Incidence Problems for Slabs

by

Sarah Tammen

Submitted to the Department of Mathematics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Mathematics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2023

Author ....
Department of Mathematics
July 11, 2023

Certified by ....
Larry Guth
Claude Shannon Professor of Mathematics
Thesis Supervisor

Accepted by ....
Davesh Maulik
Chair, Department Committee on Graduate Studies

Incidence Problems for Slabs

by

Sarah Tammen

Submitted to the Department of Mathematics
on July 11, 2023, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Mathematics

Abstract

In this thesis, I prove incidence estimates for slabs which are formed by intersecting small neighborhoods of well-spaced hyperplanes in $\mathbb{R}^d$ with the unit cube $[0,1]^d$. My work is an analogue of a theorem of Guth, Solomon, and Wang, who proved a version of the Szemerédi-Trotter theorem for thin tubes that satisfy a certain strong spacing condition. My proof uses induction on scales and the high-low method of Vinh, along with new geometric insights.

Thesis Supervisor: Larry Guth
Title: Claude Shannon Professor of Mathematics

This thesis is dedicated to the memory of my grandfather, Ralph Robert Stevens Jr.

Acknowledgments

First and foremost, I would like to thank my advisor, Larry Guth, for his patience and kindness. Larry, thank you for being a great teacher and mentor. I am so blessed that I have experienced firsthand the way you think about math, the ways you approach teaching and writing, and your constant striving to make beautiful ideas accessible to many people. I am continually impressed by the way you bring together ideas from different branches of mathematics in a new way, and I really appreciate your sharing these connections.

I would like to thank Gigliola Staffilani and Yufei Zhao for being on my thesis committee as well as being on my qualifying exam committee early in my time at MIT. I was fortunate to take classes with Gigliola and Yufei and to meet frequently with Gigliola during my first few years here. Gigliola, thanks being so generous with your time and ideas, and thank you both for teaching me about some really interesting topics.

I have also had the privilege of discussing and learning math with many other analysis graduate students at MIT, including Lingxian Zhang, Shengwen Gan, Yuqiu Fu, Alex Ortiz and Alex Cohen. I owe another thank you to Larry for advising all of us - simultaneously! (It often seems to me that Larry must be superhuman.) I also want to thank Hong Wang, Dominique Maldague, Changkeun Oh, and Marjorie Drake for introducing me and the rest of the analysis group to interesting problems from your research. I am so glad I have had the chance to learn with and from all of you.

I am immensely grateful for the continued support of my mathematical family at UGA and for many other friends in Athens who have become like family to me. Betsy and Kevin Weinrich deserve special mention, as do Beth and Randy Beck. Thank you all for helping me during a difficult time and for always believing in me, even and especially when I did not believe in myself.

I am also immensely grateful to my family for their love and support. I am grateful that my parents have always nurtured my love of learning and have always taken such a great interest in my many activities. I am also especially grateful for the friendship of my sister Sylvia, which has only grown stronger during the time that I have been in graduate school and she has been in college. I will be defending my thesis one year after Sylvia played her senior recital for her piano performance degree. It is very meaningful to me that some of the work that appears in this thesis comes from ideas I was thinking about on my trip to

Georgia for her recital. Her recital program has also comprised part of my writing playlist for the sessions in which I've been writing up my results. Sylvia, thank you for sharing your great love of learning and music, for taking an avid interest in all parts of my ever-developing life story, and for being someone I can talk to about anything.

Finally, I am thankful to God, the Author of Life, for orchestrating a beautiful life story for me - one that has been filled with math, music, and special moments with my friends and family. I am eager to see what the next chapter brings.

Contents

Chapter 1

Introduction and Preliminaries

In this thesis, we consider a fixed subdivision of the unit box $[0,1]^d \subset \mathbb{R}^d$ into grid boxes of side length $\delta < 1$, and we consider incidences between a set of these grid boxes and a set of δ-slabs. Here, a δ-slab refers to a(n) (approximate) box $S \subset \mathbb{R}^d$ with dimensions $\sim 1 \times \cdots \times 1 \times \delta$. When we say a length is $\sim x$, we mean that it is at least $c_d\, x$ and at most $C_d\, x$ for some fixed dimensional constants $c_d < 1$ and $C_d > 1$. When we say S is an 'approximate box,' we are referring to the fact that S contains a dilate by c_d of a $1 \times \cdots \times 1 \times \delta$ box and is contained in a dilate by C_d of $1 \times \cdots \times 1 \times \delta$ box, but S may not itself be a box.

If a δ-box intersects a δ-slab in a set of volume at least $\delta^d/10$, we say that the box is essentially contained in the slab. (The requirement that the intersection has volume at least $\delta^d/10$ means that the intersection must account for at least one tenth of the volume of the δ-box.) If Γ is a collection of δ-slabs and q is a grid box that is essentially contained in r different slabs, we say that q is r-rich for Γ. We let $P_r(\Gamma)$ denote the set of r-rich δ-boxes for Γ.

The question of what if means for two slabs to be 'different' from each other is an important question that motivates us to restrict our attention to collections of slabs that satisfy certain strong spacing conditions. First of all, we require that any two δ-slabs with a substantial intersection must have normal vectors that are separated by an angle $\gtrsim \delta$ (cf. Corollary 1.1.4). It turns out that this requirement alone is not enough to prove strong bounds for the size of $P_r(\Gamma)$. We consider two other options for a further constraint on the collection Γ. Each of these constraints has to do with the distribution of the δ-slabs within larger slabs of dimensions $\sim 1 \times \cdots \times 1 \times W^{-1}$ for a parameter W with $1 \leq W \leq \delta^{-1}$.

Under our first spacing condition, we allow for thick slabs of thickness $\sim W^{-1}$ whose normal directions are δ-separated. We also allow a W^{-1}-slab R to contain more than one slab of Γ but require that the number of δ-slabs inside R that share a normal direction with R is (approximately) the same for each R. (See Definition 1.4.1.) The second spacing condition requires the thick slabs under consideration to have normal directions that are W^{-1}-separated and requires that each thick slab contain at most one δ-slab from our collection Γ.

Under the first spacing condition, we prove a bound for $|P_r(\Gamma)|$ that depends on the parameter W. Specifically, we show that if Γ satisfies the first spacing condition and r is sufficiently large (see Theorem 1.4.2), then

$$|P_r(\Gamma)| \lesssim_\varepsilon \delta^{-\varepsilon}|\Gamma|^d r^{-d} W^{-(d-1)}. \tag{1.1}$$

That is, for every $\varepsilon > 0$, there is a constant C_ε so that if Γ is a collection of δ-slabs in $[0,1]^d$ that satisfies our yet-to-be stated condition and r is sufficiently large, then

$$|P_r(\Gamma)| \leq C_\varepsilon \delta^{-\varepsilon}|\Gamma|^d r^{-d} W^{-(d-1)}.$$

Under the second spacing condition, we consider circumstances under which we can prove an estimate of the form

$$|P_r(\Gamma)| \lesssim_\varepsilon \delta^{-\varepsilon}\frac{|\Gamma|^d}{r^{d+1}}. \tag{1.2}$$

My Theorem 1.4.2 is an analogue of Theorem 1.2 from [10], which discusses how to estimate the number of r-rich δ-boxes for a set of well-spaced tubes. My work uses methods developed by Guth, Solomon, and Wang in [10], appropriately adapted for slabs. A bound of the form (1.2) would be an analogue of a result of Elekes and Tóth [5], who estimate point-hyperplane incidences in $\mathbb{R}^d$. In the rest of this chapter, we give a survey of the results in [10] and [5], and we discuss some properties of δ-slabs. Finally, we give precise formulations of the spacing conditions in order to state a theorem about incidences under the first spacing conditions and a conjecture about incidences under the second.

1.1 Definitions, Notation, and Conventions

1.1.1 Notation in Inequalities and Chains of Inequalities, Part I

Given two positive real-valued functions f and g, we write

$$f \lesssim g$$

to refer to the statement "$f \leq Cg$ for some positive constant C."

If we add a subscript to the $\lesssim$ symbol, this indicates that the implicit constant C depends on the parameter in the subscript. If the implicit constant depends only d, the ambient dimension, we will often omit the d subscript, as we did in (1.1) above.

Throughout this thesis, if f and g are two positive real-valued functions on $(0,1)$, then an inequality of the form

$$f(\delta) \lessapprox_{\log} g(\delta)$$

will be shorthand for the statement, "There is a positive constant C so that

$$f(\delta) \lesssim (\log(\delta^{-1}))^C g(\delta) \tag{1.3}$$

for all δ in $\left(0, \frac{1}{e}\right)$.

This is to be contrasted with the usage of the symbol $\lessapprox$, which indicates a relationship of the form

$$f(\delta) \leq \delta^{-\varepsilon} g(\delta). \tag{1.4}$$

or

$$f(\delta) \leq C_\varepsilon \delta^{-\varepsilon} g(\delta). \tag{1.5}$$

One should note that if f and g satisfy (1.3) then they satisfy (1.4) for all δ sufficiently small; specifically if (1.3) holds for all $\delta \in (0, \frac{1}{e})$, then we can find $c_\varepsilon < 1$ so that (1.4) holds for all $\delta \leq c_\varepsilon$. Then, provided that f and g have sufficient regularity, we can choose a constant C_ε to ensure that (1.5) holds for all $\delta \in (0,1)$.

In [10] the $\lessapprox$ symbol is occasionally used when a stronger statement with $\lessapprox_{\log}$ actually holds, but then the stronger meaning is used later in the paper. Any results I state in this thesis that were originally proved in other sources, will always be stated in the form that I use them.

1.1.2 Notation in Inequalities and Chains of Inequalities, Part II

Each of the inequality symbols above has a counterpart. We write

$$f \gtrsim g$$

to refer to a relationship of the form $g \lesssim f$. If both $f \lesssim g$ and $g \lesssim f$ hold, we write $g \sim f$.

If f, g are positive real-valued functions on $(0, 1)$, then the inequality $f \gtrapprox_{\log} g$ refers to a relationship of the form $g \lessapprox_{\log} f$. That is, we write $f \gtrapprox_{\log} g$ to indicate that there is some positive constant C so that for all $\delta \in \left(0, \frac{1}{e}\right)$, we have that

$$g(\delta) \lesssim (\log(\delta^{-1}))^C f(\delta),$$

or equivalently, that

$$f(\delta) \gtrsim (\log(\delta^{-1}))^{-C} g(\delta).$$

Consequently, if $f \lessapprox_{\log} g$, we have that $\frac{1}{f} \gtrapprox_{\log} \frac{1}{g}$, as one might expect. This fact is frequently used throughout our work with no justification inline.

Finally, we comment on transitivity: if $f \lessapprox_{\log} g$ and $g \lessapprox_{\log} h$, then $f \lessapprox_{\log} h$. If we have functions $f_1, \ldots, f_m$ on $(0, 1)$ with $f_1 \lessapprox_{\log} f_2$, $f_2 \lessapprox_{\log} f_3$, and so on, then we can conclude that $f_1 \lessapprox_{\log} f_m$, provided that the value of m did not depend on δ.

1.1.3 Other Notational Conventions

Throughout this thesis, we use vertical bars to denote the cardinality of a finite set or the d-dimensional Lebesgue measure of an infinite set. For $k \leq d$, we will let $\text{Vol}^k(\cdot)$ denote k-dimensional volume of a set in $\mathbb{R}^d$. If the superscript k is omitted, it is to be assumed that $k = d$. (Thus, if $A \subset \mathbb{R}^d$ is infinite, then $|A|$ and $\text{Vol}(A)$ are used interchangeably.)

If $A \subset \mathbb{R}^d$ is a convex symmetric region with center $\boldsymbol{c}_A$ and $b > 0$, we let bA denote the dilation of A by a factor of b about the same center. That is, $bA = \{b(\boldsymbol{x} - c_A) + c_A : \boldsymbol{x} \in A\}$.

1.1.4 Geometric Facts, Essential Distinctness, and Essential Containment

We can associate to any δ-slab S a normal vector $\boldsymbol{n}$. If S is a true rectangular box, then we can view S as a neighborhood of width $\sim \delta$ of a rectangle R which is the image of the $(d-1)$-dimensional unit box $[0,1]^{d-1} \times \{0\}$ under an affine transformation $A : \mathbb{R}^d \to \mathbb{R}^d$. If we write $A(x) = Tx + b$ for some matrix $T \in GL(d, \mathbb{R})$, we must have $|\det(T)| \sim 1$. Although the determinant of T may not be precisely 1, we will always assume that the last column is of length 1 so that it lies on the surface of the unit sphere $\mathbb{S}^{d-1}$.

Strictly speaking, the normal vector of a slab is not uniquely defined, as a single slab could be contained within a small neighborhood of two different image sets $A_1([0,1]^{d-1} \times 0)$ and $A_2([0,1]^{d-1} \times 0)$. However, all of the candidate normal vectors for a δ-slab lie within a **cap** of radius $\lesssim \delta$, where a cap is defined as the intersection of the sphere $\mathbb{S}^{d-1}$ with a ball around a point x on the sphere, or more generally, as a subset of $\mathbb{S}^{d-1}$ which contains and is contained in such an intersection. More precisely, given $0 < \rho \ll 1$, we call a set $\theta \subset \mathbb{S}^{d-1}$ a ρ-cap if there is some $x \in \mathbb{S}^{d-1}$ so that

$$B(x, \tfrac{1}{10}\rho) \cap \mathbb{S}^{d-1} \subset \theta \subset B(x, 10\rho) \cap \mathbb{S}^{d-1}.$$

We will also speak of **subdividing** the unit sphere into caps, by which we mean writing it as a union

$$\mathbb{S}^{d-1} = \bigcup \theta,$$

where the sets θ are disjoint or they are finitely overlapping with $O(1)$-many sets intersecting any point on the sphere.

We are often concerned with the volume of the intersection of a pair of slabs or of the d-fold intersection of slabs $S_1, \ldots, S_d$. We can express both of these quantities in terms of (multilinear expressions of) the normal vectors of the slabs. Our intersection formulae are not exact equalities but allow for multiplication by a constant factor which accounts for the non-uniqueness of the normal vector of a slab.

Lemma 1.1.1. *Suppose that $S_1, S_2 \subset \mathbb{R}^d$ are δ-slabs with respective normal vectors $\boldsymbol{n_1}$ and $\boldsymbol{n_2}$. If the angle between $\boldsymbol{n_1}$ and $\boldsymbol{n_2}$ is at least α then the volume of the intersection $S_1 \cap S_2$*

satisfies

$$|S_1 \cap S_2| \lesssim \frac{\delta^2}{\alpha}.$$

Moreover, if S_1 and S_2 were formed by taking the δ-neighborhoods of rectangles centered at a common point (e.g. $\mathbf{0}$), then

$$|S_1 \cap S_2| \sim \frac{\delta^2}{\alpha}.$$

Lemma 1.1.2. *Suppose that $S_1, \ldots, S_d \subset \mathbb{R}^d$ are δ-slabs with respective normal vectors $\boldsymbol{n_1}, \ldots, \boldsymbol{n_d}$, Then the volume of their d-fold intersection satisfies*

$$|S_1 \cap \cdots \cap S_d| \lesssim \min\left\{\frac{\delta^d}{|\boldsymbol{n_1} \wedge \cdots \wedge \boldsymbol{n_d}|}, \delta\right\},$$

where $\boldsymbol{n_1} \wedge \cdots \wedge \boldsymbol{n_d}$ denotes the determinant of the $d \times d$ matrix that has $\boldsymbol{n_1}, \ldots, \boldsymbol{n_d}$ as its rows. Moreover, if $S_1, \ldots, S_d$ were formed by taking the δ-neighborhoods of rectangles centered at a common point (e.g. $\mathbf{0}$) and $\frac{\delta^d}{|\boldsymbol{n_1} \wedge \cdots \wedge \boldsymbol{n_d}|} \leq \delta$, then

$$|S_1 \cap \cdots \cap S_d| \sim \frac{\delta^d}{|\boldsymbol{n_1} \wedge \cdots \wedge \boldsymbol{n_d}|}.$$

We defer the proofs of Lemmas 1.1.1 and 1.1.2 to an appendix. Meanwhile, from Lemma 1.1.1 we can deduce the following corollary.

Corollary 1.1.3. *If S_1, S_2 are essentially distinct δ-slabs through $\mathbf{0}$ with respective normal vectors $\boldsymbol{n_1}$ and $\boldsymbol{n_2}$, then the angle between $\boldsymbol{n_1}$ and $\boldsymbol{n_2}$ is $\gtrsim \delta$.*

If Γ is a set of essentially distinct slabs through a common point, then their normal vectors must form a δ-separated set on $\mathbb{S}^{d-1}$, i.e. it must be the case that the distance between any two of the normal vectors must be $\gtrsim \delta$. As a result, we can deduce an upper bound on the size of Γ.

Corollary 1.1.4. *Let Γ_0 be a set of essentially distinct δ-slabs in $\mathbb{R}^d$, which are all centered at the origin. Then*

$$|\Gamma_0| \lesssim \delta^{-(d-1)}.$$

Proof. We consider the normal vectors of the slabs in Γ_0. Corollary 1.1.3 implies that if we subdivide the unit sphere $\mathbb{S}^{d-1}$ into caps of radius $\sim \delta$, then at most one of these normal vectors lies in each cap. The maximum number of δ-caps in such a subdivision is bounded by the quotient

$$\frac{\mathrm{Vol}^{d-1}(\mathbb{S}^{d-1})}{\mathrm{Vol}^{d-1}(\delta\text{-cap})},$$

which is $\sim \delta^{-(d-1)}$. □

Finally, we record for reference our definitions of essential intersection and essential containment, and we introduce the symbols we use to denote these relationships.

Definition 1.1.5. *(Essential intersection) Given two sets $A_1, A_2 \subset \mathbb{R}^d$, we say that A_1 essentially intersects A_2, denoted $A_1 \cap_{ess} A_2$ if the volume of $A_1 \cap A_2$ is at least half as big as the maximum intersection of rigid transformations of A_1 and A_2. That is, $A_1 \cap_{ess} A_2$ if*

$$|A_1 \cap A_2| \geq \frac{1}{2} \max_{\sigma_1, \sigma_2 \in SL_d(\mathbb{R}^d)} \left(|\sigma_1(A_1) \cap \sigma_2(A_2)| \right).$$

Definition 1.1.6. *(Essential containment) Suppose that $A_1, A_2 \subset \mathbb{R}^d$ with $|A_2| \geq |A_1|$. We say that A_1 is essentially contained in A_2, denoted $A_1 \subsetapprox A_2$ if*

$$|A_1 \cap A_2| \geq \frac{1}{10}|A_1|.$$

1.1.5 Incidences and Richness

If P is a set of points in $\mathbb{R}^d$ and $\mathcal{A}$ is a collection of measurable sets in $\mathbb{R}^d$, we let $I(P, \mathcal{A})$ denote the set of **incidences** between P and $\mathcal{A}$, defined by

$$I(P, \mathcal{A}) := \{(p, A) : p \in A\}.$$

If we instead let P denote a set of δ-boxes in $\mathbb{R}^d$ and $\mathcal{A}$ is a collection of measurable sets of positive volume, we define the set of incidences between P and $\mathcal{A}$ in terms of essential intersection:

$$I(P,\mathcal{A}) := \{(p,A) : p \cap_{\text{ess}} A\}. \tag{1.6}$$

If $\mathcal{L}$ is a collection of lines in $\mathbb{R}^d$ and r is a positive integer, we say that a point $p \in \mathbb{R}^d$ is r-rich for a collection of lines if it lies on at least r of the lines. We denote the collection of such points by $P_r(\mathcal{L})$. Similarly, if $\mathcal{A}$ is a collection of subsets of $\mathbb{R}^d$, each of volume $\gtrsim \delta^d$, we say that a δ-ball $p \subset \mathbb{R}^d$ is r-rich for $\mathcal{A}$ if it is essentially contained in at least r of the sets of $\mathcal{A}$. Given a fixed subdivision of $[0,1]^d$ into δ-cubes and a collection $\mathcal{A}$ of subsets of $[0,1]^d$, we let $P_r(\mathcal{A})$ denote the set of these δ-cubes which lie in at least r-many sets of $\mathcal{A}$. It is important to note that $P_r(\mathcal{A})$ does not just include δ-boxes which are essentially contained in exactly r-many sets $A \in \mathcal{A}$, but also includes those δ-boxes of P that have richness strictly greater than r.

1.2 Discrete Incidence Theorems

The fundamental result in Incidence Geometry is the Szeméredi-Trotter theorem, originally proved in [12], which concerns incidences of points and lines in $\mathbb{R}^2$. The theorem has two variants, which are equivalent to each other.

Theorem 1.2.1. *If $P \subset \mathbb{R}^2$ is a set of n points in the plane, and $\mathcal{L}$ is a set of m lines in the plane, the the number of incidences between P and $\mathcal{L}$ satisfies*

$$|I(P,\mathcal{L})| \lesssim n^{2/3}m^{2/3} + n + m.$$

Theorem 1.2.2. *Let $\mathcal{L}$ be a set of lines in the plane with $|\mathcal{L}| = m$. For any $r \in \mathbb{N}$, the number of r-rich points for $\mathcal{L}$ satisfies*

$$|P_r(\mathcal{L})| \lesssim \frac{m^2}{r^3} + \frac{m}{r}.$$

For a discussion of why each variant implies the other, see, e.g. Adam Sheffer's excellent blog post [11].

If $d > 2$, the Szeméredi-Trotter theorem - either variant - holds for collections of points and lines in $\mathbb{R}^d$ but stronger conclusions are possible if one assumes more hypotheses to rule out situations in which many of the lines are concentrated in a hyperplane.

In [5], Elekes and Tóth consider incidences of points and hyperplanes in $\mathbb{R}^d$ under similar non-degeneracy hypotheses. Given $\alpha \in (0,1)$ and a set of points $P \subset \mathbb{R}^d$, $d \geq 3$, they say that a hyperplane H is **α-degenerate** (for P) if $H \cap P$ is non-empty and at most $\alpha|H \cap P|$ points of $H \cap P$ lie in any $(d-2)$-flat within H. Under this definition, if $\alpha_1 < \alpha_2$ and H is α_1-degenerate for P, then it is α_2-degenerate for P. Thus, if we say a hyperplane H is α-degenerate for P, this means that H is at worst α-degenerate.

Under this non-degeneracy condition, Elekes and Tóth prove the following theorem.

Theorem 1.2.3. *(Elekes-Tóth, [5]) For every $d \geq 3$, there is a constant $\alpha_d \in (0,1)$ such if P is a set of n points in $\mathbb{R}^d$ and Γ is a set of m α-degenerate hyperplanes, $\alpha < \alpha_d$, then*

$$|I(P,\Gamma)| \lesssim n^{\frac{d}{d+1}} m^{\frac{d}{d+1}} + nm^{\frac{d-2}{d-1}} + m. \tag{1.7}$$

1.3 Continuous Incidence Theorems for Tubes

In [10], Guth, Solomon and Wang proved an analogue of the Szeméredi-Trotter theorem for δ-tubes that obey certain strong conditions.

Theorem 1.3.1. *([10], Theorem 1.1)*

Suppose that $1 \leq W \leq \delta^{-1}$. Suppose that $\mathbb{T}$ is a set of $\sim W^2$ δ-tubes in $[0,1]^2$ with at most one δ-tube of $\mathbb{T}$ in each $W^{-1} \times 1$ rectangle.

If $r > \max(\delta^{1-\varepsilon}|\mathbb{T}|, 1)$,

then $|P_r(\mathbb{T})| \leq C(\varepsilon)\delta^{-\varepsilon} r^{-3} |\mathbb{T}|^2$

for all $\varepsilon > 0$. Here, $C(\varepsilon)$ is a constant only depending on ε, in particular, independent of W.

We note that the exponents from Theorem 1.3.1 match the exponents in the first term of the bound in Theorem 1.2.2.

Guth, Solomon, and Wang proved an analogous bound for tubes in $\mathbb{R}^3$, which they then applied to solve a variant of Falconer's distance problem using the Elekes-Sharir framework introduced in [4].

Theorem 1.3.2. *([10], Theorem 1.3)*

Suppose that $1 \leq W \leq \delta^{-1}$. *Suppose that* $\mathbb{T}$ *is a set of* $\sim W^4$ δ*-tubes in* $[0,1]^3$ *with at most one* δ*-tube of* $\mathbb{T}$ *in any tube of radius* W^{-1} *and length* 1.

If $r > \max(\delta^{2-\varepsilon}|\mathbb{T}|, 1)$,

then $|P_r(\mathbb{T})| \leq C(\varepsilon)\delta^{-\varepsilon}r^{-2}|\mathbb{T}|^{3/2}$

for all $\varepsilon > 0$. *Here,* $C(\varepsilon)$ *is a constant only depending on* ε, *in particular, independent of* W.

The theorems for $d = 2$ and $d = 3$ can be combined into a more general theorem that the authors stated for tubes in $\mathbb{R}^d$, with a bound that depends on d. They prove that the bound holds for $d = 2$ and $d = 3$, but it is conjectured that the bound holds for larger values of d as well.

Theorem 1.3.3. *([10], Theorem 4.1)*

Let $1 \leq W \leq \delta^{-1}$. *Let* $\mathbb{T}$ *be a collection of distinct* δ*-tubes in* $B^d(0,2)$, *for* $d = 2$ *or* $d = 3$. *If* $\mathbb{T}$ *is a set of* $\sim W^{2(d-1)}$ δ*-tubes with at most one tube of* $\mathbb{T}$ *in each* $\frac{1}{W}$*-tube, then for* $r > \max\left(\delta^{d-1-\varepsilon/4}|\mathbb{T}|, 1\right)$, *the number of* r*-rich* δ*-balls is bounded by*

$$|P_r(\mathbb{T})| \lesssim_\varepsilon \delta^{-\varepsilon}\frac{|\mathbb{T}|^{\frac{d}{d-1}}}{r^{\frac{d+1}{d-1}}}.$$

All of the theorems from [10] that we have stated so far satisfy a spacing condition involving an auxiliary collection of essentially distinct W^{-1}-tubes, where $0 < \delta \leq W^{-1}$. Guth, Solomon and Wang require $\mathbb{T}$ to have $\sim W^d$ tubes, with at most one tube in each W^{-1}-tube from this auxiliary collection. Any two wide tubes in the auxiliary collection that have a substantial intersection must have long axes separated by an angle $\gtrsim W$ because of our assumption that they are pairwise essentially distinct.

In another theorem from the same paper, the authors consider W^{-1}-tubes in $[0,1]^2$ whose axis directions come from a set which is δ-separated but not necessarily W^{-1}-separated. The W^{-1} tubes with a common direction are distinct from one another. We can think of the set of directions as a collection of arcs on the circle $\mathbb{S}^1$. The theorem estimates the number of r-rich points for a collection $\mathbb{T}$ of δ-tubes that is distributed so that any two W^{-1}-tubes whose axis directions are in the same arc θ must contain about the same number of tubes in $\mathbb{T}_\theta$, where $\mathbb{T}_\theta$ consists of those tubes of $\mathbb{T}$ with axis direction in the arc θ; moreover, this common count must be the same size for each θ.

Theorem 1.3.4. *([10], Theorem 1.2)*

Let $1 \leq W \leq \delta^{-1}$, and $1 \leq N_1 \leq (W\delta)^{-1}$. Divide the circle into arcs θ of length δ. For each θ and each $1 \leq j \leq W$, let $T_{\theta,j} \subset [0,1]^2$ be a δ-tube. Suppose that for each θ and each W^{-1}-rectangle in direction θ, there are uniformly $\sim N_1$ tubes $T_{\theta,j}$ in the rectangle. Let $\mathbb{T}$ be the set of all the tubes $T_{\theta,j}$. Then for any $\varepsilon > 0$,

if $r \geq C_1(\varepsilon)\delta^{1-\varepsilon}|\mathbb{T}|$,

then $|P_r(\mathbb{T})| \leq C_2(\varepsilon)\delta^{-\varepsilon}W^{-1}r^{-2}|\mathbb{T}|^2$,

where $C_1(\varepsilon)$ and $C_2(\varepsilon)$ are two constants only depending on ε, in particular, independent of W and N_1.

It is this theorem that I generalize in my Theorem 1.4.2, stated below. Before stating my theorem, I introduce spacing conditions for slabs which I have formulated by analogy with the hypotheses of the theorems of [10].

1.4 My Work

I consider collections of slabs δ-slabs under two spacing conditions, each of which concerns the δ-slabs contained inside a larger slab of thickness $\sim W^{-1}$ with $1 \leq W \leq \delta^{-1}$. Under one spacing condition, we are only concerned with the number of δ-slabs inside a W^{-1}-slab that have the same normal direction as the W^{-1}-slab. Other the other spacing condition, we count *all* of the δ-slabs inside a W^{-1}-slab. It turns out, though, that saying we count the slabs, plural, is a bit misleading as the second spacing condition amounts to the requirement that this count is at most 1 for any W^{-1}-slab.

Definition 1.4.1. *(Spacing Conditions)*

Fix δ, W with $0 < \delta \leq W^{-1} \leq 1$. Let Γ be a collection of δ-slabs. Let Θ_δ be a subdivision of $\mathbb{S}^{d-1}$ into δ-caps, and let $\Theta_{W^{-1}}$ be a subdivision of $\mathbb{S}^{d-1}$ into W^{-1}-caps.

*We say that Γ **satisfies the first spacing condition with parameter W** if there exists $N \leq (W\delta)^{-1}$ so that for any $\theta \in \Theta_\delta$ and any W^{-1}-slab R with $\boldsymbol{n} \in \theta$, we have*

$$\#\left\{ S\in\Gamma : \begin{array}{c} \boldsymbol{n}(S)\in\theta \\ S \underset{\sim}{\subsetneq} R \end{array} \right\} \sim N. \tag{WS-1}$$

That is, any W^{-1}-slab R (essentially) contains $\sim N$-many δ-slabs with normal direction in the same δ-cap as $\boldsymbol{n}(S)$.

*We say that Γ **satisfies the second spacing condition with parameter W** if, for any $\theta\in\Theta_{W^{-1}}$ and any W^{-1}-slab R with $\boldsymbol{n}(R)\in\theta$,*

$$\#\left\{ S\in\Gamma : \begin{array}{c} \boldsymbol{n}(S)\in\theta \\ S \underset{\sim}{\subsetneq} R \end{array} \right\} \leq 1. \tag{WS-2}$$

That is, for any W^{-1}-slab R, there is at most one δ-slab $S\in\Gamma$ with $S \underset{\sim}{\subsetneq} R$.

Under the first spacing condition I can prove the following estimate for $P_r(\Gamma)$.

Theorem 1.4.2. *For any $\varepsilon>0$ sufficiently small (relative to d), there exists a constant $C_\varepsilon = C(\varepsilon,d)>1$ so that if $\delta\in(0,1)$ and $1\leq W\leq\delta^{-1}$, then*

$$|P_r(\Gamma)| \leq C_\varepsilon \delta^{-\varepsilon} W^{-(d-1)} r^{-d} |\Gamma|^d \tag{1.8}$$

for any collection Γ of δ-slabs that satisfies the first spacing condition for the parameter W and any r with

$$r \geq \delta^{-\varepsilon/4}\delta|\Gamma|. \tag{1.9}$$

Most of my thesis is devoted to the proof of this theorem, which is sharp up to ε-loss, as demonstrated in Example 4.2.2. I also discuss work toward proving an estimate of the form

$$|P_r(\Gamma)| \lesssim_\varepsilon \delta^{-\varepsilon}\frac{|\Gamma|^d}{r^{d+1}}. \tag{1.10}$$

for any collection Γ of size $\sim W^d$ that satisfies the second spacing condition with parameter W and any r sufficiently large.

The bound (1.10) corresponds to Theorem 1.2.3 in the same way that Theorem 1.3.1

corresponds to Theorem 1.2.1. To make this more precise, note that if we start with a set Γ consisting of m hyperplanes, and let $P \subset P_r(\Gamma)$, then Theorem 1.2.3 implies that

$$r|P| \lesssim |I(P,\Gamma)| \lesssim |P|^{\frac{d}{d+1}} m^{\frac{d}{d+1}} + |P| m^{\frac{d-2}{d-1}} + m.$$

If we assume that the first term on the right-hand side dominates (i.e. that it exceeds the other two), then we will have

$$r|P| \lesssim |I(P,\Gamma)| \lesssim |P|^{\frac{d}{d+1}} m^{\frac{d}{d+1}}.$$

We can rearrange to give

$$|P| \lesssim \frac{m^d}{r^{d+1}},$$

which matches our conjectured bound.

For the conjectured bound (1.10) to hold under the spacing condition (WS-2), it is necessary to impose some additional hypotheses. As in Theorem 1.4.2, we will want a lower bound on r, but it turns out that this alone is not enough. My current conjecture involves a property that I call broadness, which I introduce in Chapter 2 after presenting an example which motivates the definition. Chapters 3 and 4 contain more example configurations, including an example that is sharp for (1.10) and one that is sharp for Theorem 1.4.2.

The remaining chapters contain ingredients of the proof of Theorem 1.4.2, culminating in an inductive proof of Theorem 1.4.2 in Chapter 9, which I reinterpret as an iterative procedure in Chapter 10. One particularly important element of the proof is Proposition 6.1.1, which is used to set up an induction on scales argument. I prove Proposition 6.1.1 using the high-low method.

Proposition 6.1.1 sets up a dichotomy to characterize a set P consisting of unit balls that all have approximately the same richness for a set of slabs of dimensions $D \times \cdots \times D \times 1$ with $D > 1$ very large. (We will later apply a rescaled version of this proposition.) We show that the set P must either consist (mostly) of balls grouped into clusters, or P must be 'small' in the sense that we can obtain a bound for $|P|$ in terms of the intended cluster size. In Chapters 7-8, I explain in detail how the bound for $|P|$ that results in this case can used to show the bound of Theorem (1.4.2) directly for δ sufficiently small. Meanwhile, in the thick case, we use an inductive argument to count the number of clusters.

Chapter 2

An Enemy Example and Notions of Broadness

In this chapter we present an example which illustrates the difficulties that arise in proving estimates for $|P_r(\Gamma)|$ under the second spacing condition and discuss candidate hypotheses to rule this example out. When formulating conjectures under the second spacing condition, I was already planning to include assumptions that $|\Gamma| \sim W^d$ and that $r \geq \max\{\delta^{-\varepsilon/4}\delta|\Gamma|, d\}$ (cf Theorem 1.3.3), but these assumptions alone are insufficient to rule out the following example.

Example 2.0.1. *Let $\boldsymbol{n}_1, \ldots, \boldsymbol{n}_d$ be unit vectors which are all contained in the subspace $\boldsymbol{e}_d^{\perp}$. Suppose that for $i = 1, \ldots, d$, the slab S_i is a $1 \times \cdots \times 1 \times \delta$ box centered at $\mathbf{0}$ with normal vector $\boldsymbol{n}_i$. Then each slab S_i contains the δ-neighborhood of the line segment from $-(0, \ldots, 0, -\frac{1}{2})$ to $(0, \ldots, 0, \frac{1}{2})$. The δ-neighborhood of this line segment can be subdivided into $\sim \delta^{-1}$-many distinct δ-boxes, so if we let $\Gamma = \{S_1, \ldots, S_d\}$, then we have*

$$|P_d(\Gamma)| \sim \delta^{-1}.$$

Meanwhile, for $|\Gamma| = d$ and $r = d$, the right-hand side of our conjectured bound (1.10) becomes

$$\delta^{-\varepsilon} d^{-1},$$

which is much smaller than δ^{-1}.

This example is not an issue when it comes to proving estimates under the first spacing condition, because it does not satisfy the first spacing condition, which mandates that each distinct W^{-1}-slab for each normal direction in Θ_δ must contain $\sim N$-many δ-slabs of Γ with that same normal direction. However, merely imposing a lower bound on Γ in terms of W (cf. Theorem 1.3.3) does not rule out Example 2.0.1 if W is small.

The main issue at play in Example 2.0.1 is that the normal vectors of the slabs of Γ all lie in a hyperplane. This suggests that we might want to impose a requirement that each d-rich δ-box for Γ has slabs through it whose normal vectors are transverse. (The reason we single out d-rich boxes here is that d is the smallest possible richness that could give rise to a collection of transverse normal vectors.) For instance, we could require that for any $q \in P_d(\Gamma)$ there are slabs $S_1, \dots S_d \in \Gamma$ with $q \subsetsim S_1, \dots, S_d$ whose normal vectors are ν-transverse, where ν was a parameter which we allowed to depend on δ or possibly on W. As in [2], when we say that $\boldsymbol{n}_1, \dots, \boldsymbol{n}_d$ are ν-transverse, we mean that

$$|\boldsymbol{n}_1 \wedge \cdots \wedge \boldsymbol{n}_d| \geq \nu.$$

One issue with this proposed hypothesis is that it depends on our choice of normal vectors. We could modify our proposed hypothesis to say something like, "for any $q \in P_r(\Gamma)$ there are slabs $S_1, \dots S_d \in \Gamma$, each essentially containing q, along with a choice of normal vectors $\boldsymbol{n}_1, \dots, \boldsymbol{n}_d$ which satisfy

$$|\boldsymbol{n}_1 \wedge \cdots \wedge \boldsymbol{n}_d| \geq \nu."$$

However, I want to focus instead focus on the ratio between $|S_1 \cap \cdots \cap S_d|$ and δ^d. This ratio represents the approximate number of δ-boxes contained in the d-fold intersection $S_1 \cap \cdots \cap S_d$.

Definition 2.0.2. *Let* $0 < \kappa < 1$*, and let* Γ *be a collection of* δ*-slabs in* $[0,1]^d$*. We say that* Γ *is* ***κ-broad****, if, for each* δ*-ball* $q \in P_d(\Gamma)$*, there exist distinct slabs* $S_1, \dots, S_d$*, each* *essentially containing q, so that*

$$\frac{|S_1 \cap \cdots \cap S_d|}{\delta^d} \leq \kappa^{-1}.$$

We say that such a d-tuple of slabs ***witnesses κ-broadness*** *for q.*

With this definition, we are ready to state our conjecture for slabs under the second spacing condition.

Conjecture 2.0.3. *For any $\varepsilon > 0$ sufficiently small relative to d, there exists a constant $C'_\varepsilon > 1$ along with an increasing function $\kappa_\varepsilon : (0,1) \to (0,\infty)$ so that if Γ is a collection of δ slabs that satisfies (WS-2) with parameter W, Γ is $\kappa_\varepsilon(W^{-1})$-broad, and $|\Gamma| \sim W^d$, then*

$$|P_r(\Gamma)| \leq C'_\varepsilon \delta^{-\varepsilon} r^{-(d+1)} |\Gamma|^d$$

for any r with

$$r \geq \max\{d, \delta^{-\varepsilon/4}\delta|\Gamma|\}.$$

I have no immediate reason to believe this conjecture is false, but it seems that it would be hard to prove by induction. This is because the inductive arguments used in the rest of this paper, like those in [10], often rely on estimating incidences of just those slabs in a particular subset of Γ, possibly after rescaling those slabs. It may be the case that Γ as a whole is κ-broad but a subset of interest is not. In particular, given a large parameter $D > 1$, we may wish to consider only those slabs of Γ which lie in a thick slab of dimensions $\sim 1 \times \cdots \times 1 \times D^{-1}$. This suggests that in order to facilitate a proof of our conjectured bound, we may wish to add yet another hypothesis which requires that the subcollection $\Gamma \cap \square$ is broad for any D^{-1}-slab $\square$. Perhaps we cannot expect the collection $\Gamma \cap \square$ to be κ-broad, but we could ask that it is κ'-broad for some $\kappa' < \kappa$. How big should this κ' be though? Should it depend on D? And how big or small should we let D be?

Chapter 3

Examples Based on Farey Fractions

In this chapter, we describe sharp examples for Theorem 1.2.2 and Theorem 1.2.3. For certain values of r and W the hyperplanes of Example 3.0.2 can be thickened to give an example that is sharp, up to ε-loss, for the bound in Conjecture 2.0.3. The constructions rely on some auxiliary lemmas, which we state and prove at the end of the chapter.

Example 3.0.1. (Slope Example/Uniformly Rich Example in $\mathbb{R}^2$) Let $N, Q \in \mathbb{N}$ with $N \geq 100\,Q$, and let $P \subset \mathbb{Z}^2$ be an $N \times N$ grid. Let $S_Q = \{s_1, \ldots, s_r\} \subset \mathbb{Q}$ be the set of rational numbers of the form $\frac{p}{q}$ with $1 \leq p \leq q \leq Q$ and $\gcd(p, q) = 1$. For example, if $Q = 4$, then $S_Q = \{1, \frac{1}{2}, \frac{1}{3}, \frac{2}{3}, \frac{1}{4}, \frac{3}{4}\}$. The members of S_Q comprise the **Farey sequence of order $\boldsymbol{Q}$**. We will think of the members of S_Q as potential slopes for lines. The index r represents the number of potential slopes. By our forthcoming Lemma 3.1.1, $r \sim Q^2$.

Let $\mathcal{L}$ be a set of lines so that for each $p \in P$ and each $s \in S_Q$, there is a line of slope s through p. This means that each point of P is r-rich, i.e. $P \subseteq P_r(\mathcal{L})$. In fact, each point of P is exactly r-rich.

There may be pairs (p, s) and (p', s') so that the line of slope s through p and the line of slope s' through p' coincide. A line that is produced by multiple pairs is counted in $|\mathcal{L}|$ only once. To estimate $|\mathcal{L}|$ we use double counting: we have that

$$|P|r \sim I(P, \mathcal{L}) = \sum_{\ell \in \mathcal{L}} \#\{P \cap \ell\}. \tag{3.1}$$

By Lemma 3.1.3, $\#\{P \cap \ell\} \sim \frac{N}{Q}$ for each $\ell \in \mathcal{L}$. Thus, resuming from (3.1), we obtain that

$$|P|r \gtrsim |\mathcal{L}| \left(\frac{N}{Q} \right) \sim \frac{|\mathcal{L}|N}{r^{1/2}} \sim \frac{|\mathcal{L}||P|^{1/2}}{r^{1/2}}$$

which implies that

$$|P| \gtrsim \frac{|\mathcal{L}|^2}{r^3}.$$

That is, this construction is sharp for the Szemerédi-Trotter bound.

To generalize Example 3.0.1, we need a substitute for the set of slopes we used in our construction. A line in $\mathbb{R}^2$ is uniquely determined by the slope of the line and a choice of a point on the line. Similarly, a hyperplane in $\mathbb{R}^d$ is uniquely determined by a normal vector and a choice of a point on the hyperplane.

We will again begin with a large integer grid $P = \{1, \ldots, N\}^d$. Through each of these points, we will place hyperplanes with normal vectors drawn from the set S_Q defined by

$$S_Q = \left\{ \left(\frac{p_1}{q}, \ldots, \frac{p_{d-1}}{q}, 1 \right) : 1 \leq p_j < q \leq Q \text{ and } \gcd(p_j, q) = 1 \right\}. \tag{3.2}$$

In words, S_Q is the set of vectors $\left(\frac{p_1}{q}, \ldots, \frac{p_{d-1}}{q}, 1 \right)$ so that each fraction $\frac{p_j}{q}$ is a fraction in lowest terms between 0 and 1 with denominator $\leq Q$.

Example 3.0.2. (Normal Example/Uniformly Rich Example in $\mathbb{R}^d$) Let $P \subset \mathbb{Z}^d$ be an $N \times \cdots \times N$ grid. Let S_Q be as in (3.2), let $r = |S_Q|$, and let Γ be a set of hyperplanes so that for each point $\mathbf{x} \in P$ and each $\mathbf{v} \in S_Q$, there is a hyperplane of Γ with with normal vector $\mathbf{v}$ that passes through $\mathbf{x}$. To estimate r, note that for each $q \leq Q$, the number of choices for a $(d-1)$-tuple $(p_1, \ldots, p_{d-1})$ is $(\varphi(q))^{d-1}$, where φ denotes Euler's totient function. Thus, by Lemma 3.1.2, $r \sim Q^d$.

By design, each point of P is r-rich for Γ. By Lemma 3.1.4 below, we have that

$$r|P| \sim |I(P, \Gamma)| = \sum_{\gamma \in \Gamma} \#\{P \cap \gamma\} \gtrsim |\Gamma| \left(\frac{N^{d-1}}{Q} \right) \sim |\Gamma| \left(\frac{|P|^{(d-1)/d}}{r^{1/d}} \right).$$

We rearrange to give

$$|P| \gtrsim \frac{|\Gamma|^d}{r^{d+1}},$$

which means that this construction is sharp for Theorem 1.2.3.

3.1 Auxiliary Lemmas

Here, we provide lemmas that fill in details of the constructions we have already presented. I first learned of the result of Lemma 3.1.1 from [7] in the proof of Proposition 5.2 of that paper. However, the result is standard, and the implicit constant can be made explicit. (See, e.g. Theorem 3.7 of [1].) When combined with Hölder, Lemma 3.1.1 can be used to estimate other moments of Euler's totient function.

3.1.1 Moments of Euler's Totient Function

Lemma 3.1.1. *We have that*

$$\sum_{q \leq Q} \varphi(q) \gtrsim Q^2.$$

Lemma 3.1.2. *We have that*

$$\sum_{q \leq Q} \varphi(q)^k \gtrsim Q^{k+1}.$$

Proof. Beginning with inequality (3.1.1) we apply Hölder's inequality for the exponents k and $k' = k/(k-1)$ to give

$$\begin{aligned} Q^2 \leq \sum_{1 \leq q \leq N} 1 \cdot \varphi(q) &\leq \left(\sum_{1 \leq q \leq Q} \varphi(q)^k \right)^{1/k} \left(\sum_{1 \leq q \leq Q} 1^{k'} \right)^{1/k'} \\ &= Q^{\frac{k-1}{k}} \left(\sum_{1 \leq q \leq Q} \varphi(q)^k \right)^{1/k} \end{aligned}$$

and then rearrange to give

$$\sum_{1 \leq q \leq Q} \varphi(q)^k \gtrsim \left(Q^{-\frac{k-1}{k}} \cdot Q^2 \right)^k = Q^{2k-(k-1)} = Q^{k+1}.$$

□

3.1.2 Counting integer points on zero sets of linear equations

Lemma 3.1.3. *Fix $N, Q \in \mathbb{N}$ with $N \geq 100\,Q$. Let $p, q \in \mathbb{Z}$ with $1 \leq p \leq q \leq Q$. Let ℓ be the line with equation of the form*

$$y = \frac{p}{q}x + \frac{\beta}{q} \tag{3.3}$$

with $\beta \in \mathbb{Z}$. Then the number of points $(x, y) \in P = \{1, \ldots, N\}^2$ that lie on the line ℓ satisfies

$$\#\{P \cap \ell\} \gtrsim \tfrac{N}{q} \gtrsim \tfrac{N}{Q}.$$

Proof. We can estimate the number of points in P on $\ell_{a,b}$ by considering the number of x for which the value of $\frac{p}{q}x + \frac{\beta}{q}$ lies in the set $\{1, \ldots, N\}$. We write

$$\begin{aligned} \#(\ell \cap P) &= \#\left\{(x, y) \in \{1, \ldots, N\}^2 : y = \frac{p}{q}x + \frac{\beta}{q}\right\} \\ &= \#\left\{x \in \{1, \ldots, N\} : \frac{p}{q}x + \frac{\beta}{q} \in \{1, \ldots, N\right\}. \end{aligned}$$

We observe that any two members of the set $\{x \in \{1, \ldots, N\} : \frac{p}{q}x + \frac{\beta}{q} \in \{1, \ldots, N\}$ must differ by a multiple of q. (To see why this is, note that if $x \neq x' \in \mathbb{Z}$ are chosen so that both $\frac{p}{q}x + \frac{\beta}{q}$ and $\frac{p}{q}x' + \frac{\beta}{q}$ are integers, then it must be the case that b divides the difference $\left(\frac{p}{q}x + \frac{\beta}{q}\right) - \left(\frac{p}{q}x' + \frac{\beta}{q}\right) = \frac{p}{q}(x - x')$. Since $\gcd(p, q) = 1$, this implies that q divides $x - x'$.)

Since any two values of x in the set $\{x \in \{1, \ldots, N\} : \frac{p}{q}x + \frac{\beta}{q} \in \{1, \ldots, N\}$ must differ by a multiple of q and must both be contained in $\{1, \ldots, N\}$, it follows that the size of this set is $\sim \frac{N}{q}$. □

Lemma 3.1.4. *Fix $N, Q \in \mathbb{N}$ with $N \geq 100\,Q$. Let $d \geq 3$, and let $P = \{1, \ldots, N\}^d$. Let γ denote the plane with equation*

$$x_d = \frac{p_1}{q}x_1 + \cdots + \frac{p_{d-1}}{q}x_{d-1} + \frac{p}{q} := f(x_1, \ldots, x_{d-1}). \tag{3.4}$$

Then

$$\#(\gamma \cap P) \gtrsim \frac{N^{d-1}}{q} \gtrsim \frac{N^{d-1}}{Q}.$$

Proof. By analogy with the proof of Lemma 3.1.3 we will count the number of $(d-1)$-tuples $\vec{x} = (x_1, \dots, x_d) \in \{1, \dots, N\}^{d-1}$ such that $f(x_1, \dots, x_d)$ is an integer in $\{1, \dots, N\}$. In the case that $f(x_1, \dots, x_d)$ is is an integer in this range, we can set $x_d = f(x_1, \dots, x_d)$ to give a point $\mathbf{x} = (x_1, \dots, x_d)$ in $P \cap \gamma$.

If we fix $x_1, \dots, x_{d-2}$ in $\{1, \dots, N\}$, then the number of ways to choose x_{d-1} with $\frac{1}{q}(p_1 x_1 + \dots + p_{d-1}x_{d-1} + p) \in \{1, \dots, N\}$ is $\sim \frac{N}{q}$. This is by a similar argument to the proof of Lemma 3.1.3 which we nevertheless record here for the sake of completeness: if we choose x_{d-1} and x'_{d-1} so that both $\frac{1}{q}(p_1 x_1 + \dots + p_{d-1}x_{d-1} + p)$ and $\frac{1}{q}(p_1 x_1 + \dots + p_{d-1}x'_{d-1} + p)$ are integers in $\{1, \dots, N\}$, then q must divide the difference $(p_1 x_1 + \dots + p_{d-1}x_{d-1} + p) - (p_1 x_1 + \dots + p_{d-1}x'_{d-1} + p) = p_{d-1}(x_{d-1} - x'_{d-1})$, which implies that q divides $x_{d-1} - x'_{d-1}$.

The above work was for a fixed $(d-2)$-tuple $(x_1, \dots, x_{d-2}) \in \{1, \dots, N\}^{d-2}$. There are N^{-2} such $(d-2)$-tuples and $\sim \frac{N}{q}$ ways to complete each to a $d-1$-tuple $(x_1, \dots, x_{d-1}) \in \{1, \dots, N\}^{d-1}$ with $f(x_1, \dots, x_d) = 1$. Hence,

$$|P \cap \gamma| \sim \frac{N^{d-1}}{q} \gtrsim \frac{N}{Q}.$$

□

Chapter 4

Star Example and Cousins

4.1 Preliminaries

Let Λ be a maximal δ-separated set on $\mathbb{S}^{d-1}$. Then $|\Lambda| \sim \delta^{-(d-1)}$. This set can be thought of a set of candidate normal directions for slabs. We first consider a set of slabs Γ_0 containing one slab with each normal direction in Λ so that each slab in Γ_0 is centered at $\mathbf{0}$. The following lemma gives an estimate for the number of these slabs that pass through a typical point on a sphere of radius σ with $0 < \sigma < 1$.

Lemma 4.1.1. *Let Λ be a maximal δ-separated set on $\mathbb{S}^{d-1}$, and let Γ_0 be a set of δ-slabs centered at $\mathbf{0}$ so that each slab of Γ_0 is contained in the δ-neighborhood of the hyperplane $\boldsymbol{v}^{\perp}$ for some $\boldsymbol{v} \in \Lambda$. If $\omega \in \mathbb{R}^d$ is a point at distance σ from 0 with $\delta < \sigma < 1$, then the number of slabs that pass through ω satisfies*

$$\#\{S \in \Gamma : \omega \in S\} \sim \frac{\delta^{-(d-2)}}{\sigma}.$$

Proof. We consider $\sigma\mathbb{S}^{d-1}$, the sphere of radius σ centered at $\mathbf{0}$. Each slab of Γ_0 intersects this sphere in a band with surface area $\sim \delta\sigma^{d-2}$. (Here, when we say 'surface area,' we are referring to $(d-1)$-dimensional volume.) Since Λ was maximal, the normal directions of the slabs are approximately equidistributed on the surface of the sphere. Thus, the number of bands passing through each $\omega \in \sigma\mathbb{S}^{d-1}$ is approximately the total surface area of all the bands over the surface area of $\sigma\mathbb{S}^{d-1}$. That is, for each $\omega \in \sigma\mathbb{S}^{d-1}$, we have

$$\#\{S \in \Gamma : \omega \in S\} \sim \frac{\sum_{S\in\Gamma_0} \mathrm{Vol}^{d-1}(\Gamma \cap \sigma\mathbb{S}^{d-1})}{\mathrm{Vol}^{d-1}(\mathbb{S}^{d-1})} \sim \frac{|\Gamma_0|\delta\sigma^{d-2}}{\sigma^{d-1}} \sim \frac{\delta^{-(d-2)}}{\sigma}.$$

□

4.2 Examples

Example 4.2.1. Let Γ_0 be a set of slabs centered at $\mathbf{0}$ so that each slab of Γ_0 is contained in the δ-neighborhood of the hyperplane $\boldsymbol{v}^\perp$ for some $\boldsymbol{v} \in \Lambda$. We note that $|\Lambda| \sim \delta^{-(d-1)}$, so $\mathbf{0}$ is $\sim \delta^{-(d-1)}$-rich for Γ_0. Meanwhile, for $r < \delta^{-(d-1)}$, we can have many r-rich δ-balls in an $A\delta$-ball around 0 for an appropriately chosen $A > 1$.

By Lemma 4.1.1, a typical point $\omega \in \sigma\mathbb{S}^{d-1}$ is r-rich if $\sigma \lesssim \delta^{-(d-2)}r^{-1}$. Motivated by this lemma, we take $A \sim \delta^{-(d-1)}r^{-1}$ so that $A\delta \sim \delta^{-(d-2)}r^{-1}$. All of the δ grid boxes in an $A\delta$-ball centered at $\mathbf{0}$ are r-rich. The $A\delta$-ball centered at $\mathbf{0}$ can be subdivided into $\sim A^d$-many δ-boxes, so

$$|P_r(\Gamma_0)| \gtrsim A^d \sim \delta^{-d(d-1)}r^{-d}.$$

Example 4.2.2. (Stacked Star Fragments) We can build a sharp example for Theorem 9.0.1 by joining together many star fragments consisting of many copies of a subset $\Gamma_0' \subset \Gamma_0$, where Γ_0' is formed by removing any slabs of Γ_0 with normal directions too close to a 'bad' direction. (What constitutes a 'bad' direction is to be made precise momentarily.)

We place star fragment centers along a line segment with direction vector $\boldsymbol{e_d}$ through the center of the unit box. We take the star centers to be a maximal collection of W^{-1}-spaced points on this line segment. Through each of these points, we place one slab of each normal direction in $\Lambda' = \Lambda_0 \backslash B_w$, where B_w is a band of width $w < 1$ (to be determined) around the equator $\boldsymbol{e_d}^\perp \cap \mathbb{S}^{d-1}$. The number of points of Λ within this band satisfies

$$\#\{\boldsymbol{v} \in \Lambda : \boldsymbol{v} \in B_w\} \sim \frac{\mathrm{Vol}^{d-1}(\mathrm{Band})}{\mathrm{Vol}^{d-1}(\delta\text{-cap})} \sim w\delta^{-(d-2)}.$$

The reason we construct star fragments without the directions near the equator is that if $\boldsymbol{v}$ is too close to $\boldsymbol{e}_{\boldsymbol{d}}^{\perp}$ then slabs with normal vector $\boldsymbol{v}$ through different star centers will coincide with each other. To determine how we should choose w, suppose that $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ are two different star fragment centers with $\operatorname{dist}(\boldsymbol{x}_1, \boldsymbol{x}_2) = W^{-1}$, suppose that $\operatorname{angle}(\boldsymbol{v}, \boldsymbol{e}_{\boldsymbol{d}}^{\perp}) = \alpha$, and suppose that S_1 and S_2 are slabs of dimensions $1 \times \cdots \times 1 \times \delta$ with normal direction $\boldsymbol{v}$, centered at $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$, respectively. A computation shows that

$$\operatorname{dist}(S_1, S_2) \gtrsim W^{-1} \sin(\alpha) \geq \delta \sin \alpha. \tag{4.1}$$

We choose w to ensure that any vector $\boldsymbol{v} \in B_w$ makes a sufficiently small angle to ${e_d}^{\perp}$ that the sine of the angle compensates for the implied constant in (4.1). We'll have that $|\Lambda'| \geq |\Lambda| - |\Lambda \backslash B_w|$. So long as w is sufficiently small (relative to the implied constant in the bound $|\Lambda| \gtrsim \delta^{-(d-1)}$), we'll have $|\Lambda'| \geq \frac{1}{2}|\Lambda|$. Say for concreteness that $w = \frac{1}{1000}$ is small enough for this to work and for (4.1) to ensure essential distinctness of our slabs.

After forming a set of slabs Γ_0' by taking one slab through $\mathbf{0}$ for each normal vector in Λ', we let Γ_1 be the union of all of the slabs from the star fragments centered at our designated points. Then we have that $|\Gamma_1| \sim W|\Lambda'|$. We note that Γ_1 does not yet satisfy the spacing condition of Theorem 9.0.1. This is because if θ is within $\frac{1}{1000}$ of the subspace $\boldsymbol{e}_{\boldsymbol{d}}^{\perp}$, then there are no δ-slabs with normal direction θ inside any W^{-1}-slab with normal direction θ. To remedy this, we add more slabs to Γ_1. Specifically, for each $\boldsymbol{v} \in \Lambda \backslash \Lambda'$, we subdivide $[0,1]^d$ into thick slabs of dimensions $\sim 1 \times \cdots \times 1 \times W^{-1}$. In each of these thick slabs we insert one δ^{-1}-slab with normal vector $\boldsymbol{v}$. We denote by Γ the resulting set of slabs (i.e. the union of the original Γ_1 with these new slabs). We note that $|\Gamma \backslash \Gamma_1| \sim W\, |\Lambda \backslash \Lambda'|$ and

$$|\Gamma| \sim |\Gamma_1| \sim W|\Lambda_1| \sim W\delta^{-(d-1)}. \tag{4.2}$$

By our work in Example 4.2.1, each star fragment contributes $\sim \delta^{-d(d-1)} r^{-d}$-many boxes to $P_r(\Gamma_1)$. Provided that r is sufficiently large (see Remark 4.2.3), the clusters of r-rich

δ-boxes do not intersect, so we have by (4.2) that

$$|P_r(\Gamma)| \geq |P_r(\Gamma_1)| \gtrsim W\left(\delta^{-(d-1)}r^{-1}\right)^d$$
$$\sim W\left(|\Gamma|W^{-1}\right)^d r^{-d}$$
$$= |\Gamma|^d r^{-d} W^{-(d-1)}$$

Remark 4.2.3. *As an addendum to Example* 4.2.2*, we note that the reason that the clusters of r-rich delta-balls do not collide is our hypothesis that $r \geq \delta^{-\varepsilon/4}\delta|\Gamma|$.*

To prevent the clusters from colliding, we need to ensure that

$$W^{-1} \gtrsim A\delta \sim \delta\left(\frac{\delta^{-(d-1)}}{r}\right) = \frac{\delta^{-d}\delta^2}{r}.$$

This occurs if and only if

$$r \gtrsim \delta^{-d}\delta^2 W. \tag{4.3}$$

Under the hypotheses of Theorem 1.4.2*, we have that*

$$r \geq \delta^{-\varepsilon/4}\delta|\Gamma|$$

and

$$|\Gamma| \sim \delta^{-(d-1)}WN$$

for some $N \geq 1$. (In Example 4.2.2*, we have $N = 1$.) Combining these inequalities gives*

$$r \gtrsim \delta^{-\varepsilon/4}\delta\left(\delta^{-(d-1)}W\right) = \delta^{-\varepsilon/4}\delta^{-(d-2)}W > \delta^{-(d-2)}W.$$

Chapter 5

The Spacing Conditions and Base Cases for Induction

In this chapter, we further discuss the spacing conditions that we introduced at the end of Chapter 1. We explain base cases for induction under each spacing condition. We explain the structure of our inductive proof of Theorem 1.4.2 and comment on modifications we would need to make to this structure to prove Conjecture 2.0.3.

One particularly important result in this chapter is Lemma 5.1.3, which demonstrates that if we are given a threshold $c_\varepsilon < 1$, then there exists C_ε so that the goal bound 1.8 holds for all $\delta \geq c_\varepsilon$. Thus, after we have proved Lemma 5.1.3, the proof of Theorem 1.4.2 reduces to showing that the goal bound holds with this same constant C_ε for all $\delta < c_\varepsilon$. For this, we use induction on scales, as described in Section 5.2.

5.1 Lemmas for Base Cases

If Γ satisfies (WS-1), then any two slabs with the same center - or, more generally, any two slabs with a substantial intersection - must have normal vectors separated by an angle $\gtrsim \delta$. A maximal set of δ-separated points on $\mathbb{S}^{d-1}$ has size $\sim \delta^{-(d-1)}$, which limits the values of r for which $P_r(\Gamma)$ is nonempty. Specifically, we have the following, lemma, which we use in one of the base cases in our inductive proof of Theorem 1.4.2.

Lemma 5.1.1. *There exists a dimensional constant $\alpha_d > 0$ so that if Γ is a collection of*

δ-slabs satisfying (WS-1) with parameter W and $r \geq \alpha_d \delta^{-(d-1)}$, then $P_r(\Gamma) = \varnothing$.

Under (WS-2), there is an even stronger angle separation condition: if Γ satisfies (WS-2), then any two slabs with the same center - or, more generally, any two slabs with a substantial intersection - must have normal vectors separated by an angle $\gtrsim W^{-1}$. Thus, we have the following lemma, which we intend to use as a base case in an inductive proof of Conjecture 2.0.3.

Lemma 5.1.2. *There exists a dimensional constant $\alpha'_d > 0$ so that if Γ is δ-slabs satisfying (WS-1) with parameter W and $r \geq \alpha'_d W^{d-1}$, then $P_r(\Gamma) = \varnothing$.*

Lemmas 5.1.1 and 5.1.2 will be used for base cases corresponding to an assumption that r is large. We will also have base cases corresponding to an assumption that δ is large.

Under the first spacing condition, there are a couple potential thresholds for largeness that interest us. If δ is 'too large' relative to W, then $P_r(\Gamma)$ must be empty. We make this observation precise in Lemma 5.1.4. Meanwhile, if δ is large in an absolute sense, then we can choose C_ε so that the bound of Theorem 1.4.2 follows from the 'trivial bound'

$$|P_r(\Gamma)| \lesssim \delta^{-d}, \tag{5.1}$$

which comes from the fact that any subdivision of $[0,1]^d$ into δ-boxes contains $\lesssim \delta^{-d}$-many boxes. Specifically, we have the following lemma.

Lemma 5.1.3. *For any $\varepsilon > 0$ and any threshold $c_\varepsilon \in (0,1)$, there exists a constant $C_\varepsilon = C(\varepsilon, d)$ so that if $1 > \delta \geq c_\varepsilon > 0$ and Γ is a collection of δ-slabs which satisfies the first spacing condition with parameter W, then*

$$|P_r(\Gamma)| \leq C_\varepsilon \delta^{-\varepsilon} r^{-d} W^{-(d-1)} |\Gamma|^d. \tag{5.2}$$

Before proving this lemma, we return to the issue of what happens when δ is 'too large' relative to W.

Lemma 5.1.4. *There exists a dimensional constant $\beta_d > 0$ so that if Γ is a collection of*

δ-slabs satisfying (WS-1) with parameter $W \geq \beta_d\, \delta^{-1+\frac{\varepsilon}{10d}}$ then then

$$P_r(\Gamma) = \varnothing.$$

for any r with

$$r \geq \delta^{-\frac{\varepsilon}{4}} \delta |\Gamma|. \tag{5.3}$$

Moreover, there exists a constant c_ε so that if $\delta \leq c_\varepsilon$ and Γ is a collection of δ-slabs satisfying (WS-1) with parameter $W \geq \delta^{-1+\frac{\varepsilon}{10d}}$, then

$$P_r(\Gamma) = \varnothing$$

*for any r satisfying (*5.3*).*

Proof. Since Γ satisfies (WS-1) with parameter $W \geq \beta_d \delta^{-1+\frac{\varepsilon}{10d}}$ we have that

$$|\Gamma| \sim \delta^{-(d-1)} W N \geq \delta^{-(d-1)} W \geq \beta_d \delta^{-1+\frac{\varepsilon}{10d}} \delta^{-(d-1)}.$$

Combining this with 5.3, the bound for r from Theorem 1.4.2, gives

$$\begin{aligned} r &\geq \delta^{-\frac{\varepsilon}{4}} \delta |\Gamma| \\ &\gtrsim \delta^{-\frac{\varepsilon}{4}} \delta \left(\beta_d \delta^{-1+\frac{\varepsilon}{10d}} \delta^{-(d-1)} \right) \\ &= \beta_d \delta^{-(d-1)} \delta^{-\frac{\varepsilon}{4}+\frac{\varepsilon}{10d}}. \end{aligned}$$

We note that $-\frac{\varepsilon}{4} + \frac{\varepsilon}{10d} > 0$, so, $\beta_d \delta^{-(d-1)} \delta^{-\frac{\varepsilon}{4}+\frac{\varepsilon}{10d}} > \beta_d \delta^{-(d-1)}$. By Lemma 5.1.1, we conclude that if β_d was sufficiently large, then $P_r(\Gamma)$ is empty.

Alternatively, if we do not use the inequality $\delta^{-(d-1)} \delta^{-\frac{\varepsilon}{4}+\frac{\varepsilon}{10d}} > \delta^{-(d-1)}$, we can conclude that if δ is sufficiently small, then $\delta^{-(d-1)} \delta^{-\frac{\varepsilon}{4}+\frac{\varepsilon}{10d}} \geq \alpha \delta^{-(d-1)}$, where α is the dimensional constant from Lemma 5.1.1. □

Having proved Lemma 5.1.4, we return to the issue of proving Lemma 5.1.3.

Proof. (Proof of Lemma 5.1.3) By Lemma 5.1.1, we may assume that $r \lesssim \delta^{-(d-1)}$, because if r exceeded the threshold of Lemma 5.1.1, then we'd have $P_r(\Gamma) = \varnothing$ (in which case 5.2 would clearly hold).

Assuming that $r \lesssim \delta^{-(d-1)}$, it suffices by (5.1) to show that there exists C_ε so that if $\delta \leq c_\varepsilon$, then

$$\delta^{-d} \lesssim C_\varepsilon \delta^{-\varepsilon} r^{-d} W^{-(d-1)} |\Gamma|^d. \tag{5.4}$$

In turn, (5.4) holds for any $\delta \leq c_\varepsilon$ if

$$c_\varepsilon^{-d} \lesssim C_\varepsilon r^{-d} W^{-(d-1)} |\Gamma|^d. \tag{5.5}$$

We now analyze the right-hand side of (5.5) in hopes of finding a lower bound for it. Since Γ satisfies (WS-1) with parameter W, we have that

$$|\Gamma| \sim \delta^{-(d-1)} W N \geq \delta^{-(d-1)} W$$

from which it follows that

$$W^{-(d-1)} |\Gamma|^d \gtrsim W^{-(d-1)} \delta^{-d(d-1)} W^d = \delta^{-d(d-1)} W.$$

We use this result along with our assumption that $r \lesssim \delta^{-(d-1)}$ to give

$$C_\varepsilon r^{-d} W^{-(d-1)} |\Gamma|^d \gtrsim C_\varepsilon \delta^{-\varepsilon} \delta^{d(d-1)} \left(\delta^{-d(d-1)} W \right) = C_\varepsilon W \geq C_\varepsilon.$$

It follows that (5.5) holds so long as

$$C_\varepsilon \geq C_d \, c_\varepsilon^{-d}$$

for a sufficiently large dimensional constant C_d. □

Though the statement of Lemma 5.1.4 gives a lower bound for W sufficient to guarantee that $P_r(\Gamma)$ is empty, this lower bound for W can be rearranged to give a lower bound for δ. Specifically, we have

$$W \geq \beta_d \, \delta^{-1+\frac{\varepsilon}{10d}}$$

if and only if

$$\delta \geq \left(\beta_d W^{-1}\right)^{\frac{10d-\varepsilon}{10d}}. \tag{5.6}$$

Depending on the values of W and ε, it might be the case that the threshold from 5.6 is larger than our eventual choice of c_ε, or it might be the case that our chosen c_ε exceeds the right-hand side of 5.6. This is why we will use both Lemma 5.1.4 and Lemma 5.1.3 as base cases for our induction.

Both Lemma 5.1.3 and Lemma 5.1.4 have analogous formulations for the second spacing condition, with proofs basically the same as the ones we gave above. Thus, proving our conjecture under the second spacing condition would also reduce to showing the statement holds for all δ sufficiently small.

5.2 Our Inductive Arguments

To prove Theorem 1.4.2, we use induction on scales. Our work towards proving Conjecture 2.0.3 also involves induction on scales. Here, we discuss at a very high level how we set up inductive arguments for each spacing condition.

5.2.1 Induction Under the First Spacing Condition

Let W and ε be fixed, and let $X_{\varepsilon,W}(r,\delta;\Gamma)$ denote the statement,
"The set Γ is a collection of δ-slabs satisfying (WS-1) with parameter W, and $r \geq \delta^{-\varepsilon/4}\delta|\Gamma|$."
We let $Y_{\varepsilon,W}(r,\delta;\Gamma)$ denote the statement,

$$|P_r(\Gamma)| \leq C_\varepsilon r^{-d} W^{-(d-1)} |\Gamma|^d. \tag{5.7}$$

Given r and δ with $0 < \delta \leq W^{-1}$, we wish to prove that if Γ is a collection of δ-slabs satisfying (WS-1) and $r \geq \delta^{-\frac{\varepsilon}{4}}\delta|\Gamma|$, then (5.7) holds. That is, we wish to prove the implication

$$X_{\varepsilon,W}(r,\delta;\Gamma) \implies Y_{\varepsilon,W}(r,\delta;\Gamma). \tag{5.8}$$

For this, we assume as an inductive hypothesis that for any pair $\left(\tilde{r}, \tilde{\delta}\right)$ with $\tilde{r} > r$ or $\tilde{\delta} > \delta$, we have the implication

$$X_{\varepsilon,W}\left(\tilde{r}, \tilde{\delta}; \tilde{\Gamma}\right) \implies Y_{\varepsilon,W}\left(\tilde{r}, \tilde{\delta}; \tilde{\Gamma}\right).$$

The idea is then to construct from Γ a collection of $\tilde{\delta}$-slabs $\tilde{\Gamma}$ by taking subsets of Γ and applying appropriate transformations to the slabs of each subset. The precise definition of $\tilde{\Gamma}$ will be different for different values of the parameters r and δ. Specifically, we'll first take cases on whether

$$\delta \geq c_\varepsilon \tag{5.9}$$

with c_ε to be defined later. If $\delta < c_\varepsilon$, then we'll take cases on whether

$$|P_r(\Gamma) \backslash P_{2r}(\Gamma)| \geq \frac{1}{10} |P_r(\Gamma)|, \tag{5.10}$$

and finally, if (5.10) holds, we'll define an integral that approximates $|I(P, \Gamma)|$, apply Plancherel's theorem, and then take cases according to whether high frequencies or low frequencies make a bigger contribution to this integral. These two outcomes correspond to the thin case and the thick case of Proposition 6.1.1.

In order to apply the inductive hypothesis, we'll want to show that $X_{\varepsilon,W}(r, \delta; \Gamma)$ implies $X_{\varepsilon,W}\left(\tilde{r}, \tilde{\delta}; \tilde{\Gamma}\right)$ and that $Y_{\varepsilon,W}(r, \delta; \Gamma)$ implies $Y_{\varepsilon,W}\left(\tilde{r}, \tilde{\delta}; \tilde{\Gamma}\right)$. This plan corresponds to proving a 'trapezoid of implications.'

$$\begin{array}{ccccc} X_{\varepsilon,W}(r, \delta; \Gamma) & & & & Y_{\varepsilon,W}(r, \delta; \Gamma) \\ \Searrow \ (a) & & & & \Nearrow \ (c) \\ & X_{\varepsilon,W}(\tilde{r}, \tilde{\delta}; \tilde{\Gamma}) & \underset{(b)}{\implies} & Y_{\varepsilon,W}(\tilde{r}, \tilde{\delta}; \tilde{\Gamma}) & \end{array}$$

Here, implication (b) is assumed as an inductive hypothesis, whereas implications (a) and (c) are to be proved.

In each case, we only prove implications (a) and (c) for $\delta \leq c_\varepsilon$, where c_ε satisfies some conditions that we will describe momentarily. As we showed in Lemma (5.1.3), there is a choice of C_ε so that (5.8) holds for any $\delta \geq c_\varepsilon$. We consider the case that $\delta \geq c_\varepsilon$ to be a base case of our induction.

We can also prove implication (5.8) directly in the thin case (i.e. the case that (5.10) holds and we have high-frequency dominance), so long as δ is 'sufficiently small' relative to ε. The precise meaning of 'sufficiently small' here is one of the factors that contributes to our choice of c_ε (and, thus, to the value of the constant C_ε in our bound).

5.2.2 Intended Inductive Argument Under the Second Spacing Condition

Unlike in the proof of Theorem 1.4.2, we *do not* consider a fixed W but instead allow W to vary from scale to scale. We let $Y'_\varepsilon(r, \delta, W^{-1}; \Gamma)$ denote the statement,

$$|P_r(\Gamma)| \leq C_\varepsilon r^{-(d+1)} |\Gamma|^d. \tag{5.11}$$

We let $X'_\varepsilon(r, \delta, W^{-1}; \Gamma)$ denote the statement, "The set Γ is a collection of δ-slabs satisfying (WS-1) with parameter W so that $|\Gamma| \sim W^d$, Γ is $\nu_\varepsilon(W^{-1})$ broad, and $r \geq \max\{\delta^{-\varepsilon/4}\delta|\Gamma|, d\}$." We now assume as an inductive hypothesis that for any triple $\left(\tilde{r}, \tilde{\delta}, \tilde{W}^{-1}\right)$ with $\tilde{r} > r$, $\tilde{\delta} > \delta$, or $\tilde{W}^{-1} > W^{-1}$, we have the implication

$$X'_{\varepsilon,W}\left(\tilde{r}, \tilde{\delta}, \tilde{W}^{-1}; \tilde{\Gamma}\right) \implies Y'_{\varepsilon,W}\left(\tilde{r}, \tilde{\delta}, \tilde{W}^{-1}; \tilde{\Gamma}\right).$$

For any c_ε, there exists a constant C_ε so that we can can deduce the implication $X'_{\varepsilon,W}\left(r, \delta, W^{-1}; \Gamma\right) \implies Y'_{\varepsilon,W}\left(\tilde{r}, \tilde{\delta}, \tilde{W}^{-1}; \tilde{\Gamma}\right)$ directly for $\delta \geq c_\varepsilon(\varepsilon)$.

For an appropriately chosen c_ε, we will consider $\delta \geq c_\varepsilon$ to be one of our base cases. We will also have more base cases corresponding to the events that W is large relative to the dimension d, that δ is large relative to W, or that r is large relative to W. For triples not falling into one of these base cases, we will (try to) prove every implication in the following trapezoid.

$$X'_\varepsilon(r, \delta, W; \Gamma) \qquad\qquad\qquad\qquad Y'_\varepsilon(r, \delta, W; \Gamma)$$

$$\Downarrow \; (a') \qquad\qquad\qquad\qquad \Uparrow \; (c')$$

$$X'_\varepsilon(\tilde{r}, \tilde{\delta}, \tilde{W}^{-1}; \tilde{\Gamma}) \underset{(b')}{\implies} Y'_{\varepsilon,W}(\tilde{r}, \tilde{\delta}, \tilde{W}^{-1}; \tilde{\Gamma})$$

One major difference between this intended proof and the proof of Theorem 9.0.1 is that we cannot deduce the implication

$$X'_{\varepsilon}(r,\delta,W;\Gamma) \implies Y'_{\varepsilon}(r,\delta,W;\Gamma) \tag{5.12}$$

directly in the case that we call the 'thin case' for the spacing condition (WS-2); instead we must use the inductive hypothesis in the thin case as well. This introduces additional difficulties, because implication $\textcircled{a'}$ is not true in the 'thin case' for (WS-2) unless r is sufficiently large relative to δ. This necessitates the introduction of a different argument for small r. I am still trying to figure out how to find an argument that works for all r too small for the thin case induction to close. There are many choices prior to this step that may affect its success: choosing a different broadness will result in a different sub-cube size in the thin case partitioning argument, and thus, a different dividing point for the r we consider 'too small.' Trying to determine the 'right' broadness (if there is indeed a broadness that works) is a project for the future.

Chapter 6

The High-Low Method and a Lemma for Finding Clusters of Rich Boxes

The high-low method refers to the idea of estimating the integral of a square or a product on $\mathbb{R}^d$ by applying Plancherel's identity and then estimating separately the integrals of the high frequency part and the low frequency part of the transformed function. For concreteness, if f, g are in the Schwartz class $\mathcal{S}(\mathbb{R}^d)$, then, given $\rho > 0$, we may write

$$\int_{\mathbb{R}^d} f(x)\overline{g(x)}, dx = \int_{\mathbb{R}^d} \hat{f}(\xi)\overline{\hat{g}(\xi)}\, d\xi = \int_{|\xi|\leq\rho} \hat{f}(\xi)\overline{\hat{g}(\xi)}\, d\xi + \int_{|\xi|>\rho} \hat{f}(\xi)\overline{\hat{g}(\xi)}\, d\xi. \tag{6.1}$$

In practice, we may want to introduce a smoothed out version of the characteristic function $1_{B(0,\rho)}$. If η_ρ is identically 1 on $B(0,\rho)$ and has support in $B(0,2\rho)$, then we can replace (6.1) with

$$\int_{\mathbb{R}^d} f(x)\overline{g(x)}, dx = \int_{\mathbb{R}^d} \hat{f}(\xi)\overline{\hat{g}(\xi)}\, d\xi = \int_{\mathbb{R}^d} \eta_\rho(\xi)\hat{f}(\xi)\overline{\hat{g}(\xi)}\, d\xi + \int_{\mathbb{R}^d} (1-\eta_\rho(\xi))\hat{f}(\xi)\overline{\hat{g}(\xi)}\, d\xi. \tag{6.2}$$

The parameter ρ may be chosen to optimize the resulting estimates.

There is also an analogous formulation of (6.1) for sums, which, to my knowledge, predates the integral version in the literature. In particular, Vinh used a version for sums in [15] which inspired a lot of later work, including the incidence estimates of Guth, Solomon, and Wang, which my work closely follows. The high-low method was also used in [8] to prove decoupling results for the moment curve.

In the ensuing work, we use a decomposition as in (6.2) to estimate an integral which counts incidences. In turn, we estimate this integral to prove a lemma concerning a collection of unit balls of uniform richness, which shows that such a collection must either satisfy a size bound a priori, or it must have rich unit balls that are grouped into clusters. Though we state and prove the lemma for a collection of unit balls, we will later use rescaled versions at multiple scales.

We state the lemma for slabs of dimensions $\sim D \times \cdots \times D \times 1$. We will eventually take D to be a power of δ^{-1}. In the statement and proof of the lemma, we use $\lessapprox_{\log}$ to indicate multiplication by an omitted power of $\log D$, which will become a power of $\log(\frac{1}{\delta})$ when we apply the lemma later.

6.1 Heavy Ball Finding Lemma for Slabs

Proposition 6.1.1. *Suppose that P is a set of unit balls in $[0, D]^d$ and Γ is a set of essentially distinct slabs of dimensions $D \times \cdots \times D \times 1$ in $[0, D]^d$. Suppose that each ball of P lies in $\sim E$ slabs of Γ; specifically, suppose that each ball of P lies in at least E-many slabs, but less than $2E$-many. Let $1 \ll \rho^{-1} \ll \lambda \ll D$ - say that $\lambda = D^{\varepsilon/(10d)}$ for a small positive ε and $\rho = D^{\varepsilon^3}\lambda^{-1}$. For D sufficiently large (relative to ε), at least one of the following occurs:*

(1). Thin case: *$|P| \lessapprox_{\log} \lambda^{d-1}E^{-2}|\Gamma|D^{d-1}$ or*

(2). Thick case: *There is a set of finitely overlapping λ-boxes Q_j ('heavy boxes') such that*

(i) $\cup_j Q_j$ contains a fraction $\gtrapprox_{\log} 1$ of the balls of P, and

(ii) Each Q_j intersects $\gtrsim \rho^{-1}E$-many slabs of Γ.

Remark 6.1.2. *Here λ replaces the parameter S from the analogous proposition in [10]. Since $\rho = D^{\varepsilon^3}\lambda^{-1}$, the the second conclusion of the thick case can alternatively be stated as*

$$\#\{S \in \Gamma : S \cap Q_j \neq \varnothing\} \gtrsim D^{-\varepsilon^3}\lambda E. \tag{6.3}$$

One might wonder about the motivation for our particular choice of λ, $\lambda = D^{\varepsilon/(10d)}$. Exam-

ining the argument in Chapter 8 *shows that for the thin bound* $|P| \lessapprox_{\log} \lambda^{d-1}E^{-2}|\Gamma|D^{d-1}$ *to imply the bound of Theorem* 1.4.2*, we want to have* λ *strictly less than* $D^{\varepsilon/(d-1)}$. *(At the very least, this is necessary to bring the 'naive' approach of Chapter* 7 *to completion.) On the other hand, we want* λ *to be large so that the thickening we perform in the thick case represents a bigger step towards our base cases. (Equivalently, choosing a larger value of* λ *corresponds to iterating the procedure described in Chapter* 10 *a smaller number of times.)*

Proof. We will approximate $I(P, \Gamma)$ by an integral $\int_{\mathbb{R}_d} fg$, where f is a sum of smoothed characteristic functions of unit boxes and g is a sum of smoothed characteristic functions of slabs. One slightly technical point is that the sum defining f will not be a sum over all the boxes in P, but rather over a large subcollection $P' \subset P$ with the property that each $q \in P'$ has approximately the same number of slabs incident to its λ-neighborhood $N_\lambda(q)$. We accomplish this by dyadic pigeonholing. Specifically, we will use dyadic pigeonholing to find a popular value for the dyadic size of the set $\{S \in \Gamma : S \cap N_\lambda(q) \neq \varnothing\}$ as q ranges over the balls of P.

For each $q \in P$, we let $W_\lambda(q)$ denote the cardinality of the set $\{S \in \Gamma : S \cap N_\lambda(q) \neq \varnothing\}$. For each $q \in P$, the set $N_\lambda(q)$ consists of $\sim \lambda^d$-many essentially distinct unit boxes. Although q is $\sim E$-rich, these unit boxes may not be. However, we can still find an upper bound for the richness of these neighbor boxes using our assumption that the slabs of Γ are essentially distinct, which implies that the normal directions of the slabs through any particular unit box must be $\gtrsim \frac{1}{D}$-separated. Thus, if q' is a neighbor of q, then

$$\#\{S \in \Gamma : q' \cap S \neq \varnothing\} \lesssim \frac{1}{(1/D)^{d-1}} = D^{d-1}. \tag{6.4}$$

Multiplying this by the (approximate) number of neighbors of q gives

$$W_\lambda(q) \lesssim \lambda^d D^{d-1}. \tag{6.5}$$

When we multiply the bound from (6.4) by the approximate number of neighbors, we obtain an estimate for $W_\lambda(q)$ in a (possibly non-feasible) worst-case scenario in which each neighbor q' contributes as many new slabs as possible to the count $\#\{S \in \Gamma : S \cap N_\lambda(q) \neq \varnothing\}$. In reality, we probably double counted some slabs that pass through multiple neighbors of q, which means that our upper bound for $W_\lambda(q)$ in (6.5) is probably not sharp. However, the upper bound is still good enough for pigeonholing purposes.

Our upper bound for $W_\lambda(q)$ implies that there are only $\sim \log(D^{d-1})$ potential dyadic sizes for $\#\{S \in \Gamma : N_\lambda(q) \cap S \neq \varnothing\}$. It follows that we can choose some $k \lesssim \log(D^{d-1})$ so that if

$$P' = \{q \in P : 2^k \leq W_\lambda(q) < 2^{k+1}\},$$

then

$$|P'| \geq \left(\log(D^{d-1})\right)^{-1} |P| = \frac{1}{d-1}\left(\frac{1}{\log D}|P|\right).$$

In particular, $|P'| \gtrapprox_{\log} |P|$. For the rest of the proof we will work with P' instead of P. We will let W_k denote the common dyadic size of the count $W_\lambda(q)$ for the cubes in P'. (The subscript k is to remind of the fact that $W_k = 2^k$ and to differentiate this value from the parameter W that appears in the spacing conditions.) Since each unit box of P is $\sim E$ rich, we have that $|I(P,\Gamma)| \sim |P|E$ and $|I(P',\Gamma)| \sim |P'|E \gtrapprox_{\log} |I(P,\Gamma)|$.

For each $q \in P'$, we define a smooth bump function ψ_q so that $\psi_q = 1$ on q and ψ_q decays rapidly outside q with $\psi_q = 0$ outside $2q$. Similarly, for each slab $S \in \Gamma$, we define a smooth bump function ψ_S which is 1 on S and decays rapidly outside S with $\psi_S = 0$ outside $2S$. We let

$$f = \sum_{q \in P'} \psi_q$$

and

$$g = \sum_{S \in \Gamma} \psi_S.$$

Since the balls of P' are unit balls, we have that

$$|I(P', \Gamma)| \sim \int_{\mathbb{R}^d} fg.$$

(If the balls were not unit balls, we would have to renormalize to account for the volume of the balls.)

Let η_0 be a bump function so that $\eta_0 \equiv 1$ on $B(0, 1)$, η_0 is rapidly decaying outside $B(0, 1)$ and $\eta_0 \equiv 0$ outside $B(0, 2)$. Define η by

$$\eta(\omega) = \eta(\rho^{-1}\omega)$$

so that $\eta \equiv 1$ on $B(0, \rho)$, η is rapidly decaying outside $B(0, \rho)$ and $\eta \equiv 0$ outside $B(0, 2\rho)$.

By Plancherel, we have that

$$E|P'| \sim |I(P', \Gamma)| \sim \int_{\mathbb{R}^d} fg = \int_{\mathbb{R}^d} \hat{f}\,\overline{\hat{g}} = \int_{\mathbb{R}^d} \eta \hat{f}\,\overline{\hat{g}} + \int_{\mathbb{R}^d} (1-\eta)\hat{f}\,\overline{\hat{g}} =: I_1 + I_2. \tag{6.6}$$

We take cases on whether I_1 or I_2 is larger.

Low Frequency Case

First, suppose that $|I_1| \geq |I_2|$. In this case, we will show that there is a set of finitely overlapping λ-balls Q_j such that:

(i) their union contains a $\gtrapprox_{\log} 1$ fraction of the balls of P;

(ii) each Q_j intersects $\gtrapprox_{\log} \rho^{-1}E$-many slabs of Γ.

Each Q_j will be realized as the λ-neighborhood of a cube $q \in P'$. This representation is not necessarily unique. For instance if $Q_j = N_\lambda(q)$, there may be a $q' \in P'$ close to q so that the neighborhoods $N_\lambda(q)$ and $N_\lambda(q')$ are not essentially distinct from each other. If we consider the collection $\{N_\lambda(q) : q \in P'\}$ and eliminate redundancies, then the resulting sets Q_j contain a $\gtrsim 1$ fraction of the balls of P', which accounts for a $\gtrsim (\log D)^{-1}$ fraction of the balls of P. It remains to show that each Q_j intersects $\gtrapprox_{\log} \lambda E$-many slabs of Γ.

We already know that each Q_j intersects $\sim W_k$-many slabs for a dyadic number $W_k = 2^k$ that was a popular value of $W_\lambda(q)$. The pigeonhole process itself tells us nothing about the size of this popular count, but we can obtain a bound for W_k by using our assumption that $|I_1| \geq |I_2|$, which gives

$$\begin{aligned} E|P'| &\sim \int_{\mathbb{R}^d} fg = \int_{\mathbb{R}^d} \hat{f}\,\overline{\hat{g}} \lesssim \int_{\mathbb{R}^d} \eta \hat{f}\,\overline{\hat{g}} \\ &= \int_{\mathbb{R}^d} f\,(g * \check{\eta}) = \sum_{q \in P'} \sum_{S \in \Gamma} \int_{\mathbb{R}^d} \psi_q\,(\psi_S * \check{\eta}) \end{aligned} \tag{6.7}$$

For each slab S, the convolution $\psi_S * \check{\eta}$ decays rapidly off the ρ^{-1}-neighborhood of S. Meanwhile, for $x \in N_{\rho^{-1}}(S)$, we have

$$|\psi_S * \check{\eta}(x)| \lesssim \rho. \tag{6.8}$$

To see this, we write out the convolution and use the triangle inequality to give

$$|\psi_S * \check{\eta}(x)| = \left| \int_{\mathbb{R}^d} \psi_S(x-y)\check{\eta}(y)\,dy \right| \leq \int_{\mathbb{R}^d} |\psi_S(x-y)|\,|\check{\eta}(y)|\,dy.$$

A change of variables argument shows that

$$|\check{\eta}(y)| \sim \rho^d \tag{6.9}$$

for $y \in B(0, \rho^{-1})$. (Here, the implicit constant depends on our initial choice of bump function η_0 for B(0,1).) Meanwhile, $\check{\eta}$ decays rapidly off of $B(0, \rho^{-1})$. Since ψ_S is a bump function approximating 1_S, we have that

$$|\psi_S * \check{\eta}(x)| \lesssim \rho^d \operatorname{Vol}((-S + x) \cap B(0, \rho^{-1}),$$

where $-S$ is the slab obtained by sending every point in S to its antipode. A $D \times \cdots \times D \times 1$

slab intersects $B(0,\rho^{-1})$ in a set of volume $\lesssim \rho^{-(d-1)}$. Combining this fact with (6.9) yields (6.8).

Even though (6.8) holds for all x in the ρ^{-1}-neighborhood of S, only a small subset of those x make a contribution to the integral

$$\int_{\mathbb{R}^d} \psi_q(x)\,(\psi_S * \check{\eta})(x)\,dx,$$

because ψ_q is identically 1 on q but vanishes outside $2q$.

Recall that we chose P' so that the set $\{S \in \Gamma : S \cap N_\lambda(q) \neq 0\}$ had the same dyadic size, namely W_k, for each $q \in P'$. We note that S intersects the λ-neighborhood of q if and only if q intersects the λ-neighborhood of S. Since $\lambda = D^{\varepsilon^3}\rho^{-1}$, the λ-neighborhood of S contains the ρ^{-1}-neighborhood of S, outside of which $\psi_S * \check{\eta}$ decays rapidly.

Thus, for each $q \in P'$,

$$\sum_{S\in\Gamma} \int_{\mathbb{R}^d} \psi_q(s)\,(\psi_S * \check{\eta})\,(x)\,dx \lesssim \rho W_\lambda(q) \sim \rho W_k.$$

Substituting this into (6.7) gives

$$E|P'| \lesssim |P'|\rho W_k,$$

so

$$W_k \gtrsim \rho^{-1}E = D^{-\varepsilon^3}\lambda E.$$

High Frequency Case

If $|I_2| \geq |I_1|$ (which corresponds to the high frequency part of fg being larger), then we use Cauchy-Schwarz to write

$$\begin{aligned}|I(P',\Gamma)| &\lesssim \left|\int_{\mathbb{R}^d}(1-\eta)\hat{f}\,\overline{\hat{g}}\right| \leq \left(\int_{\mathbb{R}^d}(1-\eta)\left|\hat{f}\right|^2\right)^{1/2}\left(\int_{\mathbb{R}^d}(1-\eta)|\hat{g}|^2\right)^{1/2} \\ &\leq \|f\|_2\left(\int_{\mathbb{R}^d}(1-\eta)|\hat{g}|^2\right)^{1/2} \sim |P|^{1/2}\left(\int_{\mathbb{R}^d}(1-\eta)|\hat{g}|^2\right)^{1/2}.\end{aligned} \tag{6.10}$$

We claim that the integral on the far right-hand side satisfies

$$\int_{\mathbb{R}^d} (1-\eta)|\hat{g}|^2 \lesssim \lambda^{d-1} D^{d-1} |\Gamma|. \tag{6.11}$$

To prove this estimate, we group the slabs by normal direction. We partition the unit sphere $\mathbb{S}^{d-1}$ into almost-caps of radius $\sim D^{-1}$. We let Θ be the collection of almost-caps. For each $\theta \in \Theta$, let Γ_θ be the set of slabs in Γ with normal vector in θ. We define functions g_θ by

$$g_\theta = \sum_{S \in \Gamma_\theta} \psi_S.$$

Then

$$g = \sum_{\theta \in \Theta} g_\theta. \tag{6.12}$$

Using (6.12), we can rewrite (6.10) as

$$|I(P', \Gamma)| \lesssim |P'|^{1/2} \left(\int_{\mathbb{R}^d} (1-\eta) \left| \sum_{\theta \in \Theta} \widehat{g_\theta} \right|^2 \right)^{1/2} = |P'|^{1/2} \left(\int_{\mathbb{R}^d} (1-\eta) \left| \sum_{\theta \in \Theta} \sum_{S \in \Gamma_\theta} \widehat{\psi_S} \right|^2 \right)^{1/2}.$$

For each $\theta \in \Theta$, let T_θ be a D^{-1}-tube of length 1 centered at $\mathbf{0}$ with direction in θ. That is, T_θ is a tube of length 1 and radius $\frac{1}{D}$ whose axis passes through the center of θ and has midpoint $\mathbf{0}$. For any slab S with $\boldsymbol{n}(S)$ in θ, the function $\widehat{\psi_S}$ approximates a multiple of the characteristic function of T_θ, which has the dual dimensions of S. In particular, we have that

$$\left| \widehat{\psi_S}(\omega) \right| \sim |S| \sim D$$

for $\omega \in T_\theta$. Meanwhile, $\widehat{\psi_S}$ has rapid decay off of T_θ.

Motivated by this, we want to focus on just those directions that make a significant contribution to the sum in the integrand. For each ω, let Θ_ω be the collection of caps θ so that ω is contained in the D^{ε^3} dilate of T_θ. We can write

$$\begin{aligned}
\int_{\mathbb{R}^d}(1-\eta(\omega))\left|\sum_{S\in\Gamma}\hat{\psi}_S(\omega)\right|^2 d\omega &= \int_{\mathbb{R}^d}(1-\eta(\omega))\left|\sum_{\theta}\widehat{g_\theta}(\omega)\right|^2 d\omega \\
&= \int_{\mathbb{R}^d}(1-\eta(\omega))\left|\sum_{\theta\in\Theta_\omega}\widehat{g_\theta}(\omega)+\sum_{\theta\notin\Theta_\omega}\widehat{g_\theta}(\omega)\right|^2 d\omega \\
&\sim \int_{\mathbb{R}^d}(1-\eta(\omega))\left(\left|\sum_{\theta\in\Theta_\omega}\widehat{g_\theta}(\omega)\right|^2+\left|\sum_{\theta\notin\Theta_\omega}\widehat{g_\theta}(\omega)\right|^2\right) d\omega \\
&= \int_{\mathbb{R}^d}(1-\eta(\omega))\left|\sum_{\theta\in\Theta_\omega}\widehat{g_\theta}(\omega)\right|^2 d\omega \\
&\quad + \int_{\mathbb{R}^d}(1-\eta(\omega))\left|\sum_{\theta\notin\Theta_\omega}\widehat{g_\theta}(\omega)\right|^2 d\omega
\end{aligned} \tag{6.13}$$

Each function $\widehat{g_\theta}$ decays rapidly away from T_θ; in particular for any N, there is a constant C_N so that for any $\omega\notin T_\theta$, we have

$$|\widehat{g_\theta}(\omega)|\le C_N(\operatorname{dist}(\omega,T_\theta))^{-N}. \tag{6.14}$$

We will use this decay to show that the second integral on the far right-hand side of (6.13) is much smaller than our claimed bound in (6.11). If you are already convinced that this second integral can be ignored, you should skip to the paragraph after equation (6.16). However, I wanted to highlight that I am not making the argument that the second integral can always be dominated by the first one; I am merely arguing that the decay guarantees that the second integral satisfies our claimed upper bound (6.11).

To see that the second integral satisfies the upper bound from (6.11), we note that for For any ω, we have that

$$\left|\sum_{\theta\notin\Theta_\omega}\widehat{g_\theta}(\omega)\right|^2\le|\Theta\backslash\Theta_\omega|\sum_{\theta\notin\Theta_\omega}|\widehat{g_\theta}(\omega)|^2\lesssim D^{d-1}\sum_{\theta\notin\Theta_\omega}|\widehat{g_\theta}(\omega)|^2. \tag{6.15}$$

For each cap $\theta \notin \Theta_\omega$, we have

$$\mathrm{dist}(\omega, T_\theta) \geq D^{\varepsilon^3}.$$

Substituting this into (6.14) gives

$$|\widehat{g_\theta}(\omega)| \leq C_N D^{-\varepsilon^3 N}.$$

If we take N sufficiently large (e.g. $N \gtrsim \varepsilon^{-4}$) and take D large enough for our negative powers of D to compensate for the constant C_N, then we can conclude that the right-hand side of (6.15) is much less than our claimed bound in (6.11).

Thus, to prove (6.11), it remains to show that

$$\int_{\mathbb{R}^d} (1 - \eta(\omega)) \left| \sum_{\theta \in \Theta_\omega} \widehat{g_\theta}(\omega) \right|^2 d\omega \lesssim \lambda^{d-1} D^{d-1} |\Gamma|. \tag{6.16}$$

By our forthcoming Lemma 6.1.3, $|\Theta_\omega| \lesssim (D^{\varepsilon^3}\rho^{-1})^{d-1}$ for any ω with $\eta(\omega) \neq 0$. Thus, by Cauchy-Schwarz, we have that

$$\begin{aligned}
\int_{\mathbb{R}^d} (1 - \eta(\omega)) \left| \sum_{\theta \in \Theta_\omega} \widehat{g_\theta}(\omega) \right|^2 d\omega &\lesssim \int_{\mathbb{R}^d} (1 - \eta(\omega)) \, |\Theta_\omega| \left(\sum_{\theta \in \Theta_\omega} |\widehat{g_\theta}(\omega)|^2 \right) d\omega \\
&\lesssim (D^{\varepsilon^3}\rho^{-1})^{d-1} \int_{\mathbb{R}^d} (1 - \eta(\omega)) \left(\sum_{\theta \in \Theta_\omega} |\widehat{g_\theta}(\omega)|^2 \right) d\omega \\
&\leq (D^{\varepsilon^3}\rho^{-1})^{d-1} \int_{\mathbb{R}^d} \sum_{\theta} |\widehat{g_\theta}(\omega)|^2 \, d\omega \\
&= (D^{\varepsilon^3}\rho^{-1})^{d-1} \int_{\mathbb{R}^d} \sum_{\theta} |g_\theta(x)|^2 \, dx
\end{aligned}$$

In the two last lines, the sum is over all θ, not just the caps θ in the significant set Θ_ω. (We temporarily ignored the θ not in Θ_ω in order to have a smaller number of terms for Cauchy-Schwarz, then put them back after the Cauchy-Schwarz was complete.) For each θ, the slabs of Γ_θ have bounded overlap, so

$$(D^{\varepsilon^3}\rho^{-1})^{d-1}\int_{\mathbb{R}^d}\sum_\theta |g_\theta(x)|^2\,dx \lesssim (D^{\varepsilon^3}\rho^{-1})^{d-1}\sum_{S\in\Gamma}\int_{\mathbb{R}^d}|\psi_S(x)|^2\,dx \sim (D^{\varepsilon^3}\rho^{-1})^{d-1}|\Gamma|D^{d-1}.$$

We recall that we had defined ρ by $\rho = D^{\varepsilon^3}\lambda^{-1}$. Thus,

$$(D^{\varepsilon^3}\rho^{-1})^{d-1} = \lambda^{d-1},$$

which establishes (6.11).

Combining this with the estimate $|P| \lessapprox_{\log} |P'|$ and with (6.6), we conclude that

$$E|P|^{1/2} \lessapprox_{\log} \left(\lambda^{d-1}D^{d-1}|\Gamma|\right)^{1/2}$$

which we rearrange to give

$$|P| \lessapprox_{\log} E^{-2}\lambda^{d-1}|\Gamma|D^{d-1}.$$

Thus, the case that $|I_2| \geq |I_1|$ corresponds to outcome (1) in the statement of the proposition.

□

Lemma 6.1.3. *Let $D > 1$, and let Θ be a partition of the unit sphere $\mathbb{S}^{d-1}$ into almost-caps θ of radius $\frac{1}{D}$. Let $\mathbb{T}$ be a collection of essentially distinct tubes through $\mathbf{0}$ with radius D^{-1} and length 1 so that for each θ, $\mathbb{T}$ contains one tube T_θ with direction in the cap θ. Let $0 < \rho \leq 1$, and let $1 < \sigma < \rho$ be a parameter which may depend on D. If $\|\omega\| \geq \rho$, then the number of caps θ so that ω is in σT_θ, the dilate of T_θ by a factor of σ, satisfies*

$$\#\left\{\theta \in \Theta : \omega \in \sigma T_\theta\right\} \lesssim \left(\frac{\sigma}{\rho}\right)^{d-1}.$$

Proof. The direction vectors of the tubes in $\mathbb{T}$ form a maximal $\frac{1}{D}$-separated set, so each point on the sphere $\rho\mathbb{S}^{d-1}$ lies in approximately the same number of dilated tubes, namely the number given by the quotient

$$\frac{\sum_{\theta\in\Theta} \mathrm{Vol}^{d-1}(\sigma T_\theta \cap \rho\mathbb{S}^{d-1})}{\mathrm{Vol}^{d-1}(\rho\mathbb{S}^{d-1})}.$$

For each θ, the enlarged tube σT_θ intersects the sphere $\rho\mathbb{S}^{d-1} \subset \mathbb{R}^d$ in a cap of radius $\sim \sigma/D$, which has $(d-1)$-dimensional volume $\sim (\frac{\sigma}{D})^{d-1}$.

Thus, we have that

$$\begin{aligned}\#\{\theta \in \Theta : \omega \in \sigma T_\theta\} &\lesssim \frac{\sum_{\theta\in\Theta} \mathrm{Vol}^{d-1}(\sigma T_\theta \cap \rho\mathbb{S}^{d-1})}{\mathrm{Vol}^{d-1}(\rho\mathbb{S}^{d-1})} \sim \frac{\sum_{\theta\in\Theta}(\sigma D^{-1})^{d-1}}{\rho^{d-1}} \\ &= \frac{|\Theta|\,(\sigma D^{-1})^{d-1}}{\rho^{d-1}} \sim \frac{D^{d-1}(\sigma D^{-1})^{d-1}}{\rho^{d-1}} = \left(\frac{\sigma}{\rho}\right)^{d-1}.\end{aligned}$$

□

Chapter 7

The Thin Bound and Our Goal Bounds

In this chapter we will describe a 'naive' or 'obvious' approach to proving bounds for $P_r(\Gamma)$ from the thin bound of Proposition 6.1.1. This approach works under the first spacing condition for all values of W but does not work under the second spacing condition unless W is very large. This is why we introduce partitioning when working with the second spacing condition.

If we substitute $E = r$ and $D = \delta^{-1}$ into the our bound from the thin case in Proposition 6.1.1, we get the bound

$$|P| \lessapprox_{\log} \lambda^{d-1} r^{-2} |\Gamma| \delta^{-(d-1)}, \tag{7.1}$$

where $\lambda = \delta^{-\frac{\varepsilon}{10d}}$. In the ensuing discussion, we will replace (7.1) with the shorthand

$$|P| \lessapprox r^{-2} |\Gamma| \delta^{-(d-1)}, \tag{7.2}$$

as we are mainly concerned with achieving upper bounds for $|P|$ that have the 'right' exponents for each of r and $|\Gamma|$.

7.1 An Analogy and a Plan of Attack

In various stages of writing up my work, I have found the following analogy helpful in organizing my thoughts. I am sharing it in hopes that it will also be helpful to you, the reader. However, if you think it is silly, or if this doesn't feel like a good time for big picture thinking, feel free to skip ahead to the subsection 'A Plan of Attack' or even to Section 7.2.

The Illustration

Imagine that you are given a collection of blocks of various sizes and you want to put them into a box and ship them to a friend. You have a box at home, and you aren't sure whether it will be big enough. Maybe you start putting blocks into your box anyway. If you weren't very sensible, you might spend a lot of time arranging the small blocks in neat rows, only to later discover that the biggest block didn't fit in the box.

Maybe the specific way you packed the small blocks prevented you from being able to fit the big block, but maybe the big block would have never fit in the box, even if the box had been otherwise empty. In this case, you would have wasted a lot of time packing the small blocks into the box. You now wish that you had first determined whether the biggest block would fit. But you didn't do it, and now it's too late!

This is a bit of a silly illustration because, barring a very tight fit, it is perhaps not so hard to tell if a block will fit in a box. However, the reality it is illustrating is not so obvious.

The Illustrated

What I intend to illustrate with the blocks and the boxes is the following situation: we want to know whether the product of many terms - each of which depends on multiple parameters - obeys a certain upper bound or not. Perhaps we have a general idea of which terms are the biggest. Say, for concreteness, that $A \gtrsim B \gtrsim C \geq 1$ and we want to know whether

$$ABC \leq U. \tag{7.3}$$

I want to emphasize that the A, B and C here are functions of multiple parameters, even though A, B, and C denote constants in other chapters.

Maybe we are sort of suspicious of this purported bound. If we were wanting to *disprove* it, then our best bet would be to try to prove that $A > U$. If it turns out, though, that $A \leq U$, then maybe we should be a little bit less suspicious and should be more willing to try to prove (7.3), even if the proof is complicated.

Determining whether term A exceeds U is analogous to determining whether the biggest block fits in the box. If we have managed to fit the biggest block into the box, it is worth trying to fit the other blocks in around it. On the other hand, if the biggest block didn't fit, it is time to go shopping for another box. That is, if $A > U$, it is time to make a new conjecture.

As one final note, I should mention that when I am thinking of estimates in this framework, it is always the case that the product ABC that I am estimating is itself an upper bound for $|P_r(\Gamma)|$ and may not be an optimal one. If the inequality $ABC \leq U$ is false, this doesn't necessarily mean that the inequality $|P_r(\Gamma)| \leq U$ is false. Instead, it might be the case that the product ABC was not a good guess for how big $P_r(\Gamma)$ was. In this case, we may be able to refine our guess - say, by replacing A with a function $A' \leq A$. (We may also replace B and C with better guesses, but if the inequality $A \leq U$ is false, then A is the obvious culprit that needs to be replaced.)

This idea of rejecting a 'bad' guess in favor of a better one is precisely what we do when trying to estimate $|P_r(\Gamma)|$ under the second spacing condition in Section 7.3. The 'bad' guess we reject is one that results from trying to follow our work under the first spacing condition as closely as possible. The better guess is one we arrive at using partitioning. The argument also requires us to assume an inductive hypothesis. And even after that assumption, all that Section 7.3 does in the metaphorical scheme of things is to prove that we can fit the biggest block into the box.

In some other settings (i.e. when Γ is not in the thin case of Proposition 6.1.1), our first guess upper bound for $P_r(\Gamma)$ may include a block of intermediate size between $\delta^{-\frac{\varepsilon}{10d}}$ and $\log(\delta^{-1})$. For instance, our work in Chapter 10 includes a block of size about $\delta^{-\varepsilon^3}$.

7.1.1 A Plan of Attack

Under the first spacing condition, we want to prove a bound of the form

$$|P_r(\Gamma)| \leq C_\varepsilon \delta^{-\varepsilon} |\Gamma|^d r^{-d} W^{-(d-1)}.$$

If $|P_r(\Gamma)\backslash P_{2r}(\Gamma)|$ accounts for most of $|P_r(\Gamma)|$ and Γ is in the thin case of Proposition 6.1.1, then we know that

$$|P_r(\Gamma)| \leq C \log(\delta^{-1})^{O(1)}\, \lambda^{d-1} r^{-2} |\Gamma| \delta^{-(d-1)}.$$

for some dimensional constant C. We want to know: For δ sufficiently small, do we have

$$C \log(\delta^{-1})^{O(1)}\, \lambda^{d-1} r^{-2} |\Gamma| \delta^{-(d-1)} \leq C_\varepsilon\, \delta^{-\varepsilon} |\Gamma|^d r^{-d} W^{-(d-1)} \tag{7.4}$$

for the specific exponent implied by the $O(1)$? Or will we have to combine the thin case estimate with some other argument?

In determining whether (7.4) holds for sufficiently small δ, we begin by estimating the product $r^{-2}|\Gamma|\delta^{-(d-1)}$ which comprises the right-hand side of (7.2).

What we show in section 7.2 is that

$$r^{-2}|\Gamma|\delta^{-(d-1)} \lesssim |\Gamma|^d r^{-d} W^{-(d-1)}.$$

In order to prove (7.4), we will still have to estimate the product

$$\log(\delta^{-1})^{O(1)}\, \lambda^{d-1}.$$

We have a $\delta^{-\varepsilon}$-worth of room in which to fit this product.

To relate this to the framework I described under the subheading 'The Illustrated,' we can set

$$A = r^{-2}|\Gamma|\delta^{-(d-1)}.$$

The expression $r^{-2}|\Gamma|\delta^{-(d-1)}$, when multiplied by the product $\log(\delta^{-1})^{O(1)}\, \lambda^{d-1}$, gives an upper bound for $P_r(\Gamma)$. We can set

$$B = \lambda^{d-1} = \delta^{-\frac{(d-1)\varepsilon}{10d}}$$

Then the block C will be a constant multiple of the $\log(\delta^{-1})^{O(1)}$ term. However, we are getting a bit ahead of ourselves, as in this chapter, we are just concerned with the biggest block, metaphorically speaking.

7.2 Thin Bound vs Goal Bound under the First Spacing Condition

We will apply Proposition 5.1 to the set

$$P := P_r(\Gamma)\backslash P_{2r}(\Gamma),$$

provided that

$$|P| \geq \frac{1}{10}|P_r|. \tag{7.5}$$

(If this condition is not met, we will induct on r; see Chapter 9.)

The r^{-2} in the thin bound (7.2) came from the fact that when the high frequency case was dominant in our incidence-counting integral, the characteristic functions for our slab duals were essentially orthogonal. We want to upgrade to a more negative power of r. This comes at the expense of having a larger positive power of $|\Gamma|$ in our estimate.

By our assumption that the slabs of Γ are essentially distinct, we can assume that $r \lesssim \delta^{-(d-1)}$, which means that $r^{-(d-2)} \gtrsim \delta^{-(d-1)(d-2)}$ Thus, we have that

$$\begin{aligned} r^{-2}\delta^{-(d-1)}|\Gamma| &= r^{-2}\delta^{(d-1)(d-2)}\left(\delta^{-(d-1)(d-2)}\delta^{-(d-1)}\right)|\Gamma| \\ &\lesssim r^{-(d+1)}\delta^{-(d-1)(d-1)}|\Gamma|. \end{aligned}$$

Meanwhile, under spacing condition (WS-1) with parameter W, we have that $|\Gamma| \sim N\delta^{-(d-1)}W$, which implies that

$$\delta^{-(d-1)} \sim N^{-1}W^{-1}|\Gamma| \leq W^{-1}|\Gamma| \tag{7.6}$$

We take both sides of (7.6) to the power $d-1$ and substitute the result into (9.8) to give

$$|P| \lessapprox r^{-(d+1)}W^{-(d-1)}|\Gamma|^d.$$

The fact that the exponents for r, Γ and W resulting from this work coincided with the exponents in the bound from the stacked stars example is what led me to conjecture that Theorem 1.4.2 held. Of course, this work alone does not constitute a proof of 1.4.2. For one thing, we must show that the ε-loss suppressed by the $\lessapprox$ symbol is acceptable; we accomplish this in Chapter 8. However, we must also address what happens if our collection of slabs was in the thick case of Proposition 6.1.1 or if we could not apply this proposition in the first place. Each of these cases requires an inductive argument, as outlined in Chapter 5. The details are in Chapter 9.

7.3 Thin Bound vs Goal Bound Under the Second Spacing Condition

Again, let

$$P := P_r(\Gamma) \backslash P_{2r}(\Gamma).$$

Assuming that

$$|P| \geq \frac{1}{10}|P_r|, \tag{7.7}$$

we will (try to) apply Proposition 6.1.1 to P with $E \sim r$ and $D \sim \delta^{-1}$.

Starting from the thin bound, we have

$$\begin{aligned} |P| &\lessapprox r^{-2}|\Gamma|\delta^{-(d-1)} \\ &= \left(r^{-2}r^{-(d-1)}\right) r^{d-1} \left(|\Gamma||\Gamma|^{d-1}\right) \Gamma^{-(d-1)}\delta^{-(d-1)} \\ &= r^{-(d+1)}|\Gamma|^d \left(r^{d-1}|\Gamma|^{-(d-1)}\delta^{-(d-1)}\right). \end{aligned}$$

We note that

$$r^{d-1}|\Gamma|^{-(d-1)}\delta^{-(d-1)} \sim r^{d-1}W^{-d(d-1)}\delta^{-(d-1)}.$$

From the second spacing condition we also have the inequality $r \lesssim W^{d-1}$. This gives

$$r^{d-1}|\Gamma|^{-(d-1)}\delta^{-(d-1)} \lesssim (W^{d-1})^{d-1}W^{-d(d-1)}\delta^{-(d-1)} = W^{-(d-1)}\delta^{-(d-1)}.$$

We note that $W^{-(d-1)}\delta^{-(d-1)} \lessapprox 1$ only if $W \gtrapprox \delta^{-1}$. Thus the 'naive' approach (namely, applying the thin bound for $D = \delta^{-1}$ and using the bound for r from the spacing condition) would work only for a very small range of W consisting of only those W which are almost as big as possible.

To remedy this, I intend to use partitioning, as described below. The partitioning argument below recovers the right exponents for $|\Gamma|$ and r under the assumption that we can apply the conjectured bound at another scale. However, checking the hypotheses at a new scale presents additional challenges.

7.3.1 The Idea of Partitioning

Instead of applying Proposition 6.1.1 directly to P, we will now subdivide $[0,1]^d$ into many sub-cubes and (attempt to) use Proposition 6.1.1 to characterize the δ-balls of P that lie within each sub-cube. Specifically, for a parameter D to be determined, we subdivide $[0,1]^d$ into sub-cubes of side length $\sim D\delta$.

For a fixed $D\delta$-cube Q, let

$$\Gamma \cap Q = \{S \cap Q : S \in \Gamma, S \cap Q \neq \varnothing\}.$$

The intersection $S \cap Q$ is approximately a rectangular prism of dimensions $\sim D\delta \times \cdots \times D\delta \times \delta$ in the sense that it contains a dilate (by a small $c < 1$) of a $D\delta \times \cdots \times D\delta \times \delta$ prism and is contained in a dilate (by a large number $C > 1$) of a $D\delta \times \cdots \times D\delta \times \delta$ prism. We call the intersection $S \cap Q$ a **slab segment**.

We note that the slab segments in $\Gamma \cap Q$ may not be essentially distinct; for instance, there may be $S_1, S_2 \in \Gamma$ so that $S_1 \cap Q$ and $S_2 \cap Q$ essentially contain each other. To eliminate redundancies, let Γ_Q be a maximal essentially distinct subset of the collection $\Gamma \cap Q$ above.

We further prune Γ_Q by pigeonholing on the (dyadic) number of slabs 'represented' by each slab segment of $\Gamma \cap Q$. For each dyadic M, let $\Gamma_{Q,M}$ denote the collection of slab segments in $\Gamma \cap Q$ which are each equivalent to the segment $S \cap Q$ for $\sim M$ many different S. We can choose M_0 to preserve a $\gtrapprox_{\log} 1$ fraction of the incidences between P and Γ, i.e. we choose M_0 so that

$$\sum_Q M_0 |I(P \cap Q, \Gamma_{Q,M_0})| \gtrsim_d (\log(\delta^{-1}))^{-1} |I(P,T)|. \tag{7.8}$$

(Here, the set $I(P \cap Q, \Gamma_{Q,M_0})$ is defined in terms of essential intersection, as in Definition 1.1.6 and Equation (1.6).)

We similarly prune P. For each sub-cube Q and each dyadic E, we let

$$P_{Q,E} := \left\{ q \in P_Q : q \lesssim\!\!\!\subset S_Q \text{ for } \sim E\text{-many } S_Q \text{ in } \Gamma_{Q,M_0} \right\},$$

where

$$P_Q := P \cap Q = \{q \in P : q \subsetsim Q\}.$$

Then we can find a dyadic E_0 so that

$$\sum_Q M_0 E_0 \, |P_{Q,E_0}| \gtrsim \left(\log\left(\delta^{-1}\right)\right)^2 |I(P,\Gamma)|. \tag{7.9}$$

Because of the angle separation in (WS-2), we must have $E_0 \lesssim D^{d-1}$. Our procedure for choosing M_0 and E_0 to satisfy (7.8) and (7.9) will also ensure that

$$M_0 E_0 \gtrsim (\log(\delta^{-1}))^2 r. \tag{7.10}$$

However, we omit the details from this discussion.

We will apply Proposition 6.1.1 to each set P_{Q,E_0}. If we let $\lambda = D^{\varepsilon/(10d)}$, this gives the following two possibilities:

1. $|P_{Q,E_0}| \lessapprox_{\log} \lambda^{d-1} E_0^{-2} D^{n-1} |\Gamma_{Q,M_0}|$; or
2. There is a collection of $2\delta\lambda$-balls whose union contains a $\gtrapprox_{\log} 1$ fraction of P_{Q,E_0} so that each of these $2\delta\lambda$-balls intersects $\gtrapprox_d \lambda E_0$-many slab segments in Γ_{Q,M_0}.

If a $D\delta$-cube Q satisfies the first possibility, we say it is thin. Otherwise, we say it is thick. We can write

$$|P| \lesssim \log D \sum_Q |P_{Q,E_0}| = \log D \left(\sum_{Q \text{ thin}} |P_{Q,E_0}| + \sum_{Q \text{ thick}} |P_{Q,E_0}| \right). \tag{7.11}$$

We take cases on whether the sum over thin cubes or the sum over thick cubes is larger. If the sum over thin cubes is larger, we say we are in the thin case.

For the rest of this section, we will discuss only the case that the sum over thin cubes is larger. In this case, we have that

$$|P| \lessapprox_{\log} \sum_{Q \text{ thin}} |P_{Q,E_0}| \lessapprox_{\log} \lambda^{d-1} E_0^{-2} D^{d-1} \sum_{Q \text{ thin}} |\Gamma_{Q,M_0}|.$$

We now want to estimate $\sum_{Q \text{ thin}} |\Gamma_{Q,M_0}|$. For this, note that if we stretched the thin direction of a slab segment, then it would become a box that was $\sim M_0$-rich for dilates of our original slabs, and we could (try to) estimate the number of rescaled slab segments using induction. However, different slab segments need to be stretched in different directions to make them into M_0-rich boxes. To sort the slab segments by the direction that needs to be stretched, we cover the sphere $\mathbb{S}^{d-1}$ by caps τ of radius $1/D$. For each cap τ, we subdivide the unit ball into cells $\square_{\tau,j}$, where each cell $\square_{\tau,j}$ is a slab of dimensions $\sim 1 \times \cdots \times 1 \times \frac{1}{D}$, with normal direction defined by τ. Here, j will range from 1 to J_τ with $J_\tau \sim D$, because the number of D^{-1}-slabs with normal direction in τ is $\sim \frac{\text{Vol}(B(\mathbf{0},1))}{\text{Vol}(D^{-1}\text{-slab})} \sim \frac{1}{1/D} = D$. Summing over all τ, we conclude that the total number of cells is $\sim D^d$.

For each cell $\square_{\tau,j}$, let $\Gamma_{\tau,j}$ denote the set of slabs $S \in \Gamma$ that are essentially contained in $\square_{\tau,j}$. After fixing a cell $\square_{\tau,j}$, we rescale $\square_{\tau,j}$ to occupy the entire unit ball. (This corresponds to a dilation by a factor of D in one direction.) In the process, $\Gamma_{\tau,j}$ is rescaled to a collection of thicker slabs with thickness $\sim D\delta$. We denote the rescaled collection by $\widetilde{\Gamma_{\tau,j}}$. If we let $\tilde{r} = M_0$ and $\tilde{\delta} = D\delta$, then rescaling a cell transforms each slab segment of Γ_{Q,M_0} that was essentially contained in the cell into a(n) (approximate) $\tilde{r}$-rich $\tilde{\delta}$-box for the collection $\widetilde{\Gamma_{\tau,j}}$. Assuming that the premises of the inductive hypothesis are met, the number of M_0-rich $\tilde{\delta}$-boxes satisfies

$$\left|P_{M_0}\left(\widetilde{\Gamma_{\tau,j}}\right)\right| \leq \tilde{\delta}^{-\varepsilon} M^{-(d+1)} |\tilde{\Gamma}|^d. \tag{7.12}$$

Summing this bound over all cells $\square_{\tau,j}$ gives

$$\begin{aligned}
\sum_{Q \text{ thin}} |\Gamma_{Q,M_0}| &\lesssim (\#\text{ of cells }) \left(\max_{\square_{\tau,j}} \left|P_{M_0}\left(\widetilde{\Gamma_{\tau,j}}\right)\right|\right) \sim D^d \left(\max_{\square_\tau} \left|P_{M_0}\left(\widetilde{\Gamma_{\tau,j}}\right)\right|\right) \\
&\lesssim D^d \left(\tilde{\delta}^{-\varepsilon} M_0^{-(d+1)} |\tilde{\Gamma}|^d\right) \\
&\sim D^d \left((D\delta)^{-\varepsilon} M_0^{-(d+1)} (|\Gamma| D^{-d})^d\right).
\end{aligned}$$

Combining this with the thin case estimate, we get

$$\begin{aligned}
|P_r(\Gamma)| &\lessapprox_{\log} \lambda^{d-1} E_0^{-2} D^{d-1} \left(\sum_{Q \text{ thin}} |\Gamma_{Q,M_0}| \right) \\
&\lesssim \lambda^{d-1} E^{-2} D^{d-1} \left(D^d D^{-\varepsilon} \delta^{-\varepsilon} M_0^{-(d+1)} |\Gamma|^d D^{-d^2} \right) \\
&\sim (D^{\varepsilon'} D^{-\varepsilon}) \delta^{-\varepsilon} |\Gamma|^d \left(E_0^{-2} M_0^{-(d+1)} D^{2d-1-d^2} \right) \\
&= (D^{\varepsilon'} D^{-\varepsilon}) \delta^{-\varepsilon} |\Gamma|^d \left(E_0^{-2} M_0^{-(d+1)} \left(D^{-(d-1)} \right)^{d-1} \right) \\
&\leq (D^{\varepsilon'} D^{-\varepsilon}) \delta^{-\varepsilon} |\Gamma|^d \left(E_0^{-2} M_0^{-(d+1)} E^{d-1} \right) \\
&= (D^{\varepsilon'} D^{-\varepsilon}) \delta^{-\varepsilon} |\Gamma|^d \left(M_0^{-(d+1)} E_0^{-(d+1)} \right) \\
&\approx_{\log} |\Gamma|^d r^{-(d+1)}.
\end{aligned} \tag{7.13}$$

Through this argument, we have achieved the 'right' exponents for r and for $|\Gamma|$. However, the argument relied on an assumption that for each τ and each $j \leq J_\tau$, the collection $\widetilde{\Gamma_{\tau,j}}$ satisfied the premises of the inductive hypothesis. To be justified in applying the inductive hypothesis, we'd have to first check that each collection of thick slabs satisfied (WS-2) with parameter $\tilde{W} := \frac{W}{D}$, that it satisfied broadness assumption of our conjecture and that $\tilde{r} := M_0$ satisfied the lower bound

$$\tilde{r} \geq \max\left\{ \tilde{\delta}^{-\varepsilon/4} \tilde{\delta} |\tilde{\Gamma}|, d \right\}.$$

One complication is that we'll have $\tilde{r} \geq d$ only if our original r was $\gtrapprox D^{d-1}$. This will necessitate a separate argument for small r, as in Section 4 of [10]. However, the narrow case in subsection 4.1 does not have an analogue for slabs. This failure of the narrow case, along with Example 2.0.1, was why we introduced broadness to the hypotheses of our conjecture under the second spacing condition.

Chapter 8

The Thin Bound and Our Goal Bounds, Part II

In the previous chapter, we showed that in the thin case under the first spacing condition, we have an estimate of the form

$$|P_r(\Gamma)| \lessapprox_{\log} \lambda^{d-1} r^{-d} |\Gamma|^d W^{-(d-1)}. \tag{8.1}$$

In this chapter, we will demonstrate how to show that (8.1) implies the bound of Theorem 1.4.2.

We can rewrite (8.1) as

$$\begin{aligned}|P_r(\Gamma)| &\leq A\left(\log(\delta^{-1})\right)^B \, (\delta^{-\frac{\varepsilon}{10d}})^{d-1} r^{-d} |\Gamma|^d W^{-(d-1)} \\ &\leq A\left(\log(\delta^{-1})\right)^B \, \delta^{\varepsilon(1-\frac{d-1}{10d})} \, \delta^{-\varepsilon} r^{-d} |\Gamma|^d W^{-(d-1)}\end{aligned}$$

for some constants A and B.

Meanwhile, Theorem 1.4.2 asserts that

$$|P_r(\Gamma)| \leq C_\varepsilon r^{-d} |\Gamma|^d W^{-(d-1)}.$$

The bound of Theorem 1.4.2 was stated with a constant depending on ε, but the presence of the constant was to ensure that the bound holds for large values of δ, as discussed in

Chapter 5. What we'll show in this chapter is that we can choose $c_\varepsilon > 0$, along with a constant C which depends only on d, so that for any $\delta \in (0, c_\varepsilon)$ and any collection of δ-slabs satisfying (WS-1) we have

$$A \left(\log(\delta^{-1})\right)^B \delta^{\varepsilon(1-\frac{d-1}{10d})} \delta^{-\varepsilon} r^{-d} |\Gamma|^d W^{-(d-1)} \leq C r^{-d} |\Gamma|^d W^{-(d-1)}. \tag{8.2}$$

Thus, for $\delta \in (0, c_\varepsilon)$ we can conclude that any set of δ-slabs in the thin case of Proposition 6.1.1) must obey the theorem bound, provided that the constant C_ε of Chapter 5 was chosen to exceed the constant C above.

To prove that (8.2) holds for all δ sufficiently small, we must prove that, for an appropriate choice of C, we have

$$\left(\log(\delta^{-1})\right)^B \delta^{\varepsilon(1-\frac{d-1}{10d})} \leq \frac{C}{A}.$$

Taking logs, we see that this is equivalent to the inequality

$$B \log\log(\delta^{-1}) \leq \log\left(\frac{C}{A}\right) + \varepsilon\left(1 - \frac{(d-1)}{10d}\right)\log(\delta^{-1}). \tag{8.3}$$

If $C > A$, then each of the inequalities

$$B \log\log(\delta^{-1}) \leq \varepsilon\left(1 - \frac{(d-1)}{10d}\right)\log(\delta^{-1}) \tag{8.4}$$

and

$$B \log\log(\delta^{-1}) \leq \log\left(\frac{C}{A}\right). \tag{8.5}$$

is a sufficient condition for (8.3). The first of these, inequality (8.4), is equivalent to

$$\frac{\log\log(\delta^{-1})}{\log(\delta^{-1})} \leq \frac{\varepsilon\left(1 - \frac{(d-1)}{10d}\right)}{B} \tag{8.6}$$

Since $\frac{\log\log(\delta^{-1})}{\log(\delta^{-1})}$ approaches 0 as $\delta \to 0$, there is a constnat c_ε so that 8.6 holds for all $\delta \leq c_\varepsilon$, which means that (8.2) also holds for all $\delta \leq c_\varepsilon$, provided that we chose $C > A$.

Chapter 9

Proof of Theorem 1.4.2

This chapter is devoted to proving Theorem 1.4.2, which we restate below for convenience.

Theorem 9.0.1. *(Theorem 1.4.2, revisited) For any $\varepsilon > 0$ sufficiently small (relative to d), there exists a constant $C_\varepsilon = C(\varepsilon, d) > 1$ so that if $\delta \in (0,1)$ and $1 \le W \le \delta^{-1}$, then*

$$|P_r(\Gamma)| \le C_\varepsilon \delta^{-\varepsilon} W^{-(d-1)} r^{-d} |\Gamma|^d \tag{9.1}$$

for any collection Γ of δ-slabs that satisfies the first spacing condition for the parameter W and any r with

$$r \ge \delta^{-\frac{\varepsilon}{4}} \delta |\Gamma|.$$

The proof is by downward induction on δ and/or r. Our inductive hypothesis is that for any pair $\left(\tilde{\delta}, \tilde{r}\right)$ with $W > \tilde{\delta} \ge \delta^{1-\frac{\varepsilon}{10d}}$ or $\tilde{r} \ge 2r$, we have

$$\left| P_{\tilde{r}}(\tilde{\Gamma}) \right| \le C_2(\varepsilon) \tilde{\delta}^{-\varepsilon} \frac{|\tilde{\Gamma}|^d}{r^d W^{d-1}} \tag{9.2}$$

for any collection of $\tilde{\delta}$-slabs that satisfies the first spacing condition for the parameter W and any $\tilde{r}$ with $\tilde{r} \ge (\tilde{\delta})^{-\frac{\varepsilon}{4}} |\Gamma|$.

It is important to note that at the inductive step, we will use the bound (9.2) for the same value of ε for which we are trying to prove (9.1). (If we wanted to use the inductive hypothesis for $\varepsilon' \ne \varepsilon$, this would introduce a factor of $C_{\varepsilon'}$ to the resulting bound for $|P_r(\Gamma)|$.

Perhaps this is something we could deal with, but it would require us to track much more carefully how C_ε varies with ε.)

The argument relies on Proposition 6.1.1 to pass between scales. We will apply this proposition for $D = \delta^{-1}$ and $P = P_r(\Gamma)\backslash P_{2r}(\Gamma)$, provided that

$$|P_r(\Gamma)\backslash P_{2r}(\Gamma)| \geq \frac{1}{10}|P_r(\Gamma)|. \tag{9.3}$$

Passing to a subset $P \subset P_r(\Gamma)$ is necessary to ensure that the hypotheses of the proposition are met. We are also implicitly assuming that $D = \delta^{-1}$ is large enough for the conclusion of the proposition to hold. We can accomplish this by adjusting the constant c_ε of the first base case if needed.

If we are in the thin case from Proposition 6.1.1, then we can prove the bound (1.8) directly, as explained in Chapters 7-8. We have included a shorter version of the argument in subsection 9.3.1.

If we are in the thick case, this means that a large fraction of the δ-balls in $P_r(\Gamma)$ occur in clusters. Each cluster consists of many δ-boxes which are all (essentially) contained inside a larger box of side length $\lambda\delta$ for $\lambda = (\delta^{-1})^{\frac{\varepsilon}{10d}}$. We use our inductive hypothesis to estimate the number of clusters.

The thin/thick dichotomy described above only works if (9.3) holds. The possibility that this condition does not hold is why we must allow in our inductive hypothesis for the case that $\tilde{r} > r$ but $\tilde{\delta}$ does not exceed δ. If (9.3) does not hold, then we apply the inductive hypothesis with $\tilde{r} = 2r$ and $\tilde{\delta} = \delta$, and, in the course of proving that (1.8) holds, we prove the intermediate result

$$|P_r(\Gamma)| \leq \frac{10}{9}|P_{2r}(\Gamma)|.$$

9.1 Base Cases

9.1.1 Base Case 1

As a first base case assume that $\delta \geq c_\varepsilon$, where c_ε is chosen to satisfy the requirements of Chapter 8. Our work in Chapter 5 shows that, so long as we chose C_ε well, we have

$$|P_r(\Gamma)| \leq C_\varepsilon \delta^{-\varepsilon} r^{-(d+1)} |\Gamma|^d$$

for all $\delta \geq c_\varepsilon$.

9.1.2 Base Case 2

By Lemma 5.1.1, there is a dimensional constant α_d so that if

$$r \geq \alpha_d \delta^{-(d-1)}$$

then $P_r(\Gamma) = \varnothing$. This is because a set of essentially distinct δ-slabs all incident to a common δ-box cannot have have a size exceeding the size of a maximal δ-separated set on $\mathbb{S}^{d-1}$.

In Section 9.2, we show that if (9.3) fails, then

$$|P_r(\Gamma)| \leq \frac{10}{9} |P_{2r}(\Gamma)|. \tag{9.4}$$

We note that if $2r$ exceeds the number of directions in a maximal δ-separated set on $\mathbb{S}^{d-1}$ but r does not, then (9.4) cannot hold. What this means is that, in fact, $P_r(\Gamma) \backslash P_{2r}(\Gamma)$ must have accounted for most of $P_r(\Gamma)$, which means that we can apply Proposition 6.1.1.

9.1.3 Base Case 3

Base Case 3a

By Lemma 5.1.4, there is a dimensional constant β_d so that if

$$W \geq \beta_d \, \delta^{-1+\frac{\varepsilon}{10d}}, \tag{9.5}$$

then $P_r(\Gamma) = \varnothing$.

Base Case 3b

If δ is sufficiently small relative to ε, then the dimensional constant β_d is unnecessary; that is, for $\delta \lesssim_\varepsilon 1$, we have $P_r(\Gamma) = \varnothing$ for any Γ that satisfies (WS-2) for

$$W \geq \beta_d \, \delta^{-1+\frac{\varepsilon}{10d}} \tag{9.6}$$

and any $r \geq \delta^{-\varepsilon} 4\delta|\Gamma|$.

We note each of (9.5) and (9.6) is equivalent to a lower bound for δ in terms of W Since the inductive procedure of section 9.3 increases δ but keeps W fixed, we can arrive in these base cases by repeatedly carrying out the inductive procedure of the thick case.

We have included Base Case 3a in addition to Base Case 1, because for some values of W and ε we will have $c_\varepsilon(\varepsilon) \geq (\beta_d^{-1} W^{-1})^{\frac{10d}{10d-\varepsilon}}$, but for others we may not. Meanwhile, Base Case 3b is used in Section 9.3.3 in verifying that the conditions of the theorem are still met when we pass to another scale.

As with Base Case 2, we do not literally iterate our inductive argument of the thick case until δ is big enough for (9.5) or (9.6) to hold. Instead, we can conclude that before we performed the inductive step 'too many times,' we would have produced a collection of slabs for which the conclusion of the thin case held.

9.2 Increasing r

Recall that $P_r(\Gamma)$ is defined to be the set of δ-balls that are *at least* r-rich for Γ. We define $P_r'(\Gamma) \subset P_r(\Gamma)$ by

$$P_r'(\Gamma) = \{q \in P_r(\Gamma) : \ q \text{ is at least } r\text{-rich but is } \textit{not } 2r\text{-rich}\}.$$

Our ability to (essentially) reverse the inequality $|P_{2r}(\Gamma)| \leq |P_r(\Gamma)|$ will depend on whether or not $P_r'(\Gamma)$ accounts for a large proportion of $P_r(\Gamma)$. Specifically, we will take cases on whether

$$|P_r'(\Gamma)| \geq \frac{1}{10}|P_r(\Gamma)| \tag{9.7}$$

If inequality (9.7) holds, then we prove the goal bound (9.1) by induction on δ, as explained in Section 9.3. If (9.7) does not hold, then we apply the inductive hypothesis for $\tilde{\delta} = \delta$ and $\tilde{r} = 2r$. This gives

$$|P_r(\Gamma)| \leq \frac{10}{9}|P_{2r}(\Gamma)| \lesssim_\varepsilon \delta^\varepsilon \frac{|\Gamma|^d}{(2r)^{d+1}} \leq \delta^\varepsilon \frac{|\Gamma|^d}{r^{d+1}}.$$

If we wanted to think of our argument in terms of iteration rather than induction, we could iterate the inequality

$$|P_r(\Gamma)| \le \frac{10}{9}|P_{2r}(\Gamma)|$$

many times to give

$$|P_r(\Gamma)| \le \frac{10}{9}|P_{2r}(\Gamma)| \le \left(\frac{10}{9}\right)^2 |P_{4r}(\Gamma)| \le \left(\frac{10}{9}\right)^3 |P_{8r}(\Gamma)| \le \cdots \le \left(\frac{10}{9}\right)^m |P_{2^m r}(\Gamma)|$$

for an appropriate m.

9.3 The Thin-Thick Dichotomy

We let P be the set of δ-balls in $[0,1]^d$ that have richness at least r, but less than $2r$. By the inductive argument in Section 9.2, we may assume that $|P| \ge \frac{1}{10}|P_r(\Gamma)|$. We apply Proposition 6.1.1 with $E \sim r$ and $D = \delta^{-1}$, and $\lambda \sim D^{\frac{\varepsilon}{10}} \sim \delta^{-\frac{\varepsilon}{10}}$ Accordingly, we will have that

$$\rho \sim D^{\varepsilon^3}\lambda^{-1} \sim \delta^{-\varepsilon^3}\lambda^{-1}.$$

This gives the following two possibilities:

1. $|P| \lesssim \lambda^{d-1}r^{-2}|\Gamma|\delta^{-(d-1)}$ or

2. There is a collection of $\lambda\delta$-balls whose union contains a $\gtrapprox_{\log} 1$ fraction of P so that each of these $\delta\lambda$-balls intersects $\gtrapprox_{\log} \delta^{-\varepsilon^3}\lambda r$-many slabs of Γ.

If the first of these possibilities holds, we say we are in the thin case. Otherwise, we say we are in the thick case.

9.3.1 Thin Case

If we are in the thin case, then we have that

$$|P| \lessapprox_{\log} \lambda^{d-1}r^{-2}\delta^{-(d-1)}|\Gamma|. \tag{9.8}$$

By Base Case 2, we can assume that $r \lesssim \delta^{-(d-1)}$, which means that $r^{-(d-2)} \gtrsim \delta^{-(d-1)(d-2)}$

Thus, we have that

$$\begin{aligned} r^{-2}\delta^{-(d-1)}|\Gamma| &= r^{-2}\delta^{(d-1)(d-2)}\left(\delta^{-(d-1)(d-2)}\delta^{-(d-1)}\right)|\Gamma| \\ &\lesssim r^{-(d+1)}\delta^{-(d-1)(d-1)}|\Gamma|. \end{aligned}$$

Meanwhile, the spacing condition of Theorem 9.0.1 implies that $|\Gamma| \sim N\delta^{-(d-1)}W$, which is equivalent to the inequality

$$\delta^{-(d-1)} \sim N^{-1}W^{-1}|\Gamma|. \tag{9.9}$$

We take both sides of (9.9) to the power $d-2$ and substitute the result into (9.8) to give

$$\begin{aligned} |P| &\lessapprox_{\log} \lambda^{d-1}r^{-(d+1)}N^{-(d-1)}W^{-(d-1)}|\Gamma|^d \\ &\le \lambda^{d-1}r^{-(d+1)}W^{-(d-1)}|\Gamma|^d \\ &\sim \left(\delta^{-\frac{\varepsilon}{10d}}\right)^{d-1} r^{-(d+1)}W^{-(d-1)}|\Gamma|^d \\ &= \left(\delta^{-1}\right)^{\frac{(d-1)\varepsilon}{10d}} r^{-(d+1)}W^{-(d-1)}|\Gamma|^d \end{aligned}$$

The exponent for δ^{-1} on the far right-hand side is strictly less than ε. The discrepancy between this exponent $\frac{(d-1)\varepsilon}{10d}$ and our goal exponent ε compensates for all of the suppressed log losses implied by the symbol $\lessapprox_{\log}$, provided that δ is sufficiently small (cf. Chapter 8). Thus, we have that

$$|P| \le C_\varepsilon\,\delta^{-\varepsilon}\,r^{-(d+1)}W^{-(d-1)}|\Gamma|^d.$$

9.3.2 Thick Case

If we are in the thick case, then most of the incidences between the slabs of Γ and the r-rich δ-boxes occur inside the cubes Q_j of side length $\sim \lambda\delta$. Each Q_j contains $\lesssim \lambda^d$-many δ-boxes of $P_r(\Gamma)$, so we have that

$$|P_r(\Gamma)| \lessapprox_{\log} \#\left\{p \in P : p \subsetsim Q_j \text{ for some } j\right\} \lesssim \lambda^d\,\#\{Q_j\}. \tag{9.10}$$

To estimate the number of cubes Q_j we will apply the inductive hypothesis for a collection of thick slabs. If we thicken each δ-slab of Γ to a $\lambda\delta$-slab, then some of the thicker slabs that result may not be essentially distinct from each other, so we will need to do a pigeonholing step before applying the inductive hypothesis.

For each dyadic $M \geq 1$, let Γ_M be the collection of slabs $S \in \Gamma$ so that the thickened version of S contains $\sim M$-many slabs of Γ. If a δ-slab is essentially contained in $\lambda S \cap [0,1]^d$, then its normal direction must be in a cap of radius $\sim \lambda\delta$ centered at $\boldsymbol{n}(S)$. We can decompose this $\lambda\delta$-cap into a union of $\sim \lambda^{d-1}$-many δ-caps of bounded overlap. We let θ be a fixed δ-cap in this decomposition, and we consider the δ-slabs within $\lambda S \cap [0,1]^d$ whose normal directions are in this particular θ. It follows by comparing volumes that there are $\lesssim \lambda$-many δ-slabs for with directions in this particular cap. Since our decomposition contained $\sim \lambda^{d-1}$-many caps, we conclude that there are $\lesssim \lambda^d$-many δ-slabs total, i.e. we must have $M \lesssim \lambda^d$ for Γ_M to be nonempty. In particular, the number of candidate dyadic sizes is $\lesssim \log(\delta^{-1})$. It follows that we can choose a particular M so that Γ_M accounts for a $\gtrapprox_{\log} 1$ fraction of the incidences between $\{Q_j\}$ and Γ. The reason we can choose such an M is as follows: recalling that the expression $S \cap_{ess} Q_j$ should be read as, "S essentially intersects Q_j," we write

$$I(\{Q_j\},\Gamma) = \sum_{S\in\Gamma} \#\{j : S \cap_{\text{ess}} Q_j\} = \sum_{M \text{ dyadic}} \left(\sum_{S\in\Gamma} \#\{j : S \cap_{\text{ess}} Q_j\} \right).$$

Then we can simply select the M for which the inner sum on the right is largest.

Throughout the rest of the this section, we will consider a fixed M for which the contribution of M is maximal. For this M, we let $\tilde{\Gamma}$ be the set of λ-dilates of the slabs of Γ_M, i.e.

$$\left|\tilde{\Gamma}\right| = \{\lambda S \cap [0,1]^d\}.$$

Each slab of $\tilde{\Gamma}$ represents $\sim M$-many slabs of Γ, so we have that

$$|\tilde{\Gamma}| \sim \frac{|\Gamma|}{M}.$$

The number of δ-slabs passing through each Q_j is $\gtrapprox_{\log} \lambda r$, so that number of slabs in $|\tilde{\Gamma}|$ that pass through Q_j is $\gtrapprox_{\log} M^{-1}\lambda r$. Thus, resuming from 9.10, we have that

$$|P_r(\Gamma)| \lessapprox_{\log} \#\{p \in P : p \subsetsim Q_j \text{ for some } j\} \leq \lambda^d \#\{Q_j\} \leq \lambda^d \left| P_{\tilde{r}}(\tilde{\Gamma}) \right|.$$

As we verify in Subsection 9.3.3, the slabs of $\tilde{\Gamma}$ meet the premises of the inductive hypothesis with $\tilde{\delta} = \lambda\delta$ and $\tilde{r} \gtrapprox_{\log} \delta^{\varepsilon^3}\lambda M^{-1}r$.

Since $\tilde{r} \gtrapprox_{\log} \delta^{\varepsilon^3}\lambda M^{-1}r$, the λ^{-d} factor from the inductive hypothesis cancels the λ^d factor from the trivial bound. Thus, we have that

$$\begin{aligned} |P_r(\Gamma)| &\lessapprox_{\log} \lambda^d \left(C_\varepsilon \left(\tilde{\delta}\right)^{-\varepsilon} (\tilde{r})^{-d} W^{-(d-1)} |\tilde{\Gamma}|^d \right) \\ &\lessapprox_{\log} C_\varepsilon (\lambda\delta)^{-\varepsilon} \lambda^d \left(M^{-1}\delta^{\varepsilon^3}\lambda r \right)^{-d} W^{-(d-1)} |\tilde{\Gamma}|^d \\ &\sim C_\varepsilon (\lambda\delta)^{-\varepsilon} \delta^{-d\varepsilon^3} r^{-d} W^{-(d-1)} M^d (M^{-1}|\Gamma|)^d \\ &\sim C_\varepsilon\, \delta^{-\varepsilon} \left(\lambda^{-\varepsilon}\delta^{d\varepsilon^3} \right) r^{-d} W^{-(d-1)} |\Gamma|^d \end{aligned} \tag{9.11}$$

Recalling that $\lambda = \delta^{-\frac{\varepsilon}{10d}}$, we write

$$\lambda^{-\varepsilon}\delta^{d\varepsilon^3} = \delta^{\varepsilon^2\left(\frac{1}{10d} - d\varepsilon\right)}.$$

Provided that ε is sufficiently small, the quantity in parentheses is positive. If we set

$$\varepsilon' = \varepsilon^2 \left(\frac{1}{10d} - d\varepsilon \right),$$

then (9.11) says that

$$P_r(\Gamma)| \lessapprox_{\log} \delta^{-\varepsilon}\delta^{\varepsilon'} r^{-d} W^{-(d-1)} |\Gamma|^d. \tag{9.12}$$

For δ sufficiently small relative to ε, the gain from multiplying by $\delta^{\varepsilon'}$ compensates for all of the suppressed log losses. (This can be proved by an argument similar to the analysis in Chapter 8.)

9.3.3 Checking Hypotheses for the Thick Case

In order to use the estimate

$$|P_{\tilde{r}}(\tilde{\Gamma}) \leq C_\varepsilon \tilde{\delta}^{-\varepsilon} |\tilde{\Gamma}|^d \tilde{r}^{-(d+1)}$$

in (9.11) above, we were assuming that the slabs of $\tilde{\Gamma}$ met the premises of the inductive hypothesis with $\tilde{\delta} = \lambda\delta$ and $\tilde{r} \gtrapprox_{\log} \delta^{\varepsilon^3}\lambda M^{-1}r$. Here, we verify this assertion. (In the language of Chapter 5, we are verifying that implication ⓐ holds.)

Verifying that the first spacing condition holds

We must check that the new slabs satisfy the first spacing condition with the same parameter W that the old slabs did. First off, note that by Base Case 3b, we may assume that $\lambda\delta \leq W^{-1}$, which means that $W^{-1} \leq \tilde{\delta}^{-1}$. We will show that there is some $\tilde{N}$ so that if we subdivide $\mathbb{S}^{d-1}$ into $\lambda\theta$-caps then, any W^{-1}-slab with normal direction in a $\lambda\theta$-cap $\tilde{\theta}$ (essentially) contains $\sim \tilde{N}$-many slabs of $\tilde{\Gamma}$ whose normal directions are in $\tilde{\theta}$. We claim that if $N > M$, then we can satisfy this requirement by setting

$$\tilde{N} \sim \lambda^{d-1}\frac{N}{M}.$$

To check that this choice of $\tilde{N}$ works for $N > M$, we begin by writing $\tilde{\theta}$ as a union $\theta_1 \cup \cdots \cup \theta_\ell$, where each θ_i is a δ-cap, and $O(1)$-many δ-caps intersect at any point on $\mathbb{S}$. We must have $\ell \sim \lambda^{d-1}$, because for each i,

$$\frac{\mathrm{Vol}^{d-1}(\tilde{\theta})}{\mathrm{Vol}^{d-1}(\theta_i)} \sim \lambda^{d-1}.$$

Now, let R be a W^{-1}-slab with normal direction in $\tilde{\theta}$. We will momentarily want to estimate the number of $\tilde{\delta}$-slabs of $\tilde{\Gamma}$ that are essentially contained in R and have normal direction in $\tilde{\theta}$, but first we consider the number of δ-slabs of Γ that are essentially contained in R and have normal direction in $\tilde{\theta}$. For this, we note that the number of slabs of Γ that are essentially contained in R and have normal direction in θ_i is $\sim N$ for each i. This is because our assumption that $\lambda\delta \leq W^{-1}$ implies that if R_i is a W^{-1}-slab with the same center as R that has normal vector in θ_i, then R_i us essentially equivalent to R, i.e. R and R_i are not essentially distinct from each other.

Summing over i gives

$$\#\left\{S \in \Gamma : \begin{array}{c} \boldsymbol{n}(S) \in \tilde{\theta} \\ S \lesssim R \end{array}\right\} \sim \sum_{i=1}^{\ell} \#\left\{S \in \Gamma : \begin{array}{c} \boldsymbol{n}(S) \in \theta_i \\ S \lesssim R \end{array}\right\} \sim \sum_{i=1}^{\ell} N \sim \lambda^{d-1}N.$$

In the process of going from Γ to $\tilde{\Gamma}$ we chose thick slabs with represented a $\gtrapprox_{\log} 1$ fraction of Γ. (Recall that we say thick slab represents a thin slab, if the thin slab is essentially contained in the thick slab.) The thick slabs we chose as our representatives each contained $\sim M$ of the thick slabs, where M was found by pigeonholing.

If we did not have $\frac{N}{M} > 1$ we take $\tilde{N} \sim \lambda^{d-1}$. In this case, we may need to add more slabs to $\tilde{\Gamma}$ to ensure that λ^{d-1} is a lower bound for the number of δ-slabs in each W^{-1}-slab as well as an upper bound.

Verifying that $\tilde{r}$ is sufficiently large

In addition to demonstrating that we can choose $\tilde{N}$ for the second spacing condition, we must also check that the inequality

$$\tilde{r} \geq (\tilde{\delta})^{-\frac{\varepsilon}{4}} \tilde{\delta} |\tilde{\Gamma}| \tag{9.13}$$

is satisfied for our collection of thick slabs, provided that we had

$$r \geq \delta^{-\frac{\varepsilon}{4}} \delta |\Gamma| \tag{9.14}$$

for our original δ-slabs.

We will assume that ε is sufficiently small that $\frac{1}{40d} - \varepsilon > 0$. Specifically assume, that $\varepsilon \geq \varepsilon_0 > 0$ with $\frac{1}{40d} - \varepsilon_0 > 0$. We will show that for such an ε, there is a $c_\varepsilon > 0$ so that for any $\delta \leq \varepsilon$, we have the implication

$$r \geq \delta^{-\frac{\varepsilon}{4}} \delta |\Gamma| \qquad \Longrightarrow \qquad \tilde{r} \geq \tilde{\delta}^{-\frac{\varepsilon}{4}} \tilde{\delta} |\tilde{\Gamma}| \tag{9.15}$$

Recall that we had chosen $\tilde{r}$ with

$$\tilde{r} \gtrapprox_{\log} \delta^{\varepsilon^3} \lambda r M^{-1}, \tag{9.16}$$

where M was a dyadic number chosen by pigeonholing.

We can rewrite this as

$$\tilde{r} \geq A^{-1} \left(\log(\delta^{-1})\right)^{-B} \delta^{\varepsilon^3} \lambda r M^{-1} \tag{9.17}$$

for some absolute constants A and B (which may depend on d but not on any other parameters).

We will use this to show that (9.15) holds for any δ sufficiently small with respect to ε. Combining (9.17) with (9.14) to give

$$\begin{aligned}
\tilde{r} &\geq A^{-1}\left(\log(\delta^{-1})\right)^{-B}\delta^{\varepsilon^3}\lambda r M^{-1} \\
&\geq A^{-1}\left(\log(\delta^{-1})\right)^{-B}\delta^{\varepsilon^3}\lambda\left(\delta^{-\frac{\varepsilon}{4}}\delta|\Gamma|\right)M^{-1} \\
&= A^{-1}\left(\log(\delta^{-1})\right)^{-B}\delta^{\varepsilon^3}\delta^{-\frac{\varepsilon}{4}}(\lambda\delta)\left(|\Gamma|M^{-1}\right) \\
&\sim A^{-1}\left(\log(\delta^{-1})\right)^{-B}\delta^{\varepsilon^3}\delta^{-\frac{\varepsilon}{10d}}\delta^{-\frac{\varepsilon}{4}}\tilde{\delta}|\tilde{\Gamma}| \\
&= A^{-1}\left(\log(\delta^{-1})\right)^{-B}\delta^{\varepsilon^3}\left(\delta^{-\frac{\varepsilon}{10d}}\right)^{\frac{\varepsilon}{4}}\left(\tilde{\delta}\right)^{-\frac{\varepsilon}{4}}\tilde{\delta}|\Gamma|.
\end{aligned} \tag{9.18}$$

To account for the one $\sim$ that appeared in the above chain of inequalities, we define a new constant A'' so that

$$\tilde{r} \geq (A'')^{-1}\left(\log(\delta^{-1})\right)^{-B}\delta^{\varepsilon^3}\left(\delta^{-\frac{\varepsilon}{10d}}\right)^{\frac{\varepsilon}{4}}\left(\tilde{\delta}\right)^{-\frac{\varepsilon}{4}}\tilde{\delta}|\Gamma|.$$

To show that (9.13) holds, it suffices to show that

$$(A'')^{-1}\left(\log(\delta^{-1})\right)^{-B}\delta^{\varepsilon^3}\delta^{-\frac{\varepsilon^2}{40d}} \geq 1,$$

or, equivalently that

$$A''\left(\log(\delta^{-1})\right)^{B} \leq (\delta^{-1})^{\varepsilon^2(\frac{1}{40d}-\varepsilon)}. \tag{9.19}$$

Taking logs, we see that (9.19) is equivalent to the inequality

$$\log(A'') + B\log\log(\delta^{-1}) \leq \varepsilon^2\left(\frac{1}{40d}-\varepsilon\right)\log(\delta^{-1}). \tag{9.20}$$

To prove (9.20) for a particular δ, it is sufficient to show that the following two inequalities hold:

$$\log(A'') \leq \frac{1}{2}\varepsilon^2\left(\frac{1}{40d}-\varepsilon_0\right)\log(\delta^{-1}); \tag{9.21}$$

and

$$B\log\log(\delta^{-1}) \leq \frac{1}{2}\varepsilon^2\left(\frac{1}{40d}-\varepsilon_0\right)\log(\delta^{-1}). \tag{9.22}$$

(Here, we are using ε_0 in the sense specified directly below inequality (9.14).)

We note that (9.21) is equivalent to the inequality

$$\log(\delta^{-1}) \geq 2\log(A'')\varepsilon^{-2}\left(\frac{1}{40d} - \varepsilon_0\right)^{-1} \tag{9.23}$$

and that (9.22) is equivalent to the inequality

$$\frac{\log\log(\delta^{-1})}{\log(\delta^{-1})} \leq \frac{\varepsilon^2\left(\frac{1}{40d} - \varepsilon_0\right)}{2B}. \tag{9.24}$$

As $\delta \to 0^+$, we have that $\log(\delta^{-1}) \to \infty$ and $\frac{\log\log(\delta^{-1})}{\log(\delta^{-1})} \to 0$. Thus, there is a $c_\varepsilon = c_\varepsilon(\varepsilon)$ so that both (9.23) and (9.24) hold for all $\delta \leq c_\varepsilon(\varepsilon)$. This means that we have the implication (9.15) for all $\delta \leq c_\varepsilon(\varepsilon)$.

Remark 9.3.1. *The reason we do not have a constant $C_1(\varepsilon)$ in our lower bound for r in the statement of Theorem* 1.4.2 *is that we are only checking that $\tilde{r}$ satisfies the lower bound of the theorem in the case that $\delta \leq c_\varepsilon$. The implication* (9.14) $\implies$ (9.13) *may fail for large δ. Allowing for a 'fudge factor' $C_1(\varepsilon)$ in the bound of the theorem statement would allow us to deduce directly for large δ that the lower bound held. However, since we are accounting for $\delta \leq c_\varepsilon$ as a base case, this 'fudge factor' is unnecessary.*

Chapter 10

An Iterative Interpretation of the Inductive Argument

Our specific choice of $\tilde{\delta}$ in the thick case sets us up well to reinterpret our induction proof as an iterative argument in which we repeatedly increase δ, the slab thickness. I was inspired to search for an iterative procedure for estimating $|P_r(\Gamma)|$ after studying other iterative proofs like the density increment argument in Szemeredi's proof of Roth's theorem.

We will ultimately build up to an another proof of Theorem 1.4.2, which we restate (again) for convenience.

Theorem 10.0.1. *(Theorem 1.4.2, re-revisited)*

For any $\varepsilon > 0$ sufficiently small (relative to d), there exists a constant $C_\varepsilon = C(\varepsilon, d) > 1$ so that if $\delta \in (0,1)$ and $1 \leq W \leq \delta^{-1}$, then

$$|P_r(\Gamma)| \leq C_\varepsilon \delta^{-\varepsilon} W^{-(d-1)} r^{-d} |\Gamma|^d \tag{10.1}$$

for any collection Γ of δ-slabs that satisfies the first spacing condition for the parameter W and any r with

$$r \geq \delta^{-\frac{\varepsilon}{4}} \delta |\Gamma|. \tag{10.2}$$

To make a certain part of the argument work, we'll need to assume that ε is sufficiently small that $1 - 10d^2\varepsilon > 0$; in particular, we'll assume that

$$1 - 10d^2\varepsilon \geq \frac{1}{2}. \tag{10.3}$$

Recall that in the thick case of our induction proof we had set

$$\tilde{\delta} = \lambda\delta = \delta^{-\frac{\varepsilon}{10d}}\,\delta = \delta^{(1-\frac{\varepsilon}{10d})}.$$

Let's momentarily forget about checking all of the hypotheses that must line up for us to be able to thicken our slabs again and just think about what happens to δ as we repeatedly thicken. We define a sequence of thicknesses $\{\delta_j\}_j$ with

$\delta_0 = \delta$,

$\delta_1 = \delta^{(1-\frac{\varepsilon}{10d})}$,

$\delta_2 = \delta^{(1-\frac{\varepsilon}{10d})^2}$,

and so on. In general, we'll have

$\delta_j = \delta^{(1-\frac{\varepsilon}{10d})^j}$.

If j is sufficiently large then $\delta_j = \delta^{(1-\frac{\varepsilon}{10d})^j}$ must exceed exceed $\frac{1}{3}$. We can arrange by choosing C_ε well that (10.1) holds for all $\delta \geq \frac{1}{3}$, so if we can thicken enough times to achieve $\delta^{(1-\frac{\varepsilon}{10d})^j} \geq \frac{1}{3}$, then showing (10.1) for r_0 and δ_0, our initial values of r and δ becomes a matter of comparing $|P_{r_0}(\delta_0)|$ to $|P_{r_j}(\delta_j)|$ for an appropriately chosen r_j. [1]

Meanwhile, if we could not thicken enough times to achieve $\delta^{(1-\frac{\varepsilon}{10d})^j} \geq \frac{1}{3}$ then we must have encountered some collection of slabs Γ_k and some candidate richness r_k so that $|P_{r_k}(\Gamma_k)\backslash P_{2r_k}(\Gamma_k)| \geq \frac{1}{10}|P_{r_k}(\Gamma_k)|$ and Γ_k was in the thin case of Proposition 6.1.1. For such a pair (Γ_k, r_k), we would have

$$|P_{r_k}(\Gamma_k)| \leq C_\varepsilon\, \delta_k^{-\varepsilon}|\Gamma_k|^d r_k^{-d} W^{-(d-1)}. \tag{10.4}$$

In Section 10.1, we find an upper bound for the smallest m with

$$\delta_m^{1-\frac{\varepsilon}{10d}} \geq \frac{1}{3}.$$

[1] Our choice of threshold $\frac{1}{3}$ is not particularly important. We could replace $\frac{1}{3}$ by any constant in $(0, 1)$ whose reciprocal exceeded e and still have the argument of section 10.1 work just as well.

The upper bound will be on the order of $\frac{\log\log(\delta^{-1})}{\varepsilon/(10d)}$. We can conclude that 10.4 holds for some k not exceeding this upper bound. We then relate $|P_{r_k}(\Gamma_k)|$ to $|P_{r_0}(\Gamma_0)|$ - the quantity we were originally trying to estimate. We first show how to relate $|P_{r_k}(\Gamma_k)|$ to $|P_{r_0}(\Gamma_0)|$ in a special case (Section 10.3) and then in full generality (Section 10.4). We assume throughout our work that $k > 0$; otherwise we would already know that (10.3) held for our original r.

Although there is loss involved in comparing $|P_{r_0}(\Gamma_0)|$ to $|P_{r_k}(\Gamma_k)|$, the fact that $\delta_k^{-\varepsilon} < \delta_0^{-\varepsilon}$ compensates for this loss, provided that our initial δ_0 was sufficiently small. Specifically, we need $\delta_0 \leq c_\varepsilon$ for a small constant c_ε satisfying properties described in Section 10.5. If our initial δ did not exceed this threshold, then our goal bound holds by the 'trivial' bound $|P_r(\Gamma)| \lesssim \delta^{-d}$.

In showing that the strict inequality $\delta_k^{-\varepsilon} < \delta_0^{-\varepsilon}$ compensates for all of the losses involved in passing from $|P_{r_0}(\Gamma_0)|$ to $|P_{r_k}(\Gamma_k)|$, we first show that it compensates for ε^3-losses (Subsection 10.5.1) and then show that it compensates for log losses (Subsection 10.5.2). To reprise our 'blocks in a box' analogy, we first show that we can largest blocks into the space afforded by the strict inequality $\delta_k^{-\varepsilon} < \delta_0^{-\varepsilon}$. Then we show that we can fit the rest of the blocks into the remaining space.

10.1 Getting to the Threshold Thickness

We let m_0 be the minimal m with

$$\frac{1}{3} \leq \delta_m^{1-\frac{\varepsilon}{10d}} = \delta^{(1-\frac{\varepsilon}{10d})^{m+1}}.$$

or equivalently, the minimal m with

$$3 \geq \delta^{-(1-\frac{\varepsilon}{10d})^{m+1}}.$$

Then we must have

$$m_0 < \frac{\log\log(\delta^{-1})}{\log\left(\frac{1}{1-\frac{\varepsilon}{10d}}\right)}. \tag{10.5}$$

To see this, we use the minimality of m_0 to give

$$3 \leq \delta^{-(1-\frac{\varepsilon}{10d})^{m_0}}.$$

We take logs two times and use the fact that $\log\log 3 > 0$ to give

$$0 < \log\log 3 \leq m_0 \log\left(1 - \frac{\varepsilon}{10d}\right) + \log\log(\delta^{-1}),$$

which we can rearrange to give (10.5). We note that

$$\log\left(\frac{1}{1-\frac{\varepsilon}{10d}}\right) \sim \frac{\varepsilon}{10d}$$

We can see this by using a Taylor expansion:

$$\log\left(\frac{1}{1-\frac{\varepsilon}{10d}}\right) = \frac{\varepsilon}{10d} + \frac{(\frac{\varepsilon}{10d})^2}{2} + \frac{(\frac{\varepsilon}{10d})^3}{3} + \frac{(\frac{\varepsilon}{10d})^4}{4} + ... \tag{10.6}$$

This series converges because $\frac{\varepsilon}{10d} < 1$. Since all of the terms are positive, we have that

$$\log\left(\frac{1}{1-\frac{\varepsilon}{10d}}\right) \geq \frac{\varepsilon}{10d}.$$

We also have that

$$\begin{aligned}
\log\left(\frac{1}{1-\frac{\varepsilon}{10d}}\right) &= \frac{\varepsilon}{10d}\left(1 + \frac{\frac{\varepsilon}{10d}}{2} + \frac{(\frac{\varepsilon}{10d})^2}{3} + \frac{(\frac{\varepsilon}{10d})^3}{4} + \ldots\right) \\
&\leq \frac{\varepsilon}{10d}\left(1 + \frac{\varepsilon}{10d} + (\frac{\varepsilon}{10d})^2 + (\frac{\varepsilon}{10d})^3 + \ldots\right) \\
&= \frac{\varepsilon}{10d}\left(\frac{1}{1-\frac{\varepsilon}{10d}}\right).
\end{aligned}$$

Assuming that $\frac{\varepsilon}{10d} \leq \frac{1}{2}$, we then have that

$$\left(\frac{1}{1-\frac{\varepsilon}{10d}}\right) \leq 2,$$

so

$$\log\left(\frac{1}{1-\frac{\varepsilon}{10d}}\right) \leq 2\left(\frac{\varepsilon}{10d}\right),$$

which means we can rewrite (10.5) as

$$m_0 < \frac{\log\log(\delta^{-1})}{\frac{\varepsilon}{10d}}. \tag{10.7}$$

10.2 Setting up the Iteration

In our work in the thick case argument of Chapter 9, we proved the following proposition as an intermediate result.

Proposition 10.2.1. *Let ε, W, r, and δ be fixed. Suppose that Γ is a collection of δ-slabs that satisfies (WS-1) with parameter W and that $r \geq \delta^{-\varepsilon/4}\delta|\Gamma|$. Suppose further that*

$$|P_r(\Gamma)\backslash P_{2r}(\Gamma)| \geq \frac{1}{10}|P_r(\Gamma)|.$$

If Γ is in the thick case of proposition 6.1.1 *then there exist a dyadic M, a candidate richness $\tilde{r} \gtrapprox_{\log} \delta^{\varepsilon^3}\lambda M^{-1}r$, and a collection $\tilde{\Gamma}$ of $\lambda\delta$-slabs so that the following hold:*

(1) Γ satisfies (WS-1) with parameter W;

(2) $|\tilde{\Gamma}| \lessapprox_{\log} M^{-1}|\Gamma|$;

(3) $|P_r(\Gamma)| \lessapprox_{\log} \lambda^d|P_{\tilde{r}}(\tilde{\Gamma})|$.

We can rewrite the bound from (3) as

$$|P_r(\Gamma)| \leq A\left(\log(\delta^{-1})\right)^B \lambda^d|P_{\tilde{r}}(\tilde{\Gamma})| \tag{10.8}$$

for some constants A and B.

If we let $P'_{\tilde{r}} = P_{\tilde{r}}(\tilde{\Gamma})\backslash P_{2\tilde{r}}$, it may or may not be the case that

$$\left|P'_{\tilde{r}}(\tilde{\Gamma})\right| \geq \frac{1}{10}\left|P_{\tilde{r}}(\tilde{\Gamma})\right|. \tag{10.9}$$

If (10.9) *does* hold, then we can apply Proposition 6.1.1 for $P = P'_{\tilde{r}}$, $E = \tilde{r}$ and $D = \tilde{\delta}^{-1}$. If we are in the thin case, then algebraic manipulation will allow us to conclude that

$$|P| \lessapprox_{\log} \tilde{\lambda}^{d-1}\tilde{r}^{-d}|\tilde{\Gamma}|^d W^{-(d-1)}.$$

If we are in the thick case, then we can pass to yet another scale and consider slabs of thickness $\sim \tilde{\lambda}\tilde{\delta}$.

Passing to a collection of slabs of thickness $\sim \tilde{\lambda}\tilde{\delta}$ involved a lot of 'ifs'; we could only do this if (10.9) held and if we were not in the thin case. This second 'if' is not a problem for rewriting our induction as an iterative argument; if we arrive in the thin case, we can terminate our iteration, which is actually exactly what we want. However, the question of whether (10.9) holds creates a complication in rewriting our induction proof. It turns out that every single time we increase δ to $\lambda\delta$ we may have to repeatedly double $\tilde{r}$ in search of some ℓ so that

$$\left|P'_{2^\ell \tilde{r}}(\tilde{\Gamma})\right| \geq \frac{1}{10}\left|P_{2^\ell \tilde{r}}(\tilde{\Gamma})\right|.$$

We will always be able to find such an ℓ, but the repeated doubling complicates the iterative procedure.

We will initially consider a simpler procedure which works for the very special case in which (10.9) immediately holds every time we are in the thick case of Proposition 6.1.1. We treat this special case in subsection 10.3 and then treat the general case in subsection 10.4, in which we develop a more complicated procedure that allows for a repeated doubling step after each thickening step.

10.3 White Lie Iterative Argument - Iterating Just the Thick Case

To prove the goal bound

$$|P_r(\Gamma)| \leq C_\varepsilon\, \delta^{-\varepsilon} \frac{|\Gamma|^d}{r^d\, W^{d-1}} \tag{10.10}$$

for a fixed collection of slabs and a fixed r, we define

$r_0 = r$;

$\delta_0 = \delta$;

$\Gamma_0 = \Gamma$.

We let $\lambda_0 = \delta_0^{-\frac{\varepsilon}{10d}}$, and we apply Proposition 10.2.1 to find a dyadic M_1 and a collection Γ_1

of thickened slabs so that

$$|P_{r_0}(\Gamma_1)| \leq A \log\left(\delta_0^{-1}\right)^B \lambda^d |P_{r_1}(\Gamma_1)|.$$

for some $r_1 \gtrapprox_{\log} \delta_0^{\varepsilon^3} \lambda_0 M_0^{-1} r_0$.

Likewise, we can find a dyadic M_1 and a collection Γ_1 of thickened slabs so that

$$|P_{r_1}(\Gamma_1)| \leq A \log\left(\delta_1^{-1}\right)^B \lambda^d |P_{r_2}(\Gamma_2)|$$

for some $r_2 \gtrapprox_{\log} \delta_1^{\varepsilon^3} \lambda_1 M_1^{-1} r_0$. Thus,

$$|P_{r_0}(\Gamma_1)| \leq A^2 \log(\delta_0^{-1}\delta_1^{-1})^B \lambda^d |P_{\tilde{r_2}}(\tilde{\Gamma_2})|$$

If we apply Proposition 10.2.1 yet again, then we will have

$$|P_{r_0}(\Gamma_1)| \leq A^3 \log(\delta_0^{-1}\delta_1^{-1}\delta_2^{-1})^B \lambda^d |P_{r_3}(\Gamma_3)|$$

Let k be a large positive integer so that each of $\Gamma_1, \ldots, \Gamma_{k-1}$ was in the thick case. After applying Proposition 10.2.1 k times we will have

$$|P_{r_0}(\Gamma_0)| \leq A^k \left(\prod_{j=0}^{k-1} \log(\delta_j^{-1})\right)^B \left(\prod_{j=0}^{k-1} \lambda_j\right)^d |P_{r_k}(\Gamma_k)|. \tag{10.11}$$

We will abbreviate this inequality to

$$|P_{r_0}(\Gamma_0)| \lessapprox_{\log^{\otimes k}} (\lambda_0 \ldots \lambda_{k-1})^d \, |P_{r_k}(\Gamma_k)|. \tag{10.12}$$

We are using the tensor product to reflect the fact that we have k-many log losses, each of which involves multiplication by $A\left(\log\left(\delta_j^{-1}\right)\right)^B$ *for a different value of* j. (Hence, the implied multiplication is by a function with k inputs, namely all of the intermediate values of δ.)

One reason to temporarily suppress log losses is that they are very small compared to the product $(\lambda_0 \dots \lambda_{k-1})^d$, so long as δ is small compared to ε. If k is sufficiently large, then

$$|P_{r_k}(\Gamma_k)| \leq C_\varepsilon \delta_k^{-\varepsilon} r_k^{-d} |\Gamma_k|^d W^{-(d-1)}. \tag{10.13}$$

Specifically, by our work in Section 10.1, inequality (10.13) must hold for some k with

$$k \lesssim \frac{\log\log(\delta^{-1})}{\frac{\varepsilon}{10d}}.$$

However, we don't need the value of k quite yet, as we will first rewrite the expression $(\lambda_0 \dots \lambda_{k-1})^d \, |P_{r_k}(\Gamma_k)|$.

Combining inequalities (10.12) and (10.13) gives

$$|P_{r_0}(\Gamma_0)| \lessapprox_{\log^{\otimes k}} (\lambda_0 \dots \lambda_{k-1})^d \left(C_\varepsilon \delta_k^{-\varepsilon} r_k^{-d} |\Gamma_k|^d W^{-(d-1)} \right).$$

For each j, the contribution of λ_j^d will be canceled when we replace r_{j+1}^{-1} by $\delta^{-\varepsilon^3} \lambda_j^{-1} M_j^1 r_j^{-1}$. Each replacement introduces a new factor of M_j, which will in turn, be canceled when we replace $|\Gamma_{j+1}|$ by $M_j^{-1} |\Gamma_j|$. Each of the replacements referenced will involve a log loss. That is, each of the replacements referenced will involve multiplication by a factor of $A \left(\log \left(\delta_j^{-1} \right) \right)^B$ for some constants A and B for some constants A and B.

Compounding these individual log losses gives a relationship

$$|P_{r_0}(\Gamma)| \lessapprox_{\log^{\otimes k}} C_\varepsilon \, \delta_k^{-\varepsilon} \left(\prod_{j=1}^{k-1} \delta_j^{\varepsilon^3} \right)^{-d} r_0^{-d} |\Gamma_0|^d \, W^{-(d-1)}. \tag{10.14}$$

Here are the details:

$$
\begin{aligned}
(\lambda_0 \dots \lambda_{k-1})^d \, |P_{r_k}(\Gamma_k)| &\leq (\lambda_0 \dots \lambda_{k-1})^d \, C_\varepsilon \, \delta_k^{-\varepsilon} r_k^{-d} |\Gamma_k|^d \, W^{-(d-1)} \\
&\lessapprox_{\log^{\otimes k}} (\lambda_0 \dots \lambda_{k-1})^d \, C_\varepsilon \delta_k^{-\varepsilon} \left(\prod_{j=1}^{k-1} \delta_j^{\varepsilon^3} \right)^{-d} \\
&\qquad (\lambda_0 \dots \lambda_{k-1})^{-d} \left(\prod_{j=0}^{k-1} M_j^{-1} \right)^{-d} \quad r_0^{-d} |\Gamma_k|^d \, W^{-(d-1)} \\
&= C_\varepsilon \, \delta_k^{-\varepsilon} \left(\prod_{j=1}^{k-1} \delta_j^{\varepsilon^3} \right)^{-d} \left(\prod_{j=0}^{k-1} M_j^{-1} \right)^{-d} \quad r_0^{-d} |\Gamma_k|^d \, W^{-(d-1)} \\
&\lessapprox_{\log^{\otimes k}} C_\varepsilon \, \delta_k^{-\varepsilon} \left(\prod_{j=1}^{k-1} \delta_j^{\varepsilon^3} \right)^{-d} \left(\prod_{j=0}^{k-1} M_j^{-1} \right)^{-d} \quad r_0^{-d} \\
&\qquad \left(\prod_{j=0}^{k-1} M_j^{-1} \right)^{d} \quad |\Gamma_0|^d \, W^{-(d-1)} \\
&= C_\varepsilon \, \delta_k^{-\varepsilon} \left(\prod_{j=1}^{k-1} \delta_j^{\varepsilon^3} \right)^{-d} \quad r_0^{-d} |\Gamma_0|^d \, W^{-(d-1)}
\end{aligned}
$$

10.4 Full Iterative Argument

If we did not have the simplifying assumption that (10.9) immediately holds every time we are in the thick case of Proposition 6.1.1, then we would still work with a sequence $\Gamma_0, \dots, \Gamma_m$ like in the white lie version of the argument. We would just have to do more work to go from Γ_j to Γ_{j+1} for each j. Specifically, before applying Proposition 6.1.1 to Γ_j, we would need to repeatedly double r_j until we found some index ℓ_j for which we had

$$
|P_{2^{\ell_j} r_j}(\Gamma_j) \backslash P_{2^{\ell_j+1} r_j}(\Gamma_j)| \geq \frac{1}{10} |P_{2^{\ell_j} r_j}(\Gamma_j)|.
$$

We iterate the procedure from Section 9.2 ℓ_j-many times to give

$$
|P_{r_j}(\Gamma_j)| \leq \frac{10}{9} |P_{2r_j}(\Gamma_j)| \leq \left(\frac{10}{9}\right)^2 |P_{4r_j}(\Gamma_j)| \leq \dots \leq \left(\frac{10}{9}\right)^{\ell_j} |P_{2^{\ell_j} r_j}(\Gamma_j)|. \tag{10.15}
$$

We can now apply Proposition 6.1.1 to Γ_j for $E = 2^{\ell_j} r_j$ and $D = \delta_j^{-1}$. So long as we are not in the thin case of Proposition 6.1.1, we will obtain a new dyadic M_j, a new set of slabs Γ_{j+1}, and a new richness r_{j+1} with

$$|\Gamma_{j+1}| \sim \frac{|\Gamma_j|}{M_j}, \tag{10.16}$$

and

$$r_{j+1} \gtrapprox_{\log} \delta_j^{\varepsilon^3} \lambda_j \left(2^{\ell_j} r_j\right) M_j^{-1},$$

so that

$$\left|P_{2^{\ell_j} r_j}(\Gamma_j)\right| \lessapprox_{\log} \lambda_j^d \left|P_{r_{j+1}}(\Gamma_{j+1})\right|.$$

Combining this with (10.15) gives

$$|P_{r_j}(\Gamma_j)| \leq \left(\frac{10}{9}\right)^{\ell_j} |P_{2_j^{\ell} r_j}(\Gamma_j)| \lessapprox_{\log} \left(\frac{10}{9}\right)^{\ell_j} \lambda_j^d \left|P_{r_{j+1}}(\Gamma_{j+1})\right| \tag{10.17}$$

Iterating inequality (10.17) represents iterating the two step process in which we first double r_j as needed and then thicken and pigeonhole, as prescribed in the thick case. The first three repetitions - if we need that many - will give

$$\begin{aligned}
|P_r(\Gamma)| &= |P_{r_0}(\Gamma_0)| \\
&\leq A \left(\log(\delta_0^{-1})\right)^B \left(\frac{10}{9}\right)^{\ell_0} \lambda_0^d \, |P_{r_1}(\Gamma_1)| \\
&\leq A^2 \left(\log(\delta_0^{-1}) \log(\delta_1^{-1})\right)^B \left(\frac{10}{9}\right)^{\ell_0+\ell_1} (\lambda_0\lambda_1)^d \, |P_{r_2}(\Gamma_2)| \\
&\leq A^3 \left(\log(\delta_0^{-1}) \log(\delta_1^{-1}) \log(\delta_2^{-1})\right)^B \left(\frac{10}{9}\right)^{\ell_0+\ell_1+\ell_2} (\lambda_0\lambda_1\lambda_2)^d \, |P_{r_3}(\Gamma_3)| \, .
\end{aligned}$$

For a general k, we will have after k iterations - if this many are needed - that

$$|P_{r_0}(\Gamma_0)| \leq A^k \left(\prod_{j=0}^{k-1} \log(\delta_j^{-1})\right)^B \left(\frac{10}{9}\right)^{\ell_1+\cdots+\ell_{k-1}} (\lambda_0 \ldots \lambda_{k-1})^d \, |P_k(\Gamma_k)|.$$

As we did in the white lie version, we will temporarily suppress our log losses with the shorthand

$$|P_{r_0}(\Gamma_0)| \lessapprox_{\log^{\otimes k}} \left(\frac{10}{9}\right)^{\ell_0+\cdots+\ell_{k-1}} (\lambda_0 \ldots \lambda_{k-1})^d \, |P_{r_k}(\Gamma_k)|. \tag{10.18}$$

If k is sufficiently large, then we will have that

$$|P_{r_k}(\Gamma_k)| \le C_\varepsilon\, \delta_k^{-\varepsilon} r_k^{-d} |\Gamma_k|^d W^{-(d-1)}.$$

For $j = 0, \dots, k-1$, we can relate r_{j+1} to r_j by

$$r_{j+1} \gtrsim \log(\delta_j^{-1})^{O(1)} \delta_j^{\varepsilon^3} \lambda_j (2^{\ell_j} r_j) M_j^{-1}.$$

Thus, we have that

$$r_k \lessapprox_{\log^{\otimes k}} 2^{\ell_0 + \cdots + \ell_{k-1}} \left(\prod_{j=0}^{k-1} \delta_j^{\varepsilon^3} \lambda_j M_j^{-1} \right) r_0,$$

or, equivalently, that

$$r_k^{-d} \lessapprox_{\log^{\otimes k}} 2^{-d(\ell_0 + \cdots + \ell_{k-1})} \left(\prod_{j=0}^{k-1} \delta_j^{\varepsilon^3} \lambda_j M_j^{-1} \right)^{-d} r_0^{-d} \tag{10.19}$$

When we combine (10.19) and (10.18), the product $\prod_j \lambda_j^{-d}$ on the right-hand side of (10.19) cancels the corresponding product on the right-hand side of (10.18). Meanwhile, the expression $2^{-d(\ell_0+\cdots+\ell_{k-1})}$ more than compensates for the product $\prod_j \left(\frac{10}{9}\right)^{\ell_j}$. Thus, we have that

$$\begin{aligned}
|P_{r_0}(\Gamma_0)| &\lessapprox_{\log^{\otimes k}} \left(\frac{10}{9}\right)^{\ell_0+\cdots+\ell_{k-1}} (\lambda_0 \dots \lambda_{k-1})^d \left(C_\varepsilon\, \delta_k^{-\varepsilon}\, r_k^{-d} |\Gamma_k|^d W^{-(d-1)} \right) \\
&\lessapprox_{\log^{\otimes k}} \left(\frac{10}{9}\right)^{\ell(0)+\cdots+\ell(k-1)} (\lambda_0 \dots \lambda_{k-1})^d\, C_\varepsilon\, \delta_k^{-\varepsilon} \left(\frac{1}{2}\right)^{d(\ell_0+\cdots+\ell_{k-1})} \\
&\qquad (\lambda_0 \dots \lambda_{k-1})^{-d} \left(\prod_{j=0}^{k-1} \delta_j^{\varepsilon^3} M_j^{-1} \right)^{-d} r_0^{-d} |\Gamma_k|^d W^{-(d-1)} \\
&\le C_\varepsilon\, \delta_k^{-\varepsilon} \left(\prod_{j=0}^{k-1} \delta_j^{\varepsilon^3} M_j^{-1} \right)^{-d} r_0^{-d}\, |\Gamma_k|^d W^{-(d-1)} \\
&= C_\varepsilon\, \delta_k^{-\varepsilon} \left(\prod_{j=0}^{k-1} \delta_j^{\varepsilon^3} \right)^{-d} \left(\prod_{j=0}^{k-1} M_j \right)^{d} r_0^{-d} |\Gamma_k|^d W^{-(d-1)}.
\end{aligned} \tag{10.20}$$

Meanwhile, for each j, we have that

$$|\Gamma_{j+1}| \sim |\Gamma_j| M_j^{-1}$$

so

$$|\Gamma_k|^d = O_k(1)\,|\Gamma_0|^d \left(\prod_{j=1} M_j^{-1}\right)^d. \tag{10.21}$$

We substitute this into (10.20) to give

$$|P_{r_0}(\Gamma_0)| \lessapprox_{\log^{\otimes k}} C_\varepsilon\,\delta_k^{-\varepsilon} \left(\prod_{j=0}^{k-1} \delta_j^{\varepsilon^3}\right)^{-d} r_0^{-d} |\Gamma_0|^d W^{-(d-1)}.$$

(Here, the $O_k(1)$ factor from 10.21 is absorbed into the implied constant of the $\lessapprox_{\log^{\otimes k}}$ symbol.)

So far, we have proved that

$$|P_{r_0}(\Gamma_0)| \lessapprox_{\log^{\otimes k}} C_\varepsilon\,\delta_k^{-\varepsilon} \left(\prod_{j=0}^{k-1} \delta_j^{\varepsilon^3}\right)^{-d} r_0^{-d} |\Gamma_0|^d W^{-(d-1)}. \tag{10.22}$$

10.5 Concluding Estimates

10.5.1 Managing ε^3-losses

We now wish to estimate the product

$$\delta_k^{-\varepsilon} \left(\prod_{j=1}^{k-1} \delta_j^{\varepsilon^3}\right)^{-d}.$$

For this, we recall that

$$\delta_j = \delta_0^{(1-\frac{\varepsilon}{10d})^j}.$$

Thus, we have that

$$\delta_k^{-\varepsilon}\left(\prod_{j=1}^{k-1}\delta_j^{\varepsilon^3}\right)^{-d} = \left(\delta^{(1-\frac{\varepsilon}{10d})^k}\right)^{-\varepsilon}\left(\prod_{j=0}^{k-1}\delta^{(1-\frac{\varepsilon}{10d})^j}\right)^{-d\varepsilon^3}$$
$$= \left(\delta^{(1-\frac{\varepsilon}{10d})^k}\right)^{-\varepsilon}\left(\delta^{\Sigma_k}\right)^{-d\varepsilon^3},$$

where

$$\Sigma_k := \sum_{j=0}^{k-1}(1-\frac{\varepsilon}{10d})^j = \frac{1-(1-\frac{\varepsilon}{10d})^k}{1-(1-\frac{\varepsilon}{10d})} = \left(1-(1-\frac{\varepsilon}{10d})^k\right)\cdot\frac{10d}{\varepsilon}.$$

That is,

$$\delta_k^{-\varepsilon}\left(\prod_{j=1}^{k-1}\delta_j^{\varepsilon^3}\right)^{-d} = \left(\delta^{(1-\frac{\varepsilon}{10d})^k}\right)^{-\varepsilon}\left(\delta^{1-(1-\frac{\varepsilon}{10d})^k}\right)^{-d\varepsilon^3\cdot\frac{10d}{\varepsilon}}$$
$$= \delta^{-\varepsilon}\delta^{\varepsilon}\left(\delta^{(1-\frac{\varepsilon}{10d})^k}\right)^{-\varepsilon}\left(\delta^{1-(1-\frac{\varepsilon}{10d})^k}\right)^{-10d^2\varepsilon^2}$$
$$= \delta^{-\varepsilon}\left(\delta^{1-(1-\frac{\varepsilon}{10d})^k}\right)^{\varepsilon}\left(\delta^{1-(1-\frac{\varepsilon}{10d})^k}\right)^{-10d^2\varepsilon^2}$$
$$= \delta^{-\varepsilon}\left(\delta^{1-(1-\frac{\varepsilon}{10d})^k}\right)^{\varepsilon\left(1-10d^2\varepsilon\right)}$$
$$= \delta^{-\varepsilon}\left(\delta^{1-(1-\frac{\varepsilon}{10d})^k}\right)^{\varepsilon\left(1-10d^2\varepsilon\right)}$$

We note that $1-(1-\frac{\varepsilon}{10d})^k > 0$. In fact, this quantity is bounded away from 0, which will be important later. By our assumption (10.3), we have that $1-10d^2\varepsilon > \frac{1}{2}$, which means that

$$\left(\delta^{1-(1-\frac{\varepsilon}{10d})^k}\right)^{\varepsilon\left(1-10d^2\varepsilon\right)} < \delta^{\frac{\varepsilon}{2}(1-(1-\frac{\varepsilon}{10d})^k)}$$

Returning to (10.22), we see that we can now update our bound to

$$|P_{r_0}(\Gamma_0)| \lessapprox_{\log^{\otimes k}} C_\varepsilon\,\delta_0^{-\varepsilon}\,\delta_0^{\frac{\varepsilon}{2}(1-(1-\frac{\varepsilon}{10d})^k)}\,r_0^{-d}|\Gamma_0|^d W^{-(d-1)}. \tag{10.23}$$

10.5.2 Managing log losses

We claim that if δ is sufficiently small (relative to ε), then the $\delta^{\frac{\varepsilon}{2}(1-(1-\frac{\varepsilon}{10d})^k)}$ factor compensates for all of the the log losses that we have been suppressing with our $\lessapprox_{\log^{\otimes k}}$ notation.

To see this, we rewrite (10.23) as

$$|P_{r_0}(\Gamma_0)| \leq A^k \left(\prod_{j=0}^{k-1} \log \left(\frac{1}{\delta_j} \right) \right)^B C_\varepsilon \, \delta^{-\varepsilon} \, \delta^{\frac{\varepsilon}{2}(1-(1-\frac{\varepsilon}{10d})^k)} \, \frac{|\Gamma|^d}{r^d W^{d-1}} \tag{10.24}$$

for some constants A and B.

To show that our goal bound (10.10) holds, it is sufficient to prove that the specific A and B of inequality (10.24) satisfy

$$A^k \leq \left(\frac{1}{\delta} \right)^{\frac{\varepsilon}{4}(1-(1-\frac{\varepsilon}{10d})^k)} \tag{10.25}$$

and

$$\left(\prod_{j=0}^{k-1} \log \left(\frac{1}{\delta_j} \right) \right)^B \leq \left(\frac{1}{\delta} \right)^{\frac{\varepsilon}{4}(1-(1-\frac{\varepsilon}{10d})^k)} \tag{10.26}$$

respectively. In turn, finding a lower bound for $1-(1-\frac{\varepsilon}{10d})^k$ will allow us to further reduce to proving inequalities that are sufficient for each of (10.25) and (10.26).

I must warn you (if you have made it this far) that the particular lower bound I use for $1-(1-\frac{\varepsilon}{10d})^k$ is potentially very lossy. In particular, I use the fact that for $k \geq 1$, we have

$$1 - \left(1 - \frac{\varepsilon}{10d}\right)^k \geq 1 - \left(1 - \frac{\varepsilon}{10d}\right) = \frac{\varepsilon}{10d}. \tag{10.27}$$

Thus, to prove inequality (10.25), it is sufficent to prove that

$$A^k \leq \left(\frac{1}{\delta} \right)^{\frac{\varepsilon}{4}\left(\frac{\varepsilon}{10d}\right)}, \tag{10.28}$$

and to prove (10.26) it is sufficient to prove that

$$\left(\prod_{j=0}^{k-1} \log \left(\frac{1}{\delta_j} \right) \right)^B \leq \left(\frac{1}{\delta} \right)^{\frac{\varepsilon}{4}\left(\frac{\varepsilon}{10d}\right)}. \tag{10.29}$$

Since

$$k \lesssim \frac{\log\log(\delta')}{\varepsilon/(10d)}$$

we can write

$$A^k \leq (A')^{\frac{\log\log(\delta^{-1})}{\varepsilon/(10d)}}$$

for some A'. By taking logs, we see that it is sufficient for (10.28) to show that

$$\frac{\log\log(\delta^{-1})}{\varepsilon/(10d)} \log(A') \leq \frac{\varepsilon}{4}\left(\frac{\varepsilon}{10d}\right) \log(\delta^{-1}), \tag{10.30}$$

which we rearrange to read

$$\frac{\log\log(\delta^{-1})}{\log(\delta^{-1})} \leq \frac{\frac{\varepsilon}{4}\left(\frac{\varepsilon}{10d}\right)^2}{\log(A')} \tag{10.31}$$

Since $\frac{\log\log(\delta^{-1})}{\log(\delta^{-1})}$ approaches 0 as δ approaches 0, it follows that inequality (10.31) holds for all values of δ that are sufficiently small relative to ε.

Finally, we turn to proving inequality (10.29). Beginning with the left-hand side of (10.29), we have that

$$\begin{aligned}
\left(\prod_{j=0}^{k-1} \log(\delta_j^{-1})\right)^B &= \left(\prod_{j=0}^{k-1} \log(\delta^{-(1-\frac{\varepsilon}{10d})^j})\right)^B = \left(\prod_{j=0}^{k-1} \left(1-\frac{\varepsilon}{10d}\right)^j \log(\delta^{-1})\right)^B \\
&= \left(\left(1-\frac{\varepsilon}{10d}\right)^{\frac{k(k-1)}{2}} \log(\delta^{-1})^k\right)^B \\
&= \left(\left(1-\frac{\varepsilon}{10d}\right)^{\frac{k-1}{2}} \log(\delta^{-1})\right)^{kB}
\end{aligned}$$

Thus, to prove (10.26), we want to show that

$$\left(\left(1-\frac{\varepsilon}{10d}\right)^{\frac{k-1}{2}} \log(\delta^{-1})\right)^{kB} \leq \left(\frac{1}{\delta}\right)^{\frac{\varepsilon}{4}\left(\frac{\varepsilon}{10d}\right)}.$$

Taking logs, we see that this is equivalent to showing that

$$Bk \log\left(\left(1-\frac{\varepsilon}{10d}\right)^{\frac{k-1}{2}} \log(\delta^{-1})\right) \leq \frac{\varepsilon}{4}\left(\frac{\varepsilon}{10d}\right) \log(\delta^{-1}). \tag{10.32}$$

We note that

$$\left(1-\frac{\varepsilon}{10d}\right)^{\frac{k-1}{2}} \log(\delta^{-1}) \leq \log(\delta^{-1}). \tag{10.33}$$

Thus, to prove (10.32) it suffices to prove that

$$Bk \log\left(\log(\delta^{-1})\right) \leq \frac{\varepsilon}{4}\left(\frac{\varepsilon}{10d}\right)\log(\delta^{-1}) \tag{10.34}$$

Since

$$k \lesssim \frac{\log\log(\delta^{-1})}{\varepsilon/(10d)}$$

we can find a constant B' (not depending on ε) so that

$$Bk \log\left(\log(\delta^{-1})\right) \leq B'\left(\frac{\log\log(\delta^{-1})}{\varepsilon/(10d)}\right)\log\left(\log(\delta^{-1})\right)$$

Thus, to prove (10.34), it suffices to prove that

$$B'\left(\frac{\left(\log\log(\delta^{-1})\right)^2}{\varepsilon/(10d)}\right) \leq \frac{\varepsilon}{4}\left(\frac{\varepsilon}{10d}\right)\log(\delta^{-1}), \tag{10.35}$$

which we rearrange to read

$$\frac{\left(\log\log(\delta^{-1})\right)^2}{\log(\delta^{-1})} \leq \frac{\frac{\varepsilon}{4}\left(\frac{\varepsilon}{10d}\right)^2}{B'}. \tag{10.36}$$

Since $\frac{(\log\log(\delta^{-1}))^2}{\log(\delta^{-1})}$ approaches 0 as δ approaches 0, it follows that inequality (10.31) holds for all values of δ that are sufficiently small relative to ε. We will not work out a precise upper bound for δ here, but we will comment that whatever bound we *would* get could probably be improved if we had not used (10.33) and could be improved even more if we had used a sharper inequality than (10.27) in our earlier reductions.

Nevertheless, we have shown that for δ sufficiently small relative to ε- say, $\delta \leq c_\varepsilon$ for concreteness - then

$$A^k\left(\prod_{j=0}^{k-1}\log\left(\frac{1}{\delta_j}\right)\right)^B \delta^{\frac{\varepsilon}{2}(1-(1-\frac{\varepsilon}{10d})^k)} \leq 1$$

which means that inequality (10.24) implies the theorem bound for $\delta \leq c_\varepsilon$.

Meanwhile, so long as C_ε is chosen to be sufficiently large relative to c_ε, we can prove the theorem bound for $\delta > c_\varepsilon$ by the 'trivial' bound

$$|P_r(\Gamma)| \lesssim \delta^{-d}.$$

Appendix A

Proofs of Geometric Lemmas

A.1 Lemmas for Intersection Formulas

As a precursor to proving Lemma 1.1.1 from Chapter 1, we consider special cases in $\mathbb{R}^2$ and $\mathbb{R}^3$. We begin by estimating the area of the intersection of two rectangles in $\mathbb{R}^2$ of dimensions precisely $1 \times \delta$.

Lemma A.1.1. *If R_1, R_2 are two $\delta \times 1$ rectangles whose long sides form angle α to each other, then*

$$\mathit{Area}(R_1 \cap R_2) \leq \frac{\delta^2}{\sin \alpha}. \tag{A.1}$$

*Moreover, if R_1 and R_2 have the same center point, then equality holds in (*A.1*)*

Proof. It suffices to treat the case that R_1 and R_2 have the same center. In this case, the intersection R_1 and R_2 is a rhombus with two angles of measure α and two of measure $\pi - \alpha$. The altitude of the rhombus is δ from which it follows that each side has length $\delta/(\sin \alpha)$, so

$$\text{Area}(R_1 \cap R_2) \leq \frac{\delta^2}{\sin \alpha}.$$

□

More generally, if we allow R_1 and R_2 to belong to a broader class of rectangles with one side of length $c\delta \leq x \leq C\delta$ and one side of length $c \leq y \leq C$ and their long sides make angle $\alpha > 0$ to each other, then we have that

$$\text{Area}(R_1 \cap R_2) \lesssim \frac{\delta^2}{\sin \alpha} \sim \frac{\delta^2}{\alpha} \tag{A.2}$$

If we stipulate that R_1 and R_2 have the same center, then we have that

$$\text{Area}(R_1 \cap R_2) \sim \frac{\delta^2}{\alpha}.$$

We can use our estimates for rectangles in $\mathbb{R}^2$ to estimate the volume of the intersection of two slabs in $\mathbb{R}^3$.

Lemma A.1.2. *Let S_1, S_2 be two δ-slabs in $\mathbb{R}^3$ with normal vectors $\boldsymbol{n}_1$ and $\boldsymbol{n}_2$ which make an angle $\alpha > 0$ to each other. Then*

$$|S_1 \cap S_2| \lesssim \frac{\delta^2}{\sin \alpha}. \tag{A.3}$$

Moreover, if S_1 and S_2 were formed by taking the δ-neighborhoods of rectangles centered at a common point (e.g. $\mathbf{0}$), then

$$|S_1 \cap S_2| \sim \frac{\delta}{\sin \alpha}.$$

Proof. We first treat the special case that

1. S_1 and S_2 are rectangular boxes of dimensions precisely $1 \times 1 \times \delta$;

2. S_1 and S_2 have a common long direction $\boldsymbol{v}$, i.e. there is a unit vector $\boldsymbol{v}$ so that each of S_1 and S_2 have a side of length 1 that is parallel to $\boldsymbol{v}$.

In this special case, we will show that we have

$$|S_1 \cap S_2| \leq \frac{\delta^2}{\sin \alpha},$$

with equality if S_1 and S_2 have the same geometric center. To accomplish this, we let

$\Pi = \Pi_0$ be the orthogonal complement of $\boldsymbol{v}$ and consider slices of S_1 and S_2 by planes Π_x parallel to Π with $\Pi_x = \Pi + x\boldsymbol{v}$.

Let $\boldsymbol{w}_1$ be a unit vector parallel to the other long side of S_1, and let $\boldsymbol{w}_2$ be a unit vector parallel to the other long side of S_1. If the cross-section $S_1 \cap \Pi_x$ is non-empty, then it is a rectangle of dimensions precisely $1 \times \delta$ with long side parallel to $\boldsymbol{w}_1$ and short side parallel to $\boldsymbol{n}_1$. Similarly, if the cross-section $S_2 \cap \Pi_x$ is non-empty, then it is a rectangle of dimensions precisely $1 \times \delta$ with long side parallel to $\boldsymbol{w}_2$ and short side parallel to $\boldsymbol{n}_2$. Let $I = [x_0, x_1]$ be the interval consisting of the values of the x for which $S_1 \cap \Pi_x$ and $S_1 \cap \Pi_x$ are both non-empty. Then $|I| \leq 1$, and for any $x \in I$, we have

$$|(S_1 \cap \Pi_x) \cap (S_2 \cap \Pi_x)| \leq \frac{\delta^2}{\sin \alpha}. \tag{A.4}$$

We integrate over I to give

$$|S_1 \cap S_2| = \int_I |(S_1 \cap \Pi_x) \cap (S_2 \cap \Pi_x)|\, dx \leq \int_I \left(\frac{\delta^2}{\sin \alpha}\right) dx \leq \frac{\delta^2}{\sin \alpha}. \tag{A.5}$$

If S_1 and S_2 have the same center, then equality holds in (A.4) for every $x \in I$ and I is of length precisely 1, so equality holds in (A.5).

Having completed our work in the case that S_1 and S_2 satisfy requirements (1)-(2) above, we now consider the general case. Suppose now that S_1 and S_2 are slabs of dimensions $\sim 1 \times 1 \times \delta$ whose normal vectors $\boldsymbol{n}_1$ and $\boldsymbol{n}_2$ make angle α to each other.

The slabs S_1 and S_2 may not share a common long direction, but we will replace them by larger slabs $S_1' \supset S_1$ and $S_2' \supset S_2$ so that the following conditions hold:

(i) each S_i' is a prism with two sides of the same length;

(ii) S_i' has normal vector $\boldsymbol{n}_i$;

(iii) S_2' is an isometric copy of S_1' with a long side parallel to one of the long sides of S_1'.

We can find replacement slabs S_i' by considering the projections $\pi_i(S_i)$, where π_i denotes orthogonal projection onto the subspace $\boldsymbol{n}_i^{\perp}$. If the original slabs S_1 and S_2 belonged to a class of slabs each containing a prism of dimensions precisely $c \times c \times c\delta$ and contained in a prism of dimensions precisely $C \times C \times C\delta$, then for each slab S_i, we can choose a 1×1 rectangle $R_i \subset \boldsymbol{n}_i^{\perp}$ so that $cR_i \subset \pi_i(S_i) \subset CR_i$. We replace R_i be a square R_i' so that $CR_i \subset R_i'$ and one side of R_i' is parallel to $\boldsymbol{n}_1 \times \boldsymbol{n}_1$. We note that we can arrange that the side length of R_i' is at most $C\sqrt{2}$, because $\sqrt{2}$ is the length of the diagonal of R_i. We take S_i' to be a translation of the $C\delta$-neighborhood of R_i'.

Our choice of S_1' and S_2' ensures that they satisfy conditions (1)-(2) of the special case, modified so that each length 1 is replaced by a length $\ell \sim 1$. Specifically, we must have $\ell \leq C\sqrt{2}$.

Since S_1 and S_2 each have two long sides of length $\ell \sim 1$ and have a shared long direction $\boldsymbol{v}$, it follows that each nonempty slice of $S_i' \cap S_2'$ by a plane Π_x parallel to $\boldsymbol{v}^{\perp}$ has area satisfying

$$|(S_1' \cap \Pi_x) \cap (S_1' \cap \Pi_x)| \lesssim \frac{\delta^2}{\sin\alpha}. \tag{A.6}$$

Let I be the interval of $x \in \mathbb{R}$ for which $|(S_1' \cap \Pi_x) \cap (S_1' \cap \Pi_x)|$ is nonempty. Then I has length ~ 1. Since $S_1 \subset S_1'$ and $S_2 \subset S_2'$, we have that

$$|S_1 \cap S_2| \leq |S_1' \cap S_2'| \lesssim \frac{\delta^2}{\sin\alpha}. \tag{A.7}$$

Moreover, our work in bounding ℓ shows that, in fact $|S_1 \cap S_2| \sim |S_1' \cap S_2'|$. If S_1 and S_2 were formed by taking δ-neighborhoods of rectangles with have the same center, then the $\lesssim$ in (A.6) may be replaced by $\sim$ for each $x \in I$, and, consequently, (A.7) becomes

$$|S_1 \cap S_2| \sim \frac{\delta^2}{\sin\alpha}$$

□

The proposition we have just proved is the $d = 3$ version of Lemma 1.1.1, which we will now prove in full. First we recall the statement.

Lemma A.1.3. *(Lemma* 1.1.1*, revisited)*

Suppose that $S_1, S_2 \subset \mathbb{R}^d$ *are* δ*-slabs with respective normal vectors* $\boldsymbol{n_1}$ *and* $\boldsymbol{n_2}$*. If the angle between* $\boldsymbol{n_1}$ *and* $\boldsymbol{n_2}$ *is at least* α *then the volume of the intersection* $S_1 \cap S_2$ *satisfies*

$$|S_1 \cap S_2| \lesssim \frac{\delta}{\sin \alpha} \sim \frac{\delta}{\alpha}.$$

Moreover, if S_1 *and* S_2 *were formed by taking the* δ*-neighborhoods of rectangles centered at a common point (e.g.* $\mathbf{0}$*), then*

$$|S_1 \cap S_2| \sim \frac{\delta}{\sin \alpha}.$$

Proof. We closely follow the proof of Lemma A.1.2, above. We first treat the special case that

1. S_1 and S_2 are of dimensions precisely $1 \times \cdots \times 1 \times \delta$;

2. S_1 and S_2 have $d-2$-many long directions in common.

If assumptions (1) and (2) hold, we let $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{d-2}$ be unit vectors so that each of S_1 and S_2 has a long side parallel to each of $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{d-2}$.

For $j = 1, \ldots, d-2$, let B_j denote projection onto V_j, the 1-dimensional subspace spanned by $\boldsymbol{v}_j$. Let $I_j = B_j(S_1 \cap S_2)$ be the projection of $S_1 \cap S_2$ onto V_j. Each interval I_j has length ≤ 1. If S_1 and S_2 have the same center, then each interval I_j has length exactly 1.

We will express $|S_1 \cap S_2|$ as an integral over $I_1 \times \cdots \times I_{d-2}$, just as we wrote $|S_1 \cap S_2|$ as an integral over the interval I in (A.5) from the $d = 3$ case. For this, let $\Pi = \mathrm{Span}(\boldsymbol{n}_1, \boldsymbol{n}_1)$. Then each of $V_1, \ldots, V_{d-2}$ is contained in the orthogonal complement of Π.

We will slice each of S_1 and S_2 by planes parallel to Π. For each $(d-2)$-tuple $\vec{x} = (x_1, \ldots, x_{d-2}) \in V_1 \times \cdots \times V_{d-2}$, we let $\Pi_{\vec{x}} = \Pi + \vec{x}$ be the translation of Π by $\vec{x}$. Each nonempty slice $S_1 \cap \Pi_{\vec{x}}$ is a rectangle of dimensions precisely $1 \times \delta$ with short side parallel to $\boldsymbol{n}_1$, and each nonempty slice $S_2 \cap \Pi_{\vec{x}}$ is a rectangle of dimensions precisely $1 \times \delta$ with short side parallel to $\boldsymbol{n}_2$. The long sides of these rectangles make angle α to each other. By Lemma A.1.1, if both slices $S_1 \cap \Pi_{\vec{x}}$ $S_2 \cap \Pi_{\vec{x}}$ are nonempty, we have that

$$|(S_1 \cap \Pi_{\vec{x}}) \cap (S_2 \cap \Pi_{\vec{x}})| \leq \frac{\delta^2}{\sin \alpha} \tag{A.8}$$

with equality if $(S_1 \cap \Pi_{\vec{x}}) \cap (S_2 \cap \Pi_{\vec{x}})$ have the same center. We note that if S_1 and S_2 have the same center, then so do any slices $S_1 \cap \Pi_{\vec{x}}$ and $S_2 \cap \Pi_{\vec{x}}$.

Integrating over $I_1 \times \cdots \times I_d$, we see that

$$|S_1 \cap S_2| = \int_{I_1 \times \cdots \times I_{d-2}} |(S_1 \cap \Pi_{\vec{x}}) \cap (S_2 \cap \Pi_{\vec{x}})| \, d\vec{x} \leq \int_{I_1 \times \cdots \times I_{d-2}} \left(\frac{\delta^2}{\sin \alpha} \right) d\vec{x} \leq \frac{\delta^2}{\sin \alpha}. \tag{A.9}$$

If S_1 and S_2 have the same center, then equality holds in (A.8) for each $\vec{x} \in I_1 \times \cdots \times I_{d-2}$, and each interval I_j is of length precisely 1, so equality holds in (A.9) as well. This completes our proof under hypotheses (1)-(2).

For the general case, we suppose that S_1 and S_2 are slabs of dimensions $\sim 1 \times \cdots \times 1 \times \delta$ whose normal vectors $\boldsymbol{n}_1$ and $\boldsymbol{n}_2$ make angle α to each other. We note that $\boldsymbol{n}_1^{\perp} \cap \boldsymbol{n}_2^{\perp}$ contains a subspace V of dimension (d-2). We note that V is the orthogonal complement of the plane $\Pi = \text{Span}\{\boldsymbol{n}_1, \boldsymbol{n}_2\}$. We let $\{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{d-2}\}$ be an orthonormal basis for V. Finally, we let $\boldsymbol{w}_1$ and $\boldsymbol{w}_2$ be unit vectors so that $\boldsymbol{w}_1 \in V^{\perp} \cap \boldsymbol{n}_1^{\perp}$ and $\boldsymbol{w}_2 \in V^{\perp} \cap \boldsymbol{n}_2^{\perp}$

We will again find slabs S_1' and S_2' containing S_1 and S_2 so that the slabs S_1' and S_2' have $(d-2)$-many common long directions. We will take S_1' and S_2' to be prisms with respective normal vectors $\boldsymbol{n}_1$ and $\boldsymbol{n}_2$. Moreover, we will arrange so that S_i occupies a large proportion of the volume of S_i'. We can accomplish this by considering the projections $\pi(S_i)$, where π_i denotes orthogonal projection onto $\boldsymbol{n}_i^{\perp}$.

If S_1 and S_2 belonged to a class of slabs which each contained a prism of dimensions precisely $c \times \cdots \times c \times c\delta$ and was contained in a prism of dimensions $C \times \cdots \times C \times C\delta$ for some $0 < c \leq 1 \leq C$, then for each S_i, we can choose a $1 \times \cdots \times 1$ cube $R_i \subset \boldsymbol{n}_{\boldsymbol{i}}^{\perp}$ so that $cR_i \subset \pi_i(S_i) \subset CR_i$. We replace R_i by a cube R_i' with sides parallel to $\boldsymbol{v}_1, \ldots \boldsymbol{v}_{d-2}, \boldsymbol{w}_i$. We choose R_i' so that R_i' has the same center as R_i and contains CR_i. We can also arrange that the side length of R_i' is at most $C\sqrt{d-1}$. We can accomplish this by taking the inscribed sphere of R_i' to be the circumscribed sphere of CR_i, which has diameter $C\sqrt{d-1}$. Having chosen R_i' in this way, we then take S_i' to be an appropriate translation of the $C\delta$-neighborhood of R_i'.

Our chosen S_1' and S_2' are prisms of dimensions precisely $\ell \times \cdots \times \ell \times C\delta$ with $\ell \leq C\sqrt{d-1}$. If $\vec{x} \in V$ and $\Pi_{\vec{x}} := \Pi_0 + \vec{x}$ has nonempty intersection with $S_1' \cap S_2'$, then the slice $|(S_1' \cap \Pi_{\vec{x}}) \cap (S_2' \cap \Pi_{\vec{x}})|$ is a rhombus of area $\sim \frac{\delta^2}{\sin\alpha}$. The set of $\vec{x} \in V$ for which $\Pi_{\vec{x}} := \Pi_0 + \vec{x}$ has nonempty intersection with $S_1' \cap S_2'$ is a rectangle $I_1 \times \cdots \times I_{d-2}$ with sides parallel to $\boldsymbol{v}_1, \dots \boldsymbol{v}_1$ so that each side has length $\leq C\sqrt{d-1}$. We integrate over this rectangle to give

$$\begin{aligned} |S_1 \cap S_2| \leq |S_1' \cap S_2'| &= \int_{I_1 \times \cdots \times I_{d-2}} |(S_1' \cap \Pi_{\vec{x}}) \cap (S_2' \cap \Pi_{\vec{x}})| \, d\vec{x} \\ &\lesssim \int_{I_1 \times \cdots \times I_{d-2}} \frac{\delta^2}{\sin\alpha} \, d\vec{x} \sim \frac{\delta^2}{\sin\alpha}. \end{aligned}$$

If S_1 and S_2 had been the respective δ-neighborhoods (or $C\delta$-neighborhoods) of rectangles centered at a common point, then the $\lesssim$ in the above chain of inequalities could be replaced by a $\sim$ to give

$$|S_1 \cap S_2| \sim \frac{\delta^2}{\sin\alpha}.$$

□

Lemma A.1.4. *(Lemma 1.1.2, revisited)*

Suppose that $S_1, \dots, S_d \subset \mathbb{R}^d$ are δ-slabs with respective normal vectors $\boldsymbol{n_1}, \dots, \boldsymbol{n_d}$, Then the volume of their d-fold intersection satisfies

$$|S_1 \cap \cdots \cap S_d| \lesssim \min\left\{ \frac{\delta^d}{|\boldsymbol{n_1} \wedge \cdots \wedge \boldsymbol{n_d}|}, \ \delta \right\}, \tag{A.10}$$

where $\boldsymbol{n_1} \wedge \cdots \wedge \boldsymbol{n_d}$ denotes the determinant of the $d \times d$ matrix that has $\boldsymbol{n_1}, \dots, \boldsymbol{n_d}$ as its rows. Moreover, if $S_1, \dots, S_d$ were formed by taking the δ-neighborhoods of rectangles *centered at a common point (e.g.* $\mathbf{0}$*) and* $\frac{\delta^d}{|\boldsymbol{n_1}\wedge\cdots\wedge\boldsymbol{n_d}|} \leq \delta$*, then*

$$|S_1 \cap \cdots \cap S_d| \sim \frac{\delta^d}{|\boldsymbol{n_1} \wedge \cdots \wedge \boldsymbol{n_d}|}. \tag{A.11}$$

Proof. If

$$\delta \leq \frac{\delta^d}{|\boldsymbol{n_1} \wedge \cdots \wedge \boldsymbol{n_d}|},$$

then (A.10) is a consequence of the following chain of inequalities:

$$|S_1 \cap \cdots \cap S_d| \leq |S_1| \lesssim \delta.$$

Henceforth, we assume that

$$\frac{\delta^d}{|\boldsymbol{n_1} \wedge \cdots \wedge \boldsymbol{n_d}|} \leq \delta.$$

If $S_1, \ldots, S_d$ are not centered at a common point, we let be $S_1', \ldots, S_d'$ be translations $S_1, \ldots, S_d$ which *are* centered at a common point. Then we will have

$$|S_1 \cap \cdots \cap S_d| \leq |S_1' \cap \cdots \cap S_d'| \sim \frac{\delta^d}{|\boldsymbol{n_1} \wedge \cdots \wedge \boldsymbol{n_d}|}$$

Hence, it suffices to consider the case that $S_1, \ldots, S_d$ were formed by taking the δ-neighborhoods of rectangles centered at a common point and to show that (A.11) holds in this case. We assume for convenience that the common point is $\mathbf{0}$ so that the slab S_i is contained in the δ-neighborhood of the subspace $\{\boldsymbol{n_i}\}^\perp$.

Let N be the matrix with rows $\boldsymbol{n_1}, \ldots, \boldsymbol{n_d}$. A vector $\boldsymbol{x}$ is in slab S_i if and only if $|\boldsymbol{x} \cdot \boldsymbol{n_i}| \leq \delta$, so the image of $S_1 \cap \cdots \cap S_d$ under multiplication by N is the box $[-\delta, \delta]^d$, i.e. we have that

$$N(S_1 \cap \ldots S_d) = [-\delta, \delta]^d.$$

Therefore,

$$|\det N||S_1 \cap \cdots \cap S_d| = \left|[-\delta, \delta]^d\right|,$$

from which it follows that

$$|S_1 \cap \cdots \cap S_d| \sim \frac{\delta^d}{|\det N|} = \frac{\delta^d}{|\boldsymbol{n_1} \wedge \cdots \wedge \boldsymbol{n_d}|}.$$

$\square$

Bibliography

[1] Tom M. Apostol, *Introduction to Analytic Number Theory*, Springer, 1976.

[2] Jean Bourgain and Ciprian Demeter, *The proof of the ℓ^2 Decoupling Conjecture*, Annals of Math. 182 (2015), no. 1, 351-389.

[3] Ciprian Demeter, Larry Guth, and Hong Wang; *Small cap decouplings*, Geom. Funct. Anal., 30(4):989-1062, 2020. With an appendix by D. R. Heath-Brown.

[4] György Elekes and Micha Sharir, *Incidences in three dimensions and distinct distances in the plane*, Combinatorics, Probability and Computing, Volume 20, Issue 4, July 2011, pp. 571 - 608.

[5] György Elekes and Csaba D. Tóth, *Incidences of not-too-degenerate hyperplanes*, Proceedings of the twenty-first annual symposium on Computational geometry, 2005.

[6] Yuqiu Fu and Kevin Ren, *Incidence estimates for α-dimensional tubes and β-dimensional balls in $\mathbb{R}^2$*, `https://arxiv.org/abs/2111.05093`

[7] Larry Guth and Nets Hawk Katz, *On the Erdös distinct distances problem in the plane*, Ann. Math. 181 (2015), 155-190

[8] Larry Guth and Dominique Maldague, *Small cap decouplings for the moment curve in $\mathbb{R}^3$*, 2022

[9] Larry Guth, Dominique Maldague, and Hong Wang, *Improved decoupling for the parabola*, 2020

[10] Larry Guth, Noam Solomon, and Hong Wang, *Incidence estimates for well-spaced tubes*, Geom. Funct. Anal. Vol. 29 (2019).

[11] Adam Sheffer, *The Two Formulations of the Szemerédi–Trotter Theorem*, https://adamsheffer.wordpress.com/2015/07/05/the-two-formulations-of-the-szemeredi-trotter-theorem/

[12] Endre Szemerédi and William T. Trotter, *Extremal problems in discrete geometry*, Combinatorica 3, 381–392 (1983).

[13] Csaba D. Tóth, *The Szemerédi-Trotter theorem in the complex plane*, Combinatorica 35, 95-126 (2015). https://doi.org/10.1007/s00493-014-2686-2

[14] Sarah Tammen and Lingxian Zhang, *Incidence estimates for tubes in complex space*, https://arxiv.org/abs/2210.05477, 2022

[15] Le Anh Vinh, *The Szemeredi-Trotter type theorem and the sum-product estimate in finite fields*, European. J. Combin. 32 (2011), no. 8, 1177-1181.