The Diameter of (Threshold) Geometric Inhomogeneous Random Graphs

Zylan Benjert
TU Delft, Netherlands

Kostas Lakis
ETH Zurich, Switzerland

Johannes Lengler
ETH Zurich, Switzerland

Raghu Raman Ravi
ETH Zurich, Switzerland

Abstract

We prove that the diameter of threshold (zero temperature) Geometric Inhomogeneous Random Graphs (GIRG) is asymptotically almost surely $\Theta(\log n)$. This has strong implications for the runtime of many distributed protocols on those graphs, which often have runtimes bounded as a function of the diameter.

The GIRG model exhibits many properties empirically found in real-world networks, and the runtime of various practical algorithms has empirically been found to scale in the same way for GIRG and for real-world networks, in particular related to computing distances, diameter, clustering, cliques and chromatic numbers. Thus the GIRG model is a promising candidate for deriving insight about the performance of algorithms in real-world instances.

The diameter was previously only known in the one-dimensional case, and the proof relied very heavily on dimension one. Our proof employs a similar Peierls-type argument alongside a novel renormalization scheme. Moreover, instead of using topological arguments (which become complicated in high dimensions) in establishing the connectivity of certain boundaries, we employ some comparatively recent and clearer graph-theoretic machinery. The lower bound is proven via a simple ad-hoc construction.

2012 ACM Subject Classification Mathematics of computing → Graph theory; Theory of computation → Distributed algorithms

Keywords and phrases GIRG, Diameter, Distributed Algorithms, Complex Networks

Digital Object Identifier 10.4230/LIPIcs.STACS.2026.11

Related Version *Full Version*: https://www.arxiv.org/abs/2510.12543 [2]

Funding K.L. and R.R.R. gratefully acknowledge support by the Swiss National Science Foundation [grant number 10003390].

1 Introduction

The diameter of a graph is the maximum graph distance over all pairs of vertices. In case the graph is disconnected, we only consider pairs in the same connected component. It has far-reaching implications on the performance of many distributed algorithms, where runtime bounds often depend on the diameter of the graph explicitly [36]. The literature is too vast to do it justice here, so we just mention two examples, leader election [24] and minimum spanning tree [36], where the runtimes of state-of-the-art algorithms depend on the diameter. Thus it is important to understand the diameter of real-world networks and their models.

43rd International Symposium on Theoretical Aspects of Computer Science (STACS 2026).
Editors: Meena Mahajan, Florin Manea, Annabelle McIver, and Nguyễn Kim Thắng
Article No. 11; pp. 11:1–11:21
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Many real-world networks commonly exhibit the following properties: power-law degree distributions [20], small-world distances (at most logarithmic) and large clustering coefficient[1] [40], low-dimensionality [21] (assuming some underlying space for the vertices), hierarchical structure [15], navigability[2] [9, 13], self-similarity [37] and so on. Motivated by this, numerous models for such networks have been introduced in the literature. One example, which we study in this work, is called *Geometric Inhomogeneous Random Graphs (GIRG)*. This model generalizes Hyperbolic Random Graphs (HRG) [11]. Roughly speaking, vertices are randomly placed in Euclidean space (usually in a torus) and sample weights from a power law distribution, see Definition 1.2 below. Then, edge probabilities increase with the product of the weights and decrease with the distance. We focus on the *threshold* case of this model (also called *zero temperature GIRG* or *T-GIRG*) in this work.

Our main motivation for studying the GIRG model stems from recent empirical works which show that the performance of various algorithms on GIRG closely matches that on real-world networks [4, 14]. This includes in particular bidirectional breadth-first search for computing shortest paths, the iFUB algorithm [16] for the diameter, the dominance rule for vertex cover, the Louvain algorithm for clustering [8], and the time required to enumerate all maximal cliques, and reduction rules for computing the chromatic number. Many of the above settings have also been analyzed theoretically on HRGs [1, 5–7]. While there are many other models for social and other complex real-world networks [18, 25, 26, 31–33], we believe that these transfer results make the GIRG model a particularly promising candidate for gaining theoretical insights into the performance of algorithms, as they may sometimes translate into practical performances on real-world networks.

Moreover, there is compelling empirical evidence for the quality of fit between GIRG and certain real-world networks at least when properties like (degree) heterogeneity, clustering, typical distances and so on are the object of focus [4, 17, 21]. For the special case of the HRG model, earlier works support the notion that real-world networks (for example a certain subset of the Internet) are amenable to embeddings in hyperbolic space [23, 31], and there is a mapping from such a space to 1-dimensional GIRG.[3]

Let us now formally define the model. We first describe a power-law distribution, to be used for the weights of the vertices, which are in expectation in the order of their degrees.[4]

▶ **Definition 1.1.** *Let* $\tau > 1$*. A continuous random variable* $X \geq 1$ *is said to follow a* power-law with exponent τ *if it satisfies* $\mathbb{P}[X \geq x] = \Theta(x^{1-\tau})$*.*

We now define the precise graph model we will work with. For simplicity, throughout we use $|u - v|$ to denote the max norm distance between the positions of vertices u and v. Any other norm would do, as it would only change the distances by a constant factor at most, and our proofs are robust enough to handle such deviations. Also, throughout the paper we will assume that the dimension d, the density λ and the power-law exponent τ are constant, while $n \to \infty$. Deviating slightly from the usual computer science convention, the expected number of vertices is λn instead of n, but this will not change any asymptotic results.

[1] Consider choosing a random vertex and then two random neighbors of it. The probability that the latter vertices share an edge is the clustering coefficient.

[2] As observed in the famous Milgram experiment [34], which became popularly known as "six degrees of separation".

[3] Roughly speaking, the angle of a vertex in the hyperbolic disk corresponds to its x coordinate, and the distance from the center corresponds (inversely) to its weight.

[4] There are more general versions of this definition, most notably incorporating so-called *slowly-varying* functions as a scaling factor [39], which allows for better statistical fit to some real-world networks. Our proofs can be extended to such a scenario using standard tools (e.g. Potter bounds [3]), but we keep the current definition for simplicity of notation.

▶ **Definition 1.2.** *Let $\tau > 2, \lambda > 0$, $d, n \in \mathbb{N}$, and let $\mathcal{D}$ be a power-law distribution on* $[1, \infty)$ *with exponent τ. A* Threshold Geometric Inhomogeneous Random Graph (T-GIRG) *is sampled as follows:*

1. *The vertex set $\mathcal{V}$ is given by a Poisson Point Process of intensity λ on the d-dimensional torus $[0, n^{1/d}]^d$ of volume n.*[5]
2. *Every vertex $u \in \mathcal{V}$ draws i.i.d. a* weight *$w_u \sim \mathcal{D}$.*
3. *For every two distinct vertices $u, v \in \mathcal{V}$, add an edge between u and v in $\mathcal{E}$ if and only if*

$$w_u w_v \geq |u - v|^d.$$

▶ Remark 1.3. This is only a special case of the GIRG model as defined in [11]. In the general case, the connection probability between two vertices scales as $(w_u w_v / |u - v|^d)^\alpha$, where the constant parameter $\alpha > 1$ is the *inverse temperature*. We use $\alpha = \infty$.

Previous works. The diameter of GIRG has already been shown to be polylogarithmic, but with an unidentified exponent [12]. The difference between this bound and ours is still substantial for many modern distributed settings. Very closely related to this work is that of Müller and Staps [35] which showed that the diameter of HRG is logarithmic, following previous polylogarithmic bounds [22, 28]. As already noted, HRG can be thought of as GIRG with $d = 1$, and in fact some part of our upper bound proof reuses elements from the HRG case. However, the proof in [35], in particular the use of blocking structures, relied heavily on dimension one. We also note that the resulting graphs for $d = 1$ are *structurally different*: for $d \geq 2$, the subgraph induced by vertices of degree at most D contains a giant (linear-sized) component if D is a sufficiently large constant. For $d = 1$, the same induced subgraph is scattered and subcritical for every constant D.

Organization. The following sections (upper bound in Sections 2 and 3 and lower bound in Section 4) contain the derivation of our results, with proof sketches and intuition provided together with any claim. Some formal proofs are omitted due to space constraints, but can be found on arXiv [2]. Altogether, we prove the following theorem. Traditionally, the most important regime is the one for $\tau \in (2, 3)$, for which we obtain a diameter of $\Theta(\log n)$.

▶ **Theorem 1.4.** *The following statements hold asymptotically almost surely.*[6]

- *If $\tau = 3$ and λ is a sufficiently large constant, then the diameter of T-GIRG is $\mathcal{O}(\log n)$.*
- *If $\tau < 3$, then the diameter of T-GIRG is $\mathcal{O}(\log n)$.*
- *If $\tau > 2$, then the diameter of T-GIRG is $\Omega(\log n)$.*

Proof. Upper bounds follow by Lemmas 2.4 and 2.17 and Corollaries 3.7 and 3.8. Lower bound follows by Theorem 4.4. ◀

2 Upper bound: large enough λ and $\tau \leq 3$

In this section we present the proof in a simplified setting, namely when the density λ is larger than some constant depending on d. This makes the proof considerably less technical and easier to follow. In the subsequent section we will then explain how this assumption can be removed (but only if $\tau < 3$). For the following, $\mathcal{V}_w := [0, n^{1/d}]^d \times [1, \infty)$ is the geometric ground space augmented with the weight dimension.

[5] Equivalently, we can sample the number of vertices from a Poisson distribution with expectation λn and draw their positions uniformly at random. This has mathematically slightly nicer properties than picking *exactly* λn vertices, in particular, the number and types of vertices in any two disjoint regions of the space are independent. However, the differences are negligible, see [30, Claims 3.2,3.3].

[6] We say that an event holds *asymptotically almost surely (aas)* if it holds with probability $1 - o(1)$.

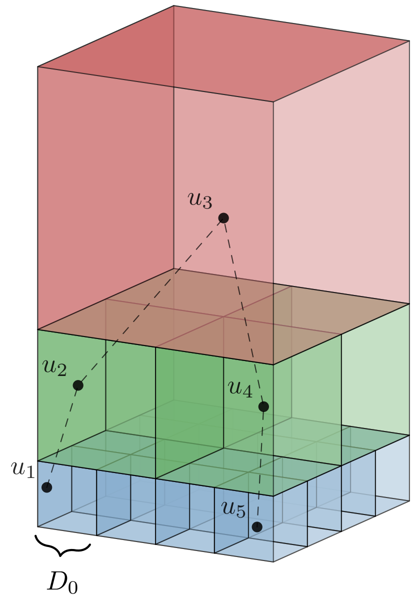

Figure 1 Visualization of the tessellation for $d = 2$, with 3 levels. The side length of lowest-level boxes is D_0. The blue boxes comprise $\mathcal{T}_0$ (weight range $[1, 2)$), the green ones are $\mathcal{T}_1$ (weight range $[2, 4)$) and the red box is $\mathcal{T}_2$ (weight range $[4, \infty)$). Also depicted is a vertex path arising from a fully active canonical box path (as in Definition 2.2). Note that this is a path in the GIRG graph, not in $\mathcal{B}$ or $\mathcal{B}^+$ from Definition 2.1. However, the *boxes containing* the vertices u_1 to u_5 are also connected by a path in $\mathcal{B}$ (and also $\mathcal{B}^+$). If u_1 were in the same box as u_5, this would no longer be true (not even for $\mathcal{B}^+$), because this box does not intersect the box containing u_2.

We follow some ideas from [35]. In particular, we tessellate $\mathcal{V}_w$ in tree-structured boxes as follows (see Figure 1). First, we fix a small constant D_0 such that $(n^{1/d})/D_0 = 2^{e_0(n)}$, with $e_0(n)$ being an integer. The lowest level of boxes is then defined as:

$$\mathcal{T}_0 = \left\{ \left(\prod_{k=1}^{d} [(j_k - 1)D_0,\ j_k D_0) \right) \times \left[1, 2^{d/2}\right) \;\middle|\; 1 \le j_k \le 2^{e_0(n)} \right\}.$$

In words, we tessellate[7] the geometric space in the natural way by boxes of side length D_0 and we take for each the product with the weight range from 1 to $2^{d/2}$. The choice of $2^{d/2}$ as the upper limit will be made clear soon.

We now similarly define the higher levels of the tessellation. The idea is that each box will have 2^d boxes of the next lowest level "directly below" it. That is, we define for $i < e_0(n)$:

$$\mathcal{T}_i = \left\{ \left(\prod_{k=1}^{d} \left[2^i (j_k - 1)D_0,\ 2^i j_k D_0\right) \right) \times \left[(2^{d/2})^i, (2^{d/2})^{i+1}\right) \;\middle|\; 1 \le j_k \le 2^{e_0(n)-i} \right\}.$$

The highest level of the hierarchy consists of a single box covering the entire geometric ground space and the leftover weight range. That is, $\mathcal{T}_{e_0(n)} = \mathcal{V} \times \left[(2^{d/2})^{e_0(n)}, \infty\right)$.

[7] To have a proper partition of the ground space, we actually replace $[(j_k - 1)D_0,\ j_k D_0)$ by $[(j_k - 1)D_0,\ j_k D_0]$ whenever j_k is the maximum permitted value. We only write this in this footnote for simplicity of the formulas.

Now, notice that any two adjacent (as in Definition 2.1 below, in any of the two interpretations) boxes B_1, B_2 in this tessellation have the following property. For any vertices x, y such that $x \in B_1$ and $y \in B_2$, we have $w_x w_y \geq |x - y|^d$ (for D_0 small enough).

This leads to the following observation. Consider any pair of vertices u and v. From the box B_u of u there is a "canonical path" of boxes to the box B_v of v, which is comprised of the two "upward" paths from B_u and B_v to their lowest common ancestor. If it were the case that all the boxes in this path contained at least one vertex, we would obtain a vertex-path (see Figure 1) from u to v of length $\mathcal{O}(\log n)$. To make this and the following more precise, we give the following definitions (here Figure 1 might be helpful again).

▶ **Definition 2.1.** *Consider the boxes of the tessellation of $\mathcal{V}_w$ as a vertex set. On this vertex set we define $\mathcal{B}$ by adding an edge between boxes B_1, B_2 if their closures intersect in a d-dimensional set.*[8] *We also define $\mathcal{B}^+$ on the same vertex set by adding edges when the closures of said boxes have a non-empty intersection.*

Of course, the previously mentioned canonical path may fail to contain a vertex in each box. Nevertheless, we can utilize this idea to construct short paths between any two vertices u and v, given that they are connected. To this end, let a box be called *active* if it contains at least one vertex in the realization of the GIRG graph. For two boxes B_1, B_2, let $L(B_1, B_2)$ be the canonical box path (not necessarily active) connecting them. More formally:

▶ **Definition 2.2** (Canonical box path)**.** *Let B_1, B_2 be two distinct boxes of the tessellation. For any box B other than the top-most box B_{top}, let $\pi(B)$ be the unique box of minimum volume under the constraint that it is a strict superset of B under projection along the weight dimension. Notice that $\{B, \pi(B)\} \in E(\mathcal{B})$. Now, the set of edges identified in this way constitutes a spanning tree $T_{\mathcal{B}}$ of $\mathcal{B}$, rooted at B_{top}. We define the canonical box path $L(B_1, B_2,)$ from B_1 to B_2 as the unique path between them in $T_{\mathcal{B}}$. We may sometimes abuse notation and use $L(u, v)$ to mean $L(B_u, B_v)$.*

It turns out that we can bound the length of a shortest path between u and v by the cardinality of the set $W(B_u, B_v)$, defined below. Intuitively this set is constructed by taking $L(B_u, B_v)$ and "expanding" along inactive boxes.

▶ **Definition 2.3** (Inactive region of canonical box path)**.** *Let B_1, B_2 be two distinct boxes of the tessellation. Let R be the subset of inactive boxes. Let $C_1^R, C_2^R, \ldots, C_c^R$ be the connected components of $\mathcal{B}^+[R]$. We define $W(B_1, B_2)$ as*

$$W(B_1, B_2) := L'(B_1, B_2) \cup \left\{ \bigcup_{C_i^R \cap L'(B_1, B_2) \neq \emptyset} C_i^R \right\},$$

where $L'(B_1, B_2) = L(B_1, B_2) \cup \{B \mid B \text{ is } \mathcal{B}^+\text{-adjacent to some } B' \text{ in } L(B_1, B_2)\}$. We may sometimes abuse notation and use $W(u, v)$ to mean $W(B_u, B_v)$. Notice that $W(B_1, B_2)$ is a $\mathcal{B}^+$-connected set of boxes. Additionally, any box $B \notin W(B_1, B_2)$ but $\mathcal{B}^+$-adjacent to some box in $W(B_1, B_2)$ must be active. We use $S(B_1, B_2)$ to refer to the set of such boxes.

To show a logarithmic upper bound on the diameter, it now suffices to show two statements. First, we will essentially show that if u and v are connected in the GIRG graph G, then G^3 contains a path from u to v where all vertices are in boxes of $W(u, v) \cup S(u, v)$.[9]

[8] Mind that $\mathcal{V}_w$ has dimension $d + 1$.
[9] G^3 is a graph where we add an edge between two vertices if and only if their distance is at most 3 in G.

As a consequence, the length of a shortest path between any pair of vertices u and v is deterministically $\mathcal{O}(|W(u,v) \cup S(u,v)|) = \mathcal{O}(|W(u,v)|)$. Second, asymptotically almost surely it holds that $|W(u,v)| = \mathcal{O}(\log n)$ for all pairs u, v. More precisely, we show the following.

▶ **Lemma 2.4.** *Assume $\tau \leq 3$ and that λ is larger than some constant depending on d. Then, asymptotically almost surely we have for some constant C that $|W(B_1, B_2)| \leq C \log n$ for all pairs B_1, B_2 of boxes.*

Proof. Fix a single pair B_1, B_2. Suppose that for these boxes, $|W(B_1, B_2)| > C \log n$. For large enough C, we can conclude that at least half the boxes in $W(u,v)$ must be inactive (some are included by virtue of being part of $L(u,v)$ and these need not be inactive). Let I refer to the set of these inactive boxes.

In [10, pp. 129–130], it is shown that in a graph with n vertices and maximum degree Δ, there are at most $n(e(\Delta-1))^k$ connected subsets of size k. We have $\Delta = 3^d - 1$. Therefore, there are at most $n(e(3^d-2))^{|W(B_1,B_2)|}$ possibilities for $W(B_1, B_2)$ that we need to consider (i.e union bound over). Once we fix such a $W(B_1, B_2)$, there are (by a gross overestimation) at most $2^{|W(B_1,B_2)|}$ ways to choose I. The probability that this choice of I is actually fully inactive is at most $p_{in}^{|W(B_1,B_2)|/2}$, where p_{in} is an upper bound on the probability of a single box not containing a vertex. Since $\tau \leq 3$, this probability is maximized at the lowest-level boxes. By increasing λ, we can assume p_{in} is as small as desired, in particular small enough such that $2e(3^d-2)p_{in}^{1/2} < 1/e$. This implies that for some large enough λ we can union bound over all possible sets $W(B_1, B_2)$ with $|W(B_1, B_2)| > C \log n$ and all possible choices for I, and thus the claim follows. ◀

To prove also the first statement, we must lay some groundwork regarding $\mathcal{B}$ and $\mathcal{B}^+$.

Boundary connectivity

Now we prove some useful structural properties about the graphs $\mathcal{B}$ and $\mathcal{B}^+$. In particular, we will establish some preconditions of theorems proven in [38]. We will define the "visible boundary" of a set C of boxes with respect to a given box B as the set of $\mathcal{B}^+$-neighbors of C that are $\mathcal{B}$-reachable from B without crossing C, see below. The goal is then to show that the visible boundaries of $\mathcal{B}^+$-connected sets are $\mathcal{B}$-connected. Applying the result for $C := W(u,v)$ will be useful later in the proof, where we construct shortcuts for subpaths that venture outside of $W(u,v) \cup S(u,v)$.

To illustrate the idea further and provide some intuition, let us consider for now such boundaries in the simplified case of $\mathbb{Z}^2$, where vertices are points of the infinite integer grid and nearest neighbors share an edge (in the counterpart of $\mathcal{B}$), as in Figure 2. For the analogue of $\mathcal{B}^+$, which we will call $\mathbb{Z}^{2*}$, we also add the edges between vertices whose x and y coordinates both differ by 1. Each vertex of this infinite graph should intuitively be thought of as a box in our tessellation, even though our tessellation does not possess the edge set of $\mathbb{Z}^2$ or $\mathbb{Z}^{2*}$. Assume now that we are presented with a blue/red coloring of the vertices of $\mathbb{Z}^2$, where blue is generally meant to signify an active box and red generally signifies an inactive box. Given such a coloring with the origin being red, we can define the "red component" C_{red} (comparable to $W(u,v)$ for our purposes) of the origin as its connected component in $\mathbb{Z}^{2*}$ induced by red vertices. This may possibly yield a red region without "holes", surrounded by a connected (in $\mathbb{Z}^2$) blue boundary (as in Figure 2, left). However, this is not the only case. For example, C_{red} could resemble a donut, in which case the blue boundary consists of two connected circles. Now, the points inside the donut hole "see" a connected part of the boundary and so do the points outside the donut. This is a crucial property which we will use in the following way. Thinking of the red region as $W(u,v)$ and of the blue boundary as $S(u,v)$, consider the graph path connecting u to v. If the path stays within the union of the red region and the blue boundary, it is already short enough and there is nothing we need to do. If, however, this walk ventures outside of $W(u,v) \cup S(u,v)$, we will show that there is a shortcut walking along the visible boundaries to obtain a new walk that avoids such excursions, except perhaps for the first and last vertices of each. More exotic cases can arise when the path successively jumps through a sequence of different holes, but these can also be tamed as we will see.

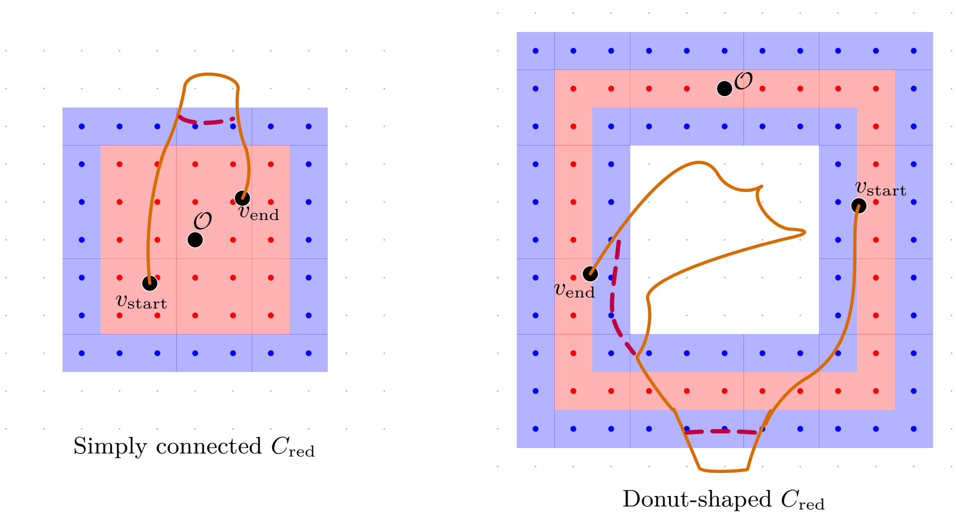

Figure 2 Two examples of a red connected component C_{red} and its blue boundary with paths that start and end in C_{red}. Left: one outward excursion; the shortcut stays inside the connected blue boundary. Right: first an outward excursion near the bottom, then a separate excursion into the hole; each has a shortcut confined to the corresponding connected part of the blue boundary.

With the previous example guiding our intuition and keeping in mind that there are new subtleties to uncover, let us now define visible boundaries in accordance with [38]. As we have mentioned, we will use these boundaries to “patch” parts of our constructed path, using their connectivity, which we will prove later.

▶ **Definition 2.5.** *For a $\mathcal{B}^+$-connected subset of C and a box B, we define the* boundary of C visible from B *as*

$$\partial_{vis(B)}(C) = \{B' \mid B' \textit{ is } \mathcal{B}^+\textit{-adjacent to some } B'' \in C \textit{ and } B' \textit{ is connected to } B \textit{ in } \mathcal{B} \setminus C\}$$

Notice that we require B' to be $\mathcal{B}^+$-adjacent to C, but we only allow paths in $\mathcal{B}$ for the connection to B. This is required for the theorems in [38] to apply, but otherwise plays little role for our analysis.

Before introducing the theorem we will use, let us describe the preconditions we need to prove. The first one is with regards to the *cycle space* of $\mathcal{B}$. The cycle space is simply the set of subgraphs of G where each vertex has an even degree. Our goal is to find a suitable generating set $\Gamma_{\mathcal{B}}$ for this space, with symmetric difference of edges being the underlying algebraic operation. For this first criterion, we regard a generating set as *suitable* if all $C \in \Gamma_{\mathcal{B}}$ are chordal in $\mathcal{B}^+$. We say that a cycle is *chordal* in $\mathcal{B}^+$ if its vertices induce a clique in $\mathcal{B}^+$. For reference, such a set for the pair $\mathbb{Z}^2, \mathbb{Z}^{2*}$ is the set of all 4-step cycles.

It is a standard fact that so-called "fundamental cycles" generate the cycle space of a graph [19, Theorem 1.9.5]. To produce a set of fundamental cycles, pick a spanning forest F. Then, for any $e \in G \setminus F$, $F \cup \{e\}$ has a unique cycle. The set of cycles found like so is referred to as a set of fundamental cycles.

It therefore suffices to find a generating set for a set of fundamental cycles. To do so, consider the natural spanning tree $T_{\mathcal{B}}$ of $\mathcal{B}$, comprised of all edges of $\mathcal{B}$ connecting boxes of different levels (see Figure 3). We will observe that fundamental cycles of this tree follow a very restricted structure which will guide us into defining $\Gamma_{\mathcal{B}}$.

▶ **Observation 2.6.** *All fundamental cycles of $T_{\mathcal{B}}$ can be expressed as the union of e, P_{B_1} and P_{B_2}, where $e = \{B_1, B_2\}$ is the added edge and P_{B_i} is the unique path in $T_{\mathcal{B}}$ from B_i to the lowest common ancestor of B_1, B_2, assuming $T_{\mathcal{B}}$ is rooted at the highest level box. Moreover, for each B'_1 and B'_2 in the same level of the tessellation and with B'_i internal in P_i, we have $\{B'_1, B'_2\} \in E(\mathcal{B})$.*

Proof. The first part of the observation follows from the fact that the described union constitutes a cycle and by uniqueness of said cycle. For the second part, notice that for any two boxes in the same level of the tessellation with $\{B'_1, B'_2\} \notin E(\mathcal{B})$, we know that the projections of the boxes along the weight dimension do not intersect, even after taking their closures. Since all descendants of these boxes have projections that are subsets of the above, no pair of them intersects either, and thus the claim follows by contraposition. ◀

We are now in a position to define our generating set.

▶ **Definition 2.7.** *Let $e = \{B_1, B_2\} \in E(\mathcal{B})$ be an edge between two boxes of the same level. Let $B'_i = \pi(B_i)$ be the parent box of B_i. If $B'_1 = B'_2 = B'$, then let $\gamma(e) = (B_1, B', B_2, B_1)$, otherwise let $\gamma(e) = (B_1, B'_1, B'_2, B_2, B_1)$. We define $\Gamma_{\mathcal{B}}$ as the set containing all such cycles.*

▶ **Lemma 2.8.** *The set $\Gamma_{\mathcal{B}}$ generates the cycle space of $\mathcal{B}$. Moreover, each element of $\Gamma_{\mathcal{B}}$ is chordal in $\mathcal{B}^+$.*

Proof. Recall that it suffices to show that $\Gamma_{\mathcal{B}}$ generates all fundamental cycles of $T_{\mathcal{B}}$. Let C be such a cycle where (see also Figure 3)

$$C = (B_{high}, B_1^L, B_2^L, \ldots, B_k^L, B_k^R, B_{k-1}^R, \ldots, B_1^R, B_{high})$$

with $\{B_k^L, B_k^R\}$ being the added edge. Recall that by Observation 2.6 we have $\{B_i^L, B_i^R\} \in E(\mathcal{B})$, meaning $\gamma(\{B_i^L, B_i^R\}) \in \Gamma_{\mathcal{B}}$ for all i. Consider $S = \sum_i \gamma(\{B_i^L, B_i^R\})$. Notice that all edges of C appear exactly once in this sum. Additionally, all other edges ($\{B_j^L, B_j^R\}$ for $j < k$) in the sum appear exactly twice. Therefore, $S = C$, yielding the first claim.

We will now show the chordality in $\mathcal{B}^+$ of elements of $\Gamma_{\mathcal{B}}$. Consider an element $\gamma(e) \in \Gamma_{\mathcal{B}}$. Because the edge e is incident to boxes of the same level, we conclude that these boxes can be characterized by vectors J_1 and J_2 as in the definition of the tessellation, where J_1 and J_2 differ by 1 in exactly one of d coordinates. Therefore, their parents also satisfy this property, or coincide. In either case, one can identify a $(d-1)$-dimensional set of points in $d+1$ dimensions, defined by restricting the weight to be the lower limit of the parents and the coordinate in which J_1 and J_2 differ to be equal to the value where the closures of the boxes of e intersect. All other coordinates are free to take any value in the intervals prescribed by the boxes of e. This set of points guarantees that the considered cycle is chordal in $\mathcal{B}^+$. ◀

With the condition on $\Gamma_{\mathcal{B}}$ satisfied for $\mathcal{B}$ and $\mathcal{B}^+$, let us present the theorem we want to use, which is an adapted version of Theorem 4 in [38].

▶ **Theorem 2.9.** *Let $\mathcal{B}^+$ be a connected graph, and $\mathcal{B}$ a connected subgraph of $\mathcal{B}^+$. Suppose that there is a generating set $\Gamma_\mathcal{B}$ for the cycle space of $\mathcal{B}$ that is chordal in $\mathcal{B}^+$, and that for every edge $e \in \mathcal{B}^+$ there is a cycle O_e (containing e) in $\mathcal{B}^+$ such that $O_e \setminus e \subset \mathcal{B}$, and O_e is chordal in $\mathcal{B}^+$. Let C be a connected subgraph of $\mathcal{B}^+$, and $x \in V(\mathcal{B}) \setminus C$. Then $\partial_{vis(x)}(C)$ is connected in $\mathcal{B}$.*

We notice the extra condition about the existence of the cycles O_e. Let us now record this condition for our particular pair $\mathcal{B}, \mathcal{B}^+$, in a way which will be useful later on as well.

▷ Claim 2.10. Let B_1, B_2 be two boxes connected by an edge e in $\mathcal{B}^+$. Consider the graph G_{local} with vertex set $V = \{B \mid d_{\mathcal{B}^+}(B, B_1) \leq 1\}$ and edge set $E(\mathcal{B}) \cap V^2 \setminus \{e\}$. Then, G_{local} contains a path from B_1 to B_2. Additionally, this path along with e constitutes a cycle O_e as prescribed by Theorem 2.9.

With Theorem 2.9 in our arsenal, we may proceed to relate $|W(u, v)|$ to the path length.

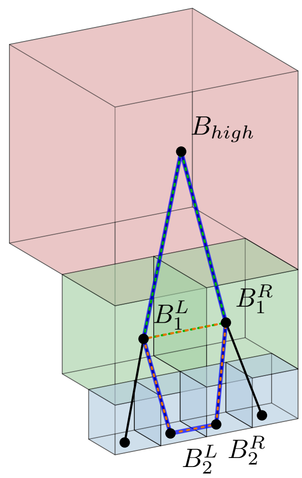

Figure 3 Part of the spanning tree $T_\mathcal{B}$. The fundamental cycle (blue) formed by adding the edge $\{B_2^L, B_2^R\}$ is the symmetric difference of the (green) triangle (B_1^L, B_{high}, B_1^R) and the (orange) quadrilateral $(B_2^L, B_1^L, B_1^R, B_2^R)$, both in $\Gamma_\mathcal{B}$.

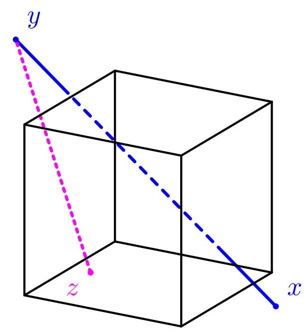

Figure 4 An active box intersected by an edge $\{x, y\}$. The vertex z inside the box is guaranteed to connect to one of the endpoints (by Lemma 2.11), in this case y. Note that the "vertical" coordinate here signifies weight.

Bounding the shortest path length by $|W(u, v)|$

As mentioned earlier, our approach will be to shortcut parts of (assumed) shortest paths which stray meaningfully out of $W(u, v)$. The following lemma is paramount for this goal. Intuitively, it allows us to "anchor" the path to the boundary of $W(u, v)$ whenever it is crossed. See Figure 4 for a visualization.

▶ **Lemma 2.11.** *Let D_0 in the definition of the tessellation be small enough. Let x, y be two vertices connected by an edge intersecting an active box B. Then, B contains a vertex z that is connected to either x or y by an edge.*

Let us now describe on a high level the rest of the proof. See also Figure 5 for a visual aid to the following arguments. Consider the sequence of boxes induced by a shortest path between u and v. That is, for each vertex in the path, write down the box in which it is contained. For later use, let us give a proper definition.

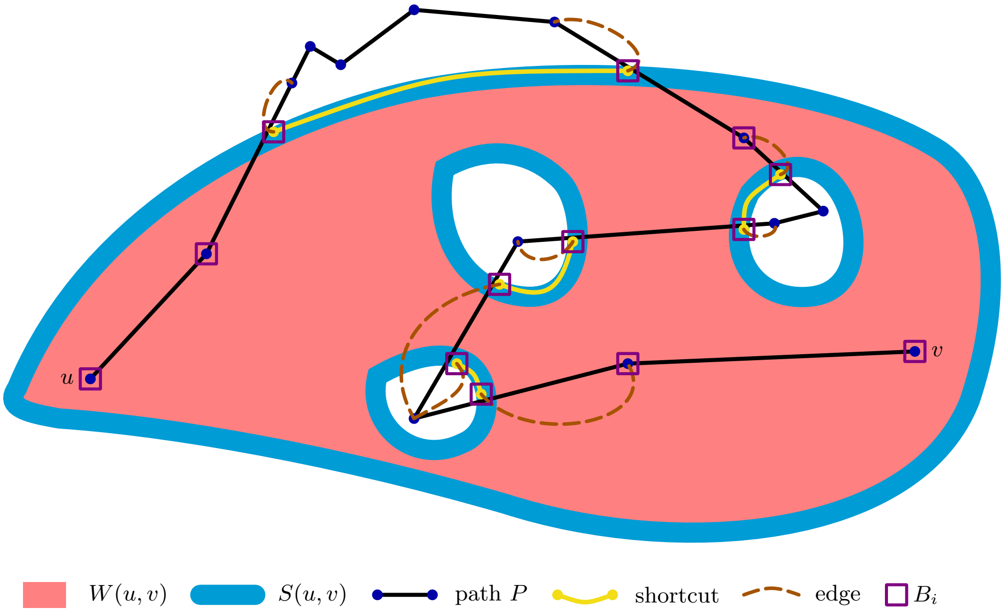

Figure 5 Pictorial overview of the proof of Lemma 2.17. The original path P initially exits $W(u,v) \cup S(u,v)$ and reenters it without visiting another hole in the meantime. For this excursion, Lemma 2.11 provides two edges (dashed orange) we can use to "anchor" the path to the outer boundary. Afterwards, the path visits three consecutive holes without intersecting $W(u,v) \cup S(u,v)$. Bridges connecting the visible boundaries of these holes are found using Claim 2.16. We finally identify a sequence of boxes $B_0, B_1, \ldots, B_k$ which are all in $W(u,v) \cup S(u,v)$ and contain a path from u to v. In more detail, any two consecutive boxes contain vertices with graph distance at most 3. For this, the worst case is in passing from the first hole to the second, where we start from the box on the first boundary, hop to one endpoint of the crossing edge, take the edge from one hole to the next and then finally hop from the other endpoint to the box on the second boundary. The boxes $B_0, B_1, \ldots, B_k$ are comprised of the violet "checkpoint" boxes and of the boxes one encounters when tracing the yellow shortcuts between them in $S(u,v)$.

▶ **Definition 2.12** (Box-footprint). *Let $P = u_0, u_1, \ldots, u_k$ be a path in the GIRG graph from u to v. We call the sequence $B_{u_0}, B_{u_1}, \ldots, B_{u_k}$ the* box-footprint *of P.*

This sequence is not necessarily a walk in $\mathcal{B}^+$ (two adjacent vertices could be in non-interescting boxes). The crucial property is that a path from u to v exists which only uses a number of vertices comparable to the number of boxes in this sequence. Our goal will be to transform this sequence in such a way that no boxes outside $W(u,v) \cup S(u,v)$ are contained in it, while preserving the path property. We have already discussed the general idea behind such modifications and depicted it in Figure 2. We will use Lemma 2.11 to establish "anchors" in the visible boundaries of $W(u,v)$ and use their connectedness to show the existence of paths between them. However, even the right panel of Figure 2 still does not impart the full picture. There, we see two excursions being patched, one in the outside region and one inside the donut hole. But the path could also jump between holes before eventually reentering $W(u,v) \cup S(u,v)$ (as in Figure 5), in which case our argument temporarily breaks down. Fortunately, this then implies that some edge of the path must cross *both* boundaries. Using Lemma 2.11 for the two boxes guaranteed to be crossed (each in one of the boundaries), we conclude that the two boundaries are actually still "connected" in the GIRG graph, in the sense that they contain vertices u_1 and u_2 respectively which are separated by a constant number of edges in the GIRG graph. So, no matter how erratic an excursion might look, the set of visited "holes" is sufficiently well-knit, and therefore we can still find a sequence of boxes preserving the path.

Let us now initiate a formal treatment of the above paragraph. We start with a definition.

▶ **Definition 2.13** (Hole). *Consider two vertices u, v. A* hole *is a connected component of $\mathcal{B} \setminus W(u,v)$. The* visible boundary *of a hole H is $\partial_{vis(x)}(W(u,v))$, for some x in H. The precise choice of x is irrelevant. Notice that all boxes in the visible boundary of H are active by construction of $W(u,v)$.*

We now wish to establish the first "anchor" when the path from u to v initially exits $W(u,v) \cup S(u,v)$. To do so, we first need another definition and a claim that will be useful throughout the rest of the proof.

▶ **Definition 2.14** (Box-shadow). *Let u, v be two vertices connected by an edge e. Let S be the line segment in $\mathcal{V}_w$ connecting the images of the vertices. We call the sequence of boxes intersected by S (which is a walk in $\mathcal{B}^+$) the* box-shadow *of $\{u,v\}$ and denote it by $\sigma(e)$.*

▷ Claim 2.15. Let B_W be a box in $W(u,v) \cup S(u,v)$ and B_H be a box in some hole H. Any walk in $\mathcal{B}^+$ from B_H to B_W intersects the visible boundary of H.

Proof. Consider the boxes $B_H = B_0, B_1, \ldots, B_k = B_W$ of such a walk. Because B_W is in $W(u,v)$ or is a box that is $\mathcal{B}^+$-adjacent to $W(u,v)$, there exists a minimum i such that $d_{\mathcal{B}^+}(B_i, W(u,v)) = 1$. We claim that B_i is in the visible boundary of H. By construction, B_i is $\mathcal{B}^+$-adjacent to $W(u,v)$. It then suffices to show that B_i is $\mathcal{B}$-connected to B_H, avoiding $W(u,v)$. Now, this follows because for every edge $\{B_j, B_{j+1}\}$ of $\mathcal{B}^+$ in the assumed walk with $j \leq i-1$, we know that $d_{\mathcal{B}^+}(B_j, W(u,v))$ is larger than 1. Therefore, B_j and B_{j+1} are connected in $\mathcal{B} \setminus W(u,v)$. This follows from Claim 2.10. Indeed, the claim provides us with a path from B_j to B_{j+1} using only edges from $\mathcal{B}$ and where all the boxes satisfy have distance in $\mathcal{B}^+$ at most 1 from B_j. Therefore, they have distance at least 1 from $W(u,v)$, i.e. they belong to the vertex set of $\mathcal{B} \setminus W(u,v)$. ◁

Thus, to establish our first anchor, we identify the first box in the box-footprint that is not in $W(u,v) \cup S(u,v)$. It therefore belongs to some hole H. Now, the box just before is in $W(u,v) \cup S(u,v)$. Therefore, the box-shadow of the corresponding edge must intersect the visible boundary of H by Claim 2.15. The box found to intersect must be active, and therefore contains a vertex connected to some endpoint of said edge by Lemma 2.11.

We have our first anchor, but we need to connect it all the way through to the next box of the box-footprint which will be in $W(u,v) \cup S(u,v)$ (such a box is guaranteed to exist because v is in $W(u,v)$). To this end, we introduce the following claim.

▷ Claim 2.16. Let u_{H_1}, u_{H_2} be two vertices connected by an edge e and suppose that the boxes containing them belong to different holes H_1, H_2. Then, there exists a vertex v_1 in some box B_1 in the visible boundary of H_1 and a vertex v_2 in some box B_2 in the visible boundary of H_2 with $d_{GIRG}(v_1, v_2) \leq 3$.

We now have all the tools necessary to prove the main lemma of this subsection.

▶ **Lemma 2.17.** *Suppose u and v are connected in the GIRG graph. There exists a sequence of boxes $B_0, B_1, \ldots, B_k$, entirely contained in $W(u,v) \cup S(u,v)$, with the following property. For every i, there is a vertex u_i in B_i with $u_0 = u$ and $u_k = v$ such that $d_{GIRG}(u_j, u_{j+1}) \leq 3$ holds for all $j \in 0, \ldots k-1$. In particular, $d_{GIRG}(u,v) \leq \mathcal{O}(|W(u,v)|)$.*

Proof. Let $P = u'_0, u'_1, \dots, u'_{k'}$ be the path assumed to exist between u and v, and let $B'_0, B'_1, \dots, B'_{k'}$ be its box-footprint. Let $i < j - 1$ be such that B'_i and B'_j are in $W(u,v) \cup S(u,v)$, but all B'_l for $i < l < j$ are not. Let H_{init} be the hole which contains B'_{i+1} and let H_{fin} be the one which contains B'_{j-1} (with the possibility that $H_{init} = H_{fin}$). We will show three things that let us replace any such subsequence $B'_{i+1}, \dots, B'_{j-1}$.

1. The visible boundary of H_{init} contains a box B_{init} with a vertex u_{init} which satisfies $d_{GIRG}(u_{init}, u'_i) \leq 3$, where u'_i is the vertex which generated B'_i.
2. A completely analogous statement holds for H_{fin} and B'_j.
3. There exists a sequence $B_{init}, \dots, B_{fin}$ satisfying the properties listed in the statement of the lemma (treating u_{init} and u_{fin} as the start and end vertices).

For the first statement, we use Claim 2.15, providing the box-shadow of the edge $\{u'_i, u'_{i+1}\}$ as the walk to be intersected. The bound on $d_{GIRG}(u_{init}, u'_i)$ follows because u_{init} is the vertex in the active box found in the intersection and therefore it connects to either u'_i or u'_{i+1} by Lemma 2.11 (actually the bound here is 2 rather than 3). The second statement follows completely analogously.

For the third statement, we employ Claim 2.16. Indeed, in the subsequence $B'_{i+1}, \dots, B'_{j-1}$, one may visit a sequence of holes $H_1, H_2, \dots, H_h$. Claim 2.16 then guarantees that for any two consecutive holes, their visible boundaries contain representative boxes that are (in 3 steps) connected through GIRG vertices. Therefore, a suitable sequence of boxes from B_{init} to B_{fin} exists, since the visible boundaries are themselves connected, comprised solely of active boxes and adjacent boxes have vertices that induce cliques in the GIRG graph. ◀

Bounding $|W(u,v)|$

The final step in bounding the diameter is controlling the size of the inactive region itself. We will do this with a simple path-counting argument. For some $W(u,v)$ to grow large, there must occur large connected sets of generally inactive boxes. However, because the degree of the box-adjacency graph is bounded, there are exponentially many (in k) such possible sets of size k and the probability of any fixed one being mostly inactive decays exponentially in k. Moreover, the strength of the decay depends on λ, whereas the scale of the combinatorial explosion does not. This Peierls-type idea is essentially the same as in [35, Lemma 13] and is the cornerstone to the proof of our Lemma 2.4.

3 Upper bound: $\tau < 3$ and arbitrary λ

In this section we will forego the assumption that λ is large enough and make do with $\tau < 3$.

Thus, we can no longer assume that each box is active with probability arbitrarily close to 1. We can now only do that for boxes above a certain level.[10] We now have to somehow deal with the boxes of lower levels. To manage this, we group these boxes into what we will call "towers".[11] A tower simply consists of the union of a box and all boxes "below" it.

[10] This is because the number of vertices in a box at level i is Poisson distributed with mean $\Theta\left(2^{id(1-(\tau-1)/2)}\right)$, which rapidly increases with growing i.

[11] This idea is similar in spirit to the renormalization carried out in [35], but requires a different approach. There, the renormalization is essentially reliant on the fact that the percolation threshold for 1D site percolation is 1, which is no longer true in higher dimensions.

▶ **Definition 3.1** (Towers). *Given a level i, we define the towers of our tessellation as:*

$$\mathbb{T} \;=\; \left\{ \left(\prod_{k=1}^{d} \left[2^i(j_k-1)D_0,\; 2^i j_k D_0\right) \right) \times \left[1, (2^{d/2})^{i+1}\right) \;\middle|\; 1 \le j_k \le 2^{e_0(n)-i} \right\}.$$

Notice that the sole difference between $\mathcal{T}_i$ and $\mathbb{T}$ is the lower end of the weight range. We now define what it means for a tower to be active.

▶ **Definition 3.2** (Active tower). *Fix some $\varepsilon > 0$. Let $T_0 \in \mathbb{T}$ with W the upper end of the weight range of T_0. Ensure that W^ε is an integer power of $(\sqrt{2})^d$ (for technicalities below) by potentially decreasing ε. Let $T_1, T_2, \ldots, T_{3^d-1}$ be the towers neighboring T, $F = \bigcup_{i=1}^{3^d-1} T_i$ and let G_F be the graph induced by vertices in F. We say that T_0 is* active, *if:*

1. *Any box (as in Definition 2.2) that is part of F and whose lower weight range is at least W^ε is active.*[12]
2. *In G_F, all vertices of weight at least W^ε are in the same connected component $\mathcal{C}$.*
3. *In G_F, for any component $\mathcal{C}' \neq \mathcal{C}$, if $u, v \in \mathcal{C}'$, then $|u-v| \le W^{\Theta(\varepsilon)}$.*

For the rest of this section, we redefine $L(u,v), W(u,v), S(u,v)$ in the following way. We simply replace the notion of boxes by elements which can be either towers or boxes of level above the tower cutoff. This still defines essentially the same graphs $\mathcal{B}, \mathcal{B}^+$, up to a reduction in the number of levels. The definition of $L(u,v)$ is still based on the spanning tree of $\mathcal{B}$, but we consider as start/end vertices in said graph either the tower containing u/v or the boxes which contain them if their weights are large enough. The definitions of $W(u,v), S(u,v)$ simply inherit the effects of the change in $L(u,v)$. The goal will now be to re-establish the "shortcuts" through $S(u,v)$ and again show that $W(u,v)$ does not grow too large.

To that end, we first show that the higher the tower, the more likely it is to be active.

▶ **Lemma 3.3.** *Let $T_0 \in \mathbb{T}$ be a tower of level i. Then,*

$$\mathbb{P}[T_0 \textit{ is active}] \ge 1 - o(1),$$

where the $o(\cdot)$ is with respect to i.

Proof. It suffices to show that properties 1 and 3 fail to be satisfied each with $o(1)$ probability. For the first one, notice that the probability for any of the boxes of interest to be inactive is maximized for the lowest-level ones. Therefore, a fixed box is inactive with probability at most $\exp\left(-W^{\Theta(\varepsilon)}\right)$. It is not hard to see that there are polynomially many (in W) boxes.[13] Thus, a union bound shows the claim.

Now, for property 3, we will expose G_F in two steps. First, reveal all vertices of weight at most W^ε and the edges between. We then reveal the higher-weight vertices and all remaining edges. Up to a $o(1)$ failure probability (from property 1), in the second step we only merge components from the first step to $\mathcal{C}$. Therefore, it suffices to show that any component $\mathcal{C}'$ from the first step violating property 2 will merge with $\mathcal{C}$ in the second step.

Indeed, let $\mathcal{C}'$ be such a component and let u, v be the vertices satisfying $|u-v| \ge W^{C\varepsilon}$ for a constant C to be determined. Consider a shortest path between u and v. This path must contain at least $W^{C\varepsilon - 2\varepsilon}$ edges, since each edge can only cover a geometric distance at most $W^{2\varepsilon}$. Consider every other vertex (i.e. the first, the third, the fifth and so on) in this path

[12] Note that this implies the next property.
[13] Because each box has 2^d "children" below it and we reach the lowest level boxes within $\mathcal{O}(\log W)$ iterations.

and place a ball of radius 1/2 around each one. If any pair of these balls would intersect, we could shortcut the path, which would contradict its minimum length. Therefore, we conclude that a subset of the vertices in the path cover a geometric area of volume $\Omega\big(W^{C\varepsilon-2\varepsilon}\big)$. Now, let A be that area. If in the second step we reveal any vertex in A, the component $\mathcal{C}'$ will merge with $\mathcal{C}$ (as we will have an edge from some vertex of the path to one vertex in $\mathcal{C}$). Now, the number of vertices of weight at least W^ε in A is Poisson distributed with expectation $\Omega\big(W^{C\varepsilon-2\varepsilon-(\tau-1)\varepsilon}\big)$. It follows that for large enough C (in particular as long as $C > \tau+1$), the probability that none are revealed is at most $\exp\big(-W^{\Theta(\varepsilon)}\big)$. This argument was for a singular $\mathcal{C}'$. However, there can be at most polynomially many (in W) candidates, so again a union bounds suffices. This bound on the number of components is because the total volume of F is polynomial in W, and each $\mathcal{C}'$ would need to disjointly occupy at least some constant volume (much like in the shortest path argument). ◀

As a first observation towards re-establishing the $S(u,v)$ "shortcuts", let us remark that a version of Lemma 2.11 still holds when we replace some box levels by towers. In more detail, if it so happens that an edge from x to y intersects an active *box* of not too low weight (i.e. just satisfying the *weight* condition of property 1 in Definition 3.2), one of these vertices must connect to a vertex in the intersected box, since it is active. This follows verbatim by Lemma 2.11.

▶ **Definition 3.4** (High-weight box). *Fix a tower level i and an $\varepsilon > 0$ as in Definition 3.2. Then, let W be as in Definition 3.2. We say that a box of the tessellation is* high-weight *if the lower limit of its weight range is at least W^ε.*

With those boxes out of the way, the only other option left for the $u-v$ path to exit $W(u,v) \cup S(u,v)$ is through the low-weight boxes of the towers. We will now show that this is also impossible without connecting to the largest component $\mathcal{C}$ of some G_F. We first deal with the case where an edge possibly starts inside $W(u,v)$ and ends up somewhere outside of $W(u,v) \cup S(u,v)$. We show that this just cannot happen without intersecting a high-weight box of $S(u,v)$.

▷ Claim 3.5. There is no edge e from a vertex x to y with x in $W(u,v)$ and y not in $W(u,v) \cup S(u,v)$ such that this edge does not intersect any high-weight box of $S(u,v)$. Moreover, the intersected box is guaranteed to be part of the visible boundary of the hole in which B_y belongs.

Proof. Assume without loss of generality that $w_x \le w_y$. Let W be as in Definition 3.2. Consider the box-shadow $\sigma(e) = B_1, B_2, \ldots, B_k$. This is a $\mathcal{B}^+$-walk from B_x to B_y. Let $i < j-1$ be a maximal pair such that $B_i \in W(u,v)$, $B_j \notin W(u,v) \cup S(u,v)$ and $B_l \in S(u,v)$ for all $i < l < j$. We must have $d_{\mathcal{B}^+}(B_i, B_j) > 1$, and so any two vertices in these boxes must be separated in d-dimensions by at least a distance $\Omega(W^{2/d})$. The part of the line segment travelling between two such vertices and through the boxes B_l stays below weight W^ε until it reaches B_j, otherwise it would intersect an active box. Therefore, this implies that the weight/distance slope of the edge's line segment is at most $\rho = C'(W^\varepsilon - w_x)/W^{2/d}$. Now, we may write $w_y = w_x + \rho|x-y|$. We know that $|x-y|^d \le w_x w_y = w_x(w_x + \rho|x-y|)$. Since $w_x\rho < 1$ for small enough ε and/or large enough W and we also have $d \ge 1$, the validity of this inequality is monotone wrt $|x-y|$, in the sense that larger values for the latter can only falsify the inequality. We know that $|x-y|$ is $\Omega\big(W^{2/d}\big)$, and so we must have for some constant $C > 0$ (depending on D_0 and importantly not on W):

$$C^{2/d}W^2 \le w_x(w_x + CW^{2/d}(W^\varepsilon - w_x)/W^{2/d}) \le w_x(w_x(1-C) + CW^\varepsilon)$$

Now, since $w_x \le W^\varepsilon$, we have a contradiction for $\varepsilon < 1$ and W large enough. ◁

The final piece of the puzzle is to handle cases where x is in $W(u,v)$ and y is in a low-weight box of $S(u,v)$. We now need to switch viewpoint and instead of actually looking at an edge between x and y we identify u_{in} as the last vertex to be in $W(u,v)$ and u_{out} to be the first vertex after u_{in} (on the path from u to v) which is not in $S(u,v)$ (similar to the definition of B_i and B_j above). Between u_{in} and u_{out}, all vertices are in $S(u,v)$. Such an arrangement must occur if the path from u to v is to ever exit $W(u,v) \cup S(u,v)$ (and in the case we are currently considering, at least one vertex between u_{in} and u_{out} must exist).

▷ Claim 3.6. It is impossible for a path to exist between a vertex $u_{in} \in W(u,v)$ to a vertex $u_{out} \notin W(u,v) \cup S(u,v)$ without some vertex in said path connecting (within its own tower) to a vertex in a high-weight box of $S(u,v)$. Moreover, the mentioned box (more accurately, the tower containing it) is guaranteed to be part of the visible boundary of the hole in which $B_{u_{out}}$ belongs.

Proof. Let $u_1, u_2, \dots, u_k$ be the vertices between u_{in} and u_{out} in the path, all of which are wlog assumed to be in $S(u,v)$. We know that $|u_{in} - u_{out}| \geq CW^{2/d}$, where W is as in Definition 3.2, by similar arguments as in the proof of Claim 3.5. Suppose first that $|u_{in} - u_1| \geq \frac{1}{3}CW^{2/d}$. In this case, notice that there must exist a vertex z in a high-weight box within distance $CW^{2/d}$ from u_1 with $w_z \geq 4^d w_{u_1}$, because u_1 is in a low-weight box of an active tower T_a. By triangle inequality, $|u_{in} - z|^d \leq 4^d |u_{in} - u_1|^d$. Therefore, z must connect to u_{in}. So, for the rest of the proof we may assume $|u_{in} - u_1| \leq \frac{1}{3}CW^{2/d}$.

Let G_F be as in Definition 3.2, with T_0 there replaced by T_a. Assume wlog that u_1 is not in the same connected component as z (the vertex in the top box of T_a) in G_F. Moreover, let v_{hop} be the first vertex among $u_2, \dots u_k, u_{out}$ which is not a vertex of G_F. One such vertex must exist, since otherwise we would conclude that $|u_{in} - u_{out}| \leq \frac{1}{3}CW^{2/d} + W^{\Theta(\varepsilon)} \leq \frac{2}{3}CW^{2/d}$. Let v_{pre} be the vertex right before v_{hop}. Then, this means that either i) the *geometric* position of v_{hop} is outside of T_a and the neighboring towers (i.e. outside of F as in Definition 3.2) or ii) $w_{v_{hop}} \geq W$ which again implies i), if we (wlog) assume that z does not connect to v_{hop}. In any case, we conclude that $|v_{pre} - v_{hop}| \geq \frac{1}{3}CW^{2/d}$. By a similar argument as in the first paragraph, this implies that z connects to v_{hop}, finishing the proof. ◁

The above claims collectively prove the following, analogous to Lemma 2.17.

▶ **Corollary 3.7.** *Suppose u and v are connected in the GIRG graph. There exists a sequence of boxes/towers $X_0, X_1, \dots, X_k$, entirely contained in $W(u,v) \cup S(u,v)$, with the following property. There exists a constant C such that for every i, there is a vertex u_i in X_i with $u_0 = u$ and $u_k = v$ such that $d_{GIRG}(u_j, u_{j+1}) \leq C$ holds for all $j \in 0, \dots k-1$. In particular, $d_{GIRG}(u,v) \leq \mathcal{O}(|W(u,v)|)$.*

Final steps in bounding the diameter

Just like we did for the case $\tau \leq 3$ with high enough λ, one can show that the inactive regions, where we now fix a large enough i and a small enough ε and consider towers instead of boxes for the low levels, do not grow too large. A crucial difference is that the events E_1 and E_2 meaning that neighboring towers T_1 and T_2 are active are *not independent*. However, in any set of k towers/boxes, we can always identify a constant fraction of them that *are* completely independent, and that fraction is roughly $1/(3^d - 1)$. This only forces us to drive the probability p_{in} in Lemma 2.4 even closer to 0. This argument proves the following, analogous to Lemma 2.4.

▶ **Corollary 3.8.** *Assume $\tau < 3$. Then, asymptotically almost surely we have for some constant C that $|W(X_1, X_2)| \leq C \log n$ for all pairs X_1, X_2 of boxes/towers.*

4 Lower bound

In this section we prove a matching logarithmic lower bound for the diameter when $\tau > 2$. The idea will be to identify $n^{\Omega(1)}$ regions in the ground space and argue that each of them has an independent chance of producing a path component of logarithmic size. The success probability per trial will be $n^{-C'}$ with a constant C' of our choice, and we show that at least one will succeed aas.

In more detail but still at an intuitive level, a central point is that the maximum weight found in the graph will aas be at most roughly $n^{1/(\tau-1)}$, which is $o(n)$ for $\tau > 2$. The paths we will try to construct will consist of vertices with constant weight. Therefore, a single trial requires only about $n^{1/(\tau-1)}$ volume of space. Inside this region, we can locally decide whether there exists any unwanted vertex connecting to our under-construction path. Outside the region, no vertex can connect to the path under the aforementioned maximum weight constraint.

Let us also provide some more details about how to decide existence and local disconnectivity of the path to the rest of the region. First of all, we will look for the path at the middle of the region. Geometrically, it will be a "snaky" path confined in a ball of radius $\Theta(\log^{1/d} n)$ but still occupying a constant fraction of its volume (see Figure 6). Assuming such a path exists, we now try to ensure that no other vertices connect to it. Fix an r and consider the geometric topos with minimum distance to any of the vertices of the path in the interval $[r, 2r]$. Since the maximum distance between any two vertices in the path is $C \log^{1/d} n$, it has volume at most $(r + C \log^{1/d} n)^d$. Now, for a vertex in this region to connect to any vertex in the path, its weight must be at least $\Omega(r^d)$. Going over r in powers of two, it can be seen that the expected number of such vertices is at most $C' \log n$, for a constant C' that can be controlled by reducing C. Since this number is Poisson-distributed, the probability that it is equal to 0 is at least $\exp(-C' \log n) = n^{-C'}$. Assuming we can attempt such a trial $n^{C''}$ times, we simply need to ensure $C' < C''$. For context, C'' will roughly be $1 - 1/(\tau - 1)$, i.e. the number of disjoint volume $n^{1/(\tau-1)}$ cubes we can find.

Now, let us describe a curve through the cubes of a uniform tessellation of a d-dimensional cube, which will serve as a "skeleton" for our path.

▶ **Definition 4.1.** *Let R be a cube-shaped region in d dimensions and tessellate it into M^d identical cubes. Each such cube can be represented as a vector with d entries from $\{0, 1, \ldots, M-1\}$, encoding the geometric position of its center. Order these vectors according to the M-ary Gray code, i.e. every two consecutive vectors differ only in a single entry and by exactly 1. We define the* Gray curve *of level M of R as the piecewise linear curve going through the centers of the cubes in the order given by the M-ary Gray code.*

▶ **Observation 4.2.** *For any two non-consecutive cubes visited by the Gray curve, the following holds. If x_1, x_2 are points on the curve in these cubes respectively, then $|x_1 - x_2| \geq c_d s$, where s is the side length of the cubes and c_d is a constant depending on the dimension d.*

Proof. Notice that each cube contains inside it one or two parts of the curve, both of which are line segments from the center of the cube to the center of some face of the cube, with length $s/2$. It suffices to show that such pairs of line segments do not come within distance smaller than $c_d s$.

If the centers of the cubes are separated by distance larger than s, the claim follows by triangle inequality (the set of possible distances is quantized). Therefore, for the rest of the proof we assume that the centers are exactly s units away from each other, i.e. the cubes share a face.

Notice that there is a limited set of directions a line segment can have (of cardinality $2d$), seen as oriented away from the center of the cube. If the segments are both oriented straight towards the other cube's center, the cubes would be consecutive. Thus, they are not and the claim follows by the law of cosines. ◀

We proceed with a claim that gives a lower bound on the probability of a cube of volume $n^{1/(\tau-1)+\varepsilon}$ to be locally good, using the object of Definition 4.1.

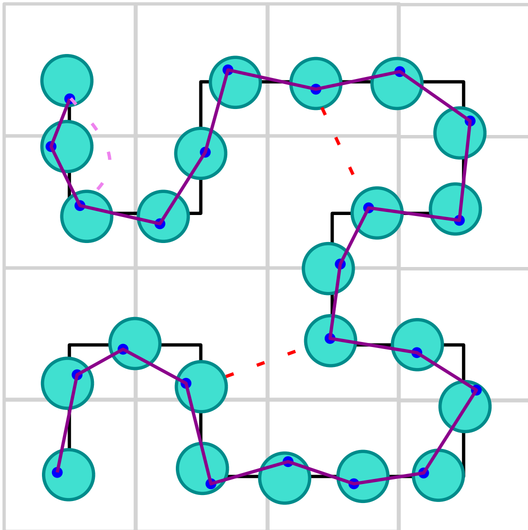

Figure 6 Logarithmic diameter lower bound construction for $d = 2$. Inside a uniform tessellation of a box R' of logarithmic volume, we consider the Gray curve (Definition 4.1) going through the centers of the boxes. Tracing along the curve, we place balls of radius 1 where we look for vertices of restricted weights. By definition, vertices in boxes not consecutive in the Gray curve cannot connect (depicted in dashed red). Other pairs of nearby balls may introduce small shortcuts (dashed pink).

▷ Claim 4.3. Let $C_1, \varepsilon > 0$. Consider a cube-shaped region R of volume $n^{1/(\tau-1)+\varepsilon}$ in the geometric space of the GIRG graph G (with $V(R)$ the set of vertices in R) and let E be the event that $G[V(R)]$ contains a component of diameter $C_1 \log n$ comprised solely of vertices of weight at most 3^d. Then $\mathbb{P}[E] \geq \exp(-C_2 \log n)$, for a C_2 that depends on C_1.

Proof. Wlog translate R such that its center is the origin. Let R' be a smaller similar region of volume $C_1' \log n$ (with C_1' sufficiently large compared to C_1), also centered at the origin. Tessellate R' into identical cubes in the finest way (i.e. largest M as in Definition 4.1) possible such that the following holds. For any two non-consecutive cubes visited by the curve and any two corresponding points on the curve (within the mentioned cubes respectively), said points have distance at least 12 from each other. Notice that by Observation 4.2, M is at least $\Omega(\log^{1/d} n)$, i.e. the cubes shrink down to constant volume.

Now, consider balls of volume 1 (i.e. also having radius 1) placed on the Gray curve whenever the distance to the previous ball's center becomes 2 (we start by placing one ball at the start of the curve). Note that each cube is assigned at least one ball, by Observation 4.2 (or more simply by demanding that the side length of the tessellation is at least 4). With at least the required probability, each such ball contains *exactly* one vertex with weight in $[2^d, 3^d]$. Now, notice that the vertices of consecutive balls connect, since their distance is at most 4 and the product of their weights is at least 4^d. Moreover, there is no edge between vertices in balls inside non-consecutive cubes (in the Gray curve ordering), since the distance between them would be at least $12 - 1 - 1 = 10$, while the product of their weights is at most 9^d. These arguments imply the existence of a *shortest* path P with length at least the number of cubes of the tessellation, which is in turn in the order of the volume of R', i.e. logarithmic.

It now remains to show that with good enough probability, this path is not shortcut when exposing the rest of the graph inside R. To this end, fix some $r = 2^l$ with l a non-negative integer. Consider the geometric topos of points with minimum distance to a vertex of P in the interval $[r, 2r]$. Since all vertices of P are in R', the volume of this topos is $\mathcal{O}(\max\{r^d, \log n\})$. Now, to connect to a vertex of P, a vertex in said topos must have weight $\Omega(r^d)$. The number of vertices in this topos with weight at least this lower bound is Poisson distributed with expectation $\mathcal{O}(r^{d(1-\tau)}(\max\{r^d, \log n\}))$. Summing over all l, we observe that the overall expected number of vertices connecting to P is $\mathcal{O}(\log n)$ and is also Poisson-distributed. Thus, the probability that it is 0 is at least as desired, as the constant in Landau notation depends on C_1. ◁

With the above claim at hand, we may now proceed to prove this section's theorem.

▶ **Theorem 4.4.** *Assume $\tau > 2$. Then, aas the diameter of GIRG is $\Omega(\log n)$.*

Proof. Fix a small enough $\varepsilon > 0$ and observe that aas the maximum weight of the graph will be at most $n^{1/(\tau-1)+\varepsilon}$. Identify $\Theta\big(n^{1-1/(\tau-1)-2\varepsilon}\big)$ disjoint cubes of volume $n^{1/(\tau-1)+2\varepsilon}$.[14] Now, if the event of Claim 4.3 succeeds for any one of these cubes, we have a witness to the claimed lower bound on the diameter, since the discovered component has large diameter within its own cube and does not connect to any vertex outside the cube, by virtue of the upper bound we have on the maximum weight and a lower bound on the distance between a vertex in said component to any vertex outside the cube (since the component is localized in R' from the proof of Claim 4.3). By setting C_2 in Claim 4.3 small enough compared to $1 - 1/(\tau - 1) - 2\varepsilon > 0$, we conclude that the number of successful cubes is binomially distributed with expectation $n^{\Omega(1)}$. This follows by the fact that the cubes are disjoint and therefore the event described in Claim 4.3 is independent across different cubes. Therefore, aas we have not only one but polynomially many components of logarithmic diameter in the graph using a Chernoff bound. ◀

▶ Remark 4.5. The above proof can be modified to also show that *the giant component* has logarithmic diameter. The only thing we need to change is to look for paths in Claim 4.3 that instead of not connecting to any other vertex in R connect to exactly one vertex of polylogarithmic weight v_p. Since all vertices above a certain polylogarithmic weight are part of the giant aas [12, Theorem 5.9], this would show the claim. However, this connection must be done carefully so as not to shortcut the original path P. One way to achieve this is the following. First, we look for a vertex v_1 of weight $\Theta(\sqrt{\log n})$ in an appropriate geometric topos of volume $\Theta(1)$. This topos is such that v_1 would connect to at least 1 vertex in P but also it would have geometric distance at least $\Omega(\log^{1/d} n)$ to some other vertex in P (a small ball close to a "corner" vertex as in Figure 6 would do the trick). The second property implies that there exists a subpath P_s (staring from the vertex distant to v_1) of P with $|P_s| \geq \Omega(|P|)$ and such that no vertices of P_s have an edge to v_1. This follows because moving along the path P one can only geometrically move a constant geometric distance per step along the

[14] The factor 2 is introduced to avoid dealing with some constants in arguing that no connections exist from R' to outside of R from Claim 4.3.

path. Notice that we succeed in finding such a v_1 with probability $(\log n)^{-\Theta(1)}$. We then look for a vertex v_2 of weight $\Theta(\log n)$ at geometric distance $\mathcal{O}((\log n)^{3/(2d)})$ from v_1 and at the same time $\Omega((\log n)^{3/(2d)})$ from any vertex in P (now the geometric topos we use can have much higher volume, but we don't actually need to use this). It is not hard to show that v_2 will not connect to any of the vertices of P. We keep increasing the weight like this until we reach a suitable vertex v_p. This modification does not impact the success probability or the independence between polynomial-volume regions of Claim 4.3 and therefore this shows that the giant component also has at least logarithmic diameter.

5 Conclusion

We have shown that the diameter of T-GIRG is asymptotically almost surely $\Theta(\log n)$. A natural next question is whether the same holds for GIRG with positive temperature. The additional edges in this case are also called *weak ties*, and they have fundamental impact on many processes on GIRG, see for example [27, 29]. Our proof method breaks down in that case, mainly since Lemma 2.11 is no longer true: a weak edge may cross an active box although both endpoints are too far away to reliably connect to any vertices in the box. Nevertheless, we conjecture that the diameter is still $\Theta(\log n)$ in this case, and a proof of this conjecture would be highly interesting. Finally, it would also be interesting to study the case where d is not a constant but grows with n, which our techniques are not equipped for. One would expect the diameter to again be logarithmic, since at least when d grows sufficiently quickly, GIRGs converge to a different random graph model (Chung-Lu random graphs), which itself has logarithmic diameter.

References

1 Samuel Baguley, Yannic Maus, Janosch Ruff, and George Skretas. Hyperbolic random graphs: Clique number and degeneracy with implications for colouring. In *42nd International Symposium on Theoretical Aspects of Computer Science*, 2025.

2 Zylan Benjert, Kostas Lakis, Johannes Lengler, and Raghu Raman Ravi. The diameter of (threshold) geometric inhomogeneous random graphs. *arXiv*, 2025. `doi:10.48550/arXiv.2510.12543`.

3 Nicholas H Bingham, Charles M Goldie, and Jef L Teugels. *Regular variation*, volume 27. Cambridge university press, 1989.

4 Thomas Bläsius and Philipp Fischbeck. On the external validity of average-case analyses of graph algorithms. *ACM Transactions on Algorithms*, 20(1):1–42, 2024. `doi:10.1145/3633778`.

5 Thomas Bläsius, Philipp Fischbeck, Tobias Friedrich, and Maximilian Katzmann. Solving vertex cover in polynomial time on hyperbolic random graphs. *Theory of Computing Systems*, 67(1):28–51, 2023. `doi:10.1007/S00224-021-10062-9`.

6 Thomas Bläsius, Cedric Freiberger, Tobias Friedrich, Maximilian Katzmann, Felix Montenegro-Retana, and Marianne Thieffry. Efficient shortest paths in scale-free networks with underlying hyperbolic geometry. *ACM Transactions on Algorithms (TALG)*, 18(2):1–32, 2022. `doi:10.1145/3516483`.

7 Thomas Bläsius, Tobias Friedrich, and Anton Krohmer. Cliques in hyperbolic random graphs. *Algorithmica*, 80(8):2324–2344, 2018. `doi:10.1007/S00453-017-0323-3`.

8 Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10):P10008, 2008.

9 Marian Boguna, Dmitri Krioukov, and Kimberly C Claffy. Navigability of complex networks. *Nature Physics*, 5(1):74–80, 2009.

10 Béla Bollobás. *The art of mathematics: Coffee time in Memphis.* Cambridge University Press, 2006.

11 Karl Bringmann, Ralph Keusch, and Johannes Lengler. Geometric inhomogeneous random graphs. *Theoretical Computer Science*, 760:35–54, 2019. doi:10.1016/J.TCS.2018.08.014.

12 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.

13 Karl Bringmann, Ralph Keusch, Johannes Lengler, Yannic Maus, and Anisur Rahaman Molla. Greedy routing and the algorithmic small-world phenomenon. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, pages 371–380, 2017. doi:10.1145/3087801.3087829.

14 Sacha Cerf, Benjamin Dayan, Umberto De Ambroggio, Marc Kaufmann, Johannes Lengler, and Ulysse Schaller. Balanced bidirectional breadth-first search on scale-free networks. *arXiv preprint*, 2024.

15 Aaron Clauset, Cristopher Moore, and Mark EJ Newman. Hierarchical structure and the prediction of missing links in networks. *Nature*, 453(7191):98–101, 2008.

16 Pilu Crescenzi, Roberto Grossi, Michel Habib, Leonardo Lanzi, and Andrea Marino. On computing the diameter of real-world undirected graphs. *Theoretical Computer Science*, 514:84–95, 2013. doi:10.1016/J.TCS.2012.09.018.

17 Benjamin Dayan, Marc Kaufmann, and Ulysse Schaller. Expressivity of geometric inhomogeneous random graphs—metric and non-metric. In *International Conference on Complex Networks*, pages 85–100. Springer, 2024.

18 Maria Deijfen, Remco Van der Hofstad, and Gerard Hooghiemstra. Scale-free percolation. In *Annales de l'IHP Probabilités et statistiques*, volume 49, pages 817–838, 2013.

19 Reinhard Diestel. *Graph Theory.* Springer Publishing Company, Incorporated, 5th edition, 2017.

20 Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the internet topology. *ACM SIGCOMM computer communication review*, 29(4):251–262, 1999. doi:10.1145/316188.316229.

21 Tobias Friedrich, Andreas Göbel, Maximilian Katzmann, and Leon Schiller. Real-world networks are low-dimensional: theoretical and practical assessment. *arXiv preprint*, 2023.

22 Tobias Friedrich and Anton Krohmer. On the diameter of hyperbolic random graphs. *SIAM Journal on Discrete Mathematics*, 32(2):1314–1334, 2018. doi:10.1137/17M1123961.

23 Guillermo García-Pérez, Antoine Allard, M Ángeles Serrano, and Marián Boguñá. Mercator: uncovering faithful hyperbolic embeddings of complex networks. *New Journal of Physics*, 21(12):123033, 2019.

24 Mohsen Ghaffari and Calvin Newport. Leader election in unreliable radio networks. In *43rd International Colloquium on Automata, Languages, and Programming*, volume 55, page 138. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.ICALP.2016.138.

25 Peter Gracar, Markus Heydenreich, Christian Mönch, and Peter Mörters. Recurrence versus transience for weight-dependent random connection models. *Electronic Journal of Probability*, 27:1–31, 2022.

26 Paul W Holland and Samuel Leinhardt. An exponential family of probability distributions for directed graphs. *Journal of the american Statistical association*, 76(373):33–50, 1981.

27 Marc Kaufmann, Kostas Lakis, Johannes Lengler, Raghu Raman Ravi, Ulysse Schaller, and Konstantin Sturm. Rumour spreading depends on the latent geometry and degree distribution in social network models. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2026. To appear. arXiv:https://arxiv.org/abs/2408.01268.

28 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.

29 Júlia Komjáthy, John Lapinskas, Johannes Lengler, and Ulysse Schaller. Four universal growth regimes in degree-dependent first passage percolation on spatial random graphs i. *arXiv preprint*, 2023.

30 Júlia Komjáthy and Bas Lodewijks. Explosion in weighted hyperbolic random graphs and geometric inhomogeneous random graphs. *Stochastic Processes and their Applications*, 130(3):1309–1367, 2020.

31 Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguná. Hyperbolic geometry of complex networks. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 82(3):036106, 2010.

32 Jure Leskovec, Deepayan Chakrabarti, Jon Kleinberg, Christos Faloutsos, and Zoubin Ghahramani. Kronecker graphs: an approach to modeling networks. *Journal of Machine Learning Research*, 11(2), 2010. `doi:10.5555/1756006.1756039`.

33 Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. Graphs over time: densification laws, shrinking diameters and possible explanations. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, pages 177–187, 2005. `doi:10.1145/1081870.1081893`.

34 Stanley Milgram et al. The small world problem. *Psychology today*, 2(1):60–67, 1967.

35 Tobias Müller and Merlijn Staps. The diameter of kpkvb random graphs. *Advances in Applied Probability*, 51(2):358–377, 2019.

36 David Peleg. *Distributed computing: a locality-sensitive approach*. SIAM, 2000.

37 Chaoming Song, Shlomo Havlin, and Hernan A Makse. Self-similarity of complex networks. *Nature*, 433(7024):392–395, 2005.

38 Ádám Timár. Boundary-connectivity via graph theory. *Proceedings of the American Mathematical Society*, 141(2):475–480, 2013.

39 Ivan Voitalov, Pim Van Der Hoorn, Remco Van Der Hofstad, and Dmitri Krioukov. Scale-free networks well done. *Physical Review Research*, 1(3):033034, 2019.

40 Duncan J Watts and Steven H Strogatz. Collective dynamics of 'small-world'networks. *nature*, 393(6684):440–442, 1998.