

Polynomial growth in degree-dependent first passage percolation on spatial random graphs

Júlia Komjáthy* John Lapinskas† Johannes Lengler‡
Ulysse Schaller§

Abstract

In this paper we study a version of (non-Markovian) first passage percolation on graphs, where the transmission time between two connected vertices is non-iid, but increases by a *penalty factor* polynomial in their expected degrees. Based on the exponent of the penalty-polynomial, this makes it increasingly harder to transmit to and from high-degree vertices. This choice is motivated by awareness or time-limitations. For the iid part of the transmission times we allow any nonnegative distribution with regularly varying behaviour at 0. For the underlying graph models we choose spatial random graphs that have power-law degree distributions, so that the effect of the penalisation becomes visible: (finite and infinite) Geometric Inhomogeneous Random Graphs, and Scale-Free Percolation. In these spatial models, the connection probability between two vertices depends on their spatial distance and on their expected degrees. We prove that upon increasing the penalty exponent, the transmission time between two far away vertices x, y sweeps through four universal phases even for a single underlying graph: explosive (tight transmission times), polylogarithmic, *polynomial but sublinear* ($|x - y|^{\eta_0 + o(1)}$ for an explicit $\eta_0 < 1$), and *linear* ($\Theta(|x - y|)$) in their Euclidean distance. Further, none of these phases are restricted to phase boundaries, and those are non-trivial in the main model parameters: the tail of the degree-distribution, a long-range parameter, and the exponent of regular variation of the iid part of the transmission times. In this paper we present proofs of lower bounds for the latter two phases and the upper bound for the linear phase. These complement the matching upper bounds for the polynomial regime in our companion paper.

Keywords: first-passage percolation; intrinsic distances; scale-free spatial random graphs; geometric inhomogeneous random graphs; scale-free percolation; hyperbolic random graphs.

MSC2020 subject classifications: 05C82; 60K35; 60K50; 82B43; 91D30; 91D25.

Submitted to EJP on October 5, 2023, final version accepted on September 24, 2024.

*Delft University of Technology, Netherlands. E-mail: j.komjathy@tudelft.nl

†University of Bristol, United Kingdom. E-mail: john.lapinskas@bristol.ac.uk

‡ETH Zürich, Switzerland. E-mail: johannes.lengler@inf.ethz.ch

§ETH Zürich, Switzerland. U.S. was supported by the Swiss National Science Foundation [grant number 200021_192079]. E-mail: ulyссе.schaller@inf.ethz.ch

1 Introduction

First passage percolation (FPP) is a natural way to understand geodesics in random metric spaces. In this paper, we investigate a natural extension of Markovian and non-Markovian first passage percolation, where the transmission time of the first passage percolation process across an edge depends on its direct surroundings in the underlying graph, in particular on the expected degrees of the sending and receiving vertex [32]. We set the transmission time through each edge uv as $L_{uv}(W_u W_v)^\mu$, where L_{uv} is a non-negative iid factor with distribution that varies regularly near 0 with exponent β , W_u, W_v are constant multiples of the expected degrees of the vertices u, v , and μ is the penalty exponent. We then let the process run according to the standard rules of first passage percolation: transmission time between x and y is the total sum of transmission time on edges on a path between x, y , minimised over all paths between x and y . We abbreviate this process as 1-FPP for short. For the underlying graph models we choose spatial models with power-law degree distributions, i.e., with highly varying degrees, so that the dependence of the transmissions on the local surroundings causes non-negligible effects. The dependence caused via setting $\mu \geq 0$ is so that the transmission time from and towards high-degree vertices is slowed down by a polynomial of their expected degrees, which makes it harder, but not impossible, to transmit to and from “superspreaders”. The 1-FPP model is inspired by degree-dependent bond percolation [29] and by topology-biased random walks [11, 20, 36, 41, 46], in which the transition probabilities from a vertex depend on the degrees of its neighbours. Those works also assume a polynomial dependence on the degrees.

This new one-dependent FPP process has several interesting features. It reveals topological features of the underlying graphs hidden from classical versions. Classical FPP tend to strongly depend on the highest degree vertices and shows the ‘explosion’ phenomenon on spatial graphs with highly varying degrees for *all* regularly varying transmission-time distributions, i.e., the process reaches infinitely many vertices in finite time [34]. However, with 1-FPP explosion can be stopped efficiently by increasing the exponent μ of the penalty polynomial above a given value [32], thus, only changing the dynamics of the process to/from outlier vertices but keeping the underlying graph fully intact. We study what happens in the sub-explosive regime. In a sequence of papers we prove that as the penalty exponent increases, the transmission time between two far away vertices x, y in the infinite component sweeps through four universal phases:

- (i) it converges to a limiting a.s. finite random variable, i.e., it grows *explosively*.
- (ii) it grows *polylogarithmically* with the Euclidean distance as $\sim (\log |x - y|)^{\Delta_0 + o(1)}$ for some $\Delta_0 \geq 1$;

Phase (i) was our previous work [32]. We treat phase (ii) and the upper bound for phase (iii) in the companion paper [33]. The focus of this paper is to prove the *lower bounds* to the following phases (iii)-(iv); and the *upper bound to phase (iv)*:

- (iii) it grows *polynomially* as $|x - y|^{\eta_0 + o(1)}$ for an explicitly given $\eta_0 < 1$;
- (iv) it grows *linearly* proportionally to $|x - y|$.

See also Table 1. Possibly the most interesting phase here is (iii): sublinear polynomial transmission times imply polynomial intrinsic ball-growth *faster than the dimension* of the underlying space. This is rare in spatial graph models, and shows that 1-FPP cannot simply be mapped and studied as a simpler process (like graph distances) on a spatial graph with different parameters. Indeed, in this paper we show that in the polynomial regime *all* long edges around x, y carry polynomial transmission times in the distance

they bridge; thus the shortest path between x, y must contain edges of polynomial cost, resulting in the lower bound $|x - y|^{\eta_0 - o(1)}$ with an explicit $\eta_0 < 1$. This is hand-in-hand with the constructive proof we give as a matching upper bound in [33], thus our result in phase (iii) is sharp. The $o(1)$ additive factor in the exponent is due to our current quite general assumptions on the iid factor L and on the connection probability, but could be possibly sharpened by imposing stronger assumptions there.

Table 1: Summary of our results. In 1-FPP, the transmission time on edge xy is $L_{xy}(W_x W_y)^\mu$ where W_x, W_y are constant multiples of the expected degrees of the vertices x, y , and L_{xy} is an iid random variable with a regularly varying distribution function near 0 with exponent $\beta \in (0, \infty]$. The parameter $\tau \in (2, 3)$ is the exponent of the power-law degree distribution, the underlying graph has doubly-logarithmic graph distances. The transmission time $d_C(0, x)$ between 0 and a far vertex x passes through four different phases as μ , the penalty exponent, increases. When the long-range parameter $\alpha \in (1, 2)$, long edges between low-degree vertices cause polylogarithmic transmission times, so increasing μ stops explosion but it has no further effect. When $\alpha > 2$, these edges are absent from the graph. Here there are three sub-explosive phases. In this paper we focus on the polynomial phases: we prove the lower bound for the strict polynomial phase, (with an explicit η_0 in (1.7)). In particular, all long edges near $0, x$ have polynomial transmission times in the distance they bridge. We also prove both the upper and lower bound for the linear distance regime, and the lower bound much more generally when $\tau > 3$. All regimes hold on proper intervals for μ .

Graph param.	1-FPP parameters	Behaviour of 1-FPP transmission times
Weak decay: $\tau \in (2, 3)$ $\alpha \in (1, 2)$	$\mu < \frac{3-\tau}{2\beta}$	Explosive: $d_C(0, x) = \Theta(1)$
	$\mu > \frac{3-\tau}{2\beta}$	Polylogarithmic: $d_C(0, x) = (\log x)^{\Delta_0 + o(1)}, \Delta_0 > 1$
Strong decay: $\tau \in (2, 3)$ $\alpha > 2$	$\mu < \frac{3-\tau}{2\beta}$	Explosive: $d_C(0, x) = \Theta(1)$
	$\mu \in \left(\frac{3-\tau}{2\beta}, \frac{3-\tau}{\beta}\right)$	Polylogarithmic: $d_C(0, x) = (\log x)^{\Delta_0 + o(1)}, \Delta_0 > 1$
	$\mu \in \left(\frac{3-\tau}{\beta}, \frac{3-\tau}{\min\{\beta, d(\alpha-2)\}} + \frac{1}{d}\right)$	Strict Polynomial: $d_C(0, x) = x ^{\eta_0 \pm o(1)}, \eta_0 < 1$
	$\mu > \frac{3-\tau}{\min\{\beta, d(\alpha-2)\}} + \frac{1}{d}$	Linear: $d \geq 2 : d_C(0, x) = \Theta(x)$ $d = 1 : \kappa x \leq d_C(0, x) \leq x ^{1+o(1)}$
Non-scale-free: $\tau > 3, \alpha > 2$	$\mathbb{P}(L > 0) = 1, \mu \geq 0$	Linear: $d_C(0, x) \geq \kappa x $

In phase (iv), we prove that the transmission time between two vertices x and y is between $\kappa_1|x - y|$ and $\kappa_2|x - y|$ for two constants $\kappa_1, \kappa_2 > 0$. The lower bound is valid in all dimensions and the upper bound valid in dimension at least 2. In dimension 1 we prove an upper bound of $|x - y|^{1+o(1)}$ in the accompanying paper [33]. Nevertheless, near-linear distances in dimension 1 are somewhat surprising, since generally long-range spatial models either do not percolate in dimension 1 or show shorter distances [9, 16, 42, 44].

An important message of our results is that the change in the phases can be obtained by *changing only the dynamics* of the process, (i.e., increasing the penalty exponent μ), while keeping the underlying graph intact. This rich behaviour arises despite the fact that the underlying graphs have doubly-logarithmic graph distances.

Growth phases of FPP and graph distances in other models. By contrast, in other models studied so far the behaviour of intrinsic distances is less rich, and the strict polynomial phase (iii) is *absent* or restricted only to phase transition boundaries. Sparse spatial graphs with finite-variance degrees (e.g. percolation, long-range percolation, random geometric graphs etc.) generally show linear graph distances/transmission times (phase (iv)) when long edges are absent [3, 4, 15, 40], or polylogarithmic distances (phase (ii)) when long edges are present [9, 10, 28]. In spatial models with infinite-variance degrees, classical FPP explodes for all edge transmission-time distributions L_{xy} with regularly varying behaviour near 0. For the process to be sub-explosive, L has to be at least doubly-exponentially flat near 0, and then transmission times are at most the same order as the doubly-logarithmic graph distances [34, 45]; in particular, there is no analogue of phases (ii)–(iv). The same happens on non-spatial models with infinite variance degrees [1, 30]. On sparse *non-spatial* graph models with finite-variance degrees, classical FPP universally shows Malthusian behaviour, i.e., exponential growth [7], (for any distribution $L > 0$ almost surely). Transmission times between two uniformly chosen vertices are then logarithmic in the graph size (phase (ii)). There is some work on one-dependent FPP on non-spatial graphs: the process either explodes [43], with the same criterion for explosion as for spatial graphs in [32], or becomes Malthusian [23]. So only phases (i) and (ii) can occur.

We generally see that the strict polynomial phase is missing from models. So far, the only known graph model that exhibits this behaviour is long-range percolation on its phase boundary $\alpha = 2$, which is recently proven to show polynomial distances [5]. Even among degenerate models (i.e., with a complete graph as the underlying graph), long-range first passage percolation [14] is the only other model where a similarly rich set of phases is proven to occur. Long-range FPP uses exponential transmission times that depend on the Euclidean distance of the edge. The Markov property and the fact that there is no fixed underlying graph makes proof techniques in [14] unavailable in the context of 1-FPP on fixed graphs. We refer the reader to the discussion in [33] for details, and also for a thorough discussion of other related literature.

Summarising, one-dependent FPP is the first process that displays a full interpolation between the four phases on a *single non-degenerate graph model*. Moreover, the phase boundaries for one-dependent FPP depend non-trivially on the main model parameters: the degree power-law exponent τ , the parameter α controlling the prevalence of long-range edges, and the behaviour of L_{xy} near 0 characterised by β , see Table 1 for our results. We leave the investigation of any phase-boundary cases in the parameter space, i.e., where μ takes the value that separates two phases of growth, for future work.

Motivation. Arguably, 1-FPP with its variety of regimes reflects spreading processes on real networks better than simpler models (e.g., iid FPP), since the speed of dissemination in social networks can range from extremely fast (e.g., when news or memes spread through social media) to rather slow, geometry-dominated patterns (e.g., as diseases spread through regions and countries), with change only in the dynamics of the process but not (much) in the underlying network. The choice of degree-dependent transmission times is not just motivated by the related work mentioned above, but comes directly from applications. Actual contacts do not scale linearly with network connectivity due to limited time or awareness [22], and this type of penalty has been used to model the sublinear impact of superspreaders as a function of contacts [25, 31, 39].

Two papers, two techniques. In our companion paper [33] we develop methods

of a very different flavour that allow us to construct paths for the upper bounds for all sub-explosive phases at once, at the cost of the $o(1)$ factor in the exponent. This construction works in the quenched setting, i.e., with revealed vertex set, then uses a multi-scale argument to construct the polylogarithmic/polynomial path. We think that both proof techniques (the ones there and the ones here) deserve their own exposition, hence we present them as two separate papers.

1.1 Main methodological contributions

For the lower bounds, we develop an (iterative) renormalisation method that allows us to treat both the strict polynomial as well as the linear regime at once. We take the renormalisation method from Berger [6] as our starting point, who proved a linear lower bound for graph distances in long-range percolation. We have to overcome several issues that occur in 1-FPP but do not occur for graph distances. Firstly, in 1-FPP the edges carry transmission times, (or costs, for short), and costs are neither bounded from below by 1, nor are independent, because costs out of a vertex carry the same penalty factor coming from the expected degree of the vertex. Secondly, the underlying graphs we consider contain an excessive amount of long edges on all scales, which is due to the presence of very high degree vertices or ‘hubs’, that make graph distances doubly-logarithmic. Thirdly, the proof in [6] uses Kingman’s subadditive ergodic theorem, while for 1-FPP this is unavailable (especially in the polynomial regime).

Nevertheless, we are able to set up a renormalisation argument that takes all the above issues into account, in particular, we quantify that long edges *do* exist when $\tau \in (2, 3)$ (which is implied by the doubly-logarithmic graph distances), but these edges are typically *very expensive*. Although we follow broadly the proof in [6], our formulation gives the lower bound directly for all vertices, while [6] only showed it along a sequence of norms for the vertices, and then used Kingman’s subadditive ergodic theorem to extend the result to all vertices. We avoid this by improving on the conditions in [6, Lemma 2] for the 1-FPP, which allows us to extend the proof to non-linear regimes, in which Kingman’s theorem is not applicable (see Proposition 2.6 and the proof on page 28). See Section 1.4 for a more detailed outline.

Proof outline for the upper bounds. The upper bound in the linear regime in dimension at least two is also proved via renormalisation, here however a single renormalisation step suffices. Namely, for a sufficiently large constant M , we introduce an auxiliary graph: let $G_M = (\mathcal{V}_M, \mathcal{E}_M)$ be formed by vertices with weights in the interval $I_M := [M, 2M]$ and edges with edge-cost at most $M^{3\mu}$ among them. The restriction $2M$ on the weights of vertices guarantees that 1-FPP restricted to vertices with weight in I_M have penalty at most $4^\mu M^{2\mu}$, and hence a typical edge among vertices in \mathcal{V}_M is also part of \mathcal{E}_M . The auxiliary graph G_M can be stochastically dominated from below by a random geometric graph, and when $\tau \in (2, 3)$, its average degree increases with M . We then use a boxing argument to renormalise to a site-bond percolation ω^* , and use a result by Antal and Pisztorá [3] that distances in the infinite component \mathcal{C}_∞^* of ω^* scale linearly with the Euclidean distance. We then generalise a local-density result of Deuschel and Pisztorá in [19], saying that the infinite cluster \mathcal{C}_∞^* of ω^* comes near every vertex $z \in \mathbb{Z}^d$, to hold also for the infinite component of quite general models of random geometric graph with high edge density (see Definition 3.4), and thus also for G_M , see Corollary 3.9. These results may be of independent interest. The local density result allows us to connect the starting vertices $0, x$ to a nearby vertex in G_M at low cost, without having to deal with dependencies between the selection of the path and the edge-costs.

When $\tau > 3$, the edge-density in G_M decreases with M , which is why infinite geometric inhomogeneous random graphs (IGIRG) on \mathbb{R}^d and scale-free percolation (SFP) on \mathbb{Z}^d with $\tau > 3$ are both non-robust under percolation, while with $\tau \in (2, 3)$ they

are robust. The upper bound thus does not hold for the $\tau > 3$ case. A proof of linear distances for $\tau > 3, \alpha > 2$ assuming high edge-density could go along a similar reasoning as above, namely renormalisation. To prove it for arbitrary supercritical edge-density is nontrivial, since the path has to avoid high-degree vertices (because of the penalties in 1-FPP), which are exactly the vertices keeping the infinite component together. We leave this for future work.

For finite models in a box, we additionally need to ensure that the constructed paths stay inside the box of the graph. For this we prove that there exist linear distance paths between any sites u and v in the infinite component which deviate “little” from the straight line segment $S_{u,v}$ (this is implicitly present in [3], and we prove it for general models of random geometric graphs with high-density, see Lemma 3.8). Hence, for random vertices u_n, v_n in the giant component of a geometric inhomogeneous random graph (GIRG), there is a.s. a cheap path π_{u_n, v_n} from u_n to v_n in the corresponding IGIRG with small deviation. Since u_n and v_n are random, they are unlikely to be close to the boundary of the box, and hence π_{u_n, v_n} is completely contained in the GIRG.

1.2 Graph models

We will consider undirected, simple graphs with vertex set $\mathcal{V} \subseteq \mathbb{R}^d$. We use standard graph notation, which we summarise along with other common terminology in Section 1.5.

We consider three random graph models: *Scale-Free Percolation* (SFP), *Infinite Geometric Inhomogeneous Random Graphs* (IGIRG)¹, and (finite) *Geometric Inhomogeneous Random Graphs* (GIRG). Since the latter model contains *Hyperbolic Random Graphs* (HypRG) as special case, our results also hold for HypRG. The main difference between the models is the vertex set \mathcal{V} . For SFP, we use $\mathcal{V} := \mathbb{Z}^d$, where $d \in \mathbb{N}$. For IGIRG, \mathcal{V} is given by a Poisson point process on \mathbb{R}^d of intensity one with respect to the Lebesgue measure. The formal definition is:

Definition 1.1 (SFP, IGIRG, GIRG). *Let $d \in \mathbb{N}$, $\tau > 2$, $\alpha \in (1, \infty)$, and $\bar{c} > \underline{c} > 0$. Let $\ell : [1, \infty) \rightarrow (0, \infty)$ be a slowly varying function, and let $h : \mathbb{R}^d \times [1, \infty) \times [1, \infty) \rightarrow [0, 1]$ be a function satisfying*

$$\underline{c} \cdot \min \left\{ 1, \frac{w_1 w_2}{|x|^d} \right\}^\alpha \leq h(x, w_1, w_2) \leq \bar{c} \cdot \min \left\{ 1, \frac{w_1 w_2}{|x|^d} \right\}^\alpha. \quad (1.1)$$

We call d the dimension, τ the power-law exponent, α the long-range parameter, and h the connection probability.

For SFP, set $\mathcal{V} := \mathbb{Z}^d$, for IGIRG, let \mathcal{V} be given by a Poisson point process on \mathbb{R}^d of intensity one.² For each $x \in \mathcal{V}$, we draw a weight W_x independently from a probability distribution on $[1, \infty)$ satisfying

$$F_W(w) = \mathbb{P}(W \leq w) = 1 - \frac{\ell(w)}{w^{\tau-1}}. \quad (1.2)$$

We denote the weight vector $(W_x)_{x \in \mathcal{V}}$ by \mathcal{W} . Conditioned on \mathcal{V} and \mathcal{W} , every edge xy is present independently with probability $h(x - y, W_x, W_y)$.

A finite GIRG is obtained by restricting an IGIRG to a cube Q_n of volume n centred at 0.

For finite GIRG models we are interested in the behaviour as $n \rightarrow \infty$. Definition 1.1 leads to a slightly less general model than those e.g. in [13] and [32]. There, the original

¹They have also been called EGIRG, where E stands for extended [34].

²If we take an IGIRG and rescale the underlying space \mathbb{R}^d by a factor λ , then we obtain a random graph which satisfies all conditions of IGIRGs except that the density of the Poisson point process is λ^{-d} instead of one. Thus it is no restriction to assume density one.

definition had a different scaling of the geometric space vs connection probabilities. However, the resulting graphs are identical in distribution after rescaling, see [32] for a comparison. Finally, [13] considered the torus topology on the cube, identifying “left” and “right” boundaries, but this does not make a difference for our results. Next we define 1-dependent FPP on these graphs.

Definition 1.2 (1-dependent first passage percolation (1-FPP)). *Let $G(\mathcal{V}, \mathcal{E})$ be a graph where to each vertex $v \in \mathcal{V}$ there is an associated vertex-weight W_v . For each edge $xy \in \mathcal{E}$, draw an iid random variable L_{xy} from distribution L , and set the transmission cost of an edge xy as*

$$\mathcal{C}(xy) := L_{xy}(W_x W_y)^\mu \tag{1.3}$$

for a fixed parameter $\mu > 0$ called the penalty strength. Setting $\mathcal{C}(xy)$ defines a cost-distance $d_{\mathcal{C}}$ between the vertices (see (1.21) below), that we call 1-dependent first passage percolation. We define the cost-ball $B_r^{\mathcal{C}}(x)$ of radius $r \geq 0$ around a vertex x to be the set of all vertices to which there is a path of cost at most r from x .

We emphasize that in the case of SFP, IGIRG, GIRG, we use the same vertex weights $(W_v)_{v \in \mathcal{V}}$ for generating the graph as well as for determining the cost of the edges, while in general the definition would allow for different choices as well. We usually assume that the cumulative distribution function (cdf) $F_L : [0, \infty) \rightarrow [0, 1]$ of L satisfies the following (exceptions of this assumption will be made explicit when applicable):

Assumption 1.3. *There exist constants $t_0, c_1, c_2, \beta > 0$ such that*

$$c_1 t^\beta \leq F_L(t) \leq c_2 t^\beta \text{ for all } t \in [0, t_0]. \tag{1.4}$$

Without much effort, Assumption 1.3 could be relaxed to $\lim_{x \rightarrow 0} \log F_L(x) / \log x = \beta$, but having the stronger form in (1.4) increases the readability of the paper. For the same reason, we also exclude extensions to $\alpha = \infty$ in Definition 1.1 and $\beta = \infty$ in (1.4) from the main body of the paper, and discuss those separately in Section 1.3.1.

We will call the set of parameters

$$\text{par} := \{d, \tau, \alpha, \mu, \beta, \underline{c}, \bar{c}, c_1, c_2, t_0\} \tag{1.5}$$

the *model parameters*. We generally take the standpoint that we consider μ to be the easiest parameter to change, e.g. by adjusting behaviour of individuals corresponding to high-degree vertices, increasing μ means gradually slowing down the spreading process around high-degree vertices. We also phrase our results from this point of view.

1.3 Results

We now present two phases (polynomial but sublinear, and linear) of the transmission times between two far away nodes. Without loss of generality we fix one of the endpoints as $0 \in \mathbb{R}^d$. For IGIRG and GIRG, this means that we condition on the vertex set \mathcal{V} containing 0, i.e., we consider the resulting Palm distribution. Recall the power-law exponent τ of vertex-weights from (1.2), the long-range parameter α from the connection probability in (1.1), and β of the transmission times from Assumption 1.3. We assume $\tau \in (2, 3)$, this ensures that the degree distribution has finite mean but infinite variance, and that graph-distance balls grow doubly-logarithmically [12]. In the companion paper [33] and in [32] we show that cost-distances are at most polylogarithmic if $\alpha < 2$ or if $\mu < (3 - \tau)/\beta$. Thus here we only consider the complementary case (neglecting phase boundaries) that $\alpha > 2$ and $\mu > (3 - \tau)/\beta$. We first define the two values that will separate the phases of growth as μ changes:

$$\mu_{\log} := \frac{3 - \tau}{\beta}, \quad \mu_{\text{pol}} := \frac{1}{d} + \frac{3 - \tau}{\min\{\beta, d(\alpha - 2)\}} =: \max\{\mu_{\log} + \frac{1}{d}, \mu_{\text{pol}, \alpha}\} \tag{1.6}$$

with $\mu_{\text{pol},\alpha} := \frac{1}{d} + \frac{3-\tau}{d(\alpha-2)} = \frac{\alpha-(\tau-1)}{d(\alpha-2)}$. These values are thus all positive when $\tau \in (2, 3)$ and $\beta > 0$, and finite under Assumption 1.3 since $\beta > 0$. When $\alpha > 2$ then $\mu_{\text{pol},\alpha}$ is above $1/d$ as well. The values are also well-defined for $\tau = 3$, see Theorem 1.12 below about results with $\tau \geq 3$. Then we define the polynomial growth exponent. For $\alpha > 2, \tau \in (2, 3)$, for all $\mu > \mu_{\text{log}}$, let

$$\eta_0 = \eta_0(d, \alpha, \tau, \beta, \mu) := \begin{cases} 1 & \text{if } \mu > \mu_{\text{pol}}, \\ \min \{d(\mu - \mu_{\text{log}}), \mu/\mu_{\text{pol},\alpha}\} & \text{if } \mu \leq \mu_{\text{pol}}, \end{cases} \quad (1.7)$$

and note that $\eta_0 > 0$ for all $\mu > \mu_{\text{log}}$, and $\eta_0 < 1$ exactly when $\mu < \mu_{\text{pol}}$ by (1.6). We present special cases when $\alpha = \infty$ or $\beta = \infty$ and extensions to $\tau \geq 3$ in Section 1.3.1 below.

Theorem 1.4 (Polynomial Lower Bound). *Consider 1-FPP on IGIRG, GIRG, or SFP of Definition 1.1 satisfying the assumptions given in (1.2)–(1.4). Assume that $\alpha > 2$, $\tau \in (2, 3)$, and $\mu > \mu_{\text{log}}$. Then for any $\varepsilon > 0$ almost surely there exists $r > 0$ (independent of n in case of finite GIRG) such that*

$$\text{for all } x \in \mathcal{V} \text{ with } |x| \geq r : d_{\mathcal{C}}(0, x) \geq |x|^{\eta_0 - \varepsilon}.$$

Note that r is random, so it depends on the instance of IGIRG. For GIRG, “independent of n ” means that for every $q > 0$ there is $r_0 = r_0(q)$ independent of n such that we can satisfy Theorem 1.4 with an r such that $\mathbb{P}(r > r_0) < q$.³ Note that the lower bound holds simultaneously for all vertices x at distance at least r from 0.

This means that we also obtain a geometric bound on the location of the cost-ball around 0. However, the matching upper bounds in [33] are not uniform over all x with $\|x\| = r$, so we can only bound the cost-ball in one direction. We rephrase Theorem 1.4 in terms of intrinsic ball growth.

Corollary 1.5. *Consider 1-FPP on IGIRG, GIRG, or SFP of Definition 1.1 satisfying the assumptions given in (1.2)–(1.4) and with $0 \in \mathcal{V}$. Assume that $\alpha > 2$, $\tau \in (2, 3)$, and $\mu > \mu_{\text{log}}$. Then for all $\varepsilon > 0$, almost surely there exists r_0 such that for all $r \geq r_0$,*

(a) $\mathcal{B}_r^{\mathcal{C}}(0) \subseteq \{x \in \mathbb{R}^d : |x| \leq r^{1/\eta_0 + \varepsilon}\}$, and

(b) $|\mathcal{B}_r^{\mathcal{C}}(0)| \leq r^{d/\eta_0 + \varepsilon}$.

Part (a) of Corollary 1.5 rephrases Theorem 1.4, and (b) follows from (a) when considering the fact that the number of vertices in a (Euclidean) ball of radius $r^{1/\eta_0 + \varepsilon}$ around 0 concentrates around the volume of this ball (uniformly for all $r \geq r_0$), both for IGIRG and SFP.

To provide context, we cite the corresponding upper bound from [33]. For $\tau \in (2, 3)$, for all edge-densities, the infinite component exists, is unique, and has positive density, but not necessarily density one.⁴ Hence for an upper bound we need to condition on 0 and x both being in the infinite component.

Theorem 1.6 (Polynomial Upper Bound [33]). *Consider 1-FPP on IGIRG or SFP of Definition 1.1 satisfying the assumptions given in (1.2)–(1.4) with $0 \in \mathcal{V}$. Assume that $\alpha > 2$, $\tau \in (2, 3)$, and $\mu > \mu_{\text{log}}$. Let \mathcal{C}_{∞} be the infinite component. Then for any $\varepsilon > 0$,*

$$\lim_{|x| \rightarrow \infty} \mathbb{P}(d_{\mathcal{C}}(0, x) \leq |x|^{\eta_0 + \varepsilon} \mid 0, x \in \mathcal{C}_{\infty}) = 1. \quad (1.8)$$

³For fixed q , the r_0 in our proof grows like $\exp(c \cdot \log(1/q) \cdot \log \log(1/q))$ for a constant $c > 0$, see Remark 2.5.

⁴In SFP, there are choices of h that enforce all nearest-neighbour edges to be present, so density one is possible, but not guaranteed.

In [33] we also prove a corresponding theorem for finite GIRGs that we omit here. Observe that Theorem 1.6 does not hold *simultaneously* for all vertices at a certain norm, rather, the convergence holds in probability. The reason for this is that we did not make any assumptions on the *tail-behaviour* of the distribution of L , and hence we cannot exclude (clusters of) vertices with all edges having very large L values so that these are reached much later, violating the upper bound in (1.8). Hence Theorem 1.6 is not strong enough to provide an analogous lower bound to Corollary 1.5(a), since for that we would need a statement over *all* vertices in distance at most r . Yet, even though some vertices at norm r might have long cost-distance to 0, linearly many of them do satisfy $d_C(0, x) \leq |x|^{\eta_0+\varepsilon}$, so we still obtain the following corollary when combining Theorems 1.4 and 1.6:

Corollary 1.7. *Consider 1-FPP on IGIRG or SFP of Definition 1.1 satisfying the assumptions given in (1.2)–(1.4) with $0 \in \mathcal{V}$. Assume that $\alpha > 2$, $\tau \in (2, 3)$, and $\mu > \mu_{\log}$. Let \mathcal{C}_∞ be the infinite component. Then for any $\varepsilon > 0$,*

$$\lim_{|x| \rightarrow \infty} \mathbb{P} \left(\left| \frac{\log d_C(0, x)}{\log |x|} - \eta_0 \right| \leq \varepsilon \mid 0, x \in \mathcal{C}_\infty \right) = 1, \tag{1.9}$$

$$\lim_{r \rightarrow \infty} \mathbb{P} \left(\left| \frac{\log |\mathcal{B}_r^C(0)|}{\log r} - \frac{d}{\eta_0} \right| \leq \varepsilon \mid 0 \in \mathcal{C}_\infty \right) = 1. \tag{1.10}$$

Note that (1.10) follows immediately from (1.9) because for the lower bound on $|\mathcal{B}_r^C(0)|$ it suffices that a constant fraction of vertices at distance at most r have cost-distance at most $r^{1/\eta_0+\varepsilon}$, which is implied by (the upper bound in) (1.9). Hence we do obtain an absolute value in (1.10) within the \mathbb{P} sign.

Our next result refines Theorem 1.4 and (1.9) in Corollary 1.7 when $\eta_0 = 1$ in (1.7). I.e., when $\mu > \mu_{\text{pol}}$, we can sharpen both upper and lower bounds if the dimension is at least 2:

Theorem 1.8. *Consider 1-FPP on IGIRG or SFP of Definition 1.1 satisfying the assumptions given in (1.2)–(1.4) with $0 \in \mathcal{V}$. Assume that $\alpha > 2$, $\tau \in (2, 3)$, $\mu > \mu_{\text{pol}}$, and additionally $d \geq 2$. Let \mathcal{C}_∞ be the infinite component. Then there exist constants $\kappa_1, \kappa_2 > 0$ depending only on the model parameters such that*

$$\lim_{|x| \rightarrow \infty} \mathbb{P} (\kappa_1|x| < d_C(0, x) < \kappa_2|x| \mid 0, x \in \mathcal{C}_\infty) = 1. \tag{1.11}$$

The lower bound also holds in dimension 1.

The lower bound is actually valid in a more general setting, see Theorem 1.12 below. Our proof of the lower bound is a generalisation of ideas from Berger [6] to the edge-weighted one-dependent setting, and indeed we recover his result on graph-distances in Long-Range Percolation and the extension to Scale-Free Percolation in [17] when we set $\alpha > 2, \tau > 3, \mu = 0, \beta = \infty$. However, we give a proof that avoids Kingman’s subadditive ergodic theorem that both papers [6, 17] use. As a corollary to our proof we obtain the following result.

Corollary 1.9. *Consider the setting of Theorem 1.8. Let $\pi_{0,x}^*$ be any path between 0 and x realising the cost-distance $d_C(0, x)$. Then there is a constant $\kappa_3 > 0$ (depending only on the model parameters) such that $\pi_{0,x}^*$ is contained in $B_{\kappa_3|x|}(0)$ asymptotically almost surely:*

$$\lim_{|x| \rightarrow \infty} \mathbb{P} \left(\pi_{0,x}^* \subset B_{\kappa_3|x|}(0) \right) = 1. \tag{1.12}$$

The statements of Theorem 1.4 and Corollary 1.5 remain valid in finite-sized GIRGs, since finite GIRGs are obtained as subgraphs of IGIRG, and hence distances can only increase in GIRG versus the surrounding infinite model. For the upper bound in Theorem 1.8, the extension to finite GIRGs requires a proof:

Theorem 1.10. Consider GIRG of Definition 1.1 satisfying the assumptions given in (1.2)–(1.4). Assume that $\alpha > 2$, $\tau \in (2, 3)$, $\mu > \mu_{\text{pol}}$, and additionally $d \geq 2$. Let $C_{\text{max}}^{(n)}$ be the largest component in Q_n . Let u_n, v_n be two vertices chosen uniformly at random from $C_{\text{max}}^{(n)}$. Then there exist constants $\kappa_1, \kappa_2 > 0$ depending only on the model parameters such that

$$\lim_{n \rightarrow \infty} \mathbb{P}(\kappa_1 |u_n - v_n| < d_C(\overline{u_n}, v_n) < \kappa_2 |u_n - v_n|) = 1. \tag{1.13}$$

Next we present extensions, in particular for the $\tau > 3, \alpha > 2$ case where we have very mild conditions on the distribution of L , see Corollary 1.13 below.

1.3.1 Limit cases and extensions

The results above naturally extend to cases/models that may informally be described as $\alpha = \infty$ or $\beta = \infty$, and to some extent to $\tau \geq 3$ as well. We start with $\alpha = \infty$. This means that we replace the condition (1.1) on the connection probability $h(\cdot)$ by

$$h(x, w_1, w_2) \begin{cases} = 0, & \text{if } \frac{w_1 w_2}{|x|^d} < c', \\ \geq \underline{c} & \text{if } \frac{w_1 w_2}{|x|^d} \geq 1, \end{cases} \tag{1.14}$$

for some constants $\underline{c}, c' \in (0, 1]$, and where $h(x, w_1, w_2)$ can take arbitrary values in $[0, 1]$ when $w_1 w_2 / |x|^d \in [c', 1)$. We use the bound $\frac{w_1 w_2}{|x|^d} \geq 1$ in the second row for convenience, as it allows us to write the proofs for different cases in a consistent way, but it could easily be replaced by $\frac{w_1 w_2}{|x|^d} \geq c''$ for any other constant $c'' \geq c'$. The connectivity function h in (1.14) covers the so-called threshold regime for *hyperbolic random graphs* when we also set $d = 1$, see [13, Theorem 2.3]. In this case, when $\tau \in (2, 3)$, we extend the definitions (1.6) and (1.7) in the natural way, since $\lim_{\alpha \rightarrow \infty} \mu_{\text{pol}, \alpha} = 1/d$:

$$\mu_{\text{log}} := \frac{3 - \tau}{\beta}, \quad \mu_{\text{pol}} := \frac{1}{d} + \frac{3 - \tau}{\beta}, \quad \eta_0 := \begin{cases} 1 & \text{if } \mu > \mu_{\text{pol}}, \\ d \cdot (\mu - \mu_{\text{log}}) & \text{if } \mu \leq \mu_{\text{pol}}. \end{cases} \tag{1.15}$$

To describe the case $\beta = \infty$, we replace (1.4) by the condition

$$\lim_{t \rightarrow 0} F_L(t)/t^\beta = 0 \text{ for all } 0 < \beta < \infty. \tag{1.16}$$

This means that the cdf of $L > 0$ is flatter than any polynomial near 0. In particular, this condition is satisfied if F_L has no probability mass around zero, for example in the case $L \equiv 1$.

In this case, we again replace (1.6)–(1.7) naturally by

$$\mu_{\text{log}} := 0, \quad \mu_{\text{pol}} := \max\{\frac{1}{d}, \mu_{\text{pol}, \alpha}\}, \quad \eta_0 := \begin{cases} 1 & \text{if } \mu > \mu_{\text{pol}}, \\ \min\{d\mu, \mu/\mu_{\text{pol}, \alpha}\} & \text{if } \mu \leq \mu_{\text{pol}}. \end{cases} \tag{1.17}$$

We mention that these definitions stay valid also for $\tau = 3$. Finally, in the case $\alpha = \beta = \infty$ we replace (1.6)–(1.7) by

$$\mu_{\text{log}} := 0, \quad \mu_{\text{pol}} := \frac{1}{d}, \quad \eta_0 := \min\{1, d\mu\}. \tag{1.18}$$

Our main results still hold for these limit regimes. We remark that the corresponding upper bounds also hold, see [33, Theorem 1.8].

Theorem 1.11 (Extension to threshold IGIRGs/GIRGs, and $\beta = \infty$).

(a) *Theorems 1.4, 1.8, and 1.10 still hold for $\alpha = \infty$ if we replace the definitions of $\mu_{\text{pol}}, \mu_{\text{log}}, \eta_0$ in (1.6)–(1.7) by their values in (1.15).*

(b) Theorems 1.4, 1.8, and 1.10 still hold for $\beta = \infty$ if we replace the definitions of $\mu_{\text{pol}}, \mu_{\text{log}}, \eta_0$ in (1.6)-(1.7) by their values in (1.17).

(c) Theorems 1.4, 1.8, and 1.10 still hold for $\alpha = \beta = \infty$ if we replace the definitions of $\mu_{\text{pol}}, \mu_{\text{log}}, \eta_0$ in (1.6)-(1.7) by their values in (1.18).

In Theorem 1.11(b), the requirement that $\mu > \mu_{\text{log}} = 0$ implies that we exclude the case $\mu = 0$. Indeed, $\mu = 0$ means the setting of classical iid first-passage percolation, where whenever $\tau \in (2, 3)$, explosion occurs under mild conditions on L , see [34, Theorems 2.11, 2.12]. Theorem 1.11(a) implies the analogous result for hyperbolic random graphs (HypRG) by setting $d = 1$ in (1.15), except for some technical details. By our Definition 1.1, the number of vertices in a finite GIRG is random, and has Poisson distribution with parameter n , while in the usual definition of HypRG [27, 35] and GIRG [13] the number of vertices is exactly n . Further, in HypRG the distribution of the vertex-weights is n -dependent $W^{(n)}$, so that $W^{(n)}$ converges to a limiting distribution W [34]. These issues can be overcome by coupling techniques e.g. presented in [34]: a model with exactly n vertices can be squeezed between two GIRGs with Poisson intensity $1 - \sqrt{4 \log n/n}$ and $1 + \sqrt{4 \log n/n}$, and one can couple n -dependent and limiting vertex-weights to each other, respectively, but we avoid spelling out the details and refer the reader to [34, Claims 3.2, 3.3].

Finally, the lower bounds in Theorems 1.4 and 1.8 also hold under much weaker, but more technical conditions. In particular, for $\tau > 3$ we can strongly relax Assumption 1.3 on L .

Theorem 1.12. *Consider 1-FPP on IGIRG or SFP of Definition 1.1 satisfying the assumptions given in (1.2)–(1.3) with $0 \in \mathcal{V}$ (but not necessarily Assumption 1.3 on L). Assume that $\alpha > 2$ and $\tau \in (2, \infty)$.*

(1) *Conditions for polynomial distance lower bound. If $\tau \in (2, 3]$ and L satisfies Assumption 1.3 with some $\beta \in (0, \infty]$ so that $\mu > \mu_{\text{log}}$, then for all $\varepsilon > 0$, almost surely there is $r > 0$ such that*

$$\text{for all } x \in \mathcal{V} \text{ with } |x| \geq r : d_{\mathcal{C}}(0, x) \geq |x|^{\eta_0 - \varepsilon}. \tag{1.19}$$

(2) *Conditions for strictly linear distance lower bound. If either $\tau > 3$ and $\mu \geq 0$, and L in (1.3) satisfies $\mathbb{P}(L > 0) = 1$, or $\tau \in (2, 3]$ and L satisfies Assumption 1.3 with some $\beta > 0$ so that $\mu > \mu_{\text{pol}}$, then there exists $\kappa > 0$ such that almost surely there is $r > 0$ such that*

$$\text{for all } x \in \mathcal{V} \text{ with } |x| \geq r : d_{\mathcal{C}}(0, x) \geq \kappa|x|. \tag{1.20}$$

Compared to Theorems 1.4 and 1.8, Theorem 1.12 covers the boundary case $\tau = 3$. More importantly, it states that the linear lower bound in (1.20) holds for $\tau > 3$ and $\alpha > 2$ under very mild conditions on L and μ , i.e., we allow arbitrary edge weight distributions L that have no probability mass at zero, and arbitrary penalty strength $\mu \geq 0$. This includes *classical first passage percolation* by setting $\mu = 0$ and L arbitrary, a.s. positive, but also the case $L \equiv 1$ and $\mu = 0$, giving *graph distances*. In this latter case we recover the result of Berger [6] for long-range percolation (LRP) and its extension by Deprez, Hazra, and Wüthrich [17] for SFP. We state the case of *classical* (iid) first passage percolation on finite variance degree models ($\tau > 3$) with long-range parameter $\alpha > 2$ as an immediate corollary.

Corollary 1.13. *Consider classical first passage percolation with iid transmission times from distribution L satisfying $\mathbb{P}(L > 0) = 1$ on IGIRG, GIRG or SFP of Definition 1.1 with $\tau > 3, \alpha > 2$. Then there exists $\kappa > 0$ such that almost surely there is $r > 0$ such that*

$$\text{for all } x \in \mathcal{V} \text{ with } |x| \geq r : d_{\mathcal{C}}(0, x) \geq \kappa|x|.$$

In Theorem 1.12, (unlike in Theorem 1.8), since we state lower bounds, we do not condition on $0, x \in \mathcal{C}_\infty$. In the case $\tau > 3$, an infinite component does not need to exist (it depends on the constants and slowly varying functions), so conditioning on $0, x \in \mathcal{C}_\infty$ would not make sense. However, for parameters that ensure an infinite component of positive density, the event $0, x \in \mathcal{C}_\infty$ has positive probability, and (1.19) and (1.20) remain true after conditioning.

Discontinuity at $\tau = 3$. Remarkably, for $\tau > 3$ there is no analogue to the strictly polynomial regime in Theorem 1.4, even though algebraically the limits of $\lim_{\tau \nearrow 3} \mu_{\text{pol}} = 1/d$ and $\lim_{\tau \nearrow 3} \eta_0 = \mu d$ exist and are in $(0, 1)$. So if we fix some $\mu < 1/d$ and let $\tau \nearrow 3$, the cost-distances grow polynomially with exponents bounded away from one (e.g., they approach $1/2$ from below for $\mu = 1/(2d)$). But as soon as $\tau > 3$ is reached, the exponent “jumps” to one. In other words, even though μ_{pol} converges to a positive limit $1/d$ as $\tau \nearrow 3$, it drops to 0 as soon as $\tau > 3$ is reached. In this sense, the parameter space is discontinuous in η_0 and μ_{pol} .

1.4 Proof outline and organisation of the paper

Proof outline of the lower bounds. For lower bounds, we need to show that every path from a vertex x to 0 is expensive. For this, we quantify the property that “most long edges are expensive” in Lemma 2.2. This allows us to generalise a renormalisation method from Berger [6]. For long-range percolation, Berger considers a growing system of boxes around the origin and defines a box Q as *good* if it does not contain edges of linear length in the box size and the same property holds recursively for most of its child-boxes, which are a certain set of non-disjoint sub-boxes that cover Q multiple times. In the setting of 1-FPP, we modify this definition since long edges *do* exist when $\tau \in (2, 3)$ (which is implied by the doubly-logarithmic graph distances), but these edges are typically very expensive. We then show inductively the deterministic statement that once a box is good, the cost-distance is “large” between any pair of vertices inside the box with Euclidean distance linear in the box size. By “large”, we mean either linear, or polynomial with an exponent less than one, in the Euclidean distance, depending on the model parameters, which is unlike the linear graph-distances in [6]. Polynomial cost-distances occur when $\mu \in (\mu_{\log}, \mu_{\text{pol}})$. For the inductive step, for any two sufficiently distant vertices u, v in the same good box Q , we use that any path $\pi_{u,v}$ between them must either (i) contain a long and thus expensive edge, or (ii) has many long disjoint sub-segments in good child-boxes of Q , whose costs we can lower-bound by induction. See Figure 1.

We give the definition of good boxes in the FPP setting in Section 2.2, after calculating the probability that long but cheap edges exist in a box in Section 2.1. In Section 2.3 we then show that there are no cheap paths *within* a good box, and in Section 2.4 that there are no cheap paths at all, excluding thus cheap paths that leave and return to the same box. We avoid Kingman’s subadditive theorem in Proposition 2.6, see also the proof on page 28. Finally, in Section 2.5 we show the analogous lower bounds for $\alpha = \infty$ and/or $\beta = \infty$ in Theorem 1.11 using coupling arguments.

The proof of the upper bound (that we sketched already in Section 1.1) can be found in Section 3.

1.5 Notation

The graphs in this paper are undirected, simple graphs with vertex set $\mathcal{V} \subseteq \mathbb{R}^d$. For a graph $G = (\mathcal{V}, \mathcal{E})$ and a set $A \subseteq \mathbb{R}^d$, we denote by $G[A]$ the induced subgraph of G with vertex set $\mathcal{V} \cap A$. For two vertices $x, y \in \mathcal{V}$, we denote the edge between them by xy . For a path π , we denote the number of edges in π by $|\pi|$.

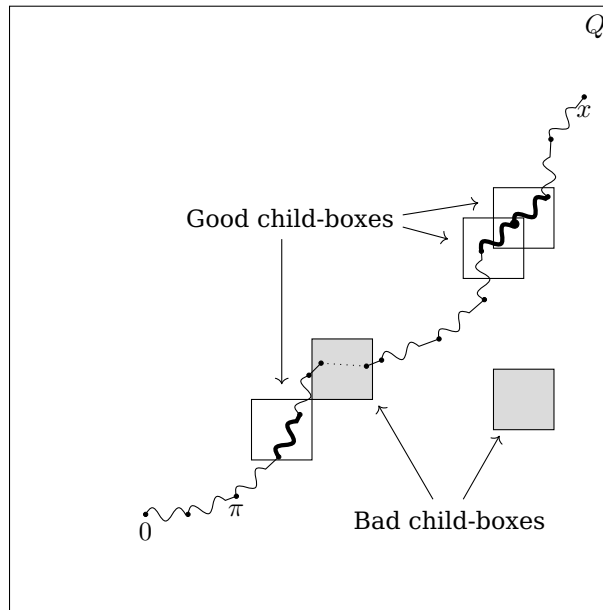


Figure 1: Illustration of the inductive step for the lower bound proof. The big box is good, hence all edges in it that are longer than one-hundredth of the smaller boxes' side-length are already too expensive. So, any cheap-enough path π between 0 and x contained in Q only uses edges shorter than that. There are only a few bad child-boxes, since Q is good, and we remove the edges of π with at least one end-point in bad child-boxes. We select from the remaining edges of π enough sub-segments such that each sub-segment is fully contained in a good child-box of Q and connects two sufficiently far vertices within that child-box. Three sub-segments are shown in bold, together with their corresponding child-boxes. By induction we have a lower bound on the costs of these sub-segments, and summing up those yields a lower bound on $\mathcal{C}(\pi)$.

Given a cost function $\mathcal{C} : \mathcal{E} \rightarrow [0, \infty]$ on the edges, the cost of a collection of edges π is the sum of the costs of the edges in the collection, $\mathcal{C}(\pi) := \sum_{e \in \pi} \mathcal{C}(e)$. For convenience we define $\mathcal{C}(xx) := 0$ for all $x \in \mathcal{V}$. We define the *cost-distance* between vertices x and y as

$$d_{\mathcal{C}}(x, y) := \inf\{\mathcal{C}(\pi) : \pi \text{ is a path from } x \text{ to } y \text{ in } G\}. \tag{1.21}$$

We define the graph distance $d_G(x, y)$ similarly, when we set all edge-costs to 1. The graph will usually be clear from the context.

We denote the Euclidean norm of $x \in \mathbb{R}^d$ by $|x|$. We denote by $B_r(x) := \{y \in \mathbb{R}^d : |x - y| \leq r\}$ the Euclidean ball with radius $r \geq 0$ around x , and by $\mathcal{B}_r(x) := \{y \in \mathcal{V} : |x - y| \leq r\} = B_r(x) \cap \mathcal{V}$ the set of vertices in this ball. The minimal notation difference is intentional, as we do not expect any confusion to arise between the two sets. The *graph-distance ball* and *cost-distance ball* (or *cost-ball* for short) around a vertex x are the vertex sets $\mathcal{B}_r^G(x) := \{y \in \mathcal{V} : d_G(x, y) \leq r\}$ and $\mathcal{B}_r^{\mathcal{C}}(x) := \{y \in \mathcal{V} : d_{\mathcal{C}}(x, y) \leq r\}$, respectively. We set $B_r := B_r(0)$, and similarly for $\mathcal{B}_r, \mathcal{B}_r^G$, and $\mathcal{B}_r^{\mathcal{C}}$, if 0 is a vertex. We denote by $\partial B_r(x) := B_r(x) \setminus \{y \in \mathbb{R}^d : |x - y| < r\}$, and similarly for $\partial \mathcal{B}_r, \partial \mathcal{B}_r^G$, and $\partial \mathcal{B}_r^{\mathcal{C}}$. In particular, $\partial \mathcal{B}_1^G(v)$ is the neighbourhood of a vertex v .

For parameters $a, b > 0$ (model parameters or otherwise), we use “for all $a \gg_* b$ ” as shortcut for “ $\forall b > 0 : \exists a_0 = a_0(b) : \forall a \geq a_0$ ”. We also say “ $a \gg_* b$ ” to mean that $a \geq a_0(b)$. We use $a \ll_* b$ analogously, and may use more than two parameters. For

example, “for $a \gg_* b, c$ ” means “ $\forall b, c > 0 : \exists a_0 = a_0(b, c) : \forall a \geq a_0$ ”. A measurable function $\ell : (0, \infty) \rightarrow (0, \infty)$ is *slowly varying* if $\lim_{x \rightarrow \infty} \ell(ax)/\ell(x) = 1$ for all $a > 0$. We denote the natural logarithm by \log . For $n \in \mathbb{N}$ we abbreviate $[n] := \{1, \dots, n\}$. For $S \subseteq \mathbb{R}^d$, we denote the Lebesgue measure of S by $\text{Vol}(S)$.

2 Lower bounds

The main part of this section is the proof of Theorem 1.12, which in turn implies the lower bounds in Theorem 1.4 and Theorem 1.8. Note that the corresponding lower bound for GIRG in Theorem 1.10 also follows trivially since GIRG is a subgraph of IGIRG. Throughout the section, we will maintain Assumption 1.3 on L unless explicitly noted otherwise.

Before we start with the proof, we note an elementary lemma stating that the product of two random variables with regularly varying tails sharing the same tail exponent again has a regularly varying tail with the same tail exponent.

Lemma 2.1. *Let X, Y be two independent positive random variables with cumulative distributions $F_X(x) = 1 - \ell_1(x)x^{-\tau}$ and $F_Y(y) = 1 - \ell_2(y)y^{-\tau}$, respectively, where ℓ_1 and ℓ_2 are slowly varying functions. Then the cumulative distribution of their product $Z := XY$ is given by $F_Z(z) = 1 - \ell^*(z)z^{-\tau}$ for some slowly varying function ℓ^* .*

Proof. This is a consequence of [21, Corollary, Page 3]. □

2.1 Upper bound on the number of long and cheap edges

As mentioned in Section 1.1, we start with developing a renormalisation argument. The argument builds on the basic observation that a box is unlikely to contain an edge that is both long in Euclidean distance and cheap, i.e., its cost is small. This follows from a straightforward but tedious calculation and is summarised in the following lemma. The more interesting part of the proof comes afterwards, when we turn this basic property into the statement that all paths between a vertex x and 0 are at least polynomially expensive.

Lemma 2.2. *Consider 1-FPP on IGIRG or SFP of Definition 1.1 satisfying the assumptions given in (1.2)–(1.4) with $0 \in \mathcal{V}$. Assume that $\tau \in (2, \infty)$ and $\beta > 0$. For all $\varepsilon > 0$, if $N > 0$ is sufficiently large relative to ε , then the following holds. Let $A > N$ and $a > 0$. Then the expected number of edges in $[-A/2, A/2]^d$ with length at least N and cost at most N^{ad} is at most*

$$\begin{cases} A^d N^\varepsilon \left(N^{-d(\alpha-1)} + N^{-d(\tau-2)} \right) & \text{if } a \geq \mu, \\ A^d N^\varepsilon \left(N^{-d(\alpha-1)} + N^{-d(\alpha-1-\frac{a}{\mu}(\alpha-(\tau-1)))} + N^{-d(\tau-2+(\mu-a)\beta)} \right) & \text{if } a < \mu. \end{cases} \quad (2.1)$$

The formula for the case $a \geq \mu$ remains valid without the restriction on the cost of the edges, and in this case L does not need to satisfy Assumption 1.3. Further, for IGIRG, the formula is also valid under the Palm measures \mathbb{P}^0 or $\mathbb{P}^{0,x}$ where we condition on having a vertex with unknown weight at $0 \in \mathbb{R}^d$ or on having two vertices with unknown weight at location $0, x \in \mathbb{R}^d$, respectively.

Before the proof, let us informally explain the formula, suppressing slowly varying factors in the discussion. The factor A^d is simply the (expected) number of vertices in $[-A/2, A/2]^d$. The term N^ε comes from applying Potter’s bound [8] to bound the slowly varying function that appears in the distribution function of the vertex weights in (1.2) when we integrate over the distribution of the products of weights $W_u W_v$ in the connectivity function in (1.1).

For the case $a \geq \mu$, the two terms in the brackets represent two types of edges. Consider a vertex v of constant weight. Then first term $N^{-d(\alpha-1)}$ counts the expected number of neighbours u in distance $\sim N$ (e.g., in distance $[N, 2N]$) of constant weight, up to negligible factors. The second term $N^{-d(\tau-2)}$ counts the number of neighbours u in distance $\sim N$ of weight $\sim N^d$, which is exactly the weight needed to get constant connection probability from (1.1). When $a \geq \mu$, the cost of such edges is typically less than N^{ad} , which is why we can ignore the cost condition in this case. There are potentially other possibilities, but their contribution to the expectation is negligible.

In the case $a < \mu$, the first term $N^{-d(\alpha-1)}$ is similar to the case $a \geq \mu$, coming from edges between vertices of constant weight, and so the cost is typically constant as well. The third term $N^{-d((\tau-2)+(\mu-a)\beta)}$ comes from edges between vertices of weights $W_v \sim 1$ and $W_u \sim N^d$. These receive a cost penalty of $N^{d\mu}$, so the probability that the edge has cost at most N^{ad} is roughly $F_L(N^{-d(\mu-a)}) \approx N^{-d(\mu-a)\beta}$. Finally, the second term $N^{-d(\alpha-1-a/\mu(\alpha-(\tau-1)))}$ comes from edges between vertices of weights $W_v \sim 1$ and $W_u \sim N^{ad/\mu}$. With this weight, the typical cost of the edge is N^{ad} , so the cost condition is satisfied. For fixed v , there are $\Theta(N^d)$ vertices u in distance $\sim N$, they have probability $N^{-(ad/\mu)(\tau-1)}$ to be of weight $W_u \sim N^{ad/\mu}$, and the probability to be adjacent to v with this weight is $(N^{ad/\mu}/N^d)^\alpha$ by (1.1). (All up to negligible terms.) Together, this yields the second term. Note that the term $(N^{ad/\mu}/N^d)^\alpha$ for the connection probability is only correct if the bracket is at most one, i.e., if $a < \mu$.

Proof of Lemma 2.2. Throughout this proof, we will denote by C_1, C_2, \dots finite positive constants depending on ε and the set par of model parameters. For readability, we allow these to appear in the middle of a calculation without necessarily being defined beforehand. Note that the statement of the lemma is stronger for smaller $\varepsilon > 0$. So without loss of generality, we can assume that we take a $\varepsilon > 0$ small enough so that

$$-d(\tau - 2) + \varepsilon/2 < 0 \quad \text{and} \quad -d(\alpha - 1) + \varepsilon/2 < 0. \tag{2.2}$$

If $\alpha - (\tau - 1) - \mu\beta < 0$, we can also assume that

$$\alpha - (\tau - 1) - \mu\beta + \varepsilon \frac{\mu}{2ad} < 0. \tag{2.3}$$

Let $E(A, N, a)$ denote the expected number of edges in $[-A/2, A/2]^d$ with length at least N and cost at most N^{ad} . We first compute $E(A, N, a)$ for SFP, having vertex set \mathbb{Z}^d . Let

$$\Lambda(r) := \mathbb{E}[(1 \wedge W_x W_y / r^d)^\alpha F_L(N^{ad}(W_x W_y)^{-\mu})]. \tag{2.4}$$

Using conditional expectation we have

$$\begin{aligned} E(A, N, a) &= \sum_{\substack{x, y \in [-A/2, A/2]^d \cap \mathbb{Z}^d \\ |x-y| \geq N}} \mathbb{E} \left[\mathbb{1}_{\{xy \text{ is an edge}\}} \cdot \mathbb{1}_{\{C(xy) \leq N^{ad}\}} \right] \\ &\leq \sum_{\substack{x, y \in [-A/2, A/2]^d \cap \mathbb{Z}^d \\ N \leq |x-y| \leq dA}} \mathbb{E} \left[\bar{c} \left(1 \wedge \frac{W_x W_y}{|x-y|^d} \right)^\alpha \cdot F_L(N^{ad}(W_x W_y)^{-\mu}) \right] \\ &= \bar{c} \sum_{x \in [-A/2, A/2]^d \cap \mathbb{Z}^d} \sum_{\substack{y \in [-A/2, A/2]^d \cap \mathbb{Z}^d \\ N \leq |x-y| \leq dA}} \Lambda(|x-y|). \end{aligned}$$

Note that the number of vertices $|\mathbb{Z}^d \cap [-A/2, A/2]^d| \leq C_1 A^d$. In order to simplify calculations, we will replace the second sum by an integral. More precisely, by usual

isoperimetric inequalities for \mathbb{Z}^d , there is a constant $C_2 = C_2(d)$ such that

$$E(A, N, a) \leq C_2 A^d \int_{r=N}^{dA} r^{d-1} \Lambda(r) dr. \tag{2.5}$$

For IGIRG, we obtain the same formula (2.5) in a simpler way. In expectation, there are A^d vertices in $[-A/2, A/2]^d$. For any fixed vertex, the expected number of neighbours in IGIRG in distance between N and dA with edge cost at most N^{ad} is given by the integral $\int_{r=N}^{dA} r^{d-1} \Lambda(r) dr$. Since dA bounds the diameter of $[-A/2, A/2]^d$, this includes all neighbours of this type in $[-A/2, A/2]^d$. Hence, the upper bound (2.5) also holds for IGIRG and its subgraph GIRG. The same bound remains true under the Palm measures $\mathbb{P}^0, \mathbb{P}^{0,x}$ by possibly increasing the constant prefactor C_2 to account for the edges emanating from the two extra vertices.

Next we bound $\Lambda(r)$ in (2.5). Defined in (2.4), $\Lambda(r)$ only depends on the product $W_x W_y =: Z$, not on the individual weights of the two vertices. By Lemma 2.1, the distribution of Z is of the form $f_Z(z) = \ell^*(z) z^{-\tau}$, where ℓ^* is a slowly varying function. For the sake of simplicity, we will work with Z having a density, but a proof using only Lebesgue-Stieltjes integration is similar. We recall from (2.5) that $r > N$ and rewrite (2.4) using law of total probability as

$$\Lambda(r) = \int_{z=1}^{\infty} \left(1 \wedge \frac{z}{r^d}\right)^\alpha \cdot F_L(N^{ad} z^{-\mu}) \cdot \frac{\ell^*(z)}{z^\tau} dz. \tag{2.6}$$

We now split into cases depending on the value of a .

Case 1: $a \geq \mu$. In this case we first show that for all $\varepsilon > 0$, for all sufficiently large r (i.e., larger than some $r_0(\varepsilon)$),

$$\Lambda(r) \leq C_3(r^{-d(\tau-1)} + r^{-d\alpha})r^{\varepsilon/2}. \tag{2.7}$$

We split the inner integral of (2.6) in two parts, at r^d . The first part is given by

$$I_1 := \int_{z=1}^{r^d} \left(1 \wedge \frac{z}{r^d}\right)^\alpha \cdot \underbrace{F_L(N^{ad} z^{-\mu})}_{\leq 1} \cdot \frac{\ell^*(z)}{z^\tau} dz \leq r^{-\alpha d} \int_{z=1}^{r^d} z^{\alpha-\tau} \ell^*(z) dz. \tag{2.8}$$

Since $r > N$, and we will later let $N \rightarrow \infty$, we can use Karamata's theorem [8, Prop. 1.5.6], and obtain that for $\alpha - (\tau - 1) > 0$,

$$I_1 \leq r^{-\alpha d} \int_{z=1}^{r^d} z^{\alpha-\tau} \ell^*(z) dz \leq r^{-\alpha d} C_4 r^{d(\alpha-(\tau-1))} \ell^*(r^d) \leq C_4 r^{-d(\tau-1)+\varepsilon/2} \tag{2.9}$$

for N (and thus r) large enough, where we used Potter's bound [8] to get $\ell^*(r) = o(r^{\varepsilon/(2d)})$ as $r \rightarrow \infty$, and thus for sufficiently large N (and hence r) we obtain that I_1 in (2.8) in the $\alpha > \tau - 1$ case is bounded from above by the first term in (2.7). If $\alpha - \tau + 1 = 0$, we again use $\ell^*(z) = o(z^{\varepsilon/(2d)})$ as $z \rightarrow \infty$ to get

$$I_1 \leq C_5 r^{-\alpha d} \int_{z=1}^{r^d} z^{\varepsilon/(2d)-1} dz \leq C_6 r^{-\alpha d + \varepsilon/2}, \tag{2.10}$$

which is the second term in (2.7). Finally, when $\alpha < \tau - 1$, we use Potter's bound [8] to get $\ell^*(z) \leq C_7 z^\varepsilon$ as $z \rightarrow \infty$, and since $\alpha - \tau + \varepsilon < -1$ we get that the integral is bounded by some constant and hence the bound (2.10) remains valid. Combining equations (2.9)–(2.10), we obtain that regardless of the relation between α and $\tau - 1$, I_1 satisfies the bound in (2.7). The second part of the inner integral in (2.6) is

$$I_2 := \int_{z=r^d}^{\infty} \left(1 \wedge \frac{z}{r^d}\right)^\alpha \cdot F_L(N^{ad} z^{-\mu}) \cdot \frac{\ell^*(z)}{z^\tau} dz \leq \int_{z=r^d}^{\infty} z^{-\tau} \ell^*(z) dz,$$

since $F_L \leq 1$ always holds. Since $r^d \rightarrow \infty$ as $N \rightarrow \infty$, Proposition 1.5.10 of [8] tells us that for N large enough

$$I_2 \leq \int_{z=r^d}^{\infty} z^{-\tau} \ell^*(z) dz \leq C_8 r^{-d(\tau-1)} \ell^*(r^d)$$

Again using Potter's bound, for sufficiently large N we have that I_2 is dominated by the first term in (2.7). This finishes the proof that (2.7) holds when $a \geq \mu$.

Returning the attention to $E(A, N, a)$ in (2.5), and using now (2.7), we have for $\tau > 2$

$$\begin{aligned} E(A, N, a) &\leq C_2 A^d \int_{r=N}^{dA} r^{d-1} C_3 (r^{-d(\tau-1)} + r^{-d\alpha}) r^{\varepsilon/2} dr \\ &= C_9 A^d \int_{r=N}^{dA} r^{-d(\tau-2)+\varepsilon/2-1} + r^{-d(\alpha-1)+\varepsilon/2-1} dr \\ &\stackrel{(2.2)}{\leq} C_{10} A^d N^{\varepsilon/2} \left(N^{-d(\tau-2)} + N^{-d(\alpha-1)} \right) \leq A^d N^\varepsilon \left(N^{-d(\tau-2)} + N^{-d(\alpha-1)} \right), \end{aligned}$$

where the last inequality holds for N large enough. Thus we have proved (2.1) when $a \geq \mu$. Observe that we used a trivial bound $F_L \leq 1$ in the proofs, hence we get that the statement also holds without any restriction on the edge-costs and without Assumption 1.3.

Case 2: $a < \mu$. In this case we also start bounding $\Lambda(r)$ in (2.6) first. We recall the constant t_0 from (1.4). Note that $t_0^{-1/\mu} N^{ad/\mu}$ is smaller than N^d (and thus r^d) for N large enough. We assume this inequality and split the integral of (2.6) into three parts. In the first part we bound the factor F_L by 1 from above:

$$\tilde{I}_1 := \int_{z=1}^{t_0^{-1/\mu} N^{ad/\mu}} \left(1 \wedge \frac{z}{r^d} \right)^\alpha \cdot F_L(N^{ad} z^{-\mu}) \cdot z^{-\tau} \ell^*(z) dz \leq r^{-d\alpha} \int_{z=1}^{t_0^{-1/\mu} N^{ad/\mu}} z^{\alpha-\tau} \ell^*(z) dz.$$

This is the same as I_1 in the previous case, except for the upper limit of integration. Using the same reasoning, we get that

$$\tilde{I}_1 \leq C_{11}(t_0) r^{-d\alpha} \left(N^{\varepsilon/2} + N^{\frac{ad}{\mu}(\alpha-\tau+1)+\varepsilon/2} \right) \tag{2.11}$$

for N large enough. In the second part the argument of F_L will be at most t_0 , hence we can use that $F_L(t) \leq c_2 t^\beta$ in this regime:

$$\begin{aligned} \tilde{I}_2 &:= \int_{z=t_0^{-1/\mu} N^{ad/\mu}}^{r^d} \left(1 \wedge \frac{z}{r^d} \right)^\alpha \cdot F_L(N^{ad} z^{-\mu}) \cdot z^{-\tau} \ell^*(z) dz \\ &\stackrel{(1.4)}{\leq} c_2 r^{-d\alpha} \int_{z=t_0^{-1/\mu} N^{ad/\mu}}^{r^d} z^{\alpha-\tau} (N^{ad} z^{-\mu})^\beta \ell^*(z) dz \\ &= c_2 N^{ad\beta} r^{-d\alpha} \int_{z=t_0^{-1/\mu} N^{ad/\mu}}^{r^d} z^{\alpha-\tau-\mu\beta} \ell^*(z) dz. \end{aligned}$$

If $\alpha - \tau - \mu\beta + 1 > 0$, we can again use Proposition 1.5.8 of [8] on the integral to get that

$$\tilde{I}_2 \leq c_2 N^{ad\beta} r^{-d\alpha} C_{12} r^{d(\alpha-\tau-\mu\beta+1)} \ell^*(r^d) \leq C_{13} N^{ad\beta} r^{-d(\tau+\mu\beta-1)+\varepsilon/2} \tag{2.12}$$

for N large enough by Potter's bound. If $\alpha - \tau - \mu\beta + 1 = 0$, we use $\ell^*(z) = o(z^{\frac{\varepsilon}{2d}})$ to get

$$\tilde{I}_2 \leq C_{14} N^{ad\beta} r^{-d\alpha} \int_{z=t_0^{-1/\mu} N^{ad/\mu}}^{r^d} z^{\frac{\varepsilon}{2d}-1} dz \leq C_{15} N^{ad\beta} r^{-d(\tau-1+\mu\beta)+\varepsilon/2}. \tag{2.13}$$

Finally, when $\alpha - \tau - \mu\beta + 1 < 0$, we use Potter's bound to get $\ell^*(z) = o(z^{\frac{\mu\varepsilon}{2ad}})$, and since in this case we assume that (2.3) holds, the integral is convergent and we have

$$\tilde{I}_2 \leq C_{16} N^{ad\beta} r^{-d\alpha} \int_{z=t_0^{-1/\mu} N^{ad/\mu}}^{r^d} z^{\alpha-\tau-\mu\beta+\frac{\mu\varepsilon}{2ad}} dz \leq C_{17} r^{-d\alpha} N^{\frac{ad}{\mu}(\alpha-\tau+1)+\varepsilon/2}. \quad (2.14)$$

Combining equations (2.12)–(2.14), and using the fact that $r \geq N$ in the inner integral, we get

$$\tilde{I}_2 \leq C_{18} \left(N^{ad\beta} r^{-d(\tau-1+\mu\beta)} + r^{-d\alpha} N^{\frac{ad}{\mu}(\alpha-\tau+1)} \right) r^{\varepsilon/2}. \quad (2.15)$$

Finally, the third part of the integral in $\Lambda(r)$ in (2.6) is given by

$$\tilde{I}_3 := \int_{z=r^d}^{\infty} \left(1 \wedge \frac{z}{r^d} \right)^\alpha \cdot F_L(N^{ad} z^{-\mu}) \cdot z^{-\tau} \ell^*(z) dz \leq c_2 N^{ad\beta} \int_{z=r^d}^{\infty} z^{-\mu\beta-\tau} \ell^*(z) dz.$$

Again, we can apply Proposition 1.5.10 of [8] to treat the inner integral and get, by Potter's bound when N is sufficiently large,

$$\tilde{I}_3 \leq C_{19} N^{ad\beta} r^{-d(\tau-1+\mu\beta)} \ell^*(r^d) \leq C_{20} N^{ad\beta} r^{-d(\tau-1+\mu\beta)+\varepsilon/2}. \quad (2.16)$$

Combining (2.11), (2.15) and (2.16), while keeping in mind that $r \geq N$, we get

$$\Lambda(r) = \tilde{I}_1 + \tilde{I}_2 + \tilde{I}_3 \leq C_{21} \left(r^{-d\alpha} + r^{-d\alpha} N^{\frac{ad}{\mu}(\alpha-\tau+1)} + N^{ad\beta} r^{-d(\tau+\mu\beta-1)} \right) r^{\varepsilon/2}.$$

Returning again to (2.5), using this bound on $\Lambda(r)$ we obtain

$$\begin{aligned} E(A, N, a) &\leq C_{22} A^d \int_{r=N}^{dA} r^{-d(\alpha-1)+\varepsilon/2-1} \left(1 + N^{\frac{ad}{\mu}(\alpha-\tau+1)} \right) + r^{-d(\tau-2+\mu\beta)+\varepsilon/2-1} N^{ad\beta} dr \\ &\stackrel{(2.2)}{\leq} C_{23} A^d N^{\varepsilon/2} \left(N^{-d(\alpha-1)} + N^{-d(\alpha-1-\frac{a}{\mu}(\alpha-\tau+1))} + N^{-d(\tau-2+(\mu-a)\beta)} \right) \\ &\leq A^d N^\varepsilon \left(N^{-d(\alpha-1)} + N^{-d(\alpha-1-\frac{a}{\mu}(\alpha-\tau+1))} + N^{-d(\tau-2+(\mu-a)\beta)} \right), \end{aligned}$$

where the last inequality holds for N large enough. Thus we have proved (2.1) when $a < \mu$. □

2.2 Good blocks

In the renormalisation scheme, we shall cover the space with blocks (i.e., boxes) that are iteratively contained in larger and larger blocks. The smallest blocks are level-1 blocks, and we group them together in level-2 blocks, and so on. A level- k block thus contains sub-blocks of level $k - 1$, which in turn contains sub-blocks of level $k - 2$, and so on until we reach the level-1 blocks. Now we turn to the definition of good blocks. Berger [6] defined them as blocks that do not contain edges of linear length (in the block size), and additionally the same property must hold for certain translates and recursively for most of its sub-blocks. In our setting, we have to modify the definition since long edges do exist. However, we can define a block to be 'good' if all edges of linear length are costly enough and arrive at the following definition.

Definition 2.3. Let $k_0 := 16 \cdot 30^d$, let $A_1 > 1$ and $u \in (0, 1)$ be constants, and define $A_k := A_1(k!)^2$ for all $k \geq k_0$.

A k -block is defined as a d -dimensional cube of side length A_k . For $k > k_0$, there is a natural partition of a k -block Q into k^{2d} many $(k - 1)$ -blocks; we call these the children of Q . We denote by Q_k the k -block centred at the origin.

We will define the notion of goodness recursively. Let $\eta \in (0, 1]$. We say that a k_0 -block Q is η -good if every edge internal to Q has cost at least u . For $k > k_0$ and a k -block Q , consider all 3^d k -blocks Q' of the form $Q + jA_{k-1}/2$ for $j \in \{-1, 0, 1\}^d$. We say that Q is η -good if for all $Q' = Q + jA_{k-1}/2$:

- (i) Every edge internal to Q' of length larger than $A_{k-1}/100$ has cost at least uA_k^η .
- (ii) All but at most 3^d children of Q' are η -good.

We will fix A_1 and then u in Definition 2.3 such that they satisfy equations (2.21), (2.24), (2.28), (2.32) and (2.36) below. We will also assume that A_1 is large enough so that the upper bound of Lemma 2.2 holds for certain choices of N and ε specified case-by-case below in (2.22), (2.23), (2.25), (2.29) and (2.30). Recall the definitions of μ_{pol} from (1.6) and η_0 from (1.7).

Proposition 2.4. Consider 1-FPP on IGIRG or SFP of Definition 1.1 satisfying the assumptions given in (1.2)–(1.3) with $0 \in \mathcal{V}$, but L not necessarily satisfying Assumption 1.3. Assume that

- 1) $\alpha > 2$, $\tau > 3$, $\mu \geq 0$ arbitrary, and the distribution of L is arbitrary satisfying $\mathbb{P}(L > 0) = 1$. Then there are choices for A_1 and u in Definition 2.3 for which a.s. there exists $k_1 \geq k_0$ such that Q_k is 1-good for all $k \geq k_1$.
- 2) $\alpha > 2$, $\tau \in (2, 3]$, $\mu > \mu_{\log}$, and L satisfies (1.4) in Assumption 1.3. Then for any $\delta > 0$ there are choices for A_1 and u in Definition 2.3 for which a.s. there exists $k_1 \geq k_0$ such that Q_k is $(\eta_0 - \delta)$ -good for all $k \geq k_1$.
- 3) $\alpha > 2$, $\tau \in (2, 3]$, $\mu > \mu_{\text{pol}}$, and L satisfies (1.4) in Assumption 1.3. Then there are choices for A_1 and u in Definition 2.3 for which a.s. there exists $k_1 \geq k_0$ such that Q_k is 1-good for all $k \geq k_1$.

Remark 2.5. Proposition 2.4 implicitly limits the type of paths present between two vertices: a path either uses short edges (of which it needs to use many if the endpoints are far away) or it uses long edges, which do have high cost. We will see that in Case 1 paths that have long edges are simply not present. In Case 3 they are present, but the cost of long edges is so high that the paths are not more efficient than paths which only use short edges (corresponding to case $\eta_0 = 1$ of linear cost-distances). Case 2 is most exotic: long edges are present, but their cost is on a polynomial scale compared to their Euclidean length, and their precise exponent will give polynomial (but nonlinear) cost-distances.

The proof we will give follows closely the one in [17, Lemma 14]. We show that

$$\mathbb{P}(Q_k \text{ is not } (\eta_0 - \delta)\text{-good}) \leq e^{-k} \tag{2.17}$$

for $k \geq k_0$, where δ is an arbitrary positive constant for Case 2 and $\delta = 0$ for Cases 1 and 3. (The same inequality holds true under both Palm measures $\mathbb{P}^0, \mathbb{P}^{0,x}$ as well). Thus

$$\mathbb{P}(\forall k \geq k' : Q_k \text{ is } (\eta_0 - \delta)\text{-good}) \geq 1 - \sum_{k \geq k'} e^{-k} \geq 1 - 2e^{-k'}.$$

The existence of k_1 (so that Q_k is $(\eta_0 - \delta)$ -good for all $k \geq k_1$) then follows from the Borel-Cantelli Lemma, and for a fixed $q > 0$ we can achieve $\mathbb{P}(k_1 \leq k') \geq 1 - q$ for $k' = \lceil \log(q/2) \rceil$. We will show in the later parts of the section that the lower bounds on cost-distances follow *deterministically* from Proposition 2.4, for all vertices in distance at least $r_0 := c \cdot A_{k_1}$ for a constant c , and thus for all vertices in distance $r \geq r_0 = c \cdot A_{\lceil \log(q/2) \rceil} = cA_1 \cdot (\lceil \log(q/2) \rceil!)^2 \geq e^{c' \cdot \log(1/q) \cdot \log \log(1/q)}$ for a constant $c' > 0$. However, we did not try to optimise these bounds.

Compared to [17] developed for graph distances, in order to show (2.17) we additionally need to derive an upper bound on the probability that property (i) of Definition 2.3 fails, i.e., to handle the cost of edges as well, not just their length. We do this using Lemma 2.2.

Proof of Proposition 2.4. In order to unify notation, we define $\eta_0 := 1$ and $\delta := 0$ in Cases 1 and 3, which is also consistent with the definition of η_0 in (1.7). So we need to show $(\eta_0 - \delta)$ -goodness in all three cases. We fix the values of A_1 and u from Definition 2.3. They will both depend on the set par of model parameters and on δ (in Case 2). We will first choose A_1 to be suitably large, and then choose u to be suitably small as a function of A_1 . We will prove the result using the Borel-Cantelli Lemma, by showing that $\sum_{k \geq k_0} \mathbb{P}(Q_k \text{ is not } (\eta_0 - \delta)\text{-good})$ is a finite sum. By writing $\mathbb{P}^{\max}(\cdot) := \max\{\mathbb{P}(\cdot), \mathbb{P}^0(\cdot), \mathbb{P}^{0,x}(\cdot)\}$, let

$$\psi_k := \mathbb{P}^{\max}(Q_k \text{ is not } (\eta_0 - \delta)\text{-good}). \tag{2.18}$$

This choice ensures that k -blocks appearing in larger k' -blocks not centred at 0 are also bad with probability at most ψ_k . We will show inductively that $\psi_k \leq e^{-k}$. Note that our choice of $k_0 = 16 \cdot 30^d$ in Definition 2.3 ensures that

$$3^d k^{4d} e^{-2(k-1)} \leq \frac{1}{2} e^{-k} \quad \text{for all } k \geq k_0. \tag{2.19}$$

Base case. We start by bounding ψ_{k_0} . Let $\mathbb{E}^0, \mathbb{E}^{0,x}$ denote expectation under the Palm measures $\mathbb{P}^0, \mathbb{P}^{0,x}$, respectively, and denote by $\mathbb{E}^{\max}(\cdot)$ the maximum of $\mathbb{E}[\cdot], \mathbb{E}^0[\cdot], \mathbb{E}^{0,x}[\cdot]$. Let $|\mathcal{E}(Q_{k_0})|$ be the number of edges inside Q_{k_0} . Then $|\mathcal{E}(Q_{k_0})|$ has finite expectation because the number of vertices in Q_k is fixed (in SFP) or has a finite second moment (in IGIRG), and the number of edges with both endpoints within Q_k is at most the square of the number of the vertices. The same is true under both Palm measures $\mathbb{P}^0, \mathbb{P}^{0,x}$. Let

$$T := \mathbb{E}^{\max}[|\mathcal{E}(Q_{k_0})|]. \tag{2.20}$$

We always assume $\mathbb{P}(L > 0) = 1$ (Assumption 1.3 is a stronger assumption), so we can choose the constant $u \in (0, 1)$ from Definition 2.3 small enough so that

$$F_L(u) \leq e^{-2k_0}/(4T). \tag{2.21}$$

Since the vertex-weights are always at least 1 and $\mu \geq 0$, for each edge it holds that $\mathcal{C}(e) \geq L_e$. Hence by Markov's inequality and a union bound,

$$\begin{aligned} \psi_{k_0} &= \mathbb{P}^{\max}(\exists e \in Q_{k_0} \text{ with } \mathcal{C}(e) < u) \leq \mathbb{P}^{\max}(\exists e \in Q_{k_0} \text{ with } L_e \leq u) \\ &\leq \mathbb{P}^{\max}(|\mathcal{E}(Q_{k_0})| > 2e^{k_0}T) + \mathbb{P}^{\max}(\exists e \in Q_{k_0} \text{ with } L_e \leq u \mid |\mathcal{E}(Q_{k_0})| \leq 2e^{k_0}T) \\ &\leq \frac{1}{2} e^{-k_0} + 2e^{k_0} \mathbb{E}^{\max}[|\mathcal{E}(Q_{k_0})|] \cdot \frac{e^{-2k_0}}{4T} = e^{-k_0}. \end{aligned}$$

Bounding the failure probability of property (i). For the induction step, we need to derive an upper bound for the probability that property (i) in the definition of $(\eta_0 - \delta)$ -goodness fails. Since we work under the conditioning that $0 \in \mathcal{V}$, (and later possibly under the presence of an extra vertex at $x \in \mathbb{R}^d$), not all blocks have the same distribution under these measures. Yet, since Lemma 2.2 is valid with or without the presence of these vertices, the same upper bound holds for all blocks when counting the expected number of edges that are long but too cheap. Using Markov's inequality and translation invariance, for any fixed $j \in \{-1, 0, 1\}^d$ we have

$$\begin{aligned} &\mathbb{P}^{\max}(Q_k + jA_{k-1}/2 \text{ fails to have property (i)}) \\ &\leq \mathbb{E}^{\max}[\{|\{vw \subset Q_k \text{ edge} : |v - w| \geq A_{k-1}/100, \mathcal{C}(vw) \leq uA_k^{\eta_0 - \delta}\}|] =: E(k). \end{aligned}$$

We will bound $E(k)$ from above using Lemma 2.2. We will specify ε, a in Lemma 2.2 later, only depending on par. Then we will set $A := A_k, N := A_{k-1}/100 < A$, and choose A_1

large enough (with respect to ε, a) so that Lemma 2.2 can be applied. When we are able to set $a \geq \mu$, with these choices of A, N then we obtain, for some constant $C > 0$ that depends only on par,

$$E(k) \leq A_k^d \cdot \frac{A_{k-1}^\varepsilon}{100^\varepsilon} \left(\frac{A_{k-1}^{-d(\alpha-1)}}{100^{-d(\alpha-1)}} + \frac{A_{k-1}^{-d(\tau-2)}}{100^{-d(\tau-2)}} \right) \tag{2.22}$$

$$\leq 100^C k^{2d} \left(A_{k-1}^{-d(\alpha-2)+\varepsilon} + A_{k-1}^{-d(\tau-3)+\varepsilon} \right),$$

where we substituted $A_k = k^2 A_{k-1}$ for the second inequality. Recall that this bound holds for any L and does not require Assumption 1.3. Observe that the second term only tends to zero when $\tau > 3$, so we can set $a \geq \mu$ only when $\tau > 3$. For $\tau \leq 3$, we will have to set $a < \mu$, again $A := A_k, N := A_{k-1}/100 < A$. By possibly increasing the constant $C > 0$, if A_1 is sufficiently large so that Lemma 2.2 can be applied then (2.1) yields

$$E(k) \leq A_k^d \cdot \frac{A_{k-1}^\varepsilon}{100^\varepsilon} \left(\frac{A_{k-1}^{-d(\alpha-1)}}{100^{-d(\alpha-1)}} + \frac{A_{k-1}^{-d(\alpha-1-\frac{a}{\mu}(\alpha-(\tau-1)))}}{100^{-d(\alpha-1-\frac{a}{\mu}(\alpha-(\tau-1)))}} + \frac{A_{k-1}^{-d(\tau-2+(\mu-a)\beta)}}{100^{-d(\tau-2+(\mu-a)\beta)}} \right) \tag{2.23}$$

$$\leq 100^C k^{2d} \left(A_{k-1}^{-d(\alpha-2)+\varepsilon} + A_{k-1}^{-d(\alpha-2-\frac{a}{\mu}(\alpha-(\tau-1)))+\varepsilon} + A_{k-1}^{-d(\tau-3+(\mu-a)\beta)+\varepsilon} \right),$$

Note that this case of Lemma 2.2 required Assumption 1.3 on L .

We now split into cases depending on the values of τ and μ to specify a and ε . Proposition 2.4 always assumes $\alpha > 2$ but we emphasise this for readability.

Case 1: $\tau > 3, \alpha > 2$, and $\mu \geq 0, L > 0$ a.s., otherwise arbitrary. In this case we have $\eta_0 = 1$ and $\delta = 0$. Let us choose $a > \max\{\mu, 1/d\}$. Then since $ad > 1 = \eta_0$ we can choose A_1 so large that

$$\left(\frac{A_{k-1}}{100} \right)^{ad} \geq A_{k-1} k^2 = A_k \quad \text{for all } k > k_0. \tag{2.24}$$

Since $\tau > 3, \alpha > 2$, we will set in (2.22)

$$\varepsilon := \frac{d}{2} \min \{ \alpha - 2, \tau - 3 \} > 0. \tag{2.25}$$

Substituting this ε into the upper bound (2.22) yields

$$E(k) \leq 2 \cdot 100^C k^{2d} A_{k-1}^{-\varepsilon} \leq 100^{C+1} A_1^{-\varepsilon} k^{2d} ((k-1)!)^{-2\varepsilon}.$$

Case 2: $\tau \in (2, 3], \alpha > 2$, and $\mu \in (\mu_{\log}, \mu_{\text{pol}}]$. In this case, $\delta > 0$ and Assumption 1.3 needs to hold for L . Define

$$a := \min \left\{ \mu - \frac{3-\tau}{\beta}, \frac{\mu(\alpha-2)}{\alpha-(\tau-1)} \right\} - \frac{\delta}{2d} = \frac{\eta_0}{d} - \frac{\delta}{2d}, \tag{2.26}$$

where the last equation can be seen by using η_0 from (1.7) and $\mu_{\log} = (3-\tau)/\beta$ from (1.6). Also note that $\eta_0 > 0$ since $\mu > \mu_{\log} = (3-\tau)/\beta$, and $\alpha > 2 \geq (\tau-1)$, and that the statement of Proposition 2.4 is stronger for smaller δ . Therefore, we can assume that $a > 0$. Note also that the first term of the minimum is at most μ and $\delta > 0$, which implies $a < \mu$. (In fact, the second term of the minimum is also at most μ .) Moreover, by the definition of a and the fact that $\alpha > 2$ implies $\alpha - (\tau - 1) > 0$, rearranging (2.26) yields that

$$3 - \tau - (\mu - a)\beta \leq -\frac{\beta\delta}{2d}, \quad 2 - \alpha + \frac{a}{\mu}(\alpha - (\tau - 1)) \leq -\frac{(\alpha - (\tau - 1))\delta}{2\mu d}. \tag{2.27}$$

We choose A_1 large enough so that

$$\left(\frac{A_{k-1}}{100}\right)^{\eta_0-\delta/2} = \left(\frac{A_k}{100k^2}\right)^{\eta_0-\delta/2} \geq A_k^{\eta_0-\delta} > uA_k^{\eta_0-\delta} \quad \text{for all } k > k_0. \quad (2.28)$$

By definition of a , cf. equation (2.26), $(A_{k-1}/100)^{ad} = (A_{k-1}/100)^{\eta_0-\delta/2} > uA_k^{\eta_0-\delta}$ as desired. Using $\alpha - (\tau - 1) > 0$ the first term in the upper bound (2.23) is dominated by the second term, and (2.23) becomes

$$E(k) \leq 2 \cdot 100^C k^{2d} \left(A_{k-1}^{-d(\alpha-2-\frac{a}{\mu}(\alpha-\tau+1))+\varepsilon} + A_{k-1}^{-d(\tau-3+(\mu-a)\beta)+\varepsilon} \right).$$

Now we set

$$\varepsilon := \frac{\delta}{4} \min \left\{ \beta, \frac{\alpha - (\tau - 1)}{\mu} \right\} > 0. \quad (2.29)$$

Using this ε , combined with (2.27), it follows that

$$E(k) \leq 100^{C+1} k^{2d} A_{k-1}^{-\varepsilon} = 100^{C+1} A_1^{-\varepsilon} k^{2d} ((k-1)!)^{-2\varepsilon}.$$

Case 3: $\tau \in (2, 3]$, $\alpha > 2$, and $\mu > \mu_{\text{pol}}$. Here again we need Assumption 1.3 to hold for L , and we have $\eta_0 = 1$ and $\delta = 0$. Remembering the definition of μ_{pol} in (1.6), rearranging $\mu > \mu_{\text{pol}}$ yields that both $d(\alpha-2) - (\alpha-\tau+1)/\mu > 0$ and $d(\tau-3+(\mu-1/d)\beta) > 0$ hold. So we will show that we can set in (2.23) the following values of ε and a :

$$\varepsilon := \frac{1}{4} \min \left\{ d(\alpha-2) - \frac{\alpha - (\tau - 1)}{\mu}, d(\tau - 3 + (\mu - 1/d)\beta) \right\} > 0, \quad (2.30)$$

and

$$a := \frac{1}{2} \left(\frac{1}{d} + \min \left\{ \frac{\mu(\alpha-2)}{\alpha - (\tau - 1)}, \mu - \frac{3 - \tau}{\beta} \right\} \right) > 0. \quad (2.31)$$

We establish first that indeed $a < \mu$ and also that $ad > 1$. To see the first, observe that a averages $1/d$ with the minimum of two different quantities. By (1.6) we have $1/d \leq \mu_{\text{pol}} < \mu$. Since $\tau \in (2, 3]$, $(\alpha - 2)/(\alpha - (\tau - 1)) \leq 1$ and so $\mu(\alpha - 2)/(\alpha - \tau + 1) \leq \mu$, and $\mu - (3 - \tau)/\beta \leq \mu$. This implies in particular that a averages one quantity strictly smaller than μ with another quantity smaller or equal to μ , showing that $a < \mu$. Since $\mu > \mu_{\text{pol}}$, rearranging also yields

$$\frac{\mu(\alpha-2)}{\alpha - (\tau - 1)} > \frac{1}{d} \quad \text{and} \quad \mu - \frac{3 - \tau}{\beta} > \frac{1}{d},$$

so $a > 1/d$. This implies $ad > 1 = \eta_0$, and thus, we can choose A_1 so large that

$$\left(\frac{A_{k-1}}{100}\right)^{ad} \geq A_k = A_k^{\eta_0-\delta} > uA_k^{\eta_0-\delta} \quad \text{for all } k > k_0. \quad (2.32)$$

Since $\alpha > 2$ and $\tau \leq 3$, we have $\alpha - (\tau - 1) > 0$ and thus, as in Case 2 the upper bound (2.23) becomes

$$E(k) \leq 2 \cdot 100^C k^{2d} \left(A_{k-1}^{-d(\alpha-2-\frac{a}{\mu}(\alpha-\tau+1))+\varepsilon} + A_{k-1}^{-d(\tau-3+(\mu-a)\beta)+\varepsilon} \right). \quad (2.33)$$

The definition of ε in (2.30) and a in (2.31) implies $2\varepsilon \leq \frac{1}{2}d(\tau - 3 + (\mu - 1/d)\beta)$ and $a \leq \frac{1}{2}(1/d + \mu - (3 - \tau)/\beta)$. From these, it is elementary to check that

$$-d(\tau - 3 + (\mu - a)\beta) \leq -2\varepsilon.$$

Likewise, from $2\varepsilon \leq \frac{1}{2}(d(\alpha - 2) - (\alpha - (\tau - 1)/\mu))$ and $a \leq \frac{1}{2}(1/d + \mu(\alpha - 2)/(\alpha - \tau + 1))$ we derive

$$-d \left(\alpha - 2 - \frac{a}{\mu}(\alpha - \tau + 1) \right) \leq -2\varepsilon,$$

so (2.33) simplifies to

$$E(k) \leq 100^{C+1} k^{2d} A_{k-1}^{-\varepsilon} = 100^{C+1} A_1^{-\varepsilon} k^{2d} ((k-1)!)^{-2\varepsilon}.$$

We remark that our approach does not give $\delta = 0$ in the boundary case $\mu = \mu_{\text{pol}}$ for $\tau \in (2, 3)$, which would yield linear distances. The reason is that for $\delta = 0$ we can not find a value of a such that $ad > 1$ holds, and which give a negative exponent in the error bound (2.1). Since Lemma 2.2 is essentially tight, this means that there exist long edges whose cost is sublinear in their length.

Induction step. Combining the three cases above, we conclude that in all cases we have for an appropriately chosen $\varepsilon > 0$ (depending on the values of μ and τ)

$$\mathbb{P}^{\max}(Q_k + jA_{k-1}/2 \text{ fails to have property (i)}) \leq E(k) \leq 100^{C+1} A_1^{-\varepsilon} k^{2d} ((k-1)!)^{-2\varepsilon}. \tag{2.34}$$

Importantly, both C and ε depend only on par . We define

$$C' = C'(A_1) := 100^{C+1} A_1^{-\varepsilon}, \tag{2.35}$$

which is a constant depending on par and A_1 .

We note that for $k > k_0$, the k -block Q_k is not $(\eta_0 - \delta)$ -good if at least one of its 3^d translations $Q_k + jA_{k-1}/2, j \in \{-1, 0, 1\}^d$ fails to have property (i) or (ii) of Definition 2.3. The failure probability of property (i) is bounded by (2.34). Therefore, translation invariance together with a union bound, and recalling property (ii), implies that

$$\psi_k \leq 3^d (C' k^{2d} ((k-1)!)^{-2\varepsilon} + \mathbb{P}(\text{at least } 3^d + 1 \text{ children of } Q_k \text{ are not } (\eta_0 - \delta)\text{-good})).$$

Observe the following: if a block has at least $3^d + 1$ bad child-blocks, one can find at least one pair of non-neighbouring bad child-blocks. Hence the event in the probability above implies that there are at least two children-blocks Q_{k-1}^* and Q_{k-1}^{**} of Q_k that are not $(\eta_0 - \delta)$ -good and whose centres have distance at least $2A_{k-1}$. In particular, any vertex in Q_{k-1}^* is at least A_{k-1} Euclidean distance away from any vertex in Q_{k-1}^{**} , which ensures that the events $\{Q_{k-1}^* \text{ is not } (\eta_0 - \delta)\text{-good}\}$ and $\{Q_{k-1}^{**} \text{ is not } (\eta_0 - \delta)\text{-good}\}$ are independent. Since there are k^{2d} many child-blocks of Q_k , there is at most $\binom{k^{2d}}{2} \leq k^{4d}$ many ways to choose Q_{k-1}^* and Q_{k-1}^{**} , that are independently bad with probability at most ψ_{k-1} . So we deduce the bound

$$\psi_k \leq 3^d (C' k^{2d} ((k-1)!)^{-\varepsilon} + k^{4d} \psi_{k-1}^2).$$

Note that $C' \rightarrow 0$ as $A_1 \rightarrow \infty$, by (2.35). Thus, we can choose A_1 large enough (and this choice only depends on par and ε) so that for all $k > k_0$,

$$\psi_k \leq \frac{1}{2} e^{-k} + 3^d k^{4d} \psi_{k-1}^2. \tag{2.36}$$

We now prove by induction on k that

$$\psi_k \leq e^{-k} \tag{2.37}$$

for all $k \geq k_0$. Indeed, we have already seen that (2.37) holds for $k = k_0$. For $k > k_0$, assuming that (2.37) holds for $k - 1$, we get

$$\psi_k \leq \frac{1}{2}e^{-k} + 3^d k^{4d} \psi_{k-1}^2 \leq \frac{1}{2}e^{-k} + 3^d k^{4d} e^{-2(k-1)} \stackrel{(2.19)}{\leq} \frac{1}{2}e^{-k} + \frac{1}{2}e^{-k} = e^{-k}.$$

This shows by induction that (2.37) holds for all $k \geq k_0$, which yields

$$\sum_{k \geq k_0} \mathbb{P}(Q_k \text{ is not } (\eta_0 - \delta)\text{-good}) = \sum_{k \geq k_0} \psi_k \leq \sum_{k \geq k_0} e^{-k} < +\infty$$

and concludes the proof of Proposition 2.4 by the Borel-Cantelli Lemma. □

2.3 Paths within good blocks

The next proposition carries out the renormalisation scheme and is an important step in proving Theorem 1.8 and Theorem 1.4.

Proposition 2.6. *Let $\eta \in (0, 1]$ and recall k_0 from Definition 2.3. There exists a constant u^* depending only on η and η such that for all $k \geq k_0$ the following holds.*

(1) *If Q is an η -good k -block and $x, y \in Q$ satisfy $|x - y| > A_k/16$, then every path from x to y within Q has cost at least $u^*|x - y|^\eta$ (deterministically).*

(2) *If the $(k - 1)$ -block Q' and the k -block Q centred at x are both η -good, and if $y \in Q$ satisfies $|x - y| > A_{k-1}/8$, then every path from x to y within Q has cost at least $u^*|x - y|^\eta/30^{d+2}$ (deterministically).*

Remark 2.7. Our proof of (1) is an adaptation of the proof of Lemma 2 in [6], with the difference that the continuous-valued edge-cost make the argument slightly more complicated. The statement of (2) allows us to prove strictly linear cost-distances in the case $\eta = 1$, avoiding Kingman’s subadditive ergodic theorem that finishes the proof in [6]. Note that since (2) is symmetric, we could also require a condition for the blocks centred at y instead of x .

Proof of Proposition 2.6. We will show the following claim by induction: there exists a constant C_1 (which depends on the same parameters as u^* does) such that for every $k \geq k_0$, if Q is an η -good k -block and $x, y \in Q$ satisfy $|x - y| > A_k/16$, then every path π from x to y within Q has cost at least

$$C(\pi) \geq C_1 \Lambda(k) |x - y|^\eta, \quad \text{where} \quad \Lambda(k) := \prod_{h=k_0}^k \left(1 - \frac{k_0}{h^2}\right). \quad (2.38)$$

Then taking $u^* = C_1 \prod_{h=k_0}^\infty (1 - k_0/h^2) > 0$ shows (1). To show (2), we will slightly modify the last step of the induction in (1). To ease notation, we will assume that $Q = Q_k$, i.e., Q is the k -block centred at the origin.

For the base case, consider $x, y \in Q_{k_0}$ that satisfy $|x - y| > A_{k_0}/16$, where Q_{k_0} is a η -good k_0 -block. Then any path between x and y contains at least one edge and since Q_{k_0} is an η -good k_0 -block, the path has cost at least u by Definition 2.3. It is here that we use the assumption that $\mathbb{P}(L > 0) = 1$, compare it to the base case in the proof of Proposition 2.4, i.e., at and below (2.21).

Since $|x - y| \leq A_{k_0} \sqrt{d}$, and $\eta \leq 1$, we define $C_1 := u/(A_{k_0} \sqrt{d})$ so that (2.38) holds for $k = k_0$. Note that $C_1 \leq u/\sqrt{d}$.

Now assume that (2.38) holds already until $k - 1$ for some $k > k_0$. Let us consider an η -good k -block Q_k , and $x, y \in Q_k$ with $|x - y| > A_k/16$. Let $\pi = (v_1, \dots, v_l)$ be a path from $v_1 = x$ to $v_l = y$ within Q_k .

Case A: If π contains one or more edges of length greater than $A_{k-1}/100$, then since Q_k is η -good, by property (i) of Definition 2.3, the path has cost at least

$$uA_k^\eta \geq \frac{u|x-y|^\eta}{\sqrt{d}^\eta} \geq C_1|x-y|^\eta \geq C_1\Lambda(k)|x-y|^\eta,$$

since $\eta \leq 1$, so we are done.

Case B: When every edge of π has length at most $A_{k-1}/100$, we will split π in subpaths as follows. Remember that for every translation $Q_k + jA_{k-1}/2$ of Q_k , $j \in \{-1, 0, 1\}^d$, at most 3^d children of $Q_k + jA_{k-1}/2$ are not η -good (call these bad), thus we have in total at most 9^d bad child-blocks in the union of all translations (some might overlap). Denote these bad children by B_1, B_2, \dots, B_p with $p \leq 9^d$, and let $B := B_1 \cup \dots \cup B_p$. If $\pi \cap B = \emptyset$, define $\pi_1 := \pi$. Otherwise, we will decompose π into *good segments* π_s followed by *bad segments* σ_t , so that π is the concatenation of the good and bad segments: $\pi = (\pi_1, \sigma_1, \pi_2, \sigma_2, \dots, \pi_{S-1}, \sigma_T, \pi_S)$ for some S, T , where some of these segments might be empty, and the last vertex of a segment is the first vertex of the next segment.

We now divide the vertices of $\pi = (v_1, \dots, v_l)$ into the segments. Intuitively, the edge-set of good segments stay fully outside of bad child-blocks. Let a_1 be the smallest index $i \leq l$ so that $v_i \in B$ and let b_1 be the index of the containing bad child-blocks: $v_{a_1} \in B_{b_1}$ (choose b_1 arbitrarily if there is more than one possibility). Then we set $\pi_1 = (v_1, \dots, v_{a_1-1})$ as the first good segment. Let z_1 be the largest value z such that $v_z \in B_{b_1}$, then $\sigma_1 = (v_{a_1-1}, \dots, v_{z_1+1})$ is the first bad segment. (Note that there may be vertices on this segment outside B_{b_1} .) Inductively, let a_{s+1} be the smallest $a > z_s$ so that $v_a \in B$, let b_{s+1} be so that $v_{a_{s+1}} \in B_{b_{s+1}}$ and let z_{s+1} be the largest z with $v_z \in B_{b_{s+1}}$. The further good segments are then defined as $\pi_2 := (v_{z_1+1}, \dots, v_{a_2-1})$ and so on, up to π_S , and the bad segments as $\sigma_1 := (v_{a_1-1}, \dots, v_{z_1+1}), \sigma_2 := (v_{a_2-1}, \dots, v_{z_2+1})$ and so on up to σ_T . Observe that the bad segment σ_t contains the two edges $v_{a_t-1}v_{a_t}$ and $v_{z_t}v_{z_t+1}$ not fully contained in B_{b_t} . So, the good segments contain edges that are *completely outside* the bad set B . Note that $S-1 \leq T \leq p \leq 9^d$, and $S-1$ may or may not equal T since two bad segments may directly follow each other.

For a path ρ , denote by $\mathcal{D}(\rho)$ the Euclidean distance between its endpoints. Moreover, for $v, w \in \rho$, let $\rho[v, w]$ be the subpath (segment) from v to w on ρ . By the triangle inequality, $|x-y| \leq \sum_{s=1}^S \mathcal{D}(\pi_s) + \sum_{t=1}^T \mathcal{D}(\sigma_t)$. Moreover, since the diameter of B_{b_t} is $\sqrt{d}A_{k-1}$ (since B_{b_t} is a level $k-1$ block), and every edge in π has length at most $A_{k-1}/100$ (by assumption of Case B), for every bad segment σ_t we have

$$\mathcal{D}(\sigma_t) \leq |v_{a_t} - v_{z_t}| + |v_{a_t-1} - v_{a_t}| + |v_{z_t} - v_{z_t+1}| \leq \sqrt{d}A_{k-1} + 2A_{k-1}/100 \leq 2dA_{k-1}. \tag{2.39}$$

Let

$$I := \{s \in \{1, \dots, S\} \mid \mathcal{D}(\pi_s) > A_{k-1}/2\} \tag{2.40}$$

be those good segments where the two endpoints span at least $A_{k-1}/2$ Euclidean distance. Then $\sum_{s \notin I} \mathcal{D}(\pi_s) \leq SA_{k-1}/2$, trivially, since there are S good segments. Since the inductive statement assumes $|x-y| > A_k/16$, this and (2.39) gives a somewhat involved pigeon-hole principle on the triangle inequality above:

$$\begin{aligned} \sum_{s \in I} \mathcal{D}(\pi_s) &\geq |x-y| - \sum_{t=1}^T \mathcal{D}(\sigma_t) - \sum_{s \notin I} \mathcal{D}(\pi_s) \geq |x-y| - (2dT + S/2)A_{k-1} \\ &\stackrel{(\star)}{\geq} |x-y| - 30^d A_{k-1} \stackrel{(\dagger)}{\geq} \left(1 - \frac{16 \cdot 30^d}{k^2}\right) |x-y| \stackrel{(\square)}{=} \left(1 - \frac{k_0}{k^2}\right) |x-y|, \end{aligned} \tag{2.41}$$

where we used that $S - 1, T \leq 9^d$ to get inequality (\star) , that $|x - y| \geq A_k/16$ and that $A_k = k^2 A_{k-1}$ to get (\dagger) , and then the definition of k_0 in Definition 2.3 to get (\square) . For later reference, we note that under the condition $|x - y| \geq 2 \cdot 30^d A_{k-1}$ (that we will assume in proving (2)), the following weaker version of (2.41) still holds:

$$\sum_{s \in I} \mathcal{D}(\pi_s) \geq |x - y| - 30^d A_{k-1} \geq \frac{1}{2}|x - y|. \tag{2.42}$$

Returning to the proof of (2.38), observe that (2.41) bounds the total Euclidean distance \mathcal{D} spanned by the endpoints of good segments π_s . Now we will work towards switching \mathcal{D} to \mathcal{C} , the cost of the segment. We claim that if $s \in I$, then

$$\mathcal{C}(\pi_s) \geq C_1 \Lambda(k - 1) \mathcal{D}(\pi_s)^\eta. \tag{2.43}$$

In order to prove (2.43), we will show inductively that we can partition $\pi_s = (v_{z_s+1}, \dots, v_{a_{s+1}-1})$ into sub-segments $\pi_{s,i} = \pi_s[x_{s,i}, y_{s,i}]$ (where $y_{s,i} = x_{s,i+1}$, i.e., the end-vertex of a sub-segment is the starting vertex of the next sub-segment) for $i = 1, \dots, q_s$, such that for all i :

- (I) $\mathcal{D}(\pi_{s,i}) > A_{k-1}/16$,
- (II) $\pi_{s,i} \subseteq B_{A_{k-1}/4}(x_{s,i})$,

i.e., the whole segment is contained in the Euclidean ball of radius $A_{k-1}/4$ centred around $x_{s,i}$ and its endpoints span enough Euclidean distance. We will construct the $\pi_{s,i}$ greedily, with the induction hypothesis that if π_s is *not* covered by the first i sub-segments (i.e. if $\pi_s \neq \cup_{j=1}^i \pi_{s,j}$), then $|y_{s,i} - v_{a_{s+1}-1}| > A_{k-1}/16$ (remembering that $v_{a_{s+1}-1}$ is the last vertex of π_s).

For the base case $i = 0$, we take $y_{s,0}$ to be the first vertex of π_s , and since $s \in I$, by (2.40), we have $|y_{s,0} - v_{a_{s+1}-1}| = \mathcal{D}(\pi_s) > A_{k-1}/2 > A_{k-1}/16$, so the induction hypothesis is satisfied.

Now assume by induction that we already have $\pi_{s,1}, \dots, \pi_{s,i}$ satisfying (I), (II) and $|y_{s,i} - v_{a_{s+1}-1}| > A_{k-1}/16$ for some $i \geq 1$, and let us construct $\pi_{s,i+1}$.

Case B1: If the segment $\pi_s[y_{s,i}, v_{a_{s+1}-1}]$ satisfies (II), that is, if $\pi_s[y_{s,i}, v_{a_{s+1}-1}] \subseteq B_{A_{k-1}/4}(y_{s,i})$ then we define $q_s := i + 1$, $\pi_{s,q_s} := \pi_s[y_{s,i}, v_{a_{s+1}-1}]$ and the procedure terminates. π_{s,q_s} also satisfies (I) since $\mathcal{D}(\pi_{s,q_s}) = |y_{s,i} - v_{a_{s+1}-1}| > A_{k-1}/16$ because of our assumption that $q_s > i$.

Case B2: If $\pi_s[y_{s,i}, v_{a_{s+1}-1}] \not\subseteq B_{A_{k-1}/4}(y_{s,i})$ we distinguish two cases depending on $|y_{s,i} - v_{a_{s+1}-1}|$.

Case B2a: $|y_{s,i} - v_{a_{s+1}-1}| \geq 5A_{k-1}/32$. Define $\pi_{s,i+1}$ to be the path obtained by following π_s from $x_{s,i+1} := y_{s,i}$ until reaching the first vertex on π_s that spans larger than $A_{k-1}/16$ Euclidean distance from $x_{s,i+1}$, and let this vertex be $y_{s,i+1}$. The vertex $y_{s,i+1}$ exists since $|x_{s,i+1} - v_{a_{s+1}-1}| = |y_{s,i} - v_{a_{s+1}-1}| > A_{k-1}/16$ was our assumption. Then $\pi_{s,i+1}$ satisfies (I) by our definition of $y_{s,i+1}$. Since every edge of π_s has length at most $A_{k-1}/100$ (since we are under Case B), and every vertex of $\pi_{s,i+1} = \pi_s[x_{s,i+1}, y_{s,i+1}]$ except $y_{s,i+1}$ is within distance $A_{k-1}/16$ of $x_{s,i+1}$, $\pi_{s,i+1}$ is also contained in a ball of radius $A_{k-1}/16 + A_{k-1}/100 < A_{k-1}/4$ centred at $x_{s,i+1}$, and thus (II) is satisfied as well. Finally, by the triangle inequality we have

$$\begin{aligned} |y_{s,i+1} - v_{a_{s+1}-1}| &\geq |y_{s,i} - v_{a_{s+1}-1}| - |y_{s,i} - y_{s,i+1}| \\ &\geq 5A_{k-1}/32 - (A_{k-1}/16 + A_{k-1}/100) > A_{k-1}/16, \end{aligned}$$

which shows the induction step for this case.

Case B2b: $|y_{s,i} - v_{a_{s+1}-1}| < 5A_{k-1}/32$. Intuitively, this case means that while $y_{s,i}$ is close to $v_{a_{s+1}-1}$ in space, the path $\pi_s[y_{s,i}, v_{a_{s+1}-1}]$ wanders far off and then comes back before reaching the last vertex $v_{a_{s+1}-1}$. Let y' be the first vertex on the path π_s starting from $x_{s,i+1} := y_{s,i}$ such that $|y' - x_{s,i+1}| \geq A_{k-1}/4$ (this vertex must exist since $\pi_s[y_{s,i}, v_{a_{s+1}-1}] \not\subseteq B_{A_{k-1}/4}(y_{s,i})$, since we are under Case B2). Then define $y_{s,i+1}$ to be the vertex right before y' on $\pi_s[x_{s,i+1}, y']$ and let $\pi_{s,i+1} := \pi_s[x_{s,i+1}, y_{s,i+1}]$. Again, since every edge of π_s has length at most $A_{k-1}/100$ (since we are under Case B), we have $\mathcal{D}(\pi_{s,i+1}) \geq |x_{s,i+1} - y'| - |y_{s,i+1} - y'| \geq A_{k-1}/4 - A_{k-1}/100 > A_{k-1}/16$ so (I) is satisfied. Moreover, $\pi_{s,i+1}$ satisfies (II) by our definition of $y_{s,i+1}$. Finally, we use the triangle inequality to get

$$\begin{aligned} |y_{s,i+1} - v_{a_{s+1}-1}| &\geq |y' - x_{s,i+1}| - |y_{s,i+1} - y'| - |v_{a_{s+1}-1} - x_{s,i+1}| \\ &\geq A_{k-1}/4 - A_{k-1}/100 - 5A_{k-1}/32 > A_{k-1}/16, \end{aligned}$$

which concludes the induction step.

The reason for requiring (II) is that the obtained partition of π_s has the property that every sub-segment $\pi_{s,i}$ is contained in a η -good $(k-1)$ -block (among the children of Q_k and its translations). This is due to a geometric statement: any ball of radius $A_{k-1}/4$ within Q_k is contained in at least one $k-1$ -level block that is a child of some shift $Q_k + jA_{k-1}/2$ for some $j \in \{-1, 0, 1\}^d$. By (I) and the induction hypothesis (2.38), we have $\mathcal{C}(\pi_{s,i}) \geq C_1\Lambda(k-1)\mathcal{D}(\pi_{s,i})^\eta$ and therefore

$$\begin{aligned} \mathcal{C}(\pi_s) &= \sum_{i=1}^{q_s} \mathcal{C}(\pi_{s,i}) \geq C_1\Lambda(k-1) \sum_{i=1}^{q_s} \mathcal{D}(\pi_{s,i})^\eta \\ &\stackrel{(\star)}{\geq} C_1\Lambda(k-1) \left(\sum_{i=1}^{q_s} \mathcal{D}(\pi_{s,i}) \right)^\eta \stackrel{(\Delta)}{\geq} C_1\Lambda(k-1)\mathcal{D}(\pi_s)^\eta, \end{aligned}$$

where we got inequality (\star) as follows: since $\eta \leq 1$ the function x^η is sublinear, i.e., $x^\eta + y^\eta \geq (x+y)^\eta$, and (Δ) is the triangle inequality applied on π_s and its sub-segments. This shows (2.43). Recall now the set I from (2.40), and the pigeon-hole argument in (2.41), which holds under the assumption $|x-y| \geq A_k/16$. By the same sublinearity argument (\star) of the function x^η , we can finally deduce that

$$\begin{aligned} \mathcal{C}(\pi) &\geq \sum_{s \in I} \mathcal{C}(\pi_s) \stackrel{(2.43)}{\geq} C_1\Lambda(k-1) \sum_{s \in I} \mathcal{D}(\pi_s)^\eta \stackrel{(\star)}{\geq} C_1\Lambda(k-1) \left(\sum_{s \in I} \mathcal{D}(\pi_s) \right)^\eta \\ &\stackrel{(2.41)}{\geq} C_1\Lambda(k-1) \left(1 - \frac{k_0}{k^2} \right)^\eta |x-y|^\eta \stackrel{\eta \leq 1}{\geq} C_1\Lambda(k)|x-y|^\eta, \end{aligned} \tag{2.44}$$

where $\Lambda(k)$ is defined in (2.38). This finishes the inductive demonstration of (2.38) and concludes the proof of part (1) of the proposition.

For part (2), we discriminate two sub-cases. If $|x-y| \geq 2 \cdot 30^d A_{k-1}$, then (2.42) holds, and we can repeat the calculation in (2.44), getting

$$\mathcal{C}(\pi) \geq C_1\Lambda(k-1) \left(\sum_{s \in I} \mathcal{D}(\pi_s) \right)^\eta \stackrel{(2.42)}{\geq} 2^{-\eta} C_1\Lambda(k-1)|x-y|^\eta,$$

which is stronger than required. So consider the remaining case, $A_{k-1}/8 < |x-y| < 2 \cdot 30^d A_{k-1}$. As before, if the path π from x to y contains an edge of length more than $A_{k-1}/100$, the claim follows from Q_k being η -good (Case A). So assume otherwise. Then there exists a vertex on π with $|x-v| \in (A_{k-1}/16, A_{k-1}/8]$. Let v be the first such vertex on π (starting from x), and let $\pi' = \pi[x, v]$. Then $v \in Q'$ (the $(k-1)$ -block centred at x)

and $\pi' \subseteq Q'$. Since Q' is good by prerequisite, we may apply part (1) to the path π' and conclude

$$\mathcal{C}(\pi) \geq \mathcal{C}(\pi') \geq u^*|x - v|^\eta \geq u^* \left(\frac{A_{k-1}}{16} \right)^\eta \geq u^* \left(\frac{|x - y|}{32 \cdot 30^d} \right)^\eta \stackrel{\eta \leq 1}{\geq} u^*|x - y|^\eta / 30^{d+2},$$

as required, finishing the proof of Proposition 2.6. \square

2.4 Proofs of lower bounds

We are now ready to prove Theorem 1.12.

Proof of Theorem 1.12. We will prove both claims (1.19) and (1.20) simultaneously by showing that the bound

$$d_C(0, x) \geq \frac{u^*|x|^\eta}{30^{d+2}\sqrt{d}} \tag{2.45}$$

holds simultaneously for all x with sufficiently large $|x| =: r$, where u^* is the constant from Proposition 2.4. The only difference between the two statements is that we prove (2.45) for all $\eta < \eta_0$ if $\mu \in (\mu_{\log}, \mu_{\text{pol}}]$, and for $\eta := \eta_0 = 1$ in the cases $\mu > \mu_{\text{pol}}$ or $\tau > 3$.

By Proposition 2.4, a.s. there is some (random) $k_1 \geq k_0$ such that for all $k \geq k_1$, the k -block Q_k centred around 0 is η -good in the sense of Definition 2.3: every edge of length larger than $A_{k-1}/100$, internal to Q_k and its shifts $Q' = Q_k + jA_{k-1}/2$ by vectors $j \in \{-1, 0, 1\}^d$, has cost at least $u^*A_k^\eta$, and there are at most 3^d bad child-blocks of Q_k and/or the Q' 's. This is the only place where we need to discriminate between the cases $\tau \in (2, 3)$ and $\mu \in (\mu_{\log}, \mu_{\text{pol}}]$, or $\tau \in (2, 3)$ and $\mu > \mu_{\text{pol}}$; or $\tau > 3$. Proposition 2.4 allows us to choose $\eta := \eta_0 = 1$ only in the latter two cases, otherwise we choose $\eta = \eta_0 - \delta$ for an arbitrarily small $\delta > 0$.

For each vertex $x \in \mathbb{R}^d$, define $k^*(x)$ to be the smallest integer $k \in \mathbb{N}$ with $x \in Q_k$. We will show that (2.45) holds for all x with $k^*(x) \geq k_1 + 1$, which in particular implies the inequality for all $|x| > \sqrt{d}A_{k_1}/2$. So let us fix some x and suppose that $k^* := k^*(x) \geq k_1 + 1$. In particular we have $|x| \leq \sqrt{d}A_{k^*}/2$. Let π be a path from 0 to x , and let k' be the smallest index such that $\pi \subseteq Q_{k'}$. Note that $k' \geq k^*$ because $x \notin Q_{k^*-1}$. Then there exists a vertex v on π in $Q_{k'} \setminus Q_{k'-1}$, and in particular $|v| > A_{k'-1}/2$. If $k' = k^*$ then we pick $v := x$, otherwise we pick any such vertex v . By the second part of Proposition 2.6, we have

$$\mathcal{C}(\pi) \geq u^*|v|^\eta / 30^{d+2}. \tag{2.46}$$

In the case $k' = k^*$, (2.45) follows from $x = v$ and thus $|v| = |x|$. In the case $k' > k^*$, it follows from $|v| > A_{k'-1}/2 \geq A_{k^*}/2 \geq |x|/\sqrt{d}$ and $\eta \leq 1$, finishing the proof of (1.19) and (1.20). \square

Theorem 1.12 contains Theorem 1.4 and the lower bound in Theorem 1.8 as special cases. Moreover, Corollaries 1.5 and 1.7 follow immediately from Theorem 1.4 and from Theorems 1.4 and 1.6, respectively. Finally, the lower bound in Theorem 1.10 is a trivial consequence of Theorem 1.8 since GIRG is a subgraph of IGIRG, which can only increase distances. We now prove Corollary 1.9, and then the limit cases of Theorem 1.11 in the next subsection.

Proof of Corollary 1.9. To prove this corollary we assume the upper bound of Theorem 1.8, which is proved in Section 3. In other words we assume that there exists a constant κ_2 such that $\lim_{|x| \rightarrow \infty} \mathbb{P}(\mathcal{C}(\pi_{0,x}^*) < \kappa_2|x|) = 1$. Given this constant κ_2 we define

$\kappa_3 := \kappa_2 30^{d+2}/u^*$. Let $\pi_{0,x}$ be any path that leaves the ball $B_{\kappa_3|x|}(0)$. We prove now using (2.46) that $\pi_{0,x}$ cannot be optimal. Indeed, let $v \in \pi_{0,x}$ be any vertex with $|v| \geq \kappa_3|x|$, and apply (2.46) to v , which gives with $\eta = 1$ that

$$\mathcal{C}(\pi_{0,x}) \geq u^*|v|/30^{d+2} \geq u^*\kappa_3|x|/30^{d+2} = \kappa_2|x|.$$

Hence, with probability tending to 1, any such path $\pi_{0,x}$ cannot be optimal. □

2.5 Limit cases, lower bounds

In this section, we prove the lower bounds in the limit cases, i.e., we prove the lower bounds in Theorem 1.11. We achieve this by coupling a model with $\alpha = \infty$ to a suitable model with $\alpha < \infty$, and likewise for $\beta = \infty$. For $\alpha = \infty$, we will keep the vertex set and the weights identical in the two coupled models, but we use a subset of edges of the $\alpha < \infty$ model to obtain the $\alpha = \infty$ model (with identical costs, where edges are kept). For $\beta = \infty$, we can use the same vertex set and weights, and the same edge set in the two coupled models with $\beta = \infty$ and $\beta < \infty$, but we decrease the edge costs.

Proof of Theorem 1.11 (lower bounds). (a) First consider the case $\alpha = \infty$. For Theorem 1.4, we need to study the case $\mu > \mu_{\log} = (3 - \tau)/\beta$ from (1.15) and show that for sufficiently large $|x|$,

$$d_{\mathcal{C}}(0, x) \geq |x|^{\eta_{0,\infty} - \varepsilon}, \tag{2.47}$$

where, recalling that $\mu_{\text{pol}} = (3 - \tau)/\beta + 1/d$ in the $\alpha = \infty$ case (see (1.15)),

$$\eta_{0,\infty} = \begin{cases} 1 & \text{if } \mu > \mu_{\text{pol}}, \\ d \cdot (\mu - \mu_{\log}) & \text{if } \mu \leq \mu_{\text{pol}}. \end{cases}$$

To show (2.47), we will show that IGIRG (and also GIRG and SFP) are *monotone* in α in the following sense. We first explain it for $\alpha < \infty$. Let $\alpha, \alpha' \in (1, \infty)$ with $\alpha' < \alpha$. Let $\bar{c} \geq \underline{c} > 0$, and consider a realisation of IGIRG, say G , with parameters $\alpha, \bar{c}, \underline{c} > 0$, and arbitrary other parameters. Set $\bar{c}' := \underline{c}' := \bar{c}$ and consider a second IGIRG G' with parameters $\alpha', \bar{c}', \underline{c}'$, and otherwise identical parameters to G . Then conditioned on the vertex set \mathcal{V} and the weight vector \mathcal{W} , by (1.1), the edge probabilities of G (of G') are given by a function h (a function h'), and h and h' satisfy the relation

$$h(x, w_1, w_2) \leq \bar{c} \cdot \min\{1, (w_1 w_2)/|x|^d\}^\alpha \leq \underline{c}' \cdot \min\{1, (w_1 w_2)/|x|^d\}^{\alpha'} \leq h'(x, w_1, w_2).$$

Therefore, we can couple G to G' such that G is a subgraph of G' . Note that G' has the same model parameters as G except for α' , and except for the parameters \bar{c}, \underline{c} whose values do not appear in Theorem 1.4, nor in any of the other theorems. For $\alpha = \infty$ and any $\alpha' \in (1, \infty)$, the same coupling is possible, as we show next. In this case, the defining equation (1.14) for h includes two parameters $\underline{c}, \underline{c}' > 0$, and requires that $h(x, w_1, w_2) = 0$ for $(w_1 w_2)/|x|^d < \underline{c}'$ and $h(x, w_1, w_2) \geq \underline{c}$ for $(w_1 w_2)/|x|^d \geq 1$. In all other cases, we still have $h(x, w_1, w_2) \leq 1$ because it is a probability. Hence by setting

$$h'(x, w_1, w_2) := \begin{cases} \underline{c}' \cdot \min\{1, (w_1 w_2)/|x|^d\}^{\alpha'} & \text{if } \frac{w_1 w_2}{|x|^d} < \underline{c}', \\ 1 & \text{if } \frac{w_1 w_2}{|x|^d} \geq \underline{c}', \end{cases}$$

we can ensure that $h'(x, w_1, w_2) \geq h(x, w_1, w_2)$ for all x, w_1, w_2 . Moreover, it can easily be checked that h' satisfies the conditions in (1.1) for α' with $\bar{c}' := \max\{1, \underline{c}, (\underline{c}')^{-\alpha'}\}$ and $\underline{c}' := \min\{1, \underline{c}\}$, i.e., we obtain an IGIRG model with parameter α' .

To resume the proof of Theorem 1.11, consider IGIRG G for $\alpha = \infty$, and for some β, μ, τ such that $\mu > \mu_{\log}$, and let $\varepsilon > 0$. We claim that there exists $\alpha' < \infty$ such that

$\eta'_0 := \eta_0(\alpha', \beta, \mu, \tau) > \eta_{0,\infty} - \frac{\varepsilon}{2}$. Indeed, this follows because we have defined $\eta_{0,\infty}$ such that $\lim_{\alpha' \rightarrow \infty} \eta_0(\alpha', \beta, \mu, \tau) = \eta_{0,\infty}$, as can be seen by comparing η_0 in (1.7) (see also (1.6)) to $\eta_{0,\infty}$ in (1.15). Now we couple G with an IGIRG G' with parameter α' (and different \underline{c}, \bar{c}), but identical parameters β, μ, τ . Theorem 1.4 applies to G' , so we obtain for sufficiently large $|x|$,

$$d_{\mathcal{C}'}^{G'}(0, x) \geq |x|^{\eta_{0,\infty} - \varepsilon}. \tag{2.48}$$

Since G is a subgraph of G' , cost-distances in G are larger, so (2.48) remains true if we replace $d_{\mathcal{C}'}^{G'}$ by $d_{\mathcal{C}}^G$.

Likewise, for the linear lower bound in Theorem 1.8 with $\alpha = \infty$, for any $\mu > \mu_{\text{pol}}$ (i.e., $\eta_{0,\infty} = 1$) we can find α' such that $\eta_0(\alpha', \beta, \mu, \tau) = 1$ with the same parameters β, μ, τ . Hence, the same coupling also shows that for sufficiently large $|x|$,

$$d_{\mathcal{C}}(0, x) > \kappa_1 |x|.$$

This implies the lower bound in Theorem 1.8 for $\alpha = \infty$.

Finally, the same argument also holds for GIRG, since for fixed $r > 0$, two randomly chosen vertices u_n, v_n in a box of volume n will a.a.s. satisfy $|u_n - v_n| \geq r$ as $n \rightarrow \infty$. Thus the lower bound in Theorem 1.10 is also implied. This concludes the case $\alpha = \infty$.

(b)&(c) Now we come to the case $\beta = \infty$. Note that we have shown part (a) already, so in particular the couplings that we construct now are also valid when $\alpha = \beta = \infty$. The idea is the same as for $\alpha = \infty$, but the coupling is much easier. Since the parameter β does not influence the graph structure, we can use the same graph, but with different transmission costs. More precisely, consider IGIRG G with $\beta = \infty$, i.e., the cumulative distribution F_L of the cost variables satisfies $\lim_{t \rightarrow 0} F_L(t)/t^\beta = 0$ for all $0 < \beta < \infty$. For any $\beta' \in (0, \infty)$, there exists $t_0 \in (0, 1)$ such that the distribution F'_L given by $F'_L(t) = t^{\beta'}$ for $t < t_0$ and $F'_L(t) = 1$ for $t \geq t_0$ dominates F_L . The distribution F'_L satisfies condition (1.4), so it yields an IGIRG model. Hence, we can couple G to an IGIRG G' with any parameter $\beta' \in (0, \infty)$ (and otherwise identical parameters except for the parameters t_0, c_1, c_2 that appear in (1.4)) such that G and G' are the same graph, and for every edge the transmission cost in G is at least as large as in G' .

We can now turn to proving the lower bound in Theorem 1.4 for $\beta = \infty$, so consider IGIRG G for $\beta = \infty$, and let $\mu > 0 = \mu_{\text{log}}$. According to (1.17), in this case we set $\eta_{0,\infty} := 1$ for $\mu > \mu_{\text{pol}}$ and $\eta_{0,\infty} := \min\{d\mu, \mu/\mu_{\text{pol},\alpha}\}$ otherwise. Since $\eta_{0,\infty} = \lim_{\beta' \rightarrow \infty} \eta_0(\alpha, \beta', \mu, \tau)$, we can find $\beta' \in (0, \infty)$ such that $\eta'_0 := \eta_0(\alpha, \beta', \mu, \tau) \geq \eta_{0,\infty} - \frac{\varepsilon}{2}$, and couple G to an IGIRG G' with parameter β' , as described above. Moreover, by choosing β' large enough, we can also ensure that $\mu > \mu_{\text{log}}(\beta', \tau, \mu) = (3 - \tau)/\beta'$. Then we can apply Theorem 1.4 to G' , and obtain that for sufficiently large $|x|$,

$$d_{\mathcal{C}'}^{G'}(0, x) \geq |x|^{\eta'_0 - \varepsilon/2}.$$

Since cost-distances in G' are less or equal than cost-distances in G , the statement remains true if we replace $d_{\mathcal{C}'}^{G'}$ by $d_{\mathcal{C}}^G$, and since $\eta'_0 - \frac{\varepsilon}{2} \geq \eta_{0,\infty} - \varepsilon$, we may replace $\eta'_0 - \frac{\varepsilon}{2}$ by $\eta_{0,\infty} - \varepsilon$. This yields the lower bound of Theorem 1.4 for G . As before, the proof for SFP is verbatim the same, and the proofs of the linear lower bound (i.e., the lower bound in Theorem 1.8) and the lower bound for GIRG (the lower bound of Theorem 1.10) are analogous. This concludes the case $\beta = \infty$, and concludes the proof of the lower bounds of Theorem 1.11. \square

3 Linear regime, upper bound

In the following section, we prove the upper bound of Theorem 1.8, that cost-distance in IGIRG and SFP scales at most linearly with Euclidean distance and the corresponding

part of Theorem 1.10 for finite GIRGs in a box. We shall make the renormalisation argument outlined on page 5 quantitative. We start by classical results on iid bond percolation, but first a definition:

Definition 3.1 (Deviation from straight line). *Given $u, v \in \mathbb{R}^d$, let $S_{u,v}$ denote the line segment between u, v . For $x \in \mathbb{R}^d$ we define the deviation $\text{dev}_{uv}(x) := \|x - S_{u,v}\|$. Given a path $\pi = x_1 \dots x_k$ in a graph G with vertices in \mathbb{R}^d , we define the deviation of π from S_{uv} as $\text{dev}_{uv}(\pi) := \max\{\text{dev}_{uv}(x_i) : i \in [k]\}$. Finally, let the deviation of π be $\text{dev}(\pi) := \max\{\text{dev}_{x_1 x_k}(x_i) : i \in [k]\}$, i.e., its deviation from the segment between the endpoints.*

The next two lemmas state two properties for bond percolation with high enough edge-density in \mathbb{Z}^d . First, the density of the infinite component is large in every ball $B_{(\log r)^{3/2}}(y)$, where y may vary through r^d vertices around the origin. Second, vertices in the infinite component can be joined by a path of linear length with small deviation. The first statement is implicit in [19], the second in [3]. We give proofs in the appendix for completeness on pages 40 and 41 respectively. Recall the notation $a \ll_* b, c$ and $a \gg_* b, c$ meaning that a needs to be respectively sufficiently small/large in terms of the parameters b, c from Section 1.5.

Lemma 3.2 (Locally high-density infinite cluster in bond-percolation). *Let $d \in \mathbb{N}$ with $d \geq 2$, let $\varepsilon, \delta, \sigma \in (0, 1)$, and let $r > 0$. Let ω^* be an iid Bernoulli bond percolation on \mathbb{Z}^d with edge-retention probability $p := 1 - \varepsilon$. Then whenever $r \gg_* \varepsilon, \delta$, and $\varepsilon \ll_* \sigma, d$, then almost surely ω^* has a unique infinite component C_∞^* with $\mathbb{P}(0 \in C_\infty^*) \geq 1 - \sigma$ and*

$$\forall x \in \mathbb{Z}^d : \mathbb{P}\left(\forall y \in B_r(x), \frac{|B_{(\log r)^{3/2}}(y) \cap C_\infty^*|}{|B_{(\log r)^{3/2}}(y) \cap \mathbb{Z}^d|} \geq 1 - \sigma\right) \geq 1 - \delta. \tag{3.1}$$

Lemma 3.3 (Linear distances in bond-percolation). *Let $d \in \mathbb{N}$ with $d \geq 2$, let $\zeta, \varepsilon, c \in (0, 1)$, and let $\kappa, r > 0$. Let ω^* be an iid Bernoulli bond percolation on \mathbb{Z}^d with edge-retention probability $p := 1 - \varepsilon$. Let C_∞^* be the infinite component of ω^* . For all $x \in \mathbb{Z}^d$, let $\mathcal{A}_{\text{linear}}^*(r, \kappa, \zeta, x)$ be the event that for all $y \in C_\infty^* \setminus B_r(x)$ there is a path π from x to y with length at most $|\pi| \leq \kappa|x - y|$ and deviation at most $\text{dev}(\pi) \leq \zeta|x - y|$. Then whenever $r \gg_* \zeta, \varepsilon$, and $\kappa, 1/\varepsilon, 1/c \gg_* d$,*

$$\mathbb{P}(\mathcal{A}_{\text{linear}}^*(r, \kappa, \zeta, x) \mid x \in C_\infty^*) \geq 1 - e^{-cr}. \tag{3.2}$$

3.1 General high-density random geometric graphs and renormalisation-coupling

In order to apply Lemmas 3.2 and 3.3 to IGIRG and SFP, we will need a coupling to bond percolation. We provide this in Lemma 3.5 for any graph model satisfying Definition 3.4 below. Given a graph $G = (\mathcal{V}, \mathcal{E})$ with vertex set $\mathcal{V} \subseteq \mathbb{R}^d$ and a set $A \subseteq \mathbb{R}^d$, we write $G[A]$ for the induced subgraph of G on $\mathcal{V}[A] := \mathcal{V} \cap A$. Two (half-open) boxes are called *neighbouring* if their closures have non-empty intersection.

Definition 3.4 (Dense geometric random graphs, generally). *Let $d \in \mathbb{N}$ and let S be a partition of \mathbb{R}^d into half-open boxes of the form $[a_1, a_1 + R) \times \dots \times [a_d, a_d + R)$ for some $a_1, \dots, a_d \in \mathbb{R}$ and side-length $R > 0$. When S is given, for all $z \in \mathbb{Z}^d$, we write S_z for the unique box in S containing $R \cdot z$. Let $G = (\mathcal{V}, \mathcal{E})$ be a random graph whose vertex set is either \mathbb{Z}^d or is given by a homogeneous PPP on \mathbb{R}^d . For all $\varepsilon > 0$, we say G is an (R, ε) -dense geometric graph with boxing scheme S if it satisfies the following properties (i)-(iii), and we call G a strong (R, D, ε) -dense geometric graph if it satisfies (i)-(iv):*

- (i) $G[S_1]$ and $G[S_2]$ are independent for any disjoint Lebesgue measurable sets S_1, S_2 ;
- (ii) for all boxes $S \in S$, $\mathbb{P}(G[S] \text{ is non-empty and connected}) \geq 1 - \varepsilon$;

(iii) for all neighbouring boxes $S_1, S_2 \in \mathcal{S}$,

$$\mathbb{P}(G[S_1 \cup S_2] \text{ contains at least one edge from } S_1 \text{ to } S_2) \geq 1 - \varepsilon;$$

(iv) for all boxes $S \in \mathcal{S}$, $\mathbb{P}(G[S] \text{ has diameter at most } D) \geq 1 - \varepsilon$.

Observe that $z \mapsto S_z$ is a bijection from \mathbb{Z}^d to \mathcal{S} ; this forms the basis of the coupling in Lemma 3.5 below, essentially acting as renormalisation. In IGIRG and SFP, the random graph G_M described at the start of the section will take the role of the dense geometric graph (see Corollary 3.9).

Lemma 3.5 (Renormalisation-coupling to bond-percolation). *Let $d \in \mathbb{N}$, $\varepsilon \in (0, 1)$, and $K, R > 0$. Suppose $K, 1/\varepsilon \gg_* R, d$ and let G be an (R, ε) -dense geometric graph with boxing scheme \mathcal{S} , and let ω^* be an iid Bernoulli bond percolation with retention probability $1 - 20d\varepsilon$. Then there exists a coupling between G and ω^* such that whenever $z_1 z_2$ is open in ω^* : $G[S_{z_1}]$ and $G[S_{z_2}]$ are non-empty and connected; $G[S_{z_1}]$ and $G[S_{z_2}]$ contain at most K vertices each; and there is an edge from $\mathcal{V}[S_{z_1}]$ to $\mathcal{V}[S_{z_2}]$ in G .*

Proof. In Definition 3.4, the vertex set is either a homogeneous PPP or \mathbb{Z}^d . If \mathcal{V} is a homogeneous PPP, there exists $K > 0$ such that for all boxes $S \in \mathcal{S}$,

$$\mathbb{P}(|\mathcal{V}[S]| \leq K) \geq 1 - \varepsilon. \tag{3.3}$$

If $\mathcal{V} = \mathbb{Z}^d$, then $|\mathcal{V}[S]| \leq (R + 1)^d$ and so (3.3) holds trivially with $K = (R + 1)^d$.

We now carry out a one-step renormalisation and define a site-bond percolation ω on \mathbb{Z}^d . Recall that \mathcal{S} in Definition 3.4 is a boxing scheme with side-length R and S_z is the box containing $R \cdot z$ for $z \in \mathbb{Z}^d$. In the renormalised lattice, we set a site $z \in \mathbb{Z}^d$ occupied in ω if $G[S_z]$ contains at most K vertices and is connected. We set two neighbouring sites z, z' connected by an open bond in ω if both sites are occupied and there is at least one edge in G between $G[S_z], G[S_{z'}]$. By Definition 3.4(i), sites are occupied independently of each other in ω with probability at least $1 - 2\varepsilon$ by (i), (ii) and (3.3), and bonds that do not share a site are also open independently by (i). However, bonds that do share a site (zz' and zz'') are not open independently since they are both influenced by $G[S_z]$.

We now couple ω to a Bernoulli bond percolation ω^* on \mathbb{Z}^d . Similar ideas have been used before — see [2, 37] in particular. By Definition 3.4(ii) and (iii) and by (3.3), every bond is open with probability at least $1 - 5\varepsilon$, where the factor of 5 comes from a union bound over the two sites being occupied and over (iii). We next show an approximate independence as follows. Consider a bond zz' in ω , and let $N(zz') := \{\text{bonds incident to } z \text{ or } z'\} \setminus \{zz'\}$, with size $|N(zz')| = 4d - 2$. For any set of bonds $S \subseteq N(zz')$,

$$\begin{aligned} \mathbb{P}(zz' \text{ open in } \omega \mid \text{all bonds in } S \text{ open}) &\geq \mathbb{P}(zz' \text{ open and all bonds in } S \text{ open}) \\ &\geq 1 - (4d - 1) \cdot 5\varepsilon \geq 1 - 20d\varepsilon =: p \end{aligned}$$

by a union bound. Using that bonds that do not share a site are independently open, one can iterate this argument to show that for every finite set S of bonds, $\mathbb{P}(\text{all bonds in } S \text{ are open}) \geq p^{|S|}$. Hence, by Strassen's theorem [38] we can couple ω with an independent bond percolation ω^* on \mathbb{Z}^d with retention-probability p , where every open bond in ω^* is open in ω . \square

Definition 3.6 (Bernoulli-induced infinite subgraph). *Let $\varepsilon \in (0, 1)$ be such that Lemma 3.2 applies with $\varepsilon_{3,2} := 20d\varepsilon$ for some $r > 0$ and some $\delta, \sigma \in (0, 1)$. Let G be an (R, ε) -dense geometric graph with boxing scheme \mathcal{S} and let ω^* be an iid Bernoulli bond percolation process with retention probability $p := 1 - 20d\varepsilon$ given by the coupling of Lemma 3.5.*

When $d \geq 2$, write \mathcal{C}_∞^* for the (unique) infinite component of ω^* guaranteed by Lemma 3.2. We define then the following infinite subgraph of G :

$$\mathcal{H}_\infty(G, \mathcal{S}, \omega^*) = \bigcup_{z \in \mathcal{C}_\infty^*} \mathcal{V}[S_z]. \tag{3.4}$$

When there is no danger of confusion, we write $\mathcal{H}_\infty := \mathcal{H}_\infty(G, \mathcal{S}, \omega^*)$ and we may not explicitly define the bond percolation process ω^* .

The graph \mathcal{H}_∞ corresponds to vertices in renormalised boxes that belong to \mathcal{C}_∞^* of the renormalised bond-percolation model, hence we “blow-up” \mathcal{C}_∞^* to correspond to boxes in G . By the connectivity of \mathcal{C}_∞ , also $G[\mathcal{H}_\infty]$ is a.s. connected by the properties of the coupling set out in Lemma 3.5. However, \mathcal{H}_∞ may only be a proper subset of some infinite component \mathcal{C}_∞ of G . Moreover, G need not have a unique infinite component. Indeed, suppose G is a random graph model with a PPP as \mathcal{V} , where each vertex is coloured red or blue independently at random, and there is no edge with endpoints of different colours. If the probability of ‘red’ is sufficiently high with respect to the box side length R , so that with probability at least $1 - \varepsilon$ a box $S \in \mathcal{S}$ contains no blue vertices, then Definition 3.4(i)–(ii) may be satisfied, and there will be an infinite component of red vertices. However, there may also be an infinite component of blue vertices with arbitrary behaviour. Of course, if ε is small, then \mathcal{H}_∞ has density larger than $1/2$ and so no other infinite component can have density larger than $1/2$, so \mathcal{H}_∞ is the uniquely determined as the infinite component with density close to 1.

3.2 Local denseness and linear distances in dense geometric random graphs

We now show in Lemma 3.7 that \mathcal{H}_∞ is “locally dense” in G , with density close to one, and in Lemma 3.8 that \mathcal{H}_∞ is a well-connected set containing short low-deviation paths between many pairs of vertices.

Lemma 3.7 (Locally high-density Bernoulli-induced infinite subgraph). *Let $d \in \mathbb{N}$ with $d \geq 2$, let $\varepsilon, \delta, \sigma \in (0, 1)$, and let $r, R > 0$. Suppose that $r \gg_* \varepsilon, \delta, R$, and that $\varepsilon \ll_* \sigma, d$. Let G be an (R, ε) -dense geometric graph with boxing scheme \mathcal{S} . Then a.s. \mathcal{H}_∞ is infinite and $G[\mathcal{H}_\infty]$ defined in (3.4) is connected, for any box $S \in \mathcal{S}$, $\mathbb{P}(\mathcal{V} \cap S \subseteq \mathcal{H}_\infty) \geq 1 - \sigma$ and moreover*

$$\forall x \in \mathbb{R}^d : \mathbb{P} \left(\forall y \in B_r(x), \frac{|B_{(\log r)^2}(y) \cap \mathcal{H}_\infty|}{|B_{(\log r)^2}(y) \cap \mathcal{V}|} \geq 1 - \sigma \right) \geq 1 - \delta. \tag{3.5}$$

The statement in (3.5) gets stronger if one decreases the radius $(\log r)^2$, since then the infinite subgraph \mathcal{H}_∞ is present already at lower radii near every vertex y inside $B_r(x)$. For our proofs later, $(\log r)^2$ is more than strong enough, but one could improve (3.5) to having $\Theta(\log r)$ as radius around y . If the vertex set of G is \mathbb{Z}^d , then Lemma 3.7 follows from Lemma 3.2 (applied with some $\sigma_{3.2}$ that is sufficiently small compared to σ in (3.5)) by the coupling in Lemma 3.5. If the vertex set comes from a Poisson point process, we use that the number of vertices in the boxes $B_{(\log r)^2}(y) \cap \mathcal{H}_\infty$ concentrates around its mean and that only at most $\sigma_{3.2}$ -fraction of boxes have too few/too many vertices in $B_{(\log r)^2}(y) \setminus \mathcal{H}_\infty$. The details are straightforward and we give a proof of Lemma 3.7 in the appendix on page 43.

Lemma 3.8 (Linear graph distances in dense geometric random graphs). *Let $d \in \mathbb{N}$ with $d \geq 2$, let $\zeta, \varepsilon, \delta \in (0, 1)$, and let $\kappa, r, R, C > 0$. Let G be an (R, ε) -dense geometric graph with boxing scheme \mathcal{S} . For $x \in \mathcal{V}$, let $\mathcal{A}_{\text{linear}}(r, \kappa, C, \zeta, x)$ be the event that for all $u \in B_r(x) \cap \mathcal{H}_\infty$ and $v \in \mathcal{H}_\infty$, there is a path from u to v in G of length at most $\kappa|u - v| + C$ and deviation at most $\zeta|u - v| + C$. Then whenever $C \gg_* r \gg_* \varepsilon, \delta, \zeta, R, d$ and $\kappa \gg_* R, d$ and $\varepsilon \ll_* d$, then $\mathbb{P}(\mathcal{A}_{\text{linear}}(r, \kappa, C, \zeta, x)) \geq 1 - \delta$.*

Proof. Let ω^* be a bond percolation process with retention probability $1 - 20d\varepsilon$ as required in Lemma 3.5, so that $\mathcal{H}_\infty = \mathcal{H}_\infty(G, \mathcal{S}, \omega^*)$ in (3.4) exists. Indeed, since $\varepsilon \ll_* d$, by Lemma 3.2 ω^* has a unique infinite component \mathcal{C}_∞^* a.s. and \mathcal{H}_∞ is well-defined. Let $K > 0$ be as in Lemma 3.5, so that K is a function of d and R , and let $\kappa' = \sqrt{\kappa}$ and $C' = \sqrt{C}$.

Denote the event $\mathcal{A}_{\text{linear}}^*(r/R, \kappa', \zeta, z)$ of Lemma 3.3 by $\mathcal{A}_1^*(z)$, so if $\mathcal{A}_1^*(z)$ occurs then there is a short low-deviation path from z to any site in $\mathcal{C}_\infty^* \setminus B_{r/R}(z)$. Let $\mathcal{A}_2^*(z)$ be the event that for all $v \in \mathcal{C}_\infty^* \cap B_{r/R}(z)$, there is a path in ω^* from z to v of length and deviation at most C' . Observe that if $\mathcal{A}_1^*(z) \cap \mathcal{A}_2^*(z)$ occurs, then for all $v \in \mathcal{C}_\infty^*$ there is a path in ω^* from z to v of length at most $C' + \kappa'|z - v|$ and deviation at most $C' + \zeta|z - v|$. Recall that for $z \in \mathbb{Z}^d$, we denote by S_z the box in \mathcal{S} that contains $R \cdot z$. We define our last event as

$$\mathcal{A}^+(x) := \bigcap_{z \in \mathbb{Z}^d : S_z \cap B_r(x) \neq \emptyset} \left((\mathcal{A}_1^*(z) \cap \mathcal{A}_2^*(z)) \cup \{z \notin \mathcal{C}_\infty^*\} \right).$$

Using the coupling of Lemma 3.5 we will prove that

$$\mathcal{A}^+(x) \subseteq \mathcal{A}_{\text{linear}}(r, \kappa, C, \zeta, x), \tag{3.6}$$

$$\sum_{z \in \mathbb{Z}^d : S_z \cap B_r(x) \neq \emptyset} \mathbb{P}(\neg \mathcal{A}_1^*(z) \cap \{z \in \mathcal{C}_\infty^*\}) \leq \delta/2, \tag{3.7}$$

$$\mathbb{P}\left(\bigcap_{z \in \mathbb{Z}^d : S_z \cap B_r(x) \neq \emptyset} (\mathcal{A}_2^*(z) \cup \{z \notin \mathcal{C}_\infty^*\}) \right) \geq 1 - \delta/2. \tag{3.8}$$

Given (3.6)–(3.8), a union bound gives that $\mathbb{P}(\mathcal{A}_{\text{linear}}(r, \kappa, C, \zeta, x)) \geq \mathbb{P}(\mathcal{A}^+) \geq 1 - \delta$, as required. We first prove (3.7), starting with

$$\sum_{z \in \mathbb{Z}^d : S_z \cap B_r(x) \neq \emptyset} \mathbb{P}(\neg \mathcal{A}_1^*(z) \cap \{z \in \mathcal{C}_\infty^*\}) \leq \sum_{z \in \mathbb{Z}^d : S_z \cap B_r(x) \neq \emptyset} \mathbb{P}(\neg \mathcal{A}_1^*(z) \mid z \in \mathcal{C}_\infty^*).$$

By Lemma 3.3, since $\kappa' \gg_* d \gg_* \varepsilon$ and $r \gg_* \varepsilon, \zeta, R, d$, there exists $c_{3.3}$ depending only on d such that each term in the sum is at most $e^{-c_{3.3}r/R}$. Since S_z is a box of side-length R , the number of terms $|\{z \in \mathbb{Z}^d : S_z \cap B_r(x)\}|$ is at most $(3r/R)^d$, and $r \gg_* \delta, R, d$ (and thus $r \gg_* c_{3.3}$),

$$\sum_{z \in \mathbb{Z}^d : S_z \cap B_r(x) \neq \emptyset} \mathbb{P}(\neg \mathcal{A}_1^*(z) \cap \{z \in \mathcal{C}_\infty^*\}) \leq (3r/R)^d e^{-c_{3.3}r/R} \leq \delta/2$$

for all sufficiently large r given δ . We next prove (3.8). The maximum length and deviation of a path between any pair of sites in $\mathcal{C}_\infty^* \cap B_{3r/R}(x)$ is a random variable which is a.s. finite, so since $C' = \sqrt{C} \gg_* \delta, r$ this maximum must be at most C' with probability at least $1 - \delta/2$, for all sufficiently large C' . Thus (3.8) follows.

Finally, we prove (3.6). Suppose $\mathcal{A}^+(x)$ occurs, let $u \in B_r(x) \cap \mathcal{H}_\infty$, and let $v \in \mathcal{H}_\infty$ two vertices in G . Let $u^- \in \mathbb{Z}^d$ be such that $u \in S_{u^-}$, and let $v^- \in \mathbb{Z}^d$ be such that $v \in S_{v^-}$; thus $u^- = \lfloor u/R \rfloor$ and $v^- = \lfloor v/R \rfloor$, and $u^-, v^- \in \mathcal{C}_\infty^*$ per definition of \mathcal{H}_∞ in (3.4). Since $\mathcal{A}^+(x)$ occurs and $u^- \in \mathcal{C}_\infty^*$, the event $\mathcal{A}_1^*(u^-) \cap \mathcal{A}_2^*(u^-)$ also occurs; thus there exists a path $\pi^* = z_1 \dots z_m$ (for some m) from u^- to v^- in ω^* with

$$|\pi^*| \leq C' + \kappa'|u^- - v^-|, \quad \text{dev}(\pi^*) \leq C' + \zeta|u^- - v^-|.$$

The statement of the lemma becomes weaker with larger κ , so we may just prove it for one fixed value $\kappa = \kappa(R, d)$, and deduce it for all larger values of κ . We may thus assume for this fixed value of κ that $C' \gg_* \kappa$ and $C' \gg_* \kappa' = \sqrt{\kappa}$. For the same reason we may

assume $C' \gg_* \zeta$. Since also $C' \gg_* R, d$ and the diameter of each box is at most $R\sqrt{d}$, we have

$$|\pi^*| \leq 2C' + \kappa'|u - v|/R, \quad \text{dev}(\pi^*) \leq 2C' + \zeta|u - v|/R. \quad (3.9)$$

By the properties of the coupling with ω^* set out in Lemma 3.5, for each $z_i \in \pi^*$, the graph $G[B_{z_i}]$ is connected and contains between 1 and K vertices, and for all $i \leq m - 1$ there is an edge between $G[B_{z_i}]$ and $G[B_{z_{i+1}}]$. Since $u \in B_{z_1}$ and $v \in B_{z_m}$, we can thus find a path π^G from u to v in $G[B_{z_1} \cup \dots \cup B_{z_m}]$. Since each box contains at most K vertices, it follows from (3.9) that

$$|\pi^G| \leq Km = K(|\pi^*| + 1) \leq 2K\kappa'|u - v|/R + K(2C' + 1).$$

Likewise, since each box has side length R and diameter at most $R\sqrt{d}$,

$$\text{dev}(\pi^G) \leq R\sqrt{d} + R \cdot \text{dev}(\pi^*) \leq \zeta|u - v| + 2RC' + R\sqrt{d}.$$

Recall that K is a function of d and R , and that $\kappa, C \gg_* d, R$; we therefore recover $|\pi^G| \leq \kappa|u - v| + C$ and $\text{dev}(\pi^G) \leq \zeta|u - v| + C$. Thus $\mathcal{A}_{\text{linear}}(r, \kappa, C, \zeta, x)$ occurs, and so $\mathcal{A}^+(x) \subseteq \mathcal{A}_{\text{linear}}(r, \kappa, C, \zeta, x)$ as claimed. \square

We will state the following corollary in a stronger form than necessary here; part (iii) and the ‘strong’ version of part (i) below is only used in our companion paper [33] but not in this paper. Recall from (1.5) that par is the set of model parameters. In order to be able to talk about a number of vertices being in the vertex set, we shall work conditional on one of two events; these conditionings might be trivial for SFP. In particular, we write $\mathcal{F} := \{x_1, \dots, x_t \in \mathcal{V}\}$ to mean that the locations $x_1, \dots, x_t \in \mathbb{R}^d$ are part of the vertex set with unknown vertex weights from distribution F_W , and $\mathcal{F}_M := \{x_1, \dots, x_t \in \mathcal{V}_M\}$ for the event that the locations x_1, \dots, x_t are part of the vertex set with unknown vertex weights in the interval $[M, 2M]$ for some large constant M . In other words, we (sometimes) work under certain Palm-measures of the underlying Poisson process.

Corollary 3.9 (Linear costs in Bernoulli-induced infinite subgraph in IGIRG and SFP). *Consider 1-FPP in Definition 1.2 on a graph $G = (\mathcal{V}, \mathcal{E})$ that is an IGIRG or SFP satisfying assumptions (1.2)–(1.4) with $\tau \in (2, 3)$. Let $\delta, \varepsilon, \sigma \in (0, 1)$, let $r, M, C, \kappa, \zeta, t > 0$, and let $R := M^{2/d}/\sqrt{d}$. Suppose $C \gg_* r \gg_* M, \zeta, \delta$, and that $\kappa \gg_* M \gg_* \varepsilon, \sigma, t, \text{par}$. Finally, suppose that $\varepsilon \ll_* \sigma, \text{par}$. Let $I_M := [M, 2M]$, and let $G_M = (\mathcal{V}_M, \mathcal{E}_M)$ be the subgraph of G formed by vertices $\mathcal{V}_M := \{v \in \mathcal{V} : W_v \in I_M\}$ and edges $\mathcal{E}_M := \{uv \in \mathcal{E} : u, v \in \mathcal{V}_M, C(uv) \leq M^{3\mu}\}$. For t as above, let $\{x_1, \dots, x_t\} \subseteq \mathbb{Z}^d$ for SFP and $\{x_1, \dots, x_t\} \subseteq \mathbb{R}^d$ for IGIRG, and let $\mathcal{F} := \{x_1, \dots, x_t \in \mathcal{V}\}$ and $\mathcal{F}_M := \{x_1, \dots, x_t \in \mathcal{V}_M\}$. Then:*

- (i) *For all dimensions $d \geq 1$, conditioned on either \mathcal{F} or \mathcal{F}_M , the graph G_M is a strong $(R, 2, e^{-M^{3-\tau-\varepsilon}})$ -dense and an (R, ε) -dense geometric graph in the sense of Definition 3.4.*
- (ii) *For all dimensions $d \geq 2$, a.s. \mathcal{H}_∞ is infinite, $G[\mathcal{H}_\infty]$ is connected, G has a unique infinite component \mathcal{C}_∞ , and $\mathcal{H}_\infty \subseteq \mathcal{V}(\mathcal{C}_\infty) \cap \mathcal{V}_M$.*
- (iii) *For all dimensions $d \geq 2$, and $x \in \mathbb{R}^d$,*

$$\mathbb{P}\left(\forall y \in B_r(x), \frac{|B_{(\log r)^2}(y) \cap \mathcal{H}_\infty|}{|B_{(\log r)^2}(y) \cap \mathcal{V}_M|} \geq 1 - \sigma \mid \mathcal{F}\right) \geq 1 - \delta. \quad (3.10)$$

- (iv) *For $x \in \mathbb{R}^d$, let $\mathcal{A}_{\text{linearcost}}(r, \kappa, C, \zeta, x)$ be the event that for all $u \in B_r(x) \cap \mathcal{H}_\infty$ and all $v \in \mathcal{H}_\infty$, there is a path from u to v of cost at most $\kappa|u - v| + C$ and deviation at most $\zeta|u - v| + C$. Then for all dimensions $d \geq 2$, $\mathbb{P}(\mathcal{A}_{\text{linearcost}}(r, \kappa, C, \zeta, z) \mid \mathcal{F}) \geq 1 - \delta$.*

Proof. Let \mathcal{S} be a boxing scheme of side length $R = M^{2/d}/\sqrt{d}$. Using Definition 3.4, we now prove that G_M is a strong $(R, 2, \varepsilon_{3.4})$ -dense geometric graph for $\varepsilon_{3.4} := e^{-M^{3-\tau-\varepsilon}}$ with boxing scheme \mathcal{S} . Note that this automatically implies that (R, ε) -dense since $M \gg_* \varepsilon$. Definition 3.4(i) follows from the definitions of the models. For the other conditions we first lower-bound the expected number of vertices in a box $S \in \mathcal{S}$. In both IGIRG and SFP, we use (1.2) to bound:

$$\begin{aligned} \mathbb{P}(W \in I_M) &= (1 - F_W(M)) - (1 - F_W(2M)) = \frac{\ell(M)}{M^{\tau-1}} - \frac{\ell(2M)}{(2M)^{\tau-1}} \\ &= M^{-(\tau-1)}\ell(M) \left(1 - 2^{-(\tau-1)} \frac{\ell(2M)}{\ell(M)} \right) \geq M^{-(\tau-1)-\varepsilon/4}, \end{aligned} \tag{3.11}$$

where to obtain the last inequality we used that ℓ is slowly-varying, so $\ell(2M)/\ell(M) \rightarrow 1$, and that $M \gg_* \varepsilon$. In IGIRG, \mathcal{V}_M follows a homogeneous PPP with intensity $\mathbb{P}(W \in I_M)$ on \mathbb{R}^d ; since S has side length R , $|\mathcal{V}_M[S]|$ is therefore a Poisson random variable with mean $\lambda_M(\mathbb{R}^d) := R^d \mathbb{P}(W \in I_M)$, which is at least $R^d M^{-(\tau-1)-\varepsilon/4} = d^{-d/2} M^{3-\tau-\varepsilon/4}$. Since $\tau < 3$, the exponent is positive since we assumed $\varepsilon \ll_* \text{par}$ small, and so $\lambda_M \geq M^{3-\tau-\varepsilon/2}$ since $M \gg_* d, \varepsilon$. Similarly, in SFP the number of vertices in $\mathcal{V}_M[S]$ follows a binomial distribution with mean $\lambda_M(\mathbb{Z}^d) := |\mathbb{Z}^d \cap S| \cdot \mathbb{P}(W \in I_M) \geq d^{-d/2} M^{3-\tau-\varepsilon/4}/2$, (where the factor of 2 suffices to account for boundary effects since M is large), which is again at least $M^{3-\tau-\varepsilon/2}$ for $M \gg_* \varepsilon$.

We now study the edge probabilities within the box S . Consider two vertices $u, v \in \mathcal{V}_M[S]$. Their distance is most the diameter of the box, $\sqrt{d}R \leq M^{2/d}$, and $W_u, W_v \in I_M = [M, 2M]$ so $W_u W_v / |u - v|^d \geq 1$ holds. Recall that the edges have cost $(W_u W_v)^\mu L_{uv} \leq 4M^{2\mu} L_{uv}$ in 1-FPP, while we keep the edge in G_M only if this edge cost is at most $M^{3\mu}$. Thus for all vertices $u, v \in \mathcal{V}_M[S]$,

$$\begin{aligned} \mathbb{P}(uv \in \mathcal{E}, \mathcal{C}(uv) \leq M^{3\mu} \mid u, v \in \mathcal{V}_M[S]) &\geq \underline{c} \left(1 \wedge \frac{W_u W_v}{|u - v|^d} \right)^\alpha \cdot F_L((W_u W_v)^{-\mu} M^{3\mu}) \\ &\geq \underline{c} F_L(4^{-\mu} M^\mu) \geq \underline{c}/2, \end{aligned} \tag{3.12}$$

where the last inequality holds since $M \gg_* \text{par}$. Note that (3.12) holds *uniformly* over the weights in I_M and locations of vertices in S , and is also valid when $\alpha = \infty$ or $\beta = \infty$. With the vertex set \mathcal{V}_M exposed, the presence of edges in $G_M[S]$ can therefore stochastically dominates an independent collection of Bernoulli($\underline{c}/2$) random variables.

Thus, the graph $G_M[S]$ dominates an Erdős-Rényi random graph with number of vertices distributed as Poisson($\lambda_M(\mathbb{R}^d)$) or binomial with mean $\lambda_M(\mathbb{Z}^d)$ with $\lambda_M(\mathbb{R}^d), \lambda_M(\mathbb{Z}^d)$ both at least $M^{3-\tau-\varepsilon/2}$, and *constant* connection probability $\underline{c}/2$. This Erdős-Rényi random graph is non-empty and connected with diameter two with probability at least $1 - \exp(-\Theta(\lambda_M)) \geq 1 - \exp(-M^{3-\tau-\varepsilon}) = 1 - \varepsilon_{3.4}$, see [24, Theorem 7.1]. Hence, G_M satisfies conditions (ii)–(iv) of Definition 3.4 with $\varepsilon_{3.4}$ and is thus a strong $(R, 2, \varepsilon_{3.4})$ -dense geometric graph, which proves that (i) of Corollary 3.9 holds unconditionally. Moreover, since $R \gg_* t$, the above argument still holds conditioned on any intersection \mathcal{F} of at most t events of the form $z \in \mathcal{V}$ or $z \in \mathcal{V}_M$; from this point on in the proof, we always condition on \mathcal{F} , and the only property of G_M we use is that it is an (R, ε) -dense geometric graph conditioned on \mathcal{F} .

When $d \geq 2$, let ω^* be an iid Bernoulli bond percolation with retention probability $1 - 20d\varepsilon$, and recall $\mathcal{H}_\infty := \mathcal{H}_\infty(G_M, \mathcal{S}, \omega^*)$ from (3.4) in Definition 3.6. Since $\varepsilon \ll_* \sigma, d$ and $r \gg_* \varepsilon, \delta, R$, Corollary 3.9(ii)–(iii) are now immediate from Lemma 3.7, except the uniqueness of \mathcal{C}_∞ , which was proved earlier in [16] for SFP and in [18] for IGIRG. Recall the definition of $\mathcal{A}_{\text{linear}}(r, \kappa, C, \zeta, x)$ from Lemma 3.8: in particular, that the *graph* distance in G_M between two vertices is κ times the Euclidean distance. Since every

edge of G_M has cost at most $M^{3\mu}$, the cost distance then is at most $\kappa M^{3\mu}$ between those vertices in G_M . Hence, $\mathcal{A}_{\text{linearcost}}(r, \kappa, C, \zeta, z) \supseteq \mathcal{A}_{\text{linear}}(r, \kappa/M^{3\mu}, C, \zeta, z)$; thus Corollary 3.9(iv) follows from Lemma 3.8. \square

The above proof works whenever $3 - \tau > 0$, which guarantees that the expected degree of vertices in G_M , i.e., $M^{3-\tau+o(1)}$, grows with M . We chose the box-size $R = M^{2/d}$ so that it corresponds to the connectivity radius of a vertex in \mathcal{V}_M . For $\tau > 3$, the expected degree in G_M tends to 0 as M increases and so G_M has no infinite component for large M (even if we removed the restriction on edge costs).

3.3 Connecting to the infinite subgraph of the auxiliary graph

Corollary 3.9(iv) guarantees linear-cost low-deviation paths within $\mathcal{H}_\infty \subseteq \mathcal{C}_\infty$ when $d \geq 2$. However, when connecting two arbitrary vertices, e.g. 0 and x , in IGIRG/SFP, we cannot assume that they fall in \mathcal{H}_∞ . We solve this issue using the following two claims.

Claim 3.10 (The infinite cluster of IGIRG/SFP is dense). *Let $G = (\mathcal{V}, \mathcal{E})$ be an IGIRG or SFP satisfying assumptions (1.2)–(1.4) with $d \geq 2$ and $\tau \in (2, 3)$. Then there exists $\rho > 0$ such that for all $u, v \in \mathbb{R}^d$ (for IGIRG) or $u, v \in \mathbb{Z}^d$ (for SFP), we have $\mathbb{P}(u, v \in \mathcal{C}_\infty \mid u, v \in \mathcal{V}) \geq \rho$.*

Claim 3.11 (Connecting to the Bernoulli-induced infinite subgraph). *Consider the setting of Corollary 3.9 with $d \geq 2$, recall $\mathcal{H}_\infty \subseteq \mathcal{V}_M$ from (3.4), and let \mathcal{C}_∞ be the infinite component of G containing \mathcal{H}_∞ . Let $u, v \in \mathbb{R}^d$ (for IGIRG) or \mathbb{Z}^d (for SFP). Then, whenever $r \gg_* M, \delta$,*

$$\mathbb{P}\left(\exists u^* \in \mathcal{H}_\infty \cap B_r(u); \exists \text{ a path } \pi_{u, u^*}^G \subseteq \mathcal{E}(G) : \mathcal{V}(\pi_{u, u^*}^G) \subseteq B_r(u), \mathcal{C}(\pi_{u, u^*}^G) \leq C \mid u, v \in \mathcal{C}_\infty\right) \geq 1 - \delta. \tag{3.13}$$

Proof of Claim 3.10. Fix two locations $u, v \in \mathbb{R}^d$ or $u, v \in \mathbb{Z}^d$, and let $M \gg_* \text{par}$, and recall from the calculation in (3.11) that

$$\mathbb{P}(u, v \in \mathcal{V}_M \mid u, v \in \mathcal{V}) = \mathbb{P}(W \in I_M)^2 \geq M^{-2(\tau-1)-\varepsilon/2} \geq M^{-2\tau}.$$

Further, by Corollary 3.9 (i), G_M is an (R, ε) -dense geometric graph even when conditioned on the event $\mathcal{F}_M = \{u, v \in \mathcal{V}_M\}$. This means that the coupling in Lemma 3.5 remains valid under the conditioning, and the high-density result $\mathbb{P}(S \cap \mathcal{V} \subseteq \mathcal{H}_\infty \mid u, v \in \mathcal{V}_M) \geq 1 - \sigma$ in Lemma 3.7 holds also conditioned on $u, v \in \mathcal{V}_M$. So, a union bound yields that

$$\begin{aligned} \mathbb{P}(u, v \in \mathcal{H}_\infty \mid u, v \in \mathcal{V}_M) &\geq 1 - \mathbb{P}(u \notin \mathcal{H}_\infty \mid u, v \in \mathcal{V}_M) - \mathbb{P}(v \notin \mathcal{H}_\infty \mid u, v \in \mathcal{V}_M) \\ &\geq 1 - 2\sigma \geq 1/2, \end{aligned}$$

whenever $\sigma \leq 1/4$. Hence

$$\mathbb{P}(u, v \in \mathcal{H}_\infty \mid u, v \in \mathcal{V}) = \mathbb{P}(u, v \in \mathcal{H}_\infty \mid u, v \in \mathcal{V}_M) \cdot \mathbb{P}(u, v \in \mathcal{V}_M \mid u, v \in \mathcal{V}) \geq 1/(2M^{2\tau}).$$

Since $\mathcal{H}_\infty \subseteq \mathcal{V}(\mathcal{C}_\infty)$ by Corollary 3.9(ii), the result follows with $\rho = 1/(2M^{2\tau})$. \square

Proof of Claim 3.11. For $r' > 0$ let $\mathcal{B}^*(r', M, u) := \mathcal{H}_\infty \cap B_{r'}(u)$. By Corollary 3.9(ii)–(iii), applied with $\delta/2$ instead of δ , \mathcal{H}_∞ has a positive density in \mathcal{V}_M around u . Hence, for $r' \gg_* M, \delta$,

$$\mathbb{P}(\mathcal{B}^*(r', M, u) = \emptyset \mid u, v \in \mathcal{C}_\infty) \leq \delta/2. \tag{3.14}$$

Fix such r' . Because $\mathcal{H}_\infty \subseteq \mathcal{C}_\infty$ connected, paths exist within \mathcal{C}_∞ to vertices in \mathcal{H}_∞ . In particular, conditioned on $u, v \in \mathcal{C}_\infty$, fix any procedure to uniquely select a path $\pi_{uu^*}^G$ from u to each $u^* \in \mathcal{B}^*(r', M, u)$ (for example, take the paths with lowest deviation, breaking ties by random coin-flips). Then the random variables

$$X := \inf\{r > 0 \mid \forall u^* \in \mathcal{B}^*(r', M, u) : V(\pi_{uu^*}^G) \subseteq B_r(u)\},$$

$$Y := \inf\{C > 0 \mid \forall u^* \in \mathcal{B}^*(r', M, u) : \mathcal{C}(\pi_{uu^*}^G) \leq C\},$$

are a.s. finite conditioned on $u, v \in \mathcal{C}_\infty$. On choosing r and C suitably large, we have

$$\mathbb{P}(X > r \text{ or } Y > C \mid u, v \in \mathcal{C}_\infty) \leq \delta/2, \tag{3.15}$$

and the result follows from a union bound over (3.14) and (3.15). \square

The next corollary combines the previous claims and lemmas and constructs a linear-cost low deviation path between two vertices in the unique infinite component of SFP/IGIRG.

Corollary 3.12 (Linear costs in 1-FPP on IGIRG/SFP). *Consider 1-FPP on IGIRG or SFP of Definition 1.1 satisfying the assumptions given in (1.2)–(1.4) with $0 \in \mathcal{V}$. Assume that the dimension $d \geq 2$, and $\alpha > 2, \tau \in (2, 3), \mu > \mu_{\text{pol}}$. Let $\kappa, r > 0$, and let $\delta, \zeta \in (0, 1)$. Suppose $r \gg_* \delta, \zeta, \text{par}$, and that $\kappa \gg_* \text{par}$. Let \mathcal{C}_∞ be the unique infinite component guaranteed by Corollary 3.9(ii). Let $u, v \in \mathbb{R}^d$ for IGIRG or \mathbb{Z}^d for SFP with $|u - v| \geq r$. Let $\mathcal{A}_{\text{linearcost}}(u, v, \kappa, \zeta)$ be the event that there is a path joining u and v in G with cost at most $\kappa|u - v|$ and deviation at most $\zeta|u - v|$. Then*

$$\mathbb{P}(\mathcal{A}_{\text{linearcost}}(u, v, \kappa, \zeta) \mid u, v \in \mathcal{C}_\infty) \geq 1 - \delta.$$

Proof. Let $\rho \in (0, 1)$ be as in Claim 3.10. Let $r_{3.9}, C_{3.9} > 0$ be as in Corollary 3.9(iv) applied with $\kappa_{3.9} = \kappa/2, \delta_{3.9} = \rho\delta/3, \zeta_{3.9} = \zeta/2$ and $t_{3.9} = 2$ (and any suitable values of ε, M, σ). Let $r_{3.11}$ and $C_{3.11}$ be as in Claim 3.11 applied with $\delta_{3.11} = \delta/3$. In Corollary 3.9(iv), the requirement $r_{3.9} \gg_* M, \zeta, \delta$ is assumed and in Claim 3.11 $r_{3.11} \gg_* M$, thus we may increase the first value to assume $r_{3.9} \geq r_{3.11}$. We may assume that $r \gg_* r_{3.9}, r_{3.11}, C_{3.9}, C_{3.11}$. We define events as follows.

- (a) Let \mathcal{A}_1 be the event that there is a path π^1 from u to some vertex $u^* \in \mathcal{H}_\infty \cap B_{r_{3.11}}(u)$ of cost at most $C_{3.11}$ and with $\mathcal{V}(\pi^1) \subseteq B_{r_{3.11}}(u)$.
- (b) Let \mathcal{A}_2 be the event that every vertex $u_1 \in \mathcal{H}_\infty \cap B_{r_{3.9}}(u)$ is joined to every vertex $u_2 \in \mathcal{H}_\infty$ by a path π_{u_1, u_2}^2 of cost at most $(\kappa/2)|u_1 - u_2| + C_{3.9}$ and deviation at most $(\zeta/2)|u_1 - u_2| + C_{3.9}$.
- (c) Let \mathcal{A}_3 be the event that there is a path π^3 from some vertex $v^* \in \mathcal{H}_\infty \cap B_{r_{3.11}}(v)$ to v of cost at most $C_{3.11}$ and with $\mathcal{V}(\pi^3) \subseteq B_{r_{3.11}}(v)$.

Observe that if $\mathcal{A}_1 \cap \mathcal{A}_2 \cap \mathcal{A}_3$ occurs, then since $|u - v| \geq r$ is large, the path $\pi^1 \pi_{u^*, v^*}^2 \pi^3$ has cost at most $\kappa|u - v|/2 + C_{3.9} + 2C_{3.11} \leq \kappa|u - v|$ and deviation at most $\zeta|u - v|/2 + C_{3.9} + 2r_{3.11} \leq \zeta|u - v|$. We must therefore prove

$$\mathbb{P}(\mathcal{A}_1 \cap \mathcal{A}_2 \cap \mathcal{A}_3 \mid u, v \in \mathcal{C}_\infty) \geq 1 - \delta. \tag{3.16}$$

By our choice of $r_{3.11}$ and $C_{3.11}$, Claim 3.11 implies that

$$\mathbb{P}(\neg \mathcal{A}_1 \mid u, v \in \mathcal{C}_\infty) \leq \delta/3, \quad \mathbb{P}(\neg \mathcal{A}_3 \mid u, v \in \mathcal{C}_\infty) \leq \delta/3.$$

Similarly, by Corollary 3.9(iv) we have $\mathbb{P}(\neg \mathcal{A}_2 \mid u, v \in \mathcal{V}) \leq \rho\delta/3$. Thus by Lemma 3.10,

$$\mathbb{P}(\neg \mathcal{A}_2 \mid u, v \in \mathcal{C}_\infty) \leq \frac{\mathbb{P}(\neg \mathcal{A}_2 \mid u, v \in \mathcal{V})}{\mathbb{P}(u, v \in \mathcal{C}_\infty \mid u, v \in \mathcal{V})} \leq \frac{\rho\delta/3}{\rho} = \delta/3.$$

Applying a union bound, (3.16) follows as required. \square

Proof of Theorems 1.8 and 1.11. We have already proven the lower bounds: namely the lower bound of Theorem 1.8 follows directly from (1.20) in Theorem 1.12, with proof on page 28, and the proof of the lower bound in Theorem 1.11 can be found on page 29. The upper bounds of both theorems follow from Corollary 3.12. \square

It remains to prove Theorem 1.10, which states that all results also hold in the finite GIRG model if u_n, v_n are chosen uniformly at random from the largest component $\mathcal{C}_{\max}^{(n)}$.

Proof of Theorem 1.10. Let $\kappa_1 > 0$ be small enough to take the role of κ in (1.20) in Theorem 1.12, and let $\kappa_2 > 0$ be large enough to take the role of κ in Corollary 3.12. Let $G_n = (\mathcal{V}_n, \mathcal{E}_n)$ be a GIRG satisfying the assumptions of the theorem statement, let $\mathcal{C}_{\max}^{(n)}$ be the largest component of G_n , and let u_n and v_n be vertices chosen independently and uniformly at random from $\mathcal{C}_{\max}^{(n)}$. For all $u, v \in \mathbb{R}^d$ and all graphs H , let $\mathcal{A}(H, u, v)$ be the event that u and v have cost-distance between $\kappa_1|u - v|$ and $\kappa_2|u - v|$ in the (sub)graph H . Then it suffices to prove that for all $\delta \in (0, 1)$, whenever $n \gg_* \delta$, $\mathbb{P}(\neg \mathcal{A}(G_n, u_n, v_n)) \leq \delta$ holds.

Let $x_n, y_n \in \mathcal{V}_n$ be chosen independently and uniformly at random from \mathcal{V}_n . Whenever $n \gg_* \delta$; it is known [32, Theorem 3.11] that $\mathcal{C}_{\max}^{(n)}$ has constant density whp, so with probability at least $1/2$, we have $|\mathcal{C}_{\max}^{(n)}| \geq \delta^{1/4}|\mathcal{V}_n|$. Thus $\mathbb{P}(x_n, y_n \in \mathcal{C}_{\max}^{(n)}) \geq \sqrt{\delta}/2$, and so

$$\begin{aligned} \mathbb{P}(\neg \mathcal{A}(G_n, u_n, v_n)) &= \mathbb{P}(\neg \mathcal{A}(G_n, x_n, y_n) \mid x_n, y_n \in \mathcal{C}_{\max}^{(n)}) \\ &\leq 2\mathbb{P}(\neg \mathcal{A}(G_n, x_n, y_n) \text{ and } x_n, y_n \in \mathcal{C}_{\max}^{(n)})/\sqrt{\delta}. \end{aligned}$$

Recall that for all $n > 0$, $Q_n := [-n^{1/d}/2, n^{1/d}/2]^d$. Let $x'_n, y'_n \in Q_n$ be random points chosen independently and uniformly at random from the Lebesgue measure in Q_n . Let \mathcal{V}'_n be a Poisson point process of unit intensity conditioned on $x'_n, y'_n \in \mathcal{V}'_n$, i.e. a Palm process. It is known that the total variation distance of $(\mathcal{V}', x'_n, y'_n)$ from (\mathcal{V}, x_n, y_n) converges to zero as $n \rightarrow \infty$, and in particular is at most $\delta^{3/2}/12$ when n is sufficiently large. Thus on taking G'_n to be a GIRG with vertex set \mathcal{V}'_n , we have

$$\mathbb{P}(\neg \mathcal{A}(G_n, u_n, v_n)) \leq (2/\sqrt{\delta})(\delta^{3/2}/12) + 2\mathbb{P}(\neg \mathcal{A}(G'_n, x'_n, y'_n) \text{ and } x'_n, y'_n \in \mathcal{C}_{\max}^{(n)})/\sqrt{\delta}.$$

We may couple G'_n to an IGIRG G^+ in such a way that $G_n = G^+[Q_n]$. Let \mathcal{C}_∞ be the infinite component of G^+ . Further, the giant component of G'_n is part of \mathcal{C}_∞ with probability tending to 1 as n tends to infinity, see [34]. So for n large enough, $\mathbb{P}(\mathcal{C}_{\max}^{(n)} \not\subseteq \mathcal{C}_\infty) \leq \delta^{3/2}/12$, and hence

$$\mathbb{P}(\neg \mathcal{A}(G_n, u_n, v_n)) \leq \delta/3 + 2\mathbb{P}(\neg \mathcal{A}(G'_n, x'_n, y'_n) \text{ and } x'_n, y'_n \in \mathcal{C}_\infty)/\sqrt{\delta}.$$

Let r be large enough for Corollary 3.12 and Theorem 1.12 to apply to $\delta^{3/2}/12$, taking $\zeta = \delta^2$ in Corollary 3.12. Let X be the set of pairs $(x, y) \in Q_{(1-d\delta^2)n}$ such that $|x - y| \geq r$, and let \mathcal{X}_n be the event that $(x'_n, y'_n) \in X$. Observe that $\mathbb{P}(\neg \mathcal{X}_n) \leq \delta^{3/2}/12$ if $\delta \ll_* d$ and $n \gg_* r$. Thus

$$\begin{aligned} \mathbb{P}(\neg \mathcal{A}(G_n, u_n, v_n)) &\leq \delta/2 + 2\mathbb{P}(\neg \mathcal{A}(G'_n, x'_n, y'_n) \text{ and } x'_n, y'_n \in \mathcal{C}_\infty \mid \mathcal{X}_n \text{ and } x'_n, y'_n \in \mathcal{V}_n)/\sqrt{\delta} \\ &\leq \delta/2 + 2 \max_{x, y \in X} \mathbb{P}(\neg \mathcal{A}(G'_n, x, y) \mid x_n = x, y_n = y, x, y \in \mathcal{C}_\infty)/\sqrt{\delta}. \end{aligned}$$

By Theorem 1.12, the lower cost bound in the event $\mathcal{A}(G'_n, x, y)$ fails with probability at most $\delta^{3/2}/8$. Moreover, when $x, y \in Q_{(1-d\delta^2)n}$, the upper cost bound in $\mathcal{A}(G'_n, x, y)$ occurs whenever x and y are joined by a path in \mathcal{C}_∞ with cost at most $\kappa|x - y|$ and deviation at most $\delta^2|x - y| \leq \delta^2nd$; thus by Corollary 3.12, the upper bound fails with probability at most $\delta^{3/2}/8$. By a union bound, it follows that $\mathbb{P}(\neg \mathcal{A}(G_n, u_n, v_n)) \leq \delta$ as required. \square

Supplementary material

A Appendix

Here we prove Lemmas 3.2, 3.3 and 3.7. For Lemma 3.2 we use the following theorem, which is a simplified version of [19, Theorem 1.1].

Theorem A.1. *For all $d \in \mathbb{N}$ with $d \geq 2$ and all $\sigma \in (0, 1/2)$ there exists $p_0 = p_0(\sigma, d) \in (0, 1)$ such that the following holds. Let ω^* be an iid Bernoulli bond percolation on \mathbb{Z}^d with edge-retention probability $p \geq p_0$. For $S \subseteq \mathbb{R}^d$, let*

$$\mathcal{A}_{\text{dense}}(S, \sigma) := \{S \text{ contains a unique cluster } \mathcal{C} \text{ with } |\mathcal{C}| \geq (1 - \sigma)\text{Vol}(S)\}. \quad (\text{A.1})$$

Let $S_n = [-n/2, n/2]^d$. Then

$$\limsup_{n \rightarrow \infty} \frac{1}{n^{d-1}} \log \mathbb{P}(\neg \mathcal{A}_{\text{dense}}(S_n, \sigma)) < 0.$$

The corresponding [19, Theorem 1.1] is stated for site percolation. Since the event $\mathcal{A}_{\text{dense}}(S, \delta)$ is monotone; the same result for bond percolation also holds via the standard domination in which we retain a vertex in a coupled site percolation process if and only if we retain all its edges in bond percolation. We now use Theorem A.1 to prove Lemma 3.2, showing that the infinite cluster ω of a high-density Bernoulli bond-percolation has locally high density, see also (3.1).

Proof of Lemma 3.2. The statement that $\mathbb{P}(0 \in \mathcal{C}_\infty^*) \geq 1 - \sigma$ as the edge-retention probability $1 - \varepsilon$ tends to 1 follows directly from the continuity of the theta-function in the supercritical regime (in particular, near 1), see [26]. We move on to showing (3.1). By translation invariance, it is enough to show the statement for $x = 0 \in \mathbb{Z}^d$. Recall $\mathcal{A}_{\text{dense}}(S_n, \sigma)$ from (A.1) and that $S_n = [-n/2, n/2]^d$. By Theorem A.1, taking $p_0 := 1 - \varepsilon$ and $\eta \ll_* \sigma$, since the dimension is $d \geq 2$,

$$\forall n \geq n_0 : \mathbb{P}(\mathcal{A}_{\text{dense}}(S_n, \sigma/2)) \geq 1 - \exp(-\eta n^{d-1}) \geq 1 - \exp(-\eta n). \quad (\text{A.2})$$

By the Borel-Cantelli lemma, almost surely $\mathcal{A}_{\text{dense}}(S_n, \sigma/2)$ occurs for all but finitely many values of n . Moreover, if $\mathcal{A}_{\text{dense}}(S_n, \sigma/2)$ and $\mathcal{A}_{\text{dense}}(S_{n+1}, \sigma/2)$ both occur for suitably large n , their respective clusters must intersect by the pigeonhole principle and hence must be equal. Thus almost surely there exists a single cluster \mathcal{C}_∞^* such that for all but finitely many n , $|\mathcal{C}_\infty^* \cap S_n| \geq (1 - \sigma/2)\text{Vol}(S_n)$. Such a cluster is necessarily both infinite and unique, as required. Hence, let

$$\mathcal{A}_\infty(S, \sigma/2) := \{|S \cap \mathcal{C}_\infty^*| \geq (1 - \sigma/2)\text{Vol}(S)\}. \quad (\text{A.3})$$

Then using (A.2), and the above intersection property when one moves from S_i to S_{i+1} , there exists $n_1 \geq n_0$ such that

$$\forall n \geq n_1 : \mathbb{P}(\mathcal{A}_\infty(S_n, \sigma/2)) \geq 1 - \sum_{i \geq n} \mathbb{P}(\neg \mathcal{A}_{\text{dense}}(S_i, \sigma/2)) \geq 1 - \exp(-\eta n/2). \quad (\text{A.4})$$

We now develop a boxing scheme. By definition, $S_{4r} = [-2r, 2r]^d$ which fully contains $B_r(0)$, and further, for each $y \in B_r(0)$, also $B_{(\log r)^{3/2}}(y)$ is fully contained in S_{4r} . For all $r > 0$, let $\ell(r) := (\log r)^{4/3}$, so that $r = \exp(\ell(r)^{3/4})$. Let \mathcal{S}_{4r} be a partition of S_{4r} into at most $(4r/\ell(r))^d$ boxes of (equal) side length $\ell \in [\ell(r), 2\ell(r)]$. We may assume that $r \gg_* \delta, \sigma, \eta$ such that $\ell(r) \geq n_1$ and furthermore:

- (i) $(4r/\ell(r))^d \cdot \exp(-\eta\ell/2) = 4^d \exp(d\ell(r)^{3/4} - \eta\ell/2)/\ell(r)^d \leq \delta$ for all $\ell \in [\ell(r), 2\ell(r)]$ and

- (ii) for all $y \in B_r(x)$, there exist disjoint boxes $S_1^{(y)}, \dots, S_t^{(y)} \in \mathcal{S}_{4r}$ such that $S_i^{(y)} \subseteq B_{(\log r)^{3/2}}(y)$ for all $i \in [t]$ and such that $\text{Vol}(S_1^{(y)} \cup \dots \cup S_t^{(y)}) \geq (1 - \sigma/2)|B_{(\log r)^{3/2}}(y) \cap \mathbb{Z}^d|$.

Here, (i) implies that the probability that there is a box S in the boxing scheme \mathcal{S}_{4r} for which $\mathcal{A}_\infty(S, \sigma/2)$ does not hold is at most δ , by combining (A.4) with a union bound over the at most $(4r/\ell(r))^d$ boxes. So, introducing $\mathcal{A}_{\text{good}}(4r, \sigma/2)$ as the event that for all $S \in \mathcal{S}_{4r}$, the event $\mathcal{A}_\infty(S, \sigma/2)$ in (A.3) occurs, then

$$\mathbb{P}(\mathcal{A}_{\text{good}}(4r, \sigma/2)) \geq 1 - (4r/\ell(r))^d \exp(-\eta\ell/2) \geq 1 - \delta. \tag{A.5}$$

Suppose that $\mathcal{A}_{\text{good}}(4r, \sigma/2)$ occurs, fix some $y \in B_r(0)$ and let $S_1^{(y)}, \dots, S_t^{(y)}$ be the boxes in (ii) contained in $B_{(\log r)^{3/2}}(y)$. Then by the uniqueness of the cluster \mathcal{C}_∞^* ,

$$\begin{aligned} \frac{|B_{(\log r)^{3/2}}(y) \cap \mathcal{C}_\infty^*|}{|B_{(\log r)^{3/2}}(y) \cap \mathbb{Z}^d|} &\geq \frac{1}{|B_{(\log r)^{3/2}}(y) \cap \mathbb{Z}^d|} \sum_{i \in [t]} |S_i^{(y)} \cap \mathcal{C}_\infty^*| \\ &\geq \frac{1}{|B_{(\log r)^{3/2}}(y) \cap \mathbb{Z}^d|} \sum_{i \in [t]} (1 - \frac{\sigma}{2}) \text{Vol}(S_i^{(y)}) \geq (1 - \frac{\sigma}{2})^2 > 1 - \sigma, \end{aligned}$$

where in the last step we used that the union of the boxes $S_1^{(y)}, \dots, S_t^{(y)}$ covers at least $1 - \sigma$ proportion of the vertices in $|B_{(\log r)^{3/2}}(y) \cap \mathbb{Z}^d|$ by (ii). Hence, (3.1) follows from (A.5). \square

We next prove Lemma 3.3, which states that the infinite component of highly supercritical Bernoulli bond percolation has linear graph distances realised by paths with sublinear deviation.

Proof of Lemma 3.3. We first recall a result on linear scaling of distances in \mathcal{C}_∞^* . Let $\{x \leftrightarrow y\}$ be the event that $x, y \in \mathbb{Z}^d$ lie in the same component of ω^* , and let $d_\star(x, y)$ denote their graph distance in ω^* . By [3, Theorem 1.1], there exists $\kappa', \varepsilon_0 > 0$ such that for all $\varepsilon \leq \varepsilon_0$ and for all $a \in \mathbb{Z}^d$,⁵

$$\limsup_{|y| \rightarrow \infty} \frac{1}{|x - y|} \log \mathbb{P}(\{x \leftrightarrow y\} \cap \{d_\star(x, y) > \kappa'|x - y|\}) < 0. \tag{A.6}$$

We say that a site $x \in \mathbb{Z}^d$ has *r-linear scaling* if $x \in \mathcal{C}_\infty^*$ and

$$d_\star(x, y) \leq \kappa'|x - y| \quad \text{holds for all } y \in \mathcal{C}_\infty^* \setminus B_r(x). \tag{A.7}$$

By a union bound over all $y \in \mathbb{Z}^d \setminus B_r(x)$, it follows from (A.6) that there exist $c_1, r_1 > 0$ (depending on d) such that for all $x \in \mathbb{Z}^d$ and $r \geq r_1$,

$$\mathbb{P}(x \text{ has } r\text{-linear scaling or } x \notin \mathcal{C}_\infty^*) \geq 1 - e^{-c_1 r}. \tag{A.8}$$

We now use (A.8) to show that there exist constants $c_2, r_2 > 0$ (depending on ζ and d) such that if $x, y \in \mathcal{C}_\infty^*$ and $|x - y| \geq r_2$, then with probability at least $1 - e^{-c_2|x-y|}$ there exists a linear-length low-deviation path from x to y . Let $\kappa := 2\sqrt{d}\kappa'$ and $K := (\kappa' + 1)\sqrt{d}/\zeta$, and let $x, y \in \mathcal{C}_\infty^*$. Cover the straight line segment $S_{x,y}$ by a sequence of $k \leq 2K$ cubes $Q^{(1)}, \dots, Q^{(k)}$ in such a way that (see also Figure 2):

- (a) each cube $Q^{(i)}$ has side length $|x - y|/K$;

⁵The formulation of [3, Theorem 1.1] allows κ' to depend on the edge-retention probability $p = 1 - \varepsilon$ because it allows p to be arbitrarily close to the criticality threshold. So it is not quite clear from the formulation that κ' is independent of ε . However, the condition $\varepsilon \leq \varepsilon_0$ means that our p is bounded away from the criticality threshold, and the proof in [3] relies on a domination argument, so we may use the same κ' for all $\varepsilon \leq \varepsilon_0$.




Figure 2: Example for $k = 3$ cubes covering the line from x to y . Any two adjacent cubes overlap by at least $4^{-d}\text{Vol}(Q^{(i)})$ and contain a cube of side-length $|x - y|/(8K)$ and any two intersections have distance at least $r/2 = |x - y|/(2K)$, where $r = |x - y|/K$ is the side-length of the cubes.

- (b) $Q^{(1)}$ is the cube covering x , $Q^{(k)}$ is the cube covering y , and $Q^{(i)} \cap Q^{(i+1)} \cap S_{x,y} \neq \emptyset$ for all $i \in [k - 1]$;
- (c) for all $i \in [k - 1]$, the intersection $Q^{(i)} \cap Q^{(i+1)}$ contains a cube of side-length $|x - y|/(8K)$;
- (d) all pair of sites a, b that fall into distinct sets in the list $\{x\}, Q^{(1)} \cap Q^{(2)}, Q^{(2)} \cap Q^{(3)}, \dots, Q^{(k-1)} \cap Q^{(k)}, \{y\}$ satisfy $|a - b| \geq |x - y|/(2K)$.

Then the following events occur with probability at least $1 - e^{-c_2|x-y|}$ if $|x - y| \geq r_2$ (we specify c_2 and r_2 below):

- (i) For all $i \in [k - 1]$, $Q^{(i)} \cap Q^{(i+1)}$ contains at least one site $s_i \in \mathcal{C}_\infty^*$;
- (ii) Both x and y have r -linear scaling with $r = |x - y|/(2K)$ conditioned on $x, y \in \mathcal{C}_\infty^*$;
- (iii) For all $i \in [k - 1]$, all sites $z \in Q^{(i)} \cap Q^{(i+1)} \cap \mathcal{C}_\infty^*$ have r -linear scaling with $r = |x - y|/(2K)$.

Indeed, since $\varepsilon \ll_\star d$, \mathcal{C}_∞^* has density at least $1/2$ and since the intersection $Q^{(i)} \cap Q^{(i+1)}$ contains a cube S of side-length $|x - y|/(8K)$, (A.4) applies, i.e., the event in (A.3) that $S \cap \mathcal{C}_\infty$ contains linearly vertices in the volume of S holds with probability $1 - \exp(-\eta|x - y|/(8K))$. This event implies that there is at least one vertex in $\mathcal{C}_\infty^* \cap S$ and with $c_3 := \eta/8K$, there exists $r_3 > 0$ depending on ζ, d such that (i) holds with probability at least $1 - e^{-c_3|x-y|}$ whenever $|x - y| \geq r_3$. Moreover, using (A.8) and a union bound over $i \in [k - 1]$ and over all sites $z \in Q^{(i)} \cap Q^{(i+1)}$ (polynomially many in $|x - y|$) and over x, y , we deduce that there exists $r_4 > 0$ depending on d such that (ii) and (iii) also hold with probability $1 - e^{-c_1|x-y|/(4K)}$ whenever $|x - y| \geq r_4$. Taking $r_2 = \max\{r_3, r_4\}$ and c_2 small enough so that $e^{-c_3r} + e^{-c_1r/(4K)} \leq e^{-c_2r}$ for all $r \geq r_2$ guarantees that the above events occur with probability at least $1 - e^{-c_2|x-y|}$ whenever $|x - y| \geq r_2$.

Suppose the events in (i)–(ii) all occur. Choose now for all $i \in [k - 1]$ sites $s_i \in \mathcal{C}_\infty^* \cap Q^{(i)} \cap Q^{(i+1)}$. Writing $s_0 = x, s_k = y$, it holds for all $i \in [k]$ that

$$|s_{i-1} - s_i| > |x - y|/(2K) \quad \text{and} \quad |s_{i-1} - s_i| \leq \sqrt{d}|x - y|/K. \tag{A.9}$$

Then the r -linear-scaling property of s_i and s_{i-1} with $r = |x - y|/(2K)$ in (A.7), combined with (A.9), implies deterministically that for all $i \in [k]$ there exists a path π_i from s_{i-1} to s_i in \mathcal{C}_∞^* of length at most $\kappa'\sqrt{d}|x - y|/K$. Since s_{i-1} has distance at most $\sqrt{d}|x - y|/K$ from the segment $S_{x,y}$ and π_i contains nearest neighbour edges, by the definition of $K = (\kappa' + 1)\sqrt{d}/\zeta$ at the beginning of the proof and Definition 3.1,

$$\text{dev}_{xy}(\pi_i) \leq \kappa'\sqrt{d}|x - y|/K + \sqrt{d}|x - y|/K = \zeta|x - y|. \tag{A.10}$$

Let $\pi = \pi_1 \dots \pi_k$; then since $k \leq 2K$, and $\kappa = 2\sqrt{d}\kappa'$, π has length at most $k\kappa'\sqrt{d}|x - y|/K \leq \kappa|x - y|$ and deviation at most $\max(\text{dev}_{xy}(\pi_i)) \leq \zeta|x - y|$ (see Definition 3.1). Therefore, we have shown that if $x, y \in \mathcal{C}_\infty^*$ and $|x - y| \geq r_2$, then with probability at least $1 - e^{-c_2|x - y|}$ there exists a path π from x to y with length at most $|\pi| \leq \kappa|x - y|$ and deviation at most $\text{dev}(\pi) \leq \zeta|x - y|$. Taking a union bound over all $y \in \mathcal{C}_\infty^* \setminus B_r(x)$ and taking $|x - y| \geq r \gg_* \zeta$ yields (3.2) with $c = c_2/2$. \square

Lastly, we prove Lemma 3.7, which uses Lemma 3.2 to obtain that the infinite subgraph \mathcal{H}_∞ of a (R, ε) -dense random geometric graph also has overall high density.

Proof of Lemma 3.7. Let ω^* be a bond percolation process with retention probability $1 - 20d\varepsilon$. We take ε small enough that Lemma 3.2 applies with $\sigma_{3.2} = \sigma^3$ and $\delta_{3.2} = \delta/3$; thus ω^* has a unique infinite component \mathcal{C}_∞^* and \mathcal{H}_∞ is well-defined in (3.4). Then, since \mathcal{H}_∞ consists of those boxes S_z for which $z \in \mathbb{Z}^d$ belong to \mathcal{C}_∞^* in ω^* , if $z \in \mathcal{C}_\infty^*$, z has at least one open adjacent bond. By the defining coupling in Lemma 3.5, such a bond is only open in ω^* if $G[S_z]$ is non-empty and connected. This implies that if $z \in \mathcal{C}_\infty^*$ then all the vertices in $G[S_z]$ are in \mathcal{H}_∞ . Since $\mathbb{P}(0 \in \mathcal{C}_\infty^*) \geq 1 - \sigma_{3.2}$ in Lemma 3.2, the translation invariance of \mathbb{Z}^d readily implies that for any box S , $\mathbb{P}(\mathcal{V} \cap S \subseteq \mathcal{H}_\infty) \geq 1 - \sigma_{3.2} \geq 1 - \sigma$.

We turn to prove (3.5). Throughout, let $\rho = (\log r)^2$ for brevity. We set out some preliminary notation and observations. Recall that for all $z \in \mathbb{Z}^d$, S_z is the unique box in \mathcal{S} such that $Rz \in S_z$ and that boxes have side-length R . For all sets $A \subseteq \mathbb{R}^d$ we write $A^* = \{z \in \mathbb{Z}^d : S_z \subseteq A\}$; for the renormalisation of A , thus $\text{Boxes}(A) := \bigcup_{z \in A^*} S_z$ is the union of all boxes fully contained in A . Finally, for all $y \in \mathbb{R}^d$, by $\lfloor y/R \rfloor$ we mean taking lower-integer part of each coordinate, here $\lfloor y/R \rfloor$ is the renormalised site, i.e., in \mathbb{Z}^d , corresponding to the (box containing) $z \in \mathbb{R}^d$. Clearly then $y \in S_{\lfloor y/R \rfloor}$. Since $r \gg_* R$, and the diameter of each box is $\sqrt{d}R$, it is not hard to see that for all $y \in \mathbb{R}^d$, and with $\rho = (\log r)^2$,

$$\begin{aligned} B_{\rho-R\sqrt{d}}(y) &\subseteq \text{Boxes}(B_\rho(y)) \subseteq B_\rho(y), \\ B_{\rho/R-2\sqrt{d}}(\lfloor y/R \rfloor) \cap \mathbb{Z}^d &\subseteq B_\rho(y)^* \subseteq B_{\rho/R+\sqrt{d}}(\lfloor y/R \rfloor), \\ B_{\rho-R\sqrt{d}}(y)^* &\subseteq B_\rho(R\lfloor y/R \rfloor)^* \cap B_\rho(y)^*. \end{aligned}$$

Thus there exists $C > 0$ depending on R and d (but not on r and thus neither on ρ) such that for all $y \in \mathbb{R}^d$,

$$|\text{Vol}(B_\rho(y)) - R^d|B_\rho(y)^*|| \leq C\rho^{d-1} \leq \sigma^4\rho^d, \tag{A.11}$$

$$||B_\rho(y)^*| - |B_{\rho/R}(\lfloor y/R \rfloor) \cap \mathbb{Z}^d|| \leq C\rho^{d-1} \leq \sigma^4\rho^d, \tag{A.12}$$

$$|B_\rho(R\lfloor y/R \rfloor)^* \setminus B_\rho(y)^*| \leq C\rho^{d-1} \leq \sigma^4\rho^d. \tag{A.13}$$

These all intuitively recover the isoperimetric properties of \mathbb{R}^d : for sufficiently large ρ , the volume of a ball of radius ρ can be well approximated by a boxing scheme of side-length R (A.11), and also the number of boxes we need both have error of order $\rho^{(d-1)/d}$ (A.12), and switching from a renormalised set around y to the renormalised set around the renormalised site corresponding to y also causes small error (A.13).

Now we relate (3.1) to (3.5). Note that (3.1) has radius $(\log r)^{3/2}$ inside the probability sign concerning overall high density in ω^* , while (3.5) has radius $(\log r)^2$ concerning overall high density of \mathcal{H}_∞ in G . Hence, given r, R , let r^* be such that $(\log r^*)^{3/2} = (\log r)^2/R$, i.e. $r^* := \exp(((\log r)^2/R)^{2/3})$, and clearly for all sufficiently large $r \gg_* R$, $r^* > 2r/R$. We will now apply Lemma 3.2 with $r = r^*$. Then, the radius inside the fraction in (3.1) is exactly $(\log r^*)^{3/2} = (\log r)^2/R$. So, for $x \in \mathbb{R}^d$, when we write the event in (3.1) for the site $\lfloor x/R \rfloor \in \mathbb{Z}^d$ and using radius r^* and $\sigma_{3.2} = \sigma^3$, we obtain

$$\mathcal{A}_{\text{dense}}^*(x, r^*) = \left\{ \forall y \in B_{r^*}(\lfloor x/R \rfloor) : \frac{|B_{(\log r)^2/R}(y) \cap \mathcal{C}_\infty^*|}{|B_{(\log r)^2/R}(y) \cap \mathbb{Z}^d|} \geq 1 - \sigma^3 \right\}. \tag{A.14}$$

Now applying Lemma 3.2 yields directly that $\mathbb{P}(\mathcal{A}_{\text{dense}}^*) \geq 1 - \delta_{3.2} = 1 - \delta/3$. When we consider any $y \in B_r(x)$ in \mathbb{R}^d , clearly $\lfloor y/R \rfloor \in B_{2r/R}(x) \subseteq B_{r^*}(\lfloor x/R \rfloor)$ because $r^* \geq 2r/R$; thus if (A.14) holds, then also

$$\forall y \in B_r(x) : \frac{|B_{(\log r)^2/R}(\lfloor y/R \rfloor) \cap \mathcal{C}_\infty^*|}{|B_{(\log r)^2/R}(\lfloor y/R \rfloor) \cap \mathbb{Z}^d|} \geq 1 - \sigma^3.$$

Note that $B_{(\log r)^2/R}(\lfloor y/R \rfloor)$, when we move from sites in \mathbb{Z}^d to boxes in \mathcal{S} , corresponds roughly to the set of boxes contained in $B_{(\log r)^2}(y)$. In fact (A.12) exactly quantifies the error being small. Hence, since r is large, it follows from (A.12) that $\mathcal{A}_{\text{dense}}^*(x, r^*)$ also implies the following event:

$$\mathcal{A}_{\text{dense}}(x, r) := \left\{ \forall y \in B_r(x) : \frac{|B_{(\log r)^2}(y)^* \cap \mathcal{C}_\infty^*|}{|B_{(\log r)^2}(y)^*|} \geq 1 - 2\sigma^3 \right\}. \tag{A.15}$$

We have shown that

$$\mathbb{P}(\mathcal{A}_{\text{dense}}(x, r)) \geq \mathbb{P}(\mathcal{A}_{\text{dense}}^*(x, r^*)) \geq 1 - \delta/3. \tag{A.16}$$

Interpreting the event on the lhs, $\mathcal{A}_{\text{dense}}$ says that sites in \mathcal{C}_∞^* form high density in $B_{(\log r)^2}(y)^*$, or equivalently that boxes containing vertices of \mathcal{H}_∞ are dense in $B_{(\log r)^2}(y)$; meanwhile, the event in (3.5) says that the vertices inside these boxes are dense in $B_{(\log r)^2}(y) \cap \mathcal{V}$. Thus we can achieve the high density in (3.5) if we can control the number of vertices per box.

We now prove a concentration bound for the number of vertices of G in a given collection of boxes; since \mathcal{C}_∞^* is random, we basically will sum the errors over all realisations of \mathcal{C}_∞^* satisfying $\mathcal{A}_{\text{dense}}(x, r)$. We again abbreviate $\rho := (\log r)^2$. Here below, we denote by A any set in $B_\rho(y)^*$ with number of vertices $|A| \leq 2\sigma^3|B_\rho(y)^*|$ so that A can essentially serve as a possible realisation of the complement of \mathcal{C}_∞^* inside the ball $B_\rho(y)^*$ when the event $\mathcal{A}_{\text{dense}}(x, r)$ holds. Then, define \mathcal{A}_{box} as

$$\mathcal{A}_{\text{box}} := \bigcap_{y \in B_r(x)} \bigcap_{\substack{A \subseteq B_\rho(y)^* \\ |A| \leq 2\sigma^3|B_\rho(y)^*|}} \left\{ \left| \bigcup_{z \in B_\rho(y)^* \setminus A} \mathcal{V}[S_z] \right| \geq (1 - \sigma/5)\text{Vol}(B_\rho(y)) \right\}, \tag{A.17}$$

i.e., that leaving out the boxes in any not-too-large set A from $B_\rho(y)^*$, the remaining boxes still contain enough vertices proportional to the volume. When $\mathcal{V} = \mathbb{Z}^d$, \mathcal{A}_{box}

always occurs. For \mathcal{V} a Poisson point process, we dominate \mathcal{A}_{box} by an intersection of events as follows. For all $y \in \mathbb{R}^d$ and $A \subseteq \mathbb{Z}^d$, let

$$\mathcal{A}_{\text{box}}(y, A) := \left\{ \left| \bigcup_{z \in B_\rho(y)^* \setminus A} \mathcal{V}[S_z] \right| \geq (1 - \sigma/10)R^d |B_\rho(y)^*| \right\}. \tag{A.18}$$

Applying (A.11) yields that $(1 - \sigma/10)R^d |B_\rho(y)^*| \geq (1 - \sigma/5)\text{Vol}(B_\rho(y))$. Moreover, by (A.13), we can lower-bound $B_\rho(y)^*$ in (A.17) by $B_\rho(R\lfloor y/r \rfloor)^* \setminus A_y$ for some set A_y with $|A_y| \leq \sigma^3 |B_\rho(y)^*|$, effectively discretising the continuous intersection over $y \in B_r(x)$. Thus

$$\begin{aligned} \mathcal{A}_{\text{box}} &\supseteq \bigcap_{y \in B_r(x)} \bigcap_{\substack{A \subseteq B_\rho(y)^* \\ |A| \leq 2\sigma^3 |B_\rho(y)^*|}} \mathcal{A}_{\text{box}}(y, A) \supseteq \bigcap_{y \in B_r(x)} \bigcap_{\substack{A \subseteq B_\rho(R\lfloor y/R \rfloor)^* \\ |A| \leq 4\sigma^3 |B_\rho(R\lfloor y/R \rfloor)^*|}} \mathcal{A}_{\text{box}}(R\lfloor y/R \rfloor, A) \\ &\supseteq \bigcap_{z \in B_{2r}(x)^*} \bigcap_{\substack{A \subseteq B_\rho(Rz)^* \\ |A| \leq 4\sigma^3 |B_\rho(Rz)^*|}} \mathcal{A}_{\text{box}}(Rz, A), \end{aligned} \tag{A.19}$$

We obtained the last row by noting that the rhs of the first row is the same event for all y with the same renormalised site $\lfloor y/R \rfloor$. When \mathcal{V} is a PPP, the number of vertices in boxes in $B_\rho(y)^* \setminus A$ in (A.18) follows a Poisson distribution with mean $R^d |B_\rho(y)^* \setminus A|$; thus when $|A| \leq 4\sigma^3 |B_\rho(y)^*|$, this mean is at least $(1 - \sigma/20)R^d |B_\rho(y)^*|$, and so by a Chernoff bound we arrive to

$$\mathbb{P}(\neg \mathcal{A}_{\text{box}}(y, A)) \leq \exp(-\sigma^2 \rho^d / 300).$$

Now some combinatorics to deal with the union when taking the complement event in (A.19): since σ is small, the number of sets $A \subseteq B_\rho(Rz)^*$ with $|A| \leq 4\sigma^3 |B_\rho(Rz)^*|$ is at most $\exp(\sigma^2 \rho^d / 600)$. Moreover, there are at most $4(r/R)^d$ choices of $z \in B_{2r}(x)^*$ in (A.19). Hence using a union bound, $\rho = (\log r)^2$, and $r \gg_* \delta, \sigma$, for Poisson \mathcal{V} ,

$$\mathbb{P}(\mathcal{A}_{\text{box}}) \geq 1 - 4(r/R)^d \exp(\sigma^2 \rho^d / 600) \cdot \exp(-\sigma^2 \rho^d / 300) \geq 1 - \delta/3. \tag{A.20}$$

Our final event, $\mathcal{A}_{\text{ball}}$, says that all radius- $\rho = (\log r)^2$ balls near x contain at most roughly the expected number of vertices, so that

$$\begin{aligned} \mathcal{A}_{\text{ball}} &= \bigcap_{y \in B_r(x)} \left\{ |B_\rho(y) \cap \mathcal{V}| \leq (1 + \sigma/5)\text{Vol}(B_\rho(y)) \right\} \\ &\supseteq \bigcap_{y \in B_r(x) \cap \mathbb{Z}^d} \left\{ |B_\rho(y) \cap \mathcal{V}| \leq (1 + \sigma/10)\text{Vol}(B_\rho(y)) \right\}. \end{aligned} \tag{A.21}$$

where in the second row we discretised the space to obtain a finite intersection at the cost of reducing the error to $\sigma/10$. It is then immediate from Chernoff bounds and a union bound over y that

$$\mathbb{P}(\mathcal{A}_{\text{ball}}) \geq 1 - \delta/3. \tag{A.22}$$

Now by a union bound on their complements in (A.16), (A.20), (A.22), the intersection $\mathcal{A}_{\text{dense}} \cap \mathcal{A}_{\text{box}} \cap \mathcal{A}_{\text{ball}}$ occurs with probability at least $1 - \delta$. Assume the intersection of the three events occur and let $y \in B_r(x)$. Since $\mathcal{A}_{\text{dense}}$ occurs, $|B_\rho(y)^* \setminus \mathcal{C}_\infty^*| \leq 2\sigma^3 |B_\rho(y)^*|$. Since \mathcal{A}_{box} also occurs in (A.15), taking $A = B_\rho(y)^* \setminus \mathcal{C}_\infty^*$ yields

$$|B_\rho(y) \cap \mathcal{H}_\infty| = \left| B_\rho(y) \cap \bigcup_{z \in \mathcal{C}_\infty^*} \mathcal{V}[S_z] \right| \geq \left| \bigcup_{z \in B_\rho(y)^* \cap \mathcal{C}_\infty^*} \mathcal{V}[S_z] \right| \geq (1 - \sigma/5)\text{Vol}(B_\rho(y)).$$

Finally, since $\mathcal{A}_{\text{ball}}$ occurs in (A.21), it follows that

$$\frac{|B_\rho(y) \cap \mathcal{H}_\infty|}{|B_\rho(y) \cap \mathcal{V}|} \geq \frac{(1 - \sigma/5)\text{Vol}(B_\rho(y))}{(1 + \sigma/5)\text{Vol}(B_\rho(y))} \geq 1 - \sigma.$$

Hence, the event in (3.5) is implied by the intersection of $\mathcal{A}_{\text{dense}} \cap \mathcal{A}_{\text{box}} \cap \mathcal{A}_{\text{ball}}$, which finishes the proof of (3.5). \square

References

- [1] Erwin Adriaans and Júlia Komjáthy, *Weighted distances in scale-free configuration models*, Journal of Statistical Physics **173** (2018), 1082–1109. MR3876918
- [2] Enrique D Andjel, *Characteristic exponents for two-dimensional bootstrap percolation*, The Annals of Probability **21** (1993), no. 2, 926–935. MR1217573
- [3] Peter Antal and Agoston Pisztora, *On the chemical distance for supercritical bernoulli percolation*, The Annals of Probability (1996), 1036–1048. MR1404543
- [4] Antonio Auffinger, Michael Damron, and Jack Hanson, *50 years of first-passage percolation*, vol. 68, American Mathematical Soc., 2017. MR3729447
- [5] Johannes Bäumlér, *Distances in $1/\|x - y\|^{2d}$ percolation models for all dimensions*, Communications in Mathematical Physics **404** (2023), no. 3, 1495–1570. MR4677584
- [6] Noam Berger, *A lower bound for the chemical distance in sparse long-range percolation models*, arXiv preprint arXiv:math/0409021 (2004).
- [7] S Bhamidi, RW van der Hofstad, and G Hooghiemstra, *Universality for first passage percolation on sparse random graphs*, The Annals of Probability **45** (2017), no. 4, 2568–2630. MR3693970
- [8] Nicholas H Bingham, Charles M Goldie, and Jef L Teugels, *Regular variation*, vol. 27, Cambridge university press, 1989. MR1015093
- [9] Marek Biskup, *On the scaling of the chemical distance in long-range percolation models*, The Annals of Probability **32** (2004), no. 4, 2938–2977. MR2094435
- [10] Marek Biskup and Jeffrey Lin, *Sharp asymptotic for the chemical distance in long-range percolation*, Random Structures & Algorithms **55** (2019), no. 3, 560–583. MR3997479
- [11] Moreno Bonaventura, Vincenzo Nicosia, and Vito Latora, *Characteristic times of biased random walks on complex networks*, Physical Review E **89** (2014), no. 1, 012803.
- [12] Karl Bringmann, Ralph Keusch, and Johannes Lengler, *Average distance in a general class of scale-free networks with underlying geometry*, arXiv preprint arXiv:1602.05712 (2016).
- [13] Karl Bringmann, Ralph Keusch, and Johannes Lengler, *Geometric inhomogeneous random graphs*, Theoretical Computer Science **760** (2019), 35–54. MR3913223
- [14] Shirshendu Chatterjee and Partha S. Dey, *Multiple phase transitions in long-range first-passage percolation on square lattices*, Communications on Pure and Applied Mathematics **69** (2016), no. 2, 203–256. MR3434612
- [15] J Theodore Cox and Richard Durrett, *Some limit theorems for percolation processes with necessary and sufficient conditions*, The Annals of Probability (1981), 583–603. MR0624685
- [16] Maria Deijfen, Remco van der Hofstad, and Gerard Hooghiemstra, *Scale-free percolation*, Annales de l’IHP Probabilités et statistiques, vol. 49,3, 2013, pp. 817–838. MR3112435
- [17] Philippe Deprez, Rajat Subhra Hazra, and Mario V Wüthrich, *Inhomogeneous long-range percolation for real-life network modeling*, Risks **3** (2015), no. 1, 1–23.
- [18] Philippe Deprez and Mario V Wüthrich, *Scale-free percolation in continuum space*, Communications in Mathematics and Statistics **7** (2019), no. 3, 269–308. MR3986777
- [19] Jean-Dominique Deuschel and Agoston Pisztora, *Surface order large deviations for high-density percolation*, Probability Theory and Related Fields (1996), 467–482. MR1384041
- [20] Cangfeng Ding and Kan Li, *Centrality ranking in multiplex networks using topologically biased random walks*, Neurocomputing **312** (2018), 263–275.
- [21] Paul Embrechts and Charles M Goldie, *On closure and factorization properties of subexponential and related distributions*, Journal of the Australian Mathematical Society **29** (1980), no. 2, 243–256. MR0566289
- [22] Joshua Feldman and Jeannette Janssen, *High degree vertices and spread of infections in spatially modelled social networks*, International Workshop on Algorithms and Models for the Web-Graph, Springer, 2017, pp. 60–74. MR3713316

- [23] Carolina Fransson, *Stochastic epidemics on random networks and competition in growth*, Ph.D. thesis, Stockholm University, Department of Mathematics, 2023, ISBN: 978-91-8014-145-1 (electronic), p. 36.
- [24] Alan Frieze and Michał Karoński, *Random graphs and networks: A first course*, Cambridge University Press, 2023. MR4567944
- [25] CV Giuraniuc, JPL Hatchett, JO Indekeu, M Leone, I Perez Castillo, Bert Van Schaeybroeck, and Carlo Vanderzande, *Criticality on networks with topology-dependent interactions*, *Physical Review E* **74** (2006), no. 3, 036108.
- [26] Geoffrey Grimmett, *Percolation*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], Springer, 1999. MR1707339
- [27] Luca Gugelmann, Konstantinos Panagiotou, and Ueli Peter, *Random hyperbolic graphs: degree sequence and clustering*, International Colloquium on Automata, Languages, and Programming, Springer, 2012, pp. 573–585.
- [28] Nannan Hao and Markus Heydenreich, *Graph distances in scale-free percolation: the logarithmic case*, *Journal of Applied Probability* **60** (2023), no. 1, 295–313. MR4546124
- [29] Hans Hooyberghs, Bert Van Schaeybroeck, André A Moreira, José S Andrade Jr, Hans J Herrmann, and Joseph O Indekeu, *Biased percolation on scale-free networks*, *Physical Review E* **81** (2010), no. 1, 011102.
- [30] Joost Jorritsma and Júlia Komjáthy, *Weighted distances in scale-free preferential attachment models*, *Random Structures & Algorithms* **57** (2020), no. 3, 823–859. MR4144085
- [31] Márton Karsai, Róbert Juhász, and Ferenc Iglói, *Nonequilibrium phase transitions and finite-size scaling in weighted scale-free networks*, *Physical Review E* **73** (2006), no. 3, 036116.
- [32] Júlia Komjáthy, John Lapinskas, and Johannes Lengler, *Stopping explosion by penalising transmission to hubs in scale-free spatial random graphs*, *Annales Institute Henri Poincaré: Probabilités et Statistiques* (2021). MR4415389
- [33] Júlia Komjáthy, John Lapinskas, Johannes Lengler, and Ulysse Schaller, *Four universal growth regimes in degree-dependent first passage percolation on spatial random graphs I*, arXiv preprint: arXiv:2309.11840 [math.PR] (2023).
- [34] Júlia Komjáthy and Bas Lodewijks, *Explosion in weighted hyperbolic random graphs and geometric inhomogeneous random graphs*, *Stochastic Processes and their Applications* **130** (2020), no. 3, 1309–1367. MR4058275
- [35] Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguná, *Hyperbolic geometry of complex networks*, *Physical Review E* **82** (2010), no. 3, 036106. MR2787998
- [36] S Lee, S-H Yook, and Y Kim, *Centrality measure of complex networks using biased random walks*, *The European Physical Journal B* **68** (2009), no. 2, 277–281.
- [37] Thomas M Liggett, Roberto H Schonmann, and Alan M Stacey, *Domination by product measures*, *The Annals of Probability* **25** (1997), no. 1, 71–95. MR1428500
- [38] Torgny Lindvall, *On Strassen’s theorem on stochastic domination*, *Electronic communications in probability* **4** (1999), 51–59. MR1711599
- [39] Giovanna Miritello, Esteban Moro, Rubén Lara, Rocío Martínez-López, John Belchamber, Sam GB Roberts, and Robin IM Dunbar, *Time as a limited resource: Communication strategy in mobile phone networks*, *Social Networks* **35** (2013), no. 1, 89–95.
- [40] Mathew Penrose et al., *Random geometric graphs*, vol. 5, Oxford university press, 2003. MR1986198
- [41] Cunlai Pu, Siyuan Li, and Jian Yang, *Epidemic spreading driven by biased random walks*, *Physica A: Statistical Mechanics and its Applications* **432** (2015), 230–239. MR3339612
- [42] Lawrence S Schulman, *Long range percolation in one dimension*, *Journal of Physics A: Mathematical and General* **16** (1983), no. 17, L639. MR0723249
- [43] Thomas Slangen, *Cost-distances in penalized scale-free random graphs*, BSc thesis, Eindhoven University of Technology, 2019.
- [44] Pieter Trapman et al., *The growth of the infinite long-range percolation cluster*, *Annals of probability* **38** (2010), no. 4, 1583–1608. MR2663638

- [45] Remco van der Hofstad and Julia Komjathy, *Explosion and distances in scale-free percolation*, arXiv preprint arXiv:1706.02597 (2017).
- [46] Vinko Zlatić, Andrea Gabrielli, and Guido Caldarelli, *Topologically biased random walk and community finding in networks*, *Physical Review E* **82** (2010), no. 6, 066109.

Acknowledgments. We thank the thorough referees for improving the presentation of the paper.