Supercritical percolation on graphs of polynomial growth

Daniel Contreras, Sébastien Martineau, Vincent Tassion

Supercritical percolation on graphs of polynomial growth

Daniel Contreras* Sébastien Martineau† Vincent Tassion*

June 14, 2022

À la mémoire de Claude Danthony (1961–2021)

Abstract

We consider Bernoulli percolation on transitive graphs of polynomial growth. In the subcritical regime ($p < p_c$), it is well known that the connection probabilities decay exponentially fast. In the present paper, we study the supercritical phase ($p > p_c$) and prove the exponential decay of the truncated connection probabilities (probabilities that two points are connected by an open path, but not to infinity). This sharpness result was established by [CCN87] on $\mathbb{Z}^d$ and uses the difficult slab result of Grimmett and Marstrand. However, the techniques used there are very specific to the hypercubic lattices and do not extend to more general geometries. In this paper, we develop new robust techniques based on the recent progress in the theory of sharp thresholds and the sprinkling method of Benjamini and Tassion. On $\mathbb{Z}^d$, these methods lead to a completely new proof of the slab result of Grimmett and Marstrand.

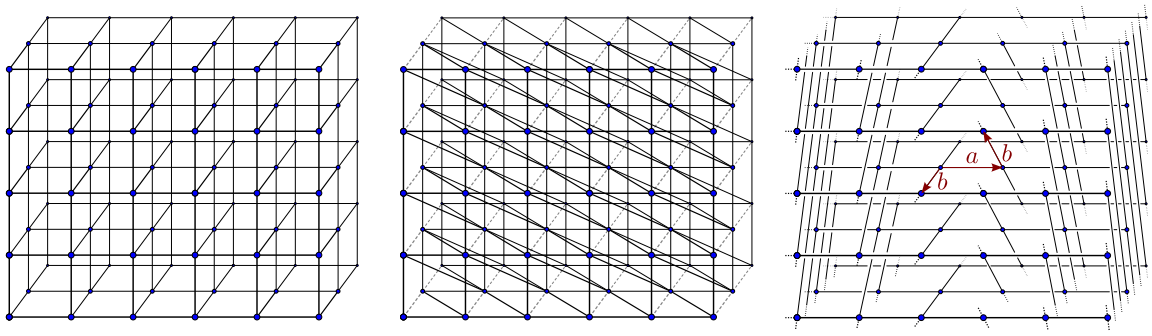

Figure 1: Lattice $\mathbb{Z}^d$ (left), Cayley graph of $\mathbb{Z}^d$ (centre), Heisenberg (right).

*ETH Zürich, Rämistrasse 101, 8092 Zurich Switzerland
†LPSM, Sorbonne Université, 4 place Jussieu, 75005 Paris France
daniel.contreras@math.ethz.ch sebastien.martineau@lpsm.paris vincent.tassion@math.ethz.ch

1 Introduction

1.1 General context

After its introduction in the sixties by Broadbent and Hammersley [BH57], percolation was mainly studied on the hypercubic lattice $\mathbb{Z}^d$, where most of the main features of the model have been rigorously described. In 1996, with their paper "Percolation beyond $\mathbb{Z}^d$, many questions and few answers" [BS96], Benjamini and Schramm initiated the systematic study of percolation on general transitive graphs, leading to a new and fascinating research area. In this generality, new questions emerged, new techniques were used to establish deep relations between the geometric properties of a graph and the behaviour of percolation processes on this graph. Interesting in their own right, these percolation results also shed new light on the theory on $\mathbb{Z}^d$. The present paper is exactly in this spirit: motivated by questions emerging in the general study of percolation on transitive graphs (such as Schramm's Locality Conjecture [BNP11]), we prove a supercritical sharpness result on transitive graphs with polynomial growth. An interested reader will also find below a new proof of the Grimmett–Marstrand Theorem, which is a central result in the study of supercritical percolation on $\mathbb{Z}^d$. The proof is robust, and we expect it to have applications to the study of more general percolation processes on $\mathbb{Z}^d$, such as FK-percolation or level sets of Gaussian processes.

Geometric framework: transitive graphs of polynomial growth Let $G = (V, E)$ be a vertex-transitive graph with a fixed origin $o \in V$ (for every $x, y \in V$, there exists a graph automorphism mapping x to y). Throughout the paper, all the graphs are assumed to be locally-finite and connected and we will always make these hypotheses without further mention. Write B_n for the ball of radius n centred at o. We say that G has polynomial growth if there exists a polynomial P such that $|B_n| \leq P(n)$ for every $n \geq 1$. A celebrated theorem of Gromov and Trofimov [Gro81, Tro85] states that such a graph is always quasi-isometric to a Cayley graph of a finitely-generated nilpotent group (see also Theorem VII.56 of [dLH00], which is there attributed to "Diximier, Wolf, Guivarc'h, Bass, and others"). This deep structure result has a long and ongoing history: see [Bas72, Gui73, Kle10, ST10, TT21]. An important consequence is that there exists an integer d, called the growth exponent of G, such that

$$\exists c > 0 \quad \forall n \geq 1 \qquad cn^d \leq |B_n| \leq \tfrac{1}{c} n^d.$$

Important examples of graphs of polynomial growth include the hypercubic lattice $\mathbb{Z}^d$, more general Cayley graphs of $\mathbb{Z}^d$, and the Heisenberg group[1]: see Figure 1.

Percolation on transitive graphs Let $G = (V, E)$ be a transitive graph. Let P_p be the Bernoulli bond percolation measure P_p on $\{0,1\}^E$, under which $\omega = (\omega(e))_{e \in E}$ is a product of i.i.d. Bernoulli random variables of parameter p (we refer to [Gri99] and

[1] The discrete Heisenberg group $\left\{ \left(\begin{smallmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{smallmatrix}\right) : x, y, z \in \mathbb{Z} \right\}$ is generated by $a = \left(\begin{smallmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{smallmatrix}\right)$ and $b = \left(\begin{smallmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{smallmatrix}\right)$.

[LP16] for general introductions to percolation). We identify ω with the subgraph of G obtained by keeping the edges such that $\omega(e) = 1$ (called open edges) and deleting the edges such that $\omega(e) = 0$ (the closed edges). The connected components of ω are called clusters. Percolation on G undergoes a phase transition at a critical parameter: there is a parameter $p_c \in (0, 1]$ such that for all $p < p_c$, there is P_p-almost surely no infinite cluster and for all $p > p_c$, there is P_p-almost surely at least one infinite cluster. It was recently proved in [DGR+20] that this phase transition is non-trivial if and only if the graph G has superlinear growth: for such graphs, we have $0 < p_c(G) < 1$.

Subcritical sharpness It was proved in [AB87] and [Men86] (see also [DT16]) that the phase transition is sharp, in the following sense: For every $p < p_c$, there exists a constant $c = c(p) > 0$ such that

$$\forall n \geq 1 \quad \mathrm{P}_p[o \leftrightarrow \partial B_n] \leq e^{-cn},$$

where $o \leftrightarrow \partial B_n$ is the event that there exists of an open path from a fixed origin o to distance n around it. The original papers [AB87] and [Men86] prove this result on the hypercubic lattice $\mathbb{Z}^d$, but the two proofs extend to general transitive graphs [AV08]. This result is central in the theory: it leads to the important notion of correlation length, and is the starting point of several finer analyses of the subcritical regime (see e.g. [Gri99] and references therein).

Supercritical sharpness There should be a supercritical counterpart of sharpness. In the supercritical regime $p > p_c$, the natural quantities to consider are the truncated connection probabilities, which encode the connectivity properties of the random graph obtained from ω by removing the infinite cluster(s). It is expected (see e.g [HH21, Conjecture 5.3]) that every $p > p_c$, there is a constant $c = c(p) > 0$ such that

$$\forall n \geq 1 \quad \mathrm{P}_p[o \leftrightarrow \partial B_n, o \not\leftrightarrow \infty] \leq e^{-cn},$$

where $o \not\leftrightarrow \infty$ denotes the event that the origin does not belong to an infinite cluster. Currently, the decay above is known for the hypercubic lattice [CCN87, GM90] and non-amenable graphs [HH21].

In general, the study in the supercritical regime is much more delicate than the subcritical regime. A first reason is geometrical: a key idea to study subcritical clusters is a domination by subcritical branching processes, whose asymptotic features do not rely on the precise geometry of the underlying G. This explains the robustness of the subcritical argument. In contrast, the study of the supercritical clusters requires to understand the infinite cluster(s), whose geometry may be very related to the underlying graph, and so far there is no robust approach to the supercritical regime. A second reason is more technical. The truncated connection events (i.e. of the form $\{A \leftrightarrow B,\ A \not\leftrightarrow \infty\}$ where $A \not\leftrightarrow \infty$ is the event that A is not connected to infinity in ω) are neither increasing nor decreasing events, which reduces considerably the size of the available "toolbox" for the study of such events.

In this paper, we obtain a clear description of the supercritical phase for transitive graphs of polynomial growth, as presented in the next section. In particular, we prove that supercritical sharpness holds for these graphs.

1.2 Main results

Let G be a transitive graph of polynomial growth, with growth exponent $d \geq 2$. Equivalently, G is taken to be a transitive graph of polynomial growth satisfying $p_c(G) < 1$. In this paper, we develop new methods that give a precise description of the supercritical phase of Bernoulli percolation on such graphs ($p > p_c$). More precisely, we build finite-size events that ensure the local existence and uniqueness of large clusters: see Proposition 1.3. This enables us to use powerful renormalisation methods, which extend several perturbative arguments (valid only for p close to 1) to the whole regime $p > p_c(G)$. Our main two results are about the geometry of the finite clusters in the supercritical regime, but our methods would also imply several other results regarding the geometry of the infinite cluster[2]. Our first result states that the radius of a finite cluster has an exponential tail in the supercritical regime.

Theorem 1.1. *Let G be a transitive graph of polynomial growth with $d \geq 2$ and $p > p_c(G)$. There exists a constant $c > 0$ such that*

$$\forall n \geq 1 \quad \mathrm{P}_p[o \leftrightarrow \partial B_n,\, o \not\leftrightarrow \infty] \leq e^{-cn}.$$

The second result is the stretch-exponential decay in volume of the finite clusters.

Theorem 1.2. *Let G be a transitive graph of polynomial growth with $d \geq 2$ and $p > p_c(G)$. Let C_o denote the cluster of o. There exists a constant $c > 0$ such that*

$$\forall n \geq 1 \quad \mathrm{P}_p[n \leq |C_o| < \infty] \leq \exp\left(-cn^{\frac{d-1}{d}}\right).$$

1.3 Comments

Previous work on hypercubic lattices Both Theorems 1.1 and 1.2 were known for the hypercubic lattice $\mathbb{Z}^d$. In dimension $d = 2$, they were proved by Kesten [Kes80]. In dimension $d \geq 3$, they were proved by [CCN87] and [KZ90], by making use of the difficult slab result of Grimmett and Marstrand [GM90]: for any $p > p_c(\mathbb{Z}^d)$, the percolation on $\mathbb{Z}^d$ restricted to a sufficiently thick slab (i.e. a graph of the form $\mathbb{Z}^2 \times \{0, \dots, k\}^{d-2}$ for k large) contains an infinite cluster. These previous methods do not extend to general graphs of polynomial growth for two main reasons. First, the proof of Grimmett and Marstrand relies strongly on the symmetries of $\mathbb{Z}^d$ (in particular reflections and rotations). Such symmetries are not available for general graphs: one can think of Cayley graphs of $\mathbb{Z}^d$ with respect to non-symmetric generating sets (see Figure 1). As for $g \mapsto g^{-1}$, it is generally

[2] A graph of polynomial growth is necessarily amenable: otherwise, the growth would be exponential. The Burton–Keane Theorem [BK89] ensures that, for amenable transitive graphs, for any p, there is at most one infinite cluster almost surely. In particular, for $p > p_c$, there is a unique infinite cluster.

not a graph automorphism for Cayley graphs of non-Abelian groups. A second obstacle is the lack of "slab"-structure for general graphs of polynomial growth: the hypercubic $\mathbb{Z}^d$ can naturally be partitioned into slabs. In contrast, the Cayley graph of the discrete Heisenberg group illustrated on Figure 1 has no natural notion of slab[3]. In our present proof, we replace the slab result by estimates on the two-point function inside "corridor" subgraphs introduced in Section 1.4. Contrary to the lattice case, our proof does not distinguish between $d = 2$ and $d \geq 3$. In this sense, this unifies the two approaches for hypercubic lattices.

Dynamic versus static renormalisation In the case of the hypercubic lattice, the Grimmett–Marstrand Theorem is proved by using a dynamic renormalisation argument (in the sense of [Gri99, Chapter 7]). In the present paper, we construct a local existence-and-uniqueness event, which allows us to directly perform a static renormalisation.

A new quantitative proof of Grimmett–Marstrand The interested reader may extract from the present paper a new proof of the Grimmett–Marstrand Theorem on $\mathbb{Z}^d$. When applying the methods of the present paper to hypercubic lattices, some simplifications occur at several places, using symmetries. For example, the whole Section 6 may be replaced by a reference to the stronger result of Cerf [Cer15], and Section 7.2 may be drastically simplified. The proof is quantitative and would give an estimate of the same order as the one obtained in [DKT].

Lower bounds The bounds in Theorems 1.1 and 1.2 are sharp at the exponential scale: This easier result is classical for the hypercubic lattice $\mathbb{Z}^d$ and the same techniques can be used to prove that, for every $p \in (p_c, 1)$, there exists a constant $c' > 0$ such that

$$\forall n \geq 1 \quad \mathrm{P}_p[o \leftrightarrow \partial B_n,\, o \not\leftrightarrow \infty] \geq e^{-c'n}.$$

and

$$\forall n \geq 1 \quad \mathrm{P}_p[n \leq |C_o| < \infty] \geq \exp\left(-c' n^{\frac{d-1}{d}}\right).$$

Remark 2.5, Lemma 2.6 and Lemma 4.2 guarantee that the geometry of G is — as far as this argument is concerned — as nice as that of $\mathbb{Z}^d$.

Locality In the context of graphs of polynomial growth (which are amenable), supercritical sharpness is related to the existence of some local existence and uniqueness event in a finite box: see Proposition 1.3. This allows us to obtain a finite volume characterisation of $p > p_c(G)$. In [CMT], we use the present work to prove Schramm's Locality Conjecture (stated in [BNP11]) in the particular case of transitive graphs of polynomial growth. This extends [MT17], with different techniques. The case of transitive graphs with (uniformly lower-bounded) exponential growth has been established in [Hut20].

[3] If we try to mimic the definition of slabs in $\mathbb{Z}^d$, we want to take a thickened version of the subgroup generated by a and b. The problem is that, in the case of the Heisenberg group, this subgroup is equal to the whole group.

Related works Several related works regarding the sharpness of the supercritical phase have been developed in the last few years. For the hypercubic lattice $\mathbb{Z}^d$, a quantitative version of the Grimmett–Marstrand Theorem was presented in [DKT]. In the case of non-amenable graphs, [HH21] proved exponential decay for the finite cluster size distribution. For the Ising Model, [DGR20] obtained exponential decay for the truncated two-point function in the supercritical regime. As for Gaussian fields, we can mention the work [DGRS] on the Gaussian free field in $\mathbb{Z}^d$ for $d \geq 3$ and [Sev] for general results on continuous Gaussian fields in $\mathbb{R}^d$ ($d \geq 2$) with correlations decaying reasonably fast.

1.4 Definitions and notation

Throughout the paper, $G = (V, E)$ denotes a fixed transitive graph of polynomial growth, whose growth exponent d satisfies $d \geq 2$. The graph G is taken to be simple: there are no multiple edges, no self-loops, no orientations on the edges. We also fix some origin $o \in V$.

Graph notation For $x \in V$, we write $B_n(x)$ for the ball of radius n centred at x, and we simply write $B_n = B_n(o)$ for the ball centred at the origin. The *boundary* ∂A of a set $A \subset V$ is defined as the set of edges having one endpoint in A and the other in $V \setminus A$. A *path of length ℓ in A* is a sequence $\gamma = (\gamma_0, \dots, \gamma_\ell)$ of vertices of A such that γ_i and γ_{i+1} are neighbours for every i.

Percolation definitions Let ω be a percolation configuration. A path γ is said to be *open* if all its edges $\{\gamma_i, \gamma_{i+1}\}$ are open. For $A \subset V$, we call *clusters in A* the connected components of the graph with vertex set A and edge set the elements of ω with both endpoints in A. For $A, B, C \subset V$, we say that A and B are *connected in C* if there exists an open path in C from A to B. We denote the corresponding event by $A \overset{C}{\leftrightarrow} B$ and its complement by $A \overset{C}{\not\leftrightarrow} B$. In the case $C = V$, we simply write $A \leftrightarrow B$ and $A \not\leftrightarrow B$.

Corridor function Let $m, n \geq 0$ and $p \in [0, 1]$. We define the *corridor function of length m and thickness n* at parameter p by

$$\kappa_p(m, n) = \min_{\gamma \,:\, \mathsf{length}(\gamma) \leq m} \mathrm{P}_p[o(\gamma) \overset{B_n(\gamma)}{\longleftrightarrow} e(\gamma)],$$

where $o(\gamma)$ and $e(\gamma)$ denote the first and last vertices of γ and $B_n(\gamma) = \bigcup_i B_n(\gamma_i)$. The quantity has two different interpretations, depending on whether $m \leq n$ ("short corridor") or $m \geq n$ ("long corridor"): In the first case, the set $B_n(\gamma)$ always looks roughly like a ball, and the quantity is similar to the two point-function restricted to a ball. In contrast, when $m \gg n$, the set $B_n(\gamma)$ takes the shape of a long corridor and the parameter m becomes relevant for its geometry. This quantity will be instrumental in our paper. First, the corridor function has an important renormalisation property, presented in Section 3. Second, it is deeply related to the local uniqueness events, defined below. Finally, it will be the central object in the proof of the main proposition of this paper, in Section 9: we

strengthen some a priori bounds on the short-corridor function into strong estimates in long corridors.

Pivotal and uniqueness events For $1 \le m \le n$ we define the *pivotal event*

$$\mathsf{Piv}(m,n) = \bigcup_{x,y \in B_m} \{x \leftrightarrow \partial B_n,\, y \leftrightarrow \partial B_n,\, x \overset{B_n}{\not\leftrightarrow} y\}.$$

In other words, the event $\mathsf{Piv}(m,n)$ occurs if there are two disjoint clusters (for the configuration restricted to B_n) connecting B_m to ∂B_n. Even if the event is not formally defined in terms of the pivotality of a set for a certain event, we use the notation Piv because the event will typically occur when the ball B_m is pivotal for large connection events. This is also known in the literature as the two-arm event (see [Cer15]). Similarly, for $x \in V$, we define $\mathsf{Piv}_x(m,n)$ as the event $\mathsf{Piv}(m,n)$, but centred at x instead of o. For $1 \le m \le n$, we define the *uniqueness event* by

$$U(m,n) = \{\text{there exists at most one cluster in } B_n \text{ intersecting } B_m \text{ and } \partial B_n\}.$$

Notice that $U(m,n) := \mathsf{Piv}(m,n)^c$. We emphasise that the event $U(m,n)$ does not require the existence of a cluster crossing from B_m to ∂B_n in B_n. When the event $U(m,n)$ occurs, there is either one or no such crossing cluster. The event $U(m,n)$ is particularly useful, as it allows us to "glue" locally macroscopic clusters.

Monotone coupling Let $\mathbf{P} = \mathcal{U}([0,1])^{\otimes E}$ be the product of Lebesgue measures on $[0,1]^E$ and consider the canonical maps $\omega_p : [0,1]^E \to \{0,1\}^E$, $(x_e) \mapsto (\mathbf{1}_{x_e \le p})$. Notice that, under $\mathbf{P}$, the random configuration ω_p has law P_p. The random configuration ω_p naturally gives rise to the notions of p-open edge, p-cluster, p-open path and p-connectivity. For $A, B, C \subset V$ we write $A \overset{C}{\underset{p}{\leftrightarrow}} B$ if A and B are p-connected in C, and $A \overset{C}{\underset{p}{\not\leftrightarrow}} B$ for the complement event. When we are looking at all the p-configurations coupled together, additional interesting events appear. For instance, we can look at how p-clusters are connected at a parameter $q \ge p$, and the following generalisation of the uniqueness event will be useful. For $p \le q$ and $n \ge m \ge 1$, we define the *sprinkled uniqueness event* by

$$U_{p,q}(m,n) = \{\text{All the } p\text{-clusters in } B_n \text{ intersecting } B_m \text{ and } \partial B_n \text{ are } q\text{-connected in } B_n\}.$$

The event above has some useful monotonicity properties: For fixed $n \ge m \ge 1$, the function $f(p,q) = \mathbf{P}[U_{p,q}(m,n)]$ is nonincreasing in p and nondecreasing in q. In contrast, the probability of the uniqueness event $\mathrm{P}_p[U(m,n)]$, which is equal to $f(p,p)$, has no clear monotonicity in p.

1.5 Organisation of the paper

The main ingredient in the proof of Theorems 1.1 and 1.2 is the following proposition, which provides local existence and uniqueness of certain crossing clusters in large boxes. From there, we apply a coarse-graining argument, presented in Section 10. In $\mathbb{Z}^d$, standard

coarse-graining arguments use the scaling property of $\mathbb{Z}^d$: the set $n\mathbb{Z}^d$ is a rescaled version of $\mathbb{Z}^d$. For general transitive graphs of polynomial growth, we circumvent scaling by using an n-independent percolation with sufficiently high marginals.

Proposition 1.3. *Let G be a transitive graph of polynomial growth with $d \geq 2$ and let $p > p_c$. Then, for all n large enough, we have*

$$\mathrm{P}_p[B_{n/10} \leftrightarrow \partial B_n, U(n/5, n/2)] \geq 1 - e^{-\sqrt{n}}.$$

The main goal of the paper is the proof of the proposition above. The proof itself is presented in Section 9. It relies on several independent arguments and intermediate results established in Sections 2–8. Section 2 is a geometric toolbox. In each of the Sections 3–8, we isolate one important ingredient. Each section may use the results from previous ones, but can be read roughly independently. Once the central Proposition 1.3 is established, the two theorems can be proved by using some adaptation of standard renormalisation arguments, which are presented in Section 10. Here is a more detailed roadmap of the forthcoming sections:

Section 2: Geometric lemmas This section provides several definitions and lemmas on the geometry of graphs of polynomial growth, in particular related to cutsets, annuli, spheres, and finding many infinite distant paths exiting B_n. The reader may want to read this section only when other sections require it.

Section 3: Renormalisation of the corridor function In this section, we reduce the proof of Proposition 1.3 to showing that, at infinitely many scales, long corridors can be crossed with probability larger than some constant. This involves a renormalisation property of the corridor function.

Section 4: Probability that two clusters meet at one point We prove that the probability $\mathrm{P}_p[\mathsf{Piv}(1, n)]$ that two clusters of radius n meet at one point decays polynomially fast in n, uniformly in p. We use an adapted version of the beautiful exploration argument of Aizenman, Kesten, and Newman [AKN87] (see also [GGR88, Cer15, Hut20]), which relates such meeting points to the deviation of the sum of i.i.d. Bernoulli(p) random variables. The argument extends to amenable transitive graphs.

Section 5: Sharp threshold results via hypercontractivity We establish a general sharp threshold result for connectivity probabilities of the form $\mathrm{P}_p[A \overset{C}{\leftrightarrow} B]$, $A, B, C \subset V$. To do so, we rely on the general inequalities on influences for Boolean functions of Talagrand [Tal94] and [KKL88]. In the spirit of recent applications of sharp threshold results to percolation [DRT18, DKT], our argument does not involve approximations by symmetric events on the torus. Instead, we only use the bounds on the influences in the "bulk" (at a sufficient distance from the set A and B) provided by Section 4.

Section 6: A priori bound on the uniqueness zone We consider the following problem: for n large, for which value of $s(n)$ can we ensure that at most one cluster in B_n crosses from $B_{s(n)}$ to ∂B_n? Ultimately, Proposition 1.3 shows that we can choose $s(n) \geq n/3$. In order to prove this, we need an a priori lower bound as an intermediate step: in Section 6, we prove that for $p > p_c$, we can choose $s(n)$ larger than any power of $\log n$. To achieve this, we have to overcome difficulties that are not present for the lattice $\mathbb{Z}^d$. In the lattice case, a result of Cerf [Cer15] shows directly that we can choose $s(n) \geq n^c$ for a positive constant $c > 0$ independent of $p \in [0, 1]$. Due to the lack of symmetry, this argument does not extend to general graphs of polynomial growth. Here we combine some arguments of [Cer15] together with a new renormalisation method.

Section 7: Sharp threshold results via Hamming distance This section presents the main new argument, which relies on a general differential inequality involving the Hamming distance on the hypercube. We prove that, in the supercritical regime, if the corridor function is small, then large annuli are crossed with high probability. In Section 9, this enables us to split the proof of Proposition 1.3 into two tractable cases: either the corridor function is large or large annuli are crossed with high probability.

Section 8: Uniqueness via sprinkling We consider the configuration in B_n and assume that the large clusters fill sufficiently well the ball B_n, in the sense that all the balls $B_k(x)$, $x \in B_n$, are p-connected to ∂B_n, for some $k \ll n$. In this case, we prove that for a small sprinkling $\delta > 0$, the sprinkled uniqueness event $U_{p,p+\delta}(n/4, n)$ occurs with high probability: all the p-clusters crossing the annulus $B_n \setminus B_{n/4}$ get connected in B_n at $p + \delta$. We use a non-trivial refinement and generalisation of [BT17].

Section 9: Proof of Proposition 1.3 This section gives the proof of Proposition 1.3. It uses all the results from Sections 3–8.

Section 10: Proof of sharpness: coarse grains without rescaling We use Proposition 1.3 to establish Theorems 1.1 and 1.2. We perform a coarse-graining argument to reduce the study of the supercritical regime to that of a suitable perturbative regime. This is not done by rescaling the graph G but by defining a k-independent percolation process on the original graph G.

1.6 Acknowledgements

We thank Romain Tessera for his availability and for sharing his vision of nilpotent geometry. We are grateful to Itai Benjamini, Hugo Duminil-Copin, Aran Raoufi, Alain-Sol Sznitman, and Matthew Tointon for stimulating discussions.

The first and third authors are supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 851565) and by the NCCR SwissMap. The second author acknowledges the support of the ERC Advanced Grant 740943 GeoBrown.

2 Geometric lemmas

In this section, we collect a certain number of geometric lemmas that will prove useful throughout the paper. This can be skipped in a first reading.

2.1 Cutsets

Percolation is primarily concerned with "connectivity events" and existence of open paths. The complement of a such an event is well understood in terms of closed cutsets, which are dual to open paths. In this section, we define and provide useful properties of cutsets.

If F and Π denote two subsets of $V(G)$, we say that Π is a *cutset* between F and ∞ if $\Pi \cap F = \varnothing$ and every infinite self-avoiding path starting in F has to intersect Π at some point. If, furthermore, no strict subset of Π is a cutset between F and ∞, we say that Π is a *minimal* cutset between F and ∞.

The following well-known lemma provides some geometric control on these *minimal* cutsets.

Lemma 2.1 (Coarse connectedness of minimal cutsets)**.** *Let G be a transitive graph of polynomial growth with $d \geq 2$. Then, there is some constant R such that the following holds.*
Let o be a vertex of G. Every minimal cutset disconnecting o and ∞ is R-connected: For every minimal cutset C, the graph having C as vertex set and with edges connecting two elements of C if their distance in G is at most R is connected.

Proof. By [Tro85], G is quasi-isometric to a Cayley graph H of some finitely generated nilpotent group Γ. As Γ is finitely generated and nilpotent, it is finitely presented (see Propositions 13.75 and 13.84 of [DK17]). Therefore, there is some constant R_H such that every finite cycle of H is a sum modulo 2 of cycles of length at most R_H. As G is quasi-isometric to H, there is some constant R_G such that every finite cycle of G is a sum modulo 2 of cycles of length at most R_G. To see this, fix $\varphi : G \to H$ and $\psi : H \to G$ quasi-isometries that are quasi-inverse of each other. To every edge $\{u, v\}$ of G (resp. H) correspond two vertices $\varphi(u)$ and $\varphi(v)$ (resp. $\psi(u)$ and $\psi(v)$), which are at bounded distance of each other. For each such edge $\{u, v\}$, select once and for all some geodesic path connecting the two corresponding vertices in the other graph. This enables us to associate with each cycle in one graph a cycle in the other. The properties of this process allow us to derive from the existence of R_H the existence of a constant R_G as above. This condition extends to infinite cycles (bi-infinite paths) if one allows infinite (locally finite) sums, because G is one-ended ($d \geq 2$). This permits us to use Theorem 5.1 of [Tim07] (see also [BB99, Tim13]), which yields the conclusion. □

The goal of this paper is to reduce the study of the whole supercritical regime to some perturbative regime, where the Peierls' argument applies. Studying this perturbative regime will require some quantitative control in the usual entropy-vs-energy spirit of the Peierls' argument. The next three lemmas will help us get such a control.

Lemma 2.2. *Let G be a transitive graph of polynomial growth with $d \geq 2$. Let R be so that the conclusion of Lemma 2.1 holds. Let Π be a cutset disconnecting o and ∞. Then, Π intersects the ball of centre o and radius $R|\Pi|$.*

Proof. Assume that Π does not intersect $B_{R|\Pi|}$. Since G is an infinite transitive graph, there is a bi-infinite geodesic path passing through o. The cutset Π has to intersect each of the two geodesic rays this path induces from o. As a result, one has $\mathrm{diam}(\Pi) > 2R|\Pi|$. Since R satisfies the conclusion of Lemma 2.1, we also have $\mathrm{diam}(\Pi) \leq R|\Pi|$, which is contradictory. □

The next lemma is related to Lemma 2.6 and [FGO15, Proposition 5]. The existence of a bi-infinite geodesic path will once again be the key of the proof.

Lemma 2.3. *Let G be an infinite transitive graph. Let F be a finite subset of $V(G)$ containing o. Let Π be a cutset disconnecting F and ∞. Then, we have $\mathrm{diam}(\Pi) \geq \frac{\mathrm{diam}(F)}{2}$.*

Proof. Let $\tilde{F} \subset V(G)$ denote the set of all vertices that can be reached from F by a path avoiding Π. If F is connected, this is simply the connected component of F in $G \setminus \Pi$. We set $m = \mathrm{diam}(\Pi)$ and $n = \mathrm{diam}(F)$. Let us prove that $m \geq \frac{n}{2}$.

First, assume that there is a vertex $v \in \tilde{F}$ such that $B_{n/4}(v) \subset \tilde{F}$. As in the proof of Lemma 2.2, the existence of a bi-infinite geodesic path passing through v yields $m > 2n/4 = n/2$, hence the desired result.

Now, assume that, on the contrary, for every vertex $v \in \tilde{F}$, one has $B_{n/4}(v) \not\subset \tilde{F}$. In particular, one can take u and v two vertices of F at distance n of each other, and then find vertices u' and v' in Π such that $d(u,u') \leq n/4$ and $d(v,v') \leq n/4$. Indeed, by definition of $\tilde{F}$, its external vertex-boundary is a subset of Π — and it is equal to Π if Π is a minimal cutset. By considering u' and v', one gets $m \geq n - 2n/4 = n/2$. □

Lemma 2.4. *Let G be a transitive graph of polynomial growth with $d \geq 1$. There exists $c > 0$ such that the following holds. Let F be a finite connected subset of $V(G)$ containing o. Let Π be a cutset disconnecting F and ∞. Then, we have $|\Pi| \geq c|F|^{\frac{d-1}{d}}$.*

Proof. Let $\tilde{F} \subset V$ be the set of vertices that can be connected to F by a path not intersecting Π. Then $F \subset \tilde{F}$ and every edge of $\partial\tilde{F}$ intersects Π. As $|B_r| \asymp r^d$ and since every transitive amenable graph is unimodular [SW90], Lemma 7.2 of [LMS08] yields the existence of $c > 0$ such that for every choice of (F, Π), we have $|\Pi| \geq c|\tilde{F}|^{\frac{d-1}{d}} \geq c|F|^{\frac{d-1}{d}}$. □

2.2 Spheres and annuli

The most basic example of a cutset is the sphere of radius r. In the hypercubic lattice of dimension $d \geq 2$, spheres are coarsely connected, in the sense of Lemma 2.1 (one can take $R = 2$). Actually, for hypercubic lattices, spheres centred at o are *minimal* cutsets disconnecting o and ∞. Such statements are not true any more for general graphs — not even for one-ended transitive graphs of polynomial growth.

In order to recover coarse connectedness of spheres, we use a notion of exposed sphere [BG18, Pet08, Tim13], that only contains the points accessible from infinity, and where

the finite "pockets" are removed. For $r \geq 0$ and x a vertex of a graph G, let $S_r^\infty(x)$ denote the set of all vertices y such that $d(x,y) = r$ and there exists an infinite self-avoiding path starting at y and intersecting $B_r(x)$ only at y. Notice that replacing the self-avoiding condition by "visiting infinitely many vertices" in this definition yields the same set $S_r^\infty(x)$. If some vertex o is fixed as a root in G, we may write S_r^∞ for $S_r^\infty(o)$.

Remark 2.5. For every $r \geq 1$, the set S_r^∞ is a *minimal* cutset between B_{r-1} and ∞. In particular, Lemma 2.1 yields coarse connectedness of these sets.

Not only does S_r^∞ disconnect B_{r-1} and ∞ but it even disconnects it from ∂B_{2r}, as stated in Lemma 2.6. This lemma corresponds to Proposition 5 in [FGO15]. For the reader's convenience, we have included below its short and nice proof.

Lemma 2.6. *Let G be an infinite transitive graph and o a vertex of G. Let $r \geq 0$ and let γ be a finite path starting in B_r and that intersects ∂B_{2r}. Then, the path γ intersects S_r^∞.*

Proof. Let us fix x and y such that x is a vertex of γ at distance $2r$ from the origin o and y is a neighbour of x satisfying $d(o,y) = 2r+1$. Since G is infinite and transitive, we can fix some bi-infinite geodesic path γ' passing through y at time 0. It is impossible for γ' to intersect B_r in both positive and negative times. Indeed, since γ' is geodesic, this would imply the existence of two points in B_r at distance larger than $2r+1$ away of each other.

Therefore, by following γ' from y in the positive or the negative direction, we get an infinite self-avoiding path γ'' that does not intersect B_r. Since γ starts in B_r and intersects ∂B_{2r}, it visits at least one vertex at distance exactly r from the origin. Take k to be the largest integer such that $d(o,\gamma_k) = r$ and set $v = \gamma_k$. This vertex v necessarily belongs to S_r^∞. Indeed, following the path γ started at time k, then the edge $\{x,y\}$, and then the path γ'' yields a path that starts at v and then leaves B_r forever (and that visits infinitely many vertices). □

While studying percolation on $\mathbb{Z}^d$, it is customary to use not only spheres but also annuli. Annuli will also be useful when G is a one-ended transitive graph of polynomial growth. The next lemma provides some control on their geometry.

Lemma 2.7 (Control on the intrinsic diameter of annuli)**.** *Let G be a transitive graph of polynomial growth with $d \geq 2$. Let $R \geq 1$ be such that the conclusion of Lemma 2.1 holds, and let $n \geq k \geq R$. Then for all $x, y \in S_n^\infty$, there exists a path from x to y within the $2k$-neighbourhood of S_n^∞, of length at most $3k|B_{3n}|/|B_k|$.*

Proof. Let $x, y \in S_n^\infty$. By definition of R and S_n^∞, we can fix some path $(v_0, \dots, v_\ell)$ from x to y that stays in the R-neighbourhood of S_n^∞. Recursively, we define some new finite sequence of vertices as follows. Set $w_0 = v_0 = x$. For every $i > 0$, let m_i be such that $v_{m_i} \in B_{2k}(w_{i-1})$ and none of the vertices $v_{m_i+1}, \dots, v_\ell$ belongs to $B_{2k}(w_{i-1})$. We then set $w_i = v_{m_i+1}$. This process is well-defined until it reaches some w_t that satisfies $d(w_t, y) \leq 2k$, and the process stops there. Notice that the sets $B_k(w_i)$ are disjoint when i ranges over $\{0, \dots, t\}$ and that all these balls are subsets of B_{3n}. Therefore, we have $t+1 \leq \frac{|B_{3n}|}{|B_k|}$. By means of paths of length at most $2k+1$ each, we can connect w_0 to w_1, …, w_{t-1} to w_t, and w_t to y. Concatenating theses paths produces a path of length at most $(2k+1)(t+1)$

that stays within the $2k$-neighbourhood of S_n^∞. As $(2k+1)(t+1) \le 3k|B_{3n}|/|B_k|$, the proof is complete. □

2.3 Crafting many distant paths

In Section 3, we will need many paths that stay away from each other in G. The purpose of the current subsection is to explain how to get such paths. The proper (deterministic) statement is given by Lemma 2.8, which makes use of the following definition.

A path $\gamma : \mathbb{N} \to V$ is said to be a *c-quasi-geodesic ray* if for any m and n in $\mathbb{N}$, we have $d(\gamma_m, \gamma_n) \ge c|m-n|$.

Lemma 2.8. *Let G be a transitive graph of polynomial growth with $d \ge 2$. Then, there exist $c > 0$ such that for every $n \ge 1/c$, for every $a \in [1, cn]$, one can find $\lceil \frac{n}{a} \rceil$ distinct c-quasi-geodesic rays that intersect B_n and that stay at distance at least ca from each other — if x belongs to some ray and y to another one, then $d(x,y) \ge ca$.*

Proof. By [Gro81, Tro85], there is a finitely generated nilpotent group Γ with one end such that G is quasi-isometric to any Cayley graph of Γ. Let S be some finite generating subset of Γ. By [HT, Lemma 3.23], we can pick a surjective homomorphism $\pi : \Gamma \to \mathbb{Z}^2$.

Notice that the conclusion of Lemma 2.8 obviously holds for the square lattice. As this conclusion is stable under quasi-isometry, it holds for the Cayley graph of $\mathbb{Z}^2$ relative to the generating subset $\pi(S)$. By using π, lifting suitable paths for this Cayley graph of $\mathbb{Z}^2$ yields suitable paths for the Cayley graph of Γ relative to S. As this graph is quasi-isometric to G, we have the desired result. □

3 Renormalisation of the corridor function

Recall that we fixed $G = (V, E)$ a transitive graph of polynomial growth with one end, and that d denotes the growth exponent of G. Namely, d is an integer larger than 1 such that $|B_n| \asymp n^d$. We also fix $R \ge 1$ some constant such that the conclusion of Lemma 2.1 holds.

Proposition 3.1. *Let G be a transitive graph of polynomial growth with $d \ge 2$. Let $p \in (0,1)$ be such that*

$$\limsup_n \kappa_p(n \log^{3d} n, n) > 0.$$

Then, for every $q \in (p, 1]$, for every n large enough, we have

$$\mathrm{P}_q[\{B_{n/10} \leftrightarrow \partial B_n\} \cap U(n/5, n/2)] \ge 1 - e^{-\sqrt{n}}. \tag{1}$$

In order to prove Proposition 3.1, we will mimic the orange-peeling argument developed for percolation on $\mathbb{Z}^d$ [Gri99, Lemma 7.89]. Since we will extend this argument in several ways, let us recall briefly how it works to obtain some lower bound on the probability of $U(n, 2n)$ for percolation on $\mathbb{Z}^d$ at a parameter $p > p_c$. First, partition $B_{2n} \setminus B_n$ into n/m annuli of the form $B_{k+m} \setminus B_k$ (m is a fixed constant here, but its role will be important in

our context so we keep it explicit). Using the slab technology [GM90], one can show that in each of these annuli, any two vertices are connected with probability larger than some positive constant δ. Now, consider two clusters $\mathcal{C}(x)$ and $\mathcal{C}(y)$ that cross the annulus $B_{2n} \setminus B_n$ and start at some fixed vertices x and y at the boundary of B_n. Explore them in a recursive manner by revealing the configuration in the annuli $B_{k+m} \setminus B_k$ the one after the other. At each such step, conditionally on the past, there is a probability at least δ that the clusters $\mathcal{C}(x)$ and $\mathcal{C}(y)$ get connected in $B_{k+m} \setminus B_k$. This shows that the probability that $\mathcal{C}(x)$ and $\mathcal{C}(y)$ reach the boundary of B_{2n} without merging is smaller than $e^{-\delta n/m}$. Summing over all the possible x and y, we get

$$\mathbb{P}[U(n,2n)] \geq 1 - |\partial B_n|^2 e^{-\delta n/m}.$$

However, two difficulties arise when extending this argument to more general graphs. First, the annuli given by the induced subgraph of $B_{k+m} \setminus B_m$ may not be connected in our setting. To overcome this, we rather work with the following annuli. For $n \geq m \geq 0$, we define

$$A(n,m) = \bigcup_{x \in S_n^\infty} B_m(x).$$

If m is large enough, Lemma 2.1 implies that $A(n,m)$ is connected and Lemma 2.7 provides us with a good control on the distances in the graph induced by $A(n,m)$.

A second and more serious difficulty is the lack of symmetry, which prevents us from using the slab technology of the Euclidean lattice, and makes the control of the two-point functions in annuli more delicate. Instead of symmetries, our approach uses a renormalisation property of the corridor function, presented in Lemma 3.3. The latter relies on a "sprinkled" version of the orange-peeling argument, presented in Lemma 3.2 below.

Lemma 3.2. *There exists a constant $c > 0$ such that the following holds. Let $p \in (0,1)$, $\delta \in (0, 1-p)$ and $n \geq m \geq R$. If*

$$\kappa_p(3m|B_n|/|B_m|, m) \geq \delta, \tag{2}$$

then, for every $\ell \geq 1$, we have

$$\kappa_{p+\delta}(\ell, n) \geq \delta^2 - \ell |B_n|^2 \exp\left(-\frac{c\delta^3 n}{m}\right). \tag{3}$$

Proof. For simplicity, we will proceed as if all the ratios (n/m, $n/10$, $c_i n/m$...) appearing in the proof were integer-valued. The more general statement can easily be obtained by appropriate "$\lceil\cdot\rceil$-and-$\lfloor\cdot\rfloor$-management". Let $c_0 > 0$ be a constant as in Lemma 2.8.

By possibly reducing the value of the constant c, one may assume without loss of generality that the ratio n/m is large. In particular, we may assume that

$$\frac{3}{c_0} m \leq \frac{c_0}{10} n \quad \text{and} \quad \frac{11}{10c_0} n \leq m \frac{|B_n|}{|B_m|}.$$

In this proof, we work with the coupled configurations $(\omega_p)_{0 \leq p \leq 1}$ under the measure $\mathbf{P}$. First, we show that the classical annulus $B_n \setminus B_{n/10}$ is crossed by at least one p-cluster

with high probability. To this end, apply Lemma 2.8 with $a = \frac{3}{c_0}m$. It guarantees that there exist $c_1 n/m$ disjoint corridors of thickness m connecting $B_{n/10}$ to ∂B_n and of length at most $\frac{11}{10c_0}\, n$. Each one of these corridors is crossed with probability at least δ, by the assumption (2). By independence, we have

$$\mathbf{P}[B_{n/10} \overset{}{\underset{p}{\not\longleftrightarrow}} \partial B_n] \le (1-\delta)^{c_1 n/m} \le \exp\left(-\frac{\delta c_1 n}{m}\right). \tag{4}$$

By automorphism-invariance, any classical annulus $B_n(x) \setminus B_{n/10}(x)$ is crossed by at least one p-cluster with high probability. The main idea to prove (3) is to consider a chain of p-clusters and show that they all get connected at $p+\delta$, with high probability (as illustrated in Fig. 2). To this end, we use an adapted version of the orange-peeling argument, and show that all the p-clusters in $B_{n/2}$ crossing from $B_{n/5}$ to $\partial B_{n/2}$ are locally $(p+\delta)$-connected to each other with high probability. More precisely, we prove that

$$\mathbf{P}[U_{p,p+\delta}(n/5, n/2)] \ge 1 - |B_n|^2 \exp\left(-\frac{\delta^3 n}{140m}\right). \tag{5}$$

To achieve this, for $0 \le i < \frac{n}{140m}$, consider the connected annuli $A_i = A(n_i, 3m)$, where $n_i = \frac{n}{5} + 4m + 7im$. Notice that $\frac{n}{5} + 4m \le n_i \le \frac{n}{4} - 3m$ and the annuli A_i are disjoint subsets of $B_{n/4} \setminus B_{n/5}$. The choice of $3m$ for the thickness of the annulus ensures that for any vertices x, y in $S^\infty_{n_i}$, one can find a corridor from x to y of thickness m and length smaller than $3m|B_{3n_i}|/|B_m|$ that fully lies inside A_i, by Lemma 2.7. Let us define $\mathcal{C}_p(x)$ to be the p-cluster of x in $B_{n/2}$. For $x, y \in B_{n/5}$, we have

$$\begin{aligned} &\mathbf{P}[x \overset{}{\underset{p}{\longleftrightarrow}} \partial B_{n/2},\ y \overset{}{\underset{p}{\longleftrightarrow}} \partial B_{n/2},\ x \overset{B_{n/2}}{\underset{p+\delta}{\not\longleftrightarrow}} y] \\ &= \sum_{C,C'} \mathbf{P}[C \overset{B_{n/2}}{\underset{p+\delta}{\not\longleftrightarrow}} C' \mid \mathcal{C}_p(x) = C,\ \mathcal{C}_p(y) = C'] \cdot \mathbf{P}[\mathcal{C}_p(x) = C,\ \mathcal{C}_p(y) = C'] \end{aligned} \tag{6}$$

where the sum is over the pairs of disjoint clusters (C, C') joining respectively x and y to $\partial B_{n/2}$. Let us fix such C and C'. Notice that under the conditional law

$$\widetilde{\mathrm{P}} = \mathbf{P}[\,\cdot \mid \mathcal{C}_p(x) = C,\ \mathcal{C}_p(y) = C'],$$

the configuration $\omega_{p+\delta}$ is an independent percolation process with marginals satisfying

$$\widetilde{\mathrm{P}}[\omega_{p+\delta}(e) = 1] \begin{cases} = 1 & \text{if } e \in C \cup C', \\ \ge \delta & \text{if } e \in \partial C \cup \partial C', \\ = p+\delta & \text{otherwise.} \end{cases}$$

By Lemma 2.6, both C and C' must intersect all the exposed spheres $S^\infty_{n_i}$ (since $n_i \le n/4$). Furthermore, for every i, Lemma 2.7 ensures that there exists a corridor included in A_i from a point of C to a point of C', of thickness m and of length smaller than $3m|B_{3n_i}|/|B_m|$. Using that $3n_i \le n$ and the bound (2), we obtain

$$\mathbf{P}[\partial C \overset{A_i}{\underset{p}{\longleftrightarrow}} \partial C' \mid \mathcal{C}_p(x) = C,\ \mathcal{C}_p(y) = C'] \ge \kappa_p\,(3m|B_n|/|B_m|, m) \ge \delta.$$

Above, we do not directly get that the two clusters C and C' are connected by a $(p+\delta)$-open path, because we are conditioning on an event where the edges at the boundary of C and C' are p-closed. Nevertheless, using that each such edge is $(p+\delta)$-open with probability larger than δ, independently of the event $\partial C \overset{A_i}{\underset{p}{\longleftrightarrow}} \partial C'$, one obtains the following lower bound:

$$\mathbf{P}[C \overset{A_i}{\underset{p+\delta}{\longleftrightarrow}} C' \mid \mathcal{C}_p(x) = C,\ \mathcal{C}_p(y) = C'] \geq \delta^3,$$

where one additional δ appears in order to leave C and another δ to reach C'. Finally, since all A_i's are disjoint, one has, by independence:

$$\mathbf{P}[C \overset{B_{n/2}}{\underset{p+\delta}{\not\longleftrightarrow}} C' \mid \mathcal{C}_p(x) = C,\ \mathcal{C}_p(y) = C'] \leq (1-\delta^3)^{n/140m}.$$

Plugging this into (6) and then summing over all pairs $x, y \in B_{n/5}$, one gets

$$\mathbf{P}[U_{p,p+\delta}(n/5, n/2)] \geq 1 - |B_n|^2 \exp\left(-\frac{\delta^3 n}{140m}\right).$$

Now, we use the two estimates (4) and (5) in order to prove that (3) holds. Let $\gamma = (v_0, v_1, \ldots, v_k)$ be a path of length $k \leq \ell$ connecting a vertex $x = v_0$ to some vertex $y = v_k$. Let K be the corridor around γ of thickness n. Consider the following three conditions:

(i) $x \leftrightarrow \partial B_n(x)$ and $y \leftrightarrow \partial B_n(y)$,

(ii) for all $j = 1, \ldots, k$, one has $B_{n/10}(v_j) \overset{}{\underset{p}{\longleftrightarrow}} \partial B_n(v_j)$,

(iii) for all $j = 1, \ldots, k$, the uniqueness event $U_{p,p+\delta}(n/5, n/2)$ centred at v_j occurs.

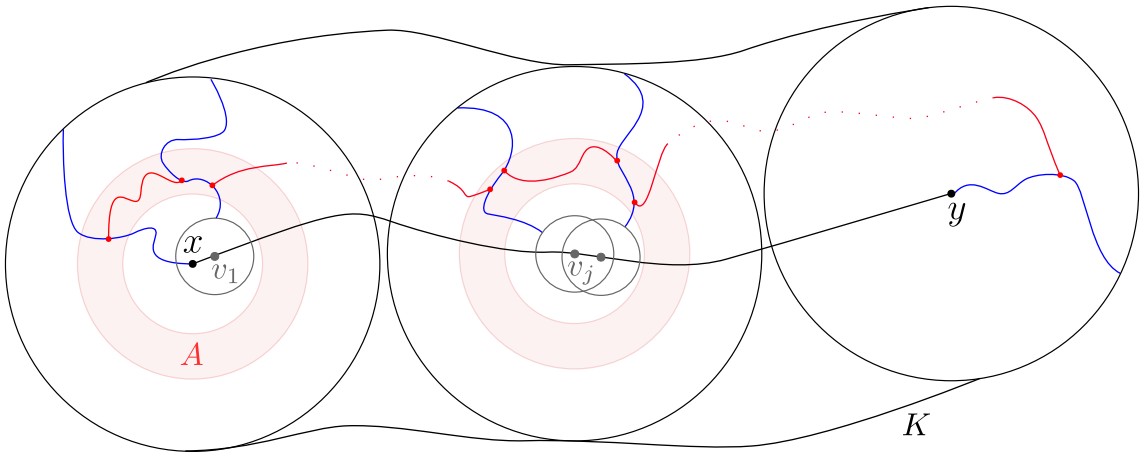

Figure 2: Gluing paths via local uniqueness.

As illustrated in Fig. 2, the simultaneous occurrence of the three events above implies that there exists a $(p+\delta)$-open path in K between x and y. The bounds $\mathbf{P}[z \leftrightarrow \partial B_n(z)] \geq \kappa_p(n, m)$ for $z = x, y$ together with the Harris–FKG Inequality imply that (i) occurs with

probability at least $\kappa_p(n,m)^2$. Therefore, by the union bound and by automorphism-invariance, we obtain

$$\begin{aligned}\mathbf{P}[x \xleftrightarrow[p+\delta]{K} y] &\geq \kappa_p(n,m)^2 - k\mathbf{P}[B_{n/10} \xleftrightarrow[p]{} \partial B_n] - k(1-\mathbf{P}[U_{p,p+\delta}(n/5,n/2)]) \\ &\overset{(4),(5)}{\geq} \delta^2 - \ell \exp\left(-\frac{\delta c_1 n}{m}\right) - \ell |B_n|^2 \exp\left(-\frac{\delta^3 n}{140m}\right),\end{aligned}$$

which concludes the proof of (3). □

Lemma 3.3 (Renormalisation property of the corridor function)**.** *Let $p>0$ and assume that*

$$\limsup_{n\to\infty} \kappa_p(n\log^{3d} n, n) > 0. \tag{7}$$

Then for all $q>p$ and $C>2d$, for all n large enough, one has

$$\kappa_q(n^C, n) \geq n^{-1/C}. \tag{8}$$

Proof. Let us fix $p>0$ such that (7) holds. Let $\eta \in (0, \frac{1-p}{2})$ and $q = p+2\eta$, let $C>2d$. The assumptions imply that

$$\delta := \min\left\{\eta, \tfrac{1}{2}\limsup_{n\to\infty} \kappa_p(n\log^{3d-3} n, n\log^{-2} n)\right\} > 0.$$

Consider a large n satisfying $\kappa_p(n\log^{3d-3} n, n\log^{-2} n) \geq \delta$ and write $m = \frac{n}{\log^2 n}$. As $|B_n| \asymp n^d$, we have $3\frac{|B_n|}{|B_m|} \leq \log^{2d+1} n$ and therefore $\kappa_p(3m\frac{|B_n|}{|B_m|}, m) \geq \delta$. Applying Lemma 3.2 to n, m and $\ell = n^{C^2}$, we obtain

$$\kappa_{p+\eta}(n^{C^2}, n) \geq \kappa_{p+\delta}(n^{C^2}, n) \geq \delta^2 - n^{C^2}|B_n|^2 \exp\big(-c\delta^3 \log^2 n\big).$$

Since one can find arbitrarily large n satisfying the equation above, we get

$$\limsup_{n\to\infty} \kappa_{p+\eta}(n^{C^2}, n) > 0. \tag{9}$$

To get (8), we perform a renormalisation argument. Let c_0 be a constant as in Lemma 3.2. Equation (9) allows us to take $m_0 \geq \max\{\frac{1+\eta}{\eta}, R\}$ such that

$$\kappa_{p+\eta}(m_0^{C^2}, m_0) \geq \frac{1}{m_0^{1/C}}$$

and, for every $m \geq m_0$,

$$m^{C^3}|B_{m^2}|^2 e^{-c_0 m^{1-3/C}} \leq \frac{1}{m^{2/C}} - \frac{1}{m} \qquad \text{and} \qquad 3m\frac{|B_{m^2}|}{|B_m|} \leq m^{C^2}.$$

Let $m_i = m_0^{C^i}$ and $p_{i+1} = p_i + 1/m_{i-1}$, with $p_0 = p+\eta$. We will prove by induction that for every $i \geq 0$,

$$\kappa_{p_i}(m_{i+2}, m_i) \geq \frac{1}{m_{i-1}}. \tag{10}$$

Notice that the relation holds for $i = 0$, by definition of m_0. Let us fix some i such that (10) holds. Since $3m_i \frac{|B_{m_i^2}|}{|B_{m_i}|} \le m_i^{C^2} = m_{i+2}$, one has

$$\kappa_{p_i}\big(3m_i|B_{m_i^2}|/|B_{m_i}|, m_i\big) \ge \frac{1}{m_{i-1}}.$$

Using Lemma 3.2 and observing that $m_{i+1} \ge m_i^2$, we get

$$\begin{aligned}\kappa_{p_{i+1}}(m_{i+3}, m_{i+1}) \ge \kappa_{p_{i+1}}(m_{i+3}, m_i^2) &\ge \frac{1}{m_{i-1}^2} - m_{i+3}|B_{m_i^2}|^2 \exp\left(-\frac{c_0 m_i^2}{m_{i-1}^3 m_i}\right) \\ &\ge \frac{1}{m_i},\end{aligned}$$

where the last inequality holds due to the choice of m_0. This completes the induction.

Now, for n large, consider j such that $m_j \le n < m_{j+1}$. Using that $p_j \le p + 2\eta$ and monotonicity, we get

$$\kappa_{p+2\eta}(n^C, n) \ge \kappa_{p_j}(m_{j+2}, m_j) \ge \frac{1}{m_j^{1/C}} \ge \frac{1}{n^{1/C}}.$$

□

Proof of Proposition 3.1. As mentioned before, the strategy of the proof is an adaptation of the orange-peeling argument used in [Gri99, Lemma 7.89]. Fix $p \in (0,1)$ and $\eta \in (0, (1-p)/2)$. We will prove that the statement (1) holds for $q = p + 2\eta$. Let $C = 8d$. By Lemma 3.3, we can fix $m_0 \ge 1$ such that for all $m \ge m_0$, we have $\kappa_q(m^C, m) \ge m^{-1/C}$. Let n be large enough, so that it is possible to choose an m satisfying

$$m_0 \le m \le n^{1/9} \text{ and } m^C \ge 3m|B_{3n}|/|B_m|. \tag{11}$$

As in the proof of Lemma 3.2, we consider disjoint annuli inside $B_{n/4} \setminus B_{n/5}$: for $0 \le i < \frac{n}{140m}$, consider the connected annuli $A_i = A(n_i, 3m)$, where $n_i = \frac{n}{5} + 4m + 7im$. The gist of the proof is to establish, for all $i < \frac{n}{140m}$ and all $x, y \in A_i$, a good lower bound for $\mathrm{P}_q[x \overset{A_i}{\longleftrightarrow} y]$. Let $i < \frac{1}{140m}$ and $x, y \in A_i$. Denote by x' a vertex of $S_{n_i}^\infty$ such that $d(x, x') = d(x, S_{n_i}^\infty)$. Define y' in the same manner, with respect to y instead of x. By Lemma 2.7, one can find a corridor from x' to y' of thickness m and length smaller than $3m|B_{3n_i}|/|B_m|$ that fully lies within A_i. Using the Harris–FKG Inequality, we get

$$\mathrm{P}_q[x \overset{A_i}{\longleftrightarrow} y] \ge \mathrm{P}_q[x \overset{B_{2m}(x')}{\longleftrightarrow} x'] \cdot \kappa_q(3m|B_{3n}|/|B_m|, m) \cdot \mathrm{P}_q[y \overset{B_{2m}(y')}{\longleftrightarrow} y']. \tag{12}$$

Then, we only need to find a good lower bound for each term. Observe that (11) and Lemma 3.3 directly yield the following lower bound for the central factor of the product above:

$$\kappa_q(3m|B_{3n}|/|B_m|, m) \ge m^{-1/C}. \tag{13}$$

For the other factors of (12), by automorphism-invariance, it is enough to get a lower bound for $\mathrm{P}_p[o \overset{B_{3m}}{\longleftrightarrow} z]$, where $z \in B_{3m}$. In order to get this bound, we will use a chaining argument, which goes as follows. For all k, set $m_k = m_0^{2^k}$. Let $j \ge 0$ be such that

$$m_0 + \cdots + m_{j-1} \le d(o, z) < m_0 + \cdots + m_j.$$

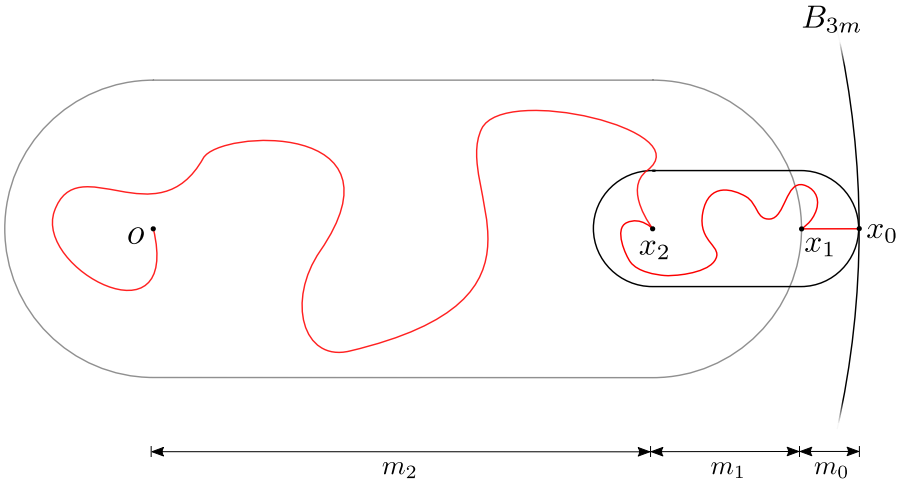

Figure 3: Illustration of the chaining argument: a chain of corridors is used to connect the centre of the ball to a boundary vertex x_0. For the corridors to be subsets of the ball, the thickness is chosen to be smaller when approaching the boundary.

Let γ be a geodesic path between z and o. Fix a sequence of points $x_0 = z, x_1, \dots, x_{j+1} = o$ in γ such that for all $0 \le i < j$, $d(x_i, x_{i+1}) = m_i$ and $d(x_j, o) < m_j$. Notice that for every $1 \le i < j$, the corridor around γ of thickness m_{i-1} between x_i and x_{i+1} is contained in the ball B_{3m}. This is not the case for $i = 0$, but in this case we use the trivial bound $\mathrm{P}_p[x_0 \overset{B_{3m}}{\longleftrightarrow} x_1] \ge p^{m_0}$. Then, using the Harris–FKG Inequality, we get

$$\mathrm{P}_q[o \overset{B_{3m}}{\longleftrightarrow} z] \ge p^{m_0} \prod_{i=1}^{j} \kappa_q(m_i, m_{i-1}) \ge \frac{p^{m_0}}{3m}. \tag{14}$$

Using (12), (13) and (14), we obtain that for all m large enough,

$$\mathrm{P}_q[x \overset{A_i}{\longleftrightarrow} y] \ge \frac{p^{2m_0}}{(3m)^2} \frac{1}{m^{1/C}} \ge \frac{1}{m^3}. \tag{15}$$

Now, we describe the adaptation of the orange-peeling argument presented in [Gri99, Lemma 7.89] to our context. Let x and y denote two distinct vertices of $B_{n/10}$. We explore their clusters step by step, starting from the inside ball $B_{n/10}$. The i^{th} step explores these clusters until they touch the annulus A_i — or until we are sure they will never do so. Denote by x_i and y_i the respective points where they touch A_i for the first time. Let x_i' and y_i' be the closest points in $S_{n_i}^{\infty}$ to x_i and y_i, respectively (see Figure 4). We know by (15) that $\mathrm{P}_q[x_i \overset{A_i}{\longleftrightarrow} y_i] \ge \frac{1}{m^3}$. By conditional independence, we obtain that for every n large enough,

$$\mathrm{P}_q[U(n/5, n/2)] \ge 1 - |B_{n/10}|^2 \left(1 - \frac{1}{m^3}\right)^{n/140m} \ge 1 - \frac{e^{-\sqrt{n}}}{2},$$

where we use that $m \le n^{1/9}$. It remains to prove that for every n large enough, we have

$$\mathrm{P}_q[B_{n/10} \leftrightarrow \partial B_n] \ge 1 - \frac{e^{-\sqrt{n}}}{2},$$

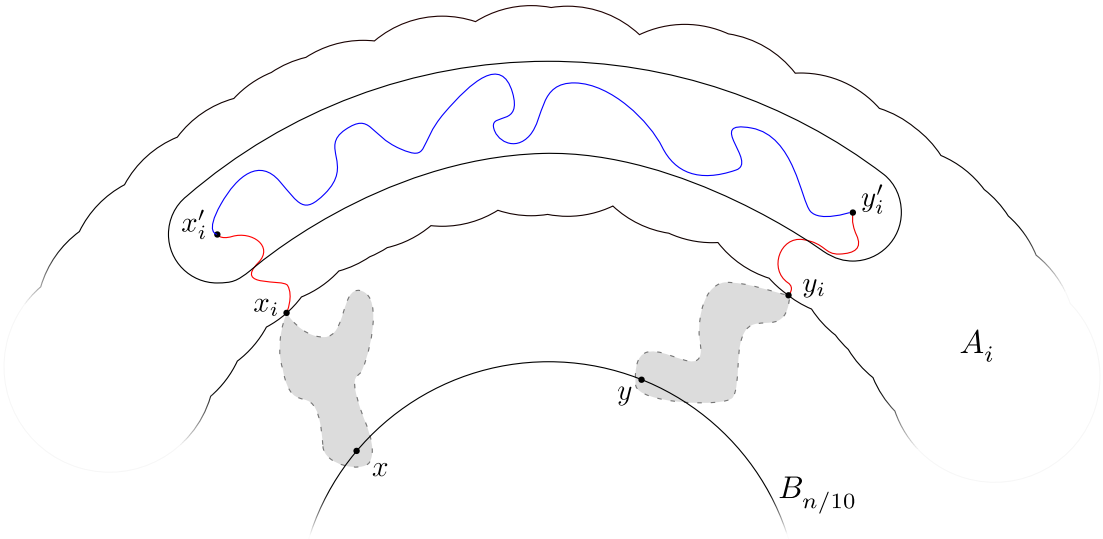

Figure 4: Orange-peeling argument: when the explorations of two clusters from $B_{n/10}$ reach the annulus A_i, they get connected inside A_i with probability at least $1/m^3$.

which can be proved as Equation (4) by using that there exist at least $cn^{1-1/C}$ disjoint corridors of thickness $n^{1/C}$ from $B_{n/10}$ to ∂B_n, and that each of these corridors is independently crossed with probability at least $n^{-1/C}$. □

4 Probability that two clusters meet at one point

This section is devoted to the proof of the following proposition, which extends the quantitative uniqueness argument of [AKN87, GGR88] to graphs of polynomial growth. The proof presented below follows the more recent presentations of [Cer15] and [Hut20].

Proposition 4.1. *Let G be a transitive graph of polynomial growth. Let $\varepsilon > 0$ and $\eta > 0$. There is a constant $c = c(G, \varepsilon, \eta)$ such that for all $p \in [\eta, 1]$ and all $n \geq 1$, we have*

$$\mathrm{P}_p[\mathsf{Piv}(1, n)] \leq cn^{-1/2+\varepsilon}.$$

The interest of Proposition 4.1 is twofold. On the one hand, it will be useful to bound the probability of pivotal edges when studying the derivative of crossing probabilities: it will be important to establish the general sharp threshold result of Proposition 5.1. On the other hand, the work [Cer15] shows that for $G = \mathbb{Z}^d$ the bound above can be strengthened into bounds on the probability $\mathrm{P}_p[\mathsf{Piv}(m, n)]$ for $m \geq 1$. For more general graphs, we will obtain a similar result in Proposition 6.1, but we need to combine the approach of [Cer15] with a new renormalisation argument in order to overcome the lack of symmetry of G.

In order to prove Proposition 4.1, we will make use of the following geometric observation.

Lemma 4.2. *Let G be a transitive graph of polynomial growth. There is a constant c such that the following holds: for every $n \geq 1$, there is an integer m such that $n \leq m < 2n$ and $|\partial B_m|/|B_m| \leq c/n$.*

Proof. Using that the sum $\sum_{m=n}^{2n-1} |\partial B_m|$ is smaller than the total number of edges in B_{2n}, we have

$$\min_{n \leq m < 2n} \frac{|\partial B_m|}{|B_m|} \leq \frac{1}{n} \sum_{m=n}^{2n-1} \frac{|\partial B_m|}{|B_m|} \leq \frac{1}{n} \cdot \frac{|E(B_{2n})|}{|B_n|}.$$

Since $|B_n| \asymp n^d$ and G has bounded degree, $\frac{|E(B_{2n})|}{|B_n|}$ is bounded from above by a constant, which concludes the proof. □

Proof of Proposition 4.1. Let $c_1 > 0$ be a constant such that the conclusion of Lemma 4.2 holds, and let n be a positive integer. Then, there is an integer m such that $n \leq m < 2n$ and

$$|\partial B_m|/|B_m| \leq c_1/n.$$

Let us fix such an m.

Following [Cer15], let us consider the random set of edges $\mathcal{H}$ defined as follows: we look at the configuration restricted to B_m and we say that e belongs to $\mathcal{H}$ if it is closed and its endpoints belong to disjoint clusters in B_m that both touch ∂B_m. Since $4n \geq 2m$, one can use the finite-energy property and automorphism-invariance to prove that for every $x \in B_{m-1}$,

$$p \cdot \mathrm{P}_p[\mathsf{Piv}(1, 4n)] \leq 2 \sum_{e \ni x} \mathrm{P}_p[e \in \mathcal{H}].$$

Summing over all $x \in B_{m-1}$, we get

$$|B_{m-1}| \cdot p \cdot \mathrm{P}_p[\mathsf{Piv}(1, 4n)] \leq 4\mathrm{E}_p[|\mathcal{H}|]. \tag{16}$$

Let us define $\mathfrak{C}$ to be the family of all the open clusters in B_m that intersect ∂B_m. Define $\overline{\mathcal{C}}$ to be the union of all these clusters. Given a subset S of B_m, we write $\mathsf{open}(S)$ (resp. $\mathsf{closed}(S)$) for the set of open (resp. closed) edges of $E(B_m)$ adjacent to at least one vertex of S. In order to bound the size of $\mathcal{H}$, we rely on the following identities,

$$\sum_{\mathcal{C} \in \mathfrak{C}} |\mathsf{open}(\mathcal{C})| = |\mathsf{open}(\overline{\mathcal{C}})| \tag{17}$$

$$\sum_{\mathcal{C} \in \mathfrak{C}} |\mathsf{closed}(\mathcal{C})| = |\mathsf{closed}(\overline{\mathcal{C}})| + |\mathcal{H}|, \tag{18}$$

which are proved by a counting argument. The first sum counts all the open edges in B_m connected to the boundary of B_m, which also corresponds to the open edges adjacent to $\overline{\mathcal{C}}$. The second sum counts all the closed edges which are connected to the boundary of B_m, except that the edges of $\mathcal{H}$ are counted twice.

Furthermore, using that the event $\{e \leftrightarrow \partial B_m\}$ is independent of the status of the

edge e, we find

$$\begin{aligned} p \cdot \mathrm{E}_p[|\mathsf{closed}(\overline{\mathcal{C}})|] &= \sum_{e \subset B_m} p \cdot \mathrm{P}[e \text{ is closed}, e \leftrightarrow \partial B_m] \\ &= \sum_{e \subset B_m} p(1-p) \cdot \mathrm{P}[e \leftrightarrow \partial B_m] \\ &= \sum_{e \subset B_m} (1-p) \cdot \mathrm{P}[e \text{ is open}, e \leftrightarrow \partial B_m] \\ &= (1-p) \cdot \mathrm{E}_p[|\mathsf{open}(\overline{\mathcal{C}})|]. \end{aligned}$$

Taking the expectation in (17) and (18), and using the computation above, we get

$$\mathrm{E}_p[|\mathcal{H}|] = \frac{1}{p} \mathrm{E}_p[\sum_{\mathcal{C} \in \mathfrak{C}} h(\mathcal{C})], \tag{19}$$

where $h(\mathcal{C}) = p|\mathsf{closed}(\mathcal{C})| - (1-p)|\mathsf{open}(\mathcal{C})|$. In the right-hand side above, the fact that the sum is over a random set makes it delicate to study. To overcome this difficulty, we will "root" each cluster at some of its vertices, and sum over the possible roots. For every $x \in B_m$, let $\mathcal{C}_x$ be the cluster of x in B_m. Using that for every cluster $\mathcal{C}$, there are exactly $|\mathcal{C}|$ vertices x such that $\mathcal{C}_x = \mathcal{C}$, we get

$$\mathrm{E}_p\left[\sum_{\mathcal{C} \in \mathfrak{C}} f(\mathcal{C})\right] = \sum_{x \in B_m} \mathrm{E}_p\left[\frac{f(\mathcal{C}_x)}{|\mathcal{C}_x|} \mathbf{1}_{\{x \leftrightarrow \partial B_m\}}\right], \tag{20}$$

for every function $f : \mathcal{P}(B_m) \to \mathbb{R}$. Applying this to $f = h$, and using the Cauchy–Schwarz inequality, we get

$$\mathrm{E}_p\left[\sum_{\mathcal{C} \in \mathfrak{C}} h(\mathcal{C})\right] \le \left(\sum_{x \in B_m} \mathrm{E}_p\left[\frac{h(\mathcal{C}_x)^2}{|\mathcal{C}_x|^{1+\varepsilon}}\right]\right)^{1/2} \left(\sum_{x \in B_m} \mathrm{E}_p\left[\frac{\mathbf{1}_{\{x \leftrightarrow \partial B_m\}}}{|\mathcal{C}_x|^{1-\varepsilon}}\right]\right)^{1/2}. \tag{21}$$

Applying (20) to $f(S) = |S|^\varepsilon$ and using Hölder's inequality with parameters $\frac{1}{\varepsilon}$ and $\frac{1}{1-\varepsilon}$ (we may and will assume that $\varepsilon < 1$), we obtain

$$\sum_{x \in B_m} \mathrm{E}_p\left[\frac{\mathbf{1}_{\{x \leftrightarrow \partial B_m\}}}{|\mathcal{C}_x|^{1-\varepsilon}}\right] = \mathrm{E}_p\Big[\sum_{\mathcal{C} \in \mathfrak{C}} |\mathcal{C}|^\varepsilon\Big] \le \mathrm{E}_p\Big[\Big(\sum_{\mathcal{C} \in \mathfrak{C}} |\mathcal{C}|\Big)^\varepsilon |\mathfrak{C}|^{1-\varepsilon}\Big] \le |B_m|^\varepsilon \, |\partial B_m|^{1-\epsilon}, \tag{22}$$

where in the third step we used that every cluster in $\mathfrak{C}$ has to touch ∂B_m. By (16), (19), (21) and (22), we find

$$\mathrm{P}_p[\mathsf{Piv}(1, 4n)] \le c_2 \left(\tfrac{1}{|B_m|} \sum_{x \in B_m} \mathrm{E}_p\left[\frac{h(\mathcal{C}_x)^2}{|\mathcal{C}_x|^{1+\varepsilon}}\right]\right)^{1/2} \left(\frac{|\partial B_m|}{|B_m|}\right)^{(1-\varepsilon)/2},$$

where c_2 is a finite constant depending only on η. To conclude the proof, it suffices to show that for every fixed $x \in B_m$, the quantity

$$\mathrm{E}_p\left[\frac{h(\mathcal{C}_x)^2}{|\mathcal{C}_x|^{1+\varepsilon}}\right]$$

is smaller than some constant $c_3 < \infty$. This will be achieved by interpreting $h(\mathcal{C}_x)$ as the value of a martingale at some stopping time, when we perform a certain exploration of the cluster $\mathcal{C}_x$.

First, fix an arbitrary deterministic ordering of the edges of B_m. Set $(O_0, C_0) = (\varnothing, \varnothing)$. Then, let e_1 be the smallest edge adjacent to x and set

$$(O_1, C_1) = \begin{cases} (\{e_1\}, \varnothing) & \text{if } e_1 \text{ is open} \\ (\varnothing, \{e_1\}) & \text{if } e_1 \text{ is closed.} \end{cases}$$

By induction, for $t \geq 2$, define e_t to be the smallest edge that is adjacent to an edge of O_{t-1} (if it exists) and

$$(O_t, C_t) = \begin{cases} (O_{t-1} \cup \{e_t\}, C_{t-1}) & \text{if } e_t \text{ is open} \\ (O_{t-1}, C_{t-1} \cup \{e_t\}) & \text{if } e_t \text{ is closed.} \end{cases}$$

If there is no such edge, we define $(O_t, C_t) = (O_{t-1}, C_{t-1})$. Observe that at this moment, (O_t, C_t) corresponds to $(\mathsf{open}(\mathcal{C}_x), \mathsf{closed}(\mathcal{C}_x))$. Let $T = |\mathsf{open}(\mathcal{C}_x)| + |\mathsf{closed}(\mathcal{C}_x)|$ be the time at which the exploration stabilises. We define

$$X_t = \sum_{k=1}^{t \wedge T} p\mathbf{1}_{\{\omega(e_k)=0\}} - (1-p)\mathbf{1}_{\{\omega(e_k)=1\}}.$$

Notice that $X_T = h(\mathcal{C}_x)$ and $T \leq \deg(o)|\mathcal{C}_x|$, since any vertex is adjacent to at most $\deg(o)$ edges in B_m. Thus,

$$\mathrm{E}_p\left[\frac{h(\mathcal{C}_x)^2}{|\mathcal{C}_x|^{1+\varepsilon}}\right] \leq \deg(o)^{1+\varepsilon}\mathrm{E}_p\left[\frac{X_T^2}{T^{1+\varepsilon}}\right]. \tag{23}$$

In order to upperbound $\mathrm{E}_p[X_T^2/T^{1+\varepsilon}]$, we use that (X_t) is a martingale with respect to the filtration generated by $\{\omega(e_t)\}_t$. By first applying Doob's maximal inequality and then using orthogonality of the increments, we obtain

$$\mathrm{E}_p[\max_{k \leq t} X_k^2] \leq 4\mathrm{E}_p[X_t^2] = 4\sum_{k=1}^{t} \mathrm{E}_p[(X_k - X_{k-1})^2] \leq t. \tag{24}$$

We conclude the proof by decomposing the expectation in the right-hand side of (23) as

$$\mathrm{E}_p\left[\frac{X_T^2}{T^{1+\varepsilon}}\right] \leq \sum_{i \geq 0} \frac{1}{2^{i(1+\varepsilon)}}\mathrm{E}_p\left[\max_{2^i \leq t \leq 2^{i+1}} X_t^2 \mathbf{1}_{\{2^i \leq T \leq 2^{i+1}\}}\right] \overset{(24)}{\leq} \sum_{i \geq 0} \frac{2}{2^{i\varepsilon}} \leq \frac{2}{1-2^{-\varepsilon}}.$$

□

5 Sharp threshold results via hypercontractivity

In this section, we establish a general sharp threshold result for connection events. Its proof involves the polynomial upper bound on the probability of $\mathsf{Piv}(1, n)$ from Section 4, together with an abstract result from the theory of Boolean functions. Given a set $A \subset V$, we write A_s for its s-thickening. Formally, A_s is defined to be the set of vertices at distance at most s from A.

Proposition 5.1. *Let G be a transitive graph of polynomial growth with $d \geq 1$. Let $\eta > 0$. There is a constant $c = c(\eta, G) > 0$ such that for every $p \in (\eta, 1-\eta)$ and $\delta \in (0, 1-\eta-p)$, the following holds. For every $C \subset V(G)$, $A, B \subset C$ and $s \geq 1$, we have*

$$\mathrm{P}_p[A \overset{C}{\leftrightarrow} B] > \frac{1}{(cs)^{\delta/20}} \quad \Longrightarrow \quad \mathrm{P}_{p+\delta}[A_s \overset{C_s}{\leftrightarrow} B_s] > 1 - \frac{1}{(cs)^{\delta/20}}.$$

Remark 5.2. The first application of an abstract sharp threshold result to percolation theory dates from the 2000's, in the paper [BR06] computing the critical value for planar Voronoi percolation. Since then, it has become a standard tool to study or prove sharpness of the phase transition for percolation processes. Standard applications involve a certain increasing event on a torus which is translation-invariant, and for which the result of [KKL88] ensures a sharp threshold phenomenon.

Proposition 5.1 above is inspired by the works [DRT18] and [DKT] where abstract sharp threshold results are used directly on $\mathbb{Z}^d$, without relying on translation-invariant events on a torus. Let us briefly expose the main difficulty we have to overcome in the present paper. An important idea in the two works mentioned above is that the event $E = \{A \overset{C}{\leftrightarrow} B\}$ is "geometrically similar" to its translate $\tau \cdot E$ by a small vector, in the sense that A, B and C are close to their translates $\tau \cdot A$, $\tau \cdot B$ and $\tau \cdot C$ respectively. This is true for $\mathbb{Z}^d$ (and more generally for Cayley graphs of *abelian* groups), but not for general transitive graphs.

Proof. In the proof below, $c_1, c_2, \ldots$ denote positive constants that may depend on η and G but are independent of everything else. Without loss of generality, we may (and will) assume that A, B and C are finite. By [BKK+92] and [Tal94, Corollary 1.2], there exists $c_1 = c_1(\eta) > 0$ such that the following inequality holds for any increasing event $\mathcal{A}$ depending on finitely many edges:

$$\sum_{e \in E} \mathrm{P}_p[e \text{ is pivotal for } \mathcal{A}] \geq c_1 \mathrm{P}_p[\mathcal{A}](1 - \mathrm{P}_p[\mathcal{A}]) f_p(\mathcal{A}), \tag{25}$$

where

$$f_p(\mathcal{A}) = \log\left(\frac{1}{\max_{e \in E} \mathrm{P}_p[e \text{ is pivotal for } \mathcal{A}]}\right).$$

Using Russo's formula together with (25), we obtain

$$\frac{d}{dp} \log\left(\frac{\mathrm{P}_p[\mathcal{A}]}{1 - \mathrm{P}_p[\mathcal{A}]}\right) \geq c_1 f_p(\mathcal{A}). \tag{26}$$

To prove the desired inequality, we will define a sequence of nested and increasing events $\mathcal{E}_i$ such that $\{A \overset{C}{\leftrightarrow} B\} \subset \mathcal{E}_i \subset \{A_s \overset{C_s}{\leftrightarrow} B_s\}$. The first step will consist in finding a uniform lower bound for $f_p(\mathcal{E}_i)$. We will do so by proving that the probability of any edge to be pivotal is uniformly small. Set $r = 2\lfloor \sqrt{s}/2 \rfloor$. Without loss of generality, we may assume that $s \geq 4$ and look for a suitable c in $(0, 1/4)$, so that $r > 0$. For $i \in \{0, \ldots, r-1\}$, consider the set $D_i = (A_{(i+1)r} \setminus A_{ir}) \cup (B_{(i+1)r} \setminus B_{ir})$, the event $\mathcal{E}_i = \{A_{ir} \overset{C_{(i+1)r}}{\longleftrightarrow} B_{ir}\}$, and

$$m_i = \max_{e \subset D_i} \mathrm{P}_p[e \text{ is a closed pivotal for } \mathcal{E}_i]. \tag{27}$$

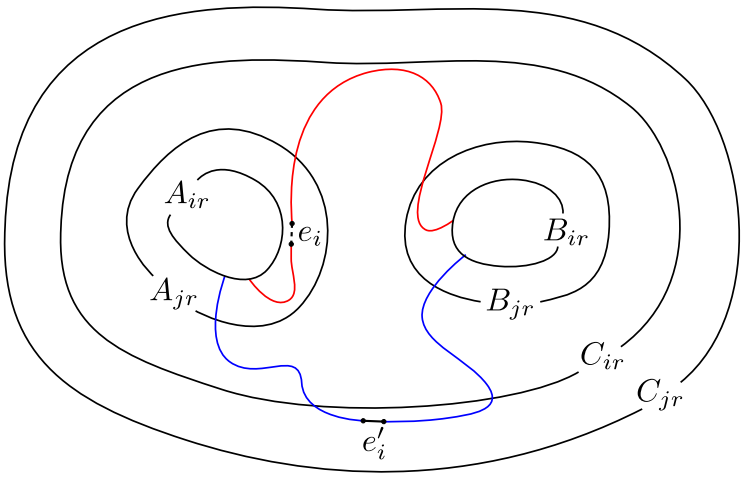

Figure 5: The red path illustrates the event $\mathcal{F}_i$ and the blue one $\mathcal{F}'_i$.

For every i, let e_i be a fixed edge that maximises the probability in (27) and define $\mathcal{F}_i$ to be the event $\{e_i \text{ is a closed pivotal for } \mathcal{E}_i\}$. We now fix $i, j \in \{0, \dots, r-1\}$ with $i < j$. First, let us assume that the event $\mathcal{F}_i$ occurs. Then the edge e_i is closed and its two endpoints are connected in $C_{(i+1)r}$ by disjoint open paths to A_{ir} and B_{ir}, respectively. This implies that there is an open path connecting A_{jr} to B_{jr} in $C_{(i+1)r}$. Thus, $\mathcal{F}_i \subset \mathcal{E}_j$. On the other hand, if $\mathcal{F}_j$ occurs, the fact that e_j is closed and pivotal implies that there is no open path connecting A_{jr} to B_{jr} in C_{jr}. Thus $\mathcal{F}_j \subset \mathcal{E}_j^c$. This leads us to a key observation, which is that $\mathcal{F}_i \cap \mathcal{F}_j = \varnothing$ for $i < j$. Therefore, we have

$$\sum_{i=0}^{r-1} m_i = \sum_{i=0}^{r-1} \mathrm{P}_p[\mathcal{F}_i] \le 1. \tag{28}$$

Similarly, for $i \in \{0, \dots, r-1\}$, we can define the set $D'_i = C_{(i+1)r} \setminus C_{ir}$ and

$$m'_i = \max_{e' \subset D'_i} \mathrm{P}_p[e' \text{ is an open pivotal for } \mathcal{E}_i]. \tag{29}$$

For every i, let e'_i be a fixed edge that maximises the probability in (29) and define $\mathcal{F}'_i$ to be the event $\{e'_i \text{ is an open pivotal for } \mathcal{E}_i\}$. Using the same reasoning as before, we can prove that for $i < j$, we have $\mathcal{F}'_i \subset \mathcal{E}_i$ and $\mathcal{F}'_j \subset \mathcal{E}_i^c$. Thus, $\mathcal{F}'_i \cap \mathcal{F}'_j = \varnothing$. This, together with (28), implies that $\sum_{i=0}^{r-1} m_i + m'_i \le 2$. From this, it follows that we can fix a set I with $|I| \ge r/2$ and such that, for all $i \in I$, we have $\max\{m_i, m'_i\} \le \frac{4}{r}$. Since the status of an edge is independent of the event that it is pivotal, we obtain

$$\max_{e \subset D_i \cup D'_i} \mathrm{P}_p[e \text{ is pivotal for } \mathcal{E}_i] \le \max\left\{\frac{1}{p}, \frac{1}{1-p}\right\} \cdot \frac{4}{r} \le \frac{4}{\eta r}. \tag{30}$$

Pick any $i \in I$. Let e be an edge in C_s such that $e \notin D_i \cup D'_i$. If $e \subset A_{ir} \cup B_{ir} \cup (C_s \setminus C_{(i+1)r})$ then $\mathrm{P}_p[e \text{ is pivotal for } \mathcal{E}_i]$ is equal to 0. Otherwise, the r-neighbourhood of e lies in $C_{(i+1)r}$. By Proposition 4.1, we get

$$\mathrm{P}_p[e \text{ is closed pivotal for } \mathcal{E}_i] \le \mathrm{P}_p[\mathsf{Piv}(1, r)] \le \frac{1}{c_2 s^{1/5}}.$$

Combined with (30), this means that for all $i \in I$,

$$\max_{e \subset C_s} \mathrm{P}_p[e \text{ is a pivotal for } \mathcal{E}_i] \leq \frac{1}{c_3 s^{1/5}}.$$

Using this, we obtain the following inequality:

$$\sum_{i=1}^{r} f_p(\mathcal{E}_i) \geq \sum_{i \in I} f_p(\mathcal{E}_i) \geq \frac{r}{2} \log(c_3 s^{1/5}). \tag{31}$$

Now, we can use (26) for all the events $\mathcal{E}_i$, where $i \in \{1, \dots, r\}$. Integrating between p and $p+\delta$ and using (31), we get

$$\sum_{i=1}^{r} \log\left(\frac{\mathrm{P}_{p+\delta}[\mathcal{E}_i]}{1-\mathrm{P}_{p+\delta}[\mathcal{E}_i]} \cdot \frac{1-\mathrm{P}_p[\mathcal{E}_i]}{\mathrm{P}_p[\mathcal{E}_i]} \right) \geq \frac{\delta r}{2} \log(c_3 s^{1/5}).$$

Therefore, there exists an index i such that

$$\log\left(\frac{\mathrm{P}_{p+\delta}[\mathcal{E}_i]}{1-\mathrm{P}_{p+\delta}[\mathcal{E}_i]} \cdot \frac{1-\mathrm{P}_p[\mathcal{E}_i]}{\mathrm{P}_p[\mathcal{E}_i]} \right) \geq \frac{\delta}{2} \log(c_3 s^{1/5}),$$

which implies

$$\frac{1}{(1-\mathrm{P}_{p+\delta}[\mathcal{E}_i])\mathrm{P}_p[\mathcal{E}_i]} \geq \frac{\mathrm{P}_{p+\delta}[\mathcal{E}_i]}{1-\mathrm{P}_{p+\delta}[\mathcal{E}_i]} \cdot \frac{1-\mathrm{P}_p[\mathcal{E}_i]}{\mathrm{P}_p[\mathcal{E}_i]} \geq (c_4 s)^{\delta/10}.$$

Since $\{A \overset{C}{\leftrightarrow} B\} \subset \mathcal{E}_i \subset \{A_s \overset{C_s}{\leftrightarrow} B_s\}$, we have

$$\mathrm{P}_{p+\delta}[A_s \overset{C_s}{\leftrightarrow} B_s] \geq \mathrm{P}_{p+\delta}[\mathcal{E}_i] \geq 1 - \frac{1}{\mathrm{P}_p[\mathcal{E}_i](c_4 s)^{\delta/10}} \geq 1 - \frac{1}{\mathrm{P}_p[A \overset{C}{\leftrightarrow} B](c_4 s)^{\delta/10}}.$$

Proposition 5.1 follows. □

Corollary 5.3. *Let G be a transitive graph of polynomial growth with $d \geq 2$. Let $p > p_c$. There exist $\delta > 0$ and $c > 0$ such that, for all $s \geq 1$, we have*

$$\mathrm{P}_p[B_s \leftrightarrow \infty] \geq 1 - \frac{1}{(cs)^{\delta}}.$$

Proof. Let $\delta > 0$ be such that $p > p_c + 20\delta$. We know that $\mathrm{P}_{p-20\delta}[o \leftrightarrow \infty] > 0$. By Proposition 5.1 applied to $A = \{o\}$, $B = \partial B_n$ and $C = V(G)$, we have $\mathrm{P}_p[B_s \leftrightarrow \partial B_{n-s}] > 1 - (cs)^{-\delta}$, where c is a positive constant independent of s and n. The proof follows by letting n tend to infinity. □

6 A priori bound on the uniqueness zone

The goal of this section is to prove Proposition 6.1 below.

Proposition 6.1. *Let G be a transitive graph of polynomial growth. Let $p > p_c$. There are some $\chi \in (0,1)$ and $c > 0$ such that the following holds: for every $q \in [p,1]$, for every $n \geq 1$, we have*

$$\mathrm{P}_q[\mathsf{Piv}(s(n),n)] \leq cn^{-1/4}, \tag{32}$$

where $s(n) = \exp((\log n)^\chi)$.

We say that the annulus of inner radius $s(n)$ and outer radius n is a *uniqueness zone* because $\mathrm{P}_q[U(s(n),n)] = 1 - \mathrm{P}_q[\mathsf{Piv}(s(n),n)]$ converges to 1 — actually at a controlled speed. The size of this uniqueness zone will in fact determine the region in which we are able to glue clusters. This will lead to the important notion of seeds, which will be instrumental in Section 7.2.

In order to prove Proposition 6.1, we will use a bootstrap argument, where the iterations consist in alternating uses of Lemma 6.2 and Lemma 6.3.

The first lemma, directly adapted from [Cer15, Lemma 7.2], provides an upper bound on the probability of $\mathsf{Piv}(u,n)$ provided some uniform lower bound on the two-point function restricted to a ball. Conversely, the second lemma deduces some lower bound on the two-point function in a ball, provided some upper bounds on the probability of $\mathsf{Piv}(u,n)$.

From there, the proof goes as follows. Proposition 4.1 and Lemma 6.2 provide us with a good upper bound on the probability of $\mathsf{Piv}(u,n)$, for some fixed and large (u,n). Then Lemma 6.3 provides us with a good lower bound on the two-point function at the scale above. Plugging this estimate in Lemma 6.2 yields a good upper bound on the probability of $\mathsf{Piv}(u,n)$ at this larger scale. Repeated inductively, this procedure leads to the quantitative estimate (32).

Notation. In this section, given two vertices a and b of B_m, we set $\tau_{p,m}(a,b) = \mathrm{P}_p\Big[a \overset{B_m}{\longleftrightarrow} b\Big]$. Besides, for $u \geq 0$, we set S_u to be the sphere of centre o and radius u, where o is the root of G that we use in the notation $B_m = B_m(o)$.

Lemma 6.2. *For all $p \in [0,1]$ and $1 < u \leq m \leq n/2$, we have*

$$\mathrm{P}_p[\mathsf{Piv}(u,n)] \leq \mathrm{P}_p[\mathsf{Piv}(1,n/2)] \cdot \frac{|S_u|^2|B_m|}{\min_{a,b\in S_u} \tau_{p,m}(a,b)}.$$

Proof. Let us fix u, m and n as in the statement of Lemma 6.2. Given a vertex $a \in B_n$, we denote by $\mathcal{C}_a$ the cluster of a in B_n. Given a second vertex $b \in B_n$, define $\mathfrak{C}_a^b$ to be the family of all connected subsets of B_n that contain a, do not contain b, and intersect ∂B_n. By the union bound, we have

$$\begin{aligned}\mathrm{P}_p[\mathsf{Piv}(u,n)] &\leq \sum_{a,b\in S_u} \mathrm{P}_p[a \leftrightarrow \partial B_n,\ b \overset{B_n\setminus\mathcal{C}_a}{\longleftrightarrow} \partial B_n] \\ &= \sum_{a,b\in S_u} \sum_{C\in\mathfrak{C}_a^b} \mathrm{P}_p[b \overset{B_n\setminus C}{\longleftrightarrow} \partial B_n]\mathrm{P}_p[\mathcal{C}_a = C],\end{aligned}$$

where we used that the two last events are independent, since they depend on disjoint

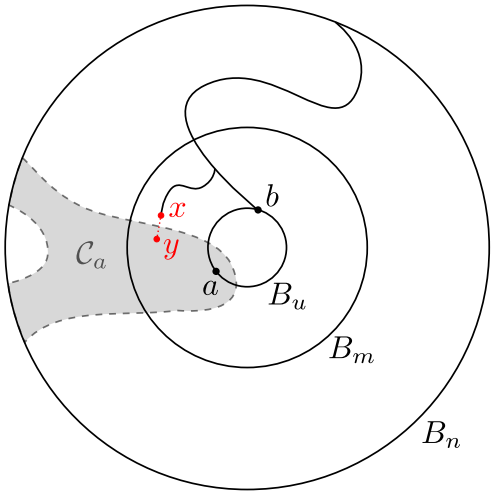

Figure 6: In the complement of $\mathcal{C}_a$, the existence of an open path inside B_m joining the cluster of b with an edge $\{x, y\} \in \partial\mathcal{C}_a$ implies that x and y are in disjoint clusters of diameter at least $n/2$.

sets of edges. Let us fix some $C \in \mathfrak{C}_a^b$. The Harris–FKG Inequality yields

$$\mathrm{P}_p[b \xleftrightarrow{B_n \setminus C} \partial B_n] \leq \frac{\mathrm{P}_p[b \xleftrightarrow{B_n \setminus C} \partial B_n,\ b \xleftrightarrow{B_m} \partial^{\mathrm{out}} C]}{\mathrm{P}_p[b \xleftrightarrow{B_m} \partial^{\mathrm{out}} C]},$$

where $\partial^{\mathrm{out}} C$ is the (deterministic) set of vertices adjacent to C but not belonging to it. Observe that for every $C \in \mathfrak{C}_a^b$, we have $\mathrm{P}_p[b \xleftrightarrow{B_m} \partial^{\mathrm{out}} C] \geq \tau_{p,m}(a, b)$. Using the bound above and independence again, we get

$$\begin{aligned}\sum_{C \in \mathfrak{C}_a^b} \mathrm{P}_p[b \xleftrightarrow{B_n \setminus C} \partial B_n]\mathrm{P}_p[\mathcal{C}_a = C] &\leq \frac{1}{\tau_{p,m}(a,b)} \sum_{C \in \mathfrak{C}_a^b} \mathrm{P}_p[\mathcal{C}_a = C,\ b \xleftrightarrow{B_n \setminus C} \partial B_n,\ b \xleftrightarrow{B_m} \partial^{\mathrm{out}} C] \\ &= \frac{1}{\tau_{p,m}(a,b)} \mathrm{P}_p[a \leftrightarrow \partial B_n,\ b \xleftrightarrow{B_n \setminus \mathcal{C}_a} \partial B_n,\ b \xleftrightarrow{B_m} \partial^{\mathrm{out}} \mathcal{C}_a].\end{aligned}$$

When the event in the last equation occurs, we can find a closed edge $e = \{x, y\} \in \partial\mathcal{C}_a$ such that $x \xleftrightarrow{B_n \setminus \mathcal{C}_a} \partial B_n$ and $y \xleftrightarrow{\mathcal{C}_a} \partial B_n$. Hence, when this event occurs, so does $\mathsf{Piv}_x(1, n - m)$: see Figure 6. Therefore, by the union bound, we have

$$\mathrm{P}_p[a \leftrightarrow \partial B_n,\ b \xleftrightarrow{B_n \setminus \mathcal{C}_a} \partial B_n,\ b \xleftrightarrow{B_m} \partial^{\mathrm{out}} \mathcal{C}_a] \leq |B_m| \mathrm{P}_p[\mathsf{Piv}(1, n/2)].$$

Putting all the equations together, we get

$$\begin{aligned}\mathrm{P}_p[\mathsf{Piv}(u, n)] &\leq \sum_{a,b \in S_u} \mathrm{P}_p[\mathsf{Piv}(1, n/2)] \frac{|B_m|}{\tau_{p,m}(a,b)} \\ &\leq \mathrm{P}_p[\mathsf{Piv}(1, n/2)] \cdot \frac{|S_u|^2 |B_m|}{\min_{a,b \in S_u} \tau_{p,m}(a,b)}.\end{aligned}$$

□

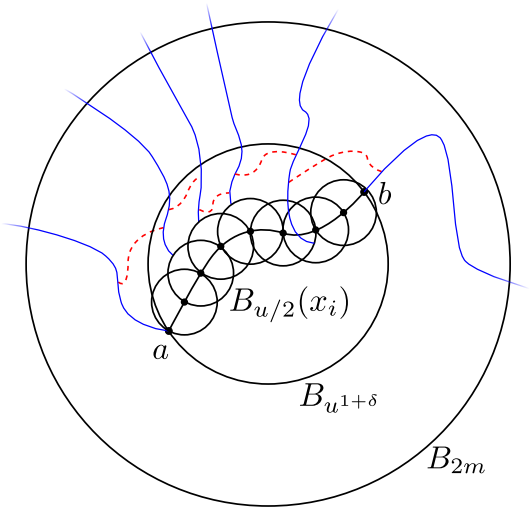

Figure 7: The event $\{a \leftrightarrow \infty,\ b \leftrightarrow \infty,\ B_{u/2}(x_i) \leftrightarrow \infty \text{ for all } i\}$ is depicted in blue. Notice that if there is no i such that $\mathsf{Piv}_{x_i}(u,m)$ holds, then we have the red dotted paths in the figure, thus creating an open path between a and b within B_{2m}.

Lemma 6.3. *Let $p > p_c$. There exist $\delta > 0$ and $u_0 \geq 1$ such that for every $u \geq u_0$ and every $m \geq u^{1+\delta}$, we have*

$$\mathrm{P}_p[\mathsf{Piv}(u,m)] \leq \frac{\delta}{u^\delta} \quad \Longrightarrow \quad \forall a,b \in B_{u^{1+\delta}} \quad \tau_{p,2m}(a,b) \geq \delta.$$

Proof. Take c and δ to satisfy the conclusion of Corollary 5.3. We may further assume that δ is small enough, so that the following inequality holds

$$\theta(p)^2 4^{-8(2/c)^\delta} \geq 9\delta. \tag{33}$$

Let $a,b \in B_{u^{1+\delta}}$. We can find vertices $x_0 = a, x_1, \dots, x_k = b$ in B_m such that $k \leq 2\frac{\lfloor u^{1+\delta} \rfloor}{\lfloor u/2 \rfloor} \leq 8u^\delta$ and $\forall i < k,\ d(x_i, x_{i+1}) \leq u/2$. Assume that $a \leftrightarrow \infty$, $b \leftrightarrow \infty$, and $B_{u/2}(x_i) \leftrightarrow \infty$ for all $i \in [\![1, k-1]\!]$. Then there are two possibilities: either $a \overset{B_{2m}}{\longleftrightarrow} b$, or there is some $i \in [\![1,k]\!]$ such that $\mathsf{Piv}_{x_i}(u,m)$ occurs. See Figure 7. Therefore, by the union bound, we get

$$\mathrm{P}_p[a \leftrightarrow \infty,\ b \leftrightarrow \infty,\ B_{u/2}(x_i) \leftrightarrow \infty \text{ for all } i] \leq \tau_{p,2m}(a,b) + \sum_{i=1}^{k} \mathrm{P}_p[\mathsf{Piv}_{x_i}(u,m)].$$

By our choice of c and δ, for every i, we have $\mathrm{P}_p[B_{u/2}(x_i) \leftrightarrow \infty] \geq 1 - \frac{1}{(cu/2)^\delta}$. By the Harris–FKG Inequality, by transitivity of G, and because P_p is invariant under graph automorphisms, we have

$$\begin{aligned}\tau_{p,2m}(a,b) &\geq \mathrm{P}_p[a \leftrightarrow \infty]\mathrm{P}_p[b \leftrightarrow \infty](1-(cu/2)^{-\delta})^{8u^\delta} - k\mathrm{P}_p[\mathsf{Piv}(u,m)] \\ &\geq \theta(p)^2(1-(cu/2)^{-\delta})^{8u^\delta} - 8u^\delta \mathrm{P}_p[\mathsf{Piv}(u,m)].\end{aligned}$$

Assuming that $u \geq u_0 := 2^{1+1/\delta}/c$ and that $\mathrm{P}_p[\mathsf{Piv}(u,m)] \leq \frac{\delta}{u^\delta}$, we get

$$\tau_{p,2m}(a,b) \geq \theta(p)^2 4^{-8(2/c)^\delta} - 8\delta.$$

By (33), the proof is complete. □

We are now ready to prove Proposition 6.1, following the strategy described at the beginning of the section.

Proof of Proposition 6.1. Let $p > p_c$. Let $\delta \in (0, 1/4)$ and $u_0 \geq 2$ be as in Lemma 6.3. By Lemma 6.2, Proposition 4.1 and the polynomial growth of G, we can fix $c_1 = c_1(G, p) \geq 1$ such that for every $1 < u \leq m \leq n/2$, we have

$$\forall q \geq p \quad \mathrm{P}_q[\mathsf{Piv}(u, n)] \leq \frac{c_1 u^{2d} m^d n^{-1/3}}{\min_{a,b \in S_u} \tau_{p,m}(a, b)}. \tag{34}$$

Applying the inequality above to $u = m = u_0$ allows us to choose $n_0 \geq \max(2(c_1/\delta^2)^{1/d}, u_0^{1+\delta})$ such that

$$\forall q \geq p \quad \mathrm{P}_q[\mathsf{Piv}(u_0, n_0)] \leq \frac{\delta}{n_0^{1/4}}.$$

Consider the sequences (u_k) and (n_k) defined by $u_{k+1} = u_k^{1+\delta}$, $n_{k+1} = n_k^{50d}$ for every $k \geq 0$. By induction, we will prove that for every $k \geq 0$,

$$\forall q \geq p \quad \mathrm{P}_q[\mathsf{Piv}(u_k, n_k)] \leq \frac{\delta}{n_k^{1/4}}. \tag{35}$$

This will conclude the proof of (32) for any $\chi < \log(1+\delta)/\log(50d)$ along the sequence (n_k). The statement for general n follows by interpolation.

Let $k \geq 0$ and assume that (35) holds. Since $n_k^{1/4} \geq u_k^{\delta}$, the quantity $\mathrm{P}_p[\mathsf{Piv}(u_k, n_k)]$ is at most δ/u_k^{δ}. Applying Lemma 6.3 to $u = u_k$ and $m = n_k \geq u_k^{1+\delta}$ gives

$$\forall a, b \in B_{u_{k+1}} \quad \tau_{p,2n_k}(a, b) \geq \delta.$$

Equation (34) applied to $u = u_{k+1}$, $m = 2n_k$ and $n = n_{k+1}$ finally gives

$$\forall q \geq p \quad \mathrm{P}_q[\mathsf{Piv}(u_{k+1}, n_{k+1})] \leq \frac{c_1 u_{k+1}^{2d} (2n_k)^d n_{k+1}^{-1/3}}{\delta} \leq \frac{\delta}{n_{k+1}^{1/4}},$$

where we used $u_{k+1} = u_k^{1+\delta} \leq n_k$ and $n_k \geq n_0 \geq 2(c_1/\delta^2)^{1/d}$ in the last inequality. □

The bound on the uniqueness zone gives us the following bound on the corridor function, which will be an important ingredient in the final proof of Proposition 1.3.

Corollary 6.4. *Let G be a graph of polynomial growth, let $p > p_c$. Let $\{s(n)\}_{n \geq 1}$ be a sequence as in Proposition 6.1. Then, for every n large enough, we have*

$$\kappa_p(s(n), n) \geq \theta(p)^2/2.$$

Proof. Notice that $x \in B_{s(n)}$ is connected to the origin by an open path in B_n if the following conditions hold:

- both o and x are connected to ∂B_n and
- there is a unique cluster crossing from $s(n)$ to n.

By the Harris–FKG Inequality and the union bound, we have

$$\forall x \in B_{s(n)} \quad \mathrm{P}_p[o \overset{B_n}{\longleftrightarrow} x] \geq \theta(p)^2 - \mathrm{P}_p[\mathsf{Piv}(s(n), n)] > \theta(p)^2/2,$$

where the last inequality holds if n is large enough, by Proposition 6.1. Since any corridor of length $s(n)$ and thickness n contains the ball B_n, the equation above implies $\kappa_p(s(n), n) > \theta(p)^2/2$ for every n large enough. □

Corollary 6.5. *Let G be a transitive graph of polynomial growth. For all $p > p_c$, $\delta \in (0, 1-p)$, and $k \geq 1$, there exists $n_0 \geq 1$ such that for every $n \geq m \geq n_0$,*

$$\kappa_p(2m, n) > \frac{1}{\log^k m} \implies \kappa_{p+\delta}(2m \log^k m, 2n) \geq \theta(p)^2/3.$$

The proof of Corollary 6.5 is analogous to the proof of Lemma 6.3. However, here, we need to apply Proposition 6.1 to get a quantitative estimate.

Proof. Let p, δ and k be as stated in Corollary 6.5, and assume that $\kappa_p(2m, n) \geq \log^{-k} m$. Let γ be a path of length at most $2m \log^k m$ starting at x and ending at y. Consider the sequence $\{s(j)\}$, where $s(j)$ is defined as in Proposition 6.1 at the parameter p. If the distance between x and y is smaller than $s(m)$, Corollary 6.4 directly concludes that x and y are connected in B_n with probability at least $\theta(p)^2/3$ (because $s(m) \leq s(n)$). We can therefore focus on the case when x and y are at distance at least $s(m)$. In that case, we may find $M \leq 2 \log^k m$ points $x_i \in \gamma$ such that $x_1 = x$, $x_M = y$ and $s(m)/2 \leq d(x_i, x_{i+1}) \leq 2m$. Let C be the corridor of thickness n around γ. Set $t = s(\lfloor s(m)/2 \rfloor)$. By our assumption on $\kappa_p(2m, n)$ and by Proposition 5.1, we have for m large enough that $s \leq m/2$ and

$$\mathrm{P}_{p+\delta}\Big[B_t(x_i) \overset{C}{\leftrightarrow} B_t(x_{i+1})\Big] \geq 1 - \frac{1}{(ct)^{\delta/20}} \geq 1 - \frac{1}{\log^k m}$$

for every i. If we assume that $x \leftrightarrow \infty$, $y \leftrightarrow \infty$ and $B_t(x_i) \overset{C}{\leftrightarrow} B_t(x_{i+1})$ for every i, then either $x \overset{C}{\leftrightarrow} y$ or there exists i such that $\mathsf{Piv}_{x_i}(t, s(m)/3)$ happens. Also, Proposition 6.1 enables us to make $\mathrm{P}_{p+\delta}[\mathsf{Piv}_{x_i}(t, s(m)/3)]$ sufficiently small. As in the proof of Lemma 6.3, we conclude by using the Harris–FKG Inequality, the union bound, the transitivity of G and the invariance of $\mathrm{P}_{p+\delta}$ under graph automorphisms. □

7 Sharp threshold results via Hamming distance

This section is devoted to the proof of the following proposition.

Proposition 7.1. *Let G be a transitive graph of polynomial growth. Let $p > p_c$. There exists n_0 such that for every $n \geq m \geq n_0$,*

$$\kappa_p(m, 2n) \leq \theta(p)^2/2 \implies \mathrm{P}_p[B_m \leftrightarrow \partial B_n] \geq 1 - e^{-\log^3 n}.$$

The statement above may seem slightly counter-intuitive at a first look: we use some negative information (the function κ_p is small) to obtain a large connection probability. Let us sketch the proof, which will be detailed in Section 7.3, and explain how this negative information can be used to our advantage.

Due to the uniqueness of the infinite cluster and the Harris–FKG Inequality, the two-point function in the whole graph is uniformly lower bounded by $\theta(p)^2$. Assume that when we restrict the connections to a box, we get something substantially smaller, in the sense that

$$\kappa_p(m, 2n) < \theta(p)^2/2.$$

By reducing the parameter from p to $p - \delta > p_c$ and using the sharp threshold results of Section 7.1 below, we strengthen this bound as follows. For $k = \log^{10}(n)$, we show that two points x and y at distance $\simeq m/k$ of each other always satisfy

$$\mathrm{P}_{p-\delta}[x \overset{B_{2n}}{\longleftrightarrow} y] \leq \frac{1}{k}. \tag{36}$$

Now, consider k points in B_m at distance at least m/k of each other. On the one hand, a proportion at least $\theta(p-\delta)$ of them are connected to the boundary of B_{2n}. On the other hand, the estimate (36) implies that all these points typically belong to different clusters of B_{2n}, which forces the paths that connect them to ∂B_n to be disjoint. It is at this point that our "negative" assumption on κ_p, combined with the "positive" assumption that $p > p_c$, yields a "positive" statement regarding connectivity of our percolation process: the expected number of disjoint paths from B_m to ∂B_{2n} is $\gtrsim k$. From this estimate, a well-known differential inequality involving the Hamming distance on the hypercube guarantees that, at parameter $p = (p-\delta) + \delta$, we have $\mathrm{P}_p[B_m \leftrightarrow \partial B_{2n}] \geq 1 - e^{-\delta k}$. This concludes the proof.

In Section 7.1, we present this new Hamming distance argument in a more general framework, since we believe it can have further applications.

7.1 Connectivity bounds via Hamming distance

In the current Section 7.1, contrary to elsewhere in the paper, G denotes any finite connected graph with vertex set $V(G)$ and edge set $E(G)$. Besides, p is arbitrary in $[0,1]$ and P_p stands for the Bernoulli bond percolation measure of parameter p on $\{0,1\}^{E(G)}$.

Proposition 7.2. *Let $A, B \subset V(G)$. Let $p \in [0,1]$ and $\theta > 0$. Assume that*

$$\min_{x \in A} \mathrm{P}_p[x \leftrightarrow B] \geq \theta \geq 2|A| \cdot \max_{\substack{x,y \in A \\ x \neq y}} \mathrm{P}_p[x \leftrightarrow y]. \tag{37}$$

Then, for every $\delta \in (0, 1-p]$, we have

$$\mathrm{P}_{p+\delta}[A \leftrightarrow B] \geq 1 - e^{-2\delta\theta|A|}.$$

Remark 7.3. The proposition above also applies to FK-percolation measures with cluster weight $q \geq 1$, and more generally to measures for which the "exponential steepness" property of [Gri06, Section 2.5] holds.

Before proving Proposition 7.2, we recall a general inequality for monotone events. The *Hamming distance* from a configuration ω to an event $\mathcal{A}$ is defined by

$$H_{\mathcal{A}}(\omega) = \inf\{H(\omega,\omega') : \omega' \in \mathcal{A}\},$$

where $H(\omega,\omega') = \sum_{e\in E(G)} |\omega_e - \omega'_e|$ is the usual Hamming distance on the hypercube $\{0,1\}^{E(G)}$. When $\mathcal{A}$ is decreasing, one can interpret $H_{\mathcal{A}}(\omega)$ as the minimal number of edges in ω that need to be closed for the event $\mathcal{A}$ to occur. Furthermore, the Hamming distance provides exponential bounds on the variation of $\mathrm{P}_p[\mathcal{A}]$ relative to p (see [Gri06, Theorem 2.53]): for every decreasing event $\mathcal{A}$ and every $p \in (0,1)$, we have

$$\frac{d}{dp} - \log \mathrm{P}_p[\mathcal{A}] \geq 4\mathrm{E}_p[H_{\mathcal{A}}].$$

By integrating the equation above and using that $H_{\mathcal{A}}$ is increasing, we get that for every $p \in [0,1]$ and every $\delta \in [0, 1-p]$,

$$\mathrm{P}_{p+\delta}[\mathcal{A}] \leq e^{-4\delta \mathrm{E}_p[H_{\mathcal{A}}]} \mathrm{P}_p[\mathcal{A}] \leq e^{-4\delta \mathrm{E}_p[H_{\mathcal{A}}]}. \tag{38}$$

Proof of Proposition 7.2. Let $A, B \subset V(G)$, let $p \in [0,1]$ and $\delta \in [0, 1-p]$. Applying (38) to the decreasing event $\mathcal{A} = \{A \not\leftrightarrow B\}$, we get

$$\mathrm{P}_{p+\delta}[A \leftrightarrow B] \geq 1 - e^{-4\delta \mathrm{E}_p[H_{\mathcal{A}}]}. \tag{39}$$

The Hamming distance $H_{\mathcal{A}}$ is clearly at least[4] the maximal number of disjoint open paths from A to B. In particular, $H_{\mathcal{A}}$ is larger than the number of disjoint clusters intersecting both A and B. By inclusion-exclusion, this number of crossing clusters can be lower bounded by

$$\sum_{x\in A} \mathbf{1}[x \leftrightarrow B] - \sum_{\substack{x,y\in A \\ x\neq y}} \mathbf{1}[x \leftrightarrow y].$$

Fixing θ as in (37) and taking the expectation above, we get

$$\mathrm{E}_p[H_{\mathcal{A}}] \geq \sum_{x\in A} \mathrm{P}_p[x \leftrightarrow B] - \sum_{\substack{x,y\in A \\ x\neq y}} \mathrm{P}_p[x \leftrightarrow y] \geq \theta|A|/2.$$

Plugging the estimate above in (39) completes the proof. □

7.2 Seeds and two-seed function

The notation G now recovers its initial meaning and denotes once again a transitive graph of polynomial growth with $d \geq 2$. In this section, we fix $p > p_c$. Let $\chi = \chi(p) \in (0,1)$ be as in Proposition 6.1. For every n, we define

$$\sigma(n) = \exp(\log^{\chi^3} n) \quad \text{and} \quad t(n) = \exp(\log^{\chi^2} n).$$

[4] Actually, if A and B are disjoint, Menger's Theorem states that this is an equality [Die17, Corollary 3.3.5] — but we only need the easy inequality.

For every positive integer n and every vertex x, set

$$S_n = B_{\sigma(n)} \quad \text{and} \quad S_n(x) = B_{\sigma(n)}(x).$$

Following the terminology introduced by Grimmett and Marstrand [GM90], we call $S_n(x)$ the *seed* of x. An important property of seeds is that they are connected to infinity with high probability. For every $n \geq 1$, define

$$\varepsilon_n = \frac{1}{\log^{10} n}.$$

Since $\sigma(n)$ is asymptotically larger than arbitrarily large powers of $\log n$, Corollary 5.3 implies that, for every n large enough, we have

$$\mathrm{P}_p[S_n \leftrightarrow \infty] \geq 1 - \tfrac{1}{100}\varepsilon_n.$$

Another important property of seeds is that they can be used to "glue" clusters. Intuitively, if two large clusters of B_n touch a certain seed $S_n(x)$, then the local uniqueness around $S_n(x)$ implies that they must be connected within B_n — which means that the two clusters are equal. Formally, we will use the following upper bound on the probability that two distinct clusters reach a fixed seed, provided by Proposition 6.1. For every n large enough we have

$$\mathrm{P}_p[\mathsf{Piv}(\sigma(n), t(n))] \leq \tfrac{1}{100}\varepsilon_n. \tag{40}$$

This follows from Proposition 6.1 together with the observations that $\sigma(n) = \exp(\log^{\chi}(t(n))$ and $c \cdot t(n)^{-1/4} \leq \frac{1}{100}\varepsilon_n$ for n large enough. Similarly, we also have

$$\mathrm{P}_p[\mathsf{Piv}(3t(n), n/2)] \leq \tfrac{1}{100}\varepsilon_n, \tag{41}$$

using that $3t(n) \leq \exp(\log^{\chi}(n/2))$ for n large.

Remark 7.4. In other works, "seed" may refer to other constructions that guarantee a high probability of connection to infinity and that can be used to glue clusters. In [GM90], this is done by defining a seed to be a large fully open box. In [MT17], one takes advantage of the fact that if an exploration reaches some vertex, then there is a long open path leading to it — in particular, see [MT17, Lemma 3.6].

We define the *two-seed function* by

$$\tau_{p,n}(x, y) = \mathrm{P}_p[S_n(x) \overset{B_n}{\longleftrightarrow} S_n(y)]$$

for every $x, y \in B_n$. Notice that for n large enough, we have $\tau_{p,n}(x, y) = 1$ whenever x and y are neighbours in B_n. As we will prove in Lemma 7.6, the two-seed function shares some features with the standard two-point function. One main advantage of replacing points by seeds is that we can make use of the following sharp threshold phenomenon.

Lemma 7.5. *Let $\delta \in (0, p - p_c)$. For every n large enough, for every $x, y \in B_n$, we have*

$$\mathrm{P}_{p-\delta}[x \overset{B_n}{\longleftrightarrow} y] \geq \varepsilon_{2n} \implies \tau_{p,2n}(x, y) \geq 1 - \varepsilon_{2n}.$$

Proof. It is a direct consequence of Proposition 5.1, together with the observation that for all $\delta > 0$, $c > 0$, for every n large enough, we have $(c\sigma(2n))^{\delta/20} \geq \varepsilon_{2n}$. □

Lemma 7.6. *For every n large enough, for all $x, y, z \in B_{n/2}$, we have*

$$\tau_{p,n}(x,z) \geq \tau_{p,n}(x,y)\tau_{p,n}(y,z) - \varepsilon_n, \tag{42}$$

$$\mathrm{P}_p[x \overset{B_n}{\longleftrightarrow} y] \geq \theta(p)^2 \cdot \tau_{p,n}(x,y) - \varepsilon_n. \tag{43}$$

Proof. We begin with the proof of (42). Without loss of generality, we may and will assume that $d(x,y) \geq d(y,z)$. We distinguish three different cases. For readability, we drop the indices p and n from the notation $\tau_{p,n}$ in this proof.

Case 1: $d(x,y) \geq d(y,z) \geq 2t(n)$.

By the Harris–FKG Inequality, we have

$$\tau(x,y)\tau(y,z) \leq \mathrm{P}_p[S_n(x) \overset{B_n}{\longleftrightarrow} S_n(y),\, S_n(y) \overset{B_n}{\longleftrightarrow} S_n(z)].$$

When there exist a cluster connecting $S_n(x)$ to $S_n(y)$ and a cluster connecting $S_n(y)$ to $S_n(z)$, then either these two clusters are connected together within $B_{t(n)}(y)$, or we observe two disjoint clusters crossing from $S_n(y)$ to the boundary of $B_{t(n)}(y)$. Therefore, by the union bound

$$\tau(x,y)\tau(y,z) \leq \tau(x,z) + \mathrm{P}_p[\mathsf{Piv}(\sigma(n), t(n))] \overset{(40)}{\leq} \tau(x,z) + \varepsilon_n.$$

Case 2: $d(x,y) \geq 2t(n) > d(y,z)$.

The probability that the seeds $S_n(z)$ and $S_n(y)$ are both connected to infinity is larger than $1 - \frac{1}{50}\varepsilon_n$. Therefore, we have

$$\tau(x,y) - \tfrac{1}{50}\varepsilon_n \leq \mathrm{P}_p[S_n(x) \overset{B_n}{\longleftrightarrow} S_n(y), S_n(y) \leftrightarrow \partial B_n, S_n(z) \leftrightarrow \partial B_n].$$

If both $S_n(y)$ and $S_n(z)$ are connected to ∂B_n, then either there exists a cluster in B_n that intersects $S_n(y)$, $S_n(z)$ and ∂B_n, or there exist two disjoint clusters crossing from $B_{3t(n)}(y)$ to $\partial B_{n/2}(y)$. Therefore, using the clusters inside $B_{t(n)}(y)$ as in Case 1, we have

$$\tau(x,y) - \tfrac{1}{50}\varepsilon_n \leq \tau(x,z) + \mathrm{P}_p[\mathsf{Piv}(\sigma(n), t(n))] + \mathrm{P}_p[\mathsf{Piv}(3t(n), n/2)] \overset{(40)+(41)}{\leq} \tau(x,z) + \tfrac{1}{50}\varepsilon_n,$$

which implies $\tau(x,z) \geq \tau(x,y) - \varepsilon_n$.

Case 3: $2t(n) > d(x,y) \geq d(y,z)$.

In this case, we use that $1 - \frac{1}{50}\varepsilon_n \leq \mathrm{P}_p[S_n(x) \leftrightarrow \partial B_n,\, S_n(z) \leftrightarrow \partial B_n]$. Reasoning as in Case 2, we get

$$1 - \tfrac{1}{50}\varepsilon_n \leq \tau(x,z) + \mathrm{P}_p[\mathsf{Piv}(3t(n), n/2)] \overset{(41)}{\leq} \tau(x,z) + \tfrac{1}{100}\varepsilon_n,$$

which implies $\tau(x,z) \geq 1 - \varepsilon_n$.

For the proof of (43), we proceed similarly. If $d(x,y) \geq 2t(n)$, we first use the Harris–FKG Inequality to show

$$\theta(p)^2\tau(x,y) \leq \mathrm{P}_p[x \leftrightarrow \partial B_n, y \leftrightarrow \partial B_n, S_n(x) \overset{B_n}{\longleftrightarrow} S_n(y)].$$

If the event estimated on the right-hand side occurs, then either x is connected to y or there is no local uniqueness around x or around y. Therefore,

$$\theta(p)^2\tau(x,y) \leq \mathrm{P}_p[x \overset{B_n}{\longleftrightarrow} y] + 2\mathrm{P}_p[\mathsf{Piv}(\sigma(n), t(n))] \overset{(40)}{\leq} \mathrm{P}_p[x \overset{B_n}{\longleftrightarrow} y] + \varepsilon_n.$$

If $d(x,y) < 2t(n)$, using the estimate (41), we directly get that $\mathrm{P}_p[x \overset{B_n}{\longleftrightarrow} y] \geq \theta(p)^2 - \varepsilon_n \geq \theta(p)^2\tau(x,y) - \varepsilon_n$. □

Lemma 7.7. *For every n large enough, the following holds. For every $u \in B_{n/2}$ satisfying*

$$\tau_{p,n}(o,u) \leq 2/3,$$

there exists a set $A \subset B_{d(o,u)}$ of cardinality at least $\log^4 n$ such that

$$\forall x, y \in A \quad \tau_{p,n}(x,y) \leq 1 - \varepsilon_n.$$

Remark 7.8. When the underlying graph G is the hypercubic lattice $\mathbb{Z}^d$, the lemma above can be easily proved using the symmetries of the graph.

Proof. Let n be a large integer. Let $u \in B_{n/2}$ and fix γ some geodesic path from the origin o to u. If a and b denote two vertices belonging to this fixed path, we denote by $[a,b]$ the set of all vertices x of the path lying between a and b, i.e. satisfying $d(a,b) = d(a,x) + d(x,b)$. Sets of this form will be called *segments* in this proof. By convention, we always assume that $d(o,a) \leq d(o,b)$ when we consider a segment $[a,b]$.

We will build the set A as a subset of the segment $I = [o,u]$, in a way reminiscent of the construction of the triadic Cantor set. We construct two suitable subsegments $I_0 = [a_0,b_0]$ and $I_1 = [a_1,b_1]$, such that the two-seed function $\tau(x,y)$, $x \in I_0$, $y \in I_1$ is well-controlled and both $\tau(a_0,b_0)$ and $\tau(a_1,b_1)$ have a nice upper bound. Then, we repeat this splitting operation in each of the two segments. After k steps, we construct 2^k intervals with a well-controlled two-seed functions between each segments. The proof is then concluded by choosing a suitable number of steps. The "splitting operation" of an interval at one step relies on the following claim.

For convenience, as in the proof of Lemma 7.6, we drop the indices p and n from the notation $\tau_{p,n}$ and ε_n.

Claim 7.9. *Let $\alpha \geq \frac{1}{2}$. Let $[a,b]$ be a segment such that $\tau(a,b) \leq \alpha$. Then, there exist two vertices $a', b' \in [a,b]$ such that*

$$\forall x \in [a,a'] \ \forall y \in [b',b] \quad \tau(x,y) \leq \alpha^{1/3} + 6\varepsilon \tag{44}$$

and

$$\max(\tau(a,a'), \tau(b,b')) \leq \alpha^{1/3}. \tag{45}$$

Proof of Claim 7.9. First, notice that if $\alpha \geq 1$, then the claim is straightforward: taking I_1 and I_2 *any* two (possibly equal) subsegments of I works. We therefore assume that $\alpha < 1$.

Observe that, provided n is large enough, $x \mapsto \tau(a,x)$ is ε-Lipschitz on $B_{n/2}$. This follows from Lemma 7.6 and the fact that for any two adjacent vertices x and y in $B_{n/2}$, we have $\tau(x,y) = 1$. Since $\tau(a,a) = 1$ and $\tau(a,b) \leq \alpha < \alpha^{1/3}$, by the discrete Intermediate Value Theorem for Lipschitz functions, we can define $a' \in [a,b]$ satisfying $\alpha^{1/3} \geq \tau(a,a') \geq \alpha^{1/3} - \varepsilon$. By choosing a' to be the closest vertex to a with this property, we additionally have

$$\forall x \in [a,a'] \quad \tau(a,x) \geq \alpha^{1/3} - \varepsilon.$$

Similarly, we take $b' \in [a,b]$ such that $\tau(b',b) \leq \alpha^{1/3}$ and

$$\forall x \in [b',b] \quad \tau(x,b) \geq \alpha^{1/3} - \varepsilon.$$

Equation (45) holds by definition and it remains to prove (44). Let $x \in [a,a']$ and $y \in [b,b']$. Since $\tau(a,b) \leq \alpha$, Lemma 7.6 implies

$$\alpha \geq \tau(a,x)\tau(x,y)\tau(y,b) - 2\varepsilon.$$

Using the inequality $\alpha \geq \frac{1}{2}$ and recalling that both $\tau(a,x)$ and $\tau(y,b)$ are larger than $\alpha^{1/3} - \varepsilon$, we obtain

$$\tau(x,y) \leq \alpha^{1/3} + 2\alpha^{-2/3}\varepsilon + 2\varepsilon \leq \alpha^{1/3} + 6\varepsilon.$$

□

Let us prove by induction that for every $k \geq 0$, we can find a family of segments $I_1, \ldots, I_{2^k}$ that satisfies the following conditions:

- for $i \neq j$, for any $x \in I_i$ and any y in I_j, we have $\tau(x,y) \leq (2/3)^{1/3^k} + 6k\varepsilon$,
- for any segment $[a,b]$ of the family, we have $\tau(a,b) \leq (2/3)^{1/3^k}$.

For $k = 0$, taking $[o,u]$ works, by hypothesis. Let $k \geq 0$ be such that the property holds at rank k, and let us prove that the property holds at rank $k+1$. Let us take $I_1, \ldots, I_{2^k}$ as above. For every i, we apply Claim 7.9 to I_i, which yields two subsegments $I_i^{(1)}$ and $I_i^{(2)}$ of I_i. For $x \in I_i^{(1)}$ and $y \in I_i^{(2)}$, we have

$$\tau(x,y) \leq \left((2/3)^{1/3^k} + 6k\varepsilon\right)^{1/3} + 6\varepsilon \leq (2/3)^{1/3^{k+1}} + 6(k+1)\varepsilon.$$

Likewise, if $[a,b]$ denotes either $I_i^{(1)}$ or $I_i^{(2)}$, we have $\tau(a,b) \leq (2/3)^{1/3^{k+1}}$. It remains to check that if $i \neq j$, then for every $x \in I_i^{(0)} \cup I_i^{(1)}$ and every $y \in I_j^{(0)} \cup I_j^{(1)}$, we have $\tau(x,y) \leq (2/3)^{1/3^{k+1}} + 6(k+1)\varepsilon$. But this is clear: since $x \in I_i$, $y \in I_j$ and $i \neq j$, we have

$$\tau(x,y) \leq (2/3)^{1/3^k} + 6k\varepsilon \leq (2/3)^{1/3^{k+1}} + 6(k+1)\varepsilon.$$

The result thus holds for all k.

Let us use this result for $k = \lceil 8 \log\log n \rceil$, which we will handle as $8 \log\log n$ for readability. Let $I_1, \dots, I_{2^k}$ be as above for this specific value of k. Let $x_1 \in I_1, \dots, x_{2^k} \in I_{2^k}$. Given $i \neq j$, we have

$$\tau(x_i, x_j) \leq (2/3)^{1/3^k} + 6k\varepsilon_n \leq e^{\log(2/3)/\log(n)^{8\log 3}} + 48\frac{\log\log n}{\log^{10} n}.$$

Since $8\log(3) < 10$, taking n large enough guarantees that $\tau(x_i, x_j) \leq 1 - \varepsilon_{n/2}$. In particular, we have $\tau(x_i, x_j) < 1$, hence $x_i \neq x_j$. We set $A = \{x_1, \dots, x_{2^k}\}$. It remains to check that A contains at least $\log^4 n$ elements, which is straightforward as $2^k = (\log n)^{8\log 2}$ and $8\log 2 > 4$. □

7.3 Proof of Proposition 7.1

Assume that $\kappa_{p+\delta}(m, 2n) \leq \theta(p)^2/2$. We can thus take $u \in B_m$ such that

$$\mathrm{P}_{p+\delta}[o \overset{B_{2n}}{\longleftrightarrow} u] \leq \theta(p)^2/2.$$

By Lemma 7.6, this implies that $\tau_{p+\delta,2n}(o, u) \leq \frac{1}{2} + \frac{\varepsilon_{2n}}{\theta(p)^2} \leq \frac{2}{3}$, provided n is taken large enough. By Lemma 7.7, we can find a set $A \subset B_m$ satisfying $\log^4 n \leq |A| \leq \log^5 n$ such that for every $x, y \in A$, we have

$$\tau_{p+\delta,2n}(x, y) \leq 1 - \varepsilon_{2n}.$$

In words, the inequality above states that the two-seed function is not too close to 1 for every pair of points of A. By decreasing the edge density from $p + \delta$ to p, we obtain that the points of A are pairwise connected with low probability. More precisely, the equation above and the sharp-threshold result of Lemma 7.5 imply

$$\forall x, y \in A \quad \mathrm{P}_p[x \overset{B_n}{\longleftrightarrow} y] \leq \varepsilon_{2n}.$$

Together with $\min_{x \in A} \mathrm{P}_p[x \leftrightarrow \partial B_n] \geq \theta(p) \geq 2|A|\varepsilon_{2n}$ (which holds for n large enough), we obtain

$$\min_{x \in A} \mathrm{P}_p[x \leftrightarrow \partial B_n] \geq \theta(p) \geq 2|A| \max_{x,y \in A} \mathrm{P}_p[x \overset{B_n}{\longleftrightarrow} y].$$

By applying Proposition 7.2 to the graph G induced by the ball B_n, we finally get

$$\mathrm{P}_{p+\delta}[A \leftrightarrow \partial B_n] \geq 1 - e^{-2\delta\theta(p)|A|},$$

which concludes the proof since $A \subset B_m$ and $|A| \geq \log^4 n$. □

8 Uniqueness via sprinkling

In this section, we establish the following proposition, which revisits the techniques of [BT17]. Recall that the coupled measure $\mathbf{P}$ and the sprinkled uniqueness event $U_{p,q}(m, n)$ are defined in Section 1.4.

Proposition 8.1. *Let G be a transitive graph of polynomial growth with $d \geq 2$. Let $R \geq 1$ be so that the conclusion of Lemma 2.1 holds. Let $0 \leq p \leq 1$. Let $\eta, \delta > 0$ be such that $p+\delta \leq 1$. Let (k, m, n) be such that $R \leq k \leq m \leq n$ and $n \geq \log |B_{2n}|$. We make the following two assumptions:*

(a) $\mathrm{P}_p[B_k \leftrightarrow \partial B_{6n}] \geq 1 - \frac{\eta}{|B_{2n}|}$,

(b) $\forall x, y \in B_{3k}\ \mathrm{P}_p[x \overset{B_m}{\longleftrightarrow} y] \geq \delta$.

Then,

$$\mathbf{P}[U_{p,p+\delta}(n, 4n)] \geq 1 - \eta - 100|B_{2n}|^2 \exp\left(-\frac{\delta^3 n}{8m \log |B_{2n}|)}\right).$$

In this section, we consider the family of coupled configuration $(\omega_p)_{p\in[0,1]}$ under the measure $\mathbf{P}$. Our goal is to show that, with high probability, all the p-clusters crossing from B_n to ∂B_{4n} are $(p+\delta)$-connected to each other within the annulus $A_{n,2n}$. In this section and contrary to the *connected* annuli of Section 3, the annulus $A_{r,s}$ is defined as the set of all edges $e \subset B_{4n}$ that intersect B_s but not B_r. If we achieve our goal, the crossing p-clusters will a fortiori get $(p+\delta)$-connected within the larger set B_{4n}.

For every $r \leq 2n$, consider the percolation configuration Y_r in B_{4n} defined as follows: for every edge $e \subset B_{4n}$,

$$Y_r(e) := \begin{cases} \omega_p(e)\mathbf{1}_{\{e \overset{}{\underset{p}{\longleftrightarrow}} \partial B_{4n}\}} & \text{if } e \text{ intersects } B_r \\ \omega_{p+\delta}(e) & \text{if } e \in A_{r,4n}. \end{cases}$$

The configuration Y_r can be explored and understood in the following manner. First, we perform the classical exploration from ∂B_{4n} of all the ω_p-clusters touching it — by doing so, we also reveal the (closed) edges at the boundary of these clusters. Then, we further reveal the $\omega_{p+\delta}$-status of every single edge included in $A_{r,4n}$. Conditionally on Y_r, the status of the edges in B_r are independent. Besides, for any such e that does not already satisfy $Y_r(e) = 1$, we have:

$$\mathbf{P}[\omega_{p+\delta} = 1 \mid Y_r] \begin{cases} \geq \delta & \text{if } e \text{ is adjacent to some } e' \text{ satisfying } Y_r(e') = 1, \\ = p + \delta & \text{otherwise.} \end{cases}$$

For every configuration ω in B_{4n} and every $r \leq 2n$, define $\mathscr{C}_r(\omega)$ to be the set of all the clusters of ω intersecting both B_r and ∂B_{4n}. We further set

$$N_r(\omega) = |\mathscr{C}_r(\omega)|.$$

We aim at proving that with high probability, we have $N_n(Y_n) = 1$. We will do so by making use of the following event, which roughly states that "large clusters grow from everywhere in B_{2n}":

$$\mathcal{E} = \{\forall x \in B_{2n}\ B_k(x) \overset{}{\underset{p}{\longleftrightarrow}} \partial B_{4n}\}.$$

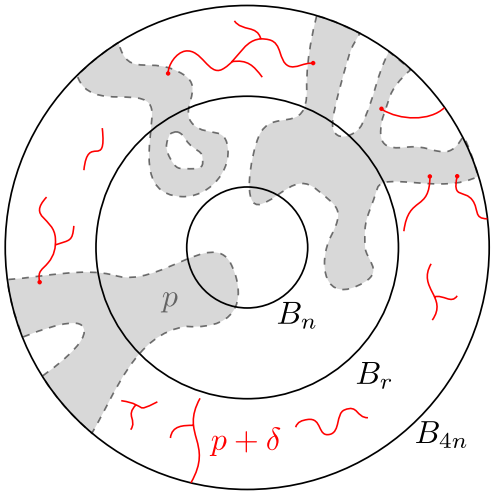

Figure 8: The configuration Y_r is the union of the grey area and the red one.

We will prove that $N_n(Y_n) = 1$ holds with high probability by proving that $\mathcal{E}$ holds with high probability and that, conditionally on this event, $N_n(Y_n) = 1$ holds with high probability.

This will be done by repeatedly using the following lemma. Lemma 8.2 guarantees that sprinkling a layer of thickness $\simeq \frac{n}{\log B_{2n}}$ is very likely to divide the number of "crossing clusters" by a factor at least 2 — except if this number is already 1. Iterating this process $\simeq \log B_{2n}$ times inwards (going from the sphere of radius $2n$ to that of radius n) will yield $N_n(Y_n) = 1$ with high probability.

Lemma 8.2. *For every $r, s \in [n, 2n]$ such that $r \le s$, we have*

$$\mathbf{P}[\mathcal{E} \cap \{N_r(Y_r) > \max(1, \tfrac{1}{2}N_s(Y_s))\}] \le 25|B_{2n}| \exp\left(-\frac{\delta^3(s-r)}{3m}\right). \tag{46}$$

Proof of Lemma 8.2. Fix $r, s \in [n, 2n]$ such that $r \le s$. The configuration Y_r is obtained from Y_s by adding some open edges in the annulus $A_{r,s}$. In particular, we have $Y_r \ge Y_s$ and our goal is to show that, with high probability, every cluster of $\mathcal{C}_s(Y_s)$ gets merged with at least one other cluster of $\mathcal{C}_s(Y_s)$ in the configuration Y_r. To analyse this "merging effect", let us introduce the configuration

$$Z = Y_r - Y_s.$$

Let us condition on the possible values for Y_s. Say that a configuration ξ in B_{4n} is *admissible* if $\mathbf{P}[Y_s = \xi] > 0$ and $B_k(x)$ is connected to ∂B_{4n} in ξ for every $x \in B_{2n}$. Writing $\mathbf{P}^\xi := \mathbf{P}[\,\cdot\,|Y_s = \xi]$ for every admissible ξ, we can rewrite the left-hand side of (46) as follows:

$$\mathbf{P}[\mathcal{E} \cap \{N_r(Y_r) > \max(1, \tfrac{1}{2}N_s(Y_s))\}] = \sum_{\xi \text{ admissible}} \mathbf{P}^\xi[N_r(\xi + Z) > \max(1, \tfrac{1}{2}N_s(\xi))]\mathbf{P}[Y_s = \xi].$$

From now on, we fix an admissible configuration ξ and the proof will be complete once we show

$$\mathbf{P}^\xi[N_r(\xi + Z) > \max(1, \tfrac{1}{2}N_s(\xi))] \le 25|B_{2n}| \exp\left(-\frac{\delta^3(s-r)}{3m}\right). \tag{47}$$

Notice that under $\mathbf{P}^\xi$, the coordinates $Z(e)$, $e \subset B_s$, are independent and satisfy

$$\mathbf{P}^\xi[Z(e)=1]\begin{cases} = p+\delta & \text{if } e \in A_{r,s} \text{ and } e \text{ is not adjacent to an edge of } \xi \\ \geq \delta & \text{if } e \in A_{r,s},\ \xi(e)=0 \text{ and } e \text{ is adjacent to an edge of } \xi, \\ = 0 & \text{if } e \text{ intersects } B_r. \end{cases}$$

If $N_r(\xi)=1$, then we also have $N_r(\xi+Z)=1$. Indeed, all the open edges of ξ intersecting B_r are already connected to ∂B_{4n} in ξ, and no open edge of Z intersects B_r: therefore, adding the edges of Z to ξ cannot create a new cluster crossing from B_r to ∂B_{4n}. As a result, if $N_r(\xi)=1$, the left-hand side of (47) is equal to 0 and there is nothing to prove. We thus assume that $N_r(\xi)>1$ and we will prove that

$$\mathbf{P}^\xi[N_r(\xi+Z) > \tfrac{1}{2}N_s(\xi)] \leq 25|B_{2n}|\exp\left(-\frac{\delta^3(s-r)}{3m}\right).$$

We will say that a cluster of $\mathscr{C}_s(\xi)$ is *merged* if it is connected to at least one other cluster of $\mathscr{C}_s(\xi)$ in the configuration $\xi+Z$. Please note that even if the cluster under consideration is taken in $\mathcal{C}_r(\xi) \subset \mathcal{C}_s(\xi)$, it is said to be merged if it is connected to at least one other cluster of $\mathscr{C}_{\underline{s}}(\xi)$ in the configuration $\xi+Z$. We will prove that typically any cluster of $\mathscr{C}_s(\xi)$ that crosses the annulus $A_{r,s}$ is merged with high probability. More precisely we will prove that

$$\forall C \in \mathscr{C}_r(\xi) \quad \mathbf{P}^\xi[C \text{ is not merged}] \leq 25\exp\left(-\frac{\delta^3(s-r)}{3m}\right). \tag{48}$$

The main idea behind (48) is that the cluster C crosses $\simeq (s-r)/m$ disjoint annuli of thickness m. In each annulus, the cluster C comes at distance $\leq k$ to another cluster of $\mathscr{C}_s(\xi)$. Morally, Hypothesis (b) can thus be used to bound the probability that the two clusters get connected by a $(\xi+Z)$-open path lying inside the considered annulus. Since these events are independent (we have disjoint annuli), we obtain the bound (48). We postpone the rigorous derivation of (48) to the end of the proof, and we now explain how to deduce (47) from (48).

The key observation is that if every cluster of $\mathscr{C}_r(\xi)$ is merged, then $N_r(\xi+Z) \leq \frac{1}{2}N_s(\xi)$. To see this, let us assume that every of $\mathscr{C}_r(\xi)$ is merged, and let us prove that every cluster of $\mathcal{C}_r(\xi+Z)$ contains at least two clusters of $\mathcal{C}_s(\xi)$. Let $C \in \mathcal{C}_r(\xi+Z)$. We have already seen that adding Z to ξ can create no new cluster crossing from B_r to ∂B_{4n}. In other words, we can fix some $C_1 \in \mathcal{C}_r(\xi)$ such that $C_1 \subset C$. Since C_1 is merged, we can take some other cluster $C_2 \in \mathcal{C}_s(\xi)$ such that $C_2 \subset C$, which ends the proof of the observation. Using this observation, the union bound, (48) and $r \leq 2n$ yields

$$\begin{aligned}\mathbf{P}^\xi[N_r(\xi+Z) > \tfrac{1}{2}N_s(\xi)] &\leq \mathbf{P}^\xi[\exists C \in \mathscr{C}_r(\xi),\ C \text{ is not merged}] \\ &\leq 25|\mathscr{C}_r(\xi)|\exp\left(-\frac{\delta^3(s-r)}{3m}\right) \leq 25|B_{2n}|\exp\left(-\frac{\delta^3(s-r)}{3m}\right),\end{aligned}$$

which is the desired inequality.

We now give the details of the proof of (48). Fix $C \in \mathscr{C}_r(\xi)$. Set

$$J := 3m\mathbb{Z} \cap [r+3m, s-3m].$$

For every $j \in J$, we will use the set S_j^∞ introduced on page 12. Notice that $|J| \geq \frac{s-r}{3m} - 3$ and that C intersects every S_j^∞, by Lemma 2.6. Define $\widetilde{C}$ to be the union of all clusters of $\mathcal{C}_s(\xi) \setminus \{C\}$. We claim that for every $j \in J$, there exists $x_j \in S_j$ such that

$$\mathrm{d}(x_j, C) \leq 3k \quad \text{and} \quad \mathrm{d}(x_j, \widetilde{C}) \leq 3k.$$

To prove this, consider a path γ from some vertex $x \in C \cap S_j$ to another vertex in $\widetilde{C} \cap S_j$ which stays in the R-neighbourhood of S_j (such path exists by Lemma 2.1). If x is at distance k or less from $\widetilde{C}$, we simply choose $x_j = x$ and we are done. Otherwise, consider $z \in \gamma$ at distance exactly $k+1$ from $\widetilde{C}$. Since $B_k(z)$ is connected to ∂B_{4n} in ξ (because ξ is admissible) and $B_k(z) \cap \widetilde{C} = \varnothing$, we must have $B_k(z) \cap C \neq \varnothing$. It now suffices to choose $x_j \in S_j \cap B_R(z)$, use the triangle inequality and remember that $R \leq k$ to get the claim.

For every $j \in J$, fix some x_j as above and introduce the event

$$\mathcal{E}_j = \left\{ C \xleftrightarrow{B_m(x_j)} \widetilde{C} \text{ in } \xi + Z \right\}.$$

By Hypothesis (b) and the standard "sprinkling" argument, we have

$$\mathbf{P}[\mathcal{E}_j] \geq \delta^2 \min_{x,y \in B_{3k}(x_j)} \mathrm{P}_p[x \xleftrightarrow{B_m(x_j)} y] \geq \delta^3.$$

As soon as one of the event $\mathcal{E}_j$ occurs, the cluster C gets connected to another cluster of $\mathscr{C}_s(\xi)$ in the configuration $\xi + Z$. Since the events $\mathcal{E}_j$ are independent, we get

$$\mathbf{P}^\xi[C \text{ is not merged}] \leq (1-\delta^3)^{|J|} \leq \exp\left(3\delta^3 - \frac{\delta^3(s-r)}{3m}\right) \leq 25 \exp\left(\frac{\delta^3(s-r)}{3m}\right).$$

□

Proof of Proposition 8.1. We wish to prove that

$$\mathbf{P}[N_n(Y_n) > 1] \leq \eta + 100|B_{2n}|^2 \exp\left(-\frac{\delta^3 n}{8m \log|B_{2n}|)}\right).$$

First, we have

$$\mathbf{P}[N_n(Y_n) > 1] \leq \mathbf{P}[\mathcal{E}^c] + \mathbf{P}[\mathcal{E} \cap \{N_n(Y_n) > 1\}]. \tag{49}$$

The first term on the right hand side can be bounded as follows, using the union bound, automorphism-invariance and Hypothesis (a):

$$\mathbf{P}[\mathcal{E}^c] \leq \sum_{x \in B_{2n}} \mathrm{P}_p[B_k(x) \not\leftrightarrow \partial B_{6n}(x)] \leq |B_{2n}| \mathrm{P}_p[B_k \not\leftrightarrow \partial B_{6n}] \leq \eta. \tag{50}$$

In order to bound the second term, we use the property established in Lemma 8.2, namely that clusters merge with high probability in annuli of the form $A_{r,r+\Delta}$ for $\Delta := \lceil \frac{n}{4\log|B_{2n}|} \rceil$, when sprinkling from p to $p+\delta$. Recall that $N_r(Y_r)$ counts the number of p-clusters crossing from B_r to ∂B_{4n}, where the clusters are identified if they are $(p+\delta)$-connected in the annulus $A_{r,4n}$. We start with the trivial bound $N_{2n}(Y_{2n}) \leq |B_{2n}|$ for the

number of clusters in Y_{2n} intersecting the ball $|B_{2n}|$. Applying Lemma 8.2 to $r = 2n - \Delta$ and $s = 2n$, we obtain

$$\mathbf{P}[\mathcal{E} \cap \{N_{2n-\Delta}(Y_{2n-\Delta}) > \max(1, \tfrac{1}{2}|B_{2n}|)\}] \leq 25|B_{2n}| \exp\left(-\frac{\delta^3\Delta}{3m}\right).$$

Then, by induction, we get for every $i \leq n/\Delta$,

$$\mathbf{P}[\mathcal{E} \cap \{N_{2n-i\Delta}(Y_{2n-i\Delta}) > \max(1, 2^{-i}|B_{2n}|)\}] \leq 25i|B_{2n}| \exp\left(-\frac{\delta^3\Delta}{3m}\right).$$

Choosing $i = \lfloor n/\Delta \rfloor$ and observing that $|B_{2n}| \leq 2^i$, we obtain

$$\mathbf{P}[\mathcal{E} \cap \{N_n(Y_n) > 1] \leq 25 \times 4 \times |B_{2n}| \log |B_{2n}| \exp\left(-\frac{\delta^3\Delta}{3m}\right). \tag{51}$$

Plugging the two bounds (50) and (51) in (49) concludes the proof. □

9 Proof of Proposition 1.3

Our goal is to prove that there exists a unique large cluster in large balls with high probability: for every n large enough, we want to show that

$$\mathrm{P}_p[B_{n/10} \leftrightarrow \partial B_n,\ U(n/5, n/2)] \geq 1 - e^{-\sqrt{n}}. \tag{52}$$

Let us start with two reductions of the problem.

First reduction: a lower bond on the two-point function in long corridors. At several places in the paper, we saw that the desired local uniqueness is intimately related to lower bounds on the two-point function inside finite regions. In particular, Proposition 3.1 asserts that Equation (52) holds if we can prove the lower bound

$$\forall p > p_c \quad \limsup_{n\to\infty} \kappa_p(n \log^{3d} n, n) > 0. \tag{53}$$

Second reduction: a lower bond on the two-point function in short corridors. In Section 5, we proved that crossing probabilities always undergo a sharp threshold phenomenon: the probabilities jump rapidly from 0 to 1 as p varies. An important consequence of this phenomenon is quantified in Corollary 6.5: a lower bound on the two-point function in a corridor at parameter p implies a lower bound in a longer corridor at a slightly larger parameter $p + \delta$. Using this corollary, one can see that Equation (53) will be established if one can show the following weaker statement

$$\forall p > p_c \quad \limsup_{n\to\infty} \kappa_p(n/\log^4 n, n) > 0.$$

Setup of the proof Let us fix $p_0 > p_c$. Let δ be such that $0 < \delta < \theta(p_0)^2/3$ and $p_0 + 2\delta < 1$. We will prove that

$$\limsup_{n\to\infty} \kappa_{p_0+2\delta}(n/\log^4 n, n) > 0. \tag{54}$$

Observe that if infinitely many values of n satisfy $\kappa_{p_0}(n/\log^4 n, n) \geq \delta$, then Equation (54) holds by monotonicity. We may thus assume that we are given some $n_0 \geq 100$ such that

$$\forall n \geq n_0 \quad \kappa_{p_0}(n/\log^4 n, n) < \delta. \tag{H}$$

The remaining of the section is devoted to the proof of (54) under the hypothesis (**H**).

High probability of connection for balls The hypothesis (**H**) is a "negative" estimate for the connection probabilities. This is useful because it allows us to apply the results of Section 7 (sharp threshold via Hamming distance), in which we showed that an upper bound on the two-point function implies a "positive" estimate, namely that balls are connected far away with high probability. In order to obtain the strongest estimate possible, for every $n \geq n_0$, we define $\alpha(n)$ to be the smallest integer satisfying

$$\kappa_{p_0}(\alpha(n), n) < \delta. \tag{55}$$

The hypothesis (**H**) and Corollary 6.4 imply that there exist $\chi > 0$ and $n_1 \geq n_0$ such that

$$\forall n \geq n_1 \quad \exp(\log^\chi n) \leq \alpha(n) \leq n/\log^4 n. \tag{56}$$

Since $\delta < \theta(p_0)^2/2$, the upper bound (55) together with Proposition 7.1 implies that we can choose $n_2 \geq n_1$ such that

$$\forall n \geq n_2 \quad \mathrm{P}_{p_0}[B_{\alpha(n)} \leftrightarrow \partial B_{n/2}] \geq 1 - e^{-\log^3(n/2)}. \tag{57}$$

In order to conclude the proof from the strong estimate above, we will apply Proposition 8.1 (uniqueness via sprinkling). To achieve this, we need to complement Equation (57) with a lower bound on the two-point function at scale $\alpha(n)$, which is the object of the next paragraph.

An improved lower bound on the two-point function via sharp-threshold Since $\alpha(n)$ was chosen to be the smallest integer satisfying the inequality (55), this inequality is not satisfied for $\alpha(n) - 1$ and we have the lower bound

$$\forall n \geq n_2 \quad \kappa_{p_0}(\alpha(n) - 1, n) \geq \delta.$$

This bound will not be sufficient, and we use a sharp threshold result to improve it. Using Corollary 6.5 and the lower bound (56) on $\alpha(n)$, we get, for every n large enough,

$$\kappa_{p_0+\delta}(\alpha(n)\log^5 n, 2n) > \delta. \tag{58}$$

Uniqueness via sprinkling We want to apply Proposition 8.1 along a certain infinite sequence of good scales, which we now define. A scale n is said to be *good* if

$$\alpha(n) \le \log^4(n) \cdot \alpha(\lceil n/\log^4 n \rceil). \tag{59}$$

It is elementary to check that there exist infinitely many good scales. Indeed, if (59) fails for every n large enough, then we have $\limsup \frac{\alpha(n)}{n} > 0$, which contradicts $\alpha(n) \le n/\log^4 n$.

Let $n \ge n_2$ be a good scale. Set $k = \alpha(n)$ and $m = \lceil n/\log^4 n \rceil$. By (57) and monotonicity, we have

$$\mathrm{P}_{p_0+\delta}[B_k \leftrightarrow \partial B_{n/2}] \ge 1 - e^{-\log^3(n/2)}. \tag{60}$$

Furthermore, Equation (58) applied to m gives $\kappa_{p_0+\delta}(\alpha(m)\log^5 m, 2m) \ge \delta$. Then, using that n is a good scale and monotonicity, we obtain

$$\kappa_{p_0+\delta}(6k, 2m) \ge \delta. \tag{61}$$

Equations (60) and (61) enable us to use Proposition 8.1, which yields

$$\mathbf{P}[U_{p_0+\delta,p_0+2\delta}(n/40, n/10)] \ge 1 - e^{-\log^2 n}. \tag{62}$$

Conclusion As in the proof of Lemma 3.2, we can use the "sprinkled" uniqueness estimate (62) to deduce a lower bound on the two-point function. The origin o is connected to another fixed vertex $x \in B_{n/40}$ by a $(p_0 + 2\delta)$-open path in B_n as soon as both o and x are $(p + \delta)$-connected to ∂B_n and the uniqueness event $U_{p_0+\delta,p_0+2\delta}(n/40, n/10)$ occurs. Using the Harris–FKG Inequality and the union bound, for every good scale $n \ge n_2$, we get

$$\begin{aligned}
\kappa_{p_0+2\delta}(n/\log^4 n, n) &\ge \kappa_{p_0+2\delta}(n/40, n) \\
&\ge \min_{x \in B_{n/40}} \mathbf{P}[\{o \xleftrightarrow[p+\delta]{} \partial B_n, x \xleftrightarrow[p+\delta]{} \partial B_n\} \cap U_{p_0+\delta,p_0+2\delta}(n/40, n/10)] \\
&\ge \theta(p_0 + \delta)^2 - e^{-\log^2 n}.
\end{aligned}$$

It follows that $\limsup_{n\to\infty} \kappa_{p_0+2\delta}(n/\log^4 n, n) > 0$, which concludes the proof. □

10 Proof of sharpness: coarse grains without rescaling

In this section, we prove Theorem 1.1 and Theorem 1.2. Both proofs involve the use of Proposition 1.3 as the building block of a Peierls-type argument.

On the hypercubic lattice $\mathbb{Z}^d$, a natural way to conclude from Proposition 1.3 is to use a standard renormalisation procedure: we look at blocks at a larger scale, and these blocks can be seen as vertices of a rescaled of $\mathbb{Z}^d$. This strategy relies on a self-similarity property of $\mathbb{Z}^d$ that does not hold for general groups. In order to overcome this difficulty, we will keep the definition of large blocks (sometimes called "coarse grains"), but we will replace the rescaling argument by the use of a k-independent[5] process with sufficiently

[5]Recall that a site percolation process X is *k-independent* if for any two sets of vertices V_1 and V_2 such that $\forall (v_1, v_2) \in V_1 \times V_2,\ d(v_1, v_2) > k$, the restrictions of X_{V_1} and X_{V_2} are independent.

high marginals. To each vertex v of G, we attach a block made of vertices at distance k around v and define a new process on the graph G without rescaling it. Contrary to the renormalisation approach, a given block overlaps $\simeq k^d$ other blocks, inducing some dependencies on the new process.

In order to apply this strategy, we rely on the following lemma, which ensures that a finite-range site percolation with sufficiently high marginals does not have large closed cutsets. Recall that a minimal cutset between o and ∞ is a subset $\Pi \subset V$ such that the connected component of o in $V \setminus \Pi$ is finite, and it is minimal for the inclusion among all sets satisfying this property. For every $n \geq 1$, write $\mathcal{P}_{\geq n}$ for the set of all minimal cutsets Π between o and ∞ that satisfy $|\Pi| \geq n$.

Lemma 10.1. *Let R be so that the conclusion of Lemma 2.1 holds. For every $k \geq R$, there exists $c = c(k, G) > 0$ such that the following holds. Let X be a $2k$-independent site percolation on G with marginals satisfying $\mathbb{P}[X(v) = 1] \geq 1 - \frac{1}{2e|B_{5k}|}$ for every $v \in V(G)$. Then, for every $n \geq 1$, we have*

$$\mathbb{P}[\exists \Pi \in \mathscr{P}_{\geq n},\ \Pi \text{ is closed in } X] \leq e^{-cn}.$$

Proof. Fix $k \geq R$ and set $D = |B_{5k}|$. Let $\Pi \in \mathscr{P}_{\geq n}$. Since R satisfies the conclusion of Lemma 2.1, the set Π is R-connected. Besides, by Lemma 2.2, Π intersects the ball of radius $R|\Pi|$.

Say that a set of vertices is r-separated if any two distinct vertices u, v of it satisfy $d(u, v) > r$. Let Π' be a $2k$-separated subset of Π that is maximal for inclusion among all such subsets. By maximality, any ball of radius $2k$ centred at a vertex of Π must contain a vertex of Π'. Therefore, the quantity $m := |\Pi'|$ is at least $|\Pi|/D$. Furthermore, we claim that Π' satisfies the following properties:

- Π' is $2k$-separated,
- $|\Pi'| = m$,
- Π' is $5k$-connected,
- Π' intersects B_{RDm+2k}.

The first two items follow from the definitions above. The third and fourth items follow from the properties of Π together with the hypothesis $k \geq R$ and the observation that Π' intersects any ball of radius $2k$ centred at a vertex of Π.

For every $m \geq 1$, write $\mathscr{P}'_m$ for the collection of all subsets of V satisfying the four properties listed above. By definition, all the elements of $\mathscr{P}'_m$ intersect B_{RDm+2k} and they are connected subsets of the graph $G' = (V, E')$ with $E' = \{\{x, y\} : d(x, y) \leq 5k\}$. Since the degree of G' is D, standard bounds on the number of connected subsets of a graph (see e.g. [Bol06, Problem 45]) give

$$|\mathscr{P}'_m| \leq |B_{RDm+2k}|(eD)^m.$$

From this bound, we get

$$\begin{aligned}\mathbb{P}[\exists \Pi \in \mathscr{P}_{\geq n},\ \Pi \text{ is closed in } X] &\leq \sum_{m \geq n/D} \mathrm{P}[\exists \Pi' \in \mathscr{P}'_m,\ \Pi' \text{ is closed in } X] \\ &\leq \sum_{m \geq n/D} |\mathscr{P}'_m| \cdot (2eD)^{-m} \qquad \text{by our assumptions on } X \\ &\leq \sum_{m \geq n/D} |B_{RDm+2k}| \cdot 2^{-m}.\end{aligned}$$

Since the volume of the ball B_{Dn} grows at most polynomially in n, the proof follows from the estimate above. □

Proof of Theorems 1.1 and 1.2 . Let $p > p_c(G)$ and let ω be a Bernoulli bond percolation of parameter p on G. By Proposition 1.3 and by polynomial growth of G, we can choose k such that

$$\mathrm{P}_p[B_{k/10} \leftrightarrow \partial B_k, U(k/5, k/2)] \geq 1 - \frac{1}{2e|B_{5k}|}.$$

From now on and until the end of the proof, we fix k as above. For each $v \in V(G)$, we will denote by $U_v = U_v(k/5, k/2)$ the event $U(k/5, k/2)$ centred at v instead of o. We define the site percolation process X on G by setting

$$X(v) = \begin{cases} 1 & \text{if both } B_{k/10}(v) \leftrightarrow \partial B_k(v) \text{ and } U_v \text{ occur,} \\ 0 & \text{otherwise.} \end{cases}$$

Notice that X is a $2k$-independent site percolation with marginals satisfying

$$\mathrm{P}_p[X(v) = 1] \geq 1 - \frac{1}{2e|B_{5k}|}.$$

Let us consider $n > |B_k|$. We denote by C_o the cluster of o in ω. Assume that $|C_o| \geq n$. In this case, the key observation about the process X is that having an infinite open path in X touching C_o entails the existence of a path between o and ∞ in ω. Let us explain this fact. Since $|C_o| \geq n > |B_k|$, for every $v \in C_o$, we have that $v \leftrightarrow \partial B_k(v)$ in C_o. Then, let us assume that there is an infinite self-avoiding path $v_1, v_2, \ldots$ with $v_1 \in C_o$ and such that $X(v_i) = 1$ for all $i \geq 1$. This implies that, for every $i \geq 1$, one can fix a cluster C_i in $\omega \cap B_k(v_i)$ that connects $B_{k/10}(v_i)$ to $\partial B_k(v_i)$. Notice that for every $i \geq 1$, C_i and C_{i+1} both connect $B_{k/5}(v_i)$ to $B_{k/2}(v_i)$. Since the event $U_{v_i}(k/5, k/2)$ holds for every i, we have that all the C_i are forced to be connected to each other in ω. If we choose C_1 to be the connected component of v_1 in $C_o \cap B_k(v_1)$, we have that o is connected to ∞ in ω. Since $\{o \leftrightarrow \partial B_n\} \subset \{|C_o| > n\}$, this observation implies that

$$\mathrm{P}_p[o \leftrightarrow \partial B_n, o \not\leftrightarrow \infty] \leq \mathrm{P}_p[C_o \overset{\not\leftrightarrow}{\underset{X}{}} \infty, n \leq \mathrm{diam}(C_o) < \infty]. \tag{63}$$

If the event on the right-hand side of (63) is satisfied, then, by Lemma 2.3, there exists an X-closed cutset Π disconnecting o and ∞ of diameter at least $n/2$. Since Π is R-connected, this yields $|\Pi| \geq \frac{n}{2R}$. One can thus use Lemma 10.1 to upper-bound the right-hand side of (63) by $\exp(-\frac{cn}{2R})$, for some constant $c > 0$. This concludes the proof of Theorem 1.1.

Regarding the proof of Theorem 1.2, we have, similarly to before, that for $n > |B_k|$,

$$\mathrm{P}_p[n \le |C_o| < \infty] = \mathrm{P}_p[C_o \overset{\not\longrightarrow}{\underset{X}{}} \infty, n \le |C_o| < \infty]. \tag{64}$$

In this case, if the event on the right-hand side of (64) is satisfied, then Lemma 2.4 yields an X-closed cutset disconnecting o and ∞ of size larger than $cn^{\frac{d-1}{d}}$. Hence, using Lemma 10.1, the right-hand side of (64) is upper bounded by $\exp(-c'n^{\frac{d-1}{d}})$ for some positive constant c'. This concludes the proof of Theorem 1.2. □

References

[AB87] Michael Aizenman and David Barsky. Sharpness of the phase transition in percolation models. *Communications in Mathematical Physics*, 108(3):489–526, 1987.

[AKN87] Michael Aizenman, Harry Kesten, and Charles Newman. Uniqueness of the infinite cluster and continuity of connectivity functions for short and long range percolation. *Communications in Mathematical Physics*, 111(4):505–531, 1987.

[AV08] Tonći Antunović and Ivan Veselić. Sharpness of the phase transition and exponential decay of the subcritical cluster size for percolation on quasi-transitive graphs. *Journal of Statistical Physics*, 130(5):983–1009, 2008.

[Bas72] Hyman Bass. The degree of polynomial growth of finitely generated nilpotent groups. *Proceedings of the London Mathematical Society*, 3(4):603–614, 1972.

[BB99] Eric Babson and Itai Benjamini. Cut sets and normed cohomology with applications to percolation. *Proceedings of the American Mathematical Society*, 127(2):589–597, 1999.

[BG18] Jérémie Brieussel and Antoine Gournay. Connectedness of spheres in Cayley graphs. *Algebra Discrete Mathematics*, 26(2):190–246, 2018.

[BH57] Simon Broadbent and John Hammersley. Percolation processes: I. Crystals and mazes. *Proceedings of the Cambridge Philosophical Society*, 53:629–641, 1957.

[BK89] Robert Burton and Michael Keane. Density and uniqueness in percolation. *Communications in Mathematical Physics*, 121(3):501–505, 1989.

[BKK+92] Jean Bourgain, Jeff Kahn, Gil Kalai, Yitzhak Katznelson, and Nathan Linial. The influence of variables in product spaces. *Israel Journal of Mathematics*, 77(1-2):55–64, 1992.

[BNP11] Itai Benjamini, Asaf Nachmias, and Yuval Peres. Is the critical percolation probability local? *Probability Theory and Related Fields*, 149(1-2):261–269, 2011.

[Bol06] Béla Bollobás. *The art of mathematics: Coffee time in Memphis*. Cambridge University Press, 2006.

[BR06] Béla Bollobás and Oliver Riordan. The critical probability for random Voronoi percolation in the plane is 1/2. *Probability Theory and Related Fields*, 136(3):417–468, 2006.

[BS96] Itai Benjamini and Oded Schramm. Percolation beyond $\mathbb{Z}^d$, many questions and a few answers. *Electronic Communications in Probability*, 1:71–82, 1996.

[BT17] Itai Benjamini and Vincent Tassion. Homogenization via sprinkling. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 53(2):997–1005, 2017.

[CCN87] Jennifer Chayes, Lincoln Chayes, and Charles Newman. Bernoulli percolation above threshold: an invasion percolation analysis. *The Annals of Probability*, 15(4):1272–1287, 1987.

[Cer15] Raphaël Cerf. A lower bound on the two-arms exponent for critical percolation on the lattice. *The Annals of Probability*, 43(5):2458–2480, 2015.

[CMT] Daniel Contreras, Sébastien Martineau, and Vincent Tassion. Locality of percolation for graphs with polynomial growth. ArXiv:2205.10253.

[DGR20] Hugo Duminil-Copin, Subhajit Goswami, and Aran Raoufi. Exponential decay of truncated correlations for the Ising model in any dimension for all but the critical temperature. *Communications in Mathematical Physics*, 374(2):891–921, 2020.

[DGR+20] Hugo Duminil-Copin, Subhajit Goswami, Aran Raoufi, Franco Severo, and Ariel Yadin. Existence of phase transition for percolation using the Gaussian Free Field. *Duke Mathematical Journal*, 169(18):3539–3563, 2020.

[DGRS] Hugo Duminil-Copin, Subhajit Goswami, Pierre-François Rodriguez, and Franco Severo. Equality of critical parameters for percolation of Gaussian free field level-sets. ArXiv:2002.07735.

[Die17] Reinhard Diestel. *Graph Theory*. Graduate Texts in Mathematics 173, 5th edition, 2017.

[DK17] Cornelia Druțu and Michael Kapovich. *Geometric group theory*. American Mathematical Society Colloquium Publications 63, 2017.

[DKT] Hugo Duminil-Copin, Gady Kozma, and Vincent Tassion. Upper bounds on the percolation correlation length. *Special volume in memory of Vladas Sidoravicius published, Birkhauser series "Progress in Probability" (to appear)*. ArXiv:1902.03207.

[dLH00] Pierre de La Harpe. *Topics in geometric group theory.* University of Chicago Press, 2000.

[DRT18] Hugo Duminil-Copin, Aran Raoufi, and Vincent Tassion. A new computation of the critical point for the planar random-cluster model with $q \geq 1$. *Annales de l'Institut Henri Poincaré Probabilités et Statistiques*, 54(1):422–436, 2018.

[DT16] Hugo Duminil-Copin and Vincent Tassion. A new proof of the sharpness of the phase transition for Bernoulli percolation and the Ising model. *Communications in Mathematical Physics*, 343(2):725–745, 2016.

[FGO15] Louis Funar, Martha Giannoudovardi, and Daniele Otera. On groups with linear sci growth. *Fundamenta Mathematicae*, 228:47–62, 2015.

[GGR88] Alberto Gandolfi, Geoffrey Grimmett, and Lucio Russo. On the uniqueness of the infinite cluster in the percolation model. *Communications in Mathematical Physics*, 114(4):549–552, 1988.

[GM90] Geoffrey Grimmett and John Marstrand. The supercritical phase of percolation is well behaved. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 430(1879):439–457, 1990.

[Gri99] Geoffrey Grimmett. *Percolation.* Grundlehren der Mathematischen Wissenschaften 321, 2nd edition, 1999.

[Gri06] Geoffrey Grimmett. *The random-cluster model.* Grundlehren der Mathematischen Wissenschaften 333, 2006.

[Gro81] Mikhaïl Gromov. Groups of polynomial growth and expanding maps (with an appendix by Jacques Tits). *Publications Mathématiques de l'IHÉS*, 53:53–78, 1981.

[Gui73] Yves Guivarc'h. Croissance polynomiale et périodes des fonctions harmoniques. *Bulletin de la société mathématique de France*, 101:333–379, 1973.

[HH21] Jonathan Hermon and Tom Hutchcroft. Supercritical percolation on non-amenable graphs: Isoperimetry, analyticity, and exponential decay of the cluster size distribution. *Inventiones mathematicae*, 224(2):445–486, 2021.

[HT] Tom Hutchcroft and Matthew Tointon. Non-triviality of the phase transition for percolation on finite transitive graphs. ArXiv:2104.05607.

[Hut20] Tom Hutchcroft. Locality of the critical probability for transitive graphs of exponential growth. *The Annals of Probability*, 48(3):1352–1371, 2020.

[Kes80] Harry Kesten. The critical probability of bond percolation on the square lattice equals $\frac{1}{2}$. *Communications in Mathematical Physics*, 74(1):41–59, 1980.

[KKL88] Jeff Kahn, Gil Kalai, and Nathan Linial. The influence of variables on Boolean functions. In *29th Annual Symposium on Foundations of Computer Science*, pages 68–80, 1988.

[Kle10] Bruce Kleiner. A new proof of Gromov's theorem on groups of polynomial growth. *Journal of the American Mathematical Society*, 23(3):815–829, 2010.

[KZ90] Harry Kesten and Yu Zhang. The probability of a large finite cluster in supercritical Bernoulli percolation. *The Annals of Probability*, 18(2):537–555, 1990.

[LMS08] Russell Lyons, Benjamin Morris, and Oded Schramm. Ends in uniform spanning forests. *Electronic Journal of Probability*, 13:1702–1725, 2008.

[LP16] Russell Lyons and Yuval Peres. *Probability on trees and networks*, volume 42 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, New York, 2016.

[Men86] Mikhail Menshikov. Coincidence of critical points in percolation problems. *Soviet Mathematics Doklady*, 33:856–859, 1986.

[MT17] Sébastien Martineau and Vincent Tassion. Locality of percolation for Abelian Cayley graphs. *The Annals of Probability*, 45(2):1247–1277, 2017.

[Pet08] Gábor Pete. A note on percolation on $\mathbb{Z}^d$: isoperimetric profile via exponential cluster repulsion. *Electronic Communications in Probability*, 13:377–392, 2008.

[Sev] Franco Severo. Sharp phase transition for Gaussian percolation in all dimensions. ArXiv:2105.05219.

[ST10] Yehuda Shalom and Terence Tao. A finitary version of Gromov's polynomial growth theorem. *Geometric and Functional Analysis*, 20(6):1502–1547, 2010.

[SW90] Paolo Soardi and Wolfgang Woess. Amenability, unimodularity, and the spectral radius of random walks on infinite graphs. *Mathematische Zeitschrift*, 205(3):471–486, 1990.

[Tal94] Michel Talagrand. On Russo's approximate zero-one law. *The Annals of Probability*, 22(3):1576–1587, 1994.

[Tim07] Ádám Timár. Cutsets in infinite graphs. *Combinatorics, Probability and Computing*, 16(1):159–166, 2007.

[Tim13] Ádám Timár. Boundary-connectivity via graph theory. *Proceedings of the American Mathematical Society*, 141(2):475–480, 2013.

[Tro85] Vladimir Trofimov. Graphs with polynomial growth. *Mathematics of the USSR-Sbornik*, 51(2):405, 1985.

[TT21] Romain Tessera and Matthew Tointon. A finitary structure theorem for vertex-transitive graphs of polynomial growth. *Combinatorica*, 41:263–298, 2021.