

Geometric Routing in Geometric Inhomogeneous Random Graphs

Yu-Cheng Chiu ✉ 

Department of Mathematics, ETH Zürich, Switzerland

Marc Kaufmann ✉ 

Department of Computer Science, ETH Zürich, Switzerland

Kostas Lakis ✉ 

Department of Computer Science, ETH Zürich, Switzerland

Ulysse Schaller ✉ 

Department of Computer Science, ETH Zürich, Switzerland

Abstract

We present the first rigorous analysis of decentralized geometric routing in Geometric Inhomogeneous Random Graphs (GIRGs), a weight-agnostic variant of the greedy routing protocol. While greedy routing in GIRGs is known to explain the *algorithmic small-world phenomenon* by finding ultra-short paths of length $\Theta(\log \log n)$, it assumes additional knowledge of vertex weights beyond geometry, an assumption that is often restrictive or unavailable. We investigate whether the underlying geometry alone is sufficient for efficient navigation. We prove that for power-law weight exponent $\tau \in (2, 3)$ and geometric decay parameter $\alpha > \tau - 1$, geometric routing succeeds with constant probability and finds ultra-short paths of length $\Theta(\log \log n)$, matching the optimal asymptotic guarantees for greedy routing. Our analysis further reveals that, upon success, both protocols follow a similar two-phase trajectory, consisting of a rapid ascent to the heavy vertices, followed by efficient navigation to the target. These results demonstrate that, in the appropriate regime, the network’s geometry alone implicitly guides the path to the target through its high-weight core.

2012 ACM Subject Classification Theory of computation → Random network models; Theory of computation → Graph algorithms analysis; Networks → Routing protocols

Keywords and phrases geometric inhomogeneous random graphs (GIRGs), hyperbolic random graphs (HRGs), greedy routing, geometric routing, navigability, small-world phenomenon, decentralized algorithms

Digital Object Identifier 10.4230/LIPIcs.WG.2026.12

Related Version Full Version: <https://arxiv.org/abs/2606.02856>

Funding Marc Kaufmann: Swiss National Science Foundation grant number 200021_192079.

Kostas Lakis: Swiss National Science Foundation grant number 10003390.

Ulysse Schaller: Swiss National Science Foundation grant number 200021_192079.

Acknowledgements We would like to thank Johannes Lengler for suggesting this problem to us and for the insightful conversations about it.

1 Introduction

The phenomenon that social networks contain surprisingly short paths, and admit decentralized navigation, was popularized by Milgram’s “small-world” experiment [36, 43]. This motivated the study of decentralized routing algorithms on random graph models of complex networks. Kleinberg’s seminal work [28, 29] provided the first theoretical explanation of the algorithmic small-world phenomenon by showing that greedy routing succeeds in $\mathcal{O}(\log^2 n)$ steps on augmented grids. Despite its influence, Kleinberg’s model has several shortcomings, such as fragile parameter choices, perfect lattice structure, and homogeneous degrees. In



© Yu-Cheng Chiu, Marc Kaufmann, Kostas Lakis, and Ulysse Schaller;
licensed under Creative Commons License CC-BY 4.0

52nd International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2026).

Editors: Jan Goedgebeur and Paweł Rzażewski; Article No. 12; pp. 12:1–12:19

Leibniz International Proceedings in Informatics



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

particular, it yields a rather large stretch compared to the average distance. While a long line of works [15–18, 35, 37] has extended Kleinberg’s framework to more general settings, none has addressed all deficiencies simultaneously. See [11] for a comprehensive discussion.

To overcome these deficiencies, we focus on *Geometric Inhomogeneous Random Graphs* (GIRGs) [9], which have emerged as a compelling framework for modeling scale-free networks. In this model, vertices are distributed independently and uniformly at random in a geometric space, each assigned a weight drawn from a power-law distribution. The connection probability combines two effects: it increases with vertex weights (favoring hubs) and decreases with geometric distance (favoring locality). A formal definition is provided in Section 2.1.

The study of *greedy routing* in GIRGs [11] provides rigorous theoretical explanations for the algorithmic small-world phenomenon. Lacking global knowledge of the network topology, the strategy is to maximize, at each step, the connection probability to the target using only local neighbor information. An intrinsic challenge in inhomogeneous geometric graphs is that an ideal neighbor must balance proximity to the target $d(v, t)$ with vertex weight w_v . In d -dimensional GIRGs this trade-off is captured by the objective function $\varphi(v) := \frac{w_v}{d(v, t)^\alpha}$, naturally derived from its edge probability. Bringmann et al. [11] obtained:

- Greedy routing succeeds with probability $\Omega(1)$.
 - Conditioned on success, the number of steps is at most $\frac{2+o(1)}{\lfloor \log(\tau-2) \rfloor} \log \log n$ a.a.s.
- Both results are asymptotically optimal. Specifically, the number of steps matches the average distance in the giant component, yielding a stretch of $1 + o(1)$.

Towards practical realizability, they further showed that greedy routing is robust under relaxations, an approximation of the objective function $\varphi(v)$ suffices for success with optimal stretch. However, in realistic settings like social networks or Internet routing, even approximate weight information might be inaccessible due to privacy concerns or limited storage. We therefore address these limitations by studying routing processes purely based on geometry.

Motivated by this, we study *geometric routing* in GIRGs. While greedy routing prioritizes connection probability, geometric routing acts greedily on geometry alone, simply forwarding to the neighbor closest to the target in the geometric space. This approach is inherently robust in real-world scenarios, circumventing issues of inaccurate popularity estimates (as in Milgram’s experiment) and adapting naturally to dynamic Internet-like graphs, where weights fluctuate over time and new nodes join, while the underlying geometry remains fixed.

Despite these advantages, geometric routing faces fundamental challenges. Without explicit weight information, the protocol must rely solely on the implicit geometry-weight correlation to find short paths via the heavy core. Consequently, its success guarantees and efficiency remain unclear. A central goal of this paper is to analyze whether this weight-agnostic approach can match the performance of greedy routing.

Our contributions

We give the first rigorous analysis of geometric routing in GIRGs. We focus on the regime with power-law exponent $2 < \tau < 3$ and decay parameter $\alpha > \tau - 1$, where larger α corresponds to a sharper geometric decay in the GIRG edge probabilities, we prove:

- Geometric routing succeeds with probability $\Omega(1)$.
- Conditioned on success, the main part takes at most $\frac{2+o(1)}{\lfloor \log(\tau-2) \rfloor} \log \log n$ steps a.a.s.

Remarkably, in this regime, geometric routing matches the asymptotically optimal performance of greedy routing, yielding a stretch of $1 + o(1)$ a.a.s., conditioned on success. We further reveal that both protocols exhibit a similar two-phase structure:

1. First phase: ascent to the heavy core.

For both protocols, vertex weights grow doubly exponentially with rate $1/(\tau - 2)$.

2. Second phase: navigation from the core to the target.

- *Geometric routing*: Requires a specific *preparation step* (combining weight growth and distance reduction) at the phase transition, before the geometric distance to the target decays doubly exponentially with rate $\tau - 2$.
- *Greedy routing*: Proceeds seamlessly without an intermediate step. The objective function $\varphi(v)$ simply continues to grow doubly exponentially with rate $\tau - 2$.

1.1 Related Works

Mean field analysis and simulations of geometric routing

The idea of weight-agnostic geometric routing¹ was previously explored by Boguná and Krioukov [5] in scale-free networks embedded in a homogeneous isotropic D -dimensional metric space. Using mean-field analysis for the first phase of the process, they concluded that, in our terminology, the navigable regime corresponds to $\alpha > 1 + 1/d$. Boguná et al. further performed numerical simulations and comparisons with real-world data [5, 6], observing that for a fixed clustering strength (governed by α), there exists a critical power-law exponent τ that determines navigability as network size grows.

This observation and other empirical findings of Boguná et al. [6] closely aligned with our theoretical results for GIRGs. In particular, their simulations demonstrated that geometric routing succeeds with a positive fraction of pairs in the strong-clustering (large α) and heavy-tailed (small τ) regime. This matches our result that geometric routing in GIRGs succeeds with constant probability when $\alpha > \tau - 1$. Moreover, they observed that within the navigable regime, the probability of hopping to a higher-degree neighbor remains high until reaching a distance-dependent critical degree, after which it decays rapidly. This observation corroborates our proposed two-phase trajectory via the heavy core. Finally, they reported that the average path length is polylogarithmic in the network size. Their simulations appear consistent with our finding that successful geometric routing in GIRGs requires $\Theta(\log \log n)$ steps, noting that the distinction between $\log \log n$ and more general polylogarithmic bounds may not be noticeable for the system sizes considered in their experiments.

Hyperbolic random graphs (HRGs) and hyperbolic embeddings

A prominent special case of GIRGs, *hyperbolic random graphs* (HRGs), have been extensively validated in experiments as a realistic model of real-world networks [6, 7, 21, 22, 32, 33, 39, 41, 45]. HRGs originated in Robert Kleinberg’s² work on routing in ad hoc wireless networks [30]. Hyperbolic geometry turned out to work well for scale-free networks [19]. Subsequent research on HRGs includes theoretical foundation [32], experimental validations of greedy processes [13, 31, 38], and empirical evidence for navigability in real-world networks [6, 7].

Bringmann et al. [9, 11] formally established that HRGs are essentially a special case of GIRGs, and the “greedy geometric processes” studied in HRGs [6, 7, 13, 31, 38] are essentially greedy routing in 1-dimensional GIRGs. In HRGs, nodes are typically embedded in a 2-dimensional hyperbolic disk³ without assigned weight. The radial and angular coordinates represent popularity and similarity, which translate naturally to weight and 1-dimensional distance in GIRGs. GIRGs extend this to arbitrary dimension and avoid the complexity of

¹ In [5], this is termed “greedy routing”. We use geometric routing to avoid confusion.

² Not to be confused with Jon Kleinberg’s small-world model [29].

³ For recent progress on higher-dimensional hyperbolic embeddings, see Jankowski et al. [25].

hyperbolic trigonometry. HRGs are known to be sparse and scale-free [24], have constant clustering coefficient [12, 24], polylogarithmic diameter [20, 27], and a giant component [4, 14]. Bringmann et al. [8–10] proved that GIRGs inherit all aforementioned properties.

Driven by the model’s realism, recent studies have focused on developing hyperbolic embeddings with specific routing-oriented properties [1–3, 21, 25, 26, 34, 42, 44]. García-Pérez et al. [21] introduced *Mercator*, a fast, accurate, and robust algorithm for embedding real-world networks into the hyperbolic plane. Bläsius et al. [2] proposed embeddings that guarantee the success of greedy routing while maintaining low stretch. More recently, Jankowski et al. [25] studied *D-Mercator*, extending *Mercator* to higher-dimensional hyperbolic space, and evaluated whether higher-dimensional embeddings are more powerful using several criteria, including the success rate of geometric routing.

Routing in generalized GIRGs

Sevilla and Fernández Anta [40] studied routing protocols in GIRGs and the generalized (κ, π) -KG model, which interpolates between Kleinberg’s model and scale-free networks. They analyzed a modified greedy objective that introduces a “critical ball” of radius r around the target, within which routing purely minimizes geometric distance. Their simulations show that increasing r improves success probability at the cost of slightly longer paths.

2 Preliminaries

2.1 Geometric Inhomogeneous Random Graphs (GIRGs)

A *Geometric Inhomogeneous Random Graph (GIRG)* is a graph $G = (V, E)$ defined over a geometric space $\mathcal{X}$ as follows:

Geometric space. We consider a d -dimensional torus $\mathcal{X}$ of volume n , which is essentially a d -dimensional cube $[0, n^{\frac{1}{d}}]^d$ with torus topology. We consider the ∞ -norm on $\mathcal{X}$, i.e. $\forall x, y \in \mathcal{X}, \|x - y\|_\infty := \max_{i \in [d]} \left\{ \min(|x_i - y_i|, n^{\frac{1}{d}} - |x_i - y_i|) \right\}$. The choice of the ∞ -norm is a matter of technical convenience, as the resulting model would remain the same if any other norm were used, as any constant factor deviation is absorbed by the asymptotic notation used in the edge probability definitions below. For the remainder of this paper, we will omit the subscript ∞ from the norm and simply write $\|x - y\|$.

Vertices. The set of vertices V is given by a Poisson point process on $\mathcal{X}$ with intensity 1. In particular, the expected number of vertices is n , and the positions of all vertices are chosen uniformly at random in $\mathcal{X}$. For each measurable $\mathcal{A} \subseteq \mathcal{X}$, let $V(\mathcal{A})$ denote the set of vertices in $\mathcal{A}$, then $|V(\mathcal{A})|$ is Poisson distributed with mean $\text{Vol}(\mathcal{A})$. Moreover, if $A, B \subseteq \mathcal{X}$ are disjoint measurable sets, then $|V(A)|$ and $|V(B)|$ are independent Poisson variables.

Weights. Let $2 < \tau < 3$ be a fixed constant. Each vertex $v \in V$ independently receives a weight w_v drawn from a probability distribution with density

$$f(w) := \begin{cases} \Theta(w^{-\tau}) & \text{if } w \geq 1, \\ 0 & \text{if } w < 1. \end{cases}$$

That is, a power-law distribution with exponent τ and minimal weight⁴ 1.

⁴ More generally, $w_{\min}$ can be a tunable scale parameter. We fix $w_{\min} = 1$ as is standard, and restoring it to any fixed constant does not introduce significant difficulties to our analysis.

Edges. Given the vertex set V , positions $\{x_v\}_{v \in V}$, and weights $\{w_v\}_{v \in V}$, for some constant $\alpha > 1$, we connect each pair of distinct vertices u, v independently with probability

$$p_{uv} := \Theta\left(\min\left\{\left(\frac{w_u w_v}{\|x_u - x_v\|^d}\right)^\alpha, 1\right\}\right). \quad (\text{EP1})$$

The parameter α governs the sharpness of the decay below the natural weight-distance threshold, as determined by the ratio $w_u w_v / \|x_u - x_v\|^d$. In (EP1), there exists a constant $c_1 > 0$ such that $p_{uv} = \Theta(1)$ whenever u and v are sufficiently close, i.e. when $\|x_u - x_v\|^d \leq c_1 w_u w_v$. In the *threshold case* $\alpha = \infty$, the connection probability drops to zero as soon as the vertices are too far apart. Instead of (EP1), we require constants $0 < c_1 \leq c_2$ such that

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

Note that there can be an interval $\|x_u - x_v\| \in [(c_1 w_u w_v)^{1/d}, (c_2 w_u w_v)^{1/d}]$ in which the behavior of p_{uv} is arbitrary.

In summary, the free parameters of the model are the expected number of vertices n , the power-law parameter τ , the dimension $d \in \mathbb{N}$, the decay parameter $\alpha > 1$, and the specific probability function p_{uv} that satisfies (EP1) or (EP2). In all subsequent theorems and proofs, we assume that all parameters except n are fixed constants. Consequently, all asymptotic notation is with respect to $n \rightarrow \infty$.

2.2 Properties of GIRGs

We next state some basic properties of GIRGs, see [9, 10] for proofs.

► **Lemma 1.** *Let u, v be two vertices of a GIRG. Then*

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

Fixing one of x_u or x_v (but not both) does not change the marginal probability by more than a constant factor. Since the vertices form a Poisson point process, this immediately implies:

► **Lemma 2.** *Let $v \in V$ be a fixed vertex with weight w_v . Then the random variable $\deg(v)$ follows a Poisson distribution with expectation $\mathbb{E}[\deg(v)] = \Theta(w_v)$.*

The next lemma from [10] shows that GIRGs contain a unique giant component of linear size, and that the average distance inside it is $\Theta(\log \log n)$. This will later be seen to match the number of steps taken by greedy and geometric routing up to a $1 + o(1)$ stretch.

► **Lemma 3.** *A.a.s. the largest component has linear size, while all other components have size $\log^{O(1)} n$. And a.a.s. the average distance in the largest component is $\frac{2 \pm o(1)}{|\log(\tau - 2)|} \log \log n$.*

2.3 Greedy Routing and Geometric Routing

Both geometric and greedy routing protocols in GIRGs follow a local greedy strategy: the current vertex v forwards the packet to a neighbor $u \in N(v)$ that optimizes a specific criterion, stopping if no neighbor improves the objective (failure) or if target t is reached (success).

Greedy routing protocol. Every vertex v knows its own address (x_v, w_v) , those of its neighbors, and t 's address is enclosed with the packet. In each hop, the protocol forwards the packet to a neighbor maximizing the connection probability to t , which is equivalent to maximizing the objective function $\varphi(u) := \frac{w_u}{\|x_u - x_t\|^a}$ for $u \in N(v)$.

In their paper, Bringmann et al. [11] established the following theorems:

► **Theorem (Success Probability).** *Greedy routing succeeds with probability $\Omega(1)$.*

► **Theorem (Greedy Path Length).** *Conditioned on success, greedy routing finds a path of length at most $\frac{2+o(1)}{|\log(\tau-2)|} \log \log n$ asymptotically almost surely.*

Comparing the length of the greedy path and the average distance in Lemma 3, we get

► **Theorem (Stretch).** *A.a.s. the stretch of a successful greedy path is $1 + o(1)$.*

A success probability of $\Theta(1)$ and stretch of $1 + o(1)$ are asymptotically best possible. In GIRGs, success is inherently limited by the model's probabilistic nature: even for vertices at infinitesimal distance, the connection probability remains $\Theta(1)$. Thus, any local routing may fail at the final hop if the current vertex is not adjacent to t , despite being arbitrarily close.

The greedy routing protocol depends on both the vertex weight w_u and geometric distance $\|x_u - x_t\|$, which poses practical restrictions. In this paper, we study the *geometric routing protocol* in GIRGs, which ignores vertex weights and relies solely on geometric positions.

Geometric routing protocol. Every vertex v knows its own position x_v , its neighbors' positions, and x_t is enclosed with the packet. In each hop, the protocol forwards the packet to a neighbor u closest to x_t .

We note that the target vertex t globally maximizes φ and minimizes the distance to x_t , providing a natural terminus for both greedy and geometric routing. The sequence of vertices obtained by applying geometric (resp. greedy) routing from source s to target t is called the *geometric path* (resp. *greedy path*).

3 Results

3.1 Geometric routing in GIRGs

Geometric routing operates on a simpler heuristic than greedy routing, yet its performance in GIRGs requires further analysis. In GIRGs, ultra-short paths of order $\Theta(\log \log n)$ typically arise from traversing a “heavy core” of high-weight vertices that act as network hubs [10]. Without knowledge of weights, it is a priori unclear whether the protocol can effectively navigate through these hubs, a process essential for efficient routing. We show that the implicit correlation between geometry and connectivity is sufficient for finding these shortcuts. Under suitable parameters, geometric routing achieves the same performance as greedy routing, succeeding with constant probability and maintaining $\Theta(\log \log n)$ path lengths.

The regime $\alpha > \tau - 1$

The success and efficiency of geometric routing are parameter-dependent, determined by the interplay between the network's geometry and its scale-free structure. Performance is governed by two key factors: (i) the power-law exponent τ for the weight distribution; and (ii) the parameter α , which controls how sharply connection probabilities decay geometrically.

Intuitively, τ dictates the prevalence of extremely heavy vertices: smaller τ corresponds to a heavier tail, while larger τ yields a more homogeneous weight distribution. The parameter α controls the influence of geometry on edge formation. In particular, it determines the fraction of “weak ties” [23] (long-range edges between low-weight nodes). Larger α makes weak ties rare (the threshold case $\alpha = \infty$ forbids them completely), resulting in highly clustered networks where long-range neighbors are predominantly heavy. Conversely, smaller α weakens geometric constraints, resulting in a more random network structure.

For geometric routing to be efficient, neighbors selected for their proximity to t should be sufficiently heavy (i.e., satisfy a weight lower bound). This property is captured by the condition $\alpha > \tau - 1$. If $\alpha < \tau - 1$, the most distant neighbors may be weak ties of low weight with non-negligible probability, which impedes progress toward the heavy core and causes routing to stall in light vertices. Henceforth, we assume $\alpha > \tau - 1$ unless specified otherwise.

Main results

Let G be a GIRG with power-law weight exponent $2 < \tau < 3$ and decay parameter $\alpha > \tau - 1$. Typically, source s and target t are sampled uniformly at random, but all results also hold if we condition on specific values for x_s, x_t, w_s, w_t . Thus, we assume s and t are fixed throughout the analysis, while the remaining vertices and all edges are drawn randomly. Our main result shows that geometric routing from s to t in G succeeds with constant probability achieves $\Theta(\log \log n)$ path lengths. The main proof ideas are given in Section 3.3; the full formal proofs are deferred to the full version.

► **Theorem 4.** *Geometric routing succeeds with probability $\Omega(1)$.*

As with greedy routing, $\Theta(1)$ success reflects the model’s inherent connectivity, being optimal in the general setting. Conditioned on success, geometric routing is efficient in traversing the main part of the routing process, as the following theorem shows.

► **Theorem 5.** *Let w_0, D_0 be sufficiently large constants. The geometric routing protocol finds a path from a vertex of weight at least w_0 to a vertex within distance D_0 of t , which constitutes the main part of the routing, whose length is upper-bounded by*

$$\frac{1 + o(1)}{|\log(\tau - 2)|} (\log \log_{w_0}(D_s^d) + \log \log_{D_0}(D_s)) + \mathcal{O}(f_0(n)), \quad (1)$$

where $D_s = \|x_s - x_t\|$ and $\omega(1) \leq f_0(n) \leq o(\log \log n)$. As $D_s \leq n^{1/d}$, the bound evaluates to

$$\frac{2 + o(1)}{|\log(\tau - 2)|} \log \log n.$$

For randomly s and t (typically having constant weights and are geometrically far apart), both greedy and geometric routing achieve path length $\frac{2+o(1)}{|\log(\tau-2)|} \log \log n$ for the main part of the routing. This demonstrates that navigation based solely on geometry is near-optimal for $\alpha > \tau - 1$, matching the efficiency of greedy routing.

3.2 Typical Trajectory

We outline the geometric routing process and the intuition behind our analysis, highlighting key differences from greedy routing. Formal statements are provided in Section 3.3, along with the main proof ideas. Complete proofs will appear in the full version.

3.2.1 Typical Trajectory of a Geometric Path

Under $\alpha > \tau - 1$, geometric proximity implicitly governs weight dynamics, guiding the path through a predictable ascent to the core and subsequent descent to t . The expected trajectory of the routing process is such that, after a few steps from the start, the current weight becomes a sufficiently large constant w_0 . We also expect that, near the end of the routing, when the distance to the target becomes some constant, only a few additional hops are needed to reach the target. We refer to the phase during which the routing process first reaches a vertex of sufficiently large constant weight w_0 as the *start phase*, and the phase during which the distance to the target are below a sufficiently large constant D_0 as the *end phase*. Our primary analysis focuses on the *main part* of the routing, which occurs in between. We will argue that, with high probability, the routing process does not fail and progresses efficiently toward the target during this main part. To this end, we leverage Lemma 2, which established $\mathbb{E}[\deg(v)] = \Theta(w_v)$ for any vertex v , implying that the weights of its neighbors follows a power-law with exponent $\tau - 1$.

First phase: implicit weight growth

For a low-weight vertex v , a.a.s. all its neighbors are spatially local and light. Since the current vertex v is still far from t , its neighbors are roughly equidistant to t , up to factors $(1 + o(1))$. The protocol thus simply selects a neighbor u that: (i) lies in the direction of t from v ; and (ii) be away from v as far as possible.

For condition (i), any neighbor of v has a constant probability to lie in the “same direction” as t from v , regardless of their weight. Condition (ii) encourages choosing a neighbor u with larger weight. Since $\alpha > \tau - 1$, weak ties are rare, most distant neighbors are “strong neighbors”, having large enough weights to connect to v with probability $\Omega(1)$. We can estimate the expected number of strong neighbors at distance $\approx r$, with weight $\Omega(r^d/w_v)$, by

$$\Theta\left(r^d \cdot \int_{r^d/w_v}^{\infty} w^{-\tau} dw\right) = \Theta(r^{d(2-\tau)} w_v^{\tau-1}).$$

Let $r_{\max}(v) := w_v^{(\tau-1)/(d(\tau-2))}$ be the maximum r where $\Omega(1)$ strong neighbors are expected. Since a.a.s. there are no weak neighbors in this distance class of $r_{\max}(v)$, the protocol typically selects a neighbor at distance $\approx r_{\max}(v)$ from v with weight of order $\Omega(r_{\max}^d/w_v) = \Omega(w_v^{1/(\tau-2)})$. In other words, besides maximizing geometric progress toward t , geometric routing exponentiates the current weight by $1/(\tau - 2)$ at each step of the first phase. This phase ends when the process hits a vertex v heavy enough to satisfy $r_{\max}(v) \approx D_v := \|x_v - x_t\|$. As a result, geometric routing reaches the heavy core within $\frac{1+o(1)}{|\log(\tau-2)|} \log \log n$ steps.

Second phase: rapid distance reduction

Once the routing reaches a node v such that $r_{\max}(v) \gtrsim D_v$, the relation between the current weight and the distance to t changes, since the length from v to its furthest neighbors can now exceed D_v . The probability that v is directly adjacent to t remains low, but v has neighbors that are significantly closer to t . Therefore, the routing protocol no longer hops simply to the furthest neighbors in the good direction. Instead, it seeks the neighbors of v in a small region around t .

To capture the neighbors at distance roughly r from t , consider an annulus $\mathcal{A}(r)$ centered at t with outer radius r and inner radius $r^{1-\mathcal{O}(\varepsilon)}$. By the structure of GIRGs, both the distribution of the maximum-weight vertex in $\mathcal{A}(r)$ and the probability that $\mathcal{A}(r)$ contains a

neighbor of v depend on the volume $\Theta(r^d)$ of $\mathcal{A}(r)$. Let r_v be the minimal radius such that $\mathcal{A}_v := \mathcal{A}(r_v)$ contains a neighbor of v . Suppose u is a neighbor of v closest to t , then $u \in \mathcal{A}_v$ and $D_u = \Theta(r_v)$ by construction. Moreover, w_u is, up to a constant factor, the maximum weight among vertices in $\mathcal{A}_v$. Indeed, if w_u were not asymptotically maximal, vertices of such weight would be abundant enough to occur in a smaller annulus of radii $r' < r_v$. Since the distance from v to these vertices in $\mathcal{A}(r')$ is of the same order as the distance to u , with high probability, v would be adjacent to at least one such vertex, contradicting the minimality of r_v . Under the power-law distribution of weights, we thus obtain

$$w_u \approx \Theta\left(r_v^{d/(\tau-1)}\right) \approx \Theta\left(D_u^{d/(\tau-1)}\right).$$

It is worth noting that this weight-distance relation need not hold for v itself, but is guaranteed for all subsequent nodes in the second phase. Specifically, for any node u in this phase (after the first hop), the routing process finds a neighbor in the annulus $\mathcal{A}_u$ with radius $r_u \approx D_u^{\tau-2}$. In other words, throughout the second phase, except for the very first hop, the distance to t decreases by an exponent of $\tau - 2$ at each step. Consequently, in the second phase, geometric routing reaches a vertex at constant distance to t within $\frac{1+o(1)}{|\log(\tau-2)|} \log \log n$ steps.

3.2.2 Comparison with Trajectory of a Greedy Path

The typical trajectories of geometric and greedy routing both exhibit a two-phase structure and achieve the same asymptotic progress, yet the underlying mechanisms are different. In particular, the two phases of geometric routing are determined by the relation between

$$r_{\max}(v) := w_v^{(\tau-1)/(d(\tau-2))} \quad \text{and} \quad D_v := \|x_v - x_t\|,$$

while in greedy routing the two phases are characterized by the relation between (see also [11])

$$\varphi(v) := w_v/D_v^d \quad \text{and} \quad w_v^{-1/(\tau-2)}.$$

Both descriptions encode the same underlying threshold, phrased in different parameters.

In the first phase, both protocols exhibit doubly exponential weight growth at rate $1/(\tau - 2)$. In the second phase, geometric routing exhibits doubly exponential decay of the distance D_v , whereas greedy routing only ensures doubly exponential growth of its objective $\varphi(v)$. Notably, geometric routing requires a single preparatory step before the desired doubly exponential decay begins, a feature absent in greedy routing. Below, we briefly outline the two-phase trajectory of greedy routing.

First phase of greedy routing

When v has small weight w_v compared to D_v , neighbors are a.a.s. equidistant to t within factors of $(1 + o(1))$. Conversely, neighbor weights fluctuate by non-negligible constant factors (so do their distances from v , but this does not affect the objective values). Consequently, weight fluctuations dominate the selection process, favoring hops to heavier neighbors.

By the $\tau - 1$ power-law tail of neighbors, the expected number of neighbors with weight at least w is $\Theta(w_v w^{2-\tau})$. Setting this to $\Theta(1)$ implies v has a neighbor of weight $w \approx w_v^{1/(\tau-2)}$, yielding a doubly exponential weight growth at rate $1/(\tau - 2)$. This stops when a node v with $\varphi(v) \gtrsim w_v^{-1/(\tau-2)}$ is reached, as the relation of the current weight and distance to t changed. Hence, this first phase needs only $\approx \log \log n / |\log(\tau - 2)|$ steps.

Second phase of greedy routing

Beyond the threshold $\varphi(v) \gtrsim w_v^{-1/(\tau-2)}$, there exist neighbors of v that are much closer to t . Despite a typically lower weight relative to v , their proximity compensates enough to increase the objective⁵ to $\approx \varphi(v)^{\tau-2}$. One can further show that the best neighbor keeps the routing in the second phase, so we can repeat the argument. Therefore, in the second phase, the objective increases by an exponent $\tau - 2$ at each step. After $\approx \log \log n / |\log(\tau - 2)|$ rounds, the routing reaches a node v with objective $\varphi(v) = \Omega(1)$.

3.3 Proof Overview

In the formal proofs, we show that w.h.p. geometric routing adheres to the typical trajectory described above. Although the best neighbor u of a single vertex v is predictable (see Section 3.3.2 for the precise statements), a naive step-by-step analysis suffers from intractable dependencies: selecting u as the best neighbor reveals that no superior candidates exist in v 's neighborhood, which biases the environment for all subsequent steps.

To bypass these dependencies, we adapt the *layer technique* of [11]: we partition the vertex space into carefully crafted layers corresponding to the two phases of geometric routing. Uncovering these layers sequentially, rather than tracking individual nodes on the path, allows us to rigorously track the trajectory's progress.

3.3.1 Formalizing the Trajectory: The Layer Technique

We now describe the layer argument used to formalize the intuitive behavior of geometric routing. The argument tracks the progress of the routing process through layers, while the local estimates in Section 3.3.2 provide the probabilistic guarantees for moving from one layer to the next.

The proof is organized around a trajectory lemma, which combines the layer construction with the local progress estimates on where to expect suitable neighbors. In this proceedings version, we state a compressed form of this lemma that captures the consequences needed for our main results. The lemma shows that, once the routing process enters the main part, the protocol forces the path to leave the current layer and enter a subsequent one with high probability, except for the transition layer, which may be visited twice. Consequently, the path does not get stuck, visits each layer only a controlled number of times, and reaches the constant-distance region around the target within the desired number of steps. This yields Theorem 5 and is the main ingredient in the proof of Theorem 4. The complete layer-exposure argument will appear in the full version.

Layers for geometric routing

To formalize our discussion, we first introduce some technical notations. Throughout, we assume $\alpha > \tau - 1$ and $2 < \tau < 3$. For any $\varepsilon > 0$, set $\gamma(\varepsilon) := \frac{1-\varepsilon}{\tau-2}$. In our proofs, ε would be sufficiently small such that $\gamma(\varepsilon) > 1$. Let $N(v)$ denote the set of neighbors of v , and $D_v := \|x_v - x_t\|$. We also define $R_v(\varepsilon) := (w_v^{1+\gamma(\varepsilon)})^{1/d} \approx w_v^{(\tau-1)/d(\tau-2)}$ for sufficiently small ε , which is the expected distance from v to its furthest neighbor.

Let $\varepsilon_1 = \varepsilon_1(\alpha, \tau)$ be a fixed constant which will be chosen sufficiently small in the proofs, in particular, $\gamma(\varepsilon_1) > 1$. We categorize each vertex v based on whether its distance to the target is greater or less than the expected distance from v to its furthest neighbor, and define:

⁵ Note that we can assume $\varphi(v) < 1$, otherwise v is incident to t with probability $\Omega(1)$.

$$V_1 := \left\{ v \in V \mid D_v \geq R_v(\varepsilon_1) \right\} \quad \text{and} \quad V_2 := \left\{ v \in V \mid D_v \leq R_v(\varepsilon_1) \right\}.$$

The vertices in V_1 and V_2 behave differently under the geometric routing protocol, where vertices in V_1 (resp. V_2) correspond to the first phase (resp. the second phase) of the routing.

Recall from discussion in Section 3.2.1, it typically takes only few steps for the source s to reach a vertex with constant weight w_0 , and similarly, for a vertex located within a ball of radius D_0 around t to reach the target t . These are referred to as the start and end phases of the routing, respectively. To this end, for any pair (w, D) of weight and distance, we define

$$\bar{V}(w, D) := \{ v \in V_1 \mid (1) : w_v \geq w; (2) : D_v \geq D \} \cup \{ v \in V_2 \mid D_v \geq D \}.$$

The main part of the routing process takes place within the set $\bar{V}(w, D)$ for some suitable constant values of w, D . In particular, we choose w_0 and D_0 sufficiently large so that the lemmata in Section 3.3.2 can be applied to $\bar{V}(w_0, D_0)$. Moreover, let $V_{\leq D_0} := \{ v \in V \mid D_v \leq D_0 \}$ characterizes the end phase that will be treated separately. Our primary goal is to show that the routing survives this main part $\bar{V}(w_0, D_0)$ with sufficiently high probability, and eventually reaches $V_{\leq D_0}$, which is a small ball around target t .

We analyze the structure of the geometric path in the main part by partitioning the set $\bar{V}(w_0, D_0)$ into layers. We will define two classes of layers which correspond to the two phases: the layers $A_{1,j}$ divide the area $V_1 \cap \bar{V}(w_0, D_0)$ and are defined via weights, whereas the layers $A_{2,j}$ divide the area $V_2 \cap \bar{V}(w_0, D_0)$ and are defined via distances. The idea is to show that, with sufficiently high probability, there exists a node u_ℓ on the geometric path P which has neighbors in $V_{\leq D_0}$, such that until reaching u_ℓ , P visits every layer at most once, except possibly the special layer $A_{2,*}$, which may be visited twice. The special layer $A_{2,*}$ is the transition layer where the geometric path first enters the second phase.

Roughly speaking, we construct a sequence $y_0 = w_0 < y_1 < \dots$ of weight thresholds and a sequence $z_0 = D_0 < z_1 < \dots$ of distance thresholds, both growing doubly exponentially:

$$y_{j+1} = y_j^{1/(\tau-2)-o(1)} \quad \text{and} \quad z_{j+1} = z_j^{1/(\tau-2)-o(1)}.$$

The first-phase layer $A_{1,j}$ consists of vertices whose weights satisfy $y_{j-1} \leq w_v < y_j$, while the second-phase layer $A_{2,j}$ consists of vertices whose distances to t satisfy $z_{j-1} \leq D_v < z_j$. The exact construction in the full proof uses two error parameters to handle the transition between constant and asymptotic scales. The resulting boundary layers contribute only an additive $\mathcal{O}(f_0(n)) = o(\log \log n)$ term and do not affect the leading constant in Theorem 5.

We can now state the trajectory lemma in the form needed below.

► **Lemma 6** (Trajectory lemma, simplified form). *Let w_0 and D_0 be sufficiently large constants, and let $M := \min\{w_0, D_0\}$. Let $\omega(1) \leq f_0(n) \leq o(\log \log n)$ be any growing function. Suppose that the geometric path P first enters $\bar{V}(w_0, D_0)$ at a vertex u_1 . Then, with probability $1 - \mathcal{O}(M^{-\Omega(1)})$, the following holds:*

P follows the layered trajectory: it visits the layers in the prescribed order, visiting each first-phase layer and each second-phase layer at most once, except possibly the layer $A_{2,}$, which may be visited twice. Moreover, before leaving the main part, P reaches a vertex u_ℓ s.t.*

$$\mathbb{E}_{<D_0} [|N(u_\ell) \cap V_{<D_0}|] = \Omega(M^{\Omega(1)}),$$

where the expectation is w.r.t. uncovering the vertices whose distance to target is at most D_0 , conditioned on the position and weight of u_ℓ . The number of hops from u_1 to u_ℓ is at most

$$\frac{1 + o(1)}{|\log(\tau - 2)|} (\log \log_{w_0}(D_{u_1}^d) + \log \log_{D_0}(D_{u_1})) + \mathcal{O}(f_0(n)).$$

In particular, since $D_{u_1} \leq n^{1/d}$, this is at most $\frac{2+o(1)}{|\log(\tau-2)|} \log \log n$.

12:12 Geometric Routing in Geometric Inhomogeneous Random Graphs

The main idea to prove the trajectory lemma is to establish the following property. Consider P_i which is the geometric path induced by the first i layers. Assuming that the i^{th} layer is not $A_{2,*}$, then, with sufficiently high probability, either

1. P_i contains no vertex from the first i layers, or
2. the first $v \in P_i$ contained in the i^{th} layer has a neighbor in some subsequent layer that is closer to t than every uncovered neighbor of v in the first i layers.

In the case where the i^{th} layer is $A_{2,*}$, the same property holds with respect to either the first or the second vertex of P_i contained in $A_{2,*}$.

This is the mechanism by which the geometric choice forces the path to leave the current layer. Counting the layers induced by the doubly exponentially growing weight and distance thresholds gives the bound for the number of hops. The expectation condition involving neighbors of u_ℓ is the preparation for routing to enter the end phase.

With this technical construction, we circumvent problematic dependencies and the trajectory lemma implies the desired length bound in Theorem 5. In Section 3.3.3, we combine the trajectory lemma with the start and end phases to obtain the constant success probability in Theorem 4.

3.3.2 Where to Expect Neighbors

We now give the technical ingredients ensuring that geometric routing progresses to subsequent layers rather than remaining within the current one. Specifically, we collect the statements and describe the proof idea that bound the weights or distances of the next neighbor chosen by the geometric routing protocol. The proofs will appear in the full version of the paper.

Recall the vertex sets:

$$V_1 := \left\{ v \in V \mid D_v \geq R_v(\varepsilon_1) \right\} \quad \text{and} \quad V_2 := \left\{ v \in V \mid D_v \leq R_v(\varepsilon_1) \right\},$$

where $R_v(\varepsilon) := (w_v^{1+\gamma(\varepsilon)})^{1/d} \approx w_v^{(\tau-1)/d(\tau-2)}$ is the expected distance from v to its furthest neighbor, and $\varepsilon_1 = \varepsilon_1(\alpha, \tau)$ is a fixed constant which will be chosen sufficiently small.

We will see the weight of the next vertex u typically increases by a fixed exponent; whereas if $v \in V_2$, the routing tends to remain within V_2 and the trajectory is dominated by decreases in distance, as described in Section 3.2.1.

Analysis of each hop in the first phase

Suppose $v \in V_1$ is the current vertex. Its weight w_v is low in the sense that the current distance to the target D_v is larger than the expected distance to its furthest neighbor. In order to describe a trajectory with doubly exponential weight growth, we classify the vertices in V which accelerate the routing as desired. Let $\zeta_1 := \frac{3}{2}$ in the case $\alpha = \infty$, and $\zeta_1 := \max\{\frac{2\alpha-1}{2\alpha+2-2\tau}, \frac{3}{2}\}$ for case $1 < \alpha < \infty$. For $v \in V_1$ and $\varepsilon = \varepsilon(n) > 0$, we define:

$$\begin{aligned} V_1^+(v, \varepsilon) &:= \left\{ u \in V \mid w_u \geq w_v^{\gamma(\zeta_1 \varepsilon)} \text{ and } D_u \leq D_v - \frac{R_v(\varepsilon)}{2} \right\}, \text{ and} \\ V_1^-(v, \varepsilon) &:= \left\{ u \in V \mid w_u < w_v^{\gamma(\zeta_1 \varepsilon)} \text{ and } D_u \leq D_v - \frac{R_v(\varepsilon)}{2} \right\}. \end{aligned} \tag{2}$$

The condition $D_u \leq D_v - R_v(\varepsilon)/2$ identifies the vertices u that are significantly closer to the target than the current vertex v . The following lemma states that $v \in V_1$ have a large expected number of good neighbors (i.e. $u \in N(v) \cap V^+(v, \varepsilon)$); whereas the expected number of bad neighbors (i.e. $u \in N(v) \cap V^-(v, \varepsilon)$) is small. Notice that for $\varepsilon = \varepsilon_1$, every vertex $v \in V_1$ satisfies $D_v \geq R_v(\varepsilon_1)$. In this case, the statement can be applied for all $v \in V_1$.

► **Lemma 7** (Neighborhoods in first phase). *Let $0 < \varepsilon \leq \varepsilon_1$ be sufficiently small and $v \in V_1$ be such that $D_v \geq R_v(\varepsilon)$. And let $w_1(\varepsilon) := c \cdot e^{d/\varepsilon}$ for c large enough. Then*

- (i) $\mathbb{E}[|N(v) \cap V^+(v, \varepsilon)|] = \Omega(w_v^\varepsilon)$, and
- (ii) *if in addition $w_v \geq w_1(\varepsilon)$, then $\mathbb{E}[|N(v) \cap V^-(v, \varepsilon)|] = \mathcal{O}(w_v^{-\Omega(\varepsilon)})$.*

The idea of the proof is to investigate an annulus around v whose inner and outer radii are both on the order of $R_v(\varepsilon)$. This allows us to lower (resp. upper) bound the number of candidates of good (resp. bad) neighbors of v . Therefore, as the routing proceeds to a neighbor that is significantly closer to target t , with high probability (w.h.p.) it also hops to a vertex with higher weight, instead of staying at low-weight vertices.

Analysis of each hop in the second phase

Once the routing reaches a vertex in V_2 , it enters the second phase. The vertices in V_2 have relatively large weight so that $v \in V_2$ has neighbors at a distance roughly D_v . We assume $D_v > 1$ throughout this section. This is valid since we are in the main part of routing, where it is guaranteed that $D_v > D_0$, where D_0 is a constant chosen sufficiently large. The goal now is to decrease the distance to t doubly exponentially at every step of the routing. Furthermore, once we reach a vertex in V_2 , we aim to remain within V_2 throughout the rest of the process.

The desired behavior is almost true for the entire second phase, except for the very first step upon entering it. To be more precise, suppose the routing reaches the first vertex $v \in V_2$ in the second phase, and let $u \in N(v)$ be the next vertex selected by the routing protocol. Ideally, we might hope that $D_u \approx D_v^{\tau-2}$, analogous to the behavior in the first phase. However, this is not the case. In fact, we have $D_u \approx D_v^{(1-\eta)(\tau-1)}$, where $\eta \in [\frac{\tau-2}{\tau-1-\varepsilon_1}, \frac{1}{\tau-1-\varepsilon_1}]$ is a random variable determined by the relation $w_v = \Theta(D_v^{d\eta})$. Here, η captures the ratio between the weight w_v and the distance D_v when the routing first enters the second phase at vertex v . That said, from u onward, we will show that the distance to the target decreases roughly by a fixed exponent $\tau - 2$ with each subsequent hop toward the target.

The idea behind analyzing the second phase is as follows: Suppose $v \in V_2$ is the current vertex, we search for the smallest annulus around the target t that contains a neighbor of v . More precisely, given $v \in V_2$ and ε , consider an annulus centered at t with outer radius r and inner radius $r^{1-\mathcal{O}(\varepsilon)}$, and we let r be the smallest value such that this annulus contains a neighbor of v . By the minimality of r , this annulus contains the neighbor of v that is closest to t , and no neighbors of v lie inside the inner circle of radius $r^{1-\mathcal{O}(\varepsilon)}$. In other words, there are neighbors of v at approximately distance r from t , but not much closer. Hence, the distance from the next vertex selected by the routing protocol to t will be of order r .

To define the minimal annulus centered at t that contains a neighbor of v . Suppose the current vertex $v \in V_2$ is such that $w_v = \Theta(D_v^{d\eta})$ for some $\eta \geq \frac{\tau-2}{\tau-1-\varepsilon_1} > \frac{\tau-2}{\tau-1}$. We may assume η is bounded away from 1 with a neglectable error probability as shown in the following:

► **Lemma 8.** *For $2 < \tau < 3$, and $\varepsilon > 0$ small enough, let $V_{\geq D_0} := \{v \in V \mid D_v \geq D_0\}$, then*

$$\mathbb{P}\left[\forall v \in V_{\geq D_0}, w_v \leq \Theta\left(D_v^{d(1-\varepsilon)}\right)\right] \geq 1 - D_0^{-\Omega(\varepsilon)}.$$

Lemma 8 ensures that we may assume $\eta \leq 1 - \varepsilon$, at a cost of an error probability that is polynomially small in ε . We adopt this assumption for the remainder of this discussion. From this point on, whenever we encounter an expression of the form $D_v^{1-\eta-x}$, with x denoting a small quantity, we implicitly assume that x is chosen small enough such that the exponent $1 - \eta - x$ remains positive.

To define the radius of the minimal annulus, let $r_v = r(v, \varepsilon) := D_v^{(1-\eta+\varepsilon)(\tau-1)}$, where $\varepsilon > 0$ is small enough such that $(1-\eta+\varepsilon)(\tau-1) < 1$, which is possible since $(1-\eta)(\tau-1) < \tau-2$ is bounded away from 1. We will also assume this throughout the rest of the discussion whenever ε is chosen sufficiently small. Moreover, let us define $D_1(\varepsilon) := c \cdot e^{d/\varepsilon}$ to be a constant depending on ε with c chosen sufficiently large. Then, $(1-\eta+\varepsilon)(\tau-1) < 1$ implies that we can assume $r_v < \frac{1}{2}D_v$ when $D_v \geq D_1(\varepsilon)$.

Finally, let us consider the annulus $\mathcal{A}_v$ centered at t , with outer radius $r_v = D_v^{(1-\eta+\varepsilon)(\tau-1)}$, and inner radius $r_I = D_v^{(1-\eta-\varepsilon)(\tau-1)} = r_v^{1-\mathcal{O}(\varepsilon)}$. Note that the area of $\mathcal{A}_v$ is $\Theta(r_v^d)$. We shall see that $\mathcal{A}_v$ is chosen such that with high probability, there is a neighbor of v in the annulus $\mathcal{A}_v$, but no neighbors of v in its inner circle.

In fact, let $\mathcal{I}_\varepsilon$ denote the inner circle of $\mathcal{A}_\varepsilon$, whose the area is $\Theta(r_I^d)$. Then we have:

► **Lemma 9** (Neighborhoods in $\mathcal{I}_\varepsilon$). *Let $v \in V_2$ with $w_v = \Theta(D_v^{d\eta})$, and let $\varepsilon > 0$ be sufficiently small. If $D_v \geq D_1(\varepsilon)$, then*

$$\mathbb{P}[N(V) \cap V(\mathcal{I}_\varepsilon) = \emptyset] \geq 1 - D_v^{-\Omega(\varepsilon)}.$$

Next, we show that the annulus $\mathcal{A}_v$ contains neighbors of v . Crucially, we lower bound their weights to prevent the weight from dropping significantly, thereby ensuring that the process does not revert to the first phase. In fact, rather than merely verifying that the next vertex u satisfies the minimum weight condition for V_2 (i.e. $w_u \geq D_u^{d(\tau-2)/(\tau-1-\varepsilon_1)}$), we establish a strictly stronger bound. This tighter bound not only ensures $u \in V_2$ but also reveals deeper insights into the routing behavior in the second phase.

To this end, analogous to the first phase, we define the sets of good and bad vertices for the second phase. Let $v \in V_2$ be such that $w_v = \Theta(D_v^{d\eta})$, and let $\varepsilon = \varepsilon(n) > 0$ be sufficiently small. Let $\zeta_2 := 1$ for $\alpha = \infty$, and $\zeta_2 := \max\{\frac{2\tau-1}{2\alpha+2-2\tau}, 1\}$ for $1 < \alpha < \infty$. We define:

$$\begin{aligned} V_2^+(v, \varepsilon) &:= \left\{ u \in V \mid w_u \geq D_u^{\frac{d(1-\eta-\zeta_2\varepsilon)}{(\tau-1)(1-\eta+\varepsilon)}} = D_u^{\frac{d}{\tau-1} \left(1 - \frac{(\zeta_2+1)\varepsilon}{1-\eta+\varepsilon}\right)} \right\}, \text{ and} \\ V_2^-(v, \varepsilon) &:= \left\{ u \in V \mid w_u < D_u^{\frac{d(1-\eta-\zeta_2\varepsilon)}{(\tau-1)(1-\eta+\varepsilon)}} = D_u^{\frac{d}{\tau-1} \left(1 - \frac{(\zeta_2+1)\varepsilon}{1-\eta+\varepsilon}\right)} \right\}. \end{aligned} \quad (3)$$

Similar to the first phase, the following lemma shows that the probability of $v \in V_2$ having a good neighbor in $\mathcal{A}_v$ is high, while the probability of such v having a bad neighbor in $\mathcal{A}_v$ is low. Therefore, as the routing proceeds, w.h.p. it hops to a neighbor with distance roughly r_v to the target t , but not significantly closer. Moreover, the weight does not drop substantially, which ensures that the routing process stays within V_2 .

► **Lemma 10** (Neighborhoods in $\mathcal{A}_v$). *Let $v \in V_2$ with $w_v = \Theta(D_v^{d\eta})$, and let $\varepsilon > 0$ be sufficiently small. If $D_v \geq D_1(\varepsilon)$, then*

- (i) $\mathbb{P}[N(V) \cap V_2^+(v, \varepsilon) \cap V(\mathcal{A}_v) \neq \emptyset] \geq 1 - e^{-D_v^{\Omega(\varepsilon)}};$
- (ii) $\mathbb{P}[N(V) \cap V_2^-(v, \varepsilon) \cap V(\mathcal{A}_v) = \emptyset] \geq 1 - D_v^{-\Omega(\varepsilon)}.$

Suppose $v \in V_2$, and $u \in N(v)$ is the next vertex selected by the routing protocol. By Lemmas 9 and 10, with high probability we obtain

$$D_v^{(1-\eta-\varepsilon)(\tau-1)} \leq D_u \leq D_v^{(1-\eta+\varepsilon)(\tau-1)}.$$

Recall that if v is the first vertex that the routing reaches upon entering the second phase, then $\eta(v)$ is a random variable lying in the interval $[\frac{\tau-2}{\tau-1-\varepsilon_1}, \frac{1}{\tau-1-\varepsilon_1}]$. From this point on, as shown in Lemma 10, w.h.p. each subsequent vertex is not a bad neighbor of its predecessor. In other words, we have the following corollary:

► **Corollary 11.** *Except for the very first vertex that the routing visits upon entering the second phase, every vertex v chosen by the routing protocol satisfies $w_v \geq D_v^{\frac{d}{\tau-1}} \left(1 - \frac{(\zeta_2+1)\varepsilon}{1-\eta+\varepsilon}\right)$, where the quantity $\left(\frac{\zeta_2+1}{1-\eta+\varepsilon} \cdot \varepsilon\right)$ can be made arbitrarily small. In particular, this implies that $v \in V_2$, and the next neighbor u satisfies $D_u \leq D_v^{(\tau-2)(1+\mathcal{O}(\varepsilon))}$, i.e. the distance to the target decreases roughly by an exponent of $\tau - 2$ with each hop thereafter.*

3.3.3 Proof of Theorem 4

We now explain how the preceding ingredients in Sections 3.3.1 and 3.3.2 imply that for all initial choices of s and t , geometric routing succeeds with constant probability. The argument is achieved in three steps, analyzing the start phase, the main part, and the end phase.

Before the formal proof, we clarify the successful outcome that marks the completion of each phase and serves as the starting condition for the next:

1. **Start phase:** Starting from an arbitrary source s , the goal is to show that the path successfully initiates its journey. This means it either enters $V_{<D_0}$ directly or finds a suitable starting vertex $u_1 \in \bar{V}(w_0, D_0)$ to initiate the main part.
2. **Main part:** Starting from u_1 , the goal is to show that the path successfully navigates the main part and arrives at a vertex u_ℓ that has in expectation $\Omega(1)$ neighbors in $V_{<D_0}$. This is precisely Lemma 6.
3. **End phase:** The goal is to show that the path successfully reaches the target t once it is positioned to enter $V_{<D_0}$. This can happen either after traversing the main part to reach a vertex u_ℓ as mentioned, or by finding such a neighbor directly during the start phase.

Throughout the process, each step succeeds with a non-zero constant probability. However, these probabilities are not large enough to allow a union bound over all error events. Instead, we compute probabilities conditioned on the success of the previous step.

► **Theorem 4.** *Geometric routing succeeds with probability $\Omega(1)$.*

Proof. Let us choose large enough constants $w_0 \geq w_1(\varepsilon_1)$ and $D_0 \geq D_1(\varepsilon_1)$ such that the preconditions of Lemma 6 are met.

Start phase. First, we analyze the initiation of the routing process. We will show that the path successfully begins its journey with constant probability. More formally, we define the successful start event:

Let $\mathcal{E}_s$ be the event that the geometric path starting at s eventually visits a vertex v that satisfies at least one of the following conditions:

- v is in the end zone, i.e. $D_v < D_0$; or
- v has $\Omega(1)$ expected neighbors in $V_{<D_0}$; or
- v is a suitable starting vertex for the main part, i.e. $v \in \bar{V}(w_0, D_0)$.

We now prove that $\mathbb{P}[\mathcal{E}_s] = \Omega(1)$ by analyzing all possible initial scenarios for the source s .

1. Case $D_s < D_0$: The event $\mathcal{E}_s$ is immediately satisfied with probability 1 by taking $v = s$. In this case, the region $V_{<D_s}$ (in fact the whole graph except s and t) remains unrevealed.
2. Case $D_s \geq D_0$: We analyze the marginal expectation $\mathbb{E}_{<D_0}[N(s) \cap V_{<D_0}]$ after fixing s and before uncovering $V_{<D_0}$.
 - (1) $\mathbb{E}_{<D_0}[N(s) \cap V_{<D_0}] \geq \frac{1}{2}$: With probability $\Omega(1)$, s has at least one neighbor in $V_{<D_0}$. If so, the best neighbor (closest to t) will be in $V_{<D_0}$, and $\mathcal{E}_s$ is again satisfied by taking $v = s$. In this case, the routing process is ready to enter the end phase, with $V_{<D_0}$ remains unrevealed.

- (2) $\mathbb{E}_{<D_0}[N(s) \cap V_{<D_0}] < \frac{1}{2}$: With probability $\Omega(1)$, s has no neighbor in $V_{<D_0}$. Conditioned on this, we consider the event where s has a unique neighbor s' in $V_{\geq D_0}$, and such that $w_{s'} \geq w_0$. Note that if this holds, s' is the best neighbor of v , and $\mathcal{E}_s$ is satisfied by taking $v = s'$.

To see this happens with probability $\Omega(1)$, note that $\deg_{V_{\geq D_0}}(s) = |N(s) \cap V_{\geq D_0}|$ is Poisson binomially distributed with expectation $\Theta(w_s) = \Theta(1)$. Therefore, with probability $\Omega(1)$, $|N(s) \cap V_{\geq D_0}| = 1$, there is a unique $s' \in N(s) \cap V_{\geq D_0}$.

Moreover, the degree distribution of a uniformly random neighbor of s follows a power-law with parameter $\tau - 1$. Therefore, with probability $\Omega(1)$, the unique s' has weight $w_{s'} \geq w_0$.

Since every scenario leads to the event $\mathcal{E}_s$ with at least a constant probability, we conclude that $\mathbb{P}[\mathcal{E}_s] = \Omega(1)$.

Main part of routing process. Our goal is to show that, conditioned on $\mathcal{E}_s$, the geometric path eventually enters $V_{<D_0}$. In particular, we focus on the case where the start phase results in a vertex $u_1 \in \bar{V}(w_0, D_0)$, and show that the routing proceeds to a vertex u_ℓ which we expect to have neighbors in $V_{<D_0}$. More formally, we consider the event

$$\mathcal{E}_m : \text{the routing on } G[V_{\geq D_0}] \text{ finds } u_\ell \text{ with } \mathbb{E}_{<D_0}[|N(u_\ell) \cap V_{<D_0}|] = \Omega(1).$$

By Lemma 6, we have $\mathbb{P}[\mathcal{E}_m \mid \exists u_1 \in \bar{V}(w_0, D_0)] = \Omega(1)$.

End phase. We now analyze the end phase, conditioned on the success of the prior stages. This means we condition on $\mathcal{E}_s$ holding, and additionally on $\mathcal{E}_m$ holding in the case where the geometric path entered the main part. Then, we either have $D_s < D_0$, or the routing reached a vertex v such that $\mathbb{E}_{<D_0}[|N(v) \cap V_{<D_0}|] = \Omega(1)$.

To show that the routing successfully reaches t with probability $\Omega(1)$, we first consider the latter case. Since the vertices in the region $V_{<D_0}$ remains unrevealed, by properties of Poisson point process, $|N(v) \cap V_{<D_0}|$ is Poisson binomially distributed with expectation $\Omega(1)$. Hence, following the same logic as the start phase, with probability $\Omega(1)$, v has a unique neighbor $v' \in V_{<D_0}$ with a sufficiently large constant weight. Since $w_{v'}$ is a sufficiently large constant and $D_{v'} < D_0$, v' is then connected to t with probability $\Omega(1)$. The independence of these two events means the total probability of finding t via the two-hop path through v' is $\Omega(1)$.

In the case where $D_s < D_0$, the region $V_{<D_s}$ remains unrevealed. Again, $|N(s) \cap V_{<D_s}|$ is Poisson binomially distributed with expectation $\Omega(1)$. By repeating the above argument with D_0 replaced by D_s , it follows that the probability of finding t via an analogous two-hop path from s is also $\Omega(1)$.

In summary, we have shown that each phase of the routing succeeds with at least constant probability conditioned on previous stages being successful. Hence, geometric routing succeeds with a probability of $\Omega(1)$. $\blacktriangleleft$

References

- 1 Royá Aliakbarisani, M Ángeles Serrano, and Marián Boguñá. Feature-enriched hyperbolic network geometry. *Physical Review Research*, 7(3):033036, 2025.
- 2 Thomas Bläsius, Tobias Friedrich, Maximilian Katzmann, and Anton Krehmer. Hyperbolic embeddings for near-optimal greedy routing, 2020.
- 3 Thomas Bläsius, Tobias Friedrich, Anton Krehmer, and Sören Laue. Efficient embedding of scale-free graphs in the hyperbolic plane. *IEEE/ACM transactions on Networking*, 26(2):920–933, 2018. doi:10.1109/TNET.2018.2810186.

- 4 Michel Bode, Nikolaos Fountoulakis, and Tobias Müller. On the largest component of a hyperbolic model of complex networks. *The Electronic Journal of Combinatorics*, pages P3–24, 2015.
- 5 Marián Boguná and Dmitri Krioukov. Navigating ultrasmall worlds in ultrashort time. *Physical review letters*, 102(5):058701, 2009.
- 6 Marián Boguná, Dmitri Krioukov, and Kimberly C Claffy. Navigability of complex networks. *Nature Physics*, 5(1):74–80, 2009.
- 7 Marián Boguná, Fragkiskos Papadopoulos, and Dmitri Krioukov. Sustaining the internet with hyperbolic mapping. *Nature communications*, 1(1):62, 2010.
- 8 Karl Bringmann, Ralph Keusch, and Johannes Lengler. Sampling geometric inhomogeneous random graphs in linear time. In *25th Annual European Symposium on Algorithms (ESA 2017)*, volume 87 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20:1–20:15, 2017. doi:10.4230/LIPIcs.ESA.2017.20.
- 9 Karl Bringmann, Ralph Keusch, and Johannes Lengler. Geometric inhomogeneous random graphs. *Theoretical Computer Science*, 760:35–54, 2019. Preliminary version in FOCS 2016. doi:10.1016/j.tcs.2018.08.014.
- 10 Karl Bringmann, Ralph Keusch, and Johannes Lengler. Average distance in a general class of scale-free networks. *Advances in Applied Probability*, 57(2):371–406, 2025. doi:10.1017/apr.2024.43.
- 11 Karl Bringmann, Ralph Keusch, Johannes Lengler, Yannic Maus, and Anisur R. Molla. Greedy routing and the algorithmic small-world phenomenon. *Journal of Computer and System Sciences*, 125:59–105, 2022. Preliminary version in PODC 2017. doi:10.1016/j.jcss.2021.11.003.
- 12 Elisabetta Candellero and Nikolaos Fountoulakis. Clustering and the hyperbolic geometry of complex networks. *Internet Mathematics*, 12(1-2):2–53, 2016. Preliminary version appeared in WAW 2014. doi:10.1080/15427951.2015.1067848.
- 13 Andrej Cvetkovski and Mark Crovella. Hyperbolic embedding and routing for dynamic graphs. In *IEEE INFOCOM 2009*, pages 1647–1655. IEEE, 2009. doi:10.1109/INFCOM.2009.5062083.
- 14 Nikolaos Fountoulakis and Tobias Müller. Law of large numbers for the largest component in a hyperbolic model of complex networks. *Annals of Applied Probability*, 28(1):607–650, March 2018. doi:10.1214/17-AAP1314.
- 15 Pierre Fraigniaud, Cyril Gavoille, and Christophe Paul. Eclecticism shrinks even small worlds. In *Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*, pages 169–178, 2004. doi:10.1145/1011767.1011793.
- 16 Pierre Fraigniaud and George Giakkoupis. On the searchability of small-world networks with arbitrary underlying structure. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 389–398, 2010. doi:10.1145/1806689.1806744.
- 17 Pierre Fraigniaud and George Giakkoupis. Greedy routing in small-world networks with power-law degrees. *Distributed computing*, 27(4):231–253, 2014. doi:10.1007/S00446-014-0210-Y.
- 18 Massimo Franceschetti and Ronald Meester. Navigation in small-world networks: A scale-free continuum model. *Journal of Applied Probability*, 43(4):1173–1180, 2006.
- 19 Tobias Friedrich. From graph theory to network science: The natural emergence of hyperbolicity. In *36th International Symposium on Theoretical Aspects of Computer Science (STACS 2019)*, pages 5:1–5:9. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.STACS.2019.5.
- 20 Tobias Friedrich and Anton Krophmer. On the diameter of hyperbolic random graphs. *SIAM Journal on Discrete Mathematics*, 32(2):1314–1334, 2018. doi:10.1137/17M1123961.
- 21 Guillermo García-Pérez, Antoine Allard, M Ángeles Serrano, and Marián Boguná. Mercator: uncovering faithful hyperbolic embeddings of complex networks. *New Journal of Physics*, 21(12):123033, 2019.

- 22 Guillermo García-Pérez, Marián Boguñá, Antoine Allard, and M. Ángeles Serrano. The hidden hyperbolic geometry of international trade: World trade atlas 1870–2013. *Scientific reports*, 6(1):1–10, 2016.
- 23 Mark S Granovetter. The strength of weak ties. *American journal of sociology*, 78(6):1360–1380, 1973.
- 24 Luca Gugelmann, Konstantinos Panagiotou, and Ueli Peter. Random hyperbolic graphs: degree sequence and clustering. In *International Colloquium on Automata, Languages, and Programming*, pages 573–585. Springer, 2012.
- 25 Robert Jankowski, Antoine Allard, Marián Boguñá, and M Ángeles Serrano. The d-mercator method for the multidimensional hyperbolic embedding of real networks. *Nature Communications*, 14(1):7585, 2023.
- 26 Hao Jiang, Lixia Li, Yuanyuan Zeng, Jiajun Fan, and Lijuan Shen. Low-complexity hyperbolic embedding schemes for temporal complex networks. *Sensors*, 22(23):9306, 2022. doi:10.3390/S22239306.
- 27 Marcos Kiwi and Dieter 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, 2014.
- 28 Jon Kleinberg. Navigation in a small world. *Nature*, 406(6798):845–845, 2000.
- 29 Jon Kleinberg. The small-world phenomenon: An algorithmic perspective. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 163–170, 2000. doi:10.1145/335305.335325.
- 30 Robert Kleinberg. Geographic routing using hyperbolic space. In *IEEE INFOCOM 2007-26th IEEE International Conference on Computer Communications*, pages 1902–1909. IEEE, 2007. doi:10.1109/INFCOM.2007.221.
- 31 Dmitri Krioukov, Fragkiskos Papadopoulos, Marián Boguñá, and Amin Vahdat. Greedy forwarding in scale-free networks embedded in hyperbolic metric spaces. *ACM SIGMETRICS Performance Evaluation Review*, 37(2):15–17, 2009. doi:10.1145/1639562.1639568.
- 32 Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguñá. Hyperbolic geometry of complex networks. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 82(3):036106, 2010.
- 33 Vince Lehman, Ashlesh Gawande, Beichuan Zhang, Lixia Zhang, Rodrigo Aldecoa, Dmitri Krioukov, and Lan Wang. An experimental investigation of hyperbolic routing with a smart forwarding plane in ndn. In *2016 IEEE/ACM 24th International Symposium on Quality of Service (IWQoS)*, pages 1–10. IEEE, 2016. doi:10.1109/IWQoS.2016.7590394.
- 34 Bolun Liu, Shane Lubold, Adrian E Raftery, and Tyler H McCormick. Bayesian hyperbolic multidimensional scaling. *Journal of Computational and Graphical Statistics*, 33(3):869–882, 2024. doi:10.1080/10618600.2024.2308219.
- 35 Chip Martel and Van 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.
- 36 Stanley Milgram et al. The small world problem. *Psychology today*, 2(1):60–67, 1967.
- 37 Van Nguyen and Chip Martel. Analyzing and characterizing small-world graphs. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 311–320, 2005.
- 38 Fragkiskos Papadopoulos, Dmitri Krioukov, Marián Boguñá, and Amin Vahdat. Greedy forwarding in dynamic scale-free networks embedded in hyperbolic metric spaces. In *2010 Proceedings IEEE Infocom*, pages 1–9. IEEE, 2010.
- 39 Fragkiskos Papadopoulos, Constantinos Psomas, and Dmitri Krioukov. Network mapping by replaying hyperbolic growth. *IEEE/ACM Transactions on Networking*, 23(1):198–211, 2014. doi:10.1109/TNET.2013.2294052.
- 40 Andrés Sevilla and Antonio Fernández Anta. Routing in generalized geometric inhomogeneous random graphs. In *International Conference on Networked Systems*, pages 368–373. Springer, 2020.

- 41 Yuval Shavitt and Tomer Tankel. Hyperbolic embedding of internet graph for distance estimation and overlay construction. *IEEE/ACM Transactions on Networking*, 16(1):25–36, 2008. doi:10.1145/1373452.1373455.
- 42 Bendegúz Sulyok and Gergely Palla. Greedy routing optimisation in hyperbolic networks. *Scientific Reports*, 13(1):23026, 2023.
- 43 Jeffrey Travers and Stanley Milgram. An experimental study of the small world problem. In *Social networks*, pages 179–197. Elsevier, 1977.
- 44 Dongsheng Ye, Hao Jiang, Ying Jiang, Qiang Wang, and Yulin Hu. Community preserving mapping for network hyperbolic embedding. *Knowledge-Based Systems*, 246:108699, 2022. doi:10.1016/J.KNSYS.2022.108699.
- 45 Xiaohan Zhao, Alessandra Sala, Haitao Zheng, and Ben Y Zhao. Efficient shortest paths on massive social graphs. In *7th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom)*, pages 77–86. IEEE, 2011. doi:10.4108/ICST.COLLABORATECOM.2011.247162.