PHD

Random interacting particle systems

Gracar, Peter

Award date:
2018

Awarding institution:
University of Bath

Link to publication

Random interacting particle systems

submitted by

Peter Gracar

for the degree of Doctor of Philosophy

of the

University of Bath

Department of Mathematical Sciences

August 2017

Abstract

Consider the graph induced by $\mathbb{Z}^d$, equipped with *uniformly elliptic* random conductances on the edges. At time 0, place a Poisson point process of particles on $\mathbb{Z}^d$ and let them perform independent simple random walks with jump probabilities proportional to the conductances. It is well known that without conductances (i.e., all conductances equal to 1), an infection started from the origin and transmitted between particles that share a site spreads in all directions with positive speed. We show that a local mixing result holds for random conductance graphs and prove the existence of a special percolation structure called the Lipschitz surface. Using this structure, we show that in the setup of particles on a uniformly elliptic graph, an infection also spreads with positive speed in any direction. We prove the robustness of the framework by extending the result to infection with recovery, where we show positive speed and that the infection survives indefinitely with positive probability.

Contents

Chapter 1

Introduction

In this thesis, we consider the behavior of particle systems on a weighted graph and how this behavior results in the spread of infection. Let $G = (\mathbb{Z}^d, E)$ be the d-dimensional integer lattice with conductances $\{\mu_{x,y}\}_{(x,y)\in E}$, which are i.i.d. non-negative weights on the edges of G. We will always consider the weights to be symmetric, so $\mu_{x,y} = \mu_{y,x}$ for all $(x, y) \in E$. We also assume that the conductances are *uniformly elliptic*: that is,

$$\begin{aligned}&\text{there exists deterministic } C_M > 0\text{, such that}\\ &\mu_{x,y} \in [C_M^{-1}, C_M] \text{ for all } (x, y) \in E,\ \mathbb{P} - a.s.\end{aligned} \tag{1.1}$$

We note that our choice of name for the constant C_M is primarily in order to stay consistent with the work of Barlow and Hambly [3], upon which we rely to show several of our results and carries no deeper meaning. We say $x \sim y$ if $(x, y) \in E$ and define $\mu_x = \sum_{y\sim x} \mu_{x,y}$. At time 0, consider a Poisson point process of particles on $\mathbb{Z}^d$, with intensity measure $\lambda(x) = \lambda_0\mu_x$ for some constant $\lambda_0 > 0$ and all $x \in \mathbb{Z}^d$. That is, for each $x \in \mathbb{Z}^d$, the number of particles at x at time 0 is an independent Poisson random variable of mean $\lambda_0\mu_x$. Then, the particles perform independent continuous-time simple random walks (sometimes called CSRW) on the weighted graph; i.e., a particle at $x \in \mathbb{Z}^d$ jumps to a neighbor $y \sim x$ at rate $\frac{\mu_{x,y}}{\mu_x}$. It follows from the thinning property of Poisson random variables that the system of particles is in stationarity so that, at any time t, the particles are distributed according to a Poisson point process with intensity measure λ. We refer to this system of particles as *Poisson random walks* on (G, μ) with intensity λ_0. We now turn our attention to the infection and the mechanism with which it spreads. Assume that at time 0 there is an infected particle at the origin. The other particles are initially uninfected, but get infected instantaneously if they occupy the same site as an infected particle.

Assume for the moment that $\mu_{x,y} = 1$ for all $(x, y) \in E$. Then, the particles perform independent continuous time simple random walks with jump rate 1. More precisely, a particle at $x \in \mathbb{Z}^d$ jumps to a neighbor $y \sim x$ at rate $\frac{1}{2d}$. This is precisely the setup that was studied by Kesten and Sidoravicius [6]. They show that in this setup, although individual particles perform random walks with no drift, the infection experiences a drift towards any given *destination site* on the graph. In order to show this, they consider *lineages* of infected particles that are moving towards the destination site. More precisely, they first define a *distinguished* particle, starting with the initially infected particle. Then, whenever the distinguished particle shares a site with other particles (which are therefore infected as well) the probability that at least one of these particles jumps in any chosen direction is larger than $\frac{1}{2d}$. This gives that with probability greater than $\frac{1}{2}$, an infected particle from this shared site jumps to a site that is closer to the destination site. When this happens, the particle that performed the jump becomes the distinguished particle and is considered part of the lineage. In their proof, Kesten and Sidoravicius develop an intricate multi-scale argument to show that for any path across space and time, there are particles in the vicinity of the path sufficiently often, hence establishing the positive speed of spread of the infection by applying this to the path traveled by the distinguished particles.

A big challenge in this setup comes from the heavily dependent structure of the model. Though particles move independently of one another, dependences do arise over time. If the infection enters a region of the graph that is empty of particles, this will slow down the spread of the infection locally. Furthermore, this low density of particles decreases the density of particles in nearby areas as time goes on, further affecting the speed of the infection. In addition to this, the model also exhibits long-range dependences. For example, if a ball of radius R centered at some vertex x of the graph turns out to have no particles at time 0, then the ball $B(x, R/2)$ of radius $R/2$ centered at x will continue to be empty of particles up to time R^2, with positive probability. In particular, the probability that the $(d + 1)$-dimensional, space-time cylinder $B(x, R/2) \times [0, R^2]$ has no particle is at least $\exp\{-cR^d\}$ for some constant c. On the other hand, one expects that, after time $t \gg R^2$, the set of particles inside the ball will become "close" to stationarity.

To deal with this complex behavior of particles, multi-scale analysis is often used. In the case of spread of infection, this kind of analysis is built on the following ideas.

While it is true that as the infection spreads it enters empty regions of arbitrarily large size, it can be shown that the existence of such regions is increasingly unlikely for larger sizes. In fact, the same can be shown even for regions where the number of particles is simply smaller than its expectation (say, by a multiplicative factor of $1-\delta$, for some fixed $\delta > 0$). On the other hand, if one takes a region of the graph where the density of particles is "typical" in the sense that particles are close to stationarity, as time goes on the density of particles remains typical with high probability even at smaller scales. This can be shown with the following *local mixing* argument. Consider a cube $Q \subseteq \mathbb{Z}^d$, tessellated into subcubes of side length $\ell > 0$. Suppose that at some time t, the configuration of particles inside Q is dense enough, in the sense that inside each subcube there are at least $c\ell^d$ particles, for some constant $c > 0$. Regardless of how the particles are distributed inside Q, as long as the subcubes are dense, we obtain that for some constant $c' > 0$, at time $t + c'\ell^2$ the particles have not only had enough time to move out of the subcubes they were in at time t, but we also obtain that the configuration of particles inside "the core" of Q (i.e., away from the boundary of Q) stochastically dominates a Poisson point process of intensity $(1-\epsilon)c\ell^d$ that is independent of the configuration of particles at time t. Moreover, the value ϵ can be taken arbitrarily close to 0 by setting c' large enough. In words, we obtain a configuration at time $t + c'\ell^2$ inside the core of Q that is *roughly* independent of the configuration at time t, and is close to the stationary distribution. Applying a mixing result like this across multiple scales and controlling for the effects of empty regions of the graph forms the basis on which such multi-scale analysis arguments are built.

Although superficially different, the study of mobile nodes by Stauffer [13] exhibits similar behavior. Here, nodes are distributed as a Poisson point process on $\mathbb{R}^d$ and move according to independent Brownian motions. An individual node detects a target if it is within a set distance $r > 0$ of the node. The author considers the problem of a target that is initially at the origin avoiding detection and shows that even a clairvoyant target cannot avoid being detected indefinitely if the density of nodes is high enough. For this problem, Stauffer resorts to fractal percolation and multi-scale analysis in order to show the result.

Multi-scale constructions such as these are often used to deal with this kind of dependences, such as in the case of interlacements by Sznitman [14], in the case of activated random walks by Sidoravicius and Teixeira [11], and other problems in the interacting particle system outlined above in further works by Kesten and Sidoravicius [8], and for the case of Brownian motion on $\mathbb{R}^d$ by Peres et al. [9] and Stauffer [13]. The main downside of these approaches is that they are normally tailored specifically to the problem under consideration and cannot be easily adapted or extended to related problems. This thesis presents a different approach that improves on this. We define a general framework that is more applicable and robust than past results, allowing us to show several new results with only minor changes needed, despite considering a more general class of graphs.

1.1 Main results

We now return to the general case where the weights $\{\mu_{x,y}\}_{(x,y)\in E}$ satisfy (1.1). Our first result establishes that even when conductances are introduced, the infection on G still spreads with positive speed.

Theorem 1.1. *Let $\{\mu_{x,y}\}_{(x,y)\in E}$ be i.i.d. satisfying (1.1). For any time $t \geq 0$, let I_t be the position of the infected particle that is furthest away from the origin. Then*

$$\liminf_{t\to\infty} \frac{\|I_t\|_1}{t} > 0 \quad \textit{almost surely.}$$

As mentioned previously, if we set $\mu_{x,y} = 1$ for all $(x, y) \in E$, we recover the setup from [6] and Theorem 1.1 agrees with the result of Kesten and Sidoravicius. However, as outlined at the beginning of the chapter, their proof relies on the infection having a drift in any chosen direction whenever an infected particle shares a site with other particles. This is not robust to changes of $\mu_{x,y}$, since one can easily come up with configurations of weights for which even a large number of particles on a chosen site does not yield a drift in the desired direction. Our result on the other hand holds for all weight configurations that satisfy (1.1) and is therefore more robust.

As discussed at the beginning, to prove the above result we develop a framework that can be easily applied to a wider range of problems. To that end, we also use our technique to analyze the spread of infection with recovery. Let the setup be as before, but now each infected particle independently recovers and becomes uninfected at rate γ for some fixed parameter $\gamma > 0$. After recovering, a particle becomes again susceptible to the infection and gets infected again whenever it shares a site with an infected particle. The next result shows that if γ is small enough, then with positive probability there will be at least one infected particle on the graph at all times (in other words, the infection does not die out). Furthermore, when this happens we also obtain that the infection spreads with positive speed.

Theorem 1.2. *Let $\{\mu_{x,y}\}_{(x,y)\in E}$ be i.i.d. satisfying (1.1). For any $\lambda_0 > 0$, there exists $\gamma_0 > 0$ such that, for all $\gamma \in (0, \gamma_0)$, with positive probability, the infection does not die* *out. Furthermore, there are constants* $c_1, c_2, c_3 > 0$ *such that*

$$\mathbb{P}[\|I_t\|_1 \geq c_1 t \textit{ for all } t \geq c_3] \geq c_2,$$

where I_t *is the position of the infected particle that is furthest away from the origin at time* t.

The first half of the theorem with $\mu_{x,y} = 1$ for all $(x, y) \in E$ was solved by Kesten and Sidoravicius [7], where they showed that if $\gamma < \gamma_0$ for some critical positive value γ_0, then the infection survives with positive probability. Although the behavior of particles remains unchanged from the case without recovery, as does the graph geometry, the construction from [6] no longer works. Therefore, the authors had to do the multi-scale analysis from scratch, considering the implications of recovery throughout their argument. As a consequence of this change, they did not make any claim about the location of the infected particles at large times. Our result improves on this aspect, since it also reveals information about the location of the infection.

In order to prove Theorems 1.1 and 1.2, which we do in Chapter 2, two results need to be obtained. The first is a *local mixing result* which allows us to recover a certain amount of independence between sets of particles on G, provided that the sets are either located in regions of the graph that are far enough apart or that a sufficient amount of time has passed between the times at which we observe the two sets. We state this result precisely in Theorem 1.3 below. The second result we require and also the main contribution of the thesis is the existence of a very special percolation structure defined on the particle system. More precisely, we consider a tessellation of space and time and using the local mixing result show that there exists a Lipschitz connected percolation structure defined on this tessellation, such that for all tessellation tiles of this structure some *local* event holds. We state the precise conditions and properties for this structure in Theorems 1.4, 1.5 and 1.6.

1.2 Mixing of particles on a conductance graph

To deal with dependences of particle systems one often resorts to a decoupling argument, showing that two local events behave roughly independently of each other, provided they are measurable according to regions in space time that are sufficiently far apart. We obtain such an argument by extending a technique called *local mixing*; we gave an example of this technique for the case when the weights $\mu_{x,y}$ are all identical to 1 at the start of this chapter. To the best of our knowledge, the idea of local mixing in such settings originated in the work of Sinclair and Stauffer [12], and was later applied by Peres et al. [9] and Stauffer [13]. This idea was then extended with the introduction of *soft local times* by Popov and Teixeira [10], and applied to other processes, such as random interlacements.

One of the goals of this thesis is to show that this local mixing result can be obtained in a larger setting than the one considered previously. It is easy to see that the proof of local mixing in [13] goes through in cases where a local central limit theorem can be established with strong enough control over the error terms. However, we establish the local mixing for settings in which a local central limit theorem result might not hold or only holds in the limit as time goes to infinity, with no good control on the convergence rate. This is precisely the situation of our setting, when the weights $\mu_{x,y}$ are not all identical to 1. To work around that, we show that local mixing can be obtained whenever a so-called *parabolic Harnack inequality* holds, and we have some good estimates on the displacement of random walks. Key results that allow us to do this come from the work of Barlow [2], where heat kernel bounds are proven for percolation clusters on the square lattice, and Barlow and Hambly [3], where the result is extended to random conductance graphs and the parabolic Harnack inequality is shown to hold for such graphs as well.

In order to make the local mixing result as robust as possible, we further relax some of the assumptions on the weights of the graph (1.1). Let p_c be the critical probability for bond percolation on $\mathbb{Z}^d$. Assume that $\mu_{x,y}$ are i.i.d. and that, for each $(x, y) \in E$, we have

$$\mathbb{P}[\mu_{x,y} = 0] < p_c \text{ and } \mu_{x,y} \text{ satisfies (1.1) whenever } \mu_{x,y} > 0. \tag{1.2}$$

Theorem 1.3. *Let $\{\mu_{x,y}\}_{(x,y)\in E}$ be i.i.d. satisfying (1.2). There exist positive constants c_1, c_2, c_3, c_4, c_5 such that the following holds. Fix $K > \ell > 0$ and $\epsilon \in (0, 1)$. Consider a cube Q of side-length K, tessellated into subcubes $(T_i)_i$ of side length ℓ with $\frac{K}{\ell} \in \mathbb{Z}_+$. Assume each subcube T_i contains at least $\beta \sum_{x\in T_i} \mu_x$ particles for some $\beta > 0$, and let $\Delta \geq c_1\ell^2\epsilon^{-c_2}$. If ℓ is large enough, then after the particles move for time Δ, we obtain that within a region $Q' \subseteq Q$ that is at least $c_3\ell\epsilon^{-c_4}$ away from the boundary of Q, the particles dominate an independent Poisson point process of intensity measure $\nu(x) = (1 - \epsilon)\beta\mu_x$, $x \in Q'$, with probability at least*

$$1 - \sum_{y\in Q'} \exp\left\{-c_5\beta\mu_y\epsilon^2\Delta^{d/2}\right\}.$$

We prove this result in Chapter 2 as well as outline the class of graphs that this result can be extended to. In essence, a similar result holds for any *locally finite* graph equipped with conductances for which the parabolic Harnack inequality holds and such that the graph allows for suitable partitioning, similar to the tessellating of the cube into subcubes in Theorem 1.3.

1.3 Lipschitz surface and multi-scale percolation

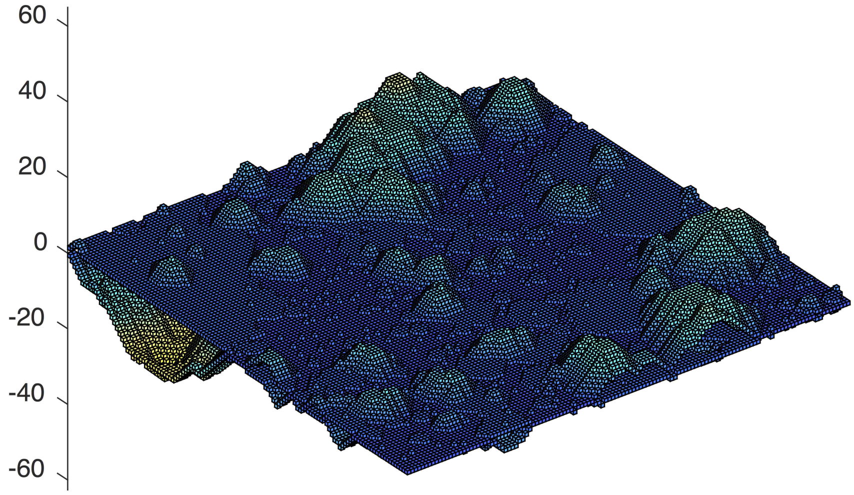

Figure 1: A two-sided Lipschitz surface for the case of $\mathbb{Z}^3$.

Consider again the case where $G = (\mathbb{Z}^d, E)$ and the weights satisfy (1.1). Although the upcoming construction does not explicitly require that $\mu_{x,y} \neq 0$ for all $(x, y) \in E$, not having this assumption introduces infinite range dependences which prevent the main object of interest, the *two-sided Lipschitz surface*, from existing. We discuss this issue in a dedicated section of Chapter 4 and proceed now to lay out the necessary steps in order to define this percolation structure.

We begin by tessellating the graph $G = (\mathbb{Z}^d, E)$ into d-dimensional cubes of side length $\ell > 0$. We index the cubes of the tessellation by integer vectors $i \in \mathbb{Z}^d$ such that the cube $i = (i_1, i_2, \dots, i_d)$ corresponds to the region $\left(\prod_{j=1}^d [i_j\ell, (i_j + 1)\ell]\right) \cap \mathbb{Z}^d$. Tessellate time into subintervals of length β. We index the subintervals by $\tau \in \mathbb{Z}$, representing the time interval $[\tau\beta, (\tau + 1)\beta]$. We refer to the pair (i, τ), representing $\prod_{j=1}^d [i_j\ell, (i_j + 1)\ell] \times [\tau\beta, (\tau + 1)\beta]$, as a *space-time cell.* We will need to consider larger, overlapping space-time cells as well. Let $\eta \geq 1$ be an integer. For each cube $i = (i_1, \dots, i_d)$ and time interval τ, define the *super cube* i as $\prod_{j=1}^d [(i_j - \eta)\ell, (i_j + \eta + 1)\ell]$ and the *super interval* τ as $[\tau\beta, (\tau + \eta)\beta]$. We define the *super cell* (i, τ) as the Cartesian product of the super cube i and the super interval τ.

We focus now on the particles and their behavior. We define a particle system on $\mathbb{Z}^d$ as a countable family of not necessarily unique elements of $\mathbb{Z}^d$, indexed by some countable set I, representing the locations of the particles belonging to the particle system. Let $(\Pi_s)_{s\geq 0}$ be a sequence of particle systems on $\mathbb{Z}^d$, with Π_s representing the locations of the particles at time s. We say a particle system Π_s is distributed according to a Poisson random measure of intensity ζ, if for every $B \subset \mathbb{Z}^d$, $N(B)$ is a Poisson random variable with intensity $\zeta(B)$, where $N(B)$ is the number of particles belonging to Π_s that lie in B. We say an event A is *increasing* for $(\Pi_s)_{s\geq 0}$ if the fact that A holds for $(\Pi_s)_{s\geq 0}$ implies that it holds for all $(\Pi'_s)_{s\geq 0}$ for which $\Pi'_s \supseteq \Pi_s$ for all $s \geq 0$. We say an event A is *restricted* to a region $X \subset \mathbb{Z}^d$ and a time interval $[t_0, t_1]$ if it is measurable with respect to the σ-field generated by all the particles that are inside X at time t_0 and their positions from time t_0 to t_1. We say a particle has displacement inside $X' \in \mathbb{Z}^d$ during a time interval $[t_0, t_1]$, if the location of the particle at all times during $[t_0, t_1]$ is inside $x + X'$, where x is the location of the particle at time t_0.

Let A be an event that is restricted to a region X and time interval $[0, t]$. Then, ν_A is called the *probability associated* to an increasing event A that is restricted to X and a time interval $[0, t]$ if, for an intensity measure ζ and a region $X' \in \mathbb{Z}^d$, $\nu_A(\zeta, X, X', t)$ is the probability that A happens given that, at time 0, the particles in X are a particle system distributed according to the Poisson random measure of intensity ζ and their motions from 0 to t are independent continuous time random walks on the weighted graph (G, μ), where the particles are conditioned to have displacement inside X' during $[0, t]$.

For each $(i, \tau) \in \mathbb{Z}^{d+1}$, let $E_{\text{st}}(i, \tau)$ be an increasing event restricted to the super cube i and the super interval τ. We also assume that $E_{\text{st}}(i, \tau)$ is invariant under space-time translations in the sense that $\mathbb{P}[E_{\text{st}}(i, \tau)] = \mathbb{P}[E_{\text{st}}(0, 0)]$ for all (i, τ) and write only E_{st} when the location of the cell does not play a role. We say that a cell (i, τ) is *good* if $E_{\text{st}}(i, \tau)$ holds and *bad* otherwise.

We will need a different way to index space-time cells, which we refer to as the *base-height index*. In the base-height index, we pick one of the d spatial dimensions and denote it as *height*, using index $h \in \mathbb{Z}$, while the remaining d space-time dimensions form the *base*, which will be indexed by $b \in \mathbb{Z}^d$. Then, a *base-height cell* will be indexed by $(b, h) \in \mathbb{Z}^{d+1}$. We analogously define the *base-height super cell* (b, h) to be the space-time super cell (i, τ), for which the base-height cell (b, h) corresponds to the space-time cell (i, τ). Similarly, we define $E_{\text{bh}}(b, h)$, the increasing event restricted to the super cell (b, h), to be the same as the event $E_{\text{st}}(i, \tau)$ for the space-time cell (i, τ) that corresponds to the base-height cell (b, h).

Finally, let a function $F : \mathbb{Z}^d \to \mathbb{Z}$ be called a *Lipschitz function* if $|F(x) - F(y)| \leq 1$

whenever $\|x-y\|_1 = 1$. With that, we are now ready to define the main object of this thesis, the *two-sided Lipschitz surface.*

Definition 1.1. A *two-sided Lipschitz surface* F is a set of base-height cells $(b,h) \in \mathbb{Z}^{d+1}$ such that for all $b \in \mathbb{Z}^d$ there are exactly two (possibly equal) integer values $F_+(b) \geq 0$ and $F_-(b) \leq 0$ for which $(b, F_+(b)), (b, F_-(b)) \in F$ and, moreover, F_+ and F_- are Lipschitz functions.

For an example of the two-sided Lipschitz surface when $d = 2$, see Figure 1. We say a space-time cell (i,τ) belongs to F if the corresponding base-height cell (b,h) belongs to F. We say a two-sided Lipschitz surface F *exists*, if for all $b \in \mathbb{Z}^d$, we have $F_+(b) < \infty$ and $F_-(b) > -\infty$. For any positive integer D, we say a two-sided Lipschitz surface *surrounds* a cell (b', h') at distance D if any path $(b', h') = (b_0, h_0), (b_1, h_1), \ldots, (b_n, h_n)$ for which $\|(b_i, h_i) - (b_{i-1}, h_{i-1})\|_1 = 1$ for all $i \in \{1, \ldots n\}$ and $\|(b_n, h_n) - (b_0, h_0)\|_1 > D$, intersects with F. For any $z \in \mathbb{Z}_+$, let $Q_z = [-z/2, z/2]^d$. We then have the following result, which establishes the existence of the two-sided Lipschitz surface.

Theorem 1.4. *Let (G, μ) be a weighted graph on the lattice $\mathbb{Z}^d$ for $d \geq 2$, satisfying (1.1). There exist positive constants c_0,c_1 and c_2 such that the following holds. Tessellate G in space-time cells and super cells as described above for some $\ell, \beta, \eta > 0$ such that the ratio $\beta/\ell^2 < c_0$. Let $E_{\mathrm{st}}(i,\tau)$ be an increasing event, restricted to the space-time super cell (i,τ). Fix $\epsilon \in (0,1)$ and fix w such that*

$$w \geq \sqrt{\frac{\eta\beta}{c_2 \ell^2} \log\left(\frac{8c_1}{\epsilon}\right)}.$$

Then, there exists a positive number α_0 that depends on ϵ, η and the ratio β/ℓ^2 so that if

$$\min\left\{C_M^{-1}\epsilon^2\lambda_0\ell^d, \log\left(\frac{1}{1-\nu_{E_{\mathrm{st}}}((1-\epsilon)\lambda, Q_{(2\eta+1)\ell}, Q_{w\ell}, \beta)}\right)\right\} \geq \alpha_0, \tag{1.3}$$

a two-sided Lipschitz surface F where $E_{\mathrm{st}}(i,\tau)$ holds for all $(i,\tau) \in F$ almost surely exists.

We now briefly explain the reasoning behind the conditions of the theorem. We normally fix β/ℓ^2 to be an arbitrary, but small constant. Next, the value of η defines the super cubes, which we use to get the required overlap between the cells of the tessellation (usually to allow information to propagate from one cell to its neighbors). Once these two parameters are fixed, we need to satisfy (1.3). First we need $C_M^{-1}\epsilon^2\lambda_0\ell^d \geq \alpha_0$. After fixing ϵ, this can be satisfied either by setting ℓ large enough (which makes the cells of the tessellation large), or by assuming that the density of particles λ_0 is large enough. Then we still need to make

$$\nu_{E_{\mathrm{st}}}((1-\epsilon)\lambda, Q_{(2\eta+1)\ell}, Q_{w\ell}, \beta) \geq 1 - \exp(-\alpha_0).$$

Usually E_{st} is a local event that becomes more and more likely by setting ℓ larger and larger or by increasing λ_0 to make the particle density higher; so having ℓ or λ_0 large enough suffices to satisfy this condition as well. The value of $\epsilon > 0$ is introduced so that in $\nu_{E_{\mathrm{st}}}$ we can consider a Poisson point process of particles of intensity measure $(1-\epsilon)\lambda$, slightly smaller than the actual intensity of particles. This slack is needed to restrict our attention to the particles that "behave well". Then the lower bound on w is to guarantee that, as particles move in $Q_{(2\eta+1)\ell}$ for time β, with high probability they stay in $Q_{(2\eta+1)\ell+w\ell}$, allowing a better control of dependences among neighboring cells of the tessellation.

In order to control the particles that "behave well", a multi-scale argument needs to be made. The basic idea of the argument is to use the local mixing result from Theorem 1.3 across multiple scales in order to control the probability of a cell not having a suitable density of particles. At some fixed large scale, Poisson bounds ensure that all cells of the relevant scale have a sufficient density of particles with high probability. Then, by using Theorem 1.3 repeatedly across the different scales, going from the largest to the smallest, one has at the smallest scale that the particles dominate a Poisson point process of intensity $(1-\epsilon)\lambda$ so that for a space-time cell (i,τ), the event $E_{\mathrm{st}}(i,\tau)$ occurs with probability $\nu_{E_{\mathrm{st}}}\big((1-\epsilon)\lambda, Q_{(2\eta+1)\ell}, Q_{w\ell}, \beta\big)$. For suitably defined events E_{st}, this probability can be made close to 1 by increasing λ_0 or ℓ. In our applications, setting ℓ large enough suffices, which allows us to derive the results on the spread of infection for *all* densities λ_0. However, in some applications, for example [13], one needs to set λ_0 large enough.

Theorem 1.4 gives the existence of the two-sided Lipschitz surface, which already enables arguments that require sequences of cells on which E_{st} holds and that are oriented in some way. More precisely, the surface allows one to choose sequences of cells where E_{st} holds, such that for any two consecutive cells (b_i, h_i) and $(b_{i'}, h_{i'})$, it holds that $(b_{i'} - b_i) = a$, where a is a d-dimensional integer vector of our choosing that satisfies $\|a\|_\infty = 1$. The surface however has further properties that can be shown. With additional work, we can establish the following.

Theorem 1.5. *Assume the conditions of Theorem 1.4 are satisfied. There exist positive* *constants c and C such that, for any sufficiently large $r > 0$, we have*

$$\mathbb{P}\left[\begin{array}{c}F \textit{ does not surround}\\ \textit{the origin at distance } r\end{array}\right] \leq \begin{cases}\sum_{s\geq r} s^d \exp\{-C\lambda_0 \frac{\ell s}{(\log \ell s)^c}\}, & \textit{for } d = 2\\ \sum_{s\geq r} s^d \exp\{-C\lambda_0 \ell s\}, & \textit{for } d \geq 3.\end{cases}$$

The proof of Theorem 1.5 also gives as a direct consequence that the parts of the two-sided Lipschitz surface where the two sides F_+ and F_- intersect not only almost surely separate the origin from infinity within the "zero-height hyperplane" $\mathbb{L} = \mathbb{Z}^d \times \{0\}$, but they even form an infinite percolation cluster within $\mathbb{L}$; we say that the two-sided Lipschitz surface percolates within $\mathbb{L}$ if the set $\mathbb{L} \setminus F$ contains only finite connected components.

Theorem 1.6. *Assume the conditions of Theorem 1.4 are satisfied. If in addition we have that ℓ is sufficiently large and $\mathbb{P}[E_{\mathrm{st}}(0,0)]$ is sufficiently large, then the zero-height cluster $F \cap \mathbb{L}$ of the two-sided Lipschitz surface F percolates within $\mathbb{L}$ almost surely.*

We now discuss the implications of Theorems 1.4, 1.5 and 1.6. First and foremost, the existence of the surface gives several new geometric properties of the infinite percolation cluster defined on increasing events as defined above. While the result of Stauffer [13] implies that the cluster of space-time cells where the increasing event E_{st} holds is infinite, no geometric properties of this percolation structure can be deduced. Stauffer shows that a clairvoyant target cannot remain undetected indefinitely by the nodes in the system by showing that any path, regardless of its direction in space and time, cannot forever avoid the cluster where E_{st} holds. One however cannot use this cluster to prove the existence of paths oriented in time, as that would require additional geometric properties of the cluster. Theorem 1.4 gives such a property. Within the cluster, there exists a two-sided Lipschitz surface which allows us to construct oriented paths in any direction within the base of the base-height index, such that E_{st} holds for all cells of the path. This is the main tool we use to prove Theorems 1.1 and 1.2, since it allows us to use an appropriately defined event E_{st} in such a way that it guarantees the infection moves from one cell to the next along a path.

A second property of the surface that we use to prove Theorems 1.1 and 1.2 is that the surface *surrounds* the origin at a finite distance. By Theorem 1.4 we already have that the surface has finite height everywhere almost surely, but this gives no control over the distance between the two sides of the surface. Theorem 1.5 improves on this by giving us precise control over how the two sides behave with respect to one another; it tells us that not only are the two sides of the surface finite as stated by Theorem 1.4, but they also coincide in such a way that the surface separates the origin from infinity at a finite distance. This guarantees that an infected particle is forced to enter the surface in finite time. Showing this is simple, since time is one of the $d+1$ space-time dimensions so the particle is guaranteed to move forward at least in that dimension and inevitably enter the two-sided Lipschitz surface.

Finally, although it is not needed to prove Theorems 1.1 and 1.2, Theorem 1.6 gives a further property of the two-sided Lipschitz surface. The two sides of the surface not only coincide in some sites almost surely, but the parts of the surface where they do percolate within $\mathbb{L}$. As with the surface itself, this is a geometric property of the infinite percolation cluster that is not implied directly by its existence.

Combined, these geometric properties of the Lipschitz surface allow us to analyze many different (global) questions pertaining to particle systems on uniformly elliptic graphs, without the need to carry out the very intricate multi-scale construction that is normally involved. Instead, we only have to come up with a local event E_{st} that combined with one or more of the geometric properties of the surface implies the desired (global) result.

1.4 Organization of the thesis

This thesis is written in the *alternative thesis format*, in accordance to *QA7 Appendix 6: Specifications for Higher Degree Theses and Portfolios*[1]. The main body of the thesis is composed of two papers, presented in Chapters 2 and 3. Both begin with a short contextualization of the corresponding paper, followed by the paper itself. Each paper is followed by closing remarks, highlighting how its results relate to the rest of the thesis.

In Chapter 2 we present the paper *Random walks in random conductances: decoupling and spread of infection*. In it, Theorem 1.3 is proved, after which Theorems 1.1 and 1.2 are proved. This is done by assuming Theorem 1.4 holds and showing that its conditions are satisfied for uniformly elliptic graphs and a particular increasing event E_{st}. In Chapter 3, the paper *Multi-scale Lipschitz percolation of increasing events for Poisson random walks* is presented. In it, we cover the construction of the two-sided Lipschitz surface and the multi-scale framework needed to establish Theorems 1.4, 1.5 and 1.6. In the final chapter the thesis concludes with a discussion of some open problems of the two-sided Lipschitz surface and some potential future areas to which the framework could be extended.

[1]Available from `http://www.bath.ac.uk/quality/documents/QA7-Appendix-6.pdf`, as of August 2017.

Chapter 2

Random walks in random conductances: decoupling and spread of infection

Introduction to the paper

In this chapter, we present the paper *Random walks in random conductances: decoupling and spread of infection*. As outlined in Chapter 1, the paper presents two important results. In Section 1 of the paper, we start with the setup of *uniformly elliptic* random conductance graphs. Next, in Section 2 several known results due to Barlow [2] and Barlow and Hambly [3] pertaining to random walks on uniformly elliptic graphs are presented. Chief among these are *heat kernel bounds*, which we use in place of a local central limit theorem. A second important result by Barlow and Hambly [3] that we use is the *parabolic Harnack inequality*, which the authors show holds for heat kernels on uniformly elliptic graphs. As a final related result we present an exit time statement for random walks on uniformly elliptic graphs from the same paper. These results allow us to prove a *local mixing* result for uniformly elliptic graphs, which we do in Section 3. More precisely, we prove Theorem 1.3 from Chapter 1. In Section 4, we prove an extension of this theorem to conditioned random walks. The key issue with these walks is that their heat kernels no longer satisfy the parabolic Harnack inequality, so an additional step is required to show that their heat kernels can be approximated by heat kernels that do satisfy the parabolic Harnack inequality. The first part of the paper concludes with further extensions of the result to locally finite graphs for which the results of Barlow and Hambly hold.

The second part of the paper, covered by Section 5, focuses on the spread of infection and applications of Theorem 1.4, which for the time being we take as given[1]. We use it with an appropriate increasing event $E_{\rm st}$ to show that the infection spreads with positive speed. Intuitively, we define $E_{\rm st}(i,\tau)$ to be as follows. Given that there is an infected particle inside the cube i at the beginning of the time interval τ, we want the infection to spread to all the neighboring cubes by the end of the time interval τ. To achieve this, we first show that in a short amount of time the infected particle spreads the infection to a large number of other particles that are located in the super-cube i at the beginning of τ. Then, these particles have enough time left in the interval τ to move into the neighboring cubes. Therefore, at the beginning of the time interval $\tau+1$, all of the neighboring cubes contain an infected particle. We conclude the paper by using the event $E_{\rm st}(i,\tau)$ and Theorem 1.4 to prove Theorem 1.1 and with only minor modifications to the event needed, also Theorem 1.2.

[1]The proof of the theorem can be found in Chapter 3.

UNIVERSITY OF BATH

Statement of Authorship (to preface each co-authored paper)

This declaration concerns the article entitled:
Random walks in random conductances: decoupling and spread of infection

Publication status (tick one)									
Draft manuscript		**Submitted**	X	**In review**		**Accepted**		**Published**	

Publication details (reference)	Stochastic Processes and their Applications

Candidate's contribution to the paper (detailed, and also given as a percentage).

The candidate was the primary contributor to:

- Formulation of ideas:
  - Use of soft local times to prove local mixing and the use of the parabolic Harnack inequality. 75%
  - Definition of the increasing event to prove spread of infections. 75%
- Design of methodology:
  - Implemented proofs of results, including all details. 100%

Statement from Candidate	This paper reports on original research I conducted during the period of my Higher Degree by Research candidature.		
Signed		**Date**	15.08.2017

Random walks in random conductances: decoupling and spread of infection

Peter Gracar* and Alexandre Stauffer†

Summary

Let (G,μ) be a *uniformly elliptic* random conductance graph on $\mathbb{Z}^d$ with a Poisson point process of particles at time $t=0$ that perform independent simple random walks. We show that inside a cube Q_K of side length K, if all subcubes of side length $\ell < K$ inside Q_K have sufficiently many particles, the particles return to stationarity after $c\ell^2$ time with a probability close to 1. We also show this result for percolation clusters on locally finite graphs. Using this mixing result and the results of [6], we show that in this setup, an infection spreads with positive speed in any direction. Our framework is robust enough to allow us to also extend the result to infection with recovery, where we show positive speed and that the infection survives indefinitely with positive probability.

Keywords and phrases: mixing, decoupling, spread of infection, heat kernel

1 Introduction

We consider the graph $G = (\mathbb{Z}^d, E)$, $d \geq 2$ to be the d-dimensional integer lattice, with edges between nearest neighbors: for $x, y \in \mathbb{Z}^d$ we have $(x,y) \in E$ iff $\|x-y\|_1 = 1$. Let $\{\mu_{x,y}\}_{(x,y)\in E}$ be a collection of i.i.d. non-negative weights, which we call *conductances*. In this paper, conductances will always be symmetric, so $\mu_{x,y} = \mu_{y,x}$ for all $(x,y) \in E$. We also assume that the conductances are uniformly elliptic: that is,

$$\begin{aligned} &\text{there exists deterministic } C_M > 0, \text{ such that} \\ &\mu_{x,y} \in [C_M^{-1}, C_M] \text{ for all } (x,y) \in E,\ \mathbb{P}-a.s. \end{aligned} \tag{1.1}$$

*Department of Mathematical Sciences, University of Bath, UK. p.gracar@bath.ac.uk

†Department of Mathematical Sciences, University of Bath, UK, a.stauffer@bath.ac.uk. Supported by a Marie Curie Career Integration Grant PCIG13-GA-2013-618588 DSRELIS, and an EPSRC Early Career Fellowship.

We say $x \sim y$ if $(x, y) \in E$ and define $\mu_x = \sum_{y \sim x} \mu_{x,y}$. At time 0, consider a Poisson point process of particles on $\mathbb{Z}^d$, with intensity measure $\lambda(x) = \lambda_0 \mu_x$ for some constant $\lambda_0 > 0$ and all $x \in \mathbb{Z}^d$. That is, for each $x \in \mathbb{Z}^d$, the number of particles at x at time 0 is an independent Poisson random variable of mean $\lambda_0 \mu_x$. Then, let the particles perform independent continuous-time simple random walks (CSRW) on the weighted graph so that a particle at $x \in \mathbb{Z}^d$ jumps to a neighbor $y \sim x$ at rate $\frac{\mu_{x,y}}{\mu_x}$. It follows from the thinning property of Poisson random variables that the system of particles is in stationarity; thus, at any time t, the particles are distributed according to a Poisson point process with intensity measure λ.

We study the spread of an infection among the particles. Assume that at time 0 there is at least one particle at the origin, all particles at the origin are infected, and all other particles are uninfected. Then an uninfected particle gets infected as soon as it shares a site with an infected particle. Our first result establishes that the infection spreads with positive speed.

Theorem 1.1. *Let $\{\mu_{x,y}\}_{(x,y)\in E}$ be i.i.d. satisfying (1.1). For any time $t \geq 0$, let I_t be the position of the infected particle that is furthest away from the origin. Then*

$$\liminf_{t\to\infty} \frac{\|I_t\|_1}{t} > 0 \quad \textit{almost surely.}$$

The above result has been established on the square lattice (i.e., $\mu_{x,y} = 1$ for all $(x, y) \in E$) by Kesten and Sidoravicius [8] via an intricate multi-scale analysis; see also [9] for a shape theorem. In a companion paper [6], we develop a framework which can be used to analyze processes in this setting without the need of carrying out a multi-scale analysis from scratch. We prove our Theorem 1.1 via this framework, showing the applicability of our technique from [6]. We also apply this technique to analyze the spread of an infection with recovery. Let the setup be as before, but now each infected particle independently recovers and becomes uninfected at rate γ for some fixed parameter $\gamma > 0$. After recovering, a particle becomes again susceptible to the infection and gets infected again whenever it shares a site with an infected particle. Our next result shows that if γ is small enough, then with positive probability there will be at least one infected particle at all times. When this happens, we also obtain that the infection spreads with positive speed.

Theorem 1.2. *Let $\{\mu_{x,y}\}_{(x,y)\in E}$ be i.i.d. satisfying (1.1). For any $\lambda_0 > 0$, there exists $\gamma_0 > 0$ such that, for all $\gamma \in (0, \gamma_0)$, with positive probability, the infection does not die out. Furthermore, there are constants $c_1, c_2, c_3 > 0$ such that*

$$\mathbb{P}[\|I_t\|_1 \geq c_1 t \textit{ for all } t \geq c_3] \geq c_2,$$

where I_t is the position of the infected particle that is furthest away from the origin at time t.

The challenge in this setup comes from the heavily dependent structure of the model. Though particles move independently of one another, dependencies do arise over time. For example, if a ball of radius R centered at some vertex x of the graph turns out to have no particles at time 0, then the ball $B(x, R/2)$ of radius $R/2$ centered at x, will continue to be empty of particles up to time R^2, with positive probability. In particular, that the probability that the $(d+1)$-dimensional, space-time cylinder $B(x, R/2) \times [0, R^2]$ has no particle is at least $\exp\{-cR^d\}$ for some constant c, which is just a stretched exponential in the volume of the cylinder. On the other hand, one expects that, after time $t \gg R^2$, the set of particles inside the ball will become "close" to stationarity.

To deal with dependences, one often resorts to a decoupling argument, showing that two local events behave roughly independently of each other, provided they are measurable according to regions in space time that are sufficiently far apart. We will obtain such an argument by extending a technique which we call *local mixing*, and which was introduced in [12]. The key observation is the following. Consider a cube $Q \subseteq \mathbb{Z}^d$, tessellated into subcubes of side length $\ell > 0$. For simplicity assume for the moment that $\mu_{x,y} = 1$ for all $(x, y) \in E$. Suppose that at some time t, the configuration of particles inside Q is dense enough, in the sense that inside each subcube there are at least $c\ell^d$ particles, for some constant $c > 0$. Regardless of how the particles are distributed inside Q, as long as the subcubes are dense, we obtain that at some time $t + c'\ell^2$, not only particles had enough time to move out of the subcubes they were in at time t, but also we obtain that the configuration of particles inside "the core" of Q (i.e., away from the boundary of Q) stochastically dominates a Poisson point process of intensity $(1-\epsilon)c\ell^d$ that is independent of the configuration of particles at time t. Moreover, the value ϵ can be made arbitrarily close to 0 by setting c' large enough. In words, we obtain a configuration at time $t + c'\ell^2$ inside the core of Q that is roughly independent of the configuration at time t, and is close to the stationary distribution. To the best of our knowledge, the idea of local mixing in such settings originated in the work of Sinclair and Stauffer [12], and was later applied in [10, 13]. This idea was then extended with the introduction of soft local times by Popov and Teixeira [11] (see also [7]), and applied to other processes, such as random interlacements.

Our second main goal in this paper is to show that this local mixing result can be obtained in a larger setting, in which a local CLT, which plays a crucial role in the proof[1] of [12, 10, 7], might not hold or only holds in the limit as time goes to infinity, with no good control on the convergence rate. This is precisely the situation in our setting, where the weights $\mu_{x,y}$ are not all identical to 1. To work around that, we will show that local mixing can be obtained whenever a so-called *Parabolic Harnack Inequality* holds, and we have some good estimates on the displacement of random walks.

For the result below, we can impose slightly weaker conditions on $\mu_{x,y}$. Let p_c be the critical probability for bond percolation on $\mathbb{Z}^d$. Assume that $\mu_{x,y}$ are i.i.d. and that, for each $(x,y) \in E$, we have

$$\mathbb{P}[\mu_{x,y} = 0] < p_c \text{ and } \mu_{x,y} \text{ satisfies (1.1) whenever } \mu_{x,y} > 0. \tag{1.2}$$

For two regions $Q' \subseteq Q \subset \mathbb{Z}^d$, we say that Q' is x away from the boundary of Q if the distance between Q' and Q^c is at least x.

Theorem 1.3. *Let $\{\mu_{x,y}\}_{(x,y)\in E}$ be i.i.d. satisfying (1.2). There exist positive constants c_1, c_2, c_3, c_4, c_5 such that the following holds. Fix $K > \ell > 0$ and $\epsilon \in (0,1)$. Consider a cube Q of side-length K, tessellated into subcubes $(T_i)_i$ of side length ℓ. Assume each subcube T_i contains at least $\beta \sum_{x \in T_i} \mu_x$ particles for some $\beta > 0$, and let $\Delta \geq c_1 \ell^2 \epsilon^{-c_2}$. If ℓ is large enough, then after the particles move for time Δ, we obtain that within a region $Q' \subseteq Q$ that is at least $c_3 \ell \epsilon^{-c_4}$ away from the boundary of Q, the particles dominate an independent Poisson point process of intensity measure $\nu(x) = (1-\epsilon)\beta\mu_x$, $x \in Q'$, with probability at least*

$$1 - \sum_{y \in Q'} \exp\left\{-c_5 \beta \mu_y \epsilon^2 \Delta^{d/2}\right\}.$$

We will prove a more detailed version of this theorem in Section 3 (see Theorem 3.1). Although we only prove the result for the case of conductances on the square lattice, Theorem 1.3 holds for more general graphs. The theorem holds for any graph G and any region Q of G that can be tessellated into subregions of diameter at most ℓ whenever the particles in each such subregion are dense enough, the so-called parabolic Harnack inequality holds for G and we have estimates on the displacement of random walks on G. We discuss some extensions in Section 4.

The structure of this paper is as follows. In Section 2, we formally define the family of graphs we consider for local mixing and present results concerning the parabolic Harnack inequality, heat kernel bounds and exit times for random walks on such graphs.

[1] The results of [12, 10] are in the setting of Brownian motions on $\mathbb{R}^d$, but can be adapted in a straightforward way to random walks on $\mathbb{Z}^d$ with $\mu_{x,y} = 1$ for all $(x,y) \in E$ by using the local CLT.

In Section 3, we state a more precise version of Theorem 1.3 and prove it. In Section 4 we prove an extension of the local mixing result to random walks whose displacement is conditioned to be bounded, which is particularly useful in applications [12, 6]. In Section 5, we use the local mixing result and results from our companion paper [6] to prove Theorems 1.1 and 1.2 for graphs satisfying (1.1).

2 Heat kernel estimates and exit times

In this section, we consider a general graph $G = (V, E)$, with uniformly bounded degrees. For $x, y \in V$, let $|x - y|$ denote the graph distance between x and y in G. In order to avoid potentially confusing notation, we allow ourselves a slight abuse of notation and also use $|x - y|$ to denote the graph distance when dealing with non-Euclidian graphs. For $x \in V$, let $B(x, r) = \{y \in V : |x - y| \leq r\}$ be the ball of radius r centered at x. We consider non-negative weights (conductances) $(\mu_{x,y})_{(x,y)\in E}$, that are symmetric. As in Section 1, we denote by $x \sim y$ whenever $x, y \in V$ are neighbors in G, and define $\mu_x = \sum_{y\sim x} \mu_{x,y}$. We also extend μ to a measure on V. The reader may think of V as $\mathbb{Z}^d$ and $\mu_{x,y}$ being i.i.d. random variables satisfying (1.2). We keep our notation in greater generality as we want to highlight the exact conditions we need for our results.

Assume the existence of $d \geq 1$ and C_U such that

$$\mu(B(x, r)) \leq C_U r^d, \quad \text{for all } r \geq 1, \text{ and } x \in V. \tag{2.1}$$

We consider a continuous time simple random walk on the weighted graph $\mathcal{G} := (G, \mu)$, which jumps from vertex x to vertex y at rate $\frac{\mu_{x,y}}{\mu_x}$. More formally, for any function $f : V \to \mathbb{R}$, let

$$\mathcal{L}f(x) = \mu_x^{-1} \sum_{y\sim x} \mu_{x,y}(f(y) - f(x)), \tag{2.2}$$

and define the random walk started at vertex x as the Markov process $Y = (Y_t, t \in [0, \infty), \mathbb{P}_x, x \in V)$ with generator $\mathcal{L}$. Its *heat kernel* on the graph is defined as

$$q_t(x, y) = \frac{\mathbb{P}_x(Y_t = y)}{\mu_y}, \text{ for any } x, y \in V. \tag{2.3}$$

We will say that a particle walks along $\mathcal{G}$ if it is a Markov process with generator $\mathcal{L}$ as defined above. We now state several definitions from [2] which we use throughout the paper.

Definition 2.1 (Very good balls)**.** Let C_V, C_P and $C_W \geq 1$ be fixed constants. We say $B(x, r)$ is $(C_V, C_P, C_W) - $*good* if:

$$\mu(B(x,r)) \geq C_V r^d,$$

and the weak Poincaré inequality

$$\sum_{y \in B(x,r)} (f(y) - \bar{f}_{B(x,r)})^2 \mu_y \leq C_P r^2 \sum_{y,z \in B(x, C_W r), z \sim y} (f(y) - f(z))^2 \mu_{yz}$$

holds for every $f : B(x, C_W r) \to \mathbb{R}$, where $\bar{f}_{B(x,r)} = \mu(B(x,r))^{-1} \sum_{y \in B(x,r)} f(y)\mu_y$ is the weighted average of f in $B(x, r)$. Furthermore, we say $B(x, R)$ is $(C_V, C_P, C_W) -$ *very good* if there exists $N_B = N_{B(x,R)} \leq R^{1/(d+2)}$ such that for all $r \geq N_B$, $B(y, r)$ is good whenever $B(y, r) \subseteq B(x, R)$. We assume that $N_B \geq 1$.

For the remainder of the paper we assume that $d \geq 2$, fix C_U, C_V, C_P and C_W and take $\mathcal{G} = (V, E, \mu)$ to satisfy (2.1).

We are now ready to present some key results from [3] that control the variation of the random walk density function. We will also present a result about random walk exit times which was initially shown in [2] for Bernoulli percolation clusters and then generalized to our setup in [3]. The first result gives Gaussian upper and lower bounds for the heat kernel for very good balls.

Proposition 2.1. *[3, Theorem 2.2] Assume the weights $\mu_{x,y}$ are i.i.d. and (1.2) holds. Fix a vertex $x \in V$. Suppose there exists $R_1 = R_1(x)$ such that $B(x, R)$ is very good with $N_{B(x,R)}^{3(d+2)} \leq R$ for every $R \geq R_1$. Then there exist positive constants c_1, c_2, c_3, c_4 such that if $t \geq R_1^{2/3}$, we obtain*

$$q_t(x,y) \leq c_1 t^{-d/2} e^{-c_2 |x-y|^2/t}, \quad \textit{for all } y \in V \textit{ with } |x-y| \leq t$$

and

$$q_t(x,y) \geq c_3 t^{-d/2} e^{-c_4 |x-y|^2/t}, \quad \textit{for all } y \in V \textit{ with } |x-y|^{3/2} \leq t.$$

Now define the space-time regions

$$\begin{aligned} Q(x,R,T) &= B(x,R) \times (0,T], \\ Q_-(x,R,T) &= B(x,\tfrac{R}{2}) \times [\tfrac{T}{4},\tfrac{T}{2}] \\ &\text{and} \\ Q_+(x,R,T) &= B(x,\tfrac{R}{2}) \times [\tfrac{3T}{4},T]. \end{aligned}$$

We denote by $t+Q(x,R,T) = B(x,R)\times(t,t+T)$. We call a function $u : V\times\mathbb{R}\to\mathbb{R}$ *caloric* on Q if it is defined on $Q = Q(x,R,T)$ and

$$\frac{\partial}{\partial t}u(x,t) = \mathcal{L}u(x,t) \quad \text{for all } (x,t)\in Q.$$

We say the *parabolic Harnack inequality* (PHI) holds with constant C_H for $Q = Q(x,R,T)$ if whenever $u = u(x,t)$ is non-negative and caloric on Q, then

$$\sup_{(x,t)\in Q_-(x,R,T)} u(x,t) \le C_H \inf_{(x,t)\in Q_+(x,R,T)} u(x,t).$$

It is well known that the heat kernel of a random walk on $\mathcal{G}$ started at x is a caloric function; in fact taking $x=0$ and $u(x,t) = q_t(0,x)$ we have

$$\begin{aligned}
\frac{\partial}{\partial t}q_t(0,x) &= \lim_{dt\to 0}\frac{1}{\mu_x}\frac{\sum_{y\sim x}\mathbb{P}_0(Y_t=y)\mathbb{P}_y(Y_{dt}=x) - \mathbb{P}_0(Y_t=x)\mathbb{P}_x(Y_{dt}\neq x)}{dt}\\
&= \frac{1}{\mu_x}\left(\sum_{y\sim x}\mathbb{P}_0(Y_t=y)\frac{\mu_{y,x}}{\mu_y} - \sum_{y\sim x}\frac{\mu_{x,y}}{\mu_x}\mathbb{P}_0(Y_t=x)\right)\\
&= \frac{1}{\mu_x}\sum_{y\sim x}\mu_{x,y}(q_t(0,y)-q_t(0,x)) = \mathcal{L}q_t(0,x).
\end{aligned}$$

The main result from [3] shows that the PHI holds in regions that are very good according to Definition 2.1.

Proposition 2.2. *[3, Theorem 3.1] Let $x_0\in V$. Suppose that $R_1\ge 16$ and $B(x_0,R_1)$ is $(C_V,C_P,C_W)-$very good with $N^{2d+4}_{B(x_0,R_1)}\le R_1/(2\log R_1)$. Then there exists a constant $C_H>0$ such that the PHI holds for $Q(x_1,R,R^2)$ for any $x_1\in B(x_0,R_1/3)$ and for R such that $R\log R = R_1$.*

A direct consequence of the PHI is the following known proposition, which when applied to the caloric function $u(x,t) = q_t(0,x)$ gives that $q_t(0,x)$ and $q_t(0,y)$ are very similar to each other when x and y are close by. This property will be crucial for our proof of local mixing, so we give the proof of this proposition for completeness.

Proposition 2.3. *Let $x_0\in V$. Suppose that there exists $s(x_0)\ge 0$ so that for all $R\ge s(x_0)$, the PHI holds with constant C_H for $Q(x_0,R,R^2)$ and that the ball $B(x_0,R)$ is (C_V,C_P,C_W) – very good. Let $\Theta = \log_2(C_H/(C_H-1))$, and for $x,y\in V$ define*

$$\rho(x_0,x,y) = s(x_0)\vee|x_0-x|\vee|x_0-y|.$$

There exists a constant $c>0$ such that the following holds. Let $r_0\ge s(x_0)$ and suppose that $u=u(x,t)$ is caloric in $Q=Q(x_0,r_0,r_0^2)$. Then for any $x_1,x_2\in B(x_0,\frac{1}{2}r_0)$ and *any* t_1, t_2 *such that* $r_0^2 - \rho(x_0, x_1, x_2)^2 \leq t_1, t_2 \leq r_0^2$ *we have*

$$|u(x_1, t_1) - u(x_2, t_2)| \leq c \left(\frac{\rho(x_0, x_1, x_2)}{r_0} \right)^{\Theta} \sup_{(t,x) \in Q_+(x_0, r_0, r_0^2)} |u(t, x)|. \tag{2.4}$$

Proof. For any integer $k \geq 0$, set $r_k = 2^{-k} r_0$, and let

$$\begin{aligned} Q(k) &= (r_0^2 - r_k^2) + Q(x_0, r_k, r_k^2), \\ Q_+(k) &= (r_0^2 - r_k^2) + Q_+(x_0, r_k, r_k^2) \\ &\textit{and} \\ Q_-(k) &= (r_0^2 - r_k^2) + Q_-(x_0, r_k, r_k^2). \end{aligned}$$

This gives that $Q_+(k) = Q(k+1)$. Now take $k \geq 1$ small enough, so that $r_k \geq s(x_0)$. If we apply the PHI to the non-negative caloric functions $-u + \sup_{Q(k)} u$ and $u - \inf_{Q(k)} u$, we get the inequalities

$$\sup_{Q(k)} u - \inf_{Q_-(k)} u \quad \leq \quad C_H (\sup_{Q(k)} u - \sup_{Q_+(k)} u)$$

and

$$\sup_{Q_-(k)} u - \inf_{Q(k)} u \quad \leq \quad C_H (\inf_{Q_+(k)} u - \inf_{Q(k)} u).$$

Adding them together and using $\sup_{Q_-(k)} u - \inf_{Q_-(k)} u \geq 0$ gives

$$\sup_{Q(k)} u - \inf_{Q(k)} u \leq C_H (\sup_{Q(k)} u - \inf_{Q(k)} u) - C_H (\sup_{Q_+(k)} u - \inf_{Q_+(k)} u).$$

Denoting by $\operatorname{Osc}(u, A) = \sup_A u - \inf_A u$ and setting $\delta = C_H^{-1}$, this gives

$$\operatorname{Osc}(u, Q_+(k)) \leq (1 - \delta) \operatorname{Osc}(u, Q(k)). \tag{2.5}$$

Next, take the largest m such that $r_m \geq \rho(x_0, x_1, x_2)$. Then, applying (2.5) repeatedly on $Q(1) \supset Q(2) \supset \ldots Q(m)$ yields, since $(x_i, t_i) \in Q(m)$,

$$|u(t_1, x_1) - u(t_2, x_2)| \leq \operatorname{Osc}(u, Q(m)) \leq (1 - \delta)^{m-1} \operatorname{Osc}(u, Q(1)).$$

Since

$$(1 - \delta)^m = 2^{-m\Theta} \leq \left(\frac{2\rho(x_0, x_1, x_2)}{r_0} \right)^{\Theta},$$

the result follows. □

We will also need to control the exit time of the random walk out of a ball of radius r, which we define as

$$\tau(x,r) = \inf\{t : Y_t \notin B(x,r)\}.$$

Proposition 2.4. *Let $x_0 \in V$ and let $B(x_0,R)$ be (C_V, C_P, C_W) − very good with $N_B^{d+2} < R$. Let $x \in B(x_0, \frac{5}{9}R)$. There exist positive constants c_1, c_2, c_3, c_4 such that if t, r satisfy*

$$0 < r \le R \quad \text{and} \quad c_1 N_B^d (\log N_B)^{1/2} r \le t \le c_2 R^2/\log R, \tag{2.6}$$

then we have

$$\mathbb{P}_x(\tau(x,r) < t) \le c_3 \exp\{-c_4 r^2/t\}. \tag{2.7}$$

Proof. The proposition was proven for percolation clusters in [2, Proposition 3.7]. The proof for more general $\mathcal{G}$ is similar and can be found in [3, Theorem 2.2a]. □

Since Propositions 2.1, 2.2 and 2.4 rely on very good balls and the related value N_B, we can assume a lower bound S such that if $R > S$, then the conditions of all three are satisfied. More formally, we assume the following.

Assumption 1. The graph G has polynomial growth; i.e., it satisfies (2.1). Furthermore, there exists a function $S : V \mapsto \mathbb{R}$ such that for all R_1 with $R_1 \log R_1 \ge S(x_0)$, the ball $B(x_0, R_1)$ is (C_V, C_P, C_W)-very good with $N_{B(x_0,R_1)}^{2d+4} \le R_1$. As a consequence, Propositions 2.1, 2.2, 2.3 and 2.4 all hold for any $R > S(x_0)$.

For i.i.d. weights as defined in Section 1, we obtain the following.

Proposition 2.5. *If $V = \mathbb{Z}^d$ and the weights $\mu_{x,y}$ are i.i.d. and satisfy (1.1) or (1.2), then Assumption 1 holds. Furthermore, we have that there exist constants $c, \gamma > 0$ such that*

$$\mathbb{P}[S(x) \ge n] \le c \exp\{-cn^\gamma\} \text{ for all } x \in \mathbb{Z}^d \text{ and } n \ge 0.$$

If the weights $\mu_{x,y}$ are i.i.d. and satisfy (1.1), then Assumption 1 holds with $S(x) = 1$ for all $x \in V$.

Proof. This has been shown in [3], following from the framework developed in [2, Theorem 2.18 and Lemma 2.19]. The tail estimate of $S(x)$ was obtained in [3, Theorem 5.7]. For the case where $\mu_{x,y}$ satisfy (1.1), the function S can be set to 1 by [3, Theorem 5.7] and the results of [5]. □

Remark 2.1. In [4] it has been shown that when the weights $\mu_{x,y}$ are i.i.d. but can assume values arbitrarily close to zero, so neither (1.1) nor (1.2) hold, it is possible to find distributions (at least in dimensions $d \ge 5$) for which Assumption 1 does not hold. Hence, even though we do not explicitly use uniform ellipticity of $\mu_{x,y}$ in our proofs, this property has a fundamental role in our analysis. Recent results (see, for example, [1]) have been derived to relax assumption (1.2), but they do not establish all the properties we need.

3 Decoupling via local mixing

In this section, we will restrict to the case $V = \mathbb{Z}^d$ and $(x, y) \in E$ if and only if $\|x - y\|_1 = 1$, but we do not assume the $\mu_{x,y}$ are i.i.d. We define a cube of side length $z > 0$ as $Q_z := [-z/2, z/2]^d$. In the remainder of the paper, we will work with the heat kernel q_t as defined in (2.3). Since we allow $\mu_{x,y} = 0$, it is possible for two sites not to be connected. To address this we require the existence of an infinite component. Formally, we assume the following.

Assumption 2. For each $(x, y) \in E$, either $\mu_{x,y} = 0$ or it satisfies (1.1) for a uniform constant C_M. Moreover, the weights $\mu_{x,y}$ are such that an infinite connected component of edges of positive weight within $\mathcal{G}$ exists and contains the origin.

With this let $\mathcal{C}_\infty$ be the infinite connected component of $\mathcal{G}$ that contains the origin and define

$$\tilde{Q}_z := Q_z \cap \mathcal{C}_\infty.$$

We note that if $\mu_{x,y}$ satisfy (1.1), then Assumption 2 is automatically satisfied. We will continue to call $\tilde{Q}_z$ as a "cube". We are now ready to state the more detailed version of Theorem 1.3.

Theorem 3.1. *Let $\mu_{x,y}$ satisfy Assumptions 1 and 2 and let $c > 0$ be an arbitrary constant. There exist constants c_0, c_1, $C > 0$ such that the following holds. Fix $K > \ell > 0$, $\epsilon \in (0, 1)$. Consider the cube Q_K tessellated into subcubes $(T_i)_i$ of side length ℓ. Let $(x_j)_j \subset \tilde{Q}_K$ be the locations at time 0 of a collection of particles, such that each subcube $\tilde{T}_i$ contains at least $\sum_{y \in \tilde{T}_i} \beta\mu_y$ particles for some $\beta > 0$. Assume that $\ell > S^{d+1}(x)$ for all $x \in \tilde{Q}_K$ and sufficiently large so that $\sum_{y \in \tilde{T}_i} \beta\mu_y \geq c$ for all subcubes $\tilde{T}_i$. Let $\Delta \geq c_0 \ell^2 \epsilon^{-4/\Theta}$ where Θ is as in Proposition 2.3. For each j denote by Y_j the location of the j-th particle at time Δ. Fix $K' > 0$ such that $K - K' \geq \sqrt{\Delta} c_1 \epsilon^{-1/d}$. Then there exists a coupling $\mathbb{Q}$ of an independent Poisson point process ψ with intensity measure $\zeta(y) = \beta(1 - \epsilon)\mu_y$, $y \in \mathcal{C}_\infty$, and $(Y_j)_j$ such that within $\tilde{Q}_{K'} \subset \tilde{Q}_K$, ψ is a subset of $(Y_j)_j$ with probability at least*

$$1 - \sum_{y \in \tilde{Q}_{K'}} \exp\left\{-C\beta\mu_y \epsilon^2 \Delta^{d/2}\right\}.$$

Note that, due to Proposition 2.5, Theorem 1.3 is a special case of Theorem 3.1, which we prove below. In order to do so, we will use something called *soft local times*, which was introduced in [11] to analyze random interlacements, following the introduction of local mixing in [12, 10, 13]; see also [7] for an application of this technique to random walks on $\mathbb{Z}^d$.

Proposition 3.1. *[11, Corollary 4.4.] Let $(Z_j)_{j\leq J}$ be a collection of J independent random particles on V distributed according to a family of density functions $g_j : V \to \mathbb{R}$, $j \leq J$. Define for all $y \in V$ the soft local time function $H_J(y) = \sum_{j=1}^J \xi_j g_j(y)$, where the ξ_j are i.i.d. exponential random variables of mean 1. Let ψ be a Poisson point process on V with intensity measure $\rho : V \to \mathbb{R}$ and define the event $E = \{$the particles belonging to ψ are a subset of $(Z_j)_{j\leq J}\}$. Then there exists a coupling between $(Z_j)_{j\leq J}$ and ψ, such that*

$$\mathbb{P}[E] \geq \mathbb{P}[H_J(y) \geq \rho(y),\ \forall y \in V].$$

We are now ready to prove Theorem 3.1.

Proof of Theorem 3.1. By Proposition 3.1, there exists a coupling $\mathbb{Q}$ of an independent Poisson point process ψ with intensity measure $\zeta(y) = \beta(1-\epsilon)\mu_y \mathbb{1}_{\{y\in\tilde{Q}_{K'}\}}$ and the locations of the particles Y_j, which are distributed according to the density functions $f_\Delta(x_j, y) := q_\Delta(x_j, y)\mu_y$, $y \in \mathcal{C}_\infty$, such that the particles belonging to ψ are a subset of $(Y_j)_j$ with probability at least

$$\mathbb{Q}[H_J(y) \geq \beta\mu_y(1-\epsilon),\ \forall y \in \tilde{Q}_{K'}],$$

where $H_J(y) = \sum_{j=1}^J \xi_j f_\Delta(x_j, y)$, $(\xi_j)_{j\leq J}$ are i.i.d. exponential random variables with parameter 1, and J is the number of particles inside $\tilde{Q}_K$. We first observe that the probability of the converse event is

$$\begin{aligned}\mathbb{Q}[\exists y \in \tilde{Q}_{K'} :\ H_J(y) < \beta\mu_y(1-\epsilon)] &\leq \sum_{y\in\tilde{Q}_{K'}} \mathbb{Q}[H_J(y) < \beta\mu_y(1-\epsilon)]\\ &\leq \sum_{y\in\tilde{Q}_{K'}} e^{\kappa\mu_y\beta(1-\epsilon)}\mathbb{E}^{\mathbb{Q}}[\exp\{-\kappa H_J(y)\}],\end{aligned}$$

where we used Markov's inequality in the last step, which is valid for any $\kappa > 0$. Let c_1 be a positive constant which we will fix later and let

$$R = \sqrt{\Delta}c_1\epsilon^{-1/d}.$$

Let J' be a subset of $\{1, 2, \ldots, J\}$ such that for each $\tilde{T}_i$, J' contains exactly $\lceil \sum_{y \in \tilde{T}_i} \beta\mu_y \rceil$ particles that are inside $\tilde{T}_i$. Define $J'(y) \subseteq J'$ to be the set of $j \in J'$ such that $|x_j - y| \le R$ and define $H'(y)$ as $H_J(y)$ but with the sum restricted to $j \in J'(y)$. Since $H_J(y) \ge H'(y)$ we get that

$$\mathbb{E}^{\mathbb{Q}}[\exp\{-\kappa H_J(y)\}] \le \mathbb{E}^{\mathbb{Q}}[\exp\{-\kappa H'(y)\}]. \tag{3.1}$$

Next, we use that the ξ_j in the definition of H are independent exponential random variables to obtain

$$\begin{aligned} \mathbb{E}^{\mathbb{Q}}[\exp\{-\kappa H'(y)\}] &= \prod_{j \in J'(y)} \mathbb{E}^{\mathbb{Q}}[\exp\{-\kappa \xi_j f_\Delta(x_j, y)\}] \\ &= \prod_{j \in J'(y)} (1 + \kappa f_\Delta(x_j, y))^{-1}. \end{aligned} \tag{3.2}$$

Using Taylor's expansion we have that $\log(1 + x) \ge x - x^2$ for $|x| \le \frac{1}{2}$. Since $\ell \ge S(x)$, we can apply Proposition 2.1, to have $q_\Delta(x, y) \le c_2 \Delta^{-d/2}$ for a constant $c_2 > 0$ and all $y \in \tilde{Q}_{K'}$ and $x \in J'(y)$. Hence if $\kappa = C\epsilon\Delta^{d/2}$ for the constant $C = (4C_U c_2)^{-1}$, then

$$\sup_{x \in B(y, R + \sqrt{d}\ell)} \kappa f_\Delta(x, y) = \sup_{x \in B(y, R + \sqrt{d}\ell)} \kappa \mu_y q_\Delta(x, y) \le C_U c_2 \kappa \Delta^{-d/2} < \frac{\epsilon}{4}.$$

For such a value of κ we have

$$\begin{aligned} \prod_{j \in J'(y)} (1 + \kappa f_\Delta(x_j, y))^{-1} &\le \prod_{j \in J'(y)} \exp\{-\kappa f_\Delta(x_j, y)(1 - \kappa f_\Delta(x_j, y))\} \\ &\le \exp\left\{ - \sum_{j \in J'(y)} \kappa f_\Delta(x_j, y) \left(1 - \sup_{x \in B(y, R + \sqrt{d}\ell)} \kappa f_\Delta(x, y) \right) \right\} \\ &\le \exp\left\{ -\kappa \sum_{j \in J'(y)} f_\Delta(x_j, y)(1 - \epsilon/4) \right\}. \end{aligned} \tag{3.3}$$

We claim that

$$\sum_{j \in J'(y)} f_\Delta(x_j, y) \ge \beta\mu_y(1 - \epsilon/2), \tag{3.4}$$

which together with (3.3), (3.2) and (3.1) give that

$$\begin{aligned} \mathbb{Q}\left[\exists y \in \tilde{Q}_{K'} :\ H_J(y) < \beta\mu_y(1 - \epsilon)\right] &\le \exp\{\kappa\mu_y\beta(1 - \epsilon) - \kappa\beta\mu_y(1 - \epsilon/2)(1 - \epsilon/4)\} \\ &\le \exp\{-\kappa\beta\mu_y\epsilon/4\}. \end{aligned}$$

Using the value of κ gives the theorem.

It remains to show (3.4). For each $\tilde{T}_i$ and each particle $x_j \in \tilde{T}_i$, let $x'_j \in \tilde{T}_i$ be such that $f_\Delta(x'_j, y) = \max_{w \in \tilde{T}_i} f_\Delta(w, y)$. Then, write

$$\sum_{j \in J'(y)} f_\Delta(x_j, y) \geq \sum_{j \in J'(y)} \left(f_\Delta(x'_j, y) - |f_\Delta(x'_j, y) - f_\Delta(x_j, y)| \right). \tag{3.5}$$

We have for each $\tilde{T}_i$

$$\begin{aligned} \sum_{\substack{j \in J'(y) \\ x_j \in \tilde{T}_i}} f_\Delta(x'_j, y) &= \max_{w \in \tilde{T}_i} f_\Delta(w, y) \sum_{\substack{j \in J'(y) \\ x_j \in \tilde{T}_i}} 1 \\ &\geq \max_{w \in \tilde{T}_i} f_\Delta(w, y) \sum_{z \in \tilde{T}_i} \beta \mu_z \\ &\geq \sum_{z \in \tilde{T}_i} \beta \mu_z f_\Delta(z, y). \end{aligned} \tag{3.6}$$

Set $R(y)$ to be the set of all sites z such that $|z - y| \leq R - \sqrt{d}\ell$. Note that if $z \in R(y)$ then for all particles x_j with $x'_j = z$ and $j \in J'$ we have $j \in J'(y)$. We observe that since $\mu_z f_\Delta(z, y) = \mu_y f_\Delta(y, z)$, we have by using (3.6) for each $\tilde{T}_i$ that

$$\begin{aligned} \sum_{j \in J'(y)} f_\Delta(x'_j, y) &\geq \sum_{z \in R(y)} \beta \mu_z f_\Delta(z, y) \\ &= \beta \mu_y \sum_{z \in R(y)} f_\Delta(y, z). \end{aligned}$$

Then, since $\ell > S^{d+1}(x)$ we have by Proposition 2.4 that there exist constants c_4 and c_5 such that

$$\begin{aligned} \sum_{j \in J'(y)} f_\Delta(x'_j, y) &\geq \beta \mu_y \mathbb{P}_y(\tau(y, R - \sqrt{d}\ell) \geq \Delta) \\ &\geq \beta \mu_y (1 - c_4 \exp\{-c_5 c_1^2 \epsilon^{-2/d}\}) \\ &\geq \beta \mu_y (1 - \epsilon/4), \end{aligned} \tag{3.7}$$

where we set c_1 large enough with respect to c_4 and c_5 for the last inequality to hold.

Now it remains to obtain an upper bound for the term $\sum_{j \in J'(y)} |f_\Delta(x'_j, y) - f_\Delta(x_j, y)|$. We define I to be the set of all i such that $\tilde{T}_i$ contains a particle x_j from the set $(x_j)_{j \in J'(y)}$. Then, since $\ell > S(x)$, there exists positive constants C_{PHI} and C_{BH} such that if we apply the PHI (cf. Proposition 2.3) with

$$r_0^2 = \Delta \geq c_0 \ell^2 \epsilon^{-4/\Theta} \tag{3.8}$$

for some constant $c_0 > d$, we obtain

$$\begin{aligned}
\sum_{j \in J'(y)} |f_\Delta(x'_j, y) - f_\Delta(x_j, y)| &= \sum_{i \in I} \sum_{\substack{j \in J'(y): \\ x_j \in \tilde{T}_i}} |f_\Delta(x'_j, y) - f_\Delta(x_j, y)| \\
&= \mu_y \sum_{i \in I} \sum_{\substack{j \in J'(y): \\ x_j \in \tilde{T}_i}} |q_\Delta(x'_j, y) - q_\Delta(x_j, y)| \\
&\leq \mu_y \sum_{i \in I} \sum_{\substack{j \in J'(y): \\ x_j \in \tilde{T}_i}} \frac{C_{PHI} \ell^\Theta}{\Delta^{\Theta/2}} C_{BH} \Delta^{-d/2} \\
&\leq \mu_y \sum_{i \in I} \sum_{x \in \tilde{T}_i} \frac{2 \max\left\{1, \frac{1}{c}\right\} \beta \mu_x C_{PHI} \ell^\Theta}{\Delta^{\Theta/2}} C_{BH} \Delta^{-d/2},
\end{aligned}$$

where in the first inequality we replaced the supremum term coming from Proposition 2.3 by its upper bound $C_{BH}\Delta^{-d/2}$ from Proposition 2.1, and used that $r_0 = \sqrt{\Delta}$ in the bound from Proposition 2.3. Then

$$\begin{aligned}
\sum_{j \in J'(y)} |f_\Delta(x'_j, y) - f_\Delta(x_j, y)| &\leq 2 \max\left\{1, \frac{1}{c}\right\} \beta \mu_y C_{PHI} C_{BH} \sum_{i \in I} \sum_{x \in \tilde{T}_i} \mu_x \ell^\Theta \Delta^{-(d+\Theta)/2} \\
&\leq 2 \max\left\{1, \frac{1}{c}\right\} \beta \mu_y C_{PHI} C_{BH} C_U R^d \ell^\Theta \Delta^{-(d+\Theta)/2} \\
&\leq \beta \mu_y \frac{\epsilon}{4},
\end{aligned} \tag{3.9}$$

where the last inequality holds by using $\Delta \geq c_0 \ell^2 \epsilon^{-4/\Theta}$ and setting $c_0 > (2 \max\left\{1, \frac{1}{c}\right\} C_{PHI} C_{BH} C_U c_1^d)^{-2/\theta}$. Note that in order to use Proposition 2.3, we need to have that each pair x_j, x'_j is contained in some ball $B(x_0, r_0/2)$. This is satisfied since $\|x_j - x'_j\| \leq \sqrt{d}\ell$ and r_0 is set sufficiently large by (3.8). Plugging (3.9) and (3.7) into (3.5) proves (3.4). □

4 Extensions

Although the estimate derived in Theorem 3.1 does not depend on the particles outside of Q_K at time 0 when $K - K'$ is sufficiently large, it still depends on the geometry of the entire graph outside of Q_K. In some applications, as in our companion paper [6], one needs to apply this coupling in many different regions of the graph simultaneously. In such cases, in order to control dependences between different regions, it is important that the coupling procedure depends only on the local structure of the graph. In order to do this, we will condition the particles to be inside some large enough, but finite region while they move for time Δ. Recall that, for any $\rho > 0$, $Q_\rho = [-\rho/2, \rho/2]^d$ is the cube of side length ρ. For any $\rho > 0$, we say that a random walk has *displacement* in Q_ρ during $[0, \Delta]$ if the random walk never exits $x + Q_\rho$ during the time interval $[0, \Delta]$, where x is the starting vertex of the random walk.

Lemma 4.1. *Let $\mu_{x,y}$ satisfy Assumptions 1 and 2. There exist constants c_1 and c_2 so that the following holds. Let $V = \mathbb{Z}^d$, $\ell > 0$ and consider the cube Q_ℓ. Assume $\ell > S(x)$ for all $x \in Q_\ell$. Let $\Delta > c_1\ell^2$ and $\rho \geq c_2\sqrt{\Delta \log \Delta}$. Consider a random walk Y that moves along $\mathcal{G}$ for time Δ conditioned on having its displacement in Q_ρ during the time interval $[0, \Delta]$. Let $x, y \in Q_\ell$ with x being the starting point of the walk, and define*

$$g(x, y) := \mathbb{P}_x \left[Y_\Delta = y \,|\, Y \text{ has displacement in } Q_\rho \text{ during } [0, \Delta] \right].$$

Then there exists a constant $C > 2$ such that for $x, y, z \in Q_\ell$ we have

$$\left| \frac{g(x, y)}{\mu_y} - \frac{g(z, y)}{\mu_y} \right| \leq C\ell^\Theta \Delta^{-(d+\Theta)/2}.$$

Remark 4.1. Note that the above bound has the same form as the one for the heat kernel of unconditioned random walks in Proposition 2.3, with the supremum being bounded above by the heat kernel bound from Proposition 2.1. This allows us to extend Theorem 3.1 to random walks conditioned to have a bounded displacement during $[0, \Delta]$.

Proof of Lemma 4.1. Denote by $p_E(\rho)$ the probability that a random walk started at x has displacement in Q_ρ during $[0, \Delta]$. From Proposition 2.4 , we have that if Δ is sufficiently big, then

$$\begin{aligned} 1 - p_E(\rho) &\leq \mathbb{P}_x[Y \text{ exits } B(x, \rho/2) \text{ during } [0, \Delta]] \\ &= \mathbb{P}_x(\tau(x, \rho/2) < \Delta) \\ &\leq c_a \exp\{-c_b\rho^2/\Delta\}. \end{aligned} \tag{4.1}$$

Next, using $h(x, y) := \mathbb{P}_x [Y_\Delta = y \,|\, Y \text{ exits } x + Q_\rho \text{ during } [0, \Delta]]$ and $f_\Delta(x, y) =$

$\mathbb{P}_x[Y_\Delta = y]$, we can write

$$f_\Delta(x,y) = g(x,y)p_E(\rho) + h(x,y)(1 - p_E(\rho)).$$

With this we have

$$g(x,y) \le f_\Delta(x,y)\frac{1}{p_E(\rho)}. \tag{4.2}$$

Then, we can write

$$\begin{aligned}
\left|\frac{g(x,y)}{\mu_y} - \frac{g(z,y)}{\mu_y}\right| &= \mathbb{1}_{\{g(x,y)>g(z,y)\}}\left(\frac{g(x,y)}{\mu_y} - \frac{g(z,y)}{\mu_y}\right) \\
&\quad +\mathbb{1}_{\{g(x,y)<g(z,y)\}}\left(\frac{g(z,y)}{\mu_y} - \frac{g(x,y)}{\mu_y}\right) \\
&\le \mathbb{1}_{\{g(x,y)>g(z,y)\}}\left(\frac{f_\Delta(x,y)}{\mu_y p_E(\rho)} - \frac{f_\Delta(z,y)}{\mu_y p_E(\rho)} + \frac{h(z,y)(1-p_E(\rho))}{p_E(\rho)\mu_y}\right) \\
&\quad +\mathbb{1}_{\{g(x,y)<g(z,y)\}}\left(\frac{f_\Delta(z,y)}{\mu_y p_E(\rho)} - \frac{f_\Delta(x,y)}{\mu_y p_E(\rho)} + \frac{h(x,y)(1-p_E(\rho))}{p_E(\rho)\mu_y}\right) \\
&\le \frac{|q_\Delta(y,x) - q_\Delta(y,z)|}{p_E(\rho)} + \frac{\max\{h(x,y),h(z,y)\}(1-p_E(\rho))}{p_E(\rho)\mu_y}.
\end{aligned}$$

Note that $h(x,y)$ can be written as $\mathbb{E}[f_{\Delta-\tau}(w,y) \mid \tau < \Delta]$, where τ is the first time Y exists $x + Q_\rho$ and w is the random vertex at the boundary of $x + Q_\rho$ where Y is at time τ. Since the weights $\mu_{x,y}$ satisfy (2.1) by Assumption 1, we have that $\frac{f_{\Delta-\tau}(w,y)}{\mu_y}$ is at most some positive constant c. This holds because either $\Delta - \tau$ is larger than $|w-y|$, which allows us to apply heat kernel bounds from Proposition 2.1, or $\Delta - \tau$ is smaller than $|w-y|$ so $f_{\Delta-\tau}(w,y)$ is bounded above by the probability that a random walk jumps at least $|w-y|$ steps in time $\Delta - \tau$, which is small enough since $|w-y|$ is large. This gives that $\frac{\max\{h(x,y),h(z,y)\}}{\mu_y}$ is at most c. With this and (4.1) we obtain that

$$\begin{aligned}
\frac{\max\{h(x,y),h(z,y)\}(1-p_E(\rho))}{\mu_y p_E(\rho)} &\le \frac{cc_a}{p_E(\rho)}\exp\left\{\frac{-c_b\rho^2}{\Delta}\right\} \\
&\le \frac{cc_a}{p_E(\rho)}\exp\left\{-c_b c_2 \log\Delta\right\}.
\end{aligned}$$

By (4.1) we can just bound $p_E(\rho)$ below by $1/2$. Then, applying Proposition 2.5 to $|q_\Delta(y,x) - q_\Delta(y,z)|$, and using Proposition 2.1 to bound the resulting supremum term, concludes the proof. □

The next theorem is an adaptation of Theorem 3.1 for conditioned random walks. Note that we need a stronger condition on $K - K'$ below than in Theorem 3.1.

Theorem 4.1. *Let $\mu_{x,y}$ satisfy Assumptions 1 and 2 and let $c > 0$ be an arbitrary constant. There exist constants c_0, c_1, $C > 0$ such that the following holds. Fix $K > \ell > 0$ and $\epsilon \in (0,1)$. Consider the cube Q_K tessellated into subcubes $(T_i)_i$ of side length ℓ.Let $(x_j)_j \subset \tilde{Q}_K$ be the locations at time 0 of a collection of particles, such that each subcube $\tilde{T}_i$ contains at least $\sum_{y\in\tilde{T}_i} \beta\mu_y$ particles for some $\beta > 0$. Assume that $\ell > S^{d+1}(x)$ for all $x \in \tilde{Q}_K$ and sufficiently large so that $\sum_{y\in\tilde{T}_i} \beta\mu_y \geq c$ for all subcubes $\tilde{T}_i$. Let $\Delta \geq c_0\ell^2\epsilon^{-4/\Theta}$, where Θ is as in Proposition 2.3. Fix $K' > 0$ such that $K - K' \geq c_1\sqrt{\Delta \log \Delta}$. For each j, denote by Y_j the location of the j-th particle at time Δ, conditioned on having displacement in $Q_{K-K'}$ during $[0,\Delta]$. Then there exists a coupling $\mathbb{Q}$ of an independent Poisson point process ψ with intensity measure $\zeta(y) = \beta(1-\epsilon)\mu_y$, $y \in \tilde{Q}_K$, and $(Y_j)_j$ such that within $\tilde{Q}_{K'} \subset \tilde{Q}_K$, ψ is a subset of $(Y_j)_j$ with probability at least*

$$1 - \sum_{y\in\tilde{Q}_{K'}} \exp\left\{-C\beta\mu_y\epsilon^2\Delta^{d/2}\right\}.$$

Proof. Using Lemma 4.1 and (4.2) when setting κ, the proof goes in the same way as the proof of Theorem 3.1. The independence from G outside of $\tilde{Q}_K$ follows from the fact that we only consider particles which have displacement in $Q_{K-K'}$ and ended in $\tilde{Q}_{K'}$, so that they never left $\tilde{Q}_K$ during $[0,\Delta]$. □

4.1 Extension to other graphs

We have shown that the local mixing result of Theorems 3.1 and 4.1 work for $\mathbb{Z}^d$, but they can easily be extended to the more general graphs defined in Section 2, as long as Assumptions 1 and 2 hold.

We start with a region $A \subseteq \mathcal{C}_\infty$ around the origin of G and tesselate it into tiles $(T_i)_{i\in I}$ of diameter at most ℓ. Let Δ be as in Theorem 3.1. Let $A' \subset A$ be all the sites in A that are at least $\sqrt{\Delta}c_1\epsilon^{-1/d} + c\ell$ away from the boundary of A. Then, if A' is not empty, using the same steps as in the proof of Theorem 3.1, if each tile T_i of A contains at least $\beta\sum_{y\in T_i}\mu_y$ particles at time 0, it holds that in the region A', there is a coupling with an independent Poisson point process ψ of intensity measure $\zeta(y) = \beta(1-\epsilon)\mu_y$ such that at time Δ the particles inside A' are contained in ψ with probability at least

$$1 - \sum_{y\in A'} \exp\left\{-C\beta\mu_y\epsilon^2\Delta^{d/2}\right\},$$

for some constant $C > 0$.

Furthermore, Theorem 4.1 can analogously be extended in the same way, if we require that A' contains only sites that are at least $c_1\sqrt{\Delta \log \Delta}$ away from the boundary of A, for some constant c_1, and if we condition the random walks to have their displacement limited to a ball of radius $c_1\sqrt{\Delta \log \Delta}$.

5 Spread of the infection

Our goal in this section will be to use Theorem 4.1 in order to show that on the graph $G = (V, E)$ with $V = \mathbb{Z}^d$ and $E = \{(x, y) : \|x - y\|_1 = 1\}$, and with $\mu_{x,y}$, $(x, y) \in E$ being i.i.d. and satisfying (1.1), information spreads with positive speed in any direction, as claimed in Theorems 1.1 and 1.2. In this setting, Proposition 2.5 guarantees that Assumption 1 holds with $S(x) \equiv 1$ and since $\mu_{x,y} \neq 0$ for all $(x, y) \in E$, we also have that Assumption 2 holds.

Recall that we assume $d \geq 2$. Tessellate $\mathbb{Z}^d$ into cubes of side length ℓ, indexed by $i \in \mathbb{Z}^d$. Next, tessellate time into intervals of length β, indexed by $\tau \in \mathbb{Z}$. With this we denote by the space-time cell $(i, \tau) \in \mathbb{Z}^{d+1}$ the region $\prod_{j=1}^{d}[i_j\ell, (i_j+1)\ell] \times [\tau\beta, (\tau+1)\beta]$. In the following, β is set as a function of ℓ so that the ratio β/ℓ^2 is fixed first to be a small constant, and then ℓ is set sufficiently large.

We will use a result from [6] that gives the existence of a Lipschitz connected surface (cf. Definitions 5.2 and 5.3 below) that surrounds the origin and which is composed of space-time cells, for which a certain local event holds. This will allow us to obtain an infinite sequence of space-time cells, such that the infection spreads from one cell to the next.

In order to obtain this result, we will need to consider overlapping space-time cells. Let $\eta \geq 1$ be an integer which will represent the amount of overlap between cells. For each cube $i = (i_1, \dots, i_d)$ and time interval τ, define the *super cube* i as $\prod_{j=1}^{d}[(i_j - \eta)\ell, (i_j + \eta + 1)\ell]$ and the *super interval* τ as $[\tau\beta, (\tau + \eta)\beta]$. We define the *super cell* (i, τ) as the Cartesian product of the super cube i and the super interval τ.

In the following we will say a particle has displacement inside X' during a time interval $[t_0, t_0 + t_1]$, if the location of the particle at all times during $[t_0, t_0 + t_1]$ is inside $x + X'$, where x is the location of the particle at time t_0. We define a particle system on $\mathbb{Z}^d$ as a countable family of not necessarily unique elements of $\mathbb{Z}^d$, indexed by some countable set I, representing the locations of the particles belonging to the particle system. Let $(\Pi_s)_{s\geq 0}$ be a sequence of particles system on $\mathbb{Z}^d$, with Π_s representing the locations of the particles at time s. We say a particle system Π_s is distributed according to a Poisson random measure of intensity ζ, if for every $B \subset \mathbb{Z}^d$, $N(B)$ is a Poisson random variable with intensity $\zeta(B)$, where $N(B)$ is the number of particles belonging to Π_s that lie in B. We say an event E is *increasing* for $(\Pi_s)_{s\geq 0}$ if the fact that E holds for $(\Pi_s)_{s\geq 0}$ implies that it holds for all $(\Pi'_s)_{s\geq 0}$ for which $\Pi'_s \supseteq \Pi_s$ for all $s \geq 0$. We say an event E is *restricted* to a region $X \subset \mathbb{Z}^d$ and a time interval $[t_0, t_1]$ if it is measurable with respect to the σ-field generated by all the particles that are inside X at time t_0 and their positions from time t_0 to t_1. For an increasing event E that is restricted to a region X and time interval $[t_0, t_1]$, we have the following definition.

Definition 5.1. ν_E is called the *probability associated* to a an increasing event E that is restricted to X and a time interval $[t_0, t_0 + t_1]$ if, for an intensity measure ζ, $\nu_E(\zeta, X, X', t_1)$ is the probability that E happens given that, at time t_0, the particles in X are a particle system distributed according to the Poisson random measure of intensity ζ and their motions from t_0 to $t_0 + t_1$ are independent continuous time random walks on the weighted graph (G, μ), where the particles are conditioned to have displacement inside X'.

For each $(i, \tau) \in \mathbb{Z}^{d+1}$, let $E_{\mathrm{st}}(i, \tau)$ be an increasing event restricted to the super cube i and the super interval τ. Here the subscript st refers to space-time. We say that a cell (i, τ) is *bad* if $E_{\mathrm{st}}(i, \tau)$ does not hold and *good* otherwise.

We will need a different way to index space-time cells, which we refer to as the *base-height index*. In the base-height index, we pick one of the d spatial dimensions and denote it as *height*, using index $h \in \mathbb{Z}$, while the remaining d space-time dimensions form the base, which we index by $b \in \mathbb{Z}^d$. In this way, for each space-time cell (i, τ) there will be $(b, h) \in \mathbb{Z}^{d+1}$ such that the base-height cell (b, h) corresponds to the space-time cell (i, τ)

We analogously define the *base-height super cell* (b, h) to be the space-time super cell (i, τ), for which the base-height cell (b, h) corresponds to the space-time cell (i, τ). Similarly, we define $E_{\mathrm{bh}}(b, h)$, the increasing event restricted to the super cell (b, h) that is the same as the event $E_{\mathrm{st}}(i, \tau)$ for the space-time super cell (i, τ) that corresponds to the base-height super cell (b, h). Here, the subscript bh refers to the base-height index.

In order to prove Theorems 1.1 and 1.2, we will need a theorem from [6], which gives the existence of a two-sided Lipschitz surface F.

Definition 5.2. A function $F : \mathbb{Z}^d \to \mathbb{Z}$ is called a *Lipschitz function* if $|F(x) - F(y)| \leq 1$ whenever $\|x - y\|_1 = 1$.

Definition 5.3. A *two-sided Lipschitz surface* F is a set of base-height cells $(b, h) \in \mathbb{Z}^{d+1}$ such that for all $b \in \mathbb{Z}^d$ there are exactly two (possibly equal) integer values $F_+(b) \geq 0$ and $F_-(b) \leq 0$ for which $(b, F_+(b)), (b, F_-(b)) \in F$ and, moreover, F_+ and F_- are Lipschitz functions.

We say a space-time cell (i, τ) belongs to F if there exists a base-height cell $(b, h) \in F$ that corresponds to (i, τ). We say a two-sided Lipschitz surface F exists, if for all $b \in \mathbb{Z}^d$, we have $F_+(b) < \infty$ and $F_-(b) > -\infty$. For a positive integer D, we say a two-sided Lipschitz surface *surrounds* a cell (b', h') at distance D if any path $(b', h') = (b_0, h_0), (b_1, h_1), \dots, (b_n, h_n)$ for which $\|(b_i, h_i) - (b_{i-1}, h_{i-1})\|_1 = 1$ for all $i \in \{1, \dots n\}$ and $\|(b_n, h_n) - (b_0, h_0)\|_1 > D$, intersects with F.

We now present the main result from our paper [6], which holds for graphs where a local mixing result, such as the one in Theorem 4.1, hold. More precisely, for a graph satisfying Assumption 1 and (1.1) (which implies Assumption 2 holds) we have that Theorem 4.1 holds (with $S(x) = 1$ for all $x \in V$), which in turn gives that the following result from [6] holds. Recall that, for any $\rho \geq 2$, Q_ρ stands for the cube $[-\rho/2, \rho/2]^d$, and that λ is the intensity measure of the Poisson point process of particles as defined in Section 1.

Theorem 5.1. *Let $\mathcal{G} = (G, \mu)$ be a graph satisfying Assumption 1 and (1.1) on the lattice $\mathbb{Z}^d$ for $d \geq 2$. There exist positive constants c_0,, c_1 and c_2 such that the following holds. Tessellate G in space-time cells and super cells as described above for some $\ell, \beta, \eta > 0$ such that the ratio $\beta/\ell^2 < c_0$. Let $E_{\mathrm{st}}(i, \tau)$ be an increasing event, restricted to the space-time super cell (i, τ). Fix $\epsilon \in (0, 1)$ and fix w such that*

$$w \geq \sqrt{\frac{\eta\beta}{c_2\ell^2} \log\left(\frac{8c_1}{\epsilon}\right)}.$$

Then, there exists a positive number α_0 that depends on ϵ, η and that ratio β/ℓ^2 so that if

$$\min\left\{C_M^{-1}\epsilon^2\lambda_0\ell^d, \log\left(\frac{1}{1 - \nu_{E_{\mathrm{st}}}((1-\epsilon)\lambda, Q_{(2\eta+1)\ell}, Q_{w\ell}, \eta\beta)}\right)\right\} \geq \alpha_0,$$

a two-sided Lipschitz surface F where $E_{\mathrm{st}}(i, \tau)$ holds for all $(i, \tau) \in F$ exists almost surely. Furthermore, the surface surrounds the origin at a finite distance almost surely.

Recall that we want to show that the infection spreads with positive speed. Given a space-time tessellation of G and a local increasing event E_{st}, Theorem 5.1 gives the existence of a Lipschitz surface F on which E_{st} holds. Let $T = \ell^{5/3}$. We will define the increasing event $E_{\mathrm{st}}(i, \tau)$ to represent a single infected particle in the middle of the super cube i at time $\tau\beta$ infecting a large number of particles in that super cube by time $\tau\beta + T$, after which the infected particles move up to time $(\tau + 1)\beta$, spreading to all of the cubes contained in the super cube.

Let (i, τ) be a space-time cell as defined previously. We consider that there is an infected particle in the center cube of the super cube i at time $\tau\beta$, that is, the particle is inside $\prod_{j=1}^d [i_j\ell, (i_j + 1)\ell]$. Starting from time $\tau\beta$, we let the infected particle move and infect sufficiently many other particles by time $\tau\beta + T$. This is given in the lemma below.

Lemma 5.1. *There exist positive constant C_1 such that the following holds for all large enough ℓ. Let $Q^* = \prod_{j=1}^d [(i_j - \eta)\ell, (i_j + \eta + 1)\ell]$ and let $(\rho(t))_{\tau\beta \le t \le \tau\beta + T}$ be the path of an infected particle that starts in $\prod_{j=1}^d [i_j\ell, (i_j+1)\ell]$ and stays inside $\prod_{j=1}^d [(i_j - \eta + 1)\ell, (i_j + \eta)\ell]$ during $[\tau\beta, \tau\beta + T]$. Assume that at time $\tau\beta$, the number of particles at each vertex $x \in Q^* \setminus \rho(\tau\beta)$ is a Poisson random variable of mean $\frac{\lambda_0}{2}\mu_x$. Let Υ be the set of these particles, and let $\Upsilon' \subset \Upsilon$ be the particles colliding with the path ρ, that is, for each particle of Υ' there exists a time $t \in [\tau\beta, \tau\beta + T]$ such that the particle is located at $\rho(t)$. Then, $|\Upsilon'|$ is a Poisson random variable of mean at least $C_1\lambda_0\ell^{1/3}$.*

Remark 5.1. We note that the statement of Lemma 5.1 is conditional on the path $(\rho(t))_{\tau\beta \le t \le \tau\beta + T}$. The bound we obtain is uniform across all such paths that we will consider later in Lemma 5.3, so we omit this in our notation.

Proof. For each time $t \in [\tau\beta, \tau\beta + T]$, let Ψ_t be the Poisson point process on V giving the locations at time t of the particles that belong to Υ. Since the particles that start in Q^* move around and can leave Q^*, we need to find a lower bound for the intensity of Ψ_t for times in $[\tau\beta, \tau\beta + T]$. Note that the infected particle we are tracking is not part of Ψ, since Ψ does not include particles located at $\rho(\tau\beta)$ at time $\tau\beta$.

We will need to apply heat kernel bounds from Proposition 2.1 to the particles in Q^*, so we need to ensure that the time intervals we consider are large enough for the proposition to hold. We will only consider times $t \in [\ell^{4/3}, T]$ so that for large enough ℓ, $t \ge \sup_{\substack{x \in Q^* \\ y \in Q^*}} \|x - y\|_1$ and so the heat kernel bounds from Proposition 2.1 hold. Then, we have that for all sites $x \in Q^*$ that are at least ℓ away from the boundary of Q^* and at any such time t the intensity of $\Psi_{\tau\beta + t}$ at vertex $x \in V$ is at least

$$\psi(x, \tau\beta + t) \ge \sum_{\substack{y \in Q^* \\ y \ne \rho(\tau\beta)}} \tfrac{\lambda_0}{2}\mu_y \cdot \mu_x q_t(y, x) = \tfrac{\lambda_0}{2}\mu_x \sum_{\substack{y \in Q^* \\ y \ne \rho(\tau\beta)}} \mathbb{P}_x[Y_t = y],$$

where we used in the last step that the heat kernel q_t is symmetric. We now use the exit time bound from Proposition 2.4 to get that

$$\sum_{y \in Q^*} \mathbb{P}_x[Y_t = y] \ge 1 - c_3 \exp\{-c_4\ell^2/t\}.$$

Next, we use that $\mathbb{P}_x[Y_t = y] = \mu_y q_t(x, y) \le C_M q_t(x, y)$, and use Proposition 2.1 to account for the particles at $\rho(\tau\beta)$, yielding

$$\sum_{\substack{y\in Q^* \\ y\neq\rho(\tau\beta)}} \mathbb{P}_x[Y_t = y] \geq 1 - c_3 \exp\left\{-c_4\ell^2/t\right\} - C_M c_5 t^{-d/2}.$$

This gives that for any $t \in [\ell^{4/3}, T]$, the intensity of $\Psi_{\tau\beta+t}$ is at least

$$\psi(x, \tau\beta + t) \geq \tfrac{\lambda_0}{2}\mu_x(1 - c_3 \exp\{-c_4\ell^2/T\} - C_M c_5 \ell^{-2d/3}).$$

Let $[\tau\beta, \tau\beta + T]$ be divided into subintervals of length $W \in (0, T]$, where we set $W = \ell^{4/3}$ so that it is large enough to allow the use of the heat kernel bounds from Proposition 2.1. Let $J = \{1, \dots, \lfloor T/W \rfloor\}$ and $t_j := \tau\beta + jW$. Then the intensity of particles that share a site with the initially infected particle at least once among times $\{t_1, t_2, \dots, t_{\lfloor T/W \rfloor}\}$ is at least

$$\sum_{j\in J} \psi(\rho(t_j), t_j)\mathbb{P}_{\rho(t_j)}[Y_{r-t_j} \neq \rho(r)\ \forall r \in \{t_{j+1}, \dots, t_{\lfloor T/W \rfloor}\}]$$
$$\geq \tfrac{\lambda_0}{2} C_M^{-1}(1 - c_3 \exp\{-c_4\ell^2/T\} - C_M c_5 \ell^{-2d/3}) \sum_{j\in J}\left(1 - \sum_{z>j} \mathbb{P}_{\rho(t_j)}[Y_{t_z - t_j} = \rho(t_z)]\right).$$

We want to make all of the terms of the sum over J positive, so we consider the term $\sum_{z>j} \mathbb{P}_{\rho(t_j)}[X_{t_z-t_j} = \rho(t_z)]$ and show that it is smaller than $\frac{1}{2}$ for large enough ℓ. To do this, we use that $\mathbb{P}_x[Y_t = y] = \mu_y q_t(x, y)$ with the heat kernel bounds from Proposition 2.1, which hold when $W \geq \ell^{4/3}$ and ℓ is large enough, to bound it from above by

$$\begin{aligned}\sum_{z>j} \mathbb{P}_{\rho(t_j)}[Y_{t_z-t_j} = \rho(t_z)] &\leq \sum_{z>j} C_M C_{HK}(t_z - t_j)^{-d/2} \\ &\leq C_M C_{HK} W^{-d/2} \sum_{z=1}^{T/W-j} z^{-d/2} \end{aligned} \tag{5.1}$$

where C_{HK} is the constant coming from Proposition 2.1. Then, (5.1) can be bounded from above by

$$C_M C_{HK} W^{-d/2}\left(2 + \sum_{z=3}^{T/W-j} z^{-d/2}\right) \leq C_M C_{HK} W^{-d/2}\left(2 + \int_2^{T/W} z^{-d/2} dz\right). \tag{5.2}$$

Let C be a constant that can depend on C_{HK}, C_M and d. Then for $d = 2$, (5.2) it is smaller than $CW^{-1}\log(T/W)$, and for $d \geq 3$ the expression in (5.2) is smaller than

$CW^{-d/2}$. Thus, setting ℓ large enough, both terms are smaller than $\frac{1}{2}$.

Then, as a sum of Poisson random variables, we get that Υ' is a Poisson random variable with a mean at least

$$\tfrac{\lambda_0}{2} C_M^{-1}(1 - c_3 \exp\{-2c_4\ell^2/T\} - C_M c_5 \ell^{-2d/3})\tfrac{T}{2W}.$$

Using that $T = \ell^{5/3}$ and setting ℓ large enough establishes the lemma, with C_1 being any constant satisfying $C_1 < \frac{C_M^{-1}}{4}$. □

Next we show that the particles from Lemma 5.1 move to nearby cells, spreading the infection.

Lemma 5.2. *Let $z = (z_1, \dots, z_d)$ with $z_j \in \{-\eta, -\eta+1, \dots, \eta\}$ for all $j \in \{1, \dots d\}$, and fix the ratio β/ℓ^2. Let $A(i, \tau, N, z)$ be the event that given a set of $N > 0$ particles in $\prod_{j=1}^d [(i_j - \eta)\ell, (i_j + \eta + 1)\ell]$ at time $\tau\beta + T$, at least one of them is in $\prod_{j=1}^d [(i_j + z_j)\ell, (i_j + z_j + 1)\ell]$ at time $(\tau+1)\beta$. Then, if ℓ is sufficiently large while keeping β/ℓ^2 fixed, we obtain*

$$\mathbb{P}[A(i, \tau, N, z)] \geq 1 - \exp\{-Nc_p\},$$

where c_p is a positive constant that is bounded away from 0 and depends only on d, η and the ratio β/ℓ^2.

Proof. Let $Q^* = \prod_{j=1}^d [(i_j - \eta)\ell, (i_j + \eta + 1)\ell]$ and $Q^{**} = \prod_{j=1}^d [(i_j + z_j)\ell, (i_j + z_j + 1)\ell]$. For $t^{2/3} \geq \sup_{\substack{x \in Q^* \\ y \in Q^{**}}} \|x - y\|_1$, define $p_t := \inf_{x \in Q^*} \sum_{y \in Q^{**}} \mathbb{P}_x[Y_t = y]$. Then, if we define $\text{bin}(N, p_t)$ to be a binomial random variable of parameters $N \in \mathbb{N}$ and $p_t \in [0, 1]$, it directly follows that

$$\mathbb{P}[A(i, \tau, N, z)] \geq \mathbb{P}[\text{bin}(N, p_t) \geq 1] \geq 1 - \exp\{-Np_t\}.$$

It remains to show that for $t = \beta - T$, we have that $p_t \geq c_p > 0$ for some constant c_p. We will use the heat kernel bounds for the pair x, y, which hold if $\|x - y\|_1^{3/2} \leq \beta - T$ for all $x \in Q^*, y \in Q^{**}$. Given the ratio β/ℓ^2, d and η, this is satisfied if ℓ is large enough. Then we have that

$$\begin{aligned} p_{\beta - T} &= \inf_{x \in Q^*} \sum_{y \in Q^{**}} \mathbb{P}_x[Y_{\beta - T} = y] \\ &\geq \inf_{x \in Q^*} C_M^{-1} \sum_{y \in Q^{**}} q_{\beta - T}(x, y) \\ &\geq \inf_{x \in Q^*} C_M^{-1} \sum_{y \in Q^{**}} c_1 \beta^{-d/2} \exp\left\{-c_2 \frac{\|x - y\|_1^2}{\beta - T}\right\}. \end{aligned}$$

Now we use that x and y can be at most $c_\eta \ell$ apart where c_η is a constant depending on d and η only, and that $\beta - T \geq \beta/2$ for ℓ large enough. Hence,

$$\begin{aligned} p_{\beta-T} &\geq \inf_{x\in Q^*} C_M^{-1} \sum_{y\in Q^{**}} c_1 \beta^{-d/2} \exp\left\{-c_2 \frac{2(c_\eta \ell)^2}{\beta}\right\} \\ &= C_M^{-1} c_1 \ell^d \left(\frac{1}{\beta}\right)^{d/2} \exp\left\{-c_2 \frac{2(c_\eta \ell)^2}{\beta}\right\} \\ &\geq c_p. \end{aligned}$$

□

In the next lemma, we will tie together the results from Lemma 5.1 and Lemma 5.2. In order to precisely describe the behavior of the particles involved, we say a particle x *collides* with particle y during a time interval $[t_0, t_1]$, if for at least one $t \in [t_0, t_1]$, x and y are at the same site.

Lemma 5.3. *Consider the super cell (i, τ). Assume that at each site $x \in \prod_{j=1}^d [(i_j - \eta)\ell, (i_j + \eta + 1)\ell]$ the number of particles at x at time $\tau\beta$ is a Poisson random variable of intensity $\frac{\lambda_0}{2}\mu_x$, and let Υ be the collection of such particles. Assume that, at time $\tau\beta$, there is at least one infected particle x_0 inside $\prod_{j=1}^d [i_j\ell, (i_j + 1)\ell]$. Let $E_{\mathrm{st}}(i, \tau)$ be the event that at time $(\tau + 1)\beta$, for all $i' \in \mathbb{Z}^d$ with $\|i - i'\|_\infty \leq \eta$, there is at least one particle from Υ in $\prod_{j=1}^d [(i'_j)\ell, (i'_j + 1)\ell]$ that collided with x_0 during $[\tau\beta, \tau\beta + T]$. If ℓ is sufficiently large for Lemmas 5.1 and 5.2 to hold, then there exists a positive constant C such that*

$$\mathbb{P}[E_{\mathrm{st}}(i, \tau)] \geq 1 - \exp\{-C\lambda_0 \ell^{1/3}\}.$$

Proof. We note that, by definition, the event $E_{\mathrm{st}}(i, \tau)$ is restricted to the super cube $\prod_{j=1}^d [(i_j - \eta)\ell, (i_j + \eta + 1)\ell]$ and time interval $[\tau\beta, (\tau + 1)\beta]$. We define the following 3 events.

- F_1: The initial infected particle x_0 never leaves $\prod_{j=1}^d [(i_j - \eta + 1)\ell, (i_j + \eta - 1)\ell]$ during $[\tau\beta, \tau\beta + T]$.
- F_2: Let C_1 be the constant from Lemma 5.1. During the time interval $[\tau\beta, \tau\beta + T]$ the initial infected particle x_0 collides with at least $\frac{C_1\lambda_0\ell^{1/3}}{2}$ different particles from Υ that are in the supercube $Q^{**} = \prod_{j=1}^d [(i_j - \eta)\ell, (i_j + \eta + 1)\ell]$ at time $\tau\beta + T$.
- F_3: Out of the $\frac{C_1\lambda_0\ell^{1/3}}{2}$ or more particles from F_2, at least one of them is in the cube $\prod_{j=1}^d [(i_j + k_j)\ell, (i_j + k_j + 1)\ell]$ at time $(\tau + 1)\beta$, for all $k = (k_1, \ldots, k_d)$ for which $\prod_{j=1}^d [(i_j + k_j)\ell, (i_j + k_j + 1)\ell] \subset Q^{**}$.

By definition of the events, we clearly have that $\mathbb{P}[E_{\mathrm{st}}(i,\tau)] \geq \mathbb{P}[F_1 \cap F_2 \cap F_3]$.

Using Proposition 2.4 we have

$$\mathbb{P}[F_1] \geq 1 - C_2 \exp\{-C_3\ell^2/T\} = 1 - C_2 \exp\{-C_3\ell^{1/3}\} \tag{5.3}$$

for some positive constants C_2 and C_3. We observe that F_1 is restricted to the super cube $\prod_{j=1}^d [(i_j - \eta)\ell, (i_j + \eta + 1)\ell]$ and the time interval $[\tau\beta, \tau\beta + T]$.

For the event F_2, we apply Lemma 5.1 to get that the intensity of the Poisson point process of particles that are in Q^{**} at time $\tau\beta$ and collide with x_0 during $[\tau\beta, \tau\beta+T]$ is at least $\lambda_0 C_1 \ell^{1/3}$ for some positive constant C_1. Since every particle that collides with x_0 enters $\prod_{j=1}^d [(i_j - \eta + 1)\ell, (i_j + \eta)\ell]$ during $[\tau\beta, \tau\beta + T]$, we can use Proposition 2.4 to bound the probability that the particle is inside of Q^{**} at time $\tau\beta + T$ by

$$1 - C_a \exp\left\{-\frac{C_b\ell^2}{T}\right\} = 1 - C_a \exp\{-C_b\ell^{1/3}\},$$

for some positive constants C_a and C_b. This term can be made as close to 1 as possible by having ℓ sufficiently large. We assume ℓ is large enough so that this term is larger than 2/3. This gives that the intensity of the process of particles from Υ that collided with x_0 during $[\tau\beta, \tau\beta + T]$ and are in Q^{**} at time $\tau\beta + T$ is at least

$$\frac{2\lambda_0 C_1 \ell^{1/3}}{3}.$$

Using Chernoff's bound (see Lemma A.1) we have that

$$\mathbb{P}[F_2] \geq 1 - \exp\{-(2/3)^2 C_1 \lambda_0 \ell^{1/3}\}. \tag{5.4}$$

Note that, by construction, F_2 is restricted to the super cube Q^{**} and the time interval $[\tau\beta, \tau\beta + T]$. Furthermore, F_2 is clearly an increasing event.

We now turn to F_3. Using Lemma 5.2, and a uniform bound across the number of cubes inside a super cube, we have that

$$\mathbb{P}[F_3] \geq 1 - (2\eta + 1)^d \exp\{-\frac{C_1\lambda_0\ell^{1/3}}{2} c_p\}, \tag{5.5}$$

where c_p is a small but positive constant. Again, the event is restricted to the super cube Q^{**} and the time interval $[\tau\beta + T, (\tau+1)\beta]$ and is an increasing event. Taking the product of the probability bounds in (5.3), (5.4) and (5.5), we see that the probability that $E_{\mathrm{st}}(i,\tau)$ holds is at least

$$1-\exp\{-C\lambda_0\ell^{1/3}\}$$

for some constant C and all large enough ℓ. □

Proof of Theorem 1.1. We start by using Theorem 5.1. Set $\eta \in \mathbb{N}$ such that $\eta \geq d$ and set $\epsilon = 1/2$. Fix the ratio β/ℓ^2 small enough so that the lower bound for w is at most $2\eta+1$, and then set $w = 2\eta+1$. Assume ℓ is large enough so that Lemma 5.3 holds.

For each $(i,\tau) \in \mathbb{Z}^{d+1}$, define $E_{\mathrm{st}}(i,\tau)$ as in Lemma 5.3. This event is increasing in the number of particles, is restricted to the super cube i and time interval $[\tau\beta,(\tau+1)\beta]$, and satisfies

$$\mathbb{P}[E_{\mathrm{st}}(i,\tau)] \geq 1-\exp\{-C\lambda_0\ell^{1/3}\},$$

for some constant C. Hence, letting $\lambda/2$ stand for the measure $\frac{\lambda}{2}(x) = \frac{\lambda_0\mu_x}{2}$, we have

$$\log\left(\frac{1}{1-\nu_{E_{\mathrm{st}}}(\frac{\lambda}{2},Q_{(2\eta+1)\ell},Q_{(2\eta+1)\ell},\eta\beta)}\right) \geq C\lambda_0\ell^{1/3},$$

which increases with ℓ, as does the term $\epsilon^2\lambda_0\ell^d$ in the condition of Theorem 5.1. Thus, setting ℓ large enough, we apply Theorem 5.1 which gives the existence of a two-sided Lipschitz surface F, on which the event $E_{\mathrm{st}}(i,\tau)$ holds. We also get that the surface is almost surely finite and that it surrounds the origin.

We now proceed to argue that the existence of the surface F implies that the infection spreads with positive speed. Since the two-sided Lipschitz surface F is finite and surrounds the origin, we have that in almost surely finite time, an infected particle started from the origin will enter some cube $\prod_{j=1}^d [i_j\ell,(i_j+1)\ell]$ for which (i,τ) is in F. We call this the *central cube* of (i,τ). Once that holds, the starting assumption of $E_{\mathrm{st}}(i,\tau)$ from Lemma 5.3 is satisfied for the super cell (i,τ), and the event $E_{\mathrm{st}}(i,\tau)$ holds. By the definition of $E_{\mathrm{st}}(i,\tau)$ this means that the initial infected particle for the super cell (i,τ) infects a large number of other particles, which spread the infection to the central cube of $(i',\tau+1)$ for all $i' \in \mathbb{Z}^d$ such that $\|i'-i\|_\infty \leq \eta$.

Let (b,h) be the base-height index of the cell $(i,\tau) \in F$. Recall that h is one of the spatial dimensions. We will also select one of the $d-1$ spatial dimensions from b and denote it b_1. Let $b' \in \mathbb{Z}^d$ be obtained from b by increasing the time dimension from τ to $\tau+1$, and by increasing the chosen spatial dimension from b_1 to b_1+1. Since $\|b-b'\|_1 = 2$, we can choose $h' \in \mathbb{Z}$ such that $(b',h') \in F$ and $|h-h'| \leq 2$, where the latter holds by the Lipschitz property of F. Therefore, there must exists $i' \in \mathbb{Z}^d$ such that $(i',\tau+1)$ is the space-time super cell corresponding to (b',h') and $\|i-i'\|_\infty \leq 1$.

Hence, at time $(\tau+1)\beta$, there is an infected particle in the central cube of the super cell i'.

We can then recursively repeat this procedure for the super cell $(i', \tau+1)$, since $E_{\rm st}(i', \tau+1)$ holds. Repeating this process we obtain that the infection spreads by a distance of at least ℓ in time β in the chosen spatial direction. Consequently

$$\liminf_{t\to\infty} \frac{\|I_t\|_1}{t} > 0 \quad \text{almost surely.}$$

□

In order to prove Theorem 1.2, we can follow the same steps as in the proof of Theorem 1.1 with the additional consideration that we have to ensure that the relevant infected particles do not recover too quickly. For that, we will require that all the particles involved do not recover for at least β.

Proof Theorem 1.2. Recall the definition of Υ and ρ from Lemma 5.1 and of $E_{\rm st}(i,\tau)$ from Lemma 5.3. Let $E'_{\rm st}(i,\tau)$ be the event that $E_{\rm st}(i,\tau)$ holds, and that the particles in Υ and the initial infected particle whose path is ρ do not recover during $[\tau\beta, (\tau+1)\beta]$. Since each such particle does not recover during $[\tau\beta, (\tau+1)\beta]$ with probability $\exp\{-\gamma\beta\}$, for Lemma 5.1 we consider that for each $x \in Q^* \setminus \rho(\tau\beta)$ the number of particles at x at time $\tau\beta$ that do not recover during $[\tau\beta, (\tau+1)\beta]$ is a Poisson random variable of intensity $\frac{\lambda_0}{2}\mu_x \exp\{-\lambda\beta\}$. Thus, once η, β and ℓ are fixed, setting γ small enough gives that $E'_{\rm st}(i,\tau)$ holds with probability at least

$$1 - (1-\exp\{-\gamma\beta\}) - \exp\{-C\lambda_0 \exp\{-\gamma\beta\}\ell^{1/3}\}$$

for some positive constant C, where the term inside the parenthesis accounts of the probability that the initial infected particles recovers during $[\tau\beta, (\tau+1)\beta]$. We now follow the same steps as in the proof of Theorem 1.1 to get that the two-sided Lipschitz surface F on which the increasing event $E'_{\rm st}(i,\tau)$ holds exists, is finite and surrounds the origin almost surely. This gives that an initially infected particle that is at the origin at time 0 has a strictly positive probability of surviving long enough to enter a cell of the two-sided Lipschitz surface. Once on the surface, the infection survives indefinitely by the definition of $E'_{\rm st}(i,\tau)$. Hence

$$\mathbb{P}\left[\|I_t\|_1 \geq c_1 t \text{ for all } t \geq c_3\right] \geq c_2.$$

□

A Appendix: Standard large deviation results

Lemma A.1 (Chernoff bound for Poisson). *Let P be a Poisson random variable with mean λ. Then, for any $0 < \epsilon < 1$,*

$$\mathbb{P}[P < (1-\epsilon)\lambda] < \exp\{-\lambda\epsilon^2/2\}$$

and

$$\mathbb{P}[P > (1+\epsilon)\lambda] < \exp\{-\lambda\epsilon^2/4\}.$$

References

[1] Andres, S., Deuschel, J.D. and Slowik, M., 2016. Harnack inequalities on weighted graphs and some applications to the random conductance model. *Probability Theory and Related Fields*, 164(3-4), pp.931–977. Available from: `http://doi.org/10.1007/s00440-015-0623-y`.

[2] Barlow, M.T., 2004. Random walks on supercritical percolation clusters. *Annals of Probability*, 32(4), pp.3024–3084. Available from: `http://doi.org/10.1214/009117904000000748`.

[3] Barlow, M.T. and Hambly, B.M., 2009. Parabolic Harnack inequality and local limit theorem for percolation clusters. *Electronic Journal of Probability*, 14, pp.1–26. Available from: `http://doi.org/10.1214/EJP.v14-587`.

[4] Berger, N., Biskup, M., Hoffman, C.E. and Kozma, G., 2008. Anomalous heat-kernel decay for random walk among bounded random conductances. *Annales de l'institut Henri Poincare (B) Probability and Statistics*, 44(2), pp.374–392. Available from: `http://doi.org/10.1214/07-AIHP126`.

[5] Delmotte, T., 1999. Parabolic Harnack inequality and estimates of Markov chains on graphs. *Revista Matematica Iberoamericana*, 15(1), pp.181–232. Available from: `http://doi.org/10.4171/RMI/254`.

[6] Gracar, P. and Stauffer, A., 2017. Multi-scale Lipschitz percolation of increasing events for Poisson random walks. *arXiv*. `1702.08748`, Available from: `https://arxiv.org/abs/1702.08748`.

[7] Hilário, M., Hollander, F. den, Sidoravicius, V., Soares dos Santos, R. and Teixeira, A., 2015. Random walk on random walks. *Electronic Journal of Probability*, 20. Available from: `http://doi.org/10.1214/EJP.v20-4437`.

[8] Kesten, H. and Sidoravicius, V., 2005. The spread of a rumor or infection in a moving population. *Annals of Probability*, 33(6), pp.2402–2462. Available from: http://doi.org/10.1214/009117905000000413.

[9] Kesten, H. and Sidoravicius, V., 2008. A shape theorem for the spread of an infection. *Annals of Mathematics*, 167(3), pp.701–766. Available from: http://doi.org/10.4007/annals.2008.167.701.

[10] Peres, Y., Sinclair, A., Sousi, P. and Stauffer, A., 2013. Mobile geometric graphs: Detection, coverage and percolation. *Probability Theory and Related Fields*, 156(1-2), pp.273–305. Available from: http://doi.org/10.1007/s00440-012-0428-1.

[11] Popov, S. and Teixeira, A., 2015. Soft local times and decoupling of random interlacements. *Journal of the European Mathematical Society*, 17(10), pp.2545–2593. Available from: http://doi.org/10.4171/JEMS/565.

[12] Sinclair, A. and Stauffer, A., 2010. Mobile geometric graphs, and detection and communication problems in mobile wireless networks. *arXiv*. 1005.1117v2, Available from: http://arxiv.org/abs/1005.1117.

[13] Stauffer, A., 2015. Space-time percolation and detection by mobile nodes. *Annals of Applied Probability*, 25(5), pp.2416–2461. Available from: http://doi.org/10.1214/14-AAP1052.

Closing remarks

Local mixing results have been used in a wide variety of constructions before, but they are typically proven by using an appropriate local central limit theorem. For random conductance graphs, a local central limit theorem was proven by Barlow and Hambly [3], but it only holds in the limit as time goes to infinity. Therefore, without good control over the convergence rate, we had to resort to a different approach. Building on the results of Barlow and Hambly [3], and using only heat kernel bounds and the parabolic Harnack inequality, we were able to prove a local mixing result for a more general class of graphs. While this result is important by itself, it serves as a means to an end for our main result, proving that infection spreads with positive speed on uniformly elliptic random conductance graphs.

The final pages of Section 5 illustrate the flexibility that using local events affords us. Traditionally, it would take a considerable amount of work in order to adapt a result that holds for infection without recovery to the case with recovery. With our framework, the result follows with only minor modifications to the definition of the increasing event E_{st}. This is in part due to how the particles behave locally, which we control with heat kernel bounds. Just as important however is the two-sided Lipschitz surface, the existence of which lets us “combine” local events into global behavior. In the next chapter, we will carefully define this percolation structure and perform a rigorous multi-scale analysis that leads to its existence.

Chapter 3

Multi-scale Lipschitz percolation of increasing events for Poisson random walks

Introduction to the paper

In this chapter, we present the paper *Multi-scale Lipschitz percolation of increasing events for Poisson random walks.* We proved the local mixing result in Chapter 2 and then applied it to prove Theorems 1.1 and 1.2; it remains to show that the object that lets us do this actually exists. First, in Sections 1 and 2 of the paper, we introduce the problem and define the tessellation of space and time into cells on which the Lipschitz surface is defined. In Section 3, we define for general site percolation on $\mathbb{Z}^{d+1}$ the *two-sided Lipschitz surface* via special sequences of cells that we call d-paths. In order to deal with the dependences between the cells, we generalize the tessellation of space and time to multiple scales and in Section 4 introduce a multi-scale version of d-paths, which we call D-paths.

The main technical work is then done in Sections 5 and 6 in order to show that the local mixing result from Chapter 2 guarantees that across multiple scales, D-paths of space-time cells for which particles do not behave "well", i.e., are not close to stationarity, are sufficiently unlikely. This in turn leads to the existence of the two-sided Lipschitz surface for which the increasing event E_{st} holds for all cells in the surface; we show this in Section 7. This is also precisely the statement of Theorem 1.4 from Chapter 1. In Section 8 of the paper, we introduce DD-paths. These paths are a generalization of

D-paths and allow for more control on long-distance behavior of the Lipschitz surface. More precisely, we use DD-paths to show that the two-sided Lipschitz surface surrounds the origin at a finite distance, as stated by Theorem 1.5 from Chapter 1. Finally, DD-paths are also used to prove Theorem 1.6.

UNIVERSITY OF BATH

Statement of Authorship (to preface each co-authored paper)

<table>
<tr><td colspan="8">This declaration concerns the article entitled:</td></tr>
<tr><td colspan="8">Multi-scale Lipschitz percolation of increasing events for Poisson random walks</td></tr>
<tr><td colspan="8">Publication status (tick one)</td></tr>
<tr><td>Draft manuscript</td><td></td><td>Submitted</td><td>X</td><td>In review</td><td></td><td>Accepted</td><td></td><td>Published</td><td></td></tr>
<tr><td>Publication details (reference)</td><td colspan="9">Annals of Applied Probability</td></tr>
<tr><td>Candidate's contribution to the paper (detailed, and also given as a percentage).</td><td colspan="9">The candidate was the primary contributor to:
- Formulation of ideas:
○ Definition and construction of the Lipschitz surface based on Lipschitz percolation. 75%
○ Multi-scale analysis to establish the existence of the Lipschitz surface. 75%
- Design of methodology:
○ Implemented proofs of results, including all details. 100%</td></tr>
<tr><td>Statement from Candidate</td><td colspan="9">This paper reports on original research I conducted during the period of my Higher Degree by Research candidature.</td></tr>
<tr><td>Signed</td><td colspan="5"></td><td colspan="2">Date</td><td colspan="2">15.08.2017</td></tr>
</table>

Multi-scale Lipschitz percolation of increasing events for Poisson random walks

Peter Gracar* and Alexandre Stauffer†

Summary

Consider the graph induced by $\mathbb{Z}^d$, equipped with *uniformly elliptic* random conductances. At time 0, place a Poisson point process of particles on $\mathbb{Z}^d$ and let them perform independent simple random walks. Tessellate the graph into cubes indexed by $i \in \mathbb{Z}^d$ and tessellate time into intervals indexed by τ. Given a local event $E(i,\tau)$ that depends only on the particles inside the space time region given by the cube i and the time interval τ, we prove the existence of a Lipschitz connected surface of *cells* (i,τ) that separates the origin from infinity on which $E(i,\tau)$ holds. This gives a directly applicable and robust framework for proving results in this setting that need a multi-scale argument. For example, this allows us to prove that an infection spreads with positive speed among the particles.

Keywords and phrases: multi-scale percolation, Lipschitz surface, spread of infection

1 Introduction

Let $G = (\mathbb{Z}^d, E)$ be the d-dimensional square lattice with edges between nearest neighbors: $(x,y) \in E$ iff $\|x-y\|_1 = 1$. Start with a collection of particles given by a Poisson point process on $\mathbb{Z}^d$ of intensity λ, and let the particles move over time as independent continuous time simple random walks on G. We refer to this system of particles as *Poisson random walks*.

Assume that at time 0 there is an infected particle at the origin, and that all other particles are uninfected. As particles move, an uninfected particle gets infected as soon as it shares a site with an infected particle. Kesten and Sidoravicius [7] showed that for

*Department of Mathematical Sciences, University of Bath, UK. p.gracar@bath.ac.uk

†Department of Mathematical Sciences, University of Bath, UK, a.stauffer@bath.ac.uk. Supported by a Marie Curie Career Integration Grant PCIG13-GA-2013-618588 DSRELIS, and an EPSRC Early Career Fellowship.

all $\lambda > 0$ the infection spreads with positive speed; that is, for all large enough t, at time t there is an infected particle at distance of order t from the origin. A main challenge in establishing this result is that, as the infection spreads, it finds empty regions (i.e., regions without particles) of arbitrarily large sizes. An empty region $A \subset \mathbb{Z}^d$ not only delays the spread of the infection locally, but also causes a decrease in the density of particles in a neighborhood around A as time goes on. A key part of the analysis in [7] is to control how often empty regions arise and how big an impact (in space and time) they cause. An additional challenge is that long-range dependences do arise. For example, if at some time the ball $B(x, r)$ of radius r centered at $x \in \mathbb{Z}^d$ is empty, then $B(x, r/2)$ is likely to remain empty for a time of order r^2. Thus, the probability that the space-time region $B(x, r/2) \times [0, r^2]$ is empty of particles is at least exponential in r^d, which is only a stretched exponential with respect to the volume of the space-time region. In [7], the effect of empty regions was controlled via an intricate *multi-scale argument.*

The problem of spread of infection among Poisson random walks is just one example where long-range dependences give rise to serious mathematical challenges, and where multi-scale arguments have been applied to great success. In fact, multi-scale arguments have proved to be very useful in the analysis of several models, including the solution of several important questions regarding Poisson random walks [7, 8, 9, 13], activated random walks [12], random interlacements [11, 14], multi-particle diffusion limited aggregation [10] and more general dependent percolation [3, 15].

However, the main problem in developing a multi-scale analysis is that the argument is quite involved and can become very technical. Also, in each of the examples above, the involved multi-scale argument had to be developed from scratch and be tailored to the specific question being analyzed. Our main goal in this paper is to develop a more robust and systematic framework that can be applied to solve questions in the model of Poisson random walks without the need of carrying out a whole multi-scale argument each time. We do this by showing that given a local event which is translation invariant and whose probability of occurrence is large enough, we can find a special percolating structure in space-time where this event holds.

We now explain our idea in a high-level way, deferring precise statements and definitions to Section 2. We tesselate space into cubes, indexed by $i \in \mathbb{Z}^d$, and tessellate time into intervals indexed by $\tau \in \mathbb{Z}$. Thus (i, τ) denotes the space-time cell of the tessellation consisting of the cube i and the time interval τ. Given any increasing, translation invariant event $E(i, \tau)$ that is local (i.e., measurable with respect to the particles that get within some fixed distance to the space-time cell (i, τ)), if the marginal distribution

$\mathbb{P}(E(i,\tau))$ is large enough, our main result gives the existence of a *two-sided Lipschitz surface* of space-time cells where $E(i,\tau)$ holds for all cells in the surface.

Once we obtain such a Lipschitz surface, instead of having to carry out a whole multi-scale analysis from scratch to analyze some question involving Poisson random walks, one is left with the much easier task of just coming up with a suitable choice of $E(i,\tau)$. For example, for the case of spread of infection mentioned above, a natural choice is to define $E(i,\tau)$ as the event that an infected particle in the cube i infects several other particles which then move to all cubes neighboring i by the end of the time interval τ. Then, the existence of the Lipschitz surface and its Lipschitz property ensures that, once the infection enters the surface, it is guaranteed to propagate through the surface.

We further illustrate the applicability of our Lipschitz surface technique in [5], where we apply the Lipschitz surface to study the spread of infection in the random conductance model.

2 Setting and precise statement of the results

Poisson random walks. We consider the graph $(\mathbb{Z}^d, E)$ with conductances $\{\mu_{x,y}\}_{(x,y)\in E}$, which are i.i.d. non-negative weights on the edges of G. In this paper, edges will always be undirected, so $\mu_{x,y} = \mu_{y,x}$ for all $(x,y) \in E$. We also assume that the conductances are *uniformly elliptic*: that is,

$$\begin{gathered}\text{there exists deterministic } C_M > 0 \text{, such that} \\ \mu_{x,y} \in [C_M^{-1}, C_M] \text{ for all } (x,y) \in E,\ \mathbb{P} - a.s.\end{gathered} \tag{1}$$

We say $x \sim y$ if $(x,y) \in E$ and define $\mu_x = \sum_{y\sim x} \mu_{x,y}$. At time 0, consider a Poisson point process of particles on $\mathbb{Z}^d$, with intensity measure $\lambda(x) = \lambda_0\mu_x$ for some constant $\lambda_0 > 0$ and all $x \in \mathbb{Z}^d$. That is, for each $x \in \mathbb{Z}^d$, the number of particles at x at time 0 is an independent Poisson random variable of mean $\lambda_0\mu_x$. Then, let the particles perform independent continuous-time simple random walks on the weighted graph; i.e., a particle at $x \in \mathbb{Z}^d$ jumps to a neighbor $y \sim x$ at rate $\frac{\mu_{x,y}}{\mu_x}$. It follows from the thinning property of Poisson random variables that the system of particles is in stationarity; that is, at any time t, the particles are distributed according to a Poisson point process with intensity measure λ. We refer to this system of particles as *Poisson random walks* on (G,μ) with intensity λ_0.

Tessellation. We now tesselate the graph $G = (\mathbb{Z}^d, E)$ into d-dimensional cubes of side length $\ell > 0$. We index the cubes of the tessellation by integer vectors $i \in \mathbb{Z}^d$ such that the cube $i = (i_1, i_2, \ldots, i_d)$ corresponds to the region $\left(\prod_{j=1}^{d}[i_j\ell, (i_j+1)\ell]\right) \cap \mathbb{Z}^d$. Tessellate time into subintervals of length β. We index the subintervals by $\tau \in \mathbb{Z}$, representing the time interval $[\tau\beta, (\tau+1)\beta]$. We refer to the pair (i, τ), representing $\prod_{j=1}^{d}[i_j\ell, (i_j+1)\ell] \times [\tau\beta, (\tau+1)\beta]$, as a *space-time cell* and define the *region of a cell* as $R_1(i, \tau) = \prod_{j=1}^{d}[i_j\ell, (i_j+1)\ell] \times [\tau\beta, (\tau+1)\beta]$.

We will need to consider larger space-time cells as well. Let $\eta \geq 1$ be an integer. For each cube $i = (i_1, \ldots, i_d)$ and time interval τ, define the *super cube* i as $\prod_{j=1}^{d}[(i_j - \eta)\ell, (i_j + \eta + 1)\ell]$ and the *super interval* τ as $[\tau\beta, (\tau+\eta)\beta]$. We define the *super cell* (i, τ) as the Cartesian product of the super cube i and the super interval τ.

Definitions for events. We define a particle system on $\mathbb{Z}^d$ as a countable family of not necessarily unique elements of $\mathbb{Z}^d$, indexed by some countable set I, representing the locations of the particles belonging to the particle system. Let $(\Pi_s)_{s\geq 0}$ be a sequence of particle systems on $\mathbb{Z}^d$, with Π_s representing the locations of the particles at time s. We say a particle system Π_s is distributed according to a Poisson random measure of intensity ζ, if for every $A \subset \mathbb{Z}^d$, $N(A)$ is a Poisson random variable with intensity $\zeta(A)$, where $N(A)$ is the number of particles belonging to Π_s that lie in A. We say an event E is *increasing* for $(\Pi_s)_{s\geq 0}$ if the fact that E holds for $(\Pi_s)_{s\geq 0}$ implies that it holds for all $(\Pi'_s)_{s\geq 0}$ for which $\Pi'_s \supseteq \Pi_s$ for all $s \geq 0$. We need the following definitions.

Definition 2.1. We say an event E is *restricted* to a region $X \subset \mathbb{Z}^d$ and a time interval $[t_0, t_1]$ if it is measurable with respect to the σ-field generated by all the particles that are inside X at time t_0 and their positions from time t_0 to t_1.

Definition 2.2. We say a particle has displacement inside X' during a time interval $[t_0, t_0 + t_1]$, if the location of the particle at all times during $[t_0, t_0 + t_1]$ is inside $x + X'$, where x is the location of the particle at time t_0.

For an increasing event E that is restricted to a region X and time interval $[0, t]$, we have the following definition.

Definition 2.3. ν_E is called the *probability associated* to an increasing event E that is restricted to X and a time interval $[0, t]$ if, for an intensity measure ζ and a region $X' \in \mathbb{Z}^d$, $\nu_A(\zeta, X, X', t)$ is the probability that E happens given that, at time 0, the particles in X are a particle system distributed according to the Poisson random measure of intensity ζ and their motions from 0 to t are independent continuous time random walks on the weighted graph (G, μ), where the particles are conditioned to have displacement inside X' during $[0, t]$.

For each $(i, \tau) \in \mathbb{Z}^{d+1}$, let $E_{\text{st}}(i, \tau)$ be an increasing event restricted to the super cube i and the super interval τ. We will assume that $E_{\text{st}}(i, \tau)$ is invariant under space-time translations. We say that a cell (i, τ) is *good* if $E_{\text{st}}(i, \tau)$ holds and *bad* otherwise.

The base-height index. We will need a different way to index space-time cells, which we refer to as the *base-height index*. In the base-height index, we pick one of the d spatial dimensions and denote it as *height*, using index $h \in \mathbb{Z}$, while the other d space-time dimensions form the *base*, which will be indexed by $b \in \mathbb{Z}^d$. Then, a *base-height cell* will be indexed by $(b, h) \in \mathbb{Z}^{d+1}$.

We analogously define the *base-height super cell* (b, h) to be the space-time super cell (i, τ), for which the base-height cell (b, h) corresponds to the space-time cell (i, τ). Similarly, we define $E_{\text{bh}}(b, h)$, the increasing event restricted to the super cell (b, h), to be the same as the event $E_{\text{st}}(i, \tau)$ for the space-time cell (i, τ) that corresponds to the base-height cell (b, h).

Two-sided Lipschitz surface. Let a function $F : \mathbb{Z}^d \to \mathbb{Z}$ be called a *Lipschitz function* if $|F(x) - F(y)| \leq 1$ whenever $\|x - y\|_1 = 1$.

Definition 2.4. A *two-sided Lipschitz surface* F is a set of base-height cells $(b, h) \in \mathbb{Z}^{d+1}$ such that for all $b \in \mathbb{Z}^d$ there are exactly two (possibly equal) integer values $F_+(b) \geq 0$ and $F_-(b) \leq 0$ for which $(b, F_+(b)), (b, F_-(b)) \in F$ and, moreover, F_+ and F_- are Lipschitz functions.

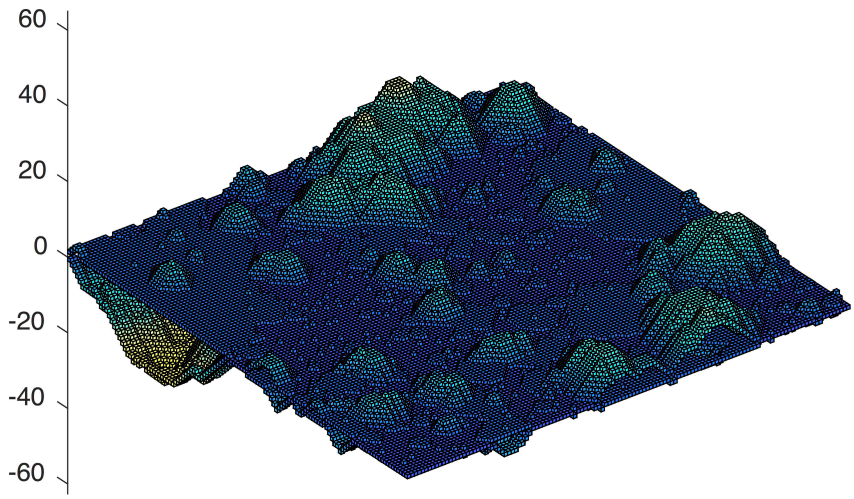

Figure 1: A two-sided Lipschitz surface for the case of $\mathbb{Z}^3$.

An illustration of F for $d = 2$ is given in Figure 1. We say a space-time cell (i, τ) belongs to F if the corresponding base-height cell (b, h) belongs to F. We say a two-sided Lipschitz surface F *exists*, if for all $b \in \mathbb{Z}^d$, we have $F_+(b) < \infty$ and $F_-(b) > -\infty$. For any positive integer D, we say a two-sided Lipschitz surface *surrounds* a cell (b', h') at distance D if any path $(b', h') = (b_0, h_0), (b_1, h_1), \ldots, (b_n, h_n)$ for which

$\|(b_i, h_i) - (b_{i-1}, h_{i-1})\|_1 = 1$ for all $i \in \{1, \dots n\}$ and $\|(b_n, h_n) - (b_0, h_0)\|_1 > D$, intersects with F.

Results. For any $z \in \mathbb{Z}_+$, let $Q_z = [-z/2, z/2]^d$. The following theorem establishes the existence of the Lipschitz surface.

Theorem 2.1. *Let (G, μ) be a uniformly elliptic conductance graph on the lattice $\mathbb{Z}^d$ for $d \geq 2$. There exist positive constants c_0, c_1 and c_2 such that the following holds. Tessellate G in space-time cells and super cells as described above for some $\ell, \beta, \eta > 0$ such that the ratio $\beta/\ell^2 < c_0$. Let $E_{\text{st}}(i, \tau)$ be an increasing event, restricted to the space-time super cell (i, τ). Fix $\epsilon \in (0, 1)$ and fix w such that*

$$w \geq \sqrt{\frac{\eta\beta}{c_2\ell^2} \log\left(\frac{8c_1}{\epsilon}\right)}.$$

Then, there exists a positive number α_0 that depends on ϵ, η and the ratio β/ℓ^2 so that if

$$\min\left\{C_M^{-1}\epsilon^2\lambda_0\ell^d, \log\left(\frac{1}{1 - \nu_{E_{\text{st}}}((1-\epsilon)\lambda, Q_{(2\eta+1)\ell}, Q_{w\ell}, \beta)}\right)\right\} \geq \alpha_0, \tag{2}$$

a two-sided Lipschitz surface F where $E_{\text{st}}(i, \tau)$ holds for all $(i, \tau) \in F$ almost surely exists.

We now briefly explain the main conditions for the establishment of the above theorem. We usually fix β/ℓ^2 to be an arbitrary, but small constant. The value of η defines the super cubes, which just model how much overlap we need between the cells of the tessellation (usually to allow information to propagate from one cell to its neighbors). Once these two parameters are fixed, we need to satisfy (2). First we need $C_M^{-1}\epsilon^2\lambda_0\ell^d \geq \alpha_0$. After fixing ϵ, this can be satisfied either by setting ℓ large enough (which makes the cells of the tessellation large), or by assuming that the density of particles λ_0 is large enough. Then we still need to make $\nu_{E_{\text{st}}}((1-\epsilon)\lambda, Q_{(2\eta+1)\ell}, Q_{w\ell}, \beta) \geq 1 - \exp(-\alpha_0)$. Usually E_{st} is a local event that becomes more and more likely by setting ℓ larger and larger; so having ℓ large enough suffices to satisfy this condition as well. The value of $\epsilon > 0$ is introduced so that in $\nu_{E_{\text{st}}}$ we can consider a Poisson point process of particles of intensity measure $(1-\epsilon)\lambda$, slightly smaller than the actual intensity of particles. This slack is needed to restrict our attention to the particles that "behave well". Then the lower bound on w is to guarantee that, as particles move in $Q_{(2\eta+1)\ell}$ for time β, with high probability they do not leave $Q_{(2\eta+1)\ell+w\ell}$, allowing a better control of dependences between neighboring cells of the tessellation. The proof of Theorem 2.1 is given in Section 7. With some additional work, which we do in Section 8, we can establish the following property of F.

Theorem 2.2. *Assume the conditions of Theorem 2.1 are satisfied. There exist positive constants c and C such that, for any sufficiently large $r > 0$, we have*

$$\mathbb{P}\left[\begin{array}{c} F \text{ does not surround} \\ \text{the origin at distance } r \end{array}\right] \leq \begin{cases} \sum_{s\geq r} s^d \exp\{-C\lambda_0 \frac{\ell s}{(\log \ell s)^c}\}, & \text{for } d = 2 \\ \sum_{s\geq r} s^d \exp\{-C\lambda_0 \ell s\}, & \text{for } d \geq 3. \end{cases}$$

The way Theorem 2.2 is proved also gives that the parts of the two-sided Lipschitz surface where the two sides F_+ and F_- intersect not only almost surely separate the origin from infinity within the "zero-height hyperplane" $\mathbb{L} = \mathbb{Z}^d \times \{0\}$, but they even percolate within $\mathbb{L}$. We say that the two-sided Lipschitz surface percolates within $\mathbb{L}$ if the set $\mathbb{L} \setminus F$ contains only finite connected components.

Theorem 2.3. *Assume the conditions of Theorem 2.1 are satisfied. If in addition we have that ℓ is sufficiently large and $\mathbb{P}[E_{\mathrm{st}}(0,0)]$ is sufficiently large, then the zero-height cluster $F \cap \mathbb{L}$ of the two-sided Lipschitz surface F percolates within $\mathbb{L}$ almost surely.*

Remark 2.1. In the definition of the base-height index, we fixed height to correspond to one of the spatial dimensions. This is the natural setting for the application of this Lipschitz surface technique to all applications we have in mind, for example the ones in [5]. However, in the definition of the surface we could have let height correspond to the *time* dimension. Then, Theorems 2.1, 2.2 and 2.3 hold for $d \geq 3$, but they no longer hold for $d = 2$. See Remark 5.2 in Section 5 for details.

The remainder of this paper is structured as follows. In Section 3 we give a construction of the two-sided Lipschitz surface for site percolation. Section 4 introduces multiple scales of the tessellation and Section 5 generalizes the paths defined in the construction from Section 3 to this multi-scale framework. Section 6 ties together the results from the previous sections, which is then applied in Section 7 to prove Theorem 2.1. Section 8 extends the results to a larger class of paths, which let us control areas where the two sides of the Lipschitz surface have non-zero height, in order to prove Theorems 2.2 and 2.3.

3 Two-sided Lipschitz surface in percolation

In this section we show how to construct the Lipschitz surface F given a realization of the events $E_{\mathrm{bh}}(b,h)$, $(b,h) \in \mathbb{Z}^{d+1}$, from Section 2. For this, we regard $(E_{\mathrm{bh}}(b,h))_{(b,h)\in\mathbb{Z}^{d+1}}$ as a site percolation process on $\mathbb{Z}^{d+1}$ so that a site $(b,h) \in \mathbb{Z}^{d+1}$ is considered to be *open* iff $E_{\mathrm{bh}}(b,h)$ holds and *closed* otherwise. We assume that the $E_{\mathrm{bh}}(b,h)$ are translation invariant. The concept of Lipschitz percolation for independent Bernoulli percolation was introduced and studied in [4, 6]. We modify their approach as we need several additional properties from the surface, such as the surface being two sided (i.e., composed of two sheets), the surface being close enough to the zero-height hyperplane $\mathbb{L}$, and the two sides of the surface intersecting in several points in $\mathbb{L}$.

The construction of F is based on the definition of a special type of paths, which we call d-paths. The definition of d-paths is based on a few rules. The first is that d-paths only start from *closed* sites at height 0 (i.e., closed sites of $\mathbb{L}$). For $x \in \mathbb{Z}$, define the set $\mathrm{Sign}(x)$ as $\{+1\}$ if $x > 0$, $\{-1\}$ if $x < 0$, and $\{-1, +1\}$ if $x = 0$. A d-path from a closed site $u \in \mathbb{L}$ to a not necessarily closed site $v \in \mathbb{Z}^{d+1}$ is any finite sequence of distinct sites $u = (b_0, 0), (b_1, h_1), \dots, (b_k, h_k) = v$ of $\mathbb{Z}^{d+1}$ such that for each $i = 1, 2, \dots, k$ we have that either (3) or (4) below hold:

$$b_i = b_{i-1},\ h_i - h_{i-1} \in \mathrm{Sign}(h_{i-1}) \text{ and } (b_i, h_i) \text{ is a closed site,} \tag{3}$$

or

$$\|b_i - b_{i-1}\|_1 = 1,\ h_{i-1} - h_i \in \mathrm{Sign}(h_{i-1}) \text{ and } h_{i-1} \neq 0. \tag{4}$$

We say the i-th move of a d-path is *vertical* if it is like (3), otherwise we say the i-th move is *diagonal*. Note that in a vertical move, the path moves away from $\mathbb{L}$, while in a diagonal move it moves towards $\mathbb{L}$. Moreover, unlike a vertical move, a diagonal move is not required to go into a closed site and cannot be performed from a site of $\mathbb{L}$.

In order to avoid issues of parity, we define for $(b, h) \in \mathbb{Z}^{d+1}$ the set of all sites that have the same base as (b, h), but are further away from $\mathbb{L}$.

$$\widehat{(b,h)} := \left\{(b, h') \in \mathbb{Z}^{d+1} :\ \tfrac{h'}{h} \geq 1\right\}.$$

For $u \in \mathbb{L}$ and $v \in \mathbb{Z}^{d+1}$, we denote by $u \rightarrowtail_d v$ the event that there is a d-path from u to at least one site of $\hat{v}$. We say v is *reachable* from u when this event holds.

We now define several sets of sites and some corresponding values, which will let us construct the desired two-sided Lipschitz surface.

Definition 3.1. The *hill* around $u \in \mathbb{L}$ is the set of all sites that are reachable from u,

$$H_u := \{v \in \mathbb{Z}^{d+1} :\ u \rightarrowtail_d v\}.$$

The *mountain* around $v \in \mathbb{L}$ is the union of all hills that contain v

$$M_v = \bigcup_{u : v \in H_u} H_u.$$

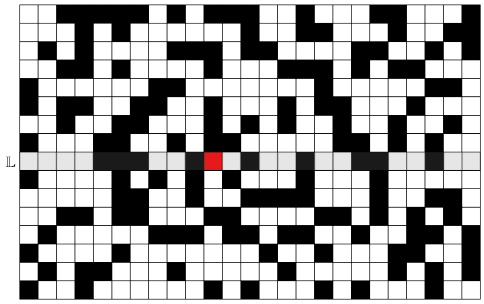

(a) Site percolation with the closed site $u \in \mathbb{L}$ marked in red.

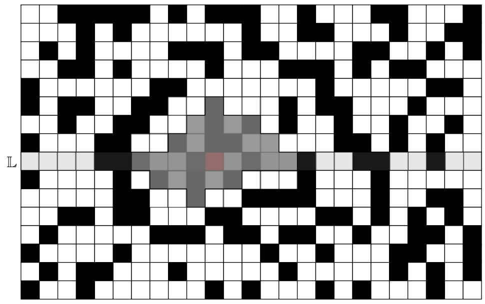

(b) The hill H_u of all sites that can be reached from u overlaid in gray.

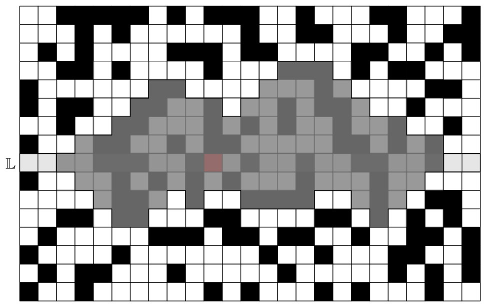

(c) The mountain M_u of all hills that contain u overlaid in gray.

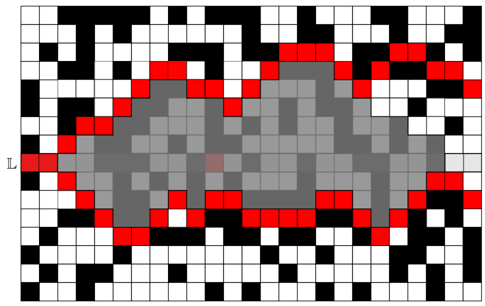

(d) A possible realization of the Lipschitz surface F in red (cf. Definition 3.3).

Figure 2: Examples of H_u and M_u of a chosen site u for site percolation on $\mathbb{Z}^2$. Open sites are white and closed sites are black.

Note that the sets H_u and M_v can be empty; in particular, $H_u = \emptyset$ if u is an open site. We define the positive and negative *depths* of a set $S \subset \mathbb{Z}^{d+1}$ at site $u = (b, h) \in \mathbb{Z}^{d+1}$ as

$$l_u^+(S) = \sup\{k : (b, h + k) \in S\},$$

and

$$l_u^-(S) = \sup\{k : (b, h - k) \in S\}.$$

Define also the *radius* of a set $S \subset \mathbb{Z}^{d+1}$ around u as

$$\mathrm{rad}_u(S) = \sup\{\|v-u\|_1 : v \in S\}.$$

We are now ready to define our two-sided Lipschitz surface; see Figure 2 for an illustration of H_u, M_u and F, and Figure 1 for an example of F in three dimensions.

Definition 3.2. For $u \in \mathbb{L}$ define

$$F_+(u) = \begin{cases} 1 + l_u^+(M_u) & \text{if } M_u \neq \emptyset \\ 0 & \text{if } M_u = \emptyset, \end{cases}$$

and

$$F_-(u) = \begin{cases} -1 - l_u^-(M_u) & \text{if } M_u \neq \emptyset \\ 0 & \text{if } M_u = \emptyset. \end{cases}$$

Definition 3.3. The two-sided Lipschitz surface F is defined as the set of sites

$$\bigcup_{b \in \mathbb{Z}^d} (b, F_-(b)) \cup (b, F_+(b)).$$

Note that the Lipschitz surface "envelops" the union of mountains $\bigcup_{u \in \mathbb{L}} M_u$. By definition, if $l_u^{\pm}(M_u)$ is infinite for some u, then it is infinite for all u (because of the diagonal moves of d-paths). Thus it is sufficient to show that $l_0^{\pm}(M_0)$ is finite almost surely in order to guarantee the existence of F. The theorem below establishes that F is finite almost surely; its proof follows along the lines of [6, Theorem 1].

Theorem 3.1. *For any $d \geq 1$, if $(E_{bh}(b,h))_{(b,h) \in \mathbb{Z}^{d+1}}$ is translation invariant and*

$$\sum_{r \geq 1} r^d \mathbb{P}\left[\mathrm{rad}_0(H_0) > r\right] < \infty, \tag{5}$$

then there exist almost surely a two sided Lipschitz surface F as in Definition 3.3. Moreover, the functions F_+ and F_- from Definition 3.2 satisfy

1. *For each $u = (b,0) \in \mathbb{L}$, the sites $(b, F_+(u))$ and $(b, F_-(u))$ are open.*
2. *For any $u, u' \in \mathbb{L}$ with $\|u - u'\|_1 = 1$, we have $|F_+(u) - F_+(u')| \leq 1$ and $|F_-(u) - F_-(u')| \leq 1$.*

Proof. We start by showing item 1. First, suppose that $M_0 \neq \emptyset$, and assume the opposite, i.e. that the site $(b, F_+(u))$ is closed. By the definition of the function l_u^+, the site $(b, l_u^+(M_u))$ belongs to M_u. Then, since $F_+(u) = 1 + l_u^+(M_u)$ and $M_u \neq \emptyset$, we can extend the d-path reaching the site $(b, l_u^+(M_u))$ with a vertical move into the closed site $(b, F_+(u))$. This gives that $(b, F_+(u)) \in M_u$, which is in contradiction with the construction of F_+. When $M_u = \emptyset$, we have $H_u = \emptyset$ and the site $(b, F_+(u)) = (b, 0)$ is open by definition. The proof for $(b, F_-(u))$ is similar.

Next, we establish item 2. Let $u, u' \in \mathbb{L}$ with $\|u - u'\|_1 = 1$. To show that $|F_+(u) - F_+(u')| \leq 1$, it is enough to show that $F_+(u') - F_+(u) \geq -1$ since the roles of u and u' are symmetric. Assume the converse, that is, that $F_+(u) \geq F_+(u') + 2$. Write $u = (b, 0)$ and $u' = (b', 0)$. We have by Definition 3.2 that $(b, F_+(u)-1) \in M_u$, so the site $(b, F_+(u)-1)$ can be reached by some d-path from $\mathbb{L}$. Extending this path by a diagonal move, we have that the site $(b', F_+(u) - 2) \in M_u$. Since $(b', F_+(u) - 2) \in \widehat{(b', F_+(u'))}$ by our assumption, we obtain that $(b', F_+(u')) \in M_u$, contradicting the construction of F_+. The proof for F_- is similar.

Finally, we prove the almost sure existence of F_+, that is, that $l_0^+(M_0)$ is almost surely finite. Because of the diagonal moves we have that $l_0(M_0) \leq \mathrm{rad}_0(M_0)$, so we only need to show that $\mathrm{rad}_0(M_0) < \infty$. By translation invariance we have

$$\begin{aligned}\mathbb{P}[\mathrm{rad}_0(M_0) \geq r] &\leq \sum_{v \in \mathbb{L}} \mathbb{P}[0 \in H_v,\, \mathrm{rad}_v(H_v) \geq r - \|v\|_1] \\ &= \sum_{v \in \mathbb{L}} \mathbb{P}[v \in H_0,\, \mathrm{rad}_0(H_0) \geq r - \|v\|_1]\end{aligned}$$

The last sum can be split into two sums depending on whether or not $\|v\|_1 \leq r/2$. In the first case, the sum is no larger than $cr^d\mathbb{P}[\mathrm{rad}_0(H_0) \geq r/2]$ for some constant c, and by (5) this term goes to 0 as r increases. Since $\{v \in H_0\} \subseteq \{\mathrm{rad}_0(H_0) \geq \|v\|_1\}$, we can bound the sum for which $\|v\|_1 > r/2$ by

$$\sum_{\substack{v \in \mathbb{L} \\ \|v\|_1 \geq r/2}} \mathbb{P}[v \in H_0] \leq \sum_{s \geq r/2} Cs^d\mathbb{P}[\mathrm{rad}_0(H_0) \geq s],$$

where $C > 0$ is a constant that depends only on d. By (5) this term also goes to 0 as r increases, which concludes the proof. □

4 Multi-scale setup

In light of Theorem 3.1, the key in establishing the existence of the Lipschitz surface is to control the radius of H_0. To do this, we look at all paths starting from 0 and the probability that they are a d-path. The challenge is that the event that a given cell (b, h) is bad is not independent of other space-time cells. To solve this problem we resort to a multi-scale approach. Before defining the multi-scale tessellation, we need a result regarding *local mixing* of particles, which we will use to link cells from one scale to the next.

4.1 Local mixing

Let $G = (\mathbb{Z}^d, E)$ be the d-dimensional square lattice equipped with conductances $(\mu_{x,y})_{(x,y)\in E}$ satisfying (1). The next theorem shows that if particles are dense enough inside a large cube $Q_K = [-K/2, K/2]^d$, then after particles move for some time, their distribution inside Q_K (but away from Q_K's boundary) dominates an independent Poisson point process.

Theorem 4.1 ([5, Theorem 4.1])**.** *Let $\mu_{x,y}$ satisfy (1) for some constant C_M and $c > 0$ be an arbitrary constant. There exist positive constants c_0, c_1, C and Θ such that the following holds. Fix $K > \ell > 0$ and $\epsilon \in (0, 1)$. Consider the cube Q_K tessellated into subcubes $(T_i)_i$ of side length ℓ. Suppose that at time 0 there is a collection of particles in Q_K with each subcube T_i containing at least $\sum_{y\in T_i} \beta\mu_y > c$ particles for some $\beta > 0$ and that ℓ is sufficiently large for this to be possible. Let $\Delta \geq c_0\ell^2\epsilon^{-4/\Theta}$. Fix $K' > 0$ such that $K - K' \geq c_1\sqrt{\Delta \log \Delta}$. For each j, denote by Y_j the location of the j-th particle of the collection at time Δ, conditioned on having displacement in $Q_{K-K'}$ during $[0, \Delta]$. Then there exists a coupling $\mathbb{Q}$ of an independent Poisson point process ψ with intensity measure $\zeta(y) = \beta(1-\epsilon)\mu_y$, $y \in Q_{K'}$, and $(Y_j)_j$ such that ψ is a subset of $(Y_j)_j$ with probability at least*

$$1 - \sum_{y\in Q_{K'}} \exp\left\{-C\beta\mu_y\epsilon^2\Delta^{d/2}\right\}.$$

4.2 Tessellation

We start by tessellating space at multiple scales. Let $m > 0$ be a sufficiently large integer, Θ be the constant from Theorem 4.1 and let $\epsilon \in (0, 1)$. To simplify the notation, we assume that $\frac{1}{\Theta}$ is an integer; otherwise we could work with $\lceil \frac{1}{\Theta} \rceil$. For each scale $k \geq 1$ we will tessellate the graph $G = (\mathbb{Z}^d, E)$ into cubes of length ℓ_k such that

$$\ell_1 = \ell \quad \text{and} \quad \ell_k = mk^a\ell_{k-1} = m^{k-1}(k!)^a\ell,$$

where a is a large integer we will set later. Set also $\ell_0 = \ell/m$.

We index the cubes by integer vectors $i \in \mathbb{Z}^d$ and denote them by $S_k(i)$. Then, for

$i = (i_1, i_2, \dots, i_d)$ we have

$$S_k(i) = \prod_{j=1}^{d} [i_j \ell_k, (i_j + 1)\ell_k].$$

This makes $S_k(i)$ the union of $(mk^a)^d$ cubes of scale $k - 1$. Next, we introduce the following hierarchy. For $k, j \geq 0$ and $i \in \mathbb{Z}^d$ we define

$$\pi_k^{(j)}(i) = i' \quad \text{iff} \quad S_k(i) \subseteq S_{k+j}(i').$$

We say $(k + 1, i')$ is the *parent* of (k, i) if $\pi_k^{(1)}(i) = i'$ and in this case also say (k, i) is a *child* of $(k + 1, i')$. We define the set of descendants of (k, i) as (k, i) and the union of all the descendants of the children of (k, i) or as only (k, i) in the case (k, i) has no children.

We introduce a new variable n that satisfies

$$n^d = \frac{m}{7\eta} \tag{6}$$

and impose the requirement on m to be large enough to yield $n > 1$. We also assume m is specified in such a way that n is an integer. Recall that η is the parameter introduced in the definition of super cells in the tessellation of Section 2, and that $\eta \geq 1$ is an integer.

We define some larger cubes based on $S_k(i)$. For $k \geq 0$ define the *base* and the *area of influence* of $S_k(i)$ as, respectively,

$$S_k^{\text{base}}(i) = \bigcup_{i' : \|i - i'\|_\infty \leq \eta mn(k+1)^a} S_k(i') \quad \text{and} \quad S_k^{\text{inf}}(i) = \bigcup_{i' : \|i - i'\|_\infty \leq 2\eta mn(k+1)^a} S_k(i').$$

For $k \geq 1$ we also define the *extended* cube

$$S_k^{\text{ext}}(i) = \bigcup_{i' : \pi_{k-1}^{(1)}(i') = i} S_{k-1}^{\text{base}}(i').$$

Observe that $S_k^{\text{ext}}(i)$ is the union of the bases of the children of (k, i), which are the $(k - 1)$-cubes contained in $S_k(i)$. We can see that $S_k(i) \subset S_k^{\text{base}}(i) \subset S_k^{\text{inf}}(i)$ and

$$S_{k+1}^{\text{ext}}(\pi_k^{(1)}(i)) = \bigcup_{i' : \pi_k^{(1)}(i') = \pi_k^{(1)}(i)} S_k^{\text{base}}(i') \supset S_k^{\text{base}}(i). \tag{7}$$

Remark 4.1. An important property derived from these definitions is that an extended cube of scale 1 has side length $\ell + 2\eta m n \ell_0 = (1+2\eta n)\ell$. Therefore, for any $i \in \mathbb{Z}^d$, the extended cube $S_1^{\text{ext}}(i)$ contains the super cube i defined in the tessellation of Section 2.

Now, we define the multi-scale tessellation of time. Let

$$\epsilon_1 = \epsilon \quad \text{and} \quad \epsilon_k = \epsilon_{k-1} - \frac{\epsilon}{k^2} \quad \text{for all } k \geq 2.$$

Define also $\epsilon_0 = 2\epsilon$ for consistency. Let

$$\beta_k = C_{\text{mix}} \frac{\ell_{k-1}^2}{(\epsilon_{k-1} - \epsilon_k)^{4/\Theta}} = C_{\text{mix}} \frac{\ell_{k-1}^2 k^{8/\Theta}}{\epsilon^{4/\Theta}} \quad \text{for all } k \geq 1, \tag{8}$$

where $C_{\text{mix}} \geq 2^{4/\Theta} c_0$, and c_0, Θ are the constants from Theorem 4.1. For $k = 1$ we set $\beta = \beta_1 = C_{\text{mix}} \frac{(\ell/m)^2}{\epsilon^{4/\Theta}}$. Given β/ℓ^2 and ϵ, m can be set sufficiently large so that $C_{\text{mix}} \geq 2^{4/\Theta} c_0$. Observe that

$$\frac{\beta_{k+1}}{\beta_k} = \frac{\ell_k^2 (k+1)^{8/\Theta}}{\ell_{k-1}^2 k^{8/\Theta}} = m^2 k^{2a - 8/\Theta} (k+1)^{8/\Theta} \quad \text{for all } k \geq 1. \tag{9}$$

Now, for scale $k \geq 1$, we tessellate time into intervals of length β_k. We index the time intervals by $\tau \in \mathbb{Z}$ and denote them by $T_k(\tau)$, where

$$T_k(\tau) = [\tau\beta_k, (\tau+1)\beta_k).$$

We allow time to be negative and note that β_{k+1}/β_k is always an integer by (9) if a is chosen larger than $4/\Theta$, which gives that a time interval of scale k is contained in a time interval of scale $k+1$. We therefore assume from now on that a is an integer and sufficiently large for

$$2a - 8/\Theta > 1 \tag{10}$$

to hold.

Let (k, τ) refer to the time interval $T_k(\tau)$. We also introduce a hierarchy over time, but which is different than the one defined for the cubes. For all k and τ let $\gamma_k^{(0)}(\tau) = \tau$, and for $j \geq 1$, define

$$\gamma_k^{(j)}(\tau) = \tau' \quad \text{if} \quad \gamma_k^{(j-1)}(\tau)\beta_{k+j-1} \in T_{k+j}(\tau'+1).$$

For the time tessellation, if $\tau' = \gamma_k^{(1)}(\tau)$, then the interval at scale $k+1$ that contains $T_k(\tau)$ is $T_{k+1}(\tau'+1)$. For any $j' \leq j$, we have $\gamma_k^{(j)} = \gamma_{k+j'}^{(j-j')}(\gamma_k^{(j')})$. Thus, for $\tau, \tau' \in \mathbb{Z}$

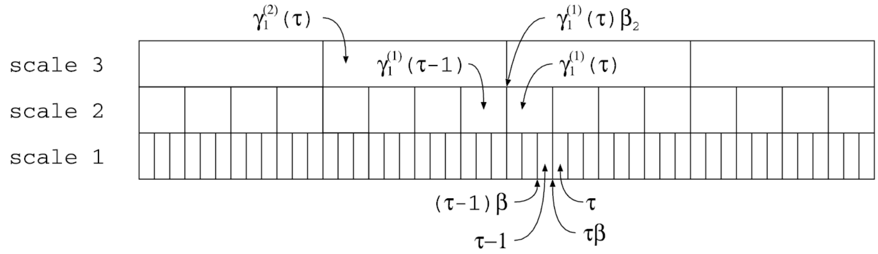

Figure 3: Time scale. The horizontal axis represents time and the vertical axis represents the scale. Note that $\gamma_1^{(1)}(\tau) = \gamma_1^{(1)}(\tau+1) = \gamma_1^{(1)}(\tau+2)$.

and $k \geq 1$ we say that $(k+1, \tau')$ is the *parent* of (k, τ), if $\gamma_k^{(1)}(\tau) = \tau'$; in this case we also say that (k, τ) is a *child* of $(k+1, \tau')$. We also define the set of descendants of (k, τ) as (k, τ) and the union of the descendants of the children of (k, τ) or only (k, τ) in the case (k, τ) has no children.

Now, for any $i \in \mathbb{Z}^d$, $k \geq 1$, $\tau \in \mathbb{Z}$, we define the space-time parallelogram

$$R_k(i, \tau) = S_k(i) \times T_k(\tau),$$

and note that these parallelograms are a tessellation of space and time. For $k = 1$ this is the same R_1 defined in the tessellation of Section 2.

We extend π and γ to a hierarchy of space and time. Then, letting (k, i, τ) refer to the space-time cell $S_k(i) \times T_k(\tau)$, we define the *descendants* of (k, i, τ) as the cells (k', i', τ') so that (k', i') is a descendant of (k, i) and (k', τ') is a descendant of (k, τ). We also say (k, i, τ) is an *ancestor* of (k', i', τ') if (k', i', τ') is a descendant of (k, i, τ).

4.3 A fractal percolation process

We now define the percolation process we will analyze. For the remainder of the paper, let $E(i, \tau) := \mathbb{1}_{E_{\mathrm{st}}(i,\tau)}$ be the indicator random variable of the increasing event $E_{\mathrm{st}}(i, \tau)$. For $k \geq 1$, define $S_k(i)$ to be *k-dense* at some time t if all $(\frac{\ell_k}{\ell_{k-1}})^d = (mk^a)^d$ cubes $S_{k-1}(i') \subset S_k(i)$ contain at least $(1-\epsilon_k)\lambda_0 \sum_{y \in S_{k-1}(i')} \mu_y$ particles at time t. For a cell (k, i, τ) let $D_k(i, \tau)$ be the indicator random variable such that

$$D_k(i, \tau) = 1 \quad \text{iff} \quad S_k(i) \text{ is } k\text{-dense at time } \tau\beta_k.$$

We also define a more restrictive indicator random variable:

> $D_k^{\text{ext}}(i,\tau) = 1$ iff, at time $\tau\beta_k$, all cubes $S_{k-1}(i')$ of scale $k-1$ contained in $S_k^{\text{ext}}(i)$ have at least $(1-\epsilon_k)\lambda_0 \sum_{y\in S_{k-1}(i')} \mu_y$ particles whose displacement throughout $[\tau\beta_k, (\tau+2)\beta_k]$ is in $Q_{\eta mnk^a \ell_{k-1}}$.

Recall the definition of the displacement of a particle from Definition 2.2. Then $D_k^{\text{ext}}(i,\tau) \le D_k(i,\tau)$ for all cells (k,i,τ).

Remark 4.2. An important property of this definition is that, when $D_k^{\text{ext}}(i,\tau) = 1$, if $(k-1,i',\tau')$ is a child of (k,i,τ), then we know that there are enough particles in $S_{k-1}^{\text{base}}(i')$ at time $\tau\beta_k$ and these particles never leave the cube $S_{k-1}^{\text{inf}}(i')$ during the interval $[\tau\beta_k, \tau'\beta_{k-1}]$. This will let us apply Theorem 4.1 to show that if $D_k^{\text{ext}}(i,\tau) = 1$, then $D_{k-1}^{\text{ext}}(i',\tau')$ is likely to be 1.

Define

> $D_k^{\text{base}}(i,\tau) = 1$ iff, at time $\gamma_k^{(1)}(\tau)\beta_{k+1}$, all cubes $S_k(i')$ of scale k inside $S_k^{\text{base}}(i)$ contain at least $(1-\epsilon_{k+1})\lambda_0 \sum_{y\in S_k(i')} \mu_y$ particles whose displacement throughout $[\gamma_k^{(1)}(\tau)\beta_{k+1}, \tau\beta_k]$ is in $Q_{\eta mn(k+1)^a \ell_k}$.

Note that if $D_{k+1}^{\text{ext}}(\pi_k^{(1)}(i), \gamma_k^{(1)}(\tau)) = 1$ then $D_k^{\text{base}}(i,\tau) = 1$. This gives that

$$D_k^{\text{base}}(i,\tau) \ge D_{k+1}^{\text{ext}}(\pi_k^{(1)}(i), \gamma_k^{(1)}(\tau)), \quad \text{for all } (k,i,\tau). \tag{11}$$

We next fix a scale κ as being the largest scale we will consider, and define

$$A_\kappa(i,\tau) = D_\kappa^{\text{ext}}(i,\tau).$$

For k satisfying $2 \le k \le \kappa - 1$, we set

$$A_k(i,\tau) = \max\left\{D_k^{\text{ext}}(i,\tau), 1 - D_k^{\text{base}}(i,\tau)\right\}.$$

For scale 1 we set

$$A_1(i,\tau) = \max\left\{E(i,\tau), 1 - D_1^{\text{base}}(i,\tau)\right\}.$$

Finally, define

$$A(i,\tau) = \prod_{k=1}^{\kappa} A_k(\pi_1^{(k-1)}(i), \gamma_1^{(k-1)}(\tau)). \tag{12}$$

Intuitively, a cell (k,i,τ) will be "well behaved" if $A_k(i,\tau) = 1$. More precisely, it follows from (11) that if $A_{k+1}(i',\tau') = 1$ and (k,i,τ) is a descendent of $(k+1,i',\tau')$, then $A_k(i,\tau) = 0$ if and only if $D_k^{\text{ext}}(i,\tau) = 0$ (or $E(i,\tau) = 0$ if $k=1$). On the other hand, $A_{k+1}(i',\tau') = 0$ implies that $D_{k+1}^{\text{ext}}(i',\tau') = 0$ and by (11) we have that $D_k^{\text{base}}(i,\tau) \ge 0$, so that $A_k(i,\tau) = 1$ if either $D_k^{\text{base}}(i,\tau) = 0$ or $D_k^{\text{ext}}(i,\tau) = 1$ (or

$E(i, \tau) = 1$ if $k = 1$). Therefore, $A_k(i, \tau)$ can be seen as the indicator of the event that the particles are "well behaved" in the cell (k, i, τ), given that they were well behaved in the ancestor cell of (k, i, τ). Finally, whenever $A_k(i, \tau) = 0$, it follows from (12) that all descendants $(1, i', \tau')$ of (k, i, τ) at scale 1 have $A(i', \tau') = 0$.

4.4 D-paths and bad clusters

Consider two distinct cells $(i, \tau), (i', \tau')$ of scale 1. We say that (i, τ) is *adjacent* to (i', τ') if $\|i - i'\|_\infty \leq 1$ and $|\tau - \tau'| \leq 1$. Also, we say that (i, τ) is *diagonally connected* to (i', τ') if there exists a sequence of cells $(i, \tau) = (b_0, h_0), (b_1, h_1), \ldots, (b_n, h_n) = (\hat{i}, \hat{\tau})$, where the indices (b_j, h_j) refer to the base-height index, such that all the following hold:

- for all $j \in \{1, \ldots, n\}$, $\|b_j - b_{j-1}\|_1 = 1$ and $h_{j-1} - h_j \in \text{Sign}(h_{j-1})$,
- $h_i h_j \geq 0$ for all $i, j \in \{0, \ldots, n\}$,
- $(\hat{i}, \hat{\tau})$ is adjacent to (i', τ') or $(\hat{i}, \hat{\tau}) = (i', \tau')$.

The definition of diagonally connected is in line with the definition of d-paths from Section 3, where paths can move diagonally towards $\mathbb{L}$ regardless of the status (open or closed) of the cells. We then define a *D-path* as a sequence of scale 1 cells where each cell is either adjacent or diagonally connected to the next cell in the sequence.

Recall also that a cell (i, τ) of scale 1 is denoted *bad* if $E_{\text{st}}(i, \tau)$ does not hold. Given a cell (i, τ) of scale 1, we define the *bad cluster* $K(i, \tau)$ as the set of cells (i', τ') of scale 1 that are bad and to which there exists a D-path from (i, τ) where all cells in the D-path are bad. We say that a cell (i, τ) of scale 1 has a *bad ancestry* if $A(i, \tau) = 0$ and in this case we define the *cluster of bad ancestries* as

$$K'(i, \tau) = \{(i', \tau') \in \mathbb{Z}^{d+1} : A(i', \tau') = 0 \text{ and } \exists \text{ a } D\text{-path from } (i, \tau) \text{ to } (i', \tau') \text{ where each cell of the path has a bad ancestry}\}.$$

Lemma 4.1. *For each cell (i, τ) of scale 1, we have that $E(i, \tau) \geq A(i, \tau)$. This implies that $K(i, \tau) \subseteq K'(i, \tau)$.*

Proof. Fix $(i, \tau) \in \mathbb{Z}^{d+1}$. Then, for $k = 1$, define $X_1 = E(i, \tau)$ and, for $k \geq 2$, define $X_k = D_k^{\text{ext}}(\pi_1^{(k-1)}(i), \gamma_1^{(k-1)}(\tau))$. Let $Y_k = D_k^{\text{base}}(\pi_1^{(k-1)}(i), \gamma_1^{(k-1)}(\tau))$. Therefore, by the definition of A in (12), we have

$$A(i,\tau) = \left(\prod_{k=1}^{\kappa-1} \max\{X_k, 1-Y_k\}\right) X_\kappa.$$

We now have $Y_k \geq X_{k+1}$ for all k. Therefore, for any $k \leq \kappa - 1$, we have

$$\max\{X_k, 1-Y_k\}X_{k+1} \leq \max\{X_k, 1-X_{k+1}\}X_{k+1} = X_k X_{k+1}.$$

Applying this repeatedly, we have

$$A(i,\tau) \leq \left(\prod_{k=1}^{\kappa-2} \max\{X_k, 1-Y_k\}\right) X_{\kappa-1}X_\kappa \leq \prod_{k=1}^{\kappa} X_k \leq X_1 = E(i,\tau).$$

□

4.5 The support of a cell

We define the *time of influence* $T_k^{\text{inf}}(\tau)$ of (k,τ) as

$T_1^{\text{inf}}(\tau) = [\gamma_1^{(1)}(\tau)\beta_2, (\tau+\max\{\eta,2\})\beta_1]$ and $T_k^{\text{inf}}(\tau) = [\gamma_k^{(1)}(\tau)\beta_{k+1}, (\tau+2)\beta_k]$ for $k \geq 2$,

and set the *region of influence* as

$$R_k^{\text{inf}}(i,\tau) = S_k^{\text{inf}}(i) \times T_k^{\text{inf}}(\tau).$$

We assume m is sufficiently large with respect to η so that $\max\{\eta,2\}\beta \leq \beta_2 = m^2 2^{8/\Theta}\beta$, which gives that

$$T_k^{\text{inf}}(\tau) \subseteq T_{k+1}(\gamma_k^{(1)}(\tau)) \cup T_{k+1}(\gamma_k^{(1)}(\tau)+1) \cup T_{k+1}(\gamma_k^{(1)}(\tau)+2) \tag{13}$$

We define the *time support* as

$$T_k^{\text{sup}}(\tau) = \bigcup_{i=0}^{8} T_{k+1}(\gamma_k^{(1)}(\tau) - 3 + i)$$

and the *spatial support* as

$$S_k^{\text{sup}}(i) = \bigcup_{i' : \|i' - \pi_k^{(1)}(i)\|_\infty \leq m} S_{k+1}(i'),$$

and, for any cell (k, i, τ), we define

$$R_k^{\sup}(i,\tau) = S_k^{\sup}(i) \times T_k^{\sup}(\tau)$$

Lemma 4.2. *For any sufficiently large m the following is true. For any cells (k, i, τ), (k', i', τ'), with $k \geq k'$, if $R_{k'}^{inf}(i',\tau') \not\subseteq R_k^{sup}(i,\tau)$ then $R_{k'}^{inf}(i',\tau') \cap R_k^{inf}(i,\tau) = \emptyset$.*

Proof. Note that, if $R_{k'}^{\inf}(i',\tau') \not\subseteq R_k^{\sup}(i,\tau)$, then either $T_{k'}^{\inf}(\tau') \not\subseteq T_k^{\sup}(\tau)$ or $S_{k'}^{\inf}(i') \not\subseteq S_k^{\sup}(i)$. We start with the case that $T_{k'}^{\inf}(\tau') \not\subseteq T_k^{\sup}(\tau)$ and show that this implies

$$T_{k'}^{\inf}(\tau') \cap T_k^{\inf}(\tau) = \emptyset,$$

which gives that $R_{k'}^{\inf}(i',\tau') \cap R_k^{\inf}(i,\tau) = \emptyset$.

Note that the interval $T_{k'}^{\inf}(\tau')$ has length at most $3\beta_{k'+1}$ by (13). Then, since $T_{k'}^{\inf}(\tau') \not\subseteq T_k^{\sup}(\tau)$,

$$T_{k'}^{\inf}(\tau') \cap [(\gamma_k^{(1)}(\tau) - 3)\beta_{k+1} + 3\beta_{k'+1}, (\gamma_k^{(1)}(\tau) + 6)\beta_{k+1} - 3\beta_{k'+1}] = \emptyset. \tag{14}$$

Using that $\beta_{k'} \leq \beta_k$, we get

$$\begin{aligned}
&[(\gamma_k^{(1)}(\tau) - 3)\beta_{k+1} + 3\beta_{k'+1}, (\gamma_k^{(1)}(\tau) + 6)\beta_{k+1} - 3\beta_{k'+1}] \\
&\quad \supseteq [\gamma_k^{(1)}(\tau)\beta_{k+1}, (\gamma_k^{(1)}(\tau) + 3)\beta_{k+1}] \\
&\quad = T_{k+1}(\gamma_k^{(1)}(\tau)) \cup T_{k+1}(\gamma_k^{(1)}(\tau) + 1) \cup T_{k+1}(\gamma_k^{(1)}(\tau) + 2) \\
&\quad \supseteq T_k^{\inf}(\tau),
\end{aligned}$$

where the last step follows from (13). This, together with (14), implies that $T_{k'}^{\inf}(\tau') \cap T_k^{\inf}(\tau) = \emptyset$.

For the spatial component, consider the case $S_{k'}^{\inf}(i') \not\subseteq S_k^{\sup}(i)$, for which we want to show that

$$S_{k'}^{\inf}(i') \cap S_k^{\inf}(i) = \emptyset.$$

Let $x_1, x_2, \ldots, x_d$ be defined so that $S_k(i) = \prod_{j=1}^d [x_j, x_j + \ell_k]$. Then, we can write

$$S_k^{\inf}(i) = \prod_{j=1}^d [x_j - 2\eta mn(k+1)^a \ell_k, x_j + \ell_k + 2\eta mn(k+1)^a \ell_k]. \tag{15}$$

Next, let $y_1, y_2, \ldots, y_d$ be defined so that $S_k^{\sup}(i) = \prod_{j=1}^d [y_j, y_j + (2m+1)\ell_{k+1}]$. Since $S_{k'}^{\inf}(i')$ is a cube of side length $(1 + 4\eta mn(k'+1)^a)\ell_{k'} \leq (1 + 4\eta mn(k+1)^a)\ell_k$ and

$S^{\inf}_{k'}(i')$ is not contained in $S^{\sup}_k(i)$, we have that

$$S^{\inf}_{k'}(i') \cap \prod_{j=1}^{d}[y_j+(1+4\eta mn(k+1)^a)\ell_k, y_j+(2m+1)\ell_{k+1}-(1+4\eta mn(k+1)^a\ell_k] = \emptyset. \quad (16)$$

Now, we use the fact that $m\ell_{k+1} \leq x_j - y_j \leq (m+1)\ell_{k+1} - \ell_k$ for all $j = 1, 2, \ldots, d$. This and (15) give

$$S^{\inf}_k(i) \subseteq \prod_{j=1}^{d}[y_j + m\ell_{k+1} - 2\eta mn(k+1)^a\ell_k, y_j + (m+1)\ell_{k+1} + 2\eta mn(k+1)^a\ell_k]. \quad (17)$$

Now, using the relation between m and n in (6), we have that

$$m\ell_{k+1} = m^2(k+1)^a\ell_k = 7\eta mn^d(k+1)^a\ell_k \geq (1 + 6\eta mn(k+1)^a)\ell_k. \quad (18)$$

Using this result in (16) we get that $S^{\inf}_{k'}(i')$ does not intersect

$$\prod_{j=1}^{d}[y_j + (1 + 4\eta mn(k+1)^a)\ell_k, y_j + (m+1)\ell_{k+1} + 2\eta mn(k+1)^a\ell_k]. \quad (19)$$

Similarly, plugging (18) into (17) we see that $S^{\inf}_k(i)$ is contained in the space-time region given by (19). These two facts establish the lemma. □

Another important property concerns the fact that the support of a cell contains all its descendants.

Lemma 4.3. *Assume $m \geq 3$. For any cell (k, i, τ), if (k', i', τ') is a descendant of (k, i, τ) then*

$$R_{k'}(i', \tau') \subseteq R^{sup}_k(i, \tau).$$

Moreover, $R^{sup}_k(i, \tau)$ contains all the neighbors of (k', i', τ').

Proof. Fix i'', τ'' such that (k', i'', τ'') is adjacent to (k', i', τ') and assume that the ancestor of (k', i'', τ'') of scale k is not (k, i, τ), otherwise the second part of the lemma follows from the first part. We prove this lemma first for space and then for time. For space, since (k', i', τ') is a descendant of (k, i, τ) we have that $S_{k'}(i') \subseteq S_k(i) \subseteq S^{\sup}_k(i)$. Also, (k', i'') is adjacent to (k', i') which implies that the ancestor of (k', i'') of scale k is adjacent to (k, i). Since $S^{\sup}_k(i)$ contains all cells of scale k that are adjacent to (k, i), it also contains $S_{k'}(i'')$.

We now prove the lemma for the time dimension. The first part corresponds to showing that $T_{k'}(\tau') \subseteq T_k^{\text{sup}}(\tau)$. Recall that $T_{k'}(\tau') = [\tau'\beta_{k'}, (\tau'+1)\beta_{k'}]$, which is contained in $[\tau\beta_k, (\tau'+1)\beta_{k'}]$ since (k', i', τ') is a descendant of (k, i, τ). Now, note that

$$\tau\beta_k = \gamma_{k'}^{(k-k')}(\tau')\beta_k \geq \gamma_{k'}^{(k-k'-1)}(\tau')\beta_{k-1} - 2\beta_k \geq \tau'\beta_{k'} - 2\sum_{i=k'+1}^{k} \beta_i.$$

Then, since $k' \geq 1$, we can use the bound

$$\sum_{i=2}^{k} \beta_i = C_{\text{mix}} \sum_{i=2}^{k} \frac{\ell_{i-1}^2 i^{8/\Theta}}{\epsilon^{4/\Theta}} = C_{\text{mix}}\epsilon^{-4/\Theta} \sum_{i=2}^{k} \ell_{i-1}^2 (i^{4/\Theta})^2 \leq C_{\text{mix}}\epsilon^{-4/\Theta} 2(k^{4/\Theta})^2 \ell_{k-1}^2 = 2\beta_k,$$

where the last inequality can be proven by induction on k. Then, we have that

$$\tau\beta_k \geq \tau'\beta_{k'} - 4\beta_k. \tag{20}$$

Since $k > k' \geq 1$, we have $k > 1$ and

$$4\beta_k + \beta_{k'} \leq 5\beta_k = 5\frac{\beta_{k+1}}{m^2 k^{2a-8/\Theta}(k+1)^{8/\Theta}} \leq \beta_{k+1}.$$

This combined with (20) gives

$$T_{k'}(\tau') \subseteq [\tau\beta_k, \tau\beta_k + 4\beta_k + \beta_{k'}] \subseteq [\tau\beta_k, \tau\beta_k + \beta_{k+1}] \subseteq T_k^{\text{sup}}(\tau).$$

This proves the first part of the lemma. To prove the second part, we use the fact that (k', τ'') is adjacent to (k', τ') and the result above, which gives

$$T_{k'}(\tau'') \subseteq [\tau\beta_k - \beta_{k'}, \tau\beta_k + \beta_{k+1} + \beta_{k'}] \subseteq T_k^{\text{sup}}(\tau). \qquad \square$$

5 Multi-scale analysis of D-paths

In order to prove our theorems we need to control the existence of D-paths of scale 1 whose cells have a bad ancestry (cf. Lemma 4.1). We will do this via a multi-scale analysis of such paths. In Section 4 we defined the multi-scale tessellation we need. Here we will use this framework to consider a multi-scale version of D-paths.

We start by defining the *extended support* of a cell. Given a cell (k, i, τ), define

$$T_k^{\text{2sup}}(\tau) = \bigcup_{i=0}^{26} T_{k+1}(\gamma_k^{(1)}(\tau) - 12 + i)$$

and

$$S_k^{2\mathrm{sup}}(i) = \bigcup_{i':\|i'-\pi_k^{(1)}(i)\|_\infty \leq 3m+1} S_{k+1}(i').$$

Then, as before, we set $R_k^{2\mathrm{sup}}(i,\tau) = S_k^{2\mathrm{sup}}(i) \times T_k^{2\mathrm{sup}}(\tau)$.

Remark 5.1. The extended support is defined in a way so that if the supports of two cells intersect, the smaller of the supports is completely contained in the extended support of the bigger of the cells.

We now extend the definition of a bad cell to multiple scales. We say that a cell (k,i,τ) is *multi-scale bad* if $A_k(i,\tau) = 0$. Recall that two cells (k,i,τ) and (k,i',τ') of the same scale are said to be adjacent if $\|i-i'\|_\infty \leq 1$ and $|\tau-\tau'| \leq 1$. Let $(k_1,i_1,\tau_1),(k_2,i_2,\tau_2)$ be two cells with $k_1 > k_2$. We say (k_1,i_1,τ_1) and (k_2,i_2,τ_2) are *adjacent* if (k_1,i_1,τ_1) is adjacent to $(k_1,\pi_{k_2}^{(k_1-k_2)}(i_2),\gamma_{k_2}^{(k_1-k_2)}(\tau_2))$. We say (k,i,τ) is *diagonally connected* to (k',i',τ') if there exists a cell $(1,\hat{i},\hat{\tau})$ that is a descendant of (k,i,τ) and a cell $(1,i'',\tau'')$ that is a descendant of (k',i',τ'), such that $(1,\hat{i},\hat{\tau})$ is diagonally connected to $(1,i'',\tau'')$.

We extend the definition of D-paths to cells of arbitrary scale by referring to a D-path as a sequence of distinct cells for which any two consecutive cells in the sequence are either adjacent or the first of the two cells is diagonally connected to the second. For any two cells (k_1,i_1,τ_1) and (k_2,i_2,τ_2) we say that they are *well separated* if $R_{k_1}(i_1,\tau_1) \not\subseteq R_{k_2}^{\mathrm{sup}}(i_2,\tau_2)$ and $R_{k_2}(i_2,\tau_2) \not\subseteq R_{k_1}^{\mathrm{sup}}(i_1,\tau_1)$. In order to ensure the cells we will look at are well separated but still not too far apart, we say that any two cells (k_1,i_1,τ_1) and (k_2,i_2,τ_2) are *support adjacent* if $R_{k_1}^{2\mathrm{sup}}(i_1,\tau_1) \cap R_{k_2}^{2\mathrm{sup}}(i_2,\tau_2) \neq \emptyset$. We say a cell (k_1,i_1,τ_1) is *support connected with diagonals* to (k_2,i_2,τ_2) if there exists a scale 1 cell contained in $R_{k_1}^{2\mathrm{sup}}(i_1,\tau_1)$ and a scale 1 cell contained in $R_{k_2}^{2\mathrm{sup}}(i_2,\tau_2)$, such that the former is diagonally connected to the latter.

Finally, define a sequence of cells $P = ((k_1,i_1,\tau_1),(k_2,i_2,\tau_2),\ldots,(k_z,i_z,\tau_z))$ to be a *support connected D-path* if the cells in P are mutually well separated and, for each $j = 1,2,\ldots,z-1$, (k_j,i_j,τ_j) is support adjacent or support connected with diagonals to $(k_{j+1},i_{j+1},\tau_{j+1})$.

For any t, define Ω_t to be the set of all D-paths of cells of scale 1 so that the first cell of the path is $(0,0)$ and the last cell of the path is the only cell not contained in the space-time region $[-t,t]^d \times [-t,t]$. Also, define $\Omega_{\kappa,t}^{\mathrm{sup}}$ as the set of all support connected D-paths of cells of scale at most κ so that the extended support of the first cell of the path contains $R_1(0,0)$ and the last cell of the path is the only cell whose extended support is not contained in $[-t,t]^d \times [-t,t]$. Then the lemma below shows that we can focus on support connected D-paths instead of D-paths with bad ancestry.

Lemma 5.1. *We have that*

$$\begin{aligned}&\mathbb{P}\left[\exists P \in \Omega_t \text{ s.t. all cells of } P \text{ have a bad ancestry}\right]\\ &\leq \mathbb{P}\left[\exists P \in \Omega_{\kappa,t}^{sup} \text{ s.t. all cells of } P \text{ are multi-scale bad}\right].\end{aligned}$$

Proof. We complete the proof in two stages. First, we show that if there exists a D-path $P \in \Omega_t$, such that each cell of P has bad ancestry, then there exists a D-path of multi-scale bad cells of arbitrary scales up to κ. Next, we show that, given the existence of such a path of multi-scale bad cells of arbitrary scales up to κ, there exists a support connected D-path of $\Omega_{\kappa,t}^{\text{sup}}$ such that all cells of the path are multi-scale bad.

Step 1: Let $\Omega_{\kappa,t}$ be the set of all D-paths of cells of scale at most κ such that the first cell of the path is an ancestor of $(0,0)$ and the last cell of the path is the only cell whose support is not contained in $[-t,t]^d \times [-t,t]$. We now establish that

$$\begin{aligned}&\mathbb{P}[\exists P \in \Omega_t \text{ s.t. all cells of } P \text{ have a bad ancestry}]\\ &\leq \mathbb{P}[\exists P \in \Omega_{\kappa,t} \text{ s.t. all cells of } P \text{ are multi-scale bad}].\end{aligned}$$

Let $P = ((1,i_1,\tau_1),(1,i_2,\tau_2),\ldots(1,i_z,\tau_z)) \in \Omega_t$ be a D-path of cells with bad ancestries; therefore $(i_1,\tau_1) = (0,0)$ and $R_1(i_z,\tau_z)$ is not contained in $[-t,t]^d \times [-t,t]$. For each j, since $A(i_j,\tau_j) = 0$, we know by definition of A in (12) that there exists a k'_j such that, if we set $i'_j = \pi_1^{(k'_j-1)}(i_j)$ and $\tau'_j = \gamma_1^{(k'_j-1)}(\tau_j)$, we obtain $A_{k'_j}(i'_j,\tau'_j) = 0$. Now construct $J \subseteq \{1,2,\ldots,z\}$, starting with $J = \{1,2,\ldots,z\}$ and removing elements from J iteratively as we go from scale $k = \kappa$ down to scale $k = 1$ using the following rule: if there exists $j \in J$ such that $k'_j = k$ and $A_{k'_j}(i'_j,\tau'_j) = 0$, then remove from J all descendants of (k'_j,i'_j,τ'_j), except for the first one; i.e. keep in J only the smallest j' for which $i'_j = \pi_1^{(k'_j-1)}(i_{j'})$ and $\tau'_j = \gamma_1^{(k'_j-1)}(\tau_{j'})$. Put differently, J contains only distinct elements of the set $\{(k'_j,i'_j,\tau'_j) : j = 1,2,\ldots z\}$ which have no ancestor within J. With this set we define

$$\tilde{P} = \{(k'_j,i'_j,\tau'_j) :\ j \in J\},$$

and show that $\tilde{P}$ is a D-path. This gives us the existence of a D-path of multi-scale bad cells of arbitrary scales starting from an ancestor of $(1,i_1,\tau_1)$ and such that the last cell $(k',i',\tau') \in \tilde{P}$ is an ancestor of $(1,i_z,\tau_z)$, which is not contained in $[-t,t]^d \times [-t,t]$. Lemma 4.3 then gives us that $R_1(i_z,\tau_z)$ is contained in $R_{k'}^{\text{sup}}(i',\tau')$ so that the union of the supports of the cells in $\tilde{P}$ is not contained in $[-t,t]^d \times [-t,t]$. We note that it is possible that the support of some other cell of $\tilde{P}$ is also not contained in $[-t,t]^d \times [-t,t]$.

In this case, we modify $\tilde{P}$ and remove from it all $j' \in J$ for which $j' > j$, where (k_j, i_j, τ_j) is the first cell of $\tilde{P}$ for which $R^{\text{sup}}_{k_j}(i_j, \tau_j)$ is not contained in $[-t,t]^d \times [-t,t]$. Furthermore, it is possible that $\tilde{P}$ might contain loops. This does not cause any issues; in fact, the procedure in step 2 will remove any loops should they exist

Now it remains to verify that $\tilde{P}$ satisfies the adjacency properties of a D-path. By construction, each cell of P has exactly one ancestor in $\tilde{P}$. If we take two adjacent cells $(1, i_j, \tau_j)$, $(1, i_{j+1}, \tau_{j+1})$ of P, they either have the same ancestor in $\tilde{P}$ or their ancestors are adjacent. This follows from the fact that two non-adjacent cells cannot have descendants of scale 1 that are adjacent. Now assume that $(1, i_j, \tau_j) \in P$ is diagonally connected to $(1, i_{j+1}, \tau_{j+1}) \in P$. In this case, the two cells either have the same ancestor in $\tilde{P}$, have ancestors that are adjacent or the ancestor of $(1, i_j, \tau_j)$ is diagonally connected to the ancestor of $(1, i_{j+1}, \tau_{j+1})$.

Step 2: Here we establish that

$$\begin{aligned}&\mathbb{P}[\exists P \in \Omega_{\kappa,t} \text{ s.t. all cells of } P \text{ are multi-scale bad}]\\ &\leq \mathbb{P}[\exists P \in \Omega^{\text{sup}}_{\kappa,t} \text{ s.t. all cells of } P \text{ are multi-scale bad}].\end{aligned}$$

Let $P = ((k_1, i_1, \tau_1), (k_2, i_2, \tau_2), \dots, (k_z, i_z, \tau_z)) \in \Omega_{\kappa,t}$ be a D-path of multi-scale bad cells. We will now show the existence of a support connected D-path P' of multi-scale bad cells. First, we order the cells of P in the following way. If two cells have the same scale, we order them by taking in the same order as they have in P. For two cells of different scales, we say the cell with the larger scale comes before the other cell. This gives us a total order of the cells of P. Next, let L be the list of cells of P following this order. We construct P' step-by-step, by adding the first element of L to P' and removing some elements from L, repeating this until L is empty. While doing this, we associate each cell of P to a cell of P', which we will later use to show that using the ordering inherited from P, P' is a support connected D-path. Assuming (k', i', τ') is the current first element of L, the steps taken to construct P' are as follows:

1. Add (k', i', τ') to P' and remove it from L. Associate (k', i', τ') in P with itself in P'.
2. Remove from L all cells (k'', i'', τ'') that are not well separated from (k', i', τ') and associate them to (k', i', τ').

We repeat these steps until L is empty. We highlight that by construction P' contains only mutually well separated cells. Let (k^*, i^*, τ^*) be the cell that (k_1, i_1, τ_1) is associated to. Note that the extended support of this cell contains $R_1(0,0)$, because $R^{\text{sup}}_{k_1}(i_1, \tau_1)$ contains $R_1(0,0)$ and is itself contained in the extended support of

(k^*, i^*, τ^*). We also obtain that

$$\bigcup_{(k',i',\tau')\in P'} R^{2\mathrm{sup}}_{k'}(i',\tau') \not\subseteq [-t,t]^d \times [-t,t].$$

This can be argued similarly as above, noting that the support of a cell in P was not contained in $[-t,t]^d \times [-t,t]$, so the extended support of the cell it is associated to cannot be contained either. Let the cell it is associated to be (k', i', τ').

Now it remains to show that there exists a subset of cells $P'' \subseteq P'$ which is a support connected D-path with diagonals and contains both (k^*, i^*, τ^*) and (k', i', τ'). To see this, we will add some cells from P' to P'', starting with (k', i', τ'). Let (k_j, i_j, τ_j) be the first cell of P that is associated to (k', i', τ'). If $j = 1$ then $(k^*, i^*, \tau^*) = (k', i', \tau')$ and P'' is just this cell. Otherwise, let (k'', i'', τ'') be the cell $(k_{j-1}, i_{j-1}, \tau_{j-1})$ is associated to and add it to P''. We will show later that

$$(k'', i'', \tau'') \text{ is support adjacent or support connected with diagonals to } (k', i', \tau'). \tag{21}$$

Now we iterate the procedure above; that is, we take the first cell $(k_\iota, i_\iota, \tau_\iota)$ of P that is associated to (k'', i'', τ'') and either finish the construction of P'' if $\iota = 1$ or continue by taking the cell that $(k_{\iota-1}, i_{\iota-1}, \tau_{\iota-1})$ is associated to and adding it to P''. Note that $\iota < j$, which guarantees that this procedure will eventually add (k^*, i^*, τ^*) to P'', thus completing the construction. It is possible that the extended support of some cell of P'' other than (k', i', τ') to not be contained in $[-t,t]^d \times [-t,t]$. As in step 1, we simply remove from P'' all cells (k_j, i_j, τ_j) that come after the first such cell in the original ordering of P.

It remains to show (21) holds. Assume for the following that $k_{j-1} \geq k_j$, the converse can be argued the same way, and recall that in P, two consecutive cells $(k_{j-1}, i_{j-1}, \tau_{j-1})$ and (k_j, i_j, τ_j) are either adjacent or the first cell is diagonally connected to the second cell. Since $k_{j-1} \geq k_j$, we have a cell $(\hat{k}, \hat{\imath}, \hat{\tau})$ at scale $\hat{k} = k_{j-1}$ that is an ancestor of (k_j, i_j, τ_j), to which $(k_{j-1}, i_{j-1}, \tau_{j-1})$ is either adjacent or diagonally connected. In the first case, we have by Lemma 4.3 that $R^{\mathrm{sup}}_{\hat{k}}(\hat{\imath}, \hat{\tau})$ contains both $R_{k_{j-1}}(i_{j-1}, \tau_{j-1})$ and $R_{k_j}(i_j, \tau_j)$. Since $\hat{k} = k_{j-1}$ and $(k_{j-1}, i_{j-1}, \tau_{j-1})$ is associated to (k'', i'', τ''), we have that $\hat{k} \leq k''$. Then, by Remark 5.1, we have that $R^{\mathrm{sup}}_{\hat{k}}(\hat{\imath}, \hat{\tau}) \subseteq R^{2\mathrm{sup}}_{k''}(i'', \tau'')$, which gives that $R^{2\mathrm{sup}}_{k''}(i'', \tau'')$ intersects $R^{2\mathrm{sup}}_{k'}(i', \tau')$. Alternatively, $(k_{j-1}, i_{j-1}, \tau_{j-1})$ is diagonally connected to (k_j, i_j, τ_j). This gives that (k'', i'', τ'') is support connected with diagonals to (k', i', τ') with the same diagonal steps that make $(k_{j-1}, i_{j-1}, \tau_{j-1})$ diagonally connected to (k_j, i_j, τ_j). □

The next lemma is a technical result bounding the probability that a random walk on a weighted graph remains inside a cube.

Lemma 5.2. *Let $\Delta > 0$ and, for any $z > 0$, define $F_\Delta(z)$ to be the event that a random walk on (G, μ) starting from the origin stays inside Q_z throughout the time interval $[0, \Delta]$. Then, on a uniformly elliptic graph, there exist constants c, c_1 and c_2 such that if $\Delta > cz$, we have*

$$\mathbb{P}[F_\Delta(z)] \geq 1 - c_1 \exp\{-c_2 z^2/\Delta\}.$$

Proof. The result is a reformulation of the exit time result for random walks on weighted graphs from [1] and [2] by taking a ball with radius $z/2$ that is contained in Q_z and using that the weights $\mu_{x,y}$ are uniformly elliptic. □

We now give a lemma that will be used to control the dependencies involving well separated cells. Let $\mathcal{F}_k(i, \tau)$ be the σ-field generated by all $A_{k'}(i', \tau')$ for which $T_{k'}^{\inf}(\tau')$ does not intersect $[\gamma_k^{(1)}(\tau)\beta_{k+1}, \infty)$ or both $\tau'\beta_{k'} \leq \tau\beta_k$ and $S_k^{\inf}(i) \cap S_{k'}^{\inf}(i') = \emptyset$. We define the following two quantities:

$$\psi_1 = \min\left\{\epsilon^2\lambda_0\ell^d C_M^{-1}, \log\left(\frac{1}{1 - \nu_{E_{\mathrm{st}}}((1-\epsilon)\lambda, Q_{(2\eta+1)\ell}, Q_{w\ell}, \eta\beta)}\right)\right\}$$
$$\psi_k = \frac{\epsilon^2\lambda_0\ell_{k-1}^d}{(k+1)^4} = \frac{\epsilon^2\lambda_0\ell^d m^{d(k-2)}((k-1)!)^{ad}}{(k+1)^4} \quad \text{for } k \geq 2. \tag{22}$$

Lemma 5.3. *Let $w \geq \sqrt{\frac{\eta\beta}{c_2\ell^2}\log\left(\frac{8c_1}{\epsilon}\right)}$ and*

$$\alpha = \min\left\{\epsilon^2\lambda_0\ell^d C_M^{-1}, \log\left(\frac{1}{1 - \nu_{E_{st}}((1-\epsilon)\lambda, Q_{(2\eta+1)\ell}, Q_{w\ell}, \eta\beta}\right)\right\}.$$

If m is sufficiently large with respect to d, β/ℓ^2, ϵ and C_M, then there are positive constants $c = c(C_M) \geq 1$ and α_0 so that, for all $\alpha \geq \alpha_0$, all cells (k, i, τ) and any $F \in \mathcal{F}_k(i, \tau)$, we have

1. $\mathbb{P}[A_k(i, \tau) = 0] \leq \exp\{-c\psi_k\}$*, for all* $k = 1, 2, \ldots, \kappa$
2. $\mathbb{P}[A_k(i, \tau) = 0 \mid F] \leq \exp\{-c\psi_k\}$*, for all* $k = 1, 2, \ldots, \kappa - 1$.

Proof. Note that A_k is defined differently for $k = 1$ and $2 \leq k \leq \kappa - 1$. We will first prove the result for $k \geq 2$ and establish part 2 of the lemma. Since

$$\mathbb{P}[A_k(i, \tau) = 0 \mid F] = \mathbb{P}[\{D_k^{\mathrm{ext}}(i, \tau) = 0\} \cap \{D_k^{\mathrm{base}}(i, \tau) = 1\} \mid F],$$

if $F \cap \{D_k^{\mathrm{base}}(i, \tau) = 1\} = \emptyset$, then the lemma holds. We now assume $F \cap \{D_k^{\mathrm{base}}(i, \tau) =$

$1\} \neq \emptyset$ and write

$$\mathbb{P}[A_k(i,\tau)=0 \mid F] \leq \mathbb{P}[\{D_k^{\text{ext}}(i,\tau)=0\} \mid F \cap \{D_k^{\text{base}}(i,\tau)=1\}].$$

Recall that $\{D_k^{\text{base}}(i,\tau)=1\}$ gives that all cubes $S_k(i')$ of scale k contained in $S_k^{\text{base}}(i)$ have at least $(1-\epsilon_{k+1})\lambda_0 \sum_{y\in S_k(i')} \mu_y$ particles at time $\gamma_k^{(1)}(\tau)\beta_{k+1}$ and the displacement of these particles throughout $[\gamma_k^{(1)}(\tau)\beta_{k+1}, \tau\beta_k]$ is in $Q_{\eta mn(k+1)^a \ell_k}$. Remember that F reveals only information about the location of these particles before time $\gamma_k^{(1)}(\tau)\beta_{k+1}$ since these particles never leave the cube $S_k^{\text{inf}}(i)$ during the whole $[\gamma_k^{(1)}(\tau)\beta_{k+1}, \tau\beta_k]$.

We now apply Theorem 4.1 and denote the variables appearing in the statement of that theorem with a bar. We apply the theorem with

$$\begin{aligned}
\bar{K} &= (1+2\eta mn(k+1)^a)\ell_k,\\
\bar{\ell} &= \ell_k,\\
\bar{\beta} &= (1-\epsilon_{k+1})\lambda_0,\\
\bar{\Delta} &= \tau\beta_k - \gamma_k^{(1)}(\tau)\beta_{k+1} \in [\beta_{k+1}, 2\beta_{k+1}],\\
\bar{K}' &\quad \text{such that } \bar{K}-\bar{K}' = \eta mn(k+1)^a\ell_k, \text{ and}\\
\bar{\epsilon} &\quad \text{such that } (1-\bar{\epsilon})(1-\epsilon_{k+1}) = \left(1-\tfrac{(\epsilon_{k+1}+\epsilon_k)}{2}\right).
\end{aligned}$$

This gives that $\bar{\epsilon} \geq \frac{\epsilon_k - \epsilon_{k+1}}{2} = \frac{\epsilon}{2(k+1)^2}$. Using these values and the fact that m is large enough, we have that

$$\bar{K}' = \ell_k + \eta mn(k+1)^a\ell_k \geq \ell_k + 2\eta mnk^a\ell_{k-1},$$

which is the side length of $S_k^{\text{ext}}(i)$. We also have $\bar{\Delta} \geq \beta_{k+1} \geq c_0\frac{\bar{\ell}^2}{\bar{\epsilon}^{4/\Theta}}$ since $C_{\text{mix}} \geq 2^{4/\Theta}c_0$ in the definition of β_{k+1}. We still have to check whether $\bar{K}-\bar{K}' \geq C\sqrt{\bar{\Delta}\log\bar{\Delta}}$, which is equivalent to checking that

$$\eta mn(k+1)^a\ell_k \geq \tilde{C}\sqrt{\beta_{k+1}\log\beta_{k+1}}$$

for some constant $\tilde{C}$. Using the definitions of ℓ_k and β_{k+1}, this inequality can be rewritten as

$$\eta mn(k+1)^a\ell_k \geq \tilde{C}\sqrt{C_{\text{mix}}}\frac{1}{\epsilon^{2/\Theta}}(k+1)^{4/\Theta}\ell_k\sqrt{\log\beta_{k+1}}.$$

Now, using the value of β_{k+1} and ℓ_k we obtain that there exists a constant C independent of k and m, but depending on ϵ such that

$$\frac{\tilde{C}\sqrt{C_{\mathrm{mix}}}}{\epsilon^{2/\Theta}}\sqrt{\log\beta_{k+1}} \leq C\sqrt{\log k + k\log m + a\log k! + \log\ell}.$$

Therefore, it remains to check that

$$\eta mn(k+1)^{a-4/\Theta} \geq C\sqrt{\log k + k\log m + a\log k! + \log\ell}.$$

Since $a-4/\Theta > \frac{1}{2}$ by (10) and $\log k! \leq k\log k$, $(k+1)^{a-4/\Theta}$ is larger than the right-hand side above for all large enough k. Then, since ϵ is fixed, setting m large enough makes the above inequality true for all $k \geq 1$.

Hence, we obtain a coupling between the particles that end up in $S_k^{\mathrm{ext}}(i)$ and an independent Poisson point process Ξ with intensity measure $\zeta(y) = (1-\bar{\epsilon})(1-\epsilon_{k+1})\lambda_0\mu_y = (1-\frac{\epsilon_k}{2}-\frac{\epsilon_{k+1}}{2})\lambda_0\mu_y$ that succeeds with probability at least

$$\begin{aligned}
&1 - \sum_{y\in S_k^{\mathrm{ext}}(i)} \exp\{-\bar{C}(1-\epsilon_{k+1})\lambda_0\mu_y\bar{\epsilon}^2\bar{\Delta}^{d/2}\} \\
&\geq 1 - \sum_{y\in S_k^{\mathrm{ext}}(i)} \exp\left\{-\bar{C}C_{\mathrm{mix}}^{d/2}\lambda_0\mu_y\frac{\epsilon^2}{4(k+1)^4}\ell_k^d\right\} \\
&\geq 1 - (\ell_k + 2\eta mnk^a\ell_{k-1})^d \exp\left\{-\bar{C}C_{\mathrm{mix}}^{d/2}\lambda_0 C_M^{-1}\frac{\epsilon^2}{4(k+1)^4}\ell_{k-1}^d\right\} \\
&\geq 1 - \tfrac{1}{2}\exp\{-c\psi_k\}
\end{aligned} \tag{23}$$

where c is constant independent of ℓ, k and ϵ, and we used that $\bar{\Delta} \geq \beta_{k+1} \geq C_{\mathrm{mix}}\ell_k^2 \geq C_{\mathrm{mix}}\ell_{k-1}^2$. The last inequality holds for large k by setting m large, since $\ell_{k-1} = m^{k-2}((k-1)!)^a\ell$. Similarly, for small $k \geq 2$ the inequality holds since $C_M^{-1}\epsilon^2\lambda_0\ell^d \geq \alpha$ is assumed large enough.

Now, for the case $k \geq 2$, define a Poisson point process Ξ' consisting of those particles of Ξ whose displacement throughout $[\tau\beta_k, (\tau+2)\beta_k]$ is in $Q_{\eta mnk^a\ell_{k-1}}$. For each particle of Ξ, this condition is satisfied with probability $\mathbb{P}[F_{2\beta_k}(\eta mnk^a\ell_{k-1})]$, independently over the particles of Ξ. Using Lemma 5.2 and the thinning property of Poisson processes, we have that Ξ' is a Poisson point process with intensity measure

$$\zeta'(y) = (1-\bar{\epsilon})(1-\epsilon_{k+1})\mathbb{P}[F_{2\beta_k}(\eta mnk^a\ell_{k-1})]\lambda_0\mu_y,$$

which is greater than

$$\begin{aligned}
&\left(1-\frac{\epsilon_k}{2}-\frac{\epsilon_{k+1}}{2}\right)\left(1-c_1\exp\left\{-c_2\frac{(\eta m n k^a \ell_{k-1})^2}{2\beta_k}\right\}\right)\lambda_0\mu_y\\
&\geq\left(1-\frac{\epsilon_k}{2}-\frac{\epsilon_{k+1}}{2}\right)\left(1-c_1\exp\left\{-c_2\frac{(\eta m n k^a)^2\epsilon^{4/\Theta}}{2C_{\mathrm{mix}}k^{8/\Theta}}\right\}\right)\lambda_0\mu_y\\
&\geq\left(1-\frac{\epsilon_k}{2}-\frac{\epsilon_{k+1}}{2}\right)\left(1-c_1\exp\left\{-c_2\frac{(\eta n)^2 k}{2(\beta/\ell^2)}\right\}\right)\lambda_0\mu_y,
\end{aligned}$$

where the first inequality follows from the definition of β_k and the second inequality follows from the condition $2a-8/\Theta>1$ in (10) and from $C_{\mathrm{mix}}=\frac{\beta\epsilon^{4/\Theta}m^2}{\ell^2}$, which is obtained by setting $\beta_1=\beta$ in (8). Setting m, and thus n, sufficiently large with respect to β, ϵ, η and the constants c_1 and c_2, we obtain that

$$\begin{aligned}
\zeta'(y)&\geq\left(1-\frac{\epsilon_k}{2}-\frac{\epsilon_{k+1}}{2}\right)\left(1-\frac{(\epsilon_k-\epsilon_{k+1})}{4}\right)\lambda_0\mu_y\\
&\geq\left(1-\frac{3\epsilon_k}{4}-\frac{\epsilon_{k+1}}{4}\right)\lambda_0\mu_y.
\end{aligned}$$

Conditioning on the coupling above, we obtain that $D_k^{\mathrm{ext}}(i,\tau)=1$ with probability at least

$$\begin{aligned}
&1-\sum_{i':S_{k-1}(i')\subseteq S_k^{\mathrm{ext}}(i)}\exp\left\{-\frac{1}{2}\left(\frac{\epsilon_k-\epsilon_{k+1}}{4}\right)^2\left(1-\frac{3\epsilon_k}{4}-\frac{\epsilon_{k+1}}{4}\right)\lambda_0\sum_{y\in S_{k-1}(i')}\mu_y\right\}\\
&\geq 1-\sum_{i':S_{k-1}(i')\subseteq S_k^{\mathrm{ext}}(i)}\exp\left\{-\frac{1}{2}\left(\frac{\epsilon^2}{16(k+1)^4}\right)\left(1-\frac{3\epsilon_1}{4}-\frac{\epsilon_2}{4}\right)\lambda_0\sum_{y\in S_{k-1}(i')}\mu_y\right\}\\
&\geq 1-(mk^a+2\eta m n k^a)^d\exp\left\{-\frac{1}{2}\left(\frac{\epsilon^2}{16(k+1)^4}\right)\left(1-\frac{15\epsilon}{16}\right)\lambda_0 C_M^{-1}\ell_{k-1}^d\right\} \qquad (24)\\
&\geq 1-\tfrac{1}{2}\exp\{-c\psi_k\}
\end{aligned}$$

for some constant c, where in the fist step we applied Chernoff's bound from Lemma A.1 with $\delta=\frac{\epsilon_k-\epsilon_{k+1}}{4}$, using that $(1-\delta)\left(1-\frac{3\epsilon_k}{4}-\frac{\epsilon_{k+1}}{4}\right)\geq 1-\epsilon_k$. In the second step, we used that ϵ_k is decreasing with k. The last inequality holds using the same argument as the one following (23). This and (23) establishes part 2 for $k\geq 2$.

For part 2 with $k=1$ we again use the Poisson point process Ξ of intensity measure

$$\zeta(y)\geq\left(1-\frac{\epsilon_k}{2}-\frac{\epsilon_{k+1}}{2}\right)\lambda_0\mu_y=\left(1-\frac{7\epsilon}{8}\right)\lambda_0\mu_y$$

over $S_1^{\text{ext}}(i)$ as defined above. We also use the fact that $E(i,\tau)$ is an event restricted to the super cell i and $S_1^{\text{ext}}(i)$ contains the super cell i (see Remark 4.1). Recall that, for the event $E(i,\tau)$, we only consider the particles of Ξ whose displacement from time $\tau\beta$ to $(\tau+\eta)\beta$ is inside $Q_{w\ell}$. Let the event that this happens for a given particle of Ξ be denoted by $F_{\eta\beta}(w\ell)$. Then, we apply Lemma 5.2 with $\Delta = \eta\beta$ and $z = w\ell$ to obtain

$$\mathbb{P}[F_{\eta\beta}(w\ell)] \geq 1 - c_1 \exp\left\{-c_2 \frac{(w\ell)^2}{\eta\beta}\right\}.$$

Using the fact that $w^2\ell^2 \geq \frac{1}{c_2}\eta\beta\log(8c_1\epsilon^{-1})$, we have that $\mathbb{P}[F_{\eta\beta}(w\ell)] \geq 1 - \frac{\epsilon}{8}$. Therefore, using thinning, we have that the particles of Ξ for which $F_{\eta\beta}(w\ell)$ hold consist of a Poisson point process with intensity at least $\left(1 - \frac{7\epsilon}{8}\right)\left(1 - \frac{\epsilon}{8}\right)\lambda_0\mu_y \geq (1-\epsilon)\,\lambda_0\mu_y$. Since $E(i,\tau)$ is increasing, we have that

$$\mathbb{P}(E(i,\tau) = 0 \mid F \cap \{D_k^{\text{base}}(i,\tau) = 1\}] \leq 1 - \nu_{E_{\text{st}}}((1-\epsilon)\lambda, Q_{(2\eta+1)\ell}, Q_{w\ell}, \eta\beta) \leq e^{-\alpha}.$$

A similar argument as above can be used to establish part 1 with $k < \kappa$. For $k = \kappa$, the argument is simpler as we do not need to carry out the coupling procedure. □

Later, in Section 6, we will use Lemma 5.3 to bound the probability that a path $P \in \Omega_{\kappa,t}^{\text{sup}}$ of multi-scale bad cells exists. We will use a uniform bound to control the probability that at least one of the space-time cell of scale κ is multi-scale bad. In the converse case, where all scale κ cells are multi-scale good (i.e. not multi-scale bad), we will need to consider paths in $\Omega_{\kappa-1,t}^{\text{sup}}$ and count how many such paths exist. To that end, we now show some bounds that hold for paths in $\Omega_{\kappa-1,t}^{\text{sup}}$.

Lemma 5.4. *Assume the conditions of Lemma 5.3 are satisfied and let $P \in \Omega_{\kappa-1,t}^{sup}$ be the support connected D-path $((k_1,i_i,\tau_1),\dots,(k_z,i_z,\tau_z))$. Then, with ψ_k as defined in (22) there exists a constant c_3, such that we have*

$$\mathbb{P}\left[\bigcap_{j=1}^{z}\{A_{k_j}(i_j,\tau_j) = 0\}\right] \leq \exp\left\{-c_3\sum_{j=1}^{z}\psi_{k_j}\right\}.$$

Proof. We derive the probability that all cells of P are multi-scale bad. Consider the following order of cells of P. First, take an arbitrary order of $\mathbb{Z}^d$. Then, we say that (k_j,i_j,τ_j) precedes $(k_{j'},i_{j'},\tau_{j'})$ in the order if $\tau_j\beta_{k_j} < \tau_{j'}\beta_{k_{j'}}$ or if both $\tau_j\beta_{k_j} = \tau_{j'}\beta_{k_{j'}}$ and i_j precedes $i_{j'}$ in the order of $\mathbb{Z}^d$. Then, for any j, we let J_j be a subset of $\{1,2,\dots,z\}$ containing all j' for which $(k_{j'},i_{j'},\tau_{j'})$ precedes (k_j,i_j,τ_j) in the order.

Using this order, we write

$$\mathbb{P}\left[\bigcap_{j=1}^{z}\{A_{k_j}(i_j,\tau_j)=0\}\right] \leq \prod_{j=1}^{z}\mathbb{P}\left[A_{k_j}(i_j,\tau_j)=0 \mid \bigcap_{j'\in J_j}\{A_{k_{j'}}(i_{k'_j},\tau_{k'_j})=0\}\right].$$

Note that, for each $j'\in J_j$, we have that (k_j,i_j,τ_j) and $(k_{j'},i_{j'},\tau_{j'})$ are well separated. Using the definition of well separated cells, we have that $R^{\inf}_{k_{j'}}(i_{j'},\tau_{j'}) \not\subseteq R^{\sup}_{k_j}(i_j,\tau_j)$ and $R^{\inf}_{k_j}(i_j,\tau_j) \not\subseteq R^{\sup}_{k_{j'}}(i_{j'},\tau_{j'})$. Hence, we obtain by Lemma 4.2 that $R^{\inf}_{k_{j'}}(i_{j'},\tau_{j'}) \cap R^{\inf}_{k_j}(i_j,\tau_j)=\emptyset$. By the ordering above, we also have $\tau_j\beta_{k_j}\geq\tau_{j'}\beta_{k_{j'}}$, which gives that the event $\bigcap_{j'\in J_j}\{A_{k_{j'}}(i_{k'_j},\tau_{k'_j})=0\}$ is measurable with respect to $\mathcal{F}_{k_j}(i_j,\tau_j)$. Then, we apply Lemma 5.3 to obtain a positive constant c_3 such that

$$\mathbb{P}\left[\bigcap_{j=1}^{z}\{A_{k_j}(i_j,\tau_j)=0\}\right] = \exp\left\{-c_3\sum_{j=1}^{z}\psi_{k_j}\right\}.$$

□

Lemma 5.5. *Let z be a positive integer and $k_1,k_2,\dots,k_z\geq 1$ be fixed. Then, if α and m are sufficiently large, the total number of support connected D-paths containing z cells of scales $k_1,k_2,\dots,k_z$ is at most $\exp\left\{\frac{c_3}{2}\sum_{j=1}^{z}\psi_{k_j}\right\}$, where c_3 is the same constant as in Lemma 5.4, ψ is defined in (22) and α is defined in Lemma 5.3.*

Proof. Recall that for two consecutive cells of a support connected D-path, they are either support adjacent or the first cell is support connected with diagonals to the second; see the beginning of this section for details. For the remainder of this proof, when a cell (k,i,τ) of a support connected D-path P is support connected with diagonals to the next cell (k',i',τ') of P, we will refer to the scale 1 cells forming the diagonal connection between a cell contained in $R^{2\sup}_{k}(i,\tau)$ and a cell contained in $R^{2\sup}_{k'}(i',\tau')$ as the *diagonal steps.* Note also that by the definition of D-paths from Section 4.4, the first cell of the diagonal steps is diagonally connected to the last cell of the diagonal steps.

We will prove the result in three steps. We will first show an upper bound for the number of support connected D-paths with no diagonal steps. Next, we will prove a bound for the number of support connected D-paths where the first and last cell of each sequence of diagonal steps is fixed and show that this bound is directly linked to the bound from the first step. Finally, we will prove an upper bound for the number of all possible arrangements of the first and last cell of the diagonal steps for each diagonal, which will then, when combined with the bound from step two, prove the lemma.

We begin with the first step, by considering the number of possible support connected D-paths when each cell of the D-path is support adjacent to the next cell, that is, there are no diagonal steps in the D-path. For any $k, k' \geq 1$, define

$$\Phi_{k,k'} = \max_{(i_1,\tau_1)\in\mathbb{Z}^{d+1}} |\{(i_2,\tau_2) \in \mathbb{Z}^{d+1} : (k, i_1, \tau_1) \text{ is support adjacent to } (k', i_2, \tau_2)\}|,$$

that is, $\Phi_{k,k'}$ is the maximum number of cells of scale k' that are support adjacent to a given cell of scale k. Let χ_k be the number of cells of scale k whose extended support contains $R_1(0, 0)$. This gives that the total number of different D-paths of z cells of scales $k_1, \dots, k_z$ with no diagonal steps can be bound above by

$$\chi_{k_1} \prod_{j=2}^{z} \Phi_{k_{j-1},k_j}.$$

Now we derive a bound for χ_k. At scale k, the number of cells that have the same extended support is $\left(\frac{\ell_{k+1}}{\ell_k}\right)^d \frac{\beta_{k+1}}{\beta_k} = m^{d+2}k^{2a-8/\Theta}(k + 1)^{8/\Theta+ad}$. Furthermore, the extended support of a cell of scale k contains exactly $27(2(3m + 1) + 1)^d$ different cells of scale $k + 1$. Thus, the number of different extended supports for a cell of scale k that contains $R_1(0, 0)$ is bounded above by

$$\chi_k \leq 27(2(3m + 1) + 1)^d m^{d+2}k^{2a-8/\Theta}(k + 1)^{8/\Theta+ad} \leq \exp\left\{\frac{c_3}{16}\psi_k\right\},$$

where the last inequality holds since m and α are large enough. To derive a bound for $\Phi_{k,k'}$, fix a cell (k, i_1, τ_1) of scale k. Now, a cell of scale k' can only be support adjacent to (k, i_1, τ_1) if it is inside the region

$$\bigcup_{x\in R_k^{2\text{sup}}(i_1,\tau_1)} \left(x + [-(3m + 2)\ell_{k'+1}, (3m + 2)\ell_{k'+1}]^d \times [-14\beta_{k'+1}, 14\beta_{k'+1}]\right). \tag{25}$$

For $k \geq k'$, let $\overline{\Phi_{k,k'}}$ be the number of cells of scale k' that lie in the region above. We then have that $\Phi_{k,k'} \leq \overline{\Phi_{k,k'}}$ and

$$\begin{aligned}\overline{\Phi_{k,k'}} &= \left(\frac{(6m+3)\ell_{k+1} + 2(3m+2)\ell_{k'+1}}{\ell_{k'}}\right)^d \left(\frac{27\beta_{k+1} + 28\beta_{k'+1}}{\beta_{k'}}\right) \\ &\leq \left((6m+3)m^{k-k'+1}\prod_{i=k'+1}^{k+1} i^a + 2(3m+2)m(k'+1)^a\right)^d \\ &\times \left(27m^{2(k-k'+1)}\prod_{i=k'}^{k} i^{2a-8/\Theta}(i+1)^{8/\Theta} + 28m^2 k'^{2a-8/\Theta}(k'+1)^{8/\Theta}\right) \\ &\leq c_4 m^{(k-k'+2)d} k^{kad} m^{2(k-k'+1)} k^{2ka} \leq c_4 m^{(d+2)(k-k'+2)} k^{(ad+2a)k},\end{aligned}$$

for some universal positive constant c_4. Note that for any constant $c > 0$, since m and α are large enough, it holds that $c\overline{\Phi_{k,k'}} \leq c\overline{\Phi_{k,1}} \leq \exp\left\{\frac{c_3}{16}\psi_k\right\}$. For $k < k'$ we set $\overline{\Phi_{k,k'}} = 2^{d+1}\overline{\Phi_{k',k}}$, which gives using (25) that $\Phi_{k,k'} \leq 2^{d+1}\overline{\Phi_{k',k}} \leq \exp\left\{\frac{c_3}{16}\psi_{k'}\right\}$.

Observe now that

$$\prod_{j=2}^{z}\overline{\Phi_{k_{j-1},k_j}} \leq \prod_{j=2}^{z}\left(\overline{\Phi_{k_{j-1},k_j}}\mathbb{1}_{k_{j-1}\geq k_j} + \mathbb{1}_{k_{j-1}<k_j}\right)\left(\overline{\Phi_{k_j,k_{j-1}}}\mathbb{1}_{k_j\geq k_{j-1}} + \mathbb{1}_{k_j<k_{j-1}}\right). \quad (26)$$

If write $k_0 = k_{z+1} = \infty$, the right hand side of (26) can be written as

$$\prod_{j=1}^{z}\left(\overline{\Phi_{k_j,k_{j-1}}}\mathbb{1}_{k_j\geq k_{j-1}} + \mathbb{1}_{k_j<k_{j-1}}\right)\left(\overline{\Phi_{k_j,k_{j+1}}}\mathbb{1}_{k_j\geq k_{j+1}} + \mathbb{1}_{k_j<k_{j+1}}\right).$$

Then, applying the bounds above for Φ and χ, we obtain

$$\chi_{k_1}\prod_{j=2}^{z}\Phi_{k_{j-1},k_j} \leq \chi_{k_1}\overline{\Phi_{k_1,1}}\prod_{j=2}^{z}\left(2^{d+1}\overline{\Phi_{k_j,1}}\right)^2 \leq \exp\left\{\frac{c_3}{8}\sum_{j=1}^{z}\psi_{k_j}\right\}.$$

We now proceed to the second step. By definition, a cell (k,i,τ) can only be support connected with diagonals to (k',i',τ') if there exists a cell $(1,i'',\tau'')$ for which $R_1(i'',\tau'') \subseteq R_k^{2\text{sup}}(i,\tau)$ that is diagonally connected to a cell $(1,\hat{i},\hat{\tau})$ for which $R_1(\hat{i},\hat{\tau}) \subseteq R_{k'}^{2\text{sup}}(i',\tau')$. Define $(\hat{i}-i'',\hat{\tau}-\tau'') \in \mathbb{Z}^{d+1}$ to be the *relative position* of the cell $(1,\hat{i},\hat{\tau})$ with respect to the cell $(1,i'',\tau'')$. For convenience we will write when (k,i,τ) is adjacent to (k',i',τ') that the relative position of $(1,\hat{i},\hat{\tau})$ with respect to $(1,i'',\tau'')$ is $(0,0)$. We will show a bound for the number of such relative positions in a D-path in step three, so we now proceed to show a bound for the number of D-paths that have fixed relative positions of $(1,\hat{i},\hat{\tau})$ with respect to $(1,i'',\tau'')$ for all consecutive pairs of cells in the path.

Let (k, i, τ) be a cell of the support connected D-path that is support adjacent or support connected with diagonals to the cell (k', i', τ') and let $(\hat{i} - i'', \hat{\tau} - \tau'') \in \mathbb{Z}^{d+1}$ be the relative position, as above. Then, for a fixed relative position $(\hat{i} - i'', \hat{\tau} - \tau'')$, define

$$\begin{aligned} \Phi^*_{k,k'} = \max_{(i_1,\tau_1)\in\mathbb{Z}^{d+1}} |\{(i_2, \tau_2) \in \mathbb{Z}^{d+1} : (k, i_1, \tau_1) &\text{ is support adjacent or} \\ &\text{support connected with diagonals to } (k', i_2, \tau_2) \\ &\text{with fixed relative position } (\hat{i} - i'', \hat{\tau} - \tau'')\}|. \end{aligned}$$

Then the number of D-paths containing z cells of scales $k_1, k_2, \ldots, k_z$ where consecutive cells are support adjacent or support connected with diagonals with fixed relative positions is smaller than

$$\chi_{k_1} \prod_{j=2}^{z} \Phi^*_{k_{j-1},k_j}. \tag{27}$$

Since a cell of a support connected D-path can either be support adjacent or support connected with diagonals to the next cell, we will consider the two cases individually. Consider first the case when the two cells are support adjacent, i.e. there are no diagonal steps between the extended supports of (k, i, τ) and (k', i', τ'). By step 1 of this proof, we have that in this case

$$\Phi^*_{k,k'} \leq \overline{\Phi_{k,k'}}.$$

Let now the relative position of $(1, \hat{i}, \hat{\tau})$ with respect to $(1, i'', \tau'')$ be different from $(0, 0)$. Then, since the relative position is fixed, $\Phi^*_{k,k'}$ can be bound by the product of the number of cells of scale 1 contained in the extended support of a cell of scale k and the number of cells of scale 1 that are contained in the extended support of a cell of scale k'. Using the bounds from step 1, this gives that

$$\Phi^*_{k,k'} \leq \overline{\Phi_{k,1}} \cdot \overline{\Phi_{k',1}}.$$

We have therefore for any fixed relative position of $(1, \hat{i}, \hat{\tau})$ with respect to $(1, i'', \tau'')$ that

$$\Phi^*_{k,k'} \leq \overline{\Phi_{k,k'}} \mathbb{1}_{(\hat{i},\hat{\tau})=(i'',\tau'')} + \overline{\Phi_{k,1}} \cdot \overline{\Phi_{k',1}} \mathbb{1}_{(\hat{i},\hat{\tau})\neq(i'',\tau'')}. \tag{28}$$

By using the bounds from step 1 and (28), we get that

$$\chi_{k_1} \prod_{j=2}^{z} \Phi^*_{k_{j-1},k_j} \leq \exp\left\{ \frac{3c_3}{8} \sum_{j=1}^{z} \psi_{k_j} \right\}. \tag{29}$$

We now move on to the third step and show a bound for the number of different relative positions that are possible in a support connected D-path of cells of scales $k_1, k_2, \ldots, k_z$. We will show that this number is smaller than

$$\exp\left\{ \frac{c_3}{8} \sum_{j=1}^{z} \psi_{k_j} \right\}, \tag{30}$$

which combined with (29) proves the lemma.

Consider two consecutive cells of the D-path and let $(1, i, \tau)$ be a cell contained in the extended support of the first cell that is diagonally connected to a cell $(1, i', \tau')$ that is contained in the extended support of the second cell. Recall from Section 2 the definition of the base-height index and from Section 4.3 the properties of the sequence of cells that make $(1, i, \tau)$ diagonally connected to $(1, i', \tau')$. Denote by x the height difference between the two cells, i.e. $x := |h - h'|$ in the base-height index, and define $A(x), x \in \mathbb{Z}$, to be the number of different cells of scale 1 that $(1, i, \tau)$ is diagonally connected to with height difference x. More precisely,

$$A(x) = \max_{(b_1,h_1)\in\mathbb{Z}^{d+1}} |\{(b_2, h_2) \in \mathbb{Z}^{d+1} : |h_2 - h_1| = x \text{ and}$$
$$(b_1, h_1) \text{ is diagonally connected to } (b_2, h_2)\}|.$$

Let H_k be the side length of the cube $S_k^{2\mathrm{sup}}(i)$ divided by the side length of the cube $S_1(i')$, that is, let $H_k = (3m + 1)m^k((k + 1)!)^a$. Recall from Section 4.3 that using the base-height index, for any two cells $(b_i, h_i), (b_j, h_j)$ of the diagonal, $h_i h_j \geq 0$ and that $h_{j-1} - h_j \in \mathrm{Sign}(h_{j-1})$ for any two consecutive cells of the diagonal. Therefore, given the z cells of scales $k_1, k_2, \ldots, k_z$, the maximum number of scale 1 diagonal steps contained in all diagonal connections between the cells of the path is at most

$$H := \sum_{i=1}^{z-1} H_{k_i}.$$

Letting x_i, for $i \in \{1, 2, \ldots, z-1\}$ be the height difference between the i-th and $(i+1)$-th cell of the path, with $x_i = 0$ if the cells are support adjacent, we have that the number of possible configurations of the diagonal steps is at most

$$\sum_{y=0}^{H} \sum_{\substack{x_1, x_2, \ldots, x_{z-1}: \\ x_1 + \cdots + x_{z-1} = y}} A(x_1+1)A(x_2+1)\cdots A(x_{z-1}+1). \tag{31}$$

See Figure 4 for an illustration of one such configuration. The +1 terms in (31) account for the fact that each diagonal either ends in a multi-scale bad cell or is adjacent to one, so by increasing the height difference by 1, we account for both possibilities at once.

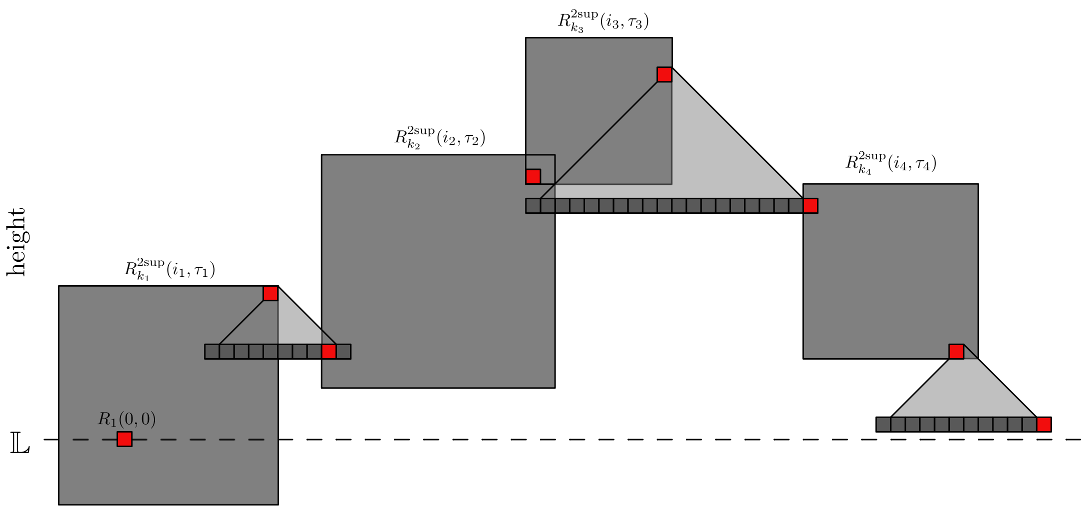

Figure 4: Example with $z = 4$ where $x_2 = 0$ and $x_1 + x_2 + x_3 + x_4 \leq H$. The red cells are the scale 1 cells used as the (fixed) beginnings and ends of the diagonals within their respective extended supports. The dark cells with the matching red cell at the bottom of the triangles represent the area containing $A(x_i)$ cells of scale 1.

Due to the properties of the diagonal steps, we have that $A(x)$ is the volume of a d-dimensional ball of radius x, so $A(x) \leq c_d x^d$ where c_d is some constant that depends on d only. It follows, e.g. by the method of Lagrange multipliers, that for $x_k \geq 0, k \in \{1, \ldots, z-1\}$ and $x_1 + \cdots + x_{z-1} = y$, we have

$$A(x_1+1)A(x_2+1)\cdots A(x_{z-1}+1) \leq \left[A\left(\frac{y}{z-1}+1\right)\right]^{z-1}.$$

Next, using the above bound and

$$\sum_{\substack{x_1, x_2, \ldots, x_{z-1}: \\ x_1 + \cdots + x_{z-1} = y}} 1 = \binom{z+y-2}{z-2},$$

we have that the sum in (31) is smaller than

$$\sum_{y=0}^{H}\binom{z+y-1}{z-1}A\left(\tfrac{y}{z-1}+1\right)^{z-1}\leq\binom{z+H}{z}A\left(\tfrac{H}{z-1}+1\right)^{z-1},$$

where the binomial inequality used can easily be proven by induction on H (using Pascal's rule).

Then, for some positive constants C and C_2 and using that $\frac{H}{z}$ is large for large α, we have that

$$\begin{aligned}\binom{z+H}{z}A\left(\tfrac{H}{z-1}+1\right)^{z-1}&\leq C\frac{(z+H)^z}{(z)!}\left(\frac{H}{z-1}+1\right)^{(z-1)d}\\&\leq C\frac{(z+H)^z}{(z/3)^z}\left(\frac{2H}{z}\right)^{zd}\\&\leq C(3+3H/z)^z\left(\frac{2H}{z}\right)^{zd}\\&\leq C\left(C_2\frac{H}{z}\right)^{2zd}.\end{aligned}$$

In order to complete the proof, it remains to show that $C\left(C_2\frac{H}{z}\right)^{2zd}\leq\exp\left\{\frac{c_3}{8}\sum_{j=1}^{z}\psi_{k_j}\right\}$, which is equivalent to showing that

$$\tilde{C}\log\left(\frac{H}{z}\right)\leq\frac{1}{z}\sum_{j=1}^{z}\psi_{k_j},\tag{32}$$

where $\tilde{C}$ is some constant. For small k, setting α (and thus ℓ) large enough gives that $H_k\leq\psi_k$, and similarly setting m large gives $H_k\leq\psi_k$ for all large k. Combined, this gives (32). □

Remark 5.2. As mentioned in Section 2 (see Remark 2.1), one can set time to be height in the base-height index. In that case all results up to and including Lemma 5.4 go through unchanged. However, an important issue arises in Lemma 5.5. In the proof of Lemma 5.5, the height H_k of the extended support of a cell becomes the length of the interval $T_k^{2\sup}(\tau)$. Then, if $d\geq 3$, the proof goes through unchanged since it still holds that $H_k\leq\psi_k$ for all $k\geq 1$ by setting m and α large enough. For $d=2$ however the lemma no longer holds, since it can happen that the number of different arrangements of diagonal steps and the z cells of a path is larger than $\exp\{\sum_{i=1}^{z}\psi_{k_i}\}$. To see this, consider the following example. Let k_1 be large and let $k_i=2$ for all $i\in\{2,\ldots,z\}$. Let z be the largest integer for which it holds that $\psi_{k_1}\geq 4\sum_{i=2}^{z}\psi_{k_i}=4(z-1)\psi_2$. Note that this gives that

$$\frac{\psi_{k_1}}{4\psi_2} \leq z \leq 1 + \frac{\psi_{k_1}}{4\psi_2} \leq \frac{\psi_{k_1}}{3\psi_2}, \tag{33}$$

where the last inequality holds for any large enough k_1. Furthermore note that since $d = 2$ we can write $H_{k_1} = a_{k_1}\psi_{k_1}$, where a_{k_1} is a term that can be made arbitrarily large by increasing k_1. Next, observe that the number of different arrangements of diagonal steps for the cells of scale 2 is at least $\binom{H_{k_1}+z-2}{z-2}$. Therefore, we want to show that for any constant $c_1 > 0$, we can set k_1 large enough to have

$$\binom{H_{k_1} + z - 2}{z - 2} \geq \exp\left\{\sum_{i=1}^{z} c_1\psi_{k_i}\right\}. \tag{34}$$

Consider first the left hand side of (34) and note that it is bigger than

$$\frac{H_{k_1}^{z-2}}{(z-2)!} \geq \left(\frac{H_{k_1}}{z-2}\right)^{z-2} \geq (3a_{k_1}\psi_2)^{z-2},$$

where in the last inequality we used the upper bound on z from (33). For the right hand side of (34), we have

$$\exp\left\{\sum_{i=1}^{z} c_1\psi_{k_i}\right\} = \exp\left\{c_1(\psi_{k_1} + (z-1)\psi_2)\right\} \leq \exp\left\{4c_1\psi_2 z + c_1\psi_2(z-1)\right\},$$

where the inequality follows from the upper bound on ψ_{k_1} obtained from the leftmost inequality in (33). Since a_{k_1} grows with k_1, we obtain (34) for large enough k_1.

For any support connected D-path $P = ((k_1, i_1, \tau_1), (k_2, i_2, \tau_2), \ldots, (k_z, i_z, \tau_z)) \in \Omega^{\text{sup}}_{\kappa-1,t}$, we define the weight of P as $\sum_{j=1}^{z} \psi_{k_j}$. The lemma below shows that, for any $P \in \Omega^{\text{sup}}_{\kappa-1,t}$, if t is large enough, then the weight of P must be large.

Lemma 5.6. *Let $t > 0$ and let $P = ((k_1, i_1, \tau_1), (k_2, i_2, \tau_2), \ldots, (k_z, i_z, \tau_z))$ be a path in $\Omega^{sup}_{\kappa-1,t}$. If α is sufficiently large and $\kappa = \mathcal{O}(\log t)$, then there exist a positive constant $c = c(C_M)$ and a value C independent of t such that*

$$\sum_{j=1}^{z} \psi_{k_j} \geq \begin{cases} C\frac{\sqrt{t}}{(\log t)^c}, & \text{for } d = 1, \\ C\frac{t}{(\log t)^c}, & \text{for } d = 2, \\ Ct, & \text{for } d \geq 3. \end{cases} \tag{35}$$

Proof. Let $\Delta_k^{2\text{sup}}$ denote the diameter of the extended support of a cell of scale k. Then, we have

$$\Delta_k^{2\sup} \le (6m+3)\ell_{k+1}\sqrt{d} + 27\beta_{k+1} = (6m+3)m(k+1)^a\ell_k\sqrt{d} + 27C_{\mathrm{mix}}\frac{\ell_k^2(k+1)^{8/\Theta}}{\epsilon^{4/\Theta}}.$$

Then, the definition of C_{mix} gives us that there exists a constant c_6 (that might depend on the ratio β/ℓ^2) such that

$$\Delta_k^{2\sup} \le (6m+3)m(k+1)^a\ell_k\sqrt{d} + 27m^2\frac{\beta}{\ell^2}\ell_k^2(k+1)^{8/\Theta} \le c_6m^2(k+1)^{2a}\ell_k^2.$$

Then, for $k \ge 2$, we have for $d = 1$ that

$$\begin{aligned}\psi_k &= \frac{\epsilon^2\lambda_0\ell_{k-1}}{(k+1)^4} = \frac{\epsilon^2\lambda_0}{(k+1)^4}\left(\frac{\ell_k}{m(k^a)}\right)\\ &\ge \frac{\epsilon^2\lambda_0}{m(k+1)^{a+4}}\left(\frac{\sqrt{c_6m^2(k+1)^{2a}\ell_k^2}}{\sqrt{c_6m^2(k+1)^{2a}}}\right)\\ &\ge \frac{\epsilon^2\lambda_0}{\sqrt{c_6}m^2(k+1)^{3a+4}}\sqrt{\Delta_k^{2\sup}}\end{aligned}$$

and for $d \ge 2$ that

$$\begin{aligned}\psi_k &= \frac{\epsilon^2\lambda_0\ell_{k-1}^{d-2}}{(k+1)^4}\left(\frac{\ell_k}{mk^a}\right)^2\\ &\ge \frac{\epsilon^2\lambda_0\ell_{k-1}^{d-2}}{m^2(k+1)^{2a+4}}\left(\frac{c_6m^2(k+1)^{2a}\ell_k^2}{c_6m^2(k+1)^{2a}}\right)\\ &\ge \frac{\epsilon^2\lambda_0\ell_{k-1}^{d-2}}{c_6m^4(k+1)^{4a+4}}\Delta_k^{2\sup}.\end{aligned}$$

Now, since $\kappa = \mathcal{O}(\log t)$, there exists a constant c_7 such that $(k+1)^b \le c_7(\log t)^b$ for all $k \le \kappa$ and any $b \ge 1$. We use this for dimensions 1 and 2. For dimension 3 and higher, we set c_7 large enough to satisfy $\frac{\ell_{k-1}^{d-2}}{(k+1)^{4a+4}} \ge \frac{\ell^{d-2}}{mc_7}$; this is possible since ℓ_k is of order $(k!)^a$. This gives

$$\psi_k \ge \begin{cases}\frac{\epsilon^2\lambda_0}{\sqrt{c_6}c_7m^2}\frac{\sqrt{\Delta_k^{2\sup}}}{(\log t)^{3a+4}}, & \text{for } d=1\\ \frac{\epsilon^2\lambda_0}{c_6c_7m^4}\frac{\Delta_k^{2\sup}}{(\log t)^{4a+4}}, & \text{for } d=2\\ \frac{\epsilon^2\lambda_0\ell^{d-2}}{c_6c_7m^5}\Delta_k^{2\sup}, & \text{for } d\ge 3.\end{cases}$$

For $k = 1$ we write $\psi_1 \ge c\sqrt{\Delta_k^{2\sup}}$ for $d = 1$ and $\psi_1 \ge c\Delta_k^{2\sup}$ for $d \ge 2$, where c is some positive value that may depend on ϵ, m, λ_0, ℓ and ν_E. Moreover, if a support connected D-path is such that $\sum_{j=1}^{z} \Delta_{k_j}^{2\text{sup}} < t/2$, the extended support of all cells of the path must be contained in $[-t,t]^{d+1}$. This is true because if there are no diagonal steps in P, then the extended supports are contained in $[-t/2,t/2]^{d+1}$, and if there are diagonal steps, they can only prolong the path by at most $\sum_{j=1}^{z} \Delta_{k_j}^{2\text{sup}}$. Therefore, for $P \in \Omega_{\kappa-1,t}^{\text{sup}}$ we have $\sum_{j=1}^{z} \Delta_{k_j}^{2\text{sup}} \geq t/2$. This implies that there exists a positive C independent of t, but depending on everything else such that

$$\sum_{j=1}^{z} \psi_{k_j} \geq \begin{cases} C\frac{\sqrt{t}}{(\log t)^{3a+4}}, & \text{for } d=1 \\ C\frac{t}{(\log t)^{4a+4}}, & \text{for } d=2 \\ Ct, & \text{for } d \geq 3. \end{cases}$$

□

We now write ψ_k, $k \geq 2$ as a multiple of ψ_2. This will be used to count the number of paths in $\Omega_{\kappa,t}^{\text{sup}}$ later. For this, set $\tilde{\psi}_2 = \psi_2 = 3^{-4}\epsilon^2\lambda_0\ell^d$, and for $j \geq 3$, define

$$\tilde{\psi}_j = 2\tilde{\psi}_2 m^{(j-2)d}((j-1)!)^{ad-3}((j-2)!)^2(j-3)!.$$

Lemma 5.7. *For all $j \geq 2$, it holds that $\tilde{\psi}_j \leq \psi_j \leq 41\tilde{\psi}_j$.*

Proof. For $j \geq 3$ we write

$$\begin{aligned}\psi_j &= \frac{\epsilon^2\lambda_0\ell^d m^{(j-2)d}((j-1)!)^{ad}}{(j+1)^4} = 3^4\tilde{\psi}_2\frac{m^{(j-2)d}((j-1)!)^{ad}}{(j+1)^4} \\ &= 3^4\tilde{\psi}_2 m^{(j-2)d}((j-1)!)^{ad-3}((j-2)!)^2(j-3)!\left(\frac{(j-1)^3(j-2)}{(j+1)^4}\right).\end{aligned}$$

This implies that $\psi_j \leq \frac{3^4}{2}\tilde{\psi}_j \leq 41\tilde{\psi}_j$. The other direction follows from the fact that $\frac{(j-1)^3(j-2)}{(j+1)^4} \geq 1/32$ for all $j \geq 3$. □

6 Size of bad clusters

For $k \geq 1$, define $\mathcal{S}_k^t$ to be the set of indices $i \in \mathbb{Z}^d$ given by

$$\mathcal{S}_k^t = \left\{ i \in \mathbb{Z}^d : \ S_k(i) \text{ intersects } [-t,t]^d \right\}.$$

Similarly, we define $\mathcal{T}_k^t$ as the set of indices τ for time intervals that have a descendent at scale 1 intersecting $[0,t]$. Formally, let

$$\mathcal{T}_k^t = \left\{ \tau \in \mathbb{Z} : \exists \tau' \text{ s.t. } \gamma_1^{(k-1)}(\tau') = \tau \text{ and } T_1(\tau') \cap [0,t] \neq \emptyset \right\}.$$

Note that an interval in $\mathcal{T}_k^t$ with $k \geq 2$ may not intersect $[0,t]$. Using these definitions define

$$\mathcal{R}_k^t = \mathcal{S}_k^t \times \mathcal{T}_k^t.$$

For the following proposition, recall from Section 4.4 the definitions of $K(0,0)$ and $K'(0,0)$.

Proposition 6.1. *For each $(i,\tau) \in \mathbb{Z}^{d+1}$, let $E_{st}(i,\tau)$ be an increasing event that is restricted to the super cube i and the super interval τ, and let $\nu_{E_{st}}$ be the probability associated to E_{st} as defined in Definition 2.3. Fix a constant $\epsilon \in (0,1)$, and integer $\eta \geq 1$ and the ratio $\beta/\ell^2 > 0$. Fix also w such that*

$$w \geq \sqrt{\frac{\eta\beta}{c_2\ell^2} \log\left(\frac{8c_1}{\epsilon}\right)},$$

for some constants c_1 and c_2 which depend on the graph. Then, there exist constants c and C, and positive numbers α_0 and t_0 that depend on ϵ, η and the ratio β/ℓ^2 such that if

$$\alpha = \min\left\{ C_M^{-1}\epsilon^2\lambda_0\ell^d, \log\left(\frac{1}{1-\nu_{E_{st}}((1-\epsilon)\lambda, Q_{(2\eta+1)\ell}, Q_{w\ell}, \eta\beta)}\right)\right\} \geq \alpha_0,$$

we have for all $t \geq t_0$ that

$$\mathbb{P}\left[K(0,0) \not\subseteq \mathcal{R}_1^t\right] \leq \begin{cases} \exp\left\{-C\lambda_0 \frac{t}{(\log t)^c}\right\} & \text{for } d = 2 \\ \exp\left\{-C\lambda_0 t\right\} & \text{for } d \geq 3. \end{cases}$$

Proof. First, for any k, note that the number of cells in $\mathcal{R}_k^t$ satisfies

$$|\mathcal{R}_k^t| \leq \left(2\left\lceil \frac{t}{\ell_k}\right\rceil\right)^d \left\lceil 1 + \frac{t}{\beta_k}\right\rceil. \tag{36}$$

Also, using Lemmas 4.1 and 5.1, we have

$$\begin{aligned}\mathbb{P}[K(i,\tau) \not\subseteq \mathcal{R}_1^t] \leq \mathbb{P}[K'(i,\tau) \not\subseteq \mathcal{R}_1^t] &= \mathbb{P}[\exists P \in \Omega_t \text{ s.t. all cells of } P \text{ have bad ancestry}] \\ &\leq \mathbb{P}[\exists P \in \Omega_{\kappa,t}^{\text{sup}} \text{ s.t. all cells of } P \text{ are multi-scale bad}].\end{aligned}$$

We note that the random variable A_κ is defined differently than other scales. It follows from Lemma 5.3, (36) and the union bound over all cells in $\mathcal{R}^t_\kappa$ that

$$\mathbb{P}[(A_\kappa(i',\tau') = 1 \text{ for all } (i,\tau) \in \mathcal{R}^t_\kappa] \geq 1 - |\mathcal{R}^t_\kappa| \exp\{-c\psi_\kappa\} \geq 1 - \exp\{-c_1 t\}, \quad (37)$$

for some positive constant c_1, where the last step follows by setting κ to be the smallest integer such that $\psi_\kappa \geq t$, which using the Lambert W function and its asymptotics gives that $\kappa = \Theta\left(\frac{\log t}{\log\log t}\right)$. Let us define H as the event that $A_\kappa(i,\tau) = 1$ for all $(i,\tau) \in \mathcal{R}^t_\kappa$. Then, we have

$$\begin{aligned}
&\mathbb{P}\left[\exists P \in \Omega^{\sup}_{\kappa,t} \text{ s.t. all cells of } P \text{ are multi-scale bad}\right] \\
&\quad \leq \mathbb{P}\left[H \cap \{\exists P \in \Omega^{\sup}_{\kappa,t} \text{ s.t. all cells of } P \text{ are multi-scale bad}\}\right] + \mathbb{P}\left[H^c\right] \\
&\quad \leq \mathbb{P}\left[\exists P \in \Omega^{\sup}_{\kappa-1,t} \text{ s.t. all cells of } P \text{ are multi-scale bad}\right] + e^{-c_1 t}.
\end{aligned}$$

To get a bound for the term above, we fix a support connected D-path

$$P = ((k_1, i_1, \tau_1), \ldots, (k_z, i_z, \tau_z)),$$

and use Lemma 5.4 to get

$$\mathbb{P}\left[\bigcap_{j=1}^z \{A_{k_j}(i_j,\tau_j) = 0\}\right] \leq \exp\left\{-c_3 \sum_{j=1}^z \psi_{k_j}\right\}.$$

We now take the union bound over all support connected D-paths with cells of scales $k_1, k_2, \ldots, k_z$ and using Lemma 5.5, we get that

$$\begin{aligned}
&\mathbb{P}\left[\exists P \in \Omega^{\sup}_{\kappa-1,t} \text{ s.t. } P \text{ has } z \text{ multi-scale bad cells of scales } k_1, k_2, \ldots, k_z\right] \\
&\qquad\qquad \leq \exp\left\{-\frac{c_3}{2} \sum_{j=1}^z \psi_{k_j}\right\}.
\end{aligned}$$

This bound depends on z and $k_1, \ldots, k_z$ only through $\sum_{j=1}^z \psi_{k_j}$, which we call the weight of the path. Let W be the set of weights for which there exists at least one path in $\Omega^{\sup}_{\kappa-1,t}$ with such a weight. Then

$$\mathbb{P}\left[\exists P \in \Omega^{\sup}_{\kappa-1} \text{ s.t. all cells of } P \text{ are multi-scale bad}\right] \leq \sum_{w \in W} \exp\left\{-\frac{c_3}{2} w\right\} M(w), \quad (38)$$

where $M(w)$ is the number of possible ways to choose z and $k_1, k_2, \ldots, k_z$ such that

$\sum_{j=1}^{z} \psi_{k_j} = w$.

Let $w = \sum_{j=1}^{z} \psi_{k_j}$ and let $w_1 = \psi_1 |\{j : k_j = 1\}|$. Let $w_2 = w - w_1$, so w_1 is the weight given by cells of scale 1 and w_2 the weight given by the other cells of the path. Note that by Lemma 5.7, $w_2 = \sum_{j:k_j \geq 2} \psi_{k_j} \geq \sum_{j:k_j \geq 2} \tilde{\psi}_{k_j} = h_2\psi_2$ for some non-negative integer h_2. Likewise, $w_2 \leq 41 h_2 \psi_2$ and $w_1 = h_1\psi_1$ for some non-negative integer h_1. Let w_0 be the lower bound on the weight of the path given by Lemma 5.6, so for all $w \in W$, we have $w \geq w_0$. Since either w_1 or w_2 has to be larger than $w_0/2$, we have that either $h_1 \geq \frac{w_o}{2\psi_1}$ or $h_2 \geq \frac{w_0}{2\cdot 41\psi_2}$. Let $M(h_1, h_2)$ be the number of ways to choose z and $k_1, \ldots, k_z$ such that there are h_1 values j with $k_j = 1$ and $\sum_{j:k_j \geq 2} \tilde{\psi}_{k_j} = h_2\psi_2$. For any such choice, we have $w = \sum_{j=1}^{z} \psi_{k_j} \geq h_1\psi_1 + h_2\psi_2$. Then, the sum in the right-hand side of (38) can be bounded above by

$$\sum_{h_1 \geq \frac{w_0}{2\psi_1}} \sum_{h_2=0}^{\infty} \exp\left\{-\frac{c_3}{2}(h_1\psi_1 + h_2\psi_2)\right\} M(h_1, h_2)$$
$$+ \sum_{h_1=0}^{\infty} \sum_{h_2 \geq \frac{w_0}{82\psi_2}} \exp\left\{-\frac{c_3}{2}(h_1\psi_1 + h_2\psi_2)\right\} M(h_1, h_2).$$

We now proceed to bound $M(h_1, h_2)$. Suppose we have h_1 blocks of size ψ_1 and h_2 blocks of size ψ_2. Consider an ordering of the blocks, such that permuting the blocks of the same size does not change the order. Then, for each block of size ψ_2, we color it either black or white, while blocks of size ψ_1 are not colored. For each choice of z and $k_1, \ldots, k_z$, we associate an order and coloring of the blocks as follows. if $k_1 = 1$, then the first block is of size ψ_1. Otherwise, the first $\tilde{\psi}_{k_1}/\psi_2$ blocks are of size ψ_2 and have black color. Then, if $k_2 = 1$, the next block is of size ψ_1, otherwise the next $\tilde{\psi}_{k_2}/\psi_2$ blocks are of size ψ_2 and have white color. We proceed in this way until k_z, where whenever $k_i \neq 1$ we use the color black if i is odd and the color white if i is even. Though there are orders and colorings that are not associated to any choice of z and $k_1, \ldots, k_z$, each such choice of z and $k_1, \ldots, k_z$ corresponds to a unique order and coloring of the blocks. Therefore, the number of ways to order and color the blocks gives an upper bound for $M(h_1, h_2)$. Note that there are $\binom{h_1+h_2}{h_1}$ ways to order the blocks and 2^{h_2} ways to color the size-ψ_2 blocks. Therefore

$$\begin{aligned}
&\mathbb{P}\left[\exists P \in \Omega^{\mathrm{sup}}_{\kappa-1,t} \text{ s.t. all cells of } P \text{ are multi-scale bad}\right] \\
&\quad\le \sum_{h_1 \ge \frac{w_0}{2\psi_1}} \sum_{h_2=0}^{\infty} \exp\left\{-\frac{c_3}{2}(h_1\psi_1 + h_2\psi_2)\right\} \binom{h_1+h_2}{h_1} 2^{h_2} \\
&\qquad+ \sum_{h_1=0}^{\infty} \sum_{h_2 \ge \frac{w_0}{82\psi_2}} \exp\left\{-\frac{c_3}{2}(h_1\psi_1 + h_2\psi_2)\right\} \binom{h_1+h_2}{h_1} 2^{h_2} \\
&\quad\le C \sum_{h1 \ge \frac{w_0}{2\psi_1}} \sum_{h_2=0}^{\infty} \exp\left\{-\frac{c_3}{3}(h_1\psi_1 + h_2\psi_2)\right\} + C \sum_{h_1=0}^{\infty} \sum_{h_2 \ge \frac{w_0}{82\psi_2}} \exp\left\{-\frac{c_3}{3}(h_1\psi_1 + h_2\psi_2)\right\} \\
&\quad\le \exp\{-cw_0\},
\end{aligned}$$

for some constants C and c, where in the second inequality we use Lemma A.2 and the fact that α is sufficiently large to write $\frac{c_3\psi_1}{2} - 1 \ge \frac{c_3\psi_1}{3}$, and similarly for ψ_2. Since we defined w_0 to be the lower bound on the weight of a path given by Lemma 5.6, the proof is complete. □

7 Proof of Theorem 2.1

Proof of Theorem 2.1. By Theorem 3.1, it suffices to show that

$$\sum_{r \ge 1} r^d \mathbb{P}\left[\mathrm{rad}_0(H_0) > r\right] < \infty.$$

We begin by noting that after tessellating space and time, $\mathcal{R}_1^t$ contains cells indexed only by (i, τ) for which $\|i\|_\infty \le \frac{t}{\ell}$ and $|\tau| \le \frac{t}{c\ell^2}$ for some positive constant c. For fixed $R > 0$, if we set $T > 0$ such that

$$\left(\frac{d}{\ell} + \frac{1}{c\ell^2}\right) T \le R,$$

then $\mathcal{R}_1^T$ is contained in $\{u \in \mathbb{Z}^{d+1} : \ \|u\|_1 < R\}$. Let $T(r) = \left(\frac{d}{\ell} + \frac{1}{c\ell^2}\right)^{-1} r$ and fix r_0 such that $T(r_0) > t_0$, where t_0 comes from Proposition 6.1. Then we have that

$$\begin{aligned}
\sum_{r \ge r_0} r^d \mathbb{P}[\mathrm{rad}_0(H_0) > r] &\le \sum_{r \ge r_0} r^d \mathbb{P}\left[H_0 \not\subseteq \mathcal{R}_1^{T(r)}\right] \\
&\le \sum_{r \ge r_0} r^d \mathbb{P}\left[K(0,0) \not\subseteq \mathcal{R}_1^{T(r)}\right],
\end{aligned}$$

where we used in the second inequality that every d-path on the space-time tessellation is also a D-path of bad cells. We now apply Proposition 6.1 with $d \geq 3$ to bound $\mathbb{P}\left[K(0,0) \not\subseteq \mathcal{R}_1^{T(r)}\right]$ for $T(r) > t_0$ and get that

$$\sum_{r \geq r_0} r^d \mathbb{P}[\mathrm{rad}_0(H_0) > r] \leq \sum_{r > r_0} r^d \exp\{-C\lambda_0 T(r)\} = \sum_{r \geq r_0} r^d \exp\left\{-C\left(\frac{d}{\ell} + \frac{1}{c\ell^2}\right)^{-1} \lambda_0 r\right\},$$

for some positive constant C, that does not depend on r. Since this expression is finite, we have by Theorem 3.1 that the Lipschitz surface exists and is a.s. finite.

For $d = 2$ we similarly get that

$$\sum_{r \geq r_0} r^d \mathbb{P}[\mathrm{rad}_0(H_0) > r] \leq \sum_{r \geq r_0} r^d \exp\left\{-C\lambda_0 \frac{\ell r}{(\log \ell r)^{-c}}\right\} < \infty.$$

□

The corollary below gives the probability that a base-height cell $(b, 0) \in \mathbb{L}$ is not part of F, i.e. $F_+(b) \neq 0$ and $F_-(b) \neq 0$, where F_+ and F_- are the two Lipschitz functions as defined in Definition 3.2.

Corollary 7.1. *Assume the setting of Theorem 2.1. There are positive constants C, c, C_3 and r_0 such that for any given $b \in \mathbb{Z}^d$, we have*

$$\mathbb{P}[F_+(b) \cdot F_-(b) \neq 0] < \begin{cases} Cr_0^d \mathbb{P}[E_{st}(0,0)^c] + \sum_{r \geq r_0} r^d \exp\{-C_3\lambda_0 \frac{\ell r}{(\log \ell r)^c}\}, & \text{for } d = 2 \\ Cr_0^d \mathbb{P}[E_{st}(0,0)^c] + \sum_{r \geq r_0} r^d \exp\{-C_3\lambda_0 \ell r\}, & \text{for } d \geq 3. \end{cases}$$

Proof. Recall first that by construction, $F_+(b) = 0$ if and only if $F_-(b) = 0$. Then, we have for a positive constant C that depends only on d that

$$\begin{aligned}
\mathbb{P}[F_+(b) \neq 0] &\leq \sum_{(x,0) \in L} \mathbb{P}[(x,0) \rightarrowtail_d (b,0)] \\
&= \sum_{\substack{(x,0) \in L \\ \|x-b\|_1 \leq r_0}} \mathbb{P}[(x,0) \rightarrowtail_d (b,0)] + \sum_{\substack{(x,0) \in L \\ \|x-b\|_1 > r_0}} \mathbb{P}[(x,0) \rightarrowtail_d (b,0)] \\
&\leq \sum_{\substack{(x,0) \in L \\ \|x-b\|_1 \leq r_0}} \mathbb{P}[E_{\mathrm{st}}^c(x,0)] + \sum_{\substack{(x,0) \in L \\ \|x\|_1 > r_0}} \mathbb{P}[(0,0) \rightarrowtail_d (x,0)] \\
&\leq Cr_0^d \mathbb{P}\left[E_{\mathrm{st}}^c(0,0)\right] + \sum_{r > r_0} Cr^d \mathbb{P}[\mathrm{rad}_0(H_0) > r].
\end{aligned}$$

The sum above can be bounded as in the proof of Theorem 2.1. □

Remark 7.1. We note that the sum in Corollary 7.1 is decreasing with ℓ and can in fact be made arbitrarily small by making ℓ large enough. This gives us that if the probability of the event $E_{\mathrm{st}}(i,\tau)$ is increasing in ℓ, the expression in Corollary 7.1 can also be made arbitrarily small.

8 Proof of Theorem 2.2

Recall from Section 3 that a hill H_u is defined as all sites in $\mathbb{Z}^{d+1}$ that can be reached by a d-path started from $u \in \mathbb{L}$. Recall also the definition of a mountain M_u as a union of all hills that contain u. By the construction of D-paths, every d-path on the space-time tessellation is also a D-path of bad cells. For this reason, as in Section 7, we will use an extension D-paths when bounding probabilities of the existence of various hills and mountains in this section.

We begin by considering a broader range of diagonally connected paths. Intuitively, these are paths that can move within sequences of hills H_u for different $u \in \mathbb{L}$. Let $u = (b, 0) \in \mathbb{L}$ be a cell of the zero-height plane. By Definition 3.2, we know a mountain touches the Lipschitz surface at $(b, F_+(u))$ and $(b, F_-(u))$, but we cannot say anything more than that. If we want to say something about the positive and negative depth of the surface F across a larger area, we therefore need to consider a large number of different mountains. Since these mountains likely intersect and are composed of some of the same hills, we need a better way to control their dependences. To that end, we will consider paths with diagonals that can be thought of as concatenations of different D-paths, where some D-paths may be taken in reverse order. In order to define these, which we will refer to as DD-paths, we will need to define the concept of a *double diagonal*, as well as slightly change the definition of two cells being diagonally connected.

As before, we say that distinct scale 1 cells (i,τ) and (i',τ') are *adjacent* if $\|i-i'\|_\infty \leq 1$ and $|\tau-\tau'| \leq 1$. Also, we say that (i,τ) is *diagonally connected* to (i',τ') if there exists a sequence of cells $(i,\tau) = (b_0,h_0),(b_1,h_1),\ldots,(b_n,h_n) = (\hat{\imath},\hat{\tau})$, where the indices (b_j,h_j) refer to the base-height index, such that all the following hold:

- for all $j \in \{1,\ldots,n\}$, $\|b_j - b_{j-1}\|_1 = 1$ and $h_{j-1} - h_j \in \mathrm{Sign}(h_{j-1})$,
- $h_i h_j \geq 0$ for all $i,j \in \{0,\ldots,n\}$,
- $(\hat{\imath},\hat{\tau})$ is adjacent to (i',τ') or $(\hat{\imath},\hat{\tau}) = (i',\tau')$.

Moreover, if $(\hat{\imath},\hat{\tau}) = (i',\tau')$ we say that (i,τ) and (i',τ') are *diagonally linked.* We say for two distinct cells (i,τ) and (i',τ') are *single diagonally connected* if (i,τ) is diagonally connected to (i', τ') or if (i', τ') is diagonally connected to (i, τ). Finally, we say two distinct cells (i, τ) and (i', τ') are *double diagonally connected*, if there exists $(\hat{i}, \hat{\tau})$ such that (i, τ) is diagonally connected to $(\hat{i}, \hat{\tau})$, (i', τ') is diagonally connected to $(\hat{i}, \hat{\tau})$, and $(\hat{i}, \hat{\tau})$ is diagonally linked to (i, τ) or (i', τ').

Note that unlike the definition from Section 4.3 of a cell (i, τ) being diagonally connected to (i', τ'), two cells being single or double diagonally connected is a symmetric relationship.

Definition 8.1. We say a sequence of cells $(i_0, \tau_0), (i_1, \tau_1), \ldots, (i_n, \tau_n)$ is a DD-path if for all $j \in \{1, \ldots, n\}$, we have that the cells (i_{j-1}, τ_{j-1}) and (i_j, τ_j) are adjacent, single diagonally connected or double-diagonally connected.

Recall from Section 4 the definition of cells of multiple scales. Now we will extend the definition of DD-paths to multiple scales, as we did in Section 5 for D-paths. We say (k, i, τ) and (k', i', τ') are *single diagonally connected* if there exists a cell $(1, \hat{i}, \hat{\tau})$ that is a descendant of (k, i, τ) and a cell $(1, i'', \tau'')$ that is a descendant of (k', i', τ'), such that $(1, \hat{i}, \hat{\tau})$ and $(1, i'', \tau'')$ are single diagonally connected. We say (k, i, τ) and (k', i', τ') are *double diagonally connected* if there exists a cell $(1, \hat{i}, \hat{\tau})$ that is a descendant of (k, i, τ) and a cell $(1, i'', \tau'')$ that is a descendant of (k', i', τ'), such that $(1, \hat{i}, \hat{\tau})$ and $(1, i'', \tau'')$ are double diagonally connected.

We refer to a *DD-path* as a sequence of distinct cells of possibly different scales for which any two consecutive cells in the sequence are either adjacent, single diagonally connected or double diagonally connected to the second.

We say two cells (k_1, i_1, τ_1) and (k_2, i_2, τ_2) are *support connected with single diagonals* if there exists a scale 1 cell contained in $R_{k_1}^{2\text{sup}}(i_1, \tau_1)$ and a scale 1 cell contained in $R_{k_2}^{2\text{sup}}(i_2, \tau_2)$ such that the two cells are single diagonally connected. We say two cells (k_1, i_1, τ_1) and (k_2, i_2, τ_2) are *support connected with double diagonals* if there exists a scale 1 cell contained in $R_{k_1}^{2\text{sup}}(i_1, \tau_1)$ and a scale 1 cell contained in $R_{k_2}^{2\text{sup}}(i_2, \tau_2)$, such that the two are double diagonally connected.

Recall from Section 5 the definitions of two cells being well separated and support adjacent. Finally, we define a sequence of cells $P = ((k_1, i_1, \tau_1), (k_2, i_2, \tau_2), \ldots, (k_z, i_z, \tau_z))$ to be a *support connected DD-path* if the cells in P are mutually well separated and, for each $j = 1, 2, \ldots, z-1$, (k_j, i_k, τ_j) and $(k_{j+1}, i_{j+1}, \tau_{j+1})$ are support adjacent, support connected with single diagonals or support connected with double diagonals.

8.1 Multi-scale analysis of DD-paths

We now follow the steps of Section 5, presenting only the parts where the statements and proofs with DD-paths differ from how they were for D-paths.

Define Ω_t to be the set of all DD-paths of cells of scale 1 such that the first cell of the path is $(0,0)$ or $(0,0)$ is single diagonally connected to the first cell, and the last cell of the path is the only cell not contained in $[-t,t]^d \times [-t,t]$. Also, define $\Omega^{\text{sup}}_{\kappa,t}$ as the set of all support connected DD-paths of cells of scale at most κ so that the extended support of the first cell of the path contains $R_1(0,0)$ or $(0,0)$ is single diagonally connected to a scale 1 cell that is contained in the extended support of the first cell of the path, and the last cell of the path is the only cell whose extended support is not contained in $[-t,t]^d \times [-t,t]$. Then the lemma below states that we can focus on support connected DD-paths instead of DD-paths with bad ancestry; the proof is identical to the one of Lemma 5.1.

Lemma 8.1. *We have that*

$$\begin{aligned}&\mathbb{P}\left[\exists P \in \Omega_t \text{ s.t. all cells of } P \text{ have a bad ancestry}\right]\\ &\leq \mathbb{P}\left[\exists P \in \Omega^{sup}_{\kappa,t} \text{ s.t. all cells of } P \text{ are multi-scale bad}\right].\end{aligned}$$

We next have to show that the bound from Lemma 5.5 holds for DD-paths as well.

Lemma 8.2. *Let z be a positive integer and $k_1, k_2, \ldots, k_z \geq 1$ be fixed. Then, if α is sufficiently large, the total number of support connected DD-paths, containing z cells of scales $k_1, k_2, \ldots, k_z$ is at most $\exp\left(\frac{c_3}{2}\sum_{j=1}^{z}\psi_{k_j}\right)$, where c_3 is the same constant as in Lemma 5.4 and ψ is as defined in (22).*

Proof. The proof follows the same steps as the proof of Lemma 5.5. The only changes are that the first cell of a DD-path need not contain $(0,0)$ and the number of different relative positions in step 3 of the proof.

For the former, we note that the extended support of the first cell of the support connected DD-path still has to contain $(0,0)$ or $(0,0)$ has to be single diagonally connected to a scale 1 cell in the extended support of the first cell. If we define χ_{k_1} as in Lemma 5.5, then the first case is already counted by χ_{k_1}. Otherwise, note that if we fix the relative position of the first and final cell of the single diagonal connecting $(0,0)$ to the extended support of the first cell, we only need to control the number of such relative positions, which is done in step 3. Therefore, it only remains to prove step 3 of the proof for DD-paths.

Consider two consecutive cells of the DD-path that are single diagonally connected and let $(1, i, \tau)$ be a cell contained in the extended support of the first cell that is single diagonally connected to a cell $(1, i', \tau')$ that is contained in the extended support of the second cell. Then, as in the proof of Lemma 5.5 we can define

$$A(x) = \max_{(b_1,h_1)\in\mathbb{Z}^{d+1}} |\{(b_2, h_2) \in \mathbb{Z}^{d+1} : |h_2 - h_1| = x \text{ and } (b_1, h_1) \text{ is diagonally connected to } (b_2, h_2)\}|.$$

Consider now two consecutive cells of the DD-path that are double diagonally connected and let $(1, i, \tau)$ be a cell contained in the extended support of the first cell that is double diagonally connected to a cell $(1, i', \tau')$ that is contained in the extended support of the second cell. Furthermore, let $(1, i'', \tau'')$ be the cell of the double diagonal that $(1, i, \tau)$ or $(1, i', \tau')$ is diagonally linked to. Then, if x is the height difference between $(1, i, \tau)$ and $(1, i'', \tau'')$ and y is the height difference between $(1, i', \tau')$ and $(1, i'', \tau'')$, we can bound the number of different relative positions of $(1, i', \tau')$ with respect to $(1, i, \tau)$, such that the height difference between $(1, i, \tau)$ and $(1, i'', \tau'')$ is x and the height difference between $(1, i', \tau')$ and $(1, i'', \tau'')$ is y by $A(x + 1)A(y + 1)$.

Let H_k be the side length of the cube $S_k^{2\text{sup}}(i)$ relative to $S_1(i)$, as in the proof of Lemma 5.5. Therefore, given the z cells of scales $k_1, k_2, \ldots, k_z$, the maximum number of scale 1 diagonal steps contained in all single and double diagonal connections between the cells of the path is at most

$$H := 2\sum_{i=1}^{z-1} H_{k_i}.$$

For notational convenience, when two consecutive cells (k, i, τ) and (k', i', τ') of the DD-path are double diagonally connected, we now consider as part of the path also the cell $(1, i'', \tau'')$ of the double diagonal that both (k, i, τ) and (k', i', τ') are diagonally connected to. Then letting x_i, for $i \in \{1, 2, \ldots, 2z-1\}$ be the height difference between two diagonally connected cells, with $x_i = 0$ if the cells are support adjacent, we have that the number of possible configurations of the diagonal steps is at most

$$\sum_{y=0}^{H} \sum_{\substack{x_1,x_2,\ldots,x_{2z-1} \\ x_1+\cdots+x_{2z-1}=y}} A(x_1 + 1)A(x_2 + 1)\cdots A(x_{2z-1} + 1). \tag{39}$$

As in the proof of Lemma 5.5, we have that

$$A(x_1 + 1)A(x_2 + 1)\cdots A(x_{2z-1} + 1) \le A\left(\tfrac{y}{2z-1} + 1\right)^{2z-1}.$$

Next, using the above bound and

$$\sum_{\substack{x_1,x_2,\ldots,x_{2z-1}:\\ x_1+\cdots+x_{2z-1}=y}} 1 = \binom{2z+y-2}{2z-2},$$

we have that the sum in (39) is smaller than

$$\sum_{y=0}^{H} \binom{2z+y-2}{2z-2} A\left(\tfrac{y}{2z-1}+1\right)^{2z-1} \leq \binom{2z+H}{2z} A\left(\tfrac{H}{2z-1}+1\right)^{2z-1},$$

where the binomial inequality used can easily be proven by induction (using Pascal's rule).

Then, for some positive constants C and C_2, we have

$$\begin{aligned}\binom{2z+H}{2z} A\left(\tfrac{H}{2z-1}+1\right)^{2z-1} &\leq C\frac{(2z+H)^{2z}}{(2z)!}\left(\frac{H}{2z-1}+1\right)^{(2z-1)d}\\ &\leq C\frac{(2z+H)^{2z}}{(2z/3)^{2z}}\left(\frac{2H}{2z}\right)^{2zd}\\ &\leq C(3+3H/(2z))^{2z}\left(\frac{H}{z}\right)^{2zd}\\ &\leq C\left(C_2\frac{H}{z}\right)^{4zd}.\end{aligned}$$

In order to complete the proof, it remains to show that $C\left(C_2\frac{H}{z}\right)^{4zd} \leq \exp\left\{\frac{c_3}{8}\sum_{j=1}^{z}\psi_{k_j}\right\}$, which is equivalent to showing that

$$\tilde{C}z\log\left(\frac{H}{z}\right) \leq \sum_{j=1}^{z}\psi_{k_j}, \tag{40}$$

where $\tilde{C}$ is some constant. Setting m and α sufficiently large, this holds using the same argument as in the proof of Lemma 5.5. □

Similar to Lemma 5.5, if $d \geq 3$ we have that Lemma 8.2 holds also when we set time to be height in the base-height index. For $d = 2$ one can construct a similar counterexample as the one outlined in Remark 5.2.

Lemma 8.3. *Let $t > 0$ and let $P = ((k_1,i_1,\tau_1),(k_2,i_2,\tau_2),\ldots,(k_z,i_z,\tau_z))$ be a path in $\Omega^{sup}_{\kappa-1,t}$. If α is sufficiently large and $\kappa = \mathcal{O}(\log t)$, then there exist a positive constant*

$c = c(C_M)$ *and a value* C *independent of* t *such that*

$$\sum_{j=1}^{z} \psi_{k_j} \geq \begin{cases} C\frac{\sqrt{t}}{(\log t)^c}, & \text{for } d = 1, \\ C\frac{t}{(\log t)^c}, & \text{for } d = 2, \\ Ct, & \text{for } d \geq 3. \end{cases} \tag{41}$$

Proof. The proof is identical to the proof of Lemma 5.6, save for one change. In Lemma 5.6, when considering the sum across the cells of the path, we require that $\sum_{j=1}^{z} \Delta_{k_j}^{2\sup} \geq t/2$. Since we now consider two diagonals per cell instead of just one, the term on the right has to be changed to $t/3$ in order for the statement to still hold. The rest of the proof is unchanged. □

We now define the analogous set of $K(i,\tau)$ for DD-paths. Given an increasing event $E_{\mathrm{st}}(i,\tau)$, let $E(i,\tau)$ be the indicator random variable of $E_{\mathrm{st}}(i,\tau)$.

Definition 8.2. Let $(i,\tau) \in \mathbb{Z}^{d+1}$. If $E(i,\tau) = 1$, define $K^*(i,\tau) = \emptyset$. Otherwise define $K^*(i,\tau)$ as the set

$$\{(i',\tau') \in \mathbb{Z}^{d+1} : E(i',\tau') = 0 \text{ and } \exists \text{ a } DD\text{-path of bad cells from } (i,\tau) \text{ to } (i',\tau')\}.$$

Proposition 8.1. *For each* $(i,\tau) \in \mathbb{Z}^{d+1}$, *let* $E_{st}(i,\tau)$ *be an increasing event that is restricted to the super cube* i *and the super interval* τ, *and let* $\nu_{E_{st}}$ *be the probability associated to* E_{st} *as defined in Definition 2.3. Fix a constant* $\epsilon \in (0,1)$, *and integer* $\eta \geq 1$ *and the ratio* $\beta/\ell^2 > 0$. *Fix also* w *such that*

$$w \geq \sqrt{\frac{\eta\beta}{c_2\ell^2} \log\left(\frac{8c_1}{\epsilon}\right)},$$

for some constants c_1 *and* c_2 *which depend only on the graph. Then, there exist constants* c *and* C, *and positive numbers* α_0 *and* t_0 *that depend on* ϵ, η *and the ratio* β/ℓ^2 *such that if*

$$\alpha = \min\left\{C_M^{-1}\epsilon^2\lambda_0\ell^d, \log\left(\frac{1}{1-\nu_{E_{st}}((1-\epsilon)\lambda, Q_{(2\eta+1)\ell}, Q_{w\ell}, \eta\beta)}\right)\right\} \geq \alpha_0,$$

we have for all $t \geq t_0$ *that*

$$\mathbb{P}\left[K^*(0,0) \not\subseteq \mathcal{R}_1^t\right] \leq \begin{cases} \exp\left\{-C\lambda_0\frac{t}{(\log t)^c}\right\} & \text{for } d = 2 \\ \exp\left\{-C\lambda_0 t\right\} & \text{for } d \geq 3. \end{cases}$$

Proof. The proof of this result proceeds in the same way as the proof of Proposition 6.1, by replacing Lemma 5.1 with Lemma 8.1, Lemma 5.5 with Lemma 8.2 and Lemma 5.6 with Lemma 8.3. □

We now argue that Proposition 8.1 implies that the Lipschitz surface not only almost surely exists as shown in Theorem 2.1, but that areas of the surface that have non-zero height are finite as well. To see why, denote with $u_i = (b_i, 0)$ sites in $\mathbb{L}$ and consider a path along the surface F. More precisely, let $\pi = \{(b_1, F_+(u_1)), (b_2, F_+(u_1)), \ldots, (b_n, F_+(u_n))\}$ be such that $F_+(u_i) \neq 0$ for all $i \in \{1, \ldots, n\}$ and $\|u_i - u_{i-1}\|_1 = 1$ for all $i \in \{2, \ldots, n\}$. If such a path exists, then both sides of the Lipschitz surface have non-zero height at least at the cells of the path, so one can follow the path $(u_1, u_2 \ldots, u_n)$ and never reach the Lipschitz surface F. Conversely, if a path π as above that leaves a ball of finite radius does not exist, a self-avoiding path will have to reach the surface in finitely many steps. Furthermore, since time is one of the $d+1$ dimensions, one cannot construct a time directed path without it containing a cell $(b, F_+(u))$ or $(b, F_-(u))$ for some $u = (b, 0) \in \mathbb{L}$ within a finite number of steps. This follows from the fact that by Theorem 2.1 the surface is a.s. finite, so a path can avoid intersecting it indefinitely only if there is always at least one way to construct a path between to the two sides of the surface. If however, paths along which the two sides of the surface have non-zero height cannot have arbitrary length, we get that avoiding the two sides indefinitely is impossible.

To simplify things, we first observe that we can limit ourselves to only the positive Lipschitz open surface, since $F_+(u) = 0$ if and only if $F_-(u) = 0$, by the definition of the two sides of the surface.

Recall from Section 3 the definition of a hill H_u. In the following, we will use H_i, $i \in \mathbb{Z}$ to differentiate between different hills without specifying a cell $u \in \mathbb{L}$ for which $H_i = H_u$. We now show that the existence of a path along the surface with only positive heights implies the existence of a sequence of hills that are pairwise intersecting or adjacent. Formally, we define the following.

Definition 8.3. We say a hill H_i is *adjacent* to a hill $H_{i'}$, if there exist a cell $u \in H_i$ and a cell $v \in H_{i'}$ such that $\|u - v\|_1 = 1$. We say H_i and $H_{i'}$ are *intersecting*, if there exists a cell u such that $u \in H_i$ and $u \in H_{i'}$.

Lemma 8.4. *Write $u_i = (b_i, 0) \in \mathbb{L}$ and let $\pi = \{(b_1, F_+(u_1)), (b_2, F_+(u_2)), \ldots, (b_n, F_+(u_n))\}$ be a path, such that for all $i \in \{1, \ldots, n\}$, $F_+(u_i) \neq 0$. Then there exists a sequence of hills $\mathcal{H} = H_1, H_2, \ldots, H_k$, $k \leq n$, such that for every u_ℓ there exists a hill $H_k \in \mathcal{H}$ that contains u_ℓ, and such that for all $i \in \{1, \ldots k\}$, there exists at least one $j \neq i$, $j \in \{1, \ldots, k\}$, for which H_i intersects with or is adjacent to H_j.*

Proof. We will prove the existence of the sequence of hills iteratively. Let $\mathcal{H} = \emptyset$ be the set of all hills that are part of the sequence already. We then add hills to $\mathcal{H}$ in the following manner. Let $u = (b, 0)$ be the first cell of the path $P = ((b_1, 0), (b_2, 0), \ldots, (b_n, 0))$ that is not contained in $\bigcup_{H \in \mathcal{H}} H$. Since $F_+(u) \neq 0$ by assumption, there has to exist at least one cell $v \in \mathbb{L}$ such that $u \in H_v$. Since the cell u is contained in H_v and it is adjacent to at least 1 cell contained in $\bigcup_{H \in \mathcal{H}} H$ (except for when $\mathcal{H} = \emptyset$), we get that H_v and at least one hill from $\mathcal{H}$ are adjacent or they intersect. We add H_v to $\mathcal{H}$, remove all cells of P that are contained in H_v from P, and repeat the procedure. After at most n steps, the recursion ends and $\mathcal{H}$ is a set of k hills for some $k \leq n$, such that every hill intersects or is adjacent to at least one other hill in the set. □

We now want to show that if the sequence of hills $\mathcal{H}$ from Lemma 8.4 exists, then a DD-path exists between any two cells contained in $\bigcup_{H \in \mathcal{H}} H$.

Lemma 8.5. *Let $\mathcal{H} = H_1, H_2, \cdots H_k$ be a sequence of hills as in Lemma 8.4. For any two $(b, 0), (b', 0) \in \bigcup_{H_i \in \mathcal{H}} H_i$, there exists a DD-path that starts in $(b, 0)$ and ends in $(b', 0)$.*

Proof. Let $u_1, u_2, \ldots, u_k \in \mathbb{L}$ be the cells such that $H_i = H_{u_i}$ for all $i \in \{1, 2, \ldots, k\}$. Next, observe that by the definition of $\mathcal{H}$, there exists a sequence of hills $H_{i_1}, H_{i_2}, \ldots, H_{i_\ell}$ such that $(b, 0) \in H_{i_1}$, $(b', 0) \in H_{i_\ell}$ and every hill in the sequence is adjacent or intersecting with the subsequent hill. For every $j \in \{1, 2, \ldots, \ell\}$, let $v_{i_j} \in H_{i_j}$ be a cell that is contained in $H_{i_{j+1}}$ or adjacent to a cell in $H_{i_{j+1}}$.

By definition of a hill, there exists a d-path P_1 from u_{i_1} to $(b, 0)$. Furthermore, there exists a d-path P_2 from u_{i_1} to v_{i_1} and a d-path P_3 from u_{i_2} to v_{i_1} or a cell that is adjacent to v_{i_1}. By repeating this, we obtain the sequence of cells

$$(b, 0), u_{i_1}, v_{i_1}, u_{i_2}, v_{i_2}, \ldots, v_{i_{\ell-1}}, u_{i_\ell}, (b', 0),$$

where there exists a d-path from the first cell to the second or from the second to the first (or a cell adjacent to it) for every consecutive pair of cells. It remains to show that this implies that there exists a DD-path from $(b, 0)$ to $(b', 0)$.

Note first that similar to D-paths, every d-path is also a DD-path. This follows directly from the fact that DD-paths are defined as an extension of D-paths. Next, note that if a sequence of cells $(w_1, w_2, \ldots, w_n) \in \mathbb{Z}^{d+1}$ is a DD-path, then the reverse sequence, i.e. $(w_n, w_{n-1}, \ldots, w_1)$ is also a DD-path. This follows trivially from the fact that being adjacent, single diagonally connected and double diagonally connected are all symmetric relationships between cells. Finally, note that if there exists a DD-path from a cell w_1 to some cell w_2 and there exists a DD-path from w_2 to w_3, then there exists at least one DD-path from w_1 to w_3. Once such path can be constructed by concatenating the two DD-paths and removing any cells in the concatenated path that would result in loops, i.e. if a site appears in the concatenated path more than once, remove from the path all sites between the first and last appearance of the site in the path, as well as the last appearance of the site.

Then, using these facts with the sequence

$$(b,0), u_{i_1}, v_{i_1}, u_{i_2}, v_{i_2}, \dots, v_{i_{\ell-1}}, u_{i_\ell}, (b',0)$$

concludes the lemma. □

We are now ready to state our main result. Theorem 2.2 is a direct consequence of the following theorem.

Theorem 8.1. *Let $\mathcal{G} = (G,\mu)$ be a graph satisfying (1) on the lattice $\mathbb{Z}^d$ for $d \geq 2$, tessellated into cubes of side-length ℓ, and indexed by $i \in \mathbb{Z}^d$. Let time be tessellated into intervals of length β, indexed by τ, let $\mathbb{L} = \{(b,h) \in \mathbb{Z}^{d+1} : \ h = 0\}$. Let $\eta \geq 1$ and let $E_{st}(i,\tau)$ be an increasing event restricted to the super cell (i,τ). Fix a constant $\epsilon \in (0,1)$ and the ratio $\beta/\ell^2 > 0$. Fix also w such that*

$$w \geq \sqrt{\frac{\eta\beta}{c_2\ell^2}\log\left(\frac{8c_1}{\epsilon}\right)},$$

for some constants c_1 and c_2 which depend on the graph. Then, there exists a positive number α_0 and r_0 so that if

$$\alpha = \min\left\{C_M^{-1}\epsilon^2\lambda_0\ell^d, \log\left(\frac{1}{1-\nu_{E_{st}}((1-\epsilon)\lambda, Q_{2(\eta+1)\ell}, Q_{w\ell}, \eta\beta)}\right)\right\} \geq \alpha_0,$$

a Lipschitz surface F such that $E_{\mathrm{bh}}(b,h)$ holds for all $(b,h) \in F$ exists a.s. and surrounds the origin at an a.s. finite distance. Furthermore, we have for $r \geq r_0$ that

$$\mathbb{P}\left[\begin{array}{c} F \textit{ does not surround} \\ \textit{the origin at distance } r \end{array}\right] \leq \begin{cases} \sum_{s\geq r} s^d \exp\{-C\lambda_0 \frac{\ell s}{(\log \ell s)^c}\}, & \textit{for } d = 2 \\ \sum_{s\geq r} s^d \exp\{-C\lambda_0 \ell s\}, & \textit{for } d \geq 3. \end{cases}$$

for some constant C.

Proof. Note that the open Lipschitz surface exists a.s. by Theorem 3.1, so we only need to show that it surrounds the origin at some finite distance.

Assume the converse. Then, for any $r > 0$ there must exist a path of adjacent cells

$(0,0) = (b_1, 0), \dots, (b_n, 0)$ with $\|(b_n, 0)\|_1 > r$, such that $F_+((b_i, 0)) \neq 0$ for all $i \in \{1, \dots, n\}$. By Lemma 8.4, this implies the existence of a sequence of hills such that the first one contains the origin and the last one contains $(b_n, 0)$. By Lemma 8.5, this gives the existence of a DD-path from the origin to $(b_n, 0)$.

Note that by Proposition 8.1, for $t \geq t_0$ we have that the probability that such a DD-path exists is smaller than

$$\mathbb{P}[K^*(0,0) \not\subseteq \mathcal{R}_1^t] \leq \begin{cases} \exp\left\{-C\lambda_0 \frac{t}{(\log t)^c}\right\} & \text{for } d = 2 \\ \exp\left\{-C\lambda_0 t\right\} & \text{for } d \geq 3. \end{cases}$$

From here, setting $t = \left(\frac{d}{\ell} + \frac{1}{c\ell^2}\right)^{-1} r$ and using the same steps as in the proof of Theorem 2.1 establishes the claim for $r \geq r_0 := \left(\frac{d}{\ell} + \frac{1}{c\ell^2}\right) t_0$. □

Next, we show that Theorem 2.3 holds. Observe first the following well known geometric property. Let B^2 be the plane spanned by any two base vectors of the base-height index. Recall also the definition of $\mathbb{L} = \{(x, 0), x \in \mathbb{Z}^d\}$, the zero-height hyperplane of $\mathbb{Z}^{d+1}$. It then holds that

$$\begin{aligned} &\mathbb{P}\left[\text{zero height cells percolate in } B^2\right] \\ &\leq \mathbb{P}\left[\text{zero height cells percolate in } \mathbb{L}\right], \end{aligned}$$

since it clearly holds that the first event implies the second. Therefore, it is enough to show that the first probability is positive for Theorem 2.3 to hold.

Corollary 8.1. *Let $d = 2$ and let $E_{st}(i, \tau)$ be an increasing event restricted to the super cell (i, τ). If ℓ is sufficiently large and $\mathbb{P}[E_{st}(0, 0)]$ is large enough, then $F \cap \mathbb{L}$ percolates within $\mathbb{L}$ with positive probability.*

Proof. Assume without loss of generality that the origin is contained in $F \cap \mathbb{L}$ and assume that the cluster of $F \cap \mathbb{L}$ that contains the origin is finite. Let $x, y \in F \cap \mathbb{L}$ be two cells of this cluster for which $\|x - y\|_1$ is largest, and let $k := \lceil \|x - y\|_1 \rceil$. Then, there exists a sequence cells $v_1, v_2, \dots, v_n \in \mathbb{L}$ for some $n \geq 2k$ such that $F_+(v_i) > 0$ for all $i \in \{1, 2, \dots, n\}$, any two consecutive cells are adjacent, i.e. $\|v_i - v_i'\|_\infty = 1$, and such that $\|v_n - v_1\|_\infty = 1$. Furthermore, each such sequence contains at least 2 cells $u, v \in \mathbb{L}$ for which $\|u - v\|_1 \geq k$. By using Lemmas 8.4 and 8.5, this gives that there exists a DD-path that begins in u and ends in v. Let r_0 be defined as in Theorem 8.1. We then have for $k \geq r_0$, by using that the probability space is space and time translation invariant that

$$\begin{aligned}
&\mathbb{P}[\text{the cluster } F \cap \mathbb{L} \text{ around the origin has diameter } k] \\
&\leq \mathbb{P}[\text{a } DD\text{-path started at the origin leaves the ball of radius } k \text{ centered at the origin}] \\
&\leq \exp\left\{-C\lambda \frac{\ell k}{(\log \ell k)^c}\right\},
\end{aligned}$$

where the second inequality follows from the same argument as in the proof of Theorem 8.1. For $k < r_0$, we can bound the probability by $Cr_0^2\mathbb{P}[E_{\text{st}}(0,0)^c]$ for some positive constant C, since a closed cell implies that the Lipschitz function F_+ is non-zero. Therefore, we get that the probability the zero-height cluster $F \cap \mathbb{L}$ at the origin is not finite is greater than

$$1 - Cr_0^2\mathbb{P}[E_{\text{st}}(0,0)^c] - \sum_{k \geq r_0} \exp\left\{-C\lambda \frac{\ell k}{(\log \ell k)^c}\right\},$$

which is positive for sufficiently large ℓ, if $\mathbb{P}[E_{\text{st}}(0,0)]$ is large enough. □

A Appendix: Standard results

Lemma A.1 (Chernoff bound for Poisson)**.** *Let P be a Poisson random variable with mean λ. Then, for any $0 < \epsilon < 1$,*

$$\mathbb{P}[P < (1-\epsilon)\lambda] < \exp\{-\lambda\epsilon^2/2\}$$

and

$$\mathbb{P}[P > (1+\epsilon)\lambda] < \exp\{-\lambda\epsilon^2/4\}.$$

Lemma A.2. *Let $x, y \in \mathbb{Z}_+$. Then, for any $c_1, c_2 > 1$, we have*

$$\binom{x+y}{x} e^{-(c_1 x + c_2 y)} \leq e^{-(c_1-1)x-(c_2-1)y}.$$

Proof. Since $\binom{x+y}{x} = \binom{x+y}{y}$, we can assume that $x \geq y$. Then we use the inequality $\binom{x+y}{x} \leq \left(\frac{(x+y)e}{x}\right)^x$ to obtain

$$\binom{x+y}{x} e^{-c_1 x - c_2 y} \leq \left(1 + \frac{y}{x}\right)^x e^{-(c_1-1)x - c_2 y} \leq e^{-(c_1-1)x-(c_2-1)y}.$$

□

References

[1] Barlow, M.T., 2004. Random walks on supercritical percolation clusters. *Annals of Probability*, 32(4), pp.3024–3084. Available from: http://doi.org/10.1214/009117904000000748.

[2] Barlow, M.T. and Hambly, B.M., 2009. Parabolic Harnack inequality and local limit theorem for percolation clusters. *Electronic Journal of Probability*, 14, pp.1–26. Available from: http://doi.org/10.1214/EJP.v14-587.

[3] Candellero, E. and Teixeira, A., 2015. Percolation and isoperimetry on transitive graphs. *arXiv*. 1507.07765, Available from: http://arxiv.org/abs/1507.07765.

[4] Dirr, N., Dondl, P.W., Grimmett, G.R., Holroyd, A.E. and Scheutzow, M.V., 2010. Lipschitz percolation. *Electronic Communications in Probability*, 15, pp.1–10. Available from: http://doi.org/10.1214/ECP.v15-1521.

[5] Gracar, P. and Stauffer, A., 2017. Random walks in random conductances: decoupling and spread of infection. *arXiv*. 1701.08021, Available from: https://arxiv.org/abs/1701.08021.

[6] Grimmett, G.R. and Holroyd, A.E., 2012. Geometry of Lipschitz percolation. *Annales de l'institut Henri Poincare (B) Probability and Statistics*, 48(2), pp.309–326. Available from: http://doi.org/10.1214/10-AIHP403.

[7] Kesten, H. and Sidoravicius, V., 2005. The spread of a rumor or infection in a moving population. *Annals of Probability*, 33(6), pp.2402–2462. Available from: http://doi.org/10.1214/009117905000000413.

[8] Kesten, H. and Sidoravicius, V., 2006. A phase transition in a model for the spread of an infection. *Illinois Journal of Mathematics*, 50(3), pp.547–634.

[9] Peres, Y., Sinclair, A., Sousi, P. and Stauffer, A., 2013. Mobile geometric graphs: Detection, coverage and percolation. *Probability Theory and Related Fields*, 156(1-2), pp.273–305. Available from: http://doi.org/10.1007/s00440-012-0428-1.

[10] Sidoravicius, V. and Stauffer, A., 2016. Multi-Particle Diffusion Limited Aggregation. *arXiv*. 1603.03218, Available from: http://arxiv.org/abs/1603.03218.

[11] Sidoravicius, V. and Sznitman, A.S., 2009. Percolation for the vacant set of random interlacements. *Communications on Pure and Applied Mathematics*, 62(6), pp.831–858. arXiv:0808.3344v1, Available from: http://doi.org/10.1002/cpa.20267.

[12] Sidoravicius, V. and Teixeira, A., 2014. Absorbing-state transition for Stochastic Sandpiles and Activated Random Walks. *arXiv*. 1412.7098, Available from: http://arxiv.org/abs/1412.7098.

[13] Stauffer, A., 2015. Space-time percolation and detection by mobile nodes. *Annals of Applied Probability*, 25(5), pp.2416–2461. Available from: http://doi.org/10.1214/14-AAP1052.

[14] Sznitman, A.S., 2012. Decoupling inequalities and interlacement percolation on G×ℤ. *Inventiones mathematicae*, 187(3), pp.645–706. Available from: http://doi.org/10.1007/s00222-011-0340-9.

[15] Teixeira, A., 2016. Percolation and local isoperimetric inequalities. *Probability Theory and Related Fields*, 165(3-4), pp.963–984. Available from: http://doi.org/10.1007/s00440-015-0653-5.

Closing remarks

The results of this chapter can be summarized as three distinct, but equally important ideas. The first is the notion that by tessellating space and time into cells of multiple scales, the problem becomes similar to *fractal percolation* in the sense that *good* cells can be further subdivided into good and bad cells of a smaller scale, whereas descendants of *bad* cells are automatically bad. More precisely, we use *multi-scale analysis* to argue the following. If some cell (k, i, τ) is part of a D-path, we know that its parent cell is necessarily good. Therefore, although the particles are "misbehaving" within (k, i, τ), they are behaving correctly at a larger scale, and so with high probability are also behaving "nicely" in nearby cells (k, i', τ'). As a result, the behavior of particles outside of (k, i, τ) is *roughly* independent of the bad behavior of particles inside (k, i, τ).

The second key idea is the two-sided Lipschitz surface. A similar percolation structure has been studied before by Dirr et al. [4] and Grimmett and Holroyd [5], but they were primarily interested in the existence of the surface for the case of Bernoulli site percolation. Our take on the surface is different. First of all, we considered a much more delicate case of site percolation, where the dependences between the sites cannot be easily controlled. Moreover, our aim goes beyond just showing the existence of the surface; instead, we primarily wanted to use the surface to infer properties of the underlying particle system.

The third and most straightforward idea is to combine the above into a cohesive framework. Defining D-paths and DD-paths lets us connect the Lipschitz surface, defined on the tessellation of space and time, with the behavior of the particles on the graph which we control through multi-scale arguments. The key step that allows us to achieve this is the *local mixing* result from Chapter 2, and careful combinatorial arguments that are needed to control the large number of D-paths and DD-paths which are possible due to the diagonal connections. The result are theorems that based on the probability of a single translation invariant local event E_{st} give rise to a global percolation structure with a wide variety of possible applications.

Chapter 4

Conclusion

To the best of our knowledge, this thesis and the two papers presented herein mark the first time anyone has applied the idea of Lipschitz percolation beyond the case of Bernoulli site percolation. In hindsight, the geometry of the surface (see Theorems 1.4, 1.5 and 1.6 from Chapter 1) provides clear benefits: the Lipschitz property gives more precise control over the shape of the infinite percolation cluster and allows for a wider variety of arguments. The downside lies primarily in proving that this surface structure exists, which becomes increasingly difficult once we move beyond Bernoulli site percolation.

It is for this reason that we have, as much as possible, done the multi-scale analysis of Chapter 3 in a way that gives the existence of the Lipschitz surface based solely on the local events E_{st}, the local behavior of the particles, and the scale of the space-time tessellation. This way, if one wants to show that the surface exists for a uniformly elliptic graph on $\mathbb{Z}^d$ and some increasing local event E_{st}, one does not have to concern oneself with the delicate multi-scale arguments about D-paths or the number of different arrangements of diagonals. Likewise, detailed knowledge of the local mixing result and its proof from Chapter 2 that allow us to do much of the work in Chapter 3 is not necessary when considering the local event E_{st} with which one wants to work. Our application of the framework in the second half of Chapter 2 provides readers with a clear and detailed example of how the conditions of Theorem 1.4 can be satisfied and how the geometry of the Lipschitz surface can be put to use.

4.1 Open problems

A careful reading of the two papers will reveal that while our framework is more robust than previous multi-scale constructions, there are two areas of our analysis where there might be room for improvement. The first is the seemingly arbitrary limitation that in two dimensions time cannot be set as height in the base-height index. The second and more interesting limitation is the choice to restrict ourselves to graphs where $\mu_{x,y} \neq 0$ for all $(x, y) \in E$, despite proving the local mixing result for a broader class of graphs. We believe that neither of these two areas can actually be improved without a significantly different approach and present our reasoning as to why below.

Existence of the Lipchitz surface in two dimensions

For all the questions we have considered, setting *height* to be one of the spatial dimensions is the natural choice. However, it could be that some applications require one to set time as the height of the surface. As shown in Lemma 5.5 and Remark 5.2 in Chapter 3, this is possible when working with three or more spatial dimensions, but not when the graph G is the two-dimensional lattice. We believe that this is a fundamental problem of the approach that cannot be avoided by simply rescaling the cubes or intervals involved. The problem arises from the nature of the multi-scale D-paths with which we work. In higher dimensions, each space-time cell decreases the probability of a path by a bigger factor than it increases the number of different arrangements of cells and diagonal steps (see Lemma 5.5 for details). In two dimensions this is no longer true. Here, as we consider paths with an increasing number of cells, the number of valid arrangements of cells and diagonal connections between them increases at a faster rate than the rate with which the probability of such paths decreases. As such, we believe that an entirely different approach would be needed in order to show the existence of a Lipschitz surface or similar percolation structure in two spatial dimensions, since any rescaling of space or time would still exhibit similar behavior to the one shown in the counterexample of Remark 5.2 in Chapter 3, albeit with different constant terms.

Existence of the Lipschitz surface on percolation clusters

In Chapter 2, we show that the local mixing results holds for uniformly elliptic graphs on the square lattice, which includes the case where the conductances $\mu_{x,y}$ are allowed to be 0, but are otherwise still bounded away from 0; see (1.2) in Chapter 1. Then, since the geometry of the tessellation of space and time remains unchanged and the mixing result still applies, one might naturally expect that the two-sided Lipschitz surface exists also for infinite percolation clusters; for example, if $\mu_{x,y} \in \{0,1\}$ and the weights satisfy (1.2), the resulting weighted graph is an infinite percolation cluster on $\mathbb{Z}^d$. That is not the case however, since allowing the weights to be zero introduces a new, infinite range dependence into the system.

As stated in Theorem 1.3 from Chapter 1, the local mixing result holds whenever the side-length of the subcubes ℓ is sufficiently large. More precisely, Proposition 2.5 from Chapter 2 gives that for weights satisfying (1.2), the probability that ℓ needs to be larger than $n \in \mathbb{N}$ satisfies

$$\mathbb{P}\left[\ell > n\right] \leq \exp\{-cn^{\gamma}\}, \tag{4.1}$$

where c is some constant and γ depends only on the distribution of the weights $\mu_{x,y}$. This is a direct consequence of the results of Barlow and Hambly [3], which require the scale at which we work to be sufficiently large for the weak Poincaré inequality to hold. Intuitively, the scale needs to be large enough to "smooth out" any significant irregularities in the graph geometry that would prevent the heat kernel bounds from holding. Therefore, when considering space-time cells of multiple scales and applying Theorem 1.3, one is confronted with various options for how to ensure that all subcubes are large enough.

The first is to set ℓ to be the smallest value for which *all* cubes of side length ℓ of the tessellation are sufficiently large for the local mixing to occur. For finite graphs this can easily be achieved with positive probability by using a uniform bound and (4.1). For graphs of infinite size, such as the above mentioned infinite percolation cluster on $\mathbb{Z}^d$, there will however almost surely always exist cubes of the tessellation for which ℓ is not sufficiently large with respect to the local configuration of the weights $\mu_{x,y}$.

The second option is to modify the fractal percolation process from Section 4.3 in Chapter 3 and consider cubes for which the side-length of the subcubes is not large enough as "geometrically bad". This however means that any space-time cell (regardless of scale) that is made up of a geometrically bad cube i and a time interval τ cannot be *good* in the sense of Chapters 2 and 3, since the local mixing result cannot be applied. We therefore have no control over the particle behavior in these cells. With this in mind, we could proceed as before and define D-paths of bad cells, where a cell is considered bad whenever it is not dense enough in the sense of Chapter 3 or when the mixing result cannot be applied due to the size of the cell. By doing this however, we introduce infinite range dependences into the construction. Consider a space-time cell (k, i, τ) where due to the configuration of weights the side-length of the subcubes is not large enough and therefore the mixing result does not apply. By (4.1) such a cell exists almost surely. Then, any other cell (k, i, τ') with $\tau' \neq \tau$ will also exhibit the same issue. Therefore, any D-path that contains a space-time cell (k, i, τ) with this property can be extended indefinitely by fixing i and increasing τ. As a consequence, the two-sided Lipschitz no longer exists almost surely.

Alternatively, we could treat these cells as "almost good" since they are not necessarily bad in the sense of Chapters 2 and 3. This would again give the almost sure existence of a two-sided Lipschitz surface. Note however that this surface might also contain space-time cells where the behavior of the particles is unknown due to the geometry of the graph. The biggest downside of such a surface is that we can no longer use the arguments from Chapter 2 to start the spread of infection once an infected particle enters the surface (since this might happen in a cell that is "geometrically bad"), nor can we guarantee that after the infection spreads from a good space-time cell, an infected particle is in at least one other good cell on the surface one time step later. We expect that in this case, the infection still spreads with positive speed the majority of the time, since these "almost good" cells are sufficiently unlikely so that they do not percolate within the two-sided Lipschitz surface. We cannot say anything about the infection's speed in the "geometrically bad" regions however, as they can both accelerate or slow down the spread, depending on the local configuration of the weights $\mu_{x,y}$.

It would take a fundamentally different construction in order to be able to avoid this problem, but we expect that such a result would not be as easily applicable since we would either have to control the geometry of the graph via the increasing events or deal with the infinite range dependences in some other way.

4.2 Future work

Existence of the Lipschitz surface for other particle systems

As seen in Chapter 3, the definition of the two-sided Lipschitz surface does not depend directly on the underlying particle system. Instead, we define the surface on the tessellation of space and time via the increasing events E_{st}. Therefore, a natural question is which other particle systems can one study with this framework. To answer this question, we need to look closely at Lemmas 5.2 and 5.3 in Chapter 3.

Intuitively, Lemma 5.2 states that it is highly unlikely that during a time interval $[0, t]$ a particle will move further than $c\sqrt{t}$ away from where it was at time 0, where c is some positive constant. Therefore, if we write Y_t for the location of a particle at time

t, at the very least, particles should satisfy

$$\mathbb{P}[|Y_t - Y_0| > z] \leq \exp\left\{-c'\frac{z^2}{t}\right\},$$

for some constant $c' > 0$. With Lemma 5.3, a close inspection of the proof reveals that the main component that makes the proof possible is Theorem 4.1 from Chapter 2. Intuitively, the existence of the Lipschitz surface depends on the existence of a stationary distribution and the probability with which particles stay close to such a distribution.

Natural candidates that satisfy these requirements are particles on $\mathbb{R}^d$ that follow Brownian motions, as well as particle systems where a local central limit theorem links their behavior to Brownian motion. Baldasso and Teixeira [1] show a similar local mixing/decoupling result for the exclusion process, so we expect that adapting the two-sided Lipschitz surface framework to that setup would be straightforward. Another likely candidate where the framework could be used are activated random walks, which are often studied with the help of multi-scale analysis arguments.

Lipschitz surface for decreasing events

Throughout this thesis we have limited ourselves to the study of increasing events, as defined in Chapter 1. Alternatively, we could try to show the existence of the two-sided Lipschitz surface defined for *decreasing* local events E_{st}, that is, events such that if $(\Pi_s)_{s\geq 0}$ is a sequence of point processes on $\mathbb{Z}^d$ and E_{st} holds for $(\Pi_s)_{s\geq 0}$, then E_{st} also holds for all $(\Pi'_s)_{s\geq 0}$ for which $\Pi'_s \subseteq \Pi_s$ for all $s \geq 0$. In order to use the local mixing result from Chapter 2, we would first have to modify it so that it would give a lower bound on the probability of an independent Poisson point process stochastically dominating the particles after they move. Note that this domination is the opposite of the one we state in Theorem 1.3 from Chapter 1. Based on our work in Chapter 2, we expect that proving such a mixing result would not be significantly different to how we prove Theorem 1.3.

The main difficulty comes from the definitions of good and bad space-time cells. In Chapter 3, we consider a cell good whenever it contains sufficiently many particles and bad otherwise. Furthermore, when considering the effects of bad cells on other cells, we treat all bad cells as if they were empty of particles. Since we are working with increasing events, this gives a "worst case" bound on the probability that the Lipschitz surface exists. For decreasing events, we would similarly need to define good cells as cells in which the number of particles is sufficiently low. The issue arises when defining bad cells. Since the number of particles inside a cell is unbounded, we cannot treat all bad cells equally. The reason for this comes from the fact that the number of particles we expect to see in cells near a bad cell increases with the number of particles in the bad cell, with no bound on this effect. We could deal with this issue by considering multiple degrees of how bad a cell is, but this significantly increases the complexity of the multi-scale analysis that needs to be done and comes with its own set of issues that have to be addressed.

4.3 Final remarks

Although questions pertaining to spread of infection are of central focus to this thesis, it is the two-sided Lipschitz surface that is our most significant contribution. We believe that the multi-scale analysis of Chapter 3 is quite beautiful in its own right and might inspire others to look at multi-scale arguments in a new light, but it is the "black box" nature of the Lipschitz surface that makes it really stand out. It allows us to separate the problem, such as the spread of infection, away from the intricacies of multi-scale analysis and instead focus on a single local event; one that can be as simple or complex as we want.

This by itself is a large departure from how results of this nature are often proven. Although previous multi-scale arguments still localized the behavior of the particles to some degree in order to show that, for example, particles were located near any path of interest in the paper by Kesten and Sidoravicius [6], no further behavior could be inferred from this. In contrast to this, our approach lets us attach as much additional behavior to the local events as wanted without having to worry if the multi-scale construction still holds. The sole limitation is that this local event needs to remain sufficiently likely. For all of the problems we considered this does not present an issue, since for the events of interest increasing either the intensity of the Poisson point process of particles or the size of the space-time cells on which the event is defined results in an increase of the probability of the event.

We believe this approach could therefore be applied to many unanswered questions that were previously intractable due to the complicated dependences that are inherent in such problems. Questions about percolation times, total infection times and others all exhibit the property that the local behavior of particles results in a global phenomenon. This is why we believe there are several areas of research where our local mixing result combined with the two-sided Lipschitz surface framework could lead to new results without the need to keep "reinventing the wheel" in order to prove them.

Bibliography

[1] Baldasso, R. and Teixeira, A., 2016. How can a clairvoyant particle escape the exclusion process? *arXiv*, p.29. 1605.07575, Available from: http://arxiv.org/abs/1605.07575.

[2] Barlow, M.T., 2004. Random walks on supercritical percolation clusters. *Annals of Probability*, 32(4), pp.3024–3084. Available from: http://doi.org/10.1214/009117904000000748.

[3] Barlow, M.T. and Hambly, B.M., 2009. Parabolic Harnack inequality and local limit theorem for percolation clusters. *Electronic Journal of Probability*, 14, pp.1–26. Available from: http://doi.org/10.1214/EJP.v14-587.

[4] Dirr, N., Dondl, P.W., Grimmett, G.R., Holroyd, A.E. and Scheutzow, M.V., 2010. Lipschitz percolation. *Electronic Communications in Probability*, 15, pp.1–10. Available from: http://doi.org/10.1214/ECP.v15-1521.

[5] Grimmett, G.R. and Holroyd, A.E., 2012. Geometry of Lipschitz percolation. *Annales de l'institut Henri Poincare (B) Probability and Statistics*, 48(2), pp.309–326. Available from: http://doi.org/10.1214/10-AIHP403.

[6] Kesten, H. and Sidoravicius, V., 2005. The spread of a rumor or infection in a moving population. *Annals of Probability*, 33(6), pp.2402–2462. Available from: http://doi.org/10.1214/009117905000000413.

[7] Kesten, H. and Sidoravicius, V., 2006. A phase transition in a model for the spread of an infection. *Illinois Journal of Mathematics*, 50(3), pp.547–634.

[8] Kesten, H. and Sidoravicius, V., 2008. A shape theorem for the spread of an infection. *Annals of Mathematics*, 167(3), pp.701–766. Available from: http://doi.org/10.4007/annals.2008.167.701.

[9] Peres, Y., Sinclair, A., Sousi, P. and Stauffer, A., 2013. Mobile geometric graphs: Detection, coverage and percolation. *Probability Theory and Related Fields*, 156(1-2), pp.273–305. Available from: http://doi.org/10.1007/s00440-012-0428-1.

[10] Popov, S. and Teixeira, A., 2015. Soft local times and decoupling of random interlacements. *Journal of the European Mathematical Society*, 17(10), pp.2545–2593. Available from: http://doi.org/10.4171/JEMS/565.

[11] Sidoravicius, V. and Teixeira, A., 2014. Absorbing-state transition for Stochastic Sandpiles and Activated Random Walks. *arXiv*. 1412.7098, Available from: http://arxiv.org/abs/1412.7098.

[12] Sinclair, A. and Stauffer, A., 2010. Mobile geometric graphs, and detection and communication problems in mobile wireless networks. *arXiv*. 1005.1117v2, Available from: http://arxiv.org/abs/1005.1117.

[13] Stauffer, A., 2015. Space-time percolation and detection by mobile nodes. *Annals of Applied Probability*, 25(5), pp.2416–2461. Available from: http://doi.org/10.1214/14-AAP1052.

[14] Sznitman, A.S., 2010. Vacant Set of Random Interlacements and Percolation. *Annals of Mathematics*, 171, pp.2039–2087.