From heavy-tailed Boolean models to scale-free Gilbert graphs

Christian Hirsch

submitted: November 26, 2014

Weierstrass Institute
Mohrenstr. 39
10117 Berlin
Germany
E-Mail: christian.hirsch@wias-berlin.de

Abstract

Define the scale-free Gilbert graph based on a Boolean model with heavy-tailed radius distribution on the d-dimensional torus by connecting two centers of balls by an edge if at least one of the balls contains the center of the other. We investigate two asymptotic properties of this graph as the size of the torus tends to infinity. First, we determine the tail index associated with the asymptotic distribution of the sum of all power-weighted incoming and outgoing edge lengths at a randomly chosen vertex. Second, we study the behavior of chemical distances on scale-free Gilbert graphs and show the existence of different regimes depending on the tail index of the radius distribution. Despite some similarities to long-range percolation and ultra-small scale-free geometric networks, scale-free Gilbert graphs are actually more closely related to fractal percolation and this connection gives rise to different scaling limits. We also propose a modification of the graph, where the total number of edges can be reduced substantially at the cost of introducing a logarithmic factor in the chemical distances.

1   Introduction

The spatial distribution of the population in a country is typically far from homogeneous, but rather exhibits fractal patterns. Specifically, this has been investigated for Great Britain and the United States [2, 3] and for Finland [17]. It is pointed out in [2, 17] that fractality has important implications for the design of wired telecommunication networks in the sense that the number and extent of various levels of hierarchy should be adapted to the fractal geometry. A trade-off is involved in determining the optimal number of levels. Using few levels has the advantage that most access points can be connected by a small number of hops. However, this comes at the cost of having to install a cable network of large total length. Indeed, a high-level node has to be connected to a large number of low-level nodes in order to guarantee connectivity. Vice versa, using a large number of hierarchies, one may be able to reduce the length of cables substantially, but this comes at the cost of increasing the number of hops it takes a low-level node to reach the topmost layer of hierarchy. A related cost analysis for hierarchical spatial networks is provided by Baccelli and Zuyev [4]. However, the network structure investigated in that paper is obtained by iteratively considering Voronoi tessellations and does not exhibit fractal geometries (see however [20] for some results in this direction).

In order to develop a fundamental understanding of the asymptotic behavior of cable lengths and chemical distances (i.e., minimal number of hops needed to connect two points) in large networks, it is important to abstract from the types of specific countries and deterministic fractals investigated in [2, 17] and move to random networks. The spatial nature of the problem calls for models based on random geometric graphs, and the fractal geometry suggests that one should look for scale-free random graphs exhibiting a power-law degree distribution. The combination of these two constraints restricts the list of appropriate choices substantially. We briefly review some of the most well-known models in literature and discuss their drawbacks with respect to modeling the kind of networks we have in mind.

One option could be to use long-range percolation, see [5, 6, 7] and the references therein. Here, one starts from the lattice $\mathbb{Z}^d$ and connects any pair of sites independently with a probability depending

only on their distance. This leads to a network with a giant connected component and power-law degree distributions. However, this model does not offer an inherently defined hierarchy of nodes. Moreover, if we think of a high-level node as having the purpose of providing access to all low-level nodes in some region, then this would suggest that the occurrence of edges should be spatially positively correlated. However, if $x, y \in \mathbb{Z}^d$ are arbitrary sites in $\mathbb{Z}^d$ and y' is close to y, then putting an edge between x and y does not influence at all the probability of seeing an edge between x and y'.

A second option could be to consider the ultra-small scale-free geometric networks on $\mathbb{Z}^d$ introduced in [21]. Two sites $x, y \in \mathbb{Z}^d$ are connected by an edge in this graph if $|x - y| \leq \min\{R_x, R_y\}$, where $\{R_z\}_{z \in \mathbb{Z}^d}$ denotes a family of iid heavy-tailed random variables. For any site $z \in \mathbb{Z}^d$ the value R_z can be thought of as the radius of influence of z, so that ultra-small scale-free geometric networks offer a natural possibility for defining the network hierarchy. Similar remarks apply to scale-free percolation [9]. Recently, also a Poisson-based continuum analogue of this model has been investigated, see [10, 11]. From a modeling point of view, this means that there is an excellent degree of connectivity between high-level nodes. However, conversely, for low-level nodes it may be very difficult to get connected to a nearby high-level node. This might pose a substantial obstruction to the build-up of a hierarchical network. Due to the drawbacks of the existing networks models, we propose two alternatives.

According to our discussion, it would be desirable to consider a variant of the ultra-small scale-free geometric network, where two sites $x, y \in \mathbb{Z}^d$ are connected by an edge if $|x - y| \leq \max\{R_x, R_y\}$ (instead of $|x - y| \leq \min\{R_x, R_y\}$). In other words, in order to connect x and y by an edge it is no longer necessary that both x lies within the radius of influence of y *and* y lies within the radius of influence of x. It suffices that x lies within the radius of influence of y *or* y lies within the radius of influence of x. We consider a spatial variant of this network model, called *scale-free Gilbert graph*, where the vertices are given by a homogeneous Poisson point process on the d-dimensional torus with side length n, for some $d, n \geq 2$. See Figure 1a for an illustration of this graph in dimension $d = 2$.

(a) Scale-free Gilbert graph   (b) Thinned scale-free Gilbert graph
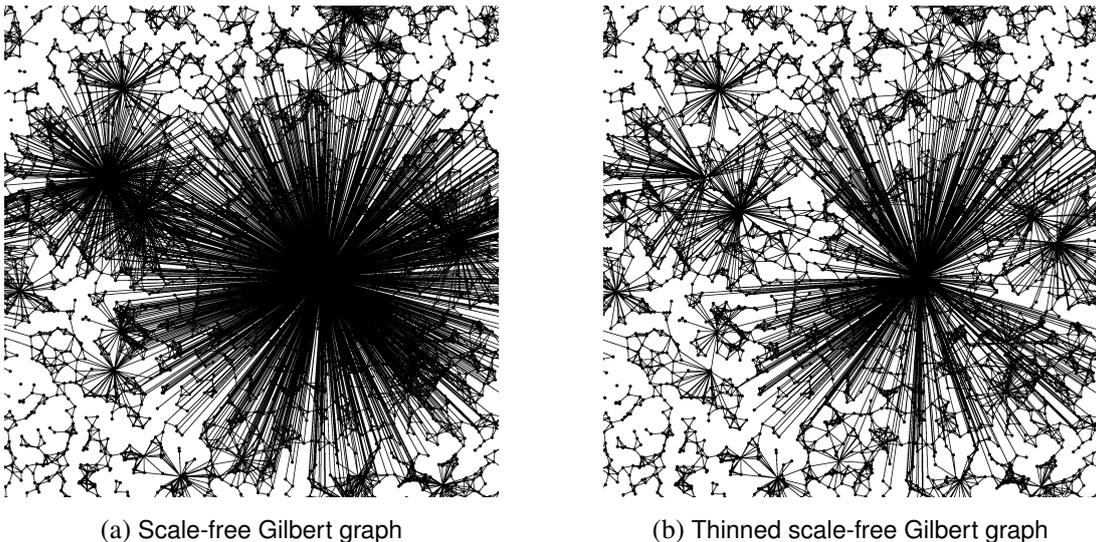
Figure 1: Planar scale-free and thinned scale-free Gilbert graphs

As $n \to \infty$, we investigate the asymptotic distribution of the power-weighted sum of all incoming and all outgoing edge lengths considered from a vertex that is picked uniformly at random, see Theorems 1 and 2. In particular, our results imply that the asymptotic out-degree and in-degree distributions admit

polynomial tails. We also investigate the growth of the expected power-weighted sum of all outgoing edge lengths as $n \to \infty$, see Theorem 3. In Section 4, we show that different scaling regimes of chemical distances emerge depending on the tail index of the radius distribution.

Since the scale-free Gilbert graph exhibits a large degree of redundancy of connections, in Section 5 we present a variant, called *thinned scale-free Gilbert graph*, that substantially decreases these redundancies by introducing a multi-layer topology. Loosely speaking, an edge in the original graph from a low-level node x to a high-level node y is removed if y can also be reached from x via an intermediate node z. See Figure 1b for an illustration of the thinned scale-free Gilbert graph. A precise definition of this graph will be given in Section 2, but we note already at this point that removing redundancies does not alter the family of connected components. Furthermore, under suitable assumptions on the tail behavior of the radius distribution, we show in Theorem 7 that the sum of power-weighted edge lengths at a randomly picked vertex decreases significantly. Of course, reducing edge lengths comes at the price of longer connection paths, but we will show in Theorem 8 that chemical distances grow at most by a logarithmic factor in the size of the torus.

The present paper is organized as follows. In Section 2, we provide precise definitions of the scale-free and thinned scale-free Gilbert graphs and state our main results. Next, in Section 3, we investigate power-weighted sums of edge lengths at a typical point and prove Theorems 1, 2 and 3. Section 4 is devoted to the analysis of chemical distances for different regimes of the tail index. Finally, in Section 5, we investigate the changes in terms of the asymptotic behavior of power-weighted sums of edge lengths and of chemical distances as we move from the scale-free Gilbert graph to its thinning.

2   Model definition and statement of main results

In Sections 2.1 and 2.2, we provide precise definitions of the network models under consideration and present our main results. In the following, $d \geq 2$ is always assumed to be an arbitrary fixed integer. Moreover, $B_r(\xi) = \{\eta \in \mathbb{R}^d : |\eta - \xi| \leq r\}$ denotes the ball of radius $r > 0$ centered at $\xi \in \mathbb{R}^d$.

2.1   Scale-free Gilbert graph

Let $X^{(n)}$ be an independently $[0, \infty)$-marked homogeneous Poisson point process with intensity 1 in the torus $\mathbb{T}_n$, where $\mathbb{T}_n$ is obtained from the cube $[-n/2, n/2]^d$ by the standard identification of its boundaries. If $n \geq 1$, $r > 0$ and $\xi \in \mathbb{T}_n$, then we write $B_r^{\mathbb{T}_n}(\xi) = \{\eta \in \mathbb{T}_n : \mathrm{d}_{\mathbb{T}_n}(\xi, \eta) \leq r\}$ for the closed ball in $\mathbb{T}_n$ with radius r centered at ξ, where $\mathrm{d}_{\mathbb{T}_n}(\xi, \eta)$ denotes the toroidal distance between ξ and η. The mark of a point from $X^{(n)}$ is interpreted as the radius of a ball centered at this point. Throughout the paper, we assume that the distribution of the typical mark R is absolutely continuous and heavy-tailed. That is, there exist $\beta, s \in (0, \infty)$ such that $\lim_{h \to \infty} h^s \mathbb{P}(R > h) = \beta$. For every $\varepsilon \in (0, 1)$, we also fix $t_0 = t_0(\varepsilon) > 0$ such that $t^s \mathbb{P}(R > t) \in (\beta(1 - \varepsilon), \beta(1 + \varepsilon))$ for all $t \geq t_0(\varepsilon)$.

For $n \geq 1$ we investigate the directed random geometric graph $G(X^{(n)})$ on the vertex set X, where an edge in $G(X^{(n)})$ is drawn from $x = (\xi, r) \in X^{(n)}$ to $y = (\eta, t) \in X^{(n)}$ if $\eta \in B_r^{\mathbb{T}_n}(\xi)$.

First, we investigate the asymptotic distributions of the power-weighted sum of all outgoing and all incoming edge lengths. To be more precise, fix $\alpha \geq 0$, $n \geq 1$ and let R^* be a copy of R that is

independent of $X^{(n)}$. Then,

$$D_{\text{out},n}^{(\alpha)} = \sum_{(\xi,r)\in X^{(n)}} |\xi|^\alpha 1_{B_{R^*}^{\mathbb{T}_n}(o)}(\xi),$$

denotes the sum of αth powers of the lengths of all outgoing edges at the node (o, R^*), where we put $|\xi| = \mathrm{d}_{\mathbb{T}_n}(\xi, o)$. Considering the limit $n \to \infty$, i.e., letting the size of the torus tend to infinity, it is intuitive (and will be shown rigorously in Theorem 1 below) that the random variables $\big\{ D_{\text{out},n}^{(\alpha)} \big\}_{n \geq 1}$ converge in distribution to the random variable

$$D_{\text{out}}^{(\alpha)} = \sum_{(\xi,r)\in X} |\xi|^\alpha 1_{B_{R^*}(o)}(\xi),$$

where X denotes an independently $[0,\infty)$-marked homogeneous Poisson point process in $\mathbb{R}^d$ with intensity 1. In addition to showing the convergence of $D_{\text{out},n}^{(\alpha)}$ to $D_{\text{out}}^{(\alpha)}$, we also investigate the behavior of the tail probabilities $p_{\text{out},t}^{(\alpha)} = \mathbb{P}\big(D_{\text{out}}^{(\alpha)} > t\big)$ as $t \to \infty$. In the following, κ_d denotes the volume of the unit ball in $\mathbb{R}^d$.

Theorem 1. *For every $\alpha > 0$ the random variables $\big\{ D_{\text{out},n}^{(\alpha)} \big\}_{n \geq 1}$ converge to the random variable $D_{\text{out}}^{(\alpha)}$ in distribution. Moreover,*

$$\lim_{t\to\infty} t^{s/(\alpha+d)} p_{\text{out},t}^{(\alpha)} = (d\kappa_d/(\alpha+d))^{s/(\alpha+d)}\beta.$$

In particular, the degree distribution of out-degrees admits asymptotically polynomial tails of order s/d. In contrast, when considering in-degrees an entirely different asymptotic behavior emerges. Indeed, in Theorem 2 below we show that for $s \leq d$ the degree distribution is asymptotically degenerate, whereas for $s > d$ it is Poissonian (with finite mean). In order to make this precise, let

$$D_{\text{in},n}^{(\alpha)} = \sum_{(\xi,r)\in X^{(n)}} |\xi|^\alpha 1_{B_r^{\mathbb{T}_n}(o)}(\xi),$$

denote the sum of αth powers of the lengths of all incoming edges at the node (o, R^*), where $\alpha \geq 0$, $n \geq 1$. As in the case of $D_{\text{out},n}^{(\alpha)}$, we will see that the random variables $D_{\text{in},n}^{(\alpha)}$ converge in distribution to a random variable $D_{\text{in}}^{(\alpha)}$ given by

$$D_{\text{in}}^{(\alpha)} = \sum_{(\xi,r)\in X} |\xi|^\alpha 1_{B_r(o)}(\xi),$$

where as before X denotes an independently $[0,\infty)$-marked homogeneous Poisson point process in $\mathbb{R}^d$ with intensity 1. The following result is devoted to the tail probabilities $p_{\text{in},t}^{(\alpha)} = \mathbb{P}\big(D_{\text{in}}^{(\alpha)} > t\big)$

Theorem 2. *For every $\alpha \geq 0$ the random variables $\big\{ D_{\text{in},n}^{(\alpha)} \big\}_{n \geq 1}$ converge to the random variable $D_{\text{in}}^{(\alpha)}$ in distribution. Moreover, $\mathbb{P}\big(D_{\text{in}}^{(\alpha)} = \infty\big) = 1$ for $s \leq d$, whereas if $s > d$, then $D_{\text{in}}^{(0)}$ is a Poissonian random variable with mean $\kappa_d \mathbb{E} R^d$. Finally, if $\alpha > 0$ and $s > d$, then*

$$\lim_{t\to\infty} t^{(s-d)/\alpha} p_{\text{in},t}^{(\alpha)} = d\kappa_d \beta(s-d)^{-1}.$$

Besides considering limit distributions, we also determine leading-order asymptotics for the expectations $\big\{ \mathbb{E}D_{\text{in},n}^{(\alpha)} \big\}_{n \geq 1}$ as $n \to \infty$. Note that here it is not necessary to distinguish between ingoing and outgoing edges, since $\mathbb{E}D_{\text{out},n}^{(\alpha)} = \mathbb{E}D_{\text{in},n}^{(\alpha)}$ for all $n \geq 1$.

Theorem 3. *Let $\alpha > 0$ be arbitrary. Then one can distinguish between three limiting regimes for the random variable $\mathbb{E}D^{(\alpha)}_{\text{in},n}$ depending on the sign of $\alpha + d - s$.*

   (i) *If $s > \alpha + d$, then $\lim_{n\to\infty} \mathbb{E}D^{(\alpha)}_{\text{in},n} = d\kappa_d(\alpha + d)^{-1}\mathbb{E}R^{\alpha+d}$.*

   (ii) *If $s = \alpha + d$, then $\lim_{n\to\infty}(\log n)^{-1}\mathbb{E}D^{(\alpha)}_{\text{in},n} = d\kappa_d\beta$.*

   (iii) *If $s < \alpha + d$, then $\lim_{n\to\infty} n^{s-\alpha-d}\mathbb{E}D^{(\alpha)}_{\text{in},n} = \beta \int_{[-1/2,1/2]^d} |\eta|^{\alpha-s}\mathrm{d}\eta$.*

Next, we investigate the asymptotic behavior of chemical distances in scale-free Gilbert graphs, where we allow a directed edge in $G(X^{(n)})$ to be traversed in both directions. For the analysis of $G(X^{(n)})$, it will be convenient to distinguish between the regimes $s < d$, $s = d$ and $s > d$. We say that a family of events occurs *with high probability (whp)* if the probabilities of the events tend to 1 as $n \to \infty$. In the case where $s < d$, whp there exists some point of $X^{(n)}$ that is connected to all other points by an edge in $G(X^{(n)})$. In particular, the diameter of $G(X^{(n)})$ is at most 2 whp. The cases $s = d$ and $s > d$ are more subtle. If $s > d$, then as $n \to \infty$ the probability that $G(X^{(n)})$ contains isolated vertices is bounded away from 0, see Proposition 10. Still, if β is sufficiently large, then it follows from continuum percolation that whp there exists a giant component containing a positive proportion of the vertices. But even inside this giant connected component the effects of the heavy-tailed nature of R are barely noticeable. This is to be understood in the sense that chemical distances (i.e., minimal number of hops) between nodes at distance n grow almost linearly in n. To be more precise, putting $\mathrm{e}_1 = (1, 0, \ldots, 0) \in \mathbb{T}_n$ and denoting by $q(x)$ the closest point of X seen from a given point $x \in \mathbb{T}_n$, we have the following result.

Theorem 4. *Assume that $s > d$ and let $\alpha > 0$ be arbitrary. Then, the chemical distance between $q(-n\mathrm{e}_1/4)$ and $q(n\mathrm{e}_1/4)$ is at least $n/(\log n)^\alpha$ whp.*

We conjecture that the sublogarithmic correction factor is only an artifact of our proof.

Conjecture 5. *Assume that $s > d$. Then, there exists a constant $c = c(\beta, d) > 0$ such that the chemical distance between $q(-n\mathrm{e}_1/4)$ and $q(n\mathrm{e}_1/4)$ is at least cn whp.*

Finally, we assume that $s = d$. On the one hand, in contrast to the case $s > d$ the heavy-tailedness of the radius distribution has a substantial effect on chemical distances. On the other hand, the connectivity structure still does not degenerate as in the case $s < d$. Therefore, from the point of view of modeling scale-free network structures, the regime $s = d$ might be considered to be the most interesting one. In Theorem 6 below, we show that if the parameter β is sufficiently large, then not only the graph $G(X^{(n)})$ is connected whp, but, moreover, its diameter remains bounded in the sense that whp it is dominated by the total progeny of a subcritical Galton-Watson process. This asymptotic behavior differs decisively from the one of either long-range percolation [5, 6] or ultra-small scale-free geometric networks [21]. As we will see in the proof of Theorem 6, this scaling regime is a consequence of the close relationship between scale-free Gilbert graphs and fractal percolation processes.

Theorem 6. *Assume that $\beta > d^{d/2}2^{2d+1}(d + 1)\log 2$. Then, the graph $G(X^{(n)})$ is connected whp, and, moreover, the diameter $\mathrm{diam}\, G(X^{(n)})$ of $G(X^{(n)})$ is stochastically dominated by an affine function in the total progeny of a subcritical Galton-Watson process whp. To be more precise, there exists a coupling between $X^{(n)}$ and the total progeny T of a subcritical Galton-Watson process such that $\mathbb{P}(\mathrm{diam}\, G(X^{(n)}) > 2 + 2^dT) \in O(n^{-1})$.*

2.2 Thinned scale-free Gilbert graph

As explained in Section 1, the scale-free Gilbert graph $G(X^{(n)})$ contains a large number of redundant edges, which can be removed without affecting its connected components. To be more precise, if $x_1 = (\xi_1, r_1), x_2 = (\xi_2, r_2), x_3 = (\xi_3, r_3) \in X^{(n)}$ are such that $r_1 > r_2 > r_3$, $\xi_2 \in B_{r_1}^{\mathbb{T}_n}(\xi_1)$ and $\xi_3 \in B_{r_1}^{\mathbb{T}_n}(\xi_1) \cap B_{r_2}^{\mathbb{T}_n}(\xi_2)$, then in $G(X^{(n)})$ the point x_3 is connected both to x_1 and x_2. However, the edge from x_3 to x_1 is redundant since one can also reach x_1 from x_3 by first moving from x_3 to x_2 and then from x_2 to x_1. See Figure 2 for an illustration of this configuration. In order to reduce the total network length, we therefore introduce a variant G' of the original graph G, where such redundancies are removed.
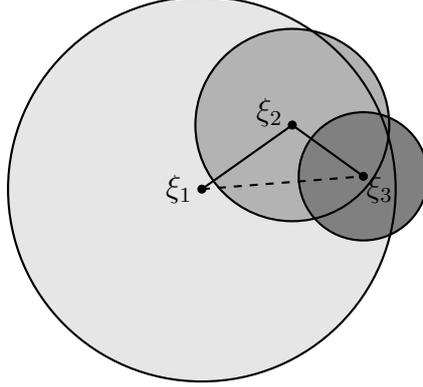

Figure 2: The dashed edge connecting x_1 and x_3 is redundant

To be more precise, for any finite subset φ of $\mathbb{T}_n \times [0, \infty)$ define the *thinned scale-free Gilbert graph* $G'(\varphi)$ as the graph on φ, where an edge is drawn from $(\xi, r) \in \varphi$ to $(\eta, t) \in \varphi$ if $(\eta, t) \in B_r^{\mathbb{T}_n}(\xi) \times [0, r)$ and there does not exist $(\zeta, w) \in \varphi \cap (B_r^{\mathbb{T}_n}(\xi) \times (t, r))$ such that $\eta \in B_w^{\mathbb{T}_n}(\zeta)$.

By definition, $G'(X^{(n)})$ is a subgraph of $G(X^{(n)})$ and we will see in Proposition 17 that the thinning does not affect the connected components of $G(X^{(n)})$. From the point of view of telecommunication networks, we can reach the same set of subscribers using a smaller cable length. Next, we investigate the question whether the reduction of cable length is substantial. For concreteness, we assume that $s = d$. We show that the leading order of the expected power-weighted sum of lengths of outgoing edges in $G'(\{(o, R^*)\} \cup X^{(n)})$ is strictly smaller when compared to the graph $G(X^{(n)})$, see Theorem 3. To be more precise, defining

$$D'^{,(\alpha)}_{\mathsf{out},n} = \sum_{(\xi, r) \in X^{(n)}} |\xi|^\alpha \mathbb{1}_{((o, R^*), (\xi, r)) \text{ is an edge in } G'(\{(o, R^*)\} \cup X^{(n)})},$$

where $\alpha \geq 0$ and $n \geq 1$ we have the following result.

Theorem 7. *Assume that $s = d$. Then, the expected out-degree $\mathbb{E}D'^{,(0)}_{\mathsf{out},n}$ is asymptotically bounded from above by a constant, i.e., $\mathbb{E}D'^{,(0)}_{\mathsf{out},n} \in O(1)$. Moreover, if $\alpha > 0$, then $\mathbb{E}D'^{,(\alpha)}_{\mathsf{out},n} \in O(n^{\alpha - \delta})$ for some $\delta > 0$.*

In Section 5, we show that removing redundancies does not destroy the property of connectivity. However, this thinning operation *does* influence the quality of connectivity, in the sense that chemical distances will increase. Indeed, instead of moving from $x \in X^{(n)}$ to $y \in X^{(n)}$ directly along an edge in the graph $G(X^{(n)})$, introducing a multilayer topology via $G'(X^{(n)})$ might force us to move through a potentially large number of layers, before we can get from x to y. Still, chemical distances increase at most by a logarithmic factor in the size of the sampling window.

Theorem 8. *There exists $c_1 > 0$ such that whp any $x, y \in X^{(n)}$ that are adjacent in $G(X^{(n)})$ can be connected by a path in $G'(X^{(n)})$ consisting of at most $c_1 \log n$ hops.*

3 Proofs of Theorems 1, 2 and 3

In the present section, we investigate the asymptotic distributions of the sum of all outgoing and of the sum of all incoming power-weighted edge lengths at a typical vertex. First, in Theorem 1, we consider the case of outgoing edges. The proof is essentially based on the observation that conditioned on R^* the random variable $D_{\mathsf{out}}^{(\alpha)}$ concentrates around its conditional mean.

Proof of Theorem 1. Using the canonical coupling between the Poisson point process $X^{(n)}$ on the torus $\mathbb{T}_n$ and the Poisson point process X on $\mathbb{R}^d$ we deduce that

$$\mathbb{P}\big(D_{\mathsf{out},n}^{(\alpha)} \neq D_{\mathsf{out}}^{(\alpha)}\big) \leq \mathbb{P}(R^* > n/2),$$

and the latter probability tends to 0 as $n \to \infty$. In order to determine the tail behavior of the random variable $D_{\mathsf{out}}^{(\alpha)}$, for any $r > 0$ we put

$$D_r = \sum_{(\xi,t) \in X} |\xi|^\alpha 1_{B_r(o)}(\xi),$$

noting that $D_{R^*} = D_{\mathsf{out}}^{(\alpha)}$. In particular,

$$t^{s/(\alpha+d)} p_{\mathsf{out},t}^{(\alpha)} = \int_0^\infty t^{s/(\alpha+d)} \mathbb{P}(D_r > t) \mathbb{P}_R(\mathrm{d}r), \tag{1}$$

where $\mathbb{P}_R$ denotes the distribution of the random variable R. Let $\varepsilon > 0$ be arbitrary. In order to analyze (1), we first consider the case where $t \geq (v(\varepsilon)^{-1}r)^{\alpha+d}$, i.e., where $r \leq v(\varepsilon)t^{1/(\alpha+d)}$, writing $v(\varepsilon) = (1+\varepsilon)^{-1}(d\kappa_d/(\alpha+d))^{-1/(\alpha+d)}$. Let N_r be a Poissonian random variable with mean $\kappa_d r^d$, where $r > 0$. Additionally, $\{U_i\}_{i \geq 1}$ be an iid sequence of random vectors that are independent of N_r and uniformly distributed in $B_1(o)$, so that $D_r = r^\alpha \sum_{i=1}^{N_r} |U_i|^\alpha$. First,

$$\mathbb{P}(D_r > t) \leq \mathbb{P}(N_r \geq n_{r,t}) + \mathbb{P}\Big(\sum_{i=1}^{n_{r,t}} r^\alpha |U_i|^\alpha \geq t\Big),$$

where $n_{r,t} = \lfloor tr^{-\alpha}\kappa_d v(\varepsilon)^{\alpha+d}(1+\varepsilon)^{(\alpha+d)/2}\rfloor$. Note that for sufficiently large values of $tr^{-\alpha}$ we have

$$n_{r,t}/(\kappa_d r^d) \geq t(v(\varepsilon)r^{-1})^{\alpha+d}(1+\varepsilon)^{(\alpha+d)/4} \geq (1+\varepsilon)^{(\alpha+d)/4},$$

so that the Poisson concentration property [19, Lemma 1.2] implies that $\sup_{r \leq v(\varepsilon)t^{1/(\alpha+d)}} \mathbb{P}(N_r \geq n_{r,t})$ decays at least exponentially fast in $t^{d/(\alpha+d)}$ as $t \to \infty$. Similarly, taking into account $\mathbb{E}|U_1|^\alpha = d/(\alpha+d)$, we obtain that

$$tr^{-\alpha}/(n_{r,t}\mathbb{E}|U_1|^\alpha) \geq (1+\varepsilon)^{(\alpha+d)/2},$$

so that using the classical theory of large deviations shows that also the expression

$$\sup_{r \leq v(\varepsilon)t^{1/(\alpha+d)}} \mathbb{P}\Big(\sum_{i=1}^{n_{r,t}} r^\alpha |U_i|^\alpha \geq t\Big)$$

decays at least exponentially fast in $t^{d/(\alpha+d)}$ as $t \to \infty$. In particular,

$$\lim_{t\to\infty} t^{s/(\alpha+d)} p^{(\alpha)}_{\mathsf{out},t} = \lim_{t\to\infty} \int_{v(\varepsilon)t^{1/(\alpha+d)}}^{\infty} t^{s/(\alpha+d)} \mathbb{P}(D_r > t) \mathbb{P}_R(\mathrm{d}r), \qquad (2)$$

provided that the latter limit exists and is finite. In order to compute the right-hand side in (2), we split the integral into three parts, which are analyzed separately. To be more precise, put

$$I_1 = \int_{v(\varepsilon)t^{1/(\alpha+d)}}^{v(-\varepsilon)t^{1/(\alpha+d)}} t^{s/(\alpha+d)} \mathbb{P}(D_r > t) \mathbb{P}_R(\mathrm{d}r),$$

$$I_{2,1} = -\int_{v(-\varepsilon)t^{1/(\alpha+d)}}^{\infty} t^{s/(\alpha+d)} \mathbb{P}(D_r \le t) \mathbb{P}_R(\mathrm{d}r),$$

and

$$I_{2,2} = \int_{v(-\varepsilon)t^{1/(\alpha+d)}}^{\infty} t^{s/(\alpha+d)} \mathbb{P}_R(\mathrm{d}r) = t^{s/(\alpha+d)} \mathbb{P}(R > v(-\varepsilon)t^{1/(\alpha+d)}).$$

First, note that $I_{2,2}$ tends to $v(-\varepsilon)^{-s}\beta$ as $t \to \infty$ and the latter expression tends to $(d\kappa_d/(\alpha + d))^{s/(\alpha+d)}$ as $\varepsilon \to 0$. Hence, it suffices to show that the integrals I_1 and $I_{2,1}$ tend to 0 as we let first $t \to \infty$ and then $\varepsilon \to 0$. Indeed,

$$I_1 \le t^{s/(\alpha+d)}\big(\mathbb{P}(R > v(\varepsilon)t^{1/(\alpha+d)}) - \mathbb{P}(R > v(-\varepsilon)t^{1/(\alpha+d)})\big),$$

and the right-hand side tends to $\beta(v(\varepsilon)^{-s} - v(-\varepsilon)^{-s})$ as $t \to \infty$, which vanishes as $\varepsilon \to 0$. Finally, we observe that

$$-I_{2,1} \le t^{s/(\alpha+d)} \mathbb{P}(D_{v(-\varepsilon)t^{1/(\alpha+d)}} \le t) \mathbb{P}(R > v(-\varepsilon)t^{1/(\alpha+d)}).$$

We conclude the proof of the theorem by noting that the expression $t^{s/(\alpha+d)} \mathbb{P}(R > v(-\varepsilon)t^{1/(\alpha+d)})$ remains bounded as $t \to \infty$, whereas $\mathbb{P}(D_{v(-\varepsilon)t^{1/(\alpha+d)}} \le t)$ tends to 0. Indeed, $D_{v(-\varepsilon)t^{1/(\alpha+d)}}$ has expectation $(1-\varepsilon)^{-(\alpha+d)}t$ and variance

$$v(-\varepsilon)^{2\alpha+d}t^{(2\alpha+d)/(\alpha+d)}v_{2\alpha}^{\alpha+d},$$

so that the latter claim follows from the Chebyshev inequality. $\qquad\square$

The proof of Theorem 2 relies on a rather delicate comparison of $D^{(\alpha)}_{\mathsf{in}}$ and the random variable $\max_{(\xi,r)\in X} |\xi|^\alpha \mathbb{1}_{B_r(o)}(\xi)$.

Proof of Theorem 2. First, observe that the in-degree of the origin $D^{(0)}_{\mathsf{in},n}$ is a Poissonian random variable with mean

$$\int_{\mathbb{T}_n} \mathbb{P}(R > |\xi|)\mathrm{d}\xi,$$

which is at least $d\kappa_d \int_0^{n/2} r^{d-1}\mathbb{P}(R > r)\mathrm{d}r$. If $s \le d$, then this lower bound tends to ∞ as $n \to \infty$. Using $D^{(\alpha)}_{\mathsf{in},n} \le D^{(\alpha)}_{\mathsf{in}}$, this shows that $\mathbb{P}(D^{(\alpha)}_{\mathsf{in}} = \infty) = 1$ and that $(D^{(\alpha)}_{\mathsf{in},n})_{n\ge1}$ converge in distribution to $D^{(\alpha)}_{\mathsf{in}}$. In the following, we may therefore assume that $s > d$. To show the assertion on the convergence in distribution, we proceed as in the proof of Theorem 2. Indeed, using the canonical coupling between the $[0,\infty)$-marked Poisson point process $X^{(n)}$ on the torus $\mathbb{T}_n$ and the $[0,\infty)$-marked Poisson point process X on $\mathbb{R}^d$, we deduce that

$$\mathbb{P}(D^{(\alpha)}_{\mathsf{in},n} \neq D^{(\alpha)}_{\mathsf{in}}) \le \mathbb{P}(X \cap S \setminus (B_{n/2}(o) \times [0,\infty)) \neq \emptyset),$$

where
$$S = \{(\xi, r) \in \mathbb{R}^d \times [0, \infty) : r > |\xi|\}.$$

Note that $\#(X \cap S \setminus (B_{n/2}(o) \times [0, \infty)))$ is Poissonian with mean

$$\int_{\mathbb{R}^d \setminus B_{n/2}(o)} \mathbb{P}(R > |\xi|) \mathrm{d}\xi = d\kappa_d \int_{n/2}^{\infty} r^{d-1} \mathbb{P}(R > r) \mathrm{d}r.$$

The assumption $s > d$ implies that the latter integral tends to 0 as $n \to \infty$. We also observe that by definition, $D_{\text{in}}^{(0)}$ is a Poissonian random variable with mean $\mathbb{E}\#(X \cap S)$. Moreover, this mean can be expressed as

$$\int_{\mathbb{R}^d} \mathbb{P}(R > |\xi|) \mathrm{d}\xi = d\kappa_d \int_0^{\infty} r^{d-1} \mathbb{P}(R > r) \mathrm{d}r = \kappa_d \mathbb{E}R^d.$$

It remains to determine the tail behavior of $D_{\text{in}}^{(\alpha)}$, for $\alpha > 0$. The derivation of the lower bound is based on the elementary relation

$$\mathbb{P}(X \cap S \setminus (B_{t^{1/\alpha}}(o) \times [0, \infty)) \neq \emptyset) = \mathbb{P}\big(\max_{(\xi,r) \in X \cap S} |\xi|^\alpha \geq t\big) \leq \mathbb{P}\big(D_{\text{in}}^{(\alpha)} \geq t\big). \qquad (3)$$

Let $\varepsilon > 0$ be arbitrary. Note that the random variable $\#(X \cap S \setminus (B_{t^{1/\alpha}}(o) \times [0, \infty)))$ is Poissonian with mean

$$\int_{\mathbb{R}^d \setminus B_{t^{1/\alpha}}(o)} \mathbb{P}(R \geq |\xi|) \mathrm{d}\xi = d\kappa_d \int_{t^{1/\alpha}}^{\infty} r^{d-1} \mathbb{P}(R \geq r) \mathrm{d}r,$$

and observe that for all $t > t_0(\varepsilon)$ the latter expression is bounded from below by

$$(1 - \varepsilon) d\kappa_d \beta \int_{t^{1/\alpha}}^{\infty} r^{d-s-1} \mathrm{d}r = (1 - \varepsilon)(s - d)^{-1} d\kappa_d \beta t^{(d-s)/\alpha}.$$

Therefore,

$$\begin{aligned}
\mathbb{P}\big(D_{\text{in}}^{(\alpha)} \geq t\big) &\geq 1 - \exp(-(1 - \varepsilon)(s - d)^{-1} d\kappa_d \beta t^{(d-s)/\alpha}) \\
&\geq (1 - 2\varepsilon)(s - d)^{-1} d\kappa_d \beta t^{(d-s)/\alpha},
\end{aligned}$$

provided that $t > 0$ is sufficiently large. For the upper bound, we need more refined arguments, since in principle $D_{\text{in}}^{(\alpha)}$ could be larger than t, even if $\max_{(\xi,r) \in X \cap S} |\xi|^\alpha \leq t$. To achieve the desired upper bound we show, loosely speaking, that the behavior of the sum $D_{\text{in}}^{(\alpha)}$ is already determined by its largest summand. Let $\varepsilon \in (0, 1)$ be arbitrary. Then our upper bound is based on the inequality

$$\begin{aligned}
\mathbb{P}(D_{\text{in}}^{(\alpha)} > t) \leq{}& \mathbb{P}\big(\max_{(\xi,r) \in X \cap S} |\xi| \geq (t(1 - \varepsilon))^{1/\alpha}\big) + \mathbb{P}(\#(X \cap S) \geq \varepsilon t^{1/4}) \\
&+ \mathbb{P}(\#(X \cap S \setminus (B_{t^{3/(4\alpha)}}(o) \times [0, \infty))) \geq 2),
\end{aligned}$$

and in the rest of the proof the three summands on the right-hand side are considered separately. Using similar bounds as in the derivation of the lower bounds, we obtain that

$$\lim_{t \to \infty} t^{(s-d)/\alpha} \mathbb{P}\big(\max_{(\xi,r) \in X \cap S} |\xi| \geq (t(1 - \varepsilon))^{1/\alpha}\big) = d\kappa_d \beta (s - d)^{-1}((1 - \varepsilon))^{(d-s)/\alpha}.$$

Furthermore, $\mathbb{P}(\#(X \cap S) \geq \varepsilon t^{1/4})$ tends to 0 exponentially fast in $t^{1/4}$ since $\#(X \cap S)$ is a Poissonian random variable with finite mean.

Finally, $\#(X \cap S \setminus (B_{t^{3/(4\alpha)}}(o) \times [0, \infty)))$ is a Poissonian random variable whose mean is at most $2d\kappa_d\beta(s-d)^{-1}t^{3(d-s)/(4\alpha)}$. Hence, the Poisson concentration property yields

$$\mathbb{P}(\#(X \cap S \setminus (B_{t^{3/(4\alpha)}}(o) \times [0, \infty))) \geq 2) \in O(t^{-3(s-d)/(2\alpha)}).$$

Since $\varepsilon > 0$ was arbitrary, this completes the proof of Theorem 2. $\qquad\square$

Theorems 1 and 2 show that on the distributional level, there is a substantial difference between power-weighted sums of outgoing and incoming edge lengths. However, when moving to the level of expectations, these differences disappear. In fact, the equality of expectations is an immediate consequence of Slivnyak's theorem, and is true for a much more general class of Poisson-based random geometric graphs. Still, for the convenience of the reader, we present some details.

Proposition 9. *Let $\alpha > 0$ be arbitrary. Then $\mathbb{E}D_{\mathrm{out},n}^{(\alpha)} = \mathbb{E}D_{\mathrm{in},n}^{(\alpha)}$.*

Proof. For $(\xi, r), (\eta, t) \in X^{(n)}$ we write $((\xi, r), (\eta, t)) \in G(X^{(n)})$ if there is a directed edge from (ξ, r) to (η, t) in the graph $G(X^{(n)})$. Then an application of Slivnyak's theorem yields

$$n^d\mathbb{E}D_{\mathrm{out},n}^{(\alpha)} = \mathbb{E}\sum_{(\xi,r)\in X^{(n)}}\sum_{(\eta,t)\in X^{(n)}} \mathrm{d}_{\mathbb{T}_n}(\xi, \eta)^\alpha 1_{((\xi,r),(\eta,t))\in G(X^{(n)})}$$

$$= \mathbb{E}\sum_{(\eta,t)\in X^{(n)}}\sum_{(\xi,r)\in X^{(n)}} \mathrm{d}_{\mathbb{T}_n}(\xi, \eta)^\alpha 1_{((\xi,r),(\eta,t))\in G(X^{(n)})}$$

$$= n^d\mathbb{E}D_{\mathrm{in},n}^{(\alpha)}.$$

$\qquad\square$

Now, we can proceed with the proof of Theorem 3.

Proof of Theorem 3. We begin with part (i). Noting that for every $n \geq 1$ the random variable $D_{\mathrm{in},n}^{(\alpha)}$ is stochastically dominated by the random variable $D_{\mathrm{in}}^{(\alpha)}$ and that the random variables $\{D_{\mathrm{in},n}^{(\alpha)}\}$ converge to $D_{\mathrm{in}}^{(\alpha)}$ in distribution, it suffices to show that $\mathbb{E}D_{\mathrm{in}}^{(\alpha)} = d\kappa_d(\alpha+d)^{-1}\mathbb{E}R^{\alpha+d}$. Indeed, Campbell's formula implies that

$$\mathbb{E}D_{\mathrm{in}}^{(\alpha)} = \int_{\mathbb{R}^d} |\xi|^\alpha\mathbb{P}(R > |\xi|)\mathrm{d}\xi = d\kappa_d\int_0^\infty r^{\alpha+d-1}\mathbb{P}(R > r)\mathrm{d}r = d\kappa_d(\alpha+d)^{-1}\mathbb{E}R^{\alpha+d}.$$

For parts (ii) and (iii) we can also proceed by applying Campbell's formula. Indeed, let $\varepsilon > 0$ be arbitrary and choose $t_0 = t_0(\varepsilon)$ as in Section 2.1. For part (ii), we first obtain that the expressions

$$(\log n)^{-1}\int_{B_{t_0}^{\mathbb{T}_n}(o)} |\xi|^\alpha\mathbb{P}(R > |\xi|)\mathrm{d}\xi$$

and

$$(\log n)^{-1}\int_{\mathbb{T}_n\setminus B_{n/2}^{\mathbb{T}_n}(o)} |\xi|^\alpha\mathbb{P}(R > |\xi|)\mathrm{d}\xi \leq 2\beta(\log n)^{-1}\int_{\mathbb{T}_1\setminus B_{1/2}^{\mathbb{T}_1}(o)} |\eta|^{\alpha-s}\mathrm{d}\eta$$

tend to 0 as $n \to \infty$. Furthermore,

$$(\log n)^{-1}\int_{B_{n/2}^{\mathbb{T}_n}(o)\setminus B_{t_0}^{\mathbb{T}_n}(o)} |\xi|^\alpha\mathbb{P}(R > |\xi|)\mathrm{d}\xi \leq (1+\varepsilon)(\log n)^{-1}d\kappa_d\beta\int_{t_0}^{n/2} r^{-1}\mathrm{d}r$$

$$= (1+\varepsilon)(\log n)^{-1}d\kappa_d\beta(\log n - \log 2 - \log t_0).$$

10

Since an analogous argument gives the corresponding lower bound, this completes the proof of part (ii). It remains to deal with part (iii). Proceeding similarly to part (ii), we first note that

$$n^{s-\alpha-d} \int_{B_{t_0}^{\mathbb{T}_n}(o)} |\xi|^\alpha \mathbb{P}(R > |\xi|) \mathrm{d}\xi$$

tends to 0 as $n \to \infty$. Furthermore,

$$n^{s-\alpha-d} \int_{\mathbb{T}_n \setminus B_{t_0}^{\mathbb{T}_n}(o)} |\xi|^\alpha \mathbb{P}(R > |\xi|) \mathrm{d}\xi \le (1+\varepsilon)\beta \int_{\mathbb{T}_1 \setminus B_{t_0/n}^{\mathbb{T}_1}(o)} |\eta|^{\alpha-s} \mathrm{d}\eta,$$

and the right-hand side tends to $(1+\varepsilon)\beta \int_{[-1/2,1/2]^d} |\eta|^{\alpha-s} \mathrm{d}\eta$ as $n \to \infty$. Again, since the lower bound can be obtained using similar arguments, this completes the proof of part (iii). $\qquad\square$

4   Chemical distances

In the present section we investigate the behavior of chemical distances (i.e., shortest-path lengths) on scale-free Gilbert graphs. We have already mentioned in the introduction that the regime, where $s < d$ might be of limited interest, as the diameter is at most 2 whp.

4.1   Regime $s > d$

In the present subsection, we consider the regime $s > d$ whose connectivity properties turn out to be rather similar to those of the Boolean model with light-tailed radii. For instance, we note that isolated vertices may occur with positive probability.

Proposition 10. *Assume that $s > d$. Then*

$$\lim_{n \to \infty} \mathbb{P}(o \text{ is isolated in } G(X^{(n)} \cup \{(o, R^*)\})) > 0.$$

Proof. First, we note that similar arguments as in Theorems 1 and 2 can be used to show that the probabilities $\mathbb{P}(o \text{ is isolated in } G(X^{(n)} \cup \{(o, R^*)\}))$ converge as $n \to \infty$. Moreover, the events $\{D_{\mathsf{in},n}^{(0)} = 0\}$ and $\{D_{\mathsf{out},n}^{(0)} = 0\}$ are both decreasing events, so that the FKG inequality (see, e.g. [14, Theorem 1.4]) implies that they are positively correlated. Hence,

$$\mathbb{P}(o \text{ is isolated in } G(X^{(n)} \cup \{(o, R^*)\})) \ge \mathbb{P}(D_{\mathsf{in},n}^{(0)} = 0)\mathbb{P}(D_{\mathsf{out},n}^{(0)} = 0),$$

so that using Theorems 1 and 2 completes the proof. $\qquad\square$

Next, we prove Theorem 4, which shows that allowing random radii with tail index $s > d$ does not reduce substantially chemical distances in comparison to the case of constant radii. To be more precise, the reduction amounts at most to a sublogarithmic factor. The key idea for the proof of Theorem 4 is to analyze the connections in the graph $G(X^{(n)})$ at different length scales. Although at every scale the presence of long edges can be used to reduce chemical distances, still such shortcuts are sufficiently rare to yield only a sublogarithmic reduction factor in comparison to the linear growth rate n. Before going into the (technical) details, we provide a rough sketch of the proof. Put $p = (s+d)/(2s)$, and subdivide the torus $\mathbb{T}_n$ into $k_1 = n^{(1-p)d}$ subcubes $Q_1, \ldots, Q_{k_1}$ of side length n^p. The probability

that there exists $(\xi, r) \in X^{(n)}$ with $r \geq n^p$ is of order at most $n^d n^{-sp} = n^{-(s-d)/2}$. On the other hand, if such a point does not exist, then the endpoints of any edge in $G(X^{(n)})$ are contained in adjacent subcubes. Hence, the number of subcubes that we need to visit if we move from $q(-n\mathrm{e}_1/4)$ to $q(n\mathrm{e}_1/4)$ is at least of order n^{1-p}.

Now, we continue by subdividing each of the subcubes $Q_1, \ldots, Q_{k_1}$ into subsubcubes of side length n^{p^2}. For each of the subcubes Q_i the probability that there exists $(\xi, r) \in X^{(n)} \cap Q_i$ with $r \geq n^{p^2}$ is of order at most $n^{pd} n^{-sp^2} = n^{-(s-d)p/2}$. The part of the path connecting $q(-n\mathrm{e}_1/4)$ and $q(n\mathrm{e}_1/4)$ in any such subcube is a nearest-neighbor path on the level of subsubcubes, so that typically n^{p-p^2} steps are needed to cross that subcube. If this was true for all $i \in \{1, \ldots, k_1\}$, then we would obtain a lower bound for the chemical distance between $q(-n\mathrm{e}_1/4)$ and $q(n\mathrm{e}_1/4)$ that is of order n^{1-p^2} and continuing in this fashion, we would obtain in fact a lower bound that is linear in n. However, at each level we have to deal with a small loss, which leads to the sublogarithmic correction term in the final lower bound.

In order to make this argument precise, it is convenient to introduce some notation similar to that used in fractal percolation, see [12]. In order to define precisely the iteration mentioned in the previous paragraph, we need to ensure that at each layer the number of subcubes is an integer. Therefore, we define $a_0 = n$, and then, inductively, $a_k = a_{k-1}/\lfloor a_{k-1}/n^{p^k} \rfloor$. We also put $b_k = a_{k-1}/a_k$. Next, in order to determine the position of subcubes in the kth layer, we use the index set

$$J_k = \{(i_1, \ldots, i_k) \in (\mathbb{Z}^d)^k : i_j \in \{0, \ldots, b_j - 1\}^d \text{ for all } j \in \{1, \ldots, k\}\}.$$

For $k \geq 1$ and $I = (i_1, \ldots, i_k) \in J_k$, we define the site $z_I = a_1 i_1 + \cdots + a_k i_k$ and the cube

$$Q_I = (-n/2, \ldots, -n/2) + z_I + [0, 1]^d a_k,$$

which is also called a k-*cube*. Note that we think of the Q_I as being embedded in the torus $\mathbb{T}_n$ so that it is possible that $Q_I \cap Q_{I'} \neq \emptyset$, even if the d_∞-distance between z_I and $z_{I'}$ is strictly larger than a_k. We say that $I, I' \in J_k$ are $*$-connected if there exist $I_1 = I, \ldots, I_m = I'$ such that $Q_{I_j} \cap Q_{I_{j+1}} \neq \emptyset$ for all $j \in \{1, \ldots, m-1\}$. In the following, we frequently consider certain neighborhoods of cubes of the form Q_I for some $I \in J_k$. To be more precise, for $r > 0$ we denote by $Q_I^r = \{\xi \in \mathbb{T}_n : \mathrm{d}_\infty^{\mathbb{T}_n}(\xi, Q_I) \leq r\}$ the subset of all $\xi \in \mathbb{T}_n$ such that the toroidal d_∞-distance from ξ to the cube Q_I is at most r.

Next, we need to capture the property that a path γ in $G(X^{(n)})$ that starts in a cube Q_I needs a large number of hops to move far away from this cube. To be more precise, for $\varepsilon > 0$ and $k \geq 1$ we introduce the notion of (ε, k)-good indices. If $I \in J_k$, then I is always (ε, k)-good. Furthermore, inductively, if $I \in J_{k'}$ is such that $0 \leq k' < k$, then we say that I is (ε, k)-good if a) $X^{(n)} \cap (Q_I \times [a_{k'+1}, a_{k'}]) = \emptyset$ and b) for every $*$-connected subset $\gamma \subset J_{k'+1}$ that is contained in $Q_I^{a_{k'}}$ and consists of at least $b_{k'+1}/4$ elements, it holds that γ contains at most $\varepsilon \# \gamma$ elements that are (ε, k)-bad. If $I = \emptyset$, then we additionally assume that $X^{(n)} \cap (Q_I \times [n, \infty)) = \emptyset$. Sometimes, we also say that the cube Q_I is (ε, k)-good if the index I has this property. Note that it would be more intuitive if condition a) required that $X^{(n)} \cap (Q_I \times [a_{k'+1}, \infty)) = \emptyset$. However, the present definition has the advantage that the (ε, k)-goodness of an index $I \in J_{k'}$ only depends on $X^{(n)} \cap (\mathbb{T}_n \times [0, a_{k'}])$. This property will be helpful in Lemma 16 below, where we establish a stochastic domination between the configuration of (ε, k)-good cubes and Bernoulli site percolation.

In the following, it will also be convenient to strengthen the notion of (ε, k)-good cubes in order to have some control over cubes in a suitable environment of a given one. To be more precise, for $u \geq 0$ and $k' \leq k$, we say that $I \in J_{k'}$ is (u, ε, k)-good if I' is (ε, k)-good for all $I' \in J_{k'}$ such that $Q_{I'} \subset Q_I^{ua_{k'}}$.

Next, in order to analyze a given self-avoiding path $\gamma = (x_1, \ldots, x_m)$ in $G(X^{(n)})$ at different scales, it is useful to introduce certain discretizations of γ. First, for any $k \geq 1$ we define a function $\mu_k : \mathbb{T}_n \times [0, \infty) \rightarrow J_k$, where $\mu_k(\xi, r)$ denotes the uniquely determined element $I \in J_k$ satisfying $\xi \in Q_I$. Note that in general applying μ_k to each element of γ results in a path that is no longer self-avoiding. A popular technique for transforming arbitrary paths into self-avoiding ones is Lawler's method of loop erasure [15]. Unfortunately, when performing loop erasure naively for discretizations at different scales, the resulting self-avoiding paths may be quite incomparable with respect to moving from one scale to another. Therefore, we consider a refinement of the standard loop-erasure method, which is adapted to dealing with different scales. For $k \geq 1$ define an ordered subset $\gamma^{(k,k,\mathsf{LE})}$ of γ, the (k, k)-*loop erasure of* γ, which can be identified with the standard loop erasure of the discretization of γ via μ_k. To be more precise, let $j \in \{1, \ldots, m\}$ be the largest integer such that $\mu_k(x_{j-1}) = \mu_k(x_1)$. Then define recursively $\gamma^{(k,k,\mathsf{LE})} = \left(x_1, \gamma_\mathsf{t}^{(k,k,\mathsf{LE})} \right)$, where $\gamma_\mathsf{t} = (x_j, \ldots, x_m)$ is the subpath of γ starting from x_j. Next, suppose that $k' < k$ and let $\gamma^{(k'+1,k,\mathsf{LE})} = (x_{m_1}, \ldots, x_{m_N})$ be the $(k' + 1, k)$-loop erasure of γ. In particular, $m_1 = 1$. Choose $j \in \{1, \ldots, N\}$ as the largest integer such that $\mu_{k'}(x_{m_{j-1}}) = \mu_{k'}(x_1)$. Then define recursively $\gamma^{(k',k,\mathsf{LE})} = \left(x_1, \gamma_\mathsf{t}^{(k',k,\mathsf{LE})} \right)$, where $\gamma_\mathsf{t} = (x_{m_j}, \ldots, x_m)$ is the subpath of γ starting from x_{m_j}. The construction of $\gamma^{(k',k,\mathsf{LE})}$ is illustrated in Figure 3.
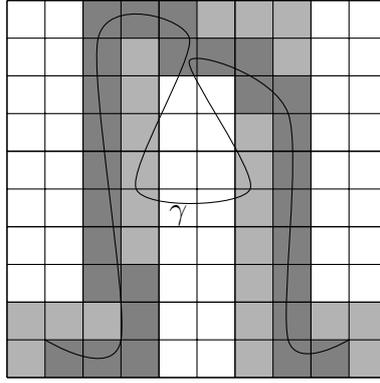

Figure 3: Construction of $\gamma^{(k,k,\mathsf{LE})}$ (dark gray) and $\gamma^{(k-1,k,\mathsf{LE})}$ (union of light and dark gray)

Next, we note that if $k' \leq k$ and γ are such that $\gamma^{(k',k,\mathsf{LE})}$ hits only (ε, k)-good $(k' - 1)$-cubes, then a large proportion of k'-cubes in $\gamma^{(k',k,\mathsf{LE})}$ are (ε, k)-good.

Lemma 11. *Let $k' \leq k$ and γ be a path in $G(X^{(n)} \cap (\mathbb{T}_n \times [0, a_{k'}]))$ hitting only (ε, k)-good $(k' - 1)$-cubes. Furthermore, assume that $\#\gamma^{(k',k,\mathsf{LE})} \geq b_{k'}/4$. Then, the number of $(3, \varepsilon, k)$-bad k'-cubes hit by $\gamma^{(k',k,\mathsf{LE})}$ is at most $147^d \varepsilon \#\gamma^{(k',k,\mathsf{LE})}$.*

Proof. Let $\gamma^{(k',k,\mathsf{LE})} = (x_{m_1}, \ldots, x_{m_N})$. We denote by γ^+ the family of all $I \in J_{k'}$ such that $Q_I \subset Q_{\mu_{k'}(x_{m_j})}^{3a_{k'}}$ for some $j \in \{1, \ldots, N\}$. In other words, γ^+ is obtained by a suitable dilation from the discretization of $\gamma^{(k',k,\mathsf{LE})}$. Then the number of $(3, \varepsilon, k)$-bad k'-cubes hit by $\gamma^{(k',k,\mathsf{LE})}$ is at most 7^d times the number of (ε, k)-bad k'-cubes in γ^+. Furthermore, for each $j \in \{1, \ldots, N\}$ denote by γ_j the $*$-connected component of $\gamma^+ \cap Q_{\mu_{k'-1}(x_{m_j})}^{a_{k'-1}}$ containing $\mu_{k'}(x_{m_j})$. Since γ hits only (ε, k)-good $(k' - 1)$-cubes, we conclude that for every $j \in \{1, \ldots, N\}$, the number of (ε, k)-bad k'-cubes in γ_j is at most $\varepsilon \#\gamma_j$. Moreover, note that for every $j_1, j_2 \in \{1, \ldots, N\}$ the components γ_{j_1} and γ_{j_2} either coincide or are disjoint. In the following, we fix a subset $S \subset \{1, \ldots, N\}$ with the property that a) for every $j \in \{1, \ldots, N\}$ there exists $s \in S$ such that $\gamma_j = \gamma_s$ and b) if $s, s' \in S$ are such that $\gamma_s = \gamma_{s'}$, then $s = s'$. Since the union $\cup_{s \in S} \gamma_s$ covers γ^+, the number of (ε, k)-bad k'-cubes in γ^+

is at most $\sum_{s \in S} \varepsilon \# \gamma_s$. Finally, noting that for each $I \in \gamma^+$ there exist at most 3^d elements $s \in S$ with $I \in \gamma_s$, we obtain that

$$\sum_{s \in S} \varepsilon \# \gamma_s \leq 3^d \varepsilon \# \gamma^+ \leq 21^d \varepsilon \# \gamma^{(k',k,\mathsf{LE})}.$$

This concludes the proof. $\qquad\square$

In the following, we use the discretizations $\{\gamma^{(k',k,\mathsf{LE})}\}_{k' \in \{1,\ldots,k\}}$ to derive suitable accurate lower bounds on the number of hops in γ. The most immediate approach to do this would start from counting the number elements in $\gamma^{(k',k,\mathsf{LE})}$ and multiplying this number by a suitable factor reflecting the scaling in the kth layer. However, if a k'-cube is hit by $\gamma^{(k',k,\mathsf{LE})}$, then this provides only very little information as to how many $(k'+1)$-cubes are hit by $\gamma^{(k'+1,k,\mathsf{LE})}$. Another approach could be to measure the Euclidean distance between the endpoints and multiply it by a suitable factor taking into account the scale of discretization. However, also this idea is problematic, since it is not clear how to determine an upper bound for the number of (ε,k)-bad cubes occurring in a path in terms of the distance of the endpoints. Therefore, in order to measure the length of a discretized path, we propose a slightly more refined approach. It is adapted to changing scales, and, moreover, the length still grows at least linearly in the number of elements of a path.

Let $k \geq 1$ and $\gamma = (x_1, \ldots, x_m)$ be a self-avoiding path in J_k. Let $D = (d_1, d_2, \ldots, d_{m'})$ be an ordered subset of $\{2, \ldots, m\}$. We say that D forms an *independent subset* of γ if $Q_{I_{d_i-1}} \cap Q_{I_{d_i+1}} = \emptyset$ for all $i \in \{1, \ldots, m'-1\}$. By $\lambda(\gamma)$ we denote the *length* of γ, which is defined as $m'_{\max} - 1$, where $m'_{\max}$ is the maximal size of an independent subset of γ. At first sight, it might seem unnatural to require $Q_{I_{d_i-1}} \cap Q_{I_{d_i+1}} = \emptyset$ instead of $Q_{I_{d_i}} \cap Q_{I_{d_i+1}} = \emptyset$. However, when considering a linear arrangement of $m \geq 2$ adjacent elements of J_k, then the second possibility would lead to a length of at most $\lceil m/2 \rceil$, whereas the first yields the more accurate value $m-2$.

In the following, we often consider the case, where $k' \leq k$ and $\gamma = (x_1, \ldots, x_m)$ is a self-avoiding path in $G(X^{(n)} \cap (\mathbb{T}_n \times [0, a_{k'}]))$. Then, an independent subset of $\mu_{k'}(\gamma^{(k',k,\mathsf{LE})})$ is also called a (k',k)-*independent subset* of γ and we write $\lambda_{k',k}(\gamma)$ for $\lambda(\mu_{k'}(\gamma^{(k',k,\mathsf{LE})}))$, which is called the (k',k)-*length* of γ. Next, we provide an affine lower bound for (k',k)-lengths in terms of number of sites.

Lemma 12. *Let $k' \leq k$ and γ be a self-avoiding path in $G(X^{(n)} \cap (\mathbb{T}_n \times [0, a_{k'}]))$. Then,*

$$\lambda_{k',k}(\gamma) + 1 \geq 10^{-d} \# \gamma^{(k',k,\mathsf{LE})}.$$

Proof. The proof proceeds by induction on $\#\gamma$. If $\lambda_{k',k}(\gamma) = 0$, then there exists $I \in J_{k'}$ such that γ remains inside $Q_I^{a_{k'}}$. In particular, $\#\gamma^{(k',k,\mathsf{LE})} \leq 3^d$. Otherwise, write $\gamma^{(k',k,\mathsf{LE})} = (x_{m_1}, \ldots, x_{m_N})$ and let $D = (d_1, \ldots, d_{m'})$ be an ordered subset of $\{1, \ldots, N\}$ such that the ordered set $(\mu_{k'}(x_{m_{d_1}}), \ldots, \mu_{k'}(x_{m_{d_{m'}}}))$ corresponds to a (k',k)-independent subset of γ of size $m' = \lambda_{k',k}(\gamma) + 1 \geq 2$. Furthermore, we also assume that among all (k',k)-independent subsets of maximal size, D is chosen as the lexicographic maximum. Next, we claim that $d_2 - d_1 \leq 9^d + 1$. Otherwise, we could choose $d' \in \{d_1+1, \ldots, d_2-1\}$ such that $Q_{\mu_{k'}(x_{m_{d'}})} \cap Q_{\mu_{k'}(x_{m_{d_1-1}})} = \emptyset$ and $Q_{\mu_{k'}(x_{m_{d'-1}})} \cap Q_{\mu_{k'}(x_{m_{d_2}})} = \emptyset$, contradicting the maximality property used to define D. Similarly, we see that $d_1 \leq 3^d + 1$. Therefore, $\#\gamma_{\mathsf{t}}^{(k,k',\mathsf{LE})} \geq \#\gamma^{(k,k',\mathsf{LE})} - 10^d$, where $\gamma_{\mathsf{t}} = (x_{m_{d_2-1}}, \ldots, x_m)$ denotes the subpath of γ starting at $x_{m_{d_2-1}}$. Furthermore, since D was chosen as the lexicographic maximum of all independent subsets of maximal size, we also see that $\lambda_{k',k}(\gamma) > \lambda_{k',k}(\gamma_{\mathsf{t}})$. Hence,

$$\lambda_{k',k}(\gamma) + 1 \geq \lambda_{k',k}(\gamma_{\mathsf{t}}) + 2 \geq 1 + 10^{-d} \# \gamma_{\mathsf{t}}^{(k',k,\mathsf{LE})} \geq 10^{-d} \# \gamma^{(k',k,\mathsf{LE})}.$$

□

Next, we show that (k', k)-lengths exhibit a good behavior with respect to changes in scale.

Lemma 13. *Let* $2 \leq k' \leq k$ *and* γ *be a path in* $G(X^{(n)} \cap (\mathbb{T}_n \times [0, a_{k'}]))$. *Then,*

$$\lambda_{k',k}(\gamma) \geq \lambda_{k'-1,k}(\gamma)b_{k'}.$$

Proof. Let $\gamma = (x_1, \ldots, x_m)$. We proceed by induction on $\#\gamma$, noting that the case $\lambda_{k'-1,k}(\gamma) = 0$ is trivial. Next, let $\gamma^{(k',k,\mathsf{LE})} = (x_{m_1}, \ldots, x_{m_N})$ and choose $u, v \in \{1, \ldots, N\}$ such that $\mu_{k'-1}(x_{m_u})$ and $\mu_{k'-1}(x_{m_v})$ are the first two elements in a maximal $(k'-1, k)$-independent subset of γ. Putting $I_1 = \mu_{k'-1}(x_{m_{u-1}})$ and $I_2 = \mu_{k'-1}(x_{m_v})$, we note that $Q_{I_1} \cap Q_{I_2} = \emptyset$. In particular, $Q_{\mu_{k'}(x_{m_{u-1}})} \not\subset Q^{a_{k'-1}}_{\mu_{k'}(x_{m_v})}$. Hence, denoting by $\gamma_{\mathsf{t}} = (x_{m_{v-1}}, \ldots, x_m)$ the subpath of γ starting from $x_{m_{v-1}}$, we obtain that

$$\lambda_{k',k}(\gamma) \geq b_{k'} + \lambda_{k',k}(\gamma_{\mathsf{t}}) \geq (1 + \lambda_{k'-1,k}(\gamma_{\mathsf{t}}))b_{k'} \geq \lambda_{k'-1,k}(\gamma)b_{k'},$$

as claimed. □

Let $\gamma = (x_1, \ldots, x_m) = ((\xi_1, r_1), \ldots, (\xi_m, r_m))$ be a self-avoiding path in $G(X^{(n)} \cap (\mathbb{T}_n \times [0, a_{k'}]))$ for some $k' < k$. In order to derive helpful lower bounds on m, we make use of the observation that inside any (ε, k)-good cube Q_I with $I \in J_{k'}$ the path γ consists of segments of length at most $a_{k'+1}$, and, moreover, most of the $(k'+1)$-cubes that hit this path are also (ε, k)-good. Hence, it is convenient to identify subpaths of γ that do not intersect (ε, k)-bad cubes. Also the subpaths should not be too short and, finally, some care should be taken to ensure compatibility with taking (k', k)-loop erasures.

To be more precise, let $k' \leq k$ and $\gamma = (x_1, \ldots, x_m)$ be a path in $G(X^{(n)} \cap (\mathbb{T}_n \times [0, a_{k'}]))$. Furthermore, let $A \subset \mu_{k'}(\gamma^{(k',k,\mathsf{LE})})$ and write $\gamma^{(k',k,\mathsf{LE})} = (x_{m_1}, \ldots, x_{m_N})$. Then, we define a family of disjoint subpaths $\Phi_A(\gamma) = \{\gamma_i\}_{i \in \{1,\ldots,\ell\}}$ inductively as follows.

- ■ If $\mu_{k'}(x_1) \notin A$, then $\Phi_A(\gamma) = \Phi_{A \setminus \{\mu_{k'}(x_1)\}}(x_{m_2}, \ldots, x_m)$. In other words, for the construction of $\Phi_A(\gamma)$ we discard the initial segment of $\gamma^{(k',k,\mathsf{LE})}$ not belonging to A.

- ■ Otherwise, let $L \in \{2, \ldots, N\}$ be the smallest integer such that $\mu_{k'}(x_{m_L}) \in A$.

    - ■ If for every $j \in \{1, \ldots, m_L\}$ there exists $I \in A$ such that $\xi_j \in Q_I^{3a_{k'}}$, then put $\ell = \ell'$, $\gamma_1 = \mathsf{concat}((x_1, \ldots, x_{m_L}), \gamma'_1)$ and $\gamma_i = \gamma'_i$ for $i \in \{2, \ldots, \ell'\}$, where $\Phi_{A \setminus \{\mu_{k'}(x_1)\}}((x_{m_L}, \ldots, x_m)) = \{\gamma'_i\}_{i \in \{1,\ldots,\ell'\}}$, and where concat denotes concatenation of paths. In other words, if until reaching x_{m_L} the path γ stays close to A, then to construct γ_1, then we proceed recursively as follows. We take the first subpath from $\Phi_{A \setminus \{\mu_{k'}(x_1)\}}((x_{m_L}, \ldots, x_m)) = \{\gamma'_i\}_{i \in \{1,\ldots,\ell'\}}$ and paste it to the subpath from x_1 to x_{m_L}.

    - ■ Otherwise, let $j \in \{1, \ldots, m_L\}$ be the smallest integer such that there does not exist $I \in A$ with $\xi_j \in Q_I^{3a_k}$ and put

        $$\Phi_A(\gamma) = \{(x_1, \ldots, x_{j-1})\} \cup \Phi_{A \setminus \{\mu_{k'}(x_1)\}}((x_{m_L}, \ldots, x_m)).$$

    In other words, γ_1 is the longest initial segment of γ that stays close to A; the other subpaths are constructed inductively.

(a) Path γ (solid line) and set A (shaded region)
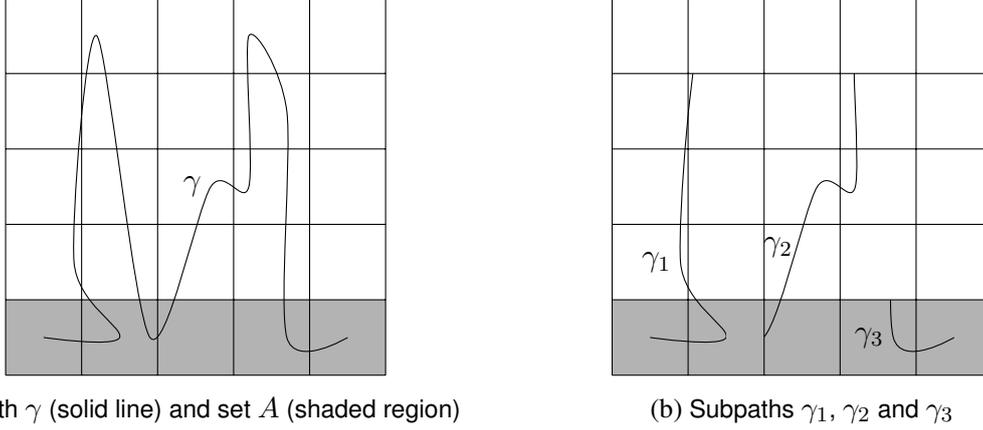

(b) Subpaths γ_1, γ_2 and γ_3

Figure 4: Construction of $\Phi_A(\gamma)$

The construction of $\Phi_A(\gamma)$ is illustrated in Figure 4.

Now, we can use the construction Φ_A to obtain (ε, k)-good cubes at smaller scales.

Lemma 14. *Let γ be a path in $G(X^{(n)} \cap (\mathbb{T}_n \times [0, a_{k'}]))$. Denote by $\{\gamma_i\}_{i=1}^{\ell}$ the collection of subpaths of γ obtained by applying the construction Φ_A with the family A as the set of those $(3, \varepsilon, k)$-good indices $I \in J_{k'}$ such that $\gamma^{(k',k,\mathsf{LE})}$ hits Q_I and $Q_I^{3a_{k'}}$ does not contain the endpoint of γ. Then, $A \subset \cup_{i=1}^{\ell} \mu_{k'}\big(\gamma_i^{(k',k,\mathsf{LE})}\big)$, and for every $i \in \{1, \ldots, \ell\}$,*

1. *the path γ_i hits only (ε, k)-good k'-cubes; in particular, γ_i is a path in $G(X^{(n)} \cap (\mathbb{T}_n \times [0, a_{k'+1}]))$,*

2. *the path γ_i starts in Q_I for some $I \in A$ and ends in $Q_{I'}$ for some I' with $Q_{I'} \cap Q_{I''} = \emptyset$ for all $I'' \in A$; in particular, $\#\gamma_i^{(k'+1,k,\mathsf{LE})} \geq b_{k'+1}/4$, and*

3. *$\sum_{i=1}^{\ell} \lambda_{k',k}(\gamma_i) \geq \lambda_{k',k}(\gamma) - \big(\#\gamma^{(k',k,\mathsf{LE})} - \#A\big)$.*

Proof. The relation $A \subset \cup_{i=1}^{\ell} \mu_{k'}\big(\gamma_i^{(k',k,\mathsf{LE})}\big)$ and claim 1 follow immediately from the definition of Φ_A. Claim 2 is shown by induction on $\#\gamma$, noting the assertion is trivial if $\ell = 0$. Next, we deal with claim 2 if $\ell > 0$. Write $\gamma^{(k',k,\mathsf{LE})} = (x_{m_1}, \ldots, x_{m_N})$. Clearly, our attention can be restricted to the case, where $\mu_{k'}(x_1) \in A$. Let $L \in \{2, \ldots, N\}$ be the smallest integer such that $\mu_{k'}(x_{m_L}) \in A$. If for every $j \in \{1, \ldots, m_L\}$ there exists $I \in A$ such that $\xi_j \in Q_I^{3a_{k'}}$, then the claim follows immediately from the induction hypothesis. Otherwise, let $j \in \{1, \ldots, m_L\}$ be the smallest integer such that there does not exist $I \in A$ with $\xi_j \in Q_I^{3a_{k'}}$. Since γ is a path in $G(X^{(n)} \cap (\mathbb{T}_n \times [0, a_{k'}]))$, we conclude that $Q_{\mu_{k'}(x_1)} \cap Q_{\mu_{k'}(x_{j-1})} = \emptyset$. This completes the proof of claim 2 for γ_1, whereas for γ_i with $i \geq 2$, we may conclude by induction.

It remains to prove claim 3. Choose an ordered subset $(d_1, \ldots, d_{m'})$ of $\{1, \ldots, N\}$ corresponding to a maximal (k', k)-independent subset of γ. Furthermore, for $i \in \{1, \ldots, \ell\}$ let D_i be the subset consisting of all $d' \in D$ such that $\mu_{k'}(x_{m_{d'}}) \in A$ and $x_{m_{d'}} \in \gamma_i$. Note that if we remove from D_i the element corresponding to the starting point of γ_i, then we obtain a (k', k)-independent subset of γ_i. Moreover, we assert that this independent subset can be enlarged by two further elements. Once this assertion is shown, we see that $\lambda_{k',k}(\gamma_i) \geq \#D_i$ and summing over all $i \in \{1, \ldots, \ell\}$ completes the proof. In order to prove the assertion, we choose $d_{i,1}$ as the largest index from D_i such that $\xi_{m_{d_{i,1}}} \in Q_I^{a_{k'}}$ for some $I \in A$. Similarly, we choose $d_{i,2}$ as the largest index such that

16

$\xi_{m_{d_{i,2}}} \in Q_I^{2a_{k'}}$ for some $I \in A$. Then, we can enlarge the independent subset D_i by adding $d_{i,1} + 1$ and $d_{i,2} + 1$. Indeed, for any $d' \in D_i$ we have $\mu_{k'}(x_{m_{d'}}) \in A$, whereas there is no $I \in A$ such that $Q_{\mu_{k'}(x_{m_{d_{i,1}+1}})} \subset Q_I^{a_{k'}}$. Similarly, $Q_{\mu_{k'}(x_{m_{d_{i,1}}})} \subset Q_I^{a_{k'}}$ for some $I \in A$, whereas there is no $I \in A$ such that $Q_{\mu_{k'}(x_{m_{d_{i,2}+1}})} \subset Q_I^{2a_{k'}}$. Therefore, $Q_{\mu_{k'}(x_{m_{d_{i,1}}})} \cap Q_{\mu_{k'}(x_{m_{d_{i,2}+1}})} = \emptyset$, and this completes the proof. □

Combining the previous auxiliary results, we now provide a lower bound for $\#\gamma$.

Lemma 15. *Let $k' \leq k$ and γ be a path in $G(X^{(n)} \cap (\mathbb{T}_n \times [0, a_{k'}]))$ hitting only (ε, k)-good $(k' - 1)$-cubes. Furthermore, assume that $\#\gamma^{(k',k,\mathsf{LE})} \geq b_{k'}/4$. Then there exists a constant $c > 0$ such that if $b_k \geq c$, then $\#\gamma \geq (1 - 1500^d\varepsilon)^{k-k'} \lambda_{k',k}(\gamma) \prod_{j=k'+1}^{k} b_j$.*

Proof. The proof proceeds via backward induction on k', the case $k' = k$ being trivial. First, let $A \subset J_{k'}$ be as in Lemma 14 and write $\Phi_A(\gamma) = \{\gamma_i\}_{i=1}^{\ell}$. By induction hypothesis and Lemma 13, we conclude that for every $i \in \{1, \dots, \ell\}$,

$$\#\gamma_i \geq (1 - 1500^d\varepsilon)^{k-k'-1} \lambda_{k'+1,k}(\gamma_i) \prod_{j=k'+2}^{k} b_j \geq (1 - 1500^d\varepsilon)^{k-k'-1} \lambda_{k',k}(\gamma_i) \prod_{j=k'+1}^{k} b_j.$$

By Lemma 14, we have $A \subset \cup_{i=1}^{m'} \mu_{k'}\big(\gamma_i^{(k',k,\mathsf{LE})}\big)$ and by Lemma 11 the number of $(3, \varepsilon, k)$-bad cubes hit by $\gamma^{(k',k,\mathsf{LE})}$ is at most $147^d\varepsilon\#\gamma^{(k',k,\mathsf{LE})}$. Hence, if b_k is sufficiently large, then part 3 of Lemma 14, Lemma 11 and Lemma 12 imply that

$$\sum_{i=1}^{\ell} \lambda_{k',k}(\gamma_i) \geq \lambda_{k',k}(\gamma) - 147^d\varepsilon\#\gamma^{(k',k,\mathsf{LE})} - 7^d \geq (1 - 1500^d\varepsilon)\lambda_{k',k}(\gamma),$$

so that

$$\sum_{i=1}^{\ell} \#\gamma_i \geq (1 - 1500^d\varepsilon)^{k-k'-1} \sum_{i=1}^{\ell} \lambda_{k',k}(\gamma_i) \prod_{j=k'+1}^{k} b_j \geq (1 - 1500^d\varepsilon)^{k-k'} \lambda_{k',k}(\gamma) \prod_{j=k'+1}^{k} b_j.$$

This completes the proof. □

Next, we show that the process of $(3, \varepsilon, k)$-good indices dominates a Bernoulli site percolation process with arbitrarily high marginal probability. This uses similar arguments as in [18, Lemma 2.2].

Lemma 16. *Let $\varepsilon > 0$ and $\rho \in (0, 1)$ be arbitrary. Furthermore, let $k = k(n)$ be such that $n^{-p^{k(n)}} \in o(1)$. Then there exists $n_0 \geq 1$ such that if $n \geq n_0$ and $k' \in \{0, \dots, k\}$, then the family of (ε, k)-good cubes in $J_{k'}$ stochastically dominates a Bernoulli site process on $J_{k'}$ with marginal probability ρ.*

Proof. In order to prove the claim, it is convenient to separate clearly the two conditions used in the definition of (ε, k)-goodness. First, an index $I \in J_{k'}$ is said to be *short-ranged* if $X^{(n)} \cap (\mathbb{T}_n \times [a_{k'+1}, a_{k'}]) = \emptyset$. Second, $I \in J_{k'}$ is said to be *iterable* if for every $*$-connected subset $\gamma \subset J_{k'+1}$ that is contained in $Q_I^{a_{k'}}$ and is of size at least $b_{k'+1}/4$, it holds that γ contains at most $\varepsilon\#\gamma$ indices that are (ε, k)-bad. By construction, the configuration of short-ranged indices in $J_{k'}$ is independent of the configuration of iterable indices in $J_{k'}$. Hence, to obtain the desired stochastic domination of a Bernoulli site process, it suffices to consider the two types of configurations separately.

To be more precise, let $\varepsilon > 0$ and $\rho \in (0,1)$ be arbitrary. Furthermore, put $q = 2^{-2\cdot 3^d/\varepsilon}$. We prove that there exist $n_0 \geq 1$ (depending only on d, ε and ρ) with the following properties, where without loss of generality we may assume that $\rho \geq 1 - q$.

1. Let $n \geq n_0$ and $k' \in \{1, \ldots, k\}$. Then the process of short-ranged indices in $J_{k'}$ stochastically dominates a Bernoulli site process on $J_{k'}$ with marginal probability given by $\sqrt{\rho}$.

2. Let $n \geq n_0$ and $k' \in \{1, \ldots, k\}$. If the process of (ε, k)-good indices in $J_{k'+1}$ dominates a Bernoulli site process on $J_{k'+1}$ with marginal probability given by $1 - q$, then the process of iterable indices in $J_{k'}$ dominates a Bernoulli site process on $J_{k'}$ with marginal probability given by $\sqrt{\rho}$.

Once these two claims are shown, we conclude by induction.

For the first claim, we note that the configuration of short-ranged elements in $J_{k'}$ is already a Bernoulli site process. Furthermore, provided that $n \geq 1$ is sufficiently large, the marginal probability of failing to be short-ranged is at most

$$2\beta a_{k'}^d a_{k'+1}^{-s} \leq 2\beta n^{p^{k'}d - p^{k'}ps} = 2\beta n^{-p^{k'}\varepsilon_1},$$

where $\varepsilon_1 = ps - d > 0$. In particular, the latter expression tends to 0 uniformly over all $k' \in \{1, \ldots, k(n)\}$ as $n \to \infty$.

It remains to prove the second claim. By assumption, the process of (ε, k)-good indices in $J_{k'+1}$ dominates a Bernoulli site process on $J_{k'+1}$ with marginal probability $1 - q$. An element $I \in J_{k'+1}$ is called *dom-bad* if it is a closed site in this Bernoulli site process. Furthermore, we say that $I \in J_{k'}$ is *dom-iterable* if for every $*$-connected subset $\gamma \subset J_{k'+1}$ that is contained in $Q_I^{a_{k'}}$ and consists of at least $b_{k'+1}/4$ elements, it holds that γ contains at most $\varepsilon\#\gamma$ elements that are dom-bad.

For any fixed $u \in \{1, \ldots, 3^d b_{k'+1}^d\}$, the cube $Q_I^{a_{k'}}$ contains at most $3^d b_{k'+1}^d 2^{(3^d-1)u}$ $*$-connected subsets consisting of precisely u elements, see [19, Lemma 9.3]. For any such $*$-connected set there exist at most 2^u possibilities to choose the location of dom-bad sites. Hence, the probability that I is not dom-iterable is at most

$$3^d b_{k'+1}^d \sum_{u \geq b_{k'+1}/4} 2^{(3^d-1)u} 2^u q^{\varepsilon u} \leq 3^d b_{k'+1}^d 2^{1-3^d\lceil b_{k'+1}/4\rceil}. \tag{4}$$

If $k' \in \{1, \ldots, k(n)\}$ and n is sufficiently large, then $b_{k'} \geq n^{p^{k'}(1-p)}/2$. Therefore, we see that (4) tends to 0 uniformly over all $k' \in \{1, \ldots, k(n)\}$ as $n \to \infty$. The proof of the second claim is concluded by invoking [16, Corollary 1.4]. Alternatively, it is also possible to apply [18, Theorem 2.1], which also yields an explicit bound for the value of n that is needed to achieve the desired domination. $\square$

Now, we have collected all preliminary results required to complete the proof of Theorem 4.

Proof of Theorem 4. Define $k = k(n) = \lfloor(\log^{(4)} n - \log^{(2)} n)/\log p\rfloor$, where for $m \geq 1$ we denote by $\log^{(m)} n$ the m-fold iterated logarithm. In particular,

$$n^{p^k} \in \big(\log^{(2)} n, (\log^{(2)} n)^{p^{-1}}\big).$$

18

Moreover, we conclude from Lemmas 15 and 16 that whp the number of hops between $q(-n\mathrm{e}_1/4)$ and $q(n\mathrm{e}_1/4)$ is at least

$$(1 - 1500^d\varepsilon)^{k-1}\lambda_{1,k}(\gamma)\prod_{j=2}^{k} b_j \geq \tfrac{1}{4}(1 - 1500^d\varepsilon)^{k-1}n^{1-p^k}$$

$$\geq \tfrac{1}{4}(1 - 1500^d\varepsilon)^{k-1}n(\log^{(2)} n)^{-p^{-1}}.$$

It thus remains to provide a suitable lower bound for $(1 - 1500^d\varepsilon)^k$. Indeed,

$$k\log(1 - 1500^d\varepsilon) \geq (-2\log(1 - 1500^d\varepsilon)/\log p)\log^{(2)} n,$$

so that choosing $\varepsilon > 0$ sufficiently close to 0 to ensure that $-\log(1 - 1500^d\varepsilon)/\log p \geq -\alpha/4$ shows that $(1 - 1500^d\varepsilon)^k \geq (\log n)^{-\alpha/2}$. $\qquad\square$

4.2 Regime $s = d$

In the present subsection, we consider the critical regime, where $s = d$ and provide a proof of Theorem 6. We first explain the main ideas before presenting all the details. If there exists $x = (\xi, r) \in X^{(n)}$ with $r \geq \sqrt{d}n$, then all points of $X^{(n)}$ are connected to x directly by an edge. On the other hand, if $\mathbb{T}_n$ is not covered by a single ball, then the torus can be subdivided into smaller subcubes, and we try to cover these subcubes by smaller balls. While some are now covered, others will stay uncovered. At this points one proceeds iteratively, subdividing the remaining subcubes into subsubcubes and aiming at covering these smaller cubes by smaller balls. We claim that this algorithm terminates with high probability, yielding a connected random geometric graph whose diameter is bounded from above by the total number of subcubes introduced in this construction.

Similar to Section 4.1, in order to make these arguments rigorous, it is useful to highlight a link to fractal percolation. However, now the fractal percolation process is considerably simpler, since each occurring subcube is subdivided into precisely 2^d subsubcubes, irrespective of the level, in which the original subcube is located. In the present setting, the total index set $\mathcal{J}$ is therefore given by $\mathcal{J} = \cup_{m \geq 0}\{0, 1\}^{md}$, where the symbol $\cup$ is interpreted as disjoint union. Furthermore, for each $I = (i_1, \ldots, i_m) \in \mathcal{J}$ we put

$$Q_I = (-n/2, -n/2, \ldots, -n/2) + n\sum_{j=1}^{m} 2^{-j}i_j + [0, n2^{-m}]^d.$$

Proof of Theorem 6. In order to make the sketch presented at the beginning of the subsection rigorous, we introduce a fractal percolation process $\{Z(I)\}_{I \in \mathcal{J}}$. For $k \geq 0$ and $I \in \{0, 1\}^{kd}$ put $Z(I) = 0$ if and only if there exists $(\xi, r) \in X$ such that $\xi \in Q_I$ and $r \in (\sqrt{d}2^{-k+1}n, \sqrt{d}2^{-k+2}n)$. The number of points $(\xi, r) \in X$ such that $\xi \in Q_I$ and $r \in (\sqrt{d}2^{-k+1}n, \sqrt{d}2^{-k+2}n)$ is Poisson distributed with parameter

$$n^d2^{-kd}(\mathbb{P}(R > \sqrt{d}2^{-k+1}n) - \mathbb{P}(R > \sqrt{d}2^{-k+2}n)).$$

Provided that $\sqrt{d}2^{-k+1}n \geq t_0(1/2)$ this expression can be bounded from below by

$$n^d2^{-kd}d^{-d/2}\beta(2^{kd-d-1}n^{-d} - 3 \cdot 2^{kd-2d-1}n^{-d}) \geq \beta d^{-d/2}2^{-2d-1}.$$

Hence, if additionally $\beta \geq d^{d/2}2^{2d+1}(d+1)\log 2$, then

$$\mathbb{P}(Z(I) = 1) \leq \exp(-\beta d^{-d/2}2^{-2d-1}) \leq 2^{-d-1}. \tag{5}$$

Since each cube Q_I gives rise to 2^d subcubes, this already provides a first strong indication for the relationship to subcritical Galton-Watson processes.

To make this precise, it is convenient to introduce some auxiliary structures. By C_k, $k \geq 0$ we denote the union of retained cubes at the kth level. That is, $C_0 = [-n/2, n/2]^d$ and if $k \geq 0$, then

$$C_{k+1} = C_k \cap \bigcup_{I \in \{0,1\}^{(k+1)d}: Z(I)=1} Q_I.$$

Furthermore, we also consider the index sets $A_k^{(0)}, A_k^{(1)} \subset \{0,1\}^{kd}$ whose associated cubes are discarded/retained at the kth step. That is $A_k^{(\sigma)} = \{I \in \{0,1\}^{kd} : Q_I \subset C_{k-1} \text{ and } Z(I) = \sigma\}$. Finally, we construct a backbone $B \subset X^{(n)}$ such that if $C_m = \emptyset$ for some $m \geq 1$, then a) any point of $X^{(n)}$ is connected by an edge in $G(X^{(n)})$ to some point in B and b) B is a connected set in $G(X^{(n)})$. To construct B, we choose for each $I \in A_k^{(0)}$ a point $x_I \in X^{(n)}$ such that $\xi \in Q_I$ and $r \in (\sqrt{d}2^{-k+1}n, \sqrt{d}2^{-k+2}n)$. Then, we denote by B the collection consisting of the points x_I, where $I \in A_k^{(0)}$ for some $k \geq 0$. See Figure 5 for an illustration of the construction of B.
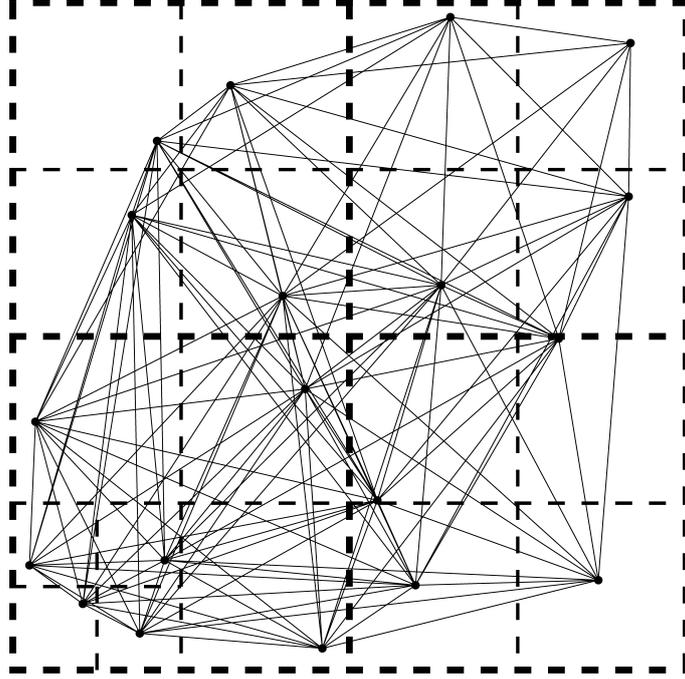

Figure 5: Backbone of a scale-free Gilbert graph

In particular,

$$\#B = \sum_{k \geq 0} \#A_{k+1}^{(0)} \leq 2^d \sum_{k \geq 0} \#A_k^{(1)}.$$

Now, we show that the set B has the desired properties. Since $C_m = \emptyset$, we conclude that the union of the cubes Q_I, with $I \in A_k^{(0)}$ for some $k \geq 0$ covers $\mathbb{T}_n$. If $x = (\xi, r) \in X^{(n)}$ is arbitrary, then by choosing $I \in A_k^{(0)}$ such that $\xi \in Q_I$ and noting that $Q_I \subset B_{r_I}^{\mathbb{T}_n}(\xi_I)$, we see that x is connected to x_I by an edge in $G(X^{(n)})$. Moreover, if $I_1 \in A_{k_1}^{(0)}$ and $I_2 \in A_{k_2}^{(0)}$ are such that $k_1 \leq k_2$

and $Q_{I_1} \cap Q_{I_2} \neq \emptyset$, then $Q_{I_1} \cup Q_{I_2} \subset B_{r_{I_1}}^{\mathbb{T}_n}(\xi_{I_1})$. Hence, x_{I_1} and x_{I_2} are connected by an edge in $G(X^{(n)})$. Since for any $I \in A_k^{(0)}$ and $I' \in A_{k'}^{(0)}$ we can find $I_1 = I, \ldots, I_m = I'$ such that $I_j \in A_{k_j}^{(0)}$ and $Q_{I_j} \cap Q_{I_{j+1}} \neq \emptyset$, this proves the second claim on B.

Hence, if there exists $m \geq 1$ such that $C_m = \emptyset$, then the diameter of $G(X^{(n)})$ is at most $2 + \#B$. It remains to compare $\#B$ to the total progeny of a subcritical Galton-Watson process; this analysis will also show that the probability that $C_m = \emptyset$ for some $m \geq 1$ tends to 1 as $n \to \infty$. Indeed, consider the subcritical Galton-Watson process whose offspring distribution is Binomial with 2^d trials and success probability 2^{-d-1}. Put $k = k_0(n) = \lfloor \log(2\sqrt{d}n/t_0(1/2))/\log 2 \rfloor$. It follows from (5) and the independence property of the Poisson point process that there exists a coupling between $X^{(n)}$ and the Galton-Watson process such that for every $k \leq k_0$ we have $T_k \geq \sum_{i=0}^{k} \#A_i^{(1)}$. Here T_k denotes the total progeny of the Galton-Watson process up to the kth step. In particular,

$$\max\{\mathbb{P}(C_{k_0} \neq \emptyset), \mathbb{P}(\#B > 2^d T_{k_0})\} \leq \mathbb{P}(T_{k_0-1} \neq T_{k_0}) \leq 2^{-k_0} \leq t_0(1/2)/(\sqrt{d}n),$$

where the second inequality uses a well-known result for Galton-Watson processes, see e.g. [13, Theorem 5.1]. $\qquad\square$

5 Using a multi-level topology for redundancy elimination

In the present section, we provide a comparison between the scale-free Gilbert graph $G(X^{(n)})$ and the thinned scale-free Gilbert graph $G'(X^{(n)})$, where some redundant edges are removed. First, we prove Theorem 7.

Proof of Theorem 7. In the following we fix $t_0 = t_0(1/2)$. We first show that we can neglect contributions $b_1(n)$ to $\mathbb{E}D_{\text{out},n}^{\prime,(\alpha)}$ coming from points $(\eta, t) \in X^{(n)}$ such that $t > \sqrt{|\eta|}$. Indeed,

$$
\begin{aligned}
b_1(n) &\leq \int_0^\infty \int_{B_r^{\mathbb{T}_n}(o)} |\eta|^\alpha \mathbb{P}(R > \sqrt{|\eta|})\mathrm{d}\eta \mathbb{P}_R(\mathrm{d}r) \\
&= \int_{\mathbb{T}_n} |\eta|^\alpha \mathbb{P}(R > |\eta|)\mathbb{P}(R > \sqrt{|\eta|})\mathrm{d}\eta \\
&= \int_{\mathbb{T}_n \setminus B_{t_0^2}^{\mathbb{T}_n}(o)} |\eta|^\alpha \mathbb{P}(R > |\eta|)\mathbb{P}(R > \sqrt{|\eta|})\mathrm{d}\eta + \int_{B_{t_0^2}^{\mathbb{T}_n}(o)} |\eta|^\alpha \mathbb{P}(R > |\eta|)\mathbb{P}(R > \sqrt{|\eta|})\mathrm{d}\eta.
\end{aligned}
$$

Since the second expression remains bounded as $n \to \infty$, it suffices to consider the first. Then,

$$
\begin{aligned}
\int_{\mathbb{T}_n \setminus B_{t_0^2}^{\mathbb{T}_n}(o)} |\eta|^\alpha \mathbb{P}(R > |\eta|)\mathbb{P}(R > \sqrt{|\eta|})\mathrm{d}\eta &\leq 2\beta^2 \int_{\mathbb{T}_n \setminus B_{t_0^2}^{\mathbb{T}_n}(o)} |\eta|^{\alpha - 3d/2}\mathrm{d}\eta \\
&= 2\beta^2 \int_{\mathbb{T}_n \setminus B_{n/2}^{\mathbb{T}_n}(o)} |\eta|^{\alpha - 3d/2}\mathrm{d}\eta + 2\beta^2 \int_{B_{n/2}^{\mathbb{T}_n}(o) \setminus B_{t_0^2}^{\mathbb{T}_n}(o)} |\eta|^{\alpha - 3d/2}\mathrm{d}\eta,
\end{aligned}
$$

and we consider the two summands separately. Clearly, the first is in $O(n^{\alpha - d/2})$. For the second we obtain that

$$
\int_{B_{n/2}^{\mathbb{T}_n}(o) \setminus B_{t_0^2}^{\mathbb{T}_n}(o)} |\eta|^{\alpha - 3d/2}\mathrm{d}\eta = d\kappa_d \int_{t_0^2}^{n/2} u^{\alpha - d/2 - 1}\mathrm{d}u,
$$

which is in $O(1)$ if $\alpha = 0$ and in $O(n^{\max\{0, \alpha - d/4\}})$ if $\alpha > 0$. Furthermore, we can clearly neglect the contributions $b_2(n)$ to $\mathbb{E}D_{\mathsf{out},n}^{\prime,(\alpha)}$ coming from points $(\eta, t) \in X^{(n)}$ with $|\eta| \leq \min\{e^{12}, t_0^2\}$.

It remains to obtain bounds for the contributions that are covered by neither $b_1(n)$ nor $b_2(n)$. For any $\gamma > 0$ and $\xi \in [-n/2, n/2]^d$ we put

$$S_{\gamma, \eta} = \{\zeta \in \mathbb{R}^d : \sqrt{|\eta|} \leq |\zeta - \eta| \leq |\eta| \text{ and } \angle(\zeta - \eta, -\eta) \in [-\gamma, \gamma]\};$$

see Figure 6 for an illustration of $S_{\gamma, \eta}$. If γ is sufficiently small, then for every $n \geq 1$ and $\eta \in [-n/2, n/2]^d$ the sector $S_{\gamma, \eta}$ is contained in $[-n/2, n/2]^d$. In the following, we fix any such value γ_0 and put $S_\eta = S_{\gamma_0, \eta}$. It will also be convenient to denote by $S'_\eta = \partial B_1(o) \cap (|\eta|^{-1}(S_\eta - \eta))$ the intersection of the unit sphere with a shifted and scaled copy of S_η. Finally, we denote by $\sigma_0 = \nu_{d-1}(S'_\eta)$ the surface area of S'_η, a quantity which is independent of η.
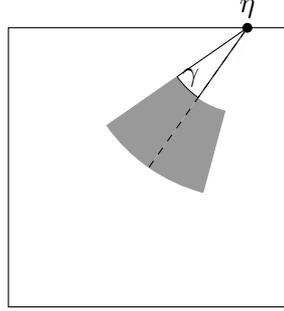

Figure 6: Construction of the set $S_{\gamma, \eta}$ (shaded)

If (η, t) is an out-neighbor of (o, R^*) in $G'(\{o, R^*\} \cup X^{(n)})$ such that $t \leq \sqrt{|\eta|}$, then $X^{(n)} \cap A_{\eta, t} = \emptyset$, where

$$A_{\eta, t} = \{(\zeta, w) \in S_\eta \times [0, R^*) : w > |\zeta - \eta|\}.$$

Furthermore, we note that conditioned on $R^* = r$ the number of points of $X^{(n)}$ contained in $A_{\eta, t}$ is a Poissonian random variable with mean

$$\int_{S_\eta} \mathbb{P}(R \in (|\zeta - \eta|, r)) \mathrm{d}\zeta = \sigma_0 \int_{\sqrt{|\eta|}}^{|\eta|} u^{d-1} \mathbb{P}(R \in (u, r)) \mathrm{d}u.$$

If $n \geq 1$ is sufficiently large and $|\eta| \geq t_0^2$, then the right-hand side is at least

$$\beta \sigma_0 \int_{\sqrt{|\eta|}}^{|\eta|} \tfrac{1}{2} u^{-1} - \tfrac{3}{2} u^{d-1} r^{-d} \mathrm{d}u \geq \beta \sigma_0 (\tfrac{1}{4} \log |\eta| - \tfrac{3}{2}) \geq \tfrac{1}{8} \beta \sigma_0 \log |\eta|.$$

Putting $b_3(n) = \mathbb{E}D_{\mathsf{out},n}^{\prime,(\alpha)} - b_1(n) - b_2(n)$, we therefore obtain that

$$\begin{aligned}
b_3(n) &\leq \int_0^\infty \int_{B_r^{\mathbb{T}_n}(o) \setminus B_{t_0^2}^{\mathbb{T}_n}(o)} |\eta|^\alpha \exp(-\tfrac{1}{8} \beta \sigma_0 \log |\eta|) \mathrm{d}\eta \, \mathbb{P}_R(\mathrm{d}r) \\
&= \int_{\mathbb{T}_n \setminus B_{t_0^2}^{\mathbb{T}_n}(o)} |\eta|^{\alpha - \beta \sigma_0/8} \mathbb{P}(R > |\eta|) \mathrm{d}\eta \\
&\leq 2\beta \int_{\mathbb{T}_n \setminus B_{n/2}^{\mathbb{T}_n}(o)} |\eta|^{\alpha - d - \beta \sigma_0/8} \mathrm{d}\eta + 2\beta \int_{B_{n/2}^{\mathbb{T}_n}(o) \setminus B_{t_0^2}^{\mathbb{T}_n}(o)} |\eta|^{\alpha - d - \beta \sigma_0/8} \mathrm{d}\eta \\
&\leq 2\beta \int_{\mathbb{T}_n \setminus B_{n/2}^{\mathbb{T}_n}(o)} |\eta|^{\alpha - d - \beta \sigma_0/8} \mathrm{d}\eta + 2\beta d \kappa_d \int_{t_0^2}^{n/2} u^{\alpha - \beta \sigma_0/8 - 1} \mathrm{d}u.
\end{aligned}$$

Observing that the last line is in $O(1)$ if $\alpha = 0$ and in $O(n^{\max\{0, \alpha - \beta\sigma_0/16\}})$ if $\alpha > 0$ completes the proof. $\qquad\square$

When passing from $G(X^{(n)})$ to $G'(X^{(n)})$ only redundant edges are removed, in the sense that connected components remain unchanged.

Proposition 17. *With probability* 1 *the graphs* $G(X^{(n)})$ *and* $G'(X^{(n)})$ *have the same connected components.*

Proof. It suffices to show that if $x = (\xi, r), y = (\eta, t) \in X^{(n)}$ are connected by an edge in $G(X^{(n)})$, then x and y are contained in the same connected component of $G'(X^{(n)})$. Suppose that this was false. Then, choose a counter-example with the property that $|r - t|$ is minimal. Without loss of generality we may assume $r \geq t$. Since $x = (\xi, r)$ and $y = (\eta, t)$ are not connected by an edge in $G'(X^{(n)})$, there exists a point $z = (\zeta, w) \in X^{(n)} \cap (B_r^{\mathbb{T}_n}(\xi) \times (t, r))$ such that $\eta \in B_w^{\mathbb{T}_n}(\zeta)$. By the minimality of the counter-example we see that both y and z as well as z and x are contained in the same connected component of $G'(X^{(n)})$. Therefore also x and y are contained in the same connected component of $G'(X^{(n)})$, yielding a contradiction to the initial assumption. $\qquad\square$

Next, we show that when moving from $G(X^{(n)})$ to $G'(X^{(n)})$ chemical distances increase at most by a logarithmic factor in the size of the torus. To achieve this goal, we make use of a variant of the descending chains concept introduced in [8]. To be more precise, let $x_1 = (\xi_1, r_1), \ldots, x_m = (\xi_m, r_m)$ be such that $r_1 > r_2 > \cdots > r_m$. Then $x_1, \ldots, x_m$ are said to form a *toroidal descending chain* if $\xi_{i+1} \in B_{r_i}^{\mathbb{T}_n}(\xi_i)$ for every $i \in \{1, \ldots, m - 1\}$. In the following result, we show that there is a close relationship between the existence of short connections in $G'(X^{(n)})$ and the absence of long toroidal descending chains.

Lemma 18. *Let* $x = (\xi, r) \in X^{(n)}$ *and* $y = (\eta, t) \in X^{(n)} \cap (B_r^{\mathbb{T}_n}(\xi) \times (0, r))$. *If* $m \geq 1$ *is such that the chemical distance between* x *and* y *in* $G'(X^{(n)})$ *is larger than* m, *then there exists a toroidal descending chain starting from* x *and consisting of more than* m *points.*

Proof. We proceed similarly to [1, Lemma 10]. Inductively, we construct points $\{x_j^{(i)}\}_{0 \leq j \leq i} = \{(\xi_j^{(i)}, r_j^{(i)})\}_{0 \leq j \leq i}$ such that for every $i \in \{1, \ldots, m\}$ we have

  (i) $r_0^{(i)} > \cdots > r_i^{(i)}$,

  (ii) $\xi_{j+1}^{(i)} \in B_{r_j^{(i)}}^{\mathbb{T}_n}(\xi_j^{(i)})$ for all $j \in \{1, \ldots, i - 1\}$, and

  (iii) there exists $j \in \{1, \ldots, i - 1\}$ such that $x_j^{(i)}$ and $x_{j+1}^{(i)}$ are not connected by an edge in $G'(X^{(n)})$.

Indeed, for the induction start we just choose $x_0^{(1)} = x$ and $x_1^{(1)} = y$. Next, suppose that $i \leq m$ and that we have constructed points $\{x_j^{(i-1)}\}_{0 \leq j \leq i}$ in $X^{(n)}$ satisfying properties (1)-(3). Choose $j_0 \in \{0, \ldots, i - 2\}$ such that $x_{j_0}^{(i-1)}$ and $x_{j_0+1}^{(i-1)}$ are not connected by an edge in $G'(X^{(n)})$. By definition of $G'(X^{(n)})$ there exists $x' = (\xi', r') \in X^{(n)} \cap \left(B_{r_{j_0}^{(i-1)}}^{\mathbb{T}_n}(\xi_{j_0}^{(i-1)}) \times \left(r_{j_0+1}^{(i-1)}, r_{j_0}^{(i-1)}\right)\right)$ such

23

that $\xi_{j_0+1}^{(i-1)} \in B_{r'}^{\mathbb{T}^n}(\xi')$. We put

$$
x_j^{(i)} = \begin{cases} x_j^{(i-1)} & \text{if } j \leq j_0, \\ x' & \text{if } j = j_0 + 1, \\ x_{j-1}^{(i-1)} & \text{if } j > j_0 + 1. \end{cases}
$$

Clearly, properties (i) and (ii) are satisfied. Since a violation of property (iii) would imply that the chemical distance between x and y in $G'(X^{(n)})$ is at most m, this completes the proof. $\qquad \square$

For the applicability of Lemma 18, it is important to show that long toroidal descending chains can occur only with small probability.

Lemma 19. *There exists a constant $c_1 > 0$ such that whp $X^{(n)}$ does not contain a toroidal descending chain consisting of more than $c_1 \log n$ elements.*

Proof. We distinguish several cases depending on the radii occurring in the chain. First, note that the probability that there exist distinct $(\xi, r), (\eta, t) \in X^{(n)}$ such that $\min\{r, t\} > n^{3/4}$ is at most $n^{2d}\mathbb{P}(R > n^{3/4})^2$, which tends to 0 as $n \to \infty$. Hence, it suffices to consider toroidal descending chains with radii bounded above by $n^{3/4}$. Next, we note that the expected number of toroidal descending chains consisting of $m+1$ steps and where all radii are bounded above by $t_0 = t_0(1/2)$ is at most $n^d(\kappa_d t_0^d)^m/m!$. Using Stirling's formula, we see that when putting $m = \log n$, this expression tends to 0 as $n \to \infty$. Finally, the expected number of toroidal descending chains consisting of $m \geq 1$ steps, and where all radii are contained in $(t_0, n^{3/4})$ is at most

$$
\int_{\mathbb{T}^n} \int_{t_0}^{n^{3/4}} \int_{B_{r_1}^{\mathbb{T}^n}(\xi_1)} \cdots \int_{t_0}^{r_{m-1}} \int_{B_{r_m}^{\mathbb{T}^n}(\xi_m)} 1 \mathrm{d}\xi_{m+1} \mathbb{P}_R(\mathrm{d}r_m) \cdots \mathrm{d}\xi_2 \mathbb{P}_R(\mathrm{d}r_1)\mathrm{d}\xi_1
$$

$$
= n^d \kappa_d^m \int_{t_0}^{n^{3/4}} \cdots \int_{t_0}^{n^{3/4}} 1_{r_1 > \cdots > r_m} r_1^d \cdots r_m^d \mathbb{P}_R(\mathrm{d}r_m) \cdots \mathbb{P}_R(\mathrm{d}r_1)
$$

$$
= \frac{n^d \kappa_d^m}{m!} \left(\int_{t_0}^{n^{3/4}} r^d \mathbb{P}_R(\mathrm{d}r) \right)^m.
$$

In order to derive an upper bound for the latter expression, we need to investigate $\mathbb{E}R^d 1_{(t_0, n^{3/4})}(R)$. We compute

$$
\mathbb{E}R^d 1_{(t_0, n^{3/4})}(R) = \int_0^{t_0^d} \mathbb{P}(R > t_0)\mathrm{d}t + \int_{t_0^d}^{n^{3d/4}} \mathbb{P}(R^d > t)\mathrm{d}t
$$

$$
\leq t_0^d + 2\beta \int_1^{n^{3d/4}} t^{-1}\mathrm{d}t
$$

$$
\leq 4d\beta \log n,
$$

provided that $n \geq 1$ is sufficiently large. In particular, the expected number of toroidal chains consisting of m steps, and where all radii are contained in $(t_0, n^{3/4})$ is at most

$$
n^d(4d\kappa_d\beta \log n)^m/m!,
$$

and Stirling's formula implies that when putting $m = c \log n$ this expression is at most

$$
\exp(d \log n + c(\log n)(\log(4de\kappa_d\beta \log n) - \log(c \log n))).
$$

Finally, the latter expression tends to 0 as $n \to \infty$, if $c > 0$ is sufficiently large. $\qquad \square$

Combining Lemmas 18 and 19 the proof of Theorem 8 is now immediate.

Proof of Theorem 8. Lemma 18 shows that if $x, y \in X^{(n)}$ are connected by an edge in $G(X^{(n)})$ and the chemical distance between x and y in $G'(X^{(n)})$ is at least $c_1 \log n$, then $X^{(n)}$ contains a toroidal descending chain consisting of at least $c_1 \log n$ elements. Lemma 19 shows that the complements of the latter events occur whp. $\square$

Acknowledgments

References

[1] D. J. Aldous. Which connected spatial networks on random points have linear route-lengths? *Arxiv preprint arXiv:0911.5296*, 2009.

[2] S. Appleby. Estimating the cost of a telecommunications network using the fractal structure of the human population distribution. *IEE Proceedings-Communications*, 142:172–178, 1995.

[3] S. Appleby. Multifractal characterization of the distribution pattern of the human population. *Geographical Analysis*, 28:147–160, 1996.

[4] F. Baccelli and S. Zuyev. Poisson-Voronoi spanning trees with applications to the optimization of communication networks. *Operations Research*, 47:619–631, 1999.

[5] M. Biskup. On the scaling of the chemical distance in long-range percolation models. *The Annals of Probability*, 32:2938–2977, 2004.

[6] M. Biskup. Graph diameter in long-range percolation. *Random Structures & Algorithms*, 39:210–227, 2011.

[7] D. Coppersmith, D. Gamarnik, and M. Sviridenko. The diameter of a long-range percolation graph. *Random Structures & Algorithms*, 21:1–13, 2002.

[8] D. J. Daley and G. Last. Descending chains, the lilypond model, and mutual-nearest-neighbour matching. *Advances in Applied Probability*, 37:604–628, 2005.

[9] M. Deijfen, R. van der Hofstad, and G. Hooghiemstra. Scale-free percolation. *Annales de l'Institut Henri Poincaré Probabilités et Statistiques*, 49:817–838, 2013.

[10] P. Deprez and M. V. Wüthrich. Networks, random graphs and percolation. *Arxiv preprint arXiv:1409.2339*, 2014.

[11] P. Deprez and M. V. Wüthrich. Poisson heterogeneous random-connection model. *Arxiv preprint arXiv:1312.1948*, 2013.

[12] K. J. Falconer and G. R. Grimmett. On the geometry of random Cantor sets and fractal percolation. *Journal of Theoretical Probability*, 5:465–485, 1992.

[13] T. E. Harris. *The Theory of Branching Processes*. Springer, Berlin, 1963.

[14] G. Last and M. D. Penrose. Poisson process Fock space representation, chaos expansion and covariance inequalities. *Probability Theory and Related Fields*, 150:663–690, 2011.

[15] G. F. Lawler. A self-avoiding random walk. *Duke Mathematical Journal*, 47:655–693, 1980.

[16] T. M. Liggett, R. H. Schonmann, and A. M. Stacey. Domination by product measures. *Annals of Probability*, 25:71–95, 1997.

[17] P. Mannersalo, A. Koski, and I. Norros. Telecommunication networks and multifractal analysis of human population distribution. *Technical Report 02, COST-257*, pages 1–15, 1998.

[18] M. V. Menshikov, S. Yu. Popov, and M. Vachkovskaia. On the connectivity properties of the complementary set in fractal percolation models. *Probability Theory and Related Fields*, 119:176–186, 2001.

[19] M. D. Penrose. *Random Geometric Graphs*. Oxford University Press, Oxford, 2003.

[20] K. Tchoumatchenko and S. Zuyev. Aggregate and fractal tessellations. *Probability Theory and Related Fields*, 121:198–218, 2001.

[21] J. E. Yukich. Ultra-small scale-free geometric networks. *Journal of Applied Probability*, 43:665–677, 2006.