

Weighted isoperimetry implies percolation

Ivailo Hartarsky¹, Franco Severo², Augusto Teixeira³

¹CNRS, IRL2924 Jean-Christophe Yoccoz, IMPA, Estrada Dona Castorina, 110 - Rio de Janeiro, Brazil, hartarsky@math.univ-lyon1.fr

²CNRS, Sorbonne Université, 4 place Jussieu, 75005 Paris, France, severo@lpsm.paris

³IMPA, Estrada Dona Castorina, 110 - Rio de Janeiro, Brazil, augusto@impa.br

October 9, 2026

Abstract

Consider an infinite edge-weighted graph satisfying an isoperimetric inequality of the type $\|\partial A\| \geq C|A|^\alpha$ for some $\alpha, C > 0$, where $\|\partial A\|$ denotes the weighted size of the edge boundary of A . We prove that, for C large enough depending on α , if each edge is open independently with probability given by its weight, then any vertex is connected to infinity with positive probability. The result also holds under weaker isoperimetric assumptions and on finite graphs.

The proof brings a new perspective on the recent proof of the Benjamini–Schramm conjecture concerning the same problem with homogeneous weights. The crucial novelty in our proof is that, rather than simply counting cutsets, we introduce a new Peierls argument which takes into account internal and external connectivity costs in addition to the cost of the blocking surface.

We provide two applications for the above result. First, we show that every non-summable long-range percolation on $\mathbb{Z}^d$, $d \geq 2$, admits a percolating truncation, solving a conjecture of Sidoravicius, Surgailis and Vares and its generalization by Friedli and de Lima. Secondly, we show that there exists a universal constant $C < \infty$ such that $p_c \leq C/\Delta$ for every transitive graph of superlinear growth and vertex degree Δ , thus proving a conjecture of Easo and Hutchcroft.

MSC2020: 82B43; 60K35; 05C80; 60C05; 05C70

Keywords: percolation; isoperimetry; cutsets; truncation; transitive graphs

1 Introduction

1.1 Main result

Let V be either a finite or an infinite countable set and $\lambda = (\lambda_{\{x,y\}})_{\{x,y\} \subset V}$ be a family of non-negative weights. One can see $G = (V, \lambda)$ as a weighted graph, where the edge set is given by $E := \{\{x, y\} : \lambda_{\{x,y\}} > 0\}$. For any set of edges $F \subset E$, we define its total weight by $\|F\|_\lambda := \sum_{f \in F} \lambda_f$. Consider the *weighted isoperimetric profile* $\psi_\lambda : \{1, \dots, \lfloor \frac{|V|}{2} \rfloor\} \rightarrow [0, \infty]$ given by

$$\psi_\lambda(n) := \inf \left\{ \|\partial S\|_\lambda : S \subset V, n \leq |S| \leq \min(2n, |V|/2) \right\}, \quad (1)$$

where $\partial S := \{\{x, y\} \in E : x \in S, y \notin S\}$ is the *edge boundary* of S . Given two vertices $x, y \in V$, we set

$$\chi_\lambda(x, y) = \inf \{\|\partial S\|_\lambda : x \in S \subset V \setminus \{y\}\}, \quad \chi_\lambda(x) = \inf \{\|\partial S\|_\lambda : x \in S \text{ finite}\}. \quad (2)$$

We also consider the associated *bond percolation* model on $G = (V, \lambda)$ where each edge $e \in E$ is declared open independently with probability $p_e := 1 - e^{-\lambda_e}$. We write $\omega \subset E$ for percolation configurations and denote the corresponding probability measure by $\mathbb{P}_\lambda$. Given two vertices $x, y \in V$, we denote by $x \leftrightarrow y$ the event that there exists a connected component C of (V, ω) such that $\{x, y\} \subset C$. We write $x \leftrightarrow \infty$ if $x \leftrightarrow y$ occurs for infinitely many y .

Our main result is the following.

Theorem 1.1 (Weighted isoperimetry implies percolation). *Let $G = (V, \lambda)$ be a weighted graph such that*

$$\Psi^{-1} := \sum_{k \geq 0} \frac{1}{\psi_\lambda(2^k)} < \frac{1}{625}. \quad (3)$$

Then the following holds, where $\varepsilon = \varepsilon(\Psi) := 25/\sqrt{\Psi} < 1$.

(i) *If V is finite, then for every $x, y \in V$, we have*

$$\mathbb{P}_\lambda(x \leftrightarrow y) \geq 1 - e^{-(1-\varepsilon)\chi_\lambda(x,y)}.$$

(ii) *If V is infinite, then for every $x \in V$, we have*

$$\mathbb{P}_\lambda(x \leftrightarrow \infty) \geq 1 - e^{-(1-\varepsilon)\chi_\lambda(x)}.$$

Remark 1.2 (Isoperimetric profile). The function ψ_λ in (1) is related to the following more standard function

$$\phi_\lambda(n) := \inf \left\{ \|\partial S\|_\lambda : S \subset V, n \leq |S| \leq |V|/2 \right\}. \quad (4)$$

Obviously, $\phi_\lambda(n) = \inf_{m \geq n} \psi_\lambda(m) \leq \psi_\lambda(n)$, so Theorem 1.1 also holds with ψ_λ replaced by ϕ_λ . However, for technical reasons, we will use the version with ψ_λ when deducing Theorem 1.7 below.

Remark 1.3 (Isoperimetric dimension). Assumption (3) is satisfied, in particular, if there exists $\alpha \in (0, 1]$ such that $\psi_\lambda(n) \geq \frac{1250}{\alpha} n^\alpha$ for all $n \geq 1$. Recall that when $\psi_\lambda(n) \geq cn^\alpha$ for some $c > 0$, we say that G satisfies a (weighted) *isoperimetric inequality of dimension $d = 1/(1 - \alpha)$* .

Remark 1.4 (Surface-order exponential decay). As it is common for perturbative results, our proof also yields surface-order exponential decay. For instance, if V is infinite and (3) holds, then for all $x \in V$

$$\mathbb{P}_\lambda(n \leq |\{z \in V : x \leftrightarrow z\}| < \infty) \leq e^{-(1-\varepsilon)\phi_\lambda(n)}.$$

Remark 1.5 (Giant and connectivity for finite graphs). For finite graphs, one may use Theorem 1.1(i) (see also Remark 1.4) to deduce an upper bound on the threshold for the emergence of a unique giant component from isoperimetric information. Furthermore, a union bound over y in Theorem 1.1(i) can be used to recover sharp full connectivity results. For many regular graphs of degree d and homogeneous weights $\lambda_0 \mathbb{1}_E$, one can obtain a giant for $\lambda_0 = C/d$, and full connectivity for $\lambda_0 = (1+\epsilon) \log |V|/d$, where the latter is often sharp as it is related to the non-existence of isolated vertices. For instance, this applies to complete graphs, hypercubes, random regular graphs, etc. We refer to [DEKK26, Joo15, DG25] for recent related works.

A homogeneous version of Theorem 1.1 (that is for $\lambda_e = \lambda_0 \mathbb{1}_{e \in E}$ for some $\lambda_0 > 0$) was recently proved by an AI [Cut26a, Cut26b], thereby settling, in particular, the “ $p_c < 1$ question” of Benjamini and Schramm [BS96, Question 2] on graphs with isoperimetric dimension $d > 1$. See also the predecessors [DGRSY20] and [EST25] where this conjecture was verified for $d > 4$ and $d > 2$, respectively.

1.2 Applications

We next gather two applications of Theorem 1.1. Let us remark that [Cut26a, Cut26b] do not suffice for the applications below, including in the case of homogeneous weights. Indeed, as we explain in Section 1.3, their proof is based on cutset counting, which necessarily degenerates as $\lambda \rightarrow 0$.

Truncation of long-range percolation. We first turn to solving the truncation problem for non-summable long-range percolation, using Theorem 1.1.

Theorem 1.6 (Percolating truncation always exists). *Fix the vertex set $V = \mathbb{Z}^d$, with $d \geq 2$. Let $\lambda = (\lambda_{\{x,y\}})_{x,y \in \mathbb{Z}^d}$ be translation-invariant family of weights, that is, $\lambda_{\{x,y\}} = \lambda_{y-x} = \lambda_{x-y}$ for all $x, y \in \mathbb{Z}^d$. Assume that $X := \{x \in \mathbb{Z}^d : \lambda_x > 0\}$ is not contained in the line $\mathbb{R}x$ for any $x \in \mathbb{R}^d$ and that $\sum_{x \in X} \lambda_x = \infty$. Then there exists $M \in [1, \infty)$ such that, setting $\lambda_{\{x,y\}}^M = \lambda_{\{x,y\}} \mathbb{1}_{\|x-y\|_1 \leq M}$, we have*

$$\mathbb{P}_{\lambda^M}(0 \leftrightarrow \infty) > 0.$$

Theorem 1.6 equivalently states that non-unidimensional non-summable λ lie in the interior of the supercritical phase for the product topology on the space $[0, \infty]^{\mathbb{Z}^d}$ of translation-invariant weights.

Theorem 1.6 was first conjectured by Sidoravicius, Surgailis and Vares [SSV99] in 1999 on $\mathbb{Z}^2$ with additional hypotheses, but remained open until now. Over the years, several particular cases of Theorem 1.6 were proved [SSV99, MSV01, FdLS04, FdL06, dLS08, CdL22, Bäu24] assuming various degrees of symmetry or structure in the weights λ and stronger non-summability or $d \geq 3$ assumptions. The most general previous result (though not encompassing the others) is the one of [Bäu24] proving Theorem 1.6 when $d \geq 3$ and λ is invariant under reflection in each of the coordinate hyperplanes. We also direct the reader to [vEdLV16, AHdLV17] for oriented versions of the problem.

Critical probability of transitive graphs. Our second application of Theorem 1.1 concerns percolation on transitive graphs of large degree (such as high-dimensional lattices, spread-out percolation). Recall that the critical point (in standard parametrization) for bond percolation on an infinite connected graph $G = (V, E)$ can be defined as

$$p_c^{\text{bond}}(G) := \inf \left\{ p \in [0, 1] : \mathbb{P}_{-\log(1-p)\mathbb{1}_E}(x \leftrightarrow \infty) > 0 \right\},$$

where the definition does not depend on $x \in V$.

Theorem 1.7 (Large average degree implies percolation). *There exists a universal constant $C < \infty$ such that if $G = (V, E)$ is a transitive graph with superlinear growth¹ and vertex degree Δ , then $p_c^{\text{bond}}(G) \leq 1 - e^{-C/\Delta} \leq C/\Delta$.*

Theorem 1.7 confirms [EH23, Conjecture 7.3]. The fact that $p_c^{\text{bond}}(G) < 1$ for transitive graphs of superlinear growth was conjectured in [BS96] and first proved in [DGRSY20]. The existence of a gap at 1 (i.e. a universal $\varepsilon > 0$ such that $p_c(G) \leq 1 - \varepsilon$ for every such graph), was first proved for Cayley graphs in [PS23] (including site and bond percolation), and for transitive graphs in [EST25] (only for bond percolation). In fact, combining the well-known bound $p_c^{\text{site}}(G) \leq 1 - (1 - p_c^{\text{bond}}(G))^{\Delta-1}$ from [GS98] with our Theorem 1.7 establishes a gap at 1 for site percolation on transitive graphs, which was open until now.

We recall the classical fact that $p_c^{\text{bond}}(G) \geq 1/(\Delta - 1)$ for any graph of maximum degree Δ . Asymptotic analogues of Theorem 1.7 were proved with a constant $C \rightarrow 1$ as $\Delta \rightarrow \infty$ for spread-out percolation on $\mathbb{Z}^d$ [Pen93, BJR07] and on transitive graphs of polynomial growth [ST24]. However, it is known [EH23, Figure 5] that C cannot be taken arbitrarily close to 1 for general transitive graphs, even as $\Delta \rightarrow \infty$.

1.3 Outline of the proof

We now discuss the main ideas in the proof of Theorem 1.1. For simplicity we assume that V is finite (the case of infinite V follows by restricting and wiring the boundary), and we aim at upper bounding $\mathbb{P}[x \not\leftrightarrow y]$ uniformly in $x, y \in V$.

In order to better appreciate the challenges and innovations of this work, we start by recalling the approach introduced in [EST25] and further developed in [Cut26a, Cut26b], which is specific to the homogeneous case $\lambda = \lambda_0 \mathbb{1}_E$. The first step, which can be traced back to the work of Peierls [Pei36], is to reduce the problem to counting (minimal) cutsets in the graph $G = (V, E)$. More precisely, first notice that the event $\{x \not\leftrightarrow y\}$ exactly corresponds to the existence of a cutset $\pi \subset E$ separating x from y (see Definition 2.1) which is ω -closed (i.e. $\omega \cap \pi = \emptyset$). In particular, $\mathbb{P}_\lambda[x \not\leftrightarrow y] \leq \sum_\pi \mathbb{P}_\lambda[\pi \text{ is } \omega\text{-closed}]$ with the sum running over inclusion-minimal cutsets. The second step is to prove that there exists a constant $C < \infty$ such that $|\{\pi : |\pi| \leq n\}| \leq e^{Cn}$ for all n . Since $\mathbb{P}[\pi \text{ is } \omega\text{-closed}] = e^{-\|\pi\|_\lambda} = e^{-\lambda_0|\pi|}$ in the homogeneous case, the result would follow readily by taking $\lambda_0 > C$ large enough. However, counting cutsets on general graphs can be a very challenging combinatorial problem. The main idea of [EST25] is to do so via a probabilistic method, namely by carefully constructing a random cutset Π

¹A graph G is said to have *superlinear growth* if $\sup_{n \geq 1} |B_n^G|/n = \infty$, where B_n^G is a ball of radius n centered at some fixed vertex.

such that $\mathbb{P}[\Pi = \pi] \geq e^{-C|\pi|}$ for all π . Such a random cutset is built in [EST25] by using the simple random walk on G and assuming that it is transient (which is known to be implied by an isoperimetric inequality of dimension $d > 2$). In [Cut26a, Cut26b], such a random cutset was obtained via a clever modification of Karger's algorithm [Kar93], and only required a summability assumption like (3). The algorithm is very simple: start with the vertex set V , where each site is seen as a singleton-cluster, then at each step select a cluster of *smallest size* out of those not containing x or y , take a uniformly chosen random edge between this cluster and its neighboring ones, and then join them into a single cluster. Iterating this process, one ends up with only two clusters (of x and y), and Π is precisely the set of edges between them.

For general inhomogeneous weights such as in Theorem 1.1, estimating the expected number of ω -closed cutsets as above has no chance to work, even in the particular contexts of Theorems 1.6 and 1.7. Indeed, given a vertex z such that $\{x, z\} \in E$, consider the cutset π_z (between x and ∞) obtained from the edges incident to either x or z , but not both. If $\lambda = \lambda^M$ as in Theorem 1.6, these cutsets satisfy $\|\pi_z\|_{\lambda^M} \leq 2\Lambda_M$, where $\Lambda_M := \sum_{\|w\|_1 \leq M} \lambda_w$. If $\lambda_w > 0$ for all w , there are at least $|B_M^{\mathbb{Z}^d}| - 1$ such cutsets, thus $\sum_{\pi} \mathbb{P}_{\lambda^M}[\pi \text{ is } \omega\text{-closed}] \geq cM^d e^{-2\Lambda_M}$, which may diverge if $\Lambda_M = o(\log M)$ as $M \rightarrow \infty$. In the context of Theorem 1.7, setting $\lambda := \frac{C}{\Delta} \mathbb{1}_E$, we have $\|\pi\|_{\lambda} \leq 2C$ for these cutsets, thus $\sum_{\pi} \mathbb{P}_{\lambda}[\pi \text{ is } \omega\text{-closed}] \geq \Delta e^{-2C}$, which diverges as $\Delta \rightarrow \infty$. One can check that such divergences may occur not only for these small cutsets, but also for arbitrarily large ones.

The main idea of our proof is to introduce the notion of *cohesion*. We say that a cutset π is ω -cohesive if, in addition to being ω -closed, none of its sides contains a small cut (i.e. with total weight < 1) which is not crossed by ω (see Definition 2.2). First, we observe that $\{x \not\leftrightarrow y\}$ happens if and only if there exists a ω -cohesive cutset (see Lemma 2.4). We then use the same algorithm as the one of [Cut26a] described above, but where edges are selected with probability proportional to their weights. This yields a random cutset Π for which we prove that (see (19))

$$\mathbb{P}[\Pi = \pi] \geq e^{(1-\varepsilon)\|\pi\|_{\lambda}} \mathbb{P}_{\lambda}[\pi \text{ is } \omega\text{-cohesive}] \quad (5)$$

for all π . Since by (2) we have $\|\pi\|_{\lambda} \geq \chi_{\lambda}(x, y)$ for all π , the bound (5) readily implies $\mathbb{P}_{\lambda}[x \not\leftrightarrow y] \leq \sum_{\pi} \mathbb{P}_{\lambda}[\pi \text{ is } \omega\text{-cohesive}] \leq \sum_{\pi} e^{-(1-\varepsilon)\|\pi\|_{\lambda}} \mathbb{P}[\Pi = \pi] \leq e^{-(1-\varepsilon)\chi_{\lambda}(x, y)}$, thus concluding the proof.

In order to prove (5), we couple the algorithm with the percolation process via the following Poissonization trick. Consider the edge-marks $(P_e)_{e \in E}$ given by independent exponential random variables with parameters $(\lambda_e)_{e \in E}$. First, note that $\omega = \mathbb{1}_{P_e \leq 1}$ has law $\mathbb{P}_{\lambda}$. Second, the law of the aforementioned algorithm can be obtained from the same $(P_e)_{e \in E}$ by merging along the edge on the boundary of the selected cluster with minimal mark. In order to preserve the exact law of the algorithm, at each step, we erase some negative information acquired on remaining edges. When estimating the probability that $\Pi = \pi$, we note that there are two kinds of steps: either the non- π boundary of the selected cluster has total weight at least 1, or smaller than 1. On the one hand, the probability of not merging along π for all steps of the first type can be estimated exactly as in [Cut26b], using the assumption (3) on the isoperimetric profile, thus leading to a lower bound of the form $e^{-\varepsilon\|\pi\|_{\lambda}} = e^{(1-\varepsilon)\|\pi\|_{\lambda}} \mathbb{P}[\pi \text{ is } \omega\text{-closed}]$. On the other hand, for steps of the second type, our construction essentially guarantees that the selected edge does not belong to π when it is ω -cohesive. As a consequence, the second type of steps generates an additional probability cost, whose total can be lower bounded by $\mathbb{P}_{\lambda}[\pi \text{ is } \omega\text{-cohesive} | \pi \text{ is } \omega\text{-closed}]$. These two estimates combined imply (5). In retrospect, this construction of the algorithm is closer to Kruskal's minimum spanning tree algorithm [Kru56], which actually inspired [Kar93].

Finally, let us mention that deducing Theorems 1.6 and 1.7 from Theorem 1.1 amounts to proving a simple isoperimetric inequality in two dimensions (Lemma 3.1) and deducing one (Lemma 4.1) on general transitive graphs from geometric group theoretic results [TT24, TT20]. These proofs are short and presented in Sections 3 and 4, respectively.

2 Proof of Theorem 1.1

In this section we prove Theorem 1.1 starting with the finite graph case. Assume V is finite and fix $x, y \in V$. For the purposes of dealing with infinite graphs later, we will work with the following




Figure 1: An illustration of a small-weight set F that disconnects V_x . Assuming that $\|F\|_\lambda < 1$, then F must intersect ω in order for π to be considered ω -cohesive.

(weaker) version of ψ avoiding x and y :

$$\psi_\lambda^*(n) = \inf \left\{ \|\partial S\|_\lambda : S \subset V \setminus \{x, y\}, n \leq |S| \leq \min(2n, |V|/2) \right\} \geq \psi_\lambda(n), \quad (6)$$

for $n \in \{1, \dots, \lfloor |V|/2 \rfloor\}$. The proof will only rely on the fact that

$$(\Psi^*)^{-1} := \sum_k 1/\psi_\lambda^*(2^k) < 1/25^2 \quad (7)$$

and $\phi_\lambda(1) = \inf_k \psi_\lambda(2^k) \geq \Psi > 25^2$, which are both consequences of our assumption (3).

Definition 2.1 (Cutset). A *cutset* (between x and y) is an inclusion-minimal subset π of E with the property that x and y are not connected in the graph $(V, E \setminus \pi)$. Given a cutset π , we denote by V_x and V_y the vertex sets of the connected components in $(V, E \setminus \pi)$ containing x and y , respectively. Let E_x and E_y denote the edge sets of the graphs induced by V_x and V_y , respectively. Note that, by minimality, for any cutset π , we have $V_x \sqcup V_y = V$.

Definition 2.2 (ω -cohesive cutset). Given $\omega \subset E$, a cutset π is said to be ω -cohesive, if the following events occur.

- $\pi \cap \omega = \emptyset$,
- for any $z \in \{x, y\}$ and $F \subset E_z$ such that $(V_z, E_z \setminus F)$ is disconnected and $\|F\|_\lambda < 1$, we have $\omega \cap F \neq \emptyset$.

In words, an ω -cohesive cutset is one which is closed and such that, in both sides of it, all small cuts meet ω (see Figure 1).

Remark 2.3 (Cohesion). To digest Definition 2.2, let us consider the homogeneous case of $\lambda_e = \lambda_0$ for all $e \in E$. If $\lambda_0 > 1$, a cutset is ω -cohesive, iff it is closed (the second condition is void). If $\lambda_0 \in (1/2, 1)$, a cutset is ω -cohesive, iff it is closed and all bridges (in the graph theoretic sense of the term) of $(V, E \setminus \pi)$ belong to ω . If G is planar and $\lambda_0 \in (1/2, 1)$, ω -cohesive cutsets correspond to *induced cycles* (also known as holes) separating x and y in the dual configuration (V^*, ω^*) , defined by $\omega_{e^*}^* = 1 - \omega_e$. Lower values of λ_0 correspond to higher order notions involving pockets accessible by multiple edges.

Lemma 2.4 (ω -cohesive cutsets exist). Assume that $\inf_k \psi_\lambda^*(2^k) \geq 2$. Let $\omega \subset E$ be such that $x \not\sim y$ in ω . Then there exists an ω -cohesive cutset.

Proof. The proof roughly says that the minimum-weight cutset not meeting ω is ω -cohesive, because any small cut in V_x allows us to switch one of its sides to V_y , yielding a cutset with smaller weight.

Let π be any minimum-weight closed cutset. More precisely, let π be some cutset for which

$$\|\pi\|_\lambda = \min\{\|\pi'\|_\lambda : \pi' \subset E \setminus \omega, \pi' \text{ cutset}\}. \quad (8)$$

The set on the right is nonempty, as we can start with $E \setminus \omega$ and delete edges one at a time, keeping x and y disconnected. Hence, π exists and $\pi \cap \omega = \emptyset$ by definition. We claim that π is ω -cohesive.

Assume by contradiction that there exists $F \subset E_x \setminus \omega$ with $G_x = (V_x, E_x \setminus F)$ disconnected and $\|F\|_\lambda < 1$. Let V'_x be the connected component of x in G_x and $W_x = V_x \setminus V'_x \neq \emptyset$. Let $\partial W_x = \{\{a, b\} \in E : a \in W_x, b \in V \setminus W_x\}$ be the boundary of W_x in G . By construction $\partial W_x \subset \pi \cup F$. Consider the set $\pi' = \pi \Delta \partial W_x \subset \pi \cup F$, where Δ is the symmetric difference. Then $\pi' \cap \omega = \emptyset$ and $x \in V'_x$ is not connected to $y \in V_y$ in $(V, E \setminus \pi')$. Hence, π' contains a cutset and (8) gives $\|\pi'\|_\lambda \geq \|\pi\|_\lambda$. However,

$$\|\pi'\|_\lambda = \|\pi\|_\lambda + 2\|\partial W_x \setminus \pi\|_\lambda - \|\partial W_x\|_\lambda \leq \|\pi\|_\lambda + 2\|F\|_\lambda - 2 < \|\pi\|_\lambda,$$

since $\|\partial W_x\|_\lambda \geq \inf_k \psi_\lambda^*(2^k) \geq 2$, as W_x is nonempty and does not contain x or y . This contradiction (and its analogue for G_y) concludes the proof. $\square$

Similarly to [Cut26a], we will construct an algorithm which starts from G and sequentially collapses edges according to some random input. This process will generate a sequence of multi-graphs, with the last one consisting of the two points $\{x, y\}$ joined by some edges. The edges present in the final multi-graph will form a (random) cutset between x and y . In order to describe a given step of this procedure, we need to setup some notation for multi-graphs with (edge and vertex) weights.

Fix a connected multi-graph $H = (V_H, E_H)$ with $x, y \in V_H$, endowed with the following functions. The first one is an edge-weight $\lambda^H : e \in E_H \mapsto \lambda_e^H \in [0, \infty)$ which will represent the percolation weight of each edge. We recall that H can be a multi-graph and, in this case, the function λ^H can assign different weights to distinct edges connecting the same endpoints. We systematically erase self-loops from multi-graphs. The second function we are given is a vertex-weight $|\cdot|_H : V_H \rightarrow \mathbb{R}$, which encodes the size of the cluster represented by each vertex. Fix also $P^H : E_H \rightarrow (0, \infty) : e \mapsto P_e^H$ which, roughly speaking, associates to each edge a “clock ring”, marking the moment at which this edge will collapse.

The next definition was one of the main inputs from [Cut26a].

Definition 2.5 (SELECT). Let $H = (V_H, E_H)$ be a multi-graph with a vertex-weight function $|\cdot|_H : V_H \rightarrow \mathbb{R}$. The procedure **SELECT** returns some $v_H \in \arg \min_{V_H \setminus \{x, y\}} |\cdot|_H$, breaking ties in an arbitrary deterministic way.

The next definition incorporates some new features with respect to [Cut26a]. It is illustrated in Figure 2.

Definition 2.6 (STEP). The procedure **STEP** takes as input $H, \lambda^H, |\cdot|_H, P^H$ and v_H as above and returns a multi-graph H' with its own $\lambda^{H'}, |\cdot|_{H'}, P^{H'}$ but possibly with $E_{H'} = \emptyset$, if $V_{H'}$ is a singleton. **STEP** is also allowed to return a **FAIL** message, while still outputting $H', \lambda^{H'}, |\cdot|_{H'}, P^{H'}$. The procedure is defined as follows (see Figure 2).

- Let $e_H = \{v_H, u_H\} \in E_H$ be such that $P_{e_H}^H = \min\{P_e^H; e \in E_H \text{ with } v_H \in e\}$. Below, e_H will be a.s. uniquely-defined, so ties may be broken arbitrarily if needed at this point.
- If $\Lambda_H := \sum_{e \in E_H; v_H \in e} \lambda_e^H < 1$ and $P_{e_H}^H > 1$, the algorithm will **FAIL**, but still continue.
- Let H' be the multi-graph obtained by identifying v_H and u_H , keeping the label u_H for the resulting vertex in $V_{H'}$.
- For $v \in V_H \setminus \{v_H, u_H\}$, we set $|v|_{H'} = |v|_H$, while we set $|u_H|_{H'} = |v_H|_H + |u_H|_H$.
- Let $\lambda_e^{H'} = \lambda_e^H$ and $P_e^{H'} = P_e^H$, if $v_H \notin e \in E_H$. Let $\lambda_{e'}^{H'} = \lambda_e^H$ and $P_{e'}^{H'} = P_e^H - P_{e_H}^H$ for $e \in E_H$ with $v_H \in e$ and $u_H \notin e$, where e' is the image of e through the quotient.




Figure 2: Illustration of **STEP** of Definition 2.6 applied to the multi-graph H above (with indices H suppressed) resulting in the multi-graph H' below. The thick red edge joining u and v is contracted. The size of u in H' is the sum of those of u and v in H .

Intuitively, the last step in the above procedure corresponds to letting the time elapse for the other edges adjacent to v_H .

Lemma 2.7 (Law of **STEP**). *Assume that, in the input of **STEP**, $(P_e^H)_{e \in E_H}$ are independent, exponential random variables with rates $(\lambda_e^H)_{e \in E_H}$. Then,*

$$\mathbb{P}(e_H = e) = \lambda_e^H / \Lambda_H \text{ for } e \in E_H \text{ with } v_H \in e, \quad (9)$$

$$\mathbb{P}(\text{FAIL}) \geq \max(0, 1 - \Lambda_H), \quad (10)$$

$$\text{FAIL is independent of } e_H, \quad (11)$$

$$\text{conditionally on } e_H \text{ and } \mathbb{1}_{\text{FAIL}}, (P_e^{H'})_{e \in E_{H'}} \text{ are also distributed as } \otimes_{e \in E_{H'}} \text{Exp}(\lambda_e^{H'}), \quad (12)$$

$$P_{e'}^{H'} \leq P_e^H, \text{ whenever } e' \text{ is the image of } e \text{ under the quotient mapping } H \text{ to } H'. \quad (13)$$

Proof. Firstly, e_H is clearly a.s. uniquely defined, since the exponential distribution is diffuse. The distribution of e_H is a standard property of Poisson processes. We have

$$\mathbb{P}(\text{FAIL}) = e^{-\Lambda_H} \cdot \mathbb{1}_{\Lambda_H < 1} \geq \max(0, 1 - \Lambda_H).$$

The fact that $(P_e^{H'})_{e \in E_{H'}}$ have the claimed conditional law follows from the memoryless property of exponential variables. Finally, the fact that $P_{e'}^{H'} \leq P_e^H$ is immediate from Definition 2.6. $\square$

Definition 2.8 (**EXPLORATION**). Let $(P_e)_{e \in E}$ be independent exponential random variables with rates $(\lambda_e)_{e \in E}$. **EXPLORATION** is the following exploration algorithm. Start with the multi-graph $H_0 = G$ with constant weight $|\cdot|_H \equiv 1$. Let v_G be given by **SELECT**. Apply **STEP** to obtain a multi-graph $H_1 = H'$. We repeat applying **SELECT** and **STEP** until we reach a multi-graph with exactly two vertices (which are necessarily labeled x and y). This defines a sequence of multi-graphs $(H_s)_{s=0}^S$. The output of **EXPLORATION** is the set $\Pi \subset E$ of edges which are mapped to the edges remaining in the final multi-graph H_S .

Note that, if, for any nonempty $A \subsetneq V$, $\|\partial A\|_\lambda \geq 1$, then, a.s., **EXPLORATION** does not fail. Moreover, the output Π is necessarily a cutset.

Remark 2.9 (Alternative construction). An important consequence of Lemma 2.7 is that it provides an alternative way to construct the sequence of multi-graphs $(H_t)_{t \geq 0}$ from **EXPLORATION**, together with the sequence of fail events without the use of the exponential random variables P_e . More precisely, we can alternate between using **SELECT** to pick a vertex v_{H_t} and then choose e_{H_t} according to (9). After the whole sequence H_t is fixed, one can decide whether each individual step t failed using independent coin tosses with parameters $e^{-\Lambda_{H_t}} \mathbb{1}_{\Lambda_{H_t} < 1}$. Note however, that this construction is not sufficient for the proof of our main theorem because it does not provide the crucial monotonicity (13).

Definition 2.10 (EXPLORATION_π). Let π be a cutset. Then EXPLORATION_π is the same as **EXPLORATION** but run with $\tilde{\lambda}$ given by $\tilde{\lambda}_e = 0$ (in which case $\tilde{P}_e = \infty$) for every $e \in \pi$ and $\tilde{\lambda}_e = \lambda_e$ otherwise. This generates a sequence of multi-graphs $(G_t)_{t=0}^T$, whose law is denoted by μ_π . Note that in this case the output is always π . We say that EXPLORATION_π succeeds if it does not **FAIL** at any step.

Lemma 2.11 (Cohesion implies success). *Let P_e be independent exponential random variables with rates $(\lambda_e)_{e \in E}$. Let $\omega = \{e \in E : P_e \leq 1\} \subset E$ (which is distributed as $\mathbb{P}_\lambda$). If a cutset π is ω -cohesive, then EXPLORATION_π succeeds a.s.*

Proof. Let P_e^t denote the exponential variable of the edge $e \in E$ after t **STEPS** of EXPLORATION_π , identifying the edge with its images under contraction until it becomes a self-loop (and is removed). By Lemma 2.7, $t \mapsto P_e^t$ is non-increasing and its initial value is at most 1 for $e \in \omega$. Therefore, we may **FAIL** only on steps when $v_H \notin \bigcup_{e \in \omega} e$. But then a failure could only be explained by $\Lambda_H < 1$, which would yield a set of edges F contradicting Definition 2.2. $\square$

Recall that $(G_t)_{t=0}^T$ is the sequence of multi-graphs obtained by EXPLORATION_π , and let $\partial_{G_t} v_{G_t}$ represent the set of edges $e \in E$ whose image in G_t under the quotient contains the selected vertex v_{G_t} (edges mapping to loops are not included). We write

$$r_t = \|\partial_{G_t} v_{G_t} \setminus \pi\|_\lambda, \quad b_t = \|\partial_{G_t} v_{G_t}\|_\lambda, \quad a_t = \|\partial_{G_t} v_{G_t} \cap \pi\|_\lambda, \quad (14)$$

so that $b_t = r_t + a_t$.

Lemma 2.12 (Law of Π). *Given a cutset π , recalling that Π denotes the cutset generated by **EXPLORATION**, we have*

$$\mathbb{P}[\Pi = \pi] = \mathbb{E}_{\mu_\pi} \left[\prod_t \frac{r_t}{b_t} \right]. \quad (15)$$

Proof. Consider a deterministic sequence of edges $e_0, e_1, \dots, e_{T-1}$ (necessarily not in π) that could be selected by EXPLORATION_π , resulting in the sequence of multi-graphs $g_0, g_1, \dots, g_T$ (necessarily with output π). The probability that this specific sequence is selected under **EXPLORATION** is

$$\prod_t \frac{\lambda_{e_t}}{b_t} = \prod_t \frac{r_t}{b_t} \prod_t \frac{\lambda_{e_t}}{r_t}, \quad (16)$$

the latter product being the probability of the trajectory under μ_π . Summing over all sequences concludes the proof. $\square$

By Lemma 2.7 (see also Remark 2.9) for any cutset π ,

$$\begin{aligned} \mathbb{P}(\text{EXPLORATION}_\pi \text{ succeeds}) &\stackrel{(10),(11)}{\leq} \mathbb{E}_{\mu_\pi} \left[\prod_t \min(r_t, 1) \right] = \mathbb{E}_{\mu_\pi} \left[\prod_t \left(\frac{r_t}{b_t} \right) / \left(\frac{\max(r_t, 1)}{b_t} \right) \right] \\ &\leq \frac{\mathbb{E}_{\mu_\pi} \left[\prod_t \frac{r_t}{b_t} \right]}{\text{ess inf}_{\mu_\pi} \prod_t \frac{\max(r_t, 1)}{b_t}} \stackrel{\text{Lemma 2.12}}{=} \frac{\mathbb{P}(\Pi = \pi)}{\text{ess inf}_{\mu_\pi} \prod_t \frac{\max(r_t, 1)}{b_t}}. \end{aligned} \quad (17)$$

In order to conclude, we need the following lemma adapted from [Cut26b], whose proof is spelled out in Appendix A for completeness.

Lemma 2.13 (Isoperimetric lower bound). *Recall the definition of Ψ^* from (7) and assume that $\phi_\lambda(1) \geq 1$. Then, for any cutset π , a.s. under μ_π , we have*

$$\prod_t \frac{\max(r_t, 1)}{b_t} \geq \exp(-\delta \|\pi\|_\lambda), \quad \delta := \left(\frac{24}{\Psi^*} + \frac{22}{\sqrt{\Psi^*}} + \frac{4 \log(e\phi_\lambda(1))}{\phi_\lambda(1)} \right)$$

Proof of Theorem 1.1. First note that by the assumption (3) we have $\min(\Psi^*, \phi_\lambda(1)) \geq \Psi > 25^2$, which readily implies $\delta \leq \varepsilon = 25/\sqrt{\Psi} < 1$. Combining (17) and Lemma 2.13, we obtain

$$e^{\varepsilon \|\pi\|_\lambda} \mathbb{P}(\Pi = \pi) \geq \frac{\mathbb{P}(\pi \cap \omega = \emptyset) \mathbb{P}(\text{EXPLORATION}_\pi \text{ succeeds})}{\mathbb{P}(\pi \cap \omega = \emptyset)}. \quad (18)$$

This expression can be estimated using the independence between the events $\{\pi \cap \omega = \emptyset\}$ and $\{\text{EXPLORATION}_\pi \text{ succeeds}\}$ with Lemma 2.11, giving

$$e^{\varepsilon \|\pi\|_\lambda} \mathbb{P}(\Pi = \pi) \geq e^{\|\pi\|_\lambda} \mathbb{P}_\lambda(\pi \text{ is } \omega\text{-cohesive}). \quad (19)$$

Using Lemma 2.4 and (19), we obtain

$$\begin{aligned} \mathbb{P}_\lambda(x \not\leftrightarrow y) &= \mathbb{P}_\lambda(\text{there exist an } \omega\text{-cohesive cutset } \pi) \leq \sum_\pi \mathbb{P}_\lambda(\pi \text{ is } \omega\text{-cohesive}) \\ &\leq \sum_\pi e^{-(1-\varepsilon)\|\pi\|_\lambda} \mathbb{P}(\Pi = \pi) \stackrel{(2)}{\leq} e^{-(1-\varepsilon)\chi_\lambda(x,y)}. \end{aligned}$$

This completes the proof of Theorem 1.1 for finite graphs.

Finally, we observe that the case where V is infinite follows readily from the case where V is finite. Indeed, for any finite subset $V' \subset V$, we can collapse all the vertices in $V \setminus V'$ into a single “ghost vertex” y . This naturally defines an associated finite weighted (multi-)graph $(V' \cup \{y\}, \lambda')$. By construction, one has $\mathbb{P}_{\lambda'}(x \leftrightarrow y) = \mathbb{P}_\lambda(x \leftrightarrow V \setminus V')$ for any $x \in V'$. Furthermore, one can check that $\psi_{\lambda'}^*(n) \geq \psi_\lambda(n)$ for all $n \in \{1, \dots, \lfloor \frac{|V'|+1}{2} \rfloor\}$ (note that this is where the relevance of (6) comes from) and $\phi_{\lambda'}(1) \geq \phi_\lambda(1) \geq (\sum_k 1/\psi_\lambda(2^k))^{-1}$. Therefore, we can apply the finite graph version of Theorem 1.1 for $(V' \cup \{y\}, \lambda')$ to deduce that

$$\mathbb{P}_{\lambda'}(x \leftrightarrow y) \geq 1 - e^{-(1-\varepsilon)\chi_{\lambda'}(x,y)} \leq 1 - e^{-(1-\varepsilon)\chi_\lambda(x)}.$$

This shows that $\mathbb{P}_\lambda(x \leftrightarrow V \setminus V') \geq 1 - e^{-(1-\varepsilon)\chi_\lambda(x)}$ for every finite $V' \subset V$ such that $x \in V'$. By taking an exhaustion of V , we obtain $\mathbb{P}_\lambda(x \leftrightarrow \infty) \geq 1 - e^{-(1-\varepsilon)\chi_\lambda(x)}$ as desired. $\square$

3 Proof of Theorem 1.6

Theorem 1.6 will follow readily from Theorem 1.1 and the following lemma.

Lemma 3.1 (Isoperimetric inequality). *If λ satisfies the assumptions of Theorem 1.6, then there exist non-negative constants $(C_M)_{M \geq 1}$ such that $C_M \rightarrow \infty$ as $M \rightarrow \infty$ and $\psi_{\lambda^M}(n) \geq C_M \sqrt{n}$ for all $M, n \geq 1$.*

We start by assuming the validity of the above lemma.

Proof of Theorem 1.6. By Lemma 3.1, we can find M such that $\psi_{\lambda^M}(n) \geq 2500n^{1/2}$. For such M we have $\sum_{k \geq 0} \frac{1}{\psi_{\lambda^M}(2^k)} \leq \frac{1}{2500} \sum_{k \geq 0} 2^{-k/2} = \frac{1}{2500} \left(\frac{\sqrt{2}}{\sqrt{2}-1} \right) < \frac{1}{625}$ and the result follows from Theorem 1.1. $\square$

Finally, we provide the missing ingredient that was postponed.

Proof of Lemma 3.1. We distinguish two cases.

Case 1. Assume that there exists $x \in \mathbb{Z}^d$ such that $\sum_{z \in \mathbb{Z}^d \cap (\mathbb{R}x)} \lambda_z = \infty$. Intuitively speaking, this corresponds to the case that, even though λ is not supported on a line, there exists a line over which λ is not summable. Then we may restrict attention to the case $d = 2$ by setting $\lambda_y = 0$ for all $y \in \mathbb{Z}^d \setminus (\mathbb{R}x)$ except one y (and its opposite, $-y$) such that $\lambda_y > 0$. Up to applying a linear transformation, we may assume x and y to be the first two unit vectors of $\mathbb{Z}^2$.

We next use a folklore proof to establish an isoperimetric inequality. Let $A \subset \mathbb{Z}^2$ and observe that

$$|A| \leq |P_1(A)| \cdot |P_2(A)|, \quad (20)$$

where $P_1(A) = \{a \in \mathbb{Z} : A \cap (\{a\} \times \mathbb{Z}) \neq \emptyset\}$ and $P_2(A) = \{b \in \mathbb{Z} : A \cap (\mathbb{Z} \times \{b\}) \neq \emptyset\}$ are the projections of A . Fix an integer $M \geq 1$ and set $\Lambda_M := \sum_{z=1}^M \lambda_{(z,0)}$ and $C_M = 4\sqrt{\Lambda_M \lambda_y}$. Looking at the right-most, left-most, top-most and bottom-most sites of A on each horizontal or vertical line, we have

$$\|\partial A\|_{\Lambda_M} \geq 2|P_1(A)|\lambda_y + 2|P_2(A)|\Lambda_M \geq C_M \sqrt{|A|},$$

the last inequality following from (20) and $(a+b)/2 \geq \sqrt{ab}$.

Case 2. Assume that, for every $x \in \mathbb{Z}^d$, we have $\sum_{y \in \mathbb{Z}^d \cap (\mathbb{R}x)} \lambda_y < \infty$. We start by noticing that $\|\partial A\|_{\lambda}$ is linear in λ . Therefore, it is enough to find a decomposition $\lambda = \sum_{i \geq 1} \lambda^i$, with each λ^i having finite support and such that, for every finite set $A \subset \mathbb{Z}^d$, we have $\|\partial A\|_{\lambda^i} \geq c_i \sqrt{|A|}$ with $C_M := \sum_{i \leq M} c_i \rightarrow \infty$ as $M \rightarrow \infty$.

The decomposition that we construct will be such that each λ^i satisfies

$$\lambda_z^i > 0 \text{ if and only if } z \in \{x_i, y_i, -x_i, -y_i\}, \quad (21)$$

$$x_i \text{ and } y_i \text{ are linearly independent and} \quad (22)$$

$$\lambda_{x_i}^i = \lambda_{y_i}^i. \quad (23)$$

We start by defining λ^1 and for this, pick some $x_1 \in \mathbb{Z}^d$ attaining the largest value of λ_x . Recalling $\sum_{y \in \mathbb{Z}^d \setminus (\mathbb{R}x)} \lambda_y = \infty$, take y_1 attaining the largest λ_y in $\mathbb{Z}^d \setminus (\mathbb{R}x)$. Finally, set $\lambda_{x_1}^1 = \lambda_{y_1}^1 = \lambda_{y_1}$ and zero otherwise. At this point, we just repeat the above procedure with λ replaced by the remainder $\lambda - \lambda^1$, inductively. This defines all the λ^i 's and the fact that $\lambda = \sum_i \lambda^i$ follows from the construction, together with the hypothesis of Case 2. We are now left with proving that for every finite set $A \subset \mathbb{Z}^d$, we have $\|\partial A\|_{\lambda^i} \geq c_i \sqrt{|A|}$ with non-summable c_i 's.

Notice that, for each $i \geq 1$, the proof of the previous case gives that the boundary of A corresponding to translates of $\{\pm x_i, \pm y_i\}$ alone satisfies $\|\partial A\|_{\lambda^i} \geq 4\sqrt{\lambda_{x_i}^i \lambda_{y_i}^i |A|} = \lambda_{y_i}^i 4\sqrt{|A|}$, for any A contained in the sub-lattice $\langle x_i, y_i \rangle_{\mathbb{Z}}$, so in fact the same holds for any $A \subset \mathbb{Z}^d$, since $\sqrt{a} + \sqrt{b} \geq \sqrt{a+b}$ for $a, b \geq 0$. The fact that the values $\lambda_{y_i}^i$ are not summable follows from $\sum_i \lambda_{y_i}^i = \frac{1}{2} \sum_x \lambda_x = \infty$. $\square$

4 Proof of Theorem 1.7

Theorem 1.7 will follow easily from Theorem 1.1 and the following lemma.

Lemma 4.1 (Isoperimetric inequality). *There exists a universal constant $c_1 > 0$ such that the following holds. If $G = (V, E)$ is a transitive graph of degree Δ and superlinear growth, then $\psi_{1_E}(n) \geq c_1 \sqrt{\Delta n}$ for all $n \geq 1$.*

Proof. By [TT24, Corollary 7.1] for $d = 1$, there exists a universal constant $\varepsilon_1 > 0$ such that every transitive graph G of degree Δ and superlinear growth satisfies $|B_n| \geq \varepsilon_1 \Delta n^2$ for all $n \geq 1$. By [TT20, Proposition 5.1], this implies the isoperimetric inequality $|\partial S| \geq \frac{\sqrt{2\varepsilon_1 \Delta}}{12} |S|^{1/2}$. Setting $c_1 = \sqrt{\varepsilon_1/72}$ concludes the proof. $\square$

Proof of Theorem 1.7. First, we observe that, for small n , we can easily obtain a better bound on ψ_{1_E} . If $S \subset V$ is such that $|S| \leq \Delta/2$, then each vertex $x \in S$ has at least $\Delta/2$ neighbours outside S and therefore $\|\partial S\|_{1_E} = |\partial S| \geq \frac{\Delta}{2} |S|$. In particular, $\psi_{1_E}(n) \geq \frac{\Delta}{2} n$ for all $n \leq \Delta/4$ (this is where we use the fact that we work with ψ instead of ϕ). Let $\lambda = \frac{C}{\Delta} \mathbb{1}_E$ for some large constant C to be chosen later.

Obviously, $\psi_\lambda = \frac{C}{\Delta} \psi_{1_E}$, so we have $\psi_\lambda(n) \geq \frac{C}{2}n$ for $n \leq \Delta/4$, and $\psi_\lambda(n) \geq c_1 C \sqrt{n/\Delta}$ for $n \geq \Delta/4$, by Lemma 4.1. Therefore,

$$\sum_{k \geq 0} \frac{1}{\psi_\lambda(2^k)} \leq \sum_{k=0}^{\lfloor \log_2(\Delta/4) \rfloor} \frac{2}{C 2^k} + \frac{\sqrt{\Delta}}{c_1 C} \sum_{k \geq \lfloor \log_2(\Delta/4) \rfloor + 1} \frac{1}{2^{k/2}} \leq \frac{4}{C} + \frac{2}{c_1 C} \left(\frac{\sqrt{2}}{\sqrt{2}-1} \right) < \frac{4(1+2c_1^{-1})}{C}.$$

By setting $C := 2500(1+2c_1^{-1})$, Theorem 1.1 gives $\mathbb{P}_\lambda[0 \leftrightarrow \infty] > 0$. Recalling the parametrization $p_e = 1 - e^{-\lambda_e}$, this readily implies that $p_c^{\text{bond}}(G) \leq 1 - e^{-C/\Delta}$. $\square$

Remark 4.2 (Optimality). The bound $\phi(n) \geq c\sqrt{\Delta n}$ from Lemma 4.1 is optimal up to constant. Indeed, consider the transitive (in fact, also Cayley) graph G with vertex set $V = \mathbb{Z}^2 \times (\mathbb{Z}/k\mathbb{Z})$ and edge set $E = \{(x, j), (y, k)\} : \text{either } x = y \text{ or } j = k \text{ and } y = x \pm e_i, i \in \{1, 2\}\}$, where $e_1 = (1, 0)$ and $e_2 = (0, 1)$. In this case, one has $\Delta = k + 3$ and the sets $S_n = [-\sqrt{n/k}, \sqrt{n/k}]^2 \times (\mathbb{Z}/k\mathbb{Z})$ satisfy $|S_n| \asymp n$ and $|\partial S_n| \asymp \sqrt{kn}$.

Acknowledgements

This project started during the XXIX Brazilian School of Probability in Rio de Janeiro, where Bernardo de Lima brought the truncation problem to our attention and later also shared helpful bibliographic pointers. We would like to thank him as well as the organizers of the event. We also thank Lyuben Lichev and Matthew Tointon for helpful reference pointers. This work was carried out in part during a visit by FS to IMPA, whose hospitality he gratefully acknowledges. AT was supported by the grants “Projeto Universal” (406250/2016-2) from CNPq, “Produtividade em Pesquisa” (304437/2018-2) from CNPq and “Cientista do Nosso Estado” (204.377/2024) from FAPERJ. For the purpose of Open Access, a CC-BY public copyright licence has been applied by the authors to the present document and will be applied to all subsequent versions up to the Author Accepted Manuscript arising from this submission.

Declaration on the use of AI. No artificial intelligence tools were used in the development of the mathematical ideas, arguments, or proofs presented in this paper, nor were they used for drafting or writing the manuscript.

A Proof of Lemma 2.13

This appendix entirely follows [Cut26b] up to minor adjustments and is presented only to keep the proof self-contained. We introduce the decreasing function

$$f(u) = \log(eu)/u, \quad u \geq 1.$$

Claim A.1. Let l be a positive integer and $(p_i)_{i=0}^l$ and $(b_i)_{i=0}^l$ be sequences of positive real numbers. Assume that $b_i \geq \max(1, p_i)$ for all i and $b_i \geq b_{i-1} + p_{i-1}/2$ for $i \in \{1, \dots, l\}$. Then

$$\sum_{i=0}^l f(b_i) \leq \frac{3\sqrt{5}}{\sqrt{2}} \left(\sum_{i=0}^l \frac{1}{p_i} \right)^{1/2}.$$

Proof. For any $u \in J_i := [b_i, b_i + p_i/2]$, we have $f(u) \geq f(3b_i/2) \geq \frac{2}{3}f(b_i)$. Therefore,

$$p_i f(b_i)^2 \leq \frac{9}{2} \int_{J_i} f^2.$$

Since the intervals J_i are disjoint, this yields

$$\sum_{i=0}^l p_i f(b_i)^2 \leq \frac{9}{2} \int_1^\infty f^2 = \frac{45}{2}.$$

The proof then follows by the Cauchy–Schwarz inequality $(\sum x_i y_i)^2 \leq (\sum x_i^2)(\sum y_i^2)$ with $x_i = \sqrt{p_i} f(b_i)$ and $y_i = 1/\sqrt{p_i}$. $\square$




Figure 3: Illustration of the proof of Lemma 2.13. The cutset π separating x and y is shown in the middle. The vertex set in V corresponding to $v_{G_{t_i}}$ and containing z is hatched. The shaded blue region is $v_{G_{t_{i+1}}} \setminus v_{G_{t_i}}$. Their common boundary is thickened and has weight at most r_{t_i} (which is the total boundary of the hatched region, excluding π). The boundary of $v_{G_{t_{i+1}}}$ in $G_{t_{i+1}}$ is dashed and has weight $r_{t_{i+1}}$.

Proof of Lemma 2.13. Recall a_t, b_t, r_t from (14) and note that we always have $b_t \geq \phi_\lambda(1) \geq 1$. Let $m_t = |v_{G_t}|_{G_t}$ be the size of the vertex selected at step t in **EXPLORATION** $_\pi$.

Let us fix $z \in V$ such that there exists $e \in \pi$ for which $z \in e$. Let $(t_i)_{i=0}^l$ be the times such that $v_{G_{t_i}} = z$.² For all $i \in \{0, \dots, l-1\}$, let

$$p_i = \min \left(\psi_\lambda^*(2^{\lfloor \log_2 m_{t_i} \rfloor}), \psi_\lambda^*(2^{\lfloor \log_2 (m_{t_{i+1}} - m_{t_i}) \rfloor}) \right).$$

Recalling (6), we have $b_{t_i} \geq p_i$. Moreover, by Definition 2.5, $m_{t_i} \geq 2m_{t_{i-1}} \geq 2^i$ for all $i \in \{1, \dots, l\}$, so the integers $\lfloor \log_2 m_{t_i} \rfloor$ are all distinct. Consequently, $\lfloor \log_2 (m_{t_{i+1}} - m_{t_i}) \rfloor$ can take the same integer value for at most two indices i , since $\lfloor \log_2 m_{t_i} \rfloor \leq \lfloor \log_2 (m_{t_{i+1}} - m_{t_i}) \rfloor \leq \lfloor \log_2 m_{t_{i+1}} \rfloor$. Hence,

$$\sum_{i=0}^{l-1} \frac{1}{p_i} \leq 3 \sum_{k \geq 0} \frac{1}{\psi_\lambda^*(2^k)} = \frac{3}{\Psi^*}. \quad (24)$$

Then each edge of the boundary of the vertex set corresponding to $v_{G_{t_{i+1}}} \setminus v_{G_{t_i}}$ in the multi-graph G_{t_i} (which has total vertex weight $m_{t_{i+1}} - m_{t_i}$) fall in one of the categories (see Figure 3):

- edges in G_{t_i} linking $v_{G_{t_i}}$ and $v_{G_{t_{i+1}}} \setminus v_{G_{t_i}}$ or
- edges in $G_{t_{i+1}} \setminus \pi$ incident to $v_{G_{t_{i+1}}}$ or
- edges in π incident to $v_{G_{t_{i+1}}} \setminus v_{G_{t_i}}$ (but not to $v_{G_{t_i}}$).

The total weight of the above three types of edges is at most r_{t_i} , $r_{t_{i+1}}$ and $a_{t_{i+1}} - a_{t_i}$ respectively. Therefore,

$$b_{t_{i+1}} - b_{t_i} \stackrel{(14)}{=} (r_{t_i} + r_{t_{i+1}} + (a_{t_{i+1}} - a_{t_i})) - 2r_{t_i} \geq \psi_\lambda^*(2^{\lfloor \log_2 (m_{t_{i+1}} - m_{t_i}) \rfloor}) - 2r_{t_i} \geq p_i - 2r_{t_i}.$$

We consider four types of indices i (recall that $b_{t_i} \geq p_i \geq \phi_\lambda(1) \geq 1$):

- If $i < l$ and $r_{t_i} \geq p_i/4$, then $\log \frac{b_{t_i}}{\max(r_{t_i}, 1)} \leq \log \frac{b_{t_i}}{r_{t_i}} = \log \left(1 + \frac{a_{t_i}}{r_{t_i}} \right) \leq \frac{a_{t_i}}{r_{t_i}} \leq 4 \frac{a_{t_i}}{p_i}$.
- If $i < l$ and $r_{t_i} < p_i/4$, then $\log \frac{b_{t_i}}{\max(r_{t_i}, 1)} \leq \log b_{t_i} \leq \log(eb_{t_i}) \leq \frac{4}{3} a_{t_i} f(b_{t_i})$ and $b_{t_{i+1}} \geq b_{t_i} + p_i/2$.
- If $i = l$, and $r_{t_l} \geq b_{t_l}/2$, then $\log \frac{b_{t_l}}{\max(r_{t_l}, 1)} \leq \log \frac{b_{t_l}}{r_{t_l}} \leq 2 \frac{a_{t_l}}{b_{t_l}} \leq 2a_{t_l} f(b_{t_l}) \leq 2a_{t_l} f(\phi_\lambda(1))$.
- If $i = l$ and $r_{t_l} < b_{t_l}/2$, then $\log \frac{b_{t_l}}{\max(r_{t_l}, 1)} \leq \log b_{t_l} \leq \frac{3}{2} a_{t_l} f(b_{t_l}) \leq 2a_{t_l} f(\phi_\lambda(1))$.

²We abused the notation by writing $v_{G_{t_i}} = z$ to represent that z is mapped to $v_{G_{t_i}}$ via the natural quotient.

Putting these bounds together and using Claim A.1, we obtain

$$\sum_{i=0}^l \frac{1}{a_{t_i}} \log \frac{b_{t_i}}{\max(r_{t_i}, 1)} \leq 4 \sum_{i=0}^{l-1} p_i^{-1} + 2\sqrt{10} \left(\sum_{i=0}^{l-1} p_i^{-1} \right)^{1/2} + 2f(\phi_\lambda(1)) \stackrel{(24)}{\leq} \frac{\delta}{2}. \quad (25)$$

Hence, recalling that $a_t = \sum_{e \in \pi, v_{G_t} \in e} \lambda_e$, we get

$$\sum_t \log \frac{b_t}{\max(r_t, 1)} = \sum_t \frac{1}{a_t} \log \frac{b_t}{\max(r_t, 1)} \sum_{e \in \pi} \lambda_e \sum_{z \in e} \mathbb{1}_{v_{G_t}=z} \leq \delta \sum_{e \in \pi} \lambda_e = \delta \|\pi\|_\lambda,$$

concluding the proof of Lemma 2.13. $\square$

References

- [AHdLV17] Caio Alves, Marcelo R. Hilário, Bernardo N.B. de Lima, and Daniel Valesin. A note on truncated long-range percolation with heavy tails on oriented graphs. *Journal of Statistical Physics*, 169(5):972–980, 2017.
- [Bäu24] Johannes Bäumlér. Truncation of long-range percolation with non-summable interactions in dimensions $d \geq 3$. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques (to appear)*, available at [arXiv:2410.00303](https://arxiv.org/abs/2410.00303), 2024.
- [BJR07] Béla Bollobás, Svante Janson, and Oliver Riordan. Spread-out percolation in $\mathbb{R}^d$. *Random Structures & Algorithms*, 31(2):239–246, 2007.
- [BS96] Itai Benjamini and Oded Schramm. Percolation beyond $\mathbf{Z}^d$, many questions and a few answers. *Electronic Communications in Probability*, 1:71–82, 1996.
- [CdL22] Alberto M. Campos and Bernardo N.B. de Lima. Truncation of long-range percolation models with square non-summable interactions. *ALEA, Latin American Journal of Probability and Mathematical Statistics*, 19(1):1025–1033, 2022.
- [Cut26a] Isoperimetric dimension > 1 implies $p_c < 1$. *AI generated*, available at <https://philipeaso.com/s/paper-t7m6.pdf>, 2026.
- [Cut26b] Addendum to “Isoperimetric dimension > 1 implies $p_c < 1$ ”: More careful bookkeeping. *AI generated*, available at <https://philipeaso.com/s/addbeterwak-pdf.pdf>, 2026.
- [DGRSY20] 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, 12 2020.
- [DEKK26] Sahar Diskin, Joshua Erde, Mihyun Kang, and Michael Krivelevich. Percolation through isoperimetry. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, 62(2):900–918, 2026.
- [DG25] Sahar Diskin and Anna Geisler. Minimum degree k and k -connectedness usually arrive together. *Discrete Mathematics*, 348(6):114453, 2025.
- [dLS08] Bernardo N.B. de Lima and Artëm Sapozhnikov. On the truncated long-range percolation on $\mathbb{Z}^2$. *Journal of Applied Probability*, 45(1):287–291, 2008.
- [EH23] Philip Easo and Tom Hutchcroft. The critical percolation probability is local. *Preprint*, available at [arXiv:2310.10983](https://arxiv.org/abs/2310.10983), 2023.
- [EST25] Philip Easo, Franco Severo, and Vincent Tassion. Counting minimal cutsets and $p_c < 1$. *Forum of Mathematics, Pi*, 13:e23, 2025.

- [FdL06] Sacha Friedli and Bernardo N.B. de Lima. On the truncation of systems with non-summable interactions. *Journal of Statistical Physics*, 122(6):1215–1236, 2006.
- [FdLS04] Sacha Friedli, Bernardo N.B. de Lima, and Vladas Sidoravicius. On long range percolation with heavy tails. *Electronic Communications in Probability*, 9:175 – 177, 2004.
- [GS98] Geoffrey R. Grimmett and Alan M. Stacey. Critical probabilities for site and bond percolation models. *The Annals of Probability*, 26(4):1788–1812, 1998.
- [Joo15] Felix Joos. Random subgraphs in sparse graphs. *SIAM Journal on Discrete Mathematics*, 29(4):2350–2360, 2015.
- [Kar93] David R. Karger. Global min-cuts in RNC, and other ramifications of a simple min-cut algorithm. In *SODA*, volume 93, pages 21–30, 1993.
- [Kru56] Joseph B. Kruskal. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical Society*, 7(1):48–50, 1956.
- [MSV01] Mikhail Menshikov, Vladas Sidoravicius, and Marina Vachkovskaia. A note on two-dimensional truncated long-range percolation. *Advances in Applied Probability*, 33(4):912–929, 2001.
- [Pei36] Rudolf Peierls. On Ising’s model of ferromagnetism. *Mathematical Proceedings of the Cambridge Philosophical Society*, 32:477–481, 1936.
- [Pen93] Mathew D. Penrose. On the spread-out limit for bond and continuum percolation. *The Annals of Applied Probability*, 3(1):253–276, 1993.
- [PS23] Christoforos Panagiotis and Franco Severo. Gap at 1 for the percolation threshold of Cayley graphs. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, 59(3):1248 – 1258, 2023.
- [SSV99] Vladas Sidoravicius, Donatas Surgailis, and Maria E. Vares. On the truncated anisotropic long-range percolation on $\mathbb{Z}^2$. *Stochastic Processes and their Applications*, 81(2):337–349, 1999.
- [ST24] Panagiotis Spanos and Matthew Tointon. Spread-out percolation on transitive graphs of polynomial growth. *Preprint*, available at arXiv:2404.17262, 2024.
- [TT20] Romain Tessera and Matthew Tointon. Sharp relations between volume growth, isoperimetry and escape probability in vertex-transitive graphs. *Preprint*, available at arXiv:2001.01467, 2020.
- [TT24] Romain Tessera and Matthew Tointon. Balls in groups: volume, structure and growth. *Preprint*, available at arXiv:2403.02485, 2024.
- [vEdLV16] Aernout C.D. van Enter, Bernardo N.B. de Lima, and Daniel Valesin. Truncated long-range percolation on oriented graphs. *Journal of Statistical Physics*, 164(1):166–173, 2016.