Percolation on Finite Cayley Graphs

Christopher Malon* and **Igor Pak**†

MIT Department of Mathematics,
Cambridge, MA 02139, USA
malon@math.mit.edu
pak@math.mit.edu

January 25, 2005

Abstract

In this paper, we study percolation on finite Cayley graphs. A conjecture of Benjamini says that the critical percolation p_c of any vertex–transitive graph satisfying a certain diameter condition can be bounded away from one. We prove Benjamini's conjecture for some special classes of Cayley graphs. We also establish a reduction theorem, which allows us to build Cayley graphs for large groups without increasing p_c.

Introduction

Percolation on finite graphs is a new subject with a classical flavor. It arose from two important and, until recently, largely independent areas of research: Percolation Theory and Random Graph Theory. The first is a classical Bernoulli percolation on a lattice, initiated as a mathematical subject by Hammersley and Morton in the 1950's, and which became a major area of research. A fundamental albeit elementary observation that the critical percolation p_c is bounded away from 1 on $\mathbb{Z}^2$ has led to a number of advanced results and quests for generalizations. Among those most relevant to this work, let us mention the Grimmett Theorem regarding the 'smallest' possible region under a graph in $\mathbb{Z}^2$ for which one still has $p_c < 1$. Similarly, percolation in finite boxes has become crucially important as a source of new questions, as well as a tool (see [13] for references and major results in the area.)

In the past decade, much attention within the subject of percolation has been devoted to the study of percolation on Cayley graphs, and, more generally, vertex–transitive graphs. A series of conjectures by Benjamini and Schramm [6] would predict an interplay of Probability Theory and Group Theory in which

*NSF Graduate Research Fellow

†Partially supported by the NSF and the NSA

the probabilistic properties of the (bond or site) percolation depend heavily on the algebraic properties of an underlying (infinite) group, but not on a particular generating set. We refer to [7] for a description of recent progress in this subject.

Motivated by the study of percolation on infinite Cayley graphs, Benjamini in [5] (see also [2]) extends the notion of critical probability to finite graphs by asking at which point the resulting graph has a large (constant proportion size) connected component. He conjectured that one can prove a new version of $p_c < 1-\varepsilon$, under a weak diameter condition. (Here and everywhere in the introduction, $\varepsilon > 0$ is a universal constant independent of the size of the graph.) In this paper we present a number of positive results toward this unexpected, and, perhaps, overly optimistic conjecture.

Our main results are of two types. First, we concentrate on special classes of groups and establish $p_c < 1-\varepsilon$ for these. We prove Benjamini's conjecture for all abelian groups with Hall bases as generating sets. We also prove that $p_c < 1-\varepsilon$ for Cayley graphs whose generating sets have enough short disjoint relations, a notion somewhat similar to that in [4]. Our most important, and technically most difficult result is the Reduction Theorem, which enables us in certain cases to obtain sharp bounds for p_c of a Cayley graph of a group G depending on those of a normal subgroup $H \lhd G$ and a quotient group G/H. While the full version of Benjamini's conjecture remains wide open, the Reduction Theorem allows us in certain cases to concentrate on finite simple groups (a sentiment expressed in [5]). By means of the classification of finite simple groups [12], and a recent series of probabilistic results relying on it (see e.g. [18]), one can hope that our results will lead to further progress towards understanding percolation on finite Cayley graphs.

Our Reduction Theorem requires that the index $[G : H]$ not be too large in relation to $|H|$. In the case where H has a complement K and G is the semidirect product $G = H \rtimes K$, this condition can be dropped, and we simply require that both $|H|$ and $|K|$ exceed some constant. Theorem 14 describes this situation.

Let us also describe a connection to Random Graph Theory. The pioneer paper [11] of Erdős and Rényi considered random graphs either as random subgraphs of a complete graph K_n, or as a result of a random graph process, in which edges are added one at a time. We use only the first model here. Although one needs the probability p of an edge to be roughly $\log n/n$ for the graph to become connected, a much smaller value $p = (1+\varepsilon)/n$ suffices for the creation of a 'giant' ($c(\varepsilon)n$ size) connected component.

The work of Erdős and Rényi led to the study of properties of random graphs, and more recently, of random subgraphs of finite graphs (see e.g. [1, 9, 14]) In the past years, connectivity and Hamiltonicity have remained the most studied properties, ever since the celebrated Margulis' Lemma, rediscovered later by Russo (see e.g. [15, 17].) One can view our work as a new treatment of the existence of a giant component in a large class of vertex–transitive graphs.

Percolation on Cayley graphs seems to resemble percolation on a more general class of vertex–transitive graphs. For infinite groups, this can be partially explained by the fact that percolation properties such as $p_c < 1-\varepsilon$ are invariant under quasi–isometry [6]. In fact, it remains an open problem whether all vertex–transitive graphs are quasi–isometric to Cayley graphs; the only potential counterexample was proposed in [10]. In this paper we restrict ourselves to Cayley graphs, as their rich group theoretic structure allows a combination of techniques to be applied.

A few words about notation: For the rest of the paper, G always will denote a finite group, Γ will denote a finite graph, and $|\Gamma|$ will denote the number of vertices in Γ. The symbol log denotes the logarithm base 2. As in [13], we sometimes write real–valued quantities in places where integers are required, in order to avoid extra notation.

1 Definitions and main results

Recall from [13] the definition of percolation on a lattice. Let L^d be the integer lattice in d dimensions, with $\mathbb{Z}^d$ its vertices and $E^d = \{((x_1, \ldots, x_d), (y_1, \ldots, y_d)) :$ for some i, $|x_i - y_i| = 1$ and for $j \neq i$, $x_j = y_j\}$ its edges. Consider the probability space with outcomes $\Omega = \prod_{e \in E^d} \{0, 1\}$ and whose measurable sets are the elements of the smallest σ–field in which the state of any finite set of edges can be tested. If $\omega \in \Omega$, we say that an edge e *remains* (or is *open*) in the outcome ω if $\omega(e) = 1$, and that e is *deleted* (or is *closed*) otherwise. Let μ_e be the Bernoulli measure on the edge e in which e remains with probability p. The product measure of the μ_e gives a measure on the probability space, which we call p–percolation.

Let Γ be a finite graph. We write its set of edges as $E(\Gamma)$, and its set of vertices (by abuse of notation) as Γ. In a *p–percolation* process on Γ, every edge $e \in E(\Gamma)$ is deleted with probability $1-p$, independently. Such a process defines a probability distribution on subgraphs of Γ, in which each subgraph $H \subset \Gamma$ is assigned the probability $p^{|E(H)|}(1-p)^{|E(\Gamma)|-|E(H)|}$, where $|\cdot|$ denotes the cardinality of a set. Later we informally refer to edges of H as 'p–percolated'.

For constants ρ, α, and p between zero and one, we let $\mathcal{L}(\rho, \alpha, p)$ denote the collection of finite graphs Γ, such that a random subgraph $H \subset \Gamma$ as above will have a connected component joining $\rho|\Gamma|$ of their vertices, with probability at least α.

Let ρ and α be fixed, and let Γ be a finite graph. Define the *critical probability* $p_c(\Gamma)$ as follows:

$$p_c(\Gamma) = p_c(\Gamma; \rho, \alpha) := \inf\big\{p : \Gamma \in \mathcal{L}(\rho, \alpha, p)\big\}.$$

From monotonicity of the percolation, $\Gamma \in \mathcal{L}(\rho, \alpha, p)$ for all $1 \geq p > p_c(\Gamma)$.

We are interested in conditions which bound the critical probability away from 1, as the size of graph Γ grows. Benjamini conjectured in [5]:

Conjecture 1. (Benjamini) *If Γ is a vertex–transitive graph with n vertices, and* $\operatorname{diam}(\Gamma) < n/\log n$*, then* $p_c(\Gamma; \rho, \alpha) < 1 - \varepsilon(\rho, \alpha)$*.*

As mentioned in the introduction, Cayley graphs are important examples of vertex–transitive graphs. From this point on, we consider only finite Cayley graphs.

Let G be a finite group and let $S = S^{-1}$ be a symmetric set of generators. A graph with vertices $g \in G$ and edges $(g, g \cdot s)$, $s \in S$ is called the *Cayley graph* $\Gamma(G, S)$ of the group G with generating set S.

Definition 2. *Suppose $s_1, \ldots, s_n$ are generators of a finite abelian group G, and let a_i be the order of s_i. We say that $s_1, \ldots, s_n$ is a* Hall basis for G *if the products $s_1^{i_1} \cdots s_n^{i_n}$ are distinct for all n–tuples $(i_1, \ldots, i_n)$, where $0 \le i_k < a_k$.*

The following result establishes Benjamini's conjecture for all Cayley graphs of abelian groups whose generating sets are Hall bases:

Theorem 3. *For any constants ρ and α between 0 and 1, there is a constant $\varepsilon = \varepsilon(\rho, \alpha) > 0$, such that for every Cayley graph $\Gamma = \Gamma(G, S)$ of any finite abelian group G and Hall basis S satisfying $\operatorname{diam}(\Gamma) < \frac{|G|}{\log |G|}$, we have $p_c(\Gamma; \rho, \alpha) < 1 - \varepsilon$.*

If the number of commuting generators is large in proportion to the diameter of the graph, for each Cayley graph in a collection, we again can bound the critical probability away from one. Precisely, let $\Gamma_n = \Gamma(G_n, R_n)$ be a sequence of Cayley graphs with diameters $d_n = \operatorname{diam}(\Gamma_n)$. For each $s \in R_n$, let $T_n(s) = \{r \in R_n : [r, s] = 1\}$.

Theorem 4. *If $d_n \to \infty$ as $n \to \infty$, and each $|T_n(s)| \ge 4 \log d_n$, then there exists $\varepsilon > 0$ such that $p_c(\Gamma(G_n, R_n); \frac{2}{3}, \frac{1}{2}) \le 1 - \varepsilon$ for all n.*

Examples satisfying the conditions of Theorem 4 are given in sections 3 and 8.

Without information about the structure or the critical probability of G/H, it still may be possible to bound the critical probability of G if the index of H in G is not too large.

Theorem 5. (Reduction Theorem) *Let $\Gamma = \Gamma(G, S)$ be a Cayley graph of a finite group, let $H \triangleleft G$ be a normal subgroup, and let ρ and α be positive constants with $\rho, \alpha < 1$ and $\rho > \frac{1}{2}$. Suppose that $R = H \cap S$ generates H, and write $p_c = p_c(\Gamma(H, R); \rho, \alpha)$ for the critical percolation of the Cayley graph of this subgroup. Suppose $p > \max(\frac{1}{\sqrt{2}}, p_c)$. There exist constants $\beta = \beta(\rho) < 1$, $\eta = \eta(\alpha)$, and $N = N(\rho, \alpha)$, so that if $\alpha > \beta$ and $[G : H] > N$, and*

$$\big(\ln [G : H] + \eta\big) [G : H] \le (2\rho - 1)\, |H| \tag{1}$$

we have $p_c(\Gamma(G, S); \rho, \alpha) \le p$.

The Reduction Theorem can be applied iteratively to groups with a composition series. Suppose we have

$$\{1\} = G_0 \triangleleft G_1 \triangleleft \ldots \triangleleft G_\ell,$$

and for each $i > 0$, equation (1) is satisfied for $G = G_{i+1}$ and $H = G_i$. If we have generating sets $S_i \subset G_i$ with $S_i \subset S_{i+1}$ for all i, then we may bound $p_c(\Gamma(G_\ell, S_\ell); \rho, \alpha) \le p_c(\Gamma(G_1, S_1); \rho, \alpha)$, under the assumptions on ρ and α in the Theorem, supposing $p_c(\Gamma(G_1, S_1)); \rho, \alpha) > \frac{\sqrt{2}}{2}$. See Section 8 for an example of such an application.

We prove these theorems in the sections that follow, and conclude with a few examples and open problems.

2 Basic Results

Large components in finite graphs are the analogues of infinite clusters in infinite graphs. The Benjamini conjecture appears to be inspired by Grimmett's Theorem (see, *e.g.*, [13], pages 304–309), which guarantees the existence of infinite clusters in certain subsets of the square lattice.

Theorem 6. (Grimmett) *Let f be a function so that $\frac{f(x)}{\log x} \to a$ as $x \to \infty$, for some positive constant a. Let $G(f)$ denote the region in the positive quadrant of the square lattice under the function $f(x)$. There exists $p < 1$ so that this region has an infinite component after p–percolation almost surely.*

The following lemma is a close version, though not a direct corollary, of the theorem. We will prove it, and use the lemma in our proof of Theorem 3.

Lemma 7. *Let Γ be an $m \times n$ box within the square grid, and let $\rho, \alpha < 1$ and $a \in \mathbb{R}^{>0}$ be constants. Then there exists $\varepsilon = \varepsilon(\rho, \alpha, a) > 0$ such that if $n \ge m > a \log n$, we have $p_c(\Gamma; \rho, \alpha) < 1 - \varepsilon$.*

The following counting lemma provides one tool with which to bound the critical probability of a vertex transitive graph. It is used in the proof of Theorem 4.

Proposition 8. *Let Γ be a vertex transitive graph undergoing p–percolation. Distinguish a vertex z. Suppose that there are constants $0 < \tau, \rho < 1$ such that for every vertex $v \in \Gamma$, the probability that z lies in the same connected component as v after percolation is at least $\tau + \rho - \tau\rho$. Then the probability that z belongs to a configuration of size at least $\rho|\Gamma|$ is at least τ.*

Proof: We prove the contrapositive: If the probability that z is in a component of size smaller than $\rho|\Gamma|$ is at least $1 - \tau$, then there exists a vertex x whose probability of being in a different component than z is at least $1 - \tau - \rho + \tau\rho$.

For each vertex $v \in \Gamma$, let $m(v)$ denote the probability that v is connected to z after percolation. Then

$$\sum_{v \in \Gamma} m(v) \le \tau|\Gamma| + (1 - \tau)\rho|\Gamma| \tag{2}$$

Indeed, even if all the graphs with ρ–size connected component were entirely connected, they would not contribute more than $\tau|\Gamma|$ to the sum, because such graphs occur with probability no more than τ. This gives the first term. The remaining graphs contribute to $m(v)$ for no more than ρ fraction of the vertices. This gives the second term. Therefore, some vertex v must have $m(v) \leq \tau + \rho - \tau\rho$. $\square$

Example 9. Let $\rho = 2/3$ and $\tau = 1/2$ in Proposition 8, and $\Gamma = \Gamma(G, S)$ be a Cayley graph undergoing p–percolation. Distinguish a vertex $x \in \Gamma$. If every $g \in \Gamma$ is connected to the identity with probability at least 5/6, then the probability that x belongs to a configuration of size at least $(2/3)\Gamma$ is at least 1/2. We use these special values to simplify the calculations that follow.

We conclude this section with the following well–known bound, which we will use repeatedly throughout what follows.

Theorem 10 (Chernoff). *(See, e.g., [8].) Let X_i, $i = 1, \ldots, n_0$, be independent Poisson trials, with outcomes* 1 *and* 0 *with probabilities p_0 and $1 - p_0$ respectively. Set $X = \sum_{i=1}^{n_0} X_i$ and $\mu_0 = \mathrm{E}[X] = n_0 p_0$. Then for every $\delta_0 > 0$, the following bound holds:*

$$\Pr(X < (1 - \delta_0)\mu_0) < e^{-\frac{\mu_0 \delta_0^2}{2}}$$

3 Commuting Generators

In this section, we prove Theorem 12, which generalizes Theorem 4 from the introduction. The following example illustrates our technique in a particularly simple case.

Let $\Gamma = \Gamma(S_n, R_n)$ be the Cayley graph for the symmetric group, with $R_n = \{(1,\ 2), (2,\ 3), \ldots, (n-1,\ n)\}$ the Coxeter transpositions. We may bound the critical probability of this Cayley graph using an idea that applies to any sequence of groups with enough generators and short disjoint relations.

Proposition 11. *There exists $\varepsilon > 0$ such that for all n, $p_c(\Gamma(S_n, R_n); \frac{2}{3}, \frac{1}{2}) \leq 1 - \varepsilon$.*

Proof: By our example following Proposition 8, it suffices to show that every element $g \in S_n$ remains connected to the identity 1 with probability at least 5/6.

Let d be the diameter of $\Gamma(S_n, R_n)$; we have $d = \binom{n}{2}$. Fix a path from 1 to g of length no more than d. Some edges of this path may be deleted by percolation.

Let us consider how to get around a deleted edge. Say the deleted edge joins a vertex x to $(i,\ i+1)x$. Observe that there are at least $n-4$ generators of the form $(j,\ j+1)$ that commute with $(i,\ i+1)$. Any of these generators allows us to replace the edge from x to $(i,\ i+1)x$ by the three–edge sequence from x

given by the word $(j,\ j+1)(i,\ i+1)(j,\ j+1)$. Under p–percolation, each such three–edge detour is unbroken with probability p^3, and since they are disjoint from each other, the probability that all $n-4$ detours break is $(1-p^3)^{n-4}$.

Even if every edge of the original path from 1 to g is deleted, we can find unbroken detours in this way around all the deleted edges with probability at least $1-d(1-p^3)^{n-4}$. Therefore, if

$$d(1-p^3)^{n-4} < \frac{1}{6},$$

the proposition is proven. The left hand side goes to zero as n goes to infinity for every $p = 1-\varepsilon$. □

We can generalize this result for other sequences of Cayley graphs as follows.

Let G be a finite group and R be a set of generators. Consider a relation of the form $r = s_1 \cdots s_m$, where $s_i \in R$. We say that its *length* is m. Two relations $r = s_1 \cdots s_m$ and $r = t_1 \cdots t_n$ are *disjoint* if, viewed as paths around the edge from e to r, they share no edges.

Theorem 12. *Let $\Gamma_n = \Gamma(G_n, R_n)$ be a sequence of Cayley graphs with diameters $d_n = \operatorname{diam}(\Gamma_n)$. Suppose that $d_n \to \infty$, and that there is a constant C such that for all n and all $s \in R_n$, there are at least $2\log d_n$ disjoint relations for s, each having length no more than C. There exists $\varepsilon > 0$ such that for all n, $p_c(\Gamma(G_n, R_n); \frac{2}{3}, \frac{1}{2}) \le 1-\varepsilon$.*

Proof: As above, we count disjoint detours around an edge $\{a, b\} \in \Gamma$. For simplicity, we may assume $a = 1$ so that $b \in S$.

For each relation $b = s_1 \cdots s_n$, we consider the detour that replaces the edge $\{1, b\}$ with the edges $\{1, s_1\}$, $\{s_1, s_1 s_2\}$, …, $\{s_1 s_2 \cdots s_{n-1}, b\}$, and apply Proposition 8 to obtain our result. Consider a path of length at most d_n from 1 to x. Apply Theorem 10 to a p–percolation process on this path, with $\delta_0 = \frac{1}{2}$ and $\mu_0 = pd_n$. With probability $1 - e^{-\frac{pd_n}{8}}$, at most δd_n edges of the path are deleted, where $\delta = 1 - \frac{p}{2}$. For each of these deleted edges $\{a, ar\}$, we have constructed at least $2\log d_n$ disjoint detours of C edges. The probability that all of these are broken is no more than $(1-p^C)^{2\log d_n}$. Thus, the total probability we cannot patch the path from 1 to x with our detours is no more than

$$e^{-\frac{pd_n}{8}} + \delta d_n (1-p^C)^{2\log d_n} \tag{3}$$

If p satisfies $p^C > \frac{1}{2}$, then $\delta d_n (1-p^C)^{2\log d_n} < \frac{\delta d_n}{d_n^2}$. Since $d_n \to \infty$ as $n \to \infty$, we have $e^{-\frac{pd_n}{8}} \to 0$. Increase p so that $\Gamma(G_n, R_n)$ has a large component in each of the finitely many graphs where the expression (3) is greater than $\frac{1}{6}$. □

Proof of Theorem 4 Take a maximal subset $T_n'(s) \subset T_n(s)$ so that $r \in T_n(s) \Rightarrow rs \notin T_n(s)$. Then $|T_n'(s)| \ge \frac{1}{2}|T_n(s)|$, and the commutation relations between s and the elements of $T_n'(s)$ are disjoint. Each commutation relation has length $C = 3$. Now apply Theorem 12. □

4 An Intersection Lemma

We will need the following lemma in the proof of the Reduction Theorem.

Let $\Gamma = \Gamma(H, R)$ be the (unpercolated) Cayley graph of a group H with generating set R. We call a subset of H *connected* if the induced subgraph of the corresponding set of vertices in Γ is connected. Let $\mathcal{A}$ be the set of connected subsets of H having cardinality exactly $\rho|H|$. Let μ_0 be the probability that Γ has a $\rho|H|$–sized component after p–percolation. If such a large component exists, choose, uniformly at random, a subset of $\rho|H|$ vertices of H that is connected after percolation, and call it A.

For $X \in \mathcal{A}$, let μ_X be the probability that when $\Gamma(H, R)$ is percolated, a $\rho|H|$–sized component exists and $A = X$. Then $X \to \frac{\mu_X}{\mu_0}$ defines a probability distribution on $\mathcal{A}$.

Lemma 13. *Let X be any fixed subset of H, and $0 < \gamma < \rho$. Then*

$$\mathrm{Pr}_{Y \in \mathcal{A}}(|X \cap Y| \geq \gamma|X|) \geq 1 - \eta \quad \text{where} \quad \eta = \frac{1-\rho}{1-\gamma}.$$

Proof: We say that two elements $Y, Y' \in \mathcal{A}$ are *equivalent* if $Y = Y'x$ for some $x \in H$, and write $\mathcal{A}/H$ for the set of equivalence classes. Because Cayley graphs are vertex transitive, for any $A \in \mathcal{A}$ and $g \in H$ we have $\mathrm{Pr}_Y(A = Y) = \mathrm{Pr}_Y(A = Yg)$. Consequently,

$$\begin{aligned}
\mathrm{Pr}_Y\big(|X \cap Y| = n\big) &= \sum_{\tilde{Y'} \in \mathcal{A}/H} \mathrm{Pr}_Y\big(|X \cap Y| = n \,|\, Y \in \tilde{Y'}\big) \cdot \mathrm{Pr}_Y\big(Y \in \tilde{Y'}\big) \\
&= \sum_{\tilde{Y'} \in \mathcal{A}/H} \mathrm{Pr}_{g \in H}\big(|X \cap Y'g| = n\big) \cdot \mathrm{Pr}_Y\big(Y \in \tilde{Y'}\big).
\end{aligned}$$

Here, Y' denotes any representative in $\mathcal{A}$ of the equivalence class $\tilde{Y'}$. Therefore, to show that $\mathrm{Pr}_Y(|X \cap Y| \geq \gamma|X|) \geq 1 - \eta$, it suffices to show for all fixed $Y \in \mathcal{A}$ that $\mathrm{Pr}_{g \in H}(|X \cap Yg| \geq \gamma|X|) \geq 1 - \eta$.

Fix $Y \in \mathcal{A}$. We have

$$\sum_{g \in H} |X \cap Yg| = |X||Y|. \tag{4}$$

Let η be the fraction of $g \in H$ for which $|X \cap Yg| < \gamma|X|$. Substituting this condition into equation 4 for these values of g, and $|X \cap Yg| \leq |X|$ for the other values of g, we obtain

$$\gamma|X| \cdot \eta|H| + |X| \cdot (1 - \eta)|H| \geq |X|\,|Y|. \tag{5}$$

Using $|Y| = \rho|H|$, equation 5 becomes

$$\gamma\eta + 1 - \eta \geq \rho \tag{6}$$

which shows

$$\eta \leq \frac{1-\rho}{1-\gamma} \tag{7}$$

as desired. This proves Lemma 13. $\square$

5 Proof of Reduction Theorem

Consider any p satisfying the hypothesis of the Theorem. Our analysis consists of three steps, in which we demonstrate that:

1. With probability $1-\epsilon_1$, there exist at least $\alpha^2\,[G:H]$ cosets with a connected component of size at least $\rho|H|$. In this event, we say that step 1 succeeded, and the cosets with the large component are called *good cosets*. Using the outcome of this step, another random process defines sets of $\rho|H|$ vertices within *every* coset, which we call the *good part* of the coset. The complementary subset in Hg_i is called the *bad part*.

2. Suppose step 1 succeeded. We show that the good parts of all the good cosets are connected to each other, with probability $1-\epsilon_2$. In this event, we say that step 2 succeeded, and observe that Γ will have a connected component of size at least $\rho\alpha^2|G|$.

3. Suppose step 1 and 2 succeeded. We show that more than $(\rho-\rho\alpha^2)|G|$ vertices that are in the bad part of some coset are attached by an edge to the good part of some good coset, with probability $1-\epsilon_3$. In this event, we say that step 3 succeeded, and observe that Γ has been p–percolated with a connected component of size $\rho|G|$ remaining.

In the sequel, we divide the p–percolation process into percolation on edges within the same coset, which we address in step 1, and percolation on edges between distinct cosets, which we address in steps 2 and 3. These percolations are independent. Therefore, we regard percolation on edges between distinct cosets as occurring "after" the definition of good cosets in step 1.

The theorem will follow once we compute values of β, η, and N as in the statement of the theorem that guarantee that $\epsilon_1+\epsilon_2+\epsilon_3<1-\alpha$.

Step 1. Let $n=[G:H]$, and write $G/H=\{Hg_i\}_{i=1}^n$. Consider each coset as a subgraph of Γ. Prior to percolation, each coset is isomorphic as a graph to $\Gamma(H,R)$. Because H is a normal subgroup, the only edges in Γ that join two vertices of a single coset in Γ come from generators in R. Our assumption implies that after the edges of each coset are percolated, each coset has a connected component of size at least $\rho|H|$, with probability at least α. Moreover, the occurrences of these large components are mutually independent. Call a coset with such a large component a "good coset." Applying the Chernoff bound with $p_0=\alpha$, $n_0=n$, and $\delta_0=1-\alpha$, we find that with probability $1-e^{-\frac{\alpha(1-\alpha)^2n}{2}}$, there are at least $\alpha^2 n$ good cosets. Thus $\epsilon_1=e^{-\frac{\alpha(1-\alpha)^2n}{2}}$.

Let $\mathcal{A}$ be the set of connected subsets of H (before percolation) having cardinality exactly $\rho|H|$. For every good coset Hg_i, choose a $\rho|H|$–size subset of the vertices of Hg_i that is connected after percolation uniformly at random, and call it A_i. The assumption that Hg_i is good ensures that at least one such choice can be made.

Let $g_1=1$ be the identity element, and let μ_0 be the probability that the identity coset is good. For $X\in\mathcal{A}$, let μ_X be the probability that when $\Gamma(G,S)$

is percolated, the identity coset H is good and $A_1 = X$. Then $X \to \frac{\mu_X}{\mu_0}$ defines a probability distribution on $\mathcal{A}$. For every coset Hg_i that is not good, select any subset A_i from $\mathcal{A} \cdot g_i$ according to this distribution. Thus, we have selected a set of vertices A_i of cardinality $\rho|H|$ for every coset of G/H.

In either case, let the set B_i be the complement of A_i in Hg_i. We will refer to A_i as the *good part* of Hg_i. For $X \in \mathcal{A}$, we have $Pr(A_i = X|Hg_i \text{ is good}) = Pr(A_i = X|Hg_i \text{ is not good})$.

Step 2. Let *pi* be the natural map $G \to G/H$. Fix a spanning tree on the (unpercolated) Cayley graph $\Gamma(G/H, \pi(S))$, and choose the root to be a good coset. Although the parent of a good coset need not be a good coset, we can take parents recursively until we reach one that is good. We call this the *good parent* of the given coset. We claim that with probability at least $1 - \frac{1}{n}$, every good parent is no more than $m = 2\log n$ levels higher in the tree. Indeed, the probability that any particular coset is good is at least $\alpha > \frac{1}{2}$. Enumerate the vertices of $\Gamma(G/H, \pi(S))$ as HG_j, for $j = 1, \ldots, n$. For a fixed Hg_j, the probability that Hg_j has its good parent more than m levels higher in the tree is no more than $(1-\alpha)^m < \frac{1}{2^m} = \frac{1}{n^2}$. (If the depth of Hg_j is less than m, then the tree's root, if nothing closer, is its good parent.) Thus, the probability that there exists a coset Hg_j whose good parent is too high in the tree, is no more than $\sum_{j=1}^{n} \frac{1}{n^2} = \frac{1}{n}$. So, suppose that every Hg_j has its good parent no more than m levels up.

The good parts of all the good cosets will be connected after percolation if each one remains connected to that of its good parent. Suppose Hg_j is the good parent of Hg_i. Then $g_j = g_i s_{i_1} \cdots s_{i_r}$ for some string of generators $s_{i_1}, \ldots, s_{i_r} \in S$ where $Hg_i s_{i_1}$ is the parent of Hg_i, *etc.* We have $r \leq m$. Right multiplication by $s_{i_1} \cdots s_{i_r}$ gives a bijection from Hg_i to Hg_j. By the inclusion–exclusion principle, at least $(2\rho - 1)|H|$ good points of Hg_i hit the good part of Hg_j. Therefore, in order for the good part of Hg_i to fail to be connected to the good part of Hg_j, we would need each of the $(2\rho-1)|H|$ paths of the form $x, xs_{i_1}, \ldots, xs_{i_1} \cdots s_{i_r}$ to break. Since these paths are all disjoint, the probability that they all break is no more than $(1-p^r)^{(2\rho-1)|H|}$. Let P_1 be the probability that some good coset fails to have its good part connected to that of its good parent. Then

$$
\begin{aligned}
P_1 &\leq n(1-p^m)^{(2\rho-1)|H|} \\
&\leq n(e^{-p^m})^{(2\rho-1)|H|} \\
&\leq n(e^{-p^{2\log n}})^{(2\rho-1)|H|} \\
&\leq n(e^{-(2\rho-1)|H|p^{2\log n}}) \\
&\leq ne^{-\frac{(2\rho-1)|H|}{n}}
\end{aligned}
$$

where the last inequality applies the hypothesis $p > \frac{1}{\sqrt{2}}$. If we take $\eta(\alpha) = \ln \frac{5}{1-\alpha}$, then the hypothesis relating $[G:H]$ and $|H|$ implies

$$\left(\ln \frac{5}{1-\alpha} + \ln n\right) n \leq (2\rho - 1)\,|H|$$

so that the failure probability $P_1 \leq \frac{1-\alpha}{5}$.

Therefore, Step 2 succeeds with probability at least $1 - \epsilon_2$, where $\epsilon_2 < \frac{1}{n} + ne^{-\frac{(2\rho-1)|H|}{n}}$.

Step 3. We build a *big forest* in G/H as follows. Fix a generator $s_1 \in S - R$, and consider the cyclic subgroup C_1 of G generated by s_1. The subgroup C_1 acts on G/H by right multiplication. In each orbit that includes a good coset, fix one particular good coset; let $Hx_1, \ldots, Hx_k$ be the good cosets chosen. Each orbit of C_1 on G/H has the same cardinality; indeed, if $gs_1^m \in Hg$ for some $g \in G$, then $s_1^m \in g^{-1}Hg = H$, so $g's_1^m \in g'H = Hg'$ for any $g' \in G$. Let m be the cardinality of each orbit.

For each good coset Hg in G/H such that $Hgs_1 \neq Hx_i$ for any $i \in \{1, \ldots, k\}$, add the vertices Hg and Hgs_1 to the big forest, and add a directed edge from Hgs_1 to Hg. There are no cycles in the big forest, because every edge connects two vertices that are adjacent within an orbit that is a cycle, and the edge of that cycle between $Hx_is_1^{-1}$ and Hx_i is not in the forest. Each tree in the big forest is a directed path, and we consider the vertex that is a target but not a source of an edge to be the root of the tree. Because the root of a tree is the target of an edge, it is a good coset. Each tree in the big forest contains at least two vertices, because a vertex is added only when it is the source or target of an edge.

We call the sources of edges in the big forest *linkable* cosets, and the targets *established* cosets; each edge points from a linkable coset Hg_j to an *established neighbor* $Hg_i = Hg_js_1^{-1}$. Every established coset is a good coset. At most k of the good cosets are not vertices in the big forest, but the good cosets $Hx_1, \ldots, Hx_k$ are roots of trees in the big forest, so at least half of the good cosets belong to the big forest. Since every path in the good forest contains at least two vertices, and every vertex of a path except its root is linkable, at least $\frac{1}{4}\alpha^2 n$ linkable cosets exist.

We claim that the sizes of the intersections $|B_j \cap A_i s_1|$ for the linkable cosets Hg_j are mutually independent. Indeed, let $I = \{j_1, \ldots, j_r\} \subset \{1, \ldots, n\}$ be a set of indices of linkable cosets. Let $i_1, \ldots, i_r$ be the indices of the established neighbors of the cosets indexed by $j_1, \ldots, j_r$ respectively. We argue by induction on $r = |I|$ that for any values $x_1, \ldots, x_r$,

$$Pr(|B_{j_1} \cap A_{i_1}s_1| \leq x_1, \ldots, |B_{j_r} \cap A_{i_r}s_1| \leq x_r) = \prod_{l=1}^{r} Pr(|B_{j_l} \cap A_{i_l}s_1| \leq x_l).$$

Here, each Pr without a subscript means the probability over the selection of A_i inside Hg_i for all i, as defined in step 1.

If $r = 1$, there is nothing to prove. If $r > 1$, reorder I so that the distance of Hg_{i_r} from the root of its tree is maximal. Then Hg_{i_r} cannot be the established neighbor of any coset indexed by I. Thus $i_r \neq j_l$ for all l. Hence

$$\begin{aligned}
&Pr\big(|B_{j_1} \cap A_{i_1} s_1| \leq x_1, \ldots, |B_{j_r} \cap A_{i_r} s_1| \leq x_r\big) \\
&= Pr\big(|B_{j_1} \cap A_{i_1} s_1| \leq x_1, \ldots, |B_{j_{r-1}} \cap A_{i_{r-1}} s_1| \leq x_{r-1}\big) \\
&\quad \times Pr_{X \in \mathcal{A}}\big(|B_{j_r} \cap X g_{j_r}| \leq x_r\big) \\
&= Pr\big(|B_{j_1} \cap A_{i_1} s_1| \leq x_1, \ldots, |B_{j_{r-1}} \cap A_{i_{r-1}} s_1| \leq x_{r-1}\big) \\
&\quad \times Pr\big(|B_{j_r} \cap A_{i_r} s_1| \leq x_r\big) \\
&= \prod_{l=1}^{r} Pr\big(|B_{j_l} \cap A_{i_l} s_1| \leq x_l\big)
\end{aligned}$$

applying the inductive hypothesis. This shows the mutual independence of the random variables $|B_j \cap A_i s_1|$.

For each j, apply Lemma 13 to $X = B_j$ with $\gamma = \frac{1}{2}$ and $\eta = 2(1-\rho)$. Using the Chernoff bound with $\delta_0 = \frac{1}{2}$ and $p_0 = 2\rho - 1$ on the $n_0 = \frac{1}{4}\alpha^2 n$ linkable cosets, we find that with probability at least $1 - \upsilon_1$, there exist at least $\frac{1}{4}(\rho - \frac{1}{2})\alpha^2 n$ linkable cosets Hg_j such that $|B_j \cap A_i s_1| > \frac{1}{2}(1-\rho)|H|$, where $\upsilon_1 = e^{-\frac{(\rho - \frac{1}{2})\alpha^2 n}{16}}$.

After step 3, the $\frac{1}{2}(1-\rho)|H|$ vertices of each of these linkable cosets attach, independently with probability p, to the large connected component of size $\rho\alpha^2|G|$ we found in step 2. Again by the Chernoff bound, using $p_0 = \frac{1}{2} < p$, $\delta_0 = \frac{1}{2}$, and $n_0 = \frac{1}{8}(\rho - \frac{1}{2})(1-\rho)\alpha^2|G|$, this adds at least $\frac{n_0}{4}$ vertices to the large component from bad parts of linkable cosets, with probability at least $1 - \upsilon_2$, where $\upsilon_2 = e^{-\frac{(\rho - \frac{1}{2})(1-\rho)\alpha^2|G|}{128}}$.

Altogether, given that step 1 and step 2 succeeded, there is a connected component of size

$$\rho\alpha^2|G| + \frac{n_0}{4} \tag{8}$$

remaining after step 3 with probability at least $1 - \epsilon_3$, where $\epsilon_3 = \upsilon_1 + \upsilon_2$. Write $\omega = \frac{n_0}{4\alpha^2|G|}$. If $\alpha > \sqrt{\frac{\rho}{\rho+\omega}}$, then expression 8 describes a component of size at least ρ, and step 3 succeeds with probability at least $1 - \epsilon_3$.

To ensure that $\epsilon_1 + \epsilon_2 + \epsilon_3 < 1 - \alpha$, we require $n > \frac{128\eta}{(\rho - \frac{1}{2})(1-\rho)\alpha^2(1-\alpha)^2}$ and apply the hypothesis relating n and $|H|$. This proves Theorem 5, with $\beta = \sqrt{\frac{\rho}{\rho+\omega}}$, $\eta = \ln\frac{5}{1-\alpha}$, and $N = \frac{128\eta}{(\rho - \frac{1}{2})(1-\rho)\alpha^2(1-\alpha)^2}$. $\square$

6 Semidirect Products

Recall the construction of a semidirect product. Let K and H be finite groups. An *action* of K on H is a homomorphism $\varphi : K \to \mathrm{Aut}(H)$. Denote $h^k = ((\varphi(k))(h)$. The *semidirect product* of H and K, denoted $H \rtimes K$, is the group defined on the set of ordered pairs $(h,k) \in H \times K$ with multiplication given by

$$(h_1,k_1)\cdot(h_2,k_2) = (h_1 \cdot h_2^{k_1}, k_1 \cdot k_2)\ .$$

The homomorphisms $h \to (h,1)$ and $k \to (1,k)$ identify H and K with subgroups of G.

Theorem 14. *Let constants ρ and α satisfy the conditions of the Reduction Theorem. There exists a constant C so that if $G = H \rtimes K$, $|H| > C$, $|K| > C$, and $p > 0$, and H and K have generating sets R and S for which $\Gamma(H,R)$ and $\Gamma(K,S)$ belong to $\mathcal{L}(\rho,\alpha,p)$, then $\Gamma(G, R\cup S) \in \mathcal{L}(\rho,\alpha,p)$.*

Proof: We may write the elements of G uniquely as $g = hk$ where $h \in H$ and $k \in K$. Given any $h \in H$, let K_h be the subgraph $\{hk : k \in K\}$, with edges joining hk to hks for $s \in S$. For $k \in K$, let H_k be the subgraph $\{hk : h \in H\}$, with edges joining hk to hkr for $r \in R$. The product structure in $H \rtimes K$ is given by

$$h_1k_1 \cdot h_2k_2 = (h_1(k_1h_2k_1^{-1}))(k_1k_2)$$

Examining this product when $k_2 = 1$ or $h_2 = 1$, we see that the sets H_k and K_h are closed under right multiplication by elements of H or K, respectively.

Clearly, for every $h \in H$, the graph of K_h is isomorphic to the Cayley graph of K. For each $k \in K$, the graph of H_k is isomorphic to the Cayley graph of H, under the isomorphism $(khk^{-1})k = kh \to h$. Indeed, if $h_1r = h_2$, then $(kh_1)r = k(h_1r) = kh_2$. Thus, each K_h and each H_k has a component of size at least $\rho|K|$ or $\rho|H|$ with probability at least α independently.

First, assume that $|H| \le |K|$. Proceed as in the proof of the Reduction Theorem, with the following change: We show that Step 2 succeeds with probability $1-\epsilon_2$, where

$$\epsilon_2 < \binom{|H|}{2}(2(1-\rho))^{a''|K|} + e^{-\frac{\alpha(1-\delta)^2|H|}{2}}. \tag{9}$$

The proof of this estimate follows.

We regard the sets H_k as the "columns" and the sets K_h as the "rows" of the Cayley graph G. If some column H_k (or row K_h) has a connected component of size $\rho|H|$ (or $\rho|K|$) considering only the generators in R (or in S), we call the column (or row) "good." In this event, we choose a subset of size exactly $\rho|H|$ (or $\rho|K|$) uniformly at random among those that remain connected after percolation, and call this subset the "good part." Step 2 succeeds if the good parts of all good columns are connected.

At the end of Step 1, we established that with probability at least $1-\varepsilon_1$, there were at least $\alpha^2|K|$ good columns. Suppose this to be the case. Pick $\delta < 1$ so that $2\alpha\delta + 2\rho > 3$. Put $a = \alpha\delta$. The Chernoff bound (Theorem 10) with $p_0 = \alpha$, $n_0 = |H|$, and $\delta_0 = 1-\delta$, shows that at least $a|H|$ good rows exist, with probability at least $1 - e^{-\frac{\alpha(1-\delta)^2|H|}{2}}$. All good columns form a single connected component in the Cayley graph of G, with high probability. Indeed, the good part of any good column intersects at least $(a-(1-\rho))|K|$ good rows. Let

$a' = a - (1-\rho)$. The good parts of any pair of good columns intersect at least

$$\begin{aligned}(2a'-1)|K| &= (2(a-(1-\rho))-1)|K| \\ &= (2a+2\rho-3)|K|\end{aligned}$$

of the same good rows. Let $a'' = 2a + 2\rho - 3$. If both columns touch the good part of such a good row, then their large components are connected in the Cayley graph of G. The probability that for some pair of good columns, this fails to happen in every such good row is no more than

$$\binom{|H|}{2}(2(1-\rho))^{a''|K|}.$$

Otherwise, the good parts of all good columns are connected in $\Gamma(G, R \cup S)$. This proves inequality 9, whose right hand side goes to zero as $|K| \to \infty$, assuming $|H| \le |K|$.

Let ϵ_1 and ϵ_3 be as in the proof of the Reduction Theorem. For C sufficiently large, $\epsilon_1 + \epsilon_2 + \epsilon_3 < 1 - \alpha$. This proves Theorem 14 in the case $|H| \le |K|$.

In the case where $|H| > |K|$, we may proceed in the same manner, with the symbols H and K interchanged. Because K might not be a normal subgroup of G, the Reduction Theorem does not formally apply. However, the proof only uses the fact that subgraph K_h is isomorphic to the Cayley graph of K, which we verified above. $\square$

7 Cayley Graphs of Abelian Groups

7.1 Correlation Length

Our proof resembles that of Theorem 6 in Grimmett [13]. To follow it, we must introduce some notation. We write $P_p(A)$ for the probability of an event A in a p–percolation process on the square lattice. Let $B(n)$ be a box inside the square lattice, centered at the origin, and with side length $2n$. Let $P_p(0 \leftrightarrow \partial B(n))$ denote the probability that there exists an open path from 0 to some point on the boundary of $B(n)$ after p–percolation on the edges of $B(n)$. Let $\xi : (0, \frac{1}{2}) \to (0, \infty)$ be the *correlation length*, *i.e.* the continuous, increasing function defined by the property that

$$\frac{\ln P_p(0 \leftrightarrow \partial B(n))}{-\frac{n}{\xi(p)}} \to 1$$

as $n \to \infty$. The function ξ converges to 0 as $p \to 0$ and converges to ∞ as $p \to \frac{1}{2}$ (see, *e.g.*, [13]).

A p–percolation process on the square lattice $\mathbb{Z}^2$ can be viewed a $(1-p)$–percolation process on the *dual lattice*, whose points are ordered pairs of the form $(a + \frac{1}{2}, b + \frac{1}{2})$ for $a, b \in \mathbb{Z}$, and whose edges run from $(a + \frac{1}{2}, b + \frac{1}{2})$ to $(a + \frac{1}{2} \pm 1, b + \frac{1}{2} \pm 1)$. Under this identification, an edge of the dual lattice is deleted ("closed") if and only if the unique edge of the square lattice intersecting it is not deleted ("open").

Lemma 15. *Let a be a positive real number, and k be a positive integer, and $p > \frac{1}{2}$. Let D_k be the box of the dual lattice with center $(k+\frac{1}{2}, \frac{1}{2})$ and side length $2a \log k$. Let E_k be the event that the vertex $(k+\frac{1}{2}, \frac{1}{2})$ is joined by a closed path of the dual to a vertex on the surface ∂D_k of D_k. Then*

$$\frac{\log P_p(E_k)}{-\frac{a}{\xi(1-p)} \log k} \to 1$$

as $k \to \infty$.

This lemma follows immediately from the definitions.

7.2 Proof of Lemma 7

Put $\gamma = \frac{1-\alpha}{2}$, and consider the function

$$f(x) = -a + a \log x.$$

By Theorem 6, there exists $\epsilon_1 > 0$ such that for $p > 1-\epsilon_1$, the subgraph $G(f)$ of the square lattice has an infinite component after p–percolation with probability more than $1 - \gamma$.

Given a positive integer x, say that "there exists an infinite path from x in $G(f)$" if there exists y such that (x, y) belongs to an infinite open path in $G(f)$. In this event, for any $x' > x$, there exist y and y' such that (x, y) is connected to (x', y') by an infinite open path in $G(f) \cap ([x, x'] \times \mathbb{Z})$. Choose N_1 so that for $p > 1 - \epsilon_1$, if $G(f)$ has an infinite component, then the probability that there exists an infinite path from N_1 in $G(f)$ is at least $1 - \gamma$.

For a p–percolation event on the square lattice, let $\theta(p)$ denote the probability that the origin resides in an infinite open cluster. Theorem 8.8 of [13] shows that θ is a continuous function on the interval $(1/2, 1]$. In particular, $\lim_{p\to 1} \theta(p) = 1$.

Suppose $\delta < 1$ and $p\delta > 1/2$. Let $B(k)$ be the box $[0, k] \times [0, k]$ inside the square lattice. Let $X_\delta(k)$ denote the event that this box contains an open cluster of size at least $\delta\theta(p)|B(k)|$, with the property that, for some x_1 and x_2, there exists an open path from $(x_1, 0)$ to (x_2, k) inside this cluster. (We call such a path an *up–down crossing* of $B(k)$ inside the cluster.) Theorem 7.61 of [13] says that $\lim_{k\to\infty} P_p(X_\delta(k)) = 1$.

Choose ϵ, δ, and $\nu < 1$ so that $\epsilon < \epsilon_1$, and for $p > 1-\epsilon$, we have $\nu\delta\theta(p) > \rho$. Fix N_2 so that for $k > a \log N_2$, we have

$$\nu P_p(X_\delta(k))\delta\theta(p) > \rho \tag{10}$$

and

$$e^{-kP_p(X_\delta(k))(1-\nu)^2} < \gamma. \tag{11}$$

Let N be the maximum of N_1 and N_2. Finally, decrease ϵ so that for the finitely many m by n boxes Γ with $a \log n \le m \le n \le N$, we have $p_c(\Gamma; \rho, \alpha) < 1 - \epsilon$.

We show that our choice of parameters satisfies the required properties in the special case where $m = a \log n$. We only need to treat the case where $n > N$. Embed Γ in the square lattice as the box $[n, 2n] \times [0, m]$. With probability at least $1 - \gamma$, there is an infinite path C in $G(f)$ from n. Suppose this event occurs.

Define events $X_\delta^{(x,y)}(k)$ for the boxes $B(k) + (kx, ky)$ as $X_\delta(k)$ was defined for $B(k)$. For fixed k, the events $X_\delta^{(x,y)}(k)$ and $X_\delta^{(x',y')}(k)$ are independent if $(x, y) \neq (x', y')$. Let $r = \lceil \frac{n}{a \log n} \rceil$, and consider the events $X_\delta^{(r-1,0)}(n)$, ..., $X_\delta^{(2r-1,0)}(n)$. By equation 11, with probability at least $1-\gamma$, at least $\frac{\rho}{\delta\theta(p)}(r+1)$ of these events occur. If the event $X_\delta^{(x,0)}(n)$ occurs, it contributes $\delta\theta(p)|B(n)|$ vertices to the large cluster of Γ, because the up–down crossing within $B(n) + (nx, 0)$ will intersect the long path C. Using equation 10, we conclude that for $p > 1 - \epsilon$, with probability $1 - 2\gamma > \alpha$, the box $[n, 2n] \times [0, a \log n]$ has an open cluster of size at least $\rho a n \log n$.

We summarize what we have shown so far as follows:

Lemma 16. *Fix $a \in \mathbb{R}^{>0}$. For every positive integer n, let Γ_n be the box $[0, n] \times [0, a \log n]$ inside the square lattice. Given α_0 and ρ_0, there exists $\epsilon_0 > 0$ so that for every n and every $p > 1 - \epsilon_0$, the probability that Γ_n has an open cluster of size at least $\rho_0 |\Gamma_n|$ containing a left–right crossing after p–percolation is at least α_0.* □

Now we handle the case where $a \log n < m < n$.

Choose ρ_0, α_0, and ν_0 so that $\nu_0 \alpha_0 \rho_0 > \rho$, and take ϵ_0 as in Lemma 16. Put $\gamma_0 = \frac{1-\alpha}{2}$. Choose a positive integer S so that

$$(1 - \frac{1}{S})\nu_0 \alpha_0 \rho_0 > \rho \tag{12}$$

and

$$e^{-S\alpha_0(1-\nu_0)^2} < \gamma_0. \tag{13}$$

There exists an integer N_0 so that for $n > N_0$, we have $\frac{n}{a \log m} > S$. Fix m and n with $n > N_0$, and embed the m by n box in the square lattice as $\Gamma = [n, 2n] \times [0, m]$ as before.

For $1 \leq k \leq \lfloor S \rfloor$, let E_k be the event that the box $[n + (a-1)k \log m, n + ak \log m] \times [0, m]$ has an up–down crossing through a ρ_0–sized component. Together, these boxes cover at least $\left(1 - \frac{a \log m}{n}\right)|\Gamma| > \left(1 - \frac{1}{S}\right)|\Gamma|$ of the box Γ. By Lemma 16, $P_p(E_k) \geq \alpha_0$. By equation 13 and the Chernoff bound (Theorem 10), with probability at least $1 - \gamma_0$, at least $\nu_0 \alpha_0 \lfloor \frac{n}{a \log m} \rfloor$ of the events E_k occur. Take $\epsilon > 0$ such that $\epsilon < \epsilon_0$ and $p_c(\Gamma; \rho, \alpha) > 1 - \epsilon$ when $a \log n \leq m \leq n \leq N_0$, and so that for $p > 1 - \epsilon$, $G(f)$ has a long path from N_0 with probability at least $1 - \gamma_0$. Altogether, for such p, with probability at least α, there is a long path connecting the large components from each of the events E_k into a component of size at least $\nu_0 \alpha_0 (1 - \frac{1}{S})\rho_0 |\Gamma| > \rho |\Gamma|$ in Γ. This is the desired result. □

7.3 Proof of Theorem 3

We wish to embed the Cayley graph $\Gamma = \Gamma(G, S)$ of an abelian group G into a two–dimensional box, so that we can apply Lemma 7.

The generators $s_1, \ldots, s_n$ in the Hall basis define a homomorphism $\varphi : \mathbb{Z}^n \to G$, given by

$$\varphi(x_1, \ldots, x_n) = s_1^{x_1} \cdots s_n^{x_n} .$$

Let the generators have orders $a_1, \ldots, a_n$. The homomorphism φ maps the box $B = [0, a_1 - 1] \times \cdots \times [0, a_n - 1]$ bijectively onto G. To flatten B into a two–dimensional box, we will select k dimensions and choose Hamiltonian paths in the Cayley graphs of a section and a cross section. Unwrapping these Hamiltonian paths each into one dimension will produce the desired two–dimensional box.

We claim that there exists $I \subset \{1, \ldots, n\}$ such that $\prod_{i \in I} a_i > \frac{\log|G|}{2}$ and $\prod_{i \notin I} a_i > \frac{\log|G|}{2}$. These constraints will allow us to apply Lemma 7 to the resulting two–dimensional box. Indeed, choose the smallest k such that $a_1 \cdots a_k > \frac{\log|G|}{2}$. If $a_1 \cdots a_k < \frac{2|G|}{\log|G|}$, then we may take $I = \{1, \ldots, k\}$, and we are done. If $k = 1$, this inequality is assured by the diameter assumption, since $\mathrm{diam}(\Gamma) = (a_1 + \cdots + a_n)/2$. If $k > 1$ and yet $a_1 \cdots a_k > \frac{2|G|}{\log|G|}$, then $a_k > \frac{4|G|}{\log^2|G|} > \frac{\log|G|}{2}$, assuming $|G|$ is large enough. The diameter condition $a_k < \frac{2|G|}{\log|G|}$ implies that $a_1 \cdots a_{k-1} a_{k+1} \cdots a_n > \frac{\log|G|}{2}$, so $I = \{k\}$ has the desired property.

Now choose Hamiltonian paths β_1 and β_2 in boxes $B_1 = \prod_{i \in I} [0, a_i]$ and $B_2 = \prod_{i \notin I} [0, a_i]$ (see, *e.g.*, [16]). One can view these paths as maps $\beta_1 : [0, x - 1] \to B_1$ and $\beta_2 : [0, y - 1] \to B_2$. Let A be the box $[0, x - 1] \times [0, y - 1]$. Observe that $\varphi \circ (g, h)$ is a graph homomorphism mapping A bijectively onto G, so that A is isomorphic to a spanning subgraph of Γ. Therefore, it suffices to show that $A \in \mathcal{L}(\rho, \alpha, 1 - \epsilon)$. This follows immediately from Lemma 7, since x and y are each at least $\frac{\log|G|}{2}$. $\square$

8 Examples

1. Our first example is a hypercube C_n, which is a Cayley graph of the group $\mathbb{Z}_2^n$ with the usual set of generators $R = \{r_1, \ldots, r_n\}$. In this case, $\mathrm{diam}(C_n) = n = o(\frac{2^n}{n})$. Therefore, $p_c(C_n) < 1 - \varepsilon$ for some $\varepsilon > 0$, by Theorem 3. Of course, this bound is much weaker than $p_c = (1 + o(1))/n$ established in [1].

2. Consider $G_n = S_n \ltimes \mathbb{Z}_2^n$, with the generating set

$$R_n = \{((i\; i+1), \overline{0}), (\mathrm{id}, r_j); i = 1, \ldots, n-1; j = 1, \ldots, n\}$$

where $\{r_1, \ldots, r_n\}$ are the usual generators for $\mathbb{Z}^n$. From the previous example, Proposition 11, and Theorem 14,

$$p_c(\Gamma(G_n, R_n)) < \max\{p_c(C_n), p_c(\Gamma(S_n, R_n))\} < 1 - \varepsilon$$

for some $\varepsilon > 0$.

3. Fix a prime power q. Let $G_n = U(n, \mathbb{F}_q)$ be the group of $n \times n$ upper triangular matrices over the finite field with q elements, with ones along the diagonal. Consider the set $L_n = \{E_{i,j}^{\pm} : 1 \le i < j \le n\}$ of all elementary transvections $E_{i,j}^{\pm}$ with ± 1 in position (i, j), ones along the diagonal, and zeros elsewhere. For each $m \le n$, let H_m be the subgroup of G_n generated by the $E_{i,j}^{\pm}$ with $j > i + (n - m)$ (consisting of matrices with zero on the first $n - m$ superdiagonals).

   For $m < \frac{n}{2}$, H_m is isomorphic to $\mathbb{F}_q^{\frac{m(m+1)}{2}}$. Therefore, $p_c(\Gamma(H_m, L_n \cap H_m)) < 1 - \varepsilon$ for some ε that is independent of m and n. For n sufficiently large and $m \ge \frac{n}{2}$, the subgroup H_{m-1} of H_m in G_n will satisfy the index condition (1), and the Reduction Theorem 5 will show that $p_c(\Gamma(H_m, L_n \cap H_m)) < 1 - \varepsilon$ for the same ε as before. Since $G_n = H_n$, this gives a bound $p_c(\Gamma(G_n, L_n) < 1 - \varepsilon$ for a value of ε that is independent of n.

4. Let $G_n = B(n, \mathbb{F}_q)$ be the set of upper triangular n by n matrices with entries in $\mathbb{F}_q$, and let $H_n = U(n, \mathbb{F}_q)$. Let R_n be any generating set for the diagonal subgroup. Then $R_n \cup L_n$ generates G_n, and equation (1) is satisfied for large n. The Reduction Theorem 5 gives $p_c(\Gamma(G_n, R_n \cup L_n)) < 1 - \varepsilon$.

5. Let $G_n = U(n, \mathbb{F}_q)$ and $R_n = \{E_{i,i+1}^{\pm} : i = 1, \ldots, n-1\}$. Theorem 4 applies in the same manner as in Proposition 11.

6. Consider S_n with the star transpositions $R_n = \{r_i = (1\ i) : i = 2, \ldots, n\}$. None of these generators commute, so we cannot apply Theorem 4. However, the short relations $(r_i\ r_j)^3 = 1$ can be used in Theorem 12 to obtain $p_c(\Gamma(S_n, R_n); \frac{2}{3}, \frac{1}{2}) < 1 - \varepsilon$.

9 Concluding Remarks

We are unable to prove the Benjamini conjecture in its full generality, even for abelian groups. It would be nice to prove the Benjamini conjecture for all generating sets of finite abelian groups.

In view of the Reduction Theorem, it is important to study simple groups with small generating sets. For example, any simple group can be generated by two elements, one of which is an involution (see [12]). The corresponding Cayley graph may provide interesting test cases for Benjamini's conjecture.

It is well known (see [3]) that all Cayley graphs Γ_n of the symmetric group S_n have a diameter $e^{o(\sqrt{n \log n})} = o\left(\frac{n!}{n \log n}\right)$. Proving Benjamini's conjecture in these cases is the ultimate challenge for the reader. Even for the generating set $\{(1\,2), (1\ 2\ \cdots n)^{\pm 1}\}$, we are unable to bound p_c away from 1.

Acknowledgements

We would like to thank Itai Benjamini for telling us about his conjecture, reading an early draft of the paper, and providing reference [2]. The extended abstract of this paper originally appeared in the proceedings of Random 2002, and we thank an anonymous referee for pointing out an error in the original draft. We thank another anonymous referee for extensive suggestions regarding a later draft.

References

[1] Ajtai, M., Komlós, J., and Szemerédi, E. Largest random component of a k-cube. *Combinatorica 2*, 1 (1982), 1–7.

[2] Alon, N., Benjamini, I., and Stacey, A. Percolation on finite graphs and isoperimetric inequalities. *math.PR/0207112* (2002).

[3] Babai, L. Automorphism groups, isomorphism, reconstruction. In *Handbook of combinatorics, Vol. 1, 2.* Elsevier, Amsterdam, 1995, pp. 1447–1540.

[4] Babson, E., and Benjamini, I. Cut sets and normed cohomology with applications to percolation. *Proc. Amer. Math. Soc. 127*, 2 (1999), 589–597.

[5] Benjamini, I. Percolation on finite graphs. *math.PR/0106022* (2001).

[6] Benjamini, I., and Schramm, O. Percolation beyond $\mathbf{Z}^d$, many questions and a few answers. *Electron. Comm. Probab. 1* (1996), no. 8, 71–82 (electronic).

[7] Benjamini, I., and Schramm, O. Recent progress on percolation beyond $\mathbf{Z}^d$. *http://www.wisdom.weizmann.ac.il/~schramm/ papers/pyond-rep* (2000).

[8] Bollobás, B. *Random graphs.* Academic Press Inc. [Harcourt Brace Jovanovich Publishers], London, 1985.

[9] Bollobás, B. *Random graphs*, second ed. Cambridge University Press, Cambridge, 2001.

[10] Diestel, R., and Leader, I. A conjecture concerning a limit of non-Cayley graphs. *J. Algebraic Combin. 14*, 1 (2001), 17–25.

[11] Erdős, P., and Rényi, A. On the evolution of random graphs. *Magyar Tud. Akad. Mat. Kutató Int. Közl. 5* (1960), 17–61.

[12] Gorenstein, D. *Finite simple groups.* Plenum Publishing Corp., New York, 1982.

[13] Grimmett, G. *Percolation*, second ed. Springer-Verlag, Berlin, 1999.

[14] Janson, S., Łuczak, T., and Rucinski, A. *Random graphs.* Wiley-Interscience, New York, 2000.

[15] Krivelevich, M., Sudakov, B., and Vu, V. H. A sharp threshold for network reliability. *Combinatorics, Probability and Computing 11*, 5 (2002), 465–474.

[16] Lovász, L. *Combinatorial problems and exercises*, second ed. North-Holland Publishing Co., Amsterdam, 1993.

[17] Pak, I., and Vu, V. H. On mixing of certain random walks, cutoff phenomenon and sharp threshold of random matroid processes. *Discrete Appl. Math. 110*, 2-3 (2001), 251–272.

[18] Shalev, A. Simple groups, permutation groups, and probability. In *Proceedings of the International Congress of Mathematicians* (1998), Documenta Mathematica, ICM Extra Volume II, pp. 129–137 (electronic).