

Greedy Routing and the Algorithmic Small-World Phenomenon

Karl Bringmann* Ralph Keusch† Johannes Lengler‡ Yannic Maus§
Anisur R. Molla¶

December 20, 2016

Abstract

The algorithmic small-world phenomenon, empirically established by Milgram’s letter forwarding experiments from the 60s [47], was theoretically explained by Kleinberg in 2000 [36]. However, from today’s perspective his model has several severe shortcomings that limit the applicability to real-world networks. In order to give a more convincing explanation of the algorithmic small-world phenomenon, we study greedy routing in a more realistic random graph model (geometric inhomogeneous random graphs), which overcomes the previous shortcomings. Apart from exhibiting good properties in theory, it has also been extensively experimentally validated that this model reasonably captures real-world networks.

In this model, we show that greedy routing succeeds with constant probability, and in case of success almost surely finds a path that is an almost shortest path. Our results are robust to changes in the model parameters and the routing objective. Moreover, since constant success probability is too low for technical applications, we study natural local patching methods augmenting greedy routing by backtracking and we show that such methods can ensure success probability 1 in a number of steps that is close to the shortest path length.

These results also address the question of Krioukov et al. [41] whether there are efficient local routing protocols for the internet graph. There were promising experimental studies, but the question remained unsolved theoretically. Our results give for the first time a rigorous and analytical answer, assuming our random graph model.

*Max-Planck-Institute for Informatics, Saarbrücken, Germany, kbringma@mpi-inf.mpg.de

†Department of Computer Science, ETH Zurich, Switzerland, rkeusch@inf.ethz.ch

‡Department of Computer Science, ETH Zurich, Switzerland, lenglerj@inf.ethz.ch

§Department of Computer Science, University of Freiburg, Germany, yannic.maus@cs.uni-freiburg.de

¶Department of Computer Science, University of Freiburg, Germany, armolla@cs.uni-freiburg.de

Contents

1 Introduction	2
1.1 Related Work and Choice of Model	3
2 Random Graph Model & Greedy Routing	6
2.1 Geometric Inhomogeneous Random Graphs (GIRGs)	6
2.2 Greedy routing	7
3 Results	8
4 Comparison with Experimental Results	11
5 Patching Algorithms: Discussion and Examples	12
6 Proof Sketches and the Typical Trajectory of a Greedy Path	15
7 Notation & Basic Properties of GIRGs	18
7.1 Notation	18
7.2 Properties of GIRGs	18
7.3 Where to Expect Neighbors & Unlikely Jumps	22
8 Proof: Success Probability and Length of Greedy Routing	30
8.1 Main Lemma	30
8.2 Proof of Theorem 3.1 (Success Probability)	36
8.3 Proof of Theorem 3.2 (Failure Probability is Exponentially Small in \mathbf{w}_{\min})	38
8.4 Proof of Theorem 3.3 (Length of Greedy Path)	42
9 Proof: Patching	44
10 Proof: Relaxations	47
11 Proof: Greedy Routing on Hyperbolic Random Graphs	52
References	55

1 Introduction

The idea that everyone is connected to everybody else through six degrees of separation, also known as the small world phenomenon, was empirically established by the classic experiments of Milgram [47, 59]. Several random people were given a target person’s name and address, and asked to send a letter to the target person by forwarding it to a personal acquaintance who is more likely to know the target person, iterating this process until the letter reached the target (or was lost). Among the about 30% successful trials, Milgram reported an average distance of about 6. This showed not only that very short paths exist between two typical persons, but also that people are able to find them without global knowledge of the network, by greedily routing the letter to an acquaintance that is closer to the target.

In the last 50 years, this phenomenon inspired researchers in various disciplines such as sociology, physics, and experimental and theoretical computer science to study networks with small-world properties. The first theoretical explanation of the algorithmic small-world phenomenon was the seminal work of Kleinberg [36] who showed that greedy routing succeeds in $O(\log^2 n)$ steps on certain random graphs. However, as we will discuss in Section 1.1, his model has several severe shortcomings, that have not been resolved in subsequent work [26, 46, 25, 45, 16]: routing succeeds only if a perfect lattice structure is contained in the network, the result is fragile with respect to small changes in the model, and the graph model is unrealistic since it is homogeneous.

In order to give a more convincing theoretical explanation for Milgram’s result, we study greedy routing on a more realistic graph model: In geometric inhomogeneous random graphs (GIRGs) [12], every node draws randomly a location in some geometric space, and a positive weight according to a power law which represents the node’s influence. Two nodes are then more likely to connect if they have larger weights and if they are geometrically close. We picked this model for the following reasons, which we substantiate in Section 1.1.

- (1) GIRGs are a natural extension of Kleinberg’s model that overcome its three main shortcomings.
- (2) Among the random graph models that have been proposed to model social networks, the GIRG model stands out because it has not only theoretically good properties, but a special case of GIRGs (hyperbolic random graphs [8]) has also been extensively experimentally validated [56, 8, 44, 50, 29]. E.g., Boguñá et al. [8] computed and studied a (heuristic) maximum likelihood fit of the internet graph into the hyperbolic random graph model, based on ideas from [38].
- (3) In general it is unclear how greedy routing generalizes to inhomogeneous geometric graphs. Since GIRGs have a specific connection probability for each pair of vertices, there is a natural interpretation of Milgram’s instruction to route to an “acquaintance who is more likely than you to know the target person” [59]: pick the neighbor most likely being adjacent to the target.
- (4) Krioukov et al. [41] identified as a major open problem in the area of routing protocols “whether we can devise routing protocols for the internet that, having no full view of the network topology, can still efficiently route messages”. This question stimulated further research on the embeddings from [8, 38], and by now routing on such graphs is experimentally well-studied [50, 5, 6, 7, 19, 42, 49, 43]. It turned out that greedy routing works experimentally surprisingly well on hyperbolic random graphs and GIRGs [42, 43, 49]. For example, the resulting greedy paths are only marginally longer than shortest paths in the network. Explaining these experimental findings theoretically was a driving question for many researchers in the theory community [52, 12, 13, 28, 35, 24, 15, 30], but it remained open. We give the first theoretical explanation for these experiments, and give (in our model) a strongly affirmative answer to Krioukov et al.’s question.

We thus believe that studying greedy routing on GIRGs, in each step maximizing the connection probability to the target, gives a convincing theoretical explanation for Milgram’s experiments, and addresses at the same time the question of Krioukov et al. We obtain the following results for any fixed source s and target t (for exact statements see Section 3).

- Greedy routing succeeds with probability $\Omega(1)$.

This is already surprising, as incorporating noisy positions in Kleinberg’s model results in a tiny success probability, as discussed in Section 1.1. The result is also best possible, as in GIRGs the target node is isolated with constant probability.

- A.a.s.¹, the stretch² of greedy routing in case of success is $1 + o(1)$. Hence, greedy routing is asymptotically optimal. In particular, since the average distance in the giant component of the GIRG model is ultra-small, i.e., $O(\log \log n)$, so are the lengths of the greedy paths.

One may speculate that $O(\log \log n)$ is closer to six degrees of separation than Kleinberg’s $O(\log^2 n)$. All previous analyses of greedy routing in small-world models yielded at least $\Omega(\log n)$ steps.

- Our results are robust in the model parameters, e.g., they hold for all parameter choices of the model, in particular for any power law exponents in the natural range $(2, 3)$, cf. Section 1.1. Moreover, we show that nodes do not need to know exact weights and positions of their neighbors, but rather rough approximations suffice for all results.

Since constant success probability is too low for technical application, it is natural to ask whether there are patching methods that use backtracking to enforce success (if source and target are in the same component). We give an affirmative answer in a very strong sense:

- Every patching method that satisfies three natural criteria has success probability 1 (conditional on s and t being in the same component), and still has a.a.s. stretch $1 + o(1)$.

Crucially, such patching does not require global knowledge of the vertices, so we can at the same time guarantee success and maintain the benefits of local protocols, with negligible additional costs.

1.1 Related Work and Choice of Model

Kleinberg. Kleinberg’s model starts from an $n \times n$ lattice graph, i.e., we embed node $v = (i, j) \in [n] \times [n]$ at position $\mathbf{x}_v = (i/n, j/n)$ in the unit square and connect any two nodes u, v in Manhattan distance $\|\mathbf{x}_u - \mathbf{x}_v\| \leq 1/n$ by an edge. Then we add for each node u a constant number of additional long-range edges, where the other endpoint v is chosen proportional to $\|\mathbf{x}_u - \mathbf{x}_v\|^{-2}$. This model has several shortcomings that limit the applicability to real-world networks and that are also present in subsequent work [26, 46, 25, 45, 16, 37]:

- **Fragile Exponent.** As shown by Kleinberg [36], his model is very fragile with respect to the exponent of the long-range probability distribution, since changing this distribution to $\|\mathbf{x}_u - \mathbf{x}_v\|^{-\alpha \cdot 2}$ for any $\alpha \neq 1$ increases the expected number of steps to $n^{\Omega(1)}$. This is reasonable for $\alpha < 1$, since in this regime most neighbors of most vertices lie very far and thus we cannot expect greedy routing to utilize locality. However, for $\alpha > 1$ there is no intuitive argument ruling out greedy routing in general. In particular, investigations of trajectories of bank notes in main cites of the USA suggest a value of $\alpha = 1.59 \pm 0.02$ [14].

¹We say that an event holds a.a.s. if it holds with probability $1 - o(1)$ with $n \rightarrow \infty$.

²The stretch is the ratio of the lengths of the greedy path and the shortest path between source and target.

- **Perfect Lattice.** Kleinberg’s model is unnatural, as it crucially assumes an underlying perfect lattice substrate to guarantee connectivity. In spirit of the classic random geometric graphs [51], in a more realistic model each vertex would choose a uniformly random position \mathbf{x}_v in the unit square, followed by the same edge sampling procedure as in Kleinberg’s model. However, it can be shown that in this adapted model with high probability greedy routing does not reach the target, essentially since in each step the current vertex has constant probability not to have any neighbor with smaller distance to the target. Thus, the unnatural assumption of a perfect lattice is crucial for Kleinberg’s result.
- **Homogeneity.** Kleinberg’s graphs are homogeneous, i.e., all vertices have the same degree³. In contrast, most real-world networks have been found to be scale-free, meaning that their degree distribution follows a power law $p(d) = d^{-\beta}$, typically with power law exponent β in the range (2, 3) [22, 2] (or at least a closely related distribution [48]).

Geometric Inhomogeneous Random Graphs (GIRGs) We now substantiate claims (1)-(4) on page 2 that explain our reasons for picking the GIRG model.

(1) Let us first describe the GIRG model, which is inspired by the classic Chung-Lu random graphs [17, 18], and compare it with Kleinberg’s model. In a GIRG, each node v is equipped with a random position \mathbf{x}_v in an underlying geometry (for simplicity, we fix the d -dimensional torus) and with a weight w_v which is drawn according to a power law $p(d) = d^{-\beta}$ and which represents the (expected) degree of vertex v and thus its “connectedness”. After drawing the positions \mathbf{x}_v and weights w_v , all edges are picked independently. Consider vertices u, v for which we already sampled the weights w_u, w_v but not yet the positions $\mathbf{x}_u, \mathbf{x}_v$. If the marginal probability of u, v forming an edge is $\Theta(\min\{w_u w_v/n, 1\})$ then the expected degree of vertex v is $\Theta(w_v)$ (as in the classic Chung-Lu random graphs). Reintroducing the positions $\mathbf{x}_u, \mathbf{x}_v$ into this picture, a short calculation shows that if the probability of u, v forming an edge is

$$p_{uv} = \Theta\left(\min\left\{\left(\frac{w_u w_v}{n \cdot \|\mathbf{x}_u - \mathbf{x}_v\|^d}\right)^\alpha, 1\right\}\right),$$

for some $\alpha > 1$ then the marginal probability is indeed $\Theta(\min\{w_u w_v/n, 1\})$ so that the expected degree of node v is $\Theta(w_v)$. Moreover, note that in terms of distances, p_{uv} is proportional to Kleinberg’s distribution $\|\mathbf{x}_u - \mathbf{x}_v\|^{-\alpha d}$, where we want to choose $\alpha > 1$ as discussed above. In this sense, the GIRG model is a natural scale-free variant of Kleinberg’s model that removes the implausible assumption of a perfect grid. As it turns out, greedy routing on this inhomogeneous model is much more robust than in Kleinberg’s model, so the problem of fragile exponents is also overcome.

(2) A large portion of the theoretical random graph literature of the last 10 years (e.g. [1, 34, 3, 23, 8, 24, 20, 13, 12, 15, 30, 28, 35, 32, 21, 9, 10, 11]) deals with the design and analysis of random graphs that have more and more of the experimentally observed properties of real-world networks, e.g., sparsity, scale-freeness, and high clustering. Most of these modern models come equipped with an underlying geometry (since locality naturally yields high clustering and, more generally, local community structure), which makes them candidates for analyzing greedy routing.

Since there are no comparative studies, it is hard to judge which of the proposed models is most realistic. However, GIRGs and their special case, hyperbolic random graphs, stand out, because this model has not only attracted theoretical research but has also been experimentally validated.

³Some variants of Kleinberg’s model have Poisson degree distribution [26], which is still almost homogeneous, as it is highly concentrated around its mean.

On the theory side, it was proven for GIRGs and hyperbolic random graphs that they are sparse and scale-free [30, 24, 13], have constant clustering coefficient [15, 30, 12], a giant component and polylogarithmic diameter [28, 35, 13], and average distance $\frac{2 \pm o(1)}{\lceil \log(\beta-2) \rceil} \log \log n$, where β is the power law exponent [13]. Furthermore, generating a random graph from these models takes time $O(n)$ and these graphs have entropy $O(n)$ [52, 12]. Recently, researchers have started to design algorithms that run fast on hyperbolic random graphs [27, 4]. In terms of practical validation, Boguñá et al. [8] computed a (heuristic) maximum likelihood fit of the internet graph into the hyperbolic random graph model and demonstrated its quality by showing that greedy routing in the underlying geometry of the fit finds near-optimal shortest paths. This successful embedding of a real-world network validates the model experimentally, and was repeated and extended in [50, 5].

(3) In general, it is not clear what a natural generalization of greedy routing is in a model with weights and position, since the ideal neighbor should optimize both objectives, being close to the target, and having large weight. Many possibilities have been suggested to resolve this conflict (e.g., in the physics community [58, 33]). However, following Milgram’s instruction, “mail this folder to a personal acquaintance who is more likely than you to know the target person” [59], there is a natural choice: We should pick the neighbor v maximizing the direct edge probability p_{vt} to the target. Thus, Milgram’s experiment has a natural analogon on GIRGs.

We remark that this type of greedy routing to maximize p_{vt} in GIRGs is equivalent to geometric routing in hyperbolic space when we consider the special case hyperbolic random graphs.

(4) There is extensive experimental work on greedy routing on GIRGs, and our results explain many of their findings. We postpone the discussion to Section 4.

Further Related Work For psychological work related to Milgram’s experiments we refer to the survey by Schnettler [55]. Due to the sheer amount of experimental and non-rigorous work on routing in real-world networks and random graphs it is impossible to give a reasonable overview; we refer to the survey by Huang et al. [33]. Regarding theoretical work, a monograph by Watts summarizes the early work [60]. Let us highlight two more recent theory papers: Watts et al. [61] and Kleinberg [37] studied hierarchical network models that are well motivated by sociological principles. This makes the model somewhat more realistic, however, they suffer from similar shortcomings as Kleinberg’s original model, e.g., both papers assume a perfect b -ary tree structure. Moreover, they do not show ultra-small routing distance.

Structure of the paper. In Section 2 we formally introduce the model, in Section 3 we list and explain our main results, and in Section 4 we discuss related experimental results. In Section 5 we discuss the patching criteria (P1)-(P3) and related literature on patching algorithms, and give two simple examples for patching algorithms which satisfy (P1)-(P3). Section 6 contains a high level discussion of the typical trajectory of greedy routing, and it also contains a sketch of the main ideas of our proofs. Section 7 contains preparations for the proofs. Thereby, in Section 7.1 we introduce the notation needed for the proofs. This section also contains a reference table for the notation used in the various proofs. In Section 8.1 we prove the main Lemma 8.1, which lies at the heart of all subsequent proofs, and finally we give the proofs of our main results in Sections 8 (basic algorithm), 9 (patching), 10 (relaxations), and 11 (hyperbolic random graphs).

2 Random Graph Model & Greedy Routing

2.1 Geometric Inhomogeneous Random Graphs (GIRGs)

In this paper we study greedy routing on Geometric Inhomogeneous Random Graphs (GIRGs). Informally speaking, a GIRG is a graph $G = (V, E)$ where both the set of vertices V and the set of edges E are random. Informally speaking, each vertex v has a random position \mathbf{x}_v in a geometric space \mathbb{T}^d and a random weight $w_v \in \mathbb{R}^+$, where the distribution of the weights follows a power law. For each pair of vertices we flip an independent coin to determine whether they are connected by an edge, where the probability to be connected increases with the weights and decreases with the distance of the vertices. We now give the formal definition of the process, following [40].

Geometric space Let $d \in \mathbb{N}$ be some (constant) parameter. As our geometric space we consider the d -dimensional torus $\mathbb{T}^d = \mathbb{R}^d / \mathbb{Z}^d$, which can be described as the d -dimensional cube $[0, 1]^d$ where opposite boundaries are identified with each other. As distance function we use the ∞ -norm on \mathbb{T}^d , i.e., for $x, y \in \mathbb{T}^d$ we define

$$\|x - y\| := \max_{1 \leq i \leq d} \{\min\{|x_i - y_i|, 1 - |x_i - y_i|\}\}.$$

The choice of the geometric space \mathbb{T}^d is in the spirit of the classical random geometric graphs [51]. We prefer the torus to the hyper-cube for technical simplicity, as it yields symmetry. However, it is not hard to replace \mathbb{T}^d by $[0, 1]^d$. Similarly, we may replace the maximum norm by any other norm and obtain the same model since we allow constant factor deviations in the definitions (EP1) and (EP2) below.

Let the set of vertices V be given by a Poisson point process on \mathbb{T}^d with intensity n , cf. [57] for a formal definition of Poisson point processes.⁴ In particular, the expected number of vertices is n , the positions of all vertices are uniformly at random in \mathbb{T}^d , and for each measurable $A \subset \mathbb{T}^d$, the number of vertices in A is Poisson distributed with mean $n \cdot \text{VOL}(A)$. Moreover, if $A, B \subseteq \mathbb{T}^d$ are disjoint measurable sets, then the number of vertices in A and in B are independent random variables.

Weights Let $2 < \beta < 3$ and $w_{\min} > 0$ be fixed constants. Then for each vertex $v \in V$ independently we draw a weight according to the probability density function

$$f(w) := \begin{cases} \Theta(w_{\min}^{\beta-1} w^{-\beta}), & w \geq w_{\min} \\ 0, & w < w_{\min}. \end{cases}$$

The density f is chosen such that every vertex takes a random weight according to a power-law with parameter β and minimal weight w_{\min} , and the probability that a random vertex has weight at least w is proportional to $w^{1-\beta}$. The choice of the power law parameter ($2 < \beta < 3$) is standard [22].

Edges Given the vertex set V and all positions and weights, we connect vertices $u \neq v$ independently with probability p_{uv} , which depends on the weights w_u, w_v and on the distance $\|\mathbf{x}_u - \mathbf{x}_v\|$. We require for some constant $\alpha > 1$ the following edge probability condition:

$$p_{uv} = \Theta\left(\min\left\{\frac{1}{\|\mathbf{x}_u - \mathbf{x}_v\|^{\alpha d}} \cdot \left(\frac{w_u w_v}{w_{\min} n}\right)^\alpha, 1\right\}\right). \quad (\text{EP1})$$

⁴A very similar model is obtained by placing n points uniformly at random in \mathbb{T}^d , e.g. [12]. We prefer the Poisson point process since then the number of vertices in disjoint regions of \mathbb{T}^d is independent.

In particular, there exists a constant $c_1 > 0$ such that $p_{uv} = \Theta(1)$ if u and v are very close to each other, i.e., $\|x_u - x_v\|^d \leq c_1 \frac{w_u w_v}{w_{\min} n}$.

The parameter α determines how fast the connection probability decays with increasing distance of the nodes. In the threshold case $\alpha = \infty$ the connection probably instantly drops to zero if the vertices are too far apart: instead of (EP1) we require that there are constants $0 < c_1 \leq c_2$ such that

$$p_{uv} = \begin{cases} \Theta(1) & \text{if } \|x_u - x_v\|^d \leq c_1 \frac{w_u w_v}{w_{\min} n} \\ 0 & \text{if } \|x_u - x_v\|^d \geq c_2 \frac{w_u w_v}{w_{\min} n}. \end{cases} \quad (\text{EP2})$$

Note that there can be an interval $[(c_1 \frac{w_u w_v}{w_{\min} n})^{1/d}, (c_2 \frac{w_u w_v}{w_{\min} n})^{1/d}]$ for $\|x_u - x_v\|$ where the behaviour of p_{uv} is arbitrary. In many settings the connection probability is 1 if u and v are sufficiently close to each other.

The formulas (EP1) and (EP2) depend on the parameter w_{\min} in such a way that the expected degree of a vertex of weight w is $\Theta(w)$. Moreover, for two vertices u and v of fixed weights, but with random positions, the connection probability is $\mathbb{E}_{x_u, x_v}[p_{uv} \mid w_u, w_v] = \Theta(\min\{\frac{w_u w_v}{w_{\min} n}, 1\})$. For details see Section 7.

Summarizing, the free parameters of the model are the intensity (i.e., the expected number of vertices) n , minimum weight w_{\min} , the power law parameter $2 < \beta < 3$, the dimension $d \in \mathbb{N}$, the decay parameter $\alpha > 1$, and the actual probability functions p_{uv} which satisfy (EP1) (if $\alpha < \infty$) or (EP2) (if $\alpha = \infty$). In all subsequent theorems and proofs, we assume that all parameters except n are constants. In particular, the O -notation may hide any dependence on the parameters d, α, β and on the hidden constants in (EP1) and (EP2), but not on w_{\min} as Theorem 3.2 will link the failure probability of the routing protocol to w_{\min} .

2.2 Greedy routing

Routing protocol We will evaluate the performance of greedy routing in a GIRG, which is the following process. A message should be sent from a starting vertex s (source) to a target vertex t . The address of every vertex v is the pair (x_v, w_v) . Every vertex has local information, i.e., it knows the address of itself and of its neighbors. In addition, the address of the target is written on the packet. Then the routing proceeds in rounds and in every hop, the packet is sent from the current vertex v to a neighbor u which maximizes a given objective function ϕ , cf. below. When a dead end is reached, the process stops and the packet is dropped. If $u = t$, then we say that the routing was successful. A pseudocode description is given in Algorithm 2.2. Note that we will analyze the process for arbitrary vertices s and t , i.e., we may explicitly choose weights and positions of s and t , while everything else of the graph is still drawn randomly.

Objective function It is natural and in spirit of Milgram's experiment [59] that in every round, the packet should be sent to a neighbor v which maximizes p_{vt} . Note that if $\alpha < \infty$, maximizing p_{vt} is equivalent to maximizing $w_v / (w_{\min} n \|x_v - x_t\|^d)$.⁵ Hence, for our analysis we use the objective function

$$\phi(v) := \frac{w_v}{w_{\min} n \|x_v - x_t\|^d}$$

⁵We keep the normalization factor n since it will turn out to be the most natural normalization. In this way, $\phi(v)$ will always increase by the same exponent in the second phase of the algorithm, cf. Section 6. We use the additional w_{\min} -factor only for technical reasons, as it simplifies several calculations.

and observe that the target vertex t globally maximizes ϕ , which is a necessary condition for any reasonable objective function. In particular, if $\{s, t\} \in E$, then the algorithm will send the packet directly to the target.

Note that nodes in the greedy protocol w.r.t. objective ϕ do not need any global knowledge except for the address of the target, which is available because it is part of the message.

Algorithm 1 Greedy Routing Algorithm

```

1: function GREEDY( $s, m$ ) //  $s$ : source,  $m$ : message (contains target information  $t$ )
2:   if  $s == t$  then deliver message
3:   else
4:      $v \leftarrow \operatorname{argmax}\{\phi(u) \mid u \in \Gamma(s)\}$ 
5:     if  $\phi(v) > \phi(s)$  then GREEDY( $v, m$ )
6:     else return failure

```

3 Results

In this Section we list and explain all our main results. Unless stated otherwise, we assume throughout the paper that s and t are fixed, while the rest of the graph is drawn randomly. That is, an adversary may pick weights and positions of s and t , while the remaining vertices and all edges are drawn randomly as described in Section 2.⁶ Then we consider greedy routing from s to t . Proof sketches and intuitive reasons for the results can be found in Section 6. The formal proofs can be found in Section 8 (basic algorithm), Section 9 (patching), Section 10 (relaxations), and Section 11 (hyperbolic random graphs).

Success Probability. As our first result we show that greedy routing from s to t succeeds with constant probability. The proofs of the following two theorem are in Section 8.2 and 8.3, respectively.

Theorem 3.1. *Greedy routing succeeds with probability $\Omega(1)$.*

In general, we cannot expect anything better than constant connection probability, since we only assumed a connection probability of $\Theta(1)$ in (EP1) and (EP2) even if two vertices are very close to each other. So we may end in a vertex which is extremely close to t , but which fails to actually connect to t . If we assume connection probability 1 in such cases, then we get a better connection probability. More precisely, assume that the constant c_1 in (EP1) or (EP2) is chosen such that the following condition holds.

$$p_{uv} = 1 \quad \text{if} \quad \|x_u - x_v\|^d \leq c_1 \frac{w_u w_v}{w_{\min} n}. \tag{EP3}$$

Then we can strengthen Theorem 3.1 by showing that the failure probability decays exponentially with w_{\min} .

⁶Equivalently, we can first choose the random graph G , then draw s and t at random, and condition the resulting probability space on x_s, x_t, w_s , and w_t .

Theorem 3.2. *If (EP3) holds in addition to the model assumptions, then:*

- (i) *Greedy routing succeeds with probability $1 - O(e^{-w_{\min}^{\Omega(1)}})$.*
- (ii) *If $\min\{w_s, w_t\} = \omega(1)$ then greedy routing succeeds with probability $1 - \min\{w_s, w_t\}^{-\Omega(1)}$. In particular, in this case greedy routing succeeds a.a.s..*

The first part of this theorem is also almost-optimal. The degree of a vertex $v \in V$ is distributed as $\text{POIS}(\Theta(w_v))$, as we will see in Lemma 7.2 below. Therefore a vertex v of weight w_v is isolated with probability $e^{-\Theta(w_v)}$. If for example $w_t = \Theta(w_{\min})$, then the success probability cannot be higher than $1 - e^{-\Theta(w_{\min})}$.

For both statements we remark that we only need (EP3) for the very last step. Even without (EP3), with sufficiently high probability greedy routing finds a vertex u such that u and t satisfy the condition of (EP3). However, without (EP3) the vertex u only has probability $\Theta(1)$ to connect to t directly, and there is a considerable chance that there is no better neighbor of u , i.e., that u is a local optimum.

Stretch. One of the most crucial efficiency measures for routing protocols is the stretch, i.e., the ratio of the routing paths length compared to the length of a shortest path. Clearly, routing can only succeed if the source s and the target t are in the same component of the graph. In [13] it was proven that a.a.s. a GIRG possesses a giant connected component of linear size in which the average distance is $\frac{2+o(1)}{\lceil \log(\beta-2) \rceil} \log \log n$. The next theorem shows that the number of steps of the greedy routing algorithm (either the target is found or the routing stops) is a.a.s. equal to shortest paths, up to a stretch of $1 + o(1)$. The proof is given in Section 8.4.

Theorem 3.3. *A.a.s., greedy routing stops after at most $\frac{2+o(1)}{\lceil \log(\beta-2) \rceil} \log \log n$ steps and either finds t or ends in a dead end. More precisely, for every $f_0(n) = \omega(1)$ a.a.s. the stronger bound*

$$\frac{1+o(1)}{\lceil \log(\beta-2) \rceil} (\log \log_{w_s} \phi(s)^{-1} + \log \log_{w_t} \phi(s)^{-1}) + O(f_0(n)). \quad (1)$$

holds.

Note that Theorems 3.1 and 3.3 together imply that even if we condition on greedy routing being successful, then still a.a.s. greedy routing needs at most $\frac{2+o(1)}{\lceil \log(\beta-2) \rceil} \log \log n$ steps. Since this agrees with the average distance in the giant, it is optimal. In particular, for two vertices s, t with random weights and positions, if greedy routing between s and t is successful then a.a.s. the stretch is $1 + o(1)$.⁷

In the most typical situation (e.g., if s and t are chosen randomly), (1) evaluates to the simpler expression $\frac{2+o(1)}{\lceil \log(\beta-2) \rceil} \log \log n$ in Theorem 3.3. For example, this is the case if s and t have constant weights and are far apart from each other, $w_s, w_t = O(1)$, and $\|x_s - x_t\| = \Omega(1)$. However, if s and t have large weights or are close to each other, then (1) gives a much better bound (but never a worse bound, since $\phi(s) = \Omega(1/n)$).

⁷Technically speaking, here we use the slightly stronger statement that a.a.s. a $1 - o(1)$ fraction of all pairs of vertices in the giant have distance $\frac{2+o(1)}{\lceil \log(\beta-2) \rceil} \log \log n$. This statement is not explicitly formulated as a lemma in [13], but is in the proof of Theorem 5.9 in [13].

Patching. As outlined in the introduction, greedy routing may have practical applications. For these applications, it is not acceptable that the algorithm simply fails if it enters a local maximum. Therefore, several patching methods have been proposed. We show that any patching algorithm A is efficient if it satisfies the following three basic conditions. For further discussions, in particular of the conditions, see Section 5.

- (P1) (Greedy choices.) If A is in some vertex v and decides to visit an unexplored neighbor then it must choose the unexplored neighbor of largest objective. If A visits a vertex v for the first time, and v has a neighbor of larger objective, then A proceeds to the neighbor of v with largest objective.
- (P2) (Poly-time exploration.) If A has explored k vertices, and at least one of these vertices has an unexplored neighbor, then A visits an unexplored vertex in $k^{O(1)}$ steps.
- (P3) (Poly-time exhaustive search.) Assume that A visits a vertex v that has larger objective than all previously visited vertices, and let S be the connected component of v in $G[V_{\geq \phi(v)}]$, i.e., S is the set of all vertices that can be reached from v without touching vertices of worse objective than v . Then A visits the vertices in S in the next $|S|^{O(1)}$ steps (or finds t in one of these steps).
If $\phi(s) = \Omega(1)$ then there exists a sufficiently slowly falling $\phi_0 = o(1)$ such that A explores the connected component S of s in $G[V_{\geq \phi_0}]$ in $|S|^{O(1)}$ steps.

Condition (P2) ensures that greedy routing is always successful if source s and target t are in the same component. The surprising result is that (P1)-(P3) are already sufficient to ensure efficient greedy routing, i.e., a.a.s. the stretch is $1 + o(1)$. The proof can be found in Section 9

Theorem 3.4. *Assume that A is a routing algorithm that satisfies (P1)-(P3). If s and t are in the same component⁸ then A always successfully routes from s to t , and a.a.s. this takes at most $\frac{2+o(1)}{|\log(\beta-2)|} \log \log n$ steps.⁹*

Relaxations. So far we assumed that the objective ϕ can be computed exactly, which may be unrealistic in practice. For example, in the Milgram experiment the participants can only estimate the objective values of their neighbors. So we also study an approximate version of greedy routing, where the greedy algorithm uses an objective function $\tilde{\phi}$ which is only an approximation of ϕ . We allow that $\tilde{\phi}$ deviates from ϕ by a bounded constant factor, or even more. This adds quite some flexibility: for any of the best $\min\{w_v, \phi(v)^{-1}\}^{o(1)}$ neighbors of a vertex v , there is a “good enough” approximation $\tilde{\phi}$ which puts this vertex to the top. In particular, it is not necessary to compute the optimal neighbor w.r.t. ϕ , which may be costly. Moreover, the routing process is also robust: for example, it is no problem if some of the edges fail during execution of the routing, since the current vertex can send the message to any other good neighbor instead.

Theorem 3.5. *Let $\tilde{\phi} : V \rightarrow \mathbb{R}$ be a function which is maximized with t such that*

$$\tilde{\phi}(v) = \Theta \left(\phi(v) \cdot \min\{w_v, \phi(v)^{-1}\}^{o(1)} \right), \quad (2)$$

holds for all $v \in V$. Then Theorems 3.1, 3.2, 3.3, and 3.4 also hold for greedy routing with respect to $\tilde{\phi}$ instead of ϕ .

⁸i.e., the a.a.s. statement holds conditioned on the event that s and t are both in the same component.

⁹In fact, a.a.s. the precise number of steps is again given by expression (1). For the sake of readability, we omit the proof of the more precise bound and restrict ourselves to the bound as claimed in the theorem.

Moreover, let $\delta > 0$ be a small constant. Then for Theorems 3.1, 3.2, and 3.3 the previous statement is still true if for vertices $v \in V$ with $\phi(v) \geq O(w_t^{-1+\delta})$ we relax Condition (2) to the weaker condition

$$\tilde{\phi}(v) = \Omega\left(w_t^{-1+\delta/2}\right). \quad (3)$$

Finally, for transferring Theorem 3.4 this weaker condition is also sufficient if t is a random vertex.

The proof of Theorem 3.5 can be found in Section 10.

A consequence of Theorem 3.5 is that the results of this paper also hold for variations of the model. For example, most of the experimental work is performed in slightly different models, e.g., they consider geometric routing on hyperbolic random graphs (see Section 4 for details). In [12] it was proven that one-dimensional GIRGs ($d = 1$) contain the model of hyperbolic random graphs as a special case, but still geometric routing in hyperbolic space is not exactly the same as greedy routing in our sense. Theorem 3.5 shows that the differences are negligible. Formally, geometric routing on hyperbolic random graphs induces an objective function ϕ_H on the corresponding GIRG, which falls into the class of objective functions considered in Theorem 3.5. The proof and further background on hyperbolic random graphs is contained in Section 11.

Corollary 3.6. *All theorems about success probability and path length apply as well for geometric routing on hyperbolic random graphs. Moreover, if t is a random vertex, then Theorem 3.4 for patching algorithms also holds for hyperbolic random graphs.*

4 Comparison with Experimental Results

There are several recent papers which experimentally study 'greedy processes' in graphs with hidden metric spaces [6, 7, 8, 19, 42, 49]. In these papers, either GIRGs are sampled for various parameter combinations, sometimes under different names; or given real networks (e.g., the internet graph) are embedded into the space $\mathbb{T}^d \times \mathcal{W}$, where \mathcal{W} is the space of weights (mostly in the special case $d = 1$, where $\mathbb{T}^d \times \mathcal{W}$ can be interpreted as hyperbolic disc [12, Theorem 6.3], cf. Corollary 3.6). Afterwards, in both settings, the authors experimentally study how greedy routing algorithms perform on these graphs. Thus our paper gives theoretical explanations for the effects that were observed in these papers, which we discuss in detail below.

In contrast, [6] and [7] study geometric greedy processes, i.e., the routing is degree-agnostic and only uses the distance in an homogeneous and isotropic space. It turns out that this geometric type of routing still works in some settings, but is far less efficient and robust (e.g., it completely fails for some values of $\beta \in [2, 3]$). This suggests that greedy routing as considered in this paper is superior to geometric routing. It remains an intriguing open task to understand theoretically how geometric routing behaves for realistic network models.

Success Probability: As the main contributions the aforementioned papers observe a large success probability of greedy routing for most studied combinations of parameters. This is explained by Theorem 3.1, which gives a general lower bound $\Omega(1)$ on the success probability. Often, the experimentally observed success probability is surprisingly large (i.e., larger than 97% in [8]). We believe that we can give the reason for this observation with Theorem 3.2 (i), where we show that the failure probability drops exponentially with the parameter w_{\min} controlling the minimum expected degree. Thus, even with very moderate values of w_{\min} it is not surprising to obtain large success probabilities.

Trajectory of a Greedy Path: In their experiments the authors in [8, 42, 43, 49] observed that a successful greedy path starting at a small degree node will first visit nodes of larger degrees and continue along this path until it reaches the core of the network; afterwards, the greedy path leaves the core towards nodes with smaller degrees but with much closer geometric distance to the target until it reaches its destination. We prove this observation formally. In our proofs, we partition our network into layers, and show that the greedy path follows closely this layer structure. In the first half of the process, the layers are characterized by increasing weights, in the second by decreasing geometric distances to the target. We prove that a.a.s. the greedy algorithm visits each layer at most once, and that it visits a $(1 - o(1))$ -fraction of all layers. Thus we characterize accurately the trajectory that greedy routing takes.

Stretch: Besides the success probability, the stretch of a greedy path, i.e., the quotient of a greedy paths length vs. the length of a shortest path, is one of the most meaningful measures to evaluate the usefulness of greedy routing in applications. Experiments revealed a surprisingly small stretch close to 1, both for embedded real graphs ([8]) and for sampled random graphs [42, 43, 49]. In Theorem 3.3 we give the theoretical explanation and prove formally that successful greedy routing has a stretch of $1 + o(1)$ in the case of success.

Patching: Patching has been experimentally found to be highly efficient, and to maintain an average stretch close to 1 [19, 33, 42, 43, 49]. However, most experiments in sampled models have been performed with a small number of vertices (e.g., $n = 1000$). Also, not all methods guarantee successful routing if s and t are in the same component. Interestingly, the gravity-pressure method considered in [19] shows good results, although it does not fall into our class of patching algorithms, and it suffers from the problems described in Section 5. So it remains open whether this algorithm is also theoretically efficient, or whether the good results were simulation artefacts (see also [54]).

As Theorem 3.4 shows, a broad class of natural patching algorithms are highly efficient, even if they boost the success probability to 100% for vertices in the same component. Most notably, we prove in Theorem 3.4 that a stretch of $1 + o(1)$ is not only achievable, but that it is guaranteed under the weak and natural assumptions (P1)-(P3).

5 Patching Algorithms: Discussion and Examples

The naive greedy protocol gets stuck if it enters a local optimum, which is not acceptable for technical applications. Several different patching strategies have been suggested to overcome this issue, and experimental results showed that many of them work quite well in practice [19, 33, 42, 43, 49, 54]. We formally prove that in fact any patching strategy works well if it satisfies some basic conditions, namely (P1) greedy choices, (P2) poly-time exploration, and (P3) poly-time exhaustive search (cf. page 10).

Note that all three conditions are natural: the first one is that the algorithm makes greedy choices whenever there is an obvious decision to be made. Note that still there are many choices left to the algorithm. For example, if it enters a local optimum, then it is free to decide if and how much it backtracks before continuing exploration. The second condition ensures that A never gets stuck (or only for a polynomial time). Finally, the third condition (P3) forces the algorithm to search exhaustively a connected set of “good” candidates in reasonable time before turning to less promising candidates. The second part of (P3) is concerned with the exceptional case $\phi(s) = \Omega(1)$ (for random s and t this only happens with probability $O(1/n)$) and is necessary to avoid trivialities: if $\phi(s) = \Omega(1)$ then there may be no (or only few) vertices of better objective

than s , and the precondition of the first part of (P3) may never be fulfilled. We remark that the requirement that ϕ_0 is “sufficiently slowly falling” can be specified to $\phi_0 = (\log \log n)^{-o(1)}$.

The last condition (P3) does exclude some algorithms like the gravity-pressure algorithm introduced in [49]. This algorithm always visits the best neighbor v' of the current search point v , even if $\phi(v') < \phi(v)$, and thus does not satisfy (P3). E.g., assume that the algorithm visits two consecutive vertices v_1, v_2 of increasing objective, and assume further that v_2 has exactly one more neighbor u of very bad objective ($\phi(u) \ll \phi(v_1)$), while v_1 has other neighbors of better objective. Then continuing the search from u before exploring the other neighbors of v_1 may be a bad strategy, and this is what (P3) rules out. For example, the gravity-pressure algorithm will prefer to visit any unexplored vertex over returning to v_1 . So if v_1 lies on the only path to t (which happens with probability $\Omega(1)$), then the algorithm may potentially explore large parts of the giant before returning to v_1 , and thus before finding t . This may explain that the algorithm is especially vulnerable in sparse networks [54].

Generally, there are two relevant types of algorithms if the global structure is unknown to the vertices: either information about the routing history is stored in the message (e.g., the protocol SMTP for emails [53]); or for every message a small amount of information is stored in each vertex, yielding rather exploration than classical routing (e.g., flooding algorithms like [31, 39], the gravity pressure algorithm [49], or tree-based approaches [54]). In both cases, it is very easy to design a routing algorithm that satisfies conditions (P1)-(P3), as we outline in the following. For simplicity, we assume that $\phi(s) = o(1)$.

For the first case, we may simply store the list of visited vertices in the message, and for each vertex v we additionally store the objective of the best unexplored incident edge in the message (i.e., the objective of the best neighbor u of v for which the algorithm did not traverse the edge uv). Compared to SMTP this only increases the required memory by a single value per visited node. With this information, a trivial way to satisfy conditions (P1)-(P3) is to use the greedy algorithm if possible (i.e., if we are not in a local optimum), and otherwise explore the best unexplored edge that goes out from any visited vertex.

For the second case, a variant of depth-first search satisfies condition (P1)-(P3), in which the message and each visited vertex only need to store a constant number of pointers and objective values. More precisely, for a vertex $v \in V$ and a value $\Phi \in \mathbb{R}^+$ consider a greedy depth first search on the subgraph $G[V_{\geq \Phi}]$ of vertices which have objective at least Φ , starting in v . “Greedy” here means that if there are several unexplored edges going out from a vertex, then the DFS algorithm picks the edge that leads to the vertex of highest objective. We call this algorithm (greedy) Φ -DFS for short. Note that since $G[V_{\geq \Phi}]$ does not need to be connected, the Φ -DFS does not necessarily visit all vertices in $G[V_{\geq \Phi}]$.

The idea is now the following. Whenever we encounter a vertex v which has strictly larger objective than all previously visited vertices, then we start a $\phi(v)$ -DFS at v . We do this recursively, so when in this $\phi(v)$ -DFS we encounter another vertex v' which has larger objective than all previously visited vertices (including the vertices visited during the $\phi(v)$ -DFS), then we pause the $\phi(v)$ -DFS, and start a $\phi(v')$ -DFS in v' . It may happen that the $\phi(v')$ -DFS is completed without success, i.e., that we return to v' after recursively exploring all its better neighbors without finding t . In this case, we simply discard the $\phi(v')$ -DFS and resume the paused $\phi(v)$ -DFS. Note that we treat all vertices visited during the $\phi(v')$ -DFS as unvisited for the resumed $\phi(v)$ -DFS. However, we do store the best objective value that we have ever seen (regardless in which DFS we have seen them) in the message. A pseudocode description is given in Algorithm 2. The algorithm is distributed, i.e, there is no shared global memory, and no global stack of function calls is required to execute the algorithm. Moreover, at each time only one vertex is active, and each vertex only needs to know positions and weights of its direct neighbors.

Algorithm 2 Example of a distributed exploration algorithm that satisfies (P1)-(P3) and that uses only a constant number of pointers and variables in the message and in each node. Whenever EXPLORE or BACKTRACK_TO are called, the parent process may terminate. For readability, we did not spell out updates of $m.last_visited_vertex$ except the special case in line 26. Furthermore we assumed that no vertex has two neighbors of equal objective.

```

1: function ROUTING( $s, m$ )                                     //  $s$ : source,  $m$ : message (contains target information)
2:    $m.best\_seen\_objective \leftarrow -\infty$                  // initialisation start
3:    $m.Phi \leftarrow -\infty$ 
4:    $m.last\_visited\_vertex \leftarrow s$ 
5:    $s.Phi \leftarrow \phi(s)$ 
6:   EXPLORE( $s, m$ )
7: function EXPLORE( $v, m$ )                                     // we visit  $v$  by an unexplored edge in current  $\Phi$ -DFS
8:   if  $v.Phi == m.Phi$  then                                  // if  $v$  already visited in current  $\Phi$ -DFS ...
9:     BACKTRACK_TO( $m.last\_visited\_vertex, m$ )                // ... then we immediately return
10:  else                                                       // if  $v$  not visited in current  $\Phi$ -DFS
11:    if  $\phi(v) > m.best\_seen\_objective$  then                // potentially start new DFS with  $\Phi = \phi(v)$ 
12:      SET_NEW_PHI( $v, m$ )
13:      INIT_VERTEX( $v, m$ )
14:      if  $\{u \in \Gamma(v) \mid \phi(u) \geq m.Phi\} \neq \emptyset$  then // if there is neighbor better than  $\Phi$ 
15:         $u \leftarrow \operatorname{argmax}\{\phi(u) \mid u \in \Gamma(v)\}$  // go to best neighbor
16:        EXPLORE( $u, m$ )
17:      else BACKTRACK_TO( $m.last\_visited\_vertex, m$ )           // no unexplored neighbors better than  $\Phi$ 
18: function BACKTRACK_TO( $v, m$ )                                 // we backtrack from  $m.last\_visited\_vertex$  to  $v$ 
19:    $S := \{u \in \Gamma(v) \mid u \neq v.parent, m.Phi \leq \phi(u) < \phi(m.last\_visited\_vertex)\}$ 
20:   if  $S \neq \emptyset$  then                                    // if there is still unexplored child better than  $\Phi$  ...
21:      $u \leftarrow \operatorname{argmax}\{\phi(u) \mid u \in S\}$  // ... go to best of them
22:     EXPLORE( $u, m$ )
23:   else                                                       // if no child left to explore
24:     if  $v.started\_new\_dfs == \text{TRUE}$  then                 // if DFS for current  $\Phi$  was unsuccessful ...
25:       RESET_TO_OLD_PHI( $v, m$ )                               // ... we resume  $\Phi$ -DFS for previous value of  $\Phi$  ...
26:        $m.last\_visited\_vertex \leftarrow v.parent$            // ... in the state where we left it, coming from  $v.parent$ 
27:       EXPLORE( $v, m$ )
28:     else                                                       // if current  $\Phi$ -DFS continues, but no child left to explore ...
29:       BACKTRACK_TO( $v.parent, m$ )                             // ... backtrack to parent
30: function SET_NEW_PHI( $v, m$ )                                  //  $v$  starts a new DFS with  $\Phi = \phi(v)$ .
31:    $m.best\_seen\_objective \leftarrow \phi(v)$ 
32:   if  $\{u \sim v \mid \phi(u) \geq \phi(v)\} \neq \emptyset$  then // we only start the new DFS if there are actually better neighbors
33:      $v.started\_new\_dfs \leftarrow \text{TRUE}$ 
34:      $v.previous\_Phi \leftarrow m.Phi$                        // store this in case we need to resume previous  $\Phi$ -DFS
35:      $m.Phi \leftarrow \phi(v)$ 
36: function RESET_TO_OLD_PHI( $v, m$ )                             // to resume  $\Phi$ -DFS for previous value of  $\Phi$ 
37:    $v.started\_new\_dfs \leftarrow \text{FALSE}$ 
38:    $m.Phi \leftarrow v.previous\_Phi$ 
39:    $v.Phi \leftarrow v.previous\_Phi$ 
40: function INIT_VERTEX( $v, m$ )                                  // we visit  $v$  for the first time in the current  $\Phi$ -DFS
41:    $v.Phi \leftarrow m.Phi$                                     //  $v$  has been visited in current  $\Phi$ -DFS
42:    $v.parent \leftarrow m.last\_visited\_vertex$                 // parent for backtracking

```

Let us argue that this algorithm only needs to store a constant number of pointers and objective values in the message and the visited vertices. During a Φ -DFS, we store at each visited vertex the identity of the parent vertex (to allow backtracking) together with the value Φ , and we call this pair the Φ -information. In the message we keep track of the current value of Φ , the identity of the last visited vertex, and the best seen objective. Note that this information already allows us to perform the Φ -DFS in a greedy depth-first manner, and also to decide whether we want to start a new DFS. If we start a new DFS in a vertex v , we need to store in v all information that is necessary to resume the previous DFS: the parent vertex, the previous value of Φ , and a flag indicating that we started a new DFS in v . So for each value of Φ , the Φ -DFS requires a constant memory in each vertex. It only remains to show that no vertex ever needs to store the Φ -information for two different values of Φ simultaneously. So assume we enter a vertex during a Φ -DFS, and it already contains the Φ' -information for some $\Phi' > \Phi$. Then we started the Φ' -DFS after the Φ -DFS, and the only way to resume the Φ -exploration is that the Φ' -DFS terminated without finding t . So in this case we can safely delete the Φ' -information. On the other hand, if we run a Φ -DFS then we claim that we can never encounter a vertex v which carries a Φ' -information for some $\Phi' < \Phi$. Indeed, this could only happen if $\phi(v) \geq \Phi$ (to visit v during the Φ -DFS), and if moreover we would have visited v during the Φ' -DFS, i.e., before the start of the Φ -DFS. But if such a vertex would exist then we wouldn't have started the Φ -DFS in the first place, since Φ wouldn't have been the best encountered objective at this point. So indeed we can not visit a vertex v which carries some Φ' -information for $\Phi' < \Phi$. Summarizing, we never need to store the Φ -information for more than one value of Φ in the same vertex. A pseudocode description is given in Algorithm 2. To enhance readability, we assumed that no vertex has two neighbors of the same objective value. Note that the algorithm is distributed, i.e, there is no shared global memory, and no global stack of function calls is required to execute the algorithm. Moreover, at each time only one vertex is active, and each vertex only needs to know positions and weights of its direct neighbors.

Finally we argue that the algorithm satisfies (P1)-(P3). It is evident that the algorithm satisfies conditions (P1) and (P2). Moreover, it also satisfies (P3): if $x \in V$ and $S \subseteq V$ is the set of all vertices of objective at least $\phi(v)$, then a $\phi(v)$ -DFS takes at most $O(|E(S, S)|) = O(S^2)$ steps. For every vertex $u \in S$, we may interrupt the $\phi(v)$ -DFS for another $\phi(u)$ -DFS, but each of these also takes at most $O(S^2)$ steps. So in total we will explore the set S in at most $O(S^3)$ steps.

6 Proof Sketches and the Typical Trajectory of a Greedy Path

In this section we describe the evolution of the basic greedy process and give a very rough outline the proof ideas. At the same time, we describe the trajectory of a typical greedy path, which is also depicted in Figure 1 below. Detailed proofs can be found in the subsequent sections.

For simplicity, assume that both the source s and the target t have constant weight and assume that $\|x_s - x_t\| = \Omega(1)$, which is the typical case. As s has small weight, a.a.s. all its neighbors (if any) have relatively small weight as well and are located in a ball of small geometric distance around x_s , which we call “region of influence” for illustrative purposes. Since the region of influence of s is small and far away from t , all points in this region have the same distance from t , up to factors $(1+o(1))$. In particular, all neighbors of s have the same distance from t , up to factors $(1+o(1))$. On the other hand, the weights of the neighbors fluctuate by non-negligible constant factors. (So do the distances from s , but they do not influence the objective values.) So the fluctuations of the weights dominate the fluctuations of the distances from t , and consequently, routing proceeds to a neighbor of s with higher weight. More precisely, we can expect that s has a neighbor of weight roughly $w_s^{1/(\beta-2)}$. We repeat the argument and observe that the routing exhibits a first phase in which the

weight of the current vertex increases by an exponent $\approx \frac{1}{\beta-2} > 1$ within every step. This phase stops when a vertex v with the property $\phi(v) \gtrsim \mathbf{w}_v^{-1/(\beta-2)}$ is reached. Since the weight increases by an exponent of $1/(\beta-2)$ with every hop, this first phase needs only $\approx \log \log n / |\log(\beta-2)|$ steps.

Once the routing reaches a vertex v such that $\phi(v) \gtrsim \mathbf{w}_v^{-1/(\beta-2)}$, the relation between the current weight and the current objective change, since now the region of influence of v contains t . That does not mean that v is adjacent to t (in fact, this is very unlikely), but v does have neighbors which are much closer to t . These neighbors typically have smaller weight than v , but the gain in distance is enough to make up for the smaller weight. In this way we can hope for a neighbor of objective $\approx \phi(v)^{\beta-2}$. (Note that $\phi(v) < 1$, so indeed the objective increases.) Moreover, the best neighbor still satisfies the relation $\phi(v) \gtrsim \mathbf{w}_v^{-1/(\beta-2)}$, so we can repeat the argument. Thus in this second phase, the objective increases by the same exponent in every step. After $\approx \log \log n / |\log(\beta-2)|$ rounds, this second phase reaches a vertex v of objective $\phi(v) = \Omega(1)$. Then, in the case $\alpha < \infty$ we have $p_{vt} \geq (\phi(v)\mathbf{w}_t)^\alpha = \Omega(1)$, so with probability $\Omega(1)$ vertex v connects directly to the target vertex t . If $\alpha = \infty$, we need a small number of additional steps to find t . This already gives an idea why greedy routing succeeds with probability $\Omega(1)$ (Theorem 3.1), and why it only needs $\frac{2+o(1)}{|\log(\beta-2)|} \log \log n$ steps (Theorem 3.3).

We thus obtain a typical trajectory of the final greedy path. In the formal proofs of our statements we show that with sufficiently high probability, the routing does not deviate too much from this typical trajectory. For a single vertex v it is not difficult to calculate the weight \mathbf{w}_u and the objective $\phi(u)$ of its best neighbor u . However, there is a technical difficulty: the greedy algorithm only proceeds to u when u is the best neighbor of v . So if we want to compute the probability that the best neighbor of u is again what we expect, then we need to compute a conditional probability, namely conditioned on v having no better neighbor. Unfortunately, this introduces dependencies and makes it much harder to estimate the probabilities.

We overcome this problem by defining layers A_i for the first phase (defined via weights) and the second phase (defined via objectives) such that the typical trajectory contains at most one vertex per layer. Then we consider the greedy path P_i induced by the first i layers and prove that with sufficiently high probability either P_i has no vertex in A_i or the first vertex $v \in A_i \cap P_i$ has a neighbor u outside the first i layers with better objective than all neighbors inside the first i layers (which we already have uncovered). Note that this is an event where we do not suffer from dependencies: the vertex v is defined only in terms of the first i layers, so we can determine v without uncovering any parts of the remaining graph. Consequently, the number of vertices of v in larger layers is Poisson distributed, and it suffices to show that this number is large. Nevertheless, it is not hard to see that the above events together imply that greedy routing succeeds throughout these layers, and that it never visits a layer twice. So we only need to apply a union bound over all layers, and obtain that with sufficiently high probability the routing algorithm succeeds with probability $\Omega(1)$ and visits at most one vertex per layer.¹⁰ This idea is formalized in our main lemma 8.1.

A more careful analysis along the same lines gives the relaxation result, Theorem 3.5. Moreover, it turns out that the failure probability of greedy routing is dominated by the first few steps (when the weight is still constant) and the last few steps (when the objective is constant). In each of these steps the algorithm has probability $\Omega(1)$ to fail. For example, in the typical case that s has a constant weight, there is a constant probability that s has no neighbors at all. However, the number of neighbors of s is Poisson distributed with mean $\Theta(\mathbf{w}_s)$. In particular, the probability to have no neighbors decays exponentially with \mathbf{w}_s . Similar considerations apply for the number of

¹⁰We do not prove that we visit every layer. However, since a.a.s. the stretch is $1 + o(1)$, we visit a.a.s. a $(1 - o(1))$ -fraction of all layers.




Figure 1: A typical greedy path. In the main lemma 8.1 we prove that if for some $\phi = o(1)$ we uncover the graph $G[V_{\leq \phi}]$ then the trajectory looks as depicted if the process gets started at all. Moreover in the last layer we find a vertex which has in expectation (over the randomness in $V_{> \phi}$) still many neighbors in the next layer. Note that the greedy path in G does not need to coincide with the greedy path on $G[V_{\leq \phi}]$, but it can only deviate by jumping to a vertex of objective larger than ϕ . The special layer $A_{1,\infty}$ will be needed only for technical reasons.

neighbors with better objective. This is the reason why the failure probability decays exponentially with w_{\min} (Theorem 3.2).

Finally, to prove the patching result (Theorem 3.4) we show three intermediate results. (i) If we have explored k vertices, starting from s , then with probability $1 - \exp\{-k^{\Omega(1)}\}$ at least one of these vertices is adjacent to a vertex of weight at least $k^{\Omega(1)}$. By condition (P2) it takes only $k^{O(1)}$ steps to explore k vertices, so after a short exploration phase of $o(\log \log n)$ steps, a.a.s. we find a vertex of weight $\omega(1)$. (ii) As in the purely greedy case, starting from a vertex of weight $\omega(1)$ a.a.s. we follow a typical trajectory of the greedy algorithm as described above, until we find a vertex v with almost constant objective (say, with objective $\phi(v) \approx (\log \log \log n)^{-1}$). This middle phase is purely greedy and does not require patching. In particular, v has better objective than any previously visited vertices. (iii) We study the graph $G_{\geq \phi(v)}$ induced by vertices of objective larger

than v and find that a.a.s. $G_{\geq\phi(v)}$ contains at most $O(\phi(v)^{-1})$ vertices and a giant component, that contains both the vertices t and v . Thus by condition (P3) the algorithm will explore the giant component of $G_{\geq\phi(v)}$ and find t in additional $o(\log \log n)$ steps.

7 Notation & Basic Properties of GIRGs

Before proving our main results, we make several preparations in this chapter. First we introduce some basic notation. Then, in Section 7.2 we give some properties of the GIRG model which hold in general and have no direct connection to greedy routing. Finally, Section 7.3 contains lemmas which are more directly related to greedy routing. In particular, there we calculate for a given vertex $v \in V$ where we expect to find its neighbors with highest objective.

7.1 Notation

Throughout the paper $G = (V, E)$ always denotes an (undirected) random graph given by the GIRG model with vertex set V and edge set E . For a vertex v , we denote its weight and position by w_v and x_v , respectively, and $\Gamma(v)$ is the set of neighbors of v in G . For $A \subseteq V$, we denote by $G[A]$ the subgraph of G induced by A . Moreover, for $\phi \in \mathbb{R}$ we denote by $V_{\leq\phi}$ and $V_{>\phi}$ the set of vertices with objective at most ϕ and larger than ϕ , respectively.

We remind the reader that the Landau notation O, Ω , etc. may hide the power law exponent β , the decay parameter α , the dimension d , and the constants in (EP1) and (EP2), but not the minimal weight w_{\min} . The Landau notation is always used with respect to $n \rightarrow \infty$. We say that an event $\mathcal{E} = \mathcal{E}(n)$ holds asymptotically almost surely (a.a.s.) if $\Pr[\mathcal{E}(n)] \rightarrow 1$ for $n \rightarrow \infty$. We refer to Table 1 on page 19 for general notation used throughout the paper.

7.2 Properties of GIRGs

In this short section we collect some basic properties of GIRGs. We start with some basic lemmas taken from [13].¹¹ First we give a formula for the marginal probability for an edge between two vertices u and v with given weights.

Lemma 7.1 ([13, Lemma 4.3]). *Let u, v be two vertices of a GIRG. Then*

$$\Pr_{x_u, x_v} [\{u, v\} \in E \mid w_u, w_v] = \Theta \left(\min \left\{ \frac{w_u w_v}{w_{\min} n}, 1 \right\} \right).$$

Note that this marginal probability is the same as for Chung-Lu graphs. This is why GIRGs can be interpreted as a geometric variant of Chung-Lu random graphs. Furthermore, the above expression does not change if we fix one of the two coordinates (but not both!). From Lemma 7.1, we can directly deduce the following statement about the degree of a given vertex.

Lemma 7.2. *Let $v \in V$ be a fixed vertex with weight w_v . Then the random variable $\deg(v)$ is distributed as $\text{POIS}(\Theta(w_v))$, and in particular $\mathbb{E}[\deg(v)] = \Theta(w_v)$.*

¹¹The model in [13] slightly differs from our version (apart from being more general, [13, Theorem 7.3]) because the vertices are not given by a Poisson point process, but rather n vertices are placed uniformly at random in \mathbb{T}^d . However, for any two given vertices u and v , the distribution of weights and positions and the connection probabilities are the same. Moreover, the models coincide if we condition our model on the number of vertices. In particular, if a statement holds a.a.s. in the model in [13] then it also holds a.a.s. in our model.

Notation	Definition	Defined on page
General		
$G = (V, E)$	random graph from the GIRG model with vertex set V and edge set E	6
$\Gamma(v)$	set of neighbors of v in G	18
$G[S]$	subgraph of G induced by $S \subseteq V$	18
$\deg(v)$	degree of vertex v in G	18
$\text{VOL}(A)$	volume of $A \subseteq \mathbb{T}^d$	6
a.a.s.	“asymptotically almost surely”, probability $\rightarrow 1$ as $n \rightarrow \infty$	18
$\text{POIS}(n)$	Poisson distribution with parameter n	6
Model		
n	expected number of vertices	6
d	dimension of the geometric space	6
β	$2 < \beta < 3$, power-law parameter	6
w_{\min}	minimal weight of a vertex	6
α	decay parameter, which balances between vertex degrees and the geographical distance in p_{uv}	6
p_{uv}	edge connection probability between u and v	6
\mathbb{T}^d	considered geometric space: d dim. torus $(\mathbb{R}^d/\mathbb{Z}^d)$	6
x_v	position of v in the geometric space \mathbb{T}^d	6
f	probability density function of the weights	6
w_v	weight of v , $w_v \in \mathbb{R}$	6
c_1, c_2	constants in the connection probability p_{uv}	6
Routing		
s, t	starting and destination vertex of the routing	6
m	message to be routed	8
ϕ	objective function	7
$V_{\leq \phi}, (V_{> \phi})$	set of vertices with objective $\leq \phi, (> \phi)$	18
$\tilde{\phi}$	approximative objective function	10
Proofs		
$\gamma(\varepsilon)$	$\frac{1-\varepsilon}{\beta-2}$, a constant	22
ε_1	$\varepsilon_1(\alpha, \beta) > 0$, a fixed global constant	22
V_1	vertex set corresponding to the first phase of routing, $\{v \in V : \phi(v) \leq w_v^{-\gamma(\varepsilon_1)}\}$	22
V_2	vertex set corresponding to the second phase of routing, $\{v \in V : \phi(v) \geq w_v^{-\gamma(\varepsilon_1)}\}$	22
ζ	$\max\{\frac{3}{2}, \frac{2\alpha-1}{2\alpha+4-2\beta}\}$ for $\alpha < \infty$, $\frac{3}{2}$ for $\alpha = \infty$	22
$V^+(v, \varepsilon)$	set of good vertices: large weight and larger objective	22
$V^-(v, \varepsilon)$	set of bad vertices: small weight and large objective	22
$w_1(\varepsilon)$	$O(e^{d/\varepsilon})$	30
$\phi_1(\varepsilon)$	$\Omega(e^{-d/\varepsilon})$	30
w -grid	grid splitting \mathbb{T}^d into $\frac{n}{w}$ equal and disjoint cubes	21
$\bar{V}(w, \phi)$	$\{v \in V_1 \mid w_v \geq w \wedge \phi(v) \leq \phi\} \cup \{v \in V_2 \mid \phi(v) \leq \phi\}$	30

Table 1: A general notation table

Proof. The fact that $\deg(v)$ is a Poisson distributed random variable follows immediately since the vertices are given by a Poisson point process. The formula $\mathbb{E}[\deg(v)] = \Theta(w_v)$ was shown in [13, Lemma 4.4], and the expectation of this model agrees with the expectation in our model, cf. Footnote 11. \square

The next lemma asserts that there is a unique giant component (i.e., a component of linear size). Moreover, the average distance in the giant component is the same as the number of steps that greedy routing takes (Theorem 3.3), up to a stretch factor of $1 + o(1)$.

Lemma 7.3 ([13, Theorem 2.2, 2.3]). *A.a.s. the largest component has linear size, while all other components have size at most $\log^{O(1)} n$. Moreover, a.a.s. the average distance between vertices in the largest component is $\frac{2+o(1)}{|\log(\beta-2)|} \log \log n$.*

For a fixed vertex v , the following lemma shows that it is unlikely for v to have a neighbor of much larger weight. We will need this auxiliary result only for the proof of Theorem 3.5, i.e., for studying the relaxations.

Lemma 7.4. *Let $\varepsilon > 0$ and let $v \in V$ be a fixed vertex of weight w_v . Then the expected number of neighbors of v with weight at least $w^+ := w_v^{(1+\varepsilon)/(\beta-2)}$ is $O(w_{\min}^{\beta-2} w_v^{-\varepsilon})$.*

Proof. We calculate the expected value with a straight-forward integral over all vertex weights larger than w^+ . The probability measure is given by the density $f(w)$, and for a vertex u of fixed weight w_u and random coordinate x_u , the probability to have an edge with v is at most $O(\frac{w_u w_v}{w_{\min} n})$ by Lemma 7.1. Then the expected number of such neighbors of v is at most

$$n \cdot \int_{w^+}^{\infty} f(w) \cdot O\left(\frac{w_u w}{w_{\min} n}\right) dw = O\left(w_{\min}^{\beta-2} w_v \int_{w^+}^{\infty} w^{1-\beta} dw\right) = O\left(w_{\min}^{\beta-2} w_v^{-\varepsilon}\right).$$

\square

Next, we compute the total number of vertices whose objective is larger than a given ϕ_0 .

Lemma 7.5. *Let ϕ_0 be a given objective, and denote by $V_{\geq \phi_0}$ the set of vertices which have objective at least ϕ_0 . Then with probability $1 - e^{-\Omega(\phi_0^{-1})}$, it holds $|V_{\geq \phi_0}| = \Theta(\phi_0^{-1})$.*

Proof. By definition of the objective function ϕ , a vertex u satisfies $\phi(u) \geq \phi_0$ if and only if $\|x_u - x_t\|^d \leq \frac{w_u}{\phi_0 w_{\min} n}$. We integrate over all weights w and see that the expected number of vertices with objective at least ϕ_0 is

$$n w_{\min}^{\beta-1} \int_{w_{\min}}^{\infty} w^{-\beta} \frac{w}{\phi_0 w_{\min} n} dw = \Theta(\phi_0^{-1}).$$

Then by a Chernoff bound, with probability $1 - e^{-\Omega(\phi_0^{-1})}$, this random variable is concentrated around its expectation and thus $\Theta(\phi_0^{-1})$. \square

The following lemma is a consequence of the Poisson point process. It will help us below to deal with events which are not independent as its statement shows how certain random variables are correlated.

Lemma 7.6. *Let $A_1, A_2 \subset \mathbb{T}^d$ and let $u_1, u_2 \notin A_1 \cup A_2$ be two vertices. Then*

$$\Pr[\Gamma(u_2) \cap A_2 = \emptyset \mid \Gamma(u_1) \cap A_1 = \emptyset] \geq \Pr[\Gamma(u_2) \cap A_2 = \emptyset].$$

Proof. Let $A_1, A_2 \subset \mathbb{T}^d$ be two subsets and let u_1, u_2 be two vertices not contained in the sets A_1, A_2 . Suppose we know that u_1 has no neighbors in A_1 . Clearly, conditioning on this event decreases the expected number of vertices in A_1 and in every subset of A_1 , hence

$$\Pr[\Gamma(u_2) \cap (A_1 \cap A_2) = \emptyset \mid \Gamma(u_1) \cap A_1 = \emptyset] \geq \Pr[\Gamma(u_2) \cap (A_1 \cap A_2) = \emptyset].$$

Furthermore, by the Poisson point process, for every set A_3 which is disjoint from A_1 , the number of vertices in A_1 and A_3 is independent. We put $A_3 := A_2 \setminus A_1$ and deduce

$$\Pr[\Gamma(u_2) \cap (A_2 \setminus A_1) = \emptyset \mid \Gamma(u_1) \cap A_1 = \emptyset] = \Pr[\Gamma(u_2) \cap (A_2 \setminus A_1) = \emptyset].$$

Combining the two inequalities gives the desired property. \square

Some of the proofs require that we split the geometric space into small subspaces. We will do this with grids which divide the torus into smaller d -dimensional cubes.

Definition 7.7. *For a given weight $w = w(n)$, a w -grid is a grid which splits the geometric ground space \mathbb{T}^d into $\frac{n}{w}$ equal and disjoint cubes of side lengths $(\frac{w}{n})^d$.*

Here, for simplifying notations, we assume that w divides n . One application of w -grids is the following lemma. For a given w -grid and a given vertex set S , it states that if S contains vertices of sufficiently many different cells of the grid, then at least one vertex $v \in S$ will have a neighbor of weight at least w .

Lemma 7.8 (Bulk lemma). *Let $w_0 = w(n) = \omega(1)$ be a weight growing in n and let S be a set of vertices of weight at most w_0 such that S contains vertices in at least w_0 cells of a fixed w_0 -grid. Then a.a.s., there exists a vertex $v \in S$ and a vertex u of weight at least w_0 such that u and v are neighbors and u and v are contained in the same cell of the w_0 -grid.*

Proof. We uncover the graph in two steps: First we only consider the vertices of weight less than w_0 and afterwards we uncover the remaining vertices. Let S be a fixed set of vertices with weight less than w_0 , and let C_1, \dots, C_{w_0} be cells of a w_0 -grid which contain at least one vertex of $v_i \in S \cap C_i$.

Now, we insert the remaining vertices of high weight at random, according to the definition of the model. Together, the w_0 cells have volume $w_0 \cdot \frac{w_0}{n}$. Therefore, a single vertex u of weight $w_u \geq w_0$ falls with probability $\frac{w_0^2}{n}$ into one of the w_0 cells. Furthermore, if u falls into such a cell C_i , then $\|x_u - x_{v_i}\|^d = O(\frac{w_0}{n})$. By definition, v_i has weight at least w_{\min} , and (EP1) or (EP2) imply

$$\Pr[\{u, v_i\} \in E \mid u, v_i \in C_i] = \Omega(1).$$

Thus every vertex of weight at least w_0 is connected to the set $\{v_1, \dots, v_{w_0}\}$ with probability at least $\Omega(\frac{w_0^2}{n})$. Note that the expected number of vertices of weight at least w_0 is $\Theta(nw_{\min}^{\beta-1}w_0^{1-\beta})$, hence the expected number of vertices of weight at least w_0 which are connected to a vertex of $\{v_1, \dots, v_{w_0}\}$ is at least $\Omega(w_{\min}^{\beta-1}w_0^{3-\beta}) = \omega(1)$. By a Chernoff bound, it follows that a.a.s., there exists a vertex v_i in a cell C_i such that v_i has at least one neighbor u with $w_u \geq w_0$. \square

Note that the vertices of a subcube A induce a random subgraph. With the next lemma we show that this induced subgraph is itself a GIRG.

Lemma 7.9. *Let $r \leq 1/4$, let A be a subcube of \mathbb{T}^d with radius r . Then with appropriate rescaling the induced subgraph $G[A]$ itself is a GIRG with intensity $n\text{VOL}(A)$, except with geometric space $[0, 1]^d$ instead of \mathbb{T}^d .*

Proof. We rescale the cube A by a factor $1/\text{VOL}(A) = (2r)^{-d}$, thus transforming it into $A' = [0, 1]^d$. Then the vertices of A' are given by a Poisson point with intensity $n_{A'} := n\text{VOL}(A)$, and the weights follow the same distribution as before. Note that by the condition $r \leq 1/4$ the distances in A are the same as the distances in A' , only scaled up by a factor $2r$. We only consider the case $\alpha < \infty$; the case $\alpha = \infty$ is analogous. For any two vertices $u, v \in A$, by (EP1) the connection probability is

$$p_{uv} = \Theta\left(\min\left\{\frac{1}{\|x_u - x_v\|_A^{\alpha d}} \cdot \left(\frac{w_u w_v}{w_{\min} n}\right)^\alpha, 1\right\}\right) = \Theta\left(\min\left\{\frac{1}{\|x_u - x_v\|_{A'}^{\alpha d}} \cdot \left(\frac{w_u w_v}{w_{\min} n_{A'}}\right)^\alpha, 1\right\}\right),$$

which is exactly (EP1) for the distances of A' and intensity n' . \square

Lemma 7.9 is useful since all results from [13] were proven for more general geometric spaces, including the space $[0, 1]^d$. Thus they apply to the induced graph $G[A']$. In particular, assume $r = \omega(n^{-d})$, so that the intensity is $n' = \omega(1)$. Then a.a.s. there is a giant component of size $\Theta(n')$ in $G[A]$ by [13, Theorem 2.2]. Moreover, by [13, Lemma 5.2] there exists a constant $c > 0$ such that a.a.s. every vertex of weight at least $(\log n')^c$ is contained in the giant connected component of such a random graph. Thus we obtain the following corollary.

Corollary 7.10. *Let $r = \omega(n^{-d})$, $r \leq 1/4$, let A be a subcube of \mathbb{T}^d with radius r , and let $G[A]$ be the random subgraph induced by all vertices in A . Then there exists a constant $c > 0$ such that all vertices in A of weight at least $(\log(r^d n))^c$ are in the same connected component of $G[A]$.*

7.3 Where to Expect Neighbors & Unlikely Jumps

We start by introducing technical notation. For all $\varepsilon > 0$, we put $\gamma(\varepsilon) := \frac{1-\varepsilon}{\beta-2}$. Next, let $\varepsilon_1 = \varepsilon_1(\alpha, \beta) > 0$ be a fixed constant which we choose sufficiently small during proofs, e.g., such that $\gamma(\varepsilon_1) > 1$. Next, we define the following two classes of vertices:

$$V_1 := \{v \in V : \phi(v) \leq w_v^{-\gamma(\varepsilon_1)}\} \quad \text{and} \quad V_2 := \{v \in V : \phi(v) \geq w_v^{-\gamma(\varepsilon_1)}\}.$$

In every hop, the process chooses a neighbor u of the current vertex v which maximizes $\phi(u)$. As discussed above in Section 6, it turns out that if $v \in V_1$, then typically w_u is by an exponent $\gamma(\varepsilon)$ larger than w_v . On the other hand, if $v \in V_2$, then we expect that $\phi(u) \sim \phi(v)^{1/\gamma(\varepsilon)}$. Therefore, vertices in V_1 will correspond to the first phase of the routing, where the current weight increases by an exponent with every hop. The vertices in V_2 correspond to the second phase where the objective is increased by an exponent with every hop (due to $\phi(v) < 1$, indeed the objective increases).

In order to find such a trajectory, we want to classify the vertices in V which accelerate the routing as desired. Let $\zeta := \max\{\frac{3}{2}, \frac{2\alpha-1}{2\alpha+4-2\beta}\}$ for $\alpha < \infty$ and $\zeta := \frac{3}{2}$ in the limit case $\alpha = \infty$. Then for $v \in V_1$ and $\varepsilon = \varepsilon(n) > 0$ we define

$$\begin{aligned} V^+(v, \varepsilon) &:= \{u \in V \mid w_u \geq w_v^{\gamma(\varepsilon)} \wedge \phi(u) \geq \phi(v)w_v^{\gamma(\varepsilon)-1}\}, \text{ and} \\ V^-(v, \varepsilon) &:= \{u \in V \mid w_u \leq w_v^{\gamma(\zeta \cdot \varepsilon)} \wedge \phi(u) \geq \phi(v)w_v^{\gamma(\varepsilon)-1}\}. \end{aligned} \tag{4}$$

In the first phase, the goal will be to increase the weight within every hop of the routing. Therefore, $\Gamma(v) \cap V^+(v, \varepsilon)$ is the set of good neighbors as every $u \in \Gamma(v) \cap V^+(v, \varepsilon)$ has significantly larger weight than v . On the other side, the set $\Gamma(v) \cap V^-(v, \varepsilon)$ contains the bad neighbors as its vertices have small weight but nevertheless large objective and therefore could force the routing to proceed with a low-weight-vertex.

In the second phase, the goal is to increase the objective by an exponent in every step of the routing. Furthermore, once we reach vertices of V_2 , we want to stay in V_2 . Let $v \in V_2$. Similarly

as above, we introduce a set $V^+(v, \varepsilon)$ of good vertices and a set $V^-(v, \varepsilon)$ of bad vertices. More precisely, for $\varepsilon = \varepsilon(n)$ we define

$$\begin{aligned} V^+(v, \varepsilon) &:= \{u \in V_2 \mid \phi(u) \geq \phi(v)^{1/\gamma(\varepsilon)}\}, \text{ and} \\ V^-(v, \varepsilon) &:= \{u \in V_1 \mid \phi(u) \geq \phi(v)^{1/\gamma(\varepsilon)}\}. \end{aligned} \tag{5}$$

In the following we collect several technical statements about the expected number of good and bad neighbors. For $0 < \varepsilon \leq \varepsilon_1$, let $w_1(\varepsilon) := O(e^{d/\varepsilon})$ and $\phi_1(\varepsilon) := \Omega(e^{-d/\varepsilon})$ with sufficiently large respectively small hidden constants. We start with the first lemma which considers vertices in V_1 and shows that the expected number of good neighbors is large whereas the expected number of bad neighbors is small. Notice that for $\varepsilon = \varepsilon_1$, every vertex $v \in V_1$ satisfies $\phi(v)\mathbf{w}_v^{\gamma(\varepsilon_1)} \leq 1$ by the definition of V_1 , and in this case, the statement can be applied for all vertices in V_1 .

Lemma 7.11 (Neighborhoods in first phase). *Let $0 < \varepsilon \leq \varepsilon_1$ and let $v \in V_1$ be a vertex such that $\phi(v)\mathbf{w}_v^{\gamma(\varepsilon)} \leq 1$. Then*

$$(i) \quad \mathbb{E}[|\Gamma(v) \cap V^+(v, \varepsilon)|] = \Omega\left(\mathbf{w}_{\min}^{\beta-2} \mathbf{w}_v^\varepsilon\right), \text{ and}$$

$$(ii) \quad \text{if in addition } \mathbf{w}_v \geq w_1(\varepsilon), \text{ then } \mathbb{E}[|\Gamma(v) \cap V^-(v, \varepsilon)|] = O\left(\mathbf{w}_{\min}^{\beta-2} \mathbf{w}_v^{-\Omega(\varepsilon)}\right).$$

Proof. Let $0 < \varepsilon \leq \varepsilon_1$ be such that $\phi(v)\mathbf{w}_v^{\gamma(\varepsilon)} \leq 1$. Let $v \in V_1$.

Proof of (i): In order to prove the first statement, we want to lower-bound the expected number of neighbors of v in $V^+(v, \varepsilon)$, i.e., the neighbors with larger weight and slightly larger objective. For this, define $\delta := \left(\frac{c_1 \mathbf{w}_v^{1+\gamma(\varepsilon)}}{\mathbf{w}_{\min} n}\right)^{1/d}$, where c_1 is the constant given by (EP1) and (EP2). Then let $A(v, \varepsilon)$ be the following set of vertices:

$$A(v, \varepsilon) := \{u \in V \mid (1) : w_u \geq \mathbf{w}_v^{\gamma(\varepsilon)}; (2) : \|\mathbf{x}_u - \mathbf{x}_t\| \leq \|\mathbf{x}_v - \mathbf{x}_t\|; (3) : \|\mathbf{x}_u - \mathbf{x}_v\| \leq \delta\}.$$

Note that $A(v, \varepsilon) \subseteq V^+(v, \varepsilon)$, because for $u \in A(v, \varepsilon)$

$$\phi(u) = \frac{\mathbf{w}_u}{\mathbf{w}_{\min} n \|\mathbf{x}_u - \mathbf{x}_t\|^d} \geq \frac{\mathbf{w}_v^{\gamma(\varepsilon)}}{\mathbf{w}_{\min} n \|\mathbf{x}_v - \mathbf{x}_t\|^d} \geq \phi(v)\mathbf{w}_v^{\gamma(\varepsilon)-1}$$

holds. Hence for proving the first statement it is sufficient to lower-bound $\mathbb{E}[|\Gamma(v) \cap A(v, \varepsilon)|]$. By the probability distribution of the weights there are in expectation $n\mathbf{w}_{\min}^{\beta-1} \mathbf{w}_v^{\gamma(\varepsilon)(1-\beta)}$ vertices satisfying Condition (1) and Condition (3) is satisfied independently with probability δ^d . Now, observe that by assumption the vertex v satisfies $\phi(v)\mathbf{w}_v^{\gamma(\varepsilon)} \leq 1$, which implies

$$\delta^d = \frac{c_1 \mathbf{w}_v^{1+\gamma(\varepsilon)}}{\mathbf{w}_{\min} n} = c_1 \phi(v)\mathbf{w}_v^{\gamma(\varepsilon)} \|\mathbf{x}_v - \mathbf{x}_t\|^d \leq c_1 \|\mathbf{x}_v - \mathbf{x}_t\|^d.$$

Hence a random vertex v fulfilling (3) satisfies with constant probability (2) as well. Furthermore, Condition (3) ensures that by (EP1) and (EP2), every vertex $u \in A(v, \varepsilon)$ is connected to v with constant probability. Combining this yields

$$\mathbb{E}[|\Gamma(v) \cap A(v, \varepsilon)|] = \Omega\left(n\mathbf{w}_{\min}^{\beta-1} \mathbf{w}_v^{\gamma(\varepsilon)(1-\beta)} \delta^d\right) = \Omega\left(\mathbf{w}_{\min}^{\beta-2} \mathbf{w}_v^{1-\gamma(\varepsilon)(\beta-2)}\right) = \Omega\left(\mathbf{w}_{\min}^{\beta-2} \mathbf{w}_v^\varepsilon\right).$$

Proof of (ii): We want to upper-bound the expected number of bad neighbors of v , i.e., the neighbors in $V^-(v, \varepsilon)$.

At first, we will show that the geometric distance between a bad node $u \in V^-(v, \varepsilon)$ and v is not much less than the geometric distance between v and t . This will imply a sufficiently small connection probability between u and v . For this purpose, define $\delta(w) := \|\mathbf{x}_v - \mathbf{x}_t\| (w \mathbf{w}_v^{-\gamma(\varepsilon)})^{1/d}$ for all weights $w \leq \mathbf{w}_v^{\gamma(\zeta \cdot \varepsilon)}$. Due to $\mathbf{w}_v \geq w_1(\varepsilon) = 2^{O(d/\varepsilon)}$, for w_1 large enough every weight $w \leq \mathbf{w}_v^{\gamma(\zeta \cdot \varepsilon)}$ satisfies $\delta(w)^d \leq \|\mathbf{x}_v - \mathbf{x}_t\|^d \mathbf{w}_v^{\gamma(\zeta \cdot \varepsilon) - \gamma(\varepsilon)} \leq (0.5 \|\mathbf{x}_v - \mathbf{x}_t\|)^d$ (the second inequality follows since $\zeta > 1$). Moreover, for a vertex $u \in V^-(v, \varepsilon)$ we observe that the condition $\phi(u) \geq \phi(v) \mathbf{w}_v^{\gamma(\varepsilon) - 1}$ implies $\|\mathbf{x}_u - \mathbf{x}_t\| \leq \delta(\mathbf{w}_u)$ and with the triangle inequality, we obtain that u fulfills $\|\mathbf{x}_u - \mathbf{x}_v\| \geq 0.5 \|\mathbf{x}_v - \mathbf{x}_t\|$.

In order to upper-bound $\mathbb{E}[|\Gamma(v) \cap V^-(v, \varepsilon)|]$, we proceed with a case distinction regarding the parameter α .

Case $\alpha = \infty$: Let $c_2 > 0$ be the constant given by (EP2), and let $w_1(\varepsilon)$ be large enough. Then

$$0.5^d \geq c_2 w_1^{-\varepsilon(\zeta-1)/(\beta-2)} \geq c_2 \mathbf{w}_v^{-\varepsilon(\zeta-1)/(\beta-2)} = c_2 \mathbf{w}_v^{\gamma(\zeta \cdot \varepsilon) - \gamma(\varepsilon)}.$$

Furthermore, $\phi(v) \mathbf{w}_v^{\gamma(\varepsilon)} \leq 1$ implies $\frac{\mathbf{w}_v^{1+\gamma(\varepsilon)}}{\mathbf{w}_{\min} n} = \phi(v) \mathbf{w}_v^{\gamma(\varepsilon)} \|\mathbf{x}_v - \mathbf{x}_t\|^d \leq \|\mathbf{x}_v - \mathbf{x}_t\|^d$. Together with the above observation we deduce

$$\|\mathbf{x}_u - \mathbf{x}_v\|^d \geq 0.5^d \|\mathbf{x}_v - \mathbf{x}_t\|^d \geq c_2 \|\mathbf{x}_v - \mathbf{x}_t\|^d \mathbf{w}_v^{\gamma(\zeta \cdot \varepsilon) - \gamma(\varepsilon)} \geq \frac{c_2 \mathbf{w}_v^{1+\gamma(\zeta \cdot \varepsilon)}}{\mathbf{w}_{\min} n} \geq \frac{c_2 \mathbf{w}_v \mathbf{w}_u}{\mathbf{w}_{\min} n}.$$

But then, by (EP2) the two vertices u and v can not be connected, hence v has deterministically no neighbor in $V^-(v, \varepsilon)$ in the case $\alpha = \infty$.

Case $\alpha < \infty$: By the geometrical property $\|\mathbf{x}_u - \mathbf{x}_v\| \geq 0.5 \|\mathbf{x}_v - \mathbf{x}_t\|$ and (EP1) the probability that a vertex $u \in V^-(v, \varepsilon)$ is connected to v is at most

$$O\left(\left(\frac{\mathbf{w}_u \mathbf{w}_v}{\|\mathbf{x}_u - \mathbf{x}_v\|^d \mathbf{w}_{\min} n}\right)^\alpha\right) = O\left(\left(\frac{\mathbf{w}_u \mathbf{w}_v}{\|\mathbf{x}_v - \mathbf{x}_t\|^d \mathbf{w}_{\min} n}\right)^\alpha\right) = O((\mathbf{w}_u \phi(v))^\alpha).$$

Furthermore, we have already seen that such a vertex of weight w needs to be in distance at most $\delta(w)$ to t , which a random vertex of weight w does with probability $\delta(w)^d$. We integrate over all weights between \mathbf{w}_{\min} and $\mathbf{w}_v^{\gamma(\zeta \cdot \varepsilon)}$, where we use the density function $f(w)$ for the distribution of the weights. We deduce that

$$\begin{aligned} \mathbb{E}[|\Gamma(v) \cap V^-(v, \varepsilon)|] &= O\left(n \int_{\mathbf{w}_{\min}}^{\mathbf{w}_v^{\gamma(\zeta \cdot \varepsilon)}} f(w) \delta(w)^d (\mathbf{w}_u \phi(v))^\alpha dw\right) \\ &= O\left(\mathbf{w}_{\min}^{\beta-2} \phi(v)^{\alpha-1} \mathbf{w}_v^{1-\gamma(\varepsilon)} \int_{\mathbf{w}_{\min}}^{\mathbf{w}_v^{\gamma(\zeta \cdot \varepsilon)}} w^{\alpha+1-\beta} dw\right) \\ &= O\left(\mathbf{w}_{\min}^{\beta-2} \phi(v)^{\alpha-1} \mathbf{w}_v^{1-\gamma(\varepsilon)+\gamma(\zeta \cdot \varepsilon)(\alpha+2-\beta)}\right). \end{aligned}$$

Next we use that $\phi(v) \mathbf{w}_v^{\gamma(\varepsilon)} \leq 1$ holds. Then the above term is at most $O(\mathbf{w}_{\min}^{\beta-2} \mathbf{w}_v^\tau)$, where

$$\tau = 1 - \alpha\gamma(\varepsilon) + \gamma(\zeta \cdot \varepsilon)(\alpha + 2 - \beta) = \frac{\alpha\varepsilon - \zeta\varepsilon(\alpha + 2 - \beta)}{\beta - 2} = -\Omega(\varepsilon)$$

by our choice of ζ . □

Next, similar to Lemma 7.11 for the first phase we give a lemma for the second phase where the routing process reached V_2 . We calculate for a vertex $v \in V_2$ the expected number of good and bad neighbors.

Lemma 7.12 (Neighborhoods in second phase). *Let $0 < \varepsilon \leq \varepsilon_1$ and let $v \in V_2$ be a vertex such that $\phi(v) \leq 1$. Then*

$$(i) \mathbb{E}[|\Gamma(v) \cap V^+(v, \varepsilon)|] = \Omega\left(\mathbf{w}_{\min}^{\beta-2} \phi(v)^{-\Omega(\varepsilon)}\right), \text{ and}$$

$$(ii) \text{ if in addition } \phi(v) \leq \phi_1(\varepsilon), \text{ then } \mathbb{E}[|\Gamma(v) \cap V^-(v, \varepsilon)|] = O\left(\mathbf{w}_{\min}^{\beta-2} \phi(v)^{\Omega(\varepsilon)}\right).$$

Proof. We prove the statement similar as we proved Lemma 7.11. Let $0 < \varepsilon \leq \varepsilon_1$.

Proof of (i): Let $v \in V_2$ such that $\phi(v) \leq 1$. In order to prove the first statement, we want to lower-bound the expected number of neighbors of v in the set $V^+(v, \varepsilon)$. In particular, such neighbors have objective at least $\phi(v)^{1/\gamma(\varepsilon)}$. Let $\delta := (\mathbf{w}_{\min}^{-1} n^{-1} \phi(v)^{-1-1/\gamma(\varepsilon)})^{1/d}$ and let $c' = \min\{1, c_1/2^d\}$, where c_1 is the constant given by (EP1) and (EP2). We consider the set

$$A(v, \varepsilon) := \{u \in V \mid (1) : \|x_u - x_t\| \leq \delta; (2) : \mathbf{w}_u \geq (c' \phi(v))^{-1} \geq \phi(v)^{-1}\}.$$

By the definition of δ and Condition (2), a vertex $u \in A(v, \varepsilon)$ satisfies

$$\phi(u) \geq \frac{\mathbf{w}_u}{\delta^d \mathbf{w}_{\min} n} \geq \mathbf{w}_u \phi(v)^{1+1/\gamma(\varepsilon)} \geq \phi(v)^{1/\gamma(\varepsilon)}.$$

On the other hand, by (2) and the assumption $\phi(v) \leq 1$ it holds $\mathbf{w}_u \geq 1$, and then we have $\mathbf{w}_u \phi(v)^{1+1/\gamma(\varepsilon)} \geq \mathbf{w}_u^{1/\gamma(\varepsilon)} \geq \mathbf{w}_u^{-\gamma(\varepsilon_1)}$. Thus, $A(v, \varepsilon) \subseteq V^+(v, \varepsilon)$.

Our next goal is to upper-bound the geometric distance between vertices of $A(v, \varepsilon)$ and the vertex v itself. We observe that the assumptions $v \in V_2$ and $\varepsilon \leq \varepsilon_1$ imply

$$\phi(v)^{-1/\gamma(\varepsilon)} \leq \mathbf{w}_v^{\gamma(\varepsilon_1)/\gamma(\varepsilon)} \leq \mathbf{w}_v,$$

and it follows

$$n \mathbf{w}_{\min} \delta^d = \phi(v)^{-1-1/\gamma(\varepsilon)} \leq \phi(v)^{-1} \mathbf{w}_v = n \mathbf{w}_{\min} \|x_v - x_t\|^d.$$

By the triangle inequality we obtain $\|x_u - x_v\| \leq 2\|x_v - x_t\|$ for all $u \in A(v, \varepsilon)$.

Next we want to verify that all vertices in $A(v, \varepsilon)$ are connected to v with constant probability. We notice that for all $u \in A(v, \varepsilon)$, by Condition (2) we have

$$\frac{c_1 \mathbf{w}_u \mathbf{w}_v}{\mathbf{w}_{\min} n} = c_1 \mathbf{w}_u \phi(v) \|x_v - x_t\|^d \geq \frac{c_1}{c'} \|x_v - x_t\|^d \geq \frac{c_1}{2^d c'} \|x_u - x_v\|^d \geq \|x_u - x_v\|^d.$$

by our choice of c' . Then, indeed by (EP1) and (EP2), all vertices of $A(v, \varepsilon)$ are connected to v with constant probability. The expected number of vertices with weight at least $(c' \phi(v))^{-1}$ is $\Omega(n(\mathbf{w}_{\min} \phi(v))^{\beta-1})$, and thus

$$\mathbb{E}[|\Gamma(v) \cap A(v, \varepsilon)|] = \Omega\left(n(\mathbf{w}_{\min} \phi(v))^{\beta-1} \delta^d\right) = \Omega\left(\mathbf{w}_{\min}^{\beta-2} \phi(v)^{\beta-2-1/\gamma(\varepsilon)}\right) = \Omega\left(\mathbf{w}_{\min}^{\beta-2} \phi(v)^{-\Omega(\varepsilon)}\right),$$

as ε is chosen small enough.

Proof of (ii): Let $v \in V_2$ such that $\phi(v) \leq \phi_1$ and let $u \in V^-(v, \varepsilon)$. We first prove that u has the property $\|x_u - x_v\| \geq 0.5\|x_v - x_t\|$. Since $u \in V^-(v, \varepsilon)$, we have $\phi(u) \geq \phi(v)^{1/\gamma(\varepsilon)}$, and by the assumption $\phi(v) \leq \phi_1$ for ε_1 small enough this is at least $\phi(v)^{1-\varepsilon} \geq 2^d \phi(v)$, where we used $\phi(v) \leq \phi_1$ in the last step. Because $u \in V_1$, it follows

$$\|x_u - x_t\|^d = \phi(u)^{-1} \frac{\mathbf{w}_u}{\mathbf{w}_{\min} n} \leq \frac{\phi(u)^{-1-1/\gamma(\varepsilon_1)}}{\mathbf{w}_{\min} n} \leq \frac{(2^d \phi(v))^{-1-1/\gamma(\varepsilon_1)}}{\mathbf{w}_{\min} n}.$$

On the other hand, $v \in V_2$, and then

$$\|\mathbf{x}_u - \mathbf{x}_t\|^d = \frac{(2^d \phi(v))^{-1-1/\gamma(\varepsilon_1)}}{\mathbf{w}_{\min} n} \leq \frac{0.5^d \phi(v) \mathbf{w}_v}{\mathbf{w}_{\min} n} = 0.5^d \|\mathbf{x}_v - \mathbf{x}_t\|^d.$$

By the triangle inequality $\|\mathbf{x}_u - \mathbf{x}_v\| \geq 0.5 \|\mathbf{x}_v - \mathbf{x}_t\|$ holds as desired. Moreover, $u \in V_1$ and thus $\mathbf{w}_u \leq \phi(u)^{-1/\gamma(\varepsilon_1)}$. Since $\phi(u) \geq \phi(v)$, for ε_1 small enough it follows

$$\mathbf{w}_u \phi(v) \leq \phi(u)^{-1/\gamma(\varepsilon_1)} \phi(v) \leq \phi(v)^{1-1/\gamma(\varepsilon)} \leq \phi(v)^\varepsilon.$$

Case $\alpha = \infty$: We obtain

$$\frac{\mathbf{w}_u \mathbf{w}_v}{\|\mathbf{x}_u - \mathbf{x}_v\|^d} \leq \frac{\mathbf{w}_u \mathbf{w}_v}{0.5^d \|\mathbf{x}_v - \mathbf{x}_t\|^d} = 2^d \mathbf{w}_u \phi(v) \mathbf{w}_{\min} n \leq 2^d \phi(v)^{\Omega(\varepsilon)} \mathbf{w}_{\min} n.$$

Together with the assumption $\phi(v) \leq \phi_1$ this implies $\frac{c_2 \mathbf{w}_u \mathbf{w}_v}{\|\mathbf{x}_u - \mathbf{x}_v\|^d \mathbf{w}_{\min} n} \leq c_2 2^d \phi_1^{\Omega(\varepsilon)}$, and by taking ϕ_1 small enough, this expression is smaller than 1. Then (EP2) yields $u \notin \Gamma(v)$ deterministically.

Case $\alpha < \infty$: We have

$$p_{uv} = O\left(\left(\frac{\mathbf{w}_u \mathbf{w}_v}{\|\mathbf{x}_u - \mathbf{x}_v\|^d \mathbf{w}_{\min} n}\right)^\alpha\right) = O\left(\left(\frac{\mathbf{w}_u \mathbf{w}_v}{\|\mathbf{x}_v - \mathbf{x}_t\|^d \mathbf{w}_{\min} n}\right)^\alpha\right) = O((\mathbf{w}_u \phi(v)^\alpha).$$

Recall that we observed already above that $\mathbf{w}_u \leq \phi(v)^{-1/\gamma(\varepsilon)}$, and note that every vertex $u \in V^-(v, \varepsilon)$ has the geometric property $\|\mathbf{x}_u - \mathbf{x}_t\|^d \leq \phi(v)^{-1/\gamma(\varepsilon)} \frac{\mathbf{w}_u}{\mathbf{w}_{\min} n}$, which a random vertex of given weight w satisfies with probability at most $\phi(v)^{-1/\gamma(\varepsilon)} \frac{w}{\mathbf{w}_{\min} n}$. We conclude that

$$\begin{aligned} \mathbb{E}[|\Gamma(v) \cap V^-(v)|] &= O\left(\int_{\mathbf{w}_{\min}}^{\phi(v)^{-1/\gamma(\varepsilon)}} n \mathbf{w}_{\min}^{\beta-1} w^{-\beta} \phi(v)^{-1/\gamma(\varepsilon)} \frac{w}{\mathbf{w}_{\min} n} (w \phi(v))^\alpha dw\right) \\ &= O\left(\mathbf{w}_{\min}^{\beta-2} \phi(v)^{-1/\gamma(\varepsilon)+\alpha} \int_{\mathbf{w}_{\min}}^{\phi(v)^{-1/\gamma(\varepsilon)}} w^{\alpha+1-\beta} dw\right) \\ &= O\left(\mathbf{w}_{\min}^{\beta-2} \phi(v)^{\frac{\tau}{1-\varepsilon}}\right), \end{aligned}$$

where $\tau = (\alpha + 2 - \beta)(1 - 1/\gamma(\varepsilon)) - \alpha\varepsilon = \Omega(\varepsilon)$ for $\varepsilon \leq \varepsilon_1$ and ε_1 small enough. This proves (ii). \square

In order to prove our main results, we want to show that with sufficiently high probability, greedy routing follows our anticipated trajectory. If $v \in V_1$ is a vertex on the greedy path, we want that the best neighbor of v is in $V^+(v, \varepsilon)$, i.e., that its weight is significantly larger than \mathbf{w}_v . By Lemma 7.11 (ii) the expected number vertices in $\Gamma(v)$ with small weight and relatively large objective is polynomially small in \mathbf{w}_v . By Markov's inequality, this will give us an error probability which is polynomially small in \mathbf{w}_v . If v has constant weight, this is not precise enough as for the proof of Theorem 3.2, we need an error probability which is exponentially small in \mathbf{w}_{\min} . With our next lemma, we improve this as follows: We consider vertices $v \in V_1$ whose objective is relatively small compared to the weight \mathbf{w}_v , and calculate the expected size of $V^-(v, \varepsilon)$ depending on $\phi(v)$. Later on, this statement can be applied in particular to vertices of constant weight with large geometric distance to the target.

Lemma 7.13. *There exists a constant $\bar{c} > 0$ such that for all $0 < \varepsilon \leq \varepsilon_1$ and all vertices $v \in V_1$ with $\mathbf{w}_v \geq w_1(\varepsilon)$ and $\phi(v) \mathbf{w}_v^{\gamma(\varepsilon)} \leq \bar{c}$, it holds*

$$\mathbb{E}[|\Gamma(v) \cap V^-(v, \varepsilon)|] = O(\phi(v)^{\Omega(1)} \mathbf{w}_v^{O(1)}).$$

Proof. Let $0 < \varepsilon < \varepsilon_1$ and let $v \in V$ be a vertex such that $\phi(v)w_v^{\gamma(\zeta \cdot \varepsilon)} \leq \bar{c} := 1/(2^d c_2)$, where c_2 is the constant given by (EP2). We want to show that with relatively high probability, every neighbor of v with small weight is located close to v on the torus. For all weights $w \leq w_v^{\gamma(\zeta \cdot \varepsilon)}$ we define a critical radius

$$r_0(w) := 0.5 \left(\frac{w_v^{1-\gamma(\zeta \cdot \varepsilon)} w}{w_{\min} n \phi(v)} \right)^{1/d}.$$

Next we define

$$B(v, \varepsilon) := \{u \in V \mid (1) : w_u \leq w_v^{\gamma(\zeta \cdot \varepsilon)}; (2) : \|x_v - x_u\| \geq r_0(w_u)\}.$$

We first show that the expected size of the set $\Gamma(v) \cap B(v, \varepsilon)$ is very small. Afterwards, we will prove that $V^-(v, \varepsilon) \subseteq B(v, \varepsilon)$ holds under the assumptions of the lemma.

With the assumption $\phi(v) \leq \bar{c} w_v^{-\gamma(\zeta \cdot \varepsilon)}$ we deduce that $r_0(w)^d \geq \frac{w w_v}{2^d \bar{c} w_{\min} n} = \frac{c_2 w w_v}{w_{\min} n}$ holds for all weights $w \leq w_v^{\gamma(\zeta \cdot \varepsilon)}$.

Case $\alpha = \infty$: By (EP2) the vertex v has deterministically no neighbors in $B(v, \varepsilon)$.

Case $\alpha < \infty$: By (EP1) a vertex $u \in B(v)$ with distance r to v is connected to v with probability $O((\frac{w_v w_u}{r^d w_{\min} n})^\alpha)$. We calculate $\mathbb{E}[|\Gamma(v) \cap B(v, \varepsilon)|]$ by integrating over all weights smaller than $w_v^{\gamma(\zeta \cdot \varepsilon)}$ and all distances $r \geq r_0(w)$. Using the density function $f(w) = \Theta(w_{\min}^{\beta-1} w^{-\beta})$ we deduce

$$\begin{aligned} \mathbb{E}[|\Gamma(v) \cap B(v, \varepsilon)|] &= O \left(\int_{w_{\min}}^{w_v^{\gamma(\zeta \cdot \varepsilon)}} n w_{\min}^{\beta-1} w^{-\beta} \int_{r_0(w)}^{1/2} \left(\frac{w_v w}{r^d w_{\min} n} \right)^\alpha r^{d-1} dr dw \right) \\ &= O \left(n w_{\min}^{\beta-1} \int_{w_{\min}}^{w_v^{\gamma(\zeta \cdot \varepsilon)}} \left(\frac{w_v w}{w_{\min} n} \right)^\alpha r_0(w)^{d(1-\alpha)} w^{-\beta} dw \right) \\ &= O \left(n w_{\min}^{\beta-1} \int_{w_{\min}}^{w_v^{\gamma(\zeta \cdot \varepsilon)}} \frac{w_v w}{w_{\min} n} w_v^{\gamma(\zeta \cdot \varepsilon)(\alpha-1)} \phi(v)^{\alpha-1} w^{-\beta} dw \right) \\ &= O \left(w_{\min}^{\beta-2} w_v^{\gamma(\zeta \cdot \varepsilon)(\alpha-1)+1} \phi(v)^{\alpha-1} \int_{w_{\min}}^{w_v^{\gamma(\zeta \cdot \varepsilon)}} w^{1-\beta} dw \right) \\ &= O \left(w_v^{\gamma(\zeta \cdot \varepsilon)(\alpha-1)+1} \phi(v)^{\alpha-1} \right). \end{aligned}$$

We finish the proof of the lemma by showing that $V^-(v, \varepsilon) \subseteq B(v, \varepsilon)$ holds under the assumption that v satisfies in addition $w_v \geq w_1(\varepsilon)$. By assumption, $v \in V_1$ and all $u \in V^-(v, \varepsilon)$ fulfill $w_u \leq w_v^{\gamma(\zeta \cdot \varepsilon)}$. Next we observe that by the definition of the set $V^-(v, \varepsilon)$ for a vertex $v \in V_1$, all $u \in V^-(v, \varepsilon)$ satisfy

$$\|x_u - x_t\|^d = \frac{w_u}{n w_{\min} \phi(u)} \leq \frac{w_v^{\gamma(\zeta \cdot \varepsilon)}}{n w_{\min} \phi(v) w_v^{\gamma(\varepsilon)-1}} = \|x_v - x_t\|^d w_v^{\gamma(\zeta \cdot \varepsilon) - \gamma(\varepsilon)} = \|x_v - x_t\|^d w_v^{-\Omega(\varepsilon)}.$$

Hence $\|x_u - x_t\| \leq \|x_v - x_t\| w_v^{-\Omega(\varepsilon)/d} \leq 0.5 \|x_v - x_t\|$ as we assume $w_v \geq w_1(\varepsilon)$. By the triangle inequality we deduce $\|x_v - x_u\| \geq 0.5 \|x_v - x_t\|$. On the other side, since $w_u \leq w_v^{\gamma(\zeta \cdot \varepsilon)}$, the definition of $r_0(w)$ implies

$$r_0(w_u)^d = 0.5^d \frac{w_v^{1-\gamma(\zeta \cdot \varepsilon)} w_u}{w_{\min} n \phi(v)} \leq 0.5^d \|x_v - x_t\|^d.$$

It follows $\|x_v - x_u\| \geq 0.5 \|x_v - x_t\| \geq r_0(w_u)$, which proves $V^-(v, \varepsilon) \subseteq B(v, \varepsilon)$. \square

We turn to the last lemma of this chapter. In order to prove Theorem 3.3 and Theorem 3.4, we will need to control the trajectory in particular at the start and at the end of the routing. For technical reasons, we require that the routing will not see any edges which increase the objective too fast. Statement (i) of this lemma states that a.a.s. there are no vertices of both high weight and high objective at all. Statements (ii) and (iii) upper-bound the probability that a vertex of small weight has an incident edge of unexpected length, i.e., a neighbor which is located geometrically far away. Finally (iv) upper-bounds this probability for a vertex of high weight.

Lemma 7.14 (Unlikely jumps). *Let $w_0 = w_0(n) = \omega(1)$ be a growing function in n and $v \in V$. Then the following statements hold.*

(i) *With prob. $1 - w_0^{-\Omega(1)}$, there exists no vertex u (except potentially t) s.t. $w_u \geq w_0$ and $\phi(u) \geq w_0^{-\Omega(1)}$.*

Let $w_v \leq w_0$:

(ii) *With prob. $1 - O(w_0^{-5})$ there exists no neighbor u of v with $w_u \leq w_0$ and $\|x_u - x_v\|^d \geq \frac{w_0^{O(1)}}{n}$.*

(iii) *Let $\phi(v) \leq w_0^{-\Omega(1)}$. With probability $1 - O(w_0^{-5})$ there exists no neighbor u of v with $w_u \geq w_0$ and $\phi(u) \leq \phi_0^{O(1)}$.*

Let $w_v \geq w_0$:

(iv) *Let $\phi(v) \leq w_0^{-1-\Omega(1)}$, and let $M := \min\{w_v, \phi(v)^{-1}\}$.*

With prob. $1 - O(M^{-\Omega(1)})$ there exists no neighbor u of v with $w_u \leq w_0$ and $\phi(u) \geq w_0^{-\Omega(1)}$.

Proof. Let $c > 0$ be a constant chosen sufficiently large, ε be a constant chosen sufficiently small, let $r_0 := (w_0^c/n)^{1/d}$ and for every weight $w \geq w_{\min}$ define $\delta(w) := (\frac{ww_0^\varepsilon}{w_{\min}n})^{1/d}$.

Proof of (i): We first observe that a vertex u of given weight w_u satisfies $\phi(u) \geq w_0^{-\varepsilon}$ if and only if $\|x_u - x_t\|^d \leq \delta(w_u)^d$. Then excluding t which is potentially fixed, the expected number of vertices with weight at least w_0 and objective at least $w_0^{-\varepsilon}$ is at most

$$\int_{w_0}^{\infty} n w_{\min}^{\beta-1} w^{-\beta} \delta(w)^d dw = \int_{w_0}^{\infty} w_{\min}^{\beta-2} w_0^\varepsilon w^{1-\beta} dw = O\left(w_{\min}^{\beta-2} w_0^{2-\beta+\varepsilon}\right) = w_0^{-\Omega(1)},$$

where we used that for ε sufficiently small it holds $\beta - 2 \geq 2\varepsilon$. By Markov's inequality with probability $1 - w_0^{-\Omega(1)}$ there exist no such vertices of weight at least w_0 .

Proof of (ii): Let v be a vertex of weight at most w_0 and define

$$A_0(v) := \{u \in V \mid w_u \leq w_0 \text{ and } \|x_u - x_v\|^d \geq r_0\}.$$

For c sufficiently large and for any vertex $u \in A_0(v)$ we have $\frac{w_u w_v}{w_{\min} n} \leq \frac{w_0^2}{w_{\min} n} = o(r_0^d)$.

Case $\alpha = \infty$: The claim follows as by (EP2) deterministically u and v do not share an edge, i.e., $\Gamma(v) \cap A_0(v) = \emptyset$.

Case $\alpha < \infty$: By (EP1) and the assumption $w_v \leq w_0$, the probability p_{uv} is at most $O((w_0^2/(n\|x_u - x_v\|^d))^\alpha)$. Thus we have

$$\mathbb{E}[|\Gamma(v) \cap A_0(v)|] = O\left(n \int_{r_0}^{1/2} r^{d-1} \left(\frac{w_0^2}{nr^d}\right)^\alpha dr\right) = O\left(w_0^{2\alpha} n^{1-\alpha} r_0^{d(1-\alpha)}\right) = O\left(w_0^{2\alpha+c(1-\alpha)}\right).$$

If we take $c \geq \frac{5+2\alpha}{\alpha-1}$, then we obtain $\mathbb{E}[|\Gamma(v) \cap A_0(v)|] = O(w_0^{-5})$, and by Markov's inequality with probability at least $1 - O(w_0^{-5})$, the set $\Gamma(v) \cap A_0(v)$ is empty.

Proof of (iii): Let v be a vertex of weight at most w_0 which satisfies $\phi(v) \leq w_0^{-\varepsilon}$, and define

$$A_1(v) := \{u \in V \mid (1) : w_u \geq w_0; (2) : \phi(u) \leq \phi(v)^c\}.$$

We want to show that a.a.s. v has no neighbor in the set $A_1(v)$. Define $\delta'(w) := (\frac{w}{\phi(v)^\varepsilon w_{\min} n})^{1/d}$ and observe that a vertex $u \in V$ with $w_u \geq w_0$ is contained in $A_1(v)$ if and only if $\delta'(w_u) \leq \|x_u - x_t\|$. However, due to $w_v \leq w_0$, for $c > 1/\varepsilon$ we have $\delta'(w_u)^d = \omega(\frac{w_u w_v}{w_{\min} n})$.

Case $\alpha = \infty$: The claim follows because (EP2) implies that, deterministically, v does not have neighbors in $A_1(v)$.

Case $\alpha < \infty$: Note that for $w \geq w_0$ we have $\delta'(w) = \omega(\|x_v - x_t\|)$, and by the triangle inequality we deduce that for a vertex $u \in A_1(v)$ it holds $\|x_u - x_v\| \geq 0.5\|x_u - x_t\|$. After this preparation, we can calculate the expected number of vertices in $\Gamma(v) \cap A_1(v)$ by integrating over the complete set $A_1(v)$. We obtain

$$\begin{aligned} \mathbb{E}[|\Gamma(v) \cap A_1(v)|] &= O\left(\int_{w_0}^{\infty} n w_{\min}^{\beta-1} w^{-\beta} \int_{\delta'(w)}^{1/2} r^{d-1} \left(\frac{w w_v}{(0.5r)^d w_{\min} n}\right)^\alpha dr dw\right) \\ &= O\left(\int_{w_0}^{\infty} n w_{\min}^{\beta-1} w^{-\beta} \left(\frac{w w_v}{w_{\min} n}\right)^\alpha \delta'(w)^{d(1-\alpha)} dw\right) \\ &= O\left(w_{\min}^{\beta-2} w_v^\alpha \phi(v)^{c(\alpha-1)} \int_{w_0}^{\infty} w^{1-\beta} dw\right) \\ &= O\left(w_{\min}^{\beta-2} w_v^\alpha \phi(v)^{c(\alpha-1)} w_0^{2-\beta}\right) = O\left(w_v^\alpha w_0^{\varepsilon c(1-\alpha)}\right). \end{aligned}$$

Since $w_v \leq w_0$ by assumption, we can choose c large enough such that this value is at most $O(w_0^{-5})$. Then Markov's inequality implies that indeed with probability $1 - O(w_0^{-5})$, the set $\Gamma(v) \cap A_1(v)$ is empty. This proves Statement (iii).

Proof of (iv): Let v be a vertex with $w_v \geq w_0$ and objective $\phi(v) \leq w_0^{-1-\varepsilon}$, and define $M := \min\{w_v, \phi(v)^{-1}\}$. We want to upper-bound the probability that v has a neighbor in

$$A_2(v) := \{u \in V \mid (1) : w_u \leq w_0; (2) : \phi(u) \geq w_0^{-\varepsilon}\}.$$

We know that every vertex $u \in A_2(v)$ satisfies

$$\|x_u - x_t\|^d \leq \delta(w_u)^d = \frac{w_u w_0^\varepsilon}{w_{\min} n} = o\left(\frac{w_u w_v}{w_{\min} n}\right).$$

On the other hand, the vertex v satisfies

$$\|x_v - x_t\|^d = \frac{w_v}{\phi(v) w_{\min} n} \geq \frac{w_v w_0^{1+\varepsilon}}{w_{\min} n} = \omega\left(\frac{w_u w_v}{w_{\min} n}\right).$$

By the triangle inequality it follows $\|x_v - x_u\| = \omega\left(\frac{w_u w_v}{w_{\min} n}\right)$.

Case $\alpha = \infty$: The claim follows by (EP2) because there exists deterministically no edge between u and v .

Case $\alpha < \infty$: We deduce $p_{uv} = O((\phi(v) w_u)^\alpha)$. By integrating over all weights $w \in [w_{\min}, w_0]$ we obtain

$$\begin{aligned} \mathbb{E}[|\Gamma(v) \cap A_2(v)|] &= O\left(\int_{w_{\min}}^{w_0} n w_{\min}^{\beta-1} w^{-\beta} \delta(w)^d (\phi(v) w)^\alpha dw\right) \\ &= O\left(w_{\min}^{\beta-2} w_0^\varepsilon (\phi(v) w_0)^\alpha \int_{w_{\min}}^{w_0} w^{1-\beta} dw\right) = O\left(w_0^\varepsilon (\phi(v) w_0)^\alpha\right). \end{aligned}$$

We continue with a case distinction. First suppose $\phi(v) \leq \mathbf{w}_v^{-1-2\varepsilon}$. Since $\alpha > 1$, $\phi(v) \leq w_0^{-1}$, and $\mathbf{w}_v \geq w_0$, in this case the above expression is at most

$$O(w_0^{1+\varepsilon}\phi(v)) = O(w_0^{1+\varepsilon}\mathbf{w}_v^{-1-2\varepsilon}) = O(\mathbf{w}_v^{-\varepsilon}) = O(M^{-\Omega(1)}).$$

On the other hand, if $\phi(v) \geq \mathbf{w}_v^{-1-2\varepsilon}$ we use

$$\phi(v) \leq w_0^{-1-\varepsilon} \Leftrightarrow w_0 \leq \phi(v)^{-1/(1+\varepsilon)}$$

and get

$$\mathbb{E}[|\Gamma(v) \cap A_2(v)|] = O(w_0^{\varepsilon+\alpha}\phi(v)^\alpha) = O(\phi(v)^{\alpha-\frac{\alpha+\varepsilon}{1+\varepsilon}}) = O(\phi(v)^{\frac{\alpha\varepsilon-\varepsilon}{1+\varepsilon}}) = O(M^{-\Omega(1)}).$$

Then again by Markov's inequality, with probability $1 - O(M^{-\Omega(1)})$ the set $\Gamma(v) \cap A_2(v)$ is empty, which finishes the proof. \square

8 Proof: Success Probability and Length of Greedy Routing

In this section we analyze the basic greedy process and prove Theorem 3.1 and Theorem 3.2 about the success probability and Theorem 3.3 about the running time. We advise the reader to first read Section 6 where we describe the high-level ideas that we use to prove our main Lemma 8.1. In a nutshell, we distinguish two phases of the process (governed by weights and by objectives, respectively), and for each of the phases we divide the space V into layers. Then we prove that with sufficiently large probability a greedy path never visits more than one vertex per layer, and a.a.s. it does not fail in any layer except for the very first or the very last ones. This observation is formalized in Lemma 8.1 and lies at the heart of all our main results, including the results for patching and for the relaxations. All algorithms follow very similar trajectories, only the starting phase and the end phase differs for the various algorithms. The layer structure is depicted in Figure 1 on page 17.

8.1 Main Lemma

Let us start by repeating the notations which we introduced in Section 7.3. For all $\varepsilon > 0$, we denote $\gamma(\varepsilon) := \frac{1-\varepsilon}{\beta-2}$. Furthermore $\varepsilon_1 = \varepsilon_1(\alpha, \beta) > 0$ is a fixed constant, chosen sufficiently small, $\zeta = \zeta(\alpha, \beta)$ is also a fixed constant, $w_1(\varepsilon) = O(e^{d/\varepsilon})$, and $\phi_1(\varepsilon) = \Omega(e^{-d/\varepsilon})$. Then we have the following two classes of vertices:

$$V_1 := \{v \in V : \phi(v) \leq \mathbf{w}_v^{-\gamma(\varepsilon_1)}\} \quad \text{and} \quad V_2 := \{v \in V : \phi(v) \geq \mathbf{w}_v^{-\gamma(\varepsilon_1)}\}.$$

Finally, for a given vertex $v \in V$ and a given $\varepsilon > 0$ we defined good sets $V^+(v, \varepsilon)$ and bad sets $V^-(v, \varepsilon)$ above in (4) and (5).

In Section 6 we claimed that the expected trajectory of the routing process is such that soon after the start the current weight is $\omega(1)$. We also expect that, at the end of the routing, when reaching rather large objectives, very few additional hops are needed to reach t . Therefore the main part of the process plays in-between. In this section we state and prove a lemma which describes the routing process in this main part. For any pair (w, ϕ) of a weight and an objective we define

$$\bar{V}(w, \phi) := \{v \in V_1 \mid (1) : \mathbf{w}_v \geq w; (2) : \phi(v) \leq \phi\} \cup \{v \in V_2 \mid \phi(v) \leq \phi\}.$$

Suppose that both w and ϕ are constants. Then the main part of the routing process will happen in this set $\bar{V}(w, \phi)$. With the following lemma, we study the routing process inside such a set. We show that with sufficiently high probability, the routing will not die out in $\bar{V}(w, \phi)$ and will be ultra-fast in traversing this set. Thereby the probabilities for the failure events will depend on w and ϕ , which can be both growing functions as well. We will apply the lemma for both the basic greedy algorithm and the patching algorithms.

Lemma 8.1 (Main lemma). *Let $w_0 \geq w_1(\varepsilon_1)$, $\phi_0 \leq \phi_1(\varepsilon_1)$, and $M := \min\{w_0, \phi_0^{-1}\}$. Furthermore let $f_0(n) = \omega(1)$ be any growing function such that $f_0(n) = o(\log \log n)$, and let A be a routing protocol which satisfies (P1), i.e., it makes greedy choices. Suppose $\phi(s) \leq \phi_0$, denote by P the greedy path induced by all vertices of objective at most ϕ_0 starting from s , and suppose that there exists a first vertex $u_1 \in P \cap \bar{V}(w_0, \phi_0)$. Then with probability*

$$1 - O\left(\mathbf{w}_{\min}^{\beta-2} M^{-\Omega(1)}\right),$$

there exists a subpath $P' = (u_1, \dots, u_\ell)$ of P starting in u_1 with the following properties:

- (i) P' is contained in $\bar{V}(w_0, \phi_0)$.
- (ii) Either $P' \subset V_1$, or $P' \subset V_2$, or there exists a vertex $u_{\ell'}$ such that $\{u_1, \dots, u_{\ell'}\} \in V_1$ and $\{u_{\ell'+1}, \dots, u_\ell\} \in V_2$.
- (iii) Let $\{u_i, u_{i+1}, u_{i+2}\}$ be three subsequent vertices on $P' \cap V_1$. Then $\mathbf{w}_{u_{i+2}} \geq \mathbf{w}_{u_i}^{\gamma(\zeta\varepsilon_1)}$.
- (iv) Let $\{u_i, u_{i+1}, u_{i+2}\}$ be three subsequent vertices on $P' \cap V_2$. Then $\phi(u_{i+2}) \geq \phi(u_i)^{1/\gamma(\varepsilon_1)}$.
- (v) The length ℓ of the path P' is at most $\frac{2+o(1)}{|\log(\beta-2)|} \log \log n$.

Moreover and more precisely, the length of P' is upper bounded by

$$\frac{1+o(1)}{|\log(\beta-2)|} \left(\log \log_{w_0}(\phi(u_1)^{-1}) + \log \log_{\phi_0^{-1}}(\phi(u_1)^{-1}) \right) + O(f_0(n)).$$

- (vi) The following holds

$$\mathbb{E}_{>\phi_0} [|\Gamma(u_\ell) \cap V^+(u_\ell, \varepsilon_1) \cap V_{>\phi_0}|] = \Omega\left(\mathbf{w}_{\min}^{\beta-2} M^{\Omega(1)}\right).$$

Here $\mathbb{E}_{>\phi_0}$ means expectation w.r.t. uncovering the vertices of objectives larger than ϕ_0 , conditioned on position and weight of u_ℓ .

Note that Lemma 8.1 naturally applies to the situation that we have uncovered all vertices of objective at most ϕ_0 , but not the rest of the graph. More precisely, conditions (i) – (v) are independent of $V_{>\phi_0}$, and (vi) makes a statement about the marginal expectation after uncovering $V_{\leq\phi_0}$ and before uncovering $V_{>\phi_0}$.

Proof. Let P be the greedy path induced by all vertices of objective at most ϕ_0 . We analyze the structure of P by partitioning the set $\bar{V}(w_0, \phi_0)$ into several small layers. Then the idea will be to show that with sufficiently high probability, there exists a vertex u_ℓ on the path P such that until reaching u_ℓ , P visits every layer at most once, and the vertex u_ℓ has neighbors of objective at least ϕ_0 . There will be two classes of layers: The layers $A_{1,j}$ divide the area $V_1 \cap \bar{V}(w_0, \phi_0)$ and are defined via weights, and the layers $A_{2,j}$ divide the area $V_2 \cap \bar{V}(w_0, \phi_0)$ and are defined via objectives.

We may assume $f_0(n) = o(\log \log n)$. Recall that we already put $\varepsilon_1 = \Theta(1)$. In addition, we put $\varepsilon_2 := (\log \log f_0(n))^{-1} = o(1)$, and we may assume $\varepsilon_2 < \varepsilon_1$. Furthermore, we define the landmarks

$$w'_0 := w_0^{(\gamma(\zeta\varepsilon_1)f_0(n))} \quad \text{and} \quad \phi'_0 := \phi_0^{(\gamma(\varepsilon_1)f_0(n))}.$$

These landmarks will be thresholds for slightly different definitions of the layers. For later reference we note that, since $\gamma(\zeta\varepsilon_1) > 1$,

$$w_0^{\Omega(\varepsilon_2)} = w_0^{\Omega(\gamma(\zeta\varepsilon_1)f_0(n)/\log \log f_0(n))} = \omega\left(w_0^{\Omega(1)}\right) \quad \text{and} \quad \phi_0^{\Omega(\varepsilon_2)} = o\left(\phi_0^{\Omega(1)}\right), \quad (6)$$

where the second equation follows analogously to the first.

Layers by weight: We define a sequence $y_0 := w_0 < y_1 < \dots < y_j < \dots$ for the weights such that this sequence grows doubly exponentially. More precisely, we require

$$y_{j+1} = \begin{cases} y_j^{\gamma(\zeta\varepsilon_1)} & \text{if } y_j < w'_0, \\ y_j^{\gamma(\varepsilon_2)} & \text{if } y_j \geq w'_0. \end{cases}$$

Then we define the layers by weight which partition the set $V'_1 := \{v \in \bar{V}(w_0, \phi_0) \mid \phi(v)\mathbf{w}_v^{\gamma(\varepsilon_2)} \leq 1\}$.

$$A_{1,j} := \{v \in V'_1 \mid y_{j-1} \leq \mathbf{w}_v < y_j\} \quad \forall j \geq 1.$$

Note that by definition, $V'_1 \subset V_1$ and no vertex of V_1 has larger weight than $(\mathbf{w}_{\min}n)^{(1+\gamma(\varepsilon_1))^{-1}}$. Thus we only need to consider sets $A_{1,j}$ for which $y_{j-1} \leq (\mathbf{w}_{\min}n)^{(1+\gamma(\varepsilon_1))^{-1}}$. Furthermore, and only for technical reasons we introduce the additional layer

$$A_{1,\infty} := \{v \in \bar{V}(w_0, \phi_0) \mid \phi(v)\mathbf{w}_v^{\gamma(\varepsilon_1)} \leq 1 \leq \phi(v)\mathbf{w}_v^{\gamma(\varepsilon_2)}\},$$

which covers the remaining vertices of $V_1 \cap \bar{V}(w_0, \phi_0)$.

Layers by objective: Similarly, define a sequence $\psi_0 = \phi_0 > \psi_1 > \dots > \psi_j > \dots$ for the objectives such that this sequence falls doubly exponentially, i.e.,

$$\psi_{j+1} = \begin{cases} \psi_j^{\gamma(\varepsilon_1)} & \text{if } \psi_j > \phi'_0, \\ \psi_j^{\gamma(\varepsilon_2)} & \text{if } \psi_j \leq \phi'_0. \end{cases}$$

Next, we define similar layers for the vertices in $V_2 \cap \bar{V}(w_0, \phi_0)$.

$$A_{2,j} := \{v \in V_2 \cap \bar{V}(w_0, \phi_0) \mid \psi_{j-1} \geq \phi(v) > \psi_j\} \quad \forall j \geq 1.$$

For this second set of layers, note that every vertex in the graph has objective at least \mathbf{w}_{\min}/n and we do not need to consider layers $A_{2,j}$ which would contain vertices of even smaller objective.

Definition of the events \mathcal{E} and $\mathcal{E}_{i,j}$: The layers $A_{i,j}$ as defined above classify all vertices of objective at most ϕ_0 . We want to show that with sufficiently high probability, P contains at most one vertex per layer $A_{i,j}$ until reaching a vertex u_ℓ for which the expected number of neighbors in $V_{>\phi}$ is large. We will do this by considering the layers in the following consecutive order:

$$A_{1,1} \prec \dots \prec A_{1,j} \prec A_{1,j+1} \prec \dots \prec A_{1,\infty} \prec \dots \prec A_{2,j} \prec A_{2,j-1} \prec \dots \prec A_{2,1}.$$

Given this ordering of the layers, we denote by $B_{i,j}$ the union of the layer $A_{i,j}$ and all previous layers, and by $P_{i,j}$ we denote the greedy path induced by the set $B_{i,j} \subset V$.

In the following, we consider for a given pair (i, j) the first vertex $v \in P_{i,j} \cap A_{i,j}$, if such a vertex exists. We will first show that with sufficiently high probability the neighbor of v with highest objective is located outside $B_{i,j}$, which we consider a “good” event.

More precisely, let $i \in \{1, 2\}$, $j \geq 1$, and let $\varepsilon \in \{\varepsilon_1, \varepsilon_2\}$ be the ε which was used for the definition of $A_{i,j}$. Moreover, let $\eta > 0$ be a sufficiently small constant, and let $E := \eta w_{\min}^{\beta-2} M^\eta$. Note that for any $c > 0$, by choosing η small enough we can achieve for all $v \in A_{i,j}$ that $c w_{\min}^{\beta-2} w_v^\varepsilon > 2E$ and $c w_{\min}^{\beta-2} \phi(v)^{-c\varepsilon} > 2E$.

Denote by $\mathcal{E}_{i,j}$ the event that either $P_{i,j} \cap A_{i,j} = \emptyset$, or the first vertex $v \in P_{i,j} \cap A_{i,j}$

- satisfies condition (vi), or
- has at least one good neighbor, i.e., a neighbor $v' \in \overline{V}(w_0, \phi_0) \setminus B_{i,j}$ with $\phi(v') \geq \phi(v)$ s.t. $\phi(v') > \phi(u)$ holds for all $u \in \Gamma(v) \cap B_{i,j}$.

Finally, we denote by $\mathcal{E} := \cap_{i,j} \mathcal{E}_{i,j}$ the intersection of all good events.

Lower bound for $\Pr[\mathcal{E}_{i,j}]$: In this part, our goal is to lower-bound $\Pr[\mathcal{E}_{i,j}]$ for every pair (i, j) in order to lower-bound $\Pr[\mathcal{E}]$. Thereby, our construction of layers gives rise to several different subcases, corresponding to different subphases of the routing.

Let (i, j) be a given pair for which we want to bound $\Pr[\mathcal{E}_{i,j}]$ from below. Clearly we have $\Pr[\mathcal{E}_{i,j} \mid P_{i,j} \cap A_{i,j} = \emptyset] = 1$, and we only need to consider $\Pr[\mathcal{E}_{i,j} \mid P_{i,j} \cap A_{i,j} \neq \emptyset]$. Let v be the first vertex on $P_{i,j} \cap A_{i,j}$. Again, if v satisfies condition (vi) then $\mathcal{E}_{i,j}$ holds and there is nothing to show, so we assume otherwise,

$$\mathbb{E}_{>\phi_0} [|\Gamma(v) \cap V^+(v, \varepsilon_1) \cap V_{>\phi_0}|] < E. \quad (7)$$

In all cases let $\varepsilon \in \{\varepsilon_1, \varepsilon_2\}$ be the same ε as used for the definition of the considered layer.

Case $i = 1$ and $j < \infty$: We are in the first phase of the routing, and the layer $A_{1,j}$ is contained in $V_1' \subset V_1$. Then by Lemma 7.11 (i) we have that the expected number of neighbors of v in $V^+(v, \varepsilon)$ is large, i.e., $\mathbb{E}[|\Gamma(v) \cap V^+(v, \varepsilon)|] = \Omega(w_{\min}^{\beta-2} w_v^\varepsilon)$. Since $\varepsilon \leq \varepsilon_1$ we have $V^+(v, \varepsilon) \subseteq V^+(v, \varepsilon_1)$, so in particular we do not expect many of these neighbors in $V_{>\phi_0}$, i.e.,

$$\mathbb{E}_{>\phi_0} [|\Gamma(v) \cap V^+(v, \varepsilon) \cap V_{>\phi_0}|] \stackrel{(7)}{<} E \leq \frac{1}{2} \mathbb{E}[|\Gamma(v) \cap V^+(v, \varepsilon)|], \quad (8)$$

where the second inequality comes from our choice of the constant η in E . Thus we expect most of these good neighbors in $V_{\leq\phi_0}$, i.e.,

$$\mathbb{E}[|\Gamma(v) \cap V^+(v, \varepsilon) \cap V_{\leq\phi_0}|] = \Omega(w_{\min}^{\beta-2} w_v^\varepsilon).$$

By a Chernoff bound, using $w_v \geq y_{j-1}$, with probability at least $1 - \exp(-\Omega(w_{\min}^{\beta-2} y_{j-1}^\varepsilon))$ there exists a neighbor v' of v in $V^+(v, \varepsilon) \cap V_{\leq\phi_0}$. By the definition of $V^+(v, \varepsilon)$, such a neighbor fulfills $v' \notin B_{1,j}$ and $\phi(v') \geq \phi(v) w_v^{\gamma(\varepsilon)-1}$.

Now, we want to show that v' actually is a good neighbor of v , i.e., every $u \in \Gamma(v) \cap B_{i,j}$ has objective less than $\phi(v) w_v^{\gamma(\varepsilon)-1}$. We observe that every vertex with objective at least $\phi(v) w_v^{\gamma(\varepsilon)-1}$ and weight at most y_j is contained in $V^-(v, \varepsilon)$, thus we want to upper-bound $\mathbb{E}[|\Gamma(v) \cap V^-(v, \varepsilon)|]$. Note that here we need to condition on that no earlier vertex on the path $P_{1,j}$ had a neighbor

$u \in V^-(v, \varepsilon)$ (because otherwise, u would be the first vertex in $A_{i,j} \cap P_{i,j}$ and not v). By Lemma 7.6, this condition decreases the expected number of vertices in $V^-(v, \varepsilon)$, and therefore also the expected number of neighbors of v in $V^-(v, \varepsilon)$. Hence, in order to give an upper bound on this value, we are allowed to neglect the dependencies. Since $y_{j-1} \geq w_1(\varepsilon) = O(e^{d/\varepsilon})$, by Lemma 7.11 (ii) we deduce

$$\mathbb{E}[|\Gamma(v) \cap V^-(v, \varepsilon)| \mid v \in A_{1,j} \cap P_{1,j}] \leq \mathbb{E}[|\Gamma(v) \cap V^-(v, \varepsilon)|] = O(w_{\min}^{\beta-2} y_{j-1}^{-\Omega(\varepsilon)}).$$

Then we apply Markov's inequality to see that with probability at least $1 - O(w_{\min}^{\beta-2} y_{j-1}^{-\Omega(\varepsilon)})$, v has no neighbor in $V^-(v, \varepsilon)$ and therefore every $u \in \Gamma(v, \varepsilon) \cap B_{1,j}$ has objective at most $\phi(u) < \phi(v) w_v^{\gamma(\varepsilon)-1} \leq \phi(v')$ as desired. It follows

$$\Pr[\mathcal{E}_{1,j}] \geq 1 - O(w_{\min}^{\beta-2} y_{j-1}^{-\Omega(\varepsilon)}). \quad (9)$$

Note that depending on how the weight y_{i-1} compares to w'_0 , we are either using ε_1 or ε_2 , and the last exponent of the above expression is either $-\Omega(1)$ or $-\Omega((\log \log f_0(n))^{-1})$.

Case $i = 1$ and $j = \infty$: Next we give a lower bound for $\Pr[\mathcal{E}_{1,\infty}]$ in exactly the same way. Let v be a vertex in this extra layer $A_{1,\infty}$, then by definition it has the property $w_v^{-\gamma(\varepsilon_2)} \leq \phi(v) \leq w_v^{-\gamma(\varepsilon_1)}$. Therefore, we can apply Lemma 7.11 (i) only with ε_1 and not with ε_2 . We observe that every vertex v' in the considered set $V^+(v, \varepsilon_1)$ satisfies

$$\phi(v') \geq \phi(v) w_v^{\gamma(\varepsilon_1)-1} \geq w_v^{\gamma(\varepsilon_1)-1-\gamma(\varepsilon_2)} \geq w_{v'}^{(\gamma(\varepsilon_1)-1-\gamma(\varepsilon_2))/\gamma(\zeta\varepsilon_1)} \geq w_{v'}^{-\gamma(\varepsilon_1)}$$

for ε_1 small enough. Then indeed $v' \notin B_{1,\infty}$, and thus $V^+(v, \varepsilon_1) \subset V_2$. This allows us to argue exactly in the same arguments as above in the case $j < \infty$. Then we apply Lemma 7.11 (ii) and deduce that

$$\Pr[\mathcal{E}_{1,\infty}] \geq 1 - O(w_{\min}^{\beta-2} w_0^{-\Omega(\varepsilon)}). \quad (10)$$

Case $i = 2$: We continue with the events $\mathcal{E}_{2,j}$ for the layers $A_{2,j}$. Here, the argument works similar: Again we can assume that the set $A_{2,j} \cap P_{2,j}$ is non-empty, and that the first vertex $v \in A_{2,j} \cap P_{2,j}$ satisfies (7). By Lemma 7.12 (i) we have $\mathbb{E}[|\Gamma(v) \cap V^+(v, \varepsilon)|] = \Omega(w_{\min}^{\beta-2} \phi(v)^{-\Omega(\varepsilon)})$, and as in (8) our choice of E ensures that also $\mathbb{E}[|\Gamma(v) \cap V^+(v, \varepsilon) \cap V_{\leq \phi_0}|] = \Omega(w_{\min}^{\beta-2} \phi(v)^{-\Omega(\varepsilon)})$. By the Chernoff bound, and since $\psi_{j-1} \geq \phi(v) > \psi_j$, with probability at least $1 - \exp(-\Omega(w_{\min}^{\beta-2} \psi_{j-1}^{-\Omega(\varepsilon)}))$ the set $\Gamma(v) \cap V^+(v, \varepsilon) \cap V_{\leq \phi_0}$ is non-empty. In this case, there exists a neighbor v' in $\bar{V}(w_0, \phi_0) \setminus B_{2,j}$ with objective larger than ψ_{j-1} . It remains to show that there exists no neighbor in V_1 with higher objective. By Lemma 7.6 and the same arguments as above we can use Lemma 7.12 (ii) to see that the expected number of neighbors of v inside $B_{2,j}$ with objective larger than ψ_{j-1} is $O(w_{\min}^{\beta-2} \psi_{j-1}^{\Omega(\varepsilon)})$. Then Markov's inequality shows that with probability $1 - O(w_{\min}^{\beta-2} \psi_{j-1}^{\Omega(\varepsilon)})$, there exists no such neighbor. Thus

$$\Pr[\mathcal{E}_{2,j}] \geq 1 - O(w_{\min}^{\beta-2} \psi_j^{\Omega(\varepsilon)}). \quad (11)$$

Here, the last exponent is either $\Omega(1)$ or $\Omega((\log \log f_0(n))^{-1})$, depending on which ε we consider.

Lower bound for $\Pr[\mathcal{E}]$: We have computed a lower bound for $\Pr[\mathcal{E}_{i,j}]$ for every layer $A_{i,j}$. Combining Equations (9),(10), and (11) yields

$$\begin{aligned} \Pr[\mathcal{E}] &\geq 1 - \sum_j \Pr[\mathcal{E}_{1,j}^c] - \Pr[\mathcal{E}_{1,\infty}] - \sum_{j \geq 1} \Pr[\mathcal{E}_{1,j}^c] \\ &= 1 - O\left(\mathbf{w}_{\min}^{\beta-2} \left(w_0^{-\Omega(1)} + w_0'^{-\Omega(\varepsilon_2)} + \phi_0^{\Omega(1)} + \phi_0'^{\Omega(\varepsilon_2)}\right)\right) \\ &\stackrel{(6)}{=} 1 - O\left(\mathbf{w}_{\min}^{\beta-2} \min\{w_0, \phi_0^{-1}\}^{-\Omega(1)}\right). \end{aligned}$$

\mathcal{E} implies (i)-(vi): Having proven that \mathcal{E} occurs with sufficiently high probability, it remains to show that this event implies all desired properties (i)-(vi). Hence suppose that \mathcal{E} occurs and that P visits at least one vertex of $\bar{V}(w_0, \phi_0)$. Then the first vertex $u_1 \in P \cap \bar{V}(w_0, \phi_0)$ is contained in some layer $A_{i,j}$. We show that there exists a subpath $P' = \{u_1, \dots, u_\ell\} \subseteq P$ which visits every layer at most once in both phases and then implies (i)-(vi).

First Phase ($i = 1$): Suppose that $u_1 \in V_1$. Then there exists a maximal subpath

$$P_1 = \{u_1, \dots, u_{\ell'}\} \subseteq P$$

starting at u_1 such that $P_1 \subset V_1$ and no vertex before $u_{\ell'}$ satisfies Condition (vi). We start with an analysis of this subpath P_1 and show by induction that every vertex of P_1 is located in a higher layer than its predecessor. Clearly, the induction hypothesis holds for the base case $\{u_1\}$. Assume that the hypothesis holds for $\{u_1, \dots, u_i\}$. Then u_i is contained in layer $A_{1,j}$, and by the induction hypothesis, u_i is the first vertex that P visits in this layer. Moreover, by the induction hypothesis it follows that $\{u_1, \dots, u_i\}$ is also a subpath of the induced greedy path $P_{1,j}$. Then by the event $\mathcal{E}_{1,j}$ we have one of the following two cases:

- u_i satisfies condition (vi):

Then set $P' = (u_1, \dots, u_\ell)$ with $u_\ell := u_i$.

- u_i has a good neighbor v' :

Let u_{i+1} be the best neighbor of u_i . Then $u_{i+1} \in V_2$, or u_{i+1} is located in a layer $A_{1,j'}$ where $j' > j$ because u_{i+1} has higher weight than every vertex in $A_{1,j}$. Since the protocol makes greedy choices, u_{i+1} follows u_i on the path P .

This proves the induction hypothesis. We see that P_1 traverses the layers $A_{1,j}$ according to our ordering and visits every layer at most once. Furthermore, applying the same argument for the last vertex $u_{\ell'} \in P_1$ shows that either $u_{\ell'}$ satisfies (vi), or it has a good neighbor in a layer $A_{2,j}$. Note that in the former case, we would choose $u_\ell := u_{\ell'}$ and $P' = P_1$, and (i), (ii), (vi) follow directly. Otherwise, we claim that we will find a path P' such that $P' \cap V_1 = P_1$.

Proof of (iii): If $\{u_i, u_{i+1}, u_{i+2}\}$ are three subsequent vertices on the subpath P_1 , then the weight increases at least by an exponent $\gamma(\zeta\varepsilon_1)$ between u_i and u_{i+2} since there exists at least one layer in-between containing u_{i+1} . Then (iii) follows as we are assuming that P' will not visit any vertex in $V_1 \setminus P_1$.

Second Phase ($i = 2$): Suppose that either the last vertex $u_{\ell'} \in P_1$ does not satisfy (vi) or $s \in V_2$. Then there exists a first vertex $u_{\ell'+1} \in P \cap V_2'$, and a maximal subpath

$$P_2 = P_1 \cup \{u_{\ell'+1}, \dots, u_\ell\} \subseteq P$$

starting at u_1 such that $(P_2 \setminus P_1) \subset V_2$ and no vertex before u_ℓ satisfies Condition (vi). We show by induction that every vertex of $P_2 \setminus P_1$ is located in a higher layer (w.r.t. the objective ϕ) than its

predecessor. Again, the hypothesis holds for $P_1 \cup \{u_{\ell'+1}\}$. Assume that the induction hypothesis holds for $P_1 \cup \{u_{\ell'+1}, \dots, u_i\}$. Then u_i is contained in a layer $A_{2,j}$, and by the induction hypothesis, u_i is the first vertex that P visits in this layer. Furthermore, $P_1 \cup \{u_{\ell'+1}, \dots, u_i\}$ is a subpath of the induced greedy path $P_{2,j}$. Then by the event $\mathcal{E}_{2,j}$, we have one of the following two cases:

- u_i satisfies condition (vi):

Then set $P' = (u_1, \dots, u_\ell)$ with $u_\ell := u_i$.

- u_i has a good neighbor v' :

Let u_{i+1} be the best neighbor of u_i . Then $u_i \in A_{2,j'}$ holds for $j' < j$. Since the protocol makes greedy choices, u_{i+1} follows u_i on the path P .

This proves the induction hypothesis, and we observe that P_2 traverses the layers $A_{2,j}$ according to our ordering and visits every layer at most once. However, by repeating the argument for the last vertex u_ℓ it follows that the best neighbor of u_ℓ can not be in V_1 . Then the only possibility is that u_ℓ satisfies Condition (vi). Hence we put $P' = P_2$, and Properties (i), (ii), and (vi) follow. In particular, we see that $P' \cap V_1 = P_1$ as claimed above.

Proof of (iv): For three subsequent vertices $\{u_i, u_{i+1}, u_{i+2}\}$ on the subpath $\{u_{\ell'+1}, \dots, u_\ell\}$, the objective increases always at least by an exponent $\gamma(\varepsilon_1)$ between u_i and u_{i+2} as there lies at least one layer in-between containing u_{i+1} . This proves (iv), and the only remaining property is (v).

Proof of (v): We already know that the path P' visits every layer at most once. Therefore we can upper-bound the length $|P'|$ by counting the total number of potentially visited layers. By construction we have $f_0(n)$ first layers $A_{1,j}$ which are defined via ε_1 , and similarly there are $f_0(n)$ final layers $A_{2,j}$ defined via ε_1 . Let L_2 denote the number of layers $A_{2,j}$ defined via $\varepsilon_2 = o(1)$ which the routing potentially visited. In order to upper-bound L_2 , we recall that between two neighboring vertices of P' , the objective is always increasing. Therefore P' contains only vertices of objective at least $\phi(u_1)$. Clearly, this implies that we only need to consider layers $A_{2,j}$ where $\psi_{j-1} \geq \phi(u_1)$. Therefore L_2 is upper-bounded by the solution of $\phi_0^{-\gamma(\varepsilon_2)L_2} = \phi(u_1)^{-1}$, and we obtain

$$L_2 = \frac{\log \log_{\phi_0^{-1}}(\phi(u_1)^{-1})}{\log \gamma(\varepsilon_2)} \leq \frac{\log \log_{\phi_0^{-1}}(\phi(u_1)^{-1})}{\log((\beta - 2)^{-1/(1+o(1))})} = \frac{1 + o(1)}{|\log(\beta - 2)|} \log \log_{\phi_0^{-1}}(\phi(u_1)^{-1}).$$

It remains to upper-bound the number L_1 of layers $A_{1,j}$, defined via ε_2 , which P' potentially visits. We observe that every vertex $v \in P'$ satisfies $\phi(v) \geq \phi(u_1)$. Thus, every vertex $v \in P'$ with weight at least $\phi(u_1)^{-1/\gamma(\varepsilon_1)} < \phi(u_1)^{-1}$ belongs to V_2 , and we only need to count layers $A_{1,j}$ which contain vertices of weight at most $\phi(u_1)^{-1/\gamma(\varepsilon_1)}$. Then L_1 is bounded from above by the solution of $w_0^{\gamma(\varepsilon_2)L_1} = \phi(u_1)^{-1}$. It follows

$$L_1 = \frac{\log \log_{w_0}(\phi(u_1)^{-1})}{\log \gamma(\varepsilon_2)} \leq \frac{\log \log_{w_0}(\phi(u_1)^{-1})}{\log((\beta - 2)^{-1/(1+o(1))})} = \frac{1 + o(1)}{|\log(\beta - 2)|} \log \log_{w_0}(\phi(u_1)^{-1}).$$

Then the length of P' is at most $L_1 + L_2 + O(f(n))$, which proves (v). \square

8.2 Proof of Theorem 3.1 (Success Probability)

In this subsection we prove that for all initial choices of s and t , greedy routing succeeds with probability $\Omega(1)$. We do this in three steps: In a first step we show that with probability $\Omega(1)$, the best neighbor u_1 of s has weight $w_{u_1} \geq w_0$, where w_0 is a given constant. In a second step we use Lemma 8.1 to see that then with probability $\Omega(1)$, from u_1 the greedy algorithm on $V_{\leq \phi_0}$ finds

a path to a vertex u_ℓ which has in expectation $\Omega(1)$ neighbors of objective at least ϕ_0 . Since the number of such neighbors is Poisson distributed, with probability $\Omega(1)$ there is at least one such neighbor, and we call the best of them v_1 . Finally, in a last step we will show that with constant probability, the routing will reach t from v_1 via at most one intermediate vertex v_2 . For all steps, the success probability is a constant larger than zero, but not large enough such that we can do a union bound over the error events. Therefore, we always compute the probabilities conditioned on the previous steps being successful.

Proof of Theorem 3.1. We assume that the parameter w_{\min} is lower-bounded by a small constant, say ε_1 . Let us start by choosing a weight $w_0 = \Theta(1)$ and an objective $\phi_0 = \Theta(1)$ such that they satisfy the preconditions of Lemma 8.1 and such that the events (i)-(vi) of Lemma 8.1 occur with constant probability $\Omega(1) > 0$.

Start of routing process: Suppose for now that $\phi(s) \leq \phi_0$. We define the event

$$\mathcal{E}_s : s \text{ itself or its best neighbor have weight at least } w_0.$$

If $w_s \geq w_0$, this happens with probability 1. Thus we assume $w_s \leq w_0$. Let $r := (\frac{c_1 w_s w_0}{w_{\min}^n})^{1/d}$, where c_1 is the constant given by (EP1) and (EP2). Since $\phi(s) \leq \phi_0$, we have

$$r^d = c_1 w_0 \phi(s) \|\mathbf{x}_s - \mathbf{x}_t\|^d \leq c_1 w_0 \phi_0 \|\mathbf{x}_s - \mathbf{x}_t\|^d = O(\|\mathbf{x}_s - \mathbf{x}_t\|^d). \quad (12)$$

Next we define a subset of V in which we want to find a neighbor of s . Let

$$A_s := \left\{ u \in V \left| \begin{array}{l} (1) : w_u \geq w_0; \\ (2) : \|\mathbf{x}_u - \mathbf{x}_t\| \leq \|\mathbf{x}_s - \mathbf{x}_t\|; \\ (3) : \|\mathbf{x}_u - \mathbf{x}_s\| \leq r \end{array} \right. \right\}.$$

By (EP1) and (EP2), every vertex $u \in A_s$ is connected to s with probability $\Omega(1)$. Furthermore, by Equation (12) a random vertex which satisfies (3) has property (2) with probability $\Omega(1)$ as well. Therefore a random vertex satisfies (2) and (3) with probability $\Omega(r^d)$, and by integrating over all weights larger than w_0 it follows that the expected number of vertices in A_s is at least

$$\Omega(n w_{\min}^{\beta-1} w_0^{1-\beta} r^d) = \Omega(w_{\min}^{\beta-2} w_0^{2-\beta} w_s) = \Omega(1).$$

Hence with probability $\Omega(1)$, there exists at least one vertex $u \in \Gamma(s) \cap A_s$. For every $u \in \Gamma(s) \cap A_s$ it holds $\phi(u) \geq \phi(s)$, as the distance to t decreases and the weight increases. Then we take u_1 as the neighbor of s in A_s with maximal objective. By definition of A_s we have $w_{u_1} \geq w_0$.

Since $w_s = O(1)$, the degree distribution of s is a Poisson random variable with rate $\Theta(1)$, and thus with probability $\Omega(1)$, s has no neighbor of weight less than w_0 with higher objective than $\phi(u_1)$. In this case, the routing indeed proceeds with the vertex u_1 and we obtain $\Pr[\mathcal{E}_s] = \Omega(1)$.

Main part of routing process: Our aim is to show that, conditioned on \mathcal{E}_s , greedy routing proceeds to a vertex of objective larger than ϕ_0 . If $\phi(u_1) > \phi_0$ then there is nothing to show, so assume otherwise. We want to show that on the graph induced by the vertex set $V_{\leq \phi_0}$, greedy routing finds a vertex u_ℓ which we expect to have neighbors in $V_{> \phi_0}$. More formally, we consider the event

$$\mathcal{E}_m : \text{the algorithm on } G[V_{\leq \phi_0}] \text{ finds a vertex } u_\ell \text{ with } \mathbb{E}_{> \phi_0}[\|\Gamma(u_\ell) \cap V_{> \phi_0}\|] = \Omega(1).$$

Then by our choice of w_0 and ϕ_0 , Lemma 8.1 (vi) gives the lower bound $\Pr[\mathcal{E}_m \mid \mathcal{E}_s] = \Omega(1)$.

End of routing process: Suppose that either both events \mathcal{E}_s and \mathcal{E}_m hold, or that already $\phi(s) \geq \phi_0$. We want to prove that with constant probability, the routing finds the target t while visiting at most two other vertices of $V_{>\phi}$. Let \bar{w} be a constant weight chosen large enough and let

$$A_t := \{u \in V \mid \mathbf{w}_u \geq \bar{w} \wedge \|\mathbf{x}_u - \mathbf{x}_t\|^d \leq (n\phi_0)^{-1}\}.$$

We claim that the expected number of vertices in the set A_t is $\Omega(1)$. Note that here, we need to condition on the events \mathcal{E}_s and \mathcal{E}_m . However, A_t is independent of \mathcal{E}_m since $A_t \subseteq V_{>\phi_0}$, and conditioning on \mathcal{E}_s only increases the expected size of A_t . Thus

$$\mathbb{E}[|A_t| \mid \mathcal{E}_s \wedge \mathcal{E}_m] \geq \mathbb{E}[|A_t|] = \Omega(nw_{\min}^{\beta-1}(\phi_0)^{\beta-1}(n\phi_0)^{-1}) = \Omega(1).$$

Now, we uncover all remaining vertices, and all edges that involve at least one vertex from $V_{\leq\phi_0}$, but not the edges within $V_{>\phi_0}$. Due to the Poisson point process, both $|A_t|$ and $|\Gamma(u_\ell) \cap V_{>\phi_0}|$ are Poisson random variables with constant mean, therefore with constant probability both sets are non-empty. In this case, there exists a vertex $v_1 \in \Gamma(u_\ell) \cap V_{>\phi_0}$ with maximal objective. Note that if $v_1 = t$, then we are already done. Furthermore, in the case $\alpha < \infty$ we can toss a coin whether the edge $\{v_1, t\}$ is present. By (EP1) it holds $p_{v_1 t} = \Omega(\min\{1, (\phi(v_1)\mathbf{w}_t)^\alpha\}) = \Omega(1)$, which is already sufficient for finding t with constant probability.

It remains the case $\alpha = \infty$. Recall that we already know that there exists at least one vertex $v_3 \in A_t$. We claim that for every $u' \in V_{>\phi_0}$ and $v' \in A_t$ it holds $p_{u'v'} = \Omega(1)$. Indeed, by the triangle inequality we have $\|\mathbf{x}_{u'} - \mathbf{x}_{v'}\| \leq 2 \max\{\|\mathbf{x}_{u'} - \mathbf{x}_t\|, \|\mathbf{x}_{v'} - \mathbf{x}_t\|\}$, and for \bar{w} large enough we deduce

$$\frac{\mathbf{w}_{u'}\mathbf{w}_{v'}}{\mathbf{w}_{\min}n\|\mathbf{x}_{u'} - \mathbf{x}_{v'}\|^d} \geq \frac{1}{2} \max\{\phi(u')\mathbf{w}_{v'}, \phi(v')\mathbf{w}_{u'}\} = \Omega(\bar{w}\phi_0) \geq \frac{1}{c_1},$$

and then by (EP2) it follows $p_{u'v'} = \Omega(1)$. In particular, it also holds $p_{tv'} = \Omega(1)$.

We can finish the proof as follows: We toss the coins for all edges in $V_{>\phi_0}$ except the edges incident to t . By the above observation, with constant probability the edge $\{v_1, v_3\}$ is present. In this case, either the best neighbor v_2 of v_1 satisfies $\phi(v_2) \geq \phi(v_1)$ or we even have $\phi(v_1) \geq \phi(v_3)$ (and put $v_2 = v_1$). We see that

$$\frac{\mathbf{w}_{v_2}\mathbf{w}_t}{\mathbf{w}_{\min}n\|\mathbf{x}_{v_2}\mathbf{x}_t\|^d} = \phi(v_2)\mathbf{w}_t \geq \phi(v_3)\mathbf{w}_t \geq \frac{1}{c_1}.$$

Finally, we uncover the edge $\{v_2, t\}$ and independently of the previous events, this edge is also present with constant probability and with probability $\Omega(1) > 0$, the routing finds t via v_2 . \square

8.3 Proof of Theorem 3.2 (Failure Probability is Exponentially Small in \mathbf{w}_{\min})

In this section, we assume in addition to (EP1) and (EP2) that whenever two vertices u and v satisfy $\|\mathbf{x}_u - \mathbf{x}_v\|^d \leq \frac{c_1\mathbf{w}_u\mathbf{w}_v}{\mathbf{w}_{\min}n}$, then the edge $\{u, v\}$ is present deterministically. This assumption is very natural and in particular satisfied by hyperbolic random graphs. We first prove the second statement of the theorem which is a direct consequence of Lemma 8.1.

Proof of Theorem 3.2 (ii). Let s and t be two vertices such that both $\mathbf{w}_s = \omega(1)$ and $\mathbf{w}_t = \omega(1)$. Then we apply Lemma 8.1 for $w_0 := \mathbf{w}_s$ and $\phi_0 := (c_1\mathbf{w}_t)^{-1}$. By Statement (vi) it follows that with probability

$$1 - O(\mathbf{w}_{\min}^{\beta-2} \min\{\mathbf{w}_s, \mathbf{w}_t\}^{-\Omega(1)}) = 1 - \min\{\mathbf{w}_s, \mathbf{w}_t\}^{-\Omega(1)}$$

the greedy routing finds a vertex u' of objective at least $\phi_0 = (c_1 w_t)^{-1}$. Let P be the final greedy path, and let u'' be its first vertex which has objective at least ϕ_0 . Then either $u'' = t$, or

$$c_1 \frac{w_u'' w_t}{w_{\min} n \|x_u'' - x_t\|^d} \geq c_1 \phi_0 w_t = 1.$$

In this case, $\{u'', t\} \in E$ and thus the routing finds t via u'' . \square

For part (i) of Theorem 3.2, note that the statement becomes trivial if w_{\min} is small. So we can assume that w_{\min} is larger than a given constant c which depends on α, β, d and ε . In particular, this allows us to assume $w_{\min} > w_1(\varepsilon_1)$.

For the proof of Theorem 3.2 (i), we can not directly apply Lemma 8.1 because this yields only a failure probability which is polynomially small in $w_0 \geq w_{\min}$. Instead, we will apply Lemma 7.13 for the very first steps, which will allow us to increase the lower-bound on the success probability. Afterwards, we will apply Lemma 8.1 for values w_0, ϕ_0 which depend exponentially in w_{\min} such that the lemma indeed yields a sufficiently small failure probability. Finally, we need to ensure that after seeing vertices of objective at least ϕ_0 , the routing does not die out and continues until it reaches vertices which connect deterministically to t .

For the first step, we will apply the following lemma.

Lemma 8.2. *Let $w_0 := e^{w_{\min}^{\Omega(1)}}$ and assume that w_{\min} is sufficiently large. Then there exists a constant $c' > 0$ such that with probability*

$$1 - O(e^{-w_{\min}^{\Omega(1)}})$$

the greedy routing visits at least one vertex u which satisfies either (i) $w_u \geq w_0$, (ii) $\phi(u) \geq w_0^{-c'}$, or (iii) $u \in V_2$.

Proof. We use similar techniques as in the proof of Lemma 8.1. Notice that if the starting vertex s satisfies at least one of the three properties, there is nothing to show. Let $\phi_0 := w_0^{-c'}$ and assume that $w_s \leq w_0$, $\phi(s) \leq \phi_0$, and $s \in V_1$. We start by defining a weight sequence $y_{(i)}$ as follows: We put $y_0 := w_s$, and for $i \geq 1$ we put $y_i := y_{i-1}^{\gamma(\zeta\varepsilon_1)}$. Notice that there exists a first y_j such that $y_j \geq w_0$. Moreover, for all $1 \leq i \leq j$ we define

$$A_i := \{v \in V_1 \mid (1) : y_{i-1} \leq w_v < y_i; (2) : \phi(v) \leq \phi_0\}$$

and

$$B_i := \{v \in V_1 \mid (1) : w_v \leq y_i; (2) : \phi(v) \leq \phi_0\}.$$

Furthermore, denote by P_i the greedy path induced by the set B_i and by \mathcal{E}_i the event that either $P_i \cap A_i$ is empty or the first vertex $v \in P_i \cap A_i$ has at least one neighbor $v' \notin B_i$ such that $\phi(v') \geq \phi(v)$ and $\phi(v') \geq \phi_0$ holds for all $u \in \Gamma(v) \cap B_i$. Let $1 \leq i \leq j$. Clearly, we have $\Pr[\mathcal{E}_i \mid P_i \cap A_i = \emptyset] = 1$, and in order to lower-bound $\Pr[\mathcal{E}_i]$ we can assume $P_i \cap A_i \neq \emptyset$. Let v be the first vertex of $P_i \cap A_i$. Then by Lemma 7.11 (i), it holds $\mathbb{E}[|\Gamma(v) \cap V^+(v, \varepsilon_1)|] = \Omega(w_{\min}^{\beta-2} w_v^{\varepsilon_1})$. By a Chernoff bound, with probability at least $1 - \exp(-\Omega(w_{\min}^{\beta-2} y_{i-1}^{\varepsilon_1}))$ there exists a neighbor v' of v in the set $V^+(v, \varepsilon_1)$. In particular $v' \notin B_i$ and $\phi(v') \geq \phi(v) w_v^{\gamma(\varepsilon_1)-1}$.

Next we observe that every vertex $u \in B_i$ with objective at least $\phi(v) w_v^{\gamma(\varepsilon_1)-1}$ is contained in $V^-(v, \varepsilon_1)$. Hence in order to verify that v' is better than every neighbor of v in B_i , it is sufficient that v has no neighbor in $V^-(v, \varepsilon_1)$. We observe that by Lemma 7.6, our conditioning on v being the first vertex in $P_i \cap A_i$ decreases the expected size of $V^-(v, \varepsilon_1)$. Furthermore, for c' large enough

we have $\phi(v)\mathbf{w}_v^{\gamma(\zeta_{\varepsilon_1})} \leq \phi_0\mathbf{w}_0^{\gamma(\zeta_{\varepsilon_1})} \leq \bar{c}$, where \bar{c} is the constant given by Lemma 7.13. Hence by Lemma 7.13, we obtain

$$\mathbb{E}[|\Gamma(v) \cap V^-(v, \varepsilon)|] = O(\phi(v)^{\Omega(1)}\mathbf{w}_v^{O(1)}) = O(\phi_0^{\Omega(1)}y_{i-1}^{O(1)}).$$

By Markov's inequality, with probability $1 - O(\phi_0^{\Omega(1)}y_{i-1}^{O(1)})$ this set is empty and in this case, $\phi(v')$ is higher than the objective of every neighbor of v in B_i . It follows that

$$\Pr[\mathcal{E}_i] \geq 1 - O(\phi_0^{\Omega(1)}y_{i-1}^{O(1)}) - \exp(-\Omega(\mathbf{w}_{\min}^{\beta-2}y_{i-1}^{\varepsilon_1})).$$

Recall that the sequence $y_{(i)}$ grows exponentially, and we obtain

$$\begin{aligned} \Pr\left[\bigwedge_{i=1}^j \mathcal{E}_i\right] &= 1 - O\left(\phi_0^{\Omega(1)}y_0^{O(1)}\right) - \exp\left(-\Omega(\mathbf{w}_{\min}^{\beta-2}y_0^{\varepsilon_1})\right) \\ &= 1 - O\left(w_0^{-c'\cdot\Omega(1)+O(1)} + \exp\left(-\Omega(\mathbf{w}_{\min}^{\beta-2}w_0^{\Omega(1)})\right)\right) \\ &= 1 - O\left(w_0^{-\Omega(1)}\right) \end{aligned}$$

if c' is chosen sufficiently large. This allows us to assume that all good events \mathcal{E}_i occur. Let us consider the final greedy path $P = \{s = v_0, v_1, v_2, \dots\}$ and suppose that all good events \mathcal{E}_i occur. By construction $s \in A_1$, and \mathcal{E}_1 implies that the routing does not die out at s (i.e., s is not isolated) and $v_1 \notin B_1$. Then we distinguish two cases: If $v_1 \notin B_j$, then we are done since v_1 would satisfy (i),(ii) or (iii). Else, v_1 is contained in a layer A_i , and in particular it is also the first vertex of $P_i \cap A_i$. Then we can repeat the argument, apply the event \mathcal{E}_i and see again that there exists a subsequent vertex v_2 on the greedy path which is either located outside B_j or contained in a layer $A_{i'}$ for $i' > i$. We can repeat this argument inductively. However, at the latest when we reach a vertex $v_k \in P \cap A_j$, the event \mathcal{E}_j implies that the best neighbor u of v_k is located outside B_j . Then u must fulfill at least one of the properties (i)-(iii), which proves the lemma. \square

Next, we study the end-phase of the routing and show that once the routing arrives at a vertex of high objective, with sufficiently high probability it continues to a vertex u such that $p_{ut} = \Omega(1)$.

Lemma 8.3. *Let $\phi_0 = e^{-\mathbf{w}_{\min}^{O(1)}}$ and assume that \mathbf{w}_{\min} is sufficiently large. Then if the greedy routing reaches a vertex v with $\phi(v) \geq \phi_0$, with probability*

$$1 - O(e^{-\mathbf{w}_{\min}^{\Omega(1)}})$$

it also visits a vertex u with $\phi(u) \geq (c_1\mathbf{w}_t)^{-1}$.

Proof. If $\phi_0 \geq \phi_1 := \frac{1}{c_1\mathbf{w}_t}$, there is nothing to show, so we can assume $\phi_0 < \phi_1$. In the following we partition the vertex set with objective between ϕ_0 and ϕ_1 in small layers. Let $\psi_0 := \phi_0$, and for $i \geq 1$ put $\psi_i := \psi_{i-1} \cdot \mathbf{w}_{\min}^{(1-1/\gamma(\varepsilon_1))/2}$. We assume w.l.o.g. that there exists an index j such that $\psi_j = \phi_1$. Then for all $1 \leq i \leq j$ we define $A_i := \{v \in V \mid \psi_{i-1} \leq \phi(v) < \psi_i\}$. Our goal is to show that with probability sufficiently high, the greedy routing visits at most one vertex per layer A_i , and if so, it will not stop at such a vertex. Therefore, we denote by P_i the greedy path induced by all vertices of objective less than ψ_i and by \mathcal{E}_i the event that either $P_i \cap A_i$ is empty or the first vertex $v \in P_i \cap A_i$ has a neighbor v' such that $\phi(v') \geq \psi_i$.

Let $1 \leq i \leq j$. If $P_i \cap A_i = \emptyset$, then \mathcal{E}_i holds with probability 1. Hence suppose that $P_i \cap A_i$ is non-empty and let v be the first vertex on $P_i \cap A_i$. We distinguish two cases: First suppose $v \in V_1$. Note that for ε_1 small enough it holds

$$\gamma(\varepsilon_1) - 1 = \frac{3 - \beta - \varepsilon_1}{\beta - 2} \geq \frac{3 - \beta - \varepsilon_1}{2 - 2\varepsilon} = \frac{1 - 1/\gamma(\varepsilon_1)}{2}.$$

Then by Lemma 7.11 (i) and a Chernoff bound, with probability at least $1 - \exp(-\Omega(w_{\min}^{\Omega(1)}))$ there exists a neighbor v' of v with objective

$$\phi(v') \geq \phi(v)w_v^{\gamma(\varepsilon_1)-1} \geq \phi(v)w_{\min}^{\gamma(\varepsilon_1)-1} \geq \phi(v)w_{\min}^{(1-1/\gamma(\varepsilon_1))/2} \geq \psi_i.$$

It remains the case $v \in V_2$. Since $w_t \geq w_{\min}$ and since we are assuming that w_{\min} is sufficiently large, it holds $\phi(v) \leq (c_1 w_t)^{-1} \leq (c_1 w_{\min})^{-1} \leq 1$. This allows us to apply Lemma 7.12 (i), and we see that the expected number of neighbors of v with objective at least $\phi(v)^{1/\gamma(\varepsilon_1)}$ is at least $\Omega(w_{\min}^{\beta-2} \phi(v)^{-\Omega(\varepsilon_1)}) = \Omega(w_{\min}^{\Omega(\varepsilon_1)})$. By a Chernoff bound, the probability that such a neighbor v' exists is at least $1 - \exp(-\Omega(w_{\min}^{\Omega(1)}))$. By definition of $V^+(v, \varepsilon_1)$ we have at least $\phi(v') \geq \phi(v)^{1/\gamma(\varepsilon_1)}$. In addition, we can use that $\phi(v) \geq \phi_1$ and w_{\min} is a sufficiently large constant. Then it follows

$$\begin{aligned} \phi(v') &\geq \phi(v)^{1/\gamma(\varepsilon_1)} = \phi(v)^{1-(3-\beta-\varepsilon)/(1-\varepsilon)} \geq \phi(v)(c_1 w_{\min})^{(3-\beta-\varepsilon)/(1-\varepsilon)} \\ &\geq \phi(v)w_{\min}^{(3-\beta-\varepsilon)/(2-2\varepsilon)} = \phi(v)w_{\min}^{(1-1/\gamma(\varepsilon_1))/2} \geq \psi_i. \end{aligned}$$

It follows that $\Pr[\mathcal{E}_i] \geq 1 - O(e^{-w_{\min}^{\Omega(1)}})$ holds for all $1 \leq i \leq j$. Note that by construction we have $\phi_0 w_{\min}^{j(1-1/\gamma(\varepsilon_1))/2} = \phi_1$, which implies $j = O(\frac{|\log \phi_0|}{\log w_{\min}}) = O(\frac{w_{\min}^{O(1)}}{\log w_{\min}})$. By a union bound it follows

$$\Pr \left[\bigwedge_{i=1}^j \mathcal{E}_i \right] = 1 - O(w_{\min}^{O(1)} e^{-w_{\min}^{\Omega(1)}}) = 1 - O(e^{-w_{\min}^{\Omega(1)}}).$$

Therefore we are allowed to assume that all good events occur and finish the proof as follows: Let P be the final greedy path and let $v \in P$ be the first vertex contained in a layer A_i . Then $v \in P_i \cap A_i$, and the event \mathcal{E}_i implies that v has a neighbor v' with higher objective than ψ_i . Then either $\phi(v') \geq \phi_1$ or $v' \in A_{i'}$ for some $i' > i$. We can repeat the argument until we reach the last layer A_j , and it follows that indeed P must contain a vertex of objective at least ϕ_1 . \square

Proof of Theorem 3.2 (i). We prove the statement by combining the two previous lemmas with Lemma 8.1. Let $c' > 0$ be a sufficiently large constant and let P be the resulting greedy path. We put $w_0 := \exp(w_{\min})$ and $\phi_0 := w_0^{-c'}$. First suppose that $\phi(s) \leq \phi_0$ and $s \notin \bar{V}(w_0, \phi_0)$. In this case, we apply Lemma 8.2 and see that if c' is large enough, then with probability

$$1 - O(e^{-w_{\min}^{\Omega(1)}})$$

there exists a vertex $u \in P$ which is either contained in the large set $\bar{V}(w_0, \phi_0)$ or satisfies $\phi(u) > \phi_0$.

Next suppose that there exists a first vertex $u_1 \in P \cap \bar{V}(w_0, \phi_0)$ (in case $s \in \bar{V}(w_0, \phi_0)$, we would put $u_1 = s$). If there exists such a vertex, by Lemma 8.1 with probability

$$1 - O(w_{\min}^{\beta-2} \min\{w_0, \phi_0^{-1}\}^{-\Omega(1)}) = 1 - O(e^{-w_{\min}^{\Omega(1)}})$$

there exists a vertex u_ℓ on the greedy path induced by $V_{\leq \phi_0}$ such that $u_\ell \in \bar{V}(w_0, \phi_0)$ and

$$\mathbb{E}_{> \phi_0} [|\Gamma(u_\ell) \cap V_{> \phi_0}|] = \Omega \left(w_{\min}^{\beta-2} \min\{w_0, \phi_0^{-1}\}^{\Omega(1)} \right) = \Omega \left(e^{w_{\min}^{\Omega(1)}} \right).$$

So far, we uncovered the graph $G[V_{\leq \phi_0}]$. We continue by tossing the coins for all remaining vertices and edges. By a Chernoff bound, with probability $1 - \exp(-\Omega(\exp(w_{\min}^{\Omega(1)})))$ the set $\Gamma(u_\ell) \cap V_{> \phi_0}$ is non-empty, and u_ℓ has a neighbor u' of objective at least ϕ_0 . Then there are two possible situations: If P visits u_ℓ , then the vertex following after u_ℓ in P has objective at least ϕ_0 . If P

does not visit u_ℓ , then there must exist an earlier vertex $u_i \in P$ from which P jumps to $G[V_{\geq \phi_0}]$. In both cases, the routing arrives at vertex u' satisfying $\phi(u') \geq \phi_0$.

Finally suppose that the routing visits a vertex u' of objective at least ϕ_0 or even starts at such a vertex. Then by Lemma 8.3, with probability $1 - O(e^{-w_{\min}^{\Omega(1)}})$ it also visits a vertex u'' such that $\phi(u'') \geq (c_1 w_t)^{-1}$.

By a union bound over all error probabilities it follows that with probability $1 - O(e^{-w_{\min}^{\Omega(1)}})$ the greedy path P arrives at a vertex u'' fulfilling $\phi(u'') \geq (c_1 w_t)^{-1}$. If $u'' = t$, we are done, otherwise we have

$$\|x_{u''} - x_t\|^d = \frac{w_{u''}}{w_{\min} n \phi(u'')} \leq \frac{c_1 w_{u''} w_t}{w_{\min} n}$$

and by our assumption on the model it follows that $\{u'', t\} \in E$ and therefore $t \in P$. \square

8.4 Proof of Theorem 3.3 (Length of Greedy Path)

Before bounding the length of the resulting greedy path, we look at the very first steps of the routing process. With the following lemma, we show for both the basic routing algorithm and its patching variants that they can not visit too many vertices of low weight before reaching a vertex of growing weight. We will use the lemma both for proving Theorem 3.3 and later in Chapter 9 for analyzing patching protocols.

Lemma 8.4. *Let $w_0 = w_0(n) = \omega(1)$ be a function growing in n , and suppose that the starting vertex s satisfies $w_s \leq w_0$ and $\phi(s) \leq e^{-w_0}$. Let A be a routing protocol which satisfies (P1), i.e., it makes greedy choices. Then a.a.s., either A visits at most $O(w_0^4)$ different vertices in total, or after visiting at most $O(w_0^4)$ different vertices, A reaches a vertex of weight at least w_0 .*

Proof. We denote by G' the subgraph induced by all vertices of weight less than w_0 and by G'' the subgraph induced by all vertices of weight less than w_0^3 . Let $r = (w_0^c/n)^{1/d}$, where we take $c > 0$ sufficiently large.

We first prove that a.a.s., during the first few hops the algorithm stays close to the starting position x_s , unless it finds a large-weight vertex very soon. Suppose that we first uncover all vertices of weight less than w_0 and see only G' . Let $s' = v_0, v_1, v_2, \dots$ be the sequence of different vertices visited by algorithm A when we run it on G' (so if a vertex is visited several times, we only list it once). By Lemma 7.14 (ii), with probability $1 - O(w_0^{-5})$ there exists no neighbor of s in G' with geometric distance at least r to x_s , i.e., there is no jump outgoing from s . For the following vertices v_1, v_2, \dots , we argue similarly: If we condition on that there was no outgoing jump on the previous vertices s, v_1, \dots, v_i , then by Lemma 7.6, the probability that there is an outgoing jump from v_{i+1} is even smaller than without conditioning, and by Lemma 7.14 (ii) for v_{i+1} this probability is at most $1 - O(w_0^{-5})$, regardless whether we condition on the same event for the previous vertices or not. We apply a union bound over the first $2w_0^4$ vertices that the routing algorithm finds in G' and see that with probability $1 - O(w_0^{-1})$, we do not see such an outgoing jump during the first $2w_0^4$ steps. Note that if the algorithm visits less vertices in G' , we can apply union bound just over all visited vertices and obtain the same result.

As a next step, let us consider first a w_0^3 -grid as given by Definition 7.7 and second the ball $B(2w_0^4 r, s)$ of radius $2w_0^4 r$ around x_s . This ball has volume $4^d w_0^{4d} r^d = 4^d w_0^{4d+c} n^{-1}$ and intersects $\Theta(w_0^{4d+c-3}) = w_0^{O(1)}$ cells of the w_0^3 -grid. Let \mathcal{C} be the collection of all these cells contained in our ball $B(2w_0^4 r, s)$. The expected number of vertices of weight at most w_0 contained in a single cell $C_i \in \mathcal{C}$ is at most w_0^3 , and by a Chernoff bound, with probability $1 - \exp(-\Omega(w_0^3))$ there are not more than $2w_0^3$ vertices in C_i . The same is true for all $C_i \in \mathcal{C}$, and by a union bound we deduce

that with probability $1 - O(w_0^{O(1)} \exp(-\Omega(w_0^3))) = 1 - o(1)$ every cell $C_i \in \mathcal{C}$ contains at most $2w_0^3$ vertices.

It follows that after visiting $2w_0^4$ different vertices, the algorithm either has escaped from the ball $B(2w_0^4 r, s)$ or it has visited at least w_0^3 different cells. However, from the above calculation we know that a.a.s. none of the first $2w_0^4$ visited vertices has a neighbor inside G' with geometric distance at least r . Thus for all $i \leq 2w_0^4$ it holds

$$\|x_s - x_{v_i}\| \leq 2w_0^4 r.$$

Therefore a.a.s. the routing does not escape the ball $B(2w_0^4 r, s)$ during the first $2w_0^4$ hops, and in this case, the algorithm running on G' either stops before visiting $2w_0^4$ vertices or it visited vertices in at least w_0^3 different cells.

In particular, we can lower-bound the geometric distance to t for all these first vertices. Observe that $\|x_s - x_t\|^d = \frac{w_s}{\phi(s)w_{\min}n} \geq \frac{e^{w_0}}{n}$ as we are assuming $\phi(s) \leq e^{-w_0}$. On the other hand it holds

$$(2w_0^4 r)^d = 2^d w_0^{O(1)} n^{-1} = o(e^{w_0} n^{-1}),$$

and by the triangle inequality it follows that for all $i \leq 2w_0^4$ we have

$$\|x_{v_i} - x_t\|^d \geq \frac{e^{w_0}}{2n}.$$

Now we add the vertices of weights between w_0 and w_0^3 to the graph in order to obtain G'' . Notice that due to the Poisson point process, the distribution of the new vertices is completely independent from what we have observed in G' . The vertices with even higher weight will be uncovered later. Then we rerun the algorithm A on G'' , yielding a sequence s'' of different visited vertices. There are three cases. First, if $|s''| < 2w_0^4$ then we are done, since then on G the algorithm either visits less than $2w_0^4$ different vertices, or it visits a vertex that is not in G'' . Similarly, if $|s''| \geq 2w_0^4$, and among the first $2w_0^4$ vertices in s'' there is at least one with weight at least w_0 , then we are also done: on G the algorithm must either visit a vertex that is not in G'' , or it visits the same first $2w_0^4$ different vertices as in s'' , thus visiting a vertex of weight at least w_0 .

It remains the case where $|s''| \geq 2w_0^4$, and the first $2w_0^4$ vertices in s'' all have weight at most w_0 . We have seen that in this case, it visits vertices in at least w_0^3 different cells of the w_0^3 -grid. By the Bulk Lemma 7.8, we deduce that when adding the heavy vertices of weight at least w_0^3 to the graph, a.a.s. there exists a smallest index $i < 2w_0^4$ such that v_i has a neighbor u with $w_u \geq w_0^3$ which is located in the same cell. We know that a.a.s. $\|x_{v_i} - x_t\|^d \geq \frac{e^{w_0}}{2n}$. Since u is in the same cell as v_i , it holds $\|x_u - x_t\|^d \leq (1 + o(1))\|x_{v_i} - x_t\|^d$ and therefore

$$\phi(u) \geq (1 - o(1))w_0^2 \phi(v_i).$$

On the other hand, we know that a.a.s. v_i has no neighbor u' satisfying both $w'_u < w_0$ and $\|x_{v_i} - x'_u\| \geq r$, therefore every neighbor u'' of v with weight less than w_0 satisfies

$$\|x_{v_i} - x''_u\| < r = o(\|x_{v_i} - x_t\|).$$

Again by the triangle inequality it follows that such a neighbor u'' satisfies

$$\|x_{u''} - x_t\| = (1 - o(1))\|x_{v_i} - x_t\|.$$

We conclude that in this case we have

$$\phi(u'') \leq (1 + o(1))w_0 \phi(v_i)$$

and indeed the objective of the vertex u is large enough such that by (P1) the algorithm will prefer u to all neighbors of v_i of weight at most w_0 . Hence, the neighbor of v_i with maximal objective has weight at least w_0 and at the latest at this point, A visits a vertex of weight at least w_0 . \square

Proof of Theorem 3.3. Let $f_0 = f_0(n) = \omega(1)$ be a growing function. Furthermore let $w_0 = w_0(n) := \max\{\log f_0(n), w_s\}$ and $\phi_0 := \min\{(c_1 w_t)^{-1}, f_0^{-1}\}$. We first assume $\phi(s) \leq \phi_0$. Suppose $w_s \leq w_0$. Then by Lemma 8.4, a.a.s. either greedy routing ends in a dead end after visiting $O(w_0^4) = o(f_0)$ vertices, or one of the first $O(w_0^4)$ vertices has weight at least w_0 . If $w_s \geq w_0$, we already start at such a vertex. Let u_1 be the first visited vertex of weight at least w_0 . Clearly $\phi(u_1) \geq \phi(s)$. Suppose $\phi(u_1) \leq \phi_0$. Then we apply Lemma 8.1 with w_0 and ϕ_0 . By Statements (v) and (vi), a.a.s. after at most

$$L = \frac{1 + o(1)}{|\log \beta - 2|} (\log \log_{w_0} (\phi(s)^{-1}) + \log \log_{\phi_0} (\phi(s)^{-1})) + O(f_0(n))$$

steps, the routing process arrives at a vertex of objective at least ϕ_0 . Notice that by construction, $w_0 \geq w_t$ and $\phi_0^{-1} \geq w_t$, thus L is not larger than the claimed length of the greedy path.

Hence we can assume that the routing reaches a vertex u' such that $\phi(u') \geq \phi_0$. Note that in the case $\phi(s) \geq \phi_0$, we would use $s = u'$. If $\phi_0 = f_0^{-1}$, then we apply Lemma 7.5 to see that a.a.s., there exist at most $O(f_0)$ vertices with objective at least ϕ_0 . Even if the algorithm visits every single vertex with this property, this needs only $O(f_0(n))$ steps. It remains to treat the case $\phi_0^{-1} = c_1 w_t$. Then

$$c_1 \frac{w_{u'} w_t}{w_{\min} n \|\mathbf{x}_{u'} - \mathbf{x}_t\|^d} = c_1 \phi(u') w_t \geq 1,$$

and we observe that every vertex u that the algorithm visits after u' has probability $p_{ut} = \Omega(1)$ to directly connect to the target t . Suppose for contradiction that from now on, the routing process visits more than $f_0(n) = \omega(1)$ additional vertices. Since each of these vertices independently connects to t with constant probability, it follows that with probability $1 - o(1)$ at least one of the f_0 additional vertices connects to t . However, in this case the algorithm indeed goes to t and stops. We conclude that after reaching a vertex of objective at least ϕ_0 , a.a.s. the routing stops before visiting $f_0(n)$ additional vertices. Adding up all steps then proves the second statement of the Theorem. For the first, general bound of $\frac{2+o(1)}{|\log \beta - 2|} \log \log n$, we pick $f_0(n) = o(\log \log n)$. Then this statement follows immediately since $\phi(s) = \Omega(n^{-1})$ and $w_s, w_t = \Omega(1)$. \square

9 Proof: Patching

Similar to our analysis of the basic routing process, we prove Theorem 3.4 in three steps. We first use Lemma 8.4 in order to show that the algorithm soon finds vertices of high weight. Next we can apply Lemma 8.1 for analyzing the process until an objective ϕ_0 is reached. For the end of the routing process, we need a preparatory result. We do this with the following lemma. It is analogue to Lemma 8.4 but considers the very last phase before the algorithm hits the target vertex t .

Lemma 9.1. *There is a constant $c > 0$ such that the following holds. Let $w_0 = w_0(n) = \omega(1)$ be a function growing in n such that $w_0(n) = O(\log \log n)$, and suppose that the target vertex t satisfies $w_t \leq w_0$ and that the start vertex satisfies $\phi(s) \leq w_0^{-c}$. Then a.a.s., either s and t are not in the same component, or there exists a path P of length $O(w_0^4)$ from t to a vertex u with $w_u \geq w_0$ such that every vertex $v \in P$ (including u) fulfills $\phi(v) \geq w_0^{-O(1)}$.*

Proof. Let G' be the subgraph induced by vertices of weight at most w_0 , and suppose that we first uncover this subgraph. Furthermore let C be the connected component of G' which contains t , and let $r := (w_0^c/n)^{1/d}$, where $c > 0$ is a sufficiently large constant. By Lemma 7.14 (ii), with probability $1 - O(w_0^{-5})$ the target vertex t has no neighbor in G' with geometric distance at least r to x_t . Next we explore C by depth first search (choosing vertices in arbitrary order) until either $2w_0^4$ vertices have been visited or C has been completely traversed. Let v be any vertex which has been explored in this process. By Lemma 7.6 and Lemma 7.14 (ii), with probability $1 - O(w_0^{-5})$ the vertex v has no neighbor in G' with geometric distance at least r to x_v . As we explored at most $2w_0^4$ vertices, by a union bound it follows that with probability $1 - O(w_0^{-1}) = 1 - o(1)$, we didn't explore any edge of geometric length at least r . Therefore, every explored vertex v satisfies $\|x_v - x_t\|^d \leq 2w_0^4 r^d = w_0^{O(1)}/n$ and thus $\phi(v) \geq w_0^{-O(1)}$.

We proceed by considering a w_0^3 -grid as given by Definition 7.7, and uncovering the vertices of weight at least w_0 . Similarly as in the proof of Lemma 8.4, we deduce that there are $w_0^{O(1)}$ cells of the grid which are in distance at most $2w_0^4 r$ to x_t , and a.a.s. all of these cells contain at most $2w_0^3$ vertices. It follows that if we explored $2w_0^4$ vertices, then they cover at least w_0 different cells. In this case, by Lemma 7.8 there exists a vertex u and a vertex v such that (i) $w_u \geq w_0$, (ii) u and v are in the same cell, (iii) v is among the explored vertices, and (iv) u and v are connected with an edge. From (ii) it follows in particular that $\phi(u) \geq w_0^{-O(1)}$, and clearly there exists a v - t -path among the explored vertices. Hence we find a u - v - t -path with all desired properties.

Suppose that the DFS stops after exploring less than $2w_0^4$ vertices, and let v be an explored vertex. Recall that $\|x_v - x_t\|^d \leq 2w_0^4 r^d = w_0^{O(1)}/n$. Then by Lemma 7.14 (iii), with probability $1 - O(w_0^{-5})$ there exists no neighbor u of v such that $\phi(u) \leq w_0^{-c}$ and $w_u \geq w_0$, for a suitable choice of c . By a union bound we deduce that a.a.s. this holds for all explored vertices. However, this is only possible if either C is itself the whole connected component of t (and in this case s is not in the same component as t) or there exists a vertex u with $w_u \geq w_0$ and $\phi(u) \geq w_0^{-O(1)}$ which has an explored neighbor v . This yields again an u - v - t path with all desired properties. \square

Proof of Theorem 3.4. We first treat the main case $\phi(s) = o(1)$, and only come to the exceptional case $\phi(s) = \Omega(1)$ at the very end.

The Case $\phi(s) = o(1)$: finding a vertex u_1 of large weight. Let $w_0 = O(\log \log \log n)$. We assume for now that the starting vertex s satisfies $\phi(s) = o(1)$, which happens a.a.s. if s is a random vertex. Note that we may choose w_0 such that $\phi(s) \leq e^{-w_0}$. Furthermore let A be an algorithm satisfying (P1) – (P3). We first show that either s is not in the same connected component as t , or A finds t quickly, or A visits soon a vertex of weight at least w_0 . If $w_s \geq w_0$, this is trivial. Otherwise, either A finds t among the first $O(w_0^4)$ (different) visited vertices, or the connected component of s has only $O(w_0^4)$ vertices, or the algorithm visits $O(w_0^4)$ without finding t . Note that by condition (P2), the algorithm A spends at most $w_0^{O(1)} = o(\log \log n)$ steps in this phase, which settles the first two cases. In the last case, A visits $O(w_0^4)$ vertices without finding t , and by Lemma 8.4 a.a.s. at least one of these vertices has weight at least w_0 . Let us call the first such vertex u_1 .

The length of middle part: application of the main lemma. As a next step, we want to apply Lemma 8.1 to bound the number of steps in the middle phase. Let $f_1(n) = \omega(1)$ be a function which grows sufficiently small in n , and let $w_1 := \log w_0$, $w_2 := \log w_1 = \log \log w_0$, and $\phi_0 := (w_2)^{-1/f_1} = o(1)$. We consider the induced graph $G[V_{\leq \phi_0}]$ of vertices of objective at most ϕ_0 , and denote by $P = \{s, \dots, u_1, \dots\}$ the path obtained from running A on $G[V_{\leq \phi_0}]$. We

apply Lemma 8.1 for w_0 and ϕ_0 . It follows that a.a.s., after u_1 the patching algorithm always finds an objective-improving neighbor until it reaches a vertex u_ℓ which has a.a.s. $\omega(1)$ neighbors in $V^+(u_\ell, \varepsilon_1)$ of objective at least ϕ_0 . Let P' be the subpath $\{u_1, \dots, u_\ell\}$. By Statement (v) of Lemma 8.1, and by our choice of w_0 and ϕ_0 , a.a.s. it holds

$$|P'| \leq \frac{1 + o(1)}{|\log(\beta - 2)|} (\log \log_{w_0} \phi(u_1)^{-1} + \log \log_{\phi_0^{-1}} \phi(u_1)^{-1}) + o(\log \log n) \leq \frac{2 + o(1)}{|\log \beta - 2|} \log \log n.$$

Existence of u_j, u_i : vertices of objective $\phi(u_j) \approx w_1^{-1}, \phi(u_i) \approx w_2^{-1}$. Next, we claim that a.a.s. there exist a vertex $u_j \in P'$ such that

$$w_1^{-\Omega(1)} \geq \phi(u_j) \geq w_1^{-1-\Omega(1)}, \quad (13)$$

and such that A also visits u_j if we let it run on $G \setminus \{t\}$ instead of $G[V_{\leq \phi_0}]$.

By Lemma 7.14 (i), there exists a constant $\varepsilon' > 0$ such that a.a.s. the graph has no vertices of weight at least w_0 and objective at least $w_0^{-\varepsilon'}$ (except potentially t , if we fix t and let w_t be large). Moreover, we already know that the vertices which A visited before u_1 have no neighbor with distance at least r and weight less than w_0 , thus for ε' small enough, they have no neighbor of objective at least $w_0^{-\varepsilon'} < w_1^{-1-\Omega(1)}$. In particular it follows that $\phi(u_1) \leq w_1^{-1-\Omega(1)}$.

Now let $u \in P'$ be a vertex such that $\phi(u) \leq w_1^{-1-\Omega(1)}$. Then by Lemma 7.14 (iv), with probability $1 - O(\min\{w_u, \phi(u)^{-1}\}^{-\Omega(1)})$ the vertex u has no neighbor of objective at least $w_1^{-\Omega(1)}$ and weight at most w_1 . Moreover, by Lemma 8.1 (iii) the weights of the vertices in P' increase doubly exponentially during the first phase, and by Lemma 8.1 (iv) the objectives increase doubly exponentially during the second phase. This allows us to apply union bound over all vertices on P' with objective at most $w_1^{-1-\Omega(1)}$. We deduce that a.a.s., no vertex $u \in P'$ with $\phi(u) \leq w_1^{-1-\Omega(1)}$ has a neighbor of objective at least $w_1^{-\Omega(1)}$ and weight at most w_1 . By Lemma 7.14 (i), we know that a.a.s. there are no vertices except t with objective at least $w_1^{-\Omega(1)}$ and weight at least w_1 . We finally see that indeed, a.a.s. the path P' contains at least one vertex u_j satisfying (13). Moreover, for sufficiently large n we have $w_1^{-\Omega(1)} \leq \phi_0$. Then, since none of the preceding vertices has a neighbor in $V_{> \phi_0}$, this vertex u_j is also visited if A runs on G . This proves that u_j exists.

Let $w_2 = \log(w_1)$. We claim that analogously to (13), a.a.s. there exist a vertex $u_i \in P'$ such that

$$w_2^{-\Omega(1)} \geq \phi(u_i) \geq w_2^{-1-\Omega(1)}, \quad (14)$$

and such that A also visits u_i if we let it run on $G \setminus \{t\}$ instead of $G[V_{\leq \phi_0}]$. Indeed, the proof is completely analogous to (13).

End phase: finding a path from u_i to t in $G[V_{> \phi(u_j)}]$. To finish the argument, we consider the graph $G[V_{> \phi(u_j)}]$. Clearly, if t is already connected to a vertex on the path P which is located before u_i on P , then the algorithm succeeds. Otherwise, we claim that it is sufficient to show that either t is not in same connected component as s , or that among all vertices of objective at least $\phi(u_j)$ there exists a path from u_j to t . Suppose we find such a path. Then u_j and t are in the same connected component S in $G[V_{> \phi(u_j)}]$. In particular, by (P3) the algorithm A visits all vertices of S (including t) during the next $|S|^{O(1)}$ steps. By Lemma 7.5, a.a.s. there are at most $O(\phi(u_j)^{-1}) = w_0^{O(1)}$ vertices of objective at least $\phi(u_j)$, therefore $|S|^{O(1)} = w_0^{O(1)} = o(\log \log n)$ as desired. Hence it remains to find a u_j - t -path in $G[V_{> \phi(u_j)}]$. Since we already know that u_i and u_j are connected by P' , it also suffices to find a path from u_i to t in $G[V_{> \phi(u_j)}]$. (Mind the two different indices i, j here.)

We start by finding a path from t to a vertex u of weight at least $w_u \geq w_2^\varepsilon$, where $\varepsilon > 0$ is a sufficiently small constant. If $w_t \geq w_2^\varepsilon$, this is trivial by setting $u = t$, otherwise we apply Lemma 9.1 with weight w_2^ε . Note that the condition $\phi(s) < w_2^{-c}$ is satisfied for sufficiently large n . Then either s and t are in different components, or there is a path from t to a vertex u with $w_u \geq w_2^\varepsilon$ such that every vertex in the path has objective at least $w_2^{-O(1)} \geq \phi(u_j)$, where the latter inequality holds for sufficiently large n . So it only remains to show that u and u_i are in the same connected component of $G[V_{>\phi(u_j)}]$. For this last step, fix a sufficiently large constant $C > 0$ and consider the set $B := \{v \in V \mid \|x_v - x_t\|^d \leq w_2^C/n\}$. Clearly, B contains t . If $u \neq t$, then by Lemma 7.14 (i), a.a.s. we have $\phi(u) \leq w_u^{-\Omega(1)}$, and thus $w_2^{-O(1)} \leq \phi(u) \leq w_u^{-\Omega(1)} \leq w_2^{-\Omega(1)}$. This bounds both w_u and $\phi(u)^{-1}$, so we get $\|x_u - x_t\|^d = \phi(u)^{-1}w_u/n \leq w_2^{O(1)}/n$. Hence, if C is large enough then $u \in B$. Similarly, using (14) we find that $u_i \in B$. The set B contains in expectation $n' = w_2^C$ vertices, and by Corollary 7.10 a.a.s. the induced graph $G[B]$ has a giant component which contains all vertices of weight at least $(\log n')^{O(1)}$. In particular, u and u_i are in the giant component of $G[B]$, so there is a path from u to u_i in $G[B]$. Moreover, every vertex v in B has objective at least $\phi(v) \geq w_u/(n\|x_v - x_t\|^d) \geq w_{\min}w_2^{-C} > \phi(u_j)$ (if n is sufficiently large), so $B \subseteq V_{>\phi(u_j)}$. Altogether, we have found a path from u_i to u in $V_{>\phi(u_j)}$, and thus a path from u_j to t in $V_{>\phi(u_j)}$. This concludes the proof under the initial assumption $\phi(s) = o(1)$.

The Case $\phi(s) = \Omega(1)$. In the remaining case $\phi(s) = \Omega(1)$, condition (P3) ensures that there exists an objective $\phi_1 = o(1)$, $\phi_1 \geq (\log \log n)^{o(1)}$, such that we explore the connected component S of s in $G[V_{>\phi_1}]$ in $|S|^{O(1)}$ steps. We note that by Lemma 7.5, a.a.s. $|V_{>\phi_1}| = O(\phi_1^{-1})$, so in particular we need at most $|S|^{O(1)} = o(\log \log n)$ steps to explore S . It only remains to show that a.a.s. either s and t are in different connected components in G , or they are in the same component in $G[V_{>\phi_1}]$. This argument resembles closely the argument in the previous paragraph, and we only repeat it briefly. By Lemma 9.1 either s and t are in different components in G , or there is path from s to t among vertices of objective at least ϕ_1 and of weight at most $\phi_1^{-\varepsilon}$, or in $G[V_{>\phi_1}]$ there are paths from s and from t to vertices u_s and u_t of weights at least $\phi_1^{-\varepsilon}$, where $u_s = s$ or $u_t = t$ are possible if w_s or w_t is large. Then by the same argument as we used for u and u_j above, a.a.s. u_s and u_t are in the same connected component of $G[V_{>\phi_1}]$. Summarizing, either there is no path from s to t in G , or there is a path from s to t in $G[V_{>\phi_1}]$, and in the latter case the path is explored in $o(\log \log n)$ steps, as required. □

10 Proof: Relaxations

In our analysis of the basic greedy algorithm and the patching algorithm we always assumed that a node is aware of the exact objectives of its neighbors and packets are forwarded always to the neighbor with largest objective. In practice the exact objective of a neighbor might not be available. In the following we show that all our theorems also hold if objective values are only approximately known. In particular, we generalize from our original objective function ϕ to a class of objective functions, and prove Theorem 3.5. Before we come to the proof, we remark without formal proof that Theorem 3.5 is essentially best possible.

Remark 10.1. *The relaxations in Theorem 3.5 are best possible. In Section 6 we have described typical trajectories of greedy routing. In particular, it turned out that there are two phases in which the weight and the objective improve by an exponent of $1/(\beta - 2)$, respectively. If we would replace the function $f_0(n)$ in Theorem 3.5 by a constant, then in each step the algorithm could pick vertices*

which would increase the weight (in the first phase) and the objective (in the second phase) only by an exponent $c < 1/(\beta - 2)$. This would increase the number of steps to at least $\frac{2+o(1)}{\log c} \log \log n$, thus making the algorithm considerably slower.

Below we show that our analysis of the greedy routing w.r.t. ϕ applies as well for such a relaxed objective function $\tilde{\phi}$ which fulfills the preconditions of Theorem 3.5. Note that all technical definitions of sets such as $V_1, V_2, V^+(v, \varepsilon), V^-(v, \varepsilon), \dots$ are still defined via ϕ and not via $\tilde{\phi}$.

Proof of Theorem 3.5. We first study how Theorems 3.1, 3.2, and 3.3 for the basic routing process can be transferred for relaxed objective functions, and defer the corresponding analysis for patching algorithms to the end of this section. Let $f_0(n) = \omega(1)$ be a growing function, let $\delta > 0$, and recall from our basic analysis that ε_1 is a constant chosen sufficiently small. Moreover we put $\varepsilon_2 = (\log \log f_0(n))^{-1}$. We want to prove for every $g_0(n) = o(1)$ that the results about greedy routing w.r.t. ϕ transfer to every relaxed objective function $\tilde{\phi}$ which is maximized at the target t and satisfies

$$\tilde{\phi}(v) = \Theta(\phi(v) \cdot \min\{w_v, \phi(v)^{-1}\}^{\pm g(n)})$$

for all v with $\phi(v) \leq \bar{c}w_t^{-1+\delta}$, where $\bar{c} > 0$ is a constant chosen sufficiently large later. Furthermore we assume that all other vertices satisfy $\tilde{\phi}(v) = \Omega(w_t^{-1+\delta/2})$. We observe that by choosing $f_0(n)$ sufficiently small compared to $g(n)$, we can assume as well that every vertex v with $\phi(v) \leq \bar{c}w_t^{-1+\delta}$ satisfies

$$\tilde{\phi}(v) = \Theta(\phi(v) \cdot \min\{w_v^{\gamma(\varepsilon_1)}, \phi(v)^{-1}\}^{\pm \varepsilon_2/c}) \quad (15)$$

for $c = 4\gamma(\varepsilon_1)$, which is a weaker relaxation and thus gives a stronger statement. Notice that for vertices $v \in V_1$, the minimum will be attained by $w_v^{\gamma(\varepsilon_1)}$, and vice-versa for vertices $v \in V_2$, the maximum will be attained by $\phi(v)^{-1}$.

Let \tilde{A} be a greedy routing algorithm which performs routing w.r.t. $\tilde{\phi}$. In the following, we show that we can modify our analysis of the routing process such that whenever \tilde{A} proceeds to a new vertex v , this vertex still has all desired properties and the routing trajectory is essentially the same. We will describe how the proofs of Theorem 3.1, 3.2, and 3.3 can be adapted for the relaxed routing algorithm \tilde{A} , but we will not prove this in detail. Recall that all these proofs basically contain three parts: First we studied the start of the routing, which contains the very first steps until a certain weight w_0 is reached. Then in the second step we applied Lemma 8.1 in order to reach a vertex of objective at least ϕ_0 , and in the final part we studied how the process finds t . When describing how our results can be transferred to the relaxed algorithm \tilde{A} , we therefore handle these three parts separately. Thereby, we assume that $\phi_0 \leq \underline{c}w_t^{-1+\delta/2}$, where $\underline{c} > 0$ is a constant chosen sufficiently small later. Below we will see that never a larger ϕ_0 for transferring the main theorems to \tilde{A} is needed. We start by studying the main part of the routing process.

Main part of routing process: In all proofs we applied Lemma 8.1 for analyzing the main part of the routing process. Hence we indicate why this main lemma still applies for the relaxed algorithm \tilde{A} . In fact, we will slightly strengthen part (vi) of the lemma such that for the end phase we can still guarantee that \tilde{A} reaches a vertex of real objective at least ϕ_0 . Recall that we put $\phi_0 \leq \underline{c}w_t^{-1+\delta/2}$. We require that \underline{c} is small enough such that for every vertex $v \in V_{\leq \phi_0}$ and every vertex u with $\phi(u) \geq \bar{c}w_t^{-1+\delta}$, it holds $\tilde{\phi}(v) < \tilde{\phi}(u)$. Notice that by the weak relaxation (3), indeed this is possible. Then we can assume that whenever we use this lemma, all vertices satisfy (15).

Lemma 8.1 treats two phases of the routing. In the first phase of the routing process, our goal was to prove for a given vertex $v \in V_1$ that its best neighbor w.r.t. ϕ has weight at least $w_v^{\gamma(\varepsilon)}$, where $\varepsilon \in \{\varepsilon_1, \varepsilon_2\}$ and this choice was depending on the weight w_v . Recall that our analysis

assumed that all good events occur. Then in particular v has no neighbor in the set $V^-(v, \varepsilon)$, and every neighbor u with weight at most $w_v^{\gamma(\zeta\varepsilon)}$ can have objective at most $\phi(v)w_v^{\gamma(\varepsilon)-1}$. Let u be such a neighbor. We claim that $\tilde{\phi}(u)$ is small enough such that u is not the best neighbor of v w.r.t. $\tilde{\phi}$. Due to the relaxation (15), the vertex u satisfies

$$\tilde{\phi}(u) \leq \phi(u)w_u^{\gamma(\varepsilon_1)\varepsilon_2/c} \leq O(\phi(v)w_v^{\gamma(\varepsilon)-1+\gamma(\zeta\varepsilon)\gamma(\varepsilon_1)\varepsilon_2/c}).$$

In order to guarantee that the routing does not proceed with such a vertex, we need to find a neighbor of weight at least $w_v^{\gamma(\zeta\varepsilon)}$ and higher relaxed objective $\tilde{\phi}$. Before, we found such good neighbors in the set $V^+(v, \varepsilon)$. Instead, we could as well consider the set $V^+(v, \varepsilon/3)$. We apply Lemma 7.11 to see that

$$\mathbb{E}[|\Gamma(v) \cap V^+(v, \varepsilon/3)|] = \Omega(w_{\min}^{\beta-2} w_v^{\Omega(\varepsilon)})$$

which is fine for all error bounds. Let $\gamma' := \frac{1+\varepsilon/3}{\beta-2}$. By Lemma 7.4, the expected number of neighbors with weight larger than $w_v^{\gamma'}$ is $O(w_{\min}^{\beta-2} w_v^{-\Omega(\varepsilon)})$. Hence we expect $\Omega(w_{\min}^{\beta-2} w_v^{\Omega(\varepsilon)})$ neighbors in $V^+(v, \varepsilon/3)$ with weight at most $w_v^{(1+\varepsilon/3)/(\beta-2)}$. By our relaxations, a vertex u' in this set satisfies

$$\tilde{\phi}(u') \geq \phi(u')w_u'^{-\gamma(\varepsilon_1)\varepsilon_2/c} \geq \Omega(\phi(v)w_v^{\gamma(\varepsilon_1/3)-1-\gamma'\gamma(\varepsilon_1)\varepsilon_2/c}).$$

Since we assume that ε_1 is small enough and $\varepsilon \leq \varepsilon_1$, we observe that

$$\gamma(\varepsilon/3) - \gamma(\varepsilon) = \frac{2\varepsilon}{3(\beta-2)} > \frac{(2+\varepsilon-\zeta\varepsilon)\varepsilon_2}{4(\beta-2)} = (\gamma(\zeta\varepsilon) + \gamma') \frac{\gamma(\varepsilon_1)\varepsilon_2}{c}$$

and indeed, $\tilde{\phi}(u')$ is larger than the relaxed objective of v itself and of every neighbor of v with weight less than $w_v^{\gamma(\zeta\varepsilon)}$. It follows that our analysis of the first phase is valid for routing w.r.t. $\tilde{\phi}$ as well and thus applies to \tilde{A} .

For the second phase in V_2 , we can apply the same arguments in order to generalize our analysis. For $v \in V_2$, we consider again $V^+(v, \varepsilon/3)$ instead of $V^+(v, \varepsilon)$. Given all good events, the set $V^-(v, \varepsilon)$ is again empty and every vertex $u \in \Gamma(v) \cap V_1$ has relaxed objective at most $\tilde{\phi}(u) = O(\phi(v)^{(1-\varepsilon_2/c)/\gamma(\varepsilon)})$. However, every vertex $u' \in \Gamma(v) \cap V^+(v, \varepsilon/3)$ with $\phi(u') \leq \bar{c}w_t^{1-\delta}$ satisfies $\tilde{\phi}(u') = \Omega(\phi(v)^{(1+\varepsilon_2/c)/\gamma(\varepsilon/3)})$ by the first condition on the relaxation $\tilde{\phi}$. Note that since $c > 4$ and $\varepsilon_2 \leq \varepsilon$, we have $\varepsilon_2/c \leq \varepsilon/3$. With a short calculation we deduce $(1+\varepsilon_2/c)(1-\varepsilon) < (1-\varepsilon_2/c)(1-\varepsilon/3)$, which is equivalent to

$$\frac{1+\varepsilon_2/c}{\gamma(\varepsilon/3)} < \frac{1-\varepsilon_2/c}{\gamma(\varepsilon)}.$$

Indeed, $\tilde{\phi}(u') > \tilde{\phi}(u)$ and the routing algorithm \tilde{A} favours vertices in $V^+(v, \varepsilon)$ as desired.

Moreover, regarding statement (vi) of the main lemma, the above arguments imply that every vertex $v \in \Gamma(u_\ell) \cap V^+(u_\ell, \varepsilon)$ with $\phi(v) \leq \bar{c}w_t^{-1+\delta}$ satisfies $\tilde{\phi}(v) > \tilde{\phi}(u_\ell)$. Furthermore, for every vertex $v \in \Gamma(u_\ell) \cap V^+(u_\ell, \varepsilon)$ with $\phi(v) > \bar{c}w_t^{-1+\delta}$ the same is true by our choice of ϕ_0 , since $u_\ell \in V_{\leq \phi_0}$. It follows

$$\{v \in \Gamma(u_\ell) \cap V^+(u_\ell, \varepsilon) \cap V_{> \phi_0} \mid \tilde{\phi}(v) > \tilde{\phi}(u_\ell)\} = \Gamma(u_\ell) \cap V^+(u_\ell, \varepsilon) \cap V_{> \phi_0},$$

and using $M := \min\{w_0, \phi_0^{-1}\}$ we obtain

$$\mathbb{E}_{> \phi_0} [|\{v \in \Gamma(u_\ell) \cap V^+(u_\ell, \varepsilon) \cap V_{> \phi_0} \mid \tilde{\phi}(v) > \tilde{\phi}(u_\ell)\}|] = \Omega(w_{\min}^{\beta-2} M^{\Omega(1)}) \quad (16)$$

for a suitable choice of $\varepsilon \in \{\varepsilon_1, \varepsilon_2\}$. Note that the only difference to part (vi) of Lemma 8.1 is the additional condition $\tilde{\phi}(v) > \tilde{\phi}(u_\ell)$. In particular, if \tilde{A} visits u_ℓ and the set described in (16) is

non-empty, then \tilde{A} proceeds to a vertex in $V_{>\phi_0}$ which has in particular higher relaxed objective than u_ℓ . Notice that it is not guaranteed that this next vertex lies in $V^+(u_\ell, \varepsilon)$, because vertices with real objective at least $\underline{c}w_t^{-1+\delta}$ can have arbitrarily large relaxed objective, thus vertices in $V^+(u_\ell, \varepsilon)$ don't need to be the best neighbors of u_ℓ w.r.t. $\tilde{\phi}$.

Start of routing process: We can assume that all considered vertices are contained in the set $V_{\leq\phi_0}$, as the analysis of the end of the routing will handle all other vertices. Let us start by considering the proof of Theorem 3.1. There, we used that if $w_s < w_1(\varepsilon_1)$, then with constant probability s has a neighbor in a superior set A_s . More precisely, A_s contains vertices of weight at least $w_1(\varepsilon_1)$, and the proof required that with constant probability, s has at least one neighbor $u_1 \in A_s$ and no neighbor with smaller weight than w_1 but higher objective than u_1 . We observe that this argument can be transferred to the analysis \tilde{A} as long as $\tilde{\phi}(u_1)$ is large enough and in particular larger than $\tilde{\phi}(s)$. However, due to (15) it holds $\tilde{\phi}(u_1) = \Omega(\phi(u_1)w_{u_1}^{-o(1)}) = \Omega(\phi(s))$. If we make the set A_s slightly smaller and require that the weight of its vertices is by a constant factor larger than $w_1(\varepsilon_1)$, then still with constant probability there exists a vertex $u_1 \in \Gamma(s) \cap A_s$, and $\tilde{\phi}(u_1)$ will be large enough. Then this starting argument applies as well for $\tilde{\phi}$.

The proof of Theorem 3.2 used Lemma 8.2 for the very first steps. There we constructed a set of layers A_i , defined via weights, and proved that with sufficiently high probability, the routing process traverses these layers and thereby the weight of the current vertex increases as desired. By the same arguments as seen above for the main lemma, we can apply Lemma 8.2 for the relaxations.

When proving Theorem 3.3, we used Lemma 8.4 for the starting phase. The crucial argument behind Lemma 8.4 was that after visiting many vertices of weight at most $w_0(n) = \omega(1)$, the routing will arrive at a vertex v which has a neighbor u in the same cell of a w_0 -grid such that $w_u \geq w_0^3$. We observe that due to (15), u still has the property $\tilde{\phi}(u) \geq (1 - o(1))w_0^{2-o(1)}\phi(v)$. On the other hand, every neighbor v' of v which has weight less than w_0 has relaxed objective $\tilde{\phi}(v') \leq (1 + o(1))w_0^{1+o(1)}\phi(v)$. We see that since $w_0 = \omega(1)$, the algorithm \tilde{A} still prefers u to v' . Therefore Lemma 8.4 can be applied as well for studying the start of the routing algorithm \tilde{A} .

End of routing process: We have seen that as long as the visited vertices have real objective at most $\phi_0 \leq \underline{c}w_t^{-1+\delta/2}$, the routing w.r.t. $\tilde{\phi}$ has the same properties as the original routing. It remains to study the end phase where the vertices have higher objective. We start by considering the proof of Theorem 3.2 (ii). There, we used $\phi_0 = \frac{1}{c_1w_t}$. Since $w_t = \omega(1)$ by assumption, we have $w_t^{-1} = o(w_t^{-1+\delta})$. Hence, there is nothing to change. For adapting the proof of Theorem 3.2 (i), we recall that we are assuming that w_{\min} is sufficiently large and thus w_t as well. In particular, we can assume that $(c_1w_t)^{-1} \leq \underline{c}w_t^{-1+\delta/2}$. Then, we slightly change this proof by setting

$$\phi_0 := \min\{(c_1w_t)^{-1}, \exp(-w_{\min}^{c'})\},$$

where $c' > 0$ is the same large constant used previously in the proof of Theorem 3.2 (i). We distinguish two cases: If $\phi_0 = (c_1w_t)^{-1}$, then our choice of ϕ_0 is valid for applying the main lemma as well with relaxations. We also know by (16) that in this case the algorithm \tilde{A} reaches a vertex u satisfying $\phi(u) \geq (c_1w_t)^{-1}$. In this case, $p_{ut} = 1$, and the routing finds u . On the other hand, if $\phi_0 = \exp(-w_{\min}^{c'})$, then it follows that $w_t = O(1)$. In this case, we need the sequence of additional layers defined via ϕ , as given in the proof of the theorem, until we reach real objective $(c_1w_t)^{-1}$. Notice that for all these vertices in-between the relaxation (15) holds. That is, every vertex v with $\phi_0 \leq \phi(v) \leq \frac{1}{c_1w_t}$ satisfies $\tilde{\phi}(v) = \Theta(\phi(v))$. As the two objective functions are very similar in this additional set of layers, all arguments can be transferred to $\tilde{\phi}$, and then Theorem 3.2 holds as well for the algorithm \tilde{A} .

For adapting the proof of Theorem 3.1, we assume again that \underline{c} is chosen sufficiently small and \bar{c} is chosen sufficiently large. Then we distinguish again two cases. If $c_1 \underline{c} w_t^{\delta/2} \geq 1$, we would like to choose $\phi_0 := (c_1 w_t)^{-1}$. First we have $\underline{c} w_t^{-1+\delta/2} \geq (c_1 w_t)^{-1}$, thus our choice is valid in terms of applying the main lemma with relaxations. Second, for \underline{c} small enough it also follows that the error probability given by the main lemma is a sufficiently small constant. Hence we can take ϕ_0 as claimed before, and by the same arguments as above we see that once the routing reaches a vertex u of real objective at least ϕ_0 , it connects to t with constant probability, which is already sufficient for the theorem. On the other hand, if $c_1 \underline{c} w_t^{\delta/2} < 1$, then $w_t = O(1)$. In this case, we choose $\phi_0 = \Omega(1)$ small enough such that (a) ϕ_0 is still a valid choice, i.e., $\phi_0 \leq \underline{c} w_t^{-1+\delta/2}$, and such that (b) the error probability when applying Lemma 8.1 is still small enough. Then we know that \tilde{A} finds at least a vertex u such that $\phi(u) \geq \phi_0 = \Omega(1)$. However, since $c_1 \underline{c} w_t^{\delta/2} < 1$, it follows that all vertices with real objective at most $\bar{c}(c_1 \underline{c})^{2(1-\delta)/\delta}$ satisfy the relaxation (15), which means that in this regime, ϕ and $\tilde{\phi}$ deviate by at most a bounded constant factor. Recall that in the proof of Theorem 3.1, we constructed a superior set A_t . By picking \bar{c} large enough, with constant probability A_t contains vertices which still satisfy (15). Then again, all arguments can be transferred to $\tilde{\phi}$.

For adapting the proof of Theorem 3.3, we slightly modify our choice of ϕ_0 compared to the original proof and use $\phi_0 := \min\{(c_1 w_t)^{-1+\varepsilon_2/c}, f_0^{-1+\varepsilon_2/c}\} = o(1)$. Note that since $\varepsilon_2 = o(1)$, we satisfy $\phi_0 \leq (\underline{c} w_t^{-1+\delta/2})$ and thus the choice is valid. If the minimum is attained with the first term, then as soon as \tilde{A} reaches a vertex of real objective at least ϕ_0 , it will never visit a vertex of real objective less than $(c_1 w_t)^{-1}$ later on, and then the argument from the proof can be adapted. On the other hand, if the minimum is attained with the second term, then again \tilde{A} will stay at vertices of real objective at least f_0^{-1} , once it reached a vertex of real objective at least ϕ_0 . And then again the proof of Theorem 3.3 can be adapted.

Patching algorithms: We first want to explain why the situation for patching algorithms is slightly different than for the basic routing algorithm, and why we can not transfer the proof of Theorem 3.4 in the case where $\tilde{\phi}$ satisfies only the weaker relaxation and t can be an arbitrary vertex. Above, we have seen that if we only assume the weak relaxation (3), then we need to take $\phi_0 \leq \underline{c} w_t^{-1+\delta/2}$ when applying the main lemma. However, our proof of Theorem 3.4 requires that $\phi_0 = o(1)$ falls arbitrarily slow in n . We see that in a model where w_t can be chosen arbitrarily large, our choice of ϕ_0 is no longer valid when considering relaxations. Vice-versa, if t is a random vertex, then a.s. the weight w_t is small enough such that this conflict does not appear. Similarly, if we assume the stronger relaxation (2) for all vertices, then it is not difficult to observe that we get rid of the restriction $\phi_0 \leq \underline{c} w_t^{-1+\delta/2}$ and our choice of ϕ_0 is also valid when considering the relaxations.

For the start, we analyzed patching algorithms by applying Lemma 8.4. Above, we have already seen that this Lemma can be transferred for relaxed objectives $\tilde{\phi}$. For the main part of the routing, we argued before that under the additional assumptions, we still can apply the main lemma. Thus it remains to study the end of the routing process. There, we proved the existence of two specific vertices u_j and u_i on the Greedy path. We required that $\phi(u_j)$ and $\phi(u_i)$ lie in a specific range. We observe that since $u_j, u_i \in V_{\leq \phi_0}$, the stronger relaxation (2) holds in the considered ranges. Therefore, $\tilde{\phi}$ deviates only by an $o(1)$ -exponent from ϕ , and we don't need to modify the arguments. Finally, we concluded the proof by finding a path from u_i to t inside a certain subset B . Clearly, this path is still present in the graph, and again the relaxations ensure that $B \subset V_{> \phi(u_j)}$. This allows us to transfer the analysis of patching algorithms to relaxed objective functions $\tilde{\phi}$. \square

11 Proof: Greedy Routing on Hyperbolic Random Graphs

Hyperbolic random graphs were introduced by Krioukov et al. [43] and have attracted a lot of attention during the last years, both in theoretical and experimental studies. Let us give a short formal explanation of the model, following the notation introduced by Gugelmann et al. [30]. It can be described by a disk H of radius R around the origin 0 , where the position of every point x is given by its polar coordinates (r_x, ν_x) . The model is isotropic around the origin. Then the hyperbolic distance between two points x and y is given by the non-negative solution $d_H = d_H(x, y)$ of the equation

$$\cosh(d_H) = \cosh(r_x) \cosh(r_y) - \sinh(r_x) \sinh(r_y) \cos(\nu_x - \nu_y).$$

Then hyperbolic random graphs are defined as follows.

Definition 11.1. *Let $\alpha_H, T_H > 0, C_H \in \mathbb{R}, n \in \mathbb{N}$, and set $R := 2 \log n + C_H$. The hyperbolic random graph $G_{\alpha_H, C_H, T_H}(n)$ is a distribution on graphs with vertex set $V = [n]$ s.t.*

- *Every vertex $v \in [n]$ independently draws random coordinates (r_v, ν_v) , where the angle ν_v is chosen uniformly at random in $[0, 2\pi)$ and the radius $r_v \in [0, R]$ is random with density $f(r) := \frac{\alpha_H \sinh(\alpha_H r)}{\cosh(\alpha_H R) - 1}$.*

- *Every potential edge $e = \{u, v\}$, $u, v \in [n]$, is independently present with probability*

$$p_H(d_H(u, v)) = \left(1 + e^{\frac{1}{2T_H}(d_H(u, v) - R)}\right)^{-1}.$$

In the limit $T_H \rightarrow 0$, we obtain the threshold hyperbolic random graph $G_{\alpha_H, C_H}(n)$, where every edge $e = \{u, v\}$ is present if and only if $d_H(u, v) \leq R$.

In [12, Theorem 6.3] it was shown that such graphs are a special case of GIRGs. However, the mapping is a bit counterintuitive, so we describe it briefly here. In particular, although hyperbolic random graphs are defined on a two-dimensional hyperbolic disc, they give rise to a one-dimensional GIRG. The reason is that in GIRGs the weights provide one additional dimension. Notice that a single point on the hyperbolic disk has measure zero, so we can assume that no vertex has radius $r_v = 0$. For the parameters, we put

$$d := 1, \quad \beta := 2\alpha_H + 1, \quad \alpha := 1/T_H, \quad \mathbf{w}_{\min} := e^{-C_H/2}.$$

Then the mapping that transforms hyperbolic random graphs into the GIRG framework is given by

$$\mathbf{w}_v := ne^{-r_v/2} \quad \text{and} \quad \mathbf{x}_v := \frac{\nu_v}{2\pi}.$$

Since this is a bijection between $H \setminus \{0\}$ and $[1, e^{R/2}) \times \mathbb{T}^1$, there exists an inverse function $g(\mathbf{w}_u, \mathbf{x}_u) = (r_u, \nu_u)$. Finally for any two vertices $u \neq v$ on the torus, we set

$$p_{uv} := p_H(d(g(u), g(v))) = p_H(d_H(g(\mathbf{w}_u, \mathbf{x}_u), g(\mathbf{w}_v, \mathbf{x}_v))).$$

This finishes our embedding. In [12] we proved that it is a special case of GIRGs, in a rigorous sense. Compared to [12], we have slightly changed the embedding such that we do not lose the freedom of choosing \mathbf{w}_{\min} .

Greedy routing on hyperbolic random graphs is completely geometric: From a vertex v , the packet is sent to the neighbor $u \in \Gamma(v)$ which minimizes the (hyperbolic) distance $d_H(u, t)$. Notice

that minimizing $d_H(u, t)$ is equivalent to maximizing $(\cosh(d_H(u, t)))^{-1/2}$, as the function \cosh is monotone increasing on positive values, and we are free to multiply this expression by any terms which are independent of u .

Let us assume that $d = 1$ and that a target vertex t is given. Then on such a GIRG, we can define the objective function

$$\phi_H(v) := \frac{n}{\mathbf{w}_t \mathbf{w}_{\min} \sqrt{\cosh(d_H(g(v), g(t)))}},$$

where $g(v)$ and $g(t)$ are the vertices on the hyperbolic disk that correspond to u and t . Suppose that the one-dimensional GIRG was obtained by embedding a hyperbolic random graph, as described above. Then, maximizing ϕ_H on the GIRG is the same as minimizing the hyperbolic distance to the target on the original hyperbolic graph. Thus routing w.r.t. ϕ_H on the GIRG is equivalent to geometric greedy routing on hyperbolic random graphs. The following lemma states that for $d = 1$, with high probability, for most vertices the objective function ϕ_H differs by at most a constant factor from ϕ .

Lemma 11.2. *Let $d = 1$ and let $\delta > 0$ be a sufficiently small constant. Then with probability $1 - n^{-\Omega(1)}$, for every vertex v with $\phi(v) \leq O(\mathbf{w}_t^{-1+\delta})$ it holds*

$$\phi_H(v) = \Theta(\phi(v)),$$

and for every vertex v with $\phi(v) = \Omega(\mathbf{w}_t^{-1+\delta})$ it holds

$$\phi_H(v) = \Omega(\mathbf{w}_t^{-1+\delta}).$$

The statement of this lemma directly implies Corollary 3.6.

Proof of Corollary 3.6. By Lemma 11.2, with probability $1 - n^{-\Omega(1)}$ the objective function ϕ_H falls into the general class of objective functions considered in Theorem 3.5. Note that the small error probability $n^{-\Omega(1)}$ coming from Lemma 11.2 is negligible compared to the error probabilities of all previous theorems. Then the statement of the corollary follows directly from Theorem 3.5, that is, our results transfer from GIRGs to hyperbolic random graphs. \square

Proof of Lemma 11.2. Let $\delta > 0$ and put $\phi_0 := O(\mathbf{w}_t^{-1+\delta})$. For the proof of this lemma, we assume that there exists no vertex with weight larger than $n^{1-2\delta}$, which happens with probability $1 - n^{-\Omega(1)}$ if δ is chosen small enough. Let $v \in V$ such that $\phi(v) \leq \phi_0$, and denote by $(r_v, \nu_v) := g(\mathbf{w}_v, \mathbf{x}_v)$ and $(r_t, \nu_t) := g(\mathbf{w}_t, \mathbf{x}_t)$ the mappings of v and t onto the hyperbolic disk, and assume without loss of generality that $\nu_t = 0$ and $\nu_v \leq \pi$. We claim that on the hyperbolic disk, the assumption $\phi(v) \leq \phi_0$ implies the geometric property

$$\nu_v \geq e^{-\min\{r_v, r_t\}}. \tag{17}$$

We first show $\nu_v \geq e^{-r_v}$. By assumption $\phi(v) = O(1)$, and then for n large enough our mapping implies

$$\nu_v = 2\pi \|\mathbf{x}_v - \mathbf{x}_t\| = \frac{2\pi \mathbf{w}_v}{\mathbf{w}_{\min} n \phi(v)} \geq \mathbf{w}_v n^{-1-2\delta}.$$

Furthermore $\mathbf{w}_v \leq n^{1-2\delta}$ as there are no vertices of higher weight, and thus

$$\nu_v \geq \mathbf{w}_v n^{-1-2\delta} \geq \mathbf{w}_v^2 n^{-2} = e^{-r_v}.$$

Next, we observe that because there are no vertices of weight higher than $n^{1-2\delta}$, we have $w_t^{1+\delta} \leq n^{(1+\delta)(1-2\delta)} = o(n)$ and thus $w_t^{1-\delta} n^{-1} = \omega(w_t^2 n^{-2})$. On the other hand we are assuming $\phi(v) = O(w_t^{-1+\delta})$. Then for n large enough, we obtain

$$\nu_v = \frac{2\pi w_v}{w_{\min} n \phi(v)} \geq \frac{2\pi}{n \phi(v)} \geq w_t^2 n^{-2} = e^{-r_t}$$

which proves (17).

In the following we estimate the hyperbolic cosine in terms of w_v , w_t and $\|x_v - x_t\|$. We apply the definition of d_h and the identity

$$\cosh(x - y) = \cosh(x) \cosh(y) - \sinh(x) \sinh(y),$$

which yields

$$\begin{aligned} \cosh(d_H(g(v), g(t))) &= \cosh(r_v) \cosh(r_t) - \sinh(r_v) \sinh(r_t) \cos(\nu_v) \\ &= \cosh(r_v - r_t) + (1 - \cos(\nu_v)) \sinh(r_v) \sinh(r_t). \end{aligned}$$

Mapping our assumption $w_v \leq n^{1-2\delta}$ on the hyperbolic geometry gives $r_v \geq 2\delta \log n$, therefore we get $\sinh(r_v) = (1 - o(1))e^{r_v}/2$. Similarly it follows $\sinh(r_t) = (1 - o(1))e^{r_t}/2$. Then

$$\begin{aligned} \cosh(d_H(g(v), g(t))) &= \cosh(r_v - r_t) + (1 - \cos(\nu_v))(1 - o(1))e^{r_v+r_t}/4 \\ &= \cosh(r_v - r_t) + \Theta(\nu_v^2 e^{r_v+r_t}), \end{aligned}$$

where we applied the Taylor approximation $1 - \cos(\nu_v) = \nu_v^2/2 - \nu_v^4/24 + O(\nu_v^6)$. Finally, Equation (17) implies $\cosh(r_v - r_t) = O(\nu_v^2 e^{r_v+r_t})$. It follows

$$\phi_H(v) = \Theta\left(\frac{e^{r_t/2}}{w_{\min} \nu_v e^{(r_v+r_t)/2}}\right) = \Theta\left(\frac{ne^{-r_v/2}}{w_{\min} n \nu_v}\right) = \Theta\left(\frac{w_v}{w_{\min} n \|x_v - x_t\|^d}\right) = \Theta(\phi(v)). \quad (18)$$

This proves the first part of the lemma. For the second part, let v be a vertex such that $\phi(v) \geq \phi_0$. We try to add an artificial vertex u of weight $w_u = w_v$ in the graph such that $\phi(u) = \phi_0$. If this is not possible because w_v is too large, then we place u such that $\|x_u - x_t\| = 1/2$. In both cases u has larger geometric distance to t than v . If $\phi(u) = \phi_0$, we know by the first part of the proof that $\phi_H(u) = \Theta(\phi_0)$. In the second case, $\nu_u = \Omega(1)$ and therefore it holds

$$\cosh(d_H(g(u), g(t))) = \cosh(r_u - r_t) + \Theta(\nu_u^2 e^{r_u+r_t}) = \Theta(\nu_u^2 e^{r_u+r_t}) = \Theta(e^{r_u+r_t}),$$

and by (18) again we obtain $\phi_H(u) = \Theta(\phi(u)) = \Omega(\phi_0)$. Considering $g(v)$ and $g(u)$, we see that $\nu_u \geq \nu_v$. Since the hyperbolic distance is monotone decreasing in ν , it follows

$$d_H(g(v), g(t)) \leq d_H(g(u), g(t))$$

and thus in both cases it holds $\phi_H(v) \geq \phi_H(u) = \Theta(\phi_0)$. □

References

- [1] W. Aiello, A. Bonato, C. Cooper, J. Janssen, and P. Prałat. A spatial web graph model with local influence regions. Internet Mathematics, 5(1-2):175–196, 2008.
- [2] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. Reviews of modern physics, 74(1):47, 2002.
- [3] A.-L. Barabási. Network science: Luck or reason. Nature, 489(7417):507–508, 2012.
- [4] T. Bläsius, T. Friedrich, and A. Krohmer. Hyperbolic random graphs: Separators and treewidth. In European Symposium on Algorithms (ESA), pages 15:1–15:16, 2016.
- [5] T. Bläsius, T. Friedrich, A. Krohmer, and S. Laue. Efficient embedding of scale-free graphs in the hyperbolic plane. In European Symposium on Algorithms (ESA), pages 16:1–16:18, 2016.
- [6] M. Boguñá and D. Krioukov. Navigating Ultrasmall Worlds in Ultrashort Time. Physical Review Letters, 102(5):058701, February 2009.
- [7] M. Boguñá, D. Krioukov, and K. C. Claffy. Navigability of complex networks. Nature Physics, 5:74–80, January 2009.
- [8] M. Boguñá, F. Papadopoulos, and D. Krioukov. Sustaining the Internet with hyperbolic mapping. Nature Communications, 1(6), September 2010.
- [9] B. Bollobás, S. Janson, and O. Riordan. The phase transition in inhomogeneous random graphs. Random Structures & Algorithms, 31(1):3–122, 2007.
- [10] A. Bonato, J. Janssen, and P. Prałat. A geometric model for on-line social networks. In 1st International Workshop on Modeling Social Media (WOSM), 2010.
- [11] M. Bradonjić, A. Hagberg, and A. G. Percus. The structure of geographical threshold graphs. Internet Mathematics, 5(1-2):113–139, 2008.
- [12] K. Bringmann, R. Keusch, and J. Lengler. Geometric inhomogeneous random graphs. Preprint available at arXiv:1511.00576, 2015.
- [13] K. Bringmann, R. Keusch, and J. Lengler. Average distance in a general class of scale-free networks with underlying geometry. arxiv, 2016.
- [14] D. Brockmann, L. Hufnagel, and T. Geisel. The scaling laws of human travel. Nature, 439(7075):462–465, 2006.
- [15] E. Candellero and N. Fountoulakis. Clustering and the hyperbolic geometry of complex networks. In 11th International Workshop on Algorithms and Models for the Web Graph (WAW), pages 1–12, 2014.
- [16] A. Chaintreau, P. Fraigniaud, and E. Lebar. Networks become navigable as nodes move and forget. In International Colloquium on Automata, Languages, and Programming, pages 133–144, 2008.
- [17] F. Chung and L. Lu. The average distances in random graphs with given expected degrees. Proceedings of the National Academy of Sciences (PNAS), 99(25):15879–15882, 2002.

- [18] F. Chung and L. Lu. Connected components in random graphs with given expected degree sequences. Annals of Combinatorics, 6(2):125–145, 2002.
- [19] A. Cvetkovski and M. Crovella. Hyperbolic embedding and routing for dynamic graphs. In INFOCOM 2009. 28th IEEE International Conference on Computer Communications, pages 1647–1655, April 2009.
- [20] M. Deijfen, R. van der Hofstad, and G. Hooghiemstra. Scale-free percolation. Annales de l’Institut Henri Poincaré, Probabilités et Statistiques, 49(3):817–838, 2013.
- [21] P. Deprez, R. S. Hazra, and M. V. Wüthrich. Inhomogeneous long-range percolation for real-life network modeling. Risks, 3(1), 2015.
- [22] S. N. Dorogovtsev and J. F. F. Mendes. Evolution of networks. Advances in Physics, 51(4):1079–1187, 2002.
- [23] L. Ferretti, M. Cortelezzi, and M. Mamino. Duality between preferential attachment and static networks on hyperbolic spaces. EPL (Europhysics Letters), 105(3):38001, 2014.
- [24] N. Fountoulakis. On the evolution of random graphs on spaces of negative curvature. Journal of Complex Networks, 2015. To appear.
- [25] P. Fraigniaud, C. Gavoille, and C. Paul. Eclecticism shrinks even small worlds. Distributed Computing, 18(4):279–291, 2006.
- [26] M. Franceschetti and R. Meester. Navigation in small-world networks: A scale-free continuum model. Journal of Applied Probability, pages 1173–1180, 2006.
- [27] T. Friedrich and A. Krohmer. Cliques in hyperbolic random graphs. In 2015 IEEE Conference on Computer Communication (INFOCOM), pages 1544–1552, 2015.
- [28] T. Friedrich and A. Krohmer. On the diameter of hyperbolic random graphs. In 42nd International Colloquium on Automata, Languages, and Programming (ICALP), Lecture Notes in Computer Science, 2015. To appear.
- [29] G. García-Pérez, M. Boguñá, A. Allard, and M. Á. Serrano. The hidden hyperbolic geometry of international trade: World trade atlas 1870–2013. Scientific Reports, 6, 2016.
- [30] L. Gugelmann, K. Panagiotou, and U. Peter. Random hyperbolic graphs: degree sequence and clustering. In 9th International Colloquium on Automata, Languages, and Programming (ICALP), pages 573–585, 2012.
- [31] K. A. Harras, K. C. Almeroth, and E. M. Belding-Royer. Delay tolerant mobile networks (dtmns): Controlled flooding in sparse mobile networks. In International Conference on Research in Networking, pages 1180–1192. Springer, 2005.
- [32] M. Heydenreich, T. Hulshof, and J. Jorritsma. Structures in supercritical scale-free percolation. Preprint available at [arXiv:1604.08180](https://arxiv.org/abs/1604.08180), 2016.
- [33] W. Huang, S. Chen, and W. Wang. Navigation in spatial networks: A survey. Physica A: Statistical Mechanics and its Applications, 393:132–154, 2014.
- [34] E. Jacob and P. Mörters. A spatial preferential attachment model with local clustering. In Algorithms and Models for the Web Graph, pages 14–25. Springer, 2013.

- [35] M. Kiwi and D. Mitsche. A bound for the diameter of random hyperbolic graphs. In 2015 Proceedings of the Twelfth Workshop on Analytic Algorithmics and Combinatorics (ANALCO), pages 26–39. SIAM, 2015.
- [36] J. Kleinberg. The small-world phenomenon: an algorithm perspective. In 32nd Annual ACM Symposium on Theory of Computing (STOC), pages 163–170, 2000.
- [37] J. Kleinberg. Small-world phenomena and the dynamics of information. Advances in neural information processing systems, 1:431–438, 2002.
- [38] R. Kleinberg. Geographic Routing Using Hyperbolic Space. In INFOCOM 2007. 26th IEEE International Conference on Computer Communications, pages 1902 –1909, May 2007.
- [39] Y.-B. Ko and N. H. Vaidya. Location-aided routing (lar) in mobile ad hoc networks. Wireless networks, 6(4):307–321, 2000.
- [40] C. Koch and J. Lengler. Bootstrap percolation on geometric inhomogeneous random graphs. In 43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, pages 147:1–147:15, 2016.
- [41] D. Krioukov, K. Fall, A. Brady, et al. On compact routing for the internet. ACM SIGCOMM Computer Communication Review, 37(3):41–52, 2007.
- [42] D. Krioukov, F. Papadopoulos, M. Boguñá, and A. Vahdat. Greedy forwarding in scale-free networks embedded in hyperbolic metric spaces. SIGMETRICS Perform. Eval. Rev., 37:15–17, October 2009.
- [43] D. Krioukov, F. Papadopoulos, M. Kitsak, A. Vahdat, and M. Boguñá. Hyperbolic geometry of complex networks. Phys. Rev. E, 82:036106, 2010.
- [44] V. Lehman, A. Gawande, B. Zhang, L. Zhang, R. Aldecoa, D. Krioukov, and L. Wang. An experimental investigation of hyperbolic routing with a smart forwarding plane in ndn. In Quality of Service (IWQoS), 2016 IEEE/ACM 24th International Symposium on, pages 1–10. IEEE, 2016.
- [45] G. S. Manku, M. Naor, and U. Wieder. Know thy neighbor’s neighbor: the power of lookahead in randomized p2p networks. In Proceedings of the thirty-sixth annual ACM Symposium on Theory of Computing, pages 54–63, 2004.
- [46] C. Martel and V. Nguyen. Analyzing kleinberg’s (and other) small-world models. In Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing, pages 179–188, 2004.
- [47] S. Milgram. The small world problem. Psychology Today, 2(1):60–67, 1967.
- [48] M. Mitzenmacher. A brief history of generative models for power law and lognormal distributions. Internet mathematics, 1(2):226–251, 2004.
- [49] F. Papadopoulos, D. Krioukov, M. Boguñá, and A. Vahdat. Greedy forwarding in dynamic scale-free networks embedded in hyperbolic metric spaces. In INFOCOM 2010. 29th IEEE International Conference on Computer Communications, pages 1 –9, March 2010.
- [50] F. Papadopoulos, C. Psomas, and D. Krioukov. Network mapping by replaying hyperbolic growth. Networking, IEEE/ACM Transactions on, 23(1):198–211, 2015.

- [51] M. Penrose. Random geometric graphs, volume 5. Oxford University Press Oxford, 2003.
- [52] U. Peter. Random Graph Models for Complex Systems. PhD thesis, ETH Zurich, 2014.
- [53] J. Postel. Simple mail transfer protocol. Information Sciences, 1982.
- [54] S. Sahhaf, W. Tavernier, D. Colle, M. Pickavet, and P. Demeester. Link failure recovery technique for greedy routing in the hyperbolic plane. Computer Communications, 36(6):698–707, 2013.
- [55] S. Schnetzler. A structured overview of 50 years of small-world research. Social Networks, 31(3):165–178, 2009.
- [56] Y. Shavitt and T. Tankel. Hyperbolic embedding of internet graph for distance estimation and overlay construction. IEEE/ACM Transactions on Networking (TON), 16(1):25–36, 2008.
- [57] R. L. Streit. Poisson Point Processes: Imaging, Tracking, and Sensing. Springer, 2010.
- [58] H. P. Thadakamalla, R. Albert, and S. R. T. Kumara. Search in spatial scale-free networks. New Journal of Physics, 9(6):190, 2007.
- [59] J. Travers and S. Milgram. An experimental study of the small world problem. Sociometry, pages 425–443, 1969.
- [60] D. J. Watts. The “new” science of networks. Annual Review of Sociology, 30:243–270, 2004.
- [61] D. J. Watts, P. S. Dodds, and M. E. J. Newman. Identity and search in social networks. science, 296(5571):1302–1305, 2002.