Citation (APA)
Komjáthy, J., Lapinskas, J., Lengler, J., & Schaller, U. (2026). Four universal growth regimes in degree-dependent first passage percolation on spatial random graphs. *Forum of Mathematics, Sigma*, *14*.
https://doi.org/10.1017/fms.2025.10161

Forum of Mathematics, Sigma (2026), **Vol. 14**:e22 1–79
doi:10.1017/fms.2025.10161

CAMBRIDGE UNIVERSITY PRESS

RESEARCH ARTICLE

Four universal growth regimes in degree-dependent first passage percolation on spatial random graphs

Júlia Komjáthy[1], John Lapinskas[2], Johannes Lengler[3] and Ulysse Schaller[4]

[1]Delft University of Technology, Netherlands; E-mail: j.komjathy@tudelft.nl (Corresponding author).
[2]University of Bristol, United Kingdom; E-mail: john.lapinskas@bristol.ac.uk.
[3]ETH Zurich, Switzerland; E-mail: johannes.lengler@inf.ethz.ch.
[4]ETH Zurich, Switzerland; E-mail: ulysse.schaller@inf.ethz.ch.

Received: 22 March 2024; **Revised:** 11 November 2025; **Accepted:** 11 November 2025

2020 Mathematics Subject Classification: *Primary* – 05C82, 60K35, 60K50, 82B43, 91D30, 91D25

Abstract
One-dependent first passage percolation is a spreading process on a graph where the transmission time through each edge depends on the direct surroundings of the edge. In particular, the classical i.i.d. transmission time L_{xy} is multiplied by $(W_x W_y)^\mu$, a polynomial of the expected degrees W_x, W_y of the endpoints of the edge xy, which we call the penalty function. Beyond the Markov case, we also allow any distribution for L_{xy} with regularly varying distribution near 0. We then run this process on three spatial scale-free random graph models: finite and infinite Geometric Inhomogeneous Random Graphs, including Hyperbolic Random Graphs, and Scale-Free Percolation. In these spatial models, the connection probability between two vertices depends on their spatial distance and on their expected degrees.

We show that as the penalty function, that is, μ increases, the transmission time between two far away vertices sweeps through four universal phases: *explosive* (with tight transmission times), *polylogarithmic*, *polynomial* but strictly sublinear, and *linear* in the Euclidean distance. The strictly polynomial growth phase is a new phenomenon that so far was extremely rare in spatial graph models. All four growth phases are robust in the model parameters and are not restricted to phase boundaries. Further, the transition points between the phases depend nontrivially on the main model parameters: the tail of the degree distribution, a long-range parameter governing the presence of long edges, and the behaviour of the distribution L near 0. In this paper we develop new methods to prove the upper bounds in all sub-explosive phases. Our companion paper complements these results by providing matching lower bounds in the polynomial and linear regimes.

Contents

1. Introduction

First passage percolation (FPP) is a natural way to understand geodesics in random metric spaces. Starting from some initial vertex at time 0, the process spreads through the underlying graph so that the transmission time between any two vertices x, y is the minimum sum of edge transmission times over all paths between x and y. In classical FPP, edge transmission times are independent and identically distributed random variables. In the recent paper [55] we introduced one-dependent FPP, where edge transmission times depend on the edge's direct surroundings in the underlying graph. There, we determined the phase transition for explosion (i.e., reaching infinitely many vertices in finite time). In this paper we study the sub-explosive regime, when explosion does not occur. We show that the process exhibits rich behaviour with several growth phases and nonsmooth phase transitions between them. This holds across a large class of scale-free spatial random graph models (namely Scale-Free Percolation (SFP), Hyperbolic Random Graphs (HypRG), and infinite and finite Geometric Inhomogeneous Random Graphs (GIRG) [26, 17, 58]), and across all Markovian and non-Markovian transmission time distributions with reasonable limiting behaviour at zero.

In SFP, the vertex set is formed by the d-dimensional lattice $\mathbb{Z}^d$. Each vertex u is then equipped with an independent and identically distributed random vertex-weight $W_u \geq 1$. Given the weighted vertex set, the edges are drawn conditionally independently. The probability of an edge between vertices u, v with weights W_u, W_v decreases with the Euclidean distance $|u - v|$ and increases with the vertex-weights, and is between constant factors of $\min(1, W_u W_v / |u-v|^d)^\alpha$, see Definition 1.3 below for full detail. The parameter α is often called the *long-range parameter*, since the model with all vertex-weights set to 1 recovers the classical long-range percolation model [72]. Instead of unit vertex-weights, here we shall rather assume that the vertex-weight distribution W follows a regularly varying tail, that is, for some $\tau \in (2, 3)$ that is called the *power-law exponent* we assume

$$\frac{\mathbb{P}(W \geq x)}{\mathbb{P}(W \geq cx)} \to c^{-(\tau-1)} \qquad \text{for all } c > 0 \text{ as } x \to \infty. \tag{1.1}$$

The heavy-tailed decay of W creates degree-inhomogeneity in the model: the vertex weight W_v of v is (up to constant factors) equal to the expected degree of v, and the degree of a high-weight vertex is concentrated around its expectation [26, Proposition 2.3]. The parameters τ, α play different roles in governing inhomogeneities in the models: while τ governs the degree distribution, a smaller α causes a heavier tail on the edge-length distribution, with $\alpha > 1$ needed for a.s. finite degrees [26]. The model GIRG follows the same construction by replacing the location of vertices by a unit-intensity Poisson point process (PPP) on $\mathbb{R}^d$. For the overview of the results, the reader may ignore this difference.

Universality classes of transmission times. In one-dependent first passage percolation, we set the transmission time through the edge $e = xy$ between vertices x, y as the product of an independent and identically distributed (i.i.d.) random factor L_{xy} and a factor depending on the weights of vertices:

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

Table 1. *Summary and brief description of the main parameters of the model.*

Parameter	Range	Eqn.	Description
τ	$(2, 3)$	(1.1)	Power-law exponent of the weight (and degree) distribution(s) of the underlying graph model. All our results focus on the infinite-variance range $2 < \tau < 3$.
α	$(1, \infty)$	(1.5)	Long-range parameter representing the influence of the latent geometry in the underlying graph model. A larger α yields fewer edges between distant vertices.
d	$\mathbb{N}$	(1.5)	Dimension of the geometric space in which the underlying graph model is embedded.
μ	$(0, \infty)$	(1.2)	Penalty strength on the edge costs. A larger μ yields longer transmission time on edges that are incident to vertices of high weight/degree.
β	$(0, \infty)$	(1.3)	Power-exponent of the random component of the edge costs around 0. The commonly used exponential distribution (with any mean) satisfies $\beta = 1$.

Definition 1.1 (1-dependent first passage percolation (1-FPP)). Consider a graph $G = (\mathcal{V}, \mathcal{E})$ where each vertex $v \in \mathcal{V}$ has an associated vertex-weight W_v. For every edge $xy \in \mathcal{E}$, draw an i.i.d. copy L_{xy} of a random variable L, and set the *(transmission) cost* of an edge xy as

$$\mathcal{C}(xy) := L_{xy}(W_x W_y)^{\mu}, \tag{1.2}$$

for a fixed parameter $\mu > 0$ called the *penalty strength*. The costs define a *cost distance* $d_{\mathcal{C}}(x, y)$ between any two vertices x and y, which is the minimal total cost of any path between x and y (see Section 1.4.1). We call $d_{\mathcal{C}}$ the 1-dependent first passage percolation.

We usually assume that the cumulative distribution function (cdf) $F_L : [0, \infty) \to [0, 1]$ of L satisfies the following assumption (with exceptions of this assumption explicitly mentioned):

Assumption 1.2. There exist constants $t_0,\ c_1,\ c_2,\ \beta > 0$ such that

$$c_1 t^{\beta} \le F_L(t) \le c_2 t^{\beta} \text{ for all } t \in [0, t_0]. \tag{1.3}$$

Without much effort, one can relax Assumption 1.2 to $\lim_{x\to 0} \log F_L(x)/\log x = \beta$, that is, regularly varying behaviour of F_L near 0. We work with (1.3) for the sake of readability. Table 1 provides an overview of the various parameters of the model.

The cost distance $d_{\mathcal{C}}(x, y)$ corresponds to the transmission time between two vertices x, y. In SFP, we use the *same* vertex weights W_x, W_y to generate the edge between x, y as well as to define the edge-cost $\mathcal{C}(xy)$. This leads to the cost of the edge to depend essentially on the expected degrees of the two involved vertices[1], however, it also leads to three layers of randomness. On the first layer, the vertex set has random vertex-weights; on the second layer, edges are drawn randomly using the randomness in the first layer, and finally, on the third layer, edge-costs depend on the presence of edges, on the vertex-weights, and on an extra source of randomness captured in L_{xy}.

When $\mu \in (0, 1)$, a high-degree vertex still causes more new infections per unit time than a low-degree vertex, but this effect is sublinear in the degree. As μ increases and/or the parameters of the underlying graph change, we prove that the following four different phases occur for the transmission time between the vertex at 0 and a far away vertex x, see Table 2 for the thresholds between the different phases:

(i) $d_{\mathcal{C}}(0, x)$ converges to a limiting distribution that is independent of the Euclidean distance $|x|$ (*explosive phase*);

This was the main result of [55]. The main result of this paper is to characterise the other phases:

(ii) $d_{\mathcal{C}}(0, x)$ grows at most *polylogarithmically* in the Euclidean distance $|x|$, without being explosive;
(iii) $d_{\mathcal{C}}(0, x)$ grows *polynomially* in $|x|$, with exponent $0 < \eta_0 < 1$;
(iv) $d_{\mathcal{C}}(0, x)$ grows *linearly* in $|x|$, that is, with exponent $\eta_0 = 1$.

[1]Using W_x instead of the actual degree of x is natural in these models. We are convinced that the same results would also hold if we took the actual degrees instead of their expectation. However, that would make the proofs more technical without giving much additional insight.

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

Table 2. Summary of our main results. In 1-FPP, edge transmission times are $L_{xy}(W_xW_y)^\mu$ where W_x, W_y are constant multiples of the expected degrees of the vertices x, y, and L_{xy} is i.i.d. with distribution function that varies regularly near 0 with exponent $\beta \in (0, \infty]$. The degree distribution follows a power law with exponent $\tau \in (2, 3)$: graph distances are doubly logarithmic in the underlying graph. The transmission time $d_{\mathcal{C}}(0, x)$ between 0 and a far away vertex x sweeps through four different phases as the penalty exponent μ increases. For long-range parameter $\alpha \in (1, 2)$, long edges between low-degree vertices maintain polylogarithmic transmission times (similar to long-range percolation), so increasing μ stops explosion but it has no further effect. When $\alpha > 2$, these edges are sparser and a larger μ slows down 1-FPP, to polynomial but sublinear transmission times in an interval of length at least $1/d$ for μ. Then, all long edges have polynomial transmission times in the distance they bridge. For even higher penalty exponent μ the behaviour becomes similar to FPP on the grid $\mathbb{Z}^d$. We give the growth exponents Δ_0 and η_0 explicitly in (1.9) and (1.10).

Graph param.	1-FPP parameters	Behaviour of 1-FPP transmission times
Weak decay: $\tau \in (2,3)\, \alpha \in (1,2)$	$\mu < \frac{3-\tau}{2\beta}$	**Explosive**: $d_{\mathcal{C}}(0,x) = \Theta(1)$
	$\mu > \frac{3-\tau}{2\beta}$	**Polylogarithmic**: $d_{\mathcal{C}}(0,x) = O((\log\lvert x\rvert)^{\Delta_0+o(1)}), \Delta_0 > 1$
Strong decay: $\tau \in (2,3)\, \alpha > 2$	$\mu < \frac{3-\tau}{2\beta}$	**Explosive**: $d_{\mathcal{C}}(0,x) = \Theta(1)$
	$\mu \in \left(\frac{3-\tau}{2\beta}, \frac{3-\tau}{\beta}\right)$	**Polylogarithmic**: $d_{\mathcal{C}}(0,x) = O((\log\lvert x\rvert)^{\Delta_0+o(1)}), \Delta_0 > 1$
	$\mu \in \left(\frac{3-\tau}{\beta}, \frac{3-\tau}{\min\{\beta, d(\alpha-2)\}} + \frac{1}{d}\right)$	**Polynomial**: $d_{\mathcal{C}}(0,x) = \lvert x\rvert^{\eta_0 \pm o(1)}, \eta_0 < 1$
	$\mu > \frac{3-\tau}{\min\{\beta, d(\alpha-2)\}} + \frac{1}{d}$	**Linear:** $d_{\mathcal{C}}(0,x) = \Theta(\lvert x\rvert)$

These phases are *highly robust* in the parameters, they are not restricted to phase boundaries in either μ or the other model parameters. Moreover, all four phases can occur on a single underlying graph by changing the penalty exponent μ only; universally across distributions of L_{xy} with regularly varying behaviour at 0, see Figure 1 for a visualisation. This rich behaviour arises despite the doubly logarithmic graph distances in the underlying spatial graph models. By contrast, in other models the behaviour of transmission times in classical FPP is less rich, see Section 1.1 for the discussion.

Precise behaviour in the four phases. In this paper we prove the upper bounds on transmission times in the sub-explosive regime (phase (i) was previous work [55]). In phase (ii), we show that the transmission time is at most $(\log\lvert x\rvert)^{\Delta_0+o(1)}$ with an explicit $\Delta_0 > 1$ which we conjecture to be tight. In phases (iii) and (iv), we show that the transmission time is precisely $\lvert x\rvert^{\eta_0\pm o(1)}$, where we give $\eta_0 < 1$ explicitly for phase (iii) and $\eta_0 = 1$ for phase (iv). The companion paper [56] contains the matching lower bounds for phases (iii)-(iv) as well as some additional results for phase (iv). We develop new techniques that allow us to treat upper bounds for all three sub-explosive phases *simultaneously*, which we expect to be of independent interest.

Motivation of the process from applications. One-dependent processes in general, and one-dependent FPP in particular, allow for more realistic modelling of real phenomena. In social networks, actual contacts and infections do not scale linearly with the degree [31, 59, 76, 54]. 1-FPP type penalisation has frequently been used to model the sublinear impact of superspreaders [37, 53, 65, 69, 77, 6], and in other contexts [16, 30, 62, 78, 3, 49]. Consistent with our model, all these applications assume a polynomial dependence with exponent in the range $\mu \in (0, 1)$, where a high-degree vertex may cause more new infections per time than a low-degree vertex, but this effect is sublinear in the degree.

While our paper is theoretical, we do believe that a model with a rich phase space can have practical implications. In the spread of physical epidemics, while some diseases spread at an exponential rate, others spread at a polynomial rate, dominated by the local geometry. Examples of the latter include HIV/AIDS, Ebola, and foot-and-mouth disease, see the survey [21] on polynomial epidemic growth. Classical epidemic models can typically only model either exponential or polynomial growth, not both.

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

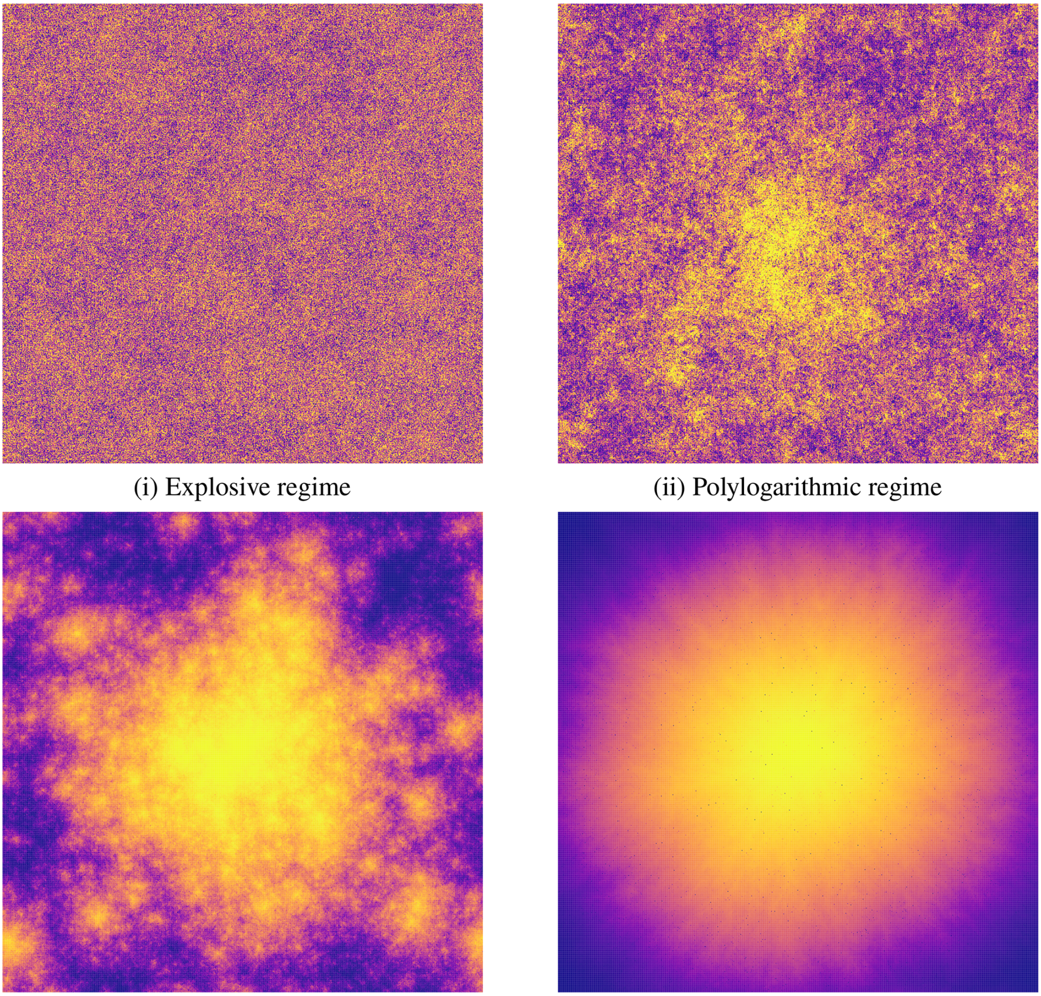

Figure 1. *Heatmaps for the four different universality classes. The vertices are sorted by their transmission times from the origin (centre vertex). The colours represent this ordering: yellow infected first, then orange, then purple. All four plots are generated on the same underlying graph (with parameters $\tau = 2.3$ and $\alpha = 5$, and edge connection probabilities $p(u, v) = (w_u w_v/(\mathbb{E}[W]\|u - v\|^2))^5 \wedge 1$), where the vertices are placed on a 750×750 grid in the 2-dimensional torus. The random factors L_{xy} associated to each edge are also identical in all four plots, and follow an exponential distribution (i.e., $\beta = 1$). The only varying parameter is the penalty exponent μ, taking values (i) $\mu = 0$ for the explosive regime, (ii) $\mu = 0.5$ for the the polylogarithmic regime, (iii) $\mu = 1$ for the polynomial regime (iv) $\mu = 2$ for the linear regime. In the linear regime, the late points are – typically – high degree vertices carrying high penalisation. We thank Zylan Benjert for generating the simulations and the pictures.*

Arguably, 1-FPP provides a natural explanation, since in 1-FPP the transition can be driven by changes only to the transmission dynamics, not to the underlying network.

New methodology: moving to quenched vertex-set to replace FKG-inequality. In this paper we develop a general technique – *nets combined with multiround exposure* – that *replaces the FKG-inequality* [32] in problems concerning vertex and/or edge-weighted graph models where this inequality does not hold. Let us explain why the FKG-inequality fails in the context of 1-FPP. Typically, for upper bounds one

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

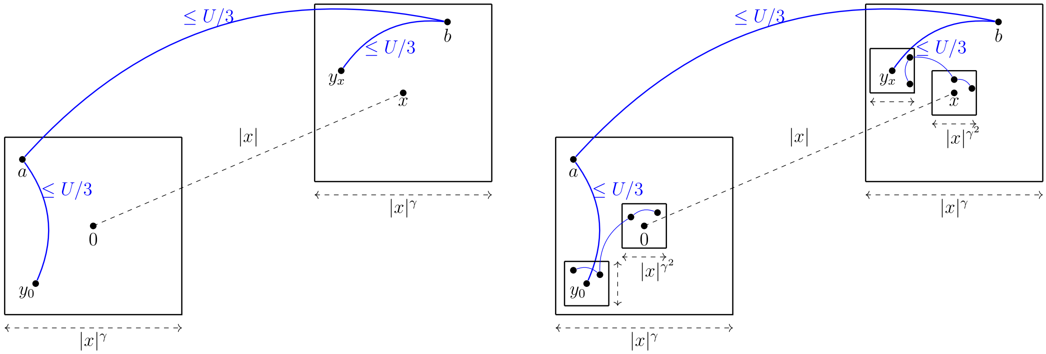

(a) (b)

Figure 2. *The budget travel plan with 3-edge bridging-paths: (a) first and (b) second iteration.*

constructs paths connecting 0 and x by revealing vertices and/or edges of the graph sequentially, which destroys the independence of edges. For graph distances in long-range percolation, the FKG-inequality resolves this problem [11], but it already needs adjustments once vertex-weights are present [41]. In 1-FPP, the existence of a long edge is positively correlated to its endpoints having large vertex weights, which is *negatively* correlated to its other outgoing edges having short transmission times. Hence, having chosen a long edge with low-cost, we lose probabilistic control over how to choose the next low-cost edge on the path connecting 0 and x. To overcome this issue, we move to the *(weighted-vertex) quenched setting* where we reveal the realisation of the whole weighted vertex set – say $(\mathcal{V}, \mathcal{W}) = (V, w_V)$ – and thus events concerning only edges become independent. We show that in a large box centred at the origin of $\mathbb{R}^d$, the proportion of realisations with behaviour 'close to what is expected' tends to 1 with the box-size. More precisely, we require that locally around a constant proportion of the vertices and uniformly across multiple scales of vertex-weights, the number of points in the weighted vertex set is close to its expectation. For this we select a subset of the vertices that we call a *net* realising this property. A net $\mathcal{N}$ is thus a subset of the weighted vertex set, such that for every not-too-small radius r, every 'reasonable' weight w, and every selected vertex $v \in \mathcal{N}$, the net has *constant density* in $B_r(v) \times [w, 2w]$, shorthand for vertices of weight in $[w, 2w]$ within Euclidean distance r of v:

$$\frac{|\mathcal{N} \cap B_r(v) \times [w, 2w]|}{\mathbb{E}\big[|\mathcal{V} \cap B_r(v) \times [w, 2w]| \mid v \in \mathcal{V}\big]} \in \Big(\frac{1}{16}, 16\Big), \tag{1.4}$$

where the expectation is taken over the randomness in the weights and location of vertices[2]. We prove via a *multiscale analysis* that as the box-size tends to infinity, asymptotically almost every realisation of the weighted vertex set contains a net $\mathcal{N}$ with total density at least $1/4$.

When we move to the quenched setting we only reveal the realisation of the weighted vertex set, but not the edges of the graph. In realisations containing a net, with a carefully chosen *multiround exposure* process we can define a coupling of the edges and their costs which lets us replace the FKG inequality needed for the construction of a low-cost path between 0 and x, see Section 3 for more details. We believe that this method is also useful for many other graph models, so we explain it streamlined now.

Budget travel plan with 3-edge bridge-paths. Switching to the quenched setting allows to prove the upper bounds in all subexponential phases (ii)–(iv) all-at-once. Our construction of a connecting path overcomes the following problem: A long edge with a short transmission time typically occurs on typical high-degree vertices and thus all other outgoing edges from the same vertices have too long transmission times. The main idea resembles a 'budget travel plan': when someone travels with a low

[2]In case of GIRG, one uses here the Palm measure by conditioning that the PPP has a point at $v \in \mathbb{R}^d$, and one takes expectation over the Poisson point process and the vertex-weights, see also the discussion before Theorem 1.4. For SFP the conditioning can be dropped, and the expectation is only over the vertex-weights, as the vertex set is deterministically $\mathbb{Z}^d$.

budget, one takes the cheapest mode of transport to the airport within a 100km radius that offers the cheapest flight landing within a 100km radius of the destination, then takes the cheapest transport to the destination city.

Formally, we put balls of radius $|x|^\gamma$ for some $\gamma \in (0,1)$ around 0 and around x, and we find a cheap 3-edge path ('bridge') $\pi_1 = y_0 a b y_x$ between these two balls *using only vertices in the net*. The net guarantees enough vertices in each vertex-weight range of interest. We find atypical high-weight vertices a, b that are connected by an atypically cheap edge, that simultaneously have an atypically cheap edge to low-weight vertices y_0, y_x, respectively. (Here we use the common terminology of fast transmission corresponding to 'cheap' cost.) Then we have replaced the task of connecting 0 and x by the two tasks of connecting 0 with y_0 and x with y_x, where the new 'gaps' $|0 - y_0|$ and $|x - y_x|$ are much smaller than $|x|$. The *multiround exposure* and the *net* on the fixed vertex set together guarantee that we can iterate this process without running out vertices in the relevant weight-ranges, and without accumulated correlations in the presence of edges along the iteration (e.g., out of y_0, y_x). Iteration yields a set of multiscale bridge-paths, which we call after Biskup a *hierarchy* [11]. The construction in [11] also uses recursion, with one-edge bridges instead of three-edge bridges, and yields polylogarithmic graph distances in long-range percolation. The techniques in [11] would not work for 1-FPP because we need to balance distances vs costs vs the penalisation on high-weight vertices in very different regimes, and at the same time deal with edge-costs dependencies. Those can only be dealt with in the quenched setting.

The cost (transmission time) of the bridge-paths π in 1-FPP are either polynomial in the distance they bridge or constant. When the cost is *polynomial* – with optimal exponent η_0 – we are in the *polynomial phase*. The cost of the first bridge π_1 then dominates the cost of the whole path, and we only carry out a constant number of iterations (irrespective of $|x|$). When bridge-paths with constant cost exist, we are in the *polylogarithmic phase*. Then, the cost of all bridges together is negligible compared to the cost of the polylogarithmic number of gaps that remain after the last iteration. Here, we iterate until we can connect the remaining gaps via essentially constant cost paths. Connecting the gaps is a nontrivial task itself since the graphs do not contain 'nearest-neighbour' edges. Solutions for filling gaps in [11] do not work in our setting due to the presence of vertex weights. Instead, we connect the gaps with *'weight-increasing paths'* that crucially use that the underlying graphs are scale-free. We give a more detailed discussion about the hierarchical construction at the beginning of Section 5 and back-of-the-envelope calculations about how to obtain the precise growth exponents in phases (ii) and (iii) at the beginning of Section 5.1 with proof sketches below Corollaries 5.2 and 5.3.

Robustness of our techniques. The technique of nets combined with multiround edge-exposure is robust, and will be applicable elsewhere, for questions concerning *first passage percolation, robustness to percolation (random deletion of edges), graph distances, SIR-type and other epidemic processes, rumour spreading*, etc. on a larger class of vertex-weighted graphs; including random geometric graphs, Boolean models with random radii, the age-dependent and the weight-dependent random connection model (mimicking spatial preferential attachment), scale-free Gilbert graph, and the models used here [2, 23, 38, 39, 40, 41, 42, 47, 50], and can also be extended to dynamical versions of the above graph models on fixed vertex sets.

Two papers, two techniques and optimality. The 'budget travel plan' together with the renormalisation group argument in [56] reveals that the strategy of polynomial paths is essentially optimal: in this phase, all long edges have polynomial transmission time in the distance they bridge. Our techniques for the lower bounds are entirely different and deserve their own exposition, hence we present them in the companion paper [56].

1.1. Related work: phases of FPP in other models.

The phase diagrams of transmission times in classical FPP are less rich. In particular, the strict polynomial phase is absent or restricted only to phase transition boundaries. Indeed, on sparse *nonspatial* graph models with finite-variance degrees, both Markovian and non-Markovian classical FPP universally show Malthusian (exponential) growth [9]. Transmission times between two uniformly chosen

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

Table 3. *Known results about the universality classes of graph-distances on long-range percolation **LRP**, scale-free percolation **SFP**, long-range first-passage percolation **LRFPP** and infinite geometric inhomogeneous random graphs **IGIRG**. The results highlighted in yellow color follow (also) from techniques in this paper.* ‡*An upper bound is only known for high enough edge-density or all nearest-neighbour edges present.*

SFP/LRP with	**Graph-distance**	**Growth**	**Upper bound**	**Lower bound**
$\tau \in (2,3)$	$\frac{(2\pm o(1))\log\log\lvert x\rvert}{\lvert\log(\tau-2)\rvert}$	doubly-logarithmic	[26, 75]	[75]
$\tau > 3$ and $\alpha \in (1,2)$	$(\log\lvert x\rvert)^{\Delta\pm o(1)}$ for some $\Delta > 0$	poly-logarithmic	SFP: [46, 60], LRP: [11, 74]	SFP: [60] LRP:[11, 12, 74]
$\tau > 3$ and $\alpha = 2$	$\lvert x\rvert^{\eta\pm o(1)}$, for some $\eta < 1$	polynomial	SFP: open, LRP: [5]	SFP: open LRP: [5]
$\tau > 3$ and $\alpha > 2$	$\Theta(\lvert x\rvert)$	linear	partly open‡ [4]	[8, 27]
LRFPP with	**Cost-distance**	**Growth**	**Upper bound**	**Lower bound**
$\alpha' < 1$	0	instantaneous		[20]
$\alpha' \in (1,2)$	$(\log\lvert x\rvert)^{\Delta_{\alpha'}\pm o(1)}$ for $\Delta_{\alpha'} = 1/(1-\log_2\alpha')$	poly-logarithmic		[20]
$\alpha' \in (2, 2+1/d)$	$\lvert x\rvert^{d(\alpha'-2)\pm o(1)}$	polynomial		[20]
$\alpha' > 2+1/d$	$\Theta(\lvert x\rvert)$	linear		[20]
IGIRG/SFP with	**Cost-distance**	**Growth**	**Upper bound**	**Lower bound**
$\tau \in (2,3)$ $\mu < \mu_{\mathrm{expl}}$	converges in distribution	explosion	[55]	[55]
$\mu \in (\mu_{\mathrm{expl}}, \mu_{\log})$ or $\alpha \in (1,2)$	$(\log\lvert x\rvert)^{\Delta_0+o(1)}$, Δ_0 as in (1.9)	poly-logarithmic	Theorem 1.4	open
$\mu \in (\mu_{\log}, \mu_{\mathrm{pol}})$ and $\alpha > 2$	$\lvert x\rvert^{\eta_0\pm o(1)}$, η_0 as in (1.10)	polynomial	Theorem 1.6	[56]
$\mu > \mu_{\mathrm{pol}}$ and $\alpha > 2$	$\Theta(\lvert x\rvert)$ $\Theta(\lvert x\rvert^{1+o(1)})$	linear	[56] for $d \geq 2$, Theorem 1.6 for $d \geq 1$	[56]

vertices are then *logarithmic* in the graph size. Sparse *spatial* graphs with finite-variance degrees (e.g., percolation, long-range percolation, random geometric graphs etc.) are typically restricted to linear graph distances/transmission times in the absence of long edges [4, 68, 25], or to polylogarithmic distances in the presence of long edges [11, 12, 46]. In both spatial and nonspatial graph models with infinite-variance degrees, classical FPP typically either explodes or exhibits a smooth transition between explosion and doubly logarithmic transmission times (which match the graph distances) [1, 51, 75]; in particular, there is no analogue of phases (ii)–(iv). For one-dependent FPP on nonspatial graphs there are strong indications that the process either explodes [73], with the same criterion for explosion as for spatial graphs in [55], or becomes Malthusian [34], the latter implying logarithmic transmission times between two uniformly chosen vertices by the universality in [9], so only two phases can occur. The only graph model to exhibit a transition from a fast-growing phase to a slow-growing phase is long-range percolation, where the polynomial phase is restricted to the phase boundary in the long-range parameter $\alpha = 2$ [5]. Even in degenerate models (where the underlying graph is complete), long-range first passage percolation [20] is the only other model where a similarly rich set of phases is known to occur. Thus one-dependent FPP is the first process that displays a full interpolation between the four phases on a *single nondegenerate graph model*. Moreover, the phase boundaries for one-dependent FPP depend nontrivially on the main model parameters: the degree power-law exponent τ, the parameter α controlling the prevalence of long-range edges, and the behaviour of L_{xy} near 0 characterised by β, see Table 2 for our results, Table 3 for phases of growth in other models, and Section 1.4 for more details on related work.

1.2. Graph Models

We consider simple and undirected graphs with vertex set $\mathcal{V} \subseteq \mathbb{R}^d$. We use standard graph notation along with other common terminology, see Section 1.4.1. We consider three random graph models: *Scale-Free Percolation* (SFP), *Infinite Geometric Inhomogeneous Random Graphs* (IGIRG)[3], and (finite) *Geometric Inhomogeneous Random Graphs* (GIRG). The latter model contains *Hyperbolic Random Graphs* (HypRG) as special case, so our results extend to HypRG, see the paragraph below Theorem 1.11. The main difference between SFP and IGIRG is the vertex set $\mathcal{V}$. For SFP, we use $\mathcal{V} := \mathbb{Z}^d$, with $d \in \mathbb{N}$. For IGIRG, a unit-intensity Poisson point process on $\mathbb{R}^d$ forms $\mathcal{V}$.

Definition 1.3 (SFP, IGIRG, GIRG). Let $d \in \mathbb{N}$, $\tau > 2$, $\alpha \in (1, \infty)$, and $\overline{c} > \underline{c} > 0$. Let $\ell : [1, \infty) \to (0, \infty)$ be function that varies slowly at infinity (see Section 1.4.1), and let $h : \mathbb{R}^d \times [1, \infty) \times [1, \infty) \to [0, 1]$ be a function satisfying

$$\underline{c} \cdot \min\left\{1, \frac{w_1 w_2}{|x|^d}\right\}^{\alpha} \le h(x, w_1, w_2) \le \overline{c} \cdot \min\left\{1, \frac{w_1 w_2}{|x|^d}\right\}^{\alpha}. \tag{1.5}$$

The vertex set and vertex-weights: For SFP, set $\mathcal{V} := \mathbb{Z}^d$, for IGIRG, let $\mathcal{V}$ be given by a Poisson point process on $\mathbb{R}^d$ of intensity one.[4] For each $v \in \mathcal{V}$, we draw a *weight* W_v independently from a probability distribution on $[1, \infty)$ satisfying

$$F_W(w) = \mathbb{P}(W \le w) = 1 - \ell(w)/w^{\tau-1}. \tag{1.6}$$

We denote $\widetilde{\mathcal{V}}(G) := (\mathcal{V}, \mathcal{W})$ the vertex set $\mathcal{V}$ together with the random weight vector $\mathcal{W}_{\mathcal{V}} := (W_v)_{v\in\mathcal{V}}$, and $(V, w_V) := (V, (w_v)_{v\in V})$ a realisation of $\widetilde{\mathcal{V}} := \widetilde{\mathcal{V}}(G)$, where $\tilde{v} := (v, w_v)$ stands for a single weighted vertex.

The edge set: Conditioned on $\widetilde{\mathcal{V}} = (V, w_V)$, consider all unordered pairs $\mathcal{V}^{(2)}$ of $\mathcal{V}$. Then every pair $xy \in \mathcal{V}^{(2)}$ is present in $\mathcal{E}(G)$ independently with probability $h(x - y, w_x, w_y)$.

Finally, a GIRG G_n is obtained as the induced subgraph $G[Q_n]$ of an IGIRG G by the set of vertices in the cube Q_n of volume n centred at 0. We call h the *connection probability*, d the *dimension*, τ the *power-law exponent*, and α the *long-range parameter*.

The above definition essentially merges the Euclidean space and the vertex-weight space by considering vertices with weights as points in $\mathbb{R}^d \times [1, \infty)$, that is, we think of each vertex as a pair $\tilde{v} = (v, w_v)$, where $v \in \mathbb{R}^d$ is its spatial location and w_v is its weight. While SFP is a somewhat simpler model due to the deterministic location of vertices, GIRGs gained significant attention in both applications and theoretical studies [13, 14, 48, 63, 64, 67], and are part of a larger class of marked random connection models [19, 42, 40]. Definition 1.3 leads to a slightly less general model than those, for example, in [17] and [55]. The original definition in [17] had a different scaling of the geometric space vs connection probabilities and (also) considered the torus topology on the unit cube, identifying 'left' and 'right' boundaries. However, the resulting finite graphs are identical in distribution after rescaling, and the torus topology vs Euclidean topology does not make a difference for the results below on cost-distances, see [55] for a comparison. We discuss extensions to $\alpha = \infty$ and $\beta = \infty$ separately in Section 1.3.1. We call the set of parameters $\texttt{par} := \{d, \tau, \alpha, \mu, \beta, \underline{c}, \overline{c}, c_1, c_2, t_0\}$ the *model parameters*. We say that a variable is *large* (or *small*) relative to a collection of other variables when it is bounded below (or above) by some finite positive function of those variables and the model parameters. We restrict to $\tau \in (2, 3)$, (explicitly stated in the theorems), which ensures that there is a unique infinite component (or linear-sized 'giant' component for finite GIRG)[5] and that graph distances between vertices x, y in the infinite/giant component grow like $d_G(x, y) \sim 2 \log\log |x - y|/|\log(\tau - 2)|$ in all three models [57, 18, 26, 75]. We

[3]They have also been called EGIRG, where E stands for extended [57].

[4]If we take an IGIRG and rescale the underlying space $\mathbb{R}^d$ by a factor λ, then we obtain a random graph which satisfies all conditions of IGIRGs except that the density of the Poisson point process is λ^{-d} instead of one. Thus it is no restriction to assume density one.

[5]For $\tau > 3$, an infinite component only exists for high enough edge density, which is captured by h in (1.5).

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

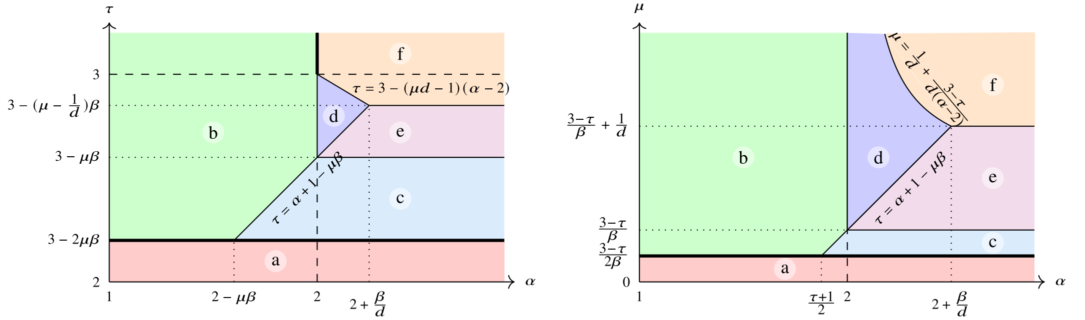

(i) (α, τ)-phase diagram of transmission times. Here, $\mu = 0.4, \beta = 1, d = 4$ are kept fixed.

(ii) (α, μ)-phase diagram of transmission times. Here, $\tau = 2.75, \beta = 1, d = 2$ are kept fixed.

Figure 3. *Phase diagrams of transmission times in one-dependent first passage percolation. On both diagrams, parameter choices falling in area (a) yield explosive spread. Parameter choices in areas (b) and (c) yield polylogarithmic transmission times* $d_{\mathcal{C}}(0,x) \leq (\log \|x\|)^{\Delta_0+o(1)}$*, where* $\Delta_0 = \Delta_\alpha = 1/(1-\log_2 \alpha)$ *on (b) and* $\Delta_0 = \Delta_\beta = 1/(1-\log_2(\tau-1-\mu\beta))$ *on (c). Parameter choices in areas (d), (e) and (f) yield polynomial transmission times,* $d_{\mathcal{C}}(0,x) = \|x\|^{\eta_0\pm o(1)}$*, where* $\eta_0 = \eta_\beta = d(\mu-(3-\tau)/\beta)$ *on (d)* $\eta_0 = \eta_\alpha = d\mu(\alpha-2)/(\alpha-(\tau-1))$ *on (e), and* $\eta_0 = 1$ *on (f). The bold lines indicate discontinuous phase transitions, while the other transitions are smooth.*

consider μ as the easiest parameter to change: increasing μ means gradually slowing down the spreading process around high-degree vertices, which corresponds to adjusting behaviour of individuals with high number of contacts. Hence, we will phrase our results from this perspective. Figure 3 shows two phase diagrams: one where μ is fixed and τ, α vary; another where τ is fixed and μ, α vary.

1.3. Results

In this paper, we focus on the sub-explosive parameter regime

$$\mu > \frac{3-\tau}{2\beta} := \mu_{\mathrm{expl}}, \tag{1.7}$$

since for $\mu < \mu_{\mathrm{expl}}$ we have shown in previous work [55] that the model is *explosive*: the cost-distance of two vertices x, y converges in distribution to an almost surely finite variable as $|x-y| \to \infty$, conditioned on x and y being in the infinite component.[6] In other words, (1.7) restricts us to the nonexplosive phase. The following two quantities define the boundaries of the new phases:

$$\mu_{\log} := \frac{3-\tau}{\beta}, \quad \mu_{\mathrm{pol}} := \frac{1}{d} + \frac{3-\tau}{\min\{\beta, d(\alpha-2)\}} = \max\left\{1/d + \mu_{\log}, \mu_{\mathrm{pol},\alpha}\right\}, \tag{1.8}$$

where we define[7] $\mu_{\mathrm{pol},\alpha} = \mu_{\mathrm{pol},\alpha}(d,\tau,\alpha) := \frac{1}{d} + \frac{3-\tau}{d(\alpha-2)} = \frac{\alpha-(\tau-1)}{d(\alpha-2)}$ and $\mu_{\mathrm{pol},\beta} = \mu_{\mathrm{pol},\beta}(d,\tau,\beta) := \frac{1}{d} + \frac{3-\tau}{\beta}$. We also define two *growth exponents*. If $\alpha \in (1,2)$ or $\mu \in (\mu_{\mathrm{expl}}, \mu_{\log})$, we define

$$\Delta_0 := \Delta_0(\alpha,\beta,\mu,\tau) := \frac{1}{1-\log_2(\min\{\alpha, \tau-1+\mu\beta\})} = \min\{\Delta_\alpha, \Delta_\beta\} > 1, \tag{1.9}$$

[6]The phase is called *explosive* since the size of the the cost-ball of radius r jumps from finite to infinite at some random finite threshold, called the *explosion time*.

[7]In $\mu_{\mathrm{pol},\alpha}(d,\tau,\alpha)$ and $\mu_{\mathrm{pol},\beta}(d,\tau,\beta)$ we consider the respective indexing α and β after 'pol' in the subscript as symbols, rather than numerical values. The same holds for the functions $\Delta_\alpha, \Delta_\beta, \eta_\alpha, \eta_\beta$ describing the growth exponents: the subscripts are meant to be considered as symbols.

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

with $\Delta_\alpha = \Delta_\alpha(\alpha) := 1/(1-\log_2\alpha)$ and $\Delta_\beta = \Delta_\beta(\tau,\mu,\beta) = 1/(1-\log_2(\tau-1+\mu\beta))$. $\Delta_0 > 1$ follows since when $\alpha \in (1,2)$ then $\Delta_\alpha > 1$, while when $\mu \in (\mu_{\mathrm{expl}}, \mu_{\log})$ then $\tau - 1 + \mu\beta > \frac{\tau+1}{2} > 1$ and also $\tau - 1 + \mu\beta < 2$, so $\log_2(\tau - 1 + \mu\beta)$ is positive but less than 1. If both $\alpha > 2$ and $\mu > \mu_{\log}$, we define

$$\eta_0 := \eta_0(\alpha,\beta,\mu,\tau) := \begin{cases} 1 & \text{if } \mu > \mu_{\mathrm{pol}}, \\ \min\{d(\mu-\mu_{\log}), \mu/\mu_{\mathrm{pol},\alpha}\} & \text{if } \mu \le \mu_{\mathrm{pol}}, \end{cases} \tag{1.10}$$

and note that $\eta_0 > 0$ for all $\mu > \mu_{\log}$, and $\eta_0 < 1$ exactly when $\mu < \mu_{\mathrm{pol}}$ by (1.8). We often write

$$\begin{aligned} \eta_\beta &= \eta_\beta(d,\tau,\mu,\beta) := d(\mu - \mu_{\log}) = d(\mu - (3-\tau)/\beta), \\ \eta_\alpha &= \eta_\alpha(d,\tau,\mu,\alpha) := \mu/\mu_{\mathrm{pol},\alpha} = \frac{\mu d(\alpha-2)}{\alpha - (\tau-1)}. \end{aligned} \tag{1.11}$$

The formulas can be naturally extended by taking limits and hold also when $\alpha = \infty$ or $\beta = \infty$, which we elaborate in Section 1.3.1 below.

We first formulate the main results for the infinite models IGIRG and SFP. We write $0 \leftrightarrow x$ for the event that there is at least one path of edges in the graph between vertices $0, x$. Whenever $\tau \in (2,3)$, these models have a unique infinite connected component with constant density, hence the event $0 \leftrightarrow x$ occurs with (uniformly) positive probability given that $0, x$ are part of the vertex set [18, 26, 55, 33]. For SFP, the vertex set is deterministic and thus $0, x \in \mathcal{V}$ holds under the assumption that $x \in \mathbb{Z}^d$. For IGIRG, we need to condition on $0, x \in \mathcal{V}$. Formally, this is achieved by switching to the *Palm measure* of the Poisson point process. The Palm measure of a Poisson point process is again a unit intensity PPP on $\mathbb{R}^d$ with the vertices $0, x$ added to the vertex set, and with all vertex-weights, edges, and edge-costs still drawn by the Equations (1.6), (1.5) and (1.2) respectively, see also the book [61]. Later in Remark 1.9 we will also give a conditional version of the following theorem given the weighted vertex set.

Theorem 1.4. *Consider* 1*-FPP in Definition 1.1 on the graphs IGIRG or SFP of Definition 1.3 satisfying the assumptions given in (1.6)–(1.3) with* $\tau \in (2,3), \alpha > 1, d \ge 1, \mu > 0$. *Assume either* $\alpha \in (1,2)$ *or* $\mu \in (\mu_{\mathrm{expl}}, \mu_{\log})$ *or both hold. For SFP, assume* $x \in \mathbb{Z}^d$. *Then for any* $\varepsilon > 0$,

$$\lim_{|x|\to\infty} \mathbb{P}\big(d_{\mathcal{C}}(0,x) \le (\log|x|)^{\Delta_0+\varepsilon} \mid 0, x \in \mathcal{V}, 0 \leftrightarrow x\big) = 1.$$

For IGIRG, due to the conditioning $0, x \in \mathcal{V}$, $\mathbb{P}$ *is the Palm version of the annealed probability measure taken over edges, edge-costs, vertex-weights and -locations.*

The result of Theorem 1.4 is also valid when $\mu < \mu_{\mathrm{expl}}$, however, then the model is explosive [55, Theorem 1.1], and the bound is not sharp. With the restriction $\mu > \mu_{\mathrm{expl}}$, we conjecture that Theorem 1.4 is actually sharp, that is, that a corresponding lower bound with exponent $\Delta_0 - \varepsilon$ also holds. The exponent $\Delta_0 > 1$ intuitively corresponds to stretched exponential ball-growth, where the number of vertices in cost-distance at most r scales as $\exp(r^{1/\Delta_0})$. Trapman in [74] showed that strictly exponential ball growth, that is, $\Delta_0 = 1$, is possible for long-range percolation when $\alpha = 1$ under additional constraints. This is consistent with our formula for Δ_0, since $\Delta_0 \to 1$ as $\alpha \to 1$. Related is the work [60] that treats polylogarithmic graph distances and classical FPP transmission times in the same model class but in a different parameter regime (finite variance degrees, that is, $\tau > 3$), however the proof techniques do not extend to infinite variance degree underlying graphs and/or to 1-FPP. We leave the lower bound in this phase for future work.

Remark 1.8. *Structure of near-optimal paths in the polylog phase.* The proof reveals two different types of paths with polylogarithmic cost-distances present in the graph. When $\alpha < 2$, randomly occurring long edges on low-weight vertices cause the existence of paths of cost at most $(\log|x|)^{\Delta_\alpha + o(1)}$ with $\Delta_\alpha = 1/(1-\log_2(\alpha))$. The closest long edge of order $|x|$ lands at distance $|x|^{\alpha/2}$ from 0 and x respectively, resulting in a polylog exponent of Δ_α after iterating. When $\mu < \mu_{\log}$, there are also paths

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

using a cheap yet long edge (of order $|x|$) between two high-weight vertices (weight roughly $|x|^{d/2}$) that lie within distance $|x|^{(\tau-1+\mu\beta)/2+o(1)}$ from 0 and x respectively, and these cause the existence of paths of cost at most $(\log|x|)^{\Delta_\beta+o(1)}$ with $\Delta_\beta = 1/(1-\log_2(\tau-1+\mu\beta))$. Δ_β is the outcome of an optimisation: we minimise the distance between the high-weight vertices to 0 and x, while maintaining that an edge with constant cost exists between them. The minimal distance possible is of order $|x|^{(\tau-1+\mu\beta)/2+o(1)}$: the tail exponent $\tau-1$ of the weight distribution (1.6), and $\mu\beta$, the penalty exponent in (1.2) times the behaviour of the cdf of L in (1.3) both play a role.

When we increase μ above $\mu_{\log}$ and α above 2, we enter a new universality class and cost distances become polynomial:

Theorem 1.6. *Consider* 1*-FPP in Definition 1.1 on the graphs IGIRG or SFP of Definition 1.3 satisfying the assumptions given in (1.6)–(1.3) with $\tau\in(2,3), d\geq 1$. When $\alpha>2$ and $\mu>\mu_{\log}$ both hold, then for any $\varepsilon>0$,*

$$\lim_{|x|\to\infty}\mathbb{P}\big(d_{\mathcal{C}}(0,x)\leq |x|^{\eta_0+\varepsilon}\mid 0,x\in\mathcal{V}, 0\leftrightarrow x\big)=1.$$

Here $\mathbb{P}$ is the annealed probability measure taken over edges, edge-costs, vertex-weights and -locations.

In the accompanying [56] we prove the corresponding lower bound, which implies:

Corollary 1.7 (Polynomial Regime). *Consider* 1*-FPP in Definition 1.1 on the graphs IGIRG or SFP satisfying the assumptions given in (1.6)–(1.3) with $\tau\in(2,3), d\geq 1$. When $\alpha>2$ and $\mu>\mu_{\log}$ both hold, then for any $\varepsilon>0$,*

$$\lim_{|x|\to\infty}\mathbb{P}\big(|x|^{\eta_0-\varepsilon}\leq d_{\mathcal{C}}(0,x)\leq |x|^{\eta_0+\varepsilon}\mid 0,x\in\mathcal{V}, 0\leftrightarrow x\big)=1.$$

Corollary 1.7 together with Theorem 1.4 implies that the phase transition is proper at $\mu_{\log}$ and at $\alpha=2$: distances increase from at most polylogarithmic to polynomial. Moreover, when $\mu>\mu_{\mathrm{pol}}$, (i.e., $\min(\eta_\alpha,\eta_\beta)>1$), and the dimension $d\geq 2$, in [56] we also prove *strictly* linear cost-distances with both upper and lower bounds. This, together with Theorem 1.6, implies that there is another phase transition at μ_{pol}, from sublinear ($\eta_0<1$) to linear ($\eta_0=1$) cost-distances. See Table 2 for a summary. We find it remarkable that 1-FPP shows polynomial distances with exponent *strictly less than one* in a spread-out parameter regime $\mu\in(\mu_{\log},\mu_{\mathrm{pol}})$. This implies polynomial ball-growth faster than the dimension for 1-FPP, which is rare in spatial models, see Section 1.4.

Remark 1.6. *Structure of near-optimal paths in the polynomial phase.* The proof reveals two different types of paths with polynomial cost-distances present in the graph. When $\mu\leq\mu_{\mathrm{pol},\alpha}$, there are a few very long edges (of order $|x|$) with endpoints polynomially near 0 and x, emanating from vertices with weight $|x|^{1/(2\mu_{\mathrm{pol},\alpha})}$, and these results in paths with cost at most $|x|^{\eta_\alpha+o(1)}$ (when (1.10) evaluates to $\mu/\mu_{\mathrm{pol},\alpha}$). Since there are only few such edges, the optimisation effect of choosing the one with smallest cost is negligible and β does not enter the formula. Further, when $\mu\leq\mu_{\mathrm{pol},\beta}$, there are many long edges (of order $|x|$) with respective endpoints polynomially near 0 and x on vertices with weight roughly $|x|^{d/2}$, and when we optimise to choose the one with cheapest cost, the effect of F_L, that is, β in (1.3), enters the formula, and we obtain a path with cost at most $|x|^{\eta_\beta+o(1)}$ (when (1.10) evaluates to $d(\mu-\mu_{\log})$). The proof of the lower bound in [56] shows that in this phase *all* long edges near $0,x$ have polynomial costs in the Euclidean distance they bridge, which explains the qualitative difference between 1-FPP and classical FPP.

Remark 1.9. From our proofs it follows that a *vertex-weighted quenched version* of Theorems 1.4 and 1.6 are also valid in the following sense: there is a (cylinder) event $\mathcal{A}_{|x|}$ measurable with respect to the sigma-algebra generated by the vertex locations and vertex weights in a box of radius $C|x|$ centred at $0\in\mathbb{R}^d$ for some constant $C>1$, that holds with probability tending to 1 as $|x|\to\infty$. For any $\delta>0$,

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

for all sufficiently large $|x|$ and all realisations $(V, w_V) \in \mathcal{A}_{|x|}$ of the weighted vertex set with $0, x \in \mathcal{V}$,

$$\begin{aligned} &\mathbb{P}\Big(d_{\mathcal{C}}(0,x) \le (\log|x|)^{\Delta_0+\varepsilon} \mid (V, w_V), 0 \leftrightarrow x\Big) \ge 1-\delta, \text{ when } \alpha \in (1,2) \text{ or } \mu \in (\mu_{\mathrm{expl}}, \mu_{\log}) \\ &\mathbb{P}\big(d_{\mathcal{C}}(0,x) \le |x|^{\eta_0+\varepsilon} \mid (V, w_V), 0 \leftrightarrow x\big) \ge 1-\delta, \text{ when } \alpha > 2, \mu > \mu_{\log}. \end{aligned} \tag{1.12}$$

Here, $\mathbb{P}$ only integrates over the randomness of the edges and their i.i.d. edge-cost variables L_{xy}.

The next theorem describes in which sense the results stay valid for finite-sized models:

Theorem 1.10. *Consider 1-FPP in Definition 1.1 on the graph GIRG of Definition 1.3 satisfying the assumptions given in (1.6)–(1.3) with $\tau \in (2,3), \alpha > 1, d \ge 1, \mu > 0$. Let $\mathrm{C}^{(n)}_{\max}$ be the largest component in Q_n. Let u_n, v_n be two vertices chosen uniformly at random from $\mathcal{V} \cap Q_n$.*
(i) *When either $\alpha \in (1,2)$ or $\mu \in (\mu_{\mathrm{expl}}, \mu_{\log})$ or both hold, then for any $\varepsilon > 0$,*

$$\lim_{n\to\infty} \mathbb{P}\Big(d_{\mathcal{C}}(u_n, v_n) \le (\log|u_n - v_n|)^{\Delta_0+\varepsilon} \;\mid\; u_n, v_n \in \mathrm{C}^{(n)}_{\max}\Big) = 1. \tag{1.13}$$

(ii) *When $\alpha > 2$ and $\mu > \mu_{\log}$ both hold, then for any $\varepsilon > 0$,*

$$\lim_{n\to\infty} \mathbb{P}\big(d_{\mathcal{C}}(u_n, v_n) \le |u_n - v_n|^{\eta_0+\varepsilon} \;\mid\; u_n, v_n \in \mathrm{C}^{(n)}_{\max}\big) = 1. \tag{1.14}$$

The size of the largest component is linear with size $n\mathbb{P}(0 \leftrightarrow \infty)(1 + o(1))$, see [52]. The lower bound in Corollary 1.7 also transfers to finite GIRGs, since GIRG is defined as a subgraph of IGIRG. We refer to [56] for details. The proofs of Theorems 1.4, 1.6, and 1.10 also reveal that the paths realising the upper bounds deviate only sublinearly from the straight line between the two vertices, cf. Definition 5.6 and Lemmas 6.2 and 6.3 for more details.

1.3.1. Limit Cases and Extensions

Theorems 1.4–1.10 can be extended to interesting cases that may informally be described as $\alpha = \infty$ or $\beta = \infty$. In the case $\alpha = \infty$, all connection probabilities are either constant or zero, and we replace the condition (1.5) by

$$h(x, w_1, w_2) \begin{cases} = 0, & \text{if } \frac{w_1 w_2}{|x|^d} < c', \\ \ge \underline{c} & \text{if } \frac{w_1 w_2}{|x|^d} \ge c'', \end{cases} \tag{1.15}$$

for some constants $\underline{c} \in (0,1]$ and $c'' \ge c' > 0$. For the sake of simplicity we will assume $c'' = 1$ in all our proofs, however the results still hold for general c''. Models satisfying (1.15) are called threshold (or zero temperature) models, and include *hyperbolic random graphs* [58] when the dimension is one. The correspondence between GIRGs and threshold hyperbolic random graphs was established in [17, Theorem 2.3]. For models where (1.15) holds, we extend the definitions (1.8)-(1.9) in the natural way to $\alpha = \infty$, since $\lim_{\alpha\to\infty} \mu_{\mathrm{pol},\alpha} = 1/d$:

$$\mu_{\log} := \frac{3-\tau}{\beta}, \quad \mu_{\mathrm{pol}} := \frac{1}{d} + \frac{3-\tau}{\beta}, \quad \eta_0 := \begin{cases} 1 & \text{if } \mu > \mu_{\mathrm{pol}}, \\ d \cdot (\mu - \mu_{\log}) & \text{if } \mu \le \mu_{\mathrm{pol}}, \end{cases} \tag{1.16}$$

and, when $\mu \in (\mu_{\mathrm{expl}}, \mu_{\log})$,

$$\Delta_0 := \frac{1}{1 - \log_2(\tau - 1 + \mu\beta)} > 0. \tag{1.17}$$

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

The case $\beta = \infty$ captures when the cdf of the edge transmission variable L in (1.3) is flatter near 0 than any polynomial, and we replace (1.3) by the condition that

$$\lim_{t\to 0} F_L(t)/t^{\beta} = 0 \text{ for all } 0 < \beta < \infty. \tag{1.18}$$

In particular, this condition is satisfied if F_L has no probability mass around zero, for example[8] when $L \equiv 1$. When $\beta = \infty$, using that $\tau \in (2,3)$ we replace (1.8)-(1.10) naturally by

$$\mu_{\mathrm{expl}} := \mu_{\log} := 0, \quad \mu_{\mathrm{pol}} := \frac{\alpha - (\tau - 1)}{d(\alpha - 2)}, \quad \eta_0 := \begin{cases} 1 & \text{if } \mu > \mu_{\mathrm{pol}}, \\ \mu/\mu_{\mathrm{pol}} & \text{if } \mu \le \mu_{\mathrm{pol}}, \end{cases} \tag{1.19}$$

and, when $\alpha \in (1,2)$,

$$\Delta_0 := \frac{1}{1 - \log_2(\alpha)} > 0. \tag{1.20}$$

Finally, when both $\alpha = \beta = \infty$ we replace (1.8) and (1.10) by

$$\mu_{\mathrm{expl}} := \mu_{\log} := 0, \quad \mu_{\mathrm{pol}} := \tfrac{1}{d}, \quad \eta_0 := \min\{1, d\mu\}, \tag{1.21}$$

and in that case we do not define Δ_0, since the polylogarithmic case is vacuous when $\alpha = \beta = \infty$ (see also below Corollary 5.2). Our main results still hold for these limit regimes. We remark that the corresponding lower bounds also hold [56, Theorem 1.10].

Theorem 1.11 (Extension to threshold GIRGs and $\beta = \infty$). (a) *Theorems 1.4, 1.6 and 1.10 still hold for $\alpha = \infty$ if we replace definitions (1.8)-(1.10) by definitions (1.16)-(1.17).*

(b) *Theorems 1.4, 1.6 and 1.10 still hold for $\beta = \infty$ if we replace definitions (1.8)-(1.10) by definitions (1.19)-(1.20).*

(c) *Theorems 1.6 and 1.10 still hold for $\alpha = \beta = \infty$ if we replace definitions (1.8)-(1.10) by definition (1.21).*

Theorem 1.11(a) implies the analogous result for hyperbolic random graphs (HypRG) by setting $d = 1$ in (1.16), except for some minor caveats. In Definition 1.3, the number of vertices in GIRG is Poisson distributed with mean n, while in the usual definition of HypRG [58, 45] and GIRG [17] the number of vertices is exactly n. In HypRG the vertex-weights have an n-dependent distribution converging to a limiting distribution [57]. However, these differences may be overcome by coupling techniques presented in, for example, [57]: a model with exactly n vertices can be squeezed between two GIRGs with Poisson intensity $1 - \sqrt{4\log n/n}$ and $1 + \sqrt{4\log n/n}$, and one can couple n-dependent and limiting vertex-weights to each other, respectively, but we avoid spelling out the details and refer the reader to [57, Claims 3.2, 3.3].

1.4. Discussion

Here we discuss our results in context with related results about (inhomogeneous) first passage percolation and graph distances on spatial random graphs.

Long-range first passage percolation. The work on long-range first passage percolation (LR-FPP) [20] is closest to our work. In that model, the underlying graph is the *complete graph* of $\mathbb{Z}^d$, and the edge transmission time on any edge uv is exponentially distributed with mean $|u-v|^{d\alpha'+o(1)}$, so $\beta = 1$, the process is Markovian, and the penalty depends on the Euclidean distance of u and v. This choice eliminates the correlations coming from the presence/absence of underlying edges, and the growth is

[8]For $\mu = 0$, $L \equiv 1$, the cost-distance $d_{\mathcal{C}}(x,y)$ then equals the *graph-distance* between x and y. [56] contains as special cases the linear lower bound on graph-distances by Berger [8] for long-range percolation (LRP) and by Deprez, Hazra, and Wüthrich [27] for SFP, see [56, Proposition 2.4].

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

strictly governed by the long-range transmission times. As α' grows, [20] finds the same sub-explosive phases for transmission times in LR-FPP that we find for 1-FPP in Table 2. The main difference is that the explosive phase is absent in LR-FPP, and is replaced by a 'super-fast' phase there where transmission times are 0 almost surely. Moreover, the behaviour on phase boundaries are different. Using the symmetries in their model, [20] proves that whenever transmission times in LR-FPP are strictly positive, then they must be at least logarithmic. In contrast, in general 1-FPP on IGIRG and SFP we also see doubly logarithmic distances, for example for graph-distances ($L \equiv 1, \mu = 0$). We summarise the results on LR-FPP in Table 3. Nevertheless, in 1-FPP, the cost function $\mathcal{C}(xy)$ in (1.2) could also depend on $|x-y|$, that is, take the form $L_{xy}(W_xW_y)^{\mu}|x-y|^{\zeta}$. The result of [55] on explosion carries through to this case without much effort [70], with the model being explosive if and only if $\mu + \zeta/d < (3-\tau)/(2\beta) = \mu_{\mathrm{expl}}$, with $\tau \in (2,3)$. In an ongoing work, we determine the full phase diagram of cost-distances with spatial penalisation also present [7], which turn out to be more complex than simply replacing μ with $\mu + \zeta/d$.

Qualitative difference between one-dependent FPP and graph distances. Some phases of 1-FPP in Table 2 are also phases for *graph-distances* in spatial models in general. However, while the polynomial phase is spread-out in 1-FPP, this phase is essentially absent for graph distances. Indeed, the polynomial phase occurs when long edges all have polynomial spreading times in the Euclidean distance they bridge, both in 1-FPP here and in LR-FPP in [20]. Thus, transmission times in 1-FPP are not equivalent to graph distances in any inhomogeneous percolation on the underlying graph. Table 3 summarises known results on 1-FPP, LR-FPP, and graph distances in spatial graphs. Now we elaborate on each phase.

The polylogarithmic phase. Theorem 1.4 proves polylogarithmic cost-distances in 1-FPP when $\tau \in (2,3)$, and either $\mu \in (\mu_{\mathrm{exp}}, \mu_{\log})$ or $\alpha \in (1,2)$. The results here, in [55] and the accompanying [56] (Corollary 1.7) together imply that μ_{exp} and $\mu_{\log}$ are true phase-transition points, separating this phase from both the explosive and the polynomial phases. Even though we do not have a matching lower bound, we conjecture that this phase is truly polylogarithmic, and the exponent Δ_0 in (1.9) is sharp. The exponent Δ_0 also depends on the product $\mu\beta$, which does not allow to match it easily to exponents for graph-distances: For long-range percolation, where each edge $(u,v) \in \mathbb{Z}^d \times \mathbb{Z}^d$ is present independently with probability $\Theta(|u-v|^{-d\alpha})$, Biskup and Lin [12] show that graph distances grow polylogarithmically with exponent $\Delta_\alpha = 1/(1-\log_2 \alpha)$ when $\alpha \in (1,2)$. This coincides with our upper bound in Theorem 1.4 if $\alpha \le \tau - 1 + \mu\beta$. The same type of paths are used in both cases, passing through only low-degree vertices (and typical edge-costs on them for 1-FPP). For SFP, Lakis *et al.* prove in [60, Theorem 1.1] that graph distances and transmission times in Markovian FPP are also polylogarithmic when $\alpha \in (1,2)$ and additionally $\tau > 3$, with exponent $\Delta_G \in [1/(1-\log_2(\min(\alpha, \tau-2))), \Delta_\alpha]$ for graph distances and $\Delta_{\mathrm{FPP}} \in [1/(1-\log_2(\min(\alpha, (\tau-1)/2))), \Delta_\alpha]$ for FPP, which improves earlier bounds [46]. The lower-bound methods in [12, 74, 60] do not transfer to 1-FPP when $\tau \in (2,3)$ since they crucially rely on finite degree-variance $\tau > 3$.

The linear phase. Linear distances are common in supercritical spatial graph models with bounded edge-lengths. For example, Random Geometric Graphs exhibit linear distances [68], and so does supercritical percolation on grids of dimension at least 2 [4]. Assuming high enough edge-density, a renormalisation argument to percolation on $\mathbb{Z}^d$ gives that SFP and LRP for $\tau > 3$ and $\alpha > 2$ also have at most linear graph-distances for $d \ge 2$. The corresponding lower bound was shown by Berger for LRP [8] and by Deprez *et al.* for SFP [27]. Our lower bound for 1-FPP contains these as special cases, and holds universally for classical FPP for any positive edge-transmission time-distribution [56, Corollary 1.12].

The strictly polynomial phase. The phase where intrinsic distances scale as $|x-y|^{\eta_0 + o(1)}$ with $\eta_0 < 1$ (the result of Theorem 1.6) is quite rare in spatial settings and we only know two examples. One is in LRP at a boundary line in the parameter space, when $\alpha = 2$ [5, 24]. Our methods do not carry through for $\alpha = 2$. The method in [5] for this setting uses the self-similarity of the model when $\alpha = 2$ and shows the sub-multiplicative structure of graph distances to obtain polynomial lower bounds. The other example is for long-range first passage percolation (LR-FPP) in [20], mentioned at the beginning of this section. There are some similarities to 1-FPP: LR-FPP is Markovian, that is, $\beta = 1$ in (1.3),

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

and has strictly polynomial growth when $\alpha' \in (2, 2+1/d)$, see Table 3. Using exponential L in 1-FPP, the length of the parameter interval $(\mu_{\log}, \mu_{\text{pol}})$ with polynomial growth is also exactly $1/d$ for μ when $\alpha > 2+1/d$, but it is longer when $\alpha < 2+1/d$, which shows that the penalty α' of LR-FPP plays a slightly different role as the long-range parameter α in (1.5) here.

Gaps at approaching the phase boundaries. Here we discuss what happens as the parameters τ, α, μ, β approach the phase boundaries of growth. Some of these are indicated on Figure 3 by bold lines.

Polylogarithmic distances with exponent Δ_0 heuristically imply *stretched exponential ball-growth*, where the number of vertices within intrinsic distance r scales as $\exp(r^{1/\Delta_0})$. Our upper bound exponent $\Delta_0 = \min\{\Delta_\alpha, \Delta_\beta\}$ in (1.9) approaches 1 as $\alpha \downarrow 1$, and so does the exponent Δ_α of LRP [11], which also partly governs SFP. This means that as $\alpha \downarrow 1$ we approach exponential growth. In LRP, strictly exponential ball growth occurs only when $\alpha = 1$ and the connectivity function has a suitably chosen slowly varying correction term $\ell(\cdot)$, that is, $p(x,y) = \ell(|x-y|)/|x-y|^{\alpha d}$, see [74]. Strictly exponential growth is a natural barrier, since (age-dependent) branching processes with finite first moments exhibit at most exponential growth, and non-Markovian FPP can be dominated by such branching processes. Interestingly, when $\alpha > 2$ and we approach the explosion phase transition by letting $\mu \downarrow \mu_{\text{expl}} = (3-\tau)/(2\beta)$, then Δ_0 in (1.9) does not converge to 1, but to $1/(2-\log_2(\tau+1)) =: \Delta_\tau$. So, for the whole range $\tau \in (2,3)$, $\Delta_\tau \geq 1/(2-\log_2(3)) > 2.4 > 1$. As $\tau \uparrow 3$, Δ_τ approaches ∞, which is natural, since graph distances are linear already when $\tau > 3$ and $\alpha > 2$ [27]. This leaves two possibilities: our upper bound Δ_0 is not sharp for $\alpha > 2$; or the ball growth jumps directly from subexponential ($\Delta_0 > 1$) into the explosive phase. If the latter is the case, it would be interesting to understand better how such a jump could happen. Such jumps at phase boundaries may occur. This paper, together with [56], proves a gap in *polynomial regime* when τ crosses the threshold 3. The limits of $\lim_{\tau\uparrow 3} \mu_{\text{pol}} = 1/d$ and $\lim_{\tau\uparrow 3} \eta_0 = \mu d$ exist and are in $(0,1)$ in (1.8) and (1.10). So if we fix some $\mu < 1/d$ and let $\tau \uparrow 3$, the cost-distances grow polynomially with exponents bounded away from one (e.g., they approach $1/2$ from below for $\mu = 1/(2d)$). But as soon as $\tau > 3$ is reached, the exponent 'jumps' to 1 and distances become strictly linear [56, Theorem 1.11]. So the parameter space is discontinuous in η_0 and μ_{pol} with respect to τ.

Some important questions are centred around such gaps. The *gap conjecture* in geometric group theory is about the ball growth of finitely generated groups: it states that there are no groups with growth between polynomial and stretched exponential of order $\exp(\Theta(\sqrt{n}))$ [43]. Although the polynomial side is understood by Gromov's theorem [44], the conjecture remains open. We find it intriguing to discover a deeper connection between phases of intrinsic growth in spatial random graphs ('stochastic lattices') and group theory ('deterministic lattices').

Organisation. We start by moving to the quenched setting. In Section 2 we develop the nets, and in Section 3 the multiround exposure of edges, with the main result in Proposition 3.9. In Section 4 we collect some connectivity-estimates that serve as our building blocks, while in Section 5 we carry out the 'budget travel plan' and build a hierarchical path that only uses vertices of the net and connects vertices $y_0^\star, y_x^\star$ very close to 0 and x, respectively. In Section 6 we connect $0, x$ to these vertices with low cost, which is a nontrivial task itself, and prove the main theorems.

1.4.1. Notation

Throughout, we consider simple and undirected graphs with vertex set $\mathcal{V} \subseteq \mathbb{R}^d$. For a graph $G = (\mathcal{V}, \mathcal{E})$ and a set $A \subseteq \mathbb{R}^d$, $G[A]$ stands for the induced subgraph of G with vertex set $\mathcal{V} \cap A$. For two vertices $x, y \in \mathcal{V}$, we denote the edge between them by xy, and for a set $V \subseteq \mathcal{V}$ we write $V^{(2)} := \{\{x,y\} : x, y \in V, x \neq y\}$ for the set of possible edges among vertices in V. For a path π, $\mathcal{E}(\pi)$ is the set of edges forming π, and $|\pi|$ is the number of edges of π. Generally the size of a discrete set S is $|S|$, while of a set $A \subseteq \mathbb{R}^d$, $\text{Vol}(A)$ is its Lebesgue measure. Given a cost function $\mathcal{C} : \mathcal{E} \to [0, \infty]$ on the edges, the cost of a set of edges $\mathcal{P}$ is $\mathcal{C}(\mathcal{P}) := \sum_{e \in \mathcal{E}(\mathcal{P})} \mathcal{C}(e)$. We define $\mathcal{C}(xx) := 0$ for all $x \in \mathcal{V}$. We define the *cost-distance* between vertices x and y as

$$d_{\mathcal{C}}(x,y) := \inf\{\mathcal{C}(\pi) : \pi \text{ is a path from } x \text{ to } y \text{ in } G\}. \tag{1.22}$$

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

We define the graph distance $d_G(x, y)$ similarly, where all edge-costs are set to 1. We denote the Euclidean norm of $x \in \mathbb{R}^d$ by $|x|$, the Euclidean ball with radius $r \geq 0$ around x by $B_r(x) := \{y \in \mathbb{R}^d : |x - y| \leq r\}$, and the set of vertices in this ball by $\mathcal{B}_r(x) := \{y \in \mathcal{V} : |x - y| \leq r\} = B_r(x) \cap \mathcal{V}$. (The minimal notation difference is intentional). The *graph-distance ball* and *cost-distance ball* (or *cost-ball* for short) around a vertex x are the vertex sets $\mathcal{B}_r^G(x) := \{y \in \mathcal{V} : d_G(x, y) \leq r\}$ and $\mathcal{B}_r^{\mathcal{C}}(x) := \{y \in \mathcal{V} : d_{\mathcal{C}}(x, y) \leq r\}$, respectively. We set $B_r := B_r(0)$, and do similarly for $\mathcal{B}_r, \mathcal{B}_r^G, \mathcal{B}_r^{\mathcal{C}}$ if 0 is a vertex. We define $\partial B_r(x) := B_r(x) \setminus \{y \in \mathbb{R}^d : |x - y| < r\}$, and use similar definitions for $\partial \mathcal{B}_r$, $\partial \mathcal{B}_r^G$ and $\partial \mathcal{B}_r^{\mathcal{C}}$. In particular, $\partial \mathcal{B}_1^G(v)$ is the set of neighbours of v.

The set of *model parameters* are $\mathtt{par} := \{d, \tau, \alpha, \mu, \beta, \underline{c}, \overline{c}, c_1, c_2, t_0\}$. For parameters $a, b > 0$ (model parameters or otherwise), we use 'for all $a \gg_\star b$' as shortcut for '$\forall b > 0 : \exists a_0 = a_0(b) : \forall a \geq a_0$'. We also say '$a \gg_\star b$' to mean that $a \geq a_0(b)$. We use $a \ll_\star b$ analogously, and may use more than two parameters. For example, 'for $a \gg_\star b, c$' means '$\forall b, c > 0 : \exists a_0 = a_0(b, c) : \forall a \geq a_0$'. A measurable function $\ell : (0, \infty) \to (0, \infty)$ is said to be *slowly varying* at infinity if $\lim_{x\to\infty} \ell(cx)/\ell(x) = 1$ for all $c > 0$. We denote by log the natural logarithm, by $\log_2$ the logarithm with base 2, and by $\log^{*k}$ the k-fold iterated logarithm, for example, $\log^{*3} x = \log\log\log x$. For $n \in \mathbb{N}$ we write $[n] := \{1, \ldots, n\}$, and for an event $\mathcal{A}$ we denote by $\mathcal{A}^{\complement}$ its complement.

2. Working conditionally on the weighted vertex-set: nets

In proving the upper bounds (Theorems 1.4 and 1.6), we will construct cheap paths along the lines of the 'budget travel plan' on page 7 in Section 1, which is an iterative scheme of finding long 3-edge bridge-paths to connect two far-away vertices. Since low-cost events in 1-FPP are not increasing, we develop a technique that replaces the FKG-inequality. Moving to the quenched setting, we will first expose all vertex positions and weights (above some threshold weight, in the case of IGIRG); then, low-cost edge existence events become independent. To be able to work with *fixed realisations* of the vertex set, we find (with high probability as $|x| \to \infty$) a 'nice' subset of the vertices, that is, that behaves regularly enough inside a box around $0, x$, as in (1.4), which we call a *net*. We formalise the notion of the nets now. We start with a less demanding notion of nets that we call *weak nets* which will suffice for the further sections of the paper. Their existence will follow from the existence of strong nets which make (1.4) precise; this may be of independent interest, and most of the section shall be devoted to proving their existence.

The vertex-weight distribution satisfies $\mathbb{P}(W \geq w) = \ell(w)w^{-(\tau-1)}$ in Definition 1.3, and consider the slowly varying function $\ell(w)$ from (1.6). We define w_0 to be the smallest integer in $[1, \infty)$ such that

$$\forall w > w_0, \forall t \in [1/2, 2] : \quad \ell(w)w^{-(\tau-1)} < 2^{-\tau-8} \quad \text{and} \quad 0.99 \leq \ell(tw)/\ell(w) \leq 1.01 \tag{2.1}$$

both hold. Note that w_0 satisfying (2.1) must exist since ℓ is a slowly varying function, and so Potter's bound [10] ensures the first inequality.

For a set $A \subset \mathbb{R}^d$ we write $\mathrm{Vol}(A)$ for its Lebesgue measure (volume), while for a discrete set $\mathcal{A} \subseteq (0, \infty)$ we write $|\mathcal{A}|$ for the cardinality (size) of the set. Recall that weighted vertices are of the form $\tilde{v} = (v, w_v) \in \mathbb{R}^d \times [1, \infty)$.

Definition 2.1 (Weak net). Let $Q \subseteq \mathbb{R}^d$ be a box of side length ξ, $\varepsilon > 0$, and $w_1 \geq w_0$. A *weak* (ε, w_1)*-net* for Q is a set $\mathcal{N} \subseteq \widetilde{\mathcal{V}} \cap Q \times [1, \infty)$ of size at least $\mathrm{Vol}(Q)/4$ such that for all $\widetilde{v} \in \mathcal{N}$, *all* $r \in [(\log\log \xi\sqrt{d})^{4/\varepsilon}, \xi\sqrt{d}]$ and all $w \in [w_1, r^{d/(\tau-1)-\varepsilon}]$:

$$|\mathcal{N} \cap (B_r(v) \times [w/2, 2w])| \geq r^{d(1-\varepsilon)} \cdot \ell(w)w^{-(\tau-1)}. \tag{2.2}$$

In a weak (ε, w_1)-net, the number of weighted vertices in balls around net-vertices are close to their expectation. Since we only require the property to hold around net vertices and not everywhere, we circumvent the issue that the vertex set may have empty/low-density areas. In a weak net, we allow an error of order $r^{-d\varepsilon}$ on the right-hand side of (2.2), and the minimal radius of the balls around net vertices grows with the size ξ of the box Q. In a strong net, we shall only allow a constant factor loss,

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

see (2.4), and the balls can have constant radii r. The result that we shall use in further sections is the following, again using the Palm measure in the case of IGIRG [61]. We condition on a few vertices to be part of the vertex set and also demand that these vertices are part of the net.

Lemma 2.2. *Consider IGIRG or SFP with $\tau > 2$ in Def. 1.3. Then for all $\varepsilon \in (0, 1/2)$, and for all ξ sufficiently large relative to ε, and $t \le \min\{1/\varepsilon, \log\log \xi\}$ the following holds. Consider a cube $Q \subseteq \mathbb{R}^d$ of side length ξ, and let $x_1, \dots, x_t \in Q$, and w_0 from (2.1), then*

$$\mathbb{P}(Q \text{ contains a weak } (\varepsilon, w_0)\text{-net } \mathcal{N}, \text{ and } x_1, \dots, x_t \in \mathcal{N} \mid x_1, \dots, x_t \in \mathcal{V}) \ge 1 - t\varepsilon.$$

The condition $t \le 1/\varepsilon$ is there to avoid a vacuous statement, and below we set $t = 2$ and replace the conditioning by $0, x \in \mathcal{V}$. Note that the condition (2.2) never counts vertices of weight less than $w_1/2$. So, we can decide whether a weak (ε, w_1)-net $\mathcal{N}$ exists by uncovering only the set of vertices of weight at least $w_1/2$ (beyond $x_1, \dots, x_t \in \mathcal{N}$). For IGIRG, this set is independent of the set of vertices of weight smaller than $w_1/2$, and we may use low-weight vertices for other purposes. This is the main reason for introducing the parameter w_1. We mention that Q does not need be a box of equal sizes, the proof also works for boxes of any finite proportions parametrised by constant multiples of ξ.

In order to prove the existence of weak nets, we will divide Q into nested sub-boxes and work inductively, taking the finest partition and lowest-weight vertices as the base case and gradually coarsening the partition and including higher-weight vertices. To make this argument work, we will need stronger control over the properties of each layer of the partition than Definition 2.1 provides; see Definition 2.10. As such, it is convenient to instead prove the existence of 'strong nets' which satisfy tighter bounds at specific scales, which we define next. We will then prove that every strong net is also a weak net, as our strong control at each layer will translate into weaker control between layers. Strong nets may also be of independent interest in cases where stronger bounds over smaller scales are required.

Recall w_0 from (2.1). For all $\delta > 0$, and $R > 0$ we define the function $f_{R,\delta}(r)$ slightly below the typical largest vertex weight in a ball of radius r (roughly $r^{d/(\tau-1)}$):

$$f_{R,\delta}(r) = r^{\frac{d}{\tau-1}}\Big(1 \wedge \inf\big\{\ell(x)\colon x \in [w_0, r^{d/(\tau-1)}]\big\}\Big)^{\frac{1}{\tau-1}} \cdot \left(\frac{1}{(2d)^{2\tau+d+8}\log(16R/\delta)}\right)^{\frac{1}{\tau-1}}. \tag{2.3}$$

Definition 2.3 (Strong net). Let $G = (\mathcal{V}, \mathcal{E})$ be an IGIRG or SFP as in Definition 1.3. Let $\mathcal{R} = (r_1, r_2, \dots, r_R) \subseteq (0, \infty)$ be an increasing list of radii with $R = |\mathcal{R}| < \infty$, let w_0 be as in (2.1) and $f_{R,\delta}(\cdot)$ be as in (2.3). Let $Q \subseteq \mathbb{R}^d$ be a box. A $(\delta, \mathcal{R})$-net for Q is a set $\mathcal{N} \subseteq \widetilde{\mathcal{V}} \cap Q \times [1, \infty)$ of size at least $\mathrm{Vol}(Q)/4$ such that for all $\tilde{v} \in \mathcal{N}$, $r \in \mathcal{R}$, and all $w \in [w_0, f_{R,\delta}(r)]$,

$$\Big|\big\{\widetilde{u} \in \mathcal{N} \cap B_r(v) \times [w/2, 2w]\big\}\Big| \ge r^d \cdot \ell(w) w^{-(\tau-1)}/(2d)^{d+\tau+5}. \tag{2.4}$$

Each $r \in \mathcal{R}$ in Definition 2.3 will correspond to a layer of the discretisation of Q alluded to above. We will require these radii to grow at least exponentially, with base depending on the number R of radii in the list. The specific condition is the following. It may seem very strong that we require (2.4) for infinitely many w. We discretise $[w_0, f_{R,\delta}(r)]$ into a finite set of subintervals $(I_j)_{j \le j_{\max}}$ in a smart way. Then we ensure that (2.4) holds with $[w/2, 2w]$ replaced by I_j on the left hand side and the $\ell(w)w^{-(\tau-1)}$ replaced by $\mathbb{P}(W \in I_j)$ on the right hand side, and then this will imply that (2.4) also holds for all values $w \in [w_0, f_{R,\delta}(r)]$.

Definition 2.4. Fix $\delta \in (0, 1)$. We say that an increasing list of radii $\mathcal{R} = (r_1, r_2, \dots, r_R) \subseteq (0, \infty)$ is *δ-well-spaced* if $R = |\mathcal{R}| < \infty$ and the following hold:

$$r_1 \ge 24d\big(\log(4R/\delta)\big)^{1/d} \vee w_0^{(\tau-1)/d} \vee \inf\{r\colon f_{R,\delta}(r) \ge w_0\}; \tag{2.5}$$

$$\frac{r_i}{r_{i-1}} \ge 6R^{1/d}\left(\frac{\log(2R/\delta)}{\delta}\right)^{1/d} \qquad \forall i \in [2, R]. \tag{2.6}$$

In Definition 2.3, it may seem very strong that we require (2.4) for infinitely many w. We discretise $[w_0, f_{R,\delta}(r)]$ into a finite set of subintervals $(I_j)_{j \le j_{\max}}$ in a smart way. Then we ensure that (2.4) holds with $[w/2, 2w]$ replaced by I_j on the left hand side and the $\ell(w)w^{-(\tau-1)}$ replaced by $\mathbb{P}(W \in I_j)$ on the right hand side, and then this will imply that (2.4) also holds for all values $w \in [w_0, f_{R,\delta}(r)]$.

We now state the main result of the section. Heuristically, a box Q contains an $(\delta, \mathcal{R})$-net with sufficiently high probability, and we can also condition on the presence of a few vertices in the net. The condition $t \le 1/\delta$ in the following is added to avoid a vacuous statement.

Proposition 2.5. *Consider IGIRG or SFP with $\tau > 2$. Let $\delta \in (0, 1/16)$, $\xi > 0$, and $Q \subseteq \mathbb{R}^d$ be a cube of side length ξ. Let $R \in \mathbb{N}$ and $\mathcal{R} = \{r_1, \dots, r_R\}$ be a δ-well-spaced increasing list of radii with $r_R = \xi\sqrt{d}$. Let $x_1, \dots, x_t \in Q$ with $t \le \min\{1/\delta, (r_1/4\sqrt{d})^d\}$, for SFP assume $x_1, \dots x_t \subset \mathbb{Z}^d$ also. Then*

$$\mathbb{P}(Q \text{ contains an } (\delta, \mathcal{R})\text{-net } \mathcal{N}) \ge 1 - \delta/R; \tag{2.7}$$

$$\mathbb{P}(Q \text{ contains an } (\delta, \mathcal{R})\text{-net } \mathcal{N}, x_1, \dots, x_t \in \mathcal{N} \mid x_1, \dots, x_t \in \mathcal{V}) \ge 1 - t\delta. \tag{2.8}$$

For IGIRG (2.8) uses the Palm measure of vertex-weights and positions for IGIRG. For SFP the conditioning can be dropped.

The proof of this proposition will take the rest of the section. Before that we show how Lemma 2.2 follows from it.

Proof of Lemma 2.2 subject to Proposition 2.5. Let $\varepsilon \in (0, 1/2)$, set $\eta := 1 - \varepsilon/2$, and define $\mathcal{R} = \{r_1, \dots, r_R\}$ as

$$R := 2 + \lfloor (\log\log \xi\sqrt{d} - \log^{*4} \xi\sqrt{d} - \log \tfrac{4}{\varepsilon})/\log(1/\eta) \rfloor, \tag{2.9}$$

$$r_i := (\xi\sqrt{d})^{\eta^{R-i}}, \quad \text{for} \quad i \in [R]. \tag{2.10}$$

We show that this choice of $\mathcal{R}$ is ε-well-spaced (Def. 2.4). Let $1 - a \in [0, 1)$ be the fractional part of the expression inside the $\lfloor \cdot \rfloor$ in (2.9). Then using that $\lfloor x \rfloor = x - 1 + a$,

$$r_1 = (\xi\sqrt{d})^{\eta^{R-1}} = (\xi\sqrt{d})^{\eta^a 4 \log^{*3} \xi\sqrt{d}/(\varepsilon \log \xi\sqrt{d})} = (\log\log \xi\sqrt{d})^{(1-\varepsilon/2)^a 4/\varepsilon}. \tag{2.11}$$

Since $a \in (0, 1]$ and $\varepsilon < 1/2$, this implies that r_1 is a strictly larger power of $\log\log \xi\sqrt{d}$ than 1 while R is of order $\log\log \xi\sqrt{d}$. Since ξ is large relative to ε, Def. 2.4 (2.5) holds for $\varepsilon = \delta$. For all $i \in [R]$, since $\eta = 1 - \varepsilon/2$:

$$r_i/r_{i-1} = (\xi\sqrt{d})^{\eta^{R-i}(1-\eta)} \ge (\xi\sqrt{d})^{\eta^{R-1}\varepsilon/2} = r_1^{\varepsilon/2} = (\log\log \xi\sqrt{d})^{2(1-\varepsilon/2)^a}.$$

Since $a \le 1$, $2(1 - \varepsilon/2)^a > 1$ for all $\varepsilon < 1/2$, so Def. 2.4 (2.6) holds even in $d = 1$, and $\mathcal{R}$ is ε-well-spaced as claimed. Moreover, since $t \le \log\log \xi$ by hypothesis, by (2.11) we have $t \le (r_1/4\sqrt{d})^d$. By Proposition 2.5, with probability at least $1 - t\varepsilon$, conditioned on $x_1, \dots, x_t \in \mathcal{V}$, Q contains a strong $(\varepsilon, \mathcal{R})$-net $\mathcal{N}$ in Def. 2.3. We now show that a strong net is also a weak net (with the same ε). Fix an arbitrary $r \in [(\log\log \xi\sqrt{d})^{4/\varepsilon}, \xi\sqrt{d}]$ and $w \in [w_0, r^{d/(\tau-1)-\varepsilon}]$ as in Def. 2.1 (this interval is nonempty since $\tau \in (2, 3)$, $\varepsilon < 1/2$), and consider a vertex $v \in \mathcal{N}$. Since $a > 0$ in (2.11), we have $r_1 \le r \le r_R$. Let r_j be such that $r \in [r_j, r_{j+1})$; thus $r^\eta \le r_j \le r$ by (2.10). Since ξ is large relative to ε, we have $w \le f_{R,\varepsilon}(r_j)$ (using (2.3)) for $f_{R,\varepsilon}$); thus by the definition of a strong $(\varepsilon, \mathcal{R})$-net in Def. 2.3, $|\mathcal{N} \cap (B_{r_j}(v) \times [w/2, 2w])| \ge \ell(w)w^{-(\tau-1)} r_j^d/(2d)^{d+\tau+5}$. Since ξ is large relative to ε and $r \ge r_j \ge r^\eta$, the required inequality in Def. 2.1 follows since

$$|\mathcal{N} \cap (B_r(v) \times [w/2, 2w])| \ge \ell(w)w^{-(\tau-1)} r^{d\eta}/(2d)^{d+\tau+5} \ge \ell(w)w^{-(\tau-1)} r^{d(1-\varepsilon)}.$$

□

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

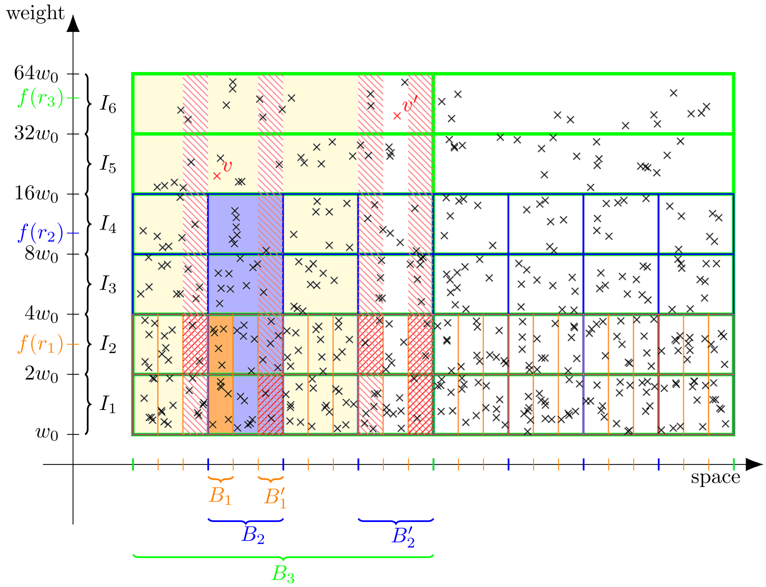

Figure 4. *Hyperrectangle-cover and definition of i-good boxes. In this figure, $d = 1$, $R = 3$, $r'_2/r'_1 = 3$, and $r'_3/r'_2 = 4$, and the requirement of (B1) for $i > 1$ is 'all but at most one sub-box $B' \in \mathcal{P}_{i-1}$ of B is good'. The hyperrectangle-cover is denoted by coloured-boundary rectangles. The spatial dimension on the x axis is covered by nested intervals, where (blue) boxes in $\mathcal{P}_2$ contain 3 level-1 (orange) boxes and (green) boxes in $\mathcal{P}_3$ contain 4 level-2 boxes. The weight dimension on the y axis is covered by a base-2-cover $I_1, \dots, I_6$. Hyperrectangles above $f(r_1)$ (e.g., $B_1 \times I_3$) and above $f(r_2)$ (e.g., $B_2 \times I_5$), are not included in $\mathcal{P}_1, \mathcal{P}_2$, respectively, since they contain too few vertices for concentration. Good boxes are shaded and bad boxes are hatched or get no colour. Box B_1 is good because its two hyper-rectangles $B_1 \times I_1$ and $B_1 \times I_2$ (filled orange) contain the right number of vertices, making all vertices in B_1 2-good, including those with weights above $I_1 \cup I_2$. Box B'_1 is bad (light hatching), since it contains too few vertices in $B_1 \times I_1$ (cross-hatching). Box B_2 is good, because it only contains one bad sub-box (B'_1) in $\mathcal{P}_1$, and because its four hyperrectangles $B_2 \times I_1, \dots, B_2 \times I_4$ (filled blue) all contain the right number of 2-good vertices in total. Since B_1 and B_2 are both good, vertex v is 3-good. Box B'_2 is 'doubly' bad (filled white): it contains two level-1 bad sub-boxes, and the hyperrectangle $B'_2 \times I_3$ contains too few 1-good vertices. Thus no vertex in B'_2 is 3-good, including v'. Still, B_3 is 3-good: it contains enough 3-good vertices in total, and only one bad level-2 sub-box (B'_2).*

The rest of this section is devoted to proving Proposition 2.5. All remaining definitions and lemmas are used only within this section.

Shortly we shall carry out a multiscale analysis. We abbreviate $f_{R,\delta} := f$ everywhere except in definitions and statements. We partition $Q \times [w_0, f(r_R)]$ into hyper-rectangles. On the weight-coordinate, we cover the interval $[w_0, f(r_R)]$ of weights with a set of disjoint intervals $(I_j)_{j=1,\dots,j_{\max}}$ so that the first interval is of length w_0, and each consecutive interval is twice as long as the previous one. On the space-marginal, we partition Q into nested boxes B. The side lengths of these nested boxes will be close to $r_1, \dots, r_R$, with some minor perturbation so that they can form a proper nested partition: we write $r'_i \approx r_i$ for the side length of the i-th level of boxes. A depiction and extended example can be found in Figure 4 on page 23 below, after the formal definition.

After fixing this partitioning of $Q \times [w_0, f(r_R)]$, we look at $\widetilde{\mathcal{V}} \cap (Q \times [w_0, f(r_R)])$. For each $i \in [R]$, we show that with probability close to 1 there is a dense subset of 'good' boxes B of side length r'_i, in the sense that $B \times I_j$ contains the right number of vertices for all I_j with $\max(I_j) \approx f(r_i)$. We choose $r'_i < r_i/\sqrt{d}$ to ensure that for all vertices v in a box of side length r'_i, the entire box will be contained in $B_{r_i}(v)$ – the ball of radius r_i around v – this will allow us to take the net to be the set of all vertices which lie in good boxes of all side lengths $r'_1, \ldots, r'_R$. We start with the space marginal and now formally define the nested boxes.

Definition 2.6. Given $\mathcal{R} = \{r_1, \ldots, r_R\}$, and Q as in Proposition 2.5, an $\mathcal{R}$*-partition* of Q is a collection of partitions $\widehat{\mathcal{P}}(\mathcal{R}) := \{\mathcal{P}_1, \ldots, \mathcal{P}_R\}$ of Q into boxes with the following properties:

(P1) For all $i \in [R]$, every box in $\mathcal{P}_i$ has the same side length r'_i with

$$r'_i \in [r_i/(2\sqrt{d}), r_i/\sqrt{d}]. \tag{2.12}$$

(P2) For all $i \in [R-1]$, every box in $\mathcal{P}_{i+1}$ is partitioned into exactly $(r'_{i+1})^d/(r'_i)^d$ boxes in $\mathcal{P}_i$.
(P3) We have $\mathcal{P}_R = \{Q\}$.

For $x \in Q, i \in [R]$, write $B^i(x)$ for the box in $\mathcal{P}_i$ containing x. A direct consequence of *(P1)* of this definition is

$$B^i(x) \subseteq B_{r_i}(x). \tag{2.13}$$

The partition $\mathcal{P}_i$ is, by *(P2)*, a refinement of the partition of $\mathcal{P}_{i+1}$, that is, every box in $\mathcal{P}_{i+1}$ can be partitioned exactly into sub-boxes in $\mathcal{P}_i$. Also, $r_R = \xi\sqrt{d}$ ensures that *(P1)* and *(P3)* can be both satisfied for $i = R$.

Lemma 2.7. *Suppose $\mathcal{R}$ and Q satisfy Proposition 2.5. Then an $\mathcal{R}$-partition $\widehat{\mathcal{P}}(\mathcal{R})$ of Q exists.*

Proof. We prove that given a δ-well-spaced $r_1 < \ldots < r_R$, there exist side lengths $r'_1, \ldots, r'_R$ that satisfy *(P1) – (P3)*, that is, that r'_{i+1}/r'_i is an integer, (2.12) holds, and $r'_R = \xi$. We proceed by induction on i, starting from $i = R$ and decreasing i. We take $r'_R := r_R/\sqrt{d} = \xi$, then (2.12) is satisfied immediately and *(P2)–(P3)* are vacuous. Suppose we have found $r'_i, \ldots, r'_R$ satisfying *(P1)–(P3)* for some $2 \le i \le R$. Let

$$r'_{i-1} = \frac{r'_i}{\lceil \sqrt{d} r'_i / r_{i-1} \rceil}. \tag{2.14}$$

This choice of r'_{i-1} divides r'_i, hence *(P2)* can be satisfied, and $r'_{i-1} \le r'_i/(\sqrt{d}r'_i/r_{i-1}) = r_{i-1}/\sqrt{d}$. Moreover,

$$r'_{i-1} \ge \frac{r'_i}{1 + \sqrt{d} r'_i/r_{i-1}} = \frac{r_{i-1}}{r_{i-1}/r'_i + \sqrt{d}}. \tag{2.15}$$

Since (2.12) holds for i (by the inductive assumption), $r'_i \ge r_i/2\sqrt{d}$. Since $\mathcal{R}$ is also well-spaced, $r_{i-1} \le r_i/2$ by (2.6); hence $r_{i-1}/r'_i \le \sqrt{d}$. so, by (2.15), $r'_{i-1} \ge r_{i-1}/2\sqrt{d}$, and so (2.12) holds also for $i-1$ and the induction is advanced.

Given these $r'_1, \ldots, r'_R$, we find an $\mathcal{R}$-partition of Q by taking $\mathcal{P}_R = \{Q\}$ and iteratively forming each layer $\mathcal{P}_{i-1}$ by taking the unique partition of each box in $\mathcal{P}_i$ into $(r'_i)^d/(r'_{i-1})^d$ sub-boxes of side length r'_{i-1}. We first define each partition box to be of the form $\prod_{j=1}^d [a_j, b_j)$, this allocates each point except d of the $d-1$-dim faces of ∂Q uniquely. Finally, we allocate the points $x \in \partial Q$ in $\mathcal{P}_i$ to the box in $\mathcal{P}_i$ that contains x in its closure, this box is unique except on $d-2$ dimensional faces. Here we again use half-open $d-1$-dim boxes to determine the $d-2$ dim boundaries, and so on until only the corner-points are left which we allocate arbitrarily (but consistently across different i). □

We continue with the weight-marginal and cover the interval $[w_0, f(r_i)]$ of weights with a collection of intervals, forming later the weight-coordinate of the hyper-rectangles:

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

Definition 2.8 (Base-2-cover). Given a closed interval $J = [a, b] \subset \mathbb{R}_+$, let $j_{\max} := \lfloor \log_2(b/a) \rfloor + 1$, and define $I_j := [2^{j-1}a, 2^j a)$ for $j \in [j_{\max}]$. Then $J \subseteq \bigcup_{j=1}^{j_{\max}} I_j$ and we call $I = \{I_j\}_{j \le j_{\max}}$ the *base-2-cover* of $[a, b]$. For each $x \in [a, b]$ we define $I(w) := \{I_j : w \in I_j\}$ to be the unique interval that contains x, and write $I(w) = [w_-, w_+)$ for its endpoints.

For each $w \in [a, b]$, $w/2 \le w_-$ and $w_+ \le 2x$, so $I(w) \subseteq [w/2, 2w]$; by the definition of $I_{j_{\max}}$, we also have $b \in I_{j_{\max}}$. We define the hyperrectangle-covering of the box Q including vertex-weights now. Recall vertex-weight distribution W from (1.6), and $f(r)$ from Def. 2.3.

Definition 2.9 (Hyperrectangles). Consider the setting of Proposition 2.5 and Definitions 2.6, 2.8. Let $\widehat{\mathcal{P}}(\mathcal{R}) := \{\mathcal{P}_1, \dots, \mathcal{P}_R\}$ be an $\mathcal{R}$-partition of the cube Q with $\mathcal{R} = \{r_1, \dots, r_R\}$, r_i' be the side-lengths in $\mathcal{P}_i$, and let $I = \{I_j\}_{j \le j_{\max}}$ be a base-2-cover of $[w_0, f_{R,\delta}(r_R)]$. Let $j_\star(i)$ be the index of the interval that contains $f_{R,\delta}(r_i)$, that is,

$$f_{R,\delta}(r_i) \in I(f_{R,\delta}(r_i)) =: I_{j_\star(i)}. \tag{2.16}$$

Then we say that the collection $\mathcal{H}(\mathcal{R}) := \big\{B_i \times I_j : B_i \in \mathcal{P}_i, 1 \le j \le j_\star(i)\big\}$ is a *hyperrectangle-cover of* $Q \times [w_0, f(r_R)]$. For all $i \in [R]$ and all $A \subset [w_0, f(r_R)]$, we define

$$\mu_i(A) := (r_i')^d \cdot \mathbb{P}(W \in A). \tag{2.17}$$

When we cover with boxes in $\mathcal{P}_i$ on the spatial coordinate, the number $j_\star(i)$ of weight intervals in $\mathcal{H}(\mathcal{R})$ depends on i. In particular, for smaller side-length we do not include intervals of very large weights. This is because there are too few (or no) vertices of large weight in a typical box of small side-length, so we cannot control their number. We illustrate a hyperrectangle cover on Figure 4 in dimension 1. In IGIRG, $\mu_i(A)$ is the expected number of vertices with weights in A in any box in $\mathcal{P}_i$. In SFP, $\mu_i(A)$ is only roughly the expectation, since, for example, a box touching the boundary ∂Q in $\mathcal{P}_i$ may not contain exactly $(r_i')^d$ vertices of $\mathbb{Z}^d$. By (2.13), all vertices in the hyperrectangle $B^i(v) \times I_j$ are within distance r_i of v. Hence once we control the number of vertices in a dense set of hyperrectangles in all partitions $i \in [R]$, we can find a net. We now define a hyperrectangle being 'good', with respect to a realisation of $\widetilde{\mathcal{V}}$. Recall that $I(w)$ denotes the interval I_j that contains w in Definition 2.8.

Definition 2.10. Consider the setting of Def. 2.13, and let a $\mathcal{H}(\mathcal{R})$ be a hyperrectangle-cover of $Q \times [w_0, f(r_R)]$. Consider a realisation of the weighted vertex set $\widetilde{\mathcal{V}} = \big((v, w_v)\big)_{v \in \mathcal{V}}$.

We recursively define when we call a vertex $\widetilde{v} \in \widetilde{\mathcal{V}}$ and a box $B \in \widehat{\mathcal{P}}(\mathcal{R})$ good. Every vertex is 1-good. For all $i \in [R]$, we say a vertex $\tilde{v} = (v, w_v) \in \widetilde{\mathcal{V}}$ is *i-good* if the boxes $B^1(v), \dots, B^{i-1}(v)$ are all good (which we define next). Denote the set of i-good (weighted) vertices by $\widetilde{\mathcal{G}}_i := \{\widetilde{v} \in \widetilde{\mathcal{V}}\ i\text{-good}\}$ and $\mathcal{G}_i := \{v : \widetilde{v} \in \widetilde{\mathcal{G}}_i\}$. We say that a box $B \in \mathcal{P}_i$ is *i-good* or simply good if the following conditions all hold:

(B1) Either $i = 1$, or B contains at least $1 - \frac{2\delta}{R} \cdot (r_i'/r_{i-1}')^d$ many $i-1$-good sub-boxes $B' \in \mathcal{P}_{i-1}$.
(B2) The total number of i-good vertices in B satisfies

$$|\mathcal{G}_i \cap B| \in \left[\left(\frac{1}{2} - \frac{2(i-1)\delta}{R}\right)(r_i')^d, 2(r_i')^d\right]. \tag{2.18}$$

(B3) For all $w \in [w_0, f_{R,\delta}(r_i)]$, the number of i-good vertices in B with weight in $I(w)$ satisfies

$$|\widetilde{\mathcal{G}}_i \cap (B \times I(w))| \in \left[\frac{1}{8}\left(1 - \frac{2i\delta}{R}\right)\mu_i(I(w)), 8\mu_i(I(w))\right]. \tag{2.19}$$

We say that the realisation $\widetilde{\mathcal{V}}$ is *good* with respect to the hyperrectangle-cover $\widetilde{\mathcal{P}}(\mathcal{R})$ if Q is R-good.

The above definition is *not circular*; the definition of i-good vertices depends only on the definition of good boxes in $\mathcal{P}_{i-1}$, that is, one level lower, and then the definition of a good box in $\mathcal{P}_i$ depends only on the number of i-good vertices in it (and their weights) and its number of good subboxes in $\mathcal{P}_{i-1}$. For $i = 1$, the (longer) definition of 1-good vertices is vacuous, so every vertex is indeed 1-good which we emphasised in the definition. Further, $\mathcal{G}_R \subseteq \mathcal{G}_{R-1} \subseteq \ldots \subseteq \mathcal{G}_1 = \mathcal{V} \cap Q$, since each i-good vertex is also $(i-1)$-good for all $i \le R$. See Figure 4 for a graphical depiction of i-good vertices and boxes.

Before we relate goodness to our overall goal of finding an $(\delta, \mathcal{R})$-net, we give some easy algebraic bounds which we need multiple times in the rest of the section. Recall that $I(w) = I_j$ iff $w \in I_j$ (cf. Def. 2.8) and that $\mu_i(A) = (r_i')^d \cdot \mathbb{P}(W \in A)$ in (2.17).

Lemma 2.11. *Consider the setting of Proposition 2.5 and Definitions 2.6, 2.13. Suppose $\widehat{\mathcal{P}}(\mathcal{R}) = \{\mathcal{P}_1, \ldots, \mathcal{P}_R\}$ is an $\mathcal{R}$-partition of Q, and for all $i \in [R]$, let r_i' be the side length of boxes in $\mathcal{P}_i$. Then for all $i \in [R]$ and all $w \in [w_0, f(r_R)]$, we have*

$$r_i^d \ell(w) w^{-(\tau-1)}/(2d)^{\tau+d+1} \le \mu_i(I(w)) = (r_i')^d \cdot \mathbb{P}(W \in I(w)) \le 2^\tau r_i^d \ell(w) w^{-(\tau-1)}, \tag{2.20}$$

$$r_i^d \ell(f(r_i)) f(r_i)^{-(\tau-1)} \ge (2d)^{2\tau+d+8} \log(16R/\delta). \tag{2.21}$$

Proof. We start by showing (2.20). We write $I(w) = [w_-, w_+)$. The definition of a base-2-cover (Definition 2.8) ensures that $w_-, w_+ \in [w/2, 2w]$ and $w_+/w_- = 2$. Thus by the lower bound (2.1) on w_0, for all $w \in [w_0, f(r_R)]$,

$$\begin{aligned}\mathbb{P}(W \in I(w)) &= \ell(w_-) w_-^{-(\tau-1)} - \ell(w_+) w_+^{-(\tau-1)} \\ &\ge \ell(w)\left(\frac{99}{100} w_-^{-(\tau-1)} - \frac{101}{100} w_+^{-(\tau-1)}\right) \\ &= \ell(w) w_-^{-(\tau-1)}\left(\frac{99}{100} - \frac{101}{100}\cdot 2^{-(\tau-1)}\right).\end{aligned}$$

Since $\tau > 2$ and $w_- \le w$, it follows that $\mu_i(I(w)) \ge (r_i')^d \ell(w) w^{-(\tau-1)}/2^{\tau+1}$. The required lower bound on $\mu_i(I(w))$ then follows by the lower bound *(P1)* in Definition 2.6 on r_i'. By a very similar argument to the lower bound, (2.1) also implies that

$$\mathbb{P}(W \in I(w)) \le \frac{101}{100} \ell(w) w_-^{-(\tau-1)} \le 2^\tau \ell(w) w^{-(\tau-1)},$$

so the required upper bound on $\mu_i(I(w))$ likewise follows by the upper bound *(P1)* on r_i'. It remains to show (2.21). First we show that $f(r_R) \ge f(r_{R-1}) \ge \ldots \ge f(r_1) \ge w_0$. Indeed, let $j \in [R]$. By the definition of f in (2.3),

$$\frac{f(r_j)}{f(r_{j-1})} = \left(\frac{r_j}{r_{j-1}}\right)^{d/(\tau-1)} \cdot \left(\frac{1 \wedge \inf\left\{\ell(x) \colon x \in [w_0, r_j^{d/(\tau-1)}]\right\}}{1 \wedge \inf\left\{\ell(x) \colon x \in [w_0, r_{j-1}^{d/(\tau-1)}]\right\}}\right)^{1/(\tau-1)}. \tag{2.22}$$

To bound the second factor, we will first observe that since $r_{j-1} \le r_j$, we have $\inf\{\ell(x) \colon x \in [w_0, r_j^{d/(\tau-1)}]\} \le \inf\{\ell(x) \colon x \in [w_0, r_{j-1}^{d/(\tau-1)}]\}$. By considering the two possible values of $1 \wedge \inf\{\ell(x) \colon x \in [w_0, r_j^{d/(\tau-1)}]\}$ separately, it follows that we can drop the minimum with 1 in the ratio:

$$\left(\frac{1 \wedge \inf\left\{\ell(x) \colon x \in [w_0, r_j^{d/(\tau-1)}]\right\}}{1 \wedge \inf\left\{\ell(x) \colon x \in [w_0, r_{j-1}^{d/(\tau-1)}]\right\}}\right)^{1/(\tau-1)} \ge \left(\frac{\inf\left\{\ell(x) \colon x \in [w_0, r_j^{d/(\tau-1)}]\right\}}{\inf\left\{\ell(x) \colon x \in [w_0, r_{j-1}^{d/(\tau-1)}]\right\}}\right)^{1/(\tau-1)}. \tag{2.23}$$

We now bound $\ell(x)$ on $[w_0, r_{j-1}^{d/(\tau-1)}]$ (in the denominator) by repeatedly applying (2.1). We write $\lg(\cdot) = \log_2(\cdot)$. Then, we have $r_j^{d/(\tau-1)} = 2^{\lg((r_j/r_{j-1})^{d/(\tau-1)})} r_{j-1}^{d/(\tau-1)}$, so we iterate the bound in (2.1) roughly $\lg((r_j/r_{j-1})^{d/(\tau-1)})$ many times to obtain that for all $x \in [r_{j-1}^{d/(\tau-1)}, r_j^{d/(\tau-1)}]$ we have

$$\ell(x) \geq \left(\frac{99}{100}\right)^{1+\lg((r_j/r_{j-1})^{d/(\tau-1)})} \ell(r_{j-1}).$$

Returning to (2.23), the ratio of the two infima in the smaller interval $[w_0, r_{j-1}^{d/(\tau-1)}]$ is one. And since $r_j \geq 2r_{j-1}$ by (R2) of Definition 2.4 (using $\delta < 1/16 < 1$), it follows that

$$\left(\frac{\inf\left\{\ell(x) : x \in [w_0, r_j^{d/(\tau-1)}]\right\}}{\inf\left\{\ell(x) : x \in [w_0, r_{j-1}^{d/(\tau-1)}]\right\}}\right)^{1/(\tau-1)} \geq \left(\frac{99}{100}\right)^{1+\lg((r_j/r_{j-1})^{d/(\tau-1)})}$$
$$\geq \left(\frac{1}{2}\right)^{\frac{1}{\tau-1}\lg((r_j/r_{j-1})^{d/(\tau-1)})} = \left(\frac{r_{j-1}}{r_j}\right)^{d/(\tau-1)^2}.$$

Combining this with (2.22), (2.23), and the fact that $\tau > 2$, we obtain

$$\frac{f(r_j)}{f(r_{j-1})} \geq \left(\frac{r_j}{r_{j-1}}\right)^{\frac{d}{\tau-1}(1-1/(\tau-1))} \geq 1.$$

Hence $f(r_R) \geq f(r_{R-1}) \geq \ldots \geq f(r_1)$, as claimed. It is now relatively easy to prove the desired lower bound. From the definition of f in (2.3), we have $f(r_i) \leq r_i^{d/(\tau-1)}$, and (2.5) in the definition of well-spacedness ensures that $f(r_i) \geq w_0$, so

$$\ell(f(r_i))f(r_i)^{-(\tau-1)} = \frac{\ell(f(r_i))}{1 \wedge \inf\{\ell(x) : x \in [w_0, r_i^{d/(\tau-1)}]\}} \cdot \frac{(2d)^{2\tau+d+8}\log(16R/\delta)}{r_i^d}$$
$$\geq \frac{(2d)^{2\tau+d+8}\log(16R/\delta)}{r_i^d}.$$

Multiplying by r_i^d finishes the statement of (2.21). □

We now show that given that the box Q is good with respect to a hyperrectangle cover, we can find an $(\delta, \mathcal{R})$-net for Q (see Def. 2.3).

Lemma 2.12. *Consider the setting of Proposition 2.5 and Definitions 2.6, 2.13, 2.10, that is, a hyperrectangle cover $\mathcal{H}(\mathcal{R})$ of $Q \times [w_0, f(r_R)]$, and a realisation of $\widetilde{\mathcal{V}}$ for which Q is R-good. Then $\widetilde{\mathcal{G}}_R$, the set of all R-good vertices, forms an $(\delta, \mathcal{R})$-net for Q.*

Proof. Suppose that Q is R-good. The side length of Q equals r'_R by *(P3)* in Def. 2.6, and $\delta \in (0, 1/16)$ in the setting of Proposition 2.5, hence we may apply *(B2)* in Def. 2.10 for $i = R$ to get

$$|\widetilde{\mathcal{G}}_R| \geq \left(\frac{1}{2} - \frac{2(R-1)\delta}{R}\right)\mathrm{Vol}(Q) \geq \left(\frac{1}{2} - 2\delta\right)\mathrm{Vol}(Q) > \mathrm{Vol}(Q)/4,$$

hence the cardinality assumption in Definition 2.3 is satisfied for $\widetilde{\mathcal{G}}_R$. Recall that $B^i(v)$ is the box in $\mathcal{P}_i$ containing v. To show that $\widetilde{\mathcal{G}}_R$ satisfies Definitions 2.3, we first show that for all $v \in \mathcal{G}_R$,

$$\mathcal{G}_i \cap B^i(v) = \mathcal{G}_R \cap B^i(v). \tag{2.24}$$

To show this, we need to show that every i-good vertex $u \in B^i(v)$ is also R-good, that is, that $B^j(u)$ is good for all $j \in [R]$. By the definition of i-goodness (Def. 2.10), $B^j(u)$ is good for all $j \leq i-1$, and $B^R(u) = Q$ is good by hypothesis. Consider now a $j \in [i+1, R-1]$. By Def. 2.6, the partition $\mathcal{P}_i$ is a

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

refinement of the partition $\mathcal{P}_j$, so $B^j(u) = B^j(v)$. Since v is R-good, it follows that $B^j(u)$ is good. So, $B^j(u)$ is good for all $j \in [R]$, so u is R-good, showing (2.24).

With $I(w) \subseteq [w/2, 2w]$ being the interval containing w in the base-2-cover of $[w_0, f(r_i)]$ in Def.2.8, we now argue that

$$|\widetilde{\mathcal{G}}_R \cap (B_{r_i}(v) \times [w/2, 2w])| \geq |\widetilde{\mathcal{G}}_R \cap (B^i(v) \times I(w))| = |\widetilde{\mathcal{G}}_i \cap (B^i(v) \times I(w))|.$$

Indeed, $B^i(v) \subseteq B_{r_i}(v)$ by (2.13) and $I(w) \subseteq [w/2, 2w]$ by Def. 2.8, and since all i-good vertices in $B^i(v)$ are also R-good by (2.24), the statement follows. Now we apply, on the right-hand side $|\widetilde{\mathcal{G}}_i \cap (B^i(v) \times I(w))|$ above, the lower bound (2.19) from Def. 2.10 *(B3)*, then (2.20) to obtain

$$\begin{aligned}\left|\widetilde{\mathcal{G}}_R \cap (B^i(v) \times [w/2, 2w])\right| &\overset{(2.19)}{\geq} \frac{1}{8}\left(1 - \frac{2i\delta}{R}\right)\mu_i(I(w)) \\ &\overset{(2.20)}{\geq} \frac{1}{8}\left(1 - \frac{2i\delta}{R}\right) r_i^d \ell(w) w^{-(\tau-1)}/(2d)^{\tau+d+1}.\end{aligned}$$

Observing that $\delta < 1/16$ and $i \leq R$ ensures that the prefactor on the right-hand side of the last row is at least $1/8 \cdot 1/2 = 1/2^4$, establishing (2.4) for all $w \leq f(r_i)$, as required. □

A lower bound on the probability that any given box in an $\mathcal{R}$-partition is good, together with Lemma 2.7 and Lemma 2.12, will yield the proof of Proposition 2.5. The bound is by induction on i together with Chernoff bounds. Recall I and $I(w)$ from Def. 2.8, applied to the interval $[w_0, f(r_R)]$ for $\mathcal{R} = \{r_1, \ldots, r_R\}$. Recall that (2.19) of Def. 2.10 *(B3)* is required only when $w \in [w_0, f(r_i)]$, and that $j_\star(i)$ in (2.16) is the index of I_j that contains $f(r_i)$. We now gradually reveal vertex-weights by giving a weight-revealment scheme. Here, we aim for unified proof that works for SFP and IGIRG simultaneously. In IGIRG we can make use the independence property of marked Poisson point processes, namely, that the number of vertices in $B \times I_j$ and $A \times I_{j'}$ are independent if $j \neq j'$, see [61]. However, in SFP this is not the case and we may run out of vertices as we gradually reveal vertices with higher and higher weights. To solve this issue, we first reveal all vertex positions only (one could think of this as a sigma-algebra, we call it $\mathcal{F}_0$ and leave vertex-weights yet unrevealed. Then, in subsequent weight-revealment steps $\mathcal{F}_1, \mathcal{F}_2, \ldots, \mathcal{F}_R$, we reveal the vertex-weights of those vertices whose weight falls into the intervals I_j, with the maximal j revealed increasing along the steps. Recall that $j_\star(i)$ is the index of the interval I_j containing $f(r_i)$ from (2.16), and that we denote the σ-algebra generated by a random variables X, Y by $\sigma(X, Y)$.

Definition 2.13. Consider the setting of Proposition 2.5 and Definitions 2.6, 2.13. Suppose $\mathcal{H}(\mathcal{R})$ is a hyperrectangle-cover of $Q \times [w_0, f(r_R)]$. Let $\widetilde{\mathcal{V}}$ be a realisation of the weighted vertices in Definition 1.3 and B be a box in one of the partitions $(\mathcal{P}_i)_{i \leq R}$. We define

$$\mathcal{F}_i(B) := \begin{cases} \sigma(\mathcal{V} \cap B) & \text{for } i = 0, \\ \sigma\Big(\mathcal{F}_0(B), \widetilde{\mathcal{V}} \cap (B \times \cup_{j \leq j_\star(i)} I_j\Big) & \text{for all } 1 \leq i \leq R. \end{cases} \tag{2.25}$$

We say that we vertices with weight in $\cup_{j \leq j_\star(i)} I_j$ have 'revealed' weight in $\mathcal{F}_i$, while the vertices in $[1, \infty) \setminus \cup_{j \leq j_\star(i)} I_j$ have 'unrevealed' weight in $\mathcal{F}_i$. We denote by $\mathcal{X}_i(B)$ the outcome of the revealment scheme $\mathcal{F}_i(B)$: $\mathcal{X}_i(B)$ contains the position of all vertices in the box B, where a subset of these vertices is marked by the revealed vertex-weight.

$\mathcal{F}_0(B)$ reveals the number and positions of vertices in B, while $\mathcal{F}_{i-1}(B)$ reveals the precise weights *only* of vertices whose weight falls in the interval $\cup_{j \leq j_\star(i-1)} I_j \supseteq [w_0, f(r_{i-1})]$. Note that the weight distribution of vertices with unrevealed weight *changes* along the revealment procedure, since after revealment step $i - 1$ it is the conditional weight distribution that the weight does not fall in the revealed interval $\cup_{j \leq j_\star(i-1)} I_j$. The fact that $\mathcal{R}$ is δ-well-spaced, means that vertex weights between $w_0 2^{j_\star(i-1)} \approx f(r_{i-1})$ and $f(r_i)$ are not revealed in $\mathcal{F}_{i-1}(B)$. Also, vertex weights in $[1, w_0)$ are not revealed at all; since w_0 is large, $\mathbb{P}(W \le w_0)$ is large and most vertex weights will not be revealed by exposing $\mathcal{F}_{i-1}(B)$. The filtration generated by $\mathcal{F}_{i-1}(B)$ thus determines whether or not boxes in $\cup_{j\le i-1}\mathcal{P}_j$ are good, and whether or not a *vertex* is i-good (see Def. 2.10), since i-goodness of a vertex only depends on boxes with index at most $i-1$ containing the vertex. So, $\mathcal{F}_{i-1}(B)$ determines whether or not $B \in \mathcal{P}_i$ satisfies Def. 2.10 *(B1) -(B2)*, but it leaves *(B3)* undecided for weights slightly below $f(r_i)$. The next lemma treats *(B3)*, with $\mathcal{F}_{i-1}(B)$ exposed.

Lemma 2.14. *Consider the setting of Proposition 2.5 and Definitions 2.6, 2.13. Let $\mathcal{H}(\mathcal{R})$ be a hyperrectangle-cover of the cube Q. Let $i \in [R]$ and let $B \in \mathcal{P}_i$. Let X be a realisation[9] of $\mathcal{X}_{i-1}(B)$ that satisfies Definition 2.10 (B1) and (B2) for B. Then independently of other boxes in $\mathcal{P}_i$, uniformly for all such X,*

$$\mathbb{P}(B \textit{ is good} \mid \mathcal{X}_{i-1}(B) = X) \ge 1 - \delta/(2R). \tag{2.26}$$

Proof. Recall I_j from Def. 2.8, and let $B \in \mathcal{P}_i$. Let

$$a(I_j) := (1 - 2i\delta/R)\mu_i(I_j)/8, \qquad b(I_j) := 8\mu_i(I_j) \tag{2.27}$$

be the required lower and upper bounds in (2.19). Let $\Xi_j(B) = |\widetilde{\mathcal{G}}_i \cap (B \times I_j)|$; thus *(B3)* holds for $B \times I_j$ iff $\Xi_j(B) \in [a(I_j), b(I_j)]$. Since B satisfies *(B1)–(B2)* on X, by a union bound,

$$\mathbb{P}\big(B \text{ is good} \mid \mathcal{X}_{i-1}(B) = X\big) \ge 1 - \sum_{I_j : j \le j_\star(i)} \mathbb{P}\big(\Xi_j(B) \notin [a(I_j), b(I_j)] \mid \mathcal{X}_{i-1}(B) = X\big). \tag{2.28}$$

We proceed by bounding each term above. By the definition of $\mathcal{F}_{i-1}(B)$ in (2.25), we already exposed $\Xi_j(B)$ when $i > 1$ and $j \le j_*(i-1)$; the latter is equivalent to $\min(I_j) \le f(r_{i-1})$.

Case 1: $i > 1$ and $j \le j_\star(i-1)$. We first show that $\Xi_j \ge a(I_j)$ holds deterministically on $\{\mathcal{X}_{i-1}(B) = X\}$. The goodness of each sub-box $B' \in \mathcal{P}_{i-1}$ of $B \in \mathcal{P}_i$ *is revealed* by $\mathcal{F}_{i-1}(B)$. If $B' \in \mathcal{P}_{i-1}$ is a good box, all vertices in $\mathcal{G}_{i-1} \cap B'$ are also i-good by Definition 2.10. So

$$\Xi_j(B) = |\widetilde{\mathcal{G}}_i \cap (B \times I_j)| \ge \sum_{B' \in \mathcal{P}_{i-1}:\, B' \text{ good}} |\widetilde{\mathcal{G}}_{i-1} \cap (B' \times I_j)|. \tag{2.29}$$

Since $j \le j_\star(i-1)$, $\min(I_j) \le f(r_{i-1})$, so we may apply (2.19) to the good subboxes:

$$|\widetilde{\mathcal{G}}_{i-1} \cap (B' \times I_j)| \ge \left(1 - \frac{2(i-1)\delta}{R}\right)\frac{\mu_{i-1}(I_j)}{8} = \left(1 - \frac{2(i-1)\delta}{R}\right)\frac{\mu_i(I_j)}{8}\left(\frac{r'_{i-1}}{r'_i}\right)^d, \tag{2.30}$$

where $\mu_i(I_j)/\mu_{i-1}(I_j) = (r'_i/r'_{i-1})^d$ follows from (2.17). By *(B1)* holding on X, B contains at least $(1 - 2\delta/R)(r'_i/r'_{i-1})^d$ good sub-boxes in $\mathcal{P}_{i-1}$. Combining that with (2.29)–(2.30) and (2.27) yields

$$\Xi_j(B) \ge \left(1 - \frac{2\delta}{R}\right) \cdot \left(1 - \frac{2(i-1)\delta}{R}\right)\frac{\mu_i(I_j)}{8} \ge \left(1 - \frac{2i\delta}{R}\right)\frac{\mu_i(I_j)}{8} = a(I_j).$$

We show $\Xi_j(B) = |\widetilde{\mathcal{G}}_i \cap (B \times I_j)| \le b(I_j)$ also holds a.s. B contains $(r'_i/r'_{i-1})^d$ sub-boxes in $\mathcal{P}_{i-1}$ by Def. 2.6 *(P2)*. If a sub-box is bad, it contains no i-good vertices. If it is good, *(B3)* holds and it contains at most $8\mu_i(I_j)(r'_{i-1}/r'_i)^d$ i-good vertices with weights in $B \times I_j$. We obtain $\Xi_j(B) \le 8\mu_i(I_j) = b(I_j)$ in (2.27). So overall we have shown that

$$\text{if } i > 1 \text{ and } j \le j_\star(i-1): \quad \mathbb{P}\big(\Xi_j(B) \notin [a(I_j), b(I_j)] \mid \mathcal{F}_{i-1}(B) = F\big) = 0. \tag{2.31}$$

[9]Thus here F contains the position of all vertices in the box B, where a subset of these vertices is marked by the revealed vertex-weight.

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

Case 2: $i > 1$ and $j > j_\star(i-1)$. Define the set

$$\mathcal{S} = \widetilde{\mathcal{G}}_i \cap \Big(B \times \big([1, w_0) \cup \bigcup\nolimits_{j > j_\star(i-1)} I_j\big)\Big), \tag{2.32}$$

the i-good vertices in B with weights *not revealed* by $\mathcal{F}_{i-1}(B)$. $\mathcal{F}_{i-1}(B)$ does reveal the positions of vertices in B, and all weighted vertices not in $\mathcal{S}$; thus $\mathcal{F}_{i-1}(B)$ reveals $|\mathcal{S}|$, and the position of vertices in $\mathcal{S}$. By Def. 1.3, vertex weights are independent of vertex positions and of each other. Further, for a vertex $v \in \widetilde{\mathcal{V}}$, conditioning on $\mathcal{F}_{i-1}(B)$ is equivalent to either exposing its weight (if $w_v \in \cup_{j \le j_\star(i-1)} I_j$) or conditioning on $w_v \notin \cup_{j \le j_\star(i-1)} I_j$ (otherwise). So for each vertex in $\mathcal{S}$, its weight distribution is the conditional distribution of W given that $W \notin \cup_{j \le j_\star(i-1)} I_j$. Since $|\mathcal{S}|$ is determined by $\mathcal{F}_{i-1}(B)$, for all $j \ge j_\star(i-1)$, $\Xi_j(B)$ given $\mathcal{F}_{i-1}(B)$ is binomially distributed with parameters

$$\Xi_j(B) \mid \mathcal{F}_{i-1}(B) \overset{d}{=} \mathrm{Bin}\big(|\mathcal{S}|, \mathbb{P}(W \in I_j)/\mathbb{P}(W \notin \cup_{j \le j_\star(i-1)} I_j)\big). \tag{2.33}$$

We next bound the conditional expectation of $\Xi_j(B) \mid \mathcal{F}_{i-1}(B)$, starting with the upper bound. Using the lower bound (2.1) on w_0, the success probability of the binomial in (2.33) is

$$\frac{\mathbb{P}(W \in I_j)}{\mathbb{P}(W \notin \cup_{j \le j_\star(i-1)} I_j)} \le \frac{\mathbb{P}(W \in I_j)}{\mathbb{P}(W \in [1, w_0))} = \frac{\mathbb{P}(W \in I_j)}{1 - \ell(w_0) w_0^{\tau-1}} \le 2\mathbb{P}(W \in I_j). \tag{2.34}$$

Since Def. 2.10 *(B2)* holds for B on $\mathcal{X}_{i-1}(B) = X$, by (2.18) there are at most $2(r_i')^d$ i-good vertices in B, so $|\mathcal{S}| \le 2(r_i')^d$. Recalling the definition of $\mu_i(I_j)$ from (2.17), and (2.27) we thus obtain

$$\mathbb{E}\big[\Xi_j(B) \mid \mathcal{X}_{i-1}(B) = X\big] \le 4(r_i')^d \mathbb{P}(W \in I_j) = 4\mu_i(I_j) = b(I_j)/2. \tag{2.35}$$

We next prove the corresponding lower bound. We start by giving a lower bound on $|\mathcal{S}|$. Clearly by (2.32), $|\mathcal{S}| = |\mathcal{G}_i \cap B| - |\cup_{j \le j_\star(i-1)} \widetilde{\mathcal{G}}_i \cap (B \times I_j)|$, both terms revealed by X: the total number of i-good vertices in B minus the ones with revealed weight. We can bound $|\mathcal{G}_i \cap B|$ from below using (2.18) in Def. 2.10 *(B2)*. By Def. 2.6 *(P2)*, B contains $(r_i'/r_{i-1}')^d$ sub-boxes in $\mathcal{P}_{i-1}$. Since there are no i-good vertices in bad boxes $B' \in \mathcal{P}_{i-1}$, we can bound $|\cup_{j \le j_\star(i-1)} \widetilde{\mathcal{G}}_i \cap (B \times I_j)|$ from above by applying (2.19) to each good sub-box $B' \in \mathcal{P}_{i-1}$ of B, yielding

$$|\mathcal{S}| \ge \left(\frac{1}{2} - \frac{2(i-1)\delta}{R}\right)(r_i')^d - \left(\frac{r_i'}{r_{i-1}'}\right)^d \cdot \sum_{j \le j_\star(i-1)} 8\mu_{i-1}(I_j). \tag{2.36}$$

Using that $I_j = [2^{j-1} w_0, w_0)$ for all $j \ge 1$, Lemma 2.11 with $w = 2^{j-1} w_0$ yields

$$\begin{aligned}(r_i'/r_{i-1}')^d \sum_{j \le j_\star(i-1)} 8\mu_{i-1}(I_j) &\le 2^{\tau+3}(r_i')^d \sum_{j=1}^{j_\star(i-1)} \ell(2^{j-1}w_0)(2^{j-1}w_0)^{-(\tau-1)} \\ &< 2^{\tau+3}(r_i')^d \sum_{j=0}^{\infty} \ell(2^j w_0)(2^j w_0)^{-(\tau-1)},\end{aligned} \tag{2.37}$$

where we switched indices in the last row. By the lower bound (2.1) on w_0 and since $\tau > 2$, for all $j \ge 0$ we have $\ell(2^{j+1}w_0)(2^{j+1}w_0)^{-(\tau-1)} < \frac{2}{3}\ell(2^j w_0)(2^j w_0)^{-(\tau-1)}$, so the sum on the right-hand side is bounded above term-wise by a geometric series. It follows from (2.36) that

$$|\mathcal{S}| \ge \left(\frac{1}{2} - \frac{2(i-1)\delta}{R} - 2^{\tau+5}\ell(w_0)w_0^{-(\tau-1)}\right)(r_i')^d \ge (r_i')^d/4,$$

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

where we used in the last step that $i \le R$ and $\delta < 1/16$, and the lower bound (2.1) on w_0. The success probability of the binomial in (2.33) is at least $\mathbb{P}(W \in I_j)$. Since $a(I_j) = (1-2i\delta/R)\mu_i(I_j)/8$ in (2.27),

$$\mathbb{E}\big[\Xi_j(B) \mid \mathcal{X}_{i-1}(B) = X\big] \ge (r_i')^d \mathbb{P}(W \in I_j)/4 = \mu_i(I_j)/4 \ge 2a(I_j). \tag{2.38}$$

Combining (2.35) with (2.38) yields that $\mathbb{E}[\Xi_j(B) \mid \mathcal{F}_{i-1}(B) = F] \in [2a(I_j), b(I_j)/2]$, which allows us to bound $\mathbb{P}(\Xi_j(B) \notin [a(I_j), b(I_j)])$ with standard Chernoff bounds. By (2.35) and Theorem A.1 applied with $\varepsilon = 1/2$, we have shown that uniformly for all realisations F satisfying *(B1)-(B2)*

$$\text{for } i > 1, j > j_\star(i-1): \quad \begin{aligned} \mathbb{P}\big(\Xi_j(B) \notin [a(I_j), b(I_j)] \mid \mathcal{F}_{i-1}(B) = F\big) &\le 2\exp(-a(I_j)/6) \\ &\le 2\exp(-\mu_i(I_j)/96). \end{aligned} \tag{2.39}$$

Case 3: $\boldsymbol{i=1}$. When $i = 1$, we set $j_\star(i-1) := 0$ naturally, since in $\mathcal{F}_0(B)$ we revealed the total number of vertices in $B \in \mathcal{P}_1$, which are all 1-good by Def. 2.10. Conditioned on $\mathcal{X}_0(B) = X$, (2.18) in *(B2)* is satisfied and $(r_1')^d/4 \le |\mathcal{S}| \le 2(r_1')^d$ directly. The rest of our previous calculations from Case 2 with $j > j_\star(0) = 0$ all carry through for estimating the left-hand side of (2.19) in *(B3)*. We obtain that the bound in (2.39) holds also for $i = 1$ and all $j \ge j_\star(0) = 0$.

Combining the cases: By (2.31), (2.39), and Case 3, for all i and $j \le j_\star(i)$ the bound in (2.39) holds. Combining that with (2.28), for all $B \in \mathcal{P}_i$ and $\mathcal{X}_{i-1}(B) = X$ satisfying *(B1),(B2)*,

$$\begin{aligned} p_i := \mathbb{P}\big(B \in \mathcal{P}_i \text{ is } i\text{-good} \mid \mathcal{X}_{i-1}(B) = X\big) &\ge 1 - 2\sum_{j \le j_\star(i)} \exp(-\mu_i(I_j)/96) \\ &\ge 1 - 2\sum_{j \le j_\star(i)} \exp\Big(-(2d)^{-(\tau+d+8)} r_i^d \ell(2^{j-1}w_0)(2^{j-1}w_0)^{-(\tau-1)}\Big), \end{aligned}$$

where the second line follows from $\mu_i(I_j) = \mu_i(I(2^{j-1}w_0))$ in Lemma 2.11. Here, the term in the sum with $j = j_\star(i)$ is maximal. We use then an reindexing $t = j_\star(i) - j, t \ge 0$ and argue that the sum can be dominated by a geometric series in t: By the lower bound (2.1) on w_0 and the fact that $\tau > 2$, for all $j \ge 1$ we have $\ell(2^{j-1}w_0)(2^{j-1}w_0)^{-(\tau-1)} \ge \frac{3}{2}\ell(2^j w_0)(2^j w_0)^{-(\tau-1)}$. Using the same geometric-sum argument as below (2.37) except now from the reversed viewpoint, writing $z := w_0 2^{j_\star(i)-1}$, the lower endpoint of $I_{j_\star(i)}$, we have

$$p_i \ge 1 - 2\sum_{t=0}^{\infty} \exp\bigg(-\frac{1}{(2d)^{\tau+d+8}} r_i^d \Big(\frac{3}{2}\Big)^t \ell(z) z^{-(\tau-1)}\bigg). \tag{2.40}$$

Since z is the lower endpoint of $I_{j_\star(i)} = I(f(r_i))$, we have $z \le f(r_i) \le 2z$ by (2.16). Hence by (2.1)

$$\ell(z) z^{-(\tau-1)} \ge 2^{-\tau}\ell(f(r_i)) f(r_i)^{-(\tau-1)}.$$

Using this bound in (2.40) and combining it with (2.21) from Lemma 2.11, we obtain that

$$p_i \ge 1 - 2\sum_{t=0}^{\infty} \exp\bigg(-\Big(\frac{3}{2}\Big)^t \log(16R/\delta)\bigg) \ge 1 - \frac{\delta}{8R} - 2\sum_{t=1}^{\infty}(\delta/16R)^t \ge 1 - \frac{\delta}{2R},$$

where we used that $(3/2)^t \ge t$ for all $t \ge 1$, and $\delta < 1/16$. Independence across boxes in the same $\mathcal{P}_i$ is immediate, since whether B satisfies *(B3)* conditioned on $\mathcal{F}_{i-1}(B) = F$ satisfying *(B1), (B2)* only depends on vertices in B. □

The next lemma gets rid of the conditioning in Lemma 2.14 on the filtration. As before, here in (2.41) below, $\mathbb{P}$ is here the Palm version of the annealed measure integrating over vertex-weights and positions for IGIRG, while for SFP the conditioning can be dropped.

Lemma 2.15. *Consider the setting of Proposition 2.5 and Definitions 2.6, 2.13. Let $\mathcal{H}(\mathcal{R})$ be a hyperrectangle-cover of the cube Q. Let $t \le (r_1/4\sqrt{d})^d$ be an integer, and for IGIRG let $x_1, \dots, x_t$ be a (possibly empty) sequence of points in $\mathbb{R}^d$. Then for each $B \in \cup_i \mathcal{P}_i$*

$$\mathbb{P}\big(B \text{ is good} \mid x_1, \dots, x_t \in \mathcal{V}\big) \ge 1 - \delta/R. \tag{2.41}$$

Proof. We prove the statement by induction on i, the base case being $i = 1$. Consider a box $B \in \mathcal{P}_1$. By Def. 2.10, *(B1)* holds vacuously. We next consider *(B2)*. Every vertex in B is 1-good, and so $|\mathcal{G}_1 \cap B| = |\mathcal{V} \cap B|$. In SFP, $\mathcal{V}$ is deterministic and $|\mathcal{G}_1 \cap B| \in [(r_1')^d/2, 2(r_1')^d]$ holds with certainty by Def. 2.6 *(P1)*. In (I)GIRG, $\mathcal{G}_1 \cap B$ is a Poisson point process and the Palm theory (see, e.g., [61]) gives that under the conditioning $x_1, \dots, x_t \in \mathcal{V}$, $|\mathcal{G}_1 \cap B| - t$ is a Poisson variable with mean $(r_1')^d$, so by a standard Chernoff bound (Theorem A.1 with $\varepsilon = 1/2$), and using $r_i' \in [r_i/(2\sqrt{d}), r_i/\sqrt{d}]$ in Def. 2.6 *(P1)*, and the bound (2.5) on r_1 in Def. 2.4,

$$p_1' := \mathbb{P}\Big(|\mathcal{G}_1 \cap B| \in [t + \tfrac{1}{2}(r_1')^d, t + \tfrac{3}{2}(r_1')^d] \mid x_1, \dots, x_t \in \mathcal{V}\Big) \ge 1 - 2\mathrm{e}^{-(r_1/24d)^d} \ge 1 - \delta/(2R),$$

so the lower bound in (2.18) in Def. 2.10 *(B2)* is satisfied. Moreover, since $t \le (r_1/4\sqrt{d})^d$, by Def. 2.6 *(P1)*, $\frac{3}{2}(r_1')^d + t \le 2(r_1')^d$, the upper bound in Def. 2.10 *(B2)* also holds, so independently for all boxes $B \in \mathcal{P}_1$, and regardless on where $x_1, \dots, x_t$ fall,

$$\mathbb{P}\big(B \in \mathcal{P}_1 \text{ satisfies } (B2) \mid x_1, \dots, x_t \in \mathcal{V}\big) \ge p_1' \ge 1 - \delta/(2R).$$

Lemma 2.14 ensures that *(B3)* holds uniformly with probability at least $1 - \delta/(2R)$ conditioned on any realisation where *(B1),(B2)* holds. A union bound proves (2.41) for $B \in \mathcal{P}_1$.

Now we advance the induction. Suppose that (2.41) holds for each $B \in \cup_{j \le i-1} \mathcal{P}_i$, and let $B \in \mathcal{P}_i$. B contains $(r_i'/r_{i-1}')^d$ sub-boxes in $\mathcal{P}_{i-1}$ (by Def. 2.6), and by induction these sub-boxes of B are good *independently* (regardless of the positions and weights of $x_1, \dots, x_t \in \mathcal{V}$), so the number of bad sub-boxes of B is binomial with mean at most $(r_i'/r_{i-1}')^d \delta/R$. Let

$$\mathcal{A}_i(B) := \Big\{\big|\{B' \in \mathcal{P}_{i-1} \colon B' \subseteq B,\ B' \text{ not } (i-1)\text{-good}\}\big| < 2(r_i'/r_{i-1}')^d \delta/R\Big\}. \tag{2.42}$$

Then, $\mathcal{A}_i(B)$ implies Def. 2.10 *(B1)*. A Chernoff bound (Theorem A.1 with $\varepsilon = 1$) yields that

$$\mathbb{P}\Big(\mathcal{A}_i(B)^{\complement} \mid x_1, \dots, x_t \in \mathcal{V}\Big) \le \exp\bigg(-\frac{\delta}{3R} \cdot \bigg(\frac{r_i'}{r_{i-1}'}\bigg)^d\bigg).$$

By Def. 2.6 *(P1)*, $(r_i'/r_{i-1}')^d \ge 2^{-d}(r_i/r_{i-1})^d$, so by Def. 2.4 (2.6), regardless of the positions of $x_1, \dots, x_t \in \mathcal{V}$, and *independently* across different boxes in $\mathcal{P}_i$:

$$\mathbb{P}\Big(\mathcal{A}_i(B)^{\complement} \mid x_1, \dots, x_t \in \mathcal{V}\Big) \le \exp\bigg(-\frac{\delta}{3R} \cdot 3^d R \cdot \frac{\log(2R/\delta)}{\delta}\bigg) \le \frac{\delta}{2R}. \tag{2.43}$$

We now show that $\mathcal{A}_i(B)$ implies *(B2)* as well, by inductively applying (2.18) to the good sub-boxes of B. Consider an $(i-1)$-good vertex v in a good sub-box $B' \in \mathcal{P}_{i-1}$ of B. Since v is $(i-1)$-good, $B^1(v), \dots, B^{i-2}(v)$ must all be good; since $B^{i-1}(v) = B'$ is also good, it follows that v is in fact i-good. Thus, for all good $B' \in \mathcal{P}_{i-1} : \mathcal{G}_{i-1} \cap B' = \mathcal{G}_i \cap B'$ holds. Since B' is $(i-1)$-good, it satisfies (2.18), and we obtain:

$$|\mathcal{G}_i \cap B| \ge \sum_{\text{good } B' \subset B} |\mathcal{G}_{i-1} \cap B'| \ge \big|\{B' \in \mathcal{P}_{i-1} \colon B' \subset B,\ B' \text{ good}\}\big| \bigg(\frac{1}{2} - \frac{2(i-2)\delta}{R}\bigg)(r_{i-1}')^d.$$

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

On $\mathcal{A}_i(B)$ in (2.42), there are $(1-2\delta/R)(r'_i/r'_{i-1})^d$ good sub-boxes of B, so a.s. on $\mathcal{A}_i(B)$

$$|\mathcal{G}_i \cap B| \geq \left(1-\frac{2\delta}{R}\right)\left(\frac{r'_i}{r'_{i-1}}\right)^d\left(\frac{1}{2}-\frac{2(i-2)\delta}{R}\right)(r'_{i-1})^d \geq \left(\frac{1}{2}-\frac{2(i-1)\delta}{R}\right)(r'_i)^d.$$

Vertices in bad sub-boxes cannot be i-good by Definition 2.10. So (2.18) similarly implies that on $\mathcal{A}_i(B)$ there are at most $2(r'_i)^d$ i-good vertices in B, and thus $\mathcal{A}_i(B)$ in (2.42) implies both *(B1)–(B2)* for B; it follows from (2.43) that

$$\mathbb{P}\big((B1) \text{ and } (B2) \text{ hold for } B \mid x_1,\dots,x_t \in \mathcal{V}\big) \geq 1-\delta/(2R). \tag{2.44}$$

By Lemma 2.14, B is good (i.e., *(B3)* also holds) with probability at least $1-\delta/(2R)$ for all $\mathcal{F}_{i-1}(B)=F$ with *(B1)–(B2)* holding for B; these events and $x_1,\dots,x_t \in \mathcal{V}$ are all determined by $\mathcal{F}_{i-1}(B)$. So, a union bound on (2.26) and (2.44) yields that independently across boxes in $\mathcal{P}_i$, and regardless of the vertices $x_1,\dots,x_t \in \mathcal{V}$, (2.41) holds. This advances the induction and finishes the proof. □

Proof of Proposition 2.5. Recall the setting of Proposition 2.5, and consider an $\mathcal{R}$-partition $\mathcal{P}_1,\dots,\mathcal{P}_R$ of Q, (which exists by Lemma 2.7), and let $\mathcal{H}(\mathcal{R})$ be the associated hyperrectangle cover of $Q \times [w_0, f(r_R)]$. By Lemma 2.15, conditioned on $x_1,\dots,x_t \in \mathcal{V}$, each box $B \in \mathcal{P}_1 \cup \dots \cup \mathcal{P}_R$ is good with probability at least $1-\delta/R$. Let $\mathcal{A}$ be the event that all boxes in $\{B^i(x_j) : i \in [R], j \in [t]\}$ are good. A union bound over $i=1,\dots,R$ and $1,\dots,t$ implies that

$$\mathbb{P}(\mathcal{A} \mid x_1,\dots,x_t \in \mathcal{V}) \geq 1-t\delta.$$

In particular, if $\mathcal{A}$ occurs then $B^R(x_1)=Q$ is also good, so by Lemma 2.12, the set $\widetilde{\mathcal{G}}_R =: \mathcal{N}$ of all R-good vertices forms an $(\delta,\mathcal{R})$-net of Q. The requirement that $\{B^i(x_j) : i \in [R], j \in [t]\}$ are good is exactly the requirement for $x_1,\dots,x_t \in \mathcal{N}$ to be R-good and this to be in this net, showing (2.8). To obtain (2.7), note that Lemma 2.15 implies that $Q \subset \mathcal{P}_R$ is good with probability at least $1-\delta/R$, and then again Lemma 2.12 finishes the proof. □

3. Multiround exposure with dependent edge-costs

Now with the nets at hand, we may reveal the realisation of the vertex set $\widetilde{\mathcal{V}}=(V,w_V)$. We shall now reveal edges adaptively to construct a fast-transmission path between $0,x$, according to the 'budget travel plan' in Section 1, see Fig. 2. In particular, we need to find low-cost edges in spatial regions which depend on the previous low-cost edges we have found. When studying graph distances in Biskup [11] this is not a major obstacle, but with the presence of edge-costs we run into conditioning issues. To overcome these, we develop a multiple-round exposure – essentially an elaborate edge-sprinkling method on the revealed vertex set – where in each round we reveal more than one edge.

In a classical random graph setting, constructing a path using multiple-round exposure would involve coupling the base graph model $\mathcal{G}$ to a suite of sparser but independent random subgraphs $H_1,\dots,H_r$, and taking the i-th edge of the path from the i-th 'round of exposure' H_i. In the edge-weighted setting we design a (slightly more restrictive) construction incorporating *independent* edge-cost variables on $H_1,\dots,H_r$ that we describe in Prop. 3.9 after some preliminary definitions. For future reusability, we formulate Prop. 3.9 in a general class of random graph models, as set out below. Recall that $V^{(2)}=\{uv : u,v \in V \text{ distinct}\}$ denotes the set of possible edges on V.

Definition 3.1 (CIRG models). A *conditionally independent edge-weighted vertex-marked random graph model (CIRG model)* $\mathcal{G}$ consists of a fixed countable weighted vertex set (V,w_V), a random edge set $\mathcal{E} \subseteq V^{(2)}$, and random edge costs $\mathcal{C}(xy)$ for each possible edge $\{x,y\} \in V^{(2)}$. All costs $\mathcal{C}(xy)$ and all events $\{xy \in \mathcal{E}\}$ are independent across $\{x,y\} \in V^{(2)}$. For brevity, we write '$G \sim \mathcal{G}$ is a CIRG' to mean that $G \sim \mathcal{G}$ and that $\mathcal{G}$ is a CIRG model. Sometimes it is convenient to speak of variables $G \sim \mathcal{G}$ without naming $\mathcal{G}$, and in this case we simply say 'G is a CIRG'.

First passage percolation (1-FPP) on SFP (Def. 1.1) and GIRG (Def. 1.3) both become CIRG models after their weighted vertex sets are exposed, with $\mathcal{C}(xy)$ in (1.2) being the edge-weights. It is easy to verify that the same is true of 1-FPP on Chung-Lu random graphs (defined in [22]) and inhomogeneous random graphs (defined in [15]) or on the subclass called the stochastic block model [71].

In a classical setting, a multiple-round exposure argument would 'split' a graph $G \sim \mathcal{G}$ into edge-disjoint graphs $G_1, \dots, G_r$, where we put each edge of G into precisely one G_i, independently across edges according to some probability distribution $(\theta_1, \dots, \theta_r)$. We would then couple $G_1, \dots, G_r$ to independent graphs $H_1, \dots, H_r$ using stochastic domination arguments. Here is the analogue of $G_1, \dots, G_r$ in our setting, where we must be careful about costs.

Definition 3.2 (θ-percolated CIRG). Let $G \sim \mathcal{G}$ be a CIRG from Definition 3.1. Then for all $\theta \in (0, 1)$ the θ*-percolation* of G is the subgraph G^θ of G which includes each $e \in \mathcal{E}(G)$ independently with probability θ, and we write $\mathcal{G}^\theta$ for its law. We call $\mathcal{G}^\theta$ the θ*-percolation* of $\mathcal{G}$, and θ the *percolation probability*.

Remark 3.3 (θ-percolated CIRGs are CIRGs). An alternative construction of $G \sim \mathcal{G}^\theta$ is to sample each edge e independently with probability $\theta\mathbb{P}(e \in \mathcal{E}(G))$. So $\mathcal{G}^\theta$ is a CIRG model, and the CIRG model class is closed under θ-percolation.

We now set out a specific coupling between a base CIRG model and percolated CIRGs, that will serve as graphs forming the rounds of exposure. Recall that $[r] := \{1, 2, \dots, r\}$.

Definition 3.4 (Exposure setting of G). Let $\mathcal{G}$ be a CIRG model from Def. 3.1 with vertex set V. Fix $r \in \mathbb{N}$ and $\theta_1, \dots, \theta_r \in [0, 1]$ satisfying $\sum_{i\in[r]} \theta_i = 1$. We define the *exposure setting* of $\mathcal{G}$ with *percolation probabilities* $\theta_1, \dots, \theta_r$ as follows. Let $(Z_{uv})_{uv\in V^{(2)}}$ be i.i.d. random variables with $\mathbb{P}(Z_{uv} = i) = \theta_i$ for all $i \in [r]$. Take $G_1^\star, \dots, G_r^\star$ to be i.i.d. CIRGs, with shared distributions $G_i^\star \sim \mathcal{G}$, and respective edge costs $\mathcal{C}_i(e)$ for $e \in \mathcal{E}(G_i^\star)$ chosen independently across $i \le r$. Let $G_i^{\theta_i}$ be the subgraph of $G_i^\star$ with edge set $\mathcal{E}(G_i^{\theta_i}) := \{e \in \mathcal{E}(G_i^\star) \colon Z_e = i\}$ and edge costs $\{\mathcal{C}_i(e) : e \in \mathcal{E}(G_i^{\theta_i})\}$.

In this definition, while the initial graphs $G_1^\star, \dots, G_r^\star$ are independent, their percolated versions $G_i^{\theta_i}$ are not, since they all use the same $(Z_{uv})_{uv\in V^{(2)}}$ collection. The following lemma reconstructs G from the percolated versions.

Lemma 3.5 (Realisation of a CIRG in the exposure setting). *Let $\mathcal{G}$ be a CIRG model from Def. 3.1 with weighted vertex set (V, w_V). Let $\theta_1, \dots, \theta_r$ be the percolation probabilities, and consider $(G_i^{\theta_i})_{i\le r}$ in Definition 3.4. Then marginally, each $G_i^{\theta_i}$ is a θ_i-percolated CIRG. Define now G as the graph with weighted vertex set (V, w_V), and with edge set $\mathcal{E}(G) := \cup_{i\in[r]}\mathcal{E}(G_i^{\theta_i})$, and with edge costs $\{\mathcal{C}(e) := \mathcal{C}_{Z_e}(e) : e \in \mathcal{E}(G)\}$. Then $G \sim \mathcal{G}$.*

Proof. That $G_i^{\theta_i}$ is marginally a θ_i-percolated CIRG, that is, that it has law $\mathcal{G}^{\theta_i}$, is immediate from the definition. To see that $G \sim \mathcal{G}$ we argue as follows. Since Z_{uv} takes a single value in $[r]$ each possible edge $e = uv$ appears in at most one of $G_1^{\theta_1}, \dots, G_r^{\theta_r}$. Hence the union $\cup_{i\in[r]}\mathcal{E}(G_i^{\theta_i}) = \mathcal{E}(G)$ is *disjoint*, and using that $G_1^\star, \dots, G_r^\star$ all have law $\mathcal{G}$,

$$\begin{aligned}\mathbb{P}\big(uv \in \mathcal{E}(G)\big) &= \sum_{i\in[r]} \mathbb{P}(Z_{uv} = i)\mathbb{P}\big(uv \in \mathcal{E}(G_i^\star)\big)\\ &= \sum_{i\in[r]} \theta_i\mathbb{P}\big(uv \in \mathcal{E}(G_1^\star)\big) = \mathbb{P}\big(uv \in \mathcal{E}(G_1^\star)\big),\end{aligned}$$

and $G_1^\star \sim \mathcal{G}$. Further, edges are present in G independently since the variables Z_e and $\mathcal{E}(G_i^\star)$ are independent. □

For a collection of variables $\mathcal{X}$ we write $\sigma(\mathcal{X})$ for the σ-algebra generated by the variables in $\mathcal{X}$. In the following definitions we formalise multiround exposure with edge-cost constraints, in the setting

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

of CIRGs with given weighted vertex-set (V, w_V), which guarantees that edge presence and edge costs are independent by Def. 3.1. These definitions are highly technical and so we provide a simplified motivating example before stating them, with further discussion to follow.

Suppose we are given a CIRG $G \sim \mathcal{G}$ with vertex set contained in $[0, \sqrt{n}]^2$, and we wish to join two fixed vertices u and v with a low-cost path in three rounds with equal edge probabilities. We therefore split G into three disjoint percolated CIRGs $G_1^{1/3} \subseteq G_1^\star$, $G_2^{1/3} \subseteq G_2^\star$ and $G_3^{1/3} \subseteq G_3^\star$ as in the exposure setting, taking $\theta_1 = \theta_2 = \theta_3 = 1/3$; by Lemma 3.5, every edge in any $G_i^{1/3}$ is an edge of G with the same cost. We bound the construction's failure probability on $G_1^{1/3}$, $G_2^{1/3}$ and $G_3^{1/3}$ by coupling it to the same construction on independent percolated CIRGs $H_1, H_2, H_3 \sim \mathcal{G}^{1/3}$; this coupling works in general and is stated later as Proposition 3.9.

In the first round, we reveal edges of G_1 (or H_1) and search for an unusually low-cost edge from an unspecified vertex u' near u to an unspecified vertex v' near v. In the second round, using edges of G_2 (or H_2) we search for a low-cost path from u to u'; and in the third round we search for a low-cost path from v' to v in G_3 (or H_3). In each round we describe the admissible object we search for and a cost constraint, such as 'any path from any $u' \in B_r(u)$ to any $v' \in B_r(v)$' and 'cost below a specific value C'). Since we reveal edges, these 'objects' are described in terms of edges. When we supply the actual graph $G_i^{1/3}$ or H_i, the round selects a concrete admissible object – for example, a fixed low-cost path $\pi_{u'v'}$ connecting two vertices u', v' in H_1. This concrete object is then used to specify admissible objects in future rounds: round 2 admissible objects depend on the value of u', which is a function of $\pi_{u',v'}$.

In defining an iterative cost construction, we consider the effect of applying it to arbitrary graphs G_1, G_2 and G_3 rather than to the specific graphs $G_i^{1/3}$ or H_i. In the definition below, these ideas are captured for the i-th round by $\mathcal{F}_i, \mathcal{U}_i$ and $\mathcal{S}_i$ respectively: $\mathcal{F}_i$ describes the 'admissible objects' among which we select one (as a list), $\mathcal{U}_i$ describes the cost-requirements on these, and then $\mathcal{S}_i$ finally reveals (part of) the edge set of G_i and selects the first admissible object in $\mathcal{F}_i$ that satisfies the cost-constraint. The next round(s) may use the outputs $\mathcal{S}_i$ of previous selection rounds – in our example, rounds 2 and 3 depend on the endpoints of the path chosen in round 1. Initially $\mathcal{F}_i$ and $\mathcal{U}_i$ will exist as 'functions' before any edge is revealed, describing all possible realisations of admissible objects and constraints on them in round i. Once $G_1, \dots, G_{i-1}$ have been specified, only a subset of these admissible objects will remain admissible; these new constraints we denote by $\mathcal{F}_i(G_1, \dots, G_{i-1})$ and $\mathcal{U}_i(G_1, \dots, G_{i-1})$. Thus each $\mathcal{F}_i$ and $\mathcal{U}_i$ is deterministic, while each $\mathcal{S}_i$ satisfying the structural constraints of $\mathcal{F}_i(G_1, \dots, G_{i-1})$ and the cost constraints of $\mathcal{U}_i(G_1, \dots, G_{i-1})$ is adapted to the natural filtration of the process. In this example, $\mathcal{F}_2$ will be a function from the set of all graphs on (V, w_V), that given any graph G_1 outputs an output which is a list of admissible objects for round 2: $\mathcal{F}_2(G_1)$ will be all possible paths between u and the specific $u' \in B_r(u)$ selected in the first round.

There is one more important detail. After choosing an edge wz in a round, its cost is exposed and cannot be redefined in future rounds. To illustrate this in our running example, suppose we are at round 2, and there is only one low cost path – say π – present between u, u' in G_2, and π uses an edge wz present in G_2 that has been selected already in round 1 as part of the path between u', v'. While wz is present independently with independent costs in H_1 and H_2, the same is not true in $G_1^{1/3}$ and $G_2^{1/3}$, and we need to couple the construction's progress on these two sets of graphs. We cannot meaningfully assign two different costs to wz in G, and so we need to adapt the definition of a construction in response.

Observe that despite the fact that we have already chosen wz as an edge, it still makes sense to use it as part of π. Indeed, we already accounted for the cost of wz when we chose it as part of round 1, and if the three paths from our three rounds overlap then we can always pass to a sub-path with lower total cost. As such, in the definition of an iterated cost construction, we artificially set the cost in G_i of every edge that has already been selected to 0. This we call the i-th *marginal cost* of the edge. Then (since $\mathcal{U}_i$ is fixed), we can require that the cost constraints $\mathcal{U}_i$ depend only on the *marginal* costs of edges in $G_i^{1/3}$ and H_i. This way we avoid exposing the same randomness twice. For our applications in later sections, the distinction between cost and marginal cost will not be important, as (like our example) we only require low total cost – and the total cost of our edge set will be precisely the sum of the marginal costs

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

over all edges we select. (See Remark 3.7.) In other applications, one may wish to modify $\mathcal{F}_i$ to enforce disjointness from some or all of already selected edges $\mathcal{S}_1, \dots, \mathcal{S}_{i-1}$.

With all this in mind, we now define the iterative cost construction in precise mathematical terms. We denote by $\mathcal{E}(G)$ the edge set of a graph G. For simplicity we restrict to simple graphs on finite vertex sets, however, extension to countable vertex sets and multigraphs is possible.

Definition 3.6 (Iterative cost construction). Fix a weighted vertex set (V, w_V), and let $r \geq 1$ be an integer. Let $V^{(2)} = \{uv \colon u, v \in V \text{ distinct}\}$ denote the set of possible edges on V, and fix an ordering on $V^{(2)}$. Let $\mathfrak{G}_{V,w_V}$ be the set of all *edge-weighted* graphs on (V, w_V). An *r-round iterative cost construction* Iter is a collection $(\mathcal{F}_1,\mathcal{U}_1), \dots, (\mathcal{F}_r,\mathcal{U}_r)$ of functions that, when applied to a sequence of graphs $(G_1, \dots, G_r)$ all in $\mathfrak{G}_{V,w_V}$, outputs a list of selected edge-lists in r rounds. We write $\mathcal{S}_i = \mathcal{S}_i(G_1, \dots, G_i) = \mathrm{Iter}_i(G_1, \dots, G_i)$ for the output in the i-th round. Iter is defined recursively as follows.

(i) For all $i \in [r]$, the domains of $\mathcal{F}_i$ and $\mathcal{U}_i$ are $\mathfrak{G}^{i-1}$ (so they are 'deterministic' admissible sets and events for $i = 1$, while for $i \geq 2$ they may depend on previously exposed graphs $G_1, \dots G_{i-1}$). The domain of $\mathcal{S}_i = \mathrm{Iter}_i$ is $\mathfrak{G}^i$ (so $\mathcal{S}_i$ depends only on $G_1, \dots G_i$).

(ii) For all $i \in [r]$ and all $(G_1, \dots, G_{i-1}) \in \mathfrak{G}^{i-1}_{V,w_V}$, $\mathcal{F}_i(G_1, \dots, G_{i-1})$ is a finite list. Each element in $\mathcal{F}_i(G_1, \dots, G_{i-1})$ is a list of pairs of vertices in $V^{(2)}$. Each list in $\mathcal{F}_i(G_1, \dots, G_{i-1})$ contains each pair in $V^{(2)}$ at most once. $\mathcal{F}_i(G_1, \dots, G_{i-1})$ represents the set of admissible objects for the output of round i, given $G_1, \dots, G_{i-1}$. Moreover, $\mathcal{F}_i(G_1, \dots G_{i-1})$ depends only on the already selected weighted edges $\mathcal{S}_1, \dots, \mathcal{S}_{i-1}$.

(iii) For all $i \in [r]$ and all $G_1, \dots, G_i \in \mathfrak{G}^i_{V,w_V}$, define the *round-i marginal cost* of an edge e by

$$\mathrm{mcost}_i(e) = \begin{cases} 0 & \text{if } e \text{ appears in any of the already-chosen lists } \mathcal{S}_1, \dots, \mathcal{S}_{i-1}, \\ \mathcal{C}_i(e) & \text{otherwise;} \end{cases} \tag{3.1}$$

note that we suppress the dependence of $\mathrm{mcost}_i(e)$ on $G_1, \dots, G_i$ for brevity.

(iv) For all $i \in [r]$ and all $(G_1, \dots, G_{i-1}) \in \mathfrak{G}^{i-1}_{V,w_V}$, $\mathcal{U}_i(G_1, \dots, G_{i-1})$ is a finite list of events. The j-th element of $\mathcal{U}_i(G_1, \dots, G_{i-1})$ is the event describing when the j-th element of $\mathcal{F}_i(G_1, \dots, G_{i-1})$ is allowed to be selected in terms of the round-i marginal costs. Each element in $\mathcal{U}_i(G_1, \dots, G_{i-1})$ is of a specific form: for all $(e_1, \dots, e_t) \in \mathcal{F}_i(G_1, \dots, G_{i-1})$, $\mathcal{U}_i(e_1, \dots, e_t)$ describes a subset $D_i(e_1, \dots, e_t)$ of $[0, \infty)^t$, and the event itself is $(\mathrm{mcost}_i(e_1), \dots, \mathrm{mcost}_i(e_t)) \in D_i(e_1, \dots, e_t)$. Moreover, $\mathcal{U}_i(G_1, \dots G_{i-1})$ depends only on the already selected weighted edges $\mathcal{S}_1, \dots, \mathcal{S}_{i-1}$.

(v) We say that a list of edges $(e_1, \dots, e_t)$ is *present in round i* if for all $j \in [t]$ we have $e_j \in \mathcal{E}(G_i) \cup \mathcal{S}_1 \cup \dots \cup \mathcal{S}_{i-1}$.

(vi) $\mathcal{S}_i = \mathrm{Iter}_i(G_1, \dots, G_i)$ is specified as follows. $\mathcal{S}_i$ is the first element $(e_1, \dots, e_t)$ of $\mathcal{F}_i(G_1, \dots, G_{i-1})$ which is present in round i and which also satisfies the corresponding event in $\mathcal{U}_i(G_1, \dots, G_r)$. If no such edge exists, we define $\mathcal{S}_i = \texttt{None}$. Observe that the value of $\mathcal{S}_i$ depends only on the values of $\mathcal{S}_1, \dots, \mathcal{S}_{i-1}$ and G_i.

We refer to $\mathcal{F}_i$ and $\mathcal{U}_i$ as the *i-th round* of Iter, we refer to $\mathcal{S}_i = \mathrm{Iter}_i(G_1, \dots, G_i)$ as the *output of the i-th round*, and we write $\mathrm{Iter}(G_1, \dots, G_r)$ for the sequence $\mathcal{S}_1, \dots, \mathcal{S}_r$. We say that Iter *succeeds* on $G_1, \dots, G_r$ if $\mathcal{S}_i \neq \texttt{None}$ for all $i \in [r]$.

Returning to our example from above Definition 3.6, the $(\mathcal{F}_i,\mathcal{U}_i)$ becomes the following. We think of paths as sequences of edges and we write $u'(\mathcal{S}_1), v'(\mathcal{S}_1)$ for the two endpoints of the selected path $\mathcal{S}_1$ (in $B_r(u), B_r(v)$, respectively). Then:

$$\begin{aligned} \mathcal{F}_1 &= \{\pi_{u'v'} : \pi_{u'v'} \text{ is a path between some } u' \in B_r(u), v' \in B_x(v)\}, \\ \mathcal{U}_1 &= \Big\{ \sum_{e \in \pi_{u'v'}} \mathrm{mcost}_1(e) \leq C/3 : \pi_{u'v'} \in \mathcal{F}_1 \Big\}, \end{aligned} \tag{3.2}$$

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

are indeed the same for any input graphs, while

$$\begin{aligned}
\mathcal{F}_2(G_1) &= \{\pi_{uu'} : \pi_{uu'} \text{ is a path from } u \text{ to } u'(\mathcal{S}_1(G_1))\},\\
\mathcal{U}_2(G_1) &= \Big\{ \sum_{e\in\pi_{uu'}} \mathrm{mcost}_2(e) \le C/3; \pi_{uu'} \in \mathcal{F}_2 \Big\},\\
\mathcal{F}_3(G_1,G_2) &= \{\pi_{vv'} : \pi_{vv'} \text{ is a path in } V \text{ from } v'(\mathcal{S}_1(G_1)) \text{ to } v\},\\
\mathcal{U}_3(G_1,G_2) &= \Big\{ \sum_{e\in\pi_{vv'}} \mathrm{mcost}_3(e) \le C/3 : \pi_{vv'} \in \mathcal{F}_3 \Big\},
\end{aligned} \tag{3.3}$$

depend on the first round and second rounds. If there is an edge overlap between the selected paths in different rounds, then we see 0 marginal cost of that edge in later rounds. Since our goal is to bound the total cost of the selected edges, this is sufficient. More generally, the following remark motivates our definition of $\mathrm{mcost}_i(e)$ in (3.1). Recall the exposure setting of a graph G from Definition 3.4. The statement is a direct consequence of Lemma 3.5 and the previous definition.

Remark 3.7. Apply an r-round iterative cost construction on $(G_1^{\theta_1}, \dots, G_2^{\theta_r})$, which form the exposure setting of G in Definition 3.4 having cost-function $\mathcal{C}$. Then

$$\sum_{i=1}^{r} \sum_{e\in\mathcal{S}_i} \mathrm{mcost}_i(e) = \sum_{i=1}^{r} \sum_{e\in\mathcal{S}_i\setminus(\mathcal{S}_1\cup\cdots\cup\mathcal{S}_{i-1})} \mathcal{C}(e) = \sum_{e\in\mathcal{S}_1\cup\cdots\cup\mathcal{S}_r} \mathcal{C}(e).$$

That is, the total marginal costs are the same as the total cost of the selected *set* of edges in G.

Our goal is to prove that iterative cost constructions behave the same way as any other multiple-round exposure argument (usually proven by applying an FKG-type inequality): If the construction succeeds whp on an independent graph sequence $(H_1, \dots H_r)$ with $H_i \sim \mathcal{G}^{\theta_i}$, then it will also succeed when a single random graph $G \sim \mathcal{G}$ is percolated into edge-disjoint copies $G_i^{\theta_i}$ (as in the exposure setting of Lemma 3.5). To this end, we first set out notation for these two situations.

Definition 3.8. Fix a weighted vertex set (V, w_V), let $r \ge 1$ be an integer, and let Iter be an r-round iterative construction consisting of $(\mathcal{F}_1,\mathcal{U}_1), \dots, (\mathcal{F}_r,\mathcal{U}_r)$. Let $\underline{\theta} = (\theta_1, \dots, \theta_r) \in [0,1]^r$ with $\sum_{i\in[r]} \theta_i = 1$, and suppose $\mathcal{G}$ is a CIRG model on (V, w_V). Let $G_1^{\theta_1}, \dots, G_r^{\theta_r}$ be as in the exposure setting; then we write $\mathrm{Iter}^{\mathrm{exp}}_{\mathcal{G},\underline{\theta}} := \mathrm{Iter}(G_1^{\theta_1}, \dots, G_r^{\theta_r})$. Let $H_1 \sim \mathcal{G}^{\theta_1}, \dots, H_r \sim \mathcal{G}^{\theta_r}$ independently; then we write $\mathrm{Iter}^{\mathrm{ind}}_{\mathcal{G},\underline{\theta}} := \mathrm{Iter}(H_1, \dots, H_r)$. We say that $\mathrm{Iter}^{\mathrm{exp}}_{\mathcal{G},\underline{\theta}}$ succeeds if Iter succeeds on $G_1^{\theta_1}, \dots, G_r^{\theta_r}$, and likewise for $\mathrm{Iter}^{\mathrm{ind}}_{\mathcal{G},\underline{\theta}}$. For all sequences of list of edges $S_1, \dots, S_i$, we introduce the shorthand notation

$$\mathcal{A}^{\mathrm{exp}}(S_1, \dots, S_i) := \bigcap_{j\in[i]} \{\mathrm{Iter}_j(G_1^{\theta_1}, \dots, G_{j-1}^{\theta_{j-1}}) = S_j\}. \tag{3.4}$$

We define $\mathcal{A}^{\mathrm{ind}}$ analogously for $H_1, \dots, H_r$. For brevity, for all $i \in [r]$ we write $\mathcal{S}_i^{\mathrm{ind}} := \mathrm{Iter}_i(H_1, \dots, H_i)$ and $\mathcal{S}_i^{\mathrm{exp}} := \mathrm{Iter}_i(G_1^{\theta_1}, \dots, G_i^{\theta_i})$, and use analogous notation for $\mathcal{F}_i$ and $\mathcal{U}_i$.

The following proposition essentially says we can lower-bound the probability of success of the exposure setting $\mathrm{Iter}^{\mathrm{exp}}_{\mathcal{G},\underline{\theta}} = \mathrm{Iter}(G_1^{\theta_1}, \dots, G_r^{\theta_r})$ by that of the much simpler independent setting $\mathrm{Iter}(H_1, \dots, H_r)$.

Proposition 3.9 (Multiround exposure). *Consider the setting of Definition 3.8. Then*

$$\mathbb{P}\big(\mathrm{Iter}^{\mathrm{exp}}_{\mathcal{G},\underline{\theta}} \textit{ succeeds}\big) \ge \min_{S_1,\dots,S_{r-1}\not\ni\mathtt{None}} \prod_{i\in[r]} \mathbb{P}\big(\mathcal{S}_i^{\mathrm{ind}} \ne \mathtt{None} \mid \mathcal{A}^{\mathrm{ind}}(S_1, \dots, S_{i-1})\big). \tag{3.5}$$

Proof of Proposition 3.9. By repeated conditioning, and taking the minimum over all successful relations, we have

$$\begin{aligned}\mathbb{P}\big(\mathrm{Iter}^{\exp}_{\mathcal{G},\underline{\theta}}\text{ succeeds}\big) &= \prod_{i\in[r]}\mathbb{P}\big(\mathcal{S}_i^{\exp}\neq \texttt{None}\mid \mathcal{S}_1^{\exp},\dots,\mathcal{S}_{i-1}^{\exp}\neq\texttt{None}\big)\\ &\geq \min_{S_1,\dots,S_{r-1}\neq\texttt{None}}\prod_{i\in[r]}\mathbb{P}\big(\mathcal{S}_i^{\exp}\neq\texttt{None}\mid \mathcal{A}^{\exp}(S_1,\dots,S_{i-1})\big).\end{aligned} \tag{3.6}$$

In order to prove (3.5), it now suffices to prove the termwise bound on the right-hand side of (3.6) that for all $i\in[r]$, for all possible non-None outcomes $S_1,\dots,S_{i-1}$ of the first $i-1$ rounds of Iter:

$$\mathbb{P}\big(\mathcal{S}_i^{\exp}\neq\texttt{None}\mid\mathcal{A}^{\exp}(S_1,\dots,S_{i-1})\big)\geq\mathbb{P}\big(\mathcal{S}_i^{\mathrm{ind}}\neq\texttt{None}\mid\mathcal{A}^{\mathrm{ind}}(S_1,\dots,S_{i-1})\big). \tag{3.7}$$

To do so, for each possible $S_1,\dots,S_{i-1}$ we will couple $G_i^{\theta_i}$ conditioned on $\mathcal{A}^{\exp}(S_1,\dots,S_{i-1})$ to H_i conditioned on $\mathcal{A}^{\mathrm{ind}}(S_1,\dots,S_{i-1})$. For all $e\in V^{(2)}$, this coupling will satisfy:

$$\{e\in\mathcal{E}(H_i)\}\cap\{e\in\mathcal{E}(G_i^{\theta_i})\}\subseteq\{\mathcal{C}_{H_i}(e)=\mathcal{C}_i(e)\}, \tag{3.8}$$

$$\{e\in\mathcal{E}(H_i)\}\subseteq\{e\in\mathcal{E}(G_i^{\theta_i})\cup S_1\cup\cdots\cup S_{i-1}\}. \tag{3.9}$$

In words, if an edge appears in both graphs H_i and $G_i^{\theta_i}$ then its cost is the same in both, and if an edge is in H_i then either it has been chosen already in previous rounds or it is also in $G_i^{\theta_i}$. Under both conditionings $\mathcal{A}^{\exp}(S_1,\dots,S_{i-1})$ and $\mathcal{A}^{\mathrm{ind}}(S_1,\dots,S_{i-1})$ in (3.7), $\mathcal{S}_j^{\exp}=\mathcal{S}_j^{\mathrm{ind}}=S_j$ for all $j\leq i-1$. Recall from Definition 3.6(ii) and (iv) that $\mathcal{F}_i^{\exp},\mathcal{F}_i^{\mathrm{ind}},\mathcal{U}_i^{\exp}$ and $\mathcal{U}_i^{\mathrm{ind}}$ only depend on the already selected edge-lists $\mathcal{S}_1,\dots,\mathcal{S}_{i-1}$, so $\mathcal{F}_i^{\exp}=\mathcal{F}_i^{\mathrm{ind}}$ and $\mathcal{U}_i^{\exp}=\mathcal{U}_i^{\mathrm{ind}}$.

Suppose we have a coupling that satisfies both (3.8) and (3.9), and $\mathcal{S}_i^{\mathrm{ind}}=:S_i\neq\texttt{None}$, that is, the independent construction returns with a list S_i when we reveal H_i. Then by (3.9) each edge $e\in S_i\subseteq\mathcal{E}(H_i)\cup S_1\cup\cdots\cup S_{i-1}$ is also contained in $\mathcal{E}(G_i^{\theta_i})\cup S_1\cup\cdots\cup S_{i-1}$, and by (3.1) and (3.8) the round-i marginal cost of e in $\mathrm{Iter}^{\exp}_{\mathcal{G},\underline{\theta}}$ equals those in $\mathrm{Iter}^{\mathrm{ind}}_{\mathcal{G},\underline{\theta}}$. Hence S_i provides a valid choice for $\mathcal{S}_i^{\exp}$ (i.e., it lies in $\mathcal{F}_i^{\exp}$ and satisfies $\mathcal{U}_i^{\exp}$), and $\mathcal{S}_i^{\exp}\neq\texttt{None}$ holds also. Thus (3.7) holds, and so the result follows from (3.6).

It remains to provide the coupling achieving (3.8) and (3.9). Recall that Definition 3.4 uses the independent graphs $G_i^\star\sim\mathcal{G}$, and obtains $G_i^{\theta_i}$ as a dependent thinning of $G_i^\star$ using $(Z_{uv})_{uv\in V^{(2)}}$ (independently across different uv). For each $uv\in V^{(2)}$ and $i\in[r]$, sample i.i.d. uniform $U_{uv}^{(i)}\sim U[0,1]$ and realise the presence of uv in H_i and respectively in $G_i^{\theta_i}$ as

$$\mathbf{1}_{uv\in H_i}=\mathbf{1}_{uv\in G_i^\star}\mathbf{1}_{U_{uv}^{(i)}\leq\theta_i},\qquad \mathbf{1}_{uv\in G_i^{\theta_i}}=\mathbf{1}_{uv\in G_i^\star}\mathbf{1}_{Z_{uv}=i}. \tag{3.10}$$

Then $H_1,\dots,H_r$ are *independent* θ_i-percolations of $G_1^\star,\dots,G_r^\star$ respectively, so $H_1,\dots,H_r$ are themselves independent as required (since $G_1^\star,\dots,G_r^\star$ are independent). Note that (3.10) is only a *partial coupling*, since we can still specify the joint distribution of $(U_{uv}^{(i)})_{i\leq r}$ and Z_{uv} over $uv\in V^{(2)}$. By this partial coupling, H_i and $G_i^{\theta_i}$ are both subgraphs of $G_i^\star$, and hence if an edge e is in both subgraphs, then $\mathcal{C}_{H_i}(e)=\mathcal{C}_i(e)$ (the cost of e in $G_i^\star$), so (3.8) holds.

We now extend (3.10) into a full coupling which satisfies (3.9). Fix i and $S_1,\dots,S_{i-1}$. We first claim the following distributional identities hold:

$$\Big(\mathcal{E}(H_i)\mid\mathcal{A}^{\mathrm{ind}}(S_1,\dots,S_{i-1})\Big)\overset{d}{=}\mathcal{E}(H_i), \tag{3.11}$$

$$\big(\mathcal{E}(G_i^\star)\mid\mathcal{A}^{\exp}(S_1,\dots,S_{i-1})\big)\overset{d}{=}\mathcal{E}(G_i^\star). \tag{3.12}$$

Indeed, by Definition 3.6 and $\mathrm{Iter}^{\mathrm{ind}}_{\mathcal{G},\underline{\theta}}$ in Definition 3.8, the event $\mathcal{A}^{\mathrm{ind}}(S_1,\dots,S_{i-1})$ is measurable with respect to $\sigma(H_1,\dots,H_{i-1})$, and $H_1,\dots,H_{i-1}$ are *independent* of H_i, so (3.11) holds. Similarly, by Definition 3.6 and (3.10), $\mathcal{A}^{\exp}(S_1,\dots,S_{i-1})$ is in $\sigma(G_1^{\theta_1},\dots,G_{i-1}^{\theta_{i-1}})\subseteq\sigma(G_1^\star,\dots,G_{i-1}^\star,(Z_e)_{e\in V^{(2)}})$, which are independent of $G_i^\star$, so (3.12) follows.

Given (3.11) and (3.12), to prove that (3.10) can be extended to a full coupling satisfying (3.9), by Strassen's theorem it now suffices to prove that for all $k\ge 1$ and all $e_1,\dots,e_k\in V^{(2)}\setminus(S_1\cup\dots\cup S_{i-1})$,

$$\begin{aligned}\mathbb{P}\big(e_1,\dots,e_k\in\mathcal{E}(G_i^{\theta_i})\mid\mathcal{A}^{\exp}(S_1,\dots,S_{i-1})\big)&\\ \ge\mathbb{P}\big(e_1,\dots,e_k\in\mathcal{E}(H_i)\mid\mathcal{A}^{\mathrm{ind}}(S_1,\dots,S_{i-1})\big).&\end{aligned}\tag{3.13}$$

We compute the right-hand side using (3.10)–(3.12):

$$\begin{aligned}&\mathbb{P}\big(e_1,\dots,e_k\in\mathcal{E}(H_i)\mid\mathcal{A}^{\mathrm{ind}}(S_1,\dots,S_{i-1})\big)=\mathbb{P}\big(e_1,\dots,e_k\in\mathcal{E}(H_i)\big)\\ &\quad=\theta_i^k\cdot\mathbb{P}\big(e_1,\dots,e_k\in\mathcal{E}(G_i^\star)\big)=\theta_i^k\cdot\mathbb{P}\big(e_1,\dots,e_k\in\mathcal{E}(G_i^\star)\mid\mathcal{A}^{\exp}(S_1,\dots,S_{i-1})\big)\end{aligned}\tag{3.14}$$

Dividing the left-hand side of (3.13) and the right-hand side here by $\mathbb{P}\big(e_1,\dots,e_k\in\mathcal{E}(G_i^\star)\mid\mathcal{A}^{\exp}(S_1,\dots,S_{i-1})\big)$ and applying (3.10) yields that showing (3.13) is equivalent to showing

$$\mathbb{P}\big(Z_{e_1}=\dots=Z_{e_k}=i\mid\{e_1,\dots,e_k\in\mathcal{E}(G_i^\star)\}\cap\mathcal{A}^{\exp}(S_1,\dots,S_{i-1})\big)\ge\theta_i^k.$$

Since $\sigma(G_i^\star)$ is independent of $\sigma(G_1^\star,\dots,G_{i-1}^\star,(Z_e)_{e\in V^{(2)}})$, we can drop the first event from the conditioning and thus proving (3.13) is equivalent to proving that for all $k\ge1$ and all $e_1,\dots,e_k\in V^{(2)}\setminus(S_1\cup\dots\cup S_{i-1})$,

$$\mathbb{P}\big(Z_{e_1}=\dots=Z_{e_k}=i\mid\mathcal{A}^{\exp}(S_1,\dots,S_{i-1})\big)\ge\theta_i^k.\tag{3.15}$$

Intuitively, this inequality holds since all chosen edges in $(S_1,\dots,S_{i-1})$ have $Z_e\le i-1$ so if none of $(e_j)_{j\le k}$ has been chosen yet, then the probability that their Z_e value was higher than $i-1$ is larger. We next make this intuition precise.

We express $\mathcal{A}^{\exp}(S_1,\dots,S_{i-1})$ in terms of simpler events, using Definition 3.6. Recall that $\mathcal{F}_i^{\exp}$ and $\mathcal{U}_i^{\exp}$ depend only on the results of the first $i-1$ rounds, which are fixed by our conditioning on $\mathcal{A}^{\exp}(S_1,\dots,S_{i-1})$. To highlight this dependence, we informally write $\mathcal{F}_i^{\exp}=:\mathcal{F}_i(S_1,\dots,S_{i-1})$ and $\mathcal{U}_i^{\exp}=:\mathcal{U}_i(S_1,\dots,S_{i-1})$. For $j\le i-1$, let $\underline{t}_{j,s}$ be the s-th element in the list $\mathcal{F}_j(S_1,\dots,S_{j-1})$, and write $S_j=:\underline{t}_{j,s_j^\star}$ so that $s_j^\star$ is the index of the outcome of $\mathrm{Iter}_j(G_1^{\theta_1},\dots,G_{j-1}^{\theta_j})$. Fix an $e_1,\dots,e_k\in V^{(2)}\setminus(S_1\cup\dots\cup S_{i-1})$, and define

$$\begin{aligned}A_{j,s}&:=\Big\{\exists e\in\Big(\underline{t}_{j,s}\setminus(S_1\cup\ldots\cup S_{j-1})\Big)\cap\{e_1,\dots,e_k\}:Z_e\ne j\Big\},\\ B_{j,s}&:=\Big\{\exists e\in\Big(\underline{t}_{j,s}\setminus(S_1\cup\ldots\cup S_{j-1})\Big)\setminus\{e_1,\dots,e_k\}:Z_e\ne j\Big\},\\ C_{j,s}&:=\big\{\underline{t}_{j,s}\text{ does not satisfy }\mathcal{U}_j^{\exp}(S_1,\dots S_{j-1})\big\}\\ &\qquad\qquad\cup\big\{\exists e\in\underline{t}_{j,s}\setminus(S_1\cup\ldots\cup S_{j-1}):e\notin\mathcal{E}(G_j^\star)\big\},\end{aligned}\tag{3.16}$$

with the idea that if $t_{j,s}\setminus(S_1\cup\ldots\cup S_{j-1})\cap\{e_1,\dots,e_k\}=\emptyset$ then $A_{j,s}=\Omega$. Since the event $Z_e\ne j$ means that the edge is not present in $G_j^{\theta_j}$ by (3.10), $A_{j,s}\cup B_{j,s}$ means that there is at least one edge in $\underline{t}_{j,s}\setminus(S_1\cup\dots\cup S_{j-1})$ that is not present in $G_j^{\theta_j}$. By Definition 3.6(v), we can only choose $t_{j,s}$ for S_j if these edges are all present. We also cannot choose $t_{j,s}$ for S_j if $C_{j,s}$ holds by Definition 3.6(v) and (vi). However, if the complement of all these events hold, then all conditions are satisfied and $t_{j,s}$

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

is choosable for S_j, and it is chosen if it is the first such element of $\mathcal{F}_i^{\exp}$. Hence

$$\mathcal{A}^{\exp}(S_1,\dots,S_{i-1}) = \bigcap_{j\le i-1}\left(A^{\complement}_{j,s_j^\star}\cap B^{\complement}_{j,s_j^\star}\cap C^{\complement}_{j,s_j^\star}\cap\bigcap_{s<s_j^\star}(A_{j,s}\cup B_{j,s}\cup C_{j,s})\right).$$

In fact, by assumption we have $e_1,\dots,e_k\notin(S_1\cup\ldots\cup S_{i-1})$ and $j\le i-1$, so the chosen lists $\underline{t}_{j,s_j^\star}=S_j$ do not have an overlap with $e_1,\dots,e_k$. So this expression becomes

$$\mathcal{A}^{\exp}(S_1,\dots,S_{i-1}) = \bigcap_{j\le i-1}\left(B^{\complement}_{j,s_j^\star}\cap C^{\complement}_{j,s_j^\star}\cap\bigcap_{s<s_j^\star}(A_{j,s}\cup B_{j,s}\cup C_{j,s})\right). \tag{3.17}$$

Now define $A := \{Z_{e_1}=\cdots=Z_{e_k}=i\}$ as in (3.15), and define $B := \bigcap_{j\le i-1}(B^{\complement}_{j,s_j^\star}\cap C^{\complement}_{j,s_j^\star})$ and $C := \bigcap_{j\le i-1}\bigcap_{s<s_j^\star}(A_{j,s}\cup B_{j,s}\cup C_{j,s})$ as in (3.17), so that

$$\mathbb{P}\big(Z_{e_1}=\cdots=Z_{e_k}=i\mid\mathcal{A}^{\exp}(S_1,\dots,S_{i-1})\big)=\mathbb{P}(A\mid B\cap C)=\frac{\mathbb{P}(A\cap B\cap C)}{\mathbb{P}(B\cap C)}.$$

We first argue that $A\subseteq\bigcap_{j\le i-1}\bigcap_{s<s_j^\star}A_{j,s}\subseteq C$. Indeed, these $A_{j,s}$ all require that $Z_e\ne j$ whenever e is both in $e_1,\dots,e_k$ and also part of another list. In particular if $A=\{Z_{e_1}=\cdots=Z_{e_k}=i\}$ occurs then $Z_e\ne j$ is satisfied for all $e\in(e_1,\dots,e_k)$. This means that $A\cap B\cap C=A\cap B$. Moreover, all events $B_{j,s}$ with $j\le i-1$ and $s\le s_j^\star$ are independent of A, as A is contained in $\sigma(Z_{e_1},\dots,Z_{e_k})$ while $B_{j,s}$ is contained in $\sigma(\{Z_{uv}: uv\in V^{(2)}\setminus\{e_1,\dots,e_k\})$. We also observe that all events $C_{j,s}$ with $j\le i-1$ are independent of A, as A is contained in $\sigma(Z_{e_1},\dots,Z_{e_k})$ while $C_{j,s}$ is contained in $\sigma(G_1^\star,\dots,G_{i-1}^\star)$. We just proved that A is independent of B, so $\mathbb{P}(A\cap B)=\mathbb{P}(A)\cdot\mathbb{P}(B)$. So,

$$\frac{\mathbb{P}(A\cap B\cap C)}{\mathbb{P}(B\cap C)}=\frac{\mathbb{P}(A\cap B)}{\mathbb{P}(B\cap C)}=\frac{\mathbb{P}(A)\mathbb{P}(B)}{\mathbb{P}(B\cap C)}=\frac{\mathbb{P}(A)}{\mathbb{P}(C\mid B)}\ge\mathbb{P}(A)=\theta_i^k,$$

where the last equality follows since $Z_{e_1},\dots,Z_{e_k}$ are independent. This yields the right-hand side of (3.15), finishing the proof. □

4. Building blocks: finding cheap edges

In this section, we return from CIRGs to GIRGs and state a few important lemmas that we shall use to construct the different parts of the low-cost path between 0 and x. We work in the quenched setting with the realisation of vertices and their weights $\widetilde{\mathcal{V}}=(V,w_V)$ *exposed*, taking the role of (V,w_V) for CIRGs of Definition 3.1, and the weighted vertex set containing a weak net as in Section 2. All lemmas here concern θ-percolated SFP/IGIRG as in Definition 3.2, so that we can later use them on the graphs H_i of the multiround exposure Proposition 3.9. We first set out some common notation for Sections 4 and 5.

Setting 4.1. (The setting) Consider 1-FPP in Definition 1.1 on the graphs IGIRG or SFP satisfying the assumptions given in (1.6)–(1.3) with $d\ge1,\alpha\in(1,\infty],\tau\in(2,3)$. Let $\underline{c}$, $\overline{c}$, h, L, c_1, c_2, and β be as in (1.6)–(1.3); we allow $\beta=\infty$ and $\alpha=\infty$. Fix a realisation (V,w_V) of $\widetilde{\mathcal{V}}$, and let $G\sim\{\mathcal{G}\mid V,w_V\}$, and for a $\theta\in(0,1]$, let G' be a θ-percolation G' of G. For brevity we write $\mathbb{P}(\cdot\mid V,w_V)$ for $\mathbb{P}(\cdot\mid\widetilde{\mathcal{V}}(G')=(V,w_V))$. Let $x\in V$, and let Q be a cube of side length ξ containing 0 and x. Let $\delta\in(0,1)$, $w_0\ge1$, and assume that (V,w_V) is such that Q contains a weak $(\delta/4,w_0)$-net $\mathcal{N}$ with $0,x\in\mathcal{N}$ given in Definition 2.1. Finally, let $\gamma\in(0,1)$.

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

We now define a function of crucial importance for the optimisation of the exponents Δ_0, η_0 in (1.9) (1.10). For all $\gamma > 0$ and all $\eta, z \geq 0$, we define

$$\Lambda(\eta, z) := 2d\gamma - \alpha(d-z) - z(\tau-1) + \big(0 \wedge \beta(\eta - \mu z)\big). \tag{4.1}$$

The first lemma ensures the existence of the 'longest' edge on Figure 2: for $\gamma < 1$ it joins two Euclidean balls of radius D^γ with a low-cost edge with endpoints in the net having specified weights in a range around D^z, for any large enough $D \in \mathbb{R}$. We will apply it on multiple scales, hence the need for a general D for the distance that the edge bridges. The function $\Lambda(\eta, z)$ gives the exponent of D in the error probability of finding the low-cost edge. Recall $\mathtt{par}$ and the notation $\gg_\star$ from Section 1.4.1.

Lemma 4.2 (Single bridge-edge). *Consider Setting 4.1. Let* $z \in [0, d]$ *satisfy* $2d\gamma > z(\tau-1)$. *Let* $c_H, \eta \geq 0$. *Suppose that* $0 < \delta \ll_\star \gamma, \eta, z, c_H, \mathtt{par}$, *and that* $D \gg_\star \eta, z, c_H, \delta, w_0$. *Assume that* $D^\gamma \in [(\log\log \xi\sqrt{d})^{16/\delta}, \xi\sqrt{d}]$ *and that* $x, y \in \mathcal{N}$ *satisfy* $|x-y| \leq c_H D$, *and let* $\underline{w} \in [w_0 \vee 4(c_H+2)^d \vee 4000c_1^{-1/(\mu\beta)}, D^\delta]$ *satisfy* $F_L((\underline{w}/4000)^\mu) \geq 1/2$. *For* $v \in \{x, y\}$, *define*

$$\mathcal{Z}(v) = \mathcal{Z}_{\gamma,z,\underline{w}}(v) := \mathcal{N} \cap \big(B_{D^\gamma}(v) \times [\underline{w}D^{z/2}/2, 2\underline{w}D^{z/2}]\big). \tag{4.2}$$

Again for each $v \in \{x, y\}$, *let* $Z_v \subseteq \mathcal{Z}(v)$ *with* $|Z_v| \geq |\mathcal{Z}(v)|/4$ *be two arbitrary subsets chosen in an arbitrary way that does not depend on the edges with one endpoint in* Z_x *and the other in* Z_y. *For* G' *a* θ*-percolation of* G, *let*

$$N_{\eta,\gamma,z,\underline{w}}(Z_x, Z_y) := \big\{(a,b) \in Z_x \times Z_y, ab \in \mathcal{E}(G'), \mathcal{C}(ab) \leq (\underline{w}/10)^{3\mu} D^\eta\big\}. \tag{4.3}$$

Then, (and also for $\alpha = \infty$ *and/or* $\beta = \infty$ *under the convention that* $\infty \cdot 0 = 0$ *in (4.1)),*

$$\mathbb{P}\big(N_{\eta,\gamma,z,\underline{w}}(Z_x, Z_y) = \emptyset \mid V, w_V\big) \leq \exp\Big(-\theta \underline{w}^{-2(\tau-1)} D^{\Lambda(\eta,z) - 2\gamma d\delta/3}\Big). \tag{4.4}$$

Note that the requirements on D and $\underline{w}$ can be simultaneously satisfied since $w_0 \vee 4(c_H+2)^d \vee 4000c_1^{-1/(\mu\beta)}$ is a large constant, while D grows at least as $\Theta(\log\log\xi)$ with ξ. The restriction of the vertex-weights to be between constant factors of $wD^{z/2}$ in (4.2) ensures that we can lower-bound edge-presence and also upper bound the cost-penalisation $(w_a w_b)^\mu$ on these edges.

Proof. We first bound the number of possible edges in $N_{\eta,\gamma,z,\underline{w}}(Z_x, Z_y)$ from below. We make use of the net property: the assumption $x, y \in \mathcal{N}$ ensures enough vertices in the sets $\mathcal{Z}(x), \mathcal{Z}(y)$. We check if all conditions are satisfied: In Def. 2.1, we will take $\varepsilon = \delta/4$, $w = \underline{w}D^{z/2}$ and $r = D^\gamma$. We have $\underline{w}D^{z/2} \geq \underline{w} \geq w_0$. Since $2d\gamma > z(\tau-1)$, we have $z/2 < d\gamma/(\tau-1)$, and $\underline{w} \leq D^\delta$ by hypothesis; for sufficiently small δ, it follows that $\underline{w}D^{z/2} < D^\delta \cdot D^{d\gamma/(\tau-1)-2\delta} \leq D^{d\gamma/(\tau-1)-\delta/4}$; thus the requirement on $\underline{w}$ of Def. 2.1 is satisfied. Also $D^\gamma \in [(\log\log\xi\sqrt{d})^{16/\delta}, \xi\sqrt{d}]$ by hypothesis. Thus (2.2) gives for $v \in \{x, y\}$:

$$|\mathcal{Z}(v)| \geq D^{\gamma d(1-\delta/4)} \ell(\underline{w}D^{z/2}) \underline{w}^{-(\tau-1)} D^{-z(\tau-1)/2}.$$

Since $\underline{w}D^{z/2} \leq D^{\delta+z/2}$ and $D \gg_\star \delta, z$, by Potter's bound $|\mathcal{Z}(v)| \geq \underline{w}^{-(\tau-1)} D^{d\gamma - z(\tau-1)/2 - 3\delta\gamma d/10}$, so $|Z_v| \geq |\mathcal{Z}(v)|/4 \geq \underline{w}^{-(\tau-1)} D^{d\gamma - z(\tau-1)/2 - 3\delta\gamma d/10}/4$ for $v \in \{x, y\}$. Accounting for the possibility of even full overlap between Z_x and Z_y, we obtain

$$\big|\{\{a,b\} : a \in Z_x, b \in Z_y\}\big| \geq \frac{(|Z_x| \wedge |Z_y|)^2}{4} \geq \underline{w}^{-2(\tau-1)} D^{2d\gamma - z(\tau-1) - 3\delta\gamma d/5}/64. \tag{4.5}$$

We now lower-bound the probability that $a \in Z_x, b \in Z_y$ forms a low-cost edge as in (4.3). By hypothesis $|x-y| \leq c_H D$, $a \in B_{D^\gamma}(x)$, $b \in B_{D^\gamma}(y)$, and $\gamma < 1$, so $|a-b| \leq c_H D + 2D^\gamma \leq (c_H+2)D$. Since

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

$w_a, w_b \in [\underline{w}D^{z/2}/2, 2\underline{w}D^{z/2}]$ by (4.2), it follows from (1.5) that

$$\mathbb{P}\big(ab \in \mathcal{E}(G') \mid V, w_V\big) \geq \theta\underline{c}\left(1 \wedge \frac{\underline{w}^2 D^z}{4(c_H+2)^d D^d}\right)^{\alpha} \geq \theta\underline{c}\big(1 \wedge \underline{w}D^{z-d}\big)^{\alpha}, \tag{4.6}$$

where we used the assumption that $\underline{w} \geq 4(c_H+2)^d$ to simplify the formula. Since $z \in [0, d]$, and δ is small relative to z, if $z - d < 0$ then we may assume $z - d \leq -2\delta$. Since $1 \leq \underline{w} \leq D^{\delta}$, the minimum in (4.6) is attained at 1 only for $z = d$. So, for all $\{a, b\} \in Z_x \times Z_y$,

$$\mathbb{P}\big(ab \in E(G') \mid V, w_V\big) \geq \begin{cases} \theta\underline{c}\mathbb{1}\{z = d\} & \text{if } \alpha = \infty, \\ \theta\underline{c}D^{\alpha(z-d)} & \text{otherwise.} \end{cases} \tag{4.7}$$

Since $w_a, w_b \leq 2\underline{w}D^{z/2}$ by (4.2),

$$\begin{aligned}\mathbb{P}\big(\mathcal{C}(ab) \leq (\underline{w}/10)^{3\mu}D^{\eta} \mid ab \in \mathcal{E}(G'), V, w_V\big) &\geq \mathbb{P}\big((4\underline{w}^2D^z)^{\mu}L_{ab} \leq (\underline{w}/10)^{3\mu}D^{\eta}\big) \\ &= F_L(4000^{-\mu}\underline{w}^{\mu}D^{\eta-\mu z}).\end{aligned}$$

If $\eta < \mu z$, then since $\delta \ll_\star z, \eta, \texttt{par}$ we may assume that $\eta - \mu z \leq -2\mu\delta$. Since $\underline{w} \leq D^{\delta}$ and $D \gg_\star \delta$, it follows that $4000^{-\mu}\underline{w}^{\mu}D^{\eta-\mu z} \leq D^{-\mu\delta} \leq t_0$ and hence using Assumption 1.2 and the assumption $\underline{w} \geq 4000c_1^{-1/(\mu\beta)}$ we get $F_L(4000^{-\mu}\underline{w}^{\mu}D^{\eta-\mu z}) \geq D^{\beta(\eta-\mu z)}$ after simplifications. If instead $\eta \geq \mu z$, then $F_L(4000^{-\mu}\underline{w}^{\mu}D^{\eta-\mu z}) \geq F_L(4000^{-\mu}\underline{w}^{\mu}) \geq 1/2$ by hypothesis. Summarising the two cases with indicators we arrive at

$$\mathbb{P}\big(\mathcal{C}(ab) \leq (\underline{w}/10)^{3\mu}D^{\eta} \mid ab \in \mathcal{E}(G'), V, w_V\big) \geq \begin{cases} \mathbb{1}\{\eta \geq \mu z\}/2 & \text{if } \beta = \infty, \\ D^{0\wedge\beta(\eta-\mu z)}/2 & \text{otherwise.} \end{cases} \tag{4.8}$$

With the convention that $\infty \cdot 0 = 0$, the second row equals the first row in both (4.7) and (4.8). Combining (4.7) and (4.8), we obtain that for all $\{a, b\} \in Z_x \times Z_y$:

$$\mathbb{P}\big(\{a, b\} \in N_{\eta,\gamma,z,\underline{w}}(x, y) \mid V, w_V\big) \geq \theta\underline{c}D^{\alpha(z-d)+(0\wedge\beta(\eta-\mu z))}/2. \tag{4.9}$$

Given V, w_V, the possible edges $\{a, b\}$ lie in $N_{\eta,\gamma,z,\underline{w}}(x, y)$ independently. Hence by (4.5) and (4.9), $|N_{\eta,\gamma,z}(x, y)|$ stochastically dominates a binomial random variable whose mean m is the product of the two equations' right-hand sides. On bounding $\underline{c}/128 \geq D^{-\delta\gamma d/15}$, we obtain

$$m \geq \theta\underline{w}^{-2(\tau-1)}D^{\Lambda(\eta,z)-2\delta\gamma d/3}.$$

Inequality (4.4) follows since this binomial variable is zero with probability at most e^{-m}. □

The next lemma finds a low-cost edge from a fix vertex in $\mathcal{N}$ with weight roughly M to some nearby vertex in $\mathcal{N}$ with weight roughly K. We will use this lemma later in two different ways, either K being much lower than M; or K being somewhat larger than M. The former corresponds to the shorter edges emanating from the longest edge on Figure 2 and ensure that the endpoints of the three-edge bridge paths; after many iterations, can be connected at low costs using 'local edges'.

Lemma 4.3 (Single cheap edge nearby). *Consider Setting 4.1. Let $M > 1$, and let $x \in \mathcal{N}$ be a vertex with $w_x \in [\frac{1}{2}M, 2M]$. Let $U, D > 0$ and $K > w_0$, and define the event*

$$\mathcal{A}_{K,D,U}(x) := \big\{\exists y \in \mathcal{N} \cap (B_D(x) \times [\tfrac{1}{2}K, 2K]) : xy \in \mathcal{E}(G'),\ \mathcal{C}(xy) \leq U\big\}. \tag{4.10}$$

Suppose that $\delta \ll_\star \texttt{par}$, that $K, M, D \gg_\star \delta, w_0$, and that

$$(\log\log \xi\sqrt{d})^{16/\delta} \leq (D \wedge (KM)^{1/d})/4^{1/d} \leq \xi\sqrt{d}, \tag{4.11}$$

$$K \leq D^{d/(\tau-1)-\delta} \wedge M^{1/(\tau-2+\delta\tau)}. \tag{4.12}$$

Then if $\beta = \infty$ *and* $U(KM)^{-\mu} \gg_\star \texttt{par}$*, or alternatively if* $\beta < \infty$*, then*

$$\mathbb{P}\big(\mathcal{A}_{K,D,U}(x) \mid V, w_V\big) \geq 1 - \exp\Big(-\theta K^{-(\tau-1)}(D^d \wedge KM)^{1-\delta}(1 \wedge (U(KM)^{-\mu})^\beta)\Big). \tag{4.13}$$

The required condition $U(KM)^{-\mu} \gg_\star \texttt{par}$ when $\beta = \infty$ ensures that when $\beta = \infty$, the minimum is at 1 in the last factor in the exponent on the right-hand side of (4.13).

Proof. Let $r = 4^{-1/d}(D \wedge (KM)^{1/d})$, and define $\mathcal{Z}'(x) := \mathcal{N} \cap (B_r(x) \times [\frac{1}{2}K, 2K])$. We will first lower-bound $|\mathcal{Z}'(x)|$ by applying Definition 2.1 with $\varepsilon = \delta/4$, $w = K$ and the same value of r, that is, $r = 4^{-1/d}(D \wedge (KM)^{1/d})$. Observe that (4.11) and the fact that $K \geq w_0$ imply all the requirements of Definition 2.1 except $K \leq r^{d/(\tau-1)-\delta/4}$, which we now prove. By (4.12), $M \geq K^{\tau-2+\tau\delta}$ and hence

$$\begin{aligned} r^{d/(\tau-1)-\delta/4} &\geq (KM/4)^{1/(\tau-1)-\delta/(4d)} \geq (K^{\tau-1+\tau\delta})^{1/(\tau-1)-\delta/4}/4 \\ &= K^{1+\delta(\tau/(\tau-1)-(\tau-1)/4-\tau\delta/4)}/4. \end{aligned}$$

Since $\tau < 3$ and $\delta \ll_\star \tau$, the coefficient of δ is positive in the exponent so the right-hand side is at least K, as required by Def. 2.1. Applying (2.2) followed by Potter's bound (since $D, K \gg_\star \delta$) yields that

$$|\mathcal{Z}'(x)| \geq \ell(K)K^{-(\tau-1)}r^{d(1-\delta/4)} \geq K^{-(\tau-1)}\big(D \wedge KM\big)^{1-\delta/2}. \tag{4.14}$$

We now lower-bound the probability that for a $y \in \mathcal{Z}'(x)$ the edge xy is present and has cost at most U, satisfying the requirements of $\mathcal{A}_{K,D,U}(x)$. Let $y \in \mathcal{Z}'(x)$. Since $w_x \in [M/2, 2M]$ and $w_y \in [K/2, 2K]$, by (1.5) and the definition of $r = D \wedge (KM/4)^{1/d}$ we have

$$\mathbb{P}\big(xy \in \mathcal{E} \mid V, w_V\big) \geq \theta\underline{c}(1 \wedge KM/(4r^d))^\alpha = \theta\underline{c}, \tag{4.15}$$

since the minimum is at the first term; also for $\alpha = \infty$. Moreover, if $\beta < \infty$, we apply (1.3); otherwise, since $U(KM)^{-\mu}$ is large, $F_L(U(4KM)^{-\mu}) \geq 1/2$, to estimate the cost

$$\begin{aligned} \mathbb{P}\big(\mathcal{C}(xy) \leq U \mid xy \in \mathcal{E}(G'), V, w_V\big) &\geq \mathbb{P}\big((4KM)^\mu L_{xy} \leq U\big) = F_L((4KM)^{-\mu}U) \\ &\geq C(1 \wedge (U(KM)^{-\mu})^\beta), \end{aligned} \tag{4.16}$$

for an appropriate choice of $C > 0$ depending only on $\texttt{par}$. Combining (4.15) and (4.16), we obtain for any $y \in \mathcal{Z}'(x)$:

$$\mathbb{P}\big(xy \in \mathcal{E}(G') \text{ with } \mathcal{C}(xy) \leq U \mid V, w_V\big) \geq \theta C\underline{c}(1 \wedge (U(KM)^{-\mu})^\beta). \tag{4.17}$$

Conditioned on (V, w_V), edges are present independently, so the number of low-cost edges between x and $\mathcal{Z}'(x)$ stochastically dominates a binomial random variable with parameters the right-hand side of (4.14) and (4.17). The mean is

$$\theta C\underline{c}K^{-(\tau-1)}(D^d \wedge KM)^{1-\delta/2}(1 \wedge (U(KM)^{-\mu})^\beta).$$

Since $K, M, D \gg_\star \delta$, we may swallow the constant factor $\theta C\underline{c}$ by increasing $\delta/2$ to δ. The result follows since for a binomial variable Z, $\mathbb{P}(Z = 0) \leq \exp(-\mathbb{E}[Z])$. □

The third lemma builds cheap *weight-increasing paths*, from a low-weight vertex in $\mathcal{N}$ to a high-weight vertex in $\mathcal{N}$. The proof is via repeated application of Lemma 4.3. The starting point of these weight-increasing paths shall be the endpoints of the 3-edge bridge paths depicted on Figure 2, and we will use them to partially fill in the 'gaps' between the 3-edge bridge paths.

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

Lemma 4.4 (Weight-increasing paths). *Consider Setting 4.1. Let $M > 1$, and let y_0 be a vertex in $\mathcal{N}$ with weight in $[\frac{1}{2}M, 2M]$. Let $K, D > 1$, $U \geq K^{2\mu}$, and let*

$$q := \left\lceil \frac{\log(\log K/\log M)}{\log(1/(\tau - 2 + 2d\tau\delta))} \right\rceil. \tag{4.18}$$

Let $\mathcal{A}_{\pi(y_0)}$ be the event that G' contains a path $\pi = y_0y_1\dots y_q$ contained in $\mathcal{N} \cap B_{qD}(y_0)$ such that $W_{y_q} \in [\frac{1}{2}K, 2K]$ and $\mathcal{C}(\pi) \leq qU$. Suppose that $\delta \ll_\star \mathtt{par}$, that $K, M, D \gg_\star \theta, \delta, w_0$, and that $M \leq K \leq D^{d/2}$, $D \leq \xi\sqrt{d}$, and $(M/2)^{2/d} \geq (\log\log \xi\sqrt{d})^{16/\delta}$. Then if $\beta = \infty$ and $U(KM)^{-\mu} \gg_\star \mathtt{par}$, or if $\beta < \infty$, then

$$\mathbb{P}\big(\mathcal{A}_{\pi(y_0)} \mid V, w_V\big) \geq 1 - \exp(-\theta M^\delta). \tag{4.19}$$

Proof. We will build π vertex-by-vertex by applying Lemma 4.3 q times. We first define a doubly exponentially increasing sequence of target weights. Let $M_0 := M$, and for all $i \in [q]$, let

$$M_i := M^{1/(\tau-2+2d\tau\delta)^i} \wedge K. \tag{4.20}$$

Since $\tau < 3$ and δ is small, $\tau - 2 + 2d\tau\delta < 1$; hence on substituting the definition of q in (4.18) into (4.20) and removing the ceiling, we obtain

$$M^{1/(\tau-2+2d\tau\delta)^q} = \exp\Big(\log M \cdot \mathrm{e}^{-q\log(\tau-2+2d\tau\delta)}\Big) \geq \exp\Big(\log M \cdot \mathrm{e}^{\log\left(\frac{\log K}{\log M}\right)}\Big) = K,$$

and hence $M_q = K$. By a very similar argument, $M_{q-1} < K$. We now define $Y_0 = y_0$, and define an arbitrary ordering on $\mathcal{N}$. For all $i \in [q]$, we define Y_i to be the first vertex in $\mathcal{N}$ in $B_D(Y_{i-1}) \times [\frac{1}{2}M_i, 2M_i]$ with the property that the edge $Y_{i-1}Y_i$ is present in G' and has cost at most U. If no such vertex exists, we define $Y_j = \mathtt{None}$ for all $j \geq i$. Let $\mathcal{A}_i$ be the event that $Y_0, \dots, Y_i \neq \mathtt{None}$. Then, if $\mathcal{A}_q$ occurs, the path $\pi = Y_0 \dots Y_q$ yields $V(\pi) \subseteq \mathcal{N} \cap B_{qD}(y_0)$ and $\mathcal{C}(\pi) \leq qU$, and that $w_{Y_q} \in [\frac{1}{2}K, 2K]$ since $M_q = K$. So, (and because $\mathcal{A}_{i-1} \subseteq \mathcal{A}_i$),

$$\mathbb{P}\big(\mathcal{A}_{\pi(y_0)} \mid V, w_V\big) \geq \mathbb{P}\big(\mathcal{A}_q \mid V, w_V\big) = \prod_{i=1}^{q} \mathbb{P}\big(\mathcal{A}_i \mid \mathcal{A}_{i-1}, V, w_V\big). \tag{4.21}$$

Our goal is to show that

$$p_i := \mathbb{P}\big(\mathcal{A}_i \mid \mathcal{A}_{i-1}, V, w_V\big) \geq 1 - \exp(-\theta M_{i-1}^{3\delta}). \tag{4.22}$$

Indeed, if this bound holds then in (4.21), we obtain that

$$\mathbb{P}\big(\mathcal{A}_q \mid V, w_V\big) \geq \prod_{i=0}^{q-1} \big(1 - \exp(-\theta M_i^{3\delta})\big) \geq 1 - \sum_{i=0}^{q-1} \exp(-\theta M_i^{3\delta}). \tag{4.23}$$

Recall that for all $2 \leq i \leq q-1$, $M_i = M_{i-1}^{1/(\tau-2+2d\tau\delta)}$, and so since δ is small and $\tau \in (2,3)$ we have $M_i \geq M_{i-1}^{1+\delta}$. Since $M_0 = M \gg_\star \delta, \theta$, we obtain $\exp(-\theta M_i^{3\delta}) \leq \frac{1}{2}\exp(-\theta M_{i-1}^{3\delta})$. It follows from (4.23) that

$$\mathbb{P}\big(\mathcal{A}_q \mid V, w_V\big) \geq 1 - \sum_{i=0}^{q-1} \exp(-\theta M_i^{3\delta}) - \exp(-M_q^{3\delta}) \geq 1 - 3\exp(-\theta M^{3\delta}) \geq 1 - \exp(-\theta M^\delta)$$

as required in (4.19), where the last step holds since $M \gg_\star \delta, \theta$.

We now set out to show (4.22). We simplify the conditioning in (4.21). For all $i \in [q]$, let

$$\begin{aligned}\mathcal{F}_i &:= \mathcal{E}(G') \cap \big\{\{x,y\} : x,y \in \mathcal{N},\ w_x \in [\tfrac{1}{2}M_{i-1}, 2M_{i-1}],\ w_y \in [\tfrac{1}{2}M_i, 2M_i]\big\}, \\ \mathcal{F}_{\le i} &:= (\mathcal{F}_1, \dots, \mathcal{F}_i).\end{aligned} \tag{4.24}$$

Thus, $\mathcal{F}_i$ reveals the edges of G' that are between vertices in the targeted $i-1$-st and i-th weight ranges. Observe that $\mathcal{A}_1, \dots, \mathcal{A}_i$ and $Y_1, \dots, Y_i$ are deterministic functions of $\mathcal{F}_{\le i}$. Moreover, if $F_{\le i-1}$ is a possible realisation of $\mathcal{F}_{\le i-1}$ such that $\mathcal{A}_{i-1}$ occurs on $F_{\le i-1}$, then this gives us the vertex $Y_{i-1}(F_{\le i-1})$. Conditioned on $\mathcal{F}_{\le i-1} = F_{\le i-1}$, $\mathcal{A}_i$ occurs iff $Y_i \neq \texttt{None}$, that is, there is an emanating edge from $Y_{i-1}(F_{\le i-1})$ leading to the next weight-range. Thus, with $\mathcal{A}_{K,D,U}$ from (4.10), (4.21) implies that for all $i \in [q]$,

$$\begin{aligned} p_i &\ge \min_{F_{\le i-1} :\, \mathcal{A}_{i-1} \text{ occurs on} F_{\le i-1}} \mathbb{P}\big(Y_i \neq \texttt{None} \mid \mathcal{F}_{\le i-1} = F_{\le i-1},\, V, w_V\big) \\ &= \min_{F_{\le i-1} :\, \mathcal{A}_{i-1}(F) \text{ occurs on} F_{\le i-1}} \mathbb{P}\big(\mathcal{A}_{M_i,D,U}(Y_{i-1}(F_{\le i-1})) \text{ occurs} \mid \mathcal{F}_{\le i-1} = F_{\le i-1},\, V, w_V\big). \end{aligned} \tag{4.25}$$

Now observe that since $\tau \in (2,3)$, δ is small, and $M \gg_\star \delta$, for all $i \in [q-2]$, we may assume $M_i = M_{i-1}^{1/(\tau-2+2d\tau\delta)^i} > 4M_{i-1}$, and moreover $M_q \ge M_{q-1}$. Therefore, the intervals $[\frac{1}{2}M_1, 2M_1], \dots, [\frac{1}{2}M_q, 2M_q]$ are all disjoint except possibly for $[\frac{1}{2}M_{q-1}, 2M_{q-1}]$ and $[\frac{1}{2}M_q, 2M_q]$. It follows that the variables $\mathcal{F}_1, \dots, \mathcal{F}_q$ are determined by disjoint sets of possible edges. Namely, in $\mathcal{F}_{q-1}$ we revealed edges between weights $[\frac{1}{2}M_{q-2}, 2M_{q-2}]$ and $[\frac{1}{2}M_{q-1}, 2M_q - 1]$, which are disjoint from edges between $[\frac{1}{2}M_{q-1}, 2M_{q-1}]$ and $[\frac{1}{2}M_q, 2M_q]$. So, the event $\mathcal{A}_{M_i,D,U}(Y_{i-1}(F))$ is independent of the edges revealed in $\mathcal{F}_{\le i-1} = F_{\le i-1}$ in (4.25) (conditioned on (V, w_V)), and hence

$$p_i \ge \min_{y :\, w_y \in [M_{i-1}/2, 2M_{i-1}]} \mathbb{P}\big(\mathcal{A}_{M_i,D,U}(y) \mid V, w_V\big). \tag{4.26}$$

We now apply Lemma 4.3 on the right-hand side: take there $\delta_{4.3} := \delta$, $\theta_{4.3} := \theta$, $M_{4.3} := M_{i-1}$, $K_{4.3} := M_i$, $D_{4.3} := D$ and $U_{4.3} := U$. Observe $K_{4.3}, M_{4.3} \ge M$; thus by hypothesis we have that $\delta_{4.3} \ll_\star \texttt{par}$ is small and that $K_{4.3}, M_{4.3}, D_{4.3} \gg_\star \delta, w_0$, as required by Lemma 4.3. Next, we have $(K_{4.3}M_{4.3})^\mu \le K^{2\mu}$, so if $\beta = \infty$ it follows that $U_{4.3}(K_{4.3}M_{4.3})^{-\mu} \ge UK^{-2\mu}$ is large as required, by the assumptions before (4.18). Next, $(D \wedge (M_iM_{i-1})^{1/d})/4^d \le D \le \xi\sqrt{d}$ by hypothesis, and $(D \wedge (M_iM_{i-1})^{1/d})/4^d \ge (M_0/2)^{2/d} \ge (\log\log \xi\sqrt{d})^{16/\delta}$ by hypothesis, so (4.11) holds. Next, $M_i \le M_q = K \le D^{d/2} \le D^{d/(\tau-1)-\delta}$ by hypothesis and because δ is small; and finally $M_i \le M_{i-1}^{1/(\tau-2+2d\tau\delta)} < M_{i-1}^{1/(\tau-2+\tau\delta)}$ by definition, so (4.12) holds. Thus the conditions of Lemma 4.3 all hold, and applying (4.13) to (4.26) yields that for all $i \in [q]$:

$$p_i \ge 1 - \exp\Big(-\theta M_i^{-(\tau-1)}\big[(D^d \wedge M_iM_{i-1})^{1-\delta}(1 \wedge (U(M_iM_{i-1})^{-\mu})^\beta\big]\Big). \tag{4.27}$$

Clearly $M_iM_{i-1} \le M_q^2 = K^2$; and since $K \le D^{d/2}$ by hypothesis, the first minimum is at M_iM_{i-1}, while the second minimum is taken at 1 on the right-hand side since $U \ge K^{2\mu}$ was assumed. Hence

$$p_i \ge 1 - \exp\Big(-\theta M_i^{-(\tau-2)-\delta} M_{i-1}^{1-\delta}\Big).$$

Since $M_i \le M_{i-1}^{1/(\tau-2+2d\tau\delta)}$ by (4.20), δ is small, and $\tau \in (2,3)$, after simplification the exponent of M_{i-1} is at least $\delta(\tau+1-2d\tau\delta)/(\tau-2+2d\tau\delta) \ge 3\delta$, so $p_i \ge 1 - \exp(-\theta M_{i-1}^{3\delta})$, showing (4.22). □

The last lemma allows us to find a common neighbour for two vertices with roughly the same weight if the distance between them is not too large with respect to their weights. This lemma will connect the weight increasing paths we built in the previous lemma (to partially fill gaps) and is thus responsible for the final connections to fill in the gaps between 3-edge bridge paths on Figure 2.

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

Lemma 4.5 (Common neighbour). *Consider Setting 4.1. Let* $\delta \ll_\star \texttt{par}$*, let* $c_H > 0$*, and let* $D \geq w_0^{2/d}$ *with* $D \gg_\star c_H, \delta$ *and* $D \in [(\log\log \xi\sqrt{d})^{16/\delta}, \xi\sqrt{d}]$*. Let* $x_0, x_1 \in \mathcal{N}$ *be vertices with* $w_{x_0}, w_{x_1} \in [D^{d/2}, 4D^{d/2}]$ *at distance* $|x_0 - x_1| \leq c_H D$*, and let* $\mathcal{A}_{x_0\star x_1}$ *be the event that* x_0 *and* x_1 *have a common neighbour in* G'*,* $v \in \mathcal{N} \cap B_D(x_0)$ *with* $\mathcal{C}(x_0v) + \mathcal{C}(vx_1) \leq D^{2\mu d}$*. Then*

$$\mathbb{P}\big(\mathcal{A}_{x_0\star x_1} \mid V, w_V\big) \geq 1 - \exp\Big(-\theta^2 D^{(3-\tau-2\delta)d/2}\Big). \tag{4.28}$$

Proof. We define a vertex $v \in \widetilde{\mathcal{V}}$ as *good* if $v \in \mathcal{N} \cap (B_D(x_0) \times [(c_H+1)^d D^{d/2}, 4(c_H+1)^d D^{d/2}])$; thus for $\mathcal{A}_{x_0\star x_1}$ to occur, it suffices that there is a good vertex v such that x_0vx_1 is a path of cost at most $D^{2\mu d}$ in G'. We call this a *good path*. We first lower-bound the number of good vertices. By assumption, $2(c_H+1)^d D^{d/2} \geq D^{d/2} \geq w_0$, and since $\tau < 3$, δ is small and $D \gg_\star c_H, \delta$ we have $2(c_H+1)^d D^{d/2} \leq D^{d/(\tau-1)-\delta/4}$. Since $\mathcal{N}$ is a weak $(\delta/4, w_0)$ net, by (2.4),

$$\begin{aligned} &\big|\mathcal{N} \cap (B_D(x_0) \times [(c_H+1)^d D^{d/2}, 4(c_H+1)^d D^{d/2}])\big| \\ &\qquad \geq D^{d(1-\delta/4)} \ell\big(2(c_H+1)^d D^{d/2}\big)\big(2(c_H+1)^d D^{d/2}\big)^{-(\tau-1)} \geq D^{(3-\tau-\delta)d/2}, \end{aligned} \tag{4.29}$$

where the last inequality follows by Potter's bound since $D \gg_\star c_H, \delta$. We now lower-bound the probability that for a good $v \in \mathcal{N}$, the edges x_0v, vx_1 are present and have cost at most $D^{3\mu d/2}$ in G'. Observe that $|x_1 - v| \leq |x_1 - x_0| + |x_0 - v| \leq (c_H+1)D$. Thus, by (1.5), and since G' is a θ-percolation, $\mathbb{P}(x_1v \in \mathcal{E}(G')|V, w_V) \geq \theta\underline{c}\big[1 \wedge (c_H+1)^d (D^{d/2})^2/((c_H+1)D)^d\big]^\alpha = \theta\underline{c}$, also when $\alpha = \infty$. Further, conditioned on the existence of the edge x_1v,

$$\begin{aligned} \mathbb{P}\big(\mathcal{C}(x_1v) \leq D^{3\mu d/2} \mid x_1v \in \mathcal{E}(G'), V, w_V\big) &\geq \mathbb{P}\big((16(c_H+1)^d D^d)^\mu L \leq D^{3\mu d/2}\big) \\ &= F_L(16^{-\mu}(c_H+1)^{-\mu d} D^{\mu d/2}) \geq 1/2, \end{aligned}$$

where the last inequality holds (including when $\beta = \infty$) since $D \gg_\star c_H$. Combining the two bounds, for all good vertices $v \in \widetilde{\mathcal{V}}$,

$$\mathbb{P}\big(x_1v \in \mathcal{E}(G'), \mathcal{C}(x_1v) \leq D^{3\mu d/2} \mid V, w_V\big) \geq \theta\underline{c}/2.$$

Since $|x_0 - v| \leq D$, the same lower bounds hold for the edge x_0v. The two events are independent conditioned on (V, w_V), and since $2D^{3\mu d/2} < D^{2\mu d}$, for all good vertices $v \in \widetilde{\mathcal{V}}$,

$$\mathbb{P}\big(x_0v, x_1v \in \mathcal{E}(G'), \mathcal{C}(x_0vx_1) \leq D^{2\mu d} \mid V, w_V\big) \geq \theta^2\underline{c}^2/4. \tag{4.30}$$

Conditioned on (V, w_V), the presence and cost of x_0vx_1 vs $x_0v'x_1$ are independent, so the number of good paths between x_0 and x_1 stochastically dominates a binomial random variable with parameters given by the right-hand side of (4.29) and that of (4.30). For a binomial variable Z, $\mathbb{P}(Z \neq 0) \geq 1 - \exp(-\mathbb{E}[Z])$, and so we obtain (4.28) by absorbing the constant $\underline{c}^2/4$ by replacing δ with 2δ in the exponent of D, using that $D \gg_\star \delta$. □

5. Budget travel plan: hierarchical bridge-paths

In this section, we present the main construction for the upper bounds in Theorems 1.4 and 1.6. This construction is a 'hierarchy' of cheap bridging paths connecting x and y that we heuristically described in Section 1 as the 'budget travel plan'. Here we elaborate more on the heuristics before diving into proofs.

Let U be either polynomial in $|x - y|$ (when proving Theorem 1.6) or sub-logarithmic in $|x - y|$ (when proving Theorem 1.4). We first find a 3-edge *bridging-path* $\pi_1 = x'aby'$ of cost at most U between two vertices x' and y' with weights $w_{x'}, w_{y'} \in [w_{H_1}, 4w_{H_1}]$, such that $|x - x'|$ and $|y - y'|$ are both at most $|x - y|^\gamma$ for some $\gamma \in (0, 1)$, see Figure 2(a). This reduces the original problem of connecting x and y to two instances of connecting two vertices at distance $|x - y|^\gamma$, at the additional cost of U. We then work recursively, applying the same procedure to find a bridging-path with endpoints near x and x' and another one with endpoints near y' and y, with all four distances at most $|x-y|^{\gamma^2}$, and both bridging-paths having cost at most U, obtaining the second level of the hierarchy, see Figure 2(b). The endpoints of the bridging paths always have weight in $[w_{H_1}, 4w_{H_1}]$, hence iteration is possible. By repeating the process R times we obtain a 'broken path' of bridging-paths of cost $U(1+2+\cdots+2^R)$ and 2^R gaps of length $|x|^{\gamma^R}$ between the bridging-paths. We call this 'broken path' a *hierarchy* after Biskup, who developed the one-edge bridge construction for graph distances in long range percolation in [11]. There are two reasons for having a *bridging-path* instead of a single bridge-edge. Firstly, a typical single bridge-edge ab has very high weights w_a, w_b and thus typically high cost, and most edges out of a and b to lower degree vertices also have high costs, which would cause high costs when filling the gaps. So instead we find an atypical bridge edge ab and take one of the cheapest edges to low-weight vertices nearby emanating from a and b, yielding a path of the form $(x'aby')$, with all three edges of cost $U/5$, and x', y' having low weight in $[w_{H_1}, 4w_{H_1}]$, giving a bridging path of length three. Secondly, to fill the 2^R gaps after R iterations whp, the failure probability of finding a connecting path has to be extremely low, $o(2^{-R})$. In most regimes this is impossible via short paths (e.g., length two) and low enough failure probability. Instead, we find weight increasing paths $\pi_{x'x''}$ and $\pi_{y'y''}$ (as in Lemma 4.4) of cost at most $U/5$ from each vertex x' and y' of the bridge paths $(x'aby')$ to respective vertices x'', y'' still near a and b but with much higher weights in $[w_{H_2}, 4w_{H_2}]$. The concatenated paths $(\pi_{x''x'}, x'aby', \pi_{y'y''})$ then themselves form a second hierarchy (now with bridging paths of more than 3 edges). Connecting all the new 2^R gaps whp is possible via paths of length two and cost U', which is polynomial in the distance $|x-y|^{\gamma^R}$, using Lemma 4.5. In the *polylogarithmic case*, U and U' are sublogarithmic, and the factor 2^R is of order $(\log|x-y|)^{\Delta_0+o(1)}$ and dominates the overall cost. The bottleneck in this regime is the number of gaps, whereas the bridge-paths have negligible costs. In the *polynomial case*, however, the cost of the first bridge $U=|x-y|^{\eta_0+o(1)}$ dominates, and all other costs (even with the factors 2^i) are negligible in comparison, causing the total cost to be polynomial in $|x-y|$. In both cases we could use and optimise level-dependent costs U_i, but that does not improve the statements of Theorems 1.4, 1.6.

The main technical result is the following proposition, that finds a path fully contained in the net that starts near 0 and ends near x. Section 6 shall connect 0 and x to this path at negligible cost compared to the one here. Recall $\Lambda(\eta, z)$ from (4.1) that determined whether a low-cost connecting edge exist between two balls. We define now the second exponent that will be crucial in determining whether low-cost edges can be found. Positivity of this function ensures that the high-weight vertices a, b above have an atypically cheap edge to a low-weight vertex nearby: For all $\eta > 0$, $z \geq 0$ we define

$$\Phi(\eta, z) := \left[d\gamma \wedge \frac{z}{2}\right] + \left[0 \wedge \beta\left(\eta - \frac{\mu z}{2}\right)\right]. \tag{5.1}$$

Recall from Setting 4.1 that G' is a θ-percolated GIRG on a vertex set (V, w_V), for some fixed $\theta > 0$, such that the vertex set contains a weak net $\mathcal{N}$.

Proposition 5.1 (Path from hierarchy). *propositionPropositionPathFromHierarchy Consider Setting 4.1, and let $y_0, y_1 \in \mathcal{N}$ with $|y_0 - y_1| = \xi$. Let $z \in [0, d], \eta \geq 0$. Let $0 < \delta \ll_\star \gamma, \eta, z, \mathtt{par}$ be such that $\Lambda(\eta, z) \geq 2\sqrt{\delta}$ and either $z = 0$ or $\Phi(\eta, z) \geq \sqrt{\delta}$. Let $\xi \gg_\star \gamma, \eta, z, \theta, \delta, w_0$. Let $R \geq 2$ be an integer satisfying $\xi^{\gamma^{R-1}} \geq (\log\log \xi\sqrt{d})^{16d/\delta^2}$ and $R \leq (\log\log\xi)^2$, let $\overline{w} := \xi^{\gamma^{R-1}d/2}$. Let $\mathcal{X}_{high\text{-}path} = \mathcal{X}_{high\text{-}path}(R, \eta, y_0, y_1)$ be the event that G' contains a path $\pi_{y_0^\star, y_1^\star}$ fully contained in $\mathcal{N}$ between some vertices $y_0^\star \in \mathcal{N} \cap (B_{c_H \xi^{\gamma^{R-1}}}(y_0) \times [\overline{w}, 4\overline{w}])$ and $y_1^\star \in \mathcal{N} \cap (B_{c_H \xi^{\gamma^{R-1}}}(y_1) \times [\overline{w}, 4\overline{w}])$ with cost*

$$\mathcal{C}(\pi_{y_0^\star y_1^\star}) \leq c_H 2^R \overline{w}^{4\mu} \xi^\eta = c_H 2^R |x|^{2\gamma^{R-1}d\mu + \eta} \tag{5.2}$$

and deviation $\mathrm{dev}_{y_0 y_1}(\pi_{y_0^\star y_1^\star}) \leq 3c_H \xi^\gamma$, for some constant c_H depending only on $\delta, \mathtt{par}$. Then

$$\mathbb{P}\big(\mathcal{X}_{high\text{-}path} \mid V, w_V\big) \geq 1 - 2\exp\left(-(\log\log\xi)^{13}\right); \tag{5.3}$$

under the convention that $\infty \cdot 0 = 0$, the statement is also valid when $\alpha = \infty$ or $\beta = \infty$.

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

The constant c_H can be found below in (5.32). We postpone the proof of this proposition and show how to obtain the cost of the optimal paths from it.

5.1. Cost optimisation of the constructed paths

In this section we apply Proposition 5.1 and optimise the cost of the path $\pi_{y_0^\star y_1^\star}$ constructed there, yielding either polylogarithmic (Corollary 5.2) or polynomial cost-distances (Corollary 5.3). The cost of $\pi_{y_0^\star y_1^\star}$ will dominate the cost of the eventual path between $0, x$. These corollaries are rather immediate: we choose appropriate values of γ, η, z, R, apply Proposition 5.1, and read off the cost of $\pi_{y_0^\star y_1^\star}$ in (5.2). There are four possible optimal choices of γ, η, z, R depending on the model parameters, and verifying that the conditions of Setting 4.1 and Proposition 5.1 hold for these choices and calculating the resulting path's cost requires some work. Thus, we defer a formal proof of Corollaries 5.2 and 5.3 to Appendix A.2, and instead focus on why these four optimisers arise and what they mean on a qualitative level.

Thus, in Proposition 5.1, disregarding constant factors, our goal is to minimise the cost bound $\mathcal{C}(\pi_{y_0^\star y_1^\star}) \le 2^R\overline{w}^{4\mu}|x|^\eta = 2^R|x|^{2\gamma^{R-1}d\mu+\eta}$ by choosing γ, η, z, R optimally. Here, R is the number of iterations in the hierarchy, and hence controls the number 2^R of gaps, while γ controls the Euclidean length of the gaps and hence also the cost of joining them, with the total cost of joining a single gap being roughly $\overline{w}^{4\mu} = |x|^{2d\mu\gamma^{R-1}}$. The exponent η controls the cost of bridge-paths in the hierarchy. From the many constraints in Proposition 5.1, the following are relevant when optimising the cost of the path. The requirement $\Lambda(\eta, z) > 0$ ensures that low-cost bridging edges exist (Lemma 4.2). The requirement that either $z = 0$ or $\Phi(\eta, z) > 0$ ensures that among the many potential bridging edges a few can be extended to low-cost 3-edge bridge-paths in Lemma 5.10. The requirement $\gamma < 1$ ensures that boxes where we search for the bridging edge shrink in size, while $z \le d$ is a formal requirement for applying Lemma 4.2 to find bridging edges, which we tolerate because increasing z above d will never be optimal. Heuristically, the effect of increasing z is to increase the probability that a given bridging edge exists at the price of increasing its expected cost; at $z = d$ the existence probability is already in the interval $[\underline{c}, \overline{c}]$ and cannot be increased further, however the penalty would increase and the number of combinatorial options decrease by increasing z, which is never optimal. The other constraints of Proposition 5.1 and Setting 4.1 (such as $2d\gamma < \tau - 1$ and $R \le (\log\log|x|)^2$) never turn out to be tight for optimal choices of η, R, γ, z. Recall $\mu_{\log}, \mu_{\mathrm{pol}}$ from (1.8).

Corollary 5.2 (Path with polylogarithmic cost). *Consider* 1*-FPP in Definition 1.1 on the graphs IGIRG or SFP satisfying the assumptions given in (1.6)–(1.3) with $d \ge 1, \alpha \in (1, \infty], \tau \in (2, 3), \mu > 0$. Let $\underline{c}, \overline{c}, h, L, c_1, c_2, \beta$ be as in (1.5)–(1.3), we allow $\beta = \infty$ and/or $\alpha = \infty$. Let $q, \varepsilon, \zeta \in (0, 1)$, let $0 < \delta \ll_\star \varepsilon, q, \mathtt{par}$, and let $w_0 > 1$. Fix a realisation (V, w_V) of $\widetilde{\mathcal{V}}$. Let $x \in V$ with $|x| \gg_\star q, \delta, \varepsilon, \zeta, w_0, \mathtt{par}$. Let Q be a cube of side length $|x|$ containing 0 and x, and assume that (V, w_V) is such that Q contains a weak $(\delta/4, w_0)$-net $\mathcal{N}$ with $0, x \in \mathcal{N}$ given in Definition 2.1. Let $G \sim \{\mathcal{G} \mid V, w_V\}$. Let $\mathcal{X}_{\mathrm{polylog}}(0, x)$ be the event that G contains a path π, fully contained in $\mathcal{N}$, with endpoints say $y_0^\star, y_x^\star$, with the following properties:*

$$w_{y_0^\star}, w_{y_x^\star} \in [\overline{w}, 4\overline{w}], \quad \textit{where} \quad \overline{w} \in [\log\log|x|, (\log|x|)^{\varepsilon}], \tag{5.4}$$

$$y_0^\star \in B_{\overline{w}^{3/d}}(0) \qquad \textit{and} \qquad y_x^\star \in B_{\overline{w}^{3/d}}(x), \tag{5.5}$$

$$\mathcal{C}(\pi) \le (\log|x|)^{\Delta_0+\varepsilon}, \qquad \textit{and} \qquad \mathrm{dev}_{0x}(\pi) \le \zeta|x|, \tag{5.6}$$

where Δ_0 is defined in (1.9), (1.17) or (1.20) depending on whether $\alpha, \beta < \infty$, $\alpha = \infty$ or $\beta = \infty$. If either $\alpha \in (1, 2)$ or $\mu \in (\mu_{\mathrm{expl}}, \mu_{\log})$ or both hold, then $\mathbb{P}(\mathcal{X}_{\mathrm{polylog}}(0, x) \mid V, w_V) \ge 1 - q$.

Sketch of proof. Corollary 5.2 covers the polylogarithmic regime, which corresponds to solutions where $\eta = 0$ is possible – such solutions exists when either $\alpha \in (1, 2)$ or $\mu < \mu_{\log}$. When $\eta = 0$, the cost of the

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

path $\pi_{y_0^\star, y_x^\star}$ is dominated by the cost $2^R|x|^{2\mu\gamma^{R-1}}$ of joining gaps. Given γ, this has minimum $2^{(1+o(1))R}$ when setting $R = (1-o(1))\log\log|x|/\log(1/\gamma)$. To minimise the cost further, we must therefore minimise $\gamma \in (0,1)$ subject to the constraints $z \in [0,d], \Lambda(0,z) > 0$, and either $z = 0$ or $\Phi(0,z) > 0$. This problem turns out to have two potentially optimal solutions corresponding to two possible strategies for finding bridging edges, with the optimal choice depending on the values of α, τ, β, μ. One possible solution – which only exists when $\alpha \in (1,2)$ – takes $\gamma = \alpha/2 + o(1)$ and $z = 0$, so that bridging edges are unusually long-range edges between pairs of low-weight vertices, yielding total path cost $(\log|x|)^{\Delta_\alpha}$ with $\Delta_\alpha = 1/(1-\log_2\alpha)$, see Claim A.5. The other possible solution – which only exists when $\mu < \mu_{\log}$ – takes $\gamma = (\tau-1+\mu\beta)/2+o(1)$ and $z = d$, so that bridging edges are unusually low-cost edges between pairs of high-weight vertices and the total path cost is $(\log x)^{\Delta_\beta}$ with $\Delta_\beta = 1/(1-\log_2(\tau-1+\mu\beta))$, see Claim A.6. The proof is in Appendix A.2.

If both $\alpha = \beta = \infty$, then the conditions of Corollary 5.2 cannot be satisfied. Indeed, when $\alpha = \infty$ then $\alpha \in (1,2)$ is not satisfied. Since $\mu_{\mathrm{expl}} = \mu_{\log} = 0$ by (1.19) when $\alpha = \beta = \infty$, so $\mu \in (\mu_{\mathrm{expl}}, \mu_{\log})$ can also not be satisfied.

Corollary 5.3 (Path with polynomial cost). *Consider* 1*-FPP in Definition 1.1 on the graphs IGIRG or SFP satisfying the assumptions given in (1.6)–(1.3) with $d \geq 1, \alpha \in (1,\infty], \tau \in (2,3), \mu > 0$. Let $\underline{c}, \overline{c}, h, L, c_1, c_2, \beta$ be as in (1.5)–(1.3), we allow $\beta = \infty$ and/or $\alpha = \infty$. Let $q, \varepsilon, \zeta \in (0,1)$, and let $0 < \delta \ll_\star \varepsilon, q, \mathtt{par}$, and $w_0 > 1$. Fix a realisation (V, w_V) of $\widetilde{\mathcal{V}}$. Let $x \in V$ with $|x| \gg_\star q, \delta, \varepsilon, \zeta, w_0, \mathtt{par}$. Let Q be a cube of side length $|x|$ containing 0 and x, and assume that (V, w_V) is such that Q contains a weak $(\delta/4, w_0)$-net $\mathcal{N}$ with $0, x \in \mathcal{N}$ given in Definition 2.1. Let $G \sim \{\mathcal{G} \mid V, w_V\}$. Let $\mathcal{X}_{\mathrm{pol}}(0,x)$ be the event that G contains a path π, fully contained in $\mathcal{N}$, with endpoints say $y_0^\star, y_x^\star$, with the following properties:*

$$w_{y_0^\star}, w_{y_x^\star} \in [\overline{w}, 4\overline{w}], \quad \textit{where} \quad \overline{w} \in [\log\log|x|, |x|^\varepsilon], \tag{5.7}$$

$$y_0^\star \in B_{\overline{w}^{3/d}}(0) \qquad \textit{and} \qquad y_x^\star \in B_{\overline{w}^{3/d}}(x), \tag{5.8}$$

$$\mathcal{C}(\pi) \leq |x|^{\eta_0+\varepsilon}, \qquad \textit{and} \qquad \mathrm{dev}_{0x}(\pi) \leq \zeta|x|, \tag{5.9}$$

where η_0 is defined in (1.10), (1.16), (1.19), or (1.21) depending on $\alpha, \beta < \infty$, $\alpha = \infty$, $\beta = \infty$, or $\alpha = \beta = \infty$. If both $\alpha > 2$ and $\mu \in (\mu_{\log}, \mu_{\mathrm{pol}}]$ hold then $\mathbb{P}(\mathcal{X}_{\mathrm{pol}}(0,x) \mid V, w_V) \geq 1-q$.

Sketch of proof. Corollary 5.3 covers the polynomial regime, which corresponds to solutions where only $\eta > 0$ is possible, that is, when $\alpha > 2$ and $\mu > \mu_{\log}$. Here, on taking R to be a suitably large constant, the cost bound on the path $2^R|x|^{2\mu\gamma^{R-1}}|x|^\eta = |x|^{\eta+o(1)}$, which is roughly the cost of the very first bridging edge. Our goal is thus to minimise η under the constraints that $\Lambda(\eta,z) > 0$, $z \in [0,d]$, $\gamma \in (0,1)$, and either $z = 0$ or $\Phi(\eta,z) > 0$. (4.1) and (5.1) show that both Φ and Λ are increasing functions of γ; thus we can take $\gamma = 1-o(1)$. As in the polylogarithmic regime, this minimisation problem has two potentially optimal solutions. One possible solution – which exists when $\mu \leq \mu_{\mathrm{pol},\beta}$ – takes $z = d$ and gives $\eta = \mu d - (3-\tau)d/\beta + o(1)$, so that bridging edges are unusually low-cost edges between pairs of high-weight vertices. The total path cost is then $|x|^{\eta_\beta+o(1)}$ with $\eta_\beta = d(\mu-(3-\tau)/\beta)$ (see Claim A.9). The other possible solution – which exists when $\mu \leq \mu_{\mathrm{pol},\alpha}$ – takes z to be as small as possible, so that bridging edges are unusually long-range edges between pairs of relatively low-weight vertices. However, when $\alpha > 2$, there are no bridging-edges between constant weight vertices, and the minimal z where bridging-edges appear is $z = d(\alpha-2)/(\alpha-(\tau-1)) + o(1) = 1/\mu_{\mathrm{pol},\alpha} + o(1)$, that is, between vertices of weight $|x|^{1/(2\mu_{\mathrm{pol},\alpha})+o(1)}$. This gives cost-exponent $\eta_\alpha := \mu/\mu_{\mathrm{pol},\alpha}$ and total cost $|x|^{\mu/\mu_{\mathrm{pol},\alpha}+o(1)}$ (see Claim A.10). Whenever a solution exists among the above two possibilities, it gives an exponent η at most 1. So, whenever $\mu \leq \max\{\mu_{\mathrm{pol},\alpha}, \mu_{\mathrm{pol},\beta}\}$, we obtain the cost bound $|x|^{\min\{\eta_\beta,\eta_\alpha\}+o(1)}$, which gives the definition of η_0 in (1.10). The proof is in Appendix A.2.

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

5.2. Constructing the hierarchy.

We now set out to prove Proposition 5.1 in several steps. We start by formally defining the concept of a *hierarchy* including edge-costs. In the rest of the paper, the symbol σ denotes an index $\sigma = \sigma_1\sigma_2\dots\sigma_R \in \{0,1\}^R$ indicating the place of a vertex in the hierarchy. This can be viewed as the Ulam-Harris labelling of the leaves of a binary tree of depth R, for example, $\sigma = 1001$ corresponds to the leaf that we reach by starting at the root and then moving to the right child, the left child twice, and the right child again. We denote the string formed by concatenating σ' to the end of σ by $\sigma\sigma'$. We 'pad' strings of length less than R by adding copies of their last digit or its complement via the T and T^c operations we now define (and discuss further below):

Definition 5.4 (Binary strings). For $\sigma = \sigma_1\dots\sigma_i \in \{0,1\}^i$ for some $i \geq 1$, we define $\sigma T := \sigma_1\dots\sigma_i\sigma_i \in \{0,1\}^{i+1}$, while $\sigma T_0 := \sigma$, and $\sigma T_k := (\sigma T_{k-1})T$ for any $k \geq 2$. Let $0_i := 0T_{i-1}$ and $1_i := 1T_{i-1}$ be the strings consisting of i copies of 0 and 1, respectively. Fix an integer $R \geq 1$. Define the *equivalence relation* $\sim_T$ on $\cup_{i=1}^R\{0,1\}^i$, where $\sigma \sim_T \sigma'$ if either $\sigma T_k = \sigma'$ or $\sigma' T_k = \sigma$ for some $k \geq 0$, with $\{\sigma\}$ be the equivalence class of σ. Let

$$\Xi_i := \{\sigma \in \cup_{j=i}^R\{0,1\}^j : \sigma_{i-1} \neq \sigma_i, \sigma_j = \sigma_i \ \forall j \geq i\}, \qquad \Xi_0 := \{\emptyset\},$$

with $\emptyset$ the empty string. We say that $\{\sigma\}$ *appears first on level* i if any (the shortest) representative of the class $\{\sigma\}$ is contained in Ξ_i.

For $\sigma = \sigma_1\dots\sigma_i \in \{0,1\}^i$ for some $i \geq 1$, we define $\sigma T^c := \sigma_1\dots\sigma_i(1-\sigma_i) \in \{0,1\}^{i+1}$. For $\sigma \in \Xi_i$, we say that $(\sigma T_{j-1})T^c \in \{0,1\}^{i+j}$ is the *level-(i+j) sibling* of $\{\sigma\}$. We say that two strings in level i are *newly appearing cousins* on level i if they are of the forms $\sigma 01$ and $\sigma 10$ for some $\sigma \in \{0,1\}^{i-2}$.

The inverse of the operator T 'cuts off' all but one of the identical last digits from a $\sigma \in \{0,1\}^R$, hence, each class $\{\sigma\}$ has exactly one representative in $\{0,1\}^R$, and the number of equivalence classes is 2^R. For $i > 1$, there are exactly 2^{i-1} equivalence classes that first appear on level i (i.e., the shortest representative of the class is in Ξ_i), and (since $0, 1 \in \Xi_1$) the total number of equivalence classes that appear until level i is 2^i. To show an example of the sibling relationship, for example, $01111 \sim 01$ belongs to Ξ_2, and the level-3 sibling of $\{01\}$ is 010, and the level-$(2+j)$ sibling of $\{01\}$ is 01_j0. Similarly, 010 and 001 are newly appearing level-3 cousins, and on level i, there are 2^{i-2} pairs of newly appearing cousins.

The hierarchy embeds each equivalence class $\{\sigma\} \in \cup_{j=1}^R\{0,1\}^R$ into the (weighted) vertex set of G so that all cousins are joined by low-cost 'bridge' paths, all siblings are close in Euclidean space, $0^R = x$ and $1^R = y$ are the vertices we start with, and the weights of all other vertices in the embedding are constrained. The Euclidean distances between siblings/cousins will decay doubly exponentially in i. We formalise the embedding in the following definition.

Definition 5.5 (Hierarchy). Consider Setting 4.1. Let $y_0, y_1 \in \widetilde{\mathcal{V}}, U, \overline{w}, c_H \geq 1$, and $R \geq 2$ be an integer. Consider a set of vertices $\{y_\sigma\}_{\sigma\in\{0,1\}^R}$, divided into *levels* $\mathcal{L}_i := \{y_\sigma : \sigma \in \Xi_i\}$ for $i \in \{1,\dots,R\}$, satisfying that $y_\sigma = y_{\sigma'}$ if $\sigma \sim_T \sigma'$. We say that $\{y_\sigma\}_{\sigma\in\{0,1\}^R} \subset \widetilde{\mathcal{V}}$ is a $(\gamma, U, \overline{w}, c_H)$*-hierarchy of depth* R with $\mathcal{L}_1 = \{y_0, y_1\}$ if it satisfies the following properties:

(H1) $W_{y_\sigma} \in [\overline{w}, 4\overline{w}]$ for all $\sigma \in \{0,1\}^R \setminus \Xi_1$.

(H2) $|y_{\sigma 0} - y_{\sigma 1}| \leq c_H|y_0 - y_1|^{\gamma^i}$ for all $\sigma \in \{0,1\}^i$, $i = 0,\dots,R-1$.

(H3) There is a set $\{P_\sigma : \sigma \in \{0,1\}^i, 0 \leq i \leq R-2\}$ of paths in G such that for all $0 \leq i \leq R-2$ and all $\sigma \in \{0,1\}^i$, P_σ connects $y_{\sigma 01}$ to $y_{\sigma 10}$. Moreover, we can partition $\bigcup_{\sigma\in\{0,1\}^R}\mathcal{E}(P_\sigma)$ into sets $\{\mathcal{E}^-(P_\sigma) : \sigma \in \{0,1\}^R\}$ in such a way that for all σ, we have $\mathcal{E}^-(P_\sigma) \subseteq \mathcal{E}(P_\sigma)$ and $\mathcal{C}(\mathcal{E}^-(P_\sigma)) \leq U$. These paths P_σ are called *bridges*.

Given a set $\mathcal{N} \subseteq \widetilde{\mathcal{V}}$, we say that a hierarchy $\{y_\sigma\}_{\sigma\in\{0,1\}^R}$ is *fully contained in* $\mathcal{N}$ if both $\{y_\sigma\}_{\sigma\in\{0,1\}^R} \subseteq \mathcal{N}$, and every vertex on the paths P_σ in (H3) lies in $\mathcal{N}$.

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

Condition *(H3)* is slightly weaker than requiring each bridge to have cost at most U. We shall construct the hierarchy via an iterative construction in Def. 3.6, using one round to embed each level $\mathcal{L}_i$. We shall use Prop. 3.9 to estimate the success probability of the whole construction, which requires that we use marginal costs, and this gives the definition of $\mathcal{E}^-(P_\sigma)$ in (3.1). Using marginal costs causes no problem, since our goal is to find a path π between y_0 and y_1 of low cost. A path π uses every edge in it once, so all the bridges P_σ together will contribute to the cost of π at most

$$\mathcal{C}(\pi) = \sum_{\sigma\in\{0,1\}^i, 0\le i\le R-2} \mathcal{C}(\mathcal{E}^-(P_\sigma)) \le (2^{R-1}-1)U.$$

Later we also need that the the hierarchy stays close to the straight line segment between the starting vertices. To track this, we have the following definition:

Definition 5.6. Given $u, v \in \mathbb{R}^d$, let $S_{u,v}$ denote the line segment between u, v. For $x \in \mathbb{R}^d$ we define the deviation $\mathrm{dev}_{uv}(x) := \min\{|x-y| : y \in S_{u,v}\}$. Given a set of vertices $\mathcal{H}$ in $\mathbb{R}^d$, we define the *deviation of* $\mathcal{H}$ *from* S_{uv} as $\mathrm{dev}_{uv}(\mathcal{H}) := \max\{\mathrm{dev}_{uv}(x) : x \in \mathcal{H}\}$. Finally, for a path $\pi = (x_1 \dots x_k)$, let the *deviation of* π be $\mathrm{dev}(\pi) := \max\{\mathrm{dev}_{x_1x_k}(x_i) : i \in [k]\}$, that is, the deviation of its vertex set from the segment between the endpoints.

Next we describe the procedure used to find the hierarchy in G. We iteratively embed the levels Ξ_i into the vertex set $\widetilde{\mathcal{V}}$. We first embed Ξ_1, by setting $0 \mapsto y_0$ and $1 \mapsto y_1$, that is, $\mathcal{L}_1 := \{y_0, y_1\}$, the two given starting vertices. Observe that this embedding trivially satisfies condition *(H2)* for $i = 0$, (i.e., $\sigma = \emptyset$ in *(H2)*) for all $c_H \ge 1$. Conditions *(H1)* and *(H3)* do not concern y_0 and y_1. In round $i+1$ we then embed all $\sigma \in \Xi_{i+1}$. Given the embedding of $\cup_{j\le i}\Xi_j$ of vertices in level $\cup_{j\le i}\mathcal{L}_j$, we will embed $\sigma \in \Xi_{i+1}$ by finding $\{y_\sigma\}_{\sigma\in\Xi_{i+1}} = \mathcal{L}_{i+1}$ as follows. For each sibling pair $\sigma0, \sigma1 \in \{0,1\}^i$, by the equivalence relation $\sim_T$ in Definition 5.4, $y_{\sigma00} = y_{\sigma0}$ and $y_{\sigma11} = y_{\sigma1}$. We then search for a pair of vertices a and b close to $y_{\sigma00}$ and $y_{\sigma11}$ respectively, so that ab is a low-cost edge (typically covering a large Euclidean distance), and both a and b have a low-cost edge to a nearby vertex with weight in $[\overline{w}, 4\overline{w}]$; we embed these latter two vertices as $y_{\sigma01}$ and $y_{\sigma10}$. The path $(y_{\sigma01}aby_{\sigma10})$ then constitutes the bridge-path P_σ required by *(H3)*. See Figure 5 for a visual explanation. We formalise our goal for this iterative construction of bridges in the following definition and lemma.

Definition 5.7 (Valid bridges). Consider Setting 4.1 and the notion of bridges in Definition 5.5, and let S be a set of edges of G. For any $D, U > 0, w \ge 1$, we say that a path $P \subseteq \mathcal{N}$ with endpoints y, y' is a (D, U, w)*-valid* bridge for x_0 and x_1 with respect to S if:

$$w_y, w_{y'} \in [w, 4w], \tag{5.10}$$

$$|x_0 - y| \le D, \quad \text{and} \quad |x_1 - y'| \le D, \tag{5.11}$$

$$\mathcal{C}(P \setminus S) \le U. \tag{5.12}$$

Lemma 5.8. *Consider Setting 4.1. Fix any ordering on $\{0,1\}^R$, and let $\{y_\sigma\}_{\sigma\in\{0,1\}^R} \subseteq \mathcal{N}$, $U > 0$, and $\overline{w} \ge 1$. For all $0 \le i \le R-2$ and $\sigma \in \{0,1\}^i$, set $y_{\sigma'} := y_\sigma$ whenever $\sigma' \sim_T \sigma$. For all $i \le R-2$, let $D_i = |y_0 - y_1|^{\gamma^i}$. Suppose that for all $0 \le i \le R-2$ and all $\sigma \in \{0,1\}^i$, there exists a bridge P_σ with endpoints $y_{\sigma01}$ and $y_{\sigma10}$ that is $(c_H D_{i+1}, U, \overline{w})$-valid for $y_{\sigma0}, y_{\sigma1} \in \{0,1\}^{i+1}$ with respect to $S = \bigcup_{\sigma'<\sigma} \mathcal{E}(P_{\sigma'})$. Then $\{y_\sigma\}_{\sigma\in\{0,1\}^R}$ is a $(\gamma, U, \overline{w}, c_H)$-hierarchy of depth R with first level $\{y_0, y_1\}$ (i.e., satisfying Definition 5.5).*

Proof. Conditions *(H1)* of Definition 5.5 is immediate from the weight constraint (5.10). *(H2)* holds for the following reason. For $\sigma0, \sigma1 \in \{0,1\}^{i+1}$, P_σ being a $(c_H D_{i+1}, U, \overline{w})$ valid bridge for $y_{\sigma0}, y_{\sigma1}$ implies by (5.11) and $y_{\sigma0} = y_{\sigma00}, y_{\sigma1} = y_{\sigma11}$ that both $|y_{\sigma00} - y_{\sigma01}|, |y_{\sigma10} - y_{\sigma11}|$ are at most $c_H|y_0 - y_1|^{\gamma^{i+1}}$ for all $\sigma \in \{0,1\}^i$. Setting now either $\sigma' := \sigma0$ or $\sigma' := \sigma1$, this is equivalent to

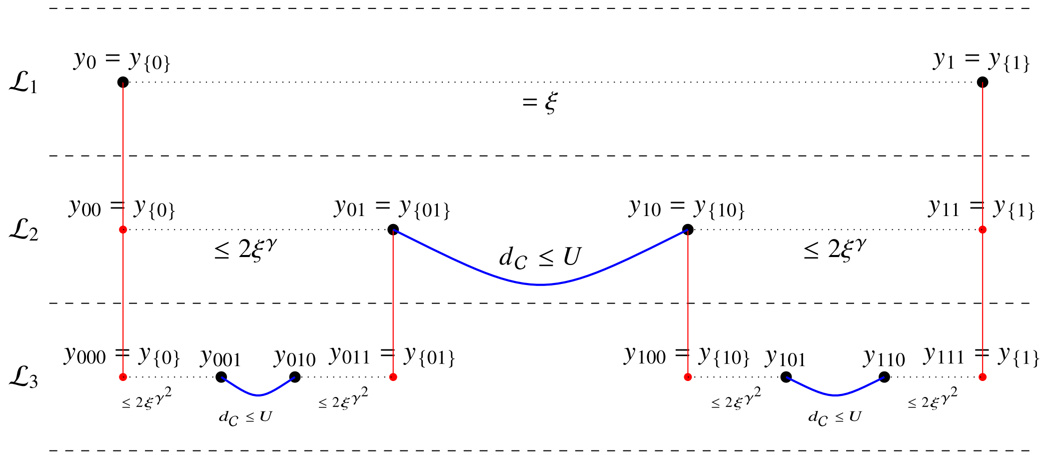

Figure 5. *A schematic representation of a $(\gamma, U, \overline{w}, 2)$-hierarchy of depth $R = 3$. The horizontal axis represents the (1-dimensional, Euclidean) distances between the vertices, while the vertical axis shows the level of the hierarchy. The weights of all vertices except y_0 and y_1 are in the interval $[\overline{w}, 4\overline{w}]$. On level 1, only the initial vertices y_0, y_1 appear and n edges. We 'push down' $y_0 = y_{00}, y_1 = y_{11}$ to level 2 (red) and we find them their respective level-2 sibling vertices y_{01} and y_{10} within Euclidean distance $2\xi^{\gamma}$, so that there is path of cost at most U between y_{01}, y_{10} (represented by the longest blue arc). Then, we 'push down' to level 3 all vertices that appeared at or before level 2, that is, $y_{000}, y_{011}, y_{100}, y_{111}$ (red), and find for each of them their level-3 siblings, that is, $y_{001}, y_{010}, y_{101}, y_{110}$, so that each vertex is within Euclidean distance $\le 2\xi^{\gamma^2}$ from its level-3 sibling, and that there is a path of cost at most U between the newly appearing cousins y_{001}, y_{010} and between y_{101}, y_{110} (represented by the two shrter blue arcs). An intuitive representation is in Figure 2.*

$|y_{\sigma'0} - y_{\sigma'1}| \le c_H |y_0 - y_1|^{\gamma^{i+1}}$ for all $\sigma' \in \{0,1\}^{i+1}, i \ge 0$, and this exactly corresponds to *(H2)*, since the inequality in *(H2)* holds for $i = 0$ trivially. Finally, condition *(H3)* follows from (5.12) by setting $\mathcal{E}^-(P_\sigma) := \mathcal{E}(P_\sigma) \setminus \bigcup_{\sigma' < \sigma} \mathcal{E}(P_{\sigma'})$. □

We now lower-bound the probability of finding a valid bridge between two fixed vertices. Recall that G' is a θ-percolation of G from Setting 4.1.

Lemma 5.9 (3-edge bridges). *Consider Setting 4.1. Let $z \in [0, d]$ and let $c_H, \eta \ge 0$. Suppose that $\delta \ll_\star \gamma, \eta, z, c_H, \mathtt{par}$ and that $D \gg_\star \gamma, \eta, z, c_H, \delta, w_0$. Suppose further that $D^\gamma \in [4^{1/d}(\log\log\xi\sqrt{d})^{16/\delta}, \xi\sqrt{d}]$ and that $\theta D^{\Lambda(\eta,z)-\sqrt{\delta}} > 1$. Suppose that $x_0, x_1 \in \mathcal{N}$ satisfy $|x_0 - x_1| \le c_H D$, and let $\underline{w} \gg_\star \delta, w_0$ satisfy $\underline{w} \in [(\log\log\xi\sqrt{d})^{16d/\delta}, D^\delta]$. Let $\mathcal{A}(x_0, x_1)$ denote the event that G' contains a bridge P that is $(2D^\gamma, 3\underline{w}^{3\mu} D^\eta, \underline{w})$-valid for x_0 and x_1 with respect to $\emptyset$, and* $\mathrm{dev}_{x_0 x_1}(P) \le 2D^\gamma$. *Finally, suppose that*

$$p(D, \underline{w}, \theta, \eta, z) := \theta \underline{w}^{-(\tau-1)} \Big(D^{d\gamma} \wedge \underline{w}^2 D^{z/2}\Big)^{1-\delta} \Big(1 \wedge \underline{w}^{\mu\beta} D^{\eta\beta - \mu\beta z/2}\Big) \ge 20^{\tau+\mu\beta}. \tag{5.13}$$

Then, with $\Lambda(\eta, z)$ from (4.1),

$$\mathbb{P}\big(\mathcal{A}(x_0, x_1) \mid V, w_V\big) \ge 1 - 3\exp\Big(-\big(\theta D^{\Lambda(\eta,z)-\sqrt{\delta}}\big)^{1/4}\Big). \tag{5.14}$$

With the convention that $\infty \cdot 0 = 0$ in (4.1), the statement is also valid when $\alpha = \infty$ or $\beta = \infty$.

When $z > 0$, the exponent of D in $p(\cdot)$ is approximating $\Phi(\eta, z)$ for small δ in (5.1). Later, $\Phi(\eta, z) > \sqrt{\delta}$ will be sufficient for the condition (5.13) to hold when $z > 0$. For $z = 0$, D does not

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

appear in the formula for $p(\cdot)$, and the exponent of $\underline{w}$ is approximating $3-\tau>0$, hence in this case the condition on $p(\cdot)$ can be satisfied by ensuring the lower bound on $\underline{w}$.

Proof of Lemma 5.9. First, when $\beta=\infty$ and $\eta<\mu z$ then $\Lambda(\eta,z)=-\infty$ in (4.1), and the condition $\theta D^{\Lambda(\eta,z)-\sqrt{\delta}}>1$ cannot hold. Hence, we can wlog assume that if $\beta=\infty$ then $\eta\ge\mu z$. We will first apply Lemma 4.3 to show that most vertices of weight roughly $\underline{w}D^{z/2}$ close to x_0 and x_1 are 'good', that is, they have a cheap edge to a vertex with weight in $[\underline{w},4\underline{w}]$. We will then apply Lemma 4.2 to find a cheap edge between some pair of good vertices.

Formally, let $I^+=[5\underline{w}D^{z/2},20\underline{w}D^{z/2}]$ and $I^-=[\underline{w},4\underline{w}]$. Note that $I_+\cap I_-=\emptyset$ for all $z\in[0,d]$. As in Lemma 4.3, for all $v\in\mathcal{N}$ let $\mathcal{A}_{2\underline{w},D^\gamma,\underline{w}^{3\mu}D^\eta}(v)=:\mathcal{A}(v)$ be the event that there is an edge of cost at most $\underline{w}^{3\mu}D^\eta$ in G' from v to a vertex $y\in\mathcal{N}\cap(B_{D^\gamma}(v)\times I^-)$. Let

$$Z_i:=\big\{v\in\mathcal{N}\cap(B_{D^\gamma}(x_i)\times I^+)\colon \mathcal{A}(v)\text{ occurs}\big\},\qquad i\in\{0,1\}. \tag{5.15}$$

The set Z_i is thus those high weight vertices near x_i that have a cheap edge to a low-weight vertex nearby. As in (4.3) of Lemma 4.2, let $N_{\eta,\gamma,z,10\underline{w}}(Z_0,Z_1)$ be the set of all edges between Z_0 and Z_1 of cost at most $\underline{w}^{3\mu}D^\eta$ and then I^+ exactly corresponds to the weight interval $[5\underline{w}D^{z/2},20\underline{w}D^{z/2}]$ as required for $Z_0\subseteq\mathcal{Z}(x_0),Z_1\subseteq\mathcal{Z}(x_1)$ in (4.2). With Z_i in (5.15), we now show that

$$\mathbb{P}\big(\mathcal{A}(x_0,x_1)\mid V,w_V\big)\ge\mathbb{P}\big(N_{\eta,\gamma,z,\underline{w}}(Z_0,Z_1)\ne\emptyset\mid V,w_V\big). \tag{5.16}$$

Indeed, suppose there exists $(a,b)\in N_{\eta,\gamma,z,10\underline{w}}(Z_0,Z_1)$. Since $a\in Z_0$, there exists $x\in\mathcal{N}\cap(B_{D^\gamma}(a)\times I^-)$ such that (x,a) is an edge of cost at most $\underline{w}^{3\mu}D^\eta$. Likewise, since $b\in Z_1$, there exists $y\in\mathcal{N}\cap(B_{D^\gamma}(b)\times I^-)$ such that (y,b) is an edge of cost at most $\underline{w}^{3\mu}D^\eta$. Since $a\in B_{D^\gamma}(x_0)$ and $b\in B_{D^\gamma}(x_1)$, by the triangle inequality, $x\in B_{2D^\gamma}(x_0)$ and $y\in B_{2D^\gamma}(x_1)$. Thus $xaby$ is a $(2D^\gamma,3\underline{w}^{3\mu}D^\eta,\underline{w})$-valid bridge with $\mathrm{dev}_{x_0x_1}\le 2D^\gamma$, as required by $\mathcal{A}(x_0,x_1)$, showing (5.16).

Now, for each $i\in\{0,1\}$, using (4.2), we set $\mathcal{Z}(x_i)=\mathcal{N}\cap(B_{D^\gamma}(x_i)\times I^+)$. For (4.4) to hold we need that $|Z_i|\ge|\mathcal{Z}(x_i)|/4$. We prove this by showing that any given vertex in $v\in\mathcal{Z}(x_i)$ lies in Z_i with probability at least $1/2$, by recalling that in (5.15), $\mathcal{A}(v)=\mathcal{A}_{2\underline{w},D^\gamma,\underline{w}^{3\mu}D^\eta}(v)=\mathcal{A}_{K,D,U}(v)$ in Lemma 4.3. Hence we set $K=2\underline{w}$, $M=10\underline{w}D^{z/2}$, $U=\underline{w}^{3\mu}D^\eta$, $D_{4.3}=D^\gamma$, and all other variables to match their current values. We check the requirements of Lemma 4.3:

By hypothesis in the statement of Lemma 5.9, δ is small; $\underline{w},D\gg_\star\delta,w_0$, and $D\gg_\star\gamma$. Since $M,K\ge\underline{w}$, it follows that $M,K,D^\gamma\gg_\star\delta,w_0$, as required above (4.11). Condition (4.11) itself holds since $(D^\gamma\wedge(20\underline{w}^2D^{z/2})^{1/d}/4^{1/d}\ge D^\gamma/4^{1/d}\wedge\underline{w}^{1/d}\ge(\log\log\xi\sqrt{d})^{16/\delta}$ by hypothesis, and $(D^\gamma\wedge(20\underline{w}^2D^{z/2})^{1/d}/4^{1/d}\le D^\gamma\le\xi\sqrt{d}$ by hypothesis. Condition (4.12) holds since $K=2\underline{w}\le 2D^\delta$ by hypothesis, so since $\delta\ll_\star\gamma$ and $D\gg_\star\delta$, we have $2\underline{w}\le D^{\gamma(d/(\tau-1)-\delta)}$, and similarly since $\tau\in(2,3)$ and $\delta\ll_\star\mathtt{par}$, $K=2\underline{w}\le(2\underline{w})^{1/(\tau-2+\tau\delta)}\le(10\underline{w}D^{z/2})^{1/(\tau-2+\delta\tau)}=M^{1/(\tau-2+\delta\tau)}$. Finally, if $\beta=\infty$, then below (4.12) we need to check $U(KM)^{-\mu}\gg_\star\mathtt{par}$. Since wlog we assumed that $\eta\ge\mu z$, clearly $\eta\ge\mu z/2$. Therefore, $U(KM)^{-\mu}=\underline{w}^{3\mu}D^\eta(20\underline{w}^2D^{z/2})^{-\mu}=20^{-\mu}\underline{w}^\mu D^{\eta-\mu z/2}\ge(\underline{w}/20)^\mu$, which is large since $\underline{w}$ is large by hypothesis. Hence, all conditions of Lemma 4.3 are met and (4.13) applies, and substituting $K=2\underline{w}$, $M=10\underline{w}D^{z/2}$, $U=\underline{w}^{3\mu}D^\eta$ there, the exponent on the right-hand side of (4.13) in our setting becomes

$$-2^{-(\tau-1)}\theta\underline{w}^{-(\tau-1)}\big(D^{d\gamma}\wedge 20\underline{w}^2D^{z/2}\big)^{1-\delta}\Big(1\wedge(\underline{w}/20)^{\mu\beta}D^{\eta\beta-\mu\beta z/2}\Big),$$

where we recognise that this matches $p(\cdot)$ from (5.13) up to a factor of at most $20^{1-(\tau+\mu\beta)}$. Since we assumed $p(\cdot)\ge 20^{\tau+\mu\beta}$ in (5.13), for any vertex $v\in\mathcal{Z}(x_i)=\mathcal{N}\cap(B_{D^\gamma}(x_i)\times I^+)$,

$$\mathbb{P}\big(v\in Z_i\mid V,w_V\big)=\mathbb{P}(\mathcal{A}(v)\mid V,w_V)\ge 1-e^{-20}>1/2. \tag{5.17}$$

Since I^+ and I^- are disjoint, the events $\mathcal{A}(v),\mathcal{A}(v')$ are functions of disjoint edge sets and are therefore mutually independent conditioned on (V,w_V). Hence, for $i\in\{0,1\}$, $|Z_i|$ is dominated below by a

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

binomial variable with mean $|\mathcal{Z}(x_i)|/2$. By the standard Chernoff bound (Theorem A.1 with $\lambda = 1/2$),

$$\mathbb{P}\big(|Z_i| < |\mathcal{Z}(x_i)|/4 \mid V, w_V\big) \le e^{-|\mathcal{Z}(x_i)|/16}. \tag{5.18}$$

To bound $|\mathcal{Z}(x_i)|$ below in (5.18), we will use that $x_0, x_1 \in \mathcal{N}$ and $\mathcal{N}$ is a weak $(\delta/4, w_0)$-net as assumed in Setting 4.1, and apply (2.2). We check if the conditions to apply (2.2) in Def. 2.1 hold. Since $\mathcal{Z}(x_i) = \mathcal{N} \cap (B_{D^\gamma}(x_i) \times [5\underline{w}D^{z/2}, 20\underline{w}D^{z/2}])$, we set there $r = D^\gamma$ and $w = 10\underline{w}D^{z/2}$, and we must bound $10\underline{w}D^{z/2}$ above and below. Recall that by hypothesis, $\theta D^{\Lambda(\eta,z)-\sqrt{\delta}} > 1$; this implies $\Lambda(\eta, z) > 0$ and hence $2d\gamma > z(\tau-1)$ using (4.1). Since $\delta \ll_\star \gamma, z$, we may therefore assume $z/2 \le d\gamma/(\tau-1) - 2\delta$. Also, we assumed $\underline{w} \le D^\delta$, so $10\underline{w}D^{z/2} \le 10D^{d\gamma/(\tau-1)-\delta} \le (D^\gamma)^{d/(\tau-1)-\delta/4}$, where the second inequality holds since $\gamma < 1$ and $D \gg_\star \delta$. Moreover, since $\underline{w} \gg_\star w_0$, we have $10\underline{w}D^{z/2} \ge w_0$. Thus all conditions in Def. 2.1 are met, and (2.2) here becomes

$$|\mathcal{Z}(x_i)| \ge D^{d\gamma(1-\delta/4)}\ell(10\underline{w}D^{z/2})(10\underline{w}D^{z/2})^{-(\tau-1)} \ge D^{d\gamma(1-\delta/4)-(\tau-1+\delta/4)(\delta+z/2)},$$

where the second inequality holds by Potter's bound since $D^\delta \ge \underline{w} \gg \delta$. The exponent of D on the right-hand side is

$$d\gamma - (\tau-1)z/2 - \delta(d\gamma/4 + z/8 + \tau - 1 + \delta/4) \ge d\gamma/2 - (\tau-1)z/2 \ge \Lambda(\eta, z)/4,$$

where we used $\delta \ll_\star \gamma$ and then the formula of $\Lambda(\eta, z)$ in (4.1). So, $|\mathcal{Z}(x_i)| \ge D^{\Lambda(\eta,z)/4}$ in (5.18), and since $D \gg_\star \delta$,

$$\mathbb{P}\big(|Z_i| < |\mathcal{Z}(x_i)|/4 \mid V, w_V\big) \le \exp(-D^{\Lambda(\eta,z)/4}/16) \le \exp(-(\theta D^{\Lambda(\eta,z)-\sqrt{\delta}})^{1/4}). \tag{5.19}$$

Returning to the event $\mathcal{A}(x_0, x_1)$ in (5.16), let $\mathcal{A}'$ be the event that $|Z_i| \ge |\mathcal{Z}(x_i)|/4$ for each $i \in \{0, 1\}$, and suppose that $\mathcal{A}'$ occurs. Observe also that the set $Z_i \subseteq \mathcal{Z}(x_i)$ were chosen independently of the edges between $\mathcal{Z}(x_0), \mathcal{Z}(x_1)$ as required in Lemma 4.2. We apply Lemma 4.2, conditioned on the values of Z_0 and Z_1, to lower-bound the right-hand side of (5.16). In the statement of Lemma 4.2, we will take $x = x_0$, $y = x_1$, $Z_x = Z_0$, $Z_y = Z_1$, $\underline{w}_{4.2} = 10\underline{w}$, and all other variables to match their current values. The event $N_{\eta,\gamma,z,10\underline{w}}(Z_0, Z_1) \neq \emptyset$ of (5.16) requires a low-cost edge between the set Z_0 and Z_1, connecting vertices with weights in I_+. Given (V, w_V), the existence of such an edge (u, v) is independent of the events $\mathcal{A}(u), \mathcal{A}(v)$ since in $\mathcal{A}(\cdot)$ the other endpoint of the edge has weight I_-, and $I_+ \cap I_- = \emptyset$. We now check the requirements of Lemma 4.2: it requires $z \in [0, d]$ that we assumed, and $2d\gamma > z(\tau-1)$. The latter holds since here we assume $\theta D^{\Lambda(\eta,z)-\sqrt{\delta}} > 1$ implying that $\Lambda(\eta, z) > 0$, so $2d\gamma > z(\tau-1)$ then follows from (4.1). Second, here we assume $\underline{w} \ge (\log\log \xi\sqrt{d})^{16d/\delta} \ge (\log\log D^\gamma)^{16d/\delta}$, and also $D \gg_\star \gamma, c_H, w_0$. So $\underline{w} \ge w_0 \vee 4(c_H+2)^d \vee 4000$ and $F_L((\underline{w}/4000)^\mu) \ge 1/2$ as required above (4.2). The requirement on D^γ here is more restrictive than in Lemma 4.2, so all requirements hold. Then, since here we have $10\underline{w}$, (4.4) turns into the following, which we then estimate by using that $\underline{w} \le D^\delta$, that $\delta \ll_\star \texttt{par}$ and that $D \gg_\star \delta$,

$$\begin{aligned}\mathbb{P}\big(N_{\eta,\gamma,z,10\underline{w}}(Z_0, Z_1) = \emptyset \mid \mathcal{A}', V, w_V\big) &\le \exp\Big(-\theta(10\underline{w})^{-2(\tau-1)}D^{\Lambda(\eta,z)-2\gamma d\delta/3}\Big)\\ &\le \exp\Big(-\theta D^{\Lambda(\eta,z)-\sqrt{\delta}}\Big) \le \exp\Big(-\big(\theta D^{\Lambda(\eta,z)-\sqrt{\delta}}\big)^{1/4}\Big),\end{aligned}$$

since we assumed $\theta D^{\Lambda(\eta,z)-\sqrt{\delta}} > 1$. Since $\mathcal{A}' = \{Z_0 \ge |\mathcal{Z}(x_0)|/4, Z_1 \ge |\mathcal{Z}(x_1)|/4\}$, combining this with (5.19) and a union bound, the result in (5.14) follows. □

We now construct a hierarchy by repeatedly applying Lemma 5.9 to find a set of valid bridges as in Lemma 5.8, using an iterative construction (Def. 3.6) to mitigate independence issues. Recall the $(\gamma, U, \overline{w}, c_H)$-hierarchy of depth R from Def. 5.5, and $\Lambda(\eta, z)$ from (4.1) and $\Phi(\eta, z)$ from (5.1).

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

Lemma 5.10 (Hierarchy with low weights $\underline{w}$). *Consider Setting 4.1, and let $y_0, y_1 \in \mathcal{N}$ with $|y_0 - y_1| = \xi$. Let $z \in [0,d]$, $\eta \geq 0$, and let $0 < \delta \ll_\star \gamma, \eta, z, \texttt{par}$ be such that $\Lambda(\eta,z) \geq 2\sqrt{\delta}$ and either $z = 0$ or $\Phi(\eta,z) \geq \sqrt{\delta}$. Let $\xi \gg_\star \gamma, \eta, z, \delta, w_0$. Let $R \geq 2$ be an integer satisfying*

$$\xi^{\gamma^{R-1}} \geq (\log\log \xi\sqrt{d})^{16d/\delta^2} \qquad \text{and} \qquad R/\theta \leq (\log\log\xi)^{1/\sqrt{\delta}}, \tag{5.20}$$

$$\text{and let} \qquad \underline{w} := \xi^{\gamma^{R-1}\delta}. \tag{5.21}$$

Let $\mathcal{X}_{low\text{-}h}(R,\eta,y_0,y_1)$ be the event that G' contains a $(\gamma, 3\underline{w}^{3\mu}\xi^\eta, \underline{w}, 2)$-hierarchy $\mathcal{H}_{low}$ of depth R with first level $\mathcal{L}_1 = \{y_0, y_1\}$, fully contained in $\mathcal{N}$, with $\mathrm{dev}_{y_0y_1}(\mathcal{H}_{low}) \leq 4\xi^\gamma$. Then

$$\mathbb{P}\big(\mathcal{X}_{low\text{-}h}(R,\eta,y_0,y_1) \mid V, w_V\big) \geq 1 - \exp\big(-(\log\log\xi)^{1/\sqrt{\delta}}\big) =: 1 - \mathrm{err}_{\xi,\delta}. \tag{5.22}$$

With the convention that $\infty \cdot 0 = 0$, the statement is also valid when $\alpha = \infty$ or $\beta = \infty$.

The lower bound on the minimal vertex weight $\underline{w}$ used in the hierarchy, and the upper bound on the number of iterations R in (5.20) jointly ensure that the thinning of edge-probabilities θ/R caused by/necessary for a multiround exposure of R rounds in Section 3 has controllable effect.

Proof. To construct a $(\gamma, 3\underline{w}^{3\mu}\xi^\eta, \underline{w}, 2)$-hierarchy in $\mathcal{N}$, we will use an iterative cost construction of $R-1$ rounds from Definition 3.6 on G'. Recall from Setting 4.1 that G' with given V, w_V is a θ-percolated CIRG. By Remark 3.3, G' is a CIRG itself (also when $\theta = \theta_n$) with distribution $\{\mathcal{G}^\theta | V, w_V\}$. In the i-th round we will construct all bridges of the i-th level of the hierarchy at once, using Lemma 5.9 2^{i-1} times to find each bridge in the level. Level 1 consists of the vertices y_0, y_1 and no edges, see also Figure 5. The first edge appears thus on level 2, hence we can start with level $i = 2$. We will use Prop. 3.9 to deal with conditioning between rounds, and union bounds to deal with conditioning within rounds. For $2 \leq i \leq R$, in the i-th round we we will set the constraints $\mathcal{F}_i$ and $\mathcal{U}_i$ so that the chosen set $\mathcal{S}_i$ in Def. 3.6(vi) consists of a $(2\xi^{\gamma^{i-1}}, 3\underline{w}^{2\mu}\xi^\eta, \underline{w})$-valid bridge for $\tilde{y}_{\sigma 0}$ and $\tilde{y}_{\sigma 1}$ for all $\sigma \in \{0,1\}^{i-2}$, where $\tilde{y}_{\sigma 0}$; and $\tilde{y}_{\sigma 1}$ are (all) endpoints of bridges from the previous levels. In other words, $\mathcal{S}_i$ will contain all the necessary bridges at the i-th level of the hierarchy for $\mathcal{X}_{\text{low-hierarchy}}(R,\eta,y_0,y_1) := \mathcal{X}_{\text{low-h}}$. Since $\mathcal{L}_1 = \{y_0, y_1\}$ contains no bridges yet, we denote the iterative construction by $(\mathcal{F}_2, \mathcal{U}_2), \dots, (\mathcal{F}_r, \mathcal{U}_r)$. Set the percolation probabilities as $\underline{\theta} := (\theta/(R-1), \dots, \theta/(R-1))$, that is, $\theta_i := \theta/(R-1)$ for $2 \leq i \leq R$ in the exposure setting of G' in Definition 3.4, and denote the outcome $\mathrm{Iter}(G_2^{\theta_2}, \dots, G_r^{\theta_r})$ by $\mathcal{S}_2, \dots, \mathcal{S}_R$.

We next inductively define the admissible edge-lists $\mathcal{F}_i$ in Def. 3.6 (ii), the cost constraints $\mathcal{U}_i$ in Def. 3.6 (iv), and vertices $\tilde{y_\sigma}$ for all $\sigma \in \{0,1\}^i$. Assume that $S_1, \dots, S_{i-1}$ is already given, that is, we constructed $(\mathcal{L}_j)_{j \leq i-1}$. For each $\sigma \in \{0,1\}^{i-2}$ consider the vertices $\tilde{y}_{\sigma 0}, \tilde{y}_{\sigma 1}$ with $\sigma 0, \sigma 1 \in \{0,1\}^{i-1}$ already found[10]. Set $D_i = \xi^{\gamma^i}$ as in Lemma 5.8, and write $\mathcal{P}(\sigma)$ for the set of all possible paths (i.e., sequence of vertices) contained in $\mathcal{N}$ between all $y \in \mathcal{N} \cap (B_{2D_{i-1}}(\tilde{y}_{\sigma 0}) \times [\underline{w}, 4\underline{w}])$ and all $y' \in \mathcal{N} \cap (B_{2D_{i-1}}(\tilde{y}_{\sigma 1}) \times [\underline{w}, 4\underline{w}])$ (so that if $\widetilde{P}_\sigma \in \mathcal{P}(\sigma)$, then $\widetilde{P}_\sigma$ satisfies both (5.10), (5.11) in Def. 5.7). Since V, w_V is given, and also $S_1, \dots, S_{i-1}$ is already determined, define now an edge-list $\underline{t}$ to be *level-i admissible* if it contains exactly one such potential path from $\mathcal{P}(\sigma)$ for each $\sigma \in \{0,1\}^{i-2}$, and let $\mathcal{F}_i(G_2^{\theta_2}, \dots, G_{i-1}^{\theta_{i-1}})$ be the list of all level-i admissible edge-lists, with an arbitrary ordering.

For each such admissible list, $\mathcal{U}_i(G_2^{\theta_2}, \dots, G_{i-1}^{\theta_{i-1}})$ describes the cost-constraint; let this be the constraint that the edges in the list have total marginal cost at most $3\underline{w}^{3\mu}\xi^\eta$ (where marginal cost is defined in (3.1)). Once we reveal the edges of $G_i^{\theta_i}$, recall that the result of round i, $\mathrm{Iter}_i(G_2^{\theta_2}, \dots, G_{i-1}^{\theta_{i-1}}, G_i^{\theta_i}) =: \mathcal{S}_i^{\exp}$ is then set by Def. 3.6(vi), and that if the construction succeeds, then $\mathcal{S}_i^{\exp}$ is the first element in $\mathcal{F}_i(G_2^{\theta_2}, \dots, G_{i-1}^{\theta_{i-1}})$ that satisfies the corresponding cost constraint in $\mathcal{U}_i(G_2^{\theta_2}, \dots, G_{i-1}^{\theta_{i-1}})$. Given $\mathcal{S}_i^{\exp}$, we define $\tilde{y}_{\sigma 00} := \tilde{y}_{\sigma 0}$, $\tilde{y}_{\sigma 11} := \tilde{y}_{\sigma 1}$, and $\tilde{y}_{\sigma 01}$ and $\tilde{y}_{\sigma 10}$ to be the endpoints of the bridge $\widetilde{P}_\sigma$ present in the chosen edge-list $\mathcal{S}_i^{\exp}$, or `None` if $\mathcal{S}_i^{\exp} = $ `None`. This gives the iterative cost construction

[10] Thus, for the initial $i = 2$ here $\sigma = \emptyset$ so we look at the vertices y_0, y_1.

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

$\mathrm{Iter}^{\mathrm{exp}} = ((\mathcal{F}_i, \mathcal{U}_i) : i \in \{2, \ldots, R\})$ applied on $G_2^{\theta_2}, \ldots, G_R^{\theta_R}$, denoted by $\mathrm{Iter}^{\mathrm{exp}}_{\{\mathcal{G}^\theta | V, w_V\}, \underline{\theta}}$ in Def. 3.8. Note that the criteria above for $\mathcal{P}(\sigma), \mathcal{F}_i, \mathcal{U}_i$ exactly matches Lemma 5.8 with $c_H = 2$ and marginal cost of each bridge $\widetilde{P}_\sigma$ at most $U = 3\underline{w}^{3\mu}\xi^\eta$, implying (5.12), that is, each chosen bridge $\widetilde{P}_\sigma \in \mathcal{S}_i^{\mathrm{exp}}$ is $(2\xi^{\gamma^{i-1}}, 3\underline{w}^{3\mu}\xi^\eta, \underline{w})$-valid for $y_{\sigma 0}, y_{\sigma 1} \in \{0,1\}^{i-1}$ with respect to the chosen edges in earlier rounds. Since all vertices in $\{P_\sigma\}_\sigma$ are contained in a $2(\xi^\gamma + \xi^{\gamma^2} + \cdots + \xi^{\gamma^{R-1}}) \le 4\xi^\gamma$ ball around y_0 and y_1, respectively, the deviation requirement is also satisfied, and so by Lemma 5.8, if $\mathrm{Iter}^{\mathrm{exp}}_{\{\mathcal{G}^\theta | V, w_V\}, \underline{\theta}}$ succeeds then $\{\tilde{y}_\sigma\}_{\sigma \in \{0,1\}^R}$ is a $(\gamma, 3\underline{w}^{3\mu}\xi^\eta, \underline{w}, 2)$-hierarchy as needed in $\mathcal{X}_{\text{low-h}}$.

Following Prop. 3.9, let $r = R - 1$ and $\theta_i \equiv 1/(R-1)$ there, and let $H_2, \ldots, H_R$ be independent $1/(R-1)$-percolations of $\{\mathcal{G}^\theta \mid V, w_W\}$, that is, with distribution $\{\mathcal{G}^{\theta/(R-1)} \mid V, w_W\}$ from Def. 3.2. Recall from Def. 3.8 the definitions of $\mathrm{Iter}^{\mathrm{ind}}_{\{\mathcal{G}^\theta | V, w_V\}, \underline{\theta}}$ and $\mathcal{A}^{\mathrm{ind}}(S_1, \ldots, S_{i-1})$. Then applying Prop. 3.9 (with an index shift, since now we start at $i = 2$), (3.5) turns into

$$\begin{aligned} \mathbb{P}(\mathcal{X}_{\text{low-h}} \mid V, w_V) &\ge \mathbb{P}(\mathrm{Iter}^{\mathrm{exp}}_{\{\mathcal{G}^\theta | V, w_V\}, \underline{\theta}} \text{ succeeds} \mid V, w_V) \\ &\ge \min_{S_2, \ldots, S_R \neq \texttt{None}} \prod_{i=2}^{R} \mathbb{P}(\mathcal{S}_i^{\mathrm{ind}} \neq \texttt{None} \mid \mathcal{A}^{\mathrm{ind}}(S_2, \ldots, S_{i-1})) \\ &\ge 1 - \sum_{i=2}^{R} \max_{S_2, \ldots, S_{i-1} \neq \texttt{None}} \mathbb{P}(\mathcal{S}_i^{\mathrm{ind}} = \texttt{None} \mid \mathcal{A}^{\mathrm{ind}}(S_1, \ldots, S_{i-1})), \end{aligned} \tag{5.23}$$

by a union bound over all rounds.

We now break the right-hand side of (5.23) down into bridge existence events under simpler conditioning. Recall Definition 5.7, in particular the notation (D, U, w)-valid bridges with respect to (already revealed edges) $S := \cup_{j \le i-1} S_j$. For each $2 \le i \le R$ and $\sigma 0, \sigma 1 \in \{0,1\}^{i-1}$, let

$$\begin{aligned} \mathcal{A}_i(\tilde{y}_{\sigma 0}, \tilde{y}_{\sigma 1}) := \{\exists \widetilde{P}_\sigma \in S_i : (2\xi^{\gamma^{i-1}}, 3\underline{w}^{2\mu}\xi^{\gamma^{i-2}\eta}, \underline{w})&\text{-valid for} \\ \tilde{y}_{\sigma 0}, \tilde{y}_{\sigma 1} &\text{ with respect to } S = \emptyset\}. \end{aligned} \tag{5.24}$$

This is a stronger condition than what is required for a $(\gamma, 3\underline{w}^{3\mu}\xi^\eta, \underline{w}, 2)$-hierarchy to exist in Lemma 5.8, since $\gamma^{i-2}\eta \le \eta$ and validity with respect to $\emptyset$ implies validity with respect to any set of edges. Conditioned on $\mathcal{A}^{\mathrm{ind}}(S_2, \ldots, S_{i-1})$ so that none of the $(S_j)_{j \le i-1}$ equals None, the event $\mathcal{S}_i^{\mathrm{ind}} = \texttt{None}$ occurs only if for some pair $\sigma 0, \sigma 1 \in \{0,1\}^{i-1}$, the complement of the event $\mathcal{A}_i(\tilde{y}_{\sigma 0}, \tilde{y}_{\sigma 1})$ occurs; hence by a union bound, (5.23) implies

$$\mathbb{P}(\mathcal{X}_{\text{low-h}} | V, w_V) \ge 1 - \sum_{i=2}^{R} 2^{i-2} \max_{\substack{\sigma \in \{0,1\}^{i-2} \\ S_2, \ldots, S_{i-1} \neq \texttt{None}}} \mathbb{P}(\mathcal{A}_i(\tilde{y}_{\sigma 0}, \tilde{y}_{\sigma 1})^{\complement} | V, w_V, \mathcal{A}^{\mathrm{ind}}(S_1, \ldots, S_{i-1})). \tag{5.25}$$

Recall that given (V, w_V), the graphs $H_2, \ldots H_{R-1}$ are i.i.d. $\{\mathcal{G}^{\theta/(R-1)} \mid V, w_W\}$. So, the events in $\mathcal{A}^{\mathrm{ind}}(S_2, \ldots, S_{i-1})$ are contained in the σ-algebra generated by $H_2, \ldots, H_{i-1}$, that is, independent of H_i and thus of the complement of $\mathcal{A}_i(\tilde{y}_{\sigma 0}, \tilde{y}_{\sigma 1})$. Hence (5.25) simplifies to

$$\mathbb{P}(\mathcal{X}_{\text{low-h}} \mid V, w_V) \ge 1 - \sum_{i=2}^{R} 2^{i-2} \max_{\tilde{y}_{\sigma 0}, \tilde{y}_{\sigma 1} \neq \texttt{None}} \mathbb{P}(\mathcal{A}_i(\tilde{y}_{\sigma 0}, \tilde{y}_{\sigma 1})^{\complement} \mid V, w_V), \tag{5.26}$$

where the maximum is taken over all possible values of $(\tilde{y}_{\sigma 0}, \tilde{y}_{\sigma 1})$ occurring in non-None S_{i-1}. Finally, we will upper-bound the probabilities on the right-hand side of (5.26) using Lemma 5.9. Let $2 \le i \le R$, let $\sigma \in \{0,1\}^{i-2}$, and let $\tilde{y}_{\sigma 0}, \tilde{y}_{\sigma 1}$ be a possible non-None realisation of the embedding. Recall $D_i = \xi^{\gamma^i}$. Then the event (5.24) requires a $(D_{i-2}^\gamma, 3\underline{w}^{3\mu}D_{i-2}^\eta, \underline{w})$-valid bridge P_σ, which formally matches Lemma 5.9 with $D := D_{i-2}, \tilde{y}_{\sigma 0} := x_0, \tilde{y}_{\sigma 1} := x_1$ there and the graph $H_i \sim \{\mathcal{G}^{\theta/(R-1)} \mid V, w_V\}$ in place of G' there, that is, with $\theta_{5.9} := \theta/(R-1)$.

We check the conditions of Lemma 5.9 in order of their appearance. $z \in [0,d], \eta, \delta > 0$ and $\delta \ll_\star \gamma, \eta, z$ is assumed both here and there. The assumption $\xi \gg_\star \gamma, \eta, z, \delta, w_0$ here implies $D_{i-2} \gg_\star \gamma, \eta, z, \delta, w_0$ since by (5.20) $D_{i-2}^\gamma \ge \xi^{\gamma^{R-1}} \ge (\log\log \xi\sqrt{d})^{16d/\delta^2}$. The latter also implies the requirement on D^γ in Lemma 5.9. Similarly, the upper bound requirement holds since $D_{i-2}^\gamma \le \xi \le \xi\sqrt{d}$. We now check whether $\theta D^{\Lambda(\eta,z)-\sqrt{\delta}} > 1$ holds in Lemma 5.9 for our choices. Since we assumed here $\Lambda(\eta,z) \ge 2\sqrt{\delta}$, and also (5.20), we estimate

$$\tfrac{\theta}{R-1} D_{i-2}^{\Lambda(\eta,z)-\sqrt{\delta}} \ge (\log\log\xi)^{-1/\sqrt{\delta}} \cdot D_{i-2}^{\sqrt{\delta}} \ge (\log\log\xi)^{15/\sqrt{\delta}} > 1. \tag{5.27}$$

Next we need to check whether $x_0 = \tilde{y}_{\sigma 0}, x_1 = \tilde{y}_{\sigma 1}$ satisfies $|x_0 - x_1| \le c_H D_{i-2}$. This is true since $\tilde{y}_{\sigma 0}, \tilde{y}_{\sigma 1}$ are possible non-None values coming from chosen tuples $S_1, \dots, S_{i-1}$; and by construction of $\mathcal{P}(\sigma)$ above, we required that $|\tilde{y}_{\sigma' 00} - \tilde{y}_{\sigma' 01}|, |\tilde{y}_{\sigma' 10} - \tilde{y}_{\sigma' 11}| \le 2D_{i-1}$ for all $\sigma' \in \{0,1\}^{i-2}$, which, when shifting indices yields exactly that $|\tilde{y}_{\sigma 0} - \tilde{y}_{\sigma 1}| \le 2D_{i-2}$ for all $\sigma \in \{0,1\}^{i-2}$. Next we check the criterion on $\underline{w}$ in Lemma 5.9. Here, $\underline{w}$ is defined in (5.21), hence, using (5.20), $\underline{w} = \xi^{\gamma^{R-1}\delta} \ge (\log\log\xi\sqrt{d})^{16d/\delta}$ as required. This also implies $\underline{w} \gg_\star \delta, w_0$ since $\xi \gg_\star \delta, w_0$. Moreover, $\underline{w} = \xi^{\gamma^{R-1}\delta} \le D_{i-2}^\delta = \xi^{\gamma^{i-2}\delta}$ holds since $i - 2 \le R - 2$ and $\gamma < 1$. Next, we check (5.13), which can be lower bounded by omitting the prefactor $\underline{w}^{\mu\beta}$ in the last factor (the minimum):

$$p(D_{i-2}, \underline{w}, \tfrac{\theta}{R-1}, \eta, z) \ge \tfrac{\theta}{R-1} \underline{w}^{-(\tau-1)} \Big(D_{i-2}^{d\gamma} \wedge \underline{w}^2 D_{i-2}^{z/2}\Big)^{1-\delta} D_{i-2}^{[0\wedge\beta(\eta-\mu z/2)]}. \tag{5.28}$$

We distinguish cases with respect to z to handle the minimum in the middle of the right-hand side. If $z = 0$, then $\underline{w}^2 D_i^{z/2} = \underline{w}^2 = \xi^{2\gamma^{R-1}\delta} \le \xi^{\gamma^{i-1}d} = D_{i-2}^{\gamma d}$, where the inequality holds because $i \le R$ and $\delta \ll_\star \gamma$. Moreover $0 \le \eta - \mu z/2$ in that case, so when $z = 0$, equation (5.28) becomes

$$p(D_{i-2}, \underline{w}, \tfrac{\theta}{R-1}, \eta, z) \ge \tfrac{\theta}{R-1} \underline{w}^{2(1-\delta)-(\tau-1)} = \tfrac{\theta}{R-1} \underline{w}^{3-\tau-2\delta} \ge \tfrac{\theta}{R-1} \underline{w}^{\sqrt{\delta}}, \tag{5.29}$$

where the last inequality holds because $\delta \ll_\star \texttt{par}$. If, however, $z \ne 0$, then we assumed that $\Phi(\eta,z) \ge \sqrt{\delta}$ in (5.1). Using again $\underline{w} \ge 1$, we lower bound (5.28) in this case

$$p(D_{i-2}, \underline{w}, \tfrac{\theta}{(R-1)}, \eta, z) \ge \tfrac{\theta}{R-1} \underline{w}^{-(\tau-1)} D_{i-2}^{(1-\delta)[d\gamma\wedge z/2]+[0\wedge\beta(\eta-\mu z/2)]} \ge \tfrac{\theta}{R-1} D_{i-2}^{\Phi(\eta,z)-\delta(\tau-1+d)},$$

where we used that $\underline{w} \le D_{i-2}^\delta$ implies $\underline{w}^{-(\tau-1)} \ge D_{i-2}^{-\delta(\tau-1)}$ and $d\gamma \wedge z/2 \le d$ (since $\gamma < 1$) to obtain the last inequality. Since δ is small, $\Phi(\eta,z) \ge \sqrt{\delta}$, and $\underline{w} \le D_{i-2}^\delta$, this implies

$$p(D_{i-2}, \underline{w}, \tfrac{\theta}{R-1}, \eta, z) \ge \tfrac{\theta}{R-1} D_{i-2}^\delta \ge \tfrac{\theta}{R-1} \underline{w} \ge \tfrac{\theta}{R-1} \underline{w}^{\sqrt{\delta}}, \tag{5.30}$$

the same lower bound as in (5.29) for $z = 0$. Thus for all $z \in [0,d]$, using (5.20) for a lower bound on θ/R and (5.21),

$$p(D_{i-2}, \underline{w}, \tfrac{\theta}{R-1}, \eta, z) \ge \tfrac{\theta}{R-1} \underline{w}^{\sqrt{\delta}} \ge \tfrac{\theta}{R-1} (\log\log\xi)^{16/\sqrt{\delta}} \ge (\log\log\xi)^{15/\sqrt{\delta}} \ge 20^{\tau+\mu\beta},$$

where the last inequality holds because $\xi \gg_\star \delta, \texttt{par}$. With this, all conditions of Lemma 5.9 are satisfied, so combining (5.14) with (5.26) and then using the lower bound in (5.27) yields

$$\begin{aligned} \mathbb{P}(\mathcal{X}_{\text{low-h}} \mid V, w_V) &\ge 1 - 3\sum_{i=2}^{R} 2^{i-2} \exp\Big(-\Big[\tfrac{\theta}{R-1} D_{i-2}^{\Lambda(\eta,z)-\sqrt{\delta}}\Big]^{1/4}\Big) \\ &\ge 1 - 3\sum_{i=2}^{R} 2^{i-2} \cdot \exp\big(-(\log\log\xi)^{3/\sqrt{\delta}}\big) \ge 1 - 2^{R+1}\exp\big(-(\log\log\xi)^{3/\sqrt{\delta}}\big). \end{aligned} \tag{5.31}$$

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

Finally, in (5.20) the estimate $R \le (\log\log\xi)^{1/\sqrt{\delta}}$ can be used to upper bound 2^{R+1}, yielding the required inequality in (5.22). □

Lemma 5.10 constructed a hierarchy with bridge endpoints $\tilde{y}_\sigma$ of weight roughly $\underline{w} = \xi^{\gamma^{R-1}\delta}$. This weight is too low to connect the final gaps (siblings) in the hierarchy via short paths. The next lemma extends this hierarchy to a new one with endpoints y_σ of weight roughly $\overline{w} := \xi^{d\gamma^{R-1}/2}$ using weight-increasing paths. At these higher weights, connecting the gaps is possible. The proof follows a very similar structure as for Lemma 5.10, with just two rounds of exposure. Recall the $(\gamma, U, \overline{w}, c_H)$-hierarchy of depth R from Def. 5.5, $\Lambda(\eta, z)$, $\Phi(\eta, z)$ from (4.1) and (5.1), respectively.

Lemma 5.11 (Hierarchy with high weights $\overline{w}$). *Consider Setting 4.1, and let $y_0, y_1 \in \mathcal{N}$ with $|y_0 - y_1| = \xi$. Let $z \in [0, d], \eta \ge 0$, and let $0 < \delta \ll_\star \gamma, \eta, z,$ `par` be such that $\Lambda(\eta, z) \ge 2\sqrt{\delta}$ and either $z = 0$ or $\Phi(\eta, z) \ge \sqrt{\delta}$. Let $\xi \gg_\star \gamma, \eta, z, \theta, \delta, w_0$. Let $R \ge 2$ be an integer satisfying $\xi^{\gamma^{R-1}} \ge (\log\log\xi\sqrt{d})^{16d/\delta^2}$ and $R \le (\log\log\xi)^2$, and set*

$$\overline{w} := \xi^{\gamma^{R-1}d/2}, \qquad c_H := 8\left(1 + \left\lceil \frac{\log(d/\delta)}{\log(1/(\tau - 2 + 2d\tau\delta))} \right\rceil\right). \tag{5.32}$$

Let $\mathcal{X}_{high\text{-}h}(R, \eta, y_0, y_1)$ be the event that G' contains a $(\gamma, c_H\overline{w}^{4\mu}\xi^\eta, \overline{w}, c_H)$-hierarchy $\mathcal{H}_{high}$ of depth R with first level $\mathcal{L}_1 = \{y_0, y_1\}$, fully contained in $\mathcal{N}$, and $\mathrm{dev}_{y_0y_1}(\mathcal{H}_{high}) \le 2c_H\xi^\gamma$. Then

$$\mathbb{P}\big(\mathcal{X}_{high\text{-}h}(R, \eta, y_0, y_1) \mid V, w_V\big) \ge 1 - \exp\left(-(\log\log\xi)^{13}\right); \tag{5.33}$$

under the convention that $\infty \cdot 0 = 0$, the statement is also valid when $\alpha = \infty$ or $\beta = \infty$.

Proof. As in Lemma 5.10, let $\underline{w} := \xi^{\gamma^{R-1}\delta}$. To construct a $(\gamma, c_H\overline{w}^{4\mu}\xi^\eta, \overline{w}, c_H)$-hierarchy in $\mathcal{N}$, we use two rounds of exposure. First we set $\theta_1 = \theta_2 = 1/2$ and construct the exposure setting of G', that is, $(G_1^{\theta_1}, G_2^{\theta_2})$. In the first round we apply Lemma 5.10 to get $\mathcal{H}_{\mathrm{low}} = \{\tilde{y}_\sigma\}$, a $(\gamma, 3\underline{w}^{3\mu}\xi^\eta, \underline{w}, 2)$-hierarchy with failure probability $\mathrm{err}_{\xi,\delta}$ in (5.22). In the second round, we use weight-increasing paths from Lemma 4.4 to connect each $\tilde{y}_\sigma \in \mathcal{H}_{\mathrm{low}}$ to a vertex y_σ of weight in $[\overline{w}, 4\overline{w}]$, transforming $\mathcal{H}_{\mathrm{low}}$ into $\mathcal{H}_{\mathrm{high}} = \{y_\sigma\}$, a $(\gamma, c_H\overline{w}^{4\mu}\xi^\eta, \overline{w}, c_H)$-hierarchy.

We now define an iterative cost construction on $G' \sim \{\mathcal{G}^\theta | V, w_V\}$. In round 1, $\mathcal{F}_1$ is the list of admissible lists of vertex pairs edges (with an arbitrary ordering): we set now a list of vertex-pairs admissible if it could form the bridges $(P_\sigma)_{\sigma\in\{0,1\}^R}$ of a $(\gamma, 3\underline{w}^{3\mu}\xi^\eta, \underline{w}, 2)$-hierarchy $\{\tilde{y}_\sigma\}_{\sigma\in\{0,1\}^R}$ fully contained in $\mathcal{N}$. For any given list in $\mathcal{F}_1$, let the corresponding event in $\mathcal{U}_1$ be specified by the set of all possible edge costs such that all P_σ satisfy *(H3)* of Definition 5.5 with $U = 3\underline{w}^{3\mu}\xi^\eta$, so that $\mathcal{H}_{\mathrm{low}}$ is indeed a valid $(\gamma, 3\underline{w}^{3\mu}\xi^\eta, \underline{w}, 2)$-hierarchy. The round-1 marginal costs in (3.1) are equal to the edge costs in $G_1^{\theta_1}$.

We move now to round 2. For each $\sigma \in \{0,1\}^R$, let $\mathcal{P}(\sigma)$ be the set of all paths $\pi_{\tilde{y}_\sigma, y_\sigma}$ in $\mathcal{N}$ connecting $\tilde{y}_\sigma$ to any vertex $y_\sigma \in \mathcal{N} \cap (B_{(c_H-2)\xi^{\gamma^{R-1}}/2}(\tilde{y}_\sigma) \times [\overline{w}, 4\overline{w}])$. Given $(V, w_V, \mathcal{S}_1)$, call a list $\underline{t}$ of vertex-pairs *admissible* in round 2 if it contains exactly one such potential path from $\mathcal{P}(\sigma)$ for each $\sigma \in \{0,1\}^R$, and let $\mathcal{F}_2(G_1^{\theta_1})$ be the collection of all admissible tuples, with an arbitrary ordering. For any given $\underline{t} \in \mathcal{F}_2(G_1^{\theta_1})$, let the corresponding event in $\mathcal{U}_2(G_1^{\theta_1})$ be specified by the set of all possible round-2 marginal costs for the edges in $\underline{t}$ which sum to at most $(c_H - 3)\overline{w}^{4\mu}\xi^\eta/2$ in (3.1). This defines an iterative cost construction $\mathrm{Iter} = ((\mathcal{F}_1, \mathcal{U}_1), (\mathcal{F}_2, \mathcal{U}_2))$ applied on $G_1^{\theta_1}, G_2^{\theta_2}$, that we denote by $\mathrm{Iter}^{\mathrm{exp}}_{\{\mathcal{G}^\theta|V,w_V\},\underline{\theta}}$. Recall from Def. 3.6(vi) that for $i \in \{1,2\}$, $\mathcal{S}_i^{\mathrm{exp}}$ is either `None` or lies in $\mathcal{F}_i^{\mathrm{exp}}$ with round-i marginal costs satisfying $\mathcal{U}_i^{\mathrm{exp}}$.

If $\mathrm{Iter}^{\mathrm{exp}}_{\mathcal{G}^\theta,\underline{\theta}}$ succeeds, then $\{y_\sigma\}_{\sigma\in\{0,1\}^R}$ is a $(\gamma, c_H\overline{w}^{4\mu}\xi^\eta, \overline{w}, c_H)$-hierarchy. Indeed, condition *(H1)* of Def. 5.5 is satisfied by construction. By the triangle inequality, *(H2)* is satisfied since for all $\sigma \in \{0,1\}^i$, $y_{\sigma 1} \in B_{(c_H-2)\xi^{\gamma^i}/2}(\tilde{y}_{\sigma 1})$ and $y_{\sigma 0} \in B_{(c_H-2)\xi^{\gamma^i}/2}(\tilde{y}_{\sigma 0})$ by construction, and $|\tilde{y}_{\sigma 1} - \tilde{y}_{\sigma 0}| \le 2\xi^{\gamma^i}$ by *(H2)* since $\tilde{y}_\sigma$ forms a $(\gamma, 3\underline{w}^{3\mu}\xi^\eta, \underline{w}, 2)$-hierarchy. This also implies that $\mathrm{dev}_{y_0y_1}(\mathcal{H}_{\mathrm{high}}) \le c_H(\xi^\gamma +$

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

$\xi^{\gamma^2} + \cdots + \xi^{\gamma^{R-1}}) \le 2c_H \xi^{\gamma}$, since $\xi \gg_\star \gamma$, as required. Finally, $\widetilde{P}_\sigma$ is a path between $\tilde{y}_{\sigma 01}$ and $\tilde{y}_{\sigma 10}$, so let P_σ be the concatenated path $\pi_{y_{\sigma 01}, \tilde{y}_{\sigma 01}} \widetilde{P}_\sigma \pi_{\tilde{y}_{\sigma 10}, y_{\sigma 10}}$. Then the total cost of P_σ is

$$\begin{aligned}\mathcal{C}(P_\sigma) &\le \mathcal{C}(\widetilde{P}_\sigma) + \mathrm{mcost}_2(\pi_{y_{\sigma 01}, \tilde{y}_{\sigma 01}}) + \mathrm{mcost}_2(\pi_{\tilde{y}_{\sigma 10}, y_{\sigma 10}}) \\ &\le 3\underline{w}^{3\mu}\xi^\eta + 2(c_H - 3)\overline{w}^{4\mu}\xi^\eta/2 \le c_H \overline{w}^{4\mu}\xi^\eta,\end{aligned}$$

since $\underline{w} = \overline{w}^{2\delta/d}$, see (5.21) vs (5.32).

As in Def. 3.8 and Prop. 3.9, we now lower-bound the probability that $\mathrm{Iter}^{\exp}_{\{\mathcal{G}^\theta | V, w_V\}, \underline{\theta}}$ succeeds by coupling to two independent percolations H_1 and H_2. With $\underline{\theta} := (\theta/2, \theta/2)$, recall the definition $\mathrm{Iter}^{\mathrm{ind}}_{\{\mathcal{G}^\theta | V, w_V\}, \underline{\theta}}$. As in Def. 3.8, let $\mathcal{A}^{\mathrm{ind}}(S_1)$ be the event that the first round returns the edge set $\mathcal{S}_1^{\mathrm{ind}} = S_1$. Then Proposition 3.9 followed by a union bound gives

$$\begin{aligned}\mathbb{P}\big(\mathcal{X}_{\text{high-h}}(R, \eta, y_0, y_1) \mid V, w_V\big) &\ge \mathbb{P}\big(\mathrm{Iter}^{\exp}_{\{\mathcal{G}^\theta | V, w_V\}, \underline{\theta}} \text{ succeeds} \mid V, w_V\big) \\ &\ge \mathbb{P}\big(\mathcal{S}_1^{\mathrm{ind}} \ne \mathtt{None} \mid V, w_V\big) \cdot \min_{S_1 \ne \mathtt{None}} \mathbb{P}\big(\mathcal{S}_2^{\mathrm{ind}} \ne \mathtt{None} \mid V, w_V, \mathcal{A}^{\mathrm{ind}}(S_1)\big) \\ &\ge 1 - \mathbb{P}\big(\mathcal{S}_1^{\mathrm{ind}} = \mathtt{None} \mid V, w_V\big) - \max_{S_1 \ne \mathtt{None}} \mathbb{P}\big(\mathcal{S}_2^{\mathrm{ind}} = \mathtt{None} \mid V, w_V, \mathcal{A}^{\mathrm{ind}}(S_1)\big).\end{aligned} \tag{5.34}$$

The event $\mathcal{S}_1^{\mathrm{ind}} \ne \mathtt{None}$ occurs precisely when the graph H_1 contains a $(\gamma, 3\underline{w}^{3\mu}\xi^\eta, \underline{w}, 2)$-hierarchy $\mathcal{H}_{\mathrm{low}} := \{\tilde{y}_\sigma\}$ of depth R fully contained in $\mathcal{N}$ with first level $\mathcal{L}_1 = \{y_0, y_1\}$. Since $H_1 \sim \{\mathcal{G}^{\theta/2} \mid V, w_V\}$ is a CIRG, and since the conditions here are stronger than those in Lemma 5.10, all requirements of Lemma 5.10 hold with θ replaced by $\theta/2$, so the first error term in (5.34) is at most $\mathrm{err}_{\xi,\delta}$ in (5.22).

It remains to upper-bound the second error term in (5.34). We use weight-increasing paths to connect the endpoints of $\mathcal{H}_{\mathrm{low}}$ to vertices of weight $[\overline{w}, 4\overline{w}]$ nearby. Let

$$q^\star := \left\lceil \frac{\log(d/\delta)}{\log(1/(\tau - 2 + 2d\tau\delta))} \right\rceil, \tag{5.35}$$

and for each $v \in \mathcal{N}$ let $\mathcal{A}_{\mathrm{path}}(v)$ be the event that H_2 contains a path $\pi_{v,v'}$ from v to some vertex $v' \in \mathcal{N} \cap (B_{4q^\star \xi^{\gamma^{R-1}}}(v) \times [\overline{w}, 4\overline{w}])$ with cost $\mathcal{C}_2(\pi_{v,v'}) \le q^\star \overline{w}^{4\mu}\xi^\eta$. The value c_H in (5.32) is chosen so that $4q^\star \le (c_H - 2)/2$ and $q^\star \le (c_H - 3)/2$ both hold; thus conditioned on $\mathcal{A}^{\mathrm{ind}}(S_1)$, the event $\mathcal{S}_2^{\mathrm{ind}} = \mathtt{None}$ occurs only if for some $\sigma \in \{0,1\}^R$ the event $\mathcal{A}_{\mathrm{path}}(\tilde{y}_\sigma)^\complement$ occurs. There are 2^R strings $\sigma \in \{0,1\}^R$, and all the events in $\mathcal{A}_{I_H}(V, w_V, S_1)$ are contained in the σ-algebra generated by H_1, which is independent of H_2 given V, w_V. So, by a union bound,

$$\begin{aligned}\max_{S_1 \ne \mathtt{None}} \mathbb{P}\big(\mathcal{S}_2^{\mathrm{ind}} = \mathtt{None} \mid \mathcal{A}^{\mathrm{ind}}(S_1)\big) &\le 2^R \cdot \max_{\substack{\sigma \in \{0,1\}^R \\ S_1 \ne \mathtt{None}}} \mathbb{P}\big(\mathcal{A}_{\mathrm{path}}(\tilde{y}_\sigma)^\complement \mid \mathcal{A}^{\mathrm{ind}}(S_1)\big) \\ &\le 2^R \cdot \max_{\tilde{y}_\sigma \ne \mathtt{None}} \mathbb{P}\big(\mathcal{A}_{\mathrm{path}}(\tilde{y}_\sigma)^\complement \mid V, w_V\big),\end{aligned} \tag{5.36}$$

where the maximum is taken over all possible values of $\tilde{y}_\sigma$ in non-None S_1. To bound (5.36), we apply Lemma 4.4 with $G' = H_2$, θ replaced by $\theta/2$, $K = 2\overline{w}$, $M = 2\underline{w}$, $D = 4\xi^{\gamma^{R-1}}$, $U = \overline{w}^{4\mu}\xi^\eta$, $y_0 = y_\sigma$, and all other variables taking their present values. Using $\underline{w}, \overline{w}$ from (5.21) and (5.32), we compute

$$\frac{\log K}{\log M} = \frac{\log(2\overline{w})}{\log(2\underline{w})} = \frac{\log 2 + \frac{1}{2}\gamma^{R-1} d \log \xi}{\log 2 + \gamma^{R-1}\delta \log \xi} \le 1 + \frac{d}{2\delta} \le \frac{d}{\delta};$$

and therefore q from (4.18) with these choices satisfies

$$q = \left\lceil \frac{\log\big(\log K/\log M\big)}{\log(1/(\tau-2+2d\tau\delta))} \right\rceil \le \left\lceil \frac{\log(d/\delta)}{\log(1/(\tau-2+2d\tau\delta))} \right\rceil = q^\star.$$

Hence the event $\mathcal{A}_{\pi(\tilde{y}_\sigma)}$ in Lemma 4.4 is contained in $\mathcal{A}_{\mathrm{path}}(\tilde{y}_\sigma)$. We now verify that the requirements of Lemma 4.4 hold in order of their appearance there. Whenever $S_1 \ne \texttt{None}$, $\tilde{y}_\sigma$ lies in $\mathcal{N}$ with weight in $[\underline{w}, 4\underline{w}] = [M/2, 2M]$ by construction, where $M = 2\underline{w} > 1$. Similarly, $K = 2\overline{w} > 1$ and $D = 4\xi^{\gamma^{R-1}} > 1$ by our choices. We check the requirement $U \ge K^{2\mu}$. By definition of $\overline{w}$ in (5.32) and the choices $U = \overline{w}^{4\mu}\xi^\eta$ and $K = 2\underline{w}$, we compute

$$UK^{-2\mu} = \overline{w}^{4\mu}\xi^\eta(2\overline{w})^{-2\mu} = 2^{-2\mu}\overline{w}^{2\mu}\xi^\eta,$$

which is larger than 1 (even if $\eta = 0$) since $\mu > 1$ and $\overline{w} \ge (\log\log\xi\sqrt{d})^{8d^2/\delta^2}$ and $\xi \gg_\star \delta$. Next, since $\delta \ll_\star \texttt{par}$ by hypothesis, $\overline{w} = \xi^{\gamma^{R-1}d/2} > \xi^{\gamma^{R-1}\delta} = \underline{w}$ and so $K \ge M$. Moreover, $K = 2\xi^{\gamma^{R-1}d/2} \le 4^{d/2}\xi^{\gamma^{R-1}d/2} = D^{d/2}$ also holds. Since $M = 2\underline{w} = 2\xi^{\gamma^{R-1}\delta} \ge 2(\log\log\xi\sqrt{d})^{16d/\delta}$, $\xi \gg_\star \theta, \delta, w_0$, and $M \le K \le D^{d/2}$, we have $K, M, D \gg_\star \theta, \delta, w_0$. Clearly $D = 4\xi^{\gamma^{R-1}} < \xi \le \xi\sqrt{d}$ since $\gamma < 1$ and ξ is large. Next, we check $(M/2)^{2/d} = (\xi^{\gamma^{R-1}\delta})^{2/d} > \xi^{\gamma^{R-1}\delta/d} \ge (\log\log\xi\sqrt{d})^{16/\delta}$ as required. Finally, if $\beta = \infty$ then we also need that $U(KM)^{-\mu} = \overline{w}^{4\mu}\xi^\eta(4\overline{w}\underline{w})^{-\mu} \ge 4^{-\mu}\overline{w}^{2\mu}\xi^\eta$ is sufficiently large. This holds even when $\eta = 0$ since $\overline{w} \gg_\star \texttt{par}$. Hence, all requirements of Lemma 4.4 are satisfied, and since θ changes to $\theta/2$ and $M = 2\underline{w} = 2\xi^{\gamma^{R-1}\delta}$ in (4.19), (5.36) can be bounded as

$$\begin{aligned}\max_{S_1 \ne \texttt{None}} \mathbb{P}\big(\mathcal{S}_2^{\mathrm{ind}} = \texttt{None} \mid \mathcal{A}^{\mathrm{ind}}(S_1)\big) &\le 2^R \exp\Big(-(\theta/2)2^\delta\xi^{\gamma^{R-1}\delta^2}\Big) \\ &\le 2^R\exp\Big(-(\log\log\xi)^{15}\Big) \le \exp\Big(-(\log\log\xi)^{14}\Big),\end{aligned} \tag{5.37}$$

where we obtained the second row from the hypotheses $\xi^{\gamma^{R-1}\delta^2} \ge (\log\log\xi)^{16}$ and $\xi \gg_\star \theta$, and then from $2^R \le e^R$ and $R \le (\log\log\xi)^2$. Combining (5.37) with (5.34) and recalling that the first error term there is at most $\exp(-(\log\log\xi)^{1/\sqrt{\delta}})$ finishes the proof of (5.33) since δ is small and ξ is large. □

The hierarchy constructed in Lemma 5.11 is a 'broken path' formed by the bridge paths between the starting vertices $y_0, y_1 \in \mathcal{N}$. Proposition 5.1 connects the endpoints of the high-hierarchy and constructs a connected path via common neighbours using Lemma 4.5, but not yet between y_0, y_1, only between $y_{0_{R-1}1}$ and $y_{1_{R-1}0}$, the closest vertices to y_0, y_1 in the hierarchy constructed in Lemma 5.11. Connecting y_0 to $y_{0_{R-1}1}$ and y_1 to $y_{1_{R-1}0}$ needs different techniques, since y_0, y_1 have typically lower weights than $\overline{w}$ in (5.32), see Section 6.

Proof of Proposition 5.1. To construct the path $\pi_{y_0^\star, y_1^\star}$, we again use two rounds of exposure. In the first round we apply Lemma 5.11 to get a $(\gamma, c_H\overline{w}^{4\mu}\xi^\eta, \overline{w}, c_H)$-hierarchy $\mathcal{H}_{\mathrm{high}} := \{y_\sigma\}$ of depth R fully contained in $\mathcal{N}$ with first level $\{y_0, y_1\}$. In the second round, we use Lemma 4.5 to connect, via a common neighbour, each pair of level-R siblings $y_{\sigma 0}, y_{\sigma 1}, \sigma \in \{0,1\}^{R-1} \setminus \{0_{R-1}, 1_{R-1}\}$. This yields a path between $y_{0_{R-1}1} =: y_0^\star$ and $y_{1_{R-1}0} =: y_1^\star$.

We now define an iterative cost construction on $G' \sim \{\mathcal{G}^\theta | V, w_V\}$. First we move to the exposure setting $G_1^{\theta_1}, G_2^{\theta_2}$ with $\theta_1 = \theta_2 = 1/2$. Let $\mathcal{F}_1$ be the list of all lists of vertex-pairs e with $\mathrm{dev}_{y_0y_1}(e) \le 2c_H\xi^\gamma$ that could form the bridges of a $(\gamma, c_H\overline{w}^{4\mu}\xi^\eta, \overline{w}, c_H)$-hierarchy $\mathcal{H}_{\mathrm{high}} = \{y_\sigma\}_{\sigma\in\{0,1\}^R}$ fully contained in $\mathcal{N}$ with first level $\{y_0, y_1\}$. Moreover, for any given admissible list in $\mathcal{F}_1$, let the corresponding event in $\mathcal{U}_1$ be the event that the costs of all P_σ satisfy *(H3)* of Definition 5.5 with $U = c_H\overline{w}^{4\mu}\xi^\eta$, so that $\mathcal{H}_{\mathrm{high}}$ is indeed a valid $(\gamma, c_H\overline{w}^{4\mu}\xi^\eta, \overline{w}, c_H)$-hierarchy with $\mathrm{dev}_{y_0y_1}(\mathcal{H}_{\mathrm{high}}) \le 2c_H\xi^\gamma$. For each $\sigma \in \{0,1\}^{R-1} \setminus \{0_{R-1}, 1_{R-1}\}$, define $\mathcal{J}(\sigma)$ to be the set of all potential paths J_σ fully contained in $\mathcal{N}$ connecting $y_{\sigma 0}$ and $y_{\sigma 1}$ with $\mathrm{dev}_{y_0y_1}(J_\sigma) \le 3c_H\xi^\gamma$. Given (V, w_V, S_1), call a list $\underline{t}$ of edges *admissible* if it contains exactly one such potential path J_σ from $\mathcal{J}(\sigma)$ for each $\sigma \in \{0,1\}^R \setminus \{0_{R-1}, 1_{R-1}\}$,

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

and let $\mathcal{F}_2(G_1^{\theta_1})$ be the collection of all admissible edge-lists, with an arbitrary ordering. For any given list $\underline{t} \in \mathcal{F}_2(G_1^{\theta_1})$, let the corresponding event in $\mathcal{U}_2(G_1^{\theta_1})$ be that the round-2 marginal costs of edges in $\underline{t}$ are such that $\mathrm{mcost}_2(J_\sigma) \le \overline{w}^{4\mu}$ in (3.1). This defines an iterative cost construction $\mathrm{Iter} = ((\mathcal{F}_1, \mathcal{U}_1), (\mathcal{F}_2, \mathcal{U}_2))$ applied on $G_1^{\theta_1}, G_2^{\theta_2}$ that we denote by $\mathrm{Iter}^{\exp}_{\{\mathcal{G}^\theta | V, w_V\}, \underline{\theta}}$ with $\underline{\theta} := (\theta/2, \theta/2)$. Recall from Def. 3.6(vi) that for $i \in \{1,2\}$, $\mathcal{S}_i^{\exp}$ is either None or lies in $\mathcal{F}_i$ with round-i marginal costs satisfying $\mathcal{U}_i$.

We claim that if $\mathrm{Iter}^{\exp}_{\{\mathcal{G}^\theta | V, w_V\}, \underline{\theta}}$ succeeds, then there is a path $\pi_{y_0^\star, y_1^\star} \subseteq \mathcal{N}$ between $y_0^\star := y_{0_{R-1}1}$ and $y_1^\star := y_{1_{R-1}0}$ with $\mathcal{C}(\pi_{y_0^\star, y_1^\star}) \le c_H 2^R \overline{w}^{4\mu} \xi^\eta$ and $\mathrm{dev}_{y_0 y_1}(\pi_{y_0^\star, y_1^\star}) \le 3 c_H \xi^\gamma$. Indeed, let us order the elements y_σ of $\mathcal{H}_{\mathrm{high}}$ lexicographically by their index σ, omitting y_0 and y_1, that is

$$y_0^\star = y_{0_{R-1}1}, y_{0_{R-2}10}, y_{0_{R-2}11}, \ldots, y_{1_{R-2}00}, y_{1_{R-2}01}, y_{1_{R-1}0} = y_1^\star,$$

and notice that $P_\sigma \in \mathcal{H}_{\mathrm{high}}$ is a path between every consecutive pair of the form $y_{\sigma 01}, y_{\sigma 10}$ while J_σ is a path between every consecutive pair of the form $y_{\sigma 00}, y_{\sigma 01}$ or $y_{\sigma 10}, y_{\sigma 11}$, so the concatenation forms a connected walk π^+. We then remove any cycles from π^+, passing to an arbitrary sub-path $\pi_{y_0^\star, y_1^\star} \in \mathcal{N}$. Since $\mathcal{H}_{\mathrm{high}}$ is a $(\gamma, c_H \overline{w}^{4\mu} \xi^\eta, \overline{w}, c_H)$ hierarchy with first level y_0, y_1, by Definition 5.5 *(H1)* $w_{y_0^\star}, w_{y_1^\star} \in [\overline{w}, 4\overline{w}]$, and by *(H2)*, the distances $|y_0 - y_0^\star| \le c_H \xi^{\gamma^{R-1}}$ and $|y_1 - y_1^\star| \le c_H \xi^{\gamma^{R-1}}$ both hold. Finally, since each edge of π^+ is contained in $\pi_{y_0^\star, y_1^\star}$ only once, its cost is at most

$$\mathcal{C}(\pi_{y_0^\star, y_1^\star}) \le \sum_{e \in E(\pi^+)} \mathcal{C}(e) \le \sum_{\sigma \in \{0,1\}^i : 0 \le i \le R-2} \mathrm{mcost}_1(P_\sigma) + \sum_{\sigma \in \{0,1\}^{R-1} \setminus \{0_{R-1}, 1_{R-1}\}} \mathrm{mcost}_2(J_\sigma).$$

The marginal cost of each P_σ is at most $c_H \overline{w}^{4\mu} \xi^\eta$ by $\mathcal{H}_{\mathrm{high}}$ (see *(H3)*), and $\mathrm{mcost}_2(J_\sigma) \le \overline{w}^{4\mu}$ by construction; since $c_H \ge 1$ it follows that

$$\mathcal{C}(\pi_{y_0^\star, y_1^\star}) \le (2^{R-1} - 1) c_H \overline{w}^{4\mu} \xi^\eta + (2^{R-1} - 2) \overline{w}^{4\mu} < c_H 2^R \overline{w}^{4\mu} \xi^\eta,$$

as required by $\mathcal{X}_{\text{high-h}}$. The deviation bound $3 c_H \xi^\gamma$ also holds since it holds individually for all J_σ and it holds for $\mathcal{H}_{\mathrm{high}}$ already by Lemma 5.11.

It remains to lower-bound the probability that $\mathrm{Iter}^{\exp}_{\{\mathcal{G}^\theta | V, w_V\}, \underline{\theta}}$ succeeds. Again as in Def. 3.8 and Prop. 3.9, let $H_1, H_2 \sim \{\mathcal{G}^{\theta/2} \mid V, w_V\}$ independently, let $I^{\mathrm{ind}}_{\{\mathcal{G}^\theta | V, w_V\}, \underline{\theta}}$ be the result of applying $((\mathcal{F}_1, \mathcal{U}_1), (\mathcal{F}_2, \mathcal{U}_2))$ to H_1 and H_2, and let $\mathcal{A}^{\mathrm{ind}}(S_1)$ be the event that the first round returns the edge set $\mathcal{S}_1^{\mathrm{ind}} = S_1$. Then Proposition 3.9 followed by a union bound gives similarly to (5.34) that

$$\begin{aligned} \mathbb{P}(\mathcal{X}_{\text{high-path}} \mid V, w_V) &\ge \mathbb{P}\big(I^{\exp}_{\{\mathcal{G}^\theta | V, w_V\}, \underline{\theta}} \text{ succeeds} \mid V, w_V\big) \\ &\ge 1 - \mathbb{P}\big(\mathcal{S}_1^{\mathrm{ind}} = \texttt{None} \mid V, w_V\big) - \max_{S_1 \ne \texttt{None}} \mathbb{P}\big(\mathcal{S}_2^{\mathrm{ind}} = \texttt{None} \mid \mathcal{A}^{\mathrm{ind}}(S_1)\big). \end{aligned} \tag{5.38}$$

We bound both errors on the right-hand side. The event $\mathcal{S}_1^{\mathrm{ind}} \ne$ None can be bounded using Lemma 5.11 with θ replaced by $\theta/2$, since $H_1 \sim \{\mathcal{G}^{\theta/2} \mid V, w_V\}$. All the requirements of Lemma 5.11 are fulfilled by hypothesis, so the first term on the right-hand side is at most $\exp\big(-(\log\log\xi)^{13}\big)$ by (5.33).

It remains to upper-bound the second term in (5.38). For each $x_0, x_1 \in \mathcal{N}$, let $\widetilde{\mathcal{A}}_{x_0 \star x_1}$ be the event that H_2 contains a two-edge path $x_0 v x_1 \subseteq \mathcal{N}$ of cost at most $\overline{w}^{4\mu}$ with $|x_0 - v| \le c_H \xi^\gamma$. If $\mathrm{dev}_{y_0 y_1}(x_0) \le 2 c_H \xi^\gamma$, this implies $\mathrm{dev}_{y_0 y_1}(v) \le 3 c_H \xi^\gamma$. Hence, conditioned on $\mathcal{A}^{\mathrm{ind}}(S_1)$, the event $\mathcal{S}_2^{\mathrm{ind}}$ = None occurs only if for some $\sigma \in \{0,1\}^{R-1} \setminus \{0_{R-1}, 1_{R-1}\}$ the complement of the event $\widetilde{\mathcal{A}}_{y_{\sigma 0} \star y_{\sigma 1}}$ occurs. Since all the events in $\mathcal{A}^{\mathrm{ind}}(S_1)$ are contained in the σ-algebra generated by H_1, which

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

is independent of H_2 given V, w_V, we get by a union bound that

$$\begin{aligned}\max_{S_1\neq \texttt{None}} \mathbb{P}\big(\mathcal{S}_2^{\mathrm{ind}} = \texttt{None} \mid \mathcal{A}^{\mathrm{ind}}(S_1)\big) &\\ \leq (2^{R-1}-2)\cdot &\max_{\substack{\sigma\in\{0,1\}^{R-1}\setminus\{0_{R-1},1_{R-1}\}\\ S_1\neq\texttt{None}}} \mathbb{P}\big(\widetilde{\mathcal{A}}^{\complement}_{y_{\sigma0}\star y_{\sigma1}} \mid \mathcal{A}^{\mathrm{ind}}(S_1)\big)\\ \leq (2^{R-1}-2)\cdot &\max_{y_{\sigma0},y_{\sigma1}\neq\texttt{None}} \mathbb{P}\big(\widetilde{\mathcal{A}}^{\complement}_{y_{\sigma0}\star y_{\sigma1}} \mid V, w_V\big),\end{aligned} \tag{5.39}$$

where the maximum is taken over all possible values of $y_{\sigma0}, y_{\sigma1}$ occurring in non-None S_1. To bound (5.39), we observe $|y_{\sigma0}-y_{\sigma1}| \leq c_H\xi^{\gamma^{R-1}}$ and $w_{y_{\sigma0}}, w_{y_{\sigma1}} \in [\overline{w}, 4\overline{w}]$ when $\sigma\in\{0,1\}^R\setminus\{0_{R-1},1_{R-1}\}$, by Def. 5.5 *(H2)* and *(H1)*, since $\mathcal{H}_{\mathrm{high}}$ is a $(\gamma, c_H\overline{w}^{4\mu}\xi^{\eta}, \overline{w}, c_H)$ hierarchy. Thus we apply Lemma 4.5 with G' replaced by H_2, θ replaced by $\theta/2$, $D=\xi^{\gamma^{R-1}}$, $x_0 = y_{\sigma0}$, $x_1 = y_{\sigma1}$ and all other variables taking their present values. We verify that the requirements of Lemma 4.5 all hold in order of their appearance there. $\delta \ll_\star \texttt{par}$ by hypothesis and $D=\xi^{\gamma^{R-1}} \geq (\log\log \xi\sqrt{d})^{16d/\delta^2} \geq (\log\log\xi\sqrt{d})^{16/\delta}$ by assumption, so in particular $D \gg_\star c_H, \delta$ and $D \geq w_0^{2/d}$ since $\xi \gg_\star \texttt{par}, \delta, w_0$. Also, clearly $\xi^{\gamma^{R-1}} \leq \xi\sqrt{d}$. Next, we check the distance and weights of $y_{\sigma0}, y_{\sigma1}$. Since $S_1\neq\texttt{None}$, they must lie in $\mathcal{N}$, and as argued already, satisfy $|y_{\sigma0}-y_{\sigma1}| \leq c_H\xi^{\gamma^{R-1}} = c_H D$ and $w_{y_{\sigma0}}, w_{y_{\sigma1}} \in [\overline{w}, 4\overline{w}] = [D^{d/2}, 4D^{d/2}]$. Finally, the cost-bound in Lemma 4.5 is $D^{2\mu d} = \xi^{2\mu d\gamma^{R-1}} = \overline{w}^{4\mu}$ exactly as we require it here, and the vertex v satisfies $|x_0 - v| \leq D = \xi^{\gamma^{R-1}} \leq c_H\xi^{\gamma}$, also as required. Lemma 4.5 applies and (5.39) can be bounded as

$$\begin{aligned}\max_{S_1\neq\texttt{None}} \mathbb{P}\big(\mathcal{S}_2^{\mathrm{ind}} = \texttt{None} \mid \mathcal{A}^{\mathrm{ind}}(S_1)\big) &\leq 2^{R-1}\exp\Big(-(\theta^2/4)\xi^{\gamma^{R-1}(3-\tau-2\delta)d/2}\Big)\\ &\leq 2^{R-1}\exp\big(-(\log\log\xi)^{15/\delta}\big),\end{aligned}$$

where for the second row we used that $\delta \ll_\star \texttt{par}$, so $(3-\tau-2\delta)d/2 \geq \delta$ and by hypothesis $\xi^{\gamma^{R-1}\delta} \geq (\log\log\xi)^{16/\delta}$ and $\xi \gg_\star \theta$. This, together with that the first error term in (5.38) was at most $\exp\big(-(\log\log\xi)^{13}\big)$ concludes the proof of (5.3). □

The goal of this section has been to prove Proposition 5.1 and Corollaries 5.2 and 5.3; now that this has been achieved, notation internal to this section will no longer be used.

6. Connecting the endpoints $0, x$ to the path

The final step is to connect the initial vertices 0 and x to the respective endpoints $y_0^\star$ and $y_x^\star$ of the path constructed in Section 5. We give a brief intuition on how we do this, then state the main result of this section, followed by the proof of the main theorems (Theorem 1.4-1.6) before the detailed proofs. When connecting the endpoints, we need to overcome the issue that the construction of the path $\pi_{y_0^\star, y_x^\star}$ already revealed information about the graph: the vertices $y_0^\star, y_x^\star$ are the outcomes of a selection procedure that might influence the graph around them. For $d\geq 2$, for some large constant M, we consider the graph G^M induced by the vertices of weight in $[M, 2M]$ restricted to edges with edge costs $\mathcal{C}_e \leq M^{3\mu}$. By a result in our companion paper [56, Corollary 3.9], this graph has an infinite component[11] C_∞^M. We connect $0, y_0^\star$ to respective nearby vertices $u_0, u_0^\star \in \mathrm{C}_\infty^M$, and then use that the cost-distance $d_{\mathcal{C}}(u_0, u_0^\star) = \Theta(|u_0 - u_0^\star|)$ within C_∞^M. We do the same for $y_x^\star$ and x. We ensure that cost-distances are linear in G^M *simultaneously* for all 'candidate' vertices for u_0 and $u_0^\star$ in Lemma 6.6 below. This overcomes the issue that $y_0^\star, y_x^\star$ are carefully chosen vertices. To obtain these results, in [56] we use a renormalisation technique to map G^M to a site-bond percolation on $\mathbb{Z}^d$ and 'pull back' density and distance results [4, 29] to G^M. In one dimension, G^M does not have an infinite component and the results of [4, 29] do not apply. So, we use

[11]Here we mean a graph-theoretical component, that is, a component with respect to edge-presence events (edge-costs ignored). Adopting standard notation we denote this infinite component by C_∞^M.

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

a finite size approach and consider the graph G^M with a value M that grows with $|x|$ to guarantee that G^M contains a large connected subgraph in the section between 0 and x. We establish the necessary density and distance bounds ourselves using paths along which the vertex-weight increase followed by renormalisation. When the graph is finite (e.g., GIRG G_n in Def. 1.3), we additionally use that (near)-shortest paths within G^M have very small deviation from the straight line (see Def. 5.6), so that when two vertices are not too close to the boundary of the box Q_n, the constructed path stays in Q_n. We define the setting of this section.

Setting 6.1. Consider 1-FPP in Definition 1.1 on the graphs IGIRG or SFP satisfying the assumptions given in (1.6)–(1.3) with $d \geq 1, \alpha \in (1, \infty], \tau \in (2, 3), \mu > 0$. Let $\underline{c}, \overline{c}, h, L, c_1, c_2, \beta$ be as in (1.5)–(1.3), we allow $\beta = \infty$ and/or $\alpha = \infty$. Let $G \sim \mathcal{G}$, let $\mathcal{F}_{0,x} := \{0, x \in \mathcal{V}\}$, and let C_∞ be the (unique) infinite component of G.

The existence of the infinite component for SFP was proved in [26], for IGIRG in [28, 57]. Uniqueness proofs exist based on adaptations of the Gandolfi-Keane-Newman argument [36] or on explicit finite-sized constructions [55, Theorem 3.11] followed by [52]. The two main results of this section are the following:

Proposition 6.2. *Consider Setting 6.1. Suppose that either $\alpha \in (1, 2)$ or $\mu \in (\mu_{\mathrm{expl}}, \mu_{\log})$ or both hold, and let Δ_0 be as defined in (1.9), (1.17), or (1.20), depending on whether $\alpha, \beta < \infty$, $\alpha = \infty$, or $\beta = \infty$. For every $q, \varepsilon, \zeta > 0$ there exists $D > 0$ such that the following holds. For any $x \in \mathbb{R}^d$ let $\mathcal{A}_{\mathrm{polylog}}$ be the event that G contains a path $\pi_{0,x}$, with endpoints 0 and x, of cost $\mathcal{C}(\pi_{0,x}) \leq (\log |x|)^{\Delta_0 + \varepsilon} + D$ and deviation* $\mathrm{dev}(\pi_{0,x}) \leq \zeta|x| + D$. *Then* $\mathbb{P}(\mathcal{A}_{\mathrm{polylog}} \mid 0, x \in \mathrm{C}_\infty) \geq 1 - q$.

Proposition 6.3. *Consider Setting 6.1. Suppose that $\alpha > 2$ and $\mu > \mu_{\log}$, and let η_0 be as defined in (1.10), (1.16), (1.19), or (1.21), depending on $\alpha, \beta < \infty$, $\beta < \alpha = \infty$, $\alpha < \beta = \infty$, or $\alpha = \beta = \infty$. For every $q, \varepsilon, \delta > 0$ there is $D > 0$ such that the following holds. For any $x \in \mathbb{R}^d$ let $\mathcal{A}_{\mathrm{pol}}$ be the event that G contains a path $\pi_{0,x}$, with endpoints 0 and x, of cost $\mathcal{C}(\pi_{0,x}) \leq |x|^{\eta_0 + \varepsilon} + D$ and deviation* $\mathrm{dev}(\pi_{0,x}) \leq \zeta|x| + D$. *Then* $\mathbb{P}(\mathcal{A}_{\mathrm{pol}} \mid 0, x \in \mathrm{C}_\infty) \geq 1 - q$.

We now explain how the proof of the main theorems follow from these propositions.

6.1. Proof of the main theorems

The proofs of Theorems 1.4 and 1.6 follow directly from Propositions 6.2 and 6.3, respectively, and so do their extensions to $\alpha = \infty$ and/or $\beta = \infty$ in Theorem 1.11. It remains to prove Theorem 1.10 treating finite graphs, including its extension to $\alpha = \infty$ and/or $\beta = \infty$.

Proof of Theorem 1.10. Following Def. 1.3, let G_n be a finite GIRG obtained by intersecting an IGIRG $G = (\mathcal{V}, \mathcal{E})$ with a finite cube Q_n of volume n, and let u_n, v_n be two vertices chosen uniformly at random from $\mathcal{V} \cap Q_n$. For the polylogarithmic case we must prove (1.13). For this, first we prove the slightly different statement that for two uniformly random *positions* $x_n, y_n \in Q_n$,

$$\lim_{n \to \infty} \mathbb{P}\big(d_{\mathcal{C}}^{G_n}(x_n, y_n) > (\log |x_n - y_n|)^{\Delta_0 + \varepsilon} \mid x_n, y_n \in \mathrm{C}_\infty\big) = 0. \tag{6.1}$$

Compared to (1.13), there are two differences. First, C_∞ replaces $\mathcal{C}_{\max}^{(n)}$ in the conditioning. By [57, Theorem 3.11] there is a constant $\rho > 0$ such that a.a.s. $|\mathcal{C}_{\max}^{(n)}| \geq \rho|\mathcal{V} \cap Q_n| \geq \rho|\mathrm{C}_\infty \cap Q_n|$, and on the other hand $\lim_{n\to\infty} \mathbb{P}\big(\mathcal{C}_{\max}^{(n)} \subseteq \mathcal{C}_\infty\big) = 1$ since C_∞ is unique. Hence, the probability of the conditions $\mathbb{P}(u_n, v_n \in \mathcal{C}_{\max}^{(n)})$ and $\mathbb{P}(u_n, v_n \in \mathrm{C}_\infty)$ differ by at most a constant factor, which means that (1.13) is equivalent to conditioning on $\{u_n, v_n \in \mathrm{C}_\infty\}$. Secondly, in (1.13) we draw two random vertices u_n, v_n from $\mathcal{V} \cap Q_n$, while in (6.1) we draw two random positions x_n, y_n and condition on those being in the vertex set. This changes the number of vertices in Q_n from Poisson(n) to Poisson(n)+2. The total variation distance of these two distributions is vanishing as $n \to \infty$, so this difference can also be ignored, and proving (6.1) implies (1.13).

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

To prove (6.1), let $C > 0$ be the constant from Proposition 6.2, let $0 < \zeta \ll_\star q, \mathtt{par}$ and consider the event $\mathcal{A}_{\mathrm{pos}}(x_n, y_n)$ that $|x_n - y_n| \geq \log n$ and that x_n, y_n have distance at least $2\sqrt{d}\zeta n^{1/d}$ from the boundary of Q_n, a box of side-length $n^{1/d}$. Then, since $\zeta \ll_\star q, \mathtt{par}$,

$$\mathbb{P}(\mathcal{A}_{\mathrm{pos}}(x_n, y_n)) \geq 1 - q/2. \tag{6.2}$$

Consider now any given realisation $x_n, y_n \in Q_n$ of the random positions that satisfy $\mathcal{A}_{\mathrm{pos}}(x_n, y_n)$. By Proposition 6.2 applied with $\varepsilon_{6.2} := \varepsilon/2, q_{6.2} := q/2$, conditional on $x_n, y_n \in \mathrm{C}_\infty$ there is a path π_{x_n,y_n} from x_n to y_n with $\mathrm{dev}(\pi_{x_n,y_n}) \leq \zeta|x_n - y_n| + C \leq 2\sqrt{d}\zeta n^{1/d}$ and cost at most $\mathcal{C}(\pi) \leq (\log|x_n - y_n|)^{\Delta_0+\varepsilon/2} + C$ with probability at least $1 - q/2$. Since $\mathcal{A}_{\mathrm{pos}}(x_n, y_n)$ holds, the deviation bound of π_{x_n,y_n} ensures that the path π_{x_n,y_n} lies fully within Q_n and thus in G_n. Moreover, since $|x_n - y_n| \geq \log n$ and n is sufficiently large, $\mathcal{C}(\pi) \leq (\log|x_n - y_n|)^{\Delta_0+\varepsilon/2} + C \leq (\log|x_n - y_n|)^{\Delta_0+\varepsilon}$. Hence, for all n large enough, whenever x_n, y_n satisfies $\mathcal{A}_{\mathrm{pos}}(x_n, y_n)$,

$$\mathbb{P}\big(d_{\mathcal{C}}^{G_n}(x_n, y_n) \leq (\log|x_n - y_n|)^{\Delta_0+\varepsilon} \mid x_n, y_n \in \mathrm{C}_\infty\big) \geq 1 - q/2. \tag{6.3}$$

Since q was arbitrary, together with (6.2), this proves (6.1) and concludes the proof for the polylogarithmic case of Theorem 1.10 (including the extensions for $\alpha = \infty$, and/or $\beta = \infty$). The polynomial case is identical except that we use Proposition 6.3 instead of Proposition 6.2. □

In the rest of the section we prove Propositions 6.2-6.3.

6.2. *Infinite weight increasing paths*

We first show a simple variant of [55, Lemma 4.3]; this lemma says that any suitably high-weight vertex is very likely to lie at the start of an infinite weight-increasing path. These weight increasing paths are necessary in dimension 1 where G^M does not percolate. Let

$$\begin{aligned} \mathcal{V}_M &:= \{v \in \mathcal{V} : w_v \in [M, 2M]\}. \\ G^M &:= (\mathcal{V}_M, \mathcal{E}_M) \quad \mathcal{E}_M := \{u, v \in \mathcal{V}_M, \mathcal{C}_{uv} \leq M^{3\mu}\}. \end{aligned} \tag{6.4}$$

Lemma 6.4. *Consider Setting 6.1 with $d = 1$. Let $\varepsilon, \delta \in (0,1)$ with $\varepsilon, \delta, \ll_\star \mathtt{par}$, and let $M_0 \gg_\star \epsilon, \delta, \mathtt{par}$. Let $z \in \mathbb{R}$ (or $\mathbb{Z}$ for SFP), and for all $i \geq 0$ define, $M_i := M_0^{(1+\varepsilon)^i}$, $R_i := M_i^{(1+\delta)(\tau-1)}$, and $I_i := [z, z + R_i]$. Let $\mathcal{A}_{\mathrm{inc}}(M_0, \varepsilon, z)$ be the event that there is an infinite path $\pi_z = z_0 z_1 \ldots$ in G starting at $z =: z_0$ such that for all $i \geq 1$ we have $z_i \in (I_i \setminus I_{i-1}) \cap \mathcal{V}_{M_i}$. Then*

$$\mathbb{P}(\mathcal{A}_{\mathrm{inc}}(M_0, \varepsilon, z)^{\complement} \mid z \in \mathcal{V}_{M_0}) \leq \exp(-M_0^{\delta(\tau-1)/4}). \tag{6.5}$$

The bound remains true if we additionally condition on $y \in \mathcal{V}$ for any $y \in \mathbb{R} \setminus \{z\}$ for GIRG.

Proof. The proof is very similar to [55, Lemma 4.3], which uses a similar construction but in more than one dimension and with less control over the weights. For all $j \geq 1$, let $\mathcal{A}_{\mathrm{inc}}^j$ be the event that there is a path $\pi_z = z_0 z_1 \ldots z_j$ in G with $z_0 := z$ such that for all $i \in [j]$ we have $z_i \in (I_i \setminus I_{i-1}) \cap \mathcal{V}_{M_i}$. Let $\mathcal{A}_{\mathrm{inc}}^0$ be the empty event. Then

$$\mathbb{P}\big(\mathcal{A}_{\mathrm{inc}}(M_0, \varepsilon, z)^{\complement} \mid z \in \mathcal{V}_{M_0}\big) = \sum_{i=1}^{\infty} \mathbb{P}\big((\mathcal{A}_{\mathrm{inc}}^i)^{\complement} \mid \mathcal{A}_{\mathrm{inc}}^{i-1} \text{ and } z \in \mathcal{V}_{M_0}\big). \tag{6.6}$$

We now bound each term in the sum of (6.6) above. Fix $i \geq 1$. Observe that $\mathcal{A}_{\mathrm{inc}}^{i-1}$ only depends on $G[I_{i-1}]$. Let G' be a possible value (realisation) of $G[I_{i-1}]$ which implies $\mathcal{A}_{\mathrm{inc}}^{i-1}$. Then we can decompose the conditioning in (6.6) by conditioning on events of the type $\mathcal{F}_{i-1} := \{G[I_{i-1}] = G'\} \cap \{z \in \mathcal{V}_{M_0}\}$ and

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

later integrating over the possible realisations G'. Given G' satisfying $\mathcal{A}_{\text{inc}}^{i-1}$, fix the vertices $z_0, \ldots, z_{i-1}$ ensuring $\mathcal{A}_{\text{inc}}^{i-1}$. Let $\mathcal{A}_{\text{vert}}^{i}$ be the event that $|(I_i \setminus I_{i-1}) \cap \mathcal{V}_{M_i}| \geq M_i^{\delta(\tau-1)/2}$; then

$$\mathbb{P}\big((\mathcal{A}_{\text{inc}}^{i})^{\complement} \mid \mathcal{F}_{i-1}\big) \leq \mathbb{P}\big((\mathcal{A}_{\text{vert}}^{i})^{\complement} \mid \mathcal{F}_{i-1}\big) + \mathbb{P}\big((\mathcal{A}_{\text{inc}}^{i})^{\complement} \mid \mathcal{A}_{\text{vert}}^{i} \cap \mathcal{F}_{i-1}\big). \tag{6.7}$$

By (1.6), the number of vertices in $(I_i \setminus I_{i-1}) \cap \mathcal{V}_{M_i}$ is independent of $\mathcal{F}_{i-1}$ and is either a Poisson variable (for IGIRG) or a binomial variable (for SFP), in both cases with mean at least

$$(R_i - R_{i-1} - 1)\left(\frac{\ell(M_i)}{M_i^{\tau-1}} - \frac{\ell(2M_i)}{(2M_i)^{\tau-1}}\right) \geq 2M_i^{\delta(\tau-1)/2},$$

where we used that ℓ is slowly varying, $\tau > 2$, and $M_i > M_0 \gg_\star \texttt{par}$ to obtain the last bound. In both IGIRG and SFP, it follows by concentration bounds (Theorem A.1) that

$$\mathbb{P}\big((\mathcal{A}_{\text{vert}}^{i})^{\complement} \mid \mathcal{F}_{i-1}\big) \leq \exp(-M_i^{\delta(\tau-1)/2}). \tag{6.8}$$

We next lower-bound the probability that z_{i-1} is connected to any given $z' \in (I_i \setminus I_{i-1}) \cap \mathcal{V}_{M_i}$. Let (V, w_V) be a possible value of $\widetilde{\mathcal{V}}$ which implies $\mathcal{A}_{\text{vert}}^{i}$, and suppose that $z' \in (I_i \setminus I_{i-1}) \cap \mathcal{V}_{M_i}$ for $\widetilde{\mathcal{V}} = (V, w_V)$. The distance between z_{i-1} and z' is at most R_i, and vertices have weight in $[M, 2M]$ in $\mathcal{V}_M$, so by (1.5) (remembering that $d = 1$),

$$\begin{aligned}\mathbb{P}(z_{i-1}z' \in \mathcal{E} \mid (V, w_V), \mathcal{F}_{i-1}) &\geq \underline{c} \cdot \min\left\{1, \tfrac{M_{i-1}M_i}{R_i}\right\}^{\alpha} \\ &= \underline{c} \cdot \min\left\{1, M_0^{(1+\varepsilon)^{i-1}[2+\varepsilon-(1+\varepsilon)(1+\delta)(\tau-1)]}\right\}^{\alpha}.\end{aligned} \tag{6.9}$$

Since $\varepsilon, \delta \ll_\star \texttt{par}$ and $\tau < 3$ we have $(1+\varepsilon)(1+\delta)(\tau-1) < 2$; thus the exponent on the right-hand side of (6.9) is positive and we obtain that the minimum is at 1 on the right hand side. By $\mathcal{A}_{\text{vert}}^{i}$ (defined above (6.7)) there are at least $M_i^{\delta(\tau-1)/2}$ such vertices z', each joined to z_{i-1} independently with probability $\underline{c} \in (0, 1)$. Thus,

$$\mathbb{P}\big((\mathcal{A}_{\text{inc}}^{i})^{\complement} \mid \mathcal{A}_{\text{vert}}^{i} \cap \mathcal{F}_{i-1}\big) \leq (1-\underline{c})^{M_i^{\delta(\tau-1)/2}} \leq \exp(-\underline{c}M_i^{\delta(\tau-1)/2}). \tag{6.10}$$

Combining (6.7), (6.8) and (6.10) and using $M_i \geq M_0 \gg_\star \texttt{par}, \delta$ yields

$$\mathbb{P}\big((\mathcal{A}_{\text{inc}}^{i})^{\complement} \mid \mathcal{F}_{i-1}\big) \leq \exp(-M_i^{\delta(\tau-1)/2}) + \exp(-\underline{c}M_i^{\delta(\tau-1)/2}) \leq \exp(-M_i^{\delta(\tau-1)/3}).$$

Substituting this bound into (6.6) and using $M_0 \gg_\star \texttt{par}, \delta$ then yields the required bound of

$$\mathbb{P}\big(\mathcal{A}_{\text{inc}}(M_0, \varepsilon, z)^{\complement} \mid z \in \mathcal{V}_{M_0}\big) \leq \sum_{i=1}^{\infty} \exp(-M_i^{\delta(\tau-1)/3}) \leq \exp(-M_0^{\delta(\tau-1)/4}).$$

The bound remains true if we additionally condition also on $y \in \mathcal{V}$: there is a unique index i so that $y \in I_i \setminus I_{i-1}$. The number of points in this interval changes by one, but the concentration bound in (6.8) still remains valid under the conditioning. □

6.3. *Being part of the infinite component*

The next lemma is a technical necessity to remove conditioning on membership of $\mathcal{C}_\infty$ later.

Lemma 6.5. *Consider Setting 6.1. There exists $\rho > 0$ such that for all distinct $a, b \in \mathbb{R}^d$, we have $\mathbb{P}(a, b \in \mathcal{C}_\infty \mid a, b \in \mathcal{V}) \geq \rho$.*

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

Proof. For $d \geq 2$, this is [56, Claim 3.10]. For $d = 1$, we instead apply Lemma 6.4. Intuitively, once the vertices have weight in $M_0, 2M_0$ they will both have weight-increasing paths whp, so ρ can be taken slightly less then the probability that both vertices have weight in this interval. Wlog, suppose $a < b$. Let $M_0 > 1$ and $0 < \delta, \varepsilon < 1$ with $\delta, \varepsilon \ll_\star \texttt{par}$ and $M_0 \gg_\star \delta, \varepsilon, \texttt{par}$. Let $\mathcal{A}_{\mathrm{path}}(a)$ be the event that $a \in \mathcal{V}$ lies in an infinite component of $G[[a, \infty)]$ (that is, G restricted to the spatial interval $[a, \infty)$), and let $\mathcal{A}_{\mathrm{path}}(b)$ be the event that b lies in an infinite component of $G[(-\infty, b]]$. By Lemma 6.4, (and the last sentence there), we have

$$\mathbb{P}\big(\mathcal{A}_{\mathrm{path}}(a)^{\complement} \mid b \in \mathcal{V}, a \in \mathcal{V}_{M_0}\big) + \mathbb{P}\big(\mathcal{A}_{\mathrm{path}}(b)^{\complement} \mid a \in \mathcal{V}, b \in \mathcal{V}_{M_0}\big) \leq 2\exp(-M_0^{\delta(\tau-1)/4}),$$

Thus by a union bound,

$$\mathbb{P}(\mathcal{A}_{\mathrm{path}}(a) \cap \mathcal{A}_{\mathrm{path}}(b) \mid a, b \in \mathcal{V}) \geq \mathbb{P}(a, b \in \mathcal{V}_{M_0} \mid a, b \in \mathcal{V}) - 2\exp(-M_0^{\delta(\tau-1)/4}). \tag{6.11}$$

By (1.6), since $M_0 \gg_\star \texttt{par}$, $\tau > 2$, and ℓ is slowly varying, we have

$$\mathbb{P}(a, b \in \mathcal{V}_{M_0} \mid a, b \in \mathcal{V}) = \left(\frac{\ell(M_0)}{M_0^{\tau-1}} - \frac{\ell(2M_0)}{(2M_0)^{\tau-1}}\right)^2 \geq \left(\frac{\ell(M_0)}{4M_0^{\tau-1}}\right)^2 \geq \frac{1}{M_0^{3(\tau-1)}}.$$

Since $M_0 \gg_\star \delta, \texttt{par}$, it follows from (6.11) that

$$\mathbb{P}(\mathcal{A}_{\mathrm{path}}(a) \cap \mathcal{A}_{\mathrm{path}}(b) \mid a, b \in \mathcal{V}) \geq M_0^{-3(\tau-1)} - 2\exp(-M_0^{\delta(\tau-1)/4}) > M_0^{-3(\tau-1)}/2;$$

since C_∞ is a.s. unique, the result therefore follows by taking $\rho := M_0^{-3(\tau-1)}/2$. □

6.4. Embedded random geometric graphs.

The next lemma, that we prove in the companion paper [56], is the main tool of the section. We first need some definitions. Recall $\mathcal{V}_M, G^M$ from (6.4). We use the notation $\pi_{a,b}$ for a path between vertices a and b. Let $r, \kappa, \zeta, C > 0$ and $z \in \mathbb{R}^d$. We say that a set of vertices $\mathcal{H} \subseteq \mathcal{V}_M$ is *r-strongly dense* around $z \in \mathbb{R}^d$ in $\mathcal{V}_M$ if the following event holds:

$$\mathcal{A}_{\mathrm{dense}}(\mathcal{H}, \mathcal{V}_M, r, z) := \Big\{\forall y \in B_r(z) : \big|B_{r^{1/3}}(y) \cap \mathcal{H}\big| \geq |B_{r^{1/3}}(y) \cap \mathcal{V}_M|/2\Big\}. \tag{6.12}$$

In words, the set $\mathcal{H}$ has local density $1/2$ around every vertex y near z. The subtlety here is that we require smaller radius around y then its distance bound from z. Consider two sets of vertices $\mathcal{H} \subseteq \mathcal{H}' \subseteq \mathcal{V}$ and a graph G on $\mathcal{V}$. We say that $\mathcal{H}$ shows *r-strongly κ-linear* distances with deviation ζ in $G[\mathcal{H}']$ around $z \in \mathbb{R}^d$ if the following event holds:

$$\begin{aligned}\mathcal{A}_{\mathrm{linear}}(\mathcal{H}, G[\mathcal{H}'], r, \kappa, \zeta, D, z) := \Big\{&\forall a \in \mathcal{B}_r(z) \cap \mathcal{H},\ \forall b \in \mathcal{H} : \exists \text{ a path } \pi_{a,b} \subseteq G[\mathcal{H}'] \text{ with} \\ &\mathcal{C}(\pi_{a,b}) \leq \kappa|a-b| + D,\ \mathrm{dev}(\pi_{a,b}) \leq \zeta|a-b| + D\Big\}.\end{aligned} \tag{6.13}$$

For $d \geq 2$, we will choose $\mathcal{H}'$ to be the vertex set of the infinite component C_∞^M. For $d = 1$, we will simply set $\mathcal{H}, \mathcal{H}'$ to $\mathcal{V}_M$, restricted to some finite interval. The meaning of the event $\mathcal{A}_{\mathrm{linear}}$ is that between *every* vertex in $\mathcal{H}$ near z and every other vertex in $\mathcal{H}$ there is a path in $G[\mathcal{H}']$ with uniform bounds on the length and deviation of these paths. Finally, we say that a set $\mathcal{H}$ is (r, D)-near to $z \in \mathcal{V}$ in G if the following event holds:

$$\begin{aligned}\mathcal{A}_{\mathrm{near}}(\mathcal{H}, G, r, D, z) := \Big\{&\exists \text{ a path } \pi_{z,a} \in G \text{ to some } a \in B_r(z) \cap \mathcal{H} \text{ with} \\ &\mathcal{C}(\pi_{z,a}) \leq D,\ \mathrm{dev}(\pi_{z,a}) \leq D\Big\}.\end{aligned} \tag{6.14}$$

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

The meaning of this event is the following: $\mathcal{H}$ is the 'good' set of vertices where distances scale uniformly linearly. $\mathcal{A}_{\mathrm{near}}$ says that the set $\mathcal{H}$ is reachable from a vertex z via a cheap path with small deviation. For $d = 2$, these events are 'typical' for a dense subset of vertices in the infinite component of G^M:

Lemma 6.6. *Consider Setting 6.1 and assume $d \geq 2$. Let $M, r_1, r_2, D, \kappa > 0$ and $q, \zeta \in (0, 1)$. Whenever $D \gg_\star r_2$ and $r_1, r_2 \gg_\star M, \zeta, q, \mathtt{par}$, and $\kappa \gg_\star M$, then a.s. a unique infinite component C_∞^M of G^M exists and there is an infinite-sized vertex set $\mathcal{H}_\infty \subseteq \mathrm{C}_\infty^M$ determined by $(\widetilde{\mathcal{V}}, \mathcal{E}(G^M))$ so that $G^M[\mathcal{H}_\infty]$ is connected, and for all $z \in \mathbb{R}^d$,*

$$\mathbb{P}(\mathcal{A}_{\mathrm{dense}}(\mathcal{H}_\infty, \mathcal{V}_M, r_1, z)) \geq 1 - q/10, \qquad \mathbb{P}(\mathcal{A}_{\mathrm{near}}(\mathcal{H}_\infty, G, r_2, D, z) \mid z \in \mathrm{C}_\infty) \geq 1 - q/10, \quad (6.15)$$

$$\mathbb{P}(\mathcal{A}_{\mathrm{linear}}(\mathcal{H}_\infty, \mathrm{C}_\infty^M, r_2, \kappa, \zeta, D, z)) \geq 1 - q/10. \quad (6.16)$$

The statement remains valid conditioned on $\mathcal{F}_{y,z} = \{y, z \in \mathcal{V}\}$; moreover, the constraints on D, r, M, κ are uniform over $\{\mathcal{F}_{y,z} : y, z \in \mathbb{R}^d\}$.

We mention that D depends on r_2 but not on r_1, this caused the need of the separate notation r_1, r_2.

Proof. The lemma follows from results in [56]. There, we show that $\mathcal{H}_\infty$ exists, and is infinite and connected in G_M in Corollary 3.9(ii), by using a renormalisation to site-bond percolation. The r_2-strong κ-linearity comes from Corollary 3.9(iv) in [56] applied with $r_{3.9} = r_2$ and $C_{3.9} = D$, and the (r_2, D)-near property comes from [56, Claim 3.11]. Moreover, we can apply [56, Corollary 3.9(iii)] with $r_{3.9} = r_1$ to get the r_1-dense property with $(\log r_1)^2$ instead of $r_1^{1/3}$ and arbitrary density $1 - \varepsilon$ instead of $1/2$ in (6.12). This is a strictly stronger statement since we can cover any ball of radius $r_1^{1/3}$ with balls of radius $(\log r_1)^2$, at the cost of increasing the fraction ε of noncovered vertices by a d-dependent factor. □

For $d = 1$, the graph G^M does not have an infinite component for any M and the proof techniques in Lemma 6.6 do not apply. Instead, we directly prove the following analogous statement for G^M in a *finite* interval. In the events $\mathcal{A}_{\mathrm{near}}, \mathcal{A}_{\mathrm{linear}}$, we replace C_∞^M by the graph G^M restricted to an M-dependent (spatial) interval.

Lemma 6.7 *Consider Setting 6.1 with $d = 1$. Let $q, \zeta \in (0, 1)$, $r_M := e^{(\log M)^2}$, $\kappa_M := M^{3\mu+2}$ and $D_M := M^{2(\tau-1)+3\mu}$. Let $z \in \mathbb{R}^d$, and $\mathcal{H}_M := B_{2r_M}(z) \cap \mathcal{V}_M$. Then whenever $M \gg_\star q, \mathtt{par}$,*

$$\mathbb{P}(\mathcal{A}_{\mathrm{dense}}(\mathcal{H}_M, \mathcal{V}_M, r_M, z)) = 1, \qquad \mathbb{P}(\mathcal{A}_{\mathrm{near}}(\mathcal{H}_M, D_M, D_M, z) \mid z \in \mathrm{C}_\infty) \geq 1 - q/10, \quad (6.17)$$

$$\mathbb{P}(\mathcal{A}_{\mathrm{linear}}(\mathcal{H}_M, G^M, r_M, \kappa_M, 0, 2\kappa_M, z)) \geq 1 - q/10. \quad (6.18)$$

The statement remains valid conditioned on $\mathcal{F}_{y,z} = \{y, z \in \mathcal{V}\}$; moreover, the constraints on r_M are uniform over $\{\mathcal{F}_{y,z} : y, z \in \mathbb{R}^d\}$.

Proof. We write $r_M =: r$. Since $\mathcal{H}_M = B_{2r}(z) \cap \mathcal{V}_M$, that is, all vertices in $\mathcal{V}_M$ in $B_{2r}(z)$ belong to $\mathcal{H}_M$, also all vertices in $B_{r^{1/3}}(y) \cap \mathcal{V}_M$ are in $\mathcal{H}_M$ for all $y \in B_r(z)$, so the event $\mathcal{A}_{\mathrm{dense}}(\mathcal{H}, \mathcal{V}_M, r, z)$ always occurs by definition, as required by (6.17).

We next prove (6.18). Here, we need to prove that for every $a \in \mathcal{B}_r(z) \cap \mathcal{H}_M$ and every $\forall b \in \mathcal{H}_M$ there is a path $\pi_{a,b} \subseteq G^M$ with

$$\mathcal{C}(\pi_{a,b}) \leq \kappa_M |a - b| + 2\kappa_M, \quad \mathrm{dev}(\pi_{a,b}) \leq 2\kappa_M.$$

We divide $B_{2r}(z)$ into sub-interval 'cells' and proving that each cell is whp both connected and joined to each of its adjacent cells in $G[\mathcal{H}_M]$. To this end, let $R := M^{2/d}/(2\sqrt{d}) = M^2/2$, let $i_{\max} := \lceil 2r/R \rceil$ and $i_{\min} := -i_{\max}$. For all $i \in [i_{\min}, i_{\max}]$, let $y_i := z + i \cdot R$ and $Q^{(i)} := [y_i, y_i + R)$; thus $Q^{(i_{\min})}, \dots, Q^{(i_{\max})}$ partition $[z - R\lceil 2r/R \rceil, z + (R + 1)\lceil 2r/R \rceil) \supset B_{2r}(z)$. Let $\mathcal{A}_{\mathrm{path}}$ be the event that $G^M[Q^{(i_{\min})}], \dots, G^M[Q^{(i_{\max})}]$ are connected graphs containing at most $2R$ vertices and that for all $i \in [i_{\min}, i_{\max} - 1]$ there is at least one edge in G^M from $Q^{(i)}$ to $Q^{(i+1)}$. If $\mathcal{A}_{\mathrm{path}}$ occurs, then for all $a, b \in B_{2r}(z) \cap \mathcal{V}_M = \mathcal{H}_M$ there is a path $\pi_{a,b}$ from a to b in G^M intersecting at most $\lfloor |a-b|/R \rfloor + 2 \le |a-b|/R + 2$ many cells; since each cell contains at most $2R$ vertices and each edge in G^M has cost penalty at most $M^{3\mu}$ by (6.4), and since $2RM^{3\mu} = \kappa_M$, it follows that

$$\mathcal{C}(\pi_{a,b}) \le (|a-b|/R + 2) \cdot 2R \cdot M^{3\mu} = 2(|a-b| + 2R) \cdot M^{3\mu} \le \kappa_M |a-b| + 2\kappa_M.$$

Moreover, since π_{ab} can leave the interval $[a, b]$ at most by the lengths of cells containing a and b, so the deviation of $\pi_{a,b}$ is at most $2R$, and $2R < \kappa_M$, that is, it does not depend on $|a-b|$. With $\zeta = 0$ and $D = 2\kappa_M$, we have just shown that

$$\mathcal{A}_{\mathrm{linear}}(\mathcal{H}_M, G^M, r_M, \kappa_M, 0, 2\kappa_M, z) \subseteq \mathcal{A}_{\mathrm{path}}. \tag{6.19}$$

We now bound $\mathbb{P}(\mathcal{A}_{\mathrm{path}})$ below. We only sketch the proof, the details can be found in [56, Corollary 3.9] Consider any two vertices in $\mathcal{V}_M$ in either the same or in neighbouring cell. We use that the weights are in $[M, 2M]$ and the distance is at most $2R = M^2$, so for all $\alpha \in [1, \infty]$, it holds that

$$\mathbb{P}\big(uv \in \mathcal{E}, \mathcal{C}(uv) \le M^{3\mu} \mid u, v \in \mathcal{V}_M[Q^{(i)} \cup Q^{(i+1)}]\big) \tag{6.20}$$

$$\ge \underline{c}\left(1 \wedge \frac{W_u W_v}{|u-v|^d}\right)^{\alpha} \cdot F_L\big((W_u W_v)^{-\mu} M^{3\mu}\big) \ge \underline{c} F_L(4^{-\mu} M^{\mu}) \ge \underline{c}/2. \tag{6.21}$$

The number of vertices in each box is Poisson with mean/deterministic $R\mathbb{P}(W \in [M, 2M]) \ge M^{3-\tau-\varepsilon}$ for some $\varepsilon \ll_\star \mathtt{par}$ with $M \gg_\star \varepsilon$. We can then couple the induced graph in each cell to an Erdős-Rényi random graph and use estimates on the probability that it forms a connected graph [35]. We use concentration of the number of low cost edges between two neighbouring cells using (6.20). So, a single cell $Q^{(i)}$ satisfies the conditions in $\mathcal{A}_{\mathrm{path}}$ with probability at least $1 - e^{-M^{3-\tau-\varepsilon}}$ for some $\varepsilon \ll_\star \mathtt{par}$ with $M \gg_\star \varepsilon$. A union bound over the at most $2 \cdot \lceil 2r/R \rceil + 1$ cells yields that

$$\mathbb{P}(\mathcal{A}_{\mathrm{path}}) \ge 1 - (2 \cdot \lceil 2r/R \rceil + 1) \cdot e^{-M^{3-\tau-\varepsilon}} \ge 1 - 5re^{-M^{3-\tau-\varepsilon}}.$$

Since $r = e^{(\log M)^2}$ and $M \gg_\star \mathtt{par}$, the $e^{-M^{3-\tau-\varepsilon}}$ term dominates, and together with (6.19) and $M \gg_\star c, q$, for any $q < 1$ we obtain $\mathbb{P}(\mathcal{A}_{\mathrm{linear}}(\mathcal{H}_M, \mathcal{H}_M, r, \kappa, 0, 2\kappa, z)) \ge 1 - q/10$, and we have proved (6.18) as required. The argument conditioned on $\mathcal{F}_{0,x}$ is identical. In dimensions $d \ge 2$ [56, Corollary 3.9(i)] also explicitly allows for planted vertices.

It remains to bound $\mathbb{P}(\mathcal{A}_{\mathrm{near}}(\mathcal{H}_M, D_M, D_M, z))$ conditioned on $z \in \mathrm{C}_\infty$, see (6.14) for the definition of $\mathcal{A}_{\mathrm{near}}$. Here, we replaced the 'usual' radius $r_M = \exp((\log M)^2)$ by $D_M = M^{2(\tau-1)+3\mu} \ll r_M$, that is, we can find a path from z to a vertex with weight M within a much smaller radius from z that r_M would give. We first dominate $\mathcal{A}_{\mathrm{near}}(\mathcal{H}_M, D_M, D_M, z)$ below by events $\mathcal{A}_1$ to $\mathcal{A}_4$ defined as follows. Let $\rho \ll_\star \mathtt{par}$ be as in Lemma 6.5, and define $M_0 > 0$ satisfying $M \gg_\star M_0 \gg_\star q, \rho, \mathtt{par}$, and let $r_0 := M_0^{2(\tau-1)}$. By Lemma 6.4 we know that a.s. C_∞ contains a vertex in $\mathcal{V}_{M_0}$, and let v_0 be an (arbitrarily chosen) closest such vertex to z in Euclidean distance. We define the following events:

(C1) $\mathcal{A}_1$: there is a path π_{z,v_0} from z to v_0 with $\mathcal{C}(\pi_{z,v_0}) \le D_M/2$ and $\mathcal{V}(\pi_{z,v_0}) \subseteq B_{D_M}(z)$;
(C2) $\mathcal{A}_2$: $B_{r_0}(z)$ contains a vertex in $\mathcal{V}_{M_0} \cap \mathrm{C}_\infty$, that is, $v_0 \in B_{r_0}(z)$;
(C3) $\mathcal{A}_3$: every vertex $x \in B_{r_0}(z) \cap \mathcal{V}_{M_0}$ has an associated path $\pi_{x\to\mathcal{V}_M}$ from x to some vertex in $\mathcal{V}_M$ with $\mathcal{V}(\pi_{x\to\mathcal{V}_M}) \subset B_{D_M}(z)$; and
(C4) $\mathcal{A}_4$: $\mathcal{A}_2$ and $\mathcal{A}_3$ both occur and $\mathcal{C}(\pi_{v_0\to\mathcal{V}_M}) \le M^{3\mu} \le D_M/2$.

Observe that if $\mathcal{A}_1$, $\mathcal{A}_2$, $\mathcal{A}_3$ and $\mathcal{A}_4$ all occur then concatenating π_{z,v_0} and $\pi_{v_0\to\mathcal{V}_M}$ yields the path required by $\mathcal{A}_{\mathrm{near}}(\mathcal{H}_M, D_M, D_M, z)$; thus

$$\mathbb{P}\big(\mathcal{A}_{\mathrm{near}}(\mathcal{H}_M, D_M, D_M, z)^{\complement} \mid z \in \mathrm{C}_\infty\big) \le \sum_{i=1}^{3} \mathbb{P}\big(\mathcal{A}_i^{\complement} \mid z \in \mathrm{C}_\infty\big) + \mathbb{P}\big(\mathcal{A}_4^{\complement} \mid \mathcal{A}_2, \mathcal{A}_3, z \in \mathrm{C}_\infty\big).$$

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

By Lemma 6.5, z is in the infinite component of G with probability at least ρ, so it follows that

$$\mathbb{P}\big(\mathcal{A}_{\text{near}}(\mathcal{H}_M, D_M, D_M, z)^\complement \mid z \in \mathrm{C}_\infty\big) \le \mathbb{P}\big(\mathcal{A}_1^\complement \mid z \in \mathrm{C}_\infty\big) + \mathbb{P}\big(\mathcal{A}_2^\complement \mid z \in \mathcal{V}\big)/\rho + \mathbb{P}\big(\mathcal{A}_3^\complement \mid z \in \mathcal{V}\big)/\rho + \mathbb{P}\big(\mathcal{A}_4^\complement \mid \mathcal{A}_2, \mathcal{A}_3, z \in \mathcal{V}\big)/\rho. \tag{6.22}$$

We first bound $\mathbb{P}\big(\mathcal{A}_1^\complement \mid z \in \mathrm{C}_\infty\big)$ in (C1). Given that we fixed v_0, let π_{z,v_0} be an (arbitrarily chosen) cheapest path from z to v_0; such a path must exist whenever $z \in \mathrm{C}_\infty$. Since $\mathcal{C}(\pi_{z,v_0})$ and $\inf\{R > 0\colon \mathcal{V}(\pi_{z,v_0}) \subseteq B_R(z)\}$ are a.s.finite random variables and since $D_M \gg_\star q, \mathtt{par}, M_0$, we can choose D_M sufficiently large so that

$$\mathbb{P}\big(\mathcal{A}_1^\complement \mid z \in \mathrm{C}_\infty\big) \le q/40. \tag{6.23}$$

We next bound $\mathbb{P}\big(\mathcal{A}_2^\complement \mid z \in \mathcal{V}\big)$ in (C2). The event $\mathcal{A}_2$ occurs if and only if $B_{r_0}(z) \cap \mathcal{V}_{M_0} \cap \mathrm{C}_\infty \neq \emptyset$. Similarly as before, $|B_{r_0}(z) \cap \mathcal{V}_{M_0}|$ is either a Poisson variable (in IGIRG) or a binomial variable (in SFP) with mean

$$\mathbb{E}[|B_{r_0}(z) \cap \mathcal{V}_{M_0}| \mid z \in \mathcal{V}] \ge r_0\left(\frac{\ell(M_0)}{M_0^{\tau-1}} - \frac{\ell(2M_0)}{(2M_0)^{\tau-1}}\right) \ge 2M_0^{(\tau-1)/2}, \tag{6.24}$$

where we used $M_0 \gg_\star \mathtt{par}$ and the value of $r_0 = M_0^{2(\tau-1)}$ for the second inequality. In particular, by Chernoff's bound,

$$\mathbb{P}(|B_{r_0}(z) \cap \mathcal{V}_{M_0}| < M_0^{(\tau-1)/2} \mid z \in \mathcal{V}) \le \exp(-M_0^{(\tau-1)/8}). \tag{6.25}$$

Let $\varepsilon, \delta \in (0,1)$ satisfy $\delta, \varepsilon{-}1 \ll_\star \mathtt{par}$ and $M_0 \gg_\star \delta, \varepsilon$. Recall the event $\mathcal{A}_{\text{inc}}(M_0, \varepsilon, x)$ about having an infinite weight-increasing path from Lemma 6.4; this event implies $\{x \in \mathrm{C}_\infty\}$. So by (6.5) in Lemma 6.4, $\mathbb{P}(x \notin \mathrm{C}_\infty \mid x \in \mathcal{V}_{M_0}) \le \exp(-M_0^{\delta(\tau-1)/4})$ for some $\delta \ll_\star \mathtt{par}$ with $M_0 \gg_\star \delta$. By translation invariance, this implies

$$\mathbb{P}(x \notin \mathrm{C}_\infty \mid x \in B_{r_0}(z) \cap \mathcal{V}_{M_0}, z \in \mathcal{V}) \le \exp(-M_0^{\delta(\tau-1)/4}).$$

Hence, the expected number of vertices in $\mathcal{V}_{M_0}$ outside the infinite component is at most

$$\begin{aligned}\mathbb{E}[|(B_{r_0}(z) \cap \mathcal{V}_{M_0}) \setminus \mathrm{C}_\infty| \mid z \in \mathcal{V}] &\le \mathbb{E}[|B_{r_0}(z) \cap \mathcal{V}_{M_0}| \mid z \in \mathcal{V}] \cdot \exp(-M_0^{\delta(\tau-1)/4}) \\ &\le \mathbb{E}[|B_{r_0}(z) \cap \mathcal{V}| \mid z \in \mathcal{V}] \cdot \exp(-M_0^{\delta(\tau-1)/4}) \le (2r_0+1)\exp(-M_0^{\delta(\tau-1)/4}),\end{aligned}$$

which is a crude upper bound. It follows by Markov's inequality that

$$\begin{aligned}\mathbb{P}(|(B_{r_0}(z) \cap \mathcal{V}_{M_0}) \setminus \mathrm{C}_\infty| \ge M_0^{(\tau-1)/2}/2 \mid z \in \mathcal{V}) &\le \frac{(2r_0+1)\exp(-M_0^{\delta(\tau-1)/4})}{M_0^{(\tau-1)/2}/2} \\ &\le \exp(-M_0^{\delta(\tau-1)/8}),\end{aligned} \tag{6.26}$$

where we used $r_0 = M_0^{2(\tau-1)}$ and $M_0 \gg_\star \delta, \mathtt{par}$. By a union bound over (6.25) and (6.26),

$$\mathbb{P}\big(\mathcal{A}_2^\complement \mid z \in \mathcal{V}\big) \le \exp(-M_0^{(\tau-1)/8}) + \exp(-M_0^{\delta(\tau-1)/8}) \le \rho q/40, \tag{6.27}$$

for all sufficiently large M_0. We next bound $\mathbb{P}\big(\mathcal{A}_3^\complement \mid z \in \mathcal{V}\big)$ in (C3). For all $x \in B_{r_0}(z) \cap \mathcal{V}_{M_0}$, let $\mathcal{A}_3(x)$ be the event that x has an associated path $\pi_{x \to \mathcal{V}_M}$ as in $\mathcal{A}_3$. We restrict this path to be a weight-increasing path as in Lemma 6.4. Let $\varepsilon, \delta \in (0,1)$ satisfy $\delta, \varepsilon \ll_\star \mathtt{par}$, and require that $M_0 \gg_\star \delta, \varepsilon$. For sufficiently

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

large M, we may choose M_0 such that $i := (\log\log M - \log\log M_0)/\log(1+\varepsilon) \le \log M$ is an integer, so $M_0^{(1+\varepsilon)^i} = M$. Then by Lemma 6.4, for any given $x \in B_{r_0}(z)$, the weight increasing path reaches a vertex of weight M at radius $R_i = M_0^{(1+\varepsilon)^i(\tau-1)(1+\delta)} = M^{(\tau-1)(1+\delta)} < M^{2(\tau-1)+3\mu} = D_M$, and so the path is contained in $B_{D_M}(z)$, and we obtain

$$\begin{aligned}\mathbb{P}\big(\mathcal{A}_3(x)^\complement \mid x \in \mathcal{V}_{M_0}, z \in \mathcal{V}\big) &\le \mathbb{P}\big(\mathcal{A}_{\mathrm{inc}}(M_0,\varepsilon,x)^\complement \mid x \in \mathcal{V}_{M_0}, z \in \mathcal{V}\big)\\ &\le \exp(-M_0^{\delta(\tau-1)/4}).\end{aligned} \tag{6.28}$$

It follows by a union bound that

$$\mathbb{P}\big(\mathcal{A}_3^\complement \mid z \in \mathcal{V}\big) \le \mathbb{P}(|B_{r_0}(z) \cap \mathcal{V}_{M_0}| \ge r_0 \mid z \in \mathcal{V}) + r_0 \exp(-M_0^{\delta(\tau-1)/4}).$$

As before $|B_{r_0}(z) \cap \mathcal{V}_{M_0}|$ is either a Poisson variable (in IGIRG) or a binomial variable (in SFP) with mean bounded from above

$$\mathbb{E}(|B_{r_0}(z) \cap \mathcal{V}_{M_0}| \mid z \in \mathcal{V}) \le 2r_0\left(\frac{\ell(M_0)}{M_0^{\tau-1}} - \frac{\ell(2M_0)}{(2M_0)^{\tau-1}}\right) \le r_0 M_0^{-(\tau-1)/2}.$$

Therefore $\mathbb{P}(|B_{r_0}(z) \cap \mathcal{V}_{M_0}| \ge r_0 \mid z \in \mathcal{V}) \le 2^{-r_0}$ and we get

$$\mathbb{P}\big(\mathcal{A}_3^\complement \mid z \in \mathcal{V}\big) \le 2^{-r_0} + r_0 \exp(-M_0^{\delta(\tau-1)/4}) \le \rho q/40, \tag{6.29}$$

where the second inequality holds because because $r_0 = M_0^{2(\tau-1)}$ and $M_0 \gg_\star \delta, \rho, q, \mathtt{par}$.

Finally we bound the last term in (6.22), (see $\mathcal{A}_4$ in (C4)). Conditioned on the realisation of G, any path $\pi_{v_0 \to \mathcal{V}_M}$ that satisfies the weight-increasing path property in (6.28) has at most $\log M$ edges and all vertex weights at most M. So, its expected cost is at most

$$\mathbb{E}[\mathcal{C}(\pi_{v_0 \to \mathcal{V}_M})] \le |\mathcal{E}(\pi_{v_0 \to \mathcal{V}_M})| \cdot M^{2\mu}\mathbb{E}[L] \le M^{2\mu}\mathbb{E}[L] \log M \le \rho q M^{3\mu}/40$$

since $M \gg_\star \rho, q, \mathtt{par}$. Thus by Markov's inequality, the probability that the cost of this path is larger than $M^{3\mu}$ is at most $\rho q/40$.

$$\mathbb{P}\big(\mathcal{A}_4^\complement \mid \mathcal{A}_2, \mathcal{A}_3, z \in \mathcal{V}\big) \le \rho q/40. \tag{6.30}$$

The result therefore follows on substituting the bounds (6.23), (6.27), (6.29), and (6.30) into (6.22). The argument conditioned on $\mathcal{F}_{y,z}$ is identical; note in particular that in applying Lemma 6.4, we may assume wlog that $y < z$ by symmetry. □

6.5. Connecting 'down' from high to low weight vertices cheaply.

We use the next lemma to connect the endpoints $y_0^\star, y_x^\star$ of the path $\pi_{y_0^\star, y_x^\star}$ obtained in Corollaries 5.2 and 5.3 to $\mathcal{H}_\infty \subseteq \mathrm{C}_\infty^M$ from Lemma 6.6 (when $d \ge 2$) and $\mathcal{H}_M$ from Lemma 6.7 (when $d = 1$). Recall $\mathcal{E}_M$ from (6.4) and that $\mathcal{F}_{0,x} = \{0, x \in \mathcal{V}\}$.

Lemma 6.8. *Consider Setting 6.1 and any $d \ge 1$. Let $w \gg_\star q, \mathtt{par}$ with $w \ge M^{8(\tau-1)}$, let $r := w^{3/d}$ and let $z \in \mathbb{R}^d$. Let $\mathcal{H} \subseteq \mathcal{V}_M$ be a random vertex set which depends only on $(V, w_V, \mathcal{E}_M)$ and which satisfies $\mathbb{P}(\mathcal{A}_{\mathrm{dense}}(\mathcal{H}, \mathcal{V}_M, r, z) \mid \mathcal{F}_{0,x}) \ge 1 - q/10$. Let*

$$\mathcal{A}_{\mathrm{down}}(w,z) := \Big\{\forall y \in \widetilde{\mathcal{V}} \cap (B_r(z) \times [w, 4w]) : \exists u \in \mathcal{H} \cap B_{r^{1/3}}(y),\ yu \in \mathcal{E},\ \mathcal{C}(yu) \le w^{2\mu}\Big\}.$$

Then for all $z \in \mathbb{R}^d$, $\mathbb{P}(\mathcal{A}_{\mathrm{down}}(w,z) \mid \mathcal{F}_{0,x}) \ge 1 - q/3$.

Informally, this lemma states that *every* vertex that has fairly high weight near z has a direct cheap edge to a nearby vertex that has weight in $[M, 2M]$ and is part of the well-connected sets C_∞^M (in dim 2), $\mathcal{H}_M$ in dim 1.

Proof. Fix $z \in \mathbb{R}^d$ and let $\mathcal{A}_1 := \mathcal{A}_{\mathrm{dense}}(\mathcal{H}, \mathcal{V}_M, r, z)$ in (6.12), so that $\mathbb{P}(\mathcal{A}_1^{\complement} \mid \mathcal{F}_{0,x}) \le q/10$ by hypothesis. Considering the definition of $\mathcal{A}_{\mathrm{dense}}$ in (6.12), let $\mathcal{A}_2$ be the event that for all $y \in B_r(z)$, $|B_{r^{1/3}}(y) \cap \mathcal{V}_M| \ge r^{d/4}$. Choose fixed points $x_1, \ldots, x_{\lceil r^d \rceil} \in B_r(z)$ such that $\{B_{r^{1/3}/2}(x_i) \colon i \le \lceil r^d \rceil\}$ covers $B_r(z)$. For all y, the ball $B_{r^{1/3}}(y)$ must contain at least one ball $B_{r^{1/3}/2}(x_i)$, so if in each ball $B_{r^{1/3}/2}(x_i)$ we find at least $r^{d/4}$ vertices from $\mathcal{H} \subseteq \mathcal{V}_M$ then the event $\mathcal{A}_2$ holds. Let c_d denote the volume of a unit-radius d-dimensional ball. By (1.6), in IGIRG $|B_{r^{1/3}/2}(x_i) \cap \mathcal{V}_M|$ is a Poisson variable with mean

$$2^{-d} c_d r^{d/3} \left(\frac{\ell(M)}{M^{\tau-1}} - \frac{\ell(2M)}{(2M)^{\tau-1}} \right) \ge 2 r^{d/3} M^{-3(\tau-1)/2} \ge 2 r^{d/4}$$

(also conditioned on $\mathcal{F}_{0,x}$), where the first inequality holds because $M \gg_\star \mathtt{par}$ and the second inequality holds since $r^{d/12} = w^{1/4} \ge M^{2(\tau-1)}$ and $M \ge 1$. Similarly, for SFP it is a binomial variable with mean greater than $2r^{d/4}$; in either case, the Chernoff bound of Theorem A.1 applies, and since $r \gg_\star q$ we have

$$\begin{aligned} \mathbb{P}(\mathcal{A}_2 \mid \mathcal{F}_{0,x}) &= \mathbb{P}\Big(\forall y \in B_r(z) : |B_{r^{1/3}}(y) \cap \mathcal{V}_M| \ge r^{d/4} \mid \mathcal{F}_{0,x}\Big) \\ &\ge \mathbb{P}(\forall i \colon |B_{r^{1/3}/2}(x_i) \cap \mathcal{V}_M| \ge r^{d/4} \mid \mathcal{F}_{0,x}) \ge 1 - \lceil r^d \rceil \cdot e^{-r^{d/4}/4} \ge 1 - q/30. \end{aligned} \tag{6.31}$$

Let $\mathcal{A}_3$ be the event that $B_r(z)$ contains at most $2(c_d r^d + 2)$ vertices. In SFP, $\mathbb{P}(\mathcal{A}_3 \mid \mathcal{F}_{0,x}) = 1$; in IGIRG, Theorem A.1 applies. In both cases, using $r \gg_\star q, \mathtt{par}$,

$$\mathbb{P}(\mathcal{A}_3 \mid \mathcal{F}_{0,x}) = \mathbb{P}(|B_r(z) \cap \mathcal{V}| \le 2(c_d r^d + 2) \mid \mathcal{F}_{0,x}) \ge 1 - e^{-c_d r^d/3} \ge 1 - q/30. \tag{6.32}$$

Since $\mathbb{P}(\mathcal{A}_1^{\complement} \mid \mathcal{F}_{0,x}) \le q/10$, a union bound with (6.31) and (6.32) yields $\mathbb{P}(\mathcal{A}_1 \cap \mathcal{A}_2 \cap \mathcal{A}_3 \mid \mathcal{F}_{0,x}) \ge 1 - q/6$. We abbreviate $\mathbb{P}(\cdot \mid V, w_V, E_M)$ when we condition on the event that $\widetilde{\mathcal{V}} = (V, w_V)$ and $\mathcal{E}_M = E_M$. The events $\mathcal{F}_{0,x}, \mathcal{A}_2, \mathcal{A}_3$, and also the set $\mathcal{H}$ and thus $\mathcal{A}_1$ are all determined by (V, w_V, E_M). Let us call the realisation (V, w_V, E_M) *good* if the event $\mathcal{A}_1 \cap \mathcal{A}_2 \cap \mathcal{A}_3 \cap \mathcal{F}_{0,x}$ holds. Then

$$\mathbb{P}(\mathcal{A}_{\mathrm{down}}(w,z) \mid \mathcal{F}_{0,x}) \ge 1 - q/6 - \max_{(V, w_V, E_M) \text{ good}} \mathbb{P}\big(\mathcal{A}_{\mathrm{down}}(w,z)^{\complement} \mid V, w_V, E_M\big). \tag{6.33}$$

Fix a good realisation (V, w_V, E_M). Following $\mathcal{A}_{\mathrm{down}}$, let $y_1, \ldots, y_k$ be the (fixed) vertices in $B_r(z)$ with weights in $[w, 4w]$, and for each $i \in [k]$ let $a_1^{(i)}, \ldots, a_{\ell_i}^{(i)}$ be the (fixed) vertices in $B_{r^{1/3}}(y_i) \cap \mathcal{H}$. Thus, by definition of $\mathcal{A}_{\mathrm{down}}$,

$$\begin{aligned} &\mathbb{P}\big(\mathcal{A}_{\mathrm{down}}(w,z)^{\complement} \mid V, w_V, E_M\big) \\ &\quad = \mathbb{P}\big(\exists i \in [k] \colon \forall j \colon y_i a_j^{(i)} \notin \mathcal{E}(G) \text{ or } \mathcal{C}(y_i a_j^{(i)}) > w^{2\mu} \mid V, w_V, E_M\big). \end{aligned} \tag{6.34}$$

Conditioned on (V, w_V, E_M), the edges $y_i a_j^{(i)}$ are present independently since $w_{y_i} \ge w \ge M^{8(\tau-1)} > 2M$ since $\tau > 2$ and $M \gg_\star \mathtt{par}$. Since $w_{y_i} \in [w, 4w]$, $w_{a_j^{(i)}} \in [M, 2M]$, and $|y_i - a_j^{(i)}| \le r^{1/3} = w^{1/d}$, we get using (1.5) and (1.2) that for all i and j,

$$\mathbb{P}(y_i a_j^{(i)} \notin \mathcal{E}(G) \text{ or } \mathcal{C}(y_i a_j^{(i)}) > w^{2\mu} \mid V, w_V, E_M) \le 1 - \underline{c}\big(1 \wedge \tfrac{wM}{w}\big)^{\alpha} + \mathbb{P}\big(L > \tfrac{w^{2\mu}}{(8wM)^{\mu}}\big) \le 1 - \tfrac{\underline{c}}{2},$$

where the last inequality holds because $w^{2\mu}/(8wM)^{\mu} \ge M^{8\mu(\tau-1)-\mu}/8^{\mu}$ tends to infinity with M and $M \gg_\star \mathtt{par}$. This computation also holds for $\alpha = \infty$ or $\beta = \infty$. Conditioned on a good realisation

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

(V, w_V, E_M), $\mathcal{A}_1 \cap \mathcal{A}_2 \cap \mathcal{A}_3$ occurs, so for each i, the number of vertices $a_j^{(i)}$ is $\ell_i \geq r^{d/4}$, and the number of vertices y_i is $k \leq 2c_d r^d + 4$. By independence across j, (6.34), and a union bound,

$$\begin{aligned}\mathbb{P}\big(\mathcal{A}_{\text{down}}(w,z)^{\complement} \mid V, w_V, E_M\big) &\leq \sum_{i\leq k} \mathbb{P}\big(\forall j\colon y_i a_j^{(i)} \notin \mathcal{E} \text{ or } \mathcal{C}(y_i a_j^{(i)}) > w^{2\mu} \mid V, w_V, E_M\big)\\ &\leq \sum_{i\leq k}(1-\underline{c}/2)^{\ell_i} \leq (2c_d r^d + 4)e^{-r^{d/4}\underline{c}/2} \leq q/6,\end{aligned}$$

where the last inequality holds since $r = w^{3/d} \gg_\star q, \texttt{par}$. The lemma then follows by (6.33). □

6.6. *Connecting* $0, x$ *to the endpoints of the hierarchy.*

We are now ready to prove the main results of this section, Propositions 6.2-6.3 that connects $0, x$ to the endpoint of the constructed paths in Corollaries 5.2-5.3, respectively. Recall $\mu_{\log}, \mu_{\text{pol}}$ from (1.8).

Proof of Proposition 6.2. We first prove the result for $d \geq 2$, then describe the necessary modifications for $d = 1$. Let $\rho > 0$ be as in Lemma 6.5 and let $\delta \ll_\star \varepsilon, q, \rho, \texttt{par}$ and $w_0 > 1$. We want to apply Corollary 5.2, which holds for sufficiently large $|x|$. Thus there exists $r_{5.2} \gg_\star q, \delta, \varepsilon, \zeta, w_0, \texttt{par}$ such that Corollary 5.2 is applicable whenever $|x| \geq r_{5.2}$. We may also assume $r_{5.2} \gg_\star \kappa$. To cover the case $|x| < r_{5.2}$, for all $v \in \mathrm{C}_\infty \cap B_{r_{5.2}}(0)$, pick the cheapest path $\pi_{0,v}$ from 0 to v. Then $R_1 := \max\{\mathcal{C}(\pi_{0,v}) : v \in \mathrm{C}_\infty \cap B_{r_{5.2}}(0)\}$ and $R_2 := \max\{\text{dev}(\pi_{0,v}) : v \in \mathrm{C}_\infty \cap B_{r_{5.2}}(0)\}$ are almost surely finite random variables, and since we may assume $D \gg_\star r_{5.2}, q, \texttt{par}$, we have $\mathbb{P}(R_1, R_2 \leq D \mid 0, x \in \mathrm{C}_\infty) \geq 1 - q$, as required. So from now on we may assume $|x| \geq r_{5.2}$.

Let $\overline{w}$ be as in (5.4) in Corollary 5.2 and let $r_1 := \overline{w}^{3/d}$. Let $M, r_2, \kappa > 0$ satisfy $\overline{w}, D \gg_\star r_2 \gg_\star M, \zeta, q, \rho, \varepsilon, \texttt{par}$ and $D \gg_\star r_{5.2} \gg_\star \kappa \gg_\star M$ as in Lemma 6.6, and note that $|x| \geq r_{5.2}$ implies $\overline{w}, r_1 \gg_\star \kappa$. Let Q be a cube of side length $|x|$ containing 0 and x, and let $\mathcal{A}_{\text{net}}$ be the event that Q contains a weak $(\delta/4, w_0)$-net (as in Definition 2.1) which contains 0 and x. Apply Corollary 5.2 with $\varepsilon_{5.2} = \varepsilon/2$ and $q_{5.2} := q\rho/5$ to obtain $\mathcal{X}_{\text{polylog}}(0, x)$. Then consider the intersection of the following events from Corollary 5.2, Lemma 6.6 (defined in (6.13), (6.14)) and Lemma 6.8:

$$\begin{aligned}\mathcal{A} &:= \mathcal{A}_{\text{net}} \cap \mathcal{X}_{\text{polylog}}(0,x)\\ &\cap \bigcap_{v\in\{0,x\}} \Big(\mathcal{A}_{\text{near}}(\mathcal{H}_\infty, r_2, D/4, v) \cap \mathcal{A}_{\text{linear}}(\mathcal{H}_\infty, \mathcal{C}_\infty^M, r_2, \kappa, \zeta, D/4, v) \cap \mathcal{A}_{\text{down}}(\overline{w}, v)\Big).\end{aligned} \tag{6.35}$$

When $\mathcal{A}$ occurs, $\mathcal{X}_{\text{polylog}}(0, x)$ gives a path between endpoints $y_0^\star, y_x^\star$ with weights in $[\overline{w}, 4\overline{w}]$ and within distance $r_1 = \overline{w}^{3/d}$ from $0, x$, with cost $\mathcal{C}(\pi_{y_0^\star, y_x^\star}) \leq (\log|x|)^{\Delta_0+\varepsilon/2}$ and deviation $\text{dev}_{0x}(\pi_{y_0^\star,y_x^\star}) \leq \zeta|x|$, respectively, see (5.4)–(5.6). Then, the events $\mathcal{A}_{\text{down}}(\overline{w}, 0), \mathcal{A}_{\text{down}}(\overline{w}, x)$ from Lemma 6.8 applied respectively to $y_0^\star, y_x^\star$ give us two paths $\pi_{y_0^\star,u_0^\star}$ and $\pi_{y_x^\star,u_x^\star}$ with $u_0^\star, u_x^\star \in \mathcal{H}_\infty$ and within respective distance $\overline{w}^{1/d}$ from $y_0^\star, y_x^\star$, and cost at most $\overline{w}^{2\mu}$. Further, the events $\mathcal{A}_{\text{near}}(\mathcal{H}_\infty, r_2, D/4, 0)$ and $\mathcal{A}_{\text{near}}(\mathcal{H}_\infty, r_2, D/4, x)$ in (6.14) also give us two paths π_{0,u_0} and π_{x,u_x}, with respective endpoints $u_0, u_x \in \mathcal{H}_\infty$ within distance $r_2 \leq \overline{w}^{3/d}$ from $0, x$ respectively, and cost at most $D/4$ each. Finally, since $u_0, u_0^\star, u_x, u_x^\star \in \mathcal{H}_\infty$, and u_0, u_x is within distance r_2 from $0, x$, respectively, the events $\mathcal{A}_{\text{linear}}(\mathcal{H}_\infty, \mathcal{C}_\infty^M, r_2, \kappa, \delta, D/4, v), v \in \{0, x\}$ in (6.13) ensure that there exist paths $\pi_{u_0,u_0^\star}$ and $\pi_{u_x,u_x^\star}$ in G that have cost at most $\kappa|u_v - u_v^\star| + D/4 \leq \kappa 3\overline{w}^{3/d} + D/4$ since $|u_v - u_v^\star| \leq 3r_1 = 3\overline{w}^{3/d}$ and deviation at most $\zeta|u_v - u_v^\star| + D/4 \leq \zeta 3\overline{w}^{3/d} + D/4$. The concatenated path is $\pi_{0,x} := \pi_{0,u_0}\pi_{u_0,u_0^\star}\pi_{u_0^\star,y_0^\star}\pi_{y_0^\star,y_x^\star}\pi_{y_x^\star,u_x^\star}\pi_{u_x^\star,u_x}\pi_{u_x,x}$. Then, since $\overline{w} \leq (\log|x|)^{\varepsilon/2}$ in (5.4), we can estimate the cost, and using that the vertices of the paths $\pi_{0,u_0}, \pi_{u_0,u_0^\star}, \pi_{y_0^\star,y_x^\star}, \pi_{y_x^\star,u_x^\star}, \pi_{u_x^\star,u_x}, \pi_{u_x,x}$ are all within distance $3r_1 = 3\overline{w}^{3/d}$ from 0 and x respectively, we can bound cost and deviation as

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

$$\begin{aligned}\mathcal{C}(\pi_{0,x}) &\le 2\cdot D/4 + 2\overline{w}^{2\mu} + 2(\kappa 3\overline{w}^{3/d} + D/4) + (\log|x|)^{\Delta_0+\varepsilon/2} \le (\log|x|)^{\Delta_0+\varepsilon} + D,\\ \mathrm{dev}(\pi_{0,x}) &\le \max\left\{\mathrm{dev}_{0,x}(\pi_{y_0^\star,y_x^\star}), 2\overline{w}^{3/d} + \zeta 3\overline{w}^{3/d} + D/4\right\} \le \zeta|x| + D,\end{aligned} \tag{6.36}$$

using $|x| \ge r_{5.2} \gg_\star \varepsilon, \kappa, \mathtt{par}$. Thus $\mathcal{A} \subseteq \mathcal{A}_{\mathrm{polylog}}$. A union bound on the complement of the events in (6.35) from Lemma 2.2 with $t = 2$ and $\varepsilon_{2.2} := \delta/4$ for $\mathcal{A}_{\mathrm{net}}$, Corollary 5.2, Lemma 6.6 with $q_{6.6} := q\rho$, and Lemma 6.8 with $q_{6.8} := q\rho$ gives

$$\begin{aligned}\mathbb{P}(\mathcal{A}^{\complement}_{\mathrm{polylog}} \mid 0,x \in \mathrm{C}_\infty) \le \mathbb{P}(\mathcal{A}^{\complement} \mid 0,x \in \mathrm{C}_\infty) &= \frac{\mathbb{P}(\mathcal{A}^{\complement} \cap \{0,x \in \mathrm{C}_\infty\} \mid 0,x \in \mathcal{V})}{\mathbb{P}(0,x \in \mathrm{C}_\infty \mid 0,x \in \mathcal{V})}\\ &\le \frac{q\rho}{\mathbb{P}(0,x \in \mathrm{C}_\infty \mid 0,x \in \mathcal{V})}.\end{aligned}$$

The result therefore follows from Lemma 6.5.

When $d = 1$, we construct $\pi_{0,x}$ in exactly the same way as below (6.35), using Lemma 6.7 with $M_{6.7} := \exp(\sqrt{(3/d)\log\overline{w}})$ (so that $r_{M_{6.7}} = \overline{w}^{3/d}$), in place of Lemma 6.6. We may assume $|x| \gg_\star \varepsilon, \zeta$ as for $d \ge 2$. Note that $M_{6.7} \le \exp(\sqrt{\log\log|x|})$ since $\overline{w} \le (\log|x|)^\varepsilon$ and $\varepsilon \ll_\star \mathtt{par}$, and in particular we may assume $\overline{w} \ge M_{6.7}^{8(\tau-1)}$ as in Lemma 6.8 since $\overline{w} \gg_\star \varepsilon, \mathtt{par}$. Using $|x| \gg_\star \varepsilon, \zeta$, this implies the costs of all our subpaths counted in (6.36) except $\pi_{y_0^\star,y_x^\star}$ are negligible compared to the $(\log|x|)^{\Delta_0+\varepsilon/2}$ cost of $\pi_{y_0^\star,y_x^\star}$, as in the $d \ge 2$ case, and likewise that the deviation of these subpaths from the line segment S_{0x} is negligible compared to $\zeta|x|$. The cost and deviation of $\pi_{y_0^\star,y_x^\star}$ are bounded using Corollary 5.2 exactly as in the $d \ge 2$ case in (6.36). □

Proof of Proposition 6.3. The proof when $\mu \in (\mu_{\log}, \mu_{\mathrm{pol}}]$ is identical to the proof of Proposition 6.2, except that the event $\mathcal{X}_{\mathrm{polylog}}(0,x)$ from Corollary 5.2 is replaced by $\mathcal{X}_{\mathrm{pol}}(0,x)$ from Corollary 5.3.

When $\mu > \mu_{\mathrm{pol}}$ and $d = 1$, we have $\eta_0 = 1$ and we can prove that the cost distance is at most $|x|^{1+\varepsilon}$ by using Lemma 6.7 directly as follows. We set $r_M = |x|$, which gives, using $r_M = \exp((\log M)^2)$, the value $M = \exp(\sqrt{\log|x|})$, which is slowly varying in $|x|$. Lemma 6.7 defines $\mathcal{H}_M := B_{2r_M} \cap \mathcal{V}_M$, and with $D_M = M^{2(\tau-1)+3\mu}$ and $\kappa_M = M^{3\mu+2}$, it states that

$$\begin{aligned}\mathbb{P}(\mathcal{A}_{\mathrm{near}}(\mathcal{H}_M, D_M, D_M, 0) &\cap \mathcal{A}_{\mathrm{near}}(\mathcal{H}_M, D_M, D_M, x)\\ &\cap \mathcal{A}_{\mathrm{linear}}(\mathcal{H}_M, \mathcal{H}_M, r_M, \kappa_M, 0, 2\kappa_M, 0) \mid 0,x \in \mathrm{C}_\infty) \ge 1 - 3q/10.\end{aligned}$$

The first two events $\mathcal{A}_{\mathrm{near}}(\mathcal{H}_M, D_M, D_M, z)$ with $z \in \{0,x\}$ guarantee that we find two paths $\pi_{0,y_0^\star}$ and $\pi_{x,y_x^\star}$ with cost at most D_M from 0 and x to respective vertices $y_0^\star, y_x^\star \in \mathcal{V}_M$, that are fully contained in $B_{D_M}(0)$ and $B_{D_M}(x)$, respectively. Here, $D_M = M^{2(\tau-1)+3\mu} \le |x|^{\varepsilon/2}$ for sufficiently large $|x|$. Then, since $y_0^\star$ is within distance $D_M < r_M$ from 0, the third event $\mathcal{A}_{\mathrm{linear}}$ guarantees a path between $y_0^\star$ and every vertex in $\mathcal{H}_M = \mathcal{V}_M \cap B_{2|x|}(0)$ with κ_M-linear cost, in particular there is such a path $\pi_{y_0^\star,y_x^\star}$ between $y_0^\star$ and $y_x^\star$. Let $\pi_{0,x} := \pi_{0,y_0^\star}\pi_{y_0^\star,y_x^\star}\pi_{y_x^\star,x}$ be the concatenation of these paths. Since the distance $|y_0^\star - y_x^\star| \le 2D_M + |x|$, the cost and deviation of this path is, using that $\kappa_M = M^{3\mu+2} \le |x|^{\varepsilon/2}$ for $|x|$ large,

$$\begin{aligned}\mathcal{C}(\pi_{0,x}) &= \mathcal{C}(\pi_{0,y_0^\star}) + \mathcal{C}(\pi_{y_x^\star,x}) + \mathcal{C}(\pi_{y_0^\star,y_x^\star}) \le 2D_M + \kappa_M|y_0^\star - y_x^\star| + 2\kappa_M\\ &\le 2|x|^{\varepsilon/2} + |x|^{\varepsilon/2}(|x| + 2|x|^{\varepsilon/2}) + 2|x|^{\varepsilon/2} \le |x|^{1+\varepsilon},\\ \mathrm{dev}(\pi_{0,x}) &\le \max\{\mathrm{dev}_{0,x}(\pi_{0,y_0^\star}), \mathrm{dev}_{0,x}(\pi_{y_0^\star,y_x^\star}), \mathrm{dev}_{0,x}(\pi_{y_x^\star},x)\}\\ &\le \max\{D_M, 0|y_0^\star - y_x^\star| + 2\kappa_M\} \le |x|^\varepsilon,\end{aligned}$$

for $|x|$ large enough. For small $|x|$ we can absorb the costs and deviation in the constant D. This proves the lemma when $d = 1$ and $\mu > \mu_{\mathrm{pol}}$ with $\eta_0 = 1$.

When $\mu > \mu_{\mathrm{pol}}$ and $d \ge 2$, a straightforward adaptation of the above proof for $d = 1$ could in principle also be used in higher dimensions. However, with some more effort one can get rid of the extra

$+\varepsilon$ in the exponent, and prove fully linear cost distances. We prove this stronger version (without the $|x|^{\varepsilon}$ factor) in [56, Theorems 1.8, 1.10]. □

A. Appendix

A.1. Concentration bounds

Theorem A.1. *(**Chernoff bounds (66, Theorems 4.4–4.5)**) Let $X_1, \dots X_k$ be independent Bernoulli distributed random variables, and define $X := \sum_{i=1}^{k} X_i$ and $m := \mathbb{E}[X]$. Then, for all $\lambda \in (0,1]$ and all $t \geq 2em$,*

$$\mathbb{P}(X \leqslant (1-\lambda)m) \leqslant e^{-m\lambda^2/2}, \qquad \mathbb{P}(X \geqslant (1+\lambda)m) \leqslant e^{-m\lambda^2/3}, \qquad \mathbb{P}(X \geq t) \leq 2^{-t}.$$

The same bounds hold when X is instead a Poisson variable with mean m.

A.2. The optimisation of total cost: proofs of Corollaries 5.2 and 5.3

Both corollaries follow from Proposition 5.1 with suitably chosen parameters. Throughout, we use the convention that $\infty \cdot 0 = 0$. In Proposition 5.1, the values of (γ, z, η, R) are not set yet (and they are not part of the model parameters $\mathtt{par}$). We will introduce constraints on these parameters below in Definition A.2 ('(K,A)-validity'), then we show in Lemma A.3 that a (K,A)-valid assignment of values in Proposition 5.1 yields a path between 0 and x of cost K with a multiplicative 'error' of at most A. Recall Λ, Φ and $\overline{w}$ from (4.1), (5.1) and (5.32), and that ξ is the side-length of the box Q in which the net exists.

Definition A.2 (Valid parameter choices). The *reduced Setting 4.1* is Setting 4.1, except without γ being defined. Consider the reduced Setting 4.1, and let $K, A > 0$. A setting of parameters (γ, z, η, R) is *(K,A)-valid* for ξ if the following conditions all hold for $\xi \gg_\star \mathtt{par}$, writing $\overline{w} := \xi^{\gamma^{R-1} d/2}$:

$$\gamma = \gamma(\mathtt{par}) \in (0,1), \quad z = z(\mathtt{par}) \in [0,d], \quad \eta = \eta(\mathtt{par}) \in [0,\infty), \tag{A.1}$$

$$R = R(\mathtt{par}, \xi) \in [2, (\log\log\xi)^2/4] \cap \mathbb{N}, \text{ with} \tag{A.2}$$

$$\overline{w}^2 \in [e^{(\log^{*3}\xi)^2}, A/\log\log\xi], \tag{A.3}$$

$$2^R \overline{w}^{4\mu} \xi^{\eta} \leq KA/\log\log\xi, \tag{A.4}$$

$$\Lambda(\eta, z) > 0 \qquad \text{and} \qquad \text{either } z = 0 \text{ or } \Phi(\eta, z) > 0. \tag{A.5}$$

Lemma A.3. *Consider the reduced Setting 4.1. Let $q, \zeta > 0$, let $0 < \delta \ll_\star q, \mathtt{par}$, and suppose that $\xi \gg_\star \delta, q, w_0, \zeta, \mathtt{par}$. Let $K, A > 0$, and suppose that (γ, z, η, R) is (K,A)-valid for ξ. Let $\mathcal{X}_{(K,A)}$ be the event that there is a path $\pi_{y_0^\star, y_x^\star}$ in G' with endpoints $y_0^\star$ and $y_x^\star$ satisfying*

$$w_{y_0^\star}, w_{y_x^\star} \in [\overline{w}, 4\overline{w}], \tag{A.6}$$

$$y_0^\star \in B_{\overline{w}^{3/d}}(0), \quad y_x^\star \in B_{\overline{w}^{3/d}}(x), \tag{A.7}$$

$$\mathcal{C}(\pi_{y_0^\star, y_x^\star}) \leq KA \textit{ and } \mathrm{dev}_{0x}(\pi) \leq \zeta|x|. \tag{A.8}$$

Then $\mathbb{P}(\mathcal{X}_{(K,A)} \mid V, w_V) \geq 1 - q$.

Proof. Let $y_0 := 0$, let $y_1 := x$, let $\xi := |x|$, and let $\theta := 1$. We first verify that the conditions of Proposition 5.1 hold. Since $\delta \ll_\star \mathtt{par}$, by (A.1) we may also assume $\delta \ll_\star \gamma, z, \eta$ as required by Proposition 5.1. Combined with (A.5) this implies that $\Lambda(\eta, z) \geq 2\sqrt{\delta}$ as required, and that either

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

$z = 0$ or $\Phi(\eta, z) \geq \sqrt{\delta}$ as required by Prop. 5.1. Since $\xi \gg_\star \delta, \mathtt{par}$, the inequalities $\xi \gg_\star \gamma, z, \eta, \delta, w_0$ and $\xi^{\gamma^{R-1}} \geq (\log\log\xi\sqrt{d})^{16d/\delta^2}$ by (A.3) and since $\overline{w} := \xi^{\gamma^{R-1}d/2}$, which is also required. Finally, $R \in [2, (\log\log\xi)^2]$ by (A.2) and $\gamma \in (0,1)$, $z \in [0,d]$ and $\eta \geq 0$ by (A.1).

Suppose that the event $\mathcal{X}_{\text{high-path}}$ of Proposition 5.1 occurs, and let π be a path as in the definition of $\mathcal{X}_{\text{high-path}}$. Then π satisfies (A.6) immediately, because $\mathcal{X}_{\text{high-path}}$ requires that the end-vertices of the path π have weights in $[\overline{w}, 4\overline{w}]$. The event also requires that the end-vertices are within distance $c_H\xi^{\gamma^{R-1}}$ from $0, x$ respectively. Since $\xi \gg_\star \mathtt{par}, \delta$, $c_H\xi^{\gamma^{R-1}} \leq \xi^{\gamma^{R-1}3/2} = \overline{w}^{3/d}$, and so π satisfies (A.7). The cost of π is at most $c_H 2^R\overline{w}^{4\mu}\xi^\eta$; by (A.4) combined with the fact that $\xi \gg_\star \mathtt{par}, \delta$, it follows that $\mathcal{C}(\pi) \leq KA$. The event $\mathcal{X}_{\text{high-path}}$ ensures that the deviation of π from the section $S_{0,x}$ is at most $3c_H\xi^\gamma$ where $\xi \gg_\star \zeta, \mathtt{par}, \delta$ and $\gamma < 1$, so (A.8) follows for any $\zeta > 0$ fixed. Thus

$$\mathbb{P}\big(\mathcal{X}_{(K,A)} \mid V, w_V\big) \geq \mathbb{P}\big(\mathcal{X}_{\text{high-path}} \mid V, w_V\big).$$

By Proposition 5.1 and the fact that $\xi \gg_\star q$, it follows that

$$\mathbb{P}\big(\mathcal{X}_{(K,A)} \mid V, w_V\big) \geq 1 - 2e^{-(\log\log\xi)^{13}} \geq 1 - q.$$

□

We shall now apply Lemma A.3 to prove Corollary 5.2 (which covers the polylogarithmic regime). Here, there are two possible choices of parameters (γ, z, η, R) for Lemma A.3: if $\alpha < 2$, then we are able to build a polylogarithmic-cost path using long-range edges between low-weight vertices (Claim A.5 below); if $\mu < \mu_{\log}$ then we are able to build a polylogarithmic-cost path using edges between high-weight vertices (Claim A.6 below). We then prove Corollary 5.2 by applying whichever parameter setting constructs a lower-cost path (in Corollary A.7). In both regimes, we need the following algebraic fact.

Claim A.4. Let

$$R = R(\xi) := \left\lceil \frac{\log\log\xi - (\log^{*4}\xi)^2}{\log\gamma^{-1}} \right\rceil. \tag{A.9}$$

Then for all $\gamma \in (1/2, 1)$ and $\xi \gg_\star \gamma$, it holds that

$$\xi^{\gamma^{R-1}} \in \big[e^{(\log^{*3}\xi)^2}, e^{\sqrt{\log\log\xi}}\big]. \tag{A.10}$$

Proof. The value of ξ is large, so using (A.9) and that $\lceil x \rceil \leq x + 1$,

$$\gamma^{R-1} \geq e^{-\log\log\xi + (\log^{*4}\xi)^2} = (\log^{*3}\xi)^{\log^{*4}\xi}/\log\xi \geq (\log^{*3}\xi)^2/\log\xi.$$

It follows that $\xi^{\gamma^{R-1}} \geq e^{(\log^{*3}\xi)^2}$, as required in (A.10). Moreover, since $\xi \gg_\star \gamma$, it holds that

$$\gamma^{R-1} \leq e^{(\log^{*4}\xi)^2}/(\gamma^2\log\xi) \leq e^{(\log^{*3}\xi)/2}/\log\xi = \sqrt{\log\log\xi}/\log\xi.$$

It follows that $\xi^{\gamma^{R-1}} \leq e^{\sqrt{\log\log\xi}}$, as required in (A.10). □

The next claim finds a (K, A)-valid parameter setting when $\alpha < 2$, for polylogarithmic cost-bound KA.

Claim A.5. Consider the reduced Setting 4.1, and fix $\varepsilon > 0$. When $\alpha < 2$, then writing $\Delta_\alpha := 1/(1 - \log_2\alpha)$, the following assignment is $((\log\xi)^{\Delta_\alpha}, (\log\xi)^\varepsilon)$-valid for $\xi \gg_\star \varepsilon, \mathtt{par}$ and $0 < \varepsilon' \ll_\star \varepsilon, \mathtt{par}$:

$$\gamma := \frac{\alpha}{2} + \varepsilon'; \qquad z := 0; \qquad \eta := 0; \qquad R := \left\lceil \frac{\log\log\xi - (\log^{*4}\xi)^2}{\log\gamma^{-1}} \right\rceil. \tag{A.11}$$

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

Proof. We check the requirements in Definition A.2 one-by-one. All the requirements of (A.1) and (A.2) are immediately satisfied except for $R \le (\log\log\xi)^2/4$, which follows from the definition since $\xi \gg_\star \gamma$ and $\gamma > 1/2$. Also since $\xi \gg_\star \gamma, \texttt{par}$, (A.3) follows from Claim A.4 since $\overline{w} = \xi^{\gamma^{R-1}d/2}$ and $A = (\log\xi)^\varepsilon$. We now prove (A.4). Since $\xi \gg_\star \gamma$, we estimate 2^R using R and γ in (A.11):

$$2^R \le 2^{\log\log\xi/\log\gamma^{-1}} = (\log\xi)^{\log 2/\log\gamma^{-1}} = (\log\xi)^{-\log 2/\log(\alpha/2+\varepsilon')}. \tag{A.12}$$

Since $\varepsilon' \ll_\star \varepsilon$, the exponent of $\log\xi$ on the right-hand side is

$$\frac{\log 2}{-\log(\alpha/2+\varepsilon')} \le \frac{\log 2}{\log(2/\alpha)} + \frac{\varepsilon}{2} = \frac{1}{1-\log_2\alpha} + \frac{\varepsilon}{2} = \Delta_\alpha + \frac{\varepsilon}{2}. \tag{A.13}$$

Moreover, since $\eta = 0$ in (A.11), the other factor $\overline{w}^{4\mu}\xi^\eta$ in (A.4) is at most (using Claim A.4 and $\xi \gg_\star \varepsilon$),

$$\overline{w}^{4\mu}\xi^\eta = \xi^{2\mu d\gamma^{R-1}} \le e^{2\mu d\sqrt{\log\log\xi}} \le (\log\xi)^{\varepsilon/2}/\log\log\xi. \tag{A.14}$$

Then (A.4) with $KA = (\log\xi)^{\Delta_\alpha+\varepsilon}$ follows from (A.12)–(A.14). We next prove (A.5). Using the formula in (4.1), with $z = 0$ and $\eta = 0$, and $\gamma = \alpha/2 + \varepsilon'$,

$$\Lambda(\eta, z) := 2d\gamma - \alpha(d-z) - z(\tau-1) + \big(0 \wedge \beta(\eta-\mu z)\big) = 2d(\alpha/2+\varepsilon') - \alpha d = 2d\varepsilon',$$

so $\Lambda(\eta, z) > 0$ as required. Since $z = 0$, (A.5) follows, so all criteria in Def. A.2 are satisfied. □

The next claim finds a (K, A)-valid parameter setting when $\mu < \mu_{\log}$, for polylogarithmic cost bound KA.

Claim A.6. Consider the reduced Setting 4.1, and fix $\varepsilon > 0$. When $\mu < \mu_{\log}$, then writing $\Delta_\beta := 1/(1-\log_2(\tau-1+\mu\beta))$, the following assignment is $((\log\xi)^{\Delta_\beta}, (\log\xi)^\varepsilon)$-valid for $\xi \gg_\star \varepsilon, \texttt{par}$ and $\varepsilon' \ll_\star \varepsilon, \texttt{par}$:

$$\gamma := \frac{\tau-1+\mu\beta}{2} + \varepsilon'; \qquad z := d; \qquad \eta := 0; \qquad R := \left\lceil \frac{\log\log\xi - (\log^{*4}\xi)^2}{\log\gamma^{-1}} \right\rceil. \tag{A.15}$$

Proof. First note that $\beta = \infty$ is not possible here, since in that case $\mu_{\log} = 0$, see (1.19). We check the requirements in Definition A.2 one-by-one. Since $\tau > 2$ and $\mu\beta \ge 0$, we obtain $\gamma > 1/2 > 0$ above, and since $\mu < \mu_{\log} = (3-\tau)/\beta$ and $\varepsilon' \ll_\star \texttt{par}$ it also holds that $\gamma < 1$; thus all the requirements of (A.1) are satisfied. It is also immediate that (A.2) is satisfied except for $R \le (\log\log\xi)^2/4$, which follows from the definition in (A.15) since $\xi \gg_\star \gamma$ and $\gamma > 1/2$. Since $\xi \gg_\star \gamma, \texttt{par}$, (A.3) follows from Claim A.4 since $\overline{w} = \xi^{\gamma^{R-1}d/2}$ and $A = (\log\xi)^\varepsilon$ as in the previous claim. We now prove (A.4). Analogously to (A.12):

$$2^R \le 2^{\log\log\xi/\log\gamma^{-1}} = (\log\xi)^{\log 2/\log(1/\gamma)}. \tag{A.16}$$

Since $\varepsilon' \ll_\star \varepsilon$, now γ is given in (A.15) and

$$\frac{\log 2}{\log(1/\gamma)} \le \frac{\log 2}{\log(2/(\tau-1+\mu\beta))} + \frac{\varepsilon}{2} = \frac{1}{1-\log_2(\tau-1+\mu\beta)} + \frac{\varepsilon}{2} = \Delta_\beta + \frac{\varepsilon}{2}. \tag{A.17}$$

Moreover, by Claim A.4 and since $\xi \gg_\star \varepsilon$, it holds that

$$\overline{w}^{4\mu}\xi^\eta = \xi^{2\mu d\gamma^{R-1}} \le e^{2\mu d\sqrt{\log\log\xi}} \le (\log\xi)^{\varepsilon/2}/\log\log\xi. \tag{A.18}$$

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

Now (A.4) follows immediately from (A.16)–(A.18). We next prove (A.5). Using the formula in (4.1), with $z = d$ and $\eta = 0$, and γ as in (A.15),

$$\Lambda(\eta, z) = 2d\gamma - \alpha(d - z) - z(\tau - 1) + \big(0 \wedge \beta(\eta - \mu z)\big) = 2d\gamma - d(\tau - 1) - d\mu\beta = 2d\varepsilon',$$

so $\Lambda(\eta, z) > 0$ as required. This also remains true for $\alpha = \infty$ both formally with $\alpha(d - z) = \infty \cdot 0 = 0$ as well as intuitively, since $z = d$ means we use edges which are present with constant probability. It remains to prove $\Phi(\eta, z) > 0$. Using the formula in (5.1), and that $\gamma \wedge 1/2 = 1/2$,

$$\Phi(\eta, z) = \Big[d\gamma \wedge \frac{z}{2}\Big] + \Big[0 \wedge \beta\Big(\eta - \frac{\mu z}{2}\Big)\Big] = d(\gamma \wedge 1/2) - \beta\mu d/2 = d(1 - \mu\beta)/2.$$

Since $\mu \le \mu_{\log} = (3 - \tau)/\beta$, it follows that $\Phi(\eta, z) \ge d(\tau - 2)/2$; since $\tau > 2$, (A.5) follows. □

Comparing the definition of Δ_0 in (1.9) to those in Claims A.5 and A.6, we recover here that

$$\Delta_0 = \frac{1}{1 - \log_2(\min\{\alpha, \tau - 1 + \mu\beta\})} = \min\{\Delta_\alpha, \Delta_\beta\}, \tag{A.19}$$

which formally remains true also when $\alpha = \infty$ or $\beta = \infty$ by (1.17), or (1.20). Combining the two claims we obtain the following corollary:

Corollary A.7. *Consider the reduced Setting 4.1, fix $\varepsilon > 0$. When either $\alpha \in (1, 2)$ or $\mu \in (\mu_{\mathrm{expl}}, \mu_{\log})$ or both hold, then there exists a setting of (γ, z, η, R) (depending on ε) which is $((\log\xi)^{\Delta_0}, (\log\xi)^{\varepsilon})$-valid for $\xi \gg_\star \varepsilon, \mathtt{par}$.*

Proof. Recall that $\mu_{\log} = (3 - \tau)/\beta$, so if $\mu_{\mathrm{expl}} < \mu < \mu_{\log}$ then $\beta < \infty$; thus we cannot have $\alpha = \beta = \infty$, and the formula (A.19) is valid whenever at least one of α, β is finite.

We show that when the minimum in the denominator is $\alpha \le \tau - 1 + \mu\beta$, (so that $\Delta_0 = \Delta_\alpha$), then also $\alpha < 2$ holds. Then, Claim A.5 directly gives a $((\log\xi)^{\Delta_\alpha}, (\log\xi)^{\varepsilon})$-valid parameter setting. There are two cases: either $\mu > \mu_{\log}$, then $\alpha < 2$ must hold by the hypothesis of the lemma; or $\mu < \mu_{\log} = (3 - \tau)/\beta$, so α being the minimum gives that $\alpha < \tau - 1 + \mu_{\log} \cdot \beta = 2$.

Similarly, we show that when the minimum in the denominator is $\tau - 1 + \mu\beta < \alpha$, (so that $\Delta_0 = \Delta_\beta$), then also $\mu < \mu_{\log}$ holds. Then, Claim A.6 directly gives a $((\log\xi)^{\Delta_\beta}, (\log\xi)^{\varepsilon})$-valid parameter setting. There are again two cases: either $\alpha \ge 2$, then $\mu < \mu_{\log}$ must hold by the hypothesis of the lemma; or $\alpha < 2$, so $\tau - 1 + \mu\beta$ being the minimum gives that $\tau - 1 + \mu\beta < 2$ and hence $\mu < (3 - \tau)/\beta = \mu_{\log}$. □

We are ready to prove Corollary 5.2 giving the polylogarithmic upper bound for the cost-distance

Proof of Corollary 5.2. Immediate from combining Lemma A.3 with Corollary A.7, where the required bounds on $\overline{w}$ in (5.4) follow from (A.3). □

We next apply Lemma A.3 to prove Corollary 5.3 that covers the polynomial regime. As with the proof of Corollary 5.2, we show that multiple possible choices of parameters are valid and choose the one which yields the lowest-cost path. We start with the case where $\alpha = \beta = \infty$. Recall the definition of η_0 from (1.10), (1.16), (1.19) and (1.21).

Claim A.8. Consider the reduced Setting 4.1, and fix $\varepsilon > 0$. When $\alpha = \beta = \infty$ and $\mu \in (\mu_{\log}, \mu_{\mathrm{pol}}]$, that is, $\eta_0 = 1 \wedge d\mu$ in (1.21), then the following setting is $(\xi^{\eta_0}, \xi^{\varepsilon})$-valid whenever $\varepsilon' \ll_\star \varepsilon, \mathtt{par}$ (with $1/(\varepsilon')^2$ an integer), and $\xi \gg_\star \varepsilon, \varepsilon', \mathtt{par}$:

$$\gamma := 1 - \varepsilon'; \qquad z := d; \qquad \eta := \eta_0 + \sqrt{\varepsilon'}; \qquad R := 1/(\varepsilon')^2. \tag{A.20}$$

Proof. Recall from (1.21) that when $\alpha = \beta = \infty$, the values $\mu_{\log} = 0$, $\mu_{\mathrm{pol}} = 1/d$. We check the requirements in Definition A.2 one-by-one. Both (A.1) and (A.2) are immediate. Since $\varepsilon' < 1/2$, it holds that $\gamma \in [e^{-2\varepsilon'}, e^{-\varepsilon'}]$; thus $\gamma^{R-1} \in [e^{-2/\varepsilon'}, e^{-1/(2\varepsilon')}]$ by the choice of R in (A.20). Since $\xi \gg_\star \varepsilon'$ and $\varepsilon' \ll_\star \varepsilon, \mathtt{par}$, it follows that $\xi^{\gamma^{R-1}d} \in [e^{(\log^{*3}\xi)^2}, \xi^{\varepsilon}/\log\log\xi]$ as required by (A.3). Moreover, using that $\overline{w} = \xi^{\gamma^{R-1}2/d}$ we estimate $2^R\overline{w}^{4\mu}\xi^{\eta} = 2^R\xi^{\eta+2\mu d\gamma^{R-1}} \le \xi^{\varepsilon/3} \cdot \xi^{\eta_0} \cdot \xi^{\varepsilon/3} \le \xi^{\eta_0+\varepsilon}/\log\log\xi$

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

and (A.4) holds. It remains to prove (A.5). Using the formula in (4.1) with γ, η, z as in (A.20), and that $\mu \le \mu_{\mathrm{pol}}$,

$$\begin{aligned}\Lambda(\eta, z) &= 2d\gamma - \alpha(d-z) - z(\tau-1) + \big(0 \wedge \beta(\eta - \mu z)\big)\\ &= 2d(1-\varepsilon') - \infty \cdot 0 - d(\tau-1) + (0 \wedge \infty) = d(3-\tau-2\varepsilon');\end{aligned}$$

since $\tau < 3$ and $\varepsilon' \ll_\star \mathtt{par}$, $\Lambda(\eta, z) > 0$ as required. Finally, using the formula in (5.1) and that $\gamma \wedge 1/2 = 1/2$, we analogously obtain that

$$\Phi(\eta, z) = \Big[d\gamma \wedge \frac{z}{2}\Big] + \Big[0 \wedge \beta\Big(\eta - \frac{\mu z}{2}\Big)\Big] = d(\gamma \wedge 1/2) + (0 \wedge \infty(d\mu/2 + \sqrt{\varepsilon'})) = d/2 > 0,$$

so (A.5) follows. Hence, all criteria in Def. A.2 are satisfied. □

When at least one of α, β is noninfinite, we can find two possible optimisers: one when $\mu < \mu_{\mu_{\mathrm{pol},\alpha}}$ and one when $\mu < \mu_{\mathrm{pol},\beta}$ hold in (1.8). We treat the two cases separately. Recall $\mu_{\mathrm{pol},\beta} = 1/d + (3-\tau)/\beta$ and let $\eta_\beta := d(\mu - \mu_{\log})$, the first term in the second row of (1.10).

Claim A.9. Consider the reduced Setting 4.1, and fix $\varepsilon > 0$. When $\alpha > 2$, $\mu \in (\mu_{\log}, \mu_{\mathrm{pol},\beta}]$, then the following setting is $(\xi^{\eta_\beta}, \xi^{\varepsilon})$-valid for $\varepsilon' \ll_\star \varepsilon, \mathtt{par}$ (with $1/(\varepsilon')^2$ an integer) and $\xi \gg_\star \varepsilon, \varepsilon', \mathtt{par}$:

$$\gamma := 1 - \varepsilon'; \qquad z := d; \qquad \eta := \eta_\beta + \sqrt{\varepsilon'}; \qquad R := 1/(\varepsilon')^2. \tag{A.21}$$

Proof. The $\alpha = \beta = \infty$ case was treated in Claim A.8 with (A.20) coinciding with (A.21). We treat the cases when at least one of α, β is finite. We check the requirements in Definition A.2 one-by-one. Both (A.1) and (A.2) are immediate. Since ε' is small we may choose it $\varepsilon' < 1/2$, implying that $\gamma \in [e^{-2\varepsilon'}, e^{-\varepsilon'}]$; thus $\gamma^{R-1} \in [e^{-2/\varepsilon'}, e^{-1/(2\varepsilon')}]$. Since $\xi \gg_\star \varepsilon'$ and $\varepsilon' \ll_\star \varepsilon, \mathtt{par}$, it follows that $\xi^{\gamma^{R-1}d} \in [e^{(\log^{*3}\xi)^2}, \xi^{\varepsilon}/\log\log\xi]$ as required by (A.3). Moreover, for (A.4) we use that $\overline{w} = \xi^{\gamma^{R-1}2/d}$ and estimate $2^R\overline{w}^{4\mu}\xi^{\eta} = 2^R\xi^{\eta + 2\mu d\gamma^{R-1}} \le \xi^{\varepsilon/3} \cdot \xi^{\eta_\beta} \cdot \xi^{\varepsilon/3} \le \xi^{\eta_\beta+\varepsilon}/\log\log\xi$ and so (A.4) holds. It remains to prove (A.5).

By their definition in (A.21), $z = d$ and $\eta = \eta_\beta + \sqrt{\varepsilon'}$ where $\eta_\beta = d(\mu - \mu_{\log}) = d(\mu - (3-\tau)/\beta)$, we compute $\eta - \mu z = \sqrt{\varepsilon'} - (3-\tau)d/\beta < 0$, since $\varepsilon' \ll_\star \mathtt{par}$. So, using the formula in (4.1) with γ, η, z as in (A.21),

$$\begin{aligned}\Lambda(\eta, z) &= 2d\gamma - \alpha(d-z) - z(\tau-1) + \big(0 \wedge \beta(\eta - \mu z)\big)\\ &= 2d(1-\varepsilon') - d(\tau-1) + \beta\sqrt{\varepsilon'} - (3-\tau)d = \beta\sqrt{\varepsilon'} - 2d\varepsilon'.\end{aligned}$$

Since $\varepsilon' \ll_\star \mathtt{par}$, it follows that $\Lambda(\eta, z) > 0$ as required by (A.5). This computation also remains valid both formally and intuitively when $\alpha = \infty$ and $\beta < \infty$, since $z = d$ and $\alpha(d-d) = 0$ reflects the fact that the edges we use appear with constant probability each. When $\alpha < \infty$ and $\beta = \infty$, $\mu_{\log} = 0$ and $\mu_{\mathrm{pol},\beta} = 1/d$, hence $\eta - \mu z = d\mu + \sqrt{\varepsilon'} - \mu d = \sqrt{\varepsilon'}$, so the minimum in $0 \wedge \beta(\eta - \mu z) = 0$. Hence when $\beta = \infty$, since $\gamma = 1 - \varepsilon'$ and $\tau - 1 < 2$,

$$\Lambda(\eta, z) = 2d\gamma - \alpha(d-d) - d(\tau-1) = d(2\gamma - (\tau-1)) > 0. \tag{A.22}$$

Finally we treat $\Phi(\eta, z) > 0$. When $\beta < \infty$, using the formula in (5.1) and that $\gamma \wedge 1/2 = 1/2$, with parameters in (A.21) and $\eta = \eta_\beta + \sqrt{\varepsilon'} = \mu d - \frac{(3-\tau)d}{\beta} + \sqrt{\varepsilon'}$, we analogously obtain that

$$\Phi(\eta, z) = \Big[d\gamma \wedge \frac{z}{2}\Big] + \Big[0 \wedge \beta\Big(\eta - \frac{\mu z}{2}\Big)\Big] = \frac{d}{2} + \Big[0 \wedge \beta\Big(\sqrt{\varepsilon'} + \frac{\mu d}{2} - \frac{(3-\tau)d}{\beta}\Big)\Big].$$

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press

In case the minimum on the right-hand side is at 0, $\Phi(\eta, z) > 0$ and so (A.5) is satisfied. In case the minimum is at the other term, we use that $\mu > \mu_{\log} = (3-\tau)/\beta$, so $\mu d/2 > (3-\tau)d/(2\beta)$, so

$$\Phi(\eta, z) \geq \frac{d}{2} - \beta \cdot \frac{(3-\tau)d}{2\beta} = \frac{(\tau-2)d}{2} > 0.$$

and so $\tau \in (2,3)$ ensures that (A.5) holds again. The computation remains valid when $\alpha = \infty$ since Φ does not depend on α. When $\alpha < \infty$ and $\beta = \infty$, the computation simplifies, and $\eta - \mu z/2 > 0$ holds since already $\eta - \mu z > 0$ see above (A.22). Hence in this case $\Phi(\eta, z) = d\gamma \wedge z/2 = d/2 > 0$. Hence, all criteria in Definition A.2 are satisfied with the choice in (A.21). □

The next claim finds minimisers whenever $\mu < \mu_{\mathrm{pol},\alpha}$. Recall that $\mu_{\mathrm{pol},\alpha} = \frac{\alpha-(\tau-1)}{d(\alpha-2)}$ from (1.8) and let $\eta_\alpha := \mu/\mu_{\mathrm{pol},\alpha}$, the second term in the second row of (1.10).

Claim A.10. Consider the reduced Setting 4.1, and fix $\varepsilon > 0$. When $\alpha > 2$, $\mu \in (\mu_{\log}, \mu_{\mathrm{pol},\alpha}]$, then the following setting is $(\xi^{\eta_\alpha}, \xi^{\varepsilon})$-valid for $\varepsilon' \ll_\star \varepsilon, \mathtt{par}$ (with $1/(\varepsilon')^2$ an integer), and $\xi \gg_\star \varepsilon, \varepsilon', \mathtt{par}$:

$$\gamma := 1 - \varepsilon'; \qquad z := (\eta_\alpha + \sqrt{\varepsilon'})/\mu; \qquad \eta := \eta_\alpha + \sqrt{\varepsilon'}; \qquad R := 1/(\varepsilon')^2. \tag{A.23}$$

Proof. We first show that $\alpha = \infty$, $\beta < \infty$ is not possible here. From (1.16) it follows that $\mu_{\mathrm{pol},\alpha} = 1/d$, while $\mu_{\log} = 1/d + (3-\tau)/\beta$, so for all $\beta > 0$ the strict inequality $\mu_{\log} > \mu_{\mathrm{pol},\alpha}$ holds and hence the interval for μ is empty when $\alpha = \infty$. Hence $\alpha < \infty$ is necessary for the conditions to be satisfied. We check the requirements of Definition A.2 one-by-one. Using the formula for $\mu_{\mathrm{pol},\alpha}$ and $\tau < 3$, we compute that $\eta_\alpha = \mu d(\alpha-2)/(\alpha-(\tau-1)) < \mu d$. Hence, since $\varepsilon' \ll_\star \mathtt{par}$ for all sufficiently small ε' the inequality $z \leq d$ holds as required by (A.1). The other conditions of (A.1) and (A.2) are immediate. Since γ and η is the same here and in Claim A.9, (A.3) and (A.4) hold by the same argument as in Claim A.9. It remains to prove (A.5). Using the formula in (4.1) with γ, η, z as in (A.23), which implies that $\eta - \mu z = 0$,

$$\begin{aligned}\Lambda(\eta, z) &= 2d\gamma - \alpha(d-z) - z(\tau-1) + \big(0 \wedge \beta(\eta - \mu z)\big)\\ &= d(2-\alpha) - 2\varepsilon' d + z(\alpha - (\tau-1)) + 0.\end{aligned} \tag{A.24}$$

This also remains valid both formally and intuitively when $\beta = \infty$ (with the convention that $\infty \cdot 0 = 0$), since $\eta - \mu z = 0$ reflects the fact that the random variable L on the edge we use is constant order. We substitute $z = (\eta_\alpha + \sqrt{\varepsilon'})/\mu$ from (A.23) and $\eta_\alpha = \mu d(\alpha-2)/(\alpha-(\tau-1))$:

$$\begin{aligned}\Lambda(\eta, z) &= d(2-\alpha) + \eta_\alpha(\alpha-(\tau-1))/\mu + \sqrt{\varepsilon'}(\alpha-(\tau-1))/\mu - 2\varepsilon' d\\ &= \sqrt{\varepsilon'}(\alpha-(\tau-1))/\mu - 2\varepsilon' d,\end{aligned}$$

since the first two terms in the first row cancelled each other. Since $\varepsilon' \ll_\star \mathtt{par}$, $\alpha > 2$ and $\tau \in (2,3)$, $\alpha - (\tau-1)$ is positive, and so is $\mu > 0$, so $\Lambda(\eta, z) > 0$ as required by (A.5). Finally, by (5.1) and since $z \leq d$,

$$\Phi(\eta, z) = \left[d\gamma \wedge \frac{z}{2}\right] + \left[0 \wedge \beta\left(\eta - \frac{\mu z}{2}\right)\right] = \frac{z}{2} + 0 > 0,$$

and so (A.5) holds. This also remains true for $\beta = \infty$ since the minimum is at 0, meaning we use edges with constant value L. Hence, all criteria in Definition A.2 are satisfied with the choice in (A.23). □

We are ready to prove Corollary 5.3 proving the upper bounds in the polynomial regime.

Proof of Corollary 5.3. Claim A.8 finds a setting of parameters that is $(|x|^{\eta_0}, |x|^{\varepsilon})$-valid whenever $\alpha = \beta = \infty$ and $\mu \leq \mu_{\mathrm{pol}} = 1/d$. When at least one of α, β is noninfinite, Claims A.9 and A.10 respectively find a setting of parameters that is $(|x|^{\eta_\beta}, |x|^{\varepsilon})$-valid whenever $\mu \leq \mu_{\mathrm{pol},\beta}$ and one that is $(|x|^{\eta_\alpha}, |x|^{\varepsilon})$-valid whenever $\mu \leq \mu_{\mathrm{pol},\alpha}$. By noting that $\eta_\beta \leq 1$ exactly when $\mu < \mu_{\mathrm{pol},\beta}$ and $\eta_\alpha \leq 1$

exactly when $\mu \leq \mu_{\text{pol},\alpha}$, we obtain that whenever $\mu \leq \max\{\mu_{\text{pol},\alpha}, \mu_{\text{pol},\beta}\}$, the two claims together find a parameter setting that is $(|x|^{\min\{\eta_\alpha,\eta_\beta\}}, |x|^{\varepsilon})$ valid. Since $\eta_0 = \min\{\eta_\alpha, \eta_\beta\}$ in (1.10), the proof from here is immediate by applying Lemma A.3, where the required bounds on $\overline{w}$ in (5.7) follow from (A.3). □

Acknowledgements. We thank Zylan Benjert for participating in writing the code for the simulations for the heatmaps in Figure 1.

Competing interests. The authors have no competing interest to declare.

Financial Support. Ulysse Schaller was supported by the Swiss National Science Foundation [grant number 200021_192079]. Open access funding provided by Delft University of Technology.

Ethical standards. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

References

[1] E. Adriaans and J. Komjáthy, 'Weighted distances in scale-free configuration models', *J. Stat. Phys.* **173** (2018), 1082–1109.
[2] W. Aiello, A. Bonato, C. Cooper, J. Janssen and P. Prałat, 'A spatial web graph model with local influence regions', *Internet Math.* **5**(1–2) (2008), 175–196.
[3] R. F. S. Andrade, J. S. AndradeJr and H. J. Herrmann, 'Ising model on the Apollonian network with node-dependent interactions', *Phys. Rev. E* **79**(3) (2009), 036105.
[4] P. Antal and A. Pisztora, 'On the chemical distance for supercritical Bernoulli percolation', *Ann. Probab.* **24**(2) (1996), 1036–1048.
[5] J. Bäumler, 'Distances in $1/\| x-y\|^{2d}$ percolation models for all dimensions', *Commun. Math. Phys.* **404**(3) (2023), 1495–1570.
[6] G. Baxter and G. Timár, 'Degree dependent transmission rates in epidemic processes', *J. Stat. Mech. Theory Exp.* 2021(10) (**2021**), 103501.
[7] Z. Benjert, J. Komjáthy, J. Lapinskas, J. Lengler and U. Schaller, 'Degree-dependent and distance-dependent contact rates interpolate between explosive, exponential and polynomial epidemic growth', In preparation, 2025+.
[8] N. Berger, 'A lower bound for the chemical distance in sparse long-range percolation models', 2004, arXiv preprint https://arxiv.org/abs/math/0409021.
[9] S. Bhamidi, R. van der Hofstad and G. Hooghiemstra, 'Universality for first passage percolation on sparse random graphs', *Ann. Probab.* **45**(4) (2017), 2568–2630.
[10] N. H. Bingham, C. M. Goldie and J. L. Teugels, *Regular Variation* (Cambridge University Press, Cambridge, 1987).
[11] M. Biskup, 'On the scaling of the chemical distance in long-range percolation models', *Ann. Probab.* **32**(4) (2004), 2938–2977.
[12] M. Biskup and J. Lin, 'Sharp asymptotic for the chemical distance in long-range percolation', *Random Struct. Algorithms* **55**(3) (2019), 560–583.
[13] T. Bläsius and P. Fischbeck, 'On the external validity of average-case analyses of graph algorithms', *ACM Trans. Algorithms* **20**(1) (2024), 1–42.
[14] T. Bläsius, T. Friedrich, M. Katzmann, U. Meyer, M. Penschuck and C. Weyand, 'Efficiently generating geometric inhomogeneous and hyperbolic random graphs', *Netw. Sci.* **10**(4) (2022), 361–380.
[15] B. Bollobás, S. Janson and O. Riordan, 'The phase transition in inhomogeneous random graphs', *Random Struct. Algorithms* **31**(1) (2007), 3–122.
[16] M. Bonaventura, V. Nicosia and V. Latora, 'Characteristic times of biased random walks on complex networks', *Phys. Rev. E* **89**(1) (2014), 012803.
[17] K. Bringmann, R. Keusch and J. Lengler, 'Geometric inhomogeneous random graphs', *Theor. Comput. Sci.* **760** (2019), 35–54.
[18] K. Bringmann, R. Keusch and J. Lengler, 'Average distance in a general class of scale-free networks', *Adv. Appl. Probab.* **760** (2024), 1–36.
[19] A. Caicedo and M. Dickson, 'Critical exponents for marked random connection models', *Electron. J. Probab.* **29** (2024), 1–57.
[20] S. Chatterjee and P. S. Dey, 'Multiple phase transitions in long-range first-passage percolation on square lattices', *Commun. Pure Appl. Math.* **69**(2) (2016), 203–256.
[21] G. Chowell, S. Bansal, L. Sattenspiel and C. Viboud, 'Mathematical models to characterize early epidemic growth: A review', *Phys. Life Rev.* **18** (2016), 66–97.
[22] F. Chung and L. Lu, 'Connected components in random graphs with given expected degree sequences', *Ann. Comb.* **6**(2) (2002), 125–145.
[23] C. Cooper, A. Frieze and P. Prałat, 'Some typical properties of the spatial preferred attachment model', *Internet Math.* **10**(1–2) (2014), 116–136.

[24] D. Coppersmith, D. Gamarnik and M. Sviridenko, 'The diameter of a long-range percolation graph', *Random Struct. Algorithms* **21**(1) (2002), 1–13.

[25] J. T. Cox and R. Durrett, 'Some limit theorems for percolation processes with necessary and sufficient conditions', *Ann. Probab.* **9**(4) (1981), 583–603.

[26] M. Deijfen, R. van der Hofstad and G. Hooghiemstra, 'Scale-free percolation', *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, **49**(3) (2013), 817–838.

[27] P. Deprez, R. S. Hazra and M. V. Wüthrich, 'Inhomogeneous long-range percolation for real-life network modeling', *Risks* **3**(1) (2015), 1–23.

[28] P. Deprez and M. V. Wüthrich, 'Scale-free percolation in continuum space', *Commun. Math. Stat.* **7**(3) (2019), 269–308.

[29] J.-D. Deuschel and A. Pisztora, 'Surface order large deviations for high-density percolation', *Probab. Theory Relat. Fields* **104** (1996), 467–482.

[30] C. Ding and K. Li, 'Centrality ranking in multiplex networks using topologically biased random walks', *Neurocomputing* **312** (2018), 263–275.

[31] J. Feldman and J. Janssen, 'High degree vertices and spread of infections in spatially modelled social networks', in *Proceedings of the 14th International Workshop on Algorithms and Models for the Web Graph* (Springer, 2017), 60–74.

[32] C. M. Fortuin, P. W. Kasteleyn and J. Ginibre, 'Correlation inequalities on some partially ordered sets', *Commun. Math. Phys.* **22** (1971), 89–103.

[33] N. Fountoulakis and T. Müller, 'Law of large numbers for the largest component in a hyperbolic model of complex networks', *Ann. Appl. Probab.* **28**(1) (2018), 607–650.

[34] C. Fransson, *Stochastic Epidemics on Random Networks and Competition in Growth*. PhD thesis, Stockholm University, Department of Mathematics, 2023. ISBN: 978-91-8014-145-1 (electronic).

[35] A. Frieze and M. Karoński, *Random Graphs and Networks: A First Course* (Cambridge University Press, Cambridge, 2023).

[36] A. Gandolfi, M. S. Keane and C. M. Newman, 'Uniqueness of the infinite component in a random graph with applications to percolation and spin glasses', *Probab. Theory Relat. Fields* **92** (1992), 511–527.

[37] C. Giuraniuc, J. Hatchett, J. Indekeu, M. Leone, I. P. Castillo, B. Van Schaeybroeck and C. Vanderzande, 'Criticality on networks with topology-dependent interactions', *Phys. Rev. E* **74**(3) (2006), 036108.

[38] P. Gracar, A. Grauer, L. Lüchtrath and P. Mörters, 'The age-dependent random connection model', *Queueing Syst.* **93**(3) (2019), 309–331.

[39] P. Gracar, A. Grauer and P. Mörters, 'Chemical distance in geometric random graphs with long edges and scale-free degree distribution', *Commun. Math. Phys.* **395**(2) (2022), 859–906.

[40] P. Gracar, M. Heydenreich, C. Mönch and P. Mörters, 'Recurrence versus transience for weight-dependent random connection models'. *Electron. J. Probab.* **27** (2022), 1–31.

[41] P. Gracar, L. Lüchtrath and C. Mönch, 'Finiteness of the percolation threshold for inhomogeneous long-range models in one dimension', *Electron. J. Probab.* **30** (2025), 1–29.

[42] P. Gracar, L. Lüchtrath and P. Mörters, 'Percolation phase transition in weight-dependent random connection models', *Adv. Appl. Probab.* **53**(4) (2021), 1090–1114.

[43] R. I. Grigorchuk, 'On growth in group theory', in *Proceedings of the International Congress of Mathematicians*, volume 1 (Mathematical Society of Japan, 1990), 325–338.

[44] M. Gromov, 'Groups of polynomial growth and expanding maps', *Publications Mathématiques de l'Institut des Hautes Études Scientifiques*, **53**(1) (1981), 53–78.

[45] L. Gugelmann, K. Panagiotou and U. Peter, Random hyperbolic graphs: degree sequence and clustering', in *Proceedings of the 39th International Colloquium on Automata, Languages, and Programming* (Springer, 2012), 573–585.

[46] N. Hao and M. Heydenreich, 'Graph distances in scale-free percolation: the logarithmic case', *J. Appl. Probab.* **60**(1) (2023), 295–313.

[47] C. Hirsch, 'From heavy-tailed Boolean models to scale-free Gilbert graphs', *Braz. J. Probab. Stat.* **31**(1) (2017), 111–143.

[48] R. v. d. Hofstad, P. van der Hoorn and N. Maitra, 'Scaling of the clustering function in spatial inhomogeneous random graphs', *J. Stat. Phys.* **190**(6) (2023), 110.

[49] H. Hooyberghs, B. Van Schaeybroeck, A. A. Moreira, J. S. AndradeJr, H. J. Herrmann and J. O. Indekeu, 'Biased percolation on scale-free networks', *Phys. Rev. E* **81**(1) (2010), 011102.

[50] E. Jacob and P. Mörters, 'Spatial preferential attachment networks: power laws and clustering coefficients', *Ann. Appl. Probab.* **25**(2) (2015), 632–663.

[51] J. Jorritsma and J. Komjáthy, 'Weighted distances in scale-free preferential attachment models', *Random Struct. Algorithms* **57**(3) (2020), 823–859.

[52] J. Jorritsma, J. Komjáthy and D. Mitsche, 'Cluster-size decay in supercritical kernel-based spatial random graphs', *Ann. Probab.* 2025+. Accepted for publication.

[53] M. Karsai, R. Juhász and F. Iglói, 'Nonequilibrium phase transitions and finite-size scaling in weighted scale-free networks', *Phys. Rev. E* **73**(3) (2006), 036116.

[54] R. Ke, C. Zitzmann, D. D. Ho, R. M. Ribeiro and A. S. Perelson, 'In vivo kinetics of SARS-CoV-2 infection and its relationship with a person's infectiousness', *Proc. Natl. Acad. Sci. USA* **118**(49) (2021), e2111477118.

[55] J. Komjáthy, J. Lapinskas and J. Lengler, 'Stopping explosion by penalising transmission to hubs in scale-free spatial random graphs', *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, **57**(4) (2021), 1968–2016.

[56] J. Komjáthy, J. Lapinskas, J. Lengler and U. Schaller, 'Polynomial growth in degree-dependent first passage percolation on spatial random graphs', *Electron. J. Probab.* **29** (2024), 1–48.
[57] J. Komjáthy and B. Lodewijks, 'Explosion in weighted hyperbolic random graphs and geometric inhomogeneous random graphs', *Stoch. Process. Their Appl.* **130**(3) (2020), 1309–1367.
[58] D. Krioukov, F. Papadopoulos, M. Kitsak, A. Vahdat and M. Boguná, 'Hyperbolic geometry of complex networks', *Phys. Rev. E* **82**(3) (2010), 036106.
[59] K. Kroy, 'Superspreading and heterogeneity in epidemics', in *Diffusive Spreading in Nature, Technology and Society* (Springer, 2023), 473–507.
[60] K. Lakis, J. Lengler, K. Petrova and L. Schiller, 'Improved bounds for graph distances in scale free percolation and related models', in *Proceedings of the Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 317 (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024), 74: 1–74:22.
[61] G. Last and M. Penrose, *Lectures on the Poisson Process* (Cambridge University Press, Cambridge, 2017).
[62] S. Lee, S.-H. Yook and Y. Kim, 'Centrality measure of complex networks using biased random walks', *Eur. Phys. J. B* **68**(2) (2009), 277–281.
[63] R. Michielan, N. Litvak and C. Stegehuis, 'Detecting hyperbolic geometry in networks: Why triangles are not enough', *Phys. Rev. E* **106**(5) (2022), 054303.
[64] R. Michielan and C. Stegehuis, 'Cliques in geometric inhomogeneous random graphs', *J. Complex Netw.* **10**(1) (2022), cnac002.
[65] G. Miritello, E. Moro, R. Lara, R. Martínez-López, J. Belchamber, S. G. Roberts and R. I. Dunbar, 'Time as a limited resource: Communication strategy in mobile phone networks', *Social Networks* **35**(1) (2013), 89–95.
[66] M. Mitzenmacher and E. Upfal, *Probability and Computing: Randomization and Probabilistic Techniques in Algorithms and Data Analysis* (Cambridge University Press, Cambridge, 2017).
[67] G. Ódor, D. Czifra, J. Komjáthy, L. Lovász and M. Karsai, 'Switchover phenomenon induced by epidemic seeding on geometric networks', *Proc. Natl. Acad. Sci. USA* **118**(41) (2021), e2112607118.
[68] M. Penrose, *Random Geometric Graphs* (Oxford University Press, Oxford, 2003).
[69] C. Pu, S. Li and J. Yang, 'Epidemic spreading driven by biased random walks', *Phys. A Stat. Mech. Appl.* **432** (2015), 230–239.
[70] R. Reubsaet, 'Topology-dependent first passage percolation on spatial scale-free networks', BSc thesis, Eindhoven University of Technology, 2021.
[71] K. Rohe, S. Chatterjee and B. Yu, 'Spectral clustering and the high-dimensional stochastic blockmodel', *Ann. Stat.* **39**(4) (2011), 1878–1915.
[72] L. S. Schulman, 'Long range percolation in one dimension', *J. Phys. A Math. Gen.* **16**(17) (1983), L639.
[73] T. Slangen, 'Cost-distances in penalized scale-free random graphs', BSc thesis, Eindhoven University of Technology, 2019.
[74] P. Trapman, 'The growth of the infinite long-range percolation cluster', *Ann. Probab.* **38**(4) (2010), 1583–1608.
[75] R. van der Hofstad and J. Komjáthy, 'Explosion and distances in scale-free percolation', arXiv preprint, 2017, arXiv:1706.02597.
[76] J. Wang, W. Xiong, R. Wang, S. Cai, D. Wu, W. Wang and X. Chen, 'Effects of the information-driven awareness on epidemic spreading on multiplex networks', *Chaos* **32**(7) (2022).
[77] R. Yang, T. Zhou, Y.-B. Xie, Y.-C. Lai and B.-H. Wang, 'Optimal contact process on complex networks', *Phys. Rev. E* **78**(6) (2008), 066109.
[78] V. Zlatić, A. Gabrielli and G. Caldarelli, 'Topologically biased random walk and community finding in networks', *Phys. Rev. E* **82**(6) (2010), 066109.

https://doi.org/10.1017/fms.2025.10161 Published online by Cambridge University Press