

Improved Bounds for Graph Distances in Scale Free Percolation and Related Models

Kostas Lakis ✉ 

ETH Zürich, Department of Computer Science, Zürich, Switzerland

Johannes Lengler ✉

ETH Zürich, Department of Computer Science, Zürich, Switzerland

Kalina Petrova ✉ 

ETH Zürich, Department of Computer Science, Zürich, Switzerland

Leon Schiller ✉

ETH Zürich, Department of Computer Science, Zürich, Switzerland

Abstract

In this paper, we study graph distances in the geometric random graph models scale-free percolation SFP, geometric inhomogeneous random graphs GIRG, and hyperbolic random graphs HRG. Despite the wide success of the models, the parameter regime in which graph distances are polylogarithmic is poorly understood. We provide new and improved lower bounds. In a certain portion of the parameter regime, those match the known upper bounds.

Compared to the best previous lower bounds by Hao and Heydenreich [19], our result has several advantages: it gives matching bounds for a larger range of parameters, thus settling the question for a larger portion of the parameter space. It strictly improves the lower bounds of [19] for all parameters settings in which those bounds were not tight. It gives tail bounds on the probability of having short paths, which imply shape theorems for the k -neighbourhood of a vertex whenever our lower bounds are tight, and tight bounds for the size of this k -neighbourhood. And last but not least, our proof is much simpler and not much longer than two pages, and we demonstrate that it generalizes well by showing that the same technique also works for first passage percolation.

2012 ACM Subject Classification Mathematics of computing → Random graphs; Mathematics of computing → Stochastic processes; Theory of computation → Random network models

Keywords and phrases Mathematics, Probability Theory, Combinatorics, Random Graphs, Random Metric Spaces

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2024.74

Category RANDOM

Related Version *Full Version*: <https://arxiv.org/abs/2405.07217>

Funding *Kostas Lakis*: gratefully acknowledges support from the John S. Latsis Public Benefit Foundation and the Onassis Foundation

Kalina Petrova: funded by SNSF grant No. CRSII5 173721

Leon Schiller: funded by SNSF grant No. 197138

Acknowledgements This research started at the joint workshop of the *Combinatorial Structures and Algorithms* and *Theory of Combinatorial Algorithms* groups of ETH Zürich held in Stels, Switzerland, January 2024. We thank the organizers for providing a very pleasant and inspiring working atmosphere.

1 Introduction and Main results

In the last years, a family of random graph models including *hyperbolic random graphs* (HRG) [26], *scale-free percolation* (SFP) [14], and *geometric inhomogeneous random graphs* (GIRG) [11], has emerged as a model for large real-world networks. They combine an



© Kostas Lakis, Johannes Lengler, Kalina Petrova, and Leon Schiller;
licensed under Creative Commons License CC-BY 4.0

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2024).

Editors: Amit Kumar and Noga Ron-Zewi; Article No. 74; pp. 74:1–74:22



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

underlying geometric space with an inhomogeneous degree distribution. This combination yields many properties that occur in real-world networks across a wide range of domains, such as strong clustering, a rich community structure, ultra-small distance, small separators, compressibility, and more [11]. The networks have been used empirically and theoretically to study algorithms like local routing protocols [7, 9, 12], bidirectional search [6] or maximum flow [8], spreading processes like bootstrap percolation [21] or SI models [22–25], and they have been used to study the effectiveness of different interventions during the Covid19 pandemic [17, 20, 27].

Despite this widespread adoption, the fundamental question of graph distances has been open for some parameter regimes. In general, the models come with two parameters: the degrees follow a power-law distribution with exponent $\tau > 1$, i.e., $\mathbb{P}(\deg(v) \geq x) \sim x^{1-\tau}$ for any fixed vertex v ;¹ and $\alpha > 1$ determines the number of *weak ties* [18] in the network, i.e. edges which are present although their geometric distance and degrees suggest otherwise. The parameter α is also called *inverse temperature*. If $\tau < 3$ then for two random vertices x, y in the giant component, with high probability² their graph distance $d_G(x, y)$ is at most doubly logarithmic in their geometric distance, i.e., $d_G(x, y) = O(\log \log |x - y|)$. For $\tau \leq 2$ we even have $d_G(x, y) = O(1)$. These regimes are very precisely understood [1, 10, 14]. On the other hand, if $\tau > 3$ and $\alpha > 2$, then it is known that the graph distance of two vertices x, y grows linearly with their geometric distance, and this is again well understood [2, 15].

However, in the *polylogarithmic regime* $\tau > 3$ and $\alpha \in (1, 2)$, the picture is incomplete. It is understood in the limiting case $\tau = \infty$, which is known as *long-range percolation* LRP, that $d_G(x, y) = (\log |x - y|)^{\Delta(\alpha) \pm o(1)}$ where $\Delta(\alpha) = 1/\log_2(2/\alpha) > 1$ [2, 3]. Since graph distances can only increase with τ , the upper bound applies for any $\tau > 1$, and this is the best known upper bound.³ On the other hand, it is easy to see that for $\tau > 3$ the k -neighbourhood can grow at most exponentially, so distances are at least logarithmic, $d_G(x, y) = \Omega(\log |x - y|)$ [14]. Hence we know that distances are polylogarithmic for $\tau > 3$ and $\alpha \in (1, 2)$. But the exponents of the upper and lower bound (Δ and 1 respectively) did not match, and this gap remained open for a long time.

Very recently, Hao and Heydenreich [19] could show an improved lower bound of $d_G(x, y) \geq (\log |x - y|)^{\Delta(\min\{\alpha, (\tau-1)/2\}) - o(1)}$, where, as before, $\Delta(x) = 1/\log_2(2/x)$. This closed the gap in the case that $\alpha < (\tau - 1)/2$, since then $\min\{\alpha, (\tau - 1)/2\} = \alpha$. Their proof had 9 pages and was a complicated application of Biskup’s hierarchy argument [3]. In this paper we give a stronger lower bound with a much simpler inductive proof of only about two pages. It is inspired by ideas of Biskup for the simpler case of LRP [4], which themselves are adaptations of those in [29]. More precisely, we show that $d_G(x, y) = \Omega((\log |x - y|)^{\Delta(\min\{\alpha, \tau-2-o(1)\})})$. This closes the gap between upper and lower bound whenever $\alpha < \tau - 2$, which comprises a strictly larger portion of the polylogarithmic regime than the bound of [19]. Moreover, throughout the polylogarithmic regime, since $\tau > 3$ implies $\tau - 2 > (\tau - 1)/2$, our bound is strictly stronger in all cases in which the bound of [19] is not tight.

¹ The distribution is often allowed to vary by constant factors or slowly varying functions, but this will not be relevant for this paper. Also, the traditional SFP parameterization uses two parameters τ and γ instead of the one parameter τ . However, one of those parameters is internal to the graph generation process and does not yield additional classes of graphs, which is why we omit the additional parameter.

² We say an event occurs *with high probability*, or w.h.p., if it occurs with probability $1 - o(1)$.

³ An improved upper bound was claimed in [19], but the proof had an issue which we consider severe, see Appendix A for details. At submission time of the camera-ready version of this paper, the problem has not been fixed, so currently we must consider this result as unproven.

Our result is also stronger in two other aspects: Firstly, we provide strong tail bounds on the probability $\mathbb{P}(d_G(x, y) \leq k)$. These yield a *shape theorem* for the geometric shape of the k -neighbourhood of a fixed vertex for growing k , in the case $\alpha < \tau - 2$ when our lower bound matches the upper bound. The shape theorem also implies that the size of the k -neighbourhood of a fixed vertex grows as $e^{k^{1/\Delta \pm o(1)}}$ as $k \rightarrow \infty$, which was not known before. Secondly, due to its simplicity, we believe that our method is also potentially easier to generalize. As demonstration, we show that a similar lower bound holds not only for graph distances, but also for first passage percolation.

In the following, we will start by formally defining the graph models and stating our precise results. The heart of the paper is Section 2, where we prove our result for the SFP model. In Section 3, we extend the proof to first-passage percolation. Finally, in Appendix A we explain why the proof of the upper bound claimed in [19] is incorrect.

1.1 Preliminaries and Random Graph Models

Our results hold for *Scale Free Percolation* (SFP), *Geometric Inhomogeneous Random Graphs* (GIRG), and *Hyperbolic Random Graphs* (HRG). We will define LRP and GIRG formally in this section. HRG has been shown to be a special case of GIRG [11], so all results proven for GIRG automatically also hold for HRG, and we do not need to formally define HRG. A formal definition of HRG together with its connection to GIRG can be found in [11].

1.1.1 Scale Free Percolation (SFP)

For SFP, we start with the infinite⁴ d -dimensional grid, which is our set of vertices. For two points $x, y \in \mathbb{Z}^d$, we define their distance $|x - y|$ via the usual Euclidean norm. Moreover, each vertex draws a *weight* w_x independently identically distributed from a power-law distribution. For our purposes, this is a Pareto distribution satisfying

$$\mathbb{P}(w_x \geq z) = z^{1-\tau}$$

for $z \geq 1$, where the parameter $\tau > 1$ is the *power-law exponent*.

We add edges in two different ways. First, we place the usual *grid edges*⁵ between points that are adjacent in \mathbb{Z}^d , i.e., we place an edge between $x = (x_1, \dots, x_d)$ and $y = (y_1, \dots, y_d)$ if there is some coordinate $1 \leq i \leq d$ such that $|x_i - y_i| = 1$ and $x_j = y_j$ for all $j \neq i$. Second, we randomly create long-range edges, also called *weak ties*, by placing an edge between $x, y \in \mathbb{Z}^d$, independently for different sets $\{x, y\}$, with probability p_{xy} , which is defined as⁶

$$p_{xy}^{(\text{SFP})} := \min \left\{ 1, \lambda \left(\frac{w_x w_y}{|x - y|^d} \right)^\alpha \right\}, \quad (1)$$

where $\lambda > 0, \alpha > 1$ are constants. We write $x \sim y$ if there is an edge between $x, y \in \mathbb{Z}^d$ and $x \approx y$ otherwise.

⁴ Traditionally, LRP is defined on an infinite vertex set while GIRG is defined on a finite vertex set, following the tradition of mathematics for LRP and of computer science for GIRG. However, both models can be defined in either a finite or an infinite version, and this does not affect graph distances, see [25] for details.

⁵ Some variants of SFP do not include grid edges. Since we prove lower bounds on graph distances, we make our result stronger by including grid edges.

⁶ In the literature, the connection probability is often defined as $1 - \exp(-\lambda|x - y|^{-\alpha d})$. We remark that this differs from our connection probability at most by a constant factor, which does not change the model substantially.

We remark that the parameterization is slightly different from the original one used in [14] and also in [19]. Compared to our formulation, they used auxiliary weights $w'_x := w_x^\alpha$, which were then drawn from a power-law distribution with exponent $\tau' = 1 + d(\tau - 1)/\alpha$. These two parameterizations are equivalent, but our formulation has the advantage that the weights correspond, up to constant factors, to the expected degree of the vertices, $\mathbb{E}[\deg(x) \mid w_x] = \Theta(w_x)$. Our formulation also saves an internal parameter of the model. Moreover, we rescaled α by a factor d to match it with the parameterization of GIRG. When comparing our results with the lower bounds in [19], the following transformations are needed, where the subscript “[19]” indicates notation from that paper, and parameters without subscript are from our paper.

$$\alpha_{[19]} = \alpha d \quad \text{and} \quad \gamma_{[19]} = \tau - 1. \quad (2)$$

The internal parameter $\tau_{[19]}$ is superfluous and does not have a correspondence in our paper.

1.1.2 Geometric Inhomogeneous Random Graphs (GIRG)

The GIRG model is sometimes also referred to as Continuum Scale Free Percolation [16]. The main difference between SFP and GIRG is that in GIRG, the positions of vertices are randomly chosen. Namely, for some large enough $n \in \mathbb{N}$, we consider a cube \mathcal{X} of volume n in \mathbb{R}^d , where d is a constant. Our vertex set V then consists of n vertices, where the position ξ_x of each vertex x is picked independently at random from the uniform distribution over \mathcal{X} . The distance between two vertices $x, y \in V$ is defined as the Euclidean norm $|\xi_x - \xi_y|$, as in SFP. Again similarly to SFP, each vertex x draws a weight w_x independently from the Pareto distribution satisfying

$$\mathbb{P}(w_x \geq z) = z^{1-\tau}$$

for $z \geq 1$, with $\tau > 1$. Finally, two different vertices x and y are connected by an edge with probability⁷

$$p_{xy}^{(\text{GIRG})} = \Theta \left(\min \left\{ 1, \left(\frac{w_x w_y}{|\xi_x - \xi_y|^d} \right)^\alpha \right\} \right), \quad (3)$$

where the hidden constants are uniform over all x, y . Formally, we require that there are two absolute constants $c_{\text{low}}, c_{\text{upp}} > 0$ independent of n such that for all n and any two different vertices $x, y \in V$, conditional on their weights $w_x, w_y \geq 1$ and positions $\xi_x, \xi_y \in \mathcal{X}$,

$$c_{\text{low}} \min \left\{ 1, \left(\frac{w_x w_y}{|\xi_x - \xi_y|^d} \right)^\alpha \right\} \leq p_{xy}^{(\text{GIRG})} \leq c_{\text{upp}} \min \left\{ 1, \left(\frac{w_x w_y}{|\xi_x - \xi_y|^d} \right)^\alpha \right\}.$$

The reason for allowing constant factor deviations is that then hyperbolic random graphs (HRG) is a special case of GIRG with $d = 1$, where the Euclidean distance is replaced by the angular distance in hyperbolic space [11]. Hence, any statement proven for this version of GIRG also holds for HRG.

Note that the constants τ and α have an analogous role here as in SFP. A notable difference to SFP is that in GIRG there are no grid edges (since positions are no longer on the grid), but that does not significantly affect our results.

⁷ The original paper [11] used a geometry that was rescaled by a factor $n^{1/d}$. I.e., they used a cube of volume one and had an additional factor n in the denominator of (3). Both variants are equivalent, but our scaling aligns better with the SFP scaling and allows GIRGs to be extended to infinite graphs if desired [25].

1.1.3 Long Range Percolation (LRP)

To put our results in an adequate context, we sometimes reference the model of *Long Range Percolation* (LRP) which is a special case and predecessor of SFP, cf. e.g. [3, 5]. We define LRP completely analogously to SFP with the only modification that every vertex has a (deterministic) weight of 1.

1.1.4 First Passage Percolation (FPP)

Our lower bounds on graph distances are (in a similar form) also applicable to *First Passage Percolation* (FPP) on the graph models SFP, GIRG, HRG, and also LRP. Here, each edge e of the graph draws a random length or cost c_e independently identically distributed (i.i.d.) from a distribution over the non-negative reals. For two vertices x, y we are then interested in the *minimal/infimal cost* of all paths from x to y . This can be done on an arbitrary finite or infinite underlying graph. In this work we restrict our attention to the case where each edge cost is sampled from an exponential distribution with rate 1. The formal definition is as follows.

► **Definition 1** (First Passage Percolation (FPP)). We call the following process *First Passage Percolation*, in short FPP. Given a graph $G = (V, E)$, assign to each edge e an i.i.d. cost c_e sampled from an exponential distribution with rate 1. For a finite path π , we define the *cost* of that path as

$$c(\pi) = \sum_{e \in \pi} c_e$$

Then, the *cost-distance* or *first passage time* between two vertices x and y is the minimum (or infimum) cost of any finite path connecting x and y , i.e.

$$d_G^{\text{cost}}(x, y) := \inf\{c(\pi) : \pi \in \mathcal{P}_{x,y}\} \quad \text{for } x, y \in V,$$

where $\mathcal{P}_{x,y}$ is the set of all finite paths between x and y .

Recently, FPP on SFP, GIRG, and HRG was studied [13, 23–25]. In particular, it was shown in [23, 25] that FPP on those graphs exhibits an *explosive* behaviour if the vertex weights have infinite variance (i.e. if $\tau < 3$). This means that the cost-distance between two vertices x, y converges in distribution against a random variable that is finite almost surely. In the infinite model SFP, this means that there are infinitely many vertices reachable within finite cost from a given vertex; in the finite models GIRG and HRG, a constant fraction of all vertices have cost-distance $\mathcal{O}(1)$. We emphasize that this is not true for *graph distances* for $\tau \in (2, 3)$ since then degrees are finite almost surely, so the number of vertices in graph distance C is finite/constant for any constant $C > 0$.

We remark that a model with a similar name, *long-range first passage percolation* was introduced and studied in [13]. This is *not* FPP on LRP, but a different model with the complete graph on \mathbb{Z}^d (all edges are present, degrees are infinite), where transmission times are penalized for edges between vertices of large Euclidean distance. Despite the differences, the models are related, and our analysis of FPP uses a coupling to a similar model and is inspired by [13].

1.1.5 Terminology

We are interested in the typical *graph distance* $d_G(x, y)$ between two vertices x, y in SFP, GIRG, and HRG, and the *cost-distance* $d_G^{\text{cost}}(x, y)$ in FPP.⁸ We are interested in the asymptotic behavior of $d_G(x, y)$ and $d_G^{\text{cost}}(x, y)$ in terms of the Euclidean distance $|x - y|$. That is, we study how $d_G(x, y)$ and $d_G^{\text{cost}}(x, y)$ scale as functions of $|x - y|$ when $|x - y| \rightarrow \infty$. We define the function $\Delta(\beta) := \frac{1}{\log_2(2/\beta)}$, which will appear in the exponent governing the polylogarithmic behavior of graph distances. Throughout, we use $\Delta = \Delta(\alpha)$ where α is the long-range parameter of the relevant model.

Previous work showed that both LRP and SFP exhibit multiple phase transitions in the asymptotic behavior of $d_G(x, y)$ depending on the model parameters, which were briefly summarized in the introduction and are further summarized (together with our results) in Table 1.

■ **Table 1** Upper and lower bounds for graph distances and cost-distances in long-range percolation (LRP), scale-free percolation (SFP) and first passage percolation (FPP) on SFP, with $\Delta(\beta) = 1/\log_2(2/\beta)$. The results on graph distance for SFP also hold for geometric inhomogeneous random graphs (GIRG) and hyperbolic random graphs (HRG). Our results are indicated in bold.

Model	$\alpha \in (1, 2)$		$\alpha > 2$	
LRP	$\Theta(\log(x - y)^{\Delta(\alpha)})$ [3, 5]		$\Theta(x - y)$ [2]	
SFP	$\Theta(\log \log(x - y))$ [14]	$\tau \in (2, 3)$	$\tau > 3$	$\Theta(\log \log(x - y))$ [14]
		$\leq \log(x - y)^{\Delta(\alpha)+o(1)}$ [3, 5] ⁹ $\geq \log(x - y)^{\Delta(\min\{\alpha, \tau-2\})-o(1)}$ Corollary 2	$\Theta(x - y)$ [2]	
FPP on SFP	$\Theta(1)$ [23]	$\leq \log(x - y)^{\Delta(\alpha)+o(1)}$ [23] $\geq \log(x - y)^{\Delta(\min\{\alpha, \frac{\tau-1}{2}\})-o(1)}$ Corollary 3	$\Theta(1)$ [23]	?

1.2 Our results

In this section, we formally state our main results. To keep the exposition simple, we state them only for SFP, not for GIRG and HRG. All results on graph distances in this section also hold for GIRG and HRG, where all constants can be chosen independently of the number n of vertices. We give details on that in the full version but omit them here due to space constraints.

We prove stronger lower bounds for graph distances in the logarithmic regimes of SFP with a much simpler proof than in [19]. Furthermore, we also obtain a *shape theorem* which sandwiches the k -neighbourhood of a given vertex between two geometric balls of similar size. The key result is a general upper bound on the probability that two vertices have graph/cost-distance at most k . For the following claims, recall that $\Delta(\beta) = 1/\log_2(2/\beta)$.

► **Theorem 1.1** (Tail Bound for Graph Distances in SFP). Consider SFP with parameters $\alpha \in (1, 2), \tau > 3$ and λ . Fix any sufficiently small $\varepsilon > 0$ and let $\Delta' = \Delta(\min\{\alpha, \tau - 2 - \varepsilon\})$. Then, there exist constants c_1, c_2, β depending on the model parameters as well as ε such that for any pair of vertices $x, y \in \mathbb{Z}^d$ and any $k \in \mathbb{N}$, we have

$$\mathbb{P}(d_G(x, y) \leq k) \leq c_2^{-1} |x - y|^{-\alpha d} (k + 1)^{-\beta} \exp\left(c_1 k^{1/\Delta'}\right).$$

⁸ Formally, the graph distance is also defined in Definition 1 by setting $c(e) := 1$ for all edges e .

⁹ In [19], the authors claim to prove an improved logarithmic upper bound with exponent $\Delta(\min\{\alpha, \tau - 2\})$, which would match our lower bounds. However, we show that their proof is wrong, see Appendix A.

The above bound is proved by induction over k in Section 2. We express the probability that a path of length $\leq k$ exists recursively by decomposing the path into two parts connected by an edge which has the longest geometric distance on the original path. Applying the inductive hypothesis and integrating over all possible endpoints and weights of the endpoints of said edge then yields the desired bound. However, this actually only works if $\alpha < \tau - 2$, which is needed to ensure convergence of some involved integrals. We remedy this and make the theorem applicable also to the case $\alpha \geq \tau - 2$ by using a coupling argument. Here (see Lemma 5 for the exact statement), we argue that decreasing α only makes the model denser, so graph distances can only become smaller. Thus, to prove the claim for some $\alpha \geq \tau - 2$, we can decrease α to some value $\alpha' := \tau - 2 - \varepsilon$ without increasing distances, and then apply Theorem 1.1 for the already settled case α' . This is the reason why we use the exponent $\Delta(\min\{\alpha, \tau - 2 - \varepsilon\})$. Using the above tail bound then directly implies a bound on typical graph distances.

► **Corollary 2** (Typical Graph Distances in SFP). *Consider SFP under the assumption $\alpha \in (1, 2)$ and $\tau > 3$. Then for every sufficiently small $\varepsilon > 0$, there is a constant $c > 0$ such that*

$$\lim_{|x-y| \rightarrow \infty} \mathbb{P} \left(d_G(x, y) \geq c \log(|x - y|)^{\Delta(\min\{\alpha, \tau - 2 - \varepsilon\})} \right) = 1.$$

Note that in the case $\alpha < \tau - 2$, the exponent is exactly $\Delta(\alpha)$ (if ε is chosen sufficiently small), which is a slightly stronger result than the one we obtain if $\alpha \geq \tau - 2$ and matches the known upper bounds in [5] up to only a constant factor in front of the log. The previously best lower bounds by Hao and Heydenreich [19] only matched these upper bounds if $2\alpha < \tau - 2$ and nonetheless were only tight if we ignore an additional additive constant of $-\varepsilon$ in the exponent. If $\alpha \geq \tau - 2$, we also have to account for such an ε in the exponent, but even in this case, our result strengthens the lower bounds in [19] and at the same time relies on a significantly simpler proof. The original proof heavily relied on so-called *hierarchies* as introduced in [3] and required complex combinatorial estimates for showing that certain structures w.h.p. do not exist. We avoid this by using the inductive proof strategy as described above instead. This not only simplifies and improves the existing lower bounds on graph distances, but the tail bound in Theorem 1.1 further yields a so-called *shape theorem* precisely characterizing the diameter and the cardinality of the k -neighbourhood of a vertex x , i.e., the set of vertices in graph distance at most k from x . To this end, we define $\mathcal{B}(x, k) := \{y \in V \mid d_G(x, y) \leq k\}$ as the set of all k -hop neighbors of a given vertex x .

► **Theorem 1.2** (Shape Theorem for k -Balls in SFP). *Consider SFP with $\alpha < \tau - 2$. Let $\Delta = \Delta(\alpha)$, fix an $\varepsilon > 0$ and let $\mathcal{X}_{low}, \mathcal{X}_{upp}$ be the set of vertices at a geometric distance of at most $q(k) = e^{k^{1/\Delta - \varepsilon}}$ and at most $r(k) = e^{k^{1/\Delta + \varepsilon}}$ from a fixed vertex x , respectively. Then, we have*

$$\lim_{k \rightarrow \infty} \mathbb{P}(\mathcal{X}_{low} \subseteq \mathcal{B}(x, k) \subseteq \mathcal{X}_{upp}) = 1. \quad (\text{A})$$

In particular,

$$\lim_{k \rightarrow \infty} \mathbb{P} \left(e^{k^{1/\Delta - \varepsilon}} \leq |\mathcal{B}(x, k)| \leq e^{k^{1/\Delta + \varepsilon}} \right) = 1. \quad (\text{B})$$

The lower bound of the theorem comes from [4], while we contribute the upper bound, which is a relatively straightforward corollary of our tail bound in Theorem 1.1. For this reason, we defer the formal proof to the full version of this paper. Note that lower bounds on graph distances correspond to upper bounds for $\mathcal{B}(x, k)$ and vice versa.

We further note that while our upper bound in Theorem 1.2 also holds for GIRG and HRG, the lower bound is not known for these models and can only hold with some caveats. Firstly, due to the lacking grid edges those graphs are not connected, and require additional constraints to ensure the existence of a giant (linear-size) connected component. Second, even if the giant component exists, a constant fraction of vertices are not in the largest component. Hence, the lower bounds in (A) and (B) can only hold conditioned on x being in the largest component, and for (A) we must intersect \mathcal{X}_{low} with the giant component. We conjecture that the lower bounds in Theorem 1.2 hold with these caveats, but this is not known.

1.3 First passage percolation on SFP

Using similar techniques and inspired from those in [13], we can prove similar statements for FPP on SFP. Analogous to Theorem 1.1, we obtain a tail bound on the probability that x, y have cost-distance at most t . We remark that the same result could be obtained analogously for GIRG (and thus HRG) as well, but we omit this for conciseness.

► **Theorem 1.3** (Tail Bound for Cost-Distances for FPP on SFP). Consider FPP on SFP and arbitrary vertices x, y . Fix any sufficiently small $\varepsilon > 0$. There exists a constant c depending only on α, ε and τ such that for $\Delta'' = \Delta(\min\{\alpha, \frac{\tau-1}{2}\} - \varepsilon)$,

$$\mathbb{P}(d_G^{\text{cost}}(x, y) \leq t) \leq |x - y|^{-\alpha d} \exp\left(ct^{1/\Delta''}\right).$$

The proof of this differs from Theorem 1.1 in some significant aspects which are formally presented in Section 3. Intuitively, the main differences are as follows. Firstly, we now have two sources of randomness: the existence of edges and the cost of an existing edge. The first step towards proving Theorem 1.3 is therefore to combine these two sources into a single one. This is achieved by coupling the model to a related model called *Complete Scale Free First Passage Percolation* or CFFP for short. Here, all possible edges on the vertex set \mathbb{Z}^d exist a priori but the cost of the edge between $x, y \in \mathbb{Z}^d$ is now drawn from an exponential distribution with rate $w_x^\alpha w_y^\alpha |x - y|^{-\alpha d}$ instead of rate 1, i.e., in CFFP the rate of an edge depends on the vertex weights and (geometric) distances of its endpoints.

The second main difference to the proof of Theorem 1.1 is that cost-distances are continuous random variables, so we cannot union-bound over all possible cost-distances before and after the longest edge of a potential path anymore like we did for SFP (in Lemma 4). Instead, we establish a continuous analog, a so called *self-bounding inequality* that relates the expected size of a k -ball to itself recursively. Another difficulty one has to overcome is that, in principle, paths of low cost-distance do not necessarily have to correspond to low graph distance as well. It could theoretically happen that many edges have very low cost and we get a low cost path which uses many edges. In such cases, we cannot use the existence of a geometrically long edge in the path, which is very central to our proof for graph distances. However, we are able to show that paths with high graph distance are actually very unlikely to have low cost-distance (see Lemma 8). Finally, a further obstacle in adapting the proof is that we have to work with the probability that a path of a certain cost exists conditioned on the weights of its endpoints at multiple points. This impacts the probabilities of edges/paths existing and thus introduces complications. To overcome this, we relate said probabilities conditional on the involved weights to their unconditional versions by employing a coupling (Proposition 12).

As a corollary of the above tail bound, we obtain lower bounds on the typical cost-distance similar to the one established for SFP (Corollary 2).

► **Corollary 3** (Typical Graph Distances in FPP on SFP). *Consider FPP on SFP under the assumption $\alpha \in (1, 2)$ and $\tau > 3$. Then, for every sufficiently small $\varepsilon > 0$*

$$\lim_{|x-y| \rightarrow \infty} \mathbb{P} \left(d_G^{\text{cost}}(x, y) \geq \log(|x-y|)^{\Delta(\min\{\alpha, \frac{\tau-1}{2}\}-\varepsilon)} \right) = 1.$$

1.3.1 Asymptotics and Probability Theory

We use standard Landau notation for indicating the asymptotic growth of a function. All asymptotic statements refer to the asymptotic behavior of a function as the distance $|x-y|$ tends to infinity, unless explicitly noted otherwise (like for the shape theorem, where we consider $k \rightarrow \infty$.) We further require a version of the Van den Berg-Kesten inequality (or BK-inequality) from [28], which allows us to bound the probability that there exist *disjoint* subpaths connecting a vertex x to u and a vertex v to y by the product of the probabilities of either path existing, as if they were independent. We refer the interested reader to the full version of this paper for further details.

2 Lower Bounds for Graph Distances in SFP

In this section, we provide the proof of our main lower bound. Our proof generally follows the structure of the proof of Theorem 3.1 in [4]. Our goal is to show that the logarithmic exponent in the distances is at least roughly $\Delta(\min\{\alpha, \tau - 2\})$. When $\alpha < \tau - 2$, this is the same as $\Delta = \Delta(\alpha)$. We will first give the proof under this condition, so we first show Lemma 4.

► **Lemma 4** (Tail Bound for Graph Distances in SFP). *Consider SFP with $\alpha \in (1, 2)$ such that $\alpha < \tau - 2$ and $\Delta = \Delta(\alpha)$. There exist constants c_1, c_2, β depending only on the model parameters, such that for any pair $x, y \in \mathbb{Z}^d$ and any $k \in \mathbb{N}$,*

$$\mathbb{P}(d_G(x, y) \leq k \mid w_x, w_y) \leq w_x^\alpha w_y^\alpha c_2^{-1} |x-y|^{-\alpha d} (k+1)^{-\beta} e^{c_1 k^{1/\Delta}}. \quad (4)$$

Proof. First of all, note that for fixed β, c_2 , and λ , the base case ($k = 1$) is true for c_1 large enough. That is because the RHS of 4 is $2^{-\beta} e^{c_1} w_x^\alpha w_y^\alpha c_2^{-1} |x-y|^{-\alpha d}$ and the actual connection probability is at most $\lambda w_x^\alpha w_y^\alpha |x-y|^{-\alpha d}$. For the inductive step, let h be the RHS of our induction hypothesis, i.e.,

$$h(r, k, w_x, w_y) := w_x^\alpha w_y^\alpha c_2^{-1} r^{-\alpha d} (k+1)^{-\beta} e^{c_1 k^{1/\Delta}}.$$

Assume that the induction hypothesis is true up to $k-1$. For x, y to be connected with at most k steps, an edge must be used with geometric distance at least $\frac{|x-y|}{k}$. This could either be the first or last edge on the path, or a so-called *internal* edge. Let us first bound the probability corresponding to this edge being internal. For this, we union bound over all possible endpoints u and v of said edge. Actually, we integrate, since constant factors are essentially immaterial for the proof. At a given distance r , there are at most $c_d r^{d-1}$ vertices, for some constant c_d . Let w_u, w_v be the weights of u, v , respectively. By the BK inequality, we have

$$\begin{aligned} & \mathbb{P}(d_G(x, y) \leq k, \text{longest edge is internal} \mid w_u, w_v, w_x, w_y) \leq \\ & \sum_{i=1}^{k-2} \mathbb{P}(d_G(x, u) \leq i \mid w_x, w_u) \mathbb{P}(u \sim v \mid w_u, w_v) \mathbb{P}(d_G(v, y) \leq k-i \mid w_v, w_y) \leq \lambda w_u^\alpha w_v^\alpha \left(\frac{|x-y|}{k}\right)^{-\alpha d} \\ & \times \underbrace{\sum_{i=1}^{k-2} \left(\int_1^\infty c_d r_u^{d-1} \min\{1, h(r_u, i, w_x, w_u)\} dr_u \right)}_{\mathcal{I}(u, i)} \underbrace{\left(\int_1^\infty c_d r_v^{d-1} \min\{1, h(r_v, k-i, w_v, w_y)\} dr_v \right)}_{\mathcal{I}(v, k-i)}. \end{aligned}$$

We will now argue that there exists a constant C_{int} (depending on the model parameters) such that

$$\begin{aligned} \mathcal{I}(u, i) & \leq C_{\text{int}} \left((i+1)^{-\frac{\beta}{\alpha}} e^{\frac{c_1}{\alpha} i^{1/\Delta}} w_u w_x c_2^{-\frac{1}{\alpha}} \right), \\ \mathcal{I}(v, k-i) & \leq C_{\text{int}} \left((k-i+1)^{-\frac{\beta}{\alpha}} e^{\frac{c_1}{\alpha} (k-i)^{1/\Delta}} w_v w_y c_2^{-\frac{1}{\alpha}} \right). \end{aligned}$$

To this end, note that that there exists a value

$$\hat{r}_u = (w_x w_u)^{\frac{1}{d}} c_2^{-\frac{1}{\alpha d}} (i+1)^{-\frac{\beta}{\alpha d}} e^{\frac{c_1}{\alpha d} (i+1)^{1/\Delta}}$$

for r_u below which the minimum inside the integral $\mathcal{I}(u, i)$ is 1. We can thus express $\mathcal{I}(u, i) = \int_1^{\hat{r}_u} f_1(r_u) dr_u + \int_{\hat{r}_u}^\infty f_2(r_u) dr_u$ where $f_1(r) = c_d r^{d-1}$ and f_2 is a polynomial in r_u with exponent smaller than -1 . Therefore, the entire integral is dominated by the value of the antiderivative of f_1 and f_2 at the splitting point \hat{r}_u . Since f_1, f_2 are polynomials, the antiderivative of f_1 is $\leq c r_u f_1(r_u)$ and the antiderivative of f_2 is $\leq c r_u f_2(r_u)$ for some constant c . Since the minimum is a continuous function, we have $f_1(\hat{r}_u) = f_2(\hat{r}_u)$ and thus, $\mathcal{I}(u, i) = \Theta(\hat{r}_u f_1(\hat{r}_u)) = \Theta(\hat{r}_u^d)$ as claimed. A similar argument holds for $\mathcal{I}(v, k-i)$.¹⁰

Plugging this in, we obtain

$$\begin{aligned} & \mathbb{P}(d_G(x, y) \leq k, \text{longest edge is internal} \mid w_u, w_v, w_x, w_y) \\ & \leq C_{\text{int}}^2 \lambda \cdot w_u^{1+\alpha} w_v^{1+\alpha} w_x^\alpha w_y^\alpha |x-y|^{-\alpha d} c_2^{-\frac{2}{\alpha}} \\ & \quad \times \underbrace{k^{\alpha d} \sum_{i=1}^{k-2} \left((i+1)^{-\frac{\beta}{\alpha}} e^{\frac{c_1}{\alpha} i^{1/\Delta}} \right) \left((k-i+1)^{-\frac{\beta}{\alpha}} e^{\frac{c_1}{\alpha} (k-i)^{1/\Delta}} \right)}_{:=S}. \end{aligned}$$

Our goal now is to show that the above term is at most $h(r, k, w_x, w_y)$. To this end, we show that

$$S \leq (k+1)^{-\beta} e^{c_1 k^{1/\Delta}} \tag{5}$$

for k and β large enough. For this, notice that for small or large i , the exponential terms in S are still quite ‘‘tame’’, due to the $\frac{1}{\alpha}$ factor. When i is around $\frac{k}{2}$, their product (which is maximized for such i due to concavity) is practically

$$\exp\left(\frac{2c_1}{\alpha} \left(\frac{k}{2}\right)^{1/\Delta}\right) = \exp\left(\frac{2c_1}{\alpha} 2^{-1/\Delta} k^{1/\Delta}\right) = \exp\left(c_1 k^{1/\Delta}\right)$$

¹⁰This observation is helpful whenever we integrate a continuous and piecewise polynomial function.

as $2^{-1/\Delta} = \alpha/2$ by definition of Δ . In fact, for $i = \frac{k}{2}$, we have actual equality and for every $1 \leq i \leq k$, we have

$$\exp\left(\frac{c_1}{\alpha}i^{1/\Delta}\right)\exp\left(\frac{c_1}{\alpha}(k-i)^{1/\Delta}\right) \leq \exp\left(c_1k^{1/\Delta}\right).$$

This is precisely the exponential term that appears in the statement we want to prove (5). However, we also need to account for the sum and the terms polynomial in k that appear in S . To this end, we use that – if $i \approx k/2$ – we gain from the product of the polynomial terms in the sum to compensate overheads. On the other hand, if i is large or small, the product of the exponential terms is much smaller than what we need, so we can compensate the other terms by using the arising gap. With this in mind, we split the sum in S into the cases where $|i - \frac{k}{2}| \leq \frac{k}{4}$ and those where this is not true. This way, we obtain

$$S \leq k^{1+\alpha d}e^{(1-\gamma)c_1k^{1/\Delta}} + k^{1+\alpha d}((k+1)/8)^{-\frac{2}{\alpha}\beta}e^{c_1k^{1/\Delta}}, \tag{6}$$

where $\gamma > 0$ is a constant depending on Δ (and therefore on α). The first term accounts for cases where i is sufficiently far from $\frac{k}{2}$, making the exponential terms merge in a tame way. When $i \in [\frac{k}{4}, \frac{3k}{4}]$, both $i+1$ and $k-i+1$ are at least $\frac{k+1}{8}$, since $k > 2$ (recall that we are analyzing the case where k edges allow for an internal edge), and this is how the other term is obtained.

Now, notice that since $\alpha < 2$, we can choose β large enough such that $(\alpha d + 1) - \frac{2\beta}{\alpha} < -\beta$. Then, the second term in 6 is at most $(k+1)^{-\beta}e^{c_1k^{1/\Delta}}$ as desired. For the first term, we notice that the same holds if k is large enough. Hence, for all k , S is at most some constant C times $(k+1)^{-\beta}e^{c_1k^{1/\Delta}}$. We use this to conclude that

$$\begin{aligned} \mathbb{P}(d_G(x, y) \leq k, \text{ longest edge is internal} \mid w_u, w_v, w_x, w_y) \\ \leq C C_{\text{int}}^2 \lambda c_2^{1-\frac{2}{\alpha}} w_u^{\alpha+1} w_v^{\alpha+1} w_x^\alpha w_y^\alpha |x-y|^{-\alpha d} (k+1)^{-\beta} e^{c_1k^{1/\Delta}} c_2^{-1} \\ = C C_{\text{int}}^2 \lambda c_2^{1-\frac{2}{\alpha}} w_u^{\alpha+1} w_v^{\alpha+1} \cdot h(|x-y|, k, w_x, w_y). \end{aligned}$$

Since we assume that $\alpha < \tau - 2$, we can integrate w_u, w_v out such that the corresponding integrals over w_u and w_v converge and only obtain another constant factor overhead. Then, we can choose c_2 large enough to compensate these constant overheads. Notice that this works since we have a factor of $c_2^{1-\frac{2}{\alpha}}$ where the exponent is negative because $\alpha < 2$. In total, we have shown that we can choose the constants β, c_1 and c_2 such that the above bound is at most $\frac{1}{3}h(|x-y|, k, w_x, w_y)$ for all k .

Now, let us also bound the probability of paths in which the longest edge is adjacent to either x or y . To this end, we sum over all possible vertices z connected to x by an edge of (geometric) length $\geq |x-y|/k$. Again, by the BK inequality, we have

$$\begin{aligned} \mathbb{P}(d_G(x, y) \leq k, \text{ longest edge incident to } x \mid w_x, w_z, w_y) \\ \leq \lambda w_x^\alpha w_z^\alpha \left(\frac{|x-y|}{k}\right)^{-\alpha d} \left(\int_1^\infty c_d r^{d-1} \min\{1, h(r, k-1, w_z, w_y)\} dr\right) \\ \leq C \lambda w_x^\alpha w_z^\alpha \left(\frac{|x-y|}{k}\right)^{-\alpha d} k^{-\frac{\beta}{\alpha}} e^{\frac{c_1}{\alpha}k^{1/\Delta}} w_z w_y c_2^{-\frac{1}{\alpha}} \\ \leq C \lambda w_z^{\alpha+1} \cdot (k+1)^\beta k^{\alpha d - \frac{\beta}{\alpha}} e^{c_1(\frac{1}{\alpha}-1)k^{1/\Delta}} c_2^{1-\frac{1}{\alpha}} \cdot h(|x-y|, k, w_x, w_y). \end{aligned}$$

Again, by integrating out w_z , we get another constant factor. We can now choose c_1 large enough so that the term exponential in k (which has a negative exponent since $\frac{1}{\alpha} - 1 < 0$) swallows the polynomial and constant terms for every k , ensuring that the factor in front

of $h(|x - y|, k, w_x, w_y)$ is at most $\frac{1}{3}$, as desired. Finally, summing the three possibilities that the longest edge is internal, the first, or the last edge on the path yields that overall, $\mathbb{P}(d_G(x, y) \leq k \mid w_x, w_y) \leq h(|x - y|, k, w_x, w_y)$ and finishes the proof. \blacktriangleleft

As promised, we now deal with cases where $\alpha \geq \tau - 2$ by coupling SFP to SFP with larger α without decreasing distances using the following lemma.

► Lemma 5. *Let α and λ be the long-range and percolation parameters of some instance of SFP. Fix the weights of all vertices and let p_{uv} refer to the probability that two vertices u and v are connected by an edge. Fix some $\alpha' < \alpha$. Then, $p_{uv} \leq \min\{1, \lambda^{\alpha'/\alpha} w_u^{\alpha'} w_v^{\alpha'} |u - v|^{-d\alpha'}\}$. In particular, this means that the original SFP graph (with parameter α) is a subgraph of the one with parameters α' and $\lambda' = \lambda^{\alpha'/\alpha}$.*

Proof. We have

$$\begin{aligned} p_{uv} &\leq \min\{1, \lambda w_u^\alpha w_v^\alpha |u - v|^{-d\alpha}\} = \left(\min\{1, \lambda^{\frac{1}{\alpha}} w_u w_v |u - v|^{-d}\}\right)^\alpha \\ &\leq \left(\min\{1, \lambda^{\frac{1}{\alpha}} w_u w_v |u - v|^{-d}\}\right)^{\alpha'} = \min\{1, \lambda^{\frac{\alpha'}{\alpha}} w_u^{\alpha'} w_v^{\alpha'} |u - v|^{-d\alpha'}\}. \end{aligned} \quad \blacktriangleleft$$

The implication of Lemma 5 is that we can artificially ensure that $\alpha < \tau - 2$ by setting $\alpha' = \tau - 2 - \varepsilon$ for an arbitrarily small ε and $\lambda' = \lambda^{\alpha'/\alpha}$. This allows us to prove Theorem 1.1 by applying Lemma 4 to this model since here, graph distances only get shorter due to Lemma 5. We defer the proof to the full version, since it is only technical and the ideas in it are already presented.

3 First Passage Percolation (FPP)

In this section we study first passage percolation (FPP) on SFP. Recall that this means that we assign a cost to every edge which is drawn independently from an exponential distribution with rate 1. For conciseness, we restrict ourselves to SFP even though the same technique would also work for GIRGs/HRGs. Note that we obtain the LRP model from SFP by informally setting $\tau = \infty$. Formally, since SFP is an increasing model in τ in terms of stochastic domination, the edge set of SFP with any finite τ stochastically dominates the edge set of LRP. Hence, all lower bounds on cost-distances from SFP also transfer to LRP.¹¹ In the following, we assume for simplicity that $\lambda = 1$; this does not affect our results.

In contrast to plain SFP, in FPP we have an additional source of randomness since not only the existence of an edge is random but also its cost. To prove lower bounds on cost-distances, it is therefore simpler (and sufficient) to consider a model in which there is only one source of randomness for the edges. We call this model *Complete Scale Free First Passage Percolation*, or CFFP for short. Here, all edges exist a priori, i.e., the graph is fixed to be the complete graph with vertex set \mathbb{Z}^d . However, we now draw the cost of each edge by sampling from an exponential distribution with rate $w_u^\alpha w_v^\alpha |u - v|^{-\alpha d}$ (i.e. a rate that depends on the weights and geometric distance between the two endpoints) instead of rate 1.

We start by showing that FPP on SFP is dominated by CFFP, i.e., that cost-distances in CFFP can only become shorter as compared to FPP on SFP. To that end, we need the following lemma that will allow us to combine the randomness of two events occurring with probability $\min\{1, \alpha\}$ and $(1 - e^{-b})$, respectively into an event occurring with probability $1 - e^{-ab}$.

¹¹The same is true for graph-distances, but here the results for LRP were already known.

► **Lemma 6.** $\min\{1, a\}(1 - e^{-b}) \leq 1 - e^{-ab}$ for all $a, b \geq 0$.

Proof. If $a \geq 1$, then the inequality is easy to see, as $(1 - e^{-b}) \leq (1 - e^{-ab})$ in this case. So, let us assume that $a < 1$ from now on. Consider the function

$$f(b) = a(1 - e^{-b}) - (1 - e^{-ab}).$$

Note that it suffices to show that $f(b) \leq 0$ for all $b \geq 0$. We can see that $f(0) = 0$ and also

$$f'(b) = a(e^{-b} - e^{-ab}) \leq 0.$$

This shows that the function $f(b)$ is non-increasing and since $f(0) = 0$, we have $f(b) \leq 0$ for all $b \geq 0$. ◀

With this, we establish a coupling between FPP on SFP and CFFP such that cost-distances in CFFP are at most as large as cost-distances in FPP on SFP.

► **Lemma 7.** Let u, v be a pair of vertices in \mathbb{Z}^d . Let further $X_{(u,v)}$ be the cost of the edge $\{u, v\}$ in FPP on SFP if it exists, and $X_{(u,v)} = \infty$ if the edge does not exist, and let $Y_{(u,v)}$ be its cost in CFFP. Then for any $t \geq 0$,

$$\mathbb{P}(X_{(u,v)} \leq t) \leq \mathbb{P}(Y_{(u,v)} \leq t).$$

Proof. For the event on the LHS to be true, the edge $\{u, v\}$ must exist and then *independently* the cost must be drawn to be at most t . The probability for the first event is $\min\{1, (w_u w_v)^\alpha |u - v|^{-\alpha d}\}$ and the probability of the latter is $1 - e^{-t}$. For the event on the RHS, one simply needs that the cost sampled from an exponential distribution with rate $(w_u w_v)^\alpha |u - v|^{-\alpha d}$ is at most t and the probability of this is exactly $1 - e^{-(w_u w_v)^\alpha |u - v|^{-\alpha d} t}$. Lemma 6 finishes the proof. ◀

Lemma 7 shows that any lower bound shown for cost-distances in CFFP will also be true for FPP on SFP. To see more clearly why this is true, note that we can couple the models in the following way. First, we sample the weights for the vertices in exactly the same way for both models. Then, conditioned on these weights, the probability space is a product space over independent one-dimensional random variables (technically, one of them can be infinite in value, but this is not a problem for our purposes) for which the inequality in Lemma 7 holds. With this in mind, we continue by establishing the lower bound for cost-distances in CFFP. We will generally follow similar arguments as the ones presented in [13], which studies a model similar to CFFP but without vertex weights. To establish an upper bound on the probability that the cost-distance between two vertices is at most t , we need a bound on the probability that the sum of exponential random variables is at most t , which is provided in the following lemma, which in turn is an adaptation of Lemma 2.1 in [13].

► **Lemma 8.** Let X_1, X_2, \dots, X_k be *i.i.d.* exponential random variables such that the rate of X_i is $(w_i w_{i+1})^\alpha |u_i - u_{i+1}|^{-\alpha d}$, for some sequence of vertices u_j with corresponding weight w_j , with $1 \leq j \leq k + 1$. The w_i are drawn from a power law with exponent τ . Assume that $2\alpha < \tau - 1$. Then, there exists a $c > 0$ depending only on α, τ such that for all $t \geq 0$,

$$\mathbb{P}\left(\sum_{i=1}^k X_i \leq t\right) \leq \left(\frac{ect}{k}\right)^k \prod_{i=1}^k |u_i - u_{i+1}|^{-\alpha d},$$

where the above probability is taken over the randomness of the weights and the X_i values.

Proof. Note that each $X_i = \frac{Y_i}{(w_i w_{i+1})^\alpha |u_i - u_{i+1}|^{-\alpha d}}$, where Y_i is an exponential random variable with rate 1. Let us use $\lambda_i = (w_i w_{i+1})^\alpha |u_i - u_{i+1}|^{-\alpha d}$ from now on. By Markov's inequality, we have

$$\mathbb{P}\left(\sum_{i=1}^k X_i \leq t\right) = \mathbb{P}\left(\exp\left(-\theta \sum_{i=1}^k X_i\right) \geq e^{-\theta t}\right) \leq e^{\theta t} \mathbb{E}\left[\exp\left(-\theta \sum_{i=1}^k X_i\right)\right].$$

Now, let us bound the expectation above. Once one fixes the weights w_j , each X_i is independent from each other. Moreover, for Y_i with rate one, it holds that $\mathbb{E}[\exp(-\theta Y_i)] = \frac{1}{1+\theta} \leq \frac{1}{\theta}$ for $\theta > 0$. So, for a fixed realization w_1, w_2, \dots, w_{k+1} of the weights, we have:

$$\begin{aligned} \mathbb{E}\left[\exp\left(-\theta \sum_{i=1}^k X_i\right) \mid w_1, w_2, \dots, w_{k+1}\right] &= \prod_{i=1}^k \mathbb{E}\left[\exp\left(-\frac{\theta}{\lambda_i} Y_i\right)\right] \\ &\leq \prod_{i=1}^k \frac{\lambda_i}{\theta} \\ &\leq \theta^{-k} \prod_{i=1}^{k+1} (w_i)^{2\alpha} \prod_{i=1}^k |u_i - u_{i+1}|^{-\alpha d}. \end{aligned}$$

The weight terms are raised to 2α , since each weight w_i enters in (at most) two λ_j as w_i^α . Integrating the weights out, we see that since they are independent, one has

$$\mathbb{E}\left[\exp\left(-\theta \sum_{i=1}^k X_i\right)\right] \leq \left[\theta^{-k} \prod_{i=1}^k |u_i - u_{i+1}|^{-\alpha d}\right] \prod_{i=1}^{k+1} \mathbb{E}[(w_i)^{2\alpha}].$$

Now, since $2\alpha - \tau < -1$, the expectations inside the rightmost product are all at most some constant c' . Let c be e.g. equal to $(c')^2$ such that $c^k \geq (c')^{k+1}$. Collecting the above bounds, we have

$$\mathbb{P}\left(\sum_{i=1}^k X_i \leq t\right) \leq e^{\theta t} \left[\theta^{-k} \prod_{i=1}^k |u_i - u_{i+1}|^{-\alpha d}\right] c^k.$$

Setting $\theta = \frac{k}{t}$ shows the desired bound. \blacktriangleleft

With Lemma 8 at hand, we show that the expected size of the t -ball around the origin grows at most exponentially with t . We define this ball $\mathcal{B}(x, t)$ as the set of vertices reachable from vertex x with a path of cost-distance at most t . Exponential growth is not enough by itself for our goal of showing a polylogarithmic lower bound on the distances but is a crucial step in doing so. To do this, we modify the proof of Lemma 2.6 and Theorem 1.2 (ii) in [13]. In the following, we only consider the growth of $\mathcal{B}(0, t)$, i.e., the t -ball around the origin, but it is easy to see that (by translation invariance) the same statements hold if we replace the origin by any vertex x .

► Theorem 9 (Exponential Ball Growth). *Let $\mathcal{B}(0, t)$ denote the set of vertices reachable with a path of cost at most t from the origin in CFFP. If $2\alpha < \tau - 1$, we have for some C depending only on α and τ ,*

$$\mathbb{E}[|\mathcal{B}(0, t)|] \leq e^{Ct}.$$

Proof. We compute

$$\begin{aligned} \mathbb{E} [|\mathcal{B}(0, t)|] &= \sum_{u \in \mathbb{Z}^d} \mathbb{P} (d_G^{\text{cost}}(0, u) \leq t) \\ &\leq 1 + \sum_{\substack{u \in \mathbb{Z}^d \\ u \neq 0}} \sum_{\substack{k=1 \\ (0,u)\text{-path} \\ \text{of length } k}}^{\infty} \mathbb{P} (\pi \text{ has cost distance at most } t) \\ &\stackrel{\text{Lemma 8}}{\leq} 1 + \sum_{\substack{u \in \mathbb{Z}^d \\ u \neq 0}} \sum_{k=1}^{\infty} \left(\frac{ect}{k}\right)^k \sum_{\substack{(0,u)\text{-path} \\ (0=u_1, u_2, \dots, u_{k+1}=u)}} \left[\prod_{i=1}^k |u_i - u_{i+1}|^{-\alpha d} \right]. \end{aligned}$$

The constant c above is as in Lemma 8. The rightmost sum above can be bounded by $b^k |u|^{-\alpha d}$ for some b depending only on α . This is done by Lemma 2.5 (c) in [13] (the quantity bounded there is the above sum and is defined in equation (2.3) in the page previous to that of Lemma 2.5). With that in mind, we have

$$\mathbb{E} [|\mathcal{B}(0, t)|] \leq 1 + \left(\sum_{u \in \mathbb{Z}^d, u \neq 0} |u|^{-\alpha d} \right) \left(\sum_{k=1}^{\infty} \left(\frac{ecbt}{k}\right)^k \right).$$

Since $\alpha > 1$, the first sum above is bounded by a constant c_1 . Moreover, note that

$$\sum_{k=1}^{\infty} \left(\frac{ecbt}{k}\right)^k \leq \sum_{k=0}^{\infty} \frac{(ecbt)^k}{k!} - 1 = e^{ecbt} - 1.$$

One can choose C large enough so that $\mathbb{E} [|\mathcal{B}(0, t)|] \leq e^{Ct}$. To see why, note that we can freely assume $c_1 \geq 1$. Then, setting $C = ecbc_1$ suffices. That is because of the following. Let $f(x) = x^{c_1} + c_1(1 - x) - 1$. This function is decreasing from 0 to 1 and increasing afterwards. Moreover, both $f(0)$ and $f(1)$ are non-negative, hence it is non-negative for all $x \geq 0$. Setting $x = e^{ecbt}$ shows that for all $t \geq 0$,

$$\mathbb{E} [|\mathcal{B}(0, t)|] \leq 1 + c_1(e^{ecbt} - 1) \leq (e^{ecbt})^{c_1} = e^{Ct}. \quad \blacktriangleleft$$

In the following, we define

$$g(t) := \mathbb{E} [|\mathcal{B}(0, t)|]$$

and note that we have already shown that $g(t)$ grows at most exponentially. But we can do better and show that in fact it grows at most *stretched* exponentially, in particular roughly as $\exp(t^{\frac{1}{\Delta}})$. This intuitively corresponds to the cost-distances between two vertices u and v growing roughly as $(\log |u - v|)^{\Delta}$, and is then also used in proving the corresponding lower bound later.

To show this improved bound, we bound the crucial quantity

$$f(r, t) = \sup_{|u|=r} \mathbb{P} (d_G^{\text{cost}}(0, u) \leq t) \in [0, 1],$$

that is, the highest possible probability with which a vertex connects to the origin with cost at most t , given that it has geometric distance r . We only consider $r, t > 0$. One can show the following bound for $f(r, t)$.

► **Lemma 10** (Towards a Self-Bounding Inequality for $g(t)$). *Consider CFFP with $2\alpha < \tau - 1$. There exist constants $c_f, \delta > 0$ depending only on α and τ such that*

$$f(r, t) \leq c_f r^{-\alpha d} h(t), \text{ where } h(t) := t^{\alpha d} \int_0^t g(t - y)(g(y) - 1)dy + e^{-\delta t}.$$

This lemma can be seen as a generalization of the technique we use to prove Lemma 4: We bound the probability that two vertices at distance r are connected by a path of cost at most t by a term that is essentially $r^{-\alpha d}$ (which is roughly the probability that the longest edge in such a path exists) times $h(t)$ which integrates over all possible y such that said edge connects the $(t - y)$ -ball around 0 and the y -ball around u . Using this, we can then derive a *self-bounding inequality* for $g(t)$, which relates g recursively to itself such that we can derive an upper bound on g by solving said recursive relation using Theorem 13 which is identical to [13, Theorem 5.3]. We derive the self-bounding inequality by summing $f(r, t)$ over all vertices and thus express $g(t)$ as a function of $h(t)$, which – in turn – depends on g . We capture this in the following lemma.

► **Lemma 11 (Self-Bounding Inequality for $g(t)$).** *Consider CFFP with $2\alpha \leq \tau - 1$. There are constants c, δ such that for all $t \geq 0$*

$$g(t)^\alpha \leq c \left(t^{\alpha d} \int_0^t g(t - y)g(y)dy + 1 \right).$$

Proof. To derive the self-bounding inequality for g using Lemma 10, we upper bound g by an expression involving $f(r, t)$ and then upper bound $f(r, t)$ using Lemma 10. Specifically, we estimate the expected size of a t -ball by integrating over all vertices times the respective probability $f(r, t)$.

$$\begin{aligned} \mathbb{E}[\mathcal{B}(0, t)] = g(t) &\leq 1 + \int_1^\infty c_d r^{d-1} \min\{1, f(r, t)\} dr \\ &= 1 + \int_1^{(c_f h(t))^{\frac{1}{\alpha d}}} c_d r^{d-1} dr + \int_{(c_f h(t))^{\frac{1}{\alpha d}}}^\infty c_d c_f r^{d-1-\alpha d} h(t) dr \end{aligned}$$

since for $r > (c_f h(t))^{\frac{1}{\alpha d}}$, the minimum is smaller than 1 by definition of $f(r, t)$ from Lemma 10. Integrating out then yields,

$$g(t) \leq 1 + c' \left(h(t)^{\frac{1}{\alpha}} + h(t)^{\frac{1}{\alpha}-1} h(t) \right) \leq 1 + ((c'' h(t))^{\frac{1}{\alpha}})$$

for some constants c', c'' that depend on α, d and τ . Therefore, we infer that

$$(g(t) - 1)^\alpha \leq c'' h(t). \tag{7}$$

It can be shown that¹² $g(t)^\alpha \leq 2^{\alpha-1}(1 + (g(t) - 1)^\alpha)$. Chaining this inequality with (7) and replacing $h(t)$ above by its definition in Lemma 10 we get the claimed recursive inequality for $g(t)$. In more detail, we have

$$g(t)^\alpha \leq 2^{\alpha-1}(1 + (g(t) - 1)^\alpha) \leq 2^{\alpha-1}(1 + c'' h(t))$$

Since $h(t)$ is bounded away from zero and since $e^{-\delta t} \leq 1$, it follows that there exists a c such that the inequality claimed in the lemma statement holds for all t . ◀

It is through this inequality that a stronger bound on $g(t)$ can be derived. For this, we use Theorem 5.3 from [13] directly which we restate as Theorem 13. It claims (among more general things) that for a given function $g(t)$, if $1 \leq g(t) \leq e^{Ct}$ for some constant C (which we have already shown) and an inequality similar to 7 holds, then one *roughly* has $g(t) \leq e^{t^{1/\Delta}}$. Now that we have motivated Lemma 10, let us prove it.

¹²One simply needs to consider the function $f(x) = 2^{\alpha-1}(1 + (x - 1)^\alpha) - x^\alpha$ restricted to $x \geq 1$, which has a global minimum of 0 at $x = 2$.

Proof of Lemma 10. Fix a “target” vertex u . Let us first focus on paths from 0 to u which contain at least ρt edges (we assume that this is an integer for simplicity) for some constant ρ that will be determined later. If P_{long} is the probability that some such path has cost-distance less than t , then by a simple union bound we have

$$\begin{aligned} P_{\text{long}} &\leq \sum_{k=\rho t}^{\infty} \sum_{\substack{(0,u)\text{-path } \pi \\ \text{of length } k}} \mathbb{P}(\pi \text{ has cost-distance at most } t) \\ &\stackrel{\text{Lemma 8}}{\leq} e^{\theta t} \sum_{k=\rho t}^{\infty} \theta^{-k} c^k \sum_{\substack{(0,u)\text{-path } (u_1, u_2, \dots, u_{k+1}) \\ \text{with } u_1=0 \text{ and } u_{k+1}=u}} \prod_{i=1}^k |u_i - u_{i+1}|^{-\alpha d} \\ &\stackrel{\text{Lemma 2.5 (c) in [13]}}{\leq} e^{\theta t} |u|^{-\alpha d} \sum_{k=\rho t}^{\infty} \left(\frac{cb}{\theta}\right)^k. \end{aligned}$$

In the second line above, we used Lemma 8 but the final step where θ is set to some value is not carried out. Furthermore, the constant b that emerges in the third line is as in Lemma 2.5 (c) of [13]. Now, setting $\theta = \rho > ecb$, we have

$$P_{\text{long}} \leq |u|^{-\alpha d} e^{\rho t} \frac{\left(\frac{cb}{\rho}\right)^{\rho t}}{1 - \frac{cb}{\rho}} \leq |u|^{-\alpha d} \frac{e}{e-1} e^{-\rho t \log \frac{\rho}{ecb}}. \quad (8)$$

Now, let us turn our attention to paths that use at most ρt edges instead and let P_{short} denote the probability that such a path has cost-distance at most t . The idea here is to notice that a geometrically long edge (u_1, u_2) must be used (similarly as in the proof of Lemma 4). In particular, this edge has to cover a distance of at least $\frac{|u|}{\rho t}$, as there are at most ρt edges used to cover a distance of $|u|$. We adapt the argumentation of the proof of Lemma 5.1 in [13], which is essentially a union bound over all the possible intermediate pairs (u_1, u_2) . More precisely, we get that

$$\begin{aligned} P_{\text{short}} &\leq \sum_{\substack{u_1, u_2 \in \mathbb{Z}^d \\ |u_1 - u_2| \geq \frac{|u|}{\rho t}}} \iint \mathbb{P}(d_G^{\text{cost}}(0, u_1) + c_{(u_1, u_2)} + d_G^{\text{cost}}(u_2, u) \leq t \mid w_{u_1}, w_{u_2}) d\mu(w_{u_1}) d\mu(w_{u_2}) \end{aligned} \quad (9)$$

where $\mu(w) = w^{1-\tau}$ is the probability measure of the weight distribution and

$$\begin{aligned} &\mathbb{P}(d_G^{\text{cost}}(0, u_1) + c_{(u_1, u_2)} + d_G^{\text{cost}}(u_2, u) \leq t \mid w_{u_1}, w_{u_2}) \\ &\leq \int_0^t d\mathbb{P}(d_G^{\text{cost}}(0, u_1) \leq s \mid w_{u_1}) \int_0^{t-s} \mathbb{P}(d_G^{\text{cost}}(u_2, u) \leq y \mid w_{u_2}) w_{u_1}^\alpha w_{u_2}^\alpha |u|^{-d\alpha} (\rho t)^{\alpha d} dy. \end{aligned} \quad (10)$$

where we use a convolution over the cost of the left and right path segment and that the density of $c_{(u_1, u_2)}$ is at most $w_{u_1}^\alpha w_{u_2}^\alpha |u|^{-d\alpha} (\rho t)^{\alpha d}$. Had there not been weights involved, the proof of Lemma 5.1 in [13] would show immediately that

$$P_{\text{short}} \leq c_{\text{short}} |u|^{-\alpha d} h(t)$$

for some constant c_{short} depending on α . In our case we first need to get rid of the weights. To this end, we use the following simple proposition whose proof we defer for now.

► **Proposition 12.** For any $u_1 \in \mathbb{Z}^d$ and any $t > 0$, we have

$$\mathbb{P}(d_G^{\text{cost}}(0, u_1) \leq t \mid w_1) \leq 2w_1^\alpha \mathbb{P}(d_G^{\text{cost}}(0, u_1) \leq t).$$

Applying this proposition to (10) yields

$$\begin{aligned} & \mathbb{P}(d_G^{\text{cost}}(0, u_1) + c_{(u_1, u_2)} + d_G^{\text{cost}}(u_2, u) \leq t \mid w_{u_1}, w_{u_2}) \\ & \leq 4w_{u_1}^{2\alpha} w_{u_2}^{2\alpha} |u|^{-d\alpha} (\rho t)^{\alpha d} \int_0^t d\mathbb{P}(d_G^{\text{cost}}(0, u_1) \leq s) \int_0^{t-s} \mathbb{P}(d_G^{\text{cost}}(u_2, u) \leq y) dy. \end{aligned}$$

Then, applying this to (9) and taking the sum into the integrals, we get

$$\begin{aligned} P_{\text{short}} & \leq \sum_{\substack{u_1, u_2 \in \mathbb{Z}^d \\ |u_1 - u_2| \geq |u|/(\rho t)}} \left(\iint 4w_{u_1}^{2\alpha} w_{u_2}^{2\alpha} d\mu(w_{u_1}) d\mu(w_{u_2}) \right) \\ & \quad \times \left(|u|^{-d\alpha} (\rho t)^{\alpha d} \int_0^t d\mathbb{P}(d_G^{\text{cost}}(0, u_1) \leq s) \int_0^{t-s} \mathbb{P}(d_G^{\text{cost}}(u_2, u) \leq y) dy \right) \\ & \leq \left(\int_1^\infty \int_1^\infty 4(\tau - 1)^2 w_{u_1}^{2\alpha - \tau} w_{u_2}^{2\alpha - \tau} dw_{u_1} dw_{u_2} \right) \\ & \quad \times \left(|u|^{-d\alpha} (\rho t)^{\alpha d} \sum_{u_1 \in \mathbb{Z}^d} \sum_{u_2 \in \mathbb{Z}^d} \int_0^t d\mathbb{P}(d_G^{\text{cost}}(0, u_1) \leq s) \int_0^{t-s} \mathbb{P}(d_G^{\text{cost}}(u_2, u) \leq y) dy \right) \end{aligned}$$

where we took the sum into the second set of parentheses and omitted the condition $|u_1 - u_2| \geq |u|/(\rho t)$ from the sum and then split it into two sums. We also replaced $\mu(w) = 1 - w^{1-\tau}$ so $d\mu(w) = (\tau - 1)w^{-\tau} dw$. Now, taking these sums into the integrals and recalling that $g(x) = \sum_{v \in \mathbb{Z}^d} \mathbb{P}(d_G^{\text{cost}}(0, v) \leq x)$ yields that

$$P_{\text{short}} \leq \left(\int_1^\infty \int_1^\infty 4w_{u_1}^{2\alpha - \tau} w_{u_2}^{2\alpha - \tau} dw_{u_1} dw_{u_2} \right) \left(|u|^{-d\alpha} (\rho t)^{\alpha d} \int_0^t dg(s) \int_0^{t-s} g(y) dy \right).$$

Here, it is easy to see that the term in the second set of parentheses is at most $c_{\text{short}} |u|^{-\alpha d} h(t)$ for some constant c_{short} as was formally shown in [13, Lemma 5.1]. The term in the first set of parentheses is some constant that only depends on α, τ since the condition $2\alpha < \tau - 1$ ensures that the integrals converge. This constant enters into c_{short} .

If we now choose $\rho > ecb$ and $\delta = \rho \log \frac{\rho}{ecb}$, then summing P_{long} from (8) and P_{short} as above yields

$$f(r, t) \leq c_f r^{-\alpha d} h(t)$$

for some constant c_f as desired. ◀

We now give the proof of the proposition deferred above.

Proof of Proposition 12. Recall that the statement of the proposition is that

$$\mathbb{P}(d_G^{\text{cost}}(0, u_1) \leq t \mid w_1) \leq 2w_1^\alpha \mathbb{P}(d_G^{\text{cost}}(0, u_1) \leq t)$$

We show this using a coupling argument. To this end, we consider a model M that resembles CFFP where instead of having the vertex u_1 with weight w_1 , we set the weight of u_1 *deterministically* equal to 1. However, to decide the costs of edges adjacent to u_1 , we take the minimum of $\lceil w_1^\alpha \rceil$ many independent samples, i.e., for the edge (u_1, v) we set its cost to the minimum of $\lceil w_1^\alpha \rceil$ independent samples from an exponential distribution with

rate $w_v^\alpha |u - v|^{-\alpha d}$. Alternatively, this can be seen as having $\lceil w_1^\alpha \rceil$ copies of u_1 (with weight 1) and then studying the minimal cost-distance from 0 to one of the copies u_1 where the minimum is taken over all copies. Note that this description diverges from the one in the previous sentence if we are interested in paths between two arbitrary vertices, but since we are only concerned with paths from 0 to u_1 , the two are equivalent.

We continue by showing that cost distances in CFFP (from 0 to u_1) stochastically dominate those in M , i.e., that

$$\mathbb{P}(d_G^{\text{cost}}(u_1, 0) \leq t \text{ in CFFP} \mid w_1) \leq \mathbb{P}(d_G^{\text{cost}}(u'_1, 0) \leq t \text{ in } M).$$

To this end, we first show that this is the case for the cost of all edges incident to u_1 . Let $e = (u_1, v)$ be an arbitrary such edge and note that

$$\begin{aligned} \mathbb{P}(c_{(u_1, v)} \geq x \text{ in } M) &= \left(e^{-x w_v^\alpha |u_1 - v|^{-\alpha d}} \right)^{\lceil w_1^\alpha \rceil} \\ &\leq e^{-x w_1^\alpha w_v^\alpha |u_1 - v|^{-\alpha d}} = \mathbb{P}(c_{(u_1, v)} \geq x \text{ in CFFP}). \end{aligned}$$

Now, let \mathcal{R} be a fixed realization of all weights and edge costs in CFFP not associated with u_1 and denote by $[d_G^{\text{cost}}(0, v)]_{\mathcal{R}}$ the cost distance from 0 to v in this realization when u_1 is removed from the underlying graph. With this, we note that – conditional on \mathcal{R} – we have

$$d_G^{\text{cost}}(u_1, 0) \leq t \iff \exists v \text{ such that } c_{(u_1, v)} \leq t - [d_G^{\text{cost}}(0, v)]_{\mathcal{R}}.$$

Since $\mathbb{P}(c_{(v, u_1)} \geq x \text{ in } M) \leq \mathbb{P}(c_{(v, u_1)} \geq x \text{ in CFFP})$ as shown above, the probability that this occurs in M is at least as large as the corresponding probability in CFFP, and since this holds conditional on any realization \mathcal{R} , it also holds unconditionally by the law of total probability. Hence, we have shown the desired stochastic domination.

It remains to be shown that

$$\mathbb{P}(d_G^{\text{cost}}(u_1, 0) \leq t' \text{ in } M) \leq 2w_1^\alpha \mathbb{P}(d_G^{\text{cost}}(u_1, 0) \leq t \text{ in CFFP}).$$

To prove this, we again consider an arbitrary but fixed realization \mathcal{R} as in the previous paragraph and recall that $d_G^{\text{cost}}(u_1, 0) \leq t$ if and only if there is some v such that $c_{(u_1, v)} \leq t - [d_G^{\text{cost}}(0, v)]_{\mathcal{R}}$. Note that by the definition of M , this occurs if and only if it happens for at least one of the $\lceil w_1^\alpha \rceil \leq 2w_1^\alpha$ copies of u_1 . Since for each copy, the probability that this happens for said copy is $\mathbb{P}(d_G^{\text{cost}}(u_1, 0) \leq t \text{ in CFFP} \mid w_{u_1} = 1, \mathcal{R})$, we get from a union bound and from the law of total probability that

$$\begin{aligned} \mathbb{P}(d_G^{\text{cost}}(u_1, 0) \leq t \text{ in } M) &\leq 2w_1^\alpha \mathbb{P}(d_G^{\text{cost}}(u_1, 0) \leq t \text{ in CFFP} \mid w_{u_1} = 1) \\ &\leq 2w_1^\alpha \mathbb{P}(d_G^{\text{cost}}(u_1, 0) \leq t \text{ in CFFP}) \quad \blacktriangleleft \end{aligned}$$

Finally, we will use Lemma 10 in conjunction with Theorem 5.3 from [13] to show the desired lower bound on cost-distances. We restate that theorem here, simplified for our use case.

► **Theorem 13** (Theorem 5.3 from [13]). *Let $g(t) : [0, \infty) \rightarrow \mathbb{R}$ be a function satisfying*

$$1 \leq g(t) \leq e^{Ct}$$

and

$$g(t)^{1/\theta} \leq c_h \left(1 + t^{\beta-1} \int_0^t g(y)g(t-y) dy \right)$$

for all $t \geq 0$ for some constants $C > 0$, $\theta \in (\frac{1}{2}, 1)$, $\beta \geq 0$, and $c_h \geq 1$. Then, there exists a constant $c_\theta > 1$ such that $g(t) \leq G(t)$ for all $t \geq 0$ where $G(t)$ is defined such that

$$\log G(t) = c_\theta (2\lambda t)^{\log_2(2^\theta)} (\log(1 + t^\beta))^{\log_2(1/\theta)} (1 + o(1)).$$

We use the above theorem to prove our main lemma for FPP, which we restate here.

► **Lemma 14** (Tail Bound for Cost-Distances in FPP). *Consider FPP on SFP with $2\alpha < \tau - 1$ and arbitrary vertices x, y . There exists a constant c depending only on α and τ such that for $\Delta = \Delta(\alpha) = 1/\log_2(2/\alpha)$,*

$$\log \mathbb{P}(d_G^{\text{cost}}(x, y) \leq t) \leq c(\log(1 + t))^{1-1/\Delta} t^{1/\Delta} (1 + o(1)) - \alpha d \log|x - y| + c.$$

Proof. As discussed, it suffices to show the claim for CFFP. Moreover, by translation invariance, we can replace x with the origin 0 and y by $u = y - x$. By Lemma 11 we have

$$g(t)^\alpha \leq c_1 \left(t^{\alpha d} \int_0^t g(t-y)g(y) dy + 1 \right) \tag{11}$$

for some $c_1, \delta > 0$ depending only on parameters of the model. We also know that $g(t) \leq e^{Ct}$ from Theorem 9 for a C with the same dependencies. Therefore, using Theorem 5.3 from [13] (stated above the current theorem) with $\theta = \frac{1}{\alpha}, \beta = \alpha d + 1$ and with the inequality $1 + t^\beta \leq (1 + t)^\beta$, we have $g(t) \leq G(t)$, where

$$\log G(t) = c(\log(1 + t))^{1-\frac{1}{\Delta}} t^{\frac{1}{\Delta}} (1 + o(1)),$$

where c depends only on parameters of the model. Now, by Lemma 10, we have

$$\begin{aligned} \log \mathbb{P}(d_G^{\text{cost}}(0, u) \leq t) &\leq \log f(|u|, t) \\ &\leq c' - \alpha d \log|u| + \log \left(t^{\alpha d} \int_0^t g(t-y)(g(y) - 1) dy + e^{-\delta t} \right) \end{aligned}$$

The constant c' above comes from Lemma 10. The final step is to notice that inequality (11) is satisfied as equality for $G(t)$, as in [13]. Then, substituting the expression for it in the resulting inequality gives the desired bound. ◀

We can use the coupling employed in Theorem 1.1 to establish the above tail bound even if $2\alpha \geq \tau - 1$ but with Δ replaced by $\Delta'' = \Delta(\min\{\alpha, \frac{\tau-1}{2}\} - \varepsilon)$ for arbitrarily small ε . We then have obtained Theorem 1.3. Since the proof is almost verbatim the same as that of Theorem 1.1, we omit it. Note that one minor technical step that is required additionally to get the bound claimed in Theorem 1.3 using the one obtained from Lemma 14 is to introduce an auxiliary constant ε' in addition to the ε from the statement of the theorem. This swallows the $(\log(1 + t))^{1-\frac{1}{\Delta}}$ factor.

References

- 1 Mohammed Amin Abdullah, Michel Bode, and Nikolaos Fountoulakis. Typical distances in a geometric model for complex networks. *Internet Mathematics*, 1:38, 2017.
- 2 Noam Berger. A lower bound for the chemical distance in sparse long-range percolation models. *arXiv preprint*, 2004. [arXiv:math/0409021](https://arxiv.org/abs/math/0409021).
- 3 Marek Biskup. On the scaling of the chemical distance in long-range percolation models. *The Annals of Probability*, 32(4):2938–2977, October 2004. doi:10.1214/009117904000000577.

- 4 Marek Biskup. Graph diameter in long-range percolation. *Random Structures & Algorithms*, 39(2):210–227, 2011. doi:10.1002/rsa.20349.
- 5 Marek Biskup and Jeffrey Lin. Sharp asymptotic for the chemical distance in long-range percolation. *Random Structures & Algorithms*, 55(3):560–583, 2019. doi:10.1002/rsa.20849.
- 6 Thomas Bläsius and Philipp Fischbeck. On the external validity of average-case analyses of graph algorithms. *ACM Transactions on Algorithms*, 20(1):1–42, 2024.
- 7 Thomas Bläsius, Tobias Friedrich, Maximilian Katzmann, and Anton Krohmer. Hyperbolic embeddings for near-optimal greedy routing. *ACM J. Exp. Algorithmics*, 25, 2020. doi:10.1145/3381751.
- 8 Thomas Bläsius, Tobias Friedrich, and Christopher Weyand. Efficiently computing maximum flows in scale-free networks. *arXiv preprint*, 2020. arXiv:2009.09678.
- 9 Marián Boguná, Fragkiskos Papadopoulos, and Dmitri Krioukov. Sustaining the internet with hyperbolic mapping. *Nature communications*, 1(1):62, 2010.
- 10 Karl Bringmann, Ralph Keusch, and Johannes Lengler. Average distance in a general class of scale-free networks with underlying geometry. *arXiv preprint*, 2016. arXiv:1602.05712.
- 11 Karl Bringmann, Ralph Keusch, and Johannes Lengler. Geometric inhomogeneous random graphs. *Theoretical Computer Science*, 760:35–54, 2019.
- 12 Karl Bringmann, Ralph Keusch, Johannes Lengler, Yannic Maus, and Anisur R Molla. Greedy routing and the algorithmic small-world phenomenon. *Journal of Computer and System Sciences*, 125:59–105, 2022.
- 13 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(2):203–256, 2016. doi:10.1002/cpa.21571.
- 14 Maria Deijfen, Remco Hofstad, and Gerard Hooghiemstra. Scale-free percolation. *Annales de l’Institut Henri Poincaré Probabilités et Statistiques*, 49:817–838, August 2013. doi:10.1214/12-AIHP480.
- 15 Philippe Deprez, Rajat Subhra Hazra, and Mario V Wüthrich. Inhomogeneous long-range percolation for real-life network modeling. *Risks*, 3(1):1–23, 2015.
- 16 Philippe Deprez and Mario V Wüthrich. Scale-free percolation in continuum space. *Communications in Mathematics and Statistics*, 7(3):269–308, 2019.
- 17 Leslie Ann Goldberg, Joost Jorritsma, Júlia Komjáthy, and John Lapinskas. Increasing efficacy of contact-tracing applications by user referrals and stricter quarantining. *Plos one*, 16(5):e0250435, 2021.
- 18 Mark S Granovetter. The strength of weak ties. *American journal of sociology*, 78(6):1360–1380, 1973.
- 19 Nannan Hao and Markus Heydenreich. Graph distances in scale-free percolation: the logarithmic case. *Journal of Applied Probability*, 60(1):295–313, 2023.
- 20 Joost Jorritsma, Tim Hulshof, and Júlia Komjáthy. Not all interventions are equal for the height of the second peak. *Chaos, Solitons & Fractals*, 139:109965, 2020.
- 21 Christoph Koch and Johannes Lengler. Bootstrap percolation on geometric inhomogeneous random graphs. *Internet Mathematics*, page 18995, 2021.
- 22 Júlia Komjáthy, John Lapinskas, and Johannes Lengler. Penalising transmission to hubs in scale-free spatial random graphs. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, 57(4):1968–2016, 2021. doi:10.1214/21-AIHP1149.
- 23 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*, 2023. arXiv:2309.11840.
- 24 Júlia Komjáthy, John Lapinskas, Johannes Lengler, and Ulysse Schaller. Four universal growth regimes in degree-dependent first passage percolation on spatial random graphs ii. *arXiv preprint*, 2023. arXiv:230911880.

- 25 Júlia Komjáthy and Bas Lodewijks. Explosion in weighted hyperbolic random graphs and geometric inhomogeneous random graphs. *Stochastic Processes and their Applications*, 130(3):1309–1367, 2020. doi:10.1016/j.spa.2019.04.014.
- 26 Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguná. Hyperbolic geometry of complex networks. *Physical Review E*, 82(3):036106, 2010.
- 27 Gergely Ódor, Domonkos Czifra, Júlia Komjáthy, László Lovász, and Márton Karsai. Switchover phenomenon induced by epidemic seeding on geometric networks. *Proceedings of the National Academy of Sciences*, 118(41):e2112607118, 2021.
- 28 David Reimer. Proof of the van den berg–kesten conjecture. *Combinatorics, Probability and Computing*, 9(1):27–32, 2000. doi:10.1017/S0963548399004113.
- 29 Pieter Trapman. The growth of the infinite long-range percolation cluster. *The Annals of Probability*, 38(4):1583–1608, 2010.

A On the upper bound claim in [19]

In this section, we briefly explain the mistake in the upper bound proof for graph distances in [19]. There, the authors add edges for some paths of length 2 in SFP, thus cutting graph distances at most in half. Let us call the resulting graph 2-SFP. Then they compare 2-SFP to LRP with different parameters. The 2-SFP graph does not dominate an LRP because the edges are correlated, but for every edge e , $\mathbb{P}(e \text{ open in 2-SFP}) \geq \mathbb{P}(e \text{ open in LRP})$.

Then they make the argument that the correlations are positive, so that for every fixed x - y -path π , by the FKG inequality:

$$\mathbb{P}(\pi \text{ open in 2-SFP}) \geq \mathbb{P}(\pi \text{ open in LRP}). \quad (12)$$

That is correct, but it does not imply the statement that we would want for a suitable k :

$$\begin{aligned} \mathbb{P}(\exists x\text{-}y\text{-path } \pi \text{ of length } \leq k: \pi \text{ open in 2-SFP}) \\ \geq \mathbb{P}(\exists x\text{-}y\text{-path } \pi \text{ of length } \leq k: \pi \text{ open in LRP}). \end{aligned} \quad (13)$$

Instead, (12) only implies by summing over all x - y -paths of length at most k :

$$\mathbb{E}[\#\text{ of open } x\text{-}y\text{-paths of length } \leq k \text{ in 2-SFP}] \quad (14)$$

$$\geq \mathbb{E}[\#\text{ of open } x\text{-}y\text{-paths of length } \leq k \text{ in LRP}]. \quad (15)$$

However, it is not hard to see that due to the correlations the left hand side of (14) is dominated by low-probability events where the number of paths is very large. In particular, the upper bound proof for LRP is centered around the concept of *hierarchies*, and the first step of a hierarchy is to find an edge of length $\Theta(|x - y|)$, where the two endpoints lie respectively close to x and y . It is easy to see that for some parameters considered in [19], with high probability such edges do exist in LRP but do not exist in 2-SFP. However, the *expected number* of such edges in 2-SFP (corresponding to the left hand side of (14)) is still large because the unlikely event of a vertex of weight $\Theta(|x - y|)$ at distance $|x - y|$ from x induces a very large number of 2-paths in SFP, which become edges in 2-SFP.

Hence, the argument in [19] shows (14) but not (13), and this is not a minor omission but a major gap. In fact, we conjecture that the upper bound statements in [19] are false, and that the exponent $\Delta(\alpha) = 1/\log_2(2/\alpha)$ is tight throughout the polylogarithmic regime, i.e. for all $\tau > 3$ and $\alpha \in (1, 2)$. The reason for this intuition is that for $\tau > 3$ there exists a constant $C > 0$ such that the induced graph of vertices of weight larger than C does not percolate. However, we do not see an obvious way to leverage this property into lower bounds for graph distances.