}{\varepsilon},$$

and similarly,

$$\deg A' = \deg A - \frac{2|E(B')|}{|V(A)|} \geq \deg A - \frac{2}{\eta}.$$

In particular, note that $\lim_{G \in \mathcal{J}} |A'| = \infty$ because $|A'| \geq \deg A'$ for all G , and $\lim_{G \in \mathcal{H}} \deg G = \infty$. So without loss of generality, let us assume that every $G \in \mathcal{I}$ satisfies

$$|A'| \geq \frac{16}{\varepsilon\delta_*} + \frac{16}{\eta\delta_*}.$$

For all $G \in \mathcal{I}$, since $|A'|$ must be a non-trivial divisor of $|A|$, we have $|A'| \leq \frac{1}{2}|A|$. Therefore,

$$\frac{\deg A'}{|A'|} \geq \frac{\deg G - \frac{2}{\varepsilon} - \frac{2}{\eta}}{\frac{1}{2}|A|} \geq 2 \frac{\deg A}{|A|} - \frac{\delta_*}{4}.$$

In particular, since A is $\frac{9}{10}\delta_*$ -dense, A' must be $2 \cdot \frac{9}{10}\delta_* - \frac{\delta_*}{2} = \frac{4}{3}\delta_*$ -dense, as required. \square

Now Theorem 6.1.3 follows immediately from Theorem 6.1.1 and row C of our table. Theorem 6.1.5 follows from $(L, 1)$ and $(L, 2)$, since the set of non-unimodular infinite transitive graphs is a closed and open subset of the set of all infinite transitive graphs. Theorem 6.1.7 follows from $(L, 3)$ and (for the converse) $(L, 4)$ and $(L, 5)$. Theorem 6.1.11 follows from row E .

EXISTENCE OF A PERCOLATION THRESHOLD ON FINITE TRANSITIVE GRAPHS

Abstract

Let (G_n) be a sequence of finite connected vertex-transitive graphs with volume tending to infinity. We say that a sequence of parameters (p_n) is a *percolation threshold* if for every $\varepsilon > 0$, the proportion $\|K_1\|$ of vertices contained in the largest cluster under bond percolation \mathbb{P}_p^G satisfies both

$$\lim_{n \rightarrow \infty} \mathbb{P}_{(1+\varepsilon)p_n}^{G_n} (\|K_1\| \geq \alpha) = 1 \quad \text{for some } \alpha > 0, \text{ and}$$

$$\lim_{n \rightarrow \infty} \mathbb{P}_{(1-\varepsilon)p_n}^{G_n} (\|K_1\| \geq \alpha) = 0 \quad \text{for all } \alpha > 0.$$

We prove that (G_n) has a percolation threshold if and only if (G_n) does not contain a particular infinite collection of pathological subsequences of dense graphs. Our argument uses an adaptation of Vanneuville's new proof of the sharpness of the phase transition for infinite graphs via couplings [Van22b] together with our recent work with Hutchcroft on the uniqueness of the giant cluster [EH21b].

7.1 Introduction

Given a graph G , build a random spanning subgraph ω by independently including each edge with a fixed probability p . This model \mathbb{P}_p^G is called *(Bernoulli bond) percolation*. In their pioneering work on random graphs, Erdős and Rényi [ER60] proved that when G is the complete graph on n vertices, percolation has a phase transition: as we increase p from $\frac{1-\varepsilon}{n}$ to $\frac{1+\varepsilon}{n}$ for any fixed $\varepsilon > 0$, a giant cluster suddenly emerges containing a positive proportion of the total vertices. Since then, there has been much interest in establishing this phenomenon for more general classes of finite graphs. However, as remarked in [Bol+10c], progress has been slow. In this paper, we solve this problem for arbitrary finite graphs that are *(vertex-)transitive*, meaning that for any two vertices u and v , there is a graph automorphism mapping u to v . This includes all Cayley graphs, and in particular, the complete graphs, hypercubes, and tori.

Our setting of percolation on finite transitive graphs places us at the intersection of two well-established fields. Loosely speaking, one of these began in combinatorics with the work of Erdős and Rényi [ER59; ER60], whereas the other began in mathematical physics with the work of Broadbent and Hammersley [BH57a]. In the former, a subgraph of a finite graph is said to

percolate if its largest cluster contains a positive proportion of the total vertices, whereas in the latter, a subgraph of an infinite transitive graph is said to *percolate* if its largest cluster is infinite. This leads to two different definitions of what it means to *have a percolation phase transition*. Let us start by making these precise.

Here are the graph-theoretic conventions we will be using throughout: The *volume* of a graph $G = (V, E)$ is simply the number of vertices $|V|$. We label the clusters (i.e. connected components) of a spanning subgraph of G in decreasing order of volume by K_1, K_2, \dots . In a slight abuse of notation, we also write K_v for the cluster containing a given vertex v . The *density* of a cluster K is $\|K\| := \frac{|K|}{|V|}$, the proportion of vertices contained in K .

Now let (G_n) be a sequence of finite graphs with volume tending to infinity. Following Bollobás, Borgs, Chayes, and Riordan [Bol+10c], we say that (G_n) has a percolation phase transition if there is a sequence of parameters (p_n) such that for every $\varepsilon > 0$, both¹

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathbb{P}_{(1+\varepsilon)p_n}^{G_n} (\|K_1\| \geq \alpha) &= 1 && \text{for some } \alpha > 0, \text{ and} \\ \lim_{n \rightarrow \infty} \mathbb{P}_{(1-\varepsilon)p_n}^{G_n} (\|K_1\| \geq \alpha) &= 0 && \text{for all } \alpha > 0. \end{aligned}$$

This is the subject of our paper: we characterise the existence of a percolation phase transition for finite transitive graphs. Let us mention that the question of a percolation phase transition on general finite graphs is attributed by the above authors of [Bol+10c] to Bollobás, Kohakayawa, and Łuksak [BKŁ92].

On the other hand, when G is an infinite (locally finite) transitive graph, we define the critical parameter

$$p_c := \sup \{p : \mathbb{P}_p^G(\text{there exists an infinite cluster}) = 0\}.$$

By Kolmogorov's zero-one law, the probability of an infinite cluster under \mathbb{P}_p^G is zero when $p < p_c$ and one when $p > p_c$. So in a trivial sense, G always has a percolation phase transition. The real question is whether $p_c < 1$. (The fact that $p_c > 0$ is obvious by a branching argument.) So in this context, we often say that G has a percolation phase transition to mean that $p_c < 1$. Hutchcroft and Tointon [HT21d] dealt with an analogue of this question for finite transitive graphs with bounded vertex degrees, i.e. the question of whether for a given sequence (G_n) of such graphs, there exists $\delta > 0$ such that $\mathbb{P}_{1-\delta}^{G_n}(\|K_1\| \geq \delta) \geq \delta$ for all n . This is not the subject of our paper. To avoid any possible confusion, when a sequence of finite graphs (G_n) with volume tending to infinity has a percolation phase transition in the above sense of [Bol+10c], we will instead say that it *has a percolation threshold*, referring to the threshold sequence of parameters (p_n) in the definition.

¹We will use the convention that $\mathbb{P}_p^G := \mathbb{P}_1^G$ if $p > 1$ and $\mathbb{P}_p^G := \mathbb{P}_0^G$ if $p < 0$.

The main result of this paper is Theorem 8.1.1 below, which characterises the existence of a percolation threshold on a sequence of finite transitive graphs in terms of the presence of an infinite collection of *molecular subsequences*. We discovered molecular sequences with Hutchcroft in [EH21b] as the only obstacles to the supercritical giant cluster being unique. Interestingly, unlike the usual story for percolation on a new family of graphs (as told in the introduction of [Bol+10c], for example), uniqueness of the supercritical giant cluster came *first*, before the existence of a percolation threshold, and the former is key to our proof of the latter. See Section 7.1 for more background. Here we will just recall the definition of a molecular sequence before stating Theorem 8.1.1.

Definition 7.1.1. Given an integer $m \geq 2$, we say that (G_n) is *m-molecular* if it is dense, meaning that $\liminf_{n \rightarrow \infty} \frac{|E(G_n)|}{|V(G_n)|^2} > 0$, and there is a constant $C < \infty$ such that for every n , there is a set of edges $F_n \subseteq E(G_n)$ satisfying the following conditions:

1. $G_n \setminus F_n$ has m connected components;
2. F_n is invariant under the action of $\text{Aut } G_n$;
3. $|F_n| \leq C |V(G_n)|$.

For example, the sequence of Cartesian products of complete graphs $(K_n \square K_m)_{n \geq 1}$ is *m-molecular*. We say that (G_n) is *molecular* if it is *m-molecular* for some $m \geq 2$.

Theorem 7.1.2. *A sequence of finite connected transitive graphs with volume tending to infinity has a percolation threshold if and only if it contains an m -molecular subsequence for at most finitely many integers m .*

The condition that a sequence (G_n) contains *m-molecular* subsequences for infinitely many integers m is extremely stringent. For example, we can rule it out if (G_n) is either sparse or dense, i.e. the edge density of G_n either tends to zero or remains bounded away from zero. Indeed, it is clear that a sparse sequence cannot contain any molecular subsequences, but also notice that since every *m-molecular* sequence (G_n) satisfies $\limsup_{n \rightarrow \infty} \frac{|E(G_n)|}{|V(G_n)|^2} \leq \frac{1}{m}$, a dense sequence can contain an *m-molecular* subsequence for at most finitely many integers m . In particular, our result implies the existence of a percolation threshold for the complete graphs, hypercubes, and tori. Our result is new even under the additional hypothesis that the graphs have uniformly bounded vertex degrees², which is particularly relevant to percolation on infinite graphs.

²One could imagine a family of sequences of such graphs that each has a percolation threshold but such that the $(1+\varepsilon)$ -supercritical giant cluster density is not bounded away from zero over the entire family. Then by diagonalising, we could construct a sequence without a percolation threshold.

Most previous work on percolation on finite graphs treated specific sequences such as the complete graphs, hypercubes, and tori. Indeed, many authors have remarked how little work has been done on more general classes of finite graphs [ABS04c; Bor+05a; Bol+10c; Ben13b]. Alon, Benjamini, and Stacey [ABS04c]³ studied percolation on expanders with bounded vertex degrees. In particular, they proved that if each graph is d -regular for some fixed integer d and has girth tending to infinity, then the sequence has a (constant) percolation threshold at $p_n := 1/(d-1)$. Borgs, Chayes, van der Hofstad, Slade, and Spencer [Bor+05a; Bor+05b] and Nachmias [Nac09] analysed the emergence of a cluster with volume of order $|V(G_n)|^{2/3}$ for percolation on finite transitive graphs satisfying the *triangle condition* or a random walk return-probability condition, respectively, both of which enforce mean-field behaviour. Frieze, Krivelevich, and Martin [FKM04] proved that if a sequence of finite regular graphs is *pseudorandom* (an eigenvalue condition forcing the graph to be like the complete graph), then it has a percolation threshold at $p_n := 1/\deg(G_n)$ where $\deg(G_n)$ is the vertex degree of the n th graph. Bollobás, Borgs, Chayes, and Riordan [Bol+10c] studied percolation on arbitrary sequences of finite graphs (G_n) that are *dense*, meaning that $\liminf_{n \rightarrow \infty} \frac{|E(G_n)|}{|V(G_n)|^2} > 0$. The following theorem from their paper will be used in our argument.

Theorem 7.1.3 (Bollobás, Borgs, Chayes, Riordan 2010). *Let (G_n) be a sequence of finite connected graphs with volume tending to infinity. Suppose that (G_n) is dense. For each n , let λ_n be the largest eigenvalue of the adjacency matrix of G_n . Then $(1/\lambda_n)_{n \geq 1}$ is a percolation threshold for (G_n) .*

Our proof of Theorem 8.1.1 does not build a percolation threshold by defining a natural candidate for the critical parameter of a finite graph in terms of, say, its vertex degrees or the largest eigenvalue of the adjacency matrix. In our setting, where we have no quantitative assumptions, we are forced to use a softer and more indirect approach. In particular, our proof of Theorem 8.1.1 says little about the rate at which the percolation probabilities tend to zero or to one. That said, as part of our proof we do obtain explicit lower bounds on the supercritical giant cluster density, which are analogous to the well-known mean-field lower bound for percolation on infinite graphs.

Corollary 7.1.4. *Let (G_n) be a sequence of finite connected transitive graphs with volume tending to infinity that does not contain an m -molecular subsequence for any $m > M$, where M is some positive integer. Let (p_n) be a percolation threshold, which exists by Theorem 8.1.1. Then for every $\varepsilon > 0$,*

$$\mathbb{P}_{(1+\varepsilon)p_n}^{G_n} \left(\|K_1\| \geq \frac{\varepsilon}{M(1+\varepsilon)} - o(1) \right) = 1 - o(1) \quad \text{as } n \rightarrow \infty.$$

³Here Alon, Benjamini, and Stacey also make several conjectures about percolation on finite transitive graphs, which may interest the reader.

Moreover, by simply bounding a percolation threshold (p_n) below by the threshold for being dominated by a subcritical branching process and above by the threshold for connectivity, we obtain the following optimal bounds on its location.

Proposition 7.1.5. *Let (G_n) be a sequence of finite connected transitive graphs with volume tending to infinity. Let d_n denote the vertex degree of G_n . If (G_n) has a percolation threshold (p_n) , then it satisfies*

$$(1 - o(1)) \frac{1}{d_n - 1} \leq p_n \leq (2 + o(1)) \frac{\log |V(G_n)|}{d_n} \quad \text{as } n \rightarrow \infty.$$

When (G_n) has a percolation threshold (p_n) and converges locally to an infinite graph G , one might ask how the location of (p_n) relates to the critical parameter p_c and the uniqueness threshold p_u of G . See Remark 1.6 in [EH21b] for a discussion of this question. Let us simply note that (p_n) typically (but not always) converges to p_c . For example, this is the case when (G_n) is a sequence of transitive expanders and G is nonamenable [BNP11b].

We conclude this discussion by explaining how our result relates to the general theory of sharp thresholds. Consider a large collection of independent random bits $b := (b_i)_{1 \leq i \leq n} \in \{0, 1\}^n$ each sampled according to the Bernoulli(p) distribution for some $p \in [0, 1]$. For many natural monotone events $A \subseteq \{0, 1\}^n$, the probability that b belongs to A has a sharp threshold: it increases from $o(1)$ to $1 - o(1)$ as p increases across an interval of width $o(p(1 - p))$. There is a general philosophy that a sharp threshold occurs if and only if A is sufficiently *symmetric/global*. For example, the event that the Erdős-Rényi random graph is connected has a sharp threshold, but the event that it contains a triangle and the event that a particular edge is present do not. Theorem 8.1.1 can be understood as a kind of extension of this principle to the existence of a phase transition in a statistical mechanics model. Indeed, Theorem 8.1.1 says that "having a giant cluster" typically has a threshold around a critical parameter p with a threshold width⁴ $o(p)$ whenever the underlying graph is *transitive*, which is a symmetry/homogeneity condition. However, we would like to emphasise that "having a giant cluster" is not a well-defined event, so it does not fall within the usual scope of sharp-threshold techniques.

Molecular sequences and the supercritical phase

Let (G_n) be a sequence of finite connected transitive graphs with volume tending to infinity. If (G_n) has a percolation threshold (p_n) , then a sequence of parameters (q_n) is called *supercritical* if

⁴The threshold width fails to be $o(1 - p)$ even for simple examples such as the sequence of cycles.

there exist $\varepsilon > 0$ and $N < \infty$ such that for all $n \geq N$ satisfying $q_n < 1$, we have

$$q_n \geq (1 + \varepsilon)p_n.$$

As in [HT21d; EH21b], we generalise this definition to the situation that there may or may not be a percolation threshold by saying that a sequence of parameters (q_n) is *supercritical* if there exist $\varepsilon > 0$ and $N < \infty$ such that for all $n \geq N$ satisfying $q_n < 1$, we have

$$\mathbb{P}_{(1-\varepsilon)q_n}(\|K_1\| \geq \varepsilon) \geq \varepsilon.$$

(The ‘ $q_n < 1$ ’ condition is a technicality that guarantees that (G_n) always admits at least one supercritical sequence of parameters, namely the sequence (p_n) with $p_n := 1$ for all n .)

In our work with Hutchcroft [EH21b], we showed that not having a molecular subsequence is the geometric counterpart to the supercritical giant cluster being unique. Here is the precise result from that paper.

Definition 7.1.6. We say that (G_n) has the *supercritical uniqueness property* if for every supercritical sequence of parameters (q_n) ,

$$\lim_{n \rightarrow \infty} \mathbb{P}_{q_n}^{G_n}(\|K_2\| \geq \alpha) = 0 \quad \text{for all } \alpha > 0.$$

Theorem 7.1.7 (Easo and Hutchcroft, 2021). *A sequence of finite connected transitive graphs with volume tending to infinity has the supercritical uniqueness property if and only if it does not contain a molecular subsequence.*

A crucial step in the proof of the above theorem is that if (G_n) does not contain a molecular subsequence, then it has the *sharp-density property*. To prove Theorem 8.1.1, we only need to recall that the sharp-density property guarantees that for every density $\alpha \in (0, 1]$, constant $\varepsilon > 0$, and supercritical sequence of parameters (p_n) ,

$$\liminf_{n \rightarrow \infty} \mathbb{P}_{p_n}^{G_n}(\|K_1\| \geq \alpha) > 0 \quad \text{implies} \quad \lim_{n \rightarrow \infty} \mathbb{P}_{(1+\varepsilon)p_n}^{G_n}(\|K_1\| \geq \alpha) = 1.$$

This implies that the event that “there exists a cluster with density at least α ” undergoes a sharp threshold for each fixed α . Notice that this does not immediately imply the existence of a percolation threshold. One obstruction could be how these thresholds are spaced: one could imagine that the α -density threshold always occurs at $\left(\frac{\alpha}{\sqrt{n}}\right)_{n \geq 1}$, say, in which case there would be no percolation threshold. The implication is not clear even if we additionally require that the G_n ’s have uniformly bounded vertex degrees: see our footnote on page 3, and see Conjecture 1.2 in [Ben+12] for an analogous situation in the context of expanders.

Proof strategy

Most of our work goes into showing that if (G_n) does not contain a molecular subsequence, then it has a percolation threshold. In this section, we outline this step. To extend this result to the case that (G_n) contains an m -molecular subsequence for at most finitely many integers m , we apply the same argument to a molecular sequence's constituent sequence of *atoms*. We will deduce that converse as an immediate consequence of a corollary from [EH21b].

Let (G_n) be a sequence of finite connected transitive graphs with volume tending to infinity. Suppose we want to prove that (G_n) has a percolation threshold. Although "having a giant cluster" is not an event, we could try taking the event $\{\|K_1\| \geq \delta_n\}$ as a proxy, where (δ_n) is a sequence tending to zero very slowly. Then as a candidate for the percolation threshold, we could take (p_n) where each p_n is defined to be the unique parameter satisfying

$$\mathbb{P}_{p_n}^{G_n} (\|K_1\| \geq \delta_n) = \frac{1}{2}.$$

As a sanity check, notice that if (G_n) does have a percolation threshold, then by a diagonal argument, there is a percolation threshold (p_n) that arises in this way.

To prove that our candidate (p_n) is in fact a percolation threshold, we need two ingredients. The first ingredient is that the emergence of a cluster of any *constant* density has a threshold about a critical parameter p with a threshold width $o(p)$. This immediately handles the subcritical half of our task: since $\lim_{n \rightarrow \infty} \delta_n = 0$, it guarantees that $\lim_{n \rightarrow \infty} \mathbb{P}_{(1-\varepsilon)p_n}^{G_n} (\|K_1\| \geq \alpha) = 0$ for every $\alpha, \varepsilon > 0$. As mentioned in Section 7.1, the existence of these constant-density thresholds is implied by the sharp-density property, which holds whenever (G_n) has no molecular subsequences.

The second ingredient is a universal lower bound on the supercritical giant cluster density. This says that for every $\varepsilon > 0$, there exists $\delta > 0$ such that for every sequence of parameters (p_n) , if $\liminf_{n \rightarrow \infty} \mathbb{E}_{p_n}^{G_n} \|K_1\| > 0$, then $\liminf_{n \rightarrow \infty} \mathbb{E}_{(1+\varepsilon)p_n}^{G_n} \|K_1\| \geq \delta$. Together with the first ingredient, this ensures that by taking (δ_n) to decay slowly enough, for every $\varepsilon > 0$, there exists $\delta > 0$ with $\lim_{n \rightarrow \infty} \mathbb{P}_{(1+\varepsilon)p_n}^{G_n} (\|K_1\| \geq \delta) = 1$. We will prove that this too holds whenever (G_n) has no molecular subsequences.

This second ingredient is reminiscent of the mean-field lower bound from the study of percolation on infinite graphs. This says, for example, that every vertex v in an infinite transitive graph G satisfies

$$\mathbb{P}_{(1+\varepsilon)p_c(G)}^G (K_v \text{ is infinite}) \geq \frac{\varepsilon}{1 + \varepsilon}.$$

Together with the exponential decay of $|K_v|$ throughout the subcritical phase, this forms what is known as the *sharpness of the phase transition* for percolation on infinite graphs. This foundational

result was first proved in [Men86; AB87a; CC87] but has recently been reproved by more modern arguments in [DT16c; DRT19; Hut20d; Van22b]. The main obstacle to adapting these proofs to our finite setting is that they concern the event that the cluster at a vertex reaches a certain distance or exceeds a certain volume, whereas we care about the cluster’s volume as a proportion of the total vertices.

Moreover, while the sharpness of the phase transition is completely general — applying to all infinite transitive graphs — there is no universal lower bound on the supercritical giant cluster density that applies to every finite transitive graph⁵. So to adapt one of these proofs to our setting, we need to include information about the underlying sequence of graphs in the argument itself, for example, that it has the supercritical uniqueness property, or equivalently, that it has no molecular subsequences.

Very recently, Vanneuville [Van22b] gave a new proof of the sharpness of the phase transition via couplings. Unlike previous arguments, this one does not rely on a differential inequality. Vanneuville’s key insight was that by using an exploration process, we can upper bound the effect of conditioning on a certain decreasing event, namely the event that a vertex’s cluster does not reach a certain distance, by the effect of slightly decreasing the percolation parameter. Our strategy is to apply this argument but with the event that a vertex’s cluster has small density.

Rather than building an exact monotone coupling of the conditioned percolation measure and the percolation measure with a smaller parameter, which is impossible, we will construct a coupling that is monotone outside of an error event. Then under the additional hypothesis that (G_n) has the supercritical uniqueness property, we will prove that this error event has probability tending to zero. Just as Vanneuville’s coupling immediately yields the mean-field lower bound, our *approximately monotone* coupling will tell us that for every sequence of parameters (p_n) and every constant $\varepsilon > 0$, if $\liminf_{n \rightarrow \infty} \mathbb{E}_{p_n}^{G_n} \|K_1\| > 0$, then

$$\liminf_{n \rightarrow \infty} \mathbb{E}_{(1+\varepsilon)p_n}^{G_n} \|K_1\| \geq \frac{\varepsilon}{1 + \varepsilon}.$$

Let us mention that for this step — establishing a universal lower bound on the supercritical giant cluster density when we have the supercritical uniqueness property — it is possible to instead adapt Hutchcroft’s proof of sharpness from [Hut20d], rather than Vanneuville’s new proof. The adaptation that we found of Hutchcroft’s proof is more involved than the argument presented here. For example, it invokes the universal tightness result from [Hut21d]. Invoking this auxiliary result

⁵Consider the sequences $(K_n \square C_m)_{n \geq 1}$ for each $m \geq 3$, where K_n is a complete graph and C_m is a cycle.

also has the consequence that the universal lower bound we ultimately obtain is weaker than the mean-field bound established here.

7.2 The coupling lemma

In this section, we control the effect on percolation of conditioning on the event that a vertex's cluster has small density. In particular, we prove that the conditioned measure *approximately* stochastically dominates percolation of a slightly smaller parameter. As mentioned in Section 7.1, the argument in this section is inspired by [Van22b].

Lemma 7.2.1. *Let $G = (V, E)$ be a finite connected transitive graph with a distinguished vertex o . Let $p \in (0, 1)$ be a parameter and $\alpha \in (0, 1)$ a density. Define*

$$\theta := \mathbb{E}_p \|K_1\|, \quad h := \mathbb{P}_p \left(\|K_1\| < \alpha \text{ or } \|K_2\| \geq \frac{\alpha}{2} \right), \quad \delta := \frac{2h^{1/2}}{1 - \theta - h},$$

and assume that $\theta + h < 1$ (so that δ is well-defined and positive). Then there is an event A with $\mathbb{P}_p(A \mid \|K_o\| < \alpha) \leq h^{1/2}$ such that

$$\mathbb{P}_{(1-\theta-\delta)p} \leq_{\text{st}} \mathbb{P}_p(\omega \cup \mathbf{1}_A = \cdot \mid \|K_o\| < \alpha),$$

where \leq_{st} denotes stochastic domination with respect to the usual partial order \preceq on $\{0, 1\}^E$, and $\mathbf{1}_A$ denotes the random configuration with every edge open on A and every edge closed on A^c .

Proof. To lighten notation, set $\hat{\mathbb{P}} := \mathbb{P}_p(\cdot \mid \|K_o\| < \alpha)$ and $q := (1 - \theta - \delta)p$. Our goal is to construct an approximately monotone coupling of \mathbb{P}_q and $\hat{\mathbb{P}}$. We will build this in the obvious way: by fixing an exploration process and building samples of \mathbb{P}_q and $\hat{\mathbb{P}}$ in terms of a common collection of E -indexed uniform random variables. The rest of the proof consists in controlling the failure of this coupling to be monotone.

Fix an enumeration of the edge set E . Recall the following standard method for exploring a configuration ω from o : Start with all edges unrevealed. Iteratively reveal the unrevealed edge of smallest index that is connected to o by an open path of revealed edges until there are none. Then iteratively reveal the unrevealed edge of smallest index from among all remaining unrevealed edges until there are none. Let $\rho_1, \dots, \rho_{|E|}$ be the sequence of edges as they are revealed by this process. Let $(\mathcal{F}_t)_{0 \leq t \leq |E|}$ be the filtration associated to this exploration, i.e.

$$\mathcal{F}_t := \sigma(\omega_{\rho_1}, \dots, \omega_{\rho_t}).$$

We now use this exploration to construct a coupling of \mathbb{P}_q and $\hat{\mathbb{P}}$. Let $(U_e)_{e \in E}$ be a collection of independent uniform-[0,1] random variables. Recursively define $\omega_{\rho_t} := \mathbf{1}_{U_{\rho_t} \leq \hat{\mathbb{P}}(\rho_t \text{ open} \mid \mathcal{F}_{t-1})}$ for

every t to obtain a configuration ω with law $\hat{\mathbb{P}}$, and simply take $(\mathbf{1}_{U_e \leq q})_{e \in E}$ for a configuration with law \mathbb{P}_q . This coupling is monotone (in the direction we want) on the edges $\rho_1, \rho_2, \dots, \rho_{\tau_{\text{fail}}}$ where τ_{fail} is the stopping time defined by

$$\tau_{\text{fail}} := \inf\{t : \hat{\mathbb{P}}(\rho_{t+1} \text{ open} \mid \mathcal{F}_t) < q\},$$

with the convention that $\inf \emptyset := |E|$. So $(\mathbf{1}_{U_e \leq q})_{e \in E} \preceq \omega$ almost surely when $\tau_{\text{fail}} = |E|$. Since $(\mathbf{1}_{U_e \leq q})_{e \in E} \preceq \mathbf{1}_{\tau_{\text{fail}} < |E|}$ holds trivially when $\tau_{\text{fail}} < |E|$, we know that

$$\mathbb{P}_q \leq_{\text{st}} \hat{\mathbb{P}}(\omega \cup \mathbf{1}_{\tau_{\text{fail}} < |E|} = \cdot).$$

So it suffices to verify that $\hat{\mathbb{P}}(\tau_{\text{fail}} < |E|) \leq h^{1/2}$. By definition of τ_{fail} , we have the upper bound

$$\hat{\mathbb{P}}(\rho_{\tau_{\text{fail}}+1} \text{ open} \mid \mathcal{F}_{\tau_{\text{fail}}}) < q \quad \text{a.s. when } \tau_{\text{fail}} < |E|. \quad (7.2.1)$$

Our first step is to prove a complementary lower bound.

Say that an edge e is *pivotal* if $\|K_o(\omega \setminus \{e\})\| < \alpha$ but $\|K_o(\omega \cup \{e\})\| \geq \alpha$. If e is open and pivotal, then $\|K_o\| \geq \alpha$, which is $\hat{\mathbb{P}}$ -almost surely impossible. So

$$\hat{\mathbb{P}}(\rho_{\tau_{\text{fail}}+1} \text{ open} \mid \mathcal{F}_{\tau_{\text{fail}}}) = \hat{\mathbb{P}}(\rho_{\tau_{\text{fail}}+1} \text{ open and not pivotal} \mid \mathcal{F}_{\tau_{\text{fail}}}) \quad \text{a.s. when } \tau_{\text{fail}} < |E|. \quad (7.2.2)$$

Suppose we reveal the edges $\rho_1, \dots, \rho_{\tau_{\text{fail}}}$ and find that $\tau_{\text{fail}} < |E|$. Note that $\rho_{\tau_{\text{fail}}+1}$ is now almost surely determined. To finish building a sample of $\hat{\mathbb{P}}$, rather than continuing our exploration process, we could first sample every unrevealed edge except $\rho_{\tau_{\text{fail}}+1}$, then sample $\rho_{\tau_{\text{fail}}+1}$ itself. The first stage will determine whether $\rho_{\tau_{\text{fail}}+1}$ is pivotal. If it is not pivotal, then conditioning on the event $\{\|K_o\| < \alpha\}$ will have no effect in the second stage, i.e. the conditional probability that $\rho_{\tau_{\text{fail}}+1}$ is open will simply be p . So we can rewrite eq. (7.2.2) as

$$\begin{aligned} \hat{\mathbb{P}}(\rho_{\tau_{\text{fail}}+1} \text{ open} \mid \mathcal{F}_{\tau_{\text{fail}}}) &= p \hat{\mathbb{P}}(\rho_{\tau_{\text{fail}}+1} \text{ not pivotal} \mid \mathcal{F}_{\tau_{\text{fail}}}) \\ &= p \left(1 - \hat{\mathbb{P}}(\rho_{\tau_{\text{fail}}+1} \text{ pivotal} \mid \mathcal{F}_{\tau_{\text{fail}}}) \right) \quad \text{a.s. when } \tau_{\text{fail}} < |E|. \end{aligned} \quad (7.2.3)$$

As in our argument for eq. (7.2.2), since it is $\hat{\mathbb{P}}$ -almost surely impossible for an edge to be both open and pivotal, this further implies that

$$\hat{\mathbb{P}}(\rho_{\tau_{\text{fail}}+1} \text{ open} \mid \mathcal{F}_{\tau_{\text{fail}}}) = p \left(1 - \hat{\mathbb{P}}(\rho_{\tau_{\text{fail}}+1} \text{ closed and pivotal} \mid \mathcal{F}_{\tau_{\text{fail}}}) \right) \quad \text{a.s. when } \tau_{\text{fail}} < |E|. \quad (7.2.4)$$

By combining inequality eq. (7.2.1) with eq. (7.2.3), we deduce that

$$\hat{\mathbb{P}}(\rho_{\tau_{\text{fail}}+1} \text{ pivotal} \mid \mathcal{F}_{\tau_{\text{fail}}}) > 0 \quad \text{a.s. when } \tau_{\text{fail}} < |E|.$$

So when $\tau_{\text{fail}} < |E|$, we can almost surely label the endpoints of $\rho_{\tau_{\text{fail}}+1}$ by v_- and v_+ such that v_- is connected to o by a path of open edges among the revealed edges $\{\rho_1, \dots, \rho_{\tau_{\text{fail}}}\}$, whereas v_+ is not.

Consider a configuration ω with $\tau_{\text{fail}} < |E|$ in which $\rho_{\tau_{\text{fail}}+1}$ is closed and pivotal. By the pigeonhole principle, since $\|K_o(\omega \cup \{\rho_{\tau_{\text{fail}}+1}\})\| \geq \alpha$,

$$\|K_{v_-}\| \geq \frac{\alpha}{2} \quad \text{or} \quad \|K_{v_+}\| \geq \frac{\alpha}{2}. \quad (7.2.5)$$

Since closing $\rho_{\tau_{\text{fail}}+1}$ disconnects $K_o(\omega \cup \{\rho_{\tau_{\text{fail}}+1}\})$, the endpoints v_- and v_+ belong to distinct clusters of ω . In particular, since $v_- \in K_o$, we know that $v_+ \notin K_o$. So eq. (7.2.5) implies that

$$\|K_o\| \geq \frac{\alpha}{2} \quad \text{or} \quad \|K_{v_+}(\omega \setminus \overline{K_o})\| \geq \frac{\alpha}{2},$$

where $\overline{K_o} := K_o \cup \partial K_o$ and ∂K_o denotes the edge boundary of K_o . Now ω was arbitrary, so by a union bound,

$$\begin{aligned} \hat{\mathbb{P}}(\rho_{\tau_{\text{fail}}+1} \text{ closed and pivotal} \mid \mathcal{F}_{\tau_{\text{fail}}}) &\leq \hat{\mathbb{P}}\left(\|K_o\| \geq \frac{\alpha}{2} \mid \mathcal{F}_{\tau_{\text{fail}}}\right) + \hat{\mathbb{P}}\left(\|K_{v_+}(\omega \setminus \overline{K_o})\| \geq \frac{\alpha}{2} \mid \mathcal{F}_{\tau_{\text{fail}}}\right) \\ &\quad \text{a.s. when } \tau_{\text{fail}} < |E|. \end{aligned} \quad (7.2.6)$$

To bound the first term in inequality (7.2.6), notice that $\|K_o\| \geq \frac{\alpha}{2}$ and $\|K_o\| < \alpha$ together imply the bad event $B := \{\|K_1\| < \alpha \text{ or } \|K_2\| \geq \frac{\alpha}{2}\}$. So because $\|K_o\| < \alpha$ occurs $\hat{\mathbb{P}}$ -almost surely,

$$\hat{\mathbb{P}}\left(\|K_o\| \geq \frac{\alpha}{2} \mid \mathcal{F}_{\tau_{\text{fail}}}\right) \leq \hat{\mathbb{P}}(B \mid \mathcal{F}_{\tau_{\text{fail}}}) \quad \text{a.s. when } \tau_{\text{fail}} < |E|. \quad (7.2.7)$$

To bound the second term in inequality (7.2.6), define a new stopping time

$$\tau_{\text{moat}} := \max\{t : \rho_t \in \overline{K_o}\}.$$

(This is defined with respect to the standard exploration described in the second paragraph of the current proof environment, not the modified exploration mentioned below eq. (7.2.2).) Just after we reveal $\rho_{\tau_{\text{moat}}}$, since conditioning on the event $\{\|K_o\| < \alpha\}$ no longer has any effect, the distribution of the configuration on the unrevealed edges is simply

$$\hat{\mathbb{P}}\left(\omega \mid_{E \setminus \overline{K_o}} = \cdot \mid \mathcal{F}_{\tau_{\text{moat}}}\right) = \mathbb{P}_p^{G \setminus \overline{K_o}} \leq_{\text{st}} \mathbb{P}_p \quad \text{a.s.} \quad (7.2.8)$$

When $\tau_{\text{fail}} < |E|$, we showed that $\rho_{\tau_{\text{fail}}+1}$ almost surely has an endpoint belonging to K_o , namely v_- . So $\tau_{\text{fail}} < |E|$ implies $\tau_{\text{fail}} < \tau_{\text{moat}}$ almost surely. By applying this observation, inequality (7.2.8),

and transitivity, we obtain

$$\begin{aligned}
\hat{\mathbb{P}}\left(\|K_{v_+}(\omega \setminus \overline{K}_o)\| \geq \frac{\alpha}{2} \mid \mathcal{F}_{\tau_{\text{fail}}}\right) &= \hat{\mathbb{E}}\left[\hat{\mathbb{P}}\left(\|K_{v_+}(\omega \setminus \overline{K}_o)\| \geq \frac{\alpha}{2} \mid \mathcal{F}_{\tau_{\text{moat}}}\right) \mid \mathcal{F}_{\tau_{\text{fail}}}\right] \\
&\leq \hat{\mathbb{E}}\left[\mathbb{P}_p\left(\|K_{v_+}\| \geq \frac{\alpha}{2}\right) \mid \mathcal{F}_{\tau_{\text{fail}}}\right] \\
&= \mathbb{P}_p\left(\|K_o\| \geq \frac{\alpha}{2}\right) \quad \text{a.s. when } \tau_{\text{fail}} < |E|.
\end{aligned}$$

When $\|K_o\| \geq \frac{\alpha}{2}$, we must have that $o \in K_1$ or $\|K_2\| \geq \frac{\alpha}{2}$. So by a union bound,

$$\begin{aligned}
\hat{\mathbb{P}}\left(\|K_{v_+}(\omega \setminus \overline{K}_o)\| \geq \frac{\alpha}{2} \mid \mathcal{F}_{\tau_{\text{fail}}}\right) &\leq \mathbb{P}_p(o \in K_1) + \mathbb{P}_p\left(\|K_2\| \geq \frac{\alpha}{2}\right) \\
&\leq \theta + h \quad \text{a.s. when } \tau_{\text{fail}} < |E|,
\end{aligned} \tag{7.2.9}$$

where $h := \mathbb{P}_p(B)$ and $\theta := \mathbb{E}_p\|K_1\|$, which satisfies $\theta = \mathbb{P}_p(o \in K_1)$ by transitivity.

Plugging inequalities (7.2.6), (7.2.7) and (7.2.9) into eq. (7.2.4) gives a lower bound on $\hat{\mathbb{P}}(\rho_{\tau_{\text{fail}}+1} \text{ open} \mid \mathcal{F}_{\tau_{\text{fail}}})$ when $\tau_{\text{fail}} < |E|$. By contrasting this with upper bound ?? and expanding the definition of q , we deduce that

$$p\left(1 - \theta - h - \hat{\mathbb{P}}(B \mid \mathcal{F}_{\tau_{\text{fail}}})\right) \leq p\left(1 - \theta - \frac{2h^{1/2}}{1 - \theta - h}\right) \quad \text{a.s. when } \tau_{\text{fail}} < |E|.$$

In particular,

$$\hat{\mathbb{P}}(B \mid \mathcal{F}_{\tau_{\text{fail}}}) \geq \frac{h^{1/2}}{1 - \theta - h} \quad \text{a.s. when } \tau_{\text{fail}} < |E|.$$

By the law of total expectation, we know that $\hat{\mathbb{E}}[\hat{\mathbb{P}}(B \mid \mathcal{F}_{\tau_{\text{fail}}})] = \hat{\mathbb{P}}(B)$. So by Markov's inequality,

$$\hat{\mathbb{P}}(\tau_{\text{fail}} < |E|) \leq \hat{\mathbb{P}}\left(\hat{\mathbb{P}}(B \mid \mathcal{F}_{\tau_{\text{fail}}}) \geq \frac{h^{1/2}}{1 - \theta - h}\right) \leq \frac{1 - \theta - h}{h^{1/2}} \hat{\mathbb{P}}(B). \tag{7.2.10}$$

By definition of $\hat{\mathbb{P}}$ and h , we have the trivial bound

$$\hat{\mathbb{P}}(B) \leq \frac{\mathbb{P}_p(B)}{\mathbb{P}_p(\|K_o\| < \alpha)} = \frac{h}{1 - \mathbb{P}_p(\|K_o\| \geq \alpha)}.$$

If $\|K_o\| \geq \alpha$, then $o \in K_1$ or $\|K_2\| \geq \alpha$. So similarly to the argument for inequality (7.2.9), a union bound gives $\hat{\mathbb{P}}(B) \leq \frac{h}{1 - \theta - h}$. Plugging this into inequality (7.2.10) yields $\hat{\mathbb{P}}(\tau_{\text{fail}} < |E|) \leq h^{1/2}$, as required. \square

7.3 Characterising the existence of a percolation threshold

In this section, we prove Theorem 8.1.1 and Corollary 7.1.4. Our first step is to establish a mean-field lower bound on the supercritical giant cluster density for sequences of graphs without molecular subsequences. This is where we use the coupling lemma from Section 7.2.

Lemma 7.3.1. *Let (G_n) be a sequence of finite connected transitive graphs with volume tending to infinity that does not contain a molecular subsequence. Fix $\varepsilon > 0$ and let (p_n) be any sequence of parameters. If $\liminf_{n \rightarrow \infty} \mathbb{E}_{p_n}^{G_n} \|K_1\| > 0$, then*

$$\liminf_{n \rightarrow \infty} \mathbb{E}_{(1+\varepsilon)p_n}^{G_n} \|K_1\| \geq \frac{\varepsilon}{1 + \varepsilon}.$$

Proof. Suppose for contradiction that $\liminf_{n \rightarrow \infty} \mathbb{E}_{p_n}^{G_n} \|K_1\| > 0$ but $\liminf_{n \rightarrow \infty} \mathbb{E}_{(1+\varepsilon)p_n}^{G_n} \|K_1\| < \frac{\varepsilon}{1 + \varepsilon}$. By passing to a suitable subsequence, we may assume that

$$\limsup_{n \rightarrow \infty} \mathbb{E}_{(1+\varepsilon)p_n}^{G_n} \|K_1\| < \frac{\varepsilon}{1 + \varepsilon}. \quad (7.3.1)$$

Pick $\alpha > 0$ such that $\mathbb{E}_{p_n}^{G_n} \|K_1\| \geq 2\alpha$ for all sufficiently large n . By Markov's inequality applied to $1 - \|K_1\|$, this implies that $\mathbb{P}_{p_n}^{G_n} (\|K_1\| \geq \alpha) \geq \alpha$ for all sufficiently large n . Now for each n , define

$$\theta_n := \mathbb{E}_{(1+\varepsilon)p_n}^{G_n} \|K_1\| \quad \text{and} \quad h_n := \mathbb{P}_{(1+\varepsilon)p_n}^{G_n} \left(\|K_1\| < \alpha \text{ or } \|K_2\| \geq \frac{\alpha}{2} \right).$$

Since (G_n) does not contain a molecular subsequence, it has the sharp-density and supercritical uniqueness properties from [EH21b]. (See Section 7.1 for more information.) So the fact that $\liminf_{n \rightarrow \infty} \mathbb{P}_{p_n}^{G_n} (\|K_1\| \geq \alpha) > 0$ implies that

$$\lim_{n \rightarrow \infty} \mathbb{P}_{(1+\varepsilon)p_n}^{G_n} (\|K_1\| < \alpha) = 0 \quad \text{and} \quad \lim_{n \rightarrow \infty} \mathbb{P}_{(1+\varepsilon)p_n}^{G_n} \left(\|K_2\| \geq \frac{\alpha}{2} \right) = 0.$$

So by a union bound, $\lim_{n \rightarrow \infty} h_n = 0$. Together with inequality (7.3.1), this guarantees that for all sufficiently large n , we have $\theta_n + h_n < 1$, allowing us to apply Lemma 7.2.1. By passing to a suitable tail of (G_n) , we may assume that this holds for every n . So for every n , Lemma 7.2.1 says that there is an event $A_n \subseteq \{0, 1\}^{E(G_n)}$ with

$$\mathbb{P}_{(1+\varepsilon)p_n}^{G_n} (A_n \mid \|K_o\| < \alpha) \leq h_n^{1/2} \quad (7.3.2)$$

such that

$$\mathbb{P}_{\left(1 - \theta_n - \frac{2h_n^{1/2}}{1 - \theta_n - h_n}\right)(1+\varepsilon)p_n}^{G_n} \leq_{\text{st}} \mathbb{P}_{(1+\varepsilon)p_n}^{G_n} (\omega \cup \mathbf{1}_{A_n} = \cdot \mid \|K_o\| < \alpha). \quad (7.3.3)$$

When $\mathbb{P}_{p_n}^{G_n} (\|K_1\| \geq \alpha) \geq \alpha$, it follows by transitivity that $\mathbb{P}_{p_n}^{G_n} (\|K_o\| \geq \alpha) \geq \alpha^2$. On the other hand, by inequalities (7.3.2) and (7.3.3), we know that

$$\mathbb{P}_{\left(1 - \theta_n - \frac{2h_n^{1/2}}{1 - \theta_n - h_n}\right)(1+\varepsilon)p_n}^{G_n} (\|K_o\| \geq \alpha) \leq h_n^{1/2}.$$

So by taking n sufficiently large that both $\mathbb{P}_{p_n}^{G_n} (\|K_1\| \geq \alpha) \geq \alpha$ and $h_n^{1/2} < \alpha^2$, we can force

$$\left(1 - \theta_n - \frac{2h_n^{1/2}}{1 - \theta_n - h_n}\right)(1 + \varepsilon)p_n \leq p_n. \quad (7.3.4)$$

However, by inequality (7.3.1) and the fact that $\lim_{n \rightarrow \infty} h_n = 0$,

$$1 - \theta_n - \frac{2h_n^{1/2}}{1 - \theta_n - h_n} > \frac{1}{1 + \varepsilon}$$

for all sufficiently large n , contradicting inequality (7.3.4). \square

By the sharp-density property, we can convert this mean-field lower bound that holds in *expectation* into one that holds *with high probability*.

Lemma 7.3.2. *Let (G_n) be a sequence of finite connected transitive graphs with volume tending to infinity that does not contain a molecular subsequence. Fix $\varepsilon > 0$ and let (p_n) be any sequence of parameters. If $\liminf_{n \rightarrow \infty} \mathbb{E}_{p_n}^{G_n} \|K_1\| > 0$, then*

$$\mathbb{P}_{(1+\varepsilon)p_n}^{G_n} \left(\|K_1\| \geq \frac{\varepsilon}{1 + \varepsilon} - o(1) \right) = 1 - o(1) \quad \text{as } n \rightarrow \infty.$$

Proof. Let $\delta \in (0, \varepsilon)$ be any constant. By Lemma 7.3.1,

$$\liminf_{n \rightarrow \infty} \mathbb{E}_{(1+\varepsilon-\delta)p_n}^{G_n} \|K_1\| \geq \frac{\varepsilon - \delta}{1 + \varepsilon - \delta}.$$

So by Markov's inequality applied to $1 - \|K_1\|$,

$$\liminf_{n \rightarrow \infty} \mathbb{P}_{(1+\varepsilon-\delta)p_n}^{G_n} \left(\|K_1\| \geq \frac{\varepsilon - \delta}{1 + \varepsilon} \right) \geq \frac{\delta(\varepsilon - \delta)}{(1 + \delta)(1 + \varepsilon - \delta)} > 0. \quad (7.3.5)$$

Since (G_n) does not contain a molecular subsequence, it has the sharp-density property from [EH21b]. (Recall our discussion in Section 7.1.) So inequality (7.3.5) implies that

$$\lim_{n \rightarrow \infty} \mathbb{P}_{(1+\varepsilon)p_n}^{G_n} \left(\|K_1\| \geq \frac{\varepsilon - \delta}{1 + \varepsilon} \right) = 1.$$

Since δ was arbitrary, the result now follows by a diagonal argument. \square

Our next step is to extend a version of this lower bound to a simple kind of molecular sequence. The idea is to break the graphs in the sequence into their constituent *atoms* then apply Lemma 7.3.2 to the sequence of atoms.

Lemma 7.3.3. *Let (G_n) be a sequence of finite connected transitive graphs with volume tending to infinity. Assume that (G_n) is m -molecular for some $m \geq 2$ but does not contain a k -molecular subsequence for any $k > m$. Fix $\varepsilon > 0$ and let (p_n) be any sequence of parameters. If $\liminf_{n \rightarrow \infty} \mathbb{E}_{p_n}^{G_n} \|K_1\| > 0$, then*

$$\mathbb{P}_{(1+\varepsilon)p_n}^{G_n} \left(\|K_1\| \geq \frac{\varepsilon}{m(1+\varepsilon)} - o(1) \right) = 1 - o(1) \quad \text{as } n \rightarrow \infty.$$

Proof. By definition of an m -molecular sequence, for every n , we can pick a set of edges $F_n \subseteq E(G_n)$ such that F_n is $\text{Aut } G_n$ -invariant, $G_n \setminus F_n$ has m connected components, and $\frac{|F_n|}{|V(G_n)|}$ is uniformly bounded. Let A_n denote one of the connected components of $G_n \setminus F_n$, which are all necessarily isomorphic to each other and transitive. Notice that (A_n) does not contain a molecular subsequence. Indeed, if (A_n) contained an r -molecular subsequence $(A_n)_{n \in I}$ for some $r \geq 2$, then $(G_n)_{n \in I}$ would be an rm -molecular subsequence of (G_n) . (Every automorphism of G_n acts on the m copies of A_n by permuting the copies and applying an automorphism of A_n to each.)

For each finite graph G , let $\lambda(G)$ denote the largest eigenvalue of the adjacency matrix for G . Recall that when G is regular, $\lambda(G) = \deg(G)$, the vertex degree of G . By Theorem 7.1.3, which is taken from [Bol+10c], since (G_n) and (A_n) are sequences of dense graphs, they have percolation thresholds at $(1/\lambda(G_n))_{n \geq 1}$ and $(1/\lambda(A_n))_{n \geq 1}$ respectively. Since $\frac{|F_n|}{|V(G_n)|}$ is uniformly bounded, $\lim_{n \rightarrow \infty} \frac{\deg A_n}{\deg G_n} = 1$, and since every A_n and G_n is regular (since transitive), this means that $\lim_{n \rightarrow \infty} \frac{\lambda(A_n)}{\lambda(G_n)} = 1$. So for any $\delta > 0$, since the sequence $((1+\delta)p_n)$ is supercritical for (G_n) , it is also supercritical for (A_n) . In particular,

$$\liminf_{n \rightarrow \infty} \mathbb{E}_{(1+\delta)p_n}^{A_n} \|K_1\| > 0.$$

Now since (A_n) has no molecular subsequences, it follows by Lemma 7.3.2 that

$$\mathbb{P}_{(1+\varepsilon)p_n}^{A_n} \left(\|K_1\| \geq \frac{\varepsilon - \delta}{1 + \varepsilon} - o(1) \right) = 1 - o(1) \quad \text{as } n \rightarrow \infty.$$

Since $\delta > 0$ was arbitrary, a diagonal argument gives

$$\mathbb{P}_{(1+\varepsilon)p_n}^{A_n} \left(\|K_1\| \geq \frac{\varepsilon}{1 + \varepsilon} - o(1) \right) = 1 - o(1) \quad \text{as } n \rightarrow \infty.$$

The result follows because A_n is a subgraph of G_n and $|V(G_n)| = m |V(A_n)|$. \square

We now extend this lower bound to sequences of graphs that have m -molecular subsequences for at most finitely many integers m .

Lemma 7.3.4. *Let (G_n) be a sequence of finite connected transitive graphs with volume tending to infinity that does not contain an m -molecular subsequence for any $m > M$, where M is some positive integer. Fix $\varepsilon > 0$ and let (p_n) be any sequence of parameters. If $\liminf_{n \rightarrow \infty} \mathbb{E}_{p_n}^{G_n} \|K_1\| > 0$, then*

$$\mathbb{P}_{(1+\varepsilon)p_n}^{G_n} \left(\|K_1\| \geq \frac{\varepsilon}{M(1+\varepsilon)} - o(1) \right) = 1 - o(1) \quad \text{as } n \rightarrow \infty.$$

Proof. It is enough to show that for every subsequence $(G_n)_{n \in I}$ of (G_n) , we can find a further subsequence $(G_n)_{n \in J}$ with $J \subseteq I$ such that

$$\mathbb{P}_{(1+\varepsilon)p_n}^{G_n} \left(\|K_1\| \geq \frac{\varepsilon}{M(1+\varepsilon)} - o(1) \right) = 1 - o(1) \quad \text{as } n \rightarrow \infty \text{ with } n \in J.$$

Let $(G_n)_{n \in I}$ be a subsequence of (G_n) . If $(G_n)_{n \in I}$ does not contain a molecular subsequence, then by Lemma 7.3.2, we know that $\mathbb{P}_{(1+\varepsilon)p_n}^{G_n} (\|K_1\| \geq \frac{\varepsilon}{1+\varepsilon} - o(1)) = 1 - o(1)$ as $n \rightarrow \infty$ with $n \in I$. On the other hand, if $(G_n)_{n \in I}$ does contain a molecular subsequence, then we can pick a subsequence $(G_n)_{n \in J}$ with $J \subseteq I$ that is m -molecular with $m \in \{2, \dots, M\}$ maximum. Then Lemma 7.3.3 tells us that $\mathbb{P}_{(1+\varepsilon)p_n}^{G_n} (\|K_1\| \geq \frac{\varepsilon}{m(1+\varepsilon)} - o(1)) = 1 - o(1)$ as $n \rightarrow \infty$ with $n \in J$. \square

We are now ready to prove that if a sequence of graphs has m -molecular subsequences for at most finitely many integers m , then it has a percolation threshold. As outlined in Section 7.1, the idea is to prove that the threshold for the emergence of a cluster of very slightly sublinear density is a percolation threshold.

Lemma 7.3.5. *If a sequence of finite connected transitive graphs with volume tending to infinity contains an m -molecular subsequence for at most finitely many integers m , then it has a percolation threshold.*

Proof. Let (G_n) be such a sequence of graphs. Let M be a positive integer such that there are no m -molecular subsequences with $m > M$. For every density $\delta \in (0, 1)$ and each index n , define $p_c^n(\delta)$ to be the unique parameter satisfying $\mathbb{P}_{p_c^n(\delta)}^{G_n} (\|K_1\| \geq \delta) = \frac{1}{2}$. Given any constants $\varepsilon, \delta \in (0, 1)$, we know by Lemma 7.3.4 that

$$\lim_{n \rightarrow \infty} \mathbb{P}_{(1+\varepsilon)p_c^n(\delta)}^{G_n} \left(\|K_1\| \geq \frac{\varepsilon}{2M} \right) = 1,$$

since $\frac{\varepsilon}{2M} < \frac{\varepsilon}{M(1+\varepsilon)}$. So by a diagonal argument, for every $\varepsilon \in (0, 1)$, there exists a sequence $(\delta_n^\varepsilon)_{n \geq 1}$ in $(0, 1)$ with $\lim_{n \rightarrow \infty} \delta_n^\varepsilon = 0$ such that

$$\lim_{n \rightarrow \infty} \mathbb{P}_{(1+\varepsilon)p_c^n(\delta_n^\varepsilon)}^{G_n} \left(\|K_1\| \geq \frac{\varepsilon}{2M} \right) = 1.$$

Now by a further diagonal argument, there exists a fixed sequence (δ_n) in $(0, 1)$ with $\lim_{n \rightarrow \infty} \delta_n = 0$ such that for every $\varepsilon \in \{\frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \dots\}$,

$$\lim_{n \rightarrow \infty} \mathbb{P}_{(1+\varepsilon)p_c^n(\delta_n)}^{G_n} \left(\|K_1\| \geq \frac{\varepsilon}{2M} \right) = 1.$$

We claim that the sequence of parameters (p_n) given by $p_n := p_c^n(\delta_n)$ has the required properties.

To verify the subcritical condition, suppose for contradiction that $\lim_{n \rightarrow \infty} \mathbb{P}_{(1-\varepsilon)p_n}^{G_n} (\|K_1\| \geq \alpha) \neq 0$ for some $\varepsilon, \alpha \in (0, 1)$. By passing to a suitable subsequence, we may assume that

$$\liminf_{n \rightarrow \infty} \mathbb{P}_{(1-\varepsilon)p_n}^{G_n} (\|K_1\| \geq \alpha) > 0.$$

Then by Lemma 7.3.4,

$$\lim_{n \rightarrow \infty} \mathbb{P}_{p_n}^{G_n} \left(\|K_1\| \geq \frac{\varepsilon}{2M} \right) = 1.$$

This contradicts the fact that for every n that is sufficiently large to ensure $\delta_n \leq \frac{\varepsilon}{2M}$,

$$\mathbb{P}_{p_n}^{G_n} \left(\|K_1\| \geq \frac{\varepsilon}{2M} \right) \leq \mathbb{P}_{p_n}^{G_n} (\|K_1\| \geq \delta_n) = \frac{1}{2}.$$

To verify the supercritical condition, fix $\varepsilon > 0$. Pick any $\varepsilon' \in \{\frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \dots\}$ with $\varepsilon' < \varepsilon$. Then by construction of (δ_n) ,

$$\mathbb{P}_{(1+\varepsilon)p_n}^{G_n} \left(\|K_1\| \geq \frac{\varepsilon'}{2M} \right) \geq \mathbb{P}_{(1+\varepsilon')p_n}^{G_n} \left(\|K_1\| \geq \frac{\varepsilon'}{2M} \right) = 1 - o(1) \quad \text{as } n \rightarrow \infty. \quad \square$$

We now verify that these sequences of graphs — those that have m -molecular subsequences for at most finitely many integers m — are the *only* sequences to have a percolation threshold. This follows from the following corollary of Theorem 7.1.7 from [EH21b] that characterises the *supercritical existence property*.

Definition 7.3.6. We say that (G_n) has the *supercritical existence property* if for every supercritical sequence of parameters (q_n) ,

$$\lim_{n \rightarrow \infty} \mathbb{P}_{q_n}^{G_n} (\|K_1\| \geq \alpha) = 1 \quad \text{for some } \alpha > 0.$$

Corollary 7.3.7 (Easo and Hutchcroft, 2021). *A sequence of finite connected transitive graphs with volume tending to infinity has the supercritical existence property if and only if it contains an m -molecular subsequence for at most finitely many integers m .*

Notice that if a sequence of finite connected transitive graphs with volume tending to infinity has a percolation threshold, then it automatically has the supercritical existence property. So the ‘only if’ direction of Corollary 7.3.7 immediately implies the following lemma.

Lemma 7.3.8. *If a sequence of finite connected transitive graphs with volume tending to infinity contains an m -molecular subsequence for infinitely many integers m , then it does not have a percolation threshold.*

Our main result, Theorem 8.1.1, now follows by combining Lemmas 7.3.5 and 7.3.8. We conclude this section by deducing Corollary 7.1.4 from Lemma 7.3.4.

Proof of Corollary 7.1.4. For every $\delta \in (0, \varepsilon)$, we know by definition of a percolation threshold that $\liminf_{n \rightarrow \infty} \mathbb{E}_{(1+\delta)p_n}^{G_n} \|K_1\| > 0$. So by Lemma 7.3.4,

$$\mathbb{P}_{(1+\varepsilon)p_n}^{G_n} \left(\|K_1\| \geq \frac{\varepsilon - \delta}{M(1+\varepsilon)} - o(1) \right) = 1 - o(1) \quad \text{as } n \rightarrow \infty.$$

Since δ was arbitrary, the result follows by a diagonal argument. \square

7.4 Bounding the threshold location

In this section, we give a proof of Proposition 7.1.5. This is simply the observation that existing arguments immediately imply bounds on the location of a percolation threshold (when it exists), which happen to be best-possible. The lower bound is a completely standard path-counting argument that we only include for completeness. The upper bound comes from bounding the percolation threshold by the connectivity threshold. As explained in [GLL21], we can estimate the connectivity threshold thanks to an upper bound by Karger and Stein on the number of *approximate* minimum cutsets in a graph [KS96].

The lower bound is sharp in the classical case of the complete graphs [ER60]. In fact, it is sharp for a very wide range of examples, including arbitrary dense regular graphs [Bol+10c], the hypercubes [AKS82a] (which are sparse yet have unbounded vertex degrees), and expanders with high girth and bounded vertex degrees [ABS04c]. However, it is less obvious that the upper bound is optimal in the sense that the constant '2' cannot be improved. For this, see the *fat cycles* construction in [GLL21].

Proof of Proposition 7.1.5. We start with the lower bound. This argument is just an optimisation of the proof of Lemma 2.8 in [EH21b]. Fix $\varepsilon \in (0, 1)$. It suffices to check that $\left(\frac{1-\varepsilon}{d_n-1}\right)_{n \geq 1}$ is not supercritical along any subsequence. There are at most $d_n(d_n - 1)^{r-1}$ simple paths of length r starting at a particular vertex o in each G_n . So for every n ,

$$\mathbb{E}_{\frac{1-\varepsilon}{d_n-1}}^{G_n} |K_o| \leq \sum_{r \geq 0} d_n(d_n - 1)^{r-1} \left(\frac{1-\varepsilon}{d_n-1}\right)^r \leq \frac{2}{\varepsilon}.$$

In particular, $\lim_{n \rightarrow \infty} \mathbb{E}_{\frac{1-\varepsilon}{d_n-1}}^{G_n} \|K_o\| = 0$, and hence by transitivity, $\lim_{n \rightarrow \infty} \mathbb{E}_{\frac{1-\varepsilon}{d_n-1}} \|K_1\| = 0$.

We now prove the upper bound. Every finite connected transitive graph has edge connectivity equal to its vertex degree [Mad71]. So by Theorem 4.1 from [GLL21],

$$\lim_{n \rightarrow \infty} \mathbb{P}_{(2+o(1))\frac{\log|V(G_n)|}{d_n}}^{G_n} (\omega \text{ is connected}) = 1.$$

In particular, $\left(\frac{2\log|V(G_n)|}{d_n}\right)_{n \geq 1}$ is not subcritical along any subsequence. \square

7.5 Acknowledgements

I thank Tom Hutchcroft for helpful comments on earlier drafts and for encouraging me to pursue this question of mine. I also thank an anonymous referee for their useful feedback and suggestions, which improved the clarity of the paper.

SHARPNESS AND LOCALITY FOR PERCOLATION ON FINITE TRANSITIVE GRAPHS

Abstract

Let $(G_n) = ((V_n, E_n))$ be a sequence of finite connected vertex-transitive graphs with uniformly bounded vertex degrees such that $|V_n| \rightarrow \infty$ as $n \rightarrow \infty$. We say that percolation on G_n has a *sharp* phase transition (as $n \rightarrow \infty$) if, as the percolation parameter crosses some critical point, the number of vertices contained in the largest percolation cluster jumps from logarithmic to linear order with high probability. We prove that percolation on G_n has a sharp phase transition unless, after passing to a subsequence, the rescaled graph-metric on G_n (rapidly) converges to the unit circle with respect to the Gromov-Hausdorff metric. We deduce that under the same hypothesis, the critical point for the emergence of a giant (i.e. linear-sized) cluster in G_n coincides with the critical point for the emergence of an infinite cluster in the Benjamini-Schramm limit of (G_n) , when this limit exists.

8.1 Introduction

Given a graph G , build a random spanning subgraph ω by independently including each edge of G with a fixed probability $p \in [0, 1]$. The law of ω is called (Bernoulli bond) percolation and is denoted by \mathbb{P}_p^G . This simple model often undergoes a phase transition: for many natural choices of the underlying graph G , as p increases past some critical value $p_c(G)$, the typical behaviour of the connected components of ω changes abruptly. The study of this phenomenon has two origins, roughly coming from mathematical physics and combinatorics, respectively.

The first origin is the 1957 work of Broadbent and Hammersley [BH57b] introducing percolation on the Euclidean lattice $G = \mathbb{Z}^d$ as a model for the spread of fluid through a porous medium. Note that Euclidean lattices are always (vertex-)transitive, meaning that for all vertices u and v , there is a graph automorphism that maps u to v . This is a way to formalise the notion that a graph is homogeneous or that its vertices are indistinguishable. For example, every Cayley graph of a finitely-generated group is transitive. In 1996, Benjamini and Schramm [BS96b] launched the systematic study of percolation on general infinite transitive graphs. A cornerstone of this theory is that percolation on an infinite transitive graph G always undergoes a *sharp* phase transition. Let us recall what this means. We will write o to denote an arbitrary vertex in G and write $|K_o|$

to denote the cardinality of its *cluster*, i.e. connected component in ω .¹ There is a trivial sense in which percolation on G always undergoes a phase transition: by Kolmogorov's 0-1 law, there exists some critical point $p_c(G) \in [0, 1]$ such that \mathbb{P}_p^G (there exists an infinite cluster) equals 0 for all $p < p_c(G)$ and equals 1 for all $p > p_c(G)$. Now the phase transition is said to be *sharp* if for all $p < p_c(G)$, not only does $\mathbb{P}_p^G(|K_o| \geq n) \rightarrow 0$ as $n \rightarrow \infty$, but in fact there exists a constant $c(G, p) > 0$ such that $\mathbb{P}_p^G(|K_o| \geq n) \leq e^{-cn}$ for every $n \geq 1$.² This was first proved in [AB87a; Men86] and now has multiple modern proofs [DT16a; DRT19; Hut20d; Van24].

The second origin is the 1960 work of Erdős and Rényi [ER60] investigating percolation on the complete graph G_n with n vertices. This is the celebrated Erdős-Rényi (or simply *random graph*) model. The fundamental result is that percolation on G_n undergoes a *sharp* phase transition around $p = 1/n$ in the sense that for any fixed $\varepsilon > 0$, the cardinality of the largest cluster of ω under \mathbb{P}_p^G jumps from being³ $\Theta(\log n)$ at $p = (1-\varepsilon)/n$ to being $\Theta(n)$ at $p = (1+\varepsilon)/n$ with high probability as $n \rightarrow \infty$.⁴ Analogous results have since been established for certain other families of finite graphs with diverging degrees. For example, Ajtai, Komlós, and Szemerédi [AKS82b] and Bollobás, Kohakayawa, and Łuksak [BKL92] investigated percolation on the hypercube $H_d = \{0, 1\}^d$, which has a sharp phase transition around $p = 1/d$. Note that every complete graph and hypercube is transitive. For a small sample of the vast literature on percolation on finite graphs, see, for example, [ABS04a; KLS20] on expanders, [FKM04] on pseudorandom graphs, [Bor+05a; Bor+05b; Bor+06; Nac09] on transitive graphs satisfying certain mean-field conditions, [Bol+10c] on dense graphs, and [Dis+24; DK24a; DK24b] on general graphs satisfying certain isoperimetric conditions.

Between these two settings lies the less-developed theory of percolation on bounded-degree finite transitive graphs. This theory, which started in 2001, was initiated by Benjamini [Ben01b] and by Alon, Benjamini, and Stacey [ABS04a]. This concerns the asymptotic properties of percolation on a finite transitive graph $G = (V, E)$ as $|V|$ becomes large while the vertex degrees of G remain bounded. As with the Erdős-Rényi model, here we are primarily interested in the phase transition for the emergence of a *giant* cluster, i.e. a cluster containing $\Theta(|V|)$ vertices, and we will call the phase transition *sharp* if the size of the largest cluster jumps from $\Theta(\log |V|)$ to $\Theta(|V|)$. (See Section 8.1 for precise definitions.) At the same time, this theory is closely related to percolation on infinite transitive graphs via the local (Benjamini-Schramm) topology on the set of all transitive

¹More generally, K_u denotes the cluster containing a vertex called u .

²Some people use sharpness to mean slightly different things e.g. the exponential decay of point-to-point connection probabilities for $p < p_c$ together with the mean-field lower bound for $p > p_c$.

³Given functions $f, g : \mathbb{N} \rightarrow (0, \infty)$, we write $f(n) = \Theta(g(n))$ to mean that there are constants $c > 0$ and $C < \infty$ such that $cg(n) \leq f(n) \leq Cg(n)$ for all n , i.e. $f(n) = O(g(n))$ and $g(n) = O(f(n))$.

⁴When $p > 1$ or $p < 1$, we define \mathbb{P}_p to be \mathbb{P}_1 or \mathbb{P}_0 respectively.

graphs. Indeed, with respect to this topology, every infinite set \mathcal{G} of finite transitive graphs with bounded degrees is relatively compact, and every graph in the boundary of \mathcal{G} is infinite.

Despite this close relation between infinite transitive graphs and bounded-degree finite transitive graphs, our understanding of percolation on infinite transitive graphs is quite far ahead. Roughly speaking, we can think of the theory of percolation on infinite transitive graphs as the theory of percolation on microscopic (i.e. $O(1)$) scales in bounded-degree finite transitive graphs. In this sense, the finite graph theory generalises the infinite graph theory. (A limitation of this maxim is that not every infinite transitive graph can be locally approximated by finite transitive graphs.) In particular, certain basic questions in the finite graph theory have no natural analogues in the infinite graph theory. For example, the uniqueness/non-uniqueness of giant clusters is not directly related to the uniqueness/non-uniqueness of infinite clusters, which is instead related to the microscopic metric distortion of giant clusters [EH21a, Remark 1.6].

In this paper we investigate the following pair of closely related questions. An affirmative answer to the second question provides a direct way to move results and conjectures about infinite transitive graphs to finite transitive graphs.

1. Does percolation on a large bounded-degree finite transitive graph G have a sharp phase transition?
2. If a finite transitive graph G and an infinite transitive graph H are close in the local sense, does the critical point for the emergence of a giant cluster in G approximately coincide with the critical point for the emergence of an infinite cluster in H ?

Unfortunately, the answer to both of these questions in general is *no*. For example, take the sequence $(\mathbb{Z}_n \times \mathbb{Z}_{f(n)})_{n=1}^\infty$ for any $f : \mathbb{N} \rightarrow \mathbb{N}$ growing fast. This sequence always converges locally to \mathbb{Z}^2 , where the critical point for the emergence of an infinite cluster is $p_c = \frac{1}{2}$. On the other hand, provided that f grows sufficiently fast, the threshold for the emergence of a giant cluster in $\mathbb{Z}_n \times \mathbb{Z}_{f(n)}$ will be as in the sequence of cycles, around $p_c = 1$. Moreover, for percolation of any fixed parameter $p \in (\frac{1}{2}, 1)$ on $\mathbb{Z}_n \times \mathbb{Z}_{f(n)}$, the order of the largest cluster will then typically be much larger than logarithmic but much smaller than linear in the total number of vertices. (See [EH23b, Example 5.1] for some more discussion of these sequences.) The problem is that these graphs are long and thin, coarsely resembling long cycles. In particular, after suitably rescaling, their graph metrics (rapidly) converge in the Gromov-Hausdorff metric to the unit circle. In this paper we prove that this is the only possible obstacle.

Locality

Question (2) above is the finite analogue of *Schramm's locality conjecture*. This conjecture was (equivalently) that for all $\varepsilon > 0$ there exists $R < \infty$ such that for every pair of infinite transitive graphs G and H that are not one-dimensional⁵, if the ball of radius R in G is isomorphic to the ball of radius R in H , then $|p_c(G) - p_c(H)| \leq \varepsilon$. This conjecture formalised the idea that the critical point of an infinite transitive graph should generally be entirely determined by the graph's small-scale, *local* geometry. By building on earlier progress, especially the work of Contreras, Martineau, and Tassion [CMT22], we verified this conjecture in our joint work with Hutchcroft [EH23b]. Schramm's locality conjecture for infinite transitive graphs also spurred research on locality in other settings, including much research on the analogue of our question (2) about locality for finite graphs but where the hypothesis that the finite graphs are transitive is replaced by the hypothesis that they are *expanders* [BNP11a; Sar21b; RS22b; ABS23].

It may be surprising, from the perspective of percolation on infinite transitive graphs, that in fact sharpness and locality for finite transitive graphs are equivalent. That is to say, if we restrict ourselves to any particular infinite set \mathcal{G} of bounded-degree finite transitive graphs, then the answers to questions (1) and (2) in the introduction will always coincide. (See Proposition 8.2.9 for a precise statement.) Indeed, if \mathcal{G} satisfies locality, then one can easily extract sharpness for \mathcal{G} from the sharpness of the phase transition on every infinite transitive graph that is a local limit of graphs in \mathcal{G} , and the converse, that sharpness implies locality, can also be established with a little more work. One reason that this equivalence may be surprising is because for infinite transitive graphs, sharpness always holds, even for \mathbb{Z} , whereas locality requires that the graphs are not one-dimensional. To make sense of this, consider that for infinite transitive graphs, locality corresponds to a version of sharpness that is *uniform* in the choice of the graph, whereas in the context of finite transitive graphs, the only meaningful notion of sharpness is necessarily uniform.

Given the similarity between locality for finite and infinite graphs, one may wonder why the present paper is necessary: why does the proof of locality for infinite transitive graphs not also imply (perhaps after some additional bookkeeping) locality and hence sharpness for finite transitive graphs? The most fundamental reason is that the approach to proving locality in [EH23b] relied inherently on the sharpness of the phase transition, which in our setting is what we are trying to prove! Let us be a little more precise. In the proof of [EH23b], we have an infinite transitive graph G and a parameter p that we want to show satisfies $p \geq p_c(G)$. The bulk of the argument in [EH23b] involves delicately propagating point-to-point connection lower bounds across larger

⁵An infinite transitive graph is one-dimensional if and only if the graph is quasi-isometric to \mathbb{Z} .

and larger scales to ultimately establish that for some function $f : \mathbb{N} \rightarrow (0, 1)$ tending to zero slower than exponentially, $\mathbb{P}_p(u \leftrightarrow v) \geq f(\text{dist}(u, v))$ for all vertices u and v . Since point-to-point connection probabilities are decaying slower than exponentially, the conclusion $p \geq p_c(G)$ then follows from the sharpness of the phase transition on infinite transitive graphs. In a finite graph adaptation of this argument, at this final stage we would need to invoke the sharpness of the phase transition for finite transitive graphs, making the argument circular. One might hope to circumvent this problem by improving the locality argument so that the function f does not tend to zero at all. Unfortunately, f tends to zero because the propagation of point-to-point lower bounds in the locality argument is *lossy*, i.e. a lower bound of ε_i at scale n_i is propagated to a lower bound of ε_{i+1} at scale n_{i+1} where $\varepsilon_{i+1} \ll \varepsilon_i$, which seems completely unavoidable to us with current technology.

We will exploit the fact that the locality argument produces an explicit choice for f that decays much slower than exponentially (even slower than algebraically). So for this final step, one only needs a weaker kind of *quasi-sharpness* of the phase transition to conclude. The new idea in the present paper is to directly establish this quasi-sharpness by applying quantitative versions of the proofs of two results that are a priori quite unrelated to locality: the uniqueness of the supercritical giant cluster [EH21a] and the existence of a percolation threshold [Eas23] on finite transitive graphs. In short, we can think of the existence of a percolation threshold as the weakest possible kind of *quasi-sharpness*. In general, if we allow graphs to have unbounded degrees (as we did in [Eas23]), then the implicit rates of convergence can be arbitrarily slow. Luckily, now assuming bounded degrees as we may in the present paper, we can plug into our argument in [Eas23] a quantitatively strong version of the uniqueness of the supercritical giant cluster from [EH21a] to get a quantitatively strong quasi-sharpness that suffices to conclude the proof of locality.

There are also quite serious obstacles to adapting to finite graphs the part of the proof of locality leading up to this application of sharpness. To illustrate, say we tried to run the locality argument on an infinite transitive graph that *is* one-dimensional. What would go wrong? We would encounter a scale where we are unable to efficiently propagate connection lower bounds because two otherwise complementary arguments simultaneously break down. The breakdown of the first argument implies that G cannot be one-ended (G is⁶ the Cayley graph of a finitely-presented group but its minimal cutsets are not coarsely connected), while the breakdown of the second implies that G must have finitely many ends (G has polynomial growth because G contains a large ball with small tripling). From this we deduce that G is two-ended, thereby successfully identifying that G was one-dimensional. On a finite transitive graph, these end-counting arguments are not applicable. This will require us to make the locality argument more finitary, even in the setting of infinite

⁶Technically this applies to a certain graph G' that approximates G .

transitive graphs, which is of independent interest. Unfortunately, this end-counting argument is so deeply embedded in the proof of [EH23b] that it will take some work to reorganise the high-level multi-scale induction in [EH23b] in order to isolate and make explicit the relevant part. Another obstacle is that the definition of *exposed spheres*, whose special connectivity properties played a pivotal role in [CMT22; EH23b], degenerates on finite transitive graphs. As part of our argument, we introduce the *exposed sphere in a finite transitive graph*, justify our definition (Lemma 8.3.19), and thereby establish that from the perspective of part of our argument, arbitrary finite transitive graphs can be treated like infinite transitive graphs that are one-ended. We hope that these basic geometric objects can be of use in future work on finite transitive graphs, analogously to their infinite counterparts.

Statement of the main result

Graphs will always be assumed to be connected, simple, countable, and locally finite. In a slight abuse of language, we identify together all graphs that are isomorphic to each other.⁷ Let \mathcal{G} be an infinite set of finite transitive graphs. Note that \mathcal{G} is countable. We will write $\lim_{G \in \mathcal{G}}$ to denote limits taken with respect to some (and hence every) enumeration of \mathcal{G} . We may omit references to G and \mathcal{G} when this does not cause confusion. Given a graph G , we will also assume by default that V and E refer to the sets of vertices and edges in G .

Given a percolation configuration ω , we write $|K_1|$ to denote the cardinality of the largest cluster. A sequence $p : \mathcal{G} \rightarrow (0, 1)$ is said to be a *percolation threshold* if for every constant $\varepsilon > 0$, we have⁸ $\lim \mathbb{P}_{(1+\varepsilon)p}(|K_1| \geq \alpha |V|) = 1$ for some constant $\alpha > 0$, whereas $\lim \mathbb{P}_{(1-\varepsilon)p}(|K_1| \geq \beta |V|) = 0$ for every constant $\beta > 0$. Note that when a percolation threshold exists, it is unique up to multiplication by $1 + o(1)$. So in this sense, we may refer to *the* percolation threshold for \mathcal{G} , when one exists. Now assume that \mathcal{G} has bounded degrees, i.e. there exists $d \in \mathbb{N}$ such that for every $G \in \mathcal{G}$, every vertex in G has degree at most d . By [Eas23], \mathcal{G} always has a percolation threshold, say p . We say that percolation on \mathcal{G} has a *sharp phase transition* if for every constant $\varepsilon > 0$, there exists a constant $A < \infty$ such that

$$\lim \mathbb{P}_{(1-\varepsilon)p}(|K_1| \geq A \log |V|) = 0.$$

Conversely, it is not hard to show in general that $\liminf p \geq \frac{1}{d-1} > 0$ (see [Eas23, Proposition 5]) and that the complementary bound on $|K_1|$ always holds in the sense that if $\liminf(1 - \varepsilon)p > 0$

⁷So a “graph” G is really a graph-isomorphism equivalence class of graphs.

⁸This equation means that for some (and hence every) enumeration $\mathcal{G} = \{G_1, G_2, \dots\}$ where $G_n = (V_n, E_n)$, we have

$$\lim_{n \rightarrow \infty} \mathbb{P}_{(1+\varepsilon)p(G_n)}^{G_n}(|K_1| \geq \alpha |V_n|) = 1.$$

then there exists $A < \infty$ such that $\lim \mathbb{P}_{(1-\varepsilon)p}(|K_1| \geq \frac{1}{A} \log |V|) = 1$ (see Proposition 8.2.6).

Given a transitive graph G , we write o to denote an arbitrary vertex, and we write B_n^G to denote the graph-metric ball of radius n centred at o , viewed as a rooted subgraph of G . We also write $\text{Gr}(n)$ for the number of vertices in B_n^G , and define S_n^G to be the sphere⁹ of radius n . The *local* (aka *Benjamini-Schramm*) topology on the set of all transitive graphs is the metrisable¹⁰ topology with respect to which a sequence (G_n) converges to G if and only if for $r \in \mathbb{N}$, the balls $B_r^{G_n}$ and B_r^G are isomorphic for all sufficiently large n . For example, the sequence of tori $(\mathbb{Z}_n^2)_{n=1}^\infty$ converges *locally* to \mathbb{Z}^2 . Given metric spaces X and Y , the *Gromov-Hausdorff* distance between X and Y , denoted $\text{dist}_{\text{GH}}(X, Y)$, is the infimum over all $\varepsilon > 0$ such that there exists a metric space Z and isometric embeddings $\phi : X \rightarrow Z$ and $\psi : Y \rightarrow Z$ such that the Hausdorff distance between the images of ϕ and ψ in Z is at most ε . Given a graph G and $r > 0$, we write rG for the rescaled graph metric of G where all distances are multiplied by r . For example, the sequence of rescaled tori $(\frac{2\pi}{n} \mathbb{Z}_n^2)_{n=1}^\infty$ Gromov-Hausdorff converges to the continuum torus $S^1 \times S^1$ with the L^1 metric, where S^1 is the unit circle. The scaling limits that arise like this, as a Gromov-Hausdorff limit of a sequence of diameter-rescaled finite transitive graphs, are explored in [BFT17].

The main result of our paper resolves the problems of sharpness and locality for all bounded-degree finite transitive graphs that are not one-dimensional in a certain coarse-geometric sense.

Theorem 8.1.1. *Let \mathcal{G} be an infinite set of finite transitive graphs with bounded degrees. Suppose that there does not exist an infinite subset $\mathcal{H} \subseteq \mathcal{G}$ such that $(\frac{\pi}{\text{diam } G} G)_{G \in \mathcal{H}}$ Gromov-Hausdorff converges to the unit circle. Then both of the following statements hold:*

1. *Percolation on \mathcal{G} has a sharp phase transition.*
2. *If \mathcal{G} converges locally to an infinite transitive graph H , then the constant sequence $p : G \mapsto p_c(H)$ is the percolation threshold for \mathcal{G} .*

In fact, if either of these two statements is false, then there exists an infinite subset $\mathcal{H} \subseteq \mathcal{G}$ such that for every $G \in \mathcal{H}$,

$$\text{dist}_{\text{GH}} \left(\frac{\pi}{\text{diam } G} G, S^1 \right) \leq \frac{e^{(\log \text{diam } G)^{1/9}}}{\text{diam } G}.$$

We interpret the upper bound on Gromov-Hausdorff distance as a bound on the rate of convergence of the large scale geometry of graphs in \mathcal{H} towards the unit circle as their diameters tend to infinity.

⁹We generalise these to non-integer n by setting $B_n^G := B_{\lfloor n \rfloor}^G$ and by defining S_n^G and $\text{Gr}(n)$ analogously.

¹⁰This topology is induced by the metric $\text{dist}(G, H) := \exp(-\max\{n : B_n^G \cong B_n^H\})$, for example.

In this sense, graphs in \mathcal{H} converge to the unit circle faster than do the tori $\{\mathbb{Z}_n \times \mathbb{Z}_{e^{(\log n)^8}}\}_{n \geq 1}$, and in particular faster than do the polynomially-stretched tori $\{\mathbb{Z}_n \times \mathbb{Z}_{n^C}\}_{n \geq 1}$ for any constant C . On the other hand, by using arguments specific to Euclidean tori [EH21a, Example 5.1], items 1 and 2 only fail once we reach exponentially-stretched tori $\{\mathbb{Z}_n \times \mathbb{Z}_{C^n}\}_{n \geq 1}$ for a constant C . So our rate is not sharp in this special case, even if we could improve the exponent $1/9$, which we did not try to optimise. Perhaps these exponentially-stretched tori are worst possible, in which case the optimal bound on the rate should be on the order of $\frac{\log \text{diam } G}{\text{diam } G}$ instead of $\frac{e^{(\log \text{diam } G)^{1/9}}}{\text{diam } G}$. There are also stronger ways that one could hope to describe the “one-dimensionality” of \mathcal{H} . (See the discussion at the end of Section 8.1.)

Previous work and strategy of the proof

Recall from our earlier discussion that sharpness and locality for finite transitive graphs are equivalent. (See Proposition 8.2.9.) In this paper we will prove sharpness directly. At a high level, our idea is to apply arguments derived from the proofs of four existing results in succession: (1) The sharpness of the phase transition for infinite transitive graphs; (2) The locality of the critical point for infinite transitive graphs; (3) The uniqueness of the supercritical giant cluster on finite transitive graphs; (4) The existence of a percolation threshold on finite transitive graphs. Below we discuss each of these works and how they feature in our argument.

To prove sharpness, we will start with a sequence $p : \mathcal{G} \rightarrow (0, 1)$ where \mathbb{P}_p has a cluster larger than a large multiple of $\log |V|$ with good probability. Then given any $\varepsilon > 0$, we will show that $\mathbb{P}_{(1+\varepsilon)p}$ has a giant cluster with high probability. Since $\inf p > 0$, we can replace $\mathbb{P}_{(1+\varepsilon)p}$ by $\mathbb{P}_{p+\varepsilon}$, or equivalently, $\mathbb{P}_{p+4\varepsilon}$. We will split the jump $p \rightarrow p + 4\varepsilon$ into four little hops $p \rightarrow p + \varepsilon \rightarrow \dots \rightarrow p + 4\varepsilon$. After each hop, we will prove something stronger about the connectivity properties of percolation at the current parameter. Each hop is the subject of one section, discussed below, and involves one of the four works listed above.

Large clusters \rightarrow local connections

The sharpness of the phase transition for infinite transitive graphs is the statement that for every infinite transitive graph G and every $p < p_c(G)$, there is constant $c(G, p) > 0$ such that $\mathbb{P}_p^G(|K_o| \geq n) \leq e^{-cn}$ for every $n \geq 1$. Some people use slightly different definitions. For example, some replace this exponential tail on $|K_1|$ by the exponential decay of connection probabilities, which is a priori weaker, and some include the mean-field lower bound $\theta((1 + \varepsilon)p_c(G)) \geq \frac{\varepsilon}{1+\varepsilon}$ as part of the definition. The analogue of the mean-field lower bound for finite transitive graphs was already established in full generality in our earlier work [Eas23, Corollary 4]. This statement of sharpness

for infinite transitive graphs looks similar to our definition of sharpness for an infinite set \mathcal{G} of finite transitive graphs with bounded degrees. Indeed, let p be the percolation threshold for \mathcal{G} . By a simple union bound, if for all $\varepsilon > 0$ there exists $C(\mathcal{G}, \varepsilon) < \infty$ such that $\mathbb{P}_{(1-\varepsilon)p(G)}^G(|K_o| \geq n) \leq Ce^{-n/C}$ for all $n \geq 1$ and $G \in \mathcal{G}$, then percolation on \mathcal{G} has a sharp phase transition. With a little more work (see Proposition 8.2.9), one can show that the converse holds too.

So a natural approach towards proving sharpness for finite transitive graphs is to try to adapt an existing proof of sharpness for infinite transitive graphs. Before explaining what is wrong with this approach, notice that *something* must go wrong because these arguments are completely general, applying to every infinite transitive graph - including \mathbb{Z} , whereas as illustrated by the sequence of stretched tori, some hypothesis on the geometry of finite transitive graphs is required for sharpness to hold. The problem is not that the arguments cannot be run, but rather that they do not address the right question. Roughly speaking, the issue is that a cluster that grows faster than every particular microscopic scale is not automatically macroscopic. Slightly more precisely, given an infinite graph G and parameter p , if $\inf_{n \geq 1} \mathbb{P}_p(|K_o| \geq n) > 0$, then under \mathbb{P}_p there is an infinite cluster almost surely. However, for an infinite set \mathcal{G} of finite graphs and a sequence of parameters p , if $\inf_{n \geq 1} \liminf_{\mathcal{G}} \mathbb{P}_p(|K_o| \geq n) > 0$, then it does not necessarily follow that under \mathbb{P}_p there is a giant cluster with high probability.

While proofs of sharpness for infinite transitive graphs do not directly yield Theorem 8.1.1, our first step is still to adapt and run one of these proofs on finite transitive graphs. We will also apply Hutchcroft's idea [Hut20a] of using his two-ghost inequality to convert point-to-sphere bounds into point-to-point lower bounds. (This was also the first step of [EH23b].) Together, this will establish that after the first hop, $\mathbb{P}_{p+\varepsilon}$ satisfies a point-to-point lower bound on a large constant scale.

In this section we will also use an elementary spanning tree argument to prove a kind of “reverse” implication that if \mathbb{P}_p was instead assumed to satisfy such a point-to-point lower bound, then it would follow that \mathbb{P}_p has a cluster much larger than $\log |V|$ with high probability. This reverse direction is not relevant to proving Theorem 8.1.1, but we will apply it to establish the equivalence of different characterisations of sharpness on finite transitive graphs and in particular to prove the equivalence of items 1 and 2 in Theorem 8.1.1.

Local connections \rightarrow global connections

Earlier we discussed the proof of the locality of the critical point for infinite transitive graphs [EH23b] and the obstructions to using the same argument to prove locality for finite transitive graphs. The primary obstruction was the application of sharpness, which we explained could be

replaced by a good enough quantitative *quasi-sharpness*. If we had organised the argument in the present paper as a direct proof of locality, rather than of sharpness, then this quasi-sharpness would be supplied by the following two hops ($p + 2\epsilon \rightarrow p + 3\epsilon \rightarrow p + 4\epsilon$). For the current hop ($p + \epsilon \rightarrow p + 2\epsilon$), we will run the part of the proof of locality for infinite graphs leading up to this application of sharpness (after dealing with the challenges that we discussed this entails) to propagate the microscopic point-to-point lower bound at $p + \epsilon$ to a global point-to-point lower bound at $p + 2\epsilon$. More precisely, we prove that if \mathcal{G} does not contain a sequence converging rapidly to the unit circle, then for some explicit and slowly-decaying function $f : \mathbb{N} \rightarrow (0, 1)$, all but finitely many graphs $G \in \mathcal{G}$ satisfy

$$\min_{u, v \in V} \mathbb{P}_{p+2\epsilon}(u \leftrightarrow v) \geq f(|V|).$$

Global connections \rightarrow unique large cluster

In the supercritical phase of percolation on a bounded-degree finite transitive graph, there is exactly one giant cluster with high probability. This had been conjectured by Benjamini and was verified in our joint work with Hutchcroft [EH21a]. It is important to note that this result actually does not rely on the existence of a percolation threshold. To make sense of this, we need a definition of the supercritical phase that is agnostic to the existence of a percolation threshold.

Let \mathcal{G} be an infinite set of bounded-degree finite transitive graphs, and let $q : \mathcal{G} \rightarrow (0, 1)$ be a sequence of parameters. If G admits a percolation threshold p , then the natural definition for q being supercritical is that $\liminf q/p > 1$. To make this independent of the existence of p , we say that q is supercritical if there exists a sequence $q' : \mathcal{G} \rightarrow (0, 1)$ and a constant $\epsilon > 0$ such that $\liminf q/q' > 1$ and $\liminf \mathbb{P}_{q'}(|K_1| \geq \epsilon |V|) \geq \epsilon$. In this language, the main result of [EH21a] is that for every supercritical sequence q , the number of vertices $|K_2|$ contained in the second largest cluster satisfies $\lim \mathbb{P}_q(|K_2| \geq \delta |V|) = 0$ for every constant $\delta > 0$.

The argument in [EH21a] is fully quantitative. In particular, if we slightly weaken the hypothesis that q is supercritical by replacing the constant $\epsilon > 0$ in the definition of “ q is supercritical” by a slowly decaying sequence $\epsilon : \mathcal{G} \rightarrow (0, 1)$, then we can still deduce that under \mathbb{P}_q the largest cluster is much larger than all other clusters with high probability. What we need is the same conclusion but with the alternative hypothesis that $\delta := \min_{u, v \in V} \mathbb{P}_{q'}(u \leftrightarrow v)$ tends to zero slowly. This is certainly possible in principle because by Markov’s inequality, the lower bound $\min_{u, v \in V} \mathbb{P}_{q'}(u \leftrightarrow v) \geq 2\epsilon$ always implies the lower bound $\mathbb{P}_{q'}(|K_1| \geq \epsilon |V|) \geq \epsilon$. Unfortunately, this approach ultimately requires that δ tends to zero extremely slowly, too slowly for our purposes. Fortunately, the argument in [EH21a] turns out to run much more efficiently if we directly supply the hypothesis

that $\min_{u,v \in V} \mathbb{P}_{q'}(u \leftrightarrow v) \geq \varepsilon$ rather than the hypothesis that $\mathbb{P}_{q'}(|K_1| \geq \varepsilon |V|) \geq \varepsilon$. Indeed, a significant loss in the proof in [EH21a] is due to the conversion of the latter into the former. We will apply this to deduce from the global point-to-point lower bound at $\mathbb{P}_{p+2\varepsilon}$ that under $\mathbb{P}_{p+3\varepsilon}$, the largest cluster is much larger than all other clusters with high probability. However, note that a priori this largest cluster might *not be giant*, i.e. we may still have $|K_1| = o(|V|)$ with high probability. In particular, our proof is not complete at this stage, which is why we need the fourth hop.

Unique large cluster \rightarrow giant cluster

Every infinite set of finite transitive graphs with bounded degrees \mathcal{G} admits a percolation threshold. We verified this in [Eas23] by combining [EH21a; Van23]. The reader may find it surprising that the uniqueness of the supercritical giant cluster comes first, before the existence of a percolation threshold. Indeed, this is opposite to the order in the classical story for the Erdős-Rényi model, for example. On the other hand, the reader may suspect that the result is obvious because standard sharp threshold techniques imply that for every sequence α , the event $\{|K_1| \geq \alpha |V|\}$ always has a sharp threshold.¹¹ The challenge is to prove that every sequence α that decays sufficiently slowly has a *common* sharp threshold.

To prove this we embedded the fact that the supercritical giant cluster is unique into Vanneuville's proof of the sharpness of the phase transition for infinite transitive graphs. In [Eas23], we did not give any explicit bounds because we were working without the hypothesis that \mathcal{G} has bounded degrees. At this level of generality, there actually exist (very particular) sequences that do not admit a percolation threshold, and even for those that do, the implicit rates of convergence can be arbitrarily bad. However, our argument is itself fully quantitative. In particular, we will explain how it can still be run under an explicit weaker version of the uniqueness of the giant cluster. This will allow us to deduce from the global two-point lower bound under $\mathbb{P}_{p+2\varepsilon}$ and the uniqueness of the largest cluster under $\mathbb{P}_{p+3\varepsilon}$ that there is a giant cluster under $\mathbb{P}_{p+4\varepsilon}$ with high probability, completing our proof of Theorem 8.1.1.

Further discussion

Let us further explore the connection between percolation on finite and infinite transitive graphs. First, let us remark on how to canonically *define* p_c for finite graphs. By [Eas23], there exists a universal function $p_c : \mathcal{F} \rightarrow (0, 1)$, where \mathcal{F} is the set of all finite transitive graphs, such that for every infinite set \mathcal{G} of finite transitive graphs with bounded degrees, the restriction $p_c|_{\mathcal{G}}$ is the

¹¹We say that a sequence of events $(A(G))_{G \in \mathcal{G}}$ has a sharp threshold if there exists a sequence p such that $\limsup q/p < 1$ implies $\mathbb{P}_q(A) = 0$ and $\liminf q/p > 1$ implies $\mathbb{P}_q(A) = 1$ for every sequence q .

percolation threshold for \mathcal{G} . Let us fix such a function p_c for the rest of this section. Now, thanks to Theorem 8.1.1, we can roughly¹² interpret this as the unique continuous extension with respect to the local topology of the usual percolation threshold p_c for infinite transitive graphs to the set of finite transitive graphs.

Let H be an infinite transitive graph, and let \mathcal{G} be an infinite set of finite transitive graphs with bounded degrees that does not contain a sequence approximating the unit circle in the sense that \mathcal{H} does in Theorem 8.1.1. Let (V_n) be an exhaustion of H by finite sets, and let K_∞ denote the set of vertices contained in infinite clusters. By a second-moment calculation, under \mathbb{P}_p^H for any p ,

$$\frac{|K_\infty \cap V_n|}{|V_n|} \rightarrow \theta^H(p) := \mathbb{P}_p^H(o \leftrightarrow \infty)$$

in probability as n tends to infinity. In this sense, $\theta^H(p)$ captures the *density* of the union of the infinite clusters. In a finite graph G , we define the giant density to be $\|K_1\| := \frac{1}{|V|} |K_1|$. In conjunction with the main result of [EH23+a], Theorem 8.1.1 implies that if \mathcal{G} converges locally to H , then for every constant $p \in (0, 1) \setminus \{p_c(H)\}$, the density $\|K_1\|$ under \mathbb{P}_p^G converges in probability to the density $\theta^H(p)$ under \mathbb{P}_p^H as we run through $G \in \mathcal{G}$. In this sense, the infinite cluster phenomenon on infinite transitive graphs is a *good model* for the giant cluster phenomenon on finite transitive graphs. Similar ideas are discussed in Benjamini's original work [Ben01b].

In light of this, our results let us easily move statements about infinite transitive graphs to the setting of finite transitive graphs. Here are three examples. For all three, remember that \mathcal{G} is assumed to be a family of graphs satisfying the hypotheses of Theorem 8.1.1. First, it is well-known that $p_c(H) < 1$ if (and only if) H is not one-dimensional [DGRSY20]. By the conclusion of Theorem 8.1.1, it immediately follows¹³ that $\sup_{G \in \mathcal{G}} p_c(G) < 1$, i.e. there exists $\varepsilon > 0$ such that $\mathbb{P}_{1-\varepsilon}(|K_1| \geq \varepsilon |V|) \geq \varepsilon$ every $G \in \mathcal{G}$. This conclusion is not new; we simply wish to illustrate how easily it follows from Theorem 8.1.1. Indeed, Hutchcroft and Tointon established this fundamental result under a weaker (essentially optimal!) version of the hypothesis that \mathcal{G} is not one-dimensional, (almost) fully resolving a conjecture of Alon, Benjamini, and Stacey [ABS04c]. Second, it is a major open conjecture that $\theta^H(\cdot)$ is continuous if (and only if) H is not one-dimensional. Following the discussion in our previous paragraph, this conjecture would immediately imply the following statement, which says that the giant cluster emerges *gradually*: Let \mathbb{P} denote the law of the standard monotone coupling $(\omega_p : p \in [0, 1])$ of the percolation measures $(\mathbb{P}_p : p \in [0, 1])$, and define

¹²It is unique (up to $o(1)$) and continuous whenever we restrict to an infinite set \mathcal{G} that is compact in the local topology and satisfies $\inf_{G \in \mathcal{G}} \text{dist}_{\text{GH}}\left(\frac{\pi}{\text{diam } G} G, S^1\right) > 0$.

¹³One just needs to verify that \mathcal{G} cannot converge locally to a one-dimensional infinite transitive graph, e.g. by Lemma 8.3.20.

$\alpha(p) := \|K_1(\omega_p)\|$. Then for all $\varepsilon > 0$ there exists $\delta > 0$ such that

$$\lim_{G \in \mathcal{G}} \mathbb{P} \left(\sup_p [\alpha(p + \delta) - \alpha(p)] \leq \varepsilon \right) = 1. \quad (8.1.1)$$

For example, thanks to [Hut16], we can already deduce from this relation between infinite and finite transitive graphs that eq. (8.1.1) holds whenever our family \mathcal{G} has exponential growth on microscopic scales, e.g. for the sequence $(\mathbb{Z}_n^3 \times G_{h(n)})$ where (G_n) is a sequence of transitive expanders and $h : \mathbb{N} \rightarrow \mathbb{N}$ tends to infinity arbitrarily slowly. (See the discussion in [EH23b, Example 5.1] of stretched tori, in which the giant cluster does not emerge gradually.) By the uniqueness of the supercritical giant cluster [EH21a], finite transitive graphs satisfying eq. (8.1.1) also automatically satisfy the conclusion of [ABS04a, Conjecture 1.1]. This links the well-known continuity conjecture for infinite transitive graphs to this conjecture about the uniqueness of the largest cluster in finite transitive graphs. Third, it is conjectured that the uniqueness threshold $p_u(H)$ satisfies $p_c(H) < p_u(H)$ if and only if H is nonamenable. Again by our discussion in previous paragraph, this conjecture would imply that if the Cheeger constant on graphs in \mathcal{G} is uniformly bounded below on microscopic scales, then percolation on \mathcal{G} has a phase in which there is a giant cluster whose metric distortion tends to infinity. (See [EH21a, Remark 1.6].) What can be said when the Cheeger constant is uniformly bounded below on larger scales? In the limit, this connects the p_c vs p_u question to the existing theory of percolation on expanders.

This opens the door to many directions for future work, adapting questions and techniques from percolation on infinite transitive graphs to finite transitive graphs. For example, what can be said about supercritical sharpness? Since the continuity conjecture for infinite transitive graphs would imply the unique giant cluster conjecture of [ABS04a, Conjecture 1.1] (possibly with a weaker one-dimensionality condition), might [ABS04a, Conjecture 1.1] be a stepping stone towards continuity that is easier to establish? Another direction for future work is to improve the rate of convergence in Theorem 8.1.1. One could also explore stronger notions of one-dimensionality. The Gromov-Hausdorff metric only considers the coarse geometry of graphs, ignoring how densely vertices are packed (i.e. the volume growth on small scales). It is natural to expect that graphs in which vertices are packed more densely can afford to have a more one-dimensional coarse geometry before percolation arguments break down. For example, consider the product of a torus with a long cycle versus the product of an expander with a long cycle. In the work of Hutchcroft and Tointon [HT21a], one-dimensionality was characterised more stringently¹⁴ in terms of the relationship of volume to diameter, for example, by requiring that $|V| \leq (\text{diam } G)^{1+\varepsilon}$ or $\frac{|V|}{\log|V|} = o(\text{diam } G)$. One

¹⁴This is indeed stronger than asking for Gromov-Hausdorff convergence to the unit circle, by the results of [BFT17].

could also investigate questions such as sharpness without bounded degrees. For example, [EH21a; Eas23; EH23+a] did not require this hypothesis, thus linking the story of percolation on infinite transitive graphs to the classical Erdős-Rényi model.

Acknowledgement

We thank Tom Hutchcroft for helpful comments on an earlier draft.

8.2 Large clusters → local connections

In this section we will adapt a proof of the sharpness of the phase transition for infinite transitive graphs to finite transitive graphs. By combining this with an idea of Hutchcroft [Hut20a] to convert volume-tail bounds into point-to-point bounds, we will prove the following proposition. This roughly says that if for some percolation parameter p , the largest cluster contains much more than $\log |V|$ vertices with good probability, then for percolation of any higher parameter $p + \eta$, we have a uniform point-to-point lower bound on a divergently large scale. Later, in Section 8.2 we will prove a kind of converse to this statement, and in Section 8.2 we will use this converse to give equivalent characterisations of sharpness for finite transitive graphs.

Proposition 8.2.1. *Let G be a finite transitive graph with degree d . Let $\eta > 0$. There exists $c(d, \eta) > 0$ such that for all $p \in (0, 1)$ and $\lambda \geq 1$,*

$$\mathbb{P}_p(|K_1| \geq \lambda \log |V|) \geq \frac{1}{c |V|^c} \implies \min_{u \in B_{c \log(\lambda) - \frac{1}{c}}} \mathbb{P}_{p+\eta}(o \leftrightarrow u) \geq \frac{\eta^2}{20}.$$

We have chosen to adapt Vanneuville's recent proof of sharpness for infinite graphs [Van24]. This involves ghost fields. Given a graph G , a ghost field of intensity $q \in (0, 1)$ is a random set of vertices $g \subseteq V$ distributed according to (Bernoulli) site percolation of parameter q .¹⁵ We denote its law by \mathbb{Q}_q and write $\mathbb{P}_p \otimes \mathbb{Q}_q$ for the joint law of independent samples $\omega \sim \mathbb{P}_p$ and $g \sim \mathbb{Q}_q$. One reason to introduce ghost fields is that it can be easier to work with the event $\{o \leftrightarrow g\}$ when $q = 1/n$ than to work with the closely related event $\{|K_o| \geq n\}$.

The following is [Van24, Theorem 2]. This can also be deduced from [Hut20d] with different constants. This says that starting from any percolation parameter p , if we decrease p by a suitable amount, then the volume of the cluster at the origin will have an exponential tail under the new parameter. This is proved by a variant of Vanneuville's stochastic comparison technique from [Van23], which we will describe in more detail in Section 8.5.

¹⁵Some authors use a slightly different parameterisation. When we write “a ghost field of intensity $q \in (0, 1)$ ”, they write “a ghost field of intensity $h > 0$ ” for the same object, where $q = 1 - e^{-h}$.

Lemma 8.2.2. *Let G be a transitive graph. Given $p \in (0, 1)$ and $h > 0$, define*

$$\mu_{p,h} := \mathbb{P}_p \otimes \mathbb{Q}_{1-e^{-h}}(o \xrightarrow{\omega} g).$$

Then for all $m \geq 1$,

$$\mathbb{P}_{(1-\mu_{p,h})p}(|K_o| \geq m) \leq \frac{\mathbb{P}_p(|K_o| \geq m)}{1 - \mu_{p,h}} e^{-hm}.$$

Vanneuville proved this lemma when G is infinite, but his proof also works verbatim when G is finite. The following easy corollary of this lemma says (contrapositively) that if the cluster at the origin is much larger than $\log |V|$ with reasonable probability, then after sprinkling, the cluster at the origin is at least mesoscopic with good probability.

Corollary 8.2.3. *Let G be a finite transitive graph. For all $\varepsilon > 0$ there exists $c(\varepsilon) > 0$ such that for all $p \in (0, 1)$ and $n, \lambda \geq 1$,*

$$\mathbb{P}_p(|K_o| \geq n) \leq \varepsilon \implies \mathbb{P}_{(1-2\varepsilon)p}(|K_o| \geq \lambda \log |V|) \leq \frac{1}{c |V|^{\frac{c\lambda}{n}}}.$$

Proof. Suppose that $\mathbb{P}_p(|K_o| \geq n) \leq \varepsilon$. We may assume that $\varepsilon < 1/2$, otherwise the result is trivial. Define $h := \frac{1}{n} \log \frac{1}{1-\varepsilon}$ and $q := 1 - e^{-h}$. By a union bound,

$$\mu_{p,h} := \mathbb{P}_p \otimes \mathbb{Q}_q(o \xrightarrow{\omega} g) \leq \mathbb{P}_p(|K_o| \geq n) + \mathbb{P}_p \otimes \mathbb{Q}_q(o \xrightarrow{\omega} g \mid |K_o| < n).$$

We now bound these two terms individually. By hypothesis, $\mathbb{P}_p(|K_o| \geq n) \leq \varepsilon$. By our choice of h and q ,

$$\mathbb{P}_p \otimes \mathbb{Q}_q(o \xrightarrow{\omega} g \mid |K_o| < n) \leq 1 - e^{-hn} = \varepsilon.$$

Therefore $\mu_{p,h} \leq 2\varepsilon$. So by Lemma 8.2.2,

$$\mathbb{P}_{(1-2\varepsilon)p}(|K_o| \geq \lambda \log |V|) \leq \frac{\mathbb{P}_p(|K_o| \geq \lambda \log |V|)}{1 - 2\varepsilon} e^{-h\lambda \log |V|} \leq \frac{1}{1 - 2\varepsilon} |V|^{-\frac{\lambda}{n} \log \frac{1}{1-\varepsilon}}.$$

So the claim holds with $c := \min\{1 - 2\varepsilon, \log \frac{1}{1-\varepsilon}\}$. \square

To convert the fact that the cluster at the origin is at least mesoscopic with good probability into a uniform point-to-point lower bound on a divergently large scale, we will apply Hutchcroft's volumetric two-arm bound [Hut20a, Corollary 1.7], stated below as Theorem 8.2.4. This applies in our setting because every finite transitive graph is unimodular, and in this case we can trivially drop the hypothesis that at least one of the clusters is finite in the definition of $\mathcal{T}_{\varepsilon,n}$. This tells us that it is always unlikely that the endpoints of a given edge belong to distinct large clusters.

Theorem 8.2.4. *Let G be a unimodular transitive graph with degree d . There exists $C(d) < \infty$ such that for all $e \in E$, $n \geq 1$, and $p \in (0, 1)$*

$$\mathbb{P}_p(\mathcal{T}_{e,n}) \leq C \left[\frac{1-p}{pn} \right]^{1/2},$$

where $\mathcal{T}_{e,n}$ is the event that the endpoints of e belong to distinct clusters, each of which contains at least n vertices, and at least one of which is finite.

Hutchcroft showed in [Hut20a] that this can be used to convert volume-tail bounds into point-to-point bounds. This was also used in [EH23b]. Here is the quantitative output of his argument, stated in the case of finite graphs.

Corollary 8.2.5. *Let G be a finite transitive graph with degree d . There exists $C(d) < \infty$ such that for all $n, r \geq 1$ and $p \in (0, 1)$,*

$$\min_{u \in B_r} \mathbb{P}_p(o \leftrightarrow u) \geq \mathbb{P}_p(|K_o| \geq n)^2 - \frac{Cr}{p^{r+1}n^{1/2}}.$$

Proof. Let $u \in B_r$. By Harris' inequality and a union bound,

$$\mathbb{P}_p(o \leftrightarrow u) \geq \mathbb{P}_p(|K_o| \geq n)^2 - \mathbb{P}_p(|K_o| \geq n \text{ and } |K_u| \geq n \text{ but } o \leftrightarrow u).$$

The second term on the right can now be bounded by [EH23b, Lemma 2.6]. (In that lemma the hypothesis that G is infinite and $p < p_c$ can be replaced by the hypothesis that G is finite.) \square

We now combine Corollary 8.2.3 and Corollary 8.2.5 to establish Proposition 8.2.1.

Proof of Proposition 8.2.1. Fix $p \in (0, 1)$, $\lambda \geq 1$, and $\eta > 0$. Let $\varepsilon := \frac{\eta}{4}$ and let $c_1(\frac{\eta}{4}) > 0$ be the corresponding constant from Corollary 8.2.3. We may assume that $\eta \leq \frac{1}{2}$, $c_1 < 1$, and $|V| > 1$. Suppose that $\mathbb{P}_p(|K_1| \geq \lambda \log |V|) \geq \frac{1}{c|V|^\varepsilon}$. By a union bound, $\mathbb{P}_p(|K_o| \geq \lambda \log |V|) \geq \frac{1}{c|V|^{c+1}}$. Let $n := \frac{c_1 \lambda}{2}$. Since $\frac{c_1 \lambda}{n} - 1 > c_1$ and $(1 - 2\varepsilon)(p + \eta) \geq p$, it follows by Corollary 8.2.3 that

$$\mathbb{P}_{p+\eta} \left(|K_o| \geq \frac{c_1 \lambda}{2} \right) \geq \frac{\eta}{4}.$$

Let $C_1(d) < \infty$ be the constant from Corollary 8.2.5. Let $r \geq 1$ be arbitrary. Then by Corollary 8.2.5,

$$\min_{u \in B_r} \mathbb{P}_{p+\eta}(o \leftrightarrow u) \geq \left(\frac{\eta}{4} \right)^2 - \frac{C_1 r}{(p + \eta)^{r+1} \left(\frac{c_1 \lambda}{2} \right)^{1/2}}.$$

Note that $r \leq \eta^{-r}$ because $\eta \leq \frac{1}{2}$. So there is a constant $C_2(d, \eta) < \infty$ such that

$$\frac{C_1 r}{(p + \eta)^{r+1} \left(\frac{c_1 \lambda}{2}\right)^{1/2}} \leq \frac{C_2}{\eta^{2r} \lambda^{1/2}}.$$

Now there exists $c_2(d, \eta) > 0$ such that $r := c_2 \log(\lambda) - \frac{1}{c_2}$ satisfies $\frac{C_2}{\eta^{2r} \lambda^{1/2}} \leq \frac{\eta^2}{80}$. Then by our above work (when $r \geq 1$, otherwise the inequality anyway holds trivially),

$$\min_{u \in B_r} \mathbb{P}_{p+\eta}(o \leftrightarrow u) \geq \left(\frac{\eta}{4}\right)^2 - \frac{\eta^2}{80} = \frac{\eta^2}{20}.$$

Therefore the claim holds with $c := \min\{c_1, c_2\}$. \square

Local connections \rightarrow large clusters

In this subsection we prove the following proposition. This implies that if there is a uniform point-to-point lower bound on a divergently large scale, then the largest cluster contains much more than $\log |V|$ vertices with high probability. This is a kind of converse to Proposition 8.2.1. This will be used in the next subsection to prove the equivalence of different notions of sharpness.

Proposition 8.2.6. *Let G be a finite transitive graph. For all $\delta > 0$ there exists $c(\delta) > 0$ such that for all $p \in (0, 1)$ and $r \geq 1$ with $|B_r| \leq |V|^{1/10}$,*

$$\min_{u \in B_r} \mathbb{P}_p(o \leftrightarrow u) \geq \delta \implies \mathbb{P}_p(|K_1| \geq c |B_r| \log |V|) \geq 1 - \frac{1}{|V|^{3/4}}.$$

The next lemma converts point-to-point connection lower bounds on one scale into volume-tail lower bounds on all scales. The idea is to approximately cover the graph by a large number of balls on which the point-to-point lower bound holds then glue together large clusters from multiple balls.

Lemma 8.2.7. *Let G be a finite transitive graph. For all $\delta > 0$ there exists $c(\delta) > 0$ such that for all $p \in (0, 1)$ and $n, r \geq 1$ satisfying $n \leq \frac{c|V|}{|B_r|}$,*

$$\min_{u \in B_r} \mathbb{P}_p(o \leftrightarrow u) \geq \delta \implies \mathbb{P}_p(|K_o| \geq n) \geq c e^{-\frac{n}{c|B_r|}}.$$

Proof. Fix $\delta > 0$, $p \in (0, 1)$, and $n, r \geq 1$. Suppose that $\min_{u \in B_r} \mathbb{P}_p(o \leftrightarrow u) \geq \delta$. Let W be a maximal (with respect to inclusion) set of vertices such that $o \in W$ and $\text{dist}_G(u, v) \geq 2r$ for all distinct $u, v \in W$. Build a graph H with vertex set W by including the edge $\{u, v\}$ if and only if $\text{dist}_G(u, v) \leq 5r$ and $u \neq v$. By maximality of W , the graph H is connected. Let T be a spanning tree for H . Let $f : W \setminus \{o\} \rightarrow W$ be a function encoding T where ‘ $f(u) = v$ ’ means that the edge $\{u, v\}$ is present in T and $\text{dist}_T(o, v) < \text{dist}_T(o, u)$. Extend this to a function $f : W \rightarrow W$ by setting $f(o) := o$. By Markov’s inequality, every $u \in V$ satisfies $\mathbb{P}_p(|K_u \cap B_r(u)| \geq \frac{\delta}{2} |B_r|) \geq \frac{\delta}{2}$.¹⁶ By

¹⁶In this proof, \mathbb{P}_p and $|B_r|$ refer to G , not to T or H .

Harris' inequality, every edge $\{u, v\}$ in H satisfies $\mathbb{P}_p(u \leftrightarrow v) \geq \delta^5$. So by Harris' inequality again, for every $u \in W$, the event A_u that $u \leftrightarrow f(u)$ and $|K_u \cap B_r(u)| \geq \frac{\delta}{2} |B_r|$ satisfies $\mathbb{P}_p(A_u) \geq \delta^5 \cdot \frac{\delta}{2} = \frac{\delta^6}{2}$.

Let $c(\delta) > 0$ be a small constant to be determined. Suppose that $n \leq \frac{c|V|}{|B_r|}$. By maximality of W , the balls $\{B_{2r}(u) : u \in W\}$ cover V , and hence $|V| \leq |W| \cdot |B_{2r}|$. So provided that c is sufficiently small,

$$|W| \geq \frac{|V|}{|B_{2r}|} \geq \frac{|V|}{|B_r|^2} \geq \frac{n}{c|B_r|} \geq \frac{2n}{\delta|B_r|}.$$

In particular, we can find a T -connected set of vertices $U \subseteq W$ such that $o \in U$ and $|U| = \left\lceil \frac{2n}{\delta|B_r|} \right\rceil$.

If A_u holds for every $u \in U$ then $|K_o| \geq \frac{\delta}{2} \cdot |B_r| \cdot |U| \geq n$. So by Harris' inequality, provided c is sufficiently small,

$$\mathbb{P}_p(|K_o| \geq n) \geq \mathbb{P}_p\left(\bigcap_{u \in U} A_u\right) \geq \left(\frac{\delta^6}{2}\right)^{\left\lceil \frac{2n}{\delta|B_r|} \right\rceil} \geq ce^{-\frac{n}{c|B_r|}}. \quad \square$$

The following is a second-moment calculation for the number of vertices contained in large clusters.¹⁷ In the proof, it will be convenient to introduce partial functions to encode partially-revealed percolation configurations. Recall that a *partial function* $f : A \rightharpoonup B$ is a function $A' \rightarrow B$ for some $A' \subseteq A$, i.e. for every $a \in A$, either $f(a) \in B$ or $f(a) = \text{'undefined'}$. We denote this set A' on which f is defined by $\text{dom}(f)$. Given partial functions f and g , the *override* $f \sqcup g$ is the partial function with $\text{dom}(f \sqcup g) = \text{dom}(f) \cup \text{dom}(g)$ that is equal to f on $\text{dom}(f)$ and is equal to g on $\text{dom}(g) \setminus \text{dom}(f)$. We write Var_p to denote the variance of a random variable under \mathbb{P}_p .

Lemma 8.2.8. *Let G be a finite transitive graph. For all $n \geq 0$ and $p \in (0, 1)$, the random set $X := \{u \in V : |K_u| \geq n\}$ satisfies*

$$\text{Var}_p |X| \leq n^2 \cdot \mathbb{E}_p |X|.$$

Proof. Let \mathbb{P} be the joint law of a uniformly random automorphism of G , denoted ϕ , and three configurations $\omega_1, \omega_2, \omega_3$ sampled according to \mathbb{P}_p , where all four of these random variables are independent. Given a configuration $\omega : E \rightarrow \{0, 1\}$, let $\hat{\omega} : E \rightharpoonup \{0, 1\}$ be the partial function encoding the edges revealed in an exploration of the cluster at o from inside (with respect to an arbitrary fixed ordering of E) that is halted as soon as the event $\{|K_o(\omega)| \geq n\}$ is determined by the states of the revealed edges. Define $\omega := (\hat{\omega}_1 \sqcup \phi(\hat{\omega}_2)) \sqcup \omega_3$. By transitivity, the law of $\phi(o)$

¹⁷We were inspired by a weaker (degree-dependent) version of this argument that arose during joint work with Hutchcroft towards [EH23+a], which was made redundant and thus did not appear in the final version of that work.

is uniform on V , and by a standard cluster-exploration argument, $\mathbb{P}(\omega = \cdot \mid \phi) = \mathbb{P}_p$ almost surely. This lets us rewrite $\text{Var}_p |X|$ as

$$\begin{aligned}\text{Var}_p |X| &= \sum_{u,v} [\mathbb{P}_p(u, v \in X) - \mathbb{P}_p(u \in X) \cdot \mathbb{P}_p(v \in X)] \\ &= |V| \mathbb{E}_p |X| \cdot \left[\frac{1}{|V|} \sum_u \mathbb{P}_p(u \in X \mid o \in X) - \mathbb{P}_p(o \in X) \right] \\ &= |V| \mathbb{E}_p |X| \cdot [\mathbb{P}(\phi(o) \in X(\omega) \mid o \in X(\omega)) - \mathbb{P}_p(o \in X)].\end{aligned}\tag{8.2.1}$$

Consider the sets of vertices $A_1 := K_o(\hat{\omega}_1)$ and $A_2 := K_o(\hat{\omega}_2)$, which are defined purely in terms of the open edges in $\hat{\omega}_1$ and $\hat{\omega}_2$ respectively, i.e. all edges with ‘undefined’ state are treated as closed. Note that $o \in X(\omega)$ if and only if $o \in X(\omega_1)$. Moreover, given that $o \in X(\omega_1)$, if $\phi(o) \in X(\omega)$ then either $o \in X(\omega_2)$ or $A_1 \cap \phi(A_2) \neq \emptyset$. So by a union bound and independence,

$$\mathbb{P}(\phi(o) \in X(\omega) \mid o \in X(\omega)) \leq \mathbb{P}_p(o \in X) + \mathbb{P}(A_1 \cap \phi(A_2) \neq \emptyset \mid o \in X(\omega_1)).\tag{8.2.2}$$

In particular, by eq. (8.2.1), it suffices to verify that

$$\mathbb{P}(A_1 \cap \phi(A_2) \neq \emptyset \mid \omega_1, \omega_2, \omega_3) \leq \frac{n^2}{|V|} \quad \text{a.s.}\tag{8.2.3}$$

Consider arbitrary deterministic sets of vertices B_1 and B_2 . By transitivity, the law of $\phi(u)$ for any fixed vertex u is uniform over V . So by a union bound,

$$\mathbb{P}(B_1 \cap \phi(B_2) \neq \emptyset) \leq \sum_{u \in B_2} \mathbb{P}(\phi(u) \in B_1) = \sum_{u \in B_2} \frac{|B_1|}{|V|} = \frac{|B_1| |B_2|}{|V|}.$$

Equation (8.2.3) now follows by applying this to the sets A_1 and A_2 , which almost surely satisfy $|A_1|, |A_2| \leq n$. \square

We now combine Lemmas 8.2.7 and 8.2.8 to prove Proposition 8.2.6.

Proof of Proposition 8.2.6. Let $\delta > 0$, $p \in (0, 1)$, and $r \geq 1$. Suppose that $|B_r| \leq |V|^{1/10}$ and $\min_{u \in B_r} \mathbb{P}_p(o \leftrightarrow u) \geq \delta$. Let $c_1(\delta) > 0$ be the constant from Lemma 8.2.7. Let $n := c |B_r| \log |V|$ for a small constant $c(\delta) > 0$ to be determined. Since $|B_r|^2 \leq |V|^{2/10}$, provided c is small,

$$n := c |B_r| \log |V| \leq \frac{c_1 |V|}{|B_r|}.$$

So Lemma 8.2.7 yields

$$\mathbb{P}_p(|K_o| \geq n) \geq c_1 e^{-\frac{n}{c_1 |B_r|}} = c_1 |V|^{-\frac{c}{c_1}} \geq c_1 |V|^{-1/100},$$

provided c is small. By transitivity, it follows that the random set $X := \{u \in V : |K_o| \geq n\}$ satisfies $\mathbb{E}_p |X| \geq c_1 |V|^{99/100}$. So by Chebychev's inequality and Lemma 8.2.8,

$$\mathbb{P}_p(|K_1| \geq n) = 1 - \mathbb{P}_p(|X| = 0) \geq 1 - \frac{\text{Var}_p |X|}{(\mathbb{E}_p |X|)^2} \geq 1 - \frac{n^2}{c_1 |V|^{99/100}}.$$

The conclusion follows because, provided c is small,

$$\frac{n^2}{c_1 |V|^{99/100}} = \frac{(c |B_r| \log |V|)^2}{c_1 |V|^{99/100}} \leq \frac{\left(c |V|^{1/10} \log |V|\right)^2}{c_1 |V|^{99/100}} \leq \frac{1}{|V|^{3/4}}. \quad \square$$

Equivalent notions of sharpness

In this subsection we apply results from earlier in Section 8.2 to prove the following proposition. In the statement and the proof, we take for granted that \mathcal{G} always admits a percolation threshold [Eas23]. Item 2 is analogous to the standard definition of sharpness for percolation on an infinite transitive graph. The fact that items 1 and 2 are equivalent is why we decided to label our version of “sharpness” for finite transitive graphs as such. Item 3 is analogous to the locality of the critical parameter for infinite transitive graphs. It is perhaps surprising that sharpness and locality are equivalent for finite graphs but not for infinite graphs. One way to make sense of this is that locality for infinite graphs is equivalent to a *uniform* (in the choice of graph) version of sharpness for infinite graphs, and for finite graphs, the only meaningful notion of sharpness is necessarily uniform.

Proposition 8.2.9. *For every infinite set \mathcal{G} of finite transitive graphs with bounded degrees, the following are equivalent:*

1. *Percolation on \mathcal{G} has a sharp phase transition.*
2. *For every subcritical sequence of parameters p , there exists a constant $C(\mathcal{G}, p) < \infty$ such that for all $G \in \mathcal{G}$ and all $n \geq 1$,*

$$\mathbb{P}_p (|K_o| \geq n) \leq C e^{-n/C}.$$

3. *If an infinite subset $\mathcal{H} \subseteq \mathcal{G}$ converges locally to an infinite transitive graph H , then the constant sequence $G \mapsto p_c(H)$ is the percolation threshold for \mathcal{H} .*

We will prove that $3 \implies 2 \implies 1 \implies 3$. For the first step, we apply Corollary 8.2.3 and compactness.

Proof that item 3 implies item 2. Assume that item 3 holds. Our goal is to prove that item 2 holds. Since \mathcal{G} has bounded degrees, \mathcal{G} is relatively compact in the local topology. In particular, we may assume without loss of generality that \mathcal{G} converges locally to some infinite transitive graph G . (If item 2 is false, then we can find an infinite subset $\mathcal{H} \subseteq \mathcal{G}$ such that item 2 is false for every sequence in \mathcal{H} .) Now fix a subcritical sequence p for \mathcal{G} . By item 3, after passing to a tail of \mathcal{G} if necessary, there exists a constant $\varepsilon > 0$ such that $p(G) \leq (1 - \varepsilon)p_c(H)$ for every $G \in \mathcal{G}$. Pick $r \geq 1$ such that $\mathbb{P}_{(1-\varepsilon/2)p_c(H)}^H(|K_o| \geq r) \leq \varepsilon/4$. By passing to a further tail of \mathcal{G} if necessary, we may assume that $B_r^G \cong B_r^H$ for every $G \in \mathcal{G}$. Then $\mathbb{P}_{(1-\varepsilon/2)p_c(H)}^G(|K_o| \geq r) \leq \varepsilon/4$ for every $G \in \mathcal{G}$. Let $c(\varepsilon) > 0$ be the constant from Corollary 8.2.3. For every $n \geq 1$ and $G \in \mathcal{G}$, Corollary 8.2.3 with $\lambda := \frac{n}{\log|V|}$ tells us that

$$\mathbb{P}_{(1-\varepsilon)p_c(H)}^G(|K_o| \geq n) \leq \mathbb{P}_{(1-\varepsilon/2)(1-2\cdot\varepsilon/4)p_c(H)}^G(|K_o| \geq n) \leq \frac{1}{ce^{cn/r}}.$$

Take $C := r/c$. The conclusion now follows by monotonicity because $p(G) \leq (1 - \varepsilon)p_c(H)$ for every $G \in \mathcal{G}$. \square

The second step is a simple union bound.

Proof that item 2 implies item 1. Given a subcritical sequence p , let $C(\mathcal{G}, p) < \infty$ be the constant guaranteed to exist by item 2. Then for every $G \in \mathcal{G}$,

$$\mathbb{P}_p(|K_o| \geq 2C \log |V|) \leq Ce^{-\frac{2C \log |V|}{C}} = \frac{C}{|V|^2},$$

and hence by a union bound,

$$\mathbb{P}_p(|K_1| \leq 2C \log |V|) \geq 1 - |V| \frac{C}{|V|^2} = 1 - \frac{C}{|V|}.$$

So $\lim \mathbb{P}_p(|K_1| \leq 2C \log |V|) = 1$, as required. \square

We now turn to the third step, $1 \implies 3$. Fix a choice of percolation threshold $p_c : \mathcal{G} \rightarrow (0, 1)$, and think of this as an extension of the usual critical points p_c for percolation on the infinite transitive graphs that make up the boundary of \mathcal{G} . Then our goal is to show that, assuming item 1, the function p_c is continuous as we approach the boundary of \mathcal{G} from the interior. We split this into two parts: upper- and lower-semicontinuity. For lower-semicontinuity, we will apply a finite graph version of an argument of Pete [Pet, Section 14.2], which was based on the mean-field lower bound for infinite transitive graphs. For upper-semicontinuity, we will combine Corollary 8.2.5 and Proposition 8.2.6.

Proof that item 1 implies item 3. Suppose for contradiction that $\mathcal{H} \subseteq G$ is an infinite subset that converges locally to some infinite transitive graph H , but the constant sequence $G \mapsto p_c(G)$ is not a percolation threshold for \mathcal{H} . By passing to a subsequence, we may assume without loss of generality that there is a constant $\varepsilon > 0$ such that either $p_c(G) \leq (1 - \varepsilon)p_c(H)$ for every $G \in \mathcal{H}$, or $p_c(G) \geq (1 + \varepsilon)p_c(H)$ for every $G \in \mathcal{H}$. Call these Case 1 and Case 2, corresponding to (a violation of) lower- and upper-semicontinuity in our discussion above.

(Case 1) Since $p_c : \mathcal{G} \rightarrow (0, 1)$ is a percolation threshold, there exists a constant $\delta > 0$ such that $\mathbb{P}_{(1+\varepsilon)p_c(G)}^G(|K_o| \geq \delta |V|) \geq \delta$ for every $G \in \mathcal{H}$. So by monotonicity, $\mathbb{P}_{(1-\varepsilon^2)p_c(H)}^G(|K_o| \geq \delta |V|) \geq \delta$ for every $G \in \mathcal{H}$. For every $r \geq 1$, there exists $G \in \mathcal{H}$ such that $\delta |V| > |B_r^H|$ and $B_r^G \cong B_r^H$, and hence

$$\mathbb{P}_{(1-\varepsilon^2)p_c(H)}^H(|K_o| \geq r) \geq \mathbb{P}_{(1-\varepsilon^2)p_c(H)}^G(o \leftrightarrow S_r) \geq \mathbb{P}_{(1-\varepsilon^2)p_c(H)}^G(|K_o| \geq \delta |V|) \geq \delta.$$

In particular, $\mathbb{P}_{(1-\varepsilon^2)p_c(H)}^H(|K_o| = \infty) > 0$, a contradiction.

(Case 2) Let $A \geq 1$ be a given arbitrary constant. It suffices to prove that the parameter $p := (1 + \varepsilon/2)p_c(H)$ satisfies $\lim_{\mathcal{G}} \mathbb{P}_p^G(|K_1| \geq A \log |V|) = 1$. Set $\delta := \mathbb{P}_p^H(o \leftrightarrow \infty) > 0$. Let d be the vertex degree of H , and note that $p_c(H) > 1/d$, as this is well-known to hold for every infinite transitive graph. So by Corollary 8.2.5, there is a constant $C(d) < \infty$ such that for all $n, r \geq 1$ and all $G \in \mathcal{H}$ with $B_n^G \cong B_n^H$,

$$\min_{u \in B_r^G} \mathbb{P}_p^G(o \leftrightarrow u) \geq \delta^2 - \frac{Crd^{r+1}}{n^{1/2}},$$

and in particular, (using that $r \leq d^r$ for all $r \geq 1$) the radius $r(n) := \log_d \left(\frac{\delta^2 n^{1/2}}{2Cd^2} \right)$ satisfies

$$\min_{u \in B_{r(n)}^G} \mathbb{P}_p^G(o \leftrightarrow u) \geq \delta^2 - \frac{\delta^2}{2} = \frac{\delta^2}{2}.$$

Let $c(\delta^2/2) > 0$ be the constant from Proposition 8.2.6. Fix n sufficiently large that $c \cdot r(n) \geq A$. By passing to a tail of \mathcal{H} if necessary, let us assume that $B_n^G \cong B_n^H$ and $|B_{r(n)}^G| \leq |V|^{1/10}$ for every $G \in \mathcal{H}$. Then by Proposition 8.2.6, for all $G \in \mathcal{H}$,

$$\mathbb{P}_p^G \left(|K_1| \geq c |B_{r(n)}^G| \log |V| \right) \geq 1 - \frac{1}{|V|^{3/4}}.$$

In particular, since $c |B_{r(n)}^H| \geq cr(n) \geq A$, we deduce that $\lim_{\mathcal{G}} \mathbb{P}_p^G(|K_1| \geq A \log |V|) = 1$ as required. \square

8.3 Local connections → global connections

In this section, we will apply the proof from [EH23b] that the critical point for percolation on (non-one-dimensional) infinite transitive graphs is local. As explained in the introduction, we need to both make this argument more finitary and adapt it to finite transitive graphs. We can roughly think of the proof of [EH23b] in two parts: First, if G does not satisfy certain geometric properties around scale n , which include that G is finitely-ended, then G must satisfy a certain statement \mathcal{I}_n about the propagation of connection bounds around scale n . Second, if G does satisfy these geometric properties and G is one-ended, then G again satisfies \mathcal{I}_n . Together, these two parts imply that if G does not satisfy \mathcal{I}_n , then G must actually be two-ended and hence one-dimensional. By looking at the proof of the second part, we can pinpoint where one-endedness is used, namely as a hypothesis in [EH23b, Lemma 5.8].

[EH23b, Lemma 5.8] concerns certain (o, ∞) -cutsets called *exposed spheres*. The lemma says that if G satisfies nice geometric properties around scale n and is one-ended, then the exposed spheres around scale n are in some sense well-connected. We took this from [CMT22, Lemma 2.1 and 2.7], where the authors deduced it from a theorem of Babson and Benjamini [BB99b]. By reading Timar's proof [Tim07] of this theorem of Benjamini and Babson, we see that if an exposed sphere is not well-connected, then not only is G multiply-ended, but this is actually witnessed by the exposed sphere itself in the sense that its removal from G would create multiple infinite components. From this we can conclude that G must in fact start to look one-dimensional from around scale n . This is how we will make this step from [EH23b] finitary. To adapt the argument to finite transitive graphs, we will additionally need to introduce the notion of the exposed sphere in a finite transitive graph and prove that finite transitive graphs can, for the purpose of part of our argument, be treated like infinite transitive graphs that are one-ended.

Unfortunately, this application of Babson-Benjamini is deeply embedded in the proof of [EH23b] as it is currently written. So it will take some work to restructure the multi-scale induction in [EH23b] to isolate the relevant part. To avoid repetition, we have deferred the details of arguments that are implicit in [EH23b] to the appendix, thereby keeping many of the arguments in this section high-level. Ultimately we will prove the following proposition, which contains this finite-graph finitary refinement of locality. While we have written this for finite graphs, the same argument yields the analogous finitary refinement for infinite graphs.

Proposition 8.3.1. *Let G be a finite transitive graph with degree d . Define*

$$\gamma := \text{dist}_{\text{GH}} \left(\frac{\pi}{\text{diam } G} G, S^1 \right) \cdot \text{diam } G \quad \text{and} \quad \gamma^+ := e^{(\log \gamma)^9}.$$

For all $\varepsilon, \eta > 0$ there exists $\lambda(d, \varepsilon, \eta) < \infty$ such that for all $p \in (0, 1)$,

$$\min_{u \in B_\lambda} \mathbb{P}_p(o \leftrightarrow u) \geq \eta \implies \min_{u \in B_{\gamma^+}} \mathbb{P}_{p+\varepsilon}(o \leftrightarrow u) \geq e^{-(\log \log \gamma^+)^{1/2}}.$$

This proof of Proposition 8.3.1 is by induction. In Section 8.3, we describe the high-level structure of this induction, which is essentially the same as in [EH23b, Section 3.2], except for two differences. The first difference is that we have reworded the induction to say “we can keep propagating until and unless we reach a scale where the geometry is bad”, with a separate lemma that says “if the geometry is bad at scale n , then G starts to look one-dimensional from around scale n ”. In contrast, the induction in the earlier work simply says “if G is not one-dimensional, then we can keep propagating forever”. The second difference is that the induction in the earlier work is slightly coarser in the sense that it groups multiple inductive steps of the argument we present here into a single inductive step. The additional detail in the present version is necessary to close the gap between the last scale from which we can propagate connection bounds and the first scale at which we can prove that the geometry “is bad”.

The individual inductive steps are all implicit in [EH23b, Sections 4 and 6]. We will justify these in the appendix. In Section 8.3, we will prove something like the “base case” of the induction. This follows by a compactness argument from some intermediary results in [EH23b] and [CMT22]. In Section 8.3, we will prove that “if the geometry is bad at scale n , then G starts to look one-dimensional from around scale n ”. This subsection is a refinement of [EH23b, Section 5], but for the reasons discussed, it will require some new ideas.

The logic of the induction

For the entirety of this subsection, fix a finite transitive graph G with degree d , and define γ as in the statement of Proposition 8.3.1. We will describe the repeated-sprinkling multi-scale induction argument used to prove Proposition 8.3.1 (which is adapted from [EH23b]) as a deterministic colouring process evolving over time. At every time $t \in \mathbb{R}$, every scale¹⁸ $n \in [3, \infty)$ can be coloured orange or green (or both, or neither — i.e. uncoloured), encoding a statement¹⁹ about the connectivity properties of percolation of parameter²⁰ $\phi(t) := 1 - 2^{-e^t}$ over distances of approximately n .

¹⁸It would have been more natural to consider scales $n \in \mathbb{N}$ rather than $n \in [3, \infty)$. We chose the latter to avoid rounding issues and so that $\log \log n$ is always positive.

¹⁹Formally, this colouring can be encoded as a function $\text{colour} : [3, \infty) \times \mathbb{R} \rightarrow \mathcal{P}(\{\text{orange}, \text{green}\})$, where $\mathcal{P}(X)$ means the powerset of X . We say “ n is green at time t ” to mean that $\text{colour}(n, t) \ni \text{orange}$. Similar statements are formalised analogously.

²⁰This choice of parameterisation appears implicitly in [EH23b] as the natural choice for arguments that involve repeated sprinkling. Indeed, our function ϕ is the function $\text{Spr}(p; \lambda)$ from [EH23b, Section 3.1] evaluated at $(1/2; t)$.

To lighten notation, let $\delta(n) := e^{-(\log \log n)^{1/2}}$ denote the standard small-quantity associated to each scale n . Now we colour a scale n *orange* at time t to mean that

$$\min_{u \in B_n} \mathbb{P}_{\phi(t)}(o \leftrightarrow u) \geq \delta(n).$$

We also define the move-right (aka increase-scale) function $R : n \mapsto e^{(\log n)^9}$, and write $R^k := R \circ \dots \circ R$ for the k -fold composition of R with itself. Now to prove Proposition 8.3.1, it suffices to prove the following lemma.

Lemma 8.3.2. *For all $\varepsilon > 0$ and $n_1 \geq 3$ there exists $n_2(d, \varepsilon, n_1) < \infty$ such that for all $t \in \mathbb{R}$ with $t \leq \frac{1}{\varepsilon}$, if $[n_1, n_2]$ is orange at time t , then $R(\gamma)$ is orange at time $t + \varepsilon$.*

Proof of Proposition 8.3.1 given Lemma 8.3.2. Fix $\varepsilon, \eta > 0$. Let $n_1(\eta)$ be the smallest integer satisfying $n_1 \geq 3$ and $\delta(n_1) \leq \eta$. Let $\alpha(\varepsilon) > 0$ be the unique real satisfying $\phi(1/\alpha) = 1 - \varepsilon$. Let $n_2(d, \alpha \wedge \varepsilon, n_1) < \infty$ be the constant that is guaranteed to exist by Lemma 8.3.2. We claim that we can take $\lambda := n_2$. Indeed, let $p \in (0, 1)$ and suppose that $\min_{u \in B_{n_2}} \mathbb{P}_p(o \leftrightarrow u) \geq \eta$. Define $t := \phi^{-1}(p)$. The claim is trivial if $p \geq 1 - \varepsilon$, so we may assume that $t \leq 1/\alpha$. By monotonicity of the function $\delta(\cdot)$, the interval $[n_1, n_2]$ is orange at time t . So by applying Lemma 8.3.2, $R(\gamma)$ is orange at time $t + \varepsilon$. Since (by calculus) ϕ is 1-Lipschitz, $R(\gamma)$ is also orange at time $\phi^{-1}(p + \varepsilon)$, which is the required conclusion. \square

We say that a set $M \subseteq \mathbb{N}$ is a certain colour if every $m \in M$ is that colour. Given a statement A about a colouring at an implicit time t , we define $s(A) := \inf\{t : A \text{ is true at time } t\}$ where $\inf \emptyset := +\infty$. For example,

$$s(\{10, 12\} \text{ is orange}) := \inf\{t : 10 \text{ and } 12 \text{ are both orange at time } t\} \in [-\infty, +\infty].$$

As a first approximation to our induction, imagine we knew that for every scale n with $n \leq \gamma$,

$$s(R(n) \text{ is orange}) \leq s(n \text{ is orange}) + \delta(n). \quad (8.3.1)$$

Suppose for simplicity that some positive integer r satisfies $R^r(n_2) = \gamma$ and that n_2 exceeds some large universal constant. Then by repeatedly applying eq. (8.3.1), we could deduce that

$$\begin{aligned} s(R(\gamma) \text{ is orange}) - s([n_1, n_2] \text{ is orange}) &\leq \sum_{k=0}^r \left[s(R^{k+1}(n_2) \text{ is orange}) - s(R^k(n_2) \text{ is orange}) \right] \\ &\leq \sum_{k=0}^r \delta(R^k(n_2)) \leq \sum_{k=0}^{\infty} e^{-(9^k \log \log n_2)^{1/2}} \leq 2\delta(n_2). \end{aligned}$$

Since $\delta(n_2) \rightarrow 0$ as $n_2 \rightarrow \infty$, this would certainly imply Lemma 8.3.2. Rather than prove something as direct as eq. (8.3.1), we will have to bring into play a new colour, *green*.

Given a finite path $\nu = (\nu_i)_{i=0}^k$, we write $\text{start}(\nu) := \nu_0$ and $\text{end}(\nu) := \nu_k$ for its start and end vertices, $|\nu| := k$ for its length, and given $r \geq 0$, we write $B_r(\nu) := \bigcup_{i=0}^k B_r(\nu_i)$ for the associated tube. Given $m, n \geq 1$ and $p \in (0, 1)$, define the *corridor function* (which we take from [CMT22]),

$$\kappa_p(m, n) := \inf_{\nu: |\nu| \leq m} \mathbb{P}_p \left(\text{start}(\nu) \xleftrightarrow{\omega \cap B_n(\nu)} \text{end}(\nu) \right).$$

Notice that we always have $\kappa_p(m, n) \leq \min_{u \in B_m} \mathbb{P}_p(o \leftrightarrow u)$. Let us also define the set of low-growth scales $\mathbb{L} := \{n \geq 3 : \text{Gr}(n) \leq e^{(\log n)^{100}}\}$. Now we colour a scale n *green* at time t to mean that n is orange and either $n \notin \mathbb{L}$ or $\kappa_{\phi(t)}(R^2(n), n) \geq \delta(R(n))$. We will use this new colour to help us propagate orange by controlling the time taken for orange scales to turn green and for green scales to turn nearby scales orange.

The next lemma says that green scales quickly turn nearby scales orange. If we see that n is green at some time t but $n \in \mathbb{L}$, then $\kappa_{\phi(t)}(R^2(n), n) \geq \delta(R(n))$, which trivially implies that $[R(n), R^2(n)]$ is already orange. So the content of this lemma is that if instead $n \notin \mathbb{L}$, then we can efficiently propagate a point-to-point connection lower bound from scale n to scales in $[R(n), R^2(n)]$. The proof of this is implicit in [EH23b, Section 4]. The argument uses some ghost-field technology that works more efficiently around scales n where $\text{Gr}(n)$ is large. See the appendix for details.

Lemma 8.3.3. *There exists $n_0(d) < \infty$ such that for all $n \geq n_0$,*

$$s([R(n), R^2(n)] \text{ is orange}) \leq s(n \text{ is green}) + \delta(n).$$

The next lemma sometimes lets us control how long it takes for a scale n to become green after turning orange. If $n \notin \mathbb{L}$, then this time is trivially zero. So it suffices to consider $n \in \mathbb{L}$. We might hope for a statement like the following: there exists $n_0(d) < \infty$ such that for all $n \geq n_0$ with $n \in \mathbb{L}$,

$$s(n \text{ is green}) \leq s(n \text{ is orange}) + \delta(n). \tag{8.3.2}$$

Our next lemma is less satisfying in two ways.²¹ First, we can only prove an upper bound like eq. (8.3.2) when n belongs to a particular distinguished subset $\mathbb{T}(c, \lambda)$ of \mathbb{L} . Second, our upper bound is in terms of a mysterious quantity Δ rather than something explicit like $\delta(n)$. So to use this lemma, we will need to: (1) Deal with scales $n \in \mathbb{L} \setminus \mathbb{T}(c, \lambda)$, and (2) Find a way to upper bound Δ explicitly. For completeness, we will now define $\mathbb{T}(c, \lambda)$ and Δ , but the reader should feel free

²¹We need to use the non-one-dimensionality hypothesis somewhere.

to skip these definitions for now because they are not necessary to follow the high-level induction argument being developed in this subsection.

Given constants $c, \lambda > 0$, let $\mathbb{T}(c, \lambda)$ be the set of scales $n \in \mathbb{L}$ such that G has (c, λ) -polylog plentiful tubes at every scale in an interval of the form $[m, m^{1+c}]$ that is contained in $[n^{1/3}, n^{1/(1+c)}]$. We will recall the definition of plentiful polylog tubes, taken from [EH23b, Section 5], in Section 8.3. Given $m, n \geq 1$, we define $\text{Piv}[m, n]$ to be the event that in the restricted configuration $\omega \cap B_n$, there are at least two clusters that each contain an open path from B_m to S_n . For each scale n and time t , let $U_t(n)$ be the *uniqueness zone* defined to be the maximum integer $b \leq \frac{1}{8}n^{1/3}$ satisfying $\mathbb{P}_{\phi(t)}(\text{Piv}[4b, n^{1/3}]) \leq (\log n)^{-1}$. The associated *cost* is

$$\Delta_t(n) := \left[\frac{\log \log n}{(\log n) \wedge \log \text{Gr}(U_t(n))} \right]^{1/4}.$$

Note that the cost is small if $\text{Gr}(U_t(n)) \geq (\log n)^C$ for a big constant C . The proof of the next lemma is implicit in [EH23b, Section 6], where, together with Hutchcroft, we used plentiful tubes to run an orange-peeling argument inspired by the one in [CMT22]. See the appendix for details.

Lemma 8.3.4. *For all $c > 0$ there exist $\lambda(d, c), n_0(d, c), K(d, c) < \infty$ such that the following holds for all $n \geq n_0$ with $n \in \mathbb{T}(c, \lambda)$. For all $t \in \mathbb{R}$, if n is orange at time t and $K\Delta_t(n) \leq 1$ then*

$$s(n \text{ is green}) \leq t + K\Delta_t(n).$$

We now turn to the problem of finding an explicit upper bound on the cost $\Delta = \Delta_t(n)$ of a scale n at a time t . Define the move-left (aka decrease-scale) function $L : n \mapsto (\log n)^{1/2}$. The next lemma provides such an upper bound if the much smaller scale $L(n)$ happens to already be green at time t . The proof of this is implicit in [EH23b, Section 6.3]. Notice that to upper bound $\Delta_t(n)$ is to lower bound $\text{Gr}(U_t(n))$. It is easy to check that in the setting of this lemma, $U_t(n) \geq L(n)$. The proof of the lemma establishes that when $L(n)$ is green and $n \in \mathbb{L}$, either $\text{Gr}(L(n))$ is big (as a function of $L(n)$) or we can find a better lower bound on $U_t(n)$ than the trivial bound that is $L(n)$. See the appendix for details.

Lemma 8.3.5. *There exists $n_0(d) < \infty$ such that the following holds for all $n \in \mathbb{L}$ with $n \geq n_0$. For all $t \in \mathbb{R}$, if $L(n)$ is green at time t then*

$$\Delta_t(n) \leq \frac{1}{\log \log n}.$$

Lemmas 8.3.4 and 8.3.5 together provide an explicit upper bound on the time it takes for a scale n that is orange to become green if the much smaller scale $L(n)$ happens to already be green, at least

until we encounter a scale $n \in \mathbb{L} \setminus \mathbb{T}(c, \lambda)$. Of course, this says nothing about how long we have to wait for at least one orange scale to become green in the first place. We will return to this shortly, but for now, consider the following method for rapidly propagating orange once we have a big interval of green. Suppose that at some time t , some interval of the form $[L(n), n]$ is green. By Lemma 8.3.3, since $R^{-1}(n) \in [L(n), n]$, we will not have to wait long for $[n, R(n)]$ to turn orange. By Lemmas 8.3.4 and 8.3.5, since $L(m) \in [L(n), n]$ for every $m \in [n, R(n)]$, we will not then have to wait long for $[n, R(n)]$ to turn green. By Lemma 8.3.3 again, we will not then have to wait long for $[R(n), R^2(n)]$ to turn orange, and so forth. We can repeat this indefinitely until and unless we encounter a scale $n \in \mathbb{L} \setminus \mathbb{T}(c, \lambda)$.

In conjunction with Lemma 8.3.4, the next lemma lets us control long it takes to get this big interval of green in the first place, starting from an even bigger interval of orange. We prove this in Section 8.3. We will use a compactness argument to reduce this to an analogous statement about an arbitrary infinite unimodular transitive graph G , which is then addressed by results in either [CMT22] or [EH23b] according to whether G has polynomial or superpolynomial growth.

Lemma 8.3.6. *For all $\varepsilon > 0$ and $n_1 \geq 3$, there exists $n_2(d, \varepsilon, n_1) < \infty$ such that the following holds for all $t \in \mathbb{R}$ with $t \leq \frac{1}{\varepsilon}$. If $[n_1, n_2]$ is orange at time t , then for some m satisfying $I := [L(m), m] \subseteq [n_1, n_2]$, we have*

$$\sup_{n \in I \cap \mathbb{L}} \Delta_{t+\varepsilon}(n) \leq \varepsilon.$$

At this point, the lemmas we have accumulated allow us to rapidly propagate orange, starting from a big interval of orange, until and unless we encounter a scale $n \in \mathbb{L} \setminus \mathbb{T}(c, \lambda)$. To prove Lemma 8.3.2, we need this propagation to keep going until we encounter the scale γ . The next lemma lets us ensure that we will encounter γ before we encounter $\mathbb{L} \setminus \mathbb{T}(c, \lambda)$. We will prove this in Section 8.3. This is the analogue in our setting of [EH23b, Section 5]. While the random walk arguments that make up [EH23b, Subection 5.2] work equally well in our setting, the geometric arguments in [EH23b, Subsection 5.1] will require some new ideas.

Lemma 8.3.7. *There exist $c(d) > 0$ such that for all $\lambda \geq 1$, there exists $n_0(d, \lambda) < \infty$ such that*

$$\inf\{n \in \mathbb{L} \setminus \mathbb{T}(c, \lambda) : n \geq n_0\} \geq \gamma.$$

Let us conclude by formalising the above sketch of the fact that Lemma 8.3.2, which we know implies Proposition 8.3.1, can be reduced to the rest of the lemmas introduced in this subsection.

Proof of Lemma 8.3.2 given Lemmas 8.3.3 to 8.3.7. Fix $\varepsilon > 0$ and $n_1 \geq 3$. We may assume that $\varepsilon < 1$. Let $c(d) > 0$ be the constant from Lemma 8.3.7. Let $\lambda(d), u_0(d), K(d)$ be the constants “ $\lambda(d, c), n_0(d, c), K(d, c)$ ” from Lemma 8.3.4 for this choice of c . We may assume that $K \geq 1$. Let $u_1(d)$ be the constant “ $n_0(d)$ ” from Lemma 8.3.3. Let $u_2(d)$ be the constant “ $n_0(d)$ ” from Lemma 8.3.5. Let $u_3(d)$ be the constant “ $n_0(d, \lambda)$ ” from Lemma 8.3.7, with the above choice of λ . Note that $\sum_i \delta(R^i(3)) < \infty$ and $\sum_i \frac{1}{\log \log R^i(3)} < \infty$. Let $i_0(d, \varepsilon)$ be the smallest non-negative integer such that

$$\sum_{i=i_0}^{\infty} \delta(R^i(3)) \leq \frac{\varepsilon}{5} \quad \text{and} \quad \sum_{i=i_0}^{\infty} \frac{1}{\log \log R^i(3)} \leq \frac{\varepsilon}{5K},$$

and set $u_4(d, \varepsilon) := R^{i_0}(3)$. Set $u_5(d, \varepsilon, n_1) := \max\{u_0, u_1, u_2, u_3, u_4, n_1\}$. Let $u_6(d, \varepsilon, n_1)$ be the constant “ $n_2(d, \frac{\varepsilon}{5K}, u_5)$ ” from Lemma 8.3.6. We claim that the conclusion holds with $n_2(d, \varepsilon, n_1) := R(u_6)$.

Let $t \in \mathbb{R}$ with $t \leq \frac{1}{\varepsilon}$, and suppose that $[n_1, n_2]$ is orange at time t . By Lemma 8.3.6, there exists m with $I := [L(m), m] \subseteq [u_5, u_6]$ such that $\sup_{n \in I \cap \mathbb{L}} K\Delta_{t+\frac{\varepsilon}{5}}(n) \leq \frac{\varepsilon}{5} \leq 1$. Consider the possibility that $I \cap (\mathbb{L} \setminus \mathbb{T}(c, \lambda)) \neq \emptyset$. Then by Lemma 8.3.7, $\gamma \leq u_6$. In particular, $R(\gamma) \leq n_2$, and hence $R(\gamma)$ is already orange at time t . Since we are trivially done in that case, let us assume to the contrary that $I \cap (\mathbb{L} \setminus \mathbb{T}(c, \lambda)) = \emptyset$. Then by Lemma 8.3.4,

$$s(I \text{ is green}) \leq t + \frac{\varepsilon}{5} + \sup_{n \in I \cap \mathbb{L}} K\Delta_{t+\frac{\varepsilon}{5}}(n) \leq t + \frac{2\varepsilon}{5}.$$

Let k be the largest non-negative integer such that $R^k(m) < \gamma$. (We may assume that such an integer exists, otherwise $\gamma \leq u_6$ and hence we are trivially done as above.) We claim that for all $i \in \{0, \dots, k-1\}$,

$$s([L(m), R^{i+1}(m)] \text{ is green}) \leq s([L(m), R^i(m)] \text{ is green}) + \delta(R^{i-1}(m)) + \frac{K}{\log \log R^i(m)}. \quad (8.3.3)$$

Indeed, fix an arbitrary index $i \in \{0, \dots, k-1\}$ and an arbitrary time $s \in \mathbb{R}$ at which $[L(m), R^i(m)]$ is green. By Lemma 8.3.3, the interval $[R^i(m), R^{i+1}(m)]$ is orange at time $s + \delta(R^{i-1}(m))$. By Lemma 8.3.5, since $L([R^i(m), R^{i+1}(m)]) \subseteq [L(m), R^i(m)]$,

$$\sup_{n \in [R^i(m), R^{i+1}(m)] \cap \mathbb{L}} \Delta_{s+\delta(R^{i-1}(m))}(n) \leq \frac{1}{\log \log R^i(m)}.$$

By Lemma 8.3.7, we know that $[R^i(m), R^{i+1}(m)] \cap (\mathbb{L} \setminus \mathbb{T}(c, \lambda)) = \emptyset$, and since $R^i(m) \geq u_4$, we know that $\frac{K}{\log \log R^i(m)} \leq \frac{\varepsilon}{5} \leq 1$. So by Lemma 8.3.4,

$$s([R^i(m), R^{i+1}(m)] \text{ is green}) \leq s + \delta(R^{i-1}(m)) + \frac{K}{\log \log R^i(m)},$$

establishing eq. (8.3.3). By repeated applying eq. (8.3.3), it follows by induction that

$$s\left(\left[L(m), R^k(m)\right] \text{ is green}\right) \leq t + \frac{2\varepsilon}{5} + \sum_{i=0}^{\infty} \left(\delta\left(R^{i-1}(m)\right) + \frac{K}{\log \log R^i(m)} \right) \leq t + \frac{4\varepsilon}{5},$$

where in the second inequality we used the fact that $R^{-1}(m) \geq u_4$. By maximality of k , we know that $R^{-1}(\gamma) \in \left[L(m), R^k(m)\right]$. So by Lemma 8.3.3,

$$\begin{aligned} s(R(\gamma) \text{ is orange}) - s\left(\left[L(m), R^k(m)\right] \text{ is green}\right) &\leq s([\gamma, R(\gamma)] \text{ is orange}) - s(R^{-1}(\gamma) \text{ is green}) \\ &\leq \delta(R^{-1}(\gamma)) \leq \delta(u_4) \leq \frac{\varepsilon}{5}. \end{aligned}$$

Therefore, as required,

$$s(R(\gamma) \text{ is orange}) \leq t + \frac{4\varepsilon}{5} + \frac{\varepsilon}{5} = t + \varepsilon. \quad \square$$

Base case of the induction

In this subsection we prove Lemma 8.3.6. By a compactness argument, we will reduce this to the following simpler statement about individual infinite transitive graphs. Although we defined $\Delta_t(n)$ and \mathbb{L} in the context of *finite* transitive graphs, let us use the exact same definitions for infinite transitive graphs.

Lemma 8.3.8. *Let G be a unimodular infinite transitive graph. For every $t \in \mathbb{R}$ with $\phi(t) > p_c(G)$,*

$$\lim_{n \rightarrow \infty} \sup_{s \geq t} \Delta_s(n) \mathbb{1}_{\mathbb{L}}(n) = 0.$$

Our first goal is to prove this lemma. As mentioned earlier, to show that $\Delta_t(n)$ is small, we need to show that $\text{Gr}(U_t(n)) \geq (\log n)^C$ for a large constant C . The following lemma²² from [EH23b, Corollary 2.4] tells us in particular that if $\text{Gr}(n)$ is not too big with respect to n , then the uniqueness zone for n is always at least of order $\log n$. The proof of this result was essentially already contained in [CMT22], which in turn was inspired by [Cer15].

Lemma 8.3.9. *Let G be a unimodular transitive graph of degree d . Fix $\eta \in (0, 1)$ and $\varepsilon \in (0, 1/2)$. There exists $c(d, \eta, \varepsilon) > 0$ such that for every $n \geq 1$ and $p \in [\eta, 1]$,*

$$\mathbb{P}_p(\text{Piv}[c \log n, n]) \leq \left(\frac{\log \text{Gr}(n)}{cn} \right)^{\frac{1}{2} - \varepsilon}.$$

²²In the version in [EH23b], we also required that n is larger than some constant depending on d, η, ε , but that is redundant.

In an infinite transitive graph, if $U_t(n) \gtrsim \log n$ then trivially $\text{Gr}(U_t(n)) \gtrsim \log n$, and hence $\Delta_t(n)$ is bounded above by a (possibly large) constant. Our goal is to improve this argument so that this constant can be made arbitrarily small. We will do this by improving either the bound on $U_t(n)$ or the bound on $\text{Gr}(U_t(n))$ given $U_t(n)$. When G has superpolynomial growth, this is easy: we can use the trivial bound on $U_t(n)$, but then use the fact that $\text{Gr}(U_t(n)) \gtrsim (U_t(n))^C$ for any particular constant C ²³. When G has polynomial growth, we will apply the following more delicate result from [CMT22, Proposition 6.1] to improve our bound on $U_t(n)$. (For background on transitive graphs of polynomial and superpolynomial growth, see [TT21a].)

Lemma 8.3.10. *Let G be an infinite transitive graph of polynomial growth. For every $p > p_c(G)$ there exist $\chi(G, p) \in (0, 1)$ and $C(G, p) < \infty$ such that for every $n \geq 1$ and $q \in [p, 1]$,*

$$\mathbb{P}_q \left(\text{Piv} \left[e^{(\log n)^\chi}, n \right] \right) \leq Cn^{-1/4}.$$

Proof of Lemma 8.3.8. Suppose that $t \in \mathbb{R}$ satisfies $\phi(t) > p_c(G)$. Note that $\phi(t) > 1/d$ because $p_c(G) \geq 1/(d-1)$ (as this holds for every infinite transitive graph). Let $c_1(d, 1/d, 1/6) > 0$ be the constant from Lemma 8.3.9. Then for every sufficiently large $n \in \mathbb{L}$,

$$\sup_{s \geq t} \mathbb{P}_{\phi(s)} \left(\text{Piv} \left[c_1 \log \left(n^{1/3} \right), n^{1/3} \right] \right) \leq \left(\frac{\log \text{Gr}(n^{1/3})}{c_1 n^{1/3}} \right)^{\frac{1}{2} - \frac{1}{6}} \leq \left(\frac{(\log n)^{100}}{c_1 n^{1/3}} \right)^{\frac{1}{3}} \leq \frac{1}{\log n},$$

and hence $\inf_{s \geq t} U_s(n) \geq \lfloor \frac{1}{4} c_1 \log(n^{1/3}) \rfloor = \frac{c_1}{13} \log n$. In particular,

$$\limsup_{n \rightarrow \infty} \sup_{s \geq t} \Delta_s(n) \mathbb{1}_{\mathbb{L}}(n) \leq \limsup_{n \rightarrow \infty} \left[\frac{\log \log n}{(\log n) \wedge \log \text{Gr} \left(\frac{c_1}{13} \log n \right)} \right]^{1/4}.$$

If G has superpolynomial growth, then $\frac{\log \text{Gr} \left(\frac{c_1}{13} \log n \right)}{\log \log n} \rightarrow \infty$ and hence $\sup_{s \geq t} \Delta_s(n) \mathbb{1}_{\mathbb{L}}(n) \rightarrow 0$ as $n \rightarrow \infty$. So we may assume to the contrary that G has polynomial growth. Let $\chi(G, \phi(t))$ and $C(G, \phi(t))$ be the constants from Lemma 8.3.10. Then for every sufficiently large $n \geq 1$,

$$\sup_{s \geq t} \mathbb{P}_{\phi(s)} \left(\text{Piv} \left[e^{(\log(n^{1/3}))^\chi}, n^{1/3} \right] \right) \leq C(n^{1/3})^{-1/4} \leq \frac{1}{\log n},$$

and hence $\inf_{s \geq t} U_s(n) \geq \frac{1}{4} e^{(\log(n^{1/3}))^\chi} \geq e^{(\log n)^{\chi/2}}$. In particular, using the trivial bound $\text{Gr}(U_s(n)) \geq U_s(n)$,

$$\limsup_{n \rightarrow \infty} \sup_{s \geq t} \Delta_s(n) \mathbb{1}_{\mathbb{L}}(n) \leq \limsup_{n \rightarrow \infty} \left[\frac{\log \log n}{(\log n) \wedge \log \left(e^{(\log n)^{\chi/2}} \right)} \right]^{1/4} = 0.$$

²³This was the idea in [EH23b, Section 3], where it sufficed to consider sequences converging to graphs of superpolynomial growth.

□

Next we will use a compactness argument to deduce Lemma 8.3.6 from Lemma 8.3.8. Given $d \in \mathbb{N}$, let \mathcal{U}_d be the space of all unimodular transitive graphs with degree d endowed with the local topology. Recall that every finite transitive graph is unimodular. By [Hut20a, Corollary 5.5], \mathcal{U}_d is a closed subset of the space \mathcal{T}_d of all transitive graphs with degree d endowed with the local topology. In particular (recalling that the local topology is metrisable), since \mathcal{T}_d is compact, so is \mathcal{U}_d .

Proof of Lemma 8.3.6. Suppose for contradiction that the statement is false. Then we can find $\varepsilon > 0$ and $n_1 \geq 3$ such that for all $N \in \mathbb{N}$ there exists $t_N \leq \frac{1}{\varepsilon}$ and a finite transitive graph G_N with degree d such that in G_N , the interval $[n_1, N]$ is orange at time t_N , but for every m with $[L(m), m] \subseteq [n_1, N]$, there exists $n \in [L(m), m] \cap \mathbb{L}(G_N)$ with $\Delta_{t_N+\varepsilon}^{G_N}(n) > \varepsilon$. (We write $\Delta^G, U^G, \mathbb{L}(G)$ to denote Δ, U, \mathbb{L} defined with respect to a specific graph G .) By compactness, there exists an infinite subset $\mathbb{M} \subseteq \mathbb{N}$ and a unimodular transitive graph G such that $G_N \rightarrow G$ as $N \rightarrow \infty$ with $N \in \mathbb{M}$.

First consider the case that G is finite. Then trivially, there exists $n_0(G) < \infty$ such that for all $n \geq n_0$ and for all $s \in \mathbb{R}$, we have $U_s^G(n) = \lfloor \frac{1}{8}n^{1/3} \rfloor$. In particular, $\lim_{n \rightarrow \infty} \sup_{s \in \mathbb{R}} \Delta_s^G(n) = 0$. So there exists m with $L(m) \geq n_1$ such that

$$\sup_{n \in [L(m), m]} \sup_{s \in \mathbb{R}} \Delta_s^G(n) \leq \varepsilon.$$

Pick $N \in \mathbb{M}$ sufficiently large that $N \geq m$ and $B_m^{G_N} \cong B_m^G$ (or even that $G \cong G_N$). Then we have a contradiction because there exists $n \in [L(m), m]$ such that $\Delta_{t_N+\varepsilon}^G(n) = \Delta_{t_N+\varepsilon}^{G_N}(n) > \varepsilon$.

So we may assume that G is infinite. We claim that

$$\liminf_{\substack{N \rightarrow \infty \\ N \in \mathbb{M}}} \phi(t_N) \geq p_c(G). \quad (8.3.4)$$

Indeed, suppose that $q \in (0, p_c(G))$. By the sharpness of the phase transition for percolation on infinite transitive graphs, there exists $C(q, G) < \infty$ such that $\mathbb{P}_q^G(o \leftrightarrow S_n) \leq Ce^{-n/C}$ for all $n \geq 1$. Pick $m \geq n_1$ such that $Ce^{-m/C} < \delta(m)$. Pick $N_0 \geq m$ such that for all $N \geq N_0$ with $N \in \mathbb{M}$, we have $B_m^{G_N} \cong B_m^G$. Then for all $N \geq N_0$ with $N \in \mathbb{N}$,

$$\min_{u \in B_m^{G_N}} \mathbb{P}_q^{G_N}(o \leftrightarrow u) \leq \mathbb{P}_q^{G_N}(o \leftrightarrow S_m) = \mathbb{P}_q^G(o \leftrightarrow S_m) < \delta(m),$$

so m is not orange for G_N at time $\phi^{-1}(q)$, and hence $q \leq \phi(t_N)$. Since q was arbitrary, this establishes eq. (8.3.4). Now by hypothesis, $t_N \leq \frac{1}{\varepsilon}$ for every $N \geq 1$. So by eq. (8.3.4), we know that $p_c(G) \leq \phi(1/\varepsilon) < 1$. (We also know that $p_c(G) > 0$ since this holds for every infinite transitive graphs.) By passing to a further subsequence, we may assume that for all $N \in \mathbb{M}$,

$$t_N + \varepsilon \geq \phi^{-1}(p_c(G)) + \frac{\varepsilon}{2} =: t.$$

Note that $\phi(t) > p_c(G)$. So by Lemma 8.3.8,

$$\lim_{n \rightarrow \infty} \sup_{s \geq t} \Delta_s^G(n) \mathbb{1}_{\mathbb{L}(G)}(n) = 0.$$

Pick m with $L(m) \geq n_1$ such that

$$\sup_{n \in [L(m), m]} \sup_{s \geq t} \Delta_s^G(n) \mathbb{1}_{\mathbb{L}(G)}(n) \leq \varepsilon. \quad (8.3.5)$$

Pick $N \in \mathbb{M}$ such that $N \geq m$ and $B_m^{G_N} \cong B_m^G$. Then $[L(m), m] \subseteq [n_1, N]$, and the same inequality as eq. (8.3.5) holds with G_N in place of G . This contradicts the existence of $n \in [L(m), m] \cap \mathbb{L}(G_N)$ satisfying $\Delta_{t_N + \varepsilon}^{G_N}(n) > \varepsilon$ because $t_N + \varepsilon \geq t$. \square

The obstacles are circles

In this subsection we prove Lemma 8.3.7. Our argument is a finitary refinement of the argument in [EH23b, Section 5]. Our first step is to isolate the part of that previous argument that needs to be improved. For this we need to introduce the definition of plentiful tubes.

Plentiful tubes Let G be a transitive graph and fix a scale $n \geq 1$. We call the r -neighbourhood $B_r(\gamma) := \bigcup_i B_r(\gamma_i)$ of a path $\gamma \in \Gamma$ a *tube*. Given constants $k, r, l \geq 1$, we say that G has (k, r, l) -*plentiful tubes* at scale n if the following always holds. Let A and B be sets of vertices such that $(A, B) = (S_n, S_{4n})$ or such that A and B both contain paths from S_n to S_{3n} . Then there is a set Γ of paths from A to B such that $|\Gamma| \geq k$, each path has length at most l , and $B_r(\gamma_1) \cap B_r(\gamma_2) = \emptyset$ for all pairs of distinct paths $\gamma_1, \gamma_2 \in \Gamma$. Note that the property of having (k, r, l) -plentiful tubes gets stronger as we increase k (the number of tubes), increase r (the thickness of tubes), or decrease l (the lengths of tubes). We will be concerned mainly with the following two-parameter subset of this three-parameter family of properties. Given constants $c, \lambda > 0$, we say that G has (c, λ) -*polylog plentiful tubes* at scale n if G has (k, r, l) -plentiful tubes at scale n with

$$(k, r, l) := \left([\log n]^{c\lambda}, n[\log n]^{-\lambda/c}, n[\log n]^{\lambda/c} \right).$$

We think of c as representing a fixed exchange rate for the tradeoff between asking for more tubes that are long and thin vs fewer tubes that are short and thick, which we can realise by varying λ . Finally, recall from Section 8.3 that $\mathbb{T}(c, \lambda)$ is defined to be the set of all scales $n \geq 3$ such that $\text{Gr}(n) \leq e^{(\log n)^{100}}$ (i.e. $n \in \mathbb{L}$) and there exists m satisfying $[m, m^{1+c}] \subseteq [n^{1/3}, n^{1/(1+c)}]$ such that G has (c, λ) -polylog plentiful tubes at every scale in $[m, m^{1+c}]$.

Now suppose in the context of proving Lemma 8.3.7 that we have a large scale $n \in \mathbb{L}$ with $n < \gamma$, and we want to build the required plentiful tubes to establish that $n \in \mathbb{T}(c, \lambda)$. We split our argument into two cases, slow growth and fast growth, according to the rate of change of Gr near n , as measured by whether $\text{Gr}(3m)/\text{Gr}(m)$ for $m \approx n$ exceeds some particular constant²⁴. The next lemma says that if G has fast growth throughout a sufficiently large interval around scale n , then for some fixed exchange rate c , we have (c, λ) -polylog plentiful tubes for every choice of λ whenever n is sufficiently large. This is [EH23b, Proposition 5.4], which was originally stated for infinite unimodular transitive graphs, but as we will justify in the appendix, exactly the same proof also works for finite transitive graphs.

Lemma 8.3.11. *Let G be a unimodular transitive graph of degree d . Suppose that*

$$\text{Gr}(m) \leq e^{(\log m)^D} \quad \text{and} \quad \text{Gr}(3m) \geq 3^5 \text{Gr}(m)$$

for every $m \in [n^{1-\varepsilon}, n^{1+\varepsilon}]$, where $\varepsilon, D, n > 0$. Then there is a constant $c(d, D, \varepsilon) > 0$ with the following property. For every $\lambda \geq 1$, there exists $n_0(d, D, \varepsilon, \lambda) < \infty$ such that if $n \geq n_0$ then G has (c, λ) -polylog plentiful tubes at scale n .

The next lemma says that if G has slow growth at some scale n , then outside of a bounded number of small problematic intervals, G has plentiful tubes with good constants (k, r, l) unless G is one-dimensional. This is equivalent to [EH23b, Proposition 5.3].

Lemma 8.3.12. *Let G be an infinite transitive graph of degree d . Suppose that $\text{Gr}(3n) \leq 3^\kappa \text{Gr}(n)$, where $n, \kappa > 0$. There exists $C(d, \kappa) < \infty$ such that the following holds if $n \geq C$:*

There is a set $A \subseteq [1, \infty)$ with $|A| \leq C$ such that for every $k \geq 1$ and every $m \in [Ckn, \infty) \setminus \bigcup_{a \in A} [a, 2ka]$, if G does not have $(C^{-1}k, C^{-1}k^{-1}m, Ck^C m)$ -plentiful tubes at scale m , then G is one-dimensional.

²⁴In [EH23b] we considered ratios of triplings $\text{Gr}(3n)/\text{Gr}(n)$ rather than of doublings $\text{Gr}(2n)/\text{Gr}(n)$ because only the former was known at the time to be sufficient to invoke the structure theory of transitive graphs of polynomial growth. Tointon and Tessera have since proved that small doublings imply small triplings [TT23], so it is now possible to work with doublings $\text{Gr}(2n)/\text{Gr}(n)$ instead, which is slightly more natural. However, since this does not significantly simplify our arguments, we have chosen to stay with triplings to avoid some repetition of work from [EH23b].

We need to improve this lemma in two ways. First, we need the conclusion to be that “ G looks one-dimensional from around scale m ”, rather than just “ G is one-dimensional”. Second, we need to allow G to be finite. Here is the modified version of Lemma 8.3.12 that we will prove.²⁵

Lemma 8.3.13. *Let G be a finite transitive graph of degree d . Suppose that $\text{Gr}(3n) \leq 3^\kappa \text{Gr}(n)$, where $n, \kappa > 0$. There exists $C(d, \kappa) < \infty$ such that the following holds if $n \geq C$:*

There is a set $A \subseteq [1, \infty)$ with $|A| \leq C$ such that for every $k \geq 1$ and every $m \in [Ckn, \infty) \setminus \bigcup_{a \in A} [a, 2ka]$, if G does not have $(C^{-1}k, C^{-1}k^{-1}m, Ck^Cm)$ -plentiful tubes at scale m , then

$$\text{dist}_{\text{GH}} \left(\frac{\pi}{\text{diam } G} G, S^1 \right) \leq \frac{Cm}{\text{diam } G}.$$

In the next two subsections we will prove Lemma 8.3.13. Before that, let us quickly check that Lemmas 8.3.11 and 8.3.13 together do imply Lemma 8.3.7. This is essentially the same as the proof of [EH23b, Proposition 5.2] given [EH23b, Propositions 5.3 and 5.4].

Proof of Lemma 8.3.7 given Lemmas 8.3.11 and 8.3.13. Let $c(d) > 0$ be a small positive constant to be determined. Let $\lambda \geq 1$ be arbitrary. Suppose that $n \geq 3$ satisfies $n \in \mathbb{L} \setminus \mathbb{T}(c, \lambda)$. We will freely (and implicitly) assume that n is large with respect to d and λ . Our goal is to show that if c is sufficiently small, it then necessarily follows that $\gamma \leq n$.

First consider the possibility that $\text{Gr}(3m) \geq 3^5 \text{Gr}(m)$ for all $m \in [n^{1/3}, n^{1/2}]$. Let $\eta := 1/100$, and let $c_1(d, 101, \eta) > 0$ be given by Lemma 8.3.11. Note that for all $m \in [n^{1/3+\eta}, (n^{1/3+\eta})^{1+\eta}]$, we have $[m^{1-\eta}, m^{1+\eta}] \subseteq [n^{1/3}, n^{1/2}]$ and $\text{Gr}(m) \leq \text{Gr}(n) \leq e^{[\log n]^{100}} \leq e^{[\log m]^{101}}$. So by construction of c_1 , we know that G has (c_1, λ) -polylog tubes at every scale $m \in [n^{1/3+\eta}, (n^{1/3+\eta})^{1+\eta}]$. In particular, if we pick $c(d) \leq c_1 \wedge \eta$, then $n \in \mathbb{T}(c, \lambda)$ - a contradiction.

So we may assume that there exists $m \in [n^{1/3}, n^{1/2}]$ such that $\text{Gr}(3m) \leq 3^5 \text{Gr}(m)$. Let $C(d, 5) < \infty$ be as given by Lemma 8.3.13. Without loss of generality, assume that C is an integer and $C \geq 2$. Let $A \subseteq [1, \infty)$ with $|A| \leq C$ be the set guaranteed to exist for our particular small-tripling scale m , and apply the conclusion of Lemma 8.3.13 with $k := (\log n)^\lambda$. Define $\varepsilon(d) := \frac{1}{3C} \log \frac{3}{2}$, and consider the sequence $(u_i : 1 \leq i \leq 3C)$ defined by

$$u_i := \left(n^{1/2} \right)^{(1+\varepsilon)^i}.$$

²⁵Although we have chosen to write everything for finite graphs, our proof also yields the analogous finitary refinement of Lemma 8.3.12 when G is infinite.

Note that thanks to our choice of ε ,

$$u_{3C} = \left(n^{1/2}\right)^{(1+\varepsilon)^{3C}} \leq \left(n^{1/2}\right)^{(e^\varepsilon)^{3C}} = n^{3/4}.$$

As we are assuming that n is large with respect to d and λ , we also have that $Ckm \leq u_1$, and for all $a \in A$, the interval $[a, 2ka]$ contains at most one of the u_i 's. So by the pigeonhole principle, there exists i such that $[u_i, u_{i+1}] \subseteq [Ckm, \infty) \setminus \bigcup_{a \in A} [a, 2ka]$. By construction of A , we know that either (1) for every scale $l \in [u_i, u_{i+1}]$, the graph G has $(C^{-1}k, C^{-1}k^{-1}m, Ck^Cm)$ -plentiful tubes at scale l , and in particular $(\frac{1}{2C}, \lambda)$ -polylog plentiful tubes at scale l , or (2) there exists $l \in [u_i, u_{i+1}]$ such that $\gamma \leq Cl$ and hence $\gamma \leq n$. If (1) holds, then by picking $c(d) \leq \varepsilon \wedge \frac{1}{2C}$, we can guarantee that $n \in \mathbb{T}(c, \lambda)$ - a contradiction. So (2) holds, i.e. $\gamma \leq n$ as required. \square

Cutsets and cycles

In this subsection we reduce Lemma 8.3.13 to Lemma 8.3.16, which is a less technical statement about cutsets and cycles. We will prove Lemma 8.3.16 in the next section.

Cutsets Let A, B, C be sets of vertices in a graph G . We write $A \xleftrightarrow{B} C$ to mean that there exists a finite path $(\gamma_k)_{k=0}^n$ such that $\gamma_0 \in A$; $\gamma_1, \dots, \gamma_{n-1} \in B$; and $\gamma_n \in C$, and we write $A \xleftrightarrow{B} \infty$ to mean that there exists an infinite self-avoiding path $(\gamma_k)_{k=0}^\infty$ such that $\gamma_0 \in A$ and $\gamma_1, \gamma_2, \dots \in B$. We write $\not\xleftrightarrow{B}$ to denote the negations of these properties. Now we say that B is an (A, C) -cutset to mean that $A \not\xleftrightarrow{B} C$, and we say that B is a *minimal* (A, C) -cutset if no proper subset of B is also an (A, C) -cutset. We extend all of these definitions in the obvious way to allow A or C to be vertices rather than set of vertices. Now suppose that G is an infinite transitive graph. Of course the spheres S_n for $n \in \mathbb{N}$ are all (o, ∞) -cutsets, but interestingly, they are not always minimal (o, ∞) -cutsets because some transitive graphs contain dead-ends, i.e. a vertex that is at least as far from o as all of its neighbours. The exposed sphere S_n^∞ is defined to be the unique minimal (o, ∞) -cutset contained in the usual sphere S_n , which is given concretely by

$$S_n^\infty = \{u \in S_n : u \xleftrightarrow{B_n^c} \infty\}.$$

Thanks to the following result of Funar, Giannoudovardi, and Otera [FGO15, Proposition 5], exposed spheres also admit the following finitary characterisation: S_n^∞ is the unique minimal (o, S_{2n+1}) -cutset contained in S_n . We have included the short and elegant proof from their paper for the reader to appreciate that it does not adapt well to finite graphs. Specifically, it does not yield Lemma 8.3.19, which is what we will need. We like to call this the *inflexible geodesic* argument.

Lemma 8.3.14. *Let G be an infinite transitive graph. Let $r \in \mathbb{N}$. Then every vertex $u \in B_{2r}^c$ satisfies $u \xrightarrow{B_r^c} \infty$.*

Proof of Lemma 8.3.14: The inflexible geodesic argument. This proof uses the well-known fact that every infinite transitive graph contains a bi-infinite geodesic $\gamma = (\gamma_n)_{n \in \mathbb{Z}}$. Here is a sketch of how to prove this: There exist geodesic segments $\gamma^N = (\gamma_n^N)_{n=0}^{2N}$ for every $N \in \mathbb{N}$. By transitivity, we can pick these with $\gamma_N^N = o$ for every N . Then γ is any local limit of these geodesic segments rooted at o , which exists by compactness.

Now fix $u \in B_{2r}^c$. Let $\gamma = (\gamma_n)_{n \in \mathbb{Z}}$ be a bi-infinite geodesic with $\gamma_0 = u$. Suppose for contradiction that there exist $s, t \in \mathbb{N}$ such that $\gamma_{-s}, \gamma_t \in B_r$. Note that $\text{dist}(\gamma_{-s}, \gamma_t) \leq \text{diam } B_r \leq 2r$. Since $\text{dist}(o, u) > 2r$, we have $\gamma_{-s}, \gamma_t \notin B_r(u)$. Since γ is a path, it follows that $s, t > r$, and in particular, $s+t > 2r$. On the other hand, since γ is a geodesic, $s+t = \text{dist}(\gamma_{-s}, \gamma_t)$. Therefore $2r < s+t \leq 2r$, a contradiction. So either $\{\gamma_n : n \geq 0\}$ or $\{\gamma_n : n \leq 0\}$ is disjoint from B_r and therefore forms a path witnessing that $u \xrightarrow{B_r^c} \infty$. \square

The usual definition of the exposed sphere is clearly inappropriate when working with finite transitive graphs. We propose that the exposed sphere in a finite transitive graphs should instead be defined according to this alternative finitary characterisation. Since Lemma 8.3.14 only applies to infinite transitive graphs, there is no reason for now that the reader should believe us that this is a good definition. We will fix this later by proving a finite graph analogue of Lemma 8.3.14, namely Lemma 8.3.19. As with (usual) spheres and balls, we extend the definition of exposed spheres to non-integer n by setting $S_n^\infty := S_{\lfloor n \rfloor}^\infty$.

Definition 8.3.15. Let G be a transitive graph, which may be finite or infinite. Let $n \in \mathbb{N}$. We define the exposed sphere S_n^∞ to be the unique minimal (o, S_{2n+1}) -cutset contained in S_n , or equivalently,

$$S_n^\infty := \{u \in S_n : u \xrightarrow{B_n^c} S_{2n+1}\}.$$

Cycles Let G be a graph. Recall that we identify spanning subgraphs of G with functions $E \rightarrow \{0, 1\}$. Pointwise addition and scalar multiplication of these functions makes the set of all spanning subgraphs into a $(\mathbb{Z}/2\mathbb{Z})$ -vector space. Recall that a cycle is finite path that starts and ends at the same vertex and visits no other vertex more than once. We identify cycles (ignoring orientation) with spanning subgraphs and hence with elements of this $(\mathbb{Z}/2\mathbb{Z})$ -vector space. Now let $\delta(G)$ be the minimal $n \in \mathbb{N}$ such that every cycle can be expressed as the linear combination of cycles having (extrinsic) diameter $\leq n$, if such an n exists, and set $\delta(G) := +\infty$ otherwise. It

is natural to ask whether cycles with diameter $\leq \delta(G)$ also generate every bi-infinite geodesic γ in the sense that there is a sequence $(\gamma_n)_{n=1}^\infty$ of cycles each having diameter $\leq \delta(G)$ such that $\gamma_n \rightarrow \gamma$ pointwise as $n \rightarrow \infty$. Notice that this is equivalent to G being one-ended. Benjamini and Babson [BB99b] (see also the proof by Timar [Tim07]) proved that if G is one-ended, then for all $u \in V$ and $v \in V \cup \{\infty\}$, every minimal (u, v) -cutset A is $\delta(G)$ -connected in the sense that $\text{dist}_G(A_1, A_2) \leq \delta(G)$ for every non-trivial partition $A = A_1 \sqcup A_2$. We will use this in the next section.

We claim that Lemma 8.3.13 can be reduced to the following statement about cutsets and cycles by applying the structure theory of groups and transitive graphs of polynomial growth. Since this step is essentially identical to the proof of [EH23b, Proposition 5.3], we have chosen to defer the details to the appendix. For the same reason, we will not give an overview of the rich theory of polynomial growth or even the definition of a virtually nilpotent group. The relevant background can be found in [EH23b; EH23d; TT21a].

Lemma 8.3.16. *Let $r, n \geq 1$. Let G be a finite transitive graph such that S_n^∞ is not r -connected. Let H be a (finite or infinite) transitive graph with $\delta(H) \leq r$ that does not have infinitely many ends. If $B_{50n}^H \cong B_{50n}^G$, then*

$$\text{dist}_{\text{GH}}\left(\frac{\pi}{\text{diam } G} G, S^1\right) \leq \frac{200n}{\text{diam } G}.$$

Solving the reduced problem

In this section we prove Lemma 8.3.16. Benjamini and Babson [BB99b] tell us that if in an infinite transitive graph H , the exposed sphere S_n^∞ is not $\delta(H)$ connected²⁶, then H must not be one-ended. If H is not infinitely-ended either, then H must in fact be two-ended and hence one-dimensional. This is how the argument (implicitly) went in [EH23b]. To make this more finitary, let us start by noting that the proof that H not one-ended actually also tells us that this is witnessed by S_n^∞ itself in the sense that $H \setminus S_n^\infty$ has multiple infinite components. Equivalently (by Lemma 8.3.14), the exposed sphere S_n^∞ disconnects²⁷ S_{2n+1} . (This alternative phrasing has the benefit that it also makes sense when H is finite.) Indeed, this follows from the next lemma with $(A, B) := (\{o\}, S_{2n+1})$. This is also an instance of Benjamini-Babson, just phrased slightly differently in terms of *sets of vertices*, *vertex cutsets*, and (extrinsic) *diameter* rather than length of generating cycles. For completeness,

²⁶meaning that there is a non-trivial partition $\delta(H) = A_1 \sqcup A_2$ with $\text{dist}(A_1, A_2) > \delta(H)$

²⁷We say that a set of vertices A disconnects another set of vertices B if there exist vertices $b_1, b_2 \in B$ such that $b_1 \xrightarrow{A^c} b_2$.

we have written Timar's proof of Benjamini-Babson with the necessary tiny adjustments in the appendix.

Lemma 8.3.17. *Let G be a graph. Let A and B be sets of vertices. Let Π be a minimal (A, B) -cutset that does not disconnect A or B . Then Π is $\delta(G)$ -connected.*

The following elementary lemma lets us conclude from this that H must begin to look one-dimensional *already from scale n* . We say that a path $\gamma = (\gamma_t : t \in I)$ is n -dense if $\sup_{v \in V} \text{dist}_G(v, \gamma) \leq n$, where $\text{dist}_G(v, \gamma) := \text{dist}_G(v, \{\gamma_t : t \in I\})$.

Lemma 8.3.18. *Let G be an infinite transitive graph. Let $n \geq 1$. If G is two-ended and $G \setminus B_n$ has two infinite components, then G contains an n -dense bi-infinite geodesic.*

Proof. Let A and B be the two infinite components of $G \setminus B_n$. For each integer $N \geq n + 1$, let $\gamma^N = (\gamma_t^N : -a_N \leq t \leq b_N)$ be a shortest path among those that start in $S_N \cap A$ and end in $S_N \cap B$, indexed such that $\gamma_0^N \in B_n$. By compactness, there exists a bi-infinite geodesic $\gamma = (\gamma_t : t \in \mathbb{Z})$ and a subsequence $(\gamma^N : N \in \mathbb{M})$ such that for every $t \in \mathbb{Z}$, we have $\gamma_t = \gamma_t^N$ for all sufficiently large $N \in \mathbb{M}$. As in the *inflexible geodesic* argument used to prove Lemma 8.3.14 (i.e. by the triangle inequality), a geodesic can never visit B_{2n}^c in between two visits to B_n . It follows that there exists t_0 such that $\gamma^- := (\gamma_{-t} : t \geq t_0)$ is entirely contained in A , and $\gamma^+ := (\gamma_t : t \geq t_0)$ is entirely contained in B .

Suppose for contradiction that γ is not n -dense. Pick $u \in V$ with $\text{dist}(u, \gamma) > n$. Since $B_n(u)$ does not intersect γ , the path γ must be entirely contained in one of the two infinite components of $G \setminus B_n(u)$, say C . Since $B_n(o)$ disconnects γ^- from γ^+ , there are at least two infinite components in $C \setminus B_n(o)$. So there are at least three infinite components in $G \setminus (B_n(o) \cup B_n(u))$, contradicting the fact that G is two-ended. \square

What happens if instead H is finite? Lemma 8.3.17 still tells us that if S_n^∞ is not $\delta(H)$ connected, then S_n^∞ disconnects S_{2n+1} . When H was infinite, this had a nice interpretation in terms of ends because we could go back to the original infinitary definition of S_n^∞ as a minimal (o, ∞) -cutset. The problem when H is finite is that we are stuck with our artificial finitary definition of S_n^∞ as a minimal (o, S_{2n+1}) -cutset. The next lemma justifies our definition by establishing that S_n^∞ is automatically a minimal (o, u) -cutset for every vertex $u \in B_{2n}^c$. Thanks to this lemma, it is simply impossible that S_n^∞ is not $\delta(H)$ -connected when H is finite.

The analogous statement for one-ended infinite transitive graphs follows from Lemma 8.3.14²⁸. In

²⁸This was the motivation for [FGO15] to prove Lemma 8.3.14.

this sense, Lemma 8.3.19 lets us treat finite transitive graphs as if they were infinite transitive graphs that are one-ended. Note that a naive finite-graph adaptation of the inflexible geodesic argument used to prove Lemma 8.3.14 would not yield Lemma 8.3.19. (It would just say that every vertex in S_{2n+1} belongs to a cluster in $G \setminus B_n$ of large diameter.) Our argument also yields a new proof of Lemma 8.3.14.

Lemma 8.3.19. *Let G be a finite transitive graph. Let $r \in \mathbb{N}$. Then B_r does not disconnect B_{2r}^c .*

Proof of Lemma 8.3.14 and Lemma 8.3.19. Suppose that B_r does disconnect B_{2r}^c . (For Lemma 8.3.19 we assume this for sake of contradiction, whereas for Lemma 8.3.14, we may assume this otherwise the conclusion is trivial.) Let C be a component of $G \setminus B_r$ intersecting B_{2r}^c . It suffices to prove that C is infinite. (For Lemma 8.3.19, this establishes the required contradiction because G is finite, whereas for Lemma 8.3.14, this is the desired conclusion.)

Suppose for contradiction that C is finite. Then we can pick a vertex $u \in C$ maximising $\text{dist}(o, u)$. Since $\text{dist}(o, u) \geq 2r + 1$ and (by transitivity) $B_r(u)$ disconnects $B_{2r}(u)^c$, there exists a vertex $v \in B_{2r}(u)^c$ such that $o \xrightarrow{B_r(u)^c} v$. Since $B_r(u) \cap B_r(o) = \emptyset$ and the subgraph induced by $B_r(o)$ is connected, $B_r(o) \xrightarrow{B_r(u)^c} v$. Since G is connected, $v \xleftarrow{B_r(o)^c} B_r(u)$. Since $B_r(u) \cap B_r(o) = \emptyset$ and the subgraph induced by $B_r(u)$ is connected, $v \xleftarrow{B_r(o)^c} u$, i.e. $v \in C$. However, since every path from o to v must visit $B_r(u)$,

$$\begin{aligned} \text{dist}(o, v) &\geq \text{dist}(o, B_r(u)) + \text{dist}(B_r(u), v) \\ &\geq (\text{dist}(o, u) - r) + (\text{dist}(u, v) - r) \geq \text{dist}(o, u) + 1, \end{aligned}$$

contradicting the maximality of $\text{dist}(o, u)$. \square

By applying our work up to this point, under the hypothesis of Lemma 8.3.16, we can prove that the graph H must be infinite and begin to look one-dimensional from around scale n . By the next lemma, it follows that G looks like a circle from around scale n .

Lemma 8.3.20. *Let G and H be transitive graphs. Suppose that G is finite whereas H contains an n -dense bi-infinite geodesic for some $n \geq 1$. If $B_{50n}^G \cong B_{50n}^H$ then*

$$\text{dist}_{GH} \left(\frac{\pi}{\text{diam } G} G, S^1 \right) \leq \frac{200n}{\text{diam } G}.$$

Proof. Let $\gamma = (\gamma_t)_{t \in \mathbb{Z}}$ be an n -dense bi-infinite geodesic in H . Without loss of generality, assume that $\gamma_0 = o_H$. We will break our proof into a sequence of small claims.

Claim. B_{2n}^H disconnects $S_{2n}^{\infty, H}$.

Proof of claim. Let $\zeta = (\zeta_t)_{t=0}^k$ be an arbitrary path from $\zeta_0 = \gamma_{-2n}$ to $\zeta_k = \gamma_{2n}$. Since γ is n -dense, for all $t \in \{0, \dots, k\}$, there exists $g_t \in \mathbb{Z}$ such that $\text{dist}(\zeta_t, \gamma_{g_t}) \leq n$. We can of course require that $g_0 := -2n$ and $g_k := 2n$. Since γ is geodesic, for all $t \in \{0, \dots, k-1\}$,

$$\begin{aligned} |g_{t+1} - g_t| &= \text{dist}(\gamma_{g_{t+1}}, \gamma_{g_t}) \\ &\leq \text{dist}(\gamma_{g_t}, \zeta_t) + \text{dist}(\zeta_t, \zeta_{t+1}) + \text{dist}(\zeta_{t+1}, \gamma_{g_{t+1}}) \\ &\leq n + 1 + n = 2n + 1. \end{aligned}$$

In particular, since $g_0 \leq -(n+1)$ but $g_k \geq n+1$, there must exist $t \in \{1, \dots, k-1\}$ such that $-n \leq g_t \leq n$. Then $\gamma_{g_t} \in B_n^H$ and hence $\zeta_t \in B_{2n}^H$. Since ζ was arbitrary, this establishes that B_{2n}^H disconnects γ_{-2n} from γ_{2n} . Since γ is a geodesic, $\gamma_s \in S_s^H$ for all $s \in \mathbb{Z}$. So the path $(\gamma_s : s \geq 2n)$ witnesses the fact that $\gamma_{2n} \in S_{2n}^{\infty, H}$. Similarly, $\gamma_{-2n} \in S_{2n}^{\infty, H}$. So B_{2n}^H disconnects $S_{2n}^{\infty, H}$. \square

Fix a non-trivial partition $S_{2n}^{\infty, H} = A \sqcup B$ such that B_{2n}^H is an (A, B) -cutset. Now suppose that there is a graph isomorphism $\psi : B_{50n}^G \rightarrow B_{50n}^H$. Note that ψ induces a bijection $S_{2n}^{\infty, G} \leftrightarrow S_{2n}^{\infty, H}$. In particular, $S_{2n}^{\infty, G} = \psi(A) \sqcup \psi(B)$. By definition of exposed sphere, B_{2n}^G does not disconnect $\psi(A)$ or $\psi(B)$ from $(B_{4n}^G)^c$. So by Lemma 8.3.19, B_{2n}^G is not a $(\psi(A), \psi(B))$ -cutset. Consider a shortest path from $\psi(A)$ to $\psi(B)$ that witnesses this, then connect the start and end of this path to o_G by geodesics. Let $\lambda = (\lambda_k)_{k \in \mathbb{Z}_l}$ be the resulting cycle, labelled such that $\lambda_0 = o_G$. We will write $|s|_l$ for the distance from s to 0 in the cycle graph \mathbb{Z}_l . The next three claims establish that λ is roughly dense and geodesic.

Claim. For all $s \in \mathbb{Z}_l$, if $|s|_l > 2n$ then $\text{dist}_G(\lambda_s, B_{2n}^G) = |s|_l - 2n$

Proof of claim. Fix $s \in \mathbb{Z}_l$ with $|s|_l > 2n$. Since B_{2n}^H is an (A, B) -cutset, the segment $(\lambda_t : |t|_l > 2n)$ must intersect $(B_{4n}^G)^c$ (it must exit the ball B_{50n} , on which G and H are isomorphic), but by construction, this segment does not intersect B_{2n}^G . So every path from γ_s to B_{2n}^G must intersect $S_{2n}^{\infty, G}$. In particular, by minimality in the construction of λ ,

$$\text{dist}_G(\gamma_s, B_{2n}^G) = \text{dist}_G(\gamma_s, \psi(A)) \wedge \text{dist}_G(\gamma_s, \psi(B)) = |s|_l - 2n. \quad \square$$

Claim. For all $s, t \in \mathbb{Z}_l$, we have $|s - t|_l - 4n \leq \text{dist}_G(\lambda_s, \lambda_t) \leq |s - t|_l$.

Proof of claim. The second inequality is trivial, and the first inequality is trivial when $|s|_l \vee |t|_l \leq 2n$. By our previous claim, if $|s|_l > 2n$ and $|t|_l \leq 2n$, then

$$\text{dist}_G(\lambda_s, \lambda_t) \geq \text{dist}_G(\lambda_s, B_{2n}^G) = |s|_l - 2n \geq |s - t|_l - 4n.$$

Similarly, the first inequality also holds if instead $|t|_l > 2n$ and $|s|_l \leq 2n$. So let us consider s and t satisfying $|s|_l \wedge |t|_l > 2n$, and fix an arbitrary path η from λ_s to λ_t . If η intersects B_{2n}^G , then by our previous claim, η has length at least

$$\text{dist}_G(\lambda_s, B_{2n}^G) + \text{dist}_G(\lambda_t, B_{2n}^G) = (|s|_l - 2n) + (|t|_l - 2n) \geq |s - t|_l - 4n.$$

If η does not intersect B_{2n}^G , then by minimality in the construction of λ , the length of η is at least $|s - t|_l$. Either way, $\text{dist}_G(\lambda_s, \lambda_t) \geq |s - t|_l - 4n$. \square

Claim. λ is $10n$ -dense

Proof of claim. Suppose for contradiction that u is a vertex with $\text{dist}_G(u, \lambda) > 10n$. Let v be a vertex in λ that is closest to u . Let z be a vertex in $S_{10n}(v)$ that lies along a geodesic from u to v . Since λ visits o_G but must exit B_{50n}^G (on which G and H are isomorphic), we know that λ has (extrinsic) diameter $> 50n$. By our previous claim, it follows that λ visits vertices x and y in $S_{10n}^G(v)$ satisfying $\text{dist}_G(x, y) \geq 2 \cdot 10n - 4n \geq 10n$. Now x, y, z are three vertices in $S_{10n}^G(v)$ such that the distance between any pair is at least $10n$. By transitivity and the fact that $B_{50n}^G \cong B_{50n}^H$, three such vertices can also be found in $S_{10n}^H(o)$, say v_1, v_2, v_3 . Since γ is n -dense, there exist integers k_1, k_2, k_3 such that $\text{dist}_H(v_i, \gamma_{k_i}) \leq n$ for each $i \in \{1, 2, 3\}$. Notice that since $\gamma_0 = o_H$ and γ is geodesic, $|k_i| \in [9n, 11n]$ for all i . So by the pigeonhole principle, either $[-9n, -11n]$ or $[9n, 11n]$ contains k_i for at least two distinct values of i . On the other hand, for all $i \neq j$, since γ is geodesic,

$$|k_i - k_j| = \text{dist}_H(\gamma_{k_i}, \gamma_{k_j}) \geq \text{dist}_H(v_i, v_j) - 2n \geq 8n.$$

So an interval of width $2n$ can never contain k_i for at least two distinct values of i , a contradiction. \square

Thanks to the previous two claims, the map $\mathbb{Z}_l \rightarrow G$ sending $t \mapsto \lambda_t$ is a $(1, 10n)$ -quasi-isometry. So (by exercise 5.10 (b) in [Pet23], for example), $\text{dist}_{\text{GH}}(\mathbb{Z}_l, G) \leq 10n$. By the obvious 1-dense isometric embedding of \mathbb{Z}_l into $\frac{l}{2\pi}S^1$, we know that $\text{dist}_{\text{GH}}(\mathbb{Z}_l, \frac{l}{2\pi}S^1) \leq 1$. Let $D := \text{diam } G$. By the previous two claims $|D - \frac{l}{2}| \leq 20n$. So by considering the identity map from G to itself,

$$\begin{aligned} \text{dist}_{\text{GH}}\left(\frac{1}{D}G, \frac{2}{l}G\right) &\leq \sup_{u, v \in V(G)} \left| \frac{1}{D} \text{dist}_G(u, v) - \frac{2}{l} \text{dist}_G(u, v) \right| \\ &\leq D \cdot \left| \frac{1}{D} - \frac{2}{l} \right| \leq \frac{40n}{l}. \end{aligned}$$

Putting these bounds together,

$$\begin{aligned}
\text{dist}_{\text{GH}}\left(\frac{\pi}{D}G, S^1\right) &\leq \text{dist}_{\text{GH}}\left(\frac{\pi}{D}G, \frac{2\pi}{l}G\right) + \text{dist}_{\text{GH}}\left(\frac{2\pi}{l}G, \frac{2\pi}{l}\mathbb{Z}_l\right) + \text{dist}_{\text{GH}}\left(\frac{2\pi}{l}\mathbb{Z}_l, S^1\right) \\
&\leq \pi \cdot \frac{40n}{l} + \frac{2\pi}{l} \cdot 10n + \frac{2\pi}{l} \cdot 1 \\
&\leq \frac{200n}{l} = 100n \cdot \frac{2}{l} \leq 100n \cdot \frac{1}{D-20n} = 5 \cdot \frac{20n/D}{1-20n/D}.
\end{aligned} \tag{8.3.6}$$

We may assume that $D \geq 40n$, other the conclusion of the lemma holds trivially because $\text{dist}_{\text{GH}}(A, B) \leq 1$ for all non-empty compact metric spaces A and B each having diameter at most 1. In particular, $\frac{20n}{D} \leq \frac{1}{2}$. Since $\frac{x}{1-x} \leq 2x$ for all $x \in [0, 1/2]$, it follows from eq. (8.3.6) that $\text{dist}_{\text{GH}}\left(\frac{\pi}{D}G, S^1\right) \leq 5 \cdot 2 \cdot 20n/D = 200n/D$ as required. \square

We now combine these lemmas to formalise this sketch of a proof of Lemma 8.3.16, thereby concluding our proof of Proposition 8.3.1.

Proof of Lemma 8.3.16. Suppose that $B_{50n}^G \cong B_{50n}^H$. Note that $S_n^{\infty, G}$ is trivially $2n$ -connected. So $r \leq 2n$. In particular, in any transitive graph, the statement “ S_n^{∞} is not r -connected” is determined by the subgraph induced by B_{50n} . So $S_n^{\infty, H}$ is not r -connected either. By definition, $S_n^{\infty, H}$ is a minimal (o_H, S_{2n+1}^H) -cutset. So by Lemma 8.3.17, since $\delta(H) \leq r$, the exposed sphere $S_n^{\infty, H}$ must disconnect S_{2n+1}^H . In particular, B_n^H disconnects $(B_{2n}^H)^c$. So by Lemma 8.3.19, H is infinite, and by Lemma 8.3.14, $H \setminus B_n^H$ contains at least two infinite components. Since H has at most finitely many ends, $H \setminus B_n^H$ must contain exactly two infinite components and H must be exactly two-ended. So by Lemma 8.3.18, H contains an n -dense bi-infinite geodesic. The conclusion follows by Lemma 8.3.20. \square

8.4 Global connections → unique large cluster

In this section we apply the methods of [EH21a]. It will be convenient to adopt the following notation from that paper: given a set of vertices A in a graph G , we define its *density* to be $\|A\| := \frac{|A|}{|V(G)|}$. In [EH21a], together with Hutchcroft, we showed that the supercritical giant cluster for percolation on bounded-degree finite transitive graphs is always unique with high probability. More precisely, for every infinite set \mathcal{G} of finite transitive graphs with bounded degrees, for every supercritical sequence of parameters p , and for every constant $\varepsilon > 0$, the density of the second largest cluster $\|K_2\|$ satisfies

$$\lim \mathbb{P}_p (\|K_2\| \geq \varepsilon) = 0.$$

The following proposition contains a quantitative version of this statement that is useful even if we slightly weaken the hypothesis that p is supercritical. We think of this as saying that if at some parameter p we have a point-to-point lower bound that is only slightly worse than constant as $|V| \rightarrow \infty$, then after passing to $p + \varepsilon$, we can still pretend that we are actually in the supercritical phase and still prove that the second largest cluster is typically much smaller than the largest cluster. Note that this largest cluster is not necessarily a giant cluster because we are not (a priori) really in the supercritical phase.²⁹

Proposition 8.4.1. *Let G be a finite transitive graph with degree d . Define $\delta := (\log |V|)^{-1/20}$. There exists $C(d) < \infty$ such that if $|V| \geq C$, then for all $p, q \in (0, 1)$ with $q - p \geq \delta$,*

$$\min_{u,v \in V} \mathbb{P}_p(u \leftrightarrow v) \geq 2\delta \implies \mathbb{P}_q\left(\|K_1\| \geq \delta \text{ and } \|K_2\| \leq \delta^2\right) \geq 1 - \delta^4.$$

In Section 8.4 we will explain why this proposition is implied by the *sandcastles*³⁰ argument of [EH21a]. In fact, [EH21a] already explicitly contains a very similar quantitative statement, namely [EH21a, Theorem 1.5]. Unfortunately this statement is not quantitatively strong enough for our purposes. One could alternatively prove a version of Proposition 8.4.1 by applying the ghost-field technology developed in [EH23b, Section 4]. (See the discussion at the end of [EH23b, Section 7.1].) This approach would be less elementary and less generalisable³¹ but quantitatively stronger.

Proof via sandcastles

Let G be a finite transitive graph. In [EH21a] we made the definition of “supercritical sequence” finitary as follows. Given a constant $\varepsilon > 0$, we say that a parameter $p \in (0, 1)$ is ε -supercritical if

$$\mathbb{P}_{(1-\varepsilon)p}(\|K_1\| \geq \varepsilon) \geq \varepsilon$$

and $|V| \geq 2\varepsilon^{-3}$, the latter being a technical condition that the reader may like to ignore. Note that a sequence of parameters is supercritical if and only if there exists a constant $\varepsilon > 0$ such that all but finitely many of the parameters are ε -supercritical. On the other hand, in the present paper the more relevant finitary notion of supercriticality concerns point-to-point connection probabilities, i.e.

$$\min_{u,v} \mathbb{P}_{(1-\varepsilon)p}(u \leftrightarrow v) \geq \varepsilon.$$

²⁹This is reminiscent of [EH23b, Section 6]. There we used the hypothesis of a point-to-point lower bound on a large scale to enable us to run arguments from [CMT22], which were ostensibly about supercritical percolation, to study subcritical percolation.

³⁰We thank Coales for suggesting this name.

³¹The ghost-field arguments ultimately rely on two-arm bounds, which are not elementary and which break down when working with graphs with rapidly diverging vertex degrees.

These properties are equivalent up to changing the constant ε . Indeed, for every parameter $p \in (0, 1)$ and every constant $\varepsilon > 0$ satisfying the technical condition $|V| \geq 2\varepsilon^{-3}$,

$$\min_{u,v} \mathbb{P}_p(u \leftrightarrow v) \geq 2\varepsilon \implies \mathbb{P}_p(\|K_1\| \geq \varepsilon) \geq \varepsilon \implies \min_{u,v} \mathbb{P}_p(u \leftrightarrow v) \geq e^{-10^5\varepsilon^{-18}}. \quad (8.4.1)$$

The first implication is an easy application of Markov's inequality, and the second implication is [EH21a, Theorem 2.1]. A version of the second implication assuming an upper bound on the degree of G is originally due to Schramm. Notice that the second implication quantitatively loses much more than the first. In this sense, we can think of the hypothesis " $\min_{u,v} \mathbb{P}_{(1-\varepsilon)p}(u \leftrightarrow v) \geq \varepsilon$ " as being quantitatively much stronger than the hypothesis " $\mathbb{P}_{(1-\varepsilon)p}(\|K_1\| \geq \varepsilon) \geq \varepsilon$ ".

Below is [EH21a, Theorem 1.5], which contains a finitary uniqueness statement similar to Proposition 8.4.1. Unfortunately, the terrible $e^{-C\varepsilon^{-18}}$ dependence on ε is not good enough for our purposes. Fortunately, it turns out that in the proof of this theorem, the source of this poor dependence is a conversion from the hypothesis of a giant cluster bound (implicit in p being ε -supercritical) into a point-to-point bound, i.e. an application of the second implication in eq. (8.4.1). This saves us because in the present setting we actually start with the "stronger" hypothesis of a point-to-point bound.

Theorem 8.4.2. *Let G be a finite transitive graph with degree d . There exists $C(d) < \infty$ such that for every $\varepsilon > 0$, every ε -supercritical parameter p , and every $\lambda \geq 1$,*

$$\mathbb{P}_p \left(\|K_2\| \geq \lambda e^{C\varepsilon^{-18}} \left(\frac{\log d}{\log |V|} \right)^{1/2} \right) \leq \frac{1}{\lambda}.$$

A key ingredient in the sandcastles argument of [EH21a] is the sharp density property, which measures the extent to which the events $\{\|K_1\| \geq \alpha\}$ for each $\alpha \in (0, 1)$ have uniformly-in- α sharp thresholds. Let $\Delta : (0, 1) \rightarrow (0, 1/2]$ be a decreasing function. For all $\alpha, \delta \in (0, 1)$, let $p_c(\alpha, \delta) \in (0, 1)$ be the parameter satisfying $\mathbb{P}_{p_c(\alpha, \delta)}(\|K_1\| \geq \alpha) = \delta$, which is unique by the strict monotonicity of this probability with respect to p . We say that G has the Δ -sharp density property if for all $\alpha \in (0, 1)$ and $\delta \in [\Delta(\alpha), 1/2]$,

$$\frac{p_c(\alpha, 1 - \delta)}{p_c(\alpha, \delta)} \leq e^\delta.$$

The following lemma establishes a sharp density property for graphs with bounded degrees. This is [EH21a, Proposition 3.2] and is an easy consequence of Talagrand's well-known sharp threshold theorem [Tal94].

Lemma 8.4.3. *Let G be a finite transitive graph with degree d . There exists $C(d) < \infty$ such that G has the Δ -sharp density property for the function $\Delta : (0, 1) \rightarrow (0, 1/2]$ given by*

$$\Delta(\alpha) := \begin{cases} \frac{1}{2} \wedge \frac{C}{(\log|V|)^{1/2}} & \text{if } \alpha \geq \left(\frac{2}{|V|}\right)^{1/3} \\ \frac{1}{2} & \text{otherwise.} \end{cases}$$

The sandcastles argument combines a sharp density property with a point-to-point bound to establish the uniqueness of the largest cluster. Here is the technical output of that argument.

Lemma 8.4.4. *Let G be a finite transitive graph. Let $\varepsilon \in (0, 1)$ and suppose that $p \in (0, 1)$ is ε -supercritical. Suppose that G satisfies the Δ -sharp density property for some decreasing function $\Delta : (0, 1) \rightarrow (0, 1/2]$. Then for all $\lambda \geq 1$,*

$$\mathbb{P}_p \left(\|K_2\| \geq \lambda \left(\frac{200\Delta(\varepsilon)}{\varepsilon^3 \tau} + \frac{25}{\varepsilon^2 \tau |V|} \right) \right) \leq \frac{\varepsilon}{\lambda},$$

where

$$\tau := \min_{u,v} \mathbb{P}_{(1-\varepsilon)p}(u \leftrightarrow v).$$

Proof. [EH21a, Theorem 3.3] is the same statement but where τ is instead defined to be

$$\tau := e^{-10^5 \varepsilon^{-18}},$$

which is the function appearing in eq. (8.4.1). We claim that the proof of [EH21a, Theorem 3.3] actually also establishes Lemma 8.4.4. First note that in the statement of [EH21a, Lemma 3.5], we can require that $q \in ((1 - \varepsilon)p, p)$ rather than just $q \in (p_c(\varepsilon, \varepsilon), p)$. Indeed, the exact same proof works, using $q_j := e^{j\Delta(\varepsilon)}(1 - \varepsilon)p$ instead of $q_j := e^{j\Delta(\varepsilon)}p_c(\varepsilon, \varepsilon)$, because $(1 - \varepsilon)p \geq p_c(\varepsilon, \varepsilon)$. So in the proof of [EH21a, Theorem 3.3], we may assume that the parameter called q , which is provided by [EH21a, Lemma 3.5], satisfies $q \geq (1 - \varepsilon)p$. In particular, when we later apply [EH21a, Theorem 2.1] to lower bound $\min_{u,v} \mathbb{P}_q(u \leftrightarrow v)$ by $e^{-10^5 \varepsilon^{-18}}$, we could instead simply lower bound $\min_{u,v} \mathbb{P}_q(u \leftrightarrow v)$ by $\min_{u,v} \mathbb{P}_{(1-\varepsilon)p}(u \leftrightarrow v)$. Running the rest of the proof of [EH21a, Theorem 3.3] exactly as written, except for the new definition “ $\tau := \min_{u,v} \mathbb{P}_{(1-\varepsilon)p}(u \leftrightarrow v)$ ” in place of “ $\tau := e^{-10^5 \varepsilon^{-18}}$ ”, yields the desired conclusion. \square

The uniqueness part of Proposition 8.4.1 will follow from Lemmas 8.4.3 and 8.4.4. The existence part will follow from the following well-known and (again) easy consequence of Talagrand’s sharp threshold theorem [Tal94].

Lemma 8.4.5. *Let G be a finite transitive graph. Let A be a non-trivial increasing event that is invariant under all graph automorphisms of G . Let $0 < p_1 < p_2 < 1$ and set $\delta := p_2 - p_1$. There exists a universal constant $c > 0$ such that*

$$\mathbb{P}_{p_1}(A) \leq \frac{1}{|V|^{c\delta}} \quad \text{or} \quad \mathbb{P}_{p_2}(A) \geq 1 - \frac{1}{|V|^{c\delta}}.$$

Proof. For every edge e , let $\text{Orb}(e)$ denote the orbit of e under the action of the automorphism group of G . By [EH21a, Theorem 3.10], there is a universal constant $c_1 > 0$ such that for all $p \in (0, 1)$, the function $f(p) := \mathbb{P}_p(A)$ satisfies

$$f'(p) \geq \frac{c_1}{p(1-p) \log \frac{2}{p(1-p)}} \cdot f(p) (1-f(p)) \cdot \log \left(2 \min_{e \in E} |\text{Orb}(e)| \right).$$

Since G is (vertex-)transitive, $|\text{Orb}(e)| \geq \frac{|V|}{2}$ for every $e \in E$. Also, by calculus, $\sup_{p \in (0,1)} p(1-p) \log \frac{2}{p(1-p)} < \infty$. Therefore, there is another universal constant $c > 0$ such that for all $p \in (0, 1)$,

$$\left[\log \frac{f}{1-f} \right]' = \frac{f'}{f(1-f)} \geq 2c \log |V|.$$

The result follows by integrating this differential inequality. \square

Proof of Proposition 8.4.1. Suppose that $q, p \in (0, 1)$ satisfy $q - p \geq \delta$ and $\min_{u,v} \mathbb{P}_p(u \leftrightarrow v) \geq 2\delta$. We will assume throughout this proof that $|V|$ is as large as we like with respect to d . Let us start with the existence of a large cluster. Let $c > 0$ be the universal constant from Lemma 8.4.5. By the first implication in eq. (8.4.1), we know that $\mathbb{P}_p(\|K_1\| \geq \delta) \geq \delta$. Since $|V|$ is large,

$$\delta = e^{-\frac{1}{20} \log \log |V|} \geq e^{-c(\log |V|)(\log |V|)^{-1/20}} = |V|^{-c\delta}. \quad (8.4.2)$$

So by Lemma 8.4.5 with $A := \{\|K_1\| \geq \delta\}$, since $\frac{\delta^4}{2} \geq |V|^{-c\delta}$ (by a calculation like eq. (8.4.2)),

$$\mathbb{P}_q(\|K_1\| \geq \delta) \geq 1 - \frac{1}{|V|^{c\delta}} \geq 1 - \frac{\delta^4}{2}. \quad (8.4.3)$$

We now turn to the uniqueness of the largest cluster. The parameter q is $(\delta/2)$ -supercritical because $|V| \geq 2(\delta/2)^{-3}$ and

$$\mathbb{P}_{(1-\delta/2)(p+\delta)}(\|K_1\| \geq \delta) \geq \mathbb{P}_p(\|K_1\| \geq \delta) \geq \delta.$$

By Lemma 8.4.3, since $\delta/2 \geq (2/|V|)^{1/3}$, there is a constant $C_1(d) < \infty$ such that G has the Δ -sharp density property for some Δ satisfying

$$\Delta(\delta/2) \leq \frac{C_1}{(\log |V|)^{1/2}}.$$

So by Lemma 8.4.4, there is a constant $C_2(d) < \infty$ such that for every $\lambda \geq 1$,

$$\mathbb{P}_q \left(\|K_2\| \geq \frac{C_2 \lambda}{\delta^4 (\log |V|)^{1/2}} \right) \leq \mathbb{P}_q \left(\|K_2\| \geq \lambda \left[\frac{200 \frac{C_1}{(\log |V|)^{1/2}}}{(\delta/2)^3 \cdot (2\delta)} + \frac{25}{(\delta/2)^2 \cdot (2\delta) \cdot |V|} \right] \right) \leq \frac{\delta/2}{\lambda}.$$

By picking λ such that $\frac{C_2 \lambda}{\delta^4 (\log |V|)^{1/2}} = \delta^2$ (which satisfies $\lambda \geq 1$ when $|V|$ is large), it follows that

$$\mathbb{P}_q (\|K_2\| \geq \delta^2) \leq \frac{C_2}{2\delta^5 (\log |V|)^{1/2}} = \frac{C_2 \delta^5}{2} \leq \frac{\delta^4}{2}. \quad (8.4.4)$$

The conclusion follows by combining eqs. (8.4.3) and (8.4.4) with a union bound. \square

8.5 Unique large cluster \rightarrow giant cluster

In this section we apply the methods of [Eas23]. We will again use the notation $\|K_1\|, \|K_2\|$ introduced in Section 8.4. Let \mathcal{G} be an infinite set of finite transitive graphs with possibly unbounded degrees. Recall that \mathcal{G} is said to have a percolation threshold if there is a fixed sequence $p_c : \mathcal{G} \rightarrow (0, 1)$ such that for every sequence $p : \mathcal{G} \rightarrow (0, 1)$, if $\limsup p/p_c < 1$ then $\lim \mathbb{P}_p (\|K_1\| \geq \varepsilon) = 0$ for all $\varepsilon > 0$, and if $\liminf p/p_c > 1$ then $\lim \mathbb{P}_p (\|K_1\| \geq \varepsilon) = 1$ for some $\varepsilon > 0$. In [Eas23] we showed that \mathcal{G} has a percolation threshold unless and only unless \mathcal{G} contains a very particular family of pathological sequences of dense graphs. This might appear to be simply a matter of proving that some nice event has a sharp threshold, perhaps by a simple application of Lemma 8.4.5 in the bounded-degree case. The subtle problem is that “ $\{K_1 \text{ is a giant}\}$ ” is not an event. Really the challenge is to prove that multiple events of the form $\{\|K_1\| \geq \alpha\}$, for different choices of α , all have sharp thresholds that in fact coincide with each other.

The bulk of our proof consisted in proving that if the supercritical giant cluster for \mathcal{G} is unique (as given by [EH21a]), then we can embed this fact into Vanneuville’s new proof of the sharpness of the phase transition for infinite transitive graphs [Van23; Van24] to deduce a kind of mean-field lower bound for the supercritical giant cluster density. This mean-field-like lower bound implies that for every $\delta > 0$ and sequence p , if there is a giant whose density exceeds some constant $\alpha > 0$ at p , i.e. $\lim \mathbb{P}_p (\|K_1\| \geq \alpha) = 1$, then there is a giant whose density exceeds some constant $c(\delta) > 0$ at $(1 + \delta)p$, i.e. $\lim \mathbb{P}_{(1+\delta)p} (\|K_1\| \geq c(\delta)) = 1$, where, crucially, $c(\delta)$ is independent of α . By a diagonalisation argument, it is clear that α can be allowed to decay slowly rather than remain constant. However, in general, the slowest allowable rate of decay can be arbitrarily slow³². This is why there was no discussion of rates of convergence in [EH23b]. Luckily, this is not the case in our restricted setting where graphs have bounded degrees.

³²Consider sequences that approximate sequences that do not have percolation thresholds.

In the following subsection we simply note that the argument in [EH23b] is fully quantitative in the sense that α can decay at any particular rate provided that we supply a sufficiently strong bound on the uniqueness of the largest (possibly non-giant) cluster. This is the content of the following proposition.

Proposition 8.5.1. *Let G be a finite transitive graph. Let $p, \varepsilon \in (0, 1)$ and $\alpha \in (0, \frac{\varepsilon}{10})$. There is a universal constant $c > 0$ such that the following holds whenever $|V|^{c\varepsilon} \geq \frac{1}{c\varepsilon}$.*

If

$$\mathbb{P}_{p-\varepsilon}(\|K_o\| \geq \alpha) \geq \alpha \quad \text{and} \quad \mathbb{P}_p\left(\|K_1\| \geq \alpha \text{ and } \|K_2\| < \frac{\alpha}{2}\right) > 1 - \alpha^2,$$

then

$$\mathbb{P}_{p+\varepsilon}(\|K_1\| \geq c\varepsilon) \geq 1 - \frac{1}{|V|^{c\varepsilon}}.$$

Proof via coupled explorations

At the heart of Vanneuville's new proof of the sharpness of the phase transition for infinite transitive graphs is a stochastic comparison lemma. This says that starting with percolation of some parameter p , decreasing from p to $p - \varepsilon$ for a certain $\varepsilon > 0$ has more of an effect than conditioning on a certain disconnection event A , roughly in the sense that

$$\mathbb{P}_{p-\varepsilon}(\omega = \cdot) \leq_{\text{st}} \mathbb{P}_p(\omega = \cdot \mid A),$$

where \leq_{st} denotes stochastic dominance with respect to the usual partial ordering $\{0, 1\}^E$. This is proved by coupling two explorations of the cluster at the origin, sampled according to each of the two laws. In [Eas23] we modified Vanneuville's argument to prove the following lemma ([Eas23, Lemma 8]). Note that here the stochastic dominance only holds approximately, i.e. only on the complement of an event with small probability.

Lemma 8.5.2. *Let G be a finite transitive graph. Let $p, \alpha \in (0, 1)$. Define*

$$\theta := \mathbb{E}_p \|K_1\|, \quad h := \mathbb{P}_p\left(\|K_1\| < \alpha \text{ or } \|K_2\| \geq \frac{\alpha}{2}\right), \quad \delta := \frac{2h^{1/2}}{1 - \theta - h},$$

and assume that $\theta + h < 1$ (so that δ is well-defined and positive). Then there is an event A with $\mathbb{P}_p(A \mid \|K_o\| < \alpha) \leq h^{1/2}$ such that

$$\mathbb{P}_{(1-\theta-\delta)p}(\omega = \cdot) \leq_{\text{st}} \mathbb{P}_p(\omega \cup \mathbf{1}_A = \cdot \mid \|K_o\| < \alpha),$$

where $\mathbf{1}_A$ denotes the random configuration with every edge open on A and every edge closed on A^c .

To prove Proposition 8.5.1, we will simply combine this lemma together with Lemma 8.4.5, which was a standard application of Russo's formula and Talagrand's inequality.

Proof of Proposition 8.5.1. Define θ and h as in Lemma 8.5.2. Suppose that $\mathbb{P}_{p-\varepsilon}(\|K_o\| \geq \alpha) \geq \alpha$ and $\mathbb{P}_p(\|K_1\| \geq \alpha \text{ and } \|K_2\| < \frac{\alpha}{2}) > 1 - \alpha^2$, i.e. $h < \alpha^2$. First consider the case that $\theta + h \geq \frac{\varepsilon}{2}$. Then by hypothesis and the fact that $\alpha \leq \frac{\varepsilon}{10}$,

$$\theta \geq \frac{\varepsilon}{2} - h \geq \frac{\varepsilon}{2} - \alpha^2 \geq \frac{\varepsilon}{4}. \quad (8.5.1)$$

Now consider the case that $\theta + h < \frac{\varepsilon}{2}$. Define δ as in Lemma 8.5.2. By Lemma 8.5.2, there is an event A such that

$$\mathbb{P}_{(1-\theta-\delta)p}(\|K_o\| \geq \alpha) \leq \mathbb{P}_p(A \mid \|K_o\| < \alpha) \leq h^{1/2}.$$

On the other hand, by our hypotheses,

$$h^{1/2} < \alpha \leq \mathbb{P}_{p-\varepsilon}(\|K_o\| \geq \alpha).$$

So by monotonicity, we must have $(1 - \theta - \delta)p \leq p - \varepsilon$. In particular, $\theta + \delta \geq \varepsilon$. We can upper bound δ by

$$\delta = \frac{2h^{1/2}}{1 - (\theta + h)} \leq \frac{2\alpha}{1 - \frac{\varepsilon}{2}} \leq \frac{2 \cdot \frac{\varepsilon}{10}}{1 - \frac{\varepsilon}{2}} \leq \frac{2\varepsilon}{5},$$

where the last inequality used the fact that $\varepsilon \in (0, 1)$. Therefore, again, $\theta \geq \varepsilon - \delta \geq \frac{\varepsilon}{4}$, as in eq. (8.5.1).

Let $c > 0$ be the constant from Lemma 8.4.5. Without loss of generality, assume that $c < \frac{1}{8}$. Suppose that $|V|^{c\varepsilon} \geq \frac{1}{c\varepsilon}$. By Markov's inequality, $\mathbb{P}_p(\|K_1\| \geq \frac{\varepsilon}{8}) \geq \frac{\varepsilon}{8}$ because $\theta \geq \frac{\varepsilon}{4}$. Therefore,

$$\mathbb{P}_p(\|K_1\| \geq c\varepsilon) \geq \mathbb{P}_p\left(\|K_1\| \geq \frac{\varepsilon}{8}\right) \geq \frac{\varepsilon}{8} > c\varepsilon \geq \frac{1}{|V|^{c\varepsilon}}.$$

So by applying Lemma 8.4.5, $\mathbb{P}_{p+\varepsilon}(\|K_1\| \geq c\varepsilon) \geq 1 - |V|^{-c\varepsilon}$, as required. \square

8.6 Proof of Theorem 8.1.1

Let \mathcal{G} be an infinite set of finite transitive graphs with bounded degrees. Suppose that for all but at most finitely many $G \in \mathcal{G}$,

$$\text{dist}_{\text{GH}}\left(\frac{\pi}{\text{diam } G} G, S^1\right) > \frac{e^{(\log \text{diam } G)^{1/9}}}{\text{diam } G}. \quad (8.6.1)$$

Our goal is to prove that both statements (1) and (2) are true. By Proposition 8.2.9, statement (1) implies statement (2). So it suffices to prove statement (1), i.e. percolation on \mathcal{G} has a sharp phase transition. We will assume without loss of generality that there exists $d \in \mathbb{N}$ such that every $G \in \mathcal{G}$ has degree exactly d . We will again adopt the notation $\|K_1\|, \|K_2\|$ from Section 8.4.

Claim 8.6.1. *For every constant $\varepsilon > 0$, there exist constants $c(\varepsilon) > 0$ and $\mu(d, \varepsilon) < \infty$ such that for every infinite subset $\mathcal{H} \subseteq \mathcal{G}$ and every sequence $p : \mathcal{H} \rightarrow (0, 1)$,*

$$\liminf_{G \in \mathcal{H}} \mathbb{P}_p(|K_1| \geq \mu \log |V|) > 0 \implies \lim_{G \in \mathcal{H}} \mathbb{P}_{p+4\varepsilon}(\|K_1\| \geq c) = 1.$$

Before proving this claim, let us explain how to conclude from it. For each $G \in \mathcal{G}$, pick a parameter $q(G) \in (0, 1)$ satisfying $\mathbb{P}_{q(G)}^G(|K_1| \geq |V|^{2/3}) = \frac{1}{2}$. We will prove that percolation on \mathcal{G} has a sharp phase transition with percolation threshold given by $q : \mathcal{G} \rightarrow (0, 1)$. First notice that $\liminf q \geq \frac{1}{2d} > 0$. Indeed, this follows from the proof of [EH21a, Lemma 2.8], but let us explain the elementary argument here for completeness. For every $G \in \mathcal{G}$ and $n \geq 1$, there are at most d^n self-avoiding paths starting from o . So by a union bound, every $G \in \mathcal{G}$ satisfies

$$\mathbb{E}_{\frac{1}{2d}} |K_o| \leq \sum_{n=0}^{\infty} \frac{d^n}{(2d)^n} = 2.$$

On the other hand, by transitivity, every $G \in \mathcal{G}$ satisfies

$$\mathbb{E}_{\frac{1}{2d}} |K_o| \geq |V|^{2/3} \mathbb{P}_{\frac{1}{2d}} \left(|K_o| \geq |V|^{2/3} \right) \geq |V|^{1/3} \mathbb{P}_{\frac{1}{2d}} \left(|K_1| \geq |V|^{2/3} \right).$$

Therefore for all but finitely many $G \in \mathcal{G}$,

$$\mathbb{P}_{\frac{1}{2d}} \left(|K_1| \geq |V|^{2/3} \right) \leq 2 |V|^{-1/3} < \frac{1}{2},$$

and hence by monotonicity, $q(G) \geq \frac{1}{2d}$.

Now fix a constant $\varepsilon > 0$. Since $\liminf q \geq \frac{1}{2d} > 0$, there exists a constant $\delta(\varepsilon, d) > 0$ such that $(1 - \varepsilon)q \leq q - \delta$ and $q + \delta \leq (1 + \varepsilon)q$ for all but finitely many $G \in \mathcal{G}$. Let $c(\delta/4) > 0$ and $\mu(d, \delta/4) < \infty$ be the constants provided by the claim. For all but finitely many $G \in \mathcal{G}$, we have $\mu \log |V| < |V|^{2/3}$. So by applying the claim with “ \mathcal{H} ” being the whole of \mathcal{G} and “ p ” being q ,

$$\lim_{G \in \mathcal{G}} \mathbb{P}_{q+\delta}(\|K_1\| \geq c) = 1.$$

On the other hand, for all but finitely many $G \in \mathcal{G}$, we have $c |V| > |V|^{2/3}$. So by applying the claim (contrapositively) with “ p ” being $q - \delta$, for every infinite subset $\mathcal{H} \subseteq \mathcal{G}$,

$$\liminf_{G \in \mathcal{H}} \mathbb{P}_{q-\delta}(|K_1| \geq \mu \log |V|) = 0.$$

Equivalently, for every infinite subset $\mathcal{H} \subseteq \mathcal{G}$ there exists a further infinite subset $\mathcal{H}' \subseteq \mathcal{H}$ such that $\lim_{G \in \mathcal{H}'} \mathbb{P}_{q-\delta}(|K_1| \geq \mu \log |V|) = 0$. Therefore,

$$\lim_{G \in \mathcal{G}} \mathbb{P}_{q-\delta}(|K_1| \geq \mu \log |V|) = 0.$$

Since $\varepsilon > 0$ was arbitrary, this establishes that percolation on \mathcal{G} has a sharp phase transition. All that remains is to verify the claim.

Proof of claim. Fix $\varepsilon > 0$. Let $c_1(d, \varepsilon) > 0$ be the constant from Proposition 8.2.1. Let $\lambda(d, \varepsilon, \frac{\varepsilon^2}{20}) < \infty$ be the constant from Proposition 8.3.1 (with “ η ” set to $\varepsilon^2/20$). Let $c_2 > 0$ be the universal constant from Proposition 8.5.1. We will prove that the claim holds with $\mu := \exp\left(\frac{\lambda}{c_1} + \frac{1}{c_1^2}\right)$ and $c := c_2\varepsilon$. Let $\mathcal{H} \subseteq \mathcal{G}$ be an infinite subset, and let $p : \mathcal{H} \rightarrow (0, 1)$ be a sequence satisfying

$$\eta := \liminf_{G \in \mathcal{H}} \mathbb{P}_p(|K_1| \geq \mu \log |V|) > 0.$$

We say that a statement A holds for *almost every* G to mean that the set $\{G \in \mathcal{H} : A \text{ does not hold for } G\}$ is finite. For almost every G ,

$$\mathbb{P}_p(|K_1| \geq \mu \log |V|) \geq \frac{\eta}{2} \geq \frac{1}{c_1 |V|^{c_1}}.$$

So by Proposition 8.2.1, noting that $c_1 \log \mu - \frac{1}{c_1} = \lambda$,

$$\min_{u \in B_\lambda} \mathbb{P}_{p+\varepsilon}(o \leftrightarrow u) \geq \frac{\varepsilon^2}{20}.$$

For each $G \in \mathcal{H}$, define $\gamma(G)$ and $\gamma^+(G)$ as in Proposition 8.3.1. Then by Proposition 8.3.1, thanks to our choice of λ , for almost every G ,

$$\min_{u \in B_{\gamma^+}} \mathbb{P}_{p+2\varepsilon}(o \leftrightarrow u) \geq e^{-(\log \log \gamma^+)^{1/2}}. \quad (8.6.2)$$

Consider a particular $G \in \mathcal{H}$ satisfying eq. (8.6.1). Then $\gamma(G) > e^{(\log \text{diam } G)^{1/9}}$, and by applying the monotone function $x \mapsto e^{(\log x)^9}$ to both sides, $\gamma^+(G) > \text{diam } G$. In particular, $B_{\gamma^+(G)}^G$ is the whole vertex set $V(G)$. We trivially have $\text{dist}_{\text{GH}}\left(\frac{1}{\text{diam } G} G, \frac{1}{\pi} S^1\right) \leq 1$, because both metric spaces involved have diameter ≤ 1 . So conversely, $\gamma(G) \leq \pi \text{diam } G$, and hence $\gamma^+(G) \leq e^{(\log(\pi \text{diam } G))^9}$. By applying these upper and lower bounds on $\gamma^+(G)$ to eq. (8.6.2), we deduce that for almost every G ,

$$\min_{u, v \in V} \mathbb{P}_{p+2\varepsilon}(u \leftrightarrow v) \geq e^{-3(\log \log(\pi \text{diam } G))^{1/2}} \geq e^{-3(\log \log(\pi |V|))^{1/2}},$$

where the second inequality follows from the trivial bound $|V| \geq \text{diam } G$.

For each $G \in \mathcal{H}$, define $\delta(G) := (\log |V|)^{-1/20}$. For every sufficiently large positive real x ,

$$2(\log x)^{-1/20} = 2e^{-\frac{1}{20} \log \log x} \leq e^{-3(\log \log(\pi x))^{1/2}}.$$

Therefore for almost every G ,

$$\min_{u,v \in V} \mathbb{P}_{p+2\varepsilon}(u \leftrightarrow v) \geq 2\delta. \quad (8.6.3)$$

By applying Proposition 8.4.1, it follows that for almost every G , (since $\delta \leq \varepsilon$)

$$\mathbb{P}_{p+3\varepsilon} \left(\|K_1\| \geq \delta \text{ and } \|K_2\| \leq \delta^2 \right) \geq 1 - \delta^4.$$

For almost every G , we have $\delta^2 < \frac{\delta}{2}$, $\delta^4 < \delta^2$, $\delta \in (0, \frac{\varepsilon}{10})$, $|V|^{c_2\varepsilon} \geq \frac{1}{c_2\varepsilon}$, and by applying Markov's inequality to eq. (8.6.3), $\mathbb{P}_{p+2\varepsilon}(\|K_o\| \geq \delta) \geq \delta$. So by Proposition 8.5.1, for almost every G ,

$$\mathbb{P}_{p+4\varepsilon} (\|K_1\| \geq c_2\varepsilon) \geq 1 - \frac{1}{|V|^{c_2\varepsilon}}.$$

In particular, $\lim_{G \in \mathcal{H}} \mathbb{P}_{p+4\varepsilon} (\|K_1\| \geq c_2\varepsilon) = 1$, as claimed. \square

Appendix: Details for some claims in Section 8.3

In this appendix, we will explain how some of the lemmas in Section 8.3 can be established by minor modifications of existing arguments.

Lemma (Lemma 8.3.3). *There exists $n_0(d) < \infty$ such that for all $n \geq n_0$,*

$$s([R(n), R^2(n)] \text{ is orange}) \leq s(n \text{ is green}) + \delta(n).$$

The following argument is essentially contained in the proof of [EH23b, Proposition 4.5].

Proof. Let $n \geq 3$ be some scale. Throughout this proof, we will assume that n is large with respect to d . Note that the result is trivial if $n > \text{diam } G$, because in that case $B_m = B_n$ for all $m \in [R(n), R^2(n)]$. So let us assume to the contrary that $S_n \neq \emptyset$. First consider the case that $n \in \mathbb{L}$. Then at any time t when n is green, we know that $\kappa_{\phi(t)}(R^2(n), n) \geq \delta(R(n))$, which implies that $[R(n), R^2(n)]$ is already orange at time t . So let us assume to the contrary that $n \notin \mathbb{L}$. Define $h := e^{-(\log n)^{100}}$, which therefore satisfies $h \geq \text{Gr}(n)^{-1}$. Pick $p_1 \in (0, 1)$ such that n is green at time $\phi^{-1}(p_1)$. Note that $p_1 \geq 1/d$ because by a union bound, using that $S_n \neq \emptyset$ and that n is large with respect to d ,

$$\min_{u \in B_n} \mathbb{P}_{1/d}(o \leftrightarrow u) \leq \mathbb{P}_{1/d}(o \leftrightarrow S_n) \leq d(d-1)^{n-1} \cdot \left(\frac{1}{d}\right)^n < \delta(n).$$

Define $p_2 := \phi(\phi^{-1}(p_1) + \delta(n))$. In the language of [EH23b, Section 3], the quantity “ $\delta(p_1, p_2)$ ” is equal to $\delta(n)$ by construction. Let $u \in B_{R^2(n)}$ be arbitrary, and let $o = u_0, u_1, \dots, u_k = u$ be a path with $k \leq R^2(n)$. Let $c_1(1), h_0(d, 1), c_2, c_3 > 0$ be the constants from [EH23b, Proposition 4.1] with $D := 1$. Since n is large with respect to d , we have $h \leq h_0, \delta(n) \leq 1$,

$$h^{c_1\delta(n)^3} = e^{-c_1(\log n)^{100}e^{-3(\log \log n)^{1/2}}} \leq \frac{c_3}{e^{(\log n)^{81}} + 1} = \frac{c_3}{R^2(n) + 1} \leq \frac{c_3}{k + 1},$$

and for all $i \in \{0, \dots, k-1\}$, by Harris’ inequality,

$$\begin{aligned} \min \{\mathbb{P}_{p_1}(x \leftrightarrow y) : x, y \in B_n(u_i) \cup B_n(u_{i+1})\} &\geq \delta(n) \cdot p_1 \cdot \delta(n) \\ &\geq \frac{1}{d} e^{-2(\log \log n)^{1/2}} \\ &\geq 4e^{-c_1(\log n)^{100}e^{-4(\log \log n)^{1/2}}} = 4h^{c_1\delta(n)^4}. \end{aligned}$$

So by [EH23b, Proposition 4.1], where the sets “ A_1, \dots, A_n ” are the balls $B_n(u_0), \dots, B_n(u_k)$,

$$\mathbb{P}_{p_2}(o \leftrightarrow u) \geq c_2\delta(n)^2 \geq \delta(R(n)).$$

Since $u \in B_{R^2(n)}$ was arbitrary, it follows that $[R(n), R^2(n)]$ is orange at time $\phi^{-1}(p_2)$, as required. \square

Lemma (Lemma 8.3.4). *For all $c > 0$ there exist $\lambda(d, c), n_0(d, c), K(d, c) < \infty$ such that the following holds for all $n \geq n_0$ with $n \in \mathbb{T}(c, \lambda)$. For all $t \in \mathbb{R}$, if n is orange at time t and $K\Delta_t(n) \leq 1$ then*

$$s(n \text{ is green}) \leq t + K\Delta_t(n).$$

This is implicit in the proof of [EH23b, Proposition 6.1]

Proof. In [EH23b], we made the following definitions: given $d \geq 1$, we wrote \mathcal{U}_d^* for the set of all infinite non-one-dimensional unimodular transitive graphs with degree d , and given $D \geq 1$ and a transitive graph G , we wrote $\mathcal{L}(G, D)$ for the set of all scales $n \geq 1$ such that $\text{Gr}(m) \leq e^{(\log m)^D}$ for all $m \in [n^{1/3}, n]$. Let us now introduce the following variants of these definitions: given $d \geq 1$, write \mathcal{W}_d for the set of all (possibly finite) unimodular transitive graphs with degree d , and given $D, \lambda \geq 1, c > 0$, and a transitive graph G , write $\mathcal{T}(G, D, \lambda, c)$ for the set all of scales $n \in \mathcal{L}(G, D)$ with $n \leq \text{diam } G$ such that G has (c, λ) -polylog plentiful tubes throughout an interval of the form $[m_1, m_2]$ with $m_2 \geq m_1^{1+c}$ satisfying $[m_1, m_2] \subseteq [n^{1/3}, n^{1/(1+c)}]$. Let [EH23b, Proposition* 6.1] be the result of modifying the statement of [EH23b, Proposition 6.1] as follows:

1. Weaken the hypothesis that $G \in \mathcal{U}_d^*$ to the hypothesis that $G \in \mathcal{W}_d$.
2. Strengthen the hypothesis that $n \in \mathcal{L}(G, D)$ to the hypothesis that $n \in \mathcal{T}(G, D, \lambda, 1/D)$.

Note that $p_c(G)$ in this statement refers to the usual percolation threshold for an infinite cluster, so in particular, $p_c(G) := 1$ if G is finite. The same proof works because the hypothesis that G was infinite and non-one-dimensional was only used to invoke [EH23b, Proposition 5.2] to establish that there is a constant $c_1(d, D) > 0$ such that for all λ , whenever n is large with respect to d, D, λ , if $n \in \mathcal{L}(G, D)$ then automatically $n \in \mathcal{T}(G, D, \lambda, c_1)$. We are just circumventing this application of [EH23b, Proposition 5.2]. Specifically, we can prove [EH23b, Proposition* 6.1] by modifying the proof of [EH23b, Proposition 6.1] as follows:

1. Strengthen the condition $n \in \mathcal{L}(G, D)$ to $n \in \mathcal{T}(G, D, \lambda, 1/D)$ in the definition of \mathcal{A} .
2. Rather than define c_1 and N to be the constants guaranteed to exist by [EH23b, Proposition 5.2], set $c_1 := 1/D$ and $N := 3$.
3. Restrict the domain of the definition of $\mathcal{P}(n)$ from all $n \in \mathcal{L}(G, D)$ to all $n \in \mathcal{T}(G, D, \lambda, 1/D)$.

4. Include the hypothesis $n \in \mathcal{T}(G, D, \lambda, 1/D)$ in the statement of [EH23b, Lemma 6.8].³³

Taking [EH23b, Proposition* 6.1] for granted, let us now explain how to prove Lemma 8.3.4. Recall that G is a finite transitive graph with degree d . Let $c > 0$ be given, and define $D := 101 \vee (1/c)$. Let $\lambda_0(d, D)$ and $c_1(d, D)$ (called “ $c(d, D)$ ”) be the constants provided by [EH23b, Proposition* 6.1]. Define $\lambda := \lambda_0 \vee (100/c_1)$. Now let $K_1(d, D, \lambda)$ and $n_0(d, D, \lambda)$ be the corresponding constants provided by [EH23b, Proposition* 6.1]. Define $K := K_1^{1/4}$. By the same argument as in our proof of Lemma 8.3.3 above, there exists $n_1(d) < \infty$ such that for all $n_1 \leq n \leq \text{diam } G$ and $t \in \mathbb{R}$, if n is orange at time t then $\phi(t) \geq 1/d$. Set $n_2 := n_0 \vee n_1 \vee e^{3^{101}}$. We claim that λ, n_2, K have the properties required of the constants called “ λ, n_0, K ” in the statement of Lemma 8.3.4.

Indeed, suppose that $t \in \mathbb{R}$ and $n \geq n_2$ with $n \in \mathbb{T}(c, \lambda)$ are such that n is orange at time t and $K\Delta_t(n) \leq 1$. Now apply [EH23b, Proposition* 6.1] with the variables called “ K, n, b, p_1, p_2 ” in that statement set to our variables $K_1, n, U_t(n), \phi(t), \phi(t + K\Delta_t(n))$. The only hypothesis that is not immediately obvious is that $n \in \mathcal{T}(G, D, \lambda, 1/D)$. To see this, first note that since $n \in \mathbb{L}$ and $n \geq e^{3^{101}}$, every $m \in [n^{1/3}, n]$ satisfies

$$\text{Gr}(m) \leq \text{Gr}(n) \leq e^{(\log n)^{100}} \leq e^{(\log(n^{1/3}))^{101}} \leq e^{(\log m)^{101}} \leq e^{(\log m)^D}.$$

So $n \in \mathcal{L}(G, D)$. Second, we may assume that $n \leq \text{diam } G$, otherwise the conclusion of Lemma 8.3.4 is trivial. Finally, since $n \in \mathbb{T}(c, \lambda)$ and $1/D < c$, and the property of having “ (x, λ) -polylog plentiful tubes” at a given scale gets weaker as we decrease x , it follows that $n \in \mathcal{T}(G, D, \lambda, 1/D)$. Therefore, by applying [EH23b, Proposition* 6.1], we deduce that

$$\kappa_{\phi(t+K\Delta_t(n))} \left(e^{(\log n)^{c_1 \lambda}}, n \right) \geq e^{-3(\log \log n)^{1/2}}.$$

In particular, since $c_1 \lambda \geq 100 \geq 81$,

$$\kappa_{\phi(t+K\Delta_t(n))} \left(R^2(n), n \right) \geq \delta(R(n)).$$

So $s(n \text{ is green}) \leq t + K\Delta_t(n)$, as required. \square

Lemma (Lemma 8.3.5). *There exists $n_0(d) < \infty$ such that the following holds for all $n \in \mathbb{L}$ with $n \geq n_0$. For all $t \in \mathbb{R}$, if $L(n)$ is green at time t then*

$$\Delta_t(n) \leq \frac{1}{\log \log n}.$$

³³While writing this paper, we noticed the following typo: [EH23b, Lemma 6.8] is missing the hypothesis that $n \in \mathcal{L}(G, D)$.

This proof is implicit in [EH23b, Section 6.3].

Proof. Suppose that $n \in \mathbb{L}$. Throughout this proof we will implicitly assume that n is large with respect to d . Let $t \in \mathbb{R}$ and assume that $L(n)$ is green at time t . We may assume that $\lfloor n^{1/3} \rfloor \leq \text{diam } G$, otherwise we trivially have $U_t(n) = \lfloor \frac{1}{8}n^{1/3} \rfloor$ and hence (since n is assumed large) $\Delta_t(n) \leq (\log n)^{-1/5}$. We split the proof into two cases according to whether $L(n) \in \mathbb{L}$.

First suppose that $L(n) \notin \mathbb{L}$. By the same argument as in our proof of Lemma 8.3.3 above, since $L(n) \leq \text{diam } G$ and $L(n)$ is green at time t (and since n is assumed large), $\phi(t) \geq 1/d$. So by [EH23b, Corollary 2.4], there exist constants $c(d) > 0$ and $C(d), n_0(d) < \infty$ such that for all $m \geq n_0(d)$,

$$\mathbb{P}_{\phi(t)}(\text{Piv}[c \log m, m]) \leq C \left(\frac{\log \text{Gr}(m)}{m} \right)^{1/3}.$$

In particular, since $4L(n) \leq c \log(n^{1/3})$ and $\text{Gr}(n^{1/3}) \leq \text{Gr}(n) \leq e^{(\log n)^{100}}$,

$$\mathbb{P}_{\phi(t)}(\text{Piv}[4L(n), n^{1/3}]) \leq C \left(\frac{(\log n)^{100}}{n^{1/3}} \right)^{1/3} \leq \frac{1}{\log n}.$$

Since we also clearly have $L(n) \leq \frac{1}{8}n^{1/3}$, it follows that $U_t(n) \geq \lfloor L(n) \rfloor$. Since $L(n) \notin \mathbb{L}$, this implies that $\text{Gr}(U_t(n)) \geq e^{(\log L(n))^{100}}$. So

$$\Delta_t(n) \leq \left(\frac{\log \log n}{(\log n) \wedge (\log L(n))^{100}} \right)^{1/4} \leq \frac{1}{\log \log n}.$$

Next suppose that $L(n) \in \mathbb{L}$. Define $b := \frac{1}{5} \left(R \circ L(n) \wedge \text{Gr}^{-1}(R^{-1}(n)) \right)$. By [EH23b, Lemma 2.3] (i.e. [CMT22, Lemma 6.2]), using the fact that $5b \leq \frac{1}{2}n^{1/3}$,

$$\mathbb{P}_{\phi(t)}(\text{Piv}[4b, n^{1/3}]) \leq \mathbb{P}_{\phi(t)}\left(\text{Piv}\left[1, \frac{1}{2}n^{1/3}\right]\right) \cdot \frac{|S_{4b}|^2 \text{Gr}(5b)}{\min_{x,y \in S_{4b}} \mathbb{P}_{\phi(t)}(x \xrightarrow{B_{5b}} y)}.$$

By [EH23b, Lem 2.1] (i.e. essentially [CMT22, Proposition 4.1]), there is a constant $C(d) < \infty$ such that

$$\mathbb{P}_{\phi(t)}\left(\text{Piv}\left[1, \frac{1}{2}n^{1/3}\right]\right) \leq C \left(\frac{\log \text{Gr}\left(\frac{1}{2}n^{1/3}\right)}{\frac{1}{2}n^{1/3}} \right)^{1/3}.$$

By hypothesis, $n \in \mathbb{L}$. So we can upper bound $\log \text{Gr}\left(\frac{1}{2}n^{1/3}\right) \leq \log \text{Gr}(n) \leq (\log n)^{100}$. Since $L(n)$ is green at time t but $L(n) \in \mathbb{L}$, then $\kappa_{\phi(t)}(R^2 \circ L(n), L(n)) \geq \delta(R \circ L(n))$. Note that $8b \leq R^2 \circ L(n)$ and (using that $L(n) \in \mathbb{L}$), $L(n) \leq b$. Therefore, $\min_{x,y \in S_{4b}} \mathbb{P}_{\phi(t)}(x \xrightarrow{B_{5b}} y) \geq \delta(R \circ L(n))$, since we

can connect any $x, y \in S_{4b}$ by a path contained in B_{4b} of length at most $8b$, and the b -thickened tube around this path is entirely contained in B_{5b} . Finally, we can upper bound $|S_{4b}| \leq \text{Gr}(5b) \leq R^{-1}(n)$ by definition of b . Therefore,

$$\mathbb{P}_{\phi(t)} \left(\text{Piv} [4b, n^{1/3}] \right) \leq C \left(\frac{(\log n)^{100}}{\frac{1}{2} n^{1/3}} \right)^{1/3} \frac{(R^{-1}(n))^3}{\delta(R \circ L(n))} \leq \frac{1}{\log n}.$$

Notice that by our choice of b , we have $b \leq \frac{1}{8} n^{1/3}$ and

$$\text{Gr}(b) \geq \text{Gr} \left(\frac{1}{5} R \circ L(n) \right) \wedge \text{Gr} \left(\frac{1}{5} \text{Gr}^{-1} (R^{-1}(n)) \right) \geq \left(\frac{1}{5} R \circ L(n) \right) \wedge (R^{-1}(n))^{1/5} = \frac{1}{5} R \circ L(n).$$

So

$$\Delta_t(n) \leq \left(\frac{\log \log n}{(\log n) \wedge \left(\log \left[\frac{1}{5} R \circ L(n) \right] \right)} \right)^{1/4} \leq \frac{1}{\log \log n}.$$

□

Lemma (Lemma 8.3.11). *Let G be a unimodular transitive graph of degree d . Suppose that*

$$\text{Gr}(m) \leq e^{(\log m)^D} \quad \text{and} \quad \text{Gr}(3m) \geq 3^5 \text{Gr}(m)$$

for every $m \in [n^{1-\varepsilon}, n^{1+\varepsilon}]$, where $\varepsilon, D, n > 0$. Then there is a constant $c(d, D, \varepsilon) > 0$ with the following property. For every $\lambda \geq 1$, there exists $n_0(d, D, \varepsilon, \lambda) < \infty$ such that if $n \geq n_0$ then G has (c, λ) -polylog plentiful tubes at scale n .

[EH23b, Lemma 5.4] is the same statement but with the additional hypothesis that G is infinite. We claim that this additional hypothesis is unnecessary.

Proof. [EH23b, Lemma 5.4] is the ultimate conclusion of [EH23b, Section 5.2]. The first result in [EH23b, Section 5.2] that requires G to be infinite is [EH23b, Lemma 5.16]. By inspecting the proof of [EH23b, Lemma 5.16], we see that this hypothesis is only used in order to apply the elementary bound $\text{Gr}(3mn) \geq n \text{Gr}(m)$ for all $m, n \geq 1$. In fact, in the language of that proof, since we may assume that the constant $c > 0$ satisfies $c \leq 1/10$, say, then the proof only invokes this elementary bound for m, n satisfying $3mn \leq \frac{1}{10} t^{1/2}$. Now this holds whenever $\text{diam } G \geq \frac{1}{10} t^{1/2}$. So [EH23b, Lemma 5.16] holds with the hypothesis “ G is infinite” replaced by the weaker hypothesis “ $\text{diam } G \geq \frac{1}{10} t^{1/2}$ ”. When [EH23b, Lemma 5.16] is applied to establish [EH23b, Lemma 5.17], the hypothesis “ $\text{diam } G \geq \frac{1}{10} t^{1/2}$ ” is already implied by the other hypothesis of [EH23b, Lemma

5.17] that $\text{Gr}(3m) \geq 3^\kappa \text{Gr}(m)$ for all $n \leq m \leq \frac{1}{2}t^{1/2}$ (and the fact that conclusion of [EH23b, Lemma 5.17] is trivial if there is no integer in $[n, \frac{1}{2}t^{1/2}]$). So in the statement of [EH23b, Lemma 5.17], we can simply drop the hypothesis that G is infinite.

We can also drop the hypothesis that G is infinite in [EH23b, Lemmas 5.18 and 5.20] because [EH23b, Lemma 5.18] is deduced from [EH23b, Lemma 5.17], and [EH23b, Lemma 5.20] is deduced from [EH23b, Lemma 5.18]. [EH23b, Lemma 5.19] already does not require G to be infinite. The ultimate proof of [EH23b, Lemma 5.4] only required G to be infinite in order to invoke [EH23b, Lemma 5.20] and (in the radial case) to know that $S_n \neq \emptyset$. The hypothesis that $S_n \neq \emptyset$ is anyway implied by the fact that $\text{Gr}(3m) \geq \text{Gr}(m)$ for some $m \in [n, n^{1+\varepsilon}]$, and as we explained, we can drop the hypothesis that G is infinite in [EH23b, Lemma 5.20]. Therefore we can drop the hypothesis that G is infinite in the statement of [EH23b, Lemma 5.4] too. \square

The next claim we will justify is that Lemma 8.3.16 implies Lemma 8.3.13. Here are the statements of these results.

Lemma (Lemma 8.3.16). *Let $r, n \geq 1$. Let G be a finite transitive graph such that S_n^∞ is not r -connected. Let H be a (finite or infinite) transitive graph with $\delta(H) \leq r$ that does not have infinitely many ends. If $B_{50n}^H \cong B_{50n}^G$, then*

$$\text{dist}_{\text{GH}} \left(\frac{\pi}{\text{diam } G} G, S^1 \right) \leq \frac{200n}{\text{diam } G}.$$

Lemma (Lemma 8.3.13). *Let G be a finite transitive graph of degree d . Suppose that $\text{Gr}(3n) \leq 3^\kappa \text{Gr}(n)$, where $n, \kappa > 0$. There exists $C(d, \kappa) < \infty$ such that the following holds if $n \geq C$:*

There is a set $A \subseteq [1, \infty)$ with $|A| \leq C$ such that for every $k \geq 1$ and every $m \in [Ckn, \infty) \setminus \bigcup_{a \in A} [a, 2ka]$, if G does not have $(C^{-1}k, C^{-1}k^{-1}m, Ck^C m)$ -plentiful tubes at scale m , then

$$\text{dist}_{\text{GH}} \left(\frac{\pi}{\text{diam } G} G, S^1 \right) \leq \frac{Cm}{\text{diam } G}.$$

The proof that Lemma 8.3.16 implies Lemma 8.3.13 is essentially the same as the proof of [EH23b, Proposition 5.3] (i.e. Lemma 8.3.12), except that G is now assumed to be a finite transitive graph rather than a non-one-dimensional infinite transitive graph. For this reason, the following proof is terse. The argument relies on the structure theory of transitive graphs of polynomial growth. See the proof of [EH23b, Proposition 5.3] for more details and [EH23b, Section 5.1] for more background.

Proof of Lemma 8.3.13 given Lemma 8.3.16. Fix $\kappa > 0$. Suppose that $\text{Gr}(3n) \leq 3^\kappa \text{Gr}(n)$ for some $n \geq 1$. We will implicitly assume that n is large with respect to d and κ . Let $H \leq \text{Aut}(G)$, $S \subseteq \Gamma := \text{Aut}(G)/H$, and $C_1(K) < \infty$ be as given by [EH23b, Theorem 5.5] (which is taken from [TT21a]) with $K := 3^\kappa$. Let $G' := \text{Cay}(\Gamma, S)$. For each $k \in \mathbb{N}$, let R_k be the set of all relations in Γ having word length at most k , let $\langle\langle R_k \rangle\rangle$ be the normal subgroup of the free group on S generated by R_k , and let $G'_k := \text{Cay}(\langle\langle S \mid R_k \rangle\rangle, S)$. By items 7 and 8 of [EH23b, Theorem 5.5],

$$\frac{\text{Gr}'(3n)}{\text{Gr}'(n)} \leq C_1^2(3 + C_1)^{C_1}.$$

In particular, by [EH23b, Theorem 5.5] again (and using that n is large), every transitive graph whose $3n$ -ball is isomorphic to the $3n$ -ball in G' is necessarily finite or infinite with polynomial growth. In particular, such graphs have at most finitely many ends. Now by [EH23d, Theorem 1.1], there exists $C_2(\kappa, d) < \infty$ such that

$$|\{i \in \mathbb{N} : i \geq \log_2 n \text{ and } \langle\langle R_{2^{i+1}} \rangle\rangle \neq \langle\langle R_{2^i} \rangle\rangle\}| \leq C_2.$$

Let $A := \{2^i : i \in \mathbb{N} \text{ and } i \geq \log_2 n \text{ and } \langle\langle R_{2^{i+10}} \rangle\rangle \neq \langle\langle R_{2^i} \rangle\rangle\}$, and note that $|A| \leq 10C_2$. Let $k \geq 1$ and $m \in [2kn, \infty) \setminus \bigcup_{a \in A} [a, 2ka]$ be arbitrary. By construction of A (and [EH23b, Lemma 5.6]), the balls of radius (say) $50n$ in $G'_{\frac{m}{k}}$ and G' are isomorphic. Note that $\delta(G'_{\frac{m}{k}}) \leq \frac{m}{k}$, and since the $3n$ -ball in $G'_{\frac{m}{k}}$ is isomorphic to the $3n$ -ball in G' , the graph $G'_{\frac{m}{k}}$ has at most finitely many ends. Consider an arbitrary pair $m_1, m_2 \in \mathbb{N}$ satisfying $\frac{m}{k} \leq m_1 \leq m_2 \leq 3m$. By Lemma 8.3.16 applied with the pair “ (G, H) ” equal to $(G', G'_{\frac{m}{k}})$, either (1) the exposed sphere $S_{m_2}^\infty(G')$ is $\lceil \frac{m}{k} \rceil$ -connected, or (2)

$$\text{dist}_{\text{GH}}\left(\frac{\pi}{\text{diam } G'} G', S^1\right) \leq \frac{200m_2}{\text{diam } G'}. \quad (8.6.4)$$

In case (1), we deduce by the proof of [CMT22, Lemma 2.7] (which was behind [EH23b, Lemma 5.8]) that for all $u, v \in S_{m_2}^\infty(G')$ there exists a path from u to v in G' that is contained in $\bigcup_{x \in S_{m_2}^\infty(G')} B_{2m_1}(x)$ and has length at most $3m_1 \text{Gr}(3m_2)/\text{Gr}(m_1)$. Now consider case (2). The existence of a $(1, C_1 n)$ -quasi-isometry from G to G' implies that $|\text{diam } G - \text{diam } G'| \leq 3C_1 n$ and $\text{dist}_{\text{GH}}(G, G') \leq C_1 n$. (For the latter, see exercise 5.10 (b) in [Pet23], for example.) We may assume without loss of generality that $\text{diam } G \geq 100C_1 n$, say, otherwise our claim is trivial. By combining these simple bounds with eq. (8.6.4), we deduce that $\text{dist}_{\text{GH}}(\frac{\pi}{\text{diam } G} G, S^1) \leq \frac{C_3 m}{\text{diam } G}$ for some constant $C_3(\kappa, d) < \infty$.

We now run the rest of the proof of [EH23b, Proposition 5.3], after the application of [EH23b, Lemma 5.8], as it is written. This establishes that there is a constant $C_4(\kappa, d) < \infty$ such that for all $k \geq 1$ and $m \in [C_4 kn, \infty) \setminus \bigcup_{a \in A} [a, 2ka]$, either (A) there exists $m_2 \in [\frac{10}{9}m, \frac{12}{9}m]$ such

that $S_{m_2}^\infty(G')$ is not $\lceil \frac{m}{k} \rceil$ -connected, or (B) G has $(C_4^{-1}k, C_4^{-1}k^{-1}, C_4k^{C_4}m)$ -plentiful tubes at scale m . (Technically, as written, the radial case of the proof of [EH23b, Proposition 5.3] invokes the existence of a bi-infinite geodesic in G . All that is really required is a geodesic of length $\geq 24m$. So it suffices to know that $\text{diam } G \geq 24m$, say, which we may anyway assume without loss of generality otherwise the conclusion holds trivially.) By above, if case (A) holds, then case (2) holds, and hence $\text{dist}_{\text{GH}}(\frac{\pi}{\text{diam } G}G, S^1) \leq \frac{C_3m}{\text{diam } G}$. Therefore the set of scales A is as required. \square

The next claim we will justify is that Timar's proof [Tim07] of Benjamini-Babson [BB99b] yields the following statement, which is phrased slightly differently to usual, in terms of *sets of vertices*, *vertex cutsets*, and (extrinsic) *diameter* rather than length of generating cycles.

Lemma (Lemma 8.3.17). *Let G be a graph. Let A and B be sets of vertices. Let Π be a minimal (A, B) -cutset that does not disconnect A or B . Then Π is $\delta(G)$ -connected.*

Proof. Suppose that $\Pi = \Pi_1 \sqcup \Pi_2$ is a non-trivial partition of Π . By minimality of Π , there exist paths γ_1 avoiding Π_2 and γ_2 avoiding Π_1 that both start in A and end in B . Let γ_A be a path from the startpoint of γ_1 to the startpoint of γ_2 that avoids Π , and let γ_B be a path from the endpoint of γ_1 to the endpoint of γ_2 that avoids Π . Let $\{C_i : i \in I\}$ be a set of cycles of diameter $\leq \delta(G)$ such that $\gamma_1 + \gamma_2 + \gamma_A + \gamma_B = \sum_{i \in I} C_i$. Let J be the set of all indices $i \in I$ such that C_i visits Π_1 , and define

$$\zeta := \gamma_1 + \sum_{i \in J} C_i = \gamma_2 + \gamma_A + \gamma_B + \sum_{i \in I \setminus J} C_i.$$

From either expression for ζ , we see that ζ has exactly two odd-degree vertices, one in A and the other in B . So ζ contains a path from A to B , and hence contains an edge incident to Π . From the second expression for ζ , we see that ζ does not contain an edge incident to Π_1 . So ζ must contain an edge incident to Π_2 . By construction, γ_1 avoids Π_2 . So by the first expression for ζ , there must exist a cycle C_i with $i \in J$ that visit Π_2 . Since this C_i also visits Π_1 (by definition of J) and has diameter at most $\delta(G)$, it follows that $\text{dist}(\Pi_1, \Pi_2) \leq \delta(G)$. \square

[AB07] Omer Angel and Itai Benjamini. “A phase transition for the metric distortion of percolation on the hypercube”. In: *Combinatorica* 27.6 (2007), pp. 645–658. ISSN: 0209-9683. doi: [10.1007/s00493-007-2241-5](https://doi.org/10.1007/s00493-007-2241-5). URL: <https://doi.org/10.1007/s00493-007-2241-5>.

[AB87a] Michael Aizenman and David J. Barsky. “Sharpness of the phase transition in percolation models”. In: *Comm. Math. Phys.* 108.3 (1987), pp. 489–526. ISSN: 0010-3616. URL: <http://projecteuclid.org/euclid.cmp/1104116538>.

[AB87b] Michael Aizenman and David J. Barsky. “Sharpness of the phase transition in percolation models”. In: *Comm. Math. Phys.* 108.3 (1987), pp. 489–526. ISSN: 0010-3616. URL: <http://projecteuclid.org/euclid.cmp/1104116538>.

[ABS04a] Noga Alon, Itai Benjamini, and Alan Stacey. “Percolation on finite graphs and isoperimetric inequalities”. In: *Ann. Probab.* 32.3A (2004), pp. 1727–1745. ISSN: 0091-1798. doi: [10.1214/009117904000000414](https://doi.org/10.1214/009117904000000414). URL: <https://doi.org/10.1214/009117904000000414>.

[ABS04b] Noga Alon, Itai Benjamini, and Alan Stacey. “Percolation on finite graphs and isoperimetric inequalities”. In: *Ann. Probab.* 32.3A (2004), pp. 1727–1745. ISSN: 0091-1798. doi: [10.1214/009117904000000414](https://doi.org/10.1214/009117904000000414). URL: <https://doi.org/10.1214/009117904000000414>.

[ABS04c] Noga Alon, Itai Benjamini, and Alan Stacey. “Percolation on finite graphs and isoperimetric inequalities”. In: *Ann. Probab.* 32.3A (2004), pp. 1727–1745. ISSN: 0091-1798. doi: [10.1214/009117904000000414](https://doi.org/10.1214/009117904000000414). URL: <https://doi.org/10.1214/009117904000000414>.

[ABS22] Yeganeh Alimohammadi, Christian Borgs, and Amin Saberi. “Algorithms using local graph features to predict epidemics”. In: *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2022, pp. 3430–3451.

[ABS23] Yeganeh Alimohammadi, Christian Borgs, and Amin Saberi. “Locality of random digraphs on expanders”. In: *The Annals of Probability* 51.4 (2023), pp. 1249–1297.

[AKN87a] M. Aizenman, H. Kesten, and C. M. Newman. “Uniqueness of the infinite cluster and continuity of connectivity functions for short and long range percolation”. In: *Comm. Math. Phys.* 111.4 (1987), pp. 505–531. ISSN: 0010-3616. URL: <http://projecteuclid.org/euclid.cmp/1104159720>.

[AKN87b] Michael Aizenman, Harry Kesten, and Charles M Newman. “Uniqueness of the infinite cluster and continuity of connectivity functions for short and long range

percolation”. In: *Communications in Mathematical Physics* 111.4 (1987), pp. 505–531.

[AKS82a] M. Ajtai, J. Komlós, and E. Szemerédi. “Largest random component of a k -cube”. In: *Combinatorica* 2.1 (1982), pp. 1–7. ISSN: 0209-9683. doi: 10.1007/BF02579276. URL: <https://doi.org/10.1007/BF02579276>.

[AKS82b] M. Ajtai, J. Komlós, and E. Szemerédi. “Largest random component of a k -cube”. In: *Combinatorica* 2.1 (1982), pp. 1–7.

[AN84] Michael Aizenman and Charles M. Newman. “Tree graph inequalities and critical behavior in percolation models”. In: *J. Statist. Phys.* 36.1-2 (1984), pp. 107–143. ISSN: 0022-4715. doi: 10.1007/BF01015729. URL: <http://dx.doi.org/10.1007/BF01015729>.

[AS04] David Aldous and J Michael Steele. “The objective method: probabilistic combinatorial optimization and local weak convergence”. In: *Probability on discrete structures*. Springer, 2004, pp. 1–72.

[AS16] Noga Alon and Joel H. Spencer. *The probabilistic method*. Fourth. Wiley Series in Discrete Mathematics and Optimization. John Wiley & Sons, Inc., Hoboken, NJ, 2016, pp. xiv+375. ISBN: 978-1-119-06195-3.

[BA91] D. J. Barsky and M. Aizenman. “Percolation critical exponents under the triangle condition”. In: *Ann. Probab.* 19.4 (1991), pp. 1520–1536. ISSN: 0091-1798.

[BB99a] E. Babson and I. Benjamini. “Cut sets and normed cohomology with applications to percolation”. In: *Proceedings of the American Mathematical Society* 127.2 (1999), pp. 589–597.

[BB99b] Eric Babson and Itai Benjamini. “Cut sets and normed cohomology with applications to percolation”. In: *Proc. Amer. Math. Soc.* 127.2 (1999), pp. 589–597. ISSN: 0002-9939. doi: 10.1090/S0002-9939-99-04995-3.

[BC12] Itai Benjamini and Nicolas Curien. “Ergodic theory on stationary random graphs”. In: *Electron. J. Probab.* 17 (2012), no. 93, 20. ISSN: 1083-6489. doi: 10.1214/EJP.v17-2401. URL: <http://dx.doi.org/10.1214/EJP.v17-2401>.

[BDKY15] Itai Benjamini et al. “Disorder, entropy and harmonic functions”. In: *Ann. Probab.* 43.5 (2015), pp. 2332–2373. ISSN: 0091-1798. doi: 10.1214/14-AOP934. URL: <https://doi.org/10.1214/14-AOP934>.

[Ben+12] Itai Benjamini et al. “Sharp threshold for percolation on expanders”. In: *Ann. Probab.* 40.1 (2012), pp. 130–145. ISSN: 0091-1798. doi: 10.1214/10-AOP610. URL: <https://doi.org/10.1214/10-AOP610>.

[Ben01a] Itai Benjamini. “Percolation on finite graphs”. In: (2001). Unpublished note. Available at <https://arxiv.org/abs/math/0106022>.

[Ben01b] Itai Benjamini. *percolation on finite graphs*. 2001. arXiv: [math/0106022](https://arxiv.org/abs/math/0106022) [math.PR]. URL: <https://arxiv.org/abs/math/0106022>.

[Ben13a] Itai Benjamini. *Coarse geometry and randomness*. Vol. 2100. Lecture Notes in Mathematics. Lecture notes from the 41st Probability Summer School held in Saint-Flour, 2011, Chapter 5 is due to Nicolas Curien, Chapter 12 was written by Ariel Yadin, and Chapter 13 is joint work with Gady Kozma, École d’Été de Probabilités de Saint-Flour. [Saint-Flour Probability Summer School]. Springer, Cham, 2013, pp. viii+129. ISBN: 978-3-319-02575-9; 978-3-319-02576-6. doi: [10.1007/978-3-319-02576-6](https://doi.org/10.1007/978-3-319-02576-6). URL: <https://doi.org/10.1007/978-3-319-02576-6>.

[Ben13b] Itai Benjamini. *Coarse geometry and randomness*. Vol. 2100. Lecture Notes in Mathematics. Lecture notes from the 41st Probability Summer School held in Saint-Flour, 2011, Springer, Cham, 2013, pp. viii+129. ISBN: 978-3-319-02575-9; 978-3-319-02576-6. doi: [10.1007/978-3-319-02576-6](https://doi.org/10.1007/978-3-319-02576-6). URL: <https://doi.org/10.1007/978-3-319-02576-6>.

[Ber02] Noam Berger. “Transience, recurrence and critical behavior for long-range percolation”. In: *Comm. Math. Phys.* 226.3 (2002), pp. 531–558. ISSN: 0010-3616. doi: [10.1007/s002200200617](https://doi.org/10.1007/s002200200617). URL: <https://doi.org/10.1007/s002200200617>.

[BFT17] Itai Benjamini, Hilary Finucane, and Romain Tessera. “On the scaling limit of finite vertex transitive graphs with large diameter”. In: *Combinatorica* 37.3 (2017), pp. 333–374. ISSN: 0209-9683. doi: [10.1007/s00493-015-2975-4](https://doi.org/10.1007/s00493-015-2975-4). URL: <https://doi.org/10.1007/s00493-015-2975-4>.

[BG11] Emmanuel Breuillard and Ben Green. “Approximate groups. I The torsion-free nilpotent case”. In: *Journal of the Institute of Mathematics of Jussieu* 10.1 (2011), pp. 37–57.

[BGN91] David J. Barsky, Geoffrey R. Grimmett, and Charles M. Newman. “Percolation in half-spaces: equality of critical densities and continuity of the percolation probability”. In: *Probab. Theory Related Fields* 90.1 (1991), pp. 111–148. ISSN: 0178-8051, 1432-2064. doi: [10.1007/BF01321136](https://doi.org/10.1007/BF01321136). URL: <https://doi.org/10.1007/BF01321136>.

[BGT11] Emmanuel Breuillard, Ben Green, and Terence Tao. “Approximate subgroups of linear groups”. In: *Geom. Funct. Anal.* 21.4 (2011), pp. 774–819. ISSN: 1016-443X. doi: [10.1007/s00039-011-0122-y](https://doi.org/10.1007/s00039-011-0122-y). URL: <https://doi.org/10.1007/s00039-011-0122-y>.

[BGT12a] Emmanuel Breuillard, Ben Green, and Terence Tao. “The structure of approximate groups”. In: *Publ. Math. Inst. Hautes Études Sci.* 116 (2012), pp. 115–221. ISSN: 0073-8301. doi: [10.1007/s10240-012-0043-9](https://doi.org/10.1007/s10240-012-0043-9). URL: <https://doi.org/10.1007/s10240-012-0043-9>.

[BGT12b] Emmanuel Breuillard, Ben Green, and Terence Tao. “The structure of approximate groups”. In: *Publ. Math. Inst. Hautes Études Sci.* 116 (2012), pp. 115–221. URL: <https://doi.org/10.1007/s10240-012-0043-9>.

[BH57a] S. R. Broadbent and J. M. Hammersley. “Percolation processes. I. Crystals and mazes”. In: *Proc. Cambridge Philos. Soc.* 53 (1957), pp. 629–641. ISSN: 0008-1981. doi: [10.1017/s0305004100032680](https://doi.org/10.1017/s0305004100032680). URL: <https://doi.org/10.1017/s0305004100032680>.

[BH57b] S. R. Broadbent and J. M. Hammersley. “Percolation processes. I. Crystals and mazes”. In: *Proc. Cambridge Philos. Soc.* 53 (1957), pp. 629–641. ISSN: 0008-1981. doi: [10.1017/s0305004100032680](https://doi.org/10.1017/s0305004100032680). URL: <https://doi.org/10.1017/s0305004100032680>.

[BHT22] Nathanaël Berestycki, Jonathan Hermon, and Lucas Teyssier. “On the universality of fluctuations for the cover time”. In: *arXiv preprint arXiv:2202.02255* (2022).

[BK89] R. M. Burton and M. Keane. “Density and uniqueness in percolation”. In: *Comm. Math. Phys.* 121.3 (1989), pp. 501–505. ISSN: 0010-3616. URL: <http://projecteuclid.org/euclid.cmp/1104178143>.

[BKŁ92] B. Bollobás, Y. Kohayakawa, and T. Łuczak. “The evolution of random subgraphs of the cube”. In: *Random Structures Algorithms* 3.1 (1992), pp. 55–90. ISSN: 1042-9832. doi: [10.1002/rsa.3240030106](https://doi.org/10.1002/rsa.3240030106). URL: <https://doi.org/10.1002/rsa.3240030106>.

[BL13] Emmanuel Breuillard and Enrico Le Donne. “On the rate of convergence to the asymptotic cone for nilpotent groups and subFinsler geometry”. In: *Proceedings of the National Academy of Sciences* 110.48 (2013), pp. 19220–19226.

[BLS99] Itai Benjamini, Russell Lyons, and Oded Schramm. “Percolation perturbations in potential theory and random walks”. In: *Random walks and discrete potential theory (Cortona, 1997)*. Sympos. Math., XXXIX. Cambridge Univ. Press, Cambridge, 1999, pp. 56–84.

[BLU07] Andrea Bonfiglioli, Ermanno Lanconelli, and Francesco Uguzzoni. *Stratified Lie groups and potential theory for their sub-Laplacians*. Springer Science & Business Media, 2007.

[BNP11a] Itai Benjamini, Asaf Nachmias, and Yuval Peres. “Is the critical percolation probability local?” In: *Probab. Theory Related Fields* 149.1-2 (2011), pp. 261–269. ISSN: 0178-8051. doi: [10.1007/s00440-009-0251-5](https://doi.org/10.1007/s00440-009-0251-5). URL: <https://doi.org/10.1007/s00440-009-0251-5>.

[BNP11b] Itai Benjamini, Asaf Nachmias, and Yuval Peres. “Is the critical percolation probability local?” In: *Probability Theory and Related Fields* 149.1 (2011), pp. 261–269. doi: [10.1007/s00440-009-0251-5](https://doi.org/10.1007/s00440-009-0251-5). URL: <https://doi.org/10.1007/s00440-009-0251-5>.

[Bol+10a] Béla Bollobás et al. “Percolation on dense graph sequences”. In: *Ann. Probab.* 38.1 (2010), pp. 150–183. issn: 0091-1798. doi: 10.1214/09-AOP478. url: <https://doi.org/10.1214/09-AOP478>.

[Bol+10b] Béla Bollobás et al. “Percolation on dense graph sequences”. In: *Ann. Probab.* 38.1 (2010), pp. 150–183.

[Bol+10c] Béla Bollobás et al. “Percolation on dense graph sequences”. In: *Ann. Probab.* 38.1 (2010), pp. 150–183. issn: 0091-1798. doi: 10.1214/09-AOP478. url: <https://doi.org/10.1214/09-AOP478>.

[Bol84] Béla Bollobás. “The evolution of random graphs”. In: *Trans. Amer. Math. Soc.* 286.1 (1984), pp. 257–274. issn: 0002-9947. doi: 10.2307/1999405. url: <https://doi.org/10.2307/1999405>.

[Bor+05a] Christian Borgs et al. “Random subgraphs of finite graphs. I. The scaling window under the triangle condition”. In: *Random Structures Algorithms* 27.2 (2005), pp. 137–184. issn: 1042-9832. doi: 10.1002/rsa.20051. url: <https://doi.org/10.1002/rsa.20051>.

[Bor+05b] Christian Borgs et al. “Random subgraphs of finite graphs. II. The lace expansion and the triangle condition”. In: *Ann. Probab.* 33.5 (2005), pp. 1886–1944. issn: 0091-1798. doi: 10.1214/009117905000000260. url: <https://doi.org/10.1214/009117905000000260>.

[Bor+06] Christian Borgs et al. “Random subgraphs of finite graphs. III. The phase transition for the n -cube”. In: *Combinatorica* 26.4 (2006), pp. 395–410. issn: 0209-9683. doi: 10.1007/s00493-006-0022-1. url: <https://doi.org/10.1007/s00493-006-0022-1>.

[Bou+92] Jean Bourgain et al. “The influence of variables in product spaces”. In: *Israel J. Math.* 77.1-2 (1992), pp. 55–64. issn: 0021-2172. doi: 10.1007/BF02808010. url: <https://doi.org/10.1007/BF02808010>.

[Bre14] Emmanuel Breuillard. “A brief introduction to approximate groups”. In: *Thin groups and superstrong approximation* 61 (2014), pp. 23–50.

[Bri02] Martin R Bridson. “The geometry of the word problem”. In: *Invitations to geometry and topology* 7 (2002), pp. 29–91.

[BS96a] I. Benjamini and Oded Schramm. “Percolation beyond \mathbb{Z}^d , many questions and a few answers”. In: *Electron. Comm. Probab.* 1 (1996), no. 8, 71–82 (electronic). issn: 1083-589X. doi: 10.1214/ECP.v1-978. url: <http://dx.doi.org/10.1214/ECP.v1-978>.

[BS96b] Itai Benjamini and Oded Schramm. “Percolation beyond \mathbb{Z}^d , many questions and a few answers”. In: *Electron. Comm. Probab.* 1 (1996), no. 8, 71–82. issn: 1083-589X. doi: 10.1214/ECP.v1-978. url: <https://doi.org/10.1214/ECP.v1-978>.

[BS96c] Itai Benjamini and Oded Schramm. “Percolation beyond \mathbb{Z}^d , many questions and a few answers”. In: *Electron. Comm. Probab.* 1 (1996), no. 8, 71–82. issn: 1083-589X. doi: [10.1214/ECP.v1-978](https://doi.org/10.1214/ECP.v1-978). url: <https://doi.org/10.1214/ECP.v1-978>.

[BT16] Emmanuel Breuillard and Matthew C. H. Tointon. “Nilprogressions and groups with moderate growth”. In: *Adv. Math.* 289 (2016), pp. 1008–1055. issn: 0001-8708. doi: [10.1016/j.aim.2015.11.025](https://doi.org/10.1016/j.aim.2015.11.025). url: <https://doi.org/10.1016/j.aim.2015.11.025>.

[BT17] Itai Benjamini and Vincent Tassion. “Homogenization via sprinkling”. In: *Ann. Inst. Henri Poincaré Probab. Stat.* 53.2 (2017), pp. 997–1005. issn: 0246-0203. doi: [10.1214/16-AIHP746](https://doi.org/10.1214/16-AIHP746). url: <https://doi.org/10.1214/16-AIHP746>.

[BZ23] Christian Borgs and Geng Zhao. *Locality via Global Ties: Stability of the 2-Core Against Misspecification*. 2023. arXiv: 2304.06170 [cs.DS]. url: <https://arxiv.org/abs/2304.06170>.

[Car15] Pietro Kreitlon Carolino. *The Structure of Locally Compact Approximate Groups*. Thesis (Ph.D.)–University of California, Los Angeles. ProQuest LLC, Ann Arbor, MI, 2015, p. 98. isbn: 978-1339-09525-7. url: http://gateway.proquest.com/openurl?url_ver=Z39.88-2004&rft_val_fmt=info:ofi/fmt:kev:mtx:dissertation&res_dat=xri:pqm&rft_dat=xri:pqdiss:3725297.

[Car85] Thomas Keith Carne. “A transmutation formula for Markov chains”. In: *Bulletin des sciences mathématiques (Paris. 1885)* 109.4 (1985), pp. 399–405.

[Cas97] J. W. S. Cassels. *An introduction to the geometry of numbers*. Classics in Mathematics. Corrected reprint of the 1971 edition. Springer-Verlag, Berlin, 1997, pp. viii+344. isbn: 3-540-61788-4.

[CC87] J. T. Chayes and L. Chayes. “The mean field bound for the order parameter of Bernoulli percolation”. In: *Percolation theory and ergodic theory of infinite particle systems (Minneapolis, Minn., 1984–1985)*. Vol. 8. IMA Vol. Math. Appl. Springer, New York, 1987, pp. 49–71. doi: [10.1007/978-1-4613-8734-3_5](https://doi.org/10.1007/978-1-4613-8734-3_5). url: https://doi.org/10.1007/978-1-4613-8734-3_5.

[Cer15] Raphaël Cerf. “A lower bound on the two-arms exponent for critical percolation on the lattice”. In: *The Annals of Probability* 43.5 (Sept. 2015). doi: [10.1214/14-aop940](https://doi.org/10.1214/14-aop940). url: <https://doi.org/10.1214%2F14-aop940>.

[CMN12] Fernando Casas, Ander Murua, and Mladen Nadinic. “Efficient computation of the Zassenhaus formula”. In: *Computer Physics Communications* 183.11 (2012), pp. 2386–2391.

[CMT21] Daniel Contreras, Sébastien Martineau, and Vincent Tassion. “Supercritical percolation on graphs of polynomial growth”. In: *arXiv preprint arXiv:2107.06326* (2021).

[CMT22] Daniel Contreras, Sébastien Martineau, and Vincent Tassion. *Supercritical percolation on graphs of polynomial growth*. 2022. arXiv: 2107.06326 [math.PR]. URL: <https://arxiv.org/abs/2107.06326>.

[CMT23a] Daniel Contreras, Sébastien Martineau, and Vincent Tassion. “Locality of percolation for graphs with polynomial growth”. In: *Electron. Commun. Probab.* 28 (2023), Paper No. 1, 9. doi: 10.1214/22-ecp508. URL: <https://doi.org/10.1214/22-ecp508>.

[CMT23b] Daniel Contreras, Sébastien Martineau, and Vincent Tassion. “Locality of percolation for graphs with polynomial growth”. In: *Electronic Communications in Probability* 28 (2023), pp. 1–9.

[CMT24] Daniel Contreras, Sébastien Martineau, and Vincent Tassion. “Supercritical percolation on graphs of polynomial growth”. In: *Duke Mathematical Journal* 173.4 (2024), pp. 745–806. doi: 10.1215/00127094-2023-0032. URL: <https://doi.org/10.1215/00127094-2023-0032>.

[CS93] Thierry Coulhon and Laurent Saloff-Coste. “Isopérimétrie pour les groupes et les variétés”. In: *Rev. Mat. Iberoamericana* 9.2 (1993), pp. 293–314. ISSN: 0213-2230. doi: 10.4171/RMI/138. URL: <https://doi.org/10.4171/RMI/138>.

[DGRSY20] Hugo Duminil-Copin et al. “Existence of phase transition for percolation using the Gaussian free field”. In: *Duke Math. J.* 169.18 (2020), pp. 3539–3563. ISSN: 0012-7094. doi: 10.1215/00127094-2020-0036. URL: <https://doi.org/10.1215/00127094-2020-0036>.

[Dis+24] Sahar Diskin et al. *Percolation through Isoperimetry*. 2024. arXiv: 2308.10267 [math.CO]. URL: <https://arxiv.org/abs/2308.10267>.

[DK24a] Sahar Diskin and Michael Krivelevich. *Components, large and small, are as they should be I: supercritical percolation on regular graphs of growing degree*. 2024. arXiv: 2408.04597 [math.CO]. URL: <https://arxiv.org/abs/2408.04597>.

[DK24b] Sahar Diskin and Michael Krivelevich. *Components, large and small, are as they should be II: supercritical percolation on regular graphs of constant degree*. 2024. arXiv: 2408.04599 [math.CO]. URL: <https://arxiv.org/abs/2408.04599>.

[DKT21] Hugo Duminil-Copin, Gady Kozma, and Vincent Tassion. “Upper bounds on the percolation correlation length”. In: *In and Out of Equilibrium 3: Celebrating Vladas Sidoravicius* (2021), pp. 347–369.

[DRT19] Hugo Duminil-Copin, Aran Raoufi, and Vincent Tassion. “Sharp phase transition for the random-cluster and Potts models via decision trees”. In: *Ann. of Math.* (2) 189.1 (2019), pp. 75–99. ISSN: 0003-486X. doi: 10.4007/annals.2019.189.1.2. URL: <https://doi.org/10.4007/annals.2019.189.1.2>.

[DST14] Hugo Duminil-Copin, Vladas Sidoravicius, and Vincent Tassion. *Absence of infinite cluster for critical Bernoulli percolation on slabs*. 2014. arXiv: 1401.7130 [math.PR]. URL: <https://arxiv.org/abs/1401.7130>.

[DST16] Hugo Duminil-Copin, Vladas Sidoravicius, and Vincent Tassion. “Absence of infinite cluster for critical Bernoulli percolation on slabs”. In: *Comm. Pure Appl. Math.* 69.7 (2016), pp. 1397–1411. ISSN: 0010-3640,1097-0312. DOI: 10.1002/cpa.21641. URL: <https://doi.org/10.1002/cpa.21641>.

[DT16a] Hugo Duminil-Copin and Vincent Tassion. “A new proof of the sharpness of the phase transition for Bernoulli percolation and the Ising model”. In: *Communications in Mathematical Physics* 343 (2016), pp. 725–745.

[DT16b] Hugo Duminil-Copin and Vincent Tassion. “A new proof of the sharpness of the phase transition for Bernoulli percolation and the Ising model”. In: *Comm. Math. Phys.* 343.2 (2016), pp. 725–745. ISSN: 0010-3616. DOI: 10.1007/s00220-015-2480-z. URL: <https://doi.org/10.1007/s00220-015-2480-z>.

[DT16c] Hugo Duminil-Copin and Vincent Tassion. “A new proof of the sharpness of the phase transition for Bernoulli percolation and the Ising model”. In: *Comm. Math. Phys.* 343.2 (2016), pp. 725–745. ISSN: 0010-3616. DOI: 10.1007/s00220-015-2480-z. URL: <https://doi.org/10.1007/s00220-015-2480-z>.

[DT19] Hugo Duminil-Copin and Vincent Tassion. “A note on Schramm’s locality conjecture for random-cluster models”. In: *Sojourns in Probability Theory and Statistical Physics-II: Brownian Web and Percolation, A Festschrift for Charles M. Newman*. Springer, 2019, pp. 123–134.

[DT22] Barbara Dembin and Vincent Tassion. *Almost sharp sharpness for Poisson Boolean percolation*. 2022. arXiv: 2209.00999 [math.PR]. URL: <https://arxiv.org/abs/2209.00999>.

[Dum+20a] Hugo Duminil-Copin et al. “Existence of phase transition for percolation using the Gaussian free field”. In: *Duke Mathematical Journal* 169.18 (Dec. 2020), pp. 3539–3563. DOI: 10.1215/00127094-2020-0036. URL: <https://doi.org/10.1215/00127094-2020-0036>.

[Dum+20b] Hugo Duminil-Copin et al. “Existence of phase transition for percolation using the Gaussian free field”. In: *Duke Math. J.* 169.18 (2020), pp. 3539–3563. ISSN: 0012-7094. DOI: 10.1215/00127094-2020-0036. URL: <https://doi.org/10.1215/00127094-2020-0036>.

[Dum+20c] Hugo Duminil-Copin et al. “Existence of phase transition for percolation using the gaussian free field”. In: *Duke Mathematical Journal* 169.18 (2020), pp. 3539–3563.

[Eas22] Philip Easo. “Existence of a percolation threshold on finite transitive graphs”. In: *arXiv preprint arXiv:2208.09501* (2022).

[Eas23] Philip Easo. “Existence of a percolation threshold on finite transitive graphs”. In: *Int. Math. Res. Not. IMRN* 21 (2023), pp. 18781–18802. issn: 1073-7928,1687-0247. doi: 10.1093/imrn/rnad222. url: <https://doi.org/10.1093/imrn/rnad222>.

[Eas24] Philip Easo. *Sharpness and locality for percolation on finite transitive graphs*. 2024. arXiv: 2408.12039 [math.PR]. url: <https://arxiv.org/abs/2408.12039>.

[EH21a] Philip Easo and Tom Hutchcroft. *Supercritical percolation on finite transitive graphs I: Uniqueness of the giant component*. 2021. arXiv: 2112.12778 [math.PR]. url: <https://arxiv.org/abs/2112.12778>.

[EH21b] Philip Easo and Tom Hutchcroft. *Supercritical percolation on finite transitive graphs I: Uniqueness of the giant component*. 2021. doi: 10.48550/ARXIV.2112.12778. url: <https://arxiv.org/abs/2112.12778>.

[EH23a] Philip Easo and Tom Hutchcroft. *The critical percolation probability is local*. 2023. arXiv: 2310.10983 [math.PR]. url: <https://arxiv.org/abs/2310.10983>.

[EH23b] Philip Easo and Tom Hutchcroft. *The critical percolation probability is local*. 2023. arXiv: 2310.10983 [math.PR].

[EH23c] Phillip Easo and Tom Hutchcroft. “The critical percolation probability is local”. In: (2023).

[EH23d] Phillip Easo and Tom Hutchcroft. *Uniform finite presentation for groups of polynomial growth*. 2023. arXiv: 2308.12428 [math.GR]. url: <https://arxiv.org/abs/2308.12428>.

[EH23+a] Philip Easo and Tom Hutchcroft. *Supercritical percolation on finite transitive graphs II: Concentration, locality, and equicontinuity of the giant’s density*. In preparation.

[EH23+b] Philip Easo and Tom Hutchcroft. *Quantitative locality and the computability of p_c* . In preparation.

[EK10] Anna Erschler and Anders Karlsson. “Homomorphisms to \mathbb{R}^d constructed from random walks”. In: *Annales de l’Institut Fourier*. Vol. 60. 6. 2010, pp. 2095–2113.

[ER59] P. Erdős and A. Rényi. “On random graphs. I”. In: *Publ. Math. Debrecen* 6 (1959), pp. 290–297. issn: 0033-3883.

[ER60] P. Erdős and A. Rényi. “On the evolution of random graphs”. In: *Magyar Tud. Akad. Mat. Kutató Int. Közl.* 5 (1960), pp. 17–61. issn: 0541-9514.

[ER61] P. Erdős and A. Rényi. “On the evolution of random graphs”. In: *Bull. Inst. Internat. Statist.* 38 (1961), pp. 343–347.

[FGO15] Louis Funar, Martha Giannoudovardi, and Daniele Ettore Otera. “On groups with linear sci growth”. In: *Fund. Math.* 228.1 (2015), pp. 47–62. issn: 0016-2736. doi: 10.4064/fm228-1-4. url: <https://doi.org/10.4064/fm228-1-4>.

[FK96] Ehud Friedgut and Gil Kalai. “Every monotone graph property has a sharp threshold”. In: *Proc. Amer. Math. Soc.* 124.10 (1996), pp. 2993–3002. issn: 0002-9939. doi: [10.1090/S0002-9939-96-03732-X](https://doi.org/10.1090/S0002-9939-96-03732-X). URL: <https://doi.org/10.1090/S0002-9939-96-03732-X>.

[FKM04] Alan Frieze, Michael Krivelevich, and Ryan Martin. “The emergence of a giant component in random subgraphs of pseudo-random graphs”. In: *Random Structures Algorithms* 24.1 (2004), pp. 42–50. issn: 1042-9832. doi: [10.1002/rsa.10100](https://doi.org/10.1002/rsa.10100). URL: <https://doi.org/10.1002/rsa.10100>.

[FKP09] David Fisher, Nets Hawk Katz, and Irine Peng. “On Freiman’s theorem in nilpotent groups”. In: *arXiv preprint arXiv:0901.1409* (2009).

[Fri98] Ehud Friedgut. “Boolean functions with low average sensitivity depend on few coordinates”. In: *Combinatorica* 18.1 (1998), pp. 27–35. issn: 0209-9683,1439-6912. doi: [10.1007/PL00009809](https://doi.org/10.1007/PL00009809). URL: <https://doi.org/10.1007/PL00009809>.

[Fri99a] Ehud Friedgut. “Sharp thresholds of graph properties, and the k -sat problem”. In: *J. Amer. Math. Soc.* 12.4 (1999). With an appendix by Jean Bourgain, pp. 1017–1054. issn: 0894-0347. doi: [10.1090/S0894-0347-99-00305-7](https://doi.org/10.1090/S0894-0347-99-00305-7). URL: <https://doi.org/10.1090/S0894-0347-99-00305-7>.

[Fri99b] Ehud Friedgut. “Sharp thresholds of graph properties, and the k -sat problem”. In: *J. Amer. Math. Soc.* 12.4 (1999). With an appendix by Jean Bourgain, pp. 1017–1054. issn: 0894-0347,1088-6834. doi: [10.1090/S0894-0347-99-00305-7](https://doi.org/10.1090/S0894-0347-99-00305-7). URL: <https://doi.org/10.1090/S0894-0347-99-00305-7>.

[GHR03] Steve M Gersten, Derek F Holt, and Tim R Riley. “Isoperimetric inequalities for nilpotent groups”. In: *Geometric & Functional Analysis GAFA* 13 (2003), pp. 795–814.

[GL17] Geoffrey Grimmett and Zhongyang Li. “Connective constants and height functions for Cayley graphs”. In: *Transactions of the American Mathematical Society* 369.8 (2017), pp. 5961–5980.

[GL18] Geoffrey R Grimmett and Zhongyang Li. “Locality of connective constants”. In: *Discrete Mathematics* 341.12 (2018), pp. 3483–3497.

[GL20] Geoffrey R Grimmett and Zhongyang Li. “Cubic graphs and the golden mean”. In: *Discrete Mathematics* 343.1 (2020), p. 111638.

[Gle52] Andrew M. Gleason. “Groups without small subgroups”. In: *Ann. of Math.* (2) 56 (1952), pp. 193–212. issn: 0003-486X. doi: [10.2307/1969795](https://doi.org/10.2307/1969795). URL: <https://doi.org/10.2307/1969795>.

[GLL21] Anupam Gupta, Euiwoong Lee, and Jason Li. “The connectivity threshold for dense graphs”. In: *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. [Society for Industrial and Applied Mathematics (SIAM)], Philadelphia,

PA, 2021, pp. 89–105. doi: [10.1137/1.9781611976465.7](https://doi.org/10.1137/1.9781611976465.7). URL: <https://doi.org/10.1137/1.9781611976465.7>.

[GM90a] G. R. Grimmett and J. M. Marstrand. “The supercritical phase of percolation is well behaved”. In: *Proc. Roy. Soc. London Ser. A* 430.1879 (1990), pp. 439–457. ISSN: 0962-8444. doi: [10.1098/rspa.1990.0100](https://doi.org/10.1098/rspa.1990.0100). URL: <https://doi.org/10.1098/rspa.1990.0100>.

[GM90b] G. R. Grimmett and J. M. Marstrand. “The supercritical phase of percolation is well behaved”. In: *Proc. Roy. Soc. London Ser. A* 430.1879 (1990), pp. 439–457. ISSN: 0962-8444. doi: [10.1098/rspa.1990.0100](https://doi.org/10.1098/rspa.1990.0100). URL: <https://doi.org/10.1098/rspa.1990.0100>.

[GP23] Agelos Georgakopoulos and Christoforos Panagiotis. “Analyticity Results in Bernoulli Percolation”. In: *Mem. Amer. Math. Soc.* 288.1431 (2023). ISSN: 0065-9266. doi: [10.1090/memo/1431](https://doi.org/10.1090/memo/1431). URL: <https://doi.org/10.1090/memo/1431>.

[Gri06] Geoffrey Grimmett. *The random-cluster model*. Vol. 333. Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Springer-Verlag, Berlin, 2006, pp. xiv+377. ISBN: 978-3-540-32890-2; 3-540-32890-4. doi: [10.1007/978-3-540-32891-9](https://doi.org/10.1007/978-3-540-32891-9). URL: <https://doi.org/10.1007/978-3-540-32891-9>.

[Gri14] Rostislav Grigorchuk. “On the gap conjecture concerning group growth”. In: *Bulletin of Mathematical Sciences* 4.1 (2014), pp. 113–128.

[Gri18] Geoffrey Grimmett. *Probability on graphs*. Vol. 8. Institute of Mathematical Statistics Textbooks. Random processes on graphs and lattices, Second edition of [MR2723356]. Cambridge University Press, Cambridge, 2018, pp. xi+265. ISBN: 978-1-108-43817-9. doi: [10.1017/9781108528986](https://doi.org/10.1017/9781108528986). URL: <https://doi.org/10.1017/9781108528986>.

[Gri99] Geoffrey Grimmett. *Percolation*. Second. Vol. 321. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Springer-Verlag, Berlin, 1999, pp. xiv+444. ISBN: 3-540-64902-6. URL: <https://doi.org/10.1007/978-3-662-03981-6>.

[Gro81a] M. Gromov. “Groups of polynomial growth and expanding maps”. In: *Publications Mathématiques de l’Institut des Hautes Études Scientifiques* 53.1 (1981), pp. 53–78.

[Gro81b] Mikhael Gromov. “Groups of polynomial growth and expanding maps”. In: *Inst. Hautes Études Sci. Publ. Math.* 53 (1981), pp. 53–73. ISSN: 0073-8301. URL: http://www.numdam.org/item?id=PMIHES_1981__53__53_0.

[Gro81c] Mikhael Gromov. “Groups of polynomial growth and expanding maps”. In: *Publ. Math. Inst. Hautes Études Sci.* 53 (1981), pp. 53–73. URL: <https://doi.org/10.1007/BF02698687>.

[Har60] T. E. Harris. “A lower bound for the critical probability in a certain percolation process”. In: *Proc. Cambridge Philos. Soc.* 56 (1960), pp. 13–20.

[Hat12a] Hamed Hatami. “A structure theorem for Boolean functions with small total influences”. In: *Ann. of Math.* (2) 176.1 (2012), pp. 509–533. issn: 0003-486X. doi: [10.4007/annals.2012.176.1.9](https://doi.org/10.4007/annals.2012.176.1.9). url: <https://doi.org/10.4007/annals.2012.176.1.9>.

[Hat12b] Hamed Hatami. “A structure theorem for Boolean functions with small total influences”. In: *Ann. of Math.* (2) 176.1 (2012), pp. 509–533. issn: 0003-486X, 1939-8980. doi: [10.4007/annals.2012.176.1.9](https://doi.org/10.4007/annals.2012.176.1.9). url: <https://doi.org/10.4007/annals.2012.176.1.9>.

[HH19] Jonathan Hermon and Tom Hutchcroft. “Supercritical percolation on nonamenable graphs: Isoperimetry, analyticity, and exponential decay of the cluster size distribution”. In: (2019). Preprint. Available at <https://arxiv.org/abs/1904.10448>.

[HH21a] Jonathan Hermon and Tom Hutchcroft. “No percolation at criticality on certain groups of intermediate growth”. In: *International Mathematics Research Notices* 2021.22 (2021), pp. 17433–17455.

[HH21b] Jonathan Hermon and Tom Hutchcroft. “Supercritical percolation on nonamenable graphs: isoperimetry, analyticity, and exponential decay of the cluster size distribution”. In: *Inventiones mathematicae* 224.2 (2021), pp. 445–486. doi: [10.1007/s00222-020-01011-3](https://doi.org/10.1007/s00222-020-01011-3).

[HH21c] Jonathan Hermon and Tom Hutchcroft. “Supercritical percolation on nonamenable graphs: isoperimetry, analyticity, and exponential decay of the cluster size distribution”. In: *Invent. Math.* 224.2 (2021), pp. 445–486. issn: 0020-9910, 1432-1297. doi: [10.1007/s00222-020-01011-3](https://doi.org/10.1007/s00222-020-01011-3). url: <https://doi.org/10.1007/s00222-020-01011-3>.

[HM22] Benjamin T. Hansen and Tobias Müller. “The critical probability for Voronoi percolation in the hyperbolic plane tends to $1/2$ ”. In: *Random Structures Algorithms* 60.1 (2022), pp. 54–67. issn: 1042-9832. doi: [10.1002/rsa.21018](https://doi.org/10.1002/rsa.21018). url: <https://doi.org/10.1002/rsa.21018>.

[Hof21] Remco van der Hofstad. *The giant in random graphs is almost local*. 2021. arXiv: [2103.11733](https://arxiv.org/abs/2103.11733) [math.PR]. url: <https://arxiv.org/abs/2103.11733>.

[HR06] Remco van der Hofstad and Frank Redig. “Maximal clusters in non-critical percolation and related models”. In: *J. Stat. Phys.* 122.4 (2006), pp. 671–703. issn: 0022-4715. doi: [10.1007/s10955-005-8012-z](https://doi.org/10.1007/s10955-005-8012-z). url: <https://doi.org/10.1007/s10955-005-8012-z>.

[Hru12] Ehud Hrushovski. “Stable group theory and approximate subgroups”. In: *J. Amer. Math. Soc.* 25.1 (2012), pp. 189–243. issn: 0894-0347. doi: [10.1090/S0894-0347-1100001-2](https://doi.org/10.1090/S0894-0347-1100001-2).

0347-2011-00708-X. URL: <https://doi.org/10.1090/S0894-0347-2011-00708-X>.

[HS93] Waldemar Hebisch and Laurent Saloff-Coste. “Gaussian estimates for Markov chains and random walks on groups”. In: *The Annals of Probability* (1993), pp. 673–709.

[HT21a] Tom Hutchcroft and Matthew Tointon. *Non-triviality of the phase transition for percolation on finite transitive graphs*. 2021. arXiv: 2104.05607 [math.PR]. URL: <https://arxiv.org/abs/2104.05607>.

[HT21b] Tom Hutchcroft and Matthew Tointon. “Non-triviality of the phase transition for percolation on finite transitive graphs”. In: *arXiv preprint arXiv:2104.05607* (2021).

[HT21c] Tom Hutchcroft and Matthew Tointon. “Non-triviality of the phase transition for percolation on finite transitive graphs”. In: (2021). arXiv: 2104.05607 [math.PR].

[HT21d] Tom Hutchcroft and Matthew Tointon. *Non-triviality of the phase transition for percolation on finite transitive graphs*. 2021. doi: 10.48550/ARXIV.2104.05607. URL: <https://arxiv.org/abs/2104.05607>.

[HT21e] Tom Hutchcroft and Matthew Tointon. “Non-triviality of the phase transition for percolation on finite transitive graphs”. In: *to appear in Journal of the European Mathematical Society* available at arXiv:2104.05607 (2021).

[Hut] Tom Hutchcroft. “On the derivation of mean-field percolation critical exponents from the triangle condition”. In: *J. Stat. Phys.* (). To appear. Available at <https://arxiv.org/abs/2106.06400>.

[Hut16] Tom Hutchcroft. “Critical percolation on any quasi-transitive graph of exponential growth has no infinite clusters”. In: *Comptes Rendus. Mathématique* 354.9 (2016), pp. 944–947.

[Hut19a] Tom Hutchcroft. “Percolation on hyperbolic graphs”. In: *Geom. Funct. Anal.* 29.3 (2019), pp. 766–810. issn: 1016-443X. doi: 10.1007/s00039-019-00498-0. URL: <https://doi.org/10.1007/s00039-019-00498-0>.

[Hut19b] Tom Hutchcroft. “Percolation on hyperbolic graphs”. In: *Geom. Funct. Anal.* 29.3 (2019), pp. 766–810. issn: 1016-443X. doi: 10.1007/s00039-019-00498-0. URL: <https://doi.org/10.1007/s00039-019-00498-0>.

[Hut20a] Tom Hutchcroft. “Locality of the critical probability for transitive graphs of exponential growth”. In: *Ann. Probab.* 48.3 (2020), pp. 1352–1371. issn: 0091-1798. doi: 10.1214/19-AOP1395. URL: <https://doi.org/10.1214/19-AOP1395>.

[Hut20b] Tom Hutchcroft. “Locality of the critical probability for transitive graphs of exponential growth”. In: *Ann. Probab.* 48.3 (2020), pp. 1352–1371. issn: 0091-1798. doi: 10.1214/19-AOP1395. URL: <https://doi.org/10.1214/19-AOP1395>.

[Hut20c] Tom Hutchcroft. “New critical exponent inequalities for percolation and the random cluster model”. In: *Probability and Mathematical Physics* (2020). To appear. Available at <https://arXiv.org/abs/1901.10363>.

[Hut20d] Tom Hutchcroft. “New critical exponent inequalities for percolation and the random cluster model”. In: *Probab. Math. Phys.* 1.1 (2020), pp. 147–165. issn: 2690-0998. doi: [10.2140/pmp.2020.1.147](https://doi.org/10.2140/pmp.2020.1.147). url: <https://doi.org/10.2140/pmp.2020.1.147>.

[Hut20e] Tom Hutchcroft. “Nonuniqueness and mean-field criticality for percolation on nonunimodular transitive graphs”. In: *J. Amer. Math. Soc.* 33.4 (2020), pp. 1101–1165. issn: 0894-0347. doi: [10.1090/jams/953](https://doi.org/10.1090/jams/953). url: <https://doi.org/10.1090/jams/953>.

[Hut20f] Tom Hutchcroft. “Nonuniqueness and mean-field criticality for percolation on nonunimodular transitive graphs”. In: *J. Amer. Math. Soc.* 33.4 (2020), pp. 1101–1165. issn: 0894-0347. doi: [10.1090/jams/953](https://doi.org/10.1090/jams/953). url: <https://doi.org/10.1090/jams/953>.

[Hut20g] Tom Hutchcroft. “Nonuniqueness and mean-field criticality for percolation on nonunimodular transitive graphs”. In: *J. Amer. Math. Soc.* 33.4 (2020), pp. 1101–1165. issn: 0894-0347. doi: [10.1090/jams/953](https://doi.org/10.1090/jams/953). url: <https://doi.org/10.1090/jams/953>.

[Hut20h] Tom Hutchcroft. “Slightly supercritical percolation on nonamenable graphs I: The distribution of finite clusters”. In: *arXiv preprint* (2020). Available at <https://arxiv.org/abs/2002.02916>.

[Hut20i] Tom Hutchcroft. “Universality of high-dimensional spanning forests and sandpiles”. In: *Probab. Theory Related Fields* 176.1-2 (2020), pp. 533–597. issn: 0178-8051. doi: [10.1007/s00440-019-00923-3](https://doi.org/10.1007/s00440-019-00923-3). url: <https://doi.org/10.1007/s00440-019-00923-3>.

[Hut21a] Tom Hutchcroft. “Power-law bounds for critical long-range percolation below the upper-critical dimension”. In: *Probability theory and related fields* 181 (2021), pp. 533–570.

[Hut21b] Tom Hutchcroft. “Power-law bounds for critical long-range percolation below the upper-critical dimension”. In: *Probab. Theory Related Fields* 181.1-3 (2021), pp. 533–570. issn: 0178-8051. doi: [10.1007/s00440-021-01043-7](https://doi.org/10.1007/s00440-021-01043-7). url: <https://doi.org/10.1007/s00440-021-01043-7>.

[Hut21c] Tom Hutchcroft. “Power-law bounds for critical long-range percolation below the upper-critical dimension”. In: *Probab. Theory Related Fields* 181.1-3 (2021), pp. 533–570. issn: 0178-8051. doi: [10.1007/s00440-021-01043-7](https://doi.org/10.1007/s00440-021-01043-7). url: <https://doi.org/10.1007/s00440-021-01043-7>.

[Hut21d] Tom Hutchcroft. “Power-law bounds for critical long-range percolation below the upper-critical dimension”. In: *Probability Theory and Related Fields* 181.1 (2021), pp. 533–570. doi: [10.1007/s00440-021-01043-7](https://doi.org/10.1007/s00440-021-01043-7). url: <https://doi.org/10.1007/s00440-021-01043-7>.

[Hut22] Tom Hutchcroft. “Slightly supercritical percolation on non-amenable graphs I: The distribution of finite clusters”. In: *Proc. Lond. Math. Soc. (3)* 125.4 (2022), pp. 968–1013. issn: 0024-6115.

[Hut23a] Tom Hutchcroft. “Transience and anchored isoperimetric dimension of supercritical percolation clusters”. In: *Electronic Journal of Probability* 28.none (2023), pp. 1–15. doi: [10.1214/23-EJP905](https://doi.org/10.1214/23-EJP905). url: <https://doi.org/10.1214/23-EJP905>.

[Hut23b] Tom Hutchcroft. “Transience and anchored isoperimetric dimension of supercritical percolation clusters”. In: *Electronic Journal of Probability* 28 (2023), pp. 1–15.

[Kes90] Harry Kesten. “Asymptotics in high dimensions for percolation”. In: *Disorder in physical systems* (1990), pp. 219–240.

[Kle10] Bruce Kleiner. “A new proof of Gromov’s theorem on groups of polynomial growth”. In: *J. Amer. Math. Soc.* 23.3 (2010), pp. 815–829. issn: 0894-0347. doi: [10.1090/S0894-0347-09-00658-4](https://doi.org/10.1090/S0894-0347-09-00658-4). url: <https://doi.org/10.1090/S0894-0347-09-00658-4>.

[KLS20] Michael Krivelevich, Eyal Lubetzky, and Benny Sudakov. “Asymptotics in percolation on high-girth expanders”. In: *Random Structures & Algorithms* 56.4 (2020), pp. 927–947.

[KS96] David R. Karger and Clifford Stein. “A new approach to the minimum cut problem”. In: *J. ACM* 43.4 (1996), pp. 601–640. issn: 0004-5411. doi: [10.1145/234533.234534](https://doi.org/10.1145/234533.234534). url: <https://doi.org/10.1145/234533.234534>.

[Laz54] Michel Lazard. “Sur les groupes nilpotents et les anneaux de Lie”. In: *Annales scientifiques de l’École Normale Supérieure*. Vol. 71. 2. 1954, pp. 101–190.

[LLMS23] Lyuben Lichev et al. “On the first and second largest components in the percolated random geometric graph”. In: *Stochastic Processes and their Applications* 164 (2023), pp. 311–336.

[LMS08] Russell Lyons, Benjamin J. Morris, and Oded Schramm. “Ends in uniform spanning forests”. In: *Electron. J. Probab.* 13 (2008), no. 58, 1702–1725. doi: [10.1214/EJP.v13-566](https://doi.org/10.1214/EJP.v13-566). url: <https://doi.org/10.1214/EJP.v13-566>.

[LMTT23] Russell Lyons et al. “Explicit universal minimal constants for polynomial growth of groups”. In: *J. Group Theory* 26.1 (2023), pp. 29–53. issn: 1433-5883. doi: [10.1515/jgth-2020-0202](https://doi.org/10.1515/jgth-2020-0202). url: <https://doi.org/10.1515/jgth-2020-0202>.

[LP16a] Russell Lyons and Yuval Peres. *Probability on Trees and Networks*. Vol. 42. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, New York, 2016, pp. xv+699. ISBN: 978-1-107-16015-6. doi: 10.1017/9781316672815. URL: <http://dx.doi.org/10.1017/9781316672815>.

[LP16b] Russell Lyons and Yuval Peres. *Probability on Trees and Networks*. Vol. 42. Cambridge Series in Statistical and Probabilistic Mathematics. Available at <http://pages.iu.edu/~rdlyons/>. Cambridge University Press, New York, 2016, pp. xv+699. ISBN: 978-1-107-16015-6. doi: 10.1017/9781316672815. URL: <http://dx.doi.org/10.1017/9781316672815>.

[LP16c] Russell Lyons and Yuval Peres. *Probability on trees and networks*. Vol. 42. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, New York, 2016, pp. xv+699. ISBN: 978-1-107-16015-6. doi: 10.1017/9781316672815. URL: <https://doi.org/10.1017/9781316672815>.

[LS99] Russell Lyons and Oded Schramm. “Indistinguishability of percolation clusters”. In: *Ann. Probab.* 27.4 (1999), pp. 1809–1836. ISSN: 0091-1798. doi: 10.1214/aop/1022677549. URL: <https://doi.org/10.1214/aop/1022677549>.

[Lyo+20] Russell Lyons et al. “Occupation measure of random walks and wired spanning forests in balls of Cayley graphs”. In: *Ann. Fac. Sci. Toulouse Math. (6)* 29.1 (2020), pp. 97–109. ISSN: 0240-2963,2258-7519. doi: 10.5802/afst.1625. URL: <https://doi.org/10.5802/afst.1625>.

[Lyo+23a] Russell Lyons et al. “Explicit universal minimal constants for polynomial growth of groups”. In: *Journal of Group Theory* 26.1 (2023), pp. 29–53. doi: doi:10.1515/jgth-2020-0202. URL: <https://doi.org/10.1515/jgth-2020-0202>.

[Lyo+23b] Russell Lyons et al. “Explicit universal minimal constants for polynomial growth of groups”. In: *Journal of Group Theory* 26.1 (2023), pp. 29–53.

[Lyo95] Russell Lyons. “Random walks and the growth of groups”. In: *Comptes rendus de l’Académie des sciences. Série 1, Mathématique* 320.11 (1995), pp. 1361–1366.

[Mad71] W. Mader. “Minimale n -fach kantenzusammenhängende Graphen”. In: *Math. Ann.* 191 (1971), pp. 21–28. ISSN: 0025-5831. doi: 10.1007/BF01433466. URL: <https://doi.org/10.1007/BF01433466>.

[Men86] M. V. Menshikov. “Coincidence of critical points in percolation problems”. In: *Dokl. Akad. Nauk SSSR* 288.6 (1986), pp. 1308–1311. ISSN: 0002-3264.

[MP01] Roman Muchnik and Igor Pak. “Percolation on Grigorchuk groups”. In: *Comm. Algebra* 29.2 (2001), pp. 661–671. ISSN: 0092-7872. doi: 10.1081/AGB-100001531. URL: <https://doi.org/10.1081/AGB-100001531>.

[MP05] B. Morris and Yuval Peres. “Evolving sets, mixing and heat kernel bounds”. In: *Probab. Theory Related Fields* 133.2 (2005), pp. 245–266. ISSN: 0178-8051. doi: [10.1007/s00440-005-0434-7](https://doi.org/10.1007/s00440-005-0434-7). URL: <https://doi.org/10.1007/s00440-005-0434-7>.

[MS19] Sébastien Martineau and Franco Severo. “Strict monotonicity of percolation thresholds under covering maps”. In: *The Annals of Probability* 47.6 (2019).

[MS96] Ronald Meester and Jeffrey E. Steif. “On the continuity of the critical value for long range percolation in the exponential case”. In: *Comm. Math. Phys.* 180.2 (1996), pp. 483–504. ISSN: 0010-3616. URL: <http://projecteuclid.org/euclid.cmp/1104287357>.

[MT17] Sébastien Martineau and Vincent Tassion. “Locality of percolation for Abelian Cayley graphs”. In: *Ann. Probab.* 45.2 (2017), pp. 1247–1277. ISSN: 0091-1798. doi: [10.1214/15-AOP1086](https://doi.org/10.1214/15-AOP1086). URL: <https://doi.org/10.1214/15-AOP1086>.

[MZ52] Deane Montgomery and Leo Zippin. “Small subgroups of finite-dimensional groups”. In: *Ann. of Math.* (2) 56 (1952), pp. 213–241. ISSN: 0003-486X. doi: [10.2307/1969796](https://doi.org/10.2307/1969796). URL: <https://doi.org/10.2307/1969796>.

[Nac09] Asaf Nachmias. “Mean-field conditions for percolation on finite graphs”. In: *Geom. Funct. Anal.* 19.4 (2009), pp. 1171–1194. ISSN: 1016-443X. doi: [10.1007/s00039-009-0032-4](https://doi.org/10.1007/s00039-009-0032-4). URL: <https://doi.org/10.1007/s00039-009-0032-4>.

[ODo14] Ryan O’Donnell. *Analysis of Boolean functions*. Cambridge University Press, New York, 2014, pp. xx+423. ISBN: 978-1-107-03832-5. doi: [10.1017/CBO9781139814782](https://doi.org/10.1017/CBO9781139814782). URL: <https://doi.org/10.1017/CBO9781139814782>.

[Oza18] Narutaka Ozawa. “A functional analysis proof of Gromov’s polynomial growth theorem”. In: *Ann. Sci. Éc. Norm. Supér.(4)* 51.3 (2018), pp. 549–556.

[Pan83] Pierre Pansu. “Croissance des boules et des géodésiques fermées dans les nilvariétés”. In: *Ergodic Theory and Dynamical Systems* 3.3 (1983), pp. 415–445.

[Pei36a] R. Peierls. “On Ising’s model of ferromagnetism.” In: *Math. Proc. Camb. Phil. Soc.* 32 (1936), pp. 477–481.

[Pei36b] Rudolf Peierls. “On Ising’s model of ferromagnetism”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 32.3 (1936), pp. 477–481.

[Pen93] Mathew D Penrose. “On the spread-out limit for bond and continuum percolation”. In: *The Annals of Applied Probability* 3.1 (1993), pp. 253–276.

[Pet] Gabor Pete. *Probability and geometry on groups*. URL: <http://www.math.bme.hu/%20gabor/PGG.pdf>.

[Pet08] Gabor Pete. “A note on percolation on Z^d : isoperimetric profile via exponential cluster repulsion”. In: *Electronic Communications in Probability* 13.none (2008), pp. 377–392. doi: 10.1214/ECP.v13-1390. url: <https://doi.org/10.1214/ECP.v13-1390>.

[Pet14] Gabor Pete. “Probability and Geometry on Groups”. <http://www.math.bme.hu/gabor/PGG.pdf>. 2014.

[Pet23] Anton Petrunin. *Pure metric geometry*. SpringerBriefs in Mathematics. Springer, Cham, [2023] ©2023, pp. viii+103. isbn: 978-3-031-39161-3; 978-3-031-39162-0. doi: 10.1007/978-3-031-39162-0. url: <https://doi.org/10.1007/978-3-031-39162-0>.

[PP96] Robin Pemantle and Yuval Peres. “On which graphs are all random walks in random environments transient?” In: *Random discrete structures (Minneapolis, MN, 1993)*. Vol. 76. IMA Vol. Math. Appl. Springer, New York, 1996, pp. 207–211. isbn: 0-387-94623-3. doi: 10.1007/978-1-4612-0719-1_14. url: https://doi.org/10.1007/978-1-4612-0719-1_14.

[PS00] Igor Pak and Tatiana Smirnova-Nagnibeda. “On non-uniqueness of percolation on nonamenable Cayley graphs”. In: *C. R. Acad. Sci. Paris Sér. I Math.* 330.6 (2000), pp. 495–500. issn: 0764-4442. doi: 10.1016/S0764-4442(00)00211-1. url: [http://dx.doi.org/10.1016/S0764-4442\(00\)00211-1](http://dx.doi.org/10.1016/S0764-4442(00)00211-1).

[PS23a] Christoforos Panagiotis and Franco Severo. “Gap at 1 for the percolation threshold of Cayley graphs”. In: *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques* 59.3 (2023), pp. 1248–1258. doi: 10.1214/22-AIHP1286. url: <https://doi.org/10.1214/22-AIHP1286>.

[PS23b] Christoforos Panagiotis and Franco Severo. “Gap at 1 for the percolation threshold of Cayley graphs”. In: *Ann. Inst. Henri Poincaré Probab. Stat.* 59.3 (2023), pp. 1248–1258. issn: 0246-0203. doi: 10.1214/22-aihp1286. url: <https://doi.org/10.1214/22-aihp1286>.

[Rag72] Madabusi S Raghunathan. *Discrete subgroups of Lie groups*. Vol. 3. Springer, 1972.

[RS22a] Michael Ren and Nike Sun. *Locality of critical percolation on expanding graph sequences*. 2022. arXiv: 2212.05943 [math.PR]. url: <https://arxiv.org/abs/2212.05943>.

[RS22b] Michael Ren and Nike Sun. *Locality of critical percolation on expanding graph sequences*. 2022. arXiv: 2212.05943 [math.PR]. url: <https://arxiv.org/abs/2212.05943>.

[Rus78] Lucio Russo. “A note on percolation”. In: *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete* 43.1 (1978), pp. 39–48. doi: 10.1007/BF00535274. url: <https://doi.org/10.1007/BF00535274>.

[Rus82] Lucio Russo. “An approximate zero-one law”. In: *Z. Wahrsch. Verw. Gebiete* 61.1 (1982), pp. 129–139. issn: 0044-3719. doi: 10.1007/BF00537230. url: <https://doi.org/10.1007/BF00537230>.

[RY17] A. Raoufi and A. Yadin. “Indicable groups and $p_c < 1$ ”. In: *Electronic Communications in Probability* 22 (2017).

[Sal95] L. Saloff-Coste. “Isoperimetric inequalities and decay of iterated kernels for almost-transitive Markov chains”. In: *Combin. Probab. Comput.* 4.4 (1995), pp. 419–442. issn: 0963-5483. doi: 10.1017/S0963548300001759. url: <https://doi.org/10.1017/S0963548300001759>.

[Sar21a] Sourav Sarkar. “A note on the local weak limit of a sequence of expander graphs”. In: *Electronic Communications in Probability* 26 (2021), pp. 1–6.

[Sar21b] Sourav Sarkar. “A note on the local weak limit of a sequence of expander graphs”. In: *Electron. Commun. Probab.* 26 (2021), Paper No. 32, 6. doi: 10.1214/21-ecp402. url: <https://doi.org/10.1214/21-ecp402>.

[Sch99] Roberto H. Schonmann. “Stability of infinite clusters in supercritical percolation”. In: *Probab. Theory Related Fields* 113.2 (1999), pp. 287–300. issn: 0178-8051,1432-2064. doi: 10.1007/s004400050209. url: <https://doi.org/10.1007/s004400050209>.

[SJ20] Christian R Scullard and Jesper Lykke Jacobsen. “Bond percolation thresholds on Archimedean lattices from critical polynomial roots”. In: *Physical Review Research* 2.1 (2020), p. 012050.

[Sla06] Gordon Slade. *The Lace Expansion and Its Applications: Ecole D’Eté de Probabilités de Saint-Flour XXXIV-2004*. Springer, 2006.

[ST10a] Yehuda Shalom and Terence Tao. “A finitary version of Gromov’s polynomial growth theorem”. In: *Geom. Funct. Anal.* 20 (2010), pp. 1502–1547.

[ST10b] Yehuda Shalom and Terence Tao. “A finitary version of Gromov’s polynomial growth theorem”. In: *Geometric and Functional Analysis* 20 (2010), pp. 1502–1547.

[SW88] Paolo Maurizio Soardi and Wolfgang Woess. *Amenability, unimodularity, and the spectral radius of random walks on infinite graphs*. Dipartimento di Matematica "F. Enriques", 1988.

[SW90] Paolo M. Soardi and Wolfgang Woess. “Amenability, unimodularity, and the spectral radius of random walks on infinite graphs”. In: *Math. Z.* 205.3 (1990), pp. 471–486. issn: 0025-5874. doi: 10.1007/BF02571256.

[Tal94] Michel Talagrand. “On Russo’s approximate zero-one law”. In: *Ann. Probab.* 22.3 (1994), pp. 1576–1587. issn: 0091-1798. url: [http://links.jstor.org/sici?&sici=0091-1798\(199407\)22:3%3C1576:ORAZL%3E2.0.CO;2-S&origin=MSN](http://links.jstor.org/sici?&sici=0091-1798(199407)22:3%3C1576:ORAZL%3E2.0.CO;2-S&origin=MSN).

[Tan19] Pengfei Tang. “Heavy Bernoulli-percolation clusters are indistinguishable”. In: *Ann. Probab.* 47.6 (2019), pp. 4077–4115. issn: 0091-1798. doi: 10.1214/19-aop1354. url: <https://doi.org/10.1214/19-aop1354>.

[Tao14] Terence Tao. *Hilbert's fifth problem and related topics*. Vol. 153. Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, 2014, pp. xiv+338. isbn: 978-1-4704-1564-8. doi: 10.1090/gsm/153. url: <https://doi.org/10.1090/gsm/153>.

[Tao17a] Terence Tao. “Inverse theorems for sets and measures of polynomial growth”. In: *Q. J. Math.* 68.1 (2017), pp. 13–57. issn: 0033-5606. doi: 10.1093/qmath/hav033. url: <https://doi.org/10.1093/qmath/hav033>.

[Tao17b] Terence Tao. “Inverse theorems for sets and measures of polynomial growth”. In: *The Quarterly Journal of Mathematics* 68.1 (2017), pp. 13–57.

[Tas22] Kenshiro Tashiro. “On the speed of convergence to the asymptotic cone for non-singular nilpotent groups”. In: *Geometriae Dedicata* 216.2 (2022), p. 18.

[Tho92] Carsten Thomassen. “Isoperimetric Inequalities and Transient Random Walks on Graphs”. In: *The Annals of Probability* 20.3 (1992), pp. 1592–1600. doi: 10.1214/aop/1176989708. url: <https://doi.org/10.1214/aop/1176989708>.

[Tim07] Á. Timár. “Cutsets in infinite graphs”. In: *Combinatorics, Probability and Computing* 16.1 (2007), pp. 159–166.

[Toi20a] Matthew C. H. Tointon. *Introduction to approximate groups*. Vol. 94. London Mathematical Society Student Texts. Cambridge University Press, Cambridge, 2020, pp. xiii+205. isbn: 978-1-08-45644-9; 978-1-08-47073-5.

[Toi20b] Matthew CH Tointon. “A brief introduction to approximate groups”. In: *European Mathematical Society Magazine* 115 (2020), pp. 12–16.

[Tro84a] V. Trofimov. “Graphs with polynomial growth”. In: *Mat. Sb. (N.S.)* 123(165).3 (1984), pp. 407–421. issn: 0368-8666.

[Tro84b] V. I. Trofimov. “Graphs with polynomial growth”. In: *Mat. Sb. (N.S.)* 123(165).3 (1984), pp. 407–421. issn: 0368-8666.

[Tro85] V. I. Trofimov. “Groups of automorphisms of graphs as topological groups”. In: *Mat. Zametki* 38.3 (1985), pp. 378–385, 476. issn: 0025-567X.

[TT17] Romain Tessera and Matthew Tointon. “Scaling limits of Cayley graphs with polynomially growing balls”. In: *arXiv preprint arXiv:1711.08295* (2017).

[TT18] Romain Tessera and Matthew Tointon. “Properness of nilprogressions and the persistence of polynomial growth of given degree”. In: *Discrete Anal.* (2018), Paper No. 17, 38. doi: 10.19086/da.5056. url: <https://doi.org/10.19086/da.5056>.

[TT20a] Romain Tessera and Matthew Tointon. “Sharp relations between volume growth, isoperimetry and escape probability in vertex-transitive graphs”. In: *Preprint* available at arXiv:2001.01467 (2020).

[TT20b] Romain Tessera and Matthew Tointon. “Sharp relations between volume growth, isoperimetry and resistance in vertex-transitive graphs”. In: *arXiv preprint arXiv:2001.01467* (2020).

[TT21a] Romain Tessera and Matthew C. H. Tointon. “A finitary structure theorem for vertex-transitive graphs of polynomial growth”. In: *Combinatorica* 41.2 (2021), pp. 263–298. issn: 0209-9683. doi: 10.1007/s00493-020-4295-6. url: <https://doi.org/10.1007/s00493-020-4295-6>.

[TT21b] Romain Tessera and Matthew CH Tointon. “A finitary structure theorem for vertex-transitive graphs of polynomial growth”. In: *Combinatorica* 41.2 (2021), pp. 263–298.

[TT23] Romain Tessera and Matthew Tointon. *Small doubling implies small tripling at large scales*. 2023. arXiv: 2310.20500 [math.CO]. url: <https://arxiv.org/abs/2310.20500>.

[TT24] Romain Tessera and Matthew Tointon. *Balls in groups: volume, structure and growth*. 2024. arXiv: 2403.02485 [math.GR]. url: <https://arxiv.org/abs/2403.02485>.

[TY16] Matthew C. H. Tointon and Ariel Yadin. “Horofunctions on graphs of linear growth”. In: *C. R. Math. Acad. Sci. Paris* 354.12 (2016), pp. 1151–1154. issn: 1631-073X. doi: 10.1016/j.crma.2016.10.015. url: <https://doi.org/10.1016/j.crma.2016.10.015>.

[Van22a] Hugo Vanneuville. “Sharpness of Bernoulli percolation via couplings”. In: *arXiv preprint arXiv:2201.08223* (2022).

[Van22b] Hugo Vanneuville. *Sharpness of Bernoulli percolation via couplings*. 2022. doi: 10.48550/ARXIV.2201.08223. url: <https://arxiv.org/abs/2201.08223>.

[Van23] Hugo Vanneuville. *Sharpness of Bernoulli percolation via couplings*. 2023. arXiv: 2201.08223 [math.PR]. url: <https://arxiv.org/abs/2201.08223>.

[Van24] Hugo Vanneuville. *Exponential decay of the volume for Bernoulli percolation: a proof via stochastic comparison*. 2024. arXiv: 2304.12110 [math.PR]. url: <https://arxiv.org/abs/2304.12110>.

[Var85] Nicholas Varopoulos. “Long range estimates for Markov chains”. In: *Bulletin des sciences mathématiques (Paris. 1885)* 109.3 (1985), pp. 225–252.

[Var86] Nicholas Th. Varopoulos. “Théorie du potentiel sur des groupes et des variétés”. In: *C. R. Acad. Sci. Paris Sér. I Math.* 302.6 (1986), pp. 203–205. issn: 0249-6291.

[Wan23] Zhiren Wang. “MATH 597 : Dynamical Systems on Nilmanifolds”. Unpublished lecture notes. Available at: <https://www.personal.psu.edu/zxw14/MATH597/index.html>. 2023.

[Yad23] Ariel Yadin. *Harmonic Functions and Random Walks on Groups*. To appear. Online draft retrieved from <https://www.math.bgu.ac.il/~yadina/> on September 25th 2023. Cambridge University Press, 2023.

[Yam50] Hidehiko Yamabe. “On an arcwise connected subgroup of a Lie group”. In: *Osaka Math. J.* 2 (1950), pp. 13–14. ISSN: 0388-0699.