Short Paths in Scale-free Percolation

Dissertation
an der Fakultät für Mathematik, Informatik und Statistik
an der Ludwig-Maximilians-Universität München

eingereicht von
Nannan Hao
07.10.2022

1. Gutachter: Prof. Dr. Markus Heydenreich
2. Gutachter: Prof. Dr. Noam Berger Steiger
3. Gutachter: Prof. Dr. Pieter Trapman

Tag der mündlichen Prüfung: 16. März 2023

Eidesstattliche Versicherung

(Siehe Promotionsordnung vom 12.07.2011, §8, Abs. 2, Pkt. 5.)

Hiermit erkläre ich an Eides statt, dass die Dissertation von mir selbstständig, ohne unerlaubte Beihilfe angefertigt ist.

München, 07.10.2022

Nannan Hao

Acknowledgement

First of all I would like to express my gratitude to my supervisor, Markus Heydenreich, who gives me the chance to explore the fantastic world of random graphs. His keen insight and profound expertise have guided me in my research throughout my PhD career.

I would like to thank Noam Berger and Pieter Trapman for accepting the invitation as referees for my dissertation. Noam, whom I already know since my Master's study, taught me much about probability theory, and I benefit quite from that in my research.

I also want to thank Johannes Lengler for the inspiring discussions with him. His patient reply to my questions helps me so much to understand the greedy routing algorithm better.

Without the support from my family I could not have finished my PhD so smoothly. I want to thank my parents for their support that encouraged me when I decided to pursue a doctorate degree, and when I met difficulties in the research.

Illustrations make the dissertation more understandable. I want to thank Junyi Lu for her advice on the simulations of several random graph models.

I would like to thank Matthew Dickson for proofreading part of my dissertation.

Despite the home office due to the Covid-19 pandemic, I could still spend some time in the institute with my kind colleagues, including but not limited to Alejandro, Anna, Felix, Jakob, Kilian, Leon, Leonid, Matija, Philipp, Simon, Stefan, Tamas, Thomas, Umberto and Vincent. I would like to thank all of them for the enjoyable coffee breaks and the exciting table tennis games.

Zusammenfassung

Seit dem soziologischen Experiment von Milgram in den 1960er Jahren steht der Graphenabstand in komplexen Netzwerken, insbesondere in sozialen Netzwerken, im Fokus der Forschung in Netzwerken. In dieser Dissertation beschäftigen wir uns mit einem räumlichen Zufallsgraph, der als skalenfreie Perkolation bekannt ist. Dieser Graph zeigt ein reiches Phasendiagramm und wir konzentrieren uns auf dessen kurzen Pfade. In diesem Modell sind $x, y \in \mathbb{Z}^d$ mit einer Wahrscheinlichkeit, die von Gewichten W_x, W_y sowie von dem euklidischen Abstand $|x - y|$ abhängt, verbunden.

Zunächst untersuchen wir asymptotische Abstände in einem Parameterregime, in dem der Graphenabstand polylogarithmisch im euklidischen Abstand ist. Mithilfe eines multiskalaren Arguments erhalten wir verbesserte Schranken für den logarithmischen Exponenten. Im Heavy-tailed-Regime zeigt die Verbesserung in der oberen Schranke eine Diskrepanz zu Long-range-Perkolation. Im Light-tailed-Regime wird der korrekte Exponent identifiziert.

Im folgenden Teil der Dissertation erforschen wir Navigationsmöglichkeiten in dem Modell. Wir untersuchen die Möglichkeit, mit ausschließlich lokalen Informationen (Gewichten und Positionen der Nachbarknoten) kurze Pfade zwischen zwei gegebenen Knoten zu finden. In dem Regime mit polylogarithmischen Graphabständen zeigen wir, dass jeder Algorithmus, der auf lokalen Informationen der Knoten basiert, mindestens polynomiell viele Schritte benötigt, um das Ziel zu finden. Im Gegensatz dazu findet ein Greedy-routing-Algorithmus in dem Parameterregime, in welchem der Graphenabstand doppelt logarithmisch im euklidischen Abstand ist, einen kurzen Pfad derselben Längenordnung.

Abstract

Graph distances in real-world networks, in particular social networks, have been always in the focus of network research since Milgram's sociological experiment in 1960s. In this dissertation we specialize in a geometric random graph known as scale-free percolation, which shows a rich phase diagram regarding graph distances, and focus on short paths in it. In this model, $x, y \in \mathbb{Z}^d$ are connected with probability depending on i.i.d weights W_x, W_y and their Euclidean distance $|x - y|$.

First we study asymptotic distances in a regime where graph distances are poly-logarithmic in Euclidean distance. With a multi-scale argument we obtain improved bounds on the logarithmic exponent. In the heavy tail regime, improvement of the upper bound shows a discrepancy with long-range percolation. In the light tail regime, the correct exponent is identified.

The following part of this dissertation investigates navigation possibility in the model. More precisely, we study the possibility to find short paths between two vertices given only local information (weights and locations of neighbors). In the regime where graph distances are poly-logarithmic we show that any algorithm based on local information takes at least polynomial steps to find the target. In contrast, in the regime where the graph distance is doubly logarithmic in the Euclidean distance, a short path with length of the same order can be found by a greedy routing algorithm.

Contents

Chapter 1

Introduction

1.1 Background

1.1.1 Real-world networks

The study of complex networks has become popular in recent years. A complex network contains typically a large number of elements, and links that represent the interactions between the elements. Many systems in real life, like social or technological networks such as Facebook or World Wide Web, can be regarded as complex networks. In order to analyse their structures and behavior, random graph models were proposed and are now commonly used to simulate these networks.

A graph consists of a set of vertices that represent the elements in the networks, and a set of edges that describe the links between the elements. For example, in a social network like Facebook, the accounts can be viewed as vertices in a graph, and two vertices of this kind are linked by an edge if they are friends of each other. A formal introduction of the terminologies for graphs can be found in the following section.

Although diverse real-world networks differ significantly, many of them share essentially several common patterns. Here we point out two most famous properties many networks in the real world possess.

- Scale-free property. A network is said to have the scale-free property, if the degree distribution follows a power law. Here the degree of a vertex is the number of vertices this vertex is connected to. In other words, a scale-free network is characterized by the existence of hubs that own many contacts.

Many real-world networks turn out to exhibit the scale-free property, e.g. the World Wide Web [5], the collaboration of movie actors in films [5], and some financial networks [27].

- Small-world property. A finite network or graph is said to have the small-world property, if the graph distance in the network is much smaller then the number of vertices. If the graph is embedded into some metric space, then small-world property means that the graph distance is negligible compared to the metric between two vertices. In the famous sociological experiment by Milgram in the 1960s, it is observed that the average graph distance between individuals in Omaha and Boston is around 6, if we view friendship and kinship as an edge between individuals, while the Euclidean distance between the two cities is around 2000 kilometers [70].

We will introduce the two properties more quantitatively in Section 1.1.2.

1.1.2 Random graph models

Graph terminologies

A graph $G = (V, E)$ consists of a countable set of vertices (or nodes), called vertex set, V and a collection of edges, called edge set, E which is a subset of $V \times V$. More precisely, $E \subseteq \{\{x, y\} : x, y \in V, x \neq y\}$. If $\{x, y\} \in E$, we say that x is *adjacent* to y, and write $x \sim y$. In this case the edge $\{x, y\}$ is called *incident* to x and y, and x, y are *neighbors* of each other. Similarly, two edges are also called *adjacent* if they share a common vertex. A graph $G = (V, E)$ is called *undirected*, if $\{x, y\} = \{y, x\}$ for all $\{x, y\} \in E$. Otherwise if $\{x, y\} \neq \{y, x\}$ for some x and y, then the graph is *directed*. In some cases we also allow *self-loops* and *multiple edges* in a graph, e.g in preferential attachment model. An edge $\{x, y\}$ is called a self-loop, if $x = y$. A graph that does not contain any self-loop or multiple edges is called *simple*. Throughout this dissertation, unless specifically mentioned otherwise, we always refer to *undirected* and *simple* graphs.

The degree d_x of a vertex x in the graph $G = (V, E)$ is defined as the number of edges incident to it, that is

$$d_x := \#\{y \in V : x \sim y\} = \sum_{y \in V} \mathbb{1}_{x \sim y}.$$

A graph $G = (V, E)$ is called locally finite, if $d_x < \infty$ for all $x \in V$. G is a complete graph, if $\{x, y\} \in E$ for all $x \neq y$. A subgraph $G' = (V', E')$ of $G = (V, E)$ is a graph such that $V' \subseteq V$ and $E' \subseteq E$. A graph $G = (V, E)$ is called the Cartesian product of $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ if $V = V_1 \times V_2$ and

$$E = \big\{\{(x_1, y_1), (x_2, y_2)\} : (x_1 = x_2, (y_1, y_2) \in E_2) \text{ or } (y_1 = y_2, (x_1, x_2) \in E_1)\big\},$$

and we denote it by $G = G_1 \times G_2$.

A *path* π in a graph $G = (V, E)$ is a sequence of distinct vertices $(x_i)_{i=0,1,\ldots,n} \subseteq V$ such that $\{x_i, x_{i+1}\} \in E$ for $i = 0, 1, \ldots, n-1$, and π is said to *join* x_0 and x_n. A graph $G = (V, E)$ is called *connected*, if for every pair (x, y) with $x \neq y$ there exists a path joining x and y. A subgraph $\mathcal{C}$ of G is called a *cluster*, if $\mathcal{C}$ is connected and is not connected to any other vertex that is not in $\mathcal{C}$. The *graph distance* $D(x, y)$ between x, y is the minimum number of edges among all paths joining x, y. If such path does not exist, i.e. x and y are in different clusters, then $D(x, y) = \infty$.

Random graphs

The concept ‘random graph’ refers to a probability distribution on a family of graphs. Typically in a random graph edges are randomly generated. Random graph was first introduced by Erdős and Rényi [39, 40, 41, 42]. In [40] they gave rather ample results about Erdős-Rényi random graph, which is one of the simplest but most instructive random graph models. Since then random graph theory was broadly extended, and varieties of random graph models have been proposed and investigated, in order to simulate the real-world networks. Alon and Spencer [3], Bollobás [15] give more details about the early literature on random graphs. Here we first introduce some basic (finite) random graph models. For more sophisticated models, especially spatial random graphs, we refer to Section 1.2 and 1.4 for more details.

As a preparation for the subsequent parts of the dissertation we reintroduce the aforementioned properties for random graphs from Section 1.1.1 in a more mathematical way.

- *Scale-free property.* A random graph is called scale-free, if the degree distribution of its vertices satisfies a power law asymptotically. More precisely, for a vertex x in the graph, there exists a constant $\kappa > 0$ such that for n large enough it holds

$$\mathbb{P}(d_x \geq n) \approx n^{-\kappa}. \tag{1.1}$$

Mind here the approximation "≈" has different interpretations in different literature. One of the most common understanding of (1.1) is that there exists a *slowly varying* function $\ell : \mathbb{R}^+ \to \mathbb{R}^+$ such that

$$\mathbb{P}(d_x \geq n) = n^{-\kappa}\ell(n), \quad \text{for all } n \in \mathbb{N}.$$

A function $\ell : \mathbb{R}^+ \to \mathbb{R}^+$ is called slowly varying, if for all $a > 0$ the following holds

$$\lim_{x\to\infty} \frac{\ell(ax)}{\ell(x)} = 1.$$

It is easy to verify that constant functions and (poly)-logarithmic functions are slowly varying.

- *Small-world property.* On finite networks, say with N vertices, "small world" means that the graph distance between two points is much shorter than a regular structure would suggest, e.g. $(\log N)^{O(1)}$ as $N \to \infty$. For infinite networks, a different interpretation to the small-world effect is given. We call an infinite subgraph $\mathcal{C} \subset \mathbb{Z}^d$ a small-world graph if the graph distance $D(x, y)$ on $\mathcal{C}$ is much smaller than the Euclidean distance, that is if for example

$$D(x, y) = \big(\log |x - y|\big)^{O(1)} \qquad \text{as } |x - y| \to \infty. \tag{1.2}$$

Besides, a sequence of events $(E_n)_{n\in\mathbb{N}}$ is called to happen *with high probability*, if it holds that

$$\lim_{n\to\infty} \mathbb{P}(E_n) = 1.$$

Sometimes it is abbreviated as *w.h.p* for brevity. In some literature, it is also called "*asymptotically almost surely*" (or *a.a.s* for short). In this dissertation we will mainly state the assertions in the former way.

Erdős-Rényi random graph

Erdős-Rényi random graph was first introduced by Erdős and Rényi [40], Gilbert [49], Austin et al. [4] with slight differences. For this model one considers $[n] := \{1, 2, \ldots, n\}$ as the vertex set V. For each pair $i \neq j$, the edge $\{i, j\}$ between the two vertices is open (or present) with probability p independently of all other pairs. As a consequence we obtain a graph with a deterministic vertex set but a random edge set, and we denote it by $\mathrm{ER}(n, p)$. An example of Erdős-Rényi random graph is

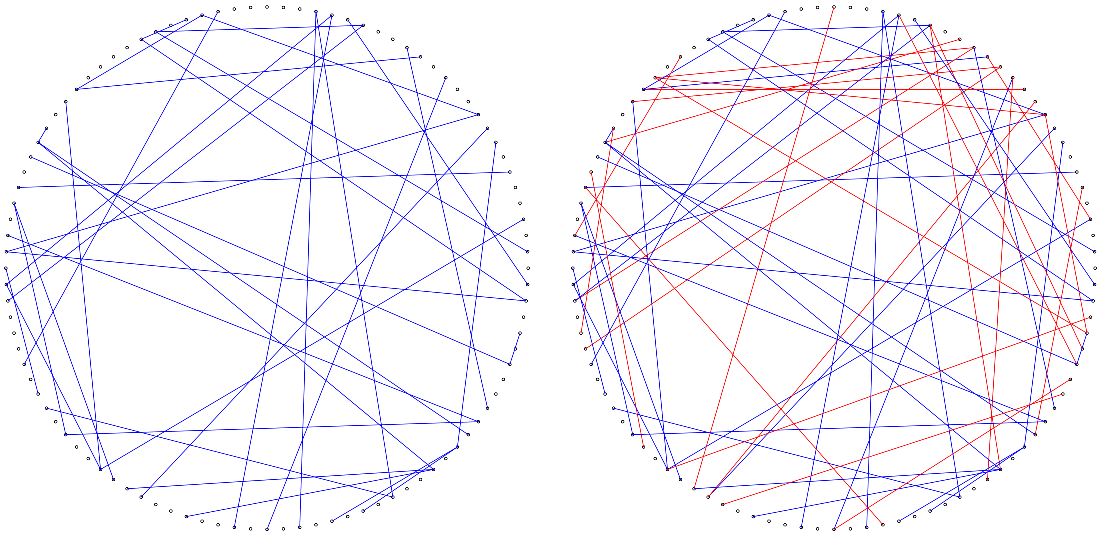

(a) $\lambda = 0.7$.　　(b) $\lambda = 1.3$.

Figure 1.1: Two realizations of Erdős-Rényi random graphs $\mathrm{ER}(n, \lambda/n)$ with $n = 100, \lambda = 0.7$ in (a) and $n = 100, \lambda = 1.3$ in (b) respectively.

illustrated in Figure 1.1. Note that in Figure 1.1 coupling is used in the simulation to show the monotonicity in the value of p.

It is easy to see that the degree of a vertex i in $\mathrm{ER}(n, p)$ follows a binomial distribution, that is,

$$\mathbb{P}(d_i = k) = \binom{n-1}{k} p^k (1-p)^{n-1-k}.$$

If we choose $p = \lambda/n$ for some positive constant λ, we see as the number of vertices n goes to infinity,

$$\mathbb{P}(d_i = k) = \binom{n-1}{k} p^k (1-p)^{n-1-k} \to e^{-\lambda} \frac{\lambda^k}{k!}.$$

Note that the limit is the mass function of a Poisson distribution. On the other hand, $\frac{\lambda^k}{k!}$ decays faster than $k^{-\tau}$ for arbitrary $\tau > 0$ as k increases, thus $\mathrm{ER}(n, \lambda/n)$ is *not* scale-free.

In $\mathrm{ER}(n, \lambda/n)$ we denote by $\mathcal{C}(i)$ the open cluster that contains vertex i, and $\mathcal{C}_{\max}$ the open cluster such that

$$|\mathcal{C}_{\max}| = \max_{i \in [n]} |\mathcal{C}(i)|.$$

Mind that the equation above can only identify the size of $\mathcal{C}_{\max}$, but not $\mathcal{C}_{\max}$ itself uniquely. If several open clusters are of maximal size, we choose the one containing the vertex with smallest index as $\mathcal{C}_{\max}$. As the parameter λ varies, the size of $\mathcal{C}_{\max}$ differs significantly, and thus exhibits a phenomenon of phase transition. In the subcrtical phase ($\lambda < 1$), as $n \to \infty$ it is shown [3, 18]

$$\frac{|\mathcal{C}_{\max}|}{\log n} \xrightarrow{\mathbb{P}} \frac{1}{\lambda - 1 - \log \lambda},$$

while in the supercritical phase ($\lambda > 1$), for every $\nu \in (1/2, 1)$ there exists $\delta = \delta(\nu, \lambda) > 0$ such that

$$\mathbb{P}\left(\big||\mathcal{C}_{\max}| - \zeta n\big| \geq n^{\nu}\right) = O\left(n^{-\delta}\right),$$

where ζ is a positive constant [3]. Around the criticality ($\lambda = 1$), $\mathcal{C}_{\max}$ is asymptotically of size $n^{2/3}$ [65, 73]. For more details about $\mathrm{ER}(n, \lambda/n)$ like connectivity, degree distribution and limit theorems we refer to [63].

Generalized random graphs

As we haven seen, Erdős-Rényi model is a *homogeneous* random graph in the sense that all the vertices play the same role. Therefore it is not a suitable model for the networks with heterogeneous nodes like Facebook community where celebrities have much more contacts than others. In order to overcome this drawback, the generalized random graph model is studied [22, 43]. In this model, we take again $[n] := \{1, 2, \ldots, n\}$ as the vertex set. For $i, j \in [n]$ with $i \neq j$, the edge $\{i, j\}$ is open independently with probability $p_{i,j}$ given as follows:

$$p_{i,j} := \frac{w_i w_j}{\ell_n + w_i w_j}, \tag{1.3}$$

where $\boldsymbol{w} = (w_i)_{i \in [n]}$ are the given *vertex weights*, and $\ell_n = \sum_{i=1}^{n} w_i$ is the total weight. We denote by $\mathrm{GRG}(n, \boldsymbol{w})$ the generalized random graph with n vertices and weights $\boldsymbol{w}$. Mind that $\mathrm{ER}(n, \lambda/n)$ is a special case of $\mathrm{GRG}(n, \boldsymbol{w})$ by taking $w_i = \frac{n\lambda}{n-\lambda}$ for all $i \in [n]$.

A more general version of $\mathrm{GRG}(n, \boldsymbol{w})$ admits one more layer of randomness in the way that it allows the weights to be independent and identically distributed random variables. More precisely, we first sample the i.i.d weights (W_i) for (w_i). Given the weights, we connect i and j independently with probability $p_{i,j}$ as in

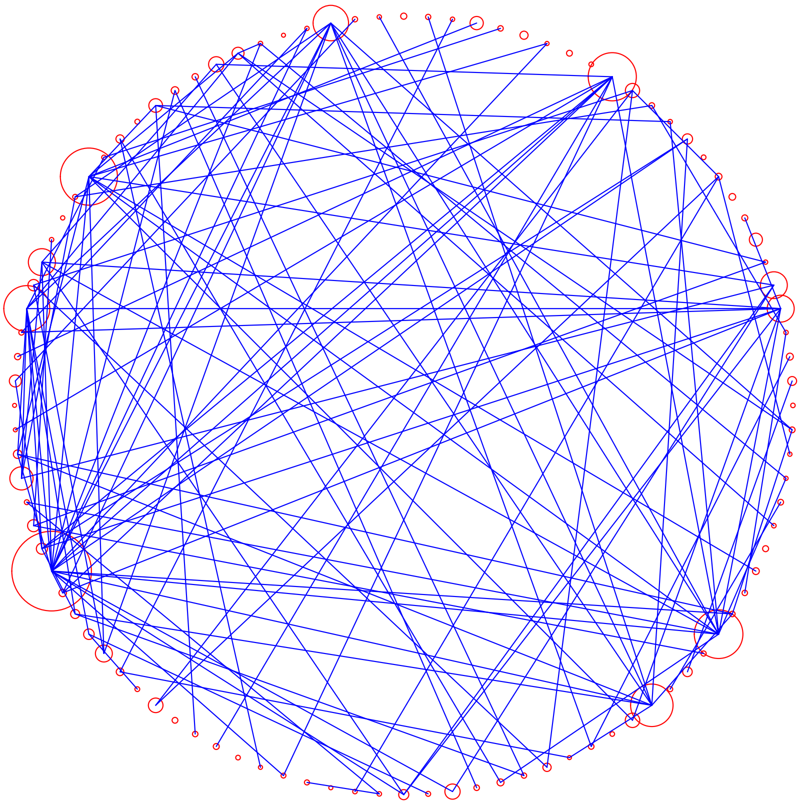

Figure 1.2: An example of generalized random graph with 100 vertices and i.i.d weights. The weight W has the tail distribution $\mathbb{P}(W > x) = x^{-1.5}, x \geq 1$. The radii of the red circles are proportional to the corresponding weights.

(1.3). Consequently, the edges are not independent anymore due to the fact that ℓ_n appears in every edge probability. Later we will encounter a similar setting in Section 1.2 when we introduce the main object of this dissertation. A realization of generalized random graph with i.i.d weights can be seen in Figure 1.2.

Preferential attachment model

Preferential attachment model (PA model) was first considered by Yule to explain the power-law distribution of the number of species per genus of flowering plants [77]. Later on Simon [75], Price [74] made use of preferential attachment to investigate other phenomena in real life. In 1999 Barabási and Albert [5] proposed the application of PA models to analyze the growth of World Wide Web. Here we introduce a classical version of PA model from [63] as a sequence of random graphs $(\mathrm{PA}_t^\delta)_{t\in\mathbb{N}}$ that admit self-loops. We start the introduction of the model by defining the graph for $t = 1$, and then construct the sequence recursively. The graph PA_1^δ

consists of a single vertex v_1 with a self-loop. Assume we have PA_t^δ with vertices $(v_i)_{i\in[t]}$. Now we describe the growth of the sequence from PA_t^δ to PA_{t+1}^δ. Given PA_t^δ, we add one more vertex v_{t+1}, choose one vertex v_i out of $(v_i)_{i\in[t+1]}$ according the following probability

$$p_i^{(t)} = \begin{cases} \frac{D_i(t)+\delta}{t(2+\delta)+(1+\delta)}, & \text{if } i \in [t], \\ \frac{1+\delta}{t(2+\delta)+(1+\delta)}, & \text{if } i = t+1, \end{cases}$$

where $D_i(t)$ is the degree of v_i in PA_t^δ, and then connect v_{t+1} and the chosen v_i. In the case $i = t+1$, v_{t+1} will be connected to itself and thus form a self-loop, leading to the isolation of v_{i+1} from PA_t^δ. As a result of the mechanism, the random graph PA_t^δ has t vertices and t edges with self-loops counted in.

A similar variation of the preferential attachment model allows multiple edges but disallows self-loops. We denote by $\left(\mathrm{PA}_t^\delta(b)\right)_{t\in\mathbb{N}}$ this variation where b stands for "model (b)" in order to distinguish from the original version. This variation of PA model can be described as follows: For $t = 2$ the graph $\mathrm{PA}_2^\delta(b)$ has two vertices v_1, v_2 with two edges between them which are multiple edges. Conditioned on $\mathrm{PA}_t^\delta(b)$, we add one more vertex v_{t+1}, choose one of the previous nodes according to the probability

$$p_i^t = \frac{D_i(t)+\delta}{t(2+\delta)}, \quad i \in [n],$$

and then connect v_{t+1} to the chosen v_i. In this way we obtain $\mathrm{PA}_{t+1}^\delta(b)$. The advantage of the model variation is that the resulting graph is always connected.

As we can see in the dynamics of preferential attachment model, the newly added vertex has a preference for those vertices with more contacts in the current graph, leading to the so-called "Matthew effect" in the resulting graph, which can be roughly summarized by the adage "the rich get richer and the poor get poorer".

One important feature of preferential attachment model is that its asymptotic degree sequence satisfies a power law [17]. Therefore it offers a possible explanation for the empirical power-law distribution observed from real-world phenomena. Due to its scale-free property, the PA model has gained the attention of random graph community, and a number of PA model variants have been proposed and studied, for example, see [16] for a directed PA model, [25] for a quite general version, [10, 11] for a competition-induced PA model, [31, 32, 33] for PA models with conditionally independent edges, and [61] for a spatial PA model.

1.1.3 Percolation

The model

Percolation was first introduced by Broadbent and Hammersley [23] in 1957, as a stochastic model for the flow of fluid through porous materials. It turns out that percolation has been of interest for not only mathematicians but also physicists, since it exhibits important properties like phase transition and critical phenomena.

Imagine we immerse some porous material like a piece of stone in the water. The water flows from some place inside the material to the other if there is an open micro-channel between them. A natural question is: what is the probability that some certain location in the material gets saturated? In other words, what is the probability that there is an open path from the surface of the material to this point?

In order to model this phenomenon, *Bernoulli bond percolation* (henceforth abbreviated as bond percolation) was proposed, and we describe now the model briefly. Consider the lattice $\mathbb{Z}^d$ as the vertex set for some integer $d \geq 1$. Let p be a number with $0 \leq p \leq 1$. With probability p a pair of nearest neighbors $\{x, y\}$ in $\mathbb{Z}^d$ is connected independently of all others, and we call the edge (or bond) $\{x, y\}$ is open, if x and y are connected. Note here x, y are nearest neighbors if $\|x - y\|_1 = \sum_{i=1}^d |x_i - y_i| = 1$. With this mechanism we obtain a locally finite random subgraph of the complete graph on $\mathbb{Z}^d$.

In the context of the modeling, the given material can be viewed as a large finite subset of $\mathbb{Z}^3$. Then the water can flow from x to y if there exists an open path between them. More generally, if we consider bond percolation on $\mathbb{Z}^d$, then the corresponding question for percolation model will be what the probability is that the origin is connected to infinity (for brevity we hence denote this probability by $\theta(p)$). Note here it makes more sense to consider the path to infinity because we have an infinite vertex set. For sure this probability depends first on the parameter p. In real world for the materials like granite, we can imagine that the value of p is close to 0, because there is hardly hole inside. In contrast, for those very porous media like sponge, it is reasonable to assume that p is very close to 1 for the modeling.

Intuitively speaking, a higher value of p results in more connectedness in bond percolation. In other words, θ is a monotonically increasing function of p. This claim can be confirmed via a coupling argument [53].

If 0 is connected to infinity, then it lies in some infinite open cluster. We denote by $\lambda(p)$ the probability that there exists an infinite open cluster. Since we have the independence of all edges, Kolmogorov's 0-1 law (see e.g Theorem 2.5.3 in [38]) ensures that $\lambda(p)$ takes values only in $\{0, 1\}$. Note that $\lambda(0) = 0, \lambda(1) = 1$ and λ is increasing in p. Therefore there exists $p_c \in [0, 1]$ such that $\lambda(p) = 0$ for all $p < p_c$ and $\lambda(p) = 1$ for all $p > p_c$. That is, there exists a phase transition in the critical value p_c, as we will see later that the behavior of percolation differs significantly for $p < p_c$ and $p > p_c$. We call the percolation is in the *supercritical phase* if $p > p_c$, and it is in the *subcritical phase* if $p < p_c$.

On the other hand, if there exists an infinite open cluster, the origin 0 is not necessarily always contained in it, and hence $\theta(p)$ can take values also in $(0, 1)$. Let $p'_c := \inf\{p : \theta(p) > 0\}$, one can show $p_c = p'_c$.

It should also be pointed out that the number of infinite open clusters is almost surely at most 1 for bond percolation on $\mathbb{Z}^d$ due to the work by Aizenman, Kesten and Newman [2]. Their result was then simplified by Burton and Keane [24], and Gandolfi, Keane and Newman [46]. A more general result about the number of infinite open clusters for transitive connected graphs can be found in [68].

The value of p_c

Apparently the value of p_c depends on the dimension d. Since an infinite open cluster can be embedded into and will stay open in higher dimensions, p_c is decreasing in d. So far the exact value of p_c is still unknown for $d \geq 3$ (see Table 1.1.3 for numerical results about p_c taken from [52, 69]). For the simplest case $d = 1$ it can be shown that $p_c = 1$. For $d = 2$ the plane square lattice it took long until the precise value of p_c was identified. In 1960 Harris [58] gave a proof that $\theta(1/2) = 0$ for $d = 2$, which implies $p_c \geq 1/2$. In 1980 Kesten proved in [66] that $p_c = 1/2$ using duality. Later, Zhang [53] gave a beautiful proof for the result $\theta(1/2) = 0$. For the other direction Duminil-Copin et al. [37] came up with a short proof in 2015. For general $d \geq 3$ the following bounds for p_c were established [23, 54]:

$$\frac{1}{2d-1} \leq p_c \leq \frac{1}{2}.$$

At the same time, some asymptotic results about p_c are known, as the dimension

dimension	3	4	5	6	7	8	9	10
p_c	0.2488	0.1601	0.1182	0.0942	0.0786	0.0677	0.0595	0.0531

Table 1.1: Numerical values of p_c for bond percolation in dimension 3-10. The results are rounded to 4 decimal places. We see p_c is decreasing in the dimension d.

d goes to infinity. In 1990s, Hara and Slade [56, 57] developed the technology known as '*lace expansion*', and obtained the asymptotics of p_c for large d:

$$p_c = \frac{1}{2d} + \frac{1}{(2d)^2} + \frac{7/2}{(2d)^3} + O\left(\frac{1}{(2d)^4}\right), \quad \text{as } d \to \infty.$$

More terms of higher orders have been identified in the physics literature for bond percolation [47, 69], as well as for site percolation [48, 69, 60].

The subcritical phase

In the subcritical phase of bond percolation there exists almost surely no infinite open cluster. In this case it is meaningful to ask how far an open path from the origin can go. More precisely, let $\partial B_n := \{x \in \mathbb{Z}^d : \|x\|_1 = n\}$. We are interested in the probability that 0 is connected with ∂B_n, that is, $\mathbb{P}(0 \leftrightarrow \partial B_n)$. In 1957 Hammersley [54] obtained the following exponential upper bound of the connection probability with branching process arguments:

$$\mathbb{P}\left(0 \leftrightarrow \partial B_n\right) \le e^{-\sigma(p)n}, \quad \text{for all } n \ge 1,$$

for some constant $\sigma(p) > 0$, if $\chi(p)$ the expected size of the open cluster containing 0 is finite. Later on, this upper bound was sharpened to an exactly exponential decay, as stated in [53]

$$\rho n^{1-d} e^{-n\phi(p)} \le \mathbb{P}(0 \leftrightarrow \partial B_n) \le \sigma n^{d-1} e^{-n\phi(p)}, \quad \text{for all } n \ge 1,$$

for $0 < p \le 1$ and some positive constants ρ, σ independent of p and some function ϕ.

The supercritical phase

In the supercritical phase of bond percolation there exists almost surely an infinite open cluster. Nevertheless we can ask a similar question as in the subcritical phase about the decay of connection probability $\mathbb{P}(0 \leftrightarrow \partial B_n)$ provided the open cluster containing the origin is finite. Grimmett gave an asymptotic tail behavior of the connection probability in [53], and we present it here:

$$\mathbb{P}(0 \leftrightarrow \partial B_n, |C(0)| < \infty) \leq A(p,d) n^d e^{-n\sigma(p)}, \quad \text{for all } n \geq 1,$$

where $C(0)$ is the open cluster that contains 0, A is a positive constant depending on p and d, and σ is a positive function of $p \in (p_c, 1]$. Furthermore, Aizenman et al. [1] gave an sub-exponential estimate for the exact size of the cluster $C(0)$:

$$\mathbb{P}(|C(0)| = n) \geq e^{-\chi(p) n^{(d-1)/d}}, \quad \text{for all } n \geq 1,$$

if $p \in (p_c, 1)$, where $\chi(p) \in (0, \infty)$ is a constant depending on p.

At criticality

It is a central question in percolation theory whether the infinite open cluster exists at criticality. Since long it is conjectured that $\theta(p_c) = 0$ for $d \geq 2$. As mentioned before, for $d = 2$, this conjecture was confirmed by Harris et. al [58, 66, 53]. With help of lace expansion, Hara and Slade [56, 57] managed to show $\theta(p_c) = 0$ for $d \geq 19$, and this result was extended to $d \geq 11$ in 2015 by Fitzner and van der Hofstad [44].

Results about the behavior of graphs at criticality for other similar models also suggest that this conjecture is very likely to be true. For example $\theta(p_c) = 0$ is verified for bond percolation on $\mathbb{Z}^d \times \mathbb{Z}^+$ by Barsky, Grimmett and Newman [6], and for $\mathbb{Z}^d \times G$ with an arbitrary finite connected graph G by Duminil-Copin, Sidoravicius and Tassion [36].

For $d \geq 2$, van der Berg and Keane [8] showed the function $\theta(p)$ is continuous for all $p \neq p_c$, and $\theta(p)$ is continuous in p_c if and only if $\theta(p_c) = 0$. So if the conjecture above holds true, θ would be continuous for all $p \in [0, 1]$.

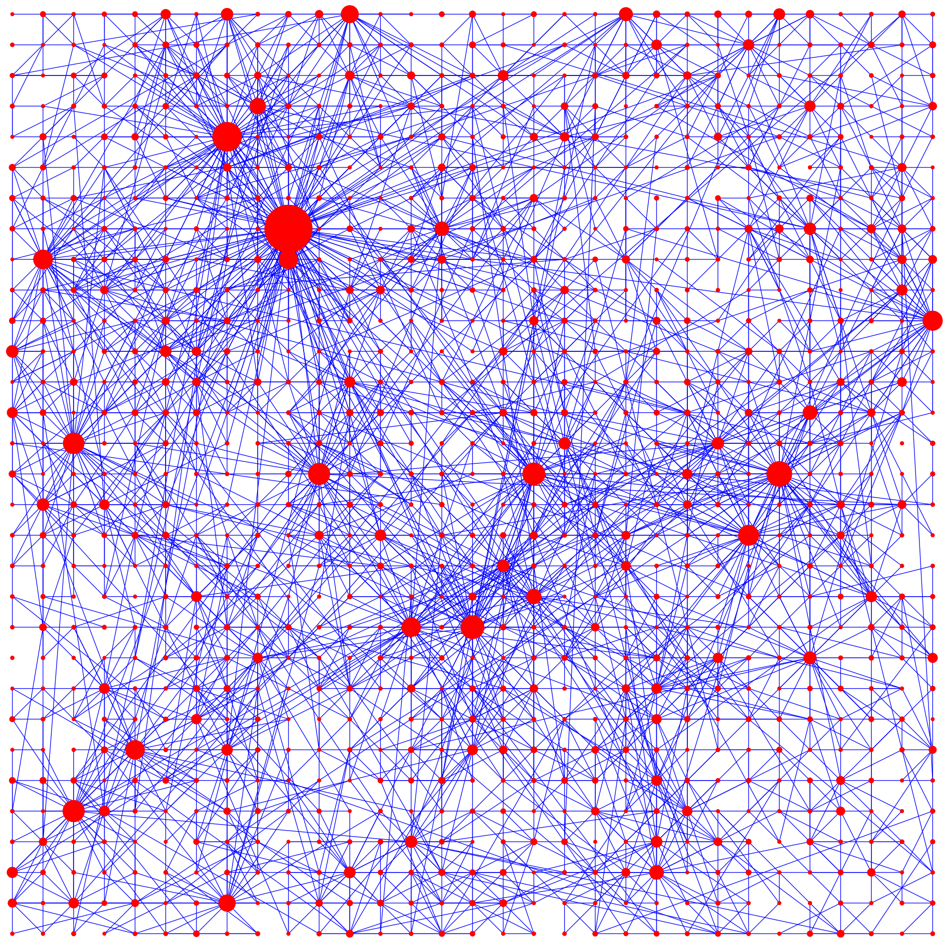

Figure 1.3: Scale-free percolation for $\lambda = 0.2, \tau = 2.5, \alpha = 3$. The radii of red balls are proportional to the square root of corresponding weights.

1.2 Scale-free percolation

The model

Now we introduce the main object of this dissertation: scale-free percolation model (also known as *heterogeneous long-range percolation*), which we henceforth abbreviate as SFP. We consider the lattice $\mathbb{Z}^d$ with fixed dimension $d \geq 1$ and construct a random subgraph of the complete graph on the vertex set $\mathbb{Z}^d$. To each vertex $x \in \mathbb{Z}^d$, we assign an i.i.d. weight W_x which follows a power-law distribution with parameter $\tau - 1$ ($\tau > 1$), that is,

$$\mathbb{P}(W_x \geq w) = w^{-(\tau-1)}, \quad w \geq 1. \tag{1.4}$$

Conditioning upon these weights, we declare an edge $\{x, y\}$ to be *open* independently of the status of other edges with probability

$$p_{x,y} = \frac{\lambda W_x W_y}{|x-y|^\alpha} \wedge 1, \tag{1.5}$$

where $|\cdot|$ denotes the Euclidean norm and $\alpha, \lambda > 0$ are further parameters of the model. Here $x \wedge y$ means the minimum of x and y. One example of scale-free percolation is illustrated in Figure 1.3. We write $x \sim y$ if the edge $\{x, y\}$ is open.

Note that other choices of connection probability are possible. For example, an alternative is

$$p_{x,y} = 1 - \exp\left(-\frac{\lambda W_x W_y}{|x-y|^\alpha}\right).$$

Unless specifically mentioned otherwise, we use (1.5) for the connection probability in scale-free percolation throughout this dissertation.

Recall a random graph is called *scale-free*, if its degree distribution follows a power-law asymptotically. In 2013 Deijfen et al. showed that the degree of vertices in SFP satisfies a power-law distribution with tail exponent

$$\gamma := \frac{\alpha(\tau-1)}{d}, \tag{1.6}$$

as we present here:

Theorem 1.1 (Theorem 2.2 in [28]). *Assume that the weight distribution in* (1.4) *satisfies $\alpha > d$ and $\gamma = \alpha(\tau-1)/d > 1$. Then there exists a function ℓ which is slowly varying at infinity such that*

$$\mathbb{P}(D_x \geq k) = k^{-\gamma}\ell(k), \quad k \in \mathbb{N},$$

where D_x is the degree of $x \in \mathbb{Z}^n$.

Therefore scale-free percolation really matches its name.

Similar to Bernoulli bond percolation, the existance of the (unique) infinite cluster is also of great interest. As they introduced this model, Deijfen et al. investigated the critical value of λ, which is defined as

$$\lambda_c := \inf\{\lambda : \theta(\lambda) > 0\}, \tag{1.7}$$

where $\theta(\lambda)$ is the probability that 0 is in the unique infinite open cluster.

Theorem 1.2 (Theorem 3.1, 4.1, 4.2 and 4.4 in [28]). *Depending on the parameters, we have following results for λ_c:*

1. *Finiteness of the critical value:*

   (a) *If $d = 1, \alpha \in (1, 2]$, then $\lambda_c < \infty$;*

   (b) *If $d \geq 2$, then $\lambda_c < \infty$.*

2. *Positivity of the critical value:*

   (a) *If $\tau > 1, \alpha > d, \gamma < 2$, then $\lambda_c = 0$;*

   (b) *If $\tau > 1, \gamma > 2$, then $\lambda_c > 0$;*

   (c) *If $\tau > 3$, then $\lambda_c > 0$.*

Graph distances

Graph distances in real-world networks, in particularly social networks, have been in the focus of network research since Milgram's experimental discovery of the small-world effect (casually phrased as "six degrees of separation"), and have also been investigated theoretically since then, e.g. [67, 70]. For a graph, the graph distance between two vertices is defined as the length of a shortest open path connecting them. If the vertices lie in different clusters (and hence such open paths do not exist), then the graph distance is ∞.

For graph distances in scale-free percolation, a rich phase diagram has been established in the literature: for two distinct vertices $x, y \in \mathbb{Z}^d$, we denote by $D(x, y)$ the graph distance between x and y. Then, conditional on x and y to be in the (unique) infinite cluster, we get that with high probability (as $|x - y| \to \infty$)

- if $\gamma \leq 1$ then $D(x, y) \leq 2$, cf. [59];
- if $\alpha < d$, then $D(x, y) \leq \lceil d/(d - \alpha) \rceil$, cf. [59];
- if $\gamma \in (1, 2)$ and $\alpha > d$, then $D(x, y) = \frac{2}{|\log(\gamma-1)|} \log \log |x - y|$, cf. [28, 64];
- if $\gamma > 2$ and $\alpha > 2d$, then $D(x, y) \gtrsim |x - y|$, cf. [29, 72].

This behaviour (together with our new results) is summarized in Figure 1.4. The results in the first three cases are referred as "ultra-small world" phenomenon, because the asymptotics are of smaller order than the requirements of (1.2). In these regimes, shortest paths are typically formed by vertices that have the highest weight in a certain neighborhood (*locally dominanting vertices* or *hubs*). In contrast, for

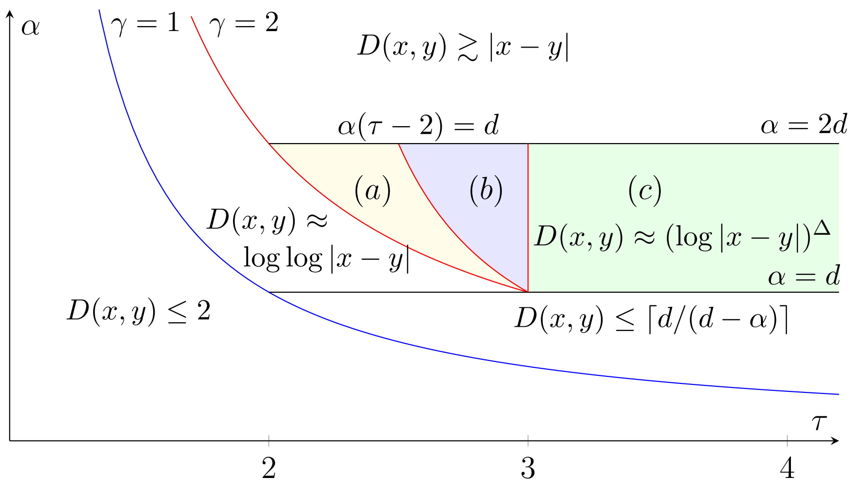

Figure 1.4: Graph distances in different regimes of scale-free percolation. The regions in shadow are those we are interested in. The areas (a), (b) and (c) represent our improved bounds established in Theorem 1.3.

$d < \alpha < 2d$ and $\gamma > 2$, the weights are more homogeneous, and it is not sufficient to consider only dominant vertices to find the shortest paths. In this regime, there is a fine interplay between weights and spatial positions of various vertices, which leads to (poly-)logarithmic upper and lower bounds on graph distances. One goal of this dissertation is to identify the right logarithmic power, thereby completing the phase diagram.

At the phase boundaries ($\gamma = 1$ and $\gamma = 2$) we expect that the graph distances depend on the precise tail behavior of the connectivity function in (1.5), so that any universality is lost.

Navigation

After we identify graph distances in SFP, a natural subsequent question will be about the *navigation possibility*. More precisely, let s and t be two arbitrary vertices in $\mathbb{Z}^d$. Is there any algorithm that finds a path between s and t of comparable length as the shortest path using only local information? In other words, we focus on algorithms with the following mechanism:

1. Information about the start s and the target t is given;
2. When the algorithm reaches some vertex x, the choice of next hop is made based on the information of x's neighbors (so called "local information").

In our context, local information includes locations and weights. In some references, for example [67, 35], such algorithm using only local information is also called a "*decentralized algorithm*", and henceforth we will also use this terminology.

Let $X_{s,t}$ be the number of steps a decentralized algorithm needs to find t from s. We aim at such algorithms with

$$X_{s,t} \approx D(s,t). \tag{1.8}$$

as $|s-t| \to \infty$. A random graph G is called *navigable*, if a decentralized algorithm exists such that (1.8) holds in the sense that $X_{s,t} = (1+o(1))D(s,t)$ with high probability. A weaker definition of navigability requires only that $X_{s,t} = \mathcal{O}\left(D(s,t)\right)$ when $|s-t|$ goes to infinity. As the reader will see in Section 1.3, scale-free percolation in the doubly logarithmic regime is navigable in the strong sense.

If we have global information about the random graph, without doubt the shortest path can be found between s and t. However, it is not always the case if we have only local information. In other words, not all random graphs are navigable. In 2000, Kleinberg showed in [67] that some finite small-world network on the 2-dimensional lattice is not navigable. Later on, Franceschetti and Meester [45], Draief and Ganesh [35] extended the results to continuum setting with Poisson points. In all the models they considered the local information contains only locations, which is a major difference to our model.

For scale-free percolation, as we will see later in Section 1.3 and Section 3.1, any decentralized algorithm fails to find the shortest paths if $\gamma > 2$ and $\alpha \in (d, 2d)$ in the sense that (1.8) are not satisfied in either strong or weak sense.

In practice, many algorithms have been proposed to solve the navigation problem on graphs. One of the frequently used decentralized algorithms is *greedy routing*. For a greedy routing algorithm we need an objective function ϕ with the local information as its input. Then the algorithm selects the neighbor with highest objective. More precisely, we have the following protocol for greedy routing:

Routing protocol: Given an objective function ϕ, when the algorithm is in x, it will jump to the neighbor of x with highest objective in the next step, and the objective of this neighbor must be larger than that of x. If such neighbor of x does not exist, T will abort.

It is easily seen that a typical greedy routing path consists of vertices with strictly increasing objectives. Mind that if greedy routing enters some local maxi-

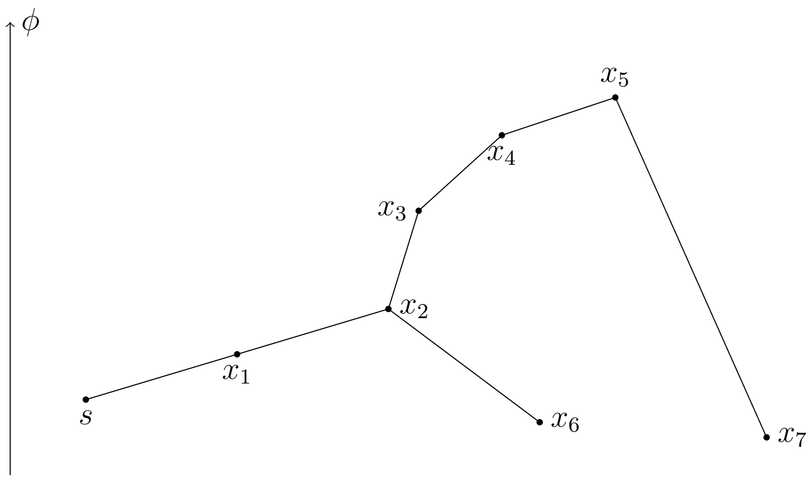

Figure 1.5: Example of running both greedy protocol and patching protocol on a graph.

mum of objective, it will abort and hence fail. To avoid failure, a patching method is proposed and will be discussed in details in Section 3.2.3. We first state it here:

Patching protocol: When the algorithm arrives at some local maximum, among all the unvisited neighbors of visited vertices, it will go to the one with highest objective and perform greedy routing from there.

Note that the patching protocol allows the algorithm to go backwards and enables it to circumvent the deadlock of local maximums. For more general conditions about patching, we refer to [20, 21].

Now we illustrate the greedy routing protocol and patching protocol with help of Figure 1.5. In the example in Figure 1.5, by the greedy routing protocol the greedy algorithm T will explore the graph from s in the following order: $s, x_1, x_2, x_3, x_4, x_5$. When T arrives at x_5, which is a local optimum, it will be trapped there and fail. Additionally with patching protocol T visits the unexplored neighbors of visited nodes $\{s, x_1, x_2, x_3, x_4, x_5\}$. In this graph the unexplored vertices are x_6 and x_7. Since $\phi(x_6) > \phi(x_7)$, T will go backwards and visit x_6. Consequently the order in which T explores the graph is: $\boldsymbol{s}, \boldsymbol{x_1}, \boldsymbol{x_2}, \boldsymbol{x_3}, \boldsymbol{x_4}, \boldsymbol{x_5}, x_4, x_3, x_2, \boldsymbol{x_6}, x_2, x_3, x_4, x_5, \boldsymbol{x_7}$. Here bold means this vertex is visited for the first time in the route.

As we can see, if T has explored k vertices, two scenarios can happen:

(1) The current vertex x has at least a good neighbor in the sense that this neighbor has a strictly larger objective than x. In this case it takes T only one step to explore a new vertex;

(2) The current vertex x is a local optimum in the sense that x has a larger objetive than all its neighbors. With the patching protocol T can go back. After at most k steps T will uncover a new vertex.

Therefore, with the patching protocol, it takes T at most $\sum_{k=1}^{n}(k-1) = \frac{n(n-1)}{2}$ to explore n different vertices. This result will be useful in Section 3.2.3.

Since it only makes sense to talk about navigation possibilities if the start and target are in the same open cluster, we assume the presence of edges between nearest neighbors throughout Chapter 3. Besides, for simplicity it is sufficient to consider the connection probability in the following form:

$$p_{xy} = \frac{W_x W_y}{|x-y|^\alpha} \wedge 1, \tag{1.9}$$

and we will use (1.9) throughout Chapter 3. Note that the connection probabilities in (1.9) already guarantee the presence of nearest neighbor edges, since $W_x \geq 1$ and $W_y \geq 1$. Therefore we can omit the percolation parameter λ in the connection probabilities.

In spirit of Milgram's experiment, a natural choice of objective function for scale-free percolation will be the following:

$$\phi(x) := \frac{W_x}{|x-t|^\alpha}, \quad x \in \mathbb{Z}^d. \tag{1.10}$$

If the edge between s and t is present, then the algorithm should explore t directly in the first hop, requiring that t should be the global maximizer of the objective function. Apparently our choice in (1.10) satisfies this requirement. In Section 3.2 we will stick to this choice of objective function.

For the other forms of connection probability, the corresponding choices of objective functions are discussed in Remark 3.6.

1.3 Main results

Before we proceed to the main results, we need to introduce some parameters.

Parameters

$$\alpha_1 := \alpha \wedge \frac{\alpha(\tau-1)}{2} = \alpha \wedge \frac{\gamma d}{2}, \quad \alpha_2 := \alpha \wedge (\alpha(\tau-1)-d) = \alpha \wedge (\gamma-1)d, \tag{1.11}$$

$$\Delta := \frac{\log 2}{\log(2d/\alpha)}, \quad \Delta_1 := \frac{\log 2}{\log(2d/\alpha_1)}, \quad \Delta_2 := \frac{\log 2}{\log(2d/\alpha_2)}. \tag{1.12}$$

Here $x \wedge y$ means the minimum of x and y.

If γ in (1.6) is larger than 2, then

$$d < \alpha_1 \leq \alpha_2 \leq \alpha < 2d.$$

As a consequence

$$1 < \Delta_1 \leq \Delta_2 \leq \Delta.$$

As showed in Theorem 1.2 that for $d < \alpha < 2d$ and $\gamma > 1$ the critical value λ_c of SFP is finite. We thus may condition on two vertices x and y to be in the same infinite cluster, if we take $\lambda > \lambda_c$. As before, let $D(x, y)$ be the graph distance between x and y, then it holds

Theorem 1.3. *For scale-free percolation with parameters $\lambda > \lambda_c, \gamma > 2$, and $d < \alpha < 2d$, we have that for any $\epsilon > 0$,*

$$\lim_{|x-y|\to\infty} \mathbb{P}\left((\log|x-y|)^{\Delta_1-\epsilon} \leq D(x,y) \leq (\log|x-y|)^{\Delta_2+\epsilon} \,\middle|\, x, y \in \mathcal{C}_\infty\right) = 1.$$

Depending on the value of γ and α, the various minima in (1.11) give rise to three different regimes. These are depicted in Figure 1.4. Writing $\mathcal{C}_\infty$ for the unique infinite cluster in the graph, we get

(a) for $\gamma > 2$, $\alpha(\tau-2) < d$ and arbitrary $\epsilon > 0$,

$$\lim_{|x-y|\to\infty} \mathbb{P}\left((\log|x-y|)^{\Delta_1-\epsilon} \leq D(x,y) \leq (\log|x-y|)^{\Delta_2+\epsilon} \,\middle|\, x, y \in \mathcal{C}_\infty\right) = 1;$$

(b) for $\tau < 3$, $\alpha(\tau-2) \geq d$ and arbitrary $\epsilon > 0$,

$$\lim_{|x-y|\to\infty} \mathbb{P}\left((\log|x-y|)^{\Delta_1-\epsilon} \leq D(x,y) \leq (\log|x-y|)^{\Delta+\varepsilon} \,\middle|\, x, y \in \mathcal{C}_\infty\right) = 1;$$

(c) for $\tau \geq 3$ and arbitrary $\epsilon > 0$,

$$\lim_{|x-y|\to\infty} \mathbb{P}\left((\log|x-y|)^{\Delta-\epsilon} \leq D(x,y) \leq (\log|x-y|)^{\Delta+\varepsilon} \,\Big|\, x,y \in \mathcal{C}_\infty\right) = 1.$$

Note that here the upper bounds in Part (b) and (c) are from [29].

The main results in Theorem 1.3 about graph distances have been published in [55]. Theorem 1.3 basically tells that with high probability the graph distance in SFP is *poly-logarithmic* in Euclidean distance in the prescribed regime. Despite the improvements in both the upper and lower bounds, the reader may observe that there is still a gap between them in case (a) and (b) in our result. Therefore, it remains open as to what the correct exponent is. The main difficulty in closing the gap between the upper and lower bounds is that we do not have a precise estimate for the probability of a path being open in scale-free percolation. Lemma 2.6 gives a nice upper bound. However, in view of Proposition 2.17, it appears that this bound is not optimal for $\tau < 3$. As shown in Proposition 2.17 as well as in Lemma 2.20, the actual asymptotics of the probability of a path being open in SFP are heterogeneous in the exponents of edges, which poses a great difficulty.

Remark 1.4. *In this dissertation, we made a specific choice for the connection probability in* (1.5)*. In fact, our methods also apply to more general forms of connection probabilities. The proofs for both lower and upper bounds in Section 2.1 and Section 2.2 only require asymptotics of the connection probability to estimate the path, for example in Lemma 2.4 and Proposition 2.17. Therefore, our results generalise to the scale-free percolation with connection probability* $p_{x,y} = \Theta\left(\frac{\lambda W_x W_y}{|x-y|^\alpha} \wedge 1\right)$ *provided that a unique infinite cluster exists.*

If we make the extra assumption that *additionally all nearest-neighbour edges are open*, then a comparison with long-range percolation (explained in the following paragraph) gives the following improvement to parts (b) and (c) above: there exists $C > 0$ such that

$$\lim_{|x-y|\to\infty} \mathbb{P}\left(D(x,y) \leq C\,(\log|x-y|)^{\Delta}\right) = 1. \tag{1.13}$$

Mind that the extra assumption ensures that $x, y \in \mathcal{C}_\infty$.

Now we state our results about navigability of scale-free percolation. In 2000 Kleinberg showed in [67] that some graph on the lattice is not navigable in the sense that any decentralized algorithm on the graph needs at least polynomially many steps to find the target, while the theoretical graph distance between the start and the target is only poly-logarithmic in the Euclidean distance. Later on, this result was extended to other models whose local information includes only locations [35, 45]. In our model, the local information does not only contain locations of neighbors, but also their weights. We manage to show the following result about navigability:

Let $X^T_{s,t}$ be the number of steps a decentralized algorithm T starting from s needs to find t and denote by N the Euclidean distance between s and t, then

Theorem 1.5. *Consider scale-free percolation with connection probability* $p_{x,y} = \frac{\lambda W_x W_y}{|x-y|^\alpha} \wedge 1$*, and parameters* $\alpha \in (d, 2d), \gamma > 2$*. Let* T *be an arbitrary decentralized algorithm. Then there exists a constant* $\delta > 0$ *such that*

$$\lim_{N\to\infty} \mathbb{P}(X^T_{s,t} \geq N^\delta) = 1.$$

Recall a decentralized algorithm is an algorithm that uses only the local information of neighbors. Theorem 1.5 tells us that in the poly-logarithmic regime that any decentralized algorithm is inefficient in the sense that it takes much more steps than the theoretical number to find the target.

In the heavy-tailed regime we show that there exists a decentralized algorithm that finds the shortest paths. In other words, scale-free percolation is navigable in this regime. We state the results as follows:

Theorem 1.6. *Consider scale-free percolation with connection probability* (1.9)*, and parameters* $\alpha > d, \gamma \in (1, 2)$*. Let* T *be the greedy routing algorithm with objective function as in* (1.10)*. Furthermore, conditional on* W_s *and* W_t*, let* L_1, L_2, L_3 *be functions of* N *such that*

$$L_i = \frac{1 + o(1)}{|\log(\gamma - 1)|} \Big(\log\log_{W_s} \left(\phi(s)^{-1}\right) + \log\log_{W_t} \left(\phi(s)^{-1}\right) \Big), \quad i = 1, 2, 3. \tag{1.14}$$

Then as $N \to \infty$*,*

(a) with at least positive constant probability, T *finds the target within* L_1 *steps;*

(b) with high probability, T *terminates after at most* L_2 *steps;*

(c) with high probability, the greedy algorithm with patching protocol finds t *within* L_3 *steps.*

Although $L_i, i = 1,2,3$ look all the same in (1.14), they differ in fact in the $o(1)$ term. We will see the exact expressions of $L_i, i = 1,2,3$ at the end of Section 3.2.1, 3.2.2 and 3.2.3 respectively.

Since $\phi(s) = W_s/N^\alpha$, we know $L_i \approx \frac{2}{|\log(\gamma-1)|}\log\log N$, as $N \to \infty$. This length coincides with the graph distance in the doubly logarithmic regime obtained in [28, 64]. In this sense the greedy routing algorithm indeed finds the shortest path, and therefore scale-free percolation is navigable in the doubly logarithmic regime.

Remark 1.7. *In Theorem 1.5 our result is for the choice of connection probability (1.9). Similar to the results about graph distances, we expect that Theorem 1.5 holds true for scale-free percolation with connection probabilities in the more general form* $p_{x,y} = \Theta\left(\frac{\lambda W_x W_y}{|x-y|^\alpha} \wedge 1\right)$.

1.4 Related models

In this section we discuss several random graph models that are closely related to scale-free percolation.

1.4.1 Long-range percolation

We first introduce a related (though easier) model named *long-range percolation*. Our analysis of graph distances in scale-free percolation is crucially based on techniques developed for long-range percolation.

Long-range percolation (henceforth LRP) is also defined on the lattice $\mathbb{Z}^d$ for fixed dimension $d \geq 1$. Independently of all the other edges, the edge $\{x, y\}$ is open with probability p^{LRP}_{xy}. A typical choice of p^{LRP}_{xy} is

$$p^{\mathrm{LRP}}_{xy} = \frac{\lambda}{|x-y|^\alpha} \wedge 1.$$

Note that p^{LRP}_{xy} is equal to p_{xy} for scale-free percolation (as defined in (1.5)) if $W_x \equiv 1$ or $\tau = \infty$. One example of long-range percolation can be found in Figure 1.6.

Biskup et al. studied the graph distances in long-range percolation and obtained sharp results.

Theorem 1.8 (Biskup [12], Trapman [76], Biskup-Lin [14]). *Consider the long-range* *percolation with connection probability* $\{p_{xy}\}$ *such that*

$$\liminf_{|x-y|\to\infty} p_{xy}^{\text{LRP}}|x-y|^{\alpha} > 0, \tag{1.15}$$

for some $\alpha > 0$*. If* $d < \alpha < 2d$ *and a unique infinite open cluster exists, then for all* $\epsilon > 0$ *one has*

$$\lim_{|x-y|\to\infty} \mathbb{P}\left((\log|x-y|)^{\Delta-\epsilon} \leq D(x,y) \leq (\log|x-y|)^{\Delta+\epsilon} \,\middle|\, x,y \in \mathcal{C}_\infty \right) = 1.$$

If, moreover, we have the stronger form of connection probability

$$p_{xy}^{\text{LRP}} = \frac{\lambda}{|x-y|^{\alpha}} \wedge 1,$$

and assume the existence of all nearest-neighbor edges, then there exist constants $C > c > 0$ *such that*

$$\lim_{|x-y|\to\infty} \mathbb{P}\left(c\,(\log|x-y|)^{\Delta} \leq D(x,y) \leq C\,(\log|x-y|)^{\Delta} \right) = 1.$$

Trapman [76], moreover, identified the growth of the balls $\{x \in \mathbb{Z}^d : D(0,x) \leq n\}$ for LRP with $d < \alpha < 2d$.

Now we can describe a coupling between LRP and SFP. To this end, we view the two models from another perspective: to each edge $\{x,y\}$ of the graph, we assign an i.i.d. Uniform$[0,1]$-distributed random variable U_{xy}. Then, for scale-free percolation model, we consider for each edge $\{x,y\}$ the event

$$A_{x,y} := \left\{ U_{xy} \leq \frac{\lambda W_x W_y}{|x-y|^{\alpha}} \wedge 1 \right\},$$

and we make the edge $\{x,y\}$ open whenever $A_{x,y}$ occurs. In the same way, for long-range percolation we consider the event

$$B_{x,y} := \left\{ U_{xy} \leq \frac{\lambda}{|x-y|^{\alpha}} \wedge 1 \right\}.$$

We have thus constructed a coupling for the two models: since $W_x \geq 1$ for all $x \in \mathbb{Z}^d$, we have

$$\frac{\lambda}{|x-y|^{\alpha}} \wedge 1 \leq \frac{\lambda W_x W_y}{|x-y|^{\alpha}} \wedge 1,$$

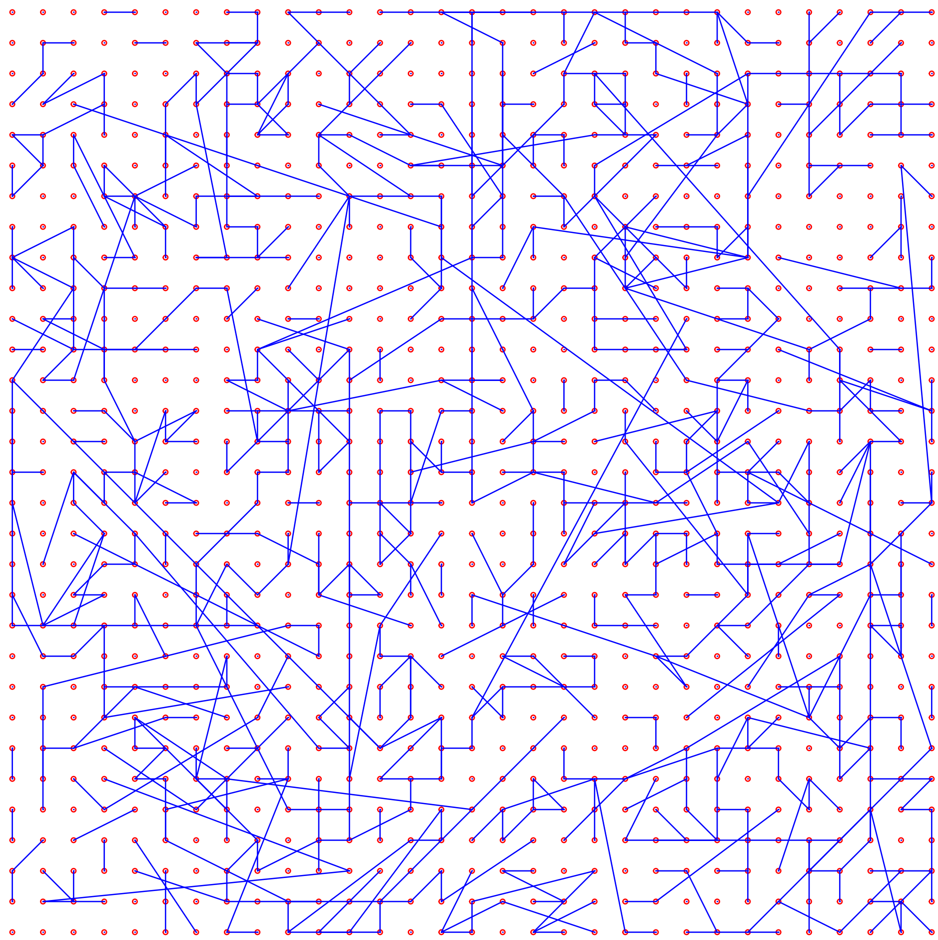

Figure 1.6: Long range percolation for $\lambda = 0.2, \alpha = 3$.

which implies $A_{x,y} \supseteq B_{x,y}$, thus scale-free percolation dominates long-range percolation in the sense that all the open edges in the LRP remain open in SFP. We therefore get that distances in LRP are an upper bound for distances in SFP and in particular get the upper bound (1.13).

For the remaining regimes, there are many rigorous results about the graph distance $D(x, y)$ as $|x - y| \to \infty$. When $\alpha < d$, Benjamini, Kesten, Peres and Schramm [7] show that $D(x, y)$ is bounded by some (explicit) constant. When $\alpha > 2d$, Berger [9] shows that $D(x, y) \geq |x - y|$. For the borderline case $\alpha = 2$ for $d = 1$, Ding and Sly [34] show that $D(x, y) \approx |x - y|^{\delta}$ for some $\delta \in (0, 1)$.

Besides, as a corollary of Theorem 1.5 we show that long-range percolation is not navigable if $d < \alpha < 2d$. We will come to this in Corollary 3.4.

1.4.2 Geometric inhomogeneous random graph

Geometric inhomogenneous random graph (henceforth GIRG) is a spatial random graph on some finite domain of $\mathbb{R}^d$ for some fixed dimension $d \geq 1$. More precisely, consider the d-dimensional torus $\mathbb{T}^d := \mathbb{R}^d/\mathbb{Z}^d$, which can be viewed as the d-dimensional unit cube $[0,1]^d$ with all opposite faces identified. Note the distance between $x, y \in \mathbb{T}^d$ can be written for example as

$$|x-y| := \sum_{i=1}^{d} \min\{|x_i - y_i|, 1 - |x_i - y_i|\},$$

with $x = (x_i)$ and $y = (y_i)$. The vertex set V of GIRG is sampled by a homogeneous Poisson process on $\mathbb{T}^d$ with intensity n.

Analogous to scale-free percolation, each vertex x in GIRG is assigned with i.i.d weight W_x satisfying a power law as in (1.4):

$$\mathbb{P}(W_x \geq w) = w^{-(\tau-1)}, \quad w \geq 1.$$

Given locations and weights, two vertices x, y in $\mathbb{T}^d$ are connected according to the following probability independently:

$$p_{x,y}^{GIRG} := \Theta\left(\frac{W_x W_y}{n|x-y|^{\alpha}} \wedge 1\right).$$

By proper rescaling some induced subgraph of the GIRG model can be viewed as the scale-free percolation in continuum in a finite domain [21]. In [19] Bringmann et al. showed that the graph distance is asymptotically $\frac{2+o(1)}{|\log(\gamma-1)|} \log\log n$ in the GRIG model, if $\gamma := \frac{\alpha(\tau-1)}{d} \in (1,2)$. Later on in [21], a greedy routing algorithm was proposed in order to find short paths between two vertices, and it turned out that with high probability this algorithm finds the target within also $\frac{2+o(1)}{|\log(\gamma-1)|} \log\log n$ steps, if it is so patched that it can circumvent the local optimum in the route.

As the readers will see, the greedy routing algorithm for scale-free percolation proceeds analogously as for the GIRG model in [21]. Here we point out the major differences between the algorithms for both models.

First, the vertex set of GIRG is generated by a homogeneous Poisson process, meaning that all vertices have random locations. In contrast, scale-free percolation has the deterministic vertex set $\mathbb{Z}^d$. This allows us to assume the presence of all nearest-neighbor edges to make sure the start and the target are in the same cluster.

Besides, since GIRG is defined on a finite domain, by the property of Poisson processes, it is a finite graph (almost surely). Consequently the patched algorithm finds the target within finitely many steps if the start and the target are in the same cluster. However, this is not the case for scale-free percolation, even if the start and the target are in the unique infinite cluster. Fortunately, as we will see in Section 3.2, this possibility can be ruled out for scale-free percolation.

1.4.3 Other related models

Various aspects of scale-free percolation have been investigated in the literature, both on the lattice $\mathbb{Z}^d$ [29, 59] as well as a continuum analogue [26, 30], where vertices are given as a Poisson point process. The results in the present dissertation have been obtained on $\mathbb{Z}^d$, but it appears that we do not make use of the lattice structure in any crucial way, so that analogue results should hold for a continuum version of the model.

It has been pointed out recently by Gracar et al. [50, 51] that scale-free percolation (in continuum), as well as many other random graphs models, can be understood as special cases of the *weight-dependent random connection models*. In the language of [50], scale-free percolation corresponds to the weight-dependent random connection model with *product kernel* and polynomial profile function. Mind that the parametrization in [50] is different, see in particular [50, Table 2].

For related recent work on spatial preferential attachment graphs we refer to the work by Hirsch and Mönch [61].

Chapter 2

Graph distances in scale-free percolation

In this chapter we prove both upper bound and lower bound in Theorem 1.3, and discuss the possibility of filling the gap between them for $\tau \in (2,3)$.

2.1 Lower bound for graph distances

Theorem 2.1 (Lower bound in Theorem 1.3)**.** *For scale-free percolation with parameters $\lambda > \lambda_c, \gamma > 2$, and $d < \alpha < 2d$, we have that for any $\epsilon > 0$,*

$$\lim_{|x-y|\to\infty} \mathbb{P}\left(D(x,y) \geq (\log|x-y|)^{\Delta_1 - \epsilon} \,\middle|\, x, y \in \mathcal{C}_\infty\right) = 1.$$

Here $D(x,y)$ is the graph distance between x and y, and $\mathcal{C}_\infty$ is the unique infinite cluster. The parameter Δ_1 is given by

$$\Delta_1 := \frac{\log 2}{\log(2d/\alpha_1)} \quad \textit{with} \quad \alpha_1 := \alpha \wedge \frac{\alpha(\tau-1)}{2} = \alpha \wedge \frac{\gamma d}{2}.$$

In order to prove the lower bound, we derive variants of Biskup's arguments [12] in the setting of scale-free percolation. Similar to [12], we split up the argument into 3 propositions.

The key difference between SFP and LRP is that adjacent edges in the former model are only *conditionally* independent. We resolve this by adjusting the definition of a *hierarchy* (below) and combine it with estimates from [28] to break up the dependence structure.

Definition 2.2. *Given an integer $n \geq 1$ and distinct vertices $x, y \in \mathbb{Z}^d$, we say that the collection*

$$\mathcal{H}_n(x,y) = \{(z_\sigma) : \sigma \in \{0,1\}^k, k = 1, 2, \ldots, n; z_\sigma \in \mathbb{Z}^d\}$$

is a hierarchy of depth n connecting x and y if

1. *$z_0 = x$ and $z_1 = y$;*
2. *$z_{\sigma 00} = z_{\sigma 0}$ and $z_{\sigma 11} = z_{\sigma 1}$ for all $k = 0, 1, \ldots, n-2$ and all $\sigma \in \{0,1\}^k$;*
3. *For all $k = 0, 1, \ldots, n-2$ and all $\sigma \in \{0,1\}^k$ such that $z_{\sigma 01} \neq z_{\sigma 10}$, the edge $\{z_{\sigma 01}, z_{\sigma 10}\}$ is open;*
4. *Each edge $\{z_{\sigma 01}, z_{\sigma 10}\}$ as specified in part 3 appears only once in $\mathcal{H}_n(x,y)$;*
5. *For $z_{\sigma_1}, z_{\sigma_2}$ in $\mathcal{H}_n(x,y)$ with $k \in \{0, 1, \ldots, n-1\}$, $\sigma_1, \sigma_2 \in \{0,1\}^{k+1}$ and $\sigma_1 \neq \sigma_2$ we have that $z_{\sigma_1} = z_{\sigma_2}$ if and only if there exists $\sigma \in \{0,1\}^k$ such that $\sigma_1 = \sigma 0$ and $\sigma_2 = \sigma 1$. In this case, we call the vertices z_{σ_1} and z_{σ_2} degenerate, otherwise non-degenerate.*

The vertices (z_σ) are called sites of the hierarchy $\mathcal{H}_n(x,y)$.

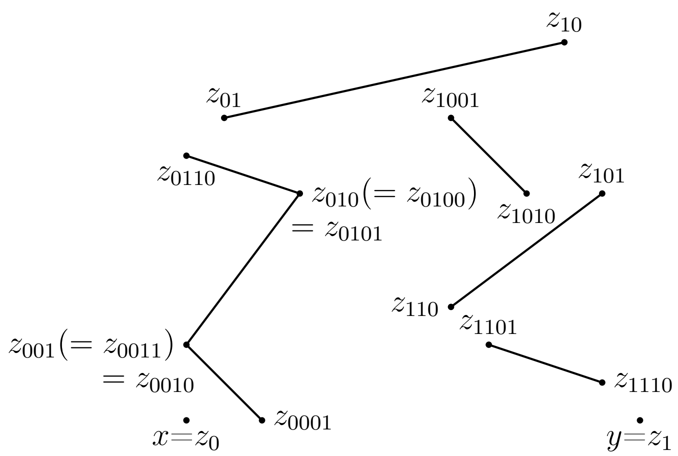

Figure 2.1: A hierarchy of depth 4 with two degenerate sites z_{001} and z_{010}

In the toy example depicted in Figure 2.1, the reader finds two overlapping sites. For $z_{001}(= z_{0011})$ and z_{0010}, there exists $\sigma = (0, 0, 1) \in \{0,1\}^3$ such that $z_{\sigma 1} = z_{\sigma 0}$. Therefore, this is a degenerate site in the sense of Condition 5. Similarly for z_{010} and z_{0101}.

Remark 2.3. *With only Conditions 1-4, our definition would coincide with Definition 2.1 in [12]. In addition, we impose Condition 5 to make sure that every element* $(z_\sigma) \in \mathcal{H}_n(x, y)$ *can be fitted into a vertex self-avoiding path connecting* x *and* y*. By adding an additional condition, one realises the set of all hierarchies here is a subset of hierarchies defined in [12], and this will be helpful when we count the eligible hierarchies.*

The hierarchy $\mathcal{H}_n(x, y)$ *is essentially a (random) subgraph of the complete graph with vertex set* $\mathbb{Z}^d$*. Condition 4 ensures that the number of open edges in this subgraph is at most* 2^{n-1}*, and Condition 5 guarantees that the degree of all vertices in* $\mathcal{H}_n$ *is no more than 2.*

Since the shortest path connecting x and y is necessarily vertex self-avoiding, meaning that the weight of a single vertex appears at most twice in the path, we can estimate the probability of such a path by the Cauchy-Schwarz inequality.

Lemma 2.4 ([28, Lemma 4.3])**.** *Let* $x, y \in \mathbb{Z}^d$ *be distinct, then for all* $\delta > 0$*, there exists a constant* $C_\delta := C(\delta, \lambda) > 1$ *such that*

$$\mathbb{E}\left[\left(\lambda \frac{W_x W_y}{|x-y|^\alpha} \wedge 1\right)^2\right]^{1/2} \leq C_\delta |x-y|^{-\alpha_1+\delta}, \tag{2.1}$$

where α_1 *is defined as in* (1.11)*.*

Proof. From the proof of Lemma 4.3 in [28] we know

$$\mathbb{E}\left[\left(\lambda \frac{W_x W_y}{|x-y|^\alpha} \wedge 1\right)^2\right] \leq C_1 \left(1 + \log|x-y|\right) |x-y|^{-2\alpha_1}.$$

for some constant $C_1 \in (0, \infty)$. Then for all $\delta > 0$, one has

$$\lim_{r\to\infty} \frac{1+\log r}{r^{2\delta}} = 0.$$

Hence there exists a constant $C_2 > 0$ such that $1 + \log r \leq C_2 r^{2\delta}$ for all $r > 0$. Then we choose $C_\delta := \sqrt{C_1 C_2} \vee 2$ as desired. □

Remark 2.5. *Actually, the estimation above can be further refined for* $\tau > 3$*. If* $\tau > 3$*, the weights* W_x *and* W_y *have finite variance. In this case, we can get rid of the* δ *in* (2.1)*. On the other hand, since we can choose* δ *arbitrarily small, the refinement does not change our result. For our purpose, we choose* δ *small enough* *that* $\alpha - \delta > d$ *and* $\alpha_1 - \delta > d$.

Now we estimate the probability that a path is open from above. Note that we call π a path of length n if there exist $n+1$ distinct vertices $x_0, \dots, x_n \in \mathbb{Z}^d$ such that $\pi = (x_0, \dots, x_n)$. We say that π is open if all the edges $\{x_{i-1}, x_i\}_{i=1,\dots,n}$ are open.

Lemma 2.6 ([28, Thm. 4.2])**.** *Let* $\pi := (z_0, z_1, \dots, z_n) \in \left(\mathbb{Z}^d\right)^{n+1}$ *be a path of length* n*. Then for all* $\delta > 0$,

$$\mathbb{P}\left(\pi \text{ is open}\right) \leq \prod_{i=1}^{n} C_\delta |z_i - z_{i-1}|^{-\alpha_1+\delta},$$

where the constant C_δ *is as in Lemma 2.4.*

The proof of Lemma 2.6 can be found in the proof of Theorem 4.2 in [28], which combines the Cauchy-Schwarz inequality with the alternating independence of the edges in the path. With Lemma 2.6, one realises immediately that SFP behaves similarly to LRP in the sense that they have similar upper bounds for the probability of a path, which also indicates that the lower bound of SFP might be treated similar to LRP.

Definition 2.7. *Let* $x, y \in \mathbb{Z}^d$ *be distinct,* $\eta \in \big(0, \alpha_1/(2d)\big)$*, and* $n \geq 2$*. We define* $\mathcal{E}_n = \mathcal{E}_n(\eta)$ *as the event that every hierarchy* $\mathcal{H}_n(x,y)$ *of depth* n *connecting* x *and* y *such that*

$$|z_{\sigma 01} - z_{\sigma 10}| \geq |z_{\sigma 0} - z_{\sigma 1}|(\log N)^{-\Delta_1} \tag{2.2}$$

holds for all $k = 0, 1, \dots, n-2$*, and all* $\sigma \in \{0,1\}^k$ *also satisfy the bounds*

$$\prod_{\sigma \in \{0,1\}^k} |z_{\sigma 0} - z_{\sigma 1}| \vee 1 \geq N^{(2\eta)^k} \quad \text{for all } k = 1, 2, \dots, n-1, \tag{2.3}$$

where $N = |x - y|$ *is the Euclidean distance between* x *and* y.

With help of Lemma 2.6 we now can estimate the probability of the event $\mathcal{E}_n$.

Proposition 2.8. *Let* $\eta \in \big(0, \alpha_1/(2d)\big)$*. Pick* $\delta > 0$ *so small that* $\alpha_1 - \delta - d > 0$ *and* $\alpha_1 - \delta \in (2d\eta, \alpha_1)$*, then there exists a constant* $c_1 > 0$ *such that for all* $x, y \in \mathbb{Z}^d$ *with* $N = |x - y|$ *satisfying* $\eta^n \log N \geq 2(\alpha_1 - \delta - d)$,

$$\mathbb{P}\left(\mathcal{E}_{n+1}^c \cap \mathcal{E}_n\right) \leq (\log N)^{c_1 2^n} N^{-(\alpha_1 - \delta - 2d\eta)(2\eta)^n},$$

and

$$\mathbb{P}\left(\mathcal{E}_2^c\right) \leq (\log N)^{c_1} N^{-(\alpha_1-\delta-2d\eta)}.$$

Proof. We modify the proof of Lemma 4.5 in [12] to fit our model.

Let $\mathcal{A}(n)$ be the set of all 2^n-tuples (z_σ) of sites (or hierarchies) such that (2.2) holds for all $\sigma \in \left\{\{0,1\}^k : k = 0, 1, \dots, n-1\right\}$ and (2.3) is true for $k = 1, 2, \dots, n-1$ but not for $k = n$. Then

$$\mathbb{P}\left(\mathcal{E}_{n+1}^c \cap \mathcal{E}_n\right) \leq \sum_{(z_\sigma)\in\mathcal{A}(n)} \mathbb{P}\left(\mathcal{H}_n(x,y) \text{ with sites } (z_\sigma)\right). \tag{2.4}$$

Here the event "$\mathcal{H}_n(x,y)$ with sites (z_σ)" means all the edges in this hierarchy with sites (z_σ) are open as in Condition 3 in Definition 2.2.

Now we fix one single hierarchy $\mathcal{H}_n(x,y)$ with sites (z_σ) and estimate its probability. Typically, a hierarchy consists of isolated edges, i.e., edges that do not share a common vertex. However, since we also allow degenerate vertices as in Condition 5 of Definition 2.2, there might be adjacent edges in the hierarchy. Nevertheless, we can decompose one hierarchy into several disjoint connected components, as exemplified in Figure 2.1. Condition 5 ensures that each of the connected components is an open path.

Example. *Consider the toy example in Figure 2.1. This hierarchy $\mathcal{H}_4(x,y)$ can be divided into 5 disjoint paths, namely*

$$\pi_1 = (z_{0110}, z_{110}, z_{001}, z_{0001}), \quad \pi_2 = (z_{01}, z_{10}),$$
$$\pi_3 = (z_{1001}, z_{1010}), \quad \pi_4 = (z_{101}, z_{110}), \quad \pi_5 = (z_{1101}, z_{1110}).$$

Now assume that the hierarchy $\mathcal{H}_n(x,y)$ can be divided into m disjoint open paths π_i, $i = 1, 2, \dots, m$, with $\pi_i = (x_{i0}, x_{i1}, \dots, x_{im_i})$ and $x_{ij} \in (z_\sigma)$. Then independence of edge occupation implies

$$\begin{aligned}\mathbb{P}\left(\mathcal{H}_n(x,y) \text{ with sites } (z_\sigma)\right) &= \mathbb{P}\left(\bigcap_{k=0}^{n-1} \bigcap_{\sigma\in\{0,1\}^k} \{z_{\sigma01} \sim z_{\sigma10}\}\right) \\ &= \mathbb{P}\left(\bigcap_{i=1}^{m} \{\pi_i \text{ is open}\}\right) = \prod_{i=1}^{m} \mathbb{P}\left(\pi_i \text{ is open}\right),\end{aligned}$$

where we rearrange the open edges in the hierarchy in the second step and use the fact that these open paths are vertex-disjoint and therefore independent in the last step. Further,

$$\begin{aligned}\mathbb{P}\left(\mathcal{H}_n(x,y) \text{ with sites } (z_\sigma)\right) &\leq \prod_{i=1}^{m}\prod_{j=1}^{m_i} C_\delta \left|x_{im_j} - x_{im_{j-1}}\right|^{-\alpha_1+\delta} \\ &= \prod_{k=0}^{n-1}\prod_{\sigma\in\{0,1\}^k} \frac{C_\delta}{(|z_{\sigma01} - z_{\sigma10}| \vee 1)^{\alpha_1-\delta}},\end{aligned}$$

where we apply Lemma 2.6 first and then bring the edges back in the original order again. In the last step we add the maximum with 1 to make sure that the denominator is not zero.

Likewise, we denote the "gaps" in the hierarchy by

$$t_\sigma := z_{\sigma0} - z_{\sigma1},$$

and $t_\emptyset := x - y$. With this notation, we rewrite condition (2.2) as

$$|z_{\sigma01} - z_{\sigma10}| \geq |t_\sigma|(\log N)^{-\Delta_1} \tag{2.5}$$

and condition (2.3) as

$$\prod_{\sigma\in\{0,1\}^k} |t_\sigma| \vee 1 \geq N^{(2\eta)^k}. \tag{2.6}$$

Let $\mathcal{B}(k)$ be the set of all collections $(t_\sigma)_{\sigma\in\{0,1\}^k}$ of vertices in $\mathbb{Z}^d$ such that (2.6) is true. Then (2.4) implies

$$\mathbb{P}\left(\mathcal{E}_{n+1}^c \cap \mathcal{E}_n\right) \leq |\mathcal{B}^c(n)| \prod_{k=0}^{n-1}\left(\sum_{(t_\sigma)\in\mathcal{B}(k)}\prod_{\sigma\in\{0,1\}^k} C_\delta\left(\frac{(\log N)^{\Delta_1}}{|t_\sigma|\vee 1}\right)^{\alpha_1-\delta}\right)$$

Note that for $k = 0$, we have $|t_\emptyset| = N$. Hence the estimation above can be written as

$$|\mathcal{B}^c(n)| \frac{\left(C_\delta(\log N)^{\Delta_1(\alpha_1-\delta)}\right)^{2^n}}{N^{\alpha_1-\delta}} \prod_{k=1}^{n-1}\left(\sum_{(t_\sigma)\in\mathcal{B}(k)}\prod_{\sigma\in\{0,1\}^k}\frac{C_\delta}{(|t_\sigma|\vee 1)^{\alpha_1-\delta}}\right), \tag{2.7}$$

For each k there are at most 2^k multipliers in the product over all $\sigma \in \{0,1\}^k$ (the number is smaller if there exist degenerate sites). Therefore, there are in total $\sum_{k=0}^{n-1} 2^k = 2^n - 1$ and we get the exponent 2^n in the numerator in the first fraction.

In addition, for $n = 2$, the event $\mathcal{E}_2^c$ means that there exists a hierarchy with sites (z_σ) of depth 2 such that

$$|z_{01} - z_{10}| \geq |z_0 - z_1| (\log N)^{-\Delta_1} = N (\log N)^{-\Delta_1},$$

and

$$|z_0 - z_{01}||z_{11} - z_1| \leq N^{2\eta}.$$

Therefore

$$\mathbb{P}(\mathcal{E}_2^c) \leq \sum_{(t_\sigma) \notin \mathcal{B}(1)} \mathbb{P}(z_{01} \sim z_{10}) \leq |\mathcal{B}^c(1)| \frac{C_\delta (\log N)^{\Delta_1(\alpha_1 - \delta)}}{N^{\alpha_1 - \delta}}. \tag{2.8}$$

In order to estimate (2.7) and (2.8), we need two lemmas from the appendix of [12]. First for $\kappa \in \mathbb{N}$ and $b > 0$, we let

$$\Theta_\kappa(b) = \left\{ (n_i) \in \mathbb{N}^\kappa : n_i \geq 1, \prod_{i=1}^{\kappa} n_i \geq b^\kappa \right\},$$

and $\Theta_\kappa^c(b)$ be its complement in $\mathbb{N}^\kappa$. Then one has the following estimates.

Lemma 2.9 (Lemma A.1 in [12]). *For each $\epsilon > 0$ there exists a constant $g_1 = g_1(\epsilon) < \infty$ such that*

$$\sum_{(n_i) \in \Theta_\kappa(b)} \prod_{i=1}^{\kappa} \frac{1}{n_i^{1+\beta}} \leq (g_1 b^{-\beta} \log b)^\kappa$$

is true for all $\beta > 0$, all $b > 1$ and all $\kappa \in \mathbb{N}$ with

$$\beta - \frac{\kappa - 1}{\kappa \log b} \geq \epsilon.$$

Lemma 2.10 (Lemma A.2 in [12]). *There exists a constant $g_2 < \infty$ such that for each $\beta > 1$, each $b \geq e/4$ and any $\kappa \in \mathbb{N}$,*

$$\sum_{(n_i) \in \Theta_\kappa^c(b)} \prod_{i=1}^{\kappa} n_i^{\beta - 1} \leq (g_2 b^\beta \log b)^\kappa.$$

Let (n_σ) be a collection of positive integers with $n_\sigma \leq |t_\sigma| \vee 1 < n_\sigma + 1$. Note that $|\{x \in \mathbb{Z}^d : n \leq |x| \vee 1 < n + 1\}| \leq c n^{d-1}$ for some positive constant $c = c(d)$

independent of n. Then for each n_σ there exists at most cn_σ^{d-1} such t_σ's. Therefore,

$$\sum_{(t_\sigma)\in\mathcal{B}(k)} \prod_{\sigma\in\{0,1\}^k} \frac{C_\delta}{(|t_\sigma|\vee 1)^{\alpha_1-\delta}} \leq \sum_{(n_\sigma)\in\Theta_{2^k}(N^{\eta^k})} \prod_{\sigma\in\{0,1\}^k} \left(cn_\sigma^{d-1}\frac{C_\delta}{n_\sigma^{\alpha_1-\delta}}\right)$$
$$\leq \frac{(C_\delta c g_1)^{2k}(\eta^k)^{2^k}(\log N)^{2^k}}{N^{\eta^k 2^k(\alpha_1-\delta-d)}}, \tag{2.9}$$

where we have applied Lemma 2.9 in the last step (2.9) with $\beta=\alpha_1-\delta-d, b=N^{\eta^k}$ and $\kappa=2^k$. Since $\eta<1$, we obtain the further bound

$$\sum_{(t_\sigma)\in\mathcal{B}(k)} \prod_{\sigma\in\{0,1\}^k} \frac{C_\delta}{(|t_\sigma|\vee 1)^{\alpha_1-\delta}} \leq \frac{(C_1\log N)^{2^k}}{N^{(\alpha_1-\delta-d)(2\eta)^k}},$$

where we choose $C_1:=cC_\delta g_1$. Now it is left to estimate the size of $\mathcal{B}^c(n)$, and this can be done with help of Lemma 2.10 as

$$\sum_{(t_\sigma)\notin\mathcal{B}^c(n)} 1 \leq (C_2\log N)^{2^n} N^{d(2\eta)^n}$$

with $\beta=d, b=N^{\eta^n}$ and $\kappa=2^n$.

Now (2.7) can be simplified to

$$(C_2\log N)^{2^n}N^{d(2\eta)^n}\frac{\left(C_\delta(\log N)^{\Delta_1(\alpha_1-\delta)}\right)^{2^n}}{N^{\alpha_1-\delta}}\prod_{i=1}^{n-1}\frac{(C_1\log N)^{2^k}}{N^{(\alpha_1-\delta-d)(2\eta)^k}}.$$
$$\leq \left(C_1C_2C_\delta(\log N)^{\Delta_1(\alpha_1-\delta)+2}\right)^{2^n} N^{-\left((\alpha_1-\delta-d)\sum_{k=1}^{n-1}(2\eta)^k+\alpha_1-\delta-d(2\eta)^n\right)}$$
$$\leq (\log N)^{c_1 2^n} N^{-(\alpha_1-\delta-2d\eta)(2\eta)^n},$$

where the last step uses the bound

$$(\alpha_1-\delta-d)\sum_{k=1}^{n-1}(2\eta)^k+\alpha_1-\delta-d(2\eta)^n \geq (\alpha_1-\delta-2d\eta)(2\eta)^n.$$

□

Our further strategy is to show that an open path with distance shorter then poly-logarithm is impossible. More precisely, we show that the existence of a shorter path is contained in some event with negligible probability. The event we use is as follows.

Definition 2.11. *Let $x, y \in \mathbb{Z}^d$ be distinct and $n \in \mathbb{N}$. We define $\mathcal{F}_n := \mathcal{F}_n(x, y)$ as the event that for every hierarchy of depth n connecting x and y and satisfying (2.2), every collection of (vertex self-avoiding and) mutually disjoint paths π_σ with $\sigma \in \{0,1\}^{n-1}$ such that π_σ connects $z_{\sigma 0}$ and $z_{\sigma 1}$ without using any vertex from the hierarchy (except for the endpoints $z_{\sigma 0}$ and $z_{\sigma 1}$) obeys the bound*

$$\sum_{\sigma \in \{0,1\}^{n-1}} |\pi_\sigma| \geq 2^n. \tag{2.10}$$

It might be instructive to look at the complement $\mathcal{F}_n^c$: this is the event that there exists such a hierarchy between x and y satisfying (2.2), but the edges filling the gaps violate (2.10). In the following proposition, we construct such a hierarchy in $\mathcal{F}_n^c$ from the shortest path.

Proposition 2.12 (Lemma 4.6 in [12])**.** *Let $\epsilon \in (0, \Delta_1)$. If $N = |x-y|$ is sufficiently large and*

$$n > \frac{\Delta_1 - \epsilon}{\log 2} \log\log N, \tag{2.11}$$

then

$$\{D(x,y) \leq (\log N)^{\Delta_1 - \epsilon}\} \cap \mathcal{F}_n = \emptyset.$$

Proof. The proof of Lemma 4.6 in [12] still holds here for the event with modified hierarchy, because the hierarchy there was constructed from the shortest path in which all the vertices have degree at most 2. For better readability the proof sketch is given here.

If $D(x,y) \leq (\log N)^{\Delta_1 - \epsilon} \leq (\log N)^{\Delta_1}$, by triangle inequality, the shortest path between x, y has at least one edge with length at least $N/(\log N)^{\Delta_1}$. Denote by z_{01} and z_{10} the endpoint of this long edges on the x-side and y-side, respectively. That is,

$$|z_{01} - z_{10}| \geq |z_0 - z_1|(\log N)^{-\Delta_1}.$$

Apparently, $D(x, z_{01})$ and $D(y, z_{10})$ are both at most $(\log N)^{\Delta_1}$. With the similar argument one finds the longest edge $\{z_{001}, z_{010}\}$ in the gap between x and z_{01}, and $\{z_{101}, z_{110}\}$ for the gap between y and z_{10}. After iterating the steps n times, we obtain a hierarchy of depth n that satisfies (2.2). (2.11) implies that

$$(\log N)^{\Delta_1 - \epsilon} \leq 2^n.$$

Therefore, the hierarchy we constructed from the short path satisfies:

$$\sum_\sigma |\pi_\sigma| < D(x,y) \leq (\log N)^{\Delta_1 - \epsilon} \leq 2^n.$$

In other words, (2.10) is violated. □

Now we start to fill the "gaps" in the hierarchy. More precisely, we relate the events $\mathcal{E}_n$ and $\mathcal{F}_n$ by the following proposition.

Proposition 2.13. *Let $\eta \in (0, \alpha_1/(2d))$. For $\delta > 0$ so small that $\alpha_1 - \delta - d > 0$ and $\alpha_1 - \delta \in (2d\eta, \alpha_1)$, there exists a constant $c_2 > 0$ such that for all distinct $x, y \in \mathbb{Z}^d$ with $N = |x-y|$ satisfying $\eta^n \log N \geq 2(\alpha_1 - \delta - d)$,*

$$\mathbb{P}\left(\mathcal{F}_n^c \cap \mathcal{E}_n\right) \leq (\log N)^{c_2 2^n} N^{-(\alpha_1 - \delta)(2\eta)^{n-1}}.$$

The idea of proof is to first fix one hierarchy with the sites (z_σ), and estimate the probability that the paths that fill the gaps of this hierarchy have a certain length. Then the gap-filling paths and the open edges in the hierarchy constitute a path connecting x and y. With help of Lemma 2.6 we get the upper bound by summing over all possible hierarchies.

Proof. Let $\mathcal{A}^*(n)$ be the set of all collections (z_σ), $\sigma \in \{0,1\}^n$, satisfying (2.2) for $k = 0, 1, \ldots, n-2$ and (2.3) for $k = 1, 2, \ldots, n-1$. Then

$$\mathbb{P}(\mathcal{F}_n^c \cap \mathcal{E}_n) = \sum_{(z_\sigma) \in \mathcal{A}^*(n)} \mathbb{P}\left(\mathcal{F}_n^c \cap \mathcal{H}_n \text{ on } (z_\sigma)\right). \tag{2.12}$$

Here "$\mathcal{F}_n^c \cap \mathcal{H}_n$ on (z_σ)" means that $\mathcal{H}_n$ with sites (z_σ) is a hierarchy satisfying $\mathcal{F}_n^c$, as we have explained after Definition 2.11.

We estimate the summands on the right hand side of (2.12) by considering all possible lengths of π_σ. More precisely, let (m_σ) be a tuple of non-negative integers for $\sigma \in \{0,1\}^{n-1}$. Then

$$\mathbb{P}\left(\mathcal{F}_n^c \cap \mathcal{H}_n \text{ on } (z_\sigma)\right) = \sum_{(m_\sigma)} \mathbb{P}\left(\mathcal{F}_n^c \cap \mathcal{H}_n \text{ on } (z_\sigma) \text{ with } (|\pi_\sigma|) = (m_\sigma)\right). \tag{2.13}$$

Note that the open path π_σ fills the gap between $z_{\sigma 0}$ and $z_{\sigma 1}$ in $\mathcal{H}_n$ for all $\sigma \in \{0,1\}^{n-1}$. All such open paths together with all the open edges $(z_{\sigma 01}, z_{\sigma 10}), \sigma \in \{0,1\}^{n-2}$, constitute a self-avoiding open path between x and y. Let $\Gamma_\sigma(m_\sigma)$ be the set of all path of length m_σ connecting $z_{\sigma 0}$ and $z_{\sigma 1}$, that is,

$$\Gamma_\sigma(m_\sigma) = \left\{\pi : \pi = (x_0, x_1, \dots, x_{m_\sigma}) \text{ with } x_0 = z_{\sigma 0} \text{ and } x_{m_\sigma} = z_{\sigma 1}\right\}.$$

Now we estimate the probability in (2.13) as

$$\begin{aligned}
&\mathbb{P}\left(\mathcal{F}_n^c \cap \mathcal{H}_n \text{ on } (z_\sigma) \text{ with } (|\pi_\sigma|) = (m_\sigma)\right) \\
=&\mathbb{E}\left[\mathbb{P}\left(\mathcal{F}_n^c \cap \mathcal{H}_n \text{ on } (z_\sigma) \text{ with } (|\pi_\sigma|) = (m_\sigma)\right) \mid (W_x)_{x\in\mathbb{Z}^d}\right] \\
=&\mathbb{E}\left[\mathbb{P}\left(\bigcap_{\sigma\in\{0,1\}^{n-1}} \{z_{\sigma 0} \overset{\pi_\sigma}{\leftrightarrow} z_{\sigma 1}\} \bigcap_{\sigma\in\{0,1\}^{n-2}} \{z_{\sigma 01} \sim z_{\sigma 10}\} \middle| (W_x)_{x\in\mathbb{Z}^d}\right)\right],
\end{aligned} \tag{2.14}$$

where $\{z_{\sigma 0} \overset{\pi_\sigma}{\leftrightarrow} z_{\sigma 1}\}$ means π_σ connects $z_{\sigma 0}$ and $z_{\sigma 1}$.
By the conditional independence of edges, we rewrite (2.14) as

$$\begin{aligned}
&\mathbb{P}\left(\mathcal{F}_n^c \cap \mathcal{H}_n \text{ on } (z_\sigma) \text{ with } (|\pi_\sigma|) = (m_\sigma)\right) \\
\le &\sum_{\substack{(\pi_\sigma):\pi_\sigma=(x_{\sigma 0},\dots,x_{\sigma m_\sigma}) \\ \text{vertex-disjoint}}} \mathbb{E}\left[\prod_{\sigma\in\{0,1\}^{n-1}} \mathbb{P}\left(\pi_\sigma \middle| (W_x)_{x\in\mathbb{Z}^d}\right) \prod_{k=0}^{n-2} \prod_{\sigma'\in\{0,1\}^k} p_{z_{\sigma' 01} z_{\sigma' 10}}\right] \\
= &\sum_{\substack{(\pi_\sigma):\pi_\sigma=(x_{\sigma 0},\dots,x_{\sigma m_\sigma}) \\ \text{vertex-disjoint}}} \mathbb{E}\left[\prod_{\sigma\in\{0,1\}^{n-1}} \prod_{i=1}^{m_\sigma} p_{x_{\sigma(i-1)}, x_{\sigma i}} \prod_{k=0}^{n-2} \prod_{\sigma'\in\{0,1\}^k} p_{z_{\sigma' 01} z_{\sigma' 10}}\right]
\end{aligned} \tag{2.15}$$

where we sum over all possible paths between $z_{\sigma 0}$ and $z_{\sigma 1}$ for all $\sigma \in \{0,1\}^{n-1}$ and p_{xy} is the connection probability as in (1.5).

In the expectation in (2.15) the probability is divided into two parts: the first double product involves the edges filling the gaps in the hierarchy while the second double product is about the open edges in the hierarchy, as depicted in Figure 2.2.

Note that all these paths (π_σ) have mutually disjoint vertices. Therefore, for fixed sites (z_σ) and fixed paths (π_σ), we obtain a self-avoiding open path starting

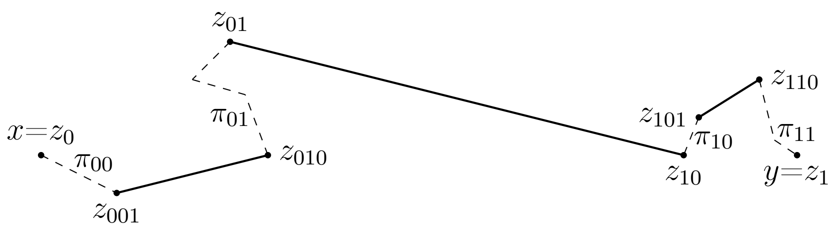

Figure 2.2: A hierarchy of depth 3 with site $(z_\sigma)_{\sigma\in\{0,1\}^3}$. The gap-filling paths are $\{\pi_\sigma\}$ with $\sigma \in \{0,1\}^2$. In this example $|\pi_{00}| = 1, |\pi_{01}| = 3, |\pi_{10}| = 1, |\pi_{11}| = 2$, and $\sum |\pi_\sigma| = 7 < 2^3 = 8$. We see that the paths here, together with the edges in the hierarchy, form a path connecting x and y.

from x and ending in y. Now we use Lemma 2.6 to bound the probability of this path, i.e. the expectation in (2.15) as

$$\begin{aligned}
&\mathbb{E}\left[\prod_{\sigma\in\{0,1\}^{n-1}}\prod_{i=1}^{m_\sigma} p_{x_{\sigma(i-1)},x_{\sigma i}} \prod_{k=0}^{n-2}\prod_{\sigma'\in\{0,1\}^k} p_{z_{\sigma'01}z_{\sigma'10}}\right]\\
\leq &\prod_{\sigma\in\{0,1\}^{n-1}}\prod_{i=1}^{m_\sigma}\frac{C_\delta}{\left(|x_{\sigma(i-1)}-x_{\sigma i}|\vee 1\right)^{\alpha_1-\delta}}\prod_{k=0}^{n-2}\prod_{\sigma'\in\{0,1\}^k}\frac{C_\delta}{|z_{\sigma'01}-z_{\sigma'10}|^{\alpha_1-\delta}}.
\end{aligned}$$

Then (2.15) becomes

$$\begin{aligned}
&\mathbb{P}\left(\mathcal{F}_n^c\cap\mathcal{H}_n \text{ on } (z_\sigma) \text{ with } (|\pi_\sigma|)=(m_\sigma)\right)\\
\leq&\sum_{(\pi_\sigma)}\prod_{\sigma\in\{0,1\}^{n-1}}\prod_{i=1}^{m_\sigma}\frac{C_\delta}{\left(|x_{\sigma(i-1)}-x_{\sigma i}|\vee 1\right)^{\alpha_1-\delta}}\prod_{k=0}^{n-2}\prod_{\sigma'\in\{0,1\}^k}\frac{C_\delta}{|z_{\sigma'01}-z_{\sigma'10}|^{\alpha_1-\delta}}\\
=&\left(\prod_{\sigma\in\{0,1\}^{n-1}} Q_{m_\sigma}(z_{\sigma0},z_{\sigma1})\right)\prod_{k=0}^{n-2}\prod_{\sigma'\in\{0,1\}^k}\frac{C_\delta}{|z_{\sigma'01}-z_{\sigma'10}|^{\alpha_1-\delta}}
\end{aligned}$$

where

$$Q_m(u,v) := \sum_{\substack{\pi=(x_0,\dots,x_m)\\ x_0=u,\ ,x_m=v}}\prod_{i=1}^{m}\frac{C_\delta}{\left(|x_{i-1}-x_i|\vee 1\right)^{\alpha_1-\delta}}.$$

Here the sum runs over self-avoiding paths π of length m, and therefore $Q_m(u,v)$ is the upper bound for the probability that u and w are connected by an open path with length m. To simplify $Q_m(u,v)$ we will need the following lemma:

Lemma 2.14. *For all $u, v \in \mathbb{Z}^d$ with $u \neq v$ and $\alpha > d$, there exists a constant*

$a \in (0,\infty)$, *independent of* u *and* v, *such that*

$$\sum_{w\in\mathbb{Z}^d, w\notin\{u,v\}} \frac{1}{|u-w|^\alpha}\frac{1}{|v-w|^\alpha} \le \frac{a}{|u-v|^\alpha}. \tag{2.16}$$

Proof. Let $A := \{w \in \mathbb{Z}^d : |u-w| \ge \frac{1}{2}|u-v|\}$ and $B := \{w \in \mathbb{Z}^d : |v-w| \ge \frac{1}{2}|u-v|\}$. By triangle inequality, for an arbitrary $w \in \mathbb{Z}^d$ we have either $w \in A$ or $w \in B$. Therefore

$$\begin{aligned}\sum_{w\in\mathbb{Z}^d, w\notin\{u,v\}} \frac{1}{|u-w|^\alpha}\frac{1}{|v-w|^\alpha} &\le \sum_{w\in A, w\neq v} \frac{1}{|u-w|^\alpha}\frac{1}{|v-w|^\alpha} + \sum_{w\in B, w\neq u} \frac{1}{|u-w|^\alpha}\frac{1}{|v-w|^\alpha}\\ &\le \sum_{w\neq v} \frac{2^\alpha}{|u-v|^\alpha}\frac{1}{|v-w|^\alpha} + \sum_{w\neq u} \frac{2^\alpha}{|u-v|^\alpha}\frac{1}{|u-w|^\alpha}\\ &\le \frac{2^{\alpha+1}}{|u-v|^\alpha}\sum_{w\neq u}\frac{1}{|u-w|^\alpha}.\end{aligned}$$

Since $\alpha > d$, we have $a := 2^{\alpha+1}\sum_{w\neq u}\frac{1}{|u-w|^\alpha} < \infty$. □

With help of Lemma 2.14 we can bound $Q_m(u,v)$ from above by applying (2.16) m times iteratively, and obtain

$$Q_m(u,v) \le \frac{(C_\delta a)^m}{(|u-v| \vee 1)^{\alpha_1-\delta}}. \tag{2.17}$$

If we now sum over all the possible combinations of (m_σ) with $\sum_\sigma m_\sigma < 2^n$, we obtain the upper bound

$$\begin{aligned}&\mathbb{P}\left(\mathcal{F}_n^c \cap \mathcal{H}_n \text{ on } (z_\sigma)\right)\\ &\le \sum_{(m_\sigma):\sum_\sigma m_\sigma<2^n}\left(\prod_{\sigma\in\{0,1\}^{n-1}} Q_{m_\sigma}(z_{\sigma0}, z_{\sigma1})\right)\prod_{k=0}^{n-2}\prod_{\sigma'\in\{0,1\}^k}\frac{C_\delta}{|z_{\sigma'01}-z_{\sigma'10}|^{\alpha_1-\delta}}\\ &\le (4aC_\delta)^{2^n}\prod_{\sigma\in\{0,1\}^{n-1}}\frac{1}{(|z_{\sigma0}-z_{\sigma1}|\vee 1)^{\alpha_1-\delta}}\prod_{k=0}^{n-2}\prod_{\sigma'\in\{0,1\}^k}\frac{C_\delta}{|z_{\sigma'01}-z_{\sigma'10}|^{\alpha_1-\delta}}\\ &\le (4aC_\delta)^{2^n}\prod_{k=0}^{n-1}\prod_{\sigma\in\{0,1\}^k}\frac{C_\delta(\log N)^{(\alpha_1-\delta)\Delta'}}{(|z_{\sigma0}-z_{\sigma1}|\vee 1)^{\alpha_1-\delta}}.\end{aligned}$$

Here we first used the estimation for $Q_m(u,v)$ in (2.17) and the fact that the number of such eligible tuples (m_σ) is at most 4^{2^n}, and subsequently used the fact that on $\mathcal{E}_n$ the lengths of open edges in the hierarchy are subject to the constrain (2.5).

We now can estimate the desired probability as

$$\mathbb{P}(\mathcal{F}_n^c \cap \mathcal{E}_n) = \sum_{(z_\sigma)\in\mathcal{A}^*(n)} (4aC_\delta)^{2^n} \prod_{k=0}^{n-1} \prod_{\sigma\in\{0,1\}^k} \frac{C_\delta(\log N)^{(\alpha_1-\delta)\Delta'}}{(|z_{\sigma0}-z_{\sigma1}|\vee 1)^{\alpha_1-\delta}}$$
$$\leq \frac{\left(C_1(\log N)^{\Delta_1(\alpha_1-\delta)}\right)^{2^n}}{N^{\alpha_1-\delta}} \prod_{k=0}^{n-1} \sum_{(t_\sigma)\in\mathcal{B}(k)} \prod_{\sigma\in\{0,1\}^k} \frac{C_\delta}{(|t_\sigma|\vee 1)^{\alpha_1-\delta}}.$$

Recall that $\mathcal{B}(k)$ is the set of all collections $(t_\sigma), \sigma \in \{0,1\}^k$, of vertices in $\mathbb{Z}^d$ such that (2.6) is true. Then by applying Lemma 2.9 again (as in (2.9)), together with

$$\alpha_1 - \delta + (\alpha_1-\delta)\sum_{k=1}^{n-1}(2\eta)^k \geq (\alpha_1-\delta)(2\eta)^{n-1},$$

the result follows. □

Proof of Theorem 1.3, lower bound. By Proposition 2.12 we can bound the probability of the event $\{D(x,y) \leq (\log N)^{\Delta_1-\epsilon}\}$ by the probability of the event $\mathcal{F}_n^c$ once Proposition 2.12 holds. That is, if the depth of the hierarchy n satisfies (2.11),

$$\mathbb{P}\left(D(x,y) \leq (\log N)^{\Delta_1-\epsilon}\right) \leq \mathbb{P}\left(\mathcal{F}_n^c\right).$$

Now we fix $\epsilon \in (0, \Delta_1 - 1)$. Since $2^{-1/\Delta_1} = \alpha_1/2d$ by (1.12), we can choose $\delta > 0$ and η such that

$$2^{-1/(\Delta_1-\epsilon)} < \eta < \frac{\alpha_1-\delta}{2d},$$

so that, in particular, $\frac{\Delta_1-\epsilon}{\log 2} < \frac{1}{\log 1/\eta}$. We further fix $\delta_1 \in (0, \alpha_1-\delta-2d\eta)$. For large N we thus find $n \in \mathbb{N}$ such that

$$\frac{\Delta_1-\epsilon}{\log 2}\log\log N < n \leq \frac{\log\log N + \log\frac{\delta_1}{c_1} - \log\log\log N}{\log 1/\eta}. \tag{2.18}$$

We henceforth assume that N is large enough that (for c_1 from Proposition 2.8)

$$(\log N)^{c_1 2^n} \leq N^{\delta_1(2\eta)^n}. \tag{2.19}$$

In this case, the right hand side of (2.18) is further bounded from above by

$$\frac{\log\log N - \log 2(\alpha_1-\delta-d)}{\log 1/\eta}.$$

Therefore, we may apply the assertions of Propositions 2.8, 2.12 and 2.13 (Proposition 2.8 even for all smaller values of n), and we thus get

$$\begin{aligned}\mathbb{P}\left(D(x,y) \leq (\log N)^{\Delta_1-\epsilon}\right) &\leq \mathbb{P}\left(\mathcal{F}_n^c\right) \leq \mathbb{P}\left(\mathcal{E}_n^c\right) + \mathbb{P}\left(\mathcal{F}_n^c \cap \mathcal{E}_n\right) \\ &\leq \sum_{k=3}^{n} \mathbb{P}(\mathcal{E}_k^c \cap \mathcal{E}_{k-1}) + \mathbb{P}(\mathcal{E}_2^c) + \mathbb{P}\left(\mathcal{F}_n^c \cap \mathcal{E}_n\right). \qquad (2.20)\end{aligned}$$

Using Proposition 2.8 and (2.18), we get for $k \leq n$ that

$$\mathbb{P}(\mathcal{E}_{k+1}^c \cap \mathcal{E}_k) \leq N^{-(\alpha_1-\delta-2d\eta-\delta_1)(2\eta)^k},$$

and Proposition 2.13 yields a similar bound for $\mathbb{P}\left(\mathcal{F}_n^c \cap \mathcal{E}_n\right)$. Since $2\eta > 1$, we thus get the right hand side of (2.20) arbitrarily close to 0 by choosing N sufficiently large.

Translation invariance and the FKG-inequality yield

$$\mathbb{P}(x, y \in \mathcal{C}_\infty) \geq \mathbb{P}(x \in \mathcal{C}_\infty)^2 > 0.$$

Therefore, we have

$$\lim_{|x-y|\to\infty} \mathbb{P}\left(D(x,y) \leq (\log|x-y|)^{\Delta_1-\epsilon} \,\middle|\, x, y \in \mathcal{C}_\infty\right) = 0,$$

as desired. □

2.2 Upper bound for graph distances

Theorem 2.15 (Upper bound in Theorem 1.3)**.** *For scale-free percolation with parameters $\lambda > \lambda_c, \gamma > 2$, and $d < \alpha < 2d$, we have that for any $\epsilon > 0$,*

$$\lim_{|x-y|\to\infty} \mathbb{P}\left(D(x,y) \leq (\log|x-y|)^{\Delta_2+\epsilon} \,\middle|\, x, y \in \mathcal{C}_\infty\right) = 1.$$

Here $D(x, y)$ is the graph distance between x and y, and $\mathcal{C}_\infty$ is the unique infinite cluster. The parameter Δ_2 is given by

$$\Delta_2 := \frac{\log 2}{\log(2d/\alpha_2)} \quad \textit{with} \quad \alpha_2 := \alpha \wedge (\alpha(\tau - 1) - d)\,.$$

The upper bound in (b) and (c) of Theorem 1.3 is already established in [29], so that we can restrict our attention here to the case $\tau \in (2, 3)$. Interestingly, for $\tau \geq 3$, the logarithmic power of upper and lower bound match, and we thus identified the correct exponent.

Unlike in long-range percolation, edges in scale-free percolation are only conditionally independent. Intuitively speaking, adjacent edges are positively correlated due to the weight of their joint vertex (see Exercise 9.45 in Chapter 9 of [62]). Here we state a more general result, which is implied by the FKG-Inequality (see e.g Theorem 2.4 in [53]).

Proposition 2.16. *Let $\pi = (x_i)_{i=0,\dots,n}$ be a path in scale-free percolation and $k \in \{1, \dots, n-1\}$, and let π_1, π_2 be two subpaths of π by cutting π at vertex x_k. That is, $\pi_1 = (x_i)_{i=0,\dots,k}$ and $\pi_2 = (x_i)_{i=k,\dots,n}$. Then*

$$\mathbb{P}\left(\pi \textit{ is open}\right) \geq \mathbb{P}\left(\pi_1 \textit{ is open}\right) \mathbb{P}\left(\pi_2 \textit{ is open}\right).$$

From Proposition 2.16 we see that two adjacent edges (or even paths) in scale-free percolation are indeed positively correlated. The next result tells us that in some cases the positive correlation is significant.

Proposition 2.17 (Probability of adjacent edges)**.** *In scale-free percolation with $\tau \in (2, 3)$ there exist $x_0 > 0$ and $c_2 > c_1 > 0$ such that for all x, y and $z \in \mathbb{Z}^d$ with $|x - y| \geq |y - z| \geq x_0$, we have*

$$c_1 |x - y|^{-\alpha} |y - z|^{-\alpha(\tau-2)} \leq \mathbb{P}(x \sim y \sim z) \leq c_2 |x - y|^{-\alpha} |y - z|^{-\alpha(\tau-2)}.$$

Proof. We start by calculating the probability of this joint occurrence as

$$\mathbb{P}(x \sim y \sim z) = \mathbb{E}\left[\left(\frac{\lambda W_x W_y}{|x - y|^\alpha} \wedge 1\right)\left(\frac{\lambda W_y W_z}{|y - z|^\alpha} \wedge 1\right)\right].$$

Now show that the two single weights W_x and W_z do not play a role in the result.

On the one hand, we know $W_x \geq 1$, therefore

$$\mathbb{E}\left[\left(\frac{\lambda W_x W_y}{|x-y|^\alpha}\wedge 1\right)\left(\frac{\lambda W_y W_z}{|y-z|^\alpha}\wedge 1\right)\right] \geq \mathbb{E}\left[\left(\frac{\lambda W_y}{|x-y|^\alpha}\wedge 1\right)\left(\frac{\lambda W_y}{|y-z|^\alpha}\wedge 1\right)\right]$$

One the other hand, the inequality $st\wedge 1 \leq s(t\wedge 1)$ for $s\geq 1$ and $t>0$, implies

$$\begin{aligned}\mathbb{E}\left[\left(\frac{\lambda W_x W_y}{|x-y|^\alpha}\wedge 1\right)\left(\frac{\lambda W_y W_z}{|y-z|^\alpha}\wedge 1\right)\right] &\leq \mathbb{E}\left[W_x\left(\frac{\lambda W_y}{|x-y|^\alpha}\wedge 1\right)\left(\frac{\lambda W_y}{|y-z|^\alpha}\wedge 1\right)W_z\right]\\ &= \mu^2\mathbb{E}\left[\left(\frac{\lambda W_y}{|x-y|^\alpha}\wedge 1\right)\left(\frac{\lambda W_y}{|y-z|^\alpha}\wedge 1\right)\right],\end{aligned}$$

where $\mu := \mathbb{E}[W_x] < \infty$ since $\tau > 2$. We thus obtain

$$\frac{1}{\mu^2}\mathbb{P}(x\sim y\sim z) \leq \mathbb{E}\left[\left(\frac{\lambda W_y}{|x-y|^\alpha}\wedge 1\right)\left(\frac{\lambda W_y}{|y-z|^\alpha}\wedge 1\right)\right] \leq \mathbb{P}(x\sim y\sim z).$$

Thus it suffices to compute the expectation

$$\begin{aligned}&\mathbb{E}\left[\left(\frac{\lambda W_y}{|x-y|^\alpha}\wedge 1\right)\left(\frac{\lambda W_y}{|y-z|^\alpha}\wedge 1\right)\right]\\ =&\int_{\mathbb{R}}\left(\frac{\lambda u}{|x-y|^\alpha}\wedge 1\right)\left(\frac{\lambda u}{|y-z|^\alpha}\wedge 1\right)dW_y(u)\\ =&\int_1^\infty\left(\frac{\lambda u}{|x-y|^\alpha}\wedge 1\right)\left(\frac{\lambda u}{|y-z|^\alpha}\wedge 1\right)(\tau-1)u^{-\tau}du\end{aligned}$$

We now split the domain of integration into following three intervals:

$$[1,|y-z|^\alpha/\lambda]\,,\quad (|y-z|^\alpha/\lambda,|x-y|^\alpha/\lambda]\,,\quad \text{and } (|x-y|^\alpha/\lambda,\infty)\,.$$

After some calculation, one obtains

$$\begin{aligned}&\mathbb{E}\left[\left(\frac{\lambda W_y}{|x-y|^\alpha}\wedge 1\right)\left(\frac{\lambda W_y}{|y-z|^\alpha}\wedge 1\right)\right]\\ =&\frac{\tau-1}{(3-\tau)(\tau-2)}\frac{\lambda}{|x-y|^\alpha}\frac{\lambda^{\tau-2}}{|y-z|^{\alpha(\tau-2)}}-\frac{\tau-1}{3-\tau}\frac{\lambda}{|x-y|^\alpha}\frac{\lambda}{|y-z|^\alpha}-\frac{\lambda^{\tau-1}}{\tau-2}\frac{1}{|x-y|^{\alpha(\tau-1)}}.\end{aligned}$$

We thus may choose $c_2 := \frac{\tau-1}{(3-\tau)(\tau-2)}\mu^2\lambda^{\tau-1}$.

For $\tau\in(2,3)$, we find that the first term dominates the sum when $|y-z|\to\infty$ (the other terms are negative, but the total sum is trivially nonnegative). Hence there exist positive constant x_0 and c_1 such that

$$\mathbb{P}(x \sim y \sim z) \geq c_1|x-y|^{-\alpha}|y-z|^{-\alpha(\tau-2)} \quad \text{for } |y-z| \geq x_0.$$

□

In fact, the weights of two end points do not contribute to the significant positive correlation in Proposition 2.17, as we formulate in the next corollary.

Corollary 2.18. *In scale-free percolation with $\tau \in (2,3)$, there exist constants $c_i = c_i(a,b) > 0$ for $i = 1,2$ and $x_0 = x_0(a,b) > 0$ such that for all x, y and $z \in \mathbb{Z}^d$ with $|x-y| \geq |y-z| \geq x_0$ we have*

$$c_1|x-y|^{-\alpha}|y-z|^{-\alpha(\tau-2)} \leq \mathbb{P}\left(x \sim y \sim z | W_x = a, W_z = b\right) \leq c_2|x-y|^{-\alpha}|y-z|^{-\alpha(\tau-2)}.$$

In particular, for constants $M > m > 0$, there exist $C_i = C_i(a,b,m,M) > 0, i = 1,2$ and $x_0' = x_0'(a,b) > 0$ such that if $|x-y|$ and $|y-z|$ are comparable in the sense

$$m|x-y| \leq |y-z| \leq M|x-y|,$$

then

$$C_1|x-y|^{-\alpha(\tau-1)/2}|y-z|^{-\alpha(\tau-1)/2} \leq \mathbb{P}\left(x \sim y \sim z | W_x = a, W_z = b\right) \leq C_2|x-y|^{-\alpha(\tau-1)/2}|y-z|^{-\alpha(\tau-1)/2},$$

for all $|x-y| \geq x_0'$.

In light of Propositions 2.16 and 2.17 and Corollary 2.18, we now aim to construct a path with edges of comparable length. Instead of connecting two vertices directly, we use an intermediate vertex as a "bridge" to connect the two vertices. For $x, y \in \mathbb{Z}^d$, $A \subset \mathbb{Z}^d$, we write

$$\{x \sim A \sim y\} = \bigcup_{z \in A} \{x \sim z \sim y\}$$

for the event that x and y are connected via an "intermediate vertex" in A.

Proposition 2.19. *For $\beta \in (0,1)$, there exist constants $N_0, K > 0$ such that for all $x, y \in \mathbb{Z}^d$ with $N := |x-y| \geq N_0$ it is true that*

$$\mathbb{P}\left(x \sim A \sim y\right) \geq \frac{K}{N^{2\alpha_1 - d\beta}},$$

where

$$A := \left(\frac{1}{2}(x+y) + \left[-N^{\beta}, N^{\beta}\right]^{d}\right) \cap \mathbb{Z}^{d}$$

is the cube with side length N^{β} centred at the middle point of the line segment between x and y.

Proof. Since $\beta < 1$, there exist constants $l = l(\beta, d)$ and $L = L(\beta, d)$ with $L > l > 0$ and $N_1 > 0$ such that

$$lN \leq |x-z| \leq LN \quad \text{and} \quad lN \leq |y-z| \leq LN,$$

for all $z \in A$ and all $N \geq N_1$. Therefore, $|x-z|$ and $|y-z|$ are comparable in the sense of Corollary 2.18. Thus we have

$$\begin{aligned}\mathbb{P}\left(x \sim A \sim y\right) &\geq \mathbb{P}\left(x \sim A \sim y | W_x = 1, W_y = 1\right) \\ &= 1 - \prod_{z \in A}\left(1 - \mathbb{P}\left(x \sim z \sim y | W_x = W_y = 1\right)\right),\end{aligned}$$

where we used the conditional independence of edges and the independence of vertex weights.

We estimate this further using Corollary 2.18 and get that there exists $N_2 > 0$, $c_1 > 0$, such that for all $N \geq N_2$,

$$\mathbb{P}\left(x \sim z \sim y | W_x = W_y = 1\right) \geq c_1 \frac{1}{|x-z|^{\alpha_1}} \frac{1}{|z-y|^{\alpha_1}} \geq \frac{c_1}{L^{2\alpha_1}} \frac{1}{N^{2\alpha_1}}. \tag{2.21}$$

Note that the right hand side of (2.21) is independent of z, which allows us to estimate

$$\mathbb{P}\left(x \sim A \sim y\right) \geq 1 - \left(1 - \frac{c_1}{L^{2\alpha_1} N^{2\alpha_1}}\right)^{N^{d\beta}}.$$

Now we use the elementary bound

$$1 - t \leq e^{-t} \leq 1 - \frac{1}{2}t \qquad (0 < t < 1) \tag{2.22}$$

to conclude that

$$\left(1 - \frac{c_1}{L^{2\alpha_1} N^{2\alpha_1}}\right)^{N^{d\beta}} \leq e^{-CN^{d\beta - 2\alpha_1}/L^{2\alpha_1}} \leq 1 - \frac{c_1}{2N^{2\alpha_1 - d\beta} L^{2\alpha_1}}.$$

Since $d\beta - 2\alpha_1 < 0$, there exists $N_3 > 0$ such that we have $c_1 N^{d\beta - 2\alpha_1}/L^{2\alpha_1} < 1$ for all $N \geq N_3$, and consequently also $c_1 N^{-2\alpha_1} L^{-2\alpha_1} < 1$. Finally, we have

$$\mathbb{P}(x \sim A \sim y) \geq \frac{c_1}{2N^{2\alpha_1 - d\beta} L^{2\alpha_1}}$$

for all $N \geq N_0 := \max\{N_1, N_2, N_3\}$ and choose $K := \frac{c_1}{2L^{2\alpha_1}}$ as desired. □

With these preparations we finally prove the upper bound.

Proof of Theorem 1.3, upper bound. Since the adjacent paths in scale-free percolation are positively correlated (by Proposition 2.16) and the probability of the compound edge "$x \sim A \sim y$" decays algebraically with exponent $2\alpha_1 - d\beta$ (by Proposition 2.19), we have that the probability of a path being open in SFP dominates that in LRP with edge probability decaying with exponent $2\alpha_1 - d\beta$ in (1.15). Therefore, the graph distance in SFP in this case is no more than twice the distance in long-range percolation with connection probability as in (1.15) but with α replaced by $2\alpha_1 - d\beta$. Since one can choose β arbitrarily close to 1, the result follows from Theorem 1.8. □

2.3 Further discussion

From the previous sections one might realize that the methods we applied to prove both upper and lower bound relies significantly on the estimates of path probability. Therefore the heterogeneity of exponents in path probability of SFP, e.g. in Proposition 2.17, leads to great difficulties in identifying the correct logarithmic exponent. In contrast, the homogeneity of exponents in long-range percolation makes the problem more tractable. As we have seen in Section 1.4.1, abundant results have been achieved for LRP [7, 9, 12, 13, 14, 34, 72, 76]. In view of this, we tried to “homogenize” the path probability in the proofs for both upper and lower bound, in order to make use of the mature techniques developed since long for LRP. In Section 2.1 we used the Cauchy-Schwarz inequality in 2.6 to obtain an upper bound with same exponent for all edges involved in a path. In Section 2.2, we constructed a path composed of purely double edges, and obtained path probability with same exponent for all double edges as well.

However, in the estimates for upper bound of path probability in Lemma 2.6, as well as in the estimates for lower bound in Proposition 2.17, it seem that we were unable to identify the exact behavior of path probability, if $\tau \in (2,3)$. Based on the information from Lemma 2.6, Proposition 2.16 and Proposition 2.17 it is reasonable to conjecture that the probability of a path with even length has the following asymptotics:

$$\mathbb{P}(\pi) \approx \prod_{i=1}^{n} (|x_{2i-2} - x_{2i-1}| \vee |x_{2i-1} - x_{2i}|)^{-\alpha} \cdot \\ (|x_{2i-2} - x_{2i-1}| \wedge |x_{2i-1} - x_{2i}|)^{-\alpha(\tau-2)}, \tag{2.23}$$

where $\pi := (z_0, z_1, \ldots, z_{2n}) \in (\mathbb{Z}^d)^{2n+1}$ is the path.

In fact it turns out that there is ample interaction between all neighboring vertices, and hence the behavior of path probability is much more sophisticated than (2.23) (the simple case $n = 2$ will be computed below), posing a major hurdle to solve the graph distance problem.

Lemma 2.20. *Let $\pi = (x_0, x_1, \ldots, x_4)$ be a path of length 4. One has the following estimation:*

1. *if $|x_0 - x_1| < |x_1 - x_2|$ and $|x_3 - x_4| < |x_2 - x_3|$, then*

   1.1. *if $|x_0 - x_1||x_2 - x_3| < |x_1 - x_2||x_3 - x_4|$, then*

$$\mathbb{P}(\pi) \approx \frac{1}{|x_0 - x_1|^{\alpha(\tau-2)}|x_1 - x_2|^{\alpha}|x_2 - x_3|^{\alpha(\tau-2)}|x_3 - x_4|^{\alpha}} \tag{2.24}$$

   1.2. *else*

$$\mathbb{P}(\pi) \approx \frac{1}{|x_0 - x_1|^{\alpha}|x_1 - x_2|^{\alpha(\tau-2)}|x_2 - x_3|^{\alpha}|x_3 - x_4|^{\alpha(\tau-2)}}$$

2. *else*

$$\mathbb{P}(\pi) \approx \frac{1}{(|x_0 - x_1| \vee |x_1 - x_2|)^{\alpha} (|x_0 - x_1| \wedge |x_1 - x_2|)^{\alpha(\tau-2)}} \cdot \\ \frac{1}{(|x_2 - x_3| \vee |x_3 - x_4|)^{\alpha} (|x_2 - x_3| \wedge |x_3 - x_4|)^{\alpha(\tau-2)}}.$$

As in Section 2.2 we want to apply the double edge method by summing up all the middle vertices x_1 and x_3. However in (2.24) if x_3 is very close to x_4, for example in the extreme case $|x_3 - x_4| = |x_0 - x_1| = 1$, then

$$\mathbb{P}(x_0 \overset{2}{\sim} x_2 \overset{2}{\sim} x_4) \gtrsim \frac{1}{|x_0 - x_2|^{\alpha}} \frac{1}{|x_2 - x_4|^{\alpha(\tau-2)}} \tag{2.25}$$

Mind that the result in (2.25) holds for $|x_0 - x_2| > |x_2 - x_4|$, and is consistent with Proposition 2.17.

In order to get rid of the exponent $\alpha(\tau - 2)$ which is too small, we would have to sum up all the middle points x_2 again to obtain an upper bound for a quadruple edge, and this is not what we want.

Chapter 3

Navigation in scale-free percolation

In the previous chapter we see that the graph distance is asymptotically poly-logarithmic for $\alpha \in (d, 2d)$ and $\gamma > 2$. However, in the proofs of both upper and lower bound, we didn't construct the shortest path. For the lower bound we see in Proposition 2.12 that $\{D(x,y) < k\}$ can be contained in some event that is easier to estimate. For the upper bound we constructed some path consisting of double edges. On the one hand, this path of even length is not necessarily the shortest path. On the other hand, we don't possess so much knowledge about this path, except that some of its adjacent edges are of comparable lengths in the sense of Corollary 2.18. Therefore, it is doubtful whether scale-free percolation is navigable in this regime.

In contrast to poly-logarithmic case, Deijfen et al. [28] constructed a path by connecting vertices with highest weights in boxes of certain size in the doubly logarithmic regime. This implies that there might be a decentralized algorithm that can follow this route, or at least in the initial steps, as we will see in Section 3.2.

Let s and t be the start and the target respectively. As mentioned in Section 1.2, we will use the following connection probability

$$p_{x,y} = \frac{W_x W_y}{|x-y|^\alpha} \wedge 1$$

in scale-free percolation throughout this chapter.

We first introduce some useful results that will be frequently used later.

Lemma 3.1 (Chernoff bound). *Let* $\{A_i\}_{i\in[n]}$ *be independent events, and* $N := \sum_{i=1}^{n} \mathbb{1}_{A_n}$. *Denote* $\mu := \mathbb{E}[N]$. *Then*

a) If $\mu \geq K$ *for some constant* $K > 0$, *then*

$$\mathbb{P}\left(N \geq 1\right) \geq 1 - e^{-K}.$$

b) Let $\delta \in [0,1]$, *then*

$$\mathbb{P}(|N-\mu| > \delta\mu) \leq 2e^{-\delta^2\mu/3}.$$

The proof of part (a) proceeds with the application of the following inequality:

$$1 - x \leq e^{-x}, \quad \forall x \in [0,1],$$

together with the independence of events. Part (b) is a result of Theorem 4.4 and Theorem 4.5 in [71].

Lemma 3.2. *Let* α *and* d *be two constants.*

a) If $\alpha > d$, *then there exists some constant* $C_1 := C_1(\alpha, d) > 0$ *such that*

$$\sum_{x\in\mathbb{Z}^d:|x|>K} \frac{1}{|x|^\alpha} \leq \frac{C_1}{K^{\alpha-d}},$$

for all $K > 0$*;*

b) If $\alpha < d$, *then there exists some constant* $C_2 := C_2(\alpha, d) > 0$ *such that*

$$\sum_{x\in\mathbb{Z}^d:0<|x|\leq K} \frac{1}{|x|^\alpha} \leq C_2 K^{d-\alpha},$$

for all $K > 0$.

The proof of Lemma 3.2 follows from the comparison with the multiple integral of $|x|^{-\alpha}$ on the corresponding domain in $\mathbb{R}^d$.

3.1 Nonnavigability in the logarithmic regime

In this section we consider the regime when $\alpha \in (d, 2d)$ and $\gamma > 2$, where the graph distance of scale-free percolation is poly-logarithmic in the Euclidean distance. We show that any decentralized algorithm fails to satisfy $X_{s,t} \approx D(s,t)$ in both strong and weak sense. More precisely, we have

Theorem 1.5. *Consider scale-free percolation with connection probability* $p_{x,y} = \frac{\lambda W_x W_y}{|x-y|^\alpha} \wedge 1$*, and parameters* $\alpha \in (d, 2d), \gamma > 2$*. Let* T *be an arbitrary decentralized algorithm. Then there exists a constant* $\delta > 0$ *such that*

$$\lim_{N\to\infty} \mathbb{P}(X^T_{s,t} \geq N^\delta) = 1.$$

Here s and t are the start and the target respectively, and $N := |s-t|$ is the Euclidean distance between s and t.

In order to proof Theorem 1.5, we need the following lemma:

Lemma 3.3. *Let* $x \in \mathbb{Z}^d$*, and* $K(w)$ *be the number of neighbors of* x *with weight at least* w*. Then conditional on the weight of* x*, we have*

$$\mathbb{P}\left[K(w) \geq 1 | W_x = u\right] \leq c u^{\frac{d}{\alpha}} w^{\frac{d(1-\gamma)}{\alpha}}$$

for some constant $c > 0$.

Proof. We consider first the expected number of such neighbors:

$$\begin{aligned}\mathbb{E}\left[K(w)|W_x = u\right] &= \sum_{y\in\mathbb{Z}^d: y\neq x} \mathbb{P}\left(x \sim y, W_y \geq w | W_x = u\right)\\ &= \sum_{y\in\mathbb{Z}^d: y\neq x} \int_w^\infty \left(\frac{uv}{|x-y|^\alpha} \wedge 1\right)(\tau-1)v^{-\tau}dv\end{aligned}$$

Depending on the position of y, one have two cases for the minimum in the integral:

- $|x-y|^\alpha \leq uw$. In this case one has

$$\int_w^\infty \left(\frac{uv}{|x-y|^\alpha} \wedge 1\right)(\tau-1)v^{-\tau}dv = \mathbb{P}(W_y \geq w) = \omega^{-\tau+1};$$

- $|x-y|^\alpha > uw$. Then

$$\begin{aligned}
&\int_w^\infty \left(\frac{uv}{|x-y|^\alpha} \wedge 1\right)(\tau-1)v^{-\tau}dv \\
=&\int_w^{\frac{|x-y|^\alpha}{u}} \frac{uv}{|x-y|^\alpha}(\tau-1)v^{-\tau}dv + \mathbb{P}\left(W_y \geq \frac{|x-y|^\alpha}{u}\right) \\
=&\frac{\tau-1}{\tau-2}\frac{uw^{2-\tau}}{|x-y|^\alpha} - \frac{1}{\tau-2}\frac{u^{\tau-1}}{|x-y|^{\alpha(\tau-1)}} \leq \frac{c_1 uw^{2-\tau}}{|x-y|^\alpha}
\end{aligned}$$

Therefore we get

$$\begin{aligned}
&\sum_{y\in\mathbb{Z}^d} \mathbb{P}(x \sim y, W_y \geq w | W_x = u) \\
=&\sum_{y\in\mathbb{Z}^d:|x-y|^\alpha\leq uw} w^{-\tau+1} + \sum_{y\in\mathbb{Z}^d:|x-y|^\alpha > uw} \frac{c_1 uw^{2-\tau}}{|x-y|^\alpha} \\
\leq& c_2 (uw)^{\frac{d}{\alpha}} w^{-\tau+1} + c_3 uw^{2-\tau}\frac{1}{(uw)^{\frac{\alpha-d}{\alpha}}} = cu^{\frac{d}{\alpha}} w^{\frac{d(1-\gamma)}{\alpha}}.
\end{aligned}$$

The result follows the from Markov's inequality. □

Proof of Theorem 1.5. Let $s =: X_0, X_1, \ldots, X_n$ be the path found by T, and $\pi_k(T)$ be the path till step k. That is, $\pi_k(T) = (X_0, X_1, \ldots, X_k)$. If $X_{s,t}^T \leq N^\delta$, then by triangle inequality, the algorithm T must have one jump that is at least $N^{1-\delta}$ long. Denote by E_k the event that such a jump happens at step k for the first time. Then,

$$\mathbb{P}(X_{s,t}^T \leq N^\delta) \leq \mathbb{P}\left(\bigcup_{k=1}^{N^\delta} E_k\right) = \sum_{k=1}^{N^\delta} \mathbb{P}(E_k).$$

Let $\Pi_k(x,y)$ be the collection of self-avoiding paths connecting x and y with length k. For an arbitrary $\epsilon > 0$, we can bound the probability of E_k in the following way:

$$\begin{aligned}
\mathbb{P}(E_k) \leq& \sum_{x,y\in\mathbb{Z}^d:|x-y|^\alpha > N^{1-\delta}} \sum_{\pi\in\Pi_{k-1}(s,x):y\notin\pi} \mathbb{P}(\pi_{k-1}(T) = \pi, X_k = y) \\
\leq& \sum_{x,y\in\mathbb{Z}^d:|x-y|^\alpha > N^{1-\delta}} \sum_{\pi\in\Pi_{k-1}(s,x):y\notin\pi} \mathbb{P}(\pi_{k-1}(T) = \pi, X_k = y, W_{X_{k-1}} \leq N^\epsilon, W_{X_k} \leq N^\epsilon) \\
&+ \sum_{x,y\in\mathbb{Z}^d:|x-y|^\alpha > N^{1-\delta}} \sum_{\pi\in\Pi_{k-1}(s,x):y\notin\pi} \mathbb{P}(\pi_{k-1}(T) = \pi, X_k = y, W_{X_{k-1}} > N^\epsilon) \\
&+ \sum_{x,y\in\mathbb{Z}^d:|x-y|^\alpha > N^{1-\delta}} \sum_{\pi\in\Pi_{k-1}(s,x):y\notin\pi} \mathbb{P}(\pi_{k-1}(T) = \pi, X_k = y, W_{X_k} > N^\epsilon)
\end{aligned}$$

Apparently the last two sums can be bounded from above by $\mathbb{P}(W_{X_{k-1}} > N^\epsilon)$ and $\mathbb{P}(W_{X_k} > N^\epsilon)$ respectively. For these two probabilities one has the following upper bound:

$$\begin{aligned}
\mathbb{P}(W_{X_k} > N^\epsilon) =& \mathbb{P}(W_{X_k} > N^\epsilon | W_{X_{k-1}} > N^\epsilon)\mathbb{P}(W_{X_{k-1}} > N^\epsilon) \\
&+ \mathbb{P}(W_{X_k} > N^\epsilon | W_{X_{k-1}} \leq N^\epsilon)\mathbb{P}(W_{X_{k-1}} \leq N^\epsilon) \\
\leq& \mathbb{P}(W_{X_{k-1}} > N^\epsilon) + c\,(N^\epsilon)^{\frac{d}{\alpha}}\,(N^\epsilon)^{\frac{d(1-\gamma)}{\alpha}} \\
\leq& \mathbb{P}(W_s > N^\epsilon) + ckN^{\frac{\epsilon d(2-\gamma)}{\alpha}} = N^{\epsilon(1-\tau)} + ckN^{\frac{\epsilon d(2-\gamma)}{\alpha}}.
\end{aligned}$$

The second last line follows from Lemma 3.3.

Now we bound the probability in the first sum from above. Since the path π does not contain y, then conditional on the weights W_x, W_y, we make use of the independence of the edge $\{x, y\}$ and the event $\pi_{k-1}(T) = \pi$. More precisely,

$$\begin{aligned}
&\mathbb{P}(\pi_{k-1}(T) = \pi, X_k = y, W_{X_{k-1}} \leq N^\epsilon, W_{X_k} \leq N^\epsilon) \\
\leq& \mathbb{P}(\pi_{k-1}(T) = \pi, x \sim y, W_x \leq N^\epsilon, W_y \leq N^\epsilon) \\
=& \int_1^{N^\epsilon}\int_1^{N^\epsilon} \mathbb{P}(\pi_{k-1}(T) = \pi | W_x = u, W_y = v) \\
&\quad \cdot \mathbb{P}(x \sim y | W_x = u, W_y = v)\mu(du)\mu(dv) \\
=& \int_1^{N^\epsilon}\int_1^{N^\epsilon} \mathbb{P}(\pi_{k-1}(T) = \pi | W_x = u, W_y = v) \\
&\quad \cdot \left(\frac{uv}{|x-y|^\alpha} \wedge 1\right)\mu(du)\mu(dv) \\
\leq& \frac{N^{2\epsilon}}{|x-y|^\alpha}\mathbb{P}\left(\pi_{k-1}(T) = \pi\right),
\end{aligned}$$

where μ is the law of W_x. Consequently we have

$$\begin{aligned}
&\sum_{x,y\in\mathbb{Z}^d:|x-y|^\alpha > N^{1-\delta}}\ \sum_{\pi\in\Pi_{k-1}(s,x):y\notin\pi} \mathbb{P}(\pi_{k-1}(T) = \pi, X_k = y, W_{X_{k-1}} \leq N^\epsilon, W_{X_k} \leq N^\epsilon) \\
\leq& \sum_{x,y\in\mathbb{Z}^d:|x-y|^\alpha > N^{1-\delta}}\ \sum_{\pi\in\Pi_{k-1}(s,x):y\notin\pi} \frac{N^{2\epsilon}}{|x-y|^\alpha}\mathbb{P}(\pi_{k-1}(T) = \pi) \\
=& \sum_{x\in\mathbb{Z}^d}\ \sum_{y\in\mathbb{Z}^d:|x-y|^\alpha > N^{1-\delta}} \frac{N^{2\epsilon}}{|x-y|^\alpha}\mathbb{P}(X_{k-1} = x) \\
\leq& \sum_{x\in\mathbb{Z}^d} \mathbb{P}(X_{k-1} = x)\frac{cN^{2\epsilon}}{N^{(\alpha-d)(1-\delta)}} = \frac{cN^{2\epsilon}}{N^{(\alpha-d)(1-\delta)}},
\end{aligned}$$

where we applied Lemma 3.2 in the second last step.

With all these preparations we can finally bound the probability of the event E:

$$\begin{aligned}\mathbb{P}(E) = \sum_{k=1}^{N^\delta} \mathbb{P}(E_k) \leq \sum_{k=1}^{N^\delta} \frac{cN^{2\epsilon}}{N^{(\alpha-d)(1-\delta)}} + 2\sum_{k=0}^{N^\delta} \mathbb{P}(W_{X_k} > N^\epsilon) \\ \leq N^\delta \cdot cN^{2\epsilon+(1-\delta)(d-\alpha)} + 2(N^\delta+1)N^{\epsilon(1-\tau)} + 2\sum_{k=1}^{N^\delta} ckN^{\frac{\epsilon d(2-\gamma)}{\alpha}} \\ \leq cN^{\delta+2\epsilon+(1-\delta)(d-\alpha)} + 4N^{\delta+\epsilon(1-\tau)} + c'N^{2\delta+\frac{\epsilon d(2-\gamma)}{\alpha}}\end{aligned}$$

Since $\alpha > d, \tau > 1$ and $\gamma > 2$, we can choose δ and ϵ so small that all the exponents on the last line above are negative, and thus conclude the proof of Theorem 1.5. ☐

Corollary 3.4. *Long-range percolation is not navigable for $\alpha \in (d, 2d)$.*

Proof. The proof of Corollary 3.4 goes similarly as for Theorem 1.5.

Let T be a decentralized algorithm on long-range percolation. Denote by s, t the start and the target respectively. Furthermore, let $N := |s-t|$ be the Euclidean distance between the start and target. Let $E(N)$ be the event that T finds the target within N^δ steps for some $\delta > 0$. If $E(N)$ happens, then among the N^δ jumps there exists at least one jump with distance at least $N^{1-\delta}$. Denote by E_i the event that this long jump happens for the first time at the i-th node x_i. By the independence of edges in long-range percolation we know

$$\mathbb{P}(E(N)) \leq \sum_{i=1}^{N^\delta} \mathbb{P}(E_i).$$

Moreover, $\mathbb{P}(E_i)$ can be bounded from above in the following way:

$$\begin{aligned}\mathbb{P}(E_i) \leq \sum_{y\in\mathbb{Z}^d:|y-x_i|\geq N^{1-\delta}} \mathbb{P}(y \sim x_i) = \sum_{y\in\mathbb{Z}^d:|y-x_i|\geq N^{1-\delta}} \frac{1}{|y-x_i|^\alpha} \\ \leq \frac{C}{N^{(1-\delta)(\alpha-d)}},\end{aligned}$$

where the last step follows from Lemma 3.2. Consequently we have the estimate for $\mathbb{P}(E)$:

$$\mathbb{P}(E(N)) \leq \sum_{i=1}^{N^\delta} \mathbb{P}(E_i) \leq \frac{N^\delta}{N^{(1-\delta)(\alpha-d)}} = \frac{1}{N^{(1-\delta)(\alpha-d)-\delta}}.$$

Since $\alpha > d$, we can choose δ so small such that $(1-\delta)(\alpha-d)-\delta > 0$. In this case it holds

$$\lim_{N\to\infty} \mathbb{P}(E(N)) = 0.$$

□

Remark 3.5. *The proof of the corollary above holds true for all α with $\alpha > d$ in long-range percolation, that is, if $\alpha > d$, then the number of steps any decentralized algorithm needs to find the target is at least polynomial in the Euclidean distance between the start and target. This is no surprise for $\alpha > 2d$, since Berger [9] showed that the graph distance $D(s,t)$ between s and t satisfies*

$$\liminf_{|s-t|\to\infty} \frac{D(s,t)}{|s-t|} > 0$$

almost surely, if for some $\alpha > 2d$ the following holds:

$$0 < \lim_{|x-y|\to\infty} \frac{p_{x,y}^{LRP}}{|x-y|^{\alpha}} < \infty.$$

3.2 Navigability in the doubly logarithmic regime

In this section we consider the regime when $\alpha > d$ and $\gamma \in (1,2)$, where scale-free percolation has doubly logarithmic graph distances.

Let T be the decentralized algorithm which obeys the greedy routing protocol in Section 1.2. We take

$$\phi(x) := \frac{W_x}{|x-t|^{\alpha}}, \quad x \in \mathbb{Z}^d$$

as our objective function for T. Besides, as mentioned before, we condition on the weights of s and t throughout this section.

Remark 3.6. *In the section, for simplicity we use the following connection proba-*

bility in scale-free percolation:

$$p_{x,y} := \frac{W_x W_y}{|x-y|^\alpha} \wedge 1,$$

and choose the corresponding objective function

$$\phi(x) := \frac{W_x}{|x-t|^\alpha}$$

for the greedy routing algorithm. For other forms of connection probability, we can also decide the corresponding objective functions. For example, if $p_{x,y} = \frac{\lambda W_x W_y}{|x-y|^\alpha} \wedge 1$, *we choose* $\phi(x) = \frac{\lambda W_x}{|x-t|^\alpha}$. *If* $p_{x,y} = 1 - \exp\left(-\frac{\lambda W_x W_y}{|x-y|^\alpha}\right)$, *one option will be* $\phi(x) = 1 - \exp\left(-\frac{\lambda W_x}{|x-t|^\alpha}\right)$.

3.2.1 Success probability of greedy routing

We first show that the greedy routing algorithm finds the target with at least positive constant probability within doubly logarithmically many steps. More precisely, we have:

Theorem 3.7 (Part (a) in Theorem 1.6)**.** *Consider scale-free percolation with connection probability* (1.9)*, and parameters* $\alpha > d, \gamma \in (1,2)$*. Let* T *be the greedy routing algorithm with objective function as in* (1.10)*. Then, conditional on* W_s *and* W_t*, with at least positive constant probability,* T *finds the target within* L_1 *steps as* $N \to \infty$*, where* L_1 *is a function of* N *given as follows:*

$$L_1 = \frac{1+o(1)}{|\log(\gamma-1)|}\left(\log\log_{W_s}\left(\phi(s)^{-1}\right) + \log\log_{W_t}\left(\phi(s)^{-1}\right)\right).$$

Given the objective function as in (1.10), one has the following heuristics for the greedy routing: At the very beginning, the vertices the greedy routing algorithm T visits have typically small weights and hence have mainly close neighbors. In this circumstance, T will first visit some vertex very close to the current one, but with much higher weight hence also higher objective, and iterate such steps until it is difficult to increase the weights. This intuition in the initial stage coincides with the path constructed in the proof of Theorem 5.1 in [28]. After this stage T already reaches some vertex with very high weight, which allows the existence of long-range connections. Then the algorithm starts to overcome the distance to the target, and ends in some vertex very close to it or even finds it.

More precisely, given w_0 and ϕ_0, the greedy routing proceeds in three stages:

- **Start stage:** The start stage consists of at most one jump. The vertex s itself has a weight at least w_0 or after the first step the greedy routing algorithm finds some neighbor x_1 of s with $\phi(x_1) \geq \phi(s)$ and $W_{x_1} \geq w_0$;

- **Main stage:** T starts in this stage with a vertex (s or x_1) with weight at least w_0. In the main stage, first the weights along the greedy path grow doubly exponentially, and then the objectives increase similarly. After the main stage, T ends in some vertex $x_{\ell+1}$ with $\phi(x_{\ell+1}) > \phi_0$;

- **End stage:** If $x_{\ell+1}$ is not t, then it is connected to t directly. In this case T finds t.

We show that all events in the three stages happen with at least positive constant probability independent of $|s-t|$.

Proposition 3.8 (Start stage)**.** *Conditional on W_s, s itself or its best neighbor x_1 has weight at least $w_0 > 1$ with at least constant positive probability independent of $N := |s-t|$. Here best neighbor means it has the largest objective among the neighbors of s.*

Proof. We henceforth assume $W_s < w_0$, otherwise the proof is trivial if $W_s \geq w_0$. Denote by A the following set

$$A := \left\{y \in \mathbb{Z}^d : |y-t| \leq |s-t|\right\}.$$

Further we denote B as the ball around s with radius $\frac{|s-t|}{2}$. That is

$$B := \left\{y \in \mathbb{Z}^d : |y-s| \leq \frac{|s-t|}{2}\right\}.$$

The relation between the sets A and B can be seen in Figure 3.1.

We show first that the best neighbor of s does not lie in $A \setminus B$ with high probability. We do it by showing that, as $|s-t|$ grows, the number of neighbors of

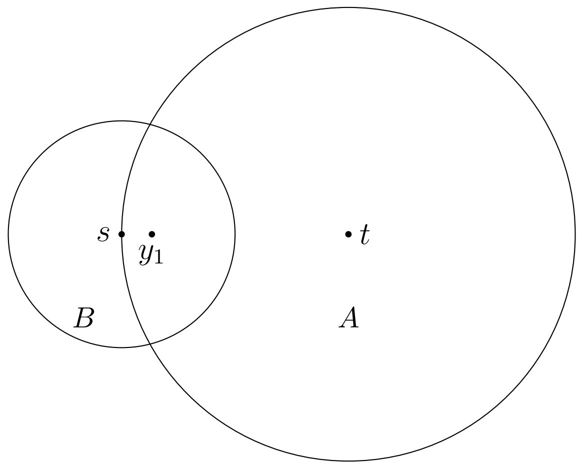

Figure 3.1: Illustration of A and B.

s in $A \setminus B$ goes to 0. Let $\mathcal{N}_s$ be the number of neighbors of s in $A \setminus B$, then

$$\begin{aligned}\mathbb{E}[\mathcal{N}_s|W_s] &= \sum_{y\in A\setminus B} \mathbb{P}\left(y \sim s|W_s\right) = \sum_{y\in A\setminus B} \mathbb{E}\left[\frac{W_sW_y}{|s-y|^\alpha} \wedge 1 \middle| W_s\right] \\ &\le \sum_{y\in A\setminus B} 2^\alpha \mathbb{E}\left[\frac{W_sW_y}{|s-t|^\alpha} \wedge 1 \middle| W_s\right] \\ &\le \sum_{y\in A\setminus B} w_0 2^\alpha \mathbb{E}\left[\frac{W_y}{|s-t|^\alpha} \wedge 1\right].\end{aligned}$$

Now we can make a case distinction to calculate the expected value.

- $\tau \in (1,2)$. Then $\mathbb{E}[W_y] = \infty$. In this case we have

$$\begin{aligned}\mathbb{E}\left[\frac{W_y}{|s-t|^\alpha} \wedge 1\right] &= \int_1^{|s-t|^\alpha} \frac{u}{|s-t|^\alpha} u^{-\tau} du + \mathbb{P}\left(W_y \ge |s-t|^\alpha\right) \\ &\le \frac{c}{|s-t|^{\alpha(\tau-1)}},\end{aligned}$$

  for some constant $c > 0$. Since $\gamma = \alpha(\tau-1)/d > 1$, we have $\alpha(\tau-1) > d$. Together with the fact that $|A \setminus B| \le C|s-t|^d$ for some constant C depending only on d, we obtain $\mathbb{E}\left[\mathcal{N}_s|W_s\right] \to 0$ as $|s-t| \to \infty$.

- $\tau = 2$.

$$\begin{aligned}\mathbb{E}\left[\frac{W_y}{|s-t|^\alpha} \wedge 1\right] &= \int_1^{|s-t|^\alpha} \frac{u}{|s-t|^\alpha} u^{-2} du + \mathbb{P}\left(W_y \ge |s-t|^\alpha\right) \\ &\le \frac{c\log|s-t|}{|s-t|^\alpha} + \frac{c}{|s-t|^\alpha},\end{aligned}$$

for some constant $c > 0$. With the same argument as in the case $\tau \in (1, 2)$ we conclude that $\mathbb{E}[\mathcal{N}_s | W_s] \to 0$ as $|s - t| \to \infty$.

- $\tau \in (2, 3)$. In this case we have $\mu := \mathbb{E}[W_y] < \infty$ and therefore

$$\mathbb{E}\left[\frac{W_y}{|s-t|^\alpha} \wedge 1\right] \leq \mathbb{E}\left[\frac{W_y}{|s-t|^\alpha}\right] = \frac{\mu}{|s-t|^\alpha},$$

Since $\alpha > d$, we have $\mathbb{E}[\mathcal{N}_s | W_s] \to 0$ as $|s - t| \to \infty$.

To summarize, for $\tau \in (1, 3)$, we have $\lim_{|s-t|\to\infty} \mathbb{E}[\mathcal{N}_s | W_s] = 0$, and therefore also $\lim_{|s-t|\to\infty} \mathbb{P}(\mathcal{N}_s = 0 | W_s) = 1$. Especially if $|s - t| > K$ for some constant $K > 0$, we have $\mathbb{P}(\mathcal{N}_s = 0 | W_s) \geq 1/2$. Note that the constant K here is independent of W_s.

Let y_1 be one nearest vertex to s in $A \cap B$. By our assumption of all nearest-neighbor edges we know $s \sim y_1$. Then

$$\mathbb{P}\left(W_{y_1} \geq 2^\alpha w_0\right) = \left(2^\alpha w_0\right)^{-\tau+1}.$$

Conditional on $W_{y_1} \geq 2^\alpha w_0$, one has $\phi(y_1) \geq \frac{2^\alpha w_0}{|s-t|^\alpha} > \phi(s)$. Let $y_{\max}$ be the best neighbor of s. Then for $|s-t| > K$ with probability at least $1/2$ we have $y_{\max} \notin A \backslash B$. In this case $|y_{\max} - t| \geq 1/2|s - t|$. By definition of best neighbor, we have

$$\phi(y_{\max}) \geq \phi(y_1) \geq \frac{2^\alpha w_0}{|s-t|^\alpha}.$$

Then

$$W_{y_{\max}} \geq \frac{2^\alpha w_0}{|s-t|^\alpha} |y_{\max} - t|^\alpha \geq w_0.$$

□

As we discussed before, the main stage is divided into two main phases and a transition phase. In the first main phase, T goes along vertices with increasing weights and hence increasing objectives; In the second phase, T visits vertices with increasing objectives, until it reaches some threshold ϕ_0 for the objective. Now we describe this procedure in more details.

We first introduce some parameters and functions.

- $\chi : \mathbb{R}^+ \to \mathbb{R}$ is a function with

$$\chi(\epsilon) := \frac{1 - \alpha\epsilon/d}{\gamma - 1}. \tag{3.1}$$

- ζ is a constant in $(1, \infty)$ such that

$$-\zeta + \frac{\alpha(\zeta - 1)}{d(\gamma - 1)} > 0. \tag{3.2}$$

  Such ζ exists because $\frac{\alpha}{d(\gamma-1)} > 1$.

- ϵ_1 is a positive constant such that $\chi(\epsilon_1) > 1$ and

$$\frac{\chi(\epsilon_1) - 1 - \chi(\epsilon)}{\chi(\epsilon_1)} + \chi(\zeta\epsilon_1) \geq 0, \tag{3.3}$$

  for all $\epsilon \in [0, \epsilon_1]$. Such ϵ_1 exists because

$$\frac{\chi(0) - 1 - \chi(0)}{\chi(0)} + \chi(0) = \frac{1}{\gamma - 1} - (\gamma - 1) > 0,$$

  since $\gamma \in (1, 2)$.

- $f_0 : \mathbb{R}^+ \to \mathbb{R}^+$ is a function with

$$\lim_{N\to\infty} f_0(N) = \infty, \quad \text{and} \quad \lim_{N\to\infty} \frac{f_0(N)}{\log\log N} = 0; \tag{3.4}$$

- $\epsilon_2 : \mathbb{R}^+ \to \mathbb{R}^+$ is a function with

$$\epsilon_2(N) := \frac{1}{\log\log f_0(N)}.$$

- $w_0' : \mathbb{R}^+ \to \mathbb{R}^+$ and $\phi_0' : \mathbb{R}^+ \to \mathbb{R}^+$ are functions with

$$w_0'(N) := w_0^{\left(\chi(\zeta\epsilon_1)^{f_0(N)}\right)}, \quad \text{and } \phi_0'(N) = \phi_0^{\left(\chi(\epsilon_1)^{f_0(N)}\right)}. \tag{3.5}$$

Since $\epsilon_2 \to 0$ as $N \to \infty$, we assume $\epsilon_2 \leq \epsilon_1$ from now on. The function χ is decreasing, and hence $\chi(\epsilon_2) \geq \chi(\epsilon_1) > 1$.

Now we define some sets of vertices in which greedy routing algorithm is mainly performed. Based on the weight configuration, we define the random sets

$$\begin{aligned}
V_{>\phi} &:= \{x \in \mathbb{Z}^d : \phi(x) > \phi\}, \\
V_1 &:= \left\{x \in \mathbb{Z}^d : \phi(x) \le W_x^{-\chi(\epsilon_1)}\right\}, \\
V_2 &:= \left\{x \in \mathbb{Z}^d : \phi(x) \ge W_x^{-\chi(\epsilon_1)}\right\}, \\
V(w,\phi) &:= \{x \in V_1 : W_x \ge w, \phi(x) \le \phi\} \cup \{x \in V_2 : \phi(x) \le \phi\} \\
V_1' &:= \left\{x \in V(w_0,\phi_0) : \phi(x) \le W_x^{-\chi(\epsilon_2)}\right\}
\end{aligned}$$

Since we choose $\epsilon_2 \le \epsilon_1$, we get $V_1' \subseteq V_1$. Later on we will see that greedy routing takes place mostly in $V(w_0, \phi_0)$. In the first phase, T explores vertices in V_1, and in the second phase, it visits V_2.

In the first phase, if $x \in V_1$, we define

$$\begin{aligned}
V_1^+(x,\epsilon) &:= \left\{y \in \mathbb{Z}^d : W_y \ge W_x^{\chi(\epsilon)}, \phi(y) \ge \phi(x) W_x^{\chi(\epsilon)-1}\right\} \\
V_1^-(x,\epsilon) &:= \left\{y \in \mathbb{Z}^d : W_y \le W_x^{\chi(\zeta\epsilon)}, \phi(y) \ge \phi(x) W_x^{\chi(\epsilon)-1}\right\}
\end{aligned}$$

In the second phase, if $x \in V_2$, we define

$$\begin{aligned}
V_2^+(x,\epsilon) &:= \left\{y \in V_2 : \phi(y) \ge \phi(x)^{1/\chi(\epsilon)}\right\} \\
V_2^-(x,\epsilon) &:= \left\{y \in V_1 : \phi(y) \ge \phi(x)^{1/\chi(\epsilon)}\right\}
\end{aligned}$$

$V_1^+(x,\epsilon)$ and $V_2^+(x,\epsilon)$ contain 'good' neighbors of x, which we expect T to visit in the next jump from x. In contrast, $V_1^-(x,\epsilon)$ and $V_2^-(x,\epsilon)$ include 'bad' neighbors that T should avoid. In addition, we denote by $\Gamma(x)$ the set of neighbors of x. Also it is reasonable to assume that $\phi(x) < 1$ for all the vertices in the greedy path. Otherwise we have $p_{x,t} = \phi(x) W_t \wedge 1 = 1$ for some x, and T jumps to t from x according to the routing protocol.

Furthermore, we divide the sets V_1 and V_2 into smaller layers:

Layers in the first phase:

Let $(z_j)_{j\in\mathbb{N}}$ be an increasing sequence with $z_0 := w_0$ defined recursively in the following way

$$z_{j+1} = \begin{cases} z_j^{\chi(\zeta\epsilon_1)} & \text{if } z_j < w_0', \\ z_j^{\chi(\zeta\epsilon_2)} & \text{otherwise,} \end{cases} \tag{3.6}$$

where w_0' is defined in (3.5). Then we define the following layers in the first phase

$$A_{1,j} := \{x \in V_1' : z_{j-1} \leq W_x < z_j\}, \quad j \geq 1, \tag{3.7}$$

Layer in the transition phase:

$$A_{1,\infty} := \left\{x \in V(w_0, \phi_0) : \phi(x) W_x^{\chi(\epsilon_1)} \leq 1 \leq \phi(x) W_x^{\chi(\epsilon_2)}\right\}. \tag{3.8}$$

Layers in the second phase:

Let $(\psi_j)_{j\in\mathbb{N}}$ be a decreasing sequence with $\psi_0 := \phi_0$ defined also recursively in the following way:

$$\psi_{j+1} = \begin{cases} \psi_j^{\chi(\epsilon_1)} & \text{if } \psi_j > \phi_0', \\ \psi_j^{\chi(\epsilon_2)} & \text{otherwise }, \end{cases} \tag{3.9}$$

where ϕ_0' is defined in (3.5). Now we define the layers in the second phase

$$A_{2,j} := \{x \in V_2 \cap V(\omega_0, \phi_0) : \psi_{j-1} \geq \phi(x) > \psi_j\}. \tag{3.10}$$

By construction of the layers one realizes

$$\bigcup_{j\in\mathbb{N}\cup\{\infty\}} A_{1,j} = V_1 \cap V(w_0, \phi_0), \quad \text{and} \bigcup_{j\in\mathbb{N}} A_{2,j} = V_2 \cap V(w_0, \phi_0). \tag{3.11}$$

Correspondingly we define two sequences $(\epsilon_1^{(j)})_{j\in\mathbb{N}\cup\{\infty\}}$ and $(\epsilon_2^j)_{j\in\mathbb{N}}$ taking values in $\{\epsilon_1, \epsilon_2\}$ based on the layers. More precisely

$$\epsilon_1^{(j)} := \begin{cases} \epsilon_1 & \text{if } z_j < w_0', \\ \epsilon_2 & \text{else.} \end{cases}$$

$$\epsilon_1^{\infty} := \epsilon_1,$$

$$\epsilon_2^{(j)} := \begin{cases} \epsilon_1 & \text{if } \psi_j > \phi_0', \\ \epsilon_2 & \text{else.} \end{cases}$$

where z_j and ψ_j are defined in (3.6) and (3.9) respectively. In other words, $\epsilon_i^{(j)}$ takes the value of $\epsilon \in \{\epsilon_1, \epsilon_2\}$ that is used in the definition of the upper bound of layer $A_{i,j+1}$.

Remark 3.9. *We will show that the greedy routing algorithm T visits each layer at* *most once, and follows the order below:*

$$A_{1,1}, \ldots, A_{1,j}, \ldots, A_{1,\infty}, \ldots, A_{2,j}, \ldots, A_{2,1}. \tag{3.12}$$

Denote $B_{i,j}$ as the union of $A_{1,1}$ to $A_{i,j}$ according to the order above, that is,

$$B_{i,j} := \begin{cases} \bigcup_{m=1}^{j} A_{1,m} & i=1, j<\infty; \\ \bigcup_{m=1}^{\infty} A_{1,m} \bigcup A_{1,\infty} & i=1, j=\infty; \\ \bigcup_{m=1}^{\infty} A_{1,m} \bigcup A_{1,\infty} \bigcup_{m=j}^{\infty} A_{2,m} & i=2. \end{cases}$$

The next two lemmas tell us it is very likely that T explores those good neighbors. We denote by $\mathbb{E}_x$ and $\mathbb{P}_x$ the expectation and the probability measure conditioned on the weight of the vertex x respectively. In other words,

$$\mathbb{E}_x[\,\cdot\,] := \mathbb{E}[\,\cdot\,|W_x], \quad \text{and} \quad \mathbb{P}_x[\,\cdot\,] := \mathbb{P}[\,\cdot\,|W_x].$$

Then we have the following results:

Lemma 3.10 (Jump in the first phase)**.** *Let x be a vertex in $A_{1,j}$ for some $j \in \mathbb{N}$. Then we have*

a) There exists some constant $c > 0$ independent of x such that

$$\mathbb{E}_x\left[\left|\Gamma(x) \cap V_1^+\left(x, \epsilon_1^{(j)}\right)\right|\right] \geq c W_x^{\epsilon_1^{(j)}};$$

b) If in addition w_0 is chosen so large that

$$w_0^{\chi(\zeta\epsilon_1)-\chi(\epsilon_1)} \leq (1/2)^\alpha, \quad \text{and} \tag{3.13}$$

$$w_0^{\left(\frac{\alpha(1-\zeta)}{d(\gamma-1)} \frac{\chi(\zeta\epsilon_1)^{f_0(N)}}{\log\log f_0(N)}\right)} \leq (1/2)^\alpha, \quad \text{for all } N \geq N_1, \tag{3.14}$$

where $N_1 > 0$ is the number such that $\frac{\chi(\zeta\epsilon_1)^{f_0(N)}}{\log\log f_0(N)} > 1$ for all $N \geq N_1$. Then we also have for all $N \geq N_1$ that

$$\mathbb{E}_x\left[\left|\Gamma(x) \cap V_1^-\left(x, \epsilon_1^{(j)}\right)\right|\right] \leq C W_x^{-\rho\epsilon_1^{(j)}} \log W_x,$$

for some positive constants $\rho := \rho(\alpha, d, \tau, \zeta)$ and $C := C(\alpha, d, \tau, \zeta)$.

Proof. a) Let $\epsilon \in \{\epsilon_1, \epsilon_2\}$ and

$$A(x,\epsilon) := \left\{ y \in \mathbb{Z}^d : W_y \geq W_x^{\chi(\epsilon)}, |y-t| \leq |x-t| \right\}.$$

For $y \in A(x,\epsilon)$ one has

$$\phi(y) = \frac{W_y}{|y-t|^\alpha} \geq \frac{W_x^{\chi(\epsilon)}}{|y-t|^\alpha} \geq \frac{W_x^{\chi(\epsilon)}}{|x-t|^\alpha} = \phi(x) W_x^{\chi(\epsilon)-1}.$$

In other words, $A(x,\epsilon) \subseteq V_1^+(x,\epsilon)$.

Then the expected size of good neighbors is

$$\begin{aligned} &\mathbb{E}_x\left[|\Gamma(x) \cap V_1^+(x,\epsilon)|\right] \geq \mathbb{E}\left[|\Gamma(x) \cap A(x,\epsilon)|\right] \\ = &\sum_{y \in \mathbb{Z}^d : |y-t| \leq |x-t|} \mathbb{P}\left(y \sim x, W_y \geq W_x^{\chi(\epsilon)}\right) \\ = &\sum_{y \in \mathbb{Z}^d : |y-t| \leq |x-t|} \int_{W_x^{\chi(\epsilon)}}^\infty \left(\frac{W_x u}{|x-y|^\alpha} \wedge 1\right) u^{-\tau} du \end{aligned}$$

Depending on the minimum in the integral above, we need to distinguish all such y's.

Let $A_1(x,\epsilon) := \left\{ y \in A(x,\epsilon) : |x-y|^\alpha \leq W_x^{\chi(\epsilon)+1} \right\}$, and $A_2(x,\epsilon) := A(x,\epsilon) \setminus A_1(x,\epsilon)$.

First for $y \in A_1(x,\epsilon)$, the integral above becomes

$$\int_{W_x^{\chi(\epsilon)}}^\infty 1 \cdot u^{-\tau} du = \frac{1}{(\tau-1) W_x^{\chi(\epsilon)(\tau-1)}}.$$

Therefore

$$\sum_{y \in A_1(x,\epsilon)} \frac{1}{(\tau-1) W_x^{\chi(\epsilon)(\tau-1)}} \geq \frac{c\left(W_x^{\frac{\chi(\epsilon)+1}{\alpha}}\right)^d}{W_x^{\chi(\epsilon)(\tau-1)}},$$

for some constant c independent of x and W_x. Here we used the following fact to estimate the volume of the intersection between two balls, as depicted in Figure 3.2.

Fact 3.11. *Let $t, x \in \mathbb{Z}^d$, and $R_1 := |x-t|$. Let $A_1 := \{y \in \mathbb{Z}^d : |y-t| \leq R_1\}$, and $A_2 := \{y \in \mathbb{Z}^d : |y-x| \leq R_2\}$. Assume $R_2 \leq R_1$, then there exists a constant* *c independent of* R_1, R_2 *such that*

$$|A_1 \cap A_2| \geq cR_2^d.$$

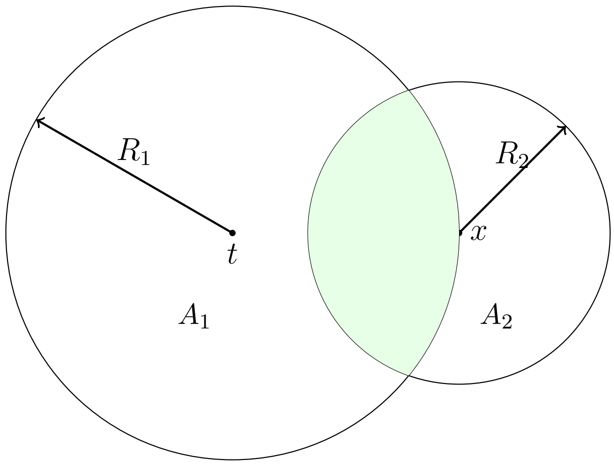

Figure 3.2: The shaded area is the intersection of A_1 and A_2.

Therefore

$$\mathbb{E}_x\left[|\Gamma(x) \cap V_1^+(x,\epsilon)|\right] \geq \frac{c\left(W_x^{\frac{\chi(\epsilon)+1}{\alpha}}\right)^d}{W_x^{\chi(\epsilon)(\tau-1)}} = cW_x^{(\chi(\epsilon)+1)d/\alpha-\chi(\epsilon)(\tau-1)} = cW_x^{\epsilon}.$$

In the last step, we plug in $\chi(\epsilon) = \frac{1-\alpha\epsilon/d}{\gamma-1}$.

b) First such $N_1 > 0$ exists because $\chi(\zeta\epsilon_1) > 1$. Besides, the existence of such w_0 is guaranteed by the fact that $\zeta > 1$. For brevity we write ϵ instead of $\epsilon_1^{(j)}$ for corresponding x in the proof of part b).

Let $B_1(x,\epsilon) := \left\{y \in \mathbb{Z}^d : |y-t|^\alpha \leq W_x^{\chi(\zeta\epsilon)+1-\chi(\epsilon)}\phi(x)^{-1}\right\}$. We have the following estimate for the volume of $B_1(x,\epsilon)$:

$$|B_1(x,\epsilon)| \leq \left(cW_x^{\chi(\zeta\epsilon)+1-\chi(\epsilon)}\phi(x)^{-1}\right)^{d/\alpha},$$

for some $c := c(d)$. Let $y \in V_1^-(x,\epsilon)$. Then we have

$$\phi(x)W_x^{\chi(\epsilon)-1}|y-t|^\alpha \leq W_y \leq W_x^{\chi(\zeta\epsilon)}.$$

Then

$$
\begin{aligned}
&\mathbb{E}_x\left[\left|\Gamma(x) \cap V_1^-(x,\epsilon)\right|\right] \\
&= \sum_{y \in B_1(x,\epsilon)} \mathbb{P}\left(\phi(x) W_x^{\gamma(\epsilon)-1}|y-t|^\alpha \leq W_y \leq W_x^{\gamma(\zeta\epsilon)}, y \sim x\right) \\
&= \sum_{y \in B_1(x,\epsilon)} \int_{\phi(x) W_x^{\gamma(\epsilon)-1}|y-t|^\alpha}^{W_x^{\gamma(\zeta\epsilon)}} \left(\frac{W_x u}{|x-y|^\alpha} \wedge 1\right) u^{-\tau} du
\end{aligned}
$$

For $y \in B_1(x,\epsilon)$, one has $|y-t|^\alpha/|x-t|^\alpha \leq W_x^{\gamma(\zeta\epsilon)-\gamma(\epsilon)}$. Now we want to show that this ratio goes to 0 if W_x grows to infinity. Depending on the value of ϵ we have the following two cases:

- $\epsilon = \epsilon_1$. In this case we know $w_0 \leq W_x < w_0'$, where w_0' is defined in (3.5). Consequently one has

$$
|y-t|^\alpha/|x-t|^\alpha \leq W_x^{\chi(\zeta\epsilon_1)-\chi(\epsilon_1)} \leq w_0^{\chi(\zeta\epsilon_1)-\chi(\epsilon_1)} \leq (1/2)^\alpha;
$$

- $\epsilon = \epsilon_2$. In this case we know $W_x \geq w_0'$. Therefore

$$
\begin{aligned}
|y-t|^\alpha/|x-t|^\alpha &\leq W_x^{\chi(\zeta\epsilon_2)-\chi(\epsilon_2)} \leq (w_0')^{\chi(\zeta\epsilon_2)-\chi(\epsilon_2)} = \left(w_0^{\left(\chi(\zeta\epsilon_1)^{f_0(N)}\right)}\right)^{\frac{\alpha(1-\zeta)\epsilon_2}{d(\gamma-1)}} \\
&= w_0^{\left(\frac{\alpha(1-\zeta)}{d(\gamma-1)} \frac{\chi(\zeta\epsilon_1)^{f_0(N)}}{\log\log f_0(N)}\right)} \leq (1/2)^\alpha.
\end{aligned}
$$

In both cases we obtain $|y-t| \leq \frac{1}{2}|x-t|$. By the triangle inequality,

$$
|x-y| \geq \frac{1}{2}|x-t|. \tag{3.15}
$$

Then we have for $u \leq W_x^{\chi(\zeta\epsilon)}$:

$$
\frac{W_x u}{|x-y|^\alpha} \wedge 1 \leq \frac{W_x u}{|x-y|^\alpha} \leq 2^\alpha \frac{W_x u}{|x-t|^\alpha},
$$

and hence

$$
\mathbb{E}_x\left[\left|\Gamma(x) \cap V_1^-\left(x,\epsilon_1^{(j)}\right)\right|\right] \leq 2^\alpha \sum_{y \in B_1\left(x,\epsilon_1^{(j)}\right)} \int_{\phi(x) W_x^{\chi\left(\epsilon_1^{(j)}\right)-1}|y-t|^\alpha}^{W_x^{\chi\left(\zeta\epsilon_1^{(j)}\right)}} \frac{W_x u}{|x-t|^\alpha} u^{-\tau} du.
$$

We calculate the integral in different cases:

i) $\tau \in (1,2)$.

$$\int_{\phi(x)W_x^{\chi(\epsilon)-1}|y-t|^\alpha}^{W_x^{\chi(\zeta\epsilon)}} \frac{W_x u}{|x-t|^\alpha} u^{-\tau} du \le \frac{\phi(x)}{2-\tau} W_x^{\chi(\zeta\epsilon)(2-\tau)}.$$

Then

$$\begin{aligned}\mathbb{E}_x\left[|\Gamma(x)\cap V_1^-(x,\epsilon)|\right] \le& \phi(x)\frac{2^\alpha}{2-\tau} W_x^{\chi(\zeta\epsilon)(2-\tau)} |B_1(x,\epsilon)| \\ \le& \phi(x)\frac{2^\alpha}{2-\tau} W_x^{\chi(\zeta\epsilon)(2-\tau)} \left(cW_x^{\chi(\zeta\epsilon)+1-\chi(\epsilon)}\phi(x)^{-1}\right)^{d/\alpha} \\ =& C\phi(x)^{1-d/\alpha} W_x^{\chi(\zeta\epsilon)(2-\tau)+d/\alpha(\chi(\zeta\epsilon)+1-\chi(\epsilon))}\end{aligned}$$

We know at the same time $\phi(x) \le W_x^{-\chi(\epsilon)}$, since $x \in A_{1,j} \subseteq V_1'$. Thus

$$\begin{aligned}\mathbb{E}\left[|\Gamma(x)\cap V_1^-(x,\epsilon)|\right] \le& CW_x^{(d/\alpha-1)\chi(\epsilon)} W_x^{\chi(\zeta\epsilon)(2-\tau)+d/\alpha(\chi(\zeta\epsilon)+1-\chi(\epsilon))} \\ =& CW_x^{\epsilon\left(\zeta+\frac{\alpha(1-\zeta)}{d(\gamma-1)}\right)}.\end{aligned}$$

Denote $\rho := -\zeta + \frac{\alpha(\zeta-1)}{d(\gamma-1)}$. By our choice of ζ in (3.2) one has $\rho > 0$.

ii) $\tau = 2$.

$$\int_{\phi(x)W_x^{\chi(\epsilon)-1}|y-t|^\alpha}^{W_x^{\chi(\zeta\epsilon)}} \frac{W_x u}{|x-t|^\alpha} u^{-\tau} du \le \chi(\zeta\epsilon)\frac{W_x \log W_x}{|x-t|^\alpha}.$$

Together with the estimate for the volume of $B_1(x,\epsilon)$ we obtain

$$\begin{aligned}\mathbb{E}_x\left[|\Gamma(x)\cap V_1^-(x,\epsilon)|\right] &\le 2^\alpha \chi(\zeta\epsilon)\frac{W_x \log W_x}{|x-t|^\alpha}\left(cW_x^{\chi(\zeta\epsilon)+1-\chi(\epsilon)}\phi(x)^{-1}\right)^{d/\alpha} \\ &\le c\phi(x)^{1-d/\alpha} W_x^{\frac{d}{\alpha}(\chi(\zeta\epsilon)+1-\chi(\epsilon))} \log W_x \\ &\le cW_x^{\frac{\gamma-\zeta}{\gamma-1}\epsilon} \log W_x.\end{aligned}$$

Note here we used the fact that for $\tau = 2$, it holds $\gamma = \alpha/d$.

iii) $\tau \in (2,3)$.

$$\int_{\phi(x)W_x^{\chi(\epsilon)-1}|y-t|^\alpha}^{W_x^{\chi(\zeta\epsilon)}} \frac{W_x u}{|x-t|^\alpha} u^{-\tau} du \le \frac{\phi(x)^{3-\tau}}{\tau-2}\frac{W_x^{(\chi(\epsilon)-1)(2-\tau)}}{|y-t|^{\alpha(\tau-2)}}$$

Since $\gamma = \frac{\alpha(\tau-1)}{d} < 2$ and $\alpha > d$, one has the following estimate (see also

Figure 1.4)

$$\alpha(\tau-2)=\alpha(\tau-1)-\alpha=\gamma d-\alpha<2d-\alpha<d.$$

Then Lemma 3.2 implies:

$$\begin{aligned}\mathbb{E}_x\left[\left|\Gamma(x)\cap V_1^-(x,\epsilon)\right|\right]\le&\sum_{y\in\mathbb{Z}^d:|y-t|^\alpha\le W_x^{\chi(\zeta\epsilon)+1-\chi(\epsilon)}\phi(x)^{-1}}\frac{2^\alpha\phi(x)^{3-\tau}}{\tau-2}\frac{W_x^{(\chi(\epsilon)-1)(2-\tau)}}{|y-t|^{\alpha(\tau-2)}}\\&\le C\phi(x)^{3-\tau}W_x^{(\chi(\epsilon)-1)(2-\tau)}\left[W_x^{\chi(\zeta\epsilon)+1-\chi(\epsilon)}\phi(x)^{-1}\right]^{\frac{d-\alpha(\tau-2)}{\alpha}}\\&\le CW_x^{(d/\alpha-1)\chi(\epsilon)}W_x^{\chi(\zeta\epsilon)(2-\tau)+d/\alpha(\chi(\zeta\epsilon)+1-\chi(\epsilon))}\\&=CW_x^{\epsilon\left(\zeta+\frac{\alpha(1-\zeta)}{d(\gamma-1)}\right)}.\end{aligned}$$

□

Lemma 3.12 (Jump in the second phase)**.** *Let x be a vertex in the layer $A_{2,j}$ for some $j\in\mathbb{N}$. Then we have*

a) There exist constants $c,c'>0$ independent of x such that

$$\mathbb{E}_x\left[\left|\Gamma(x)\cap V_2^+\left(x,\epsilon_2^{(j)}\right)\right|\right]\ge c\phi(x)^{-c'\epsilon_2^{(j)}};$$

b) If, in addition, ϕ_0 is chosen so small that

$$\phi_0^{\left(1+\frac{1}{\chi(\epsilon_1)}\right)\left(1-\frac{1}{\chi(\epsilon_1)}\right)}\le\frac{1}{2^\alpha},\quad\text{and}\tag{3.16}$$

$$\left(\phi_0^{\left(\chi(\epsilon_1)^{f_0(N)}\right)}\right)^{\left(1+\frac{1}{\chi(\epsilon_1)}\right)\left(1-\frac{1}{\chi(\epsilon_2)}\right)}\le\frac{1}{2^\alpha},\text{for all } N\ge N_2,\tag{3.17}$$

where N_2 is a constant such that $f_0(N)\ge 1$ for all $N\ge N_2$. Then we also have for all $N\ge N_2$ that

$$\mathbb{E}_x\left[\left|\Gamma(x)\cap V_2^-\left(x,\epsilon_2^{(j)}\right)\right|\right]\le C\phi(x)^{\epsilon_2^{(j)}}\log\left(\phi(x)^{-1}\right),$$

for some constant $C:=C(\alpha,d,,\tau,\epsilon_1)$.

Proof. a) Let $\epsilon \in \{\epsilon_1, \epsilon_2\}$, and define the set

$$A(x,\epsilon) := \left\{y \in \mathbb{Z}^d : |y-t|^\alpha \le \phi(x)^{-1-1/\chi(\epsilon)}, W_y \ge \phi(x)^{-1}\right\}.$$

For $y \in A(x,\epsilon)$, one has

$$\phi(y) = \frac{W_y}{|y-t|^\alpha} \ge \frac{\phi(x)^{-1}}{\phi(x)^{-1-1/\chi(\epsilon)}} = \phi(x)^{1/\chi(\epsilon)}.$$

In order to show that $y \in V_2^+(x,\epsilon)$, we still need to show $y \in V_2$. Since $W_y \ge 1$, $\chi(\epsilon_1) \ge 1$ and $\chi(\epsilon) \ge 1$, one has

$$W_y^{1+\chi(\epsilon_1)} \ge W_y^{1+1/\chi(\epsilon)} \ge \phi(x)^{-1-1/\chi(\epsilon)} \ge |y-t|^\alpha.$$

This means $\phi(y) \ge W_y^{-\chi(\epsilon_1)}$. Therefore $y \in V_2$ and $A(x,\epsilon) \subseteq V_2^+(x,\epsilon)$.

In addition, for $y \in A(x,\epsilon)$, one has

$$|y-t|^\alpha \le \phi(x)^{-1-1/\chi(\epsilon)}, \quad \text{and} \quad \phi(x) \ge W_x^{-\chi(\epsilon_1)} \ge W_x^{-\chi(\epsilon)}.$$

Then $|y-t|^\alpha \le \phi(x)^{-1} W_x^{\chi(\epsilon)1/\chi(\epsilon)} = |x-t|^\alpha$. We know by triangle inequality that $|x-y| \le 2|x-t|$.

Now we can estimate the size of the set of good neighbors in the second phase:

$$\begin{aligned}
&\mathbb{E}_x\left[|\Gamma(x) \cap V_2^+(x,\epsilon)|\right] \ge \mathbb{E}_x\left[|\Gamma(x) \cap A(x,\epsilon)|\right] \\
=&\sum_{y\in\mathbb{Z}^d:|y-t|^\alpha\le\phi(x)^{-1-1/\chi(\epsilon)}} \int_{\phi(x)^{-1}}^{\infty}\left(\frac{W_x u}{|x-y|^\alpha}\wedge 1\right)u^{-\tau}du \\
\ge&\sum_{y\in\mathbb{Z}^d:|y-t|^\alpha\le\phi(x)^{-1-1/\chi(\epsilon)}} \frac{1}{2^\alpha}\int_{\phi(x)^{-1}}^{\infty}\left(\frac{W_x u}{|x-t|^\alpha}\wedge 1\right)u^{-\tau}du \\
=&\sum_{y\in\mathbb{Z}^d:|y-t|^\alpha\le\phi(x)^{-1-1/\chi(\epsilon)}} \frac{1}{2^\alpha}\int_{\phi(x)^{-1}}^{\infty} u^{-\tau}du \\
\ge& c\left(\phi(x)^{-1-1/\chi(\epsilon)}\right)^{d/\alpha}\phi(x)^{\tau-1} = c\phi(x)^{-\frac{\gamma-1}{1-\alpha\epsilon/d}\epsilon} \\
\ge& c\phi(x)^{-(\gamma-1)\epsilon}.
\end{aligned}$$

b) It is clear that such N_2 exists, because $\lim_{N\to\infty} f_0(N) = \infty$. By the construction of the function χ in (3.1), we know there is a constant $K > 0$ such that

$$\left(-1 - \frac{1}{\chi(\epsilon_1)}\right)\left(\frac{1}{\chi(\epsilon_2)} - 1\right) \geq K.$$

Consequently for $N \geq N_2$ one has

$$\left(\phi_0^{\left(\chi(\epsilon_1)^{f_0(N)}\right)}\right)^{\left(-1-\frac{1}{\chi(\epsilon_1)}\right)\left(\frac{1}{\chi(\epsilon_2)}-1\right)} \leq \phi_0^{K\chi(\epsilon_1)}.$$

Therefore such ϕ_0 satisfying (3.16) and (3.17) exists. Again, for brevity we write ϵ instead of $\epsilon_2^{(j)}$ for corresponding x in the proof of part b).

Let $y \in V_2^-(x, \epsilon)$. Then

$$\phi(y) \geq \phi(x)^{1/\chi(\epsilon)}, \quad \phi(y) \leq W_y^{-\chi(\epsilon_1)}.$$

Consequently we have the following estimate for the distance between y and t:

$$|y-t|^\alpha = \phi(y)^{-1}W_y \leq \phi(y)^{-1}\phi(y)^{-1/\chi(\epsilon_1)} \leq \left(\phi(x)^{-1}\phi(x)^{-1/\chi(\epsilon_1)}\right)^{1/\chi(\epsilon)}$$

Now we want to show that $\left(\phi(x)^{-1}\phi(x)^{-1/\chi(\epsilon_1)}\right)^{1/\chi(\epsilon)-1}$ can be small by choosing ϕ_0 properly.

Depending on the value of ϵ we have two possible cases:

- $\epsilon = \epsilon_1$. In this case we know $\phi_0' < \phi(x) < \phi_0$, where ϕ_0' is defined in (3.5). Hence

$$\begin{aligned}\left(\phi(x)^{-1}\phi(x)^{-1/\chi(\epsilon_1)}\right)^{1/\chi(\epsilon)-1} &= \phi(x)^{\left(-1-\frac{1}{\chi(\epsilon_1)}\right)\left(\frac{1}{\chi(\epsilon_1)}-1\right)} \\ &\leq \phi_0^{\left(-1-\frac{1}{\chi(\epsilon_1)}\right)\left(\frac{1}{\chi(\epsilon_1)}-1\right)} \leq \frac{1}{2^\alpha};\end{aligned}$$

- $\epsilon = \epsilon_2$. In this case we have $\phi(x) \leq \phi_0'$. By plugging ϕ_0' in we obtain

$$\begin{aligned}\left(\phi(x)^{-1}\phi(x)^{-1/\chi(\epsilon_1)}\right)^{1/\chi(\epsilon)-1} &= \phi(x)^{\left(-1-\frac{1}{\chi(\epsilon_1)}\right)\left(\frac{1}{\chi(\epsilon_2)}-1\right)} \\ \leq \phi_0'^{\left(-1-\frac{1}{\chi(\epsilon_1)}\right)\left(\frac{1}{\chi(\epsilon_2)}-1\right)} &= \left(\phi_0^{\left(\chi(\epsilon_1)^{f_0(N)}\right)}\right)^{\left(-1-\frac{1}{\chi(\epsilon_1)}\right)\left(\frac{1}{\chi(\epsilon_2)}-1\right)} \leq \frac{1}{2^\alpha}.\end{aligned}$$

In both cases we know

$$\begin{aligned}\left(\phi(x)^{-1}\phi(x)^{-1/\chi(\epsilon_1)}\right)^{1/\chi(\epsilon)} &\le \frac{1}{2^\alpha}\phi(x)^{-1}\phi(x)^{-1/\chi(\epsilon_1)} \\ &\le (1/2)^\alpha\phi(x)^{-1}W_x = (1/2|x-t|)^\alpha.\end{aligned}$$

In the last line we used the fact that $x \in V_2$ and hence $\phi(x) \ge W_x^{-\chi(\epsilon_1)}$. By triangle inequality, $|x-y| \ge 1/2|x-t|$.

Denote $B(x,\epsilon) := \left\{y \in \mathbb{Z}^d : \phi(x) \le |y-t|^{-\frac{\chi(\epsilon)\chi(\epsilon_1)\alpha}{1+\chi(\epsilon_1)}}\right\}$. Then

$$\begin{aligned}\mathbb{E}_x\left[|\Gamma(x)\cap V_2^-(x,\epsilon)|\right] &= \sum_{y\in B(x,\epsilon)} \mathbb{P}(y\sim x, \phi(x)^{1/\chi(\epsilon)}|y-t|^\alpha \le W_y \le |y-t|^{\frac{\alpha}{1+\chi(\epsilon_1)}}) \\ &= \sum_{y\in B(x,\epsilon)} \int_{\phi(x)^{1/\chi(\epsilon)}|y-t|^\alpha}^{|y-t|^{\frac{\alpha}{1+\chi(\epsilon_1)}}} \left(\frac{W_x u}{|x-y|^\alpha}\wedge 1\right) u^{-\tau} du \\ &\le \sum_{y\in B(x,\epsilon)} 2^\alpha \int_{\phi(x)^{1/\chi(\epsilon)}|y-t|^\alpha}^{|y-t|^{\frac{\alpha}{1+\chi(\epsilon_1)}}} \left(\frac{W_x u}{|x-t|^\alpha}\wedge 1\right) u^{-\tau} du\end{aligned}$$

For $y \in B(x,\epsilon)$ one has $\phi(x) \le |y-t|^{-\frac{\chi(\epsilon)\chi(\epsilon_1)\alpha}{1+\chi(\epsilon_1)}}$. Consequently $|y-t|^\alpha \le \phi(x)^{-\frac{1+\chi(\epsilon_1)}{\chi(\epsilon)\chi(\epsilon_1)}}$, and

$$\phi(x)|y-t|^{\frac{\alpha}{1+\chi(\epsilon_1)}} \le \phi(x)\phi(x)^{-\frac{1}{\chi(\epsilon)\chi(\epsilon_1)}} \le 1.$$

Then we can get rid of the minimum in the integrand above, and the integral becomes

$$\int_{\phi(x)^{1/\chi(\epsilon)}|y-t|^\alpha}^{|y-t|^{\frac{\alpha}{1+\chi(\epsilon_1)}}} \phi(x)u^{-\tau+1}du$$

For different values of τ we have the following 3 cases:

- $\tau \in (1,2)$. In this case the integral becomes:

$$\int_{\phi(x)^{1/\chi(\epsilon)}|y-t|^\alpha}^{|y-t|^{\frac{\alpha}{1+\chi(\epsilon_1)}}} \phi(x)u^{-\tau+1}du \le \frac{1}{2-\tau}\phi(x)|y-t|^{\frac{\alpha(2-\tau)}{1+\chi(\epsilon_1)}}.$$

Lemma 3.2 implies:

$$\begin{aligned}\mathbb{E}_x\left[|\Gamma(x)\cap V_2^-(x,\epsilon)|\right] &\le \frac{2^\alpha}{2-\tau}\sum_{y\in B(x,\epsilon)}\phi(x)|y-t|^{\frac{\alpha(2-\tau)}{1+\chi(\epsilon_1)}}\\ &\le C\phi(x)\left[\phi(x)^{-\frac{1+\chi(\epsilon_1)}{\chi(\epsilon)\chi(\epsilon_1)\alpha}}\right]^{\frac{\alpha(2-\tau)}{1+\chi(\epsilon_1)}+d}\\ &= C\phi(x)^{f(\epsilon,\epsilon_1)}.\end{aligned}$$

where $f(\epsilon,\epsilon_1) := 1-\frac{1+\chi(\epsilon_1)}{\chi(\epsilon)\chi(\epsilon_1)\alpha}\left(\frac{\alpha(2-\tau)}{1+\chi(\epsilon_1)}+d\right)$.

$f(\epsilon,\epsilon_1)$ is a continuous function of ϵ and ϵ_1, and

$$f(0,0) = (2-\gamma)\left[1+(2-\tau)(\gamma-1)\right] > 0.$$

Therefore we can choose ϵ_1 so small that $f(\epsilon,\epsilon_1)\ge\epsilon$ for all $\epsilon\in(0,\epsilon_1]$. Since $\phi(x)\le 1$, we have $\mathbb{E}_x\left[|\Gamma(x)\cap V_2^-(x,\epsilon)|\right]\le C\phi(x)^\epsilon$.

- $\tau = 2$. We compute the integral in the following way:

$$\begin{aligned}\int_{\phi(x)^{1/\chi(\epsilon)}|y-t|^\alpha}^{|y-t|^{\frac{\alpha}{1+\chi(\epsilon_1)}}}\phi(x)u^{-\tau+1}du &= \phi(x)\int_{\phi(x)^{1/\chi(\epsilon)}|y-t|^\alpha}^{|y-t|^{\frac{\alpha}{1+\chi(\epsilon_1)}}}u^{-1}du\\ &\le c\phi(x)\log|y-t|.\end{aligned}$$

Therefore we have the following estimate:

$$\begin{aligned}\mathbb{E}_x\left[|\Gamma(x)\cap V_2^-(x,\epsilon)|\right] &\le c\phi(x)\phi(x)^{-\frac{d}{\alpha}\frac{1+\chi(\epsilon_1)}{\chi(\epsilon)\chi(\epsilon_1)}}\log|y-t|\\ &\le \frac{c}{\chi(\epsilon)}\phi(x)^{f(\epsilon,\epsilon_1)}\log\frac{1}{\phi(x)},\end{aligned}$$

where $f(\epsilon,\epsilon_1) := 1-\frac{d}{\alpha}\frac{1+\chi(\epsilon_1)}{\chi(\epsilon)\chi(\epsilon_1)}$. Note for $\tau=2$ we have $\gamma=\alpha/d$, and

$$f(0,0)=\gamma\left[1-(\gamma-1)^2\right] > 0.$$

With the same argument as in the case $\tau\in(1,2)$ we can choose ϵ_1 so small that $f(\epsilon,\epsilon_1)\ge\epsilon$ for all $\epsilon\in(0,\epsilon_1]$. Together with the fact that $\chi(\epsilon)\ge\chi(\epsilon_1)$, we obtain the following upper bound

$$\mathbb{E}_x\left[|\Gamma(x)\cap V_2^-(x,\epsilon)|\right]\le C\phi(x)^\epsilon\log\left(\phi(x)^{-1}\right).$$

- $\tau \in (2,3)$. The integral can be simplified as follows:

$$\int_{\phi(x)^{1/\chi(\epsilon)}|y-t|^{\alpha}}^{|y-t|^{\frac{\alpha}{1+\chi(\epsilon_1)}}} \phi(x)u^{-\tau+1}du \le c\phi(x)\phi(x)^{-\frac{\tau-2}{\chi(\epsilon)}}|y-t|^{-\alpha(\tau-2)}.$$

Lemma 3.2 allows us to make following estimation:

$$\begin{aligned}\mathbb{E}_x\left[|\Gamma(x)\cap V_2^-(x,\epsilon)|\right] &\le \sum_{y\in B(x,\epsilon)} c\phi(x)\phi(x)^{-\frac{\tau-2}{\chi(\epsilon)}}|y-t|^{-\alpha(\tau-2)} \\ &\le c'\phi(x)^{1-\frac{\tau-2}{\chi(\epsilon)}}\left[\phi(x)^{-\frac{1+\chi(\epsilon_1)}{\alpha\chi(\epsilon)\chi(\epsilon_1)}}\right]^{-\alpha(\tau-2)+d} \\ &= c'\phi(x)^{f(\epsilon,\epsilon_1)},\end{aligned}$$

where $f(\epsilon,\epsilon_1) := 1-\frac{1+\chi(\epsilon_1)}{\chi(\epsilon)\chi(\epsilon_1)\alpha}\left(\frac{\alpha(2-\tau)}{1+\chi(\epsilon_1)}+d\right)$.

Again we play the trick as in the case $\tau \in (1,2)$ and can find some ϵ_1 so small that $f(\epsilon,\epsilon_1)\ge\epsilon$ for all $\epsilon\in(0,\epsilon_1]$. Therefore we have $\mathbb{E}_x\left[|\Gamma(x)\cap V_2^-(x,\epsilon)|\right] \le C\phi(x)^{\epsilon}$.

To summarize, in all three cases we have the following upper bound for the expected size of 'bad' neighbors:

$$\mathbb{E}_x\left[|\Gamma(x)\cap V_2^-(x,\epsilon)|\right] \le C\phi(x)^{\epsilon}\log\left(\phi(x)^{-1}\right).$$

□

Lemma 3.10 and Lemma 3.12 suggest that the objectives of vertices in the greedy path grow doubly exponentially, that is, given a vertex x in the greedy path with objective $\phi(x)$, the k-th vertex x_k after x has objective

$$\phi(x_k)\approx\phi(x)^{a^k},$$

for some constant $a<1$. We will come to this in Part (c) and (d) in Proposition 3.13.

Now we describe the typical trajectory of greedy routing. Note that every hop of T depends highly on the current vertex, leading to a dependence structure between consecutive jumps. To overcome the dependence, we introduce so-called 'layers' in $\mathbb{Z}^d$, and show that with high probability, the greedy path will traverse each layer at most once. Depending on the phases, we introduce different layers as follows:

The next proposition tells us that after traversing at most $\mathcal{O}(\log\log N)$ layers as defined above, the greedy routing algorithm reaches some vertex with objective at least ϕ_0. In light of Proposition 3.8, we may assume that $\phi(x_1) < \phi_0$, where x_1 is the vertex found in Proposition 3.8 with $W_{x_1} \geq w_0$. Otherwise the algorithm skips the main part stage and starts with a vertex of objective larger than ϕ_0.

For the proposition about the main part stage we make two choices of w_0 and ϕ_0:

- w_0 and ϕ_0 are positive constants satisfying (3.13)(3.14) and (3.16)(3.17) respectively. Or
- w_0 and ϕ_0 are functions of N such that

$$\lim_{N\to\infty} w_0(N) = \infty, \quad \lim_{N\to\infty} \phi_0(N) = 0. \tag{3.18}$$

Note that Lemma 3.10 and 3.12 are true if w_0 and ϕ_0 satisfy (3.18) respectively. As we will see, Proposition 3.13 is valid for both choices of w_0 and ϕ_0. In this section we use the proposition with the first choice of w_0 and ϕ_0 (see the proof of Proposition 3.14). The second choice is applied in Section 3.2.2 and Section 3.2.3.

Proposition 3.13 (Main stage)**.** *Let f_0 be a function as in* (3.4)*. Then there exist constants κ and N_0 such that the following statement holds for all $N \geq N_0$:*

Let w_0 and ϕ_0 be either positive constants such that Lemma 3.10 and 3.12 hold true, or positive functions such that conditions in (3.18) *are satisfied. Furthermore, let $G_{\leq\phi_0}$ be the subgraph of scale-free percolation on $\mathbb{Z}^d$ induced by vertices of objective at most ϕ_0, and P_{ϕ_0} be the greedy path on $G_{\leq\phi_0}$ starting in s. Assume there exists a vertex $x_1 \in P_{\phi_0} \cap V(w_0, \phi_0)$, then there exists a positive constant δ, as well as $C_\delta > 0$ depending on δ such that with probability*

$$1 - C_\delta M^{-\delta}$$

the set P_{ϕ_0} contains a subpath $P' = (x_1, x_2, \dots, x_\ell)$ such that

(a) $P' \subseteq V(w_0, \phi_0)$;

(b) $P' \subseteq V_1$ or $P' \subseteq V_2$ or there exists $k \in \{2, \dots, \ell-1\}$ such that $(x_1, \dots, x_k) \subseteq V_1$ and $(x_{k+1}, \dots, x_\ell) \subseteq V_2$;

(c) If $\{x_i, x_{i+1}, x_{i+2}\}$ *be three subsequent vertices in* $P' \cap V_1$ *and* $x_i \in A_{1,j}$ *for some* j*, then* $W_{x_{i+2}} \geq W_{x_i}^{\chi(\zeta\epsilon_1)}$*;*

(d) If $\{x_i, x_{i+1}, x_{i+2}\}$ *be three subsequent vertices in* $P' \cap V_2$ *and* $x_i \in A_{2,j}$ *for some* j*, then* $\phi(x_{i+2}) \geq \phi(x_i)^{1/\chi(\epsilon_1)}$*;*

(e) The length ℓ *of* P' *satisfies*

$$\ell \leq 2f_0(N) + \frac{\log\log_{w_0}\left(\phi(x_1)^{-1}\right)}{\log \chi(\zeta\epsilon_2)} + \frac{\log\log_{\phi_0^{-1}}\left(\phi(x_1)^{-1}\right)}{\log \chi(\epsilon_2)};$$

(f) If $x_\ell \in A_{i,j}$ *for some pair* (i, j)*, then*

$$\mathbb{E}_{x_\ell}\left[\left|\Gamma(x_\ell) \cap V_i^+\left(x_\ell, \epsilon_i^{(j)}\right) \cap V_{>\phi_0}\right|\right] \geq \kappa M^\kappa. \tag{3.19}$$

Proof. The proof of Proposition 3.13 is divided into three steps

(1) Construction of some event E that implies $(a) - (e)$;

(2) Existence of a constant $\kappa > 0$ for (f);

(3) Estimation of $\mathbb{P}(E)$ from below.

We first define some event that satisfies $(a) - (e)$ in the proposition. Denote by $P_{i,j}$ the greedy path in $B_{i,j}$. Let $E_{i,j}$ be the event that satisfies

i) $P_{i,j} \cap A_{i,j} = \emptyset$, or

ii) the first vertex $x \in P_{i,j} \cap A_{i,j}$ satisfies (3.19), or

iii) the first vertex $x \in P_{i,j} \cap A_{i,j}$ has at least one good neighbor. x' is called *a good neighbor* of x if $x' \in \Gamma(x)$ and satisfies

$$x' \in V(w_0, \phi_0) \setminus B_{i,j} \text{ with } \phi(x') \geq \phi(x)$$

and

$$\phi(x') > \phi(y) \text{ for all } y \in \Gamma(x) \cap B_{i,j}.$$

Remark. *The three cases in* $E_{i,j}$ *correspond to three situations in the routing:*

i) The greedy routing path does not go through the layer $A_{i,j}$*;*

ii) The greedy routing path goes through the layer $A_{i,j}$ *and jumps to some vertex with objective larger than* ϕ_0 *after this step;*

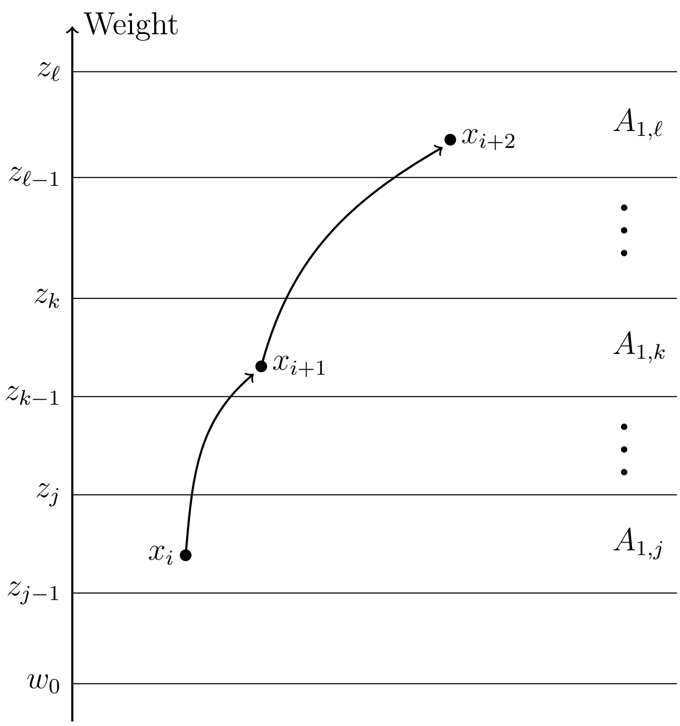

Figure 3.3: Three subsequent vertices in $P' \cap V_1$. Note that they don't lie necessarily in consecutive layers and $j < k < \ell$.

iii) The greedy routing path goes through $A_{i,j}$ and it continues in the main stage.

First we show that the event $E := \bigcap_{i,j} E_{i,j}$ implies $(a) - (e)$. More precisely, we will find a path $P' = (x_i)$ based on the events $(E_{i,j})$. By assumption in the proposition, we have $x_1 \in P_{\phi_0} \cap V(w_0, \phi_0)$ and hence it lies in some layer $A_{i,j}$ because of (3.11). Since the event $E_{i,j}$ holds true, the greedy routing algorithm T finds $x_2 \in A_{i',j'}$ as a good neighbor of x_1 with $A_{i,j}$ before $A_{i',j'}$ in the prescribed order of (3.12). By iterating this procedure we build a path P' with at most one vertex in each layer, until (3.19) is satisfied for some $x_\ell \in P'$.

Clearly, (a) and (b) are true for the path P' due to the construction of layers $(A_{i,j})$ in (3.7), (3.8) and (3.10).

For (c), if $\{x_i, x_{i+1}, x_{i+2}\}$ are three subsequent vertices in $P' \cap V_1$ and $x_i \in A_{1,j}$ for some j, as in Figure 3.3, then we have

$$W_{x_{i+2}} \geq z_{j+1} \geq z_j^{\chi(\zeta\epsilon_1)} \geq W_{x_i}^{\chi(\zeta\epsilon_1)}.$$

For (d) the argument is similar to (c). If $\{x_i, x_{i+1}, x_{i+2}\}$ are three subsequent vertices in $P' \cap V_2$ and $x_i \in A_{2,j}$ for some j, as illustrated in Figure 3.4, then we

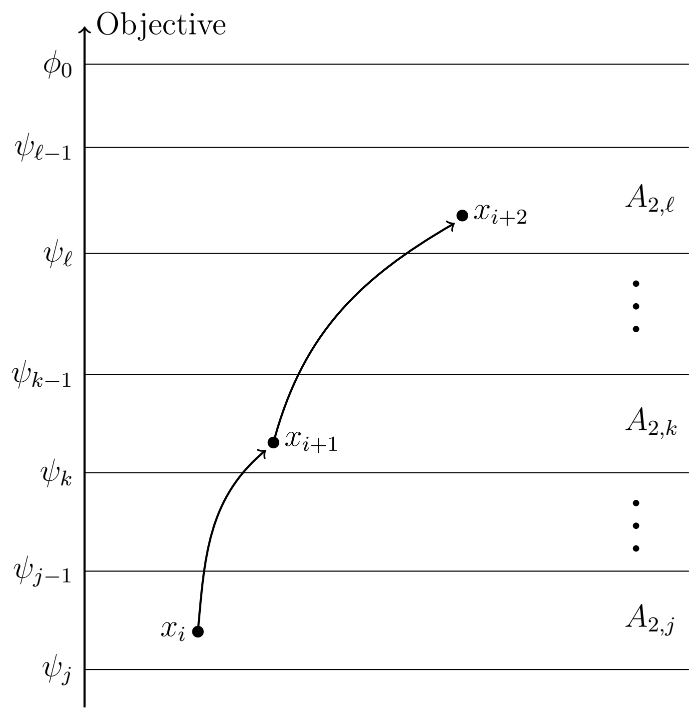

Figure 3.4: Three subsequent vertices in $P' \cap V_2$. Note that they don't lie necessarily in consecutive layers and $j > k > \ell$.

have

$$\phi(x_{i+2}) \geq \psi_{j-2} \geq \psi_{j-1}^{1/\chi(\epsilon_1)} \geq \phi(x_i)^{1/\chi(\epsilon_1)}.$$

For (e) let L_1, L_2 be the number of layers visited by P' in V_1 and V_2 respectively.

In the first phase, as we have seen in (3.6) and (3.7), P' visits possibly two kinds of layers. Let $L_{1,1}$ and $L_{1,2}$ be the number of layers defined with ϵ_1 and ϵ_2 respectively that are visited by P'. By definition of w_0' in (3.5), we obtain $L_{1,1} \leq f_0(N)$. For $L_{1,2}$ we observe that by the routing protocol all the vertices in P' has objetive at least $\phi(x_1)$. Therefore, if $x \in P'$ has a weight larger than $\phi(x_1)^{-1}$, then $x \in V_2$ because

$$W_x^{\chi(\epsilon_1)} \phi(x) \geq \phi(x_1)^{-\chi(\epsilon_1)} \phi(x_1) \geq 1.$$

So it is sufficient to count how many layers are there with weight between w_0 and $\phi(x_1)^{-1}$. This simplifies to solve the following equation

$$w_0^{\left(\chi(\zeta\epsilon_2)^{L_{1,2}}\right)} = \phi(x_1)^{-1}.$$

Therefore we obtain

$$L_{1,2} = \frac{\log \log_{w_0} (\phi(x_1)^{-1})}{\log \chi(\zeta\epsilon_2)}.$$

For the second phase we have a similar argument. Let $L_{2,1}$ and $L_{2,2}$ be the number of layers visited by P' that are defined with ϵ_1 and ϵ_2 respectively in the second phase. Similarly we have $L_{2,1} = f_0(N)$. For $L_{2,2}$ it is sufficient to consider the number of layers of objective between $\phi(x_1)$ and ϕ_0. That is to solve the equation

$$\phi_0^{\left(\chi(\epsilon_2)^{L_{2,2}}\right)} = \phi(x_1).$$

Consequently we get

$$L_{2,2} = \frac{\log \log_{\phi_0^{-1}} (\phi(x_1)^{-1})}{\log \chi(\epsilon_2)}.$$

Therefore we obtain

$$\begin{aligned}\ell = L_1 + L_2 &= L_{1,1} + L_{1,2} + L_{2,1} + L_{2,2} \\ &\leq 2f_0(N) + \frac{\log \log_{w_0} (\phi(x_1)^{-1})}{\log \chi(\zeta\epsilon_2)} + \frac{\log \log_{\phi_0^{-1}} (\phi(x_1)^{-1})}{\log \chi(\epsilon_2)},\end{aligned}$$

and this finishes the first step.

Now we try to find a constant κ satisfying (3.19). More precisely, given $c > 0$ as a constant, we find a constant $\kappa > 0$ such that the following holds: there exists $M > 0$, and for all $N \geq M$ one has

$$cW_x^{\epsilon_i^{(j)}} > 2\kappa M^\kappa, \quad \text{and } c\phi(x)^{-c\epsilon_i^{(j)}} > 2\kappa M^\kappa, \tag{3.20}$$

for all $x \in A_{i,j}$, $i \in 1, 2$ and $j \geq 1$.

In order to find such κ we need to consider three cases:

- $i = 1, j < \infty$. Let j_0 be the number such that if $j \leq j_0$, ϵ_1 is used for $A_{1,j}$, and otherwise ϵ_2 is used.

  If $x \in A_{1,j}$ and hence $x \in V_1'$ for some $j \leq j_0$, then $W_x \geq w_0$ and $\phi(x) \leq$

$W_x^{-\chi(\epsilon_2)}$. We can choose κ so small that

$$\begin{cases} cW_x^{\epsilon_1} > 2\kappa w_0^\kappa \geq 2\kappa M^\kappa, \\ c\phi(x)^{-c\epsilon_1} \geq cW_x^{c\epsilon_1\chi(\epsilon_2)} \geq cW_x^{c\epsilon_1} > 2\kappa M^\kappa. \end{cases}$$

If $x \in A_{1,j}$ for some $j > j_0$, then $W_x \geq w_0'$ and $\phi(x) \leq W_x^{-\chi(\epsilon_2)}$. Since $\lim_{N\to\infty} f_0(N) = \infty$, there exists $M_1 > 0$ such that for all $N \geq M_1$ one has

$$\frac{\chi(\zeta\epsilon_1)^{f_0(N)}}{\log\log f_0(N)} \geq c,$$

where c is the constant given in (3.20). In this case, we can choose κ with $\kappa < c^2 \wedge \frac{1}{2}c$. Then

$$\begin{cases} cW_x^{\epsilon_2} > cw_0'^{\epsilon_2} = cw_0^{\frac{\chi(\zeta\epsilon_1)^{f_0(N)}}{\log\log f_0(N)}} > 2\kappa w_0^\kappa \geq 2\kappa M^\kappa, \\ c\phi(x)^{-c\epsilon_2} \geq cW_x^{c\epsilon_2} \geq cw_0^{c^2} > 2\kappa M^\kappa. \end{cases}$$

- $i = 1, j = \infty$. Since $x \in A_{1,\infty}$, one has $x \in V_1 \cap V(w_0, \phi_0)$ and therefore $W_x \geq w_0$ and $\phi(x) \leq \phi_0$. For $c > 0$ we can choose such a small κ that

$$\begin{cases} cW_x^{\epsilon_1} > 2\kappa w_0^\kappa \geq 2\kappa M^\kappa \\ c\phi(x)^{-c\epsilon_1} > 2\kappa\phi_0^{-\kappa} \geq 2\kappa M^\kappa. \end{cases}$$

- $i = 2, j < \infty$. Let j_1 be the number such that for $j \leq j_1$, ϵ_1 is used for the layer $A_{2,j}$, and else ϵ_2 is used.

  If $x \in A_{2,j}$ for some $j \leq j_1$, then $\phi(x) \leq \phi_0$ and $W_x^{-\chi(\epsilon_1)} \leq \phi(x)$. We can choose κ so small that

$$\begin{cases} c\phi(x)^{-c\epsilon_1} > 2\kappa\phi_0^{-\kappa} \geq 2\kappa M^\kappa, \\ cW_x^{\epsilon_1} \geq c\phi(x)^{-\frac{\epsilon_1}{\chi(\epsilon_1)}} > 2\kappa\phi_0^{-\kappa} \geq 2\kappa M^\kappa. \end{cases}$$

  If $x \in A_{2,j}$ for some $j > j_1$, then $\phi(x) \leq \phi_0'$ and $W_x^{-\chi(\epsilon_1)} \geq \phi(x)$. Since $\lim_{N\to\infty} f_0(N) = \infty$, there exists $M_2 > 0$ such that for $N > M_2$ one has

$$\frac{\chi(\epsilon_1)^{f_0(N)}}{\log\log f_0(N)} > 1.$$

In this case, we can choose κ so small that

$$\begin{cases} c\phi(x)^{-c\epsilon_2} \geq c{\phi'_0}^{-c\epsilon_2} = c\phi_0^{-\frac{c\chi(\epsilon_1)^{f_0(N)}}{\log\log f_0(N)}} > 2\kappa\phi_0^{-\kappa} \geq 2\kappa M^{-\kappa}, \\ cW_x^{\epsilon_2} \geq c\phi(x)^{-\frac{\epsilon_2}{\chi(\epsilon_1)}} \geq c{\phi_0}^{\prime -\frac{\epsilon_2}{\chi(\epsilon_1)}} > 2\kappa\phi_0^{-\kappa} \geq 2\kappa M^{-\kappa}. \end{cases}$$

By choosing the minimum of all κ's in all three phases above we obtain the desired constant κ, and hence finish the second step.

As we have see, E is indeed the good event. Now we need to estimate the probability of E from below. Since we have

$$\mathbb{P}(E) \geq 1 - \sum_{j\geq 1} \mathbb{P}(E_{1,j}^c) - \mathbb{P}(E_{1,\infty}^c) - \sum_{j\geq 1} \mathbb{P}(E_{2,j}^c),$$

and

$$\begin{aligned} \mathbb{P}(E_{i,j}) =& \mathbb{P}\left(E_{i,j} | P_{i,j} \cap A_{i,j} = \emptyset\right) \mathbb{P}\left(P_{i,j} \cap A_{i,j} = \emptyset\right) \\ &+ \mathbb{P}\left(E_{i,j} | P_{i,j} \cap A_{i,j} \neq \emptyset\right) \mathbb{P}\left(P_{i,j} \cap A_{i,j} \neq \emptyset\right) \\ =& \mathbb{P}\left(P_{i,j} \cap A_{i,j} = \emptyset\right) + \mathbb{P}\left(E_{i,j} | P_{i,j} \cap A_{i,j} \neq \emptyset\right) \mathbb{P}\left(P_{i,j} \cap A_{i,j} \neq \emptyset\right) \\ \geq& \mathbb{P}\left(E_{i,j} | P_{i,j} \cap A_{i,j} \neq \emptyset\right), \end{aligned}$$

it is sufficient to give a lower bound for the conditional probability in the last line.

Depending on the phases we have the following estimates for $\mathbb{P}(E_{1,j})$:

- $i = 1, j < \infty$. Since we condition on $P_{i,j} \cap A_{1,j} \neq \emptyset$, let y be the first vertex in $P_{i,j} \cap A_{1,j}$. Assume (3.19) does not hold for y. By our choice of κ we have

$$\mathbb{E}_y\left[\left|\Gamma(y) \cap V_1^+\left(y, \epsilon_1^{(j)}\right) \cap V_{>\phi_0}\right|\right] < \kappa M^\kappa < \frac{1}{2} c W_y^{\epsilon_1^{(j)}}.$$

By Lemma 3.10 we know

$$\mathbb{E}_y\left[\left|\Gamma(y) \cap V_1^+\left(y, \epsilon_1^{(j)}\right)\right|\right] \geq c W_y^{\epsilon_1^{(j)}}.$$

As a consequence we obtain

$$\mathbb{E}_y\left[\left|\Gamma(y) \cap V_1^+\left(y, \epsilon_1^{(j)}\right) \cap V_{\leq\phi_0}\right|\right] \geq \frac{1}{2} c W_y^{\epsilon_1^{(j)}}.$$

Since $y \in A_{1,j}$, we know $W_y \geq z_{j-1}$. By Chernoff's inequality in Lemma

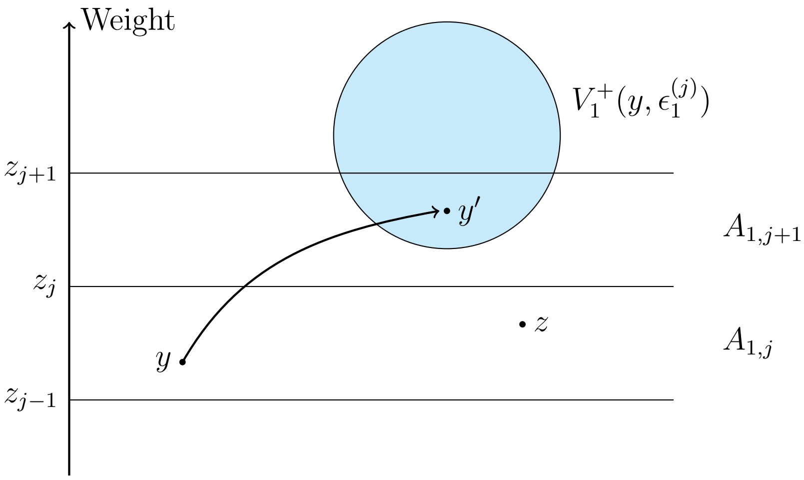

Figure 3.5: Jump in the first phase.

3.1, with probability at least $1 - \exp\left(-1/2cz_{j-1}^{\epsilon_1^{(j)}}\right)$, y has a neighbor $y' \in V_1^+(y, \epsilon_1^{(j)}) \cap V_{<\phi_0}$, as in Figure 3.5. If there exist several such neighbors, we choose the one with highest objective to make sure that it is the target of the greedy algorithm T. Now we show that y' is a good neighbor of y_l.

First, since $y' \in V_1^+(y, \epsilon_1^{(j)})$, we know $W_{y'} \geq W_y^{\chi(\epsilon_1^{(j)})} \geq z_{j-1}^{\chi(\epsilon_1^{(j)})} \geq z_j$. Therefore $y' \in V(w_0, \phi_0) \setminus B_{1,j}$. Besides, $\phi(y') \geq \phi(y) W_y^{\chi(\epsilon_1^{(j)})-1}$. In order to show y' is a good neighbor, we still need to show that $\phi(y) W_y^{\chi(\epsilon_1^{(j)})-1} > \phi(z)$ for all $z \in \Gamma(y) \cap B_{1,j}$.

Note if $\phi(z) \geq \phi(y) W_y^{\chi(\epsilon_1^{(j)})-1}$ for some $z \in B_{1,j}$, then $z \in V_1^-\left(y, \epsilon_1^{(j)}\right)$, because $W_z \leq z_j = z_{j-1}^{\chi(\zeta\epsilon_1^{(j)})} \leq W_y^{\chi(\zeta\epsilon_1^{(j)})}$. Therefore it is sufficient to estimate the size of $V_1^-\left(y, \epsilon_1^{(j)}\right)$. By Lemma 3.10 we have

$$\mathbb{E}_y\left[\left|\Gamma(y) \cap V_1^-\left(y, \epsilon_1^{(j)}\right)\right|\right] \leq C W_y^{-\rho\epsilon_1^{(j)}} \log W_y \leq C z_{j-1}^{-\rho\epsilon_1^{(j)}} \log z_j.$$

Note that the right hand side is independent of W_y. Therefore, we can replace $\mathbb{E}_y$ by $\mathbb{E}$ on the left hand side. By Markov's inequality, with probability at least $1 - C z_{j-1}^{-\rho\epsilon_1^{(j)}} \log z_j$, y has no neighbor in $V_1^-\left(y, \epsilon_1^{(j)}\right)$. In this case y' is a good neighbor of y. As a consequence

$$\begin{aligned}\mathbb{P}(E_{1,j}) &\geq \mathbb{P}(E_{1,j}|P_{1,j}\cap A_{1,j}\neq\emptyset)\\ &\geq \left(1-\exp\left(-1/2cz_{j-1}^{\epsilon_1^{(j)}}\right)\right)\left(1-Cz_{j-1}^{-\rho\epsilon_1^{(j)}}\log z_j\right)\\ &\geq 1-\exp\left(-1/2cz_{j-1}^{\epsilon_1^{(j)}}\right)-Cz_{j-1}^{-\rho\epsilon_1^{(j)}}\log z_j\\ &\geq 1-C'z_{j-1}^{-\rho\epsilon_1^{(j)}}\log z_j.\end{aligned}$$

- $i=1, j=\infty$. Since $y\in A_{1,\infty}$ satisfies $\phi(y)W_y^{\chi(\epsilon_1)}\leq 1\leq \phi(y)W_y^{\chi(\epsilon_2)}$, for $x\in V_1^+(y,\epsilon_1)$ one has $\phi(x)\geq\phi(y)W_y^{\chi(\epsilon_1)-1}$ and $W_x\geq W_y^{\chi(\epsilon_1)}$. Then

$$\phi(x)\geq\phi(y)W_y^{\chi(\epsilon_1)-1}\geq W_y^{-\chi(\epsilon_2)}W_y^{\chi(\epsilon_1)-1}\geq W_x^{\frac{\chi(\epsilon_1)-1-\chi(\epsilon_2)}{\chi(\epsilon_1)}}\geq W_x^{-\chi(\epsilon_1)}.$$

The last step above is true due to the choice of ϵ_1 in (3.3) and the fact that $\zeta>1$. Therefore $x\in V_2$ and $x\notin B_{1,\infty}$. This means once the algorithm T visits the layer of transition $A_{1,\infty}$, it is very likely to jump to a vertex in V_2 and consequently enter the second phase of greedy routing. Recall that $B_{i,j}$ is the union from $A_{1,1}$ to $A_{i,j}$ according to the order in Remark 3.9.

Now let y be the first vertex in $P_{1,\infty}\cap A_{1,\infty}$, and as before assume (3.19) does not hold for y. Then by the choice of κ we have

$$\mathbb{E}_y\left[\left|\Gamma(y)\cap V_1^+(y,\epsilon_1)\cap V_{>\phi_0}\right|\right]<\kappa M^\kappa<\frac{1}{2}cW_y^{\epsilon_1}.$$

Lemma 3.10 implies

$$\mathbb{E}_y\left[\left|\Gamma(y)\cap V_1^+(y,\epsilon_1)\right|\right]\geq cW_y^{\epsilon_1}.$$

Therefore

$$\mathbb{E}_y\left[\left|\Gamma(y)\cap V_1^+(y,\epsilon_1)\cap V_{<\phi_0}\right|\right]>1/2cW_y^{\epsilon_1}.$$

By Chernoff's bound in Lemma 3.1, with at least probability $1-\exp\left(-1/2cw_0^{\epsilon_1}\right)$, y has at least one neighbor y' in $V_1^+(y,\epsilon_1)\cap V_{<\phi_0}$. Consequently $y'\notin B_{1,\infty}$ and $y'\in V_2$.

We show now y' is a good neighbor of y. That is, we need to show

1. $\phi(y')\geq\phi(y)$;

2. $\phi(y') \geq \phi(x)$ for all $x \in \Gamma(y) \cap B_{1,\infty}$.

Since $y' \in V_1^+(y, \epsilon_1)$, one has $\phi(y') \geq \phi(y) W_y^{\chi(\epsilon_1)-1} \geq \phi(y)$. Therefore the first condition is satisfied.

For the second condition, let $z \in B_{1,\infty}$ with $\phi(z) \geq \phi(y) W_y^{\chi(\epsilon_1)-1}$. Then

$$W_z^{\chi(\epsilon_1)} \leq \phi(z)^{-1} \leq \phi(y)^{-1} W_y^{-\chi(\epsilon_1)+1} \leq W_y^{\chi(\epsilon_2)-\chi(\epsilon_1)+1} \leq W_y^{\chi(\epsilon_1)\chi(\zeta\epsilon_1)},$$

where we used (3.3) in the last estimation. As a result we obtain $z \in V_1^-(y, \epsilon_1)$. By Lemma 3.10 we have for arbitrary $\delta' > 0$

$$\mathbb{E}_y \left[|\Gamma(y) \cap V_1^-(y, \epsilon_1)| \right] \leq C W_y^{-\rho\epsilon_1} \log W_y \leq C_{\delta'} W_y^{-\rho\epsilon_1+\delta'} \leq C_{\delta'} w_0^{-\rho\epsilon_1+\delta'},$$

for some positive constant $C_{\delta'} := C_\delta(C, \delta')$.

Hence with probability at least $1 - C_{\delta'} w_0^{-\rho\epsilon_1+\delta'}$, y has no neighbor in $V_1^-(y, \epsilon_1)$. In this case all $x \in \Gamma(y) \cap B_{1,\infty}$ satisfies

$$\phi(x) < \phi(y_l) W_y^{\chi(\epsilon_1)-1} \leq \phi(y'),$$

which means y' is a good neighbor of y.

To summarize, we have

$$\begin{aligned} \mathbb{P}(E_{1,\infty}) &\geq (1 - \exp(-1/2cw_0^{\epsilon_1})) \left(1 - C_{\delta'} w_0^{-\rho\epsilon_1+\delta'}\right) \\ &\geq 1 - \exp(-1/2cw_0^{\epsilon_1}) - C_{\delta'} w_0^{-\rho\epsilon_1+\delta'} \\ &\geq 1 - C'_{\delta'} w_0^{-\rho\epsilon_1+\delta'}. \end{aligned}$$

- $i = 2, j < \infty$. Let y be the first vertex in $P_{2,j} \cap A_{2,j}$. Assume (3.19) does not hold for y, that is

$$\mathbb{E}_y \left[\left| \Gamma(y) \cap V_2^+(y, \epsilon_2^{(j)}) \cap V_{>\phi_0} \right| \right] < \kappa M^\kappa < \frac{c}{2} \phi(y)^{-c\epsilon_2^{(j)}}.$$

By Lemma 3.12 one knows

$$\mathbb{E}_y \left[\left| \Gamma(y) \cap V_2^+(y, \epsilon_2^{(j)}) \right| \right] \geq c\phi(y)^{-c\epsilon_2^{(j)}}.$$

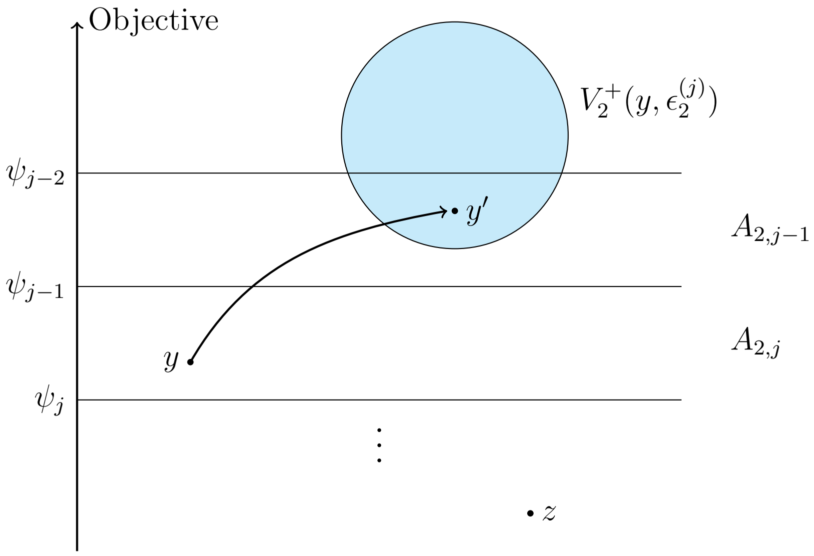

Figure 3.6: Jump in the second phase.

It follows that

$$\mathbb{E}_y\left[\left|\Gamma(y)\cap V_2^+(y,\epsilon_2^{(j)})\cap V_{<\phi_0}\right|\right] > \frac{c}{2}\phi(y)^{-c\epsilon_2^{(j)}}.$$

By Chernoff's bound in Lemma 3.1, together with the fact that $\psi_{j-1}\geq\phi(y)>\psi_j$, with probability at least $1-\exp\left(-c/2\psi_{j-1}^{-c\epsilon_2^{(j)}}\right)$, there will be at least one vertex in the set $\Gamma(y)\cap V_2^+(y,\epsilon_2^{(j)})\cap V_{<\phi_0}$. Let y' be the vertex in this set with highest objective, as showed in Figure 3.6. Because $y'\in V_2^+(y,\epsilon_2^{(j)})$, we have

$$\phi(y')\geq\phi(y)^{1/\chi(\epsilon_2^{(j)})}>\psi_j^{1/\chi(\epsilon_2^{(j)})}=\psi_{j-1}.$$

Therefore $y'\notin B_{2,j}$. Next we show y' is a good neighbor of y.

To be a good neighbor of y, y' must satisfy the following conditions:

1. $\phi(y')\geq\phi(y)$;
2. $\phi(y')\geq\phi(z)$ for all $z\in\Gamma(y)\cap B_{2,j}$.

The first condition is clearly satisfied, because $\phi(y')>\psi_{j-1}\geq\phi(y)$. For the second condition, we will show that all the vertices in $B_{2,j}$ have objective less than ψ_{j-1}. Consider $z\in B_{2,j}$ with $\phi(z)\geq\phi(y)^{1/\chi(\epsilon_2^{(j)})}\geq\psi_{j-1}$, then $z\in B_{1,\infty}$. That is, $y\in V_1$ and consequently $y\in V_2^-(y,\epsilon_2^{(j)})$. By Lemma 3.12 we know

$$\mathbb{E}_y\left[\left|\Gamma(y)\cap V_2^-(y,\epsilon_2^{(j)})\right|\right]\leq C\phi(y)^{\epsilon_2^{(j)}}\log\left(\phi(y)^{-1}\right)\leq C\psi_{j-1}^{\epsilon_2^{(j)}}\log\psi_j^{-1}.$$

Therefore, by Markov inequality, with probability at least $1 - C\psi_{j-1}^{\epsilon_2^{(j)}} \log \psi_j^{-1}$ such vertex z above does not exist. In this case, y' is a good neighbor of y. Then

$$\begin{aligned}\mathbb{P}\left(E_{2,j}\right) &\geq \left(1 - \exp\left(-c/2\psi_{j-1}^{-c\epsilon_2^{(j)}}\right)\right)\left(1 - C\psi_{j-1}^{\epsilon_2^{(j)}} \log \psi_j^{-1}\right)\\ &\geq 1 - C'\psi_{j-1}^{\epsilon_2^{(j)}} \log \psi_j^{-1}.\end{aligned}$$

With all the preparations we can estimate the probability of E:

$$\begin{aligned}\mathbb{P}(E) &\geq 1 - \sum_{j=1}^{\infty} \mathbb{P}(E_{1,j}^c) - \mathbb{P}(E_{1,\infty}^c) - \sum_{j=1}^{\infty} \mathbb{P}(E_{2,j}^c)\\ &\geq 1 - \sum_{j=1}^{\infty} C' z_{j-1}^{-\rho\epsilon_1^{(j)}} \log z_j - C'_{\delta'} w_0^{-\rho\epsilon_1+\delta'} - \sum_{j=1}^{\infty} C'\psi_{j-1}^{\epsilon_2^{(j)}} \log \psi_j^{-1}. \end{aligned} \tag{3.21}$$

For the first sum in (3.21), we consider two cases, depending on the value of $\epsilon_1^{(j)}$. Let $j_0 = \min\{j \in \mathbb{N} : z_j > w'_0\}$. That is, j_0 is the smallest index such that z_{j_0+1} is defined with ϵ_2 in (3.6). Then

$$\sum_{j=1}^{\infty} z_{j-1}^{-\rho\epsilon_1^{(j)}} \log z_j = \sum_{j=1}^{j_0} z_{j-1}^{-\rho\epsilon_1} \log z_j + \sum_{j=j_0+1}^{\infty} z_{j-1}^{-\rho\epsilon_2} \log z_j. \tag{3.22}$$

For the first part of (3.22) we have the following estimation:

$$\begin{aligned}\sum_{j=1}^{j_0} z_{j-1}^{-\rho\epsilon_1} \log z_j &= \sum_{j=1}^{j_0} \left(w_0^{\chi(\zeta\epsilon_1)^{j-1}}\right)^{-\rho\epsilon_1} \chi(\zeta\epsilon_1)^j \log w_0\\ &\leq \log w_0 \left(\gamma(\zeta\epsilon_1) w_0^{-\rho\epsilon_1} + \sum_{j=1}^{j_0-1} \left(\gamma(\zeta\epsilon_1) w_0^{-\rho\epsilon_1\chi(\zeta\epsilon_1)}\right)^j\right)\\ &\leq \left(\gamma(\zeta\epsilon_1) w_0^{-\rho\epsilon_1} + 2w_0^{-\rho\epsilon_1\chi(\zeta\epsilon_1)}\right) \log w_0\\ &\leq 2\gamma(\zeta\epsilon_1) w_0^{-\rho\epsilon_1} \log w_0.\end{aligned}$$

For the second part of (3.22) we apply a similar argument:

$$\sum_{j=j_0+1}^{\infty} z_{j-1}^{-\rho\epsilon_2} \log z_j \leq \sum_{j=0}^{\infty} \left((w_0')^{\chi(\zeta\epsilon_2)^j} \right)^{-\rho\epsilon_2} \chi(\zeta\epsilon_2)^2 \log w_0'$$
$$\leq \chi(\zeta\epsilon_2)^2 \log w_0' \left((w_0')^{-\rho\epsilon_2} + \sum_{j=1}^{\infty} \left((w_0')^{-\rho\epsilon_2\chi(\zeta\epsilon_2)} \right)^j \right).$$

Note the base in the geometric series above satisfies

$$(w_0')^{-\rho\epsilon_2\chi(\zeta\epsilon_2)} = \left(w_0^{\chi(\zeta\epsilon_2)\frac{\chi(\zeta\epsilon_1) f_0(N)}{\log\log f_0(N)}} \right)^{-\rho} \to 0, \quad \text{as } N \to \infty.$$

Therefore for N large one has

$$\sum_{j=1}^{\infty} \left((w_0')^{-\rho\epsilon_2\chi(\zeta\epsilon_2)} \right)^j \leq 2\,(w_0')^{-\rho\epsilon_2\chi(\zeta\epsilon_2)} \leq 2\,(w_0')^{-\rho\epsilon_2}.$$

As a result, we obtain for N large,

$$\sum_{j=1}^{\infty} z_{j-1}^{-\rho\epsilon_1^{(j)}} \log z_j \leq 2\chi(\zeta\epsilon_1) w_0^{-\rho\epsilon_1} \log w_0 + 3\chi(\zeta\epsilon_2)^2 (w_0')^{-\rho\epsilon_2} \log w_0'$$
$$\leq C_\delta w_0^{-\delta} \leq C_\delta M^{-\delta},$$

where δ is a constant in $(0, \rho\epsilon_1)$ and $M = w_0 \wedge \phi_0^{-1}$. Note that the second step holds because $(w_0')^{\epsilon_2} \gg w_0$.

For the second sum in (3.21) we use exactly the same method and obtain

$$\sum_{j=1}^{\infty} C' \psi_{j-1}^{\epsilon_2^{(j)}} \log \psi_j^{-1} \leq C_\delta \phi_0^\delta \leq C_\delta M^{-\delta}.$$

This finishes the proof of Proposition 3.13. □

Proposition 3.14 (End stage)**.** *Assume now T arrives at some vertex x_ℓ that satisfies the condition (3.19). Then there exists a positive constant $\mu \in (0,1]$ such that the greedy algorithm T starting from x_ℓ ends in the target t within 2 steps with probability at least μ.*

Proof. If $t \in \Gamma(x_\ell)$, then we are done. Otherwise consider the set $\Gamma(x_\ell) \cap V_{>\phi_0}$. Since the vertex x_ℓ satisfies (3.19), we know

$$\mathbb{E}\left[|\Gamma(x_\ell) \cap V_{>\phi_0}|\right] \geq \kappa M^\kappa.$$

By the Chernoff bound in Lemma 3.1, $\Gamma(x_\ell) \cap V_{>\phi_0}$ is non-empty with at least probability $\nu > 0$. In this case, let $x_{\ell+1}$ be the vertex in $\Gamma(x_\ell) \cap V_{>\phi_0}$ with highest objective. Then $x_{\ell+1}$ is linked to t with probability

$$\mathbb{P}(x_{\ell+1} \sim t) = \mathbb{E}\left[1 \wedge \frac{W_{x_{\ell+1}} W_t}{|x_{\ell+1} - t|^\alpha}\right] = \mathbb{E}\left[1 \wedge \phi(x_{\ell+1}) W_t\right] \geq \phi_0.$$

If $x_{\ell+1}$ is connected with t, since t is the global maximizer of the objective function, T jumps from $x_{\ell+1}$ to t. We choose $\mu := \nu\phi_0$ as the desired the constant. Note in the proof we apply Proposition 3.13 with w_0, ϕ_0 as constants. This finishes the proof. □

Proof of Theorem 3.7. With Proposition 3.8, 3.13 and 3.14 we come to the conclusion that with at least constant probability the greedy routing algorithm T starting from s will find the target t successfully within L steps, and L satisfies

$$\begin{aligned} L &\leq 1 + 2f_0(N) + \frac{\log \log_{w_0}\left(\phi(x_1)^{-1}\right)}{\log \chi(\zeta\epsilon_2)} + \frac{\log \log_{\phi_0^{-1}}\left(\phi(x_1)^{-1}\right)}{\log \chi(\epsilon_2)} + 2 \\ &\leq \frac{1 + o(1)}{|\log(\gamma - 1)|}\Bigg(\log \log_{W_s}\left(\phi(s)^{-1}\right) + \log \log_{W_t}\left(\phi(s)^{-1}\right)\Bigg). \end{aligned}$$

In this case, we choose w_0 and ϕ_0 as constants for Proposition 3.13. □

3.2.2 Length of greedy routing paths

In Section 3.2.1 we see that the greedy algorithm starting from s finds the target t within $\mathcal{O}(\log \log |s - t|)$ steps with at least constant positive probability. The following result tells us that no matter whether the algorithm finds the target or not, it terminates after at most $\mathcal{O}(\log \log |s - t|)$ steps.

Theorem 3.15 (Part (b) in Theorem 1.6)**.** *Consider scale-free percolation with connection probability* $p_{x,y} = \frac{W_x W_y}{|x-y|^\alpha} \wedge 1$, *and parameters* $\alpha > d, \gamma \in (1, 2)$. *Let* T *be the greedy routing algorithm with objective function* $\phi(x) = \frac{W_x}{|x-t|^\alpha}$ *as in* (1.10)*. Then, conditional on* W_s *and* W_t, *with high probability,* T *terminates within* L_2 *steps* *as $N \to \infty$, where L_2 is a function of N given as follows:*

$$L_2 = \frac{1 + o(1)}{|\log(\gamma - 1)|} \Bigg(\log\log_{W_s} \left(\phi(s)^{-1}\right) + \log\log_{W_t} \left(\phi(s)^{-1}\right) \Bigg).$$

In the previous discussions we have seen that the greedy routing algorithm T will either find the target t or be trapped in some local optimum. In this section we will show that with high probability in both cases the length of greedy path will be at most doubly logarithmic in the Euclidean distance, as stated in Theorem 3.15.

Proposition 3.16 (Unlikely jumps). *Let $c > 1$ be a constant and $x \in \mathbb{Z}^d$, and $w_0, \epsilon > 0$. Then*

(a) There exists some constant $c_1 > 0$ such that with probability $1 - c_1 w_0^{c_1(1+\epsilon-\gamma)}$ there exists no vertex y (except possibly t) such that $W_y \geq w_0$ and $\phi(y) \geq w_0^{-\epsilon}$;

If $W_x \leq w_0$, then for arbitrary $\delta > 0$, there exists $C_\delta := C_\delta(\delta)$ such that

(b) With probability $1 - C_\delta w_0^{1+\delta-c(\alpha-d)/d}$, x has no such neighbor y that $W_y \leq w_0$ and $|x - y|^d \geq w_0^c$;

Proof. The main idea for the proof of the assertions in the proposition is to first estimate the expected size of corresponding sets and then to obtain the bounds for the probabilities by Markov's inequality.

(a) Let $A_1 := \{x \in \mathbb{Z}^d : W_x \geq w_0, \phi(x) \geq w_0^{-\epsilon}\}$. We now estimate the expected size of A_1:

$$\begin{aligned}
\mathbb{E}[|A_1|] &= \sum_{x \in \mathbb{Z}^d} \mathbb{P}\left(W_x \geq w_0, \phi(x) \geq w_0^{-\epsilon}\right) \\
&= \sum_{x \in \mathbb{Z}^d} \mathbb{P}\left(W_x \geq w_0, W_x \geq w_0^{-\epsilon}|x - t|^\alpha\right) \\
&= \sum_{x \in S} \mathbb{P}\left(W_x \geq w_0\right) + \sum_{x \in S^c} \mathbb{P}\left(W_x \geq w_0^{-\epsilon}|x - t|^\alpha\right),
\end{aligned}$$

where $S := \{x \in \mathbb{Z}^d : w_0^{1+\epsilon} \geq |x - t|^\alpha\}$. For the first sum it is sufficient to estimate the size of S, and we know $|S| \leq c w_0^{\frac{(1+\epsilon)d}{\alpha}}$ for some constant $c := c(d)$. Therefore

$$\sum_{x \in S} \mathbb{P}\left(W_x \geq w_0\right) \leq c w_0^{\frac{(1+\epsilon)d}{\alpha}} w_0^{-\tau+1}.$$

For the second sum we apply Lemma 3.2, and obtain

$$\sum_{x\in S^c} \mathbb{P}\left(W_x \geq w_0^{-\epsilon}|x-t|^{\alpha}\right) = \sum_{x\in\mathbb{Z}^d:|x-t|\geq w_0^{\frac{1+\epsilon}{\alpha}}} \frac{w_0^{\epsilon(\tau-1)}}{|x-t|^{\alpha(\tau-1)}}$$
$$\leq \frac{Cw_0^{\epsilon(\tau-1)}}{w_0^{\frac{1+\epsilon}{\alpha}(\alpha(\tau-1)-d)}} = Cw_0^{\frac{(1+\epsilon)d}{\alpha}} w_0^{-\tau+1}.$$

Here we used the fact that $\alpha(\tau-1) > d$ because $\gamma = \frac{\alpha(\tau-1)}{d} > 1$. Hence we have $\mathbb{E}[|A_1|] \leq c' w_0^{\frac{d}{\alpha}(1+\epsilon-\gamma)}$. By Markov's inequality the result follows.

(b) Let $A_2 := \{y \in \mathbb{Z}^d : x \sim y, W_y \leq w_0, |x-y|^d \geq w_0^c\}$. Then

$$\mathbb{E}[|A_2|] = \sum_{y\in\mathbb{Z}^d:|x-y|^d\geq w_0^c} \mathbb{P}(y\sim x, W_y \leq w_0)$$
$$= \sum_{y\in\mathbb{Z}^d:|x-y|^d\geq w_0^c} \int_1^{w_0} u^{-\tau}\left(\frac{W_x u}{|x-y|^{\alpha}} \wedge 1\right) du$$
$$\leq \sum_{y\in\mathbb{Z}^d:|x-y|^d\geq w_0^c} \frac{w_0}{|x-y|^{\alpha}} \int_1^{w_0} u^{-\tau+1} du$$
$$\leq \frac{Cw_0}{w_0^{c(\alpha-d)/d}} \int_1^{w_0} u^{-\tau+1} du$$

where we applied Lemma 3.2 in the last step.

- If $\tau \in (1,2)$, then $\int_1^{w_0} u^{-\tau+1} du \leq \frac{w_0^{2-\tau}}{2-\tau}$, and therefore
$$\mathbb{E}[|A_2|] \leq C_1 w_0^{3-\tau-c(\alpha-d)/d}.$$
- If $\tau = 2$, we have $\int_1^{w_0} u^{-\tau+1} du \leq \log w_0$. And
$$\mathbb{E}[|A_2|] \leq Cw_0^{3-\tau-c(\alpha-d)/d} \log w_0 \leq C_\delta w_0^{3-\tau+\delta-c(\alpha-d)/d},$$
for any $\delta > 0$.
- If $\tau \in (2,3)$, we have $\int_1^{w_0} u^{-\tau+1} du \leq \frac{1}{\tau-2}$. Then
$$\mathbb{E}[|A_2|] \leq C_1 w_0^{1-c(\alpha-d)/d}.$$

To summarize, for $\delta > 0$ there exists a constant $C_\delta := C_\delta(\delta) > 0$ such that $\mathbb{E}[|A_2|] \leq C_\delta w_0^{1+\delta-c(\alpha-d)/d}$.

□

Definition 3.17 (w-grid)**.** *A w-grid is a partition of $\mathbb{Z}^d$ into hypercubes with side length $w^{1/d}$. The hypercubes in the w-grid are called cells.*

By definition each cell in a w-grid has volume w.

Analogous to the proof in Section 3.2.1, the greedy path can also be divided into three parts.

1. **Start stage**: Starting from s, with high probability T aborts or finds some vertex x_1 with weight at least w_0 in several steps;
2. **Main stage**: From x_1 with high probability T arrives at some vertex x_ℓ with given objective ϕ_0 within doubly logarithmic number of steps;
3. **End stage**: From x_ℓ with high probability T terminates in a few steps.

For the main stage we just apply Proposition 3.13 with w_0, ϕ_0 satisfying (3.18). The following proposition deals with the start stage and ensures that with high probability greedy algorithm reaches some vertex with weight at least w_0 or just stops within several steps.

Proposition 3.18 (Start stage)**.** *Let $w_0 = w_0(N)$ be a function of N such that $\lim_{N\to\infty} w_0(N) = \infty$. Further we assume the starting vertex s satisfies $W_s \leq w_0$ and $\phi(s) \leq e^{-w_0}$. Then, with high probability, the greedy routing algorithm terminates within $w_0^{6d/\alpha}$ steps, or after visiting $w_0^{6d/\alpha}$ different vertices, it reaches some vertex with weight at least w_0.*

Proof. Denote by G the graph generated by scale-free percolation in $\mathbb{Z}^d$. Let G' be the subgraph of G induced by all the vertices with weights less than w_0.

First we let the greedy routing run on G'. Consider the first hop from $x_0 := s$. By choosing c properly in Proposition 3.16 (b) we know with probability at least $1-w_0^{-1-6d/\alpha}$ all neighbors of s have a distance at most w_0^c to s. In this case, the greedy routing will visit the next vertex x_1 in the path with $|x_1 - s| \leq w_0^c$. By repeating the step for $w_0^{6d/\alpha}$ times, together with the subadditivity of probability, we obtain with probability at least $1 - w_0^{-1}$ a greedy routing path $P' := \{x_0, x_1, \ldots, x_{w_0^{6d/\alpha}}\}$ that satisfies $|x_i - x_{i-1}| \leq w_0^c$. Therefore $|s - x_i| \leq w_0^{c+6d/\alpha}$ for all $i \in [w_0^{6d/\alpha}]$.

Then we let the greedy routing run on G for $w_0^{6d/\alpha}$ steps, and get the real greedy path P. The following scenarios may happen:

- $P \neq P'$. In this case we are done because $P \neq P'$ means the greedy routing on G has gone out of G', which implies the algorithm has reached some vertex in G with weight at least w_0.
- $P = P'$. Consider a $w_0^{3d/\alpha}$-grid on $\mathbb{Z}^d$. By definition each $w_0^{3d/\alpha}$-cell contains $w_0^{3d/\alpha}$ vertices. Then greedy routing either stops before it visits $w_0^{6d/\alpha}$ vertices, or after exploring for $w_0^{6d/\alpha}$ steps, it has visited at least $w_0^{3d/\alpha}$ different cells. Let $C_1, \dots, C_{w_0^{3d/\alpha}}$ be the first $w_0^{3d/\alpha}$ cells in the grid the greedy routing algorithm goes through, and choose $y_i \in P \cap C_i$. Note here y_i is not necessarily x_i because the path may have more vertices in a single cell. Further let $M_i := \{x \in C_i | W_x \geq w_0^3, x \sim y_i\}$ and $M := \bigcup M_i$. Then

$$\begin{aligned}\mathbb{E}\left[|M|\right] &= \sum_i \mathbb{E}\left[|M_i|\right] = w_0^{3d/\alpha}\mathbb{E}\left[|M_i|\right] \\ &\geq w_0^{3d/\alpha} w_0^{3d/\alpha} w_0^{-3(\tau-1)} \frac{w_0^3}{\left(w_0^{3/\alpha}\right)^\alpha} = w_0^{\frac{3d}{\alpha}(2-\gamma)}.\end{aligned}$$

  Here we use the fact that if x and y_i are in the same cell, we have $|x-y_i| \leq w_0^{3/\alpha}$. Since $\gamma \in (1,2)$ we know $\mathbb{E}[|M|] \to \infty$ as $N \to \infty$. By the Chernoff bound from Lemma 3.1, with high probability there exists a smallest index $i \in [w_0^{3d/\alpha}]$ such that there is a vertex y in the cell C_i with $W_y \geq w_0$ and $y \sim y_i$.

It remains to show that the vertex y we find above has a large enough objective. Let y' be a neighbor of y_i with $W_{y'} < w_0$. By Proposition 3.16 (b) we know with high probability $|y' - y_i| \leq w_0^c$ and hence $|y' - t| \geq |y_i - t| - w_0^c$. Now we consider the locations of the vertices. Since $\phi(s) \leq e^{-w_0}$ one has $|s-t|^\alpha \geq e^{w_0}$. On the other hand, $|y_i - s| \leq w_0^{c+6d/\alpha} \ll |s-t|$ for N large enough. Then

$$|y_i - t| \geq \frac{1}{2}|s-t| \geq \frac{1}{2}e^{w_0} \gg w_0^k$$

for any $k > 0$ and N large. Consequently for N large, one has $|y_i - t| + w_0^{3/\alpha} \leq 2\left(|y_i - t| - w_0^c\right)$. Then

$$\begin{aligned}\phi(y) &= \frac{W_y}{|y-t|^\alpha} \geq \frac{w_0^3}{\left(|y_i - t| + w_0^{3/\alpha}\right)^\alpha} \geq \frac{w_0^3}{2^\alpha\left(|y_i - t| - w_0^c\right)^\alpha} \\ &> \frac{w_0}{\left(|y_i - t| - w_0^c\right)^\alpha} > \phi(y').\end{aligned}$$

Therefore the greedy routing will not visit the neighbors of y_i with weights less than w_0 since y_i has at least one neighbor with higher objective than all these vertices. □

Proof of Theorem 3.15. Let f_0 be a function as in (3.4), and w_0, ϕ_0 be as follows:

$$w_0(N) := \max\{\log f_0(N), W_s\}, \qquad \text{and } \phi_0(N) := \min\{W_t^{-1}, f_0^{-1}(N)\}.$$

Apparently w_0 and ϕ_0 satisfy (3.18). Assume now $\phi(s) \leq \phi_0$, otherwise we skip both start and main stage. Furthermore, for N large, we have

$$\phi(s) = \frac{W_s}{N^\alpha} \leq \frac{w_0}{N^\alpha} \ll e^{-w_0}.$$

Considering the value of w_0 we have two possible cases:

i) If $w_0 = W_s$, then we already start with some vertex with weight at least w_0, and hence the start stage will be skipped.

ii) If $w_0 = \log f_0(N)$, then by Proposition 3.18, with high probability, T reaches after at most $w_0^{6d/\alpha}$ steps some vertex x_1 with $W_{x_1} \geq w_0$.

By Proposition 3.13 we know with high probability starting from x_1 within $\ell+1$ steps T visits a vertex $x_{\ell+1}$ with objective at least ϕ_0, where ℓ is bounded as follows:

$$\ell \leq 2f_0(N) + \frac{\log\log_{w_0}(\phi(s)^{-1})}{\log\chi(\zeta\epsilon_2)} + \frac{\log\log_{\phi_0^{-1}}(\phi(s)^{-1})}{\log\chi(\epsilon_2)}.$$

Assume now T reaches $x_{\ell+1}$ with $\phi(x_{\ell+1}) \geq \phi_0$. For ϕ_0 we have two cases:

i) If $\phi_0 = W_t^{-1}$. Then we know

$$p_{x_{\ell+1},t} = \mathbb{E}\left[\frac{W_{x_{\ell+1}}W_t}{|x_{\ell+1}-t|^\alpha} \wedge 1\right] = \mathbb{E}\left[(\phi(x_{\ell+1})W_t) \wedge 1\right] = 1.$$

In this case, T jumps to t from $x_{\ell+1}$ with probability 1.

ii) If $\phi_0 = f_0^{-1}$. Let N_{ϕ_0} be the number of vertices in $\mathbb{Z}^d$ with objective at least ϕ_0. Then

$$\begin{aligned}\mathbb{E}(N_{\phi_0}) &= \sum_{x\in\mathbb{Z}^d, x\neq t} \mathbb{P}(\phi(x) \geq \phi_0) = \sum_{x\in\mathbb{Z}^d, x\neq t} \mathbb{P}(W_x \geq |x-t|^\alpha \phi_0) \\ &= \sum_{x\in\mathbb{Z}^d, x\neq t} \frac{\phi_0^{1-\tau}}{|x-t|^{\alpha(\tau-1)}} = C\phi_0^{-\tau+1},\end{aligned}$$

where $C := \sum_{x\in\mathbb{Z}^d, x\neq t} \frac{1}{|x-t|^{\alpha(\tau-1)}} < \infty$ because $\gamma = \frac{\alpha(\tau-1)}{d} > 1$. By Markov's inequality, with probability at most $C\phi_0^{3-\tau}$ we have $N_{\phi_0} \geq \phi_0^{-2}$. Since $\tau \in$

$(1,3)$, one has $\lim_{N\to\infty} C\phi_0^{3-\tau} = 0$. Therefore with high probability we have $N_{\phi_0} \le \phi_0^{-2}$. Conditioned on the event $N_{\phi_0} \le \phi_0^{-2}$, from $x_{\ell+1}$ greedy routing will continue with at most ϕ_0^{-2} steps because the routing protocol only admits vertices with higher objective.

To summerize, with high probability, the length of greedy path L satisfies

$$\begin{aligned}
L &\le w_0^{6d/\alpha} + \ell + \phi_0^{-2} \\
&= \frac{\log\log_{w_0}(\phi(s)^{-1})}{\log\chi(\zeta\epsilon_2)} + \frac{\log\log_{\phi_0^{-1}}(\phi(s)^{-1})}{\log\chi(\epsilon_2)} + f(N) \\
&= \frac{\log\log_{w_0}(\phi(s)^{-1}) + \log\log_{\phi_0^{-1}}(\phi(s)^{-1})}{|\log(\gamma-1)| + o(1)} + f(N) \\
&\le \frac{1+o(1)}{|\log(\gamma-1)|}\left(\log\log_{W_s}\left(\phi(s)^{-1}\right) + \log\log_{W_t}\left(\phi(s)^{-1}\right)\right) + f(N),
\end{aligned}$$

as $N\to\infty$, where

$$f(N) := (\log f_0(N))^{6d/\alpha} + f_0(N)^2 + 2f_0(N) = o(\log\log N)$$

by our choice of f_0 in (3.4). □

3.2.3 A patching method

In this section we propose a patching solution such that even if the greedy algorithm reaches some local optimum, it can still continue. The patching protocol goes as follows:

Let $V^{(i)}$ be the set of vertices the greedy algorithm T has explored after i steps, and $V_N^{(i)}$ be the set of unexplored neighbors of all vertices in $V^{(i)}$. That is,

$$V_N^{(i)} := \left\{ y \in \bigcup_{x\in V^{(i)}} \Gamma(x) \,\middle|\, y \text{ is unexplored} \right\}.$$

When T arrives at some local optimum at step i, then it will go to the vertex in $V_N^{(i)}$ with highest objective in the next step and resume the greedy routing from there.

In other words,

$$x_{i+1} = \arg\max_{x \in V_N^{(i)}} \phi(x),$$

where x_{i+1} is the $i+1$-th vertex in the routing path.

For the patched greedy routing algorithm we have the following result:

Theorem 3.19 (Part (c) in Theorem 1.6). *Consider scale-free percolation with connection probability (1.9), and parameters $\alpha > d, \gamma \in (1,2)$. Let T be the greedy routing algorithm with objective function as in (1.10). Furthermore, we assume T admits the patching protocol. Then, conditional on W_s and W_t, with high probability, T finds the target t within L_3 steps as $N \to \infty$, where L_3 is a function of N given as follows:*

$$L_3 = \frac{1+o(1)}{|\log(\gamma-1)|}\Bigg(\log\log_{W_s}\left(\phi(s)^{-1}\right) + \log\log_{W_t}\left(\phi(s)^{-1}\right)\Bigg).$$

Theorem 3.19 tells that with high probability the greedy routing algorithm with patching protocol finds the target successfully within doubly logarithmic steps. Analogous to Section 3.2.1 and Section 3.2.2, the routing process is divided into three stages. By Proposition 3.13 with high probability the abortion of T does not occur in the main part, therefore we don't state here the proposition for this phase repeatedly.

Proposition 3.20 (Start of patching). *Let $w_0 = w_0(N)$ be a function fulfilling*

$$\lim_{N\to\infty} w_0(N) = \infty, \quad \textit{and} \quad \limsup_{N\to\infty} \frac{w_0(N)}{\log\log\log N} \le C$$

for some positive constant C. Assume $\phi(s) \le e^{-w_0}$. Then starting from s, the patched greedy algorithm T finds either the target t or some vertex with weight at least w_0 within $o(\log\log N)$ steps.

Proof. The proof is trivial, if $W_s \ge w_0$. Otherwise we let T run on the graph from s. Two scenarios may happen:

1. T stops before it visits $w_0^{6d/\alpha}$ different vertices. In this case, with the patching protocol the only reason is that T already finds the target t and therefore finishes the routing ahead. Assume now T has visited k different vertices. By the patching protocol it takes at most k jumps to reach the next unexplored vertex. Therefore,

T needs at most $\sum_{k=1}^{n} k = n(n+1)/2 \leq n^2$ steps to visit n different vertices. In this scenario, the length L of a greedy path is at most $w_0^{12d/\alpha} = o(\log\log N)$ as $N \to \infty$;

2. T visits $w_0^{6d/\alpha}$ different vertices. In this case we know from Proposition 3.18 with high probability T finds some vertex u with weight $W_u \geq w_0$ within $\left(w_0^{6d/\alpha}\right)^2 = o(\log\log N)$ steps.

This finishes the proof. □

Proposition 3.21 (End of patching)**.** *Let $w_0 = w_0(N)$ be a positive function with $\lim_{N\to\infty} w_0(N) = \infty$. Further let k be a positive constant fulfilling*

$$k > \frac{\tau - 1}{d}. \tag{3.23}$$

Suppose the target t has weight $W_t \leq w_0$. Then with high probability, there exists an open path P of length at most w_0^k from t to some vertex v with $W_v \geq w_0$ such that for each $y \in P$ it holds that $\phi(y) \geq w_0^{-k\alpha}$.

Proof. Due to the existence of nearest edges, we can prove this proposition in a easy way. Let B be the ball around t with volume w_0^{kd}. Then every vertex in B is joined with t. We consider the probability of the event E that there exists a vertex v in B with weight at least w_0.

$$\begin{aligned}\mathbb{P}(E^c) &= \mathbb{P}\left(\bigcap_{x\in B}\{W_x < w_0\}\right) = \left(1 - w_0^{-(\tau-1)}\right)^{w_0^{kd}} \\ &= \left[\left(1 - \frac{1}{w_0^{\tau-1}}\right)^{w_0^{\tau-1}}\right]^{w_0^{kd-(\tau-1)}}.\end{aligned}$$

Therefore if $kd > \tau - 1$ we get $\mathbb{P}(E) \to 1$ as $N \to \infty$. Let v be the vertex in B with $W_v \geq w_0$, then $|v - t| \leq w_0^k$ and hence

$$\phi(v) = \frac{W_v}{|v-t|^\alpha} \geq \frac{w_0}{w_0^{k\alpha}} = w_0^{1-k\alpha} \geq w_0^{-k\alpha}.$$

Besides, let P be an open shortest path joining t and v using only nearest edges, then the length of P is at most w_0^k. For $y \in P$ with $y \neq t$ and $y \neq v$ it holds that

$|y-t| \le |v-t|$ and

$$\phi(y) = \frac{W_y}{|y-t|^\alpha} \ge \frac{1}{w_0^{k\alpha}} = w_0^{-k\alpha}.$$

□

Proof of Theorem 3.19. Let $w_0 = w_0(N)$ be a function as in Proposition 3.20 and $\phi_0 = w_0^{-1/2}$. If $W_s \ge w_0$, the greedy algorithm T already finds some vertex with weight at least w_0. Otherwise

$$\phi(s) = \frac{W_s}{|s-t|^\alpha} \le \frac{w_0}{N^\alpha} \le e^{-w_0}, \text{ for } N \text{ large.}$$

By Proposition 3.20, with high probability T finds the target t or some vertex u with weight at least w_0 within $o(\log\log N)$ steps. In the former case we are done. So it is sufficient to consider the second case.

Assume now the patched algorithm T visits u with $W_u \ge w_0$. By Proposition 3.13, with high probability T finds some vertex u_1 with $\phi(u_1) \ge \phi_0$ within $\ell + 1$ steps, and ℓ satisfies

$$\begin{aligned}\ell &= 2f_0(N) + \frac{\log\log_{w_0}(\phi(x_1)^{-1})}{\log\chi(\zeta\epsilon_2)} + \frac{\log\log_{\phi_0^{-1}}(\phi(x_1)^{-1})}{\log\chi(\epsilon_2)} \\ &\le 2f_0(N) + \frac{\log\log_{w_0}(\phi(s)^{-1})}{\log\chi(\zeta\epsilon_2)} + \frac{\log\log_{\phi_0^{-1}}(\phi(s)^{-1})}{\log\chi(\epsilon_2)}.\end{aligned}$$

Denote by $G_{\ge\phi}$ the subgraph of scale-free percolation on $\mathbb{Z}^d$ induced by vertices with objective at least ϕ. We first consider the expected size of $G_{\ge\phi}$:

$$\begin{aligned}\mathbb{E}\left[|G_{\ge\phi}|\right] &= \sum_{x\in\mathbb{Z}^d : x\neq t} \mathbb{P}(\phi(x)\ge\phi) + 1 = \sum_{x\in\mathbb{Z}^d : x\neq t} \mathbb{P}(W_x \ge \phi|x-t|^\alpha) + 1 \\ &= \sum_{x\in\mathbb{Z}^d : x\neq t} \phi^{-\tau+1}|x-t|^{-\alpha(\tau-1)} + 1 = C\phi^{-\tau+1}.\end{aligned}$$

for some positive constant $C > 0$. By the Chernoff bound in Lemma 3.1, with high probability, $|G_{\ge\phi}| \le 2C\phi^{-(\tau-1)}$ if $\phi := \phi(N) \to 0$ as $N \to \infty$. Now we take $k = \frac{2}{\alpha}$. Note that this choice of k satisfies (3.23) since it holds $\gamma := \frac{\alpha(\tau-1)}{d} < 2$ in the doubly logarithmic regime. We consider $G_{\ge w_0^{-2}}$. Since $\phi(u_1) \ge \phi_0 = w_0^{-1/2} \ge w_0^{-2}$, we have $u_1 \in G_{\ge w_0^{-2}}$. Let $\mathcal{C}$ be the cluster of $G_{\ge w_0^{-2}}$ in which u_1 lies. We continue with the patched greedy algorithm from u_1. If $t \in \mathcal{C}$, we estimate the number of steps in order to find t. Assume now T has explored k vertices in $\mathcal{C}$. To visit an unexplored vertex it takes at most k steps by the patching protocol. As a result, T needs at most $\sum_{k=1}^{|\mathcal{C}|} k \leq |\mathcal{C}|^2 \leq |G_{\geq w_0^{-2}}|^2 = o(\log\log N)$ steps to find t. So it is sufficient to show that with high probability t is also in the cluster $\mathcal{C}$.

By Proposition 3.21 we know with high probability there exists an open path that joins t and v for some $v \in \mathbb{Z}^d$ with $W_v \geq w_0$. Denote $r := |v-t|$. We have two cases depending on the distance between u_1 and t:

i) $|u_1 - t| \geq r$. In this case u_1 is still far away from the target. We consider the probability that u_1 is connected to v:

$$\begin{aligned}\mathbb{P}(u_1 \sim v) &= \mathbb{E}\left[\frac{W_{u_1}W_v}{|u_1 - v|^\alpha} \wedge 1\right] \geq \mathbb{E}\left[\frac{W_{u_1}W_v}{(|u_1 - t| + r)^\alpha} \wedge 1\right] \\ &\geq \mathbb{E}\left[\frac{W_{u_1}W_v}{2^\alpha |u_1 - t|^\alpha} \wedge 1\right] \geq \frac{\phi_0 w_0}{2^\alpha} \wedge 1 = \frac{w_0^{1/2}}{2^\alpha} \wedge 1 = 1.\end{aligned}$$

The last step holds for N large, since $\lim_{N\to\infty} w_0(N) = \infty$.

ii) $|u_1 - t| \leq r$. From the proof of Proposition 3.21 we know $|u_1 - t| \leq r \leq w_0^{2/\alpha}$. Therefore $u_1 \in B$ with $B := \{x \in \mathbb{Z}^d : |x - t| \leq w_0^{2/\alpha}\}$. It is clear that all vertices in B are joined with t and have objective at least w_0^{-2}. This means $B \subseteq \mathcal{C}$, and in particular $t \in \mathcal{C}$.

To summarize, with high probability, the patched greedy algorithm finds the target t within L steps, where L is subject to the following bound:

$$\begin{aligned}L &\leq w_0^{12\alpha/d} + 2f_0(N) + \frac{\log\log_{w_0}(\phi(s)^{-1})}{\log \chi(\zeta\epsilon_2)} + \frac{\log\log_{\phi_0^{-1}}(\phi(s)^{-1})}{\log\chi(\epsilon_2)} + 2Cw_0^{2(\tau-1)} \\ &\leq \frac{1+o(1)}{|\log(\gamma-1)|}\left(\log\log_{W_s}\left(\phi(s)^{-1}\right) + \log\log_{W_t}\left(\phi(s)^{-1}\right)\right).\end{aligned}$$

□

Bibliography

[1] M. Aizenman, F. Deylon, and B. Souillard, *Lower bounds on the cluster size distribution*, J. Stat. Phys. **23** (1980), 267–280.

[2] M. Aizenman, H. Kesten, and C. M. Newman, *Uniqueness of the infinite cluster and continuity of connectivity functions for short and long range percolation*, Comm. Math. Phys. **111** (1987), no. 4, 505–531.

[3] N. Alon and J. Spencer, *The probabilistic method*, 2 ed., John Wiley & Sons, 2000.

[4] T. L. Austin, R. E. Fagen, W. F. Penney, and J. Riordan, *The number of components in random linear graphs*, Ann. Math. Statist. **30** (1959), 747–754.

[5] A.-L. Barabási and R. Albert, *Emergence of scaling in random networks*, Science **286** (1999), no. 5439, 509–512.

[6] D. J. Barsky, G. Grimmett, and C. M. Newman, *Percolation in half-spaces: Equality of critical densities and continuity of the percolation probability*, Probab. Theory Relat. Fields **90** (1991), no. 1, 111–148.

[7] I. Benjamini, H. Kesten, Y. Peres, and O. Schramm, *Geometry of the uniform spanning forest: Transitions in dimensions 4, 8, 12,...*, Ann. Math. **160** (2004), no. 2, 465–491.

[8] J. van der Berg and M. Keane, *On the continuity of the percolation probability function*, Conference on Modern Analysis and Probability, vol. 26, Contemporary Mathematics, 1984, pp. 61–65.

[9] N. Berger, *A lower bound for the chemical distance in sparse long-range percolation models*, arXiv:math/0409021 (2004).

[10] N. Berger, C. Borgs, J. T. Chayes, R. M. D'Souza, and R. D. Kleinberg, *Competition-induced preferential attachment*, Automata, languages and pro-

gramming, vol. 3142, Springer, 2004, pp. 208–221.

[11] ———, *Degree distribution of competition-induced preferential attachment graphs.*, Combin. Probab. Comput. **14** (2005), no. 5-6, 697–721.

[12] M. Biskup, *On the scaling of the chemical distance in long-range percolation models*, Ann. Probab. **32** (2004), no. 4, 2938–2977.

[13] ———, *Graph diameter in long-range percolation*, Random Struct. Algorithms **39** (2011), no. 2, 210–227.

[14] M. Biskup and J. Lin, *Sharp asymptotic for the chemical distance in long-range percolation*, Random Struct. Algorithms **55** (2019), no. 3, 560–583.

[15] B. Bollobás, *Random graphs*, 2 ed., Cambridge Unversity Press, 2001.

[16] B. Bollobás, C. Borgs, J. Chayes, and O. Riordan, *Directed scale-free graphs*, Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (Baltimore, MD, 2003)., ACM, 2003, pp. 132–139.

[17] B. Bollobás, O. Riordan, J. Spencer, and G. Tusnády, *The degree sequence of a scale-free random graph process*, Random Struct. Algorithms **18** (2001), no. 3, 279–290.

[18] C. Borgs, J. Chayes, R. van der Hofstad, and J. Spencer, *Random subgraphs of finite graphs. I. The scaling window under the triangle condition*, Random Struct. Algorithms **27** (2005), no. 2, 137–184.

[19] K. Bringmann, R. Keusch, and J. Lengler, *Average distance in a general class of scale-free networks with underlying geometry*, arXiv:1602.05712.

[20] K. Bringmann, R. Keusch, J. Lengler, Y. Maus, and A. R. Molla, *Greedy routing and the algorithmic small-world phenomenon*, Proceedings of the 36th ACM symposium on principles of distributed computing, PODC '17, Washington, DC, USA, July 25–27, 2017, New York, NY: Association for Computing Machinery (ACM), 2017, pp. 371–380.

[21] K. Bringmann, R. Keusch, J. Lengler, Y. Maus, and A. R. Molla, *Greedy routing and the algorithmic small-world phenomenom*, J. Comput. System Sci. **125** (2022), 59–105.

[22] T. Britton, M. Deijfen, and A. Martin-Löf, *Generating simple random graphs with prescribed degree distribution*, J. Stat. Phys. **124** (2006), no. 6, 1377–1397.

[23] S. R. Broadbent and J. M. Hammersley, *Percolation processes I. Crystals and mazes*, Proc. Cambr. Phil. Soc **53** (1957), 629–641.

[24] R. M. Burton and M. Keane, *Density and uniqueness in percolation*, Comm. Math. Phys. **121** (1989), no. 3, 501–505.

[25] C. Cooper and A. Frieze, *A general model of web graphs*, Random Struct. Algorithms **22** (2003), no. 3, 311–335.

[26] J. Dalmau and M. Salvi, *Scale-free percolation in continuum space: quenched degree and clustering coefficient*, J. Appl. Probab. **58** (2021), no. 1, 106–127.

[27] G. De Masi, G. Iori, and G. Caldarelli, *Fitness model for the Italian interbank money market*, Phys. Rev. E **74** (2006), no. 6, 066112.

[28] M. Deijfen, R. van der Hofstad, and G. Hooghiemstra, *Scale-free percolation*, Ann. Inst. H. Poincaré Probab. Statist. **49** (2013), no. 3, 817–838.

[29] P. Deprez, R. S. Hazra, and M. V. Wüthrich, *Inhomogeneous long-range percolation for real-life network modeling*, Risks **3** (2015), no. 1, 1–23.

[30] P. Deprez and M.V. Wüthrich, *Scale-free percolation in continuum space*, Commun. Math. Stat. **7** (2019), no. 3, 269–308.

[31] S. Dereich and P. Mörter, *Random networks with sublinear preferential attachment: Degree evolutions*, Electron. J. Probab. **14** (2009), 1222–1267.

[32] ______, *Random networks with concave preferential attachment rule*, Jahresber. Dtsch. Math.-Ver. **113** (2011), no. 1, 21–40.

[33] ______, *Random networks with sublinear preferential attachment: the giant component*, Ann. Probab. **41** (2013), no. 1, 329–384.

[34] J. Ding and A. Sly, *Distances in critical long range percolation*, arXiv:1303.3995 (2018).

[35] M. Draief and A. Ganesh, *Efficient routeing in Poisson small-world networks*, J. Appl. Probab. **43** (2006), no. 3, 678–686.

[36] H. Duminil-Copin, V. Sidoravicius, and V. Tassion, *Absence of infinite cluster for critical Bernoulli percolation on slabs*, Comm. Pure Appl. Math. **69** (2016), no. 7, 1397–1411.

[37] H. Duminil-Copin and V. Tassion, *A new proof of the sharpness of the phase transition for Bernoulli percolation on* $\mathbb{Z}^d$, Enseign. Math. **62** (2016), no. 1,

199–206.

[38] R. Durrett, *Probability: Theory and examples*, 5 ed., Cambridge University Press, 2019.

[39] P. Erdős and A. Rényi, *On random graphs. I.*, Publ. Math. Debrecen **6** (1959), 290–297.

[40] ______, *On the evolution of random graphs*, Magyar Tud. Akad. Mat. Kutató Int. Közl **5** (1960), 17–61.

[41] ______, *On the evolution of random graphs*, Bull. Inst. Internat. Statist **38** (1961), 343–347.

[42] ______, *On the strength of connectedness of a random graph*, Acta Math. Acad. Sci. Hungar **12** (1961), 261–267.

[43] H. van den Esker, R. van der Hofstad, and G. Hooghiemstra, *Universality for the distance in finite variance random graphs*, J. Stat. Phys. **133** (2008), no. 1, 169–202.

[44] R. Fitzner and R. van der Hofstad, *Mean-field behavior for nearest-neighbor percolation in $d > 10$*, Electron. J. Probab. **22** (2017), no. 43, 1–65.

[45] M. Franceschetti and R. Meester, *Navigation in small-world Networks: a scale-free continuum model*, J. Appl. Probab. **43** (2006), no. 4, 1173–1180.

[46] A. Gandolfi, M. S. Keane, and C. M. Newman, *Uniqueness of the infinite component in a random graph with applications to percolation and spin glasses*, Probab. Theory Relat. Fields **92** (1992), no. 4, 521–527.

[47] D. S. Gaunt and H. Ruskin, *Bond percolation processes in d-dimensions*, J. Phys. A: Math. Gen. **11** (1978), no. 7, 1369–1380.

[48] D. S. Gaunt, M. S. Sykes, and H. Ruskin, *Percolation processes in d-dimensions*, J. Phys. A: Math. Gen. **9** (1976), no. 11, 1899–1911.

[49] E. N. Gilbert, *Random graphs*, Ann. Math. Statist. **30** (1959), 1141–1144.

[50] P. Gracar, M. Heydenreich, C. Mönch, and P Mörters, *Recurrence versus Transience for Weight-Dependent Random Connection Models*, Electron. J. Probab. **27** (2022), 1–31.

[51] P. Gracar, L. Lüchtrath, and P. Mörters, *Percolation phase transition in weight-dependent random connection models*, Adv. Appl. Probab. **53** (2021), no. 4,

1090–1114.

[52] P. Grassberger, *Critical percolation in high dimensions*, Phys. Rev. E **67** (2003), no. 3, 036101.

[53] G. Grimmett, *Percolation*, 2 ed., Springer Verlag, 1999.

[54] J. M. Hammersley, *Percolation processes: Lower bounds for the critical probability*, Ann. Math. Statist. **28** (1957), no. 3, 790–795.

[55] N. Hao and M. Heydenreich, *Graph distances in scale-free percolation: the logarithmic case*, J. Appl. Probab. **60** (2023), no. 1, 295–313.

[56] T. Hara and G. Slade, *Mean-field critical behaviour for percolation in high dimensions*, Commun. Math. Phys. **128** (1990), no. 2, 333–391.

[57] ———, *The Self-Avoiding-Walk and Percolation Critical Points in High Dimensions*, Comb. Probab. Comput. **4** (1995), no. 3, 197–215.

[58] T. E. Harris, *A lower bound for the critical probability in a certain percolation process*, Proc. Cambr. Phil. Soc. **56** (1960), 13–20.

[59] M. Heydenreich, T. Hulshof, and J. Jorritsma, *Structures in supercritical scale-free percolation*, Ann. Appl. Probab. **27** (2017), no. 4, 2569–2604.

[60] M. Heydenreich and K. Matzke, *Expansion for the critical point of site percolation: The first three terms*, Comb. Probab. Comput. **31** (2022), no. 3, 430 – 454.

[61] C. Hirsch and C. Mönch, *Distances and large deviations in the spatial preferential attachment model*, Bernoulli **26** (2020), no. 2, 927–947.

[62] R. van der Hofstad, *Random graphs and complex networks: Volume 2*, Preliminary version available at https://www.win.tue.nl/~rhofstad/.

[63] ———, *Random graphs and complex networks: Volume 1*, 1 ed., Cambridge University Press, 2016, a digital version available at https://www.win.tue.nl/~rhofstad/NotesRGCN.pdf.

[64] R. van der Hofstad and J. Kompathy, *Explosion and distances in scale-free percolation*, arXiv:1706.02597 (2017).

[65] S. Janson and J. Spencer, *A point process describing the component sizes in the critical window of the random graph evolution*, Combin. Probab. Comput. **16** (2008), no. 4, 631–658.

[66] H. Kesten, *The critical probability of bond percolation on the square lattice equals* 1/2, Comm. Math. Phys. **74** (1980), no. 1, 41–59.

[67] J. Kleinberg, *The small-world phenomenon: An algorithmic perspective*, Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, STOC '00, Association for Computing Machinery, 2000, pp. 163–170.

[68] R. Lyons and Y. Peres, *Probability on trees and networks*, 1 ed., Cambridge University Press, 2017.

[69] S. Mertens and C. Moore, *Percolation thresholds and Fisher exponents in hypercubic lattices*, Phys. Rev. E **98** (2018), no. 2, 022120.

[70] S. Milgram, *The small world problem*, Psychol. Today **1** (1967), no. 1, 61–67.

[71] M. Mitzenmacher and E. Upfal, *Probability and Computing: Randomization and Probabilistic Techniques in Algorithms and Data Analysis*, 2 ed., Cambridge University Press, 2017.

[72] C. M. Newman and L. S. Schulman, *One dimensional* $1/|x-y|^s$ *percolation models: the existence for a transition for* $s \leq 2$, Commun. Math. Phys. **104** (1986), no. 4, 547–571.

[73] B. Pittel, *On the largest component of the random graph at a nearcritical stage*, J. Combin. Theory Ser. B **82** (2001), no. 2, 237–269.

[74] D. D. S. Price, *A general theory of bibliometric and other cumulative advantage processes*, J. Amer. Soc. Inform. Sci. **27** (1976), no. 5, 292–306.

[75] H. A. Simon, *On a class of skew distribution functions*, Biometrika **42** (1955), no. 3–4, 425–440.

[76] P. Trapman, *The growth of the infinite long-range percolation cluster*, Ann. Probab. **38** (2010), no. 4, 1583–1608.

[77] U. Yule, *A mathematical theory of evolution, based on the conclusions of Dr. J. C. Willis, F. R. S*, Philos. Trans. R. Soc. B **213** (1929), 402–410.