A Macroscopic View of Two Discrete Random Models

by

Alma Sarai Hernandez Torres

B.Sc. Mathematics, Universidad de Guanajuato, 2014

M.Sc. Mathematics, The University of British Columbia, 2016

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES

(Mathematics)

The University of British Columbia

(Vancouver)

August 2020

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the thesis entitled:

A Macroscopic View of Two Discrete Random Models

submitted by **Alma Sarai Hernandez Torres** in partial fulfillment of the requirements for the degree of **Doctor of Philosophy** in **Mathematics**.

Examining Committee:

Omer Angel, Mathematics
Supervisor

Martin Barlow, Mathematics
Supervisor

Mathav Murugan, Mathematics
Supervisory Committee Member

Lior Silberman, Mathematics
University Examiner

Nick Harvey, Computer Science
University Examiner

Russell Lyons, Indiana University
External Examiner

Abstract

This thesis investigates the large-scale behaviour emerging in two discrete models: the uniform spanning tree on $\mathbb{Z}^3$ and the chase-escape with death process.

Uniform spanning trees

We consider the uniform spanning tree (UST) on $\mathbb{Z}^3$ as a measured, rooted real tree, continuously embedded into Euclidean space. The main result is on the existence of sub-sequential scaling limits and convergence under dyadic scalings. We study properties of the intrinsic distance and the measure of the sub-sequential scaling limits, and the behaviour of the random walk on the UST. An application of Wilson's algorithm, used in the study of scaling limits, is also instrumental in a related problem. We show that the number of spanning clusters of the three-dimensional UST is tight under scalings of the lattice.

Chase-escape with death

Chase-escape is a competitive growth process in which red particles spread to adjacent uncoloured sites while blue particles overtake adjacent red particles. We propose a variant of the chase-escape process called chase-escape with death (CED). When the underlying graph of CED is a d-ary tree, we show the existence of critical parameters and characterize the phase transitions.

Lay Summary

Statistical mechanics states that natural phenomena arise as the average behaviour of a large number of particles with random interactions. A central endeavour in probability theory is to establish a mathematical foundation for this paradigm. Our objective is to obtain precise relations between the microscopic and macroscopic descriptions of a phenomenon. This thesis is a contribution to the task. In particular, we are interested in the macroscopic properties emerging in two discrete random models. In this work, we study the "uniform spanning tree" and the "chase-escape with death" process. The first one is a combinatorial model that provides insights into other models in statistical mechanics. In a different setting, "chase-escape with death" mimics the behaviour of predators chasing prey on space, or the spread of a rumor throughout a social network.

Preface

Part I is the introduction for this thesis. Chapter 1 is an overview, while Chapters 2, 3 and 4 are surveys on background material.

Part II presents original research on uniform spanning trees. Chapter 5 and Chapter 6 are based on the preprints *"Scaling limits of the three-dimensional uniform spanning tree and associated random walk"* [11] and *"The number of spanning clusters of the uniform spanning tree in three dimensions"* [10], respectively. Our work in [10] will appear in the proceedings of "The 12th Mathematical Society of Japan, Seasonal Institute (MSJ-SI) Stochastic Analysis, Random Fields and Integrable Probability", while [11] is under review for publication. The research leading to these was an equal collaboration between Omer Angel, David Croydon, Daisuke Shiraishi, and myself. The writings of [11] and [10] were done in equal parts between Omer Angel, David Croydon, Daisuke Shiraishi, and myself.

Part III is original work on competitive growth processes. The research and writing was conducted in equal collaboration with Erin Beckman, Keisha Cook, Nicole Eikmeier and Matthew Junge. Chapter 7 is based on *"Chase-escape with death on trees"* [25] and has been submitted for publication.

Table of Contents

List of Tables

List of Figures

List of Symbols

$\sim$ Asymptotically equivalent. 12

τ_A First time a simple random walk hits A. 12

τ_A^+ First positive time a simple random walk hits A. 12

ξ_A First exit time of simple random walk from A. 13

$\mathbb{D}$ Unit ball (disk) on the complex plane. 26

d_H Hausdorff metric. 39, 95

σ Local state space [Chapter 5]. 49

$\lambda_c(G)$ Critical parameter for chase-escape. 73

$\mathcal{U}$ Uniform spanning tree on $\mathbb{Z}^3$. 78

$\mathbf{P}$ Law of the uniform spanning tree $\mathcal{U}$. 78

$d_{\mathcal{U}}$ Intrinsic metric on the graph $\mathcal{U}$. 78

$\mu_{\mathcal{U}}$ Counting measure on $\mathcal{U}$. 78

$\phi_{\mathcal{U}}$ Continuous embedding of $\mathcal{U}$ into $\mathbb{R}^3$. 78

$\rho_{\mathcal{U}}$ Root of $\mathcal{U}$. It is equal to 0. 78

β Growth exponent for the three-dimensional loop-erased random walk. 80

$\mathbf{P}_\delta$ Law of $\left(\mathcal{U}, \delta^\beta d_{\mathcal{U}}, \delta^3 \mu_{\mathcal{U}}, \delta\phi_{\mathcal{U}}, \rho_{\mathcal{U}}\right)$. 80

$B_{\mathcal{T}}(x,r)$ Ball in the metric space $\mathcal{T}$ of radius r around x. 82

$\mathcal{T}$ Limit metric space of the scaled uniform spanning tree $\mathcal{U}$. 82

d_f Fractal dimension of $\mathcal{U}$. 82

$\hat{\mathbf{P}}$ Law of the limit space $(\mathcal{T}, d_{\mathcal{T}}, \mu_{\mathcal{T}}, \phi_{\mathcal{T}}, \rho_{\mathcal{T}})$. 83

$X^{\mathcal{U}}$ Simple random walk on $\mathcal{U}$. 85

$P_x^{\mathcal{U}}$ Quenched law of the simple random walk on $\mathcal{U}$ started at x. 85

$\mathbb{P}^{\mathcal{U}}$ Annealed law of the simple random walk on $\mathcal{U}$. 85

$R_{\mathcal{U}}$ Effective resistance on $\mathcal{U}$. 85

$\mathbb{T}$ Collection of measure, rooted, spatial trees. 91

$B_\delta(x, r)$ Discrete Euclidean ball on $\delta\mathbb{Z}^3$ of radius r around x. 97, 202

$D_\delta(x, r)$ Discrete cube on $\delta\mathbb{Z}^3$ of side-length $2r$ centred at x. 97

$\partial_i A$ Inner boundary of $A \subset \mathbb{Z}^3$. 97

d_γ^S Schramm metric on a parameterized curve γ. 101

d_γ Intrinsic metric on a parameterized curve γ. 101

$\bar{\gamma}$ Loop-erased random walk endowed with its β-parameterization. 103

γ_∞^x Infinite loop-erased random walk on $\mathbb{Z}^3$ starting at x. 103

$\mathcal{P}$ Dyadic polyhedron. 109

$\mathscr{T}$ Parameterized tree. 158

$\mathscr{T}^K$ Space of parameterized trees with K leaves. 159

$\Gamma^e(\mathscr{T})$ Space of parameterized trees with K leaves. 164

$\mathcal{U}_\delta$ Uniform spanning tree on $\delta\mathbb{Z}^3$. 198

$\mathbb{B}$ Hypercube $[0, 1]^d$ 199

$\mathbb{T}_d$ d-ary tree 213

$\mathcal{R}$ Set of sites that are ever coloured red. 213

$\mathcal{B}$ Set of sites that are ever coloured blue. 213

$C_k^{\lambda,\rho}$ Weighted Catalan number 216

$g(z)$ Generating function of weighted Catalan numbers [Chapter 7]. 216, 224

$f(z)$ Continued fraction equal to $g(z)$ [Chapter 7]. 224

M Radius of convergence of g centred at the origin [Chapter 7]. 224

Acknowledgements

My deepest gratitude goes to my supervisors Martin Barlow and Omer Angel. I thank Martin for his support and guidance since my arrival at UBC as a master's student, and Omer for the countless hours talking about mathematics in coffee shops. I owe you my understanding and appreciation of concepts, heuristics, and methods; I am excited to continue thinking on these topics as I move forward in my career.

I thank Mathav Murugan, Lior Silberman, Nick Harvey, and Russell Lyons for reading this thesis, for their precise comments, and their insightful questions during the defence. It was an honour to have you in my examining committee.

I thank Daisuke Shiraishi, David Croydon, Erin Beckman, Keisha Cook, Matthew Junge, Nicole Eikmeier, and Tom Hutchcroft for the pleasure of working together. I am especially indebted to Tom for connecting Omer and me with Daisuke and David.

It was a privilege to have been part of the probability group at UBC. A special mention goes to my friends and colleagues Thomas Hughes and Guillermo Martinez Dibene. I think of them as my academic brothers. I would also like to acknowledge Ed Perkins for teaching several of my graduate courses in probability theory. These courses were inspiring, and my lecture notes from those days are a treasure.

I am fortunate to have found good friends during my time in Vancouver. I thank Bernardo, Delphin, Donají, Felipe, Hugo, Javier, Juan Camilo, Manuel, Marybel, Myrto, Nishant, Raimundo, Tara, and Thomas Budzinski. I was blessed to count on the fellowship and support of Arista, Christina, Jessie, Kaori, Michael and Amy Weidman, all the Navigators, and my favourite Chinese-Colombian-Venezuelan friends: Karen, Kally, and Cynthia Wong, and Gaby and Carlos García. I am happy to have kept close friends in Guadalajara: Jeanette, Kike, Manuel, Natalia, Poncho, Sarai, and Tonalli. I am also glad that Carmen and Malors have always been only a text message away.

I am deeply grateful for the friendship of Daniel and Karina. I appreciate their constant encouragement, their help at all times, and their wise and rational advice–especially when I have been stubborn. I cherish the good times that we have had together, and I look forward to more of them.

I am thankful for my parents Alejandro and Alma, and my sister Lizette. Your unconditional love has been an anchor in my life. I also thank all my family in Mexico City for their love and support.

This thesis was finished during a strange period of isolation, and I had a unique opportunity to walk around a quiet city. I treasure the view of the cherry blossoms, the turn at 8th Avenue, the mountains overlooking Jericho Beach, and of the gardens around Kitsilano and West Point Grey. These places guard sweet memories of friendship and peacefulness. I thank God for the gift that these years in Vancouver have been to me.

Finally, I am grateful for the funding from the Mexican National Council for Science and Technology (Consejo Nacional de Ciencia y Tecnología, CONACYT) in support of my graduate studies.

Dedication

Con amor a mi familia.

Part I

Introduction

Chapter 1

Discrete and Continuous Probability Models

A modern scientific paradigm is that natural phenomena arise from the collective behaviour of random microscopic interactions. This principle gained widespread acceptance with the introduction of Brownian motion to physics. Brownian motion was first described by and named after, the botanist Robert Brown in 1827. Brown observed the irregular movement of particles of pollen immersed in water. In 1905, Albert Einstein explained the physical mechanism for this motion as the result of random interactions at the molecular level [65]. This explanation gave evidence for the discrete nature of matter. Jean Perrin verified Einstein's predictions experimentally and hence confirmed molecular-kinetic theory in 1909 [145]. With experimental evidence firmly established and the pioneering theoretical work of James Clerk Maxwell, Ludwig Boltzmann, and J. Willard Gibbs, statistical mechanics gained a central place within modern physics. Since then, physics has thrived with successful applications of statistical mechanics. Its influence has spread to all the sciences. Statistical physics has driven ground-breaking developments in chemistry [70, 71], mathematical biology [12, 47, 142, 163] and theoretical computer science [137, 158], to cite some examples. The success of such applications, including further advancements in physics, requires a precise mathematical understanding of the relation between discrete models and their continuous counterparts.

In the last fifty years, mathematicians have reformulated problems from classical statistical physics (without quantum mechanics) within the framework of probability theory. Roland Dobrushin and Frank Spitzer started the study of interacting particle systems (see [75, 139] for surveys of their respective contributions), while Simon Broadbent, John Hammersley, and Dominic Welsh introduced percolation and first-passage percolation, respectively [37, 81]. These models have a simple description in terms of particle interactions. Even though the model is simple at a local (microscopic) level, it gains complexity when we consider a large number of particles. The latter is the most relevant case since it corresponds to the phenomena at a *macroscopic scale*. Research in the interface of probability and mathematical physics has flourished, and it has involved insights and methods from analysis and combinatorics. Nevertheless, some central questions in the area result in challenging mathematical problems. The work in this thesis is part of this continuing endeavour.

In this thesis, we study two discrete models: the *uniform spanning tree* and a competitive growth process called *chase-escape with death*. Our objective is to understand their *large scale behaviour*. For the uniform spanning forest, our main result is the existence of its *scaling limits*. The three-dimensional case is particularly interesting since it exhibits non-Gaussian behaviour. In the case of chase-escape with death, we study its *phase transitions* as we vary the model parameters. Our results touch on two main topics in statistical physics, namely scaling limits and phase transitions. In the next two sections, we present the probabilistic approach to these concepts.

1.1 Scaling limits as weak convergence of probability measures

A scaling limit is a formal connection between discrete and continuous probability models. In line with the interpretation from statistical physics, a phenomenon may be described either by a discrete or a continuous model.

The discrete model represents a microscopic scale and is usually defined over a graph, whereas the continuous model reflects a macroscopic scale and it is defined on $\mathbb{R}^d$. Some properties are understood more easily in the discrete setting, where combinatorial tools are at hand, but the continuous model usually presents symmetries absent in the discrete space. Examples of these symmetries are scale and rotational invariance or conformal invariance in the two-dimensional case. (See Propositions 2.2.10, 2.2.11, and 2.4.3.) We remark that the physics community was the first to observe symmetries on scaling limits, e.g. conformal invariance was predicted by Belavin, Polyakov, and Zamolodchikov [26].

Let us describe the general framework of a scaling limit. Consider a discrete model with a parameter describing its size (for example, the number of vertices of the underlying graph). We obtain a tractable problem by choosing a meaningful object associated with the discrete model. We say that the scaling limit exists when, after appropriate normalization, the chosen object converges as we increase the size parameter.

The archetypical example is the convergence of simple random walk on $\mathbb{Z}$ to Brownian motion. In this case, we scale the simple random walk by defining the processes on $\delta\mathbb{Z} := \{\delta v \colon v \in \mathbb{Z}\}$. Note that the distance between nearest-neighbour vertices on $\delta\mathbb{Z}$ decreases as $\delta \to 0$; the geometric effect is a zoom-out of the space. The corresponding size parameter is the number of vertices on $[0, 1]$, while the meaningful object for the scaling limit is the curve defined by interpolation of the random walk path. We have the convergence of this curve with respect to the space of continuous curves $\mathcal{C}[0, 1]$ endowed the supremum norm, and we thus say that the scaling limit exists. Chapter 2 expands our discussion on the simple random walk and Brownian motion.

Now, let us we specify the type of convergence of these random objects. Recall that we choose a representative object for studying the scaling limit of a given model. With this choice, we determine a Polish space E where these objects are defined. We thus get a probability measure μ_n, valued on E, associated with each size parameter n. The precise meaning of the existence of the scaling limit is that μ_n **converges weakly** to μ as $n \to \infty$, i.e.

$$\lim_{n\to\infty} \int f \, d\mu_n = \int f d\mu \qquad \text{for all } f \in C_c(E),$$

where $\mathcal{C}_c(E)$ is the space of continuous functions on E with compact support.

The scaling limit operation is applicable to a large variety of discrete models. In some situations, the limit object is deterministic, but it may also be random. The latter corresponds to phenomena exhibiting fluctuations at every scale. This behaviour is typical of *critical phenomena*, which we present in the following section.

1.2 Phase transitions

The models studied in statistical physics depend on a set of parameters. Among others, these could be the dimension of the underlying space, temperature, pressure, or rate of change. A **phase** is determined by a set of parameters with common *qualitative* properties. We observe a **phase transition** when we move between different phases as we modify the model parameters. The phases of water provide a well-known example. As we vary temperature or pressure, water abruptly transitions between solid, liquid, and gaseous states.

Branching processes provide a simple example of a phase transition. They are also a fundamental piece of our analysis in Chapter 7. Here we follow [157]; this reference provides more illustrations of critical phenomena for the interested reader. Branching processes are a model for population growth, where we represent individuals with identical particles ordered in a genealogical tree. For our purposes in this section, we restrict to branching processes with binomial offspring distribution. We remark that this a fundamental model in probability theory and its definition is more general (see [85, Chapter 3], [60, Section 2.1]).

We start with one particle occupying the root of a d-ary tree, while the rest of the vertices on the tree are vacant. Any particle reproduces only once in its lifetime. The offspring of a particle occupies some of the children nodes on the d-ary tree, leaving the rest of the children nodes vacant. Hence, the number of descendants of each individual follows a binomial distribution with parameters d (maximum number of children) and p (reproduction probability). The number of individuals (or occupied vertices) at generation t is a *Galton-Watson process* with offspring distribution $\mathrm{Bin}(d, p)$. Let $Z_d(p)$ be the total number of individuals.

One fundamental question for this model is on the size of the branching process: is it finite, or does the process generate an infinite number of generations? Let

$$\theta_d(p) = P\left(Z_d(p) = \infty\right)$$

be the *survival probability*, and let

$$\xi_d(p) = E\left(Z_d(p)\right)$$

be the *average family size*. A classical theorem states that the survival probability has a phase transition (see [85, 157]):

$$\theta_d(p) = \begin{cases} 0 & \text{if } p \leq \frac{1}{d}, \\ s & \text{if } p > \frac{1}{d}, \end{cases}$$

where $s > 0$ is the non-trivial solution to $s = ((1-p)+ps)^b$. The value $p_c := \frac{1}{d}$ is known as the **critical parameter**, since the model changes between phases at that point. Accordingly, a branching process is **subcritical** if $p < \frac{1}{d}$, **critical** if $p = p_c$, and **supercritical** if $p > d$. A simple calculation shows that the average family size also exhibits a phase transition around the critical parameter:

$$\xi_d(p) = \begin{cases} \frac{1}{1-dp} & \text{if } p < \frac{1}{d}, \\ \infty & \text{if } p \geq \frac{1}{d}. \end{cases}$$

We remark that at the critical parameter, the expected family size is infinite, even when survival probability at criticality is 0.

A crucial observation is that universal exponents govern the asymptotic behaviour around the critical point. For each $d \geq 2$, there exists constants

$C_1(d) > 0$, $C_2(d) > 0$ depending on d such that

$$\begin{aligned}\theta_d(p) &\sim C_1(d)(p - p_c)^{\beta}, \qquad p \to p_c^+,\\ \xi_c(p)(p) &\sim C_2(d)(p_c - p)^{-\gamma}, \qquad p \to p_c^-.\end{aligned}$$

In the asymptotic formulas above, β and γ are known as the *critical exponents*. They take the values $\beta = 1$ and $\gamma = 1$ for all d-ary trees. The independence of the critical exponents from the parameter d is an instance of *universality*.

The different phases of a discrete model explain the qualitative properties of the modelled phenomenon. However, quantitative conclusions depend on the details of the model, for example, the dimension of the underlying graph. In the case of the branching process, a qualitative property is the positivity of the survival probability, but the values of $\theta_d(p)$ and $\xi_d(p)$ depend on the parameters d and p. Note that the parameter d is a significant restriction on the model since it establishes a maximum number of offspring. Nevertheless, this restriction is irrelevant around the critical parameter. The **principle of universality** suggests that, at a critical point, the mathematical model approximates the physical reality. Therefore **critical exponents** determine the behaviour of physical phenomena at criticality. This principle justifies the value of understanding the simple discrete model. Moreover, several systems converge to the same behaviour as they approach criticality; we obtain a division of these systems into **universality classes**.

A common hypothesis in modelling is a convergence to the *Gaussian universality class*. When a model has enough (stochastic) independence among its different components, the central limit theorem applies, and its statistics converge to Gaussian random variables. However, several models have strong intrinsic dependencies, and their limit behaviour is non-Gaussian. Understanding non-Gaussian limits is a challenge for modern probability theory.

1.3 Structure of this thesis

The rest of the chapters in Part I introduce background material for this thesis. Chapter 2 and Chapter 3 are concise surveys on random walks and uniform spanning forests, respectively. These two chapters focus on essential definitions for Part II. Chapter 4 is an introduction to competitive growth processes, which are the main topic in Part III.

Part II and Part III report on original work on uniform spanning trees and competitive growth processes, respectively.

Chapter 8 presents the conclusions. In the concluding chapter, we summarize the contributions included in this thesis and future research directions.

Chapter 2

Random Walks

In this chapter, we define the simple random walk and the loop-erased random walk on $\mathbb{Z}^d$ and present their relevant properties for this work.

2.1 Notation

We begin with some notation that we will use through this thesis.

Following standard set notation $\mathbb{N}, \mathbb{Z}, \mathbb{R}$, and $\mathbb{C}$ represent the natural, integer, real and complex numbers, respectively. $\mathbb{R}^d$ is the d-dimensional real space, $\mathbb{R}_+ = \{x \in \mathbb{R} : \ z \geq 0\}$, $\mathbb{Z}_+ := \{z \in \mathbb{Z} : \ z \geq 0\}$, and $\mathbb{N}_0 := \mathbb{Z}_+$. The *indicator function* $\mathbf{1}\{x \in P\} : \mathbb{R}^d \to \{0, 1\}$ is defined as

$$\mathbf{1}\{B\}(x) := \begin{cases} 1, & \text{if } x \text{ satisfies property } P, \\ 0, & \text{otherwise.} \end{cases}$$

2.1.1 Subsets

For $x \in \mathbb{Z}^d$ and $z \in \mathbb{R}^d$, the **discrete ℓ^2 Euclidean ball** and the **Euclidean ℓ^2 ball** are the sets

$$B(x, r) := \left\{ y \in \mathbb{Z}^d : \ |x - y| < r \right\}, \qquad B_E(x, r) := \left\{ y \in \mathbb{R}^d : \ |x - y| < r \right\},$$

respectively. We use the abbreviations $B(R) = B(0, R)$ and $B_E(0, R) = B_E(R)$. The **discrete cube** (or ℓ^∞ ball of radius r) with side-length $2r$ centred at x is defined to be the set

$$D(x, r) := \left\{ y \in \mathbb{Z}^d : \|x - y\|_\infty < r \right\}.$$

Similarly to the definitions above, but with ℓ^∞ balls, $D_E(x, r)$ denotes the Euclidean cube. We further write $D(R) = D(0, R)$. The Euclidean distance between a point x and a set A is given by

$$\mathrm{dist}(x, A) := \inf \{|x - y| : y \in A\}.$$

The Euclidean *diameter* and the *radius* of A are

$$\mathrm{diam}\, A := \sup_{x,y \in A} |x - y|, \qquad \mathrm{rad}\, A := \min\{n \in \mathbb{N} : A \subset B(n)\}.$$

If $A \subset \mathbb{Z}^d$, we denote by ∂A the *discrete boundary* of A. It is defined as the set

$$\{x \notin A : \text{there exists } y \in A \text{ such that } x \text{ and } y \text{ are nearest-neighbours }\}.$$

2.1.2 Paths and curves

A **path** in $\mathbb{Z}^d$ is a finite or infinite sequence of vertices $[v_0, v_1, \ldots]$ such that v_{i-1} and v_i are nearest neighbours, i.e. $|v_{i-1} - v_i| = 1$, for all $i \in \{1, 2, \ldots\}$. The *length* of a finite path $\gamma = [v_0, v_1, ..., v_m]$ will be denoted $\mathrm{len}(\gamma)$ and is defined to be the number of steps taken by the path, that is $\mathrm{len}(\gamma) = m$.

A **(parameterized) curve** is a continuous function $\gamma : [0, T] \to \mathbb{R}^d$. For a curve $\gamma : [0, T] \to \mathbb{R}^3$, we say that $T < \infty$ is its *duration*, and will sometimes use the notation $T(\gamma) := T$. The curve γ is *simple* if it is an injective function. When the specific parameterization of a curve γ is not important, then we might consider only its trace, which is the closed subset of $\mathbb{R}^3$ given by $\mathrm{tr}\, \gamma = \{\gamma(t) : t \in [0, T]\}$. To simplify notation, we sometimes write γ, instead of $\mathrm{tr}\, \gamma$, where the meaning should be clear.

The space of parameterized curves of finite duration, $\mathcal{C}_f$, will be endowed with a metric ψ, as defined by

$$\psi(\gamma_1, \gamma_2) := |T_1 - T_2| + \max_{0\le s\le 1} |\gamma_1(sT_1) - \gamma_2(sT_2)|\,,$$

where $\gamma_i : [0, T_i] \to \mathbb{R}^3$, $i = 1, 2$ are elements of $\mathcal{C}_f$. Alternatively, consider the metric

$$\rho_{\mathcal{C}}(\gamma_1, \gamma_2) := \inf \sup_{t\in[0,1]} |\gamma_1 \circ \theta_1(t) - \gamma_2 \circ \theta_2(t)|,$$

where the infimum is over all the reparameterizations $\theta_1 : [0, 1] \to [0, T_1]$ and $\theta_2 : [0, 1] \to [0, T_2]$. In the literature, $\rho_{\mathcal{C}}$ is known as the metric of the *space of unparameterized paths.*

A continuous map $\gamma^\infty : [0, \infty) \to \mathbb{R}^d$ is a **transient curve** if $|\gamma^\infty(t)| \to \infty$ as $t \to \infty$. Let $\mathcal{C}$ be the set of transient curves, and endow $\mathcal{C}$ with the metric χ given by

$$\chi(\gamma_1^\infty, \gamma_2^\infty) = \sum_{k=1}^{\infty} 2^{-k} \left(1 \wedge \max_{t\le k} |\gamma_1^\infty(t) - \gamma_2^\infty(t)|\right).$$

2.1.3 Constants and asymptotic notation

We denote constants with the letters C, C_n c and c_n, with $n \in \mathbb{N}$. The values of these constants change from line to line, and we indicate their dependencies.

Let f, g, h be real valued functions with $f, g, h \ge 0$. We write $f(x) = O\left(g(x)\right)$ to indicate that there exists a constant $C > 0$ such that

$$f(x) \le Cg(x), \qquad \text{for all } x.$$

Similarly, we write $f(x) = h(x) + O(g(x))$ to indicate that

$$|f(x) - h(x)| \le Cg(x), \qquad \text{for all } x.$$

If $f(x) \preceq g(x)$, it means that there exists C such that

$$f(x) \leq Cg(x), \qquad \text{for all } x.$$

Similarly, we write $f(x) \asymp g(x)$ if there exist $c_1, c_2 > 0$ such that

$$c_1 g(x) \leq f(x) \leq c_2 g(x), \qquad \text{for all } x.$$

Finally, if f and g are positive functions, we write $f \sim g$ if

$$\lim_{x \to \infty} \frac{f(x)}{g(x)} = 1.$$

2.2 Simple random walk

The simple random walk is a random path on a given graph, where each step is chosen uniformly at random. For this work, we delimit our discussion to simple random walks on $\mathbb{Z}^d$. Consider the set of directions on $\mathbb{Z}^d$, $\mathcal{E} = \{\pm e_1, \ldots, \pm e_n\}$, where $e_k(j) = \mathbf{1}\{k = j\}$. Let $(\eta_j)_{j \in \mathbb{N}}$ be independent random variables, each one with uniform distribution over $\mathcal{E}$. For $x \in \mathbb{Z}^d$, the **simple random walk** $S = (S_n)_{n \in \mathbb{N}}$ started at x is

$$S_0 := x, \qquad S_n := S_0 + \sum_{j=1}^{n} \eta_j.$$

We denote by P^x the probability measure of S. The distribution of η_j is called the *step distribution.* Thorough this work, we may also use de notation $S(n) = S_n$.

The simple random walk is a Markovian process. Our interest is in the geometry of the random walk, and hence the most relevant stopping times are those related to exit and hitting times. We define the **hitting** and **positive hitting times** of S by

$$\tau_A := \inf\{n \geq 0 \;:\; S_n \in A\}, \quad \text{and} \qquad \tau_A^+ := \inf\{n > 0 \;:\; S_n \in A\}. \quad (2.1)$$

We write τ_m and τ_m^+ for the hitting times of the ball $B(m)$. A related stopping time is the escape from a set. We write

$$\xi_A := \inf\{n \geq 0 \colon S_n \in A^c\} \tag{2.2}$$

for the **escape time** from A. If $A = B(m)$, we write

$$\xi_m := \inf\{n \geq 0 \colon S_n \in B(m)^c\}. \tag{2.3}$$

2.2.1 Recurrence and transience

A random walk S is **recurrent** if

$$P(S_n = 0 \text{ i.o }) = 1,$$

otherwise, we say that S is **transient**. The recurrence of the simple random walk on $\mathbb{Z}^d$ depends on the dimension d.

Theorem 2.2.1. *The simple random walk on $\mathbb{Z}^d$ is recurrent in $d = 1, 2$ and transient in $d \geq 3$.*

A simple proof of Theorem 2.2.1 is in [61, Subsection 5.4]. The basis for the proof is the following characterization of recurrence in terms of *hitting probabilities.*

Proposition 2.2.2 ([61, Theorem 5.4.3])**.** *For a simple random walk S on $\mathbb{Z}^d$, the following are equivalent:*

(i) *S is recurrent,*

(ii) *$P^0(\tau_{\{0\}}^+ < \infty) = 1$, and*

(iii) $\displaystyle\sum_{m=0}^{\infty} P^0(S_n = 0) = \infty.$

The next proposition is a quantified version of Proposition 2.2.2 (ii) in the transient case.

Proposition 2.2.3 ([117, Proposition 6.4.2])**.** *Let $d \geq 3$. For $x \in \mathbb{Z}^d \setminus B(m)$*

$$P^x\left(\tau_m^+ < \infty\right) = \left(\frac{m}{|x|}\right)^{d-2} [1 + O(m^{-1})].$$

Throughout this work, we often make a distinction between dimensions $d = 2$ and $d \geq 3$ of the lattice $\mathbb{Z}^d$. This distinction is due to the difference between recurrent and transient behaviour.

2.2.2 Harmonic measure and hitting probabilities

Estimates on hitting probabilities of a random walk lie at the core of this work. We review hitting probabilities, harmonic measure, and capacity to provide some background. We follow [113, 117], where more details are available.

Let $A \subset \mathbb{Z}^d$ be a finite set. The **harmonic measure** of A is defined as the limit

$$\mathrm{hm}_A(y) := \lim_{|x|\to\infty} P^x\left(S_{\tau_A^+} = y : \tau^+ < \infty\right).$$

In the two-dimensional case, the simple random walk is recurrent and the harmonic measure is simply defined by

$$\mathrm{hm}_A(y) := \lim_{|x|\to\infty} P^x\left(S_{\tau_A^+} = y\right).$$

We refer to [113, Theorem 2.1.3] for a proof of the existence of this limit.

The following result gives bounds on the harmonic measure of straight segments.

Theorem 2.2.4 ([113, Section 2.4])**.** *Let $L \subset \mathbb{Z}^d$ be the line segment on the x-axis from $(0, \ldots, 0)$ to $(n, 0, \ldots, 0)$. Then*

$$\mathrm{hm}_L(0) \asymp \begin{cases} cn^{-1/2}, & d = 2, \\ c(\log n)^{1/2} n^{-1}, & d = 3. \end{cases}$$

We define below the capacity of a set and relate it to the hitting probability of a random walk and the harmonic measure.

Capacity in the transient case

On $\mathbb{Z}^d$, and for $d \geq 3$, the **capacity** of a finite set A is defined as

$$\operatorname{cap}(A) := \lim_{m \to \infty} \sum_{x \in A} \mathbb{P}^x \left(\tau_A^+ > \xi_m \right) = \sum \mathbb{P}^x(\tau_A^+ = \infty).$$

It follows that we can write the harmonic measure as

$$\operatorname{hm}_A(x) = \frac{\mathbb{P}^x \left(\tau_A^+ = \infty \right)}{\operatorname{cap}(A)}.$$

The capacity of a set A indicates how much "*hittable*" is A by a random walk starting at a large distance. We formally state this relation in the proposition below.

Proposition 2.2.5 ([117, Proposition 6.5.1])**.** *Assume that $A \subset B(n)$ and $\|x\| \geq 2n$, then*

$$\mathbb{P}^x \left(\tau_A^+ < \infty \right) = C_d \|x\|^{2-d} \operatorname{cap}(A) \left[1 + O\left(\frac{n}{\|x\|} \right) \right],$$

where C_d is a constant depending on the dimension.

An example to keep in mind is the capacity of the closed ball of radius n. This is

$$\operatorname{cap}(\bar{B}(n)) = a_d^{-1} n^{d-2} + O\left(n^{d-3} \right);$$

where the constant a_d takes the value

$$a_d = \frac{d}{2} \Gamma\left(\frac{d}{2} - 1 \right) \pi^{-d/2} = \frac{2}{(d-2)\omega_d}.$$

In the expression above, Γ is the Gamma function and ω_d is the volume of the unit ball in $\mathbb{R}^d$ [113, (2.16) and Theorem 1.5.4]. In comparison, for an arbitrary connected set, we limit our calculations to bounds over the harmonic measure, e.g. Theorem 2.2.7 gives an upper bound.

The recurrence of the two-dimensional random walk entails a different definition of capacity for $\mathbb{Z}^2$. We will not use this definition, and instead, refer the reader to [113]. In exchange, we state below a general theorem for the hitting probability of a connected subset of $\mathbb{C}$.

Beurling estimate

The Beurling projection theorem is a classical result for the hitting probabilities of a two-dimensional Brownian motion. Consider a Brownian motion on $\mathbb{C}$, started at 0 and stopped when it hits the unit circle; and let $\mathcal{A}$ be the collection of connected subsets of $\mathbb{C}$ containing the origin and intersecting the unit circle. The Beurling projection theorem states that, among all subsets in $\mathcal{A}$, this Brownian motion is most likely to avoid the straight line $[0, 1]$ (see [27, 166]). This theorem is a consequence of Beurling's theorem [32], formulated originally in potential theory.

We have a discrete analogue for simple random walks. In this case, we consider $A \subset \mathbb{Z}^d$ *path connected*, meaning that there is a path between any two points in A. In terms of hitting probabilities, the statement is the following.

Theorem 2.2.6 (Beurling estimate [115, Theorem 6.8.1])**.** *Let $A \subseteq \mathbb{Z}^2$ be an infinite and path-connected set containing the origin. Then, for a simple random walk starting at the origin*

$$P\left(\xi_n < \tau_A^+\right) \leq \frac{c}{n^{1/2}}.$$

In terms of the harmonic measure, the Beurling estimate states that the harmonic measure of a line (hm_L in Theorem 2.2.4) is an upper bound of the harmonic measure of any path-connected set. We state this version of the Beurling estimate as follows.

Theorem 2.2.7 ([113, Theorem 2.5.2])**.** *Let $A \subset \mathbb{Z}^d$ be a path-connected set of radius n containing 0. Then*

$$\mathrm{hm}_A(0) \leq \begin{cases} cn^{-1/2}, & d = 2, \\ c(\log n)^{1/2} n^{-1}, & d = 3, \\ cn^{-1}, & d \geq 4. \end{cases}$$

Kesten proved the two-dimensional case in [97], while the argument in [113] appeared first in [112]. The proof in [112] uses the following lower bound on the capacity of connected sets with radius n.

Proposition 2.2.8 ([113, Lemma 2.5.4])**.** *Let $A \subset \mathbb{Z}^d$ be a connected set of radius n containing* 0*. Then*

$$\mathrm{cap}(A) \geq \begin{cases} cn\,(\log n)^{-1}, & d = 3, \\ cn, & d \geq 4. \end{cases}$$

In Chapter 5, we will require a better estimate in the three-dimensional case. In that case, the set to hit is the trace of a loop-erased random walk and we use properties specific to the loop-erased random walk. However, we still refer to such results as a *Beurling-type estimate* (see Subsection 2.3.3).

2.2.3 Scaling limit of the simple random walk

For simplicity, consider a simple random walk $S = (S(n))$ in the line. If we interpolate between $S(0), S(1), \ldots, S(n)$, we obtain a continuous function in $\mathbb{R}$. Denote by $S(t)$, $t \geq 0$ the function defined by this interpolation. In this sense, the random walk S is a model for a discrete random function.

The continuous analogue for a random function is *Brownian motion*. In this subsection, we introduce Brownian motion and its relation to the simple random walk through the scaling limit.

Brownian motion

Let $(\Omega, \mathcal{F}, P)$ be a probability space. We say that a random variable X follows a normal distribution with mean 0 and variance σ^2, denoted by $\mathcal{N}(0, \sigma^2)$ if for any Borel set $A \subset \mathcal{B}(\mathbb{R})$

$$P(X \in A) = P\left(\omega \in \Omega \colon X(\omega) \in A\right) = \frac{1}{\sqrt{2\pi\sigma}} \int_A \exp\left(-\frac{x^2}{2\sigma^2}\right).$$

Definition 2.2.9. A Brownian motion in $\mathbb{R}$ (or linear Brownian motion) is a collection of random variables $W := (W(t,\omega) \ : \ t \geq 0, \omega \in \Omega)$ satisfying the following properties:

(i) The distribution at time 0 is identically 0, i.e. $W(0,\omega) = 0$, for all $\omega \in \Omega$.

(ii) For any $0 \leq t < s$, the random variable $W_t - W_s := W(t,\cdot) - W(s,\cdot)$ follows a normal distribution $\mathcal{N}(0, t-s)$.

(iii) For any $0 \leq t_1 \leq t_2 \leq \ldots \leq t_n$, the increments $W_{t_n} - W_{t_{n-1}}, W_{t_{n-1}} - W_{t_{n-2}}, \ldots, W_{t_2} - W_{t_1}$ are independent random variables.

(iv) For P-almost all $\omega \in \Omega$, the function $t \mapsto W(t,\omega)$ is continuous.

The extension to higher dimensions is straightforward. Let $W^1, \ldots, W^d$ be d independent linear Brownian motions. The collection of random variables $(B(t)\colon t \geq 0)$ given by

$$B(t) = (W_t^1, \ldots, W_t^d)^T$$

is a **d-dimensional Brownian motion**. In the case $d = 2$, we call B a **planar Brownian motion**. An equivalent definition for the d-dimensional Brownian motion $(B(t) : t \geq 0, \omega \in \Omega)$ is the analogue of Definition 2.2.9, but we exchange condition (ii) by the requirement that $B(t) - B(s)$ follows a d-dimensional normal distribution with mean 0 and covariance matrix $(t-s)I_d$.

We assume above that Brownian motion starts at 0, but the initial point may be any $z \in \mathbb{R}^d$. In this case, we change condition (i) for $B(0) = z$ with probability one. We thus say that B starts at z and denote the corresponding probability measure by $P^z(\cdot)$ We interpret B as a random continuous function. That is, for each $\omega \in \Omega$ we get a continuous function $B = B(\cdot, \omega) : [0, \infty) \to \mathbb{R}^d$.

We will see below that Brownian motion is a limit object. Accordingly, Brownian motion satisfies translation, scale, and rotation invariance. We refer to [140] for the proof of Proposition 2.2.10 and Proposition 2.2.11.

Proposition 2.2.10. *Let $B = (B(t) : t \geq 0)$ be a d-dimensional Brownian motion. For $a > 0$, $z \in \mathbb{R}^d$ and an orthogonal linear transformation $L : \mathbb{R}^d \to \mathbb{R}^d$, the processes $\left(\widetilde{B}(t) : t \geq 0\right)$ and $\left(\bar{B}(t) : t \geq 0\right)$ given by*

$$\widetilde{B}(t) = \frac{1}{a} B(a^2 t) + z, \qquad \bar{B}(t) = L(B(t)),$$

are also Brownian motions, started at z and at 0, respectively.

In the two-dimensional case, the planar Brownian motion satisfies conformal invariance. We consider planar Brownian in the complex plane by setting

$$B(t) = W_1 + iW_2,$$

where W_1 and W_2 are two independent linear Brownian motions. For a domain $D \subset \mathbb{C}$, let

$$\xi_D^B = \inf\{t \geq 0 \,:\, B(t) \notin D\}$$

be the exist time of the Brownian motion from the domain D

Proposition 2.2.11. *Let $B = (B(t) : t \geq 0)$ be a planar Brownian motion. For a conformal map $\varphi : D \to \hat{D}$, the process $\left(\hat{B}(t) : t \geq 0\right)$ given by*

$$\hat{B}(t) = \varphi(B(t))$$

is a time-changed Brownian motion such that $\xi_{\hat{D}}^{\hat{B}} = \varphi(\xi_D^B)$.

Convergence of the simple random walk to Brownian motion

We begin with a discussion of the one-dimensional case. Consider the function in $[0,1]$ given by

$$S_n^*(t) = \frac{S(nt)}{\sqrt{n}}, \qquad \text{for all } t \in [0,1].$$

With this normalization, the central limit theorem implies that $S_n^*(1)$ converges in distribution to $\mathcal{N}(0,1)$. On the other hand, note that $W(1) =$

$\mathcal{N}(0,1)$ in distribution. In general, for each fixed time t,

$$S_n^*(t) \overset{D}{\Rightarrow} W(t), \qquad \text{as } n \to \infty,$$

and we see that S_n^* converges weakly to W pointwise. However, if we think of S^* as a continuous curve, pointwise convergence is unsatisfactory. It corresponds to the convergence of finite-dimensional distributions. Donsker's invariance principle extends this convergence to the space of continuous functions $\mathcal{C}[0,1]$, endowed with the supremum norm.

Donker's invariance principle holds in all dimensions. In the general case, we scale a d-dimensional simple random walk by

$$S_n^{*,d} = \frac{\sqrt{d}}{\sqrt{n}} S^d([nt]), \qquad t \in [0,1].$$

We state it below, and we refer to [66, Chapter 5, Theorem 1.2] for a proof.

Theorem 2.2.12 (Donsker's Invariance Principle)**.** *$S_n^{*,d}$ converges weakly to a standard Brownian motion on $\mathbb{R}^d$ in the space of continuous functions $\mathcal{C}_{\mathbb{R}^d}[0,1]$.*

We thus say that *Brownian motion is the scaling limit of the simple random walk.*

As we have stated above, we are mainly interested in the simple random walk on $\mathbb{Z}^d$. We remark that the convergence of the simple random walk to Brownian motion holds in a large class of graphs. For example, finite range, symmetric and irreducible random walks on $\mathbb{Z}^d$ converge to Brownian motion [117, Chapter 3]. The convergence also holds for centred step distributions with finite variance, [102, Theorem 21.43] gives a proof in the one-dimensional case. The fact that this scaling limit holds in such generality is an instance of the *universality* phenomenon, and we thus say that Brownian motion is a *universal object.*

2.3 Loop-erased random walks

The loop-erased random walk is a model for simple curves motivated by the self-avoiding walk (SAW) model [110]. In $\mathbb{Z}^d$, for $d \geq 2$, consider a curve $\gamma : \{0, 1, \ldots, n\} \to \mathbb{Z}^d$. We define the **loop-erasure** of γ as a curve created by deleting the loops of γ in chronological order. Let

$$s_0 := \sup\{j \colon \gamma(j) = \gamma(0)\},$$

and for $i > 0$,

$$s_i := \sup\{j \colon \gamma(j) = \gamma(s_{i-1} + 1)\}.$$

The length of the loop-erasure is $m = \inf\{i \colon s_i = n\}$. Then, the loop-erasure of γ is

$$\mathrm{LE}(\gamma) := [\gamma(s_0), \gamma(s_1), \ldots, \gamma(s_m)].$$

Let D be a subset of $\mathbb{Z}^d$. Consider a simple random walk S on $\mathbb{Z}^d$ starting at $x \in D$. The **loop-erased random walk** on D is defined as the loop-erasure of S up to its first exit from D

$$\gamma = \mathrm{LE}\, S[0, \xi_D], \tag{2.4}$$

where ξ_D is the escape time defined in (2.3).

2.3.1 The infinite loop-erased random walk

The infinite loop-erased random walk is the loop-erasure of a simple random walk without an stopping condition. The latter statement has an immediate interpretation when the simple random walk is transient, which is the case of $\mathbb{Z}^d$ with $d \geq 3$. In $\mathbb{Z}^2$, the simple random walk is recurrent, but we can define the loop-erased random walk as a weak limit. We discuss both cases below.

Transient case

Let $S = (S_n)_{n \geq 0}$ be a simple random walk on $\mathbb{Z}^d$. We assume that the dimension is $d \geq 3$. In this case, the simple random walk is transient, and the loop-erasure of S is well-defined, with probability one. Similarly to the finite case, we set

$$s_0 := \sup\{j \colon S(j) = S(0)\},$$

and for $i > 0$,

$$s_i := \sup\{j \colon \gamma(j) = \gamma(s_{i-1} + 1)\}.$$

We note that s_i is finite with probability one due to the transience of the simple random walk. Then, the **infinite loop-erased random walk** (ILERW) is the transient path

$$\mathrm{LE}(S) := [\gamma(s_0), \gamma(s_1), \ldots,].$$

Two-dimensional case

For each $\ell \geq 1$, let Ω_ℓ be the set of simple paths $\omega = [0, \omega_1, \ldots, \omega_k]$ from 0 to the boundary of B_ℓ i.e. $\omega_1, \ldots, \omega_{j-1} \in B_\ell$ and $\omega_j \in \partial B_\ell$. Let γ_m be a loop-erased random walk on B_m and $\gamma_m|^\ell$ be the restriction of γ_m up to its first exit from B_ℓ. We denote by

$$\nu_{m,\ell}(\omega) = P(\gamma_m|^\ell = \omega), \qquad \omega \in \Omega_\ell$$

the probability measure on Ω_ℓ induced by γ_m.

Proposition 2.3.1 (Lawler [113, Proposition 7.4.2])**.** *Let $\omega \in \Omega_\ell$. If γ_m is a loop-erased random walk on B_m, then*

$$\lim_{m \to \infty} P(\gamma_m|^\ell = \omega) = \lim_{m \to \infty} \nu_{m,\ell}(\omega) = \hat{\nu}_\ell(\omega)$$

exists.

The collection $\{\hat{\nu}_\ell\}_{\ell \geq 1}$ is consistent and defines a measure $\hat{\nu}$ on infinite paths. The *two-dimensional infinite loop-erased random walk* is the random infinite path with measure $\hat{\nu}$.

Restrictions of infinite loop-erased random walks

The LERW and ILERW are different objects. However, the definition of the ILERW suggests that their respective measures are comparable within a small ball.

Proposition 2.3.2 (Masson [136, Corollary 4.5])**.** *Let $\ell \geq 1$ and $n \geq 4$. Let K be a subset containing $B(n\ell)$ and such that, for the escape time defined in* (2.2)*, $P^0(\xi_K < \infty) = 1$. If γ^∞ is an infinite loop-erased random walk and γ^K is a loop-erased random walk on K and $\omega \in \Omega_\ell$ then*

$$P(\gamma^\infty[0, \xi_\ell^\infty] = \omega) = \begin{cases} \left[1 + O\left(\frac{1}{\log n}\right)\right] P(\gamma^K[0, \xi_\ell^K] = \omega) & d = 2, \\ \left[1 + O\left(n^{2-d}\right)\right] P(\gamma^K[0, \xi_\ell^K] = \omega) & d \geq 3, \end{cases}$$

where ξ_ℓ^∞ and ξ_ℓ^K are the escape times from the ball $B(\ell)$ of γ^∞ and γ^K, respectively.

2.3.2 Growth exponent

The **growth exponent** of the d-dimensional loop-erased random walk is the asymptotic number of steps necessary to reach Euclidean distance n. In a sense, it indicates the efficiency of the random path to reach a macroscopic distance.

It is convenient to compare the growth exponent of the LERW with two examples. A growth exponent equal to 1 indicates linear growth, and a straight line provides an example. The second example is a simple random walk. Its growth exponent is 2 since the loops increase the number of steps in the path. It is intuitively clear that the growth exponent depends on the dimension of the lattice.

Let $S = (S(t))$ be a simple random walk in $\mathbb{Z}^d$ started at the origin and let $\xi_n = \inf\{t \geq 0\colon \|S(t)\| > n\}$ be the first exit time from the Euclidean ball of radius n. Then

$$M_d(n) := |\operatorname{LE}(S[0, \xi_n])|,$$

is the number of steps it takes to the LERW to exit a ball of radius n in the d-dimensional space.

We define the **growth exponent of the loop-erased random walk** as

$$\beta_d := \lim_{n\to\infty} \frac{\log \mathbb{E}(M_d(n))}{\log n}, \tag{2.5}$$

provided that the limit exists. In this case, we write that $\mathbb{E}(M_d(n)) \approx n^{\beta_d}$

The following theorem summarizes results on the existence of the growth exponent for the LERW. Kenyon determined the planar case in [95]. Shiraishi established the existence in $d = 3$ [155, Theorem 1.4], and the upper and lower bounds come from work in [114].

Theorem 2.3.3. *The growth exponent β_d for the LERW on $\mathbb{Z}^d$ exists for all $d \geq 2$. The growth exponent takes the following values in each dimension:*

(a) $\beta_2 = \frac{5}{4}$,

(b) $\beta_3 \in (1, 5/3]$,

(c) $\beta_d = 2$, *for* $d \geq 4$.

Further work has obtained a more precise asymptotic behaviour for the planar LERW and in higher dimensions $d \geq 4$. We present below some of these results.

In the two-dimensional case, Lawler obtained the asymptotic probability that the path of a LERW contains the edge $[0, 1]$ while it crosses a square of length $2n$ [116, Theorem 1.1]. This estimate gives a precise asymptotic for the growth exponent. We refer to [18, Corollary 3.15] for details on the connection between the crossing probability and the growth exponent.

Theorem 2.3.4 (Lawler [116])**.** *There exist absolute constant $c_1, c_2 > 0$ such that*

$$c_1 n^{5/4} \leq \mathbb{E}(M_2(n)) \leq c_2 n^{5/4}. \tag{2.6}$$

The result extends to more general planar graphs. Proposition 2.3.5 shows that the growth exponent is a function of the dimension, and it does not depend on the particularities of $\mathbb{Z}^2$. This is an instance of the universality of the growth exponent. Let $\bar{S}$ be an irreducible bounded symmetric random walk, starting at the origin, on a two-dimensional lattice. As in $\mathbb{Z}^2$, we set $\bar{\xi}_n = \inf\{t \geq 0\colon \|S(t)\| > n\}$ and $\bar{M}(n) = |\operatorname{LE}(\bar{S}[0, \bar{\xi}_n])$.

Proposition 2.3.5 (Masson [136])**.** *The limit*

$$\lim_{n\to\infty} \frac{\log \mathbb{E}(\bar{M}(n))}{\log n} = \frac{5}{4}$$

and hence the growth exponent for the LERW on a two-dimensional lattice is $\frac{5}{4}$.

In the critical dimension $d = 4$, Lawler obtained the logarithmic corrections for a related exponent. The physics community has predicted these logarithmic corrections. Let $K_n := |\operatorname{LE}(S_n)|$, that is, the number of points kept in the loop-erasure of a simple random walk of n steps. It was proved on [114] that for loop-erased random walks in $d = 4$

$$\lim_{n\to\infty} \frac{\mathbb{E}(K_n)}{c_4 n (\log n)^{-1/2}} = 1.$$

In comparison, $\mathbb{E}(K_n) \sim c_d n$, for $d \geq 5$.

In higher dimensions, the behaviour is Gaussian, and precise asymptotics are available (see [114]). For $d \geq 5$, there exist $c_1, c_2 > 0$ such that

$$c_1 n^2 \leq \mathbb{E}(M_d(n)) \leq c_2 n^2.$$

2.3.3 Hitting probabilities

General estimates on the harmonic measure are enough for the study of the two-dimensional loop-erased random walk. In three-dimensions, we require a result specific for loop-erased random walks.

Theorem 2.3.6 (Sapozhnikov-Shiraishi [148, Theorem 3.1])**.** *Let γ be a loop-erased random walk on $B(n)$. There exist $\eta > 0$ and an absolute con-*

stant $C < \infty$ such that for all $\varepsilon > 0$ and $n \geq 1$,

$$P\left(\begin{array}{c} \text{For all } x \in B(n) \text{ with } \operatorname{dist}(x,\gamma) \leq \varepsilon^2 n \\ P^x\left(\xi^S_{B(x,\sqrt{\varepsilon}n)} < \left(\tau^S_\gamma\right)^+\right) \leq \varepsilon^\eta \end{array}\right) \geq 1 - C\varepsilon,$$

where S is an independent simple random walk on $\mathbb{Z}^3$ started at x.

Subsection 5.3.4 presents variants of Theorem 2.3.6 for infinite loop-erased random walks and loop-erased random walks stooped at a random boundary.

2.4 Scaling limits of loop-erased random walks

In this section, we discuss the scaling limits of loop-erased random walks. In contrast with the case of simple random walks, the limit process depends on the dimension of the space.

2.4.1 Two dimensions

We begin with the planar case $\mathbb{Z}^2$. The complex plane reveals a rich structure for two-dimensional processes. In this subsection, we work on $\mathbb{C}$ and denote the unit ball (or disk) by $\mathbb{D} = B(0,1)$.

We have an explicit description of the scaling limit of the two-dimensional loop-erased random walk. This is the Schramm-Loewner evolution with parameter 2 (SLE(2)). We introduce SLE in this specific case for comparison with the scaling limits in higher dimensions.

Radial SLE

The Schramm-Loewner evolution (SLE(κ)) is a one-parameter family of conformally invariant scaling limits of two-dimensional discrete models. The following results are known as we take the scaling limit.

- The loop-erased random walk converges to SLE(2) [122, 149].
- The interface of the planar critical Ising model converges to SLE(3) [46].

- The harmonic explorer converges to SLE(4) [150].

- SLE(6) corresponds to the scaling limit of critical percolation on the triangular lattice (proof outlined in [159, 160] and completed in [41, 42]).

- The Peano curve of the uniform spanning tree converges to SLE(8) [122].

SLE(8/3) is the conjectured limit of the self-avoiding random walk; close relation between Brownian motion and SLE(8/3) supports this conjecture [121].

SLE(κ) is defined over domains $D \subset \mathbb{C}$. We distinguish two points on D, where the process starts and finishes. *Radial* SLE corresponds to $\mathbb{D}$ with the process starting at a point in the boundary and finishing at the origin. On the other hand, *chordal* SLE refers to the process on the upper-half plane $\mathbb{H}$ starting at 0 and ending at ∞.

Let us describe radial SLE(2). We follow the construction in [122] and refer the reader to proofs in [30]. The proofs in [30] are for the chordal case, but they also apply to radial SLE after a conformal transformation.

We say that K is a $\mathbb{D}-$*hull* if K is a compact subset of $\bar{\mathbb{D}}$ and $\mathbb{D} \setminus K$ is a simply connected domain. There is a one-to-one correspondence between $\mathbb{D}$-hulls and conformal homeomorphisms

$$g_K : \mathbb{D} \setminus K \to \mathbb{D} \tag{2.7}$$

satisfying $g_K(0) = 0$ and $g'_K(0) \geq 0$. The Riemann mapping theorem and the Schwarz lemma provide this bijection (see [30, Corollary 1.4]). We will look at families of $\mathbb{D}$-hulls. We say that $(K_t : t \geq 0)$ is *increasing* if $K_t \subsetneq K_s$ for $t < s$. Moreover, a family $(K_t : t \geq 0)$ satisfies the *local growth property* if

$$\operatorname{diam}(K_{t,t+h}) \to 0 \qquad \text{as } h \downarrow 0 \text{ uniformly on compacts in } t,$$

where $K_{t,t+h} = g_{K_t}(K_{t+h} \setminus K_t)$. Simple continuous curves provide the most relevant example of a family of $\mathbb{D}$-hulls. If $\eta : [0, \infty) \to \bar{\mathbb{D}}$ is a continuous simple curve with $\eta(0) \in \partial\mathbb{D}$, $\lim_{t\to\infty} \eta(t) = 0$ and $\eta(0,\infty) \subset \mathbb{D}$, then $K_t = \eta[0,t]$ defines an increasing family of $\mathbb{D} - hulls$ with the local growth property.

We also have a correspondence between continuous functions

$$W : [0,\infty) \to \partial\mathbb{D}$$

and increasing families of $\mathbb{D}$-hulls satisfying the local growth property

$$(K_t :\ t \geq 0) \tag{2.8}$$

with $K_0 \in \partial\mathbb{D}$, $K_t \setminus K_0 \subset \mathbb{D}$ for $t \in (0,\infty)$, and the assumption that 0 is in the closure of $\cup_{t\geq 0} K_t$.

Given $(K_t :\ t \geq 0)$ satisfying (2.8), let $g_t := g_{K_t}$ be as in (2.7) for each $t \geq 0$. We further assume that the conformal maps are parameterized so that $g_t'(0) = \exp(t)$. Then, for all $t \geq 0$, there exist a unique real number in $\bar{K}_{t,t+h}$ for all $h > 0$. We have that

$$W(t) := \lim_{h\to 0} \bar{K}_{t,t+h} = W(t) \tag{2.9}$$

exists and $W : [0,\infty) \to \mathbb{R}$ defines a real-valued continuous function [30, Proposition 7.1].

Loewner's Slit Mapping Theorem [131] provides a crucial observation to reverse the construction.

Theorem 2.4.1 (Loewner [131])**.** *The conformal homeomorphism g_t satisfies the differential equation*

$$\partial_t g_t(z) = -g_t(z) \frac{g_t(z) + W(t)}{g_t(z) - W(t)}, \tag{2.10}$$

where W is the continuous map (2.9)*. Clearly $g_0(z) = z$, for all $z \in \mathbb{D}$.*

In light of (2.10), we call W the **driving function** of $(K_t :\ t \geq 0)$.

Now, we start with a continuous function $W : [0,\infty) \to \partial\mathbb{D}$. We define a conformal map g_t as the solution of the ODE (2.10) with initial value

$g_0(z) = z$ up to some time $\tau(z) \in (0, \infty]$. For $t \geq \tau(z)$, the solution to (2.10) does not exist. We then define the hull at time t as

$$K_t = \{z \in \bar{\mathbb{D}} : \tau(z) \leq t\},$$

and $D_t := \mathbb{D} \setminus K_t$, so the domain of g_t is D_t and maps onto $\mathbb{D}$. Then $(K_t : t \geq 0)$ is an increasing family of $\mathbb{D}$-hulls with the local growth-property and W is its driving function. [30, Proposition 8.2].

Schramm defined the Schramm-Loewner evolution (first known as the stochastic Loewner evolution) in his influential work on scaling limits of the loop-erased random walk [149].

Definition 2.4.2. Radial Schramm-Loewner evolution with parameter k (SLE(k)) is the process of random $\mathbb{D}$-hulls $(K_t, t \geq 0)$ with driving function

$$W(t) = \exp(iB(kt)),$$

where $B : [0, \infty) \to \mathbb{R}$ is a Brownian motion.

We define radial SLE similarly in any simply connected domain. If D is a simply connected domain containing 0, the radial SLE curves in D start at ∂D and converge to 0 as $t \to \infty$. A fundamental property of SLE is its *conformal invariant.*

Proposition 2.4.3 (Conformal invariance)**.** *Let D be a simply connected domain containing* 0 *and let $x \in \partial D$. Let $\eta^{\mathbb{D},1,0}(\kappa)$ denote the law of* SLE(κ) *in $\mathbb{D}$ between* 1 *and* 0*, and let $\nu^{D,x,0}$ be the law of* SLE(κ) *in D between x and* 0*. If $g : \mathbb{D} \mapsto D$ is the unique conformal map between D and $\mathbb{D}$ with $g(1) = x$ and fixing* 0*, then*

$$\nu^{D,x,0} = g \circ \eta^{\mathbb{D},1,0}(\kappa).$$

Convergence of LERW to SLE

Let $D \subsetneq \mathbb{C}$ be a simply connected domain with $0 \in D$ and let $D_\delta = \delta\mathbb{Z}^2 \cap D$. Let ν_δ^D be the law of a loop-erased random walk on D_δ, started at 0 and stopped when it hits ∂D_δ. Let η^D be the law of a radial SLE_2 path from 0 to the boundary of D.

Theorem 2.4.4 (Lawler-Schramm-Werner [122, Theorem 1.1])**.** *The measures ν_δ^D converge weakly to η^D as $\delta \to 0$, in the space of unparameterized curves $(\bar{\mathcal{C}}, \rho_C)$.*

The planar case is well-understood for the scaling limit of the simple random walk. Here, let $\mathcal{W}$ denote the law of a planar Brownian motion started at 0 and stopped on its first exist from the disk $\mathbb{D} = B(0,1)$ on the complex plane. Let G be a planar graph such that the simple random walk on it is irreducible. We denote by μ_δ the law of the simple random walk on the scaled graph δG, started at 0 and stopped when it exits $\mathbb{D}$, and $\nu_\delta^{\mathbb{D},G}$ is the law of the loop-erased random walk on $\delta G \cap \mathbb{D}$.

Theorem 2.4.5 ([171, Theorem 1.1])**.** *Let $(\delta_n)_{n\in\mathbb{N}}$ be a sequence converging to 0. If μ_{δ_n} converges weakly to $\mathcal{W}$ as $n \to \infty$, then $\nu_{\delta_n}^{\mathbb{D},G}$ converges weakly to $\eta^{\mathbb{D}}$.*

Lawler and Viklund proved the convergence of [122] for the *natural parametrization.* They considered a planar LERW parameterized by its renormalized length, and showed that these curves converge to SLE(2) parameterized by its Minkowski content [119].

2.4.2 Three dimensions

In the three-dimensional case, Kozma proved the existence of the scaling limit of the loop-erased random walk in a polyhedral domain, along the scaling subsequence 2^{-n} [107].

For a set $D \subset \mathbb{R}^3$ and $a \in \mathbb{R}^3$, we write $D_{2^{-n}} := D \cap 2^{-n}\mathbb{Z}^3$ and a_{2^n} for the closest point to a in $2^{-n}\mathbb{Z}^3$.

Theorem 2.4.6 (Kozma [107, Theorem 6 and Subsection 6.1])**.** *Let $D \subset \mathbb{R}^3$ be a polyhedron and let $a \in D$. Let $\mathcal{L}_{2^n}$ be a loop-erased random walk on D_{2^n}, started at a_{2^n} and stopped at ∂D_{2^n}. Then the law of $\mathcal{L}_{2^n}$ converges weakly as $n \to \infty$, with respect to the Hausdorff topology.*

Moreover, if $\mathcal{K}$ is a sample of the scaling limit of $\mathcal{L}_{2^n}$, then $\mathcal{K}$ is invariant under dilations and rotations.

Sapozhnikov and Shiraishi proved further topological properties.

Theorem 2.4.7 (Sapozhnikov-Shiraishi [148, Theorem 1.2])**.** *The scaling limit $\mathcal{K}$ is a simple path, almost surely.*

Heuristically, the growth exponent gives an estimate of the number of boxes (or cells in $\mathbb{Z}^3$), which are hit by the loop-erased random walk. The number of such boxes is related to the Hausdorff dimension of the LERW. Shiraishi proved that, in the scaling limit of loop-erased random walks, the growth exponent is the Hausdorff dimension. The next theorem builds up from an upper bound in [148].

Theorem 2.4.8 (Shiraishi [156, Theorem 1.1.1])**.** *The Hausdorff dimension of $\mathcal{K}$ is equal to β_3 (as defined in (2.5)), almost surely.*

2.4.3 Four and higher dimensions

In dimensions $d \geq 4$, the loop-erased random walk converges to Brownian motion. In this case, the exponents present logarithmic corrections $d = 4$ [110, 111]. The scaling is the following.

Theorem 2.4.9 (Lawler [113, Theorem 7.7.6])**.** *Let $d = 4$. Consider a simple random walk S on $\mathbb{Z}^4$ and denote its loop-erasure by $L[0,n] = \mathrm{LE}\, S[0,n]$. There exists a non-negative sequence (a_n)*

$$L_n^*(t) = \frac{\sqrt{d}\sqrt{a_n} L([nt])}{\sqrt{n}}, \qquad t \in [0,1].$$

converges weakly to a 4-dimensional Brownian motion with respect to the the space of continuous $C^d[0,1]$.

The scaling is easier for $d \geq 5$. In these high dimensions, the simple random walk intersects itself infrequently and in relatively small loops. After loop-erasure, the LERW preserves a positive fraction of the points in the simple random walk. We denote this fraction by a. Moreover, the erased loops are negligible when we re-scale the space. Hence, as we take the scaling limit, the loop-erased random walk behaves like a random walk scaled by a. Therefore, a high-dimensional loop-erased random walk converges to Brownian motion, in the scaling limit. Lawler proved this convergence in [110]; for a concise proof, we refer to [113].

Theorem 2.4.10 (Lawler [113, Theorem 7.7.6]). *Let S be a simple random walk on $\mathbb{Z}^d$, with $d \geq 5$. Denote its loop-erasure by L_n, $L[0,n] = \mathrm{LE}\, S[0,n]$, and write*

$$L_n^*(t) = \frac{\sqrt{d}\sqrt{a}L([nt])}{\sqrt{n}}, \qquad t \in [0,1].$$

Then L_n^* converges weakly to a d-dimensional Brownian motion, in the space of continuous functions $C^d[0,1]$, endowed with the supremum norm.

Chapter 3

Uniform Spanning Forests

The uniform spanning forest on $\mathbb{Z}^d$ arises from the infinite-volume limit of uniform spanning tree measures of a growing sequence of boxes. Within probability theory, Pemantle was the first to study uniform spanning forests [143], while the work of Benjamini, Lyons, Peres, and Schramm brought the field to maturity [29]. In this chapter, we first introduce the definition and first properties of the uniform spanning forest in Section 3.1. A remarkable feature of the USF is its close relation to other probabilistic models. Section 3.2 describes the relation of uniform spanning forests with other models in statistical mechanics, while Section 3.3 includes more connections in the form of sampling algorithms. We finish the chapter with a survey of results on scaling limits of uniform spanning forests in Section 3.5.

3.1 Definition and basic properties

We follow the definition of uniform spanning forests in [132, Chapter 10]. In a finite and connected graph $G = (V, E)$, a *spanning tree* of G is a connected subgraph $\mathcal{T}$ such that, for any pair of vertices $v, w \in G$, there is a unique path in $\mathcal{T}$ connecting v and w. The **uniform spanning tree** (UST) of G is a uniform sample over the collection of spanning trees of G.

For an infinite graph, we take a weak limit of the uniform spanning tree measure over an increasing sequence of graphs. Let G be an infinite connected and locally finite graph. An *exhaustion* of G is a sequence $G_n = (V_n, E_n)$ of induced subgraphs of G that are finite and connected, and such that $V_n \subset V_{n+1}$ and $V = \cup_n V_n$. Let $\mathrm{UST}^F_{G_n}$ be the uniform spanning tree measure on G_n. We add a superscript F to indicate *free boundary conditions*. Alternatively, for each induced subgraph G_n, we let G_n^W be G_n with a wired boundary. We denote the uniform spanning tree measure of G_n^W by $\mathrm{UST}^W_{G_n}$. The superscript W indicates *wired boundary conditions*.

Let $\Omega = \{0,1\}^E$ be the space of subgraphs of the infinite graph G. Each element $\omega = (\omega_e)_{e\in E} \in \Omega$ represents a subgraph of G, under the correspondence that an edge $e \in E$ is present in the associated graph if, and only if, $\omega_e = 1$. We endow Ω with the product topology, and $\mathbb{B}_\Omega$ denotes the corresponding Borel sets. Let B be a finite set of edges of E and let $\mathcal{T}$ be a random spanning tree, then the limits

$$\lim_{n\to\infty} \mathrm{USF}^F_{G_n}(B \subset \mathcal{T}), \qquad \lim_{n\to\infty} \mathrm{USF}^W_{G_n}(B \subset \mathcal{T})$$

exist and do not depend on the exhaustion G_n (see [132, Section 10.1]). We thus define the **free uniform spanning forest** measure (FUSF) and the **wired uniform spanning forest** (WUSF) measure of G as the weak limits

$$\mathrm{UST}^F_{G_n} \xrightarrow[n\to\infty]{} \mathrm{FUSF}, \qquad \mathrm{UST}^W_{G_n} \xrightarrow[n\to\infty]{} \mathrm{WUSF},$$

respectively.

In this thesis, we study uniform spanning forests of $\mathbb{Z}^d$. In $\mathbb{Z}^d$ the free and the wired uniform spanning forests coincide WUSF = FUSF, so we refer to both as the **uniform spanning forest** measure of $\mathbb{Z}^d$ (USF).

Theorem 3.1.1 (Pemantle [143])**.** *The support of the uniform spanning forest measure of $\mathbb{Z}^d$ is on disconnected subgraphs in dimensions $d \geq 5$, and connected subgraphs in dimensions $d \leq 4$.*

We refer to a random subgraph $\mathcal{U}$ of $\mathbb{Z}^d$ with the USF measure as a uniform spanning forest, for $d \geq 5$. In dimensions, $d = 2, 3$, and 4, $\mathcal{U}$ is simply called a uniform spanning tree.

Given that a uniform spanning tree $\mathcal{T}$ connects vertices without creating cycles, the presence of an edge in $\mathcal{T}$ depends on other edges. We state this intuition as the *negative correlation* property.

Proposition 3.1.2. *Let $\mathcal{T}$ be a uniform spanning forest. For two different edges $e, f \in E$,*

$$P(e \in \mathcal{T} | f \in T) \leq P(e \in \mathcal{T}).$$

We refer to [77, Theorem 2.1] for a proof of Proposition 3.1.2 for uniform spanning trees on a finite graph. It extends to uniform spanning forests by taking limits.

3.2 Relation to other models

A remarkable feature of uniform spanning forests is its deep relation to other probabilistic models. These include electric networks [29, 40, 101], the random-cluster model [76–78], the Gaussian free field [31, 118], the bi-Laplacian Gaussian field [123], domino tilings [95], the Abelian sandpile model [15, 89, 91, 92, 134], and the rotor-router model [44, 45, 87]. A different type of connection is through sampling algorithms, as it is the case for the simple random walk [5, 38], the loop-erased random walk [29, 169], and the interlacement process [88]. We review these sampling algorithms in Section 3.3.

We then refer to [132, Chapter 2, 4] and [90, Section 4] for the connection between electric networks and uniform spanning trees. The lecture notes [167, Chapter 2] present a concise explanation of the relation between uniform spanning trees and the discrete Gaussian Free Field. In [90, Section 10], we find the Majumdar-Dhar correspondence between Abelian sandpiles and uniform spanning trees. In this section, we will focus on the connection to the random cluster-model, following [77].

3.2.1 The random-cluster model

The random cluster model unifies percolation, the Ising model, and the Potts model in a single framework. The uniform spanning tree is a limit case, in the sense of Theorem 3.2.1.

Let $G = (V, E)$ be a finite graph with configuration space $\Omega = \{0,1\}^E$. An edge $e \in E$ on state 1 is *open* while state 0 indicates a *closed* edge. For each $\omega \in \Omega$, let $\eta(\omega) = \{e \in E : \omega(e) = 1\}$ and $k(\omega)$ is the number of connected components of the subgraph $(V, \eta(\omega))$. The **random-cluster measure** on Ω with parameters $p \in [0,1]$, $q \in (0,\infty)$ is given by

$$\phi_{p,q}(\omega) = \frac{1}{Z}\left\{\prod_{e\in E} p^{\omega(e)}(1-p)^{1-\omega(e)}\right\} q^{k(\omega)}, \qquad \omega \in \Omega,$$

where Z is the normalizing constant, also known as partition function.

Theorem 3.2.1. *The random-cluster measure converges $\phi_{p,q}$ weakly to the uniform spanning tree measure* UST *as $q \to 0$, under the condition that $p \to 0$ and $q/p \to 0$.*

3.3 Sampling algorithms

In the subsections above, presented above the simple random walk, loop-erased random walk, the interlacement process, and the continuum random tree. Our next task is to show the connection of these models to the uniform spanning forest. In the case of the interlacement process, simple random walk and loop-erased random walks, these connections appear as sampling algorithms.

3.3.1 Wilson's algorithm

Wilson's algorithm is an essential theoretical tool in the study of USF [132]. Since it gives an explicit connection between loop-erased random walks and spanning trees, we can translate questions about uniform spanning forests to questions about loop-erased random walks.

Algorithm. Let G be a graph with a finite number of vertices $V = \{v_i\}$.

- Set v_1 as the root and $T_1 = \{v_1\}$.
- For $i = 1, \ldots, |V|$, given a subtree T_i, let γ be loop-erased random walk starting at v_{i+1} and finishing at T_i. Then we set $T_{i+1} = T_i \cup \{\gamma\}$.

Theorem 3.3.1 (Wilson [169]). *The tree $T_{|V|}$ is a uniform spanning tree of G, and T_i is the subtree of $T_{|V|}$ spanned by $\{v_1, \ldots, v_i\}$.*

Wilson's algorithm also works for infinite recurrent graphs. In the case of an infinite transient graph, Benjamini, Lyons, Peres, and Schramm extended Wilson's algorithm in [29] as follows. Note that in Wilson's algorithm, we consider a loop-erased random walk until it hits the root v_0. In the extension for transient graphs, we let the loop-erased random walk continue until it "hits infinity". For this reason, we call the extension **Wilson's algorithm rooted at infinity**.

Algorithm (Wilson's algorithm rooted at infinity). Let G be an infinite transient graph and let $V = \{v_0, v_1, \ldots\}$ be an enumeration of its vertices.

- Let γ_0 be an infinite loop-erased random walk starting at 0. Let $T_0 = \gamma_0$ and set v_0 as its root.
- Given T_i, let γ_i be a loop-erased random walk starting at v_{i+1}. This loop-erased random walk can be either infinite, or stopped when it hits T_i. Then set $T_{i+1} = T_i \cup \gamma_i$.

Theorem 3.3.2 (Benjamini-Lyons-Peres-Schramm, [29]). *T_k is the subtree of the uniform spanning forest of G spanned by $\{v_0, \ldots, v_k\}$.*

3.3.2 Aldous-Broder algorithm

The Aldous-Broder algorithm samples a uniform spanning tree over a finite graph. It was proposed, simultaneously, by Aldous and Broder [5, 38]. Given a finite and connected graph G, let R be a simple random walk on G. For each vertex $v \in G$, let $\tau(v) := \tau_{\{v\}}$ the hitting time of v, as defined in (2.1). Then the oriented edge

$$e(v) := \{R(\tau(v) - 1)\}$$

is the *first entrance edge* of v.

Theorem 3.3.3 (Aldous-Broder [5, 38]). *The set of first entry edges*

$$\{-e(v) \colon v \in G\}$$

has the distribution of a uniform spanning tree on G, oriented towards the root.

3.3.3 Interlacement Aldous-Broder algorithm

The interlacement Aldous-Broder algorithm extends the classic algorithm to infinite graphs. Instead of taking the first entrance edges in the simple random walk, the *interlacement Aldous-Broder algorithm* takes the first entrance edges in an interlacement process [88].

Process (Interlacement Aldous-Broder). Let $\mathscr{I}$ be an interlacement process on $\mathbb{Z}^d$. For each $v \in \mathbb{Z}^d$, we define the first hitting time of the vertex v as

$$\tau^t(v) := \inf\{s \geq t \colon \exists (W, s) \in \mathscr{I} \text{ such that } v \in W\}. \tag{3.1}$$

Let $e^t(v)$ be the oriented edge of $\mathbb{Z}^d$ that is traversed by the trajectory $W_{\tau^t(v)}$, as it enters v for the first time. For each $t \in \mathbb{R}$, let

$$\mathrm{AB}^t := \{-e^t(v) \colon v \in \mathbb{Z}^d\}. \tag{3.2}$$

Theorem 3.3.4 (Hutchcroft [88, Theorem 1.1]). *The set AB^t has the law of the uniform spanning forest of $\mathbb{Z}^d$, oriented towards the root.*

3.4 Encodings of uniform spanning trees

The uniform spanning tree of $\mathbb{Z}^d$ has a natural embedding on the space $\mathbb{R}^3$. As we take the scaling limit, the number of vertices of the UST within any neighbourhood increases, and eventually fills the space. The scaling limit is no longer a graph. Therefore, the study of these scaling limits requires encoding of the uniform spanning tree that carries properties on to the limit.

3.4.1 Paths

Schramm proposed an encoding for the uniform spanning tree in terms of the collection of its paths [149]. Let $\mathcal{U}_n$ be the UST in $2^{-n}\mathbb{Z}^3$, with the point at ∞ added to get a closed set in $\mathbb{S}^3$. For $a, b \in \mathcal{U}_n$, $\omega_n^{a,b}$ denotes the path in $\mathcal{U}_n$ from a to b. The **paths ensemble**

$$\mathcal{I}_n := \left\{(a, b, \omega_n^{a,b}) : a, b \in \mathcal{U}_n\right\}$$

is the collection of all paths in the UST of $2^{-n}\mathbb{Z}^3$. Note that $\mathcal{I}_n$ is a subset of $\mathscr{P} := \mathbb{S}^3 \times \mathbb{S}^3 \times \mathcal{H}(\mathbb{S}^3)$, where $\mathcal{H}(\mathbb{S}^3)$ is the collection of closed subsets of $\mathbb{S}^3$. $\mathscr{P}$ is a compact space and we endow it with the Hausdorff topology:

$$d_H(A, B) = \inf\{r \geq 0 \,:\, A \subset B_r,\ B \subseteq A_r\}, \qquad A, B \in \mathscr{P},$$

where $B_r = \{x \in X \,:\, d(x, B) \leq r\}$ is the r-expansion of B.

The topology of paths ensemble quantifies the shape difference among paths between vertices. In particular, it does not take into account the length of these paths inherited from the graph distance. This path-length is known as *intrinsic distance*. A simple approach for studying the convergence of the intrinsic distance is in terms of **finite-dimensional distributions**. For each fixed $k \in \mathbb{N}$, we consider the joint distribution of the distance between k vertices chosen uniformly at random. This approach gives insight into the structure of the typical sub-tree spanned by k vertices.

3.4.2 Graphs as metric spaces

Two compact metric spaces (X, d^X) and (Y, d^Y) are *isometrically equivalent* if there exists an isometric map $\varphi : X \to Y$ Let $\mathbf{M}$ be the space of isometry classes of compact metric space. We endow $\mathbf{M}$ with the **Gromov-Hausdorff** distance d_{GH}, defined by

$$d_{\mathrm{GH}}(X, Y) = \inf_{\varphi, \tilde{\varphi}, Z} d_{\mathrm{H}}^Z(\varphi(X), \tilde{\varphi}(Y)), \qquad X, Y \in \mathbf{M},$$

where the infimum is over all metric spaces Z and isometries $\varphi : X \to Z$ and $\tilde{\varphi} : Y \to Z$.

Additionally to the metric structure, we consider a measure endowed to a compact metric space. Recall that a Polish space is a metric is separable and completely metrizable topological space. For a Polish space X, we denote by $\mathcal{M}_f(X)$ the set of all non-negative Borel measures on X. The Prohorov metric is defined as

$$d_P^X(\mu,\nu) = \inf\left\{\varepsilon > 0 : \begin{array}{c} \mu(A) \leq \nu(A^\varepsilon) + \varepsilon \text{ and} \\ \nu(A) \leq \mu(A^\varepsilon) + \varepsilon \text{ for all Borel sets } A \end{array}\right\}.$$

Let $\mathcal{X} = (X, d^X, \mu^X, \rho^X)$ and $\mathcal{Y} = (Y, d^Y, \mu^Y, \rho^Y)$ be two compact rooted and measure metric space. We define the **Gromov-Hausdorff-Prohorov** distance by

$$d_{\mathrm{GHP}}(\mathcal{X},\mathcal{Y}) = \inf_{\Phi,\tilde{\Phi},Z}\left\{\begin{array}{c} d^Z(\Phi(\rho^X), \tilde{\Phi}(\rho^Y)) + \\ d_{\mathrm{H}}^Z(\Phi(X), \tilde{\Phi}(Y)) + d_P^Z(\Phi_*\mu^X, \tilde{\Phi}_*\mu^Y) \end{array}\right\},$$

where the infimum is over all Polish spaces (Z, d^Z) and isometries $\Phi : X \to Z$ and $\tilde{\Phi} : Y \to Z$.

We are mainly interested in compact metric spaces with a tree-like structure. A **real tree** (or $\mathbb{R}$-tree) (T, d_T) is a metric space satisfying the following conditions for any $x, y \in T$ with $D = d_T(x, y)$

(i) there exists a unique isometric map $\gamma^{(x,y)} : [0, D] \to T$ such that $\gamma^{(x,y)}(0) = x$ and $\gamma^{(x,y)}(D) = y$.

(ii) If $\varphi : [0, 1] \to T$ is a continuous injective map such that $\varphi(0) = x$ and $\varphi(1) = y$ then $\varphi([0, 1]) = \gamma^{(x,y)}([0, D])$.

Let $\mathbf{T}$ be the space of real trees and $\mathbf{T}_c$ denotes the subspace of compact real trees. We write $\mathbf{T}_c^{\mathrm{GH}}$ and $\mathbf{T}_c^{\mathrm{GHP}}$ for the corresponding isometry classes under the Gromov-Hausdorff and Gromov-Hausdorff-Prohorov metric, respectively.

Theorem 3.4.1 ([67, Theorem 1], [1, Corollary 3.2]). *The metric spaces* $(\mathbf{T}_c^{GH}, d_{GH})$ *and* $(\mathbf{T}_c^{GHP}, d_{GHP})$ *are Polish.*

A graph G is easily described by its set of vertices V and edges E. In the scaling limit, it is convenient to consider any graph as a metric space. If G is a connected graph, we endow the set of vertices with the discrete metric given by

$$d_G(x, y) = \inf_{\lambda(x,y)} \operatorname{len}(\lambda(x, y)), \qquad x, y \in V;$$

where the infimum is taken over all paths $\lambda(x, y)$ on G between x and y. We thus say that (G, d_G) is a metric space. Furthermore, we endow G with **counting measure** μ_G. This measure is uniform over the set of vertices.

If G is a finite graph and $\mathcal{U}$ is a uniform spanning tree of G, it is immediate that $(\mathcal{U}, d_{\mathcal{U}}, \mu_{\mathcal{U}})$ is a measured real tree. In Chapter 5 we consider a uniform spanning tree of $\mathbb{Z}^3$ as a locally compact metric space. We extend the topological framework introduced in Subsection 5.2.

3.5 Scaling limits of uniform spanning trees

In this section, we overview different results for the scaling limit of uniform spanning trees. We begin discussing the case of uniform spanning trees on finite graphs, and then we pass to the results known for infinite graphs.

3.5.1 Finite graphs

In the mean-field case, the scaling limit of the uniform spanning tree of a finite graph is the Brownian continuum random tree. We begin with the presentation of this universal object.

The Brownian continuum random tree

The Brownian continuum random tree (CRT) is a compact rooted real tree. Aldous introduced the Brownian CRT as the scaling limit of critical Galton-Watson trees with finite variance and conditioned to have a large number of vertices [6]. We can also describe the Brownian CRT in terms of a Brownian excursion. A Brownian excursion $\mathbf{e} : [0, 1] \to [0, \infty)$ is a Brownian motion conditioned to be positive in $[0, 1)$ and to take the value $\mathbf{e}(1) = 0$.

Definition 3.5.1 ([124, Theorem 2.2])**.** Let

$$d(s,t) = 2\mathbf{e}(s) + 2\mathbf{e}(t) - 2 \inf_{r \in [s \wedge t, s \vee t]} 2\mathbf{e}(r) \qquad \text{for } s,t \geq 0.$$

We define T as the quotient space of $[0,1]$, where we identify points s and t with $d(s,t) = 0$. The CRT is the compact real tree (T, d_T).

Let p be the canonical projection of $[0,1]$ onto T. The interval $[0,1]$ is endowed with the Lebesgue measure $\mathcal{L}$. Then we define the uniform measure μ_T on T is the push-forward measure given by

$$\mu_T(B) := p_\sharp(\mathcal{L})(B) = \mathcal{L}(p^{-1}(B)),$$

where B is a Borel subset of T.

Alternatively, we can sample a Brownian CRT with the stick-breaking algorithm. This description corresponds to the original definition of the Brownian CRT in [6].

Algorithm (Stick-breaking algorithm)*.* Let R be an inhomogeneous Poisson process on $[0,\infty)$ with intensity measure tdt. We write $(R_n)_{n\in\mathbb{N}}$ for the location of the points of the Poisson process in increasing order. Denote the length of the sticks by $L_1 = R_1$, and $L_n = R_n - R_{n-1}$, and consider the following sequence of compact real trees.

1. T_1 is the closed line segment with length L_1. Label the ends of T_1 as z_1 and z_2. The point z_1 is the root of T_1.

2. For $n > 1$, let x be a uniform point in T_{n-1}. Then, attach at x a closed line segment with length L_n. We label by z_{n+1} the end of the segment L_n on the other side of x.

Theorem 3.5.2 (Aldous [7, Corollary 22])**.** *The tree T_n in the stick-breaking algorithm is equal in distribution to a subtree of the Brownian CRT spanned by $n+1$ leaves independently sampled from its uniform measure μ_T. Moreover, the closure of the set $\bigcup_{n\in\mathbb{N}} T_n$ has the distribution of the Brownian CRT T.*

For each $k \in \mathbb{N}$, let $w_1, \dots, w_k$ be independent random points with law μ_T. We thus define the finite-dimensional distribution of the Brownian CRT F_k as the joint distribution

$$(d_T(w_i, w_j))_{1 \le i < j \le k}. \tag{3.3}$$

We remark that $d_T(w_1, w_2)$ corresponds to the length of T_1 in Algorithm 3.5.1, then

$$P\left(d_T(w_1, w_2) > \lambda\right) = \exp\left(-\frac{\lambda^2}{2}\right).$$

In general, the Brownian CRT is the scaling limit of trees arising in combinatorial models. Among others, uniform random finite trees and random uniform unordered trees converge to the Brownian CRT [135].

Convergence of uniform spanning trees to the Brownian CRT

The scaling limit of the complete graph exhibits the mean-field behaviour of a model. In the case of the uniform spanning tree, the scaling limit is the Brownian CRT. The first theorem on this direction is due to Aldous, in the sense of convergence finite-dimensional distributions. Recall that we define the joint distribution of the distance between k leaves on the Brownian CRT in (3.3) as F_k.

Theorem 3.5.3 (Aldous [6])**.** *Let K_n be the complete graph on n vertices and let $\mathcal{U}_n$ be the uniform spanning tree of K_n. For a fixed $k \in \mathbb{N}$, let $x_1, \dots, x_k$ be vertices chosen uniformly at random from K_n and let $d_n(x_i, x_j)$ be the distance between x_i and x_j on $\mathrm{UST}(K_n)$. Then*

$$\left(\frac{d_n(x_i, x_j)}{\sqrt{n}}\right)_{1 \le i < j \le k} \to F_k \text{ as } n \to \infty$$

in distribution.

Mean-field behaviour is characteristic of high-dimensional graphs. The threshold for a high dimension depends on the model. In the case of uniform spanning trees on $\mathbb{Z}^d$, dimensions $d = 4$ is critical and $d \geq 5$ is a high-dimension. This characterization is related to the scaling limit of the loop-erased random walk. As we described in Subsection 2.4.3, the LERW on $\mathbb{Z}^d$ with $d \geq 4$ converges to Brownian motion.

Mean-field behaviour holds for a larger family of connected graphs. Let us introduce some quantities of a graph relevant to the characterization of mean-field behaviour. Here we follow [138]. Let $G = (V, E)$ be a finite connected graph. Let $\hat{\delta}$ be the ratio of the maximum to the minimum degree of G. The *lazy random walk* $X = (X_t)_{t\geq 0}$ is defined on the set of vertices V with transition probability $\mathbf{p}^t(\cdot,\cdot)$. Given X_t, with probability $\frac{1}{2}$ X_t stays on the same vertex and $X_{t+1} = X_t$. With probability $\frac{1}{2}$, the walk changes its position, and X_{t+1} is one of the nearest-neighbours chosen uniformly at random. We denote by $\mathbf{p}^t(u,v) = P_u(X_t = v)$ and by π the stationary distribution of the lazy random walk. We define the *uniform mixing time* of the lazy random walk on G by

$$t_{\text{mix}}(G) := \min\left\{ t \geq 0 : \max_{u,v\in V} |\frac{\mathbf{p}^t(u,v)}{\pi(v)} - 1| \leq \frac{1}{2} \right\}.$$

The *bubble sum* of G is

$$\mathcal{B}(G) := \sum_{t=0}^{t_{\text{mix}}} (t+1) \sup_{v\in V} \mathbf{p}^t(v,v).$$

We say that

(i) G is *D-balanced* if $\hat{d}(G) \leq D$,

(ii) G is *α-mixing* if $t_{\text{mix}}(G) \leq n^{1/2-\alpha}$, and

(iii) G is θ-escaping if $\mathcal{B}(G) \leq \theta$.

The assumptions (i), (ii) and (iii) are proposed by Michaeli, Nachmias and Shalev in [138] as a characterization of mean-field behaviour for finite graphs (with respect to the UST). They also show that these assumptions are sharp on Theorem 3.5.4 for the diameter of the UST of G, which we denote by diam(UST(G)).

Theorem 3.5.4 (Michaeli-Nachmias-Shalev [138, Theorem 1.1]). *For every $D, \alpha, \theta, \varepsilon > 0$, there exists a constant $C = C(D, \alpha, \theta, \varepsilon)$ satisfying the following. Let G be a connected graph on n-vertices and assume that it is D-balanced, α-mixing and θ-escaping then*

$$P\left(C^{-1}\sqrt{n} \leq \mathrm{diam}(\mathrm{UST}(G)) \leq C\sqrt{n}\right) \geq 1 - \varepsilon.$$

Graphs satisfying (i), (ii) and (iii) include the d-dimensional torus $\mathbb{Z}_m^d$, the hypercube $\{0,1\}^m$ and expander graphs. A version of the assumptions (ii) and (iii) was proposed first in [144]. However, instead of (i), the results in [144] assume vertex transitivity. The transitivity hypothesis holds for the d-dimensional torus and the hypercube, and we thus state convergence for finite-dimensional distribution on these cases.

Theorem 3.5.5 (Peres-Revelle [144, Theorem 1.2]). *Let $d \geq 5$ and let (G_n) be either the sequence of d-dimensional torus $\mathbb{Z}_m^d$ on n vertices, the sequence of hypercubes $\{01\}^m$ on n vertices or a d-regular expander family. For a fixed $k \in \mathbb{N}$, let $y_1, \ldots y_k$ be points chosen uniformly at random on G_n We denote by d_n the intrinsic distance on $\mathrm{UST}(G_n)$. Then there exists a sequence of constants (β_n) bounded away from 0 and infinity such that the joint distribution of the distances*

$$\left(\frac{d_n(y_i, y_j}{\beta_n |G_n|^{1/2}}\right)_{1 \leq i < j \leq k} \to F_k$$

in distribution as $n \to \infty$.

The corresponding result for the finite torus $\mathbb{Z}_m^4$ includes logarithmic corrections. These are expected for the critical dimension $d = 4$.

Theorem 3.5.6 (Schweinsberg [151, Theorem 1.1]). *Let (G_n) be either the sequence of d-dimensional torus $\mathbb{Z}_m^d$ on n vertices and we denote by d_n the intrinsic distance on $\mathrm{UST}(G_n)$. For a fixed $k \in \mathbb{N}$, let $z_1, \ldots, z_k$ be points chosen uniformly at random on G_n. There exists a sequence of constants*

(γ_n) bounded away from 0 and infinity, such that

$$\left(\frac{d_n(z_i, z_j}{\gamma_n n^{1/2}(\log n)^{1/6}}\right)_{1\leq i<j\leq k} \to F_k$$

in distribution as $n \to \infty$.

3.5.2 Infinite graphs

In [4], Aizenman, Burchard, Newman, and Wilson described scaling limits of random trees in terms of their collection of subtrees. Following a different approach, Schramm studied in [149] the paths ensembles of UST as defined in Subsection 3.4.1. Both [4] and [149] prove existence of sub-sequential scaling limits in their respective topologies. We present here the sub-sequential scaling limit in the paths ensemble topology. Tightness is an immediate consequence of the definition of the topological space.

Theorem 3.5.7 (Schramm [149, Theorem 1.6])**.** *Let $\mathcal{I}_\delta$ be the paths ensemble of the UST on $\delta\mathbb{S}^2$. If μ_δ is the law of $\mathcal{I}_\delta$, then there is a sub-sequential weak limit $\mu_\delta \to \mu$ with respect to the space $\mathcal{H}\left(\mathbb{S}^2 \times \mathbb{S}^2 \times \mathcal{H}(\mathbb{S}^2)\right)$ as $\delta \to 0$.*

Although [149] does not prove the full convergence of the paths-ensemble, Schramm introduced a characterization of the limit object. He proposed the Schramm-Loewner evolution (SLE) as the conjectured scaling limit of loop-erased random walks, the Peano curve of the UST, and other conformally invariant processes in the plane. As we mentioned in Chapter 2, [122] establishes the convergence of the LERW to SLE(2). Building up from the work in [149], [122] proves the existence of scaling limits of wired and free UST in bounded domains with smooth boundary, with respect to their paths ensembles. One of the main results in [122] is on the conformal invariance of the scaling limit. With this in mind, we define the scaling limit of the UST on a domain of the complex plane.

Let $D \subsetneq \mathbb{C}$ be a simple domain. We consider both the wired UST on $\delta\mathbb{Z}^2 \cap D$, denoted by WT_δ^D; and the free UST on $\delta\mathbb{Z}^2 \cap D$, denoted by FT_δ^D. Similarly, as in the infinite volume case, we consider their paths ensembles.

Let $\mathcal{WI}_\delta^D$ and $\mathcal{FI}_\delta^D$ be the paths ensembles of WT_δ^D and FT_δ^D, respectively. These paths ensembles are elements of the space $\mathcal{H}(\bar{D} \times \bar{D} \times \mathcal{H}(\bar{D}))$ with the Hausdorff topology. We denote their laws by $\mu_\delta^{W,D}$ and $\mu_\delta^{F,D}$.

Theorem 3.5.8 (Lawler-Schramm-Werner [122, Corollary 1.2])**.** *Let $D \subsetneq \mathbb{C}$ be a simple domain such that ∂D is a C^1-smooth simple closed curve. Then the weak limit of the wired UST and the free UST on D exists:*

$$\mu_\delta^{W,D} \to \mu^{W,D}, \qquad \mu_\delta^{F,D} \to \mu^{F,D}$$

as $\delta \to 0$. Moreover, the scaling limits $\mu^{W,D}$ and $\mu^{F,D}$ are conformally invariant.

The paths ensemble allows us to study the scaling limit as a subset of $\mathbb{R}^2$. Indeed, Schramm gave a complete topological description of the scaling limit of the planar UST in [149]. Nevertheless, the topology of the paths ensemble is inadequate to study other properties, such as the intrinsic metric, the uniform measure, and the simple random walk over the limit object. For this latter purpose, it is more convenient to consider the UST as a measured metric space; in fact, the UST is a real tree.

Barlow, Croydon, and Kumagai considered the uniform spanning tree as a quintuple $\mathcal{T} = (\mathcal{U}, d_\mathcal{U}, \mu_\mathcal{U}, \phi_\mathcal{U}, 0)$, where $\mathcal{U}$ is a uniform spanning tree of $\mathbb{Z}^2$, $d_\mathcal{U}$ is intrinsic metric, $\mu_\mathcal{U}$ is the uniform measure, $\phi_\mathcal{U}$ is an embedding of $\mathcal{U}$ into $\mathbb{R}^2$, and 0 indicates that the embedded tree is rooted at the origin. Their work in [24] proves the existence of subsequential scaling limits of the UST in a Gromov-Hausdorff-Prohorov type topology. It includes the convergence of the embedding $\phi_\mathcal{U}$. Recall that β_2 denotes the growth exponent of the loop-erased random walk.

Theorem 3.5.9 (Barlow-Croydon-Kumagai [24, Theorem 1.1])**.** *Let $\mathbf{P}_\delta$ the law of $(\mathcal{U}, \delta^{-\beta_2} d_\mathcal{U}, \delta^2 \mu_\mathcal{U}, \delta\phi_\mathcal{U}, 0)$. Then the collection $(\mathbf{P}_\delta)_{\delta \in (0,1)}$ is tight with respect to a Gromov-Hausdorff-Prohorov topology.*

We remark that the results in [24] rely on a detailed understanding of the growth properties of the two-dimensional loop-erased random walk. In particular, [24] applies results of [116, 136].

The sub-sequential limit was extended to full convergence in the Gromov-Hausdorff-Prohorov topology by Holden and Sun [86]. In [86], the authors prove the existence of the scaling limit of contour functions of the UST; the convergence is for the space of continuous functions endowed with the topology of uniform convergence on compact sets. This topology is sufficiently strong to imply the convergence of the corresponding real trees (see [1]).

Theorem 3.5.10 (Holden-Sun [86, Theorem 1.1, Remark 1.2]). *The law of the sequence of measured, rooted spatial trees*

$$(\mathcal{U}, \delta^{-\beta_2} d_{\mathcal{U}}, \delta^2 \mu_{\mathcal{U}}, \delta\phi_{\mathcal{U}}, 0)$$

converges as $\delta \to 0$ *with respect to a Gromov-Hausdorff-Prohorov-type topology.*

Chapter 4

Competitive Growth Processes

An interacting particle system is a random model of spatial configurations evolving. The spatial structure is given by a connected graph $G = (V, E)$. We often refer to the vertices of G as **sites**. At any given time, each site is in a **state**, which is an element of a the **local state space** σ. σ may be a set of numbers or letters. They represent the presence (or absence) of different types of particles. A set of local rules governs the interaction of these particles. These interactions induce changes in the state of a site. The **global state (or configuration)** of the system is given by a Markov chain $X = (X_t)_{t\geq 0}$ such that $X_t = (X_t(v))_{v\in V}$ takes values in the collection of functions σ^V.

The simple random walk $S = (S_n : n \geq 0)$ provides an elementary example of a discrete interacting particle system on $\mathbb{Z}^d$. In this case, we set $\sigma = \{0, 1\}$ as the state space. Here 0 represents a *vacant* site, while 1 is a site *occupied* by the random walk. Then, at each integer time $n \in \mathbb{Z}_+$, the Markov chain for the global configuration is defined as

$$X_n(z) = \begin{cases} 1 & \text{if } S_n = X_n(z) \\ 0 & \text{otherwise,} \end{cases} \qquad \forall\, z \in \mathbb{Z}^d.$$

We may think of an occupied site as the position of a single particle. Under this interpretation, at each time $n \in \mathbb{N}$, a particle simultaneously produces an identical child and dies. The child-particle occupies a neighbouring site chosen uniformly.

Our next example is the branching random walk. A branching random walk represents a growing population of identical particles, where each of them reproduces and moves randomly around space, independently from others. A branching process (introduced in Section 1.2) drives the reproduction mechanism. We restrict to a binomial offspring distribution to simplify the exposition and refer to [153, 154] for surveys on this model.

Let $\mathbb{T}$ be the genealogical tree of a branching process with offspring distribution $\mathrm{Bin}(m, p)$. $\mathbb{T}$ is rooted at ρ and has a set of edges E. Consider a collection of random variables $(\zeta_e)_{e \in E}$ indexed by the edges of $\mathbb{T}$. This collection is independent and uniformly distributed over the set of directions on $\mathbb{Z}^d$, $\mathcal{E} = \{\pm e_1, \ldots, \pm e_d\}$, where $e_k(j) = \mathbf{1}\{k = j\}$. For a vertex x in $\mathbb{T}$, the *ancestral line* $[\![\rho, x]\!]$ is the unique path of edges from the root ρ to the vertex x. We denote by $\|x\|$ the generation of x, defined as the number of elements in $[\![\rho, x]\!]$. With this notation, $\|\rho\| = 0$. Let

$$V(x) := \sum_{e \in [\![\rho, x]\!]} \zeta_e$$

be the sum of increments associated with the ancestors of x. The **branching random walk** on $\mathbb{Z}^d$, with offspring distribution $B(m, p)$, is the collection of random variables $(V(x) : \, x \in \mathbb{T})$. For each generation n, we obtain a finite point process

$$(V(x_{1,n}), \ldots V(x_{N,n}) : \, i = 1, \ldots, N)\,,$$

where $\{x_{i,n}\}_{1 \le i \le N}$ is the set of vertices in the n-th generation. Note that $N \ge 0$ is random. The number of vertices at generation n is a Galton-Watson process. We can also describe the branching random walk like a (discrete-time) interacting particle system. In this case, the local state space is $\sigma = \{0, 1, 2, \ldots\}$. The state of a site $z \in \mathbb{Z}^d$ indicates the number of particles at that position. For each integer time $n \geq 0$, the global configuration is

$$X_n(z) = \sum_{x:\, \|x\|=n} \mathbf{1}\{V(x) = z\}, \qquad \forall\, z \in \mathbb{Z}^d.$$

At each time $n \in \mathbb{Z}_+$, a particle at site z dies and simultaneously gives birth to a random number of particles. The offspring occupies positions uniformly distributed among the nearest-neighbours of z.

For the random walk and the branching random walk, particles evolve at a discrete time. In the following section, we consider interacting particle systems in continuous time. We need to introduce a framework to justify that the Markov chain $X = (X_t)_{t \geq 0}$ is well-defined. This construction is our main task in Section 4.2. A second limitation in the branching random walk model lies in the independence between different particles. If we think of these systems as a spatial population model, competition for resources or predator-prey behaviour are natural assumptions. These assumptions lead to competitive growth processes. As the first example of a competitive growth process, we present the two-type Richardson model, and the original one-type Richardson model, in Section 4.3. The motivation for the study of the Richardson model is related to first passage percolation and serves as an inspiration for further questions. We thus present first passage percolation and its connections to Richardson models in Section 4.4. We finish the chapter with the second example of a competitive growth process, called chase-escape. We overview this model in Section 4.5. This process is closely related to the model that we study in Chapter 7.

4.1 Markov processes

We begin with standard background on Markov process. For this section we follow [102] and [129].

Let E be a Polish space with Borel σ-algebra $\mathcal{B}(E)$. We denote the collection of continuous functions $f : E \to \mathbb{R}$ by $C(E)$. A map $\kappa : \Omega_1 \times \mathcal{A}_2 \to [0, \infty]$ is a **stochastic kernel** between the measure spaces $(\Omega_1, \mathcal{A}_1)$ and $(\Omega_2, \mathcal{A}_2)$ if

(i) $\omega_1 \mapsto \kappa(\omega_1, A_2)$ is $\mathcal{A}_1$-measurable for any $A_2 \in \mathcal{A}_2$, and

(ii) $A_2 \mapsto \kappa(\omega_1, A_2)$ is a probability measure on $(\Omega_2, \mathcal{A}_2)$ for any $\omega_1 \in \Omega_1$.

Let $(X_t)_{t\geq 0}$ be an stochastic process and we write $(\mathcal{F}_t)_{t\geq 0}$ for the filtration generated by X. The stochastic process $X = (X_t)_{t\geq 0}$ is a time-homogeneous **Markov process** with distributions $(\mathbf{P}_\eta)_{\eta\in E}$ if:

(i) For every $\eta \in E$, X is a stochastic process on an abstract probability space $(\Omega, \Sigma, \mathbf{P}_\eta)$ with $\mathbf{P}_\eta\left(X_0 = \eta\right) = 1$.

(ii) The map $\kappa : E \times \mathcal{B}(E)^{\otimes \mathbb{R}_+} \to [0, 1]$ defined by

$$(\eta, B) \mapsto \mathbf{P}_\eta\left(X \in B\right)$$

is a stochastic kernel. For every $t \geq 0$, the *transition kernel* κ_t : $E \times \mathcal{B}(E) \to [0, 1]$ is defined by

$$\kappa_t(\eta, A) := \kappa\Big(\eta, \{y \in E^{\mathbb{R}_+} :\ y(t) \in A\}\Big) = \mathbf{P}_\eta\left(X_t \in A\right).$$

(iii) X satisfies the Markov property: for every $A \in \mathcal{B}(E)$, $\eta \in E$ and $s, t \geq 0$,

$$\mathbf{P}_\eta\left(X_{t+s} \in A \,|\, \mathcal{F}_s\right) = \mathbf{P}_{X_s}\left(X_t \in A\right), \qquad \mathbf{P}_\eta - a.s.$$

A function $f : [0, \infty) \to E$ is **càdlàg** if it is continuous from the right and the left limits exist: for every $t \geq 0$

$$f(t) = \lim_{s\downarrow t} f(s), \text{ and } \lim_{s\uparrow t} f(s) \text{ exists and is finite.}$$

We let $D[0, \infty)$ be the collection of càdlàg functions $X : \mathbb{R}_+ \to E$ and $\mathcal{D}$ is the Borel σ-algebra generated by the evaluation maps $X \mapsto X_t$ for $t \geq 0$. Then $(D[0, \infty), \mathcal{D})$ is the canonical path space for a stochastic process X, and in particular for the Markov processes that we define below. We denote the expectation corresponding to $\mathbf{P}_\eta$ by $\mathbf{E}_\eta$.

Now we additionally assume that the space E is locally compact. The **Markov semigroup** $\{P_t : t \geq 0\}$ associated to the Markov process $X = (X_t)_{t\geq 0}$ starting at $\eta \in E$ is defined as

$$P_t f(\eta) := \mathbf{E}_\eta \left(f(X_t) \right) := \int_E f(\eta_t) d\mathbf{P}_\eta$$

for bounded functions $f : E \to \mathbb{R}$. A Markov process $X = (X_t)_{t\geq 0}$ is a **Feller process** if the Markov semigroup maps the collection of continuous functions $C(E)$ into itself, i.e. for every $f \in C(E)$, $P_t f \in C(E)$ for all $t \geq 0$. The Markov processes arising from interacting particle systems are Feller processes. For our purposes, the fundamental property of a Feller process is that we can define it in terms of its Markov semigroup (see [129, Theorem 1.5]).

The **infinitesimal generator** (or **generator**) of the Markov semigroup $\{P_t : t \geq 0\}$ is defined as the operator

$$Gf := \lim_{t \downarrow 0} \frac{P_t f - f}{t}, \tag{4.1}$$

where f belongs to a subset of $C(E)$ where the limit (4.1) exists. The Hille-Yosida theorem (see [129, Theorem 2.9]) establishes a one-to-one correspondence between infinitesimal generators on $C(E)$ and Markov semigroups on $C(E)$. This property is crucial for the definition of interacting particle systems, as we define them in terms of the corresponding generators. Such construction requires additional work. In the next section, we present sufficient conditions for our interacting particle systems and refer to [129, Chapter 1] for details.

4.2 Interacting particle systems

Let G be a connected graph with a countable set of sites V. Recall that we write $\sigma = \{s_1, \ldots, s_n\}$ for the local state space. Our interest is on the global configuration of V, where each site has a local state. The configuration space is σ^V, which is the collection of functions $\eta : V \to \sigma$. We endow σ^V

with the product topology and denote the space of real-valued continuous functions on σ^V by $C(\sigma^V)$.

For the description of an interacting particle system, we require a set of possible transitions between global states, and rates at which these transitions occur. Following [129, 130], we describe these two elements as

(i) a set of **local maps** between global configurations

$$\mathcal{G} = \{\eta^T : \sigma^V \to \sigma^V : T \subset V, |T| < \infty\},$$

where the index T indicates the finite subset of sites where the map η^T changes values on an element $\eta \in \sigma^V$; and

(ii) a collection of non-negative **transition rates**

$$\{c(T, \eta) : \eta^T \in \mathcal{G}\}.$$

Another common notation is to write $c(\eta, \eta^T)$ for $c(T, \eta)$. We will use $c(T, \eta)$ in the construction of interacting particle systems, and $c(\eta, \eta^T)$ for our examples. We assume that the function c is non-negative, uniformly bounded, and continuous as a function of η.

An **interacting particle system** is a continuous-time Markov process $X = (X_t)_{t \geq 0}$ on the configuration space σ^V. Under a suitable set of conditions over the local maps and the rates (see (4.5) and (4.6)), the generator

$$Gf(x) = \sum_{\eta^T \in \mathcal{G}} c(T, \eta) \left(f(\eta^T) - f(\eta) \right), \qquad \eta \in \sigma^G, \ f \in C(\sigma^V) \tag{4.2}$$

defines the Markov process X.

The dynamics in our examples change at one or, at most, two sites, at the same time. Then the collection of local maps and transition rates are

$$\mathcal{G} = \{\eta^x, \eta^{x,y} : x, y \in V\}, \qquad \{c(x, \eta), c(x, y, \eta) : x, y \in V\}. \tag{4.3}$$

We also assume that

$$c(x,y,\eta) = \begin{cases} p(x,y) & \text{if } \eta(x) = 1, \eta(y) = 0, \\ 0 & \text{otherwise.} \end{cases} \tag{4.4}$$

for some non-negative sequence of real numbers $(p(x,y))_{x,y \in V}$. With these assumptions, the generator of X takes the form

$$Gf(x) = \sum_x c(x,\eta) \left(f(\eta^x) - f(\eta) \right) + \sum_{x,y} c(x,y,\eta) \left(f(\eta^{x,y}) - f(\eta) \right),$$

for $\eta \in \sigma^V$ and $f \in C(\sigma^V)$.

As pointed out above, we require some assumptions over the local maps and their rates. For an interacting particle system of the form (4.3) and satisfying (4.4), it suffices that

$$\sup_{x \in V} \sum_{u \in V} \sup_{\eta \in \sigma^V} |c(x,\eta) - c(x,\eta_u)| < \infty, \tag{4.5}$$

and

$$\sup_{y \in V} \sum_{x \in V} p(x,y) < \infty. \tag{4.6}$$

Theorem 4.2.1 (Liggett, [130, Theorem B3])**.** *Consider the description of an interacting particle system of the form* (4.3) *and satisfying* (4.4)*. If* (4.5) *and* (4.6) *are also satisfied, then the closure of* (4.2) *is the generator of a Feller Markov process $X = (X_t)_{t \geq 0}$ on the space of global configurations σ^V.*

A general construction for a finite local state space σ is in [162, Chapter 4], while [129, Theorem 3.9] gives general conditions for the existence of particle systems with a countable local state space.

The interacting particle systems in this work satisfy the hypothesis of Theorem 4.2.1. In particular, (4.5) is a consequence of the finite range of the rates in our examples. We say that a rate $c(x,\eta)$ has finite range if there exists a constant C such that $c(x,\eta)$ depends on η through at most C coordinates of η. Conditions (4.3), (4.4) are part of the construction and (4.6) is easy to verify.

Example 4.2.2. Let us define the **continuous-time random walk** as an interacting particle system. The underlying graph is $\mathbb{Z}^d$ and the state space is $\sigma = \{0, 1\}$. Similarly to the example above of a simple random walk, 0 indicates a vacant site, while 1 indicates a site occupied by the random walk. We define the local map $\eta^{x,y} : \sigma^{\mathbb{Z}^d} \to \sigma^{\mathbb{Z}^d}$ for $x, y \in \mathbb{Z}^d$ nearest-neighbours (i.e. $|x - y| = 1$) by

$$\eta^{x,y}(z) = \eta^{x,y}(\eta)(z) = \begin{cases} \eta(z) & \text{if } z \neq x, y, \\ \eta(y) & \text{if } z = x, \\ \eta(x) & \text{if } z = y, \end{cases}$$

and the transition rate of this map is

$$c(\eta, \eta^{x,y}) = \mathbf{1}\{|x - y| = 1\}. \tag{4.7}$$

The rate of any other local map is 0. The map $\eta^{x,y}$ exchanges the position of the particle if (and only if) one of these sites is occupied. The meaning of (4.7) is that change in positions occurs after a random time T, where $T \sim \text{Exp}(1)$. For the continuous-time random walk on $\mathbb{Z}$, a simple way to indicate the local maps and their transition rates is by writing

$$01 \xrightarrow{1} 10, \qquad 10 \xrightarrow{1} 01.$$

An alternative way to define an interacting particle system is with a Poisson process. Let us define the Poisson process and then discuss the construction of interacting particle systems with an example.

A **Poisson process** on $\mathbb{R}_+$ with intensity λ is a continuous-time Markov process $(N_t,\ t \geq 0)$, valued on $\mathbb{Z}_+$ and satisfying the following conditions.

(i) $N_0 = 0$

(ii) For any finite collection of indices $0 = t_0 < t_1 < \ldots < t_n$, the family of increments $(N_{t_i} - N_{t_{i-1}} :\ i = 1, \ldots, n)$ is independent

(iii) For $t > s \geq 0$, the random variable $N_t - N_s$ follows a Poisson distribution with parameter $\lambda(t-s)$.

We refer to [102, Chapter 5] for a proof of the existence of the Poisson process on $\mathbb{R}_+$ (and in more general spaces). A classic and beautiful reference on the general theory of Poisson processes is [100].

Example 4.2.3. We continue Example 4.2.2 with an alternative construction. Let $(R_t : t \geq 0)$ and $(L_t : t \geq 0)$ be two independent Poisson processes on $\mathbb{R}_+$. Then the continuous-time symmetric random walk on $\mathbb{Z}$ is equal in distribution to

$$X_t = R_t - L_t.$$

Poisson representations are available for more general interacting particle systems. These constructions are in [129, Chapter 3, Section 6] and [162, Theorem 4.14].

4.3 Richardson models

Our first example of competitive growth processes is the Richardson model. They model populations spreading uniformly through a graph. We represent the population growth with the occupancy of vacant vertices. The dynamics are similar to the growth of cells or infections. If an individual is at a given site, it will "conquer" a vacant nearest-neighbour site (chosen uniformly) after an exponential waiting period. We consider two variants: the one-type Richardson model for the growth of a (single) population, and the multiple type Richardson model. The first model considers unobstructed growth, while the multiple-type model corresponds to different species competing for resources.

In the one-type Richardson model, the population is homogeneous. In this case, we only have two states for a site. It is either vacant or occupied. Since our graphs are connected, any site will be occupied after some random time. The main question for this model is on the shape of the occupied sites in the long run.

The two-type Richardson model considers two species in competition for vacant spaces. We distinguish these species as red and blue individuals.

In contrast to the situation for the one-type Richardson model, it is not necessarily true that both species occupy an infinite number of sites. For example, if red particles no longer have vacant sites between their nearest-neighbours (because blue particles are in those sites) then red will not be able to reproduce any longer. We interpret that situation as an *extinction event* for the red particles. A similar event is possible for blue. On the contrary, if both red and blue particles occupy an infinite number of sites, we interpret this as *coexistence*. Our approach to the two-type Richardson models focusses on sufficient conditions for a coexistence event.

In the final subsection, we consider the multiple type Richardson model. It is a generalization of the two-type model for k competing species.

4.3.1 The one-type Richardson model

The one-type Richardson model is an interacting particle system on $\mathbb{Z}^d$ with two states $\sigma = \{0, 1\}$. We refer to state 0 as *vacant* and to 1 as *occupied*. If y is vacant then it becomes occupied at a rate proportional to the number of occupied nearest-neighbours. An occupied site remains that way for the rest of the process.

For a precise definition, we follow the notation in Section 4.2. For each $x \in \mathbb{Z}^d$, we define the *occupation map* (or infection map) $I^x : \sigma^{\mathbb{Z}^d} \to \sigma^{\mathbb{Z}^d}$ by

$$I^x(\eta)(z) = \begin{cases} 1 & \text{if } z = x, \\ \eta(z) & \text{otherwise,} \end{cases} \qquad \eta \in \sigma^{\mathbb{Z}^d}.$$

The **one-type Richardson model** is the Markov process taking values in $\{0, 1\}^{\mathbb{Z}^d}$ with transition rates

$$c(\eta, I^x(\eta)) = \sum_{\substack{y \in \mathbb{Z}^d \\ |x-y|=1}} \mathbf{1}\{\eta(y) = 1\},$$

with the initial configuration $\eta(0) = 1$ and $\eta(z) = 0$ for any other $z \in \mathbb{Z}^d$. Richardson introduced this interacting particle system in [147] as a model for cell-growth. We remark that the original definition was for a discrete time Markov process.

We are interested in the asymptotic shape of the occupied sites. For each $t \geq 0$, we let

$$B(t) := \{z \in \mathbb{Z}^d : \ X_t(z) = 1\} \tag{4.8}$$

denote the occupied vertices at time $t \geq 0$. The main theorem in Richardson's seminal work [147] is on the asymptotic shape of $B(t)$. It is shown in [147] that there exists a convex and compact deterministic set $\mathcal{B}_R$ such that, for any $\varepsilon > 0$, the probability of the event

$$(1-\varepsilon)\mathcal{B}_R \subset \frac{B(t)}{t} \subset (1+\varepsilon)\mathcal{B}_R \tag{4.9}$$

tends to 1 as $t \to \infty$. The *shape theorem* in [147] if the first result of its kind. Cox and Durrett observed in [48] that a lemma suggested by Kesten in [39] improves Richardson's theorem: the event in (4.9) holds almost surely for $t \geq 0$ large enough. In [48], Cox and Durrett generalized the shape theorem to the setting of first passage percolation. We present this generalization as Theorem 4.4.2. We discuss the relation between the Richardson model and first passage percolation in Section 4.4.1.

A model similar to the one-type Richardson model is the Eden model [63]. The Eden model has a simple construction on $\mathbb{Z}^d$. For simplicity, we describe the process in terms of cell growth, as in its original formulation. We start with a cell at the origin. This cell divides into an identical daughter, and the newborn cell occupies one of the neighbouring sites, chosen uniformly at random. The process continues its reproduction in the same way. To describe the evolution of the Eden process, for each $n \geq \mathbb{Z}_+$, we let $\mathcal{A}(0) = \{0\}$ and define $\mathcal{A}(n)$ as the set of vertices after the n-th reproduction.

The one-type Richardson model and the Eden model are the same up to a suitable time scale. Here we follow [17]. From the collection of discrete balls $(B(t))_{t\geq 0}$ in (4.8), we construct a sequence of random times $\{N_k\}_{k\in\mathbb{Z}_+}$. We define $N_0 = 0$ and

$$N_k = \inf\{t \geq 0 : \ B(t) \text{ contains } k+1 \text{ points of } \mathbb{Z}^d\}.$$

Richardson observed that the collection of subsets $\{A(k) : \ k \geq 1\}$ and $\{B(N_k) : \ k \geq 1\}$ have the same distribution [147, Example 9].

4.3.2 The two-type Richardson model

In the two-type Richardson model we have three states: $\sigma = \{w, r, b\}$ for the vertices of $\mathbb{Z}^d$, for $d \geq 3$. Similarly to the one-type version, only sites at state w may flip on a rate depending on its number of occupied neighbours. Once a site reaches states r or b, it remains on that state for the rest of the process. We identify the sites with state r as *red particles*, sites with state b correspond to *blue particles*, while the state w represents a *vacant* site. The blue and red particles represent two different species competing for space. Häggström and Pemantle defined this variant of the one-type Richardson model in [79], as a tool for understanding infinite geodesics in first passage percolation. We expand on the discussion of this connection in Subsection 4.4.2. For a formal definition of this process, we define the local maps acting on the model. For each $x \in \mathbb{Z}^d$, $R^x : \sigma^{\mathbb{Z}^d} \to \sigma^{\mathbb{Z}^d}$ is the *red-occupation map* defined by

$$R^x(\eta)(z) = \begin{cases} r & \text{if } z = x \text{ and } \eta(x) = w, \\ \eta(z) & \text{otherwise,} \end{cases} \qquad \eta \in \sigma^{\mathbb{Z}^d}.$$

The *blue-occupation map* $B^x : \sigma^{\mathbb{Z}^d} \to \sigma^{\mathbb{Z}^d}$ is given by

$$B^x(\eta)(z) = \begin{cases} b & \text{if } z = x \text{ and } \eta(x) = w, \\ \eta(z) & \text{otherwise,} \end{cases} \qquad \eta \in \sigma^{\mathbb{Z}^d}.$$

The **two-type Richardson model** with parameter λ is the Markov process taking values in $\{w, r, b\}^{\mathbb{Z}^d}$ with transition rates

$$c(\eta, R^x(\eta)) = \lambda \sum_{\substack{y \in \mathbb{Z}^d \\ |x-y|=1}} \mathbf{1}\{\eta(y) = r\}, \qquad c(\eta, B^x(\eta)) = \sum_{\substack{y \in \mathbb{Z}^d \\ |x-y|=1}} \mathbf{1}\{\eta(y) = b\}.$$

The initial condition is the configuration

$$\eta(x) = \begin{cases} r & \text{if } x = 0, \\ b & \text{if } x = (1, 0, \dots, 0), \\ w & \text{otherwise.} \end{cases} \tag{4.10}$$

We denote the probability measure associated to this process by P_λ.

It is reasonable to think that λ rules the coexistence behaviour of red and blue particles. When $\lambda = 1$, the red and blue particles spread at the same "speed" and, intuitively, one would expect both reach an infinite number of sites. Let $\mathcal{A}$ be set of sites occupied by red particles at some time during the process. We define $\mathcal{B}$ for the blue particles in an analogous way. The **coexistence event** is defined as

$$E = \{|\mathcal{A}| = \infty,\ |\mathcal{B}| = \infty\}. \tag{4.11}$$

Positive probability for the coexistence event was first proved in $\mathbb{Z}^2$ by Häggström and Pemantle in [79]. The d-dimensional case was proved independently by Garet and Marchand [72] and Hoffman [83].

Theorem 4.3.1 (Häggström-Pemantle [79, Theorem 1.2], Garet-Marchand [72, Theorem 3.1], Hoffman [83, Theorem 2])**.** *For the two-type Richardson process on $\mathbb{Z}^d$ with $\lambda = 1$, the coexistence event E has positive probability.*

Häggström and Pemantle conjectured in [79] and in [80, Conjecture 1.1] that the converse of Theorem 4.3.1 holds, and the coexistence event E has probability zero whenever $\lambda \neq 1$ for all $\mathbb{Z}^d$ with $d \geq 2$. The article [80] gives a partial result, but the general case is an open question.

Theorem 4.3.2 (Häggström-Pemantle [80, Theorem 1.2])**.** *For the two-type Richardson process on $\mathbb{Z}^d$,*

$$P_\lambda(E) = 0$$

for all $\lambda \in \mathbb{R}_+ \setminus \Lambda$, and the cardinality of Λ is at most countable.

A variation is to study the probability of the coexistence event under different initial conditions. If two sites other than 0 and $(1, 0, \dots, 0)$ are occupied at $t = 0$, then the situation is equivalent to the initial conditions in (4.10) [79, Proposition 1.1]. The same is true for any initial condition with a finite number of occupied sites [55, Theorem 1]. We have a change in the model when a infinite number of particles are present at the beginning of the process. Consider the initial configuration on $\mathbb{Z}^d$, for $d \geq 2$:

$$\eta^{\mathcal{H}}(x) = \begin{cases} r & \text{if } x = 0, \\ b & \text{if } r_1 = 0,\, x \neq 0, \\ w & \text{otherwise,} \end{cases}$$

where $x = (r_1, \ldots, r_d)$. We denote the probability measure associated with this process by $P_\lambda^{\mathcal{H}}$. Recall that λ corresponds to the transition rate of the red particles.

Theorem 4.3.3 (Deijfen-Häggström [57, Theorem 1.1])**.** *For the two-type Richardson process on $\mathbb{Z}^d$, with $d \geq 2$,*

$$P_\lambda^{\mathcal{H}}(E) > 0 \textit{ if, and only if } \lambda > 1.$$

4.3.3 Multiple type Richardson model

An immediate generalization of the two-type Richardson model is to consider k different species. The process evolves on $\mathbb{Z}^d$, and the local state space is $\{0, \ldots, k\}$. When a site has state $j \geq 1$, it means that it has been occupied by a particle of type j. It remains at that state for the rest of the process. Otherwise, the site has state $j = 0$, meaning that it is vacant. For $j = 1, \ldots, k$, the I_j^x*-occupation map* is the analogue of the blue occupation maps defined for the two-type Richardson model. In words, the j-type occupies an adjacent vacant vertex after an exponential time Exp(1).

Let $x_1, \ldots, x_k$ be different sites in $\mathbb{Z}^d$. The k-type Richardson model with initial conditions $(x_1, \ldots, x_k)$ is the Markov process $X^k = (X_t^k)_{t \geq 0}$ taking values in $\{0, 1, \ldots, k\}^{\mathbb{Z}^d}$, with rates

$$c(\eta, I_j^x(\eta)) = \sum_{\substack{y \in \mathbb{Z}^d \\ |x-y|=1}} \mathbf{1}\{\eta(y) = j\},$$

and initial conditions

$$\eta(z) = \begin{cases} j & \text{if } z = x_j, \\ 0 & \text{otherwise.} \end{cases}$$

For $j = 1, \ldots, k$, we write

$$\mathcal{B}_j = \{z \in \mathbb{Z}^d : \ X^k(t) = j \text{ for some } t \geq 0\}$$

for the set of sites that eventually become of type j. The *coexistence event* for the k-type Richardson model with initial conditions $(x_1, \ldots, x_k)$ is

$$E(x_1, \ldots, x_k) := \{|\mathcal{B}_j| = \infty : \ j = 1, \ldots, k\}. \tag{4.12}$$

In line with Theorem 4.3.1, Hoffman proved the next theorem as a tool to obtain a lower bound on the number of infinite geodesics in first passage percolation (c.f. Theorem 4.4.4).

Theorem 4.3.4 (Hoffman [84, Theorem 1.6])**.** *Consider the* 4*-type Richardson model on* $\mathbb{Z}^2$*. For any* $\varepsilon > 0$ *there exist* x_1*,* x_2*,* x_3 *and* x_4 *in* $\mathbb{Z}^2$ *such that*

$$P(E(x_1, x_2, x_3, x_4))) > 1 - \varepsilon.$$

4.4 First passage percolation

Hammersley and Welsh introduced first passage percolation (FPP) as an extension of the percolation model [81]. The traditional example of FPP is the phenomenon of fluid moving through a random medium. While percolation studies the sites where the fluid arrives (occupied sites), FPP incorporates the time of arrival to each site. In this section, we present the definition and main theorems in the theory of FPP. We follow [17] in the presentation of these preliminaries.

Let $G = (V, E)$ be a connected graph and let F be a probability distribution on $(0, \infty)$. In particular, in this work we only consider $F(0) = 0$. The collection of **edge weights** $(\tau_e)_{e \in E}$ is a family of independent random variables with common distribution F. Recall that a *path* on G between the vertices x and y is a collection of vertices $[v_1, \ldots, v_n]$ such that $(v_i, v_{i+1}) \in E$, $v_1 = x$ and $v_n = y$. We define the passage time of a path $\gamma = [v_1, \ldots, v_n]$ by

$$T_F(\gamma) = \sum_{i=1}^{n-1} \tau_{(v_i, v_{i+1})}.$$

For any $x, y \in V$, we define the **passage time between x and y** by

$$T_F(x, y) := \inf\{T(\gamma) \colon \gamma \in \Gamma(x, y)\}, \tag{4.13}$$

where $\Gamma(x, y)$ is the collection of (finite) paths on G between x and y.

Proposition 4.4.1. *If $F(0) = 0$, then T_F defines a (random) metric on V almost surely.*

If $G = \mathbb{Z}^d$, we extend the passage time to a metric for $\mathbb{R}^d$. We define the metric $T : \mathbb{R} \to (0, \infty)$ by

$$T_F(x, y) = T_F(x', y'),$$

where $x' \in \mathbb{Z}^d$ is the closest vertex to x (in case of a tie, we choose the smallest vertex with respect to the lexicographical order). We choose y' similarly.

We denote by $B_F(t)$ the random open ball centred at the origin in the metric T_F

$$B_F(t) := \{y \in \mathbb{R}^d \colon T_F(0, y) < t\}.$$

One of the main results in the theory of FPP is on the shape of the ball $B_F(t)$ at large scales.

Theorem 4.4.2 (Cox-Durrett [48])**.** *There exists a deterministic, convex* *and compact set $\mathcal{B}_F \subset \mathbb{R}^d$ such that for each $\epsilon > 0$,*

$$(1-\epsilon)\mathcal{B}_F \subset \frac{B_F(t)}{t} \subset (1+\epsilon)\mathcal{B}_F \text{ for all large } t$$

almost surely.

4.4.1 First passage competition models

Competitive growth models are interacting particle systems, but we can define equivalent processes in terms of first passage percolation. We usually refer to these models as *first passage competition models* or *competing first passage percolation.*

We begin with a construction of the one-type Richardson model, defined in Subsection 4.3.1. Consider FPP on $\mathbb{Z}^d$ with exponential (with parameter 1) edge weights and let $T_{\text{Exp}} : \mathbb{Z}^d \times \mathbb{Z}^d \to [0, \infty)$ be the corresponding passage time metric. The random discrete ball

$$B(t) = \{x \in \mathbb{Z} : \, T_{\text{Exp}}(0, x) \leq t\}$$

represents the sites occupied at time t. We thus consider the Markov process $X = (X_t)_{t \geq 0}$ taking values in $\{0, 1\}^{\mathbb{Z}^d}$ and given by

$$X_t(z) = \begin{cases} 1 & \text{if } z \in B(t), \\ 0 & \text{otherwise.} \end{cases}$$

It is well-known that the one-type Richardson model is equal in distribution to $(X_t)_{t \geq 0}$. The process $(X_t)_{t \geq 0}$ is also known as the edge representation of the one-type Richardson model.

Our next example is the two-type Richardson model. For simplicity, we focus on the case $\lambda = 1$. As above, let $T_{\text{Exp}} : \mathbb{Z}^d \times \mathbb{Z}^d \to [0, \infty)$ be the passage time metric defined by FPP, with Exp(1) passage times. Recall that, for the two-type Richardson process at time $t = 0$, the origin $x_1 = 0$ is red, while the site $x_2 = (1, 0, \ldots, 0)$ is blue. The idea is to compare the passage times $T_{\text{Exp}}(x_1, z)$ and $T_{\text{Exp}}(x_2, z)$, which are the passage times from the initial red and blue sites, respectively. The colour with a smaller passage time conquers the site z and occupies it for the rest of the process. This colour is well-defined since the distribution of the passage times is continuous and hence $\{T_{\mathrm{Exp}}(x_1, z), T_{\mathrm{Exp}}(x_2, z)\}$ has a unique minimum. Then, the two-type Richardson model has the distribution of the Markov process $Y = (Y_t)_{t\geq 0}$ with $Y_t(x_1) = r$ and $Y_t(x_2) = b$ for all $t \geq 0$; and for $z \in \mathbb{Z}^d$ different to x_1 and x_2

$$Y_t(z) = \begin{cases} r & \text{if } T_{\mathrm{Exp}}(x_1, z) \leq t \text{ and } T_{\mathrm{Exp}}(x_1, z) < T_{\mathrm{Exp}}(x_2, z), \\ b & \text{if } T_{\mathrm{Exp}}(x_2, z) \leq t \text{ and } T_{\mathrm{Exp}}(x_1, z) > T_{\mathrm{Exp}}(x_2, z). \end{cases}$$

The edge representation of a two-type Richardson model with rate $\lambda \neq 1$ is more complex and requires additional care. We need to enforce the condition that, for every red site v, there is a nearest-neighbour path of red sites from 0 to v; and similarly for each blue site. In particular, if the set of blue sites is surrounded by red sites, then blue will not longer reproduce (and vice versa). We refer to [55, 80] for the construction of the edge representation of the asymmetrical two-type Richardson model.

In the following subsection, we will see that the construction of Richardson models with first passage percolation also gives insights on FPP.

4.4.2 Geodesics

If a finite path γ between x and y satisfies $T_F(\gamma) = T_F(v, w)$ (it achieves the minimum) we thus say that γ is a **geodesic** from x to y.

Wierman and Reh proved the existence of geodesics in full generality for $\mathbb{Z}^2$ [168, Corollary 1.3]. In higher dimensions, we require assumptions over the passage time distribution F. We refer the reader to [17, Section 4.1] for general conditions, but the next theorem is sufficient for us.

Theorem 4.4.3 (Kesten [96, (9.23)]). *For a continuous distribution F, there exists a unique geodesic between any two points of $\mathbb{Z}^d$ almost surely.*

From now on, we assume that the distribution F is continuous. We remark that a crucial point for existence in Theorem 4.4.3 is that $F(0) = 0$.

Note that the exponential distribution Exp satisfies this condition. Let $G(x,y)$ be the unique geodesic between x and y, and let

$$\mathcal{T}_F = \bigcup_{z\in\mathbb{Z}} \{G(0,z)\}$$

be the **tree of infection** rooted at 0, where F indicates the distribution of the passage times. Since the geodesics $G(0,x)$ are unique, then $\mathcal{T}_F$ is a tree.

An infinite self-avoiding path $[v_1, v_2, \ldots]$ is a **geodesic ray** if $[v_1, \ldots, v_k]$ is a geodesic from v_1 to v_k for every $k \geq 1$. The number of geodesics rays is equal to the number of (topological) ends of the tree of infection $\mathcal{T}_F$. With this equivalence in mind, we denote the **number of geodesic rays** by $K(\mathcal{T}_F)$ A simple compactness argument shows the existence of a sub-sequential limit for $(G(0,(n,0,\ldots,0)))_{n\in\mathbb{N}}$, and then $K(\mathcal{T}_F) \geq 1$.

Newman conjectured in [141] that for a continuous distribution F (with additional hypothesis over distribution tails and exponential moments),

$$|K(\mathcal{T}_F)| = \infty, \qquad \text{almost surely.} \tag{4.14}$$

This question is still open, and (4.14) for continuous distributions is stated as Question 27 in [17]. Progress on (4.14) has been close to the two-type Richardson model for growth-cell. We defined this model in Subsection 4.3.2 and discussed its relation to FPP in Subsection 4.4.1. Now we present the consequences of the Theorems in Subsection 4.3.2 for $K(\mathcal{T}_F)$.

Häggström and Pemantle observed that the coexistence event for the two-type Richardson model (defined in (4.11)) implies the existence of two disjoint geodesic rays starting at 0. This property follows from a compactness argument. They showed that Theorem 4.3.1 implies $K(\mathcal{T}_{\text{Exp}}) \geq 2$ with positive probability on $\mathbb{Z}^2$ [79, Theorem 1.1]. Similarly, independent work of Garet and Marchand [72] and Hoffman [83] for the two-type Richardson model implied for FPP on $\mathbb{Z}^d$, $d \geq 2$, that

$$P\left(K(\mathcal{T}_F) \geq 2\right) > 0.$$

This positive probability holds for a large family of distributions F, in particular for the exponential distribution.

Hoffman proved in [84] a more general result for $K(\mathcal{T}_F)$. Hoffman's theorem is in terms of the geometry of the limit ball $\mathcal{B}_F$ appearing in the shape theorem (Theorem 4.4.2). The distributions considered in [84] constitute a large family. Here, for simplicity, we state the result for continuous distributions F with finite $2+\alpha$ moment.

If $\mathcal{B}_F$ is a polygon, we define Sides_F as the number of sides of $\partial\mathcal{B}_F$. If $\mathcal{B}_F$ is not a polygon, then Sides_F is equal to infinity. By symmetry, $\mathrm{Sides}_F \geq 4$ in $\mathbb{Z}^2$. Theorem 4.4.4 is a consequence of a version of Theorem 4.3.4 with distribution F for the passage times.

Theorem 4.4.4 (Hoffman [84, Theorem 1.2 and Theorem 1.4])**.** *Let F be a continuous distribution with $\mathbb{E}\left(\tau_e^{2+\alpha}\right) < \infty$ for some $\alpha > 0$. For any $\varepsilon > 0$ and $k \leq \mathrm{Sides}_F$ there exist $x_1, \ldots, x_k$ such that*

$$P\left(\textit{There exist } k \textit{ disjoint geodesics starting at } x_1, \ldots, x_k\right) > 1 - \varepsilon.$$

Moreover, if $k \leq \mathrm{Sides}_F/2$,

$$P(K(\mathcal{T}_F) \geq k) = 1.$$

A consequence of Theorem 4.4.4 is the existence of infinite geodesics when the limit shape $\mathcal{B}_F$ is not polygonal. This has been proved by Auffinger and Damron for measures ν (for the distribution of the passage time τ_e) which satisfy

$$\mathrm{supp}(\nu) \subset [1, \infty), \text{ and } \nu(\{1\}) = p \geq \vec{p}_c, \tag{4.15}$$

where $\vec{p_c}$ is the critical parameter for oriented percolation on $\mathbb{Z}^2$.

Theorem 4.4.5 (Auffinger-Damron [16, Theorem 2.3])**.** *Consider FPP on $\mathbb{Z}^2$ and let ν be the law for the passage times τ_e. If ν is a measure satisfying* (4.15) *with $p \geq \vec{p}_c$ on $\mathbb{Z}^2$, then*

$$P(K(\mathcal{T}_\nu) = \infty) = 1.$$

We have used that the coexistence event on the Richardson model implies the existence of geodesic rays. Nevertheless, the equivalence between coexistence, for a k-type Richardson model, and existence, of k ends for the infection tree, is not immediate. Such equivalence corresponds to a result announced recently. Recall that $E(x_1, \ldots, x_k)$ was defined in (4.12) as the coexistence event. We cite the theorem for two dimensions, and remark that there is a d-dimensional version of Theorem 4.4.6 with additional hypothesis over distribution moments and the number Sides_F.

Theorem 4.4.6 (Ahlberg [2, Theorem 1])**.** *Let F be a continuous distribution with $\mathbb{E}\left(\tau_e^{2+\alpha}\right) < \infty$ for some $\alpha > 0$. For any $k \in \mathbb{N} \cup \{\infty\}$, and $\varepsilon > 0$:*

(i) If $P(E(x_1, \ldots, x_k)) > 0$ for some $x_1, \ldots, x_k \in \mathbb{Z}^2$, then

$$P(K(\mathcal{T}_F) \geq k) = 1.$$

(ii) If $P(K(\mathcal{T}_F) \geq k) > 0$, then

$$P(E(x_1, \ldots, x_k)) > 1 - \varepsilon$$

for some $x_1, \ldots, x_k \in \mathbb{Z}^2$.

4.5 Chase-escape process

Chase-escape is an interacting particle system on a connected graph G. It emulates the behaviour of predators chasing prey that are escaping from them. On the graph $G = (V, E)$ the local state space is $\sigma = \{w, r, b\}$. We refer to states r and b as red and blue, respectively. The state w is white or vacant. The transitions on chase-escape are

$$rw \xrightarrow{\lambda} rr, \qquad br \xrightarrow{1} bb. \tag{4.16}$$

This means that red particles spread to adjacent uncoloured sites according to a Poisson process with rate λ. Meanwhile, blue particles overtake adjacent red particles at rate 1. For a formal definition we follow [35]. We define the maps $R^x : \sigma^V \to \sigma^V$ and $B^x : \sigma^V \to \sigma^V$ by

$$R^x(\eta)(z) = \begin{cases} r & \text{if } z = x \text{ and } \eta(x) = w, \\ \eta(z) & \text{otherwise,} \end{cases} \qquad \eta \in \sigma^{\mathbb{Z}^d},$$

$$B^x(\eta)(z) = \begin{cases} b & \text{if } z = x \text{ and } \eta(x) = r, \\ \eta(z) & \text{otherwise,} \end{cases} \qquad \eta \in \sigma^{\mathbb{Z}^d}.$$

The **chase-escape model** with parameter λ is the Markov process taking values in $\{w, r, b\}^{\mathbb{Z}^d}$ with transition rates

$$c(\eta, R^x(\eta)) = \lambda\, \mathbf{1}\{\eta(x) = w\} \sum_{\substack{y \in \mathbb{Z}^d \\ |x-y|=1}} \mathbf{1}\{\eta(y) = r\},$$

$$c(\eta, B^x(\eta)) = \mathbf{1}\{\eta(x) = r\} \sum_{\substack{y \in \mathbb{Z}^d \\ |x-y|=1}} \mathbf{1}\{\eta(y) = b\}.$$

Alternative interpretations of the chase-escape process include the spread of a rumour or an infection. In that case, the white particles represent the susceptible population (S), the red particles represent infected individuals (I) and the blue particles are recovered ones (R). Under the chase-escape dynamics, an individual can recover from an infection when it is in contact with someone already in state R. In terms of a rumour, this recovery corresponds to the spread of facts. The properties of the graph G correspond to the social network of the population.

In the context of tree, chase escape is equivalent to the escape process and the rumour-scotching process. Escape and chase-escape were both proposed by Kordzakhia [104]. They are similar except that, in the **escape process**, blue particles are allowed to conquer both vacant and red sites. A second variant, introduced by Bordenave, is the **rumour-scotching process**. We may think of the rumour-scotching process as a directed version of chase-escape. In this case, blue particles only conquer red sites, but such spread is performed only through edges that have spread red particles at some previous time. This restriction for the blue particles is associated with usual social dynamics for the transmission of a rumour. Under this model, a red site represents an individual believing a rumour, and the blue sites represent individuals who have additional information that denies such belief. A white site stands for susceptible individuals who have not heard anything. The red sites are prone to gossip, and they spread the rumour to their nearest-neighbours. When a site turns blue, they want to scotch the rumour that they propagated. However, they only turn to those with whom they shared the rumour before. Some results for chase-escape on trees were proved first in the context of the escape or the rumour-scotching processes.

4.5.1 Phase transitions

We distinguish two phases on the chase-escape process, depending on the cardinality of the set of sites eventually occupied. In the first scenario, the predator consumes all prey, and then the predator cannot move any further. On the contrary, if the prey advances fast enough, then there is prey alive at any time in the process before it ends. (We specify below the end of the process for finite graphs.) It is intuitively clear that these two phases depend on the "speed" of the red particles relative to the "speed" of the blue particles. The parameter λ controls such "speed". Recall the assumption in (4.16) that the transition rate for blue particles is 1. This choice is just a normalization of the parameters. We denote by P_λ the probability measure of chase-escape with parameter λ for the spread rate of red particles. The phase transitions and the existence of a critical parameter λ have been analysed for d-ary trees, and Galton-Watson trees, and complete graphs. Let us overview the main results for these graphs.

Finite graphs

Let K_n be the complete graph on n vertices. At time $t = 0$, there is one blue site, one red site, and $n - 2$ vacant sites. We add two absorbing states to the dynamics. The chase-escape process on K_n, with $n \geq 3$, finishes if

(a) there are no red particles left; or

(b) there are no vacant sites left i.e. all sites are occupied by a blue or a red particle.

Let $\mathcal{A}_n$ be the set of vertices of K_n that are red, at any time, before the chase-escape process stops. Let

$$A_n = \{|\mathcal{A}| = n - 1\}$$

be the **coexistence event**, where all vacant sites of K_n were coloured red at some finite time (with the addition of the initial red site). Note that the event A_n is equivalent to the absorbing state (b). We call $P_\lambda(A_n)$ the coexistence probability.

Theorem 4.5.1 (Kortchemski [105, Theorem 1])**.** *Let $(K_n)_{n\geq 3}$ be a growing sequence of complete graphs. Then*

$$\lim_{n\to\infty} P_\lambda(A_n) = \begin{cases} 0 & \textit{if } \lambda < 1, \\ \frac{1}{2} & \textit{if } \lambda = 1, \\ 1 & \textit{if } \lambda > 1. \end{cases}$$

In particular, $\lambda = 1$ is the critical parameter.

We see that on the complete graph, the coexistence probability is asymptotically positive if $\lambda \geq 1$. Otherwise, the coexistence probability is asymptotically 0. The first scenario corresponds to an **coexistence phase**, while the second is the **extinction phase** on the finite graph. The transition between these two phases occurs at the critical parameter $\lambda = 1$.

Infinite graphs

For an infinite graph G with root ρ_G, we assume that the process evolves on a modified version of G. We add an additional vertex $\hat{\rho}_G$ attached to ρ_G. Then the initial conditions of chase-escape at time $t = 0$ are

$$X_0(\hat{\rho}_G) = b, \qquad X_0(\rho_G) = r$$

and the rest of the vertices are on state w, i.e. these sites are vacant.

Let $\mathcal{B}$ denote the set of sites that are blue at some time in the process, and let

$$B = \{|\mathcal{B}| = \infty\}.$$

In the **coexistence phase** there is a positive probability that red particles occupy infinitely many sites so $P_\lambda(B) > 0$. We define the **extinction phase** as $P_\lambda(B) = 0$, in the case both types occupy only finitely many sites almost surely. We define the critical parameter

$$\lambda_c(G) := \inf\{\lambda : \ P_\lambda(B) > 0\}. \tag{4.17}$$

With this notation, we emphasize the dependence of the critical parameter on the underlying graph G.

We consider first the case of chase-escape on the ray $\mathbb{N}$. In this case, the first-guess answer is the correct one.

Proposition 4.5.2 ([62, Proposition 2.1])**.** *For chase-escape on* $\mathbb{N}$

$$\lambda_c(\mathbb{N}) = 1, \tag{4.18}$$

and chase-escape is in the extinction phase at criticality.

A tree is a natural generalization from $\mathbb{N}$, since we may consider the tree is the union of an infinite number of rays. A d-ary tree $\mathbb{T}_d$ is a rooted infinite tree were all vertices have d children.

Theorem 4.5.3 (Kordzakhia [104, Theorem 1])**.** *There exists a critical value*

$$\lambda_c(\mathbb{T}_d) = 2d - 1 - 2\sqrt{d^2 - d} \tag{4.19}$$

such that, for chase-escape on a d-ary tree with parameter λ,

(i) the process is in the extinction phase if $0 < \lambda < \lambda_c(\mathbb{T}_d)$*, and*

(ii) coexistence occurs if $\lambda > \lambda_c(\mathbb{T}_d)$*.*

Note that

$$\lambda_c \sim \frac{1}{4d} \text{ as } d \uparrow \infty.$$

Bordenave extended the previous result and determined the behaviour at criticality.

Proposition 4.5.4 (Bordenave [35, Corollary 1.5])**.** *Extinction happens for chase-escape on a d-ary tree at the critical parameter* $\lambda_c(\mathbb{T}_d)$*.*

Durrett, Junge and Tang presented in [62] a simple probabilistic argument for Theorem 4.5.3 that included the behaviour at criticality of Proposition 4.5.4. The base of the proof in [62] is a comparison between chase-escape on trees and $\mathbb{N}$. The arguments in [35] are analytical, but they apply to a more general setting: Galton-Watson trees with an additional assumption over the growth rate. The assumptions over the growth rate on [35] follow, almost surely, from conditioning T on being infinite. We state the following theorem under this hypothesis.

Theorem 4.5.5 (Bordenave [35, Theorem 1.1., Corollary 1.5])**.** *Let* $\lambda_c(\mathbb{T}_d)$ *be as in* (4.19)*. Let* T *be a realization of a Galton-Watson tree with mean number of offsprings* $d > 1$ *and conditioned to be infinite. The following holds* T*-almost surely for chase-escape on* T *with parameter* λ*.*

(i) If $\lambda \leq \lambda_c(\mathbb{T}_d)$ *then we have extinction.*

(ii) If $\lambda > \lambda_c(\mathbb{T}_d)$ *, then coexistence occurs.*

Kortchemski extended the study of chase-escape on Galton-Watson trees. In [106], he introduced a coupling on Galton-Walton trees of the chase-escape dynamics and branching random walks killed at 0 [106, Theorem 1]. With this method, Kortchemski obtained a shorter probabilistic proof of [35, Corollary 1.5] for super-critical Galton-Watson trees [106, Proposition 3] and asymptotics on the tail of the distribution of the number of prey [106, Theorem 4].

The examples above have a simple geometric structure. The complete graph represents the mean-field behaviour, while all the paths on trees are self-avoiding. Studying chase-escape on other lattices has been a challenging problem without significant advancements.

Kordzakhia and Martin have conjectured that coexistence is possible for

$$\lambda_c(G) < 1 \tag{4.20}$$

on two-dimensional lattices. Durrett, Junge, and Tang proved (4.20) for a variation of chase-escape on high-dimensional oriented lattices [62]. The variation considered in [62] allowed for infinite passage times for the red particles. Blue is allowed to cross that edge if the opposite vertex, on that edge, was reached by red already. To our knowledge, a similar proposition has not been proved for a non-oriented graph. The main challenge in the analysis of two-dimensional lattices is the presence of cycles. Some simple graphs have been candidates for satisfying (4.20). This was the case of the graph $\mathbb{N} \times [0, 1]$. Durrett, Junge, and Tang discarded this graph as an example for (4.20) by proving that $\lambda_c(\mathbb{N} \times [0, 1]) = 1$.

Tang, Kordzakhia, and Lalley have obtained simulations in support of (4.20). Their simulations show that $\lambda_c(\mathbb{Z}^2) \approx \frac{1}{2}$ and $\lambda_c(\Lambda) < 1$ for different two-dimensional lattices Λ, including the hexagon, triangle and 8-directional lattices. Furthermore, the simulations in [164] show fractal behaviour for the shape of occupied sites in all these latices when the parameter λ is close to its corresponding critical value. This property is typical in critical phenomena and gives further support to the following conjecture.

Conjecture 4.5.6 (Kordzakhia-Martin)**.** *The critical parameter for chase-escape on the square lattice is* $\lambda_c\left(\mathbb{Z}^2\right) = \frac{1}{2}$.

Part II

Scaling Limits of Uniform Spanning Trees

Chapter 5

Scaling Limit of the Three-Dimensional Uniform Spanning Tree and the Associated Random Walk[1]

Summary of this chapter

We show that the law of the three-dimensional uniform spanning tree (UST) is tight under rescaling in a space whose elements are measured, rooted real trees, continuously embedded into Euclidean space. We also establish that the relevant laws actually converge along a particular scaling sequence. The techniques that we use to establish these results are further applied to obtain various properties of the intrinsic metric and measure of any limiting space, including showing that the Hausdorff dimension of such is given by $3/\beta$, where $\beta \approx 1.624\ldots$ is the growth exponent of three-dimensional

[1] Joint work with Omer Angel, David Croydon, and Daisuke Shiraishi.

Acknowledgements. DC would like to acknowledge the support of a JSPS Grant-in-Aid for Research Activity Start-up, 18H05832 and a JSPS Grant-in-Aid for Scientific Research (C), 19K03540. DS is supported by a JSPS Grant-in-Aid for Early-Career Scientists, 18K13425.

loop-erased random walk. Additionally, we study the random walk on the three-dimensional uniform spanning tree, deriving its walk dimension (with respect to both the intrinsic and Euclidean metric) and its spectral dimension, demonstrating the tightness of its annealed law under rescaling, and deducing heat kernel estimates for any diffusion that arises as a scaling limit.

5.1 Introduction

Remarkable progress has been made in understanding the scaling limits of two-dimensional statistical mechanics models in recent years, much of which has depended in a fundamental way on the asymptotic conformal invariance of the models in question that has allowed many powerful tools from complex analysis to be harnessed. See [122, 149, 160] for some of the seminal works in this area, and [115] for more details. By contrast, no similar foothold for studying analogous problems in the (physically most relevant) case of three dimensions has yet been established. It seems that there is currently little prospect of progress for the corresponding models in this dimension.

Nonetheless, in [107], Kozma made the significant step of establishing the existence of a (subsequential) scaling limit for the trace of a three-dimensional loop-erased random walk (LERW). Moreover, in work that builds substantially on this, the time parametrisation of the LERW has been incorporated into the picture, with it being demonstrated that (again subsequentially) the three-dimensional LERW converges as a stochastic process, see [127] and the related articles [128, 155]. The aim of this work is to apply the latter results in conjunction with the fundamental connection between uniform spanning trees (USTs) and LERWs – specifically that paths between points in USTs are precisely LERWs [143, 169] – to determine the scaling behaviour of the three-dimensional UST (see Figure 5.1) and the associated random walk.

Before stating our results, let us introduce some of our notation. We follow closely the presentation of [24], where similar results were obtained in the two-dimensional case. Henceforth, we will write $\mathcal{U}$ for the UST on $\mathbb{Z}^3$, and $\mathbf{P}$ the probability measure on the probability space on which this is built (the corresponding expectation will be denoted $\mathbf{E}$). We refer the reader to [143] for Pemantle's construction of $\mathcal{U}$ in terms of a local limit of the USTs on the finite boxes $[-n,n]^3 \cap \mathbb{Z}^3$ (equipped with nearest-neighbour bonds) as $n \to \infty$, and proof of the fact that the resulting graph is indeed a spanning tree of $\mathbb{Z}^3$. We will denote by $d_{\mathcal{U}}$ the intrinsic (shortest path) metric on the graph $\mathcal{U}$, and $\mu_{\mathcal{U}}$ the counting measure on $\mathcal{U}$ (i.e., the measure which places a unit mass at each vertex). Similarly to [24], in describing a scaling limit for $\mathcal{U}$, we will view $\mathcal{U}$ as a measured, rooted spatial tree. In particular, in addition to the metric measure space $(\mathcal{U}, d_{\mathcal{U}}, \mu_{\mathcal{U}})$, we will also consider the embedding $\phi_{\mathcal{U}} : \mathcal{U} \to \mathbb{R}^3$, which we take to be simply the identity on vertices; this will allow us to retain information about $\mathcal{U}$ in the Euclidean topology. Moreover, it will be convenient to suppose the space $(\mathcal{U}, d_{\mathcal{U}})$ is rooted at the origin of $\mathbb{R}^3$, which we will write as $\rho_{\mathcal{U}}$. To fit the framework of [24], we extend $(\mathcal{U}, d_{\mathcal{U}})$ by adding unit line segments along edges, and linearly interpolate $\phi_{\mathcal{U}}$ between vertices.

5.1.1 Scaling limits of the three-dimensional UST

We have defined a random quintuplet $(\mathcal{U}, d_{\mathcal{U}}, \mu_{\mathcal{U}}, \phi_{\mathcal{U}}, \rho_{\mathcal{U}})$. Our main result (Theorem 5.1.1 below) is the existence of a certain subsequential scaling limit for this object in an appropriate Gromov-Hausdorff-type topology, the

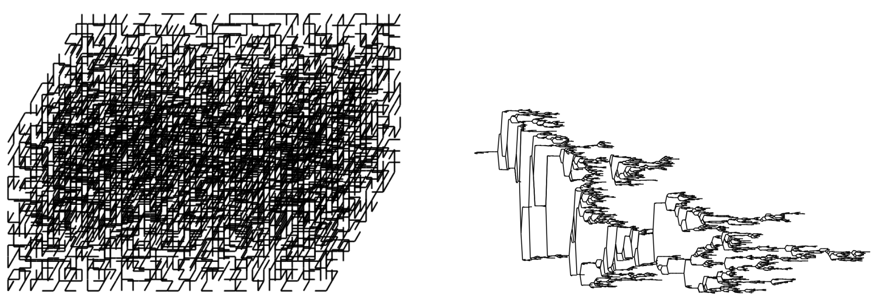

Figure 5.1: A realisation of the UST in a three-dimensional box, as embedded into $\mathbb{R}^3$ (left), and drawn as a planar graph tree (right). Source code adapted from two-dimensional version of Mike Bostock.

precise definition of which we postpone to Section 5.2. Moreover, the result incorporates the statement that the laws of the rescaled objects are tight even without taking the subsequence. One further quantity needed to state the result precisely is the growth exponent of the three-dimensional LERW. Let M_n be the number of steps of the LERW on $\mathbb{Z}^3$ until its first exit from a ball of radius n. The **growth exponent** is defined by the limit:

$$\beta := \lim_{n\to\infty} \frac{\log \mathbf{E} M_n}{\log n},$$

(equivalently, $\mathbf{E}M_n = n^{\beta+o(1)}$). The existence of this limit was proved in [155]. Whilst the exact value of β is not known, rigourously proved bounds are $\beta \in (1, \frac{5}{3}]$, see [114]. Numerical estimates suggest that $\beta = 1.624\dots$, see [170]. We remark that in two dimensions the corresponding exponent is 5/4, first proved by Kenyon [95], and in dimension 4 or more its value is 2. In three dimensions there is no conjecture for an exact value of β.

The exponent β determines the scaling of $d_{\mathcal{U}}$. Specifically, let $\mathbf{P}_\delta$ be the law of the measured, rooted spatial tree

$$\left(\mathcal{U}, \delta^\beta d_{\mathcal{U}}, \delta^3 \mu_{\mathcal{U}}, \delta\phi_{\mathcal{U}}, \rho_{\mathcal{U}}\right), \tag{5.1}$$

when $\mathcal{U}$ has law $\mathbf{P}$. For the rooted measured metric space $(\mathcal{U}, d_{\mathcal{U}}, \mu_{\mathcal{U}}, \rho)$ we consider the local Gromov-Hausdorff-Prohorov topology. This is extended with the locally uniform topology for the embedding $\phi_{\mathcal{U}}$. As a straightforward consequence of our tightness and scaling results with respect to this Gromov-Hausdorff-type topology, we also obtain the corresponding conclusions with respect to Schramm's path ensemble topology. The latter topology was introduced in [149] as an approach to taking scaling limits of two-dimensional spanning trees. Roughly speaking this topology observes the set of all macroscopic paths in an object, in the Hausdorff topology. See Section 5.2 for detailed definitions of these topologies.

Theorem 5.1.1. *The collection $(\mathbf{P}_\delta)_{\delta\in(0,1]}$ is tight with respect to the local Gromov-Hausdorff-Prohorov topology with locally uniform topology for the embedding, and with respect to the path ensemble topology. Moreover the* *limit of $\mathbf{P}_\delta$ exists as $\delta = 2^{-n} \to 0$ exists in both topologies.*

Remark. The reason we only state convergence along the subsequence (2^{-n}) stems from the fact that our argument fundamentally depends on Kozma's original work on the scaling of three-dimensional LERW, where a similar restriction was imposed [107]. There is no reason to believe that this is an essential requirement for the result to hold. (Indeed, Theorem 5.1.1 shows that subsequential limits exist) This is the only place in our proof where we require $\delta = 2^{-n}$. If one were to generalise Kozma's result to an arbitrary sequence of δs, the natural extension of the above theorem would immediately follow.

Remark. An important open problem, for both the LERW and UST in three dimensions, is to describe the limiting object directly in the continuum. In two dimensions, there are connections between the LERW and SLE_2, as well as between the UST and SLE_8, see [86, 122, 149], which give a direct construction of the continuous objects. In the three-dimensional case, there is as yet no parallel theory. The development of such a representation would be a significant advance in three-dimensional statistical mechanics.

Before continuing, we briefly outline the strategy of proof for the convergence part of the above result, for which there are two main elements. The first of these is a finite-dimensional convergence statement: Theorem 5.7.2 states that the part of $\mathcal{U}$ spanning a finite collection of points converges under rescaling. Appealing to Wilson's algorithm [169], which gives the means to construct $\mathcal{U}$ from LERW paths, this finite-dimensional result extends the scaling result for the three-dimensional LERW of [127]. Here we encounter a central hurdle: after the first walk, Wilson's algorithm requires us to take a LERW in an rough subdomain of $\mathbb{Z}^3$, namely the complement of the previous LERWs. Existing results in [107, 127] on scaling limits of LERWs require subdomains with smooth boundary, and some care is needed to extend the existence of the scaling limit. We resolve this difficulty by proving that we can approximate the rough subdomain with a simpler one, and showing the corresponding LERWs are close to each other as parametrized curves.

Secondly, to prove tightness, we need to check that the trees spanning a finite collection of points give a sufficiently good approximation of the entire UST $\mathcal{U}$, once the number of points is large. For this, we need to know that LERWs started from the remaining lattice points hit the trees spanning a finite collection of points quickly. In two dimensions, such a property was established using Beurling's estimate, which says that a simple random walk hits any given path quickly if it starts close to it in Euclidean terms, see [97]. In three dimensions, Beurling's estimate does not hold. In its place, we have a result from [148], which yields that a simple random walk hits a typical LERW path quickly if it starts close to it. Thus, although the intuition in the three-dimensional case is similar, it requires us to remember much more about the structure of the part of the UST we have already constructed as Wilson's algorithm proceeds.

5.1.2 Properties of the scaling limit

While uniqueness of the scaling limit is as yet unproved, the techniques we use to establish Theorem 5.1.1 allow us to deduce some properties of any possible scaling limit. These are collected below. NB. For the result, the scaling limits we consider are with respect to the Gromov-Hausdorff-type topology on the space of measured, rooted spatial trees, see Section 5.2 below. The one-endedness of the limiting space matches the corresponding result in the discrete case, [143, Theorem 4.3]. We use $B_{\mathcal{T}}(x, r)$ to denote the ball in the limiting metric space $\mathcal{T} = (\mathcal{T}, d_{\mathcal{T}})$ of radius r around x. It is natural to expect that the scaling limit will have dimension

$$d_f := \frac{3}{\beta}.$$

Moreover, one would expect that a ball of radius r in the limiting object has measure of order $r^{3/\beta}$. The following theorem establishes uniform bounds of this magnitude for all small balls in the limiting tree, with a logarithmic correction for arbitrary centres and with iterated logarithmic corrections for a fixed centre, which may be fixed to be ρ. We use $f \preceq g$ to denote that $f \leq Cg$ for some absolute (i.e. deterministic, and not depending on the particular subsequence) constant C. We denote by $\gamma_{\mathcal{T}}(x, y)$ the path in the topological tree $\mathcal{T}$ between points x and y. We write $\mathcal{L}$ to represent Lebesgue measure on $\mathbb{R}^3$. The definition of the 'Schramm distance' below is inspired by [149, Remark 10.15].

Theorem 5.1.2. *Let* $\hat{\mathbf{P}}$ *be a subsequential limit of* $\mathbf{P}_\delta$ *as* $\delta \to 0$, *and the random measured, rooted spatial tree* $(\mathcal{T}, d_{\mathcal{T}}, \mu_{\mathcal{T}}, \phi_{\mathcal{T}}, \rho_{\mathcal{T}})$ *have law* $\hat{\mathbf{P}}$. *Then the following statements hold* $\hat{\mathbf{P}}$*-a.s.*

(a) The tree $\mathcal{T}$ *is one-ended (with respect to the topology induced by the metric* $d_{\mathcal{T}}$*).*

(b) Every ball in $(\mathcal{T}, d_{\mathcal{T}})$ *has Hausdorff dimension* d_f.

(c) There exists an absolute constant $C < \infty$ *so that: for any* $R > 0$, *there exists a random* $r_0(\mathcal{T}) > 0$ *such that*

$$\begin{aligned} r^{d_f}(\log r^{-1})^{-C} &\preceq \inf_{x \in B_{\mathcal{T}}(\rho, R)} \mu_{\mathcal{T}}\left(B_{\mathcal{T}}(x, r)\right) \\ &\leq \sup_{x \in B_{\mathcal{T}}(\rho, R)} \mu_{\mathcal{T}}\left(B_{\mathcal{T}}(x, r)\right) \preceq r^{d_f}(\log r^{-1})^{C}, \end{aligned}$$

for all $r < r_0$.

(d) For some absolute $C < \infty$, *there exists a random* $r_0(\mathcal{T}) > 0$ *such that*

$$r^{d_f}(\log\log r^{-1})^{-C} \preceq \mu_{\mathcal{T}}\left(B_{\mathcal{T}}(\rho, r)\right) \preceq r^{d_f}(\log\log r^{-1})^{C}, \qquad \forall r < r_0.$$

(e) The metric $d_{\mathcal{T}}$ *is topologically equivalent to the 'Schramm metric'* $d_{\mathcal{T}}^S$ *on* $\mathcal{T}$, *defined by*

$$d_{\mathcal{T}}^S(x, y) := \operatorname{diam}\left(\phi_{\mathcal{T}}(\gamma_{\mathcal{T}}(x, y))\right), \tag{5.2}$$

where diam *is the diameter in the Euclidean metric.*

(f) $\mu_{\mathcal{T}} = \mathcal{L} \circ \phi_{\mathcal{T}}$.

Differences from the two-dimensional case

Analogues for the properties described in Theorem 5.1.2 (and others) were proved in the two-dimensional case in [24], see also the related earlier work [149]. There are, however, several notable differences in three dimensions. Following Schramm [149], consider the **trunk** of the tree $\mathcal{T}$, denoted $\mathcal{T}^\circ$, which is the set of all points of $\mathcal{T}$ of degree greater than 1, where the degree of x is the number of connected components of $\mathcal{T} \setminus \{x\}$. In the two-dimensional case, it is known that the restriction of the continuous map $\phi_\mathcal{T}$ to the trunk is a homeomorphism between $\mathcal{T}^\circ$ (equipped with the induced topology from $\mathcal{T}$) and its image $\phi_\mathcal{T}(\mathcal{T}^\circ)$ (equipped with the induced Euclidean topology). Thus the image of the trunk, which is dense in $\mathbb{R}^2$, determines its topology. We do not expect the same to be true in three-dimensions. Indeed, due to the greater probability that three LERWs started from adjacent points on the integer lattice escape to a macroscopic distance before colliding, we expect that the image of the trunk $\phi_\mathcal{T}(\mathcal{T}^\circ)$ is no longer a topological tree in $\mathbb{R}^3$, see Figure 5.2. We aim to establish this as a result in a forthcoming work.

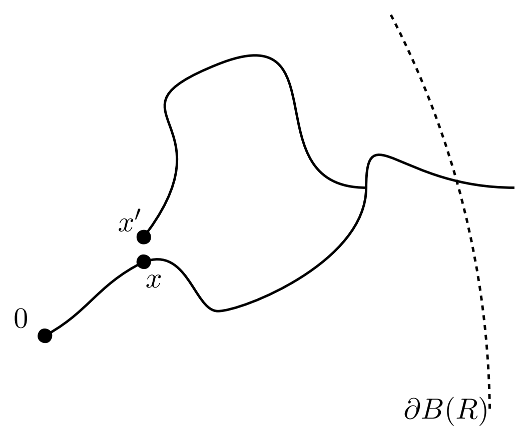

Figure 5.2: In the above sketch, $|x - x'| = 1$, but the path in the UST between these points has Euclidean diameter greater than $R/3$. We expect that such pairs of points occur with positive probability, uniformly in R.

Secondly, for the two-dimensional UST, it was shown in [24] that the maximal degree in $\mathcal{T}$ is 3, and that $\mu_{\mathcal{T}}$ is supported on the leaves of $\mathcal{T}$, i.e. the set of points of degree 1. We can show that the same is true in three dimensions, though we also postpone these results to a separate paper, since they are significantly harder than in two dimensional case. Indeed, as well as appealing to the homeomorphism between the trunk and its embedding, the two-dimensional arguments in the literature depend on a duality argument that does not extend to three dimensions. We replace this with a more technical direct argument. The aforementioned homeomorphism and duality also allow it to be shown that in two dimensions $\max_{x\in\mathbb{R}^3}|\phi^{-1}(x)| = 3$ (where we write $|A|$ to represent the cardinality of a set A), and, although not mentioned explicitly in [24, 149], it is also easy to deduce the Hausdorff dimension of the set of points with given pre-image size. Our forthcoming work will explore the corresponding results in the three dimensional case.

5.1.3 Scaling the random walk on $\mathcal{U}$

The metric-measure scaling of $\mathcal{U}$ yields various consequences for the associated simple random walk (SRW), which we next introduce. For a given realisation of the graph $\mathcal{U}$, the SRW on $\mathcal{U}$ is the discrete time Markov process $X^{\mathcal{U}} = ((X_n^{\mathcal{U}})_{n\geq 0}, (P_x^{\mathcal{U}})_{x\in\mathbb{Z}^3})$ which at each time step jumps from its current location to a uniformly chosen neighbour in $\mathcal{U}$. For $x \in \mathbb{Z}^3$, the law $P_x^{\mathcal{U}}$ is called the **quenched** law of the simple random walk on $\mathcal{U}$ started at x. We then define the **annealed** or **averaged** law for the process started from $\rho_{\mathcal{U}}$ as the semi-direct product of the environment law $\mathbf{P}$ and the quenched law $P_0^{\mathcal{U}}$ by setting

$$\mathbb{P}^{\mathcal{U}}(\cdot) := \int P_0^{\mathcal{U}}(\cdot)d\mathbf{P}.$$

We use $\mathbb{E}^{\mathcal{U}}$ for the corresponding annealed expectation.

The behaviour of the random walk on a graph is fundamentally linked to the associated electrical resistance. We refer the reader to [19, 59, 126, 132] for introductions to this connection, including the definition of effective resistance in particular. For the three-dimensional UST, we will write $R_{\mathcal{U}}$ for the effective resistance on $\mathcal{U}$, considered as an electrical network with unit resistors placed along each edge.

As noted above, the typical measure of $B_{\mathcal{U}}(\rho, R)$ is of order R^{d_f}. We show below that the effective resistance to the complement of the ball is typically of order R (it is trivially at most R). In light of these, and following [109], we define the set of well-behaved scales with parameter λ by

$$J(\lambda) := \left\{ R \in [1,\infty) : \begin{array}{c} R^{-d_f}\mu_{\mathcal{U}}\left(B_{\mathcal{U}}(\rho,R)\right) \in [\lambda^{-1},\lambda] \\ \text{and } R_{\mathcal{U}}\left(\rho, B_{\mathcal{U}}(\rho,R)^c\right) \geq \lambda^{-1}R \end{array} \right\}.$$

In particular, for R to be in $J(\lambda)$, we require good control over the volume of the intrinsic ball centred at the root of $\mathcal{U}$ of radius R, and control over the resistance from the root to the boundary of this ball. As our next result, we show that the these events hold with high probability, uniformly in R.

Theorem 5.1.3. *There exist constants $c, c_1, c_2 \in (0,\infty)$ such that: for all $R, \lambda > 1$,*

$$\mathbf{P}\left(R \in J(\lambda)\right) \geq 1 - ce^{-c_1\lambda^{c_2}}.$$

The motivation for Theorem 5.1.3 is provided by the general random walk estimates presented by Kumagai and Misumi in [109]. (Which builds on the work [23].) More specifically, Theorem 5.1.3 establishes the conditions for the main results of [109], which yield several important exponents governing aspects of the behaviour of the random walk. Indeed, as is made precise in the following corollary, we obtain that the walk dimension with respect to intrinsic distance is given by

$$d_w := 1 + d_f = \frac{3+\beta}{\beta},$$

the walk dimension with respect to extrinsic (Euclidean) distance d_E is given by $\beta d_w = 3+\beta$ (this requires a small amount of additional work to the tools of [109]), and the spectral dimension is given by

$$d_s := \frac{2d_f}{d_w} = \frac{6}{3+\beta}. \tag{5.3}$$

Various further consequences for the random walk on $\mathcal{U}$ also follow from the

	General form	$d=2$	$d=3$
LERW growth exponent	β	$5/4 = 1.25$	1.62
Fractal dimension of $\mathcal{U}$	$d_f = d/\beta$	$8/5 = 1.60$	1.85
Intrinsic walk dimension of $\mathcal{U}$	$d_w = 1 + d_f$	$13/5 = 2.60$	2.85
Extrinsic walk dimension of $\mathcal{U}$	βd_w	$13/4 = 3.25$	4.62
Spectral dimension of $\mathcal{U}$	$2d_f/d_w$	$16/13 = 1.23$	1.30

Table 5.1: Exponents associated with the LERW and UST in two and three dimensions. The two-dimensional exponents are known rigorously from [20, 21, 24, 95]. The three-dimensional values are based on the results of this study, together with the numerical estimate for the growth exponent of the three-dimensional LERW from [170].

results of [109], but rather than simply list these here, we refer the interested reader to that article for details. Table 5.1 summarises the numerical estimates for the three-dimensional random walk exponents that follow from the above formulae, together with the numerical estimate for β from [170], and compares these with the known exponents in the two-dimensional model.

Corollary 5.1.1. (a) For $\mathbf{P}$-a.e. realisation of $\mathcal{U}$ and all $x \in \mathcal{U}$,

$$\lim_{R\to\infty} \frac{\log E_x^{\mathcal{U}} \tau_{x,R}^{\mathcal{U}}}{\log R} = d_w, \tag{5.4}$$

where $\tau_{x,R}^{\mathcal{U}} := \inf\{n \geq 0 :\ d_{\mathcal{U}}(x, X_n^{\mathcal{U}}) > R\}$,

$$\lim_{R\to\infty} \frac{\log E_x^{\mathcal{U}} \tau_{x,R}^{E}}{\log R} = \beta d_w, \tag{5.5}$$

where $\tau_{x,R}^{E} := \inf\{n \geq 0 :\ d_E(x, X_n^{\mathcal{U}}) > R\}$, and

$$-\lim_{n\to\infty} \frac{2 \log p_{2n}^{\mathcal{U}}(x,x)}{\log n} = d_s. \tag{5.6}$$

(b) For $\mathbb{P}^{\mathcal{U}}$-a.e. realisation of $X^{\mathcal{U}}$,

$$\lim_{R\to\infty}\frac{\log\tau^{\mathcal{U}}_{0,R}}{\log R}=d_w,\qquad \lim_{n\to\infty}\frac{\log\max_{0\le m\le n}d_{\mathcal{U}}(0,X^{\mathcal{U}}_m)}{\log n}=\frac{1}{d_w},\tag{5.7}$$

$$\lim_{R\to\infty}\frac{\log\tau^{E}_{0,R}}{\log R}=\beta d_w,\qquad \lim_{n\to\infty}\frac{\log\max_{0\le m\le n}d_{E}(0,X^{\mathcal{U}}_m)}{\log n}=\frac{1}{\beta d_w}.\tag{5.8}$$

(c) It holds that

$$\lim_{R\to\infty}\frac{\log\mathbb{E}^{\mathcal{U}}\left(\tau^{\mathcal{U}}_{0,R}\right)}{\log R}=d_w,\tag{5.9}$$

$$\lim_{R\to\infty}\frac{\log\mathbb{E}^{\mathcal{U}}\left(\tau^{E}_{0,R}\right)}{\log R}=\beta d_w,\tag{5.10}$$

where $\mathbb{E}^{\mathcal{U}}$ is the expectation under $\mathbb{P}^{\mathcal{U}}$, and

$$-\lim_{n\to\infty}\frac{2\log\mathbf{E}\left(p^{\mathcal{U}}_{2n}(0,0)\right)}{\log n}=d_s.\tag{5.11}$$

Remark. In part (c) of the previous result, we do not provide averaged results for the distance travelled by the process up to time n with respect to either the intrinsic or extrinsic metrics. In the two-dimensional case, the corresponding results were established in [21], with the additional input being full off-diagonal annealed heat kernel estimates. Since the latter require a substantial amount of additional work, we leave deriving such as an open problem.

Finally, it is by now well-understood how scaling limits of discrete trees transfer to scaling limits for the associated random walks on the trees, see [14, 24, 50, 52–54]. We apply these techniques in our setting to deduce a (subsequential) scaling limit for $X^{\mathcal{U}}$. As we will explain in Section 5.10, the limiting process can be written as $(\phi_{\mathcal{T}}(X^{\mathcal{T}}_t))_{t\ge0}$, where $((X^{\mathcal{T}}_t)_{t\ge0},(P^{\mathcal{T}}_x)_{x\in\mathcal{T}})$ is the canonical Brownian on the limit space $(\mathcal{T},d_{\mathcal{T}},\mu_{\mathcal{T}})$. This Brownian motion is constructed in [13, 98]. Moreover, the volume estimates of Theorem 5.1.2, in conjunction with the general heat kernel estimates of [49], yield sub-diffusive transition density bounds for the limiting diffusion. Modulo the different exponents, these are of the same sub-Gaussian form as established for the Brownian continuum random tree in [51], and for the two-dimensional UST in [24]. Note in particular that our results imply that the spectral dimension of the continuous model, defined analogously to (5.6), is equal to the value d_s given at (5.3).

Theorem 5.1.4. *If $(\mathbf{P}_{\delta_n})_{n\geq 0}$ is a convergent sequence with limit $\hat{\mathbf{P}}$, then the following statements hold.*

(a) The annealed law of $(\phi_{\mathcal{T}}(X_t^{\mathcal{T}}))_{t\geq 0}$, where $X^{\mathcal{T}}$ is Brownian motion on $(\mathcal{T}, d_{\mathcal{T}}, \mu_{\mathcal{T}})$ started from $\rho_{\mathcal{T}}$, i.e.

$$\mathbb{P}^{\mathcal{T}}(\cdot) := \int P_{\rho_{\mathcal{T}}}^{\mathcal{T}} \circ \phi_{\mathcal{T}}^{-1}(\cdot) d\hat{\mathbf{P}},$$

is a well-defined probability measure on $C(\mathbb{R}_+, \mathbb{R}^3)$.

(b) Let $(X_t^{\mathcal{U}})_{t\geq 0}$ be the simple random walk on $\mathcal{U}$ started from $\rho_{\mathcal{U}}$, then the annealed laws of the rescaled processes

$$\left(\delta_n X^{\mathcal{U}}_{t\delta_n^{-(3+\beta)}}\right)_{t\geq 0}$$

converge to the annealed law of $(\phi_{\mathcal{T}}(X_t^{\mathcal{T}}))_{t\geq 0}$.

(c) $\hat{\mathbf{P}}$-a.s., the process $X^{\mathcal{T}}$ is recurrent and admits a jointly continuous transition density $(p_t^{\mathcal{T}}(x,y))_{x,y\in\mathcal{T},t>0}$. Moreover, it $\hat{\mathbf{P}}$-a.s. holds that, for any $R > 0$, there exist random constants $c_1(\mathcal{T}), c_2(\mathcal{T}), c_3(\mathcal{T}), c_4(\mathcal{T})$ and $t_0(\mathcal{T}) \in (0,\infty)$ and deterministic constants $\theta_1, \theta_2, \theta_3, \theta_4 \in (0,\infty)$ (not depending on R) such that

$$\begin{aligned} p_t^{\mathcal{T}}(x,y) \leq & c_1 t^{-d_s/2} \ell(t^{-1})^{\theta_1} \\ & \cdot \exp\left\{ -c_2 \left(\frac{d_{\mathcal{T}}(x,y)^{d_w}}{t} \right)^{\frac{1}{d_w - 1}} \ell(d_{\mathcal{T}}(x,y)/t)^{-\theta_2} \right\}, \end{aligned}$$

$$p_t^{\mathcal{T}}(x,y) \geq c_3 t^{-d_s/2}\ell(t^{-1})^{-\theta_3} \cdot \exp\left\{-c_4\left(\frac{d_{\mathcal{T}}(x,y)^{d_w}}{t}\right)^{\frac{1}{d_w-1}}\ell(d_{\mathcal{T}}(x,y)/t)^{\theta_4}\right\},$$

for all $x,y \in B_{\mathcal{T}}(\rho_{\mathcal{T}},R)$, $t \in (0,t_0)$, *where* $\ell(x) := 1 \vee \log x$.

(d) *(i)* $\hat{\mathbf{P}}$*-a.s., there exists a random* $t_0(\mathcal{T}) \in (0,\infty)$ *and deterministic* $c_1, c_2, \theta_1, \theta_2 \in (0,\infty)$ *such that*

$$c_1 t^{-d_s/2}(\log\log t^{-1})^{-\theta_1} \leq p_t^{\mathcal{T}}(\rho_{\mathcal{T}},\rho_{\mathcal{T}}) \leq c_2 t^{-d_s/2}(\log\log t^{-1})^{\theta_2},$$

for all $t \in (0,t_0)$.

(i) There exist constants $c_1, c_2 \in (0,\infty)$ *such that*

$$c_1 t^{-d_s/2} \leq \hat{\mathbf{E}} p_t^{\mathcal{T}}(\rho_{\mathcal{T}},\rho_{\mathcal{T}}) \leq c_2 t^{-d_s/2},$$

for all $t \in (0,1)$.

Organization of this chapter

The remainder of the chapter is organised as follows. In Section 5.2, we introduce the topologies that provide the framework for Theorem 5.1.1, and set out three conditions that imply tightness in this topology. Then, in Section 5.3, we collect together the properties of loop-erased random walks that will be useful for this article. After these preparations, the three tightness conditions are checked in Section 5.4, and the volume estimates contained within this are strengthened in Sections 5.5 and 5.6 in a way that yields more detailed properties concerning the limit space and simple random walk. In Section 5.7, we demonstrate our finite-dimensional convergence result for subtrees of $\mathcal{U}$ that span a finite number of points. The various pieces for proving Theorem 5.1.1 are subsequently put together in Section 5.8, and the properties of the limiting space are explored in Section 5.9, with Theorem 5.1.2 being proved in this part of the article. Finally, Section 5.10 covers the results relating to the simple random walk and its diffusion scaling limit.

5.2 Topological framework

In this section, we introduce the Gromov-Hausdorff-type topology on measured, rooted spatial trees with respect to which Theorem 5.1.1 is stated. This topology is metrizable, and for completeness sake we include a possible metric (see Proposition 5.2.1). Moreover, we provide a sufficient criterion (Assumptions 1,2, and 3 below) for tightness of a family of measures on measured, rooted spatial trees in the relevant topology (see Lemma 5.2.2). This will be applied in order to prove tightness under scaling of the three-dimensional UST. In the first part of the section, we follow closely the presentation of [24].

Define $\mathbb{T}$ to be the collection of quintuplets of the form

$$\underline{\mathcal{T}} = (\mathcal{T}, d_{\mathcal{T}}, \mu_{\mathcal{T}}, \phi_{\mathcal{T}}, \rho_{\mathcal{T}}),$$

where: $(\mathcal{T}, d_{\mathcal{T}})$ is a complete and locally compact real tree (for the definition of a real tree, see [125, Definition 1.1], for example); $\mu_{\mathcal{T}}$ is a locally finite Borel measure on $(\mathcal{T}, d_{\mathcal{T}})$; $\phi_{\mathcal{T}}$ is a continuous map from $(\mathcal{T}, d_{\mathcal{T}})$ into a separable metric space (M, d_M); and $\rho_{\mathcal{T}}$ is a distinguished vertex in $\mathcal{T}$. (In this article, the image space (M, d_M) we consider is $\mathbb{R}^3$ equipped with the Euclidean distance.) We call such a quintuplet a **measured, rooted, spatial tree**. We will say that two elements of $\mathbb{T}$, $\underline{\mathcal{T}}$ and $\underline{\mathcal{T}}'$ say, are equivalent if there exists an isometry $\pi : (\mathcal{T}, d_{\mathcal{T}}) \to (\mathcal{T}', d'_{\mathcal{T}})$ for which $\mu_{\mathcal{T}} \circ \pi^{-1} = \mu'_{\mathcal{T}}$, $\phi_{\mathcal{T}} = \phi'_{\mathcal{T}} \circ \pi$ and also $\pi(\rho_{\mathcal{T}}) = \rho'_{\mathcal{T}}$.

We now introduce a variation on the Gromov-Hausdorff-Prohorov topology on $\mathbb{T}$ that also takes into account the mapping $\phi_{\mathcal{T}}$. In order to introduce this topology, we start by recalling from [24] the metric Δ_c on $\mathbb{T}_c$, which is the subset of elements of $\mathbb{T}$ such that $(\mathcal{T}, d_{\mathcal{T}})$ is compact. In particular, for two elements of $\mathbb{T}_c$, we set $\Delta_c\left(\underline{\mathcal{T}}, \underline{\mathcal{T}}'\right)$ to be equal to

$$\inf_{\substack{Z,\psi,\psi',\mathcal{C}:\\ (\rho_{\mathcal{T}},\rho'_{\mathcal{T}})\in\mathcal{C}}} \left\{ \begin{array}{c} d_P^Z\left(\mu_{\mathcal{T}} \circ \psi^{-1}, \mu'_{\mathcal{T}} \circ \psi'^{-1}\right) + \\ \sup_{(x,x')\in\mathcal{C}} \left(d_Z\left(\psi(x), \psi'(x')\right) + d_M\left(\phi_{\mathcal{T}}(x), \phi'_{\mathcal{T}}(x')\right)\right) \end{array} \right\} \tag{5.12}$$

where the infimum is taken over all metric spaces $Z = (Z, d_Z)$, isometric embeddings $\psi : (\mathcal{T}, d_{\mathcal{T}}) \to Z$, $\psi' : (\mathcal{T}', d'_{\mathcal{T}}) \to Z$, and correspondences $\mathcal{C}$ between $\mathcal{T}$ and $\mathcal{T}'$, and we define d_P^Z to be the Prohorov distance between finite Borel measures on Z. Note that, by a correspondence $\mathcal{C}$ between $\mathcal{T}$ and $\mathcal{T}'$, we mean a subset of $\mathcal{T} \times \mathcal{T}'$ such that for every $x \in \mathcal{T}$ there exists at least one $x' \in \mathcal{T}'$ such that $(x, x') \in \mathcal{C}$ and conversely for every $x' \in \mathcal{T}'$ there exists at least one $x \in \mathcal{T}$ such that $(x, x') \in \mathcal{C}$. (Except for the term involving ϕ and ϕ', this is the usual metric for the Gromov-Hausdorff-Prohorov topology.)

Given the definition of Δ_c at (5.12), we then define a pseudo-metric Δ on $\mathbb{T}$ by setting

$$\Delta\left(\underline{\mathcal{T}}, \underline{\mathcal{T}}'\right) := \int_0^\infty e^{-r}\left(1 \wedge \Delta_c\left(\underline{\mathcal{T}}^{(r)}, \underline{\mathcal{T}}'^{(r)}\right)\right) dr, \tag{5.13}$$

where $\underline{\mathcal{T}}^{(r)}$ is obtained by taking the closed ball in $(\mathcal{T}, d_{\mathcal{T}})$ of radius r centred at $\rho_{\mathcal{T}}$, restricting $d_{\mathcal{T}}$, $\mu_{\mathcal{T}}$ and $\phi_{\mathcal{T}}$ to $\mathcal{T}^{(r)}$, and taking $\rho_{\mathcal{T}}^{(r)}$ to be equal to $\rho_{\mathcal{T}}$. We have the following result, and it is the corresponding topology that provides the framework for Theorem 5.1.1.

Proposition 5.2.1 ([24, Proposition 3.4])**.** *The function Δ defines a metric on the equivalence classes of $\mathbb{T}$. Moreover, the resulting metric space is separable.*

We next present a criterion for tightness of a sequence of random measured, rooted spatial trees. This is a probabilistic version of [24, Lemma 3.5] (which adds the spatial embedding to the result of [1, Theorem 2.11]) Recall the definition of **stochastic equicontinuity**: Suppose for some index set $\mathcal{A}$ there are random metric spaces (X_i, d_i) and random functions $\phi_i : X_i \to M$ for a metric space (M, d_M). The functions are stochastically equicontinuous if their moduli of continuity converge to 0 uniformly in probability, i.e. for every $\varepsilon > 0$,

$$\lim_{\eta \to 0} \sup_{i \in \mathcal{A}} \mathbf{P}\left(\sup_{\substack{x,y \in X_i: \\ d_i(x,y) \le \eta}} d_M\left(\phi_i(x), \phi_i(y)\right) > \varepsilon \right) = 0.$$

Lemma 5.2.2. *Suppose* (M, d_M) *is proper (i.e. every closed ball in* M *is compact), and* ρ_M *is a fixed point in* M*. Let* $\underline{\mathcal{T}}_\delta = (\mathcal{T}_\delta, d_{\mathcal{T}_\delta}, \mu_{\mathcal{T}_\delta}, \phi_{\mathcal{T}_\delta}, \rho_{\mathcal{T}_\delta})$*,* $\delta \in \mathcal{A}$ *(where* $\mathcal{A}$ *is some index set), be a collection of random measured, rooted spatial trees. Moreover, assume that for every* $R > 0$*, the following quantities are tight:*

(i) For every $\varepsilon > 0$*, the number* $N(\underline{\mathcal{T}}_\delta, R, \varepsilon)$ *of balls of radius* ε *required to cover the ball* $\mathcal{T}_\delta^{(R)}$*,*

(ii) The measure of the ball: $\mu_{\mathcal{T}_\delta}\left(\mathcal{T}_\delta^{(R)}\right)$*;*

(iii) The distances $d_M(\rho_M, \phi_{\mathcal{T}_\delta}(\rho_{\mathcal{T}_\delta}))$*.*

And additionally the restrictions of $\phi_{\mathcal{T}_\delta}$ *to* $\mathcal{T}_\delta^{(R)}$ *are stochastically equicontinuous. Then the laws of* $(\underline{\mathcal{T}}_\delta)_{\delta\in\mathcal{A}}$*, form a tight sequence of probability measures on the space of measured, rooted spatial trees.*

For convenience in applying Lemma 5.2.2 to the three-dimensional UST, we next summarise the conditions that we will check for this example. Since these are of a different form to those given above, we complete the section by verifying their sufficiency in Lemma 5.2.3. We recall that the notation $B_{\mathcal{U}}(x, r)$ is used for balls in $(\mathcal{U}, d_{\mathcal{U}})$.

Assumption 1. For every $R \in (0, \infty)$, it holds that

$$\lim_{\lambda\to\infty} \limsup_{\delta\to 0} \mathbf{P}\left(\delta^3 \mu_{\mathcal{U}}\left(B_{\mathcal{U}}(0, \delta^{-\beta} R)\right) > \lambda\right) = 0.$$

Assumption 2. For every $\varepsilon, R \in (0, \infty)$, it holds that

$$\lim_{\eta\to 0} \limsup_{\delta\to 0} \mathbf{P}\left(\inf_{x \in B_{\mathcal{U}}(0,\delta^{-\beta} R)} \delta^3 \mu_{\mathcal{U}}\left(B_{\mathcal{U}}(x, \delta^{-\beta}\varepsilon)\right) < \eta\right) = 0.$$

Assumption 3. For every $\varepsilon, R \in (0, \infty)$, it holds that

$$\lim_{\eta\to 0} \limsup_{\delta\to 0} \mathbf{P}\left(\inf_{\substack{x,y \in B_{\mathcal{U}}(0,\delta^{-\beta} R): \\ \delta d_E(x,y) > \varepsilon}} \delta^{\beta} d_{\mathcal{U}}(x, y) < \eta\right) = 0.$$

Lemma 5.2.3. *If Assumptions 1, 2 and 3 hold, then so does the tightness claim of Theorem 5.1.1.*

Proof. We first check that if Assumptions 1 and 2 hold, then, for every $\varepsilon, R \in (0,\infty)$,

$$\lim_{\lambda\to\infty} \limsup_{\delta\to 0} \mathbf{P}\left(N_{\mathcal{U}}\left(\delta^{-\beta}R, \delta^{-\beta}\varepsilon\right) > \lambda\right) = 0, \tag{5.14}$$

where $N_{\mathcal{U}}(\delta^{-\beta}R, \delta^{-\beta}\varepsilon)$ is the minimal number of intrinsic balls of radius $\delta^{-\beta}\varepsilon$ needed to cover $B_{\mathcal{U}}(0, \delta^{-\beta}R)$. Towards proving this, suppose that

$$\delta^3 \mu_{\mathcal{U}}\left(B_{\mathcal{U}}(0, \delta^{-\beta}(R+\varepsilon/2))\right) \le \lambda\eta, \tag{5.15}$$

and also

$$\inf_{x\in B_{\mathcal{U}}(0,\delta^{-\beta}R)} \delta^3 \mu_{\mathcal{U}}\left(B_{\mathcal{U}}(x, \delta^{-\beta}\varepsilon/2)\right) \ge \eta. \tag{5.16}$$

Set $x_1 = 0$, and choose

$$x_{i+1} \in B_{\mathcal{U}}(0, \delta^{-\beta}R) \backslash \cup_{j=1}^{i} B_{\mathcal{U}}(x_j, \delta^{-\beta}\varepsilon),$$

stopping when this is no longer possible, to obtain a finite sequence $(x_i)_{i=1}^M$. The construction ensures that $\cup_{i=1}^M B_{\mathcal{U}}(x_i, \delta^{-\beta}\varepsilon)$ contains $B_{\mathcal{U}}(0, \delta^{-\beta}R)$, and so $M \ge N_{\mathcal{U}}(\delta^{-\beta}R, \delta^{-\beta}\varepsilon)$. Moreover, since $d_{\mathcal{U}}(x_i, x_j) \ge \delta^{-\beta}\varepsilon$ for $i \ne j$, it is the case that the balls $(B_{\mathcal{U}}(x_i, \delta^{-\beta}\varepsilon/2))_{i=1}^M$ are disjoint. Putting these observations together with (5.15) and (5.16), we find that

$$\begin{aligned}
N_{\mathcal{U}}(\delta^{-\beta}R, \delta^{-\beta}\varepsilon) &\le M \\
&\le \eta^{-1}\sum_{i=1}^{M} \delta^3 \mu_{\mathcal{U}}\left(B_{\mathcal{U}}(x_i, \delta^{-\beta}\varepsilon/2)\right) \\
&= \eta^{-1}\delta^3 \mu_{\mathcal{U}}\left(\cup_{i=1}^M B_{\mathcal{U}}(x_i, \delta^{-\beta}\varepsilon/2)\right) \\
&\le \eta^{-1}\delta^3 \mu_{\mathcal{U}}\left(B_{\mathcal{U}}(0, \delta^{-\beta}(R+\varepsilon/2))\right) \\
&\le \lambda.
\end{aligned}$$

From this, we conclude that

$$\begin{aligned}
&\mathbf{P}\left(N_{\mathcal{U}}\left(\delta^{-\beta}R,\delta^{-\beta}\varepsilon\right)>\lambda\right)\\
&\quad\leq\ \mathbf{P}\left(\delta^3\mu_{\mathcal{U}}\left(B_{\mathcal{U}}(0,\delta^{-\beta}(R+\varepsilon/2))\right)>\lambda\eta\right)\\
&\qquad+\mathbf{P}\left(\inf_{x\in B_{\mathcal{U}}(0,\delta^{-\beta}R)}\delta^3\mu_{\mathcal{U}}\left(B_{\mathcal{U}}(x,\delta^{-\beta}\varepsilon/2)\right)<\eta\right),
\end{aligned}$$

and so (5.14) follows by letting $\delta\to 0$, $\lambda\to\infty$ and then $\eta\to 0$.

Second, we show that if Assumption 3 holds, then, for every $\varepsilon, R\in(0,\infty)$,

$$\lim_{\eta\to 0}\limsup_{\delta\to 0}\mathbf{P}\left(\sup_{\substack{x,y\in B_{\mathcal{U}}(0,\delta^{-\beta}R):\\ d_{\mathcal{U}}(x,y)<\delta^{-\beta}\eta}}\delta d_E(x,y)>\varepsilon\right)=0. \tag{5.17}$$

Indeed, this follows from the elementary observation that

$$\mathbf{P}\left(\sup_{\substack{x,y\in B_{\mathcal{U}}(0,\delta^{-\beta}R):\\ d_{\mathcal{U}}(x,y)<\delta^{-\beta}\eta}}\delta d_E(x,y)>\varepsilon\right)\leq\mathbf{P}\left(\inf_{\substack{x,y\in B_{\mathcal{U}}(0,\delta^{-\beta}R):\\ \delta d_E(x,y)>\varepsilon}}\delta^{\beta}d_{\mathcal{U}}(x,y)<\eta\right).$$

Given (5.14), Assumption 1, the fact that $\delta\phi_{\mathcal{U}}(\rho_{\mathcal{U}})=0$, and (5.17), the result is a straightforward application of Lemma 5.2.2. □

5.2.1 Path ensembles

Finally, we also define the path ensemble topology used in Theorem 5.1.1. This topology was introduced by Schramm [149] in the context of scaling of two-dimensional uniform spanning trees, and a related topology (based on quad-crossings) have been used in the context of scaling limits of critical percolation. Recall that $\gamma_{\mathcal{T}}(x,y)$ is the unique path from x to y in a topological tree $\mathcal{T}$.

We denote by $\mathcal{H}(X)$ the Hausdorff space of compact subsets of a metric space X, endowed with the Hausdorff topology. This is generated by the

Hausdorff distance, given by

$$d_H(A,B) = \inf\{r \geq 0 \,:\, A \subset B_r,\ B \subseteq A_r\},$$

where $B_r = \{x \in X \,:\, d(x,B) \leq r\}$ is the r-expansion of B.

We shall consider the sphere $\mathbb{S}^3$ as the one-point compactification of $\mathbb{R}^3$, on which also consider the one-point compactification of a uniform spanning tree of $\mathbb{Z}^3$. For concreteness, fix some homeomorphism from $\mathbb{R}^3$ to $\mathbb{S}^3$ and endow it with the Euclidean metric on the sphere. Given a compact topological tree $\mathcal{T} \subset \mathbb{S}^3$, we consider the set $\Gamma_{\mathcal{T}} \subset \mathbb{S}^3 \times \mathbb{S}^3 \times \mathcal{H}(\mathbb{S}^3)$

$$\Gamma_{\mathcal{T}} = \{(x, y, \gamma_{\mathcal{T}}(x,y)) \,:\, x, y \in \mathcal{T}\}.$$

Thus $\Gamma_{\mathcal{T}}$ consists of a pair of points and the path between them. We call $\Gamma_{\mathcal{T}}$ the **path ensemble** of the tree $\mathcal{T}$. Clearly $\Gamma_{\mathcal{T}}$ is a compact subset of $\mathbb{S}^3 \times \mathbb{S}^3 \times \mathcal{H}(\mathbb{S}^3)$. Since each tree corresponds to a compact subset of $\mathbb{S}^3 \times \mathbb{S}^3 \times \mathcal{H}(\mathbb{S}^3)$, the Hausdorff topology on this product space induces a topology on trees. Theorem 5.1.1 states that the laws of the uniform spanning on $\delta\mathbb{Z}^3$ are tight and have a subsequential weak limit with respect to this topology (in addition to the Gromov-Hausdorff-type topology described above).

5.3 Loop-erased random walks

As noted in the introduction, the fundamental connection between loop-erased random walks (LERWs) and uniform spanning tree (USTs) will be crucial to this study. In this section, we recall the definition of the LERW, and collect together a number of properties of the three-dimensional LERW that hold with high probability. These properties will be useful in our study of the three-dimensional UST. We start by introducing some general notation and terminology.

5.3.1 Notation for Euclidean subsets

The **discrete ℓ^2 Euclidean ball** will be denoted by

$$B(x,r) := \left\{ y \in \mathbb{Z}^3 \,:\, |x-y| < r \right\},$$

where we write $|x-y| = d_E(x,y)$ for the Euclidean distance between x and y. (We will use the notation $|x-y|$ and $d_E(x,y)$, interchangeably.) A **δ-scaled discrete ℓ^2 ball**, for $\delta > 0$, will be denoted by

$$B_\delta(x,r) := \left\{ y \in \delta\mathbb{Z}^3 \,:\, |x-y| < r \right\},$$

and the Euclidean ℓ^2 ball is

$$B_E(x,r) := \left\{ y \in \mathbb{R}^3 \,:\, |x-y| < r \right\}.$$

We will also use the abbreviation $B(r) = B(0,r)$, similarly for B_δ and B_E. We also write $B_n(0,r) = B_{2^{-n}}(r)$. The **discrete cube** (or ℓ^∞ ball of radius r) with side-length $2r$ centred at x is defined to be the set

$$D(x,r) := \left\{ y \in \mathbb{Z}^3 \,:\, \|x-y\|_\infty < r \right\}.$$

Similarly to the definitions above, but with ℓ^∞ balls, $D_\delta(x,r)$ denotes the δ-scaled discrete cube and $D_E(x,r)$ the Euclidean cube. We further write $D(R) = D(0,R)$ and $D_n(r) = D_{2^{-n}}(0,r)$. The Euclidean distance between a point x and a set A is given by

$$\mathrm{dist}(x,A) := \inf \left\{ |x-y| \,:\, y \in A \right\}.$$

For a subset A of $\mathbb{Z}^3$, the **inner boundary** $\partial_i A$ is defined by

$$\partial_i A := \left\{ x \in A \,:\, \exists y \in \mathbb{Z}^3 \setminus A \text{ such that } |x-y| = 1 \right\}.$$

5.3.2 Notation for paths and curves

We introduce definitions related to paths and curves. Some concepts were also defined in Section 2.1.

A path in $\mathbb{Z}^3$ is a finite or infinite sequence of vertices $[v_0, v_1, \ldots]$ such that v_{i-1} and v_i are nearest neighbours, i.e. $|v_{i-1} - v_i| = 1$, for all $i \in \{1, 2, \ldots\}$. The **length** of a finite path $\gamma = [v_0, v_1, ..., v_m]$ will be denoted $\mathrm{len}(\gamma)$ and is defined to be the number of steps taken by the path, that is $\mathrm{len}(\gamma) = m$.

A **(parameterized) curve** is a continuous function $\gamma : [0, T] \to \mathbb{R}^3$. For a curve $\gamma : [0, T] \to \mathbb{R}^3$, we say that T is its **duration**, and will sometimes use the notation $T(\gamma) := T$. When the specific parameterization of a curve γ is not important, then we might consider only its trace, which is the closed subset of $\mathbb{R}^3$ given by $\mathrm{tr}\,\gamma = \{\gamma(t) : \ t \in [0, T]\}$. To simplify notation, we sometimes write γ for instead of $\mathrm{tr}\,\gamma$ where the meaning should be clear. A curve is **simple** if γ is an injective function. All curves in this chapter are assumed to be simple, often implicitly.

The space of parameterized curves of finite duration, $\mathcal{C}_f$, will be endowed with a metric ψ, as defined by

$$\psi(\gamma_1, \gamma_2) = |T_1 - T_2| + \max_{0 \le s \le 1} |\gamma_1(sT_1) - \gamma_2(sT_2)|,$$

where $\gamma_i : [0, T_i] \to \mathbb{R}^3$, $i = 1, 2$ are elements of $\mathcal{C}_f$.

We say that a continuous function $\gamma^\infty : [0, \infty) \to \mathbb{R}^3$ is a **transient (parameterized) curve** if $\lim_{t \to \infty} |\gamma^\infty(t)| = \infty$. We let $\mathcal{C}$ be the set of transient curves, and endow $\mathcal{C}$ with the metric χ given by

$$\chi(\gamma_1^\infty, \gamma_2^\infty) = \sum_{k=1}^{\infty} 2^{-k} \left(1 \wedge \max_{t \le k} |\gamma_1^\infty(t) - \gamma_2^\infty(t)| \right).$$

The concatenation of two curves $\gamma_1 : [0, T_1] \to \mathbb{R}^3$ and $\gamma_2 : [0, T_2] \to \mathbb{R}^3$ with $\gamma_1(T_1) = \gamma_2(0)$ is the curve $\gamma_1 \oplus \gamma_2$ of length $T_1 + T_2$ given by

$$\gamma_1 \oplus \gamma_2(t) := \begin{cases} \gamma_1(t) & \text{if } 0 \le t \le T_1 \\ \gamma_2(t - T_1) & \text{if } T_1 < t \le T_1 + T_2. \end{cases}$$

The time-reversal of $\gamma : [0,T] \to \mathbb{R}^3$ is the curve $\vec{\gamma} : [0,T] \to \mathbb{R}^3$ defined by

$$\vec{\gamma}(t) := \gamma(T-t), \qquad t \in [0,T].$$

We define several kinds of restrictions for a curve $\gamma : [0,T] \to \mathbb{R}^3$. Analogous restrictions are defined for transient curves. The restriction of γ to an interval $[a,b] \subseteq [0,T]$ is the curve $\gamma|_{[a,b]} : [0, b-a] \to \mathbb{R}^3$ defined by setting

$$\gamma|_{[a,b]}(t) = \gamma(t+a), \qquad 0 \le t \le b-a.$$

Similarly, if γ is a simple parametrized curve, and $x, y \in \operatorname{tr}\gamma$ and x appears before y in γ, then we define the restriction of γ between x and y to be the curve $\gamma(x,y)$, where

$$\gamma(x,y)(t) = \gamma(t+t_x), \qquad 0 \le t \le t_y - t_x,$$

with $t_x \le t_y$ satisfying $\gamma(t_x) = x$ and $\gamma(t_y) = y$. (Note that the simplicity of γ ensures that t_x and t_y are well-defined.) Finally, the restriction of γ to the Euclidean ball of radius R, with $R > 0$, is the curve $\gamma|^R := \gamma|_{[0,\xi_R \wedge T]}$, where $\xi_R = \inf\{t \in [0,T] : |\gamma(t)| \ge R\}$ is the time γ exits the ball of radius R.

Proposition 5.3.1. *Let $(\gamma_n)_{n\in\mathbb{N}} \subset \mathcal{C}_f$ be a sequence of curves. Assume that $\gamma_n \to \gamma \in \mathcal{C}_f$. Then, the convergence is preserved in $\mathcal{C}_f$ under the following operations.*

(a) Time reversal: for the sequence of curves under time-reversal

$$\vec{\gamma}_n \to \vec{\gamma} \qquad \text{as } n \to \infty.$$

(b) Restriction: for $0 \le a < b < T(\gamma)$, the restrictions

$$\gamma_n|_{[a,b]} \to \gamma|_{[a,b]} \qquad \text{as } n \to \infty,$$

where the sequence above is defined for n large enough.

(c) *Concatenation: if $\tilde{\gamma}_n \to \tilde{\gamma}$ in $\mathcal{C}_f$, then*

$$\gamma_n \oplus \tilde{\gamma}_n \to \gamma \oplus \tilde{\gamma} \qquad \text{as } n \to \infty.$$

Proof. In this proof, we write $T_n = T(\gamma_n)$ and $T = T(\gamma)$. The convergence after a time-reversal is immediate from the definition and we get (a). For (b), we consider the case $a = 0$. Let $r_n, r \in [0,1]$ be such that $b = r_n T_n$ and $b = rT$. Then

$$\begin{aligned}\psi(\gamma_n|_{[0,b]}, \gamma|_{[0,b]}) &= \max_{0 \le s \le 1} |\gamma_n(sb) - \gamma(sb)| = \max_{0 \le s \le 1} |\gamma_n(sr_nT_n) - \gamma(srT)| \\ &\le \max_{0 \le s \le 1} |\gamma_n(sr_nT_n) - \gamma(sr_nT)| + \max_{0 \le s \le 1} |\gamma(sr_nT) - \gamma(srT)|.\end{aligned}$$

The convergence of $\gamma_n \to \gamma$ implies that the first term above goes to 0 as $n \to \infty$. Note that $|r_n - r| = b|T_n^{-1} - T^{-1}| \to 0$, and hence the convergence of the last term above follows from uniform continuity of γ. The convergence of γ_n under time-reversal gives the general when $a > 0$.

Next we prove (c). We write $\tilde{T}_n = T(\tilde{\gamma}_n)$, $\tilde{T} = T(\tilde{\gamma})$ and $\delta_n = |T_n + \tilde{T}_n - (\tilde{T} + T)|$. Note that $\delta_n \to 0$ as $n \to \infty$. For $0 \le s \le 1$, when we compare the times that we compare for ψ, $|s(T_n + \tilde{T}_n) - s(T + \tilde{T})| \le \delta$. Then $\psi(\gamma_n \oplus \tilde{\gamma}_n, \gamma \oplus \tilde{\gamma})$ is bounded above by

$$\begin{aligned}&\delta_n + \max_{\substack{|r-s| \le \delta_n \\ r \le T_n \vee \tilde{T}_n,\, s \le T \vee \tilde{T}}} |\gamma_n \oplus \tilde{\gamma}_n(r) - \gamma \oplus \tilde{\gamma}(s)| \\ &\le \delta_n + \max_{\substack{|r-s| \le \delta_n \\ r \le T_n,\, s \le T}} |\gamma_n(r) - \gamma(s)| + \max_{\substack{|r-s| \le \delta_n \\ r \le \tilde{T}_n,\, s \le \tilde{T}}} |\tilde{\gamma}_n(r) - \tilde{\gamma}(s)| + \delta_n M_n,\end{aligned}$$

where M_n is an upper bound for the norms of $\gamma_n, \tilde{\gamma}_n, \gamma$, and $\tilde{\gamma}$. The last term in the inequality above comes from comparisons between γ_n and $\tilde{\gamma}$ (or between $\tilde{\gamma}_n$ and γ) close to the concatenation point. The convergence of $\gamma_n \to \gamma$ and $\tilde{\gamma}_n \to \tilde{\gamma}$, and the uniform continuity of each curve give the desired result. □

Proposition 5.3.2. *Let $(\gamma_n^\infty)_{n \in \mathbb{N}} \subset \mathcal{C}$ be a sequence of parameterized curves with limit $\gamma_n^\infty \to \gamma^\infty$ in $(\mathcal{C}, \chi)$. The convergence is preserved under the* *operations below.*

(a) *Restriction: for any* $b > 0$

$$\gamma_n^\infty|_{[0,b]} \to \gamma^\infty|_{[0,b]} \qquad \text{as } n \to \infty,$$

in the space $\mathcal{C}_f$.

(b) *Concatenation: if* $(\gamma_n)_{n\in\mathbb{N}} \subset \mathcal{C}_f$ *converges to a finite parameterized curve* γ *as* $n \to \infty$, *then*

$$\gamma_n \oplus \gamma_n^\infty \to \gamma \oplus \gamma^\infty \qquad \text{as } n \to \infty,$$

in $\mathcal{C}$.

(c) *Evaluation: if* $t_n \to t$ *then*

$$\gamma_n^\infty(t_n) \to \gamma^\infty(t) \qquad \text{as } n \to \infty.$$

Proof. The convergence in (a) follows from the definition of the metric χ. Similarly, (b) is a consequence of Proposition 5.3.1 (c) and the definition of χ. Finally, (c) follows from the uniform continuity of $\gamma_n|_{[0,k]}$. $\square$

If γ is a parameterized (simple) curve and $x, y \in \operatorname{tr}\gamma$, we define the **Schramm metric** (cf. (5.2)) by setting

$$d_\gamma^S(x,y) := \operatorname{diam}\operatorname{tr}\gamma(x,y). \tag{5.18}$$

The **intrinsic distance** between x and y is given by

$$d_\gamma(x,y) := T(\gamma(x,y)) = t_y - t_x, \tag{5.19}$$

where $\gamma(t_x) = x$ and $\gamma(t_y) = y$, i.e. this is the time duration of the curve segment between x and y. Formally, both (5.18) and (5.19) are only defined when x comes before y in γ, but the definition is extended symmetrically in the obvious way.

5.3.3 Definition and parameterization of loop-erased random walks

We will now define the loop-erased random walk. Let $S = [v_0, \ldots, v_m]$ be a path in some graph (which we take to be $\mathbb{Z}^3$ or $\delta\mathbb{Z}^3$). By erasing the cycles (or loops) in S in chronological order, we obtain a simple path from v_0 to v_m. This operation is called **loop-erasure**, and is defined as follows. Set $T(0) = 0$ and $\widetilde{v_0} = v_0$. Inductively, we set $T(j)$ according to the last visit time to each vertex:

$$T(j) = 1 + \sup\{n \colon v_n = \widetilde{v}_j\}, \qquad \tilde{v}_j = v_{T(j)}. \tag{5.20}$$

We continue until $\tilde{v}_l = v_m$, at which time $T(j) = m + 1$ and there is no additional vertex $\tilde{v}_j$. The **loop-erased random walk (LERW)** is the simple path $\mathrm{LE}(S) = [\tilde{v}_0, \ldots, \tilde{v}_l]$.

The exact same definition also applies to an infinite, transient path S. Since the path S is transient, the times $T(j)$ in (5.20) are finite, almost surely, for every $j \in \mathbb{N}$. In this case $\mathrm{LE}(S)$ is an infinite simple path.

The loop-erased random walk is just what the name implies: the loop erasure of a random walk. In $\mathbb{Z}^3$ (or $\delta\mathbb{Z}^3$) we can take S_∞ to be an infinite random walk. S_∞ is almost surely transient, so the path $L(S_\infty)$, called the **infinite loop-erased random walk (ILERW)**, is a.s. well defined. We will also need loop-erased random walks in a domain $D \subset \mathbb{R}^3$. We will write $\hat{D} = \mathbb{Z}^3 \cap D$ for the subset of vertices of $\mathbb{Z}^3$ inside D. Moreover, the (inner vertex) boundary of $\hat{D}$ is the set $\partial\hat{D}$ defined as the collection of vertices $v \in D$ for which v is connected to $v_1 \in \mathbb{Z}^3 \setminus \hat{D}$. In this case, for a given starting vertex v_0, we may take S to be a simple random walk up to the stopping time m when it first hits $\partial\hat{D}$. (We will apply this to bounded domains, so that m is almost surely finite, though the definition is valid even if $m = \infty$.) Examples of domains of a loop-erased random walk include the family of L^2 balls $\{B(R)\}_{R>0}$ and of L^∞ balls $\{D(R)\}_{r>0}$.

A discrete simple path $\gamma = (v_i)$ may naturally be considered as a curve by setting $\gamma(i) = v_i$, for $i \in \mathbb{N}$, and linearly interpolating between $\gamma(i)$ and $\gamma(i+1)$. With this parameterization, the length of γ (as a path) is equal to its duration as a curve: $\text{len}(\gamma) = T(\gamma)$. If γ is a loop-erased random walk on $\delta\mathbb{Z}^3$, its length as $\delta \to 0$, and the curve needs to be reparameterized. To obtain a macroscopic curve in the scaling limit, we reparameterize loop-erased random walks by β-**parameterization**:

$$\bar{\gamma}(t) := \gamma(\delta^{-\beta} t), \qquad \forall t \in [0, \delta^{-\beta} \, \text{len}(\gamma)],$$

where β is the LERW growth exponent. Similarly, for an infinite loop-erased random walk $\gamma_\infty = [v_0, v_1, \dots]$, we consider its associated curve γ_∞ by linearly interpolating between integer times, and its β-parameterization is given by

$$\bar{\gamma}_\infty(t) = \gamma_\infty(\delta^{-\beta} t), \qquad \forall t \geq 0.$$

In this article, we will sometimes consider the ILERW restricted to a finite domain. Specifically, if γ_∞ is an ILERW starting at the origin, we denote its restriction to a ball of radius $r > 0$ by $\gamma_\infty|^r = \text{LE}(S^\infty)|_{[0, \xi_r(LE(S^\infty))]}$, where $\xi_r(\text{LE}(S^\infty))$ is the first time $\text{LE}(S^\infty)$ exits $B(r)$. Note that this is a different object to a LERW started at the origin and stopped at the first hitting time of $\partial B(r)$. However, the two are closely related, see [136, Corollary 4.5].

5.3.4 Path properties of the infinite loop-erased random walk

In this section, we summarize some path properties of the ILERW that hold with high probability. Typically the events will involve some property that holds on the appropriate scale in a neighbourhood of radius $R\delta^{-1}$ about the starting point of the ILERW, for δ the scaling parameter, and for some fixed $R \geq 1$. Since the results hold uniformly in the scaling parameter $\delta \in (0, 1]$, they will also be useful in the scaling limit. As for notation, for $x \in \mathbb{Z}^3$, we let γ_∞^x be an ILERW on $\mathbb{Z}^3$ starting at x. If $x = 0$, then we simply write γ_∞. We highlight that in this section the space is not rescaled.

Quasi-loops

A path γ is said to have an (r, R)–**quasi-loop** if it contains two vertices $v_1, v_2 \in \gamma$ such that $|v_1 - v_2| < r$, but $\gamma(v_1, v_2) \not\subseteq B(v_1, R)$. (Up to changing the parameters slightly, this is almost the same as $d^S_\gamma(x, y) \geq R$.) We denote the set of (r, R)–quasi-loops of γ by $\mathrm{QL}(r, R; \gamma)$. Estimates on probabilities of quasi-loops in LERWs were central to Kozma's work [107]. The following bound on the probability of quasi-loops for the ILERW was established in [148] for loop-erased random walks. The extension to the infinite case follows from [148, Theorem 6.1] in combination with [136, Corollary 4.5].

Proposition 5.3.3 (cf. [148, Theorem 6.1])**.** *For every $R \geq 1$, there exist constants $C, M < \infty$, and $\tilde{\eta} > 0$ such that for any $\delta, \varepsilon \in (0, 1)$,*

$$\mathbf{P}\left(\mathrm{QL}(\varepsilon^M \delta^{-1}, \sqrt{\varepsilon}\delta^{-1}; \gamma_\infty|^{R\delta^{-1}}) \neq \emptyset\right) \leq C\varepsilon^{\tilde{\eta}}.$$

Intrinsic length and diameter

Let ξ_n be the first time that the loop-erased walk γ_∞ exits the ball $B(n)$ (i.e. the number of steps after the loop erasure). The next result is a quantitative tightness result for $n^{-\beta}\xi_n$. It is a combination of the exponential tail bounds of [155], together with the estimates on the expected value of ξ_n from [128]. We note that the result in [155] is for the LERW, but the proof is the same for the ILERW.

Proposition 5.3.4 ([155, Theorems 1.4 and 8.12] and [128, Corollary 1.3])**.** *There exist constants $C, c_1, c_2 \in (0, \infty)$ such that: for all $\lambda \geq 1$ and $n \geq 1$,*

$$\mathbf{P}\left(\xi_n \leq \lambda n^\beta\right) \geq 1 - 2e^{-c_1\lambda},$$
$$\mathbf{P}\left(\xi_n \geq \lambda^{-1} n^\beta\right) \geq 1 - Ce^{-c_2\lambda^{1/2}}.$$

While any possible pattern appears in γ_∞, the scaling relation (given by β) between the intrinsic distance and the Euclidean distance holds uniformly along the path of γ_∞. We quantify this relation in terms of equicontinuity.

Let $R \geq 1$, $\delta \in (0, 1)$ and $\lambda \geq 1$. We say that γ_∞ is **λ-equicontinuous** in the ball $B(R\delta^{-1})$ (with exponents $0 < b_1, b_2 < \infty$) if the following event holds:

$$E^*_\delta(\lambda, R) = \left\{ \begin{array}{c} \forall x, y \in \gamma_\infty|^{R\delta^{-1}}, \\ \text{if } d_{\gamma_\infty}(x,y) \leq \lambda^{-b_1}\delta^{-\beta}, \text{ then } |x-y| < \lambda^{-b_2}\delta^{-1} \end{array} \right\}.$$

The bound for the ILERW was proved in [127].

Proposition 5.3.5 (cf. [127, Proposition 7.1])**.** *There exist constants $0 < b_1, b_2 < \infty$ such that the following is true. Given $R \geq 1$, there exists a constant C such that: for all $\delta \in (0,1)$and $\lambda \geq 1$,*

$$\mathbf{P}(E^*_\delta(\lambda, R)) \geq 1 - C\lambda^{-b_2}.$$

A partial converse bounds the intrinsic distance in terms of the Schramm distance, where we recall that the Schramm distance was defined at (5.18). For $\delta, r \in (0,1]$, $\lambda \geq 1$, set

$$S^*_\delta(\lambda, r) := \left\{ \begin{array}{c} \forall x, y \in \gamma_\infty|^{\lambda r\delta^{-1}}, \\ \text{if } d^S_{\gamma_\infty}(x,y) < r\delta^{-1}, \text{ then } d_{\gamma_\infty}(x,y) < \lambda r^\beta \delta^{-\beta} \end{array} \right\}.$$

The following result follows from [127, (7.51)].

Proposition 5.3.6. *There exist constants $0 < c, C < \infty$ such that: for any $\delta, r \in (0,1]$ and $\lambda \geq 1$,*

$$\mathbf{P}(S^*_\delta(\lambda, r)) \geq 1 - C\lambda^3 e^{-c\lambda}.$$

Proof. For $u \in \mathbb{Z}^3$, let B_u be the box of side length $3r\delta^{-1}$ centred at u, and let $X_u = |\gamma_\infty|^{\lambda r\delta^{-1}} \cap B_u|$ be the number of points in B_u hit by $\gamma_\infty|^{\lambda r\delta^{-1}}$. We recall [127, equation (7.51)], which states that for some absolute c, C and any u,

$$\mathbf{P}\left(X_u \geq \lambda r^\beta \delta^{-\beta}\right) \leq Ce^{-c\lambda}.$$

Cover the ball $B(0, \lambda r\delta^{-1})$ by boxes of side length $r\delta^{-1}$ centred at some $\{u_1, \dots, u_N\}$ with $N \asymp \lambda^3$. If some pair x, y violates the event S^*_δ, and x is in the box of side length $r\delta^{-1}$ around u_i, then the segment $\gamma_\infty(x,y)$ is in the thrice larger box around the same u_i, and so $X_{u_i} \geq \lambda r^{\beta} \delta^{-\beta}$. A union bound gives the conclusion. $\square$

Capacity and hittability

As noted in the introduction, one of the key differences from the two-dimensional case is that in three dimensions it is much easier for a random walk to avoid a LERW. The electrical capacity of a connected path of diameter r in $\mathbb{Z}^3$ can be as large as Cr, but can also be as low as $O(r/\log r)$ (see Proposition 2.2.8 for the lower bound). However, the latter occurs only when the path is close to a smooth curve (see Subsection 2.2.2). The fractal nature of the scaling limit of LERWs suggests that a segment of LERW has capacity comparable to its diameter, and consequently, is likely to be hit by a second random walk starting nearby.

Let $R \geq 1$ and $r \in (0,1)$, and γ_∞^x a LERW started at x and stopped when exiting $B(0, \delta^{-1}R)$. In this subsection, we give bounds on the hitting probability of γ_∞^x by a random walk started from a point y. The hitting bounds are uniformly over the starting points $y \in B := B(x, R\delta^{-1})$ with $\mathrm{dist}(y, \gamma_\infty^x) < r\delta^{-1}$. More precisely, denote by P_S^y the probability measure of a random walk S starting at y, which is independent of γ_∞^x. We say that γ_∞^x is η-**hittable** in B if the following event holds:

$$A_\delta(x, R, r; \eta) := \left\{ \begin{array}{c} \forall y \in B(x, R\delta^{-1}) \text{ with } \mathrm{dist}(y, \gamma_\infty^x) \leq r\delta^{-1}, \\ P_S^y \left(S\left[0, \xi_S(B(y, r^{1/2}\delta^{-1}))\right] \cap \gamma_\infty^x = \emptyset \right) \leq r^\eta \end{array} \right\},$$

where $\xi_S(B(y, r^{1/2}\delta^{-1}))$ is the first time that S exits from $B(y, r^{1/2}\delta^{-1})$. (Recall $\mathrm{dist}(\cdot,\cdot)$ stands for the Euclidean distance between a point and a set.) A local version of this event, restricted to starting points near x, is given by

$$G_\delta(x, r; \eta) = \left\{ \begin{array}{c} \forall y \in B(x, r\delta^{-1}), \\ P_S^y \left(S\left[0, \xi_S(B(y, r^{1/2}\delta^{-1}))\right] \cap \gamma_\infty^x = \emptyset \right) \leq r^\eta \end{array} \right\}.$$

The next result, which was established in [148], indicates that γ_∞^x is η-

hittable with high probability.

Proposition 5.3.7 (cf. [148, Lemma 3.2 and Lemma 3.3])**.** *There exists a constant $\hat{\eta} \in (0,1)$ such that the following is true. Given $R \geq 1$, there exists a constant C such that: for all $\delta, r \in (0,1)$,*

$$\mathbf{P}\left(A_\delta(x, R, r; \hat{\eta})\right) \geq 1 - Cr.$$

In particular, $\mathbf{P}(G_\delta(x, r; \hat{\eta})) \geq 1 - Cr$.

In terms of capacity, Proposition 5.3.7 implies that, with high probability, the capacity of a connected segment of γ^∞ is comparable to its diameter.

We write $\mathbf{P}_S^{x,y}$ for the joint probability law of γ_∞^x and an independent simple random walk S starting at y. Working on the joint probability space, together with a change of variable, Proposition 5.3.7 implies the following result. This result is well-know and simply states that a simple random walk hits a ILERW almost surely.

Proposition 5.3.8 (cf. [133][Theorem 1.1, Corollary 5.3])**.** *For $x, y \in \mathbb{Z}^3$ we have that, for all $R > 0$,*

$$\inf_{\delta \in (0,1]} \mathbf{P}_S^{x,y}\left(S[0, \xi_S(B(y, R\delta^{-1}))] \cap \gamma_\infty^x = \emptyset\right) = 0.$$

Hittability of sub-paths

The main result of this subsection, Proposition 5.3.9, is crucial for obtaining exponential tail bounds on the volume of balls in the UST in Section 5.5. It establishes that the path $\gamma_\infty = \mathrm{LE}\,(S[0,\infty))$, i.e. the infinite LERW, has hittable sections across a range of distances from its starting point.

For $1 \leq \lambda < R$, consider a sequence of boxes $D_i = D\left(\frac{iR}{\lambda}\right)$, $i = 1, 2, \ldots, \lambda$, where $D(r)$ was defined in Subsection 5.3.1. Let t_i be the first time that γ_∞ exits D_i. We denote $x_i = \gamma_\infty(t_i)$, and write

$$\sigma_i = \inf\left\{n \geq t_i \mid \gamma_\infty(n) \notin B\left(x_i, \frac{R}{2\lambda}\right)\right\}.$$

For each $i = 1, 2, \dots, \lambda$, we define the event A_i by

$$A_i = \left\{ \begin{array}{c} P^z\left(R^z[0,\xi_i] \cap \gamma_\infty[t_i,\sigma_i] \cap D_{\frac{R}{2\lambda}}(x_i) \neq \emptyset\right) \geq c_0 \\ \text{for all } z \in B\left(x_i, \frac{R}{16\lambda}\right) \end{array} \right\}, \tag{5.21}$$

where: R^z is a simple random walk started at z, independent of γ_∞, with law denoted P^z; ξ_i is the first time that R^z exits $B(x_i, \frac{R}{2\lambda})$; and $D_{\frac{R}{2\lambda}}(x_i)$ is the box centered on the infinite half line started at x_i that does not intersect D_i and is orthogonal to the face of D_i containing x_i, with centre at distance $R/4\lambda$ from x and radius $\frac{R}{2{,}000\lambda}$, see Figure 5.3.

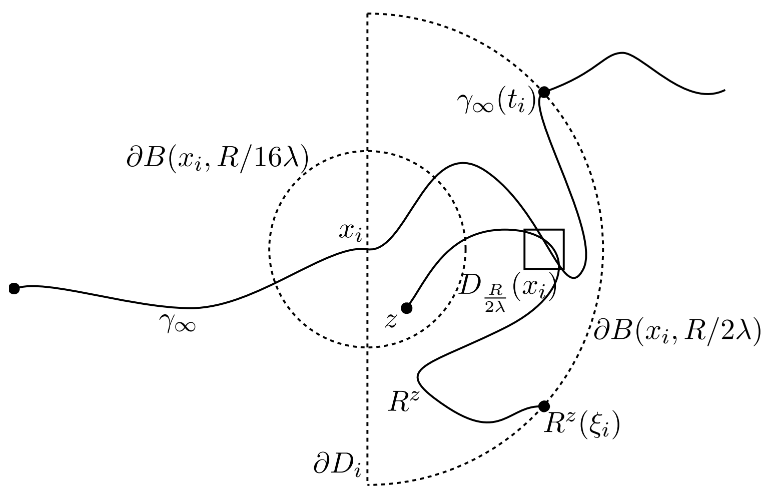

Figure 5.3: On the event A_i, as defined at (5.21), the above configuration occurs with probability greater than c_0 for any $z \in B(x_i, R/16\lambda)$.

Now, for fixed $a \in (0,1)$, we consider a sequence of subsets of the index set $\{1, 2, \dots, \lambda\}$ as follows. Let $q = \lfloor \lambda^{1-a}/3 \rfloor$. For each $j = 0, 1, \dots, q$, define the subset I_j of the set $\{1, 2, \dots, \lambda\}$ by setting

$$I_j := \{\lfloor 2j\lambda^a + 1 \rfloor, \lfloor 2j\lambda^a + 2 \rfloor, \dots, \lfloor (2j+1)\lambda^a \rfloor\}, \tag{5.22}$$

and the event F_j by

$$F_j = F_j^a = \bigcup_{i \in I_j} A_i, \tag{5.23}$$

i.e. F_j is the event that there exists at least one index $i \in I_j$ such that $\gamma_\infty[t_i, \sigma_i]$ is a hittable set in the sense that A_i holds. The next proposition shows that with high probability the event F_j holds for all $j = 1, 2, \ldots, q$. We will prove it in the following subsection.

Proposition 5.3.9. *Define the events F_j as in (5.23). There exists a universal constant $c_1 > 0$ such that*

$$\mathbf{P}\left(\bigcap_{j=1}^{q} F_j \right) \geq 1 - \lambda^{1-a} e^{-c_1 \lambda^a}. \tag{5.24}$$

Remark. (i) The reason that we decompose the ILERW γ_∞ using the sequence of random times t_i as in the above definition is that we need to control the future path $\gamma_\infty[t_i, \sigma_i]$ uniformly on the given past path $\gamma_\infty[0, t_i]$ via [155, Proposition 6.1].
(ii) We expect that each $\gamma_\infty[t_i, \sigma_i]$ is a hittable set not only with positive probability as in Proposition 5.3.16 below, but also with high probability in the sense of [148, Theorem 3.1]. However, since Proposition 5.3.9 is enough for us, we choose not to pursue this point further here.

5.3.5 Loop-erased random walks on polyhedrons

We defined that a loop-erased random walk on a domain $\hat{D} \subset \mathbb{Z}^3$ starts at an interior vertex of $\hat{D}$ and ends with its first hitting time to the boundary $\partial \hat{D}$. As we have discussed above, the geometry of the domain $\hat{D}$ affects the path properties of loop-erased random walks on it. In this subsection we will see that the results in [107, 127, 148] hold for a collection of scaled polyhedrons, which we define below. Similarly to Subsection 5.3.4, and under the assumption that the polyhedrons are scaled with a large parameter, the proofs in the aforementioned papers carry without major modifications to our setting. For clarity, we comment on the differences between the work in [127, 148] and this subsection.

A **dyadic polyhedron** on $\mathbb{R}^3$ is a connected set $\mathcal{P}$ of the form

$$\mathcal{P} = \bigcup_{j=1}^{m} C_j,$$

where each $C_j \subset \mathbb{R}^3$ is a closed cube of the form $[a_1, b_1] \times [a_2, b_2] \times [a_3, b_3]$ with $a_i, b_i \in \mathbb{Z}^3$ (cf. (5.83), where we scale the lattice instead of the polyhedron). We say that a polyhedron $\mathcal{P}$ is bounded by R if $\mathcal{P} \subset B(R)$. Let us assume that $0 \in \mathcal{P}$ and write

$$2^n \mathcal{P} := \{z \in \mathbb{Z}^3 : \ 2^{-n} z \in \mathcal{P}\},$$

for the 2^n-expansion of the polyhedron $\mathcal{P}$. In this subsection we restrict our scaling to powers of 2, and note that $2^n\mathcal{P}$ is a dyadic polyhedron as well. If $\mathcal{P}$ is bounded by R, then $B(0,1) \subset 2^n\mathcal{P} \subset B(0, 2^n R)$, for all $n \geq 1$.

Let S be a simple random walk starting at 0 and let $\xi_{\partial \mathcal{P}}$ be the exit time of the random walk from the polyhedron. In this section we study the path properties of the loop-erased random walk

$$\gamma_n^{\mathcal{P}} = \mathrm{LE}(S[0, \xi_{\partial 2^n \mathcal{P}}]). \tag{5.25}$$

Note that the index n indicates a 2^n-expansion of $\mathcal{P}$ (cf. (5.84)).

We say that $\gamma_n^{\mathcal{P}}$ is η-**hittable** if the following event holds:

$$A_n^{\mathcal{P}}(r;\eta) := \left\{ \begin{array}{c} \forall y \in 2^n\mathcal{P} \text{ with } \operatorname{dist}(y, \gamma_n^{\mathcal{P}}) \leq r2^n, \\ P_S^y \left(S\left[0, \xi_S(B(y, r^{1/2}2^n))\right] \cap \gamma_n^{\mathcal{P}} = \emptyset \right) \leq r^{\eta} \end{array} \right\},$$

where $\xi_S(B(y, r^{1/2}2^n))$ is the first time that S exits from $B(y, r^{1/2}2^n)$.

Proposition 5.3.10 (cf. Proposition 5.3.7)**.** *Fix $R \geq 1$, let $\mathcal{P}$ be a dyadic polyhedron containing* 0 *and bounded by R, and let $\gamma_n^{\mathcal{P}}$ be the loop-erased random walk in* (5.25)*. There exists a constant $\hat{\eta} \in (0,1)$ such that there exists a constant C (depending on R) and $N \geq 1$ for which the following is* *true: for all* $r \in (0,1)$ *and* $n \geq N$*,*

$$\mathbf{P}\left(A_n^{\mathcal{P}}(r;\hat{\eta})\right) \geq 1 - Cr.$$

Proposition 5.3.10 follows from [148, Lemma 3.2] and [148, Lemma 3.3], using the argument for the proof of [148, Theorem 3.1]. The argument for Proposition 5.3.10 considers two cases, depending on the starting point of the simple random walk $S(0) = y$. For some $\varepsilon > 0$, either $y \in B(0, \varepsilon n)$ or $y \in \mathcal{P} \setminus B(0, \varepsilon n)$. For the first case we apply [148, Lemma 3.2], and here we use that $\gamma_n^{\mathcal{P}}$ is a "large" path when n is large enough. If $y \in \mathcal{P} \setminus B(0, \varepsilon n)$, we then consider a covering of $\mathcal{P}$ with a collection of balls $\{B(v_i, \varepsilon^2 n)\}_{1 \leq i \leq L}$, with $v_1, \ldots v_L \in \mathcal{P} \setminus B(0, \varepsilon n)$ and $L \leq 10R^3 \lfloor \varepsilon \rfloor^{-6}$. We then use [148, Lemma 3.3] on each one of these balls and a union bound gives the desired result.

Recall the definition of (r, R)–quasi-loop in Subsection 5.3.4 and that $\mathrm{QL}(r, R; \gamma)$ denotes the set of (r, R)–quasi-loops of γ. Proposition 5.3.3 indicates the the ILERW does not have quasi-loops with high probability. A similar statement holds for a polyhedral domain. The proof makes use of Proposition 5.3.10 and we use modifications over the stopping times and the covering of the domain (as in Proposition 5.3.10). Indeed, the proof of [148, Theorem 6.1] is divided in three cases. If the LERW has a quasi-loop at a vertex v, then either v is close to the starting point of the LERW, or v is close to the boundary, or v is in an intermediate region. The probability of the first two cases is bounded by escape probabilities for random walks. We can use the same bounds in [148, Theorem 6.1] as long as the scale n is large enough (as we assume in Proposition 5.3.11). The bound for the third case follows from a union bound over a covering of the domains. We can use this argument because $\mathcal{P}$ has a regular boundary.

Proposition 5.3.11 (cf. Proposition 5.3.3)**.** *Fix* $R \geq 1$ *and let* $\mathcal{P}$ *be a dyadic polyhedron containing* 0 *and bounded by* R*, and let* $\gamma_n^{\mathcal{P}}$ *be the loop-erased random walk in* (5.25)*. There exist constants* $C, M < \infty$*,* $N \geq 1$ *and* $\tilde{\eta} > 0$ *such that for any* $\varepsilon \in (0,1)$ *and* $n \geq N$*,*

$$\mathbf{P}\left(\mathrm{QL}(\varepsilon^M 2^n, \sqrt{\varepsilon} 2^n; \gamma_n^{\mathcal{P}}) \neq \emptyset\right) \leq C\varepsilon^{\tilde{\eta}}.$$

Since Propositions 5.3.10 and 5.3.11 hold for scaled dyadic polyhedrons, we can follow the argument in [127] leading to the proof of the the scaling limit of the LERW. From this argument we obtain control of the paths and the scaling limit for the LERW $\gamma_n^{\mathcal{P}}$ with β-parameterization. We finish this sections stating these three results.

For a LERW $\gamma_n^{\mathcal{P}}$, $n \geq 1$ and $\lambda \geq 1$, the path $\gamma_n^{\mathcal{P}}$ is λ-equicontinuous (with exponents $0 < b_1, b_2 < \infty$) if

$$E_n^{\mathcal{P}}(\lambda, R) := \left\{ \forall x, y \in \gamma_n^{\mathcal{P}},\ \text{if } d_\gamma(x,y) \leq \lambda^{-b_1} 2^{\beta},\ \text{then } |x-y| < \lambda^{-b_2} 2^n \right\}.$$

The partial converse is the event:

$$S_n^{\mathcal{P}}(\lambda, r) := \left\{ \forall x, y \in \gamma_n^{\mathcal{P}},\ \text{if } d_\gamma^S(x,y) < r 2^n,\ \text{then } d_\gamma(x,y) < \lambda r^{\beta} 2^{\beta} \right\}.$$

Proposition 5.3.12 (cf. Proposition 5.3.5)**.** *There exist constants $0 < b_1, b_2 < \infty$ such that the following is true. Given $R \geq 1$, there exist constants $0 < C < \infty$ and $N \geq 1$ such that: for all $\lambda \geq 1$ and $n \geq N$,*

$$\mathbf{P}(E_n^{\mathcal{P}}(\lambda, R)) \geq 1 - C\lambda^{-b_2}.$$

Proposition 5.3.13 (cf. Proposition 5.3.6)**.** *There exist constants $0 < c, C < \infty$ and $N \geq 1$ such that: for any $r \in (0,1]$, $\lambda \geq 1$ and $n \geq N$,*

$$\mathbf{P}(S_n^{\mathcal{P}}(\lambda, r)) \geq 1 - C\lambda^3 e^{-c\lambda}.$$

Proposition 5.3.14 (cf. [127, Theorem 1.4])**.** *Let $\mathcal{P}$ be a dyadic polyhedron containing 0 and bounded by R and let $\gamma_n^{\mathcal{P}}$ be the loop-erased random walk in* (5.25)*. The β-parameterization of this loop-erased random walk is the curve given by*

$$\bar{\gamma}_n^{\mathcal{P}}(t) = \gamma_n^{\mathcal{P}}(2^{\beta n} t), \qquad t \in [0, 2^{-\beta n}\operatorname{len}(\gamma_n^{\mathcal{P}})]$$

and let $\bar{\gamma}_n^{\mathcal{P}}$ be the β-parameterization of the loop-erased random walk in (5.25)*. Then the law of $\bar{\gamma}_n^{\mathcal{P}}$ converges as $n \to \infty$ with respect to the metric space $(\mathcal{C}_f, \psi)$.*

5.3.6 Proof of Proposition 5.3.9

In this subsection we show that sub-paths of the ILERW are hittable in the sense required for the event (5.21) to hold, see Proposition 5.3.16 below. The latter result leads to the proof of Proposition 5.3.9. With this objective in mind, we first study a conditioned LERW. We begin with a list of notation.

- Recall that $D(R)$ is the cube of radius R centered at 0, as defined in Subsection 5.3.1.

- Take positive numbers m, n. Let $x \in \partial D(m)$ be a point lying in a "face" of $D(m)$ (we denote the face containing x by F). Write ℓ for the infinite half line started at x which lies in $D(m)^c$ and is orthogonal to F. We let y be the unique point which lies in ℓ and satisfies $|x-y| = n/2$. We set $D_n(x) := D(y, n/1000)$ for the box centered at y with side length $n/500$. (Cf. the definition of $D_{\frac{R}{2\lambda}}(x_i)$ above.)

- Suppose that $m, n, x, D_n(x)$ are as above. Take $K \subseteq D(m) \cup \partial D(m)$. Let X be a random walk started at x and conditioned that $X[1, \infty) \cap K = \emptyset$. We set $\eta = \mathrm{LE}\left(X[0, \infty)\right)$ for the loop-erasure of X, and σ for the first time that η exits $B(x, n)$. Finally, we denote the number of points lying in $\eta[0, \sigma] \cap D_n(x)$ by $J^K_{m,n,x}$. This is an analogue of [155, Definition 8.7].

- Suppose that X is the conditioned random walk as above. We write $G^X(\cdot, \cdot)$ for Green's function of X.

This setup is illustrated in Figure 5.4 (cf. [155, Figure 3]).

We will give one- and two-point function estimates for η in the following proposition.

Proposition 5.3.15. *Suppose that $m, n, x, K, X, \eta, \sigma$ are as above. There exists a universal constant c such that for all $z, w \in D_n(x)$ with $z \neq w$,*

$$\mathbf{P}\left(z \in \eta[0, \sigma]\right) \geq c n^{-3+\beta}, \tag{5.26}$$

$$\mathbf{P}\left(z, w \in \eta[0, \sigma]\right) \leq \frac{1}{c} n^{-3+\beta} |z - w|^{-3+\beta}. \tag{5.27}$$

Proof. The inequality (5.26) follows from [155, (8.29)] and [128, Corollary 1.3]. So, it remains to prove (5.27). We first recall [155, Proposition 8.1], the setting of which is as follows. Take $z_1, z_2 \in D_n(x)$ with $z_1 \neq z_2$. We set $z_0 = x$, and write $l = |z_1 - z_2|$. Note that $1 \leq l \leq n/100$. For $i = 0, 1, 2$, we let X^i be independent versions of X with $X^i(0) = z_i$. We write σ^i_w for the first time that X^i hits w. For $i = 0, 1$, let Z^i be X^i conditioned on the event $\{\sigma^i_{z_{i+1}} < \infty\}$, and also let $Z^2 = X^2$. Also for $i = 0, 1$, write $u(i)$ for the last time that Z^i passes through z_{i+1}, and set $u(2) = \infty$. Define the event $F^\eta_{z_1,z_2}$ by

$$F^\eta_{z_1,z_2} = \left\{ \begin{array}{c} \text{There exist } 0 < t_1 < t_2 < \infty \\ \text{such that } \eta(t_1) = z_1 \text{ and } \eta(t_2) = z_2 \end{array} \right\},$$

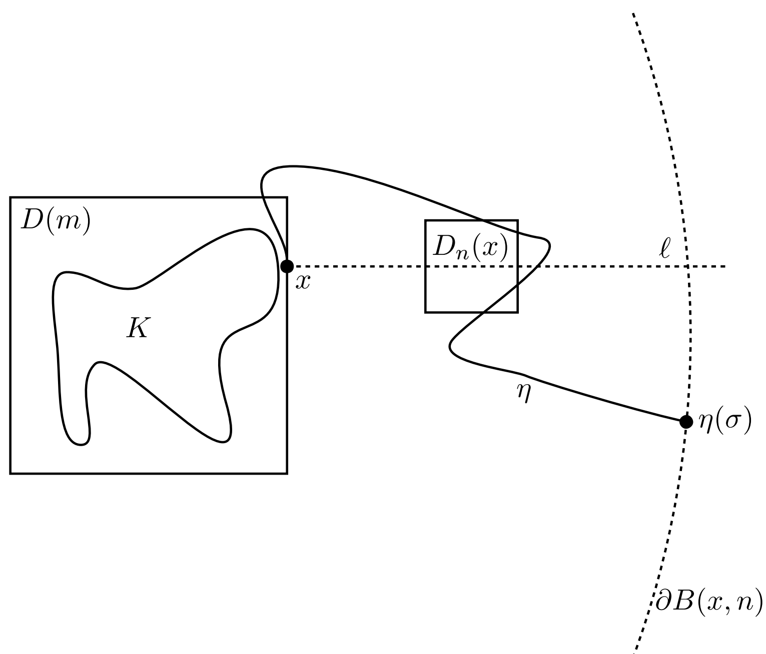

Figure 5.4: Notation used in the proof of Proposition 5.3.9.

and non-intersection events F_1 and F_2 by

$$F_1 = \left\{ \mathrm{LE}\left(Z^0[0, u(0)]\right) \cap \left(Z^1[1, u(1)] \cup Z^2[1, \infty)\right) = \emptyset \right\},$$
$$F_2 = \left\{ \mathrm{LE}\left(Z^1[0, u(1)]\right) \cap Z^2[1, \infty) = \emptyset \right\}.$$

Then [155, Proposition 8.1] shows that

$$\mathbf{P}\left(F^{\eta}_{z_1, z_2}\right) = G^X(z_0, z_1) G^X(z_1, z_2) \mathbf{P}\left(F_1 \cap F_2\right).$$

Now, in the proof of [155, Lemma 8.9], it is shown that

$$G^X(z_0, z_1) \leq \frac{C}{n}, \quad G^X(z_1, z_2) \leq \frac{C}{l},$$

and so it suffices to estimate $\mathbf{P}(F_1 \cap F_2)$. To do this, we consider four balls

$$B_1 = B(z_1, l/8), \ B_2 = B(z_2, l/8), \ B_1' = B(z_1, 2l), \ B_1'' = B(z_1, n/16).$$

Note that $B_1 \cup B_2 \subset B_1' \subset B_1''$ and $B_1 \cap B_2 = \emptyset$. For $i = 0, 1$, let $Y^i = \left(Z^i[0, u(i)]\right)^R$ be the time reversal of $Z^i[0, u(i)]$ where for a path $\lambda = [\lambda(0), \lambda(1), \dots, \lambda(u)]$, we write $(\lambda)^R = [\lambda(u), \lambda(u-1), \dots, \lambda(0)]$ for its time reversal. By the time reversibility of LERW (see [113, Lemma 7.2.1] for the time reversibility), we see that $\mathbf{P}(F_1 \cap F_2) = \mathbf{P}(F_1' \cap F_2')$, where the events F_1' and F_2' are defined by

$$F_1' = \left\{ \mathrm{LE}\left(Y^0[0, \sigma(0)]\right) \cap \left(Y^1[0, \sigma(1) - 1] \cup Z^2[1, \infty)\right) = \emptyset \right\},$$
$$F_2' = \left\{ \mathrm{LE}\left(Y^1[0, \sigma(1)]\right) \cap Z^2[1, \infty) = \emptyset \right\}.$$

Here $\sigma(i)$ is the first time that Y^i hits z_i. We define several random times as follows:

- s_0 is the first time that $\mathrm{LE}\left(Y^0[0, \sigma(0)]\right)$ exits B_1;
- s_2 is the first time that $\mathrm{LE}\left(Y^0[0, \sigma(0)]\right)$ exits B_1'';
- s_1 is the last time up to s_2 that $\mathrm{LE}\left(Y^0[0, \sigma(0)]\right)$ exits B_1';

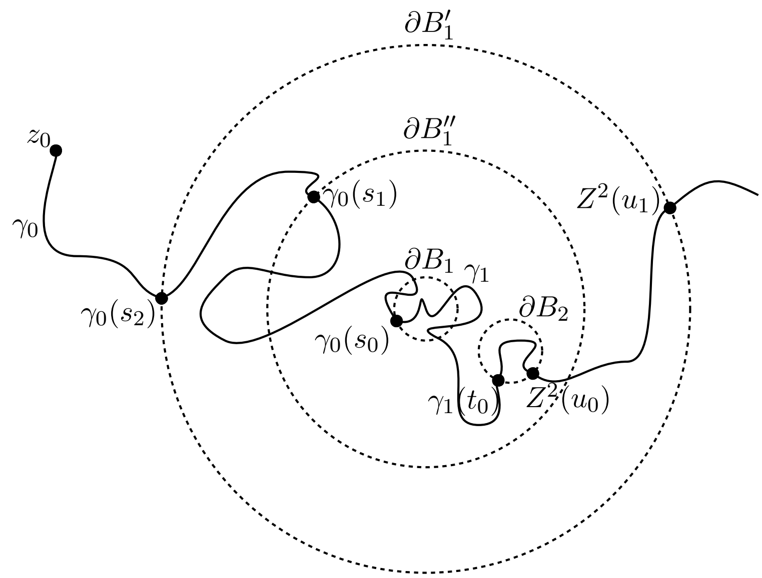

Figure 5.5: The random times s_0, s_1, s_2, t_0, u_0, u_1.

- t_0 is the first time that $\text{LE}\left(Y^1[0, \sigma(1)]\right)$ exits B_2;
- t_1 is the last time up to $\sigma(1)$ that $Y^1[0, \sigma(1)]$ hits ∂B_1;
- u_0 is the first time that Z^2 exits B_2;
- u_1 is the first time that Z^2 exits B_1''.

See Figure 5.5 for an illustration showing these random times. If we write $\gamma_i = \text{LE}\left(Y^i[0, \sigma(i)]\right)$ for $i = 0, 1$, we see that $\mathbf{P}(F_1' \cap F_2') \leq \mathbf{P}(H_1 \cap H_2 \cap H_3)$, where the events H_1, H_2, H_3 are defined by

$$\begin{aligned} H_1 &= \left\{ \gamma_0[0, s_0] \cap Y^1[t_1, \sigma(1) - 1] = \emptyset \right\}, \\ H_2 &= \left\{ \gamma_1[0, t_0] \cap Z^2[1, u_0] = \emptyset \right\}, \\ H_3 &= \left\{ \gamma_0[s_1, s_2] \cap Z^2[u_0, u_1] = \emptyset \right\}. \end{aligned}$$

Since $\text{dist}\left(D(m), B_1''\right) \geq n/4$, it follows from the discrete Harnack principle (see [113, Theorem 1.7.6], for example) that the distribution of $Z^2[0, u_1]$ is comparable to that of $R_2[0, u_1']$, assuming $R_2(0) = z_2$ where R_2 is a simple random walk, and u_1' is the first time that R_2 exits B_1''. More precisely, there exist universal constants $c, C \in (0,\infty)$ such that for any path λ

$$c\mathbf{P}\left(R_2[0,u_1'] = \lambda\right) \le \mathbf{P}\left(Z^2[0,u_1] = \lambda\right) \le C\mathbf{P}\left(R_2[0,u_1'] = \lambda\right).$$

Also, since $\gamma_0[s_1,s_2] \subseteq (B_1')^c$, using the Harnack principle again, we see that

$$\mathbf{P}(H_1 \cap H_2 \cap H_3) \asymp E_{Y^0,Y^1}\left\{ \begin{array}{c} \mathbf{1}_{H_1} P^{z_2}_{R_2}\left(\gamma_1[0,t_0] \cap R_2[1,u_0'] = \emptyset\right) \\ \cdot P^{z_1}_{R_2}\left(\gamma_0[s_1,s_2] \cap R_2[0,u_1'] = \emptyset\right) \end{array} \right\}, \quad (5.28)$$

where u_0' is the first time that R_2 exits B_2 and E_{Y^0,Y^1} stands for the expectation with respect to the probability law of (Y^0,Y^1).

Another application of the Harnack principle tells that $\gamma_1[0,t_0]$ and $Y^1[t_1,\sigma(1)-1]$ are "independent up to constant" (see [127, Lemma 4.3]). Namely, there exist universal constants $c, C \in (0,\infty)$ such that for any paths λ_1, λ_2

$$\begin{aligned} &c\mathbf{P}\left(\gamma_1[0,t_0] = \lambda_1\right)\mathbf{P}\left(Y^1[t_1,\sigma(1)-1] = \lambda_2\right) \\ &\le \mathbf{P}\left(\gamma_1[0,t_0] = \lambda_1,\ Y^1[t_1,\sigma(1)-1] = \lambda_2\right) \\ &\le C\mathbf{P}\left(\gamma_1[0,t_0] = \lambda_1\right)\mathbf{P}\left(Y^1[t_1,\sigma(1)-1] = \lambda_2\right). \end{aligned}$$

This implies that given Y^0, $\mathbf{1}_{H_1}$ and $P^{z_2}_{R_2}\left(\gamma_1[0,t_0] \cap R_2[1,u_0'] = \emptyset\right)$ are independent up to constant. Also, it is proved in [136, Propositions 4.2 and 4.4] that the distribution of $\gamma_1[0,t_0]$ is comparable with that of the ILERW started at z_2 until it exits B_2. Using the discrete Harnack principle again, we see that the distribution of the time reversal of $Y^1[t_1,\sigma(1)-1]$ coincides with that of the SRW started at z_1 until it exits B_1. Therefore, if we write R_1 and R_3 for independent SRWs, the right hand side of (5.28) is comparable to

$$E_{Y^0}\left\{P^{z_1}_{R_2}\left(\gamma_0[s_1,s_2]\cap R_2[0,u_1']=\emptyset\right)P^{z_1}_{R_1}\left(\gamma_0[0,s_0]\cap R_1[1,\sigma_1']=\emptyset\right)\right\}$$
$$\times P^{z_2,z_2}_{R_2,R_3}\left(R_2[1,u_0']\cap \mathrm{LE}\left(R_3[0,\infty)\right)[0,t_3']=\emptyset\right), \tag{5.29}$$

where σ_1' is the first time that R_1 exits B_1, and t_3' is the first time that $\mathrm{LE}\left(R_3[0,\infty)\right)$ exits B_2. Moreover, it follows from [155, Proposition 6.7] and [128, Corollary 1.3] that

$$P^{z_2,z_2}_{R_2,R_3}\left(R_2[1,u_0']\cap \mathrm{LE}\left(R_3[0,\infty)\right)[0,t_3']=\emptyset\right)\asymp l^{-2+\beta}.$$

Finally, let R_0 be a SRW started at z_1 and $\gamma_0'=\mathrm{LE}\left(R_0[0,\infty)\right)$ be the ILERW. Similarly to above, define:

- s_0' to be the first time that γ_0' exits B_1;
- s_2' to be the first time that γ_0' exits B_1'';
- s_1' to be the last time up to s_2' that γ_0' exits B_1'.

We then have from [136, Propositions 4.2 and 4.4] that the distribution of $\gamma_0[0,s_2]$ is comparable with that of $\gamma_0'[0,s_2']$. Moreover, [136, Proposition 4.6] ensures that $\gamma_0'[0,s_0']$ and $\gamma_0'[s_1',s_2']$ are independent up to a constant. Therefore the expectation with respect to Y^0 in (5.29) is comparable to

$$\begin{aligned}&P^{z_1,z_1}_{R_0,R_1}\left(\gamma_0'[0,s_0']\cap R_1[1,\sigma_1']=\emptyset\right)P^{z_1,z_1}_{R_0,R_2}\left(\gamma_0'[s_1',s_2']\cap R_2[0,u_1']=\emptyset\right)\\&\asymp \mathrm{Es}(l)\mathrm{Es}(l,n),\end{aligned} \tag{5.30}$$

where we use the notation Es defined in [155]. Finally, by [128, Corollary 1.3], it holds that the right hand side of (5.30) is comparable to $n^{-2+\beta}$. This gives (5.27) and finishes the proof. $\square$

Definition 5.3.1. Suppose that m,n,x,K,X,η,σ are as above. For each $z\in B(x,n/8)$, let R^z be a SRW on $\mathbb{Z}^3$ started at z, independent of X. Write

ξ for the first time that R^z exits $B(x,n)$, and let

$$N_z = |R^z[0,\xi] \cap \eta[0,\sigma] \cap D_n(x)|$$

be the number of points in $D_n(x)$ hit by both $R^z[0,\xi]$ and $\eta[0,\sigma]$. Furthermore, define the (random) function $g(z)$ by setting

$$g(z) := P^z\left(N_z > 0\right) = P^z\left(R^z[0,\xi] \cap (\eta[0,\sigma] \cap D_n(x)) \neq \emptyset\right),$$

where P^z stands for the probability law of R^z. Note that $g(z)$ is a measurable function of $\eta[0,\sigma]$, and that, given $\eta[0,\sigma]$, $g(\cdot)$ is a discrete harmonic function in $D_n(x)^c$.

The next proposition says that with positive probability (for η), $g(z)$ is bounded below by some universal positive constant for all $z \in B(x, n/8)$.

Proposition 5.3.16. *Suppose that the function $g(z)$ is defined as in Definition 5.3.1. There exists a universal constant $c_0 > 0$ such that*

$$\mathbf{P}\left(g(z) \geq c_0 \text{ for all } z \in B(x, n/8)\right) \geq c_0. \tag{5.31}$$

Proof. We claim that it suffices to show that

$$\mathbf{P}\left(g(x) \geq c_0\right) \geq c_0 \tag{5.32}$$

for some $c_0 > 0$. The reason for this is as follows. Suppose that (5.32) is true and the event $\{g(x) \geq c_0\}$ occurs. Since $\mathrm{dist}\left(B(x, n/8), D_n(x)\right) \geq n/4$, using the Harnack principle, there exists a universal constant $c_1 > 0$ such that $g(z) \geq c_1 g(x) \geq c_1 c_0$ for all $z \in B(x, n/8)$. Thus we have $\mathbf{P}(g(z)) \geq c_0 c_1$ for all $z \in B(x, n/8)) \geq c_0$, which gives (5.31).

We will prove (5.32). Recall the definition of N_z from Definition 5.3.1. By (5.26), we see that

$$\mathbf{E}(N_x) = \sum_{w \in D_n(x)} \mathbf{P}\left(w \in \eta[0,\sigma]\right) P^x\left(w \in R^x[0,\xi]\right) \geq c n^{-1+\beta}$$

for some $c > 0$. On the other hand, by (5.27), we have

$$\begin{aligned}\mathbf{E}(N_x^2) &= \sum_{w_1, w_2 \in D_n(x)} \mathbf{P}\left(w_1, w_2 \in \eta[0,\sigma]\right) P^x\left(w_1, w_2 \in R^x[0,\xi]\right) \\ &\leq C n^{-4+\beta} \sum_{w_1, w_2 \in D_n(x)} |w_1 - w_2|^{-4+\beta} \\ &\leq C n^{-4+\beta} n^{2+\beta} \\ &= C n^{-2+2\beta}.\end{aligned}$$

This gives $\mathbf{E}(N_x^2) \leq C\{\mathbf{E}(N_x)\}^2$. Therefore, the second moment method tells us that $\mathbf{E}(g(x)) \geq c_2$ for some universal constant $c_2 > 0$. This implies $\mathbf{P}\left(g(x) \geq c_2/2\right) \geq c_2/3$, which gives (5.32). $\square$

Proof of Proposition 5.3.9. We will prove that for each $j = 1, 2, \ldots, q$

$$\mathbf{P}(F_j^c) \leq (1 - c_0)^{\lambda^a}, \tag{5.33}$$

where c_0 is the constant of Proposition 5.3.16. Since $q \leq \lambda^{1-a}$, the inequality (5.33) gives the desired inequality (5.24). Take $j \in \{1, 2, \ldots, q\}$. Suppose that F_j^c occurs. This implies that for every $i \in I_j$, the event A_i does not occur. Setting $l = 2j\lambda^a$, we need to estimate

$$\mathbf{P}\left(\bigcap_{i=l+1}^{l+\lambda^a} A_i^c\right) = \mathbf{P}\left(A_{l+\lambda^a}^c \,\middle|\, \bigcap_{i=l+1}^{l+\lambda^a-1} A_i^c\right) \mathbf{P}\left(\bigcap_{i=l+1}^{l+\lambda^a-1} A_i^c\right).$$

Note that the event $\bigcap_{i=l+1}^{l+\lambda^a-1} A_i^c$ is measurable with respect to $\gamma[0, t_{l+\lambda^a}]$ while the event $A_{l+\lambda^a}^c$ is measurable with respect to $\gamma[t_{l+\lambda^a}, \sigma_{l+\lambda^a}]$. Therefore, using the domain Markov property of γ (see [113, Proposition 7.3.1]), Proposition 5.3.16 tells us that

$$\mathbf{P}\left(A_{l+\lambda^a}^c \,\middle|\, \bigcap_{i=l+1}^{l+\lambda^a-1} A_i^c\right) \leq 1 - c_0,$$

where we apply Proposition 5.3.16 with $m = \frac{(l+\lambda^a)R}{\lambda}$, $n = \frac{R}{2\lambda}$, $x = \gamma\left(t_{l+\lambda^a}\right)$

and $K=\gamma[0,t_{l+\lambda^a}]$. Thus we have that

$$\mathbf{P}\left(\bigcap_{i=l+1}^{l+\lambda^a} A_i^c\right) \leq (1-c_0)\mathbf{P}\left(\bigcap_{i=l+1}^{l+\lambda^a-1} A_i^c\right).$$

Repeating this procedure λ^a times, we obtain (5.33), and thereby finish the proof. $\square$

5.4 Checking the assumptions sufficient for tightness

The aim of this section is to check Assumptions 1, 2 and 3, as set out in Section 5.2. In what follows, we let $\gamma_{\mathcal{U}}(x,y)$ be the unique injective path in $\mathcal{U}$ between x and y. In particular, $\gamma_{\mathcal{U}}(x,y)(k)$ is the location at kth step of the path. Note that $\gamma_{\mathcal{U}}(x,y)(0)=x$ and $\gamma_{\mathcal{U}}(x,y)\left(d_{\mathcal{U}}(x,y)\right)=y$. Given a subset A of $\mathbb{Z}^3$, we define a α**-net** of A as the minimal set of lattice points such that $A \subseteq \bigcup_{z\in D_k} B\left(z,\alpha\right)$.

5.4.1 Assumption 1

The first assumption follows from the following proposition.

Proposition 5.4.1. *For every $R\in(0,\infty)$, there exist universal constants $\lambda_0>1$, and constants $c_1,c_2\in(0,\infty)$ depending only on R such that: for all $\lambda\geq\lambda_0$ and $\delta\in(0,1)$ small enough,*

$$\mathbf{P}\left(B_{\mathcal{U}}\left(0,R\delta^{-\beta}\right)\subseteq B\left(\lambda\delta^{-1}\right)\right)\geq 1-c_1\lambda^{-c_2}.$$

Proof. Fix $R\in(0,\infty)$. We may assume that $\delta>0$ is sufficiently small so that

$$\frac{\delta^{-1}}{-2\log_2\delta+2}\geq 10, \tag{5.34}$$

and also that $\lambda\geq 2$. For each $k\geq 1$, let $\varepsilon_k=\lambda^{-1}2^{-k}$, $\eta_k=(2k)^{-1}$ and

$$A_k=B(\delta^{-1})\setminus B((1-\eta_k)\delta^{-1}).$$

Write k_0 for the smallest integer satisfying $\delta^{-1}\varepsilon_{k_0} < 1$. We remark that the condition at (5.34) ensures that $(1-\eta_{k_0})\delta^{-1} \leq \delta^{-1} - 10$. Thus the inner boundary $\partial_i B(0, \delta^{-1})$ is contained in A_{k_0}. (We defined the inner boundary in Subsection 5.3.1.)

Let D_k be a "$\delta^{-1}\varepsilon_k$-net" of A_k. The minimality assumption implies that $|D_k| \leq C\varepsilon_k^{-3}$. Since $\delta^{-1}\varepsilon_{k_0} < 1$ and $\partial_i B(0, \delta^{-1}) \subseteq A_{k_0}$, it follows that $\partial_i B(0, \delta^{-1}) \subseteq D_{k_0}$.

Now, to construct $\mathcal{U}$, we perform Wilson's algorithm rooted at infinity (see [29, 169]) as follows:

- Consider the infinite LERW $\gamma_\infty = \mathrm{LE}(S[0,\infty))$, where $S = (S(n))_{n\geq 0}$ is a SRW on $\mathbb{Z}^3$ started at the origin. We think of $\mathcal{U}_0 = \gamma_\infty$ as the "root" in this algorithm.

- Consider a SRW started at a point in D_1, and run until it hits $\mathcal{U}_0$; we add its loop-erasure to $\mathcal{U}_0$, and denote the union of them by $\mathcal{U}_1^1$. We next consider a SRW from another point in D_1 until it hits $\mathcal{U}_1^1$; let $\mathcal{U}_1^2$ be the union of $\mathcal{U}_1^1$ and the loop-erasure of the second SRW. We continue this procedure until all points in D_1 are in the tree. We consider each loop-erased random walk as branches of the tree. Write $\mathcal{U}_1$ for the output random tree.

- We now repeat the above procedure for D_2. Namely, we think of $\mathcal{U}_1$ as a root and add a loop-erasure of SRWs from each point in D_2. Let $\mathcal{U}_2$ be the output tree. We continue inductively to define $\mathcal{U}_3, \mathcal{U}_4, \dots, \mathcal{U}_{k_0}$.

- Finally, we perform Wilson's algorithm for all points in $\mathbb{Z}^3 \setminus \mathcal{U}_{k_0}$ to obtain $\mathcal{U}$.

Note that, by construction, $\mathcal{U}_k \subseteq \mathcal{U}_{k+1}$, and also $\partial_i B(0, \delta^{-1}) \subseteq \mathcal{U}_{k_0}$.

We proceed with an estimate on the number of steps that γ_∞ takes to exit an extrinsic ball. Specifically, we define the event $F := \{\xi \geq \lambda^{-9}\delta^{-\beta}\}$, where ξ is the first time that γ_∞ exits $B(0, \lambda^{-4}\delta^{-1})$. From Proposition 5.3.4 we have that

$$\mathbf{P}(F) \geq 1 - Ce^{-c\sqrt{\lambda}}$$

for some universal constants $c, C \in (0, \infty)$.

Next, for each $x \in \mathbb{Z}^3$, let $t_x = \inf\{k \geq 0 : \gamma_{\mathcal{U}}(x,0)(k) \in \mathcal{U}_0\}$ be the first time that $\gamma_{\mathcal{U}}(x,0)$ hits $\mathcal{U}_0$. We write $\gamma_{\mathcal{U}}(x,\mathcal{U}_0) = \gamma_{\mathcal{U}}(x,0)[0,t_x]$ for the path in $\mathcal{U}$ connecting x and $\mathcal{U}_0$. We remark that $t_x = 0$ and $\gamma_{\mathcal{U}}(x,\mathcal{U}_0) = \{x\}$ when $x \in \mathcal{U}_0$. We consider the event G defined by

$$G = \left\{ \gamma_{\mathcal{U}}(x,\mathcal{U}_0) \cap B\left(0, \lambda^{-4}\delta^{-1}\right) = \emptyset \text{ for all } x \in D_1 \right\}.$$

Suppose that the event G does not occur, and that there exists an $x \in D_1$ such that $\gamma_{\mathcal{U}}(x,\mathcal{U}_0)$ hits $B\left(0, \lambda^{-4}\delta^{-1}\right)$. This implies that in Wilson's algorithm, as described above, the SRW R started at x enters into $B\left(0, \lambda^{-4}\delta^{-1}\right)$ before it hits γ_∞. Since $\delta^{-1}/2 \leq |x| \leq \delta^{-1}$, it follows from [113, Proposition 1.5.10] that

$$P_R^x\left(R[0,\infty) \cap B\left(0, \lambda^{-4}\delta^{-1}\right) \neq \emptyset\right) \leq C\lambda^{-4}$$

for some universal constant $C < \infty$, where R is a SRW started at x, with P_R^x denoting the law of the latter process. Taking the sum over $x \in D_1$ (recall that the number of points in D_1 is comparable to λ^3), we find that

$$\mathbf{P}(G) \geq 1 - C\lambda^{-1}.$$

To complete the proof, we will consider several "good" events that ensure $\gamma_\infty \cup \gamma_{\mathcal{U}}(x,\mathcal{U}_0)$ with $x \in D_k$ ($k = 1, 2, \ldots, k_0$) is a "hittable" set in the sense that if we consider another independent SRW R whose starting point is close to $\gamma_\infty \cup \gamma_{\mathcal{U}}(x,\mathcal{U}_0)$, then, with high probability for $\gamma_\infty \cup \gamma_{\mathcal{U}}(x,\mathcal{U}_0)$, it is likely that R intersects $\gamma_\infty \cup \gamma_{\mathcal{U}}(x,\mathcal{U}_0)$ quickly. Such hittability of LERW paths was studied in [148, Theorem 3.1]. With this in mind, for $k \geq 1$ and $\zeta > 0$, we define the event $H(k,\zeta)$ by setting

$$H(k,\zeta) = \left\{ \begin{array}{c} \forall x \in D_k,\ y \in B\left(x, \varepsilon_k\delta^{-1}\right): \\ P_R^y\left(R\left[0, T_R\left(x, \sqrt{\varepsilon_k}\delta^{-1}\right)\right] \cap \left(\gamma_\infty \cup \gamma_{\mathcal{U}}(x,\mathcal{U}_0)\right) = \emptyset\right) \leq \varepsilon_k^\zeta \end{array} \right\}, \tag{5.35}$$

where R is a SRW, independent of γ_∞, P_R^y stands for its law assuming that $R(0) = y$, and $T_R(x,r)$ is the first time that R exits $B(x,r)$. (For convenience, we omit the dependence of $H(k,\zeta)$ on δ.) Note that the event $H(k,\zeta)$ roughly says that when $R(0)$ is close to $\gamma_\infty \cup \gamma_{\mathcal{U}}(x,\mathcal{U}_0)$, it is likely for R to intersect with $\gamma_\infty \cup \gamma_{\mathcal{U}}(x,\mathcal{U}_0)$ before it travels very far, see Figure 5.6. From [148, Lemma 3.2], we see that the probability of the event $H(k,\zeta)$ is greater than $1-C\varepsilon_k^2$ if we take ζ sufficiently small. The reason for this is as follows. Suppose that the event $H(k,\zeta)$ does not occur, which means that there exist $x \in D_k$ and $y \in B(x,\varepsilon_k\delta^{-1})$ such that the probability considered in (5.35) is greater than ε_k^ζ. The existence of those two points $x \in D_k$ and $y \in B(x,\varepsilon_k\delta^{-1})$ implies the occurrence of the event $I(x,k,\zeta)$, as defined by

$$I(x,k,\zeta) = \left\{ \begin{array}{c} \exists y \in B\left(x,\varepsilon_k\delta^{-1}\right) \text{ such that} \\ P_R^y\left(R\left[0,T_R\left(x,\sqrt{\varepsilon_k}\delta^{-1}\right)\right] \cap \gamma_\infty^x = \emptyset\right) > \varepsilon_k^\zeta \end{array} \right\},$$

where we write γ_∞^x for the unique infinite path started at x in $\mathcal{U}$ (notice that $\gamma_\infty^0 = \gamma_\infty$). Namely, we have

$$H(k,\zeta)^c \subseteq \bigcup_{x\in D_k} I(x,k,\zeta).$$

We mention that the distribution of γ_∞^x coincides with that of the infinite LERW started at x. With this in mind, applying [148, Lemma 3.2] with $s=\varepsilon_k\delta^{-1}$, $t=\sqrt{\varepsilon_k}\delta^{-1}$ and $K=10$, it follows that there exist universal constants $\zeta_1>0$ and $C<\infty$ such that for all $k\geq 1$, $\lambda \geq 2$, $\delta\in(0,1)$ and $x\in D_k$,

$$\mathbf{P}\left(I(x,k,\zeta_1)\right) \leq C\varepsilon_k^5.$$

Since the number of points in D_k is comparable to ε_k^{-3}, we see that

$$\mathbf{P}\left(H(k,\zeta_1)\right) \geq 1-C\varepsilon_k^2, \tag{5.36}$$

as desired.

Set $A_1' := F\cap G\cap H(1,\zeta_1)$. For λ sufficiently large, the event A_1' is non-empty. We have already proved that $\mathbf{P}(A_1') \geq 1-C\lambda^{-1}$. Moreover, we note that on the event A_1', we have $d_{\mathcal{U}}(0,y) \geq \lambda^{-9}\delta^{-\beta}$ for all $y\in\gamma_\infty\cap B\left(0,\delta^{-1}/3\right)^c$. We also have on A_1' that the event G holds, and then the branch $\gamma_{\mathcal{U}}(x, \mathcal{U}_0)$ does not intersect with $B\left(0, \lambda^{-4}\delta^{-1}\right)$ for any $x \in D_1$. Since the event $F \subset A'_1$, we get that $d_{\mathcal{U}}(0, y) \geq \lambda^{-9}\delta^{-\beta}$ for all $y \in \gamma_{\mathcal{U}}(x, \mathcal{U}_0)$ with $x \in D_1$. Recall that $\mathcal{U}_1$ is the union of γ_∞ and all branches $\gamma_{\mathcal{U}}(x, \mathcal{U}_0)$ with $x \in D_1$. Conditioning $\mathcal{U}_1$ on the event A'_1, we perform Wilson's algorithm for points in D_2. It is convenient to think of $\mathcal{U}_1$ as deterministic sets in this algorithm. Adopting this perspective, we take $y \in D_2$ and consider the SRW R started at y until it hits $\mathcal{U}_1$. Suppose that R hits $B(0, \delta^{-1}/2)$ before it hits $\mathcal{U}_1$. Since the number of "$\sqrt{\varepsilon_1}\delta^{-1}$-displacements" of R until it hits $B(0, \delta^{-1}/2)$ is bigger than $10^{-1}\varepsilon_1^{-1/2}$, the hittability condition $H(1, \zeta_1)$ ensures that

$$P_R^y\left(R \text{ hits } B\left(0, \delta^{-1}/2\right) \text{ before it hits } \mathcal{U}_1\right) \leq \varepsilon_1^{\frac{c\zeta_1}{\sqrt{\varepsilon_1}}}, \tag{5.37}$$

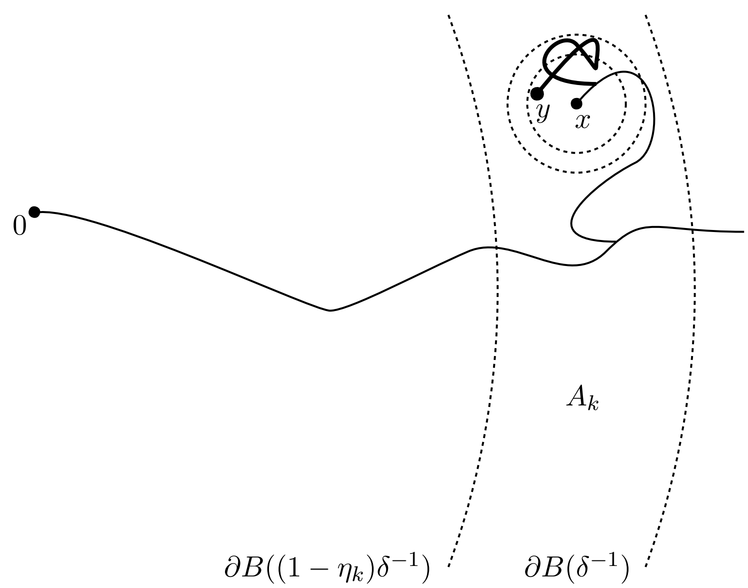

Figure 5.6: On the event $H(k, \zeta)$, as defined at (5.35), the above configuration occurs with probability greater than $1 - \varepsilon_k^{\zeta}$ for any $x \in D_k$, $y \in B(x, \varepsilon_k\delta^{-1})$. The circles shown are the boundaries of $B(x, \varepsilon_k\delta^{-1})$ and $B(x, \sqrt{\varepsilon_k}\delta^{-1})$. The non-bold paths represent $\gamma_\infty \cup \gamma_{\mathcal{U}}(x, \mathcal{U}_0)$, and the bold path $R[0, T_R(x, \sqrt{\varepsilon_k}\delta^{-1})]$.

for some universal constant $c > 0$. Define the event B_2 by

$$B_2 = \left\{ \forall y \in D_2, \gamma_{\mathcal{U}}(y, \mathcal{U}_1) \cap B\left(0, \delta^{-1}/2\right) = \emptyset \right\},$$

where $\gamma_{\mathcal{U}}(y, \mathcal{U}_1)$ denotes the branch between y and $\mathcal{U}_1$ in $\mathcal{U}$. Taking the sum over $y \in D_2$, the conditional probability (recall that we condition $\mathcal{U}_1$ on the event A_1') of the event B_2 satisfies

$$\mathbf{P}(B_2) \geq 1 - C\varepsilon_1^{-3}\varepsilon_1^{\frac{c\zeta_1}{\sqrt{\varepsilon_1}}}.$$

Thus, letting $A_2' := A_1' \cap B_2 \cap H(2, \zeta_1)$, it follows that

$$\mathbf{P}(A_2') \geq 1 - C\varepsilon_1^2,$$

where we also use that ε_1 is comparable to ε_2, and that the number of points in D_2 is comparable to ε_2^{-3}. We also note that A_2' is non-empty for λ large enough.

Conditioning $\mathcal{U}_2$ on the event A_2', we can do the same thing as above for a SRW started at $z \in D_3$. Hence if we define the event B_3 by setting

$$B_3 = \left\{ \forall z \in D_3, \gamma_{\mathcal{U}}(z, \mathcal{U}_2) \cap B\left(0, \delta^{-1}/2\right) = \emptyset \right\},$$

then the conditional probability of the event B_3 satisfies

$$\mathbf{P}(B_3) \geq 1 - C\varepsilon_2^{-3}\varepsilon_2^{\frac{c\zeta_1}{\sqrt{\varepsilon_2}}}.$$

So, letting $A_3' := A_2' \cap B_3 \cap H(3, \zeta_1)$, it follows that

$$\mathbf{P}(A_3') \geq 1 - C\varepsilon_2^2,$$

and we continue this until we reach the index k_0. In particular, if we define B_k and A_k' for each $k = 2, 3, \ldots, k_0$ by

$$B_k = \left\{ \forall z \in D_k, \gamma_{\mathcal{U}}(z, \mathcal{U}_{k-1}) \cap B\left(0, \delta^{-1}/2\right) = \emptyset \right\},$$

and $A'_k := A'_{k-1} \cap B_k \cap H(k, \zeta_1)$, we can conclude that

$$\begin{aligned} \mathbf{P}\left(A'_{k_0}\right) &= \mathbf{P}\left(A'_1\right) \prod_{k=2}^{k_0} \mathbf{P}\left(A'_k | A'_{k-1}\right) \\ &\geq \left(1 - C\lambda^{-1}\right) \prod_{k=1}^{\infty} \left(1 - C\varepsilon_k^2\right) \geq 1 - C\lambda^{-1}. \end{aligned} \tag{5.38}$$

We take a universal constant λ_0 for which (5.38) is positive for all $\lambda \geq \lambda_0$.

On the event A'_{k_0}, it is easy to see that:

- $d_{\mathcal{U}}(0, y) \geq \lambda^{-9}\delta^{-\beta}$ for all $y \in \gamma_\infty \cap B\left(0, \delta^{-1}/3\right)^c$,
- $d_{\mathcal{U}}(0, y) \geq \lambda^{-9}\delta^{-\beta}$ for all $y \in \gamma_{\mathcal{U}}(x, \mathcal{U}_0)$ and all $x \in D_k$ with $k = 1, 2, \ldots, k_0$.

Since $\partial_i B(0, \delta^{-1}) \subseteq \mathcal{U}_{k_0}$, this implies that $d_{\mathcal{U}}(0, y) \geq \lambda^{-9}\delta^{-\beta}$ for all $y \in B(0, \delta^{-1})^c$ on the event A'_{k_0}. Therefore, it follows that

$$\mathbf{P}\left(B_{\mathcal{U}}\left(0, \lambda^{-9}\delta^{-\beta}\right) \subseteq B(0, \delta^{-1})\right) \geq 1 - C\lambda^{-1}.$$

Reparameterizing this, we have

$$\mathbf{P}\left(B_{\mathcal{U}}\left(0, R\delta^{-\beta}\right) \subseteq B\left(0, \lambda\delta^{-1}\right)\right) \geq 1 - CR^{\frac{1}{9}}\lambda^{-\frac{\beta}{9}},$$

for some universal constant $C < \infty$. This finishes the proof. □

5.4.2 Assumption 2

We will prove the following variation on Assumption 2. Given Proposition 5.4.1, it is easy to check that this implies Assumption 2. The restriction of balls to the relevant Euclidean ball will be useful in the proof of the scaling limit part of Theorem 5.1.1.

Assumption 4. For every $\varepsilon, R \in (0, \infty)$, it holds that

$$\lim_{\eta \to 0} \limsup_{\delta \to 0} \mathbf{P}\left(\inf_{x \in B(\delta^{-1}R)} \delta^3 \mu_{\mathcal{U}}\left(B_{\mathcal{U}}(x, \delta^{-\beta}\varepsilon) \cap B(\delta^{-1}R)\right) < \eta\right) = 0.$$

We begin with the following warm-up lemma, which gives a lower bound on the volume of $B_{\mathcal{U}}(0,\theta\delta^{-\beta})$ for each fixed $\theta \in (0,1]$.

Lemma 5.4.2. *There exist constants $\lambda_0 > 1$, c_1, c_2 and c_3, such that: for all $\lambda \geq \lambda_0$, $\delta \in (0,1)$ and $\theta \in (0,1]$*

$$\mathbf{P}\left(\mu_{\mathcal{U}}\left(B_{\mathcal{U}}\left(0,\theta\delta^{-\beta}\right)\right) < \lambda^{-1}\delta^{-3}\right) \leq c_1\theta^{-c_2}\lambda^{-c_3}. \tag{5.39}$$

Proof. We will first deal with the case that $\theta = 1$, and then prove (5.39) for general $\theta \in (0,1]$ by reparameterizing. We may assume that $\lambda \geq 2$ and $\delta > 0$ sufficiently small. Let σ be the first time that the infinite LERW γ_∞ exits $B(\lambda^{-1/3}\delta^{-1})$. Define the event F^* by setting

$$F^* = \left\{\gamma_\infty[\sigma,\infty) \cap B\left(0,\lambda^{-1/2}\delta^{-1}\right) = \emptyset,\ \sigma \leq \lambda^{-1/4}\delta^{-\beta}\right\}.$$

Suppose that γ_∞ returns to the ball $B(0,\lambda^{-1/2}\delta^{-1})$ after time σ. Then so does the SRW that defines γ_∞ after the first time that it exits $B(0,\lambda^{-1/3}\delta^{-1}$. The probability of such a return by the SRW is, by [113, Proposition 1.5.10], smaller than $C\lambda^{-1/6}$ for some universal constant $C < \infty$. On the other hand, combining [155, Theorem 1.4] with [128, Corollary 1.3], it follows that the probability that σ is greater than $\lambda^{-1/4}\delta^{-\beta}$ is bounded above by $C\exp\left\{-c\lambda^{1/12}\right\}$ for some universal constants $c, C \in (0,\infty)$. Thus we have

$$\mathbf{P}\left(F^*\right) \geq 1 - C\lambda^{-1/6}. \tag{5.40}$$

Note that on the event F^*, the number of steps (in γ_∞) between the origin and $x \in \gamma_\infty \cap B\left(0,\lambda^{-1/2}\delta^{-1}\right)$ is smaller than $\lambda^{-1/4}\delta^{-\beta}$.

Next we introduce an event G^*, which ensures hittability of γ, similarly to the event $H(k,\zeta)$ defined at (5.35). Namely, for $\zeta > 0$, we set

$$G^*(\zeta) = \left\{\begin{array}{c} \forall x \in B\left(0,2\lambda^{-1}\delta^{-1}\right), \\ P_R^x\left(R\left[0,T_R\left(0,\lambda^{-1/2}\delta^{-1}\right)\right] \cap \gamma_\infty = \emptyset\right) \leq \lambda^{-\zeta} \end{array}\right\}.$$

From [148, Lemma 3.2], we have that there exist universal constants $C < \infty$

and $\zeta_2 > 0$ such that: for all $\lambda \geq 2$ and $\delta > 0$

$$\mathbf{P}\left(G^*(\zeta_2)\right) \geq 1 - C\lambda^{-1}.$$

Moreover, we consider the following net, which is again similar to the version appearing in the proof of Proposition 5.4.1. Here is the list of notation that we need.

- For each $k \geq 1$, let $\varepsilon_k^* = \lambda^{-\left(1+\frac{\zeta_2}{6}\right)} 2^{-k}$ and $\eta_k^* = 1/(2k)$.
- Write k_0^* for the smallest integer satisfying $\varepsilon_{k_0^*}^* \delta^{-1} < 1$.
- Set $A_k^* = B(0, (1+\eta_k^*)\lambda^{-1}\delta^{-1}) \setminus B(0, (1-\eta_k^*)\lambda^{-1}\delta^{-1})$ and let $D_k^* \subseteq \mathbb{Z}^3$ be a $\varepsilon_k^* \delta^{-1}$-net of A_k^* in the sense that the number of points in D_k^* is smaller than $C\lambda^{-3}(\varepsilon_k^*)^{-3}$ and A_k^* is contained in the union of all balls $B(z, \varepsilon_k^* \delta^{-1})$ with $z \in D_k^*$.

Note that since we take $\delta > 0$ sufficiently small, it follows that both boundaries $\partial_i B(0, \lambda^{-1}\delta^{-1})$ and $\partial B(0, \lambda^{-1}\delta^{-1})$ are contained in $D_{k_0^*}^*$.

Now we perform Wilson's algorithm as follows:

- The root of the algorithm is $\mathcal{U}_0^* := \gamma_\infty$, the infinite LERW started at the origin.
- We run sequentially LERWs from each point in D_1^* until they hit the part of the tree already constructed, and let $\mathcal{U}_1^*$ be the union of those branches and $\mathcal{U}_0^*$.
- We define $\mathcal{U}_k^*$ inductively for $k = 2, 3, \ldots, k_0^*$ by adding all branches starting from every point in D_k^* to $\mathcal{U}_{k-1}^*$.
- Finally, we consider LERW's starting from $\mathbb{Z}^3 \setminus \mathcal{U}_{k_0^*}^*$ to obtain $\mathcal{U}$.

We condition the root γ_∞ on the event $F^* \cap G^*(\zeta_2)$ and think of it as a deterministic set. Since the number of points in D_1^* is bounded above by $C\lambda^{\zeta_2/2}$, it follows that with high (conditional) probability, every branch

$\gamma_{\mathcal{U}}(x, \mathcal{U}_0^*)$ with $x \in D_1^*$ is contained in $B(0, \lambda^{-1/2}\delta^{-1})$. Namely, if we define the event H^* by

$$H^* = \left\{ \begin{array}{c} \gamma_{\mathcal{U}}(x, \mathcal{U}_0^*) \subseteq B(0, \lambda^{-1/2}\delta^{-1}) \\ \text{and } d_{\mathcal{U}}(x, \mathcal{U}_0^*) \leq \lambda^{-1/4}\delta^{-\beta} \text{ for all } x \in D_1^* \end{array} \right\},$$

where $d_{\mathcal{U}}(x, \mathcal{U}_0^*)$ stands for the number of steps of the branch $\gamma_{\mathcal{U}}(x, \mathcal{U}_0^*)$, then the condition of the event $G^*(\zeta_2)$, [155, Theorem 1.4] and [128, Corollary 1.3] ensure that the conditional probability of the event H^* satisfies

$$P(H^*) \geq 1 - C\lambda^{-\zeta_2/2}.$$

If we define the event $I^*(k, \zeta)$ by

$$I^*(k, \zeta) =$$
$$\left\{ \begin{array}{c} \forall x \in D_k^*, y \in B(x, \varepsilon_k^* \delta^{-1}), \\ P_R^y \left(R \left[0, T_R \left(x, (\varepsilon_k^*)^{1 - \frac{\zeta_2}{1000}} \delta^{-1} \right) \right] \cap (\gamma_\infty \cup \gamma_{\mathcal{U}}(x, \mathcal{U}_0^*)) = \emptyset \right) \leq (\varepsilon_k^*)^{\zeta} \end{array} \right\},$$

then a similar technique to used to deduce the inequality at (5.36) gives that there exist universal constants $\zeta_3 > 0$ and $C < \infty$ such that: for all $\lambda \geq 2$, $\delta > 0$ and $k = 1, 2, \ldots, k_0^*$,

$$\mathbf{P}\left(I^*(k, \zeta_3)\right) \geq 1 - C(\varepsilon_k^*)^2.$$

Now we define $L_1^* := F^* \cap G^*(\zeta_2) \cap H^* \cap I^*(1, \zeta_3)$. Note that on the event L_1^*, it follows that for any $y \in \left(\gamma_\infty \cap B(0, \lambda^{-1/2}\delta^{-1})\right) \cup (\cup_{x \in D_1^*} \gamma_{\mathcal{U}}(x, \mathcal{U}_0^*))$, we have

$$d_{\mathcal{U}}(0, y) \leq 2\lambda^{-1/4}\delta^{-\beta}.$$

We inductively define L_k^* for $k \geq 2$ in the following way. Let

$$H^*(k) = \left\{ \begin{array}{c} \gamma_{\mathcal{U}}\left(x, \mathcal{U}_{k-1}^*\right) \subseteq A_{k-1}^* \\ \text{and } d_{\mathcal{U}}(x, \mathcal{U}_{k-1}^*) \leq 2^{-k/8}\lambda^{-1/4}\delta^{-\beta} \text{ for all } x \in D_k^* \end{array} \right\}.$$

We define $L_k^* := L_{k-1}^* \cap I^*(2, \zeta_3) \cap H^*(k)$ for $k \geq 2$. Suppose that we condition

$\mathcal{U}^*_{k-1}$ on the event L^*_{k-1}. Since each branch $\gamma_\infty \cup \gamma_\mathcal{U}(x, \mathcal{U}^*_0)$ with $x \in D^*_{k-1}$ is a hittable set, by using a similar iteration argument to that used for (5.37), as well as [155, Theorem 1.4] and [128, Corollary 1.3], we see that the conditional probability of $H^*(k)$ is bounded above by $C\exp\{-c2^{k/4}\lambda^{1/2}\}$. With this in mind, we let

$$L^* = \bigcap_{k=1}^{k_0^*} L^*_k.$$

As at (5.38), we have

$$\begin{aligned}\mathbf{P}\left(L^*\right) &= \mathbf{P}\left(L^*_1\right)\prod_{k=2}^{k_0^*}\mathbf{P}\left(L^*_k | L^*_{k-1}\right)\\ &\geq \left(1 - C\lambda^{-\zeta_2/2}\right)\prod_{k=1}^{\infty}\left(1 - C\left(\varepsilon^*_k\right)^2\right) \geq 1 - C\lambda^{-\zeta_2/2}.\end{aligned}$$

The hard part of the proof is now complete. Indeed, on the event L^*, it is easy to check that

$$d_\mathcal{U}(0,y) \leq C\lambda^{-1/4}\delta^{-\beta},$$

as long as $y \in (\gamma_\infty \cap B(0, \lambda^{-1/2}\delta^{-1})) \cup \mathcal{U}^*_{k_0^*}$. Since the subtree $\mathcal{U}^*_{k_0}$ contains $\partial_i B(0, \lambda^{-1}\delta^{-1})$, using [155, Theorem 1.4] and [128, Corollary 1.3] again, we see that

$$\mathbf{P}\left(d_\mathcal{U}(0,y) \leq C\lambda^{-1/4}\delta^{-\beta} \text{ for all } y \in B(0, \lambda^{-1}\delta^{-1})\right) \geq 1 - C\lambda^{-c}. \quad (5.41)$$

for some universal constants $c, C \in (0,1)$. This implies that

$$\mathbf{P}\left(\mu_\mathcal{U}\left(B_\mathcal{U}\left(0, \delta^{-\beta}\right)\right) < \lambda^{-1}\delta^{-3}\right) \leq C\lambda^{-c}.$$

for some universal constants $c, C \in (0,\infty)$. Reparameterizing this, it follows that for all $\lambda \geq 1$, $\delta \in (0,1)$ and $\theta \in (0,1]$,

$$\mathbf{P}\left(\mu_\mathcal{U}\left(B_\mathcal{U}\left(0, \theta\delta^{-\beta}\right)\right) < \lambda^{-1}\delta^{-3}\right) \leq C\theta^{-\frac{3c}{\beta}}\lambda^{-c},$$

which finishes the proof. $\square$

Let γ^v_∞ be the infinite simple path in $\mathcal{U}$ started at v. When $v = 0$, we write $\gamma^0_\infty = \gamma_\infty$. Fix a point v. The next lemma gives a lower bound on $B_{\mathcal{U}}\left(x, \theta\delta^{-\beta}\right) \cap B(0, \delta^{-1})$ uniformly in $x \in \gamma^v_\infty \cap B(0, \delta^{-1})$.

Lemma 5.4.3. *There exist universal constants $\lambda_0 > 1$ and b_0, c_6, c_7, c_8 in $(0, \infty)$ such that for all $\lambda \geq \lambda_0$, $\delta \in (0,1)$, $\theta \in (0,1]$ and all $v \in B\left(0, \lambda^{b_0}\delta^{-1}\right)$,*

$$\mathbf{P}\left(\begin{array}{c} \exists x \in \gamma^v_\infty \cap B(0, \delta^{-1}) \text{ such that} \\ \mu_{\mathcal{U}}\left(B_{\mathcal{U}}\left(x, \theta\delta^{-\beta}\right) \cap B(0, \delta^{-1})\right) < \lambda^{-1}\delta^{-3} \end{array} \right) \leq c_6 \theta^{-c_7} \lambda^{-c_8}. \quad (5.42)$$

Proof. Again, by reparameterizing, it suffices to show the inequality (5.42) the case that $\theta = 1$. Also, we may assume that $\lambda \geq 2$ is sufficiently large and that $\delta > 0$ is sufficiently small. We recall that we proved at (5.41) that there exist universal constants $a_1, a_2 \in (0, \infty)$ such that

$$\mathbf{P}(A^1) := \mathbf{P}\left(B(0, \lambda^{-1}\delta^{-1}) \subseteq B_{\mathcal{U}}\left(0, a_2 \lambda^{-1/4}\delta^{-\beta}\right)\right) \geq 1 - a_2 \lambda^{-a_1}. \quad (5.43)$$

Similarly to previously, we need to deal with the hittability of γ_∞. To this end, define the event $A(\zeta)$ by

$$A(\zeta) = \left\{ \sup_{\substack{x \in B(0, \lambda\delta^{-1}): \\ \mathrm{dist}(x, \gamma_\infty) \leq \lambda^{-\frac{a_1}{10}}\delta^{-1}}} P^x_R\left(R\left[0, T_R\left(x, \lambda^{-\frac{a_1}{20}}\delta^{-1}\right)\right] \cap \gamma_\infty = \emptyset\right) \leq \lambda^{-\zeta a_1} \right\}.$$

From [148, Lemmas 3.2 and 3.3], it follows that there exist universal constants $\zeta_4 \in (0,1)$ and $C < \infty$ such that

$$\mathbf{P}\left(A(\zeta_4)\right) \geq 1 - C\lambda^{-a_1}. \quad (5.44)$$

Now we let $b_0 = a_1\zeta_4/5000$ and take $v \in B\left(0, \lambda^{b_0}\delta^{-1}\right)$, henceforth in this proof only, we write $\gamma^v_\infty = \gamma_\infty$ to simplify notation. We also write ρ_0 for the first time that γ_∞ exits $B(0, \delta^{-1})$ (we set $\rho_0 = 0$ if $v \notin B(0, \delta^{-1})$), set

$R = \lambda^{\frac{\zeta_4 a_1}{100}} \delta^{-1}$, and define ρ to be the first time that γ_∞ exits $B(0, R)$. Then a similar argument used to deduce (5.40) gives that

$$\mathbf{P}(A^2) := \mathbf{P}\left(\gamma_\infty[\rho, \infty) \cap B(0, \delta^{-1}) = \emptyset,\ \rho \leq \lambda^{\frac{a_1}{50}} \delta^{-\beta}\right) \geq 1 - C\lambda^{-\frac{\zeta_4 a_1}{100}}.$$

So, it suffices to deal with the event that there exists an $x \in \gamma_\infty[0, \rho] \cap B(0, \delta^{-1})$ for which the volume of $B_{\mathcal{U}}\left(x, \delta^{-\beta}\right) \cap B(0, \delta^{-1})$ is less than $\lambda^{-1}\delta^{-3}$.

Given these preparations, and moreover writing $r = \lambda^{-\frac{\zeta_4 a_1}{100}} \delta^{-1}$, we decompose the path $\gamma_\infty[0, \rho]$ in the following way.

- Let $\tau_0 = 0$. For $l \geq 1$, define τ_l by $\tau_l = \inf\{j \geq \tau_{l-1} : |\gamma_\infty(j) - \gamma_\infty(\tau_{l-1})| \geq r\}$.
- Let N be the unique integer such that $\tau_{N-1} < \rho_0 \leq \tau_N$.
- Set $\tau_1' = \inf\{j \geq \rho_0 : |\gamma_\infty(j) - \gamma_\infty(\rho_0)| \geq r\}$.
- For $l \geq 1$, if $\rho_{l-1} < \rho$, then we define $\tau_l' = \inf\{j \geq \rho_{l-1} : |\gamma_\infty(j) - \gamma_\infty(\rho_{l-1})| \geq r\}$ and set $\rho_l = \inf\{j \geq \tau_l' : \gamma_\infty(j) \in B(0, \delta^{-1})\} \wedge \rho$. Otherwise, we let $\tau_l' = \infty$ and $\rho_l = \rho$.
- Let N' be the smallest integer l such that $\rho_l = \rho$.
- For $0 \leq l \leq N' - 1$, we let $\tau_l'' = \max\{j \leq \rho_l : |\gamma_\infty(j) - \gamma_\infty(\rho_l)| \geq r\}$ if it is the case that $\{j \leq \rho_l : |\gamma_\infty(j) - \gamma_\infty(\rho_l)| \geq r\} \neq \emptyset$. Otherwise, we set $\tau_l'' = \rho_l$.

Notice that we don't consider the sequence $\{\tau_l\}$ if $v \notin B(0, \delta^{-1})$ since $\tau_0 = \rho_0 = 0$ in that case. (Namely, if $v \notin B(0, \delta^{-1})$, we only consider the sequence $\{\tau_l'\}$.) We also note that for any $x \in \gamma_\infty[\rho_0, \rho] \cap B(0, \delta^{-1})$, there exists $0 \leq l < N'$ such that $x \in \gamma_\infty[\rho_l, \tau_{l+1}']$.

Our first observation is that by considering the same decomposition for the corresponding SRW, it follows that the probability that $N + N' \geq \lambda^{\zeta_4 a_1/10}$ is smaller than $C\exp\{-c\lambda^c\}$. Furthermore, applying [155, Theorem 1.4] together with [128, Corollary 1.3], with probability at least $1 -$

$C\exp\{-c\lambda^c\}$, it holds that $\tau_l - \tau_{l-1} \leq \lambda^{-\frac{\zeta_4 a_1}{200}}\delta^{-\beta}$ for all $l = 1, 2, \ldots, N$, and that $\tau'_l - \rho_{l-1} \leq \lambda^{-\frac{\zeta_4 a_1}{200}}\delta^{-\beta}$ for all $l = 1, 2, \ldots, N'$. Consequently,

$$\mathbf{P}(A^3) \geq 1 - C\exp\{-c\lambda^c\},$$

where the event A^3 is defined by setting

$$A^3 = \left\{ \begin{array}{c} N + N' \leq \lambda^{\frac{\zeta_4 a_1}{10}},\ \tau_l - \tau_{l-1} \leq \lambda^{-\frac{\zeta_4 a_1}{200}}\delta^{-\beta} \text{ for all } l = 1, 2, \ldots, N \\ \text{and } \tau'_{l+1} - \tau''_l \leq \lambda^{-\frac{\zeta_4 a_1}{200}}\delta^{-\beta} \text{ for all } l = 0, 1, \ldots, N' - 1 \end{array} \right\}.$$

Replacing the constant ζ_4 by a smaller constant if necessary, [148, Theorem 6.1] (see Proposition 5.3.3) guarantees that γ_∞ has no "quasi-loops". Namely, it follows that

$$\mathbf{P}(A^4) \geq 1 - C\lambda^{-ca_1},$$

where the event A^4 is defined by setting

$$A^4 = \left\{ \begin{array}{c} B\left(\gamma_\infty(\tau_l), \lambda^{-\frac{a_1}{30}}\delta^{-1}\right) \cap (\gamma_\infty[0, \tau_{l-1}] \cup \gamma_\infty[\tau_{l+1}, \infty)) = \emptyset \\ \forall\, l = 1, 2, \ldots, N \text{ and} \\ B\left(\gamma_\infty(\rho_l), \lambda^{-\frac{a_1}{30}}\delta^{-1}\right) \cap \left(\gamma_\infty[0, \tau''_l] \cup \gamma_\infty[\tau'_{l+1}, \infty)\right) = \emptyset \\ \forall\, l = 0, 2, \ldots, N' - 1 \end{array} \right\}.$$

We now consider a $\lambda^{-\frac{a_1}{10}}\delta^{-1}$-net of $B(R)$, which we denote by D. We may assume that for each $y \in D \cap B(\delta^{-1})$, it holds that $B(y, 2\lambda^{-1}\delta^{-1}) \subseteq B(\delta^{-1})$. Notice that the number of the points in D is bounded above by $C\lambda^{a_1/3}$. For each $1 \leq l \leq N = 2$, we can find a point $x_l \in D \cap B(\delta^{-1})$ satisfying $|x_l - \gamma_\infty(\tau_l)| \leq \lambda^{-\frac{a_1}{10}}\delta^{-1}$. Also, for each $0 \leq l \leq N' - 1$, there exists a point $x'_l \in D \cap B(\delta^{-1})$ satisfying $|x'_l - \gamma_\infty(\rho_l)| \leq \lambda^{-\frac{a_1}{10}}\delta^{-1}$. (Here, note that we can find x'_l in $B(\delta^{-1})$ since $\gamma_\infty(\rho_l) \in B(\delta^{-1})$.)

We perform Wilson's algorithm as follows.

- The root of the algorithm is γ_∞.
- Consider the SRW R^1 started from x_1, and run until it hits γ_∞. We

let $\mathcal{U}^1$ be the union of γ_∞ and $\mathrm{LE}(R^1)$. Next, we consider the SRW R^2 started at x_2 until it hits $\mathcal{U}^1$; the union of it and $\mathcal{U}^1$ is denoted by $\mathcal{U}^2$. We define $\mathcal{U}^l$ for $l = 3, 4, \ldots, N-2$ similarly.

- Consider the SRW Z^0 starting from x'_0 until it hits $\mathcal{U}^{N-2}$. We let $\tilde{\mathcal{U}}^0$ be the union of $\mathcal{U}^{N-2}$ and $\mathrm{LE}(Z^0)$. Next, we consider the SRW Z^1 started at x'_1 until it hits $\tilde{\mathcal{U}}^0$; the union of $\mathrm{LE}(Z^1)$ and $\tilde{\mathcal{U}}^0$ is denoted by $\tilde{\mathcal{U}}^1$. We define $\tilde{\mathcal{U}}^l$ for $l = 2, 3, \ldots, N'-1$ similarly.
- Finally, run sequentially LERWs from every point in $\mathbb{Z}^3 \setminus \tilde{\mathcal{U}}^{N'-1}$ to obtain $\mathcal{U}$.

Define $F^0 := A^1 \cap A^2 \cap A^3 \cap A^4 \cap A(\zeta_4)$ as a "good" event for γ_∞. Conditioning γ_∞ on the event F^0, we consider all simple random walks $R^1, R^2, \ldots, R^{N-2}$, $Z^0, Z^1, \ldots, Z^{N'-1}$ starting from $x_1, x_2, \ldots, x_{N-2}, x'_0, x'_1, \ldots, x'_{N'-1}$ respectively. The event $A(\zeta_4)$ ensures that the probability that R^l (respectively Z^l) exits $B(x_l, \lambda^{-\frac{a_1}{20}}\delta^{-1})$ (resp. $B(x'_l, \lambda^{-\frac{a_1}{20}}\delta^{-1})$ before hitting γ_∞ is smaller than $\lambda^{-\zeta_4 a_1}$ for each l. Moreover, the event A^4 says that the endpoint of R^l (resp. Z^l) lies in $\gamma_\infty[\tau_{l-1}, \tau_{l+1}]$ (resp. $\gamma_\infty[\tau''_l, \tau'_{l+1}]$) for each l. On the other hand, the number of SRW's $N + N' - 2$ is less than $\lambda^{\frac{\zeta_4 a_1}{10}}$ by the event A^3. Also, we can again appeal to [155, Theorem 1.4] and [128, Corollary 1.3] to see that with probability at least $1 - C\exp\{-c\lambda^c\}$, the length of the branch $\mathrm{LE}(R^l)$ (resp. $\mathrm{LE}(Z^l)$) is less than $\lambda^{-\frac{a_1}{40}}\delta^{-\beta}$ for each $l = 1, 2, \ldots, N-2$ (respectively $l = 0, 1, \ldots, N'-1$). Thus, taking the sum over l, we see that

$$\mathbf{P}(F^1) \geq 1 - C\lambda^{-\frac{\zeta_4 a_1}{2}},$$

where the event F^1 is defined by setting

$$F^1 = \left\{ \begin{array}{c} \mathrm{LE}(R^l) \subseteq B\left(x_l, \lambda^{-\frac{a_1}{20}} \delta^{-1}\right), \\ \text{the endpoint of } R^l \text{ lies in } \gamma_\infty[\tau_{l-1}, \tau_{l+1}], \\ \text{and the length of } \mathrm{LE}(R^l) \text{ is smaller than } \lambda^{-\frac{a_1}{40}} \delta^{-\beta} \\ \text{for all } l = 1, 2, \dots, N-2 \end{array} \right\}$$
$$\cap \left\{ \begin{array}{c} \mathrm{LE}(Z^l) \subseteq B\left(x'_l, \lambda^{-\frac{a_1}{20}} \delta^{-1}\right), \\ \text{the endpoint of } Z^l \text{ lies in } \gamma_\infty[\tau''_l, \tau'_{l+1}], \\ \text{and the length of } \mathrm{LE}(Z^l) \text{ is smaller than } \lambda^{-\frac{a_1}{40}} \delta^{-\beta} \\ \text{for all } l = 0, 1, \dots, N'-1 \end{array} \right\}.$$

Recall that for each $y \in D \cap B(\delta^{-1})$, it holds that $B(y, 2\lambda^{-1}\delta^{-1}) \subseteq B(\delta^{-1})$. Since the number of the points in D is bounded above by $C\lambda^{a_1/3}$, the translation invariance of $\mathcal{U}$ and (5.43) tell that

$$\mathbf{P}(F^2) \geq 1 - a_2 \lambda^{-\frac{2a_1}{3}},$$

where the event F^2 is defined by

$$F^2 = \left\{ B(x_l, \lambda^{-1}\delta^{-1}) \subseteq B_{\mathcal{U}}\left(x_l, a_2 \lambda^{-1/4} \delta^{-\beta}\right) \text{ for all } l = 1, 2, \dots, N-2 \right\}$$
$$\cap \left\{ B(x'_l, \lambda^{-1}\delta^{-1}) \subseteq B_{\mathcal{U}}\left(x'_l, a_2 \lambda^{-1/4} \delta^{-\beta}\right) \text{ for all } l = 0, 1, \dots, N'-1 \right\}.$$

We set $F^3 := F^0 \cap F^1 \cap F^2$. Suppose that the event F^3 occurs. Take a point $x \in \gamma_\infty[0, \rho_0]$. We can then find $l \in \{0, 1, \dots, N-2\}$ such that $x \in \gamma_\infty[\tau_l, \tau_{l+2}]$. Let y_l be the endpoint of R^l. Since y_l lies in $\gamma_\infty[\tau_{l-1}, \tau_{l+1}]$, and the event A^3 holds, we see that $d_{\mathcal{U}}(x, y_l) \leq \tau_{l+2} - \tau_{l-1} \leq 3\lambda^{-\frac{\zeta_4 a_1}{200}} \delta^{-\beta}$. However, the event F^1 says that $d_{\mathcal{U}}(y_l, x_l) \leq \lambda^{-\frac{a_1}{40}} \delta^{-\beta}$. Finally, the event F^2 ensures that for every point $z \in B(x_l, \lambda^{-1}\delta^{-1})$, we have $d_{\mathcal{U}}(x_l, z) \leq a_2 \lambda^{-1/4} \delta^{-\beta}$. So, the triangle inequality tells that $d_{\mathcal{U}}(x, z) \leq 5\lambda^{-\frac{\zeta_4 a_1}{200}} \delta^{-\beta}$ for all $z \in B(x_l, \lambda^{-1}\delta^{-1}) \subseteq B(\delta^{-1})$.

We next consider a point $x \in \gamma_\infty[\rho_0, \rho] \cap B(\delta^{-1})$. There then exists $0 \leq l < N'$ such that $x \in \gamma_\infty[\rho_l, \tau'_{l+1}]$. Let y'_l be the endpoint of Z^l. Since y'_l lies in $\gamma_\infty[\tau''_l, \tau'_{l+1}]$, and the event A^3 holds, we see that $d_{\mathcal{U}}(x, y'_l) \leq \tau'_{l+1} -$

$\tau_l'' \le \lambda^{-\frac{\zeta_4 a_1}{200}}\delta^{-\beta}$. However, the event F^1 says that $d_{\mathcal{U}}(y_l', x_l') \le \lambda^{-\frac{a_1}{40}}\delta^{-\beta}$. Finally, the event F^2 ensures that for every point $z \in B(x_l', \lambda^{-1}\delta^{-1})$, we have $d_{\mathcal{U}}(x_l', z) \le a_2\lambda^{-1/4}\delta^{-\beta}$. So, the triangle inequality tells that $d_{\mathcal{U}}(x, z) \le 3\lambda^{-\frac{\zeta_4 a_1}{200}}\delta^{-\beta}$ for all $z \in B(x_l', \lambda^{-1}\delta^{-1}) \subseteq B(\delta^{-1})$.

This implies that for all $x \in \gamma_\infty[0, \rho] \cap B(\delta^{-1})$,

$$\mu_{\mathcal{U}}\left\{B_{\mathcal{U}}\left(x, 5\lambda^{-\frac{\zeta_4 a_1}{200}}\delta^{-\beta}\right) \cap B(\delta^{-1})\right\} \ge c\lambda^{-3}\delta^{-3}. \tag{5.45}$$

Reparameterizing this, we finish the proof. □

Assumption 4 immediately follows from the next lemma.

Lemma 5.4.4. *There exist constants c_1, c_2, c_3 such that: for all $\lambda \ge 1$, $\delta \in (0,1)$ and $\theta \in (0,1]$,*

$$\mathbf{P}\left(\inf_{x \in B(\delta^{-1})} \delta^3 \mu_{\mathcal{U}}\left(B_{\mathcal{U}}(x, \theta\delta^{-\beta}) \cap B(\delta^{-1})\right) < \lambda^{-1}\right) \le c_1\theta^{-c_2}\lambda^{-c_3}.$$

Proof. We will only consider the case that $\theta = 1$. We also assume that $\lambda \ge 2$ is sufficiently large and that $\delta > 0$ is sufficiently small, similarly to the proof of the previous lemma. Moreover, we will use the same notation as in the proof of Lemma 5.4.3. Recall that the constants a_1 and ζ_4 appeared at (5.43) and (5.44), and that we defined $b_0 := a_1\zeta_4/5000$ and $R := \lambda^{\frac{\zeta_4 a_1}{100}}\delta^{-1}$. For $v \in B(\lambda^{b_0}\delta^{-1})$, ρ was defined to be the first time that γ_∞^v exits $B(R)$ ($\rho = 0$ if $v \notin B(\delta^{-1})$). In the proof of Lemma 5.4.3, we proved that for each $v \in B(\lambda^{b_0}\delta^{-1})$,

$$P\left(\begin{array}{c}\mu_{\mathcal{U}}\left\{B_{\mathcal{U}}\left(x, \lambda^{-b_1}\delta^{-\beta}\right) \cap B(0, \delta^{-1})\right\} \ge c\lambda^{-3}\delta^{-3} \\ \text{for all } x \in \gamma_\infty^v[0, \rho] \cap B(\delta^{-1})\end{array}\right) \ge 1 - C\lambda^{-b_1}, \tag{5.46}$$

for some $b_1 > 0$, see (5.45). Let $b_2 = \frac{\zeta_4 a_1}{10^8} \wedge \frac{b_1}{10^8}$. We consider a $\lambda^{-b_2}\delta^{-1}$-net $D' = (x^l)_{l=1}^M$ of $B(0, 2\delta^{-1})$. Note the number of points in D', which is denoted by M, can be assumed to be smaller than $C\lambda^{3b_2}$.

Now we perform Wilson's algorithm as follows:

- The root of the algorithm is $\gamma_\infty = \gamma_\infty^0$.

- Consider the SRW R^1 started at $x^1 \in D'$, and run until it hits γ_∞. Let $\mathcal{U}_1$ be the union of $\mathrm{LE}(R^1)$ and γ_∞. We then consider the SRW R^l started from $x^l \in D'$, and run until it hits $\mathcal{U}_{l-1}$; add $\mathrm{LE}(R^l)$ to $\mathcal{U}_{l-1}$ – this union is denoted by $\mathcal{U}_l$. Since $M \leq C\lambda^{3b_2}$, by applying (5.46) for each x^l, we have
$$\mathbf{P}(V^1) \geq 1 - C\lambda^{-b_1/2},$$
where the event V^1 is defined by setting
$$V^1 := \left\{ \begin{array}{c} \mu_{\mathcal{U}} \left\{ B_{\mathcal{U}} \left(x, \lambda^{-b_1}\delta^{-\beta} \right) \cap B(\delta^{-1}) \right\} \geq c\lambda^{-3}\delta^{-3}, \\ \text{for all } x \in \mathcal{U}_M \cap B(\delta^{-1}) \end{array} \right\}.$$

- Taking $a > 0$ such that $a \sum_{j=1}^{\infty} j^{-2} = 1/2$, we let $a_k = a \sum_{j=1}^{k} j^{-2}$, and consider a $2^{-k}\lambda^{-b_2}\delta^{-1}$-net $D^k = (x_i^k)_i$ of $B((2 - a_k)\delta^{-1})$, where the number of points in D^k is bounded above by $C2^{3k}\lambda^{3b_2}$. Let k_0 be the smallest integer k such that $2^{-k}\lambda^{-b_2}\delta^{-1} \leq 1$.

- Perform Wilson's algorithm for all points in D^1 adding new branches to $\mathcal{U}_M$; the output tree is denoted by $\hat{\mathcal{U}}_1$. Then perform Wilson's algorithm for points D^k ($k = 2, 3, \ldots, k_0$) inductively; the output trees are denoted by $\hat{\mathcal{U}}_2, \ldots, \hat{\mathcal{U}}_{k_0}$. Note that $B(\delta^{-1}) \subseteq \hat{\mathcal{U}}_{k_0}$.

Since every branch generated in the procedure above is a hittable set, we can prove that there exist universal $0 < b_3 < b_2$ and $C > 0$ such that
$$\mathbf{P}(V^2) \geq 1 - C\lambda^{b_3}, \tag{5.47}$$
where the event V^2 is defined by
$$V^2 := \left\{ \forall x \in \hat{\mathcal{U}}_{k_0}, d_{\mathcal{U}}(x, x(M)) \leq \lambda^{-b_3}\delta^{-\beta} \text{ and } x(M) \in B(\delta^{-1}) \right\}.$$

Here, for each x, we write $x(M) \in \mathcal{U}_M$ for the point such that $d_{\mathcal{U}}(x, x(M)) = d_{\mathcal{U}}(x, \mathcal{U}_M)$. The inequality (5.47) can be proved in a similar way to the proof of Proposition 5.4.1, so the details are left to the reader.

Suppose that the event $V^1 \cap V^2$ occurs. Since $B(\delta^{-1}) \subseteq \hat{\mathcal{U}}_{k_0}$, this implies that for any $x \in B(\delta^{-1})$, we have

$$\mu_{\mathcal{U}}\left\{B_{\mathcal{U}}\left(x, 2\lambda^{-b_3}\delta^{-\beta}\right) \cap B(0, \delta^{-1})\right\} \geq c\lambda^{-3}\delta^{-3}.$$

A simple reparameterization completes the proof. $\square$

Combining Proposition 5.4.1 and Lemma 5.4.4, we have the following.

Corollary 5.4.1. Assumptions 2 and 4 hold.

5.4.3 Assumption 3

In this subsection, we will prove the following proposition.

Proposition 5.4.5. *Assumption 3 holds.*

Proof. In [127], it is proved that there exist universal constants $b_3, b_4 \in (0, \infty)$ such that for all $\delta \in (0, 1)$ and $\lambda \geq 1$,

$$\mathbf{P}(J_1) \geq 1 - b_4\lambda^{-b_3}, \tag{5.48}$$

where the event J_1 is defined by setting

$$J_1 = \left\{ \begin{array}{c} \forall x, y \in \gamma_\infty \cap B\left(\lambda^{b_3}\delta^{-1}\right) \\ \text{with } d_{\mathcal{U}}(x, y) \leq \lambda^{-b_4}\delta^{-\beta}, |x - y| \leq \lambda^{-b_3}\delta^{-1} \end{array} \right\},$$

see [127, (7.19)] in particular. We also need the hittability of γ_∞ as follows. For $\zeta > 0$, define the event $J(\zeta)$ by setting

$$J(\zeta) = \left\{ \begin{array}{c} P_R^x\left(R\left[0, T_R\left(x, \lambda^{b_3/2}\delta^{-1}\right)\right] \cap \gamma_\infty = \emptyset\right) \leq \lambda^{-\zeta b_3} \\ \text{for all } x \in B\left(\lambda^{b_3/4}\delta^{-1}\right) \end{array} \right\}.$$

It follows from [148, Lemma 3.2 and Lemma 3.3] that there exist universal constants $C < \infty$ and $\zeta_5 \in (0, 1)$ such that for all $\delta > 0$ and $\lambda \geq 1$,

$$\mathbf{P}\left(J(\zeta_5)\right) \geq 1 - C\lambda^{-b_3}.$$

With this in mind, we set $b_5 = \frac{\zeta_5 b_3}{1000}$ and $R_1 = \lambda^{b_5}\delta^{-1}$. Let $D'' = (z_l)_l$ be a $\lambda^{-b_5}\delta^{-1}$-net of $B(R_1)$. The number of points M'' of D'' can be assumed to be smaller than $C\lambda^{6b_5}$. We perform Wilson's algorithm as follows. The root of the algorithm is γ_∞ as usual. Then we consider the loop-erasure of the SRWs $R^1, R^2, \dots, R^{M''}$ started from $z_1, z_2, \dots, z_{M''}$ respectively; we denote the output tree by $\mathcal{U}_{M''}$. Finally, we consider LERW's starting from all points in $\mathbb{Z}^3 \setminus \mathcal{U}_{M''}$.

Conditioning γ_∞ on the event $J_1 \cap J(\zeta_5)$, for each $l = 1, 2, \dots, M''$, the probability that R^l exits $B\left(z_l, \lambda^{b_3/2}\delta^{-1}\right)$ before hitting γ_∞ is, on the event $J(\zeta_5)$, bounded above by $\lambda^{-\zeta_5 b_3}$. Taking the sum over l, we see that if

$$J_2 := \left\{ R^l \left[0, T_{R^l}\left(x, \lambda^{b_3/2}\delta^{-1}\right)\right] \cap \gamma_\infty \neq \emptyset \text{ for all } l = 1, 2, \dots, M'' \right\},$$

then

$$\mathbf{P}(J_2) \geq 1 - C\lambda^{-b_5}.$$

On the other hand, if we define

$$J_3^l = \left\{ \begin{array}{c} \forall x, y \in \gamma_\infty^{z_l} \cap B\left(z_l, \lambda^{b_3}\delta^{-1}\right) \\ \text{with } d_{\mathcal{U}}(x,y) \leq \lambda^{-b_4}\delta^{-\beta}, |x-y| \leq \lambda^{-b_3}\delta^{-1} \end{array} \right\},$$

for each $l = 1, 2, \dots, M''$, (recall that γ_∞^x stands for the unique infinite path in $\mathcal{U}$ starting from x,) by the translation invariance of $\mathcal{U}$ and (5.48), it follows that $\mathbf{P}(J_3^l) \geq 1 - b_4\lambda^{-b_3}$ for all l. Thus, letting

$$J_3 = \bigcap_{l=1}^{M''} J_3^l,$$

we have $\mathbf{P}(J_3) \geq 1 - \lambda^{-b_3}$.

Now, suppose that the event $J := J_1 \cap J(\zeta_5) \cap J_2 \cap J_3$ occurs. The triangle inequality tells that on the event J, for all $x, y \in \mathcal{U}_{M''} \cap B(\lambda^{\frac{2b_3}{3}}\delta^{-1})$ with $d_{\mathcal{U}}(x,y) \leq \lambda^{-b_4}\delta^{-\beta}$, we have $|x-y| \leq 3\lambda^{-b_3}\delta^{-1}$. Thus

$$\mathbf{P}\left(\begin{array}{c} \forall x, y \in \mathcal{U}_{M''} \cap B\left(\lambda^{\frac{2b_3}{3}}\delta^{-1}\right) \\ \text{with } d_{\mathcal{U}}(x,y) \leq \lambda^{-b_4}\delta^{-\beta}, |x-y| \leq 3\lambda^{-b_3}\delta^{-1} \end{array} \right) \geq 1 - C\lambda^{-b_5}.$$

By the translation invariance of $\mathcal{U}$ again, we can prove that each branch $\gamma_\infty^{z_l}$ is also a hittable set with high probability. Namely, if we let

$$J^l(\zeta) = \left\{ \begin{array}{c} P_R^x\left(R\left[0, T_R\left(x, \lambda^{-b_5/2}\delta^{-1}\right)\right] \cap \gamma_\infty^{z_l} = \emptyset\right) \leq \lambda^{-\zeta b_5} \\ \text{for all } x \in B(z_l, \lambda^{-b_5}\delta^{-1}) \end{array} \right\}$$

for each $l = 1, 2, \ldots, M''$, then by using [148, Lemma 3.2], we see that there exist universal constants $\zeta_6 \in (0,1)$ and $C < \infty$ such that for all $\delta \in (0,1)$, $\lambda \geq 1$ and $l = 1, 2, \ldots M''$,

$$\mathbf{P}\left(J^l(\zeta_6)\right) \geq 1 - C\lambda^{-100b_5}.$$

With this in mind, we let

$$J_4 := \bigcap_{l=1}^{M''} J^l(\zeta_6),$$

so that $\mathbf{P}(J_4) \geq 1 - C\lambda^{-b_5}$.

Conditioning $\mathcal{U}_{M''}$ on the event $J \cap J_4$, we perform Wilson's algorithm for all points in $B(R_1/2) \setminus \mathcal{U}_{M''}$, considering finer and finer nets there as in the proof of Proposition 5.4.1. The event J_4 ensures that every SRW starting from a point w in $B(R_1/2)$ hits $\mathcal{U}_{M''}$ before it exits $B(w, \lambda^{-b_5/3}\delta^{-1})$ with probability at least $1 - C\lambda^{-\zeta_6 b_5 \lambda^{\frac{b_5}{6}}}$. Thus we can conclude that with probability at least $1 - C\lambda^{-b_5}$, we have $\operatorname{diam}(\gamma_{\mathcal{U}}(w, \mathcal{U}_{M''})) \leq \lambda^{-b_5/3}\delta^{-1}$ and $d_{\mathcal{U}}(w, \mathcal{U}_{M''}) \leq \lambda^{-b_5/4}\delta^{-\beta}$ for all $w \in B\left(0, R_1/2\right)$. Therefore, by the triangle inequality again, it follows that

$$\mathbf{P}\left(\begin{array}{c} \forall x, y \in \mathcal{U} \cap B\left(R_1/2\right) \\ \text{with } d_{\mathcal{U}}(x,y) \leq \lambda^{-b_4}\delta^{-\beta}, |x-y| \leq \lambda^{-b_3/5}\delta^{-1} \end{array} \right) \geq 1 - C\lambda^{-b_5}. \quad (5.49)$$

Finally, Proposition 5.4.1 shows that with probability $1 - C\lambda^{-cb_5}$, the intrinsic ball $B_{\mathcal{U}}(0, L\delta^{-\beta}) \subseteq B(R_1/2)$ for each fixed L. Combining this with (5.49) completes the proof. $\square$

5.5 Exponential lower tail bound on the volume

In Lemma 5.4.3, we established a polynomial (in λ) lower tail bound on the volume of a ball. In this section, we will improve this bound to an exponential one, see Theorem 5.5.2 below. We start by proving the following analogue of [20, Theorem 3.4] in three dimensions. The proof strategy is modelled on that of the latter result, though there is a key difference in that the Beurling estimate used there (see [113, Theorem 2.5.2]) is not applicable in three dimensions, and we replace it with the hittability estimate of Proposition 5.3.9.

Theorem 5.5.1. *There exist constants $\lambda_0 > 1$ and $c, C, b \in (0, \infty)$ such that: for all $R \geq 1$ and $\lambda \geq \lambda_0$,*

$$\mathbf{P}\left(\mu_{\mathcal{U}}\left(B_{\mathcal{U}}(0,R)\right) \leq \lambda^{-1} R^{\frac{3}{\beta}}\right) \leq C \exp\left\{-c\lambda^{b}\right\}. \tag{5.50}$$

Proof. We begin by describing the setting of the proof. We assume that $\lambda \geq 1$ is sufficiently large, and let $a = \frac{99}{100}$. Let $q = [\lambda^{(1-a)}/3]$ be the number of subsets $I_0, I_1, \dots, I_q$ of the index set $\{1, 2, \dots, \lambda\}$, as defined in (5.22). Note that for all $0 \leq j_1 < j_2 \leq q$ and all $i_1 \in I_{j_1}, i_2 \in I_{j_2}$ we have

$$\operatorname{dist}\left(\partial D_{i_1}, \partial D_{i_2}\right) \geq \lambda^{a-1} R. \tag{5.51}$$

For each $j = 0, 1, \dots, q$, recall that the event F_j stands for the event that there exists a "good" index $i \in I_j$ in the sense that $\gamma[t_i, \sigma_i]$ is a hittable set. By Proposition 5.3.9, with probability at least $1 - \lambda^{1-a} \exp\{-c_1 \lambda^a\}$, the ILERW γ has a good index in I_j for every $j = 0, 1, \dots, q$. Let

$$F = \bigcap_{j=1}^{q} F_j, \tag{5.52}$$

and suppose that the event F occurs. It then holds that, for each $j = 0, 1, \dots, q$, we can find a good index $i_j \in I_j$ such that the event A_{i_j} occurs. We will moreover fix deterministic nets $W^p = (w^p_k)_k$, $p = 1, 2, 3$, of $B(2R)$

satisfying

$$B(2R) \subseteq \bigcup_k B\left(w_k^p, \frac{R}{10^2\lambda^{2p}}\right) \quad \text{and} \quad |w_k^p - w_{k'}^p| \geq \frac{R}{10^4\lambda^{2p}} \text{ for all } k \neq k'.$$

Note that we may assume that $|W^p| \asymp \lambda^{6p}$.

From now on, we assume that the event F occurs whenever we consider γ. We also highlight the correspondence between our setting and that of [20, Theorem 3.4]. In the proof of [20, Theorem 3.4], k points $z_1, z_2, \ldots, z_k$ were chosen on the ILERW. Here the points $x_{i_0} = \gamma(t_{i_0}), x_{i_1} = \gamma(t_{i_1}), \ldots, x_{i_q} = \gamma(t_{i_q})$ correspond to those points. Setting $n = \frac{R}{2\lambda}$, we write $B_j = B(x_{i_j}, n)$ for $j = 0, 1, \ldots, q$. Note that for each $j_1 \neq j_2$

$$\operatorname{dist}(B_{j_1}, B_{j_2}) \geq \frac{\lambda^{a-1}R}{2} \tag{5.53}$$

by (5.51).

As in [20, (3.18) and (3.19)], we define the events F^1, F^2 by setting

$$F^1 = \{\gamma[T_{2R}, \infty) \text{ hits more than } q/2 \text{ of } B_0, B_1, \ldots B_q\}, \tag{5.54}$$

$$F^2 = \left\{T_{2R} \geq \lambda^{a'} R^{\beta}\right\},$$

where T_r is the first time that γ exits $B(r)$, and $a' = \frac{1}{1000}$, see Figure 5.7. Here we also need to introduce the event F^3, as given by

$$F^3 = \left\{\exists w_k^1 \in W^1 \text{ such that } N_k^1 \geq \lambda^5\right\},$$

where N_k^1 is equal to

$$\left| \left\{ \begin{array}{c} w_l^2 \in W^2 \ : \\ B\left(w_l^2, \frac{R}{10^2\lambda^4}\right) \subseteq B\left(w_k^1, \frac{R}{10^2\lambda^2}\right) \text{ and } B\left(w_l^2, \frac{R}{10^2\lambda^4}\right) \cap \gamma[0,\infty) \neq \emptyset \end{array} \right\} \right|,$$

i.e., N_k^1 stands for the number of balls of the net W^2 contained in the ball $B(w_k^1, \frac{R}{10^2\lambda^2})$ and hit by ILERW γ.

We will first show that $\mathbf{P}(F^1)$ is exponentially small in λ. Let Γ_r be the set of paths ζ satisfying $\mathbf{P}(\gamma[0,T_r] = \zeta) > 0$. Namely, Γ_r stands for the set of all possible candidates for $\gamma[0,T_r]$. Take $\zeta \in \Gamma_{2R}$, and let $z = \zeta(\mathrm{len}(\zeta))$ be the endpoint of ζ. Write Y for the random walk started at z and conditioned on the event that $Y[1,\infty) \cap \zeta = \emptyset$. The domain Markov property (see [113, Proposition 7.3.1]) yields that the distribution of $\gamma[T_{2R},\infty)$ conditioned on the event $\{\gamma[0,T_{2R}] = \zeta\}$ coincides with that of $\mathrm{LE}(Y[0,\infty))$. Therefore, we have

$$\mathbf{P}(F^1) \leq \sum_{\zeta \in \Gamma_{2R}} \mathbf{P}(H_\zeta)\mathbf{P}\left(\gamma[0,T_{2R}] = \zeta\right),$$

where the event H_ζ is defined by

$$H_\zeta = \{Y \text{ hits more than } q/2 \text{ of } B_0, B_1 \ldots, B_q\}.$$

Recall that R^z stands for the simple random walk started at z. We remark that $\mathrm{dist}(z, B_j) \geq R/4$ for all $j \in \{0,1,\ldots,q\}$. Let τ be the first time that

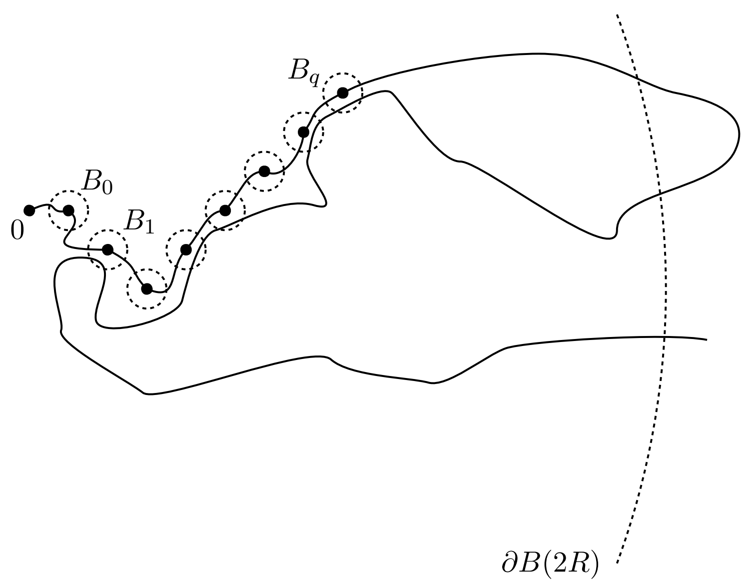

Figure 5.7: A typical realisation of γ on the event F^1, as defined at (5.54).

R^z exits $B(z, R/8)$, and observe that

$$\mathbf{P}\left(R^z[1,\infty) \cap \zeta = \emptyset\right) \asymp \mathbf{P}\left(R^z[1,\tau] \cap \zeta = \emptyset\right). \tag{5.55}$$

Indeed, it is clear that the left-hand side is bounded above by the right-hand side. To see the opposite inequality, we note that [155, Proposition 6.1] (see also [148, Claim 3.4]) yields that

$$\begin{aligned}
&\mathbf{P}\left(R^z[1,\infty) \cap \zeta = \emptyset\right) \\
&\geq \mathbf{P}\left(R^z[1,\tau] \cap \zeta = \emptyset,\ \operatorname{dist}\left(B(2R), R^z(\tau)\right) \geq \frac{R}{16},\ R^z[\tau,\infty) \cap B(2R) = \emptyset\right) \\
&\geq c\mathbf{P}\left(R^z[1,\tau] \cap \zeta = \emptyset,\ \operatorname{dist}\left(B(2R), R^z(\tau)\right) \geq \frac{R}{16}\right) \\
&\geq c'\mathbf{P}\left(R^z[1,\tau] \cap \zeta = \emptyset\right),
\end{aligned}$$

which gives (5.55). Consequently, we obtain that

$$\begin{aligned}
\mathbf{P}(H_\zeta) &\leq \frac{C\mathbf{P}\left(R^z[1,\tau] \cap \zeta = \emptyset,\ R^z \text{ hits more than } q/2 \text{ of } B_0, B_1 \ldots, B_q\right)}{\mathbf{P}\left(R^z[1,\tau] \cap \zeta = \emptyset\right)} \\
&\leq C \max_{z' \in B(z,R/8)} \mathbf{P}\left(R^{z'} \text{ hits more than } q/2 \text{ of } B_0, B_1 \ldots, B_q\right).
\end{aligned}$$

Take $z' \in B(z, R/8)$, and note that $\operatorname{dist}(z', B_j) \geq R/8$ for all $j = 0, 1, \ldots, q$. We define a sequence of stopping times $u_1, u_2, \ldots$ as follows. Let

$$u_1 = \inf\left\{t \geq 0 \ :\ R^{z'}(t) \in \bigcup_{j=0}^{q} B_j\right\},$$

and j^1 be the unique index such that $R^{z'}(u_1) \in B_{j^1}$. For $l \geq 2$, we define u_l by setting

$$u_l = \inf\left\{t \geq u_{l-1} \ :\ R^{z'}(t) \in \left(\bigcup_{j=0}^{q} B_j\right) \setminus B_{j^{l-1}}\right\},$$

and write j^l for the unique index such that $R^{z'}(u_l) \in B_{j^l}$. Since the distance between two different balls is bigger than $\lambda^{a-1}R/2$ by (5.53), and each ball has radius $n = R/2\lambda$, it follows from [113, Proposition 1.5.10] that

$$\mathbf{P}\left(u_l < \infty \,|\, u_{l-1} < \infty\right) \leq C\lambda^{-a}\lambda^{1-a} = C\lambda^{-\frac{49}{50}},$$

for all l. Thus, taking λ sufficiently large so that $C\lambda^{-\frac{49}{50}} < 1/2$, it holds that

$$\mathbf{P}(H_\zeta) \leq C(1/2)^{q/2} \leq C\exp\left\{-c\lambda^{\frac{1}{100}}\right\},$$

which gives

$$\mathbf{P}(F^1) \leq C\exp\left\{-c\lambda^{\frac{1}{100}}\right\}. \tag{5.56}$$

As for the event F^2, we have from Proposition 5.3.4 that

$$\mathbf{P}(F^2) \leq C\exp\left\{-c\lambda^{a'}\right\}. \tag{5.57}$$

Finally, we will deal with the event F^3. Define

$$M_k^1 = \left| \left\{ w_l^2 \in W^2 \;:\; \begin{array}{c} B\left(w_l^2, \frac{R}{10^2\lambda^4}\right) \subseteq B\left(w_k^1, \frac{R}{10^2\lambda^2}\right) \\ \text{and } B\left(w_l^2, \frac{R}{10^2\lambda^4}\right) \cap S[0,\infty) \neq \emptyset \end{array} \right\} \right|,$$

in other words, M_k^1 stands for the number of balls of the net W^2 contained in $B\left(w_k^1, \frac{R}{10^2\lambda^2}\right)$ and hit by SRW $S[0,\infty)$. It is clear that $N_k^1 \leq M_k^1$. Thus, on the event F^3, there exists $w_k^1 \in W^1$ such that $M_k^1 \geq \lambda^5$. However, for each k, it is easy to see that $\mathbf{P}(M_k^1 \geq \lambda^5) \leq Ce^{-c\lambda}$. Therefore, since $|W^1| \asymp \lambda^6$, we see that

$$\mathbf{P}(F^3) \leq C\exp\left\{-c\lambda^{1/2}\right\}. \tag{5.58}$$

We are now ready to follow the proof of [20, Theorem 3.4]. If the event $F^c \cup F^1 \cup F^2 \cup F^3$ (recall that the event F is defined at (5.52)) occurs, we terminate the algorithm with a 'Type 1' failure. Otherwise, for each $j = 0, 1, \ldots, q$, we can find $z_j \in W^3 \cap B(x_{i_j}, n/8)$ such that $B(z_j, \lambda^{-4}) \cap \gamma[0,\infty) = \emptyset$. Using this point z_j, we write

$$B'_j = B\left(z_j, \lambda^{-4}R\right), \qquad B''_j = B\left(z_j, \lambda^{-6}R\right).$$

Let $\mathcal{U}_0 = \gamma[0,\infty)$. Suppose that the event $F \cap \bigcap_{k=1}^3 (F^k)^c$ occurs. We consider the SRW R^{z_0} until it hits $\mathcal{U}_0$. Let $\gamma_0 = \mathrm{LE}(R^{z_0})$ be its loop-erasure which is the branch on $\mathcal{U}$ between z_0 and $\mathcal{U}_0$. Define the event G_1^0 by $G_1^0 = \{R^{z_0} \not\subseteq B_0\}$. Since γ satisfies the event F, we see that

$$\mathbf{P}\left((G_1^0)^c\right) \geq c_0. \tag{5.59}$$

Suppose that the event G_1^0 occurs. We mark the ball B_j as 'bad' if $R^{z_0} \cap B_j \neq \emptyset$. Otherwise, we define the event $G_2^0 := \{\mathrm{len}(\gamma_0) \geq \lambda^{-1/2} R^\beta\} \cap \{R^{z_0} \subseteq B_0\}$. If the event G_2^0 occurs, we also mark B_0 as 'bad' (we only mark B_0 in this case). By Proposition 5.3.4, it holds that

$$\mathbf{P}\left(G_2^0\right) \leq C \exp\left\{-c\lambda^{1/2}\right\}. \tag{5.60}$$

If the event $(G_1^0)^c \cap (G_2^0)^c$ occurs, we use Wilson's algorithm to fill in the reminder of B_0''. Define the event G_3^0 by setting

$$G_3^0 = \left\{ \begin{array}{c} \exists v \in B_0'' \text{ such that } \gamma_{\mathcal{U}}\left(v, \gamma_0 \cup \mathcal{U}_0\right) \not\subseteq B_0' \\ \text{or } \mathrm{len}\left(\gamma_{\mathcal{U}}\left(v, \gamma_0 \cup \mathcal{U}_0\right)\right) \geq \lambda^{-2} R^\beta \end{array} \right\} \cap (G_1^0)^c \cap (G_2^0)^c,$$

where we recall that $\gamma_{\mathcal{U}}(v, A)$ stands for the branch on $\mathcal{U}$ between v and A. Modifying the proof of Lemma 5.4.2, we see that

$$\mathbf{P}\left(G_3^0\right) \leq C\lambda^{-c} \tag{5.61}$$

for some universal constants $c, C \in (0, \infty)$. Suppose that the event G_3^0 occurs. We again mark the ball B_j as 'bad' if S^v hits B_j for some $v \in B_0''$ in the algorithm above. If the event $(G_1^0)^c \cap (G_2^0)^c \cap (G_3^0)^c$ occurs, we label this first 'ball step' as successful and we terminate the whole algorithm. In this case, for all $v \in B_0''$

$$d_{\mathcal{U}}(0, v) \leq \left(\lambda^{a'} + \lambda^{-1/2} + \lambda^{-2}\right) R^\beta \leq C\lambda^{a'} R^\beta,$$

and so

$$\mu_{\mathcal{U}}\left(B_{\mathcal{U}}\left(0, C\lambda^{a'} R^\beta\right)\right) \geq c\lambda^{-18} R^3. \tag{5.62}$$

If the event $G_1^0 \cup G_2^0 \cup G_3^0$ occurs, we denote the number of bad balls by N_0^B. Using a similar idea used to establish (5.56), we see that

$$\mathbf{P}\left(N_0^B \geq \sqrt{q}/4\right) \leq C \exp\left\{-c\lambda^{1/200}\right\}. \tag{5.63}$$

If $N_0^B \geq \sqrt{q}/4$, we terminate the whole algorithm as 'Type 2' failure. If $N_0^B < \sqrt{q}/4$, we can choose B_j which is not bad and perform the second 'ball step', replacing B_0 with B_j in the above. We terminate this ball step algorithm whenever we get a successful ball step or we have Type 2 failure. We write F^4 for the event that some ball step ends with a Type 2 failure. Since we perform at most $q^{1/2}$ ball steps, it follows from (5.63) that

$$\mathbf{P}(F^4) \leq C \exp\left\{-c\lambda^{1/400}\right\}. \tag{5.64}$$

Finally, we let F^5 be the event that we can perform the jth ball step for all $j = 1, 2 \dots, q^{1/2}$ without Type 2 failure and success. By combining (5.59), (5.60) and (5.61), taking λ sufficiently large, we see that each ball step has a probability at least $c_0/2$ of success. Therefore, we have

$$\mathbf{P}(F^5) \leq C \exp\left\{-c\lambda^{1/200}\right\}. \tag{5.65}$$

Once we terminate the ball step algorithm with a success, we end up with a good volume estimate as in (5.62). Combining (5.56), (5.57), (5.58), (5.64), (5.65) with Proposition 5.3.9, we conclude that

$$\mathbf{P}\left(\mu_{\mathcal{U}}\left(B_{\mathcal{U}}\left(0, C\lambda^{a'}R^{\beta}\right)\right) \geq c\lambda^{-18}R^3\right) \geq 1 - C\exp\left\{-c\lambda^{a'}\right\}.$$

Reparameterizing this gives the desired result. $\square$

We are now ready to derive the main result of this section, which gives exponential control of the volume of balls, uniformly over spatial regions.

Theorem 5.5.2. *There exist constants $\lambda_0 > 1$ and $c, C, b \in (0, \infty)$ such* *that: for all $R \geq 1$ and $\lambda \geq \lambda_0$,*

$$\mathbf{P}\left(\inf_{x\in B(R^{1/\beta})} \mu_{\mathcal{U}}\left(B_{\mathcal{U}}\left(x, \lambda^{-b'} R\right)\right) \leq \lambda^{-1} R^{\frac{3}{\beta}}\right) \leq C \exp\left\{-c\lambda^{b}\right\}. \tag{5.66}$$

Proof. We will follow the strategy used in the proofs of Lemmas 5.4.3 and 5.4.4. Theorem 5.5.1 tells us that if $A_1 := \{|B_{\mathcal{U}}(0, \lambda^{-1}R)| \leq \lambda^{-4} R^{\frac{3}{\beta}}\}$, then

$$\mathbf{P}(A_1) \leq C \exp\left\{-c\lambda^{b}\right\}, \tag{5.67}$$

We may assume that $b \in (0,1)$. We also let $b_1 = b/1000$. Applying Proposition 5.3.7 with $s = \exp\{-\lambda^{b_1}\}R^{1/\beta}$, $r = \exp\{-\lambda^{b_1}/2\}R^{1/\beta}$ and $K = 100$, we find that there exists universal constants $\eta \in (0,1)$ and $C > 0$ such that

$$\mathbf{P}(A_2) \leq C \exp\left\{-\lambda^{b_1}\right\},$$

where A_2 is defined to be the event

$$\left\{\begin{array}{c} \exists v \in B(5R^{1/\beta}) \text{ such that} \\ \operatorname{dist}(v,\gamma) \leq e^{-\lambda^{b_1}} R^{1/\beta} \text{ and } P^v\left(R^v[0,t_v] \cap \gamma = \emptyset\right) \geq e^{-\eta\lambda^{b_1}} \end{array}\right\}.$$

Here, γ represents the ILERW started at the origin, R^v stands for a SRW started at v, the probability law of which is denoted by P^v, and t_v stands for the first time that R^v exits $B(v, \exp\{-\lambda^{b_1}/2\}R^{1/\beta})$. We next use Proposition 5.3.3 to conclude that there exists universal constants $C, M < \infty$ such that

$$\mathbf{P}(A_3) \leq C \exp\left\{-\lambda^{b_1}/M\right\},$$

where the event A_3 is defined by

$$A_3 = \left\{\begin{array}{c} \exists v \in B(5R^{1/\beta}) \text{ and } i < j \text{ such that} \\ \gamma(i), \gamma(j) \in B\left(v, 10\exp\{-\lambda^{b_1}/2\}R^{1/\beta}\right) \\ \text{and } \gamma[i,j] \not\subseteq B\left(v, 10^{-1}\exp\{-\lambda^{b_1}/M\}R^{1/\beta}\right) \end{array}\right\},$$

Namely, the event A_3 says that γ has a quasi-loop in $B(5R^{1/\beta})$. We next let

$$\delta = 10^{-3} \min\{\eta, 1/M\}, \tag{5.68}$$

and define a sequence of random times $s_1, s_2, \dots$ by setting $s_0 = 0$,

$$s_1 = \inf\left\{t \geq 0 \,:\, \gamma(t) \notin B\left(\exp\{-\delta\lambda^{b_1}\}R^{1/\beta}\right)\right\},$$
$$s_i = \inf\left\{t \geq s_{i-1} \,:\, \gamma(t) \notin B\left(\gamma(s_{i-1}), \exp\{-\delta\lambda^{b_1}\}R^{1/\beta}\right)\right\}, \qquad \forall i \geq 2.$$

Let $x_i = \gamma(s_i)$, write

$$I = \left\{i \geq 1 \,:\, (\gamma[s_{i-1}, s_i] \cup \gamma[s_i, s_{i+1}] \cup \gamma[s_{i+1}, s_{i+2}]) \cap B(4R^{1/\beta}) \neq \emptyset\right\},$$

and set $N = |I|$. By considering the number of balls of radius $\exp\{-\delta\lambda^{b_1}\}R^{\frac{1}{\beta}}$ crossed by a SRW before ultimately leaving $B(4R^{1/\beta})$, we see that

$$\mathbf{P}(A_4) \leq C \exp\left\{-ce^{\delta\lambda^{b_1}}\right\},$$

where $A_4 := \{N \geq \exp\{3\delta\lambda^{b_1}\}\}$. A similar argument to that used in [127, (7.51)] yields that

$$\mathbf{P}(A_5) \leq C \exp\left\{-ce^{\frac{\delta\lambda^{b_1}}{2}}\right\},$$

where A_5 is the event that there exists an $i \in I$ such that $s_i - s_{i-1} \geq \exp\{-\delta\lambda^{b_1}/2\}R$. Thus, defining the event A by setting

$$A = \bigcap_{i=1}^{5} A_i^c,$$

combining the above estimates gives us that

$$\mathbf{P}(A) \geq 1 - C \exp\left\{-\lambda^{b_1}/M\right\}. \tag{5.69}$$

We now fix a net $W = (w_j)_j$ of $B(5R^{1/\beta})$ such that

$$B(5R^{1/\beta}) \subseteq \bigcup_j B\left(w_j, \exp\{-\lambda^{b_1}\}R^{1/\beta}\right)$$

and $|W| \asymp \exp\{3\lambda^{b_1}\}$. For $i \in I$, let $w_i \in W$ be such that $|x_i - w_i| \leq \exp\{-\lambda^{b_1}\} R^{1/\beta}$. We now use Wilson's algorithm for all points w_i. On A_2^c, it holds that, for each $i \in I$,

$$P^{w_i}\left(R^{w_i}[0, t_{w_i}] \cap \gamma = \emptyset\right) \leq \exp\{-\eta \lambda^{b_1}\}.$$

Therefore we have

$$\mathbf{P}(B_1) \leq C \exp\{-\eta \lambda^{b_1}\} \exp\{3\delta \lambda^{b_1}\} \leq C \exp\{-\eta \lambda^{b_1}/2\}, \tag{5.70}$$

where B_1 is the event that there exists $i \in I$ such that $R^{w_i}[0, t_{w_i}] \cap \gamma = \emptyset$. Suppose that the event B_1^c occurs. For $i \in I$, write u_i for the first time that R^{w_i} hits γ, and let $z_i = R^{w_i}(u_i)$. On A_3^c, we have that $z_i \in \gamma[s_{i-1}, s_i] \cup \gamma[s_i, s_{i+1}]$, because otherwise γ has a quasi-loop. We define the events B_2 and B_3 by setting

$$B_2 = \left\{\exists i \in I \text{ such that } \operatorname{len}\left(\operatorname{LE}\left(R^{w_i}[0, u_i]\right)\right) \geq \exp\{-\lambda^{b_1}/4\} R\right\},$$

$$B_3 = \left\{\exists i \in I \text{ such that } \left|B_{\mathcal{U}}\left(w_i, \lambda^{-1} R\right)\right| \leq \lambda^{-4} R^{\frac{3}{\beta}}\right\}.$$

Combining the translation invariance of $\mathcal{U}$ with Proposition 5.3.4 ensures that

$$\mathbf{P}(B_2) \leq C \exp\left\{-c e^{\lambda^{b_1}/4}\right\}. \tag{5.71}$$

Moreover, by (5.67) and the translation invariance of the UST again, we have

$$\mathbf{P}(B_3) \leq C e^{-c\lambda^b} \times e^{3\lambda^{b_1}} \leq C e^{-c\lambda^b/2},$$

where we use the fact that $|W| \asymp e^{3\lambda^{b_1}}$ and $b_1 = b/1000$. Defining

$$B = \bigcap_{j=1}^{3} B_j^c,$$

we have proved that

$$\mathbf{P}(A \cap B) \geq 1 - C \exp\left\{-\delta \lambda^{b_1}\right\},$$

where we recall that $\delta > 0$ was defined as in (5.68).

Next, suppose that the event $A \cap B$ occurs. Take $x \in \gamma \cap B(4R^{1/\beta})$. We can then find some $i \in I$ such that $x \in \gamma[s_i, s_{i+1}]$. On A_5^c, we have

$$d_{\mathcal{U}}(x_{i-1}, x_i) \le \exp\{-\delta\lambda^{b_1}/2\}R \text{ and } d_{\mathcal{U}}(x_i, x_{i+1}) \le \exp\{-\delta\lambda^{b_1}/2\}R.$$

Furthermore, on $B_1^c \cap A_3^c \cap B_2^c$, it holds that

$$z_i \in \gamma[s_{i-1}, s_i] \cup \gamma[s_i, s_{i+1}] \text{ and } d_{\mathcal{U}}(w_i, z_i) \le \exp\{-\lambda^{b_1}/4\}R.$$

This implies that $d_{\mathcal{U}}(w_i, x) \le \exp\{-\delta\lambda^{b_1}/4\}R$. If B_3^c also holds, it follows that we have

$$\mu_{\mathcal{U}}\left(B_{\mathcal{U}}\left(x, 2\lambda^{-1}R\right)\right) \ge \lambda^{-4}R^{\frac{3}{\beta}}.$$

Consequently, we have proved that there exist universal constants $C, \delta, b_1 \in (0, \infty)$ such that for all R and λ

$$\mathbf{P}\begin{pmatrix} \mu_{\mathcal{U}}\left(B_{\mathcal{U}}\left(x, \lambda^{-1}R\right)\right) \ge \lambda^{-5}R^{\frac{3}{\beta}} \\ \text{for all } x \in \gamma \cap B(4R^{1/\beta}) \end{pmatrix} \ge 1 - C\exp\left\{-\delta\lambda^{b_1}\right\}. \tag{5.72}$$

Finally, once we get (5.72), the proof of (5.66) can be completed by following the strategy used to prove Lemma 5.4.4 given Lemma 5.4.3. Indeed, thanks to (5.72), we can use a net whose mesh size is exponentially small in λ, which guarantees the exponential bound as in (5.50). The simple modification is left to the reader. □

5.6 Exponential upper tail bound on the volume

Complementing the main result of the previous section, we next establish an exponential tail upper bound on the volume, see Theorem 5.6.2, which improves the polynomial tail upper bound on the volume proved in Proposition 5.4.1. We begin with the following proposition.

Proposition 5.6.1. *There exist constants $\lambda_0 > 1$ and $c, C, a \in (0, \infty)$ such* *that: for all $R \geq 1$ and $\lambda \geq \lambda_0$,*

$$\mathbf{P}\left(B_{\mathcal{U}}\left(0, \lambda^{-1} R^{\beta}\right) \not\subseteq B(R)\right) \leq C \exp\{-c\lambda^{a}\}.$$

In particular, it holds that

$$\mathbf{P}\left(\mu_{\mathcal{U}}\left(B_{\mathcal{U}}\left(0, \lambda^{-1} R^{\beta}\right)\right) \geq R^{3}\right) \leq C \exp\{-c\lambda^{a}\}.$$

Proof. The second inequality immediately follows from the first one. Thus it remains to prove the first inequality. We follow the strategy used in the proof of [20, Theorem 3.1]. We may assume that λ is sufficiently large. It follows from Proposition 5.3.4 that there exist constants C, c and $a_0 > 0$ such that

$$\mathbf{P}\left(T_{R/8} < \lambda^{-1} R^{\beta}\right) \leq C \exp\{-c\lambda^{a_0}\},$$

where again T_r stands for the first time that the ILERW γ exits $B(r)$. Setting $a_1 = a_0/10$, we define a sequence of nets D_k as follows. For $k \geq 1$, set $\delta_k = 2^{-k} \exp\{-\lambda^{a_1}\}$, $\eta_k = (2k)^{-1}$, and k_0 be the smallest integer such that $\delta_{k_0} R < 1$. Defining

$$A_k := B(R) \setminus B\left((1-\eta_k)R\right),$$

let D_k be a set of points in A_k satisfying $|D_k| \asymp \delta_k^{-3}$ and also that

$$A_k \subseteq \bigcup_{w \in D_k} B\left(w, \delta_k R\right).$$

We then perform Wilson's algorithm as follows.

- Let $\mathcal{U}_0 = \gamma$ be the ILERW, which is the root of the algorithm.
- Take $w \in D_1$, and consider the SRW R^w started at w, and run until it hits $\mathcal{U}_0$. We add $\mathrm{LE}(R^w)$ to $\mathcal{U}_0$. We choose another point $w' \in D_1$ and add the loop-erasure of $R^{w'}$, a SRW started at w' and run until it hits the part of the tree already constructed. We perform the same procedure for every point in D_1. Let $\mathcal{U}_1$ be the output tree.

- We perform the same algorithm as above for all points in D_2. Let $\mathcal{U}_2$ be the output tree. Similarly, we define $\mathcal{U}_k$.
- We perform Wilson's algorithm for all points in $\mathcal{U}_{k_0}^c$.

Since $\delta_{k_0} R < 1$, we note that $\partial_i B(R) \subseteq A_{k_0} \subseteq \mathcal{U}_{k_0}$.

Now, take $w \in D_1$, and let N_w be the first time that $\gamma_{\mathcal{U}}(0, w)$ exits $B(R/8)$. Using [136, Proposition 4.4], we see that

$$\mathbf{P}\left(N_w < \lambda^{-1} R^{\beta}\right) \leq C \mathbf{P}\left(T_{R/8} < \lambda^{-1} R^{\beta}\right) \leq C \exp\{-c\lambda^{a_0}\}$$

for each $w \in D_1$. Thus if we define the event F_1 by setting

$$F_1 = \left\{T_{R/8} < \lambda^{-1} R^{\beta}\right\} \cup \bigcup_{w \in D_1} \left\{N_w < \lambda^{-1} R^{\beta}\right\},$$

then it follows that

$$\mathbf{P}(F_1) \leq C \delta_1^{-3} \exp\{-c\lambda^{a_0}\} \leq C \exp\{-c' \lambda^{a_0}\},$$

where we have used the fact that $|D_1| \asymp \delta_1^{-3} \asymp \exp\{3\lambda^{a_1}\}$ and that $a_1 = a_0/10$.

Next, for $b > 0$, we define $G_1^w(b)$ to be the event

$$\left\{ \begin{array}{c} \exists v \in B(2R) \text{ with } \operatorname{dist}\left(v, \gamma_{\mathcal{U}}(w, \infty)\right) \leq \delta_1 R \\ \text{such that } P^v\left(R^v[0, \xi] \cap \gamma_{\mathcal{U}}(w, \infty) = \emptyset\right) \geq \delta_1^b \end{array} \right\},$$

where ξ is the first time that R^v exits $B\left(v, \sqrt{\delta_1} R\right)$. Applying Proposition 5.3.7 to the case that $K = 100$, it holds that there exists $b_0 > 0$ such that

$$\mathbf{P}(G_1^w) := \mathbf{P}\left(G_1^w(b_0)\right) \leq C \delta_1^{50}. \tag{5.73}$$

So, if we define the event $G_1 := \cup_{w \in D_1} G_1^w$, then

$$\mathbf{P}(G_1) \leq C \delta_1^{47}.$$

Suppose that the event $F_1^c \cap G_1^c$ occurs, and perform Wilson's algorithm (see [169]) from all points in D_2. For $w \in D_2$, define

$$H_2^w := \{\gamma_{\mathcal{U}}(w,0) \text{ enters } B(R/2) \text{ before it hits } \mathcal{U}_1\},$$

and let $H_2 = \cup_{w \in D_2} H_2^w$. The event H_2^w implies that R^w enters $B(R/2)$ without hitting $\mathcal{U}_1$. Since the event G_1^c occurs, we see that

$$\mathbf{P}(H_2^w) \leq \left(\delta_1^{b_0}\right)^{c\delta_1^{-1/2}},$$

and thus we have

$$\mathbf{P}(H_2) \leq C\delta_1^{10}.$$

For $w \in D_2$, we then define $G_2^w = G_2^w(b_0)$ to be the event

$$\left\{ \begin{array}{c} \exists v \in B(2R) \text{ with } \operatorname{dist}(v, \gamma_{\mathcal{U}}(w,\infty)) \leq \delta_2 R \\ \text{such that } P^v\left(R^v[0,\xi] \cap \gamma_{\mathcal{U}}(w,\infty) = \emptyset\right) \geq \delta_2^{b_0} \end{array} \right\},$$

where ξ is the first time that R^v exits $B(v, \sqrt{\delta_2}R)$, and b_0 is the constant defined as above (see (5.73) for b_0). Using [148, Lemma 3.2 and Lemma 3.3] once again (with $r = \sqrt{\delta_2}R$ and $s = \delta_2 R$), we have

$$\mathbf{P}(G_2^w) \leq C\delta_2^{50}.$$

Importantly, we can take b_0 depending only on $K = 100$. Define the event G_2 by setting $G_2 := \cup_{w \in D_2} G_2^w$, and then

$$\mathbf{P}(G_2) \leq C\delta_2^{47}.$$

Defining H_k and G_k, $k \geq 3$ similarly, it follows that

$$\mathbf{P}(H_k \cup G_k) \leq C\delta_k^{47}.$$

Finally, we define

$$J = F_1^c \cap G_1^c \cap \bigcap_{k=2}^{k_0} (H_k^c \cap G_k^c).$$

On the event J, we have the following.

- For all $k = 1, 2, \dots k_0$ and every $w \in D_k$, the first time that $\gamma_{\mathcal{U}}(0, w)$ exits $B(R/8)$ is greater than $\lambda^{-1} R^{\beta}$.
- The set D_{k_0} disconnects 0 and $B(R)^c$.

Thus, on the event J, it holds that $B_{\mathcal{U}}\left(0, \lambda^{-1} R^{\beta}\right) \subseteq B(R)$. Since

$$\mathbf{P}(J^c) \le C \exp\{-c' \lambda^{a_0}\} + C \sum_{k=1}^{k_0} \delta_k^{47} \le C \exp\{-\lambda^{a_1}\},$$

we have thus completed the proof. $\square$

We are now ready to establish the main result of the section.

Theorem 5.6.2. *There exist constants* $\lambda_0 > 1$ *and* $c', C', a' \in (0, \infty)$ *such that: for all* $R \ge 1$ *and* $\lambda \ge \lambda_0$,

$$\mathbf{P}\left(\max_{z \in B(R^{1/\beta})} \mu_{\mathcal{U}}\left(B_{\mathcal{U}}\left(z, \lambda^{a'} R\right)\right) \ge \lambda R^{3/\beta}\right) \le C' \exp\{-c' \lambda^{a'}\}. \tag{5.74}$$

Proof. Since the proof is very similar to that of Theorem 5.5.2, we will only explain how to modify it here. Also, we will use the same notation used in the proof of Theorem 5.5.2. Proposition 5.6.1 tells that there exist constants $c, C, b \in (0, \infty)$ such that

$$\mathbf{P}(A_1') \le C \exp\{-c\lambda^b\}, \tag{5.75}$$

where $A_1' := \{\mu_{\mathcal{U}}(B_{\mathcal{U}}(0, \lambda R)) \le \lambda^{10} R^{\frac{3}{\beta}}\}$. In this proof, we choose the constant b in this way, and let $b_1 = b/1000$. Using this constant b_1, we define the events $A_2, \dots, A_5$ as in the proof of Theorem 5.5.2. Let $A = (A_1')^c \cap (\cap_{i=2}^5 A_i^c)$ so that

$$\mathbf{P}(A) \ge 1 - C \exp\left\{-\lambda^{b_1}/M\right\},$$

see (5.69). We also recall the events B_1 and B_2 defined in the proof of

Theorem 5.5.2, for which

$$\mathbf{P}(B_1 \cup B_2) \leq C \exp\{-\eta\lambda^{b_1}/2\},$$

see (5.70) and (5.71). Moreover, let

$$B_3' = \left\{\exists i \in I \text{ such that } \mu_{\mathcal{U}}\left(B_{\mathcal{U}}\left(w_i, \lambda R\right)\right) \geq \lambda^{10} R^{\frac{3}{\beta}}\right\}.$$

Combining (5.75) with the translation invariance of the UST, we have

$$\mathbf{P}(B_3') \leq Ce^{-c\lambda^b} \times e^{3\lambda^{b_1}} \leq Ce^{-c\lambda^b/2},$$

where we have also used the fact that $|W| \asymp e^{3\lambda^{b_1}}$ and $b_1 = b/1000$. Setting $B = B_1^c \cap B_2^c \cap (B_3')^c$, we then have that

$$\mathbf{P}(B) \geq 1 - C\exp\{-\eta\lambda^{b_1}/4\}.$$

Now, suppose that the event $A \cap B$ occurs, and let $x \in \gamma \cap B(4R^{1/\beta})$. We can then find some $i \in I$ such that $x \in \gamma[s_i, s_{i+1}]$. Since A_5^c holds, we have

$$d_{\mathcal{U}}(x_{i-1}, x_i) \leq \exp\{-\delta\lambda^{b_1}/2\}R \text{ and } d_{\mathcal{U}}(x_i, x_{i+1}) \leq \exp\{-\delta\lambda^{b_1}/2\}R.$$

Furthermore, since $B_1^c \cap A_3^c \cap B_2^c$ holds, we have that

$$z_i \in \gamma[s_{i-1}, s_i] \cup \gamma[s_i, s_{i+1}] \text{ and } d_{\mathcal{U}}(w_i, z_i) \leq \exp\{-\lambda^{b_1}/4\}R.$$

This implies that $d_{\mathcal{U}}(w_i, x) \leq \exp\{-\delta\lambda^{b_1}/4\}R$. Given $(B_3')^c$ also holds, we therefore have

$$\mu_{\mathcal{U}}\left(B_{\mathcal{U}}\left(x, \lambda R/2\right)\right) \leq \lambda^{10} R^{\frac{3}{\beta}}.$$

Consequently, we have proved that there exist universal constants $C, \delta, b_1 \in (0, \infty)$ such that: for all R and λ,

$$\mathbf{P}\left(\begin{array}{c} \mu_{\mathcal{U}}\left(B_{\mathcal{U}}\left(x, \lambda R/2\right)\right) \leq \lambda^{10} R^{\frac{3}{\beta}} \\ \text{for all } x \in \gamma \cap B(4R^{1/\beta}) \end{array}\right) \geq 1 - C\exp\left\{-\delta\lambda^{b_1}\right\}. \tag{5.76}$$

Similarly to the comment at the end of the proof of Theorem 5.5.2, given (5.76), the proof of (5.74) follows by applying the same strategy as that used to prove Lemma 5.4.4 given Lemma 5.4.3. Indeed, given (5.76), we can use a net whose mesh size is exponentially small in λ, which guarantees the exponential bound as in (5.74). The simple modification is again left to the reader. $\square$

5.7 Convergence of finite-dimensional distributions

As noted in the introduction, the existence of a scaling limit for the three-dimensional LERW was first demonstrated in [107]. The work in [107] established the result in the Hausdorff topology, and this was recently extended in [127] to the uniform topology for parameterized curves. Whilst the latter seems a particularly appropriate topology for understanding the scaling limit of the LERW, the results in [107, 127] are restrictive when it comes to the domain upon which the LERW is defined. More specifically, we say that a LERW is defined in a domain D if it starts in an interior point of D and ends when it reaches the boundary of D. The assumptions in [107] cover the case of LERWs defined in domains with a polyhedral boundary, while [127] requires the domain to be a ball or the full space.

In this section, we extend the existence of the scaling limit to LERWs defined in the domain $\mathbb{R}^3 \setminus \cup_{j=1}^{K} \operatorname{tr} \mathcal{K}_j$, where each $\mathcal{K}_j$ is itself a path of the scaling limit of a LERW. Once we gain this level of generality, we use Wilson's algorithm to obtain the convergence in distribution of rescaled subtrees of the UST (see Figure 5.8 for an example realisation of the subtree spanning a finite collection of points). This will be crucial for establishing the convergence part of Theorem 5.1.1. We begin by introducing some notation for subtrees.

5.7.1 Parameterized trees

A parameterized tree is an encoding for an infinite tree embedded in the closure of $\mathbb{R}^3$. This encoding is specialized for infinite trees with a finite

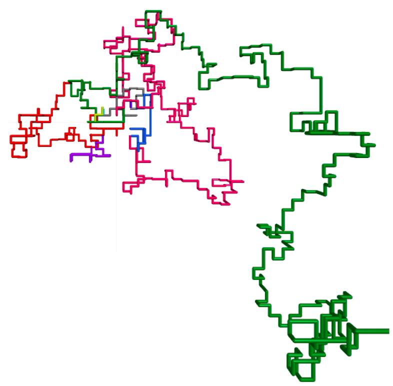

Figure 5.8: A realisation of a subtree of the UST of $\delta\mathbb{Z}^3$ spanned by 0 and the corners of the cube $[-1,1]^3$. The tree includes part of its path towards infinity (in green). Colours indicate different LERWs used in Wilson's algorithm.

number of spanning points and one end. More precisely, a **parameterized tree** $\mathscr{T}$ with K spanning points is defined as $\mathscr{T} = (X, \Gamma)$ where:

1. $X = \{x(1), \ldots, x(K)\} \subset \mathbb{R}^3$ are the spanning (or distinguished) points; and

2. $\gamma^{x(i)}$ is a transient parameterized (simple) curve starting at $x(i)$, and

$$\Gamma = \{\gamma^{x(i)} \colon 0 \leq i \leq K\}.$$

   We require that for any pair i, j there exist **merging times** $s^{i,j}, s^{j,i} \geq 0$ satisfying

   (a) $\gamma^{x(i)}|_{[s^{i,j},\infty)} = \gamma^{x(j)}|_{[s^{j,i},\infty)}$; and

   (b) $\operatorname{tr} \gamma^{x(i)}|_{[0,s^{i,j})} \cap \operatorname{tr} \gamma^{x(j)}|_{[0,s^{j,i})} = \emptyset$.

Let $\mathscr{F}^K$ be the space of parameterized trees with K distinguished points. We endow $\mathscr{F}^K$ with the distance

$$d_{\mathscr{F}^K}\left(\mathscr{T}, \tilde{\mathscr{T}}\right) := \max_{1 \leq i \leq K} \left\{\chi\left(\gamma^{x(i)}, \tilde{\gamma}^{\tilde{x}(i)}\right)\right\} + \max_{1 \leq i,j \leq K} \{|s^{i,j} - \tilde{s}^{i,j}|\},$$

for $\mathscr{T} = (X, \Gamma)$, $\tilde{\mathscr{T}} = (\tilde{X}, \tilde{\Gamma}) \in \mathscr{F}^K$.

We write

$$\operatorname{tr} \mathscr{T} = \bigcup_{\gamma \in \Gamma} \operatorname{tr} \gamma$$

for the trace of a parameterized tree.

Proposition 5.7.1. *Let $\mathscr{T}$ be a parameterized tree. Then $\operatorname{tr} \mathscr{T}$ is a topological tree with one end. Additionally, for any $z, w \in \operatorname{tr} \mathscr{T}$ there exists a unique curve from z to infinity on $\mathscr{T}$, denoted by γ^z and a unique curve from z to w in $\mathscr{T}$ denoted by $\gamma^{z,w}$.*

Proof. The set $\operatorname{tr} \mathscr{T}$ is path-connected as a consequence of condition (2a) in the definition of a parameterized tree. It is also one-ended, since $\cap_{i=1}^K \gamma^{x(i)}$ is a single parameterized curve towards infinity.

The main task in this proof is to show that there cannot be cycles embedded in $\operatorname{tr} \mathscr{T}$. We proceed by contradiction. Let S^1 be the circle and assume that $\varphi : S^1 \to \operatorname{tr} \mathscr{T}$ is an injective embedding. Since every curve in Γ is simple and φ is injective, then $\varphi(S^1)$ intersects at least two different curves, say $\gamma^{x(i)}$ and $\gamma^{x(j)}$. From the definition of merging times, we see that $T^2 = \operatorname{tr} \gamma^{x(i)} \cup \operatorname{tr} \gamma^{x(j)}$ is homeomorphic to $([0, \infty) \times \{0\}) \cup (\{1\} \times [0, 1])$, but the latter space cannot contain a embedding of S^1. It follows that $\varphi(S^1)$ intersects at least a third curve $\gamma^{x(\ell)}$. We assume that $\varphi(S^1)$ is contained in $T^3 = \operatorname{tr}(\gamma^{x(i)}) \cup \operatorname{tr}(\gamma^{x(j)}) \cup \operatorname{tr}(\gamma^{x(\ell)})$. Under the last assumption, it is necessary that $\gamma^{x(\ell)}$ intersects $\gamma^{x(i)}$ and $\gamma^{x(j)}$ before these last two curves merge (otherwise the case is similar to T^2). Denote the intersection times by $t^{\ell,i}$ and $t^{i,\ell}$, so $\gamma^{x(\ell)}(t^{\ell,i}) = \gamma^{x(i)}(t^{i,\ell})$. We use the same notation for $\gamma^{x(j)}$. Then, we have that $t^{i,\ell} < s^{i,j}$ and $t^{j,\ell} < s^{j,i}$. Without loss of generality, $t^{\ell,i} < t^{\ell,j}$. However, it is easy to verify that $t^{\ell,i}$ is not the merging time $s^{\ell,i}$, since $\gamma^{x(\ell)}|_{[t^{\ell,i},\infty)}$ does not merge with $\gamma^{x(i)}$ at that point. Therefore $s^{\ell,i}$ does not exist, and this conclusion contradicts the definition of Γ. It follows that $\varphi(S^1)$ is not contained in T^3, but it intersects more curves, e.g. all of them in $\operatorname{tr} \mathscr{T} = \cup_{i=1}^K \operatorname{tr}((\gamma^{x(i)}))$. However, the argument that we used for T^3 also applies to $\operatorname{tr} \mathscr{T}$. We conclude that the embedding φ does not exist.

Finally, observe that $\operatorname{tr} \mathscr{T}$ is one ended and all curves in Γ are parameterized towards infinity. It is then straightforward to define γ^z and $\gamma^{z,w}$. $\square$

A corollary of Proposition 5.7.1 is that the intrinsic distance in $\operatorname{tr} \mathscr{T}$ is well-defined. It is given by

$$d_{\mathscr{T}}(z,w) := T(\gamma^{z,w}), \qquad z,w \in \operatorname{tr} \mathscr{T}$$

where $T(\cdot)$ is the duration of a curve.

We will consider restrictions of parameterized trees to balls centred at the origin. For a parameterized tree $\mathscr{T} = (X,\Gamma)$, let $R \geq 1$ be large enough so that $X \subseteq B_E(R)$. We restrict each curve in Γ to $\gamma^{x(i)}|^R$ (where the restriction to the ball of radius $R > 0$ is in the sense described in Subsection 5.3.2), and define the restriction of a parameterized tree to $B_E(R)$ as the subset of $\mathbb{R}^3$

$$\mathscr{T}|^R := \bigcup_{\gamma \in \Gamma} \operatorname{tr} \gamma^{x(i)}|^R.$$

Note that $\mathscr{T}|^R$ may not be connected for some values of $R > 0$. But for R large enough, $\mathscr{T}|^R$ is a topological tree (Figure 5.9 gives an example of both cases).

5.7.2 The scaling limit of subtrees of the UST

We introduce the main results of this section.

Let $\mathcal{U}_n$ be the uniform spanning tree on $2^{-n}\mathbb{Z}^3$. We are interested in subtrees of $\mathcal{U}_n$ spanned by K distinguished points. Let $x(1), \ldots, x(K)$ be different points in $\mathbb{R}^3$ and let $X_n = \{x_n(1), \ldots, x_n(K)\}$ be a subset of $2^{-n}\mathbb{Z}^3$ such that $x_n(i) \to x(i)$ as $n \to \infty$, for each $i = 1, \ldots, n$. Denote by $\gamma_n^{x_n(i)}$ the transient path in $\mathcal{U}_n$ starting at $x_n(i)$ and parameterized by path length. We set $\bar{\gamma}_n^{x(i)}$ to be the β-parameterization of $\gamma_n^{x_n(i)}$ and $\Gamma_n = \{\gamma_n^{x_n(i)} \colon 1 \leq i \leq K\}$. Then $\mathscr{S}_n^K = (X, \Gamma_n)$ is the parameterized tree corresponding to the subtree of $\mathcal{U}_n$ spanned by $x_n(1), \ldots, x_n(K)$ and the point at infinity.

Theorem 5.7.2. *The sequence of parameterized trees $(\mathscr{S}_n^K)_{n \in \mathbb{N}}$ converges weakly to $\hat{\mathscr{S}}^K$ in the space $\mathscr{F}^K$ as $n \to \infty$.*

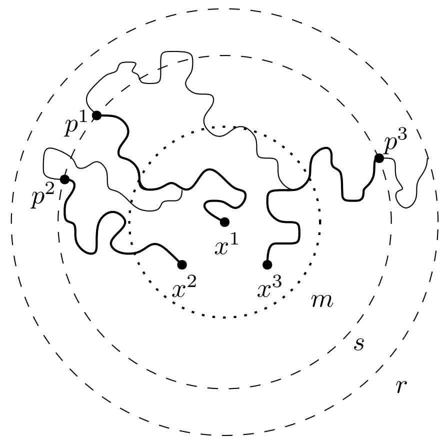

Figure 5.9: $\mathscr{S}$ is a parameterized tree with spanning points $x(1), x(2)$ and $x(3)$. The restriction $\mathscr{S}|^s$ is the union of the paths between $x(i)$ and p^i, with $i = 1, 2, 3$. In this example, $\mathscr{S}|^r$ and $\mathscr{S}|^s$ are different inside the radius m. A crucial difference between these two sets is that $\mathscr{S}|^r$ is connected, but $\mathscr{S}|^s$ is disconnected.

The proof of Theorem 5.7.2 relies on the convergence of the branches of the uniform spanning tree as they appear in Wilson's algorithm. In the next section, we control the behaviour of a LERW before it hits an approximation of a parameterized tree. These loop-erased random walks correspond to the branches of the UST. Then Proposition 5.7.8 shows that convergence of such branches implies convergence of parameterized trees. After these arguments, we are prepared for the proof of Theorem 5.7.2. We present it in Subsection 5.7.5.

Conversely, Proposition 5.7.6 shows that convergence of parameterized trees implies the convergence of the intrinsic distance. We thus get the following corollary of Theorem 5.7.2.

Corollary 5.7.1. Let $(x_\delta(i))_{i=1}^K$ be a collection of points in $\delta\mathbb{Z}^3$ such that $x_\delta(i) \to x(i)$, for all $i = 1, \dots, K$, for some collection of distinct points $(x(i))_{i=1}^K$ in $\mathbb{R}^3$. Along the subsequence $\delta_n = 2^{-n}$, it holds that

$$\left(\delta_n^\beta d_{\mathcal{U}}\left(x_{\delta_n}(i), x_{\delta_n}(j)\right)\right)_{i,j=1}^K$$

converges in distribution.

5.7.3 Parameterized trees and random walks

We begin with a property on the hittability of a parameterized tree. We say that a parameterized tree $\mathscr{T}_\delta$ is on the scaled lattice $\delta\mathbb{Z}^3$ if each one of its curves defines a path on $\delta\mathbb{Z}^3$.

Definition 5.7.3. Let $\delta \in (0,1)$, $R \geq 1$ and $\varepsilon \in (0,1)$ and let $\mathscr{T}_\delta$ be a parameterized tree on $\delta\mathbb{Z}^3$. We say that $\mathscr{T}_\delta$ is η**-hittable** in $B_\delta(0,R)$ if the following event occurs:

$$H(\mathscr{T}_\delta, \varepsilon; \eta) := \left\{ \begin{array}{c} \forall x \in B_\delta(0,R) \text{ with } \operatorname{dist}(x, \mathscr{T}_\delta) \leq \varepsilon^2, \\ P^x\left(S^x\left[0, \xi_S(B_\delta(x, \varepsilon^{1/2}))\right] \cap \mathscr{T}_\delta = \emptyset\right) \leq \varepsilon^\eta \end{array} \right\}.$$

In the definition above, recall that $\xi_S(B_\delta(x,\varepsilon^{1/2}))$ stands for the first exit time from the δ-scaled discrete ball $B_\delta(x,\varepsilon^{1/2})$.

Proposition 5.7.4. *There exist constants $\eta > 0$ and $C < \infty$ such that: if $\mathscr{S}_\delta^K$ is a parameterized subtree of the uniform spanning tree on $\delta\mathbb{Z}^3$ with K spanning points for all $\delta \in (0,1)$, $R \geq 1$ and $\varepsilon > 0$,*

$$\mathbf{P}\left(H(\mathscr{S}_\delta^K|^R, \varepsilon; \eta)\right) \geq 1 - CKR^3\varepsilon.$$

Proof. Recall that any path towards infinity in the uniform spanning tree is equal, in distribution, to a ILERW. Then, the probability that $x \in B_\delta(0,R)$ hits the tree $\mathscr{S}_\delta^K|^R$ is at least the probability that x hits a restricted ILERW, where such restriction is up to the first exit of the LERW from $B_\delta(0,R)$. Then Proposition 5.7.1 is a consequence of Proposition 5.3.7. □

Remark. The proof of Proposition 5.7.9 can be generalized to any subset of $\mathbb{R}^3$ that is η-hittable with high probability. We restrict to the case of parameterized subtrees for clarity, and because it is the most relevant for our purposes. To further increase the clarity of the proof of Proposition 5.7.9, the reader can think of the subtree $\mathscr{S}_n^K$ as consisting of a single ILERW.

5.7.4 Essential branches of parameterized trees

Let $\mathscr{T} = (X, \Gamma)$ be a parameterized tree. For a leaf $x(i) \in X$ with $i > 1$, let

$$y(i) := \operatorname{tr} \gamma^{x(i)} \cap \bigcup_{j=1}^{i-1} \operatorname{tr} \gamma^{x(j)} \tag{5.77}$$

be the intersection point of $\gamma^{x(i)}$ with any of the curves with an smaller index. We define $y(1)$ to be the point at infinity and say $y(i)$ is a **branching point**. When we compare (5.77) with conditions (2a) and (2b) in the definition of parameterized tree, we see that

$$y(i) = \gamma^{x(i)}(s^{i,m(j)}),$$

where $s^{i,m(j)} = \min_{j<i}\{s^{i,j}\}$ is the first merging time.

The parameterized curves $\gamma^{x(i),y(i)}$ are called **essential branches** for $i = 1, \ldots, K$. Note that $\gamma^{x(1),y(1)}$ is the transient curve $\gamma^{x(1)} \in \mathcal{C}$, while $\gamma^{x(i),y(i)} \in \mathcal{C}_f$ for $i = 1, \ldots, K$. We denote the set of essential branches by $\Gamma^e(\mathscr{T}) := \{\gamma^{x(i),y(i)}\}_{1 \le i \le K}$.

Proposition 5.7.5. *Assume that $\mathscr{T}_n \to \mathscr{T}$ in the space of parameterized trees $\mathscr{F}^K$. Then*

$$\gamma_n^{x_n(1),y(1)} \to \gamma^{x_n(1),y(1)} \qquad \text{as } n \to \infty$$

in the space $\mathcal{C}$. For $i = 2, \ldots, K$, the essential branches and the curves between branching points converge:

$$\gamma_n^{x_n(i),y_n(i)} \to \gamma^{x(i),y(i)}, \qquad \gamma_n^{y_n(i),y_n(j)} \to \gamma^{y(i),y(j)} \qquad \text{as } n \to \infty$$

in the space of finite parameterized curves $\mathcal{C}_f$.

Proof. The convergence of the first essential branch is immediate from the definition of the metric $d_{\mathscr{F}^K}$, since $\gamma_n^{x_n(1),y(1)} = \gamma_n^{x_n(1)}$.

To prove the convergence of the other essential branches, and the curves between branching points, we first need to show that spanning and branching points converge.

Each $x_n(i) \in X$ is the initial point of a curve in Γ and hence convergence in the space of parameterized trees implies that $x_n \to x$ as $n \to \infty$. Now we consider a branching point $y_n(i)$, with $i = 2, \ldots, K$. Recall that $y_n(i) = \gamma_n^{x(i)}(s_n^{i,m(j)})$, where $s_n^{i,m(j)} = \min_{j<i}\{s_n^{i,j}\}$. Since convergence of the parameterized trees $\mathscr{T}$ imply convergence of the merging times $s_n^{i,j} \to s^{i,j}$ as $n \to \infty$, then, for the sequence of minima, $s_n^{i,m(j)} \to s^{i,m(j)}$. With an application of Proposition 5.3.2 (c), we get convergence of the branching points $y_n(i) = \gamma_n^{x(i)}(s_n^{i,m(j)}) \to \gamma^{x(i)}(s^{i,m(j)}) = y(i)$.

With convergence of both the spanning and branching points, Proposition 5.3.1 and Proposition 5.3.2 imply that the corresponding restrictions of $\gamma^{x(i)}$ converge. □

Proposition 5.7.6. *Assume that $\mathscr{T}_n = (X_n, \Gamma_n)$ converges to $\mathscr{T}$ in the space $\mathscr{F}^K$. If the corresponding collections of spanning points are $X_n = \{x_n(1), \ldots, x_n(K)\}$ and $X = \{x(1), \ldots, x(K)\}$, then*

$$(d_{\mathscr{T}_n}(x_n(i), x_n(j)))_{1\le i,j\le K} \to (d_{\mathscr{T}}(x(i), x(j)))_{1\le i,j\le K} \tag{5.78}$$

as $n \to \infty$.

Proof. Proposition 5.7.1 shows that restriction, concatenation and time-reversal of the curves in Γ_n define $\gamma_n^{x_n(i),x_n(j)}$. In fact,

$$\gamma_n^{x_n(i),x_n(j)} = \gamma_n^{x_n(i),y_n(i)} \oplus \gamma_n^{y_n(\ell_1),y_n(\ell_2)} \oplus \ldots \oplus \gamma_n^{y_n(\ell_{m-1}),y_n(\ell_m)} \oplus \gamma_n^{y_n(j)x_n(j)}, \tag{5.79}$$

where $\ell_1 = i$ and $\ell_m = j$. Then Proposition 5.7.5 implies the convergence of each essential branch, and Proposition 5.3.1 and Proposition 5.3.2 imply the convergence of $(\gamma_n^{x_n(i),x_n(k)})_{n\in\mathbb{N}}$. In particular, the duration of each curve in (5.79) converges and we get (5.78). □

Conversely, we can reconstruct a tree from a set of essential branches.

Proposition 5.7.7. *Let $X = \{x(1), \ldots, x(K)\} \subset \mathbb{R}^3$ and consider a collection of curves with the following conditions:*

(a) Let $\gamma^{x(1),\tilde{y}(1)}$ be a transient parameterized curve starting at $x(1)$; recall that $\tilde{y}(1)$ denotes the point at infinity.

(b) For $i = 2, \dots n$, $\gamma^{x(i),\tilde{y}(i)}$ is a parameterized curve starting at $x(i)$ and ending at $\tilde{y}(i)$, where the endpoint $\tilde{y}(i)$ is the first hitting point to $\bigcup_{j=1}^{i-1} \operatorname{tr} \gamma^{x(i),\tilde{y}(i)}$.

Then $\{\gamma^{x(i),\tilde{y}(i)}\}_{1\le i\le K}$ defines a set of transient curves $\Gamma = \{\gamma^{x(i)}\}_{1\le i\le K}$ and a parameterized tree $\mathscr{T} = (X, \Gamma)$.

Proof. First we to construct Γ from the collection of curves $\{\gamma^{x(i),\tilde{y}(i)}\}_{1\le i\le K}$. Note that $\gamma^{x(1),\tilde{y}(1)}$ is already a transient curve starting at $x(1)$. We construct the other elements in Γ recursively. Assume that $\gamma^{x(1)}, \dots \gamma^{x(i-1)}$ have been defined and satisfy conditions (2a) and (2b) in the definition of parameterized tree. Recall that the endpoint of $\gamma^{x(i),\tilde{y}(i)}$ is $\tilde{y}(i)$, and this point intersects some $\gamma^{x(j)}$ with $j < i$. Then

$$\gamma^{x(i)} = \gamma^{x(i),\tilde{y}(i)} \oplus \gamma^{x(j)}|_{[\tilde{y}(j),\infty)}.$$

Since the endpoint of $\gamma^{x(i),\tilde{y}(i)}$ is the first hitting point to $\bigcup_{j=1}^{i-1} \operatorname{tr} \gamma^{x(i),\tilde{y}(i)}$, we have that $(\operatorname{tr} \gamma^{x(i)}|_{[x(i),\tilde{y}^i)} \cap \operatorname{tr} \gamma^{x(j)}) = \emptyset$ for $j < i$. This construction ensures that $\gamma^{x(i)}$ satisfies conditions (2a) and (2b), when we compare it against curves with smaller indexes. We continue with this construction for $i = 2, \dots K$ to define Γ. Therefore $\mathscr{T} = (X, \Gamma)$ is a parameterized tree. Finally, note that $\tilde{y}(i)$ satisfies (5.77) and hence $\tilde{y}(i) = y(i)$, for $i = 2, \dots, K$. □

Proposition 5.7.8. *Let $(\mathscr{T}_n)_{n\in\mathbb{N}}$ be a sequence of parameterized trees with essential branches $\Gamma^e(\mathscr{T}_n) = \{\gamma_n^{x_n(i),y_n(i)}\}_{1\le i\le K}$. Assume that*

$$\left(\gamma_n^{x_n(i),y_n(i)}\right)_{1\le i\le K} \to \left(\gamma^{x(i),y(i)}\right)_{1\le i\le K} \tag{5.80}$$

in the product topology as $n \to \infty$ and $\{\gamma^{x(i),y(i)}\}$ satisfy the conditions in Proposition 5.7.7. Then $(\mathscr{T}_n)_{n\in\mathbb{N}}$ converges in the metric space $\mathscr{F}^K$ to a parameterized tree $\mathscr{T}$ for which $\Gamma^e(\mathscr{T}) = \{\gamma^{x(i),y(i)}\}_{0\le i\le K}$ is a set of essential branches.

Proof. Convergence of $(\gamma_n^{x_n(i),y_n(i)})_{0\le i\le K}$ in the product topology implies that each element in Γ_n^e converges. Proposition 5.7.7 shows that every curve $\gamma_n^{x_n(i)}$ is the concatenation of sub-curves of Γ^e. Moreover, $\{\gamma^{x(i),y(i)}\}_{i=1\dots K}$

satisfy the conditions in Proposition 5.7.7 and hence they define a parameterized tree $\mathscr{T}$ with $\Gamma = (\gamma^{x(i)})$. Finally, (5.80) implies the convergence of the branching points $y_n(i)$, and from here we get convergence of the merging times. Then, Proposition 5.3.1 and Proposition 5.3.2 imply that $\chi(\gamma_n^{x_n(i)}, \gamma^{x(i)}) \to 0$ as $n \to \infty$. Therefore $d_{\mathscr{F}^K}(\mathscr{T}_n, \mathscr{T}) \to 0$ as $n \to \infty$. $\square$

5.7.5 Proof of Theorem 5.7.2

The proof of Theorem 5.7.2 is by mathematical induction. The convergence in the scaling limit of the ILERW provides the base case. We state the inductive step in Proposition 5.7.9.

Proposition 5.7.9. *Let $\mathcal{U}_n$ be the uniform spanning tree on $2^{-n}\mathbb{Z}^3$. Let $(x_n(i))_{i=1,\dots,K+1}$ be a set of vertices in $2^{-n}\mathbb{Z}^3$ and assume that $x_n(i)$ converges to $x(i) \in \mathbb{R}^3$ as $n \to \infty$. Let $\bar{\gamma}_n^{x_n(i)}$ be the β-parameterization of the transient path in $\mathcal{U}_n$ starting at $x_n(i)$ and directed towards infinity. Assume that $(\bar{\gamma}_n^{x_n(i)})_{i=1,\dots,K}$ converges weakly as a parameterized tree to $\hat{\mathscr{S}}^K$. Then $(\bar{\gamma}_n^{x_n(i)})_{i=1,\dots,K+1}$ converges weakly to a parameterized tree $\hat{\mathscr{S}}^{K+1}$, with respect to the metric $\mathscr{F}^{K+1}$ for parameterized trees.*

We devote the rest of this section to the proof of Proposition 5.7.9. It is based in Proposition 5.7.8. According to the latter proposition, it suffices to prove convergence of the essential branches with respect to the product topology. We then shift our attention to the essential branches of an infinite subtree of the uniform spanning tree. Wilson's algorithm provides a natural construction of them; and we present it below. Subsection 5.7.6 develops the arguments for the proof of Proposition 5.7.9.

Let $\mathcal{U}_n$ be the uniform spanning tree on $2^{-n}\mathbb{Z}^3$. Let $x_n(i) \in \mathbb{Z}^3$ and $x(i) \in \mathbb{R}^3$ as in the statement of Proposition 5.7.9, so $x_n(i) \to x(i)$ as $n \to \infty$ for $i = 1, \dots, K+1$. Now we apply Wilson's algorithm on the scaled lattice $2^{-n}\mathbb{Z}^3$.

- Let γ_n^1 be an ILERW starting at $x_n(1)$, and

$$\bar{\gamma}_n^{x(1),y(1)}(t) = \gamma_n^1(2^{\beta n}t), \qquad \forall t \geq 0.$$

be its β-parameterization. Note that we omit the sub-index n on $x(1)$ and $y(1)$ to ease the notation. This transient curve is the first branch of the parameterized tree.

- Let γ_n^i be the loop-erased random walk started at $x_n(i)$, and stopped when it hits any of the previous loop-erased random walks $\gamma_n^1, \ldots, \gamma_n^{i-1}$. Let $y_n(i) \in 2^{-n}\mathbb{Z}^3$ be the hitting point, and set

$$\bar{\gamma}_n^{x(i),y(i)} = \gamma_n^i(2^{\beta n} t), \qquad \forall t \in [0, 2^{-\beta n} \operatorname{len}(\gamma_n^i)].$$

  The duration of the curve $\bar{\gamma}_n^{x(i),y(i)}$ is $2^{-\beta n} \operatorname{len}(\gamma_n^i)$, i.e. the length of the path γ_n^i with the appropriate scaling. We also omit the sub-index n on $x(i)$ and $y(i)$ when they appear in the curve $\bar{\gamma}_n^{x(i),y(i)}$.

Set $X_n = \{x_n(1), \ldots, x_n(K)\}$ and $\Gamma_n^e = \{\bar{\gamma}_n^{x(i),y(i)} : i = 1, \ldots, K\}$. By Proposition 5.7.7, X_n and Γ_n^e determine a parameterized tree $\mathscr{S}_n^K$, and Wilson's algorithm shows that $\operatorname{tr} \mathscr{S}_n^K$ is equal in distribution to the subtree of $\mathcal{U}_n$ spanned by X_n and the point at infinity.

As part of the proof of Theorem 5.7.2, we will show that the limit of parameterized trees $\hat{\mathscr{S}}^K$ has the following (formal) representation (see Lemma 5.7.19). The next construction is Wilson's algorithm, but in this case, the branches have the distribution of the scaling limit of the ILERW.

- Let $\hat{\gamma}^{x(1),y(1)} \in \mathcal{C}$ be the scaling limit of ILERW starting at $x(1)$, endowed with the natural parameterization, see [127].

- Let $\hat{\gamma}^{x(i),y(i)} \in \mathcal{C}_f$ be the scaling limit of LERW started at $x(i)$, and stopped when it hits any of $\hat{\gamma}^{x(1),y(1)}, \ldots, \hat{\gamma}^{x(i-1),y(i-1)}$. (Our construction will give that this hitting time is finite, see Lemma 5.7.18.) Here we denote the hitting point by $y(i)$.

Set $X = \{x(1), \ldots, x(K)\}$ and $\hat{\Gamma}^e = \{\hat{\gamma}^{x(i),y(i)}\}_{1 \le i \le K}$. Proposition 5.7.7 defines the parameterized tree $\hat{\mathscr{S}}^K = (X, \hat{\Gamma})$ with set of essential branches $\Gamma^e(\mathscr{S}^K) = \hat{\Gamma}^e$.

5.7.6 Proof of Proposition 5.7.9

The next proposition allows us to work with restrictions of parameterized trees, when we compare them within a smaller subset.

Proposition 5.7.10. *Let $\mathscr{S}_\delta^K$ be a parameterized subtree of the uniform spanning tree on $\delta\mathbb{Z}^3$ with K spanning points. Assume that $|x| < m$ for each $x \in K$. Then for $r > s \geq m^2 > 0$,*

$$\mathbf{P}\left((\mathscr{S}_\delta^K|^r \triangle \mathscr{S}_\delta^K|^s) \cap B_E(m) \neq \emptyset\right) \leq K\delta m^{-1}[1 + O(m^{-1})]. \tag{5.81}$$

Proof. The restrictions $\mathscr{S}_\delta^K|^r$ and $\mathscr{S}_\delta^K|^s$ are different inside $B_E(m)$ when a path returns to $B_E(m)$ after its first exit from $B_E(r)$; we refer to Figure 5.9 as an example of this situation. By virtue of Wilson's algorithm and a union bound, the probability on (5.81) is bounded above by the probability of return to $B(m\delta^{-1})$ of K simple random walks on $\mathbb{Z}^3$:

$$K \sup_{x \in \partial B(s\delta^{-1})} P_S^x\left(\tau_S(B(m\delta^{-1})) < \infty\right),$$

where P_S^x indicates the probability measure of a simple random walk on $\mathbb{Z}^3$ started at x, and $\tau_S(B(m\delta^{-1}))$ is the first time that the random walk S hits the ball $B(m\delta^{-1})$. Therefore, the upper bound in (5.81) follows from well-known estimates on the return probability for the simple random walk, see e.g. [117, Proposition 6.4.2]. $\square$

The proof of Proposition 5.7.9 is divided into a sequence of lemmas, and these are grouped into five steps. The final and sixth step finishes the proof.

Step 1: set-up.

We begin with the set-up of the proof. First note that the assumptions of Proposition 5.7.9 indicate that $(\bar{\gamma}_n^{x(i)})_{1 \leq i \leq K}$ converges in distribution. From now on, we work in the coupling given by Skorohod's embedding theorem where $(\bar{\gamma}_n^{x(i)})_{1 \leq i \leq K}$ converges to a collection of continuous curves $(\hat{\gamma}^{x(i)})_{1 \leq i \leq K}$, almost surely.

Let $S_n = (S_n(t))_{t\in\mathbb{N}}$ be an independent random walk on $\delta_n\mathbb{Z}^3$ starting at $x_n(K+1)$. Consider the hitting time of the parameterized tree $\mathscr{S}_n^K$, as given by

$$\xi_n^{\mathscr{S}} = \inf\left\{t \geq 0 : S_n(t) \cap \operatorname{tr} \mathscr{S}_n^K \neq \emptyset\right\}.$$

We let $\gamma_n = \operatorname{LE}\left(S_n[0, \xi^{\mathscr{S}}]\right)$ be the corresponding LERW from $x_n(K+1)$ to $\mathscr{S}_n^K$, and set

$$\bar{\gamma}_n(t) = \gamma_n(2^{\beta n} t), \qquad \forall t \in [0, 2^{-\beta n} \operatorname{len}(\gamma_n)].$$

We want to show that $\bar{\gamma}_n$ converges to a scaling limit. Since the domain $\mathbb{Z}^3 \setminus \cup_{i=1}^K \operatorname{tr} \bar{\gamma}^{x(i)}$ does not have a polyhedral boundary, we cannot use [127, Theorem 1.3] directly. To get around this obstacle, we approximate with a simpler domain. Furthermore, to gain some control over the paths of the loop-erased random walks $(\bar{\gamma}_n^{x(i)})_{1\leq i\leq K}$, we also need to work within a bounded domain.

We write $D_n(R) = D_{2^{-n}}(R)$ to denote an scaled discrete box with side length $R \geq 1$ around the origin. Since the points $x(1), \ldots, x(K+1)$ are fixed, we can take R large enough so that $x(1), \ldots, x(K+1) \in D_n(R)$. For each curve $\bar{\gamma}_n^{x(i)} \in \Gamma$ and for the parameterized tree $\mathscr{S}_n^K$, we denote its restriction to the closed box $\bar{D}_n(R)$ with a super-index, as $\gamma_n^{x(i),R}$ and $\mathscr{S}_n^{K,R}$, respectively. We also consider the ILERW in the domain $D_n(R) \setminus \operatorname{tr} \mathscr{S}_n^{K,R}$. The exit time from such domain for the random walk S_n is

$$\xi_n^{\mathscr{S},R} = \inf\left\{t \geq 0 : S_n(t) \cap \left(\partial D_n(R) \cup \operatorname{tr} \mathscr{S}_n^{K,R}\right) \neq \emptyset\right\}.$$

The curve $\gamma_n^R = \operatorname{LE}\left(S_n[0, \xi^{\mathscr{S},R}]\right)$ is the LERW from $x_n(K+1)$ to either $\mathscr{S}_n^{K,R}$ or the boundary of $D_n(R)$; and we set

$$\bar{\gamma}_n^R(t) = \gamma_n^R(2^{\beta n} t), \qquad \forall t \in [0, 2^{-\beta n} \operatorname{len}(\gamma_n^R)]. \tag{5.82}$$

Note that we omit $x_n(K+1)$ as a super-index of γ_n^R to simplify the notation. We emphasize that $\bar{\gamma}_n^R$ is not necessarily the same as $\bar{\gamma}_n|^R$, where the latter is the restriction of the ILERW to the box $D_n(R)$, For each integer u, the cubes at scale u are closed cubes with vertices in $2^{-u}\mathbb{Z}^3$ and side length 2^{-u}. For $u \leq n$, let $A_n^{u,R}$ be the **u-dyadic approximation** to $\operatorname{tr}\mathscr{S}_n^{K,R}$, defined by

$$A_n^{u,R} := \bigcup_j \{C^{(u)} \colon \operatorname{dist}(C^{(u)}, \operatorname{tr}\mathscr{S}_n^{K,R}) \leq 2^{-2u}\}. \tag{5.83}$$

Proposition 5.7.11. *With fixed $n \in \mathbb{N}$, the sequence of sets $(A_n^{u,R})_{u\in\mathbb{N}}$ converges to $\operatorname{tr}\mathscr{S}_n^{K,R}$ in the Hausdorff topology. If $u \in \mathbb{N}$ is fixed, then the sequence $(A_n^{u,R})_{n\in\mathbb{N}}$ is eventually constant, almost surely.*

Proof. We begin with $n \in \mathbb{N}$ fixed. The construction of an u-dyadic approximation provides that $d_H(A_n^{u,R}, \operatorname{tr}\mathscr{S}_n^{K,R}) \leq 2^{-(u-2)}$. From here, it follows the convergence of $(A_n^{u,R})_{v\in\mathbb{N}}$ in the Hausdorff topology.

Next we consider $(A_n^{u,R})_{n\in\mathbb{N}}$ with u fixed. In this case, note that the a.s. convergence of $\mathscr{S}_n^{K,R}$ implies the a.s. convergence of $\operatorname{tr}\mathscr{S}_n^{K,R}$ in the Hausdorff topology. Then, for N large enough, $d_H(\operatorname{tr}\mathscr{S}_n^{K,R}, \operatorname{tr}\mathscr{S}_m^{K,R}) < 2^{-4u}$, if $n, m \geq N$, almost surely. It follows that $(A_n^{u,R})_{n\geq N}$ is constant almost surely. $\square$

We denote the constant limit of $A_n^{u,R}$, as $n \to \infty$, by $A^{u,R}$.

For each $n, u \in \mathbb{N}$ with $u \leq n$, consider the loop-erasure of the random walk S_n started from $x_n(K+1)$, and stopped when it exits $D_n(R) \setminus \operatorname{tr}\mathscr{S}_n^{K,R}$. We denote the latter hitting time by $\xi_u^{\mathcal{P},R}$, and the corresponding LERW by

$$\gamma_n^{u,R} = \operatorname{LE}(S_n[0, \xi_u^{\mathcal{P},R}]). \tag{5.84}$$

This curve has the β-parameterization

$$\bar{\gamma}_n^{u,R}(t) := \gamma_n^{u,R}(2^{\beta n} t), \qquad \forall t \in [0, 2^{-\beta n} \operatorname{len}(\gamma_n^{u,R})]. \tag{5.85}$$

The weak convergence of (5.85) is an immediate consequence of [127, Theorem 1.4] and Proposition 5.7.11. We state this observation below as Lemma 5.7.15.

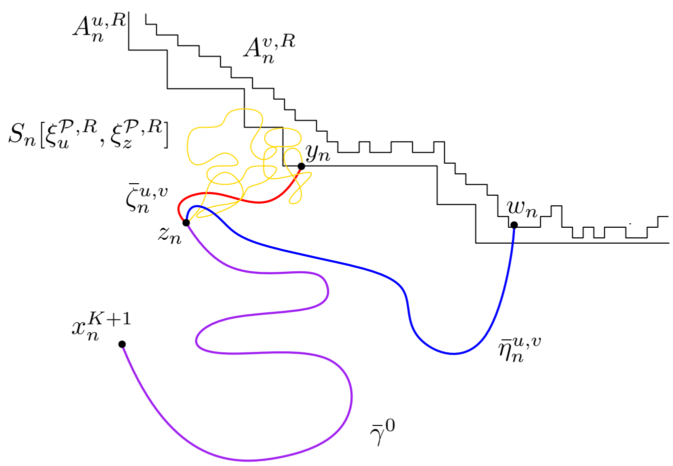

Figure 5.10: The figure shows the decomposition of the curves $\bar{\gamma}_n^{u,R}$ and $\bar{\gamma}_n^{v,R}$ used in the proof of Proposition 5.7.9. The curve $\bar{\gamma}_n^{u,R}$ is the concatenation of $\bar{\gamma}_0$ (in purple) and $\bar{\zeta}_n^{u,v}$ (in red). The curve $\bar{\gamma}_n^{v,R}$ is the concatenation of $\bar{\gamma}_0$ and $\bar{\eta}_n^{u,v}$ (in blue). The figure also shows a restriction of the random walk S_n from y_n to z_n (in yellow). In this case, S_n avoids hitting $A_n^{v,R}$ when it is close to y_n.

Step 2: comparing $\bar{\gamma}_n^{u,R}$ and $\bar{\gamma}_n^{v,R}$

Let $\bar{\gamma}_n^{u,R}$ and $\bar{\gamma}_n^{v,R}$ be the loop-erased random walks defined in (5.85). Our aim is to bound the distance $\psi(\bar{\gamma}_n^{u,R}, \bar{\gamma}_n^{v,R})$ for large values of n, u and v. Let us consider the event where this distance is large. More precisely, for $\varepsilon > 0$ and integers n, u and v, with $u, v < n$, we define

$$\mathcal{E}_n^{u,v}(\varepsilon) := \left\{ \psi(\bar{\gamma}_n^{u,R}, \bar{\gamma}_n^{v,R}) \geq \varepsilon \right\}.$$

Recall that we use the same random walk on $2^{-n}\mathbb{Z}^3$ to generate $\bar{\gamma}_n^{u,R}$ and $\bar{\gamma}_n^{v,R}$. Typically, these two curves have a segment in common, $\bar{\gamma}^0$ (see Figure 5.10). We claim that $\bar{\gamma}_n^{u,R} \setminus \bar{\gamma}^0$ and $\bar{\gamma}_n^{v,R} \setminus \bar{\gamma}^0$ have a small effect on $\psi(\bar{\gamma}_n^u, \bar{\gamma}_n^v)$.

Towards proving the preceding claim, we start by introducing some further notation for elements in the curves $\bar{\gamma}_n^{u,R}$ and $\bar{\gamma}_n^{v,R}$; Figure 5.10 serves as a reference. For clarity, and without loss of generality, we elaborate our arguments on the event where the random walk S_n hits the boundary of $\mathcal{P}^{u,R}$ first, that is

$$\mathcal{F}^u := \left\{\xi_u^{\mathcal{P},R} \leq \xi_v^{\mathcal{P},R}\right\},$$

and hence it generates $\bar{\gamma}_n^{u,R}$ before $\bar{\gamma}_n^{v,R}$. The symmetric event is $\mathcal{F}^v := \{\xi_v^{\mathcal{P},R} \leq \xi_u^{\mathcal{P},R}\}$. We will consider the restriction of the random walk S_n:

$$S_n^{u,v} := S_n|_{[\xi_u^{\mathcal{P},R}, \xi_v^{\mathcal{P},R}]}.$$

Denote the endpoint of $\bar{\gamma}_n^{u,R}$ by $y_n := S_n(\xi_u^{\mathcal{P},R})$. To simplify notation, we denote the durations of $\bar{\gamma}_n^{u,R}$ and $\bar{\gamma}_n^{v,R}$ by

$$T^u = T(\bar{\gamma}_n^{u,R}), \qquad T^v = T(\bar{\gamma}_n^{v,R}),$$

respectively. The last time that $S_n^{u,v}$ hits its past $\bar{\gamma}_n^{u,R}$ determines the endpoint of $\bar{\gamma}^0$. Let

$$\xi_z^{\mathcal{P},R} := \sup\{t \leq \xi_v^{\mathcal{P},R} \colon S(t) \in \bar{\gamma}_n^{u,R}\}$$

and set $z_n := S(\xi_z^{\mathcal{P},R})$. Let T_{z} be such that $\bar{\gamma}_n^{v,R}(T_{\mathrm{z}}) = \bar{\gamma}_n^{v,R}(T_{\mathrm{z}}) = z_n$. We then have for the common curve $\bar{\gamma}^0 = \bar{\gamma}_n^{u,R}[0, T_{\mathrm{z}}] = \gamma_n^{v,R}[0, T_{\mathrm{z}}]$. The difference between $\bar{\gamma}_n^{u,R}$ and $\bar{\gamma}_n^{v,R}$ are the curves

$$\bar{\zeta}_n^{u,v} := \bar{\gamma}_n^{u,R}[T_{\mathrm{z}}, T^u], \qquad \bar{\eta}_n^{u,v} = \bar{\gamma}_n^{v,R}[T_{\mathrm{z}}, T^v].$$

Note that the range of $\bar{\eta}_n^{u,v}$ is a subset of $S_n^{u,v}$.

We now compare the shapes of $\bar{\gamma}_n^{v,R}$ and $\bar{\gamma}_n^{u,R}$. In particular, we note that the respective traces of these curves can be significantly different if one of the next two bad events occur. The first event controls the diameter of $\bar{\eta}_n^{u,v}$, while the second event imposes a limit on the size of $\bar{\zeta}_n^{u,v}$.

- Since $\bar{\eta}_n^{u,v}$ is a subset of $S_n^{u,v}$, $\bar{\eta}_n^{u,v}$ has a diameter larger than ε_0 only if $S_n^{u,v}$, the segment of the random walk S_n between the hitting times $\xi_u^{\mathcal{P},R}$ and $\xi_v^{\mathcal{P},R}$, has a similarly large diameter. We denote this event

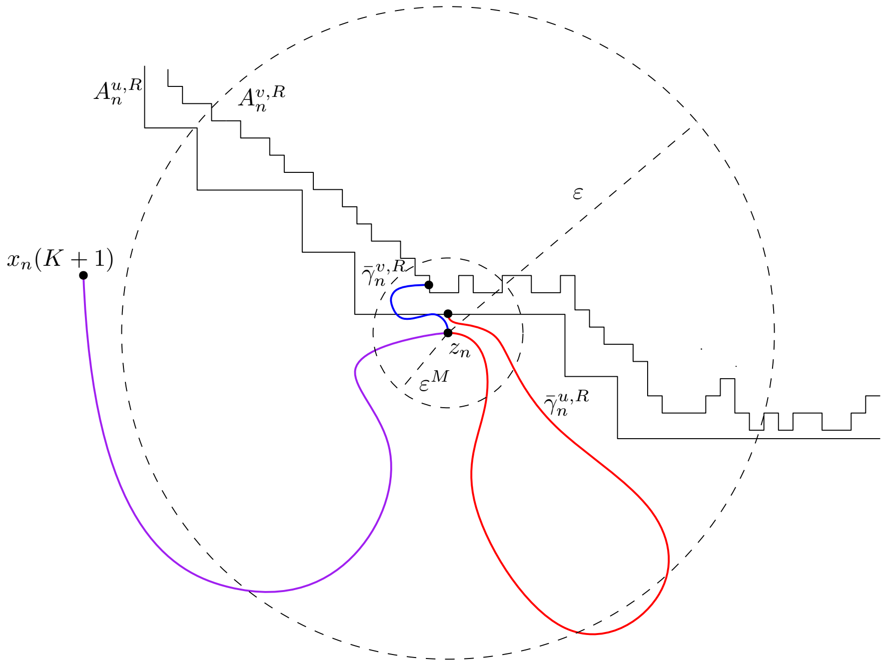

Figure 5.11: A realization of the event $\mathcal{D}_n^{u,v}(\varepsilon)^c \cap \mathcal{Q}(\varepsilon^M, \varepsilon)$. In the figure, $\bar\gamma_n^{u,R}$ is the concatenation of the purple and blue curves, while $\bar\gamma_n^{v,R}$ is the concatenation of the purple and red curves.

by

$$\mathcal{D}_n^{u,v}(\varepsilon_0) := \{\mathrm{diam}(S_n^{u,v}) \geq \varepsilon_0\}\,.$$

- On the complementary event $\mathcal{D}_n^{u,v}(\varepsilon_0)^c$, the curve $\bar\zeta_n^{u,v}$ has diameter larger than ε only if $\bar\gamma_n^{u,R}$ has an $(\varepsilon_0, \varepsilon)$–quasi-loop. Figure 5.11 shows an example of this situation. Let

$$\mathcal{Q}(\varepsilon_0, \varepsilon; \gamma) := \{\gamma \text{ has an } (\varepsilon_0, \varepsilon)\text{-quasi-loop}\},$$

and $\mathcal{Q}_n(\varepsilon_0, \varepsilon) = \mathcal{Q}(\varepsilon_0, \varepsilon; \bar\gamma_n^u) \cup \mathcal{Q}(\varepsilon_0, \varepsilon; \bar\gamma_n^v)$.

Combining the definitions above, we introduce a bad event for the shape by setting

$$\mathcal{B}_n^{u,v}(\varepsilon) := \mathcal{D}_n^{u,v}(\varepsilon^M) \cup \mathcal{Q}_n(\varepsilon^M, \varepsilon),$$

noting that we have taken $\varepsilon_0 = \varepsilon^M$, with $M > 1$ being the exponent of Proposition 5.3.11. We highlight that, on the event $(\mathcal{B}_n^{u,v}(\varepsilon))^c$, it holds that $d_H(\bar{\gamma}_n^{u,R}, \bar{\gamma}_n^{v,R}) \leq \varepsilon$.

The following result establishes that, on $(\mathcal{B}_n^{u,v}(\varepsilon))^c$, $\bar{\gamma}_n^{u,R}$ and $\bar{\gamma}_n^{v,R}$ are also close as parameterized curves. The issue here is that even if the traces of two curves may be close in shape, they may take a large number of steps in a small diameter. We will compare the Schramm and intrinsic distances, as defined at (5.18) and (5.19), on the event $(\mathcal{B}_n^{u,v}(\varepsilon))^c$ where the shapes are close to each other. The Schramm and intrinsic distances of $\bar{\gamma}_n^{u,R}$ are comparable on the events $S^\dagger_{2^{-n}}(R,\varepsilon)$ and $E^\dagger_{2^{-n}}(R,\varepsilon)$. These events are introduced in Subsection 5.3.5. To simplify notation, we write $S^\dagger_n(R,\varepsilon) = S^\dagger_{2^{-n}}(R,\varepsilon)$ and $E^\dagger_n(R,\varepsilon) = E^\dagger_{2^{-n}}(R,\varepsilon)$.

Lemma 5.7.12. *Fix $R \geq 1$ and let $\varepsilon \in (0,1)$. On the event $(\mathcal{B}_n^{u,v}(\varepsilon))^c \cap S^\dagger_n(R\varepsilon^{-1},\varepsilon)$, we have that*

$$T(\bar{\eta}_n^{u,v}) \leq R\varepsilon^{\beta-1}, \qquad T(\bar{\zeta}_n^{u,v}) \leq R\varepsilon^{\beta-1}.$$

Proof. In this proof, we write $G = (\mathcal{B}_n^{u,v}(\varepsilon))^c \cap S^\dagger(R\varepsilon^{-1},\varepsilon)$. We begin with an upper bound for the duration of $\bar{\eta}_n^{u,v}$. On G, the random walk $S_n^{u,v}$ is localized in a neighbourhood around y_n. Indeed, on $\mathcal{D}_n^{u,v}(\varepsilon^M)^c$ we have that $\operatorname{diam}(S_n^{u,v}) \leq \varepsilon^M$. Since $\bar{\eta}_n^{u,v}$ is a subset of $S_n^{u,v}$, it follows that $\operatorname{diam}(\bar{\eta}_n^{u,v}) < \varepsilon^M$, and, in particular, for the endpoints of $\bar{\eta}_n^{u,v}$, z_n and w_n say (as in Figure 5.10), we have that $d^S_{\bar{\gamma}_n^v}(z_n,w_n) = d^S_{\bar{\eta}}(z_n,w_n) \leq \varepsilon^M < \varepsilon$. On $G \subseteq S^\dagger_n(R\varepsilon^{-1},\varepsilon)$, this implies that

$$T(\bar{\eta}_n^{u,v}) = d_{\bar{\gamma}_n^v}(z_n,w_n) \leq R\varepsilon^{\beta-1}.$$

Next we bound the duration of $\bar{\zeta}_n^{u,v}$ on the event G. We have that the endpoints of $\bar{\zeta}_n^{u,v}$ are in $S_n^{u,v}$. Indeed, $y_n = S_n^{u,v}(0)$ and $z_n \in \bar{\zeta}_n^{u,v} \subseteq S_n^{u,v}$. Thus

$$|z_n - y_n| < \varepsilon^M. \tag{5.86}$$

On the event $G \subseteq \mathcal{Q}(\varepsilon^M,\varepsilon)^c$, the loop-erased random walk $\bar{\gamma}_n^{u,R}$ does not have $(\varepsilon^M,\varepsilon)$–quasi-loops, and so (5.86) implies that $d^S_{\bar{\gamma}_n^u}(z_n,y_n) < \varepsilon$. The argument used for $\bar{\eta}_n^{u,v}$ also gives $T(\bar{\zeta}_n^{u,v}) = d_{\bar{\gamma}_n^u}(z_n, y_n) < R\varepsilon^{\beta-1}$. $\square$

We finish this step by showing that $\mathcal{E}_n^{u,v}(\varepsilon)$ can be contained in the events already described.

Lemma 5.7.13. *Let $R \geq 1$ and $\varepsilon > 0$. On the event $\mathcal{B}_n^{u,v}(\varepsilon)^c \cap S_n^\dagger(R\varepsilon^{-1}, \varepsilon) \cap E_n^\dagger(R, \varepsilon)$, we have that*

$$\psi(\bar{\gamma}_n^{u,R}, \bar{\gamma}_n^{v,R}) \leq CR^{b_3}\varepsilon^{\beta'},$$

where $0 < b_3, C < \infty$ are universal constants, and $\beta' = b_3(\beta - 1)$.

Proof. It suffices to show that on the event $G_2 = \mathcal{B}_n^{u,v}(\varepsilon)^c \cap S^\dagger(R\varepsilon^{-1}, \varepsilon) \cap E^\dagger(R, \varepsilon)$,

$$\psi(\bar{\gamma}_n^u, \bar{\gamma}_n^v) = |T^u - T^v| + \max_{0 \leq s \leq 1} |\bar{\gamma}_n^u(sT^u) - \bar{\gamma}_n^v(sT^v)| \leq CR^{b_3}\varepsilon^{\beta'}. \tag{5.87}$$

Lemma 5.7.12 gives $|T^u - T^v| < 2R\varepsilon^{\beta-1}$. Next we bound the second term in (5.87). Let $a = \bar{\gamma}_n^{u,R}(sT^u)$ and $b = \bar{\gamma}_n^{v,R}(sT^v)$ and assume that one of these points belongs to the common path, say $b \in \bar{\gamma}^0$. In this case, $sT^v \leq T^u$ and we can re-write $b = \bar{\gamma}_n^{u,R}((s(T^v/T^u))T^u)$. Then, with respect to the intrinsic metric of $\bar{\gamma}_n^{u,R}$, we compare points within distance

$$d_{\bar{\gamma}_n^u}(a, b) \leq |sT^u - (sT^v/T^u)\, T^u| \leq 2R\varepsilon^{\beta-1}.$$

We introduce

$$N^u = \sup\left\{|a - b| : a, b \in \operatorname{tr} \bar{\gamma}_n^{u,R}, d_{\bar{\gamma}_n^u}(x, y) \leq 2R\varepsilon^{\beta-1}\right\},$$

define N^v similarly from $\bar{\gamma}_n^v$, and also introduce notation for the diameter of the segments $\bar{\eta}_n^{u,v}$ and $\bar{\zeta}_n^{u,v}$ by setting

$$N^\eta = \sup_{0 \leq t \leq \operatorname{len}(\bar{\eta}_n^{u,v})} |\bar{\eta}_n^{u,v}(t) - z_n|, \qquad N^\zeta = \sup_{0 \leq t \leq \operatorname{len}(\bar{\zeta}_n^{u,v})} |\bar{\zeta}_n^{u,v}(t) - z_n|.$$

It readily holds that we have the following bound:

$$\max_{0\leq s\leq 1} |\bar{\gamma}_n^u(sT^u) - \bar{\gamma}_n^v(sT^v)| \leq N^u + N^v + N^\eta + N^\xi. \tag{5.88}$$

Lemma 5.7.12 implies that $d_\eta(a,b) < R\varepsilon^{\beta-1}$ for all $a \in \operatorname{tr}\eta$. Let $b_3 = \frac{b_2}{b_1}$, where b_1 and b_2 are the constants of Proposition 5.3.5. On the event $E_n^\dagger(R,\varepsilon)$, we thus have that $N^\eta \leq R^{b_3}\varepsilon^{b_3(\beta-1)}$, and similarly for N^ξ. On the event $E_n^\dagger(R,\varepsilon)$, the loop-erased random walks $\bar{\gamma}_n^u$ and $\bar{\gamma}_n^v$ are uniformly equicontinuous, so that $N^u \leq CR^{b_3}\varepsilon^{b_3(\beta-1)}$, and the same bound holds for N^v. Adding the upper bounds for N^u, N^v N^η and N^ξ in (5.88), we get (5.87). □

Step 3: bounding $\mathbf{P}(\mathcal{E}_n^{u,v}(\varepsilon))$

In this step, we give an upper bound on the probability of the bad event $\mathcal{E}_n^{u,v}(\varepsilon)$. The key is that, given u, and v, this estimate is uniform over all n for u and v large enough.

Lemma 5.7.14. *Fix $R \geq 1$. For each $\varepsilon \in (0,1)$, there exists $U = U(\varepsilon)$ such that for all $n \geq u, v \geq U(\varepsilon)$*

$$\mathbf{P}\left(\mathcal{E}_n^{u,v}(\varepsilon)\right) \leq C\varepsilon^\theta,$$

for constants $C = C(R) > 0$ and $\theta = \theta(R) > 0$, depending only on R.

Proof. Lemma 5.7.13 gives that

$$\begin{aligned}\mathbf{P}\left(\mathcal{E}_n^{u,v}(CR\varepsilon^{\beta'})\right) \leq \mathbf{P}\left(\mathcal{D}_n^{u,v}(\varepsilon^M)\right) + \mathbf{P}\left(\mathcal{Q}(\varepsilon^M,\varepsilon)\right) \\ + \mathbf{P}\left(\left(S_n^\dagger(R\varepsilon^{-1},\varepsilon)\right)^c\right) + \mathbf{P}\left(\left(E_n^\dagger(R,\varepsilon)\right)^c\right).\end{aligned} \tag{5.89}$$

Proposition 5.3.11 implies

$$\mathbf{P}(\mathcal{Q}(\varepsilon^M,\varepsilon)) \leq \mathbf{P}(\mathcal{Q}(\varepsilon^M,\varepsilon^2)) \leq CR^3\varepsilon^{\hat{b}_2},$$

and Propositions 5.3.12 and 5.3.13 give upper bounds for the last two terms of (5.89). Thus we are left to bound the probability of $\mathcal{D}_n^{u,v}(\varepsilon^M)$. For this, we need $U(\varepsilon)$ large enough so that, by Proposition 5.7.11, we have that

$$d_H(A_n^{u,R}, A_n^{v,R}) < \varepsilon^{4M} \qquad \text{for all } u, v \geq U(\varepsilon). \tag{5.90}$$

On $\mathcal{F}^u$, $\bar{\gamma}_n^{u,R}$ is the first walk LERW to stop, and we call its endpoint $y_n \in \partial \mathcal{P}^{u,R}$. From (5.90), we have $\operatorname{dist}(y_n, \partial \mathcal{P}^{v,R}) < \varepsilon^{4M}$. But, along $S_n^{u,v}$, the random walk S_n reaches distance ε^M before hitting $\partial \mathcal{P}^{v,R}$. The same argument on the complement of $\mathcal{F}^u$, i.e. on $\mathcal{F}^v$. Hence Proposition 5.7.4 implies that $\mathbf{P}(\mathcal{D}_n^{u,v}(\varepsilon^M)) \leq CKR^3\varepsilon^{2M} + \varepsilon^{2M\hat{\eta}}$. In conjunction with the aforementioned bounds,

$$\sup_{\substack{n,u,v: \\ n \geq u,v \geq U}} \mathbf{P}\left(\mathcal{E}_n^{u,v}(R^{b_3}\varepsilon^{\beta'})\right) \leq CKR^3\varepsilon^{2M} + \varepsilon^{2M\hat{\eta}} + C\varepsilon^{\tilde{\eta}} + C\left(\frac{R}{\varepsilon}\right)^3 e^{-c\left(\frac{R}{\varepsilon}\right)^a} + C\varepsilon^{b_2}.$$

The dominant term above is $\varepsilon^{\theta(R)}$, and a reparameterization completes the proof. $\square$

Step 4: the scaling limit of a loop-erased random walk

Recall that $\bar{\gamma}_n^R$ is the LERW on $\bar{D}_n(R) \backslash \mathscr{S}_n^{K,R}$ defined in (5.82). In (5.85), we defined $\bar{\gamma}_n^{u,R}$, for $u \leq n$, as the β-parameterization of the loop-erased random walk $\operatorname{LE}\left(S_n[0, \xi_m^{\mathcal{P},R}]\right)$, where $\xi_m^{\mathcal{P},R}$ is the first exit time from the dyadic polyhedron $\mathcal{P}^{m,R}$. In this step, we establish that $\bar{\gamma}_n^R$ and $\bar{\gamma}_n^{u,R}$ converge to the same limit. We take limits on each variable in the following order. For $\bar{\gamma}_n^{u,R}$, we first take $n \to \infty$. The limit object is a curve on the bounded and polyhedral domain $D_E(R) \setminus A^{u,R} \subset \mathbb{R}^3$, where $A^{u,R}$ is the polyhedral domain of Proposition 5.7.11. Then we take $u \to \infty$, and the limit is a curve on the bounded set $D_E(R) \setminus \operatorname{tr} \mathscr{S}^{K,R}$. In Step 5, we take $R \to \infty$, and we thus define $\hat{\gamma}$ as a limit curve on the full space $\mathbb{R}^3 \setminus \operatorname{tr} \mathscr{S}^K$.

Lemma 5.7.15. *Fix $R \geq 1$. For each $u \in \mathbb{N}$, the law of $\bar{\gamma}_n^{u,R}$ converges with respect to the metric ψ, as $n \to \infty$.*

Proof. Proposition 5.7.11 shows that the domain $\bar{\gamma}_n^{u,R}$ is the polyhedron

$D_E(R) \setminus A^{u,R}$, for n large enough. Then, the weak convergence of $\{\bar{\gamma}_n^{u,R}\}_{n\in\mathbb{N}}$ is an immediate consequence of Proposition 5.3.14. $\square$

We denote by $\hat{\gamma}^{u,R}$ a curve with the limit law of Lemma 5.7.15.

Lemma 5.7.16. *Fix $R \geq 1$. Let $(\hat{\gamma}^{u,R})_{u\in\mathbb{N}}$ be the sequence of limit elements from Lemma 5.7.15. It is then the case that $(\hat{\gamma}^{u,R})_{u\in\mathbb{N}}$ converges in distribution in the metric ψ as $u \to \infty$.*

Proof. Denote the laws of $\bar{\gamma}_n^{u,R}$ and $\hat{\gamma}^{u,R}$ by $\mathcal{L}(\bar{\gamma}_n^{u,R})$ and $\mathcal{L}(\hat{\gamma}^{u,R})$, respectively. Since $(\mathcal{C}_f, \psi)$ is a complete and separable metric space (see [94, Section 2.4]), to prove weak convergence it suffices to show that $(\mathcal{L}(\hat{\gamma}^{u,R}))_{u\in\mathbb{N}}$ is a Cauchy sequence in the Prohorov metric $d_{\mathbf{P}}$. Let $u, v \in \mathbb{N}$. By the triangle inequality, for $n \geq u, v$,

$$\begin{aligned} d_{\mathbf{P}}(\mathcal{L}(\hat{\gamma}^{u,R}), \mathcal{L}(\hat{\gamma}^{v,R})) \leq\ & d_{\mathbf{P}}(\mathcal{L}(\hat{\gamma}^{u,R}), \mathcal{L}(\bar{\gamma}_n^{u,R})) + d_{\mathbf{P}}(\mathcal{L}(\hat{\gamma}^{v,R}), \mathcal{L}(\bar{\gamma}_n^{v,R})) \\ & + \sup_{n\geq u,v} d_{\mathbf{P}}(\mathcal{L}(\bar{\gamma}_n^{u,R}), \mathcal{L}(\bar{\gamma}_n^{v,R})). \end{aligned} \tag{5.91}$$

Letting $n \to \infty$, the first two terms on the right hand side of (5.91) converge to 0 by Lemma 5.7.15. Moreover, Lemma 5.7.14 shows that the last term of (5.91) converges to 0 as $u, v \to \infty$. Therefore $(\mathcal{L}(\hat{\gamma}^{u,R}))_{u\in\mathbb{N}}$ is a Cauchy sequence in the Prohorov metric. It follows that $(\hat{\gamma}^{u,R})_{u\in\mathbb{N}}$ converges weakly. $\square$

We denote by $\hat{\gamma}^R$ a curve with the limit law of Lemma 5.7.16. The random curve $\hat{\gamma}^R$ is the limit of dyadic approximations. We see below that it is also the limit of the LERWs stopped when they hit $\mathscr{S}^{K,R}$.

Lemma 5.7.17. *Fix $R \geq 1$. Then $\bar{\gamma}_n^R \to \hat{\gamma}^R$ in distribution as $n \to \infty$, with respect to the metric ψ.*

Proof. Since $\bar{\gamma}_n^{u,R} \to \hat{\gamma}^{u,R}$ in distribution as $n \to \infty$, and $\bar{\gamma}^{u,R} \to \hat{\gamma}^R$ in distribution as $u \to \infty$, to complete the proof it suffices to notice that, for $\varepsilon > 0$,

$$\lim_{u\to\infty} \limsup_{n\to\infty} \mathbf{P}\left(\psi(\bar{\gamma}_n^{u,R}, \bar{\gamma}_n^R) > \varepsilon\right) = 0,$$

see [33, Theorem 3.2], for example. However, since $\bar{\gamma}_n^R = \bar{\gamma}_n^{n,R}$, the above statement readily follows from Lemma 5.7.14. $\square$

Step 5: taking $R \to \infty$

Until this point, we have only considered LERW inside a box $D_E(R)$. Indeed, $\bar{\gamma}_n^R$ was defined as a LERW in $D_n(R) \setminus \mathscr{S}_n^K$, and its scaling limit $\hat{\gamma}^R$ is within $D_E(R)$. In this final step, we will take $R \to \infty$ to consider the tree $\hat{\mathscr{S}}^K$ and the random walk S_n in the full space.

Lemma 5.7.18. *Let $\left(\hat{\gamma}^R\right)_{R \geq 1}$ be the sequence of limit elements from Lemma 5.7.16 and $\hat{\mathscr{S}}^K$ is the parameterized tree in Proposition 5.7.9. There exists a random element $\hat{\gamma} \in \mathcal{C}_f$ such that $\hat{\gamma}^R$ converges in distribution to $\hat{\gamma}$ in the metric ψ as $R \to \infty$. Moreover, the intersection of $\operatorname{tr} \hat{\gamma}$ and $\operatorname{tr} \hat{\mathscr{S}}^K$ is the endpoint of $\hat{\gamma}$.*

Proof. Denote the laws of $\bar{\gamma}_n^R$ and $\hat{\gamma}^R$ by $\mathcal{L}(\bar{\gamma}_n^R)$ and $\mathcal{L}(\hat{\gamma}^R)$, respectively. This proof is similar to the one for Lemma 5.7.16 as we will show that $\left(\mathcal{L}(\hat{\gamma}^R)\right)_{R \geq 1}$ is a Cauchy sequence in the Prohorov metric $d_{\mathbf{P}}$. For two integers $r > s > 0$, the triangle inequality yields

$$d_{\mathbf{P}}\left(\mathcal{L}(\hat{\gamma}^r, \hat{\gamma}^s)\right) \leq d_{\mathbf{P}}\left(\mathcal{L}(\hat{\gamma}^r, \bar{\gamma}_n^r)\right) + d_{\mathbf{P}}\left(\mathcal{L}(\hat{\gamma}^s, \bar{\gamma}_n^s)\right) + \sup_n d_{\mathbf{P}}\left(\mathcal{L}(\bar{\gamma}_n^r, \bar{\gamma}_n^s)\right). \tag{5.92}$$

Letting $n \to \infty$, the first two terms on the right hand side of (5.92) converge to 0 by Lemma 5.7.17. Then we are left to bound $\sup_n d_{\mathbf{P}}\left(\mathcal{L}(\bar{\gamma}_n^r, \bar{\gamma}_n^s)\right)$.

$$d_{\mathbf{P}}\left(\mathcal{L}(\hat{\gamma}^r, \hat{\gamma}^s)\right) \leq \sup_n d_{\mathbf{P}}\left(\mathcal{L}(\bar{\gamma}_n^r, \bar{\gamma}_n^s)\right). \tag{5.93}$$

Recall that we sample $\bar{\gamma}_n^r$ and $\bar{\gamma}_n^r$ as loop-erasures of the simple random walk S_n. On the event that $\mathscr{S}_n^{K,r} \triangle \mathscr{S}_n^{K,s} \cap B(s^{1/2}) = \emptyset$, $\bar{\gamma}_n^r = \bar{\gamma}_n^r$ (as parameterized curves) whenever S_n hits $\operatorname{tr} \mathscr{S}_n^{K,r}$ before reaching the boundary of $B_n(s^{1/2})$, and so

$$\begin{aligned} \mathbf{P}\left(\bar{\gamma}_n^r \neq \bar{\gamma}_n^s\right) \leq & \mathbf{P}\left((\mathscr{S}_n^{K,r} \triangle \mathscr{S}_n^{K,s}) \cap B_n(s^{1/2}) \neq \emptyset\right) \\ & + \mathbf{P}\left(S_n[0, \xi_S(B_n(s^{1/2}))] \cap \operatorname{tr} \mathscr{S}_n^{K,r} \neq \emptyset\right). \end{aligned}$$

Proposition 5.7.10 gives $\mathbf{P}\left((\mathscr{S}_n^{K,r} \triangle \mathscr{S}_n^{K,s}) \cap B_n(s^{1/2}) \neq \emptyset\right) \to 0$ as $s \to \infty$. Recall that $\mathscr{S}_n^K$ is a subtree of the uniform spanning tree, including a path to infinity. Then, Proposition 5.7.4 implies that

$$\mathbf{P}\left(S[0, \xi_S(B_n(s^{1/2}))] \cap \operatorname{tr} \mathscr{S}_n^{K,r} \neq \emptyset\right) \to 0 \text{ as } r, s \to \infty.$$

Therefore, (5.93) converges to 0 as $r, s \to \infty$. It follows that $\left(\mathcal{L}(\hat{\gamma}^R)\right)_{R \geq 1}$ is a Cauchy sequence in the Prohorov metric. Since $d_{\mathbf{P}}$ is a complete metric, we conclude that $\left(\mathcal{L}(\hat{\gamma}^R)\right)$ converges weakly. Such limit is a random element $\hat{\gamma}$ taking values in $\mathcal{C}_f$, and in particular $\hat{\gamma}$ has finite duration.

On the space of finite curves $(\mathcal{C}_f, \psi)$, the evaluation of the endpoint defines a continuous function $E : \mathcal{C}_f \to \mathbb{R}^3$. Therefore, as we take $n \to \infty$, the endpoint of $E(\bar{\gamma}_n^R) \in \operatorname{tr} \mathscr{S}_n^K$ converges to $E(\hat{\gamma}^R)$ (see [33, Theorem 5.1], for example). Proposition 5.3.8 implies that, with probability one, $E(\bar{\gamma}_n^R) \in \operatorname{tr} \mathscr{S}_n^K$ for R large enough. Additionally, note that $\mathscr{S}_n^K$ converges weakly to $\mathscr{S}^K$ as a parameterized tree, when $n \to \infty$. It follows that the law of $E(\hat{\gamma}^R)$ is supported on $\mathscr{S}^K$. $\square$

Lemma 5.7.19. *The collection of curves $\Gamma^e(\hat{\mathscr{S}}^K) \cup \{\hat{\gamma}\}$ define a parameterized tree $\hat{\mathscr{S}}^{K+1}$. This tree coincides with the description in Section 5.7.5.*

Proof. Lemma 5.7.18 shows that $\Gamma^c(\mathscr{S}^K) \cup \{\hat{\gamma}\}$ satisfies the conditions of Proposition 5.7.7. It follows that $\Gamma^c(\mathscr{S}^K) \cup \{\hat{\gamma}\}$ is the set of essential branches for a parameterized tree $\hat{\mathscr{S}}^{K+1}$.

Finally, note that Lemma 5.7.17 shows that $\hat{\gamma}$ is the limit of scaled loop-erased random walks, stopped when they hit the previous limit element $\operatorname{tr} \mathscr{S}^K$, and such hitting time is finite. Therefore $\hat{\mathscr{S}}^{K+1}$ is the tree of Section 5.7.5. $\square$

Step 6: the scaling limit of parameterized trees

Proof of Proposition 5.7.9. First let us describe the probability measure induced by $(\mathscr{S}_n^K, \bar{\gamma}_n)$. Let μ_n be the probability measure on $\mathscr{F}^K$ induced by $\mathscr{S}_n^K$. For each $\mathscr{S}_n^K \in \mathscr{F}^K$, let $\nu_n^{\gamma_n}$ be the probability measure on $(\mathcal{C}_f, \psi)$

induced by the loop-erased random walk $\bar{\gamma}_n$; recall that $\bar{\gamma}_n$ is stopped when it exits $(\mathbb{R}^3 \setminus \operatorname{tr} \mathscr{S}_n^K) \cap \mathbb{Z}^3$. The measure $\nu_n^{\gamma_n}$ defines the stochastic kernel

$$K_n(\mathscr{S}_n^K, A) = \nu_n^{\gamma_n}(A), \qquad \forall \mathscr{S}_n^K \in \mathscr{F}^K, A \in \mathcal{B}(\mathcal{C}_f),$$

where $\mathcal{B}(\mathcal{C}_f)$ is the Borel σ-algebra corresponding to $(\mathcal{C}_f, \psi)$. That is, the probability measure induced by $(\mathscr{S}_n^K, \bar{\gamma}_n)$, $\mu_n \otimes K_n$ say, is the unique probability measure such that

$$\mu_n \otimes K_n(A_1 \times A_2) = \int_{A_1} K_n(\mathscr{S}_n^K, A_2)\mu_n(d\mathscr{S}_n^K),$$

for Borel sets $A_1 \in \mathcal{B}(\mathscr{F}^K)$ and $A_2 \in \mathcal{B}(\mathcal{C}_f)$.

Now, recall we are supposing that we have a coupling so that $\mathscr{S}_n^K \to \hat{\mathscr{S}}^K$, almost-surely. In what follows, we write $\mathbf{P}^*$ for the corresponding probability measure. From Lemma 5.7.18, we obtain that, $\mathbf{P}^*$-a.s., $\nu_n^{\gamma_n} \to \nu^{\hat{\gamma}}$ as $n \to \infty$, where $\nu^{\hat{\gamma}}$ is the law of $\hat{\gamma}$. Hence $\nu^{\hat{\gamma}}$ is $\mathbf{P}^*$-measurable, and, in particular, so is $\nu^{\hat{\gamma}}(A)$ for all $A \in \mathcal{B}(C_f)$. As a consequence, the integral

$$\mu \otimes K\,(A_1 \times A_2) := \int \mathbf{1}_{A_1}(\hat{\mathscr{S}}^K)\nu^{\hat{\gamma}}(A_2)d\mathbf{P}^*$$

is well-defined for every $A_1 \in \mathcal{B}(\mathscr{F}^K)$, $A_2 \in \mathcal{B}(\mathcal{C}_f)$. Moreover, $\mu \otimes K$ is readily extended to give a measure on the product space $\mathscr{F}^K \times \mathcal{C}_f$. Finally, let $A_1 \in \mathcal{B}(\mathscr{F}^K)$, $A_2 \in \mathcal{B}(\mathcal{C}_f)$ be continuity sets for $\mu \otimes K$, in the sense that $\mu \otimes K(\partial A_1 \times \mathcal{C}_f) = 0 = \mu \otimes K(\mathscr{F}^K \times \partial A_2)$. We then have that, $\mathbf{P}^*$-a.s., $\mathbf{1}_{A_1}(\mathscr{S}_n^K)\nu_n^{\hat{\gamma}}(A_2) \to \mathbf{1}_{A_1}(\hat{\mathscr{S}}^K)\nu^{\hat{\gamma}}(A_2)$. An application of the dominated convergence theorem thus yields

$$\mu_n \otimes K_n(A_1 \times A_2) \to \mu \otimes K\,(A_1 \times A_2)\,,$$

which is enough to establish that $\mu \otimes K$ is a measure on $(\mathcal{F}^K, \mathcal{C}_f)$ (see [33, Theorem 2.8]). Lemma 5.7.19 shows that $\mu \otimes K$ defines a measure on the spaces of parameterized trees $\mathcal{F}^{K+1}$. $\square$

5.8 Proof of tightness and subsequential scaling limit

Given the preparations in the previous sections, we are now in a position to establish the first main result of this article, namely Theorem 5.1.1.

Proof of Theorem 5.1.1. We start by establishing the parts of the result concerning the Gromov-Hausdorff-type topology. Applying Lemma 5.2.3, the tightness claim follows from Proposition 5.4.1, Corollary 5.4.1 and Proposition 5.4.5. It remains to check the distributional convergence of $\underline{\mathcal{U}}_n$ as $n \to \infty$, where we write $\underline{\mathcal{U}}_n$ for the random measured, rooted spatial tree at (5.1), indexed by $\delta_n = 2^{-n}$. By the first part of the theorem and Prohorov's theorem (see [93, Theorem 16.3], for example), we know that every subsequence $(\underline{\mathcal{U}}_{n_i})_{i\geq 1}$ admits a convergent subsubsequence $(\underline{\mathcal{U}}_{n_{i_j}})_{j\geq 1}$. Thus we only need to establish the uniqueness of the limit.

Now, suppose $(\underline{\mathcal{U}}_{n_i})_{i\geq 1}$ is a convergent subsequence, and write $\underline{\mathcal{T}} = (\mathcal{T}, d_{\mathcal{T}}, \mu_{\mathcal{T}}, \phi_{\mathcal{T}}, \rho_{\mathcal{T}})$ for the limiting random element in $\mathbb{T}$. To show that the convergence specifies the law of $\underline{\mathcal{T}}$ uniquely, we will start by considering finite restrictions of $\underline{\mathcal{U}}_{n_i}$, $i \geq 1$. In particular, for $R \in (0, \infty)$, set $\underline{\mathcal{U}}_{n_i}^{(R)}$ as

$$\left(B(\delta_{n_i}^{-1}R), \delta_{n_i}^{\beta} d_{\mathcal{U}}|_{B(\delta_{n_i}^{-1}R)\times B(\delta_{n_i}^{-1}R)}, \delta_{n_i}^{3}\mu_{\mathcal{U}}\left(\cdot \cap B(\delta_{n_i}^{-1}R)\right), \delta_{n_i}\phi_{\mathcal{U}}|_{B(\delta_{n_i}^{-1}R)}, \rho_{\mathcal{U}}\right),$$

i.e. the part of $\underline{\mathcal{U}}_{n_i}$ contained inside $B(\delta_{n_i}^{-1}R)$. (We acknowledge this notation clashes with that used in Section 5.2 for restrictions to balls with respect to the tree metric.) Note that, by (5.13), we have that

$$\begin{aligned}
&\lim_{R\to\infty} \limsup_{i\to\infty} \mathbf{P}\left(\Delta\left(\underline{\mathcal{U}}_{n_i}^{(R)}, \underline{\mathcal{U}}_{n_i}\right) > \varepsilon\right) \\
&\leq \lim_{R\to\infty} \limsup_{i\to\infty} \left(\mathbf{1}_{\{e^{-\lambda^{-1}R^{\beta}} > \varepsilon\}} + \mathbf{P}\left(B_{\mathcal{U}}(0, \lambda^{-1}\delta_{n_i}^{-\beta}R^{\beta}) \not\subseteq B(\delta_{n_i}^{-1}R)\right)\right) \\
&\leq Ce^{-c\lambda^{a}}
\end{aligned}$$

for any $\varepsilon > 0$ and $\lambda \geq 1$, where we have applied Proposition 5.6.1 to deduce the final bound. In particular, since λ can be taken arbitrarily large in the above estimate, we obtain that

$$\lim_{R\to\infty}\limsup_{i\to\infty}\mathbf{P}\left(\Delta\left(\underline{\mathcal{U}}_{n_i}^{(R)},\underline{\mathcal{U}}_{n_i}\right)>\varepsilon\right)=0. \tag{5.94}$$

As a consequence, to prove the uniqueness of the law of $\underline{\mathcal{T}}$, it will be enough to show that, for each $R\in(0,\infty)$, $(\underline{\mathcal{U}}_{n_i}^{(R)})_{i\geq 1}$ converges in distribution to a uniquely specified limit. Indeed, if $\underline{\mathcal{T}}^{(R)}$ is the limit of $\underline{\mathcal{U}}_{n_i}^{(R)}$, then, since $\underline{\mathcal{U}}_{n_i}\xrightarrow{d}\underline{\mathcal{T}}$ (as $i\to\infty$) and (5.94) both hold, we have that $\underline{\mathcal{T}}^{(R)}\xrightarrow{d}\underline{\mathcal{T}}$ as $R\to\infty$.

Next, for given n_i and R, consider the measure $\pi_{n_i}^{(R)}$ on $B(\delta_{n_i}^{-1}R)\times\mathbb{R}^3$ given by

$$\pi_{n_i}^{(R)}(dxdy)=\frac{\mu_{\mathcal{U}}(dx)\delta_{\delta_{n_i}\phi_{\mathcal{U}}(x)}(dy)}{\mu_{\mathcal{U}}(B(\delta_{n_i}^{-1}R))},$$

where $\delta_z(\cdot)$ is the probability measure on $\mathbb{R}^3$ placing all its mass at z. We will check that the triple

$$\left(B(\delta_{n_i}^{-1}R),\delta_{n_i}^{\beta}d_{\mathcal{U}}|_{B(\delta_{n_i}^{-1}R)\times B(\delta_{n_i}^{-1}R)},\pi_{n_i}^{(R)}\right) \tag{5.95}$$

converges in the marked Gromov-weak topology of [58, Definition 2.4]; a characterisation of this convergence that will be relevant to us is given in the following paragraph. Towards establishing tightness, we first note that the projections of $\pi_{n_i}^{(R)}$ onto the sets $B(\delta_{n_i}^{-1}R)$ and $\mathbb{R}^3$ are simply the uniform probability measures on $B(\delta_{n_i}^{-1}R)$ and $\delta_{n_i}B(\delta_{n_i}^{-1}R)$, respectively. Since the latter measure clearly converges to the uniform probability measure on $B_E(R)$, by [58, Theorem 4] (see also [74, Theorem 3]), the desired tightness is implied by the following two conditions.

(a) The distributions of

$$\delta_{n_i}^{\beta}d_{\mathcal{U}}\left(\xi_1^{n_i,R},\xi_2^{n_i,R}\right),\qquad i\geq 1,$$

are tight, where $\xi_1^{n_i,R}$ and $\xi_2^{n_i,R}$ are independent uniform random variables on $B(\delta_{n_i}^{-1}R)$, independent of $\mathcal{U}$.

(b) For every $\varepsilon > 0$, there exists an $\eta > 0$ such that

$$\mathbf{E}\left(\delta_{n_i}^3 \mu_{\mathcal{U}}\left(\left\{x \in B(\delta_{n_i}^{-1}R) : \mu_{\mathcal{U}}\left(B_{\mathcal{U}}(x, \delta_{n_i}^{-\beta}\varepsilon) \cap B(\delta_{n_i}^{-1}R)\right) \leq \eta\right\}\right)\right) \leq \varepsilon.$$

The fact that (b) holds readily follows from the mass lower bound of Corollary 5.4.1. As for (a), this is a simple consequence of Corollary 5.7.1. Moreover, if we write $(\xi_j^{n_i,R})_{j\geq 1}$ for a sequence of independent uniform random variables on $B(\delta_{n_i}^{-1}R)$, independent of $\mathcal{U}$, then Corollary 5.7.1 further implies that

$$\left(\left(\delta_{n_i}^{\beta} d_{\mathcal{U}}\left(\xi_j^{n_i,R}, \xi_k^{n_i,R}\right)\right)_{j,k\geq 1}, \left(\xi_j^{n_i,R}\right)_{j\geq 1}\right) \tag{5.96}$$

converges in distribution. This enables us to deduce, by applying [58, Theorem 5, see also Remark 2.7], that the triple at (5.95) in fact converges in distribution in the marked Gromov-weak topology. We denote the limit by $(\mathcal{T}^{(R)}, d_{\mathcal{T}^{(R)}}, \pi_{\mathcal{T}^{(R)}})$, where $(\mathcal{T}^{(R)}, d_{\mathcal{T}^{(R)}})$ is a complete, separable metric space, and $\pi_{\mathcal{T}^{(R)}}$ is a probability measure on $\mathcal{T}^{(R)} \times \mathbb{R}^3$ such that $\pi_{\mathcal{T}^{(R)}}(\cdot \times \mathbb{R}^3)$ has full support on $\mathcal{T}^{(R)}$. In addition, by combining (5.43) with Proposition 5.4.5, we have the following adaptation of Assumption 3: there exists a continuous, increasing function $h(\eta)$ with $h(0) = 0$ such that

$$\lim_{\eta\to 0} \liminf_{\delta\to 0} \mathbf{P}\left(\sup_{\substack{x,y\in B(\delta^{-1}R):\\ \delta^{\beta} d_{\mathcal{U}}(x,y)<\eta}} \delta\left|\phi_{\mathcal{U}}(x) - \phi_{\mathcal{U}}(y)\right| \leq h(\eta)\right) = 1.$$

This allows us to apply [103, Theorem 3.7] to deduce that

$$\pi_{\mathcal{T}^{(R)}}(dxdy) = \mu_{\mathcal{T}^{(R)}}(dx)\delta_{\phi_{\mathcal{T}^{(R)}}(x)}(dy),$$

where $\mu_{\mathcal{T}^{(R)}}$ is a probability measure on $\mathcal{T}^{(R)}$ of full support, and $\phi_{\mathcal{T}^{(R)}} : \mathcal{T}^{(R)} \to \mathbb{R}^3$ is a continuous function.

As a consequence of the convergence described in the previous paragraph and the separability of the marked Gromov-weak topology (see [58, Theorem 2]), we can assume that all the random objects are built on the same probability space with probability space with probability measure $\mathbf{P}^*$ such that, $\mathbf{P}^*$-a.s.,

$$\left(B(\delta_{n_i}^{-1}R), \delta_{n_i}^{\beta} d_{\mathcal{U}}|_{B(\delta_{n_i}^{-1}R)\times B(\delta_{n_i}^{-1}R)}, \pi_{n_i}^{(R)}\right) \to \left(\mathcal{T}^{(R)}, d_{\mathcal{T}^{(R)}}, \pi_{\mathcal{T}^{(R)}}\right).$$

By [58, Lemma 3.4], this implies that, $\mathbf{P}^*$-a.s., there exists a complete and separable metric space (Z, d_Z) and isometric embeddings

$$\psi_{n_i} : (B(\delta_{n_i}^{-1}R), \delta_{n_i}^{\beta} d_{\mathcal{U}}) \to (Z, d_Z), \qquad \psi : (\mathcal{T}^{(R)}, d_{\mathcal{T}^{(R)}}) \to (Z, d_Z)$$

such that

$$\pi_{n_i}^{(R)} \circ (\tilde{\psi}_{n_i})^{-1} \to \pi_{\mathcal{T}^{(R)}} \circ \tilde{\psi}^{-1} \tag{5.97}$$

weakly as probability measures on $Z \times \mathbb{R}^3$, where $\tilde{\psi}_{n_i}(x, y) = (\psi_{n_i}(x), y)$ and $\tilde{\psi}(x, y) = (\psi(x), y)$. From our initial assumption that $(\underline{\mathcal{U}}_{n_i})_{i\geq 1}$ is distributionally convergent in $\mathbb{T}$, Corollary 5.4.1 and (5.43), we further have the existence of a deterministic subsequence $(n_{i_j})_{j\geq 1}$ such that, $\mathbf{P}^*$-a.s., $\underline{\mathcal{U}}_{n_{i_j}} \to \underline{\mathcal{T}}$ in $\mathbb{T}$,

$$\inf_{j\geq 1} \delta_{n_{i_j}}^3 \inf_{x\in B(\delta_{n_{i_j}}^{-1}R)} \mu_{\mathcal{U}}\left(B_{\mathcal{U}}(x, \delta_{n_{i_j}}^{-\beta}\delta)\right) > 0, \qquad \forall \delta > 0, \tag{5.98}$$

and also

$$\sup_{x\in B(\delta_{n_{i_j}}^{-1}R)} \delta_{n_{i_j}}^{\beta} d_{\mathcal{U}}(0, x) \to \Lambda \in (0, \infty). \tag{5.99}$$

Now, taking projections onto Z and rescaling, we readily obtain from (5.97) that

$$\delta_{n_i}^3 \mu_{\mathcal{U}}\left((\psi_{n_i})^{-1}(\cdot) \cap B(\delta_{n_i}R)\right) \to c\mu_{\mathcal{T}^{(R)}} \circ \psi^{-1} \tag{5.100}$$

weakly as probability measures on Z, where the constant c is the Lebesgue measure of $B_E(R)$. Moreover, appealing again to the mass lower bound of (5.98), we also obtain the subsequential convergence of measure supports, i.e.

$$\psi_{n_{i_j}}\left(B(\delta_{n_{i_j}}R)\right) \to \psi\left(\mathcal{T}^{(R)}\right)$$

with respect to the Hausdorff topology on compact subsets of Z (cf. the argument of [14, Theorem 6.1], for example). That $\mathcal{T}^{(R)}$ is indeed compact is established as in [14], and that it is a real tree follows from [67, Lemma 2.1]. In particular, if we define a sequence of correspondences by setting

$$\mathcal{C}_{n_{i_j}} := \left\{ \begin{array}{c} (x,x') \in B(\delta_{n_{i_j}} R) \times \mathcal{T}^{(R)} : \\ d_Z\left(\psi_{n_{i_j}}(x), \psi(x')\right) \leq 2 d_H^Z \left(\psi_{n_{i_j}}(B(\delta_{n_{i_j}} R)), \psi(\mathcal{T}^{(R)})\right) \end{array} \right\},$$

where d_H^Z is the Hausdorff distance on Z, then we have that

$$\sup_{(x,x') \in \mathcal{C}_{n_{i_j}}} d_Z\left(\psi_{n_{i_j}}(x), \psi(x')\right) \to 0. \tag{5.101}$$

Given that $\underline{\mathcal{U}}_{n_{i_j}} \to \underline{\mathcal{T}}$ in $\mathbb{T}$ and (5.99) holds, it is a straightforward application of [24, Lemmas 3.5 and 5.1] to also check that, $\mathbf{P}^*$-a.s.,

$$\lim_{\eta \to 0} \limsup_{j \to \infty} \sup_{\substack{x,y \in B(\delta_{n_{i_j}}^{-1} R): \\ \delta_{n_{i_j}}^{\beta} d_{\mathcal{U}}(x,y) < \eta}} \delta_{n_{i_j}} \left|\phi_{\mathcal{U}}(x) - \phi_{\mathcal{U}}(y)\right| = 0,$$

and, applying this equicontinuity in conjunction with (5.97), this yields in turn that

$$\sup_{(x,x') \in \mathcal{C}_{n_{i_j}}} \left|\phi_{\mathcal{U}}(x) - \phi_{\mathcal{T}^{(R)}}(x')\right| \to 0. \tag{5.102}$$

Finally, although not included in the framework of [58, 74, 103], it is not difficult to include the convergence of roots in the above arguments, i.e. we may further suppose that

$$d_Z\left(\psi_{n_{i_j}}(\rho_{\mathcal{U}}), \psi(\rho_{\mathcal{T}^{(R)}})\right) \to 0 \tag{5.103}$$

for some $\rho_{\mathcal{T}^{(R)}} \in \mathcal{T}^{(R)}$ with $\phi_{\mathcal{T}^{(R)}}(\rho_{\mathcal{T}^{(R)}}) = 0$. Recalling the definition of Δ_c from (5.12), combining (5.100), (5.101), (5.102) and (5.103) yields that

$$\Delta_c(\underline{\mathcal{U}}_{n_{i_j}}^{(R)}, \underline{\mathcal{T}}^{(R)}) \to 0 \qquad \mathbf{P}^*\text{-a.s.},$$

where $\underline{\mathcal{T}}^{(R)} := (\mathcal{T}^{(R)}, d_{\mathcal{T}^{(R)}}, \mu_{\mathcal{T}^{(R)}}, \phi_{\mathcal{T}^{(R)}}, \rho_{\mathcal{T}^{(R)}})$. Since the distribution of

$\underline{\mathcal{T}}^{(R)}$ is uniquely specified by (5.96), and the same limit can be deduced for some subsubsequence of any subsequence of $(n_i)_{i\geq 1}$, we obtain that $\underline{\mathcal{U}}_{n_i}^{(R)} \to \underline{\mathcal{T}}^{(R)}$ in distribution in $\mathbb{T}$, and thus the part of the proof concerning the Gromov-Hausdorff-type topology is complete.

As for the path ensemble topology, we know from [24, Lemma 3.9] that convergence of compact measured, rooted spatial trees with respect to our Gromov-Hausdorff-type implies the corresponding path ensemble statement. To extend from this to the desired conclusion, we can proceed exactly as in the proof of [24, Lemma 5.5], with the additional inputs required being provided by (5.43) and the coupling lemma that is stated below at Lemma 5.9.3. □

5.9 Properties of the limiting space

The aim of this section is to prove Theorem 5.1.2. To this end, we present several preparatory lemmas. In the first of these, we check that for large enough annuli there is only one disjoint crossing by a path in $\mathcal{U}$. Precisely, for $r < R$, we introduce the event $\mathcal{C}_{\mathcal{U}}^{E}(r, R)$ by setting

$$\mathcal{C}_{\mathcal{U}}^{E}(r, R) = \{\exists x, y \in B(R)^c \text{ such that } \gamma_{\mathcal{U}}(x, y) \cap B(r) \neq \emptyset\},$$

and show that the probability of this occurring decays as the ratio R/r increases.

Lemma 5.9.1. *There exist universal constants $\lambda_0 > 0$ and $a, b, C \in (0, \infty)$ such that for all $\delta \in (0, 1)$ and $\lambda \geq \lambda_0$,*

$$\mathbf{P}\left(\mathcal{C}_{\mathcal{U}}^{E}\left(\lambda^{-a}\delta^{-1}, \delta^{-1}\right)\right) \leq C\lambda^{-b}.$$

Proof. This is essentially established in the proof of Proposition 5.4.1. We will use the same notation as in that proof here. First, suppose that the event A'_{k_0}, as defined in the proof of Proposition 5.4.1, occurs. It then holds that: for every point $x \in \partial B(\delta^{-1})$,

$$\gamma_{\mathcal{U}}\left(x, \gamma_\infty\right) \cap B\left(\lambda^{-4}\delta^{-1}\right) = \emptyset,$$

where γ_∞ is the unique infinite simple path in $\mathcal{U}$ started at the origin, and $\gamma_{\mathcal{U}}(x, \gamma_\infty)$ is shortest path in $\mathcal{U}$ from x to a point of γ_∞. Note that we have already proved that $\mathbf{P}(A'_{k_0}) \geq 1 - C\lambda^{-1}$. Second, let u be the first time that γ_∞ exits $B(\lambda^{-4}\delta^{-1})$, and define

$$W = \left\{ \gamma_\infty[u, \infty) \cap B(\lambda^{-5}\delta^{-1}) = \emptyset \right\}.$$

By Proposition 1.5.10 of [113], it holds that $\mathbf{P}(W) \geq 1 - C\lambda^{-1}$. Finally, suppose that the event $A'_{k_0} \cap W$ occurs. For $x, y \in B(\delta^{-1})^c$, let $x', y' \in \gamma_\infty$ be such that $\gamma_{\mathcal{U}}(x, \gamma_\infty) = \gamma_{\mathcal{U}}(x, x')$ and $\gamma_{\mathcal{U}}(y, \gamma_\infty) = \gamma_{\mathcal{U}}(y, y')$. We then have that $\gamma_{\mathcal{U}}(x, x') \cap B(\lambda^{-4}\delta^{-1}) = \emptyset$ and $\gamma_{\mathcal{U}}(y, y') \cap B(\lambda^{-4}\delta^{-1}) = \emptyset$. Also, it holds that $x', y' \in \gamma_\infty[u, \infty)$. In particular, it follows that $\gamma_{\mathcal{U}}(x, y) \cap B(\lambda^{-5}\delta^{-1}) = \emptyset$ for all $x, y \in B(\delta^{-1})^c$. This completes the proof of the result with $a = 5$ and $b = 1$. $\square$

We next establish a result which essentially gives the converse of Assumption 3. In particular, we define the event $\mathcal{D}(a, b, c)$ by

$$\mathcal{D}(a, b, c) = \left\{ \exists x, y \in B(a) \text{ such that } d^{\mathrm{S}}_{\mathcal{U}}(x, y) < b \text{ and } d_{\mathcal{U}}(x, y) > c \right\},$$

where we define the Schramm metric $d^S_{\mathcal{U}}$ on $\mathcal{U}$ analogously to (5.2), and check the following.

Lemma 5.9.2. *There exist universal $\lambda_0 > 0$ and $a_1, \dots, a_4, C \in (0, \infty)$ such that for all $\delta \in (0,1)$ and $\lambda \geq \lambda_0$,*

$$\mathbf{P}\left(\mathcal{D}\left(\lambda^{a_1}\delta^{-1}, \lambda^{-a_2}\delta^{-1}, \lambda^{-a_3}\delta^{-\beta} \right) \right) \leq C\lambda^{-a_4}.$$

Proof. Consider the event $\hat{\mathcal{D}}(a, b, c)$ given by

$$\hat{\mathcal{D}}(a, b, c) = \left\{ \exists x, y \in B(a) \cap \gamma_\infty \text{ such that } d^{\mathrm{S}}_{\mathcal{U}}(x, y) < b \text{ and } d_{\mathcal{U}}(x, y) > c \right\}.$$

We first prove that there exist universal $a_1, \dots, a_4, C \in (0, \infty)$ such that for all $\delta \in (0,1)$ and $\lambda \geq 1$,

$$\mathbf{P}\left(\hat{\mathcal{D}}\left(\lambda^{a_1}\delta^{-1}, \lambda^{-a_2}\delta^{-1}, \lambda^{-a_3}\delta^{-\beta} \right) \right) \leq C\lambda^{-a_4}. \tag{5.104}$$

To do this, let $a_1 = 10^{-4}$, $a_2 = 1$ and $a_3 = 1/2$. Moreover, let $D = (w_k)_{k=1}^{M}$ be a $\lambda^{-a_2}\delta^{-1}$-net of $B(\lambda^{a_1}\delta^{-1})$ such that $B(\lambda^{a_1}\delta^{-1}) \subseteq \bigcup_{k=1}^{M} B(w_k, \lambda^{-a_2}\delta^{-1})$ and $M \asymp \lambda^{3(a_1+a_2)}$. Suppose that the event $\hat{\mathcal{D}}\left(\lambda^{a_1}\delta^{-1}, \lambda^{-a_2}\delta^{-1}, \lambda^{-a_3}\delta^{-\beta}\right)$ occurs. Then there exists $w_k \in D$ such that $|\gamma_\infty \cap B(w_k, \lambda^{-a_2}\delta^{-1})| \geq c\lambda^{-a_3}\delta^{-\beta}$ for some universal $c > 0$. Now, it follows from [127, (7.51)] that

$$\mathbf{P}\left(\exists w_k \in D \text{ such that } \left|\gamma_\infty \cap B(w_k, \lambda^{-a_2}\delta^{-1})\right| \geq c\lambda^{-a_3}\delta^{-\beta}\right) \leq Ce^{-c'\lambda^{1/2}},$$

for some universal $c', C \in (0, \infty)$. Thus, the inequality (5.104) holds when we let $a_4 = 100$.

We next consider a $\lambda^{-4}\delta^{-1}$-net $D' = (x_i)_{i=1}^{N}$ of the ball $B(\lambda^{a_1}\delta^{-1})$ for which $B(\lambda^{a_1}\delta^{-1})$ is a subset of $\bigcup_{i=1}^{N} B(x_i, \lambda^{-4}\delta^{-1})$ and $N \asymp \lambda^{3(a_1+4)}$. We perform Wilson's algorithm as follows:

- Consider a subtree spanned by $D' = (x_i)_{i=1}^{N}$. The output random tree is denoted by $\mathcal{U}_1$.
- Perform Wilson's algorithm for all remaining points $\mathbb{Z}^3 \setminus D'$ to generate $\mathcal{U}$.

We define the event L by

$$L = \bigcap_{i=1}^{N} \hat{\mathcal{D}}\left(\lambda^{a_1}\delta^{-1}, \lambda^{-a_2}\delta^{-1}, \lambda^{-a_3}\delta^{-\beta}; i\right)^c,$$

where the event $\hat{\mathcal{D}}(a, b, c; i)$ is defined by

$$\hat{\mathcal{D}}(a, b, c; i) = \left\{ \begin{array}{c} \exists x, y \in B(a) \cap \gamma_\infty^{x_i} \\ \text{such that } d_{\mathcal{U}}^{\mathrm{S}}(x, y) < b \text{ and } d_{\mathcal{U}}(x, y) > c \end{array} \right\},$$

with γ_∞^x standing for the unique infinite simple path in $\mathcal{U}$ started at x. By (5.104), we have $\mathbf{P}(L) \geq 1 - C\lambda^{-80}$. Furthermore, if we define

$$J = \left\{ \begin{array}{c} \forall x \in B(\lambda^{a_1}\delta^{-1}), \\ \operatorname{diam}(\gamma_{\mathcal{U}}(x, \mathcal{U}_1)) < \lambda^{-2}\delta^{-1} \text{ and } d_{\mathcal{U}}(x, \mathcal{U}_1) < \lambda^{-2}\delta^{-\beta} \end{array} \right\},$$

then applying the hittability of each branch of $\mathcal{U}$ as in the proof of Proposition 5.4.1 guarantees that $\mathbf{P}(J) \geq 1 - C\lambda^{-10}$. Finally, suppose that the event $L \cap J$ occurs. The event L ensures that for all $x, y \in \mathcal{U}_1$ with $d^{\mathrm{S}}_{\mathcal{U}}(x,y) < \lambda^{-a_2}\delta^{-1}$, we have $d_{\mathcal{U}}(x,y) < 2\lambda^{-a_3}\delta^{-\beta}$. Also, the event J guarantees that for all $x, y \in B(\lambda^{a_1}\delta^{-1})$ with $d^{\mathrm{S}}_{\mathcal{U}}(x,y) < \frac{1}{2}\lambda^{-a_2}\delta^{-1}$, we have $d_{\mathcal{U}}(x,y) < 3\lambda^{-a_3}\delta^{-\beta}$. Thus the proof is complete, establishing the result with $a_1 = 10^{-4}$, $a_2 = 1$, $a_3 = 1/2$ and $a_4 = 10$. □

For the remainder of the section, including in the proof of Theorem 5.1.2, we fix a sequence $\delta_n \to 0$ such that $(\mathbf{P}_{\delta_n})_{n\geq 1}$ converges weakly (as measures on $(\mathbb{T}, \Delta)$), and write $\mathcal{U}_{\delta_n} = (\mathcal{U}, \delta_n^{\kappa} d_{\mathcal{U}}, \delta_n^2 \mu_{\mathcal{U}}, \delta_n \phi_{\mathcal{U}}, 0)$. Letting $\hat{\mathbf{P}}$ be the relevant limiting law, we denote by $\underline{\mathcal{T}} = (\mathcal{T}, d_{\mathcal{T}}, \mu_{\mathcal{T}}, \phi_{\mathcal{T}}, \rho_{\mathcal{T}})$ a random element of $\mathbb{T}$ with law $\hat{\mathbf{P}}$. A key ingredient to the proof of Theorem 5.1.2 is the following coupling between the discrete and continuous models, which is a ready consequence of this convergence assumption. Since the proof of the corresponding result in [24] was not specific to the two-dimensional case, we omit the proof here.

Lemma 5.9.3 (cf. [24, Lemma 5.1]). *There exist realisations of $(\mathcal{U}_{\delta_n})_{n\geq 1}$ and $\underline{\mathcal{T}}$ built on the same probability space, with probability measure $\mathbf{P}^*$ say, such that: for some subsequence $(n_i)_{i\geq 1}$ and divergent sequence $(r_j)_{j\geq 1}$ it holds that, $\mathbf{P}^*$-a.s.,*

$$D_{i,j} := \Delta_c \left(\mathcal{U}_{\delta_{n_i}}^{(r_j)}, \underline{\mathcal{T}}^{(r_j)} \right) \to 0$$

as $i \to \infty$, for every $j \geq 1$.

Proof of Theorem 5.1.2. We start by checking the measure bounds of parts (c) and (d), and we also remark that part (b) is an elementary consequence of (c) (see [64, Proposition 1.5.15], for example). The uniform bound of (c) will follow from the estimates: for $R > 0$, there exist constants $c_i \in (0, \infty)$

such that, for every $r \in (0,1)$,

$$\hat{\mathbf{P}}\left(\inf_{x\in B_{\mathcal{T}}(\rho_{\mathcal{T}},R)} \mu_{\mathcal{T}}\left(B_{\mathcal{T}}(x,r)\right) \leq c_1 r^{d_f}(\log r^{-1})^{-c_2}\right) \leq c_3 r^{c_4}, \quad (5.105)$$

$$\hat{\mathbf{P}}\left(\sup_{x\in B_{\mathcal{T}}(\rho_{\mathcal{T}},R)} \mu_{\mathcal{T}}\left(B_{\mathcal{T}}(x,r)\right) \geq c_5 r^{d_f}(\log r^{-1})^{c_6}\right) \leq c_7 r^{c_8}. \quad (5.106)$$

Indeed, given these, applying Borel-Cantelli along the subsequence $r_n = 2^{-n}$, $n \in \mathbb{N}$, yields the result. By appealing to the coupling of Lemma 5.9.3, the above inequalities readily follow from the following discrete analogues:

$$\limsup_{\delta\to\infty} \mathbf{P}\left(\begin{array}{c} \delta^3 \min_{x\in B_{\mathcal{U}}(\rho_{\mathcal{U}},\delta^{-\beta}R)} \mu_{\mathcal{T}}\left(B_{\mathcal{U}}(x,\delta^{-\beta}r)\right) \\ \leq c_1 r^{d_f}(\log r^{-1})^{-c_2} \end{array}\right) \leq c_3 r^{c_4}, \quad (5.107)$$

$$\limsup_{\delta\to\infty} \mathbf{P}\left(\begin{array}{c} \delta^3 \max_{x\in B_{\mathcal{U}}(\rho_{\mathcal{U}},\delta^{-\beta}R)} \mu_{\mathcal{T}}\left(B_{\mathcal{U}}(x,\delta^{-\beta}r)\right) \\ \geq c_5 r^{d_f}(\log r^{-1})^{c_6} \end{array}\right) \leq c_7 r^{c_8}. \quad (5.108)$$

To establish these, we start by noting that Proposition 5.6.1 implies that the probability in (5.107) is bounded above by

$$Ce^{-cz^a} + \mathbf{P}\left(\delta^3 \min_{x\in B(\delta^{-1}R^{1/\beta}z)} \mu_{\mathcal{T}}\left(B_{\mathcal{U}}(x,\delta^{-\beta}r)\right) \leq c_1 r^{d_f}(\log r^{-1})^{-c_2}\right)$$

for any $z \geq 1$. Moreover, applying a simple union bound and Theorem 5.5.2 (with $R = \delta^{-\beta}r$, $\lambda = c_1^{-1}\log(r^{-1})^{c_2}$), we can bound this in turn by

$$Ce^{-cz^a} + \frac{C' R^{d_f} z^3}{r^{d_f}} e^{-c' c_1^{-a'} \log(r^{-1})^{a' c_2}}.$$

Choosing $z = (c^{-1}\log(r^{-1}))^{1/a}$, c_1 small enough so that $c' c_1^{-a} > d_f$, and $c_2 = 1/a'$, the above is bounded above by $C'' r^{c''}$, as desired. The proof of (5.108) is similar, with Theorem 5.6.2 replacing Theorem 5.5.2. As for (d), this follows from a Borel-Cantelli argument and the following estimates:

there exist constants $c_i \in (0,\infty)$ such that

$$\hat{\mathbf{P}}\left(\mu_{\mathcal{T}}\left(B_{\mathcal{T}}(\rho_{\mathcal{T}},r)\right) \geq \lambda r^{d_f}\right) \leq c_1 e^{-c_2\lambda^{c_3}}, \tag{5.109}$$

$$\hat{\mathbf{P}}\left(\mu_{\mathcal{T}}\left(B_{\mathcal{T}}(\rho_{\mathcal{T}},r)\right) \leq \lambda^{-1} r^{d_f}\right) \leq c_4 e^{-c_5\lambda^{c_6}}, \tag{5.110}$$

for all $r>0$, $\lambda \geq 1$. Similarly to the proof of the uniform estimates (5.105) and (5.106), applying the coupling of Lemma 5.9.3, these readily follow from Theorem 5.5.1 and Proposition 5.6.1.

For part (a), since $(\mathcal{U}, d_{\mathcal{U}})$ has infinite diameter, we immediately find that $(\mathcal{T}, d_{\mathcal{T}})$ has at least one end at infinity. Thus we need to show that there can be no more than one end at infinity. Given Lemma 5.9.3 and the inclusion results of (5.43) and Proposition 5.6.1, this can be proved exactly as in the two-dimensional case. In particular, as in [24], it follows from the following crossing estimate: for $r>0$,

$$\lim_{R\to\infty}\limsup_{\delta\to 0}\mathbf{P}\left(C^E_{\mathcal{U}}(\delta^{-1}r,\delta^{-1}R)\right)=0,$$

which is given by Lemma 5.9.1.

For part (e), we can proceed exactly as in the proof of [24, Lemma 5.4]. Given Lemma 5.9.3, the one additional ingredient we need to do this is the estimate corresponding to [24, (5.12)]: for every $r,\eta>0$,

$$\lim_{\varepsilon\to 0}\limsup_{\delta\to 0}\mathbf{P}\left(\inf_{\substack{x,y\in B_{\mathcal{U}}(0,\delta^{-\beta}r):\\ d_{\mathcal{U}}(x,y)\geq\delta^{-\beta}\eta}} d^S_{\mathcal{U}}(x,y) < \delta^{-1}\varepsilon\right)=0,$$

and this was established in Lemma 5.9.2 (when viewed in conjunction with Proposition 5.6.1).

Given Lemma 5.9.3 and (5.6.1), the proof of part (f) is identical to that of [24, Lemma 5.2]. $\square$

5.10 Simple random walk and its diffusion limit

In this section, we complete the article with the proofs of Theorem 5.1.3, Corollary 5.1.1 and Theorem 5.1.4.

Proof of Theorem 5.1.3. On the event

$$\left\{\inf_{x\in B_{\mathcal{U}}(0,R)} \mu_{\mathcal{U}}\left(B_{\mathcal{U}}(x,R/8)\right) \geq \lambda^{-1}R^{d_f},\ \mu_{\mathcal{U}}\left(B_{\mathcal{U}}(0,2R)\right) \leq \lambda R^{d_f}\right\}, \quad (5.111)$$

one can find a cover $(B_{\mathcal{U}}(x_i,R/4))_{i=1}^N$ of $B_{\mathcal{U}}(0,R)$ of size $N \leq \lambda^2$ (cf. [49, Lemma 9], for example). Following the argument of [22, Lemma 2.4] (see alternatively [108, Lemma 4.1]), it holds that on the event at (5.111),

$$R_{\mathcal{U}}\left(0, B_{\mathcal{U}}(0,R)^c\right) \geq \frac{R}{\lambda^2}.$$

Hence the result is a consequence of Theorem 5.5.2 and Proposition 5.6.1. □

Proof of Corollary 5.1.1. By Theorem 5.1.3, parts (1) and (4) of [109, Assumption 1.2] hold. Moreover, since $R_{\mathcal{U}}(0,B_{\mathcal{U}}(0,R)^c) \leq R+1$, we also have that part (2) of [109, Assumption 1.2] holds. Hence (5.4), (5.6), (5.7), (5.9) and (5.11) follow from [109, Proposition 1.4 and Theorem 1.5]. It remains to prove the claims involving the Euclidean distance. To this end, note that by (5.43) and Proposition 5.6.1,

$$\mathbf{P}\left(B_{\mathcal{U}}(0,\lambda^{-1}R^{\beta}) \subseteq B(R) \subseteq B_{\mathcal{U}}(0,\lambda R^{\beta})\right) \geq 1 - c_1\lambda^{-c_2}.$$

Hence, by Borel-Cantelli, if $R_n := 2^n$ and $\lambda_n := n^{2/c_2}$, then

$$B_{\mathcal{U}}(0,\lambda_n^{-1}R_n^{\beta}) \subseteq B(R_n) \subseteq B_{\mathcal{U}}(0,\lambda_n R_n^{\beta})$$

for all large n, $\mathbf{P}$-a.s. Combining this with the results at (5.4) and (5.7), we obtain (5.5) and (5.8). As for (5.10), the lower bound follows from Jensen's inequality, Fatou's lemma and (5.8). Indeed,

$$\begin{aligned}\liminf_{R\to\infty}\frac{\log \mathbb{E}^{\mathcal{U}}\left(\tau^E_{0,R}\right)}{\log R} &\geq \liminf_{R\to\infty}\mathbb{E}^{\mathcal{U}}\left(\frac{\log \tau^E_{0,R}}{\log R}\right)\\ &\geq \mathbb{E}^{\mathcal{U}}\left(\liminf_{R\to\infty}\frac{\log \tau^E_{0,R}}{\log R}\right)=\beta d_w.\end{aligned}$$

As for the upper bound, a standard estimate for exit times (see [19, Corollary 2.66], for example) gives that

$$E_0^{\mathcal{U}}\tau^E_{0,R}\leq R^3 R_{\mathcal{U}}(0,B(R)^c)\leq R^3\xi_R,$$

where ξ_R is defined above Proposition 5.3.4. The latter result thus yields

$$\mathbb{E}^{\mathcal{U}}\left(\tau^E_{0,R}\right)\leq R^3\mathbf{E}\left(\xi_R\right)\leq cR^{3+\beta}=cR^{\beta d_w},$$

which gives (a stronger statement than) the desired conclusion. □

Proof of Theorem 5.1.4. The result can be proved by a line-by-line modification of [24, Theorems 1.4 and 7.2], and so we omit the details. However, as an aid to the reader, we summarise the key steps. As per the construction of [98], $\hat{\mathbf{P}}$-a.s., there is a 'resistance form' $(\mathcal{E}_{\mathcal{T}},\mathcal{F}_{\mathcal{T}})$ on $(\mathcal{T},d_{\mathcal{T}})$, characterised by

$$d_{\mathcal{T}}(x,y)^{-1}=\inf\left\{\mathcal{E}_{\mathcal{T}}(f,f):\ f\in\mathcal{F}_{\mathcal{T}},\ f(x)=0,\ f(y)=1\right\},$$

for all $x,y\in\mathcal{T}$, $x\neq y$. Moreover, by taking

$$\mathcal{D}_{\mathcal{T}}:=\overline{\mathcal{F}_{\mathcal{T}}\cap C_0(\mathcal{T})},$$

where $C_0(\mathcal{T})$ are the compactly supported continuous functions on $(\mathcal{T},d_{\mathcal{T}})$, and the closure is taken with respect to $\mathcal{E}_{\mathcal{T}}(f,f)+\int_{\mathcal{T}}f^2d\mu_{\mathcal{T}}$, we obtain a regular Dirichlet form $(\mathcal{E}_{\mathcal{T}},\mathcal{D}_{\mathcal{T}})$ on $L^2(\mathcal{T},\mu_{\mathcal{T}})$ (see [13, Remark 1.6] or [99, Theorem 9.4]). Moreover, since $(\mathcal{T},d_{\mathcal{T}})$ is complete and has one end at infinity (by Theorem 5.1.2(a)), the naturally associated stochastic process $((X_t^{\mathcal{T}})_{t\geq 0},(P_x^{\mathcal{T}})_{x\in\mathcal{T}})$ is recurrent (see [13, Theorem 4]). And, from [99, Theorem 10.4], we have that the process admits a jointly continuous transition density $(p_t^{\mathcal{T}}(x,y))_{x,y\in\mathcal{T},t>0}$.

Next, by appealing to the Skorohod representation theorem, it is possible to construct realisations of $(\mathcal{U}, \delta_n^\beta d_{\mathcal{U}}, \delta_n^3 \mu_{\mathcal{U}}, \delta_n \phi_{\mathcal{U}}, \rho_{\mathcal{U}})$, $n \geq 1$, and the limit $(\mathcal{T}, d_{\mathcal{T}}, \mu_{\mathcal{T}}, \phi_{\mathcal{T}}, \rho_{\mathcal{T}})$ on the same probability space with probability measure $\mathbf{P}^*$ such that

$$(\mathcal{U}, \delta_n^\beta d_{\mathcal{U}}, \delta_n^3 \mu_{\mathcal{U}}, \delta_n \phi_{\mathcal{U}}, \rho_{\mathcal{U}}) \to (\mathcal{T}, d_{\mathcal{T}}, \mu_{\mathcal{T}}, \phi_{\mathcal{T}}, \rho_{\mathcal{T}}) \qquad \mathbf{P}^*\text{-a.s.}$$

Moreover, applying Theorem 5.1.3 in a simple Borel-Cantelli argument allows one to deduce that, $\mathbf{P}^*$-a.s.,

$$\lim_{R\to\infty} \liminf_{n\to\infty} \delta_n^\beta R_{\mathcal{U}} \left(0, B_{\mathcal{U}}(0, R\delta_n^{-\beta})^c\right) = \infty.$$

Hence we can apply [54, Theorem 7.1] to deduce that, $\mathbf{P}^*$-a.s.,

$$P_0^{\mathcal{U}} \left(\left(\delta_n X^{\mathcal{U}}_{t\delta_n^{-(3+\beta)}} \right)_{t\geq 0} \in \cdot \right) \to P_{\rho_{\mathcal{T}}}^{\mathcal{T}} \circ \phi_{\mathcal{T}}^{-1} \tag{5.112}$$

weakly as probability measures on $C(\mathbb{R}_+, \mathbb{R}^3)$. Since the left-hand side above is $\mathbf{P}^*$-measurable, so is the right-hand side. Moreover, for any measurable set $B \subseteq C(\mathbb{R}_+, \mathbb{R}^3)$, we have that

$$P_{\rho_{\mathcal{T}}}^{\mathcal{T}} \circ \phi_{\mathcal{T}}^{-1}(B) = \mathbf{E}^* \left(P_{\rho_{\mathcal{T}}}^{\mathcal{T}} \circ \phi_{\mathcal{T}}^{-1}(B) \middle| \mathcal{T} \right),$$

where $\mathbf{E}^*$ is the expectation under $\mathbf{P}^*$, and so $P_{\rho_{\mathcal{T}}}^{\mathcal{T}} \circ \phi_{\mathcal{T}}^{-1}$ is in fact $\hat{\mathbf{P}}$-measurable, as is required to prove part (a). For part (b), we apply (5.112) and integrate out with respect to $\mathbf{P}^*$.

As for the heat kernel estimates, we note that the measure bounds of Theorem 5.1.2(c) are enough to apply the arguments of [49] to deduce part (c) (for further details, see the proof of [24, Theorem 1.4(c)]). As for the on-diagonal estimates of part (d), similarly to the proof of [24, Theorem 7.2] (cf. [51, Theorems 1.6 and 1.7]), these follow from the distributional estimates on the measures of balls at (5.109) and (5.110), together with the following resistance estimate

$$\mathbf{P}\left(R_{\mathcal{T}}(\rho_{\mathcal{T}}, B_{\mathcal{T}}(\rho_{\mathcal{T}}, R)^c) \leq \lambda^{-1}R\right) \leq Ce^{-c\lambda^a}, \tag{5.113}$$

where $R_{\mathcal{T}}$ is the resistance associated with $(\mathcal{E}_{\mathcal{T}}, \mathcal{F}_{\mathcal{T}})$. As in the proof of Theorem 5.1.3, to check (5.113), it is enough to combine (5.109) with the bound

$$\hat{\mathbf{P}}\left(\inf_{x\in B_{\mathcal{T}}(\rho_{\mathcal{T}},R)} \mu_{\mathcal{T}}\left(B_{\mathcal{T}}(x, R/8)\right) \leq \lambda^{-1}R^{d_f}\right) \leq Ce^{-c\lambda^a},$$

which is again a ready consequence of the discrete analogue (see Theorem 5.5.2 and Proposition 5.6.1). □

Chapter 6

The Number of Spanning Clusters of the Uniform Spanning Tree in Three Dimensions[1]

Summary of this chapter

Let $\mathcal{U}_\delta$ be the uniform spanning tree on $\delta\mathbb{Z}^3$. A spanning cluster of $\mathcal{U}_\delta$ is a connected component of the restriction of $\mathcal{U}_\delta$ to the unit cube $[0,1]^3$ that connects the left face $\{0\} \times [0,1]^2$ to the right face $\{1\} \times [0,1]^2$. In this note, we will prove that the number of the spanning clusters is tight as $\delta \to 0$, which resolves an open question raised by Benjamini in [28].

[1] Joint work with Omer Angel, David Croydon, and Daisuke Shiraishi.
Acknowledgements. DC would like to acknowledge the support of a JSPS Grant-in-Aid for Research Activity Start-up, 18H05832 and a JSPS Grant-in-Aid for Scientific Research (C), 19K03540. DS is supported by a JSPS Grant-in-Aid for Early-Career Scientists, 18K13425 and JSPS KAKENHI Grant Number 17H02849 and 18H01123.

6.1 Introduction

Given a finite connected graph $G = (V, E)$, a spanning tree T of G is a subgraph of G that is a tree (i.e. is connected and contains no cycles) with vertex set V. A uniform spanning tree (UST) of G is obtained by choosing a spanning tree of G uniformly at random. This is an important model in probability and statistical physics, with beautiful connections to other subjects, such as electrical potential theory, loop-erased random walk and Schramm-Loewner evolution. See [29] for an introduction to various aspects of USTs.

Fix $\delta \in (0,1)$ and $d \in \mathbb{N}$. In [143] it was shown that, by taking the local limit of the uniform spanning trees on an exhaustive sequence of finite subgraphs of $\delta\mathbb{Z}^d$, it is possible to construct a random subgraph $\mathcal{U}_\delta$ of $\delta\mathbb{Z}^d$. Whilst the resulting graph $\mathcal{U}_\delta$ is almost-surely a forest consisting on an infinite number of disjoint components that are trees when $d \geq 5$, it is also the case that $\mathcal{U}_\delta$ is almost-surely a spanning tree of $\delta\mathbb{Z}^d$ with one topological end for $d \leq 4$, see [143]. In the latter low-dimensional case, $\mathcal{U}_\delta$ is commonly referred to as the UST on $\delta\mathbb{Z}^d$.

In this note, we study a macroscopic scale property of $\mathcal{U}_\delta$, namely the number of its spanning clusters, as previously studied by Benjamini in [28]. To be more precise, let us proceed to introduce some notation. Write

$$\mathbb{B} = [0,1]^d = \Big\{(x_1, x_2, \cdots, x_d) \in \mathbb{R}^d \,:\, 0 \leq x_i \leq 1,\ i = 1, 2, \cdots, d\Big\} \tag{6.1}$$

for the unit hypercube in $\mathbb{R}^d$. Also, set

$$F = \Big\{(x_1, x_2, \cdots, x_d) \in \mathbb{R}^d \;:\; x_1 = 0\Big\} \tag{6.2}$$

and

$$G = \Big\{(x_1, x_2, \cdots, x_d) \in \mathbb{R}^d \;:\; x_1 = 1\Big\} \tag{6.3}$$

for the hyperplanes intersecting the 'left' and 'right' sides of the hypercube $\mathbb{B}$. Given a subgraph $U = (V, E)$ of $\delta\mathbb{Z}^d$, we write $U' = (V', E')$ for the restriction of U to the cube $\mathbb{B}$, i.e. we set $V' = V \cap \mathbb{B}$ and $E' = \{\{x, y\} \in E : x, y \in V'\}$. A connected component of U' is called a cluster of U. Moreover, following [28], a spanning cluster of U is a cluster of U containing vertices x and y such that $\operatorname{dist}(x, F) < \delta$ and $\operatorname{dist}(y, G) < \delta$, where $\operatorname{dist}(z, A) := \inf_{w \in A} |z - w|$ is the Euclidean distance between a point $z \in \mathbb{R}^d$ and subset $A \subseteq \mathbb{R}^d$. That is, a cluster of U is called spanning when it connects F to G (at the level of discretization being considered).

Concerning the number of spanning clusters of $\mathcal{U}_\delta$, it was proved in [28] that:

- for $d \geq 4$, the expected number of spanning clusters of $\mathcal{U}_\delta$ grows to infinity as $\delta \to 0$;
- for $d = 2$, the number of spanning clusters of $\mathcal{U}_\delta^+$ is tight as $\delta \to 0$, where $\mathcal{U}_\delta^+$ denotes the uniform spanning tree of the square $\mathbb{B} \cap \delta\mathbb{Z}^2$ when all the vertices on the right side of the square are identified to a single point $\mathfrak{w}$ ($\mathcal{U}_\delta^+$ is called the right wired uniform spanning tree in [28]). We also consider two spanning clusters of $\mathcal{U}_\delta^+$ different if they are disjoint on $\mathcal{U}_\delta^+ \setminus \mathfrak{w}$. Figure 6.1 shows the spanning cluster of a realisation of (an approximation to) $\mathcal{U}_\delta$ on $\delta\mathbb{Z}^2$.

The case $d = 3$ was left as an open question in [28]. The main purpose of this note is to resolve it by showing the following theorem.

Theorem 6.1.1. *Let $d = 3$. It holds that the number of spanning clusters of $\mathcal{U}_\delta$ is tight as $\delta \to 0$.*

Remark. The proof for Theorem 6.1.1 can be adapted to show tightness of the number of spanning clusters of $\mathcal{U}_\delta$ on $\delta\mathbb{Z}^2$. This is an improvement over the result in [28], which required right-wired boundary conditions.

Remark. Part of Benjamini's motivation for studying the number of spanning clusters came from percolation. Indeed, for critical Bernoulli percolation in $\mathbb{Z}^d$, it is conjectured that the number of spanning clusters is tight when $d \leq 6$, while the expected number of spanning clusters grows to infinity as the mesh size goes to zero for $d > 6$, see for instance [3, 36, 43]. Putting our main conclusion together with the results obtained by Benjamini in [28], the corresponding qualitative picture is proved for the uniform spanning tree.

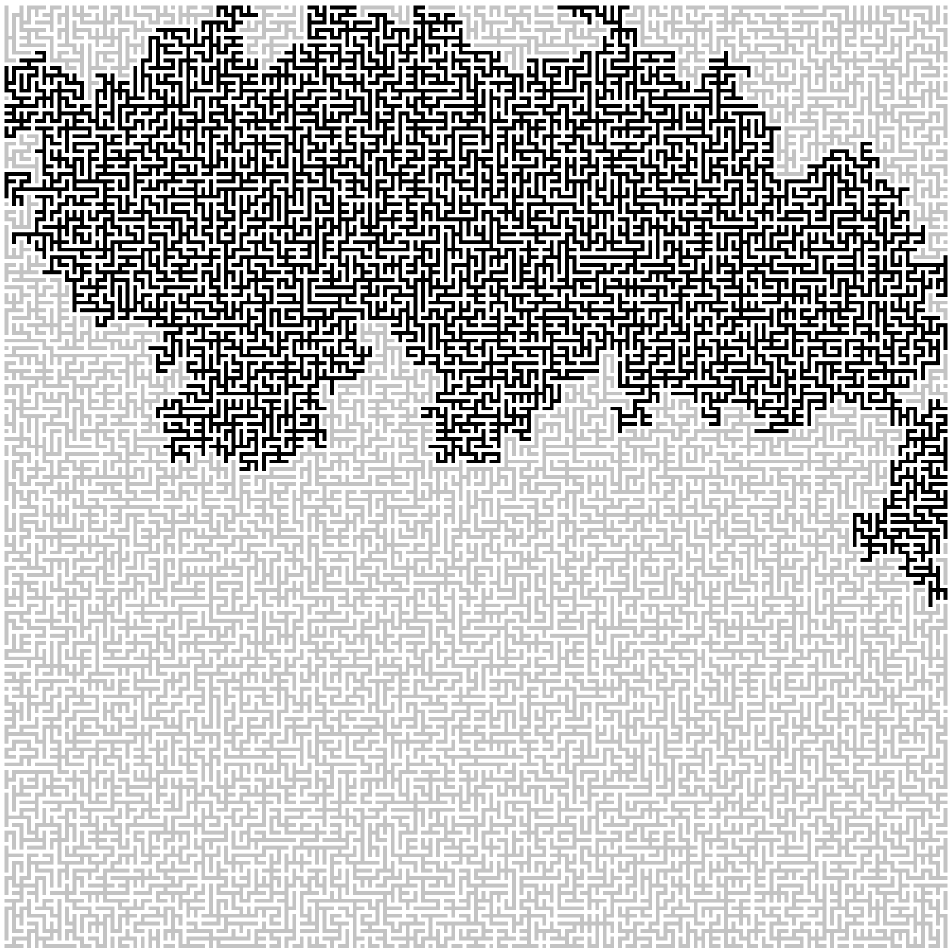

Figure 6.1: Part of a UST in a two-dimensional box; the part shown is the central 115×115 section of a UST on a 229×229 box. The single cluster spanning the two sides of the box is highlighted.

Remark. In [11], we establish a scaling limit for the three-dimensional UST in a version of the Gromov-Hausdorff topology, at least along the subsequence $\widehat{\delta}_n := 2^{-n}$. The corresponding two-dimensional result is also known (along an arbitrary sequence $\delta \to 0$), see [24] and [86, Remark 1.2]. In both cases, we expect that the techniques used to prove such a scaling limit can be used to show that the number of spanning clusters of $\mathcal{U}_\delta$ actually converges in distribution. We plan to pursue this in a subsequent work that focusses on the topological properties of the three-dimensional UST.

The organization of the remainder of the paper is as follows. In Section 2, we introduce some notation that will be used in the paper. The proof of Theorem 6.1.1 is then given in Section 3.

6.2 Notation

In this section, we introduce the main notation needed for the proof of Theorem 6.1.1. We write $|\cdot|$ for the Euclidean norm on $\mathbb{R}^3$ and, as in the introduction, $\mathrm{dist}(\cdot,\cdot)$ for the Euclidean distance between a point and a subset of $\mathbb{R}^3$. Given $\delta \in (0,1)$, if $x \in \delta\mathbb{Z}^3$ and $r > 0$, then we write

$$B_\delta(x,r) = \Big\{ y \in \delta\mathbb{Z}^3 : \ |x-y| < r \Big\}$$

for the lattice ball of centre x and radius r (we will commonly omit dependence on δ for brevity). Let $\mathbb{B}$, F and G be defined as at (6.1), (6.2) and (6.3) in the case $d = 3$.

For $\delta \in (0,1)$, a sequence $\lambda = (\lambda(0), \lambda(1), \cdots, \lambda(m))$ is said to be a path of length m if $\lambda(i) \in \delta\mathbb{Z}^3$ and $|\lambda(i) - \lambda(i+1)| = \delta$ for every i. A path λ is simple if $\lambda(i) \neq \lambda(j)$ for all $i \neq j$. For a path $\lambda = (\lambda(0), \lambda(1), \cdots, \lambda(m))$, we define its loop-erasure $\mathrm{LE}(\lambda)$ as follows. Firstly, let

$$s_0 = \max\left\{ j \le m : \ \lambda(j) = \lambda(0) \right\},$$

and for $i \ge 1$, set

$$s_i = \max\left\{ j \le m : \ \lambda(j) = \lambda(s_{i-1}+1) \right\}.$$

Moreover, write $n = \min\{ i : \ s_i = m \}$. The loop-erasure of λ is then given by

$$\mathrm{LE}(\lambda) = \left(\lambda(s_0), \lambda(s_1), \cdots, \lambda(s_n) \right).$$

We write $\mathrm{LE}(\lambda)(k) = \lambda(s_k)$ for each $0 \le k \le n$. Note that the vertices hit by $\mathrm{LE}(\lambda)$ are a subset of those hit by λ, and that $\mathrm{LE}(\lambda)$ is a simple path such that $\mathrm{LE}(\lambda)(0) = \lambda(0)$ and $\mathrm{LE}(\lambda)(n) = \lambda(m)$. Although the loop-erasure of λ has so far only been defined in the case that λ has a finite length, it is clear that we can define $\mathrm{LE}(\lambda)$ similarly for an infinite path λ if the set $\{k \geq 0 : \lambda(j) = \lambda(k)\}$ is finite for each $j \geq 0$. Additionally, when the path λ is given by a simple random walk, we call $\mathrm{LE}(\lambda)$ a loop-erased random walk (see [114] for an introduction to loop-erased random walks).

Again given $\delta \in (0,1)$, write $\mathcal{U}_\delta$ for the uniform spanning tree on $\delta\mathbb{Z}^3$. As noted in the introduction, this object was constructed in [143], and shown to be a tree with a single end, almost-surely. The graph $\mathcal{U}_\delta$ can be generated from loop-erased random walks by a procedure now referred to as Wilson's algorithm rooted at infinity (the name is after [169], while the version for infinite graphs was proved in [29]), which is described as follows.

- Let $(x_i)_{i\geq 1}$ be an arbitrary, but fixed, ordering of $\delta\mathbb{Z}^3$.
- Write R^{x_1} for a simple random walk on $\delta\mathbb{Z}^3$ started at x_1. Let $\gamma_{x_1} = \mathrm{LE}(R^{x_1})$ be the loop-erasure of R^{x_1} – this is well-defined since R^{x_1} is transient. Set $\mathcal{U}^1 = \gamma_{x_1}$. We refer to γ_{x_1} as a branch of $\mathcal{U}^1$.
- Given $\mathcal{U}^i$ for $i \geq 1$, let $R^{x_{i+1}}$ be a simple random walk (independent of $\mathcal{U}^i$) started at x_{i+1} and stopped on hitting $\mathcal{U}^i$. Then $\mathrm{LE}(R^{x_{i+1}})$ is a branch of the tree and we let $\mathcal{U}^{i+1} = \mathcal{U}^i \cup \mathrm{LE}(R^{x_{i+1}})$.

It is then the case that the output random tree $\cup_{i=1}^{\infty}\mathcal{U}^i$ has the same distribution as $\mathcal{U}_\delta$. In particular, the distribution of the output tree does not depend on the ordering of points $(x_i)_{i\geq 1}$.

Similarly to above, for $z \in \delta\mathbb{Z}^3$, we will write γ_z for the infinite simple path in $\mathcal{U}_\delta$ starting from z. Given a point $z \in \delta\mathbb{Z}^3$, it follows from the construction of $\mathcal{U}_\delta$ explained hitherto that the distribution of γ_z coincides with that of $\mathrm{LE}(R^z)$, where R^z is a simple random walk on $\delta\mathbb{Z}^3$ started at z.

Furthermore, as we explained in the introduction, we will write $\mathcal{U}'_\delta$ for the restriction of $\mathcal{U}_\delta$ to the cube $\mathbb{B}$. A connected component of $\mathcal{U}'_\delta$ is called a cluster. Also, as we defined previously, a spanning cluster is a cluster connecting F to G. We let N_δ be the number of spanning clusters of $\mathcal{U}_\delta$.

Finally, we will use c, C, c_0, etc. to denote universal positive constants which may change from line to line.

6.3 Proof of the main result

In this section, we will prove the following theorem, which incorporates Theorem 6.1.1.

Theorem 6.3.1. *There exists a universal constant C such that: for all $M < \infty$ and $\delta > 0$,*

$$\mathbf{P}\left(N_\delta \geq M\right) \leq CM^{-1}. \tag{6.4}$$

In particular, the laws of $(N_\delta)_{\delta\in(0,1)}$ form a tight sequence of probability measures on $\mathbb{Z}_+$.

Remark. In [3], Aizenman proved that for critical percolation in two dimensions, the probability of seeing M distinct spanning clusters is bounded above by Ce^{-cM^2}. We do not expect that the polynomial bound in (6.4) is sharp, but leave it as an open problem to determine the correct tail behaviour for number of spanning clusters of the UST in three dimensions, and, in particular, ascertain whether it also exhibits Gaussian decay.

Proof. Let $\delta \in (0,1)$, and suppose $M \geq 1$ is such that $\delta \leq M^{-1}$. For $r \in [0,1]$, we let

$$A(r) = \{(x_1, x_2, x_3) \in \mathbb{B} : \ x_1 = r\}.$$

We also define

$$A = [-1,2]^3, \quad \mathbb{B}' = \{(x_1, x_2, x_3) \in \mathbb{B} : \ x_1 \leq 2/3\}.$$

Moreover, let $(z_i)_{i=1}^L$ be a sequence of points in $A \cap \delta\mathbb{Z}^3$ such that $A \subseteq \cup_{i=1}^L B(z_i, 1/M)$ and $L \leq 10^5 M^3$.

To construct $\mathcal{U}_\delta$, we first perform Wilson's algorithm rooted at infinity for $(z_i)_{i=1}^L$ (see Section 6.2). Namely, we consider

$$\mathcal{U}^1 := \bigcup_{i=1}^L \gamma_{z_i},$$

which is the subtree of $\mathcal{U}_\delta$ spanned by $(z_i)_{i=1}^L$. (Recall that for $z \in \delta\mathbb{Z}^3$ we denote the infinite simple path in $\mathcal{U}_\delta$ starting from z by γ_z.) The idea of the proof is then as follows. Crucially, each branch of $\mathcal{U}^1$ is a 'hittable' set, in the sense that for a simple random walk R whose starting point is close to $\mathcal{U}^1$, it is likely that R hits $\mathcal{U}^1$ before moving far away. As a result, Wilson's algorithm guarantees that, with high probability, the spanning clusters of $\mathcal{U}_\delta$ correspond to those of $\mathcal{U}^1$ when M is sufficiently large. So, the problem boils down to the tightness of the number of spanning clusters of $\mathcal{U}^1$, which is not difficult to prove.

To make the above argument rigorous, we introduce the following two "good" events for $\mathcal{U}^1$:

$$H_i = H_i(\xi) := \left\{ \begin{matrix} \text{For any } x \in B(0,4) \cap \delta\mathbb{Z}^3 \text{ with } \operatorname{dist}(x, \gamma_{z_i}) \leq 1/M, \\ P_R^x\left(R[0,T] \cap \gamma_{z_i} = \emptyset\right) \leq M^{-\xi} \end{matrix} \right\},$$

$$I_i := \left\{ \begin{matrix} \text{The number of crossings of } \gamma_{z_i} \text{ between } A(0) \text{ and } A(2/3) \\ \text{in } \mathbb{B}' \text{ is smaller than } M \end{matrix} \right\},$$

for $1 \leq i \leq L$, where

- R is a simple random walk which is independent of γ_{z_i}, the law of which is denoted by P_R^x when we assume $R(0) = x$;
- T is the first time that R exits $B(x, 1/\sqrt{M})$;
- a crossing of γ_{z_i} between $A(0)$ and $A(2/3)$ in $\mathbb{B}'$ is a connected component of the restriction of γ_{z_i} to $\mathbb{B}'$ that connects $A(0)$ to $A(2/3)$.

Namely, the event H_i guarantees that the branch γ_{z_i} is a hittable set (see Figure 6.2), and the event I_i controls the number of crossings of γ_{z_i}.

Now, [148, Theorem 3.1] ensures that there exist universal constants $\xi_0, C > 0$ such that

$$\mathbf{P}\left(\bigcap_{i=1}^{L} H_i(\xi_0)\right) \geq 1 - CM^{-10}.$$

Thus, with high probability (for $\mathcal{U}^1$), each branch of $\mathcal{U}^1$ is a hittable set.

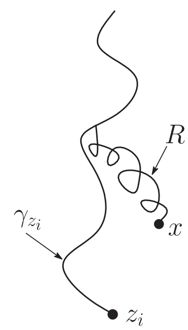

Figure 6.2: Conditional on the event H_i, for any $x \in B(0,4) \cap \delta\mathbb{Z}^3$ with $\mathrm{dist}(x, \gamma_{z_i}) \leq 1/M$, the above configuration occurs with probability at least $1 - M^{-\xi}$.

The probability of the event I_i is easy to estimate. Indeed, suppose that the event I_i does not occur. This implies that the number of "traversals" of S^{z_i} from $A(0)$ to $A(2/3)$ or vice versa must be bigger than M, where S^{z_i} stands for a simple random walk starting from z_i. Notice that there exists a universal constant $c_0 > 0$ such that for any point $w \in A(0)$ (respectively $w \in A(2/3)$), the probability that S^w hits $A(2/3)$ (respectively $A(0)$) is smaller than $1 - c_0$ (see [113, Proposition 1.5.10], for example). Thus, the probability of the event I_i is bounded below by $1 - (1 - c_0)^M =: 1 - e^{-aM}$, where $a > 0$. Taking sum over $1 \leq i \leq L$, we find that

$$\mathbf{P}\left(\bigcap_{i=1}^{L} I_i\right) \geq 1 - Le^{-aM}.$$

To put the above together, let

$$J = \bigcap_{i=1}^{L} H_i(\xi_0) \cap I_i.$$

For $1 \leq i \leq L$, set $\mathcal{U}_i^1 = \cup_{j=1}^{i} \gamma_{z_j}$ so that $\mathcal{U}^1 = \mathcal{U}_L^1$. As above, by a spanning cluster of $\mathcal{U}_i^1$ between $A(0)$ and $A(2/3)$ in $\mathbb{B}'$ we mean a connected component of the restriction of $\mathcal{U}_i^1$ to $\mathbb{B}'$ which connects $A(0)$ to $A(2/3)$. We write n_i for the number of spanning clusters of $\mathcal{U}_i^1$ between $A(0)$ and $A(2/3)$ in $\mathbb{B}'$.

On the event J, we have that

$$n_i \leq iM + i - 1,$$

for all $1 \leq i \leq L$, since $n_{i+1} - n_i$ is at most $M+1$ for each $i \geq 1$. In particular, we see that the number of spanning clusters of $\mathcal{U}^1$ between $A(0)$ and $A(2/3)$ in $\mathbb{B}'$ is bounded above by $L(M+1)$, which is comparable to M^4.

We next consider a sequence of subsets of A as follows. Let $a^* > 0$ be the positive constant such that

$$a^* \sum_{k=1}^{\infty} k^{-2} = 10^{-1}. \tag{6.5}$$

Set $\eta_1 = 0$, and $\eta_k = a^* \sum_{j=1}^{k-1} j^{-2}$ for $k \geq 2$. Finally, for $k \geq 1$, let

$$A_k = [-1 + \eta_k, 2 - \eta_k]^3.$$

Notice that $A_{k+1} \subseteq A_k$ and $[-1/2, 3/2]^3 \subseteq A_k$ for all $k \geq 1$, and moreover $\operatorname{dist}(\partial A_k, \partial A_{k+1}) = a^* k^{-2}$. We further introduce sequences $(z_i^k)_{i=1}^{L_k}$ consisting of points in $A_k \cap \delta\mathbb{Z}^3$ such that

$$A_k \subseteq \bigcup_{i=1}^{L_k} B\left(z_i^k, \delta_k\right),$$

and

$$L_k \leq 10^5 \delta_k^{-3}, \text{ where } \delta_k := M^{-1} 2^{-(k-1)}. \tag{6.6}$$

Note that we may assume that $L_1 = L$ and $(z_i^1)_{i=1}^{L_1} = (z_i)_{i=1}^{L}$.

For $\xi > 0$, we set

$$H_i^k = H_i^k(\xi) := \left\{ \begin{array}{c} \text{For any } x \in B(0,4) \cap \delta\mathbb{Z}^3 \text{ with } \operatorname{dist}\left(x, \gamma_{z_i^k}\right) \leq \delta_k, \\ P_R^x\left(R[0, T^k] \cap \gamma_{z_i^k} = \emptyset\right) \leq \delta_k^{\xi} \end{array} \right\},$$

where R is a simple random walk that is independent of $\gamma_{z_i^k}$, with law denoted by P_R^x when we assume $R(0) = x$, and T^k is the first time that R

exits $B(x, \sqrt{\delta_k})$. By [148, Theorem 3.1] again, there exist universal constants $\xi_1, C > 0$ (which do not depend on k) such that

$$\mathbf{P}\left(\bigcap_{i=1}^{L_k} H_i^k(\xi_1)\right) \geq 1 - C\delta_k^{10},$$

for all $k = 1, 2, \cdots, k_0$, where k_0 is the smallest integer k such that $\delta_k < \delta$. Thus if we write

$$H^k = \bigcap_{i=1}^{L_k} H_i^k(\xi_1)$$

and

$$J' = J \cap \bigcap_{k=1}^{k_0} H^k,$$

we have

$$\mathbf{P}(J') \geq 1 - CM^{-10}.$$

Given the above setup, we perform Wilson's algorithm rooted at infinity as follows:

- recall that $\mathcal{U}^1$ is the tree spanned by $(z_i^1)_{i=1}^{L_1} = (z_i)_{i=1}^{L}$;
- next perform Wilson's algorithm for $(z_i^2)_{i=1}^{L_2}$ – for each z_i^2, run a simple random walk $R^{z_i^2}$ from z_i^2 until it hits the part of the tree that has already been constructed, and adding its loop-erasure as a new branch – the output tree is denoted by $\mathcal{U}^2$;
- repeat the previous step for $(z_i^k)_{i=1}^{L_k}$ to construct $\mathcal{U}^k$ for $k = 1, 2, \cdots, k_0$.

Now, condition $\mathcal{U}^1$ on the event J above. We will show that, with high (conditional) probability, every new branch in $\mathcal{U}^2 \setminus \mathcal{U}^1$ has diameter smaller than $M^{-1/4}$. To this end, for $1 \leq i \leq L_2$, we write d_i^2 for the Euclidean diameter of the path from z_i^2 to $\mathcal{U}^1$ in $\mathcal{U}^2$, and define the event W_i^2 by setting

$$W_i^2 = \left\{ d_i^2 \geq M^{-1/4} \right\}.$$

Suppose that the event W_i^2 occurs. By Wilson's algorithm, the simple random walk $R^{z_i^2}$ must not hit the tree $\mathcal{U}^1$ until it exits $B(z_i^2, M^{-1/4})$.

Since $\mathrm{dist}(z_i^2, \partial A) \geq a^*$ (for the constant a^* defined at (6.5)), it holds that $B(z_i^2, M^{-1/4}) \subseteq A$. With this in mind, we set $u_0 = 0$, and

$$u_m = \inf \left\{ j \geq u_{m-1} : \left| R^{z_i^2}(j) - R^{z_i^2}(u_{m-1}) \right| \geq M^{-1/2} \right\}$$

for $m \geq 1$. We then have that

$$R^{z_i^2}[u_{m-1}, u_m] \cap \mathcal{U}^1 = \emptyset$$

for all $1 \leq m \leq M^{1/4}$. Since $A \subseteq \cup_{i=1}^L B(z_i, 1/M)$, it follows that for each $1 \leq m \leq M^{1/4}$, there exists a z_i such that $R^{z_i^2}(u_{m-1}) \in B(z_i, 1/M)$. Thus the event $H_i(\xi_0)$ guarantees that

$$\mathbf{P}\left(R^{z_i^2}[u_{m-1}, u_m] \cap \mathcal{U}^1 = \emptyset \text{ for all } 1 \leq m \leq M^{1/4} \right) \leq M^{-\xi_0 M^{1/4}}.$$

Consequently, the conditional probability of $\cup_{i=1}^{L_2} W_i^2$ is bounded above by $L_2 M^{-\xi_0 M^{1/4}}$, which is smaller than $C M^{-\xi_0 M^{1/4}+3}$ for some universal constant $C \in (0, \infty)$ (see (6.6), for the definition of L_2). Replacing constants if necessary, this implies that, with probability at least $1 - Ce^{-cM^{1/4}}$, every new branch in $\mathcal{U}^2 \setminus \mathcal{U}^1$ has diameter smaller than $M^{-1/4}$. Notice that once each new branch has such a small diameter, the event J guarantees that the number of spanning clusters of $\mathcal{U}^2$ between $A(0)$ and $A(2/3 + M^{-1/4})$ in $\mathbb{B}$ is bounded above by $L(M+1) \leq 10^6 M^4$.

Essentially the same argument is valid for $\mathcal{U}^k$. Indeed, conditioning $\mathcal{U}^k$ on the good event $J \cap \cap_{l=1}^k H^l$ as above, it holds that, with probability at least $1 - Ce^{-c\delta_k^{-1/4}}$ every new branch in $\mathcal{U}^{k+1} \setminus \mathcal{U}^k$ has diameter smaller than $\delta_k^{1/4}$. Notice that $\sum_k \delta_k^{1/4} \leq 10 M^{-1/4} < 10^{-2}$ when M is large. Therefore, with probability at least $1 - CM^{-10}$, the number of spanning clusters of $\mathcal{U}^{k_0}$ between $A(0)$ and $A(3/4)$ in $\mathbb{B}$ is bounded above by $L(M+1) \leq CM^4$ for some universal constant C.

Finally, we perform Wilson's algorithm for all of the remaining points in $\delta\mathbb{Z}^3$ to construct $\mathcal{U}_\delta$. Since k_0 is the smallest integer k such that $\delta_k < \delta$, it follows that the restriction of $\mathcal{U}_\delta$ to $\mathbb{B}$ coincides with that of $\mathcal{U}^{k_0}$. Note that $k_0 = 1$ when $\delta = M^{-1}$. Thus we conclude that there exists a universal constant C such that: for all $M < \infty$ and $\delta \in (0, M^{-1}]$,

$$\mathbf{P}\left(N_\delta \geq M\right) \leq CM^{-2}. \tag{6.7}$$

For the case that $\delta \in (M^{-1}, M^{-1/2}]$, we apply (6.7) to $M^{1/2}$ and monotonicity implies $\mathbf{P}\left(N_\delta \geq M\right) \leq \mathbf{P}\left(N_\delta \geq M^{1/2}\right) < CM^{-1}$. Finally, if $\delta > M^{-1/2}$, we use that $N_\delta \leq \delta^{-2} < M$. Combining these three bounds, we readily obtain the bound at (6.4). $\square$

Part III

Competitive Growth Processes

Chapter 7

Chase-Escape with Death on Trees[1]

Summary of this chapter

Chase-escape is a competitive growth process in which red particles spread to adjacent uncoloured sites, while blue particles overtake adjacent red particles. We introduce the variant in which red particles die and describe the phase diagram for the resulting process on infinite d-ary trees. A novel connection to weighted Catalan numbers makes it possible to characterize the critical behaviour.

7.1 Introduction

Chase-escape (CE) is a model for predator-prey interactions in which expansion of predators relies on but also hinders the spread of prey. The spreading dynamics come from the Richardson growth model [147]. Formally, the pro-

[1] Joint work with Erin Beckman, Keisha Cook, Nicole Eikmeier, and Matthew Junge. Acknowledgements. Thanks to David Sivakoff and Joshua Cruz for helpful advice and feedback. We are also grateful to Sam Francis Hopkins for pointing us to a reference about weighted Catalan numbers. Feedback during the review process greatly improved the final version. This work was partially supported by NSF DMS Grant #1641020 and was initiated during the 2019 AMS Mathematical Research Community in Stochastic Spatial Systems.

cess takes place on a graph in which vertices are in one of the three states $\{w, r, b\}$. Adjacent vertices in states (r, w) transition to (r, r) according to a Poisson process with rate λ. Adjacent (b, r) vertices transition to (b, b) at rate 1. The standard initial configuration has a single vertex in state r with a vertex in state b attached to it. All other sites are in state w. These dynamics can be thought of as prey r-particles "escaping" to empty w-sites while being "chased" and consumed by predator b-particles. We will refer to vertices in states r, b, and w as being red, blue, and white, respectively.

We introduce the possibility that prey dies for reasons other than predation in a variant which we call *chase-escape with death* (CED) . This is CE with four states $\{w, r, b, \dagger\}$ and the additional rule that vertices in state r transition to state $\dagger$ at rate $\rho > 0$. We call such vertices *dead*. Dead sites cannot be reoccupied.

7.1.1 Results

We study CED on the infinite rooted d-ary tree $\mathbb{T}_d$—the tree in which each vertex has $d \geq 2$ children—with an initial configuration that has the root red, one extra blue vertex $\mathfrak{b}$ attached to it, and the rest of the vertices white. Let $\mathcal{R}$ be the set of sites that are ever coloured red. Similarly, let $\mathcal{B}$ be the set of sites that are ever coloured blue. Denote the events that red and blue occupy infinitely many sites by $A = \{|\mathcal{R}| = \infty\}$ and $B = \{|\mathcal{B}| = \infty\}$. Since $\mathcal{B} - \{\mathfrak{b}\} \subseteq \mathcal{R}$ deterministically, we also have $B \subseteq A$. We will typically write P and E in place of $P_{\lambda,\rho}$ and $E_{\lambda,\rho}$ for probability and expectation when the rates are understood to be fixed. There are three possible phases for CED:

- *Coexistence* $P(B) > 0$.
- *Escape* $P(A) > 0$ and $P(B) = 0$.
- *Extinction* $P(A) = 0$.

For each fixed d and λ, we are interested in whether or not these phases occur and how the process transitions between them as ρ is varied. Accordingly, we define the critical values

$$\rho_c = \rho_c(d, \lambda) = \inf\{\rho \colon P_{\lambda,\rho}(B) = 0\}, \tag{7.1}$$

$$\rho_e = \rho_e(d, \lambda) = \inf\{\rho \colon P_{\lambda,\rho}(A) = 0\}. \tag{7.2}$$

One feature of CE, and likewise CED, that makes it difficult to study on graphs with cycles is that there is no known coupling that proves $P(B)$ increases with λ. On trees the coupling is clear, which makes analysis more tractable. It follows from [35] that $P(B) > 0$ in CE on a d-ary tree if and only if $\lambda > \lambda_c^-$ with

$$\lambda_c^- = 2d - 1 - 2\sqrt{d^2 - d} \sim (4d)^{-1}. \tag{7.3}$$

For CED on $\mathbb{T}_d$, $P(B)$ is no longer monotonic with λ. As λ increases, blue falls further behind and so the intermediate red particles must live longer for coexistence to occur. This lack of monotonicity makes

$$\lambda_c^+ = 2d - 1 + 2\sqrt{d^2 - d} \sim 4d$$

also relevant, because we will see that when $\lambda \geq \lambda_c^+$, the gap between red and blue is so large that coexistence is impossible for any $\rho > 0$.

Suppressing the dependence on d, let $\Lambda = (\lambda_c^-, \lambda_c^+)$. Unless stated otherwise, we will assume that $d \geq 2$ is fixed. Our first result describes the phase structure of CED (see Figure 7.1).

Theorem 7.1.1. *Fix* $\lambda > 0$*.*

(i) If $\lambda \in \Lambda$*, then* $0 < \rho_c < \rho_e = \lambda(d-1)$ *with escape occurring at* $\rho = \rho_c$*, and extinction at* $\rho = \rho_e$*.*

(ii) If $\lambda \notin \Lambda$*, then* $0 = \rho_c < \rho_e = \lambda(d-1)$ *with extinction occurring at* $\rho = \rho_e$*.*

Our next result concerns the behaviour of $E|\mathcal{B}|$ at and above criticality.

Theorem 7.1.2. *Fix* $\lambda \in \Lambda$*.*

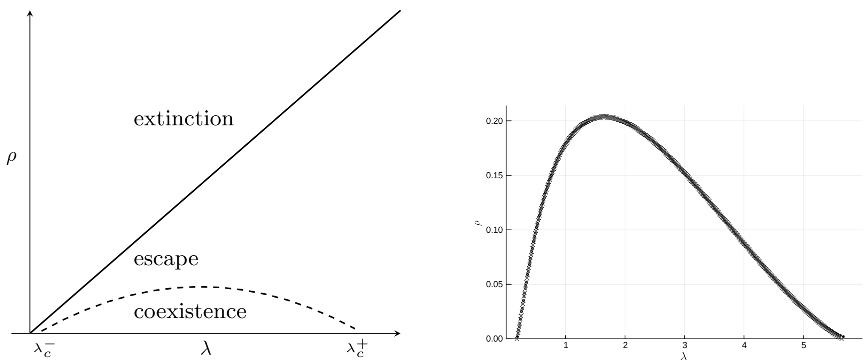

Figure 7.1: (Left) The phase diagram for fixed d. The dashed line is ρ_c and the solid line ρ_e. (Right) A rigorous approximation of ρ_c when $d = 2$. The approximations for larger d have a similar shape.

(i) If $\rho > \rho_c$, then $E|\mathcal{B}| < \infty$.

(ii) If $\rho = \rho_c$, then $E|\mathcal{B}| = \infty$.

Theorem 7.1.2 (ii) is particularly striking because it is known that $E|\mathcal{B}| < \infty$ in CE with $\lambda = \lambda_c^-$ (see [35, Theorem 1.4]). Hence the introduction of death changes the critical behaviour. The reason for this comes down to singularity analysis of a generating function associated to CED and is discussed in more detail in Remark 7.3.2.

We prove three further results about ρ_c concerning: the asymptotic growth in d, smoothness in λ, and the approximate value for a given d and λ (see Figure 7.1).

Theorem 7.1.3. *Fix $\lambda > 0$, $c < \sqrt{\lambda/2}$, and $C > \sqrt{2\lambda}$. For all d large enough*

$$c\sqrt{d} \leq \rho_c \leq C\sqrt{d}.$$

Theorem 7.1.4. *The function ρ_c is infinitely differentiable in $\lambda \in \Lambda$.*

Theorem 7.1.5. *Fix $\lambda \in \Lambda$. For every $\rho \neq \rho_c$, there is a finite runtime algorithm to determine if $\rho < \rho_c$ or if $\rho > \rho_c$.*

7.1.2 Proof methods

Theorems 7.1.1 and 7.1.3 are proven by relating coexistence in CED to the survival of an embedded branching process that renews each time blue is within distance one of the farthest red. Describing the branching process comes down to understanding how CED behaves on the nonnegative integers with 0 initially blue and 1 initially red. In particular, we are interested in the event $\mathfrak{R}_k$ that at some time k is blue, $k+1$ is red, and all further sites are white.

The probability of $\mathfrak{R}_k$ can be expressed as a weighted Catalan number. These are specified by non-negative weights $u(j)$ and $v(j)$ for $j \geq 0$. Given a lattice path γ consisting of unit rise and fall steps, each rise step from (x, j) to $(x+1, j+1)$ has weight $u(j)$, while fall steps from $(x, j+1)$ to $(x+1, j)$ has weight $v(j)$. The *weight* $\omega(\gamma)$ of a Dyck path γ is defined to be the product of the rise and fall step weights along γ.

The corresponding *weighted Catalan number* is $C_k^{u,v} = \sum \omega(\gamma)$ where the sum is over all Dyck paths γ of length $2k$ (nonnegative paths starting at $(0,0)$ consisting of k rise and k fall steps). See Figure 7.2 for an example. For CED, we define $C_k^{\lambda,\rho}$ as the weighted Catalan number with weights

$$u(j) = \frac{\lambda}{1+\lambda+(j+1)\rho} \text{ and } v(j) = \frac{1}{1+\lambda+(j+2)\rho}. \tag{7.4}$$

At (7.8) we explain why $P(\mathfrak{R}_k) = C_k^{\lambda,\rho}$.

Returning to CED on $\mathbb{T}_d$, self-similarity ensures that the expected number of renewals in the embedded branching process is equal to the generating function $g(z) = \sum_{k=0}^{\infty} C_k^{\lambda,\rho} z^k$ evaluated at $z = d$. We prove in Proposition 7.4.1 that ρ_c is the value at which the radius of convergence of g is equal to d. We characterize the radius of convergence using a continued fraction representation of g, which leads to the proofs of Theorems 7.1.2, 7.1.4, and 7.1.5.

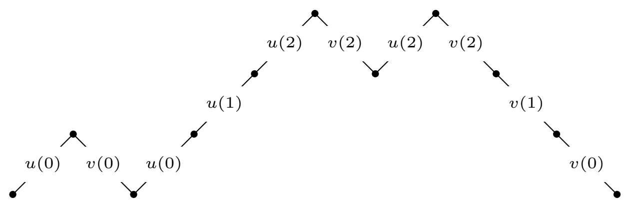

Figure 7.2: A Dyck path of length 10. The weight of this path is $u(0)^2v(0)^2u(1)v(1)u(2)^2v(2)^2$.

7.1.3 History and context

The forebearer of CE is the Richardson growth model for the spread of a single species [147]. In our notation, this process corresponds to the setting with $\lambda = 1$, $\rho = 0$, and no blue particles. Many basic questions remain unanswered for the Richardson growth model on the integer lattice [17], as well as for the competitive version [56].

James Martin conjectured that coexistence occurs in CE on lattices when red spreads at a slower rate than blue. Simulation evidence from Tang, Kordzakhia, Lalley in [164] suggested that, on the two-dimensional lattice, red and blue coexist with positive probability so long as $\lambda \geq 1/2$. Durrett, Junge, and Tang proved in [62] that red and blue can coexist with red stochastically slower than blue on high-dimensional oriented lattices with spreading times that resemble Bernoulli bond percolation.

The first rigorous result we know of for CE is Kordzakhia's proof that the phase transition occurs at λ_c^- for CE on regular trees [104]. Later, Kortchemski considered the process on the complete graph as well as trees with arbitrary branching number [105, 106]. An alternative perspective of CE as scotching a rumor was studied by Bordenave in [34]. The continuous limit of rumor scotching was studied many years earlier by Aldous and Krebs [8]. Looking to model malware spread and suppression through a device network, Hinsen, Jahnel, Cali, and Wary studied CE on Gilbert graphs [82].

To the best of our knowledge, CED has not been studied before. From the perspective of modelling species competition, it seems natural for prey to die from causes other than being consumed, and, in rumor scotching, for holders to cease spread because of fading interest. Furthermore, CED has new mathematical features. The existence of an escape phase, the fact that $E|\mathcal{B}| = \infty$ at criticality, and the lack of monotonicity of $P(B)$ in λ are all different than what occurs in CE.

The perspective we take on weighted Catalan numbers also appears to be novel. Typically they are studied with integer weights [9, 146, 152]. We are interested in the fractional weights at (7.4). Flajolet and Guilleman observed that fractionally weighted lattice paths describe birth and death chains for which the rates depend on the population size [68]. The distance between the rightmost red and rightmost blue for CED on the nonnegative integers is a birth and death chain in which mass extinction may occur. Since we are interested in CED on trees, we analyse the radius of convergence of the generating function of weighted Catalan numbers, which we believe has not been studied before.

7.2 CED on the line

Let CED^+ denote the process with CED dynamics on the nonnegative integers for which the vertices 0 and 1 are initially blue and red, respectively. All other vertices are initially white. Let $s_t(n) \in \{w, r, b, \dagger\}$ indicate the state of vertex n at time t. Define the processes $B_t = \max\{n \colon s_t(n) = b\}$, $R_t = \max\{n \colon s_t(n) = r\}$, and the random variable

$$Y = \sup\{B_t \colon t \geq 0\}. \tag{7.5}$$

If $s_t(n) \neq r$ for all n, then define $R_t = -\infty$. Let $\partial_t = (s_t(n))_{n=B_t}^{R_t}$ be the state of the interval $[B_t, R_t]$. One can think of this as the boundary of the process. Note that this interval only makes sense when $R_t > -\infty$. *Renewal times* are times $t \geq 0$ such that $\partial_t = (b, r)$. For $k \geq 0$ let

$$\mathfrak{R}_k = \{B_t = k \text{ for some renewal time } t\} \tag{7.6}$$

be the event that there is a renewal when blue occupies site k. Also define the event $A_t = \{s_t(n) \neq \dagger \text{ for all } n\}$ that none of the red sites have died up to time t.

Let $S_t = R_t - B_t$ be the distance between the rightmost blue and red particles at time t. Define the collection of *jump times* as $\tau(0) = 0$ and for $i \geq 1$

$$\tau(i) = \inf\{t \geq \tau(i-1)\colon S_t \neq S_{\tau(i-1)} \text{ or } \mathbb{1}(A_t) = 0\}.$$

The *jump chain* $J = (J_i)$ of S_t is given by

$$J_i = \begin{cases} S_{\tau(i)}, & \mathbb{1}(A_{\tau(i)}) = 1 \\ 0, & \text{otherwise} \end{cases}.$$

This is a Markov chain with an absorbing state 0 corresponding to blue no longer having the potential to reach infinity. The transition probabilities for $j > 0$ are:

$$p_{j,j+1} = \frac{\lambda}{1+\lambda+j\rho}, \qquad p_{j,j-1} = \frac{1}{1+\lambda+j\rho}, \qquad p_{j,0} = \frac{j\rho}{1+\lambda+j\rho}. \quad (7.7)$$

Call a jump chain $(J_0, \dots, J_n)$ *living* if $J_i > 0$ for all $0 \leq i \leq n$. Translating the set of Dyck paths of length $2k$ up by one vertical unit gives the jump chains corresponding to $\mathfrak{R}_k$. Notice that $p_{j,j+1} = u(j-1)$ and $p_{j,j-1} = v(j-2)$ with u and v as in (7.4). Thus, it is easy to see that for all $k \geq 0$ we have

$$P_{\lambda,\rho}(\mathfrak{R}_k) = C_k^{\lambda,\rho}, \qquad (7.8)$$

with $C_k^{\lambda,\rho}$ the weighted Catalan number defined in Section 7.1.2.

Lemma 7.2.1. *For any $\epsilon > 0$ there exists $\rho' > 0$ such that for all $\rho \in [0, \rho')$ and sufficiently large k*

$$C_k^{\lambda,\rho} \geq \left(\frac{(4-\epsilon)\lambda}{(1+\lambda)^2} \right)^k.$$

Proof. Let $C_{k,m}$ be number of Dyck paths of length $2k$ that never exceed height m. We first claim that for any $\delta > 0$, there exists m_δ such that

$$C_{k,m_\delta} \geq (4-\delta)^k \tag{7.9}$$

for sufficiently large k.

The mth Catalan number $C_m := C_{m,\infty}$ counts the number of Dyck paths of length $2m$. Consider any sequence of $\lfloor k/m \rfloor$ Dyck paths of length $2m$. If we concatenate these paths, we have a path of length $2m\lfloor k/m \rfloor \geq 2k - 2m$ which stays below height m. We extend this path into a Dyck path of length $2k$ by concatenating the necessary number of up and down steps to the end in any manner. Since each of the $\lfloor k/m \rfloor$ Dyck paths of length $2m$ can be chosen independently of one another, we have $C_{k,m} \geq (C_m)^{\lfloor k/m \rfloor}$.

Using the standard asymptotic relation $C_m \sim (1/\sqrt{\pi}) m^{-3/2} 4^m$ (see [161]), we have for large enough m, k

$$C_{k,m} \geq \left(\frac{4^m}{2\sqrt{\pi} m^{3/2}} \right)^{\lfloor k/m \rfloor} \geq \left(\frac{4^m}{2\sqrt{\pi} m^{3/2}} \right)^{-1} \left(\frac{4}{(2\sqrt{\pi} m^{3/2})^{1/m}} \right)^k .$$

It is easy to verify that $(2\sqrt{\pi} m^{3/2})^{1/m} \to 1$ as $m \to \infty$. Thus, we can choose m large enough so that

$$\frac{4}{(2\sqrt{\pi} m^{3/2})^{1/m}} > 4 - \frac{\delta}{2}.$$

We then have

$$C_{k,m} \geq C \left(\frac{4^m}{2\sqrt{\pi} m^{3/2}} \right)^{-1} (4 - (\delta/2))^k.$$

This is true for all m, k sufficiently large, and we can see that if we fix an m_δ big enough so that this inequality holds, then we can increase k enough such that we have the claimed inequality at (7.9).

Using the weights at (7.4), each path γ counted by $C_{k,m}$ satisfies

$$\omega(\gamma) \geq \left(\frac{\lambda}{(1 + \lambda + m\rho)^2} \right)^k ,$$

because γ has length $2k$ but never exceeds height m. Summing over just the

Dyck paths counted by $C_{k,m}$ gives

$$C_k^{\lambda,\rho} \geq C_{k,m} \frac{\lambda^k}{(1+\lambda+m\rho)^{2k}}. \tag{7.10}$$

Fix an $\epsilon > 0$. We can choose $\delta > 0$ small enough so that $(4-\delta)(1-\delta) > 4-\epsilon$. By (7.9), pick m_δ large enough to have $C_{k,m_\delta} > (4-\delta)^k$ for all sufficiently large k. Finally, choose $\rho' > 0$ small enough so that

$$\frac{\lambda}{(1+\lambda+m_\delta\rho')^2} > \frac{\lambda}{(1+\lambda)^2}(1-\delta).$$

Since the $C_k^{\lambda,\rho}$ are decreasing in ρ, applying these choices to (7.10) gives the desired inequality for all $\rho \in [0,\rho')$:

$$C_k^{\lambda,\rho} \geq \left((4-\delta)(1-\delta)\frac{\lambda}{(1+\lambda)^2}\right)^k \geq \left(\frac{(4-\epsilon)\lambda}{(1+\lambda)^2}\right)^k.$$

□

Recall the definition of Y at (7.5). We conclude this section by proving that $P(Y \geq k)$ can be bounded in terms of $P(\mathfrak{R}_k)$. The difficulty is that the event $\{Y \geq k\}$ includes *all* realizations for which blue reaches k, while $\mathfrak{R}_k$ only includes realizations which have a renewal at k.

Lemma 7.2.2. *For $\rho > 0$, there exists $C > 0$, which is a function of λ, ρ, such that $P(Y \geq k) \leq Ck^{1+\lambda/\rho}P(\mathfrak{R}_k)$ for all $k \geq 1$.*

Proof. Given a living jump chain $J = (J_0, J_1, \ldots, J_m)$, define the *height profile* of J to be $h(J) = (h_1(J), \ldots, h_{m+1}(J))$, where $h_i(J)$ are the number of entries J_ℓ in J with $\ell < m$ for which $J_\ell = i$. These values correspond to the total number of times that blue is at distance i from red. Suppose that red takes $r(J)$ many steps in a jump chain J. It is straightforward to show that

$$p(J) = \lambda^{r(J)} \prod_{j=1}^{m+1} \left(\frac{1}{1+\lambda+j\rho}\right)^{h_j(J)} \tag{7.11}$$

is the probability that the process follows the living jump chain J.

We view realizations leading to outcomes in $\{Y \geq k\}$ as having two distinct stages. In the *first stage*, the rightmost red particle reaches k. In the *second stage*, we ignore red and only require that blue advances until it reaches k. The advantage of this perspective is that we can partition outcomes in $\{Y \geq k\}$ by their behaviour in the first stage and then restrict our focus to the behaviour of the process in the interval $[0, k]$ for the second stage.

For $2 \leq \ell \leq k$, define Γ_ℓ to be the set of all living jump chains of length $2k - \ell - 1$ which go from $(0, 1)$ to $(2k - \ell - 1, \ell)$ with the last step being a rise step (see Figure 7.3). These are the jump chains from the first stage. Now we describe the second stage. Assume that blue is at $k - \ell$ when red reaches k. For blue to reach k, the red sites in $[k - \ell + 1, k]$ must stay alive long enough for blue to advance another ℓ steps. This has probability $\sigma(\ell) := \prod_{i=1}^{\ell}(1 + i\rho)^{-1}$. Given $\gamma \in \Gamma_\ell$, the formula at (7.11) implies that the probability $Y \geq k$ and the first $2k - \ell - 1$ steps of the process jump chain follow γ is

$$q(\gamma) := p(\gamma)\sigma(\ell) = \lambda^{k-1}\sigma(\ell)\prod_{j=1}^{k+1}(1 + \lambda + j\rho)^{-h_j(\gamma)}. \tag{7.12}$$

Let $q_\ell = \sum_{\gamma \in \Gamma_\ell} q(\gamma)$. Notice that

$$q_2 \frac{\lambda}{(1 + \lambda + \rho)(1 + \lambda + 2\rho)^2} \leq P(\mathfrak{R}_k). \tag{7.13}$$

This is because a subset of $\mathfrak{R}_k$ is the collection of processes which follow jump chains in Γ_2 and for which the next three steps have blue advance by one, then red advance by one, followed by blue advancing one. We will further prove that there exists C_0 (independent of ℓ) such that

$$q_\ell \leq C_0 \ell^{\lambda/\rho} q_2. \tag{7.14}$$

The claimed inequality then follows from (7.13), (7.14), and the partitioning $P(Y \geq k) = \sum_{\ell=2}^{k} q_\ell$.

To prove (7.14), notice that by inserting $\ell - 2$ fall steps right before the final upward step of each path in $\gamma \in \Gamma_\ell$, we obtain a path $\tilde{\gamma} \in \Gamma_2$ (see Figure 7.3). Since the paths γ and $\tilde{\gamma}$ agree for the first $2k - \ell - 2$ steps, we have from (7.12)

$$q(\gamma) = \frac{\sigma(\ell)}{\sigma(2)\prod_{i=1}^{\ell-2}(1+\lambda+i\rho)^{-1}} q(\tilde{\gamma}) \tag{7.15}$$

$$= \frac{(1+\lambda+\rho)(1+\lambda+2\rho)}{(1+\ell\rho)(1+(\ell-1)\rho)} q(\tilde{\gamma}) \prod_{i=3}^{\ell-2}\left(1+\frac{\lambda}{1+i\rho}\right). \tag{7.16}$$

Rewriting as a sum and using integral bounds, one can verify that

$$\prod_{i=3}^{\ell-2}\left(1+\frac{\lambda}{1+i\rho}\right) \leq C_1 \ell^{\lambda/\rho},$$

for some C_1 that depends on λ and ρ. This gives $q(\gamma) \leq C_0 \ell^{\lambda/\rho} q(\tilde{\gamma})$. Thus, $q_\ell \leq \sum_{\gamma \in \Gamma_\ell} C_0 \ell^{\lambda/\rho} q(\tilde{\gamma})$. If we restrict to paths $\gamma \in \Gamma_\ell$, the map $\gamma \mapsto \tilde{\gamma}$ is injective and hence

$$q_\ell \leq \sum_{\gamma \in \Gamma_\ell} C_0 \ell^{\lambda/\rho} q(\tilde{\gamma}) \leq C_0 \ell^{\lambda/\rho} \sum_{\gamma \in \Gamma_2} q(\gamma) = C_0 \ell^{\lambda/\rho} q_2.$$

This yields (7.14) and completes the lemma. □

7.3 Properties of weighted Catalan numbers

7.3.1 Preliminaries

Given a sequence $(c_n)_{n \geq 0}$, define the formal continued fraction

$$K[c_0, c_1, \ldots] := \cfrac{c_0}{1 - \cfrac{c_1}{1 - \ddots}}. \tag{7.17}$$

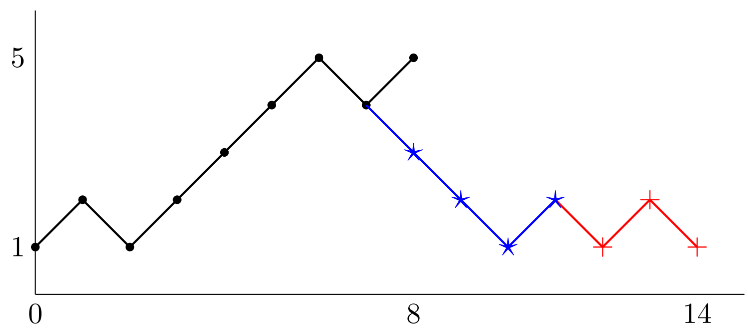

Figure 7.3: Let $k = 7$. The black line with dots is a path $\gamma \in \Gamma_5$. The blue line with stars is the modified path $\tilde{\gamma} \in \Gamma_2$. The red line with pluses is the extension of $\tilde{\gamma}$ to a jump chain in $\mathfrak{R}_7$.

The c_i may be fixed numbers, or possibly functions. Also, whenever we write x we mean a nonnegative real number, and z represents an arbitrary complex number.

The discussion in [73, Chapter 5] tells us that, for general weighted Catalan numbers, $g(z) := \sum_{k=0}^{\infty} C_k^{u,v} z^k$ is equal to the function

$$f(z) := K[1, a_0 z, a_1 z, \ldots] \tag{7.18}$$

for all $|z| < M$, where M is the radius of convergence of g centred at the origin, and $a_j = u(j)v(j)$. M is the modulus of the nearest singularity of g to the origin, or by the Hadamard-Cauchy theorem

$$M = \frac{1}{\limsup_{k\to\infty} (C_k^{u,v})^{1/k}}. \tag{7.19}$$

Let $u(j)$ and $v(j)$ be as in (7.4) so that, unless stated otherwise, we have

$$a_j = \frac{\lambda}{(1 + \lambda + (j+1)\rho)(1 + \lambda + (j+2)\rho)}. \tag{7.20}$$

When necessary, we denote the dependence on λ and ρ by $f_{\lambda,\rho}$.

7.3.2 Properties of f and M

Our proof that f is meromorphic relies on a classical convergence criteria for continued fractions [165, Theorem 10.1].

Theorem 7.3.1 (Worpitzky Circle Theorem)**.** *Let $c_j \colon D \to \{|w| < 1/4\}$ be a family of analytic functions over a domain $D \subseteq \mathbb{C}$. Then*

$$K[1, c_0(z), c_1(z), \ldots]$$

converges uniformly for z in any compact subset of D and the limit function takes values in $\{|z - 4/3| \leq 2/3\}$.

Corollary 7.3.1. If $\rho > 0$, then f is a meromorphic function on $\mathbb{C}$.

Proof. We will prove that f is meromorphic for all $z \in \Delta = \{|z| < r_0\}$ with $r_0 > 0$ arbitrary. Let $h_j(z) := K[a_j z, a_{j+1} z, \ldots]$ be the tail of the continued fraction so that $f(z) = K[1, a_0 z, \ldots, a_{j-1} z, h_j(z)]$. Since $\rho > 0$ we have $|a_j| \downarrow 0$ as $j \to \infty$. It follows that for some $j = j(r_0)$ large enough, $|a_k z| \leq 1/4$ for all $k \geq j$ and $z \in \Delta$. Theorem 7.3.1 ensures that $|h_j(z)| < \infty$ and the partial continued fractions $K[a_j z, \ldots, a_n z]$ are analytic (again by Theorem 7.3.1) and converge uniformly to h_j for $z \in \Delta$. Thus, h_j is a uniform limit of analytic functions and is therefore analytic on Δ. We can then write $f(z) = K[1, a_0 z, \ldots, a_{j-1} z, h_j(z)]$. Since each $a_i z$ is a linear function in z, f is a quotient of two analytic functions. $\square$

Next we show that the exponential growth rate of the $C_k^{\lambda,\rho}$ responds continuously to changes in ρ.

Lemma 7.3.2. *Fix $\lambda > 0$.*

(i) Fix $\rho > 0$. Given $0 \leq \rho' < \rho$, there exists $\epsilon > 0$ such that

$$C_k^{\lambda,\rho}(1+\epsilon)^k \leq C_k^{\lambda,\rho'}. \tag{7.21}$$

(ii) Fix $\rho > 0$. Given $\epsilon' > 0$, there exists a $\delta > 0$ such that for $\rho' \in (\rho - \delta, \rho + \delta)$

$$(1-\epsilon')^k \le \frac{C_k^{\lambda,\rho}}{C_k^{\lambda,\rho'}} \le (1+\epsilon')^k. \tag{7.22}$$

(iii) If $\rho = 0$, then given $\epsilon' > 0$, there exists a $\delta > 0$ such that for $\rho' \in [0, \delta)$ it holds that

$$1 \le \frac{C_k^{\lambda,\rho}}{C_k^{\lambda,\rho'}} \le (1+\epsilon')^k. \tag{7.23}$$

Proof. Let γ be a Dyck path of length $2k$. From the definition of $C_k^{\lambda,\rho}$, we have $w(\gamma)$ is a product of some combination of exactly k of the a_j terms. Let a'_j be the weights corresponding to ρ'. It is a basic calculus exercise to show that, when $\rho' < \rho$ we have the ratio $a'_j/a_j > 1$ is an increasing function in j. Let $a'_0/a_0 = 1+\epsilon$. Using the fact that $w(\gamma)$ and $w'(\gamma)$ have the same number of each weight, we can directly compare their ratio using the worst case lower bound

$$(1+\epsilon)^k = (a'_0/a_0)^k \le \frac{w'(\gamma)}{w(\gamma)}. \tag{7.24}$$

Cross-multiplying then summing over all paths γ gives (i).

Towards (ii), suppose $0 < \rho' < \rho$. Then we have $1 \le a'_j/a_j \le (\rho/\rho')^2$. Choose δ_1 such that $(\rho - \delta_1)/\rho = \sqrt{1-\epsilon'}$. Using the same logic as above, we have for $\rho' > \rho - \delta_1$ that

$$(1-\epsilon')^k \le \left(\frac{\rho'}{\rho}\right)^{2k} \le \frac{w(\gamma)}{w'(\gamma)} \le 1^k \le (1+\epsilon)^k.$$

Now suppose $0 < \rho < \rho'$. Choose δ_2 such that $(\rho+\delta_2)/\rho = \sqrt{1+\epsilon'}$. Then, following the same steps as above, we see that for $\rho' < \rho + \delta_2$

$$(1-\epsilon')^k \leq (1)^k \leq \frac{w(\gamma)}{w'(\gamma)} \leq \left(\frac{\rho'}{\rho}\right)^{2k} \leq (1+\epsilon')^k.$$

The same reasoning as with (7.24) and choosing $\delta = \min(\delta_1, \delta_2)$ gives (ii).

We now prove (iii). Notice that $C_k^{\lambda,0} = C_k(\lambda/(1+\lambda)^2)^k$ with C_k the kth Catalan number, since the transition probabilities at (7.7) for the associated jump chain are homogeneous when $\rho = 0$. By Lemma 7.2.1, we have for a fixed $\tilde{\epsilon} > 0$ and the bound $C_k \leq 4^k$, there exists a $\rho' > 0$ such that

$$1 \leq \frac{C_k^{\lambda,0}}{C_k^{\lambda,\rho'}} \leq \frac{C_k(\lambda/(1+\lambda)^2)^k}{(4-\tilde{\epsilon})^k\,(\lambda/(1+\lambda)^2)^k} = \frac{C_k}{(4-\tilde{\epsilon})^k} \leq \left(\frac{4}{4-\tilde{\epsilon}}\right)^k. \tag{7.25}$$

Choosing $\tilde{\epsilon}$ such that $4/(4-\tilde{\epsilon}) = 1+\epsilon'$ and letting δ be small enough so that (7.25) holds for all $\rho' < \delta$, we obtain an identical bound to the one in Lemma 7.3.2 (ii) but for $\rho' \in [0, \delta)$.

□

Lemma 7.3.3. *M is a continuous strictly increasing function for $\rho \in [0,\infty)$ satisfying $M(0) = (1+\lambda)^2/(4\lambda)$.*

Proof. First note that for any ρ, $M(\rho) \leq (a_0)^{-1} < \infty$. This is because $C_k^{\lambda,\rho} \geq \omega(\gamma_k) = (a_0)^k$ where γ_k is the Dyck path consisting of k alternating rise and fall steps.

That M is increasing follows immediately from the fact that the $C_k^{\lambda,\rho}$ are decreasing in ρ. To see that M is strictly increasing we use Lemma 7.3.2 (i) in combination with the definition of M at (7.19). Indeed, the lemma implies that for any $0 \leq \rho' < \rho$ there exists a $\delta > 0$ such that

$$M(\rho') = \frac{1}{\limsup_{k\to\infty}(C_k^{\lambda,\rho'})^{1/k}} \leq \frac{1}{(1+\delta)\limsup_{k\to\infty}(C_k^{\lambda,\rho})^{1/k}} = \frac{M(\rho)}{1+\delta}.$$

So M is strictly increasing.

To show continuity for $\rho > 0$, we use (7.19). Fix $\rho > 0$ and $\epsilon > 0$. Let $\epsilon' = \epsilon/M(\rho)$ and let $\delta > 0$ be as in Lemma 7.3.2 (ii). For $\rho' \in (\rho-\delta, \rho+\delta)$ we have

$$\liminf_{k\to\infty}\left(\frac{C_k^{\lambda,\rho}}{C_k^{\lambda,\rho'}}\right)^{1/k} \le \frac{M(\rho')}{M(\rho)} \le \limsup_{k\to\infty}\left(\frac{C_k^{\lambda,\rho}}{C_k^{\lambda,\rho'}}\right)^{1/k}.$$

Using the bound in Lemma 7.3.2 (ii) results in

$$1-\epsilon' \le \frac{M(\rho')}{M(\rho)} \le 1+\epsilon'.$$

Simplifying, then replacing out ϵ' gives $|M(\rho')-M(\rho)|<\epsilon$. Thus, M is continuous at $\rho>0$. Continuity for $\rho=0$ uses a similar argument along with Lemma 7.3.2 (iii). The explicit formula for $M(0)$ comes from the formula $C_k^{\lambda,\rho}=C_k(\lambda/(1+\lambda)^2)^k$ and the fact that the generating function for the Catalan numbers has radius of convergence 1/4. $\square$

Lemma 7.3.4. *If $\lambda\in\Lambda$ then there exists a unique value $\rho_d>0$ such that $M(\rho_d)=d$; Moreover, if $\lambda\notin\Lambda$, then $M>d$ for all $\rho\ge 0$.*

Proof. Fix λ. To signify the dependence on ρ, let $g_\rho(z)$ be the generating function of the $C_k^{\lambda,\rho}$. Using the continuity and strict monotonicity of M in Lemma 7.3.3, to show the first statement, it suffices to prove that $g_\rho(d)<\infty$ for ρ large enough, and that $g_0(d)=\infty$. It is easy to see that if $\rho>\lambda(d-1)$ then the branching process of red spreading with no blue particles has finite expected size, and thus $g_\rho(d)<\infty$ for such ρ.

Using the formula for $M(0)$ from Lemma 7.3.3 and rearranging, we can see that when $\rho=0$, $M<d$ whenever $\lambda d/(1+\lambda^2)>1/4$. The set of λ for which this occurs is by definition Λ. Therefore, $g_0(d)=\infty$, proving the first claim. The claim that $M>d$ for $\lambda\notin\Lambda$, follows from Lemma 7.3.3. The explicit formula for $M(0)$ is easily shown to satisfy $M(0)>d$, and since M is increasing, this inequality holds for all $\rho\ge 0$. $\square$

Our next lemma requires a old theorem from complex variable theory (see [69, Theorem IV.6] for example).

Theorem 7.3.5 (Pringsheim's Theorem)**.** *If $f(z)$ is representable at the origin by a series expansion that has non-negative coefficients and radius of convergence M, then the point $z=M$ is a singularity of $f(z)$.*

Lemma 7.3.6. *Let $\rho > 0$. Then, $M \leq d$ if and only if $g(d) = \infty$.*

Proof. We first note that the implication "$M < d$ implies $g(d) = \infty$" as well as the reverse direction "$g(d) = \infty$ implies $M \leq d$" both follow immediately from the definition of the radius of convergence. It remains to show that $M = d$ implies $g(d) = \infty$. Corollary 7.3.1 proves that f is a meromorphic function. Since $g = f$ for $|z| < M$, $f(x) > 0$ for $x \in (0, d)$, and Theorem 7.3.5 gives $z = d$ is a pole, we have for $x \in \mathbb{R}$ that

$$g(d) = \lim_{x \uparrow d} g(x) = \lim_{x \uparrow d} f(x) = \infty.$$

$\square$

Remark. *CE and CED differ in their behaviour at $\lambda = \lambda_c^-$ (see Theorem 7.1.2). In particular, the argument that $g(d) = \infty$ when $M = d$ does not apply to CE. Why is this the case? The difference is that when $\rho = 0$, we cannot use Corollary 7.3.1. Instead, we can write a closed expression for the function:*

$$g(z) = \sum_{k=0}^{\infty} C_k^{\lambda,0} = \sum_{k=0}^{\infty} C_k \left(\frac{\lambda}{(1+\lambda)^2} \right)^k z^k = \frac{1 - \sqrt{1 - 4z\lambda/(1+\lambda)^2}}{2z},$$

we find that g has a branch cut rather than isolated poles. [69, Theorem IV.10] ensures that the growth of the C_k^{λ,ρ_c} is determined by the orders of the singularities of f at distance $M = d$. Formally,

$$C_k^{\lambda,\rho_c} = \sum_{|\alpha_j| = d} \alpha_j^{-k} \pi_j(k) + o(d^{-k}) \tag{7.26}$$

where the α_j are the poles of f at distance d and the π_j are polynomials with degree equal to the order of the pole of f at α_j minus one. When evaluating at the radius of convergence $M = \lambda_c^-$, this gives a summable pre-factor of $k^{-3/2}$ in (7.26).

7.4 M and CED

The main results from this section are that for $\lambda \in \Lambda$ we have $\rho_c = \rho_d$ and that $P(B)$ is continuous at ρ_d. The lemmas in this section will also be useful for characterizing the phase behaviour of CED.

Proposition 7.4.1. *For $\lambda \in \Lambda$ it holds that $\rho_d = \rho_c$.*

Proof. Lemmas 7.4.5 and 7.4.6 give that whenever $\lambda \in \Lambda$ we have $\rho_d = \inf\{\rho \colon P_{\lambda,\rho}(B) = 0\}$, which is the definition of ρ_c. $\square$

Lemma 7.4.2. *$M > d$ if and only if $E|\mathcal{B}| < \infty$.*

Proof. Letting Y be as at (7.5), we first observe that

$$E|\mathcal{B}| = \sum_{k=0}^{\infty} P(Y = k)d^k \geq \sum_{k=1}^{\infty} P(\mathfrak{R}_k)d^k = g(d).$$

Hence $E|\mathcal{B}| < \infty$ implies $g(d) < \infty$, which gives $M > d$ by Lemma 7.3.6.

For the other direction, using the comparison in Lemma 7.2.2 gives

$$E|\mathcal{B}| = \sum_{k=0}^{\infty} P(Y = k)d^k \leq 1 + \sum_{k=1}^{\infty} Ck^{1+\lambda/\rho} P(\mathfrak{R}_k)d^k. \tag{7.27}$$

Lemma 7.3.3 tells us that M is continuous and since $M > d$, the sum on the right still converges with the polynomial prefactor. $\square$

Call a vertex $v \in \mathbb{T}_d$ a *tree renewal vertex* if at some time it is red, the parent vertex is blue, and all vertices one level or more down from v are white. Note that this definition of renewal is a translation one unit left of the definition of renewals in the CED$^+$. This makes it so each vertex at distance k from the root is a tree renewal vertex with probability $P(\mathfrak{R}_k)$. We call a tree renewal vertex v a *first tree renewal vertex* if none of the non-root vertices in the shortest path from v to the root are renewal vertices. Let Z be the number of first tree renewal vertices in CED.

Lemma 7.4.3. *$EZ \geq 1$ if and only if $g(d) = \infty$.*

Proof. Using self-similarity of $\mathbb{T}_d$, the number of renewal vertices is a Galton-Watson process with offspring distribution Z. Since each of the d^k vertices at distance k have probability $P(\mathfrak{R}_k)$ of being a tree renewal vertex, the expected size of the Galton-Watson process is $g(d)$. Standard facts about branching processes imply the claimed equivalence. $\square$

Lemma 7.4.4. *If $\lambda \in \Lambda$ and $\rho = \rho_d$, then $EZ \leq 1$.*

Proof. Since M is strictly increasing in ρ (Lemma 7.3.3) we have $M > d$ for $\rho > \rho_d$ and, so, Lemma 7.4.3 implies that $EZ < 1$. Let $\mathfrak{R}'_k$ be the event that a fixed, but arbitrary vertex at distance k is a first tree renewal vertex. Fatou's lemma gives

$$E_{\lambda,\rho_d} Z = \sum_{k=1}^{\infty} d^k P_{\lambda,\rho_d}(\mathfrak{R}'_k) = \sum_{k=1}^{\infty} d^k \liminf_{\rho \downarrow \rho_d} \mathbb{P}(\mathfrak{R}'_k) \tag{7.28}$$

$$\leq \liminf_{\rho \downarrow \rho_d} \sum_{k=1}^{\infty} d^k \mathbb{P}(\mathfrak{R}'_k) = \liminf_{\rho \downarrow \rho_d} E_{\lambda,\rho} Z \leq 1. \tag{7.29}$$

To see the second equality, notice that $\mathbb{P}(\mathfrak{R}'_k)$ is continuous in ρ at ρ_d, as $\mathfrak{R}'_k$ consists of finitely many jump chains.

$\square$

Lemma 7.4.5. *If $\lambda \in \Lambda$ and $\rho < \rho_d$, then $P(B) > 0$.*

Proof. It follows from Lemmas 7.3.3 and 7.4.3, along with the easily observed fact that EZ is strictly decreasing in ρ that $EZ > 1$ for $\rho < \rho_d$. Thus, the embedded branching process of renewal vertices is infinite with positive probability. As $|\mathcal{B}|$ is at least as large as the number of renewal vertices, this gives $P(B) > 0$. $\square$

Lemma 7.4.6. *If $\lambda \in \Lambda$ and $\rho = \rho_d$, then $P(B) = 0$.*

Proof. The probability blue survives along a path while remaining at least distance $L > 0$ from the most distant red site from the root for k consecutive steps is bounded by the probability none of the intermediate red particles die, a quantity which is in turn bounded by:

$$\left(\frac{1+\lambda}{1+\lambda+L\rho}\right)^k. \tag{7.30}$$

Choose L_0 large enough so that (7.30) is smaller than $(2d)^{-k}$. Let $D_k(v)$ be the event that, in the subtree rooted at v, the front blue particle moves from vertex v to a vertex at level k without ever coming closer than L_0 to the front red particle. Applying a union bound over all vertices at distance k from v and the estimate (7.30) we have $P(\cup_{|v|=k} D_k(v)) \le 2^{-k}$. As these probabilities are summable, the Borel-Cantelli lemma implies that almost surely only finitely many D_k occur.

On a given vertex self-avoiding path from the root to ∞, let B_t (respectively, R_t) be the distance between the furthest blue (respectively, the furthest red) and the root. In order for B to occur there must exist a path such that $B_t - R_t = m$ infinitely often for some fixed $m < L_0$. Suppose $B_t - R_t \ge m$ for all times and $B_t - R_t = m$ infinitely often. Self-similarity ensures that the vertices at which $B_t - R_t = m$ form a branching process. Using monotonicity of the jump chain probabilities at (7.7), this branching process is dominated by the embedded branching process of renewal vertices (i.e., when $m = 1$). Lemma 7.4.4 ensures that the embedded branching process of renewal vertices is almost surely finite (since it is critical), and thus for each fixed $m < L_0$ the associated branching process is also almost surely finite. As blue can only reach infinitely many sites if either infinitely many D_k occur, or one of the m-renewing branching processes survives, we have $P(B) = 0$. □

7.5 Proofs of Theorems 7.1.1, 7.1.2 , and 7.1.3

Lemma 7.5.1. *If $\lambda \in \Lambda$ then $\rho_c < \rho_e$.*

Proof. Consider the function

$$f_U(\lambda) = \frac{1}{4}\left(\sqrt{(32d+2)\lambda + \lambda^2 + 1} - 3\lambda - 3\right).$$

In the proof of Theorem 7.1.3, we verify that $\rho_c < f_U(\lambda)$. One can further check that $f_U(\lambda_c^-) = 0 = f_U(\lambda_c^+)$ are the only zeros of f_U, so that $f_U(\lambda) > 0$ for $\lambda \in \Lambda$. Moreover, we have

$$\rho_e - f_U(\lambda) = \lambda(d-1) - \frac{1}{4}\left(\sqrt{(32d+2)\lambda + \lambda^2 + 1} - 3\lambda - 3\right).$$

Some algebra gives $\rho_e - f_U(\lambda) = 0$ is equivalent to solving $-8 + \lambda(8d+8) + \lambda^2(8\lambda - 16\lambda^2) = 0$, which has no real solutions for $d \geq 2$. As $\rho_e - f_U(\lambda)$ is continuous in λ and positive at any choice of $d \geq 2$ and λ, we must have $\rho_e - f_U(\lambda) > 0$. We thus arrive at the claimed relation $\rho_c < f_U(\lambda) < \rho_e$. $\square$

Proof of Theorem 7.1.1. First we prove that $\rho_e = \lambda(d-1)$. In CED on $\mathcal{T}_d$ with no blue particles present, red spreads as a branching process with mean offspring $d\lambda/(\lambda+\rho)$. This is supercritical whenever $\rho < \lambda(d-1)$. It is easy to see that $P(R) > 0$ for such ρ, since with positive probability a child of the root becomes red, then the red particle at the root dies. At this point, blue cannot spread, and red spreads like an unchased branching process. Since the red branching process is not supercritical for $\rho \geq \lambda(d-1)$, we have $P(R) = 0$ for such ρ, which proves that ρ_e is as claimed.

Now we prove (i). Suppose that $\lambda \in \Lambda$. Lemma 7.3.4 and Proposition 7.4.1 ensure that $\rho_c > 0$, and Lemma 7.5.1 implies $\rho_c < \lambda(d-1) = \rho_e$. For $0 \leq \rho < \rho_c$, Lemma 7.4.5 implies that coexistence occurs with positive probability. For $\rho_c \leq \rho < \rho_e$, Lemma 7.4.6 and Lemma 7.4.5 imply that $P(B) = 0$, but the red branching process survives so $P(R) > 0$. Thus, escape occurs. For $\rho \geq \rho_e$, red cannot survive in CED with no blue, so extinction occurs.

We end by proving (ii). Suppose that $\lambda \notin \Lambda$. By Lemma 7.3.4 we have $M > d$ for all $\rho \geq 0$. It then follows from Lemma 7.4.2 that $E|\mathcal{B}| < \infty$ and so $P(B) = 0$ for all such ρ. Similar arguments as in (i) that only consider the behaviour of red after it separates from blue can show that the escape and extinction phases occur for $0 \leq \rho < \rho_e$ and $\rho \geq \rho_e$, respectively. $\square$

Proof of Theorem 7.1.2. Lemma 7.3.3 and Proposition 7.4.1 give that $\rho > \rho_c$ implies $M > d$, which, by Lemma 7.4.2, implies that $E|\mathcal{B}| < \infty$. This gives (i). The claim at (ii) holds because Lemma 7.3.4 and Proposition 7.4.1 imply that $M = d$ when $\rho = \rho_c$. Lemma 7.3.6 gives that $M = d$ implies $g(d) = \infty$. Since $E|\mathcal{B}| \geq g(d)$, we obtain the claimed behaviour. $\square$

Proof of Theorem 7.1.3. Define the functions

$$f_L(\lambda) = \frac{1}{4}\left(\sqrt{(8d+2)\lambda + \lambda^2 + 1} - 3\lambda - 3\right),$$
$$f_U(\lambda) = \frac{1}{4}\left(\sqrt{(32d+2)\lambda + \lambda^2 + 1} - 3\lambda - 3\right).$$

We will prove that

$$f_L(\lambda) \leq \rho_c(\lambda) < f_U(\lambda). \tag{7.31}$$

Upon establishing this, it follows immediately that for large enough d we have $c\sqrt{d} \leq \rho_c \leq C\sqrt{d}$ for any $c < \sqrt{\lambda/2}$ and $C > \sqrt{2\lambda}$.

On a given path from the root, the probability that red advances one vertex and then blue advances is given by

$$p = \frac{\lambda}{(1+\lambda+\rho)(1+\lambda+2\rho)}. \tag{7.32}$$

For each vertex v at distance k from the root, let G_v be the event that red and blue alternate advancing on the path from the root to the parent vertex of v, after which red advances to v, and then does not spread to any children of v before the parent of v is coloured blue. Letting

$$c = \frac{\lambda}{1+\lambda+\rho}\frac{1}{1+d\lambda+2\rho}, \tag{7.33}$$

it is easy to see that $P(G_v) = cp^{k-1}$. A renewal occurs at each v for which G_v occurs. It is straightforward to verify that $pd > 1$ whenever $\rho < f_L(\lambda)$. Accordingly, we can choose K large enough so that $cp^{K-1}d^K > 1$, and thus the branching process of these renewals is infinite with positive probability, which implies that $P(B) > 0$.

To prove the upper bound we observe that monotonicity of the transition probabilities at (7.7) in j ensures that the maximal probability jump chain of length $2k$ is the sawtooth path that alternates between height 1 and height 2. This path occurs as a living jump chain with probability p^k with p as in (7.32). As this is the maximal probability jump chain of the C_k many Dyck paths counted by $\mathfrak{R}_k$, we have $g(d) \le \sum_{k=1}^{\infty} C_k p^k d^k$. The radius of convergence for the generating function of the Catalan numbers is 1/4 with convergence also holding at 1/4. Thus, when ρ is large enough that $pd \le 1/4$, we have $g(d) < \infty$. Lemma 7.3.6 and Lemma 7.4.2 then imply that $E|\mathcal{B}| < \infty$ and so $\rho > \rho_c$. It is easily checked that the above inequality holds whenever $\rho \ge f_U(\lambda)$. Thus, $\rho_c < f_U(\lambda)$.

□

7.6 Proof of Theorem 7.1.4

Recall our notation for continued fractions at (7.17). In this section we use the definition of a_j at (7.20) and also define $b_j := a_j d$. The following lemma shows that the nested continued fractions in the definition of f at (7.18) are decreasing when $z = M$.

Lemma 7.6.1. *$K[a_i M, a_{i+1} M, \ldots] \le K[a_{i-1} M, a_i M, \ldots]$ for all $i \ge 1$.*

Proof. Let $f_i(x) = K[a_i x, a_{i+1} x, \ldots]$. Since the f_i are analytic for $x < M$, we must have $f_i(x) < 1$ for all $x < M$ (otherwise f would have a singularity with modulus smaller than M). For any fixed n we can the use monotonicity of K when we change the entries of K one by one and use the fact that $a_j < a_{j-1}$ for all $j \ge 1$ to conclude that

$$K[a_i x, a_{i+1} x, \ldots, a_n x] \le K[a_{i-1} x, a_i x, a_{i+2} x, \ldots, a_{n-1} x]. \tag{7.34}$$

Taking the limit as $n \to \infty$ gives $f_i(x) \le f_{i-1}(x)$ for all $x < M$. Letting $x \uparrow M$, these inequalities continue to hold. □

Note that since $f(x)$ has a pole at M and $f_i(M) \le f_{i-1}(M)$, we must have that $f_0(M) = 1$ and $f_i(M) < 1$ for all $i \ge 0$.

Proof of Theorem 7.1.4. When $\rho = \rho_c$, it follows from Proposition 7.4.1 and Theorem 7.3.5 that a pole of f occurs at $z = d$. Let $K(\lambda, \rho) = K[b_0, b_1, \ldots]$. Due to Lemma 7.6.1 and the equality $f(z) = (1 - K[a_0 z, a_1 z, \ldots])^{-1}$, the singularity at $z = d$ occurs as a result of $K(\lambda, \rho_c(\lambda)) = 1$.

We can use a similar argument as in Corollary 7.3.1 to view f as a meromorphic function in the complex variables λ, ρ and z. Thus, when fixing $z = d$ and considering λ and ρ as nonnegative real numbers, the function $K(\lambda, \rho) = f_{\lambda,\rho}(d)$ is real analytic. Moreover, since K is easily seen to be strictly decreasing in ρ, we have $\partial K/\partial \rho \neq 0$ at $(\lambda, \rho_c(\lambda))$. As $K(\lambda, \rho_c(\lambda)) \equiv 1$, and K is infinitely differentiable, it follows from the implicit function theorem that $\rho_c(\lambda)$ is smooth. $\square$

7.7 Proof of Theorem 7.1.5

We begin this section by describing lower and upper bounds on $C_k^{\lambda,\rho}$. These are easier to analyze than $C_k^{\lambda,\rho}$ and lend insight into the local behaviour of ρ_c. In particular, we obtain if and only if conditions to have $\rho < \rho_c$ (Lemma 7.7.1) and $\rho > \rho_c$ (Lemma 7.7.3). We use these bounds to prove Theorem 7.1.5.

7.7.1 A lower bound on $C_k^{\lambda,\rho}$

The idea is to assign weight 0 to rise and fall steps above a fixed height $m \geq 1$. Accordingly, we introduce the weights

$$\hat{u}(j) = \begin{cases} u(j), & j \leq m \\ 0, & j > m \end{cases}, \qquad \hat{v}(j) = \begin{cases} v(j), & j \leq m \\ 0, & j > m \end{cases}. \tag{7.35}$$

Here $u(j)$ and $v(j)$ are as in (7.4). Let $\hat{C}_k^{\lambda,\rho}$ be the corresponding weighted Catalan numbers. Since $\hat{u}(j) \leq u(j)$ and $\hat{v}(j) \leq v(j)$ for all j, we have $\hat{C}_k^{\lambda,\rho} \leq C_k^{\lambda,\rho}$.

Let $\hat{g}_m(z) = \sum_{k=0}^{\infty} \hat{C}_k^{\lambda,\rho} z^k$. The truncation ensures that

$$\hat{g}_m(z) = K[1, a_0 z, \ldots, a_m z]$$

is a rational function with radius of convergence $\hat{M} \geq M$. Note that $\hat{M}$ depends on λ, ρ and m. For nonnegative real values x we have $\hat{g}_m(x) < \hat{g}_{m+1}(x)$. It follows from the monotone convergence theorem that $\hat{g}_m(x) \to g(x)$ as $m \to \infty$ and thus

$$\lim_{m\to\infty} \hat{M} = M. \tag{7.36}$$

Lemma 7.7.1. *$\rho < \rho_c$ if and only if $K[b_i, \ldots, b_m] > 1$ for some $m \geq 1$ and $0 \leq i \leq m$.*

Proof. Suppose that $\rho < \rho_c$. Lemma 7.3.3 and Proposition 7.4.1 imply that $M < d$. By (7.36), we have $\hat{M} < d$ for a large enough choice of m. Since $\hat{g}_m(x)$ is a rational function, its singularities occur wherever one of the partial denominators $1 - K[a_i x, \ldots, a_m x] = 0$. Let i^* be the largest index such that there is $x_0 < d$ with

$$1 - K[a_{i^*} x_0, \ldots, a_m x_0] = 0.$$

Since i^* is the maximum index for which the equation above holds, we have $K[a_j x, \ldots, a_m x] < 1$ for all $j > i^*$ and $x \in (x_0, d)$. This ensures that $K[a_{i^*} x, \ldots, a_m x]$ does not have a singularity and hence it is a strictly increasing function for $x \in (x_0, d)$. Thus, $K[b_{i^*}, \ldots, b_m] > 1$.

Suppose that $K[b_i, \ldots, b_m] > 1$ for some $m \geq 1$ and $0 \leq i \leq m$. This implies that $\hat{g}_m$ has a singularity of modulus strictly less than d. Thus, $\hat{M} < d$. Since $M \leq \hat{M}$, Lemma 7.4.5 implies that $\rho < \rho_c$. ☐

7.7.2 An upper bound on $C_k^{\lambda,\rho}$

Since $u(j)$ and $v(j)$ are decreasing in j, we obtain an upper bound by assigning weight $u(m)$ and $v(m)$ to all rise and fall steps above height m. More precisely, set

$$\tilde{u}(j) = \begin{cases} u(j), & j < m \\ u(m), & j \geq m \end{cases}, \qquad \tilde{v}(j) = \begin{cases} v(j), & j < m \\ v(m), & j \geq m \end{cases}. \tag{7.37}$$

Let $\tilde{C}_k^{\lambda,\rho}$ be the weighted Catalan number using the weights $\tilde{u}$ and $\tilde{v}$. Also, let $\tilde{g}$ and $\tilde{M}$ be the corresponding generating function and radius of convergence. Since $u(j) \leq \tilde{u}(j)$ and $v(j) \leq \tilde{v}(j)$ for all $j \geq 0$, we have $C_k^{\lambda,\rho} \leq \tilde{C}_k^{\lambda,\rho}$ for all $k \geq 0$. It follows that the corresponding generating functions and radii of convergence satisfy $g(x) \leq \tilde{g}(x)$, and $M \geq \tilde{M}$.

There is a finite procedure for bounding $\tilde{M}$. We say that the function $K[1, c_0, c_1, \ldots, c_k]$ is *good* if all of the partial continued fractions are smaller than 1:

$$c_k < 1, \quad K[c_{k-1}, c_k] < 1, \quad \cdots \quad , \quad K[c_0, \ldots, c_k] < 1. \tag{7.38}$$

Define the continued fraction

$$K_m(x) := K[1, a_0x, \ldots, a_{m-2}x, a_{m-1}x\psi(a_mx)] \tag{7.39}$$

where $\psi(x) = K[1, x, x, \ldots]$ is the generating function of the usual Catalan numbers. By (7.18), so long as $x < \tilde{M}$ we have

$$\tilde{g}(x) = K_m(x).$$

We now prove that when a partial continued fraction is good, it is good in a neighbourhood.

Lemma 7.7.2. *If $K[\mathbf{c}] := K[c_0, c_1, \ldots, c_k]$ is good, then there exists an $\epsilon > 0$ such that $K[\tilde{\mathbf{c}}] := K[\tilde{c}_0, \tilde{c}_1, \ldots, \tilde{c}_k]$ is good when $\tilde{c}_j \leq c_j(1 + \varepsilon)$ for all $j = 0, \ldots, k$.*

Proof. If $K[\mathbf{c}]$ is good, then each partial fraction in (7.38) is < 1. Note that each of those partial fractions is a decreasing function of the c_j which appear in the fraction. Therefore, if $\tilde{c}_j \leq c_j$, $K[\tilde{\mathbf{c}}]$ is good. Since the inequalities are strict and $K[\mathbf{c}]$ is continuous in each c_j, we can extend to have $K[\tilde{\mathbf{c}}]$ good for all $\tilde{c}_j \leq (1 + \epsilon)c_j$ where ϵ is chosen small enough so that none of the partial fractions in (7.38) equal 1. □

Lemma 7.7.3. *$\rho > \rho_c$ if and only if $b_m < 1/4$ and $K_m(d)$ is good for some $m \geq 1$.*

Proof. Suppose that $\rho > \rho_c$. Then Lemma 7.3.3 and Proposition 7.4.1 say that g has radius of convergence $M > d$. For $|z| < M$, we write

$$g(z) = K[1, a_0 z, \ldots, a_{m-2}z, a_{m-1}z g_m(z)] \tag{7.40}$$

with $g_m(z) := K[1, a_m z, a_{m+1}z, \ldots]$ the tail of the continued fraction.

Because $M > d$, $g_m(d) < \infty$, we can see, using a quick argument by contradiction, that $g(d)$ is good. Suppose that one of the partial fractions in (7.38) is strictly larger than 1. But because these finite continued fractions are continuous in the inputs, this would imply that there is a singularity at some $|z| < d$, contradicting the fact that $M > d$.

As the continued fraction representation of $g(d)$ is good, we then apply Lemma 7.7.2 to say that there exists an $\epsilon > 0$ such that

$$K[1, b_0, \ldots, b_{m-2}, b_{m-1}g_m(d)(1+\epsilon)] \tag{7.41}$$

is also good. Notice that $\psi(a_m x) \geq g_m(x) \geq 1$ by definition. Since $|a_m z| \to 0$ as $m \to \infty$, we can use the explicit formula for ψ to directly verify that $\psi(a_m x) \to 1$ as $m \to \infty$. Choose m large enough so that $b_m < 1/4$ and $\psi(b_m) < 1 + \epsilon$. Then $\psi(b_m) < (1+\epsilon)g_m(d)$ and it follows from (7.41) that $K_m(d)$ is good.

Now, suppose that $b_m < 1/4$ and $K_m(d)$ is good for some $m \geq 1$. Our definitions of $\tilde{u}$ and $\tilde{v}$ ensure that we can write

$$K[1, \tilde{a}_m x, \tilde{a}_{m+1}x, \ldots] = K[1, a_m x, a_m x, \ldots] = \psi(a_m x)$$

for all x with $a_m x < 1/4$. Since $b_m = a_m d < 1/4$, this ensures that this is true for all $x < d(1+\epsilon')$ for some $\epsilon' > 0$. Similar reasoning as Corollary 7.3.1 then gives K_z is a meromorphic function for $|z| \leq d(1+\epsilon')$. Moreover, Theorem 7.3.5 ensures that the first pole occurs at the smallest x for which some partial continued fraction of K_m is equal to 1. By Lemma 7.7.2, $K_m(d)$ being good implies that there exists $0 < \epsilon \leq \epsilon'$ such that K_m is good for all $x \leq d(1+\epsilon)$. Hence, there are no poles of K_m within distance $d(1+\epsilon)$ of the origin. It follows that $\tilde{g} = K_m$ for all such x, and thus $\tilde{M} \geq d(1+\epsilon)$.

Since $M \geq \tilde{M}$, we have $M > d$. This gives $\rho > \rho_c$ by Lemma 7.4.2 and Proposition 7.4.1.

$\square$

7.7.3 A finite runtime algorithm

Proof of Theorem 7.1.5. Suppose we are given $\rho \neq \rho_c$. Lemmas 7.7.1 and 7.7.3 give a finite set of conditions to check whether $\rho < \rho_c$ or $\rho > \rho_c$, respectively. To decide which is true, we increase m and the algorithm terminates once the conditions holds. $\square$

Chapter 8

Conclusions

We finish with a summary of the contributions in this work and further research directions.

8.1 Contributions

We divide this discussion between the two main topics in this dissertation.

8.1.1 Uniform spanning trees

In Chapter 5, we show that the scaling limit of the UST on $\mathbb{Z}^3$ exists in the space of measured and rooted spatial trees. Let $\mathcal{U}$ be the uniform spanning tree on $\mathbb{Z}^3$, let $d_{\mathcal{U}}$ denote the intrinsic metric on $\mathcal{U}$, let $\mu_{\mathcal{U}}$ be the counting measure on $\mathcal{U}$ and let $\phi_{\mathcal{U}} : \mathcal{U} \to \mathbb{R}^3$ be a continuous map defined as the identity on the vertices of $\mathcal{U}$ and a linear interpolation between its edges. Recall that we write β for the growth exponent of the loop-erased random walk. Theorem 5.1.1 shows that, if $\mathbf{P}_\delta$ is the law of

$$(\mathcal{U}, \delta^\beta d_{\mathcal{U}}, \delta^3 \mu_{\mathcal{U}}, \delta \phi_{\mathcal{U}}). \tag{8.1}$$

then the collection $(\mathbf{P}_\delta)_{(0,1]}$ is tight. Moreover,

$$(\mathcal{U}, 2^{-\beta n} d_{\mathcal{U}}, 2^{-\beta 3n} \mu_{\mathcal{U}}, 2^{-n} \phi_{\mathcal{U}})$$

converges to $(\mathcal{T}, d_\mathcal{T}, \mu_\mathcal{T}, \phi_\mathcal{T})$ as $n \to \infty$, with respect to a type of Gromov-Hausdorff topology, as defined in Section 5.2. To our knowledge, this is the first result on scaling limits for three-dimensional uniform spanning trees.

The techniques used to prove Theorem 5.1.1 are also useful to study geometric properties of the limit space. If $(\mathcal{T}, d_\mathcal{T})$ is the scaling limit of uniform spanning trees on $\mathbb{Z}^3$, in Theorem 5.1.2 we prove that

(i) the tree $(\mathcal{T}, d_\mathcal{T})$ is one-ended;

(ii) the Hausdorff dimension of $(\mathcal{T}, d_\mathcal{T})$ is $d_f = 3/\beta$;

(iii) the Schramm distance is a well defined metric on $\mathcal{T}$ and it is topologically equivalent to the intrinsic metric $d_\mathcal{T}$.

Theorem 5.1.2 also states upper and lower matching bounds for the measure of the intrinsic balls of $(\mathcal{T}, d_\mathcal{T})$. Theorem 5.1.1 has consequences on the simple random walk defined over the uniform spanning tree. In particular,

(i) we find exponents of the simple random walk in terms of β (see Theorem 5.1.3 and Corollary 5.1.1);

(ii) the annealed law of the simple random walk on $\mathcal{U}$ is tight under rescaling (see Theorem 5.1.4); and

(iii) we obtain kernel estimates for any diffusion that arises as a scaling limit (see Theorem 5.1.4).

The general technique Chapter 5 also serves to study macroscopic properties of the uniform spanning tree. Let us denote by $\mathcal{U}_\delta$ the uniform spanning tree on $\delta\mathbb{Z}^3$. By a spanning cluster, we refer to a connected component of the uniform spanning tree intersecting two opposite faces of the unit cube $[0,1]^3$. Theorem 6.1.1 states that the number of spanning clusters of $\mathcal{U}_\delta$ is tight as $\delta \to 0$.

Theorem 6.1.1 verifies a conjecture by Benjamini [28]. For models below its critical dimension, we expect to have tightness on the number of generating clusters, whereas it grows to infinity on high dimensions. This behaviour has been verified, in all dimensions, only for uniform spanning trees. The affirmative answer that we find for USTs supports the corresponding conjecture for percolation on $\mathbb{Z}^d$ with $d < 6$ (see [3]).

8.1.2 Competitive growth process

Chapter 7 introduces a novel competitive growth process, *chase-escape with death (CED)*. It is a natural generalization for predator-prey competition models.

Chase-escape with death is a competitive growth process on a graph G. It resembles the behaviour of predators and prey, including the possibility of prey dying for reasons unrelated to the predators. We define chase-escape with death as an interacting particle system with local state space $\{w, r, b, \dagger\}$. We call these states white, red, blue, and dead, respectively. We also refer to red and blue as particles. White states indicate vacant sites. Predators, prey, and dead prey may occupy these sites as time advances. A red site indicates the presence of prey, blue is the colour for predators, and † corresponds to a dead site. In the initial configuration, a red particle occupies one vertex, a blue particle occupies a neighbour vertex, and the rest of the graph is white. Red particles spread to adjacent white sites according to a Poisson process with rate λ. Meanwhile, blue particles overtake adjacent red particles at rate 1. Red particles die, and turn to a death state, at an independent rate ρ. Once a site reached a dead state, it stays dead for the rest of the process.

Depending on the parameters λ and ρ, the process reaches either a finite or an infinite number of vertices.

(i) Coexistence phase: both particle types occupy infinitely many sites with positive probability.

(ii) Extinction phase: both types occupy only finitely many sites almost surely.

(iii) Escape phase: red particles occupy infinitely many sites with positive probability, but blue particles reach a finite number of vertices almost surely.

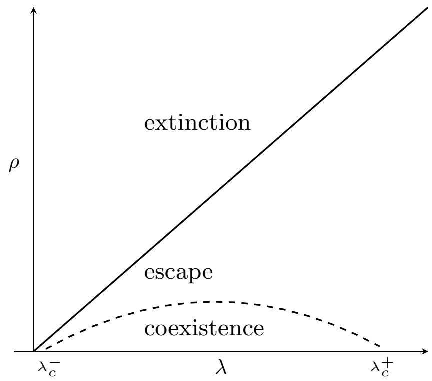

Figure 8.1: The two-dimensional phase diagram with d fixed. The horizontal axis is for values of λ and the vertical axis is for values of ρ. The dashed black line depicts ρ_c and the solid black line ρ_e.

In Chapter 7, we analyse CED when the underlying graph G is a d-ary tree $\mathbb{T}_d$. The main result was Theorem 7.1.1. It states that, for fixed λ within an explicit interval Λ, there exists an interval of death rates $[\rho_c, \rho_e)$ where the escape phase occurs. Coexistence occurs for $\rho \in [0, \rho_c)$ and extinction for $\rho \geq \rho_e$. If $\lambda \notin \Lambda$, then the process transitions between escape and extinction phases. We obtain a characterization of the phase diagram for the spreading parameter λ and the death rate ρ (see Figure 8.1).

In Theorem 7.1.2, we analyse the behaviour at criticality of the average growth of blue particles. If $\rho > \rho_c$, then the expected number of sites that are ever coloured blue is finite. If $\rho = \rho_c$, this expected value is infinite.

A closed formula is not available for ρ_c, but Theorems 7.1.3, 7.1.4 describe the behaviour of this critical parameter. For large d, $\rho_c \asymp_\lambda \sqrt{d}$. Moreover, as as function of $\lambda \in \Lambda$, ρ_c is a smooth function. We can approximate the graph of ρ_c with an algorithm proposed in Theorem 7.1.5.

8.2 Future directions

The common obstacle for addressing open problems on three-dimensional uniform spanning trees and chase-escape models is a lack of tools. Three-dimensional models in statistical mechanics lack properties available in two dimensions, such as conformal invariance and planarity. The results in two-dimensions guide conjectures, but the proofs will require new methods. The situation is similar for chase-escape. We have results for chase-escape and chase-escape with death on $\mathbb{N}$ and trees $\mathbb{T}_d$, and $\mathbb{N} \times \{0, 1\}$. However, chase-escape on $\mathbb{Z}^2$ seems out of reach.

In general, obstacles for solving a mathematical problem maintain the vitality of our science. At first, one may regret the failure of a known technique. But this momentarily defeat demands creativity for a deeper understanding of the problem in hand. Following this spirit, we conclude this thesis with a collection of problems that are not resolved in this work but stimulate further research.

8.2.1 Uniform spanning trees

One the main tools to analyse uniform spanning trees and forests is Wilson's algorithm. As we have seen in this work, questions on uniform spanning trees get an answer after their analogue finds a solution for loop-erased random walks. The problems that we present here follow this direction and correspond to properties of the loop-erased random walk.

Characterization of the scaling limit of the LERW on $\mathbb{Z}^3$

Let $\mathcal{K}$ be the scaling limit of the loop-erased random walk on $\mathbb{Z}^3$ started at the origin and stopped at its first exit from the unit ball. Kozma proved the existence of $\mathcal{K}$ in [107], but a construction of $\mathcal{K}$ without passing through the limit is missing. In [148], Sapozhnikov and Shiraishi compared the traces of $\mathcal{K}$, Brownian motion, and a Brownian loop-soup. Denote by $\mathcal{B}$ the trace of Brownian motion from the origin to the boundary of the unit ball. Let BS be the Brownian loop soup [120], and let $\mathcal{L}$ be the set of loops of BS intersecting $\mathcal{K}$. If we set

$$\widehat{\mathcal{K}} = \mathcal{K} \cup \mathcal{L},$$

then $\mathcal{B}$ is equal in law to $\widehat{\mathcal{K}}$ [148].

Question 8.2.1. *Does the law of $\widehat{\mathcal{K}}$ determine the law of $\mathcal{K}$?*

Universality of the scaling limit of the LERW

We have a good understanding of the scaling limit of the loop-erased random walk in the planar case. For any planar graph where the simple random walk converges to Brownian motion, the LERW converges in the scaling limit to SLE(2) [171]. In contrast, we know little on the three-dimensional scaling limit.

Question 8.2.2. *Let G be a periodic three-dimensional graph different to $\mathbb{Z}^3$. Prove the existence of the scaling limit of the LERW on G.*

8.2.2 Competitive growth models

We already presented an intriguing open question for chase-escape on $\mathbb{Z}^2$ in Conjecture 4.5.6. The simpler version of the conjecture of Kordzakhia Martin is also open.

Question 8.2.3. *Does there exist a graph non-oriented G for which the critical parameter on chase-escape is $\lambda_c(G) < 1$?*

Natural starting points for Question 8.2.3 is a binary tree with one edge added. Such an edge creates a cycles. On a different direction, we can consider Question 8.2.3 on Cayley graphs, such like $\mathbb{Z} * \mathbb{Z} * \mathbb{Z}/3$ or $\mathbb{Z} * \mathbb{Z} * \mathbb{Z}^2$.

The Birth-and-Assassination (BA) process is the scaling limit of chase-escape on a d-ary tree. Aldous and Krebs defined the BA process process in [8]. Bordenave proved this convergence, in the scaling limits, for the rumor-scotching process on trees [34]. The equivalence between the rumor-scotching process chase-escape, on d-ary trees, extends the result to the latter. A natural question is the nature of the scaling limit on other graphs at criticality. Simulations in [164] indicate that such limit exists for $\mathbb{Z}^2$ at $\lambda = \frac{1}{2}$.

Question 8.2.4. *Does the chase-escape process at criticality on a graph G converge to a scaling limit? Does there exist a scaling limit for chase-escape with death?*

Bibliography

[1] R. Abraham, J.-F. Delmas, and P. Hoscheit. A note on the Gromov-Hausdorff-Prokhorov distance between (locally) compact metric measure spaces. *Electron. J. Probab.*, 18(14):1–21, 2013. doi: 10.1214/EJP.v18-2116. → pages 40, 48, 92

[2] D. Ahlberg. Existence and coexistence in first-passage percolation. ArXiv e-prints: 2001.04142. URL https://arxiv.org/abs/2001.04142. → page 69

[3] M. Aizenman. On the number of incipient spanning clusters. *Nuclear Phys. B*, 485(3):551–582, 1997. doi: 10.1016/S0550-3213(96)00626-8. → pages 200, 204, 243

[4] M. Aizenman, A. Burchard, C. M. Newman, and D. B. Wilson. Scaling limits for minimal and random spanning trees in two dimensions. *Random Struct. Algorithms*, 15(3-4):319–367, 1999. doi: 10.1002/(SICI)1098-2418(199910/12)15:3/4<319:: AID-RSA8>3.0.CO;2-G. → page 46

[5] D. Aldous. The random walk construction of uniform spanning trees and uniform labelled trees. *J. SIAM J. Discret. Math.*, 3(4):450–465, 1990. doi: 10.1137/0403039. → pages 35, 37, 38

[6] D. Aldous. The continuum random tree. I. *Ann. Probab.*, 19(1):1–28, 1991. doi: 10.1214/aop/1176990534. → pages 41, 42, 43

[7] D. Aldous. The continuum random tree. III. *Ann. Probab.*, 21(1): 248–289, 1993. doi: 10.1214/aop/1176989404. → page 42

[8] D. Aldous and W. B. Krebs. The birth-and-assassination process. *Statistics & Probability Letters*, 10(5):427–430, 1990. doi: 10.1016/0167-7152(90)90024-2. → pages 217, 246

[9] E. Allen and I. Gheorghiciuc. A weighted interpretation for the super catalan numbers. *J. Integer Seq.*, 17(9), 2014. → page 218

[10] O. Angel, D. A. Croydon, S. Hernandez-Torres, and D. Shiraishi. The number of spanning clusters of the uniform spanning tree in three dimensions. ArXiv e-prints: 2003.04548, . URL https://arxiv.org/abs/2003.04548. → page v

[11] O. Angel, D. A. Croydon, S. Hernandez-Torres, and D. Shiraishi. Scaling limits of the three-dimensional uniform spanning tree and associated random walk. ArXiv e-prints: 2003.09055, . URL https://arxiv.org/abs/2003.09055. → pages v, 201

[12] M. A. Aon, S. Cortassa, and B. O'Rourke. Percolation and criticality in a mitochondrial network. *Proc. Natl. Acad. Sci. U. S. A.*, 101(13): 4447–4452, 2004. ISSN 00278424. doi: 10.1073/pnas.0307156101. → page 2

[13] S. Athreya, M. Eckhoff, and A. Winter. Brownian motion on $\mathbb{R}$-trees. *Trans. Amer. Math. Soc.*, 365(6):3115–3150, 2013. → pages 89, 195

[14] S. Athreya, W. Löhr, and A. Winter. Invariance principle for variable speed random walks on trees. *Ann. Probab.*, 45(2):625–667, 2017. doi: 10.1214/15-AOP1071. → pages 88, 187

[15] S. R. Athreya and A. A. Járai. Infinite volume limit for the stationary distribution of abelian sandpile models. *Commun. Math. Phys.*, 249(1):197–213, 2004. doi: 10.1007/S00220-004-1080-0. → page 35

[16] A. Auffinger and M. Damron. Differentiability at the edge of the percolation cone and related results in first-passage percolation. *Probab. Theory Relat. Fields*, 156:193–227, 2013. doi: 10.1007/s00440-012-0425-4. → page 68

[17] A. Auffinger, M. Damron, and J. Hanson. *50 years of first-passage percolation*, volume 68 of *University lecture series.* American Mathematical Soc., 2017. → pages 59, 63, 66, 67, 217

[18] M. T. Barlow. *Loop Erased Walks and Uniform Spanning Trees*, volume 34 of *MSJ Memoirs*, pages 1–32. The Mathematical Society of Japan, Tokyo, Japan, 2016. doi: 10.2969/msjmemoirs/03401C010. → page 24

[19] M. T. Barlow. *Random walks and heat kernels on graphs*, volume 438 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 2017. → pages 85, 195

[20] M. T. Barlow and R. Masson. Spectral dimension and random walks on the two dimensional uniform spanning tree. *Comm. Math. Phys.*, 305(1):23–57, 2011. doi: 10.1007/s00220-011-1251-8. → pages xi, 87, 142, 143, 146, 153

[21] M. T. Barlow, D. A. Croydon, and T. Kumagai. Quenched and averaged tails of the heat kernel of the two-dimensional uniform spanning tree. In preparation. → pages xi, 87, 88

[22] M. T. Barlow, T. Coulhon, and T. Kumagai. Characterization of sub-Gaussian heat kernel estimates on strongly recurrent graphs. *Comm. Pure Appl. Math.*, 58(12):1642–1677, 2005. doi: 10.1002/cpa.20091. → page 194

[23] M. T. Barlow, A. A. Járai, T. Kumagai, and G. Slade. Random walk on the incipient infinite cluster for oriented percolation in high dimensions. *Comm. Math. Phys.*, 278(2):385–431, 2008. doi: 10.1007/s00220-007-0410-4. → page 86

[24] M. T. Barlow, D. A. Croydon, and T. Kumagai. Subsequential scaling limits of simple random walk on the two-dimensional uniform spanning tree. *Ann. Probab.*, 45(1):4–55, 2017. doi: 10.1214/15-AOP1030. → pages xi, 47, 78, 79, 84, 85, 87, 88, 89, 91, 92, 187, 188, 191, 193, 195, 196, 201

[25] E. Beckman, K. Cook, N. Eikmeier, S. Hernandez-Torres, and J. M. Chase-escape with death on trees. ArXiv e-prints: 1909.01722. URL https://arxiv.org/abs/1909.01722. → page v

[26] A. A. Belavin, A. M. Polyakov, and A. B. Zamolodchikov. Infinite conformal symmetry in two-dimensional quantum field theory. *Nucl. Physics, Sect. B*, 241(2):333–380, 1984. ISSN 05503213. doi: 10.1016/0550-3213(84)90052-X. → page 4

[27] C. Beneš. Some estimates for planar random walk and brownian motion. ArXiv e-prints: 0611127. URL https://arxiv.org/abs/math/0611127. → page 16

[28] I. Benjamini. Large scale degrees and the number of spanning clusters for the uniform spanning tree. In *Perplexing problems in probability*, volume 44 of *Progr. Probab.*, pages 175–183. Birkhäuser Boston, 1999. doi: 10.1007/978-1-4612-2168-5. → pages 198, 199, 200, 242

[29] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm. Uniform spanning forests. *Ann. Probab.*, 29(1):1–65, 2001. doi: 10.1214/aop/1008956321. → pages 33, 35, 37, 122, 199, 203

[30] N. Berestycki and J. Norris. *Lectures on Schramm–Loewner Evolution*. 2014. → pages 27, 28, 29

[31] N. Berestycki, B. Laslier, and G. Ray. Dimers and imaginary geometry. *Ann. Probab.*, 48(1):1–52, 2020. doi: 10.1214/18-aop1326. → page 35

[32] A. Beurling. *Études sur un problème de majoration*. Imprimerie Almquist and Wiksell, 1933. → page 16

[33] P. Billingsley. *Convergence of Probability Measures*. Wiley-Interscience, 2nd edition, 1999. → pages 180, 181, 182

[34] C. Bordenave. On the birth-and-assassination process, with an application to scotching a rumor in a network. *Electron. J. Probab.*, 13(66):2014–2030, 2008. doi: 10.1214/EJP.v13-573. → pages 217, 246

[35] C. Bordenave. Extinction probability and total progeny of predator-prey dynamics on infinite trees. *Electron. J. Probab*, 19(20): 1–33, 2014. doi: 10.1214/EJP.v19-2361. → pages 70, 74, 214, 215

[36] C. Borgs, J. Chayes, H. Kesten, and J. Spencer. Uniform boundedness of critical crossing probabilities implies hyperscaling. *Random Struct. Alg.*, 15(3-4):368–413, 1999. doi: 10.1002/(SICI)1098-2418(199910/12)15:3/4<368::AID-RSA9>3.0.CO;2-B. → page 200

[37] S. R. Broadbent and J. M. Hammersley. Percolation processes: I. Crystals and mazes. *Math. Proc. Cambridge Philos. Soc.*, 1957. doi: 10.1017/S0305004100032680. → page 3

[38] A. Broder. Generating random spanning trees. pages 442–447, 1989. doi: 10.1109/SFCS.1989.63516. → pages 35, 37, 38

[39] D. L. Burkholder, D. Daley, H. Kesten, P. Ney, F. Spitzer, J. M. Hammersley, and J. F. C. Kingman. Discussion on professor Kingman's paper. *Ann. Probab.*, 1(6):900–909, 1973. doi: 10.1214/aop/1176996799. → page 59

[40] R. Burton and R. Pemantle. Local characteristics, entropy and limit theorems for spanning trees and domino tilings via transfer-impedances. *Ann. Probab.*, 21(3):1329–1371, 1993. doi: 10.1214/aop/1176989121. → page 35

[41] F. Camia and C. M. Newman. Two-dimensional critical percolation: The full scaling limit. *Commun. Math. Phys.*, 268(1):1–38, 2006. doi: 10.1007/s00220-006-0086-1. → page 27

[42] F. Camia and C. M. Newman. Critical percolation exploration path and SLE 6: A proof of convergence. *Probab. Theory Relat. Fields*, 139(3–4):473–519, 2007. doi: 10.1007/s00440-006-0049-7. → page 27

[43] J. Cardy. The number of incipient spanning clusters in two-dimensional percolation. *J. Phys. A*, 31(5):L105–L110, feb 1998. doi: 10.1088/0305-4470/31/5/003. → page 200

[44] S. H. Chan. Infinite-step stationarity of rotor walk and the wired spanning forest. ArXiv e-prints: 1909.13195. URL https://arxiv.org/abs/1909.13195. → page 35

[45] S. H. Chan. Rotor walks on transient graphs and the wired spanning forest. *SIAM J. Discret. Math.*, 33(4):2369 – 2393, 2019. doi: 10.1137/18M1217139. → page 35

[46] D. Chelkak, H. Duminil-Copin, C. Hongler, A. Kemppainen, and S. Smirnov. Convergence of Ising interfaces to Schramm's SLE curves. *Comptes Rendus Math.*, 352(2):157 – 161, 2014. doi: 10.1016/j.crma.2013.12.002. → page 26

[47] E. A. Codling, M. J. Plank, and S. Benhamou. Random walk models in biology. *Journal of the Royal Society, Interface*, 5(25):813–834, August 2008. ISSN 1742-5689. doi: 10.1098/rsif.2008.0014. → page 2

[48] J. T. Cox and R. Durrett. Some limit theorems for percolation processes with necessary and sufficient conditions. *Ann. Probab.*, 9 (4):583–603, 1981. doi: 10.1214/aop/1176994364. → pages 59, 64

[49] D. A. Croydon. Heat kernel fluctuations for a resistance form with non-uniform volume growth. *Proc. Lond. Math. Soc. (3)*, 94(3): 672–694, 2007. doi: 10.1112/plms/pdl025. → pages 89, 194, 196

[50] D. A. Croydon. Convergence of simple random walks on random discrete trees to Brownian motion on the continuum random tree. *Ann. Inst. Henri Poincaré Probab. Stat.*, 44(6):987–1019, 2008. doi: 10.1214/07-AIHP153. → page 88

[51] D. A. Croydon. Volume growth and heat kernel estimates for the continuum random tree. *Probab. Theory Related Fields*, 140(1-2): 207–238, 2008. doi: 10.1007/s00440-007-0063-4. → pages 89, 196

[52] D. A. Croydon. Hausdorff measure of arcs and Brownian motion on Brownian spatial trees. *Ann. Probab.*, 37(3):946–978, 2009. doi: 10.1214/08-AOP425. → page 88

[53] D. A. Croydon. Scaling limits for simple random walks on random ordered graph trees. *Adv. in Appl. Probab.*, 42(2):528–558, 2010. doi: 10.1017/S0001867800004183.

[54] D. A. Croydon. Scaling limits of stochastic processes associated with resistance forms. *Ann. Inst. Henri Poincaré Probab. Stat.*, 54(4): 1939–1968, 2018. doi: 10.1214/17-AIHP861. → pages 88, 196

[55] M. Deijfen and O. Häggström. The initial configuration is irrelevant for the possibility of mutual unbounded growth in the two-type Richardson model. *Comb. Probab. Comput.*, 15(3):345–353, 2006. doi: 10.1017/S0963548305007315. → pages 62, 66

[56] M. Deijfen and O. Häggström. *Analysis and Stochastics of Growth Processes and Interface Models*, chapter The pleasures and pains of studying the two-type Richardson model, pages 39 – 55. Oxford University Press, 2008. doi: 10.1093/acprof:oso/9780199239252.001.0001. → page 217

[57] M. Deijfen and O. Häggström. The two-type Richardson model with unbounded initial configurations. *Ann. Appl. Probab.*, 17(5-6): 1639–1656, 2007. doi: 10.1214/07-AAP440. → page 62

[58] A. Depperschmidt, A. Greven, and P. Pfaffelhuber. Marked metric measure spaces. *Electron. Commun. Probab.*, 16:174–188, 2011. doi: 10.1214/ECP.v16-1615. → pages 184, 185, 186, 187

[59] P. G. Doyle and J. L. Snell. *Random walks and electric networks*, volume 22 of *Carus Mathematical Monographs*. Mathematical Association of America, 1984. doi: 10.4169/j.ctt5hh804. → page 85

[60] R. Durrett. *Random graph dynamics*, volume 20 of *Cambridge series in statistical and probabilistic mathematics*. Cambridge University Press, 2007. ISBN 9780511546594. doi: 10.1017/CBO9780511546594. → page 5

[61] R. Durrett. *Probability: theory and examples*, volume 49 of *Cambridge series in statistical and probabilistic mathematics*. Cambridge University Press, 5th edition, 2019. ISBN 9781108591034. doi: 10.1017/9781108591034. → page 13

[62] R. Durrett, M. Junge, and S. Tang. Coexistence in chase-escape. *Electron. Commun. Probab.*, 25(22):1 – 14, 2020. doi: 10.1214/20-ECP302. → pages 73, 74, 75, 217

[63] M. Eden. A two-dimensional growth process. In *Proc. 4th Berkeley Symposium on Mathematics Statistics and Probability*, volume 4, pages 223–239, 1961. URL https://projecteuclid.org/euclid.bsmsp/1200512888. → page 59

[64] G. A. Edgar. *Integral, probability, and fractal measures*. Springer-Verlag, New York, 1998. doi: 10.1007/978-1-4757-2958-0. → page 191

[65] A. Einstein. Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung. *Ann. Phys.*, 322(8):549 – 560, 1905. URL https://einsteinpapers.press.princeton.edu/vol2-doc/. → page 2

[66] S. N. Ethier and T. G. Kurtz. *Markov processes: characterization and convergence*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons Inc., New York, 1986. ISBN 0-471-08186-8. → page 20

[67] S. N. Evans, J. Pitman, and A. Winter. Rayleigh processes, real trees, and root growth with re-grafting. *Probab. Theory Related Fields*, 134(1):81–126, 2006. doi: 10.1007/s00440-004-0411-6. → pages 40, 187

[68] P. Flajolet and F. Guillemin. The formal theory of birth-and-death processes, lattice path combinatorics and continued fractions.

Advances in Applied Probability, 32(3):750–778, 2000. doi: 10.1017/S0001867800010247. → page 218

[69] P. Flajolet and R. Sedgewick. *Analytic Combinatorics.* Cambridge University Press, 2009. ISBN 9781139477161. doi: 10.1017/CBO9780511801655. → pages 228, 229

[70] P. J. Flory. *Principles of polymer chemistry.* Cornell Univ. Press, 1953. ISBN 9788184120134. doi: 10.1126/science.119.3095.555-a. → page 2

[71] P. J. Flory. *Statistical mechanics of chain molecules.* Interscience Publishers, 1969. ISBN 9780470264959. → page 2

[72] O. Garet and R. Marchand. Coexistence in two-type first-passage percolation models. *Ann. Appl. Probab.*, 15(1A):298–330, 2005. doi: 10.1214/105051604000000503. → pages 61, 67

[73] I. P. Goulden and D. M. Jackson. *Combinatorial enumeration.* Wiley-Interscience series in discrete mathematics. Wiley, 1983. → page 224

[74] A. Greven, P. Pfaffelhuber, and A. Winter. Convergence in distribution of random metric measure spaces (Λ-coalescent measure trees). *Probab. Theory Related Fields*, 145(1-2):285–322, 2009. doi: 10.1007/s00440-008-0169-3. → pages 184, 187

[75] D. Griffeath. Frank Spitzer's Pioneering Work on Interacting Particle Systems. *Ann. Probab.*, 21(2):608 – 621, 1993. doi: 10.1214/aop/1176989258. → page 3

[76] G. Grimmett. *The Random-Cluster Model*, volume 333 of *Institute of Mathematical Statistics Textbooks.* Springer-Verlag Berlin Heidelberg, 2006. doi: 10.1007/978-3-540-32891-9. → page 35

[77] G. Grimmett. *Probability on Graphs: Random Processes on Graphs and Lattices.* Institute of Mathematical Statistics Textbooks. Cambridge University Press, 2010. doi: 10.1017/CBO9780511762550. → page 35

[78] O. Häggström. Random-cluster measures and uniform spanning trees. *Stoch. Process. Their Appl.*, 59(2):267 – 275, 1995. doi: 10.1016/0304-4149(95)00042-6. → page 35

[79] O. Häggström and R. Pemantle. First passage percolation and a model for competing spatial growth. *Journal of Applied Probability*, 35(3):683–692, 1998. doi: 10.1017/S0021900200016338. → pages 60, 61, 62, 67

[80] O. Häggström and R. Pemantle. Absence of mutual unbounded growth for almost all parameter values in the two-type Richardson model. *Stoch. Process. Their Appl.*, 90(2):207–222, 2000. doi: 10.1016/S0304-4149(00)00042-9. → pages 61, 66

[81] J. M. Hammersley and D. J. A. Welsh. First-Passage Percolation, Subadditive Processes, Stochastic Networks, and Generalized Renewal Theory. In N. J. and L. C. L.M., editors, *Bernoulli 1713 Bayes 1763 Laplace 1813*, pages 61–110. Springer, Berlin, Heidelberg, 1965. doi: 10.1007/978-3-642-99884-3_7. → pages 3, 63

[82] A. Hinsen, B. Jahnel, E. Cali, and J.-P. Wary. Phase transitions for chase-escape models on gilbert graphs. ArXiv e-prints: 1911.02622, 2019. URL https://arxiv.org/abs/1911.02622. → page 217

[83] C. Hoffman. Coexistence for Richardson type competing spatial growth models. *Ann. Appl. Probab.*, 15(1B):739–747, 2005. doi: 10.1214/105051604000000729. → pages 61, 67

[84] C. Hoffman. Geodesics in first passage percolation. *The Annals of Applied Probability*, 18(5):1944–1969, 2008. doi: 10.1214/07-AAP510. → pages 63, 68

[85] R. v. d. Hofstad. *Random Graphs and Complex Networks*, volume 1 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, 2016. doi: 10.1017/9781316779422. → pages 5, 6

[86] N. Holden and X. Sun. SLE as a mating of trees in Euclidean geometry. *Comm. Math. Phys.*, 364(1):171–201, 2018. doi: 10.1007/s00220-018-3149-1. → pages 48, 81, 201

[87] A. E. Holroyd, L. Levine, K. Mészáros, Y. Peres, J. Propp, and D. B. Wilson. Chip-Firing and Rotor-Routing on Directed Graphs. In *In and Out of Equilibrium 2*, volume 60 of *Progress in Probability*, pages 331–364. Birkhäuser Basel, 2008. doi: 10.1007/978-3-7643-8786-0_17. → page 35

[88] T. Hutchcroft. Interlacements and the wired uniform spanning forest. *Ann. Probab.*, 46(2):1170–1200, 2018. doi: 10.1214/17-AOP1203. → pages 35, 38

[89] T. Hutchcroft. Universality of high-dimensional spanning forests and sandpiles. *Probab. Theory Relat. Fields*, 176(1-2):533–597, 2020. doi: 10.1007/s00440-019-00923-3. → page 35

[90] A. A. Járai. The uniform spanning tree and related models, 2009. URL https://people.bath.ac.uk/aj276/teaching/USF/USFnotes.pdf. → page 35

[91] A. A. Járai and F. Redig. Infinite volume limit of the Abelian sandpile model in dimensions $d \geq 3$. *Probab. Theory Relat. Fields*, 141(1):181–212, 2008. doi: 10.1007/s00440-007-0083-0. → page 35

[92] A. A. Járai and N. Werning. Minimal Configurations and Sandpile Measures. *J. Theor. Probab.*, 27(1):153–167, 2014. doi: 10.1007/s10959-012-0446-z. → page 35

[93] O. Kallenberg. *Foundations of modern probability.* Probability and its Applications. Springer-Verlag, 2nd edition, 2002. doi: 10.1007/b98838. → page 183

[94] I. Karatzas and S. E. Shreve. *Brownian motion and stochastic calculus*, volume 113 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2nd edition, 1991. → page 179

[95] R. Kenyon. The asymptotic determinant of the discrete Laplacian. *Acta Math.*, 185(2):239–286, 2000. doi: 10.1007/BF02392811. → pages xi, 24, 35, 80, 87

[96] H. Kesten. Aspects of first passage percolation. In *École d'été de probabilités de Saint-Flour, XIV—1984*, volume 1180 of *Lecture Notes in Math.*, pages 125–264. Springer, Berlin, 1986. doi: 10.1007/BFb0074919. → page 66

[97] H. Kesten. Hitting probabilities of random walks on $\mathbf{Z}^d$. *Stochastic Process. Appl.*, 25(2):165–184, 1987. → pages 17, 82

[98] J. Kigami. Harmonic calculus on limits of networks and its application to dendrites. *J. Funct. Anal.*, 128(1):48–86, 1995. doi: 10.1006/jfan.1995.1023. → pages 89, 195

[99] J. Kigami. Resistance forms, quasisymmetric maps and heat kernel estimates. *Mem. Amer. Math. Soc.*, 216(1015):vi+132, 2012. → page 195

[100] J. Kingman. *Poisson Processes*. Oxford science publications. Clarendon Press, 1993. ISBN 9780198536932. → page 57

[101] G. Kirchhoff. Ueber die auflösung der gleichungen, auf welche man bei der untersuchung der linearen vertheilung galvanischer ströme geführt wird. *Ann. Phys. Chem.*, 148(12):497–508, 1847. doi: 10.1002/andp.18471481202. → page 35

[102] A. Klenke. *Probability Theory*. Springer-Verlag London, 2008. doi: 10.1007/978-1-84800-048-3. → pages 20, 51, 57

[103] S. Kliem and W. Löhr. Existence of mark functions in marked metric measure spaces. *Electron. J. Probab.*, 20(73):1–24, 2015. doi: 10.1214/EJP.v20-3969. → pages 185, 187

[104] G. Kordzakhia. The escape model on a homogeneous tree. *Electron. Commun. Probab.*, 10:113–124, 2005. doi: 10.1214/ECP.v10-1140. → pages 70, 73, 217

[105] I. Kortchemski. A predator-prey SIR type dynamics on large complete graphs with three phase transitions. *Stoch. Process. Their Appl.*, 125(3):886–917, 2015. doi: 10.1016/j.spa.2014.10.005. → pages 72, 217

[106] I. Kortchemski. Predator–prey dynamics on infinite trees: A branching random walk approach. *Journal of Theoretical Probability*, 29(3):1027–1046, 2016. doi: 10.1007/s10959-015-0603-2. → pages 74, 217

[107] G. Kozma. The scaling limit of loop-erased random walk in three dimensions. *Acta Math.*, 199(1):29–152, 2007. doi: 10.1007/s11511-007-0018-8. → pages 30, 78, 81, 104, 109, 158, 245

[108] T. Kumagai. Heat kernel estimates and parabolic Harnack inequalities on graphs and resistance forms. *Publ. Res. Inst. Math. Sci.*, 40(3):793–818, 2004. doi: 10.2977/prims/1145475493. → page 194

[109] T. Kumagai and J. Misumi. Heat kernel estimates for strongly recurrent random walk on random media. *J. Theoret. Probab.*, 21(4): 910–935, 2008. doi: 10.1007/s10959-008-0183-5. → pages 86, 87, 194

[110] G. F. Lawler. A self-avoiding random walk. *Duke Math. J.*, 47(3): 655–693, 1980. doi: 10.1215/S0012-7094-80-04741-9. → pages 21, 31, 32

[111] G. F. Lawler. Gaussian behavior of loop-erased self-avoiding random walk in four dimensions. *Duke Math. J.*, 53(1):249–269, 1986. doi: 10.1215/S0012-7094-86-05317-2. → page 31

[112] G. F. Lawler. Intersections of random walks with random sets. *Isr. J. Math.*, 65(2):113–132, 1989. doi: 10.1007/BF02764856. → page 17

[113] G. F. Lawler. *Intersections of random walks*. Probability and its Applications. Birkhäuser Boston, 1991. doi: 10.1007/978-1-4757-2137-9. → pages 14, 15, 16, 17, 22, 31, 32, 115, 116, 120, 123, 128, 142, 144, 146, 189, 206

[114] G. F. Lawler. Loop-erased random walk. In *Perplexing problems in probability*, volume 44 of *Progr. Probab.*, pages 197–217. Birkhäuser Boston, 1999. doi: 10.1007/978-1-4612-2168-5. → pages 24, 25, 80, 203

[115] G. F. Lawler. *Conformally Invariant Processes in the Plane*, volume 14 of *Mathematical Surveys and Monographs*. American Mathematical Society, 2005. → pages 16, 78

[116] G. F. Lawler. The probability that planar loop-erased random walk uses a given edge. *Electron. Commun. Probab.*, 19, 2014. doi: 10.1214/ECP.v19-2908. → pages 24, 47

[117] G. F. Lawler and V. Limic. *Random walk: a modern introduction*, volume 123 of *Cambridge studies in advanced mathematics*. Cambridge University Press, 2010. ISBN 978-0-521-51918-2. doi: 10.1017/CBO9780511750854. → pages 14, 15, 20, 169

[118] G. F. Lawler and J. Perlman. *Random Walks, Random Fields, and Disordered Systems*, chapter Loop Measures and the Gaussian Free Field, pages 211–235. Springer International Publishing, Cham, 2015. ISBN 978-3-319-19339-7. doi: 10.1007/978-3-319-19339-7_5. → page 35

[119] G. F. Lawler and F. Viklund. Convergence of loop-erased random walk in the natural parametrization. ArXiv e-prints: 1603.05203. URL https://arxiv.org/abs/1603.05203. → page 30

[120] G. F. Lawler and W. Werner. The Brownian loop soup. *Probab. Theory Relat. Fields*, 128(4):565–588, 2004. doi: 10.1007/s00440-003-0319-6. → page 245

[121] G. F. Lawler, O. Schramm, and W. Werner. Conformal restriction: The chordal case. *Journal of the American Mathematical Society*, 16 (4):917–955, 2003. doi: 10.1090/S0894-0347-03-00430-2. → page 27

[122] G. F. Lawler, O. Schramm, and W. Werner. Conformal Invariance of Planar Loop-Erased Random Walks and Uniform Spanning Trees. *Ann. Probab.*, 32(1):939–995, 2004. doi: 10.1214/aop/1079021469. → pages 26, 27, 30, 46, 47, 78, 81

[123] G. F. Lawler, X. Sun, and W. Wu. Four-Dimensional Loop-Erased Random Walk. *Ann. Probab.*, 47(6):3866–3910, 2019. doi: 10.1214/19-AOP1349. → page 35

[124] J.-F. Le Gall. Random trees and applications. *Probab. Surv.*, 2: 245–311, 2005. doi: 10.1214/154957805100000140. → page 42

[125] J.-F. Le Gall. Random real trees. *Ann. Fac. Sci. Toulouse Math. (6)*, 15(1):35–62, 2006. → page 91

[126] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, Providence, RI, 2009. With a chapter by James G. Propp and David B. Wilson. → page 85

[127] X. Li and D. Shiraishi. Convergence of three-dimensional loop-erased random walk in the natural parametrization. ArXiv e-prints: 1811.11685. URL https://arxiv.org/abs/1811.11685. → pages 78, 81, 105, 109, 112, 117, 139, 150, 158, 168, 170, 171, 190

[128] X. Li and D. Shiraishi. One-point function estimates for loop-erased random walk in three dimensions. *Electron. J. Probab.*, 24:Paper No. 111, 46, 2019. doi: 10.1214/19-EJP361. → pages 78, 104, 114, 118, 128, 130, 131, 133, 135

[129] T. M. Liggett. *Interacting Particle Systems*. Springer-Verlag Berlin Heidelberg, 1985. doi: 10.1007/b138374. → pages 51, 53, 54, 55, 57

[130] T. M. Liggett. *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes*. Springer, Berlin, Heidelberg, 1999. doi: 10.1007/978-3-662-03990-8. → pages 54, 55

[131] K. Löwner. Untersuchungen über schlichte konforme Abbildungen des Einheitskreises. I. *Math. Ann.*, 89:103–121, 1923. doi: 10.1007/BF01448091. → page 28

[132] R. Lyons and Y. Peres. *Probability on trees and networks*, volume 42 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, New York, 2016. doi: 10.1017/9781316672815. → pages 33, 34, 35, 36, 85

[133] R. Lyons, Y. Peres, and O. Schramm. Markov chain intersections and the loop-erased walk. *Ann. Inst. Henri Poincaré Probab. Stat.*, 39:779–791, 2003. doi: 10.1016/S0246-0203(03)00033-5. → page 107

[134] S. N. Majumdar and D. Dhar. Equivalence between the Abelian sandpile model and the $q \to 0$ limit of the Potts model. *Phys. A Stat. Mech. its Appl.*, 1992. doi: 10.1016/0378-4371(92)90447-X. → page 35

[135] J. F. Marckert and G. Miermont. The CRT is the scaling limit of unordered binary trees. *Random Struct. Algorithms*, 38(4):467–501, 2011. doi: 10.1002/rsa.20332. → page 43

[136] R. Masson. The growth exponent for planar loop-erased random walk. *Electron. J. Probab.*, 14:no. 36, 1012–1073, 2009. doi: 10.1214/EJP.v14-651. → pages 23, 25, 47, 103, 104, 117, 118, 154

[137] M. Mézard and A. Montanari. *Information, Physics, and Computation*. Oxford University Press, 2009. ISBN 9780191718755. doi: 10.1093/acprof:oso/9780198570837.001.0001. → page 2

[138] P. Michaeli, A. Nachmias, and M. Shalev. The diameter of uniform spanning trees in high dimensions. ArXiv e-prints: 1911.12319. URL https://arxiv.org/abs/1911.12319. → pages 44, 45

[139] R. A. Minlos. R.L. Dobrushin - one of the founders of modern mathematical physics. *Russ. Math. Surv.*, 1997. ISSN 0036-0279. doi: 10.1070/rm1997v052n02abeh001772. → page 3

[140] P. Mörters and Y. Peres. *Brownian motion*, volume 30 of *Cambridge series in statistical and probabilistic mathematics*. Cambridge University Press, 2010. doi: 10.1017/CBO9780511750489. → page 18

[141] C. M. Newman. A surface view of first-passage percolation. In *Proceedings of the International Congress of Mathematicians (Zürich, 1994)*, volume 1,2, page 10171023. Birkhäuser, 1995. → page 67

[142] M. E. Newman. Spread of epidemic disease on networks. *Phys. Rev. E - Stat. Physics, Plasmas, Fluids, Relat. Interdiscip. Top.*, 66(1 Pt 2), 2002. ISSN 1063651X. doi: 10.1103/PhysRevE.66.016128. → page 2

[143] R. Pemantle. Choosing a spanning tree for the integer lattice uniformly. *Ann. Probab.*, 19(4):1559–1574, 1991. doi: 10.1214/aop/1176990223. → pages 33, 34, 78, 79, 82, 199, 203

[144] Y. Peres and D. Revelle. Scaling limits of the uniform spanning tree and loop-erased random walk on finite graphs. ArXiv e-prints: 0410430. URL https://arxiv.org/abs/math/0410430. → page 45

[145] J. Perrin. Mouvement brownien et réalité moléculaire. *Ann. Chim. Phys.*, 18, 1909. → page 2

[146] A. Postnikov and B. E. Sagan. What power of two divides a weighted catalan number? *Journal of Combinatorial Theory, Series A*, 114(5): 970 – 977, 2007. doi: 10.1016/j.jcta.2006.09.007. → page 218

[147] D. Richardson. Random growth in a tessellation. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 74, pages 515–528. Cambridge University Press, 1973. doi: 10.1017/S0305004100077288. → pages 58, 59, 60, 212, 217

[148] A. Sapozhnikov and D. Shiraishi. On Brownian motion, simple paths, and loops. *Probab. Theory Relat. Fields*, 172(3):615–662, 2018. doi: 10.1007/s00440-017-0817-6. → pages 25, 31, 82, 104, 106, 107, 109, 111, 123, 124, 128, 132, 134, 139, 141, 145, 155, 205, 208, 245

[149] O. Schramm. Scaling limits of loop-erased random walks and uniform spanning trees. *Israel J. Math.*, 118:221–288, 2000. doi: 10.1007/BF02803524. → pages 26, 29, 39, 46, 47, 78, 80, 81, 83, 84, 85, 95

[150] O. Schramm and S. Sheffield. Harmonic explorer and its convergence to sle 4. *Ann. Probab.*, 33(6):2127–2148, 2005. doi: 10.1214/009117905000000477. → page 27

[151] J. Schweinsberg. The loop-erased random walk and the uniform spanning tree on the four-dimensional discrete torus. *Probab. Theory Relat. Fields*, 144(3):319–370, 2009. doi: 10.1007/s00440-008-0149-7. → page 45

[152] S. Shader and M. G. Liu. Weighted Catalan numbers and their divisibility properties. *PhD. Thesis*, 2015. → page 218

[153] Z. Shi. Random Walks and Trees. In *ESAIM Proc.*, volume 31, pages 1–39, 2011. doi: 10.1051/proc/2011002. → page 50

[154] Z. Shi. *Branching Random Walks*. École d'Été de Probabilités de Saint-Flour. Springer International Publishing, 2012. doi: 10.1007/978-3-319-25372-5. → page 50

[155] D. Shiraishi. Growth exponent for loop-erased random walk in three dimensions. *Ann. Probab.*, 46(2):687–774, 2018. doi: 10.1214/16-AOP1165. → pages 24, 78, 80, 104, 109, 113, 114, 115, 118, 128, 130, 131, 133, 135, 145

[156] D. Shiraishi. Hausdorff dimension of the scaling limit of loop-erased random walk in three dimensions. *Ann. Inst. H. Poincaré Probab. Statist.*, 55(2):791–834, 2019. doi: 10.1214/18-AIHP899. → page 31

[157] G. Slade. Probabilistic models of critical phenomena. In T. Gowers, J. Barrow-Green, and I. Leader, editors, *The Princeton Companion to Mathematics*. Princeton University Press, 2008. → pages 5, 6

[158] A. Sly. Computational transition at the uniqueness threshold. In *Proc. - Annu. IEEE Symp. Found. Comput. Sci. FOCS*, 2010. doi: 10.1109/FOCS.2010.34. → page 2

[159] S. Smirnov. Critical percolation in the plane. ArXiv e-prints: 0909.4499. URL https://arxiv.org/abs/0909.4499. → page 27

[160] S. Smirnov. Critical percolation in the plane: conformal invariance, Cardy's formula, scaling limits. *C. R. Acad. Sci. Paris*, 333(3): 239–244, 2001. doi: 10.1016/S0764-4442(01)01991-7. → pages 27, 78

[161] R. P. Stanley. *Catalan Numbers*. Cambridge: Cambridge University Press, 2015. doi: 10.1017/CBO9781139871495. → page 220

[162] J. Swart. A Course in Interacting Particle Systems. ArXiv e-prints: 1703.10007v2, 2020. URL https://arxiv.org/abs/1703.10007v2. → pages 55, 57

[163] A. Szabó and R. M. Merks. Cellular Potts modeling of tumor growth, tumor invasion, and tumor evolution. *Front. Oncol.*, 3(87), 2013. ISSN 2234943X. doi: 10.3389/fonc.2013.00087. → page 2

[164] S. Tang, G. Kordzakhia, and S. P. Lalley. Phase Transition for the Chase-Escape Model on 2D Lattices. *ArXiv e-prints: 1807.08387*, 2018. → pages 75, 217, 246

[165] H. Wall. *Analytic Theory of Continued Fractions*. Dover Books on Mathematics. Dover Publications, 2018. ISBN 9780486830445. → page 225

[166] W. Werner. Beurling's projection theorem via one-dimensional Brownian motion. *Math. Proc. Cambridge Philos. Soc.*, 119(4): 729–738, 1996. doi: 10.1017/S0305004100074557. → page 16

[167] W. Werner and E. Powell. Lecture notes on the gaussian free field. ArXiv e-prints: 2004.04720. URL https://arxiv.org/abs/2004.04720. → page 35

[168] J. C. Wierman and W. Reh. On conjectures in first passage percolation theory. *Ann. Probab.*, 6(3):388–397, 1978. doi: 10.1214/aop/1176995525. → page 66

[169] D. B. Wilson. Generating random spanning trees more quickly than the cover time. In *Proceedings of the Twenty-eighth Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996)*, pages 296–303. ACM, New York, 1996. doi: 10.1145/237814.237880. → pages 35, 37, 78, 81, 122, 155, 203

[170] D. B. Wilson. Dimension of loop-erased random walk in three dimensions. *Phys. Rev. E*, 82(6):062102, 2010. doi: 10.1103/PhysRevE.82.062102. → pages xi, 80, 87

[171] A. Yadin and A. Yehudayoff. Loop-erased random walk and Poisson kernel on planar graphs. *Ann. Probab.*, 39(4):1243–1285, 2011. doi: 10.1214/10-AOP579. → pages 30, 246