Kakeya-type Sets, Lacunarity, and Directional Maximal Operators in Euclidean Space

by

Edward Kroc

A thesis submitted in partial fulfilment 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)

April 2015

Abstract

Given a Cantor-type subset Ω of a smooth curve in $\mathbb{R}^{d+1}$, we construct random examples of Euclidean sets that contain unit line segments with directions from Ω and enjoy analytical features similar to those of traditional Kakeya sets of infinitesimal Lebesgue measure. We also develop a notion of finite order lacunarity for direction sets in $\mathbb{R}^{d+1}$, and use it to extend our construction to direction sets Ω that are sublacunary according to this definition. This generalizes to higher dimensions a pair of planar results due to Bateman and Katz [4], [3]. In particular, the existence of such sets implies that the directional maximal operator associated with the direction set Ω is unbounded on $L^p(\mathbb{R}^{d+1})$ for all $1 \leq p < \infty$.

Preface

Much of the proceeding document is adapted from two research papers authored by myself and Malabika Pramanik, currently unpublished. These materials are used with permission. Chapters 6 through 11 form the main content of [31], *Kakeya-type sets over Cantor sets of directions in* $\mathbb{R}^{d+1}$, while Chapters 2, 3.7, and 12 through 19 are adapted from [32], *Lacunarity, Kakeya-type sets and directional maximal operators.* The first of these two manuscripts has recently been conditionally accepted for publication in the Journal of Fourier Analysis and Applications.

Table of Contents

List of Figures

Acknowledgements

Above all others, I would like to thank my supervisor Malabika Pramanik. This work would not have been remotely possible without her unwavering support and tireless commitment. Her advice, her guidance, and her enthusiasm have been invaluable. I am grateful for the financial support she has granted me throughout my graduate studies at UBC, support that has allowed me to fulfill a personal dream of studying mathematics professionally, while simultaneously pursuing my myriad other interests and ambitions. Malabika has generously supported me in all my scientific and professional endeavours, regardless of their bearing on my status as a mathematician. She has guided me through the sometimes existential perils of life in academia, providing truly life changing direction and support. She has been there for me at each and every step of my academic career over the past six years, always happy to help in any way possible. I will be forever grateful for her kindness, her assistance, and her wisdom. I consider myself immeasurably lucky to have had the chance to work with and learn from her. She is my mentor, my colleague, and my friend.

I would also like to thank Paul Gustafson, Derek Bingham, Tim McMurry, and Ahmed Zayed. They have each provided unique and invaluable support at various stages in my career as a graduate student. They are significantly responsible for shaping the type of scientist I am today, and have provided much of the model for the type I would like to become. Paul and Derek in particular have helped guide me through some of the roughest parts of this process, providing abundant encouragement and inspiration when these were most needed. I will be forever indebted to them for their generosity.

I want to thank the members of my committee, Izabella Łaba, Ozgur Yilmaz, and Akos Magyar, for their advice, their support, and their time.

Beyond vocational or academic matters, I want to thank the Mathematics department administration and staff. No other department on campus has such exceptional staff who work so hard at making their department function smoothly. They are consummate professionals and wonderful people. In particular, I want to thank Lee Yupitun for her friendship and for all her help navigating the graduate program at UBC.

I would be extremely remiss not to acknowledge the unshakable support of my partner, my parents, and my friends throughout the drafting of this document, as well as throughout my entire graduate student career. I would never have pursued the dream of a Ph.D., especially not in such a wonderful place as Vancouver, British Columbia, without the encouragement of my partner, Marybeth Welty. She has always kept me afloat in the worst of storms, and has made all the best parts of this process orders of magnitude better. My parents have helped me financially, mentally, and emotionally throughout this process, and have always been happy to provide a temporary but comfortable escape from it. For that, I am especially grateful. My best and oldest friends Fred Drueck and Darin Hatcher remain incomparable in their ability to keep me grounded and thankful for what I have.

Finally, I want to thank my graduate colleagues in and outside the Mathematics department at UBC, in particular those who spend most of their time boxed up in the Auditorium Annex A offices. The open and easy atmosphere they create makes a day at the office feel like an escape amongst friends.

Chapter 1

Introduction

The main focus of this document is a study of Kakeya-type sets in Euclidean space of general dimensions. Such sets have much in common with the now classic Besicovitch construction of Kakeya sets with zero Lebesgue measure [6]. Interest in the analytical properties of these sets arises in both the traditional study of Kakeya sets in Euclidean space, as well as from a well-known connection between the existence of such sets and the boundedness properties of various maximal operators over L^p-space, notably the generalized one-dimensional Hardy-Littlewood maximal operator; see (1.8) and (1.2).

Definition 1.1. *Fix a set of directions $\Omega \subseteq \mathbb{S}^d$. We say a cylindrical tube is oriented in direction $\omega \in \Omega$ if the principal axis of the cylinder is parallel to ω. If for some fixed constant $A_0 \geq 1$ and any choice of integer $N \geq 1$, there exist*

- *a number $0 < \delta_N \ll 1$, $\delta_N \searrow 0$ as $N \nearrow \infty$, and*
- *a collection of tubes $\{P_t^{(N)}\}$ with orientations in Ω, length at least 1 and cross-sectional radius at most δ_N*

obeying

$$\lim_{N \to \infty} \frac{|E_N^*(A_0)|}{|E_N|} = \infty, \qquad \textit{with} \quad E_N := \bigcup_t P_t^{(N)}, \quad E_N^*(A_0) := \bigcup_t A_0 P_t^{(N)}, \tag{1.1}$$

then we say that Ω admits Kakeya-type sets. *Here,* $|\cdot|$ *denotes* $(d+1)$*-dimensional Lebesgue measure, and* $A_0 P_t^{(N)}$ *denotes the tube with the same centre, orientation and cross-sectional radius as* $P_t^{(N)}$ *but* A_0 *times its length. The tubes that constitute* E_N *may have variable dimensions subject to the restrictions mentioned above. We refer to* $\{E_N : N \geq 1\}$ *as sets of* Kakeya type*.*

We will explore the motivations and the history behind this definition in Section 1.5, but for now it suffices to reiterate that sets of Kakeya type, according to this Definition 1.1, share a critical feature with known constructions of zero measure Kakeya sets. This feature is intimately related to the notion of *stickiness*, arising from the study of classical Kakeya sets (see e.g. [46], [27]). We will describe these connections and discuss their history in Sections 1.3–1.5, but roughly speaking stickiness means the following: the tubes $\{P_t^{(N)}\}$ have considerable overlap, while their dilates $\{A_0 P_t^{(N)}\}$ are comparatively disjoint.

1.1 Summary of results

In brief, the new results contained in this document generalize the work of Bateman and Katz [4] and of Bateman [3] in the plane to an arbitrary number of dimensions. Our main body of work is devoted to establishing the existence of Kakeya-type sets, as defined in Definition 1.1, for certain direction sets $\Omega \subseteq \mathbb{S}^d$. These results will have immediate consequences for the L^p-boundedness of two often studied maximal operators (see (1.2) and (1.3) below). Indeed, as we will explain in Section 1.4, if $\Omega \subset \mathbb{S}^d$ admits Kakeya-type sets, then both of these operators are unbounded on $L^p(\mathbb{R}^{d+1})$ for all $p \in [1, \infty)$.

Given a set of directions Ω, the *directional maximal operator* D_Ω is defined by

$$D_\Omega f(x) := \sup_{\omega \in \Omega} \sup_{h>0} \frac{1}{2h} \int_{-h}^{h} |f(x + \omega t)| dt, \tag{1.2}$$

where $f : \mathbb{R}^{d+1} \to \mathbb{C}$ is a function that is locally integrable along lines. Also, for any locally integrable function f on $\mathbb{R}^{d+1}$, we consider the *Kakeya-Nikodym maximal* *operator* M_Ω defined by

$$M_\Omega f(x) := \sup_{\omega\in\Omega} \sup_{\substack{P\ni x \\ P\|\omega}} \frac{1}{|P|}\int_P |f(y)|dy, \tag{1.3}$$

where the inner supremum is taken over all cylindrical tubes P containing the point x, oriented in the direction ω. The tubes are taken to be of arbitrary length l and have circular cross-section of arbitrary radius r, with $r \le l$.

Bateman and Katz [4] show that $\Omega \subset \mathbb{S}^1$ admits Kakeya-type sets when $\Omega = \{(\cos\theta, \sin\theta) : \theta \in \mathcal{C}_{1/3}\}$ and $\mathcal{C}_{1/3}$ is the ordinary middle-third Cantor set on $[0, 1]$. We extend this result to a general $(d+1)$-dimensional setting, using the following notion of a Cantor set of directions in $(d+1)$ dimensions.

Fix some integer $M \ge 3$. Construct an arbitrary Cantor-type subset of $[0, 1)$ as follows.

- Partition $[0, 1]$ into M subintervals of the form $[a, b]$, all of equal length M^{-1}. Among these M subintervals, choose any two that are not adjacent (i.e., do not share a common endpoint); define $\mathcal{C}_M^{[1]}$ to be the union of these chosen subintervals, called first stage basic intervals.

- Partition each first stage basic interval into M further (second stage) subintervals of the form $[a, b]$, all of equal length M^{-2}. Choose two non-adjacent second stage subintervals from each first stage basic one, and define $\mathcal{C}_M^{[2]}$ to be the union of the four chosen second stage (basic) intervals.

- Repeat this procedure *ad infinitum*, obtaining a nested, non-increasing sequence of sets. Denote the limiting set by $\mathcal{C}_M$:

$$\mathcal{C}_M = \bigcap_{k=1}^{\infty} \mathcal{C}_M^{[k]}.$$

We call $\mathcal{C}_M$ a *generalized Cantor-type set* (with base M).

While conventional uniform Cantor sets, such as the Cantor middle-third set, are special cases of generalized Cantor-type sets, the latter may not in general look like the former. In particular, sets of the form $\mathcal{C}_M$ need not be self-similar. It is well-known (see [16, Chapter 4]) that such sets have Hausdorff dimension at most $\log 2/\log M$. By choosing M large enough, we can thus construct generalized Cantor-type sets of arbitrarily small dimension.

In Chapters 8–11, we prove the following [31].

Theorem 1.2. *(Kroc, Pramanik) Let $\mathcal{C}_M \subset [0,1]$ be a generalized Cantor-type set described above. Let $\gamma : [0,1] \to \{1\} \times [-1,1]^d$ be an injective map that satisfies a bi-Lipschitz condition*

$$\forall\ x, y,\ \ c|x-y| \leq |\gamma(x) - \gamma(y)| \leq C|x-y|, \tag{1.4}$$

for some absolute constants $0 < c < 1 < C < \infty$. Set $\Omega = \{\gamma(t) : t \in \mathcal{C}_M\}$. Then

(i) the set Ω admits Kakeya-type sets;

(ii) the operators D_Ω and M_Ω are unbounded on $L^p(\mathbb{R}^{d+1})$ for all $1 \leq p < \infty$.

Part (ii) of the theorem follows directly from part (i) as we will see shortly in Section 1.3. The condition in Theorem 1.2 that γ satisfies a bi-Lipschitz condition can be weakened, but it will help in establishing some relevant geometry. It is instructive to envision γ as a smooth curve on the plane $\{x_1 = 1\}$, and we recommend the reader does this to aid in visualization. Our underlying direction set of interest $\Omega = \gamma(\mathcal{C}_M)$ is essentially a Cantor-type subset of this curve.

After working through the details of Theorem 1.2, we generalize many of the ideas considerably to establish the following [32].

Theorem 1.3. *(Kroc, Pramanik) Let $d \geq 1$. If the direction set $\Omega \subseteq \mathbb{R}^{d+1}$ is sublacunary in the sense of Definition 2.7, then*

(i) Ω admits Kakeya-type sets;

(ii) the operators D_Ω and M_Ω are unbounded on $L^p(\mathbb{R}^{d+1})$ for all $1 \leq p < \infty$.

Again, part (ii) of the theorem follows from part (i) by the same mechanism to be described in Section 1.3. Precise definitions of *lacunarity* and *sublacunarity* needed in this document are deferred to Chapter 2, but the general idea is easy to describe. In one dimension, a relatively compact set $\{a_i\}$ is lacunary of order 1 if there is a point $a \in \mathbb{R}$ and some positive $\lambda < 1$ such that $|a_{i+1} - a| \leq \lambda |a_i - a|$ for all i. Such a set has traditionally been referred to as a lacunary sequence with lacunarity constant (at most) λ. A lacunary set of order 2 consists of a single (first-level) lacunary sequence $\{a_i\}$, along with a collection of disjoint (second-level) lacunary sequences; a second-level sequence is squeezed between two adjacent elements of $\{a_i\}$. The lacunarity constants of all sequences are uniformly bounded by some positive $\lambda < 1$. See Figure 1.1 for an illustration of a lacunary set of directions in the plane of order 2.

Roughly speaking, a set on the real line is lacunary of finite order if there is a decomposition of the real line by points of a lacunary sequence such that the restriction of the set to each of the resulting subintervals is lacunary of lower order. All lacunarity constants implicit in the definition are assumed to be uniformly bounded away from unity. A set is then said to be *sublacunary* if it does not admit a finite covering by lacunary sets of finite order. A Cantor-like set is a particular kind of sublacunary set according to these definitions.

As noted before, we will devise a formal definition of sublacunarity in Chapter 2. However, for analytical purposes, we will see that it is in fact more convenient to recast the lacunarity (or lack therof) of a set in terms of its encoded *tree structure*. This idea was first formally noted in the work of Bateman and Katz [4], and subsequently played a role in the later work of Bateman [3]. We will discuss the tree structure of sets in Chapter 3, and see that the inherent tree structure of a sublacunary (in particular, Cantor-like) subset of Euclidean space is, in a quantifiable sense, as full as that of any nonzero Lebesgue measure Euclidean set; see Propositions 3.2 and 3.6. In the meantime, it is recommended that the reader keep in mind that sublacunary sets are, in a rather fundamental sense to be explored in Section 3.6, fully generalized Cantor-like subsets of Euclidean space.

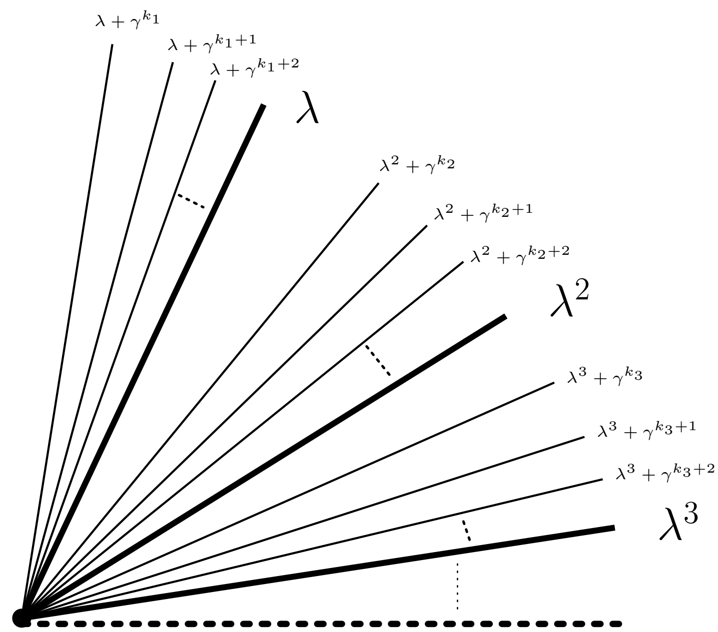

Figure 1.1: A direction set in the plane, represented as a collection of unit vectors, with parameters $0 < \gamma < \lambda < 1/2$. The set of angles made by these vectors with the positive horizontal axis is $\{(\lambda^j + \gamma^k) : k \geq j\}$, which is lacunary of order 2.

1.2 Notations, conventions, and structure of the document

Unless otherwise stated, for $A \subseteq \mathbb{R}^d$, with d minimal, the notation $|A|$ will always be used to denote the d-dimensional Lebesgue measure of A. For a countable, possibly infinite set Q, the notation $\#(Q)$ will denote cardinality of the set Q. If X_N and Y_N are quantities depending on N, then we will write $X_N \lesssim Y_N$ to mean that there exists a constant C independent of N such that $X_N \leq CY_N$. If both $X_N \lesssim Y_N$ and $Y_N \lesssim X_N$ hold, then we will write $X_N \sim Y_N$.

As in Definition 1.1, if P is a tube with some arbitrary centre, orientation, cross-sectional radius, and length l, and if C_0 is a positive constant, then C_0P will denote the tube with the same centre, orientation, and cross-sectional radius as P, but with length C_0l.

This document is divided into twenty main chapters, and each chapter falls in one of four expository groups, each with a single, broad goal. Chapters 1–5 set definitions and discuss the necessary background, Chapters 6–11 form the proof of Theorem 1.2, and Chapters 12–19 deal with the proof of Theorem 1.3; the final Chapter 20 discuss potential future work that could arise from this document. The flow charts below diagram the logical ordering of the individual chapters contained within this document.

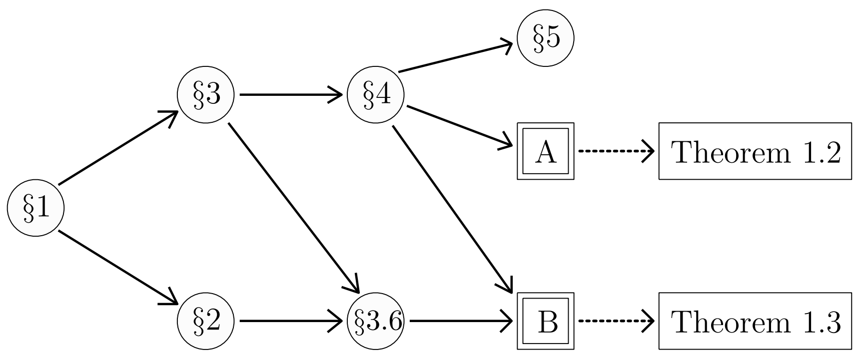

Figure 1.2: Diagram illustrating the approximate dependence structure between chapters in this paper, with respect to the proofs of Theorems 1.2 and 1.3. The chapter structure of groups A and B are detailed in the two maps below.

Chapter 1 is introductory and discusses the history of the problems considered in this document, as well as the motivation for their study. Chapter 2 discusses the notion of lacunarity. We lay out precise definitions, check for consistency with the established literature, and are able to properly state Theorem 1.3 with respect to these definitions. This chapter does not directly pertain to the proof of Theorem 1.2 and can be skipped until the reader begins the chapters of group B.

Chapter 3 introduces the critical idea of a tree encoding a set in Euclidean space, and Section 3.7 discusses the idea of lacunarity on trees (not needed until the chapters of group B). The so-called *splitting number* of a tree, as defined in [3], is then shown to be the critical concept that allows us to recast the notion of (admissible) finite order lacunarity of a set into an equivalent and more tractable form for the purposes of our proof. Chapter 4 reviews the relevant literature pertaining to percolation on trees. Chapter 5 describes the known results about Kakeya-type sets in the plane,

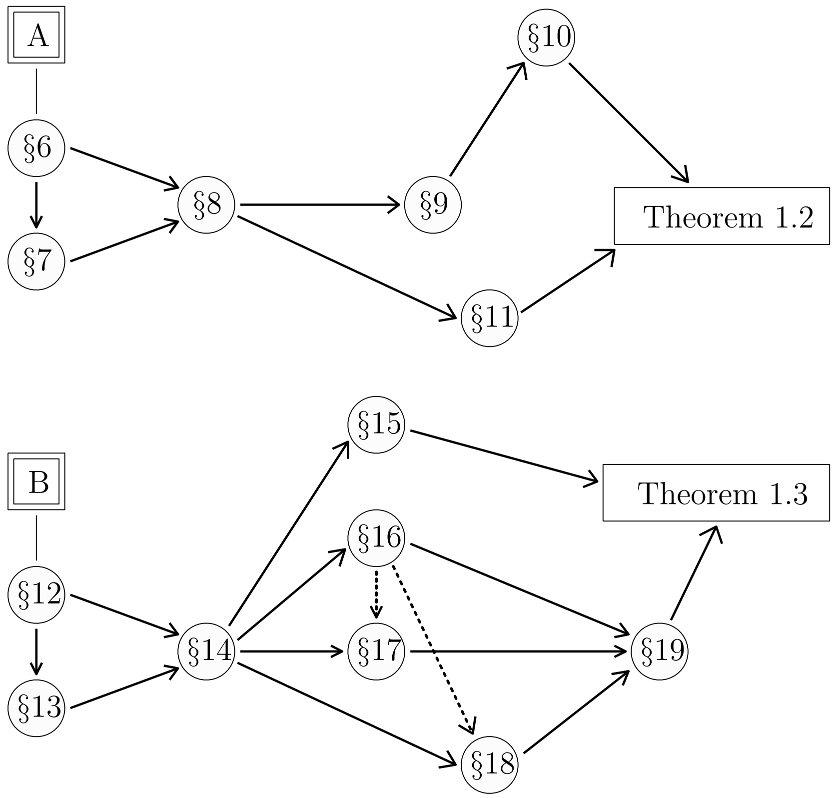

Figure 1.3: Diagram illustrating the approximate dependence structure between chapters in groups A and B of Figure 1.2. The chapters in group A detail the proof of Theorem 1.2, while the chapters in group B constitute the proof of Theorem 1.3. Dotted arrows indicate a dependence in terms of definitions and notation only.

and presents the proof of Bateman and Katz [4] in detail. The material in this chapter is not directly required in subsequent chapters and can be skipped if desired.

Chapters 6–11 form the proof of Theorem 1.2. Chapter 6 sets up the construction of Kakeya-type sets over a Cantor set of directions, and presents a probabilistic version of our main theorems. Chapter 7 explores the relevant geometry of the intersection of two tubes in Euclidean space, recording facts for later use. Chapter 8 combines results of the previous two chapters to explicitly describe the probabilistic mechanism in play. We also reformulate Theorem 1.2 in terms of quantitative upper and lower bounds on the sizes of a typical Kakeya-type set E_N and its principal dilate $E_N^*(A_0)$ as described in Definition 1.1 (see Proposition 8.4). From here, the remaining chapters in group A split into more or less two disjoint expositions, each one charged with establishing one of these two probabilistic and quantitative bounds.

Chapters 9 and 10 combine to establish the quantitative lower bound. Chapter 9 in particular explores the analytical implications of the structure imposed on the position and slope trees of a collection of two, three, or four δ-tubes, certain pairs of which are required to intersect at a given location in space. In Chapter 11 we prove the quantitative upper bound using an argument similar to [4].

Chapter 12 begins the program of proving Theorem 1.3; i.e., of constructing Kakeya-type sets over sublacunary direction sets. The structure of the chapters of group B mirrors that of group A just described.

We use the language of trees developed in Chapter 3 to extract a convenient subset of an arbitrary sublacunary direction set, denoted by Ω_N. Chapter 13 expands on the geometry of the intersection of two tubes as initiated in Chapter 7 and the implications of this geometry for the structure of trees encoding the sets of orientations and positions of a given collection of thin δ-tubes.

Chapter 14 combines results from the previous two chapters to describe the actual mechanism we use to assign slopes in Ω_N to δ-tubes affixed to a prescribed set of points in Euclidean space. Again, we reformulate Theorem 1.3 in terms of quantitative upper and lower bounds on the sizes of a typical Kakeya-type set E_N and its principal dilate $E_N^*(A_0)$ (see Proposition 14.2).

In Chapter 15 we prove the quantitative upper bound previously prescribed. Chapters 16–19 combine to establish the corresponding lower bound. Chapter 19 details the actual estimation, utilizing all the smaller pieces developed in Chapters 16–18. These three chapters revolve around a central theme of ideas, notably the structure imposed on the position and slope trees of a collection of two, three, or four δ-tubes, certain pairs of which are required to intersect at a given location in space.

1.3 Early history of Kakeya sets

A *Kakeya set* (also called a *Besicovitch set*) in $\mathbb{R}^d$ is a set that contains a unit line segment in every direction. The study of such sets spans nearly one hundred years, originating independently from the works of Abram Besicovitch and Sōichi Kakeya. Besicovitch was originally interested in the question of whether or not for any given Riemann integrable function in $\mathbb{R}^2$, an orthogonal coordinate system always exists such that the two-dimensional integral is equal to the iterated integral, which must thus be well-defined. Kakeya meanwhile was concerned with determining the smallest area of a planar region in which a unit line segment (a "needle") could be continuously rotated through 180°, the so-called *Kakeya needle problem*. Surprisingly, the answers to these questions were not so disparate and provided the genesis of a long line of inquiry that continues to this day.

Besicovitch realized that his question could be answered (in the negative) if it was possible to construct a zero Lebesgue measure planar set containing a line segment in every direction [33]. Indeed, suppose such a set E exists and fix a coordinate system in $\mathbb{R}^2$. Define a function f so that $f(x, y) = 1$ if $(x, y) \in E$ and if at least one of x or y is rational, and $f(x, y) = 0$ otherwise. The function f is clearly Riemann-integrable in two dimensions since its points of discontinuity comprise a set of two-dimensional Lebesgue measure zero. However, after possibly translating E so that the x- and y-coordinates of the line segments in E parallel to the fixed coordinate axes are irrational, we see that f is not Riemann-integrable as a function of one variable regardless of the direction along which we choose to integrate. In fact, by the structure of E, we are guaranteed that for any direction in the plane, there is a cross-section of the function f that behaves like the characteristic function of the rationals along the real line.

The main body of Besicovitch's work in [5] was to actually construct the required set E. The idea is summarized in Figure 1.4 below; in words, the idea is as follows. Begin with an equilateral triangle ABC and notice that it contains line segments of all slopes between those of AB and AC. Now, fix some large integer n and subdivide the base BC into 2^n equispaced pieces. Cut ABC into 2^n thin, tall triangles with one vertex at A and the others at subsequent points resulting from the previous subdivision of BC. Slide these triangles along their bases to create a new figure with smaller area, yet that still contains line segments with all the slopes of the original ABC triangle.

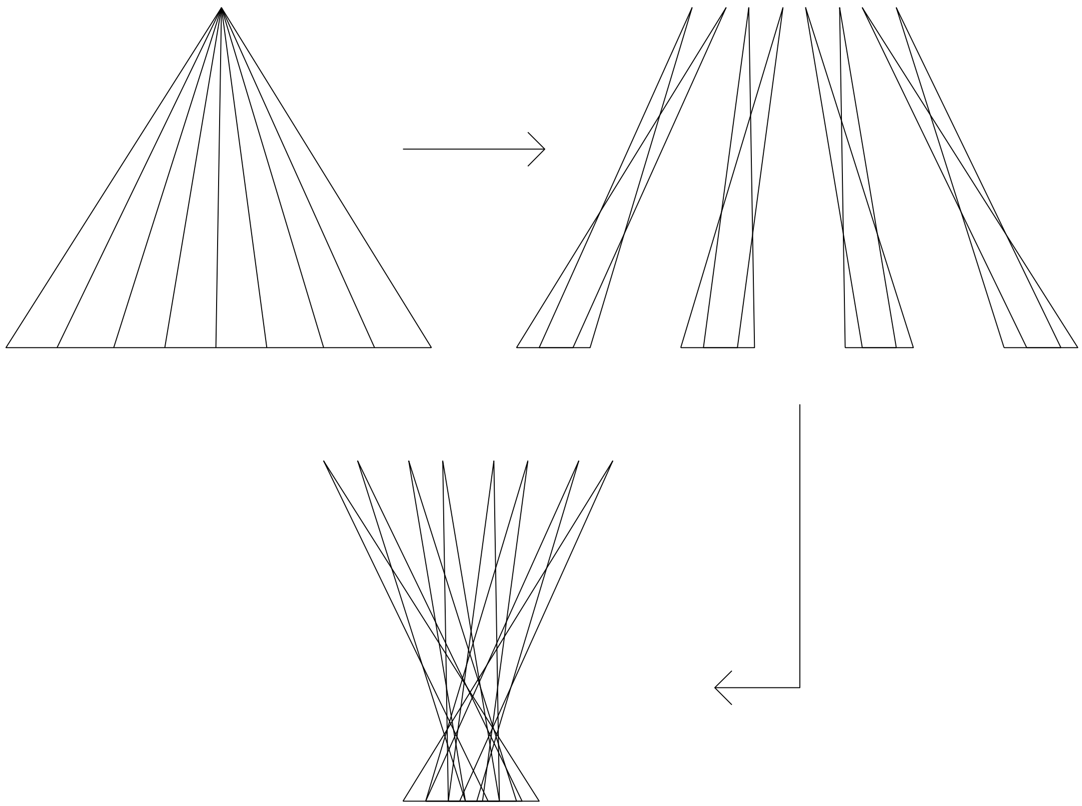

Figure 1.4: Diagram of the first iteration of a Besicovitch-style construction of a zero measure Kakeya set in the plane.

Besicovitch optimized the number of cuts and the resultant reconfiguration process to show that for any $\varepsilon > 0$, this procedure can produce a set with area less than ε, while retaining all the requisite slopes from the original ABC triangle. Iterating the construction and taking the limit, Besicovitch produced a set with zero (planar) Lebesgue measure. Applying this construction individually to the union of six appropriately rotated copies of the original equilateral triangle results in a zero

Lebesgue measure set containing a line segment in *every* direction.

Once such a construction is complete in the plane, it is easy to produce a zero Lebesgue measure set containing a line segment in every direction in $\mathbb{R}^d$, for any $d \geq 3$. Indeed, if E denotes the zero measure planar set constructed previously, then simply considering the set $E \times [0,1]^{d-2}$ yields a valid example. Observe that such a set is compact, and in fact always containable in a ball centred at the origin of radius, say, 2.

Besicovitch's original construction has been simplified over the years, and the typical construction is now quite streamlined (see Stein [42] for example). Several alternate constructions have also been developed, most notably one by Kahane [23] which utilizes Cantor sets in the plane, and another by Besicovitch himself [7].

Besicovitch did not connect the construction of his set to the Kakeya needle problem until nearly a decade after his initial paper on the subject appeared. In fact, he was unaware of the problem due to the 1917 civil war taking place around him in Russia at the time, which effectively silenced all communication between scientists inside the country and the rest of the scientific community [33]. In this meantime, Kakeya [24] and Fujiwara and Kakeya [18] worked on the needle problem in Japan without knowledge of Besicovitch's work. They conjectured that the smallest convex planar set within which a unit line segment could be rotated continuously through $180°$ was the equilateral triangle of height 1; they also realized that one could do better by removing the convexity assumption. Their conjecture was soon verified by Pál [38].

In fact, it was Pál who pointed Besicovitch to Kakeya's needle problem and who also noted how Besicovitch's original construction could be modified to solve it [33]. This lead to Besicovitch publishing the surprising solution in 1928 [6]:

- *For any $\varepsilon > 0$, there is a planar region of area less than ε within which a unit line segment can be rotated through a full $180°$.*

To construct such a set given an $\varepsilon > 0$, simply iterate Besicovitch's original construction enough times so that the resulting set E has area less than $\varepsilon/2$. Notice that any tall, thin translated triangle at any iteration of the construction has at least one side parallel to another "partner" triangle from the same stage of the construction (refer to Figure 1.4). Denote these two triangles by $A_1B_1C_1$ and $A_2B_2C_2$, as in Figure 1.5. Since A_1B_1 and A_2B_2 are parallel, we can take two radii D_1A_1 and D_2A_2 which form an arbitrarily small angle. Where these radii intersect, we take a sector OR_1R_2 of radius 1. Then we can continuously move a unit line segment from inside the triangle $A_1B_1C_1$, along the line D_1O, through the sector OR_1R_2, along the line D_2O, and through the triangle $A_2B_2C_2$. Choosing a small enough angle for the two radii D_1O and D_2O, we can add enough sectors to the original Besicovitch construction to ensure continuous movement of a needle throughout the full range of slopes, while retaining a total area of no more than ε.

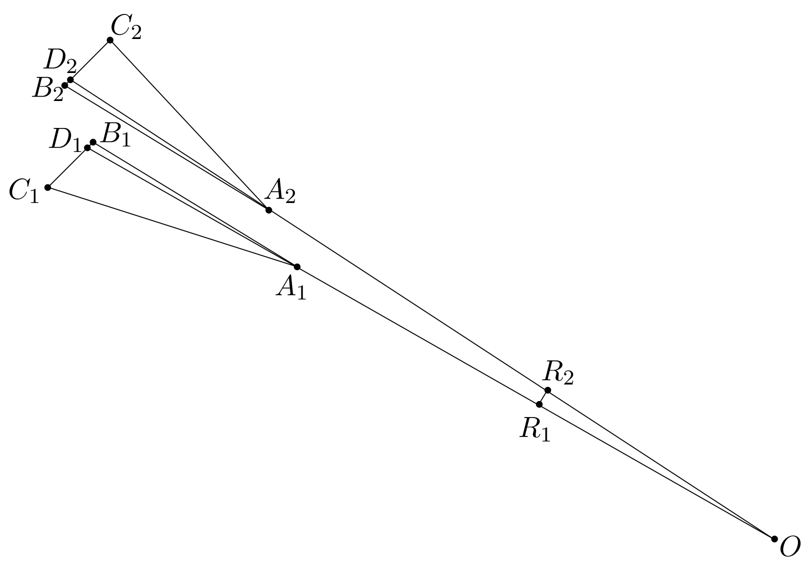

Figure 1.5: The needle sliding component of Besicovitch's solution to the Kakeya needle problem [6], as noted by Pál.

After the publication of Besicovitch's 1928 paper, the study of Kakeya sets focused on their geometry and on the construction of similar pathological objects; for example, sets of measure zero containing translates of all circles were constructed by Besicovitch and Rado [8], and independently by Kinney [29]. The subject was developed over the next forty-odd years by these and other mathematicians, notably

Perron, Ward, and Marstrand.

Then in 1971, two seminal results appeared almost concurrently. One, due to Cunningham [13], provided the most economic Kakeya-type construction and refined Besicovitch's answer to the original Kakeya needle problem. He showed that it was possible to rotate a unit line segment 180° inside a simply connected subset of the unit circle of arbitrarily small measure. This provided a remarkable strengthening of Besicovitch's solution to the Kakeya needle problem, as Cunningham's construction did not require the use of the needle sliding components depicted in Figure 1.5, which result in sets of increasingly large diameter. The other result, due to Davies [14], showed that any Kakeya set in the plane must have full Hausdorff dimension, a natural conjecture that had existed in the research community for some time. In higher dimensions, the analogous conjecture remains open to this day. This well-known conjecture is still most often dubbed *the Kakeya conjecture*.

Davies' proof of the Kakeya conjecture in the plane relied strictly on geometric considerations. Soon though, it was realized that this conjecture and related problems could be attacked via more analytical means, most notably by proving certain L^p estimates on suitably defined maximal operators. Inspired largely by the work of Fefferman, notably [17], Córdoba reproved Davies' result using just such a method [11], deriving an $L^2(\mathbb{R}^2)$ estimate on an appropriate maximal operator. His result hinged critically on the following geometric property (see [11] page 7, the proof of Proposition 1.2).

Fact 1.4 (Córdoba)**.** *Let $T_e^\delta(a)$ denote the tube of length 1 and cross-sectional radius δ, centred at a and oriented parallel to e. For any pair of directions $e_k, e_l \in \mathbb{S}^{d-1}$, and any pair of points $a, b \in \mathbb{R}^d$, we have the estimates*

$$diam(T_{e_k}^\delta(a) \cap T_{e_l}^\delta(b)) \lesssim \frac{\delta}{|e_k - e_l|},$$

and

$$|T_{e_k}^\delta(a) \cap T_{e_l}^\delta(b)| \lesssim \frac{\delta^n}{|e_k - e_l|}. \tag{1.5}$$

This property has seen repeated application in the subsequent literature (see [9],

[46], [4], [3], for example), and has since become known colloquially as the *Córdoba estimate*.

The maximal operator Córdoba considered was equivalent to the two-dimensional version of the following operator:

$$M_\delta f(x) = \sup_{\substack{x \in P \\ P \in \mathcal{P}_\delta}} \frac{1}{|P|} \int_P |f(y)| dy, \tag{1.6}$$

where f is locally integrable, and $\mathcal{P}_\delta$ is the collection of all tubes in $\mathbb{R}^d$ with length 1 and cross-sectional radius δ. This function has gone by many different names in the literature over the years, including both the Kakeya maximal function and the Nikodym maximal function. In the interest of compromise and historical accuracy, this function will be referred to as the *restricted Kakeya-Nikodym maximal function* in the remainder. It is instructive to compare this function with the more general form of the Kakeya-Nikodym maximal function introduced in (1.3).

The study of these kinds of maximal functions, where one considers the average of an arbitrary, locally-integrable function f over a certain set of geometric objects, dates back at least to the now classical *Lebesgue differentiation theorem*:

Theorem 1.5. *(Lebesgue) If $f : \mathbb{R}^d \to \mathbb{R}$ is integrable, then for almost every x we have*

$$\lim_{r \to 0} \frac{1}{|B(x,r)|} \int_{B(x,r)} f(y) dy = f(x), \tag{1.7}$$

where $B(x,r)$ denotes the ball of radius r centred at $x \in \mathbb{R}^d$.

This result is of paramount significance, since it is both analogous to and a generalization of the Fundamental Theorem of Calculus. Closely related to this, the idea of studying a *maximal* average of locally integrable functions over a certain set of geometric objects dates back at least to the canonical work of Hardy and Littlewood [21], near the time of Besicovitch and Kakeya's original works previously discussed. Their eponymous operator, the *Hardy-Littlewood maximal function*, is defined as

$$M_{HL} f(x) = \sup_{r > 0} \frac{1}{|B(x,r)|} \int_{B(x,r)} |f(y)| dy. \tag{1.8}$$

Hardy and Littlewood's inaugural result established that this operator was bounded on $L^p(\mathbb{R})$ for all $1 < p \leq \infty$, and was also of weak-type (1,1) on $\mathbb{R}$. This result was extended to any dimension by Wiener in a seminal 1939 work [45]; i.e. Wiener proved, among other things, that $||M_{HL}f||_p \leq C_{p,d}||f||_p$ for all $1 < p \leq \infty$, and that $|\{x \in \mathbb{R}^d : |M_{HL}f(x)| > \lambda\}| \leq C_d\lambda^{-1}||f||_1$. These estimates can be used to deduce the Lebesgue differentation theorem with little additional work. The argument justifying this transition is quite robust; with minor modifications, L^p estimates on a suitably defined maximal operator can lead to differentiation theorems over a prescribed set of geometric objects.

Indeed, the balls in (1.7) and (1.8) can be replaced by cubes, by parallelopipeds with sides parallel to the coordinate axes, or any other suitably symmetric class of objects without changing the L^p-boundedness properties of the analogous maximal functions. The existence of zero measure Kakeya sets however implies that no such bounds can hold for the Kakeya-Nikodym maximal function defined in (1.3) when Ω is a set of directions with nonempty interior, for $p < \infty$. Moreover, the existence of zero measure Kakeya sets also implies that no differentiation theorem analogous to Theorem 1.5 can hold over the collection of cylindrical tubes in $\mathbb{R}^d$ with arbitrary orientations and cross-sectional radius smaller than their given length L.

To see this, let $0 < \varepsilon < L$, and let E_ε be an ε-neighbourhood of a zero measure Kakeya set $E \subset \mathbb{R}^d$ constructed as the limit of sets in Besicovitch's procedure. This type of Kakeya set exhibits a property called *stickiness*, and we say that the set E is *sticky*. We will return to this concept in detail later, but for now we attempt to give just an impression of the idea.

Observe that if l is any unit line segment in E, say oriented along the direction ω, then if we extend l along its direction ω by at least 4 units, then this extended segment $\widetilde{l}$ will fall outside the ball of radius 2 centred at the origin (recall that E could be constructed to be entirely contained within such a ball). Moreover, if l and l' are two different unit line segments in E with different orientations ω and ω', then their analogous extensions $\widetilde{l}$ and $\widetilde{l'}$ along these directions must be disjoint. Denote by $\widetilde{E}$ the extended Kakeya set created by applying this extension procedure to every unit line segment $l \subset E$, and let $\widetilde{E}_\varepsilon$ denote its ε-neighbourhood.

Now, for any $x \in \widetilde{E}_\varepsilon$, note that

$$M_\Omega \mathbf{1}_{E_\varepsilon}(x) \geq \frac{1}{8\varepsilon^{d-1}} \int_{-4}^{4} \int_{(-\varepsilon,\varepsilon)^{d-1}} \mathbf{1}_{E_\varepsilon}(x + t\omega + s\omega^\perp) ds dt \gtrsim \frac{1}{8},$$

where the implicit constant is allowed to depend on L. Therefore,

$$||M_\Omega||_{p\to p} \gtrsim \frac{||\mathbf{1}_{\widetilde{E}_\varepsilon}||_p}{8||\mathbf{1}_{E_\varepsilon}||_p} \gtrsim \left(\frac{\left|\widetilde{E}_\varepsilon\right|}{|E_\varepsilon|} \right)^{\frac{1}{p}}, \tag{1.9}$$

and this holds for any $0 < \varepsilon < L$. Now, Besicovitch's construction yields $|E_\varepsilon| < \varepsilon$, while at the same time $\left|\widetilde{E}_\varepsilon\right| \geq c_0$ for some $c_0 > 0$ independent of ε. Thus, if $p < \infty$ and $\Omega = \mathbb{S}^{d-1}$, the Kakeya-Nikodym operator M_Ω is unbounded on $L^p(\mathbb{R}^d)$ by (1.9).

A vast amount of work has been poured into finding δ-dependent L^p bounds for the restricted Kakeya-Nikodym maximal operator. Such estimates have provided myriad applications to all sorts of analytical problems; most notably, these bounds can have implications for the Kakeya conjecture, see [9], [46], [28], [27] for example. We will not explore this deep and exciting work here, but our work does focus on another set of problems that were born out of the study of Kakeya sets and the maximal operators in (1.3) and (1.6).

1.4 Background: maximal averages over lines with prescribed directions

Consider the restricted Kakeya-Nikodym maximal function, given by (1.6). What if we were to let $\delta \to 0$? Clearly, the limit would not make immediate sense, but notice what this operation tries to accomplish. Instead of averaging a function over a thin tube with length 1 and cross-sectional radius δ, we would be wanting to consider averages simply over line segments of length 1. Our object of study is then a "lower dimensional" maximal function, the so-called *directional maximal operator*, defined in (1.2).

The directional maximal function is a natural generalization of the classic Hardy-Littlewood function. In fact, when $d = 1$, the directional maximal function reduces to (1.8). The boundedness of D_Ω for finite Ω, in any dimension, follows directly from the boundedness of the Hardy-Littlewood maximal operator on L^p, $1 < p \leq \infty$. Considerable work has been done on the size of the corresponding norms when Ω is a finite but otherwise arbitrary set, especially when $d = 2$, $p = 2$, most notably in [11], [44], [15], [25], [26]. In the mid and late 1990's, Katz showed [26] that D_Ω is bounded on $L^2(\mathbb{R}^2)$ with the bound $\lesssim \sqrt{1 + \log N}$, where N is the cardinality of Ω, while at the same time the lower bound $\gtrsim \sqrt{\log \log N}$ holds [25], where $N = 3^n$ and Ω is the nth iterate of the Cantor set on $[0, 1]$. These results establish that L^p-bounds independent of the cardinality of Ω cannot be expected to hold in general. Furthermore, these bounds together imply the unboundedness of D_Ω as an operator on $L^2(\mathbb{R}^2)$ when Ω contains a Cantor set by the same argument employed immediately preceding and within (1.9). This result was later generalized by Hare [22] to include sets Ω containing any Cantor set of positive Hausdorff dimension. These results came as somewhat of a surprise, as it had already been conjectured that cardinality independent bounds on D_Ω should exist for Cantor sets [44]. This belief seems to have been motivated by the fact that such cardinality independent bounds do in fact exist for certain subsets of L^2-functions, notably positive radial functions with weak Fourier decay [44].

Additionally, two famous results establish that D_Ω is bounded as an operator on any $L^p(\mathbb{R}^d)$, $1 < p \leq \infty$, in arbitrary dimension d, for certain sets of directions Ω with infinite cardinality. As to be expected from our discussion in Section 1.1, these sets are required to exhibit a certain degree of *lacunarity*. The original results in this direction came from Córdoba and Fefferman [12], Stromberg [43], and the classic work of Nagel, Stein, and Wainger [37]. The latter authors considered lacunary sets of the form $\Omega = \{(\theta_j^{m_1}, \ldots, \theta_j^{m_d}) : j \geq 1\}$, where $0 < m_1 < \cdots < m_d$ are fixed constants and $\{\theta_j\}$ is a lacunary sequence with lacunarity constant $0 < \lambda < 1$; i.e., $0 < \theta_{j+1} \leq \lambda \theta_j$. For such sets Ω, they showed in their 1978 paper that D_Ω is bounded on all $L^p(\mathbb{R}^d)$, $1 < p \leq \infty$. In addition to the usual Littlewood-Paley theory, a main feature of their proof is the use of an almost-orthogonality principle that allows them to control the degree of overlap between a group of averages in Fourier space (see Lemma 1, [37]). Such almost-orthogonality results reappear as a crucial part of the analysis in all subsequent work on the subject of boundedness of these operators.

In 1988, Carbery [10] utilized similar methods, substantially generalized, to show that directional maximal operators arising from coordinate-wise lacunary sets of the form $\Omega = \{(r^{k_1}, \ldots, r^{k_d}) : k_1, \ldots, k_d \in \mathbb{Z}^+\}$ for some $0 < r < 1$ are bounded on all $L^p(\mathbb{R}^d)$, $1 < p \leq \infty$. This result paired with the one due to Nagel, Stein, and Wainger seemed to decide the question of L^p-boundedness of the directional maximal operator in $\mathbb{R}^d$ over sets Ω adhering to a certain intuitive notion of lacunarity. Extending these ideas to their limit however requires a more formal notion of generalized lacunarity, in addition to the appropriate almost-orthogonality results.

In the plane, this notion was sufficiently generalized by Sjögren and Sjölin [41] in 1981, establishing that so-called *lacunary sets of finite order* gave rise to bounded directional maximal operators on $L^p(\mathbb{R}^2)$-space, $1 < p \leq \infty$. On the real line, they defined a lacunary set of order 0 as a singleton, and a lacunary set of order $N \geq 1$ as a *successor* of a lacunary set of order $N-1$, where the successor of a set is defined as follows. For any closed sets $A, A' \subset \mathbb{R}$ of Lebesgue measure zero, A' is a successor of A if there exists a constant $c > 0$ such that $x, y \in A'$, with $x \neq y$, implies that $|x - y| \geq c \cdot \mathrm{dist}_A(x)$. Here, $\mathrm{dist}_A(x)$ denotes the usual distance between two sets. Using such a definition, sets of the form

$$\left\{ \sum_{k=1}^{N} \lambda_k^{i_k} : i_k \in \mathbb{Z},\ 1 \leq k \leq N \right\} \cup \{0\} \tag{1.10}$$

are lacunary of order N. Once this notion is defined on the real line, it may be lifted and applied directly to sets of directions in the plane, $\Omega \subseteq \mathbb{S}^1$. In this way, Sjögren and Sjölin were able to generalize the results of [37], as well as to prove that L^2-boundedness cannot be expected to hold for direction sets $\Omega \subset \mathbb{S}^1$ corresponding to a classical Cantor set.

In 2003, Alfonseca, Soria, and Vargas published an L^2-almost-orthogonality result for directional maximal operators in the plane [2], shortly thereafter generalized by

Alfonseca to any L^p-space [1]. Essentially, what was recognized was that if $\Omega_0 \subset \Omega \subseteq [0, \frac{\pi}{4})$ is an ordered subset of angles,

$$\frac{\pi}{4} = \theta_0 > \theta_1 > \theta_2 > \dots,$$

then the L^p-norm of the directional maximal operator over Ω is controlled by the L^p-norm of D_{Ω_0} and the supremum of the L^p-norms of D_{Ω_j}, where for each $j \geq 1$, one defines $\Omega_j = \Omega \cap [\theta_j, \theta_{j-1})$. These results substantially generalized the original almost-orthogonality principle proved by Nagel, Stein, and Wainger and played a critical role in the work of Bateman [3].

These almost-orthogonality results and the study of lacunarity were not the only foci of research on directional maximal operators. As previously mentioned, concurrent work had been taking place that studied the behaviour of the L^p-norms of D_Ω when Ω is given by a Cantor subset of directions [15], [44], [25], [26]. This research aimed to uncover the middle ground between lacunary directional maximal operators, and those arising from a set of directions Ω with nonempty interior. Indeed, due to the existence of zero measure Kakeya sets, and by the same argument that appears in and preceding inequality (1.9), it is an immediate consequence that all operators D_Ω arising from Ω with nonempty interior are unbounded on $L^p(\mathbb{R}^d)$, $p \neq \infty$.

In the plane, Bateman and Katz finally settled the issue of L^p-boundedness of directional maximal operators arising from a Cantor set of directions Ω for a generic $p \in [1, \infty)$ in their 2008 publication [4]. They showed that not only is such an operator D_Ω unbounded on $L^p(\mathbb{R}^2)$, $p \neq \infty$, but more strongly that such a set of directions Ω admits Kakeya-type sets. Precisely, they showed that for any $n \geq 1$, there is a union of n parallelograms in $\mathbb{R}^2$ of dimensions $1 \times \frac{1}{n}$, and with slopes of the longest sides contained in the standard middle-thirds Cantor set, so that

$$\left| \bigcup_{j=1}^{n} P_j \right| \lesssim \frac{1}{\log n}, \text{ while } \left| \bigcup_{j=1}^{n} 2P_j \right| \gtrsim \frac{\log \log n}{\log n}. \tag{1.11}$$

Notice that a set that obeys (1.11) is indeed of Kakeya-type according to our Definition 1.1.

Bateman and Katz's construction is probabilistic. In particular, they show that the first inequality in (1.11) holds for a typical collection of suitable parallelograms with respect to a certain uniformly distributed probability measure. This result relies on an ingenious application of a seminal theorem from the theory of percolation on tree structures due to Russ Lyons [34], [35]. Their work represents the first time such an idea has been applied to the study of Kakeya-type sets and related objects. The new results that appear in this document will rely critically on the same idea (see Chapters 11 and 15).

Shortly after [4] appeared, Bateman published a generalization of their result [3] that provided a complete characterization of the L^p-boundedness of directional maximal operators in the plane for an arbitrary set of directions $\Omega \subseteq \mathbb{S}^1$. The proof utilized the same techinques as in [4], with the addition of some necessary complexities to account for the totally general structure of the direction set Ω. Introduced originally perhaps as a convenient way to formally apply Lyons' percolation result, the inherent "tree structure" naturally associated to an arbitrary set of directions played a more vital role in Bateman's work. This type of encoding allows for a rich and flexible structure capable of accommodating both the generic nature of the problem and the specific combinatorics and analysis required to establish (1.11) in the general context.

1.5 Kakeya-type sets and the property of stickiness

If we consider once again Besicovitch's construction of a zero measure Kakeya set, described near the beginning of Section 1.3 and partially depicted in Figure 1.4, an interesting observation can be made; indeed, we have already appealed to this observation when proving the unboundedness of the Kakeya-Nikodym maximal operator over $\mathbb{S}^d$. Although the many thin triangles in that construction are translated as to overlap near a common core, by virtue of the disparate directions they each contain, these triangles are disjoint sufficiently far away from that core. Put more generically, if many thin tubes with distinct directions intersect at a common locus, then these tubes will become mutually disjoint outside a large enough ball surrounding this locus. The necessary size of such a ball is dependent on both the cross-sectional width of these tubes, and the amount of separation between their directions. A collection of tubes that exhibits this property are said to be *sticky*, referring to the idea that they should all "stick together" near a common core. Such a collection is pictured in Figure 1.6.

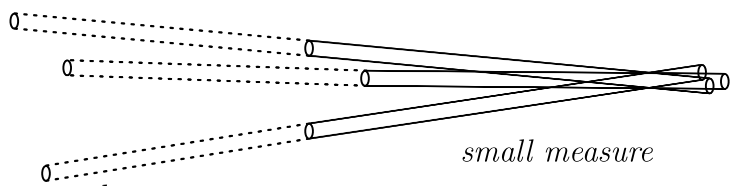

Figure 1.6: A sticky collection of tubes $\bigcup P_t$, along with their extensions $\bigcup A_0 P_t$ in $\mathbb{R}^3$.

Wolff was the first to formally recognize that zero measure Kakeya sets must enjoy this property; see Property $(*)$ in [46]. More specifically, if Kakeya sets are to have small measure, then a δ-thickening of the set must contain many pieces that look like the "small measure," right hand side of Figure 1.6; i.e., contain many tubes that overlap for much of their length. This can only be so if both the centres of these tubes and their orientations are nearly the same. In this case, as long as the tubes are given distinct orientations, an extension of these tubes by an appropriate amount must result in a set of relatively large measure, as in the left hand side of Figure 1.6. Indeed, if we again consider the Besicovitch construction in Figure 1.4, although the original triangle can be chosen to have, say, unit area, the resulting set after n iterations of the construction can be made to have area less than about $1/\log n$. However, extending each subdivided triangle along their respective diameters will result in a set with approximately unit area, independent of n.

The reader will recognize that the phenomenon we have been describing has already been laid down in quantitative terms in Definition 1.1, and exploited analytically to show why the existence of Kakeya-type sets implies unboundedness of various associated maximal operators, as in the argument leading to (1.9). Indeeed, the condition (1.1) is a generic formulation of the idea of stickiness for a family of tubes with orientations arising from a given set of directions Ω. We see then that the idea of a Kakeya-type set is motivated by the property of stickiness exhibited by traditional Kakeya sets of zero measure; Kakeya-type sets in fact preserve this property as their defining feature.

The idea of stickiness is not restricted to collections of tubes. Although similar notions can be defined for essentially any collection of common, thickened lower dimensional objects in Euclidean space, sticky collections of curves, circles, and spheres have been of particular interest [8], [29], [40], [30]. We will discuss further the idea of stickiness and its interplay with work developed in this document in Chapter 20.

Chapter 2

Finite order lacunarity

As noted in Section 1.4, the concept of finite order lacunarity plays a fundamental role in the study of planar Kakeya-type sets and associated directional maximal operators. The results of Nagel, Stein, Wainger [37], and Carbery [10] suggest that it continues to play a similar role in higher dimensional space. The existing literature on the subject embodies several different notions of Euclidean lacunarity both in single and general dimensions, see in particular [3, 10, 39, 41]. The present chapter is devoted to a discussion of the definitions to be used in the remainder of the paper. The concepts introduced here will be revisited in Section 3.7, using the language of trees. The interplay of these two perspectives is essential to the proof of Theorem 1.3.

2.1 Lacunarity on the real line

Definition 2.1 (Lacunary sequence)**.** *Let* $A = \{a_1, a_2, \ldots\}$ *be an infinite sequence of points contained in a compact subset of* $\mathbb{R}$*. Given a constant* $0 < \lambda < 1$*, we say that* A *is a* lacunary sequence *converging to* α *with* constant of lacunarity *at most* λ*, if*

$$|a_{j+1} - \alpha| \leq \lambda |a_j - \alpha| \quad \textit{for all } j \geq 1.$$

Definition 2.2 (Lacunary sets)**.** *In* $\mathbb{R}$*, a* lacunary set of order 0 *is a set of cardinality* *at most 1; i.e., either empty or a singleton. Recursively, given a constant* $0 < \lambda < 1$ *and an integer* $N \geq 1$*, we say that a relatively compact subset* U *of* $\mathbb{R}$ *is a* lacunary set of order at most N *with lacunarity constant at most* λ*, and write* $U \in \Lambda(N, \lambda)$*, if there exists a lacunary sequence* A *with lacunarity constant* $\leq \lambda$ *with the following properties:*

- $U \cap [\sup(A), \infty) = \emptyset$, $U \cap (-\infty, \inf(A)] = \emptyset$,
- *For any two elements* $a, b \in A$, $a < b$ *such that* $(a, b) \cap A = \emptyset$*, the set* $U \cap [a, b) \in \Lambda(N-1, \lambda)$.

The order of lacunarity *of* U *is exactly* N *if* $U \in \Lambda(N, \lambda) \setminus \Lambda(N-1, \lambda)$*. A lacunary sequence* A *obeying the conditions above will be called a* special sequence *and its limit will be termed a* special point *for* U*.*

For any fixed N and λ, the class $\Lambda(N, \lambda)$ is closed under containment, scalar addition and multiplication; these properties, summarized in the following lemma, are easy to verify.

Lemma 2.3. *Let* $U \in \Lambda(N, \lambda)$*. Then*

(i) $V \in \Lambda(N, \lambda)$ *for any* $V \subseteq U$.

(ii) $c_1 U + c_2 \in \Lambda(N, \lambda)$ *for any* $c_1 \neq 0$, $c_2 \in \mathbb{R}$.

Proof. For (i), $U \in \Lambda(N, \lambda)$ means that there exists a lacunary sequence A with lacunarity constant $\leq \lambda$ such that $U \cap [\sup(A), \infty) = \emptyset$, $U \cap (-\infty, \inf(A)] = \emptyset$, and $U \cap [a, b) \in \Lambda(N-1, \lambda)$ for any $a, b \in A$ with $a < b$ and $(a, b) \cap A = \emptyset$. Since $V \subseteq U$, the same set of conditions is satisfied exactly with U replaced by V.

To prove (ii), we note that $c_1 A + c_2$ can serve as special sequence for $c_1 U + c_2$, if A is a special sequence for U. Since $U \cap [a, b) \in \Lambda(N-1, \lambda)$ if and only if $(c_1 U + c_2) \cap [c_1 a + c_2, c_1 b + c_2) \in \Lambda(N-1, \lambda)$ with $a, b \in A$, $a < b$, $(a, b) \cap A = \emptyset$, the lemma follows. □

The sets of interest to us are those that are generated by finite unions of sets of the form described in Definition 2.2.

Definition 2.4 (Admissible lacunarity of finite order and sublacunarity)**.** *We say that a relatively compact set $U \subseteq \mathbb{R}$ is an* admissible lacunary set of finite order *if there exist a constant $0 < \lambda < 1$ and integers $1 \leq N_1, N_2 < \infty$ such that U can be covered by N_1 lacunary sets of order $\leq N_2$, each with lacunarity constant $\leq \lambda$. If U does not satisfy this criterion, we call it* sublacunary*.*

2.1.1 Examples

(a) A standard example of a lacunary set of order 1 and lacunarity constant $\lambda \in (0, 1)$ is $U = \{\lambda^j : j \geq 1\}$, or any nontrivial subsequence thereof. Indeed U is itself a lacunary sequence, and hence its own special sequence.

A general lacunary set of order 1 need not always be a lacunary sequence. For example $\{2^{-2j} \pm 4^{-2j} : j \geq 1\}$ is lacunary of order 1 relative to the special sequence $\{2^{-j} : j \geq 1\}$. Despite this, lacunary sequences are in a sense representative of the class $\Lambda(1, \lambda)$, since any set in $\Lambda(1, \lambda)$ can be written as the union of at most four lacunary sequences with lacunarity constant $\leq \lambda$. By Lemma 2.3, the set $\{a\lambda^j + b : j \geq 1\}$ is lacunary of order at most 1 for any unit vector (a, b).

(b) In general, given an integer $k \geq 1$ and constants $M_1 \leq M_2 \leq \cdots \leq M_k$ with $M_1 \geq \max(2, k - 1)$, the set

$$U = \left\{M_1^{-j_1} + M_2^{-j_2} + \cdots + M_k^{-j_k} : 0 \leq j_1 \leq j_2 \leq \cdots \leq j_k\right\}$$

is lacunary of order k and has lacunarity constant $\leq M_1^{-1}$. The special sequence can be chosen to be $A = \{M_1^{-j} : j \geq 1\}$.

(c) A set that is dense in some nontrivial interval, however small, is sublacunary. For example, dyadic rationals of the form $\{\frac{k}{2^m} : 0 \leq k < 2^m\}$ for a fixed m can be written as a finite union of lacunary sequences with a given lacunarity λ, but the number of sequences in the union grows without bound as $m \to \infty$. By Lemma 2.3, a set that contains an affine copy of $\{\frac{k}{2^m} : 0 \leq k < 2^m\}$ for every m is sublacunary.

(d) The set $U = \{2^{-j} + 3^{-k} : j, k \geq 0\}$ can be covered by a finite union of sets in $\Lambda(2, \frac{1}{2})$. For instance the two subsets of U where $k \ln 3 \leq (j-1) \ln 2$ and $k \ln 3 \geq j \ln 2$ respectively are each lacunary of order 2, with $\{3^{-k}\}$ and $\{2^{-j}\}$ being their respective special sequences. The complement, where $(j-1)\ln 2 \leq k \ln 3 \leq j \ln 2$, contains at most one k per j, and is a finite union of lacunary sets of order 1.

(e) A slight variation of the above example: $\{2^{-j} + (q_j - 2^{-j})3^{-k} : j, k \geq 0\}$, where $\{q_j\}$ is an enumeration of the rationals in $[\frac{9}{10}, 1]$, leads to a very different conclusion. This set contains $\{q_j\}$, and is hence sublacunary, even though the set may be viewed as a special sequence $\{2^{-j}\}$ with collections of lacunary sequences converging to every point of it. This example illustrates the relevance of the requirement that the lower order components of $\Lambda(N, \lambda)$ lie in disjoint intervals of $\mathbb{R}$.

(f) Given any $0 < \lambda < 1$ and $m > 0$, there is a constant $C = C(\lambda, m)$ such that for any unit vector (a, b), the set $U_{a,b} = \{a\lambda^j + b\lambda^{mj} : j \geq 1\}$ can be covered by C sets in $\Lambda(1, \lambda)$. We prove an analogous statement along these lines in Section 2.2.1, see example (a).

(g) Given any $0 < \lambda < 1$, $m \in \mathbb{Q} \cap (0, \infty)$, there is a constant $C = C(\lambda, m)$ such that for any unit vector (a, b), the set

$$U_{a,b} = \{u_{jk} = a\lambda^j + b\lambda^{mk} : j, k \geq 1\}$$

can be covered by at most C lacunary sets of order at most 2. This is clear for $(a, b) = (1, 0)$ or $(0, 1)$, with the order of lacunarity being 1. For $ab \neq 0$, there are four possibilities concerning the signs of a and b. We deal with $a > 0$ and $b < 0$, the treatment of which is representative of the general case. The set $U_{a,b}$ is decomposed into three parts:

$$V_{a,b} = \left\{u_{jk} \in U : a\lambda^j + b\lambda^{mk} \geq a\lambda^{j+1}\right\},$$
$$W_{a,b} = \left\{u_{jk} \in U : a\lambda^j + b\lambda^{mk} < b\lambda^{m(k+1)}\right\},$$

$$Z_{a,b} = U_{a,b} \setminus \big[V_{a,b} \cup W_{a,b}\big].$$

Then for every fixed j, the set $V_{a,b} \cap [a\lambda^{j+1}, a\lambda^j)$ is an increasing lacunary sequence with constant $\leq \lambda^m$, converging to $a\lambda^j$. An analogous conclusion holds for $W_{a,b} \cap [b\lambda^{mk}, b\lambda^{m(k+1)})$. Thus $V_{a,b}$ and $W_{a,b}$ are both lacunary of order 2, with their special sequences being $A = \{a\lambda^j\}$ and $A = \{b\lambda^{mk}\}$ respectively. For $u_{jk} \in Z_{a,b}$, the indices j and k obey the inequality

$$-\frac{a}{b}(1-\lambda) < \lambda^{mk-j} \leq -\frac{a}{b}(1-\lambda^m)^{-1}.$$

Since m is rational, the values of $mk - j$ range over rationals of a fixed demonimator (same as that of m). The inequality above therefore permits at most C solutions of $mk - j$, the constant C depending on λ and m, but independent of (a, b). Thus $Z_{a,b}$ is covered by a C-fold union of subsets, each consisting of elements $u_{jk} = \lambda^j(a + b\lambda^{mk-j})$ for which $mk - j$ is held fixed at one of these solutions. Each such set is lacunary of order 1 with lacunarity $\leq \lambda$.

2.1.2 Non-closure of finite order lacunarity under algebraic sums

An important aspect of the class of admissible lacunary sets of finite order is that it is not closed under set-algebraic operations, as we establish in the example furnished below. This feature, perhaps initially counterintuitive, is the main inspiration for the definition of higher dimensional lacunarity provided in the next subsection.

Example: Let $N_j \nearrow \infty$ be a fast growing sequence, and $M_j = 2^{m_j}$ a slower growing one, so that

$$M_j < N_j - N_{j-1}. \tag{2.1}$$

For instance, $N_j = 2^{j^2}$ and $M_j = 2^j$ will do. For every $j \geq 1$ and $1 \leq k \leq M_j = 2^{m_j}$, set $q_{jk} = 2^{-N_j}(1 + k2^{-m_j})$, and define

$$U_j = \{2^{-N_j+k} + q_{jk} : 1 \leq k \leq M_j\}, \quad U = \bigcup_{j=1}^{\infty} U_j, \quad V = \{-2^{-j} : j \geq 1\}.$$

An element of U_j of the form $2^{-N_j+k}+q_{jk}$ lies in the dyadic interval $[2^{-N_j+k}, 2^{-N_j+k+1})$, and for a given k, is the only element of U_j in this interval. Further, $U_j \subseteq [2^{-N_j+1}, 2^{-N_j+M_j+1})$, hence by the relation (2.1), $U_j \cap U_{j'} = \emptyset$ if $j \neq j'$. Thus $U \in \Lambda(1, \frac{1}{2})$, since for any $i \geq 1$, the set $U \cap [2^{-i}, 2^{-i+1})$ is either empty or a single point. Clearly V is a lacunary sequence, hence $V \in \Lambda(1, \frac{1}{2})$ as well, being its own special sequence. On the other hand,

$$U + V \supseteq \bigcup_{j=1}^{\infty} \{q_{jk} : 1 \leq k \leq M_j\}.$$

In other words, $U + V$ contains an affine copy of the dyadic rationals of the form $\{k2^{-m_j} : 1 \leq k \leq 2^{m_j}\}$ in $[0, 1]$, for every j. As discussed in example (c) in Section 2.1.1, $U + V$ is sublacunary.

The counterexample above illustrates the sensitivity of lacunarity on ambient coordinates, and precludes a higher dimensional generalization of this notion that relies on componentwise extension. For instance, the two-dimensional set $U \times V$ (with U, V as above) has lacunary coordinate projections in the current system of coordinates, but there are other directions of projection, for instance the line of unit slope, along which the projection of this set is much more dense.

2.2 Finite order lacunarity in general dimensions

Let $\mathbb{V}$ be a d-dimensional affine subspace of an Euclidean space $\mathbb{R}^n$, $n \geq d$. Given a base point $\mathbf{a}$ of $\mathbb{V}$ and an orthonormal basis $\mathcal{B} = \{\mathbf{v}_1, \ldots, \mathbf{v}_d\}$ of the linear subspace

$\mathbb{V} - a$, we define the projection maps

$$\pi_j = \pi_j[\mathbf{a}, \mathcal{B}] : \mathbb{V} \to \mathbb{R}, \quad \text{via} \quad x = \mathbf{a} + \sum_{i=1}^{d} x_i \mathbf{v}_i \to x_i = \pi_i(x), \quad 1 \leq i \leq d. \quad (2.2)$$

Definition 2.5 (Admissible lacunarity and sublacunarity of Euclidean sets)**.** *Let U be a relatively compact subset of $\mathbb{V}$.*

(i) We say that the set U is admissible lacunary of order at most N (as an Euclidean subset of $\mathbb{V}$) *with lacunarity constant at most $\lambda < 1$ if there exists an integer $R \geq 1$ satisfying the following property: for any choice of basis $\mathcal{B}$ and base point $\mathbf{a}$, and each $1 \leq j \leq d$, the projected set*

$$\pi_j(U) = \{\pi_j(x) : x \in U\} \subseteq \mathbb{R}$$

can be covered by R members of $\Lambda(N, \lambda)$, with the class $\Lambda(N, \lambda)$ as described in Definition 2.2. The projection π_j depends on $\mathbf{a}$ and $\mathcal{B}$ via (2.2). The collection of sets U that obey these conditions for a given choice of N, λ and R will be denoted by $\Lambda_d(N, \lambda, R; \mathbb{V})$.

(ii) The set U is called sublacunary in $\mathbb{V}$ *if it is not admissible lacunary of finite order; i.e., if for any $\lambda < 1$ and integers $N, R \geq 1$ there exists a choice of basis $\mathcal{B}$ and an index $1 \leq j \leq d$ such that $\pi_j(U)$ cannot be covered by any R-fold union of one-dimensional lacunary sets of order at most N and lacunarity constant at most λ.*

Remarks:

- An equivalent formulation of the definition of $U \in \Lambda_d(N, R, \lambda; \mathbb{V})$ is that for any line L in $\mathbb{V}$ (and indeed in $\mathbb{R}^{d+1}$ as we will soon see in Lemma 2.6), the projection of U onto L is coverable by at most R sets in $\Lambda(N, \lambda)$.

- The choice of base point in $\mathbb{V}$ is not important in this definition, since $\pi_j[\mathbf{a}, \mathcal{B}](U)$ is a translate of $\pi_j[\mathbf{a}', \mathcal{B}](U)$ for any $\mathbf{a}, \mathbf{a}' \in \mathbb{V}$. Thus $\pi_j[\mathbf{a}, \mathcal{B}](U) \in \Lambda(N, \lambda)$ if and only if $\pi_j[\mathbf{a}', \mathcal{B}](U) \in \Lambda(N, \lambda)$.

- The definition is also invariant under rotation in $\mathbb{R}^n$; if O is an orthogonal transformation of $\mathbb{R}^n$, then $U \in \Lambda_d(N, \lambda, R; \mathbb{V})$ if and only if $O(U) \in \Lambda_d(N, \lambda, R; O(\mathbb{V}))$.

- The choice of rotation $\mathcal{B}$ within $\mathbb{V}$ is however critical. It is not possible to have necessary and sufficient implications like the ones above for two arbitrary choices of bases $\mathcal{B}$ and $\mathcal{B}'$. We provide examples below. Henceforth, we will refer to the choice of a pair $\varphi = (\mathbf{a}, \mathcal{B})$ as a system of coordinates, with the main focus on $\mathcal{B}$.

- Loosely speaking, the order of lacunarity N may be viewed as the number of independent parameters required to describe Ω. The examples in the next section will substantiate this statement.

- If $U \in \Lambda_d(N, \lambda, R; \mathbb{V})$, then the order of lacunarity of $\pi_L U$ cannot exceed N for generic lines L, where π_L denotes the projection onto L. However, there may exist many lines onto which the projections are much sparser than what the order of lacunarity suggests, see for example (2.8).

Before proceeding to examples, we check the definition for consistency if U is a subset of several affine subspaces.

Lemma 2.6. *Let $U \subseteq \mathbb{V}$ be as above. Then for any choice of N, R, λ, the set $U \in \Lambda_d(N, R, \lambda; \mathbb{V})$ if and only if $U \in \Lambda_n(N, R, \lambda; \mathbb{R}^n)$.*

Proof. The "if" implication is clear, so we consider the converse. Without loss of generality, we may choose $\mathbb{V} = \{1\} \times \mathbb{R}^{n-1}$. Given any unit vector $\omega = (\omega_1, \cdots, \omega_n) \in \mathbb{R}^n$ with $0 < |\omega_1| < 1$, let $\mathbb{L}$ denote the line through the origin in $\mathbb{R}^n$ pointing in the direction of ω. Let $\mathbb{L}'$ denote the projection of $\mathbb{L}$ on $\mathbb{V}$, so that $\mathbb{L}' = \{e_1 + s\omega' : s \in \mathbb{R}\}$, where e_1 is the first canonical basis vector in $\mathbb{R}^n$, and $\omega' = (0, \omega_2, \cdots, \omega_n)$. The desired conclusion follows from the claim that

$$\text{the sets } \pi(U) \text{ and } \pi'(U) \text{ are affine copies of each other,} \tag{2.3}$$

where $\pi(U)$ and $\pi'(U)$ denote the scalar projections onto $\mathbb{L}$ and $\mathbb{L}'$, measured from the origin and from $(1, 0, \cdots, 0)$ respectively. Indeed, Lemma 2.3 then permits us to extend known lacunarity features of the former directly to the latter.

To establish (2.3), it suffices to note that for any $x \in \mathbb{R}^n$,

$$\pi(x) = (x \cdot \omega)\omega, \quad \text{and} \quad \pi'(x) = \frac{x \cdot \omega'}{|\omega'|^2}\omega'.$$

The choice of $\mathbb{V}$, ω and ω' yield the relations $(x-y)\cdot\omega = (x-y)\cdot\omega'$ for any $x, y \in U$, hence the above expressions imply that

$$|\pi(x) - \pi(y)| = |\omega'||\pi'(x) - \pi'(y)|,$$

which is the desired conclusion. □

2.2.1 Examples of admissible lacunary and sublacunary sets in $\mathbb{R}^d$

(a) A set of the form considered by Nagel, Stein and Wainger [37], such as

$$U = \{\gamma(\theta_j) : j \geq 1\}, \quad \text{where} \quad \gamma(t) = (t^{m_1}, \cdots, t^{m_d}) \tag{2.4}$$

is admissible lacunary of order 1. Here $0 < m_1 < \cdots < m_d$ are fixed constants, and $0 < \theta_{j+1} \leq \lambda\theta_j$, for some $0 < \lambda < 1$ and all j. Critical to this verification are the following two properties of U appearing in [37, Lemma 4]:

- There is a constant $C_1 = C_1(m_1, \cdots, m_d)$ obeying the following requirement. For any unit vector $\xi = (\xi_1, \cdots, \xi_d)$ in $\mathbb{R}^d$, the set $\mathbb{N}$ of positive integers can be decomposed into C_1 disjoint consecutive intervals $\{\mathbb{N}_s\}$; for every s, there exists $r(s) \in \{1, \cdots, d\}$ such that

$$\max_{1 \leq r \leq d} |\theta_j^{m_r}\xi_r| = \left|\theta_j^{m_{r(s)}}\xi_{r(s)}\right| \quad \text{for all } j \in \mathbb{N}_s. \tag{2.5}$$

  The composition of $\mathbb{N}_s$ depends on ξ.

- Further for any $c > 0$, there is a constant $C_2 = C_2(c, m_1, \cdots, m_d)$ independent of ξ and $\mathbb{N}_s$ so that

$$\max_{\substack{r\in\{1,\cdots,d\}\\ r\neq r(s)}} \left|\theta_j^{m_r}\xi_r\right| < c\left|\theta_j^{m_{r(s)}}\xi_{r(s)}\right| \tag{2.6}$$

for all but C_2 integers $j \in \mathbb{N}_s$.

Assuming these two facts, the claim of lacunarity is established as follows. Using the definition of $\mathbb{N}_s$ in (2.5), the set U can be decomposed into C_1 pieces U_s, where $U_s = \{\gamma(\theta_j) : j \in \mathbb{N}_s\}$. Fix a constant R such that $2d\lambda^{m_1 R-1} < 1$. If $j' > j$ are two integers in $\mathbb{N}_s$ that are at least R-separated and for both of which (2.6) holds with $c = \frac{1}{2d}$, then

$$\begin{aligned}\Big|\sum_{r=1}^{d}\xi_r\theta_{j'}^{m_r}\Big| &\leq d\left|\xi_{r(s)}\theta_{j'}^{m_{r(s)}}\right| \leq d\big(\lambda^{j'-j}\big)^{m_{r(s)}}\left|\xi_{r(s)}\theta_j^{m_{r(s)}}\right| \\ &\leq 2d\big(\lambda^R\big)^{m_1}\Big|\sum_{r=1}^{d}\xi_r\theta_j^{m_r}\Big| < \lambda\Big|\sum_{r=1}^{d}\xi_r\theta_j^{m_r}\Big|.\end{aligned} \tag{2.7}$$

Thus each U_s is the union of at most R lacunary sequences of lacunarity $< \lambda$, together with the C_2 points where (2.6) fails.

(b) A set of the form considered by Carbery [10], i.e.,

$$U = \{\Gamma_{\mathbf{k}} = (\lambda^{k_1}, \cdots, \lambda^{k_d}) : \mathbf{k} = (k_1, \cdots, k_d) \in \mathbb{N}^d\} \tag{2.8}$$

is admissible lacunary of order d. We prove this by induction on d. The initializing step for $d = 2$ has been covered in example (g) of Section 2.1.1. For a general d and after splitting U into $d!$ pieces, we may assume that $k_1 \leq k_2 \leq \cdots \leq k_d$. Given any unit vector $\xi = (\xi_1, \cdots, \xi_d) \in \mathbb{R}^d$, we write

$$U = \bigcup_{s=1}^{d} U_s \quad \text{with} \quad U_s = \{\Gamma_{\mathbf{k}} : \mathbf{k} \in \mathbb{N}_s^d\}, \text{ where}$$
$$\mathbb{N}_s^d = \left\{\mathbf{k} \in \mathbb{N}^d : \left|\lambda^{k_s}\xi_s\right| = \max_{1\leq r\leq d}\left|\lambda^{k_r}\xi_r\right|\right\}.$$

Depending on the signs of $\lambda^{k_s}\xi_s$ and $\Gamma_k \cdot \xi - \lambda^{k_s}\xi_s$, each $\mathbb{N}_s^d$ can be decomposed into four parts. Their treatments are similar with trivial adjustments, so we focus on the subset of $\mathbb{N}_s^d$ where

$$\lambda^{k_s}\xi_s > 0 \quad \text{and} \quad \sum_{r \neq s} \lambda^{k_r}\xi_r \geq 0,$$

continuing to call this subset $\mathbb{N}_s^d$ to ease notational burden. One last splitting is needed; for a constant A to be specified shortly, we write

$$\mathbb{N}_s^d = \mathbb{N}_{s,1}^d \cup \mathbb{N}_{s,2}^d, \text{ where } \mathbb{N}_{s,1}^d = \{\mathbf{k} \in \mathbb{N}_s^d : \lambda^{k_s}\xi_s > A|\lambda^{k_r}\xi_r| \text{ for all } r \neq s\}.$$

For $\mathbf{k} \in \mathbb{N}_{s,1}^d$,

$$\lambda^{k_s}\xi_s \leq \Gamma_{\mathbf{k}} \cdot \xi < \lambda^{k_s}\xi_s\big(1 + dA^{-1}\big) < \lambda^{k_s - 1}\xi_s, \tag{2.9}$$

where the last inequality follows for a suitable choice of A. We argue that $\{\xi_s\lambda^{k_s} : k_s \geq 1\}$ may be viewed as a special sequence for $\{\Gamma_{\mathbf{k}} \cdot \xi : \mathbf{k} \in \mathbb{N}_{s,1}^d\}$. Indeed, if k_s is fixed, then (2.9) shows that

$$\begin{aligned} \{\Gamma_{\mathbf{r}} \cdot \xi : \mathbf{r} \in \mathbb{N}_{s,1}^d\} \cap [\xi_s\lambda^{k_s}, \xi_s\lambda^{k_s-1}) &= \{\Gamma_{\mathbf{r}} \cdot \xi : \mathbf{r} \in \mathbb{N}_{s,1}^d,\ r_s = k_s\} \\ &\subseteq \xi_s\lambda^{k_s} + \Big\{\sum_{r \neq s} \lambda^{k_r}\xi_r : k_r \in \mathbb{N},\ r \neq s\Big\}. \end{aligned}$$

By the induction hypothesis, there is a constant R independent of ξ such that the set on the right hand side above is coverable by at most R sets in $\Lambda(d-1;\lambda)$. Hence $\{\Gamma_{\mathbf{k}} : \mathbf{k} \in \mathbb{N}_{s,1}^d\}$ is admissible lacunary of order d.

We turn to the complementary set $\mathbb{N}_{s,2}^d$. After decomposing $\mathbb{N}_{s,2}^d$ into $(d-1)$ subsets, we may fix an index ℓ such that

$$|\lambda^{k_\ell}\xi_\ell| \leq \lambda^{k_s}\xi_s \leq A|\lambda^{k_\ell}\xi_\ell| \tag{2.10}$$

on $\mathbb{N}_{s,2}^d$. Without loss of generality let $\ell \geq s$. The number of possible values of $k_\ell - k_s$ obeying (2.10) is at most a fixed constant C depending on A (hence λ

and d), but independent of ξ. Thus $\mathbb{N}^d_{s,2}$ may be written as the C-fold union of subsets indexed by c, where the subset identified by c contains all $\mathbf{k} \in \mathbb{N}^d_{s,2}$ with the property that $k_\ell - k_s = c \geq 0$. For $\mathbf{k}$ in such a subset,

$$\Gamma_{\mathbf{k}} \cdot \xi = (\xi_s + \lambda^c \xi_\ell)\lambda^{k_s} + \sum_{r \neq \ell, s} \lambda^{k_r} \xi_r.$$

Since the number of summands in the linear combination above is $(d-1)$, the induction hypothesis dictates that $\{\Gamma_{\mathbf{k}} : \mathbf{k} \in \mathbb{N}^d_{s,2}\}$ is admissible lacunary of order $(d-1)$, completing the proof.

(c) A curve in $\mathbb{R}^d$ is sublacunary. So is a Cantor-like subset of it, by Theorem 1.2.

(d) If U and V are the lacunary sets of order 1 constructed in Section 2.1.2, the set $U \times V$ is sublacunary. Indeed, after a rotation of angle $\frac{\pi}{4}$ one of the coordinate projections turns out to be a constant multiple of $U+V$. We have seen in Section 2.1.2 that this last set is sublacunary on $\mathbb{R}$.

2.3 Finite order lacunarity for direction sets

Given two sets $\Omega_1, \Omega_2 \subseteq \mathbb{R}^{d+1} \setminus \{0\}$, we say that $\Omega_1 \sim \Omega_2$ if

$$\left\{ \frac{\omega}{|\omega|} : \omega \in \Omega_1 \right\} = \left\{ \frac{\omega}{|\omega|} : \omega \in \Omega_2 \right\}.$$

The binary relation $\sim$ is clearly an equivalence relation among sets in $\mathbb{R}^{d+1} \setminus \{0\}$. An equivalence class of $\sim$ is, by definition, a *direction set*. By a slight abuse of nomenclature, we will refer to a set $\Omega \subseteq \mathbb{R}^{d+1} \setminus \{0\}$ as a direction set to mean the equivalence class of $\sim$ that contains Ω. Clearly the maximal operators D_Ω and M_Ω, as well as the admittance of Kakeya-type sets (as in Definition 1.1), remain unchanged for all members of this equivalence class.

Certain modifications are necessary to extend the notion of lacunarity from Euclidean sets to direction sets, in view of the latter's scale invariance. Given a direction set $\Omega \subseteq \mathbb{R}^{d+1} \setminus \{0\}$, we denote by $\mathcal{C}_\Omega$ the cone generated by this set of directions, namely

$$\mathcal{C}_\Omega := \{r\omega : r > 0,\ \omega \in \Omega\}. \tag{2.11}$$

Definition 2.7. *Let $\Omega \subseteq \mathbb{R}^{d+1} \setminus \{0\}$ be a direction set, with $\mathcal{C}_\Omega$ as in* (2.11)*.*

(i) Given an integer N and a positive constant $\lambda < 1$, we say that Ω is admissible lacunary as a direction set *with order at most N and lacunarity at most λ if there exists an integer R such that $U \in \Lambda_d(N, \lambda, R; \mathbb{V})$ in the sense of Definition 2.5, for every hyperplane $\mathbb{V}$ at unit distance from the origin and every relatively compact subset U of $\mathcal{C}_\Omega \cap \mathbb{V}$. The collection of direction sets that obey these conditions for a given N, λ and R will be denoted by $\Delta_d(N, \lambda, R)$.*

(ii) A direction set $\Omega \subseteq \Omega_0$ failing this property is termed a sublacunary direction set*. Thus Ω is sublacunary as a direction set if for any choice of integers N, R and positive constant $\lambda < 1$ there is a tangential hyperplane $\mathbb{V}$ of the unit sphere, a relatively compact subset U of $\mathcal{C}_\Omega \cap \mathbb{V}$ and a line L in $\mathbb{V}$ such that the projection of U along L cannot be covered by any R-fold union of sets in $\Lambda(N, \lambda)$.*

2.3.1 Examples of admissible lacunary and sublacunary direction sets

(a) A direction set Ω of the form considered by Nagel, Stein and Wainger [37],

$$\Omega = \{u_j = (\gamma(\theta_j), 1) : j \geq J\}$$

is admissible lacunary of order 1. Here the function γ and the sequence θ_j are as described in example (a) of Section 2.2.1. Thus Ω is parameterized by the positive constants $m_1 < m_2 < \cdots < m_d$. We set $m_{d+1} = 0$. To verify the claim, we choose $\mathbb{V} = \{x \in \mathbb{R}^{d+1} : x \cdot \eta = 1\}$ for some unit vector η, so that

$$\mathcal{C}_\Omega \cap \mathbb{V} = \left\{ v_j = \frac{u_j}{u_j \cdot \eta} : u_j \in \Omega \right\}.$$

Fix a unit vector $\omega = (\omega', \omega_{d+1}) \in \mathbb{R}^{d+1}$, and let π_ω denote the scalar projection onto ω; i.e., $\pi_\omega(v) = v \cdot \omega$. As required by Definitions 2.7 and 2.5 and in view of Lemma 2.6, we aim to show that that there is a large constant R (independent of $\mathbb{V}$) for which any relatively compact subset of $\pi_\omega(\mathcal{C}_\Omega \cap \mathbb{V})$ can be covered by R members of $\Lambda(1; \lambda)$. By the property (2.5) of Ω, we first decompose the integers into a bounded number C_1 of disjoint intervals (C_1 independent of ω and η), on each of which there exists an index $1 \leq r \leq d+1$ such that

$$\max_{1 \leq i \leq d+1} |\theta_j^{m_i} \eta_i| = |\theta_j^{m_r} \eta_r|. \tag{2.12}$$

Let us denote by $\mathbb{N}_r[\eta]$ one of the subintervals for which (2.12) holds. For $j \in \mathbb{N}_r[\eta]$,

$$\pi_\omega(v_j) - \frac{\omega_r}{\eta_r} = \frac{\xi \cdot u_j}{\eta_r\ (\eta \cdot u_j)}, \qquad \text{where} \qquad \xi = (\xi_1, \cdots, \xi_{d+1}) \in \mathbb{R}^{d+1} \tag{2.13}$$

with $\xi_k = \omega_k \eta_r - \omega_r \eta_k$, so that $\xi_r = 0$. Our goal is to show that for $j \in \mathbb{N}_r[\eta]$, the sequence on the right hand side above can be covered by an R-fold union of lacunary sequences converging to 0.

Using (2.5) again, we decompose $\mathbb{N}_r[\eta]$ into at most C_1 pieces, of the form $\mathbb{N}_{rs}[\eta, \xi] = \mathbb{N}_r[\eta] \cap \mathbb{N}_s[\xi]$. Since $\xi_r = 0$, we conclude that $\mathbb{N}_r[\xi] = \emptyset$; hence $s \neq r$. By property (2.6), for every $c > 0$, there are at most a bounded number $C_2 = C_2(c)$ indices $j \in \mathbb{N}_{rs}[\eta, \xi]$ for which at least one of the inequalities

$$\max_{i \neq r} |\theta_j^{m_i} \eta_i| < c|\theta_j^{m_r} \eta_r|, \qquad \max_{i \neq s} |\theta_j^{m_i} \xi_i| < c|\theta_j^{m_s} \xi_s| \tag{2.14}$$

fails.

First suppose $s > r$. Choosing two integers $j, j' \in \mathbb{N}_{rs}[\eta, \xi]$ with $j' - j \geq R$ for both of which the constraints in (2.14) hold, we follow the steps laid out in (2.7), obtaining from (2.13)

$$\begin{aligned}\left[\left|\pi_\omega(v_{j'}) - \frac{\omega_r}{\eta_r}\right|\right] \left[\left|\pi_\omega(v_j) - \frac{\omega_r}{\eta_r}\right|\right]^{-1} &= \frac{\xi \cdot u_{j'}}{\xi \cdot u_j} \cdot \frac{\eta \cdot u_j}{\eta \cdot u_{j'}} \\ &\leq \left[\frac{d|\xi_s|\theta_{j'}^{m_s}}{\frac{1}{2}|\xi_s|\theta_j^{m_s}}\right] \cdot \left[\frac{d|\eta_r|\theta_j^{m_r}}{\frac{1}{2}|\eta_r|\theta_{j'}^{m_r}}\right] \\ &\leq 4d^2 \left(\frac{\theta_{j'}}{\theta_j}\right)^{m_s - m_r} \leq 4d^2 \lambda^{R(m_s - m_r)}.\end{aligned}$$

If R is selected large enough to satisfy $4d^2\lambda^{R(m_s-m_r)} < \lambda$, then for $j \in \mathbb{N}_{rs}[\eta, \xi]$ the sequence on the right hand side of (2.13) can be covered by the union of R lacunary sequences converging to zero, excluding the C_2 points where (2.14) fails. For $s < r$, the same calculation above can be replicated for $j' < j$ with $j' - j < -R$. Thus in this case the sequence in (2.13) grows as j increases, and hence has to be finite by the assumption of relative compactness. Nonetheless, this finite sequence is still coverable by a lacunary sequence going to zero, this time in reverse order of j. In either event, we have decomposed the set $\{\pi_\omega(v_j) : j \in \mathbb{N}_{rs}[\eta, \xi]\}$ into R lacunary sequences of lacunarity λ, proving the claim.

(b) A direction set of the type studied in [10], namely

$$\Omega = \{(\Gamma_{\mathbf{k}}, 1) : 0 \leq k_1 \leq k_2 \leq \cdots \leq k_d\},$$

(with $\Gamma_{\mathbf{k}}$ as in (2.8)) is admissible lacunary of order d. This is proved along lines similar to the example above, using methods already explained in examples (g) and (b) of Section 2.1.1 and 2.2.1 respectively; we omit the details here.

(c) A curve in $\mathbb{R}^{d+1}$ is sublacunary as a direction set.

(d) For sets U, V as constructed in Section 2.1.2, the direction set $\Omega = \{1\} \times U \times V$ is sublacunary, since $U \times V$ is sublacunary as an Euclidean set (see example (d) in Section 2.2.1).

(e) Let $\{q_\ell : \ell \geq 1\}$ be an enumeration of the rationals on any nontrivial interval, say on $[\frac{1}{2}, \frac{2}{3}]$. A direction set of the type considered by Parcet and Rogers [39, Example 1 on page 4], such as

$$\Omega = \{(q_\ell 2^{-\ell}, 2^{-\ell}, 1) : \ell \geq 1\}$$

is sublacunary, even though the one-dimensional coordinate projections in the current coordinate system are lacunary of order at most 1. Choosing $\mathbb{V} = \{x_2 = 1\}$, we find that

$$\mathcal{C}_\Omega \cap \mathbb{V} = \{(q_\ell, 1, 2^\ell) : \ell \geq 1\}.$$

The order of lacunarity of the x_1-projection grows without bound as we choose increasingly large compact subsets of $\mathcal{C}_\Omega \cap \mathbb{V}$.

(f) We also mention another example considered by Parcet and Rogers [39, Example 2 on page 4]. Given the canonical orthonormal basis $\{e_1, e_2, e_3\}$ of $\mathbb{R}^3$, let us fix another orthonormal basis $\{e_1, e_2', e_3'\}$ with $\text{span}\{e_2, e_3\} = \text{span}\{e_2', e_3'\}$ and e_3' lying in the first quadrant determined by e_2 and e_3. The direction set under consideration is $\Omega = \{u_\ell : \ell \geq 1\}$, where u_ℓ is a sequence of vectors satisfying $u_\ell \cdot e_2' = q_\ell u_\ell \cdot e_1$ for some enumeration of rationals $\{q_\ell\}$ in an interval. The last condition does not completely specify u_ℓ, hence the direction set so defined is not unique (further restrictions are imposed in [39]), but regardless of any subsequent choice Ω is sublacunary. Choosing $\mathbb{V} = \{x_1 = 1\}$, we observe that

$$\mathcal{C}_\Omega \cap \mathbb{V} = \left\{ \frac{u_\ell}{u_\ell \cdot e_1} : \ell \geq 1 \right\}.$$

Projecting $\mathcal{C}_\Omega \cap \mathbb{V}$ in the direction e_2', we find that the projected set is $\{q_\ell : \ell \geq 1\}$, which is not lacunary of finite order.

Chapter 3

Rooted, labelled trees

As in Bateman and Katz's work [4], [3], the language of rooted, labelled trees remains the vehicle of choice for construction of Kakeya-type sets. We explore the basic terminology of trees and state the relevant facts in Sections 3.1 and 3.6 below. This constitutes only a very small part of the literature on such objects, but will suffice for our purposes. These objects are treated comprehensively in the text by Lyons and Peres [36].

3.1 The terminology of trees

An undirected graph $\mathcal{G} := (\mathcal{V}, \mathcal{E})$ is a pair, where $\mathcal{V}$ is a set of vertices and $\mathcal{E}$ is a symmetric, nonreflexive subset of $\mathcal{V} \times \mathcal{V}$, called the edge set. By symmetric, here we mean that the pair $(u, v) \in \mathcal{E}$ is unordered; i.e. the pair (u, v) is identical to the pair (v, u). By nonreflexive, we mean $\mathcal{E}$ does not contain the pair (v, v) for any $v \in \mathcal{V}$.

A path in a graph is a sequence of vertices such that each successive pair of vertices is a distinct edge in the graph. A finite path (with at least one edge) whose first and last vertices are the same is called a cycle. A graph is connected if for each pair of vertices $v \neq u$, there is a path in $\mathcal{G}$ containing v and u. We define a *tree* to be a connected undirected graph with no cycles.

All our trees will be of a specific structure. A *rooted, labelled tree* $\mathcal{T}$ is one whose vertex set is a nonempty collection of finite sequences of nonnegative integers such that if $\langle i_1, \ldots, i_n \rangle \in \mathcal{T}$, then

(i.) for any k, $0 \le k \le n$, $\langle i_1, \ldots, i_k \rangle \in \mathcal{T}$, where $k = 0$ corresponds to the empty sequence, and

(ii.) for every $j \in \{0, 1, \ldots, i_n\}$, we have $\langle i_1, \ldots, i_{n-1}, j \rangle \in \mathcal{T}$.

We say that $\langle i_1, \ldots, i_{n-1} \rangle$ is the *parent* of $\langle i_1, \ldots, i_{n-1}, j \rangle$ and that $\langle i_1, \ldots, i_{n-1}, j \rangle$ is the $(j+1)$*th child* of $\langle i_1, \ldots, i_{n-1} \rangle$. If u and v are two sequences in $\mathcal{T}$ such that u is a child of v, or a child's child of v, or a child's child's child of v, etc., then we say that u is a *descendant* of v (or that v is an *ancestor* of u), and we write $u \subset v$ (see the remark below). If $u = \langle i_1, \ldots, i_m \rangle \in \mathcal{T}$, $v = \langle j_1, \ldots, j_n \rangle \in \mathcal{T}$, $m \le n$, then the *youngest common ancestor* of u and v is the vertex in $\mathcal{T}$ defined by

$$D(u,v) = D(v,u) := \begin{cases} \emptyset, & \text{if } i_1 \neq j_1 \\ \langle i_1, \ldots, i_k \rangle & \text{if } k = \max\{l : i_l = j_l\}. \end{cases} \tag{3.1}$$

One can similarly define the youngest common ancestor for any finite collection of vertices.

Remark: At first glance, using the notation $u \subset v$ to denote when u is a descendant of v may seem counterintuitive, since u is a descendant of v precisely when v is a subsequence of u. However, we will soon be identifying vertices of rooted labelled trees with certain nested families of cubes in $\mathbb{R}^d$. Consequently, as will become apparent in the next two sections, u will be a descendant of v precisely when the cube associated with u is contained within the cube associated with v.

We designate the empty sequence $\emptyset$ as the *root* of the tree $\mathcal{T}$. The sequence $\langle i_1, \ldots, i_n \rangle$ should be thought of as the vertex in $\mathcal{T}$ that is the $(i_n + 1)th$ child of the $(i_{n-1} + 1)th$ child,..., of the $(i_1 + 1)th$ child of the root. All unordered pairs of the form $(\langle i_1, \ldots, i_{n-1} \rangle, \langle i_1, \ldots, i_{n-1}, i_n \rangle)$ describe the edges of the tree $\mathcal{T}$. We say that the edge originates at the vertex $\langle i_1, \ldots, i_{n-1} \rangle$ and that it terminates at the vertex $\langle i_1, \ldots, i_{n-1}, i_n \rangle$. Note that every vertex in the tree that is not the root is uniquely identified by the edge terminating at that vertex. Consequently, given an edge $e \in \mathcal{E}$, we define $v(e)$ to be the vertex in $\mathcal{V}$ at which e terminates. The vertex $\langle i_1, \ldots, i_n \rangle \in \mathcal{T}$ also prescribes a unique path, or *ray*, from the root to this vertex:

$$\emptyset \to \langle i_1 \rangle \to \langle i_1, i_2 \rangle \to \cdots \to \langle i_1, i_2, \ldots, i_n \rangle.$$

We let $\partial\mathcal{T}$ denote the collection of all rays in $\mathcal{T}$ of maximal (possibly infinite) length. For a fixed vertex $v = \langle i_1, \ldots, i_m \rangle \in \mathcal{T}$, we also define *the subtree (of $\mathcal{T}$) generated by the vertex v* to be the maximal subtree of $\mathcal{T}$ with v as the root, i.e. it is the subtree

$$\{\langle i_1, \ldots, i_m, j_1, \ldots, j_k \rangle \in \mathcal{T} : k \geq 0\}.$$

The *height* of the tree is taken to be the supremum of the lengths of all the sequences in the tree. Further, we define the *height* $h(\cdot)$, or *level*, of a vertex $\langle i_1, \ldots, i_n \rangle$ in the tree to be n, the length of its identifying sequence. All vertices of height n are said to be members of the nth *generation* of the root, or interchangeably, of the tree. More explicitly, a member vertex of the nth generation has exactly n edges joining it to the root. The height of the root is always taken to be zero.

If $\mathcal{T}$ is a tree and $n \in \mathbb{Z}^+$, we write the *truncation* of $\mathcal{T}$ to its first n levels as $\mathcal{T}_n = \{\langle i_1, \ldots, i_k \rangle \in \mathcal{T} : 0 \leq k \leq n\}$. This subtree is a tree of height at most n. A tree is called *locally finite* if its truncation to every level is finite, i.e. consists of finitely many vertices. All of our trees will have this property. In the remainder of this document, when we speak of a *tree* we will always mean a *locally finite, rooted labelled tree.*

Roughly speaking, two trees are isomorphic if they have the same collection of rays. To make this precise we define a special kind of map between trees; this definition will be very important for us later.

Definition 3.1. *Let $\mathcal{T}$ and $\mathcal{T}'$ be two trees with equal (possibly infinite) heights. A map $\sigma : \mathcal{T} \to \mathcal{T}'$ is called* **sticky** *if*

- *for all $v \in \mathcal{T}$, $h(v) = h(\sigma(v))$, and*

- *$u \subset v$ implies $\sigma(u) \subset \sigma(v)$ for all $u, v \in \mathcal{T}$.*

We often say that σ is sticky if it preserves heights and lineages.

A one-to-one and onto sticky map between two trees, when it exists, is said to be an *isomorphism* and the two trees are said to be *isomorphic*. Two isomorphic trees can and will be treated as essentially identical objects.

Comparing this definition of stickiness with the one described in Section 1.5, the two notions do not immediately appear compatible. Indeed, the notion of stickiness described here in Definition 3.1 is, in some ways, far more restrictive than the geometric one discussed in Section 1.5. We will return to these observations shortly, in Section 3.4, after we have established how to encode Euclidean sets by trees in general.

3.2 Encoding bounded subsets of the unit interval by trees

The language of rooted labelled trees is especially convenient for representing bounded sets in Euclidean spaces. This connection is well-studied in the literature [36].

We start with $[0, 1) \subset \mathbb{R}$. Fix any positive integer $M \geq 2$. We define an M-adic rational as a number of the form i/M^k for some $i \in \mathbb{Z}$, $k \in \mathbb{Z}^+$, and an M-adic interval as $[i \cdot M^{-k}, (i+1) \cdot M^{-k})$. For any nonnegative integer i and positive integer k such that $i < M^k$, there exists a unique representation

$$i = i_1 M^{k-1} + i_2 M^{k-2} + \cdots + i_{k-1} M + i_k, \tag{3.2}$$

where the integers $i_1, \ldots, i_k$ take values in $\mathbb{Z}_M := \{0, 1, \ldots, M-1\}$. These integers should be thought of as the "digits" of i with respect to its base M expansion. An easy consequence of (3.2) is that there is a one-to-one and onto correspondence between M-adic rationals in $[0, 1)$ of the form i/M^k and finite integer sequences $\langle i_1, \ldots, i_k \rangle$ of length k with $i_j \in \mathbb{Z}_M$ for each j. Naturally then, we define the tree of infinite height

$$\mathcal{T}([0,1);M) = \{\langle i_1, \dots, i_k \rangle : k \geq 0,\ i_j \in \mathbb{Z}_M\}. \tag{3.3}$$

The tree thus defined depends of course on the base M; however, once the base M has been fixed, we will omit its usage in our notation, denoting the tree $\mathcal{T}([0,1);M)$ by $\mathcal{T}([0,1))$ instead.

Identifying the root of the tree defined in (3.3) with the interval $[0,1)$ and the vertex $\langle i_1, \dots, i_k \rangle$ with the interval $[i \cdot M^{-k}, (i+1) \cdot M^{-k})$, where i and $\langle i_1, \dots, i_k \rangle$ are related by (3.2), we observe that the vertices of $\mathcal{T}([0,1);M)$ at height k yield a partition of $[0,1)$ into M-adic subintervals of length M^{-k}. This tree has a self-similar structure: every vertex of $\mathcal{T}([0,1);M)$ has M children and the subtree generated by any vertex as the root is isomorphic to $\mathcal{T}([0,1);M)$. In the sequel, we will refer to such a tree as a *full M-adic tree*.

Any $x \in [0,1)$ can be realized as the intersection of a nested sequence of M-adic intervals, namely

$$\{x\} = \bigcap_{k=0}^{\infty} I_k(x),$$

where $I_k(x) = [i_k(x) \cdot M^{-k}, (i_k(x)+1) \cdot M^{-k})$ is the unique M-adic interval in the kth M-adic partition of $[0,1)$ containing the point x. The point x should be visualized as the destination of the infinite ray

$$\emptyset \to \langle i_1(x) \rangle \to \langle i_1(x), i_2(x) \rangle \to \cdots \to \langle i_1(x), i_2(x), \dots, i_k(x) \rangle \to \cdots$$

in $\mathcal{T}([0,1);M)$. Conversely, every infinite ray

$$\emptyset \to \langle i_1 \rangle \to \langle i_1, i_2 \rangle \to \langle i_1, i_2, i_3 \rangle \cdots$$

identifies a unique $x \in [0,1)$ given by the convergent sum

$$x = \sum_{j=1}^{\infty} \frac{i_j}{M^j}.$$

Thus the tree $\mathcal{T}([0,1);M)$ can be identified with the interval $[0,1)$ exactly. Any subset $E \subseteq [0,1)$ is then given by a subtree $\mathcal{T}(E;M)$ of $\mathcal{T}([0,1);M)$ consisting of all infinite rays that identify some $x \in E$. As before, we will drop the notation for the base M in $\mathcal{T}(E;M)$ once this base has been fixed.

Any *truncation* of $\mathcal{T}(E;M)$, say up to height k, will be denoted by $\mathcal{T}_k(E;M)$ and should be visualized as a covering of E by M-adic intervals of length M^{-k}. More precisely, $\langle i_1, \ldots, i_k \rangle \in \mathcal{T}_k(E;M)$ if and only if $E \cap [i \cdot M^{-k}, (i+1) \cdot M^{-k}) \neq \emptyset$, where i and $\langle i_1, \ldots, i_k \rangle$ are related via (3.2).

3.3 Encoding higher dimensional bounded subsets of Euclidean space by trees

The approach to encoding a bounded subset of Euclidean space by a tree extends readily to higher dimensions. For any $\mathbf{i} = \langle j_1, \ldots, j_d \rangle \in \mathbb{Z}^d$ such that $\mathbf{i} \cdot M^{-k} \in [0,1)^d$, we can apply (3.2) to each component of $\mathbf{i}$ to obtain

$$\frac{\mathbf{i}}{M^k} = \frac{\mathbf{i}_1}{M} + \frac{\mathbf{i}_2}{M^2} + \cdots + \frac{\mathbf{i}_k}{M^k},$$

with $\mathbf{i}_j \in \mathbb{Z}_M^d$ for all j. As before, we identify $\mathbf{i}$ with $\langle \mathbf{i}_1, \ldots, \mathbf{i}_k \rangle$.

Let $\phi : \mathbb{Z}_M^d \to \{0, 1, \ldots, M^d - 1\}$ be an enumeration of $\mathbb{Z}_M^d$. Define the full M^d-adic tree

$$\mathcal{T}([0,1)^d; M, \phi) = \left\{ \langle \phi(\mathbf{i}_1), \ldots, \phi(\mathbf{i}_k) \rangle : k \geq 0,\ \mathbf{i}_j \in \mathbb{Z}_M^d \right\}. \tag{3.4}$$

The collection of kth generation vertices of this tree may be thought of as the d-fold Cartesian product of the kth generation vertices of $\mathcal{T}([0,1);M)$. For our purposes, it will suffice to fix ϕ to be the lexicographic ordering, and so we will omit the notation for ϕ in (3.4), writing simply, and with a slight abuse of notation,

$$\mathcal{T}([0,1)^d; M) = \left\{ \langle \mathbf{i}_1, \ldots, \mathbf{i}_k \rangle : k \geq 0,\ \mathbf{i}_j \in \mathbb{Z}_M^d \right\}. \tag{3.5}$$

As before, we will refer to the tree in (3.5) by the notation $\mathcal{T}([0,1)^d)$ once the base M has been fixed.

By a direct generalization of our one-dimensional results, each vertex $\langle \mathbf{i}_1, \ldots, \mathbf{i}_k \rangle$ of $\mathcal{T}([0,1)^d; M)$ at height k represents the unique M-adic cube in $[0,1)^d$ of sidelength M^{-k}, containing $\mathbf{i} \cdot M^{-k}$, of the form

$$\left[\frac{j_1}{M^k}, \frac{j_1+1}{M^k}\right) \times \cdots \times \left[\frac{j_d}{M^k}, \frac{j_d+1}{M^k}\right).$$

As in the one-dimensional setting, any $x \in [0,1)^d$ can be realized as the intersection of a nested sequence of M-adic cubes. Thus, we view the tree in (3.5) as an encoding of the set $[0,1)^d$ with respect to base M. As before, any subset $E \subseteq [0,1)^d$ then corresponds to a subtree of $\mathcal{T}([0,1)^d; M)$.

3.4 Stickiness as a kind of mapping between trees

Throughout this document, we will be concerned with studying certain mappings between pairs of given trees of equal height and comparable base. The first tree $\mathcal{T}$ will represent a set in Euclidean space that *roots* a collection of thin tubes; i.e. each point in the set will identify the location in Euclidean space of one particular tube in this collection. The second tree $\mathcal{S}$, to which we will map, will encode a set of directions or *slopes*, one for each tube in our collection.

Let $\sigma : \mathcal{T} \to \mathcal{S}$ be a height-preserving transformation that maps full-length rays in $\mathcal{T}$ into full-length rays in $\mathcal{S}$, and let

$$K_\sigma := \bigcup_{t \in \mathcal{T}} P_{t,\sigma},$$

where $P_{t,\sigma}$ is a thin tube rooted at t and oriented in the direction $\sigma(t)$. If σ is sticky in the sense of Definition 3.1, then the collection of tubes K_σ is, on average, sticky in the geometric sense of the term discussed in Section 1.5. This will be made precise later in Section 5.1, but for now it suffices to imagine the following heuristic.

Fix a pair of trees $\mathcal{T}$ and $\mathcal{S}$ with the same height H and base M. Consider a pair of tubes rooted at $t_1, t_2 \in \mathcal{T}$ respectively. By our work in the previous two sections, we know that $|t_1 - t_2| \lesssim M^{-h(u)}$, where $u = D(t_1, t_2)$ is the youngest common ancestor of t_1 and t_2 in $\mathcal{T}$.

Suppose $\sigma : \mathcal{T} \to \mathcal{S}$ is sticky in the sense of Definition 3.1. Then since sticky mappings preserve heights and lineages, we must have that $h(w) \geq h(u)$ where $w = D(\sigma(t_1), \sigma(t_2))$. Naturally then, we have $|\sigma(t_1) - \sigma(t_2)| \leq M^{-h(w)} \leq M^{-h(u)}$. In this way, we see that if in fact $|t_1 - t_2| \sim M^{-h(u)}$, then the separation in the slopes of $P_{t_1,\sigma}$ and $P_{t_2,\sigma}$ can be no greater than the separation in their roots. Geometrically, what this means is that if, say, $|t_1 - t_2| \sim M^{-H}$, and if the tubes $P_{t_1,\sigma}$, $P_{t_2,\sigma}$ are to intersect, then that intersection must occur about a unit distance from their roots. Such a collection of tubes exhibits the geometric quality of stickiness illustrated in Figure 1.6.

It turns out that this notion of stickiness as a mapping between trees is sufficient to construct Kakeya-type sets over Cantor sets of directions, both in the plane [4] and in general dimensions [31]; see Chapters 5 and 6–11. However, this notion is a bit too restrictive when our set of directions is chosen to simply be sublacunary.

To see why, note that it is often the case that two points may lie close together in Euclidean space, while their M-adic separation as members of a tree is quite large. To take a concrete example, consider the set $\mathbb{D}$ of dyadic rationals on $[0, 1]$ and its corresponding tree encoding $\mathcal{S} = \mathcal{S}(\mathbb{D}; 2)$. Fix $\varepsilon > 0$. Since $\mathbb{D}$ is dense in $[0, 1]$, we may find $d_1, d_2 \in \mathbb{D}$ such that $d_1 \in (\frac{1}{2} - \varepsilon, \frac{1}{2})$ and $d_2 \in (\frac{1}{2}, \frac{1}{2} + \varepsilon)$. Then we see that the dyadic separation between d_1 and d_2 is $2^{-h(D(d_1,d_2))} = 1$, while their Euclidean separation is $|d_1 - d_2| \leq 2\varepsilon$.

Now consider a sticky mapping σ of a collection of root vertices $\mathcal{T} = \mathcal{T}([0, \frac{1}{2}) \cup \{1\}; 2)$ into the slope tree $\mathcal{S}$. Suppose $|t_1 - t_2| \sim \varepsilon$. Then it is impossible to have $\sigma(t_1) = d_1$, $\sigma(t_2) = d_2$, else we would require $h(D(d_1, d_2)) \geq h(D(t_1, t_2))$. But $h(D(t_1, t_2)) \geq 1$, so this is not possible by the definition of d_1 and d_2. Thus, any sticky mapping between $\mathcal{T}$ and $\mathcal{S}$ cannot make the pair of tubes $P_{t_1,\sigma}$, $P_{t_2,\sigma}$ sticky in the sense of Figure 1.6. Of course, the sticky slope assignment σ may generate other sticky pairs of tubes, but the fact remains that we are undoubtedly neglecting many other pairs. Clearly, such an assignment is not optimal, in the sense of exploiting the geometry of stickiness.

The notion of tree stickiness is sufficient to establish the theorem of Bateman and Katz (see Chapter 5), as well as our Theorem 1.2. However, once we turn our attention to Theorem 1.3, we will require a more flexible notion that takes fuller advantage of the geometric idea of sticky tubes. This will motivate us to define the notion of a *weakly sticky* mapping between trees (see Definition 13.6). Such a mapping will turn out to sufficiently capture the geometric essence of stickiness for the purposes of constructing Kakeya-type sets over generic sublacunary sets of directions.

3.5 The tree structure of Cantor-type and sublacunary sets

We now state and prove a key structural result about our sets of interest for Theorem 1.2, the generalized Cantor sets $\mathcal{C}_M$.

Proposition 3.2. *Fix any integer $M \geq 3$. Define $\mathcal{C}_M$ as in Section 1.1. Then*

$$\mathcal{T}(\mathcal{C}_M; M) \cong \mathcal{T}([0,1); 2).$$

That is, the M-adic tree representation of $\mathcal{C}_M$ is isomorphic to the full binary tree, illustrated in Figure 3.1.

Proof. Denote $\mathcal{T} = \mathcal{T}(\mathcal{C}_M; M)$ and $\mathcal{T}' = \mathcal{T}([0,1); 2)$. We must construct a bijective sticky map $\psi : \mathcal{T} \to \mathcal{T}'$. First, define $\psi(v_0) = v_0'$, where v_0 is the root of $\mathcal{T}$ and v_0' is the root of $\mathcal{T}'$.

Now, for any $k \geq 1$, consider the vertex $\langle i_1, i_2, \ldots, i_k \rangle \in \mathcal{T}$. We know that $i_j \in \mathbb{Z}_M$ for all j. Furthermore, for any fixed j, this vertex corresponds to a kth level subinterval of $\mathcal{C}_M^{[k]}$. Every such k-th level interval is replaced by exactly two arbitrary $(k+1)$-th level subintervals in the construction of $\mathcal{C}_M^{[k+1]}$. Therefore, there

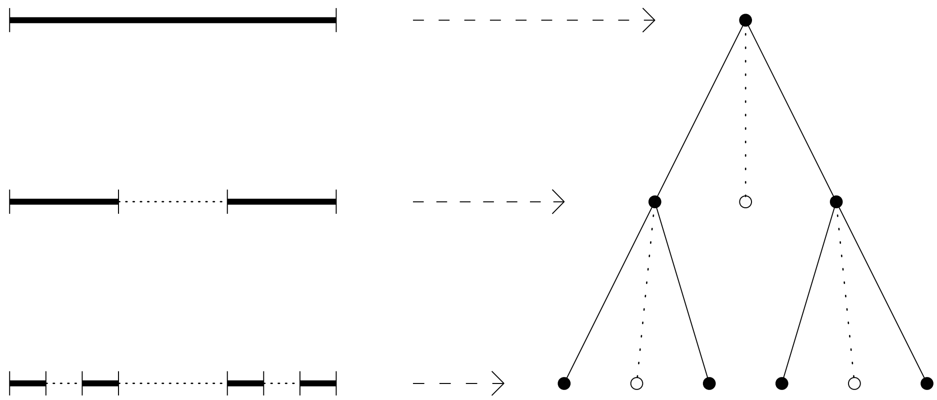

Figure 3.1: A pictorial depiction of the isomorphism between a standard middle-thirds Cantor set and its representation as a full binary subtree of the full base $M = 3$ tree.

exists $N_1 := N_1(\langle i_1, \ldots, i_k \rangle)$, $N_2 := N_2(\langle i_1, \ldots, i_k \rangle) \in \mathbb{Z}_M$, with $N_1 < N_2$, such that $\langle i_1, \ldots, i_k, i_{k+1} \rangle \in \mathcal{T}$ if and only if $i_{k+1} = N_1$ or N_2. Consequently, we define

$$\psi(\langle i_1, i_2, \ldots, i_k \rangle) = \langle l_1, l_2, \ldots, l_k \rangle \in \mathcal{T}', \tag{3.6}$$

where

$$l_{j+1} = \begin{cases} 0 & \text{if } i_{j+1} = N_1(\langle i_1, \ldots, i_j \rangle), \\ 1 & \text{if } i_{j+1} = N_2(\langle i_1, \ldots, i_j \rangle). \end{cases}$$

The mapping ψ is injective by construction and surjectivity follows from the binary selection of subintervals at each stage in the construction of $\mathcal{C}_M$. Moreover, ψ is sticky by (3.6). $\square$

The tree structure of a Cantor-type set is easy to quantify via the above isomorphism. However, as will soon see, the key properties of this structure are not dependent on the heights of the vertices in the tree, but rather upon the *lineages* of those vertices. Imagine we consider some tree $\mathcal{T}$, possibly infinite, and define $\mathcal{D}^* \subseteq \mathcal{T}$ to be the collection of all splitting vertices of $\mathcal{T}$. Then if every vertex $v \in \mathcal{T}$ of height J for some $J > 0$ contains at least N splitting vertices along its lineage, i.e., $\#(\{d \in \mathcal{D}^* : v \subset d\}) \geq N$, the tree $\mathcal{T}_J$ has essentially as rich of structure as the full binary tree of height N. Heuristically, sublacunary sets exhibit this feature for *any* choice of $N > 0$. This idea will be rigorously developed in terms of a quantity called the *splitting number* of a tree in the next section. It plays a critical role in both the proof of our Theorem 1.3 and Bateman's planar analogue [3].

3.6 The splitting number of a tree

There are many ways to quantify the "size" or "spread" of a tree (see [36]). Of these, the concept of a *splitting number* proved to be the most relevant in the planar characterization of directions that admit Kakeya-type sets [3]. Not surprisingly, it will turn out to be equally important for us. As we will see in Section 3.7, this notion actually provides a way of encoding the *lacunarity*, or lack thereof, of a Euclidean or direction set. As such, its importance in establishing the results of Section 1.1 will prove critical.

We say that a vertex $v \in \mathcal{T}$ *splits in* $\mathcal{T}$ if it has at least two children in $\mathcal{T}$. When it is clear to which tree we are referring, we will just say that v *splits*, and we will call v a *splitting vertex.* Define $\text{split}_{\mathcal{T}}(R)$, the *splitting number of a ray R in $\mathcal{T}$* to be the number of splitting vertices in $\mathcal{T}$ along that ray. The *splitting number of a vertex v with respect to a tree $\mathcal{T}$* is defined to be

$$\text{split}_{\mathcal{T}}(v) := \max_{\mathcal{S}_v \subseteq \mathcal{T}} \min_{R_v \in \partial \mathcal{S}_v} \text{split}_{\mathcal{S}_v}(R_v), \tag{3.7}$$

where the maximum is taken over all subtrees $\mathcal{S}_v \subseteq \mathcal{T}$ rooted at v, and the minimum is taken over all rays R_v in $\mathcal{S}_v$ that originate at the vertex v. Finally, the *splitting number of the tree $\mathcal{T}$* is defined as

$$\text{split}(\mathcal{T}) := \max_{v \in \mathcal{T}} \text{split}_{\mathcal{T}}(v). \tag{3.8}$$

To take an easy example, consider the set $\Omega = \{2^{-j} : j \geq 1\}$. Then, as is clear from Figure 3.2 below, we can quickly observe that $\text{split}(\mathcal{T}(\Omega; 2)) = 1$. Similarly, if we consider $\Omega_m = \{\frac{k}{2^m} : 0 \leq k < 2^m\}$, then we have $\text{split}(\mathcal{T}(\Omega_m; 2)) = m$ (see Figure 3.2). Consequently, the tree depicting all dyadic rationals must have infinite splitting number, at least when encoded using a base of 2. In fact, this claim holds true regardless of the base M chosen to encode the tree. Indeed, consider $\Omega_{\kappa m}$ for any integer $\kappa \geq 1$. Then $\text{split}(\mathcal{T}(\Omega_{\kappa m}; 2)) = \kappa m$, while $\text{split}(\mathcal{T}(\Omega_{\kappa m}; 2^\kappa)) = m$. Thus, the tree depicting all dyadic rationals with base $M = 2^\kappa$ must have infinite splitting number. For any base $M \neq 2^\kappa$ for some integer $\kappa \geq 1$, note that M must still be on the order of some $2^{\widetilde{\kappa}}$. Consequently, every M-adic interval on $[0, 1)$ encoded by each vertex of $\mathcal{T}(\Omega_m; M)$ at height j contains and is contained in a fixed number of $2^{\widetilde{\kappa}}$-adic intervals of length $2^{-\widetilde{\kappa} j}$, independent of the height j. This implies $\text{split}(\mathcal{T}(\Omega_m; M)) \sim \text{split}(\mathcal{T}(\Omega_m; 2^{\widetilde{\kappa}}))$.

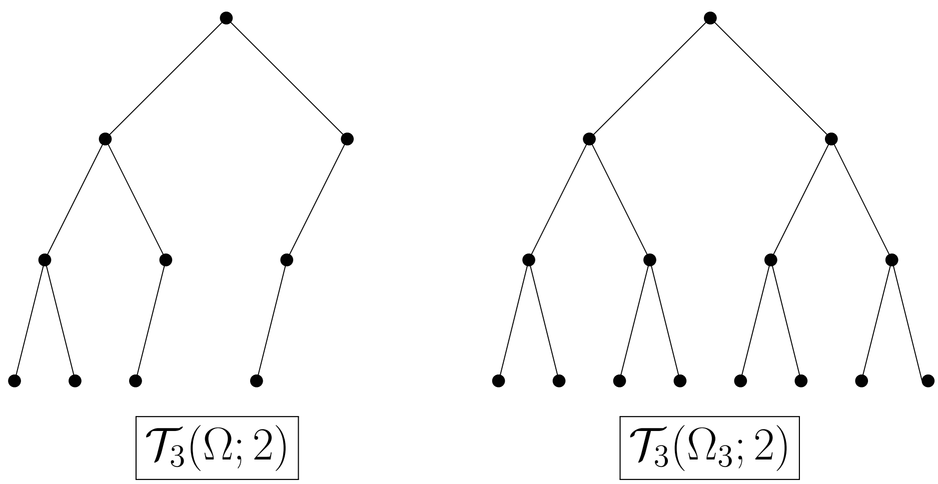

Figure 3.2: A diagram of the trees representing the Euclidean sets $\Omega = \{2^{-j} : j \geq 1\}$ and $\Omega_m = \{\frac{k}{2^m} : 0 \leq k < 2^m\}$ for $m = 3$, up to height 3. Notice that both trees exhibit a kind of self-similar structure, making the calculation of their respective splitting numbers particularly easy. The tree encoding Ω is self-similar with respect to the full subtree rooted at any one of the vertices on the leftmost ray. The tree encoding Ω_m is self-similar with respect to the full subtree rooted at any vertex of the tree.

3.6.1 Preliminary facts about splitting numbers

The calculation of the splitting number of a tree was not particularly difficult for the two special examples above. However, given an arbitrary tree, it is not immediately obvious how to best go about calculating its corresponding splitting number. The following lemmas will provide the necessary facts to simplify this calculation.

Lemma 3.3. *Let $u, v \in \mathcal{T}$ with $u \subseteq v$. Then $\mathit{split}_{\mathcal{T}}(u) \leq \mathit{split}_{\mathcal{T}}(v)$.*

Proof. Let $\mathcal{S}_u$ be a subtree of $\mathcal{T}$ rooted at u. Define $\mathcal{S}_{v\to u}$ to be the union of the tree $\mathcal{S}_u$ with the path in $\mathcal{T}$ connecting v to u. This is a subtree of $\mathcal{T}$ rooted at v. Since v does not split in $\mathcal{S}_{v\to u}$ and there are no splitting vertices in $\mathcal{S}_{v\to u}$ between v and u, we find that for any ray R in $\mathcal{S}_u$,

$$\mathrm{split}_{\mathcal{S}_u}(R) = \mathrm{split}_{\mathcal{S}_{v\to u}}(R_v), \tag{3.9}$$

where R_v is the ray in $\mathcal{S}_{v\to u}$ rooted at v obtained by extending R to v. Conversely, if R_v is a ray in $\mathcal{S}_{v\to u}$, then (3.9) holds for $R = R_v \cap \mathcal{S}_u$. Maximizing over all subtrees $\mathcal{S} \subseteq \mathcal{T}$ rooted at u, we have that

$$\begin{aligned}
\mathrm{split}_{\mathcal{T}}(u) &= \max_{\mathcal{S}_u \subseteq \mathcal{T}} \min_{R \in \partial\mathcal{S}_u} \mathrm{split}_{\mathcal{S}_u}(R) \\
&= \max_{\mathcal{S}_{v\to u} \subseteq \mathcal{T}} \min_{R_v \in \partial\mathcal{S}_{v\to u}} \mathrm{split}_{\mathcal{S}_{v\to u}}(R_v) \\
&\leq \mathrm{split}_{\mathcal{T}}(v).
\end{aligned}$$

The last inequality is a consequence of (3.7), since the class of subtrees of the form $\mathcal{S}_{v\to u}$ is a subcollection of trees rooted at v. □

Lemma 3.3 says that splitting numbers (of vertices) are monotone nonincreasing in lineages. An immediate consequence of this fact is that $\mathrm{split}(\mathcal{T}) = \mathrm{split}_{\mathcal{T}}(v_0)$, where v_0 is the root of $\mathcal{T}$. Our next result says that splitting numbers of trees are also monotonic in an appropriate sense.

Lemma 3.4. *Let $\mathcal{S} \subseteq \mathcal{T}$. Then $\mathit{split}(\mathcal{S}) \leq \mathit{split}(\mathcal{T})$.*

Proof. By Lemma 3.3, $\text{split}(\mathcal{S}) = \text{split}_{\mathcal{S}}(v_0)$, where v_0 is the root of $\mathcal{S}$. Since $v_0 \in \mathcal{S} \subseteq \mathcal{T}$ and any subtree of $\mathcal{S}$ is also a subtree of $\mathcal{T}$, we find that

$$\begin{aligned} \text{split}_{\mathcal{S}}(v_0) &= \max_{\mathcal{S}_{v_0} \subseteq \mathcal{S}} \min_{R_{v_0} \in \partial \mathcal{S}_{v_0}} \text{split}_{S_{v_0}}(R_{v_0}) \\ &\leq \max_{\mathcal{S}_{v_0} \subseteq \mathcal{T}} \min_{R_{v_0} \in \partial \mathcal{S}_{v_0}} \text{split}_{S_{v_0}}(R_{v_0}) \\ &\leq \text{split}_{\mathcal{T}}(v_0) \\ &\leq \text{split}(\mathcal{T}), \end{aligned}$$

where the last two inequalities are implied by (3.7) and (3.8) respectively. Lemma 3.4 follows. □

A feature of trees with finite splitting number, originally observed in [3, Lemma 5], is that all vertices with largest split occur along a single ray. This specialized ray will turn out to be critical in the detection of lacunary limits.

Lemma 3.5. *Let $\mathcal{T}$ be a tree with split($\mathcal{T}$) $= N$. Then there exists a ray R in $\mathcal{T}$ (of finite or infinite length) such that a vertex v lies on R if and only if split$_{\mathcal{T}}(v) = N$, provided the latter collection contains more than one element.*

Proof. We prove by contradiction. Suppose there are two vertices $u, v \in \mathcal{T}$ with $\text{split}_{\mathcal{T}}(u) = \text{split}_{\mathcal{T}}(v) = N$, $u \not\subseteq v$, $v \not\subseteq u$. Then their youngest common ancestor $D(u, v)$ is neither u nor v. By Lemma 3.3, we know that $\text{split}_{\mathcal{T}}(D(u, v)) \geq N$. Since $u \neq v$, the vertex $D(u, v)$ is actually a splitting vertex. Therefore, $\text{split}_{\mathcal{T}}(D(u, v)) \geq N + 1$. But this contradicts the requirement that $\text{split}(\mathcal{T}) = N$, establishing our claim. □

3.6.2 A reformulation of Theorem 1.3

The dichotomy between trees with finite versus infinite splitting number will prove to be our main distinction of interest. Roughly speaking, a tree that has infinite splitting number in some coordinate system must encode a "large" subset of Euclidean space, the threshold of size being determined by sublacunarity. However, the splitting number of a tree encoding a set is sensitive to the coordinates used to represent the set. For example, let U and V be the sets constructed in Section 2.1.2. Then $\text{split}(\mathcal{T}(U \times V); 2) = 2$, while $\text{split}(\mathcal{T}(\varphi(U \times V); 2)) = \infty$ for the coordinate transformation $\varphi(u, v) = (u+v, u-v)$. More strongly, notice that even the finiteness of the splitting number could be affected by the choice. This consideration features prominently in the following proposition which, when taken together with Proposition 3.7, furnishes a restatement of Theorem 1.3 that we will exploit in subsequent chapters.

Our proof of Theorem 1.3 will follow a two-step route.

Proposition 3.6. *Fix a dimension $d \geq 2$ and an integer $M \geq 2$. If a direction set $\Omega \subseteq \mathbb{R}^{d+1} \setminus \{0\}$ is sublacunary (in the sense of Definition 2.7), then*

$$\sup_{\mathbb{V}} \sup_{W_\Omega} \sup_{\varphi} \textit{split}(\mathcal{T}(\varphi(W_\Omega); M)) = \infty. \tag{3.10}$$

Here $\mathbb{V}$ ranges over the collection of all hyperplanes at unit distance from the origin. For a fixed $\mathbb{V}$, the set W_Ω ranges over all relatively compact subsets of $\mathcal{C}_\Omega \cap \mathbb{V}$, and the innermost supremum is taken over all coordinate choices $\varphi = (\mathbf{a}, \mathcal{B})$ on $\mathbb{V}$, where $\mathbf{a} \in \mathbb{V}$ is the point closest to the origin and $\mathcal{B} = \{\mathbf{v}_1, \cdots, \mathbf{v}_d\}$ is any orthonormal basis of $\mathbb{V} - \mathbf{a}$. In other words, φ represents a rotation in $\mathbb{V}$ centred at $\mathbf{a}$, with

$$\varphi(\mathcal{C}_\Omega \cap \mathbb{V}) = \Big\{(x_1, \cdots, x_d) : x = \mathbf{a} + \sum_{j=1}^{d} x_j \mathbf{v}_j \in \mathcal{C}_\Omega \cap \mathbb{V}\Big\}.$$

Thus for every $N \geq 1$, there exists a hyperplane $\mathbb{V}_N$, a relatively compact subset W_N of $\mathcal{C}_\Omega \cap \mathbb{V}_N$, and a coordinate system φ_N on $\mathbb{V}_N$ such that

$$\textit{split}(\mathcal{T}(\varphi_N(W_N); M)) > N. \tag{3.11}$$

Proposition 3.7. *If a direction set Ω obeys (3.10) for some $M \geq 2$, then Ω admits Kakeya-type sets.*

Proposition 3.7 will be the subject of Chapters 15 – 19. We prove Proposition 3.6

in Section 3.7 below.

3.7 Lacunarity on trees

A distinctive feature in the planar characterization of Kakeya-type sets [3] is the observation that the lacunarity of a set is reflected in the structure of its tree. Following the ideas developed there, we recast the concept of finite order lacunarity of a one-dimensional set using the structure of the splitting vertices of its tree. This provides a tool of convenience in the proof of Proposition 3.6, the main objective of this section.

Lemma 3.8. *For any $M \geq 2$, $N \geq 1$, there is a constant $C = C(N, M)$ with the following property. If a relatively compact set $U \subseteq \mathbb{R}$ is such that split$(\mathcal{T}(U; M)) = N$, then U can be covered by the C-fold union of sets in $\Lambda(N, M^{-1})$ as described in Definition 2.2.*

The proof of this lemma will be presented later in this section. Assuming this, the proof of the proposition is completed as follows.

Proof of Proposition 3.6. We prove the contrapositive, starting with the assumption that

$$\sup_{\mathbb{V}} \sup_{W_\Omega} \sup_{\varphi} \operatorname{split}(\mathcal{T}(\varphi(W_\Omega; M)) = N < \infty. \tag{3.12}$$

Fix an arbitrary coordinate system $\varphi = (\mathbf{a}, \mathcal{B})$ of $\mathbb{V}$ and let π_j denote the projection maps defined in (2.2) with respect to this choice. For the remainder of this proof, we will assume that $\mathbb{V}$ is represented in these coordinates, so that π_j may be thought of as the coordinate projections. Let $W = W_\Omega$ be an arbitrary relatively compact subset of $\mathcal{C}_\Omega \cap \mathbb{V}$. Since the tree encoding a set matches that of its closure, we may suppose without loss of generality that $W = W_\Omega$ is compact in $\mathbb{V}$.

For any $1 \leq j \leq d+1$, we create a subset $W_j \subseteq W$ that contains for every $x_j \in \pi_j(W)$ a unique point $x \in W$ with $\pi_j(x) = x_j$. For concreteness, x could be chosen to be minimal in $\pi_j^{-1}(x_j) \cap W$ with respect to the lexicographic ordering. In other words, π_j restricted to W_j is a bijection onto $\pi_j(W)$. We claim that

$$\operatorname{split}(\mathcal{T}(\pi_j(W); M)) \leq \operatorname{split}(\mathcal{T}(W_j; M)). \tag{3.13}$$

Assuming this for the moment, we obtain from the hypothesis (3.12) and Lemma 3.4 that $\operatorname{split}(\mathcal{T}(\pi_j(W); M) \leq \operatorname{split}(\mathcal{T}(W; M)) \leq N$. Applying Lemma 3.8 to $U = \pi_j(W)$, we see that there is a constant C (uniform in $\mathbb{V}$, φ, j and W) such that the projections $\pi_j(W)$ can be covered by the C-fold union of one-dimensional lacunary sets of order $\leq N$ and lacunarity $\leq M^{-1}$. Thus, $W = W_\Omega$ is admissible lacunary of order at most N according to Definition 2.5. Hence Ω is admissible lacunary of finite order as a direction set by Definition 2.7.

It remains to establish (3.13). Any infinite ray $R = R(x_j)$ in $\partial\mathcal{T}(\pi_j(W); M)$ corresponds to a point $x_j \in \pi_j(W)$. Let $R^* = R^*(x) \in \partial\mathcal{T}(W_j; M)$ denote the ray that represents $\pi_j^{-1}(x_j) = x$. This establishes a bijection between the collection of rays in the two trees. Let v_0 and v_0^* denote the roots of the trees $\mathcal{T}(\pi_j(W); M)$ and $\mathcal{T}(W; M)$ respectively, so that $\pi_j(v_0^*) = v_0$. If $\mathcal{S}$ is a subtree of $\mathcal{T}(\pi_j(W); M)$ rooted at v_0, let us denote by $\mathcal{S}^*$ the subtree of $\mathcal{T}(W_j; M)$ rooted at v_0^* generated by all rays R^* such that R is a ray of $\mathcal{S}$. It is clear that if a vertex v on $R(x_j)$ splits in $\mathcal{S}$, then there are two points $x_j \neq x_j'$ in $\pi_j(W)$ lying in distinct children of v. This implies that $x = \pi_j^{-1}(x_j)$ and $x' = \pi_j^{-1}(x_j')$ lie in distinct children of v^*, which denotes the vertex of height $h(v)$ on $R^*(x)$. This makes v^* a splitting vertex of $\mathcal{S}^*$. Thus every splitting vertex of $\mathcal{S}$ lying on R generates a splitting vertex of $\mathcal{S}^*$ lying on R^* at the same height. As a result, $\operatorname{split}_{\mathcal{S}}(R) \leq \operatorname{split}_{\mathcal{S}^*}(R^*)$. Combining these facts with the definition of the splitting number of a tree, we obtain

$$\begin{aligned}
\operatorname{split}(\mathcal{T}(\pi_j(W); M)) &= \max_{\mathcal{S}} \min_{R \in \partial\mathcal{S}} \operatorname{split}_{\mathcal{S}}(R) \\
&\leq \max_{\mathcal{S}^*} \min_{R^* \in \partial\mathcal{S}^*} \operatorname{split}_{\mathcal{S}^*}(R^*) \\
&\leq \operatorname{split}(\mathcal{T}(W_j; M)).
\end{aligned}$$

In view of Lemma 3.3, the maxima in the first and second lines above are taken over all subtrees $\mathcal{S}$ and $\mathcal{S}^*$ rooted at v_0 and v_0^* respectively. This completes the proof of

(3.13) and hence of Proposition 3.6. $\square$

We now turn to the proof of the lemma on which the argument above was predicated.

Proof of Lemma 3.8. We apply induction on N. The base case $N = 1$ will be treated momentarily in Lemma 3.9. Proceeding to the induction step, let R^* denote an infinite ray of the tree $\mathcal{T} = \mathcal{T}(U; M)$ that contains all the vertices $\{v^* : \text{split}_{\mathcal{T}}(v^*) = N\}$. The existence of such a ray has been established in Lemma 3.5. For every vertex v in $\mathcal{T}(U; M)$ which does not lie on R^* but whose parent does, we define a set U_v as follows: $\mathcal{T}_v = \mathcal{T}(U_v; M)$, where $\mathcal{T}_v$ denotes the maximal subtree of $\mathcal{T}$ rooted at v. The definition of the ray R^* dictates that each U_v has the property that $\text{split}(\mathcal{T}(U_v; M)) \leq N - 1$. By the induction hypothesis, there exists a constant $C = C(N-1, M)$ such that each U_v is covered by the C-fold union of sets in $\Lambda(N-1; M^{-1})$. The set U can therefore be covered by the C-fold union of sets $U^{[i]}$, where each $U^{[i]}$ shares a tree structure similar to U: it contains the point identified by R^*, with the additional feature that now $U_v^{[i]} \in \Lambda(N-1; M^{-1})$ for every $v \in \mathcal{V}^{[i]}$, where

$$\mathcal{V}^{[i]} := \left\{ v \in \mathcal{T}(U^{[i]}; M) : v \notin R^* \text{ but parent of } v \text{ is in } R^* \right\}.$$

For every vertex $v \in \mathcal{V}^{[i]}$, let a_v denote the left hand endpoint of the M-adic interval represented by v. The tree encoding the collection of points $A = \{a_v : v \in \mathcal{V}^{[i]}\}$ contains the ray R^*; indeed the only splitting vertices of $\mathcal{T}(A; M)$ lie on R^*. Therefore $\text{split}(\mathcal{T}(A; M)) = 1$. Hence, by Lemma 3.9, A is at most a C-fold union of monotone lacunary sequences with lacunarity M^{-1}, each converging to the point identifying R^*. Let us continue to denote by A one such monotone (say decreasing) sequence. If $a = a_v$ and b are two successive elements of this sequence with $a < b$, then $U^{[i]} \cap [a, b) = U_v^{[i]}$, which is in $\Lambda(N-1; M^{-1})$. Thus $U^{[i]}$ is in $\Lambda(N; M^{-1})$ according to Definition 2.2, completing the proof. $\square$

Lemma 3.9. *Fix $M \geq 2$, and let $A \subseteq \mathbb{R}$ be a relatively compact set with the property that split$(\mathcal{T}(A; M)) = 1$. Then A can be written as the union of at most $6M$ lacunary sequences (defined in Definition 2.1) each with lacunarity constant $\leq M^{-1}$.*

Proof. The argument here closely follows the line of reasoning in [3, Remark 2, page 60]. By Lemma 3.5, there is a ray R^* in $\mathcal{T}(A; M)$ of infinite length such that all the splitting vertices of $\mathcal{T}(A; M)$ lie on it. The ray R^* uniquely identifies a point in $\mathbb{R}$, say $a^* = \alpha(R^*)$. Any ray that is not R^* but is rooted at a vertex of R^* is therefore non-splitting. Thus for every $j = 0, 1, 2, \cdots$ there exists at most $M - 1$ rays R_j in $\mathcal{T}(A; M)$ whose M-adic distance from R^* is j. In other words, if $a_j = \alpha(R_j)$ is the point in A identified by R_j, then there are at most $M - 1$ distinct points $a_j \neq a^*$ such that

$$h(D(a^*, a_j)) = h(D(\alpha(R^*), \alpha(R_j))) = j. \tag{3.14}$$

We define two subsets $A_\pm$ of A, containing respectively points $a \geq a^*$ and $a \leq a^*$. This decomposes $\mathcal{T}(A; M)$ into two subtrees $\mathcal{T}(A_\pm; M)$. Let us focus on $\mathcal{T}(A_+; M)$, the treatment for the other tree being identical. We decompose $\mathcal{T}(A_+; M)$ as the union of at most M trees $\mathcal{T}(A_{i+}; M)$, $i \in \mathbb{Z}_M$, constructed as follows. The tree $\mathcal{T}(A_{i+}; M)$ contains the ray R^*, and for every vertex v in R^* the ray in $\mathcal{T}(A_+; M)$, if any, descended from the ith child of v. In view of the discussion in the preceding paragraph, if there exists an integer j for which a ray R_j in $\mathcal{T}(A_{i+}; M)$ obeys (3.14), then such a ray must be unique.

We now fix $i \in \mathbb{Z}_M$ and proceed to cover A_{i+} by a threefold union of lacunary sequences converging to a^*. Let $\{n_1 < n_2 < \cdots\}$ be the subsequence of integers with the property that $R_j \in \mathcal{T}(A_{i+}; M)$ if and only if $j = n_k$ for some k. The important observation is that if n_{k+2} is a member of this subsequence, then

$$a_{n_k} - a^* \geq \frac{1}{M^{n_{k+2}}}. \tag{3.15}$$

We will return to the proof of this statement in a moment, but a consequence of it and (3.14) is that for any $k \geq 0$ and fixed $\ell = 0, 1, 2$,

$$a_{n_{3(k+1)+\ell}} - a^* \leq M^{-n_{3k+3+\ell}} = M^{-n_{3k+3+\ell}+n_{3k+2+\ell}} M^{-n_{3k+2+\ell}} \leq M^{-1}(a_{n_{3k+\ell}} - a^*).$$

Thus for every fixed $\ell = 0, 1, 2$, the sequence $\mathfrak{A}_\ell = \{a_{n_{3k+\ell}} : k \geq 0\}$ is covered by a lacunary sequence with constant $\leq M^{-1}$ converging to a^*. Since A_{i+} is the union of

$\{\mathfrak{A}_\ell : \ell = 0, 1, 2\}$, the result follows.

It remains to settle (3.15), which is best explained by Figure 3.3. If I_j is the M-

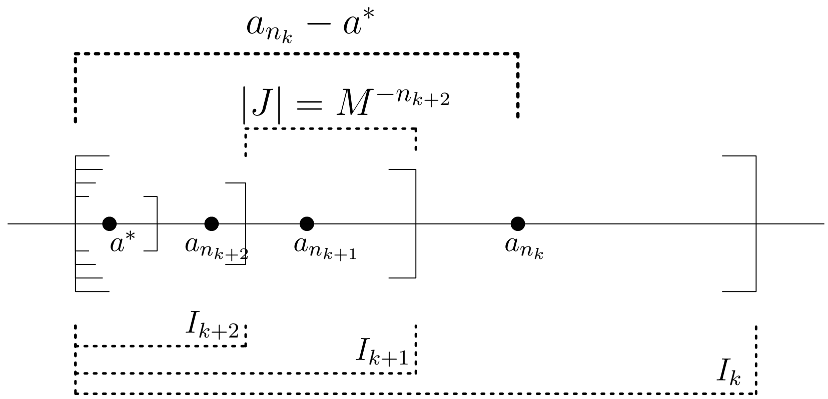

Figure 3.3: A figure explaining inequality (3.15) when $M = 2$ and $n_k = k$.

adic interval of length M^{-n_j} containing a^*, then I_{k+2} cannot share a right endpoint with I_{k+1}, since this would prevent the existence of a point $a_{n_{k+1}} \geq a^*$ obeying (3.14) with $j = n_{k+1}$. Thus a^* (in I_{k+2}) and a_{n_k} (which is to the right of I_{k+1}) must lie on opposite sides of J, the rightmost M-adic subinterval of length $M^{-n_{k+2}}$ in I_{k+1}. This implies $a_{n_k} - a^* \geq |J|$, which is the conclusion of (3.15). $\square$

Chapter 4

Electrical circuits and percolation on trees

As mentioned in Sections 1.4 and 1.5, a key component in the work of Bateman and Katz [4], [3], as well as in the work that forms the backbone of this document [31], [32], is a rather famous estimate from the theory of percolation processes on trees. The literature on this and related percolation items is massive; e.g. see [20], [36]. We will discuss only a very small piece of this vast topic, sufficient for our purposes.

4.1 The percolation process associated to a tree

The special probabilistic process of interest to us is called a bond percolation on trees. Imagine a liquid that is poured on top of some porous material. How will the liquid flow - or *percolate* - through the holes of the material? How likely is it that the liquid will flow from hole to hole in at least one uninterrupted path all the way to the bottom? The first question forms the intuition behind a formal percolation process, whereas the second question turns out to be of critical importance to the proof of Theorems 1.2 and 1.3.

Although it is possible to speak of percolation processes in far more general terms (see [20]), we will only be concerned with a percolation process on a tree. Accordingly, given some tree $\mathcal{T}$ with vertex set $\mathcal{V}$ and edge set $\mathcal{E}$, we define an *edge-dependent Bernoulli (bond) percolation process* to be any collection of random variables $\{X_e : e \in \mathcal{E}\}$, where X_e is Bernoulli(p_e) with $p_e < 1$. The parameter p_e is called the *survival probability of the edge e*. We will always be concerned with a particular type of percolation on our trees: we define a *standard Bernoulli*(p) *percolation* to be one where the random variables $\{X_e : e \in \mathcal{E}\}$ are mutually independent and identically distributed Bernoulli(p) random variables, for some $p < 1$. In fact, for our purposes, it will suffice to consider only standard Bernoulli($\frac{1}{2}$) percolations.

Rather than imagining a tree with a percolation process as the behaviour of a liquid acted upon by gravity in a porous material, it will be useful to think of the percolation process as acting more directly on the mathematical object of the tree itself. Given some percolation process on a tree $\mathcal{T}$, we will think of the event $\{X_e = 0\}$ as the event that we *remove* the edge e from the edge set $\mathcal{E}$, and the event $\{X_e = 1\}$ as the event that we *retain* this edge; denote the random set of retained edges by $\mathcal{E}^*$. Notice that with this interpretation, after percolation there is no guarantee that $\mathcal{E}^*$, the subset of edges that remain after percolation, defines a subtree of $\mathcal{T}$. In fact, it can be quite likely that the subgraph that remains after percolation is a union of many disconnected subgraphs of $\mathcal{T}$.

For a given edge $e \in \mathcal{E}$, we think of $p = \Pr(X_e = 1)$ as the probability that we retain this edge after percolation. The probability that at least one uninterrupted path remains from the root of the tree to its bottommost level is given by the *survival probability* of the corresponding percolation process. More explicitly, given a percolation on a tree $\mathcal{T}$, the survival probability after percolation is the probability that the random variables associated to all edges of at least one ray in $\mathcal{T}$ take the value 1; i.e.

$$\Pr\left(\text{survival after percolation on } \mathcal{T}\right) := \Pr\left(\bigcup_{R \in \partial\mathcal{T}} \bigcap_{e \in \mathcal{E} \cap R} \{X_e = 1\}\right). \tag{4.1}$$

Estimation of this probability is what will be utilized in the proofs of Theorems 1.2 and 1.3. This estimation will require reimagining a tree as an electrical network.

4.2 Trees as electrical networks

Formally, an electrical network is a particular kind of weighted graph. The weights of the edges are called *conductances* and their reciprocals are called *resistances*. In his seminal works on the subject, Lyons visualizes percolation on a tree as a certain electrical network. In [34], he lays the groundwork for this correspondence. While his results hold in great generality, we describe his results in the context of standard Bernoulli percolation on a locally finite, rooted labelled tree only.

A percolation process on the truncation of any given tree $\mathcal{T}$ is naturally associated to a particular electrical network. To see this, we truncate the tree $\mathcal{T}$ at height N and place the positive node of a battery at the root of $\mathcal{T}_N$. Then, for every ray in $\partial\mathcal{T}_N$, there is a unique terminating vertex; we connect each of these vertices to the negative node of the battery. A resistor is placed on every edge e of $\mathcal{T}_N$ with resistance R_e defined by

$$\frac{1}{R_e} = \frac{1}{1-p_e} \prod_{\emptyset \supset v(e') \supseteq v(e)} p_{e'}. \tag{4.2}$$

Notice that the resistance for the edge e is essentially the reciprocal of the probability that a path remains from the root of the tree to the vertex $v(e)$ after percolation. For standard Bernoulli($\frac{1}{2}$) percolation, we have

$$R_e = 2^{h(v(e))-1}. \tag{4.3}$$

One fact that will prove useful for us later is that connecting any two vertices at a given height by an ideal conductor (i.e. one with zero resistance) only decreases the overall resistance of the circuit. This will allow us to more easily estimate the total resistance of a generic tree.

Proposition 4.1. *Let $\mathcal{T}_N$ be a truncated tree of height N with corresponding electrical network generated by a standard Bernoulli($\frac{1}{2}$) percolation process. Suppose at height $k < N$ we connect two vertices by a conductor with zero resistance. Then the resulting electrical network has a total resistance no greater than that of the original network.*

Proof. Let u and v be the two vertices at height k that we will connect with an ideal conductor. Let R_1 denote the resistance between u and $D(u, v)$, the youngest common ancestor of u and v; let R_2 denote the resistance between v and $D(u, v)$. Let R_3 denote the total resistance of the subtree of $\mathcal{T}_N$ generated by the root u and let R_4 denote the total resistance of the subtree of $\mathcal{T}_N$ generated by the root v. These four connections define a subnetwork of our tree, depicted in Figure 4.1(a). The connection of u and v by an ideal conductor, as pictured in Figure 4.1(b), can only change the total resistance of this subnetwork, as that action leaves all other connections unaltered. It therefore suffices to prove that the total resistance of the subnetwork comprised of the resistors R_1, R_2, R_3 and R_4 can only decrease if u and v are joined by an ideal conductor.

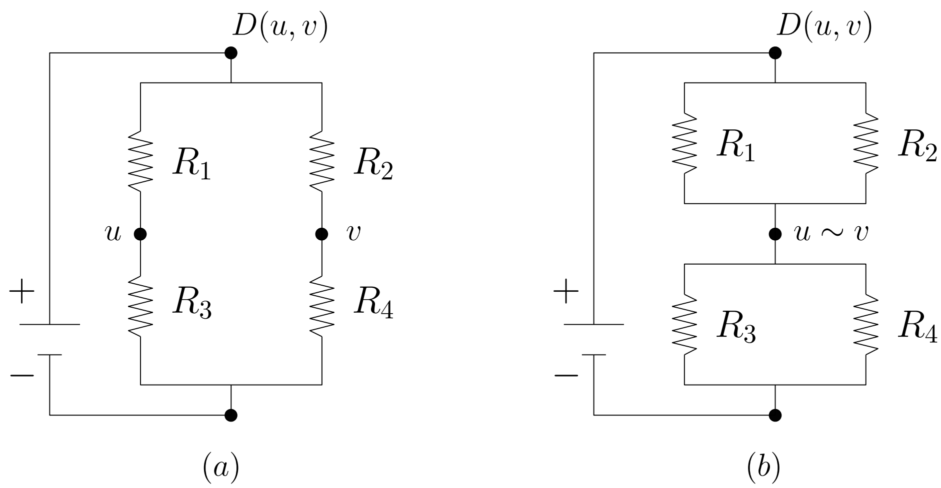

Figure 4.1: (a) The original subnetwork with the resistors R_1, R_3 and R_2, R_4 in series; (b) the new subnetwork obtained by connecting vertices u and v by an ideal conductor.

In the original subnetwork, the resistors R_1 and R_3 are in series, as are the resistors R_2 and R_4. These pairs of resistors are also in parallel with each other. Thus, we calculate the total resistance of this subnetwork, R_{original}:

$$R_{\text{original}} = \left(\frac{1}{R_1 + R_3} + \frac{1}{R_2 + R_4} \right)^{-1}$$

$$= \frac{(R_1+R_3)(R_2+R_4)}{R_1+R_2+R_3+R_4}. \tag{4.4}$$

After connecting vertices u and v by an ideal conductor, the structure of our subnetwork is inverted as follows. The resistors R_1 and R_2 are in parallel, as are the resistors R_3 and R_4, and these pairs of resistors are also in series with each other. Therefore, we calculate the new total resistance of this subnetwork, R_{new}, as

$$\begin{aligned} R_{\text{new}} &= \left(\frac{1}{R_1}+\frac{1}{R_2}\right)^{-1} + \left(\frac{1}{R_3}+\frac{1}{R_4}\right)^{-1} \\ &= \frac{R_1R_2(R_3+R_4)+R_3R_4(R_1+R_2)}{(R_1+R_2)(R_3+R_4)}. \end{aligned} \tag{4.5}$$

We claim that (4.4) is greater than or equal to (4.5). To see this, simply cross-multiply these expressions. After cancellation of common terms, our claim reduces to

$$R_1^2R_4^2 + R_2^2R_3^2 \geq 2R_1R_2R_3R_4.$$

But this is trivially satisfied since $(a-b)^2 \geq 0$ for any real numbers a and b. □

The main consequence of this observation that we draw upon in Lemmas 15.4 and 11.3 is given by the following corollary.

Corollary 4.2. *Given a subtree $\mathcal{T}_N$ of height N contained in the full d-dimensional M-adic tree, let $R(\mathcal{T}_N)$ denote the total resistance of the electrical network that corresponds to standard Bernoulli$(\frac{1}{2})$ percolation on this tree, in the sense of the theorem of Lyons as given in Theorem 4.3. Then*

$$R(\mathcal{T}_N) \geq \sum_{k=1}^{N} \frac{2^{k-1}}{n_k}, \tag{4.6}$$

where n_k denote the number of its kth generation vertices in $\mathcal{T}_N$.

Proof. To show this, we construct an auxiliary electrical network from the one naturally associated to our tree $\mathcal{T}_N$, as follows. For every $k \geq 1$, we connect all vertices at height k by an ideal conductor to make one node V_k. Call this new circuit E.

The resistance of E cannot be greater than the resistance of the original circuit, by Proposition 4.1.

Fix k, $1 \le k \le N$, and let R_k denote the resistance in E between V_{k-1} and V_k. The number of edges between V_{k-1} and V_k is equal to the number n_k of kth generation vertices in $\mathcal{T}_N$, and each edge is endowed with resistance 2^{k-1} by (4.3). Since these resistors are in parallel, we obtain

$$\frac{1}{R_k} = \sum_{1}^{n_k} \frac{1}{2^{k-1}} = \frac{n_k}{2^{k-1}}.$$

This holds for every $1 \le k \le N$. Since the resistors $\{R_k\}_{k=1}^{N}$ are in series, $R(\mathcal{T}_N) \ge R(E) = \sum_{k=1}^{N} R_k$, establishing inequality (11.8). $\square$

4.3 Estimating the survival probability after percolation

We now present Lyons' pivotal result linking the total resistance of an electrical network and the survival probability under the associated percolation process.

Theorem 4.3 (Lyons, Theorem 2.1 of [35])**.** *Let $\mathcal{T}$ be a tree with mutually associated percolation process and electrical network, and let $R(\mathcal{T})$ denote the total resistance of this network. If the percolation is Bernoulli, then*

$$\frac{1}{1+R(\mathcal{T})} \le Pr(\mathcal{T}) \le \frac{2}{1+R(\mathcal{T})},$$

where $Pr(\mathcal{T})$ denotes the survival probability after percolation on $\mathcal{T}$.

We will not require the full strength of this theorem. A reasonable upper bound on the survival probability coupled with the result of Proposition 4.1 will suffice for our applications. The sufficient simpler version of Theorem 4.3 that we state and prove below was essentially formulated by Bateman and Katz [4].

Proposition 4.4. *Let $M \geq 2$ and let $\mathcal{T}$ be a subtree of a full M-adic tree. Let $R(\mathcal{T})$ and $Pr(\mathcal{T})$ be as in Theorem 4.3. Then under Bernoulli percolation, we have*

$$Pr(\mathcal{T}) \leq \frac{2}{1+R(\mathcal{T})}. \tag{4.7}$$

Proof. We will only focus on the case when $R(\mathcal{T}) \geq 1$, since otherwise (4.7) holds trivially. We prove this by induction on the height of the tree N. When $N = 0$, then (4.7) is trivially satisfied. Now suppose that up to height $N-1$, we have

$$\Pr(\mathcal{T}) \leq \frac{2}{1+R(\mathcal{T})}.$$

Suppose $\mathcal{T}$ is of height N. We can view the tree $\mathcal{T}$ as its root together with at most M edges connecting the root to the subtrees $\mathcal{T}_1, \ldots, \mathcal{T}_M$ of height $N-1$ generated by the terminating vertices of these edges. If there are $k < M$ edges originating from the root, then we take $M-k$ of these subtrees to be empty. Note that by the induction hypothesis, (4.7) holds for each $\mathcal{T}_j$. To simplify notation, we denote

$$\Pr(\mathcal{T}_j) = P_j \text{ and } R(\mathcal{T}_j) = R_j,$$

taking $P_j = 0$ and $R_j = \infty$ if $\mathcal{T}_j$ is empty.

Using independence and recasting $\Pr(\mathcal{T})$ as one minus the probability of *not* surviving after percolation on $\mathcal{T}$, we have the formula:

$$\Pr(\mathcal{T}) = 1 - \prod_{k=1}^{M} \left(1 - \frac{1}{2}P_k\right).$$

Note that the function $F(x_1, \ldots, x_M) = 1 - (1-x_1/2)(1-x_2/2)\cdots(1-x_M/2)$ is monotone increasing in each variable on $[0,2]^M$. Now define

$$Q_j := \frac{2}{1+R_j}.$$

Since resistances are nonnegative, we know that $Q_j \leq 2$ for all j. Therefore,

$$\begin{aligned}\Pr(\mathcal{T}) &= F(P_1, \ldots, P_M) \\ &\leq F(Q_1, \ldots, Q_M) \\ &\leq \frac{1}{2} \sum_{k=1}^{M} Q_k.\end{aligned}$$

Here, the first inequality follows by monotonicity and the induction hypothesis. Plugging in the definition of Q_k, we find that

$$\Pr(\mathcal{T}) \leq \sum_{k=1}^{M} \frac{1}{1+R_k}.$$

But since each resistor R_j is in parallel, we know that

$$\frac{1}{R(\mathcal{T})} = \sum_{k=1}^{M} \frac{1}{1+R_k}.$$

Combining this formula with the previous inequality and recalling that $R(\mathcal{T}) \geq 1$, we have

$$\Pr(\mathcal{T}) \leq \frac{1}{R(\mathcal{T})} \leq \frac{2}{1+R(\mathcal{T})},$$

as required.

$\square$

Chapter 5

Kakeya-type sets over a Cantor set of directions in the plane

The construction of Kakeya-type sets over a Cantor set of directions $\Omega = \mathcal{C}_M$ is considerably easier than the general case, due to the very particular and fixed structure of the corresponding slope tree, $\mathcal{T}(\mathcal{C}_M; M)$, as established in Proposition 3.2. This statement holds true regardless of dimension considered, and as such we will begin our discussion in this simplified setting. Starting with the work of Bateman and Katz [4], we will outline their construction in the plane and point out the additional challenges that appear in the higher dimensional case. We will proceed into a discussion of the added difficulties that arise when the set of directions is generalized to a sublacunary set. In Chapter 6, we will begin the details of our general construction in any dimension. The framework of the analysis in all subsequent chapters will have its foundation in the arguments of the present chapter.

5.1 The result of Bateman and Katz

In [4], Bateman and Katz consider what happens when thin tubes in the plane are assigned directions that arise in a natural way from the standard middle-third Cantor set on $[0, 1)$, denoted here by $\mathcal{C}$. For every $n \geq 1$, they define a random sticky mapping

$$\sigma_n : \mathcal{T}_n([0,1);3) \to \mathcal{T}_n(\mathcal{C};3). \tag{5.1}$$

Recall that such a mapping between trees is set to preserve heights and lineages, per Definition 3.1. Note that as an immediate consequence of Proposition 3.2, we have

$$\mathcal{T}_n(\mathcal{C};3) \cong \mathcal{T}_n([0,1);2) \text{ for all } n \geq 1. \tag{5.2}$$

The random mechanism of the mapping σ_n then assigns one edge in $\mathcal{T}_n([0,1);3)$ to one edge in $\mathcal{T}_n(\mathcal{C};3)$, independently and with equal probability. The collection of all sticky mappings is thus uniformly distributed on the edge set of $\mathcal{T}_n([0,1);3)$.

Now fix $n \geq 1$; for each $t \in \mathcal{T}_n([0,1);3)$, we define a random tube P_{t,σ_n} in $\mathbb{R}^2$ with principal axis given by the line segment from $(0,t)$ to $(1, t+\sigma_n(t))$ and cross-sectional diameter $3^{-(n+1)}$. Here, as in the sequel, we abuse notation slightly and identify $t \in \mathcal{T}_n([0,1);3)$ and $\sigma_n(t) \in \mathcal{T}_n(\mathcal{C};3)$ with their naturally associated real numbers on $[0,1)$, via (3.2) and (3.3). Consequently, we consider the random collection of tubes

$$K_{\sigma_n} := \bigcup_{\substack{t \in \mathcal{T}_n([0,1);3) \\ h(t)=n}} P_{t,\sigma_n}, \tag{5.3}$$

with the slope of each tube a member of the standard middle-third Cantor set at its nth stage of construction.

Bateman and Katz prove the following two estimates on the random sets K_{σ_n}:

$$\text{for any sticky map } \sigma_n,\ |K_{\sigma_n}| \gtrsim \frac{\log n}{n}, \tag{5.4}$$

and

$$\text{there exists a sticky } \sigma_n \text{ such that } \left| K_{\sigma_n} \cap \left(\left[\frac{1}{3}, 1\right] \times \mathbb{R} \right) \right| \lesssim \frac{1}{n}. \tag{5.5}$$

Notice that these estimates provide precisely the Kakeya-type set condition (1.1) for any $n \geq 1$. More precisely, in the notation of Definition 1.1, we set $K_{\sigma_n} = E_n^*(A_0)$ and $K_{\sigma_n} \cap \left[\frac{1}{3}, \frac{2}{3}\right] = E_n$, where $A_0 = 3$. (Note that the tubes P_{t,σ_n} just defined in (5.3) are not the same tubes $P_t^{(n)}$ referred to in Definition 1.1.) Then the estimates

(5.4) and (5.5) yield

$$\lim_{n\to\infty} \frac{|E_n^*(A_0)|}{|E_N|} = \infty,$$

which is the requirement of (1.1). Consequently, we conclude that the set of directions $\Omega = \{\omega \in \mathbb{S}^1 : \tan\omega \in \mathcal{C}\}$ admits Kakeya-type sets.

5.1.1 Proof of inequality (5.4)

Estimate (5.4) is proven by combining a general observation about intersections of like sets with the Córdoba estimate (1.5) applied on a proper partition of the random set K_{σ_n}. The estimate (5.5) follows from a probabilistic argument that relies on Lyons' theorem about percolation on trees, (see Theorem 4.3). We will sketch these proofs in a bit more detail, as they will provide the analytical foundation for the more general cases to follow. Since n is fixed in what follows, to simplify notation we will write σ in place of σ_n; we will also write $\mathcal{T}_n$ in place of $\mathcal{T}_n([0,1);3)$ and $\mathcal{C}_n$ instead of $\mathcal{T}_n(\mathcal{C};3)$.

For the lower bound estimate (5.4), we begin with the following lemma.

Lemma 5.1 (Bateman and Katz)**.** *Suppose* $(X, \mathcal{M}, \mu)$ *is a measure space and* $A_1, \ldots, A_m$ *are sets with* $\mu(A_i) = \alpha$ *for every* i*. Let* $L > 0$*, and suppose that*

$$\sum_{j=1}^{m}\sum_{i=1}^{m} \mu(A_i \cap A_j) \le L.$$

Then

$$\mu\left(\bigcup_{i=1}^{m} A_i\right) \ge \frac{m^2\alpha^2}{16L}.$$

Proof. By the pigeonhole principle, there exists a set $M \subset \{1, \ldots, m\}$ with $\#(M) \ge \frac{m}{2}$ such that whenever $i \in M$, we have

$$\sum_{j=1}^{m} \mu(A_i \cap A_j) \le \frac{2L}{m}.$$

For any such i, we rewrite this inequality as

$$\frac{1}{\alpha}\int_{A_i}\sum_{j=1}^{m}\mathbf{1}_{A_j}(x)d\mu(x)\leq\frac{2L}{m\alpha}.$$

Since $\sum \mathbf{1}_{A_j}(x)$ is a nonnegative function, the only way its average value over A_i can be bounded by $\frac{2L}{m\alpha}$ is if there exists a subset $B_i\subseteq A_i$ with $\mu(B_i)\geq\frac{\alpha}{2}$ such that for every $x\in B_i$,

$$\sum_{j=1}^{m}\mathbf{1}_{A_j}(x)\leq\frac{4L}{m\alpha}.$$

Now we divide through by the quantity on the right-hand side and integrate over $\bigcup_{i\in M} B_i$:

$$\frac{m\alpha}{4L}\int_{\bigcup_{i\in M}B_i}\sum_{j=1}^{m}\mathbf{1}_{A_j}(x)d\mu(x)\leq\mu\left(\bigcup_{i\in M}B_i\right).$$

Remembering that $B_i\subseteq A_i$ for all $i\in M$, and interchanging the sum and the integral on the left-hand side, we conclude

$$\mu\left(\bigcup_{i=1}^{m}A_i\right)\geq\frac{m\alpha}{4L}\sum_{i\in M}\mu(B_i)\geq\frac{m^2\alpha^2}{16L},$$

using the trivial bounds $\mu(B_i)\geq\frac{\alpha}{2}$ and $\#(M)\geq\frac{m}{2}$. □

Proof of (5.4). For $0\leq j<\frac{1}{2}\log n$, we define $P_{t,\sigma,j}:=P_{t,\sigma}\cap[3^{-j},3^{-(j-1)}]$, and note that $|P_{t,\sigma,j}|\sim 3^{-(j+n)}$. Then the collection $\{P_{t,\sigma,j}\}_j$ is disjoint over j. In light of Lemma 5.1, we would like to show that the measure of the union of these pieces is smaller than some appropriate quantity for all $j<\frac{1}{2}\log n$. Consequently, it suffices to show that

$$\sum_{\substack{t_1\in\mathcal{T}_n\\h(t_1)=n}}\sum_{\substack{t_2\in\mathcal{T}_n\\h(t_2)=n}}|P_{t_1,\sigma,j}\cap P_{t_2,\sigma,j}|\lesssim\frac{n}{3^{2j}}.\tag{5.6}$$

The height restrictions $h(t_1)=h(t_2)=n$ will hold throughout this section, but we will henceforth suppress this in the notation to save on clutter. Once we have (5.6) for every $0 \le j < \frac{1}{2}\log n$, using Lemma 5.1, we calculate

$$\left|\bigcup_{t\in\mathcal{T}_n} P_{t,\sigma,j}\right| \gtrsim \frac{1}{n}$$

for each $j < \frac{1}{2}\log n$. Thus,

$$\left|K_\sigma \cap [n^{-c}, 1)\right| \gtrsim \frac{\log n}{n},$$

which implies (5.4).

Notice that the diagonal term of the sum in (5.6) satisfies the proper bound since $j < \frac{1}{2}\log n$, and that in fact this is the best we can do. Now if $t_1 \neq t_2$ and $P_{t_1,\sigma,j} \cap P_{t_2,\sigma,j} \neq \emptyset$, then

$$|\sigma(t_1) - \sigma(t_2)| \gtrsim 3^j |t_1 - t_2|, \tag{5.7}$$

since $\theta \gtrsim \sin\theta$ for θ small. Pairing this bound on the difference in the slopes $\sigma(t_1)$ and $\sigma(t_2)$ with the Córdoba estimate (1.5), we have

$$|P_{t_1,\sigma,j} \cap P_{t_2,\sigma,j}| \lesssim \frac{1}{3^{2n+j}|t_1 - t_2|}. \tag{5.8}$$

Recall that the notation $D(t_1, t_2)$ represents the youngest common ancestor of t_1 and t_2 in the tree $\mathcal{T}_n$. Now, the stickiness of σ and (5.7) imply that

$$|I_{D(t_1,t_2)}| \gtrsim 3^j |t_1 - t_2|, \tag{5.9}$$

where I_u denotes the unique triadic interval corresponding to the vertex u (see Section 3.2). We will shortly prove the following useful counting lemma.

Lemma 5.2. *Let $t_1, t_2 \in \mathcal{T}_n([0,1);3)$ and set $u = D(t_1, t_2)$. Suppose $|I_u| \gtrsim 3^j|t_1 - t_2|$, and define the set $A_{k,j}(u) := \{(t_1,t_2) \in D^{-1}(u) : |t_1 - t_2| \sim 3^{-j-k}|I_u|\}$. Then*

$$\#(A_{k,j}(u)) \lesssim 3^{2n-2j-2k-2h(u)}.$$

Note that by definition each pair (t_1, t_2) belongs to exactly one $A_{k,j}(u)$. Thus, after plugging in the bound (5.8), we rewrite the off-diagonal part of the sum in (5.6) as

$$\begin{aligned}
\sum_{t_1 \in \mathcal{T}_n} \sum_{\substack{t_2 \in \mathcal{T}_n \\ t_2 \neq t_1}} |P_{t_1,\sigma,j} \cap P_{t_2,\sigma,j}| &\lesssim \sum_{u \in \mathcal{T}_n} \sum_{\substack{(t_1,t_2) \in D^{-1}(u) \\ 3^j|t_1 - t_2| \lesssim |I_u|}} \frac{1}{3^{2n+j}|t_1 - t_2|} \\
&\lesssim \sum_{u \in \mathcal{T}_n} \sum_{k \geq 0} \sum_{(t_1,t_2) \in A_{k,j}(u)} \frac{1}{3^{2n+j}|t_1 - t_2|} \\
&\lesssim \sum_{u \in \mathcal{T}_n} \sum_{k \geq 0} \sum_{(t_1,t_2) \in A_{k,j}(u)} \frac{1}{3^{2n-k}|I_u|}.
\end{aligned}$$

Using Lemma 5.2, we complete the estimation as

$$\sum_{u \in \mathcal{T}_n} \sum_{k \geq 0} \frac{3^{2n-2j-2k-2h(u)}}{3^{2n-k}|I_u|} \lesssim \sum_{u \in \mathcal{T}_n} \sum_{k \geq 0} \frac{3^{-k}|I_u|}{3^{2j}} \lesssim \frac{n}{3^{2j}},$$

where we have used the fact that 3^{-k} is summable, and that the intervals I_u form a partition of the unit interval over all u of fixed height; since $\mathcal{T}_n$ has $n + 1$ distinct heights, (5.6) follows. □

It remains to prove Lemma 5.2. We will accomplish this by counting the number of permissible vertices t_2 at height n after fixing a vertex t_1. Then we will count the number of t_1.

Proof of Lemma 5.2. If $(t_1, t_2) \in A_{k,j}(u)$, then the lineages of t_1 and t_2 split at height $h(u)$, yet remain close enough in the tree $\mathcal{T}_n([0, 1); 3)$ so that $|t_1 - t_2| \sim 3^{-j-k-h(u)}$. To count the number of permissible pairs (t_1, t_2), let u_l denote the lth child of u. We will restrict attention to the pairs (t_1, t_2) that are descendants of u_1 and u_2 respectively. Counting over all options will introduce no more than a constant factor of $\binom{3}{2}$ to the final count.

Fix $t_1 \subseteq u_1$. Notice that the only permissible t_2 now lie in the intersection of the interval associated to u_2 with an interval of length $3^{-j-k-h(u)}$ containing t_1. Thus, to bound the number of t_2 with such a t_1 fixed, it suffices to simply 3^{-n}-separate the interval of length $3^{-j-k-h(u)}$. Similarly, if we now let t_1 vary, t_1 can only be chosen to lie no more than a distance of $\lesssim 3^{-j-k-h(u)}$ from the triadic interval associated to u_2; hence, we again 3^{-n}-separate the interval of length $3^{-j-k-h(u)}$ to bound the number of t_1. Thus, we have the size estimate $\#(A_{k,j}(u)) \lesssim 3^{2n-2j-2k-2h(u)}$ as claimed. $\square$

The lower bound (5.4) holds for any sticky collection of tubes, but notice that the only time we require the fact that σ is sticky is for the estimate (5.9). Without such an estimate, there would be no restriction placed on the height of the youngest common ancestor of the two points t_1 and t_2 that generate intersecting tubes. As in our generic discussion in Section 1.5, stickiness is a property that provides a quantitative connection between the location of tubes in space and their orientations. By requiring $P_{t_1,\sigma,j} \cap P_{t_2,\sigma,j} \neq \emptyset$ in the proof above, we are fixing a region in space where these tubes can overlap. This naturally places a restriction on the permissible slopes that can be assigned to any two base points t_1 and t_2 in order for a nonempty intersection to occur in the $[3^{-j}, 3^{-(j-1)}] \times \mathbb{R}$ strip. But it is our enforced notion of stickiness, formulated in the context of a mapping between trees, that provides the necessary quantitative information to effectively exploit the geometry and establish (5.4). This will prove to be sufficient for the purposes of our Theorem 1.2, but Theorem 1.3 will require a more flexible notion of stickiness between root and slope trees to yield to these methods.

5.1.2 Proof of inequality (5.5)

In contrast to (5.4), it is quite clear that estimate (5.5) cannot be expected to hold for an arbitrary sticky σ; for example, consider the sticky map that assigns to every base point $t \in \mathcal{T}_n$ the same slope, say $0 \in \mathcal{C}_n$. However, the percolation argument exploited by Bateman and Katz shows that (5.5) does hold for a typical sticky σ, as we shall now see.

Proof of (5.5). Fix a point $(x, y) \in (\frac{1}{3}, 1) \times \mathbb{R}$. We begin by observing that for every base point $t \in \mathcal{T}_n$, there is at most one slope $c_{(x,y)}(t) \in \mathcal{C}_n$ such that $(x, y) \in P_{t,c_{(x,y)}(t)}$ (here, we have abused notation slightly and let $P_{t,c}$ represent the tube with principal axis given by the line segment $(0,t)$ to $(1,t+c)$ with cross-sectional diameter $3^{-(n+1)}$). Indeed, suppose two distinct slopes $c_1 \neq c_2$ existed such that $(x,y) \in P_{t,c_1} \cap P_{t,c_2}$. Then there must exist $a_1, a_2 \in \mathbb{R}$ with $|a_1|, |a_2| \leq \frac{1}{2} \cdot 3^{-(n+1)}$, the cross-sectional radius of a tube, such that

$$y = t + a_1 + c_1 x = t + a_2 + c_2 x.$$

But since $|c_2 - c_1| \geq 3^{-n}$ by construction, it must follow that $x \leq \frac{1}{3}$, a contradiction.

Let $\text{Poss}(x,y)$ denote the subtree of $\mathcal{T}_n$ where $t \in \text{Poss}(x,y)$ if and only if $c_{(x,y)}(t)$ exists. This is a deterministic object, dependent only on the choice of (x,y). Now notice that the event $(x,y) \in K_\sigma$ happens only if $\sigma(t) = c_{(x,y)}(t)$ for some $t \in \text{Poss}_n(x,y)$; i.e. for some $t \in \text{Poss}(x,y)$ with $h(t) = n$. Equivalently, this can only happen if $\sigma(t_k) = c_{(x,y)}(t_k)$ for every ancestor t_k of $t \in \text{Poss}_n(x,y)$, $0 \leq k \leq n$. Recall the definition of σ as given by (5.1). By the structure of the slope tree given in (5.2), we see that defining a random sticky map σ is equivalent to defining a collection of $3^{n+1} - 1$ independent Bernoulli($\frac{1}{2}$) random variables $\{X_e : e \in \mathcal{E}(\mathcal{T}_n)\}$, one corresponding to each edge in the tree $\mathcal{T}_n$. More precisely, by (5.2), and in order to preserve stickiness, each edge $e \in \mathcal{E}(\mathcal{T}_n)$ can be mapped to only one of two possible edges in $\mathcal{C}_n$. We declare that e maps to the first of these two possibilities if $X_e = 0$, and that e maps to the second otherwise.

Now, if $(x,y) \in K_\sigma$, then we have already seen that there must exist a $t \in \text{Poss}_n(x,y)$ so that $\sigma(t_k) = c_{(x,y)}(t_k)$ for every ancestor $t_k \supseteq t$, $0 \leq k \leq n$. As in Section 4.1, we will equate the event $\{X_e = 0\}$ with the action of removing the edge e from the edge set $\mathcal{E}(\mathcal{T}_n)$, and the event $\{X_e = 1\}$ with the action of retaining this edge. Thus, we see that the event $(x,y) \in K_\sigma$ occurs only if a path from the root to the bottommost level remains after percolation on the tree $\text{Poss}(x,y)$. Letting $\Pr(\text{Poss}(x,y))$ denote the probability of this event, Lyons' Theorem 4.3 tells us that

$$\Pr(\text{Poss}(x,y)) \leq \frac{2}{1 + R(\text{Poss}(x,y))}, \tag{5.10}$$

where $R(\text{Poss}(x,y))$ is the resistance of the electrical network associated to the tree $\text{Poss}(x,y)$, as defined in Section 4.2.

We claim that $R(\mathrm{Poss}(x,y)) \gtrsim n$. Notice that the number of vertices at height k in $\mathrm{Poss}(x,y)$ is bounded above by $\lesssim 2^k$. Indeed, there are exactly 2^k different vertices at height k in the slope tree $\mathcal{C}_n$; thus, using a similar argument as the one presented in the first paragraph of this proof, for any fixed $c \in \mathcal{C}_n$ at height k, there are a finite number of $t \in \mathcal{T}_n$ of height k that could possibly yield $c_{(x,y)}(t) = c$.

Now consider the electrical network associated to $\mathrm{Poss}(x,y)$. By Proposition 4.1, connecting all vertices in $\mathrm{Poss}(x,y)$ at height k with an ideal conductor can only reduce the total resistance of the circuit; make this transformation at each height $k > 0$ and denote the resulting tree by $\mathcal{T}(x,y)$. Thus, there are $\lesssim 2^k$ resistors, each with resistance $\sim 2^k$, connected in parallel between height $k-1$ and height k in $\mathcal{T}(x,y)$. This gives a resistance $\gtrsim 1$ between heights $k-1$ and k, and consequently a total resistance of the circuit $\gtrsim n$. This paired with (5.10) establishes the bound $\Pr((x,y) \in K_\sigma) \lesssim \frac{1}{n}$.

Finally, we use this bound to estimate the expected measure of $K_\sigma \cap ([\frac{1}{3}, 1] \times \mathbb{R})$. Observing that we must have $y \in [0,2]$ if $(x,y) \in K_\sigma$, we calculate

$$\begin{aligned}
\mathbb{E}_\sigma \left(\left| K_\sigma \cap \left(\left[\frac{1}{3}, 1 \right] \times \mathbb{R} \right) \right| \right) &= \int \left(\int_{\frac{1}{3}}^{1} \int_0^2 \mathbf{1}_{K_\sigma}(x,y) dy dx \right) d\sigma \\
&= \int_{\frac{1}{3}}^{1} \int_0^2 \Pr((x,y) \in K_\sigma) dy dx \\
&\lesssim \frac{1}{n}.
\end{aligned}$$

Thus, there is a choice of sticky map σ such that (5.5) holds. Moreover, since the random variable σ is distributed uniformly over all sticky maps (this is by definition of our random slope assignment), this argument shows that in fact most sticky σ must satisfy (5.5). □

5.2 Points of distinction between the construction of Kakeya-type sets over Cantor directions in the plane and over arbitrary sublacunary sets in any dimension

We have already pointed out some of the added difficulties that arise when trying to use similar arguments to the ones outlined in the previous section in a more general setting. We summarize these points here and draw the readers attention to several others.

The proof of Theorem 1.2 is modeled on the proof of Bateman and Katz's theorem presented in Section 5.1, with several important distinctions. Overall, our goal in Theorem 1.2 remains essentially the same: to construct a family of tubes rooted on the hyperplane $\{0\} \times [0, 1)^d$, the union of which will eventually give rise to a Kakeya-type set. The slopes of the constituent tubes will be assigned from Ω via a random mechanism involving stickiness akin to the one developed by Bateman and Katz and described in Section 5.1.2. We develop this random mechanism in detail in Chapter 8, with the requisite geometric considerations collected in Chapter 7.

As we have already noted, the notion of Bernoulli percolation on trees plays an important role in the proof of our Theorem 1.2, as it did in the two-dimensional setting. The higher-dimensional structure of Ω does however result in minor changes to the argument, as we will see in Chapter 9. But of the two estimates analogous to (5.4) and (5.5) necessary for the Kakeya-type construction, the first is where the greatest amount of new work is to be done. The bound (5.4) used in [4] is deterministic, providing a bound on the size of *any* sticky collection of tubes as defined in (5.3). However, the counting argument that led to this bound fails to produce a tight enough estimate in higher dimensions; instead, we replace it by a probabilistic statement that suffices for our purposes.

More precisely, the issue is the following. A large lower bound on a union of tubes follows if they do not have significant pairwise overlap amongst themselves; i.e. if the total size of pairwise intersections is small. In dimension two, a good upper bound on the size of this intersection was available uniformly in every sticky slope assignment. Although the argument that provided this bound is not transferable to general dimensions, it is still possible to obtain the desired bound with large probability. A probabilistic statement similar to but not as strong as (5.4) can be derived relatively easily via an estimate on the first moment of the total size of random pairwise intersections. Unfortunately, this is still not sharp enough to yield the disparity in the sizes of the tubes and their translated counterparts necessary to claim the existence of a Kakeya-type set. Indeed, since we prove the straight analogue of the already probabilistic bound (5.5), in order to claim the existence of a single set satisfying both probabilistic estimates simulatenously, we will need knowledge of the variance in sizes of collections of tubes over a not necessarily uniform probability space of sticky maps. Thus, for our higher dimensional setting of Theorem 1.2, we need a second moment estimate on the pairwise intersections of tubes.

Both moment estimates share some common features. For instance, they both exploit Euclidean distance relations between roots and slopes of two intersecting tubes, and combine this knowledge with the relative positions of the roots and slopes within the respective trees in which they live, which affects the slope assignments. However, the technicalities are far greater for the second moment compared to the first. In particular, for the second moment we are naturally led to consider not just pairs, but triples and quadruples of tubes, and need to evaluate the probability of obtaining pairwise intersections among these. Not surprisingly, this probability depends on the structure of the root tuple within its ambient tree. It is the classification of these root configurations, computation of the relevant probabilities and their subsequent application to the estimation of expected intersections that form the novel pieces of the proof of Theorem 1.2 and distinguish it from the planar case.

These added complications remain present in the general treatment when we prove Theorem 1.3. However, in this more general setting, we have to adjust our notion of a sticky mapping between root and slope trees to take better advantage of the potentially very sparse structure of the given sublacunary slope tree. This will lead naturally to a type of mapping between trees that we call *weakly sticky*; see Section 13.2. It is this more flexible notion of stickiness between trees that will allow us to exploit the same general methods developed in the proof of Bateman and Katz's result presented in this chapter, as well as the methods that we will develop during the proof of Theorem 1.2.

Chapter 6

Setup of construction of Kakeya-type sets in $\mathbb{R}^{d+1}$ over a Cantor set of directions: a reformulation of Theorem 1.2

With this chapter, we begin the program of directly proving our Theorems 1.2 and 1.3. Chapter 6 through Chapter 11 cover the construction of Kakeya-type sets over a Cantor set of directions in an arbitrary number of dimensions, Theorem 1.2, while Chapters 12 through 19 treat the general case of Theorem 1.3.

To begin, we choose some integer $M \geq 3$ and a generalized Cantor-type set $\mathcal{C}_M \subseteq [0, 1)$ as described in Section 1.1, and fix these items for the remainder. We also fix an injective map $\gamma : [0, 1] \to \{1\} \times [-1, 1]^d$ satisfying the bi-Lipschitz condition in (1.4). These objects then define a fixed set of directions $\Omega = \{\gamma(t) : t \in \mathcal{C}_M\} \subseteq \{1\} \times [-1, 1]^d$.

Next, we define the collection of tubes that will comprise our Kakeya-type set. Let

$$\mathcal{Q}(n) := \{t \in \mathcal{T}(\{0\} \times [0, 1)^d; M) : h(t) = n\}, \tag{6.1}$$

be the collection of disjoint d-dimensional cubes of sidelength M^{-n} generated by the lattice $M^{-n}\mathbb{Z}^d$ in the set $\{0\} \times [0,1)^d$. More specifically, each $t \in \mathcal{Q}(n)$ is of the form

$$t = \{0\} \times \prod_{l=1}^{d} \left[\frac{j_l}{M^n}, \frac{j_l + 1}{M^n} \right), \tag{6.2}$$

for some $\mathbf{j} = (j_1, \ldots, j_d) \in \{0, 1, \cdots, M^n - 1\}^d$, so that $\#(\mathcal{Q}(n)) = M^{nd}$. For technical reasons, we also define Q_t to be the c_d-dilation of t about its center point, where c_d is a small, positive, dimension-dependent constant. The reason for this technicality, as well as possible values of c_d, will soon emerge in the sequel, but for concreteness choosing $c_d = d^{-2d}$ will suffice.

Fix an arbitrarily large integer $N \geq 1$, typically much bigger than M. For the sake of establishing Theorem 1.2, we will set $n = N$ in most of what follows, through Chapter 11. We will however prove some more generic facts along the way, for example in Chapter 7; thus, in these instances we will work with an arbitrary integer n. This will allow us to easily apply these facts when we treat the general case of Theorem 1.3.

Recall that the Nth iterate $\mathcal{C}_M^{[N]}$ of the Cantor construction is the union of 2^N disjoint intervals each of length M^{-N}. We choose a representative element of $\mathcal{C}_M$ from each of these intervals, calling the resulting finite collection $\mathcal{D}_M^{[N]}$. Clearly $\text{dist}(x, \mathcal{D}_M^{[N]}) \leq M^{-N}$ for every $x \in \mathcal{C}_M$. Set

$$\Omega_N := \gamma(\mathcal{D}_M^{[N]}), \tag{6.3}$$

so that $\text{dist}(\omega, \Omega_N) \leq CM^{-N}$ for any $\omega \in \Omega$, with C as in (1.4). The following fact is an immediate corollary of Proposition 3.2 that will naturally prove vital in establishing Theorem 1.2.

Fact 6.1. *With the set $\mathcal{D}_M^{[N]}$ defined as above, $\mathcal{T}_N(\mathcal{D}_M^{[N]}; M) \cong \mathcal{T}_N([0,1); 2)$.*

For any $t \in \mathcal{Q}(N)$ and any $\omega \in \Omega_N$, we define

$$\mathcal{P}_{t,\omega} := \{r + s\omega : r \in Q_t,\ 0 \leq s \leq 10C_0\}, \tag{6.4}$$

where C_0 is a large constant to be determined shortly (for instance, $C_0 = d^d c^{-1}$ will work, with c as in (1.4)). Thus the set $\mathcal{P}_{t,\omega}$ is a cylinder oriented along ω. Its (vertical) cross-section in the plane $x_1 = 0$ is the cube Q_t. We say that $\mathcal{P}_{t,\omega}$ is *rooted* at t. While $\mathcal{P}_{t,\omega}$ is not strictly speaking a tube as defined in the introduction, the distinction is negligible, since $\mathcal{P}_{t,\omega}$ contains and is contained in constant multiples of δ-tubes with $\delta = c_d \cdot M^{-N}$. By a slight abuse of terminology but no loss of generality, we will henceforth refer to $\mathcal{P}_{t,\omega}$ as a tube.

If a slope assignment $\sigma : \mathcal{Q}(N) \to \Omega_N$ has been specified, we set $P_{t,\sigma} := \mathcal{P}_{t,\sigma(t)}$. Thus $\{P_{t,\sigma} : t \in \mathcal{Q}(N)\}$ is a family of tubes rooted at the elements of an M^{-N}-fine grid in $\{0\} \times [0,1)^d$, with essentially uniform length in t that is bounded above and below by fixed absolute constants. Two such tubes are illustrated in Figure 6.1. For the remainder, we set

$$K_N(\sigma) := \bigcup_{t \in \mathcal{Q}(N)} P_{t,\sigma}. \tag{6.5}$$

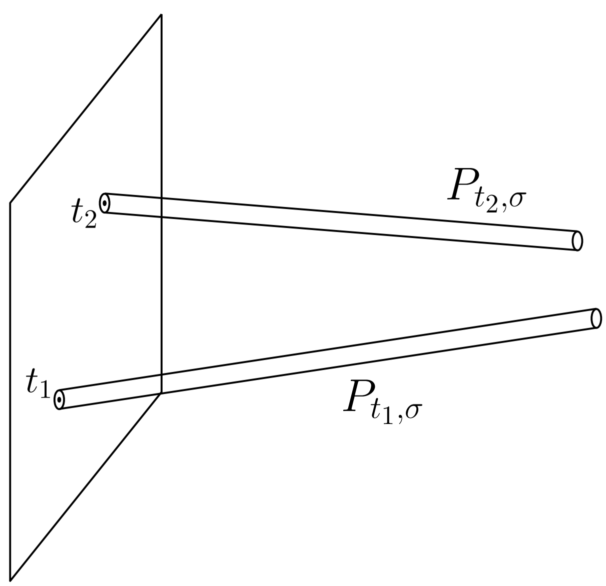

Figure 6.1: Two typical tubes $P_{t_1,\sigma}$ and $P_{t_2,\sigma}$ rooted respectively at t_1 and t_2 in the $\{x_1 = 0\}$–coordinate plane.

For a certain choice of sticky slope assignment σ, this collection of tubes will be shown to generate a Kakeya-type set in the sense of Definition 1.1. This particular slope assignment will not be explicitly described, but rather inferred from the contents of the following proposition.

Proposition 6.2. *For any $N \geq 1$, let Σ_N be a finite collection of slope assignments from the lattice $\mathcal{Q}(N)$ to the direction set Ω_N. Every $\sigma \in \Sigma_N$ generates a set $K_N(\sigma)$ as defined in (6.5). Denote the power set of Σ_N by $\mathfrak{P}(\Sigma_N)$.*

Suppose that $(\Sigma_N, \mathfrak{P}(\Sigma_N), Pr)$ is a discrete probability space equipped with the probability measure Pr, for which the random sets $K_N(\sigma)$ obey the following estimates:

$$Pr\big(\{\sigma : |K_N(\sigma) \cap [0,1] \times \mathbb{R}^d| \geq a_N\}\big) \geq \frac{3}{4}, \tag{6.6}$$

and

$$\mathbb{E}_\sigma |K_N(\sigma) \cap [C_0, C_0+1] \times \mathbb{R}^d| \leq b_N, \tag{6.7}$$

where $C_0 \geq 1$ is a fixed constant, and $\{a_N\}$, $\{b_N\}$ are deterministic sequences satisfying

$$\frac{a_N}{b_N} \to \infty, \quad \textit{as} \quad N \to \infty.$$

Then Ω admits Kakeya-type sets.

Proof. Fix any integer $N \geq 1$. Applying Markov's Inequality to (6.7), we see that

$$\Pr\big(\{\sigma : |K_N(\sigma) \cap [C_0, C_0+1] \times \mathbb{R}^d| \geq 4b_N\}\big) \leq \frac{\mathbb{E}_\sigma |K_N(\sigma) \cap [C_0, C_0+1] \times \mathbb{R}^d|}{4b_N} \leq \frac{1}{4},$$

so,

$$\Pr\big(\{\sigma : |K_N(\sigma) \cap [C_0, C_0+1] \times \mathbb{R}^d| \leq 4b_N\}\big) \geq \frac{3}{4}. \tag{6.8}$$

Combining this estimate with (6.6), we find that

$$\begin{aligned}
&\Pr\Big(\big\{\sigma : |K_N(\sigma) \cap [0,1] \times \mathbb{R}^d| \geq a_N\big\} \bigcap \big\{\sigma : |K_N(\sigma) \cap [C_0, C_0+1] \times \mathbb{R}^d| \leq 4b_N\big\}\Big) \\
&\geq \Pr\big(\big\{|K_N(\sigma) \cap [0,1] \times \mathbb{R}^d| \geq a_N\big\}\big) + \Pr\big(\big\{|K_N(\sigma) \cap [C_0, C_0+1] \times \mathbb{R}^d| \leq 4b_N\big\}\big) - 1 \\
&\geq \frac{3}{4} + \frac{3}{4} - 1 = \frac{1}{2}.
\end{aligned}$$

We may therefore choose a particular $\sigma \in \Sigma_N$ for which the size estimates on $K_N(\sigma)$ given by (6.6) and (6.8) hold simultaneously. Set

$$E_N := K_N(\sigma) \cap [C_0, C_0 + 1] \times \mathbb{R}^d, \quad \text{so that} \quad E_N^*(2C_0 + 1) \supseteq K_N(\sigma) \cap [0, 1] \times \mathbb{R}^d.$$

Then E_N is a union of δ-tubes oriented along directions in $\Omega_N \subset \Omega$ for which

$$\frac{|E_N^*(2C_0 + 1)|}{|E_N|} \geq \frac{a_N}{4b_N} \to \infty, \quad \text{as} \quad N \to \infty,$$

by hypothesis. This shows that Ω admits Kakeya-type sets, per condition (1.1). $\square$

Proposition 6.2 proves our Theorem 1.2. The following five chapters are devoted to establishing a proper randomization over slope assignments Σ_N that will then allow us to verify the hypotheses of Proposition 6.2 for suitable sequences $\{a_N\}$ and $\{b_N\}$. We return to a more concrete formulation of the required estimates in Proposition 8.4.

Chapter 7

Families of intersecting tubes

In this chapter, we will take the opportunity to establish some geometric facts about two intersecting tubes in Euclidean space. These facts will be used in several instances within the proof of Theorem 1.2, as well as in our more general Theorem 1.3. Nonetheless they are really general observations that are not limited to our specific arrangement or description of tubes.

Lemma 7.1. *For $v_1, v_2 \in \Omega_N$ and $t_1, t_2 \in \mathcal{Q}(n)$, $t_1 \neq t_2$, let $\mathcal{P}_{t_1,v_1}$ and $\mathcal{P}_{t_2,v_2}$ be the tubes defined as in (6.4). If there exists $x = (x_1, \cdots, x_{d+1}) \in \mathcal{P}_{t_1,v_1} \cap \mathcal{P}_{t_2,v_2}$, then the inequality*

$$\big| cen(t_2) - cen(t_1) + x_1(v_2 - v_1) \big| \leq 2c_d\sqrt{d}M^{-N}, \tag{7.1}$$

holds, where $cen(t)$ denotes the centre of the cube t.

Proof. The proof is described in the diagram below. If $x \in \mathcal{P}_{t_1,v_1} \cap \mathcal{P}_{t_2,v_2}$, then there exist $y_1 \in Q_{t_1}$, $y_2 \in Q_{t_2}$ such that $x = y_1 + x_1 v_1 = y_2 + x_1 v_2$; i.e., $x_1(v_2 - v_1) = y_1 - y_2$. The inequality (7.1) follows since $|y_i - \text{cen}(t_i)| \leq c_d\sqrt{d}M^{-n}$ for $i = 1, 2$. □

The inequality in (7.1) provides a valuable tool whenever an intersection takes place. For the reader who would like to look ahead, Lemma 7.1 will be used along with Corollary 7.2 to establish Lemma 10.4. The following Corollary 7.3 will be needed for the proofs of Lemmas 10.5 and 10.9.

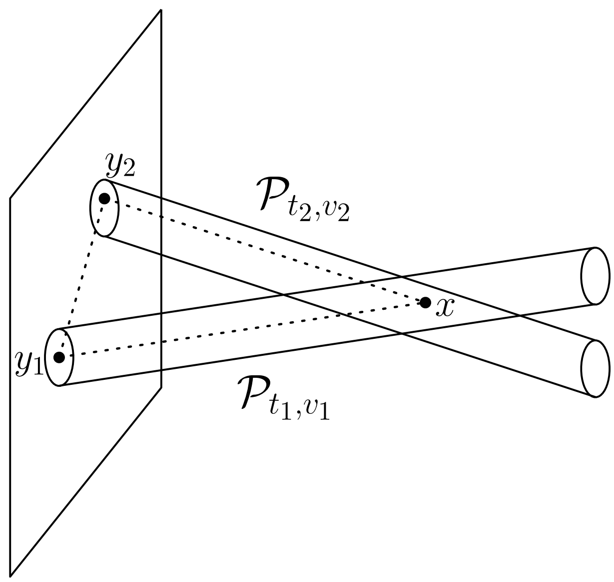

Figure 7.1: A simple triangle is defined by two rooted tubes, $\mathcal{P}_{t_1,v_1}$ and $\mathcal{P}_{t_2,v_2}$, and any point x in their intersection.

Corollary 7.2. *Under the hypotheses of Lemma 7.1 and for $c_d > 0$ suitably small,*

$$|x_1(v_2 - v_1)| \geq \frac{1}{2} \cdot M^{-n}. \tag{7.2}$$

Proof. Since $t_1 \neq t_2$, we must have $|\text{cen}(t_1) - \text{cen}(t_2)| \geq M^{-n}$. Thus an intersection is possible only if

$$x_1|v_2 - v_1| \geq |\text{cen}(t_2) - \text{cen}(t_1)| - 2c_d\sqrt{d}M^{-n} \geq (1 - 2c_d\sqrt{d})M^{-n} \geq \frac{1}{2} \cdot M^{-n},$$

where the first inequality follows from (7.1) and the last inequality holds provided c_d is chosen to satisfy $2c_d\sqrt{d} \leq \frac{1}{2}$. □

Corollary 7.3. *If $t_1 \in \mathcal{Q}(n)$, $v_1, v_2 \in \Omega_N$ and a cube $Q \subseteq \mathbb{R}^{d+1}$ of sidelength C_1M^{-n} with sides parallel to the coordinate axes are given, then there exists at most $C_2 = C_2(C_1)$ choices of $t_2 \in \mathcal{Q}(n)$ such that $\mathcal{P}_{t_1,v_1} \cap \mathcal{P}_{t_2,v_2} \cap Q \neq \emptyset$.*

Proof. As $x = (x_1, \cdots, x_{d+1})$ ranges in Q, x_1 ranges over an interval I of length C_1M^{-n}. If $x \in \mathcal{P}_{t_1,v_1} \cap \mathcal{P}_{t_2,v_2} \cap Q$, the inequality (7.1) and the fact $\text{diam}(\Omega) \leq$

$\operatorname{diam}(\{1\} \times [-1,1]^d) = 2\sqrt{d}$ implies

$$\begin{aligned}\left|\operatorname{cen}(t_2) - \operatorname{cen}(t_1) + \operatorname{cen}(I)(v_2 - v_1)\right| &\leq |(x_1 - \operatorname{cen}(I))(v_2 - v_1)| + 2c_d\sqrt{d}M^{-n} \\ &\leq 2\sqrt{d}(C_1 + c_d)M^{-n},\end{aligned}$$

restricting $\operatorname{cen}(t_2)$ to lie in a cube of sidelength $2\sqrt{d}(C_1+c_d)M^{-n}$ centred at $\operatorname{cen}(t_1) - \operatorname{cen}(I)(v_2 - v_1)$. Such a cube contains at most C_2 subcubes of the form (6.2), and the result follows. □

A recurring theme in the proof of Theorem 1.2 is the identification of a criterion that ensures that a specified point lies in the Kakeya-type set $K_N(\sigma)$ defined in (6.5). With this in mind, we introduce for any $x = (x_1, x_2, \cdots, x_{d+1}) \in [0, 10C_0] \times \mathbb{R}^d$ a set

$$\operatorname{Poss}(x) := \left\{t \in \mathcal{Q}(N) : \text{ there exists } v \in \Omega_N \text{ such that } x \in \mathcal{P}_{t,v}\right\}. \tag{7.3}$$

This set captures all the possible M^{-N}-cubes of the form (6.2) in $\{0\} \times [0,1)^d$ such that a tube rooted at one of these cubes has the potential to contain x, provided it is given the correct orientation. Note that $\operatorname{Poss}(x)$ is independent of any slope assignment σ. Depending on the location of x, $\operatorname{Poss}(x)$ could be empty. This would be the case if x lies outside a large enough compact subset of $[0, 10C_0] \times \mathbb{R}^d$, for example. Even if $\operatorname{Poss}(x)$ is not empty, an arbitrary slope assignment σ may not endow *any* t in $\operatorname{Poss}(x)$ with the correct orientation.

In the next lemma, we list a few easy properties of $\operatorname{Poss}(x)$ that will be helpful later, particularly during the proof of Lemma 11.3. Lemma 7.4 establishes the main intuition behind the $\operatorname{Poss}(x)$ set, as we give a more geometric description of $\operatorname{Poss}(x)$ in terms of an affine copy of the direction set Ω_N. This is illustrated in Figure 7.2 for a particular choice of directions Ω_N.

Lemma 7.4. *Suppose a slope assignment $\sigma : \mathcal{Q}(n) \to \Omega$ has been specified.*

(a) Then we have the containment

$$\left\{t \in \mathcal{Q}(n) : x \in P_{t,\sigma}\right\} \subseteq \mathit{Poss}(x).$$

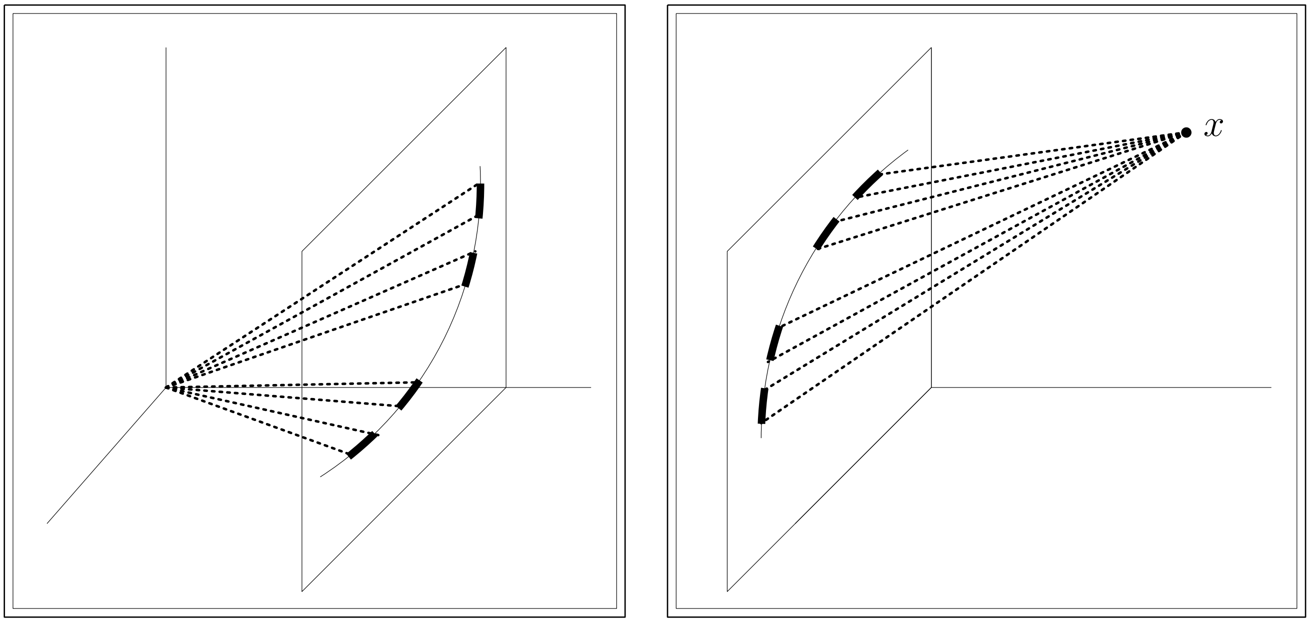

(a) (b)

Figure 7.2: Figure (a) depicts the cone generated by a second stage Cantor construction, Ω_2, on the set of directions given by the curve $\{(1, s, s^2) : 0 \le s \le C\}$ in the $\{1\} \times \mathbb{R}^2$ plane. In Figure (b), a point $x = (x_1, x_2, x_3)$ has been fixed and the cone of directions has been projected backward from x onto the coordinate plane, $x - x_1\Omega_2$. The resulting Poss(x) set is thus given by all cubes $t \in \mathcal{Q}(N)$ such that Q_t intersects a subset of the curve $\{(0, x_2 - x_1 s, x_3 - x_1 s^2) : 0 \le s \le C\}$.

(b) Further, for any $x \in [0, 10C_0] \times \mathbb{R}^d$,

$$Poss(x) = \big\{t \in \mathcal{Q}(n) : Q_t \cap (x - x_1\Omega) \neq \emptyset\big\} \tag{7.4}$$

$$\subseteq \{t \in \mathcal{Q}(n) : t \cap (x - x_1\Omega) \neq \emptyset\}. \tag{7.5}$$

Note that the set in (7.4) could be empty, but the one in (7.5) is not.

Proof. If $x \in P_{t,\sigma}$, then $x \in \mathcal{P}_{t,\sigma(t)}$ with $\sigma(t)$ equal to some $v \in \Omega$. Thus $\mathcal{P}_{t,v}$ contains x and hence $t \in$ Poss(x), proving part (a). For part (b), we observe that $x \in \mathcal{P}_{t,v}$ for some $v \in \Omega$ if and only if $x - x_1 v \in Q_t$; i.e., $Q_t \cap (x - x_1\Omega) \neq \emptyset$. This proves the relation (7.4). The containment in (7.5) is obvious. □

We will also need a bound on the cardinality of Poss(x) within a given cube, and on the cardinality of possible slopes that give rise to indistinguishable tubes passing through a given point x sufficiently far away from the root hyperplane. Contained in Lemma 7.5, these results will prove critical throughout Chapter 11, specifically in the proofs of Lemmas 11.1 and 11.2. We will also have need to formulate a version of this lemma in the language of trees: see Lemma 8.3. Not surprisingly, the Cantor-like construction of Ω plays a role in all these estimates.

Lemma 7.5. *There exists a constant $C_0 \geq 1$ with the following properties.*

(a) For any $x \in [C_0, C_0 + 1] \times \mathbb{R}^d$ and $t \in \mathcal{Q}(N)$, there exists at most one $v \in \Omega_N$ such that $p \in \mathcal{P}_{t,v}$. In other words, for every Q_t in Poss(p), there is exactly one δ-tube rooted at t that contains p.

(b) For any p as in (a), and Q_t, $Q_{t'} \in$ Poss(p), let $v = \gamma(\alpha)$, $v' = \gamma(\alpha')$ be the two unique slopes in Ω_N guaranteed by (a) such that $p \in \mathcal{P}_{t,v} \cap \mathcal{P}_{t',v'}$. If k is the largest integer such that Q_t and $Q_{t'}$ are both contained in the same cube $Q \subseteq \{0\} \times [0, 1)^d$ of sidelength M^{-k} whose corners lie in $M^{-k}\mathbb{Z}^d$, then α and α' belong to the same kth stage basic interval in the Cantor construction.

Proof. (a) Suppose $v, v' \in \Omega_N$ are such that $p \in \mathcal{P}_{t,v} \cap \mathcal{P}_{t,v'}$. Then $p - p_1 v$ and $p - p_1 v'$ both lie in Q_t, so that $p_1 |v - v'| \leq c_d \sqrt{d} M^{-N}$. Since $p_1 \geq C_0$ and (1.4) holds, we find that

$$|\alpha - \alpha'| \leq \frac{c_d \sqrt{d}}{c C_0} M^{-N} < M^{-N},$$

where the last inequality holds if C_0 is chosen large enough. Let us recall from the description of the Cantor-like construction in Section 1.1 that any two basic rth stage intervals are non-adjacent, and hence any two points in $\mathcal{C}_M$ lying in distinct basic rth stage intervals are separated by at least M^{-r}. Therefore the inequality above implies that both α and α' belong to the same basic Nth stage interval in $\mathcal{C}_M^{[N]}$. But $\mathcal{D}_M^{[N]}$ contains exactly one element from each such interval. So $\alpha = \alpha'$ and hence $v = v'$.

(b) If $p \in \mathcal{P}_{t,v} \cap \mathcal{P}_{t',v'}$, then $p_1|v - v'| \leq \text{diam}(\widetilde{Q}_t \cup \widetilde{Q}_{t'}) \leq \text{diam}(Q) = \sqrt{d}M^{-k}$. Applying (1.4) again combined with $p_1 \geq C_0$, we find that $|\alpha - \alpha'| \leq \frac{\sqrt{d}}{cC_0}M^{-k} < M^{-k}$, for C_0 chosen large enough. By the same property of the Cantor construction as used in (a), we obtain that α and α' lie in the same kth stage basic interval in $\mathcal{C}_M^{[k]}$.

□

We should point out that we will require a direct analogue of Lemma 7.5 when treating the general case of sublacunary direction sets in Theorem 1.3. We defer the statement and proof of this analogue until Section 13.1, Lemma 13.1, as it is instructive to first understand the pruning mechanism that defines our direction set Ω_N. While a general lemma could easily be stated that encompasses both Lemma 7.5 and Lemma 13.1, keeping them separate avoids some unhelpful abstraction.

Chapter 8

The random mechanism and sticky collections of tubes in $\mathbb{R}^{d+1}$ over a Cantor set of directions

As we have seen in the planar context, the construction of a Kakeya-type set with orientations given by Ω will require a certain random mechanism. We now describe this mechanism in detail when Ω arises from a Cantor set of directions in an arbitrary number of dimensions.

In order to assign a slope $\sigma(\cdot)$ to the tubes $P_{t,\sigma} := \mathcal{P}_{t,\sigma(t)}$ given by (6.4), we want to define a collection of random variables $\{X_{\langle i_1,\ldots,i_k\rangle} : \langle i_1,\ldots,i_k\rangle \in \mathcal{T}([0,1)^d; M)\}$, one on each edge of the tree used to identify the roots of these tubes. The tree $\mathcal{T}_1([0,1)^d)$ consists of all first generation edges of $\mathcal{T}([0,1)^d)$. It has exactly M^d many edges and we place (independently) a Bernoulli($\frac{1}{2}$) random variable on each edge: $X_{\langle 0\rangle}, X_{\langle 1\rangle}, \ldots, X_{\langle M^d-1\rangle}$. Now, the tree $\mathcal{T}_2([0,1)^d)$ consists of all first and second generation edges of $\mathcal{T}([0,1)^d)$. It has M^d+M^{2d} many edges and we place (independently) a new Bernoulli($\frac{1}{2}$) random variable on each of the M^{2d} second generation edges. We label these $X_{\langle i_1,i_2\rangle}$ where $0 \leq i_1, i_2 < M^d$. We proceed in this way, eventually assigning an ordered collection of independent Bernoulli($\frac{1}{2}$) random variables to the tree

$\mathcal{T}_N([0,1)^d)$:

$$\mathbb{X}_N := \left\{X_{\langle i_1,\ldots,i_k\rangle} : \langle i_1,\ldots,i_k\rangle \in \mathcal{T}_N([0,1)^d),\ 1 \le k \le N\right\},$$

where $X_{\langle i_1,\ldots,i_k\rangle}$ is assigned to the unique edge identifying $\langle i_1, i_2, \cdots, i_k\rangle$, namely the edge joining $\langle i_1, i_2, \cdots, i_{k-1}\rangle$ to $\langle i_1, i_2, \ldots, i_k\rangle$. Each realization of $\mathbb{X}_N$ is a finite ordered collection of cardinality $M^d + M^{2d} + \cdots + M^{Nd}$ with entries either 0 or 1.

We will now establish that every realization of the random variable $\mathbb{X}_N$ defines a sticky map between the truncated position tree $\mathcal{T}_N([0,1)^d)$ and the truncated *binary* tree $\mathcal{T}_N([0,1);2)$, as defined in Definition 3.1. Fix a particular realization $\mathbb{X}_N = \mathbf{x} = \{x_{\langle i_1,\cdots,i_k\rangle}\}$. Define a map $\tau_{\mathbf{x}} : \mathcal{T}_N([0,1)^d) \to \mathcal{T}_N([0,1);2)$, where

$$\tau_{\mathbf{x}}(\langle i_1, i_2, \ldots, i_k\rangle) = \left\langle x_{\langle i_1\rangle}, x_{\langle i_1,i_2\rangle}, \ldots, x_{\langle i_1,i_2,\ldots,i_k\rangle}\right\rangle. \tag{8.1}$$

We then have the following key proposition.

Proposition 8.1. *The map $\tau_{\mathbf{x}}$ just defined is sticky for every realization $\mathbf{x}$ of $\mathbb{X}_N$. Conversely, any sticky map τ between $\mathcal{T}_N([0,1)^d)$ and $\mathcal{T}_N([0,1);2)$ can be written as $\tau = \tau_{\mathbf{x}}$ for some realization $\mathbf{x}$ of $\mathbb{X}_N$.*

Proof. Recalling Definition 3.1, we need to verify that $\tau_{\mathbf{x}}$ preserves heights and lineages. By (8.1), any finite sequence $v = \langle i_1, i_2, \cdots, i_k\rangle$ in $\mathcal{T}([0,1)^d)$ is mapped to a sequence of the same length in $\mathcal{T}([0,1);2)$. Therefore $h(v) = h(\tau_{\mathbf{x}}(v))$ for every $v \in \mathcal{T}([0,1)^d)$.

Next suppose $u \supset v$. Then $u = \langle i_1, \ldots, i_{h(u)}\rangle$, with $h(u) \le k$. So again by (8.1),

$$\tau_{\mathbf{x}}(u) = \left\langle x_{\langle i_1\rangle}, \ldots, x_{\langle i_1,\ldots,i_{h(u)}\rangle}\right\rangle \supset \left\langle x_{\langle i_1\rangle}, \ldots, x_{\langle i_1,\ldots,i_{h(u)}\rangle}, \ldots, x_{\langle i_1,\ldots,i_k\rangle}\right\rangle = \tau_{\mathbf{x}}(v).$$

Thus, $\tau_{\mathbf{x}}$ preserves lineages, establishing the first claim in Proposition 8.1.

For the second, fix a sticky map $\tau : \mathcal{T}_N([0,1)^d) \to \mathcal{T}_N([0,1);2)$. Define $x_{\langle i_1\rangle} := \tau(\langle i_1\rangle)$, $x_{\langle i_1,i_2\rangle} := \pi_2 \circ \tau(\langle i_1, i_2\rangle)$, and in general

$$x_{\langle i_1,\cdots,i_k\rangle} := \pi_k \circ \tau(\langle i_1, i_2, \cdots, i_k\rangle), \qquad k \ge 1,$$

where π_k denotes the projection map whose image is the kth coordinate of the input sequence. The collection $\mathbf{x} = \{x_{\langle i_1, i_2, \cdots, i_k \rangle}\}$ is the unique realization of $\mathbb{X}_N$ that verifies the second claim. $\square$

8.1 Slope assignment algorithm

Recall from Section 1.1 and Chapter 6 that $\Omega := \gamma(\mathcal{C}_M)$ and $\Omega_N := \gamma(\mathcal{D}_M^{[N]})$, where $\mathcal{C}_M$ is the generalized Cantor-type set and $\mathcal{D}_M^{[N]}$ a finitary version of it. In order to exploit the binary structure of the trees $\mathcal{T}(\mathcal{C}_M) := \mathcal{T}(\mathcal{C}_M; M)$ and $\mathcal{T}(\mathcal{D}_M^{[N]}) := \mathcal{T}(\mathcal{D}_M^{[N]}; M)$ advanced in Proposition 3.2 and Fact 6.1, we need to map traditional binary sequences onto the subsequences of $\{0, \ldots, M-1\}^\infty$ defined by $\mathcal{C}_M$ or $\mathcal{D}_M^{[N]}$.

Proposition 8.2. *Every sticky map τ as in* (8.1) *that maps $\mathcal{T}_N([0,1)^d; M)$ to $\mathcal{T}_N([0,1); 2)$ induces a natural mapping $\sigma = \sigma_\tau$ from $\mathcal{T}_N([0,1)^d)$ into Ω_N. The maps σ_τ obey a uniform Lipschitz-type condition: for any $t, t' \in \mathcal{T}_N([0,1)^d)$, $t \neq t'$,*

$$\left|\sigma_\tau(t) - \sigma_\tau(t')\right| \leq CM^{-h(D(\tau(t), \tau(t')))}, \tag{8.2}$$

where C is as in (1.4)*.*

Remark: While the choice of $\mathcal{D}_M^{[N]}$ for a given $\mathcal{C}_M^{[N]}$ is not unique, the mapping $\tau \mapsto \sigma_\tau$ is unique given a specific choice. Moreover, if $\mathcal{D}_M^{[N]}$ and $\overline{\mathcal{D}}_M^{[N]}$ are two selections of finitary direction sets at scale M^{-N}, then the corresponding maps σ_τ and $\overline{\sigma}_\tau$ must obey

$$\left|\sigma_\tau(v) - \overline{\sigma}_\tau(v)\right| \leq CM^{-h(v)} \quad \text{for every } v \in \mathcal{T}_N([0,1)^d), \tag{8.3}$$

where C is as in (1.4). Thus given τ, the slope in Ω that is assigned by σ_τ to an M-adic cube in $\{0\} \times [0,1)^d$ of sidelength M^{-N} is unique up to an error of $O(M^{-N})$. As a consequence P_{t,σ_τ} and $P_{t,\overline{\sigma}_\tau}$ are comparable, in the sense that each is contained in a $O(M^{-N})$-thickening of the other.

Proof. There are two links that allow passage of τ to σ. The first of these is the isomorphism ψ constructed in Proposition 3.2 that maps $\mathcal{T}(\mathcal{C}_M; M)$ onto $\mathcal{T}([0,1); 2)$.

Under this isomorphism, the pre-image of any k-long sequence of 0's and 1's is a vertex w of height k in $\mathcal{T}(\mathcal{C}_M; M)$, in other words one of the 2^k chosen M-adic intervals of length M^{-k} that constitute $\mathcal{C}_M^{[k]}$. The second link is a mapping $\Phi : \mathcal{T}_N(\mathcal{C}_M; M) \to \mathcal{D}_M^{[N]}$ that sends every vertex w to a point in $\mathcal{C}_M \cap w$, where, as usual, we have also let w denote the particular M-adic interval that it identifies. While the choice of the image point, i.e., $\mathcal{D}_M^{[N]}$ is not unique, any two candidates Φ, $\overline{\Phi}$ satisfy

$$\left|\Phi(w) - \overline{\Phi}(w)\right| \leq \operatorname{diam}(w) = M^{-h(w)} \quad \text{for every } w \in \mathcal{T}_N(\mathcal{C}_M; M). \tag{8.4}$$

We are now ready to describe the assignment $\tau \mapsto \sigma = \sigma_\tau$. Given a sticky map $\tau : \mathcal{T}_N([0,1)^d; M) \to \mathcal{T}_N([0,1); 2)$ such that

$$\tau(\langle i_1, i_2, \cdots, i_k \rangle) = \langle X_{\langle i_1 \rangle}, \cdots, X_{\langle i_1, i_2, \cdots, i_k \rangle} \rangle,$$

the transformed random variable

$$Y_{\langle i_1, i_2 \ldots, i_k \rangle} := \gamma \circ \Phi \circ \psi^{-1} \left(\langle X_{\langle i_1 \rangle}, X_{\langle i_1, i_2 \rangle}, \ldots, X_{\langle i_1, i_2, \ldots, i_k \rangle} \rangle \right)$$

associates a random direction in $\Omega_N = \gamma(\mathcal{D}_M^{[N]})$ to the sequence $t = \langle i_1, \ldots, i_k \rangle$ identified with a unique vertex $t \in \mathcal{T}_N([0,1)^d)$. Thus, defining

$$\sigma := \gamma \circ \Phi \circ \psi^{-1} \circ \tau \tag{8.5}$$

gives the appropriate (random) mapping claimed by the proposition. The weak Lipschitz condition (8.2) is verified as follows,

$$\begin{aligned}
\left|\sigma_\tau(t) - \sigma_\tau(t')\right| &= \left|\gamma \circ \Phi \circ \psi^{-1} \circ \tau(t) - \gamma \circ \Phi \circ \psi^{-1} \circ \tau(t')\right| \\
&\leq C\left|\Phi \circ \psi^{-1} \circ \tau(t) - \Phi \circ \psi^{-1} \circ \tau(t')\right| \\
&\leq C M^{-h(D(\psi^{-1}\circ\tau(t), \psi^{-1}\circ\tau(t')))} \\
&= C M^{-h(D(\tau(t), \tau(t')))}.
\end{aligned}$$

Here the first inequality follows from (1.4), the second from the definition of Φ.

The third step uses the fact that ψ is an isomorphism, so that $h(D(\tau(t),\tau(t'))) = h(D(\psi^{-1}\circ\tau(t),\psi^{-1}\circ\tau(t')))$. Finally, any non-uniqueness in the definition of σ comes from Φ, hence (8.3) follows from (8.4) and (1.4). □

The stickiness of the maps $\tau_{\mathbf{x}}$ is built into their definition (8.1). The reader may be interested in observing that there is a naturally sticky map that we have already introduced, which should be viewed as the inspiration for the construction of τ and σ_τ. We refer to the geometric content of Lemma 7.5, which in the language of trees has a particularly succinct reformulation. We record this below.

Lemma 8.3. *For C_0 obeying the requirement of Lemma 7.5 and $x \in [C_0, C_0+1]\times\mathbb{R}^d$, let Poss$(x)$ be as in* (7.3). *Let Φ and ψ be the maps used in Proposition 8.2. Then the map $t \mapsto \beta(t)$ which maps every $t \in$ Poss(x) to the unique $\beta(t) \in [0,1)$ such that*

$$x \in \mathcal{P}_{t,v(t)} \quad \textit{where} \quad v(t) = \gamma\circ\Phi\circ\psi^{-1}\circ\beta(t), \tag{8.6}$$

extends as a well-defined sticky map from $\mathcal{T}_N(\text{Poss}(x); M)$ to $\mathcal{T}_N([0,1); 2)$.

Proof. By Lemma 7.5(a), there exists for every $t \in \text{Poss}(x)$ a unique $v(t) \in \Omega_N$ such that $x \in \mathcal{P}_{t,v(t)}$. Let us therefore define for $1 \le k \le N$,

$$\beta(\pi_1(t),\cdots,\pi_k(t)) = (\pi_1\circ\beta(t),\cdots,\pi_k\circ\beta(t)) \tag{8.7}$$

where $\beta(t)$ is as in (8.6) and as always π_k denotes the projection to the kth coordinate of an input sequence. More precisely, $\pi_k(t)$ represents the unique kth level M-adic cube that contains t. Similarly $\pi_k(\beta(t))$ is the kth component of the N-long binary sequence that identifies $\beta(t)$. The function β defined in (8.7) maps $\mathcal{T}_N(\text{Poss}(x); M)$ to $\mathcal{T}_N([0,1); 2)$, and agrees with β as in (8.6) if $k = N$.

To check that the map is consistently defined, we pick $t \ne t'$ in Poss(x) with $u = D(t,t')$ and aim to show that $\beta(\pi_1(t),\cdots,\pi_k(t)) = \beta(\pi_1(t'),\cdots,\pi_k(t'))$ for all k such that $k \le h(u)$. But by definition (8.6), $v(t)$ and $v(t')$ have the property that $x \in \mathcal{P}_{t,v(t)} \cap \mathcal{P}_{t',v(t')}$. Hence Lemma 7.5(b) asserts that $\alpha(t) = \gamma^{-1}(v(t))$ and $\alpha(t') = \gamma^{-1}(v(t'))$ share the same basic interval at step $h(u)$ of the Cantor construction.

Thus $\beta(t) = \psi \circ \Phi^{-1} \circ \alpha(t)$ and $\beta(t') = \psi \circ \Phi^{-1} \circ \alpha(t')$ have a common ancestor in $\mathcal{T}_N([0,1);2)$ at height $h(u)$, and hence $\pi_k(\beta(t)) = \pi_k(\beta(t'))$ for all $k \le h(u)$, as claimed. Preservation of heights and lineages is a consequence of the definition (8.7), and stickiness follows. □

8.2 Construction of Kakeya-type sets revisited

As τ ranges over all sticky maps $\tau_{\mathbf{x}} : \mathcal{T}_N([0,1)^d) \to \mathcal{T}_N([0,1);2)$ with $\mathbf{x} \in \mathbb{X}_N$, we now have for every vertex $t \in \mathcal{T}_N([0,1)^d)$ with $h(t) = N$ a random sticky slope assignment $\sigma(t) \in \Omega_N$ defined as above. For all such t, this generates a randomly oriented tube $P_{t,\sigma}$ given by (6.4) rooted at the M-adic cube identified by t, with sidelength $c_d \cdot M^{-N}$ in the $\{x_1 = 0\}$ plane. We may rewrite the collection of such tubes from (6.5) as

$$K_N(\sigma) := \bigcup_{\substack{t \in \mathcal{T}_N([0,1)^d) \\ h(t)=N}} P_{t,\sigma}. \tag{8.8}$$

On average, a random collection of tubes with the above described sticky slope assignment will comprise a Kakeya-type set, as per (1.1). Specifically, we will show in the next chapter that the following proposition holds. In view of Proposition 6.2, this will suffice to prove Theorem 1.2.

Proposition 8.4. *Suppose $(\Sigma_N, \mathfrak{P}(\Sigma_N), Pr)$ is the probability space of sticky maps described above, equipped with the uniform probability measure. For every $\sigma \in \Sigma_N$, there exists a set $K_N(\sigma)$ as defined in (8.8), with tubes oriented in directions from $\Omega_N = \gamma(\mathcal{D}_M^{[N]})$. Then these random sets obey the hypotheses of Proposition 6.2 with*

$$a_N = c_M \frac{\sqrt{\log N}}{N} \qquad \text{and} \qquad b_N = \frac{C_M}{N}, \tag{8.9}$$

where c_M and C_M are fixed positive constants depending only on M and d. The content of Proposition 6.2 allows us to conclude that Ω admits Kakeya-type sets.

Chapter 9

Slope probabilities and root configurations, Cantor case

Having established the randomization method for assigning slopes to tubes, we are now in a position to apply this toward the estimation of probabilities of certain events that will be of interest in the next chapter. Roughly speaking, we wish to compute conditional probabilities that one or more cubes on the root hyperplane are assigned prescribed slopes, provided similar information is available for other cubes.

Lemma 9.1. *Let us fix $v_1, v_2 \in \Omega_N$, so that $v_1 = \gamma(\alpha_1)$ and $v_2 = \gamma(\alpha_2)$ for unique $\alpha_1, \alpha_2 \in \mathcal{D}_M^{[N]}$. We also fix $t_1, t_2 \in \mathcal{T}_N([0,1)^d)$, $h(t_1) = h(t_2) = N$, $t_1 \neq t_2$. Let us denote by $u \in \mathcal{T}_N([0,1)^d)$ and $\alpha \in \mathcal{T}_N(\mathcal{D}_M^{[N]})$ the youngest common ancestors of (t_1, t_2) and (α_1, α_2) respectively; i.e., $u = D(t_1, t_2)$, $\alpha = D(\alpha_1, \alpha_2)$. Then*

$$Pr\big(\sigma(t_2) = v_2 \big| \sigma(t_1) = v_1\big) = \begin{cases} 2^{-(N-h(u))} & \text{if } h(u) \leq h(\alpha), \\ 0 & \text{otherwise.} \end{cases} \tag{9.1}$$

Proof. Keeping in mind the slope assignment as described in (8.5), and the stickiness of the map τ as given in Proposition 8.1, the proof can be summarized as in Figure 9.1. Since t_1 and t_2 must map to $v_1 = \gamma(\alpha_1)$ and $v_2 = \gamma(\alpha_2)$ under $\sigma = \sigma_\tau$, the sticky map $\psi^{-1} \circ \tau$ must map t_1 and t_2 to the Nth stage basic intervals in the

Cantor construction containing α_1 and α_2 respectively. Since sticky maps preserve heights and lineages, we must have $h(\alpha) \geq h(u)$. Assuming this, we simply count the number of distinct edges on the ray defining t_2 that are not common with t_1. The map τ generating $\sigma = \sigma_\tau$ is defined by a binary choice on every edge in $\mathcal{T}_N([0,1)^d)$, and the rays given by t_1 and t_2 agree on their first $h(u)$ edges, so we have exactly $N - h(u)$ binary choices to make. This is precisely (9.1).

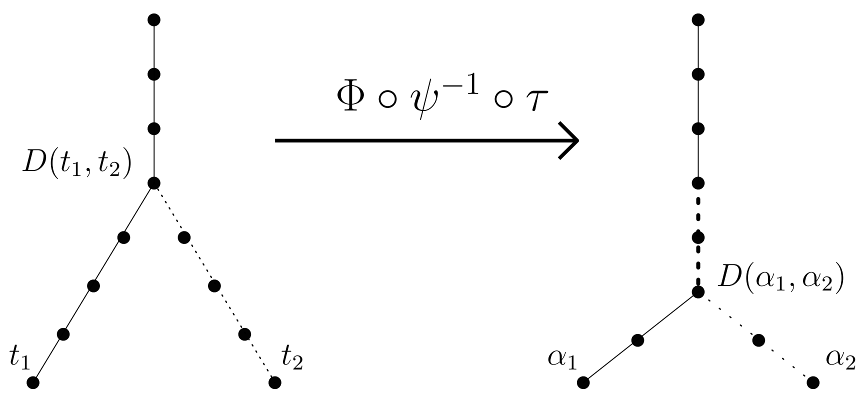

Figure 9.1: Diagram of the sticky assignment between the two rays defining $t_1, t_2 \in \mathcal{T}_N([0,1)^d)$ and the two rays defining their assigned slopes $\alpha_1, \alpha_2 \in \mathcal{D}_M^{[N]}$. The bold edges defining t_1 are fixed to map to the corresponding bold edges at the same height defining α_1. This leaves a binary choice to be made at each of the dotted edges along the path between $D(t_1, t_2)$ and t_2. We see that t_2 is assigned the slope v_2 under σ if and only if these dotted edges are assigned via $\Phi \circ \psi^{-1} \circ \tau$ to the dotted edges on the ray defining α_2.

More explicitly, if $t_1 = \langle i_1, i_2, \cdots, i_N \rangle$ and $t_2 = \langle j_1, \cdots, j_N \rangle$, then

$$\langle i_1, \cdots, i_{h(u)} \rangle = \langle j_1, \cdots, j_{h(u)} \rangle. \tag{9.2}$$

The event of interest may therefore be recast as

$$\begin{aligned}&\big\{\sigma(t_2) = v_2 \big| \sigma(t_1) = v_1\big\} \\ &\qquad = \Big\{\tau(j_1, \cdots, j_N) = \psi \circ \Phi^{-1}(\alpha_2) \Big| \tau(i_1, \cdots, i_N) = \psi \circ \Phi^{-1}(\alpha_1)\Big\}\end{aligned}$$

$$= \left\{ \langle X_{\langle j_1 \rangle}, \cdots, X_{\langle j_1, \cdots, j_N \rangle} \rangle = \psi \circ \Phi^{-1}(\alpha_2) \middle| \langle X_{\langle i_1 \rangle}, \cdots, X_{\langle i_1, \cdots, i_N \rangle} \rangle = \psi \circ \Phi^{-1}(\alpha_1) \right\}$$
$$= \left\{ X_{\langle j_1, \cdots, j_k \rangle} = \pi_k \circ \psi \circ \Phi^{-1}(\alpha_2) \text{ for } h(u) + 1 \le k \le N \right\},$$

where π_k denotes the kth component of the input sequence. At the second step above we have used (8.1) and Proposition 8.2, and the third step uses (9.2). The last event then amounts to the agreement of two $(N - h(u))$-long binary sequences, with an independent, 1/2 chance of agreement at each sequential component. The probability of such an event is $2^{-(N-h(u))}$, as claimed. □

The same idea can be iterated to compute more general probabilities. To exclude configurations that are not compatible with stickiness, let us agree to call a collection

$$\{(t, \alpha_t) : t \in A,\ h(t) = h(\alpha_t) = N\} \subseteq \mathcal{T}_N([0,1)^d) \times \mathcal{D}_M^{[N]} \tag{9.3}$$

of point-slope combinations *sticky-admissible* if there exists a sticky map τ such that $\psi^{-1} \circ \tau$ maps t to α_t for every $t \in A$. Notice that existence of a sticky τ imposes certain consistency requirements on a sticky-admissible collection (9.3); for example $h(D(\alpha_t, \alpha_{t'})) \ge h(D(t, t'))$, and more generally $h(D(\alpha_t : t \in A')) \ge h(D(A'))$ for any finite subset $A' \subseteq A$.

For sticky-admissible configurations, we summarize the main conditional probability of interest below.

Lemma 9.2. *Let A and B be finite disjoint collections of vertices in $\mathcal{T}_N([0,1)^d)$ of height N. Then for any choice of slopes $\{v_t = \gamma(\alpha_t) : t \in A \cup B\} \subseteq \Omega_N$ such that the collection $\{(t, \alpha_t) : t \in A \cup B\}$ is sticky-admissible, the following equation holds:*

$$Pr\big(\sigma(t) = v_t \text{ for all } t \in B \bigm| \sigma(t) = v_t \text{ for all } t \in A\big) = \left(\frac{1}{2}\right)^{k(A,B)},$$

where $k(A, B)$ is the number of distinct edges in the tree identifying B that are not common with the tree identifying A. If $\{(t, \alpha_t) : t \in A \cup B\}$ is not sticky-admissible, then the probability is zero.

For the remainder of this chapter, we focus on some special events of the form dealt with in Lemma 9.2 that will be critical to the proof of (6.6). In all these cases of interest $\#(A), \#(B) \leq 2$. As is reasonable to expect, the configuration of the root cubes within the tree $\mathcal{T}_N([0,1)^d)$ plays a role in determining $k(A,B)$. While there is a large number of possible configurations, we isolate certain structures that will turn out to be generic enough for our purposes.

9.1 Four point root configurations

Definition 9.3. *Let $\mathbb{I} = \{(t_1, t_2); (t'_1, t'_2)\}$ be an ordered tuple of four distinct points in $\mathcal{T}_N([0,1)^d)$ of height N such that*

$$h(u) \leq h(u') \quad \textit{where } u = D(t_1, t_2),\ u' = D(t'_1, t'_2). \tag{9.4}$$

We say that $\mathbb{I}$ is in type 1 configuration if exactly one of the following conditions is satisfied:

(a) either $u \cap u' = \emptyset$, or

(b) $u' \subsetneq u$, or

(c) $u = u' = D(t_i, t'_j)$ for all $i, j = 1, 2$

If $\mathbb{I}$ satisfying (9.4) is not of type 1, we call it of type 2. An ordered tuple $\mathbb{I}$ not satisfying the inequality in (9.4) is said to be of type $j = 1, 2$ if $\mathbb{I}' = \{(t'_1, t'_2); (t_1, t_2)\}$ is of the same type.

The different structural possibilities are listed in Figure 9.2. The advantage of a type 1 configuration is that, in addition to being overwhelmingly popular, it allows (up to permutations) an easy computation of the quantity $k(A,B)$ described in Lemma 9.2 if $\#(A) = \#(B) = 2$, $A \cup B = \{t_1, t'_1, t_2, t'_2\}$ and $\#(A \cap \{t_1, t_2\}) = \#(B \cap \{t_1, t_2\}) = 1$.

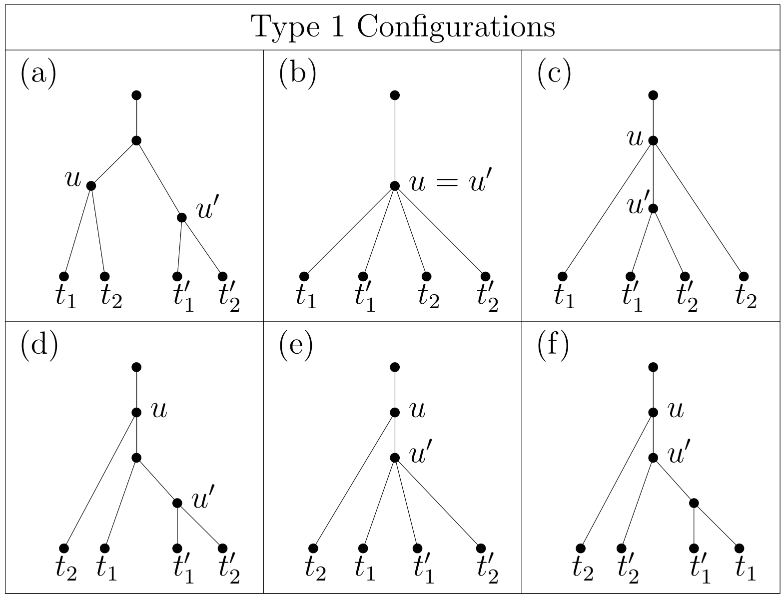

Figure 9.2: All possible four point configurations of type 1, up to permutations.

Lemma 9.4. *Let* $\mathbb{I} = \{(t_1, t_2); (t_1', t_2')\}$ *obeying* (9.4) *be in type 1 configuration. Let* $v_i = \gamma(\alpha_i)$*,* $v_i' = \gamma(\alpha_i')$*,* $i = 1, 2$*, be (not necessarily distinct) points in* Ω_N*. Then there exist two permutations* $\{i_1, i_2\}$ *and* $\{j_1, j_2\}$ *of* $\{1, 2\}$ *such that*

$$Pr\big(\sigma(t_{i_2}) = v_{i_2}, \sigma(t_{j_2}') = v_{j_2}' \big| \sigma(t_{i_1}) = v_{i_1}, \sigma(t_{j_1}') = v_{j_1}'\big) = \left(\frac{1}{2}\right)^{2N-h(u)-h(u')}.$$

provided the collection $\{(t_i, \alpha_i), (t_i', \alpha_i'); i = 1, 2\}$ *is sticky-admissible. If the admissibility requirement is not met, then the probability is zero.*

Proof. The proof is best illustrated by referring to the above diagram, Figure 9.2. If $u \cap u' = \emptyset$, then any two permutations will satisfy the conclusion of the lemma, Figure 9.2(a). In particular, choosing $i_1 = j_1 = 1$, $i_2 = j_2 = 2$, we see that the number of edges in $B = \{t_2, t_2'\}$ not shared by $A = \{t_1, t_1'\}$ is $k(A, B) = (N - h(u)) + (N - h(u')) = 2N - h(u) - h(u')$. The same argument applies if $u = u' = D(t_i, t_j')$ for all $i, j = 1, 2$, Figure 9.2(b).

We turn to the remaining case where $u' \subsetneq u$. Here there are several possiblities for the relative positions of t_1, t_2. Suppose first that there is no vertex w on the ray joining u and u' with $h(u) < h(w) < h(u')$ such that w is an ancestor of t_1 or t_2. This means that the rays of t_1, t_2 and u' follow disjoint paths starting from u, so any choice of permutation suffices, Figure 9.2(c). Suppose next that there is a vertex w on the ray joining u and u' with $h(u) < h(w) < h(u')$ such that w is an ancestor of exactly one of t_1, t_2, but no descendant of w on this path is an ancestor of either t_1 or t_2, Figure 9.2(d). In this case, we choose t_{i_1} to be the unique element of $\{t_1, t_2\}$ whose ancestor is w. Note that the ray for t_{i_2} must have split off from u in this case. Any permutation of $\{t'_1, t'_2\}$ will then give rise to the desired estimate. If neither of the previous two cases hold, then exactly one of $\{t_1, t_2\}$, say t_{i_1}, is a descendant of u'. If $u' = D(t_{i_1}, t'_j)$ for both $j = 1, 2$, then again any permutation of $\{t'_1, t'_2\}$ works, Figure 9.2(e). Thus the only remaining scenario is where there exists exactly one element in $\{t'_1, t'_2\}$, call it t'_{j_1}, such that $h(D(t_{i_1}, t'_{j_1})) > h(u')$. In this case, we choose $A = \{t_{i_1}, t'_{j_1}\}$ and $B = \{t_{i_2}, t'_{j_2}\}$, Figure 9.2(f). All cases now result in $k(A, B) = 2N - h(u) - h(u')$, completing the proof. □

Lemma 9.5. *Let* $\mathbb{I} = \{(t_1, t_2); (t'_1, t'_2)\}$ *obeying* (9.4) *be in type 2 configuration. Then there exist permutations* $\{i_1, i_2\}$ *and* $\{j_1, j_2\}$ *of* $\{1, 2\}$ *for which we have the relations*

$$\begin{aligned} &u_1 \subseteq u,\ u_2 \subsetneq u \text{ with } h(u) \le h(u_1) \le h(u_2), \quad \text{where} \\ &u_1 = D(t_{i_1}, t'_{j_1}),\ u_2 = D(t_{i_2}, t'_{j_2}), \end{aligned}$$

and for which the following equality holds:

$$Pr\big(\sigma(t_{i_1}) = v_{i_1}, \sigma(t'_{j_1}) = v'_{j_1} \;\big|\; \sigma(t_{i_2}) = v_{i_2}, \sigma(t'_{j_2}) = v'_{j_2}\big) = \left(\frac{1}{2}\right)^{2N - h(u) - h(u_1)}$$

for any choice of slopes $v_1, v'_1, v_2, v'_2 \in \Omega_N$ *for which* $\{(t_i, \alpha_i), (t'_i, \alpha'_i); i = 1, 2\}$ *is sticky-admissible.*

Proof. Since $\mathbb{I}$ is of type 2, we know that $u = u'$, and hence all pairwise youngest common ancestors of $\{t_1, t'_1, t_2, t'_2\}$ must lie within u, but that there exist $i, j \in \{1, 2\}$

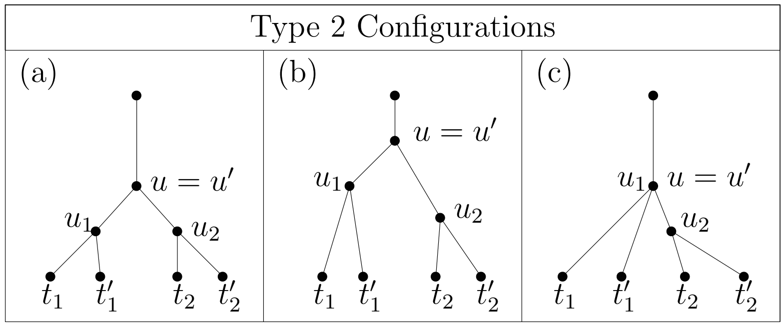

Figure 9.3: All possible four point configurations of type 2, up to permutations.

such that $h(D(t_i, t'_j)) > h(u)$. Let us set (i_2, j_2) to be a tuple for which $h(D(t_{i_2}, t'_{j_2}))$ is maximal. The height inequalities and containment relations are now obvious, and Figure 9.3 shows that $k(A, B) = (N - h(u)) + (N - h(u_1))$ if $A = \{t_{i_2}, t'_{j_2}\}$ and $B = \{t_{i_1}, t'_{j_1}\}$. □

9.2 Three point root configurations

The arguments in the previous section simplify considerably when there are three root cubes instead of four. Since the proofs here are essentially identical to those presented in Lemmas 9.4 and 9.5, we simply record the necessary facts with the accompanying diagram of Figure 9.4.

Definition 9.6. *Let $\mathbb{I} = \{(t_1, t_2); (t_1, t'_2)\}$ be an ordered tuple of three distinct points in $\mathcal{T}_N([0, 1)^d)$ of height N such that $h(u) \leq h(u')$, where $u = D(t_1, t_2)$, $u' = D(t_1, t'_2)$. We say that $\mathbb{I}$ is in type 1 configuration if exactly one of the following two conditions holds:*

(a) $u' \subsetneq u$, or

(b) $u = u' = D(t_2, t'_2)$.

Else $\mathbb{I}$ is of type 2, in which case one necessarily has $u = u'$ and $u_2 = D(t_2, t'_2)$ obeys $u_2 \subsetneq u$. If $h(u) > h(u')$, then the type $\mathbb{I}$ is the same as that of $\mathbb{I}' = \{(t_1, t'_2); (t_1, t_2)\}$.

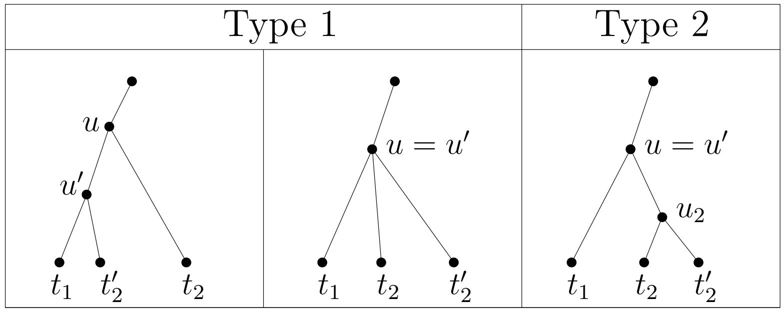

Figure 9.4: Structural possibilities for three point root configurations

Lemma 9.7. *Let* $\mathbb{I} = \{(t_1, t_2); (t_1, t_2')\}$ *be any three-point configuration with* $h(u) \leq h(u')$ *in the notation of Definition 9.6, and let* $v_1 = \gamma(\alpha_1)$*,* $v_2 = \gamma(\alpha_2)$ $v_2' = \gamma(\alpha_2')$ *be slopes in* Ω_N*. Then*

$$Pr\big(\sigma(t_2) = v_2,\ \sigma(t_2') = v_2' \big| \sigma(t_1) = v_1\big) = \begin{cases} \left(\frac{1}{2}\right)^{2N-h(u)-h(u')} & \text{if } \mathbb{I} \text{ is of type 1,} \\ \left(\frac{1}{2}\right)^{2N-h(u)-h(u_2)} & \text{if } \mathbb{I} \text{ is of type 2,} \end{cases}$$

provided the point-slope combination $\{(t_1, \alpha_1), (t_2, \alpha_2), (t_2', \alpha_2')\}$ *is sticky-admissible.*

Chapter 10

Proposition 8.4: proof of the lower bound (6.6)

To establish the appropriate lower bound (6.6), we will exploit the general measure theoretic fact introduced by Bateman and Katz in the planar case (see Lemma 5.1 of this document). Recall that this fact quantifies the notion that if a collection of many thin tubes is to have a large volume, then the intersection of most pairs of tubes should be small. In light of this fact, we will see that the derivation of inequality (6.6) with the a_N specified in (8.9) reduces to the following proposition.

Throughout this chapter, all probability statements are understood to take place on the probability space $(\Sigma_N, \mathfrak{P}(\Sigma_N), \mathrm{Pr})$ identified in Proposition 8.4.

Proposition 10.1. *Fix integers N and R with $N \gg M$ and $N - \frac{1}{10}\log_M N \leq R \leq N - 10$. Define $P^*_{t,\sigma,R}$ to be the portion of $P_{t,\sigma}$ contained in the vertical slab $[M^{R-N}, M^{R+1-N}] \times \mathbb{R}^d$. Then*

$$\mathbb{E}_\sigma\Big[\sum_{t_1 \neq t_2} \big|P^*_{t_1,\sigma,R} \cap P^*_{t_2,\sigma,R}\big|\Big] \lesssim N M^{-2N+2R}, \tag{10.1}$$

where the implicit constant depends only on M and d.

If one can show that with large probability and for all R specified in Proposition 10.1, the quantity $\sum_{t_1 \neq t_2} \big|P^*_{t_1,\sigma,R} \cap P^*_{t_2,\sigma,R}\big|$ is bounded above by the right hand side of (10.1), then Lemma 5.1 would imply (6.6) with $a_N = \sqrt{\log N}/N$. Unfortunately, (10.1) only shows this on average for every R, and hence is too weak a statement to permit such a conclusion. However, with some additional work we are able to upgrade the statement in Proposition 10.1 to a second moment estimate, given below. While still not as strong as the statement mentioned above, this suffices for our purposes with a smaller choice of a_N.

Proposition 10.2. *Under the same hypotheses as Proposition 10.1, there exists a constant $C_{M,d} > 0$ such that*

$$\mathbb{E}_\sigma\Big[\Big(\sum_{t_1\neq t_2}\big|P^*_{t_1,\sigma,R}\cap P^*_{t_2,\sigma,R}\big|\Big)^2\Big] \le C^2_{M,d}\Big(NM^{-2N+2R}\Big)^2. \tag{10.2}$$

Corollary 10.3. *Proposition 10.2 implies (6.6) with a_N as in (8.9).*

Proof. Fix a small constant $c_1 > 0$ such that $2c_1 < \frac{1}{10}$. By Chebyshev's inequality, (10.2) implies that there exists a large constant $C_{M,d} > 0$ such that for every R with $c_1 \log N \le N - R \le 2c_1 \log N$,

$$\begin{aligned}\Pr\Big(\Big\{\sigma : \sum_{t_1\neq t_2}\big|P^*_{t_1,\sigma,R}\cap P^*_{t_2,\sigma,R}\big| \ge 2C_{M,d}N\sqrt{\log N}M^{-2N+2R}\Big\}\Big)&\\ \le \frac{\mathbb{E}_\sigma\Big[\Big(\sum_{t_1\neq t_2}\big|P^*_{t_1,\sigma,R}\cap P^*_{t_2,\sigma,R}\big|\Big)^2\Big]}{\big(2C_{M,d}N\sqrt{\log N}M^{-2N+2R}\big)^2}&\\ \le \frac{1}{4\log N}.&\end{aligned}$$

Therefore,

$$\begin{aligned}\Pr\Big(\bigcup_{N-R=c_1\log N}^{2c_1\log N}\Big\{\sigma : \sum_{t_1\neq t_2}\big|P^*_{t_1,\sigma,R}\cap P^*_{t_2,\sigma,R}\big| \ge C_{M,d}N\sqrt{\log N}M^{-2N+2R}\Big\}\Big)&\\ \le \frac{c_1\log N}{4\log N} < \frac{1}{4}.&\end{aligned}$$

In other words, for a class of σ with probability at least $\frac{3}{4}$,

$$\sum_{t_1 \neq t_2} \left| P^*_{t_1,\sigma,R} \cap P^*_{t_2,\sigma,R} \right| \leq C_{M,d} N \sqrt{\log N} M^{-2N+2R}$$

for every $N - R \in \left[c_1 \log N, 2c_1 \log N\right]$. For such σ and the chosen range of R, we apply Lemma 5.1 with $A_t = P^*_{t,\sigma,R}$, $m = M^{Nd}$, for which $\alpha = C_d M^{R-N} M^{-Nd}$, and

$$\begin{aligned}
\sum_{t_1,t_2} \left| P^*_{t_1,\sigma,R} \cap P^*_{t_2,\sigma,R} \right| &= \Big[\sum_{t_1 = t_2} + \sum_{t_1 \neq t_2} \Big] \left| P^*_{t_1,\sigma,R} \cap P^*_{t_2,\sigma,R} \right| \\
&\leq \alpha n + C_{M,d} N \sqrt{\log N} M^{-2N+2R} \\
&\lesssim M^{R-N} + N \sqrt{\log N} M^{-2N+2R} \\
&\lesssim N \sqrt{\log N} M^{-2N+2R} =: L.
\end{aligned}$$

The last step above uses the specified range of R. Lemma 5.1 now yields that

$$\Big| \bigcup_t P^*_{t,\sigma,R} \Big| \gtrsim \frac{(M^{R-N})^2}{L} \sim \frac{1}{N \sqrt{\log N}}$$

for every $N - R \in \left[c_1 \log N, 2c_1 \log N\right]$. Since $\{\cup_t P^*_{t,\sigma,R} : R \geq 0\}$ is a disjoint collection, we obtain

$$\left| K_N(\sigma) \cap [0,1] \times \mathbb{R}^d \right| \geq \sum_{R = N - 2c_1 \log N}^{N - c_1 \log N} \Big| \bigcup_t P^*_{t,\sigma,R} \Big| \gtrsim \log N \frac{1}{N \sqrt{\log N}} = a_N,$$

which is the desired conclusion (6.6). □

10.1 Proof of Proposition 10.1

Thus, we are charged with proving Proposition 10.2. We will prove Proposition 10.1 first, since it involves many of the same ideas as in the proof of the main proposition, but in a simpler setting. We will need to take advantage of several geometric facts, counting arguments and probability estimates prepared in Chapters 7 and 9 that will be described shortly. For now, we prescribe the main issues in establishing the bound in (10.1).

Proof. Given N and R as in the statement of the proposition, we decompose the slab $[M^{R-N}, M^{R+1-N}] \times \mathbb{R}^d$ into thinner slices Z_k, where

$$Z_k := \left[\frac{k}{M^N}, \frac{k+1}{M^N}\right] \times \mathbb{R}^d, \qquad M^R \le k \le M^{R+1} - 1.$$

Setting $P_{t,\sigma,k} := P_{t,\sigma} \cap Z_k$, we observe that $P^*_{t,\sigma,R}$ is an essentially disjoint union of $\{P_{t,\sigma,k}\}$. Since $P^*_{t,\sigma,R}$ is transverse to Z_k, we arrive at the estimate

$$\begin{aligned} \sum_{t_1 \neq t_2} \left|P^*_{t_1,\sigma,R} \cap P^*_{t_2,\sigma,R}\right| &= \sum_{M^R \le k < M^{R+1}} \sum_{t_1 \neq t_2} |P_{t_1,\sigma,k} \cap P_{t_2,\sigma,k}| \\ &\lesssim M^{-(d+1)N} \sum_{M^R \le k < M^{R+1}} \sum_{t_1 \neq t_2} T_{t_1 t_2}(k) \qquad (10.3) \\ &\lesssim M^{-(d+1)N} \sum_{M^R \le k < M^{R+1}} \sum_{\substack{u \in \mathcal{T}_N([0,1)^d) \\ h(u) < N}} \sum_{(t_1,t_2) \in \mathcal{S}_u} T_{t_1 t_2}(k), \qquad (10.4) \end{aligned}$$

where $T_{t_1 t_2}(k)$ is a random variable that equals one if $P_{t_1,\sigma,k} \cap P_{t_2,\sigma,k} \neq \emptyset$, and is zero otherwise. At the last step in the above string of inequalities, we have further stratified the sum in (t_1, t_2) in terms of their youngest common ancestor $u = D(t_1, t_2)$ in the tree $\mathcal{T}_N([0,1)^d)$, with the index set $\mathcal{S}_u$ of the innermost sum being defined by

$$\mathcal{S}_u := \left\{(t_1, t_2) : t_1, t_2 \in \mathcal{T}_N([0,1)^d),\ h(t_1) = h(t_2) = N,\ D(t_1, t_2) = u\right\}.$$

We will prove below in Lemma 10.7 that

$$\mathbb{E}_\sigma\Big[\sum_{(t_1,t_2) \in \mathcal{S}_u} T_{t_1 t_2}(k)\Big] \lesssim M^{R-N} M^{-dh(u)+Nd} = M^{R-dh(u)+N(d-1)}. \qquad (10.5)$$

Plugging this expected count into the last step of (10.4) and simplifying, we obtain

$$\mathbb{E}_\sigma\Big[\sum_{t_1\neq t_2}\big|P^*_{t_1,\sigma,R}\cap P^*_{t_2,\sigma,R}\big|\Big] \lesssim \sum_{M^R\leq k<M^{R+1}-1} M^{R-2N}\sum_{\substack{u\in\mathcal{T}_N([0,1)^d)\\ h(u)<N}} M^{-dh(u)}$$

$$\lesssim \sum_{M^R\leq k<M^{R+1}-1} M^{R-2N}N \lesssim NM^{2R-2N},$$

which is the estimate claimed by Proposition 10.1. At the penultimate step, we have used the fact that there are M^{dr} vertices u in $\mathcal{T}_N([0,1)^d)$ of height r, resulting in

$$\sum_u M^{-dh(u)} = \sum_{0\leq r<N} M^{-dr}M^{dr} = N. \tag{10.6}$$

□

10.2 Proof of Proposition 10.2

Proof. To establish (10.2), we take a similar route, with some extra care in summing over the (now more numerous) indices. Squaring the expression in (10.3), we obtain

$$\Big[\sum_{t_1\neq t_2}\big|P^*_{t_1,\sigma,R}\cap P^*_{t_2,\sigma,R}\big|\Big]^2 \leq M^{-2(d+1)N}\sum_{k,k'\in[M^R,M^{R+1})}\sum_{\substack{t_1\neq t_2\\ t_1'\neq t_2'}} T_{t_1t_2}(k)T_{t_1't_2'}(k')$$

$$\leq \mathfrak{S}_2+\mathfrak{S}_3+\mathfrak{S}_4,$$

where the index i in $\mathfrak{S}_i$ corresponds to the number of distinct points in the tuple $\{(t_1,t_2);(t_1',t_2')\}$. More precisely, for $i=2,3,4$,

$$\mathfrak{S}_i := M^{-2(d+1)N}\sum_{k,k'}\sum_{\mathbb{I}\in\mathfrak{I}_i} T_{t_1t_2}(k)T_{t_1't_2'}(k'), \quad \text{where} \tag{10.7}$$

$$\mathfrak{I}_i := \left\{\mathbb{I}=\{(t_1,t_2);(t_1',t_2')\}\,\middle|\, \begin{array}{l} t_j,t_j'\in\mathcal{T}_N([0,1)^d), h(t_j)=h(t_j')=N\ \forall j=1,2,\\ t_1\neq t_2,\ t_1'\neq t_2',\ \#(\{t_1,t_1',t_2,t_2'\})=i \end{array}\right\}.$$

The main contribution to the left hand side of (10.2) will be from $\mathbb{E}_\sigma(\mathfrak{S}_4)$, and we will discuss its estimation in detail. The other terms, whose treatment will be briefly sketched, will turn out to be of smaller size.

We decompose $\mathfrak{I}_4 = \mathfrak{I}_{41} \cup \mathfrak{I}_{42}$, where $\mathfrak{I}_{4j}$ is the collection of 4-tuples of distinct points $\{(t_1, t_2); (t_1', t_2')\}$ that are in configuration of type $j = 1, 2$, as explained in Definition 9.3. This results in a corresponding decomposition $\mathfrak{S}_4 = \mathfrak{S}_{41} + \mathfrak{S}_{42}$. For $\mathfrak{S}_{41}$, we further stratify the sum in terms of $u = D(t_1, t_2)$ and $u' = D(t_1', t_2')$, where we may assume without loss of generality that $h(u) \leq h(u')$. Thus,

$$\mathbb{E}_\sigma\big(\mathfrak{S}_{41}\big) = \sum_{k,k'} \sum_{\substack{u,u' \in \mathcal{T}_N([0,1)^d) \\ h(u) \leq h(u') < N}} \mathbb{E}_\sigma\big(\mathfrak{S}_{41}(u, u'; k, k')\big) \quad \text{where} \tag{10.8}$$

$$\mathfrak{S}_{41}(u, u'; k, k') := M^{-2(d+1)N} \sum_{\mathbb{I} \in \mathfrak{I}_{41}(u,u')} T_{t_1 t_2}(k) T_{t_1' t_2'}(k'), \text{ and}$$

$$\mathfrak{I}_{41}(u, u') := \{\mathbb{I} \in \mathfrak{I}_{41} : u = D(t_1, t_2), u' = D(t_1', t_2')\}.$$

In Lemma 10.8 below, we will show that

$$\begin{aligned} \mathbb{E}_\sigma\big[\mathfrak{S}_{41}(u, u'; k, k')\big] &\lesssim M^{-2(d+1)N} M^{2R - d(h(u)+h(u')) + 2N(d-1)} \\ &= M^{2R-4N-d(h(u)+h(u'))}. \end{aligned} \tag{10.9}$$

Inserting this back into (10.8), we now follow the same summation steps that led to (10.1) from (10.5). Specifically, applying (10.6) twice, we obtain

$$\begin{aligned} \mathbb{E}_\sigma(\mathfrak{S}_{41}) &\lesssim M^{2R-4N} \sum_{k,k'} \sum_{u,u'} M^{-d(h(u)+h(u'))} \\ &\lesssim \sum_{k,k'} N^2 M^{2R-4N} \lesssim N^2 M^{4R-4N}, \end{aligned}$$

which is the right hand side of (10.2).

Next we turn to $\mathfrak{S}_{42}$. Motivated by the configuration type, and after permutations of $\{t_1, t_2\}$ and of $\{t_1', t_2'\}$ if necessary (so that the conclusion of Lemma 9.5 holds), we stratify this sum in terms of $u = u' = D(t_1, t_2) = D(t_1', t_2')$, $u_1 = D(t_1, t_1')$,

$u_2 = D(t_2, t_2')$, writing

$$\mathfrak{S}_{42} = \sum_{k,k'} \sum_{\substack{u,u_1,u_2 \in \mathcal{T}_N([0,1)^d) \\ u_1,u_2 \subseteq u}} \mathfrak{S}_{42}(u, u_1, u_2; k, k'), \text{ where}$$

$$\mathfrak{S}_{42}(u, u_1, u_2; k, k') := M^{-2(d+1)N} \sum_{\mathbb{I} \in \mathfrak{I}_{42}(u,u_1,u_2)} T_{t_1 t_2}(k) T_{t_1' t_2'}(k'), \text{ and}$$

$$\mathfrak{I}_{42}(u, u_1, u_2) := \left\{ \mathbb{I} \in \mathfrak{I}_{42} \middle| \begin{array}{l} u = D(t_1, t_2) = D(t_1' t_2'), \\ u_1 = D(t_1, t_1'), \ u_2 = D(t_2, t_2') \end{array} \right\} \tag{10.10}$$

for given $u_1, u_2 \subseteq u$ with $h(u) \leq h(u_1) \leq h(u_2)$. For such u, u_1, u_2, we will prove in Lemma 10.9 below that

$$\mathbb{E}_\sigma\big(\mathfrak{S}_{42}(u, u_1, u_2; k, k')\big) \lesssim M^{-2N-2dh(u_2)}. \tag{10.11}$$

Accepting this estimate for the time being, we complete the estimation of $\mathbb{E}_\sigma(\mathfrak{S}_{42})$ as follows,

$$\begin{aligned}
\mathbb{E}_\sigma(\mathfrak{S}_{42}) &\lesssim \sum_{k,k'} \sum_{u,u_1,u_2} M^{-2N-2dh(u_2)} \\
&\lesssim M^{-2N} \sum_{k,k'} \sum_{u} \sum_{u_2 \subseteq u} M^{-2dh(u_2)} \sum_{\substack{u_1 \subseteq u \\ h(u_1) \leq h(u_2)}} 1 \\
&\lesssim M^{-2N} \sum_{k,k'} \sum_{u} \sum_{u_2 \subseteq u} M^{-2dh(u_2)} \Big[M^{d(h(u_2)-h(u))} \Big] \qquad (10.12) \\
&\lesssim M^{-2N} \sum_{k,k'} \sum_{u} M^{-dh(u)} \sum_{u_2 \subseteq u} M^{-dh(u_2)} \\
&\lesssim N M^{-2N} \sum_{k,k'} \sum_{u} M^{-2dh(u)} \qquad (10.13) \\
&\lesssim N M^{2R-2N}. \qquad (10.14)
\end{aligned}$$

For the range $N - R \leq \frac{1}{2} \log_M N$ assured by Proposition 10.2, the last quantity above is smaller than $(NM^{2R-2N})^2$. The string of inequalities displayed above involve repeated applications of the fact used to prove (10.6), namely that there are $M^{dj-dh(u)}$

cubes of sidelength M^{-j} contained in u. Thus the estimates

$$\sum_{\substack{u_1 \subseteq u \\ h(u_1) \leq h(u_2)}} 1 \lesssim \sum_{j=h(u)}^{h(u_2)} M^{d(j-h(u))} \lesssim M^{d(h(u_2)-h(u))},$$

$$\sum_{u_2 \subseteq u} M^{-dh(u_2)} \lesssim \sum_{N \geq j \geq h(u)} M^{-dj} M^{d(j-h(u))} \lesssim NM^{-dh(u)}, \text{ and}$$

$$\sum_{u} M^{-2dh(u)} = \sum_{j=0}^{N} M^{dj} M^{-2dj} = \sum_{j=0}^{N} M^{-dj} \lesssim 1$$

were used in (10.12) (10.13) and (10.14) respectively, completing the estimation of $\mathbb{E}(\mathfrak{S}_4)$.

Arguments similar to and in fact simpler than those above lead to the following estimates for $\mathbb{E}(\mathfrak{S}_3)$ and $\mathbb{E}(\mathfrak{S}_2)$, where $\mathfrak{S}_3$ and $\mathfrak{S}_2$ are as defined in (10.7):

$$\begin{aligned} \mathbb{E}(\mathfrak{S}_3) &= \mathbb{E}(\mathfrak{S}_{31}) + \mathbb{E}(\mathfrak{S}_{32}) \\ &\lesssim NM^{3R-3N} + M^{3R-3N} \lesssim NM^{3R-3N}, \text{ and} \qquad (10.15) \\ \mathbb{E}(\mathfrak{S}_2) &\lesssim NM^{3R-(d+3)N}. \qquad (10.16) \end{aligned}$$

Here without loss of generality and after a permutation if necessary, we have assumed that $\mathbb{I} = \{(t_1, t_2); (t_1, t_2')\} \in \mathfrak{I}_3$, with $h(D(t_1, t_2)) \leq h(D(t_1, t_2'))$. The subsum $\mathfrak{S}_{3i}$ then corresponds to tuples $\mathbb{I}$ that are in type i configuration in the sense of Definition 9.6. There is only one possible configuration of pairs in $\mathfrak{I}_2$. The derivation of the expectation estimates (10.15) and (10.16) closely follow the estimation of $\mathfrak{S}_4$, with appropriate adjustments in the probability counts; for instance, (10.15) uses Lemma 9.7 and (10.16) uses Lemma 9.1. To avoid repetition, we leave the details of (10.15) and (10.16) to the reader, noting that the right hand term in each case is dominated by $(NM^{2R-2N})^2$ by our conditions on R. $\square$

10.3 Expected intersection counts

It remains to establish (10.5), (10.9) and (10.11). The necessary steps for this are laid out in the following sequence of lemmas. Unless otherwise stated, we will be using the notation introduced in the proofs of Propositions 10.1 and 10.2.

Lemma 10.4. *Fix Z_k. Let us define $\mathcal{A}_u = \mathcal{A}_u(k)$ to be the (deterministic) collection of all $t_1 \in \mathcal{T}_N([0,1)^d)$, $h(t_1) = N$ that are contained in the cube u and whose distance from the boundary of some child of u is $\lesssim kM^{-N-h(u)}$.*

For $t_1 \in \mathcal{A}_u$, let $\mathcal{B}_{t_1} = \mathcal{B}_{t_1}(k)$ denote the (also deterministic) collection of $t_2 \in \mathcal{T}_N([0,1)^d)$ with $h(t_2) = N$ and $D(t_1,t_2) = u$ such that the distance between the centres of t_1 and t_2 is $\lesssim kM^{-N-h(u)}$.

(a) Then for any slope assignment σ, the random variable $T_{t_1t_2}(k) = 0$ unless $t_1 \in \mathcal{A}_u$ and $t_2 \in \mathcal{B}_{t_1}$. In other words,

$$\sum_{(t_1,t_2)\in\mathcal{S}_u} T_{t_1t_2}(k) = \sum_{t_1\in\mathcal{A}_u}\sum_{t_2\in\mathcal{B}_{t_1}} T_{t_1t_2}(k), \text{ so that}$$
$$\mathbb{E}_\sigma\Big[\sum_{(t_1,t_2)\in\mathcal{S}_u} T_{t_1t_2}(k)\Big] = \sum_{t_1\in\mathcal{A}_u}\mathbb{E}_\sigma\Big[\sum_{t_2\in\mathcal{B}_{t_1}} T_{t_1t_2}(k)\Big]. \tag{10.17}$$

(b) The description of $\mathcal{A}_u$ yields the following bound on its cardinality:

$$\#(\mathcal{A}_u) \lesssim \Big(\frac{k}{M^N}\Big)M^{d(N-h(u))} \lesssim M^{R-dh(u)+(d-1)N}.$$

Proof. We observe that $T_{t_1t_2}(k) = 1$ if and only if there exists a point $p = (p_1,\cdots,p_{d+1}) \in Z_k$ and $v_1, v_2 \in \Omega_N$ such that $p \in \mathcal{P}_{t_1,v_1}\cap\mathcal{P}_{t_2,v_2}$, and $\sigma(t_1) = v_1$, $\sigma(t_2) = v_2$. By Lemma 7.1, this implies that

$$|\text{cen}(t_1) - \text{cen}(t_2) + p_1(\sigma(t_1) - \sigma(t_2))| \leq 2c_d\sqrt{d}M^{-N}, \tag{10.18}$$

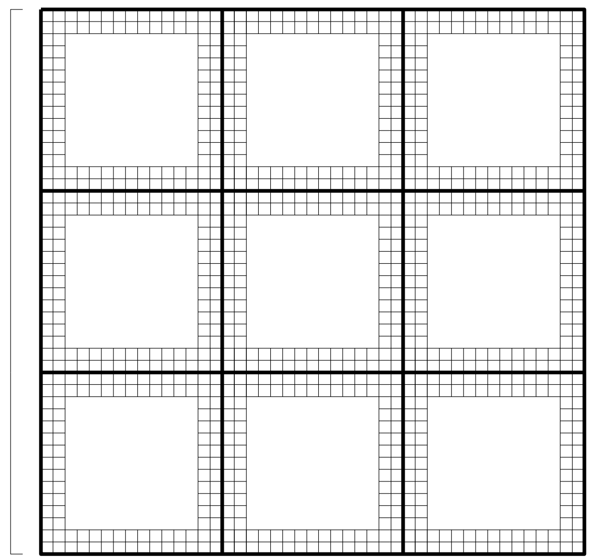

Figure 10.1: A diagram of $\mathcal{A}_u$ when $d = 2$, $M = 3$. Here the largest square is u. The thatched area depicts $\mathcal{A}_u$. The finest squares are the root cubes contained in $\mathcal{A}_u$.

where $\text{cen}(t_i)$ denotes the centre of the cube t_i. For $p \in Z_k$, (10.18) yields

$$\begin{aligned}|\text{cen}(t_1) - \text{cen}(t_2)| &\le p_1|\sigma(t_1) - \sigma(t_2)| + 2c_d\sqrt{d}M^{-N} \lesssim p_1|\sigma(t_1) - \sigma(t_2)| \\ &\lesssim \left(\frac{k+1}{M^N}\right)|\sigma(t_1) - \sigma(t_2)| \lesssim \left(\frac{k}{M^N}\right)M^{-h(D(\tau(t_1),\tau(t_2)))} \\ &\lesssim kM^{-N-h(u)}. \end{aligned} \tag{10.19}$$

The second inequality in the steps above follows from Corollary 7.2, the third from the definition of Z_k and the fourth from the property (8.2) of the slope assignment. Here τ is the unique sticky map that generates σ, as specified in Proposition 8.2. Since τ preserves heights and lineages, $h(D(\tau(t_1), \tau(t_2))) \ge h(D(t_1, t_2)) = h(u)$, and the last step follows.

The inequality in (10.19) implies that $T_{t_1t_2}(k) = 0$ unless $t_2 \in \mathcal{B}_{t_1}$. Further, t_1, t_2

lie in distinct children of u, so t_1 must satisfy

$$\mathrm{dist}(t_1, \partial u') \lesssim \frac{k}{M^N} M^{-h(u)} \quad \text{for some child } u' \text{ of } u,$$

to allow for the existence of some t_2 obeying (10.19). This means $t_1 \in \mathcal{A}_u$, proving (a).

For (b) we observe that u has M^d children. The Lebesgue measure of the set

$$\bigcup_{u'} \Big\{ x \in u' : \mathrm{dist}(x, \partial u') \lesssim k M^{-N-h(u)}, u' \text{ is a child of } u \Big\} \tag{10.20}$$

is therefore $\lesssim (M^d) k M^{-N-h(u)} M^{-(d-1)h(u)}$. The cardinality of $\mathcal{A}_u$ is comparable to the number of M^{-N}-separated points in the set (10.20), and (b) follows. □

Our next task is to make further reductions to the expression on the right hand side of (10.17) that will enable us to invoke the probability estimates from Chapter 9. To this end, let us fix Z_k, $t_1 \in \mathcal{A}_u(k)$, $v_1 = \gamma(\alpha_1) \in \Omega_N$, and define a collection of point-slope pairs

$$\mathcal{E}_u(t_1, v_1; k) := \left\{ (t_2, v_2) \left| \begin{array}{l} t_2 \in \mathcal{T}_N([0,1)^d) \cap \mathcal{B}_{t_1},\ v_2 = \gamma(\alpha_2) \in \Omega_N, \\ h(t_2) = h(\alpha_2) = N,\ u = D(t_1, t_2), \\ \mathcal{P}_{t_1, v_1} \cap \mathcal{P}_{t_2, v_2} \cap Z_k \neq \emptyset,\ h(D(\alpha_1, \alpha_2)) \geq h(u) \end{array} \right. \right\}. \tag{10.21}$$

Thus $\mathcal{E}_u(t_1, v_1; k)$ is non-random as well. The significance of this collection is clarified in the next lemma.

Lemma 10.5. *For $(t_2, v_2) \in \mathcal{E}_u(t_1, v_1; k)$ described as in (10.21), define a random variable $\overline{T}_{t_2 v_2}(t_1, v_1; k)$ as follows:*

$$\overline{T}_{t_2 v_2}(t_1, v_1; k) := \begin{cases} 1 & \text{if } \sigma(t_2) = v_2, \\ 0 & \text{otherwise.} \end{cases} \tag{10.22}$$

(a) *The random variables* $T_{t_1t_2}(k)$ *and* $\overline{T}_{t_2v_2}(t_1, v_1; k)$ *are related as follows: given* $\sigma(t_1) = v_1$,

$$T_{t_1t_2}(k) = \sup\bigl\{\overline{T}_{t_2v_2}(t_1, v_1; k) : (t_2, v_2) \in \mathcal{E}_u(t_1, v_1; k)\bigr\}. \tag{10.23}$$

In particular under the same conditional hypothesis $\sigma(t_1) = v_1$, *one obtains the bound*

$$T_{t_1t_2}(k) \leq \sum_{\substack{v_2 \in \Omega_N \\ (t_2,v_2) \in \mathcal{E}_u(t_1,v_1;k)}} \overline{T}_{t_2v_2}(t_1, v_1; k), \tag{10.24}$$

which in turn implies

$$\mathbb{E}_\sigma\Bigl[\sum_{t_2 \in \mathcal{B}_{t_1}} T_{t_1t_2}(k)\Big|\sigma(t_1) = v_1\Bigr] \leq \sum_{(t_2,v_2) \in \mathcal{E}_u(t_1,v_1;k)} Pr\bigl(\sigma(t_2) = v_2\big|\sigma(t_1) = v_1\bigr). \tag{10.25}$$

(b) *The cardinality of* $\mathcal{E}_u(t_1, v_1; k)$ *is* $\lesssim 2^{N-h(u)}$.

Proof. We already know from Lemma 10.4 that $T_{t_1t_2}(k) = 0$ unless $t_2 \in \mathcal{B}_{t_1}$. Further, if $\sigma(t_1) = v_1$ is known, then it is clear that $T_{t_1t_2}(k) = 1$ if and only if there exists $v_2 \in \Omega_N$ such that $\mathcal{P}_{t_1,v_1} \cap \mathcal{P}_{t_2,v_2} \cap Z_k \neq \emptyset$ and $\sigma(t_2) = v_2$. But this means that the sticky map τ that generates σ must map t_2 to the N-long binary sequence that identifies α_2. Stickiness dictates that $h(D(\alpha_1, \alpha_2)) = h(D(\tau(t_1), \tau(t_2))) \geq h(D(t_1, t_2)) = h(u)$, explaining the constraints that define $\mathcal{E}_u(t_1, v_1; k)$. Rephrasing the discussion above, given $\sigma(t_1) = v_1$, the event $T_{t_1t_2}(k) = 1$ holds if and only if there exists $v_2 \in \Omega_N$ such that $(t_2, v_2) \in \mathcal{E}_u(t_1, v_1; k)$ and $\sigma(t_2) = v_2$. This is the identity claimed in (10.23) of part (a). The bound in (10.24) follows easily from (10.23) since the supremum is dominated by the sum. The final estimate (10.25) in part (a) follows by taking conditional expectation of both sides of (10.24), and observing that $\mathbb{E}_\sigma(\overline{T}_{t_2v_2}(t_1, v_1; k)|\sigma(t_1) = v_1) = \Pr(\sigma(t_2) = v_2\big|\sigma(t_1) = v_1\bigr)$.

We turn to (b). If $v_2 \in \Omega_N$ is fixed, then it follows from Corollary 7.3 (taking Q in that corollary to be the cube of sidelength $O(M^{-N})$ containing $\mathcal{P}_{t_1,v_1} \cap Z_k$) that there exist at most a constant number of choices of t_2 such that $(t_2, v_2) \in \mathcal{E}_u(t_1, v_1; k)$.

But by Fact 6.1 the number of points $\alpha_2 \in \mathcal{D}_M^{[N]}$ (and hence slopes $v_2 \in \Omega_N$) that obey $h(D(\alpha_1, \alpha_2)) \geq h(u)$ is no more than $2^{N-h(u)}$, proving the claim. □

The same argument above applied twice yields the following conclusion, the verification of which is left to the reader.

Corollary 10.6. *Given $t_1 \in \mathcal{A}_u(k)$, $t_1' \in \mathcal{A}_{u'}(k')$, $v_1, v_1' \in \Omega_N$, define $\mathcal{E}_u(t_1, v_1; k)$ and $\mathcal{E}_{u'}(t_1', v_1'; k')$ as in* (10.21) *and the random variables $\overline{T}_{t_2 v_2}(t_1, v_1; k)$, $\overline{T}_{t_2' v_2'}(t_1', v_1'; k')$ as in* (10.22)*. Then given $\sigma(t_1) = v_1$ and $\sigma(t_1') = v_1'$,*

$$\sum_{\substack{t_2 \in \mathcal{B}_{t_1} \\ t_2' \in \mathcal{B}_{t_1'}}} T_{t_1 t_2}(k) T_{t_1' t_2'}(k') \leq \sum^{*} \overline{T}_{t_2 v_2}(t_1, v_1; k) \overline{T}_{t_2' v_2'}(t_1', v_1'; k'),$$

where the notation $\sum^{*}$ represents the sum over all indices $\{(t_2, v_2); (t_2', v_2')\} \in \mathcal{E}_u(t_1, v_1; k) \times \mathcal{E}_{u'}(t_1', v_1'; k')$.

We are now ready to establish the key estimates in the proofs of Propositions 10.1 and 10.2.

Lemma 10.7. *The estimate in* (10.5) *holds.*

Proof. We combine the steps outlined in Lemmas 10.4, 10.5 and 9.1. By Lemma 10.4(a),

$$\begin{aligned} \mathbb{E}_\sigma\Big[\sum_{(t_1,t_2) \in \mathcal{S}_u} T_{t_1 t_2}(k)\Big] &= \sum_{t_1 \in \mathcal{A}_u} \mathbb{E}_\sigma\Big[\sum_{t_2 \in \mathcal{B}_{t_1}} T_{t_1 t_2}(k)\Big] \\ &= \sum_{t_1 \in \mathcal{A}_u} \mathbb{E}_{v_1} \mathbb{E}_\sigma\Big[\sum_{t_2 \in \mathcal{B}_{t_1}} T_{t_1 t_2}(k) \Big| \sigma(t_1) = v_1\Big]. \end{aligned} \tag{10.26}$$

Applying (10.25) from Lemma 10.5 followed by Lemma 9.1, we find that the inner expectation above obeys the bound

$$\mathbb{E}_\sigma\Big[\sum_{t_2 \in \mathcal{B}_{t_1}} T_{t_1 t_2}(k) \big| \sigma(t_1) = v_1\Big] \leq \sum_{(t_2, v_2) \in \mathcal{E}_u(t_1, v_1; k)} \Pr(\sigma(t_2) = v_2 | \sigma(t_1) = v_1)$$

$$\leq \#(\mathcal{E}_u(t_1, v_1; k)) \times \underbrace{2^{-N+h(u)}}_{\text{Lemma 9.1}}$$
$$\lesssim \underbrace{2^{N-h(u)}}_{\text{Lemma 10.5(b)}} \times 2^{-N+h(u)} \lesssim 1,$$

uniformly in v_1. Inserting this back into (10.26), we arrive at

$$\mathbb{E}_\sigma\Big[\sum_{(t_1,t_2)\in\mathcal{S}_u} T_{t_1t_2}(k)\Big] \lesssim \#(\mathcal{A}_u),$$

which according to Lemma 10.4(b) is the bound claimed in (10.5). □

Lemma 10.8. *The estimate in* (10.9) *holds.*

Proof. The proof of (10.9) shares many similarities with that of Lemma 10.7, except that there are now two copies of each of the objects appearing in the proof of (10.5) and the probability estimate comes from Lemma 9.4 instead of Lemma 9.1. We outline the main steps below.

In view of Lemma 9.4 and after a permutation of (t_1, t_2) and of (t'_1, t'_2) if necessary, we may assume that for every $\mathbb{I} = \{(t_1, t_2); (t'_1, t'_2)\} \in \mathfrak{I}_{41}(u, u')$,

$$\Pr\big(\sigma(t_2) = v_2,\ \sigma(t'_2) = v'_2 | \sigma(t_1) = v_1,\ \sigma(t'_1) = v'_1\big) = \left(\frac{1}{2}\right)^{2N-h(u)-h(u')}. \tag{10.27}$$

Now,

$$\begin{aligned}
&\mathbb{E}_\sigma\big(\mathfrak{S}_{41}(u, u'; k, k')\big) \\
&\quad\leq M^{-2(d+1)N}\mathbb{E}_\sigma\Big[\sum_{\mathbb{I}\in\mathfrak{I}_{41}(u,u')} T_{t_1t_2}(k)T_{t'_1t'_2}(k')\Big] \\
&\quad= M^{-2(d+1)N} \sum_{\substack{t_1\in\mathcal{A}_u(k)\\ t'_1\in\mathcal{A}_{u'}(k')}} \mathbb{E}_{v_1,v'_1}\mathbb{E}_\sigma\Big[\sum_{\substack{t_2\in\mathcal{B}_{t_1}\\ t'_2\in\mathcal{B}_{t'_1}}} T_{t_1t_2}(k)T_{t'_1t'_2}(k')\Big|\sigma(t_1) = v_1, \sigma(t'_1) = v'_1\Big] \\
&\quad\lesssim M^{-2(d+1)N} \underbrace{\left(\frac{kk'}{M^{2N}} M^{d(2N-h(u)-h(u'))}\right)}_{\#(t_1,t'_1) \text{ from Lemma 10.4}} \lesssim M^{2R-4N-d(h(u)+h(u'))},
\end{aligned}$$

since according to Corollary 10.6

$$\begin{aligned}
\mathbb{E}_\sigma\Big[\sum_{(t_2,t_2')\in\mathcal{B}_{t_1}\times\mathcal{B}_{t_1'}} & T_{t_1t_2}(k)T_{t_1't_2'}(k')\Big|\sigma(t_1)=v_1,\sigma(t_1')=v_1'\Big] \\
&\leq \mathbb{E}_\sigma\Big[\sum^{*}\overline{T}_{t_2v_2}(t_1,v_1;k)\overline{T}_{t_2',v_2'}(t_1',v_1';k')\Big|\sigma(t_1)=v_1,\ \sigma(t_1')=v_1'\Big] \\
&\lesssim \sum^{*}\Pr(\sigma(t_2)=v_2,\ \sigma(t_2')=v_2' \mid \sigma(t_1)=v_1,\sigma(t_1')=v_1') \\
&\lesssim \underbrace{(2^{N-h(u)})}_{\#(\mathcal{E}_u(t_1,v_1;k))} \times \underbrace{(2^{N-h(u')})}_{\#(\mathcal{E}_{u'}(t_1',v_1';k'))} \times \underbrace{(2^{-2N+h(u)+h(u')})}_{(10.27)\text{ via Lemma 9.4}} \\
&\lesssim 1, \quad \text{uniformly in } v_1,v_1'.
\end{aligned}$$

The proof is therefore complete. □

Lemma 10.9. *The estimate in* (10.11) *holds.*

Proof. The proof of (10.11) is similar to (10.9), and in certain respects simpler. But the configuration type dictates that we set up a different class $\mathcal{E}^*$ of point-slope tuples that will play a role analogous to $\mathcal{E}_u(t_1,v_1;k)$ in the preceding lemmas. Recall the structure of a type 2 configuration from Figure 9.3 and the definition of $\mathfrak{I}_{42}(u,u_1,u_2)$ from (10.10). Given root cubes t_2,t_2', and $u,u_1,u_2\in\mathcal{T}_N([0,1)^d)$ with the property that

$$u_1\subseteq u,\ u_2\subsetneq u, \quad u_2=D(t_2,t_2'), \quad h(u)\leq h(u_1)\leq h(u_2)\leq N=h(t_2)=h(t_2'),$$

and slopes $v_2=\gamma(\alpha_2)$, $v_2'=\gamma(\alpha_2')\in\Omega_N$, we define $\mathcal{E}^*$ (depending on all these objects) to be the following collection of root-slope tuples:

$$\mathcal{E}^* := \left\{ \{(t_1,v_1);(t_1',v_1')\} \middle| \begin{array}{c} \mathbb{I}=\{(t_1,t_2);(t_1',t_2')\}\in\mathfrak{I}_{42}(u,u_1,u_2), \\ v_1=\gamma(\alpha_1),\ v_1'=\gamma(\alpha_1') \text{ for some } \alpha_1,\alpha_1'\in\mathcal{D}_M^{[N]}, \\ \mathcal{P}_{t_1,v_1}\cap\mathcal{P}_{t_2,v_2}\cap Z_k\neq\emptyset,\ \mathcal{P}_{t_1',v_1'}\cap\mathcal{P}_{t_2',v_2'}\cap Z_{k'}\neq\emptyset, \\ \{(t_i,\alpha_i),(t_i',\alpha_i'):i=1,2\} \text{ is sticky-admissible }. \end{array} \right\} \tag{10.28}$$

The relevance of $\mathcal{E}^*$ is this: if $\sigma(t_2) = v_2$ and $\sigma(t_2') = v_2'$ are given, then $T_{t_1t_2}(k)T_{t_1't_2'}(k') = 0$ unless there exist $v_1, v_1' \in \Omega_N$ with $\{(t_1, v_1); (t_1', v_1')\} \in \mathcal{E}^*$ and $\sigma(t_1) = v_1$, $\sigma(t_1') = v_1'$.

We first set about obtaining a bound on the size of $\mathcal{E}^*$ that we will need momentarily. Stickiness dictates that $h(D(\alpha_1, \alpha_2)) \geq h(u)$, and that α_1 is an Nth level descendant of α, the ancestor of α_2 at height $h(u)$. Thus the number of possible α_1 (and hence v_1) is $\leq 2^{N-h(u)}$, by Fact 6.1. Again by stickiness, $h(D(\alpha_1, \alpha_1')) \geq h(u_1)$, so for a given α_1, the number of α_1' (hence v_1') is no more than the number of possible descendants of α^*, the ancestor of α_1 at height $h(u_1)$. This number is thus $\leq 2^{N-h(u_1)}$. Once v_1, v_1' have been fixed (recall that v_2, v_2', t_2, t_2' are already fixed), it follows from Corollary 7.3 that the number of t_1, t_1' obeying the intersection conditions in (10.28) is $\lesssim 1$. Combining these, we arrive at the following bound on the cardinality of $\mathcal{E}^*$:

$$\#(\mathcal{E}^*) \lesssim \left(2^{N-h(u)}\right)\left(2^{N-h(u_1)}\right) = 2^{2N-h(u)-h(u_1)}. \tag{10.29}$$

We use this bound on the size of $\mathcal{E}^*$ to estimate a conditional expectation, essentially the same way as in the previous two lemmas.

$$\begin{aligned} &\mathbb{E}_\sigma\Big[\sum_{\substack{t_1, t_1' \\ \mathbb{I} \in \mathfrak{I}_{42}(u, u_1, u_2)}} T_{t_1t_2}(k)T_{t_1't_2'}(k')\big|\sigma(t_2) = v_2, \sigma(t_2') = v_2'\Big] \\ &\qquad\qquad = \sum_{\mathcal{E}^*} \Pr(\sigma(t_1) = v_1, \sigma(t_1') = v_1' | \sigma(t_2) = v_2, \sigma(t_2') = v_2') \\ &\qquad\qquad \lesssim \#(\mathcal{E}^*) \left(\frac{1}{2}\right)^{2N-h(u)-h(u_1)} \lesssim 1, \end{aligned} \tag{10.30}$$

where the last step follows by combining Lemma 9.5 with (10.29). As a result, we obtain

$$\begin{aligned} &\mathbb{E}_\sigma\big(\mathfrak{S}_{42}(u, u_1, u_2; k, k')\big) \\ &\quad = M^{-2(d+1)N} \mathbb{E}_\sigma\Big[\sum_{\mathbb{I} \in \mathfrak{I}_{42}(u, u_1, u_2)} T_{t_1t_2}(k)T_{t_1't_2'}(k')\Big] \end{aligned}$$

$$\leq M^{-2(d+1)N} \sum_{t_2, t_2' \subseteq u_2} \mathbb{E}_{v_2, v_2'} \mathbb{E}_\sigma \Big[\sum_{\substack{t_1, t_1' \\ \mathbb{I} \in \mathfrak{I}_{42}(u, u_1, u_2)}} T_{t_1 t_2}(k) T_{t_1' t_2'}(k') \Big| \sigma(t_2) = v_2, \sigma(t_2') = v_2' \Big]$$

$$\lesssim M^{-2(d+1)N} \sum_{t_2, t_2' \subseteq u_2} 1$$

$$\lesssim M^{-2(d+1)N} \big(M^{-dh(u_2)+Nd} \big)^2,$$

where the estimate from (10.30) has been inserted in the third step above. The final expression is the bound claimed in (10.11). □

Chapter 11

Proposition 8.4: proof of the upper bound (6.7)

Using the theory developed in Chapter 4, we can establish inequality (6.7) with $b_N = C_M/N$ as in Proposition 8.4 with relative ease. For $x \in \mathbb{R}^{d+1}$, we write $x = (x_1, \overline{x})$, where $\overline{x} = (x_2, \ldots, x_{d+1})$. Since the Kakeya-type set defined by (8.8) is contained in the parallelepiped $[C_0, C_0 + 1] \times [-2C_0, 2C_0]^d$, we may write

$$\begin{aligned}\mathbb{E}_\sigma \left| K_N(\sigma) \cap [C_0, C_0+1] \times \mathbb{R}^d \right| &= \mathbb{E}_\sigma \left(\int_{C_0}^{C_0+1} \int_{[-2C_0,2C_0]^d} \mathbf{1}_{K_N(\sigma)}(x_1, \overline{x}) d\overline{x} dx_1 \right) \\ &= \int_{C_0}^{C_0+1} \int_{[-2C_0,2C_0]^d} \mathbb{E}_\sigma \left(\mathbf{1}_{K_N(\sigma)}(x_1, \overline{x}) \right) \, d\overline{x} dx_1 \\ &= \int_{C_0}^{C_0+1} \int_{[-2C_0,2C_0]^d} \Pr(x) \, d\overline{x} dx_1, \qquad (11.1)\end{aligned}$$

where $\Pr(x)$ denotes the probability that the point $(x_1, \overline{x})$ is contained in the set $K_N(\sigma)$. To establish inequality (6.7) then, it suffices to show that this probability is bounded by a constant multiple of $1/N$, the constant being uniform in $x \in [C_0, C_0 + 1] \times \mathbb{R}^d$.

Let us recall the definition of $\text{Poss}(x)$ from (7.3). We would like to define a certain percolation process on the tree $\mathcal{T}_N(\text{Poss}(x))$ whose probability of survival can majorize $\Pr(x)$. By Lemma 7.5(a), there corresponds to every $t \in \text{Poss}(x)$ exactly one $v(t) \in \Omega_N$ such that $\mathcal{P}_{t,v(t)}$ contains x. Let us also recall that $v(t) = \gamma(\alpha(t))$ for some $\alpha(t) \in \mathcal{D}_M^{[N]}$. By Fact 6.1, $\alpha(t)$ is uniquely identified by $\beta(t) := \psi(\alpha(t))$, which is a deterministic sequence of length N with entries 0 or 1. Here ψ is the tree isomorphism described in Lemma 3.2.

Given a slope assignment $\sigma = \sigma_\tau$ generated by a sticky map $\tau : \mathcal{T}_N([0,1)^d) \to \mathcal{T}_N([0,1); 2)$ as defined in Proposition 8.2 and a vertex $t = \langle i_1, \cdots, i_N \rangle \in \mathcal{T}_N(\text{Poss}(x))$ with $h(t) = N$, we assign a value of 0 or 1 to each edge of the ray identifying t as follows. Let e be the edge identified by the vertex $\langle i_1, i_2, \cdots, i_k \rangle$. Set

$$Y_e := \begin{cases} 1 \text{ if } \pi_k(\tau(t)) = \pi_k(\beta(t)), \\ 0 \text{ if } \pi_k(\tau(t)) \neq \pi_k(\beta(t)). \end{cases} \tag{11.2}$$

To clarify the notation above, recall that both $\tau(t)$ and $\beta(t)$ are N-long binary sequences, and π_k denotes the kth component of the input. Though the definition of Y_e suggests a potential conflict for different choices of t, our next lemma confirms that this is not the case.

Lemma 11.1. *The description in* (11.2) *is consistent in t; i.e., it assigns a uniquely defined binary random variable Y_e to each edge of $\mathcal{T}_N(Poss(x))$. The collection $\{Y_e\}$ is independent and identically distributed as Bernoulli($\frac{1}{2}$) random variables.*

Proof. Let $t, t' \in \mathcal{T}_N(\text{Poss}(x))$, $h(t) = h(t') = N$. Set $u = D(t, t')$, the youngest common ancestor of t and t'. In order to verify consistency, we need to ascertain that for every edge e in $\mathcal{T}_N(\text{Poss}(x))$ leading up to u and for every sticky map τ, the prescription (11.2) yields the same value of Y_e whether we use t or t'. Rephrasing this, it suffices to establish that

$$\pi_k(\tau(t)) = \pi_k(\tau(t')) \quad \text{and} \quad \pi_k(\beta(t)) = \pi_k(\beta(t')) \quad \text{for all } 0 \leq k \leq h(u). \tag{11.3}$$

Both equalities are consequences of the height and lineage-preserving property of sticky maps, by virtue of which

$$h(D(t,t')) \le \min\bigl[h(D(\tau(t),\tau(t'))), h(D(\beta(t),\beta(t')))\bigr].$$

Of these, stickiness of τ has been proven in Proposition 8.1. The unambiguous definition and stickiness of β has been verified in Lemma 8.3.

For the remainder, we recall from Chapter 8 (see the discussion preceding Proposition 8.1) that for $t = \langle i_1, i_2, \cdots, i_N \rangle$, the projection $\pi_k(\tau(t)) = X_{\langle i_1,\cdots,i_k\rangle}$ is a Bernoulli($\frac{1}{2}$) random variable, so $\Pr(Y_e = 1) = \frac{1}{2}$. Further the random variables Y_e associated with distinct edges e in $\mathcal{T}_N(\text{Poss}(x))$ are determined by distinct Bernoulli random variables of the form $X_{\langle i_1,\cdots,i_k\rangle}$. The stated independence of the latter collection implies the same for the former. $\square$

Thus the collection $\mathbb{Y}_N = \{Y_e\}_{e\in\mathcal{E}}$ defines a Bernoulli percolation on $\mathcal{T}_N(\text{Poss}(x))$, where $\mathcal{E}$ is the edge set of $\mathcal{T}_N(\text{Poss}(x))$. As described in Section 4.1, the event $\{Y_e = 0\}$ corresponds to the removal of the edge e from $\mathcal{E}$, and the event $\{Y_e = 1\}$ corresponds to retaining this edge.

Lemma 11.2. *Let $Pr(x) = Pr\{\tau : x \in K_N(\sigma_\tau)\}$ be as in* (11.1), *and $\{Y_e\}$ as in* (11.2).

(a) For any $x \in [C_0, C_0+1] \times \mathbb{R}^d$, the event $\{\tau : x \in K_N(\sigma_\tau)\}$ is contained in

$$\{\tau : \exists \text{ a full-length ray in } \mathcal{T}_N(Poss(x)) \text{ that survives percolation via } \{Y_e\}\}. \tag{11.4}$$

(b) As a result,

$$Pr(x) \le Pr\bigl(\text{survival after percolation on } \mathcal{T}_N(Poss(x))\bigr).$$

Proof. It is clear that $x \in K_N(\sigma_\tau)$ if and only if there exists $t \in \text{Poss}(x)$ such that $\sigma_\tau(t) = v(t)$, where $v(t)$ is the unique slope in Ω_N prescribed by Lemma 7.5(a) for which $x \in \mathcal{P}_{t,v(t)}$. In other words, we have

$$\begin{aligned}\{\tau : x \in K_N(\sigma_\tau)\} &= \bigcup \{\sigma(t) = v(t) : t \in \mathrm{Poss}(x)\} \\ &= \bigcup \{\tau(t) = \beta(t) : t \in \mathrm{Poss}(x)\}, \qquad (11.5)\end{aligned}$$

where the last step follows from the preceding one by unraveling the string of bijective mappings γ^{-1}, Φ^{-1} and ψ (described in Proposition 8.2) that leads from $\sigma(t)$ to $\tau(t)$, and which incidentally also generates $\beta(t) = \langle j_1, \cdots, j_N\rangle \in \mathcal{T}([0,1);2)$ from $v(t)$. Since t is identified by some sequence $\langle i_1, i_2, \ldots, i_N\rangle$, we have its associated random binary sequence

$$\tau(t) = \langle X_{\langle i_1\rangle}, X_{\langle i_1,i_2\rangle}, \ldots, X_{\langle i_1,i_2,\ldots,i_N\rangle}\rangle \in \mathcal{T}_N([0,1);2).$$

Using this, we can rewrite (11.5) as follows:

$$\begin{aligned}&\bigcup_{t\in\mathrm{Poss}(x)} \{\sigma(t) = v(t)\} \\ &= \bigcup_{t\in\mathrm{Poss}(x)} \left\{\langle X_{\langle i_1\rangle}, X_{\langle i_1,i_2\rangle}, \ldots, X_{\langle i_1,i_2,\ldots,i_N\rangle}\rangle = \langle j_1, j_2, \ldots, j_N\rangle\right\} \\ &= \bigcup_{t\in\mathrm{Poss}(x)} \bigcap_{k=1}^{N} \{X_{\langle i_1,\ldots,i_k\rangle} = j_k\} \\ &= \bigcup_{\mathcal{R}\leftrightarrow\langle i_1,\cdots,i_N\rangle\in\partial\mathcal{T}} \ \bigcap_{e\leftrightarrow\langle i_1,\ldots,i_k\rangle\in\mathcal{E}\cap\mathcal{R}} \{X_{\langle i_1,\ldots,i_k\rangle} - j_k = 0\} \\ &= \bigcup_{\mathcal{R}\in\partial\mathcal{T}} \bigcap_{e\in\mathcal{E}\cap\mathcal{R}} \{Y_e = 1\}. \qquad (11.6)\end{aligned}$$

In the above steps we have set $\mathcal{T} := \mathcal{T}_N(\mathrm{Poss}(x))$ for brevity and let $\mathcal{E}$ be the edge set of $\mathcal{T}$. The last step uses (11.2), and the final event is the same as the one in (11.4). Using (11.6), we have

$$\Pr(x) \le \Pr\left(\bigcup_{\mathcal{R}\in\partial\mathcal{T}} \bigcap_{e\in\mathcal{E}\cap\mathcal{R}} \{Y_e = 1\}\right). \qquad (11.7)$$

This last expression is obviously equivalent to the right hand side of (4.1), verifying the second part of the lemma. □

Our next task is therefore to estimate the survival probability of $\mathcal{T}_N(\text{Poss}(x))$ under Bernoulli($\frac{1}{2}$) percolation. For this purpose and in view of the discussion in Section 4.3, we should visualize $\mathcal{T}_N(\text{Poss}(x))$ as an electrical circuit, the resistance of an edge terminating at a vertex of height k being 2^{k-1}, per equation (4.2). Let us denote by $R(\text{Poss}(x))$ the resistance of the entire circuit. In light of the theorem of Lyons, restated in the form of Proposition 4.4, it suffices to establish the following lemma.

Lemma 11.3. *With the resistance of Poss(x) defined as above, we have*

$$R(Poss(x)) \gtrsim N. \tag{11.8}$$

Proof. We follow the same argument as Bateman and Katz [4]. Recalling the containment (7.5) from Lemma 7.4, we find that N_k, the number of kth generation vertices of $\mathcal{T}_N(\text{Poss}(x))$, is bounded above by $\overline{N}_k$, the number of kth level vertices in $\mathcal{T}_N(\{0\} \times [0,1)^d \cap (x - x_1\Omega_N))$. We will shortly prove in Lemma 11.4(b) that $\overline{N}_k \lesssim 2^k$, where the implicit constant is uniform in $x \in [C_0, C_0+1] \times [-2C_0, 2C_0]^d$. Thus, by Corollary 4.2,

$$R(\text{Poss}(x)) \geq \sum_{k=1}^{N} \frac{2^{k-1}}{N_k} \gtrsim \sum_{k=1}^{N} \frac{2^k}{\overline{N}_K} \gtrsim N,$$

establishing inequality (11.8). □

The remaining Lemma 11.4 is an easy observation that follows simply from the connection between sets and the trees used to encode them.

Lemma 11.4. *Let Ω_N be the set defined in* (6.3).

(a) Given Ω_N, there is a constant $C_1 > 0$ (depending only on d and C, c from (1.4)*) such that for any $1 \leq k \leq N$, the number of kth generation vertices in $\mathcal{T}_N(\Omega_N; M)$ is $\leq C_1 2^k$.*

(b) *For any compact set* $\mathbb{K} \subseteq \mathbb{R}^{d+1}$*, there exists a constant* $C(\mathbb{K}) > 0$ *with the following property. For any* $x = (x_1, \cdots, x_{d+1}) \in \mathbb{K}$*, and* $1 \le k \le N$*, the number of* k*th generation vertices in* $\mathcal{T}_N(E(x); M)$ *is* $\le C(\mathbb{K})2^k$*, where* $E(x) := (x - x_1\Omega_N) \cap \{0\} \times [0, 1)^d$*.*

Proof. There are exactly 2^k basic intervals of level k that comprise $\mathcal{C}_M^{[k]}$. Under γ, each such basic interval maps into a set of diameter at most CM^{-k}. Since $\Omega_N = \gamma(\mathcal{D}_M^{[N]}) \subseteq \gamma(\mathcal{C}_M^{[k]})$, the number of kth generation vertices in $\mathcal{T}_N(\Omega_N; M)$, which is also the number of kth level M-adic cubes needed to cover Ω_N, is at most $C_1 2^k$. This proves (a).

Let Q be any kth generation M-adic cube such that $Q \cap \Omega_N \neq \emptyset$. Then on one hand, $(x - x_1 Q) \cap (x - x_1\Omega_N) \neq \emptyset$; on the other hand, the number of kth level M-adic cubes covering $(x - x_1 Q)$ is $\le C(\mathbb{K})$, and part (b) follows. □

Combining Lemmas 11.2 and 11.3 with Proposition 4.4 gives us the desired bound of $\lesssim 1/N$ on (11.1). This completes the proof of inequality (6.7), and so too Proposition 8.4. Our first Theorem 1.2 is therefore established.

Chapter 12

Construction of Kakeya-type sets in $\mathbb{R}^{d+1}$ over an arbitrary sublacunary set of directions

With Theorem 1.2 proven, we turn our attention to the more general Theorem 1.3. Recall that this means we will be occupied with an arbitrary set of sublacunary directions, of which a Cantor-type set of directions is a particular example, as described in Definition 2.7.

The method of proof will be analogous to that of the Cantor case, although many of the details will require a necessarily more technical treatment. Most notably, it will require substantially more work to define a coherent and flexible enough random mechanism on the slope assignments for a collection of root cubes. We will naturally want to exploit the geometric idea of stickiness, but to get the full strength of our Theorem 1.3, we will have to weaken the notion of tree stickiness that was so integral to the proof of Theorem 1.2.

It will be convenient to work with a pruned subset of a sublacunary direction set, one that is guaranteed to have certain useful structural properties that we can then exploit in our subsequent calculations.

12.1 Pruning of the slope tree

Recall that to establish Theorem 1.3 it suffices to prove Proposition 3.7; we focus our attention thusly. Fix a base integer $M \geq 2$ and a sublacunary direction set $\Omega \subseteq \mathbb{R}^{d+1}$ (obeying the conclusion of Proposition 3.6). We also fix an absolute constant $C_0 \geq 1$, which will remain unchanged for the rest of the proof, and whose value will be specified later ($C_0 = 10$ will do). Given any integer N, however large, Proposition 3.6 (see (3.11)) supplies a hyperplane $\mathbb{V}_N$ at unit distance from the origin, a coordinate system φ_N on $\mathbb{V}_N$, and a relatively compact subset $W_N \subseteq \mathcal{C}_\Omega \cap \mathbb{V}_N$ for which $\text{split}(\mathcal{T}(\varphi_N(W_N); M)) > (N+1)(2C_0+1)^d$. The choice of N, and hence $\mathbb{V}_N$, W_N and φ_N will stay fixed during the analysis in Chapters 13–19. The existence of Kakeya-type sets, which is the goal of Proposition 3.7, relies on the ability to conduct this analysis for arbitrarily large N. The constant C_0, on the other hand, does not change with N.

Without loss of generality we will assume that $\mathbb{V}_N = \{1\} \times \mathbb{R}^d$ and that φ_N is the ambient coordinate system in $\mathbb{V}_N$ (and hence in all hyperplanes parallel to $\mathbb{V}_N$). The use of φ_N will be dropped in the sequel, and we will simply write $\text{split}(\mathcal{T}(W_N; M)) > (N+1)(2C_0+1)^d$. We will also assume that $W_N \subseteq \{1\} \times [0,1)^d$; indeed if $W_N \subseteq \{1\} \times [0, M^L)^d$ for some large L, then we scale by a factor of M^{-L} in directions perpendicular to $e_1 = (1, 0, \cdots, 0)$, leaving the direction e_1 unchanged. The tree corresponding to the scaled version of W_N has the same splitting number as the original tree. Further, a union E_N of tubes pointing in the scaled directions can be rescaled back to tubes with orientations in W_N, with the ratio $|E_N^*|/|E_N|$ (as explained in (1.1)) unchanged. From this point onwards, our direction set will be an appropriately chosen subset of $W_N \subseteq \{1\} \times [0,1)^d$ for a fixed N. We rename W_N as Ω, since this will not cause any confusion in the sequel.

We now prune our direction set Ω down to a representative tree enjoying special structural properties, in terms of M-adic and Euclidean distances between certain vertices. The essential features of this trimming process and the modified direction set are summarized below in the main result of this section.

Proposition 12.1. *Let $M \geq 2$ be a base integer, $C_0 \geq 1$ a fixed constant, and $N \gg 1$*

a large parameter as described above. Let $\Omega \subseteq \{1\} \times [0,1)^d$ *be a direction set obeying the hypothesis split*$(\mathcal{T}(\Omega; M)) > (N+1)(2C_0+1)^d$. *Then there exist*

- *a finite subset* $\Omega_N \subseteq \Omega$ *of cardinality* 2^N, *and*
- *an integer* $J = J(\Omega, N) \geq N$

such that the following properties hold for the tree $\mathcal{T}_J(\Omega_N; M)$ *of height* J *encoding* Ω_N*:*

(i) Every ray in $\mathcal{T}_J(\Omega_N; M)$ *splits exactly* N *times.*

(ii) Every splitting vertex in $\mathcal{T}_J(\Omega_N; M)$ *has exactly two children.*

(iii) For any splitting vertex v *of* $\mathcal{T}_J(\Omega_N; M)$, *there exists an integer* $h_v^* > h(v)$ *obeying the following constraints:*

- *None of the descendants of the two children of* v *specified in part (ii) split at any height strictly smaller than* h_v^*.
- *If* $w_1(v), w_2(v)$ *are the two descendants of* v *at height* h_v^*, *then the Euclidean distance between the cubes* w_1 *and* w_2 *obeys the relation*

$$C_0 M^{-h_v^*} \leq \mathit{dist}(w_1(v), w_2(v)) \leq (2C_0 + \sqrt{d}) M^{-h_v^*+1}. \tag{12.1}$$

In fact, h_v^* *is the smallest integer exceeding* $h(v)$ *with this property.*

The integer J *can be chosen to ensure that the following additional condition is met:*

(iv) $C_0 M^{-J} \leq \min\{|\omega - \omega'| : \omega \neq \omega',\ \omega, \omega' \in \Omega_N\}$.

Notice that the tree $\mathcal{T}_N(\mathcal{D}_M^{[N]}; M)$ encoding a Cantor-type set of directions that we considered in Chapters 6 to 11 already satisfies the requirements of Proposition 12.1 by virtue of Fact 6.1; i.e. the slope tree $\mathcal{T}_N(\mathcal{D}_M^{[N]}; M)$ is already pruned. In this way, we see in what sense a pruned tree behaves like a full binary tree of height N.

The pruning process leading to the outcome claimed in the proposition is based on an iterative algorithm. The building block of the iteration is contained in Lemma 12.3 below, with Lemma 12.2 supplying an easy but necessary intermediate step.

Lemma 12.2. *Fix integers $r \geq 0$ and $C_0 \geq 1$. A collection of cubes of cardinality $\geq (2C_0+1)^d + 1$ consisting of M-adic cubes of sidelength M^{-r} and must contain at least two cubes whose Euclidean separation is $\geq C_0 M^{-r}$.*

Proof. We first treat the case $r = 0$. The cube $Q_0 = [0, 2C_0+1)^d$ contains exactly $(2C_0+1)^d$ subcubes of unit sidelength with vertices in $\mathbb{Z}^d$. The central subcube Q maintains a minimum distance of C_0 from the boundary of Q_0. Rephrasing this after a translation, any cube Q with vertices in $\mathbb{Z}^d$ and of sidelength 1 admits at most $(2C_0+1)^d$ similar cubes whose distance from itself is $\leq C_0$. The case of a general $r \geq 0$ follows by scaling Q_0 by a factor of M^{-r}. □

Lemma 12.3. *Fix a constant integer $C_0 \geq 1$, an integer $N_0 \geq (2C_0+1)^d$ and a vertex v_0 of the full M^d-adic tree $\mathcal{T}(\{1\} \times [0,1)^d; M)$. Let $\mathcal{T}_{[0]}$ rooted at v_0 be a subtree with the property that every ray in $\mathcal{T}_{[0]}$ splits at least N_0 times. Then there exist an integer $K^* = K^*(v_0) \geq 1$ and a subtree $\mathcal{T}_{[1]}$ of $\mathcal{T}_{[0]}$ rooted at v_0 and of height K^* such that:*

(i) The root v_0 has exactly two descendants v_1 and v_2 of height K^* in $\mathcal{T}_{[1]}$. Thus $\mathcal{T}_{[1]}$ has exactly one splitting vertex.

(ii) The integer K^* is the smallest with the property that dist$(v_1, v_2) \geq C_0 M^{-K^*-h(v_0)}$. In particular, dist$(v_1, v_2) \leq (C_0 + 2\sqrt{d}) M^{-K^*-h(v_0)+1}$.

(iii) If $\mathcal{T}_{[0]}(v_i)$ is the maximal subtree of $\mathcal{T}_{[0]}$ rooted at v_i then each ray in $\mathcal{T}_{[0]}(v_i)$ splits at least $N_0 - (2C_0+1)^d$ times.

Proof. Each ray in $\mathcal{T}_{[0]}$ splits at least N_0 times, so there exists a generation in this tree consisting of at least 2^{N_0} vertices. Since $2^{N_0} \gg (2C_0+1)^d$, let us define K_0 to be the smallest height in $\mathcal{T}_{[0]}$ such that the number of vertices at that height exceeds $(2C_0+1)^d$. Choose K^* to be the smallest integer with the property that there exist vertices v_1 and v_2 of $\mathcal{T}_{[0]}$ at height $K^* + h(v_0)$ obeying the relation $\text{dist}(v_1, v_2) \geq C_0 M^{-K^*-h(v_0)}$. By Lemma 12.2, this property holds at height K_0, hence K^* exists and is at most K_0.

The subtree $\mathcal{T}_1$ of height K^* rooted at v_0 and generated by v_1, v_2 clearly obeys conditions (i) stated in Lemma 12.3. The lower bound on $\text{dist}(v_1, v_2)$ required in part (ii) is built into the construction; see Figure 12.1. To obtain the upper bound, let v_i' denote the parent of v_i. It follows from the minimality of K^* that $\text{dist}(v_1', v_2') < C_0 M^{-K^*-h(v_0)+1}$. Thus,

$$\text{dist}(v_1, v_2) \le \text{diam}(v_1') + \text{diam}(v_2') + \text{dist}(v_1', v_2') \le (2\sqrt{d} + C_0) M^{-K^*-h(v_0)+1}.$$

It remains to complete the proof of part (iii). Let us recall from the definition of K_0 that the number of elements of $\mathcal{T}_{[0]}$ at height $K^* - 1$ is $\le (2C_0 + 1)^d$. Thus any ray of $\mathcal{T}_{[0]}$ rooted at v_0 contains at most $(2C_0 + 1)^d - 1$ splitting vertices of height $\le K^* - 2$, since each splitting vertex of height $\le K^* - 2$ gives rise to at least one new element (different among themselves and distinct from the terminating vertex of the ray) at height $K^* - 1$. Since every ray of $\mathcal{T}_{[0]}$ contained at least N_0 splitting vertices to begin with, at most $(2C_0 + 1)^d$ of which may be lost by height $K^* - 1$, we are left with at least $N_0 - (2C_0 + 1)^d$ splitting vertices per ray rooted at v_i, which is the conclusion claimed in (iii). □

With the preliminary steps out of the way, we are ready to prove the main proposition.

Proof of Proposition 12.1. We know that $\text{split}(\mathcal{T}(\Omega; M)) > (N+1)(2C_0+1)^d$. Given any $N \ge 1$, we can therefore fix a subtree $\overline{\mathcal{T}}$ of $\mathcal{T}(\Omega; M)$ of infinite height in which every ray splits at least $(N + 1)(2C_0 + 1)^d$ times. The pruning is executed on the subtree $\overline{\mathcal{T}}$ as follows.

In the first step we apply Lemma 12.3 with

$$\mathcal{T}_{[0]} = \overline{\mathcal{T}}, \quad v_0 = \{1\} \times [0, 1)^d \quad \text{and} \quad N_0 = (N + 1)(2C_0 + 1)^d.$$

This yields a subtree $\mathcal{T}_{[1]}$ rooted at $\{1\} \times [0, 1)^d$ of height $K^*(v_0) = i_0$ consisting of two vertices w_1 and w_2 at the bottom-most level. Let us denote by $\overline{\mathcal{T}}(w_i)$ the maximal subtree of $\overline{\mathcal{T}}$ rooted at w_i. By Lemma 12.3 any ray of $\overline{\mathcal{T}}(w_i)$ splits at least

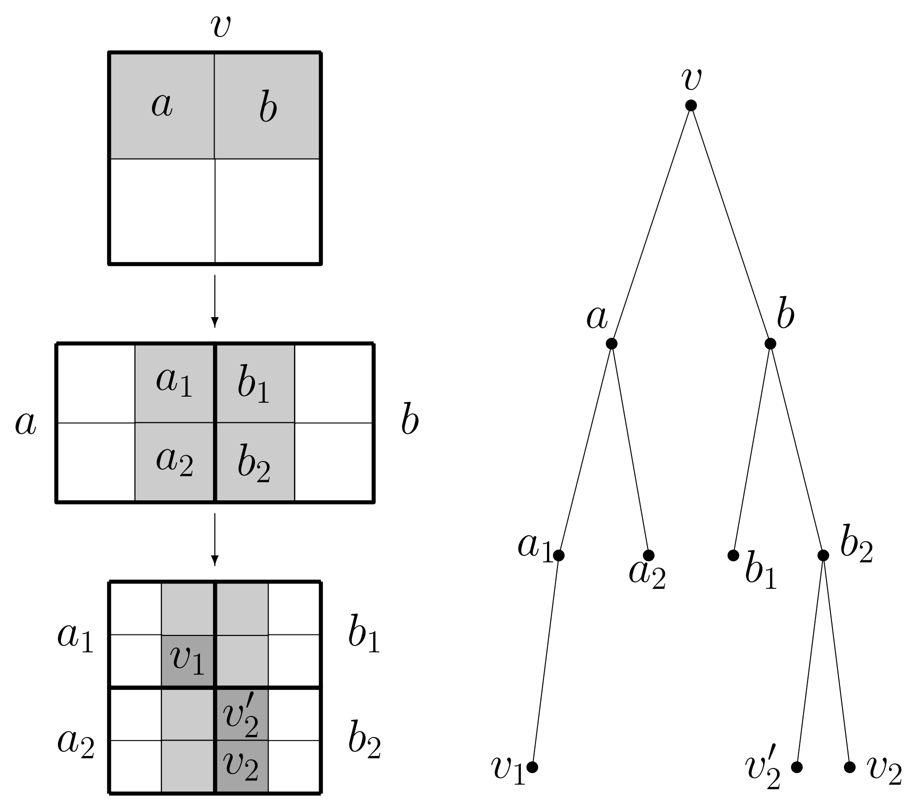

Figure 12.1: An illustration of the procedure generating the forced Euclidean separation between the descendants v_1 and v_2 of $v \in \mathcal{T}$, in $\mathbb{R}^2$ when $M = 2$.

$N(2C_0+1)^d$ times. There is exactly one splitting vertex v in $\mathcal{T}_{[1]}$, and its descendants w_1, w_2 obey $C_0 M^{-i_0} \leq \text{dist}(w_1, w_2) \leq (C_0 + 2\sqrt{d})M^{-i_0}$. Set $\mathbb{W}_1 := \{w_1, w_2\}$.

At the second step we invoke Lemma 12.3 twice, resetting the parameters in that lemma to be

$$\mathcal{T}_{[0]} = \overline{\mathcal{T}}(w_i), \quad v_0 = w_i, \quad N_0 = N(2C_0 + 1)^d$$

for $i = 1, 2$ respectively, and obtaining two subtrees as a consequence. Appending these two newly pruned subtrees of $\overline{\mathcal{T}}(w_i)$ to $\mathcal{T}_{[1]}$ from the previous step, we arrive at a tree $\mathcal{T}_{[2]}$ rooted at $\{1\} \times [0, 1)^d$ of finite height but with rays of possibly variable length, in which every ray splits exactly twice, and every splitting vertex has exactly two children. If v is a splitting vertex of $\mathcal{T}_{[2]}$ not already considered in the first step, then v must either be equal to or a descendant of some $w \in \mathbb{W}_1$. Suppose that

$w_1(v)$ and $w_2(v)$ are the two descendants of w of maximal height in $\mathcal{T}_{[2]}$. Then part (ii) of Lemma 12.3 implies that (12.1) holds with $h_v^* = i_0 + K^*(w)$. This verifies the requirements (i)-(iii) of Proposition 12.1 for $N = 2$. Let us denote by $\mathbb{W}_2$ the collection of four vertices of maximal lineage in $\mathcal{T}_{[2]}$ obtained at the conclusion of this step.

In general at the end of the kth step we have a tree $\mathcal{T}_{[k]}$ of finite height, but with rays of potentially variable length, obeying the requirements (i)-(iii) for $N = k$. The collection of vertices of highest lineage in $\mathcal{T}_{[k]}$ is termed $\mathbb{W}_k$. We have that $\#(\mathbb{W}_k) = 2^k$. The collection $\mathbb{W}_k$ can be decomposed as

$$\mathbb{W}_k = \bigcup \{\mathbb{W}_k(v) : v \text{ is a splitting descendant of some } w \in \mathbb{W}_{k-1}\},$$

where $\mathbb{W}_k(v) = \{w_1(v), w_2(v)\}$ consists of the two descendants of v that lie in $\mathbb{W}_k$. In the $(k+1)$th step, Lemma 12.3 is applied 2^k times in succession. In each application, the values of $\mathcal{T}_{[0]}$, v_0, N_0 are reset to

$$\mathcal{T}_{[0]} = \overline{\mathcal{T}}(w), \quad v_0 = w, \quad N_0 = (N - k + 1)(2C_0 + 1)^d$$

respectively for some $w \in \mathbb{W}_k$. The resulting tree $\mathcal{T}_{[k+1]}$, obtained by appending the 2^k newly constructed trees to $\mathcal{T}_{[k]}$ at the appropriate roots w, clearly obeys (ii) and also (i) with $N = k + 1$. Part (iii) only needs to be verified for the splitting vertices v descended from some $w \in \mathbb{W}_k$, since the splitting vertices of older generations have been dealt with in previous steps. But this follows from part (ii) of Lemma 12.3, with $h_v^* = K^*(w) + h(w)$, which is a maximal height in the tree $\mathcal{T}_{[k+1]}$.

In view of the number of splitting vertices per ray in the original subtree $\overline{\mathcal{T}}$, the process described above can be continued at least N steps. The tree $\mathcal{T}_{[N]}$ of finite height but variable ray lengths obtained at the conclusion of the Nth step satisfies the conditions (i)-(iii). We pick from every vertex of maximal lineage in $\mathcal{T}_{[N]}$ exactly one point of Ω, calling the resulting collection of 2^N chosen points Ω_N. Set $\delta := \min\{|\omega - \omega'| : \omega, \omega' \in \Omega_N, \ \omega \neq \omega'\} > 0$. The rays in $\mathcal{T}_{[N]}$ are now extended as rays representing the points in Ω_N (and hence without introducing any further splits) to a uniform height J that satisfies $M^{-J} \leq C_0^{-1}\delta$, thereby meeting the criterion in part (iv). □

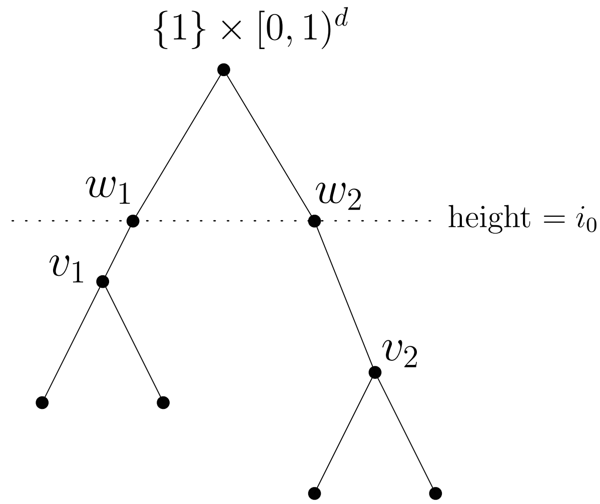

Figure 12.2: An illustration of a pruned tree at the second step of pruning. Note that, in general, we may have $v_1 = w_1$ and/or $v_2 = w_2$.

12.2 Splitting and basic slope cubes

The pruned slope tree $\mathcal{T}_J(\Omega_N; M)$ produced by Proposition 12.1 looks like an elongated version of the full binary tree of height N. Rays in this tree may have long segments with no splits. However only the splitting vertices of $\mathcal{T}_J(\Omega_N; M)$ and certain other vertices related to these are of central importance to the subsequent analysis. With this in mind and to aid in quantification later on, we introduce the class of splitting vertices

$$\mathcal{G} = \mathcal{G}(\Omega_N) := \bigcup_{j=1}^{N} \mathcal{G}_j(\Omega_N), \text{ where for every } 1 \leq j \leq N \tag{12.2}$$

$$\mathcal{G}_j(\Omega_N) := \left\{ \gamma : \begin{array}{l} \text{there exists } v \in \Omega_N \text{ such that } \gamma \text{ is the } j\text{th splitting} \\ \text{vertex on the ray identifying } v \text{ in } \mathcal{T}_J(\Omega_N; M) \end{array} \right\}. \tag{12.3}$$

The vertices in $\mathcal{G}_j(\Omega_N)$ will be termed *the jth splitting vertices.* As dictated by the pruning mechanism, such vertices γ may occur at different heights of the tree $\mathcal{T}_J(\Omega_N; M)$, and hence could represent M-adic cubes of varying sizes. Thus the index j, which encodes the number of splitting vertices on the ray leading up to and including γ, should not be confused with the height of γ in $\mathcal{T}_J(\Omega_N; M)$. Given $\gamma \in \mathcal{G}(\Omega_N)$, we write

$$\nu(\gamma) = j \quad \text{ if } \gamma \in \mathcal{G}_j(\Omega_N), \tag{12.4}$$

and refer to $\nu(\gamma)$ as the *splitting index* of γ. Indeed $N - \nu(\gamma)$ is the splitting number of γ with respect to $\mathcal{T}_J(\Omega_N; M)$, defined as in (3.7). Note that $\mathcal{G}_1(\Omega_N)$ consists of a single element, namely the unique splitting vertex of $\mathcal{T}(\Omega_N; M)$ of minimal height. In general $\#(\mathcal{G}_j(\Omega_N)) = 2^{j-1}$; i.e., there are exactly 2^{j-1} splitting vertices of index j. We declare $\mathcal{G}_{N+1}(\Omega_N) \equiv \Omega_N$.

Another related quantity of importance is the one mentioned in part (iii) of Proposition 12.1. In view of its ubiquitous occurrence in the sequel, we set up the following notation. For $\gamma \in \mathcal{G}_j(\Omega_N)$, $1 \leq j \leq N-1$, we denote

$$\lambda(\gamma) = \lambda_j(\gamma) := h^*_\gamma \text{ defined as in Proposition 12.1.} \tag{12.5}$$

Thus $\lambda_j(\gamma) > h(\gamma)$ is the smallest height for which the Euclidean separation condition (12.1) can be ensured for the two descendants of γ. We refer to an element of $\{\lambda_j(\gamma) : \gamma \in \mathcal{G}_j(\Omega_N)\}$ as a *jth fundamental height* of Ω_N. There could be at most 2^{j-1} such heights. The collection of all fundamental heights will be denoted by $\mathcal{R}$; it will play a vital role in the remainder of the article, specifically in the random construction outlined in Chapter 14. The two descendants of $\gamma \in \mathcal{G}_j(\Omega_N)$ at height $\lambda_j(\gamma)$, are called the *jth basic slope cubes.* The entirety of jth basic slope cubes as γ ranges over $\mathcal{G}_j(\Omega_N)$ is termed $\mathcal{H}_j(\Omega_N)$. More precisely,

$$\mathcal{H}_j(\Omega_N) := \left\{\theta \; : \; \begin{array}{l} \text{there exists } \omega \in \Omega_N \text{ and } \gamma_j \in \mathcal{G}_j(\Omega_N) \\ \text{such that } \omega \in \theta \subsetneq \gamma_j \text{ and } h(\theta) = \lambda(\gamma_j) \end{array}\right\}. \tag{12.6}$$

Note that every jth basic slope cube θ is either itself a $(j+1)$th splitting vertex

$\gamma_{j+1} \in \mathcal{G}_{j+1}(\Omega_N)$, or uniquely identifies such a vertex in the sense that there exists a non-splitting ray in the slope tree rooted at θ that terminates at γ_{j+1}. In either event, we say that *$\gamma_{j+1} \in \mathcal{G}_{j+1}(\Omega_N)$ is identified by $\theta \in \mathcal{H}_j(\Omega_N)$*. Since every $\gamma \in \mathcal{G}_j(\Omega_N)$ contributes exactly two cubes to $\mathcal{H}_j(\Omega_N)$, it follows that $\#(\mathcal{H}_j(\Omega_N)) = 2^j$. We declare $\mathcal{H}_0(\Omega_N) = \mathcal{G}_1(\Omega_N)$ and $\mathcal{H}_N(\Omega_N) = \Omega_N$.

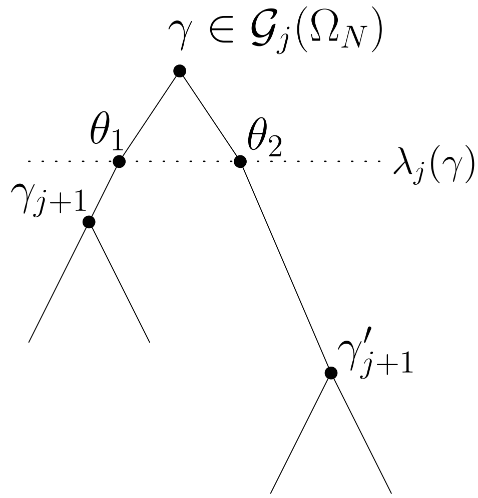

Figure 12.3: Two basic slope cubes $\theta_1, \theta_2 \in \mathcal{H}_j(\Omega_N)$ and their parent vertex $\gamma \in \mathcal{G}_j(\Omega_N)$. Notice that $\gamma_{j+1} = \theta_1$ and γ'_{j+1} are both members of $\mathcal{G}_{j+1}(\Omega_N)$.

The following implication of the Euclidean separation condition (12.1) will be convenient for later use.

Corollary 12.4. *Given a splitting vertex γ of $\mathcal{T}_J(\Omega_N)$, define*

$$\rho_\gamma := \sup\{|a-b| : a \in \gamma_1 \cap \Omega_N,\ b \in \gamma_2 \cap \Omega_N\}, \tag{12.7}$$

$$\delta_\gamma := \inf\{|a-b| : a \in \gamma_1 \cap \Omega_N,\ b \in \gamma_2 \cap \Omega_N\}, \tag{12.8}$$

where γ_1 and γ_2 are the two children of γ in $\mathcal{T}_J(\Omega_N; M)$. Then, the two quantities ρ_γ and δ_γ are comparable; i.e., $\delta_\gamma \le \rho_\gamma \le (1 + 2\sqrt{d}C_0^{-1})\delta_\gamma$. Moreover, $\rho_\gamma \le C_1 M^{-\lambda(\gamma)}$ for some constant C_1 depending only on C_0 and d.

Proof. Using part (iii) of Proposition 12.1 and the notation set up in (12.5), we observe that $\gamma_i \cap \Omega_N \subseteq w_i$ where w_i is the only descendant of γ_i at height $\lambda(\gamma)$, so that $\delta_\gamma = \text{dist}(w_1, w_2) \geq C_0 M^{-\lambda(\gamma)}$. Let a_i, b_i be points in the closures of $\gamma_i \cap \Omega_N$, $i = 1, 2$ such that $\delta_\gamma = |a_1 - a_2|$, $\rho_\gamma = |b_1 - b_2|$. Then

$$\begin{aligned}
\rho_\gamma = |b_1 - b_2| &\leq |a_1 - b_1| + |a_2 - a_1| + |a_2 - b_2| \\
&\leq |a_2 - a_1| + \text{diam}(w_1) + \text{diam}(w_2) \\
&\leq |a_2 - a_1| + 2\sqrt{d} M^{-\lambda(\gamma)} \\
&\leq \delta_\gamma + 2\sqrt{d} C_0^{-1} \delta_\gamma,
\end{aligned}$$

where the third inequality above follows from the fact that w_i is itself a cube of sidelength $M^{-\lambda(\gamma)}$. This shows that ρ_γ and δ_γ are comparable. In order to obtain the stated upper bound on ρ_γ, we observe that $\delta_\gamma \leq (C_0 + 2\sqrt{d}) M^{-\lambda(\gamma)+1}$ by (12.1). □

12.3 Binary representation of Ω_N

The classes of basic slope cubes $\mathcal{H}_j(\Omega_N)$ allow us to represent each element in Ω_N in terms of a unique N-long binary sequence as follows. Since every splitting vertex of $\mathcal{T}_J(\Omega_N)$ has exactly two children, one of them must be larger (or older) than the other in the lexicographic ordering. Let us agree to call the older (respectively younger) child of a vertex v its 0th (respectively 1st) offspring. For $1 \leq j \leq N$, we define a bijective map $\Psi_j : \{0, 1\}^j \to \mathcal{H}_j(\Omega_N)$ inductively as follows. For $j = 1$,

$$\Psi_1(i) := \begin{cases} \text{the unique element of } \mathcal{H}_1(\Omega_N) \\ \text{descended from the } i\text{th child of } \gamma_1 \ , \end{cases} \tag{12.9}$$

where $i = 0, 1$, and γ_1 is the single element in $\mathcal{H}_0(\Omega_N) = \mathcal{G}_1(\Omega_N)$. In general if Ψ_j has been defined, then for $\bar{\epsilon} \in \{0, 1\}^j$ and $i = 0, 1$, we set

$$\Psi_{j+1}(\bar{\epsilon}, i) := \begin{cases} \text{the unique element of } \mathcal{H}_{j+1}(\Omega_N) \\ \text{descended from the } i\text{th child of } \gamma_{j+1}, \end{cases} \tag{12.10}$$

where γ_{j+1} is the unique element of $\mathcal{G}_{j+1}(\Omega_N)$ identified by $\Psi_j(\bar{\epsilon})$.

The map Ψ_N provides the claimed bijection of $\{0,1\}^N$ onto Ω_N. In fact, the discussion above yields the following stronger conclusion, the verification of which is straightforward.

Proposition 12.5. *Let $\mathcal{H}_j(\Omega_N)$ be as in* (12.6)*.*

(i) The collection of vertices

$$\mathcal{H}(\Omega_N) := \bigcup_{j=1}^{N} \left\{ (\theta_1, \cdots, \theta_j) \middle| \begin{array}{l} \exists\, \omega \in \Omega_N, \text{ such that } \omega \in \theta_k, \\ \theta_k \in \mathcal{H}_k(\Omega_N),\ 1 \le k \le j \end{array} \right\} \bigcup \{\gamma_1\} \tag{12.11}$$

is a tree rooted at $\gamma_1 \in \mathcal{H}_0(\Omega_N)$ of height N, in which $(\theta_1, \cdots, \theta_j, \theta_{j+1})$ is a vertex of height $(j+1)$ and a child of $(\theta_1, \cdots, \theta_j)$. Every element $\theta_j \in \mathcal{H}_j(\Omega_N)$ identifies a vertex $(\theta_1, \cdots, \theta_j)$ of the jth generation in this tree.

(ii) Let $\mathcal{B}_N$ denote the full binary tree of height N, namely the tree $\mathcal{T}_N([0,1);2)$. The map $\Psi : \mathcal{B}_N \to \mathcal{H}(\Omega_N)$ defined by

$$\begin{aligned} \Psi(\emptyset) &= \text{ the unique element } \gamma_1 \in \mathcal{H}_0(\Omega_N), \\ \Psi(\bar{\epsilon}) &= \Psi_j(\bar{\epsilon}) \quad \text{ if } \quad \bar{\epsilon} \in \{0,1\}^j,\ 1 \le j \le N, \end{aligned} \tag{12.12}$$

with Ψ_j as in (12.10) *is a tree isomorphism in the sense of Definition 3.1.*

Although we will not need to use it, an analogous argument shows that the class of splitting vertices $\mathcal{G}(\Omega_N)$ is isomorphic to $\mathcal{B}_{N-1}$.

Chapter 13

Families of intersecting tubes, revisited

The finite set of directions Ω_N created in Proposition 12.1 forms the basis of the construction of Kakeya-type sets. Predictably, the sets of interest that verify the conclusion of Theorem 1.3 will be the union of a family of tubes, with each tube assigned a slope from Ω_N. Each tube is based on a suitably fine subcube of the d-dimensional *root hyperplane*, $\{0\} \times [0,1)^d$. The tree depicting the root hyperplane, more precisely the full M-adic tree of dimension d and height J will be termed the *root tree*.

As in Chapter 6, for $0 \leq k \leq J$, let $\mathcal{Q}(k)$ be the collection of all vertices of height k in the root tree; i.e.,

$$\mathcal{Q}(k) := \left\{ Q \;:\; Q \in \mathcal{T}(\{0\} \times [0,1)^d; M),\ h(Q) = k \right\}. \tag{13.1}$$

Geometrically, and in view of the discussion in Section 3.4, a member Q of $\mathcal{Q}(k)$ is an M-adic cube of sidelength M^{-k} of the form

$$Q = \{0\} \times \prod_{\ell=1}^{d} \left[\frac{j_\ell}{M^k}, \frac{j_{\ell+1}}{M^k} \right), \quad \text{where} \;\; (j_1, j_2, \cdots j_d) \in \{0, 1, \cdots, M^k - 1\}^d, \tag{13.2}$$

so that $\#(\mathcal{Q}(k)) = M^{kd}$. In view of the above, and for the purpose of distinguishing vertices of the root and the slope trees, a vertex in the root tree is termed a *spatial cube*. For reasons to be made clear in a moment, an element of $\mathcal{Q}(J)$ (i.e., a youngest vertex of the root tree) is of added significance and will be called a *root cube*.

Given a fixed constant $A_0 \geq 1$, and for $t \in \mathcal{Q}(J)$, $\omega \in \Omega_N$, we define a *tube rooted at t with orientation ω* to be the set

$$P_{t,\omega} := \widetilde{Q}_t + [0, 10A_0]\omega = \big\{s + r\omega : s \in \widetilde{Q}_t,\ 0 \leq r \leq 10A_0\big\}. \tag{13.3}$$

Here $\widetilde{Q}_t$ denotes the c_d-dilate of the cube t; i.e., the cube with the same centre as t but with c_d times its sidelength, for a small positive constant c_d specified in Corollary 7.2. For instance, the choice $c_d = d^{-2d}$ will suffice. Thus $P_{t,\omega}$ is essentially a $(d+1)$-dimensional cylinder of constant length and with cubical cross-section of sidelength $c_d M^{-J}$ perpendicular to the x_1-axis. An algorithm σ that assigns to every root $t \in \mathcal{Q}(J)$ a slope $\sigma(t) \in \Omega_N$ produces, according to the prescription (13.3), a family of tubes of cardinality M^{Jd}, and a corresponding set

$$\mathbb{K}(\sigma) = \mathbb{K}(\sigma; N, J) := \bigcup \{P_{t,\sigma(t)} : t \in \mathcal{Q}(J)\}. \tag{13.4}$$

While this definition is quite general, in our applications the slope assignment map σ will be chosen to be weakly sticky in the sense of Definition 13.6 and as a mapping between the trees representing roots and slopes respectively; specifically,

$$\sigma : \mathcal{T}_J(\{0\} \times [0,1)^d; M) \to \mathcal{T}_J(\Omega_N; M).$$

Random slope assignment algorithms will be prescribed in the next section, but for now we record some properties of general sets of the form $\mathbb{K}(\sigma)$ generated by an arbitrary σ.

13.1 Tubes and a point

A crucial component of the proof of Proposition 3.7, amplified in Chapter 15, is to identify when a given point x belongs to a union of tubes of the form (13.4). In our applications, the set $\mathbb{K}(\sigma)$ in (13.4) will be probabilistically generated by random weakly sticky maps, and we will need to estimate the likelihood of such an inclusion. But many major ingredients of the argument pertain to general sets $\mathbb{K}(\sigma)$ generated by an arbitrary weakly sticky σ. We discuss these features here.

Directly analogous to Lemma 7.5 (a), we have the following lemma.

Lemma 13.1. *Let $x \in \mathbb{R}^{d+1}$, $A_0 \le x_1 \le 10A_0$. If the parameter C_0 used in the pruning of the slope tree $\mathcal{T}(\Omega; M)$ (see Proposition 12.1) is chosen sufficiently large relative to the constant A_0 in* (13.3)*, then the following property holds: for any $t \in \mathcal{Q}(J)$, there exists at most one $v(t) \in \Omega_N$ such that $x \in P_{t,v(t)}$.*

Proof. If there exist slopes $v, v' \in \Omega_N$ such that $x \in P_{t,v} \cap P_{t,v'}$, then the points $x - x_1 v$ and $x - x_1 v'$ must both lie in t. In other words,

$$|x_1(v - v')| = |(x - x_1 v) - (x - x_1 v')| \le \sqrt{d} M^{-J}.$$

Since $x_1 \ge A_0$, this implies that $|v - v'| \le A_0^{-1}\sqrt{d}M^{-J}$, which is $\le \frac{C_0}{2} M^{-J}$ for a choice of C_0 sufficiently large. Comparing with part (iv) of Proposition 12.1, we find this is possible in Ω_N only if $v = v'$. $\square$

The lemma above motivates the following familiar definition: for $x \in \mathbb{R}^{d+1}$ with $A_0 \le x_1 \le 10A_0$,

$$\text{Poss}(x) := \left\{ t \in \mathcal{Q}(J) : \begin{array}{l} \text{there exists } v(t) = v(t;x) \in \Omega_N \\ \text{such that } x \in P_{t,v(t)} \end{array} \right\}. \tag{13.5}$$

Lemma 13.2. *The set Poss(x) introduced in* (13.5) *can also be characterized as follows:*

$$\textit{Poss}(x) = \{t \in \mathcal{Q}(J) : t \cap (x - x_1\Omega_N) \neq \emptyset\}. \tag{13.6}$$

Thus Poss(x) is contained in an $O(M^{-J})$-neighborhood of an affine copy of Ω_N in the root hyperplane $\{0\} \times [0,1)^d$.

Proof. This lemma is just a restatement of Lemma 7.4, whose proof goes through without alteration. $\square$

The mapping

$$v : \text{Poss}(x) \to \Omega_N \quad \text{which sends} \quad t \mapsto v(t) \quad \text{with} \quad x \in P_{t,v(t)} \tag{13.7}$$

is uniquely defined by Lemma 13.1. It captures for every $t \in \text{Poss}(x)$ the "correct slope" that ensures that a tube rooted at t with that slope contains x. A purely deterministic object driven by Ω_N, this map has a certain structure that is critical to the subsequent analysis. To formalize this property, let us recall the definitions of $\mathcal{G}_j(\Omega_N)$ and $\mathcal{H}_j(\Omega_N)$ from (12.3) and (12.6). We denote for every $\omega \in \Omega_N$ and $1 \leq j \leq N$,

$$\eta_j(\omega) := h(\theta) \quad \text{where} \quad \omega \subseteq \theta \in \mathcal{H}_j(\Omega_N). \tag{13.8}$$

In other words, $\eta_j(\omega)$ is the height of the jth basic slope cube on the ray identifying ω in $\mathcal{T}_J(\Omega_N; M)$. We note that $\eta_N(\omega) \equiv J$ for all $\omega \in \mathcal{H}_N(\Omega_N) = \Omega_N$.

The quantity η_j is used to define the following objects:

$$\mathcal{N}_x := \{\Phi_j(t) \, : \, t \in \text{Poss}(x), \; 0 \leq j \leq N\}, \tag{13.9}$$

$$\mathcal{M}_x := \{\Theta_j(t) \, : \, t \in \text{Poss}(x), \; 0 \leq j \leq N\}, \quad \text{where} \tag{13.10}$$

$$\Phi_j(t) := \begin{cases} \{0\} \times [0,1)^d & \text{for } j = 0 \\ \big(Q_1^*(t), \cdots, Q_j^*(t)\big) & \text{for } j \geq 1, \quad \text{and} \end{cases} \tag{13.11}$$

$$\Theta_j(t) := \begin{cases} \{1\} \times [0,1)^d & \text{for } j = 0 \\ \big(\theta_1(t), \cdots, \theta_j(t)\big) & \text{for } j \geq 1. \end{cases} \tag{13.12}$$

Here for $j \geq 1$, the cube $Q_j^*(t)$ is a cube in the root hyperplane containing t. In contrast, $\theta_j(t)$ is a vertex in $\mathcal{H}_j(\Omega_N)$, hence a cube in $\{1\} \times [0,1)^d$, containing the point $v(t) \in \Omega_N$. Furthermore, both cubes are located at the same height in their respective trees and obey the defining properties

$$t \subseteq Q_j^*(t), \quad v(t) \in \theta_j(t), \quad \text{and} \quad h(Q_j^*(t)) = h(\theta_j(t)) = \eta_j(v(t)). \tag{13.13}$$

We pause briefly to clarify the definitions (13.11) and (13.12) (see Figure 13.1). Given any $t \in \text{Poss}(x)$, we pick on the ray identifying t the vertices that lie at the same height as the basic slope cubes of $v(t)$. The entries of the vector $\Phi_j(t)$ are the first j chosen vertices on this ray. On the other hand, $\Theta_j(t)$ consists of the first j basic slope cubes containing $v(t)$. The vectors $\Phi_N(t)$ and $\Theta_N(t)$ identify t and $v(t)$ respectively. For reasons to emerge shortly in Lemma 13.4, we view the collection $\mathcal{N}_x$ as a tree, in which $\Phi_j(t)$ is a vertex of height j, and $\Phi_{j+1}(t)$ is a child of $\Phi_j(t)$. As we have already noted, the set $\text{Poss}(x)$, and hence the youngest generation of $\mathcal{N}_x$, contains all possible roots that could support tubes with directions in Ω_N containing x. For an arbitrary σ, it is therefore natural to phrase a necessary criterion for the inclusion $x \in \mathbb{K}(\sigma)$ in terms of $\mathcal{N}_x$. For this reason we choose to call $\mathcal{N}_x$ the *reference tree*, and its defining cubes $Q_j^*(t)$ as *reference cubes*. The collection $\mathcal{M}_x$ should be thought of as the "image" of $\mathcal{N}_x$ on the slope side, and hence a tree as well, with $\Theta_j(t)$ being a vertex of the jth generation and the parent of $\Theta_{j+1}(t)$. In fact, $\mathcal{M}_x$ is a subtree of $\mathcal{H}(\Omega_N)$ defined as in (12.11). In view of Proposition 12.5, any vertex $\Theta_j(t)$ of height $j \geq 1$ in $\mathcal{M}_x$ is identified with the j-long binary sequence $\Psi^{-1}(\Theta_j(t))$.

Given the constraints of our pruning mechanism in Proposition 12.1, the "correct slope" map $t \mapsto v(t)$ need not be sticky as a mapping from $\mathcal{T}_J(\text{Poss}(x); M)$ to $\mathcal{T}_J(\Omega_N; M)$. It does however possess a weak variant of the stickiness property that we specify in the next lemma. As we will see in Lemma 13.4, this milder substitute is able to achieve two goals that are of fundamental relevance to this study. First, it assigns a tree structure to $\mathcal{N}_x$ and $\mathcal{M}_x$. Second, it is strong enough to lift v as a sticky map from $\mathcal{N}_x \to \mathcal{M}_x$.

Lemma 13.3. *There is a sufficiently large choice of the parameter C_0 in Proposition 12.1 for which the following conclusion holds. Let $x \in \mathbb{R}^{d+1}$ with $A_0 \leq x_1 \leq 10A_0$, $t, t' \in$ Poss(x) and $u = D(t, t')$. Set $w = D(v(t), v(t'))$, so that $w \in \mathcal{G}(\Omega_N)$, the class*

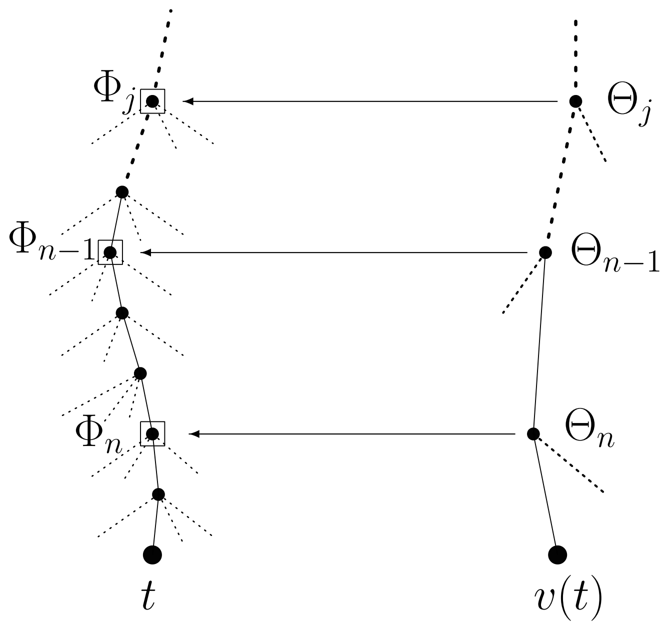

Figure 13.1: The pull-back mechanism used to define $\mathcal{N}_x$, for $M = d = 2$.

of splitting vertices defined in (12.2)*. Then*

$$h(u) < \lambda(w), \tag{13.14}$$

with λ *defined as in* (12.5)*.*

Remark: If v defined in (13.7) was indeed a sticky map, one would have access to the inequality $h(u) \le h(w)$. We know however that $\lambda(w) > h(w)$, and hence (13.14) should be viewed as a weak version of stickiness.

Proof. If $x \in P_{t,v(t)} \cap P_{t',v(t')}$, then by the inequality (7.1) in Lemma 7.1,

$$\begin{aligned} A_0|v(t) - v(t')| &\le |x_1||v(t) - v(t')| \\ &\le |\mathrm{cen}(t') - \mathrm{cen}(t)| + 2\sqrt{d}M^{-J} \\ &\le 2\sqrt{d}M^{-h(u)} + 2\sqrt{d}M^{-J} \le 4\sqrt{d}M^{-h(u)}, \\ \text{and thus } |v(t) - v(t')| &\le 4\sqrt{d}A_0^{-1}M^{-h(u)}. \end{aligned} \tag{13.15}$$

On the other hand, $v(t)$ and $v(t')$ each lie in distinct children of w, which must be a splitting vertex of $\mathcal{T}_J(\Omega_N; M)$. If $w \in \mathcal{G}_j(\Omega_N)$ and if γ, γ' denote the $(j+1)$th splitting vertices descended from w, then each of γ and γ' contains exactly one of $v(t)$ and $v(t')$. By Proposition 12.1(iii),

$$|v(t) - v(t')| \geq \mathrm{dist}(\gamma, \gamma') \geq C_0 M^{-\lambda_j(w)}. \tag{13.16}$$

Combining (13.15) and (13.16) we obtain

$$C_0 M^{-\lambda_j(w)} \leq 4\sqrt{d} A_0^{-1} M^{-h(u)}.$$

If the constant C_0 is chosen larger than $4\sqrt{d}A_0^{-1}$, then the inequality above implies (13.14), as claimed. □

Lemma 13.4. *The collection of vertex tuples $\mathcal{N}_x$, $\mathcal{M}_x$ defined in* (13.9)*,* (13.10) *are well-defined as trees rooted at $\{0\} \times [0,1)^d$ and $\{1\} \times [0,1)^d$ respectively, with the ancestry relation as described in the discussion leading up to Lemma 13.3. More precisely, the map v defined in* (13.7) *meets the following consistency requirements:*

(i) Let $t, t' \in Poss(x)$, $u = D(t,t')$. If the index j satisfies $\eta_j(v(t)) \leq h(u)$ then we also have $\eta_j(v(t')) \leq h(u)$, in which case $\Phi_j(t) = \Phi_j(t')$ and $\Theta_j(t) = \Theta_j(t')$.

(ii) The map from $\mathcal{N}_x \to \mathcal{M}_x$ that sends $\Phi_j(t) \mapsto \Theta_j(t)$ is well-defined and sticky.

Proof. Let $\gamma_j(t) \in \mathcal{G}_j(\Omega_N)$ denote the jth splitting vertex on the ray identifying $v(t)$. Then $\eta_j(v(t)) = \lambda_j(\gamma_j(t))$. If $\eta_j(v(t)) = \lambda_j(\gamma_j(t)) \leq h(u)$, then Lemma 13.3 implies that $\eta_j(v(t)) = \lambda_j(\gamma_j(t)) < \lambda(w)$, where $w = D(v(t), v(t'))$. Unravelling the implication of this inequality, we see that the height of the first splitting descendant of $\gamma_j(t)$ is strictly smaller than the corresponding quantity for w. Since both $\gamma_j(t)$ and w are splitting vertices lying on the ray of $v(t)$, this means that $\gamma_j(t)$ is an ancestor of w of strictly lesser height. In other words, $w \subseteq \gamma_{j+1}(t)$. Since the rays for $v(t)$ and $v(t')$ agree up to and including height $h(w)$, we conclude that their first $(j+1)$ splitting vertices are identical; i.e.,

$$\gamma_k(t) = \gamma_k(t') \text{ for } k \leq j+1. \tag{13.17}$$

Hence $\eta_k(v(t)) = \lambda_k(\gamma_k(t)) = \lambda_k(\gamma_k(t')) = \eta_k(v(t'))$ for all such k, implying one of the desired conclusions in part (i). Since

$$h(w) \geq h(\gamma_{j+1}(t)) = h(\gamma_{j+1}(t')) \geq \eta_j(v(t)) = \eta_j(v(t')),$$

the vectors $v(t)$ and $v(t')$ must agree at height η_j. Thus $\Theta_j(t) = \Theta_j(t')$. Of course if $\eta_j(v(t)) = \eta_j(v(t')) \leq h(u)$, then $\Phi_j(t) = \Phi_j(t')$. This completes the proof of the first part of the lemma.

Part (ii) is essentially a restatement of the result in part (i). To ascertain that the map is well-defined we choose $t, t' \in \text{Poss}(x)$ with $u = D(t,t')$ and $\Phi_j(t) = \Phi_j(t')$ and aim to show that $\Theta_j(t) = \Theta_j(t')$. The hypothesis $\Phi_j(t) = \Phi_j(t')$ implies that $\eta_j(v(t)) = \eta_j(v(t')) \leq h(u)$, and part (i) implies that the images match. Stickiness is a by-product of the definitions. □

Lemma 13.4 permits the unambiguous assignment of an "ideal image" (namely an edge in $\mathcal{M}_x$) to every edge of the tree $\mathcal{N}_x$, in the following sense: if every edge in the ray leading up to $\Phi_N(t)$ receives its ideal image, then $x \in P_{t,v(t)}$. To make this quantitatively precise, let us define the *reference slope function* κ as follows: for every edge e in $\mathcal{N}_x$ joining the vertices $\Phi_j(t)$ to $\Phi_{j+1}(t)$, we define a binary counter $\kappa(e)$ through the defining equation

$$\Psi^{-1} \circ \Theta_{j+1}(t) = (\Psi^{-1} \circ \Theta_j(t), \kappa(e)) \tag{13.18}$$

where Ψ is the tree isomorphism defined in Proposition 12.5. In other words, $\kappa(e)$ is zero (respectively one) if and only if the ray identifying $\Theta_{j+1}(t)$ in $\mathcal{T}_J(\Omega_N)$ passes through the 0th (respectively 1st) child of the $(j+1)$th splitting vertex identified by $\Theta_j(t)$.

Corollary 13.5. *The reference slope function κ described in* (13.18) *is well-defined, and assigns to each edge of $\mathcal{N}_x$ a unique value of 0 or 1.*

Proof. If there exist $t \neq t'$ in $\text{Poss}(x)$ such that the terminating vertex of e could be represented either as $\Phi_{j+1}(t)$ or as $\Phi_{j+1}(t')$, then Lemma 13.4 guarantees that

$\Theta_k(t) = \Theta_k(t')$ for all $k \leq j+1$, proving that $\kappa(e)$ given by (13.18) is a well-defined function on the edge set of $\mathcal{N}_x$. □

The reader may find it helpful to visualize the edges of the reference tree $\mathcal{N}_x$ with an overlay of model binary values assigned by κ, against which any other slope assignment will be tested. This intuition is made precise below.

13.2 Weakly sticky maps

Motivated by Lemma 13.3, we introduce a general notion of weak stickiness. This is a property that each random slope assignment σ prescribed in Chapter 14 will enjoy (see Lemma 14.1).

Definition 13.6. *Let σ be a height-preserving function that maps every full-length ray of $\mathcal{T}_J(\{0\} \times [0,1)^d; M)$ to a full-length ray in $\mathcal{T}_J(\Omega_N; M)$. We say that σ is* weakly sticky *if for any $t, t' \in \mathcal{Q}(J)$, $t \neq t'$, one has the relation $h(u) < \lambda(w)$, where $u = D(t,t')$ and $w = D(\sigma(t), \sigma(t'))$.*

Given a fixed point x and a union of tubes $\mathbb{K}(\sigma)$ of the form (13.4) generated by a weakly sticky slope map σ, we obtain in Lemma 13.7 below a criterion governed by the reference slope function κ for verifying whether $x \in \mathbb{K}(\sigma)$. Indeed for such σ, we can define $\mathcal{N}_x(\sigma)$ and $\mathcal{M}_x(\sigma)$ akin to (13.9) and (13.10), but using the given slope map $t \mapsto \sigma(t)$ instead of the naturally generated v given by (13.7). More precisely, we set

$$\mathcal{N}_x(\sigma) := \{\Phi_j(t;\sigma) \,:\, t \in \text{Poss}(x), 0 \leq j \leq N\}, \tag{13.19}$$

$$\mathcal{M}_x(\sigma) := \{\Theta_j(t;\sigma) \,:\, t \in \text{Poss}(x),\ 0 \leq j \leq N\}, \tag{13.20}$$

where for $j \geq 1$, both $\Phi_j(t;\sigma)$ and $\Theta_j(t;\sigma)$ are j-long vectors whose ith components are M-adic cubes of identical size, containing t in the root hyperplane and $\sigma(t)$ in the slope tree respectively. For $\Theta_j(t;\sigma)$, the ith entry is required to lie in $\mathcal{H}_i(\Omega_N)$, which uniquely specifies both vectors. In light of the preceding results in this chapter, it is not surprising that the collections (13.19) and (13.20) are trees and that σ extends to a map between these trees.

Lemma 13.7. *The following conclusions hold:*

(i) The collections $\mathcal{N}_x(\sigma)$ and $\mathcal{M}_x(\sigma)$ as in (13.19) *and* (13.20) *are well-defined as trees rooted respectively at $\{0\} \times [0,1)^d$ and $\{1\} \times [0,1)^d$. The tuples $\Phi_j(t;\sigma)$ and $\Theta_j(t;\sigma)$ are deemed vertices of generation j, and parents of $\Phi_{j+1}(t;\sigma)$ and $\Theta_{j+1}(t;\sigma)$ respectively. The map $\Phi_j(t;\sigma) \mapsto \Theta_j(t;\sigma)$ from $\mathcal{N}_x(\sigma) \to \mathcal{M}_x(\sigma)$ is well-defined and sticky.*

(ii) If e denotes the edge connecting $\Phi_j(t;\sigma)$ and $\Phi_{j+1}(t;\sigma)$ in $\mathcal{N}_x(\sigma)$, then the quantity $\iota_\sigma(e)$ defined by

$$\Psi^{-1} \circ \Theta_{j+1}(t;\sigma) = (\Psi^{-1} \circ \Theta_j(t;\sigma), \iota_\sigma(e)) \tag{13.21}$$

gives rise to a well-defined binary function on the edge set of $\mathcal{N}_x(\sigma)$.

(iii) If $x \in \mathbb{K}(\sigma)$, then there exists $t \in Poss(x)$ such that $\Theta_N(t;\sigma) = \Theta_N(t)$. In particular, this implies that

$$\Phi_j(t;\sigma) = \Phi_j(t) \quad \textit{for all} \quad 1 \le j \le N, \tag{13.22}$$

and hence that $\mathcal{N}_x$ and $\mathcal{N}_x(\sigma)$ share a common ray R identifying t with the property

$$\iota_\sigma(e) = \kappa(e) \quad \textit{for every edge } e \textit{ in } R. \tag{13.23}$$

Proof. Not surprisingly, the proof of part (i) is a verbatim reproduction of the proof of Lemma 13.4 with v replaced by σ. The relation (13.14) which played a critical role in the proof of Lemma 13.4 is now ensured by the assumption of weak stickiness of σ. Part (ii) is an easy consequence of part (i) and follows exactly the same way as Corollary 13.5 was deduced from Lemma 13.4. Finally, if $x \in \mathbb{K}(\sigma)$, then there is some $t \in \text{Poss}(x)$ such that $\sigma(t) = v(t)$. Since the chain of basic slope cubes containing any $v \in \Omega_N$ is unique, this implies that $\Theta_N(t;\sigma) = \Theta_N(t)$, and hence

$\Phi_j(t;\sigma) = \Phi_j(t)$ for all $1 \leq j \leq N$. The last equality says that t is identified by the same sequence of vertices and hence the same ray in both $\mathcal{N}_x$ and $\mathcal{N}_x(\sigma)$. If $e_1, e_2, \cdots e_N$ are the successive edges in this ray, with e_{j+1} connecting $\Phi_j(t)$ with $\Phi_{j+1}(t)$, then a consequence of the definitions (13.18), (13.21) of κ and ι_σ is that

$$(\iota_\sigma(e_1), \cdots, \iota_\sigma(e_N)) = \Psi^{-1} \circ \sigma(t) = \Psi^{-1} \circ v(t) = (\kappa(e_1), \cdots, \kappa(e_N)),$$

where Ψ is the tree isomorphism defined in Proposition 12.5, and part (iii) follows. □

We end this section with a bound on the number of vertices of the reference tree at a given height, a result that will be useful for probability computations later. In view of the characterization (13.6) of Poss(x) given in Lemma 13.2, and our construction of Ω_N, this is intuitively clear.

Lemma 13.8. *There exists a positive constant C depending on d and A_0, but uniform in $x \in [A_0, A_0+1] \times \mathbb{R}^d$, such that the number of vertices of height j in $\mathcal{N}_x$ is bounded above by $C2^j$.*

Proof. Let $n_j(x)$ denote the number of vertices of height j in $\mathcal{N}_x$. In view of the relations (13.9) and (13.13) defining $\mathcal{N}_x$, the cardinality $n_j(x)$ equals the number of spatial cubes in the collection

$$\{Q_j^*(t) : \; t \in \text{Poss}(x), \; t \subseteq Q_j^*(t), \; h(Q_j^*(t)) = \eta_j(v(t))\}, \tag{13.24}$$

so we proceed to count the number of such cubes $Q_j^*(t)$. Let us recall from the definition (13.5) of Poss(x) that $x \in P_{t,v(t)}$. This implies that $x - x_1 v(t) \in t$, and hence for $\theta_j(t)$ as in (13.13),

$$\begin{aligned}
|\text{cen}(Q_j^*(t)) &- x + x_1 \text{cen}(\theta_j(t))| \\
&\leq |\text{cen}(Q_j^*(t)) - \text{cen}(t)| + |x_1||\text{cen}(\theta_j(t)) - v(t)| + |\text{cen}(t) - x + x_1 v(t)| \\
&\leq \sqrt{d} M^{-\eta_j(v(t))} + (A_0+1)\sqrt{d} M^{-\eta_j(v(t))} + \sqrt{d} M^{-J} \\
&\leq 4A_0 \sqrt{d} M^{-\eta_j(v(t))}.
\end{aligned}$$

Let us unravel the geometric implications of the inequality above. For a given $\theta_j(t)$ containing $v(t)$, there are at most a constant number $C(d, A_0)$ of M-adic cubes of sidelength same as $\theta_j(t)$ (hence candidates for $Q_j^*(t)$) whose centres are within distance $4A_0\sqrt{d}M^{-\eta_j(v(t))}$ of $x - x_1\text{cen}(\theta_j(t))$. On the other hand, each $\theta_j(t) \in \mathcal{H}_j(\Omega_N)$, and hence the total number of possible $\theta_j(t)$ as t ranges over $\text{Poss}(x)$ is at most $\#(\mathcal{H}_j(\Omega_N)) = 2^j$, by Proposition 12.5. Since $n_j(x)$ is the cardinality of the collection in (13.24), the observations above lead to the bound $n_j(x) = O(2^j)$ as claimed. $\square$

Chapter 14

Random construction of Kakeya-type sets

Motivated by the generalities laid out in the previous chapter, specifically Lemmas 13.4 and 13.7, we now proceed to describe a randomized algorithm for generating a class of weakly sticky slope assignments σ. Let us recall the class $\mathcal{R}$ of fundamental heights of Ω_N defined in (12.5) and the discussion thereafter.

We start with a collection of independent and identically distributed Bernoulli($\frac{1}{2}$) random variables

$$\mathbb{X} := \{X_Q : Q \in \mathcal{Q}(k),\ k \in \mathcal{R}\}, \tag{14.1}$$

with $\mathcal{Q}(k)$ defined as in (13.1). The collection $\mathbb{X}$ therefore assigns, for every fundamental height k an independent binary random variable to every M-adic cube of sidelength M^{-k} in the root hyperplane. We use $\mathbb{X}$ as the randomization source for our construction.

Let h_0 denote the height of the single element $\theta_0 \in \mathcal{G}_1(\Omega_N) = \mathcal{H}_0(\Omega_N)$, in other words, the first splitting vertex of $\mathcal{T}_J(\Omega_N; M)$. We define $\sigma(Q_0) \equiv \theta_0$ for all $Q_0 \in \mathcal{Q}(h_0)$. At the first step of the randomization process, each $Q_0 \in \mathcal{Q}(h_0)$ is decomposed into subcubes Q_1 of sidelength M^{-h_1} where $h_1 = \lambda_1(\theta_0) > h_0$. We call these subcubes the *first basic spatial cubes*. Each first basic spatial cube Q_1 receives from the Bernoulli collection $\mathbb{X}$ defined in (14.1) a value of X_{Q_1}, which is either zero or one. Recalling from (12.9) that

$$\Psi_1(X_{Q_1}) \in \mathcal{H}_1(\Omega_N), \quad \text{and that} \quad h(\Psi_1(X_{Q_1})) = h_1,$$

we define

$$\sigma(Q_1) = \sigma_{\mathbb{X}}(Q_1) = \Psi_1(X_{Q_1})$$

for any first basic spatial cube Q_1. Each element of $\mathcal{H}_1(\Omega_N)$, and hence each $\sigma(Q_1)$, is either a second splitting vertex of Ω_N or the identifier of one. If the root cube Q_1 already maps into a second splitting vertex under σ, no further action is needed for it in step one. Now, suppose there exists $\gamma \in \mathcal{G}_2(\Omega_N)$ such that $h(\gamma) > h_1$. Then for any cube Q_1 for which $\Psi_1(X_{Q_1})$ is the unique ancestor of γ at height h_1, we decompose Q_1 into subcubes Q_1' of sidelength $M^{-h(\gamma)}$ and set $\sigma(Q_1') = \gamma$ for all such $Q_1' \subsetneq Q_1$. Thus, at the end of the first step,

(a) we have obtained a partition of the root hyperplane into first basic spatial cubes, and randomly assigned each such cube a first basic slope cube in $\mathcal{H}_1(\Omega_N)$ of the same height, namely $\lambda_1(\theta_0) = h_1$.

(b) If the vertices in $\mathcal{G}_2(\Omega_N)$ occur at different heights, then predicated on the random assignment in part (a) certain first basic spatial cubes could subdivide further to generate a different partition of the root hyperplane, say $\{\mathcal{Q}_1(\gamma) : \gamma \in \mathcal{G}_2(\Omega_N)\}$. Each cube $Q_1' \in \mathcal{Q}_1(\gamma)$ is of height $h(\gamma)$ and is mapped to γ. We will refer to Q_1' as *a spatial cube of second splitting height*. Thus a first basic spatial cube is either itself a spatial cube of second splitting height, or is uniformly partitioned into a disjoint union of such cubes.

In general, the jth step of the construction generates a random and possibly non-uniform partition of the root hyperplane into spatial cubes Q_j' of $(j+1)$th splitting height. Each Q_j' is the terminal member of a descending chain

$$Q_j' \subseteq Q_j \subsetneq Q_{j-1}' \subseteq Q_{j-1} \subsetneq \cdots \subsetneq Q_1' \subseteq Q_1, \tag{14.2}$$

where for every $k \leq j$, Q_k is a kth basic spatial cube, and Q_k' is a spatial cube

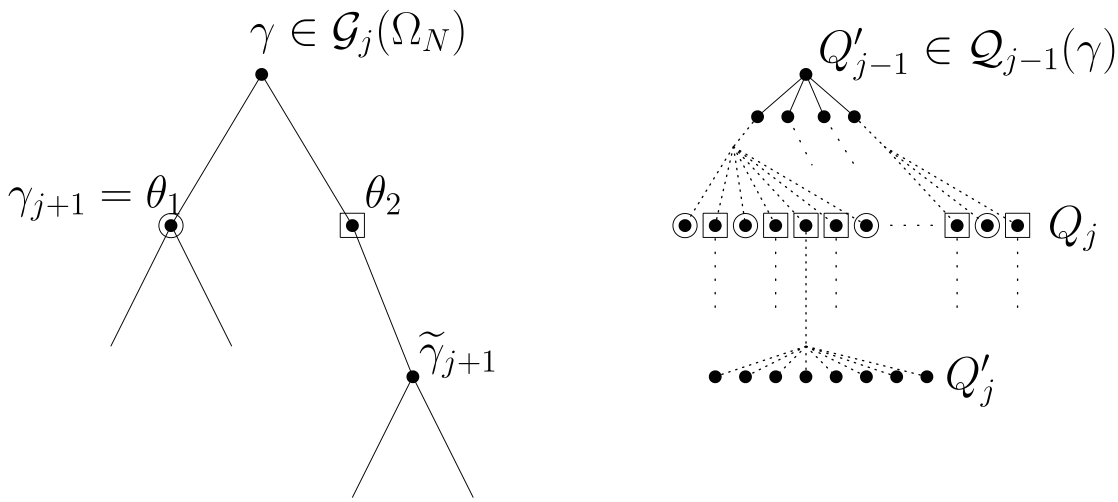

Figure 14.1: A pictorial representation of the basic slope and root cubes and a typical slope assignment. Vertices Q_j for which $X_{Q_j} = 0$ are indicated by a circle and assigned θ_1; others are indicated by squares and assigned θ_2. For the squared vertices, a further slope assignment is made at a finer level.

of $(k+1)$th splitting height. Each Q_k is mapped by σ to a kth basic slope cube in $\mathcal{H}_k(\Omega_N)$, whereas Q'_k is mapped to a splitting vertex in $\mathcal{G}_{k+1}(\Omega_N)$. All such assignments preserve heights and satisfy the following relation for a sequence of cubes as in (14.2),

$$\sigma(Q'_j) \subseteq \sigma(Q_j) \subsetneq \sigma(Q'_{j-1}) \subseteq \cdots \subsetneq \sigma(Q'_1). \tag{14.3}$$

In this way, lineages of sequences of basic spatial cubes and spatial cubes of successive splitting heights are preserved. However, notice that the full lineage of an arbitrary vertex in the root tree need not be preserved under σ.

We may classify the spatial cubes at $(j+1)$th splitting height as follows:

$$\mathcal{Q}_j(\gamma) := \{Q'_j : \sigma(Q'_j) = \gamma\}, \qquad \gamma \in \mathcal{G}_{j+1}(\Omega_N). \tag{14.4}$$

At the $(j+1)$th step each Q'_j from the collection $\mathcal{Q}_j(\gamma)$ is decomposed into subcubes Q_{j+1} of height $\lambda_{j+1}(\gamma) > h(\gamma)$. These are the $(j+1)$th *basic spatial cubes.*

Each spatial cube Q_{j+1} is assigned the binary value $X_{Q_{j+1}}$ from the Bernoulli collection $\mathbb{X}$ in (14.1). Combined with the random assignments that the basic ancestors of Q_{j+1} have received, this produces an image of Q_{j+1} under σ:

$$\sigma_{\mathbb{X}}(Q_{j+1}) := \Psi_{j+1}(X_{Q_1}, \cdots, X_{Q_{j+1}}) \in \mathcal{H}_{j+1}(\Omega_N), \qquad Q_{j+1} \subsetneq \cdots \subsetneq Q_1. \tag{14.5}$$

Each $\sigma(Q_{j+1})$ is the unique identifier of some $\gamma \in \mathcal{G}_{j+2}(\Omega_N)$. We decompose Q_{j+1} into subcubes Q'_{j+1} of height $h(\gamma)$ (in some cases no further decomposition may be needed) and set $\sigma(Q'_{j+1}) = \gamma$. This results in a newer and finer partition of the root hyperplane into spatial cubes Q'_{j+1} of $(j+1)$th splitting height, producing an analogue of (14.4) for the $(j+2)$th step and allowing us to carry the induction forward.

Continuing the procedure described above for N steps, we obtain a decomposition of the root hyperplane into a family of basic cubes of order N, each of which is of sidelength M^{-J}, and hence is by definition a root cube. Every such cube $t = Q_N(t)$ is contained in a unique chain of basic spatial cubes of lower order:

$$t = Q_N(t) \subsetneq Q_{N-1}(t) \subsetneq \cdots Q_2(t) \subsetneq Q_1(t) \tag{14.6}$$

and is assigned a slope $\sigma_{\mathbb{X}}(t) = \Psi_N(X_{Q_1}, \cdots, X_{Q_N})$ in $\mathcal{H}_N(\Omega_N) = \Omega_N$. We will shortly expand on further structural properties of the slope map $t \mapsto \sigma_{\mathbb{X}}(t)$, but first observe that it gives rise to a random set

$$K_N(\mathbb{X}) := K(\sigma_{\mathbb{X}}; N, J) \tag{14.7}$$

according to the prescription (13.4).

14.1 Features of the construction

We pause briefly to summarize the important features of the construction above:

- Randomization only occurs for cubes in the root hyperplane that correspond to the fundamental heights, though all cubes of a given fundamental height need not receive a random assignment.

- The only cubes that receive a random binary assignment from $\mathbb{X}$ are by definition the basic spatial cubes. *Unlike the basic slope cubes that constitute $\mathcal{H}_j(\Omega_N)$, a basic spatial cube Q_j is a random quantity.* For instance, the size of a jth basic spatial cube Q_j always ranges in the set $\{h(\theta) : \theta \in \mathcal{H}_j(\Omega_N)\} \subseteq \mathcal{R}$, but the exact value of the size depends on the binary assignment $X_{Q_1}, \cdots, X_{Q_{j-1}}$ received by its basic ancestors. Similarly, a spatial cube Q_j' of jth splitting height is random, though of course a splitting vertex in $\mathcal{G}_j(\Omega_N)$ is not.

- On the other hand, the random variable X_{Q_j} that a basic spatial cube Q_j *receives* is independent of all random variables used in previous or concurrent steps of the process, by virtue of our choice of (14.1). In other words,

$$\text{The collection of random variables } \left\{X_{Q_j} : Q_j \text{ basic}\right\} \text{ is independent.} \tag{14.8}$$

  This fact is vital in computing slope assignment probabilities in Chapters 15 and 16.

- Thus far, σ has been prescribed only for basic cubes and their subcubes of splitting heights. Having achieved this, it is not difficult to extend σ as a weakly sticky map between the root tree and the slope tree. We address this in the next lemma.

Lemma 14.1. *For every realization of $\mathbb{X}$, there exists a weakly sticky map*

$$\sigma_{\mathbb{X}} : \mathcal{T}_J(\{0\} \times [0,1)^d; M) \to \mathcal{T}_J(\Omega_N; M)$$

in the sense of Definition 13.6 that agrees with the slope assignment algorithm prescribed in (14.5).

Proof. The prescriptions made in (14.5) show that $\sigma_{\mathbb{X}}$ assigns a full ray in $\mathcal{T}_J(\{0\} \times [0,1)^d; M)$ to one in $\mathcal{T}_J(\Omega_N; M)$. We aim to show that $h(u) < \lambda(w)$ for $u = D(t,t')$ and $w = D(\sigma_{\mathbb{X}}(t), \sigma_{\mathbb{X}}(t'))$. Since w is by definition a splitting vertex in $\mathcal{T}_J(\Omega_N; M)$, let j denote the index such that $w \in \mathcal{G}_j(\Omega_N)$. Indeed, if $h(u) \geq \lambda(w)$, let Q be the common ancestor of t and t' at height $\lambda(w)$. Then Q is a jth basic spatial cube, which is mapped under $\sigma_{\mathbb{X}}$ to a jth basic slope cube θ. Thus θ is a common ancestor of $\sigma_{\mathbb{X}}(t)$ and $\sigma_{\mathbb{X}}(t')$ at height $\lambda(w) > h(w)$, contradicting the definition of w. □

14.2 Theorem 1.3 revisited

We will now invest our efforts into proving that with positive probability the sets $K_N(\mathbb{X})$ just created in (14.7) are of Kakeya type.

Proposition 14.2. *There exist positive absolute constants $c = c(d, M)$ and $C = C(d, M)$ obeying the property described below. For every $N \geq 1$ and Ω_N as in Proposition 12.1, the random set $K_N(\mathbb{X})$ defined in (14.7) satisfies the following inequalities:*

$$Pr\Big(\Big\{\mathbb{X} : |K_N(\mathbb{X}) \cap [0,1] \times \mathbb{R}^d| \geq c\frac{\sqrt{\log N}}{N}\Big\}\Big) \geq \frac{3}{4}, \tag{14.9}$$

$$\mathbb{E}_{\mathbb{X}}\big|K_N(\mathbb{X}) \cap [A_0, A_0+1] \times \mathbb{R}^d\big| \leq \frac{C}{N}. \tag{14.10}$$

The proof of the proposition will occupy the remainder of the main document, with the estimates (14.10) and (14.9) established in Chapters 15 and 19 respectively. Before launching into them, let us observe that these two estimates combine to generate the Kakeya-type set whose existence is claimed in Theorem 1.3 and subsequently reformulated in Proposition 3.7.

Corollary 14.3. *Given Proposition 14.2, the statement of Proposition 3.7 follows. Specifically, for every $N \geq 1$ there exists a realization of $\mathbb{X}$ for which the union of tubes defined by*

$$E_N := K_N(\mathbb{X}) \cap [A_0, A_0+1] \times \mathbb{R}^d \quad \textit{obeys} \quad \frac{|E_N^*(2A_0+1)|}{|E_N|} \underset{N\to\infty}{\longrightarrow} \infty. \tag{14.11}$$

In other words, Ω admits Kakeya-type sets.

Proof. The proof is identical to that of Proposition 6.2, where we set $a_N = c\frac{\sqrt{\log n}}{N}$ and $b_N = \frac{C}{N}$. $\square$

Chapter 15

Proof of the upper bound (14.10)

Proposition 15.1. *There exists a positive constant C possibly depending on d and M but uniform in $x \in [A_0, A_0+1]\times\mathbb{R}^d$ such that the probability $Pr(x) := Pr(x \in K_N(\mathbb{X}))$ obeys the estimate*

$$Pr(x) \leq \frac{C}{N}. \tag{15.1}$$

As a consequence, (14.10) *holds.*

Proof. The proof of (15.1) is a consequence of the three lemmas stated and proved below in this chapter. In Lemma 15.3 and following the direction laid out in [4, 3], we establish that $\mathrm{Pr}(x)$ is bounded above by the probability that the reference tree $\mathcal{N}_x$ survives a Bernoulli($\frac{1}{2}$) percolation, as described in Chapter 4. The details of the specific percolation criterion that permit this correspondence are described in Lemma 15.2. Using general facts about percolation collected in Chapter 4 and information on $\mathcal{N}_x$ observed in Chapter 13, we compute in Lemma 15.4 a bound on the survival probability that is uniform in x to obtain the claimed estimate (15.1).

Given (15.1), the upper bound in (14.10) follows easily. Since $\Omega_N \subseteq \{1\}\times[0,1)^d$, any tube, and hence $K_N(\mathbb{X})$, is contained in the compact set $[0, 10A_0]^{d+1}$. Thus

$$K_N(\mathbb{X}) \cap [A_0, A_0+1]\times\mathbb{R}^d = K_N(\mathbb{X}) \cap [A_0, A_0+1]\times[0, 10A_0]^d,$$

and hence

$$\begin{aligned}\mathbb{E}_{\mathbb{X}}\big|K_N(\mathbb{X}) \cap [A_0, A_0+1] \times \mathbb{R}^d\big| &= \mathbb{E}_{\mathbb{X}} \int_{[A_0,A_0+1]\times[0,10A_0]^d} 1_{K_N(\mathbb{X})}(x)\,dx \\ &= \int_{[A_0,A_0+1]\times[0,10A_0]^d} \mathbb{E}_{\mathbb{X}}\big(1_{K_N(\mathbb{X})}(x)\big)\,dx \\ &= \int_{[A_0,A_0+1]\times[0,10A_0]^d} \Pr(x)\,dx \\ &\le \frac{C}{N},\end{aligned}$$

completing the proof. □

Much of the groundwork for Lemma 15.3 has already been established in Section 13.1. In particular, let us recall the definition of the reference tree $\mathcal{N}_x$ and reference cubes $Q_j^*(t)$ from (13.9), (13.11), and (13.13). We will also need the reference slope function κ as in (13.18) defined on the edges of $\mathcal{N}_x$. Motivated by Lemma 13.7(iii), we define a random variable for each edge of $\mathcal{N}_x$:

$$Y_e = Y_e(\mathbb{X}) := \begin{cases} 1 & \text{if } X_{Q_{j+1}^*(t)} = \kappa(e), \\ 0 & \text{otherwise}, \end{cases} \tag{15.2}$$

where as usual e denotes the edge in $\mathcal{N}_x$ joining $\Phi_j(t)$ and $\Phi_{j+1}(t)$. As described in Chapter 4, we use Y_e to determine whether to retain or to remove the edge e in $\mathcal{N}_x$, the value zero corresponding to removal. We emphasize that a reference cube $Q_{j+1}^*(t)$ is a deterministic vertex of the tree representing the root hyperplane, and need not in general coincide with the $(j+1)$th basic spatial cube $Q_{j+1}(t)$ described in (14.6). The important point, as we will see in Lemma 15.3, is that if $x \in K_N(\mathbb{X})$, then these two cubes do match for some t and for every j.

Lemma 15.2. *The retention-removal criterion described in* (15.2) *gives rise to a well-defined Bernoulli*($\frac{1}{2}$) *percolation on* $\mathcal{N}_x$.

Proof. Since $Q_{j+1}^*(t)$ identifies the terminating vertex of the edge e, any two representations $\Phi_{j+1}(t) = \Phi_{j+1}(t')$ of this vertex gives rise to $Q_{j+1}^*(t) = Q_{j+1}^*(t')$. So

$X_{Q^*_{j+1}(t)}$ is consistently defined on the edges. We have already seen in Corollary 13.5 that κ is a well-defined function on the edge set of $\mathcal{N}_x$, hence so is Y_e. The probability that Y_e equals one is clearly $1/2$ since it is given by the Bernoulli($\frac{1}{2}$) random variable $X_{Q^*_{j+1}(t)}$. Finally, any two distinct edges e and e' must have distinct terminating vertices, and therefore end in distinct reference cubes. The random variable assignments for such cubes are independent by our assumption on $\mathbb{X}$. Hence the events of retention and removal are independent for different edges, and the result follows. □

Lemma 15.3. *Let x be a point in $[A_0, A_0+1] \times \mathbb{R}^d$. If $x \in K_N(\mathbb{X})$, then there is at least one ray of full length in $\mathcal{N}_x$ all of whose edges are retained after the percolation described by $Y_e(\mathbb{X})$. As a result, the probability $Pr(x)$ defined in Proposition 15.1 admits the bound*

$$Pr(x) \le p^*(x), \tag{15.3}$$

where $p^*(x)$ denotes the survival probability of $\mathcal{N}_x$ under the Bernoulli($\frac{1}{2}$) percolation given in (15.2).

Proof. If $x \in K_N(\mathbb{X})$, then by Lemma 13.7(iii) there exists $t \in \mathrm{Poss}(x)$ such that the ray identifying t is common to $\mathcal{N}_x$ and $\mathcal{N}_x(\sigma_{\mathbb{X}})$. Restating (13.22), this means that $\Phi_j(t) = \Phi_j(t;\sigma)$ for all $1 \le j \le N$. But the left hand side of the preceding equality identifies the (deterministic) jth reference cube containing t, whereas the right hand side represents the (random) jth basic spatial cube containing t. In other words, we find that $Q_j(t;\mathbb{X}) = Q^*_j(t)$ for all $1 \le j \le N$, and hence

$$\iota_\sigma(e) = X_{Q_{j+1}(t;\mathbb{X})} = X_{Q^*_{j+1}(t)}.$$

Combined with (15.2) and (13.23), this implies the existence of an entire ray in $\mathcal{N}_x$ (namely the one identifying t) that survives the percolation given by Y_e. Summarizing, we obtain that

$$\{\mathbb{X} : x \in K_N(\mathbb{X})\} \subseteq \left\{ \mathbb{X} : \begin{array}{c} \mathcal{N}_x \text{ survives the Bernoulli}\big(1/2\big) \\ \text{percolation dictated by } Y_e(\mathbb{X}) \end{array} \right\},$$

from which (15.3) follows. □

Lemma 15.4. *There is a positive constant* C *that is uniform in* $x \in [A_0, A_0+1] \times \mathbb{R}^d$ *such that the survival probability* $p^*(x)$ *of* $\mathcal{N}_x$ *under Bernoulli*$(\frac{1}{2})$ *percolation is* $\leq \frac{C}{N}$.

Proof. In view of Corollary 4.2, $p^*(x)$ is bounded above by

$$\left[\sum_{j=1}^{N} \frac{2^j}{n_j(x)}\right]^{-1} \text{ where } n_j(x) = \text{ number of vertices in } \mathcal{N}_x \text{ of height } j.$$

But Lemma 13.8 gives that $n_j(x) \leq C2^j$, which leads to the stated bound. □

Chapter 16

Probability estimates for slope assignments

We now turn to (14.9), where we need to establish that with high probability, the volume of space close to the root hyperplane is much more widely populated by the random set $K_N(\mathbb{X})$ than away from it. As we have already seen in Section 9 during the Cantor directions case, the proof requires detailed knowledge of the probability that a given subset of root cubes receives prescribed slope assignments. We establish the necessary probabilistic estimates in this section for easy reference in the proof of (14.9), which is presented in Chapter 19.

16.1 A general rule

To get started, let us recall from Chapter 14 that a slope assigned to a root is not completely arbitrary and has to obey the requirement of weak stickiness. The definition below, introduced to avoid vacuous root-slope combinations, draws attention to this constraint.

Definition 16.1. *Let A be a collection of root cubes and* $\Gamma_A = \{\alpha(t) : t \in A\} \subseteq \Omega_N$

a collection of slopes indexed by A. We say that the collection of root-slope pairs

$$\{(t, \alpha(t)) \,:\, t \in A \subseteq \mathcal{Q}(J),\ \alpha(t) \in \Gamma_A \subseteq \Omega_N\} \tag{16.1}$$

is sticky-admissible if there exists a realization of $\mathbb{X}$ as in (14.1) *for which the weakly sticky map $\sigma_{\mathbb{X}}$ described in Chapter 14 has the property that*

$$\sigma_{\mathbb{X}}(t) = \alpha(t) \quad \textit{for all } t \in A. \tag{16.2}$$

Given a sticky-admissible collection (16.1), we first prescribe a general algorithm for computing the probability of the event (16.2). Preparatory to stating the result, let us define two collections consisting of tuples of vertices from the root tree and the slope tree respectively:

$$\mathbb{N}(A;\alpha) := \{\Phi_j(t;\alpha) : t \in A,\ 0 \le j \le N\}, \tag{16.3}$$

$$\mathbb{M}(A;\alpha) := \{\Theta_j(t;\alpha) : t \in A,\ 0 \le j \le N\}. \tag{16.4}$$

These objects are analogous to the trees (13.19) and (13.20) introduced earlier, with the usual interpretation of $\Phi_j(t;\alpha)$ and $\Theta_j(t;\alpha)$ following those definitions. Namely, for $j \ge 1$, the element $\Theta_j(t;\alpha)$ is a vector with j entries, whose ith component represents the ith basic slope cube in $\mathcal{T}_J(\Omega_N)$ containing $\alpha(t)$. The vector $\Phi_j(t;\alpha)$ is also a j-long sequence. Its ith entry represents the unique cube containing t located at the same height as the ith entry of $\Theta_j(t;\alpha)$. This common height is $\eta_i(\alpha(t))$ defined as in (13.8). Not surprisingly, for a choice A and α that gives rise to a sticky-admissible collection (16.1), the collections $\mathbb{N}(A;\alpha)$ and $\mathbb{M}(A;\alpha)$ are indeed trees (with the 0th generations removed) that contain the information required for computing the probability of the event (16.2). This is the content of Lemma 16.2 below, which forms the computational framework for all the probability estimates in this section.

Lemma 16.2. *Let $A \subseteq \mathcal{Q}(J)$ and $\Gamma_A = \{\alpha(t) : t \in A\} \subseteq \Omega_N$ be sets for which the collection given in* (16.1) *is sticky-admissible. Then the following conclusions hold.*

(i) The collections $\mathbb{N}(A;\alpha)$ *and* $\mathbb{M}(A;\alpha)$ *defined in* (16.3) *and* (16.4) *are well-defined trees in which* $\Phi_j(t;\alpha)$ *and* $\Theta_j(t;\alpha)$ *are deemed vertices of height* j*, and parents of* $\Phi_{j+1}(t;\alpha)$ *and* $\Theta_{j+1}(t;\alpha)$ *respectively.*

(ii) If $n(A;\alpha)$ *denotes the total number of vertices in* $\mathbb{N}(A;\alpha)$ *not counting the root, then*

$$Pr(\sigma_{\mathbb{X}}(t) = \alpha(t) \text{ for all } t \in A) = 2^{-n(A;\alpha)}. \tag{16.5}$$

Proof. The proof of the first claim follows the same line of reasoning as in Lemmas 13.3 and 13.7 and is hence omitted. We turn to the proof of (16.5). Let us write

$$\Phi_j(t;\alpha) = \bigl(Q_1^*(t;\alpha), \cdots, Q_j^*(t;\alpha)\bigr) \quad \text{and} \quad \Theta_j(t;\alpha) = \bigl(\theta_1(t;\alpha), \cdots, \theta_j(t;\alpha)\bigr). \tag{16.6}$$

In order to describe the event of interest, we need to recall from (14.6) the definition of basic spatial cubes $Q_j(t)$ containing t, their role in the random construction as explained in Chapter 14, and also the definition of the maps Ψ_j and Ψ from (12.10) and Proposition 12.5. Putting these together we find that

$$\begin{aligned}
&\bigl\{\sigma_{\mathbb{X}}(t) = \alpha(t) \text{ for all } t \in A\bigr\} \\
&\quad = \Bigl\{\sigma_{\mathbb{X}}(Q_j(t)) = \theta_j(t;\alpha) \text{ for all } 1 \le j \le N \text{ and all } t \in A\Bigr\} \\
&\quad = \Bigl\{\Psi_N(X_{Q_1(t)}, \cdots, X_{Q_N(t)}) = \alpha(t) \text{ for all } t \in A\Bigr\} \\
&\quad = \bigcap_{j=1}^{N} \bigcap_{t \in A} \Bigl\{X_{Q_j(t)} = \pi_j \circ \Psi^{-1} \circ \alpha(t)\Bigr\} \\
&\quad = \bigcap_{j=1}^{N} \bigcap_{t \in A} \Bigl\{Q_j(t) = Q_j^*(t;\alpha) \text{ and } X_{Q_j^*(t;\alpha)} = \pi_j \circ \Psi^{-1} \circ \alpha(t)\Bigr\}. \quad (16.7)
\end{aligned}$$

Here π_j denotes the projection onto the jth component of an input sequence. In the first two steps of the string of equations above, we have used the definition (14.5) of σ and its weak stickiness as ensured by Lemma 14.1. To justify the last step we observe that $Q_1(t) = Q_1^*(t;\alpha)$ is non-random; further if it is given that $Q_\ell(t) = Q_\ell^*(t;\alpha)$ for all $\ell \leq j$, then the additional requirement

$$X_{Q_j(t)} = \pi_j \circ \Psi^{-1} \circ \alpha(t) \quad \text{implies} \quad Q_{j+1}(t) = Q^*_{j+1}(t;\alpha),$$

leading to the conclusion in (16.7). By virtue of our assumption of sticky-admissibility, the event described above is of positive probability; in particular the value assignment to the random variables in $\mathbb{X}$ as prescribed in (16.7) is consistent; i.e., for $t \neq t'$,

$$\pi_j \circ \Psi^{-1} \circ \alpha(t) = \pi_j \circ \Psi^{-1} \circ \alpha(t') \quad \text{whenever} \quad Q^*_j(t;\alpha) = Q^*_j(t';\alpha).$$

In view of our assumption (14.1) on the distribution of $\mathbb{X}$, the probability of the event in (16.7) is half raised to a power that equals the number of distinct cubes in the collection $\{Q^*_j(t;\alpha); 1 \leq j \leq N,\ t \in A\}$, in other words $n(A;\alpha)$. □

16.2 Root configurations

Application of Lemma 16.2 requires explicit knowledge of the structure of the trees $\mathbb{N}(A;\alpha)$ and $\mathbb{M}(A;\alpha)$, from which $n(A;\alpha)$ can be computed. These objects depend in turn on the trees depicting A and Γ_A. We now proceed to compute $n(A;\alpha)$ in some simple situations where $\#(A) \leq 4$. On one hand, the small size of A permits the classification of possible root configurations into relatively few categories, each of which gives rise to a specific $n(A;\alpha)$. On the other hand, these cases cover all the probabilistic estimates that we will need in Chapter 19.

While each root configuration requires distinct consideration, it is recommended that the reader focus on the cases when $\#(A) = 2$, and when $\#(A) = 4$ with the four roots in what we call a type 1 configuration (see Definition 16.7). These cases contain many of the main ideas needed to push through the proof of the lower bound on the size of a typical $K_N(\mathbb{X})$ claimed in (14.9), Proposition 14.2. A thorough treatment of all distinct cases when $\#(A) \leq 4$ is needed to completely establish Proposition 14.2, but focusing on the two recommended cases should make the arguments far easier to absorb upon a first reading. When $\#(A) = 2$ in particular, the reader may focus attention on Lemmas 16.3, 17.3, 18.1 and 18.2, and the application of these lemmas in the proof of Proposition 19.1. The treatment of the case of four distinct roots in type 1 configuration has been carried out on Lemmas 16.8, 17.6, 18.1 and 18.2, with the application of these lemmas occurring in the proof of Proposition 19.2, for which this is the generic case.

16.3 Notation

Throughout this chapter the following notation will be used, in conjunction with the terminology of root hyperplane, root tree and root cube already set up in Chapter 13, page 140. Since any vertex $\Phi_j(t;\alpha)$ in $\mathbb{N}(A;\alpha)$ is uniquely identified by its last component $Q_j^*(t;\alpha)$ defined as in (16.6), we write

$$\Phi_j(t;\alpha) \cong Q_j^*(t;\alpha), \tag{16.8}$$

often opting to describe the left hand side by the right. In particular if $j = N$, then $\Phi_N(t;\alpha) \cong Q_N^*(t;\alpha) = t$, in which case the latter notation is used instead of the (more cumbersome) former.

Given a vertex u in the root tree, a vertex $\omega \in \mathcal{T}_J(\Omega_N)$ and a positive integer k such that $k \le h(u) \le \lambda(w)$, we also define

$$\theta(\omega,k) := \text{ the basic slope cube containing } \omega \text{ of maximal height} \le k, \text{ and} \tag{16.9}$$

$$\begin{aligned} \mu(\omega,k) &:= j \text{ if } \theta(\omega,k) \in \mathcal{H}_j(\Omega_N) \\ &= \text{ number of basic slope cubes of height} \le k \text{ that contain } \omega, \quad \text{and} \end{aligned} \tag{16.10}$$

$$Q_u[\omega,k] := \text{ ancestor of } u \text{ in the root tree at height } h(\theta(\omega,k)). \tag{16.11}$$

Figure 16.1 on page 168 depicts these quantities. If $\omega' \subseteq \omega$ and/or $u' \subseteq u$, then it follows from the definitions above that

$$\theta(\omega,k) = \theta(\omega',k), \quad \mu(\omega,k) = \mu(\omega',k), \text{ and}$$

$$Q_u[\omega, k] = Q_{u'}[\omega, k] = Q_u[\omega', k] = Q_{u'}[\omega', k].$$

These facts will be frequently used in the sequel without further reference.

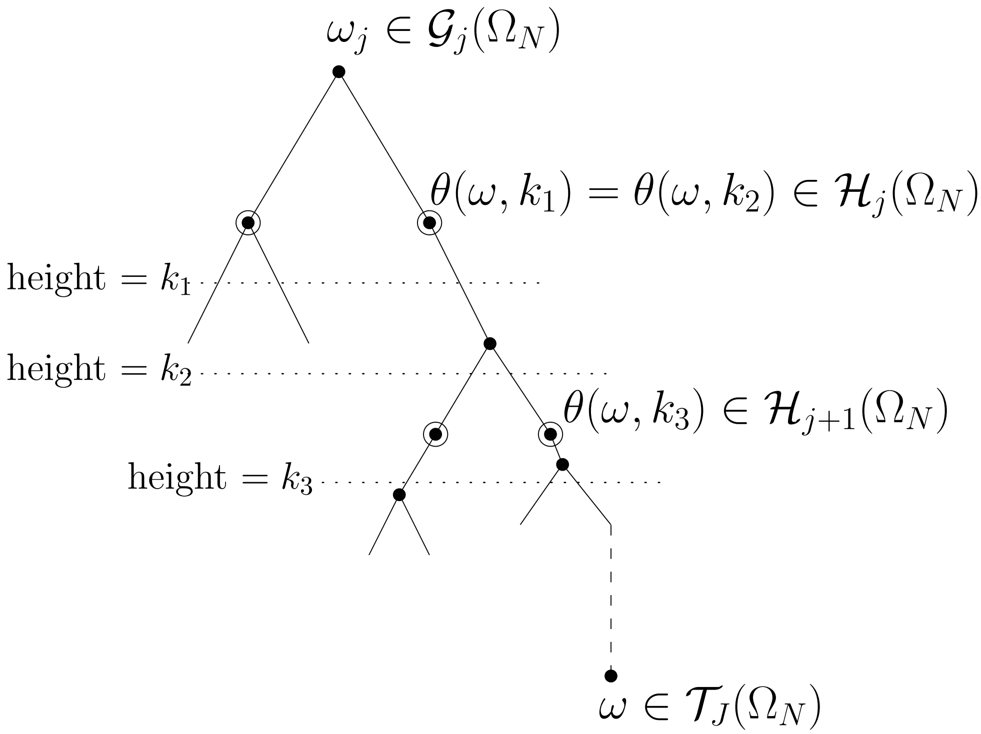

Figure 16.1: Given $\omega \in \Omega_N$ and a set of heights k_i, $i = 1, 2, 3$, the basic slope cubes $\theta(\omega, k_i)$ are identified. Here $\mu(\omega, k_1) = \mu(\omega, k_2) = j$ and $\mu(\omega, k_3) = j+1$. All vertices depicting basic slope cubes are circled.

16.4 The case of two roots

We start with the simplest case when A consists of two root cubes.

Lemma 16.3. *Let $A = \{t_1, t_2\}$ be two distinct root cubes and $\Gamma_A = \{\alpha(t_1) = v_1, \alpha(t_2) = v_2\} \subseteq \Omega_N$ be a subset of (not necessarily distinct) slopes such that $\{(t_1, v_1), (t_2, v_2)\}$ is sticky-admissible. If $u = D(t_1, t_2)$, $\omega = D(v_1, v_2)$ and $k = h(u)$,* *then $k < \lambda(\omega)$, and*

$$Pr\big(\sigma(t_1) = v_1, \sigma(t_2) = v_2\big) = \left(\frac{1}{2}\right)^{2N-\mu(\omega,k)}. \tag{16.12}$$

Proof. Since there exists a weakly sticky map σ such that $\sigma(t_i) = v_i$ for $i = 1, 2$, we see that

$$\lambda(\omega) = \lambda\big(D(v_1, v_2)\big) = \lambda\big(D(\sigma(t_1), \sigma(t_2))\big) \geq h\big(D(t_1, t_2)\big) = h(u) = k. \tag{16.13}$$

In order to establish (16.12) we invoke Lemma 16.2. The tree $\mathbb{N}(A; \alpha)$ consists of two rays terminating at $\Phi_N(t_1; \alpha) \cong t_1$ and $\Phi_N(t_2; \alpha) \cong t_2$ respectively, according to the notational rule prescribed in (16.8). Letting $u_{\mathbb{N}} = D_{\mathbb{N}}(t_1, t_2)$ denote the youngest common ancestor of t_1 and t_2 in $\mathbb{N}(A; \alpha)$, we observe that $u_{\mathbb{N}} \cong Q_u[\omega, k]$, with $Q_u[\omega, k]$ defined as in (16.11). Thus $u_{\mathbb{N}}$ lies at height $\mu(\omega, k)$ in $\mathbb{N}(A; \alpha)$. This allows us to compute $n(A; \alpha)$ as follows: $n(A; \alpha) = \mu(\omega, k) + 2(N - \mu(\omega, k)) = 2N - \mu(\omega, k)$. □

16.5 The case of three roots

Next we turn to the slightly more complex event where three distinct root cubes receive prescribed slopes. Here for the first time we observe the dependence of slope assignment probabilities on configuration types of the roots.

Definition 16.4. *Let t_1, t_2, t_2' be three distinct root cubes. We say that the ordered tuple $\mathbb{I} = \{(t_1, t_2); (t_1, t_2')\}$ with*

$$u = D(t_1, t_2), \quad u' = D(t_1, t_2'), \quad u' \subseteq u \tag{16.14}$$

is in type 1 configuration if exactly one of the following conditions hold:

(a) $u' \subsetneq u$, or

(b) $u = u' = D(t_2, t_2')$.

A tuple $\mathbb{I}$ *that obeys* (16.14) *but is not of type 1 is said to be of type 2. Thus for* $\mathbb{I}$ *of type 2, one must have* $u = u'$ *and additionally* $t = D(t_2, t_2')$ *satisfies* $t \subsetneq u$. *If* $\mathbb{I} = \{(t_1, t_2); (t_1, t_2')\}$ *with the same definitions of* u *and* u' *does not meet the containment relation required by* (16.14), *i.e., if* $u \subsetneq u'$, *then we declare* $\mathbb{I}$ *to be of the same type as* $\mathbb{I}' = \{(t_1, t_2'); (t_1, t_2)\}$.

The different structural possibilities are shown in Figure 16.2.

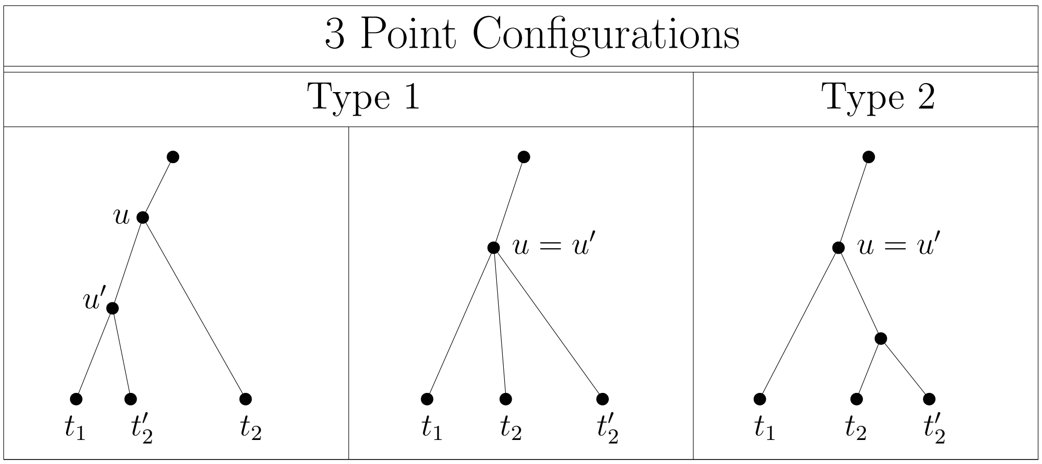

Figure 16.2: All possible configurations of three distinct root cubes.

As in Lemma 16.3, the quantity μ defined in (16.10) when evaluated at certain vertices of the slope tree dictated by $A = \{t_1, t_2, t_2'\}$ provides the value of $n(A; \alpha)$ necessary for estimating the probability in (16.5).

Lemma 16.5. *Let* $A = \{t_1, t_2, t_2'\}$ *be three distinct root cubes such that the ordered tuple* $\mathbb{I} = \{(t_1, t_2); (t_1, t_2')\}$ *obeys* (16.14) *and is of type 1. Set*

$$k = h(u), \quad k' = h(u').$$

Suppose that $\Gamma_A = \{\alpha(t_1) = v_1,\ \alpha(t_2) = v_2,\ \alpha(t_2') = v_2'\} \subseteq \Omega_N$ *is a subset of (not necessarily distinct) directions such that the collection* $\{(t_1, v_1); (t_2, v_2); (t_2', v_2')\}$ *is sticky-admissible. Then the vertices defined by*

$$\omega = D(v_1, v_2), \quad \omega' = D(v_1, v_2')$$

must satisfy the height relations

$$k \leq \lambda(\omega), \quad k' \leq \lambda(\omega') \tag{16.15}$$

and the following equality holds:

$$Pr\big(\sigma(t_1) = v_1,\ \sigma(t_2) = v_2,\ \sigma(t_2') = v_2'\big) = \left(\frac{1}{2}\right)^{3N-\mu(\omega,k)-\mu(\omega',k')}.$$

Proof. The inequalities in (16.15) are proved exactly as in Lemma 16.3; we omit these. The probability is again computed using Lemma 16.2, via counting $n(A;\alpha)$. The tree $\mathbb{N} = \mathbb{N}(A;\alpha)$ now consists of three rays, terminating at $\Phi_N(t_1;\alpha)$, $\Phi_N(t_2;\alpha)$ and $\Phi_N(t_2';\alpha)$, which are identified with t_1, t_2 and t_2' respectively. Let us recall from the proof of Lemma 16.3 that $u_{\mathbb{N}} = D_{\mathbb{N}}(t_1, t_2)$ denotes the M-adic cube specifying the youngest common ancestor of t_1 and t_2 in $\mathbb{N}(A;\alpha)$. The vertex $u_{\mathbb{N}}' = D_{\mathbb{N}}(t_1, t_2')$ is defined similarly. Then using the notation (16.8),

$$u_{\mathbb{N}}' \cong Q_{u'}[\omega', k'] = Q_{t_1}[v_1, k'], \quad \text{and} \quad u_{\mathbb{N}} \cong Q_u[\omega, k] = Q_{t_1}[v_1, k]. \tag{16.16}$$

Since $k \leq k'$, it follows from (16.16) above that $u_{\mathbb{N}}' \subseteq u_{\mathbb{N}}$. If $h_{\mathbb{N}}(\cdot)$ denotes the height of a vertex within the tree $\mathbb{N}(A;\alpha)$, then (16.16) also yields

$$h_{\mathbb{N}}\big(u_{\mathbb{N}}\big) = \mu(\omega, k) \quad \text{and} \quad h_{\mathbb{N}}\big(u_{\mathbb{N}}'\big) = \mu(\omega', k'), \quad \text{so that} \quad \mu(\omega, k) \leq \mu(\omega', k').$$

Using these relations and referring to Figure 16.2, we compute $n(A;\alpha)$ as follows,

$$\begin{aligned} n(A;\alpha) &= \underbrace{h_{\mathbb{N}}\big(u_{\mathbb{N}}\big) + \big[N - h_{\mathbb{N}}\big(u_{\mathbb{N}}\big)\big]}_{\text{vertices on the ray of } t_2 \text{ in } \mathbb{N}} + \underbrace{\big[h_{\mathbb{N}}\big(u_{\mathbb{N}}'\big) - h_{\mathbb{N}}\big(u_{\mathbb{N}}\big)\big]}_{\text{vertices between } u_{\mathbb{N}} \text{ and } u_{\mathbb{N}}'} + \underbrace{2\big[N - h_{\mathbb{N}}(u_{\mathbb{N}}')\big]}_{\text{vertices below } u_{\mathbb{N}}'} \\ &= \mu(\omega, k) + \big[N - \mu(\omega, k)\big] + \big[\mu(\omega', k') - \mu(\omega, k)\big] + 2\big[N - \mu(\omega', k')\big] \\ &= 3N - \mu(\omega, k) - \mu(\omega', k'), \end{aligned}$$

which leads to the desired probability estimate by Lemma 16.2. $\square$

Lemma 16.6. *Let* $A = \{t_1, t_2, t_2'\}$ *be three distinct root cubes such that the ordered tuple* $\mathbb{I} = \{(t_1, t_2); (t_1, t_2')\}$ *obeys* (16.14) *and is of type 2. Set*

$$k = h(u) = h(u'), \quad \text{and} \quad \ell = h(t) \quad \text{where} \quad t = D(t_2, t_2') \subsetneq u = u'.$$

If $\{(t_1, v_1); (t_2, v_2); (t_2', v_2')\}$ *is a sticky-admissible collection, then the vertices*

$$\omega = D(v_1, v_2), \quad \omega' = D(v_1, v_2'), \quad \vartheta = D(v_2, v_2')$$

must satisfy the relations

$$k \leq \min\{\lambda(\omega), \lambda(\omega')\}, \quad \ell \leq \lambda(\vartheta), \quad \mu(\omega, k) = \mu(\omega', k), \tag{16.17}$$

and the following equality holds:

$$Pr\big(\sigma(t_1) = v_1,\ \sigma(t_2) = v_2,\ \sigma(t_2') = v_2'\big) = \left(\frac{1}{2}\right)^{3N-\mu(\omega,k)-\mu(\vartheta,\ell)}. \tag{16.18}$$

Proof. The first two inequalities in (16.17) are consequences of weak stickiness, since there exists a weakly sticky map σ that assigns $\sigma(t_1) = v_1$, $\sigma(t_2) = v_2$, $\sigma(t_2') = v_2'$. Thus the first inequality in (16.17) is proved as in (16.13), while the second one also follows a similar route:

$$\lambda(\vartheta) = \lambda(D(v_2, v_2')) = \lambda\big(D(\sigma(t_2), \sigma(t_2'))\big) \geq h(D(t_2, t_2')) = h(t) = \ell.$$

For the last identity in (16.17), we observe that both ω and ω' lie on the ray identifying v_1. Thus $\theta(\omega, k) = \theta(\omega', k)$ and hence $\mu(\omega, k) = \mu(\omega', k)$ by the first inequality in (16.17).

We now turn to the counting of $n(A; \alpha)$, which leads to the probability estimate (16.18) via Lemma 16.2. Using the notation introduced in the proof of Lemma 16.5, the pairwise youngest common ancestors of the last generation vertices in $\mathbb{N}(A; \alpha)$

are seen to satisfy the following:

$$u_{\mathbb{N}} = D_{\mathbb{N}}(t_1, t_2) \cong Q_u[\omega, k] = Q_{t_2}[v_2, k],$$
$$t_{\mathbb{N}} = D_{\mathbb{N}}(t_2, t_2') \cong Q_t[\vartheta, \ell] = Q_{t_2}[v_2, \ell].$$

Since the type of $\mathbb{I}$ guarantees that $k < \ell$, the relations above imply

$$t_{\mathbb{N}} \subseteq u_{\mathbb{N}} \quad \text{and hence} \quad h_{\mathbb{N}}(u_{\mathbb{N}}) = \mu(\omega, k) \le h_{\mathbb{N}}(t_{\mathbb{N}}) = \mu(\vartheta, \ell).$$

This enables us to compute, with the aid of Figure 16.2,

$$\begin{aligned} n(A;\alpha) &= \underbrace{\mu(\omega,k) + \big[N - \mu(\omega,k)\big]}_{\text{vertices on the ray of } t_1 \text{ in } \mathbb{N}} + \underbrace{\big[\mu(\vartheta,\ell) - \mu(\omega,k)\big]}_{\text{vertices between } u_{\mathbb{N}} \text{ and } t_{\mathbb{N}}} + \underbrace{2\big[N - \mu(\vartheta,\ell)\big]}_{\text{vertices below } t_{\mathbb{N}}} \\ &= 3N - \mu(\omega,k) - \mu(\vartheta,\ell). \end{aligned}$$

This is the exponent claimed in (16.18). $\square$

16.6 The case of four roots

Finally we turn our attention to four point root configurations. Depending on the relative positions of root cubes within the root tree, we can classify the configuration types as follows. Let $\mathbb{I} = \{(t_1, t_2); (t_1', t_2')\}$ be an ordered tuple of four distinct root cubes, for which

$$u = D(t_1, t_2) \text{ and } u' = D(t_1', t_2') \text{ obey } h(u) \le h(u'). \tag{16.19}$$

Then exactly one of the following conditions must hold:

$$u \cap u' = \emptyset, \tag{16.20}$$
$$u = u' = D(t_i, t_j') \text{ for all } i, j = 1, 2, \tag{16.21}$$
$$u' \subsetneq u, \tag{16.22}$$

$$u = u', \text{ and } \exists \text{ indices } 1 \le i, j \le 2 \text{ such that } D(t_i, t'_j) \subsetneq u. \tag{16.23}$$

Definition 16.7. *For an ordered tuple* $\mathbb{I} = \{(t_1, t_2); (t'_1, t'_2)\}$ *of four distinct root cubes meeting the requirement of* (16.19)*, we say that* $\mathbb{I}$ *is of*

(a) type 1 if exactly one of (16.20) *or* (16.21) *holds,*

(b) type 2 if (16.22) *holds, and*

(c) type 3 if (16.23) *holds.*

If $\mathbb{I}$ *does not meet the height relation in* (16.19)*, then* $\mathbb{I}' = \{(t'_1, t'_2); (t_1, t_2)\}$ *does, and the type of* $\mathbb{I}$ *is said to be the same as that of* $\mathbb{I}'$*.*

Several different structural possibilities for the root quadruple exist within the confines of a single type, excluding permutations within and between the pairs $\{t_1, t_2\}$ and $\{t'_1, t'_2\}$. These have been listed in Figure 16.3. We note in passing that the type definition above is slightly different from the Cantor case of Chapter 9. Here, the main motivation for the nomenclature is the classification of the unconditional probabilities of slope assignment as exemplified in (16.5), whereas in Chapter 9 a simpler analysis involving conditional probabilities only was possible.

We now proceed to analyze how the configuration types affect the slope assignment probabilities.

Lemma 16.8. *Let* $A = \{t_1, t_2, t'_1, t'_2\}$ *be a collection of four distinct root cubes such that* $\mathbb{I} = \{(t_1, t_2); (t'_1, t'_2)\}$ *obeys* (16.19) *and is of type 1. Let* $\Gamma_A = \{v_1, v_2, v'_1, v'_2\} = \{\alpha(t_i) = v_i,\ \alpha(t'_i) = v'_i,\ i = 1, 2\} \subseteq \Omega_N$ *be a choice of slopes such that the collection* $\{(t_i, v_i); (t'_i, v'_i); i = 1, 2\}$ *is sticky-admissible. Set*

$$z = D(u, u'), \quad k = h(u), \quad k' = h(u'), \quad \ell = h(z),$$

so that

$$\begin{cases} u, u' \subsetneq z, \text{ and hence } \ell < k \le k' \text{ if } (16.20) \text{ holds}, \\ u = u' = z, \text{ and hence } \ell = k = k' \text{ if } (16.21) \text{ holds.} \end{cases}$$

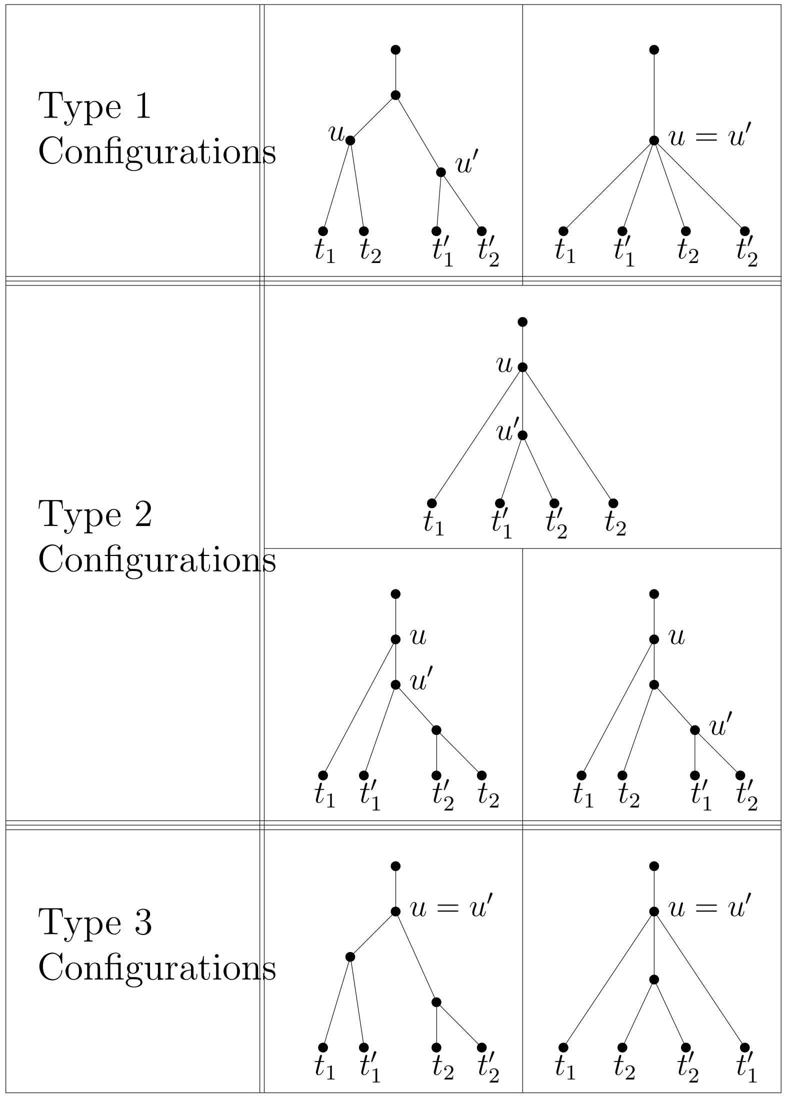

Figure 16.3: Configurations of four root cubes, up to permutations.

Then the vertices

$$\omega = D(v_1, v_2), \quad \omega' = D(v_1', v_2'), \quad v = D(\omega, \omega'),$$

must satisfy $k \le \lambda(\omega)$, $k' \le \lambda(\omega')$ *and* $\ell \le \lambda(v)$, *and the following equality holds:*

$$Pr\big(\sigma(t_i) = v_i,\ \sigma(t'_i) = v'_i,\ i = 1, 2\big) = \left(\frac{1}{2}\right)^{4N-\mu(\omega,k)-\mu(\omega',k')-\mu(v,\ell)}. \tag{16.24}$$

Proof. The proofs of the height relations may be reproduced verbatim from the previous lemmas in this chapter, so we focus only on the probability estimate. As before,

$$\begin{aligned} &u_{\mathbb{N}} = D_{\mathbb{N}}(t_1, t_2) \cong Q_u[\omega, k], \quad u'_{\mathbb{N}} = D_{\mathbb{N}}(t'_1, t'_2) \cong Q_{u'}[\omega', k'] \\ &z_{\mathbb{N}} = D_{\mathbb{N}}(u_{\mathbb{N}}, u'_{\mathbb{N}}) \cong Q_z[v, \ell] = Q_u[w, \ell] = Q_{u'}[\omega', \ell], \text{ and hence} \\ &h_{\mathbb{N}}(u_{\mathbb{N}}) = \mu(\omega, k), \quad h_{\mathbb{N}}(u'_{\mathbb{N}}) = \mu(\omega', k'), \quad h_{\mathbb{N}}(z_{\mathbb{N}}) = \mu(v, \ell). \end{aligned} \tag{16.25}$$

Since $\ell \le k \le k'$, (16.25) implies

$$u_{\mathbb{N}} \cup u'_{\mathbb{N}} \subseteq z_{\mathbb{N}}, \quad \text{and thus} \quad \mu(v, \ell) \le \min\big[\mu(\omega, k), \mu(\omega', k')\big].$$

It is important to keep in mind that $\mathbb{N}(A; \alpha)$ need not inherit the same type of structure as A. For example, if (16.20) holds, it need not be true that $u_{\mathbb{N}} \cap u'_{\mathbb{N}} = \emptyset$; indeed the vertices $u_{\mathbb{N}}$, $u'_{\mathbb{N}}$ and $z_{\mathbb{N}}$ could be distinct or (partially) coincident depending on the structure of the slope tree. Nonetheless the information collected above is sufficient to compute the number of vertices in $\mathbb{N}(A; \alpha)$ (see Figure 16.3):

$$\begin{aligned} n(A; \alpha) = &\underbrace{\mu(v, \ell)}_{\text{vertices above } z_{\mathbb{N}}} + \underbrace{\big[\mu(\omega, k) - \mu(v, \ell)\big]}_{\text{vertices between } z_{\mathbb{N}} \text{ and } u_{\mathbb{N}}} + \underbrace{\big[\mu(\omega', k') - \mu(v, \ell)\big]}_{\text{vertices between } z_{\mathbb{N}} \text{ and } u'_{\mathbb{N}}} \\ &+ \underbrace{2\big[N - \mu(\omega, k)\big]}_{\substack{\text{ancestors of } t_1 \text{ and } t_2 \\ \text{in } \mathbb{N} \text{ descended from } u_{\mathbb{N}}}} + \underbrace{2\big[N - \mu(\omega', k')\big]}_{\substack{\text{ancestors of } t'_1 \text{ and } t'_2 \\ \text{in } \mathbb{N} \text{ descended from } u'_{\mathbb{N}}}} \\ = &\ 4N - \mu(\omega, k) - \mu(\omega', k') - \mu(v, \ell). \end{aligned}$$

Combined with Lemma 16.2, this leads to (16.24). □

Lemma 16.9. *Let* $A = \{t_i, t'_i; i = 1, 2\}$ *be a collection of four distinct root cubes* *such that* $\mathbb{I} = \{(t_1, t_2); (t'_1, t'_2)\}$ *obeys* (16.19) *and is of type 2. Suppose that* $\Gamma_A = \{\alpha(t_i) = v_i,\ \alpha(t'_i) = v'_i\ ; i = 1, 2\}$ *is a choice of slopes such that the collection* $\{(t_i, v_i); (t'_i, v'_i); i = 1, 2\}$ *is sticky-admissible. Set*

$$\omega = D(v_1, v_2), \quad \omega' = D(v'_1, v'_2), \quad k = h(u), \quad k' = h(u'),$$

so that $k < k'$*. Then the following inequalities hold:* $k \leq \lambda(\omega)$*,* $k' \leq \lambda(\omega')$*. Further, there exist permutations* $\{i_1, i_2\}$ *and* $\{j_1, j_2\}$ *of* $\{1, 2\}$ *for which the quantities*

$$\vartheta = D(v_{i_2}, v'_{j_2}), \quad t = D(t_{i_2}, t'_{j_2}), \quad \ell = h(t)$$

obey the relation $\ell \leq \lambda(\vartheta)$*, and for which the probability of slope assignment can be computed as follows:*

$$Pr\big(\sigma(t_i) = v_i, \sigma(t'_i) = v'_i \text{ for } i = 1, 2\big) = \left(\frac{1}{2}\right)^{4N - \mu(\omega,k) - \mu(\omega',k') - \mu(\vartheta,\ell)}. \tag{16.26}$$

Proof. The definition of the configuration type dictates that u' is strictly contained in u, but depending on other properties of the ray joining u and u' we are led to consider several cases. If there does not exist any vertex in the root tree that is strictly contained in u and also contains t_i for some $i = 1, 2$, then any permutation of the root pairs $\{t_1, t_2\}$ and $\{t'_1, t'_2\}$ works. In particular, it suffices to choose $i_1 = j_1 = 1$, $i_2 = j_2 = 2$. In this case $t = u$, hence $\ell = k$. In particular this implies

$$\theta(\omega, k) = \theta(v_2, k) = \theta(v_2, \ell) = \theta(\vartheta, \ell), \quad \text{hence} \quad \mu(\omega, k) = \mu(\vartheta, \ell). \tag{16.27}$$

Further

$$u'_{\mathbb{N}} = Q_{u'}[\omega', k'] = Q_{t'_1}[v'_1, k'] \subseteq Q_{t'_1}[v'_1, k] = Q_u[\omega, k] = u_{\mathbb{N}}, \text{ and}$$
$$h_{\mathbb{N}}(u_{\mathbb{N}}) = \mu(\omega, k), \quad h_{\mathbb{N}}(u'_{\mathbb{N}}) = \mu(\omega', k').$$

Referring to Figure 16.3 we find that

$$\begin{aligned} n(A;\alpha) &= \mu(\omega,k) + 2\big[N - \mu(\omega,k)\big] + \big[\mu(\omega',k') - \mu(\omega,k)\big] + 2\big[N - \mu(\omega',k')\big] \\ &= 4N - 2\mu(\omega,k) - \mu(\omega',k') \\ &= 4N - \mu(\omega,k) - \mu(\omega',k') - \mu(\vartheta,\ell), \end{aligned}$$

where the last step uses one of the equalities in (16.27).

Suppose next that the previous case does not hold, and also that none of the descendants of u' lying on the rays of t'_1, t'_2 is an ancestor of t_1 or t_2. Then there is a vertex, let us call it t, such that $u' \subseteq t \subsetneq u$, and t is of maximal height in this class subject to the restriction that it is an ancestor of some t_i, which we call t_{i_2}. Thus t_{i_1} is the unique element in $\{t_1, t_2\}$ that is not a descendant of t. In this case, any permutation of $\{t'_1, t'_2\}$ works, and we can keep $j_1 = 1$, $j_2 = 2$. Then $t = D(t_{i_2}, t'_{j_2})$, $k < \ell \le k'$, and

$$\begin{aligned} u'_{\mathbb{N}} &= D_{\mathbb{N}}(t'_1, t'_2) \cong Q_{u'}[\omega', k'] = Q_{u'}[v'_{j_2}, k'], \\ t_{\mathbb{N}} &= D_{\mathbb{N}}(t_{i_2}, t'_{j_2}) \cong Q_t[\vartheta, \ell] = Q_{u'}[v'_{j_2}, \ell], \text{ and} \\ u_{\mathbb{N}} &= D_{\mathbb{N}}(t_1, t_2) \cong Q_u[\omega, k] = Q_u[\omega_0, k] = Q_{u'}[v'_{j_2}, k], \end{aligned}$$

where the last line uses the fact that $u = D(t_1, t_2, t'_1, t'_2)$, so that the second equality in that line holds $\omega_0 = D(v_1, v_2, v'_1, v'_2)$. These relations imply that

$$u'_{\mathbb{N}} \subseteq t_{\mathbb{N}} \subseteq u_{\mathbb{N}} \text{ with } h_{\mathbb{N}}(u_{\mathbb{N}}) = \mu(\omega,k),\ h_{\mathbb{N}}(u'_{\mathbb{N}}) = \mu(\omega',k'),\ h_{\mathbb{N}}(t_{\mathbb{N}}) = \mu(\vartheta,\ell). \quad (16.28)$$

Using this, we compute $n(A;\alpha)$ as follows,

$$\begin{aligned} n(A;\alpha) = {}& \underbrace{\mu(\omega,k) + \big[N - \mu(\omega,k)\big]}_{\text{vertices of } t_{i_1} \text{ in } \mathbb{N}} + \underbrace{\big[\mu(\vartheta,\ell) - \mu(\omega,k)\big] + \big[N - \mu(\vartheta,\ell)\big]}_{\text{vertices of } t_{i_2} \text{ in } \mathbb{N} \text{ below } u_{\mathbb{N}}} \\ & + \underbrace{\big[\mu(\omega',k') - \mu(\vartheta,\ell)\big]}_{\text{vertices between } t_{\mathbb{N}} \text{ and } u'_{\mathbb{N}}} + \underbrace{2\big[N - \mu(\omega',k')\big]}_{\substack{\text{vertices of } t'_1 \text{ and } t'_2 \\ \text{in } \mathbb{N} \text{ below } u'_{\mathbb{N}}}} \end{aligned}$$

$$= 4N - \mu(\omega, k) - \mu(\omega', k') - \mu(\vartheta, \ell),$$

which is the required exponent.

The last case, complementary to the ones already considered is when there exists a pair of indices, denoted $i_2, j_2 \in \{1, 2\}$ such that $t = D(t_{i_2}, t'_{j_2}) \subsetneq u'$. In this case we leave the reader to verify by the usual means that

$$t_{\mathbb{N}} \subseteq u'_{\mathbb{N}} \subseteq u_{\mathbb{N}},$$

with their heights given by the same expressions as in (16.28). Accordingly,

$$\begin{aligned} n(A; \alpha) = & \underbrace{N}_{\substack{\text{vertices of } t_{i_1} \\ \text{in } \mathbb{N}}} + \underbrace{\big[N - \mu(\omega, k)\big]}_{\substack{\text{vertices on } t'_{j_1} \\ \text{in } \mathbb{N} \text{ below } u_{\mathbb{N}}}} + \underbrace{\big[\mu(\vartheta, \ell) - \mu(\omega', k')\big]}_{\text{vertices between } t_{\mathbb{N}} \text{ and } u'_{\mathbb{N}}} + \underbrace{2\big[N - \mu(\vartheta, \ell)\big]}_{\substack{\text{vertices of } t_{i_2} \text{ and } t'_{j_2} \\ \text{in } \mathbb{N} \text{ below } t_{\mathbb{N}}}} \\ = & 4N - \mu(\omega, k) - \mu(\omega', k') - \mu(\vartheta, \ell). \end{aligned}$$

Thus, despite structural differences, all the cases give rise to the same value of $n(A; \alpha)$ that agrees with the exponent in (16.26), completing the proof. □

We pause for a moment to record a few properties of the youngest common ancestors of the roots and slopes that emerged in the proof of Lemma 16.9.

Corollary 16.10. *Let A and Γ_A be as in Lemma 16.9.*

(i) The possibly distinct vertices u, u' and t, as described in Lemma 16.9, are linearly ordered in terms of ancestry, i.e., there is some ray of the root tree that they all lie on. Depending on A, the vertex t may lie above or below u', but always in u.

(ii) The splitting vertices $\omega, \omega', \vartheta$ in the slope tree also obey certain inclusions; namely, for each of the pairs (ω, ϑ) and (ω', ϑ), one member of the pair is contained in the other.

Proof. Both u' and t lie on the ray identifying t'_{j_2} by definition, and u lies on the ray of u' by the assumption on the type of the root configuration. This establishes the first claim. The definitions also imply that $v_{i_2} \subseteq \omega \cap \vartheta$ and $v'_{j_2} \subseteq \omega' \cap \vartheta$, hence both intersections are non-empty. The second conclusion then follows from the nesting property of M-adic cubes. $\square$

Lemma 16.11. *Let $A = \{t_i, t'_i; i = 1, 2\}$ be a collection of four distinct root cubes such that $\mathbb{I} = \{(t_1, t_2); (t'_1, t'_2)\}$ is of type 3. Suppose that $\Gamma_A = \{\alpha(t_i) = v_i,\ \alpha(t'_i) = v'_i\ ; i = 1, 2\}$ is a choice of slopes such that the collection $\{(t_i, v_i); (t_i, v'_i); i = 1, 2\}$ is sticky-admissible. Set*

$$\omega = D(v_1, v_2), \quad \omega' = D(v'_1, v'_2), \quad k = h(u) = h(u').$$

Then the following relations must hold: $k \leq \lambda(\omega)$, $k \leq \lambda(\omega')$, $\mu(\omega, k) = \mu(\omega', k)$. Further, there exist permutations $\{i_1, i_2\}$ and $\{j_1, j_2\}$ of $\{1, 2\}$ such that the quantities

$$\begin{aligned} s_1 &= D(t_{i_1}, t'_{j_1}), \quad s_2 = D(t_{i_2}, t'_{j_2}), \quad \ell_1 = h(s_1), \\ \vartheta_1 &= D(v_{i_1}, v'_{j_1}), \quad \vartheta_2 = D(v_{i_2}, v'_{j_2}), \quad \ell_2 = h(s_2) \end{aligned}$$

satisfy

$$s_1 \subseteq u, \quad s_2 \subsetneq u, \quad k \leq \ell_1 \leq \ell_2, \quad \ell_i \leq \lambda(\vartheta_i) \text{ for } i = 1, 2, \tag{16.29}$$

and for which

$$Pr\big(\sigma(t_i) = v_i, \sigma(t'_i) = v'_i \text{ for } i = 1, 2\big) = \left(\frac{1}{2}\right)^{4N - \mu(\omega,k) - \mu(\vartheta_1,\ell_1) - \mu(\vartheta_2,\ell_2)}. \tag{16.30}$$

Proof. Since $\mathbb{I}$ is of type 3, $u = u'$ is the youngest common ancestor of the four elements in $\mathbb{I}$. If ω_0 is the youngest common ancestor of the slopes $\{v_i, v'_i : i = 1, 2\}$, then $\lambda(\omega_0) \geq h(u) = k$ by sticky admissibility. Thus $\theta(\omega, k) = \theta(\omega_0, k) = \theta(\omega', k)$, and therefore $\mu(\omega, k) = \mu(\omega', k)$, as claimed.

We turn to (16.29) and the probability estimate. The configuration type dictates that there exist indices $(i, j) \in \{1, 2\}^2$ such that $D(t_i, t'_j) \subsetneq u$. Among all such pairs (i, j), we pick one for which $D(t_i, t'_j)$ is of maximal height. Let us call this pair (i_2, j_2), so that $h(D(t_{i_2}, t'_{j_2})) \geq h(D(t_i, t'_j))$ for all $1 \leq i, j \leq 2$. The first three relations in (16.29) are now immediate. The last one follows from sticky admissibility and is left to the reader.

It remains to compute $n(A; \alpha)$. The structure of $\mathbb{N}(A; \alpha)$ gives that

$$\begin{aligned}
&u_{\mathbb{N}} = u'_{\mathbb{N}} = D_{\mathbb{N}}(t_1, t_2) = D_{\mathbb{N}}(t'_1, t'_2) \\
&u_{\mathbb{N}} = Q_u[\omega, k] = Q_{u'}[\omega', k] = Q_{t_1}[v_1, k] = Q_{t_2}[v_2, k], \\
&s_{i\mathbb{N}} = D_{\mathbb{N}}(t_i, t'_i) = Q_{s_i}[\vartheta_i, \ell_i] = Q_{t_i}[v_i, \ell_i], \\
&s_{i\mathbb{N}} \subseteq u_{\mathbb{N}} = D_{\mathbb{N}}(s_{1\mathbb{N}}, s_{2\mathbb{N}}) \;\text{ for } i = 1, 2, \;\text{ so that} \\
&h_{\mathbb{N}}(u_{\mathbb{N}}) = \mu(\omega, k) \leq h_{\mathbb{N}}(s_{i\mathbb{N}}) = \mu(\vartheta_i, \ell_i), \; i = 1, 2.
\end{aligned}$$

Putting these together, the number of vertices in $\mathbb{N}(A; \alpha)$ is obtained as follows,

$$\begin{aligned}
n(A; \alpha) &= \underbrace{\mu(\omega, k)}_{\text{vertices up to } u_{\mathbb{N}}} + \sum_{i=1}^{2} \underbrace{\big[\mu(\vartheta_i, \ell_i) - \mu(\omega, k)\big]}_{\text{vertices between } u_{\mathbb{N}} \text{ and } s_{i\mathbb{N}}} + \sum_{i=1}^{2} \underbrace{2\big[N - \mu(\vartheta_i, \ell_i)\big]}_{\text{vertices below } s_{i\mathbb{N}}} \\
&= 4N - \mu(\vartheta_1, \ell_1) - \mu(\vartheta_2, \ell_2) - \mu(\omega, k).
\end{aligned}$$

The probability estimate claimed in (16.30) now follows from Lemma 16.2. □

Corollary 16.12. *Let $\omega, \omega', \vartheta_1, \vartheta_2$ be as in Lemma 16.9. Then each of the pairs (ω, ϑ_1), (ω, ϑ_2), (ω', ϑ_1) and (ω', ϑ_2) has the property that one member of the pair is contained in the other.*

Proof. Since $v_{i_1} \subseteq \omega \cap \vartheta_1$, $v_{i_2} \subseteq \omega \cap \vartheta_2$, $v'_{j_1} \subseteq \omega' \cap \vartheta_1$ and $v'_{j_2} \subseteq \omega' \cap \vartheta_2$, all four intersections are nonempty, and the desired conclusion follows from the nesting property of M-adic cubes. □

As the reader has noticed, the classification of probability estimates in this section is predicated on the configuration types of the roots, *not* the slopes. Of course, such definitions of type apply equally well to slope tuples $\{(v_1, v_2); (v'_1, v'_2)\}$. Indeed, a point worth noting is that configuration types are not preserved under even sticky maps; see for example the diagram in Figure 16.4 below, where a four tuple of roots of type 1 maps to a sticky image of type 3. In view of these considerations, we shall refrain for the most part from using any type properties of slopes. In the rare instances where structural properties of slopes are relevant, a case in point being Section 19.2.3, we need to consider all possible configurations.

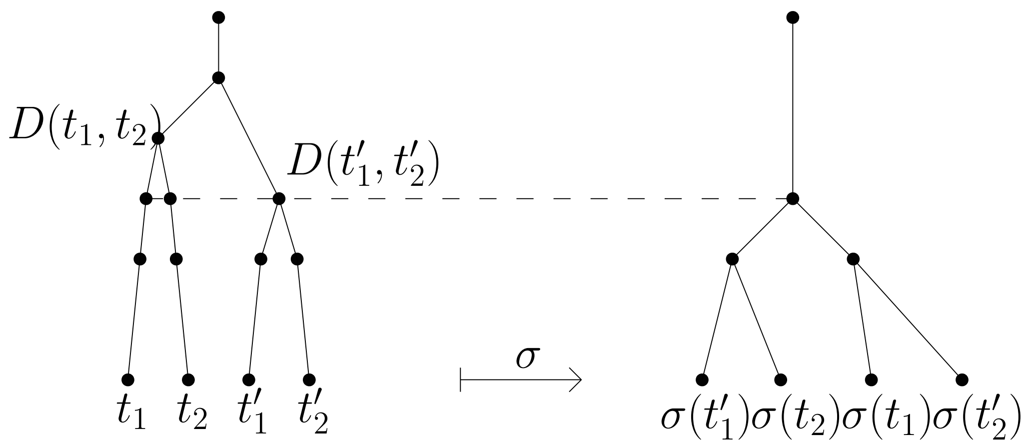

Figure 16.4: An example of a four tuple of roots of type 1 mapping to a sticky image of type 3. Notice that $D(\sigma(t_1), \sigma(t_2)) = D(\sigma(t_1'), \sigma(t_2'))$.

Chapter 17

Tube counts

A question of considerable import, the full significance of which will emerge in Chapter 19, is the following: what is the maximum possible cardinality of a sticky-admissible collection of tube tuples that admit certain pairwise intersections in a pre-fixed segment of space? The answer depends, among other things, on the size and configuration type of the roots of the tubes. In this section, we discuss these size counts for collections that are simple enough in the sense that an element in the collection is either a pair, a triple or at most a quadruple of tubes, so that the configuration type of the roots has to fall in one of the categories described in Section 16.2.

17.1 Collections of two intersecting tubes

Let us start with the case where the collection consists of pairs of tubes. To phrase the question above in more refined terms we define a collection of root-slope tuples $\mathcal{E}_2[u, \omega; \varrho]$, where u is a vertex of the root tree, ω is a splitting vertex of the slope tree, and $\varrho \in [M^{-J}, 10A_0]$ is a constant that represents the (horizontal) distance from the root hyperplane to where the intersection takes place.

$$\mathcal{E}_2[u,\omega;\varrho] := \left\{ \underset{\text{sticky-admissible}}{\{(t_1,v_1),(t_2,v_2)\}} \;\middle|\; \begin{array}{l} t_1,t_2 \in \mathcal{Q}(J),\ u = D(t_1,t_2),\ t_1 \neq t_2, \\ v_1,v_2 \in \Omega_N,\ \omega = D(v_1,v_2), \\ P_{t_1,v_1} \cap P_{t_2,v_2} \cap [\varrho, C_1\varrho] \times \mathbb{R}^d \neq \emptyset \end{array} \right\}. \tag{17.1}$$

In this context the question at the beginning of this section can be restated as: what is the cardinality of $\mathcal{E}_2[u,\omega;\varrho]$? We answer this question in Lemma 17.3 of this chapter, splitting the necessary work between two intermediate lemmas whose content will also be used in later counting arguments. To be specific, Lemma 17.1 obtains a uniform bound on a t_2-slice of $\mathcal{E}_2[u,\omega;\varrho]$ for fixed t_1, v_1 and v_2. The cardinality of the projection of $\mathcal{E}_2[u,\omega;\varrho]$ onto the t_1 coordinate is obtained in Lemma 17.2.

Lemma 17.1. *Let* $\mathcal{E}_2[u,\omega;\varrho]$ *be the collection defined in* (17.1)*, and let* $\rho_\omega = \sup\{|a-b| : a,b \in \Omega_N,\ D(a,b) = \omega\}$ *be the quantity defined in* (12.7).

(i) If $\mathcal{E}_2[u,\omega;\varrho]$ *is nonempty, then* $2C_1\varrho\rho_\omega \geq M^{-J}$.

(ii) Given a constant $C_1 > 0$ *used to define* $\mathcal{E}_2[u,\omega;\varrho]$*, there exists a constant* $C_2 = C_2(d, M, C_0, A_0, C_1) > 0$ *with the following property. For any fixed choice of* $t_1 \in \mathcal{Q}(J)$ *and* $v_1, v_2 \in \Omega_N$ *the following estimate holds:*

$$\#\big\{t_2 \in \mathcal{Q}(J) \;:\; \{(t_1,v_1),(t_2,v_2)\} \in \mathcal{E}_2[u,\omega;\varrho]\big\} \leq C_2\varrho\rho_\omega M^J. \tag{17.2}$$

Proof. The proof is illustrated in Figure 17.1. If $\{(t_1,v_1),(t_2,v_2)\}$ is a tuple that lies in $\mathcal{E}_2[u,\omega;\varrho]$, then there exists $x \in [\varrho, C_1\varrho] \times \mathbb{R}^d$ such that x also belongs to $P_{t_1,v_1} \cap P_{t_2,v_2}$. By Lemma 7.1, an appropriate version of inequality (7.1) must hold, i.e., there exists $x_1 \in [\varrho, C_1\varrho]$ such that

$$|\text{cen}(t_2) - \text{cen}(t_1) + x_1(v_2 - v_1)| \leq 2c_d\sqrt{d}M^{-J}. \tag{17.3}$$

In conjunction with Corollary 7.2, this leads to the inequality

$$M^{-J} \leq |\text{cen}(t_2) - \text{cen}(t_1)| \leq |x_1||v_2 - v_1| + 2c_d\sqrt{d}M^{-J}$$

$$\leq 2|x_1||v_2 - v_1| \leq 2C_1\varrho\rho_\omega,$$

which is the conclusion of part (i). The inequality (17.3) also implies that $\text{cen}(t_2)$ is constrained to lie in a $O(M^{-J})$ neighborhood of the line segment

$$\text{cen}(t_1) - s(v_2 - v_1), \qquad \varrho \leq s \leq C_1\varrho. \tag{17.4}$$

The length of this segment is at most $C_1\varrho|v_2 - v_1| \leq C_1\varrho\rho_\omega$, since v_1 and v_2 must lie in distinct children on ω. In view of part (i), the number of possible choices for M^{-J}-separated points $\text{cen}(t_2)$, and hence for t_2, lying within this neighborhood is $O(\varrho\rho_\omega M^J)$, as claimed in part (ii). □

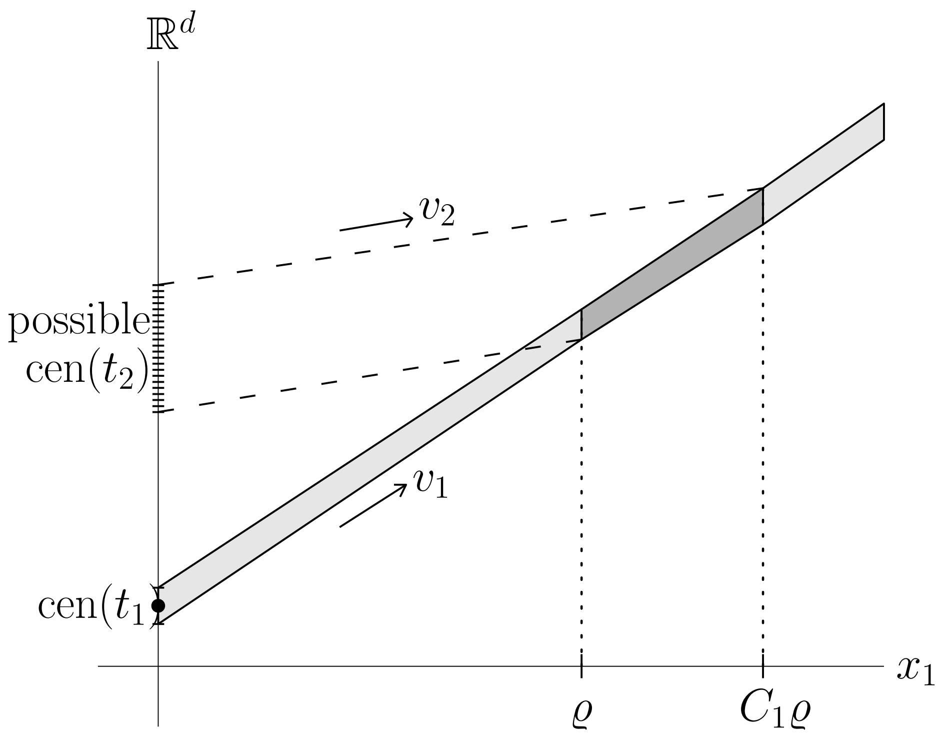

Figure 17.1: An illustration of the proof of Lemma 17.1.

Lemma 17.2. *Given $C_1 > 0$, there exists a positive constant $C_2 = C_2(d, M, A_0, C_1)$ with the following property. For any $\mathcal{E}_2[u, \omega; \varrho]$ defined as in* (17.1)*, the following* *estimate holds:*

$$\#\left\{ t_1 \in \mathcal{Q}(J) \middle| \begin{array}{c} \exists t_2 \in \mathcal{Q}(J) \text{ and } v_1, v_2 \in \Omega_N \text{ such} \\ \text{that } \{(t_1, v_1); (t_2, v_2)\} \in \mathcal{E}_2[u, \omega; \varrho] \end{array} \right\} \leq C_2 \varrho \rho_\omega M^{-(d-1)h(u)+dJ}, \tag{17.5}$$

where ρ_ω *is as in* (12.7)*.*

Proof. The proof is illustrated in Figure 17.2. If $\{(t_1, v_1), (t_2, v_2)\} \in \mathcal{E}_2[u, \omega; \varrho]$, then there exists $x = (x_1, \cdots, x_{d+1}) \in P_{t_1,v_1} \cap P_{t_2,v_2}$ with $\varrho \leq x_1 \leq C_1\varrho$. Combining inequality (17.3) obtained from Lemma 7.1 along with Corollary 7.2 as we did in Lemma 17.1, we obtain

$$\begin{aligned} |\text{cen}(t_2) - \text{cen}(t_1)| &\leq |x_1||v_2 - v_1| + 2c_d\sqrt{d}M^{-J} \\ &\leq (1 + 4c_d\sqrt{d})|x_1||v_1 - v_2| \\ &\leq C_1(1 + 4c_d\sqrt{d})\varrho\rho_\omega = C\varrho\rho_\omega, \end{aligned} \tag{17.6}$$

where the last step follows from the definition of ω. This means that $\text{cen}(t_1)$ and $\text{cen}(t_2)$ must be within distance $C\varrho\rho_\omega$ of each other. On the other hand, it is known as part of the definition of $\mathcal{E}_2[u, \omega; \varrho]$ that $u = D(t_1, t_2)$, so $\text{cen}(t_1)$ and $\text{cen}(t_2)$ must lie in distinct children of u. This forces the location of $\text{cen}(t_1)$ to be within distance $C\varrho\rho_\omega$ of the boundary of some child of u, to allow for the existence of a point $\text{cen}(t_2)$ contained in a different child and obeying the constraint of (17.6). In other words, $\text{cen}(t_1)$ belongs to the set

$$\mathcal{A}_u = \left\{ s \in u : \text{dist}(s, \text{bdry}(u')) \leq C\varrho\rho_\omega \text{ for some child } u' \text{ of } u \right\}, \tag{17.7}$$

which is the union of at most dM parallelepipeds of dimension d, with length $M^{-h(u)}$ in $(d-1)$ "long" directions and $C\varrho\rho_\omega$ in the remaining "short" direction. Note that $\rho_\omega \leq M^{-h(\omega)} \leq M^{-h(u)}$ by sticky-admissibility, hence $\varrho\rho_\omega = O(M^{-h(u)})$, which justifies this description.

Since $C\varrho\rho_\omega \geq M^{-J}$ by Lemma 17.1(i), the constituent parallelpipeds of $\mathcal{A}_u$ as described above are thick relative to the finest scale M^{-J} in all directions. The volume of $\mathcal{A}_u$ is then easily computed as

$$|\mathcal{A}_u| \leq C\varrho\rho_\omega M^{-(d-1)h(u)}.$$

Therefore the number of M^{-J} separated points $\text{cen}(t_1)$, and hence the number of possible root cubes t_1, contained in $\mathcal{A}_u$ is at most $C_2\varrho\rho_\omega M^{-(d-1)h(u)+dJ}$, as claimed. □

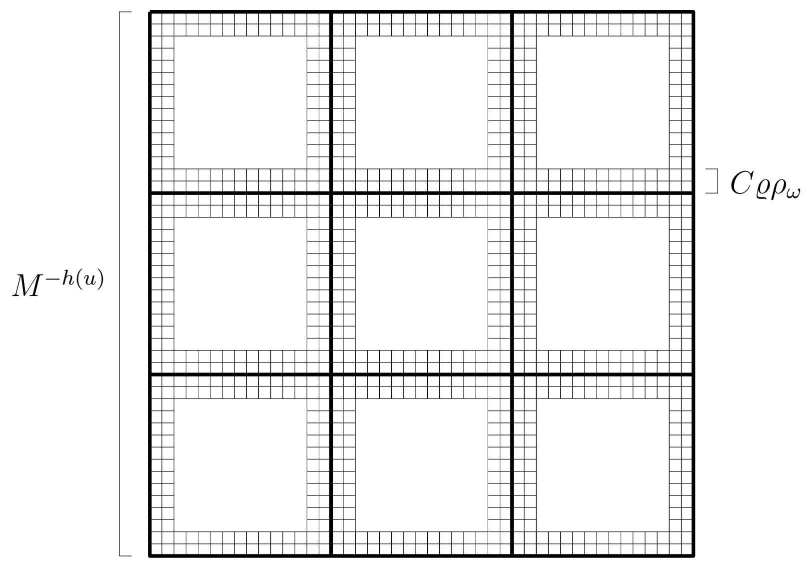

Figure 17.2: Proof of Lemma 17.2 illustrated, with $d = 2$ and $M = 3$. The outermost square is u, and the smallest squares depict the root cubes in $\mathcal{A}_u$.

Lemma 17.3. *Let* $\mathcal{E}_2[u,\omega;\varrho]$ *be the collection of pairs of tubes defined in* (17.1). *Then*

$$\#(\mathcal{E}_2[u,\omega;\varrho]) \leq C\big(\varrho\rho_\omega\big)^2 2^{2(N-\nu(\omega))} M^{-(d-1)h(u)+(d+1)J}.$$

Here $\nu(\omega)$ *denotes the index of the splitting vertex* ω, *as defined in* (12.4).

Proof. We combine the counts from Lemmas 17.1 and 17.2. For fixed t_1, v_1 and v_2, the number of possible t_2 such that $\{(t_1, v_1), (t_2, v_2)\} \in \mathcal{E}_2[u, \omega; \varrho]$ is bounded above by the quantity on the right hand side of (17.2). The number of possible t_1 is at most the right hand side of (17.5), whereas the number of possible v_1, hence also v_2, is $2^{N-\nu(\omega)}$ due to the binary nature of Ω_N as discussed in Section 12.3. The claimed size estimate of $\mathcal{E}_2[u, \omega; \varrho]$ is simply the product of all the quantities mentioned above. □

17.2 Counting slope tuples

Variations of the arguments presented in Section 17.1 also apply to more general collections. For the proof of the lower bound (14.9), we will need to estimate, in addition to the above, the sizes of collections consisting of tube triples and tube quadruples with certain pairwise intersections. The collections of tube tuples whose cardinalities are of interest are analogues of $\mathcal{E}_2[u, \omega; \varrho]$ of greater complexity, and their constructions share the common feature that the probability of slope assignment for any tube tuple within a collection is constant and falls into one of the categories classified in Chapter 16. As we have seen in that chapter, the probability depends, among other things, on certain splitting vertices of the slope tree occurring as pairwise youngest common ancestors. In particular, which subset of pairwise youngest common ancestors has to be considered, whether for root or slope, is dictated by the root configuration type. An important component of tube-counting is therefore to estimate how many possible slope tuples can be generated from a given set of such splitting vertices. Before moving on to the main counting arguments in this chapter presented in Sections 17.3 and 17.4, we observe a few facts that help in counting tuples of slopes, given some information about their ancestry.

Lemma 17.4. *(i) Given any $\Gamma \subseteq \Omega_N$, $\#(\Gamma) \leq 4$, there exist at most three distinct vertices $\{\varpi_i : i = 1, 2, 3\} \subseteq \mathcal{G}(\Omega_N)$ with the properties*

$$h(\varpi_1) \leq h(\varpi_2) \leq h(\varpi_3), \quad \varpi_2, \varpi_3 \subseteq \varpi_1, \tag{17.8}$$

such that $D(w, w') \in \{\varpi_i : i = 1, 2, 3\}$ for any $w \neq w'$, $w, w' \in \Gamma$. Of course, *the containment relations in* (17.8) *imply that* $\lambda(\varpi_2), \lambda(\varpi_3) \le \lambda(\varpi_1)$.

(ii) Suppose now that we are given $\{\varpi_i : i = 1, 2, 3\}$, *possibly distinct splitting vertices of the slope tree obeying* (17.8). *Define*

$$m = m[\varpi_1, \varpi_2, \varpi_3] := \begin{cases} 2(\nu(\varpi_3) + \nu(\varpi_2)) & \text{if } \varpi_3 \not\subseteq \varpi_2, \\ 2\nu(\varpi_3) + \nu(\varpi_2) + \nu(\varpi_1) & \text{if } \varpi_3 \subseteq \varpi_2. \end{cases} \tag{17.9}$$

Fix three distinct pairs of indices $\{(i_k, j_k) : \ i_k \ne j_k, \ 1 \le k \le 3\} \subseteq \{1, 2, 3, 4\}^2$ *with the property that* $\bigcup\{i_k, j_k : k = 1, 2, 3\} = \{1, 2, 3, 4\}$. *Then*

$$\#\big\{(w_1, w_2, w_3, w_4) \in \Omega_N^4 : D(w_{i_k}, w_{j_k}) = \varpi_k, \ 1 \le k \le 3\big\} \le C2^{4N-m}.$$

Proof. If Γ is given, we arrange all the pairwise youngest common ancestors of Γ, i.e, the vertices in $\mathcal{D}_\Gamma := \{D(w, w') : w \ne w', \ w, w' \in \Gamma\}$, in increasing order of height, where distinct vertices of the same height can be arranged in any way, say according to the lexicographic ordering. We define ϖ_3 to be a vertex of maximal height in $\mathcal{D}_\Gamma$, and ϖ_2 to be a vertex of maximal height in $\mathcal{D}_\Gamma \setminus \{\varpi_3\}$. Due to maximality of height and the binary nature of the slope tree as ensured by Proposition 12.1, ϖ_3 has exactly two descendants in Γ, say w_1 and w_2.

If $\varpi_3 \not\subseteq \varpi_2$, then there is no overlap among the descendants of these two vertices. Thus the two descendants w_3 and w_4 of ϖ_2 must be distinct from w_1, w_2, thus accounting for all the elements of Γ. In this case the conclusion of the lemma holds with $\varpi_1 = D(\varpi_2, \varpi_3)$. If $\varpi_3 \subsetneq \varpi_2$, then again by maximality of height ϖ_2 can contribute exactly one member of Γ that is neither w_1 nor w_2. Let us call this new member w_3. If $\#(\Gamma) = 3$, then the proof is completed by setting $\varpi_1 = D(\varpi_2, \varpi_3) = \varpi_2$. If $\#(\Gamma) = 4$, we call the remaining child w_4, which is not descended from ϖ_2, and set $\varpi_1 = D(\varpi_2, w_4)$. This selection meets (17.8), and also accounts for all the pairwise youngest common ancestors of Γ, as required by part (i) of the lemma.

A very similar argument can be used to prove part (ii). Since the total number of slopes in Ω_N generated by ϖ_3 is exactly $2^{N-\nu(\varpi_3)+1}$, this is the maximum number of possible choices for each of w_{i_3} and w_{j_3}. If $\varpi_3 \not\subseteq \varpi_2$, then $\{i_2, j_2\} \cap \{i_3, j_3\} = \emptyset$. Since each of w_{i_2} and w_{j_2} admits at most $2^{N-\nu(\varpi_2)+1}$ possibilities by the same reasoning, the size of possible four tuples (w_1, w_2, w_3, w_4) in this case is at most 2 raised to the power $2(N-\nu(\varpi_3)+1)+2(N-\nu(\varpi_2)+1)$, which gives the claimed estimate. If $\varpi_3 \subseteq \varpi_2 \subseteq \varpi_1$, then by our assumptions on i_k, j_k, there exist indices $\ell_2 \in \{i_2, j_2\} \setminus \{i_3, j_3\}$ and $\ell_1 \in \{i_1, j_1\} \setminus \{i_3, j_3, \ell_2\}$. Since i_3, j_3, ℓ_1, ℓ_2 are distinct indices and the number of possible choices of w_{i_3}, w_{j_3}, w_{ℓ_1} and w_{ℓ_2} are at most $2^{N-\nu(\varpi_3)}$, $2^{N-\nu(\varpi_3)}$, $2^{N-\nu(\varpi_1)}$ and $2^{N-\nu(\varpi_2)}$ respectively, the result follows. □

Minor modifications of the argument above yield the following analogue for slope triples. The proof is left to the reader.

Lemma 17.5. *(i) Given a collection $\Gamma \subseteq \Omega_N$, $\#(\Gamma) \leq 3$, it is possible to rearrange the collection of vertices $\{D(w, w'); w \neq w',\ w, w' \in \Gamma\}$ as $\{\varpi_1, \varpi_2\}$ with $\varpi_2 \subseteq \varpi_1$.*

(ii) Given a pair $\{\varpi_1, \varpi_2\} \subseteq \mathcal{G}(\Omega_N)$ with $\varpi_2 \subseteq \varpi_1$, define

$$\widehat{m} = \widehat{m}[\varpi_1, \varpi_2] := 2\nu(\varpi_2) + \nu(\varpi_1). \tag{17.10}$$

Let $(i_1, j_1) \neq (i_2, j_2)$ be two pairs of indices such that $\{i_1, j_1, i_2, j_2\} = \{1, 2, 3\}$. Then the following estimate holds:

$$\#\{(w_1, w_2, w_3) : D(w_{i_1}, w_{j_1}) = \varpi_1,\ D(w_{i_2}, w_{j_2}) = \varpi_2\} \leq 2^{3N-\widehat{m}}.$$

17.3 Collections of four tubes with at least two pairwise intersections

17.3.1 Four roots of type 1

We start with the simplest and generic situation, when the root quadruple is of type 1. Motivated by the expression of the probability obtained in (16.24), let us first fix two vertex triples (u, u', z) and (ω, ω', v) in the root tree and slope tree respectively that satisfy the height and containment relations prescribed in Lemma 16.8. For such a selection and with $\varrho \in [M^{-J}, 10A_0]$, we define a collection $\mathcal{E}_{41} = \mathcal{E}_{41}[u, u', z; \omega, \omega', v; \varrho]$ of sticky-admissible tube quadruples of the form $\{(t_1, v_1), (t_2, v_2), (t_1', v_1'), (t_2', v_2')\}$, obeying the following restrictions:

$$\left\{ \begin{aligned} &\mathbb{I} = \{(t_1, t_2); (t_1', t_2')\} \text{ is of type 1}, \, t_1 \neq t_2, \; t_1' \neq t_2', \; u = D(t_1, t_2), \\ &u' = D(t_1', t_2'), \; z = D(u, u'), \; \omega = D(v_1, v_2), \; \omega' = D(v_1', v_2'), \; v = D(\omega, \omega'), \\ &P_{t_1, v_1} \cap P_{t_2, v_2} \cap [\varrho, C_1 \varrho] \times \mathbb{R}^d \neq \emptyset, \; P_{t_1', v_1'} \cap P_{t_2', v_2'} \cap [\varrho, C_1\varrho] \times \mathbb{R}^d \neq \emptyset. \end{aligned} \right\} \tag{17.11}$$

The result below provides a bound on the size of $\mathcal{E}_{41}$.

Lemma 17.6. *There exists a constant $C > 0$ such that*

$$\#\big(\mathcal{E}_{41}\big) \leq C \big(\varrho^2 \rho_\omega \rho_{\omega'}\big)^2 2^{4N - 2\left(\nu(\omega) + \nu(\omega')\right)} M^{-(d-1)\left[h(u) + h(u')\right] + 2(d+1)J}.$$

Proof. Since the intersection and ancestry conditions imply that

$$\mathcal{E}_{41}[u, u', z; \omega, \omega', v] \subseteq \mathcal{E}_2[u, \omega; \varrho] \times \mathcal{E}_2[u', \omega'; \varrho],$$

the stated size bound for $\mathcal{E}_{41}$ is the product of the sizes of the two factors on the right. These are obtained from Lemma 17.3 in Section 17.1, applied twice. □

17.3.2 Four roots of type 2

The treatment of this case follows a similar route, though with certain important variations. The main distinction from Section 17.3.1 is that the intersection and type requirements place greater constraints on the selection of the roots and slopes, and hence on the number of tube quadruples. Thus better bounds are possible, compared to the trivial ones exploited in Lemma 17.3.

Let (u, u', t) and $(\omega, \omega', \vartheta)$ be vertex triples in the root tree and slope tree respectively that meet the requirement of Corollary 16.10. In other words, the vertices u, u', t are linearly ordered in terms of ancestry, and obey $u' \subsetneq u$, while $\omega \cap \vartheta \neq \emptyset$ and $\omega' \cap \vartheta \neq \emptyset$. Holding these fixed, define $\mathcal{E}_{42} = \mathcal{E}_{42}[u, u', t; \omega, \omega', \vartheta; \varrho]$ to be the collection of all sticky-admissible tuples of the form $\{(t_i, v_i), (t_i', v_i') : i = 1, 2\}$ obeying the properties:

$$\left\{\begin{array}{l} \mathbb{I} = \{(t_1, t_2); (t_1', t_2')\} \text{ is of type 2, } u' = D(t_1', t_2') \subsetneq u = D(t_1, t_2), \\ t = D(t_2, t_2'),\ \omega = D(v_1, v_2),\ \omega' = D(v_1', v_2'),\ \vartheta = D(v_2, v_2'), \\ P_{t_1,v_1} \cap P_{t_2,v_2} \cap [\varrho, C_1\varrho] \times \mathbb{R}^d \neq \emptyset,\ P_{t_1',v_1'} \cap P_{t_2',v_2'} \cap [\varrho, C_1\varrho] \times \mathbb{R}^d \neq \emptyset. \end{array}\right\} \quad (17.12)$$

The vertex triple $(\omega, \omega', \vartheta)$ obeys the hypothesis of Lemma 17.4(ii), permitting the application of this lemma in the counting argument presented in Lemma 17.8.

Lemma 17.7. *If the vertex pairs (ω, ϑ) and (ω', ϑ) both have the property that one member of the pair is contained in the other, then there exists a rearrangement of $\{\omega, \omega', \vartheta\}$ as $\{\varpi_1, \varpi_2, \varpi_3\}$ that meets the requirement* (17.8).

Proof. If $\omega \cap \omega' = \emptyset$, then ϑ must contain both ω and ω'. In this case, we rename ϑ as ϖ_1 and call ϖ_3 the element of $\{\omega, \omega'\}$ with greater height. If $\omega \cap \omega' \neq \emptyset$, then the inclusion requirements imply that there must be a ray which contains all three vertices. Since the vertices are linearly ordered, we rename them based on height. □

Lemma 17.7 above allows us to define the quantity m as in (17.9), which by a slight abuse of notation we denote by $m[\omega, \omega', \vartheta]$. We are now in a position to state the main result of this subsection, namely the size estimate for $\mathcal{E}_{42}$. The location of t relative to u, u' affects the size estimate of $\mathcal{E}_{42}$, even though we have seen that the probability estimate in (16.26) remains unchanged with respect to this property.

Lemma 17.8. *The following conclusions hold:*

(i) *If $u' \subseteq t \subseteq u$, then $\mathcal{E}_{42}$ is non-empty only if $dist(t, bdry(u_*)) \leq C\varrho\rho_\omega$. Here u_* is defined to be the unique child of u containing t if $t \subsetneq u$ and is set to be equal to u if $t = u$. In either case,*

$$\#\big(\mathcal{E}_{42}\big) \leq C\big(\varrho^3 \rho_{\omega'}^2 \rho_\omega\big) \min\big[\varrho\rho_\omega, M^{-h(t)}\big] 2^{4N - m[\omega,\omega',\vartheta]} \times M^{-(d-1)\big(h(t)+h(u')\big)+2(d+1)J}.$$

(ii) If $t \subsetneq u' \subsetneq u$, then $\mathcal{E}_{42}$ is non-empty only if

$$\mathit{dist}(t, \mathit{bdry}(u_*)) \le C\varrho\rho_\omega \quad \text{and} \quad \mathit{dist}(t, \mathit{bdry}(u'_*)) \le C\varrho\rho_{\omega'},$$

where u_*, u'_* are the children of u, u' respectively that contain t. In this case,

$$\#\big(\mathcal{E}_{42}\big) \le C\big(\varrho^2\rho_\omega\rho_{\omega'}\big)2^{4N-m[\omega,\omega',\vartheta]}\min\big[\varrho\rho_\omega, M^{-h(t)}\big] \\ \times \min\big[\varrho\rho_{\omega'}, M^{-h(t)}\big]\, M^{-2(d-1)h(t)+2(d+1)J}.$$

Proof. Both statements in the lemma involve similar arguments. We only prove part (i) in detail, and leave a brief sketch for the other part. The argument here follows the basic structure of Lemma 17.3, since we still have the trivial containment

$$\mathcal{E}_{42}[u, u', t; \omega, \omega', \vartheta; \varrho] \subseteq \mathcal{E}_2[u, \omega; \varrho] \times \mathcal{E}_2[u', \omega'; \varrho], \tag{17.13}$$

but with a few modifications resulting from the more refined information about the roots and slopes available from t and ϑ. For instance, combining the defining assumptions that $t_2 \subseteq t$ and $u = D(t_1, t_2)$ with the intersection inequality $|\text{cen}(t_2) - \text{cen}(t_1)| \le 2C_1\varrho\rho_\omega$ derived from (17.3) in Lemma 17.1, we deduce that t has to lie within distance $2C_1\varrho\rho_\omega$ of the boundary of u_*. This is the first conclusion of part (i). For the size bound, we reason as follows. By Lemma 17.1(ii), the number of t_1 and t'_1, if everything else is held fixed, is $\le C(\varrho\rho_\omega M^J)(\varrho\rho_{\omega'}M^J) \le C\varrho^2\rho_\omega\rho_{\omega'}M^{2J}$. Turning to slope counts, we apply Lemma 17.4(ii), the use of which has already been justified in Lemma 17.7, to deduce that the number of possible slope quadruples (v_1, v_2, v'_1, v'_2) is 2^{4N-m}. It remains to compute the size of the t_2 and t'_2 projections of $\mathcal{E}_{42}$. In view of (17.13), a bound on the size of the t'_2 projection is given by the right hand side of (17.5) with u replaced by u'. On the other hand, t_2 is restricted to lie within t and within distance $2C_1\varrho\rho_\omega$ from the boundary of t if $t \subsetneq u$. This places a nontrivial spatial restriction on t_2 only if $2C_1\varrho\rho_\omega < M^{-h(t)}$. If $t = u$, the argument leading up to (17.5) shows that t_2 lies in $\mathcal{A}_u$ defined in (17.7). In either event the volume of the region where t_2 can range is at most $C\min(\varrho\rho_\omega, M^{-h(t)})M^{-(d-1)h(t)}$, hence the cardinality of the t_2 projection is at most M^{dJ} times this quantity (see

Figure 17.3). Combining all these counts yields the bound on the size of $\mathcal{E}_{42}$ given in part (i).

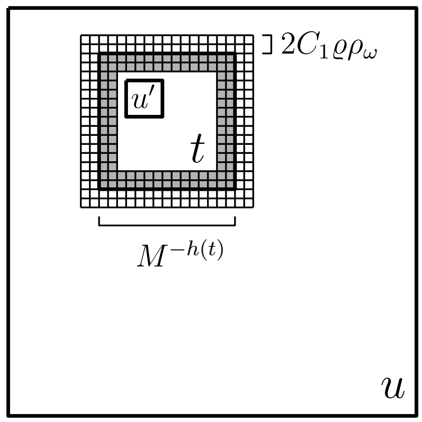

Figure 17.3: Illustration of the spatial restriction on t_2 imposed by the conditions $u' \subset t \subset u$, $t_2 \subset t$, $\text{dist}(t_1, t_2) \le 2C_1\varrho\rho_\omega < M^{-h(t)}$. Here, t_2 must lie within the shaded region along the boundary of t, with t_1 falling just outside this boundary in the unshaded thatched region.

For part (ii), the size estimate of $\mathcal{E}_{42}$ is a product of a number of factors analogous to the ones already considered, the origins of which are indicated below.

$$\left.\begin{aligned}&\#(t_1 \text{ given } v_1, v_2, t_2) \le C\varrho\rho_\omega M^J,\\ &\#(t_1' \text{ given } v_1', v_2', t_2') \le C\varrho\rho_{\omega'} M^J,\end{aligned}\right\} \quad \text{(Lemma 17.1(ii))}$$

$$\left.\begin{aligned}&\#(t_2) \le C\min\big[\varrho\rho_\omega, M^{-h(t)}\big]M^{-(d-1)h(t)+dJ},\\ &\#(t_2') \le C\min\big[\varrho\rho_{\omega'}, M^{-h(t)}\big]M^{-(d-1)h(t)+dJ},\end{aligned}\right\} \quad \text{(arguments similar to part (i))},$$

$$\#(v_1, v_2, v_1', v_2') \le 2^{4N-m[\omega,\omega',\vartheta]} \quad \text{(from Lemma 17.4(ii))}.$$

We omit the details. □

17.3.3 Four roots of type 3

To complete the discussion of size for collections consisting of intersecting tube quadruples, it remains to consider the case where the root configuration is of type 3.

Motivated by the conclusions of Lemma 16.11 and Corollary 16.12, we fix two vertex tuples (u, s_1, s_2) and $(\omega, \omega', \vartheta_1, \vartheta_2)$ in the root tree and the slope tree respectively, with the properties that $s_1, s_2 \subseteq u$, $h(u) \leq h(s_1) \leq h(s_2)$, $\omega \cap \vartheta_i \neq \emptyset$, and $\omega' \cap \vartheta_i \neq \emptyset$ for $i = 1, 2$. For such a selection, we define $\mathcal{E}_{43}[u, s_1, s_2; \omega, \omega', \vartheta_1, \vartheta_2; \varrho]$ to be the collection of all sticky-admissible tuples $\{(t_i, v_i), (t'_i, v'_i) : i = 1, 2\}$ that satisfy the list of conditions below:

$$\left\{ \begin{array}{l} \mathbb{I} = \{(t_1, t_2); (t'_1, t'_2)\} \text{ is of type 3, } u = D(t_1, t_2) = D(t'_1, t'_2), \\ \omega' = D(v'_1, v'_2) \subseteq \omega = D(v_1, v_2), \; s_i = D(t_i, t'_i), \; \vartheta_i = D(v_i, v'_i), \; i = 1, 2, \\ P_{t_1, v_1} \cap P_{t_2, v_2} \cap [\varrho, C_1 \varrho] \times \mathbb{R}^d \neq \emptyset, \; P_{t'_1, v'_1} \cap P_{t'_2, v'_2} \cap [\varrho, C_1 \varrho] \times \mathbb{R}^d \neq \emptyset. \end{array} \right\} \quad (17.14)$$

Since $\mathbb{I}$ is of type 3, interchanging (t_1, t_2) and (t'_1, t'_2) leaves u unchanged. Hence we may assume without loss of generality that $\rho_\omega \leq \rho_{\omega'}$. Further, Lemma 17.4(i) dictates that for $\mathcal{E}_{43}$ to be non-empty, at most three out of the four vertices $\omega, \omega', \vartheta_1, \vartheta_2$ can be distinct. We leave the reader to verify that Lemma 17.7 can be applied to any triple of these four vertices. Thus for any choice of an eligible tuple $\{\omega, \omega', \vartheta_1, \vartheta_2\}$, there exists a rearrangement of its entries as $\{\varpi_1, \varpi_2, \varpi_3\}$ obeying the hypothesis and hence the conclusion of Lemma 17.4(ii). This permits a consistent definition of the quantity $m[\omega, \omega', \vartheta_1, \vartheta_2]$ as in (17.9), which we use in the statement of the lemma below. (Note that $\#(\{\omega, \omega', \vartheta_1, \vartheta_2\}) \leq 3$.)

Lemma 17.9. *If $s_i \subsetneq u$, let u_i denote the child of u that contains s_i. Set $\Delta := \min[\varrho \rho_\omega, \varrho \rho_{\omega'}]$.*

(i) The collection $\mathcal{E}_{43}$ is nonempty only if

$$\sum_{i=1}^{2} dist\big(s_i, bdry(u_i)\big) \leq C\Delta, \quad (17.15)$$

where $dist(s_1, bdry(u_1))$ is defined to be zero if $u = s_1$.

(ii) If $\Delta \leq M^{-h(s_1)}$ and $\mathcal{E}_{43}$ is nonempty, then in addition to (17.15)*, one of the following two conditions must hold:*

1. $s_2 \subsetneq s_1 = u$, *in which case* s_2 *lies within distance* $C\Delta$ *of the boundary of some child of* $s_1 = u$.
2. $s_2 \cap s_1 = \emptyset$, *in which case* $dist(s_2, bdry(s_1)) \le C\Delta$.

In either case, s_2 *is constrained to lie in the union of at most* $2^d M$ *slab-like parallepipeds, each with* $(d-1)$ *"long" directions of sidelength* $M^{-h(s_1)}$ *and one "short" direction of sidelength* Δ.

(iii) *If* $\Delta \ge M^{-h(s_1)}$ *and* $\mathcal{E}_{43}$ *is nonempty, then in addition to* (17.15), s_2 *has to lie within a thin tube-like parallelepiped of length* $\varrho \min(M^{-h(\omega)}, M^{-h(\omega')})$ *in one "long" direction and thickness* $CM^{-h(s_1)}$ *in the remaining* $(d-1)$ *"short" directions; more precisely, both the following inequalities must hold:*

$$|cen(s_2) - cen(s_1) + x_1(cen(\omega \cap \vartheta_2) - cen(\omega \cap \vartheta_1))| \le CM^{-h(s_1)}, \text{ and} \tag{17.16}$$

$$|cen(s_2) - cen(s_1) + x_1'(cen(\omega' \cap \vartheta_2) - cen(\omega' \cap \vartheta_1))| \le CM^{-h(s_1)} \tag{17.17}$$

for some $x_1, x_1' \in [\varrho, C_1\varrho]$. *Here* $cen(t)$ *denotes the centre of the cube* t.

(iv) *In all cases, if* $\mathcal{E}_{43}$ *is nonempty,*

$$\#(\mathcal{E}_{43}) \le C2^{4N - m[\omega,\omega',\vartheta_1,\vartheta_2]} M^{-2(d-1)h(s_2)+2(d+1)J} \times \prod_{i=1}^{2} \Big[\min\big[\varrho\rho_\omega, M^{-h(s_i)}\big] \min\big[\varrho\rho_{\omega'}, M^{-h(s_i)}\big]\Big].$$

Proof. Let us fix a tuple $\{(t_i, v_i); (t_i', v_i') : i = 1, 2\}$ in $\mathcal{E}_{43}$. As in previous proofs such as Lemma 17.2 (applications of which have appeared in the counting arguments of Lemma 17.3 and 17.8), the key elements are the inequalities

$$|\text{cen}(t_2) - \text{cen}(t_1)| \le C\varrho\rho_\omega \quad \text{and} \quad |\text{cen}(t_2') - \text{cen}(t_1')| \le C\varrho\rho_{\omega'}. \tag{17.18}$$

They are proved exactly in the same way as (17.6) follows from (17.3), resulting from the nontrivial intersection conditions that define $\mathcal{E}_{43}$. Combined with the set inclusion relations $u = D(t_1, t_2) = D(t_1', t_2')$ and $t_1, t_1' \subseteq s_1$ and $t_2, t_2' \subseteq s_2$ that are guaranteed by the type assumption on the roots, this yields that

$$\begin{aligned}\operatorname{dist}(s_i, \operatorname{bdry}(u_i)) &\leq \inf\bigl[\operatorname{dist}(t_i, \operatorname{bdry}(u_i)), \operatorname{dist}(t_i', \operatorname{bdry}(u_i))\bigr] \\ &= \inf\bigl[\operatorname{dist}(t_i, u_i^c), \operatorname{dist}(t_i', u_i^c)\bigr] \\ &\leq \inf\bigl[\operatorname{dist}(t_1, t_2), \operatorname{dist}(t_1', t_2')\bigr] \\ &\leq C \min[\varrho\rho_\omega, \varrho\rho_{\omega'}] = C\Delta,\end{aligned}$$

leading to the distance constraints in (17.15). Incidentally, the inequalities (17.18) also prove the relation in part (ii) if $s_2 \cap s_1 = \emptyset$. On the other hand, if $s_2 \subseteq s_1$, then $s_1 = u$ and the desired inequality is simply a restatement of the one in (17.15). For part (iii), we refer again to the intersection inequality (17.3), using it to deduce that

$$\begin{aligned}&\Bigl|\operatorname{cen}(s_2) - \operatorname{cen}(s_1) + x_1\bigl(\operatorname{cen}(\omega \cap \vartheta_2) - \operatorname{cen}(\omega \cap \vartheta_1)\bigr)\Bigr| \\ &\quad \leq \sum_{i=1}^{2} |\operatorname{cen}(s_i) - \operatorname{cen}(t_i)| + |x_1| \sum_{i=1}^{2} \bigl|\operatorname{cen}(\omega \cap \vartheta_i) - v_i\bigr| + |\operatorname{cen}(t_2) - \operatorname{cen}(t_1) + x_1(v_2 - v_1)| \\ &\quad \leq \sqrt{d} \sum_{i=1}^{2} M^{-h(s_i)} + C_1 \varrho \sqrt{d} \sum_{i=1}^{2} M^{-h(\vartheta_i)} + 2c_d\sqrt{d} M^{-J} \leq C M^{-h(s_1)}.\end{aligned}$$

Here we have also used the height and inclusion relations associated with the root configuration type established in Lemma 16.11; namely,

$$t_i \subseteq s_i, \quad v_i \in \omega \cap \vartheta_i, \quad h(s_i) \leq h(\vartheta_i), \quad h(s_1) \leq h(s_2), \quad i = 1, 2.$$

The inequality above implies that s_2 has to lie within distance $O(M^{-h(s_1)})$ of a line segment of length at most $\varrho|\operatorname{cen}(\omega \cap \vartheta_2) - \operatorname{cen}(\omega \cap \vartheta_1)| \leq \varrho M^{-h(\omega)}$. The inequality (17.17) is proved in an identical manner, using t_i', v_i', ω' instead of t_i, v_i, ω. The first statement in part (iii) is a consequence of both these inequalities.

The bound on the size of $\mathcal{E}_{43}$ uses the same machinery as in the proof of Lemma 17.8, so we simply indicate the breakdown of the contributions from the different sources:

$$\#(\mathcal{E}_{43}) \le \underbrace{C\min\big[\varrho\rho_\omega, M^{-h(s_1)}\big]M^J}_{\#(t_1) \text{ with } t_2, v_1, v_2 \text{ fixed}} \times \underbrace{C\min\big[\varrho\rho_{\omega'}, M^{-h(s_1)}\big]M^J}_{\#(t_1') \text{ with } t_2', v_1', v_2' \text{ fixed}} \times \underbrace{2^{4N-m[\omega,\omega',\vartheta_1,\vartheta_2]}}_{\substack{\#(v_1,v_2,v_1',v_2') \\ \text{from Lemma 17.4(ii)}}}$$
$$\times \underbrace{C\min\big[\varrho\rho_\omega, M^{-h(s_2)}\big]M^{-(d-1)h(s_2)+dJ}}_{\#(t_2\text{-projection})} \times \underbrace{\min\big[\varrho\rho_{\omega'}, M^{-h(s_2)}\big]M^{-(d-1)h(s_2)+dJ}}_{\#(t_2'\text{-projection})},$$

which leads to the stated estimate. □

17.4 Collections of three tubes with at least two pairwise intersections

For the sake of completeness and book-keeping, we record in this section the cardinality of collections consisting of intersecting tube triples. No new ideas are involved in the proofs, which are in fact simpler than the ones in Section 17.3. These are left to the interested reader.

Using the notation set up in Lemmas 16.5 and 16.6, we define the collections $\mathcal{E}_{31} = \mathcal{E}_{31}[u, u'; \omega, \omega'; \varrho]$ and $\mathcal{E}_{32} = \mathcal{E}_{32}[u, t; \omega, \omega', \vartheta; \varrho]$ in exactly the same way $\mathcal{E}_{4i}$ were defined. Namely, $\mathcal{E}_{3i}$ consists of all sticky-admissible tuples of the form $\{(t_1, v_1), (t_2, v_2), (t_2', v_2')\}$ such that $\mathbb{I} = \{t_1, t_2, t_2'\}$ is of type i and

$$P_{t_1,v_1} \cap P_{t_2,v_2} \cap [\varrho, C_1\varrho] \times \mathbb{R}^d \neq \emptyset, \quad P_{t_1,v_1} \cap P_{t_2',v_2'} \cap [\varrho, C_1\varrho] \times \mathbb{R}^d \neq \emptyset.$$

In addition, the members of $\mathcal{E}_{3i}$ must satisfy

$$u = D(t_1, t_2), \quad u' = D(t_1, t_2'), \quad \omega = D(v_1, v_2), \quad \omega' = D(v_1, v_2'),$$

with $u = u'$, $t = D(t_2, t_2')$ and $\vartheta = D(v_2, v_2')$ if $i = 2$. We also define the quantities $\widehat{m}[\omega, \omega']$ for $\mathcal{E}_{31}$ and $\widehat{m}[\omega, \omega', \vartheta]$ for $\mathcal{E}_{32}$; both are expressed using the formula (17.10), where $\{\varpi_1, \varpi_2\}$ with $\varpi_2 \subseteq \varpi_1$ is a rearrangement of $\{\omega, \omega'\}$ for $\mathcal{E}_{31}$ and of $\{\omega, \omega', \vartheta\}$ for $\mathcal{E}_{32}$, by virtue of Lemma 17.5. With this notation in place, the size estimates on

$\mathcal{E}_{3i}$ are as follows.

Lemma 17.10. *(i) Set* $\Delta := \min(\varrho\rho_\omega, \varrho\rho_{\omega'})$. *Then*

$$\#(\mathcal{E}_{31}) \leq C\Delta\varrho^2\rho_\omega\rho_{\omega'}2^{3N-\widehat{m}[\omega,\omega']}M^{-(d-1)(h(u)+h(u'))+(2d+1)J}.$$

(ii) With the same definition of Δ *as in part (i),*

$$\#(\mathcal{E}_{32}) \leq C\Delta\min[\varrho\rho_\omega, M^{-h(t)}]\min[\varrho\rho_{\omega'}, M^{-h(t)}]2^{3N-\widehat{m}[\omega,\omega',\vartheta]}M^{-2(d-1)h(t)+(2d+1)J}.$$

Chapter 18

Sums over root and slope vertices

A recurrent feature of the proof of (14.9), as we will soon see in Chapter 19, is the use of certain sums over specific subsets of vertices in the root and slope trees. We record the outcomes of these summation procedures in this section for easy reference later.

Lemma 18.1. *Fix a vertex $\varpi_0 \in \mathcal{G}(\Omega_N)$, i.e. ϖ_0 is a splitting vertex of the slope tree. Then the following estimates hold.*

(i) For any $\alpha \in \mathbb{R}$,

$$\sum_{\substack{\varpi \in \mathcal{G}(\Omega_N) \\ \varpi \subseteq \varpi_0}} 2^{-\alpha\nu(\varpi)} \leq \begin{cases} C_\alpha 2^{-\alpha\nu(\varpi_0)} & \text{if } \alpha > 1, \\ N 2^{-\nu(\varpi_0)} & \text{if } \alpha = 1, \\ C_\alpha 2^{-\alpha\nu(\varpi_0) + N(1-\alpha)} & \text{if } \alpha < 1. \end{cases}$$

(ii) For $M \geq 2$, $\beta > 0$ and $\alpha \geq 1$,

$$\sum_{\substack{\varpi \in \mathcal{G}(\Omega_N) \\ \varpi \subseteq \varpi_0}} M^{-\beta\lambda(\varpi)} 2^{-\alpha\nu(\varpi)} \leq C_{\alpha,\beta} M^{-\beta\lambda(\varpi_0)} 2^{-\alpha\nu(\varpi_0)}.$$

Proof. By Proposition 12.5, the number of splitting vertices descended from ϖ_0 with the property that $\nu(\varpi) = \nu(\varpi_0) + j$ is 2^j. Since j can be at most N, we see that

$$\sum_{\substack{\varpi \in \mathcal{G}(\Omega_N) \\ \varpi \subseteq \varpi_0}} 2^{-\alpha\nu(\varpi)} \leq \sum_j 2^{-\alpha(\nu(\varpi_0)+j)} 2^j \leq 2^{-\alpha\nu(\varpi_0)} \sum_{j=1}^{N} 2^{j(1-\alpha)},$$

from which part (i) follows. On the other hand, if $\nu(\varpi) = \nu(\varpi_0) + j$, then $\lambda(\varpi) - \lambda(\varpi_0) \geq \nu(\varpi) - \nu(\varpi_0) = j$. Thus, a similar computation shows that

$$\begin{aligned} \sum_{\substack{\varpi \in \mathcal{G}(\Omega_N) \\ \varpi \subseteq \varpi_0}} M^{-\beta\lambda(\varpi)} 2^{-\alpha\nu(\varpi)} &= \sum_j M^{-\beta\left(\lambda(\varpi_0)+j\right)} 2^{-\alpha\left(\nu(\varpi_0)+j\right)} 2^j \\ &\leq M^{-\beta\lambda(\varpi_0)} 2^{-\alpha\nu(\varpi_0)} \sum_{j=1}^{\infty} M^{-\beta j} 2^{-(\alpha-1)j}. \end{aligned}$$

The last sum in the displayed expression is convergent, establishing the desired conclusion in part (ii). □

Lemma 18.2. *Fix a vertex y in the root tree and a splitting vertex ϖ in the slope tree such that $h(y) \leq h(\varpi)$. Given a constant β, one of the following estimates holds for*

$$\mathfrak{s}(\beta) := \sum_z{}' M^{-\beta h(z)} 2^{\mu(\varpi, h(z))},$$

where the sum $\sum'$ takes place over all vertices z of the root tree such that $z \subseteq y$ and $h(z) \leq \lambda(\varpi)$.

(i) If $\beta < d$, then $\mathfrak{s}(\beta) \leq C_\beta 2^{\nu(\varpi)} M^{(d-\beta)\lambda(\varpi) - dh(y)}$.

(ii) If $\beta = d$, then $\mathfrak{s}(d) \leq C 2^{\nu(\varpi)} h(\varpi) M^{-dh(y)}$.

(iii) If $\beta > d$, then $\mathfrak{s}(\beta) \leq C_\beta 2^{\nu(\varpi)} M^{-\beta h(y)}$.

(iv) If $\beta > d$ is large enough so that $2M^d < M^\beta$, then $\mathfrak{s}(\beta) \leq C_\beta M^{-dh(y)}$.

Proof. Since ϖ is a splitting vertex of the slope tree, there exists an integer $j \in [1, N]$ such that $\varpi \in \mathcal{G}_j(\Omega_N)$, i.e., $\nu(\varpi) = j$. By definition, every jth splitting vertex is either itself a $(j-1)$th basic slope cube or is contained in one. Let $\varpi_\ell \in \mathcal{H}_\ell(\Omega_N)$ be the ℓth slope cube that contains ϖ, so that

$$\varpi_1 \supsetneq \varpi_2 \supsetneq \cdots \supsetneq \varpi_{j-1} \supseteq \varpi.$$

If z is a vertex of the root tree such that $h(\varpi_{\ell-1}) \leq h(z) < h(\varpi_\ell)$ for some $\ell \leq j-1$, then $\mu(\varpi, h(z)) = \ell - 1$; on the other hand, if $h(\varpi_{j-1}) \leq h(z) \leq \lambda(\varpi)$, then $\mu(\varpi, h(z)) = j-1$. This suggests decomposing the sum defining $\mathfrak{s}(\beta)$ according to the heights of the slope cubes containing ϖ. Implementing this and recalling that $\#\{z : z \subseteq y,\ h(z) = k\} = M^{dk - dh(y)}$, we obtain

$$\begin{aligned}
\mathfrak{s}(\beta) &= \sum_{\ell=1}^{j-1} \sum_{k=h(\varpi_{\ell-1})}^{h(\varpi_\ell)-1} 2^{\ell-1} \sideset{}{'}\sum_{z:h(z)=k} M^{-\beta k} + \sum_{k=h(\varpi_{j-1})}^{\lambda(\varpi)} 2^{j-1} \sideset{}{'}\sum_{z:h(z)=k} M^{-\beta k} \\
&\leq C \Big[\sum_{\ell=1}^{j} \sum_{k=h(y)}^{\lambda(\varpi)} 2^{\ell-1} M^{-\beta k} M^{dk-dh(y)} \Big] \\
&\leq C M^{-dh(y)} \sum_{\ell=1}^{j} 2^{\ell-1} \sum_{k=h(y)}^{\lambda(\varpi)} M^{(d-\beta)k} \\
&\leq C M^{-dh(y)} 2^j \sum_{k=h(y)}^{\lambda(\varpi)} M^{(d-\beta)k} \\
&\leq C 2^{\nu(\varpi)} M^{-dh(y)} \begin{cases} M^{(d-\beta)\lambda(\varpi)} & \text{if } \beta < d, \\ \lambda(\varpi) & \text{if } \beta = d, \\ M^{-(\beta-d)h(y)} & \text{if } \beta > d. \end{cases}
\end{aligned}$$

Upon simplification, these are the estimates claimed in parts (i)-(iii) of the lemma. Part (iv) follows from the observation that $\mu(\varpi, h(z)) \leq h(z)$, hence

$$\mathfrak{s}(\beta) \leq \sideset{}{'}\sum_{z} M^{-\beta h(z)} 2^{h(z)} \leq \sum_k 2^k M^{-\beta k + d(k-h(y))}$$

$$\le M^{-dh(y)} \sum_k \left(\frac{2M^d}{M^\beta}\right)^k \le C_\beta M^{-dh(y)}.$$

□

In view of spatial constraints on the ancestors of root cubes as encountered in Lemmas 17.8 and 17.9, occasionally the sums that we consider take place over more restricted ranges of vertices than the one in Lemma 18.2, even though the summands may retain the same form. The next result makes this quantitatively precise. Let ϖ be a splitting vertex of the slope tree, and $\mathcal{R}$ a fixed parallelepiped in the root hyperplane with sidelength β in $(d-r)$ directions and γ in the remaining r directions, where $1 \le r \le d-1$ and $\beta \ge \gamma \ge M^{-J}$. Given constants $\epsilon \ge M^{-\lambda(\varpi)}$ and $\alpha \in \mathbb{R}$, we define

$$\mathfrak{s}_\pm = \mathfrak{s}_\pm(\alpha, \epsilon, \mathcal{R}, \varpi) := \sum_{z \in \mathcal{Z}_\pm} M^{-\alpha h(z)} 2^{\mu(\omega, h(z))}, \tag{18.1}$$

where the index sets $\mathcal{Z}_\pm$ are collections of vertices of the root tree defined as follows:

$$\begin{aligned} \mathcal{Z} &:= \{z \subseteq \mathcal{R} : h(z) \le \lambda(\varpi),\ M^{-h(z)} \le \epsilon\}, \\ \mathcal{Z}_+ &:= \mathcal{Z} \cap \{z : M^{-h(z)} \ge \gamma\}, \\ \mathcal{Z}_- &:= \mathcal{Z} \cap \{z : M^{-h(z)} \le \gamma\}. \end{aligned}$$

Lemma 18.3. *The following estimates hold for $\mathfrak{s}_\pm$ defined in* (18.1).

(i) If $\alpha > d-r$ and $\epsilon \ge \gamma$ then $\mathfrak{s}_+ \le C 2^{\nu(\varpi)} \beta^{d-r} \epsilon^{\alpha-d+r}$.

(ii) If $\alpha > d$, then $\mathfrak{s}_- \le C 2^{\nu(\varpi)} \beta^{d-r} \gamma^r \min(\epsilon, \gamma)^{\alpha-d}$.

Proof. We have already established in the proof of Lemma 18.2 that $\mu(\varpi, h(z)) \le \nu(\varpi) - 1$. Further if $M^{-k} \ge \gamma$, then there can be at most a constant number of possible choices of kth generation M-adic cubes z that are contained in $\mathcal{R}$ and intersect with a slice of $\mathcal{R}$ that fixes coordinates in the $(d-r)$ long directions. Thus we only need to count the number of possible z in the long directions, obtaining

$$\#\{z \in \mathcal{Q}(k) : z \subseteq \mathcal{R}\} \le C_r \beta^{d-r} M^{(d-r)k}. \tag{18.2}$$

Taking this into account, we obtain

$$\mathfrak{s}_+ \leq 2^{\nu(\varpi)} \sum_{k:\gamma \leq M^{-k} \leq \epsilon} M^{-\alpha k} \beta^{d-r} M^{(d-r)k} \leq C 2^{\nu(\varpi)} \beta^{d-r} \epsilon^{\alpha-d+r},$$

as claimed in part (i). Part (ii) follows in an identical manner; the only difference is that now all directions of $\mathcal{R}$ are thick relative to the scale of z, hence (18.2) has to be replaced by

$$\#\{z \in \mathcal{Q}(k) : z \subseteq \mathcal{R}\} \leq C \gamma^r \beta^{d-r} M^{dk}.$$

□

Chapter 19

Proof of the lower bound (14.9)

We are now in a position to complete the proof of Proposition 14.2 by verifying the probabilistic statement on the lower bound of $K_N(\mathbb{X})$ claimed in (14.9). The two propositions stated below are the main results of this chapter and allow passage to this final step.

Proposition 19.1. *Fix integers N and R with $N \gg M$ and $10 \leq R \leq \frac{1}{10}\log_M N$. Define*

$$P^*_{t,\sigma,R} := P_{t,\sigma(t)} \cap [M^{-R}, M^{-R+1}] \times \mathbb{R}^d,$$

where $\sigma = \sigma_{\mathbb{X}}$ is the randomized weakly sticky map described in Section 14. Then there exists a constant $C = C(M,d) > 0$ such that

$$\mathbb{E}_{\mathbb{X}}\Big[\sum_{t_1 \neq t_2} \big|P^*_{t_1,\sigma,R} \cap P^*_{t_2,\sigma,R}\big|\Big] \leq CNM^{-2R}. \tag{19.1}$$

Proposition 19.2. *Under the same hypotheses as Proposition 19.1,*

$$\mathbb{E}_{\mathbb{X}}\Big[\Big(\sum_{t_1 \neq t_2} \big|P^*_{t_1,\sigma,R} \cap P^*_{t_2,\sigma,R}\big|\Big)^2\Big] \leq C\big(NM^{-2R}\big)^2. \tag{19.2}$$

Propositions 19.1 and 19.2 should be viewed as the direct generalizations of Propositions 10.1 and 10.2 for arbitrary direction sets. These are proved below in Sections 19.1 and 19.2 respectively. Of the two results, Proposition 19.2 is of direct interest, since it leads to (14.9), as we will see momentarily in Corollary 19.3. Proposition 19.1, while not strictly speaking relevant to (14.9), nevertheless provides a context for presenting the core arguments within a simpler framework.

Corollary 19.3. *Proposition 19.2 implies* (14.9).

Proof. The argument here is identical to Corollary 10.3, and is omitted. □

19.1 Proof of Proposition 19.1

Proof. We first recast the sum on the left hand side of (19.1) in a form that brings into focus its connections with the material in Chapters 16 and 17. By the Córdoba estimate, inequality (1.5),

$$\sum_{t_1 \neq t_2} \left| P^*_{t_1,\sigma,R} \cap P^*_{t_2,\sigma,R} \right| \leq \sum_1 \frac{C_d M^{-(d+1)J}}{|\sigma(t_1) - \sigma(t_2)| + M^{-J}}, \tag{19.3}$$

where $\sum_1$ denotes the sum over all root pairs (t_1, t_2) such that $t_1 \neq t_2$ and $P^*_{t_1,\sigma,R} \cap P^*_{t_2,\sigma,R} \neq \emptyset$. Unravelling the implications of the intersection we find that

$$\begin{aligned} &\left\{(t_1,t_2) : t_1 \neq t_2,\, P^*_{t_1,\sigma,R} \cap P^*_{t_2,\sigma,R} \neq \emptyset\right\} \\ &\qquad \subseteq \left\{ (t_1,t_2) \,\middle|\, \begin{array}{c} \exists \text{ a unique pair } (v_1,v_2) \in \Omega_N^2 \text{ such that} \\ P_{t_1,v_1} \cap P_{t_2,v_2} \cap [M^{-R}, M^{-R+1}] \times \mathbb{R}^d \neq \emptyset, \\ \text{and } \sigma(t_1) = v_1,\ \sigma(t_2) = v_2,\ t_1 \neq t_2 \end{array} \right\}. \end{aligned} \tag{19.4}$$

For a given root pair (t_1, t_2), there may exist more than one slope pair (v_1, v_2) that meets the intersection criterion in (19.4). But only one pair will also satisfy, for a given σ, the requirement $\sigma(t_1) = v_1$, $\sigma(t_2) = v_2$, which explains the uniqueness claim in (19.4). Using this, the expression on the right hand side of (19.3) can be expanded as follows,

$$\sum_{t_1 \neq t_2} \left| P^*_{t_1,\sigma,R} \cap P^*_{t_2,\sigma,R} \right| \leq \sum_1 \frac{C_d M^{-(d+1)J}}{|\sigma(t_1) - \sigma(t_2)|}$$

$$\leq \sum_{2} \frac{C_d M^{-(d+1)J}}{|v_1 - v_2|} T((t_1, v_1), (t_2, v_2))$$
$$\leq \sum_{u,\omega} \frac{C_d M^{-(d+1)J}}{\delta_\omega} \sum_{3} T((t_1, v_1), (t_2, v_2)), \tag{19.5}$$

where the notation $\sum_2$ in the second step denotes summation over the collection in (19.4), and $T((t_1, v_1), (t_2, v_2))$ is a binary (random) counter given by

$$T((t_1, v_1), (t_2, v_2)) = \begin{cases} 1 & \text{if } \sigma(t_1) = v_1 \text{ and } \sigma(t_2) = v_2, \\ 0 & \text{otherwise.} \end{cases} \tag{19.6}$$

In the last step (19.5) of the string of inequalities above, we have rearranged the sum in terms of the youngest common ancestors $u = D(t_1, t_2)$ and $\omega = D(v_1, v_2)$ in the root tree and in the slope tree respectively. The summation $\sum_3$ takes place over *all* sticky-admissible tube pairs $\{(t_1, v_1), (t_2, v_2)\}$ in the deterministic collection $\mathcal{E}_2[u, \omega; \varrho]$ defined in (17.1), with $\varrho = \varrho_R = M^{-R}$ and $C_1 = M$. Incidentally, the requirement of sticky-admissibility restricts u and ω to obey the height relation $h(u) < \lambda(\omega)$. The quantity δ_ω has been defined in (12.8), and is therefore $\leq |v_1 - v_2|$.

With this preliminary simplification out of the way, we proceed to compute the expected value of the expression in (19.5), combining the geometric facts and counting arguments from Chapter 17 with appropriate probability estimates from Chapter 16. Accordingly, we get

$$\begin{aligned}
\mathbb{E}_{\mathbb{X}}\Big[\sum_{t_1 \neq t_2} \big|P^*_{t_1,\sigma,R} \cap P^*_{t_2,\sigma,R}\big|\Big] &\leq \sum_{u,\omega} \frac{C_d M^{-(d+1)J}}{\delta_\omega} \sum_{\mathcal{E}_2[u,\omega;\varrho]} \mathbb{E}_{\mathbb{X}}\big[T((t_1, v_1), (t_2, v_2))\big] \\
&\leq \sum_{u,\omega} \frac{C_d M^{-(d+1)J}}{\underbrace{C_2 \rho_\omega}_{\text{Corollary 12.4}}} \sum_{\mathcal{E}_2[u,\omega;\varrho]} \Pr\big(\sigma(t_1) = v_1,\ \sigma(t_2) = v_2\big) \\
&\leq C M^{-(d+1)J} \sum_{u,\omega} \rho_\omega^{-1} \#(\mathcal{E}_2[u, \omega; \varrho]) \underbrace{\left(\frac{1}{2}\right)^{2N - \mu(\omega, h(u))}}_{\text{Lemma 16.3}}
\end{aligned}$$

$$\leq CM^{-(d+1)J} \sum_{u,\omega} \rho_\omega^{-1} \underbrace{(\varrho\rho_\omega)^2 2^{2(N-\nu(\omega))} M^{-(d-1)h(u)+(d+1)J}}_{\text{Lemma 17.3}} \times \left(\frac{1}{2}\right)^{2N-\mu(\omega,h(u))}$$

$$\leq CM^{-2R} \sum_{u,\omega} M^{-\lambda(\omega)-(d-1)h(u)} 2^{\mu(\omega,h(u))-2\nu(\omega)},$$

where the last step uses the fact that $\rho_\omega \leq C_1 M^{-\lambda(\omega)}$. To establish the conclusion claimed in (19.1), it remains to show that the last expression in the displayed steps above is bounded by CN. This follows from a judicious use of the summation results proved in Chapter 18; namely,

$$\begin{aligned}
\sum_{u,\omega} M^{-\lambda(\omega)-(d-1)h(u)} 2^{\mu(\omega,h(u))-2\nu(\omega)} &= \sum_{\omega \in \mathcal{G}} M^{-\lambda(\omega)} 2^{-2\nu(\omega)} \sum_u M^{-(d-1)h(u)} 2^{\mu(\omega,h(u))} \\
&\leq C \sum_{\omega \in \mathcal{G}} M^{-\lambda(\omega)} 2^{-2\nu(\omega)} \left[2^{\nu(\omega)} M^{\lambda(\omega)} \right] \\
&\leq C \sum_{\omega \in \mathcal{G}} 2^{-\nu(\omega)} \leq CN,
\end{aligned}$$

where the second and last steps are consequences, respectively, of Lemma 18.2(i) with $\varpi = \omega$ and $\beta = d-1$ and of Lemma 18.1(i) with $\alpha = 1$. In both applications, y and ϖ_0 have been chosen to be the unit cube, in the root tree and the slope tree respectively. □

19.2 Proof of Proposition 19.2

We are now ready to prove the main Proposition 19.2.

Proof. As in the proof of Proposition 19.1, an initial processing of the sum on the left hand side of (19.2) is needed before embarking on the evaluation of the expectation.

Accordingly, we decompose and simplify the quantity of interest as follows,

$$\begin{aligned}\Big[\sum_{t_1\neq t_2}\big|P^*_{t_1,\sigma,R}\cap P^*_{t_2,\sigma,R}\big|\Big]^2 &= \sum_{\substack{t_1\neq t_2\\ t_1'\neq t_2'}}\big|P^*_{t_1,\sigma,R}\cap P^*_{t_2,\sigma,R}\big|\times\big|P^*_{t_1',\sigma,R}\cap P^*_{t_2',\sigma,R}\big|\\ &= \mathfrak{S}_2+\mathfrak{S}_3+\mathfrak{S}_4,\end{aligned}$$

where for $i=2,3,4$,

$$\begin{aligned}\mathfrak{S}_i &:= \sum_{\mathfrak{I}_i}\big|P^*_{t_1,\sigma,R}\cap P^*_{t_2,\sigma,R}\big|\times\big|P^*_{t_1',\sigma,R}\cap P^*_{t_2',\sigma,R}\big|, \text{ and}\\ \mathfrak{I}_i &:= \left\{\mathbb{I}=\{(t_1,t_2);(t_1',t_2')\}\;\middle|\;\begin{matrix}t_1,t_2,t_1',t_2'\in\mathcal{Q}(J), t_1\neq t_2,\ t_1'\neq t_2',\\ i = \text{number of distinct elements in } \mathbb{I}\end{matrix}\right\}.\end{aligned}$$

Without loss of generality, by interchanging the pairs (t_1,t_2) and (t_1',t_2') if necessary, we may assume that $h(D(t_1,t_2))\leq h(D(t_1',t_2'))$ for all quadruples $\mathbb{I}\in\mathfrak{I}_i$. We will continue to make this assumption for the treatment of all three terms $\mathfrak{S}_i$.

The claimed inequality in (19.2) is a consequence of the three main estimates below:

$$\mathbb{E}_{\mathbb{X}}(\mathfrak{S}_2)\leq CNM^{-2R-dJ}, \tag{19.7}$$

$$\mathbb{E}_{\mathbb{X}}(\mathfrak{S}_3)\leq CNM^{-3R-J}, \text{ and} \tag{19.8}$$

$$\mathbb{E}_{\mathbb{X}}(\mathfrak{S}_4)\leq CN^2M^{-4R}. \tag{19.9}$$

We will prove (19.9) in full detail, since this clearly makes the primary contribution among the three terms mentioned above. The other two estimates follow analogous and in fact simpler routes using the machinery developed in Chapters 16 and 17. We leave their verification to the reader.

The configuration type of the quadruple $\mathbb{I}=\{(t_1,t_2);(t_1',t_2')\}$ of distinct roots, as introduced in Section 16.6, plays a decisive role in the estimation of (19.9). Recalling the type definitions from that section, we decompose $\mathfrak{I}_4$ as

$$\mathfrak{I}_4 = \bigsqcup_{i=1}^{3} \mathfrak{I}_{4i} \text{ where } \mathfrak{I}_{4i} := \left\{ \mathbb{I} \in \mathfrak{I}_4 \middle| \begin{array}{l} \mathbb{I} \text{ is of type } i \text{ in the} \\ \text{sense of Definition 16.7} \end{array} \right\}.$$

This results in a corresponding decomposition of $\mathfrak{S}_4$:

$$\mathfrak{S}_4 = \mathfrak{S}_{41} + \mathfrak{S}_{42} + \mathfrak{S}_{43}, \quad \text{where} \quad \mathfrak{S}_{4i} = \sum_{\mathfrak{I}_{4i}} \left| P^*_{t_1,\sigma,R} \cap P^*_{t_2,\sigma,R} \right| \times \left| P^*_{t_1',\sigma,R} \cap P^*_{t_2',\sigma,R} \right|.$$

We will prove in Sections 19.2.1–19.2.3 below that

$$\mathbb{E}_{\mathbb{X}}\left[\mathfrak{S}_{4i}\right] \leq CN^2 M^{-4R} \quad \text{for } i = 1, 2, 3. \tag{19.10}$$

□

19.2.1 Expected value of $\mathfrak{S}_{41}$

We start with $\mathfrak{S}_{41}$, simplifying it initially along the same lines as in Proposition 19.1. As before, a summand in $\mathfrak{S}_{41}$ is nonzero if and only if the tuple $\{(t_1, t_2); (t_1', t_2')\}$ lies in the set

$$\left\{ \{(t_1, t_2); (t_1', t_2')\} \in \mathfrak{I}_{41} \middle| \begin{array}{l} P^*_{t_1,\sigma,R} \cap P^*_{t_2,\sigma,R} \neq \emptyset \\ P^*_{t_1',\sigma,R} \cap P^*_{t_2',\sigma,R} \neq \emptyset \end{array} \right\}, \tag{19.11}$$

which in turn is contained in

$$\left\{ \{(t_1, t_2); (t_1', t_2')\} \in \mathfrak{I}_{41} \middle| \begin{array}{l} \exists \text{ a unique tuple } (v_1, v_2, v_1', v_2') \in \Omega_N^4 \ni \\ P_{t_1,v_1} \cap P_{t_2,v_2} \cap [M^{-R}, M^{-R+1}] \times \mathbb{R}^d \neq \emptyset, \\ P_{t_1',v_1'} \cap P_{t_2',v_2'} \cap [M^{-R}, M^{-R+1}] \times \mathbb{R}^d \neq \emptyset, \\ \sigma(t_i) = v_i, \ \sigma(t_i') = v_i', \ i = 1, 2 \end{array} \right\}. \tag{19.12}$$

Incorporating this information into the simplification of the sum, we obtain

$$
\begin{aligned}
\mathfrak{S}_{41} &\leq \sum_{1} \underbrace{\frac{C_d M^{-(d+1)J}}{|\sigma(t_1) - \sigma(t_2)|} \times \frac{C_d M^{-(d+1)J}}{|\sigma(t_1') - \sigma(t_2')|}}_{\text{Lemma 1.5}} \\
&\leq CM^{-2(d+1)J} \sum_{2} \frac{T((t_1, v_1), (t_2, v_2))}{|v_1 - v_2|} \times \frac{T((t_1', v_1'), (t_2', v_2'))}{|v_1' - v_2'|} \\
&\leq CM^{-2(d+1)J} \sum_{\substack{u,u',z \\ \omega,\omega',v}} \frac{1}{\delta_\omega \delta_{\omega'}} \sum_{3} T((t_1, v_1), (t_2, v_2)) T((t_1', v_1'), (t_2', v_2')), \qquad (19.13) \\
&:= \overline{\mathfrak{S}}_{41}
\end{aligned}
$$

where the summations $\sum_1$ and $\sum_2$ range over the root quadruples in (19.11) and (19.12) respectively. The notation $T((t_1, v_1), (t_2, v_2))$ and δ_ω represent the same quantities as they did in Proposition 19.1, with their definitions in (19.6) and (12.8) respectively. Following the same reasoning that led to (19.4), in the last step we have stratified the sum in terms of the root vertices $u = D(t_1, t_2)$, $u' = D(t_1', t_2')$, $z = D(u, u')$ and the (splitting) slope vertices $\omega = D(v_1, v_2)$, $\omega' = D(v_1', v_2')$, $v = D(\omega, \omega')$, so that the summation $\sum_3$ takes place over the tube tuples in the collection $\mathcal{E}_{41} = \mathcal{E}_{41}[u, u', z; \omega, \omega, v; \varrho]$ defined in (17.11), with $\varrho = M^{-R}$, $C_1 = M$. We are now in a position to compute the expected value of $\mathfrak{S}_{41}$.

Lemma 19.4. *The estimate in* (19.10) *holds for* $i = 1$.

Proof. Let us refer to the bound $\overline{\mathfrak{S}}_{41}$ on $\mathfrak{S}_{41}$ defined by (19.13) that we obtained from the preliminary simplification. Assembling the various components of the estimation from the previous chapters, the expected value of $\mathfrak{S}_{41}$ is estimated as follows,

$$
\begin{aligned}
\mathbb{E}_{\mathbb{X}}\big(\mathfrak{S}_{41}\big) &\leq \mathbb{E}_{\mathbb{X}}\big(\overline{\mathfrak{S}}_{41}\big) \\
&\leq CM^{-2(d+1)J} \sum_{\substack{u,u',z \\ \omega,\omega',v}} \underbrace{\frac{1}{\rho_\omega \rho_{\omega'}}}_{\text{Corollary 12.4}} \sum_{3} \Pr\big(\sigma(t_i) = v_i,\ \sigma(t_i') = v_i',\ i = 1, 2\big)
\end{aligned}
$$

$$
\begin{aligned}
&\le CM^{-2(d+1)J} \sum_{\substack{u,u',z\\ \omega,\omega',v}} \frac{\#(\mathcal{E}_{41})}{\rho_\omega \rho_{\omega'}} \underbrace{\left(\frac{1}{2}\right)^{4N-\mu(\omega,h(u))-\mu(\omega',h(u'))-\mu(v,h(z))}}_{\text{(16.24) from Lemma 16.8}} \\
&\le CM^{-2(d+1)J} \sum_{\substack{u,u',z\\ \omega,\omega',v}} \underbrace{\left(\varrho^2 \rho_\omega \rho_{\omega'}\right)^2 2^{4N-2(\nu(\omega)+\nu(\omega'))} M^{-(d-1)(h(u)+h(u'))+2(d+1)J}}_{\text{bound on the size of } \mathcal{E}_{41} \text{ from Lemma 17.6}} \\
&\qquad\qquad\qquad\qquad \times \frac{1}{\rho_\omega \rho_{\omega'}} \times 2^{-4N+\mu(\omega,h(u))+\mu(\omega',h(u'))+\mu(v,h(z))} \\
&\le CM^{-4R} \mathfrak{S}^*_{41}, \quad \text{where}
\end{aligned}
$$

$$
\begin{aligned}
\mathfrak{S}^*_{41} := \sum_{\omega,\omega',v} 2^{-2(\nu(\omega)+\nu(\omega'))} M^{-[\lambda(\omega)+\lambda(\omega')]} \sum_{u,u',z} & 2^{\mu(\omega,h(u))+\mu(\omega',h(u'))+\mu(v,h(z))} \\
& \times M^{-(d-1)[h(u)+h(u')]}.
\end{aligned} \tag{19.14}
$$

It remains to use the appropriate summation results in Chapter 18 to show that $\mathfrak{S}^*_{41}$ is bounded above by a constant multiple of N^2. We start with the inner sum.

$$
\begin{aligned}
&\sum_{u,u',z} M^{-(d-1)(h(u)+h(u'))} 2^{\mu(\omega,h(u))+\mu(\omega',h(u'))+\mu(v,h(z))} \\
&\quad \le \sum_z 2^{\mu(v,h(z))} \underbrace{\Big[\sum_{u\subseteq z} M^{-(d-1)h(u)} 2^{\mu(\omega,h(u))}\Big]}_{\text{apply Lemma 18.2(i), } \beta=d-1} \underbrace{\Big[\sum_{u'\subseteq z} M^{-(d-1)h(u')} 2^{\mu(\omega',h(u'))}\Big]}_{\text{apply the same lemma again}} \\
&\quad \le \sum_z 2^{\mu(v,h(z))} \Big[M^{-dh(z)+\lambda(\omega)} 2^{\nu(\omega)}\Big] \Big[M^{-dh(z)+\lambda(\omega')} 2^{\nu(\omega')}\Big] \\
&\quad \le CM^{\lambda(\omega)+\lambda(\omega')} 2^{\nu(\omega)+\nu(\omega')} \underbrace{\Big[\sum_z 2^{\mu(v,h(z))} M^{-2dh(z)}\Big]}_{\text{apply Lemma 18.2(iv), } h(y)=0} \\
&\quad \le CM^{h(\omega)+h(\omega')} 2^{\nu(\omega)+\nu(\omega')}.
\end{aligned} \tag{19.15}
$$

Note that Lemma 18.2(iv) applies with $\beta = 2d$ since $2M^d < M^{2d}$ for $M \ge 2$ and $d \ge 2$. Inserting the expression in (19.15) into the inner sum of (19.14), we proceed to complete the outer sum in $\mathfrak{S}^*_{41}$.

$$\begin{aligned}
\mathfrak{S}^*_{41} &\leq C \sum_{\omega,\omega',v\in\mathcal{G}} M^{-\lambda(\omega)-\lambda(\omega')} 2^{-2(\nu(\omega)+\nu(\omega'))} \Big[M^{\lambda(\omega)+\lambda(\omega')} 2^{\nu(\omega)+\nu(\omega')} \Big] \\
&\leq C \sum_{\omega,\omega',v\in\mathcal{G}} 2^{-\nu(\omega)-\nu(\omega')} \\
&\leq C \sum_{v\in\mathcal{G}} \underbrace{\Big[\sum_{\omega\in\mathcal{G},\omega\subseteq v} 2^{-\nu(\omega)} \Big]}_{\text{apply Lemma 18.1(i), } \alpha=1} \times \underbrace{\Big[\sum_{\omega'\in\mathcal{G},\omega'\subseteq v} 2^{-\nu(\omega')} \Big]}_{\text{same lemma again}} \\
&\leq C \sum_{v\in\mathcal{G}} \big[N 2^{-\nu(v)} \big]^2 \leq C N^2 \sum_{v\in\mathcal{G}} 2^{-2\nu(v)} \leq C N^2,
\end{aligned}$$

where at the last step we have again used Lemma 18.1 (i) with $\alpha = 2$, and $\nu(\varpi_0) = 0$. This completes the proof of the lemma. □

19.2.2 Expected value of $\mathfrak{S}_{42}$

We turn to $\mathfrak{S}_{42}$ next. After the usual preliminary simplification similar to that of $\mathfrak{S}_{41}$, we find that $\mathfrak{S}_{42}$ is bounded by a sum $\overline{\mathfrak{S}}_{42}$ of the form (19.13), where

$$\overline{\mathfrak{S}}_{42} := C M^{-2(d+1)J} \sum{}' \frac{1}{\delta_\omega \delta_{\omega'}} \sum_3 T((t_1, v_1), (t_2, v_2)) T((t'_1, v'_1), (t'_2, v'_2)). \qquad (19.16)$$

In view of Lemma 16.9 we may assume, after a permutation of (t_1, t_2) and of (t'_1, t'_2) if necessary, that the outer sum $\sum'$ in (19.16) is over all vertex tuples (u, u', t) and $(\omega, \omega', \vartheta)$ in the root tree and the slope tree respectively, such that u, u', t lies on a single ray with $u' \subsetneq u$, while $\omega, \omega', \vartheta \in \mathcal{G}(\Omega_N)$, $\omega \cap \vartheta \neq \emptyset$, $\omega' \cap \vartheta \neq \emptyset$. The inner sum $\sum_3$ in $\overline{\mathfrak{S}}_{42}$ ranges over the collection $\mathcal{E}_{42} = \mathcal{E}_{42}[u, u', t; \omega, \omega', \vartheta; \varrho]$ defined in (17.12) with the usual $\varrho = M^{-R}$ and $C_1 = M$.

Lemma 19.5. *The estimate in* (19.10) *holds for* $i = 2$.

Proof. As in Lemma 19.4, the evaluation of the expectation requires a combination of the appropriate probabilistic estimate from Section 16.2 (specifically Lemma 16.9), size estimate of $\mathcal{E}_{42}$ from Section 17.3.2 (specifically Lemma 17.8) and the summation results from Chapter 18. Putting these together, we obtain

$$\begin{aligned}
\mathbb{E}_{\mathbb{X}}\big(\mathfrak{S}_{42}\big) &\leq \mathbb{E}_{\mathbb{X}}\big(\overline{\mathfrak{S}}_{42}\big) \\
&\leq CM^{-2(d+1)J} \sum{}' \frac{1}{\rho_\omega \rho_{\omega'}} \sum_{3} \Pr\big(\sigma(t_i) = v_i,\ \sigma(t_i') = v_i',\ i = 1,2\big) \\
&\leq CM^{-2(d+1)J} \sum{}' \frac{\#\big(\mathcal{E}_{42}\big)}{\rho_\omega \rho_{\omega'}} \underbrace{\left(\frac{1}{2}\right)^{4N - \mu(\omega, h(u)) - \mu(\omega', h(u')) - \mu(\vartheta, h(t))}}_{\text{(16.26) from Lemma 16.9}} \\
&\leq CM^{-4R}\big[\mathfrak{S}_{42}^* + \mathfrak{S}_{42}^\circ\big],
\end{aligned}$$

where the closed form expressions for $\mathfrak{S}_{42}^*$ and $\mathfrak{S}_{42}^\circ$ at the last step are obtained from the count on the size of $\mathcal{E}_{42}$ from Lemma 17.8, and reflect the two complementary cases considered therein. To be precise,

$$\mathfrak{S}_{42}^* := \varrho^{-1} \sideset{}{'}\sum_{u' \subseteq t \subseteq u} \rho_{\omega'} \min\big[\varrho\rho_\omega, M^{-h(t)}\big] M^{-(d-1)\big(h(t)+h(u')\big)} \times 2^{\mu(\omega,h(u)) + \mu(\omega',h(u')) + \mu(\vartheta,h(t)) - m[\omega,\omega',\vartheta]}, \quad \text{and} \tag{19.17}$$

$$\mathfrak{S}_{42}^\circ := \varrho^{-2} \sideset{}{'}\sum_{t \subsetneq u' \subsetneq u} \min\big[\varrho\rho_\omega, M^{-h(t)}\big] \min\big[\varrho\rho_{\omega'}, M^{-h(t)}\big] M^{-2(d-1)h(t)} \times 2^{\mu(\omega,h(u)) + \mu(\omega',h(u')) + \mu(\vartheta,h(t)) - m[\omega,\omega',\vartheta]}, \tag{19.18}$$

where the notation $\sum'_{\mathcal{P}}$ indicates the subsum of $\sum'$ subject to the additional requirement $\mathcal{P}$. These two quantities are estimated via the usual channels. Lemma 17.8 places certain restrictions on the spatial location of t, but for a large part of the proof the full strength of these statements will not be needed. For instance, replacing $\min(\varrho\rho_\omega, M^{-h(t)})$ in (19.17) by $\varrho\rho_\omega$, we arrive at the following bound for $\mathfrak{S}_{42}^*$:

$$\mathfrak{S}_{42}^* \leq \sideset{}{'}\sum_{u' \subseteq t \subseteq u} \rho_\omega \rho_{\omega'} M^{-(d-1)\big(h(t)+h(u')\big)} 2^{\mu(\omega,h(u)) + \mu(\omega',h(u')) + \mu(\vartheta,h(t)) - m[\omega,\omega',\vartheta]}$$

$$\leq \sum_{\omega,\omega',\vartheta} \rho_\omega \rho_{\omega'} 2^{-m[\omega,\omega',\vartheta]} \mathfrak{S}^*_{42}(\text{inner}), \tag{19.19}$$

where the inner expression $\mathfrak{S}^*_{42}(\text{inner})$ is a sequence of three summations in root vertices, the computation of each requiring a suitable form of Lemma 18.2. Precisely,

$$\begin{aligned}
\mathfrak{S}^*_{42}(\text{inner}) &:= \sum_{\substack{u,u',t \\ u'\subseteq t\subseteq u}} M^{-(d-1)\left(h(t)+h(u')\right)} 2^{\mu(\omega,h(u))+\mu(\omega',h(u'))+\mu(\vartheta,h(t))} \\
&= \sum_{\substack{(t,u) \\ t\subseteq u}} M^{-(d-1)h(t)} 2^{\mu(\omega,h(u))+\mu(\vartheta,h(t))} \Big[\sum_{u'\subseteq t} M^{-(d-1)h(u')} 2^{\mu(\omega',h(u'))} \Big] \\
&\leq C \sum_{\substack{(t,u) \\ t\subseteq u}} M^{-(d-1)h(t)} 2^{\mu(\omega,h(u))+\mu(\vartheta,h(t))} \underbrace{\Big[2^{\nu(\omega')} M^{-dh(t)+\lambda(\omega')} \Big]}_{\text{from Lemma 18.2(i), } \beta = d-1} \\
&\leq C 2^{\nu(\omega')} M^{\lambda(\omega')} \sum_{u} 2^{\mu(\omega,h(u))} \sum_{t:t\subseteq u} M^{-(2d-1)h(t)} 2^{\mu(\vartheta,h(t))} \\
&\leq C 2^{\nu(\omega')} M^{\lambda(\omega')} \sum_{u} 2^{\mu(\omega,h(u))} \underbrace{\Big[M^{-(2d-1)h(u)} 2^{\nu(\vartheta)} \Big]}_{\text{from Lemma 18.2(iii), } \beta = 2d-1} \\
&\leq C 2^{\nu(\omega')+\nu(\vartheta)} M^{\lambda(\omega')} \sum_{u} 2^{\mu(\omega,h(u))} M^{-(2d-1)h(u)} \\
&\leq C 2^{\nu(\omega')+\nu(\vartheta)+\nu(\omega)} M^{\lambda(\omega')},
\end{aligned} \tag{19.20}$$

where the summation in u in the last step also follows from Lemma 18.2(iii), since $\beta = 2d-1 > d$. Inserting the estimate (19.20) of $\mathfrak{S}^*_{42}(\text{inner})$ into (19.19), we proceed to simplify the outer sum. Let us recall from Lemma 17.7 that $\{\omega, \omega', \vartheta\}$ can be rearranged as $\{\varpi_1, \varpi_2, \varpi_3\}$ satisfying (17.8), and that $m[\omega, \omega', \vartheta]$ is defined as in (17.9). Since the definition of m involves two possibilities, we write $\sum^{[a]}$ and $\sum^{[b]}$ to denote the sum over vertex triples $(\omega, \omega', \vartheta)$ for which $\varpi_3 \not\subseteq \varpi_2$ and $\varpi_3 \subseteq \varpi_2$ respectively. This means that

$$\mathfrak{S}^*_{42} \leq C \sum_{\omega,\omega',\vartheta} \rho_\omega \rho_{\omega'} 2^{-m[\omega,\omega',\vartheta]} \Big[2^{\nu(\omega')+\nu(\vartheta)+\nu(\omega)} M^{\lambda(\omega')} \Big]$$

$$\leq C\Big[\sum^{[a]} + \sum^{[b]}\Big]\rho_\omega \rho_{\omega'} 2^{-m[\omega,\omega',\vartheta]}\big[2^{\nu(\omega')+\nu(\vartheta)+\nu(\omega)} M^{\lambda(\omega')}\big].$$

Using the bounds

$$\rho_{\omega'} M^{\lambda(\omega')} \leq C \quad \text{and} \quad \rho_\omega \leq C M^{-\lambda(\omega)} \leq C M^{-\lambda(\varpi_1)},$$

the estimation is completed as follows,

$$\begin{aligned}
&\sum^{[a]} \rho_\omega \rho_{\omega'} 2^{-m[\omega,\omega',\vartheta]}\Big[2^{\nu(\omega')+\nu(\vartheta)+\nu(\omega)} M^{\lambda(\omega')}\Big] \\
&\qquad \leq C \sum_{\varpi_1} M^{-\lambda(\varpi_1)} 2^{\nu(\varpi_1)} \Big[\sum_{\substack{\varpi_2 \\ \varpi_2 \subseteq \varpi_1}} 2^{-\nu(\varpi_2)}\Big] \times \Big[\sum_{\substack{\varpi_3 \\ \varpi_3 \subseteq \varpi_1}} 2^{-\nu(\varpi_3)}\Big] \\
&\qquad \leq C \sum_{\varpi_1} M^{-\lambda(\varpi_1)} 2^{\nu(\varpi_1)} \underbrace{\big(N 2^{-\nu(\varpi_1)}\big)^2}_{\text{Lemma 18.1 (i) twice}} \\
&\qquad \leq C N^2 \underbrace{\sum_{\varpi_1} M^{-\lambda(\varpi_1)} 2^{-\nu(\varpi_1)}}_{\text{apply Lemma 18.1(ii)}} \leq C N^2.
\end{aligned}$$

The same bound holds for $\sum^{[b]}$, and is proved along similar lines:

$$\begin{aligned}
&\sum^{[b]} \rho_\omega \rho_{\omega'} 2^{-m[\omega,\omega',\vartheta]}\Big[2^{\nu(\omega')+\nu(\vartheta)+\nu(\omega)} M^{\lambda(\omega')}\Big] \\
&\qquad \leq C \sum_{\substack{\varpi_1,\varpi_2 \\ \varpi_2 \subseteq \varpi_1}} M^{-\lambda(\varpi_1)} \sum_{\substack{\varpi_3 \\ \varpi_3 \subseteq \varpi_2}} 2^{-\nu(\varpi_3)} \leq C \sum_{\substack{\varpi_1,\varpi_2 \\ \varpi_2 \subseteq \varpi_1}} M^{-\lambda(\varpi_1)} \underbrace{\big[N 2^{-\nu(\varpi_2)}\big]}_{\text{Lemma 18.1 (i)}} \\
&\qquad \leq C N \sum_{\varpi_1} M^{-\lambda(\varpi_1)} \sum_{\varpi_2 \subseteq \varpi_1} 2^{-\nu(\varpi_2)} \leq C N \sum_{\varpi_1} M^{-\lambda(\varpi_1)} \underbrace{\big[N 2^{-\nu(\varpi_1)}\big]}_{\text{Lemma 18.1 (i)}} \\
&\qquad \leq C N^2 \underbrace{\sum_{\varpi_1} M^{-\lambda(\varpi_1)} 2^{-\nu(\varpi_1)}}_{\text{apply Lemma 18.1(ii)}} \leq C N^2.
\end{aligned}$$

This completes the estimation of $\mathfrak{S}^*_{42}$.

We briefly remark on the analysis of $\mathfrak{S}^\circ_{42}$. For $d \geq 3$, replacing the minima in

(19.18) by the trivial bounds $\varrho\rho_\omega$ and $\varrho\rho_{\omega'}$ results in an expression analogous to that of $\mathfrak{S}^*_{42}$:

$$\mathfrak{S}^\circ_{42} \leq \sum_{t\subsetneq u'\subsetneq u}{}' \rho_\omega\rho_{\omega'} M^{-2(d-1)h(t)} 2^{\mu(\omega,h(u))+\mu(\omega',h(u'))+\mu(\vartheta,h(t))-m[\omega,\omega',\vartheta]}.$$

This term is estimated exactly the same way as $\mathfrak{S}^*_{42}$, since Lemma 18.2(iii) applies as before with $\beta = 2(d-1) > d$ per our choice of d. The bound obtained is a constant multiple of N. These details are omitted to avoid repetition. We only present the case $d = 2$, where Lemma 18.2 does not give the desired consequence, and the treatment of which exhibits a slight departure from the norm so far. For $d = 2$, inserting the bound $\min(\varrho\rho_\omega, M^{-h(t)}) \leq \varrho\rho_\omega$ into (19.18) yields

$$\mathfrak{S}^\circ_{42} \leq \sum_{t\subsetneq u'\subsetneq u}{}' 2^{\mu(\omega,h(u))+\mu(\omega',h(u'))+\mu(\vartheta,h(t))-m[\omega,\omega',\vartheta]} \times \begin{cases} \rho_\omega\rho_{\omega'} M^{-2h(t)} & \text{if } M^{-h(t)} \geq \varrho\rho_{\omega'}, \\ \varrho^{-1}\rho_\omega M^{-3h(t)} & \text{if } M^{-h(t)} < \varrho\rho_{\omega'}. \end{cases} \tag{19.21}$$

Further, Lemma 17.8(ii) prescribes that t cannot be arbitrarily placed inside u', but must lie within the union of at most $2M$ thin rectangles of dimension $\varrho\rho_{\omega'} \times M^{-h(u')}$ each. Using this information, we sum the expression (19.21) in t as follows: if $\sum_1$ and $\sum_2$ denote the summations in t with $t \subseteq u'$ and $\mathcal{E}_{42} \neq \emptyset$ subject to the conditions $M^{-h(t)} \geq \varrho\rho_{\omega'}$ and $M^{-h(t)} < \varrho\rho_{\omega'}$ respectively, then

$$\begin{aligned} &\rho_\omega\rho_{\omega'} \sum_1 M^{-2h(t)} 2^{\mu(\vartheta,h(t))} + \varrho^{-1}\rho_\omega \sum_2 M^{-3h(t)} 2^{\mu(\vartheta,h(t))} \\ &\quad\leq C\rho_\omega\rho_{\omega'} \left[2^{\nu(\vartheta)} M^{-2h(u')}\right] + C\varrho^{-1}\rho_\omega \left[2^{\nu(\vartheta)} M^{-h(u')} \big(\varrho\rho_{\omega'}\big) \min\big[\varrho\rho_{\omega'}, M^{-h(u')}\big]\right] \\ &\quad\leq C\rho_\omega\rho_{\omega'} 2^{\nu(\vartheta)} M^{-2h(u')}, \end{aligned}$$

where both sums have been evaluated using Lemma 18.3 with $d = 2$, $r = 1$, $\varpi = \vartheta$, $\beta = \epsilon = M^{-h(u')}$ and $\gamma = \varrho\rho_{\omega'}$. In particular, $\sum_1$ appeals to part (i) of this lemma with $\alpha = 2$ while $\sum_2$ uses part (ii) with $\alpha = 3$. Incorporating this into (19.21), we find that

$$\mathfrak{S}^\circ_{42} \le \sum_{\omega,\omega',\vartheta} \rho_\omega \rho_{\omega'} 2^{\nu(\vartheta) - m[\omega,\omega',\vartheta]} \mathfrak{S}^\circ_{42}(\text{inner}), \qquad \text{where} \tag{19.22}$$

$$\mathfrak{S}^\circ_{42}(\text{inner}) := \sum_u 2^{\mu(\omega,h(u))} \underbrace{\sum_{u' \subseteq u} M^{-2h(u')} 2^{\mu(\omega',h(u'))}}_{\text{apply Lemma 18.2(ii)}} \tag{19.23}$$

$$\le C 2^{\nu(\varpi_2)} \lambda(\varpi_2) \underbrace{\sum_u M^{-2h(u)} 2^{\mu(\omega,h(u))}}_{\text{apply the same lemma again}} \tag{19.24}$$

$$\le C 2^{\nu(\varpi_2)+\nu(\varpi_1)} \lambda(\varpi_2)\lambda(\varpi_1), \tag{19.25}$$

We pause for a moment to explain these steps. In the first application of Lemma 18.2(ii) in (19.23) above we have used, in addition to $h(u') \le \lambda(\omega')$, the fact that

$$h(u') = h(D(t_1', t_2')) \le h(t) = h(D(t_2, t_2')) \le \lambda(D(v_2, v_2')) = \lambda(\vartheta),$$

which is a consequence of weak stickiness. Since one of ω' and ϑ is contained in the other, this implies that $\mu(\omega', h(u')) = \mu(\vartheta, h(u'))$. Hence Lemma 18.2(ii), applied once with $\varpi = \omega'$ and again with $\varpi = \vartheta$, yields

$$\sum_{u' \subseteq u} M^{-2h(u')} 2^{\mu(\omega',h(u'))} \le C M^{-2h(u)} \min\left[2^{\nu(\vartheta)}\lambda(\vartheta), 2^{\nu(\omega')}\lambda(\omega')\right]$$

$$\le C\lambda(\varpi_2) 2^{\nu(\varpi_2)} M^{-2h(u)}.$$

The second application of Lemma 18.2(ii) in (19.24) uses a similar argument relying on the fact that $h(u) \le \lambda(\varpi_1)$. Inserting (19.25) into (19.22), the estimation of $\mathfrak{S}^\circ_{42}$ can now be completed in the same way as for $\mathfrak{S}^*_{42}$:

$$\mathfrak{S}^\circ_{42} \le C \sum_{\omega,\omega',\vartheta} \rho_\omega \rho_{\omega'} \lambda(\varpi_2)\lambda(\varpi_1) 2^{\nu(\vartheta)+\nu(\varpi_1)+\nu(\varpi_2) - m[\omega,\omega',\vartheta]}$$

$$
\begin{aligned}
&\leq C \sum_{\omega,\omega',\vartheta} M^{-\lambda(\varpi_1)-\lambda(\varpi_2)}\lambda(\varpi_1)\lambda(\varpi_2)2^{\nu(\vartheta)+\nu(\varpi_1)+\nu(\varpi_2)-m[\omega,\omega',\vartheta]} \\
&\leq C \sum^{[a]} M^{-\frac{1}{2}\lambda(\varpi_1)-\frac{1}{2}\lambda(\varpi_2)}2^{-\nu(\varpi_3)-\nu(\varpi_2)+\nu(\varpi_1)} + \sum^{[b]} M^{-\frac{1}{2}\lambda(\varpi_1)-\frac{1}{2}\lambda(\varpi_2)}2^{-\nu(\varpi_3)} \\
&\leq CN,
\end{aligned}
$$

where the symbols $\sum^{[a]}$ and $\sum^{[b]}$ carry the same meaning as they did in the estimation of $\mathfrak{S}^*_{42}$ and the last step involves several summations all of which have used appropriate parts of Lemma 18.1. The estimation of $\mathfrak{S}_{42}$ is complete. □

19.2.3 Expected value of $\mathfrak{S}_{43}$

Lemma 19.6. *The estimate in* (19.10) *holds for* $i = 3$.

Proof. After the usual initial processing of $\mathfrak{S}_{43}$ which we omit, we reduce to the following estimate:

$$
\begin{aligned}
\mathbb{E}_{\mathbb{X}}\big(\mathfrak{S}_{43}\big) &\leq CM^{-2(d+1)J}\sum{}' \frac{\#(\mathcal{E}_{43})}{\rho_\omega\rho_{\omega'}}\left(\frac{1}{2}\right)^{4N-\mu(\omega,h(u))-\mu(\vartheta,h(s_1))-\mu(\vartheta_2,h(s_2))} \\
&\leq C\sum{}' \big(\rho_\omega\rho_{\omega'}\big)^{-1}\big(\varrho\rho_{\omega'}\big)^2 M^{-2(d-1)h(s_2)}\prod_{i=1}^{2}\Big[\min[\varrho\rho_\omega, M^{-h(s_i)}]\Big] \\
&\qquad\qquad \times 2^{-m[\omega,\omega',\vartheta_1,\vartheta_2]+\mu(\omega,h(u))+\mu(\vartheta_1,h(s_1))+\mu(\vartheta_2,h(s_2))} \\
&\leq CM^{-4R}\big[\mathfrak{S}^*_{43}+\mathfrak{S}^\circ_{43}\big],
\end{aligned}
$$

where $\sum'$ denotes the sum over all tuples (u, s_1, s_2) in the root tree and $(\omega, \omega', \vartheta_1, \vartheta_2)$ in the slope tree such that $s_1, s_2 \subseteq u$, $h(u) \leq h(s_1) \leq h(s_2)$, $\rho_\omega \leq \rho_{\omega'}$ and for which $\mathcal{E}_{43}$ is nonempty. The second inequality displayed above uses the estimate on $\#(\mathcal{E}_{43})$ obtained in Lemma 17.9, with an additional simplification resulting from $\min(\varrho\rho_{\omega'}, M^{-h(s_i)}) \leq \varrho\rho_{\omega'}$. The quantities $\mathfrak{S}^*_{43}$ and $\mathfrak{S}^\circ_{43}$ refer to the subsum of $\sum'$ under the additional constraints of $M^{-h(s_1)} \geq \varrho\rho_\omega$ and $M^{-h(s_1)} < \varrho\rho_\omega$ respectively.

Thus

$$\begin{aligned}
\mathfrak{S}_{43}^{*} &= \varrho^{-1} \sideset{}{'}\sum_{M^{-h(s_1)} \geq \varrho\rho_\omega} \rho_{\omega'} \min[\varrho\rho_\omega, M^{-h(s_2)}] M^{-2(d-1)h(s_2)} \\
&\qquad\qquad \times 2^{-m[\omega,\omega',\vartheta_1,\vartheta_2]+\mu(\omega,h(u))+\mu(\vartheta_1,h(s_1))+\mu(\vartheta_2,h(s_2))} \\
&=: \varrho^{-1} \sum_{\omega,\omega',\vartheta_1,\vartheta_2} \rho_{\omega'} 2^{-m[\omega,\omega',\vartheta_1,\vartheta_2]} \mathfrak{S}_{43}^{*}(\text{inner}), \qquad \text{and} \qquad (19.26) \\
\mathfrak{S}_{43}^{\circ} &= \varrho^{-2} \sideset{}{'}\sum_{M^{-h(s_1)} < \varrho\rho_\omega} \frac{\rho_{\omega'}}{\rho_\omega} \min[\varrho\rho_\omega, M^{-h(s_2)}] M^{-2(d-1)h(s_2)-h(s_1)} \\
&\qquad\qquad \times 2^{-m[\omega,\omega',\vartheta_1,\vartheta_2]+\mu(\omega,h(u))+\mu(\vartheta_1,h(s_1))+\mu(\vartheta_2,h(s_2))} \\
&=: \varrho^{-2} \sum_{\omega,\omega',\vartheta_1,\vartheta_2} \frac{\rho_{\omega'}}{\rho_\omega} 2^{-m[\omega,\omega',\vartheta_1,\vartheta_2]} \mathfrak{S}_{43}^{\circ}(\text{inner}). \qquad (19.27)
\end{aligned}$$

For the purpose of simplifying $\mathfrak{S}_{43}^{*}(\text{inner})$, we recall from Lemma 17.9(ii) that $s_2 \subsetneq u$ has sidelength no more than $M^{-h(s_1)}$, and moreover, is constrained to lie in the union of at most $2^d M$ parallelepipeds with $(d-1)$ long directions and one short direction, of dimensions $M^{-h(s_1)}$ and $\varrho\rho_\omega$ respectively. Denoting by $\sum_{s_2}^{*}$ the summation over all such cubes s_2, we find that

$$\begin{aligned}
&\sideset{}{^*}\sum_{s_2} 2^{\mu(\vartheta_2,h(s_2))} M^{-2(d-1)h(s_2)} \min\left[\varrho\rho_\omega, M^{-h(s_2)}\right] \\
&\leq \varrho\rho_\omega \sideset{}{^*}\sum_{M^{-h(s_2)} \geq \varrho\rho_\omega} M^{-2(d-1)h(s_2)} 2^{\mu(\vartheta_2,h(s_2))} + \sideset{}{^*}\sum_{M^{-h(s_2)} < \varrho\rho_\omega} M^{-(2d-1)h(s_2)} 2^{\mu(\vartheta_2,h(s_2))} \\
&\leq \varrho\rho_\omega \mathfrak{s}_+ + \mathfrak{s}_- \\
&\leq C\left[\varrho\rho_\omega 2^{\nu(\vartheta_2)} M^{-2(d-1)h(s_1)} + 2^{\nu(\vartheta_2)} \left(\varrho\rho_\omega\right)^d M^{-(d-1)h(s_1)}\right] \\
&\leq C\varrho\rho_\omega 2^{\nu(\vartheta_2)} M^{-2(d-1)h(s_1)}, \qquad (19.28)
\end{aligned}$$

where $\mathfrak{s}_\pm$ are defined as in (18.1), and estimated according to Lemma 18.3, with the parameters being set at $\epsilon = \beta = M^{-h(s_1)}$, $\gamma = \varrho\rho_\omega$, $\varpi = \vartheta_2$ for both. The value of α is $2(d-1)$ for $\mathfrak{s}_+$ and $(2d-1)$ for $\mathfrak{s}_-$. A similar argument applies for the summation in s_1 with $M^{-h(s_1)} \geq \varrho\rho_\omega$. According to Lemma 17.9(i), s_1 has to lie in u and within a distance at most $C\Delta$ from the boundary of some child of u. Hence the range of s_1 lies within the union of at most dM parallelepipeds, each of dimension $M^{-h(u)}$ in $(d-1)$ directions and $C\Delta$ in the remaining one. Denoting by $\sum_{s_1}^*$ the relevant sum, and applying Lemma 18.3 again with $\alpha = 2(d-1)$, $r = 1$, $\epsilon = \beta = M^{-h(u)}$, $\gamma = \varrho\rho_\omega$, $\varpi = \vartheta_1$,

$$\sum_{s_1}^{*} M^{-2(d-1)h(s_1)} 2^{\mu(\vartheta_1,h(s_1))} \leq \mathfrak{s}_+ \leq 2^{\nu(\vartheta_1)} M^{-2(d-1)h(u)}. \tag{19.29}$$

Inserting the estimates (19.28) and (19.29), we arrive at the following bound on $\mathfrak{S}_{43}^*(\text{inner})$:

$$\begin{aligned}
\mathfrak{S}_{43}^*(\text{inner}) &= \sum_u \sum_{s_1}^{*} 2^{\mu(\omega,h(u))+\mu(\vartheta_1,h(s_1))} \\
&\qquad\qquad \times \Big[\sum_{s_2}^{*} 2^{\mu(\vartheta_2,h(s_2))} M^{-2(d-1)h(s_2)} \min\big[\varrho\rho_\omega, M^{-h(s_2)}\big]\Big] \\
&\leq C \sum_{u,s_1} 2^{\mu(\omega,h(u))+\mu(\vartheta_1,h(s_1))} \Big[2^{\nu(\vartheta_2)} \varrho\rho_\omega M^{-2(d-1)h(s_1)}\Big] \\
&\leq \varrho\rho_\omega 2^{\nu(\vartheta_2)} \sum_u 2^{\mu(\omega,h(u))} \sum_{s_1}^{*} M^{-2(d-1)h(s_1)} 2^{\mu(\vartheta_1,h(s_1))} \\
&\leq \varrho\rho_\omega 2^{\nu(\vartheta_2)} \sum_u 2^{\mu(\omega,h(u))} \Big[M^{-2(d-1)h(u)} 2^{\nu(\vartheta_1)}\Big] \\
&\leq \varrho\rho_\omega 2^{\nu(\vartheta_2)+\nu(\vartheta_1)} \sum_u 2^{\mu(\omega,h(u))} M^{-2(d-1)h(u)} \\
&\leq \varrho\rho_\omega 2^{\nu(\vartheta_2)+\nu(\vartheta_1)+\nu(\varpi_1)} \lambda(\varpi_1),
\end{aligned} \tag{19.30}$$

where ϖ_1 is the youngest common ancestor of $\omega, \omega', \vartheta_1, \vartheta_2$, and hence $\lambda(\varpi_1) \geq h(u)$. The last estimate follows from Lemma 18.2, invoking part (iii) if $d \geq 3$ and part(i) if $d = 2$. An analogous sequence of steps, the details of which are left to the reader, can be executed to estimate $\mathfrak{S}_{43}^{\circ}(\text{inner})$, the only distinction being that the space restrictions are now dictated by Lemma 17.9(iii), so that the summation in s_2 invokes Lemma 18.3 with $r = d-1$, $\beta = \varrho \min(M^{-\lambda(\omega)}, M^{-\lambda(\omega')})$, $\gamma = M^{-h(s_1)}$. The outcome of this is that

$$\mathfrak{S}_{43}^{\circ}(\text{inner}) \le \varrho^2 \rho_\omega \min(M^{-\lambda(\omega)}, M^{-\lambda(\omega')}) 2^{\nu(\vartheta_2)+\nu(\vartheta_1)+\nu(\varpi_1)} \lambda(\varpi_1). \tag{19.31}$$

Substituting (19.30) into (19.26) and (19.31) into (19.27) leads to the following simpler sum over slope vertices:

$$\mathfrak{S}_{43}^{*} + \mathfrak{S}_{43}^{\circ} \le C \sum_{\omega,\omega',\vartheta_1,\vartheta_2} M^{-\lambda(\omega)-\lambda(\omega')} 2^{-m[\omega,\omega',\vartheta_1,\vartheta_2]+\nu(\varpi_1)+\nu(\vartheta_1)+\nu(\vartheta_2)}.$$

In order to complete the summation, let us recall that the sum, ostensibly over four parameters, in fact ranges over at most three vertices $\{\varpi_1, \varpi_2, \varpi_3\}$, which is a rearrangement of the quadruple $\{\omega, \omega', \vartheta_1, \vartheta_2\}$ satisfying (17.8). However, it is not apriori possible to assign a unique correspondence between these two sets of vertices. Indeed, as already indicated in the last paragraph of Chapter 16, the configuration type of the slopes (which does not in general mimic the configuration type of the roots) dictates which vertex or vertices of the quadruple $\{\omega, \omega', \vartheta_1, \vartheta_2\}$ represents ϖ_i after the rearrangement. A careful analysis of the possible structures of $\omega, \omega', \vartheta_1, \vartheta_2$, as depicted in Figure 19.1, shows that

$$\begin{aligned} &M^{-\lambda(\omega)-\lambda(\omega')} \lambda(\varpi_1) 2^{-m[\omega,\omega',\vartheta_1,\vartheta_2]+\nu(\varpi_1)+\nu(\vartheta_1)+\nu(\vartheta_2)} \\ &\qquad \le M^{-2\lambda(\varpi_1)} \lambda(\varpi_1) \times \begin{cases} 2^{-\nu(\varpi_3)-\nu(\varpi_2)+\nu(\varpi_1)} & \text{if } \varpi_3 \not\subseteq \varpi_2 \\ 2^{-\nu(\varpi_3)} & \text{if } \varpi_3 \subseteq \varpi_2. \end{cases} \end{aligned}$$

The expression on the right hand side is of the type already considered in the estimation of $\mathfrak{S}_{42}^{*}$ and $\mathfrak{S}_{42}^{\circ}$. In particular, it is summable in $\varpi_1, \varpi_2, \varpi_3$ using repeated applications of Lemma 18.1 and yields the desired bound of CN^2.

□

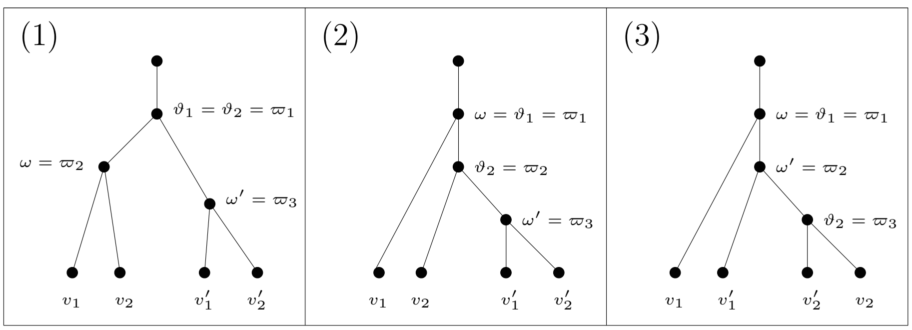

Figure 19.1: A partial list of 4-slope configurations for 4 roots of type 3, with distinct $\{\varpi_1, \varpi_2, \varpi_3\}$. Other configurations (where partial coincidences may arise) are possible after permutation of $\{v_1, v_1', v_2, v_2'\}$ in these diagrams.

Chapter 20

Future work

In this final chapter, we discuss some potential ideas for future work to come more or less directly out of the work presented in the current document. Some of these ideas have yet to be seriously considered, while others, specifically those in Section 20.2 are presently being developed in collaboration with Dimitrios Karslidis and Malabika Pramanik. All material presented in this chapter should be considered work in progress and not established fact.

20.1 Maximal functions over other collections of sticky objects

We briefly mentioned at the end of Section 1.5 that the notion of stickiness is not unique to collections of tubes. Indeed, curves and spheres are among two of the most well studied objects that can also be grouped into sticky collections; see [8], [29], [40], [30] for example. Besicovitch and Rado [8] proved that there exist measure zero sets in $\mathbb{R}^d$ that contain a sphere of every radius (Kinney [29] proved the same concurrently in two dimensions). Following the same argument used to deduce the unboundedness of the Kakeya-Nikodym maximal operator, presented in Section 1.3, the existence of such sets means that the *circular maximal function* defined as

$$M^{\mathrm{circ}} f(x) := \sup_{r\in\mathbb{R}^+} \frac{1}{|C(x,r)|} \int_{C(x,r)} |f(y)| dy, \tag{20.1}$$

must be unbounded as an operator on $L^p(\mathbb{R}^d)$ for $p \in [1,\infty)$. Here, $C(x,r)$ denotes the sphere centred at x with radius r. Naturally, we then consider the *restricted circular maximal function* defined as

$$M^{\mathrm{circ}}_\delta f(x) := \sup_{r\in\mathbb{R}^+} \frac{1}{|C_\delta(x,r)|} \int_{C_\delta(x,r)} |f(y)| dy, \tag{20.2}$$

where $C_\delta(x,r)$ denotes the δ-thickening of the circle $C(x,r)$. Kolasa and Wolff derived δ-dependent bounds on the size of the norm of an equivalent operator. Interestingly, their bounds turn out to be optimal in dimensions $d \geq 3$ [30].

Like their classical counterparts, zero measure circular Kakeya sets seem to exhibit a stickiness quality in that many spheres positioned close to each other in space must have comparable radii. This motivates the notion of circular Kakeya-type sets.

Definition 20.1. *Fix a set of radii $\Phi \subseteq \mathbb{R}^+$. If for some fixed constant $A_0 \geq 1$ and any choice of integer $N \geq 1$, there exist*

- *a number $0 < \delta_N \ll 1$, $\delta_N \searrow 0$ as $N \nearrow \infty$, and*
- *a collection of spheres $\{S_t^{(N)}\}$ with radii in Φ and thickness at most δ_N*

obeying

$$\lim_{N\to\infty} \frac{|E_N^*(A_0)|}{|E_N|} = \infty, \quad \textit{with} \quad E_N := \bigcup_t S_t^{(N)}, \quad E_N^*(A_0) := \bigcup_t A_0 S_t^{(N)}, \tag{20.3}$$

then we say that Φ admits circular Kakeya-type sets*. Here, $A_0 S_t^{(N)}$ denotes the sphere with the same centre and thickness as $S_t^{(N)}$ but A_0 times the radius.*

The existence of such sets for a given $\Phi \subseteq \mathbb{R}^+$ should imply the L^p-unboundedness of a corresponding maximal operator analogous to (20.1), where the supremum is now taken over all radii $r \in \Phi$. While the geometry is considerably different in this context, it is perhaps reasonable to hope that the construction of circular Kakeya-type sets could be approached using some of the ideas developed in the work of Bateman and Katz, as well as in this document.

Beside for spheres, it may also be worthwhile to consider if the techniques developed in this document can be applied to the construction of Kakeya-type sets over flat slabs. Flat slabs are essentially thickened circles and can be defined as the intersection of a hyperplane with a circular cylinder whose principal axis sits normal to the hyperplane. We say that such a slab is oriented in the direction given by this normal vector.

Definition 20.2. *Fix a set of directions $\Omega \subseteq \mathbb{S}^d$. If for some fixed constant $A_0 \geq 1$ and any choice of integer $N \geq 1$, there exist*

- *a number $0 < \delta_N \ll 1$, $\delta_N \searrow 0$ as $N \nearrow \infty$, and*
- *a collection of δ_N-thickened slabs $\{R_t^{(N)}\}$ with orientations lying inside Ω and diameter at least 1*

obeying

$$\lim_{N\to\infty} \frac{|E_N^*(A_0)|}{|E_N|} = \infty, \quad \text{with} \quad E_N := \bigcup_t R_t^{(N)}, \quad E_N^*(A_0) := \bigcup_t A_0 R_t^{(N)}, \tag{20.4}$$

then we say that Ω admits Kakeya-type sets over slabs*. Here, $A_0 R_t^{(N)}$ denotes the slab with the same centre, orientation and thickness as $R_t^{(N)}$ but A_0 times its diameter.*

Again, the existence of such sets would have implications for the L^p behaviour of a corresponding maximal operator. How much the techniques discussed in this document can be used to study these and other like sets is not currently known, but the potential for further applications appears viable.

20.2 A characterization of the $L^p(\mathbb{R}^{d+1})$-boundedness of directional maximal operators over an arbitrary set of directions

Ultimately, one would like to establish a necessary and sufficient condition for the boundedness of directional maximal operators in general Euclidean space. One would hope that this goal could be recast to completely characterize those direction sets Ω that admit Kakeya-type sets according to Definition 1.1. We aim to combine the result of Theorem 1.3 with others from the literature, most notably [3], [1], [39], to obtain the following necessary and sufficient conditions for Kakeya-type sets to exist.

Conjecture 20.3. *(Kroc, Pramanik) For any $d \geq 1$, the following are equivalent:*

(1) The direction set $\Omega \subseteq \mathbb{S}^d$ is sublacunary in the sense of Definition 2.7.

(2) The set of directions Ω admits Kakeya-type sets in the sense of Definition 1.1.

(3) The maximal operators D_Ω and M_Ω defined in (1.2) and (1.3) respectively are unbounded on $L^p(\mathbb{R}^{d+1})$ for every $p \in (1, \infty)$.

To clarify, the implication (3) $\implies$ (1) for $d = 1$ is in [1], expanding on the work started in [37], [41], [10], [2]. For $d \geq 2$, our notion of lacunarity paired with the result of Parcet and Rogers [39] suggests a possible bridge, although substantial details remain to be resolved to fully connect their work with ours; see below. The proof of (1) $\implies$ (2) is sketched in [3] for $d = 1$. Theorem 1.3 is exactly the implication (1) $\implies$ (2) for all $d \geq 1$. The implication (2) $\implies$ (3) is easy and is established in the argument presented in the paragraph of (1.9) in all dimensions.

Some of the implications above are known to admit stronger variants. For instance, (2) implies (3) even when $p = 1$, as the argument leading to (1.9) shows. Further, it does not appear to be necessary to know that the operator D_Ω is unbounded on all $L^p(\mathbb{R}^{d+1})$, $p \in (1, \infty)$, in order to conclude that Ω is sublacunary. We will sketch a possible path next in Section 20.2.1 toward the weaker requirement

(3') The maximal operator D_Ω is unbounded on $L^p(\mathbb{R}^{d+1})$ for some $p \in (1, \infty)$,

which suffices to establish (1). Thus D_Ω enjoys an interesting dichotomy in that it is either bounded on all or none of the Lebesgue spaces L^p with $p \in (1, \infty)$.

20.2.1 Boundedness of directional maximal operators, sketch

In this subsection, we provide a sketch of the implication (3) $\implies$ (1) of Theorem 20.3, relying heavily on the result of Parcet and Rogers [39]. Let us recall from (1.2) and (1.3) the relevant definitions.

Conjecture 20.4. *Given positive integers N, R, a positive constant $\lambda < 1$ and any exponent $p \in (1, \infty]$, there exists a positive finite constant $C_p = C_p(N, \lambda, R)$ with the following property. Any admissible lacunary direction set $\Omega \subseteq \mathbb{R}^{d+1}$ of finite order that obeys Definition 2.7(i) with the prescribed values of N, λ and R also satisfies*

$$||M_\Omega||_{p\to p} \leq C_p \quad \textit{and} \quad ||D_\Omega||_{p\to p} \leq C_p. \tag{20.5}$$

Proof. (*Sketch*) We first argue that the boundedness of D_Ω on any $L^p(\mathbb{R}^{d+1})$ implies the same for M_Ω. Without loss of generality, we may assume that $\Omega \subseteq (-\epsilon, \epsilon)^d \times \{1\}$ for some small constant $\epsilon > 0$. Let us define for any $x \in \mathbb{R}^{d+1}$ the vectors

$$v_j(x) = x_{d+1}e_j - x_j e_{d+1}, \quad 1 \leq j \leq d,$$

where $\{e_1, \cdots e_{d+1}\}$ denotes the canonical orthonormal basis in $\mathbb{R}^{d+1}$. For $\omega = (\omega_1, \cdots, \omega_d, 1) \in \Omega$, the collection $\{v_1(\omega), \cdots, v_d(\omega)\}$ spans $\omega^\perp$. Then

$$\begin{aligned} M_\Omega f(x) &\leq C_d \sup_{\omega\in\Omega} \sup_{0<r\leq h} \frac{1}{r^d h} \int_{\substack{|t|\leq h\\|s|\leq r}} \big|f\big(x - t\omega - \sum_{j=1}^d s_j v_j(\omega)\big)\big|\, dt\, ds \\ &\leq C_d \sup_{\omega\in\Omega} \sup_{r>0} \frac{1}{r^d} \int D_\Omega f\big(x - \sum_{j=1}^d s_j v_j(\omega)\big)\, ds \\ &\leq C_d \sup_{\omega\in\Omega} \sup_{r>0} \frac{1}{r^{d-1}} \int D_{\Omega_1} \circ D_\Omega f\big(x - \sum_{j=2}^d s_j v_j(\omega)\big)\, ds_2 \cdots, ds_d \\ &\leq \cdots \leq C_d D_{\Omega_d} \circ D_{\Omega_{d-1}} \circ \cdots \circ D_{\Omega_1} \circ D_\Omega f(x), \end{aligned}$$

where $\Omega_j = \{v_j(\omega) : \omega \in \Omega\}$, $1 \le j \le d$. The relation

$$\frac{v_j(\omega) \cdot \xi}{v_j(\omega) \cdot \eta} = \frac{\omega \cdot v_j(\xi)}{\omega \cdot v_j(\eta)} \qquad \text{for all } \xi, \eta \in \mathbb{R}$$

implies that if Ω is admissible lacunary of order at most N as a direction set, then so is Ω_j for every j. Thus a bound on the L^p operator norm of M_Ω would follow if the second conclusion (for directional maximal operators) in (20.5) is known to hold for all such direction sets. We will henceforth concentrate only on D_Ω, with Ω being admissible lacunary of finite order N.

As mentioned before, the L^p-boundedness of D_Ω can be seen to follow from the main result in [39], modulo considerable connecting detail. We sketch these details now and then indicate where more rigour is required.

The proof is by induction on the order of lacunarity N. The initializing step $N = 0$ is a consequence of the one-dimensional Hardy-Littlewood maximal theorem. To set up the induction step, we observe that

$$||D_\Omega||_{p \to p} = ||D_{T(\Omega)}||_{p \to p} \quad \text{ for all } p \in [1, \infty]$$

if $T(\Omega)$ is one of the following two types:

- $T(\Omega) = \{A\omega : \omega \in \Omega\}$ for any nonsingular linear transformation A, or

- $T(\Omega) = \{c_\omega \omega : \omega \in \Omega\}$ for any collection on nonzero scalars $\{c_\omega\}$.

Clearly, Ω and $T(\Omega)$ also have the same order of lacunarity as direction sets. We say that Ω' is a *representative* of Ω if there is a finite sequence of transformations T of the types described above that maps Ω to Ω'.

Set $\Sigma := \{(j, k) : 1 \le j < k \le d + 1\}$. We will prove in Lemma 20.5 below that after a decomposition into at most $C(d, R)$ pieces, it is possible to find a representative Ω' of the direction set Ω obeying certain desirable properties. Specifically, for any $\sigma = (j, k) \in \Sigma$, define

$$\pi_\sigma(\Omega') := \left\{ \frac{\omega_k}{\omega_j} : (\omega_1, \cdots, \omega_{d+1}) \in \Omega' \right\},$$

which is the projection of $\mathcal{C}_{\Omega'} \cap \{\omega_j = 1\}$ onto the kth coordinate axis. Then there exists an integer $0 \leq N_\sigma \leq N$ such that $\pi_\sigma(\Omega') \in \Lambda(N_\sigma; \lambda) \setminus \Lambda(N_\sigma - 1; \lambda)$. Set $\Sigma^* = \Sigma^*(\Omega') := \{\sigma \in \Sigma : N_\sigma \geq 1\}$. In order to apply the main result of [39], we require that for any $\sigma \in \Sigma^*$,

$$\text{the special point of } \pi_\sigma(\Omega'), \text{ as defined in Definition 2.1, is } 0. \tag{20.6}$$

For $\sigma \in \Sigma^*$, let $\{\theta_{\sigma,i} : i \in \mathbb{Z}\}$ denote a monotone decreasing non-negative lacunary sequence with lacunary constant $\leq \lambda$ that converges to 0 and ∞ as $i \to \infty$ and $i \to -\infty$ respectively, and which serves as a special sequence for $\pi_\sigma(\Omega')$. Define

$$\Omega'_{\sigma,i} := \left\{ \omega = (\omega_1, \cdots, \omega_{d+1}) \in \Omega' : \theta_{\sigma,i+1} < \left| \frac{\omega_k}{\omega_j} \right| \leq \theta_{\sigma,i} \right\}. \tag{20.7}$$

This puts us in the framework of [39], where the authors prove that

$$||D_\Omega||_{p\to p} = ||D_{\Omega'}||_{p\to p} \leq C \sup_{\sigma\in\Sigma^*} \sup_{i\geq 1} ||D_{\Omega'_{\sigma,i}}||_{p\to p}.$$

(In fact, [39] addresses the generic and nontrivial case of $\Sigma^* = \Sigma$, but the proof goes through with trivial modifications after a reduction to lower dimensions even when $\Sigma^* \subsetneq \Sigma$). Lemma 20.5 ensures that each set $\Omega'_{\sigma,i}$ is lacunary of order $N - 1$ as a direction set, allowing us to carry the induction forward. □

Lemma 20.5. *Given integers $N, R \geq 1$, and a constant $0 < \lambda < 1$, let Ω_0 be a direction set that is admissible lacunary of finite order, obeying Definition 2.7(i) with these values of N, R, λ. Then there exists a decomposition of Ω_0 into at most $C(d, R)$ subsets of directions $\{\Omega\}$, each of which has the following property: there exists a representative Ω' of Ω such that*

(i) Ω' is contained in a hyperplane at unit distance from the origin,

(ii) the condition (20.6) *holds,*

(iii) for any choice of $\sigma \in \Sigma^*$ and $i \in \mathbb{Z}$, the direction set $\Omega'_{\sigma,i}$ given by (20.7) *is lacunary of order at most $N - 1$.*

Proof. (*Sketch*) We use a generic linear transformation A_0 of $\mathbb{R}^{d+1}$ to fix an ambient coordinate system, in which we represent Ω_0. By decomposing Ω_0 into at most $R^{d(d-1)/2}$pieces $\{\Omega\}$ if necessary, we may assume that $\pi_\sigma(\Omega) \in \Lambda(N;\lambda)$ for every $\sigma \in \Sigma$. We now fix such a direction set Ω, and aim to apply on it a sequence of transformations of the types mentioned in the proof of Theorem 20.2.1, eventually reaching a representative that meets the specified criteria.

After a splitting into $d!$ subsets and permuting coordinates so that $\Omega \subseteq \{|x_1| \geq |x_2| \geq \cdots \geq |x_{d+1}|\}$, we start with the representative $\Omega^{[0]} = \mathcal{C}_\Omega \cap \{x_1 = 1\}$ of Ω. If a_k denotes the special point of $\pi_{(1,k)}(\Omega^{[0]})$, we set $\Omega^{[1]} = T(\Omega^{[0]})$, where T is the linear transformation

$$T(x_1, \cdots, x_{d+1}) = (x_1, x_2 - a_2 x_1, \cdots, x_{d+1} - a_{d+1} x_1).$$

Then $\pi_\sigma(\Omega^{[1]})$ has 0 as its special point for all σ of the form $\sigma = (1,k)$. We set $\Gamma_g^{[1]} = \{1\}$ and $\Gamma_b^{[1]} = \{2, 3, \cdots, (d+1)\}$. These collections represent the subsets of "good" and "bad" coordinates respectively at the end of the first step.

Inductively, we define representatives $\Omega^{[m]}$ of Ω such that for each $1 \leq m \leq d$, the disjoint collections of good and bad indices

$$\Gamma_g^{[m]} = \{1, \cdots, m\} \quad \text{and} \quad \Gamma_b^{[m]} = \{m+1, \cdots, d+1\}$$

have the following significance:

- $\Omega^{[m]} \subseteq \{\omega \in \mathbb{R}^{d+1} : \omega_m = 1\}$.
- For any $\sigma = (j,k) \in \Sigma$ with $j, k \in \Gamma_g^{[m]}$, the special point of the set $\pi_\sigma(\Omega^{[m]})$ is 0.
- For $\sigma = (j,k) \in \Sigma$ with $j \in \Gamma_g^{[m]}$, $k \in \Gamma_b^{[m]}$, the set $\pi_\sigma(\Omega^{[m]})$ has 0 as its special point as well.

Given a representative $\Omega^{[m]}$ with these properties, we permute coordinates so that after a finite decomposition $\Omega^{[m]} \subseteq \{|\omega_{m+1}| \geq \cdots \geq |\omega_{d+1}|\}$. The scaling transformation

$$(\omega_1, \cdots, \omega_m = 1, \cdots, \omega_{d+1}) \in \Omega^{[m]} \mapsto (\zeta_1, \cdots, \zeta_{d+1}) \in \Theta, \text{ where } \zeta_j = \frac{\omega_j}{\omega_{m+1}},$$

leads to the inclusion $\Theta \subseteq \{\zeta_{m+1} = 1\}$. On the other hand, by the induction hypothesis $\pi_\sigma(\Theta) = \pi_\sigma(\Omega^{[m]})$ has 0 as its special point for every $\sigma = (j,k) \in \Sigma$, $1 \le j,k \le m+1$. We next apply the linear transformation $\Omega^{[m+1]} = T(\Theta)$ where $\eta = T(\zeta)$ is given by

$$\eta_k = \begin{cases} \zeta_k & \text{if } k \le m+1, \\ \zeta_k - b_k \zeta_{m+1} & \text{if } k > m+1, \end{cases}$$

and b_k denotes the special point of $\pi_\sigma(\Theta)$ for $\sigma = (m+1,k) \in \Sigma$. For $1 \le j \le m+1 < k \le d+1$, the induction hypothesis again yields that the sets $\{\eta_j^{-1} : (\eta_1, \cdots, \eta_{d+1}) \in \Omega^{[m+1]}\}$ and $\{\eta_k : (\eta_1, \cdots, \eta_{d+1}) \in \Omega^{[m+1]}\}$ both have their special point at zero. The definition of finite order lacunarity of Ω then implies that the same conclusion continues to hold true for $\pi_\sigma(\Omega^{[m+1]})$, with $\sigma = (j,k)$. Thus $\Omega^{[m+1]}$ satisfies the requirements listed above with m replaced by $(m+1)$.

Proceeding in this manner for d steps, we arrive at a representation $\Omega' = \Omega^{[d]}$ of Ω, for which $\Gamma_g^{[d]} = \{1, \cdots, d\}$, and which therefore verifies the requirements (i) and (ii) of the lemma. For $\Omega'_{\sigma,i}$ defined as in (20.7), we claim that $\pi_\sigma(\Omega'_{\sigma,i}) \in \Lambda(N-1;\lambda)$. Let us recall that $\Omega'_{\sigma,i}$ is lacunary of finite order as a direction set (being the subset of a representative of Ω), and that $\pi_\sigma(\Omega'_{\sigma,i})$ is the projection of $\mathcal{C}_{\Omega'_{\sigma,i}} \cap \{\zeta_j = 1\}$ onto the ζ_k-axis. Hence for a generic choice of coordinates based on the selection of A_0 at the beginning of this proof (see also the last remark on page 31), $\Omega'_{\sigma,i}$ is lacunary as a direction set of order at most $N-1$. $\square$

20.2.2 Boundedness of directional maximal operators, a detailed example

The main gap that remains to be filled in the previous subsection's sketch is the claim that the projections onto the coordinate axes of any hyperplane $\{x_j = 1\}$ of the sets $\Omega'_{\sigma,i}$ are lacunary of order strictly less than N. Such a statement is absolutely necessary if we hope to verify boundedness via an inductive argument. In this subsection, we will indicate how to justify such a claim in three dimensions when $N = 2$. This will already indicate where much of the additional work lies.

Suppose Ω is admissible lacunary of order 2 as a direction set in $\mathbb{R}^3$; denote this by $\Omega \in \Delta(2, \lambda; R)$ so that Ω is coverable by R lacunary direction sets of order 2. Furthermore, suppose that $C_\Omega \cap \{x_1 = 1\}$ has the property that $\exists\, a \in \{x_1 = 1\}$ such that if L is a line in $\{x_1 = 1\}$ passing through a, then $\pi_L(C_\Omega \cap \{x_1 = 1\}) = \bigcup_{i=1}^{R} U_i$, with $U_i \in \Lambda(2, \lambda)$ and special point a for all i. Without loss of generality, we may assume $a = (1, 0, 0)$. We will say that such a projection obeying the above property is a member of $\Lambda(2, \lambda; R)$ *relative to the origin*.

Write $C_\Omega \cap \{x_1 = 1\} = \{(1, \omega_2, \omega_3) \in \Omega\}$. Then using the projection notation of the previous subsection, we see that $\pi_\sigma(C_\Omega \cap \{x_1 = 1\}) \in \Lambda(2, \lambda; R)$ relative to the origin for $\sigma = (1, 2)$ and $\sigma = (1, 3)$. In order to apply the result of Parcet and Rogers, we must also be able to conclude the same when $\sigma = (3, 2)$. It is true that

$$\pi_{(3,2)}(C_\Omega \cap \{x_1 = 1\}) = \left\{ \frac{\omega_2}{\omega_3} : (1, \omega_2, \omega_3) \in \Omega \right\} \in \Lambda(2, \lambda; R)$$

since $\Omega \in \Delta(2, \lambda; R)$ by hypothesis, but the lacunarity of this projection is not necessarily relative to a single point. That is, *a priori*, there is nothing to rule out the possibility that the projection is only coverable by R different Euclidean sets, lacunary of order 2, all with potentially different special points.

This requires a decomposition of the original direction set Ω. We have that

$$\left\{ \frac{\omega_2}{\omega_3} : (1, \omega_2, \omega_3) \in \Omega \right\} = \bigcup_{i=1}^{R} V_i,$$

with $V_i \in \Lambda(2,\lambda)$ and c_i the special point of V_i. Split Ω into R pieces Ω_i such that

$$C_{\Omega_i} \cap \{x_1 = 1\} = \left\{(1,\omega_2,\omega_3) \in \Omega_i : \frac{\omega_2}{\omega_3} \in V_i\right\}.$$

Henceforth, we fix i (which fixes V_i and c_i) and rename Ω_i as Ω.

Now we apply a linear transformation T to Ω, $T(\Omega) := \{(1,\omega_2 - c\omega_3,\omega_3) : (1,\omega_2,\omega_3) \in \Omega\}$. By hypothesis, we have $\{\omega_2 - c\omega_3\} \in \Lambda(2,\lambda;R)$. Consequently, $\pi_\sigma(T(\Omega)) \in \Lambda(2,\lambda;R)$ relative to the origin for all $\sigma \in \{(1,2),(1,3),(3,2)\}$.

We are now in a position to apply the result of [39]:

$$||M_\Omega||_{p\to p} = ||M_{T(\Omega)}||_{p\to p} \lesssim \sup_\sigma \sup_{i\geq 1} ||M_{\Omega'_{\sigma,i}}||_{p\to p}, \tag{20.8}$$

where $\Omega'_{\sigma,i}$ is as defined in (20.7). The symmetry of the following argument will allow us to fix $\sigma = (1,2)$ and concentrate on the quantity $\sup_{i\geq 1} ||M_{\Omega'_{\sigma,i}}||_{p\to p}$. We claim the following:

$$\exists\, C = C(R) < \infty \text{ such that } \Omega'_{(1,2),i} \in \Lambda(1,\lambda;C)\ \forall\, i \text{ as a Euclidean set.} \tag{20.9}$$

What we want of course is for the angular sectors $\Omega'_{(1,2),i}$ to actually be admissible lacunary of order 1 as *direction sets*, but it seems that this is too much to hope for after only a single application of the inequality in [39]. However, there is nothing to prevent us from decomposing these sectors further and applying inequality (20.8) again.

Fix $i \geq 1$. Assuming (20.9), we note that both $\{\omega_2'\} \in \Lambda(1,\lambda;C)$ and $\{\omega_3'\} \in \Lambda(1,\lambda;C)$ where $\Omega'_{(1,2),i} = \{(1,\omega_2',\omega_3')\}$. By decomposing $\Omega'_{(1,2),i}$ into at most C^2 many pieces $\{\Omega''_{(1,2),i}\}$, we can ensure that $\{\omega_2'\} \in \Lambda(1,\lambda)$ with special point α_i, and $\{\omega_3'\} \in \Lambda(1,\lambda)$ with special point β_i. Applying another linear transformation, we have the set

$$\Omega'''_{(1,2),i} = \{(1,\omega_2' - \alpha_i,\omega_3' - \beta_i) : (1,\omega_2',\omega_3') \in \Omega''_{(1,2),i}\}.$$

This allows us to conclude that the projections have lower order lacunarity with special points at the origin for $\sigma = (1,2)$ and $\sigma = (1,3)$, but when $\sigma = (3,2)$ we know only that

$$\frac{\omega_2' - \alpha_i}{\omega_3' - \beta_i} \in \Lambda(2, \lambda; R).$$

Decompose $\Omega'''_{(1,2),i}$ into at most R pieces to ensure that $\dfrac{\omega_2' - \alpha_i}{\omega_3' - \beta_i}$ has a unique special point γ_i. After another linear transformation, we define the following set:

$$\widetilde{\Omega}_{(1,2),i} := \{(1, (\omega_2' - \alpha_i) - \gamma_i(\omega_3' - \beta_i), \omega_3' - \beta_i) : (1, \omega_2', \omega_3') \in \Omega''_{(1,2),i}\}. \qquad (20.10)$$

Now by construction, we see that $\pi_\sigma(\widetilde{\Omega}_{(1,2),i}) \in \Lambda(2, \lambda; R)$ relative to the origin for all $\sigma \in \{(1,2), (1,3), (3,2)\}$. Thus, the result of Parcet and Rogers applies and we once again apply inequality (20.8):

$$||M_{\widetilde{\Omega}_{(1,2),i}}||_{p\to p} \lesssim \sup_{\sigma} \sup_{j\geq 1} ||M_{\widetilde{\Omega}_{\sigma,j}}||_{p\to p},$$

where we have suppressed the fixed subscripts $(1,2)$ and i on the right hand side and once again $\widetilde{\Omega}_{\sigma,j}$ is defined as in (20.7).

Now for $\sigma = (1,2)$, the dimension of the direction set $\widetilde{\Omega}_{\sigma,j}$ drops, allowing us to then induct on the dimension. Indeed, writing $\widetilde{\Omega}_{(1,2),j} = \{(1, x_2, x_3)\}$, we see that there can exist no more than R many distinct values of x_2. Of course, there may exist many values of x_3 so that $(1, x_2, x_3) \in \widetilde{\Omega}_{(1,2),j}$ for one particular x_2, but then the dimension of the overall set drops. The same holds if $\sigma = (1,3)$.

For $\sigma = (3,2)$, we claim that since $\{x_2\}, \{x_3\} \in \Lambda(1, \lambda)$ with special point the origin, and $\left\{\dfrac{x_2}{x_3}\right\} \in \Lambda(1, \lambda)$, possibly after a finite decomposition, then $\widetilde{\Omega}_{(3,2),j} \in \Delta(1, \lambda; R)$. This allows us to then induct on the lacunarity order of the direction set.

Clearly, many details remain to be fleshed out, most notably in the claim of the preceding paragraph and of (20.9). Additionally, the simplifying assumption with which we began this example, that the projection of $C_\Omega \cap \{x_1 = 1\}$ onto every line passing through the origin in $\{x_1 = 1\}$ must be coverable by R many lacunary sets of order 2 relative to the origin, must be relaxed. Such a simplification is rather immediate in the base case when $N = 1$, but more work is required to reduce matters in this way for higher order lacunary sets. These details and many others are currently being explored in collaboration with Dimitrios Karslidis and Malabika Pramanik.

Bibliography

[1] A. Alfonseca. *Strong type inequalities and an almost-orthogonality principle for families of maximal operators along directions in* $\mathbb{R}^2$. J. London Math. Soc. (2) **67**, No. 1, 208-218 (2003).

[2] A. Alfonseca, F. Soria, A. Vargas. *A remark on maximal operators along directions in* $\mathbb{R}^2$. Math. Res. Lett. **10**, No. 1, 41-49 (2003).

[3] M. Bateman. *Kakeya sets and directional maximal operators in the plane.* Duke Math. J. **147**, No. 1, 55-77 (2009).

[4] M. Bateman, N.H. Katz. *Kakeya sets in Cantor directions.* Math. Res. Lett. **15**, No. 1, 73-81 (2008).

[5] A.S. Besicovitch. *Sur Deux Questions de l'Intgrabilit des fonctions.* J. Société des Math. et de Phys. à l'Univ. à Perm, **2**, 105-123 (1919).

[6] A.S. Besicovitch. *On Kakeya's problem and a similar one.* Mat. Zeitschrift, **27**, No. 1, 312-320 (1928).

[7] A.S. Besicovitch. *On fundamental geometric properties of plane line-sets.* J. London Math. Soc., **39**, 441-448 (1964).

[8] A.S. Besicovitch, R. Rado. *A plane set of measure zero containing circumferences of every radius.* J. London Math. Soc. **43**, 717-719 (1968).

[9] J. Bourgain. *Besicovitch type maximal operators and applications to Fourier analysis.* Geometric and Functional Analysis **1**, 147-187 (1991).

[10] A. Carbery. *Differentiation in lacunary directions and an extension of the Marcinkiewicz multiplier theorem.* Ann. Inst. Four. **38**, No. 1, 157-168 (1988).

[11] A. Córdoba. *The Kakeya maximal function and spherical summation multipliers.* Amer. J. Math. **99**, No. 1, 1-22 (1977).

[12] A. Córdoba, R. Fefferman. *On differentiation of integrals.* Proc. Nat. Acad. Sci. US **74**, 2211-2213 (1977).

[13] F. Cunningham. *The Kakeya problem for simply connected and for star-shaped sets.* Amer. Math. Monthly **78**, 114-129 (1971).

[14] R.O. Davies. *Some remarks on the Kakeya problem.* Proc. Cambridge Phil. Soc. **69**, 417-421 (1971)

[15] J. Duoandikoetxea, A. Vargas. *Directional operators and radial functions on the plane.* Ark. Mat. **33**, 281-291 (1995).

[16] K. Falconer. *Fractal Geometry: mathematical foundations and applications*, 2nd edition. John Wiley & Sons (2003).

[17] C. Fefferman. *The multiplier problem for the ball.* Ann. Math. (2) **94**, 330-336 (1971).

[18] M. Fujiwara, S. Kakeya. *On some problems of maxima and minima for the curve of constant breadth and the in-revolvable curve of the equilateral triangle.* Tôhoku Math. Journal, **11**, 92-110 (1917).

[19] L. Grafakos. *Modern Fourier analysis.* Springer, 2008.

[20] G. Grimmett. *Percolation*, 2nd edition. Grundlehren der math. Wissenschaften, Vol. 321 (1999).

[21] G.H. Hardy, J.E. Littlewood. *A maximal theorem with function-theoretic applications.* Acta Math. **54**, 81-116 (1930).

[22] K.E. Hare. *Maximal operators and Cantor sets.* Canad. Math. Bul. **43**, 330-342 (2000).

[23] J.P. Kahane *Trois notes sur les ensembles parfaites lineaires.* L'enseignement Mathematique Revue Internationale, **15**, 185-192 (1969).

[24] S. Kakeya. *Some problems on maximum and minimum regarding ovals.* Tôhoku Sci. Reports, **6**, 71-88 (1917).

[25] N.H. Katz. *A counterexample for maximal operators over a Cantor set of directions.* Mat. Res. Lett. **3**, 527-536 (1996).

[26] N.H. Katz. *Remarks on maximal operators over arbitrary sets of directions.* Bul. London Math. Soc. **31**, 700-710 (1999).

[27] N.H. Katz, I. Łaba, T. Tao. *An improved bound on the Minkowski dimension of Besicovitch sets in* $\mathbb{R}^3$. Ann. of Math. (2) **152**, 383-446 (2000).

[28] N.H. Katz, T. Tao. *New bounds for Kakeya problems.* J. Anal. Mat. **87**, 231-263 (2002).

[29] J.R. Kinney. *A thin set of circles.* Amer. Math. Monthly **75**, 1077-1081 (1968).

[30] L. Kolasa, T. Wolff. *On some variants of the Kakeya problem.* Pacific J. Math., **190**, 111-154 (1999).

[31] E. Kroc, M. Pramanik. *Kakeya-type sets over Cantor sets of directions in* $\mathbb{R}^{d+1}$. [preprint available at http://arxiv.org/abs/1404.6235]

[32] E. Kroc, M. Pramanik. *Lacunarity, Kakeya-type sets and directional maximal operators.* [preprint available at http://arxiv.org/abs/1404.6241]

[33] I. Łaba. *From harmonic analysis to arithmetic combinatorics.* Bull. Amer. Math. Soc., **45**, 77-115 (2008).

[34] R. Lyons. *Random walks and percolation on trees.* Ann. Prob., **18**, 931-958 (1990).

[35] R. Lyons. *Random walks, capacity, and percolation on trees.* Ann. of Prob. **20**, (1992).

[36] R. Lyons, Y. Peres. *Probability on Trees and Networks.* Cambridge Univ. Press, mypage.iu.edu/∼rdlyons/prbtree/prbtree.html (in preparation).

[37] A. Nagel, E.M. Stein, S. Wainger. *Differentiation in lacunary directions.* Proc. Natl. Acad. Sci. US, **75** (3), 1060-1062 (1978).

[38] J. Pál. *Ein Minimumproblem fr Ovale.* Math. Ann., **83**, 311-319 (1921).

[39] J. Parcet, K.M. Rogers. *Differentiation of integrals in higher dimensions.* Proc. Natl. Acad. Sci. US, **110** (13), 4941-4944 (2013).

[40] H. Pecher. *Nonlinear small data scattering for the wave and Klein-Gordon equation.* Math. Zeit., **185**, 261-270 (1984).

[41] P. Sjögren, P. Sjölin. *Littlewood-Paley decompositions and Fourier multipliers with singularities on certain sets.* Ann. Inst. Four., **31**, 157-175 (1981).

[42] E. Stein. *Harmonic Analysis: real-variable methods, orthogonality and oscillatory integrals.* Princeton mathematical series; 43. Princeton University Press, 1995.

[43] J. Stromberg. *Maximal functions for rectangles with given directions.* Thesis, Mittag-Leffler Inst., Sweden (1976).

[44] A. Vargas. *A remark on a maximal function over a Cantor set of directions.* Rend. Circ. Mat. Palermo **44**, 273-282 (1995).

[45] N. Wiener. *The ergodic theorem.* Duke Math. J. **5**, 1-18 (1939).

[46] T. Wolff. *An improved bound for Kakeya type maximal functions.* Rev. Math. Iberoamericana **11**, No. 3, 651-674 (1995).