

**From the holdings of the Department of Special Collections , Princeton University Library,
Princeton, NJ**

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyright law. We do not charge any permission or usage fees for the publication of images of material in our collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff. For more information regarding copyright, credit, and citation please review the information on the Copyright, Credit and Citation Guidelines page of our website.

<https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines>. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. <https://library.princeton.edu/special-collections/ask-us>

Doctoral Dissertations, 2011-2026

Mathematics

2025

Combinatorics with a view towards number theory

Kravitz, Noah

Princeton, NJ : Princeton University

<http://arks.princeton.edu/ark:/88435/dsp017m01bq11x>

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. <https://www.copyright.gov/title17/title17.pdf>

COMBINATORICS
WITH A VIEW TOWARDS
NUMBER THEORY

NOAH KRAVITZ

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
MATHEMATICS
ADVISER: NOGA ALON

MAY 2025

© Copyright by Noah Kravitz, 2025.

All Rights Reserved

Abstract

We study several problems in arithmetic combinatorics, loosely grouped around the themes of “combinatorial number theory” and “higher-order Fourier analysis”. The problems related to the former theme are the Lonely Runner Problem (from Diophantine approximation), questions about minimal additive complements, and a conjecture of Graham about rearrangements of subsets of $\mathbb{F}_p^\times$. The problems related to the latter theme concern quantitative bounds for some instances of the Bergelson–Leibman Theorem (about subsets of $\mathbb{N}^d$ avoiding polynomial progressions).

Acknowledgements

This thesis is a culmination of my work over the past five years at Princeton University. I am grateful to have had the opportunity to participate in such a vibrant mathematical community.

The first thanks go to my adviser, Noga Alon, who has taught and supported me throughout my PhD journey. I would not have reached this point without Noga's patience, advice, technical expertise, sense of humor, and willingness to think about the many strange open problems that I would bring to our weekly meetings. Noga somehow struck the perfect balance between gently nudging me in productive mathematical directions and letting me find my own faltering way.

I am fortunate to have had several unofficial mentors at Princeton. I am particularly grateful to Peter Sarnak, for making me an honorary number theorist; Sarah Peluse, for initiating me into higher-order Fourier analysis; Zeev Dvir, for teaching me about discrete geometry; and Matija Bucić, for engaging in interesting conversations on a variety of topics.

One advantage of graduate study at Princeton is the constant inspiration of working alongside so many talented, motivated peers. Let me highlight just two of the many. I learned how to read and digest math papers during the long-term reading projects that Katy Woo and I undertook in our first three years at Princeton. And I owe a great deal to Anshul Advé for patiently explaining so many things to me at the whiteboard.

Another important mathematical community for me has been the Duluth REU program. I cannot sufficiently express my gratitude to Joe Gallian and Colin Defant for allowing me to continue to play a role in the program. Working with the wonderful students in Duluth is what cemented my desire to pursue a career in academic mathematics.

I am grateful to my previous mathematical mentors for their continued support, enthusiasm, and advice: Stefan Steinerberger, my undergraduate adviser at Yale; and Andy Lipps, my high school calculus teacher.

Of course, the work on which this thesis is based would have been impossible without all of my brilliant collaborators, both at Princeton and elsewhere.

I also wish to thank my non-mathematical mentors and colleagues at Princeton—there is, after all, a limit to how many hours per day one can spend on math. Alan Feinberg has been like a second adviser to me, and our weekly piano lessons have been a mainstay of my time at Princeton. On the other side of the golf course, I have derived great joy from playing the carillon in Cleveland Tower under the guidance of Lisa Lonie. Closer to Fine Hall, I am grateful to the Near Eastern Studies department for welcoming me with open arms into classes, seminars, and informal reading groups. I have particularly benefitted from studying with George Kiraz, Johannes Haubold, and Lara Harb. And no list of extracurricular activities would be

complete without a shout-out to the Princeton Running Club.

Finally, I thank my family for supporting me through all of these years and humoring my mathematical (mis)adventures. Even if you never fully get the hang of how mathematicians use “epsilon” in ordinary conversation, I appreciate the effort.

My graduate work has been financially supported by a Graduate Research Fellowship from the National Science Foundation (grant DGE-203965) and by Frances Lane and Charlotte Elizabeth Procter Fellowships from Princeton.

To Alan Reinberg, inventeur extraordinaire,
whose creative impulses I hope are visible in this thesis.

Contents

Abstract	3
Acknowledgements	4
Foreword	12
0.1 Arithmetic combinatorics	12
0.2 Combinatorial number theory	13
0.2.1 Lonely runners	13
0.2.2 Minimal complements	14
0.2.3 Rearrangements	15
0.3 Higher-order Fourier analysis	16
0.3.1 Polynomial configurations	16
0.3.2 A pattern of interest	17
0.3.3 Quantitative concatenation	18
0.3.4 Polynomial corners	19
0.4 Sources	20
I Combinatorial number theory	22
1 Lonely runners: the qualitative view	23
1.1 Introduction	23
1.1.1 Lonely runners	23
1.1.2 Maximum loneliness spectra	24
1.1.3 Accumulation points	25
1.1.4 View-obstruction, subtori, and subgroups	25
1.1.5 Main result	26

1.1.6	Checking small speeds suffices	27
1.1.7	Organization	28
1.2	Relation to similar problems	28
1.2.1	Abelian covers	29
1.2.2	The Markoff spectrum	30
1.3	Geometry of cubes	30
1.4	Quantitative Kronecker Theorem	34
1.5	Producing accumulation points	37
1.6	Putting everything together	39
1.7	Checking small speeds suffices	40
1.8	Calculations for small (co)dimension	41
1.9	Remarks and open problems	43
2	Lonely runners: the quantitative view	45
2.1	Introduction	45
2.1.1	The Lonely Runner Conjecture	45
2.1.2	Relative Lonely Runner spectra	46
2.1.3	Main result	48
2.1.4	Comparison with previous work	49
2.1.5	Computing relative spectra	49
2.1.6	Future directions	50
2.1.7	Organization	51
2.2	Proof of Theorem 2.1	51
2.2.1	Proof strategy	51
2.2.2	Intersections of subtori and subgroups	52
2.2.3	Approximating minima	56
2.2.4	Coprime pairs in arithmetic progressions	63
2.2.5	Remarks on the proof	67
2.3	Identifying candidate subtori U	69
2.3.1	General setup	70
2.3.2	Subtori $U \subset (\mathbb{R}/\mathbb{Z})^3$ with $D(U) = 1/10$	70
2.3.3	Subtori $U \subset (\mathbb{R}/\mathbb{Z})^4$ with $D(U) = 1/4$	72
2.4	The top of the spectrum $\mathcal{S}_1(4)$	75

2.4.1	Computing $\mathcal{S}_1(U^1)$	75
2.4.2	Computing $\mathcal{S}_1(U^2)$	78
2.5	The spectrum $\mathcal{S}_1(3)$ to the second accumulation point	83
2.5.1	Computing $\mathcal{S}_1(U^3)$	84
2.5.2	Computing $\mathcal{S}_1(U^4)$	85
2.5.3	Computing $\mathcal{S}_1(U^5)$	87
2.5.4	Computing $\mathcal{S}_1(U^6)$	89
2.6	A new progression in the spectrum $\mathcal{S}_1(6)$	89
2.7	Finite and infinite relative spectra	94
3	Minimal additive complements	97
3.1	Introduction	97
3.1.1	Background	97
3.1.2	Main results	98
3.1.3	Open questions	100
3.2	Preliminaries	101
3.3	Small sets are minimal complements	104
3.4	Many large sets are not minimal complements	107
3.5	Maximal supplements	110
4	Graham's rearrangement conjecture	114
4.1	Introduction	114
4.1.1	Main result	114
4.1.2	Proof sketch and organization	115
4.2	Notation and parameters	116
4.3	Structure theorem	117
4.4	Ordering P and N	122
4.5	Splitting the dissociated sets	127
4.6	Randomizing the dissociated sets	132
4.7	Remarks and open problems	137
II	Higher-order Fourier analysis	138
5	Quantitative concatenation	139

5.1	Introduction	139
5.2	Proof strategy	147
5.2.1	Comparison with previous approaches	152
5.3	Definitions and preliminaries	154
5.3.1	Basic notation	154
5.3.2	Van der Corput Inequalities	156
5.3.3	Box norms	157
5.4	PET induction scheme	164
5.4.1	Example of PET	165
5.4.2	Setting up the induction scheme	169
5.4.3	Bounds by averages of box seminorms	178
5.5	Model example: concatenation of degree-1 box norms	180
5.5.1	Sketch of the proof of Proposition 5.24	183
5.6	Concatenation of averages of higher-degree seminorms	188
5.7	Equidistribution estimates for concatenation	196
5.8	Concatenation along polynomials	209
5.9	Proof of Theorem 5.6	214
5.10	One more PET result	215
6	Polynomial corners	218
6.1	Introduction	218
6.1.1	Future directions	222
6.2	Proof strategy	223
6.2.1	Organization	228
6.3	Definitions and preliminaries	229
6.3.1	Basic notation	229
6.3.2	W -trick notation	230
6.3.3	Van der Corput Inequalities	231
6.3.4	Box norms	232
6.3.5	Dual-difference interchange	234
6.3.6	Fourier analysis	236
6.4	Degree-lowering	236
6.4.1	Proof of Proposition 6.12: preliminary maneuvers	239

6.4.2	Examining the structure of a	243
6.4.3	Auxiliary norm control on f_1	246
6.4.4	Completing the proof of Proposition 6.12	249
6.5	Reducing Proposition 6.13 to a nilsequence problem	252
6.6	Nilsequence application	257
6.6.1	Completing the proof of Proposition 6.13	261
6.7	Finishing the argument	265
6.8	Inverse theorems for box norms	271
6.8.1	U^1 -inverse theorems	271
6.8.2	Unnormalized box norms	273
6.8.3	U^2 -inverse theorems	274
6.8.4	$U^1 \times U^1$ -inverse theorems	275
6.8.5	$U^2 \times U^1$ -inverse theorem	276
6.9	Nilsequence definitions	277
6.10	Fraction-comparison estimates	281
6.11	Sárközy-type results	286
6.11.1	Circle method preliminaries	286
6.11.2	Applications to Sárközy-type estimates	291
6.12	Fourier uniformity of difference of weights	295
6.13	Supersaturation	299

Bibliography	308
---------------------	------------

Foreword

0.1 Arithmetic combinatorics

A drab but descriptive title for my thesis would have been “Topics in arithmetic combinatorics”. The title that I eventually settled on—less drab, even if equally un-catchy—is a nod to Einsiedler and Ward’s *Ergodic Theory with a View Towards Number Theory*, the first math book that I ever read cover-to-cover. As advertised, this book introduces the central notions of ergodic theory and dynamics, with an emphasis on applications to objects of number-theoretic significance. In the same vein, I, primarily a combinatorialist by training, have attempted in this thesis to bring hands-on combinatorial approaches to bear on problems number-theoretic in nature.

The field of combinatorics primarily concerns the study of patterns. The subfield of arithmetic combinatorics aims to understand patterns built from the arithmetic operations of addition and multiplication. The objects of study are number-theoretic, the questions have a combinatorial flavor, and the methods draw on many other areas of math, including harmonic analysis, ergodic theory, discrete geometry, and Ramsey theory. Problems in arithmetic combinatorics are often easy to state but surprisingly difficult to solve.

Although there are some big guiding problems in the field, much progress comes from looking at the objects of study from a variety of perspectives and asking concrete, highly focused questions that give glimpses into different parts of their structure. Of particular interest to me are cases where one “type” of structure implies the existence of another “type” of structure.

In this thesis, I will highlight two themes that together encompass much of my work in arithmetic combinatorics. The first theme, “combinatorial number theory”, traditionally refers to combinatorially-oriented questions about properties of natural numbers. The second theme, “higher-order Fourier analysis”, refers to the outgrowth of the tools originally developed by Gowers in his new proof of Szemerédi’s Theorem. The former is disposed towards short, creative arguments, while the latter involves the depolyment of more “high-tech” tools where much of the effort lies in working out technical details.

0.2 Combinatorial number theory

The first half of my thesis concerns three different problems, each of which gives you a set of numbers and asks you to find “something else” that makes the numbers behave in a particular way. The Lonely Runner Conjecture gives you a set of runner speeds and asks you to find times when the runners are all far away from the start line. The problem of minimal additive complements gives you a subset of an abelian group and asks you to find some other subset such that the two sets together “cover” the whole group efficiently. Graham’s rearrangement conjecture gives you a set of residues modulo a prime and asks you to find an ordering of the residues so that the partial sums are all distinct.

The techniques for making progress on these problems are varied, but they are all “elementary”, in the sense that one builds one’s tools from scratch, with bare hands. In this style, once one alights on the correct approach, the execution of the strategy tends to be fairly quick.

0.2.1 Lonely runners

I begin with two chapters about the Lonely Runner Conjecture of Wills and Cusick. This conjecture originated in Diophantine approximation but has a purely combinatorial formulation: Consider n runners standing at the start line of a unit-length circular track, and suppose that they simultaneously begin running around the track, each at a nonzero constant speed. The Lonely Runner Conjecture predicts that regardless of the runners’ speeds, there will always be some time t when all of the runners are a distance at least $1/(n+1)$ away from the start line. (We can imagine that one extra runner, who stays at the start line, gets “lonely” when the other runners are all far away.) The quantity $1/(n+1)$ would be tight, with equality achieved, for instance, when the runners have the speeds $1, 2, \dots, n$.

More formally, let $\|x\|$ denote the distance from a real number x to the nearest integer. For natural numbers $v_1, \dots, v_n$, define the *maximum loneliness* to be the quantity

$$\text{ML}(v_1, \dots, v_n) := \max_{t \in \mathbb{R}} \min_{1 \leq i \leq n} \|tv_i\|,$$

where the inner minimum records the distance from the start line to the closest runner at time t ; in this language, the Lonely Runner Conjecture is the assertion that $\text{ML}(v_1, \dots, v_n) \geq 1/(n+1)$ for all choices of $v_1, \dots, v_n$.¹ In my undergraduate thesis, I suggested studying not only the minimum value of $\text{ML}(v_1, \dots, v_n)$ for each n but also the entire *Lonely Runner spectrum* $S(n)$ consisting of all of the values of $\text{ML}(v_1, \dots, v_n)$. This perspective opens the vista for asking many new types of questions, for instance, about the behavior

¹That it suffices to consider only natural number speeds is an old result of Bohman, Holzman, and Kleitman.

of the near-minimal elements of $S(n)$.

The first chapter of this thesis, based on joint work with Vikram Giri, contains a “qualitative” study of the accumulation points of Lonely Runner spectra. In the “geometric view-obstruction” formulation of the Lonely Runner Problem, one associates each tuple $(v_1, \dots, v_n)$ of natural numbers with the 1-dimensional subtorus $T = \langle (v_1, \dots, v_n) \rangle_{\mathbb{R}}$ of the n -dimensional torus $(\mathbb{R}/\mathbb{Z})^n$. A short calculation reveals that $\text{ML}(v_1, \dots, v_n)$ is equal to $1/2$ minus the L^∞ -distance from T to the point $(1/2, \dots, 1/2)$ at the “center” of $(\mathbb{R}/\mathbb{Z})^n$. The innovation in our work is considering not only 1-dimensional subtori of $(\mathbb{R}/\mathbb{Z})^n$ but also higher-dimensional subtori, as well as subgroups. Our main result relates the accumulation points of $S(n)$ to the maximum loneliness values (defined via the L^∞ -distance to $(1/2, \dots, 1/2)$) of higher-dimensional subtori of $(\mathbb{R}/\mathbb{Z})^n$. This result bears on an “inductive” relationship among the Lonely Runner spectra $S(n)$ for different values of n .

The next chapter, based on joint work with Vanshika Jain, presents a “quantitative” complement which characterizes the universal structure of accumulation points in Lonely Runner spectra. For a 2-dimensional subtorus U of $(\mathbb{R}/\mathbb{Z})^n$, consider the *relative Lonely Runner spectrum* $S(U)$ consisting of all values of $\text{ML}(v_1, \dots, v_n)$ such that $\langle (v_1, \dots, v_n) \rangle_{\mathbb{R}}$ is contained in U . In general, the set $S(U)$ has a unique accumulation point at the maximum loneliness value of U . Since of course $S(n)$ contains every such $S(U)$, this gives a way to produce accumulation points of $S(n)$; the upshot of the previous chapter is that this is actually the “main” source of accumulation points of $S(n)$. The second chapter shows that, up to finitely many exceptional values, each relative Lonely Runner spectrum $S(U)$ is a finite union of sets of the form

$$\{\text{ML}(U) - 1/(\alpha s + \beta) : s \in \mathbb{N}\},$$

where $\text{ML}(U)$ is the maximum loneliness of U and α, β are positive rationals. This structure theorem exhibits a surprising arithmetic rigidity in the behavior of (relative) Lonely Runner spectra near accumulation points. Moreover, the structure constants α, β can be calculated explicitly for each U . Using this approach, we are able provide explicit descriptions of several “chunks” of Lonely Runner spectra for small n , and we produce several families of interesting new examples of maximum loneliness values.

0.2.2 Minimal complements

The third chapter, based on joint work with Noga Alon and Matt Larson, concerns minimal additive complements in abelian groups. This topic is related to Sidon sets, additive bases, and covering codes. Let W

be a subset of an abelian group G . A set C is a *complement* for W if $C + W = G$, where

$$C + W := \{c + w : c \in C, w \in W\}$$

denotes the Minkowski sum. A complement C is *minimal* for W if no proper subset of C is a complement for W . These notions were first introduced by Nathanson, who studied which sets W have minimal complements. We consider matters from the other side and study the “inverse” problem (first raised by Kwon) of determining which sets C can appear as minimal complements when W ranges over the subsets of G . Our main result shows that when G is finite, every subset C of size at most around $|G|^{1/3}/(\log |G|)^{2/3}$ is a minimal complement for some W ; one consequence is that when G is infinite, every finite subset C is a minimal complement for some W . Our proofs use probabilistic arguments to produce sets W for which our given C is a minimal complement. We also establish a converse statement, namely, that every finite abelian group G has a subset C of size around $|G|^{3/4}$ that is not a minimal complement for any subset W ; the construction is again probabilistic.

0.2.3 Rearrangements

The first part of this thesis concludes with a chapter, based on joint work with Benjamin Bedert, about Graham’s rearrangement conjecture. For p a (large) prime, let $\mathbb{F}_p$ denote the finite field with p elements. A conjecture of Graham, later repeated by Erdős, asserts that if A is any subset of $\mathbb{F}_p \setminus \{0\}$, then there is an ordering $a_1, a_2, \dots, a_{|A|}$ of the elements of A such that the partial sums

$$a_1, \quad a_1 + a_2, \quad a_1 + a_2 + a_3, \quad \dots, \quad a_1 + \dots + a_{|A|}$$

are all distinct. Previous work had established this conjecture when the set A is very large (of size at least $p - p^{1/1000}$) or very, very small (of size at most 12). The work of this chapter is to establish Graham’s conjecture when A is small-ish, namely, of size at most $e^{(\log p)^{1/4}}$.

A simple, but previously overlooked, “rectification” argument shows that to establish the conjecture when A has size at most $\log p / \log \log p$, it suffices to solve the analogous problem in the integers; the latter is fairly straightforward. The key new idea for reaching the quasi-polynomial threshold $e^{(\log p)^{1/4}}$ is a structure theorem that decomposes any small-ish subset A of $\mathbb{F}_p$ into several large-ish dissociated sets (sets with “little” additive structure) and a rectifiable leftover piece. The remainder of the argument combines probabilistic tools with a careful analysis of the integer version of the conjecture.

0.3 Higher-order Fourier analysis

The second half of my thesis concerns a cluster of problems in arithmetic Ramsey theory, where a typical problem asks which arithmetic structures inevitably appear in all sufficiently large subsets of the natural numbers (or of $\mathbb{N}^d$). We will be interested in the setting where the “arithmetic structures” in question are multidimensional patterns defined by polynomials and the “sufficiently large” sets are those with positive upper Banach density.

0.3.1 Polynomial configurations

Let us be a bit more precise. Given vectors $\mathbf{v}_1, \dots, \mathbf{v}_\ell \in \mathbb{Z}^d$ and integer-coefficient polynomials $P_1, \dots, P_\ell \in \mathbb{Z}[z]$, consider the pattern

$$\mathbf{x}, \quad \mathbf{x} + P_1(z)\mathbf{v}_1, \quad \dots, \quad \mathbf{x} + P_\ell(z)\mathbf{v}_\ell, \tag{*}$$

for $\mathbf{x} \in \mathbb{Z}^d$ and $z \in \mathbb{Z}$. For example, when $d = 1$, $\mathbf{v}_1 = \dots = \mathbf{v}_\ell = 1$, and $P_j(z) = jz$, we get the $(\ell + 1)$ -term arithmetic progression

$$x, \quad x + z, \quad x + 2z, \quad \dots, \quad x + \ell z.$$

In general, let $\lambda(N) = \lambda_{\mathbf{v}_1, \dots, \mathbf{v}_\ell; P_1, \dots, P_\ell}(N)$ denote the maximum possible size of a subset of $[N]^d = \{1, 2, \dots, N\}^d$ that *avoids* containing copies of the pattern $(*)$.² One says that a pattern $(*)$ is *inevitable* if $\lambda(N) = o(N^d)$, i.e., if $\lambda(N)/N^d$ tends to 0 as N tends to infinity.

Some patterns fail to be inevitable because of modular constraints called “local obstructions”. For example, the set of odd numbers up to N has constant density $1/2$ in $[N]$ but does not contain any copy of the pattern $x, x + 1$, or any copy of the pattern $x, x + 100z^2 - 7$, so for each of these patterns we have $\lambda(N) \geq N/2$, which is not $o(N)$. The Polynomial Szemerédi Theorem of Bergelson and Leibman says that this is just about the only thing that can go wrong: If every P_j satisfies $P_j(0) = 0$, then the associated pattern $(*)$ is inevitable. This extremely general theorem contains as special cases the celebrated results of Roth (when the pattern is a 3-term arithmetic progression), Szemerédi (when the pattern is a k -term arithmetic progression for any fixed k), and Furstenberg and Sárközy (when the pattern is $x, x + z^2$).

A limitation of the Bergelson–Leibman Theorem is that, because of its reliance on methods from ergodic theory, it delivers only qualitative bounds. That is, the theorem tells us that $\lambda(N)/N^d$ tends to 0 as N grows, but it gives no hint of how quickly. The problem of obtaining “reasonable quantitative bounds” for the decay

²Typically one wants to stipulate that the occurrence of the pattern be *nontrivial*, in the sense that the $P_j(z)$ ’s are all nonzero.

of $\lambda(N)/N^d$ —even for particular patterns of interest—has assumed a life of its own in recent decades. The main approaches to this research program are based on the Gowers uniformity norms and the concomitant tools of “higher-order Fourier analysis” (HOFA).

HOFA is markedly non-elementary, and there is much material to master before one dives into open problems; this is the area of math in which I have acquired the most “specialized technical skill” during my PhD. The methods behind HOFA rely on a diverse array of techniques involving objects such as Gowers norms and nilsequences that are too complicated to be defined in this overview. Proofs usually require the careful fitting-together of intricate details and the implementation (and adaptation) of tools from the literature, even after one has identified a promising overarching proof strategy.

0.3.2 A pattern of interest

The focus of the second half of this thesis, based on joint work with Borys Kuca and James Leng, is using HOFA methods to establish quantitative bounds for subsets of $[N]^2$ avoiding the 2-dimensional “polynomial corner configuration”

$$(x, y), \quad (x + P(z), y), \quad (x, y + P(z)),$$

where $P \in \mathbb{Z}[z]$ is any fixed polynomial with a simple integer root (a fully representative example is $P(z) = z^2 - z$). The Bergelson–Leibman Theorem tells us that $\lambda(N) = o(N^2)$ for this pattern, and our main result upgrades this qualitative statement to the quantitative bound

$$\lambda(N) = O_P \left(\frac{N^2}{(\log \log \log N)^c} \right), \tag{*}$$

for some explicit constant $c > 0$. There are several reasons to be interested in quantitative bounds for the polynomial corner configuration:

- Our result is the first, along with concurrent work of Peluse, Prendiville, and Shao, to establish quantitative bounds for a pattern that is *both* multidimensional and non-linear.
- One can think of a corner as the more challenging cousin of the 3-term arithmetic progression—the “arithmetic pattern” *par excellence*—and on this view our result extends recent work of Peluse, Sah, and Sawhney on the pattern $x, x + z^2 - 1, x + 2(z^2 - 1)$.
- Most importantly, our proof required the development of several new technical tools, some of which have already found applications both elsewhere in higher-order Fourier analysis and in other fields.

The proof breaks down roughly into two stages, each of which is the topic of one chapter in this thesis.

The first stage is the development of general “quantitative concatenation” machinery for box norms, which pertains to a broad class of problems. The second stage is a “transference argument” for polynomial corners using the method of “degree-lowering”.

0.3.3 Quantitative concatenation

The story behind this chapter begins with Roth’s Theorem on 3-term arithmetic progressions. Suppose one wishes to show that every “sufficiently large” subset A of $[N]$ contains a 3-term arithmetic progression. Let $1_A : [N] \rightarrow \mathbb{C}$ be the indicator function of A , which equals 1 on A and vanishes otherwise, and let $f_A := 1_A - \alpha$ be the “balanced” indicator function, where $\alpha := |A|/N$ is the density of the set A . Using algebraic manipulations, one can express the number of 3-term arithmetic progressions in A in terms of α and the Fourier transform $\widehat{f_A}$ of the function 1_A . Roth’s key insight was noticing a “randomness versus structure” dichotomy depending on whether or not $\widehat{f_A}$ is ever large: If $\widehat{f_A}$ is pointwise everywhere small (the “random” case), then the number of 3-term arithmetic progressions in A is roughly the same as the expected number of 3-term arithmetic progressions in a random subset of $[N]$ of size αN ; in particular, there are many such progressions. If instead $\widehat{f_A}$ is somewhere large (the “structured” case), then we can “zoom in” on an arithmetic progression in $[N]$ on which the set A has increased density $\alpha' > \alpha$; we then repeat the argument from the beginning, using this new “denser” set, and iterate until we find a 3-term arithmetic progression.

For $k > 3$, one can still express the number of k -term arithmetic progressions in A in terms of α and $\widehat{f_A}$, but it is no longer true that the pointwise smallness of $\widehat{f_A}$ is enough to deliver a good estimate on the number of k -term arithmetic progressions. The seminal work of Gowers introduced a new way of measuring the randomness versus structure dichotomy, under which notion a “random-like” set A does have roughly the same number of k -term arithmetic progressions as a truly random subset of $[N]$ of the same size. This method of measuring uses what are now called the Gowers uniformity norms.³

One can attempt to set up Gowers’s machinery for a more complicated configuration such as the polynomial corner configuration, using what it called “PET induction”. The trouble is that one often obtains not a “nice” uniformity norm but rather an unruly average of “box norms”. Concatenation, originally introduced by Tao and Ziegler, is a general principle that lets one replace such an average of box norms by a single better-behaved box norm. In the past, instances of the concatenation principle have been worked out on an ad-hoc basis as needed. Our contribution is establishing a general quantitative concatenation result that covers all averages of box norms arising from multidimensional polynomial configurations. This general technical tool will be useful for future work on the Bergelson–Leibman Theorem, and it has already found

³The Gowers uniformity norm related to 3-term arithmetic progressions is essentially the ℓ^4 -norm of $\widehat{f_A}$. By analogy, the study of the Gowers norms for longer arithmetic progressions is termed “higher-order” Fourier analysis.

some applications outside of arithmetic Ramsey theory.

A crucial step in our proof is establishing nearly-optimal anticoncentration for certain large systems of multilinear polynomial expressions. For example, in the simplest toy case, one might consider the expression hk , where h, k both range over $[H]$ for some large H . There are H^2 pairs (h, k) , and the expression hk is certainly contained in $[H^2]$. If hk were distributed roughly uniformly over the range $[H^2]$, then we would expect that for each integer n , there are at most $O(H^2/H^2) = O(1)$ pairs (h, k) such that $hk = n$. Unfortunately, this is not the case, since n 's with many divisors occur disproportionately many times. Thankfully, there is a trick, based on the Cauchy–Schwarz Inequality, that lets us replace a system of multilinear polynomials with another system in which some of the variables have been “doubled”. In our toy example, we can replace hk by $hk + h'k'$ (not yet good enough) and then by

$$h_1k_1 + h_2k_2 + h_3k_3 + h_4k_4,$$

where still the h_i 's and k_i 's range over $[H]$. It turns out that now for each n , there are $O(H^8/H^2) = O(H^6)$ choices of the h_i 's and k_i 's such that this last expression equals n . The main difficulty is generalizing this procedure when hk is replaced by a more complicated polynomial and (especially) when we must simultaneously massage several different polynomials into convenient forms.

0.3.4 Polynomial corners

In the final chapter, we return to the polynomial corner configuration and establish the bound $(\star)$. The main idea, called *transference*, is that we wish to compare the number of polynomial corners in a given set $A \subseteq [N]^2$ with the number of *ordinary* corners

$$(x, y), (x + z, y), (x, y + z)$$

in the same set; we can then conclude by invoking a result of Shkredov that provides an upper bound on the size of a corner-free subset of $[N]^2$. To execute this transference strategy, we draw on a wide range of tools and ideas from HOFA, including the W -trick, our PET induction and quantitative concatenation machinery, degree-lowering, equidistribution of polynomial sequences on nilmanifolds, and the circle method.

The first important observation is that, even after suitable normalizations, the number of polynomial corners cannot be directly compared to the number of ordinary corners. The issue is that the image of the polynomial $P(z)$ (the side length for our polynomial corners) grows sparser at larger scales and can be biased modulo small primes, whereas the image of the identity function z (the side length of our ordinary

corners) does not share either of these features. To address the first discrepancy, we count ordinary corners with a new “weight” that discounts corners with large side length. To address the second discrepancy, we replace the polynomial $P(z)$ by a so-called “ W -tricked” version $\tilde{P}(z)$ which behaves mostly like $P(z)$ but has better equidistribution properties modulo small primes. Miraculously, these two small changes suffice to let us compare the counts of polynomial corners and ordinary corners.

The central object of study is the *polynomial corner counting operator*

$$\Lambda(f_0, f_1, f_2) := \sum_{x, y, z \in \mathbb{Z}} f_0(x, y) f_1(x + P(z), y) f_2(x, y + P(z)),$$

for functions $f_0, f_1, f_2 : [N]^2 \rightarrow \mathbb{C}$ with pointwise absolute value at most 1. When $f_0 = f_1 = f_2 = 1_A$, the quantity $\Lambda(1_A, 1_A, 1_A)$ is precisely the number of polynomial corners in the set A , but it is important to work in the setting of arbitrary functions. After some fairly standard maneuvers, what we need to show is that the operator Λ is “controlled” by suitable box norms. The required statement is roughly that if $|\Lambda(f_0, f_1, f_2)|$ is large (say, at least a constant δ times $\Lambda(1_{[N]}, 1_{[N]}, 1_{[N]})$), then the $U^1 \times U^1$ box norm

$$\|f_0\|_{U^1 \times U^1[N]}^4 := \sum_{x, y, h, k} f_0(x, y) \overline{f_0(x + h, y)} \overline{f_0(x, y + k)} f_0(x + h, y + k)$$

is also large (say, at least some function of δ times $\|1_{[N]}\|_{U^1 \times U^1[N]}^4$). The output of PET induction and quantitative concatenation (from the previous chapter) is that we can conclude the largeness of a higher-degree box norm $\|f_0\|_{U^t \times U^t[N]}^{2^{2t}}$, for some natural number t (depending only on the polynomial P).

The main work consists of gradually decreasing t all the way to 1; this strategy, first introduced by Peluse, is called “degree-lowering”. Previous degree-lowering procedures in the literature worked with Gowers uniformity norms, and our procedure faces new challenges because we must work instead with multidimensional box norms. Our argument is quite delicate. Here I will say only that it runs the full gamut of HOFA tools and tricks, including direct work with nilsequences.

0.4 Sources

The work presented in this thesis is based on the following publications, with only topical changes:

- Chapter 1: V. Giri and N. Kravitz, The structure of Lonely Runner spectra. *Preprint* arXiv:2411.12684v2 (2025).
- Chapter 2: V. Jain and N. Kravitz, Relative Lonely Runner spectra. *Preprint* ArXiv:2304.01462v4 (2025).

- Chapter 3: N. Alon, N. Kravitz, and M. Larson, Inverse problems for minimal complements and maximal supplements. *J. Number Theory*, Vol. 223, pp. 307-324 (2021).
- Chapter 4: B. Bedert and N. Kravitz, Graham's rearrangement conjecture beyond the rectification barrier. *Israel J. Math.*, to appear (arXiv:2409.07403v2, 2025).
- Chapter 5: N. Kravitz, B. Kuca, and J. Leng, Quantitative concatenation for polynomial box norms. *Preprint* arXiv:2407.08636v1 (2024).
- Chapter 6: N. Kravitz, B. Kuca, and J. Leng, Corners with polynomial side length. *Preprint* arXiv:2407.08637v1 (2024).

I thank my coauthors for permitting me to include our joint work in this thesis. See the original papers for further acknowledgements.

As the various chapters of this thesis come from stand-alone papers, they are self-contained and can be read independently. In particular, even though Chapters 1 and 2 are best appreciated together, as are Chapters 5 and 6, I have repeated the necessary background and definitions in each chapter.

Part I

Combinatorial number theory

Chapter 1

Lonely runners: the qualitative view

1.1 Introduction

1.1.1 Lonely runners

The Lonely Runner Problem of Wills [113] and Cusick [28] is based on the following setup. Suppose $n + 1$ runners start at the same position on a unit-length circular track and then begin running around the track at pairwise distinct constant speeds. We say that a runner is *lonely* at a particular time if their distance from every other runner is at least $1/(n + 1)$. The Lonely Runner Conjecture predicts that each runner will be lonely at some time (allowing different times for different runners).

Consider the frame of reference of a single runner, and suppose that the n other runners have speeds $v_1, \dots, v_n$ relative to our fixed runner. Define the *maximum loneliness* for this set of speeds to be

$$\text{ML}(v_1, \dots, v_n) := \sup_{t \in \mathbb{R}} \min_{1 \leq i \leq n} \|tv_i\|_{\mathbb{R}/\mathbb{Z}},$$

where $\|x\|_{\mathbb{R}/\mathbb{Z}}$ denotes the distance from x to the nearest integer. Reformulated in this language, the Lonely Runner Conjecture asserts that

$$\text{ML}(v_1, \dots, v_n) \geq 1/(n + 1)$$

for all nonzero real numbers $v_1, \dots, v_n$. Using Kronecker's Equidistribution Theorem, Bohman, Holzman, and Kleitman [17] showed that it suffices to prove the Lonely Runner Conjecture for integer speeds; we will revisit this perspective below the fold.

The Lonely Runner Problem has attracted substantial attention since Wills first posed it in 1967. Due to the work of many authors [4, 10, 12, 17, 30, 92], the Lonely Runner Conjecture is known to hold for $n \leq 6$;

unfortunately, none of the approaches to these small- n cases seem to generalize to an arbitrary number of runners. The trivial bound $\text{ML}(v_1, \dots, v_n) \geq 1/(2n)$ has also been improved slightly (see [108] and the references therein), and the conjecture has been proven in the case where all of the runners have slow speeds (e.g., [18, 89, 95, 108]). For other work on and around the Lonely Runner Problem, see the references in [65].

1.1.2 Maximum loneliness spectra

In this chapter, we will be concerned with the perspective introduced in the author's thesis [65]. The new motivating question was: What can we say about the set $\tilde{\mathcal{S}}(n)$ of *all* of the possible values of $\text{ML}(v_1, \dots, v_n)$? For example, what is the second-smallest element after (conjecturally) $1/(n+1)$? Can we describe all of $\tilde{\mathcal{S}}(n)$ explicitly? More qualitatively, what do the accumulation points of $\tilde{\mathcal{S}}(n)$ look like? Notice that each $\tilde{\mathcal{S}}(n)$ is a subset of the interval $(0, 1/2]$. We will see later that $\tilde{\mathcal{S}}(n)$ is a closed subset of the rationals.

The focus of the work [65] was the following conjecture about the “bottom” part of the set $\tilde{\mathcal{S}}(n)$.

Conjecture 1.1 (Loneliness Spectrum Conjecture ([65], Conjecture 1.2)). *For every natural number $n \geq 2$, we have*

$$\tilde{\mathcal{S}}(n) \cap (0, 1/n) = \{s/(ns+1) : s \in \mathbb{N}\}.$$

One should think of the quantity $s/(ns+1)$ as the result of “rounding down” $1/n$ to the nearest multiple of $1/(ns+1)$. An explicit construction from [65] shows that all of the values $s/(ns+1)$ are contained in $\tilde{\mathcal{S}}(n)$. The conjecture has no content for $n=1$ and is easily verified for $n=2$. One of the main results of [65] is a proof of the conjecture for $n=3$.

Using computer experiments, Fan and Sun [39] discovered a family of examples that disproves Conjecture 1.1 for $n=4$. In particular, they showed that

$$\text{ML}(8, 4r+3, 4r+11, 4r+19) = \frac{2r+7}{4(2r+7)+2}$$

for every integer $r \geq 0$. This provides a second infinite family of maximum loneliness values limiting to $1/n$ (at least for $n=4$) in a structured way. Following the “ $1/n$ rounded down” philosophy, Fan and Sun proposed the following natural weakening of Conjecture 1.1.

Conjecture 1.2 (Amended Loneliness Spectrum Conjecture ([39], Conjecture 1.3)). *For every natural number $n \geq 2$, we have*

$$\tilde{\mathcal{S}}(n) \cap (0, 1/n) \subseteq \{s/(ns+k) : s \in \mathbb{N}, 1 \leq k \leq n\}.$$

1.1.3 Accumulation points

On a qualitative level, both Conjecture 1.1 and its weakening Conjecture 1.2 predict that $1/n$ is the smallest accumulation point of $\tilde{\mathcal{S}}(n)$. It is no coincidence that this number $1/n$ coincides with the conjectural smallest value of $\tilde{\mathcal{S}}(n-1)$. Recall that, for a subset $\mathcal{S}$ of $\mathbb{R}$, the number $x \in \mathbb{R}$ is an *upper accumulation point* if $\mathcal{S} \cap (x, x + \varepsilon) \neq \emptyset$ for all $\varepsilon > 0$; we define *lower accumulation points* analogously. We write $\text{acc}(\mathcal{S})$ for the set of (all) accumulation points of $\mathcal{S}$. The following theorem forms the starting point for our investigations.

Theorem 1.3 ([65], Theorem 6.5). *Let $n \geq 2$ be a natural number. Then the set of lower accumulation points of $\tilde{\mathcal{S}}(n)$ contains $\tilde{\mathcal{S}}(n-1)$.*

Also in [65], we raised the questions of whether upper accumulation points are impossible, and whether the set containment in Theorem 1.3 is actually an equality. The goal of the present chapter is to give an affirmative answer to the first part of this question and make partial progress towards the second part. Along the way, our arguments will show that the sets $\tilde{\mathcal{S}}(n)$ and some closely related sets have a “hierarchical” structure; this information is potentially useful for inductive approaches to the Lonely Runner Conjecture, which so far have seemed inaccessible because of a failure to relate the instances of the conjecture for different numbers of runners. Before stating our results precisely, we need to introduce a few geometric notions.

1.1.4 View-obstruction, subtori, and subgroups

It will be convenient for us to work with the “view-obstruction” formulation of the Lonely Runner Problem, as popularized by Cusick [28]. Associate each tuple $(v_1, \dots, v_n)$ of nonzero integers with the 1-dimensional subtorus

$$T := \pi(\langle (v_1, \dots, v_n) \rangle_{\mathbb{R}}) = \{t(v_1, \dots, v_n) : t \in \mathbb{R}\} / \mathbb{Z}^n$$

of the torus $(\mathbb{R}/\mathbb{Z})^n$, where $\pi : \mathbb{R}^n \rightarrow (\mathbb{R}/\mathbb{Z})^n$ is the standard quotient map. Notice that T is not contained in any of the coordinate hyperplanes of $(\mathbb{R}/\mathbb{Z})^n$; in general, we say that a subtorus is *proper* if it is not contained in the union of the coordinate hyperplanes. Let $D(T)$ denote the L^∞ -distance from T to the point $(1/2, \dots, 1/2)$ in the “center” of $(\mathbb{R}/\mathbb{Z})^n$, so that

$$D(T) = 1/2 - \text{ML}(v_1, \dots, v_n).$$

Now define the *n-th Lonely Runner spectrum* to be the set $\mathcal{S}(n)$ of all values achieved by $D(T)$ as T ranges over the 1-dimensional proper subtori of $(\mathbb{R}/\mathbb{Z})^n$; we have

$$\mathcal{S}(n) = 1/2 - \tilde{\mathcal{S}}(n).$$

In this new language, yet another reformulation of the Lonely Runner Conjecture is the assertion that $\mathcal{S}(n) \subseteq [0, 1/2 - 1/(n+1)]$. Notice that lower accumulation points for $\tilde{\mathcal{S}}(n)$ correspond to upper accumulation points for $\mathcal{S}(n)$, and vice versa.

The more geometric line of inquiry arising from view-obstruction has been fruitfully related to the study of polytopes, and the “zonotope” version of the Lonely Runner Conjecture has received significant attention (see, e.g., [5, 6, 11, 55] and the references therein).

The function D and the notion of properness also make sense for arbitrary closed subsets $X \subseteq (\mathbb{R}/\mathbb{Z})^n$. In particular, define $D(X)$ to be the L^∞ -distance from X to the point $(1/2, \dots, 1/2)$, and say that X is *proper* if it is not contained in the union of the coordinate hyperplanes.¹ We will have occasion to study the values assumed by $D(X)$ not only when X is a 1-dimensional proper subtorus but also when X is a higher-dimensional subtorus or a closed subgroup.

To this end, for $1 \leq k \leq n$, let $\mathcal{S}_k(n)$ be the set of all values achieved by $D(T)$ as T ranges over the k -dimensional proper subtori of $(\mathbb{R}/\mathbb{Z})^n$. Notice that $\mathcal{S}(n) = \mathcal{S}_1(n)$ in this notation. We will also need to work with closed subgroups of $(\mathbb{R}/\mathbb{Z})^n$; in this chapter, “subgroup” always means “closed subgroup”. Every subgroup $\Delta \subseteq (\mathbb{R}/\mathbb{Z})^n$ is the direct product of a k -dimensional subtorus of $(\mathbb{R}/\mathbb{Z})^n$ and a finite subgroup of $(\mathbb{R}/\mathbb{Z})^n$, and we say that the dimension of Δ is k . Now, for $0 \leq k \leq n$, let $\mathcal{S}_k^*(n)$ denote the set of all values achieved by $D(\Delta)$ as Δ ranges over the k -dimensional proper subgroups of $(\mathbb{R}/\mathbb{Z})^n$. We make the convention $\mathcal{S}_0^*(0) := \{0\}$, and we write $\mathcal{S}^*(n)$ for $\mathcal{S}_1^*(n)$. Since every (proper) k -dimensional subtorus is also a (proper) k -dimensional subgroup, we have the inclusion $\mathcal{S}_k(n) \subseteq \mathcal{S}_k^*(n)$, which in general is strict (see below).

Finally, we remark that one can think of $\mathcal{S}_k(n)$ as a multiset by keeping track of the values of $D(T)$ for all k -dimensional proper subtori $T \subseteq (\mathbb{R}/\mathbb{Z})^n$; we will denote this multiset by $\mathcal{S}_{k,\text{mult}}(n)$. The multiset $\mathcal{S}_{k,\text{mult}}^*(n)$ is obtained from $\mathcal{S}_k^*(n)$ analogously. The *density points* of a multiset $\mathcal{S}$ are defined to be the accumulation points of $\mathcal{S}$ together with the infinite-multiplicity elements of $\mathcal{S}$. We write $\text{den}(\mathcal{S})$ for the set of density points of $\mathcal{S}$.

1.1.5 Main result

We are finally ready to state our main result, which is summarized by the following relationships among the sets $\mathcal{S}_k(n), \mathcal{S}_k^*(n)$.

Theorem 1.4. *Let $n \geq 2$ be a natural number, and let $1 \leq k < n$ and $0 \leq k' < n$. Then the sets $\mathcal{S}_k(n), \mathcal{S}_{k'}^*(n)$*

¹This condition is equivalent to $D(X) < 1/2$, and it is slightly stronger than saying that X is not contained in any single coordinate hyperplane.

have only upper accumulation points. On the level of sets we have the inclusions

$$\text{acc}(\mathcal{S}_k(n)) \subseteq \mathcal{S}_{k+1}(n) \subseteq \mathcal{S}_{k+1}^*(n) = \mathcal{S}_0^*(n - k - 1),$$

and on the level of multisets we have the equalities

$$\text{den}(\mathcal{S}_{k,\text{mult}}(n)) = \mathcal{S}_{k+1}(n) \quad \text{and} \quad \text{den}(\mathcal{S}_{k',\text{mult}}^*(n)) = \mathcal{S}_{k'+1}^*(n) = \mathcal{S}_{k'}^*(n - 1).$$

Moreover, when $k = 1$, we also have the inclusion $\mathcal{S}(n - 1) \subseteq \text{acc}(\mathcal{S}(n))$, which implies that

$$\mathcal{S}(n - 1) \subseteq \text{acc}(\mathcal{S}(n)) \subseteq \mathcal{S}_2(n) \subseteq \mathcal{S}_2^*(n) = \mathcal{S}^*(n - 1). \quad (1.1)$$

The natural question, of course, is which of the inclusions (1.4) are actually equalities. Equality holds in (1.4) for $n = 2$ trivially; it also holds for $n = 3$ because (see Section 1.8 below) we can explicitly compute

$$\mathcal{S}(2) = \mathcal{S}^*(2) = \{0\} \cup \{1/(4s + 2) : s \in \mathbb{N}\}.$$

Equality cannot hold in general, however; we will show in Proposition 1.17 that $7/50 \in \mathcal{S}^*(3) \setminus \mathcal{S}(3)$. Nonetheless, we conjecture that equality holds in Theorem 1.3.

Conjecture 1.5. *For any natural number $n \geq 2$, we have $\text{acc}(\mathcal{S}(n)) = \mathcal{S}(n - 1)$.*

To prove this conjecture, it would suffice to show that $\mathcal{S}_2(n) = \mathcal{S}(n - 1)$, i.e, that the first two inclusions in (1.4) are equalities. More generally, we pose the following problem, which we suspect has an affirmative answer.

Problem 1.6. *Is it the case that $\mathcal{S}_k(n) = \mathcal{S}_{k-\ell}(n - \ell)$ for all natural numbers $1 \leq \ell < k \leq n$?*

Theorem 1.4 has a couple of nice consequences. One is that, due to the trivial containments $\mathcal{S}_k(n) \subseteq \mathcal{S}_{k-1}(n)$ and $\mathcal{S}_{k'}^*(n) \subseteq \mathcal{S}_{k'-1}^*(n)$, the sets $\mathcal{S}(n), \mathcal{S}^*(n)$ are closed for all n . Another consequence is that each of $\mathcal{S}(n), \mathcal{S}^*(n)$ is a well-ordered set with order type $\omega^{n-1} + 1$.

1.1.6 Checking small speeds suffices

The maximum loneliness function is easy to compute, and it is natural to ask whether the Lonely Runner Conjecture can be verified by checking that there are no counterexamples with all speeds below some explicit threshold. Tao [108] showed that this question can be answered in the affirmative, and that it suffices to check all speeds up to n^{Cn^2} for some explicitly computable constant C . Recently, Malikiosis, Santos, and

Schymura [76] used the zonotope formulation of the Lonely Runner Conjecture to improve this threshold to $(n(n+1)/2)^{n-1} \approx n^{2n}$. The new perspective afforded by Theorem 1.4 leads to a simple and transparent proof of the same result with the slightly weaker bound $n^{(5/2)n}$ (see below for our precise bound). It is interesting to note that Minkowski’s Theorem on successive minima plays a role both for us and for Malikiosis, Santos, and Schymura.

1.1.7 Organization

In Section 1.2, we put the Lonely Runner spectra in their proper context alongside other “bass note spectra” such as the Markoff spectrum, and we sketch an interpretation in terms of abelian Bloch wave theory; this perspective was suggested to us by Peter Sarnak. We prove our main result Theorem 1.4 over the course of the following four sections: Section 1.3 includes some geometric lemmas about cubes; Section 1.4 establishes a quantitative Kronecker-type result about high-volume subtori; Section 1.5 shows how to produce accumulation points in Lonely Runner spectra; and Section 1.6 assembles the pieces. In Section 1.7 we explain the computability result mentioned in Section 1.1.6, and in Section 1.8 we provide several examples and characterizations in low-dimensional cases. Finally, we raise some prospects for future work in Section 1.9.

1.2 Relation to similar problems




Figure 1.1: These three sketches illustrate the behaviour of various bass-note spectra. Top: The Lonely Runner spectrum $\mathcal{S}(2)$ is a “rigid” spectrum which is infinite and discrete with the unique accumulation point 0. Middle: The Markoff spectrum has a continuous (“flexible”) bottom part, a discrete (“rigid”) top part, and a fractal transition part. Bottom: The Lonely Runner spectrum $\mathcal{S}(3)$ looks complicated, but it exhibits the “hierarchical” structure that its set of accumulation points is $\mathcal{S}(2)$.

1.2.1 Abelian covers

There is a natural generalization of the setup described in Section 1.1.4. For any continuous function $f : (\mathbb{R}/\mathbb{Z})^n \rightarrow \mathbb{R}$, we can define $D_f(T)$ to be the minimum value of f assumed on the subtorus $T \subseteq (\mathbb{R}/\mathbb{Z})^n$. Our $D(T)$ from above corresponds to the function $f(x) = \|x - (1/2, \dots, 1/2)\|_\infty$. One can further define $\mathcal{S}_{f,k}(n)$ to be the set of all values assumed by $D_f(T)$ where T is a k -dimensional proper subtorus of $(\mathbb{R}/\mathbb{Z})^n$, and it is interesting to study the qualitative structure of such sets. (For some choices of f , it may make more sense to drop the “properness” condition.)

Abelian Bloch wave theory provides a fertile source of functions f . Let M be a geometric object such as a manifold or a graph, and let $\widetilde{M}$ be its universal abelian cover. Denote by $\pi_1(M)$ the fundamental group of M . The *character torus* of M is

$$T_M := \{\text{characters } \chi : \pi_1(M) \rightarrow \mathbb{C} \text{ such that } \chi(\gamma) = 1 \text{ for all } \gamma \in \pi_1(\widetilde{M})\}$$

(which really is a torus in the usual sense of the term). Finite abelian covers of M correspond to finite subgroups of T_M , and infinite abelian covers correspond to closed subgroups of T_M of dimension at least 1. The universal abelian cover of M corresponds to the entire character torus. The overarching principle of Bloch wave theory (or Floquet theory) is that spectral properties of an operator on an abelian cover M' of M can be understood in terms of “twisted” versions of the operator on the base space M ; the subgroup of T_M corresponding to M' dictates which character twists to consider. In many cases of interest (such as when the operator is the Laplacian), these twisted operators are perturbations of the original operator, and Kato’s perturbation theory (see [60]) tells us that quantities of interest vary continuously as we move along the character torus. We can then define the function $f(\chi)$ to be the value of such a quantity for the twisted operator corresponding to the point χ . For example, taking $f(\chi)$ to be the spectral gap of the χ -twisted Laplacian on M results in the “bass-note spectrum” of M . See Section 3 of [63] for an overview of how Bloch wave theory applies to the study of graph spectra; see also [98] and the work [1] of Adve and Giri for applications to spinor spectra of closed Riemann surfaces; and see [41] for a standard treatment of Bloch wave theory from a physics perspective.

Our results and techniques in the context of the Lonely Runner Problem may also provide a path to understanding similar qualitative phenomena in spectra of differential operators on abelian covers. Even if there turns out not to be a formal connection, it is helpful to situate the Lonely Runner Problem in this broader context.

1.2.2 The Markoff spectrum

In Section 1.2.1 we described a question about (simple) geodesics in the flat torus $(\mathbb{R}/\mathbb{Z})^n$. One can generalize this setup to geodesics in other locally symmetric spaces.

One particularly nice example comes from the locally symmetric space $X = \mathbb{H}^2/\Gamma$, where $\mathbb{H}^2 := \{x + iy : y > 0\}$ is the complex upper half-plane with the hyperbolic metric and Γ is a finite-index subgroup of the modular group $\mathrm{PSL}_2(\mathbb{Z})$ (which acts on $\mathbb{H}^2$ by Möbius transformations). Given a continuous function f on X , we can define the function $D_{X,f}(\gamma)$ to be the minimum value of f achieved on the closed (simple) geodesic γ , and one can ask about the set $\mathcal{S}_{X,f}$ of values of $D_{X,f}(\gamma)$.

The top part of the famous *Markoff spectrum* (which is related to the Lagrange spectrum from Diophantine approximation) can be described in this language: It equals $\mathcal{S}_{X,f}$ for a certain choice of a congruence subgroup Γ and a function f for which $D_{X,f}(\gamma)$ measures how far γ travels into the “cusp” of X (i.e., the “distance from γ to the point at infinity”). The structure of the Markoff spectrum is quite complicated (see Figure 1.1) and remains far from understood. The top part of the spectrum is discrete and the bottom consists of a continuous interval (the so-called Hall–Freiman ray); there is a “fractal” region in the middle. See [101] and Chapter 7 of [29] for overviews.

One can also consider the locally symmetric spaces $\mathrm{PGL}_n(\mathbb{R})/\mathrm{PGL}_n(\mathbb{Z})$ for $n \geq 3$. In many number-theoretic applications of this setup, it is natural to consider collections of geodesics (so-called “packets”) or certain families of closed torus orbits rather than individual geodesics. The *measure rigidity conjectures* would imply that any infinite sequence of distinct such packets must eventually go arbitrarily deep into the cusp, and this would imply the rigidity of the corresponding spectra. See [35–37] for more in this direction.

1.3 Geometry of cubes

We begin with some short “combinatorial” lemmas about how subtori and subgroups of $(\mathbb{R}/\mathbb{Z})^n$ can intersect cubes centered at $(1/2, \dots, 1/2)$.

Lemma 1.7. *Let $n \geq 2$ and $0 \leq k \leq n$ be natural numbers. Then all elements of $\mathcal{S}_k^*(n)$ (and a fortiori of $\mathcal{S}_k(n)$ for $k \geq 1$) are rational numbers. Moreover, for each k -dimensional proper subgroup Δ of $(\mathbb{R}/\mathbb{Z})^n$, there is a point $q \in \Delta$ with all rational coordinates such that $D(\Delta) = \|q - (1/2, \dots, 1/2)\|_\infty$.*

Proof. Let Δ be a k -dimensional proper subgroup of $(\mathbb{R}/\mathbb{Z})^n$. Then there exist vectors $u_1, \dots, u_k \in \mathbb{Z}^n$ and $v_1, \dots, v_\ell \in \mathbb{Q}^n$ such that

$$\Delta = \bigcup_{1 \leq j \leq \ell} \pi(\mathbb{R}u_1 + \dots + \mathbb{R}u_k + v_j).$$

The point $p = \alpha_1 u_1 + \dots + \alpha_k u_k + v_j$ has image $\pi(p) = p - [p]$ (with $[p]$ taken coordinate-wise) in the fundamental domain $[0, 1]^n$ of $(\mathbb{R}/\mathbb{Z})^n$, and hence the L^∞ -distance from p to $(1/2, \dots, 1/2)$ in $(\mathbb{R}/\mathbb{Z})^n$ is

$$\|p - [p] - (1/2, \dots, 1/2)\|_{L^\infty(\mathbb{R}^n)}.$$

Since the expression inside the L^∞ -norm is unchanged if the α_i 's are shifted by integers (recall that the u_i 's have all integer coordinates), we have

$$D(\Delta) = \min_{1 \leq j \leq \ell} \min_{\alpha_1, \dots, \alpha_k \in [0, 1]} \left\| \sum_i \alpha_i u_i + v_j - \left\lfloor \sum_i \alpha_i u_i + v_j \right\rfloor - (1/2, \dots, 1/2) \right\|_{L^\infty(\mathbb{R}^n)}.$$

As j and $\alpha_1, \dots, \alpha_k$ range, the quantity $\lfloor \sum_i \alpha_i u_i + v_j \rfloor$ ranges over finitely many vectors in $\mathbb{Z}^n$, say, $m_1, \dots, m_R \in \mathbb{Z}^n$. Then we can write

$$D(\Delta) = \min_{1 \leq r \leq R} \min_{1 \leq j \leq \ell} \min_{\alpha_1, \dots, \alpha_k \in [0, 1]} \left\| \sum_i \alpha_i u_i + v_j - m_r - (1/2, \dots, 1/2) \right\|_{L^\infty(\mathbb{R}^n)}.$$

The expression inside the L^∞ -norm is a rational affine-linear function of the α_i 's for each choice of r, j , so we conclude that the set of α_i 's attaining the minimum is a finite union of rational polytopes (i.e., polytopes defined by linear equalities and inequalities with rational coefficients). Hence some minimizer has all rational coordinates, and it follows that $D(\Delta)$ is also rational. $\square$

It is clear that $\mathcal{S}_{k-\ell}^*(n-\ell) \subseteq \mathcal{S}_k^*(n)$ for all $1 \leq \ell \leq k \leq n$: If Δ is a $(k-\ell)$ -dimensional proper subgroup of $(\mathbb{R}/\mathbb{Z})^{n-\ell}$ with $D(\Delta) = d$, then $\Delta' := \Delta \times (\mathbb{R}/\mathbb{Z})^\ell$ is a k -dimensional proper subgroup of $(\mathbb{R}/\mathbb{Z})^n$ with $D(\Delta') = d$. Somewhat surprisingly, this inclusion is an equality. It is this “hierarchical” structure that makes the function D particularly nice (among all of the possible functions D_f described in Section 1.2.1). One should think of the equality $\mathcal{S}_k^*(n) = \mathcal{S}_{k-\ell}^*(n-\ell)$ as saying that the set of attainable values of $D(\Delta)$ depends on only the codimension of the subgroups Δ under consideration.

Lemma 1.8. *For natural numbers $1 \leq \ell \leq k \leq n$, we have $\mathcal{S}_k^*(n) = \mathcal{S}_{k-\ell}^*(n-\ell)$.*

Proof. The result is trivial for $k = n$, so suppose that $k < n$. By the previous remark, it suffices to show that $\mathcal{S}_k^*(n) \subseteq \mathcal{S}_{k-\ell}^*(n-\ell)$. It is clear that $0 \in \mathcal{S}_{k-\ell}^*(n-\ell)$, as witnessed by, for instance, the subgroup

$$\underbrace{\{(0, \dots, 0)\}}_{n-k}, \underbrace{\{(1/2, \dots, 1/2)\}}_{n-k} \times (\mathbb{R}/\mathbb{Z})^{k-\ell} \subseteq (\mathbb{R}/\mathbb{Z})^{n-\ell}.$$

Now let $d \in \mathcal{S}_k^*(n)$ with $d > 0$, and let Δ be a k -dimensional proper subgroup of $(\mathbb{R}/\mathbb{Z})^n$ such that $D(\Delta) = d$.

Let B denote the box in $(\mathbb{R}/\mathbb{Z})^n$ consisting of the points whose L^∞ -distance from $(1/2, \dots, 1/2)$ is d . The properness of Δ ensures that $d < 1/2$, so B really is a “box”. We claim that Δ intersects a face of B of dimension at most $n - k - 1$. Indeed, let F be a lowest-dimensional face of B that intersects Δ . If $\dim(F) > n - k$, then the intersection of F and Δ contains a line segment by naive dimension-counting, and we can follow this line segment to find an intersection of Δ with a lower-dimensional face of B . If $\dim(F) = n - k$, then again the intersection of F and Δ contains a line segment, since otherwise Δ would intersect F transversely and pass into the interior of B , which is impossible. So we conclude that $\dim(F) \leq n - k - 1$.

Let p be an intersection point of Δ with a face of B of dimension at most $n - k - 1$. Then p has at least $k + 1$ coordinates with values in $\{1/2 + d, 1/2 - d\}$; without loss of generality, we may assume that the first $k + 1 \geq \ell + 1$ coordinates of p equal $1/2 + d$. Recall from Lemma 1.8 that d is rational and that the set of points $x \in \Delta$ with $\|x - (1/2, \dots, 1/2)\|_\infty = d$ is a finite union of rational polytopes. If we impose the additional constraint that the first $\ell + 1$ coordinates all equal $1/2 + d$, then we again obtain a finite union of rational polytopes, and we know that it is nonempty since it contains p . Hence it also contains some point q with all rational coordinates. Let Δ' be the subgroup of Δ obtained by intersecting Δ with the subspace

$$\{(x_1, \dots, x_n) : x_1 = \dots = x_{\ell+1}\}.$$

Notice that Δ' is a subgroup of dimension at least $k - \ell$, and it is proper because it contains the point q . Now, take any $(k - \ell)$ -dimensional subgroup of Δ' containing q , and let $\Delta'' \subseteq (\mathbb{R}/\mathbb{Z})^{n-\ell}$ denote its the projection onto the last $n - \ell$ coordinates of $(\mathbb{R}/\mathbb{Z})^n$. Then $\Delta'' \subseteq (\mathbb{R}/\mathbb{Z})^{n-\ell}$ is a $(k - \ell)$ -dimensional proper subgroup with $D(\Delta'') = D(\Delta') = D(\Delta) = d$, so $d \in \mathcal{S}_{k-\ell}^*(n - \ell)$, as desired. $\square$

Let us now describe what happens when one works with subtori instead of subgroups. We still have the trivial inclusion $\mathcal{S}_{k-\ell}(n - \ell) \subseteq \mathcal{S}_k(n)$ for $1 \leq \ell < k \leq n$, although one cannot hope for any relationship when $\ell = k$ since $\mathcal{S}_0(n - \ell) = \emptyset$, as there are no 0-dimensional proper subtori.

Matters around the reverse inclusion are more complicated. For $\ell < k$, one can run the first paragraph of the argument from the proof of Lemma 1.8 on a k -dimensional proper subtorus U . The trouble comes in the second paragraph, where the intersection of U with the subspace $\{(x_1, \dots, x_n) : x_1 = \dots = x_{\ell+1}\}$ may be a *disconnected* subgroup rather than a connected subtorus. Since $\ell < k$, one could also consider the intersections of U with other subspaces of the form

$$\{(x_1, \dots, x_n) : x_{i_1} = \dots = x_{i_{\ell+1}}\}$$

for $1 \leq i_1 < \dots < i_{\ell+1} \leq k + 1$. Unfortunately, even this additional flexibility is insufficient to guarantee that

the intersection is connected; a concrete counterexample with $\ell = 1, k = 2, n = 3$ is given by the subtorus $\pi(\langle (0, 7, -5), (5, -93, 0) \rangle_{\mathbb{R}}) \subseteq (\mathbb{R}/\mathbb{Z})^3$.

It is still possible to say a little bit. A special case of the following lemma appeared, in somewhat different language, as Lemma 8 of the paper of Bohman, Holzman, and Kleitman [17]. One should think of the lemma as as a characterization of $\max \mathcal{S}_k(n)$ by the “codimension” $n - k$.

Lemma 1.9. *For natural numbers $1 \leq \ell < k \leq n$, we have $\max \mathcal{S}_k(n) = \max \mathcal{S}_{k-\ell}(n - \ell)$.*

Proof. By induction, it suffices to establish the $\ell = 1$ case of the lemma. The trivial inclusion $\mathcal{S}_{k-1}(n-1) \subseteq \mathcal{S}_k(n)$ gives $\max \mathcal{S}_k(n) \geq \max \mathcal{S}_{k-1}(n-1)$. It remains to show the reverse inequality.

Let $d \in \mathcal{S}_k(n)$, and let $U \subseteq (\mathbb{R}/\mathbb{Z})^n$ be a k -dimensional proper subtorus with $D(U) = d$. Write $U = \pi(\langle w_1, \dots, w_{k-2}, u, v \rangle_{\mathbb{R}})$ for nonzero vectors $w_1, \dots, w_{k-2}, u = (u_1, \dots, u_n), v = (v_1, \dots, v_n) \in \mathbb{Z}^n$. The properness of U guarantees that for each $1 \leq i \leq n$, at least one of the k generators of U has its i -th coordinate nonzero. Thus, by replacing u with u plus suitable integer multiples of the other generators, we may assume that u has all nonzero coordinates. Moreover, by changing coordinates $x_i \mapsto -x_i$ for the i 's with $u_i < 0$, we may assume that u has all coordinates strictly positive. Finally, by permuting the coordinates x_i , we may assume that the quantities

$$v_1/u_1, v_2/u_2, \dots, v_n/u_n$$

are non-decreasing. Since u, v are non-parallel, there is some index $1 \leq i \leq n-1$ such that $v_i/u_i < v_{i+1}/u_{i+1}$.

Now, let T denote the identity component of the intersection $U \cap \{x_i = -x_{i+1}\}$. Note that T is a subtorus of dimension either $k-1$ or k . We claim that T is proper. To this end, the key observation is that T contains the 1-dimensional proper subtorus

$$\pi(\langle (v_i + v_{i+1})u - (u_i + u_{i+1})v \rangle_{\mathbb{R}}).$$

Indeed, the i -th and $(i+1)$ -th coordinates of $(v_i + v_{i+1})u - (u_i + u_{i+1})v$ sum to zero by direct calculation, so this is indeed a subtorus of T . For properness, it suffices to show that $(v_i + v_{i+1})u - (u_i + u_{i+1})v$ has all coordinates nonzero. To see this, notice that the j -th coordinate is equal to

$$(v_i + v_{i+1})u_j - (u_i + u_{i+1})v_j,$$

which vanishes if and only if

$$\frac{v_j}{u_j} = \frac{v_i + v_{i+1}}{u_i + u_{i+1}}.$$

Since $u_i, u_{i+1} > 0$ and $v_i/u_i < v_{i+1}/u_{i+1}$, the quantity $(v_i + v_{i+1})/(u_i + u_{i+1})$ lies strictly between v_i/u_i and v_{i+1}/u_{i+1} , and in particular it is not equal to v_j/u_j for any $1 \leq j \leq n$.

Take T' to be a $(k-1)$ -dimensional proper subtorus of T , and let T'' denote the projection of T' onto all but the i -th coordinate. Then T'' is a $(k-1)$ -dimensional proper subtorus of $(\mathbb{R}/\mathbb{Z})^{n-1}$, and $D(U) \leq D(T) \leq D(T') = D(T'') \leq \max \mathcal{S}_{k-1}(n-1)$, as desired. $\square$

We remark that we could have intersected U with the subspace where $x_1 = x_n$ instead of the subspace where $x_i = -x_{i+1}$; there is some (limited) flexibility in this choice.

It is instructive to compare what the proofs of Lemmas 1.8 and 1.9 tell us about the intersections $U' := U \cap \{x_i = x_j\}$ where $U \subseteq (\mathbb{R}/\mathbb{Z})^n$ is a fixed 2-dimensional proper subtorus and $1 \leq i < j \leq n$ range. The proof of Lemma 1.8 provides indices $i < j$ such that U' satisfies $D(U') = D(U)$ but may be disconnected. The proof of Lemma 1.9 provides indices $i < j$ such that U is connected but may satisfy $D(U') > D(U)$. The example described between the two lemmas shows that for some choices of U it is impossible to find indices $i < j$ such that U' both is connected and satisfies $D(U') = D(U)$.

1.4 Quantitative Kronecker Theorem

We will require some technical results on the distribution of subtori of $(\mathbb{R}/\mathbb{Z})^n$. In this setting, every k -dimensional subtorus ($1 \leq k \leq n$) is of the form $T = \pi(W)$ for some k -dimensional subspace W of $\mathbb{R}^n$ that is “rational” in the sense that $\Lambda := \mathbb{Z}^n \cap W$ is a lattice of rank k . We write $\text{vol}_k(T)$ for the k -dimensional volume of a k -dimensional subtorus T . Note that the volume of the torus $T = \pi(W)$ is simply the covolume of its associated lattice Λ .

We will require a simple lemma showing that there are only finitely many subtori with volume smaller than any constant. Schmidt obtained this result with a sharp bound in [100], but we will also include a short proof here (with a worse bound) to keep our arguments self-contained. We first recall a standard fact from reduction theory.

Theorem 1.10 ([104], Sections X.5–6). *Let $\Lambda \subset \mathbb{R}^k$ be a lattice of full rank. Then there is a basis $b_1, \dots, b_k$ of Λ such that*

$$\|b_1\|_2 \cdots \|b_k\|_2 \leq 2^k (3/2)^{k(k-1)/2} \text{covol}_k(\Lambda) / \omega_k,$$

where $\omega_k := \pi^{k/2} / \Gamma(k/2 + 1)$ is the volume of the L^2 -unit ball in $\mathbb{R}^k$.

This theorem is a variation of Minkowski’s Second Theorem from the geometry of numbers. We cannot apply Minkowski’s Second Theorem directly because the elements of Λ achieving the successive minima need not form a $\mathbb{Z}$ -basis (as discussed in [104], Section X.5).

Lemma 1.11. *Let $0 \leq k \leq n$ be nonnegative integers, and let $V > 0$ be a positive real number. Then there are only finitely many k -dimensional subtori of $(\mathbb{R}/\mathbb{Z})^n$ with volume at most V .*

Proof. Every k -dimensional subtorus T of $(\mathbb{R}/\mathbb{Z})^n$ with volume at most V is of the form $\pi(W)$ for some k -dimensional subspace W of $\mathbb{R}^n$ that is “rational” in the sense that $\Lambda := \mathbb{Z}^n \cap W$ is a rank- k lattice in W ; identifying W with $\mathbb{R}^k$, we see that T is isomorphic to $\mathbb{R}^k/\Lambda$. Theorem 1.10 provides a basis $b_1, \dots, b_k$ of Λ such that

$$\|b_1\|_2 \cdots \|b_k\|_2 \leq 2^k (3/2)^{k(k-1)/2} \text{vol}_k(T) / \omega_k.$$

Since every nonzero element of $\mathbb{Z}^n$ has length at least 1, we see that each

$$\|b_i\|_2 \leq 2^k (3/2)^{k(k-1)/2} V / \omega_k := \ell(k, V).$$

Since $\mathbb{Z}^n$ has only finitely many elements of length at most $\ell(k, V)$, we conclude that there are only finitely many choices for the b_i ’s and hence (since T is determined by W , which is determined by Λ) only finitely many choices for T . $\square$

We also need the analogous result for subgroups.

Lemma 1.12. *Let $0 \leq k \leq n$ be nonnegative integers, and let $V > 0$ be a positive real number. Then there are only finitely many k -dimensional subgroups of $(\mathbb{R}/\mathbb{Z})^n$ of volume at most V .*

Proof. Every subgroup Δ of $(\mathbb{R}/\mathbb{Z})^n$ with volume at most V can be written as a direct sum $\Delta = T \oplus H$, where T is the identity component of Δ and H is a finite subgroup of $(\mathbb{R}/\mathbb{Z})^n$; note that T is a k -dimensional subtorus and $H \subseteq (\mathbb{Q}/\mathbb{Z})^n$. From $\text{vol}_k(T) \geq 1$ and

$$\text{vol}_k(T) \cdot |H| = \text{vol}_k(\Delta) \leq V,$$

we see that $\text{vol}_k(T), |H| \leq V$. Lemma 1.12 tells us that there are only finitely many choices for T . To see that there are only finitely many choice for H , note that H is contained in the set of elements of $(\mathbb{R}/\mathbb{Z})^n$ of order at most V and this set is finite. Since Δ is determined by T, H , we conclude that there are only finitely many choices for Δ . $\square$

We can now prove our Kronecker-type result. The classical version of Kronecker’s Equidistribution Theorem says that every irrational orbit in $(\mathbb{R}/\mathbb{Z})^n$ equidistributes in some subtorus of dimension at least 2. Our version can be understood as a quantitative, finitary analogue for closed subtori and subgroups.

Lemma 1.13. *Let $0 \leq k \leq n$ be nonnegative integers, and let $\varepsilon > 0$ be a positive real number. Then there exist a constant $C^* = C^*(n, k, \varepsilon) > 0$ and a finite list $L^* = L^*(n, k, \varepsilon)$ of subgroups of $(\mathbb{R}/\mathbb{Z})^n$ such that for each k -dimensional subgroup Δ of $(\mathbb{R}/\mathbb{Z})^n$, one of the following holds:*

1. Δ has volume at most C .
2. Δ is contained in one of the subgroups $\Gamma \in L^*$ and is ε -dense (with respect to the $L^2((\mathbb{R}/\mathbb{Z})^n)$ -norm) in Γ . If Δ is a subtorus, then Γ can also be taken to be subtorus.

Moreover, each subgroup in L^* has dimension at least $k + 1$ if $k < n$.

Notice that $(\mathbb{R}/\mathbb{Z})^n$ will always be one of the elements of our list L^* ; this corresponds to the situation where Δ is dense in $(\mathbb{R}/\mathbb{Z})^n$.

Proof. We fix n and proceed by downward induction on k . For the base case $k = n$, the only subgroup Δ is $\Delta = (\mathbb{R}/\mathbb{Z})^n$; take $L^*(n, n, \varepsilon) = \{(\mathbb{R}/\mathbb{Z})^n\}$ and $C^*(n, n, \varepsilon) = 1$ (say), and there is nothing to prove.

We proceed to the induction step. For visualizing the argument, the reader may find it helpful to keep the $k = n - 1$ case in mind. Let Δ be a k -dimensional subgroup of volume V , and let T be the identity component of Δ . Note that T is itself a k -dimensional subtorus, and write $T = \pi(W)$ with W a k -dimensional rational subspace of $\mathbb{R}^n$. Moreover, we can write $\Delta = T \oplus H$ for some finite group $H \subseteq (\mathbb{R}/\mathbb{Z})^n$, and there is a unique finite set $R \subseteq [0, 1]^n$ such that $\pi(R) = H$ bijectively; note that in fact $R \subseteq \mathbb{Q}^n$. Let X denote the tubular neighborhood of Δ of radius $\varepsilon/2$. If

$$V\omega_{n-k}(\varepsilon/2)^{n-k} > 1,$$

then the $(n - k)$ -dimensional “orthogonal slices” of X around the points of Δ cannot all be disjoint. Lifting this picture to $\mathbb{R}^n$ gives us some $x \in (\mathbb{Z}^n + R) \setminus W$ such that the $\varepsilon/2$ -tubular neighborhoods around W and $W + x$ intersect. In particular, the orthogonal complement of W intersects $W + x$ at some (necessarily nonzero) point p of length at most ε . The point p has all rational coordinates since it is determined by a rational system of linear equations (involving W and x). Let W' be the $\mathbb{R}$ -subspace of $\mathbb{R}^n$ spanned by W and p , and let $T' := \pi(W')$. Notice that T' is a $(k + 1)$ -dimensional subtorus (because of the rationality of p) and that the $\varepsilon/2$ -neighborhood of Δ contains T' . It follows that Δ is $\varepsilon/2$ -dense in the subgroup $\Delta' := T' + H$.

We now have a dichotomy depending on whether the volume of Δ' is small or large. Set

$$L^*(n, k, \varepsilon) := L^*(n, k + 1, \varepsilon/2) \cup \{(k + 1)\text{-dim'l subgroups of volume at most } C^*(n, k + 1, \varepsilon/2)\}.$$

If the volume of Δ' is at most $C^*(n, k + 1, \varepsilon/2)$, then $\Delta' \in L^*(n, k, \varepsilon)$ and we are done since Δ is ε -dense in Δ' . If instead the volume of Δ' is larger than $C^*(n, k + 1, \varepsilon/2)$, then there is some $\Delta'' \in L^*(n, k + 1, \varepsilon/2) \subseteq$

$L^*(n, k, \varepsilon)$ such that Δ' is $\varepsilon/2$ -dense in Δ'' , and again we are done since Δ is ε -dense in Δ'' by the Triangle Inequality. This proves the lemma with the choice

$$C^*(n, k, \varepsilon) := \frac{1}{\omega_{n-k}(\varepsilon/2)^{n-k}}.$$

Lemma 1.12 implies that $L^*(n, k, \varepsilon)$ is a finite set. The second part of (2) is clear from tracing through the proof with only subtori under consideration. $\square$

1.5 Producing accumulation points

Our next task is showing how elements of certain spectra lead to accumulation points in other spectra. We begin by producing density points in multiset spectra.

Proposition 1.14. *For natural numbers $1 \leq k < n$, we have $\mathcal{S}_{k+1}(n) \subseteq \text{den}(\mathcal{S}_{k, \text{mult}}(n))$. For nonnegative integers $0 \leq k < n$, we have $\mathcal{S}_{k+1}^*(n) \subseteq \text{den}(\mathcal{S}_{k, \text{mult}}^*(n))$.*

Proof. For the first statement, we must show that for every $(k+1)$ -dimensional proper subtorus $U \subseteq (\mathbb{R}/\mathbb{Z})^n$, the value $D(U)$ is a density point of the multiset $\mathcal{S}_{k, \text{mult}}(n)$. Fix such a U . Notice that every k -dimensional subtorus $T \subseteq U$ has $D(T) \geq D(U)$ by the definition of D . If moreover T is ε -dense in U with respect to the L^∞ -norm, then the Triangle Inequality also gives $D(T) \leq D(U) + \varepsilon$. Hence, it suffices to exhibit an infinite sequence of k -dimensional proper subtori $T_1, T_2, \dots$ that are contained in U and become $o(1)$ -dense in U . Notice that the properness condition comes for free if $T_1, T_2, \dots$ become $o(1)$ -dense in U , since non-proper subtori all have D -value $1/2$.

Now, in fact, any infinite sequence of k -dimensional subtori of U becomes $o(1)$ -dense in U (by the codimension-1 cases of Lemmas 1.11 and 1.13). We can also produce such a sequence explicitly by writing $U = \pi(\langle u_1, \dots, u_{k+1} \rangle_{\mathbb{R}})$ for some nonzero vectors $u_1, \dots, u_{k+1} \in \mathbb{Z}^n$ and setting

$$T_j := \pi(\langle u_1, u_2, \dots, u_{k-1}, u_k + ju_{k+1} \rangle_{\mathbb{R}}) \subseteq U$$

for each $j \in \mathbb{N}$. The $o(1)$ -denseness of the T_j 's in U follows from the $o(1)$ -denseness of the subtori $\pi(\langle (1, j) \rangle_{\mathbb{R}})$ in $(\mathbb{R}/\mathbb{Z})^2$, which is geometrically obvious.

The second statement goes in the same way. Let $\Delta \subseteq (\mathbb{R}/\mathbb{Z})^n$ be a $(k+1)$ -dimensional proper subgroup. As in the previous paragraph, it suffices to find an infinite sequence $\Gamma_1, \Gamma_2, \dots$ of k -dimensional subgroups of Δ that become $o(1)$ -dense in Γ . We can write Δ as a direct sum $\Delta = U \oplus G$, where U is a $(k+1)$ -dimensional subtorus (not necessarily proper) and G is a finite subgroup (also not necessarily proper). It

suffices to produce a sequence of k -dimensional subgroups $H_1, H_2, \dots$ of U that become $o(1)$ -dense in U , since then $\Gamma_1 := H_1 \oplus G, \Gamma_2 := H_2 \oplus G, \dots$ will be $o(1)$ -dense in Δ . If $k = 0$, then U is a 1-dimensional subtorus and we can take H_j to be the unique subgroup of U of order j . If $k \geq 1$, then we can take H_j to be the subtorus T_j constructed in the previous paragraph. $\square$

The situation becomes more subtle when we work with set spectra instead of multiset spectra. The proof strategy for Proposition 1.14 completely breaks down. In the context of the first statement (about subtori), it is possible that $D(T) = D(U)$ for all but finitely many k -dimensional proper subtori T of U ; in this case, the D -values of the k -dimensional proper subtori of U witness $D(U)$ as a density point of $\mathcal{S}_{k,\text{mult}}(n)$ but do not witness $D(U)$ as a genuine accumulation point of $\mathcal{S}_k(n)$. In the following chapter, we will see a simple geometric criterion for when a 2-dimensional proper subtorus $U \subseteq (\mathbb{R}/\mathbb{Z})^n$ has infinitely many 1-dimensional proper subtori T with $D(T) > D(U)$: This occurs if and only if the locus where U attains its D -value is contained in a finite union of parallel line segments.² Using this criterion, they showed that, for example, the 2-dimensional proper subtorus

$$U = \pi(\langle (1, 2, 3, 2, 0, 0, 0), (0, 0, 0, 2, 1, 2, 3) \rangle_{\mathbb{R}}) \subseteq (\mathbb{R}/\mathbb{Z})^7$$

has only finitely many 1-dimensional proper subtori T with $D(T) > D(U)$. It is even easier to construct such examples in the context of the second statement of Proposition 1.14 (about subgroups). For instance, the 1-dimensional proper subgroup

$$\Delta = \mathbb{R}/\mathbb{Z} \times \{0, 1/3, 2/3\} \subseteq (\mathbb{R}/\mathbb{Z})^2$$

has only finitely many discrete subgroups Γ with $D(\Gamma) > D(\Delta)$.

An important special case occurs when U is of the form $U = U' \times \mathbb{R}/\mathbb{Z}$ for U' a 1-dimensional proper subtorus. Since U' achieves its D -value at only finitely many points, we see that U achieves its D -value on only finitely many “vertical” line segments. The above criterion guarantees that U contains infinitely many 1-dimensional proper subtori T with $D(T) > D(U) = D(U')$. In particular, there exist such T ’s with volume tending to infinity, in which case $D(T)$ approaches $D(U)$ from above due to the codimension-1 cases of Lemmas 1.11 and 1.13. We conclude that, as observed more directly in [65], we have the inclusion

$$\mathcal{S}(n-1) \subseteq \text{acc}(\mathcal{S}(n))$$

²The same argument shows that a $(k+1)$ -dimensional proper subtorus U has infinitely many k -dimensional proper subtori T with $D(T) > D(U)$ if and only if the locus where U attains its D -value is contained in a finite union of parallel k -dimensional disks.

for every natural number $n \geq 2$. It is an interesting open problem to determine whether or not

$$\mathcal{S}_k(n-1) \subseteq \text{acc}(\mathcal{S}_k(n))$$

for all $1 \leq k < n$; we suspect that this is likely the case, and that in fact equality holds.

1.6 Putting everything together

We have nearly all of the pieces that comprise Theorem 1.4. Before we can complete the proof, we need to use our Kronecker-type result to relate $\text{acc}(\mathcal{S}_k(n))$ to $\mathcal{S}_{k+1}(n)$.

Lemma 1.15. *Let $1 \leq k < n$ be natural numbers. Then the accumulation points of $\mathcal{S}_k(n)$ are all accumulation points only from above, and $\text{den}(\mathcal{S}_{k,\text{mult}}(n)) \subseteq \mathcal{S}_{k+1}(n)$.*

Proof. We first claim that if $1 \leq \ell \leq n-1$ and $T_1, T_2, \dots$ is a sequence of distinct ℓ -dimensional proper subtori of $(\mathbb{R}/\mathbb{Z})^n$, then there exist a subtorus T of dimension at least $\ell+1$ and a subsequence $T_{i_1}, T_{i_2}, \dots$ such that each T_{i_j} is contained in T and the T_{i_j} 's become $o(1)$ -dense in T as $j \rightarrow \infty$. We proceed by downward induction on ℓ . Notice that the volume of the T_i 's tends to infinity as i grows. The base case $\ell = n-1$ now follows immediately from Lemma 1.13, which tells us that the T_i 's are $o(1)$ -dense in $(\mathbb{R}/\mathbb{Z})^n$. For the induction step, suppose that the T_i 's are ℓ -dimensional with $\ell < n-1$. Lemma 1.13 tells us that there are subtori T'_i , each of dimension at least $\ell+1$, such that each T_i is contained in T'_i and the T_i 's are $o(1)$ -dense in their respective T'_i 's as $i \rightarrow \infty$. By passing to a subsequence, we may assume that all of the T'_i 's have the same dimension, say, $\ell' \geq \ell+1$. If there are only finitely many distinct T'_i 's, then one of them appears infinitely often and we can conclude by letting T be such a T'_i . If instead there are infinitely many distinct T'_i 's, then (by the induction hypothesis) there exist a subtorus T of dimension at least $\ell'+1$ and a subsequence $T'_{i_1}, T'_{i_2}, \dots$ such that each T'_{i_j} is contained in T and the T'_{i_j} 's are $o(1)$ -dense in T . Then we also see that each T_{i_j} is contained in T and the T_{i_j} 's are $o(1)$ -dense in T (by the Triangle Inequality), as desired.

We now prove the lemma by (upward) induction on k . Let $T_1, T_2, \dots$ be a sequence of k -dimensional proper subtori with $\lim_{i \rightarrow \infty} D(T_i) = d$. The claim from the previous paragraph tells us that, after passing to a subsequence, we may assume that all of the T_i 's are contained in a single subtorus T of dimension at least $k+1$ and that the T_i 's are $o(1)$ -dense in T . It follows that every $D(T_i) \geq D(T)$ and that $\lim_{i \rightarrow \infty} D(T_i) = D(T)$, so $D(T) = d$. In particular, $D(T_i)$ cannot limit to d from below. This establishes the first statement of the lemma, namely, that $\mathcal{S}_k(n)$ has only upper accumulation points.

We know that T has dimension at least $k+1$. We can find a $(k+1)$ -dimensional subtorus T' such that $d = D(T) = D(T') \in \mathcal{S}_{k+1}(n)$: Just take any $(k+1)$ -dimensional subtorus T' of T that passes through a

rational point of T with L^∞ -distance d from the point $(1/2, \dots, 1/2)$ (using Lemma 1.7). This establishes the second statement of the lemma, namely, that $\text{acc}(\mathcal{S}_k(n)) \subseteq \mathcal{S}_{k+1}(n)$. $\square$

The following subgroup-version of this lemma can be proven by an identical argument.

Lemma 1.16. *Let $0 \leq k < n$ be nonnegative integers. Then the accumulation points of $\mathcal{S}_k^*(n)$ are all accumulation points only from above, and $\text{den}(\mathcal{S}_{k,\text{mult}}^*(n)) \subseteq \mathcal{S}_{k+1}^*(n)$.*

We remark that Lemmas 1.15 and 1.16 do not use anything special about the fact that $D(T)$ measures the L^∞ distance from T to $(1/2, \dots, 1/2)$: We would get the same “qualitative” result for any $D_f(T)$.

Combining all of our results so far yields Theorem 1.4, as follows.

- Lemmas 1.15 and 1.16 tell us that $\mathcal{S}_k(n), \mathcal{S}_{k'}^*(n)$ have only upper accumulation points.
- For the long string of set inclusions in Theorem 1.4: We obtain $\text{acc}(\mathcal{S}_k(n)) \subseteq \mathcal{S}_{k+1}(n)$ from Lemma 1.15. The inclusion $\mathcal{S}_{k+1}(n) \subseteq \mathcal{S}_{k+1}^*(n)$ is trivial. And we obtain $\mathcal{S}_{k+1}^*(n) = \mathcal{S}_0^*(n - k - 1)$ from Lemma 1.8.
- For the multiset equalities: We obtain $\mathcal{S}_{k+1}(n) \subseteq \text{den}(\mathcal{S}_{k,\text{mult}}(n))$ and $\mathcal{S}_{k'+1}^*(n) \subseteq \text{den}(\mathcal{S}_{k',\text{mult}}^*(n))$ from Proposition 1.14. And we obtain $\text{den}(\mathcal{S}_{k,\text{mult}}(n)) \subseteq \mathcal{S}_{k+1}(n)$ and $\text{den}(\mathcal{S}_{k',\text{mult}}^*(n)) \subseteq \mathcal{S}_{k'+1}^*(n)$ from Lemmas 1.15 and 1.16. The final equality $\mathcal{S}_{k'+1}^*(n) = \mathcal{S}_{k'}^*(n - 1)$ follows from Lemma 1.8.
- For the final assertion, we obtain $\mathcal{S}(n - 1) \subseteq \text{acc}(\mathcal{S}(n))$ from Theorem 1.3, and Lemma 1.8 gives us $\mathcal{S}_2^*(n) = \mathcal{S}^*(n - 1)$.

1.7 Checking small speeds suffices

In this short section, we explain how to deduce the computability result described in Section 1.1.6. The argument is a straightforward modification of the approach taken in the proof of Lemma 1.13. Suppose we already know that the Lonely Runner Conjecture holds for $n - 1$ runners and we want to determine whether or not it holds for n runners, simply by checking the maximum loneliness for all n -tuples of small natural-number speeds. It would suffice to show that if there exists a counterexample to the Lonely Runner Conjecture with n runners, then there exists such a counterexample with all slow runners. We will show a stronger statement, namely, that every n -tuple with large sum of squares of speeds must satisfy the Lonely Runner Conjecture. Recall that the tuple of speeds $(v_1, \dots, v_n)$ corresponds to the 1-dimensional proper subtorus $T := \pi(\langle (v_1, \dots, v_n) \rangle_{\mathbb{R}}) \subseteq (\mathbb{R}/\mathbb{Z})^n$, which has length $\text{vol}_1(T) = \sqrt{v_1^2 + \dots + v_n^2}$ if $\gcd(v_1, \dots, v_n) = 1$ (which of course we may assume).

Let $\varepsilon := 1/n - 1/(n+1) = 1/(n(n+1))$, and let T be a 1-dimensional proper subtorus of $(\mathbb{R}/\mathbb{Z})^n$ with volume V . The argument from the proof of Lemma 1.13 shows that if

$$V\omega_{n-1}\varepsilon^{n-1} > 1,$$

then there is some 2-dimensional proper subtorus T' such that T is ε -dense in T' . In particular, $D(T) \leq D(T') + \varepsilon$, and Lemma 1.9 (together with our “induction” hypothesis) ensures that the latter quantity is at most $1/2 - 1/(n+1)$. Hence T satisfies the Lonely Runner Conjecture once

$$V > \frac{1}{\omega_{n-1}\varepsilon^{n-1}} = \frac{\Gamma((n+1)/2)(n(n+1))^{n-1}}{\pi^{(n-1)/2}},$$

and Stirling’s Approximation shows that this threshold is smaller than $n^{(5/2)n}$.

We remark that we can run the same argument if the Lonely Runner Conjecture for $n-1$ runners fails but we still have $\max \mathcal{S}(n-1) = \max \mathcal{S}_2(n) \leq 1/2 - 1/(n+1) - \delta$ for some $\delta > 0$; the threshold for V will just be corresponding larger.

The ideas in this section also bear on the problem of characterizing *tight instances* of the Lonely Runner Problem, i.e., 1-dimensional proper subtori $T \subseteq (\mathbb{R}/\mathbb{Z})^n$ with the (conjecturally) largest possible value $D(T) = 1/2 - 1/(n+1)$. The above proof shows that if the Lonely Runner Conjecture holds for $n-1$ runners, then there are only finitely many tight instances for n runners, and one could easily extract an explicit upper bound on their volume (and, by extension, their number).

1.8 Calculations for small (co)dimension

When the number of runners is small, we can provide some more precise information about the relation between Lonely Runner spectra and subgroup Lonely Runner spectra; this leads to more general observations about (subgroup) Lonely Runner spectra of small codimension.

We start with $\mathcal{S}_1(1) = \{0\}$. It is also easy to compute (see, e.g., [65, Theorem 4.1] or the remark on page 5 of [12]) that

$$\mathcal{S}_1(2) = \{0\} \cup \{1/(4s+2) : s \in \mathbb{N}\}.$$

In particular,

$$\text{acc}(\mathcal{S}_1(2)) = \{0\} = \mathcal{S}_1(1),$$

which provides an affirmative answer to the $k = n$ (“codimension-0”) case of Problem 1.6.

Next, we calculate that

$$\mathcal{S}_0^*(1) = \{0\} \cup \{1/(4s+2) : s \in \mathbb{N}\} = \mathcal{S}_1(2).$$

Indeed, each proper discrete subgroup of $\mathbb{R}/\mathbb{Z}$ is equal to $\pi(\langle 1/q \rangle_{\mathbb{Z}})$ for some natural number $q \geq 2$. If q is even, then $1/2 \in \pi(\langle 1/q \rangle_{\mathbb{Z}})$ and $D(\pi(\langle 1/q \rangle_{\mathbb{Z}})) = 0$. If instead $q = 2s+1$ is odd, then the closest points in $\pi(\langle 1/q \rangle_{\mathbb{Z}})$ to $1/2$ are $\pm s/q$, so

$$D(\pi(\langle 1/q \rangle_{\mathbb{Z}})) = 1/2 - s/q = 1/(4s+2).$$

Notice that the degenerate Lonely Runner spectrum $\mathcal{S}_0(1) = \emptyset$ is emphatically not equal to $\mathcal{S}_0^*(1)$.

For each $n \geq 2$ we have the chain of inclusions

$$\mathcal{S}_1(2) \subseteq \mathcal{S}_{n-1}(n) \subseteq \mathcal{S}_{n-1}^*(n) = \mathcal{S}_0^*(1) = \mathcal{S}_1(2),$$

so equality must hold. This proves the $k = n-1$ (“codimension-1”) case of Problem 1.6 and implies that $\text{acc}(\mathcal{S}_1(3)) = \mathcal{S}_1(2)$, as expected.

The subgroup Lonely Runner spectrum $\mathcal{S}_0^*(2)$ is more complicated since there are many discrete subgroups of $(\mathbb{R}/\mathbb{Z})^2$. Of course, we have the inclusion

$$\mathcal{S}_1(3) \subseteq \mathcal{S}_1^*(3) = \mathcal{S}_0^*(2).$$

This inclusion, however, is strict.

Proposition 1.17. *We have $7/50 \in \mathcal{S}_0^*(2) \setminus \mathcal{S}_1(3)$.*

Proof. A direct calculation shows that $D(\pi(\langle (12/25, 9/25) \rangle_{\mathbb{Z}})) = 7/50 \in \mathcal{S}_0^*(2)$; see Figure 1.2. It remains to show that $7/50 \notin \mathcal{S}_1(3)$. Let $T \subseteq (\mathbb{R}/\mathbb{Z})^3$ be a 1-dimensional proper subtorus. We will show that $D(T) \neq 7/50$. First, suppose that $\text{vol}(T) > 199$; note that $199 \cdot ((1/25 - \varepsilon)^2 \pi) > 1$ for some $\varepsilon > 0$. Then the tubular neighborhood argument from the proof of Lemma 1.13, with a tube of radius $1/25 - \varepsilon$, shows that T is $(1/25 - \varepsilon)$ -dense in some proper subtorus $U \subseteq (\mathbb{R}/\mathbb{Z})^3$ of dimension at least 2. Thus

$$D(U) \leq D(T) \leq D(U) + 1/25 - \varepsilon.$$

We know from the characterization of $\mathcal{S}_2(3)$ that either $D(U) = 1/6$ or $D(U) \leq 1/10$. In the former case,

the inequality $D(T) \geq D(U) = 1/6$ implies that $D(T) \neq 7/50$. In the latter case, the inequality

$$D(T) \leq D(U) + 1/25 - \varepsilon \leq 1/10 + 1/25 - \varepsilon = 7/50 - \varepsilon$$

implies that $D(T) \neq 7/50$. Finally, an exhaustive computer search (taking about an hour on a standard laptop) shows that $D(T) \neq 7/50$ when $\text{vol}(T) \leq 199$. $\square$




Figure 1.2: The black dots show elements of the discrete subgroup $\pi(\langle (12/25, 9/25) \rangle_{\mathbb{Z}}) \subseteq (\mathbb{R}/\mathbb{Z})^2$, which witnesses the value $7/50 \in \mathcal{S}_0^*(2)$, and the blue square shows the L^∞ -ball of radius $7/50$ centered at the point $(1/2, 1/2)$. Notice that there are black dots on all four edges of the blue square; this is a necessary feature of any such example.

We believe that $7/50 \notin \mathcal{S}_2(4)$, but we do not have the computing power to carry out an exhaustive computer search analogous to what we did in Proposition 1.17. In any case, the example in Proposition 1.17 shows that proving $\mathcal{S}_2(4) = \mathcal{S}_1(3)$ will necessarily be more difficult than proving $\mathcal{S}_2(3) = \mathcal{S}_1(2)$. In particular, one would have to show that not all 1-dimensional proper subgroups of $(\mathbb{R}/\mathbb{Z})^3$ can arise when one applies the slicing argument of Lemma 1.8 to 2-dimensional proper subtori in $(\mathbb{R}/\mathbb{Z})^4$.

1.9 Remarks and open problems

We conclude with a few comments and questions.

1. Our proof of Theorem 1.4 shows that each element of $\mathcal{S}_1(n) \setminus \mathcal{S}_2(n)$ appears with finite multiplicity (but not uniformly bounded multiplicity, even for $n = 2$) in $\mathcal{S}_{1,\text{mult}}(n)$. In particular, as we mentioned in Section 1.7, if the Lonely Runner Conjecture is true for n and $n - 1$ runners, then there are only finitely many tight instances of the Lonely Runner Conjecture for n runners, and all such instances are low-volume. It would be feasible (and very useful) to carry out numerical calculations in this direction.
2. For a proper subtorus $U \subseteq (\mathbb{R}/\mathbb{Z})^n$ of dimension at least 2, one can define the *Lonely Runner spectrum relative to U* to be the set $\mathcal{S}(U)$ of all values of $D(T)$ as T ranges over 1-dimensional proper subtori contained in U . Our proof of Theorem 1.4 shows that all accumulation points of $\mathcal{S}(n)$ are “due” to

relatively Lonely Runner spectra, in the sense that for each $d \in \text{acc}(\mathcal{S}(n))$ there are some small $\varepsilon > 0$ and a finite list $U_1, \dots, U_t$ with $D(U_1) = \dots = D(U_t) = d$ such that

$$\mathcal{S}(n) \cap (d, d + \varepsilon) = \left(\bigcup_{j=1}^t \mathcal{S}(U_j) \right) \cap (d, d + \varepsilon).$$

The following chapter contains a study of such relatively Lonely Runner spectra for 2-dimensional subtori U and shows that the sets $\mathcal{S}(U)$ have surprisingly rigid arithmetical properties.

3. Our results show that

$$\text{acc}(\mathcal{S}(3)) = \mathcal{S}(2) = \{1/(4s+2) : s \in \mathbb{N}\} \cup \{0\}.$$

The main result of [65] explicitly determines $\mathcal{S}(3) \cap [1/6, 1/2]$, i.e., the set $\mathcal{S}(3)$ “up to the first accumulation point $1/6$ ” (see Conjecture 1.1). The work of the following chapter yields a description of $\mathcal{S}(3)$ up to the second accumulation point $1/10$ as well as chunks of some other Lonely Runner spectra. It would be interesting to obtain more such characterizations.

4. Is the set $\cup_{n \geq 1} \mathcal{S}(n)$ dense in the interval $[0, 1/2]$? If not, does $\overline{\cup_{n \geq 1} \mathcal{S}(n)}$ have any “nice” self-similarity properties?

Chapter 2

Lonely runners: the quantitative view

2.1 Introduction

2.1.1 The Lonely Runner Conjecture

The Lonely Runner Conjecture of Wills [113] and Cusick [28] originally arose as a possible converse to Dirichlet's theorem on Diophantine approximation. There are now several equivalent formulations of the Lonely Runner Problem. In the most popular version, due to Bienia, Goddyn, Gvozdzjak, Sebő, and Tarsi [12], we have n runners standing at the start line of a unit-length circular track. The runners simultaneously begin running around the track, each at a constant nonzero integer speed. The *maximum loneliness* of a set of speeds is defined to be the largest L such that there is a time when all of the runners are a distance at least L from the start line. The Lonely Runner Conjecture asserts that the maximum loneliness is always at least $1/(n+1)$, regardless of the choice of the n speeds. (Equality holds for the speeds $1, 2, \dots, n$, for instance.) The combined work of many authors over the last several decades has established the Lonely Runner Conjecture for $n \leq 6$ (see [4, 10, 12, 17, 30, 92]), but it remains open for $n \geq 7$.

In this chapter, we will work with the following *view-obstruction* formulation of the Lonely Runner Problem, as introduced by Cusick [28]. For a subtorus $T \subseteq (\mathbb{R}/\mathbb{Z})^n$, let $D(T)$ denote the L^∞ -distance from T to the point $(1/2, \dots, 1/2)$. For $1 \leq k \leq n$, let $\mathcal{S}_k(n)$ denote the set of all of the values of $D(T)$ as T ranges over the k -dimensional proper subtori of $(\mathbb{R}/\mathbb{Z})^n$; here we say that a subtorus is *proper* if it is not contained in the union of the coordinate hyperplanes of $(\mathbb{R}/\mathbb{Z})^n$. The set $\mathcal{S}_1(n)$ is called the *n -th Lonely Runner spectrum*. Note that each $\mathcal{S}_k(n)$ is a subset of $[0, 1/2)$.

It is easy to check that if $v_1, \dots, v_n$ are nonzero integers, then the maximum loneliness of this set of

speeds equals $1/2 - D(\langle(v_1, \dots, v_n)\rangle_{\mathbb{R}})$.¹ The condition on the v_i 's being all nonzero is equivalent to the condition on the torus $\langle(v_1, \dots, v_n)\rangle_{\mathbb{R}}$ being proper. Hence, in the language of view-obstruction, the Lonely Runner Conjecture asserts that

$$\max \mathcal{S}_1(n) = 1/2 - 1/(n+1).$$

The equality case mentioned at the end of the first paragraph gives $D(\langle(1, 2, \dots, n)\rangle_{\mathbb{R}}) = 1/2 - 1/(n+1) \in \mathcal{S}_1(n)$, and the difficult part of the Lonely Runner Conjecture is showing that for every 1-dimensional proper subtorus $T \subseteq (\mathbb{R}/\mathbb{Z})^n$, we have $D(T) \leq 1/2 - 1/(n+1)$.

Until recently, work on the Lonely Runner Conjecture focused narrowly on the maximum value of each Lonely Runner spectrum $\mathcal{S}_1(n)$. In 2020, the author [65] proposed studying the entire set $\mathcal{S}_1(n)$: For instance, one can try to characterize not only its maximum value but all of its near-maximal values, and one can study other properties of $\mathcal{S}_1(n)$, such as its accumulation points. With regard to the first question, the thesis [65] conjectured that

$$\mathcal{S}_1(n) \cap \left(\frac{1}{2} - \frac{1}{n}, \frac{1}{2} \right] = \left\{ \frac{1}{2} - \frac{1}{n} + \frac{1}{n^2s + n} : s \in \mathbb{N} \right\} \quad (2.1)$$

(where for us $\mathbb{N} = \{1, 2, 3, \dots\}$). The D -values on the right-hand side are achieved by the subtori $\langle(1, 2, \dots, n-1, ns)\rangle_{\mathbb{R}}$. The appeal of this conjecture is that it provides a “good” explanation for the appearance of the quantity $1/2 - 1/(n+1)$ in the Lonely Runner Conjecture: This quantity is the endpoint (corresponding to $s = 1$) of a highly arithmetically structured sequence of D -values. One of the main results of [65] was a proof of this conjecture for $n = 2, 3$. Fan and Sun [39] later discovered an infinite family of counterexamples for $n = 4$, namely,

$$D(\langle(8, 4s+3, 4s+11, 4s+19)\rangle_{\mathbb{R}}) = \frac{1}{2} - \frac{1}{4} + \frac{1}{16s+60}. \quad (2.2)$$

These elements of $\mathcal{S}_1(4)$ are different from the values in (2.1) because $60 \not\equiv 4 \pmod{16}$. Nonetheless, this sequence of large elements of $\mathcal{S}_1(4)$ is highly arithmetically structured, just like the sequence in (2.1). The main result of the present chapter will show that this is no coincidence.

2.1.2 Relative Lonely Runner spectra

We now discuss the qualitative structure of Lonely Runner spectra and the role of higher-dimensional subtori. For a set S of real numbers, let $\text{acc}(S)$ denote the set of accumulation points of S . One can distinguish between accumulation points from above and from below, denoted $\text{acc}^+(S), \text{acc}^-(S)$ (respectively). In [65], the author showed that

$$\mathcal{S}_1(n-1) \subseteq \text{acc}^+(\mathcal{S}_1(n)) \quad (2.3)$$

¹In this chapter, all $\mathbb{R}$ -spans of rational vectors in $\mathbb{R}^n$ should be interpreted as subtori in $(\mathbb{R}/\mathbb{Z})^n$; in particular, we will dispense with explicitly writing the quotient $\pi : \mathbb{R}^n \rightarrow (\mathbb{R}/\mathbb{Z})^n$ from the previous chapter.

and then asked whether equality holds and whether $\text{acc}^-(\mathcal{S}_1(n)) = \emptyset$. The previous chapter presented some progress towards understanding these questions. One of the main results was that indeed $\text{acc}^-(\mathcal{S}_1(n)) = \emptyset$. We saw that

$$\text{acc}(\mathcal{S}_1(n)) \subseteq \mathcal{S}_2(n), \quad (2.4)$$

and that equality holds if one interprets $\mathcal{S}_1(n)$ suitably as a multiset. We also saw the conjecture that

$$\mathcal{S}_2(n) = \mathcal{S}_1(n-1),$$

which would imply that equality holds in (2.3). Easy calculations verify this conjecture for $n = 2, 3$, as discussed in the previous chapter

The thrust of the argument for (2.4) is that if $T_1, T_2, \dots$ is a sequence of 1-dimensional proper subtori with $D(T_i) \rightarrow d$, then, once we pass to a subsequence of the T_i 's, there is proper subtorus U of dimension at least 2 such that $D(U) = d$, the T_i 's are all contained in U , and T_i 's become $o(1)$ -dense in U as $i \rightarrow \infty$. Since $D(T_i) \geq D(U)$, we see that $D(T_i)$ approaches d from above, and it follows that $\text{acc}^-(\mathcal{S}_1(n)) = \emptyset$. Only D -values of 2-dimensional subtori appear on the right-hand side of (2.4) because of the trivial containments

$$\mathcal{S}_n(n) \subseteq \mathcal{S}_{n-1}(n) \subseteq \dots \subseteq \mathcal{S}_2(n).$$

We also saw some further results involving “subgroup Lonely Runner spectra” in which k -dimensional subtori are replaced by k -dimensional subgroups.

To make full use of the results of the preceding chapter, it is useful to have the notion of relative Lonely Runner spectra. For a subtorus $U \subseteq (\mathbb{R}/\mathbb{Z})^n$ and a natural number $k \leq \dim(U)$, define $\mathcal{S}_k(U)$, the *order- k Lonely Runner spectrum relative to U* , to be the set of all values of $D(T)$ as T ranges over the k -dimensional proper subtori of U . The special case $\mathcal{S}_1((\mathbb{R}/\mathbb{Z})^n) = \mathcal{S}_1(n)$ corresponds to the Lonely Runner spectra defined above.

The natural output of the aforementioned argument from the last chapter is that for each natural number n and real number $d \in \text{acc}(\mathcal{S}_1(n))$, there are $\varepsilon > 0$ and a finite list of proper subtori $U_1, \dots, U_t \subseteq (\mathbb{R}/\mathbb{Z})^n$, with each U_j satisfying $\dim(U_j) \geq 2$ and $D(U_j) = d$, such that

$$\mathcal{S}_1(n) \cap (d, d + \varepsilon) = \left(\bigcup_{j=1}^t \mathcal{S}_1(U_j) \right) \cap (d, d + \varepsilon). \quad (2.5)$$

In other words, the accumulation of $\mathcal{S}_1(n)$ at the point d is completely described by the (finitely many) relative spectra $\mathcal{S}_1(U_j)$, and so studying accumulation points essentially amounts to studying relative spectra. (The

equation (2.5) also holds trivially for $d \notin \text{acc}(\mathcal{S}_1(n))$ since taking ε sufficiently small makes the left-hand side empty.)

2.1.3 Main result

Our main result is a structural characterization of order-1 relative spectra for 2-dimensional subtori, which can be understood as a “quantitative” complement to the “qualitative” results described in the previous section. We show that these relative spectra are highly arithmetically structured, in a sense that captures the infinite families described in (2.1) and in (2.2). In other words, we show that convergence to accumulation points in Lonely Runner spectra is subject to surprisingly rigid constraints. Let

$$\text{Prog}(\alpha, \beta) := \{\alpha s + \beta : s \in \mathbb{Z}_{\geq 0}\}$$

denote the infinite one-sided arithmetic progression with common difference α and shift β .

Theorem 2.1. *Let $n \geq 2$ be a natural number, and let $U \subseteq (\mathbb{R}/\mathbb{Z})^n$ be a 2-dimensional proper subtorus. Then there are $N \in \mathbb{Z}_{\geq 0}$ and positive rationals $\alpha_1, \dots, \alpha_N, \beta_1, \dots, \beta_N$ such that $\mathcal{S}_1(U)$ has finite symmetric difference with the set*

$$\bigcup_{i=1}^N (D(U) + 1/\text{Prog}(\alpha_i, \beta_i)).$$

We make three remarks about the statement of this theorem. First, the finite symmetric difference arises because a few 1-dimensional subtori of U can be too “small” to see the structure leading to the sets $D(U) + 1/\text{Prog}(\alpha_i, \beta_i)$; below we give an explicit example where this finite symmetric difference is necessary. Second, our argument actually shows the ostensibly slightly stronger statement that there are only finitely many 1-dimensional proper subtori T of U with $D(T)$ lying outside of the set $\{D(U)\} \cup [\cup_{i=1}^N (D(U) + 1/\text{Prog}(\alpha_i, \beta_i))]$. Third, the set $D(U) + 1/\text{Prog}(\alpha_i, \beta_i)$ can be understood as a simple transformation of the set $\mathcal{S}_1(2) = 1/\text{Prog}(4, 6)$.

Notice that all of the subtori U_j in (2.5) are necessarily 2-dimensional for an accumulation point $d \notin \mathcal{S}_3(n)$; this is the case, for instance, for $d = \max \text{acc}(\mathcal{S}_1(n)) = \max \mathcal{S}_1(n-1)$. Here, Theorem 2.1 immediately gives the following pleasant characterization.

Corollary 2.2. *Let $n \geq 2$ be a natural number, and let $d \in \text{acc}(\mathcal{S}_1(n)) \setminus \mathcal{S}_3(n)$ be a real number. Then there are $\varepsilon > 0$, $N \in \mathbb{Z}_{\geq 0}$, and positive rationals $\alpha_1, \dots, \alpha_N, \beta_1, \dots, \beta_N$ such that $\mathcal{S}_1(n) \cap (d, d + \varepsilon)$ has finite*

symmetric difference with the set

$$\left(\bigcup_{i=1}^N (d + 1/\text{Prog}(\alpha_i, \beta_i)) \right) \cap (d, d + \varepsilon).$$

2.1.4 Comparison with previous work

Section 6 of [65] was concerned with the “one very fast runner” setting of the Lonely Runner Problem, in which we fix nonzero integers $v_1, \dots, v_{n-1}$ and consider the maximum loneliness values associated with the speeds $v_1, \dots, v_{n-1}, v_n$ for v_n large. This essentially amounts to studying the relative spectrum

$$\mathcal{S}_1(\langle (v_1, \dots, v_{n-1}) \rangle_{\mathbb{R}} \times (\mathbb{R}/\mathbb{Z})).$$

Our Theorem 2.1 can be understood as an extension of this study to general 2-dimensional subtori.

For other computation-related results on the Lonely Runner Problem, see, e.g., [76, 108].

2.1.5 Computing relative spectra

The proof of Theorem 2.1 is quite explicit and provides an algorithm (albeit not so simple) for calculating the structure parameters α_i, β_i for any given 2-dimensional subtorus U . As a proof of concept, we carry out this calculation for several particular choices of U .

As a first example, we determine $\mathcal{S}_1(4)$ up to the first accumulation point, namely, $1/2 - 1/4 = 1/4$. We show that this portion of the fourth Lonely Runner spectrum consists (up to finitely many exceptions) of only the infinite progressions in (2.1) and in (2.2); along the way, our argument provides a new (and more systematic) proof of the identity (2.2) from [39].

Theorem 2.3. *The set $\mathcal{S}_1(4) \cap (1/4, 1/2]$ has finite symmetric difference with the set*

$$1/4 + 1/\text{Prog}(8, 12) = 1/4 + 1/4 \text{Prog}(2, 3).$$

Characterizing this finite symmetric difference is itself a finite calculation that we have not attempted to carry out (but numerical experiments suggest that there are no exceptional elements).

For our second example, we determine $\mathcal{S}_1(3)$ up to the second accumulation point, namely, $1/2 - 2/5 = 1/10$. Recall from [65] that $\mathcal{S}_1(3) \cap (1/6, 1/2] = 1/6 + 1/\text{Prog}(9, 12) = 1/6 + 1/3 \text{Prog}(3, 4)$.

Theorem 2.4. *The set $\mathcal{S}_1(3) \cap (1/10, 1/6]$ has finite symmetric difference with the set*

$$(1/10 + 4/5 \text{Prog}(5, 7)) \cup (1/10 + 3/5 \text{Prog}(5, 9)).$$

Our third example, of a slightly different flavor, is the discovery of a new infinite family in $\mathcal{S}_1(6)$ limiting to the first accumulation point, namely, $1/2 - 1/6 = 1/3$.

Theorem 2.5. *The set $\mathcal{S}_1(6) \cap (1/3, 1/2]$ contains the set*

$$1/3 + 1/6 \text{Prog}(6, 11).$$

One could in principle calculate $\mathcal{S}_1(6) \cap (1/3, 1/2]$ up to finite symmetric difference, following the model of the proof of Theorem 2.3, but this would be fairly lengthy and we have not attempted it.

In general, before using our methods to determine $\mathcal{S}_1(n)$ up to the first accumulation point, one must first determine all tight instances for $\mathcal{S}_1(n-1)$, i.e., proper 1-dimensional subtori $T \subseteq (\mathbb{R}/\mathbb{Z})^{n-1}$ with $D(T) = \max \mathcal{S}_1(n-1)$; see Section 2.3 below. Determining all tight instances for $\mathcal{S}_1(n)$ is a (large) finite calculation for any particular n , thanks to the results of the previous chapter, but at the present it is available in the literature only for $n \in \{1, 2, 3, 5\}$. This is one reason why studying accumulation points of $\mathcal{S}_1(6)$ is currently more convenient than studying accumulation points of $\mathcal{S}_1(5)$. (Numerical simulations suggest that (2.1) holds for $n = 5$.)

In a related direction, we give a necessary and sufficient geometric condition (Proposition 2.16) for the relative spectrum of a 2-dimensional subtorus being finite, i.e., having no progressions.

2.1.6 Future directions

There are several appealing directions for future work. The most obvious is computing more examples of relative spectra for low-volume 2-dimensional subtori. Such calculations could lead to new families of counterexamples to (2.1); perhaps one could even discover such families for infinitely many values of n or make a “reasonable” conjecture constraining $\mathcal{S}_1(n)$ up to the first accumulation point. It also seems that a thorough understanding of $\mathcal{S}_1(3)$ is within reach; this set is well-ordered with order type $\omega^2 + 1$, and, using relative spectra, one could hope to characterize it up to an exceptional set of order type at most ω .

Our arguments for Theorem 2.1 would likely lead to an analogous structure theorem for $\mathcal{S}_{m-1}(U)$ whenever U is an m -dimensional proper subtorus, but the details would be messier and we have not pursued this.

Leaving the codimension-1 regime, a more ambitious open problem is obtaining a structure theorem for order-1 relative spectra of 3-dimensional subtori. One new difficulty is that the 1-dimensional subgroups of $(\mathbb{R}/\mathbb{Z})^3$ are much more complicated than the 1-dimensional subgroups of $(\mathbb{R}/\mathbb{Z})^2$ (see Section 2.2.2). It is conceivable that for a 3-dimensional proper subtorus U , the relative spectrum $\mathcal{S}_1(U)$ might “look like” a

finite union of sets obtained in a simple way from $\mathcal{S}_1(3)$; presumably an improved understanding of $\mathcal{S}_1(3)$ would be essential here.

2.1.7 Organization

We prove Theorem 2.1 in Section 2.2. In Section 2.3 we describe some techniques for determining all of the 2-dimensional subtori with a given D -value and carry out this determination for the cases that will figure in our explicit examples. We then prove Theorems 2.3, 2.4, and 2.5 in Sections 2.4, 2.5, and 2.6. In Section 2.7 we establish a criterion for the finiteness of a relative spectrum and give an example of a finite, non-empty relative spectrum.

2.2 Proof of Theorem 2.1

2.2.1 Proof strategy

Before diving into the technical details of the proof of Theorem 2.1, we sketch the main points of the argument. Fix a choice of $u, v \in \mathbb{Z}^2$ such that $U = \langle u, v \rangle_{\mathbb{R}}$. 1-dimensional subtori of U are of the form $T = \langle Au + Bv \rangle_{\mathbb{R}}$ for coprime integers A, B satisfying $A \geq 0$. The main goal is to express $D(T)$ in terms of A, B . An easy perturbation argument [65, Proposition 2.1] shows that

$$D(T) = \min_{i,j,\epsilon} D(T_{i,j,\epsilon}),$$

where for $1 \leq i < j \leq n$ and $\epsilon \in \{+, -\}$, we have set

$$T_{i,j,\epsilon} := T \cap \{x_i = \epsilon x_j\}.$$

We will analyze each $D(T_{i,j,\epsilon})$ individually.

An important observation is that (besides some trivial cases) $T_{i,j,\epsilon}$ is a finite cyclic subgroup of

$$U_{i,j,\epsilon} := U \cap \{x_i = \epsilon x_j\},$$

which is embedded in $(\mathbb{R}/\mathbb{Z})^n$ as a copy of $(\mathbb{R}/\mathbb{Z}) \times (\mathbb{Z}/K\mathbb{Z})$ for some natural number K . In Section 2.2.2, we recall the structure of finite cyclic subgroups of $(\mathbb{R}/\mathbb{Z}) \times (\mathbb{Z}/K\mathbb{Z})$ and then show that the “structure parameters” for the cyclic group $T_{i,j,\epsilon}$ depend in a linear fashion on A, B once we account for the residues of A, B with respect to some large modulus M (depending on U).

Consider the function $D : (\mathbb{R}/\mathbb{Z})^n \rightarrow [0, 1/2]$ given by $D(x_1, \dots, x_n) := \max_k |x_k - 1/2|$. The restriction of D to each connected component $U_{i,j,\epsilon,\ell}$ of $U_{i,j,\epsilon}$ is continuous and piecewise-linear. The work of Section 2.2.2 essentially reduces the task of computing $D(T_{i,j,\epsilon,\ell})$ to the task of computing

$$\min_{x \in \langle 1/q \rangle} f(x)$$

where q is a natural number and $f : \mathbb{R}/\mathbb{Z} \rightarrow \mathbb{R}$ is a continuous and piecewise-linear function assuming its minimum value at some points $\tau_1, \dots, \tau_H$. So it suffices to understand how well the τ_h 's can be approximated by multiples of $1/q$. This work is carried out in Section 2.2.3, and it is here that the expressions $1/\text{Prog}(\alpha, \beta)$ appear. The output of this section (ignoring some important edge cases) is a sector decomposition of $\mathbb{R}_{\geq 0} \times \mathbb{R}$ such that if we partition the pairs (A, B) according to the residues of A, B modulo some M and the sector containing (A, B) , then on each part of the partition we obtain a formula of the form

$$D(T) = D(U) + 1/(EA + FB)$$

for some rationals E, F .

At this point, a change of variables gives that on each part of our partition, $D(T)$ takes values in the set $D(U) + 1/\text{Prog}(\alpha, \beta)$ for some positive rationals α, β depending on E, F . In order to check that the value $D(U) + 1/(\alpha s + \beta)$ is in fact attained for all sufficiently large $s \in \mathbb{N}$, we must verify that there are coprime A, B satisfying $EA + FB = \alpha s + \beta$. This last step, which boils down to a brief analytic number theory argument that we provide in Section 2.2.4, requires us to put a few modular constraints on s , which is harmless.

2.2.2 Intersections of subtori and subgroups

The first set of lemmas concerns the possible intersections of subtori and subspaces in a 2-dimensional torus. In the following discussion, a *torus* is any topological group that is isomorphic to $(\mathbb{R}/\mathbb{Z})^n$ for some natural number n , and by a *k-dimensional subgroup* of a torus we mean a closed k -dimensional Lie subgroup. A *subtorus* is a subgroup that is also a torus.

Lemma 2.6. *Let U be a 2-dimensional torus. Then every 1-dimensional subgroup Γ of U is of the form $U' \times H$ where U' is a 1-dimensional subtorus of U and H is a finite cyclic group. In particular, every such Γ is isomorphic (as a Lie group) to $(\mathbb{R}/\mathbb{Z}) \times (\mathbb{Z}/K\mathbb{Z})$ for some natural number K .*

Proof. Let Γ be a 1-dimensional subgroup of U . Let U' denote the identity component of Γ . Then U' is a 1-dimensional subtorus of U , and Γ/U' is a discrete subgroup of $U/U' \cong \mathbb{R}/\mathbb{Z}$. Hence, considered in $\mathbb{R}/\mathbb{Z}$,

the group Γ/U' is equal to $\langle 1/K \rangle$ for some natural number K . Let $g \in \Gamma$ be such that gU' is identified with $1/K$. Then $Kg \in U'$, and (since $U' \cong \mathbb{R}/\mathbb{Z}$ is a divisible group) there is some $g' \in U'$ such that $Kg = Kg'$. Now the element $g'' := g - g'$ has order K , and $g''U'$ is still identified with $1/K$. Let $H := \langle g'' \rangle$. Then $|H| = K$ and $\Gamma = U' \times H$, as desired. $\square$

Notice that every such isomorphism $\Gamma \cong (\mathbb{R}/\mathbb{Z}) \times (\mathbb{Z}/K\mathbb{Z})$ preserves the set of points with all coordinates rational. In particular, there is a rational change of coordinates on $(\mathbb{R}/\mathbb{Z})^2$ such that in the new coordinates we have $\Gamma = (\mathbb{R}/\mathbb{Z}) \times \langle 1/K \rangle_{\mathbb{Z}} \subseteq (\mathbb{R}/\mathbb{Z})^2$. We will now describe the intersection of such a Γ with an arbitrary 1-dimensional subtorus.

Lemma 2.7. *Let K be a natural number, and let A', B' be coprime integers with $B' \neq 0$. Then*

$$((\mathbb{R}/\mathbb{Z}) \times \langle 1/K \rangle_{\mathbb{Z}}) \cap \langle (A', B') \rangle_{\mathbb{R}} = \bigcup_{\ell=0}^{K-1} ((a\ell/(Kq) + \langle 1/q \rangle_{\mathbb{Z}}) \times \{\ell/K\}),$$

where $q := |B'|$, and $0 \leq a < K$ satisfies $a \equiv A' \pmod{K}$ if $B' > 0$ and satisfies $a \equiv -A' \pmod{K}$ if $B' < 0$.

Proof. Write

$$G := ((\mathbb{R}/\mathbb{Z}) \times \langle 1/K \rangle_{\mathbb{Z}}) \cap \langle (A', B') \rangle_{\mathbb{R}}.$$

Since $B' \neq 0$, we know that the 1-dimensional subtorus $\langle (A', B') \rangle_{\mathbb{R}}$ is not equal to $(\mathbb{R}/\mathbb{Z}) \times \{0\}$. Hence G is a discrete group, and we can compute

$$\begin{aligned} G &= \{t(A', B') : t \in \mathbb{R}, tB' \in \langle 1/K \rangle_{\mathbb{Z}}\} \\ &= \langle (A'/(KB'), 1/K) \rangle_{\mathbb{Z}} \\ &= \langle (A'\delta/(Kq), 1/K) \rangle_{\mathbb{Z}}, \end{aligned}$$

where $q = |B'|$ and $\delta := B'/|B'| \in \{-1, 1\}$. Since A', B' are coprime, we have

$$G \cap ((\mathbb{R}/\mathbb{Z}) \times \{0\}) = \langle (A'\delta/q, 0) \rangle_{\mathbb{Z}} = \langle 1/q \rangle_{\mathbb{Z}} \times \{0\}.$$

Finally, for each $0 \leq \ell < K$, the point $(A'\delta\ell/(Kq), \ell/K) \in G$ tells us that

$$G \cap ((\mathbb{R}/\mathbb{Z}) \times \{\ell/K\}) = (A'\delta\ell/(Kq), \ell/K) + \langle 1/q \rangle_{\mathbb{Z}} \times \{0\} = ((a\ell/(Kq) + \langle 1/q \rangle_{\mathbb{Z}}) \times \{\ell/K\}),$$

where $a \equiv A'\delta \pmod{K}$, as desired.

□

We will apply the first lemma with $\Gamma = U_{i,j,\epsilon}$ and (after a change of coordinates) the second lemma with $\langle(A', B')\rangle_{\mathbb{R}} = T$. We first require a bit of setup.

As in the sketch in Section 2.2.1, suppose that $U \subseteq (\mathbb{R}/\mathbb{Z})^n$ ($n \geq 2$) is a proper 2-dimensional subtorus that is not contained in any subspace of the form $\{x_i = x_j\}$ or $\{x_i = -x_j\}$ for $1 \leq i < j \leq n$. Then U is the image modulo $\mathbb{Z}^n$ of the subspace $\langle u, v \rangle_{\mathbb{R}}$ for some vectors $u, v \in \mathbb{Z}^n$, and we can choose u, v such that

$$\langle u, v \rangle_{\mathbb{R}} \cap \mathbb{Z}^n = \langle u, v \rangle_{\mathbb{Z}}; \quad (2.6)$$

fix such a choice of u, v . Every 1-dimensional subtorus T of U is of the form

$$T = \langle Au + Bv \rangle_{\mathbb{R}},$$

i.e., T is the image modulo $\mathbb{Z}^n$ of $\langle Au + Bv \rangle_{\mathbb{R}}$, for some coprime integers A, B with $A \geq 0$ and $(A, B) \notin \{(0, 0), (0, -1)\}$. This correspondence between T and (A, B) is one-to-one. For ease of reference, let $\mathcal{T}$ denote the set of all such pairs (A, B) .

For the remainder of this subsection, fix a choice of $1 \leq i < j \leq n$ and a choice of a sign $\epsilon \in \{+, -\}$. The subgroup

$$U_{i,j,\epsilon} := U \cap \{x_i = \epsilon x_j\}$$

is a 1-dimensional subgroup of $U \cong (\mathbb{R}/\mathbb{Z})^2$ (recall our assumption that U does not lie in the subspace $\{x_i = \epsilon x_j\}$), so Lemma 2.6 tells us that $U_{i,j,\epsilon}$ is isomorphic to the cartesian product of $\mathbb{R}/\mathbb{Z}$ and $\mathbb{Z}/K\mathbb{Z}$ for some natural number $K = K_{i,j,\epsilon}$. In order to apply Lemma 2.7, we need to write down an explicit isomorphism $U \cong (\mathbb{R}/\mathbb{Z})^2$ in coordinates “adapted” to $U_{i,j,\epsilon}$. To this end, fix a primitive vector $u' = u'_{i,j,\epsilon} \in \mathbb{Z}^n$ such that the identity component of $U_{i,j,\epsilon}$ equals $\langle u' \rangle_{\mathbb{R}}$. Now fix a vector $v' = v'_{i,j,\epsilon} \in \mathbb{Z}^n$ such that

$$\langle u', v' \rangle_{\mathbb{Z}} = \langle u, v \rangle_{\mathbb{Z}};$$

such a v' exists by the primitivity assumption on u' . Hence the linear map $\psi = \psi_{i,j,\epsilon} : \langle u, v \rangle_{\mathbb{R}} \rightarrow \mathbb{R}^2$ given by $\psi(u') := (1, 0)$ and $\psi(v') := (0, 1)$ descends to an isomorphism from U to $(\mathbb{R}/\mathbb{Z})^2$. Since ψ sends the identity component of $U_{i,j,\epsilon}$ to $\mathbb{R}/\mathbb{Z} \times \{0\}$, we see that

$$\psi(U_{i,j,\epsilon}) = (\mathbb{R}/\mathbb{Z}) \times \langle 1/K \rangle_{\mathbb{Z}}$$

is a concrete witness to the isomorphism described earlier in this paragraph.

Write

$$\psi(T) = \psi(\langle Au + Bv \rangle_{\mathbb{R}}) = \langle A\psi(u) + B\psi(v) \rangle_{\mathbb{R}}.$$

Since $\langle u', v' \rangle_{\mathbb{Z}} = \langle u, v \rangle_{\mathbb{Z}}$, the vectors $\psi(u), \psi(v)$ have integer coordinates, say,

$$\psi(u) = (z_1, z_2) \quad \text{and} \quad \psi(v) = (z_3, z_4),$$

with $z_1, z_2, z_3, z_4 \in \mathbb{Z}$ (depending on i, j, ϵ). So we write

$$\psi(T) = \langle A\psi(u) + B\psi(v) \rangle_{\mathbb{R}} = \langle (Az_1 + Bz_3, Az_2 + Bz_4) \rangle_{\mathbb{R}}. \quad (2.7)$$

Since $\psi(u), \psi(v)$ are linearly independent, we have $z_1 z_4 \neq z_2 z_3$. Now we apply Lemma 2.7 in order to obtain the main result of this section, as follows.

Proposition 2.8. *Let $n, U, u, v, i, j, \epsilon, K = K_{i,j,\epsilon}, \psi = \psi_{i,j,\epsilon}, z_1 = z_{1,i,j,\epsilon}, z_2 = z_{2,i,j,\epsilon}, z_3 = z_{3,i,j,\epsilon}, z_4 = z_{4,i,j,\epsilon}$ be as above. Then there is a natural number M (depending on all of the parameters introduced so far) such that for each choice of $0 \leq \mathfrak{x}, \mathfrak{z} < M$ and each choice of a sign $\delta \in \{+, -\}$, the following holds: There are rational numbers $E = E_{\mathfrak{x}, \mathfrak{z}, \delta}, F = F_{\mathfrak{x}, \mathfrak{z}, \delta}$ (not both zero) and a nonnegative integer $a = a_{\mathfrak{x}, \mathfrak{z}, \delta} < K'_{\mathfrak{x}, \mathfrak{z}, \delta}$ such that for $T = \langle Au + Bv \rangle_{\mathbb{R}}$, we have*

$$\psi(T_{i,j,\epsilon}) = \bigcup_{\ell=0}^{K-1} ((a\ell/(Kq) + \langle 1/q \rangle_{\mathbb{Z}}) \times \{\ell/K\})$$

with

$$q = EA + FB$$

whenever $(A, B) \in \mathcal{T}$ satisfies $A \equiv \mathfrak{x} \pmod{M}$ and $B \equiv \mathfrak{z} \pmod{M}$, and

$$\delta(Az_2 + Bz_4) > 0.$$

Proof. Let

$$\omega := \gcd(Az_1 + Bz_3, Az_2 + Bz_4).$$

Then the integers $A' := (Az_1 + Bz_3)/\omega$ and $B' := (Az_2 + Bz_4)/\omega$ are coprime, and we can apply Lemma 2.7 to find that

$$\psi(T_{i,j,\epsilon}) = \bigcup_{\ell=0}^{K-1} ((a\ell/(Kq) + \langle 1/q \rangle_{\mathbb{Z}}) \times \{\ell/K\}),$$

with $q := \delta B'$ and $0 \leq a < K$ satisfying $a \equiv \delta A' \pmod{K}$ (by the definition of δ). In particular, we will obtain the conclusion of the lemma with the choices

$$E := \delta z_2 / \omega \quad \text{and} \quad F := \delta z_4 / \omega.$$

To complete the proof of the proposition, it remains to show that the quantities A', B' are linear functions of A, B , under appropriate modular conditions. In particular, it suffices to show that ω depends only on the residues of A, B modulo some fixed natural number M to be specified shortly.

Notice that ω divides

$$z_4(Az_1 + Bz_3) - z_3(Az_2 + Bz_4) = (z_1z_4 - z_2z_3)A$$

and

$$z_1(Az_2 + Bz_4) - z_2(Az_1 + Bz_3) = (z_1z_4 - z_2z_3)B.$$

Since A, B are coprime by assumption, we conclude that ω divides

$$M := |z_1z_4 - z_2z_3|,$$

which is nonzero by the remark following (2.7). It follows that ω depends only on the residues of $Az_1 - Bz_3, Az_2 + Bz_4$ modulo M , which in turn depend only on the residues of A, B modulo M . $\square$

2.2.3 Approximating minima

The next set of lemmas concerns approximating minima of continuous piecewise-linear functions on $\mathbb{R}/\mathbb{Z}$; see Figure 2.1.

For $\tau, b \in \mathbb{Q}$ and $q \in \mathbb{N}$, define

$$\text{Approx}^-(\tau, b; q) := \min\{\tau - (b/q + r/q) : r \in \mathbb{Z}, b/q + r/q \leq \tau\}$$

and

$$\text{Approx}^+(\tau, b; q) := \min\{(b/q + r/q) - \tau : r \in \mathbb{Z}, b/q + r/q \geq \tau\}$$

to be the errors in best approximations to τ (from below and above) by numbers of the form $b/q + r/q$, for $r \in \mathbb{Z}$. Since $\text{Approx}^-(\tau, b; q) = \text{Approx}^-(\tau + 1/q, b; q)$, this quantity is also well-defined for $\tau \in \mathbb{Q}/\mathbb{Z}$ (and likewise for $\text{Approx}^+(\tau, b; q)$). The next lemma shows that once we specify the residue class of q with respect

to some modulus depending on τ, b , the functions $\text{Approx}^-(\tau, b; q)$, $\text{Approx}^+(\tau, b; q)$ have a simple form. This calculation is the source of the expressions $1/\text{Prog}(\alpha, \beta)$ in Theorem 2.1.

Lemma 2.9. *Let $\tau, b \in \mathbb{Q}$, and write $\tau = w/x$ and $b = y/z$ with w, x, y, z nonnegative integers. Let q be a natural number. Then*

$$\text{Approx}^-(\tau, b; q) = \frac{R^-}{xzq} \quad \text{and} \quad \text{Approx}^+(\tau, b; q) = \frac{R^+}{xzq},$$

where $0 \leq R^-, R^+ < xz$ satisfy

$$R^- \equiv wzq - xy \pmod{xz} \quad \text{and} \quad R^+ \equiv xy - wzq \pmod{xz}.$$

Notice that R^-, R^+ depend on only the residue of q modulo xz , and that the parameters w, x, y, z depend on only τ, b . Note also that $R^- + R^+ = xz$ except when $R^- = R^+ = 0$.

Proof. We prove the lemma only for $\text{Approx}^-(\tau, b; q)$ since the argument for $\text{Approx}^+(\tau, b; q)$ is identical. In the definition of $\text{Approx}^-(\tau, b; q)$, the minimum is achieved for $r = \lfloor q(\tau - b/q) \rfloor = \lfloor q\tau - b \rfloor$. We compute

$$q\tau - b = wq/x - y/z = (wzq - xy)/(xz).$$

The fractional part of $(wzq - xy)/(xz)$ is $R^-/(xz)$, so

$$r = (q\tau - b) - R^-/(xz)$$

and

$$\text{Approx}^-(\tau, b; q) = \tau - (b/q + r/q) = R^-/(xzq),$$

as desired. □

We will now apply this lemma where τ ranges over the points at which a piecewise-linear function achieves its minimum value.

Lemma 2.10. *Let $f : \mathbb{R}/\mathbb{Z} \rightarrow \mathbb{R}$ be a continuous, piecewise-linear function with finitely many pieces, each with rational slope and rational endpoints, and let b be a fixed rational number. Then there is a natural number M (depending on f, b) such for each integer $0 \leq Q < M$, the following holds: There is a nonnegative*

rational γ_Q (depending also on f, b) such that

$$\min_{t \in b/q + \langle 1/q \rangle} f(t) = \min_{t \in \mathbb{R}/\mathbb{Z}} f(t) + \frac{\gamma_Q}{q} \quad (2.8)$$

for all sufficiently large natural numbers q satisfying $q \equiv Q \pmod{M}$.

Proof. Write $m := \min_{t \in \mathbb{R}/\mathbb{Z}} f(t)$. We first dispose of the trivial case where $f(t) = m$ on some interval $I \subseteq \mathbb{R}/\mathbb{Z}$ of strictly positive length. Since $b/q + \langle 1/q \rangle$ is $1/q$ -dense in $\mathbb{R}/\mathbb{Z}$, we see that

$$\min_{t \in b/q + \langle 1/q \rangle} f(x) = m$$

for all $q \geq |I|$ (with no mention of modular restrictions), and the conclusion of the lemma holds with $M = 1$ and $\gamma_0 = 0$.

Now suppose that f achieves the value m at only finitely many points $\tau_1, \dots, \tau_H$. Then there are some $\delta, \varepsilon > 0$ such that $f(t) > m + \delta$ whenever t does not lie within ε of some τ_h . Since f is continuous and $b/q + \langle 1/q \rangle$ is $1/q$ -dense, we know that $\min_{t \in b/q + \langle 1/q \rangle} f(t) < m + \delta$ for all q sufficiently large (in terms of f). In particular, the minimum in (2.8) is achieved at a point t that lies within ε of some τ_h ; for the remainder of the proof we will assume that q is large enough for this to be the case.

For each τ_h , let $-\lambda_h^-, \lambda_h^+$ denote the (rational by assumption) slopes of the pieces of f directly to the left and right (respectively) of τ_h ; notice that $\lambda_h^-, \lambda_h^+ > 0$. After possibly shrinking ε further, we find that

$$f(\tau_h - t) = m + \lambda_h^- t \quad \text{and} \quad f(\tau_h + t) = m + \lambda_h^+ t$$

for every h and for all $0 \leq t < \varepsilon$. Hence $\min_{t \in b/q + \langle 1/q \rangle} f(t)$ is equal to the minimum of the values

$$m + \lambda_h^- \text{Approx}^-(\tau_h, b; q) \quad \text{and} \quad m + \lambda_h^+ \text{Approx}^+(\tau_h, b; q)$$

over $1 \leq h \leq H$. Write $b = y/z$ and $\tau_h = w_h/x_h$ as fractions, and let $M := \text{lcm}(z, x_1, \dots, x_H)$. Then for each h we have

$$m + \lambda_h^- \text{Approx}^-(\tau_h, b; q) = m + \frac{\lambda_h^- R_h^-}{x_h z q} \quad \text{and} \quad m + \lambda_h^+ \text{Approx}^+(\tau_h, b; q) = m + \frac{\lambda_h^+ R_h^+}{x_h z q},$$

where R_h^-, R_h^+ are the parameters output by Lemma 2.9 for τ_h, b . It is important that R_h^-, R_h^+ depend on only the residue of q modulo $x_h z$, which divides M . For each $0 \leq Q < M$, consider the case corresponding to $q \equiv Q \pmod{M}$, and let γ_Q denote the minimum of the values $\lambda_h^- R_h^- / (x_h z), \lambda_h^+ R_h^+ / (x_h z)$ over all h . (Since

$\lambda_h^-, \lambda_h^+ > 0$, we have $\gamma_Q \geq 0$.) Then for all sufficiently large q congruent to Q modulo M , we have

$$\min_{t \in b/q + \langle 1/q \rangle} f(t) = m + \frac{\gamma_Q}{q},$$

as desired. $\square$

We now apply this lemma where f is the restriction of D to a connected component of $U_{i,j,\epsilon}$ and the quantities b, q depend “nicely” on the subtorus T , as described in Proposition 2.8.




Figure 2.1: This figure illustrates some of the quantities in Lemmas 2.9 and 2.10. Here we are approximating the minimum at $\tau = 1/6$ using the coset $1/12 + \langle 1/4 \rangle_{\mathbb{Z}}$, and the slopes of the pieces to the left and right of τ are $-\lambda^- = -1$ and $\lambda^+ = 2$.

Proposition 2.11. *Let $n \geq 2$ be a natural number, and let $U \subseteq (\mathbb{R}/\mathbb{Z})^n$ be a 2-dimensional proper subtorus that is not contained in any subspace of the form $\{x_i = x_j\}$ or $\{x_i = -x_j\}$ for $1 \leq i < j \leq n$. Let $u, v \in \mathbb{Z}^n$ be such that $U = \langle u, v \rangle_{\mathbb{R}}$ and (2.6) is satisfied. Then there are a natural number M' and a finite set $\mathcal{L}$ of rational half-lines in $\mathbb{R}_{\geq 0} \times \mathbb{R}$ (depending on the parameters introduced so far) such that for each choice of $0 \leq \mathfrak{x}, \mathfrak{z} < M'$ the following holds:*

There are a partition of $(\mathbb{R}_{\geq 0} \times \mathbb{R}) \setminus \mathcal{L}$ into sectors $\sigma_1 = \sigma_{\mathfrak{x}, \mathfrak{z}, 1}, \dots, \sigma_X = \sigma_{\mathfrak{x}, \mathfrak{z}, X(\mathfrak{x}, \mathfrak{z})}$; and $\kappa = \kappa_{\mathfrak{x}, \mathfrak{z}, k} \in \{0, 1\}$

and rationals $E' = E'_{\mathfrak{x}, \mathfrak{z}, k}, F' = F'_{\mathfrak{x}, \mathfrak{z}, k}$ for each $1 \leq k \leq X(\mathfrak{x}, \mathfrak{z})$, such that for $T = \langle Au + Bv \rangle_{\mathbb{R}}$, we have

$$D(T) = D(U) + \frac{\kappa}{E'A + F'B}$$

whenever $(A, B) \in \mathcal{T} \setminus \mathcal{L}$ satisfies $(A, B) \equiv (\mathfrak{x}, \mathfrak{z}) \pmod{M'}$, $(A, B) \in \sigma_k$, and $A^2 + B^2$ is sufficiently large. Moreover, for each half-line $L \in \mathcal{L}$, there are $\kappa = \kappa_{\mathfrak{x}, \mathfrak{z}, L} \in \{0, 1\}$ and rationals $E' = E'_{\mathfrak{x}, \mathfrak{z}, L}, F' = F'_{\mathfrak{x}, \mathfrak{z}, L}$ such that for $T = \langle Au + Bv \rangle_{\mathbb{R}}$, we have

$$D(T) = D(U) + \frac{\kappa}{E'A + F'B}$$

whenever $(A, B) \in \mathcal{T} \cap L$ satisfies $(A, B) \equiv (\mathfrak{x}, \mathfrak{z}) \pmod{M'}$ and $A^2 + B^2$ is sufficiently large.

Proof. Let M'' be the lcm of all of the values of $M = M_{i,j,\epsilon}$ output by Proposition 2.8 as we range over indices $1 \leq i < j \leq n$ and signs $\epsilon \in \{+, -\}$. For the remainder of the proof, assume that the residues of A, B modulo M'' are fixed. Notice that this choice fixes the residues of A, B modulo each individual M output by Proposition 2.8. Our partition of $\mathbb{R}_{\geq 0} \times \mathbb{R}$ into sectors will be the least common refinement of partitions $\mathcal{P}_{i,j,\epsilon}$ and $\mathcal{P}_{i,j,\epsilon,\ell;i',j',\epsilon',\ell'}$ defined later in the proof. Each of these partitions will have a finite number of sectors, so their refinement will also have a finite number of sectors. Since $\mathcal{T}$ contains at most one point on each line through the origin, the intersection of $\mathcal{T}$ with the boundary lines defining our sectors will be finite. By taking $\sqrt{A^2 + B^2}$ sufficiently large, we can ensure that (A, B) never lies on the boundary of a sector, and we will tacitly use this assumption in what follows.

Consider now a fixed choice of indices $1 \leq i < j \leq n$ and a sign $\epsilon \in \{+, -\}$, and take $K, \psi, z_1, z_2, z_3, z_4$ (all depending on i, j, ϵ) as in the statement of Proposition 2.8. Let $\mathcal{P}_{i,j,\epsilon}$ be the partition of $\mathbb{R}_{\geq 0} \times \mathbb{R}$ into two sectors according to the sign of the expression $xz_2 + yz_4$ (for $(x, y) \in \mathbb{R}_{\geq 0} \times \mathbb{R}$). Fix one of the sectors $\sigma \in \mathcal{P}_{i,j,\epsilon}$, and assume that $(A, B) \in \sigma$. Then Proposition 2.8 supplies integers $0 \leq a < K$ and rationals E, F (depending on i, j, ϵ, σ) such that

$$\psi(T_{i,j,\epsilon}) = \bigcup_{\ell=0}^{K-1} ((a\ell/(Kq) + \langle 1/q \rangle_{\mathbb{Z}}) \times \{\ell/K\}),$$

where $q = EA + FB$. Now write $U_{i,j,\epsilon}$ as the disjoint union

$$U_{i,j,\epsilon} = \bigcup_{\ell=0}^{K-1} U_{i,j,\epsilon,\ell},$$

such that

$$\psi(U_{i,j,\epsilon,\ell}) = (\mathbb{R}/\mathbb{Z}) \times \{\ell/K\}$$

for each $0 \leq \ell < K$. Also define

$$T_{i,j,\epsilon,\ell} := T \cap U_{i,j,\epsilon,\ell} = T_{i,j,\epsilon} \cap U_{i,j,\epsilon,\ell},$$

and notice that

$$\psi(T_{i,j,\epsilon,\ell}) = (a\ell/(Kq) + \langle 1/q \rangle_{\mathbb{Z}}) \times \{\ell/K\}.$$

Since the function D is continuous and piecewise linear with finitely many pieces, so is its restriction to $U_{i,j,\epsilon,\ell}$. It follows that the function $D_{i,j,\epsilon,\ell} : \mathbb{R}/\mathbb{Z} \rightarrow \mathbb{R}$ defined by $D_{i,j,\epsilon,\ell}(x) := (D \circ \psi^{-1})(x, \ell/K)$ satisfies the hypotheses of Lemma 2.10; let M, γ_Q (for $0 \leq Q < M$) be as in the conclusion of that lemma with $b = a\ell/K$. The residue of q modulo M depends only on the residues of A, B modulo the product of M with the denominators of E, F ; call this natural number $M_{i,j,\epsilon,\sigma,\ell}$. Suppose (as we will ensure shortly) that the residue classes of A, B modulo $M_{i,j,\epsilon,\sigma,\ell}$ are fixed. This also fixes the residue class of q modulo M , say, $q \equiv Q \pmod{M}$; write γ for the γ_Q from Lemma 2.10.

Now Lemma 2.10 tells us that

$$D(T_{i,j,\epsilon,\ell}) = D_{i,j,\epsilon,\ell}(a\ell/(Kq) + \langle 1/q \rangle_{\mathbb{Z}}) = D(U_{i,j,\epsilon,\ell}) + \gamma/q$$

whenever q is sufficiently large, say, at least some $q_0 = q_{0,i,j,\epsilon,\sigma,\ell}$. Since q is a rational linear combination of A, B , we see that the set of pairs (A, B) corresponding to $q < q_0$ lies on a finite set $\mathcal{L}_{i,j,\epsilon,\sigma,\ell}$ of rational half-lines.² Suppose now that $(A, B) \notin \mathcal{L}_{i,j,\epsilon,\sigma,\ell}$. By dividing through by γ if $\gamma \neq 0$, we can write the previous centered equation as

$$D(T_{i,j,\epsilon,\ell}) = D(U_{i,j,\epsilon,\ell}) + \frac{\kappa}{E'A + F'B}, \quad (2.9)$$

where $\kappa \in \{0, 1\}$ and E', F' are still rational (and now depend on ℓ in addition to i, j, ϵ).

We pause to update a few parameters. Let $\mathcal{L}$ be the union of the sets $\mathcal{L}_{i,j,\epsilon,\sigma,\ell}$ from the previous paragraph; for this paragraph and the two that follow, assume that (A, B) does not lie on $\mathcal{L}$, so that (2.9) holds for all i, j, ϵ, ℓ . Let M' be the lcm of M'' and the natural numbers $M_{i,j,\epsilon,\sigma,\ell}$ from the previous paragraph; for the remainder of the proof, assume that $(A, B) \equiv (\mathfrak{x}, \mathfrak{z}) \pmod{M'}$ for fixed $0 \leq \mathfrak{x}, \mathfrak{z} < M'$. Let $\mathcal{P}'$ (depending on $\mathfrak{x}, \mathfrak{z}$) be the least common refinement of the partitions $\mathcal{P}_{i,j,\epsilon}$ from the previous paragraph; for the remainder of the proof, assume that (A, B) lies in a fixed sector σ of $\mathcal{P}_{i,j,\epsilon}$. We will later refine the partition $\mathcal{P}'$ further.

²This set of pairs (A, B) lies naturally on a finite set of rational lines. Any such non-vertical line intersects $\mathbb{R}_{\geq 0} \times \mathbb{R}$ in a half-line, and any such vertical line can be arbitrarily broken into two half-lines.

It is a basic fact (see, e.g., [65, Proposition 2.1]) that

$$D(T) = \min_{i,j,\epsilon} D(T_{i,j,\epsilon}).$$

Hence

$$D(T) = \min_{i,j,\epsilon,\ell} D(T_{i,j,\epsilon,\ell}).$$

Let Y be the set of quadruples (i, j, ϵ, ℓ) such that $D(U_{i,j,\epsilon,\ell}) = D(U)$, and note that

$$D(T) = \min_{(i,j,\epsilon,\ell) \in Y} D(T_{i,j,\epsilon,\ell})$$

for $A^2 + B^2$ sufficiently large since the right-hand side tends to $D(U)$ and each $(i, j, \epsilon, \ell) \notin Y$ has $D(T_{i,j,\epsilon,\ell})$ uniformly bounded above $D(U)$. We first dispose of the special case in which there is some $(i, j, \epsilon, \ell) \in Y$ with $\kappa_{i,j,\epsilon,\ell} = 0$. If this occurs, then

$$D(T) = D(U)$$

for all (A, B) under consideration, and we obtain the conclusion of the proposition with $\kappa := 0$ and (say) $E' = 1, F' = 0$. It remains to consider the case where $\kappa_{i,j,\epsilon,\ell} = 1$ for all $(i, j, \epsilon, \ell) \in Y$.

For each pair of distinct tuples $(i, j, \epsilon, \ell), (i', j', \epsilon', \ell') \in Y$, let us compare the expressions in (2.9) appearing for these two tuples. There are rationals E_1, E_2, F_1, F_2 such that

$$D(T_{i,j,\epsilon,\ell}) = D(U) + \frac{1}{E_1 A + F_1 B} \quad \text{and} \quad D(T_{i',j',\epsilon',\ell'}) = D(U) + \frac{1}{E_2 A + F_2 B}.$$

Which of these two expressions is smaller is determined by the sign of the expression

$$(E_2 A + F_2 B) - (E_1 A + F_1 B) = (E_2 - E_1)A + (F_2 - F_1)B.$$

Hence there is a partition $\mathcal{P}_{i,j,\epsilon,\ell;i',j',\epsilon',\ell'}$ of $\mathbb{R}_{\geq 0} \times \mathbb{R}$ into at most two sectors such that on the each sector, the above expression is either always non-negative or always non-positive. Let $\mathcal{P}$ be the least common refinement of $\mathcal{P}'$ (from above) and the partitions $\mathcal{P}_{i,j,\epsilon,\ell;i',j',\epsilon',\ell'}$. Hence each sector σ of $\mathcal{P}$ determines a tuple (i, j, ϵ, ℓ) , with its associated rationals E', F' , such that

$$D(T) = D(T_{i,j,\epsilon,\ell}) = D(U) + \frac{1}{E' A + F' B}$$

whenever $(A, B) \in (\mathcal{T} \setminus \mathcal{L}) \cap \sigma$. This establishes the first conclusion of the proposition.

It remains to analyze what happens when (A, B) lies on $\mathcal{L}$. Fix a half-line $L \in \mathcal{L}$, and consider pairs $(A, B) \in L$. If $L \cap \mathcal{T}$ is finite, then there is nothing to show, so assume that this intersection is infinite. Let $Y' = Y'_{\mathbf{x}, \mathbf{z}, L}$ denote the set of all quadruples $(i, j, \epsilon, \ell) \in Y$ whose corresponding values of q tend to infinity as (A, B) tends to infinity along L . We claim that Y' is nonempty. It was shown in the previous chapter that the volume of the torus $T = \langle Au + Bv \rangle_{\mathbb{R}}$ tends to infinity with $A^2 + B^2$ and that therefore T becomes $o(1)$ -dense in U as $A^2 + B^2$ tends to infinity; in particular, $D(T)$ tends to $D(U)$. Since each quadruple $(i, j, \epsilon, \ell) \notin Y'$ has $D(T_{i,j,\epsilon,\ell})$ uniformly bounded away from $D(U)$ from above, we see that Y' is nonempty, as claimed. Now we repeat the argument from the previous two paragraphs, with Y replaced by Y' , and we note that, sufficiently far from the origin, the half-line L lies entirely in a single sector of $\mathcal{P}$. This establishes the second conclusion of the proposition. $\square$

2.2.4 Coprime pairs in arithmetic progressions

The third stage of our preparatory lemmas concerns when long arithmetic progressions in $\mathbb{Z}^2$ are guaranteed to contain some points with coprime coordinates. The first lemma will help us find coprime pairs (A, B) on intersections of rational lines with the sectors from Proposition 2.11.

Lemma 2.12. *Let $C \geq 0$ be a real number, and let a_1, a_2 be nonzero integers. For N a parameter, let b_1, b_2 be integers with $|b_1|, |b_2| \leq CN$, and let I be an interval (of integers) of length N . Suppose that $\gcd(a_1, a_2, b_1, b_2) = 1$ and $a_1 b_2 \neq a_2 b_1$. Then*

$$|\{x \in I : \gcd(a_1 x + b_1, a_2 x + b_2) = 1\}|$$

tends to infinity with N (uniformly in b_1, b_2, I).

Proof. Note that $\gcd(a_1 x + b_1, a_2 x + b_2)$ divides

$$Z := |a_1(a_2 x + b_2) - a_2(a_1 x + b_1)| = |a_1 b_2 - a_2 b_1|,$$

which by assumption is nonzero and has size $Z = O(N)$. Hence, we wish to count $x \in I$ such that $\gcd(a_1 x + b_1, a_2 x + b_2)$ is coprime to Z . Let p be a prime divisor of Z . The forms $a_1 x + b_1, a_2 x + b_2$ cannot both be identically zero modulo p by the assumption that $\gcd(a_1, a_2, b_1, b_2) = 1$, so there is at most a single residue class x_p modulo p where both forms vanish. If there is no such residue class, then pick x_p arbitrarily.

By the Chinese Remainder Theorem, there is some integer R such that $R \equiv x_p \pmod{p}$ for each prime p dividing Z , and we have $\gcd(a_1 x + b_1, a_2 x + b_2) = 1$ whenever $x - R$ is coprime to Z . So it suffices to

lower-bound the number of elements of $J := I - R$ that are coprime to Z . Let μ denote the Möbius function, and recall that $\sum_{d|m} \mu(d)$ equals 1 if $m = 1$ and equals 0 if $m > 1$. Using this identity, we compute

$$\begin{aligned}
|\{x \in J : \gcd(x, Z) = 1\}| &= \sum_{x \in J} \mathbb{1}_{\gcd(x, Z)=1} \\
&= \sum_{x \in J} \sum_{d|x, Z} \mu(d) \\
&= \sum_{d|Z} \mu(d) \cdot |\{x \in J : d|x\}| \\
&= \sum_{d|Z} \mu(d) (|J|/d + O(1)) \\
&= (\varphi(Z)/Z)N + O(\tau(Z)),
\end{aligned}$$

where φ is Euler's totient function and τ is the divisor function. Recalling that $Z = O(N)$, we apply the standard bounds $\varphi(Z)/Z \gg 1/\log \log(Z)$ and $\tau(Z) \ll e^{O(\log(Z)/\log \log(Z))}$ to find that

$$|\{x \in J : \gcd(x, Z) = 1\}| \gg N/\log \log(N),$$

which certainly tends to infinity with N . □

The next lemma will help us characterize the coprime pairs (A, B) on the exceptional half-lines from Proposition 2.11.

Lemma 2.13. *Let a_1, b_1, a_2, b_2 be integers with $a_1 > 0$. Then there is a natural number M such that for x a sufficiently large natural number, whether or not $\gcd(a_1x + b_1, a_2x + b_2) = 1$ depends only on the residue of x modulo M .*

Proof. As in the proof of the previous lemma, note that $\gcd(a_1x + b_1, a_2x + b_2)$ divides

$$a_1(a_2x + b_2) - a_2(a_1x + b_1) = a_1b_2 - a_2b_1.$$

Suppose first that $a_1b_2 - a_2b_1 = 0$. For each $i \in \{1, 2\}$, let

$$a'_i := a_i / \gcd(a_i, b_i) \quad \text{and} \quad b'_i := b_i / \gcd(a_i, b_i),$$

so that $\gcd(a'_i, b'_i) = 1$ and we have

$$a_i x + b_i = \gcd(a_i, b_i)(a'_i x + b'_i).$$

Now $a_1 b_2 = a_2 b_1$ implies that $a'_1 b'_2 = a'_2 b'_1$, and from $\gcd(a'_1, b'_1) = \gcd(a'_2, b'_2) = 1$ we conclude that $(a'_1, b'_1) = \pm(a'_2, b'_2)$. It follows that $a_1 x + b_1, a_2 x + b_2$ are both integer multiples of $a'_1 x + b'_1$. For sufficiently large x , we have $a'_1 x + b'_1 > 1$ and hence $\gcd(a_1 x + b_1, a_2 x + b_2) \geq a'_1 x + b'_1 \neq 1$, so the conclusion of the lemma holds with $M = 1$ (say).

Now, suppose that $a_1 b_2 - a_2 b_1 \neq 0$. Then $\gcd(a_1 x + b_1, a_2 x + b_2)$ depends only on the residue of x modulo $M := |a_1 b_2 - a_2 b_1|$, as desired. $\square$

We are finally ready to prove Theorem 2.1.

Proof of Theorem 2.1. If U is contained in some subspace of the form $\{x_i = \epsilon x_j\}$, then the projection $\pi(U)$ onto all but the i -th coordinate clearly satisfies

$$\mathcal{S}_1(U) = \mathcal{S}_1(\pi(U)),$$

and we can apply induction on n . Now assume that U is not contained in any subspace of the form $\{x_i = \epsilon x_j\}$, and apply Proposition 2.11. In the language of Proposition 2.11, fix a choice of residues $0 \leq \mathfrak{N}, \mathfrak{Z} < M'$, and assume that $(A, B) \in \mathcal{T}$ has $A^2 + B^2$ sufficiently large and satisfies $A \equiv \mathfrak{N} \pmod{M'}$ and $B \equiv \mathfrak{Z} \pmod{M'}$.

We start with the pairs (A, B) not lying on $\mathcal{L}$: Fix a sector σ from Proposition 2.11, and assume that $(A, B) \in (\mathcal{T} \setminus \mathcal{L}) \cap \sigma$. For such (A, B) , there are $\kappa \in \{0, 1\}$ and rationals E', F' such that

$$D(T) = D(U) + \frac{\kappa}{E'A + F'B}.$$

If $\kappa = 0$, then $D(T) = D(U)$ contributes only a single value to $\mathcal{S}_1(U)$. Now suppose that $\kappa = 1$. Write $E' = y/w$ and $F' = z/w$ where y, z are coprime integers and w is a rational number, so that

$$\frac{1}{E'A + F'B} = \frac{w}{yA + zB}.$$

Notice that $yA + zB \equiv y\mathfrak{N} + z\mathfrak{Z} \pmod{M}$ for all (A, B) under consideration.

For s' an integer, define the level set

$$V(s') := \{(A', B') \in ((\mathfrak{N} + M'\mathbb{Z}) \times (\mathfrak{Z} + M'\mathbb{Z})) \cap \sigma : yA' + zB' = s'M' + y\mathfrak{N} + z\mathfrak{Z}\}.$$




Figure 2.2: This figure illustrates how level sets intersect a sector. Here, the sector $\sigma = \{x/3 \leq y \leq 2x\}$ is shaded. Filled-in circles indicate elements of $\mathcal{T} \cap \sigma$ (i.e., points with coprime coordinates), and outlined circles indicate other lattice points in σ . The gray diagonal lines with slope -1 indicate the various level sets $V(s')$ of the expression $A + B$ (where there are no modular constraints placed on A, B).

We have not required A', B' to be coprime in the definition of $V(s')$; the expression $yA + zB$ assumes the value $s'M' + y\mathfrak{x} + z\mathfrak{z}$ if and only if $V(s')$ contains a point of $\mathcal{T} \setminus \mathcal{L}$. For s' sufficiently large, the set $V(s')$ intersects each half-line of $\mathcal{L}$ in at most one point, so $V(s')$ contains a point of $\mathcal{T} \setminus \mathcal{L}$ whenever $V(s')$ contains a sufficiently large number of points with coprime coordinates. Since y, z are coprime, there are natural numbers b'_1, b'_2 such that $yb'_1 + zb'_2 = 1$. Hence

$$y(s'M'b'_1 + \mathfrak{x}) + z(s'M'b'_2 + \mathfrak{z}) = s'M' + y\mathfrak{x} + z\mathfrak{z},$$

and $V(s')$ consists of the points

$$(zM'x + s'M'b'_1 + \mathfrak{x}, -yM'x + s'M'b'_2 + \mathfrak{z})$$

for x in an appropriate interval $I(s')$ of integers. See Figure 2.2 for an illustration.

Notice that $I(s')$ is empty if $s'M' + y\mathfrak{x} + z\mathfrak{z} < 0$. The argument from the last paragraph of the proof of Proposition 2.11 shows that the quantity $E'A + F'B$ tends to infinity as $A^2 + B^2$ tends to infinity. In particular, each $I(s')$ is finite, and the length of $I(s')$ grows linearly with s' as s' tends to infinity. We wish to apply Lemma 2.12 with the parameters

$$I := I(s'), \quad a_1 := zM', \quad a_2 := -yM', \quad b_1 := s'M'b'_1 + \mathfrak{x}, \quad b_2 := s'M'b'_2 + \mathfrak{z};$$

the parameter $N := |I|$ grows linearly with s' (as observed above), and the constant C in the statement of Lemma 2.12 can be chosen appropriately (depending on M', y, z, b'_1, b'_2). We check that

$$a_1 b_2 - a_2 b_1 = (M')^2 s' + M'(y\mathfrak{x} + z\mathfrak{z})$$

is nonzero for s' sufficiently large. It remains to check whether or not $\gcd(a_1, a_2, b_1, b_2) = 1$. It is always the case that $\gcd(a_1, a_2, b_1, b_2)$ divides $\gcd(a_1, a_2) = M'$. In particular, whether or not $\gcd(a_1, a_2, b_1, b_2) = 1$ depends only on the residue class of s' modulo M' . Fix some $0 \leq r < M'$, and suppose that $s' = M's + r$ for s an integer. If r is such that $\gcd(a_1, a_2, b_1, b_2) > 1$, then $V(s')$ does not contain any points with coprime coordinates. Suppose instead that r is such that $\gcd(a_1, a_2, b_1, b_2) = 1$, in which case Lemma 2.12 ensures that the number of points of $V(s')$ with coprime coordinates tends to infinity with s' . By the discussion in the previous paragraph, this ensures that $D(T)$ assumes all of the values

$$D(U) + \frac{w}{(M')^2 s + M'r + y\mathfrak{x} + z\mathfrak{z}}$$

for s sufficiently large. Setting $\alpha := (M')^2/w$ and $\beta := (M'r + y\mathfrak{x} + z\mathfrak{z})/w$, we find that $D(T)$ assumes all but finitely many elements of the set

$$D(U) + \frac{1}{\text{Prog}(\alpha, \beta)}.$$

We obtain only finitely many such progressions because there were finitely many choices of $\mathfrak{x}, \mathfrak{z}, \sigma, r$.

It remains to treat the exceptional half-lines in $\mathcal{L}$. Fix a half-line $L \in \mathcal{L}$ from Proposition 2.11, and consider pairs $(A, B) \in \mathcal{T} \cap L$, still with $A^2 + B^2$ sufficiently large and with $A \equiv \mathfrak{x} \pmod{M'}$ and $B \equiv \mathfrak{z} \pmod{M'}$. As above, Proposition 2.11 provides $\kappa \in \{0, 1\}$ and rationals E', F' such that

$$D(T) = D(U) + \frac{\kappa}{E'A + F'B}$$

for such (A, B) . We again restrict our attention to the case $\kappa = 1$ and clear denominators, after which we can immediately apply Lemma 2.13 instead of Lemma 2.12; this leads, as before, to finitely many progressions of the form $D(U) + 1/\text{Prog}(\alpha, \beta)$. This completes the proof of Theorem 2.1. $\square$

2.2.5 Remarks on the proof

We now briefly mention a few further pieces of information that can be extracted from the proof of Theorem 2.1; some of this additional structure will streamline the calculations in the remaining sections of the chapter.

Sectors and exceptional half-lines

In the proof above, the relative spectrum $\mathcal{S}_1(U)$ is obtained as the union of contributions from sectors and contributions from exceptional half-lines. One can show (as follows) that in fact the progressions in $\mathcal{S}_1(U)$ come either entirely from sectors or entirely from exceptional half-lines, according to whether U achieves its D -value at finitely or infinitely many points. This observation significantly shortens calculations in concrete examples.

First, suppose that U achieves its D -value at infinitely many points. (This is always the case, for instance, when U is of the form $U = U' \times (\mathbb{R}/\mathbb{Z})$.) Then there is some quadruple (i, j, ϵ, ℓ) such that $D(U_{i,j,\epsilon,\ell}) = D(U)$ and $D_{i,j,\epsilon,\ell}$ is equal to $D(U)$ on an interval of positive length. It follows that $D(T) = D(U)$ except when (A, B) lies on one of the exceptional half-lines for the quadruple (i, j, ϵ, ℓ) , so the entire relative spectrum $\mathcal{S}_1(U)$ comes from these half-lines. In particular, we do not have to consider contributions from the various sectors. For an example, see Section 2.4.1.

Second, suppose that U achieves its D -value at only finitely many points. We claim that the contribution to $\mathcal{S}_1(U)$ from exceptional half-lines is already contained in the contribution from sectors and hence can be ignored. It suffices to show that the first conclusion of Proposition 2.11 continues to hold on exceptional half-lines. Fix a choice of residues of A, B modulo M' as in the proof of that proposition. Let L be an exceptional half-line, and let Y'' denote the set of quadruples (i, j, ϵ, ℓ) for which L is exceptional; these are precisely the quadruples for which the corresponding value $q_{i,j,\epsilon,\ell} = |T_{i,j,\epsilon,\ell}|$ is constant on L . It follows that for each $(i, j, \epsilon, \ell) \in Y''$, the quantity $D(T_{i,j,\epsilon,\ell})$ is equal to some constant, say, $D(T_{i,j,\epsilon,\ell}) = D(U) + \eta_{i,j,\epsilon,\ell}$, for all $(A, B) \in L$. If $\eta_{i,j,\epsilon,\ell} = 0$ for some $(i, j, \epsilon, \ell) \in Y''$, then $D(T) = D(U)$ and L contributes only the single value $D(U)$ to $\mathcal{S}_1(U)$.

Now consider the case where $\eta_{i,j,\epsilon,\ell} > 0$ for every $(i, j, \epsilon, \ell) \in Y''$. For $(i, j, \epsilon, \ell) \in Y''$, the parameter $q = q_{i,j,\epsilon,\ell}$ (which is constant) may be too small for the formula (2.9) to hold; in particular, the formula (2.9) may erroneously predict that $D(T_{i,j,\epsilon,\ell}) = D(U) + \eta'_{i,j,\epsilon,\ell}$ for some $\eta'_{i,j,\epsilon,\ell} > 0$ that is different from $\eta_{i,j,\epsilon,\ell}$. We know (from the previous chapter) that $D(T) \rightarrow D(U)$ as $A^2 + B^2 \rightarrow \infty$. Thus, for $A^2 + B^2$ sufficiently large, the first minimum over Y in the proof of Proposition 2.11 must be achieved by some quadruple $(i, j, \epsilon, \ell) \in Y$ with $D(T_{i,j,\epsilon,\ell}) \rightarrow D(U)$. In particular, this minimum is achieved on $Y \setminus Y''$, and it makes no difference if the η 's are replaced by the η' 's on Y'' . The rest of the proof then goes through verbatim.

Other symmetries and reductions

In the proof of Proposition 2.11, we apply Lemma 2.9 (via Lemma 2.10) for each choice of $(i, j, \epsilon, \ell), \tau, \delta$, with $q := \delta(EA + FB)$ and $b := \delta a \ell / K$ for some parameters E, F, a, K depending on (i, j, ϵ, ℓ) and the

residues of A, B modulo M'' . Lemma 2.9 then provides formulas for the residues of R^-, R^+ modulo M as certain linear combinations of A, B (with coefficients depending on all of the parameters introduced so far). In these formulas, changing the sign δ changes the signs of q, b (as linear combinations of A, B), which has the effect of swapping R^-, R^+ . In particular, the sign δ does not affect whether or not $\min\{R^-, R^+\} = 0$ or, by extension, whether or not $D(T_{i,j,\epsilon,\ell}) = 0$. Thus, when we make our sector decomposition in the proof of Proposition 2.11, either $D(T) = D(U)$ on all sectors or $D(T) > D(U)$ on all sectors: If $D(T) = D(U)$ on some sector, then the $(i, j, \epsilon, \ell), \tau$ realizing this D -value (corresponding to $\min\{R^-, R^+\} = 0$) will also realize the same D -value on the rest of $\mathbb{R}_{\geq 0} \times \mathbb{R}$. This can be seen in Figure 2.8.

We conclude with a few minor remarks about conventions and computational redundancies. First, in the conclusion of Lemma 2.9, we can divide through all of R^-, R^+, xz by the common factor $\gcd(x, z)$; this reduces the length of some of our later numerical calculations. Second, in Section 2.2, we restricted (A, B) to the half-space $\{A \geq 0\}$, but any other half-space would work equally well; sometimes in examples it will be convenient to restrict to different half-spaces.

Third, we observe that since the function D on $(\mathbb{R}/\mathbb{Z})^n$ is an even function, its pushforward by each $\psi = \psi_{i,j,\epsilon}$ is also even. It follows that when we apply Lemmas 2.9 and 2.10 in the proof of Proposition 2.11, the critical values of τ come in pairs (except for 2-torsion points). More precisely, suppose that $U_{i,j,\epsilon}$ has K components. Then, in the notation of the proof of Proposition 2.11, we have $D_{i,j,\epsilon,\ell}(x) = D_{i,j,\epsilon,-\ell}(-x)$ (with the index $-\ell$ interpreted modulo K). Since the subgroup $T_{i,j,\epsilon}$ is also obviously invariant under the map $x \mapsto -x$, we see that $D_{i,j,\epsilon,\ell}(\tau_0) = D(U)$ if and only if $D_{i,j,\epsilon,-\ell}(-\tau_0) = D(U)$; in this case, the best approximations to these two minima by elements of $\psi(T_{i,j,\epsilon})$ make the same contribution to the expression of $D(T_{i,j,\epsilon})$ as a minimum over approximations to all critical points τ . It follows that in order to compute $D(T_{i,j,\epsilon})$ (and, by extension, $D(T)$), it suffices to consider one point from each pair described in the previous sentence. For example, we could consider only values of τ in $[0, 1/2]$; or we could consider only values of τ coming from $0 \leq \ell \leq k/2$ (and, for $\ell \in \{0, k/2\}$, look at only $\tau \in [0, 1/2]$). We will use this reduction in all of our computational examples below.

2.3 Identifying candidate subtori U

We now turn to our computational results Theorems 2.4 and 2.3, which characterize $\mathcal{S}_1(3) \cap (1/10, 1/6]$ and $\mathcal{S}_1(4) \cap (1/4, 1/2]$. The first step is identifying the relevant 2-dimensional subtori whose relative spectra contribute to these sets. In particular, we must find all proper subtori $U \subseteq (\mathbb{R}/\mathbb{Z})^3$ satisfying $\dim(U) = 2$ and $D(U) = 1/10$, and all proper subtori $U \subseteq (\mathbb{R}/\mathbb{Z})^4$ satisfying $\dim(U) = 2$ and $D(U) = 1/4$. We will begin with some general considerations that aid the identification of 2-dimensional subtori with D -value above a

given threshold.

2.3.1 General setup

We can reduce the length of our calculations by taking advantage of natural symmetries in the Lonely Runner Problem. The function D on $(\mathbb{R}/\mathbb{Z})^n$ is preserved by the automorphisms of $(\mathbb{R}/\mathbb{Z})^n$ given by negating and permuting coordinates. These automorphisms also permute the proper subtori of $(\mathbb{R}/\mathbb{Z})^n$, so, for the purposes of this section, it suffices to consider a single subtorus U in each orbit of this automorphism group; we will say that subtori in the same orbit are equivalent “up to symmetry”.

Suppose that $U \subseteq (\mathbb{R}/\mathbb{Z})^n$ is a 2-dimensional proper subtorus. Then there are nonparallel vectors $u = (u_1, \dots, u_n), v = (v_1, \dots, v_n) \in \mathbb{Z}^n$ such that $U = \langle u, v \rangle_{\mathbb{R}}$. The linear independence of u, v implies that there are indices $i < j$ such that $(u_i, u_j), (v_i, v_j) \in \mathbb{Z}^2$ are linearly independent; up to symmetry, we may assume that $(i, j) = (1, 2)$. Then, by replacing u, v with

$$(v_2 - v_1)u + (u_1 - u_2)v, \quad (-v_2 - v_1)u + (u_2 + u_1)v,$$

we can assume that $u_1 = u_2$ and $v_1 = -v_2$. For the remainder of this section, we will work with subtori U of the form

$$U = \langle (u_1, u_1, u_3, \dots, u_n), (v_1, -v_1, v_3, \dots, v_n) \rangle_{\mathbb{R}}, \quad (2.10)$$

where $u_1, u_3, \dots, u_n, v_1, v_3, \dots, v_n \in \mathbb{Z}$. We will split our analyses into several cases depending on which of these entries are zero.

2.3.2 Subtori $U \subset (\mathbb{R}/\mathbb{Z})^3$ with $D(U) = 1/10$.

In this section we identify all 2-dimensional proper subtori $U \subseteq (\mathbb{R}/\mathbb{Z})^3$ with $D(U) = 1/10$. We will show that, up to symmetry, the only such subtori are

$$\begin{aligned} U^3 &:= \langle (0, 1, 4), (1, 0, 0) \rangle_{\mathbb{R}}, & U^4 &:= \langle (0, 2, 3), (1, 0, 0) \rangle_{\mathbb{R}}, \\ U^5 &:= \langle (0, 1, 3), (1, 0, 1) \rangle_{\mathbb{R}}, & U^6 &:= \langle (0, 1, 2), (1, 1, 0) \rangle_{\mathbb{R}}. \end{aligned}$$

The arguments in this subsection generalize in a straightforward way to let one find all 2-dimensional proper subtori $U \subset (\mathbb{R}/\mathbb{Z})^3$ with $D(U) \geq M$ for any fixed $M > 0$. The only modification is that the list in Equation (2.11) below must be enlarged to contain all pairs $\{x, 2r+1-x\}$ with $4r+2 \leq 1/M$ and $1 \leq x \leq r$ ($r, x \in \mathbb{N}$).

For the remainder of this subsection, we assume that U is a subtorus of the form

$$U = \langle (a, a, b), (c, -c, d) \rangle_{\mathbb{R}}$$

(with $a, b, c, d \in \mathbb{Z}$) such that $D(U) = 1/10$. Without loss of generality we may assume that $\gcd(a, b) = \gcd(c, d) = 1$. We condition on which of a, b, c, d vanish.

No zero entries

Suppose first that a, b, c, d are all nonzero. Note that

$$1/10 = D(U) \leq D(\langle (a, a, b) \rangle_{\mathbb{R}}) = D(\langle (|a|, |b|) \rangle_{\mathbb{R}}).$$

Recall from [65, Theorem 4.1] that $D(\langle (|a|, |b|) \rangle_{\mathbb{R}})$ equals 0 if a, b are both odd and equals $1/2(a + b)$ if a, b have different parities. Hence we must have

$$\{|a|, |b|\} \in \{\{1, 2\}, \{1, 4\}, \{2, 3\}\}, \quad (2.11)$$

and likewise for $\{|c|, |d|\}$.

This leads to finitely many possibilities for U , and we can compute $D(U)$ for each. We conclude that, up to symmetry, U must be equal to one of the subtori

$$\begin{aligned} \langle (1, 1, 4), (-1, 1, 4) \rangle_{\mathbb{R}} &= \langle (0, 1, 4), (1, 0, 0) \rangle_{\mathbb{R}} = U^3, \\ \langle (2, 2, 3), (-2, 2, 3) \rangle_{\mathbb{R}} &= \langle (0, 2, 3), (1, 0, 0) \rangle_{\mathbb{R}} = U^4, \\ \langle (1, 1, 4), (1, -1, -2) \rangle_{\mathbb{R}} &= \langle (0, 1, 3), (1, 0, 1) \rangle_{\mathbb{R}} = U^5, \end{aligned}$$

as desired.

A single zero entry

Next, suppose that exactly one of a, b, c, d vanishes. If $a = 0$, then we have

$$D(U) = D(\langle (0, 0, 1), (c, -c, d) \rangle_{\mathbb{R}}) = D(\langle (c, -c) \rangle_{\mathbb{R}}) = 0,$$

which contradicts the assumption that $D(U) = 1/10$. So we may assume that $a \neq 0$, and likewise that $c \neq 0$. The remaining case (up to symmetry) is where $b = 0$. Then

$$U = \langle (1, 1, 0), (c, -c, d) \rangle_{\mathbb{R}},$$

where we know that $\{|c|, |d|\} \in \{\{1, 2\}, \{1, 4\}, \{2, 3\}\}$ (by the argument above). This again leads to only finitely many possibilities for U , and computing $D(U)$ for each leads to the conclusion that, up to symmetry, U must be equal to

$$\langle (1, 1, 0), (1, -1, -4) \rangle = \langle (0, 1, 2), (1, 1, 0) \rangle_{\mathbb{R}} = U^6.$$

This concludes the second subcase.

Two or more zero entries

Finally, suppose that at least two of a, b, c, d vanish. Since $(a, a, b), (c, -c, d)$ are both nonzero, we see that exactly two of a, b, c, d vanish. We cannot have $a = c = 0$ or $b = d = 0$ since we assumed that U is proper. Up to symmetry, it remains only to consider $a = d = 0$, in which case

$$D(U) = D(\langle (0, 0, 1), (1, -1, 0) \rangle_{\mathbb{R}}) = 0$$

gives a contradiction. This completes the proof.

2.3.3 Subtori $U \subset (\mathbb{R}/\mathbb{Z})^4$ with $D(U) = 1/4$

In this subsection, we will describe all proper subtori $U \subseteq (\mathbb{R}/\mathbb{Z})^4$ with $\dim(U) \geq 2$ and $D(U) = 1/4$. More precisely, we will show that, up to symmetry, the only such subtori are

$$U^1 := \langle (0, 1, 2, 3), (1, 0, 0, 0) \rangle_{\mathbb{R}}, \quad U^2 := \langle (1, 0, 1, 1), (1, 1, 0, 2) \rangle_{\mathbb{R}}.$$

We remark that U^1 already appeared in [65] but U^2 would have been more difficult to guess.

We do not have to consider subtori of dimension larger than 2 because the considerations of the previous chapter tell us that $\mathcal{S}_3(4) = \mathcal{S}_1(2)$ has maximum value $1/6$. For the remainder of this subsection, we assume that U is a subtorus of the form

$$U = \langle (a, a, b, c), (d, -d, e, f) \rangle_{\mathbb{R}}$$

(with $a, b, c, d, e, f \in \mathbb{Z}$) such that $D(U) = 1/4$.

We will repeatedly use the fact that if a, b, c are all nonzero, then

$$1/4 = D(U) \leq D(\langle (a, a, b, c) \rangle_{\mathbb{R}}) = D(\langle (|a|, |b|, |c|) \rangle_{\mathbb{R}}).$$

We may assume without loss of generality that $\gcd(|a|, |b|, |c|) = 1$, and then the characterization of equality for the Lonely Runner Conjecture with 3 runners [65, Section 5] implies that

$$\{|a|, |b|, |c|\} = \{1, 2, 3\};$$

this is useful because it completely determines one of the generators of U .

We also note that if $d = 0$ (i.e., U is contained in the subspace $\{x_1 = x_2\}$), then we have

$$1/4 = D(U) = D(\langle (a, a, b, c), (0, 0, e, f) \rangle_{\mathbb{R}}) = D(\langle (a, b, c), (0, e, f) \rangle_{\mathbb{R}}) \in \mathcal{S}_2(3).$$

But this is impossible since the largest value of $\mathcal{S}_2(3) = \mathcal{S}_1(2)$ is $1/6$, so we can always assume that $d \neq 0$.

The same argument lets us assume that $a \neq 0$ (and more generally that U is not contained in any subspace $\{x_i = \epsilon x_j\}$). We proceed by a case analysis as in the previous subsection.

No zero entries

Suppose that a, b, c, d, e, f are all nonzero. Then, after possibly scaling the generators of U , we have

$$\{|a|, |b|, |c|\} = \{|d|, |e|, |f|\} = \{1, 2, 3\}.$$

These finitely many possibilities lead, up to symmetry, to only

$$\langle (1, 1, 2, 3), (-1, 1, 2, 3) \rangle = \langle (0, 1, 2, 3), (1, 0, 0, 0) \rangle_{\mathbb{R}} = U^1.$$

A single zero entry

Suppose that exactly one of a, b, c, d, e, f vanishes. Up to symmetry, the only possibility is $b = 0$ (recall that $a, d \neq 0$). After possibly scaling the second generator of U , we may assume that $\{|d|, |e|, |f|\} = \{1, 2, 3\}$; fix one such choice of values for d, e, f . A change of basis, with an eye on the third and fourth coordinates, gives

$$U = \langle (cd + a(e - f), -cd + a(e - f), ce, ce), (cd - a(e + f), -cd - a(e + f), ce, -ce) \rangle_{\mathbb{R}}.$$

This change of basis has determinant $2ce$, which is nonzero by the assumption that $c, e \neq 0$. If all of the entries of these new generators are nonzero, then we land in the case that we already considered in the first subsection (with the third and fourth coordinates playing the role of the first and second coordinates). It remains to consider the case where some entry vanishes. Recall that d, e, f are already fixed. Since each entry is a linear combination of a, c , the vanishing of an entry determines the ratio of a, c , which in turn determines the vector $(a, a, 0, c)$ up to scaling. Computing $D(U)$ for each possibility, we find that, up to symmetry, U must equal

$$\langle (1, 1, 2, 3), (1, -1, 0, -1) \rangle_{\mathbb{R}} = \langle (1, 0, 1, 1), (1, 1, 0, 2) \rangle_{\mathbb{R}} = U^2.$$

Two or more zero entries

Suppose that at least two of a, b, c, d, e, f vanish, and recall from above that we may assume that $a, d \neq 0$. We cannot have $b = e = 0$ or $c = f = 0$ because of the assumption that U is proper. So exactly two of b, c, e, f vanish. If $b = c = 0$, then we have

$$U = \langle (a, a, 0, 0), (d, -d, e, f) \rangle_{\mathbb{R}} = \langle (1, 1, 0, 0), (d, -d, e, f) \rangle_{\mathbb{R}},$$

and we can take $\{|d|, |e|, |f|\} = \{1, 2, 3\}$; this leads to finitely many possibilities for U , none of which leads to any new possibilities for U besides U^1 and U^2 . The case $e = f = 0$ is identical.

It remains to consider the cases where $b = f = 0$ and where $c = e = 0$; these are equivalent up to symmetry, so we will consider only the former. Suppose that $b = f = 0$, so that

$$U = \langle (a, a, 0, c), (d, -d, e, 0) \rangle_{\mathbb{R}}.$$

Replacing these two generators with linear combinations to make the third and fourth coordinates more agreeable yields

$$U = \langle (cd + ae, -cd + ae, ce, ce), (cd - ae, -cd - ae, ce, -ce) \rangle_{\mathbb{R}}$$

(again this change of coordinates has determinant $2ce \neq 0$). If all of the entries of these new generators are nonzero, then we land in the case that we already considered in the previous subsection, with the third and fourth coordinates playing the role of the first and second coordinates.

It remains to consider the case where some entry vanishes, i.e., $cd = \pm ae$ (recall $ce \neq 0$). These cases are

again equivalent up to symmetry, so suppose that the former occurs and

$$U = \langle (a, a, 0, c), (a, -a, c, 0) \rangle_{\mathbb{R}}.$$

Another change of basis, this time focusing on the first and third coordinates, gives

$$U = \langle (ac, ac - 2a^2, ac, c^2 - ac), (ac, ac + 2a^2, -ac, ac + c^2) \rangle_{\mathbb{R}}$$

(this change of coordinates has determinant $-2ac \neq 0$). We are done if all of the entries of these new generators are nonzero, so it remains only to check the cases where some entry vanishes, viz., $c = \pm 2a$ or $c = \pm a$ (recall $a, c \neq 0$). Since we are free to rescale the generators of U , we may assume that $a = 1$. Then each of the four cases gives a single value for c ; this in turn determines the ratio $d/e = a/c$, and so we can take $d = a$ and $e = c$. Computing $D(U)$ in each case, we do not find any new possibilities for U besides U^1, U^2 . This completes the proof.

2.4 The top of the spectrum $\mathcal{S}_1(4)$

As an illustration of Theorem 2.1, we will now prove Theorem 2.3, which states that the set $\mathcal{S}_1(4) \cap (1/4, 1/2]$ has finite symmetric difference with the set $1/4 + 1/4 \text{Prog}(2, 3)$. In the previous section we showed that, up to symmetry,

$$U^1 = \langle (0, 1, 2, 3), (1, 0, 0, 0) \rangle_{\mathbb{R}} \quad \text{and} \quad U^2 = \langle (1, 0, 1, 1), (1, 1, 0, 2) \rangle_{\mathbb{R}}$$

are the only 2-dimensional subtori $U \subseteq (\mathbb{R}/\mathbb{Z})^4$ with $D(U) = 1/4$. Now (2.5) tells us that to prove Theorem 2.3, it remains only to use Theorem 2.1 to analyze $\mathcal{S}_1(U^1)$ and $\mathcal{S}_1(U^2)$ individually. Notice that these choices of generators for U^1, U^2 satisfy the condition (2.6).

2.4.1 Computing $\mathcal{S}_1(U^1)$

The objective of this subsection is to prove the following proposition.

Proposition 2.14. *The set $\mathcal{S}_1(U^1)$ has finite symmetric difference with the set $1/4 + 1/4 \text{Prog}(4, 5)$.*

We parameterize 1-dimensional subtori of U^1 as

$$T = \langle A(0, 1, 2, 3) + B(1, 0, 0, 0) \rangle_{\mathbb{R}} = \langle (B, A, 2A, 3A) \rangle_{\mathbb{R}},$$

for coprime integers A, B with $B \geq 0$ (see Section 2.2.5). We begin by computing the intersections $U_{i,j,\epsilon} =$

Figure 2.3: Intersections of U^1 with the subspaces $\{x_i = \epsilon x_j\}$. The last column gives (where applicable) the values of ℓ for which $D(U_{i,j,\epsilon,\ell}) = 1/4$ and the values of τ where this D -value is achieved; we write “ $\ell : S$ ” to indicate that $U_{i,j,\epsilon,\ell}$ achieves the D -value $1/4$ on the set S (and we omit ℓ when $U_{i,j,\epsilon}$ is connected). In the case of the antepenultimate row, the D -value $1/4$ is achieved by $\tau \in [1/4, 3/4]$ for both $\ell = 1$ and $\ell = 3$.

Subspace	$U_{i,j,\epsilon}$	$D(U_{i,j,\epsilon})$	τ achieving $1/4$
$\{x_1 = x_2\}$	$\langle(1, 1, 2, 3)\rangle_{\mathbb{R}}$	$1/4$	$\{1/4, 3/4\}$
$\{x_1 = -x_2\}$	$\langle(-1, 1, 2, 3)\rangle_{\mathbb{R}}$	$1/4$	$\{1/4, 3/4\}$
$\{x_1 = x_3\}$	$\langle(2, 1, 2, 3)\rangle_{\mathbb{R}}$	$1/4$	$\{1/4, 3/4\}$
$\{x_1 = -x_3\}$	$\langle(-2, 1, 2, 3)\rangle_{\mathbb{R}}$	$1/4$	$\{1/4, 3/4\}$
$\{x_1 = x_4\}$	$\langle(3, 1, 2, 3)\rangle_{\mathbb{R}}$	$1/4$	$\{1/4, 3/4\}$
$\{x_1 = -x_4\}$	$\langle(-3, 1, 2, 3)\rangle_{\mathbb{R}}$	$1/4$	$\{1/4, 3/4\}$
$\{x_2 = x_3\}$	$\langle(1, 0, 0, 0)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_2 = -x_3\}$	$\cup_{\ell=0}^2 \langle(1, 0, 0, 0)\rangle_{\mathbb{R}} + (0, \ell/3, 2\ell/3, 0)$	$1/2$	—
$\{x_2 = x_4\}$	$\cup_{\ell=0}^1 \langle(1, 0, 0, 0)\rangle_{\mathbb{R}} + (0, \ell/2, 0, \ell/2)$	$1/2$	—
$\{x_2 = -x_4\}$	$\cup_{\ell=0}^3 \langle(1, 0, 0, 0)\rangle_{\mathbb{R}} + (0, \ell/4, 2\ell/4, 3\ell/4)$	$1/4$	$1, 3 : [1/4, 3/4]$
$\{x_3 = x_4\}$	$\langle(1, 0, 0, 0)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_3 = -x_4\}$	$\cup_{\ell=0}^4 \langle(1, 0, 0, 0)\rangle_{\mathbb{R}} + (0, \ell/5, 2\ell/5, 3\ell/5)$	$3/10$	—

$U^1 \cap \{x_i = \epsilon x_j\}$; see Figure 2.3.

Carrying out the reductions outlined in Section 2.2.5, we quickly dispose of all pairs (A, B) except for those with $|A| = 1$ and $B \equiv 0 \pmod{4}$. The key idea is to leverage the fact that the restriction of D to $U_{2,4,-,1}$ and $U_{2,4,-,3}$ equals $1/4$ on intervals of positive length; we will find that T inevitably intersects these intervals unless (A, B) is of the special form described in the previous sentence.

Consider the isomorphism $\psi = \psi_{2,4,-} : \langle(0, 1, 2, 3), (1, 0, 0, 0)\rangle_{\mathbb{R}} \rightarrow \mathbb{R}^2$ given by $\psi(0, 1, 2, 3) = (0, 1)$ and $\psi(1, 0, 0, 0) = (1, 0)$; since $\psi^{-1}(\langle(1, 0), (0, 1)\rangle_{\mathbb{Z}}) = \langle(0, 1, 2, 3), (1, 0, 0, 0)\rangle_{\mathbb{R}} \cap \mathbb{Z}^4$, the map ψ descends to an isomorphism $U^1 \rightarrow (\mathbb{R}/\mathbb{Z})^2$. In these new coordinates, we compute that

$$\psi(U_{2,4,-}) = (\mathbb{R}/\mathbb{Z}) \times \langle 1/4 \rangle_{\mathbb{Z}}$$

and

$$\psi(T) = \psi(\langle A(0, 1, 2, 3) + B(1, 0, 0, 0) \rangle_{\mathbb{R}}) = \langle (B, A) \rangle_{\mathbb{R}}.$$

One can check that $D_{2,4,-}$ assumes the constant value $1/4$ on the intervals $[1/4, 3/4] \times \{1/4, 3/4\}$. As described in Proposition 2.8, the subgroup

$$\psi(T_{2,4,-}) = \langle (B, A) \rangle_{\mathbb{R}} \cap ((\mathbb{R}/\mathbb{Z}) \times \langle 1/4 \rangle_{\mathbb{Z}}) = \bigcup_{\ell=0}^3 \left(\left(\frac{B\ell}{4A} + \left\langle \frac{1}{|A|} \right\rangle_{\mathbb{Z}} \right) \times \left\{ \frac{\ell}{4} \right\} \right)$$

intersects each coset $(\mathbb{R}/\mathbb{Z}) \times \{\ell/4\}$ in exactly $|A|$ equally-spaced points (since A, B are coprime). In particular, if $|A| \geq 2$, then $\psi(T_{2,4,-})$ intersects (both of) the aforementioned intervals where D assumes the

constant value $1/4$, which implies that $D(T) \leq 1/4$; since of course $D(T) \geq D(U^1) = 1/4$, we conclude that $D(T) = 1/4$, and so all such pairs (A, B) contribute only the value $1/4$ to the set $\mathcal{S}_1(U^1)$. Henceforth, we will restrict our attention to the case $|A| = 1$. In the language of Section 2.2, the pairs $(1, B)$ and $(-1, B)$ lie on two exceptional half-lines. Now, when $B \not\equiv 0 \pmod{4}$, we have

$$\psi(T_{2,4,-,1}) = \{(B/4A, 1/4)\} \in [1/4, 3/4] \times \{1/4\}, \quad (2.12)$$

which implies as above that $D(T) = 1/4$. Thus it remains only to treat the case $B \equiv 0 \pmod{4}$.³

Assume that $|A| = 1$, that $B \equiv 0 \pmod{4}$, and that $B \geq 0$. The next step is computing $D(T_{i,j,\epsilon,\ell})$ for each quadruple (i, j, ϵ, ℓ) with $D(U_{i,j,\epsilon,\ell}) = 1/4$. It follows from (2.12) that $D(T_{2,4,-,1}) = 1/2$, and likewise $D(T_{2,4,-,3}) = 1/2$. For each remaining (i, j, ϵ, ℓ) under consideration, we have $\ell = 0$ and the critical values of τ are $1/4, 3/4$. As described in Section 2.2.5, it suffices to analyze the behavior around $\tau = 1/4$.

As an example, let us write down some details of the calculation for $(i, j, \epsilon, \ell) = (1, 2, +, 0)$; the other five cases are completely analogous. This time, our isomorphism $\psi = \psi_{1,2,+}$ is given by $\psi(1, 1, 2, 3) = (1, 0)$ and $\psi(1, 0, 0, 0) = (0, 1)$, so that

$$\psi(T) = \psi(\langle A(0, 1, 2, 3) + B(1, 0, 0, 0) \rangle_{\mathbb{R}}) = \psi(\langle A(1, 1, 2, 3) + (B - A)(1, 0, 0, 0) \rangle_{\mathbb{R}}) = \langle (A, B - A) \rangle_{\mathbb{R}}.$$

Since $\psi(U_{1,2,+,0}) = (\mathbb{R}/\mathbb{Z}) \times \{0\}$, we have

$$\psi(T_{1,2,+,0}) = \langle 1/|B - A| \rangle_{\mathbb{Z}} \times \{0\}.$$

Notice that $B - A \geq 4 - 1 > 0$ except for the single case $(A, B) = (\pm 1, 0)$, which we can safely ignore. Then

$$\text{Approx}^+(1/4; 0, B - A) = \frac{\lceil (B - A)/4 \rceil}{B - A} - \frac{1}{4} = \frac{R^+}{4(B - A)},$$

where $R^+ \in \{0, 1, 2, 3\}$ satisfies $R^+ \equiv A - B \pmod{4}$. Likewise, $\text{Approx}^-(1/4; 0, B - A) = R^-/4(B - A)$, where $R^- \in \{0, 1, 2, 3\}$ satisfies $R^- \equiv B - A \pmod{4}$. The function $D_{1,2,+,0}(t) = D(t(1, 1, 2, 3))$ assumes its minimum value at $\tau = 1/4$ (and at $3/4$). By computing

$$D_{1,2,+,0}(1/4 + \varepsilon) = D((1/4 + \varepsilon, 1/4 + \varepsilon, 1/2 + 2\varepsilon, 3/4 + 3\varepsilon)) = 1/4 + \max\{-\varepsilon, 3\varepsilon\},$$

³We remark that the reductions in this paragraph can also be explained in terms of runners, using the language of “pre-jumps” (see [12]). Consider runners with speeds $B, A, 2A, 3A$. At times of the form $1/(4A) + s/A$ (for $s \in \mathbb{Z}$), the runners are at the positions $B/4A + Bs/A, 1/4, 1/2, 3/4$; the point is that only the first of these positions actually changes with s , since the pre-jump by $1/A$ fixes the other three positions. The first position ranges over the set $B/4A + \langle 1/|A| \rangle_{\mathbb{Z}}$ as s varies; in particular, it lies in $[1/4, 3/4]$ for some s unless $|A| = 1$ and $B \equiv 0 \pmod{4}$.

Figure 2.4: Approximations to $\tau = 1/4$ for U^1 in the case $A = 1$, $B > 0$, $B \equiv 0 \pmod{4}$. In each case, the values of R^+ , R^- were calculated as the unique values in $\{0, 1, 2, 3\}$ congruent modulo 4 to certain integer linear combinations of A, B .

i, j, ϵ, ℓ	$\psi_{i,j,\epsilon}(T_{i,j,\epsilon,\ell})$	λ^+	λ^-	R^+	R^-	Approx $^+$	Approx $^-$
1, 2, +, 0	$\langle \frac{1}{ B-A } \rangle_{\mathbb{Z}} \times \{0\}$	3	1	$A - B \rightsquigarrow 1$	$B - A \rightsquigarrow 3$	$\frac{1}{4(B-1)}$	$\frac{3}{4(B-1)}$
1, 2, -, 0	$\langle \frac{1}{ B+A } \rangle_{\mathbb{Z}} \times \{0\}$	3	1	$-B - A \rightsquigarrow 3$	$A + B \rightsquigarrow 1$	$\frac{3}{4(B+1)}$	$\frac{1}{4(B+1)}$
1, 3, +, 0	$\langle \frac{1}{ B-2A } \rangle_{\mathbb{Z}} \times \{0\}$	3	1	$2A - B \rightsquigarrow 2$	$B - 2A \rightsquigarrow 2$	$\frac{2}{4(B-2)}$	$\frac{2}{4(B-2)}$
1, 3, -, 0	$\langle \frac{1}{ B+2A } \rangle_{\mathbb{Z}} \times \{0\}$	3	1	$-B - 2A \rightsquigarrow 2$	$2A + B \rightsquigarrow 2$	$\frac{2}{4(B+2)}$	$\frac{2}{4(B+2)}$
1, 4, +, 0	$\langle \frac{1}{ B-3A } \rangle_{\mathbb{Z}} \times \{0\}$	3	1	$3A - B \rightsquigarrow 3$	$B - 3A \rightsquigarrow 1$	$\frac{3}{4(B-3)}$	$\frac{1}{4(B-3)}$
1, 4, -, 0	$\langle \frac{1}{ B+3A } \rangle_{\mathbb{Z}} \times \{0\}$	3	1	$-B - 3A \rightsquigarrow 1$	$3A + B \rightsquigarrow 3$	$\frac{1}{4(B+3)}$	$\frac{3}{4(B+3)}$

we see that $\lambda^+ = 3$ and $\lambda^- = 1$ determine the slopes of the pieces of $D_{1,2,+,0}$ directly adjacent to the point $\tau = 1/4$. (Recall that the slope of the piece to the left is defined to be $-\lambda^-$.) Figure 2.4 shows the output of these calculations, in the case $A = 1$, for all of the relevant quadruples (i, j, ϵ, ℓ) .

When B is sufficiently large, we have

$$D(T_{i,j,\epsilon,\ell}) = 1/4 + \min\{\lambda^+ \text{Approx}^+, \lambda^- \text{Approx}^-\}$$

for each quadruple (i, j, ϵ, ℓ) from Figure 2.4, and then $D(T)$ is the minimum of these six quantities. The result of these comparisons is that $D(T) = 1/4 + 1/4(B+1)$ (coming from $T_{1,2,-,0}$). Substituting $B = 4s$ —note that A, B are automatically coprime since $A = 1$ —we obtain the progression $1/4 + 1/4 \text{Prog}(4, 5)$. The analogous calculation for the half-line with $A = -1$ yields the same progression. Hence $\mathcal{S}_1(U^1)$ has finite symmetric difference with $1/4 + 1/4 \text{Prog}(4, 5)$, as desired.

2.4.2 Computing $\mathcal{S}_1(U^2)$

Our calculation of $\mathcal{S}_1(U^2)$ is more involved and can be summarized as follows; Theorem 2.3 is then a quick consequence since $\text{Prog}(4, 5) \subseteq \text{Prog}(2, 3)$.

Proposition 2.15. *The set $\mathcal{S}_1(U^2)$ has finite symmetric difference with the set $1/4 + 1/4 \text{Prog}(2, 3)$.*

We parameterize 1-dimensional subtori of U^2 as

$$T = \langle A(1, 0, 1, 1) + B(1, 1, 0, 2) \rangle_{\mathbb{R}} = \langle (A+B, B, A, A+2B) \rangle_{\mathbb{R}},$$

for coprime A and B with $A \geq 0$. As before, we intersect U^2 with the subspaces $\{x_i = \epsilon x_j\}$ to get the subgroups $U_{i,j,\epsilon}$, and then we compute the values $D(U_{i,j,\epsilon})$; see Figure 2.5. In contrast to the example in the previous subsection, the D -value $1/4$ is achieved by only finitely many points of U^2 (as opposed to intervals

Figure 2.5: Intersections of U^2 with the subspaces $\{x_i = \epsilon x_j\}$.

Subspace	$U_{i,j,\epsilon}$	$D(U_{i,j,\epsilon})$	τ achieving $1/4$
$\{x_1 = x_2\}$	$\langle(1, 1, 0, 2)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_1 = -x_2\}$	$\langle(1, -1, 2, 0)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_1 = x_3\}$	$\langle(1, 0, 1, 1)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_1 = -x_3\}$	$\langle(1, 2, -1, 3)\rangle_{\mathbb{R}}$	$1/4$	$\{1/4, 3/4\}$
$\{x_1 = x_4\}$	$\langle(1, 0, 1, 1)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_1 = -x_4\}$	$\langle(-1, 2, -3, 1)\rangle_{\mathbb{R}}$	$1/4$	$\{1/4, 3/4\}$
$\{x_2 = x_3\}$	$\langle(2, 1, 1, 3)\rangle_{\mathbb{R}}$	$1/4$	$\{1/4, 3/4\}$
$\{x_2 = -x_3\}$	$\langle(0, 1, -1, 1)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_2 = x_4\}$	$\langle(0, 1, -1, 1)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_2 = -x_4\}$	$\langle(2, -1, 3, 1)\rangle_{\mathbb{R}}$	$1/4$	$\{1/4, 3/4\}$
$\{x_3 = x_4\}$	$\cup_{\ell=0}^1 \langle(1, 0, 1, 1)\rangle_{\mathbb{R}} + (\ell/2, \ell/2, 0, 0)$	$1/4$	$1 : \{1/4, 3/4\}$
$\{x_3 = -x_4\}$	$\cup_{\ell=0}^1 \langle(0, 1, -1, 1)\rangle_{\mathbb{R}} + (\ell/2, 0, \ell/2, \ell/2)$	$1/4$	$1 : \{1/4, 3/4\}$

of positive length). The considerations of Section 2.2.5 tell us that it will suffice to study only sectors, rather than exceptional half-lines (as in the previous subsection). By Section 2.2.5, it suffices to work with $\tau = 1/4$.

Let us write out some of the details for approximating $\tau = 1/4$ in $U_{3,4,+,1}$; the other cases are similar to this case or to what we demonstrated in the previous subsection. The isomorphism $\psi = \psi_{3,4,+}$ is given by $\psi(1, 0, 1, 1) = (1, 0)$ and $\psi(1, 1, 0, 2) = (0, 1)$, and $\psi(T) = \langle(A, B)\rangle_{\mathbb{R}}$. Following Proposition 2.8, we have

$$\psi(T_{3,4,+,1}) = \left(\frac{a}{2|B|} + \left\langle \frac{1}{|B|} \right\rangle_{\mathbb{Z}} \right) \times \left\{ \frac{1}{2} \right\},$$

where $a \equiv A \pmod{2}$. (Here a is independent of the sign of B since $A \equiv -A \pmod{2}$.) Lemma 2.9 tells us that the best approximations of $\tau = 1/4$ by elements of $\frac{a}{2|B|} + \left\langle \frac{1}{|B|} \right\rangle_{\mathbb{Z}}$ have errors

$$\text{Approx}^- \left(\frac{1}{4}, \frac{a}{2}; |B| \right) = \frac{R^-}{4|B|}, \quad R^- \equiv |B| - 2A \pmod{4},$$

$$\text{Approx}^+ \left(\frac{1}{4}, \frac{a}{2}; |B| \right) = \frac{R^+}{4|B|}, \quad R^+ \equiv 2A - |B| \pmod{4},$$

where $R^+, R^- \in \{0, 1, 2, 3\}$. Finally, to compute the slopes λ^+ and λ^- , we write $D_{3,4,+,1}(t) = D(t(1, 0, 1, 1) + (1/2, 1/2, 0, 0))$ and compute

$$D_{3,4,+,1}(1/4 + \varepsilon) = D((3/4 + \varepsilon, 1/2, 1/4 + \varepsilon, 1/4 + \varepsilon)) = 1/4 + \max\{-\varepsilon, \varepsilon\}$$

for small ε . Thus $\lambda^+ = \lambda^- = 1$. See Figure 2.6 for these calculations for all quadruples (i, j, ϵ, ℓ) .

For each (i, j, ϵ, ℓ) , we partition $\mathbb{R}_{\geq 0} \times \mathbb{R}$ into two sectors, corresponding to the sign of the expression in the absolute value in the second column of Figure 2.6; we use $\delta_{i,j,\epsilon,\ell} \in \{+, -\}$ to denote this sign. The

Figure 2.6: Approximations to $\tau = 1/4$ for U^2 . In the fifth row, $a \equiv A \pmod{2}$, and in the sixth row, $a \equiv B \pmod{2}$. The values of R^+, R^- lie in $\{0, 1, 2, 3\}$. The headings for the last four columns indicate the sign of δ in parentheses.

i, j, ϵ, ℓ	$\psi(T_{i,j,\epsilon,\ell})$	λ^+	λ^-	$R^+ (+)$	$R^- (+)$	$R^+ (-)$	$R^- (-)$
1, 3, -, 0	$\langle \frac{1}{ 2A+B } \rangle_{\mathbb{Z}} \times \{0\}$	3	1	$-2A - B$	$2A + B$	$2A + B$	$-2A - B$
1, 4, -, 0	$\langle \frac{1}{ 2A+3B } \rangle_{\mathbb{Z}} \times \{0\}$	3	1	$2A + B$	$2A + 3B$	$2A + 3B$	$2A + B$
2, 3, +, 0	$\langle \frac{1}{ B-A } \rangle_{\mathbb{Z}} \times \{0\}$	3	1	$A - B$	$B - A$	$B - A$	$A - B$
2, 4, -, 0	$\langle \frac{1}{ A+3B } \rangle_{\mathbb{Z}} \times \{0\}$	3	1	$B + 3A$	$A + 3B$	$A + 3B$	$B + 3A$
3, 4, +, 1	$(\frac{a}{2 B } + \langle \frac{1}{ B } \rangle_{\mathbb{Z}}) \times \{\frac{1}{2}\}$	1	1	$2A - B$	$B - 2A$	$B - 2A$	$2A - B$
3, 4, -, 1	$(\frac{a}{2 A+B } + \langle \frac{1}{ A+B } \rangle_{\mathbb{Z}}) \times \{\frac{1}{2}\}$	1	1	$B - A$	$A - B$	$A - B$	$B - A$

common refinement of these partitions gives a partition of $\mathbb{R}_{\geq 0} \times \mathbb{R}$ into several sectors, each of which must be analyzed separately.

For example, consider the sector $\{(x, y) : 0 \leq x \leq y\}$, corresponding to the case where all of the $\delta_{i,j,\epsilon,\ell}$'s are +. For each choice of residues $(A, B) \pmod{4}$ and each quadruple (i, j, ϵ, ℓ) , we have

$$D(T_{i,j,\epsilon,\ell}) = 1/4 + \min\{\lambda^+ \text{Approx}^+, \lambda^- \text{Approx}^-\},$$

and this quantity equals $1/4 + \gamma/q$, where $q = |T_{i,j,\epsilon,\ell}|$ is a linear combination of A, B . Figure 2.7 shows the outcome of this calculation for all six quadruples (i, j, ϵ, ℓ) and the various choices of residues $(A, B) \pmod{4}$ (all in the sector $\{(x, y) : 0 \leq x \leq y\}$).

For each choice of residues $(A, B) \pmod{4}$, we now compare the quantities $D(T_{i,j,\epsilon,\ell})$, which will make us further partition our sector $\{(x, y) : 0 \leq x \leq y\}$. Here we will illustrate this procedure in the case $(A, B) \equiv (1, 3) \pmod{4}$. We wish to find the minimum among the quantities

$$\frac{1}{2A+B}, \quad \frac{3}{2A+3B}, \quad \frac{2}{B-A}, \quad \frac{2}{A+3B}, \quad \frac{1}{B}, \quad \frac{2}{A+B}.$$

Since $0 \leq A \leq B$, the first of these quantities is always smaller than or equal to the second, third, fifth, and sixth. It remains to compare $1/(2A+B)$ and $2/(A+3B)$: The former is smaller when $B < 3A$, and the latter is smaller when $B > 3A$, so we accordingly use the slope 3-ray through the origin to subdivide the sector $\{(x, y) : 0 \leq x \leq y\}$ into two sectors. See Figure 2.8 for the outcomes of these calculations for other choices of the $\delta_{i,j,\epsilon,\ell}$'s and pairs of residues $(A, B) \pmod{4}$, and see Figure 2.9 for a visual representation of our running example $(A, B) \equiv (1, 3) \pmod{4}$.

For each entry in the third column of Figure 2.8, we must determine which values are assumed by the denominator as (A, B) ranges over pairs of coprime integers satisfying the modular conditions from the

Figure 2.7: The values of 4γ in the sector $\{0 \leq A \leq B\}$ for U^2 . Each grid represents a quadruple (i, j, ϵ, ℓ) . For each, we show how γ depends on $(A, B) \pmod{4}$; note that we have excluded the cases where A, B are both even, due to our standing assumption that A, B are coprime. For the reader's convenience, we have reiterated the value of $q = |T_{i,j,\epsilon,\ell}|$ in each case.

$B \backslash A$	0	1	2	3
0	-	2	-	2
1	1	3	1	3
2	-	0	-	0
3	3	1	3	1

(a) $(1, 3, -, 0)$ and $q = 2A + B$.

$B \backslash A$	0	1	2	3
0	-	2	-	2
1	3	1	3	1
2	-	0	-	0
3	1	3	1	3

(b) $(1, 4, -, 0)$ and $q = 2A + 3B$.

$B \backslash A$	0	1	2	3
0	-	3	-	1
1	1	0	3	2
2	-	1	-	3
3	3	2	1	0

(c) $(2, 3, +, 0)$ and $q = B - A$.

$B \backslash A$	0	1	2	3
0	-	1	-	3
1	3	0	1	2
2	-	3	-	1
3	1	2	3	0

(d) $(2, 4, -, 0)$ and $q = A + 3B$.

$B \backslash A$	0	1	2	3
0	-	2	-	2
1	1	1	1	1
2	-	0	-	0
3	1	1	1	1

(e) $(3, 4, +, 1)$ and $q = B$.

$B \backslash A$	0	1	2	3
0	-	1	-	1
1	1	0	1	2
2	-	1	-	1
3	1	2	1	0

(f) $(3, 4, -, 1)$ and $q = A + B$.

first column and the linear inequalities from the second column. To illustrate this step, consider the sector $\{(x, y) : 0 \leq 3x \leq y\}$ with $(A, B) \equiv (1, 3) \pmod{4}$, which corresponds to the D -value offset $2/(A + 3B)$. Let us parameterize A, B as $A = 4r + 1$, $B = 4s + 3$, so that $A + 3B = 4r + 12s + 10$. This last expression is always equivalent to 2 modulo 4, so we wish to determine the set of integers t such that there are valid choices of r, s (in the sense that A, B are coprime and satisfy $B \geq 3A \geq 0$) with

$$A + 3B = 4r + 12s + 10 = 4t + 2,$$

i.e., $r + 3s + 2 = t$. Such pairs (r, s) can in turn be parameterized as

$$(r, s) = (t - 2 + 3x, -x)$$

for x an integer in a suitable interval (of length growing with t), and this leads to

$$(A, B) = (4t - 7 + 12x, -4x + 3).$$

Figure 2.8: D -values for sectors for U^2 . Each row corresponds to a pair of residues $(A, B) \pmod{4}$, as recorded in the first column. The second column gives the slopes of the dividing rays for the corresponding sector decomposition of $\mathbb{R}_{\geq 0} \times \mathbb{R}$. In the third column we record (proceeding through the sectors clockwise) the corresponding values of the quantity $4(D(T) - 1/4)$, which we term the “offset” for brevity. The last column shows the values achieved by the offsets for coprime A, B .

(A, B)	Slopes of dividing rays	Offsets	Scaled progressions
$(0, 1)$	$-3/4$	$\frac{1}{2A+B}, \frac{1}{-A-3B}$	$\frac{1}{4t+1}, \frac{1}{4t+1}$
$(0, 3)$	$-1/4$	$\frac{1}{2A+3B}, \frac{1}{A-B}$	$\frac{1}{4t+1}, \frac{1}{4t+1}$
$(1, 0)$	$0, -4$	$\frac{1}{A+3B}, \frac{1}{A-B}, \frac{2}{-2A-3B}$	$\frac{1}{4t+1}, \frac{1}{4t+1}, \frac{1}{2t+1}$
$(1, 1)$	(none)	0	—
$(1, 2)$	(none)	0	—
$(1, 3)$	$3, -1$	$\frac{2}{A+3B}, \frac{1}{2A+B}, \frac{1}{-2A-3B}$	$\frac{1}{2t+1}, \frac{1}{4t+1}, \frac{1}{4t+1}$
$(2, 1)$	$1/2, -1/2, -3/2$	$\frac{1}{A+3B}, \frac{1}{2A+B}, \frac{1}{A-B}, \frac{1}{-2A-3B}$	$\frac{1}{4t+1}, \frac{1}{4t+1}, \frac{1}{4t+1}, \frac{1}{4t+1}$
$(2, 3)$	$-1/2$	$\frac{1}{2A+3B}, \frac{1}{-A-3B}$	$\frac{1}{4t+1}, \frac{1}{4t+1}$
$(3, 0)$	$0, -4/7$	$\frac{2}{2A+3B}, \frac{2}{2A+B}, \frac{1}{-A-3B}$	$\frac{1}{2t+1}, \frac{1}{2t+1}, \frac{1}{4t+1}$
$(3, 1)$	$-3/7, -1$	$\frac{1}{2A+3B}, \frac{2}{A-B}, \frac{2}{-A-3B}$	$\frac{1}{4t+1}, \frac{1}{2t+1}, \frac{1}{2t+1}$
$(3, 2)$	(none)	0	—
$(3, 3)$	(none)	0	—




Figure 2.9: This figure illustrates the sector decomposition for U^2 when $(A, B) \equiv (1, 3) \pmod{4}$. Each sector is labeled with the progression of its corresponding offset (see Figure 2.8).

Lemma 2.12 tells us that since $\gcd(4t - 7, 12, -4, 3) = 1$, there are many choices of x making (A, B) coprime (with $B \geq 3A \geq 0$) when t is large. In particular, every sufficiently large value of t is achieved (even if one insists on throwing out exceptional half-lines, which is unnecessary in light of Section 2.2.5), and the expression $2/(A + 3B)$ assumes the values

$$2/(4t + 2) = 1/(2t + 1).$$

The last column of Figure 2.8 shows the outcome of the analogous calculations for the other entries in the third column of the table. Since $\{4t + 1 : t \in \mathbb{N}\} \subseteq \{2t + 1 : t \in \mathbb{N}\}$, we conclude that $\mathcal{S}_1(U^2)$ has finite symmetric difference with the single progression $1/4 + 1/4 \text{Prog}(2, 3)$. This proves Proposition 2.15.

Let us return to the progression (2.2) found by Fan and Sun [39]. This family of T 's corresponds to the case $A + 4s + 3$, $B = 8$ in the setting currently under consideration. In the $(A, B) \equiv (3, 0) \pmod{4}$ row of Figure 2.8, these choices of (A, B) fall in the first sector (and do not coincide with an exceptional half-line). Since A, B are coprime, we see that indeed

$$D(T) = \frac{1}{4} + \frac{2}{4(2A + 3B)} = \frac{1}{4} + \frac{2}{4(2(4s + 3) + 3(8))} = \frac{1}{4} + \frac{1}{16s + 60};$$

this provides a systematic way of carrying out the calculation of (2.2).

2.5 The spectrum $\mathcal{S}_1(3)$ to the second accumulation point

We will now prove Theorem 2.4, which states that $\mathcal{S}_1(3) \cap (1/10, 1/6]$ has finite symmetric difference with the union of $1/10 + 4/5 \text{Prog}(5, 7)$ and $1/10 + 3/5 \text{Prog}(5, 9)$. Recall that [65] determined $\mathcal{S}_1(3)$ up to the first accumulation point, viz., $\mathcal{S}_1(3) \cap (1/6, 1/2] = 1/6 + 1/6 \text{Prog}(6, 7)$. Theorem 2.4 determines $\mathcal{S}_1(3)$ up to the second accumulation point $(1/10)$, up to finitely many exceptions. We reiterate (see the comment at the beginning of Section 2.3.2) that our method here would let us compute $\mathcal{S}_1(3) \cap (\varepsilon, 1/2]$, up to finitely many exceptions, for any $\varepsilon > 0$: We can compute $\mathcal{S}_1(3)$ up to as many accumulation points as we want, and these accumulation points approach 0.

In Section 2.3.2, we showed that, up to symmetry, the only 2-dimensional proper subtori $U \subseteq (\mathbb{R}/\mathbb{Z})^3$ with $D(U) = 1/10$ are

$$\begin{aligned} U^3 &= \langle (0, 1, 4), (1, 0, 0) \rangle_{\mathbb{R}}, & U^4 &= \langle (0, 2, 3), (1, 0, 0) \rangle_{\mathbb{R}}, \\ U^5 &= \langle (0, 1, 3), (1, 0, 1) \rangle_{\mathbb{R}}, & U^6 &= \langle (0, 1, 2), (1, 1, 0) \rangle_{\mathbb{R}}. \end{aligned}$$

Figure 2.10: Intersections of U^3 with the subspaces $\{x_i = \epsilon x_j\}$.

Subspace	$U_{i,j,\epsilon}$	$D(U_{i,j,\epsilon})$	τ achieving $D(U_{i,j,\epsilon})$
$\{x_1 = x_2\}$	$\langle(1, 1, 4)\rangle_{\mathbb{R}}$	$1/10$	$\{2/5, 3/5\}$
$\{x_1 = -x_2\}$	$\langle(-1, 1, 4)\rangle_{\mathbb{R}}$	$1/10$	$\{2/5, 3/5\}$
$\{x_1 = x_3\}$	$\langle(4, 1, 4)\rangle_{\mathbb{R}}$	$1/10$	$\{2/5, 3/5\}$
$\{x_1 = -x_3\}$	$\langle(-4, 1, 4)\rangle_{\mathbb{R}}$	$1/10$	$\{2/5, 3/5\}$
$\{x_2 = x_3\}$	$\cup_{\ell=0}^2 \langle(1, 0, 0)\rangle_{\mathbb{R}} + (0, \ell/3, \ell/3)$	$1/6$	—
$\{x_2 = -x_3\}$	$\cup_{\ell=0}^4 \langle(1, 0, 0)\rangle_{\mathbb{R}} + (0, \ell/5, 4\ell/5)$	$1/10$	$2, 3 : [2/5, 3/5]$

Notice that each of these choices of generators satisfies the condition (2.6).

All that remains for the proof of Theorem 2.4 is computing $S_1(U^i)$ for $i \in \{3, 4, 5, 6\}$ individually. The argument is similar to the argument in Section 2.4. In each of the first two subsections, we will be able to quickly dispose of all but finitely many exceptional lines; in each of the last two subsections, we will deal with full sector decompositions of $\mathbb{R}_{\geq 0} \times \mathbb{R}$. One topically new (but unimportant) feature is that we have different values of τ for different tuples (i, j, ϵ, ℓ) . We will abbreviate parts of the exposition that are similar to Section 2.4.

2.5.1 Computing $\mathcal{S}_1(U^3)$

We parametrize 1-dimensional subtori of U^3 as

$$T = \langle A(0, 1, 4) + B(1, 0, 0) \rangle_{\mathbb{R}}$$

for coprime integers A, B with $B \geq 0$. We compute the intersections $U_{i,j,\epsilon} = U^3 \cap \{x_i = \epsilon x_j\}$ in Figure 2.10. Note that on $U_{2,3,-}$, the D -value $1/10$ is achieved on intervals of positive length; we will use this to restrict our attention to a small number of exceptional half-lines.

Consider the isomorphism $\psi = \psi_{2,3,-}$ given by $\psi((1, 0, 0)) = (1, 0)$ and $\psi((0, 1, 4)) = (0, 1)$. Then Proposition 2.8 gives

$$\psi(T_{2,3,-}) = \bigcup_{\ell=0}^4 \left(\left(\frac{a\ell}{(5|A|)} + \left\langle \frac{1}{|A|} \right\rangle_{\mathbb{Z}} \right) \times \left\{ \frac{\ell}{5} \right\} \right),$$

where δ is the sign of A and $a \equiv \delta B \pmod{5}$. One can check $D_{2,3,-}$ achieves the constant values $1/10$ on the set $[2/5, 3/5] \times \{2/5, 3/5\}$. If $|A| \geq 5$, then $\psi(T_{2,3,-})$ certainly intersects $[2/5, 3/5] \times \{2/5, 3/5\}$, which immediately implies that $D(T) = 1/10$. It remains only to consider $|A| \leq 4$.

For each value of $m := |A|$, there are several equivalence classes of B modulo $5m$ that give $D(T) = 1/10$ and hence can be disregarded. For $m = 1$, for instance, this occurs exactly when $2a/5 \in [2/5, 3/5]$ (or, equivalently, $3a/5 \in [2/5, 3/5]$), i.e., $a = \delta B \equiv 1, 4 \pmod{5}$; so it remains only to consider $B \equiv 0, 2, 3$

Figure 2.11: Approximations to $\tau = 2/5$ for U^3 . The values of R^+, R^- lie in $\{0, 1, 2, 3, 4\}$. We treat only the sector where every $\delta = +$, as this is the region containing all of the exceptional half-lines of interest.

i, j, ϵ, ℓ	$\psi(T_{i,j,\epsilon,\ell})$	λ^+	λ^-	$R^+ (+)$	$R^- (+)$
$1, 2, +, 0$	$\langle \frac{1}{ B-A } \rangle_{\mathbb{Z}} \times \{0\}$	4	1	$3(B-A)$	$2(B-A)$
$1, 2, -, 0$	$\langle \frac{1}{ B+A } \rangle_{\mathbb{Z}} \times \{0\}$	4	1	$3(B+A)$	$2(B+A)$
$1, 3, +, 0$	$\langle \frac{1}{ B-4A } \rangle_{\mathbb{Z}} \times \{0\}$	4	1	$3(B-4A)$	$2(B-4A)$
$1, 3, -, 0$	$\langle \frac{1}{ B+4A } \rangle_{\mathbb{Z}} \times \{0\}$	4	1	$3(B+4A)$	$2(B+4A)$

(mod 5). Likewise, for $m = 2$, we can ignore the cases where

$$2a/10 = a/5 \in [2/5, 3/5] \cup (1/2 + [2/5, 3/5]);$$

after also accounting for the condition that A, B are coprime (i.e., B is odd), we are left with checking only $B \equiv 1, 9 \pmod{10}$. Similarly, for $m = 3$, we need to consider only $B \equiv 5, 10 \pmod{15}$, and we can completely ignore the case $m = 4$.

Each of the remaining cases from the previous paragraph corresponds to an exceptional line; as B increases, the parameter $q = |A|$ remains fixed. By symmetry (see Section 2.2.5), it suffices to study the behavior around the point $\tau = 2/5$ for each tuple $(i, j, \epsilon) \neq (2, 3, -)$ with $D(U_{i,j,\epsilon}) = 1/10$; see Figure 2.11. All of our exceptional lines eventually lie within the sector containing the positive B -axis, which corresponds to the case where every $\delta_{i,j,\epsilon} = +$.

For each $(m, x) \in \{(1, 0), (1, 2), (1, 3), (2, 1), (2, 9), (3, 5), (3, 10)\}$, we consider pairs (A, B) with $|A| = m$ and $B \equiv x \pmod{5m}$; after parametrizing $B = 5|A|t + x$ (for $t \in \mathbb{N}$), we compute $D(T)$ as a function of t in the usual way. This results in several progressions in $\mathcal{S}_1(U^3)$, which we verify are contained in the progressions from the statement of Theorem 2.4. Figure 2.12 presents the results of these calculations for the half-lines with $A > 0$; the half-lines with $A < 0$ turn out to yield the same progressions, so we have omitted these calculations.

2.5.2 Computing $\mathcal{S}_1(U^4)$

We parametrize 1-dimensional subtori of U^4 as

$$T = \langle A(0, 2, 3) + B(1, 0, 0) \rangle_{\mathbb{R}}$$

for coprime A, B with $B \geq 0$. We compute the intersections $U_{i,j,\epsilon} = U^4 \cap \{x_i = \epsilon x_j\}$ in Figure 2.13. Since $U_{2,3,-}$ achieves the D -value $1/10$ on intervals of positive length, we restrict our attention to a small number

Figure 2.12: D -value progressions for U^3 with fixed $1 \leq A \leq 3$ and fixed residue class $B \equiv x \pmod{5A}$ ($x \in \{1, 2, \dots, 5A\}$). We parametrize $B = 5At + x$ ($t \in \mathbb{N}$). The second column gives the corresponding progression $5(D(T) - 1/10)$. These values (as t ranges) form subprogressions of the two progressions appearing in the statement of Theorem 2.4; the last column provides the changes of variable witnessing these containments.

(A, B)	Scaled progression	Theorem progression	Change of variable
$A = 1, B \equiv 0 \pmod{5}$	$2/(5t + 1)$	$4/(5s + 2)$	$s = 2t$
$A = 1, B \equiv 2 \pmod{5}$	$1/(5t + 3)$	$3/(5s + 4)$	$s = 3t + 1$
$A = 1, B \equiv 3 \pmod{5}$	$3/(5t + 4)$	$3/(5s + 4)$	$s = t$
$A = 2, B \equiv 1 \pmod{10}$	$1/(10t + 3)$	$3/(5s + 4)$	$s = 6t + 1$
$A = 2, B \equiv 9 \pmod{10}$	$2/(10t + 11)$	$4/(5s + 2)$	$s = 4t + 4$
$A = 3, B \equiv 5 \pmod{15}$	$1/(15t + 8)$	$4/(5s + 2)$	$s = 12t + 6$
$A = 3, B \equiv 10 \pmod{15}$	$1/(15t + 13)$	$4/(5s + 2)$	$s = 12t + 10$

Figure 2.13: Intersections of U^4 with the subspaces $\{x_i = \epsilon x_j\}$.

Subspace	$U_{i,j,\epsilon}$	$D(U_{i,j,\epsilon})$	τ achieving $1/10$
$\{x_1 = x_2\}$	$\langle(2, 2, 3)\rangle_{\mathbb{R}}$	$1/10$	$\{1/5, 4/5\}$
$\{x_1 = -x_2\}$	$\langle(-2, 2, 3)\rangle_{\mathbb{R}}$	$1/10$	$\{1/5, 4/5\}$
$\{x_1 = x_3\}$	$\langle(3, 2, 3)\rangle_{\mathbb{R}}$	$1/10$	$\{1/5, 4/5\}$
$\{x_1 = -x_3\}$	$\langle(-3, 2, 3)\rangle_{\mathbb{R}}$	$1/10$	$\{1/5, 4/5\}$
$\{x_2 = x_3\}$	$\langle(1, 0, 0)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_2 = -x_3\}$	$\cup_{i=0}^4 \langle(1, 0, 0)\rangle_{\mathbb{R}} + (0, 2i/5, 3i/5)$	$1/10$	$1, 4 : [2/5, 3/5]$

of exceptional half-lines, as in the previous subsection. This analysis leaves us with the cases $|A| = 1$ and $B \equiv 0, 1, 4 \pmod{5}$; $|A| = 2$ and $B \equiv 3, 7 \pmod{10}$; and $|A| = 3$ and $B \equiv 5, 10 \pmod{15}$. For each case, we approximate the critical value $\tau = 1/5$ for the remaining tuples (i, j, ϵ) ; see Figure 2.14.

Each of our exceptional half-lines is eventually contained in the sector containing the positive B -axis, which corresponds to the case where every $\delta = +$. Each exceptional line generates a progression, again contained in one of the progressions from the statement of Theorem 2.4; see Figure 2.15 for $A > 0$ (as in the previous subsection, $A < 0$ gives the same results and these calculations are omitted).

Figure 2.14: Approximations to $\tau = 1/5$ for U^4 . The values of R^+, R^- lie in $\{0, 1, 2, 3, 4\}$. We consider the case where every $\delta = +$ as that corresponds to the region of interest.

i, j, ϵ, ℓ	$\psi(T_{i,j,\epsilon,\ell})$	λ^+	λ^-	$R^+ (+)$	$R^- (+)$
$1, 2, +, 0$	$\langle \frac{1}{ B-2A } \rangle_{\mathbb{Z}} \times \{0\}$	3	2	$2A - B$	$B - 2A$
$1, 2, -, 0$	$\langle \frac{1}{ B+2A } \rangle_{\mathbb{Z}} \times \{0\}$	3	2	$-2A - B$	$2A + B$
$1, 3, +, 0$	$\langle \frac{1}{ B-3A } \rangle_{\mathbb{Z}} \times \{0\}$	3	2	$3A - B$	$B - 3A$
$1, 3, -, 0$	$\langle \frac{1}{ B+3A } \rangle_{\mathbb{Z}} \times \{0\}$	3	2	$-3A - B$	$3A + B$

Figure 2.15: D -value progressions for U^4 with fixed $1 \leq A \leq 3$ and fixed residue class $B \equiv x \pmod{5A}$ ($x \in \{1, 2, \dots, 5A\}$). We parametrize $B = 5At + x$ ($t \in \mathbb{N}$). The second column gives the corresponding progression $5(D(T) - 1/10)$. These values (as t ranges) form subprogressions of the two progressions appearing in the statement of Theorem 2.4.

(A, B)	Scaled progression	Theorem progression	Change of variable
$A = 1, B \equiv 0 \pmod{5}$	$4/(5t + 2)$	$4/(5s + 2)$	$s = t$
$A = 1, B \equiv 1 \pmod{5}$	$3/(5t + 4)$	$3/(5s + 4)$	$s = t$
$A = 1, B \equiv 4 \pmod{5}$	$2/(5t + 6)$	$4/(5s + 2)$	$s = 2t + 2$
$A = 2, B \equiv 3 \pmod{10}$	$3/(10t + 9)$	$3/(5s + 4)$	$s = 2t + 1$
$A = 2, B \equiv 7 \pmod{10}$	$2/(10t + 11)$	$4/(5s + 2)$	$s = 4t + 4$
$A = 3, B \equiv 5 \pmod{15}$	$2/(15t + 11)$	$4/(5s + 2)$	$s = 6t + 4$
$A = 3, B \equiv 10 \pmod{15}$	$2/(15t + 16)$	$4/(5s + 2)$	$s = 6t + 6$

Figure 2.16: Intersections of U^5 with subspaces $\{x_i = \epsilon x_j\}$.

Subspace	$U_{i,j,\epsilon}$	$D(U_{i,j,\epsilon})$	τ achieving $1/10$
$\{x_1 = x_2\}$	$\langle(1, 1, 0)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_1 = -x_2\}$	$\langle(-1, 1, -4)\rangle_{\mathbb{R}}$	$1/10$	$\{2/5, 3/5\}$
$\{x_1 = x_3\}$	$\langle(2, 1, 2)\rangle_{\mathbb{R}}$	$1/6$	—
$\{x_1 = -x_3\}$	$\langle(-2, -3, 2)\rangle_{\mathbb{R}}$	$1/10$	$\{1/5, 4/5\}$
$\{x_2 = x_3\}$	$\langle(3, 2, 2)\rangle_{\mathbb{R}}$	$1/10$	$\{1/5, 4/5\}$
$\{x_2 = -x_3\}$	$\langle(-1, -2, 2)\rangle_{\mathbb{R}}$	$1/6$	—

2.5.3 Computing $\mathcal{S}_1(U^5)$

For U^5 we will have to work with sector decompositions instead of only exceptional half-lines. We parametrize 1-dimensional subtori of U^5 as

$$T = \langle A(0, 1, 3) + B(1, 0, 1) \rangle_{\mathbb{R}}$$

for coprime A, B with $A \geq 0$. We compute the intersections $U_{i,j,\epsilon} = U^5 \cap \{x_i = \epsilon x_j\}$ and their corresponding D -values in Figure 2.16. In Figure 2.17, we show the calculations for approximating the critical points τ . Finally, for each pair of residues for $(A, B) \pmod{5}$, we partition $\mathbb{R}_{\geq 0} \times \mathbb{R}$ into sectors, calculate $D(T)$ on each sector, and give the corresponding progressions; see Figure 2.18.

Figure 2.18 tells us that the expression $5(D(T) - 1/10)$ assumes values in the progressions

$$\frac{1}{5t + 3}, \quad \frac{2}{5t + 1}, \quad \frac{3}{5t + 4}, \quad \frac{4}{5t + 2}.$$

Regarding the first progression, we observe that $\{1/(5t + 3) : t \in \mathbb{N}\} \subseteq \{4/(5s + 2) : s \in \mathbb{N}\}$ (via $s = 4t + 2$) and $\{1/(5t + 3) : t \in \mathbb{N}\} \subseteq \{3/(5s + 4) : s \in \mathbb{N}\}$ (via $s = 3t + 1$). In fact, $\{1/(5t + 3) : t \in \mathbb{N}\}$ is precisely these two progressions containing it. Likewise $\{2/(5t + 1) : t \in \mathbb{N}\} \subseteq \{4/(5s + 2) : s \in \mathbb{N}\}$ (via $s = 2t$). Thus, we obtain only the two progressions from the statement of Theorem 2.4.

Figure 2.17: Approximations to $\tau = 1/5, 2/5$ for U^5 . The values of R^+, R^- lie in $\{0, 1, 2, 3, 4\}$.

i, j, ϵ	τ	$\psi(T_{i,j,\epsilon})$	λ^+	λ^-	$R^+ (+)$	$R^- (+)$	$R^+ (-)$	$R^- (-)$
1, 2, -	2/5	$\langle \frac{1}{ A+2B } \rangle_{\mathbb{Z}}$	4	1	$3A + B$	$2A + 4B$	$2A + 4B$	$3A + B$
1, 3, -	1/5	$\langle \frac{1}{ 3A+B } \rangle_{\mathbb{Z}}$	3	2	$2A + 4B$	$3A + B$	$3A + B$	$2A + 4B$
2, 3, +	1/5	$\langle \frac{1}{ B-2A } \rangle_{\mathbb{Z}}$	3	2	$2A + 4B$	$3A + B$	$3A + B$	$2A + 4B$

Figure 2.18: D -values in sectors for U^5 . The third column records the offsets $5(D(T) - 1/10)$.

(A, B)	Slopes of dividing rays	Offsets	Scaled progressions
(0, 1)	-1	$\frac{2}{3A+B}, \frac{1}{-A-2B}$	$\frac{2}{5t+1}, \frac{1}{5t+3}$
(0, 2)	1, -1	$\frac{2}{A+2B}, \frac{3A+B}{4}, \frac{-A-2B}{3}$	$\frac{2}{5t+4}, \frac{5t+2}{4}, \frac{2}{5t+1}$
(0, 3)	0, -2	$\frac{A+2B}{1}, \frac{2A-B}{2}, \frac{-A-2B}{3}$	$\frac{5t+1}{1}, \frac{5t+2}{2}, \frac{5t+4}{3}$
(0, 4)	0	$\frac{A+2B}{2}, \frac{2A-B}{4}$	$\frac{5t+3}{2}, \frac{5t+1}{4}$
(1, 0)	0, -2	$\frac{A+2B}{1}, \frac{2A-B}{2}, \frac{-A-2B}{3}$	$\frac{5t+1}{1}, \frac{5t+2}{2}, \frac{5t+4}{3}$
(1, 1)	0	$\frac{A+2B}{1}, \frac{2A-B}{2}$	$\frac{5t+3}{1}, \frac{5t+1}{2}$
(1, 2)	(none)	0	-
(1, 3)	-1	$\frac{2}{3A+B}, \frac{1}{-A-2B}$	$\frac{2}{5t+1}, \frac{1}{5t+3}$
(1, 4)	0, -1	$\frac{A+2B}{2}, \frac{2A-B}{4}, \frac{-A-2B}{3}$	$\frac{5t+4}{2}, \frac{5t+2}{4}, \frac{2}{5t+1}$
(2, 0)	-1	$\frac{2}{3A+B}, \frac{1}{-A-2B}$	$\frac{2}{5t+1}, \frac{1}{5t+3}$
(2, 1)	1, -1	$\frac{A+2B}{2}, \frac{3A+B}{4}, \frac{-A-2B}{3}$	$\frac{5t+4}{2}, \frac{5t+2}{4}, \frac{2}{5t+1}$
(2, 2)	0, -2	$\frac{A+2B}{1}, \frac{2A-B}{2}, \frac{-A-2B}{3}$	$\frac{5t+1}{1}, \frac{5t+2}{2}, \frac{5t+4}{3}$
(2, 3)	0	$\frac{A+2B}{1}, \frac{2A-B}{2}$	$\frac{5t+3}{1}, \frac{5t+1}{2}$
(2, 4)	(none)	0	-
(3, 0)	0	$\frac{1}{A+2B}, \frac{2}{2A-B}$	$\frac{1}{5t+3}, \frac{2}{5t+1}$
(3, 1)	(none)	0	-
(3, 2)	-1	$\frac{2}{3A+B}, \frac{1}{-A-2B}$	$\frac{2}{5t+1}, \frac{1}{5t+3}$
(3, 3)	1, -1	$\frac{A+2B}{2}, \frac{3A+B}{4}, \frac{-A-2B}{3}$	$\frac{5t+4}{2}, \frac{5t+2}{4}, \frac{2}{5t+1}$
(3, 4)	0, -2	$\frac{A+2B}{2}, \frac{2A-B}{4}, \frac{-A-2B}{3}$	$\frac{5t+1}{2}, \frac{5t+2}{4}, \frac{5t+4}{3}$
(4, 0)	1, -1	$\frac{A+2B}{3}, \frac{3A+B}{4}, \frac{-A-2B}{2}$	$\frac{5t+4}{3}, \frac{5t+2}{4}, \frac{2}{5t+1}$
(4, 1)	0, -2	$\frac{A+2B}{2}, \frac{3A+B}{4}, \frac{-A-2B}{3}$	$\frac{5t+4}{2}, \frac{5t+2}{4}, \frac{5t+1}{3}$
(4, 2)	0	$\frac{A+2B}{1}, \frac{2A-B}{2}, \frac{-A-2B}{3}$	$\frac{5t+1}{1}, \frac{5t+2}{2}, \frac{5t+4}{3}$
(4, 3)	(none)	0	-
(4, 4)	-1	$\frac{2}{3A+B}, \frac{1}{-A-2B}$	$\frac{2}{5t+1}, \frac{1}{5t+3}$

Figure 2.19: Intersections of U^6 with the subspaces $\{x_i = \epsilon x_j\}$.

Subspace	$U_{i,j,\epsilon}$	$D(U_{i,j,\epsilon})$	τ achieving $1/10$
$\{x_1 = x_2\}$	$\langle(1, 1, 4)\rangle_{\mathbb{R}}$	$1/10$	$\{2/5, 3/5\}$
$\{x_1 = -x_2\}$	$\langle(1, -1, -2)\rangle_{\mathbb{R}}$	$1/6$	—
$\{x_1 = x_3\}$	$\langle(1, 0, 1)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_1 = -x_3\}$	$\langle(-3, 2, 3)\rangle_{\mathbb{R}}$	$1/10$	$\{1/5, 4/5\}$
$\{x_2 = x_3\}$	$\langle(-2, 1, 1)\rangle_{\mathbb{R}}$	$1/6$	—
$\{x_2 = -x_3\}$	$\langle(-4, 1, -1)\rangle_{\mathbb{R}}$	$1/10$	$\{2/5, 3/5\}$

Figure 2.20: Approximations to $\tau = 1/5, 2/5$ for U^6 . The values of R^+, R^- lie in $\{0, 1, 2, 3, 4\}$.

i, j, ϵ	τ	$\psi(T_{i,j,\epsilon})$	λ^+	λ^-	$R^+ (+)$	$R^- (+)$	$R^+ (-)$	$R^- (-)$
$1, 2, +$	$2/5$	$\langle \frac{1}{ B-A } \rangle_{\mathbb{Z}}$	4	1	$2A + 3B$	$3A + 2B$	$3A + 2B$	$2A + 3B$
$1, 3, -$	$1/5$	$\langle \frac{1}{ 3A+2B } \rangle_{\mathbb{Z}}$	3	2	$2A + 3B$	$3A + 2B$	$3A + 2B$	$2A + 3B$
$2, 3, -$	$2/5$	$\langle \frac{1}{ 4A+B } \rangle_{\mathbb{Z}}$	4	1	$2A + 3B$	$3A + 2B$	$3A + 2B$	$2A + 3B$

2.5.4 Computing $\mathcal{S}_1(U^6)$

We proceed as in the previous subsection. We parametrize 1-dimensional subtori of U^6 as

$$T = \langle A(0, 1, 2) + B(1, 1, 0) \rangle_{\mathbb{R}}$$

for coprime A, B with $A \geq 0$. As before, we intersect U^6 with the subspaces $\{x_i = \epsilon x_j\}$; see Figure 2.19. Then we approximate critical values of τ ; see Figure 2.20. Finally, we partition $\mathbb{R}_{\geq 0} \times \mathbb{R}$ into sectors for each pair of equivalence classes for (A, B) , calculate the corresponding values of $D(T)$, and represent these values as progressions; see Figure 2.21. As before we get exactly the two progressions from the statement of Theorem 2.4.

2.6 A new progression in the spectrum $\mathcal{S}_1(6)$

The goal of this section is to prove Theorem 2.5, which states that $1/3 + 1/6 \text{Prog}(6, 11) \subseteq \mathcal{S}_1(6)$. This example is noteworthy because it provides a new progression, beyond the progression $1/3 + 1/6 \text{Prog}(6, 7)$ predicted by (2.1). The progression $1/3 + 1/6 \text{Prog}(6, 11)$ appears in the relative spectrum of

$$U^7 := \langle (1, 0, 1, 2, 3, 3), (0, 1, 1, 1, 1, 2) \rangle_{\mathbb{R}}.$$

We remark that this choice of generators bears some resemblance to the generators of $U^2 \subseteq (\mathbb{R}/\mathbb{Z})^4$, which (after permuting the coordinates) can be written as $(1, 0, 1, 1), (0, 1, 1, 2)$ (but we were not able to find an

Figure 2.21: D -values in sectors for U^6 .

(A, B)	Slopes of dividing rays	Offsets	Scaled progressions
$(0, 1)$	-2	$\frac{2}{4A+B}, \frac{3}{A-B}$	$\frac{2}{5t+1}, \frac{3}{5t+4}$
$(0, 2)$	0	$\frac{3}{3A+2B}, \frac{1}{A-B}$	$\frac{5t+4}{3}, \frac{5t+3}{1}$
$(0, 3)$	-3	$\frac{4A+B}{3}, \frac{-3A-2B}{3}$	$\frac{5t+3}{3}, \frac{5t+4}{2}$
$(0, 4)$	-1	$\frac{4A+B}{3}, \frac{A-B}{2}$	$\frac{5t+4}{3}, \frac{5t+1}{2}$
$(1, 0)$	-1	$\frac{4A+B}{3}, \frac{A-B}{2}$	$\frac{5t+4}{3}, \frac{5t+1}{2}$
$(1, 1)$	(none)	0	$-$
$(1, 2)$	-2	$\frac{2}{4A+B}, \frac{3}{A-B}$	$\frac{2}{5t+1}, \frac{3}{5t+4}$
$(1, 3)$	0	$\frac{3}{3A+2B}, \frac{1}{A-B}$	$\frac{5t+4}{3}, \frac{5t+3}{1}$
$(1, 4)$	-3	$\frac{4A+B}{3}, \frac{-3A-2B}{3}$	$\frac{5t+3}{3}, \frac{5t+4}{2}$
$(2, 0)$	-3	$\frac{4A+B}{3}, \frac{-3A-2B}{3}$	$\frac{5t+3}{3}, \frac{5t+4}{2}$
$(2, 1)$	-1	$\frac{4A+B}{3}, \frac{A-B}{2}$	$\frac{5t+4}{3}, \frac{5t+1}{2}$
$(2, 2)$	(none)	0	$-$
$(2, 3)$	-2	$\frac{2}{4A+B}, \frac{3}{A-B}$	$\frac{2}{5t+1}, \frac{3}{5t+4}$
$(2, 4)$	0	$\frac{3}{3A+2B}, \frac{1}{A-B}$	$\frac{5t+4}{3}, \frac{5t+3}{1}$
$(3, 0)$	0	$\frac{3}{3A+2B}, \frac{1}{A-B}$	$\frac{5t+4}{3}, \frac{5t+3}{1}$
$(3, 1)$	-3	$\frac{4A+B}{3}, \frac{-3A-2B}{3}$	$\frac{5t+3}{3}, \frac{5t+4}{2}$
$(3, 2)$	-1	$\frac{4A+B}{3}, \frac{A-B}{2}$	$\frac{5t+4}{3}, \frac{5t+1}{2}$
$(3, 3)$	(none)	0	$-$
$(3, 4)$	-2	$\frac{2}{4A+B}, \frac{3}{A-B}$	$\frac{2}{5t+1}, \frac{3}{5t+4}$
$(4, 0)$	-2	$\frac{4A+B}{2}, \frac{A-B}{3}$	$\frac{5t+1}{2}, \frac{5t+4}{3}$
$(4, 1)$	0	$\frac{3}{3A+2B}, \frac{1}{A-B}$	$\frac{5t+4}{3}, \frac{5t+3}{1}$
$(4, 2)$	-3	$\frac{4A+B}{3}, \frac{-3A-2B}{3}$	$\frac{5t+3}{3}, \frac{5t+4}{2}$
$(4, 3)$	-1	$\frac{4A+B}{3}, \frac{A-B}{2}$	$\frac{5t+4}{3}, \frac{5t+1}{2}$
$(4, 4)$	(none)	0	$-$

Figure 2.22: Intersections of U^7 with the subspaces $\{x_i = \epsilon x_j\}$.

Subspace	$U_{i,j,\epsilon}$	$D(U_{i,j,\epsilon})$	τ achieving $1/3$
$\{x_1 = x_2\}$	$\langle(1, 1, 2, 3, 4, 5)\rangle_{\mathbb{R}}$	$1/3$	$\{1/6, 5/6\}$
$\{x_1 = -x_2\}$	$\langle(1, -1, 0, 1, 2, 1)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_1 = x_3\}$	$\langle(1, 0, 1, 2, 3, 3)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_1 = -x_3\}$	$\langle(1, -2, -1, 0, 1, -1)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_1 = x_4\}$	$\langle(1, -1, 0, 1, 2, 1)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_1 = -x_4\}$	$\langle(1, -3, -2, -1, 0, -3)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_1 = x_5\}$	$\langle(1, -2, -1, 0, 1, -1)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_1 = -x_5\}$	$\langle(1, -4, -3, -2, -1, -5)\rangle_{\mathbb{R}}$	$1/3$	$\{1/6, 5/6\}$
$\{x_1 = x_6\}$	$\cup_{\ell=0}^1 \langle(1, -1, 0, 1, 2, 1)\rangle_{\mathbb{R}} + (0, \frac{\ell}{2}, \frac{\ell}{2}, \frac{\ell}{2}, \frac{\ell}{2}, 0)$	$1/3$	$1 : \{1/6, 1/3, 2/3, 5/6\}$
$\{x_1 = -x_6\}$	$\cup_{\ell=0}^1 \langle(1, -2, -1, 0, 1, -1)\rangle_{\mathbb{R}} + (0, \frac{\ell}{2}, \frac{\ell}{2}, \frac{\ell}{2}, \frac{\ell}{2}, 0)$	$1/3$	$1 : \{1/6, 1/3, 2/3, 5/6\}$
$\{x_2 = x_3\}$	$\langle(0, 1, 1, 1, 1, 2)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_2 = -x_3\}$	$\langle(2, -1, 1, 3, 5, 4)\rangle_{\mathbb{R}}$	$1/3$	$\{1/6, 5/6\}$
$\{x_2 = x_4\}$	$\cup_{\ell=0}^1 \langle(0, 1, 1, 1, 1, 2)\rangle_{\mathbb{R}} + (\frac{\ell}{2}, 0, \frac{\ell}{2}, 0, \frac{\ell}{2}, \frac{\ell}{2})$	$1/3$	$1 : \{1/6, 1/3, 2/3, 5/6\}$
$\{x_2 = -x_4\}$	$\cup_{\ell=0}^1 \langle(1, -1, 0, 1, 2, 1)\rangle_{\mathbb{R}} + (0, \frac{\ell}{2}, \frac{\ell}{2}, \frac{\ell}{2}, \frac{\ell}{2}, 0)$	$1/3$	$1 : \{1/6, 1/3, 2/3, 5/6\}$
$\{x_2 = x_5\}$	$\cup_{\ell=0}^2 \langle(0, 1, 1, 1, 1, 2)\rangle_{\mathbb{R}} + (\frac{\ell}{3}, 0, \frac{\ell}{3}, \frac{2\ell}{3}, 0, 0)$	$1/3$	$1, 2 : \{1/6, 5/6\}$
$\{x_2 = -x_5\}$	$\langle(2, -3, -1, 1, 3, 0)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_2 = x_6\}$	$\langle(-1, 3, 2, 1, 0, 3)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_2 = -x_6\}$	$\cup_{\ell=0}^2 \langle(1, -1, 0, 1, 2, 1)\rangle_{\mathbb{R}} + (0, \frac{\ell}{3}, \frac{\ell}{3}, \frac{\ell}{3}, \frac{\ell}{3}, \frac{2\ell}{3})$	$1/3$	$1 : \{1/6, 1/2\};$ $2 : \{1/2, 5/6\}$
$\{x_3 = x_4\}$	$\langle(0, 1, 1, 1, 1, 2)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_3 = -x_4\}$	$\langle(2, -3, -1, 1, 3, 0)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_3 = x_5\}$	$\cup_{\ell=0}^1 \langle(0, 1, 1, 1, 1, 2)\rangle_{\mathbb{R}} + (\frac{\ell}{2}, 0, \frac{\ell}{2}, 0, \frac{\ell}{2}, \frac{\ell}{2})$	$1/3$	$1 : \{1/6, 1/3, 2/3, 5/6\}$
$\{x_3 = -x_5\}$	$\cup_{\ell=0}^1 \langle(1, -2, -1, 0, 1, -1)\rangle_{\mathbb{R}} + (0, \frac{\ell}{2}, \frac{\ell}{2}, \frac{\ell}{2}, \frac{\ell}{2}, 0)$	$1/3$	$1 : \{1/6, 1/3, 2/3, 5/6\}$
$\{x_3 = x_6\}$	$\langle(1, -2, -1, 0, 1, -1)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_3 = -x_6\}$	$\langle(3, -4, -1, 2, 5, 1)\rangle_{\mathbb{R}}$	$1/3$	$\{1/6, 5/6\}$
$\{x_4 = x_5\}$	$\langle(0, 1, 1, 1, 1, 2)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_4 = -x_5\}$	$\langle(2, -5, -3, -1, 1, -4)\rangle_{\mathbb{R}}$	$1/3$	$\{1/6, 5/6\}$
$\{x_4 = x_6\}$	$\langle(1, -1, 0, 1, 2, 1)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_4 = -x_6\}$	$\langle(3, -5, -2, 1, 4, -1)\rangle_{\mathbb{R}}$	$1/3$	$\{1/6, 5/6\}$
$\{x_5 = x_6\}$	$\langle(1, 0, 1, 2, 3, 3)\rangle_{\mathbb{R}}$	$1/2$	—
$\{x_5 = -x_6\}$	$\cup_{\ell=0}^2 \langle(1, -2, -1, 0, 1, -1)\rangle_{\mathbb{R}} + (0, \frac{\ell}{3}, \frac{\ell}{3}, \frac{\ell}{3}, \frac{\ell}{3}, \frac{2\ell}{3})$	$1/3$	$1 : \{1/2, 5/6\};$ $2 : \{1/6, 1/2\}$

Figure 2.23: Approximations to the various values of τ for U^7 . Throughout, we consider the sector in which the half-line $\{6\} \times \mathbb{R}_{>0}$ is (eventually) contained; this sector has $\delta = -$ for the first row and $\delta = +$ for all other rows. In rows three, four, eight, and nine, $a \equiv B \pmod{2}$; in rows five and six, $a \equiv A \pmod{2}$; in rows ten and eleven, $a \equiv \delta B \pmod{3}$; in rows twelve and thirteen $a \equiv -\delta B \pmod{3}$; and in rows seventeen and eighteen, $a \equiv \delta A \pmod{3}$.

i, j, ϵ, ℓ	τ	$\psi(T_{i,j,\epsilon,\ell})$	λ^+	λ^-	R^+	R^-	$D(T_{i,j,\epsilon,\ell}) - \frac{1}{3}$
1, 2, +, 0	1/6	$\langle \frac{1}{ A-B } \rangle_{\mathbb{Z}} \times \{0\}$	5	1	$A + 5B \rightsquigarrow 1$	$5A + B \rightsquigarrow 5$	$\frac{5}{6(B-A)}$
1, 5, -, 0	1/6	$\langle \frac{1}{ 4A+B } \rangle_{\mathbb{Z}} \times \{0\}$	5	1	$2A + 5B \rightsquigarrow 1$	$4A + B \rightsquigarrow 5$	$\frac{5}{6(4A+B)}$
1, 6, +, 1	1/6	$\left(\frac{a}{2 A+B } + \langle \frac{1}{ A+B } \rangle_{\mathbb{Z}} \right) \times \{\frac{1}{2}\}$	2	1	$5A + 2B \rightsquigarrow 4$	$A + 4B \rightsquigarrow 2$	$\frac{2}{6(A+B)}$
1, 6, +, 1	1/3	$\left(\frac{a}{2 A+B } + \langle \frac{1}{ A+B } \rangle_{\mathbb{Z}} \right) \times \{\frac{1}{2}\}$	1	2	$4A + B \rightsquigarrow 5$	$2A + 5B \rightsquigarrow 1$	$\frac{2}{6(A+B)}$
1, 6, -, 1	1/6	$\left(\frac{a}{2 2A+B } + \langle \frac{1}{ 2A+B } \rangle_{\mathbb{Z}} \right) \times \{\frac{1}{2}\}$	2	1	$A + 5B \rightsquigarrow 1$	$5A + B \rightsquigarrow 5$	$\frac{2}{6(2A+B)}$
1, 6, -, 1	1/3	$\left(\frac{a}{2 2A+B } + \langle \frac{1}{ 2A+B } \rangle_{\mathbb{Z}} \right) \times \{\frac{1}{2}\}$	1	2	$5A + 4B \rightsquigarrow 2$	$A + 2B \rightsquigarrow 4$	$\frac{2}{6(2A+B)}$
2, 3, -, 0	1/6	$\langle \frac{1}{ A+2B } \rangle_{\mathbb{Z}} \times \{0\}$	5	1	$5A + 4B \rightsquigarrow 2$	$A + 2B \rightsquigarrow 4$	$\frac{4}{6(A+2B)}$
2, 4, +, 1	1/6	$\left(\frac{a}{2 A } + \langle \frac{1}{ A } \rangle_{\mathbb{Z}} \right) \times \{\frac{1}{2}\}$	2	1	$5A + 3B \rightsquigarrow 3$	$A + 3B \rightsquigarrow 3$	$\frac{3}{6A}$
2, 4, +, 1	1/3	$\left(\frac{a}{2 A } + \langle \frac{1}{ A } \rangle_{\mathbb{Z}} \right) \times \{\frac{1}{2}\}$	1	2	$4A + 3B \rightsquigarrow 3$	$2A + 3B \rightsquigarrow 3$	$\frac{3}{6A}$
2, 5, +, 1	1/6	$\left(\frac{a}{3 A } + \langle \frac{1}{ A } \rangle_{\mathbb{Z}} \right) \times \{\frac{1}{3}\}$	1	1	$5A + 2B \rightsquigarrow 4$	$A + 4B \rightsquigarrow 2$	$\frac{2}{6A}$
2, 5, +, 1	5/6	$\left(\frac{a}{3 A } + \langle \frac{1}{ A } \rangle_{\mathbb{Z}} \right) \times \{\frac{1}{3}\}$	1	1	$A + 2B \rightsquigarrow 4$	$5A + 4B \rightsquigarrow 2$	$\frac{2}{6A}$
2, 6, -, 1	1/6	$\left(\frac{a}{3 A+B } + \langle \frac{1}{ A+B } \rangle_{\mathbb{Z}} \right) \times \{\frac{1}{3}\}$	1	1	$5A + 3B \rightsquigarrow 3$	$A + 3B \rightsquigarrow 3$	$\frac{3}{6(A+B)}$
2, 6, -, 1	1/2	$\left(\frac{a}{3 A+B } + \langle \frac{1}{ A+B } \rangle_{\mathbb{Z}} \right) \times \{\frac{1}{3}\}$	1	1	$3A + B \rightsquigarrow 5$	$3A + 5B \rightsquigarrow 1$	$\frac{1}{6(3A+B)}$
3, 6, -, 0	1/6	$\langle \frac{1}{ 4A+3B } \rangle_{\mathbb{Z}} \times \{0\}$	5	1	$2A + 3B \rightsquigarrow 3$	$4A + 3B \rightsquigarrow 3$	$\frac{3}{6(4A+3B)}$
4, 5, -, 0	1/6	$\langle \frac{1}{ 5A+2B } \rangle_{\mathbb{Z}} \times \{0\}$	5	1	$A + 4B \rightsquigarrow 2$	$5A + 2B \rightsquigarrow 4$	$\frac{4}{6(5A+2B)}$
4, 6, -, 0	1/6	$\langle \frac{1}{ 5A+3B } \rangle_{\mathbb{Z}} \times \{0\}$	5	1	$A + 3B \rightsquigarrow 3$	$5A + 3B \rightsquigarrow 3$	$\frac{3}{6(5A+3B)}$
5, 6, -, 1	1/2	$\left(\frac{a}{3 2A+B } + \langle \frac{1}{ 2A+B } \rangle_{\mathbb{Z}} \right) \times \{\frac{1}{3}\}$	1	1	$2A + 3B \rightsquigarrow 3$	$4A + 3B \rightsquigarrow 3$	$\frac{3}{6(2A+B)}$
5, 6, -, 1	5/6	$\left(\frac{a}{3 2A+B } + \langle \frac{1}{ 2A+B } \rangle_{\mathbb{Z}} \right) \times \{\frac{1}{3}\}$	1	1	$4A + B \rightsquigarrow 5$	$2A + 5B \rightsquigarrow 1$	$\frac{1}{6(2A+B)}$

“analogous” torus in $(\mathbb{R}/\mathbb{Z})^8$.

We will show that $\mathcal{S}_1(U^7)$ contains the set $1/3 + 1/6 \text{Prog}(6, 11)$. We parametrize 1-dimensional subtori of U^7 as

$$T = \langle A(1, 0, 1, 2, 3, 3) + B(0, 1, 1, 1, 1, 2) \rangle_{\mathbb{R}},$$

where A, B are coprime and $A \geq 0$, and we check that this choice of generators satisfies (2.6). As before, we compute the intersections $U_{i,j,\epsilon} = U^7 \cap \{x_i = \epsilon x_j\}$; see Figure 2.22.

This computation differs from previous ones in several ways. First, there are several pairs of triples (i, j, ϵ) and (i', j', ϵ') such that $U_{i,j,\epsilon}$ and $U_{i',j',\epsilon'}$ are identical; obviously $D(T_{i,j,\epsilon}) = D(T_{i',j',\epsilon'})$ in such cases, and we will omit these redundant computations. Second, there are some quadruples (i, j, ϵ, ℓ) with several values of τ which must be analyzed individually, even after we account for the usual $x \mapsto -x$ symmetries. This is visible in Figure 2.23, where the triples $(1, 6, +, 1)$, $(1, 6, -, 1)$, and $(2, 4, +, 1)$ each have two rows corresponding to different values of τ . For the triples $(2, 6, -)$ and $(5, 6, -)$ (which correspond to subgroups $U_{i,j,\epsilon}$ with three components), the values of τ no longer come in pairs $(\tau, -\tau)$ on each connected component $U_{i,j,\epsilon,\ell}$ (although of course a value τ on $U_{i,j,\epsilon,\ell}$ still corresponds to a value $-\tau$ on $U_{i,j,\epsilon,-\ell}$). For triples (i, j, ϵ) where $U_{i,j,\epsilon}$ has 3 components, the symmetry discussed in Section 2.2.5 guarantees that $D(U_{i,j,\epsilon,1}) = D(U_{i,j,\epsilon,2})$, so in such cases we will analyze only $\ell = 1$.

Recall that we are concerned only with finding a single progression in $\mathcal{S}_1(U^7)$, not with computing the entire relative spectrum. So we will restrict our attention to the half-line yielding this progression, namely, the half-line consisting of pairs (A, B) with $A = 6$ and $B \equiv 5 \pmod{6}$ (with $B > 6$); note that such A, B are always coprime. We analyze only the sector containing this half-line. Comparing the expressions in the last column of Figure 2.23, we find that

$$D(T) = \frac{1}{3} + \frac{1}{6(2A + B)},$$

and substituting $A = 6$ and $B = 6s + 5$ (for $s \in \mathbb{N}$) gives

$$D(T) = \frac{1}{3} + \frac{1}{6(6s + 17)}.$$

The choice $(A, B) = (5, 1)$ leads to $D(T) = 1/3 + 1/(6 \cdot 11)$, and the choice $(A, B) = (6, 5)$ leads to $D(T) = 1/3 + 1/(6 \cdot 17)$. Altogether, these values comprise the progression $1/3 + 1/6 \text{Prog}(6, 11)$; this proves Theorem 2.5.

We remark that $\mathcal{S}_1(U^7)$ also contains the progression $1/3 + 1/6 \text{Prog}(6, 7)$; this gives a new way to achieve the progression in (2.1). This progression comes from the sector $\{0 \leq A \leq B\}$ with the modular constraints

$(A, B) \equiv (1, 0) \pmod{6}$. A further analysis along the lines of the previous two sections (not shown in full detail here) reveals that $\mathcal{S}_1(U^7)$ does not contain any other progressions.

2.7 Finite and infinite relative spectra

In this final section, we make a few brief remarks about when relative spectra are finite. Let $n \geq 3$, and let $U \subseteq (\mathbb{R}/\mathbb{Z})^n$ be a 2-dimensional proper subtorus. Write

$$Z(U) := \{x \in U : D(x) = D(U)\}$$

for the locus of points where U achieves its D -value. Notice that $Z(U)$ is the intersection of U with the hollow cube $\{x \in (\mathbb{R}/\mathbb{Z})^n : D(x) = D(U)\}$. We can express $Z(U)$ as the union of the intersections of U with the various facets of this hollow cube. Each such intersection has dimension at most 1 (otherwise U would be contained in a translate of a coordinate hyperplane), and we conclude that $Z(U)$ consists of a finite union of points and (closed) line segments. Of course these points and the endpoints of these segments all have rational coordinates.

Our criterion for the finiteness of $\mathcal{S}_1(U)$ concerns whether or not $Z(U)$ contains non-parallel line segments; notice that the “parallel-ness” of line segments is the same in $(\mathbb{R}/\mathbb{Z})^n$ and in U (considered as a 2-dimensional torus).

Proposition 2.16. *Let $n \geq 3$ be a natural number, and let $U \subseteq (\mathbb{R}/\mathbb{Z})^n$ be a 2-dimensional proper subtorus. The relative spectrum $\mathcal{S}_1(U)$ is finite if and only if the set $Z(U)$ contains non-parallel line segments.*

Proof. First, suppose that $Z(U)$ contains two non-parallel line segments. To show that $\mathcal{S}_1(U)$ is finite, it suffices to show that all but finitely many 1-dimensional subtori $T \subseteq U$ intersect the union of these line segments and thus satisfy $D(T) = D(U)$. Fix an identification of U with $(\mathbb{R}/\mathbb{Z})^2$, and consider the images s_1, s_2 of the two non-parallel line segments. Let ℓ be the minimum of the lengths of s_1, s_2 , and let L_1, L_2 denote the lines containing s_1, s_2 (respectively). Let $0 < \theta \leq \pi/2$ be the measure of the angle formed by L_1, L_2 .

Now, let $T' \subseteq (\mathbb{R}/\mathbb{Z})^2$ be a 1-dimensional subtorus. It forms an angle of at least $\theta/2$ with some L_i . Then the height of s_i in the direction perpendicular to T' is at least

$$\ell \sin(\theta/2).$$

The perpendicular distance between consecutive leaves of T' is $1/\text{vol}(T')$, where $\text{vol}(T')$ denotes the volume

(i.e., length) of T' . Thus, if $\text{vol}(T') \geq 1/(\ell \sin(\theta/2))$, then T' must intersect s_i , and we conclude that the corresponding 1-dimensional subtorus $T \subseteq U$ has $D(T) = D(U)$. Since there are only finitely many 1-dimensional subtori with volume below any given threshold, it follows that all but finitely many 1-dimensional subtori $T \subseteq U$ have $D(T) = D(U)$, as desired.

Second, suppose that $Z(U)$ does not contain any non-parallel line segments. To show that $\mathcal{S}_1(U)$ is infinite, it suffices to show that there are infinitely many 1-dimensional subtori $T \subseteq U$ that are disjoint from $Z(U)$: Indeed, such T 's have $D(T) > D(U)$, and there are infinitely many distinct such values since $D(T) \rightarrow D(U)$ as $\text{vol}(T) \rightarrow \infty$. By assumption, $Z(U)$ is contained in a finite union of rational parallel lines and hence is contained in a 1-dimensional subgroup Γ of U .

We claim that U contains at least two non-parallel 1-dimensional subtori that are contained in the coordinate hyperplanes of $(\mathbb{R}/\mathbb{Z})^n$. Indeed, for each $1 \leq i \leq n$, the intersection $U \cap \{x_i = 0\}$ is a 1-dimensional subgroup, and the identity components of these intersections cannot all be identical because then this common identity component would consist of only the point 0. Thus, there is a 1-dimensional subtorus T^* of U that is contained in a coordinate hyperplane of $(\mathbb{R}/\mathbb{Z})^n$ and is not parallel to Γ . Fix an identification of U with $(\mathbb{R}/\mathbb{Z})^2$ such that T^* is identified with the x_2 -axis; let Γ' denote the image of Γ under this identification. Since $D(T^*) = 1/2 > D(U)$, we know that $Z(U)$ is disjoint from T^* , and it follows that the image of $Z(U)$ is contained in $\Gamma' \setminus (\{0\} \times (\mathbb{R}/\mathbb{Z}))$. Hence it suffices to find infinitely many 1-dimensional subtori $T \subseteq (\mathbb{R}/\mathbb{Z})^2$ whose intersections with Γ' are contained in $\{0\} \times (\mathbb{R}/\mathbb{Z})$.

The choice of T^* guarantees that Γ' has some finite slope b/a , where a, b are coprime integers with $a > 0$. Let $q \in \mathbb{N}$ be such that the leaves of Γ' are spaced out $1/q$ apart vertically. Now, for each $s \in \mathbb{N}$, consider the 1-dimensional subtorus

$$T_s := \langle (a, b + a/(qs)) \rangle_{\mathbb{R}}.$$

As we travel along T_s starting at $(0, 0)$, the first intersection with Γ' occurs at the point

$$(s/a)(a, b + 1/(qs)) = (s, sb/a + 1/q) \in \{0\} \times (\mathbb{R}/\mathbb{Z}),$$

so in fact $T_s \cap \Gamma' \subseteq \{0\} \times (\mathbb{R}/\mathbb{Z})$, as desired. □

We remark that [65] established $\mathcal{S}_1(n-1) \subset \text{acc } \mathcal{S}_1(n)$ by showing that $\mathcal{S}_1(U' \times (\mathbb{R}/\mathbb{Z}))$ is infinite for every proper 1-dimensional subtorus $U' \subseteq (\mathbb{R}/\mathbb{Z})^{n-1}$; the construction in the second part of the proof of Proposition 2.16 is a generalization of this argument.

It is not a priori obvious that the geometric condition for finiteness in Proposition 2.16 is ever satisfied. In fact, one can show by ad-hoc arguments that it is never satisfied for $n \leq 6$. For $n = 7$, however, there are

examples like the following. Consider the 2-dimensional proper subtorus

$$U^8 := \langle (1, 2, 3, 2, 0, 0, 0), (0, 0, 0, 2, 1, 2, 3) \rangle_{\mathbb{R}}.$$

A tedious but routine calculation gives that $D(U^8) = 3/10$, and that $Z(U^8)$ consists of the points $a(1, 2, 3, 2, 0, 0, 0) + b(0, 0, 0, 2, 1, 2, 3)$ where (a, b) either lies on one of the eight line segments

$$\{2/5, 3/5\} \times ([1/5, 4/15] \cup [11/15, 4/5]), \quad ([1/5, 4/15] \cup [11/15, 4/5]) \times \{2/5, 3/5\}.$$

or is one of the four points $(1/5, 1/5), (2/5, 2/5), (3/5, 3/5), (4/5, 4/5)$. It follows from Proposition 2.16 that $\mathcal{S}_1(U^8)$ is finite, i.e., $\mathcal{S}_1(U^8)$ does not contain any progressions. Using a computer to calculate $D(T)$ for low-volume proper 1-dimensional subtori $T \subseteq U^8$ reveals that $\mathcal{S}_1(U^8)$ contains (at least) the values $7/22, 11/34, 17/54, 23/74$ in addition to $3/10$. Hence, for $\mathcal{S}_1(U^8)$, the finite symmetric difference in Theorem 2.1 is necessary.

Chapter 3

Minimal additive complements

3.1 Introduction

The *Minkowski sum* of two subsets A, B of an (additive) abelian group G is given by

$$A + B = \{a + b : a \in A, b \in B\}.$$

Given a subset $W \subseteq G$, we say that $C \subseteq G$ is an *(additive) complement* for W in G if $W + C = G$ (that is, if every $g \in G$ can be expressed at least once as $g = w + c$, for $w \in W, c \in C$). If, moreover, no proper subset $C' \subset C$ is a complement for W , then we say that C is a *minimal (additive) complement* for W in G . Similarly, we say that $C \subseteq G$ is an *(additive) supplement* for W in G if the translations of W by elements of C are pairwise disjoint (that is, if every $g \in G$ can be expressed at most once as $g = w + c$, for $w \in W, c \in C$). If, moreover, no proper superset $C' \supset C$ is a supplement for W , then we say that C is a *maximal (additive) supplement* for W in G . The project of this chapter is to investigate which sets C arise as minimal complements and maximal supplements not for a specific subset W but rather for some W . We study these questions in both finite and infinite abelian groups.

3.1.1 Background

The study of minimal complements in infinite abelian groups was initiated by Nathanson [79], who showed that if W is a finite non-empty subset of $\mathbb{Z}$, then every complement C for W contains a minimal complement C' for W . The question of determining which subsets W of $\mathbb{Z}$ admit minimal complements has received considerable attention. Chen and Yang [24] showed that $W \subseteq \mathbb{Z}$ has a minimal complement if it is unbounded both above and below. It is easy to check that, for instance, $\mathbb{N}$ does not have a minimal complement in $\mathbb{Z}$.

Kiss, Sándor, and Yang [62] studied the existence of minimal complements for “eventually periodic” subsets of $\mathbb{Z}$. There has been some progress in infinite abelian groups other than $\mathbb{Z}$: Biswas and Saha [15] generalized Nathanson’s result by showing that if G is any abelian group and W is a finite non-empty subset of G , then any complement to W contains a minimal complement. The same authors [13] later studied minimal complements in $\mathbb{Z}^r$ more closely.

The study of minimal complements and maximal supplements is related to problems in coding theory. Indeed, when W is a ball in a Hamming space, the complements for W correspond to covering codes. The possible sizes of such covering codes have been extensively studied in various settings (see, e.g., [57]). Similarly, maximal supplements for W can be interpreted as sphere packings and thus are related to results such as Hamming bounds on block codes.

We remark that a 1995 paper of Habsieger and Ruzsa [52] treated the closely related matter of finding $C \subset \mathbb{Z}$ such that $W + C$ contains a given interval $[m] = \{1, \dots, m\}$. Additive bases and asymptotic additive bases, the study of which dates back at least to Lagrange’s Four-Square Theorem, are of course also related to additive complements. Objects similar to maximal supplements have been studied in the context of Sidon sets (see, e.g., [94] and the references therein).

The natural “inverse problem” for minimal complements is determining whether a given set $C \subseteq G$ is a minimal complement for some set W (in which case we say simply that C is a *minimal complement* in G). This study was initiated by Kwon [70], who showed that every finite non-empty subset of $\mathbb{Z}$ is a minimal complement. Biswas and Saha [14] generalized this result to $\mathbb{Z}^r$.

In infinite groups, the inverse problem for minimal complements can behave quite differently from the non-inverse problem. Nathanson [79] showed that if W is a finite subset of $\mathbb{Z}$, then every complement C for W contains a minimal complement $C' \subseteq C$ for W . The same is not true, however, for the inverse problem: for instance $C = \{0, 1, 2\}$ is a complement for $W = 2\mathbb{Z}$, but there is no $W' \subseteq 2\mathbb{Z}$ for which C is a minimal complement. (Of course, C is a minimal complement for other sets, such as $3\mathbb{Z}$.)

3.1.2 Main results

In Section 3.2, we gather several basic facts about minimal complements, including that in a finite group G , no subset $C \subset G$ of size $2|G|/3 < |C| < |G|$ can be a minimal complement. We give a description of the minimal complements in G with size greater than $3|G|/5$ and the minimal complements that are arithmetic progressions.

In Section 3.3, we show with a probabilistic argument that every “small” subset of a finite group is a minimal complement. Given a finite abelian group G , let $T(G)$ be the greatest integer T such that every

subset $C \subseteq G$ with size at most T is a minimal complement. Moreover, given a natural number n , let $T(n)$ be the minimum value of $T(G)$ as G ranges over all abelian groups of order n . In other words, $T(n)$ is the greatest integer T such that for every group G of order n , every subset $C \subseteq G$ with size at most T is a minimal complement.

Theorem 3.1. *Let G be a group of order n , and let $C \subset G$ be subset of size $|C| = k$. If*

$$0 < k \leq \frac{2^{2/3} n^{1/3}}{(3e \log n)^{2/3}},$$

then C is a minimal complement in G . In other words, $T(n) \geq \frac{n^{1/3}}{2(\log n / \log 2)^{2/3}}$.

Since this growth rate for $T(n)$ can likely be improved, we have not optimized the constant in this theorem. We remark that this bound (like many in the chapter) uses the estimate $|C - C| \leq k^2$; when we know that $C - C$ is small, we can often obtain slightly better dependence of n on k . Next, we use this theorem to obtain a generalization of the result of Biswas and Saha [13] that every nonempty finite subset of $\mathbb{Z}^r$ is a minimal complement.

Theorem 3.2. *Let G be an infinite abelian group, and let C be a finite non-empty subset of G . Then C is a minimal complement in G .*

In Section 3.4, we investigate upper bounds on $T(n)$. By examining groups G with subgroups of certain sizes, we derive (as a corollary to Proposition 3.17) that

$$\liminf_{n \rightarrow \infty} \frac{T(n)}{\sqrt{n}} \leq \sqrt{2}.$$

We also obtain an upper bound on $T(G)$ for all finite groups G .

Theorem 3.3. *For every $\varepsilon > 0$, we have that $T(G) = O(|G|^{3/4+\varepsilon})$. In particular, we have that $T(n) = O(n^{3/4+\varepsilon})$.*

The proof of Theorem 3.3 goes by showing that, with high probability, a random subset formed by including independently each element of G with probability $p = |G|^{\varepsilon-1/4}$ is not a minimal complement. The proof of Theorem 3.3 works for any $p > |G|^{\varepsilon-1/4}$ that is bounded away from 1. In particular, taking $p = 1/2$, we see that a subset chosen uniformly at random is with high probability not a minimal complement.

Say that a subset of an abelian group G is *H-invariant* for a subgroup H of G if it is invariant under translation by elements of H . Burcroff and Luntzara [22] have observed that an H -invariant subset $S \subseteq G$ is a minimal complement if and only if the image of S in G/H is a minimal complement. Thus, Theorems 3.1 and 3.3 have immediate consequences for H -invariant subsets when H has finite index.

Corollary 3.4. *When ordered by minimal period, almost all periodic subsets of $\mathbb{Z}$ (i.e., subsets that are $n\mathbb{Z}$ -invariant for some $n > 0$) are not minimal complements.*

Indeed, the proof of Theorem 3.3 shows that almost all $n\mathbb{Z}$ -invariant subsets are not minimal complements. An $n\mathbb{Z}$ -invariant subset of $\mathbb{Z}$ with minimal period n/k is determined by its intersection with $\{0, 1, \dots, n/k-1\}$, so we see that all but at most $n2^{n/2}$ $n\mathbb{Z}$ -invariant subsets have minimal period n , which implies the corollary.

In Section 3.5, we study “dual” problems about maximal supplements. (For a fixed subset W , the problems of finding the smallest complement for W and the largest supplement for W are given by dual integer programs.) It is tempting to hope that, in analogy with the minimal complement setting, every sufficiently small subset of a finite group is a maximal supplement. This is not quite the case because every minimal supplement must satisfy a natural “solidity” condition: we say that a subset C of an abelian group G is *solid* if C is not properly contained in any set D such that $D - D = C - C$. It is in fact true that every sufficiently small solid subset is a maximal supplement.

Theorem 3.5. *There is an absolute constant $b > 0$ such that the following holds: Let G be a finite abelian group, and let $C \subseteq G$ be a non-empty solid subset. If*

$$|C| \leq b \left(\frac{|G|}{\log |G|} \right)^{1/4},$$

then C is a maximal supplement in G .

We also show that every finite non-empty solid subset of an infinite abelian group is a maximal supplement; the proof is completely different from (and in many ways simpler than) the proof for the finite setting.

Theorem 3.6. *Let G be an infinite abelian group, and let C be a finite non-empty solid subset of G . Then C is a maximal supplement.*

3.1.3 Open questions

Our results on minimal complements and maximal supplements raise several natural questions. It would be interesting to find the optimal bound in Theorem 3.1 and understand the asymptotic behavior of $T(n)$. Say that $f(n) = \tilde{\Theta}(g(n))$ if $f = \Omega(g(n)(\log n)^a)$ and $f(n) = O(g(n)(\log n)^b)$ for some integers a, b .

Conjecture 3.7. *We have that $T(n) = \tilde{\Theta}(\sqrt{n})$. More generally, we have that $T(G) = \tilde{\Theta}(\sqrt{|G|})$.*

Our best constructions of small subsets of finite groups that are not minimal complements rely on the existence of subgroups of certain orders. In particular, our best upper bound on $T(p)$ for p prime is $O(p^{3/4+\varepsilon})$.

Question 3.8. *Is the behavior of $T(n)$ different when n is prime? How much does $T(G)$ depend on the structure of G in addition to the order of G ?*

We can also ask the analogous question for maximal supplements. Theorem 3.5 shows that non-empty solid subsets of size $O((n/\log n)^{1/4})$ are maximal supplements, but it is not clear that this bound is anywhere near optimal.

Question 3.9. *What is the optimal function in Theorem 3.5?*

It is also interesting to consider the case of a random set where each element is included with probability p and determine when such a set is a minimal complement. Although the property of being a minimal complement is not monotone, Theorems 3.1 and 3.3 (and their proofs) show that for small values of p a random set is with high probability a minimal complement, and that for large values of p a random set is with high probability not a minimal complement.

Question 3.10. *How does the probability that a random subset of a finite group G with parameter $p = p(|G|)$ is a minimal complement (or maximal supplement) vary with p ? Is there a (sharp) threshold?*

3.2 Preliminaries

We begin with a few general results on the behavior of minimal complements under subgroups and quotients, which will be useful in the sequel. Note first of all that the collection of minimal complements in a fixed group G is invariant under translations by elements of G and applications of automorphisms of G .

Lemma 3.11. *Let G be an abelian group, and let H be a subgroup of G . If the subset $C \subseteq H$ is a minimal complement in H , then it is also a minimal complement in G .*

Proof. Suppose C is a minimal complement for W in H . Let K contain exactly one element from each coset of H . Then C is a minimal complement for $W + K$ in G : indeed,

$$(W + K) + C = K + (W + C) = K + H = G$$

shows that C is a complement for $W + K$, and minimality follows from the observation that $((W + K) + C) \cap H$ is the translate of $W + C$ by the unique element of $K \cap H$. $\square$

Lemma 3.12. *Let C be a subset of an abelian group G , and let $\pi : G \rightarrow H$ be a surjective group homomorphism such that the restriction of π to C is injective. If $\pi(C)$ is a minimal complement (in H), then C is a minimal complement (in G).*

Proof. Let $W \subseteq H$ be a subset for which $\pi(C)$ is a minimal complement. Then we claim that C is a minimal complement for $\pi^{-1}(W)$. Indeed, it is clear that $C + \pi^{-1}(W) = G$ (since this set intersects every coset of $\ker(\pi)$ and is invariant under $\ker(\pi)$). To see that C is minimal, note that any proper subset $D \subset C$ with $D + \pi^{-1}(W) = G$ would give a proper subset $\pi(D) \subset \pi(C)$ that is a minimal complement for W . $\square$

We now turn to the possible sizes of “large” minimal complements in finite groups.

Proposition 3.13. *Let G be a finite abelian group, and let $W \subseteq G$. If C is a minimal complement for W , then*

$$|C| \leq |G| \frac{|W|}{2|W| - 1}.$$

In particular, if $C \neq G$, then $|C| \leq 2|G|/3$.

Proof. The minimality of C guarantees that for each $c \in C$, there is an element $x(c) \in G$ such that $x(c) \in c + W$ but $x(c) \notin c' + W$ for any $c' \neq c$. (Note that these $x(c)$ ’s are pairwise distinct.) Let $X = \cup_{c \in C} \{x(c)\}$. Each of the remaining $|G| - |C|$ elements $g \in G \setminus X$ can be expressed as the sum of an element of C and an element of W in at most $|W|$ ways. We thus have

$$|C| \cdot |W| \leq |C| + (|G| - |C|)|W|,$$

and rearranging gives the result. $\square$

In particular, we see that any C of size $3|G|/5 < |C| \leq 2|G|/3$ can be a minimal complement only for a W of size 2, and it is possible to deduce an explicit structure theorem for minimal complements of this size. A set C is a minimal complement for some $W = \{w_1, w_2\}$ if and only if it is also a complement for $W - w_1 = \{0, w_1 - w_2\}$, so we may restrict our attention to sets W containing the identity. The following characterization of minimal complements of $W = \{0, a\}$ holds in all abelian groups G but is most interesting when G is finite.

Proposition 3.14. *Let G be an abelian group, and let $W = \{0, a\}$ for some non-zero $a \in G$. The set $C \subseteq G$ is a minimal complement for W if and only if for every $g \in G$, the set C contains at least one of g and $g + a$ but does not contain all of g , $g + a$, and $g + 2a$. This condition is equivalent to requiring that the intersection of C with any coset of the cyclic group generated by a neither miss any two consecutive elements nor contain any three consecutive elements.*

Proof. Let H be the subgroup generated by a . Since $W \subseteq H$, we know that C is a minimal complement for W if and only if the following holds for each coset $g + H$: first, $W + (C \cap (g + H)) = g + H$; and second, the

translate $(C \cap (g + H)) - g$ is a minimal complement for W in H . The first condition is satisfied if and only if C contains at least one of g and $g - a$ for every $g \in G$.

For the second condition, consider a complement D for W in H . If there exists $x \in H$ such that $x, x + a, x + 2a \in D$, then $D \setminus \{x + a\}$ is also a complement for W in H (which means that D is not minimal). Suppose, on the contrary, that no such x exists. Let $d \in D$. Then either $d + a$ or $d - a$ is not in D . In the former case, the sum $d + a$ is uniquely expressed in $D + W$; in the latter case, the sum $d + 0$ is uniquely expressed in $D + W$; so we conclude that D is minimal. This completes the proof. $\square$

Finally, we mention that short arithmetic progressions in finite groups are minimal complements. Recall from Proposition 3.13 that G does not contain any minimal complements of size strictly between $2|G|/3$ and $|G|$; the following proposition, together with this fact, tells us that the sizes of the minimal complements in $\mathbb{Z}/n\mathbb{Z}$ are exactly $1, 2, \dots, \lfloor 2n/3 \rfloor, n$.

Before proceeding to the proof, we introduce a useful piece of terminology. Suppose C is a complement for W in some group G . We say that an element $c \in C$ is *essential* (for this W) if $C \setminus \{c\}$ is no longer a complement for W . Clearly, C being a minimal complement for W is equivalent to every element $c \in C$ being essential.

Proposition 3.15. *Let G be an abelian group of order n , and let $C \subseteq G$ be an arithmetic progression of length k . Let H be the subgroup generated by C , and write $|H| = m$. The set C is a minimal complement if and only if*

$$k \leq \frac{2nm}{2n + m} \quad \text{or} \quad k = m.$$

Proof. The case $k = m$ is trivial, so we restrict our attention to $k < m$. We first treat the special case $G = H = \mathbb{Z}/n\mathbb{Z}$. Since the property of being a minimal complement is invariant under translation and automorphisms, we may assume that $C = \{0, 1, \dots, k - 1\}$.

If $k \leq n/2$, then we claim that C is a minimal complement for $W = \{0, k, k + 1, \dots, n - k\}$. Indeed, $W + C = \mathbb{Z}/n\mathbb{Z}$, and minimality follows from the observation that for every proper subset $D \subset C$, the sumset $D + W$ fails to cover some element of the interval $0, 1, \dots, k - 1$. If instead $n/2 < k \leq 2n/3$, then we claim that C is a minimal complement for $W = \{0, k\}$. Again, it is clear that $W + C = \mathbb{Z}/n\mathbb{Z}$, and minimality is only slightly harder. For $0 \leq r < n - k$, the element $k + r$ is uniquely represented in $W + C$ as $k + r$, which means that r is essential. For $n - k \leq r \leq k - 1$, the element r is uniquely represented in $W + C$ as $0 + r$, which again means that r is essential, so we conclude that C is minimal. Proposition 3.13 tells us that the bound $k \leq 2n/3$ is in fact tight.

We now treat the general case. Since H is clearly cyclic, we may identify it with $\mathbb{Z}/m\mathbb{Z}$. By the translation

invariance of minimal complements, we may assume that $C = \{0, 1, \dots, k-1\}$ as a subset of H . If $k \leq 2m/3$, then the previous paragraph ensures that C is a minimal complement in H . By Lemma 3.11, C is then also a minimal complement in G . We thus restrict our attention to $k > 2m/3$.

We construct W by carefully including two elements of each coset of H . Choose representatives $g_1 + H, \dots, g_{n/m} + H$ for G/H , so that $G = \cup_i (g_i + H)$. Note that

$$(W \cap (g_i + H)) + C \subseteq g_i + H.$$

Suppose $W_i = W \cap (g_i + H) = \{g_i, g_i + s_i\}$ for some choice of elements s_i . We ensure that $W_i + C = g_i + H$ (i.e., $\{0, s_i\} + C = H$) by picking s_i to satisfy $m - k \leq s_i \leq k$. Writing

$$\{0, s_i\} + C = \{0, 1, \dots, k-1\} \cup \{s_i, s_i + 1, \dots, s_i + k-1\},$$

we see that every $r \in C$ satisfying $s_i + k - m \leq r \leq s_i - 1$ or $k - s_i \leq r \leq m - s_i - 1$ is essential. Now, take $s_i = i(m - k)$ for every $1 \leq i \leq \left\lceil \frac{k}{2(m-k)} \right\rceil$ (where the latter quantity is at most n/m by the hypothesis on the size of k in the statement of the proposition), and set every remaining $s_i = 1$ for larger i . It is now easy to see that every $r \in C$ is essential, so we conclude that C is a minimal complement for this W , as desired.

Proposition 3.17 (proven in Section 3.4) gives the converse. $\square$

3.3 Small sets are minimal complements

We prove Theorem 3.1 by establishing the following more precise statement. Theorem 3.1 follows from setting $s = \lceil (3/2) \log n \rceil$. In the interest of readability, we have not optimized the bound in the theorem below as being more careful does not give a better exponent in Theorem 3.1. The proof is probabilistic and bears some resemblance to the argument in [2] that every sufficiently large subset of $\mathbb{Z}/p\mathbb{Z}$ is a sumset of the form $A + A$.

Theorem 3.16. *Let G be an abelian group of order n , and let $C \subseteq G$ be a non-empty subset of order k . If there exists a positive integer s such that*

$$\frac{s^2 k^3}{n} + \frac{e^s k^{3s}}{n^{s-1}} + k \left(\frac{s^2 k^3}{n} \right)^s < 1, \quad (3.1)$$

then C is a minimal complement in G .

Proof. Write $C = \{c_1, c_2, \dots, c_k\}$, and fix some s such that (3.1) holds. For each $i \in [k]$, choose s (not

necessarily distinct) elements $w_i^{(1)}, w_i^{(2)}, \dots, w_i^{(s)}$ uniformly and independently at random from G . For each $i \in [k]$ and $p \in [s]$, set

$$g_i^{(p)} = w_i^{(p)} + c_i.$$

We will find a subset W for which C is a minimal complement as follows. First, for each i , we will choose one element $w_i^{(p)}$ to include in W (as described below). Next, we will include in W every $w \in G$ such that $w + C$ does not contain any of the s chosen elements $g_i^{(p)} = w_i^{(p)} + c_i$. We will have chosen the $w_i^{(p)}$'s so as to make the corresponding $g_i^{(p)}$'s pairwise distinct; this condition will guarantee the minimality of C because $W + (C \setminus \{c_i\})$ does not contain $g_i^{(p)}$. The main part of the proof is showing that $W + C = G$ with strictly positive probability.

Let E_1 be the event that there are distinct pairs (i, p) and (j, q) so that $g_i^{(p)} \in w_j^{(q)} + C$. There are fewer than $(sk)^2$ choices for the pairs (i, p) and (j, q) , and for each fixed choice the probability that $g_i^{(p)} \in w_j^{(q)} + C$ is exactly k/n , so

$$\mathbb{P}(E_1) \leq \frac{s^2 k^3}{n}. \quad (3.2)$$

Next, let E_2 be the event that there exist an element $z \in G$ and s distinct pairs (i, p) such that each $(g_i^{(p)} - C) \cap (z - C) \neq \emptyset$. There are

$$n \cdot \binom{sk}{s} \leq n(ek)^s$$

choices for the element z and the s distinct pairs of indices (i, p) . For each such choice, the probability that $(g_i^{(p)} - C) \cap (z - C) \neq \emptyset$ for all of the chosen $g_i^{(p)}$'s is at most $(k^2/n)^s$: the choices of the elements $g_i^{(p)}$ are independent, and for each fixed (i, p) the probability of such a non-empty intersection is the probability that $g_i^{(p)} \in z - C + C$, where this set has cardinality smaller than k^2 . Therefore

$$\mathbb{P}(E_2) \leq \frac{e^s k^{3s}}{n^{s-1}}. \quad (3.3)$$

Now, order the elements of each translate $z - C$ of C according to the numbering of the elements of C , that is, $z - c_1, z - c_2, \dots, z - c_k$. Let E_3 be the event that there is an index $i \in [k]$ such that the following holds: for each $p \in [s]$, there are $z \in G$, $j \neq i$ in $[k]$, and $q \in [s]$ such that both

$$\frac{k}{s} < \left| (g_i^{(p)} - C) \cap (z - C) \right| < k \quad (3.4)$$

and

$$(g_j^{(q)} - C) \text{ contains the first element of } (z - C) \setminus (g_i^{(p)} - C). \quad (3.5)$$

There are k possibilities for the index i . Fix such an i , and expose the random elements $g_j^{(q)}$ for all $j \neq i$ and for all q . The union of all such translates $g_j^{(q)} - C$ (call it Y) has cardinality smaller than sk^2 . For each index p and each choice of the element $g = g_i^{(p)}$, let $Z = Z(g)$ denote the set of elements $z \in G$ for which (3.4) holds. For each $z \in Z(g)$, let $u = u(g)$ be the first element of $z - C$ (in the ordering fixed above) not covered by $g - C$, and let $U = U(g)$ be the set of all such first elements as z ranges over $Z(g)$. When $z \in G$ is chosen uniformly at random, the expected size of $(g - C) \cap (z - C)$ is k^2/n , so Markov's Inequality tells us that $Z(g)$ (and hence also $U(g)$) has size at most sk . Note that for any $g_1, g_2 \in G$, we have the shifts $Z(g_2) = Z(g_1) + g_2 - g_1$ and $U(g_2) = U(g_1) + g_2 - g_1$.

The existence of $j \neq i$, z and q satisfying (3.4) and (3.5) for the index p is equivalent to the condition that the random set $U(g_i^{(p)})$ intersects the set Y . Since Y is a fixed set of size at most sk^2 (the elements $g_j^{(q)}$ already having been exposed) and $U(g_i^{(p)})$ is a random shift of a set of size at most sk , the probability that they intersect is at most s^2k^3/n . The events corresponding to distinct indices p are independent, so the probability of finding such an intersection for all values of $p \in [s]$ is at most $(s^2k^3/n)^s$. As there are k choices for i , we have

$$\mathbb{P}(E_3) \leq k \left(\frac{s^2k^3}{n} \right)^s. \quad (3.6)$$

It follows from the assumption (3.1) and the inequalities (3.2), (3.3), and (3.6) that with positive probability none of these three events holds. Fix a choice of the random elements $w_i^{(p)}$ for which none of these events holds. For each $i \in [k]$, pick a $p \in [s]$ such that no choice of $z \in G$, $j \neq i$ in $[k]$ and $q \in [s]$ satisfies (3.4) and (3.5). (Such a p exists because the event E_3 does not occur). We now define W to contain these k elements $w_i^{(p)}$, together with every element $w \in G$ that is not contained in the union of the chosen $g_i^{(p)} - C$'s (i.e., the translates $w_i^{(p)} + c_i - C$). Since the event E_1 does not occur, we see that each $g_i^{(p)}$ is uniquely expressible in $W + C$ as $w_i^{(p)} + c_i$, which (as discussed above) guarantees the minimality of C . It remains only to show that $W + C = G$.

Let $z \in G$ be an arbitrary element. We have to show that W intersects $z - C$. If $z - C = g_i^{(p)} - C$ for one of our k chosen elements $g_i^{(p)}$, then

$$w_i^{(p)} = g_i^{(p)} - c_i \in g_i^{(p)} - C = z - C,$$

as needed. Otherwise, no single translate $g_i^{(p)} - C$ covers $z - C$. We will show that in this case, the k sets $g_i^{(p)} - C$ do not cover $z - C$. Since the event E_2 does not occur, there are at most s sets $g_i^{(p)} - C$ that intersect $z - C$. We are done unless there is some particular $g_i^{(p)} - C$ that covers at least k/s elements of $z - C$. But then, because the event E_3 does not occur, we see that there is no $g_j^{(q)}$ such that $g_j^{(q)} - C$ covers

the first element of $(z - C) \setminus (g_i^{(p)} - C)$, so $z - C$ is not covered by the translates $g_i^{(p)} - C$. We conclude that $z \in W + C$. Since this holds for every $z \in G$, we have $W + C = G$, which completes the proof. $\square$

We now deduce Theorem 3.2 from Theorem 3.1.

Proof of Theorem 3.2. Let G be an infinite abelian group, and let $C \subset G$ be a finite subset. First, suppose G is a torsion group. Then we can find a finite subgroup H containing C such that $|H| > 100|C|^4$. By Theorem 3.1, C is a minimal complement in H ; Lemma 3.11 then tells us that C is also a minimal complement in G .

Second, suppose G is not a torsion group, and let $x \in G$ be an element of infinite order. Then the subgroup H generated by the set $C \cup \{x\}$ is finitely generated and has infinite order; write $H = K \times \mathbb{Z}^r$, where K is a finite group and $r \geq 1$. Choose some $M > 100|C|^4$ such that C is contained in the “cube” $K \times [-M, M]^r$, and let $\pi : H \rightarrow K \times (\mathbb{Z}/2M\mathbb{Z})^r$ be the canonical projection map. Note that π is injective on C . By Theorem 3.1, $\pi(C)$ is a minimal complement in $K \times (\mathbb{Z}/2M\mathbb{Z})^r$; by Lemma 3.12, C is a minimal complement in H ; by Lemma 3.11, C is a minimal complement in G , as desired. $\square$

We remark that one can directly prove a slightly stronger version of Theorem 3.2 by performing the construction of the proof of Theorem 3.1 in the infinite group and using transfinite induction. Taking $s = 1$, an examination of the proof shows that one can iteratively construct W by adding $c_i + w_i$ for each $c_i \in C$, and then adding all other elements that do not destroy the minimality of C . When we try to find the next w_i to add to W , we make sure that the events E_1, E_2 , and E_3 do not occur. This will be true if the chosen w_i is not contained in a certain union of translates of $C - C$ (determined by the elements that we already added to W), and the number of such translates is less than $|C|$ or is finite. We can therefore find an appropriate w_i as long as $S + C - C \neq G$, for certain subsets S , with the cardinality of S equal to the number of elements of W already chosen. Therefore this procedure will succeed as long as $S + C - C$ is a proper subset of G whenever S has cardinality strictly smaller than that of C or is finite. See also the proof of Theorem 3.6.

3.4 Many large sets are not minimal complements

We now turn to the problem of finding upper bounds on $T(G)$. Our first result shows that $T(n) = O(\sqrt{n})$ under certain divisibility conditions on n . We then proceed with the probabilistic proof of Theorem 3.3, which extends to all values of n at the cost of increasing the exponent from $1/2$ to $3/4 + \varepsilon$.

Proposition 3.17. *Let G be an abelian group of order n with a subgroup H of order m , and let k be an*

integer satisfying

$$\frac{2nm}{m+2n} < k < m.$$

Then no subset $C \subset H$ of size k is a minimal complement in G .

Proof. Assume (for the sake of contradiction) that C is a minimal complement for some W in G . Since $C + W = G$ and C is properly contained in H , we see that W must contain at least two elements of each of the n/m cosets of H . Consider a coset $g + H$, and let w_1, w_2 be distinct elements of $W \cap (g + H)$. The translates $w_1 + C$ and $w_2 + C$ each have size k and are contained in $g + H$. In the sum $\{w_1, w_2\} + C$, each element of $g + H$ is represented at most twice, so (since there are exactly $2k$ total sums) at least $2k - m$ elements are represented twice. It follows that in the sum

$$(W \cap (g + H)) + C = g + H,$$

there are at most $2m - 2k$ elements that are uniquely represented. Since this is true of each of the n/m cosets of H , there are at most

$$\frac{2n(m-k)}{m} < k$$

elements of G that are uniquely represented in $W + C$. Thus, fewer than k elements of C are essential, which contradicts the minimality of C . $\square$

Corollary 3.18. *We have that $\liminf_{n \rightarrow \infty} \frac{T(n)}{\sqrt{n}} \leq \sqrt{2}$.*

Proof. Let $G = \mathbb{Z}/n\mathbb{Z}$, where $n = (k+1)(\lceil k/2 \rceil - 1)$. Then the element $\lceil k/2 \rceil - 1$ generates a subgroup of order $m = k+1$, and Proposition 3.17 implies that $T(\mathbb{Z}/n\mathbb{Z}) \leq k$. $\square$

We now transition to the proof of Theorem 3.3. Let G be an abelian group of order n , and let A be a random subset of G obtained by including each $g \in G$ independently with probability $p = n^{-1/4+\varepsilon}$. In the following discussion, we say that an event holds *with high probability (whp)* if it holds with probability tending to 1 as n tends to infinity.

Lemma 3.19. *Let G and A be as above. Then whp $|A| = (1 + o(1))n^{3/4+\varepsilon}$ and every complement B for A (that is, a subset $B \subseteq G$ with $A + B = G$) satisfies*

$$|B| > \frac{1}{4p} \log n = \frac{1}{4} n^{1/4-\varepsilon} \log n.$$

Proof. The fact that $|A| = (1 + o(1))n^{3/4+\varepsilon}$ whp is trivial. Fix a subset B with $m = \frac{1}{4p} \log n$ elements. There is a subset $T \subseteq G$ with at least $n/m^2 = \tilde{\Theta}(n^{1/2+2\varepsilon})$ elements such that $t_1 - B$ and $t_2 - B$ are disjoint

for all distinct elements $t_1, t_2 \in T$ (i.e., $(T - T) \cap (B - B) = \{0\}$). For each fixed $t \in T$, the probability of $t - B$ being disjoint from A is precisely $(1 - p)^m \geq n^{-1/4+o(1)}$. Thus, the probability that $t \in B + A$ (i.e., $t - B$ intersects A) is at most $1 - n^{-1/4+o(1)}$. Since the events corresponding to distinct elements of T are mutually independent, the choice of T ensures that $T \subseteq A + B$ with probability at most

$$(1 - n^{-1/4+o(1)})^{|T|} \leq e^{-n^{1/4+\varepsilon}}.$$

This is clearly an upper bound for the probability that $A + B = G$. As there are only $\binom{n}{m} = e^{\tilde{\Theta}(n^{1/4-\varepsilon})}$ sets B of size m , it follows that whp there is no such B satisfying $A + B = G$. $\square$

Lemma 3.20. *Let G and A be as above. Then whp every subset $D \subseteq G$ of size*

$$|D| \geq (1 - \varepsilon) \frac{1}{4p} \log n = (1 - \varepsilon) \frac{1}{4} n^{1/4-\varepsilon} \log n$$

satisfies $|G \setminus (A + D)| \leq n^{3/4+\varepsilon/2}$.

Proof. It clearly suffices to prove the statement for every D of cardinality exactly $(1 - \varepsilon) \frac{1}{4p} \log n$. Fix such a set D . Let H be the Cayley graph of G with respect to $(D - D) \setminus \{0\}$: the graph on the vertex set G where distinct $g_1, g_2 \in G$ are adjacent if and only if there are $d_1, d_2 \in D$ such that $g_1 = g_2 + d_1 - d_2$ (i.e., the sets $g_1 - D$ and $g_2 - D$ intersect). The maximum degree of this graph is smaller than $|D|^2$. Therefore, by a well known theorem of Hajnal and Szemerédi [53] (which is convenient to use here but can also be avoided), the graph H has a proper coloring with at most $|D|^2 = \tilde{\Theta}(n^{1/2-2\varepsilon})$ colors where the color classes differ in size by at most 1.

For each color class C , the number of elements $t \in C$ that are not covered by $A + D$ is a binomial random variable with parameters $|C| = \tilde{\Theta}(n^{1/2+2\varepsilon})$ and $q = (1 - p)^{|D|} = n^{-1/4(1-\varepsilon+o(1))}$. By the standard estimates for binomial distributions (see, e.g., [3], Appendix A), the probability the value of this random variable exceeds its expectation by a factor of 2 is exponentially small in the expectation, which in turn is larger than $n^{1/4+\varepsilon}$. Therefore, whp this does not happen for any of the color classes and any choice of D , as there are only $e^{\tilde{\Theta}(n^{1/4-\varepsilon})}$ choices for D . But this means that whp $|G \setminus (A + D)| \leq n^{3/4+\varepsilon/2}$ holds for every such D . $\square$

Using these two lemmas, we proceed with the proof of Theorem 3.3.

Proof of Theorem 3.3. Let A be a set of size $(1 + o(1))n^{3/4+\varepsilon}$ satisfying the conclusions of Lemma 3.19 and Lemma 3.20. We claim that A is not a minimal complement. Assume (for contradiction) that A is a minimal complement for some subset $B \subseteq G$. Lemma 3.19 tells us that $|B| \geq \frac{1}{4p} \log n$. Partition B arbitrarily into

$t = \lceil 1/\varepsilon \rceil$ (pairwise disjoint) parts $B = B_1 \cup B_2 \cup \dots \cup B_t$ of equal cardinalities (up to a difference of 1). For each $i \in [t]$, let $D_i = B \setminus B_i$. Then each D_i has cardinality

$$|D_i| \geq (1 - \varepsilon)|B| \geq (1 - \varepsilon) \frac{1}{4p} \log n.$$

Since A satisfies the conclusion of Lemma 3.20, each of the t sets $A + D_i$ covers all but at most $n^{3/4+\varepsilon/2}$ of the elements of G . So in the sum $A + B$, all but at most $tn^{3/4+\varepsilon/2} < |A|$ of the elements of G are covered by every set $A + D_i$ and are hence covered at least twice. In particular, there are strictly fewer than $|A|$ elements that are covered exactly once, so not every element of A can be essential, and A is not in fact minimal as a complement for B . This contradiction completes the proof. $\square$

3.5 Maximal supplements

Problems about large supplements are “dual” to problems about small complements in a natural way: given a subset W of a finite abelian group, the problem of finding the minimum size of a complement for W and the problem of finding the maximum size of a supplement for W are dual as integer programs. Of course, not all minimal complements have the minimum size, and not all maximal complements have the maximum size, but this connection nonetheless motivates studying the two problems together.

We can characterize being a maximal supplement as follows. C is a supplement to W if and only if $c_1 + w_1 = c_2 + w_2$ has only the trivial solutions for $c_1, c_2 \in C$, $w_1, w_2 \in W$. In particular, this condition is equivalent to $(C - C) \cap (W - W) = \{0\}$. The maximality condition on C is that for every $d \in G \setminus C$, the translate $d - C$ intersects $W - W$ nontrivially, i.e., $d \in C + W - W$. Since $C \subseteq C + W - W$ trivially, this maximality condition can be expressed as $C + W - W = G$. So, putting everything together, we have that C is a maximal supplement in G if and only if there is a subset W satisfying

$$(C - C) \cap (W - W) = \{0\} \quad \text{and} \quad C + W - W = G.$$

One might ask if inverse results about minimal complements carry over to the setting of maximal supplements. In particular, one might ask if every finite subset of $\mathbb{Z}$ is a maximal supplement and if every sufficiently small subset of a finite abelian group is a maximal supplement. The answer to each of these questions is “no”, for a simple reason. Recall that a subset C of an abelian group G is *solid* if C is not properly contained in any set D such that $D - D = C - C$. We show that every maximal supplement is solid.

Proposition 3.21. *Let G be an abelian group. If the subset $C \subseteq G$ is a maximal supplement, then C must be solid.*

Proof. Assume (for the sake of contradiction) that C is a maximal supplement to some W but C is not solid, and let D be a set properly containing C such that $D - D = C - C$. Then

$$(D - D) \cap (W - W) = (C - C) \cap (W - W) = \{0\}$$

shows that D is also a supplement to W , which contradicts the maximality of C . $\square$

Now, we may ask if solidity is sufficient for a set to be a maximal supplement. For finite subsets of infinite groups, the answer is “yes”. We first require a lemma that establishes a sufficient condition for C to be a maximal supplement for a particular W .

Lemma 3.22. *Let G be an abelian group, and let $C \subseteq G$ be a non-empty solid subset. If the subset $W \subseteq G$ satisfies*

$$G \setminus (W - W) = (C - C) \setminus \{0\},$$

then C is a maximal supplement for W .

Proof. It is immediate that

$$(C - C) \cap (W - W) = \{0\}.$$

Suppose (for the sake of contradiction) that $C + W - W \neq G$, and let $x \in G \setminus (C + W - W)$. Then $x - C$ is disjoint from $W - W$, which means that

$$x - C \subseteq (C - C) \setminus \{0\}.$$

Note that $x \notin C$ since $x - C$ cannot contain the element 0. Now, the set $D = C \cup \{x\}$ witnesses that C is not solid:

$$(C \cup \{x\}) - (C \cup \{x\}) \subseteq (C - C) \cup (x - C) \cup (C - x) \subseteq C - C.$$

This contradiction lets us conclude that in fact C is a maximal supplement to W . $\square$

We now prove the promised result about finite subsets of infinite groups.

Theorem 3.23. *thm:suppinfinite Let G be an infinite abelian group, and let $C \subset G$ be a finite non-empty solid subset. Then C is a maximal supplement in G .*

Proof. We construct a set W for which C is a maximal supplement. Using the well-ordering principle, fix a bijection of $(G \setminus (C - C)) \cup \{0\}$ with the minimal ordinal of cardinality $|G|$. This gives an ordering $\{g_\lambda\}_{\lambda \in \Lambda}$ on $(G \setminus (C - C)) \cup \{0\}$ where 0 is the minimal element of Λ and the set $\{\alpha : \alpha < \lambda\}$ has cardinality strictly smaller than that of G for every $\lambda \in \Lambda$. We may take $g_0 = 0$. We now use transfinite induction to construct a sequence of (increasing) nested subsets $\{W_\lambda\}_{\lambda \in \Lambda}$ such that each W_λ satisfies

$$(C - C) \cap (W_\lambda - W_\lambda) = \{0\} \quad \text{and} \quad g_{\lambda'} \in W_\lambda - W_\lambda \text{ for all } \lambda' < \lambda.$$

Then it is clear by Lemma 3.22 that C is a maximal supplement for

$$W = \bigcup_{\lambda \in \Lambda} W_\lambda,$$

since $G \setminus (W - W) = (C - C) \setminus \{0\}$. For the base cases, set $W_0 = W_1 = \{0\}$. For a limit element α , we set $W_\alpha = \bigcup_{\lambda < \alpha} W_\lambda$. For a successor element $\lambda + 1$, suppose we already have W_λ . We wish to find some $x \in G$ such that we may “safely” define

$$W_{\lambda+1} = W_\lambda \cup \{x, x + g_\lambda\};$$

in order to ensure that $W_{\lambda+1} - W_{\lambda+1}$ intersects $C - C$ only at 0, it suffices to check that $W_\lambda - \{x, x + g_\lambda\} = W_\lambda - \{0, g_\lambda\} - \{x\}$ is disjoint from $C - C$ (since $C - C$ is symmetric). Indeed, the induction hypothesis ensures that $W_\lambda - W_\lambda$ intersects $C - C$ only at 0, and the elements $\pm g_\lambda$ of $W_{\lambda+1} - W_{\lambda+1}$ are not in $C - C$ by assumption. So we must find x such that

$$x \notin W_\lambda - \{0, g_\lambda\} - C + C.$$

Since $C - C$ is finite, the set $W_\lambda - \{0, g_\lambda\} - C + C$ is finite if W_λ is finite and has cardinality equal to W_λ if W_λ is infinite. As the cardinality of W_λ is strictly smaller than that of G , there always exists some such choice of x . $\square$

This argument actually gives the following slightly more general statement.

Theorem 3.24. *Let G be an infinite abelian group, and let $C \subset G$ be a non-empty solid subset such that $G \neq S + C - C$ whenever $|S| < |G|$. Then C is a maximal supplement in G .*

As was the case for additive complements, the finite group setting is more complex. Our next aim is showing that every sufficiently small solid subset of a finite abelian group is a maximal supplement. We again apply Lemma 3.22. An old result of Alon [2] shows that every sufficiently large symmetric subset containing

0 of a finite abelian group is a difference set of the form $A - A$. The following appears as Theorem 4.1 in [2]. (As discussed in that paper, the bound is essentially tight.)

Theorem 3.25 (Alon [2]). *There exists an absolute constant $c > 0$ such that the following holds: Let G be an abelian group of order n , and let $V = -V$ be a symmetric subset of G containing 0. If*

$$|V| \geq n - c\sqrt{n/\log n},$$

then there is some subset $A \subseteq G$ such that $V = A - A$.

Combining these two pieces shows that every sufficiently small solid subset is a maximal supplement.

Proof of Theorem 3.5. We have $|C - C| \leq b^2\sqrt{n/\log n}$, so that $V = (G \setminus (C - C)) \cup \{0\}$ is a symmetric subset containing 0 and $|V| \geq n - b^2\sqrt{n/\log n}$. By Theorem 3.25, a sufficiently small choice of b guarantees the existence of some $W \subseteq G$ such that $V = W - W$. Since this W satisfies the conditions of Lemma 3.22, we conclude that C is a maximal supplement for W . $\square$

As in the proof of Theorem 3.2, it is possible to deduce Theorem 3.6 from Theorem 3.5 by showing that the property of being a maximal supplement lifts from certain subquotients.

Chapter 4

Graham's rearrangement conjecture

4.1 Introduction

4.1.1 Main result

Let A be a finite subset of an abelian group. We say that an ordering $a_1, \dots, a_{|A|}$ of A is *valid* if the partial sums $a_1, a_1 + a_2, \dots, a_1 + a_2 + \dots + a_{|A|}$ are all distinct. In 1971, Graham conjectured that every set of non-zero elements of $\mathbb{F}_p$ has a valid ordering.

Conjecture 4.1 ([45]). *Let p be a prime. Then every subset $A \subseteq \mathbb{F}_p \setminus \{0\}$ has a valid ordering.*

This conjecture also appeared in a 1980 book of Erdős and Graham [38], and a very similar conjecture for finite cyclic groups is due to Alspach (see [16]).

The main avenue of attack on Graham's conjecture has been to show that its conclusion holds when A is small. Until recently, the published world record had established Graham's conjecture for sets A of size at most 12 (see, e.g., the discussion in [27, 66]). Earlier this year, the author [66] used a simple rectification argument to show that Graham's conjecture holds for all sets A of size $|A| \leq \log p / \log \log p$; Will Sawin [99] had independently proven a comparable bound, using roughly similar ideas, in a 2015 MathOverflow post. The purpose of the present chapter is to prove Graham's conjecture for sets A of up to quasi-polynomial size.

Theorem 4.2. *The following holds for every constant $c > 0$. Let p be a large prime. Then every subset $A \subseteq \mathbb{F}_p \setminus \{0\}$ of size*

$$|A| \leq e^{c(\log p)^{1/4}}$$

has a valid ordering.

We have not made a serious effort to optimize the exponent $1/4$, but the quasi-polynomial shape of this bound does appear as a natural barrier in several parts of our argument. We also mention that the conclusion of Theorem 4.2 still holds, with a nearly identical proof, if $\mathbb{F}_p$ is replaced by any abelian group with no non-zero elements of order strictly smaller than p .

The proof strategy for Theorem 4.2 is motivated by the argument in [66]. (The argument in [99] seems less well-suited to generalization.) Two new ingredients are the theory of dissociated sets (from additive combinatorics) and probabilistic tools. One of our intermediate results (see Theorem 4.6 below) is a structure theorem involving dissociated sets, which may be of independent interest.

Let us say a few more words about related work. One natural approach is establishing Graham’s conjecture for sets A with size very close to p . See [56] and the references therein for early results on the case where $|A| = p - O(1)$. More recently, work of Müyesser and Pokrovskiy [78] on the Hall–Paige Conjecture has proven Graham’s conjecture for set A of size $|A| \geq p - p^c$ for some explicit constant $c > 0$. In a different direction, Bucić, Frederickson, Müyesser, Pokrovskiy, and Yepremyan [21] have established an “approximate” version of Graham’s conjecture, namely, that for any nonempty subset $A \subseteq \mathbb{F}_p \setminus \{0\}$, there is a subset $A' \subseteq A$ of size $|A'| = (1 - o(1))|A|$ which has a valid ordering. (The $o(1)$ asymptotic is with respect to $p \rightarrow \infty$.)

4.1.2 Proof sketch and organization

We say that an ordering of A is *two-sided valid* if no proper nonempty subinterval sums to zero; this condition is slightly stronger than A being valid. As in [66], we will prove Theorem 4.2 with two-sided valid orderings.

Let us briefly recall the main ideas of [66]. Let $A \subseteq \mathbb{F}_p \setminus \{0\}$ be a subset of size $|A| \leq \log p/2 \log \log p$. Using the pigeonhole principle, one can find some $\lambda \in \mathbb{F}_p^\times$ such that the dilate $\lambda \cdot A$ is contained in the interval $(-p/|A|, p/|A|)$. Since sums of elements of $\lambda \cdot A$ have no “wrap-around”, we can interpret $\lambda \cdot A$ as a subset of $\mathbb{Z} \setminus \{0\}$; this process is known as “rectification”. Finally, in the integer setting, one can use induction on $|A|$ to find a two-sided valid ordering in which all of the positive elements appear before all of the negative elements.

Our proof of Theorem 4.2 proceeds in four main steps. The first step is showing that every subset of $\mathbb{F}_p$ can be decomposed into a union of large dissociated sets and a rectifiable residual set. (A dissociated set is a set all of whose subset sums are distinct; see below.) The residual set can be broken into “positive” and “negative” sets. We will aim to find a two-sided valid ordering consisting of the positive elements, then the elements of the dissociated sets, then the negative elements.

The second step is ordering the positive and negative elements. Following [66], we inductively construct these orderings in order to avoid zero-sum intervals that begin in the positive region and end in the negative

region. We take advantage of some flexibility in the argument from [66] in order to prepare for “potential” zero-sum intervals with one endpoint in the positive region or negative region and the other endpoint very close to one of the edges of the dissociated region.

The third and fourth steps concern ordering the elements of the dissociated sets. The main idea is that in a uniformly random ordering of a dissociated set of size R , the sum of the first k elements is uniformly distributed on $\binom{R}{k}$ different values. Since the probability of assuming any particular value is very small, the probability of this initial segment forming the end of a zero-sum interval is also very small. This naïve random strategy essentially works for handling sets A of size up to $(\log p)^{3/2}$ (which breaks the “rectification barrier” of [66]), but we must employ a more elaborate random procedure in order to reach the threshold in Theorem 4.2. In particular, it becomes important to distinguish between the “borders” and “interiors” of the orderings of the dissociated sets. The third step of the proof is randomly splitting and then reordering the dissociated sets, and the fourth step is choosing a (suitably) random ordering for the elements within each dissociated set.

We carry out these four steps in Sections 4.3, 4.4, 4.5, and 4.6, respectively, and then we make some concluding remarks and pose several open problems in Section 4.7.

4.2 Notation and parameters

Before jumping into the proofs, we set a few pieces of notation.

- We use $\mathbb{F}_p$ to denote the field with p elements, and for us p will always be a large prime.
- We denote dilation by $\lambda \cdot A := \{\lambda a : a \in A\}$.
- We denote the restricted sumset by $B \hat{+} B := \{b + b' : b, b' \in B \text{ and } b \neq b'\}$.
- Let $\sum_{=M}(S) := \{\sum_{s \in S'} s : S' \subseteq S, |S'| = M\}$ denote the set of all sums of exactly M elements of S . Likewise, let $\sum_{\leq M}(S) := \{\sum_{s \in S'} s : S' \subseteq S, |S'| \leq M\}$ denote the set of all sums of at most M elements of S , and let $\sum_{\geq M}(S) := \{\sum_{s \in S'} s : S' \subseteq S, |S'| \geq M\}$ denote the set of all sums of at least M elements of S .
- For a sequence $\mathbf{b} = b_1, \dots, b_r$, let $\text{IS}(\mathbf{b}) := \{b_1 + \dots + b_j : 0 \leq j \leq r\}$ denote the set of initial segment sums of $\mathbf{b}$, and let $\overline{\mathbf{b}} := b_r, \dots, b_1$ denote the reverse of $\mathbf{b}$.

We use standard asymptotic notation. We write $f = O(g)$ or $f \ll g$ if there is a universal constant $C > 0$ such that $|f| \leq Cg$. If f is non-negative and $f = O(g)$, then we also write $g = \Omega(f)$. We write $f \asymp g$ when $f \ll g$ and $g \ll f$. Finally, we write $f(p) = o(g(p))$ if $\lim_{p \rightarrow \infty} f(p)/g(p) = 0$.

When there is no risk of confusion, we sometimes omit floor functions in calculations for typographical clarity.

Let us also record a few parameters that we will carry through our proofs.

- Our set A will have size $|A| \leq e^{c(\log p)^{1/4}}$ for some absolute constant $c > 0$.
- We define the *rectification threshold* for a subset $A \subseteq \mathbb{F}_p$ to be

$$R = R(A) := c_1 \max \left((\log p)^{1/2}, \frac{\log p}{\log |A|} \right),$$

where $c_1 > 0$ is a sufficiently small absolute constant.

- The *border width* is $K := c_2 R^{1/3}$, for yet another absolute constant $c_2 > 0$.
- We will use s (and later u) to denote the number of dissociated sets in our decomposition of the set A . The precise values of s, u , which are of no importance (besides the trivial bound $s, u \leq |A|$), will vary over the course of the proofs.
- We will always use δ_j to denote the sum of the elements of the set D_j , and we will always use τ_j to denote the sum of the elements of the set T_j ; when applicable, we will also write $\delta := \sum_j \delta_j$.

Finally, we reiterate that a sequence $b_1, \dots, b_t$ is *two-sided valid* if

$$b_i + \dots + b_j \neq 0 \quad \text{for all } 1 \leq i < j \leq t \text{ with } (i, j) \neq (1, t).$$

4.3 Structure theorem

Let G be an abelian group. A subset $D = \{d_1, \dots, d_r\} \subseteq G$ is *dissociated* if

$$\epsilon_1 d_1 + \dots + \epsilon_r d_r \neq 0 \quad \text{for all } (\epsilon_1, \dots, \epsilon_r) \in \{-1, 0, 1\}^r \setminus \{(0, \dots, 0)\}.$$

Equivalently, D is dissociated if all of the $2^{|D|}$ subset sums of D are distinct. The *dimension* of a subset $B \subseteq G$, written $\dim(B)$, is the size of the largest dissociated set contained in B . One should think of sets of small dimension as being highly “constrained”.

Lemma 4.3. *Let $B \subseteq G$ be a finite subset of an abelian group. If D is a maximal dissociated subset of B , then*

$$B \subseteq \text{span}(D) := \left\{ \sum_{d \in D} \epsilon_d d : \epsilon_d \in \{-1, 0, 1\} \right\}.$$

Proof. The maximality of D ensures that for every element $b \in B \setminus D$, the set $\{b\} \cup D$ is not dissociated; rearranging then gives the desired expression for b as an element of $\text{span}(D)$. $\square$

The following lemma says that sets of sufficiently small dimension can always be “rectified”. To make this precise, we define for each (nonempty) subset $A \subseteq \mathbb{F}_p$ the parameter

$$R = R(A) := c_1 \max \left((\log p)^{1/2}, \frac{\log p}{\log |A|} \right), \quad (4.1)$$

where c_1 is a sufficiently small absolute constant.

Lemma 4.4. *If $B \subseteq \mathbb{F}_p$ is a nonempty subset of dimension $\dim(B) < R = R(B)$, then there is some $\lambda \in \mathbb{F}_p^\times$ such that the dilate $\lambda \cdot B$ is contained in the interval $(-\frac{p}{100|B|}, \frac{p}{100|B|})$.*

Proof. Let D be a maximal dissociated subset of B , so that $|D| = \dim(B)$. Lemma 4.3 tells us that $B \subseteq \text{span}(D)$. Consider the set

$$\{(\lambda d/p)_{d \in D} : \lambda \in \mathbb{F}_p^\times\} \subseteq (\mathbb{R}/\mathbb{Z})^{\dim(B)}.$$

The pigeonhole principle provides some distinct $\lambda_1, \lambda_2 \in \mathbb{F}_p^\times$ such that $\|\lambda_1 d/p - \lambda_2 d/p\|_{\mathbb{R}/\mathbb{Z}} \leq p^{-1/\dim(B)}$ for all $d \in D$. Set $\lambda := \lambda_1 - \lambda_2 \in \mathbb{F}_p^\times$, so that $\lambda d \in [-p^{1-1/\dim(B)}, p^{1-1/\dim(B)}]$ for all $d \in D$. Since $B \subseteq \text{span}(D)$, we have

$$\lambda \cdot B \subseteq [-\dim(B)p^{1-1/\dim(B)}, \dim(B)p^{1-1/\dim(B)}].$$

It remains only to show that $\dim(B)p^{1-1/\dim(B)} < p/(100|B|)$, i.e., that $100|B|\dim(B) < p^{1/\dim(B)}$, as long as c_1 is chosen to be sufficiently small. When $\log |B| < (\log p)^{1/2}$, this inequality follows from $\dim(B) \leq R = c_1 \log p / \log |B|$. When $\log |B| \geq (\log p)^{1/2}$, the desired inequality follows from $\dim(B) \leq R = c_1 (\log p)^{1/2}$ and $|B| \leq 3^{\dim(B)}$. $\square$

Remark 4.5. For applications in this chapter, we will always work with sets of size at most $e^{c(\log p)^{1/4}}$, in which case the previous lemma says that every set B of size at most $e^{c(\log p)^{1/4}}$ with $\dim(B) < c_1 (\log p)^{3/4}$ is rectifiable. We opted to prove Lemma 4.4 for arbitrary sets $B \subseteq \mathbb{F}_p$, however, so that we could state the structural results in the rest of this section in full generality. These results are nontrivial for sets B of all sizes since the rectification threshold always satisfies $R(B) \gg (\log p)^{1/2}$.

We can combine these two lemmas to obtain a decomposition of *any* subset of $\mathbb{F}_p$ into large dissociated sets and a residual set that (after suitable dilation) is contained in a small interval around 0. We shall from

now on simply write R for $R(A)$. The following theorem bears many similarities to an argument of Bourgain [19] from a different context.

Theorem 4.6. *Every subset $A \subseteq \mathbb{F}_p$ can be partitioned as*

$$A = D_1 \cup \dots \cup D_s \cup E, \quad (4.2)$$

where the following holds:

- (i) each D_j is a dissociated set of size $|D_j| \asymp R$;
- (ii) $|E| \geq R/2$ if $s > 0$;
- (iii) there is some $\lambda \in \mathbb{F}_p^\times$ such $\lambda \cdot (E \cup \{\delta\}) \subseteq (-\frac{p}{90(|E|+1)}, \frac{p}{90(|E|+1)})$, where $\delta := \sum_{j=1}^s \sum_{d \in D_j} d$ is the sum of all of the elements in the dissociated sets.

Proof. Start with $E = A$. As long as $\dim(E) \geq R$, iteratively remove a dissociated subset of size $R/2$, so that at each step the set E of remaining elements has size $|E| \geq R/2$. Once we reach a residual set E of dimension smaller than R , Lemma 4.4 (applied to $E \cup \{\delta\}$, where δ is the sum of all of the dissociated elements removed) provides the desired $\lambda \in \mathbb{F}_p^\times$. $\square$

We will, of course, apply this theorem to the set A for which we are trying to find a two-sided valid ordering. If the number s of dissociated sets happens to be 0, then the entire set A is rectifiable and therefore has a two-sided valid ordering by [66] (see the discussion in the proof sketch). Thus, we will restrict our attention to the case where $s \geq 1$ (so in particular $|E| \gg R$ from (ii)). The presence of a large dissociated set allows us to obtain a more detailed structural result.

Proposition 4.7. *For every nonempty subset $A \subseteq \mathbb{F}_p \setminus \{0\}$, there is some $\lambda \in \mathbb{F}_p^\times$ such that $\lambda \cdot A$ can be partitioned as*

$$\lambda \cdot A = P \cup N \cup (\cup_{j=1}^s D_j),$$

where

- (i) the “positive” set P is contained in $(0, \frac{p}{4|P \cup N|})$, the “negative set” N is contained in $(-\frac{p}{4|P \cup N|}, 0)$, and the element $\delta := \sum_{j=1}^s \sum_{d \in D_j} d$ is contained in $(-\frac{p}{4}, \frac{p}{4})$;

and the following also holds if $s > 0$:

- (ii) $P \cup N$ is nonempty, and each D_j is a dissociated set of size $|D_j| \asymp R$, where the implied constant is absolute;

(iii) $\delta \notin \{0\} \cup -P \cup -N$, and moreover $\delta \neq -\sum_{p \in P} p$ if N is nonempty and $\delta \neq -\sum_{n \in N} n$ if P is nonempty;

(iv) $D_1 \cup D_s \cup \{\delta\}$ is a dissociated set;

(v) $|D_1| = |D_s|$.

Before proving this proposition, we make a simple but powerful observation about absorbing elements into dissociated sets.

Lemma 4.8. *Let G be an abelian group, and let $D_1 \cup D_2$ be a partition of a dissociated subset of G . For every element $x \in G \setminus \{0\}$, either $D_1 \cup \{x\}$ or $D_2 \cup \{x\}$ is dissociated.*

Proof. Assume for the sake of contradiction that neither $D_1 \cup \{x\}$ nor $D_2 \cup \{x\}$ is dissociated. Since D_1 is dissociated, the failure of $D_1 \cup \{x\}$ to be dissociated implies that $x \in \text{span}(D_1)$; similarly, $x \in \text{span}(D_2)$. So $\text{span}(D_1) \cap \text{span}(D_2)$ contains a non-zero element, contradicting the assumption that $D_1 \cup D_2$ is dissociated. $\square$

Iterating this observation, we find that if B is a set of size t and $D_1 \cup \dots \cup D_{t+1}$ is a partition of a dissociated set, then it is always possible to add the elements of B to the dissociated sets D_j in such a way that the sets remain dissociated.

We will also make use of the trivial lower bound for the size of a restricted sumset in $\mathbb{Z}$: If $B \subseteq \mathbb{Z}$ is a finite set, then $|B \hat{+} B| \geq 2|B| - 3$. We are now ready to prove Proposition 4.7. The choice of numerical constants appearing in the proof is not important.

Proof of Proposition 4.7. To start, Theorem 4.6 provides some $\lambda \in \mathbb{F}_p^\times$ and a decomposition

$$\lambda \cdot A = D_1 \cup \dots \cup D_s \cup E,$$

where each D_j is a dissociated set of size $\asymp R$ and we have $E \cup \{\delta\} \subseteq (-\frac{p}{90(|E|+1)}, \frac{p}{90(|E|+1)})$, for $\delta := \sum_{j=1}^s \sum_{d \in D_j} d$. Set $P := E \cap (0, p/4|E|)$ and $N := E \cap (-p/4|E|, 0)$. If $s = 0$, then we have already obtained the desired decomposition of $\lambda \cdot A$, so for the remainder of the proof we assume that $s \geq 1$. By replacing λ with $-\lambda$ if necessary, we may assume that $|P| \geq |N|$. In particular, since $|E| \gg R$, this implies that $|P| \gg R$.

We remark that once we have a decomposition satisfying conditions (i)–(iii), we can modify the decomposition to satisfy (iv) and (v) as follows. Split D_1 into 2 parts $D_1^{(1)}, D_1^{(2)}$ each of size $\asymp R$. Lemma 4.8 ensures that either $D_1^{(1)} \cup \{\delta\}$ or $D_1^{(2)} \cup \{\delta\}$ is dissociated; without loss of generality, assume that $D_1^{(1)} \cup \{\delta\}$

is dissociated. Then further split $D_1^{(1)}$ into 2 parts $D_1^{(3)}, D_1^{(4)}$ each of size $\lfloor |D_1^{(1)}|/2 \rfloor \asymp R$, add the leftover element of $D_1^{(1)}$ to $D_1^{(2)}$ if $|D_1^{(1)}|$ was odd, and replace the sequence of sets $D_1, \dots, D_s$ by the sequence $D_1^{(3)}, D_1^{(2)}, D_2, D_3, \dots, D_s, D_1^{(4)}$. This new sequence satisfies (iv) and (v). The remainder of the proof is devoted to finding a decomposition satisfying conditions (i)–(iii).

We will later apply sumset inequalities involving P, N , and we will need P, N to be not-too-small so that we have “room” for sumsets to expand. In anticipation of this, we begin by reducing to the case where N is either empty or of size at least 10. Suppose that $0 < |N| < 10$. Note that $\sum_{n \in N} n \in (-p/90, p/90)$. Split D_1 into $|N| + 1 \leq 10$ sets each of size $\asymp R$; the remark before the proof ensures that we can absorb all of the elements of N into these dissociated sets, and this procedure changes the value of δ by at most $p/90$. Notice that each newly formed D_j still has size $\asymp R$, and that we still have $P \subseteq (0, \frac{p}{40|P \cup N|})$, $N \subseteq (-\frac{p}{40|P \cup N|}, 0)$, and $\delta \in (-p/40, p/40)$.

We now consider two cases depending on the size of N . First, suppose that $N = \emptyset$, and recall that $|P| \gg R$. Since P is rectifiable (i.e., Freiman-isomorphic to a subset of $\mathbb{Z}$), the trivial lower bound for restricted sumsets in integers gives

$$|P \hat{+} P| \geq 2|P| - 3 > |P| + 2,$$

and hence we can find distinct $p_1, p_2 \in P$ such that $p_1 + p_2 + \delta \notin \{0, -\sum_{n \in N} n\} \cup -P$. Splitting D_1 and absorbing p_1, p_2 with the help of Lemma 4.8 as above yields a new decomposition of A where the sum of all of the elements in the dissociated sets is $p_1 + p_2 + \delta$ and hence conditions (ii) and (iii) are satisfied. This procedure also changes the value of δ by at most $p/90$ (say), so we obtain the desired decomposition of $\lambda \cdot A$.

Finally, suppose that $|N| \geq 10$, and recall that we also have $|P| \geq |N| \geq 10$. By absorbing 5 arbitrary elements of N into D_1 , we may assume that $|P| \geq |N| - 5$. Since we still have $|N| > 1$, there is some $n_1 \in N$ such that $\delta + n_1 \neq -\sum_{p \in P} p$; as above, we absorb n_1 into D_1 , so that the final part of condition (iii) is satisfied. Now we have

$$|P \hat{+} P| \geq 2|P| - 3 > |P| + |N| + 2,$$

so there are distinct $p_1, p_2 \in P$ such that

$$\delta + p_1 + p_2 \notin \{0, -\sum_{n \in N} n\} \cup -P \cup -N.$$

Absorbing p_1, p_2 into D_1 gives the desired decomposition of $\lambda \cdot A$. (For the final part of condition (iii), note that this last step preserves the property $\delta + \sum_{p \in P} p \neq 0$.) $\square$

4.4 Ordering P and N

With Proposition 4.7 in hand, we can say a bit more about the remainder of the proof of Theorem 4.2. We will aim to find orderings $\mathbf{p}$ of P , $\mathbf{n}$ of N , and $\mathbf{d}$ of $\cup_j D_j$ such that $\bar{\mathbf{p}}, \mathbf{d}, \mathbf{n}$ is a two-sided valid ordering of A . Of course, we will need each of the three orderings to be two-sided valid on its own, and we will need to avoid creating zero-sum intervals when we concatenate them.

Condition (i) from Proposition 4.7 means that the problem of constructing $\mathbf{p}$ and $\mathbf{n}$ naturally lives in the integers rather than in $\mathbb{F}_p$, as follows. Identify δ and the elements of $P \cup N$ with elements of $(-p/4, p/4) \subseteq \mathbb{Z}$ in the natural way, and note that sums of these elements can be computed equivalently in $\mathbb{F}_p$ and in $(-p/2, p/2) \subseteq \mathbb{Z}$ because the sums in $\mathbb{F}_p$ do not exhibit any wrap-around. Likewise, for any ordering $\mathbf{d}$ of $\cup_j D_j$, we can identify $\text{IS}(\mathbf{d})$ and $\text{IS}(\bar{\mathbf{d}})$ with subsets of $(-p/2, p/2) \subseteq \mathbb{Z}$. Now we observe that the ordering $\bar{\mathbf{p}}, \mathbf{d}, \mathbf{n}$ is two-sided valid if and only if the ordering $\bar{\mathbf{p}}, \delta, \mathbf{n}$ is two-sided valid in the integers, the ordering $\mathbf{d}$ is two-sided valid in $\mathbb{F}_p$, and $\text{IS}(\mathbf{p}) \cap -\text{IS}(\mathbf{d}) = \text{IS}(\mathbf{n}) \cap -\text{IS}(\bar{\mathbf{d}}) = \emptyset$; the key point is that the first condition lives entirely in the integers, the second condition does not concern $\mathbf{p}$ and $\mathbf{n}$, and the third condition lives in the integers for each fixed choice of $\mathbf{d}$.

We will later choose $\mathbf{d}$ randomly, but it turns out that we can model $-\text{IS}(\mathbf{d}), -\text{IS}(\bar{\mathbf{d}})$ by somewhat larger deterministic sets that encode all of the “potentially important” intersections with $\text{IS}(\mathbf{p}), \text{IS}(\mathbf{n})$ (respectively); it will suffice to ensure that $\text{IS}(\mathbf{p}), \text{IS}(\mathbf{n})$ have fairly small intersections with these deterministic sets. With this in mind, the main result of this section is as follows.

Proposition 4.9. *Let $P \subseteq (0, \infty)$ and $N \subseteq (-\infty, 0)$ be finite sets of integers, and let $\delta > 0$ be a positive integer not contained in $-N$; moreover, assume that $\delta \neq -\sum_{n \in N} n$ if $P \neq \emptyset$. Let $Y_1^+, \dots, Y_m^+, Y_1^-, \dots, Y_m^- \subseteq \mathbb{Z}$ be finite sets. Then there are orderings $\mathbf{p}$ of P and $\mathbf{n}$ of N such that $\bar{\mathbf{p}}, \delta, \mathbf{n}$ is two-sided valid and we have*

$$|\text{IS}(\mathbf{p}) \cap Y_j^+| \leq \inf_{L \in \mathbb{N}} \left(\frac{|Y_j^+|}{L} + L + 4 + 4 \sum_{i=1}^{j-1} |Y_i^+| \right) \quad (4.3)$$

and

$$|\text{IS}(\mathbf{n}) \cap Y_j^-| \leq \inf_{L \in \mathbb{N}} \left(\frac{|Y_j^-|}{L} + L + 4 + 4 \sum_{i=1}^{j-1} |Y_i^-| \right) \quad (4.4)$$

for all $1 \leq j \leq m$.

In [66], the author presented a simple algorithm for inductively constructing a two-sided valid ordering of any finite subset of $\mathbb{Z} \setminus \{0\}$. Let us quickly review this algorithm since it forms the basis for the proof of Proposition 4.9. Suppose $P \subseteq (0, \infty)$ and $N \subseteq (-\infty, 0)$ are finite sets of integers; we want to produce

orderings $\mathbf{p}$ of P and $\mathbf{n}$ of N such that $\overline{\mathbf{p}}, \mathbf{n}$ is two-sided valid. We will construct the sequences $\mathbf{p} = p_1, \dots, p_{|P|}$ and $\mathbf{n} = n_1, \dots, n_{|N|}$ from the larger indices to the smaller indices. For the first step, consider the sign of $\sum_{n \in N} n + \sum_{p \in P} p$. Suppose that this sum is non-negative; we will choose the value of $p_{|P|}$ as follows. There is some $p^* \in P$ such that $\sum_{n \in N} n + \sum_{p \in P \setminus \{p^*\}} p \neq 0$, and we choose this p^* to be our $p_{|P|}$. This choice ensures that if $\overline{\mathbf{p}}, \mathbf{n}$ is a two-sided valid ordering of the remaining elements $P \setminus \{p_{|P|}\}, N$, then $p_{|P|}, \overline{\mathbf{p}}, \mathbf{n}$ is the desired two-sided valid ordering of P, N : Any interval containing $p_{|P|}$ and not containing any of $\mathbf{n}$ clearly has strictly positive sum; any proper interval containing both $p_{|P|}$ and some elements of $\mathbf{n}$ must contain all of P and hence has strictly positive sum by our assumption that $\sum_{n \in N} n + \sum_{p \in P} p \geq 0$; the intervals strictly contained in $\overline{\mathbf{p}}, \mathbf{n}$ are all non-zero-sum by assumption; and the interval consisting of all of $\overline{\mathbf{p}}, \mathbf{n}$ has non-zero sum by our choice of p^* . If instead $\sum_{n \in N} n + \sum_{p \in P} p < 0$, then we choose $n_{|N|}$ analogously. With this first step complete, we throw out the already-chosen element $p_{|P|}$ or $n_{|N|}$ and repeat this process with the remaining elements. This procedure produces the desired orderings $\mathbf{p}, \mathbf{n}$.

The above algorithm has a lot of slack, in the sense that at each step there are many possible choices. To exploit this slack and prove Proposition 4.9, we will employ a modified algorithm that greedily avoids partial sums of $\mathbf{p}$ lying in the Y_j^+ 's and partial sums of $\mathbf{n}$ lying in the Y_j^- 's.

Proof of Proposition 4.9. We will construct the sequences $\mathbf{p} = p_1, \dots, p_{|P|}$ and $\mathbf{n} = n_1, \dots, n_{|N|}$ from the larger indices to the smaller indices. Suppose that we have already chosen the values of $p_{|P|}, p_{|P|-1}, \dots, p_{k+1}$ and $n_{|N|}, n_{|N|-1}, \dots, n_{\ell+1}$. At the next step, we will choose the value of either p_k or n_ℓ depending on the sign of the sum of all of the remaining elements. Let

$$P_k := P \setminus \{p_{|P|}, \dots, p_{k+1}\} \quad \text{and} \quad N_\ell := N \setminus \{n_{|N|}, \dots, n_{\ell+1}\}$$

be the sets of remaining elements of P and N , and define the quantities

$$\pi_k := \sum_{p \in P_k} p \quad \text{and} \quad \nu_\ell := \sum_{n \in N_\ell} n.$$

As in the algorithm from [66], we will succeed in constructing a two-sided valid ordering as long as p_k, n_ℓ avoid a few particular potential values (for more details, see Claim 4.11 below). If we are choosing p_k , then we want p_k not to be equal to $\pi_k + \delta + \nu_\ell$, since this choice of p_k would lead to a zero-sum interval $p_{k-1} + \dots + p_1 + \delta + n_1 + \dots + n_\ell = 0$. Likewise, if we are choosing n_ℓ , then we want n_ℓ not to be equal to either $\delta + \nu_\ell$ or $\pi_k + \delta + \nu_\ell$, since these choices of n_ℓ would lead to zero-sum intervals $\delta + n_1 + \dots + n_{\ell-1} = 0$

and $p_k + \dots + p_1 + \delta + n_1 + \dots + n_{\ell-1} = 0$. With this in mind, we define the sets

$$P'_k := P_k \setminus \{\pi_k + \delta + \nu_\ell\} \quad \text{and} \quad N'_\ell := N_\ell \setminus \{\delta + \nu_\ell, \pi_k + \delta + \nu_\ell\}$$

of “allowable” choices for p_k and n_ℓ . Due to the assumptions in Proposition 4.9, it will always transpire that the set P'_k or N'_ℓ under consideration is nonempty.

Again as in the algorithm from [66], the sign of the quantity $\pi_k + \delta + \nu_\ell$ will determine whether we choose the value of p_k or the value of n_ℓ next:¹

1. Suppose that $\pi_k + \delta + \nu_\ell \geq 0$ and $k > 0$. Then we will choose $p_k \in P'_k$ as follows. If there is some $p^* \in P'_k$ such that $\pi_k - p^* \notin \cup_j Y_j^+$, then choose p_k to be this p^* and say that the current step is a *skip-step* for P .

Now, consider the case where $\pi_k - P'_k \subseteq \cup_j Y_j^+$. Let i be minimal such that $\pi_k - P'_k$ intersects Y_i^+ , and say that the current step is an *i-step* for P . If $\pi_k - P'_k \subseteq Y_i^+$, then let p_k be the largest element of P'_k . If $\pi_k - P'_k \not\subseteq Y_i^+$, then let p_k be the largest $p^* \in P'_k$ such that $\pi_k - p^* \notin Y_i^+$.

2. Suppose that $\pi_k + \delta + \nu_\ell \geq 0$, $k = 0$ (i.e., we have already chosen all of $\mathbf{p}$), and $\ell > 0$, or that $\pi_k + \delta + \nu_\ell < 0$ and $\ell > 0$. Then we will choose $n_\ell \in N'_\ell$ as follows. If there is some $n^* \in N'_\ell$ such that $\nu_\ell - n^* \notin \cup_j Y_j^-$, then choose n_ℓ to be this n^* and say that the current step is a *skip-step* for N .

Now, consider the case where $\nu_\ell - N'_\ell \subseteq \cup_j Y_j^-$. Let i be minimal such that $\nu_\ell - N'_\ell$ intersects Y_i^- , and say that the current step is an *i-step* for N . If $\nu_\ell - N'_\ell \subseteq Y_i^-$, then let n_ℓ be the smallest (i.e., most negative) element of N'_ℓ . If $\nu_\ell - N'_\ell \not\subseteq Y_i^-$, then let n_ℓ be the smallest (i.e., most negative) $n^* \in N'_\ell$ such that $\nu_\ell - n^* \notin Y_i^-$.

We begin with $(k, \ell) = (|P|, |N|)$ and run the above procedure until we reach $(k, \ell) = (0, 0)$. To establish Proposition 4.9, we must show three things: that the algorithm actually runs and produces orderings $\mathbf{p}$ of P and $\mathbf{n}$ of N ; that the resulting ordering $\mathbf{p}, \delta, \bar{\mathbf{n}}$ is two-sided valid; and that $\text{IS}(\mathbf{p})$ and $\text{IS}(\mathbf{n})$ have small intersections with the Y_j^+ 's and Y_j^- 's (respectively).

Claim 4.10. *The above algorithm runs all the way to $(k, \ell) = (0, 0)$ and produces orderings $\mathbf{p}$ of P and $\mathbf{n}$ of N .*

Proof. Note that as long as $(k, \ell) \neq (0, 0)$, we fall into one of the two cases. Indeed, when $\pi_k + \delta + \nu_\ell \geq 0$, we fall into case (1) or case (2) according to whether $k > 1$ or $k = 0$. When $\pi_k + \delta + \nu_\ell < 0$, we must have $\ell > 0$ since $\delta > 0$ and $\pi_k \geq 0$, so we fall into case (2).

¹The apparent asymmetry in the cases arises from the assumption in Proposition 4.9 that $\delta > 0$.

It remains to show that the sets P'_k, N'_ℓ are always nonempty when needed. First, consider case (1). It is clear that $P'_k \neq \emptyset$ as long as $k > 1$. When $k = 1$, the set P_1 consists of a single element p_1 , and we have $\pi_1 = p_1$. We must show that $\pi_1 + \delta + \nu_\ell \neq p_1$, i.e., that $\delta \neq -\nu_\ell$. When $\ell = |N|$, this is precisely the assumption in Proposition 4.9 that $\delta \neq -\sum_{n \in N} n$. For $\ell < |N|$, recall that $n_{\ell+1}$ was chosen to be an element of $N'_{\ell+1}$, which by construction does not contain $\delta + \nu_{\ell+1}$. It follows that $\nu_\ell = \nu_{\ell+1} - n_{\ell+1} \neq \nu_{\ell+1} - (\delta + \nu_{\ell+1}) = -\delta$, as desired.

Now, consider case (2). We begin with the subcase where $\pi_k + \delta + \nu_\ell \geq 0$ and $k = 0$. Note that $\pi_0 = 0$ and hence $\pi_0 + \delta + \nu_\ell = \delta + \nu_\ell$. It is clear that $N'_\ell \neq \emptyset$ as long as $\ell > 1$. When $\ell = 1$, the set N_1 consists of a single element n_1 , and we have $\nu_1 = n_1$. The assumption $\delta > 0$ ensures that $\delta + \nu_1 = \delta + n_1 \neq n_1$, so $N'_1 \neq \emptyset$.

Finally, we treat the subcase where $\pi_k + \delta + \nu_\ell < 0$. It is clear that $N'_\ell \neq \emptyset$ as long as $\ell > 2$. When $\ell = 2$, the set N_2 consists of two elements n_1, n_2 , and we have $\nu_2 = n_1 + n_2$. Then $\delta + \nu_2 = \delta + n_1 + n_2 \notin N_2$ by the assumption in Proposition 4.9 that $\delta \notin -N$ so neither of n_1, n_2 is equal to $-\delta$, and it follows that $N'_2 \neq \emptyset$. When $\ell = 1$, the set N_1 consists of a single element n_1 , and we have $\nu_1 = n_1$. Then $\delta + \nu_1 = \delta + n_1 \neq n_1$ since $\delta > 0$, and $\pi_k + \delta + \nu_1 = \pi_k + \delta + n_1 \neq n_1$ since $\pi_k + \delta > 0$. Thus $N'_1 \neq \emptyset$, and this concludes the proof. $\square$

Claim 4.11. *The ordering $\bar{\mathbf{p}}, \delta, \mathbf{n}$ is two-sided valid.*

Proof. Since any zero-sum interval must contain both positive and negative numbers, we can restrict our attention to intervals of the form $p_k + \cdots + p_1 + \delta + n_1 + \cdots + n_\ell$ (with sum $\pi_k + \delta + \nu_\ell$) and $\delta + n_1 + \cdots + n_\ell$ (with sum $\delta + \nu_\ell$).

Let us first consider the sums $\pi_k + \delta + \nu_\ell$. Note that we do not need to worry about $(k, \ell) = (|P|, |N|)$, since the corresponding interval is the entire sequence $\bar{\mathbf{p}}, \delta, \mathbf{n}$, so we may assume that either $k < |P|$ or $\ell < |N|$. Let (k^*, ℓ^*) be the earliest step in the algorithm where $k^* \leq k$ and $\ell^* \leq \ell$. Then the previous step in the algorithm was either $(k^* + 1, \ell^*)$ or $(k^*, \ell^* + 1)$; without loss of generality assume that it was the former, since the argument for the latter is identical. Then $k^* = k$ and $\ell^* \leq \ell$. If $\ell^* = \ell$, then

$$\pi_k + \delta + \nu_\ell = (\pi_{k+1} - p_{k+1}) + \delta + \nu_\ell$$

is nonzero because we chose $p_{k+1} \in P'_{k+1}$ and the set P'_{k+1} does not contain $\pi_{k+1} + \delta + \nu_\ell$. If instead $\ell^* < \ell$, then there is some $k' > k$ such that n_ℓ was chosen at step (k', ℓ) . It follows that

$$\pi_k + \delta + \nu_\ell < \pi_{k'} + \delta + \nu_\ell < 0,$$

so $\pi_k + \delta + \nu_\ell$ is nonzero, as desired.

Let us now consider the sums $\delta + \nu_\ell$. We can again quickly dispose of the case $\ell = |N|$. Indeed, if $P = \emptyset$, then the corresponding interval is the entire sequence $\bar{\mathbf{p}}, \delta, \mathbf{n}$, and if $P \neq \emptyset$, then $\delta \neq -\sum_{n \in N} n = -\nu_{|N|}$ by assumption. So we may assume that $\ell < |N|$, and we conclude by noting that $\delta + \nu_\ell = \delta + (\nu_{\ell+1} - n_{\ell+1}) \neq 0$ since $N'_{\ell+1}$ does not contain $\delta + \nu_{\ell+1}$. $\square$

Claim 4.12. *The orderings $\mathbf{p}$ and $\mathbf{n}$ satisfy*

$$|\text{IS}(\mathbf{p}) \cap Y_j^+| \leq \inf_{L \in \mathbb{N}} \left(\frac{|Y_j^+|}{L} + L + 2 + 4 \sum_{i=1}^{j-1} |Y_i^+| \right)$$

and

$$|\text{IS}(\mathbf{n}) \cap Y_j^-| \leq \inf_{L \in \mathbb{N}} \left(\frac{|Y_j^-|}{L} + L + 2 + 4 \sum_{i=1}^{j-1} |Y_i^-| \right)$$

for all $1 \leq j \leq m$.

Proof. We will prove the statement only for $|\text{IS}(\mathbf{n}) \cap Y_j^-|$ since the argument for $|\text{IS}(\mathbf{p}) \cap Y_j^+|$ is essentially identical. Recall that $\text{IS}(\mathbf{n}) = \{\nu_\ell = \sum_{i=1}^\ell n_i : 0 \leq \ell \leq |N|\}$. For $0 \leq \ell < |N|$, write $\nu_\ell = \nu_{\ell+1} - n_{\ell+1}$. This quantity can lie in Y_j^- only when the choice of $n_{\ell+1}$ is a j -step or an i -step for some $i < j$. We will bound these two contributions separately. Note that skip-steps and i -steps for $i > j$ never contribute.

We first consider the contribution of j -steps. Notice that the partial sums ν_ℓ are strictly increasing (becoming less negative) as ℓ decreases. Suppose that the choice of $n_{\ell+1}$ is a j -step and $\nu_\ell = \nu_{\ell+1} - n_{\ell+1} \in Y_j^-$. Then we must have $\nu_{\ell+1} - N'_{\ell+1} \subseteq Y_j^-$. Since $n_{\ell+1}$ is the smallest (most negative) element of $N'_{\ell+1}$, the other $|N'_{\ell+1} \setminus \{n_{\ell+1}\}| \geq \ell + 1 - 3 = \ell - 2$ elements of $\nu_{\ell+1} - N'_{\ell+1} \subseteq Y_j^-$ lie in the interval $(\nu_{\ell+1}, \nu_\ell)$; it follows that these elements are “skipped” and can never appear in $\text{IS}(\mathbf{n})$. In particular, from such j -steps with $\ell \geq L + 1$ we obtain at most $|Y_j^-|/L$ elements of $\text{IS}(\mathbf{n}) \cap Y_j^-$. From j -steps with $\ell \leq L$ we trivially obtain at most $L + 1$ elements of $\text{IS}(\mathbf{n}) \cap Y_j^-$.

We now consider the contribution of i -steps with $i < j$. We will trivially bound this contribution by the total number of i -steps with $i < j$. We claim that the number of i -steps is at most $4|Y_i^-|$ for each i . For each i -step ℓ , let $y(\ell)$ denote the largest (least negative) element of $(\nu_\ell - N'_\ell) \cap Y_i^-$. It suffices to show that each $y \in Y_i^-$ appears as $y(\ell)$ for at most 4 different i -steps ℓ . If $y(\ell)$ is not the largest element of $\nu_\ell - N'_\ell$, then it is distinct from $y(\ell')$ for all $\ell' < \ell$ since

$$\nu_{\ell'} \geq \nu_{\ell-1} = \nu_\ell - n_\ell > y(\ell)$$

by the definition of an i -step. If $y(\ell)$ is the largest element of $\nu_\ell - N'_\ell$, then it is one of the three largest

elements of $\nu_\ell - N_\ell$. Notice that the largest element of $\nu_\ell - N_\ell$ is strictly increasing as ℓ decreases, the second-largest element of $\nu_\ell - N_\ell$ is strictly increasing as ℓ decreases, and the third-largest element of $\nu_\ell - N_\ell$ is strictly increasing as ℓ decreases; it follows that each y can appear at most three times as one of the three largest elements of $\nu_\ell - N_\ell$. Thus we have shown that each i -step ℓ in the algorithm is associated with some number $y(\ell) \in Y_i^-$ and moreover that any given $y \in Y_i^-$ appears as $y(\ell)$ for at most 4 different i -steps ℓ , so we conclude that the total number of i -steps is at most $4|Y_i^-|$. This establishes the claim.

Combining these contributions (and adding 1 for $\nu_{|N|}$) gives the desired upper bound.² $\square$

These three claims together imply Proposition 4.9. $\square$

4.5 Splitting the dissociated sets

In this section we manipulate the dissociated sets D_j in order to make their sums suitably generic; this will avoid “bad” scenarios in the random orderings of the D_j ’s that we will consider in the next section. Recall that if D is a dissociated set, then all of the subset sums of D are distinct. In particular, if we choose a uniformly random partition of D into parts $D^{(1)}, D^{(2)}, D^{(3)}, D^{(4)}$ of equal size (up to rounding), then (omitting floor functions) for each $1 \leq i \leq 4$ the $\binom{|D|}{|D|/4}$ possible values of $\sum_{d \in D^{(i)}} d$ are all achieved with equal probability; likewise, each of the quantities $\sum_{d \in D^{(1)} \cup D^{(2)}} d, \sum_{d \in D^{(2)} \cup D^{(3)}} d, \sum_{d \in D^{(3)} \cup D^{(4)}} d$ is uniformly distributed on $\binom{|D|}{|D|/2}$ possible values, and each of the quantities $\sum_{d \in D^{(1)} \cup D^{(2)} \cup D^{(3)}} d, \sum_{d \in D^{(2)} \cup D^{(3)} \cup D^{(4)}} d$ is uniformly distributed on $\binom{|D|}{3|D|/4}$ possible values. Since $\binom{|D|}{|D|/4}, \binom{|D|}{|D|/2}, \binom{|D|}{3|D|/4}$ are all $e^{\Omega(|D|)}$, we obtain very strong anti-concentration for the sums under consideration. We record this simple but important fact in the following lemma.

Lemma 4.13. *Let $D \subset G$ be a dissociated set, and let $D = D^{(1)} \cup D^{(2)} \cup D^{(3)} \cup D^{(4)}$ be a uniformly random partition of D into four sets of equal size (up to rounding). Then for every nonempty proper interval $I \subseteq [4]$ and every $x \in G$, we have*

$$\mathbb{P} \left(\sum_{i \in I} \sum_{d \in D^{(i)}} d = x \right) \leq e^{-\Omega(|D|)}.$$

Let $D_1, \dots, D_s$ be the dissociated sets appearing in the structural decomposition of A from Proposition 4.7. We will split and reorder these dissociated sets as follows. For each $j \in [1, s]$, we partition $D_j = \cup_{i=1}^4 D_j^{(i)}$ into four sets of equal size uniformly at random as in Lemma 4.13, and we require that $|D_1^{(1)}| = |D_s^{(4)}|$. We

²To bound the intersection between $\text{IS}(\mathbf{p})$ and Y_j^+ , one simply interchanges “smaller” and “larger” throughout the proof. Since we have $|P_k \setminus P'_k| \leq 1$ instead of $|N_\ell \setminus N'_\ell| \leq 2$, we could replace $\ell - 2$ with $k - 1$ in the second paragraph and replace $4|Y_j^-|$ with $3|Y_j^+|$ in the third paragraph to obtain even a slightly tighter bound.

do all of these splittings independently. Next, we place these newly formed dissociated sets in the order

$$D_1^{(1)}, D_1^{(2)}, D_2^{(1)}, D_2^{(2)}, \dots, D_s^{(1)}, D_s^{(2)}, D_1^{(3)}, D_1^{(4)}, D_2^{(3)}, D_2^{(4)}, \dots, D_s^{(3)}, D_s^{(4)} \quad (4.5)$$

and note that of course the decomposition

$$\lambda \cdot A = P \cup N \cup (\cup_{i=1}^4 \cup_{j=1}^s D_j^{(i)})$$

still holds (with the same value of δ). For notational convenience, write $T_1, T_2, \dots, T_u$ (with $u = 4s$) for the new sequence of dissociated sets in (4.5), and let $\tau_j := \sum_{t \in T_j} t$.

Let us pause at this point and describe the remainder of the strategy for proving Theorem 4.2. We will eventually construct a two-sided valid ordering of A of the form

$$\bar{\mathbf{p}}, \mathbf{t}_1, \dots, \mathbf{t}_u, \mathbf{n},$$

where each $\mathbf{t}_i$ is an ordering of T_i chosen randomly according to a certain distribution. Our task will be to show that such an ordering $a_1, \dots, a_{|A|}$ is likely to avoid zero-sum subintervals, namely, proper nonempty intervals $I \subset [|A|]$ with $\sum_{i \in I} a_i = 0$. For the remainder of the chapter, we will refer to proper nonempty intervals as simply “intervals”. We divide such intervals I into two “types”, which we will treat using different arguments. Recall that $K = c_2 R^{1/3}$.

Definition 4.14. Let $I \subset [|A|]$ be a proper nonempty interval. We say that I is *Type II* if it contains between K and $|T_j| - K$ elements of some T_j , and otherwise we say that it is *Type I*.

We will refer to the first K elements in an ordering $\mathbf{t} = t_1, \dots, t_m$ as its left border and to the final K elements as its right border. The remaining elements $t_{K+1}, \dots, t_{m-K}$ make up the interior region of $\mathbf{t}$. In this language (and ignoring intervals contained in a single T_j , which can never be zero-sum), a Type II interval is an interval with at least one endpoint in the interior region of one of the orderings $\mathbf{t}_j$, and a Type I interval is an interval with each endpoint in $\bar{\mathbf{p}}, \mathbf{n}$, or a border region of some $\mathbf{t}_j$. One should think of Type II intervals as generic and of Type I intervals as exceptional. (Obviously the identification of intervals $I \subset [|A|]$ as Type I and Type II does not depend on the random choices of the $\mathbf{t}_j$'s).

The main benefit of the above splitting-and-rearranging procedure is that it lets us dispose of nearly all Type I intervals even before we choose the random orderings $\mathbf{t}_j$. The following lemma makes this precise. We say that an event holds *with high probability* if it holds with probability tending to 1 as p tends to infinity.

Lemma 4.15. *Let $c > 0$ be any constant. Let $1 \leq s \leq e^{c(\log p)^{1/4}}$, and let $D_1, \dots, D_s \subseteq \mathbb{F}_p$ be dissociated*

sets each of size $\asymp R$, with the property that $D_1 \cup D_s \cup \{\delta\}$ is dissociated. Let $\mathbf{p}$ and $\mathbf{n}$ be sequences over $\mathbb{F}_p$ each of length at most $e^{c(\log p)^{1/4}}$, and assume that $\bar{\mathbf{p}}, \delta, \mathbf{n}$ is a two-sided valid ordering. If the sequence $T_1, \dots, T_u$ of dissociated sets is chosen randomly as described above, then each $|T_j| \asymp R$ and each $T_{2j-1} \cup T_{2j}$ is dissociated, and the following holds with high probability:

(i) for each proper nonempty interval $I = [i, j] \subseteq [u]$, we have that

$$0 \notin \left(\sum_{\leq K} (T_{i-1}) \cup - \sum_{\leq K} (T_i) \right) + \tau_i + \dots + \tau_j + \left(- \sum_{\leq K} (T_j) \cup \sum_{\leq K} (T_{j+1}) \right)$$

(with the convention that $T_0 = T_{u+1} = \emptyset$);

(ii) for each $1 \leq j \leq u-1$, we have that

$$0 \notin \text{IS}(\mathbf{p}) + \tau_1 + \dots + \tau_j + \left(- \sum_{\leq K} (T_j) \cup \sum_{\leq K} (T_{j+1}) \right);$$

and for each $2 \leq j \leq u$, we have that

$$0 \notin \text{IS}(\mathbf{n}) + \tau_u + \dots + \tau_j + \left(- \sum_{\leq K} (T_j) \cup \sum_{\leq K} (T_{j-1}) \right);$$

(iii) the ordering $\bar{\mathbf{p}}, \tau_1, \dots, \tau_u, \mathbf{n}$ is two-sided valid.

Three remarks are in order before we proceed to the proof.

1. To see how this lemma pertains to Type I intervals containing nearly all (i.e., at least $|T_j| - K$ elements) of some T_j , simply note the identity $\sum_{\geq |T_j| - K} (T_j) = \tau_j - \sum_{\leq K} (T_j)$.

2. Items (i)–(iii) handle all Type I intervals except for the following:

- intervals fully contained in a single T_j ;
- intervals starting in the left border of $\mathbf{t}_1$ and ending in the right border of $\mathbf{t}_u$;
- intervals beginning in the right border of $\mathbf{t}_j$ and ending in the left border of $\mathbf{t}_{j+1}$ for some j ;
- intervals with one endpoint in $\bar{\mathbf{p}}$ or $\mathbf{n}$ and the other endpoint in the left border of $\mathbf{t}_1$ or the right border of $\mathbf{t}_u$.

Moreover, the first case cannot lead to zero-sum intervals because each T_j is dissociated; likewise, there cannot be zero-sum intervals in the second case because of the assumption that $D_1 \cup D_s \cup \{\delta\}$ (and a

fortiori $T_1 \cup T_u \cup \{\delta\}$ is dissociated. In the third case, we never have to worry about zero-sum intervals with j odd since each $T_{2k-1} \cup T_{2k}$ is dissociated.

3. The lemma would continue to hold with K as large as a small constant times R , but we will not have occasion to make use of this fact.

Proof of Lemma 4.15. We begin with the crucial observation that if $I \subset [u]$ is any proper nonempty subinterval and $x \in \mathbb{F}_p$ is any element, then we have the anti-concentration inequality

$$\mathbb{P}\left(\sum_{i \in I} \tau_i = x\right) = e^{-\Omega(R)}.$$

Indeed, there is some $j \in [s]$ such that $\{T_i : i \in I\}$ contains at least one but not all of $D_j^{(1)}, \dots, D_j^{(4)}$. Suppose that it contains $D_j^{(1)}$ but none of $D_j^{(2)}, D_j^{(3)}, D_j^{(4)}$ (the remaining cases are analogous). Since the splitting of D_j is independent of the splittings of the other D_k 's, Lemma 4.13 gives

$$\begin{aligned} \mathbb{P}\left(\sum_{i \in I} \tau_i = x\right) &= \sum_{z \in \mathbb{F}_p} \mathbb{P}\left(\sum_{i \in I \setminus \{2j-1\}} \tau_i = z \quad \text{and} \quad \tau_{2j-1} = x - z\right) \\ &= \sum_{z \in \mathbb{F}_p} \mathbb{P}\left(\sum_{i \in I \setminus \{2j-1\}} \tau_i = z\right) \mathbb{P}\left(\sum_{d \in D_j^{(1)}} d = x - z\right) \\ &\leq \sum_{z \in \mathbb{F}_p} \mathbb{P}\left(\sum_{i \in I \setminus \{2j-1\}} \tau_i = z\right) e^{-\Omega(|D_j|)} = e^{-\Omega(R)}. \end{aligned}$$

With this observation in hand, we proceed to the main body of the proof. Note that (iii) holds whenever (i) and (ii) hold since $0 \in \sum_{\leq K}(T_j)$ and as we assumed that $\bar{\mathbf{p}}, \delta, \mathbf{n}$ is two-sided valid. So, by the union bound, it suffices to show that each of (i) and (ii) holds with high probability.

We begin with (i). Fix some proper nonempty interval $I = [i, j] \subseteq [u]$. The assertion of (i) for this I is that

$$\tau_i + \dots + \tau_j \notin \left(-\sum_{\leq K}(T_{i-1}) \cup \sum_{\leq K}(T_i)\right) + \left(\sum_{\leq K}(T_j) \cup -\sum_{\leq K}(T_{j+1})\right).$$

The set on the right-hand side has size at most

$$\left(\left|\sum_{\leq K}(T_{i-1})\right| + \left|\sum_{\leq K}(T_i)\right|\right) \left(\left|\sum_{\leq K}(T_j)\right| + \left|\sum_{\leq K}(T_{j+1})\right|\right) \leq e^{O(R \cdot H(O(K/R)))},$$

where $H(x) := -x \log_2(x) - (1-x) \log_2(1-x)$ is the binary entropy function. The definition of K ensures that $H(O(K/R)) = o(1)$ (with a lot of room to spare), and then the observation from the beginning of the

proof tells us that (i) fails for I with probability at most $e^{o(R)-\Omega(R)} = e^{-\Omega(R)}$. A union bound over the (at most u^2) choices of I shows that (i) fails with probability at most

$$u^2 e^{-\Omega(R)} \leq e^{2c(\log p)^{1/4} - \Omega(c_1(\log p)^{3/4})} = o(1),$$

again with plenty of room to spare.

The proof of (ii) is nearly identical and we omit it; we remark that the bounds $|\text{IS}(\mathbf{p})|, |\text{IS}(\mathbf{n})| \leq |A| \leq e^{c(\log p)^{1/4}}$ hold because we fixed $\mathbf{p}$ and $\mathbf{n}$ in advance.

□

As noted in remark (2) following Lemma 4.15, there remain two sorts of Type I intervals to address. The first is Type I intervals contained in $T_{2k} \cup T_{2k+1}$ for some k . We can avoid zero-sums here by picking the orderings $\mathbf{t}_{2k}, \mathbf{t}_{2k+1}$ according to a suitable joint distribution which we will describe in section 4.6. The second is Type I intervals with one endpoint in $\bar{\mathbf{p}}$ or $\mathbf{n}$ and the other endpoint in the left border of $\mathbf{t}_1$ or the right border of $\mathbf{t}_u$. The crucial ingredient for dealing with these will turn out to be the last part of Proposition 4.9. Since Proposition 4.9 must be applied prior to the random splitting procedure described in this section, it is a bit of a nuisance that the input sets Y_j^+, Y_j^- must be described in terms of the sets D_j rather than the sets T_j . The following lemma will let us remedy this issue.

Lemma 4.16. *Let $T_1 = D_1^{(1)}$ and $T_u = D_s^{(4)}$ be the random sets from (4.5). Then with probability at least $1/2$, we have for all $1 \leq j \leq K$ that*

$$\begin{aligned} \left| \sum_{=j} (T_1) \cap (-\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))) \right| &\leq 4K \frac{\binom{|T_1|}{j}}{\binom{|D_1|}{j}} \left| \sum_{=j} (D_1) \cap (-\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))) \right|, \\ \left| \sum_{=j} (T_u) \cap (-\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))) \right| &\leq 4K \frac{\binom{|T_u|}{j}}{\binom{|D_s|}{j}} \left| \sum_{=j} (D_s) \cap (-\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))) \right|. \end{aligned}$$

Proof. Since D_1 is dissociated, the quantity $\left| \sum_{=j} (D_1) \cap (-\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))) \right|$ simply counts the subsets $S \subseteq D_1$ of size $|S| = j$ with $\sum_{d \in S} d \in -\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))$, and likewise for T_1 . As $T_1 = D_1^{(1)}$ is chosen uniformly from all subsets of D_1 of size $|D_1|/4$, we have

$$\mathbb{E} \left(\left| \sum_{=j} (T_1) \cap (-\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))) \right| \right) = \frac{\binom{|T_1|}{j}}{\binom{|D_1|}{j}} \left| \sum_{=j} (D_1) \cap (-\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))) \right|,$$

and Markov's Inequality implies that the first bound in the conclusion of the lemma fails for each j with probability at most $1/4K$. The same argument applies with D_s, T_u in place of D_1, T_1 , and the conclusion of

the lemma follows from a union bound over $1 \leq j \leq K$. $\square$

4.6 Randomizing the dissociated sets

We are finally ready to describe how we will construct a two-sided valid ordering of A . Suppose that $A \subseteq \mathbb{F}_p \setminus \{0\}$ has size $|A| \leq e^{c(\log p)^{1/4}}$. After we replace A by a suitable dilate (which is harmless with regard to finding two-sided valid orderings), Proposition 4.7 provides a decomposition

$$A = P \cup N \cup (\cup_{j=1}^s D_j)$$

satisfying conditions (i)-(v) of that proposition and $\delta = \sum_j \sum_{d \in D_j} d > 0$. Now, with $K = c_2 R^{1/3}$ for a suitably small constant $c_2 > 0$, set

$$Y_j^+ := - \sum_{=j} (D_1) \cup \left(-\delta + \sum_{=j} (D_s) \right) \quad \text{and} \quad Y_j^- := - \sum_{=j} (D_s) \cup \left(-\delta + \sum_{=j} (D_1) \right) \quad (4.6)$$

for each $1 \leq j \leq K$, and apply Proposition 4.9. This provides orderings $\mathbf{p}$ of P and $\mathbf{n}$ of N such that the sequence $\bar{\mathbf{p}}, \delta, \mathbf{n}$ is two-sided valid and such that (4.3), (4.4) hold. Finally, we can use Lemmas 4.15 and 4.16 to obtain dissociated sets $T_1, \dots, T_u$ from $D_1, \dots, D_s$ such that $A = P \cup N \cup (\cup_i T_i)$ and the conclusions of these two lemmas are simultaneously satisfied; fix such a choice of $T_1, \dots, T_u$. Recall that $\tau_i = \sum_{t \in T_i} t$, and write $m_i := |T_i|$. The two-sided valid ordering of A that we will construct will be of the form

$$\bar{\mathbf{p}}, \mathbf{t}_1, \dots, \mathbf{t}_u, \mathbf{n},$$

where the $\mathbf{t}_i$'s are orderings of the T_i 's chosen randomly according to certain distributions, whose description and analysis occupies the remainder of this section.

Recall that if T is a dissociated set, then all of the subset sums of T are distinct. In particular, in a uniformly random ordering of the elements of T , the sum of the first k elements is uniformly distributed on $\binom{|T|}{k}$ values. As long as k is not too close to 0 or $|T|$, this sum is very anti-concentrated, and so with very high probability it will avoid any particular small set of values. It follows that uniformly random orderings $\mathbf{t}_i$ would with high probability avoid zero-sum Type II intervals. We can ignore most Type I intervals due to Lemma 4.15, but the remaining Type I intervals, as described in remark (2) following that lemma, still cause issues. We will show that each of these potential zero-sum Type I intervals can be avoided “locally” by introducing some non-uniformity into the distributions determining the orderings $\mathbf{t}_i$.

We begin with the orderings $\mathbf{t}_1, \mathbf{t}_u$. Say that an ordering $t_1, \dots, t_{m_1}$ of T_1 is *acceptable* if

$$t_1 + \dots + t_k \notin -\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n})) \quad \text{for all } 1 \leq k \leq K,$$

and say that an ordering $t_1, \dots, t_{m_u}$ of T_u is *acceptable* if

$$t_1 + \dots + t_k \notin -\text{IS}(\mathbf{n}) \cup (\delta + \text{IS}(\mathbf{p})) \quad \text{for all } 1 \leq k \leq K.$$

Using Proposition 4.9 and the fact that T_1, T_u satisfy the conclusion of Lemma 4.16, we can show that uniformly random orderings of T_1, T_u are acceptable with large probability.

Lemma 4.17. *A uniformly random ordering of T_1 is acceptable with probability at least 0.98, and a uniformly random ordering of T_u is acceptable with probability at least 0.98.*

Proof. We prove only the statement for T_1 since the argument for T_u is identical. Let $t_1, \dots, t_{|T_1|}$ be our uniformly random ordering of T_1 . By the union bound, it suffices to show that $\mathbb{P}(t_1 \in -\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))) \leq 0.01$ and that

$$\mathbb{P}(t_1 + \dots + t_k \in -\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))) \leq 0.01K^{-1}$$

for each $2 \leq k \leq K$. Fix some $1 \leq k \leq K$. Then the quantity $t_1 + \dots + t_k$ is uniformly distributed on the set $\sum_{=k}(T_1)$, which has size $\binom{|T_1|}{k}$. Recall that we applied Proposition 4.9 with the sets Y_j^+, Y_j^- as in (4.6). Since $|D_1| = |D_s|$ and $D_1 \cup D_s \cup \{\delta\}$ is dissociated by Proposition 4.7, we have $|Y_j^+| = |Y_j^-| = 2\binom{|D_1|}{j}$ for all j . Then the conclusion of Proposition 4.9, with $L := \lfloor |Y_k^+|^{1/2} \rfloor$, gives

$$\left| \sum_{=k}(D_1) \cap (-\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))) \right| \ll |Y_k^+|^{1/2} + 1 + \sum_{j < k} |Y_j^+|.$$

For $k = 1$, this gives (recall that $|D_1| \asymp R$)

$$\left| \sum_{=1}(D_1) \cap (-\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))) \right| \ll |Y_1^+|^{1/2} \ll \binom{|D_1|}{1} \cdot R^{-1/2},$$

and for $2 \leq k \leq K$ (recall that $K = c_2 R^{1/3}$) it gives

$$\left| \sum_{=k}(D_1) \cap (-\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))) \right| \ll \binom{|D_1|}{k} \cdot \frac{K}{|D_1|} \ll \binom{|D_1|}{k} \cdot c_2^3 K^{-2}.$$

Since the conclusion of Lemma 4.16 also holds, we can “transfer” this bound from D_1 to T_1 . In particular,

we obtain that

$$\left| \sum_{=1} (T_1) \cap (-\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))) \right| \ll 4K \binom{|T_1|}{1} \cdot R^{-1/2}$$

and that

$$\left| \sum_{=k} (T_1) \cap (-\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))) \right| \ll 4K \binom{|T_1|}{k} \cdot \frac{K}{|D_1|} \ll 4K \binom{|T_1|}{k} \cdot c_2^3 K^{-2}$$

for $2 \leq k \leq K$. It follows that

$$\mathbb{P}(t_1 \in -\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))) \ll R^{-1/6}$$

is certainly at most 0.01, and for $2 \leq k \leq K$ we see that

$$\mathbb{P}(t_1 + \dots + t_k \in -\text{IS}(\mathbf{p}) \cup (\delta + \text{IS}(\mathbf{n}))) \ll c_2^3 K^{-1}$$

is at most $0.01K^{-1}$ as long as c_2 is sufficiently small. This completes the proof. $\square$

We now choose $\mathbf{t}_1, \mathbf{t}_u$ independently such that $\mathbf{t}_1, \overline{\mathbf{t}_u}$ are uniformly random acceptable orderings of T_1, T_u , respectively. We deduce from Lemma 4.17 that the random variables $\mathbf{t}_1, \mathbf{t}_u$ are highly anti-concentrated in the sense that the probability of $\mathbf{t}_1$ assuming any particular ordering is at most $\ll 1/m_1!$, and likewise for $\mathbf{t}_u$. Notice that the constraint that $\mathbf{t}_1, \overline{\mathbf{t}_u}$ are acceptable precisely guarantees the absence of zero-sum Type I intervals with one endpoint in $\overline{\mathbf{p}}$ or $\mathbf{n}$ and the other endpoint in the left border of $\mathbf{t}_1$ or the right border of $\mathbf{t}_u$.

For each $1 \leq j \leq u/2 - 1$, we choose the pair of orderings $\mathbf{t}_{2j}, \mathbf{t}_{2j+1}$ as follows. Recall that $|T_{2j}| = m_{2j}$ and $|T_{2j+1}| = m_{2j+1}$ both have size $\asymp R$. Say that a pair of partial orderings $t_1, \dots, t_k$ of T_{2j} and $t'_1, \dots, t'_\ell$ of T_{2j+1} is *permissible* if

$$t_1 + \dots + t_i + t'_1 + \dots + t'_j \neq 0 \quad \text{for all } (i, j).$$

Let $N(k, \ell)$ denote the number of permissible pairs with lengths (k, ℓ) . Note that each permissible pair with lengths (k, ℓ) can be extended to at least $m_{2j} - k - \ell$ permissible pairs of lengths $(k+1, \ell)$ and to at least $m_{2j+1} - k - \ell$ permissible pairs of lengths $(k, \ell+1)$. It follows that

$$N(k, k) \geq (m_{2j})(m_{2j+1} - 1)(m_{2j} - 2)(m_{2j+1} - 3) \cdots (m_{2j} - 2k + 2)(m_{2j+1} - 2k + 1).$$

The choice of K ensures that $N(K, K) \geq m_{2j}^K m_{2j+1}^K / 2$ (say), which means that the permissible pairs comprise at least a constant fraction of the total pairs. (In fact, this bound would continue to hold with K as large

as a small constant times $R^{1/2}$.)

We now choose $\mathbf{t}_{2j}, \mathbf{t}_{2j+1}$ to be a uniformly random pair of orderings of T_{2j}, T_{2j+1} conditional on the length- K prefixes of $\overline{\mathbf{t}_{2j}}, \mathbf{t}_{2j+1}$ forming a permissible pair of length (K, K) . Equivalently, we let $t_1, \dots, t_K$ and $t'_1, \dots, t'_K$ be a uniformly random permissible pair of orderings of T_{2j}, T_{2j+1} , and then we let $\mathbf{t}_{2j}$ be a uniformly random ordering of T_{2j} conditional on $\overline{\mathbf{t}_{2j}}$ beginning with $t_1, \dots, t_K$, and we let $\mathbf{t}_{2j+1}$ be a uniformly random ordering of T_{2j+1} conditional on $\mathbf{t}_{2j+1}$ beginning with $t'_1, \dots, t'_K$. We make these choices independently for different values of j , and independently of the choices of $\mathbf{t}_1, \mathbf{t}_u$. The following lemma shows that even though the random variables $\mathbf{t}_{2j}, \mathbf{t}_{2j+1}$ are dependent, they are “conditionally anti-concentrated” in the sense that if we condition on $\mathbf{t}_{2j}$ being any particular ordering, then the probability of $\mathbf{t}_{2j+1}$ assuming any particular ordering is still very small, and vice versa.

Lemma 4.18. *Choose $\mathbf{t}_{2j}, \mathbf{t}_{2j+1}$ according to the distribution described above. Then, conditional on $\mathbf{t}_{2j}$ assuming any particular ordering, the probability of $\mathbf{t}_{2j+1}$ assuming any particular ordering is $\ll 1/m_{2j+1}!$; likewise, conditional on $\mathbf{t}_{2j+1}$ assuming any particular ordering, the probability of $\mathbf{t}_{2j}$ assuming any particular ordering is $\ll 1/m_{2j}!$.*

Proof. We prove only the first statement. Let $\mathbf{u}_{2j}$ be any fixed ordering of T_{2j} . Let $\mathbf{u}_{2j}^{(K)}$ denote the ordering consisting of the first K elements of $\mathbf{u}_{2j}$. The number of permissible pairs $\mathbf{u}_{2j}^{(K)}, \mathbf{u}_{2j+1}^{(K)}$ with $\mathbf{u}_{2j+1}^{(K)}$ of length K is at least

$$(m_{2j+1} - K)^K \gg m_{2j+1}^K,$$

by our choices of R, K (see the above discussion of $N(k, \ell)$). Thus, the number of orderings $\mathbf{u}_{2j+1}$ of T_{2j+1} such that $\mathbf{u}_{2j}^{(K)}, \mathbf{u}_{2j+1}^{(K)}$ is a permissible pair is

$$\gg m_{2j+1}^K \cdot (m_{2j+1} - K)! \gg m_{2j+1}!.$$

The lemma now follows since each of these $\gg m_{2j+1}!$ orderings of T_{2j+1} is equally likely to occur as $\mathbf{t}_{2j+1}$, after we condition on $\mathbf{t}_{2j} = \mathbf{u}_{2j}$. $\square$

Notice that the constraints on the pairs $\mathbf{t}_{2j}, \mathbf{t}_{2j+1}$ guarantee the absence of zero-sum Type I intervals beginning in the right border of $\mathbf{t}_{2j}$ and ending in the left border of $\mathbf{t}_{2j+1}$.

We will show that if the orderings $\mathbf{t}_1, \dots, \mathbf{t}_u$ of $T_1, \dots, T_u$ are chosen randomly as above, then with high probability the ordering

$$a_1, \dots, a_{|A|} := \overline{\mathbf{p}}, \mathbf{t}_1, \dots, \mathbf{t}_u, \mathbf{n} \tag{4.7}$$

of A is two-sided valid, i.e., we have $\sum_{i \in I} a_i \neq 0$ for every nonempty proper interval $I \subseteq [|A|]$. The output of

Lemma 4.15 and the constraints on the random orderings $\mathbf{t}_j$ together guarantee that there are no zero-sum Type I intervals in the ordering (4.7); the reader can refer to remark (2) following Lemma 4.15 to see how we have covered all possible cases. It remains to verify that with high probability there are no zero-sum Type II intervals. The key point is that the sum $\sum_{i \in I} a_i$ for each Type II interval I is highly anti-concentrated because there is still enough randomness in the orderings $\mathbf{t}_j$; the following lemma makes this observation precise.

Lemma 4.19. *Let $I \subset [1, |A|]$ be a Type II interval, and let $a_1, \dots, a_{|A|} = \bar{\mathbf{p}}, \mathbf{t}_1, \dots, \mathbf{t}_u, \mathbf{n}$ be the random ordering (4.7) of A . Then for every $x \in \mathbb{F}_p$ we have*

$$\mathbb{P}\left(\sum_{i \in I} a_i = x\right) \leq e^{-\Omega(K \log R)}.$$

Proof. By definition, there exists some j such that I contains exactly k elements of T_j , where $K \leq k \leq |T_j| - K$. As in the first step of the proof of Lemma 4.15, we break the sum over I into the sum over the part intersecting T_j and the part not intersecting T_j and condition on $\mathbf{t}_i$ for all $i \neq j$. Lemma 4.18 ensures that even after this conditioning, the probability of $\mathbf{t}_j$ assuming any particular ordering is $\ll 1/m_j!$. Since the k -element subsets of T_j all have distinct sums, we see that the sum over the part of I intersecting T_j assumes each particular value with probability

$$\ll \binom{m_j}{k}^{-1} \leq \binom{m_j}{K}^{-1} \leq e^{-\Omega(K \log R)},$$

and the lemma follows. □

Recall that $K = c_2 R^{1/3}$ and that $R = R(A) \gg c_1 (\log p)^{3/4}$ holds when $|A| \leq e^{c(\log p)^{1/4}}$ (see (4.1)). Since the number of Type II intervals is trivially at most $|A|^2 \leq e^{2c(\log p)^{1/4}}$, Lemma 4.19 and the union bound imply that the probability of (4.7) containing a zero-sum Type II interval is at most

$$e^{2c(\log p)^{1/4} - \Omega(K \log R)} = o(1),$$

again with plenty of room to spare. From this and the above observations about the absence of zero-sum Type I intervals, we conclude that (4.7) is two-sided valid with high probability; in particular, for p sufficiently large (in terms of c), there is at least one two-sided valid ordering of A . This proves Theorem 4.2.

One can in fact take c to grow as, e.g., $\asymp \log \log p$, but we are not concerned with such lower-order terms since we have not even seriously optimized the exponent $1/4$ in Theorem 4.2.

4.7 Remarks and open problems

We make a couple of remarks about our proof of Theorem 4.2.

- The union bound in Lemma 4.16 is one of the main bottlenecks for the value of the exponent $1/4$ in Theorem 4.2. Improving the argument around this lemma would likely let one take K to be a larger power of R , which in turn would let one increase $1/4$ (perhaps to $1/3$) in Theorem 4.2.
- In proposition 4.7, we can also obtain the extra property that each of P, N is either empty or of size at least $100s$ (say), by splitting each dissociated set into 201 parts and then absorbing up to 100 elements of each of P, N if P, N are small. This property was useful in an earlier version of our proof and may be of interest in the future.

Our work also leads to several open problems for future inquiry:

- The most obvious open problem is improving the bound in Theorem 4.2; a natural next goal would be a polynomial threshold (of the form p^c). Even if our methods can be adapted to improve the exponent $1/4$ in Theorem 4.2, it seems that neither our probabilistic toolbox nor our dissociated set machinery is suited for sets of polynomial size, so substantial new inputs would be necessary to reach a polynomial threshold.
- Our arguments for Theorem 4.2 show not only that there is some two-sided valid ordering of A but that there are many such orderings. It would be interesting to estimate the minimum possible number of two-sided valid orderings as a function of $|A|$ (and perhaps also p).
- The main result of [66] applies not only to the group $\mathbb{F}_p$ but also to all groups of the form $H_1 \times H_2$, where H_1 is an abelian group such that every subset of $H_1 \setminus \{0\}$ has a two-sided valid ordering and H_2 is an abelian group with no non-zero elements of order strictly smaller than p ; one example is the group $\mathbb{Z}/2p\mathbb{Z} \cong \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/p\mathbb{Z}$. One could try to extend Theorem 4.2 to such groups.
- Finally, we mention that nonabelian versions of Graham's conjecture, particularly in dihedral groups, have received some attention. As in the abelian case, work prior to [66] concerned sets of size at most 12. Costa and Della Fiore [26] then adapted the ideas of [66] to obtain results for sets of nearly logarithmic size in dihedral and dicyclic groups. It seems more difficult to transfer the proof of Theorem 4.2 to nonabelian settings, and this could be a fruitful topic for future research.

Part II

Higher-order Fourier analysis

Chapter 5

Quantitative concatenation

5.1 Introduction

One of the starting points of arithmetic combinatorics is Roth’s Theorem [93], which says that every subset of the natural numbers with positive upper Banach density contains a nontrivial 3-term arithmetic progression $x, x + z, x + 2z$. Szemerédi [106] famously extended this result to ℓ -term arithmetic progressions for all ℓ . The farthest-reaching generalization of this line of work is the Polynomial Szemerédi Theorem of Bergelson and Leibman [9], which asserts that positive-upper-Banach-density subsets of $\mathbb{Z}^D$ always contain nontrivial “polynomial progressions” under certain weak local conditions on the polynomials involved.

The quantitative bounds from Roth’s original proof of his theorem have been improved many times, most recently by Kelley and Meka [61]. Szemerédi provided effective, albeit enormous, bounds for his theorem, and Gowers’ work [42] on obtaining more “reasonable” bounds inaugurated the field of higher-order Fourier analysis, where recent developments have led to improved bounds for Szemerédi’s Theorem [75]. The proof of Bergelson and Leibman [9], by contrast, uses ergodic-theoretic methods and is completely ineffective. Obtaining quantitative (and “reasonable”) bounds for the Bergelson–Leibman Theorem is a difficult problem that was highlighted by Gowers e.g. in [44].

In an early development, Sárközy [96,97] obtained quantitative bounds for nonlinear polynomial progressions of length two. Green [48] and Prendiville [90] later studied other particular examples of polynomial progressions, but their methods do not have obvious applications to larger families of polynomial progressions. Building on work of Peluse [84], Peluse and Prendiville [86] gave an effective polynomial Szemerédi theorem for the progression $x, x + z, x + z^2$. Peluse [83] then extended this method to obtain reasonable bounds for polynomial progressions involving polynomials with distinct degrees. Quite a few quantitative

polynomial Szemerédi theorems [8, 54, 67–69, 72, 74, 88] have been proven in recent years, and the result of the following chapter can be seen as progress in this direction.

Most arguments in this line of research have begun by using the PET induction scheme of Bergelson [7] to show that the counting operator for the associated polynomial progression is bounded by an appropriate Gowers norm or, more generally, an appropriate box norm.¹ Green and Tao [50] showed that iterated applications of the Cauchy–Schwarz inequality in the style of Gowers [42, 43] suffice to give Gowers-norm-control for positive-complexity linear configurations in any number of dimensions. For 1-dimensional polynomial progressions in the finite-field setting, a naïve application of the PET argument produces the required norm-control. For example, if the counting operator

$$\sum_{x, z \in \mathbb{Z}/N\mathbb{Z}} f_0(x) f_1(x+z) f_2(x+z^2)$$

is large for some 1-bounded functions $f_0, f_1, f_2 : \mathbb{Z}/N\mathbb{Z} \rightarrow \mathbb{C}$ (where N is a large prime), then a short PET argument gives that the Gowers norm $\|f_2\|_{U^3(\mathbb{Z}/N\mathbb{Z})}$ is also large, with polynomial dependences. In the integer setting, by contrast, the PET argument alone does not suffice. If the counting operator

$$\sum_{x, z \in \mathbb{Z}} f_0(x) f_1(x+z) f_2(x+z^2)$$

is large for some 1-bounded functions $f_0, f_1, f_2 : \mathbb{Z} \rightarrow \mathbb{C}$ supported on $[N] := \{1, \dots, N\}$ for some large $N \in \mathbb{N}$, then PET gives only the largeness of the quantity

$$\sum_{h_1, h_2 \in [N^{1/2}]} \sum_{m_1, m_2, m_3 \in [N^{1/2}]} \sum_x \Delta_{h_1 m_1, h_2 m_2, (h_1 + h_2) m_3} f_2(x).$$

This expression is essentially an average of many “arithmetic box norms” along progressions with common differences

$$h_1, h_2, h_1 + h_2. \tag{5.1}$$

For more complicated 1-dimensional polynomial progressions, the common differences arising are multilinear polynomials in parameters $h_1, h_2, \dots$. At this point, it is desirable to “concatenate” these norms into a single Gowers norm of higher degree, so that one can use the numerous techniques involving the inverse theory of the Gowers norms.

¹For the definitions of the concepts appearing in the introduction, including Gowers norms, box norms and multiplicative derivatives, we refer the reader to Section 5.3. The details of the PET induction scheme are elucidated in Section 5.4.

The first such concatenation result was proven by Tao and Ziegler [110]. Their approach was infinitary, however, and gave no quantitative bounds on the losses incurred while passing from an average of box norms to a single concatenated Gowers norm. In [81, 86], Peluse and Prendiville developed a different approach that could provide quantitative bounds². They showed, among other things, that the averages of box norms arising from applying PET to 1-dimensional polynomial progressions

$$x, x + P_1(z), \dots, x + P_\ell(z) \quad (5.2)$$

can be concatenated into a single Gowers norm that control the counting operator for (5.2) with only polynomial losses. This concatenation result was instrumental in their proof of quantitative bounds for subsets of $[N]$ lacking progressions (5.2) in the case where $P_1, \dots, P_\ell \in \mathbb{Z}[z]$ have distinct degrees.

The situation is even more complicated in higher dimensions. Consider, for example, the *multidimensional* polynomial pattern

$$(x_1, x_2), (x_1 + z^2 + z, x_2), (x_1, x_2 + z^2 + z). \quad (5.3)$$

This time, we would like to show that if

$$\sum_{x_1, x_2, z \in \mathbb{Z}} f_0(x_1, x_2) f_1(x_1 + z^2 + z, x_2) f_2(x_1, x_2 + z^2 + z) \quad (5.4)$$

is large for some 1-bounded functions $f_0, f_1, f_2 : \mathbb{Z} \rightarrow \mathbb{C}$ supported on $[N]^2$, then so is³

$$\sum_{h_1, \dots, h_{2t} \in \mathbb{Z}} \sum_{x_1, x_2 \in \mathbb{Z}} \Delta_{h_1 \mathbf{e}_2, \dots, h_t \mathbf{e}_2, h_{t+1}(\mathbf{e}_2 - \mathbf{e}_1), \dots, h_{2t}(\mathbf{e}_2 - \mathbf{e}_1)} f_2(x_1, x_2)$$

for some positive integer $t = O(1)$, where $\mathbf{e}_1, \mathbf{e}_2$ are the usual coordinate vectors in $\mathbb{Z}^2$. (From here, one can hope to use degree-lowering arguments to decrease t to its optimal value.) The usual PET argument gives control of (5.4) by an average of degree-7 box norms of f_2 in the directions

$$\begin{aligned} &2(h_2 + h_3)(\mathbf{e}_2 - \mathbf{e}_1) + 2h_1 \mathbf{e}_2, \quad 2h_2(\mathbf{e}_2 - \mathbf{e}_1) + 2h_1 \mathbf{e}_2, \quad 2h_3(\mathbf{e}_2 - \mathbf{e}_1) + 2h_1 \mathbf{e}_2, \\ &2h_1 \mathbf{e}_2, \quad 2(h_2 + h_3)(\mathbf{e}_2 - \mathbf{e}_1), \quad 2h_2(\mathbf{e}_2 - \mathbf{e}_1), \quad 2h_3(\mathbf{e}_2 - \mathbf{e}_1). \end{aligned} \quad (5.5)$$

In contrast to the single-dimensional examples described above, this example requires us to concatenate an average of many box norms involving linear combinations of *two* principal directions $\mathbf{e}_2, \mathbf{e}_2 - \mathbf{e}_1$ into a

²See also [87], in which Peluse, Prendiville and Shao extend the quantitative concatenation approach from [86] to obtain Gowers-norm-control for the configuration $(x_1, x_2), (x_1 + z, x_2), (x_1, x_2 + z^2)$.

³Naturally, we also want to obtain the largeness of certain averages of box norms of f_0, f_1 , in which the role of $\{\mathbf{e}_2, \mathbf{e}_2 - \mathbf{e}_1\}$ is played by $\{\mathbf{e}_1, \mathbf{e}_2\}$ and $\{\mathbf{e}_1, \mathbf{e}_1 - \mathbf{e}_2\}$ respectively.

single box norm involving these two principal directions. More generally, in the study of multidimensional polynomial progressions

$$\mathbf{x}, \mathbf{x} + \mathbf{P}_1(z), \dots, \mathbf{x} + \mathbf{P}_\ell(z) \quad (5.6)$$

for some choice of $\mathbf{P}_1, \dots, \mathbf{P}_\ell \in \mathbb{Z}^D[z]$, one is confronted with the task of concatenating averages of box norms along

$$\mathbf{C}_1(h_1, \dots, h_r), \dots, \mathbf{C}_s(h_1, \dots, h_r), \quad (5.7)$$

for some *multilinear* polynomials $\mathbf{C}_1, \dots, \mathbf{C}_s \in \mathbb{Z}^D[h_1, \dots, h_r]$ that arise from the application of PET to (5.6); the goal is to obtain control by a single box norm, which usually involves several principal directions.

Let us compare the directions in (5.1) with the directions in (5.5). The polynomial directions in (5.1), as examined by Peluse and Prendiville, are all multiples of the same principal direction in $\mathbb{Z}$. The polynomial directions in (5.6), by contrast, are linear combinations of the two principal directions $\mathbf{e}_2 - \mathbf{e}_1, \mathbf{e}_2$ that come from the leading coefficients of the polynomials in (5.3) and their pairwise differences. Thus the problem of quantitative concatenation of box norms gains a layer of difficulty when one passes from the 1-dimensional setting of Peluse and Prendiville to the multidimensional setting.

The goal of this chapter is to prove a new quantitative concatenation result that is robust enough to handle the averages of polynomial box norms output by the PET argument for multidimensional polynomial progressions. This results in the following box-norm-control for general multidimensional polynomial progressions.

Theorem 5.1. *Let $d, \ell \in \mathbb{N}$. There exists a positive integer $t = O_{d, \ell}(1)$ such that the following holds. Let $D, N \in \mathbb{N}$ and $\delta > 0$ satisfy $N \gg_{d, D, \ell} \delta^{-O_{d, \ell}(1)}$, and let $\mathbf{P}_1, \dots, \mathbf{P}_\ell \in \mathbb{Z}^D[z]$ be polynomials of the form $\mathbf{P}_j(z) = \sum_{i=0}^d \beta_{ji} z^i$ with coefficients of size $O(1)$. Then for all 1-bounded functions $f_0, \dots, f_\ell : \mathbb{Z}^D \rightarrow \mathbb{C}$ supported on $[N]^D$, the bound*

$$\left| \sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{z \in [N^{1/d}]} f_0(\mathbf{x}) \cdot f_1(\mathbf{x} + \mathbf{P}_1(z)) \cdots f_\ell(\mathbf{x} + \mathbf{P}_\ell(z)) \right| \geq \delta N^D \quad (5.8)$$

implies that

$$\|f_j\|_{\{(\beta_{jd_{jj'}} - \beta_{j'd_{jj'}}) \cdot [\pm N^{d_{jj'}/d}]: j' \in [0, \ell], j' \neq j\}^t}^{2^{\ell t}} \gg_{d, D, \ell} \delta^{O_{d, \ell}(1)} N^D \quad (5.9)$$

for all $j \in [0, \ell]$, where $d_{jj'} = \deg(\mathbf{P}_j - \mathbf{P}_{j'})$ and $\mathbf{P}_0 = \mathbf{0}$.

Informally, Theorem 5.1 asserts that every multidimensional polynomial progression can be controlled, with polynomial losses, by a single box norm of bounded degree along the “correct” boxes. Let us explain in words why the box norm in (5.9) is what one “should” expect, and specifically why it involves intervals of length $\asymp N^{d_{jj'}/d}$ in the directions $\beta_{jd_{jj'}} - \beta_{j'd_{jj'}}$. Consider the special case where all of the f_i ’s except for $f_j, f_{j'}$ are the constant function 1.⁴ Shifting $\mathbf{x} \mapsto \mathbf{x} - \mathbf{P}_{j'}(z)$, we find that (5.8) simplifies to

$$\sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{z \in [N^{1/d}]} f_{j'}(\mathbf{x}) f_j(\mathbf{x} + \mathbf{P}_j(z) - \mathbf{P}_{j'}(z)) \geq \delta N^D.$$

The polynomial $\mathbf{P}_j - \mathbf{P}_{j'}$ has degree $d_{jj'}$ and leading coefficient $\beta_{jd_{jj'}} - \beta_{j'd_{jj'}}$. Applying the classical van der Corput inequality $d_{jj'}$ times removes the lower-order terms of $\mathbf{P}_j - \mathbf{P}_{j'}$ and (more or less) results in the inequality

$$\sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{h_1, \dots, h_{d_{jj'}} \in [N^{1/d}]} f_j(\mathbf{x}) f_j(\mathbf{x} + d!(\beta_{jd_{jj'}} - \beta_{j'd_{jj'}})h_1 \cdots h_{d_{jj'}}) \gg_{d_{jj'}} \delta^{O_{d_{jj'}}(1)} N^D.$$

One can use iterated applications of the Cauchy–Schwarz inequality to replace the monomial $h_1 \cdots h_{d_{jj'}}$ with a single variable ranging over $[N^{d_{jj'}/d}]$, and this morally gives the largeness of a box norm of f_j along the progression $(\beta_{jd_{jj'}} - \beta_{j'd_{jj'}}) \cdot [\pm N^{d_{jj'}/d}]$. The large multiplicity t appearing in Theorem 5.1 is an artifact of the large number of applications of the Cauchy–Schwarz inequality in both the PET induction scheme and the concatenation arguments.

For $D > 1$, the results of Theorem 5.1 and the more general Theorem 5.6 below are completely novel. For $D = 1$, they recover Peluse’s concatenation result [81, Theorem 6.1] and also address some configurations that are not covered by Peluse’s work.

Example 5.2. As an example of a 1-dimensional configuration not covered by [81, Theorem 6.1], take $x, x + x^2, x + 2z^2, x + 2z^2 + z$. If

$$\left| \sum_{x \in \mathbb{Z}} \mathbb{E}_{z \in [N^{1/2}]} f_0(x) f_1(x + z^2) f_2(x + 2z^2) f_3(x + 2z^2 + z) \right| \geq \delta N,$$

then Theorem 5.1 (combined with standard properties of box norms listed in Lemma 5.15) gives a positive integer $t = O(1)$ for which

$$\|f_0\|_{U^t([\pm N])}^{2^t}, \|f_1\|_{U^t([\pm N])}^{2^t} \gg \delta^{O(1)} N \quad \text{and} \quad \|f_2\|_{U^t([\pm N^{1/2}])}^{2^t}, \|f_3\|_{U^t([\pm N^{1/2}])}^{2^t} \gg \delta^{O(1)} N.$$

⁴Obviously, such functions do not satisfy the assumption of being supported on $[N]^D$, but this example can be modified slightly to fit the assumptions of the more general Theorem 5.6 below.

The disparity in the scales of these Gowers norms comes from the polynomials $2z^2$, $2z^2 + z$ having the same leading coefficient: Since the difference of these polynomials is a linear polynomial that ranges over an interval of length $\asymp N^{1/2}$ rather than $\asymp N$, the shifting parameters in the Gowers norms controlling f_2 and f_3 must be at the smaller scale $N^{1/2}$. For configurations like this in which several polynomials have the same leading coefficients, [81, Theorem 6.1] does not give any nontrivial box-norm-control, so our result is novel in this case. We remark as a side note that such restricted Gowers norms are often difficult to use in applications because the inverse theorem for the $U^t([\pm N^{1/2}])$ -norm involves correlations with $(t-1)$ -step nilsequences on short intervals of length $\asymp N^{1/2}$, with no discernible global (i.e., scale- N) structure.

When the polynomials $\mathbf{P}_j$ have distinct degrees, the conclusion of Theorem 5.1 radically simplifies.

Corollary 5.3. *Let $d, \ell \in \mathbb{N}$. There exists a positive integer $t = O_{d,\ell}(1)$ such that the following holds. Let $D, N \in \mathbb{N}$ and $\delta > 0$ satisfy $N \gg_{d,D,\ell} \delta^{-O_{d,\ell}(1)}$, let $\mathbf{v}_1, \dots, \mathbf{v}_\ell \in \mathbb{Z}^D$ be nonzero vectors of size $O(1)$, and let $P_1, \dots, P_\ell \in \mathbb{Z}[z]$ be polynomials with coefficients of size $O(1)$ and degrees*

$$1 \leq \deg P_1 < \dots < \deg P_\ell = d.$$

Then for all 1-bounded functions $f_0, \dots, f_\ell : \mathbb{Z}^D \rightarrow \mathbb{C}$ supported on $[N]^D$, the bound

$$\left| \sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{z \in [N^{1/d}]} f_0(\mathbf{x}) \cdot f_1(\mathbf{x} + \mathbf{v}_1 P_1(z)) \cdots f_\ell(\mathbf{x} + \mathbf{v}_\ell P_\ell(z)) \right| \geq \delta N^D$$

implies that

$$\|f_\ell\|_{U^t(\mathbf{v}_\ell \cdot [\pm N])}^{2^t} \gg_{d,D,\ell} \delta^{O_{d,\ell}(1)} N^D.$$

Example 5.4. In particular, Corollary 5.3 recovers the Gowers-norm-control obtained by Peluse and Prendiville in [86]: It shows that if

$$\left| \sum_{x \in \mathbb{Z}} \mathbb{E}_{z \in [N^{1/2}]} f_0(x) f_1(x+z) f_2(x+z^2) \right| \geq \delta N,$$

then $\|f_2\|_{U^t([\pm N])}^{2^t} \gg \delta^{O(1)} N$ for some positive integer $t = O(1)$.

Example 5.5. Corollary 5.3 also recovers the Gowers-norm-control obtained by Peluse, Prendiville and

Shao in [87]: It shows that if

$$\left| \sum_{x_1, x_2 \in \mathbb{Z}} \mathbb{E}_{z \in [N^{1/2}]} f_0(x_1, x_2) f_1(x_1 + z, x_2) f_2(x_1, x_2 + z^2) \right| \geq \delta N,$$

then $\|f_2\|_{U^t(\mathbf{e}_2 \cdot [\pm N])}^{2^t} \gg \delta^{O(1)} N^2$ for some positive integer $t = O(1)$.

It is often desirable for technical reasons to consider polynomial progressions with “ W -tricked” polynomials whose coefficients grow (in a controlled way) with the parameter N . We in fact prove the following more general version of Theorem 5.1 which is flexible enough to accommodate W -tricks.

Theorem 5.6. *Let $d, \ell \in \mathbb{N}$. There exists a positive integer $t = O_{d, \ell}(1)$ such that the following holds. Let $D, K, N, V \in \mathbb{N}$ and $\delta > 0$ satisfy*

$$\delta^{-O_{d, \ell}(1)} \ll_{d, D, \ell} K \ll_{d, D, \ell} (N/V)^{1/d},$$

and let $\mathbf{P}_1, \dots, \mathbf{P}_\ell \in \mathbb{Z}^D[z]$ be polynomials of the form $\mathbf{P}_j(z) = \sum_{i=0}^d \beta_{ji} z^i$ with coefficients at most V .⁵ Then for all 1-bounded functions $f_0, \dots, f_\ell : \mathbb{Z}^D \rightarrow \mathbb{C}$ supported on $[N]^D$, the bound

$$\left| \sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{z \in [K]} f_0(\mathbf{x}) \cdot f_1(\mathbf{x} + \mathbf{P}_1(z)) \cdots f_\ell(\mathbf{x} + \mathbf{P}_\ell(z)) \right| \geq \delta N^D$$

implies that

$$\|f_j\|_{\{(\beta_{jd_{jj'}}, -\beta_{j'd_{jj'}}) \cdot [\pm K^{d_{jj'}}] : j' \in [0, \ell], j' \neq j\}^t}^{2^t} \gg_{d, D, \ell} \delta^{O_{d, \ell}(1)} N^D \quad (5.10)$$

for all $j \in [0, \ell]$, where $d_{jj'} = \deg(\mathbf{P}_j - \mathbf{P}_{j'})$ and $\mathbf{P}_0 = \mathbf{0}$.

When $\mathbf{P}_1, \dots, \mathbf{P}_\ell$ are all $\mathbb{Z}^D$ -multiples of a single integer polynomial, Theorem 5.6 reduces to the following.

Corollary 5.7. *Let $d, \ell \in \mathbb{N}$. There exists a positive integer $t = O_{d, \ell}(1)$ such that the following holds. Let $D \in \mathbb{N}$, let $\mathbf{v}_0, \dots, \mathbf{v}_\ell \in \mathbb{Z}^D$ be nonzero vectors of size $O(1)$ (with $\mathbf{v}_0 = \mathbf{0}$), and let $P \in \mathbb{Z}[z]$ be a degree- d polynomial with leading coefficient β_d . Then for all $N \in \mathbb{N}$ and 1-bounded functions $f_0, \dots, f_\ell : \mathbb{Z}^D \rightarrow \mathbb{C}$ supported on $[N]^D$, the bound*

$$\left| \sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{z \in [(N/\beta_d)^{1/d}]} f_0(\mathbf{x}) \cdot f_1(\mathbf{x} + \mathbf{v}_1 P(z)) \cdots f_\ell(\mathbf{x} + \mathbf{v}_\ell P(z)) \right| \geq \delta N^D$$

⁵In the sense that for every j, i , each coordinate of β_{ji} is at most V in absolute value, or equivalently that $\beta_{ji} \in [\pm V]^D$.

implies that

$$\|f_j\|_{\{\beta_d(\mathbf{v}_j - \mathbf{v}_{j'}) \cdot [\pm N/\beta_d] : j' \in [0, \ell], j' \neq j\}^t}^{2\ell t} \gg_{d, D, \ell} \delta^{O_{d, \ell}(1)} N^D$$

for all $j \in [0, \ell]$.

Example 5.8. If $D = 1$ and $\mathbf{v}_j = j$, Corollary 5.7 tells us that

$$\left| \sum_{x \in \mathbb{Z}} \mathbb{E}_{z \in [(N/\beta_d)^{1/d}]} f_0(x) f_1(x + P(z)) \cdots f_\ell(x + \ell P(z)) \right| \geq \delta N$$

implies that

$$\min_{j \in [0, \ell]} \|f_j\|_{U^t(\beta_d \cdot [\pm N/\beta_d])}^{2t} \gg_{d, \ell} \delta^{O_{d, \ell}(1)} N$$

for some positive integer $t = O_{d, \ell}(1)$, where we recall that β_d is the leading coefficient of P . This improves on the Gowers-norm-control obtained by Prendiville in [90] and recovers the global Gowers-norm-control used by Peluse, Sah and Sawhney in [88].

Example 5.9. If $\ell = D = 2$ and $\mathbf{v}_j = \mathbf{e}_j$, Corollary 5.7 tells us that

$$\left| \sum_{x_1, x_2 \in \mathbb{Z}} \mathbb{E}_{z \in [(N/\beta_d)^{1/d}]} f_0(x_1, x_2) f_1(x_1 + P(z), x_2) f_2(x_1, x_2 + P(z)) \right| \geq \delta N^2$$

implies that

$$\begin{aligned} \|f_0\|_{(\mathbf{e}_1 \beta_d \cdot [\pm N/\beta_d])^t, (\mathbf{e}_2 \beta_d \cdot [\pm N/\beta_d])^t}^{2t} &\gg_d \delta^{O_d(1)} N^2 \\ \|f_1\|_{(\mathbf{e}_1 \beta_d \cdot [\pm N/\beta_d])^t, ((\mathbf{e}_1 - \mathbf{e}_2) \beta_d \cdot [\pm N/\beta_d])^t}^{2t} &\gg_d \delta^{O_d(1)} N^2 \\ \|f_0\|_{(\mathbf{e}_2 \beta_d \cdot [\pm N/\beta_d])^t, ((\mathbf{e}_2 - \mathbf{e}_1) \beta_d \cdot [\pm N/\beta_d])^t}^{2t} &\gg_d \delta^{O_d(1)} N^2 \end{aligned}$$

for some positive integer $t = O(1)$. This result serves as the starting point for the following chapter.

Our new concatenation result (the key ingredient in the proof of Theorem 5.6) is the following.

Theorem 5.10. *Let $r, s \in \mathbb{N}$. There exists a positive integer $t = O_{r, s}(1)$ such that the following holds. Let $d, D, H, M, N, V \in \mathbb{N}$ and $\delta > 0$ satisfy*

$$d \leq r, \quad \delta^{-O_{\ell, r, s}(1)} \ll_{D, r, s} H \leq M \quad \text{and} \quad H^d M V \leq N,$$

and let $\mathbf{C}_1, \dots, \mathbf{C}_s \in \mathbb{Z}^D[h_1, \dots, h_r]$ be multilinear polynomials $\mathbf{C}_j(\underline{h}) = \sum_{\underline{u} \in \{0,1\}^r} \gamma_{j\underline{u}} \underline{h}^{\underline{u}}$ of degree at most d and with coefficients at most V . Then for all 1-bounded functions $f : \mathbb{Z}^D \rightarrow \mathbb{C}$ supported on $[N]^D$, the bound

$$\mathbb{E}_{\underline{h} \in [\pm H]^r} \|f\|_{\mathbf{C}_1(\underline{h}) \cdot [\pm M], \dots, \mathbf{C}_s(\underline{h}) \cdot [\pm M]}^{2^s} \geq \delta N^D$$

implies that

$$\|f\|_{E_1^t, \dots, E_s^t}^{2^{st}} \gg_{D,r,s} \delta^{O_{r,s}(1)} N^D,$$

where

$$E_j = \sum_{\underline{u} \in \{0,1\}^d} \gamma_{j\underline{u}} \cdot [\pm H^{|\underline{u}|+1}].$$

We conclude this introduction by mentioning a few places where our results have already found applications. The first two applications lie in ergodic theory: Our arguments inspired work of Donoso, Koutsogianis, Kuca, Sun and Tsinas [33] on ergodic averages along Hardy sequences, and Kosz, Mirek, Peluse, and Wright [64] have used our main PET result (Proposition 5.23 below). In analytic number theory, Green and Sawhney [49] used our main concatenation result to handle Type II sums in their study pairs of primes (p, q) such that $p^2 + 4q^2$ is also prime.

5.2 Proof strategy

Because the proofs of our main results are rather technical, we will begin with an informal discussion of how our techniques apply to a few simplified examples.

The proof of Theorem 5.6 begins with PET induction, a complexity-reduction argument that allows us to control the counting operator of a polynomial progression by an average of box norms. While PET arguments are now standard and have been carried out in many places (see, e.g., [7, 9, 32, 81, 90]), we need to keep careful track of the structure of the polynomial differencing parameters $\mathbf{C}_j$ that are output by PET. In general, the input of PET is the counting operator lower bound

$$\sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{z \in [N^{1/d}]} f_0(\mathbf{x}) \cdot f_1(\mathbf{x} + \mathbf{P}_1(z)) \cdots f_\ell(\mathbf{x} + \mathbf{P}_\ell(z)) \geq \delta N^D,$$

for some polynomials $\mathbf{P}_1, \dots, \mathbf{P}_\ell \in \mathbb{Z}^D[z]$ of degree at most d and some 1-bounded functions $f_0, \dots, f_\ell$

supported on $[N]^D$. The output of PET is then that the average of box norms

$$\mathbb{E}_{\underline{h} \in [\pm N^{1/d}]^r} \|f\|_{\mathbf{C}_1(\underline{h}) \cdot [\pm N^{1/d}], \dots, \mathbf{C}_s(\underline{h}) \cdot [\pm N^{1/d}]}^{2^s} \gg \delta^{O(1)} N^D \quad (5.11)$$

is also large. (Here and elsewhere, we allow all $O(1)$ terms to depend on d, ℓ .) Two crucial properties of this output are that the polynomials $\mathbf{C}_1, \dots, \mathbf{C}_s \in \mathbb{Z}^D[h_1, \dots, h_r]$ appearing above are multilinear and that their coefficients are $O(1)$ -multiples of the leading coefficients of $\mathbf{P}_1, \dots, \mathbf{P}_\ell$ and of their pairwise differences. This PET argument is carried out in detail in Section 5.4.

Sections 5.5–5.8 are then devoted to the proof of Theorem 5.10, our main concatenation result, which we will focus on in the remainder of this sketch. To differentiate between the ranges of the various parameters, let us rewrite (5.11) as

$$\mathbb{E}_{\underline{h} \in [\pm H]^r} \|f\|_{\mathbf{C}_1(\underline{h}) \cdot [\pm M], \dots, \mathbf{C}_s(\underline{h}) \cdot [\pm M]}^{2^s} \gg \delta N^D, \quad (5.12)$$

where we have set $H := M := N^{1/d}$ and replaced $\delta^{O(1)}$ with δ . (It will later be convenient to be able to consider slightly different values of H, M .)

We begin with the very simplest example, namely, where $d = 2$, $D = s = r = 1$, and $\mathbf{C}_1(h) = h$, so that our hypothesis is

$$\mathbb{E}_{h \in [\pm H]} \|f\|_{h \cdot [\pm M]}^2 \gg \delta N.$$

Ignoring smoothing terms, we can expand this as

$$\mathbb{E}_{h \in [\pm H]} \mathbb{E}_{m \in [\pm M]} \sum_{x \in \mathbb{Z}} f(x) \overline{f(x + hm)}.$$

If the quantity hm ranged over $[\pm HM]$ uniformly as h, m ranged, then we could replace hm by a new variable t and conclude that $\|f\|_{[\pm HM]}^2 \gg \delta N$. This is of course not the case, but our goal will be to reach a similar situation. Exchanging the order of summation and applying the Cauchy–Schwarz inequality twice, we obtain

$$\mathbb{E}_{h_1, h_2, h_3, h_4 \in [\pm H]} \mathbb{E}_{m_1, m_2, m_3, m_4 \in [\pm M]} \sum_{x \in \mathbb{Z}} f(x) \overline{f(x + h_1 m_1 + h_2 m_2 + h_3 m_3 + h_4 m_4)} \gg \delta^4 N.$$

An elementary number-theoretic argument (see Corollary 5.32 below) tells us that the quantity $h_1 m_1 + h_2 m_2 + h_3 m_3 + h_4 m_4$ assumes each value $t \in [\pm 4HM]$ at most $O(H^3 M^3)$ times as $h_1, h_2, h_3, h_4, m_1, m_2, m_3, m_4$ range; in other words, this quantity is nearly uniformly distributed, in the sense of being very anti-concentrated. The idea behind the lemma is that for most fixed choices of the h_i 's, the expression is well-distributed as the

m_i 's range; more precisely, this is the case unless $\gcd(h_1, h_2, h_3, h_4)$ is big or all of the h_i 's are small, and there are very few such “bad” choices. Together with the triangle inequality, this fact implies (after some work) that

$$\mathbb{E}_{t \in [\pm 4HM]} \left| \sum_{x \in \mathbb{Z}} f(x) \overline{f(x+t)} \right| \gg \delta^{O(1)} N.$$

Using the Cauchy–Schwarz inequality to remove the absolute value gives

$$\mathbb{E}_{t \in [\pm 4HM]} \sum_{x, s \in \mathbb{Z}} f(x) \overline{f(x+t)} f(x+s) \overline{f(x+s+t)} \gg \delta^{O(1)} N^2.$$

Since f is supported on $[N]$, we can restrict the s -variable to $[-N, N]$ (say), and renormalizing gives

$$\|f\|_{[\pm 4HM], [\pm N]}^4 = \mathbb{E}_{t \in [\pm 4HM]} \mathbb{E}_{s \in [\pm N]} \sum_{x \in \mathbb{Z}} f(x) \overline{f(x+t)} f(x+s) \overline{f(x+s+t)} \gg \delta^{O(1)} N.$$

So we have successfully concatenated the differencing sets $[\pm H], [\pm M]$ into a single differencing set $[\pm 4HM]$, at the cost of introducing an additional differencing over $[\pm N]$. Using monotonicity properties of box norms (see Lemma 5.15 below), we can “trim” $[\pm 4HM], [\pm N]$ and replace each of them with a copy of $[\pm HM]$.

For our next example, take $d = 3$, $r = 2$, $D = s = 1$, and $\mathbf{C}_1(h_1, h_2) = h_1 h_2$, so that our hypothesis is

$$\mathbb{E}_{h_1, h_2 \in [\pm H]} \|f\|_{h_1 h_2 \cdot [\pm M]}^2 \gg \delta N.$$

Expanding the left-hand side and using the Cauchy–Schwarz inequality to duplicate the differencing parameters, as in the first example, gives

$$\mathbb{E}_{h_{11}, \dots, h_{14} \in [\pm H]} \mathbb{E}_{h_{21}, \dots, h_{24} \in [\pm H]} \mathbb{E}_{m_1, \dots, m_4 \in [\pm M]} \sum_{x \in \mathbb{Z}} \overline{f(x) f(x + h_{11} h_{21} m_1 + h_{12} h_{22} m_2 + h_{13} h_{23} m_3 + h_{14} h_{24} m_4)} \gg \delta^4 N.$$

If the quantity $h_{11} h_{21} m_1 + h_{12} h_{22} m_2 + h_{13} h_{23} m_3 + h_{14} h_{24} m_4$ were nearly uniformly distributed on its range $[\pm H^2 M]$, then we could replace it by a single new variable ranging over $[\pm H^2 M]$ and conclude as in the first example. But this quantity fails to be sufficiently anti-concentrated, essentially because it is too often divisible by squares of small primes. The solution to this problem is to apply the Cauchy–Schwarz inequality several more times in order to replace each term $h_{2j} m_j$ with the expression

$$h_{2j1} m_{j1} + h_{2j2} m_{j2} + h_{2j3} m_{j3} + h_{2j4} m_{j4}$$

(at the cost of degrading the power of δ , which we do not care about). As the variables range, this expression is nearly uniformly distributed on $[\pm 4HM]$ (as in the previous example), and so with an acceptably small loss we can replace it by a new variable t_j that ranges over $[\pm 4HM]$. In turn, the expression

$$h_{11}t_1 + h_{12}t_2 + h_{13}t_3 + h_{14}t_4$$

is nearly uniformly distributed on $[\pm 16H^2M]$, and we conclude as in the first example.

This strategy of constructing large “expanding” polynomials with nested equidistribution properties suffices to handle the general case where $s = 1$ and $\mathbf{C}_1$ is a monomial. When $\mathbf{C}_1$ is not a monomial, we must show that several multilinear expressions are simultaneously nearly uniformly distributed. For example, with $d = 3$, $r = D = 2$, $s = 1$, and $\mathbf{C}_1(h_1, h_2) = h_1h_2\mathbf{e}_1 + h_1\mathbf{e}_2 = (h_1h_2, h_1)$, our hypothesis

$$\mathbb{E}_{h_1, h_2 \in [\pm H]} \|f\|_{(h_1h_2\mathbf{e}_1 + h_1\mathbf{e}_2) \cdot [\pm M]}^2 \gg \delta N$$

leads to the bound

$$\mathbb{E}_{\substack{h_i \in [\pm H]: \\ i \in [4]}} \mathbb{E}_{\substack{h_{ij} \in [\pm H]: \\ i, j \in [4]}} \mathbb{E}_{\substack{m_{ij} \in [\pm M]: \\ i, j \in [4]}} \sum_{\mathbf{x} \in \mathbb{Z}^2} \overline{f(\mathbf{x})f\left(\mathbf{x} + \left(\sum_{i \in [4]} h_i \sum_{j \in [4]} h_{ij} m_{ij}\right)\mathbf{e}_1 + \left(\sum_{i \in [4]} h_i \sum_{j \in [4]} m_{ij}\right)\mathbf{e}_2\right)} \gg \delta^{O(1)} N^2 \quad (5.13)$$

after the usual Cauchy–Schwarz maneuvers, so we wish to show that the $\mathbf{e}_1$ - and $\mathbf{e}_2$ -coefficients are simultaneously equidistributed. The first step is showing that for each $i \in [4]$, the expressions

$$\sum_{j \in [4]} h_{ij} m_{ij}, \quad \sum_{j \in [4]} m_{ij}$$

are simultaneously equidistributed. Dropping the i ’s from the indices, we need to show that for each pair $(n_0, n_1) \in [\pm 4M] \times [\pm 4HM]$, the system of equations

$$m_1 + m_2 + m_3 + m_4 = n_0, \quad h_1m_1 + h_2m_2 + h_3m_3 + h_4m_4 = n_1$$

has at most $O(H^3M^2)$ solutions. (Here $H^3M^2 = H^4M^4/(HM \cdot M)$.) The new idea here is that we can solve for m_1 in the second equation and make a substitution into the first equation, which (after regrouping) yields

$$(h_2 - h_1)m_2 + (h_3 - h_1)m_3 + (h_4 - h_1)m_4 = n_1 - h_1n_0.$$

Our previous arguments imply that for each of the $O(H)$ choices for h_1 , this equation has at most $O(H^2 M^2)$ solutions $(h_2, h_3, h_4, m_2, m_4, m_4)$. Such a solution can be extended to at most one choice of m_1 (sometimes there are no choices for m_1 because we constrained it to lie in $[\pm M]$), so in total we have at most $O(H^3 M^2)$ solutions, as desired.

Using the first step, we replace $\sum_{j \in [4]} h_{ij} m_{ij}, \sum_{j \in [4]} m_{ij}$ with new variables t_i, t'_i ranging over $[\pm 4HM], [\pm 4M]$, respectively, and it remains to show that the expressions

$$\sum_{i \in [4]} h_i t_i, \quad \sum_{i \in [4]} h_i t'_i$$

are simultaneously equidistributed. These expressions are individually identical to the expressions that we discussed above, but they are coupled in the sense that the variables h_i appear in both expressions. It turns out that the few “bad” choices of h_1, h_2, h_3, h_4 are now too costly because they lead to concentration in both expressions. The solution to this problem is to eliminate the bad choices before counting solutions: We can throw out the bad tuples (h_1, h_2, h_3, h_4) in (5.13) at the cost of degrading the implicit constant in the inequality.

To apply this type of multiscale analysis to multilinear expressions whose degrees differ by more than 1, we introduce dummy expressions at the intermediate degrees and then run the above argument for pairs of expressions with consecutive degrees. The genericity assumption on the h ’s is an essential technical element. The details, which are quite delicate, are carried out in Section 5.7.

Further difficulties arise for $s > 1$. In the $s = 1$ case, we iteratively used the Cauchy–Schwarz inequality to double certain variables. When $s > 1$, this maneuver also increases the number of multilinear expressions that must be simultaneously equidistributed—which in turn potentially requires further applications of the Cauchy–Schwarz inequality to replace these new expressions by ones with good equidistribution properties. Using arguments originally developed by Kuca [68] for concatenation in the finite-field setting, we show in Section 5.6 that this process can be satisfactorily controlled. This part of our proof uses no information about the “polynomial” nature of our box-norm parameters, and it ends up being convenient to carry out the argument in the more general setting of averages of box norms

$$\mathbb{E}_{i \in I} \|f\|_{H_{1i}, \dots, H_{si}}^{2^s}, \quad (5.14)$$

where $H_{1i}, \dots, H_{si}$ for $i \in I$ are arbitrary finite multisets of a finitely generated abelian group G that are contained in a bounded dilate of the support of f . We also believe that this level of generality might be useful for future extensions of our concatenation argument to other settings.

The key idea for handling (5.14), elucidated in Lemma 5.27, is to use the Cauchy–Schwarz and Gowers–Cauchy–Schwarz Inequalities in order to control (5.14) by

$$\mathbb{E}_{i_0, i_1 \in I} \|f\|_{H_{1i_0} - H_{1i_1}, H_{2i_0}, H_{2i_1}, \dots, H_{si_0}, H_{si_1}}^{2^{2s-1}};$$

we see that the multiset H_{1i} got replaced by a potentially much better distributed multiset $H_{1i_0} - H_{1i_1}$ at the cost of doubling the number of remaining multisets. Iterating this procedure many times, we can control (5.14) by an analogous expression in which each multiset H_{ji} is replaced by many iterated multisets of the form $H_{ji_1} + \dots + H_{ji_\ell}$. In the first example discussed above, this process corresponds to replacing hm by $h_1m_1 + h_2m_2 + h_3m_3 + h_4m_4$, and we have already established that the latter expression enjoys better equidistribution properties.

The last step of our argument, described in Section 5.8, specializes this general concatenation result to our “polynomial” setting and shows how it can be massaged to lead to the types of systems of multilinear equations described above. The results from Section 5.6 allow us to replace the boxes $\mathbf{C}_j(h) \cdot [\pm M]$ in (5.12) by iterated sums

$$\sum_{i_1, \dots, i_r=1}^{\ell} \mathbf{C}_j(h_{1i_1} \dots h_{ri_1 \dots i_r}) \cdot [\pm M]. \quad (5.15)$$

Each such box generates a system of multilinear equations, corresponding to the monomials in the polynomial $\mathbf{C}_j$. The arguments from Section 5.7 then give that if $\mathbf{C}_j(\underline{h}) = \sum_{\underline{u} \in \{0,1\}^r} \gamma_{j\underline{u}} \underline{h}^{\underline{u}}$, then (5.15) is uniformly distributed in the box $\sum_{\underline{u} \in \{0,1\}^d} \gamma_{j\underline{u}} \cdot [\pm H^{|\underline{u}|+1}]$. This concludes the proof of Theorem 5.10.

In Section 5.9, we combine the conclusions of Theorem 5.10 with the output of PET to prove Theorem 5.6. In doing so, we use the important observation from PET that the coefficients of the polynomials $\mathbf{C}_j$ are $O(1)$ -multiples of the leading coefficients of the polynomials $\mathbf{P}_1, \dots, \mathbf{P}_\ell$ and their pairwise differences.

Figure 5.1 provides a visual overview of the main steps of our argument.

5.2.1 Comparison with previous approaches

Our argument deviates in several ways from previous concatenation arguments in the literature: the quantitative concatenation arguments of Peluse and Prendiville [81, 86, 91] for 1-dimensional polynomial progressions over $\mathbb{Z}$; Kuca’s quantitative concatenation arguments for multidimensional polynomial progressions over finite fields [68]; and the original qualitative argument of Tao and Ziegler [110], together with its later development in the context of ergodic theory (e.g., [32]).

One difference between our argument and the approach of Peluse and Prendiville [81, 86, 91] lies in




Figure 5.1: The main steps and logical implications in the proof of Theorem 5.6.

the order of carrying out two major maneuvers, namely, duplicating variables using the Cauchy–Schwarz inequality, and replacing a multilinear polynomial like $h_1m_1 + h_2m_2 + h_3m_3 + h_4m_4$ by a new variable using equidistribution properties. As described in the sketch above, our argument clearly separates these two operations: First, we apply the Cauchy–Schwarz inequality many times to convert the original average (5.12) into one where each arithmetic progression $\mathbf{C}_j(\underline{h}) \cdot [\pm M]$ can be replaced by a generalized arithmetic progression $\mathbf{C}_j(\underline{h}_1) \cdot [\pm M] + \cdots + \mathbf{C}_j(\underline{h}_\ell) \cdot [\pm M]$ (interpreted as a multiset). Only then do we show that the massive system of multilinear forms arising out of this procedure is equidistributed in its range. In [81, 86, 91], these two steps are intertwined: Applications of the Cauchy–Schwarz inequality are followed by equidistribution results for intermediate systems of bilinear forms, then more applications of the Cauchy–Schwarz inequality, and so on. While the approach of Peluse and Prendiville has the advantage of using “simpler” equidistribution inputs (only systems of bilinear forms appear), our approach has the advantages of clearly separating the aforementioned two steps. Consequently, we obtain quantitative concatenation results (Corollaries 5.29 and 5.30) for averages of arbitrary box norms in arbitrary finitely-generated abelian groups; these more closely resemble the qualitative statements in the work of Tao and Ziegler [110].

Another important point of difference is the choice of induction scheme. In [81], Peluse begins the induction with degree-2 box norms as the base case (see [81, Lemma 5.5]), and she treats degree-1 box norms separately ([81, Lemma 5.4]). The strategy in [86, 91] is organized the same way. Here, in order to concatenate averages of degree-2 box norms, Peluse and Prendiville apply an inverse theorem for 2-dimensional box norms and then remove the resulting structured functions using their periodicity properties. For us, degree-1 box norms serve as the base case, which we analyze using a simple application of the Cauchy–Schwarz inequality. Consequently, our approach is more elementary and requires neither inverse theory nor periodicity properties.

Furthermore, dealing with multidimensional progressions forces us to keep careful track—more so than

in [81]—of how polynomial coefficients change during the PET procedure. Compared with the approach of [81], our coefficient-tracking scheme⁶ records more information about lower-degree terms of polynomials. It is precisely this extra information that allows Theorems 5.1 and 5.6 to extend Peluse’s result even in the single-dimensional setting (see Example 5.2).

Last but not least, whereas Peluse and Prendiville use basic Fourier analysis on several occasions, our argument is entirely carried out on the physical side. In particular, we do not use Fourier analysis in replacing equidistributed expressions like $h_1m_1 + h_2m_2 + h_3m_3 + h_4m_4$ by a new variable.

The relationship between our argument and the recent quantitative concatenation result over the finite fields by Kuca [68] is more complex. As already mentioned, our concatenation argument consists of two big steps: handling averages of box norms (5.14) along arbitrary multisets by means of the Cauchy–Schwarz inequality, and establishing equidistribution results for systems of multilinear equations. The first of these steps is largely based on an analogous argument from [68], and its adaptation to the integer setting requires only minor technical modifications. The two approaches deviate in the second step, in which we replace each monomial $h_1 \cdots h_r m$ by a new variable. This task is much simpler in the finite-field setting because the monomial expression in question is already very uniformly distributed; the integer setting requires several new ideas.

Like Tao and Ziegler [110], we ultimately prove a concatenation result for general averages of box norms (5.14) (see Corollary 5.29). The methods, however are completely different. Tao and Ziegler proceed by qualitatively analyzing the interactions among the dual functions implicit in the definition of box norms. Our proof of Corollary 5.29, by contrast, gives some truth to the stereotype that higher-order Fourier analysis largely consists of clever applications of the Cauchy–Schwarz inequality. Tao and Ziegler obtain better control on the degree of the “concatenated” box norms, but their conclusion is purely qualitative.

5.3 Definitions and preliminaries

5.3.1 Basic notation

We write $\mathbb{N}, \mathbb{N}_0, \mathbb{Z}, \mathbb{R}, \mathbb{C}$ for the sets of positive integers, nonnegative integers, integers, real numbers, and complex numbers (respectively). For reals $a < b$, we set $[a, b] = \{[a], a + 1, \dots, [b]\}$, abbreviating $[1, N]$ as $[N]$ and $[-N, N]$ as $[\pm N]$.

To facilitate the reading of the chapter, we fix the uses of various letters. The letter d denotes the maximum degree of polynomials under consideration; D is the dimension of the ambient space $\mathbb{Z}^D$ in which

⁶Our scheme is strongly inspired by a similar scheme from [32], but it improves on the latter in that we are able to identify fewer (and more relevant) properties that need to be tracked throughout PET.

we work; K measures the range of the polynomial variable z ; ℓ is the length of the polynomial family, and later on it also denotes the number of summands in concatenated box norms; N refers to the size of the support of the functions; r is the number of parameters $h_1, \dots, h_r$, and H concerns their range; s, t denote the degrees of box norms appearing in the arguments; and V bounds the size of polynomial coefficients.

We write $\mathbf{x} = (x_1, \dots, x_D)$ for an element of $\mathbb{Z}^D$ (for some fixed $D \in \mathbb{N}$) as x for an element of $\mathbb{Z}$. When the context is clear, we often write a sum over $\mathbb{Z}^D$ or $\mathbb{Z}$ as $\sum_{\mathbf{x}}$ or $\sum_x$ instead of $\sum_{\mathbf{x} \in \mathbb{Z}^D}$ or $\sum_{x \in \mathbb{Z}}$.

We usually write $\underline{h} = (h_1, \dots, h_r)$ for a tuple of length r , and we denote its ℓ^1 -norm by $|\underline{h}| := |h_1| + \dots + |h_r|$. For $\underline{u} \in \mathbb{N}_0^r$, we let $\text{supp}(\underline{u}) := \{i \in [r] : u_i > 0\}$ and get the monomial $\underline{h}^{\underline{u}} := h_1^{u_1} \dots h_r^{u_r}$. Given $\underline{h}, \underline{h}' \in \mathbb{Z}^r$ and $\underline{\epsilon} \in \{0, 1\}^r$, we also set $\underline{h}^{\underline{\epsilon}} := (h_1^{\epsilon_1}, \dots, h_r^{\epsilon_r})$, where $h_i^{\epsilon_i} := h_i$ if $\epsilon_i = 0$ and $h_i^{\epsilon_i} := h'_i$ if $\epsilon_i = 1$.

We frequently work with multisets.

Definition 5.11 (Multisets). For a finite multiset X , we write $|X|$ for the total number of elements in X , counted with multiplicity. We use $\text{supp}(X)$ to denote the set of elements appearing in X with strictly positive multiplicity. With regard to containment of multisets, we write $X \subseteq Y$ if the set $\text{supp}(X)$ is contained in the set $\text{supp}(Y)$, and we denote $X \leq Y$ if each element of $\text{supp}(X)$ has multiplicity at least as large in Y as in X . Lastly, for multisets X, Y , we define the sumset $X + Y$ to be the Minkowski sum of X, Y with multiplicities: The multiplicity of the element z in $X + Y$ is $|\{(x, y) \in X \times Y : z = x + y\}|$. Equivalently, the weight function of $X + Y$ is the convolution of the weight functions of X, Y . Note that $|X + Y| = |X| \cdot |Y|$. The difference set $X - Y$ is defined analogously.

Definition 5.12 (Indicator function). We use $\mathbf{1}(E)$ to denote the indicator function of an event E . For instance, $\mathbf{1}(a + b = c)$ is the function (of a, b, c) that equals 1 when $a + b = c$ and equals 0 when $a + b \neq c$. For a multiset E , we also abbreviate $\mathbf{1}(x \in \text{supp}(E))$ by $\mathbf{1}_E(x)$.

Let X be a nonempty multiset. If X is finite, we let $\mathbb{E}_{x \in X} := \frac{1}{|X|} \sum_{x \in X}$ be the average over X .

We call a function $f : X \rightarrow \mathbb{C}$ *1-bounded* if $\|f\|_{\infty} := \sup_{x \in X} |f(x)| \leq 1$.

We use the following standard asymptotic notation. If $f, g : X \rightarrow \mathbb{C}$, with g taking positive real values, we denote $f = O(g)$, $f \ll g$, $g \gg f$, or $g = \Omega(f)$ if there exists $C > 0$ such that $|f(x)| \leq Cg(x)$ for all $x \in X$. If the constant C depends on a parameter, we record this dependence with a subscript. Lastly, we write $f \asymp g$ if $f \ll g$ and $f \gg g$.

We let $\mathcal{C} : \mathbb{C} \rightarrow \mathbb{C}$ denote the complex conjugation operator $\mathcal{C}z = \bar{z}$.

5.3.2 Van der Corput Inequalities

If G is a finitely generated abelian group and $E \subseteq G$ is a finite and nonempty multiset, we define the *Fejér kernel*

$$\mu_E(h) := \mathbb{E}_{h_1, h_2 \in E} \mathbf{1}(h = h_1 - h_2),$$

and also $\mu_E(h_1, \dots, h_r) = \mu_E(h_1) \cdots \mu_E(h_r)$. Summing over h , we easily observe that $\sum_h \mu_E(h) = 1$. Moreover, if there is some $C > 0$ such that every element of E has multiplicity at most $C \frac{|E|}{|\text{supp}(E)|}$, then we also have the pointwise bound $0 \leq \mu_E(h) \leq \mathbf{1}_{E-E}(h) \cdot \frac{C}{|\text{supp}(E)|}$.⁷ When $E = [\pm H]$, we use the shorthand notation

$$\mu_H(h) := \mu_{[\pm H]}(h) = \mathbb{E}_{h_1, h_2 \in [\pm H]} \mathbf{1}_{h=h_1-h_2} = \frac{1}{2H+1} \left(1 - \frac{|h|}{2H+1} \right)_+.$$

This definition differs slightly from those in [81, 85, 86, 90] in that we take the ranges of h_1, h_2 to be $[\pm H]$ rather than $[H]$. Making the range of h_1, h_2 symmetric around the origin will be convenient for us later.

We record several versions of the standard van der Corput inequality. The proof of the following basic version is standard (see, e.g., [90, Lemma 3.2]).

Lemma 5.13 (van der Corput inequality, version 1). *Let $f : \mathbb{Z} \rightarrow \mathbb{C}$ be 1-bounded, and let $H, K \in \mathbb{N}$. Then:*

1. (*Symmetric van der Corput inequality*)

$$\left| \mathbb{E}_{z \in [K]} f(z) \right|^2 \leq \frac{K+2H}{K} \mathbb{E}_{h, h' \in [\pm H]} \frac{1}{K} \sum_{z \in ([K]-h) \cap ([K]-h')} f(z+h) \overline{f(z+h')}.$$

2. (*Asymmetric van der Corput inequality*)

$$\left| \mathbb{E}_{z \in [K]} f(z) \right|^2 \leq \frac{K+2H}{K} \sum_h \mu_H(h) \frac{1}{K} \sum_{z \in [K] \cap ([K]-h)} f(z) \overline{f(z+h)}.$$

Of course, it does not make sense to take $H \geq K$ since $[K] \cap ([K]-h)$ is empty whenever $|h| \geq K$.

In applications, we will use the following corollary of Lemma 5.13.

⁷The lower bound is obvious; for the upper bound, write

$$\mu_E(h) = \mathbf{1}_{E-E}(h) \cdot \frac{1}{|E|^2} \sum_{h_1 \in E} \sum_{h_2 \in E} \mathbf{1}(h_2 = h_1 + h),$$

and notice that the sum over h_2 is bounded by the maximum multiplicity of an element of E .

Lemma 5.14 (van der Corput inequality, version 2). *Let $0 < \delta \leq 1$ and $H, K \in \mathbb{N}$ satisfy $H \leq \frac{\delta^2}{4}K$. If $f : \mathbb{Z} \rightarrow \mathbb{C}$ is a 1-bounded function satisfying*

$$\left| \mathbb{E}_{z \in [K]} f(z) \right| \geq \delta, \quad (5.16)$$

then

$$\mathbb{E}_{h, h' \in [\pm H]} \mathbb{E}_{z \in [K]} f(z+h) \overline{f(z+h')} \geq \frac{\delta^2}{4}$$

and

$$\sum_h \mu_H(h) \mathbb{E}_{z \in [K]} f(z) \overline{f(z+h)} \geq \frac{\delta^2}{4}.$$

Proof. We prove only the second inequality since the proof of the first is almost identical. Since $K - |[K] \cap ([K] - h)| \leq |h| \leq H \leq \frac{\delta^2}{4}K$, we have

$$\left| \sum_h \mu_H(h) \sum_{z \in [K]} f(z) \overline{f(z+h)} - \sum_h \mu_H(h) \sum_{z \in [K] \cap ([K] - h)} f(z) \overline{f(z+h)} \right| \leq \frac{\delta^2}{4}K.$$

It then follows Lemma 5.13 and the assumption (5.16) that

$$\sum_h \mu_H(h) \mathbb{E}_{z \in [K]} f(z) \overline{f(z+h)} \geq \frac{K}{K+2H} \left| \mathbb{E}_{z \in [K]} f(z) \right|^2 - \frac{\delta^2}{4} \geq \frac{\delta^2}{4},$$

as desired. □

5.3.3 Box norms

In this section, we summarize the theory of box norms on an arbitrary finitely generated abelian group G ; throughout, we write $\sum_x := \sum_{x \in G}$. Given $h, h' \in G$ and a finitely supported function $f : G \rightarrow \mathbb{C}$, we define the (multiplicative) discrete derivatives $\Delta_h f(x) := f(x) \overline{f(x+h)}$ and $\Delta'_{(h, h')} f(x) := f(x+h) \overline{f(x+h')}$. Similarly, compositions of such discrete derivatives are given by

$$\begin{aligned} \Delta_{h_1, \dots, h_s} f(x) &= \Delta_{h_1} \cdots \Delta_{h_s} f(x) = \prod_{\underline{\epsilon} \in \{0,1\}^s} C^{|\underline{\epsilon}|} f(x + \underline{\epsilon} \cdot \underline{h}), \\ \Delta'_{(h_1, h'_1), \dots, (h_s, h'_s)} f(x) &= \Delta'_{(h_1, h'_1)} \cdots \Delta'_{(h_s, h'_s)} f(x) = \prod_{\underline{\epsilon} \in \{0,1\}^s} C^{|\underline{\epsilon}|} f(x + (\underline{1} - \underline{\epsilon}) \cdot \underline{h} + \underline{\epsilon} \cdot \underline{h'}), \end{aligned}$$

where again $\mathcal{C}z = \bar{z}$ for $z \in \mathbb{C}$. We will later also use the notation

$$\Delta'_{\alpha_1, \dots, \alpha_s; (\underline{h}, \underline{h}') := \Delta'_{(\alpha_1 h_1, \alpha_1 h'_1), \dots, (\alpha_s h_s, \alpha_s h'_s)}$$

for $\underline{h}, \underline{h}' \in G^s$. For finite multisets $E_1, \dots, E_s \subseteq G$, we define the *box norm* of f along $E_1, \dots, E_s$ to be

$$\|f\|_{E_1, \dots, E_s} := \left(\mathbb{E}_{h_1, h'_1 \in E_1} \cdots \mathbb{E}_{h_s, h'_s \in E_s} \sum_x \Delta'_{(h_1, h'_1), \dots, (h_s, h'_s)} f(x) \right)^{1/2^s}.$$

We can equivalently express this as

$$\|f\|_{E_1, \dots, E_s} = \left(\sum_{x, h_1, \dots, h_s} \mu_{E_1}(h_1) \cdots \mu_{E_s}(h_s) \Delta_{h_1, \dots, h_s} f(x) \right)^{1/2^s} \quad (5.17)$$

after shifting $x \mapsto x - h'_1 - \cdots - h'_s$ and removing the extraneous averages over $h'_1, \dots, h'_s$. If a multiset E repeats s times, we often write E^s instead of listing it s times, and we also write $\|f\|_{U^s(E)} = \|f\|_{E^s}$, calling it the degree- s *Gowers norm* of f along E .

For functions $f_{\underline{\epsilon}}$ indexed by $\underline{\epsilon} \in \{0, 1\}^s$, we also define the *Gowers–Cauchy–Schwarz inner product* (or *box inner product*) along $E_1, \dots, E_s$ to be

$$\begin{aligned} \langle (f_{\underline{\epsilon}})_{\underline{\epsilon} \in \{0, 1\}^s} \rangle_{E_1, \dots, E_s} &:= \mathbb{E}_{h_1, h'_1 \in E_1} \cdots \mathbb{E}_{h_s, h'_s \in E_s} \sum_x \prod_{\underline{\epsilon} \in \{0, 1\}^s} \mathcal{C}^{|\underline{\epsilon}|} f_{\underline{\epsilon}}(x + (1 - \underline{\epsilon}) \cdot \underline{h} + \underline{\epsilon} \cdot \underline{h}') \\ &= \sum_{x, h_1, \dots, h_s} \mu_{E_1}(h_1) \cdots \mu_{E_s}(h_s) \prod_{\underline{\epsilon} \in \{0, 1\}^s} \mathcal{C}^{|\underline{\epsilon}|} f_{\underline{\epsilon}}(x + \underline{\epsilon} \cdot \underline{h}). \end{aligned}$$

The classical *Gowers–Cauchy–Schwarz inequality*⁸ then states that

$$\langle (f_{\underline{\epsilon}})_{\underline{\epsilon} \in \{0, 1\}^s} \rangle_{E_1, \dots, E_s} \leq \prod_{\underline{\epsilon} \in \{0, 1\}^s} \|f_{\underline{\epsilon}}\|_{E_1, \dots, E_s}. \quad (5.18)$$

Box norms satisfy a number of well known properties, listed in the following lemma.

Lemma 5.15 (Properties of box norms). *Let $\delta > 0$, let $d, s \in \mathbb{N}$, let G be a finitely generated abelian group, and let $E_1, \dots, E_s \subseteq G$ be finite and nonempty multisets. Let $f : G \rightarrow \mathbb{C}$ be 1-bounded and supported on a finite nonempty set B . Then the following properties hold:*

⁸Even though this result is usually stated only for sets, exactly the same proof works for multisets.

1. (Inductive formula) For any $k \in [s]$, we have

$$\|f\|_{E_1, \dots, E_s}^{2^s} = \mathbb{E}_{h_1, h'_1 \in E_1} \cdots \mathbb{E}_{h_k, h'_k \in E_k} \left\| \Delta'_{(h_1, h'_1), \dots, (h_k, h'_k)} f \right\|_{E_{k+1}, \dots, E_s}^{2^{s-k}}.$$

2. (Permutation invariance) For any permutation $\sigma : [s] \rightarrow [s]$, we have

$$\|f\|_{E_1, \dots, E_s} = \|f\|_{E_{\sigma(1)}, \dots, E_{\sigma(s)}}.$$

3. (Monotonicity) If $|\text{supp}(B - E_s)| \leq C|B|$ for some $C > 0$, then

$$\|f\|_{E_1, \dots, E_{s-1}}^{2^{s-1}} \geq \delta|B| \implies \|f\|_{E_1, \dots, E_s}^{2^s} \geq \delta^2|B|/C. \quad (5.19)$$

4. (Enlarging the sets) If $s \geq 2$ and $E'_i \subseteq G$ is a finite multiset satisfying $E'_i \geq E_i$ for each $i \in \{1, \dots, s\}$, then

$$\|f\|_{E_1, \dots, E_s} \leq \left(\frac{|E'_1| \cdots |E'_s|}{|E_1| \cdots |E_s|} \right)^{1/2^{s-1}} \|f\|_{E'_1, \dots, E'_s}.$$

5. (Trimming the dimensions of boxes)⁹ Let $d'_i \in [d_i]$ for $i \in [s]$, and let

$$E_i = \sum_{j \in [d_i]} \beta_{ij} \cdot [\pm H_{ij}] \quad \text{and} \quad E'_i = \sum_{j \in [d'_i]} \beta_{ij} \cdot [\pm H_{ij}]$$

for some $\beta_{ij} \in G$ and $H_{ij} \in \mathbb{N}_0$. Then

$$\|f\|_{E_1, \dots, E_s}^{2^s} \geq \delta|B| \implies \|f\|_{E'_1, \dots, E'_s}^{2^s} \geq \delta^{2^s}|B|.$$

6. (Trimming the lengths of boxes) Let

$$E_i = \sum_{j \in [d]} \beta_{ij} \cdot [\pm H_{ij}] \quad \text{and} \quad E'_i = \sum_{j \in [d]} \beta_{ij} \cdot [\pm \kappa_{ij} H_{ij}]$$

for some $\beta_{ij} \in G$, $\kappa_{ij} > 0$, and $H_{ij} \in \mathbb{N}_0$. Then

$$\|f\|_{E_1, \dots, E_s}^{2^s} \geq \delta|B| \implies \|f\|_{E'_1, \dots, E'_s}^{2^s} \gg_{d,s} \delta^{2^s} \min(1, \delta^{2^s}/\kappa)^{2ds}|B|$$

⁹Although in this and the next property we have taken a single parameter d , setting some H_{ij} 's to be 0 lets us apply these results to the case where different boxes E_i have different dimensions.

for $\kappa = \max_{i,j} \kappa_{ij}$.

7. (Passing to sub-arithmetic progressions) Suppose that

$$E_i = \sum_{j \in [d]} \beta_{ij} \cdot [\pm H_{ij}] \quad \text{and} \quad E'_i = \sum_{j \in [d]} q_{ij} \beta_{ij} \cdot [\pm H_{ij}/q_{ij}]$$

for some $\beta_{ij} \in G$ and $q_{ij}, H_{ij} \in \mathbb{N}$ satisfying $q_{ij} \ll_{d,s} \delta^{2^s} H_{ij}$ for all i, j . Then

$$\|f\|_{E_1, \dots, E_s}^{2^s} \geq \delta |B| \implies \|f\|_{E'_1, \dots, E'_s}^{2^s} \gg_{d,s} \delta |B|.$$

In applications, we will take $G = \mathbb{Z}^D$, $B = [N]^D$, and E_i, E'_i will be generalized arithmetic progressions of bounded dimension (interpreted appropriately as multisets).

Proof. The first two properties are standard, and we omit the proofs. For the third property, we expand the definition of $\|f\|_{E_1, \dots, E_{s-1}}$ and introduce an extra averaging over E_s to obtain

$$\|f\|_{E_1, \dots, E_{s-1}}^{2^{s-1}} = \mathbb{E}_{h_1, h'_1 \in E_1} \cdots \mathbb{E}_{h_{s-1}, h'_{s-1} \in E_{s-1}} \sum_x \mathbb{E}_{h_s \in E_s} \Delta'_{(h_1, h'_1), \dots, (h_{s-1}, h'_{s-1})} f(x + h_s).$$

The term $f(x + h_1 + \cdots + h_s)$ in $\Delta'_{(h_1, h'_1), \dots, (h_{s-1}, h'_{s-1})} f(x + h_s)$ vanishes unless $x + h_1 + \cdots + h_{s-1} \in B - E_s$, so we can introduce the indicator function $\mathbf{1}_{B-E_s}(x + h_1 + \cdots + h_{s-1})$ into the expression above without altering its value. Applying the Cauchy-Schwarz inequality then gives

$$\begin{aligned} \delta^2 |B|^2 &\leq \|f\|_{E_1, \dots, E_{s-1}}^{2^s} \leq \mathbb{E}_{h_1, h'_1 \in E_1} \cdots \mathbb{E}_{h_{s-1}, h'_{s-1} \in E_{s-1}} \sum_x \mathbf{1}_{B-E_s}(x + h_1 + \cdots + h_{s-1}) \\ &\quad \times \mathbb{E}_{h_1, h'_1 \in E_1} \cdots \mathbb{E}_{h_s, h'_s \in E_s} \sum_x \Delta'_{(h_1, h'_1), \dots, (h_s, h'_s)} f(x) \\ &\leq C |B| \cdot \|f\|_{E_1, \dots, E_s}^{2^s}, \end{aligned}$$

and the claim follows.

For the fourth property, we write

$$\|f\|_{E_1, \dots, E_s}^{2^s} = \mathbb{E}_{h_1, h'_1 \in E_1} \cdots \mathbb{E}_{h_{s-1}, h'_{s-1} \in E_{s-1}} \sum_x \left| \mathbb{E}_{h_s \in E_s} \Delta'_{(h_1, h'_1), \dots, (h_{s-1}, h'_{s-1})} f(x + h_s) \right|^2.$$

Using the assumption $s \geq 2$ and nonnegativity, we can extend the averages over $E_1, \dots, E_{s-1}$ to $E'_1, \dots, E'_{s-1}$

and upper-bound the previous expression by

$$\left(\frac{|E'_1| \cdots |E'_{s-1}|}{|E_1| \cdots |E_{s-1}|} \right)^2 \cdot \mathbb{E}_{h_1, h'_1 \in E'_1} \cdots \mathbb{E}_{h_{s-1}, h'_{s-1} \in E'_{s-1}} \sum_x \left| \mathbb{E}_{h_s \in E_s} \Delta'_{(h_1, h'_1), \dots, (h_{s-1}, h'_{s-1})} f(x + h_s) \right|^2.$$

We can replace average over E_s with an average over E'_s using the same maneuver.

To prove the fifth property, we start with the case $s = 1$. Dropping the index i , we can write

$$\|f\|_E^2 = \sum_x \mathbb{E}_{\substack{h'_j \in [\pm H_j]: \\ j \in [d]}} f \left(x + \sum_{j \in [d]} \beta_j h'_j \right) \overline{\mathbb{E}_{\substack{h_j \in [\pm H_j]: \\ j \in [d]}} f \left(x + \sum_{j \in [d]} \beta_j h_j \right)}.$$

Using the Cauchy–Schwarz inequality to double the variables $(h_j)_{j \in [d']}$, we bound

$$\|f\|_E^4 \leq |B| \cdot \mathbb{E}_{\substack{h_j \in [\pm H_j]: \\ j \in [d'+1, d]}} \sum_x \mathbb{E}_{\substack{h_j, h'_j \in [\pm H_j]: \\ j \in [d']}} f \left(x + \sum_{j \in [d'+1, d]} \beta_j h_j + \sum_{j \in [d']} \beta_j h_j \right) \overline{f \left(x + \sum_{j \in [d'+1, d]} \beta_j h_j + \sum_{j \in [d']} \beta_j h'_j \right)}.$$

Shifting $x \mapsto x - \sum_{j \in [d'+1, d]} \beta_j h_j$, we can ignore the contribution of $j \in [d' + 1, d]$ altogether, and we obtain

$$\|f\|_E^4 \leq |B| \cdot \sum_x \mathbb{E}_{\substack{h_j, h'_j \in [\pm H_j]: \\ j \in [d']}} f \left(x + \sum_{j \in [d']} \beta_j h_j \right) \overline{f \left(x + \sum_{j \in [d']} \beta_j h'_j \right)} = |B| \cdot \|f\|_{E'}^2.$$

We now consider the case $s > 1$. The induction formula for box norms, the Cauchy–Schwarz inequality, and the $s = 1$ case give

$$\begin{aligned} \|f\|_{E_1, \dots, E_s}^{2^{s+1}} &\leq \mathbb{E}_{h_1, h'_1 \in E_1} \cdots \mathbb{E}_{h_{s-1}, h'_{s-1} \in E_{s-1}} \left\| \Delta'_{(h_1, h'_1), \dots, (h_{s-1}, h'_{s-1})} f \right\|_{E_s}^4 \\ &\leq |B| \cdot \mathbb{E}_{h_1, h'_1 \in E_1} \cdots \mathbb{E}_{h_{s-1}, h'_{s-1} \in E_{s-1}} \left\| \Delta'_{(h_1, h'_1), \dots, (h_{s-1}, h'_{s-1})} f \right\|_{E'_s}^2 \\ &= |B| \cdot \|f\|_{E_1, \dots, E_{s-1}, E'_s}^{2^s}. \end{aligned}$$

Performing the same procedure with E_{s-1} in place of E_s gives

$$\|f\|_{E_1, \dots, E_s}^{2^{s+2}} \leq |B|^2 \cdot \|f\|_{E_1, \dots, E_{s-1}, E'_s}^{2^{s+1}} \leq |B|^3 \cdot \|f\|_{E_1, \dots, E_{s-2}, E'_{s-1}, E'_s}^{2^s},$$

and continuing in this manner gives

$$\|f\|_{E_1, \dots, E_s}^{4^s} \leq |B|^{2^s - 1} \cdot \|f\|_{E'_1, \dots, E'_s}^{2^s},$$

from which the result promptly follows.

To prove the sixth property, we first set

$$E''_i = \sum_{j \in [d]} \beta_{ij} \cdot [\pm \varepsilon_{ij} H_{ij}]$$

for parameters $\varepsilon_{ij} > 0$ to be chosen later. As in the proof of the fifth property, we begin with the $s = 1$ case and drop the index i . Note that for every $\varepsilon > 0$ and 1-bounded function $g : \mathbb{Z} \rightarrow \mathbb{C}$, we have

$$\left| \mathbb{E}_{h \in [\pm H]} g(h) - \mathbb{E}_{h \in [\pm H]} \mathbb{E}_{h' \in [\pm \varepsilon H]} g(h + h') \right| \leq 2\varepsilon. \quad (5.20)$$

Hence we can approximate

$$\begin{aligned} \|f\|_E^2 &= \sum_x \mathbb{E}_{\substack{h_j, h'_j \in [\pm H_j] \\ j \in [d]}} f \left(x + \sum_{j \in [d]} \beta_j h'_j \right) \overline{f \left(x + \sum_{j \in [d]} \beta_j h_j \right)} \\ &= \sum_x \mathbb{E}_{\substack{h_j, h'_j \in [\pm H_j] \\ j \in [d]}} f \left(x + \sum_{j \in [d]} \beta_j h'_j \right) \mathbb{E}_{\substack{h''_j \in [\pm \varepsilon_j H_j] \\ j \in [d]}} \overline{f \left(x + \sum_{j \in [d]} \beta_j (h_j + h''_j) \right)} + O_d(\varepsilon |B|), \end{aligned}$$

where $\varepsilon := \max_j \varepsilon_j$. By the Cauchy–Schwarz inequality in all variables except h''_j , we get

$$\begin{aligned} \|f\|_E^4 &\leq |B| \cdot \sum_x \mathbb{E}_{\substack{h_j \in [\pm H_j], \\ h'_j, h''_j \in [\pm \varepsilon_j H_j] : j \in [d]}} f \left(x + \sum_{j \in [d]} \beta_j (h_j + h'_j) \right) \\ &\quad \overline{f \left(x + \sum_{j \in [d]} \beta_j (h_j + h''_j) \right)} + O_d(\varepsilon |B|^2). \end{aligned}$$

Shifting the variable $x \mapsto x - \sum_{j \in [d]} \beta_j h_j$ gives

$$\begin{aligned} \|f\|_E^4 &\leq |B| \cdot \sum_x \mathbb{E}_{\substack{h'_j, h''_j \in [\pm \varepsilon_j H_j] \\ j \in [d]}} f \left(x + \sum_{j \in [d]} \beta_j h'_j \right) \overline{f \left(x + \sum_{j \in [d]} \beta_j h''_j \right)} + O_d(\varepsilon |B|^2) \\ &= |B| \cdot \|f\|_{E''}^2 + O_d(\varepsilon |B|^2). \end{aligned}$$

Iterating this argument for higher-degree norms as in the proof of (5) gives

$$\|f\|_{E_1, \dots, E_s}^{4^s} \leq |B|^{2^s-1} \cdot \|f\|_{E_1'', \dots, E_s''}^{2^s} + O_{d,s}(\varepsilon |B|^{2^s}),$$

where $\varepsilon := \max_{i,j} \varepsilon_{ij}$, so

$$\|f\|_{E_1, \dots, E_s}^{2^s} \geq \delta |B| \implies \|f\|_{E_1'', \dots, E_s''}^{2^s} \geq (\delta^{2^s} - C_{d,s} \varepsilon) |B|$$

for some constant $C_{d,s} > 0$. If the κ_{ij} 's are sufficiently small, namely, if $\kappa := \max_{i,j} \kappa_{ij} \leq \delta^{2^s} / (2C_{d,s})$, then we apply the above bound directly with $\varepsilon_{ij} = \kappa_{ij}$. Otherwise, we first decrease the lengths of boxes, setting $\varepsilon_{ij} = \min(\delta^{2^s} / (2C_{d,s}), \kappa_{ij})$ for all i, j , and then use (4) to enlarge the boxes from E_i'' to E_i' , which gives

$$\|f\|_{E_1', \dots, E_s'}^{2^s} \geq \left(\prod_{i \in [s], j \in [d]} \frac{\varepsilon_{ij}^2}{\kappa_{ij}^2} \right) \|f\|_{E_1'', \dots, E_s''}^{2^s} \gg_{d,s} \delta^{2^s} \cdot (\delta^{2^s} / \kappa)^{2ds} |B|.$$

Combining these two cases yields the desired bound.

For the last property, we first observe that for positive integers H, q, u , the sets $[\pm H]$ and $\bigcup_{u=0}^{q-1} (u+q \cdot [\pm H/q])$ differ in at most $2q$ “boundary” elements. Hence, if $(x_i)_{i \in \mathbb{Z}}$ is a 1-bounded sequence of complex numbers, then

$$\begin{aligned} \left| \mathbb{E}_{h \in [\pm H]} x_h - \mathbb{E}_{u \in [0, q-1]} \mathbb{E}_{h \in [\pm H/q]} x_{qh+u} \right| \\ \leq \frac{2q}{2H+1} + \left| \frac{1}{2H+1} - \frac{1}{(2 \lfloor H/q \rfloor + 1)q} \right| (2 \lfloor H/q \rfloor + 1)q \leq \frac{3q+1}{2H+1}. \end{aligned}$$

For the proof proper of the last property, we start with the case $s = 1$. Approximating each box $\beta_{ij} \cdot [\pm H_{ij}]$ by the union of boxes $\beta_{ij} \cdot (u_j + q_j \cdot [\pm H_j/q_j])$ indexed by $u_j \in [0, q_j - 1]$ and applying the estimate above, we get

$$\begin{aligned} \|f\|_E^2 = \mathbb{E}_{\substack{u_j, u'_j \in [0, q_j-1]: \\ j \in [d]}} \sum_x \mathbb{E}_{\substack{h_j, h'_j \in [\pm H_j/q_j]: \\ j \in [d]}} f \left(x + \sum_{j \in [d]} \beta_j(q_j h_j + u_j) \right) \\ \overline{f \left(x + \sum_{j \in [d]} \beta_j(q_j h'_j + u'_j) \right)} + O_d(\varepsilon |B|) \end{aligned}$$

for $\varepsilon := \max_{j \in [d]} \frac{q_j}{H_j}$. We pigeonhole in u_j, u'_j and apply the Cauchy–Schwarz inequality to get

$$\|f\|_E^4 \leq \left\| f\left(\cdot + \sum_{j \in [d]} \beta_j u_j\right) \right\|_{E'}^2 \cdot \left\| f\left(\cdot + \sum_{j \in [d]} \beta_j u'_j\right) \right\|_{E'}^2 + O_d(\varepsilon |B|^2).$$

The translation-invariance of box norms then gives $\|f\|_{E'}^2 \geq (\delta - O_d(\sqrt{\varepsilon}))|B|$. The statement for general s follows inductively, using the inductive formula for box norms. $\square$

5.4 PET induction scheme

In this section, we begin the study of the counting operators

$$\sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{z \in [K]} f_0(\mathbf{x}) \cdot f_1(\mathbf{x} + \mathbf{P}_1(z)) \cdots f_\ell(\mathbf{x} + \mathbf{P}_\ell(z)) \quad (5.21)$$

associated with families $\mathcal{P} = (\mathbf{P}_1, \dots, \mathbf{P}_\ell)$ of essentially distinct polynomials in $\mathbb{Z}^D[z]$ (where “essentially distinct” means that all polynomials and their pairwise differences are nonconstant). For instance, in this notation the configuration

$$(x_1, x_2), (x_1 + z^2 + z, x_2), (x_1, x_2 + z^2 + z)$$

can be rewritten as $\mathbf{x}, \mathbf{x} + \mathbf{e}_1(z^2 + z), \mathbf{x} + \mathbf{e}_2(z^2 + z)$ and corresponds to the family $\mathcal{P} = \{\mathbf{e}_1(z^2 + z), \mathbf{e}_2(z^2 + z)\}$.

We present here the standard PET induction scheme which allows us to bound the counting operator (5.21) by an average of box norms. These box norms can be seen as “local” in the sense that the differencing parameters lie in intervals of length $\asymp N^{1/d}$, where $d = \max_j \deg \mathbf{P}_j$. This is the first step in showing that multidimensional polynomial progressions are controlled by a single “global” box norm. Combining arguments from [32, 111], we obtain a precise description of the directions of the box norms; in particular, we show that they come from the coefficients of the polynomials $\mathbf{P}_j$ and their differences. We start with a simple example in Section 5.4.1 that will motivate the argument in the general case. In Section 5.4.2, we set up the formalism that we need for this purpose and derive various properties of the polynomial families that appear in the PET induction scheme. We then assemble bits and pieces from Section 5.4.2 in order to prove Proposition 5.23, the main result of this section.

The material in this subsection is not novel, and various parts of it can be found in previous works on the subject (in [32, 81, 90, 111], *inter alia*).

5.4.1 Example of PET

Consider the counting operator

$$\Lambda(f_0, f_1, f_2) := \sum_{\mathbf{x} \in \mathbb{Z}^2} \mathbb{E}_{z \in [K]} f_0(\mathbf{x}) f_1(\mathbf{x} + \beta_{12} z^2 + \beta_{11} z) f_2(\mathbf{x} + \beta_{22} z^2 + \beta_{21} z),$$

and suppose that

$$\Lambda(f_0, f_1, f_2) \geq \delta N^2$$

for some 1-bounded functions $f_0, f_1, f_2 : \mathbb{Z}^2 \rightarrow \mathbb{C}$ supported on $[N]^2$. To avoid degenerate cases, we assume that $\beta_{12}, \beta_{12} - \beta_{22} \neq 0$, and that $(\beta_{22}, \beta_{21}) \neq (0, 0)$.

Our goal in this section is to work out the PET argument for this specific example before delving into it more abstractly. The goal of the PET argument is to pass from the counting operator for a polynomial progression to an average of box norms of one of the functions in the configuration (say, f_1). This is achieved in a finite number of steps, in each of which we apply the Cauchy–Schwarz inequality (to remove terms that do not depend on z), the van der Corput inequality (to double the variable z) and a change of variables (to bring the progression into a form where some functions are independent of z). That this procedure ends in finite time follows from the fact that at each step, we reduce the (appropriately measured) complexity, or *type*, of the polynomial progression.

For this particular example, we will need three applications of the aforementioned procedure to reduce the progression to an average of box norms of f_1 .

Step 1

We apply the Cauchy–Schwarz inequality in $\mathbf{x}$; since f_0 is independent of z and has L^2 -norm at most N , we can bound

$$\sum_{\mathbf{x}} \left| \mathbb{E}_{z \in [K]} f_1(\mathbf{x} + \beta_{12} z^2 + \beta_{11} z) f_2(\mathbf{x} + \beta_{22} z^2 + \beta_{21} z) \right|^2 \geq \delta^2 N^2.$$

For the $\Omega(\delta^2 N^2)$ values of $\mathbf{x} \in [N]^2$ for which the absolute value above is at least $\delta^2 N^2/2$, we then apply the van der Corput inequality (Lemma 5.14) with $H \ll \delta^8 K$; this has the effect of doubling z and gives

$$\begin{aligned} \sum_{\mathbf{x}, h_1} \mu_H(h_1) \mathbb{E}_{z \in [K]} f_1(\mathbf{x} + \beta_{12} z^2 + \beta_{11} z) \overline{f_1(\mathbf{x} + \beta_{12}(z + h_1)^2 + \beta_{11}(z + h_1))} \\ f_2(\mathbf{x} + \beta_{22} z^2 + \beta_{21} z) \overline{f_2(\mathbf{x} + \beta_{22}(z + h_1)^2 + \beta_{21}(z + h_1))} \gg \delta^2 N. \end{aligned}$$

We now want to bring the configuration into a form where one copy of f_2 has an argument independent of z ; then this copy of f_2 can be removed at the subsequent step of the argument. Shifting $\mathbf{x} \mapsto \mathbf{x} - \beta_{22}z^2 + \beta_{21}z$, we obtain

$$\sum_{\mathbf{x}, h_1} \mu_H(h_1) \mathbb{E}_{z \in [K]} f_1(\mathbf{x} + (\beta_{12} - \beta_{22})z^2 + (\beta_{11} - \beta_{21})z) \\ \overline{f_1(\mathbf{x} + (\beta_{12} - \beta_{22})z^2 + (2\beta_{12}h_1 + \beta_{11} - \beta_{21})z + \beta_{12}h_1^2 + \beta_{11}h_1)} \\ f_2(\mathbf{x}) \overline{f_2(\mathbf{x} + 2\beta_{22}h_1z + \beta_{22}h_1^2 + \beta_{21}h_1)} \gg \delta^2 N^2.$$

This new progression (or, rather, family of progressions indexed by h_1) may not look simpler than the original one, but it is: Whereas the original progression involved two quadratic¹⁰ polynomials with distinct leading coefficients β_{12}, β_{22} , the new progression involves two quadratic polynomials with the same leading coefficient $\beta_{12} - \beta_{22}$ (which is nonzero by assumption), as well as a linear polynomial. The decrease in the number of quadratic polynomials with distinct leading coefficients is what makes this progression “less complex” than the original one.

Step 2

Once again, we apply the Cauchy–Schwarz inequality in $\mathbf{x}, h_1$. This time, this has the effect of removing one of the two instances of f_2 , giving

$$\sum_{\mathbf{x}, h_1} \mu_H(h_1) \left| \mathbb{E}_{z \in [K]} f_1(\mathbf{x} + (\beta_{12} - \beta_{22})z^2 + (\beta_{11} - \beta_{21})z) \right. \\ \overline{f_1(\mathbf{x} + (\beta_{12} - \beta_{22})z^2 + (2\beta_{12}h_1 + \beta_{11} - \beta_{21})z + \beta_{12}h_1^2 + \beta_{11}h_1)} \\ \left. \overline{f_2(\mathbf{x} + 2\beta_{22}h_1z + \beta_{22}h_1^2 + \beta_{21}h_1)} \right|^2 \gg \delta^4 N^2.$$

The van der Corput inequality (Lemma 5.14) then gives

$$\sum_{\mathbf{x}, h_1, h_2} \mu_H(h_1, h_2) \mathbb{E}_{z \in [K]} f_1(\mathbf{x} + (\beta_{12} - \beta_{22})z^2 + (\beta_{11} - \beta_{21})z) \\ \overline{f_1(\mathbf{x} + (\beta_{12} - \beta_{22})(z + h_2)^2 + (\beta_{11} - \beta_{21})(z + h_2))} \\ \overline{f_1(\mathbf{x} + (\beta_{12} - \beta_{22})z^2 + (2\beta_{12}h_1 + \beta_{11} - \beta_{21})z + \beta_{12}h_1^2 + \beta_{11}h_1)} \\ f_1(\mathbf{x} + (\beta_{12} - \beta_{22})(z + h_2)^2 + (2\beta_{12}h_1 + \beta_{11} - \beta_{21})(z + h_2) + \beta_{12}h_1^2 + \beta_{11}h_1) \\ \overline{\Delta_{2\beta_{22}h_1h_2} f_2(\mathbf{x} + 2\beta_{22}h_1z + \beta_{22}h_1^2 + \beta_{21}h_1)} \gg \delta^4 N^2.$$

¹⁰in z (the degree of a polynomial in this section always refers to the degree in z)

We want to make the remaining instances of f_2 independent of z . To this end we shift¹¹ $\mathbf{x} \mapsto \mathbf{x} - 2\beta_{22}h_1z$, which gives

$$\begin{aligned} \sum_{\mathbf{x}, h_1, h_2} \mu_H(h_1, h_2) & \mathbb{E}_{z \in [K]} f_1(\mathbf{x} + (\beta_{12} - \beta_{22})z^2 + (-2\beta_{22}h_1 + \beta_{11} - \beta_{21})z) \\ & \overline{f_1(\mathbf{x} + (\beta_{12} - \beta_{22})(z + h_2)^2 + (-2\beta_{22}h_1 + \beta_{11} - \beta_{21})z + \mathbf{a}_1(h_1, h_2))} \\ & \overline{f_1(\mathbf{x} + (\beta_{12} - \beta_{22})z^2 + (2(\beta_{12} - \beta_{22})h_1 + \beta_{11} - \beta_{21})z + \mathbf{a}_2(h_1, h_2))} \\ & f_1(\mathbf{x} + (\beta_{12} - \beta_{22})(z + h_2)^2 + (2(\beta_{12} - \beta_{22})h_1 + \beta_{11} - \beta_{21})z + \mathbf{a}_3(h_1, h_2)) \\ & \overline{\Delta_{2\beta_{22}h_1h_2} f_2(\mathbf{x} + \mathbf{a}_4(h_1, h_2))} \gg \delta^4 N^2 \end{aligned}$$

for some polynomials $\mathbf{a}_j$ in h_1, h_2 , the exact form of which is irrelevant for us. Even though the configuration has increased in length (this is an unfortunate aspect of the PET argument), it has become less complex in the sense that all of the quadratic polynomials now have the same leading coefficient, and there is no linear polynomial left (as contrasted with the previous step, where we still had one polynomial linear in z). We are ready for the last step, in which we will remove f_2 altogether.

Step 3

Proceeding as before, we apply the Cauchy–Schwarz inequality in $\mathbf{x}, h_1, h_2$ to get

$$\begin{aligned} \sum_{\mathbf{x}, h_1, h_2} \mu_H(h_1, h_2) & \left| \mathbb{E}_{z \in [K]} f_1(\mathbf{x} + (\beta_{12} - \beta_{22})z^2 + (-2\beta_{22}h_1 + \beta_{11} - \beta_{21})z) \right. \\ & \overline{f_1(\mathbf{x} + (\beta_{12} - \beta_{22})(z + h_2)^2 + (-2\beta_{22}h_1 + \beta_{11} - \beta_{21})z + \mathbf{a}_1(h_1, h_2))} \\ & \overline{f_1(\mathbf{x} + (\beta_{12} - \beta_{22})z^2 + (2(\beta_{12} - \beta_{22})h_1 + \beta_{11} - \beta_{21})z + \mathbf{a}_2(h_1, h_2))} \\ & \left. f_1(\mathbf{x} + (\beta_{12} - \beta_{22})(z + h_2)^2 + (2(\beta_{12} - \beta_{22})h_1 + \beta_{11} - \beta_{21})z + \mathbf{a}_3(h_1, h_2)) \right|^2 \gg \delta^8 N^2. \end{aligned}$$

An application of Lemma 5.14 then gives

$$\begin{aligned} \sum_{\mathbf{x}, h_1, h_2, h_3} \mu_H(h_1, h_2, h_3) & \mathbb{E}_{z \in [K]} f_1(\mathbf{x} + (\beta_{12} - \beta_{22})z^2 + (-2\beta_{22}h_1 + \beta_{11} - \beta_{21})z) \\ & \overline{f_1(\mathbf{x} + (\beta_{12} - \beta_{22})(z + h_3)^2 + (-2\beta_{22}h_1 + \beta_{11} - \beta_{21})(z + h_3))} \\ & \overline{f_1(\mathbf{x} + (\beta_{12} - \beta_{22})(z + h_2)^2 + (-2\beta_{22}h_1 + \beta_{11} - \beta_{21})z + \mathbf{a}_1(h_1, h_2))} \\ & f_1(\mathbf{x} + (\beta_{12} - \beta_{22})(z + h_2 + h_3)^2 + (-2\beta_{22}h_1 + \beta_{11} - \beta_{21})(z + h_3) + \mathbf{a}_1(h_1, h_2)) \end{aligned}$$

¹¹If $\beta_{22} = \mathbf{0}$, then this term can be removed before applying the Cauchy–Schwarz inequality, and we can proceed directly to the change of variables performed in Step 3 below.

$$\begin{aligned}
& \overline{f_1(\mathbf{x} + (\beta_{12} - \beta_{22})z^2 + (2(\beta_{12} - \beta_{22})h_1 + \beta_{11} - \beta_{21})z + \mathbf{a}_2(h_1, h_2))} \\
& f_1(\mathbf{x} + (\beta_{12} - \beta_{22})(z + h_3)^2 + (2(\beta_{12} - \beta_{22})h_1 + \beta_{11} - \beta_{21})(z + h_3) + \mathbf{a}_2(h_1, h_2)) \\
& f_1(\mathbf{x} + (\beta_{12} - \beta_{22})(z + h_2)^2 + (2(\beta_{12} - \beta_{22})h_1 + \beta_{11} - \beta_{21})z + \mathbf{a}_3(h_1, h_2)) \\
& \overline{f_1(\mathbf{x} + (\beta_{12} - \beta_{22})(z + h_2 + h_3)^2 + (2(\beta_{12} - \beta_{22})h_1 + \beta_{11} - \beta_{21})(z + h_3) + \mathbf{a}_3(h_1, h_2))} \\
& \gg \delta^8 N^2.
\end{aligned}$$

Shifting

$$\mathbf{x} \mapsto \mathbf{x} - (\beta_{12} - \beta_{22})z^2 - (-2\beta_{22}h_1 + \beta_{11} - \beta_{21})z,$$

we end up with

$$\begin{aligned}
& \sum_{\mathbf{x}, h_1, h_2, h_3} \mu_H(h_1, h_2, h_3) \mathbb{E}_{z \in [K]} \overline{f_1(\mathbf{x}) f_1(\mathbf{x} + 2(\beta_{12} - \beta_{22})h_3z + \mathbf{a}_1(h_1, h_2, h_3))} \\
& \overline{f_1(\mathbf{x} + 2(\beta_{12} - \beta_{22})h_2z + \mathbf{a}_2(h_1, h_2, h_3))} f_1(\mathbf{x} + 2(\beta_{12} - \beta_{22})(h_2 + h_3)z + \mathbf{a}_3(h_1, h_2, h_3)) \\
& \overline{f_1(\mathbf{x} + 2\beta_{12}h_1z + \mathbf{a}_4(h_1, h_2, h_3))} f_1(\mathbf{x} + (2(\beta_{12} - \beta_{22})h_3 + 2\beta_{12}h_1)z + \mathbf{a}_5(h_1, h_2, h_3)) \\
& f_1(\mathbf{x} + (2(\beta_{12} - \beta_{22})h_2 + 2\beta_{12}h_1)z + \mathbf{a}_6(h_1, h_2, h_3)) \\
& \overline{f_1(\mathbf{x} + (2(\beta_{12} - \beta_{22})(h_2 + h_3) + 2\beta_{12}h_1)z + \mathbf{a}_7(h_1, h_2, h_3))} \gg \delta^8 N^2.
\end{aligned}$$

for some polynomials $\mathbf{a}_j$ in h_1, h_2, h_3 . The important point about this new expression is that it is an average, indexed by (h_1, h_2, h_3) , of counting operators for the 8-point linear (in z) configurations

$$\begin{aligned}
& \mathbf{x}, \mathbf{x} + 2(\beta_{12} - \beta_{22})h_3z, \mathbf{x} + 2(\beta_{12} - \beta_{22})h_2z, \mathbf{x} + 2(\beta_{12} - \beta_{22})(h_2 + h_3)z, \\
& \mathbf{x} + 2\beta_{12}h_1z, \mathbf{x} + (2(\beta_{12} - \beta_{22})h_3 + 2\beta_{12}h_1)z, \mathbf{x} + (2(\beta_{12} - \beta_{22})h_2 + 2\beta_{12}h_1)z, \\
& \mathbf{x} + (2(\beta_{12} - \beta_{22})(h_2 + h_3) + 2\beta_{12}h_1)z;
\end{aligned}$$

as such, each of these configurations can be controlled individually by a certain degree-7 box norm as in Lemma 5.19 below, using Gowers' original argument for controlling arithmetic progressions by Gowers norms [42, 43]. Gowers' argument, which we abstain from showing here, consists of 7 more applications of the Cauchy–Schwarz and van der Corput Inequalities, together with changes of variables, and its output is

the norm-control

$$\sum_{h_1, h_2, h_3} \mu_H(h_1, h_2, h_3) \|f_1\|_{2(\beta_{12}-\beta_{22})h_3 \cdot [\pm M], 2(\beta_{12}-\beta_{22})h_2 \cdot [\pm M], 2(\beta_{12}-\beta_{22})(h_2+h_3) \cdot [\pm M], 2\beta_{12}h_1 \cdot [\pm M], (2(\beta_{12}-\beta_{22})h_3+2\beta_{12}h_1) \cdot [\pm M], (2(\beta_{12}-\beta_{22})h_2+2\beta_{12}h_1) \cdot [\pm M], (2(\beta_{12}-\beta_{22})(h_2+h_3)+2\beta_{12}h_1) \cdot [\pm M]}}^{2^7} \gg \delta^{2^{10}} N^2$$

for any $M \ll \delta^{2^{10}} K$. The important points about this average of box norms are that the directions

$$\begin{aligned} &2(\beta_{12} - \beta_{22})h_3, 2(\beta_{12} - \beta_{22})h_2, 2(\beta_{12} - \beta_{22})(h_2 + h_3), \\ &2\beta_{12}h_1, 2(\beta_{12} - \beta_{22})h_3 + 2\beta_{12}h_1, 2(\beta_{12} - \beta_{22})h_2 + 2\beta_{12}h_1, \\ &2(\beta_{12} - \beta_{22})(h_2 + h_3) + 2\beta_{12}h_1 \end{aligned}$$

in the average of box norms above depend only on the leading coefficients of the polynomials $\mathbf{P}_1, \mathbf{P}_1 - \mathbf{P}_2$, and that they are linear polynomials in h_1, h_2, h_3 (for a general polynomial family $\mathcal{P}$, the directions will be multilinear polynomials in h_i of degree one smaller than the degree of $\mathcal{P}$). Understanding these directions for a general polynomial family is an important (even if tedious) task, on which we embark now.

5.4.2 Setting up the induction scheme

In order to carry out the PET argument in general, we first need to set up some notation that allows us to talk about general multidimensional polynomial progressions. Let $d, D, \ell, r, s \in \mathbb{N}$, and let $\underline{h} := (h_1, \dots, h_r)$ be a tuple of length r . All of the polynomial families that we shall consider will take the following form.

Definition 5.16 (Normal polynomial family). Let $\mathcal{Q} = (\mathbf{Q}_1, \dots, \mathbf{Q}_s)$ be a family of polynomials in $\mathbb{Z}^D[z, \underline{h}]$ of degree at most d . We call $\mathcal{Q}$ *normal* if the polynomials $\mathbf{Q}_1, \dots, \mathbf{Q}_s$ are nonzero and distinct, $\mathbf{Q}_1$ has highest degree (with respect to the variable z), and $\mathbf{Q}_1(0, \underline{h}), \dots, \mathbf{Q}_s(0, \underline{h})$ are the zero polynomial for all $\underline{h} \in \mathbb{Z}^r$.

Let $\mathcal{P} = (\mathbf{P}_1, \dots, \mathbf{P}_\ell)$ be a normal family of polynomials in $\mathbb{Z}^D[z]$ of degree at most d . Without loss of generality, we can assume that $\deg \mathbf{P}_1 = d$. We denote the coefficients of the polynomials $\mathbf{P}_j$ by $\mathbf{P}_j(z) = \sum_{i=1}^d \beta_{ji} z^i$, and we also set $\mathbf{P}_0 := \mathbf{0}$ and $\beta_{0i} := \mathbf{0}$. During the PET argument, we will subject $\mathcal{P}$ to an iterated application of the following operation.

Definition 5.17 (van der Corput operation). Let $\mathcal{Q} = (\mathbf{Q}_1, \dots, \mathbf{Q}_s)$ be a normal family of polynomials in $\mathbb{Z}^D[n, \underline{h}]$. Given $m \in [s]$, we define the new polynomial family

$$\partial_m \mathcal{Q} := (\sigma_{h_{r+1}} \mathbf{Q}_1 - \mathbf{Q}_m, \dots, \sigma_{h_{r+1}} \mathbf{Q}_s - \mathbf{Q}_m, \mathbf{Q}_1 - \mathbf{Q}_m, \dots, \mathbf{Q}_s - \mathbf{Q}_m)^*,$$

where $\sigma_{h_{r+1}} \mathbf{Q}(z, \underline{h}) := \mathbf{Q}(z + h_{r+1}, \underline{h}) - \mathbf{Q}(h_{r+1}, \underline{h})$ and the $*$ operation removes all zero polynomials and all but the first copy of each polynomial that appears multiple times.

We refer to the operation $\mathcal{Q} \mapsto \partial_m \mathcal{Q}$ as a *van der Corput operation* on $\mathcal{Q}$ since the new family $\partial_m \mathcal{Q}$ will arise from $\mathcal{Q}$ via an application of the van der Corput inequality. It is through this operation that we obtain a new polynomial family from $\mathcal{Q}$ at each step of the PET induction scheme.

The following lemma shows how we can use the Cauchy–Schwarz and van der Corput Inequalities, together with changes of variables, to pass from the family $\mathcal{Q}$ to the family $\partial_m \mathcal{Q}$ at each step of PET.

Lemma 5.18. *Let $d, D, m, \ell, r, s \in \mathbb{N}$, and let $\mathcal{Q} = \{\mathbf{Q}_1, \dots, \mathbf{Q}_s\}$ be a normal family of polynomials in $\mathbb{Z}^D[z, \underline{h}]$ of degree at most d . Suppose that $\deg \mathbf{Q}_1 > 1$. Let $\delta > 0$ and $H, K, N \in \mathbb{N}$ be parameters satisfying $H \leq \frac{\delta^2}{4}K$. Then for any 1-bounded functions $f_0, \dots, f_\ell : \mathbb{Z}^D \rightarrow \mathbb{C}$ supported on $[N]^D$ and any $\underline{h} \in \mathbb{Z}^s$, the lower bound*

$$\left| \sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{z \in [K]} f_0(\mathbf{x}) f_1(\mathbf{x} + \mathbf{Q}_1(z, \underline{h})) \cdots f_s(\mathbf{x} + \mathbf{Q}_s(z, \underline{h})) \right| \geq \delta N^D$$

implies that

$$\sum_{h_{r+1}} \mu_H(h_{r+1}) \sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{z \in [K]} \prod_{\mathbf{Q} \in \partial_m \mathcal{Q}} f_{\mathbf{Q}}(\mathbf{x} + \mathbf{Q}(z, \underline{h}, h_{r+1})) \geq \frac{\delta^2}{4} N^D$$

for some 1-bounded functions $f_{\mathbf{Q}} : \mathbb{Z} \rightarrow \mathbb{C}$ with $f_{\partial_{h_{r+1}} \mathbf{Q}_1 - \mathbf{Q}_m} = f_1$.

The property $f_{\partial_{h_{r+1}} \mathbf{Q}_1 - \mathbf{Q}_m} = f_1$ is crucial for letting us keep track of the function f_1 over the course of the PET procedure.

Proof. First, we make the important observation that all the polynomials

$$\sigma_{h_{r+1}} \mathbf{Q}_2 - \mathbf{Q}_m, \dots, \sigma_{h_{r+1}} \mathbf{Q}_s - \mathbf{Q}_m, \mathbf{Q}_1 - \mathbf{Q}_m, \dots, \mathbf{Q}_s - \mathbf{Q}_m \quad (5.22)$$

are distinct from $\sigma_{h_{r+1}} \mathbf{Q}_1 - \mathbf{Q}_m$ since all of the polynomials in $\mathcal{Q}$ are distinct from $\mathbf{Q}_1$. In making this deduction, we crucially use the assumption that $\mathbf{Q}_1$ is nonlinear in z , since otherwise $\mathbf{Q}_1 - \mathbf{Q}_m$ and $\sigma_{h_{r+1}} \mathbf{Q}_1 - \mathbf{Q}_m$ would be the same (note that the σ operation removes terms independent of z , so were $\mathbf{Q}_1$ linear, it would remove entirely the shift by h_{r+1}). Consequently, when we apply the $*$ operation, we do not delete any polynomial from (5.22) on account of its agreeing with $\sigma_{h_{r+1}} \mathbf{Q}_1 - \mathbf{Q}_m$. The conclusion of the lemma then follows from an application of the Cauchy–Schwarz and van der Corput Inequalities, together with the change of variables $\mathbf{x} \mapsto \mathbf{x} - \mathbf{Q}_m(z) - \mathbf{Q}_1(h_{r+1}, \underline{h})$, which ensures that $f_{\partial_{h_{r+1}} \mathbf{Q}_1 - \mathbf{Q}_m} = f_1$. $\square$

We deal with the linear case differently.

Lemma 5.19. *Let $\ell, D \in \mathbb{N}$, and let $\alpha_1, \dots, \alpha_\ell \in \mathbb{Z}^D$ be nonzero vectors with $\alpha_0 = \mathbf{0}$. Let $\delta > 0$ and $K, M, N \in \mathbb{N}$ be parameters with $M \ll_\ell \delta^{O_\ell(1)} K$. Then for any 1-bounded functions $f_0, \dots, f_\ell : \mathbb{Z}^D \rightarrow \mathbb{C}$ supported on $[N]^D$, the lower bound*

$$\left| \sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{z \in [K]} f_0(\mathbf{x}) \cdot \prod_{j \in [\ell]} f_j(\mathbf{x} + \alpha_j z) \right| \geq \delta N^D$$

implies that

$$\mathbb{E}_{\underline{m}, \underline{m}' \in [\pm M]^\ell} \sum_{\mathbf{x} \in \mathbb{Z}^D} \Delta'_{\alpha_j - \alpha_0, \dots, \alpha_j - \alpha_{j-1}, \alpha_j - \alpha_{j+1}, \dots, \alpha_j - \alpha_\ell; (\underline{m}, \underline{m}')} f_j(\mathbf{x}) \geq \frac{\delta^{2^\ell}}{4^{2^\ell - 1}} N^D$$

for each $j \in [0, \ell]$.

Proof. Since the hypothesis of the lemma is invariant under permutations of the functions $f_1, \dots, f_\ell$, it suffices to consider the case $j = 1$. We proceed by induction on ℓ . Applying the Cauchy–Schwarz and van der Corput Inequalities, we get

$$\mathbb{E}_{m_1, m'_1 \in [\pm M]} \sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{z \in [K]} \prod_{j \in [\ell]} \Delta'_{(\alpha_j m_1, \alpha_j m'_1)} f_j(\mathbf{x} + \alpha_j z) \geq \frac{\delta^2}{4} N^D.$$

After translating $\mathbf{x} \mapsto \mathbf{x} - \alpha_\ell z$, we obtain the result immediately in the case $\ell = 1$, and for $\ell > 1$ we invoke the case $\ell - 1$ after observing that

$$(\alpha_1 - \alpha_\ell) - (\alpha_j - \alpha_\ell) = \alpha_1 - \alpha_j.$$

□

Our next goal is to describe the properties of $\mathcal{P}$ that are preserved by the van der Corput operation. Let $\mathcal{Q} = (\mathbf{Q}_1, \dots, \mathbf{Q}_s)$ be a normal family of polynomials in $\mathbb{Z}^D[z, \underline{h}]$, and write

$$\mathbf{Q}_j(z, \underline{h}) = \sum_{i=1}^d \gamma_{ji}(\underline{h}) z^i \tag{5.23}$$

for some $\gamma_{ji} \in \mathbb{Z}[\underline{h}]$. In other words, the γ_{ji} ’s are the coefficients of $\mathbf{Q}_j$ as a polynomial in z , and they are themselves polynomials in $\underline{h}$. For $j, j' \in [0, s]$, also define

$$d_{jj'} := \max\{i \in [d] : \gamma_{ji}(\underline{h}) \neq \gamma_{j'i}(\underline{h})\}$$

to be the degree of $\mathbf{Q}_j - \mathbf{Q}_{j'}$ as a polynomial in z , where we set $\mathbf{Q}_0 := \mathbf{0}$ and $\gamma_{0i} := \mathbf{0}$ for all $i \in \mathbb{N}_0$. We want to show that if $\mathcal{Q}$ is a family obtained from $\mathcal{P}$ at some stage of the PET induction scheme, then the polynomials γ_{ji} preserve several useful properties.

Definition 5.20 (Descendent family of polynomials). A normal family $\mathcal{Q} = (\mathbf{Q}_1, \dots, \mathbf{Q}_s)$ of polynomials in $\mathbb{Z}^D[z, \underline{h}]$ *descends from* $\mathcal{P}$ if for each $j \in [s]$, the polynomials γ_{ji} in (5.23) can be expressed as¹²

$$\gamma_{ji}(\underline{h}) = \sum_{\substack{\underline{u} \in \mathbb{N}_0^r, \\ |\underline{u}| \leq d-i}} \binom{u_1 + \dots + u_r + i}{u_1, \dots, u_r} (\beta_{w_{j\underline{u}}(|\underline{u}|+i)} - \beta_{w_{\underline{u}}(|\underline{u}|+i)}) \underline{h}^{\underline{u}} \quad (5.24)$$

for some indices $w_{\underline{u}}, w_{j\underline{u}} \in \{0, \dots, \ell\}$ satisfying the following properties:

1. (Coefficients of $\mathbf{Q}_1$) $w_{1\underline{u}} = 1$ for all $\underline{u}$.
2. (Multilinearity) For each j , the polynomial $\gamma_{1d_{1j}} - \gamma_{jd_{1j}}$ (the leading coefficient of $\mathbf{Q}_1 - \mathbf{Q}_j$ as a polynomial in z) is multilinear in $\underline{h}$.
3. (Leading coefficients) For each j and $\underline{u}$, the vector $\beta_{1(|\underline{u}|+d_{1j})} - \beta_{w_{j\underline{u}}(|\underline{u}|+d_{1j})}$ is nonzero only if it is the leading coefficient of $\mathbf{P}_1 - \mathbf{P}_{w_{j\underline{u}}}$.

The rationale behind recording the particular properties listed in Definition 5.20 is as follows. The normality of the family ensures that the progression does not become degenerate¹³ and that the polynomial with index 1 continues to have the highest degree (which is essential for the function originally evaluated at $\mathbf{x} + \mathbf{P}_1(z)$ not to disappear at an intermediate step of PET). The general form of the polynomials γ_{ji} ensures that the coefficients of the $\mathbf{Q}_j$'s and their pairwise differences are bounded multiples of the coefficients of the $\mathbf{P}_j$'s and their differences. The property (1) guarantees that we do not lose information about the coefficients of $\mathbf{P}_1$ while carrying out PET. The multilinearity of the polynomials $\gamma_{1d_{1j}} - \gamma_{jd_{1j}}$ given in (2) will be crucial in our concatenation argument since, as we show in Section 5.7, bounding the number of solutions to systems of multilinear equations is *much* simpler than obtaining bounds for general systems of polynomial equations. Lastly, property (3) ensures that the coefficients of $\gamma_{1d_{1j}} - \gamma_{jd_{1j}}$ come from only the *leading* coefficients of $\mathbf{P}_1, \mathbf{P}_1 - \mathbf{P}_2, \dots, \mathbf{P}_1 - \mathbf{P}_\ell$, and so the end product of PET will not depend on the lower-degree coefficients.

We note several consequences of Definition 5.20:

¹²The multinomial coefficient

$$\binom{u_1 + \dots + u_r + i}{u_1, \dots, u_r} = \frac{(u_1 + \dots + u_r + i)!}{u_1! \dots u_r! \cdot i!}$$

is the coefficient of $\underline{h}^{\underline{u}} z^i = h_1^{u_1} \dots h_r^{u_r} z^i$ in $(z + h_1 + \dots + h_r)^{i+u_1+\dots+u_r}$.

¹³A degenerate family would be one in which the difference of some two polynomials is independent of z ; such a family is “bad” because successive applications of the Cauchy–Schwarz and van der Corput Inequalities would not allow us to separate the functions corresponding to these polynomials.

1. If $\mathcal{Q}$ descends from $\mathcal{P}$, then the formulas (5.23) and (5.24) jointly ensure that the polynomials in $\mathcal{Q}$ cannot have their total degree in $z, h_1, \dots, h_r$ greater than d .
2. $\mathcal{P}$ descends from itself: Since $r = 0$, the only choice of $\underline{u}$ is $\underline{0}$, and one can set $w_{j\underline{0}} = j$ and $w_{\underline{0}} = 0$ for each $j \in [\ell]$.
3. The multilinearity in (2) tells us that only $\underline{u}$'s in $\{0, 1\}^r$ can contribute to $\gamma_{1d_{1j}} - \gamma_{jd_{1j}}$, so the multinomial coefficients simplify and we have

$$\gamma_{1d_{1j}}(\underline{h}) - \gamma_{jd_{1j}}(\underline{h}) = \sum_{\substack{\underline{u} \in \{0,1\}^r, \\ |\underline{u}| \leq d-d_{1j}}} \frac{(|\underline{u}| + d_{1j})!}{d_{1j}!} \cdot (\beta_{1(|\underline{u}|+d_{1j})} - \beta_{w_{j\underline{u}}(|\underline{u}|+d_{1j})}) \underline{h}^{\underline{u}}.$$

4. As a special case of (3), if the leading coefficient of $\mathbf{P}_1$ differs from the leading coefficients of $\mathbf{P}_2, \dots, \mathbf{P}_\ell$, then for each j , the polynomial $\gamma_{1d_{1j}} - \gamma_{jd_{1j}}$ is a multilinear, homogeneous polynomial of degree $d - d_{1j}$, and the expression for the polynomials $\gamma_{1d_{1j}} - \gamma_{jd_{1j}}$ further simplifies to

$$\gamma_{1d_{1j}}(\underline{h}) - \gamma_{jd_{1j}}(\underline{h}) = \sum_{\substack{\underline{u} \in \{0,1\}^r, \\ |\underline{u}| = d-d_{1j}}} \frac{d!}{d_{1j}!} \cdot (\beta_{1d} - \beta_{w_{j\underline{u}}d}) \underline{h}^{\underline{u}}.$$

5. If the leading coefficient of $\mathbf{P}_1$ is the same as the leading coefficient of some $\mathbf{P}_j$ for $2 \leq j \leq \ell$, then the polynomial $\gamma_{1d_{1j}} - \gamma_{jd_{1j}}$ ceases to be homogeneous and of degree $d - d_{1j}$; for instance, if $\mathcal{P} = \{z^2 + z, z^2\}$, then $\partial_2 \mathcal{P} = \{(2h+1)z, 2hz, z\}$ descends from $\mathcal{P}$, but $2h+1$ fails to be homogeneous and the polynomial 1 has degree $0 < 1 = d - d_{12}$.

The descendent family obtained during the van der Corput operation admits more properties: For instance, the coefficients $w_{j\underline{u}}, w_{\underline{u}}$ can be shown to satisfy $w_{j\underline{u}} = w_{j\underline{u}'}$ and $w_{\underline{u}} = w_{\underline{u}'}$ whenever $\text{supp}(\underline{u}) = \text{supp}(\underline{u}')$ (here, $\text{supp}(\underline{u}) := \{i \in [r] : u_i > 0\}$; this property has been observed in [34, Definition 5.5 (P3)] and used in this paper in the proof of [34, Proposition 5.7]). However, the use of this observation has been superseded in our work by the use of property (2) from Definition 5.20; this property previously appeared implicitly in [81, 111].

The following lemma ascertains that the van der Corput operation preserves the abovementioned properties inherited by $\mathcal{Q}$ from $\mathcal{P}$ by virtue of descending from $\mathcal{P}$.

Proposition 5.21. *Let $d, D, \ell, r, s \in \mathbb{N}$, and let $\mathcal{P} = (\mathbf{P}_1, \dots, \mathbf{P}_\ell)$ be a normal family of polynomials in $\mathbb{Z}^D[z]$ with coefficients given by $\mathbf{P}_j(z) = \sum_{i=1}^d \beta_{ji} z^i$. Suppose that the family $\mathcal{Q} = (\mathbf{Q}_1, \dots, \mathbf{Q}_s)$ of polynomials in $\mathbb{Z}^D[z, \underline{h}]$ descends from $\mathcal{P}$. Then the family $\partial_m \mathcal{Q}$ also descends from $\mathcal{P}$ for each $m \in [s]$ for which $\sigma_{h_{r+1}} \mathbf{Q}_1 - \mathbf{Q}_m$ has the highest degree (as a polynomial in z) among the elements of $\partial_m \mathcal{Q}$.*

Proof. We first note that the family $\partial_m \mathcal{Q}$ is normal: The application of the $*$ operation ensures that the polynomials of the family are nonzero and distinct, and the assumption in the proposition ensures that the first polynomial in the family has the highest degree (in z).

We now check that $\partial_m \mathcal{Q}$ satisfies the three properties of a descendent family. Since $\mathcal{Q}$ is a descendent family, the polynomials $\mathbf{Q}_j$ satisfy the formulas (5.23) and (5.24) for some indices $w_{j\underline{u}}, w_{\underline{u}}$ with $w_{1\underline{u}} = 1$. Using the binomial theorem, we can expand

$$\begin{aligned} \mathbf{Q}_j(z + h_{r+1}) &= \sum_{i=0}^d \sum_{\substack{\underline{u} \in \mathbb{N}_0^r, \\ |\underline{u}| \leq d-i}} \binom{u_1 + \dots + u_r + i}{u_1, \dots, u_r} (\beta_{w_{j\underline{u}}(|\underline{u}|+i)} - \beta_{w_{\underline{u}}(|\underline{u}|+i)}) \underline{h}^{\underline{u}} (z + h_{r+1})^i \\ &= \sum_{i=0}^d \sum_{\substack{\underline{u} \in \mathbb{N}_0^r, \\ |\underline{u}| \leq d-i}} \binom{u_1 + \dots + u_r + i}{u_1, \dots, u_r} (\beta_{w_{j\underline{u}}(|\underline{u}|+i)} - \beta_{w_{\underline{u}}(|\underline{u}|+i)}) \underline{h}^{\underline{u}} \sum_{l=0}^i \binom{i}{l} h_{r+1}^{i-l} z^l. \end{aligned}$$

Changing the order of summation of i and l , setting $u_{r+1} := i - l$, and using the identity

$$\binom{u_1 + \dots + u_{r+1} + l}{u_1, \dots, u_{r+1}} = \binom{u_1 + \dots + u_r + u_{r+1} + l}{u_1, \dots, u_r} \binom{u_{r+1} + l}{l},$$

we obtain

$$\begin{aligned} \mathbf{Q}_j(z + h_{r+1}) &= \sum_{l=0}^d z^l \sum_{\substack{(\underline{u}, u_{r+1}) \in \mathbb{N}_0^{r+1}, \\ |\underline{u}| + u_{r+1} \leq d-l}} \binom{u_1 + \dots + u_{r+1} + l}{u_1, \dots, u_{r+1}} \\ &\quad (\beta_{w_{j\underline{u}}(|\underline{u}|+u_{r+1}+l)} - \beta_{w_{\underline{u}}(|\underline{u}|+u_{r+1}+l)}) \underline{h}^{\underline{u}} h_{r+1}^{u_{r+1}}. \end{aligned}$$

Hence

$$\begin{aligned} \mathbf{Q}_j(z + h_{r+1}) - \mathbf{Q}_m(z) &= \sum_{l=0}^d z^l \sum_{\substack{(\underline{u}, u_{r+1}) \in \mathbb{N}_0^{r+1}, \\ |\underline{u}| + u_{r+1} \leq d-l}} \binom{u_1 + \dots + u_{r+1} + l}{u_1, \dots, u_{r+1}} \\ &\quad (\beta_{w_{j(\underline{u}, u_{r+1})}(|\underline{u}|+u_{r+1}+l)} - \beta_{w_{(\underline{u}, u_{r+1})}(|\underline{u}|+u_{r+1}+l)}) \underline{h}^{\underline{u}} h_{r+1}^{u_{r+1}}, \quad (5.25) \end{aligned}$$

where we set $w_{j(\underline{u}, u_{r+1})} := w_{j\underline{u}}$ and

$$w_{(\underline{u}, u_{r+1})} := \begin{cases} w_{\underline{u}}, & u_{r+1} > 0, \\ w_{m\underline{u}}, & u_{r+1} = 0. \end{cases}$$

In particular, we have $w_{1(\underline{u}, u_{r+1})} = w_{1\underline{u}} = 1$ by the induction hypothesis, and so $\partial_m \mathcal{Q}$ satisfies property (1).

To show properties (2) and (3), we need to compare $\sigma_{h_{r+1}} \mathbf{Q}_1(z, \underline{h}) - \mathbf{Q}_m(z, \underline{h})$ with other polynomials in $\partial_m \mathcal{Q}$ and examine the leading coefficients of these differences (as polynomials in z). There are three cases to investigate. First, we check (2) and (3) for differences of the form

$$\begin{aligned} & (\sigma_{h_{r+1}} \mathbf{Q}_1(z, \underline{h}) - \mathbf{Q}_m(z, \underline{h})) - (\sigma_{h_{r+1}} \mathbf{Q}_j(z, \underline{h}) - \mathbf{Q}_m(z, \underline{h})) \\ &= \sum_{i=1}^{d_{1j}} (\gamma_{1d_{1j}}(\underline{h}) - \gamma_{jd_{1j}}(\underline{h}))(z + h_{r+1})^i - (\mathbf{Q}_1(h_{r+1}, \underline{h}) - \mathbf{Q}_j(h_{r+1}, \underline{h})), \end{aligned} \quad (5.26)$$

where the terms $\mathbf{Q}_m(z, \underline{h})$ cancel out. The leading coefficient of this expression (interpreted as a polynomial in z) is $\gamma_{1d_{1j}}(\underline{h}) - \gamma_{jd_{1j}}(\underline{h})$, the same as the leading coefficient of the polynomial $\mathbf{Q}_1 - \mathbf{Q}_m$, and this is multilinear by assumption. For (3), we need to show that the nonzero vectors $\beta_{1(|\underline{u}|+u_{r+1}+d_{1j})} - \beta_{w_j(\underline{u}, u_{r+1})}(|\underline{u}|+u_{r+1}+d_{1j})$ are leading coefficients of the polynomials $\mathbf{P}_1 - \mathbf{P}_{w_j(\underline{u}, u_{r+1})}$. If $u_{r+1} = 0$, then this follows immediately from the formula

$$\beta_{1(|\underline{u}|+u_{r+1}+d_{1j})} - \beta_{w_j(\underline{u}, u_{r+1})}(|\underline{u}|+u_{r+1}+d_{1j}) = \beta_{1(|\underline{u}|+d_{1j})} - \beta_{w_j \underline{u}}(|\underline{u}|+d_{1j})$$

and the assumption on $\mathcal{Q}$, while if $u_{r+1} = 1$, then

$$\beta_{1(|\underline{u}|+u_{r+1}+d_{1j})} - \beta_{w_j(\underline{u}, u_{r+1})}(|\underline{u}|+u_{r+1}+d_{1j}) = \beta_{1(|\underline{u}|+d_{1j}+1)} - \beta_{w_j \underline{u}}(|\underline{u}|+d_{1j}+1) = \mathbf{0}$$

because the leading coefficient $\gamma_{1d_{1j}}(\underline{h}) - \gamma_{jd_{1j}}(\underline{h})$ is independent of h_{r+1} .

Second, we let d_1 denote the degree of $\mathbf{Q}_1$ in z and expand

$$\begin{aligned} & (\sigma_{h_{r+1}} \mathbf{Q}_1(z, \underline{h}) - \mathbf{Q}_m(z, \underline{h})) - (\mathbf{Q}_1(z, \underline{h}) - \mathbf{Q}_m(z, \underline{h})) \\ &= \sum_{i=1}^{d_1} \gamma_{1i}(\underline{h})((z + h_{r+1})^i - z^i) - \mathbf{Q}_1(h_{r+1}, \underline{h}) \\ &= \sum_{i=1}^{d_1} \sum_{l=1}^{i-1} \gamma_{1i}(\underline{h}) \binom{i}{l} h_{r+1}^{i-l} z^l \\ &= \sum_{l=1}^{d_1-1} z^l \sum_{i=l+1}^{d_1} \gamma_{1i}(\underline{h}) \binom{i}{l} h_{r+1}^{i-l}; \end{aligned} \quad (5.27)$$

the leading coefficient here is $d_1 h_{r+1} \gamma_{1d_1}(\underline{h})$, corresponding to $l = d_1 - 1, i = d_1$. The polynomial $\gamma_{1d_1}(\underline{h})$ is the leading coefficient of $\mathbf{Q}_1$, so it is multilinear by assumption, and hence $d_1 h_{r+1} \gamma_{1d_1}(\underline{h})$ is multilinear as well. Thus (2) holds. To show property (3), we first observe that the coefficient of $\underline{h}^{\underline{u}}$ in $\gamma_{1d_1}(\underline{h})$ is $\frac{(|\underline{u}|+d_1)!}{d_1!} (\beta_{1(|\underline{u}|+d_1)} - \beta_{w_{\underline{u}}}(|\underline{u}|+d_1))$ (where we use implicitly property (1)). Therefore, the coefficient of $\underline{h}^{\underline{u}} h_{r+1}$

in $d_1 h_{r+1} \gamma_{1d_1}(\underline{h})$ is

$$\begin{aligned} & \frac{(|\underline{u}| + d_1)!}{(d_1 - 1)!} (\beta_{1(|\underline{u}|+d_1)} - \beta_{w_{\underline{u}}(|\underline{u}|+d_1)}) \\ &= \frac{(|(\underline{u}, 1)| + (d_1 - 1))!}{(d_1 - 1)!} (\beta_{1(|(\underline{u}, 1)|+(d_1-1))} - \beta_{w_{(\underline{u}, 1)}(|(\underline{u}, 1)|+(d_1-1))}). \end{aligned}$$

By the induction hypothesis, if $\beta_{1(|\underline{u}|+d_1)} - \beta_{w_{\underline{u}}(|\underline{u}|+d_1)}$ is nonzero, then it is the leading coefficient of $\mathbf{P}_1 - \mathbf{P}_{w_{\underline{u}}}$, and hence

$$\beta_{1(|(\underline{u}, 1)|+(d_1-1))} - \beta_{w_{(\underline{u}, 1)}(|(\underline{u}, 1)|+(d_1-1))} = \beta_{1(|\underline{u}|+d_1)} - \beta_{w_{\underline{u}}(|\underline{u}|+d_1)}$$

is also nonzero and the leading coefficient of $\mathbf{P}_1 - \mathbf{P}_{w_{(\underline{u}, 1)}} = \mathbf{P}_1 - \mathbf{P}_{w_{\underline{u}}}$.

It remains to consider

$$(\sigma_{h_{r+1}} \mathbf{Q}_1(z, \underline{h}) - \mathbf{Q}_m(z, \underline{h})) - (\mathbf{Q}_j(z, \underline{h}) - \mathbf{Q}_m(z, \underline{h}))$$

for $j \neq 1$; we split this case into two subcases. If $\gamma_{1d_1}(\underline{h}) \neq \gamma_{jd_1}(\underline{h})$, then the leading coefficient equals $\gamma_{1d_1}(\underline{h}) - \gamma_{jd_1}(\underline{h}) = \gamma_{1d_1}(\underline{h}) - \gamma_{jd_1}(\underline{h})$. This is the leading coefficient of $\mathbf{Q}_1 - \mathbf{Q}_j$, so it is multilinear by assumption, and property (3) holds for the same reason as it held for (5.26).

Otherwise, $\gamma_{1d_1}(\underline{h}) = \gamma_{jd_1}(\underline{h})$, in which case the leading coefficient equals

$$d_1 h_{r+1} \gamma_{1d_1}(\underline{h}) + \gamma_{1(d_1-1)}(\underline{h}) - \gamma_{j(d_1-1)}(\underline{h}), \quad (5.28)$$

where

$$\gamma_{1(d_1-1)}(\underline{h}) - \gamma_{j(d_1-1)}(\underline{h}) = \begin{cases} \gamma_{1d_1}(\underline{h}) - \gamma_{jd_1}(\underline{h}), & \text{if } d_{1j} = d_1 - 1, \\ \mathbf{0}, & \text{if } d_{1j} < d_1 - 1. \end{cases}$$

The term $d_1 h_{r+1} \gamma_{1d_1}(\underline{h})$ is multilinear by the same argument as in the previous case. The term $\gamma_{1(d_1-1)}(\underline{h}) - \gamma_{j(d_1-1)}(\underline{h})$ is trivially multilinear when $d_{1j} < d_1 - 1$, and its multilinearity follows from the assumption on $\mathcal{Q}$ when $d_{1j} = d_1 - 1$ (since then it is the leading coefficient of $\mathbf{Q}_1 - \mathbf{Q}_j$). Hence (5.28) is multilinear as well.

It remains only to prove property (3) in this very last case, namely, to show that if the coefficient of the monomial $\underline{h}^{\underline{u}} h_{r+1}^{u_{r+1}}$ in (5.28) is nonzero, then it is an appropriate scalar multiple of the leading coefficient of the polynomial $\mathbf{P}_1 - \mathbf{P}_{w_{j(\underline{u}, u_{r+1})}} = \mathbf{P}_1 - \mathbf{P}_{w_{j\underline{u}}}$. If $u_{r+1} = 1$, then the coefficient of $\underline{h}^{\underline{u}} h_{r+1}^{u_{r+1}}$ in (5.28) is the same as in $d_1 h_{r+1} \gamma_{1d_1}(\underline{h})$, and the claim follows by the same reasoning as in the case of (5.27). Otherwise,

$u_{r+1} = 0$, and the coefficient $\underline{h}^u h_{r+1}^{u_{r+1}}$ in (5.28) comes fully from $\gamma_{1d_{1j}}(\underline{h}) - \gamma_{jd_{1j}}(\underline{h})$. The result then follows by the same argument as one used to prove property (3) for (5.26). $\square$

The next result ensures that we reach a family of linear polynomials after a bounded number of van der Corput operations; this fact is the foundation of all PET induction schemes.

Proposition 5.22. *Let $d, D, r, s \in \mathbb{N}$, let $\mathcal{P}$ be a normal family of polynomials in $\mathbb{Z}^D[z]$ of degree at most d , and let $\mathcal{Q}$ be a normal family of polynomials in $\mathbb{Z}^D[z, \underline{h}]$ of length s that descends from $\mathcal{P}$. There exist a nonnegative integer $r' = O_{d,s}(1)$ and positive integers $m_1, \dots, m_{r'}$ such that the family $\mathcal{Q}_i := \partial_{m_i} \cdots \partial_{m_1} \mathcal{Q}$ descends from $\mathcal{P}$ for each $i \in [r']$ and, moreover, the elements of $\mathcal{Q}_{r'}$ are linear in z .*

We remark that r' is independent of the dimension D of the underlying space.

Proof. Let $\mathcal{Q} := (\mathbf{Q}_1, \dots, \mathbf{Q}_s)$. As usual, we expand each $\mathbf{Q}_j$ as in (5.23) and say that its leading coefficient is $\gamma_{ji} \in \mathbb{Z}^D[\underline{h}]$ for the largest $i \in [d]$ with $\gamma_{ji} \neq \mathbf{0}$. Let d' be the maximum degree of $\mathbf{Q}_j, \dots, \mathbf{Q}_s$ as a polynomial in z ; since $\mathcal{Q}$ descends from $\mathcal{P}$ and is normal, necessarily $1 \leq d' \leq d$. For $l \in [d']$, let w_l denote the number of distinct leading coefficients of the degree- l polynomials in $\mathcal{Q}$. We then define $w = (w_1, \dots, w_{d'})$ to be the *type* of $\mathcal{Q}$, and we order the types colexicographically.¹⁴

Since the family $\mathcal{Q}$ is normal and has positive length, we have $w > (0, \dots, 0)$; let l be the smallest index for which $w_l \neq 0$. We define a parameter m_1 as follows. If $l < d'$, then we choose an arbitrary $m_1 \in [s]$ with $\deg \mathbf{Q}_{m_1} = l$. If $l = d'$ (i.e., all of the polynomials in $\mathcal{Q}$ have degree d' as polynomials in z) and $w_{d'} > 1$, then we choose an arbitrary $m_1 \in [s]$ such that $\mathbf{Q}_{m_1}$ and $\mathbf{Q}_1$ have different leading coefficients. Finally, if $l = d'$ and $w_{d'} = 1$ (i.e., all of the polynomials have the same leading term), then we choose $m_1 = 1$. A standard calculation shows that in each case the new family $\mathcal{Q}_1 = \partial_{m_1} \mathcal{Q}$ has type

$$w' = (w'_1, \dots, w'_{l-1}, w_l - 1, w_{l+1}, \dots, w_{d'}) < w,$$

for some $0 \leq w'_1, \dots, w'_{l-1} \leq 2s$, and the polynomial $\partial_h \mathbf{Q}_1 - \mathbf{Q}_{m_1}$ has the maximum degree among the polynomials in $\mathcal{Q}_1$. It follows from Proposition 5.21 that $\mathcal{Q}_1$ descends from $\mathcal{P}$.

We now iterate this algorithm. Since each decreasing sequence of types $w^0 > w^1 > w^2 > \dots$ with $w^0 = w$ eventually reaches a type $w^{r'} = (w_1^{r'}, 0, \dots, 0)$ for some $r' = O_{d',s}(1) = O_{d,s}(1)$ and $w_1^{r'} = O_{r',s}(1) = O_{d,s}(1)$, the algorithm terminates after $O_{d,s}(1)$ steps, and we obtain the conclusion of the proposition. $\square$

We remark that in the existing literature, the type of a polynomial family is often expressed as a matrix rather than a (row) vector, with different coordinates in $\mathbb{Z}^D$ giving rise to different rows in the matrix (see

¹⁴That is, we have $w > w'$ if $w_l > w'_l$ for the largest l with $w_l \neq w'_l$.

e.g. [9, 25, 34]). This bookkeeping is excessive, however; as the above proof illustrates, the type-reduction proceeds in the same way regardless of D , and the number of steps does not depend on D . The argument is even more robust: The number of steps would in fact be the same if we considered polynomial families over any integral domains with characteristic either equal to 0 or sufficiently large in terms of the degrees of the polynomials and the order of their coefficients.¹⁵

In the example discussed in Section 5.4.1, the original family has type $(0, 2)$ (or $(1, 1)$ if $\beta_{22} = 0$; we ignore this case for now). The new families obtained in steps 1, 2, and 3 have types $(1, 1)$, $(0, 1)$ and $(7, 0)$ (respectively).

5.4.3 Bounds by averages of box seminorms

Having presented the PET induction scheme in detail, we are ready now to show that the counting operator (5.21) is controlled by an average of box norms whose directions can be described in terms of the coefficients of the polynomials appearing in the multidimensional polynomial progression and the coefficients of differences of these polynomials. We no longer assume that $\mathbf{P}_1$ has the highest degree among $\mathbf{P}_1, \dots, \mathbf{P}_\ell$.

Proposition 5.23 (PET bound for polynomials). *Let $d, \ell \in \mathbb{N}$. There exist positive integers $r, s = O_{d, \ell}(1)$ such that the following holds. Let $\delta > 0$ and $D, H, K, M, N \in \mathbb{N}$ satisfy $H, M \ll_{d, D, \ell} \delta^{O_{d, \ell}(1)} K$. Let $\mathbf{P}_1, \dots, \mathbf{P}_\ell \in \mathbb{Z}^D[z]$ be essentially distinct polynomials with degrees at most d and coefficients $\mathbf{P}_j(z) = \sum_{i=0}^d \beta_{ji} z^i$. Then for any 1-bounded functions $f_0, \dots, f_\ell : \mathbb{Z}^D \rightarrow \mathbb{C}$, the lower bound*

$$\left| \sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{z \in [K]} f_0(\mathbf{x}) \cdot f_1(\mathbf{x} + \mathbf{P}_1(z)) \cdots f_\ell(\mathbf{x} + \mathbf{P}_\ell(z)) \right| \geq \delta N^D$$

implies that

$$\mathbb{E}_{\underline{h} \in [\pm H]^r} \|f_1\|_{\mathbf{C}_1(\underline{h}) \cdot [\pm M], \dots, \mathbf{C}_s(\underline{h}) \cdot [\pm M]}^{2^s} \gg_{d, D, \ell} \delta^{O_{d, \ell}(1)} N^D,$$

where $\mathbf{C}_1, \dots, \mathbf{C}_s \in \mathbb{Z}^D[\underline{h}]$ are nonzero multilinear polynomials depending only on $\mathbf{P}_1, \dots, \mathbf{P}_\ell$, and they take the form

$$\mathbf{C}_j(\underline{h}) = \sum_{\substack{\underline{u} \in \{0, 1\}^r \\ |\underline{u}| \leq d-1}} (|\underline{u}| + 1)! \cdot (\beta_{1(|\underline{u}|+1)} - \beta_{w_{j\underline{u}}(|\underline{u}|+1)}) \underline{h}^{\underline{u}} \quad (5.29)$$

¹⁵In the small-characteristic case, there may be additional degeneracies. For instance, a term $\beta_d z^d$ may reduce eventually to a term $\beta_d dz^{d-1}h$, and if the underlying ring has characteristic dividing $\beta_d d$, then this term vanishes, which requires a more refined treatment. Other complications arise if the underlying ring has zero divisors. See [8] for an example of a PET argument dealing with the latter issue.

for some indices $w_{j\underline{u}} \in [0, \ell]$ (with $\beta_{0(|\underline{u}|+1)} := \mathbf{0}$). Moreover, each nonzero vector $\beta_{1(|\underline{u}|+1)} - \beta_{w_{j\underline{u}}(|\underline{u}|+1)}$ from (5.29) is the leading coefficient of $\mathbf{P}_1 - \mathbf{P}_{w_{j\underline{u}}}$.

We typically apply Proposition 5.23 with $K \asymp_d (N/V)^{1/d}$ for $V := \max_{j,i} |\beta_{ji}|$, as this provides the maximum possible range for which $\mathbf{P}_j(z) \in [\pm N]^D$ for all $j \in [\ell]$ and $z \in [K]$. Since we do not place any special assumptions on $\mathbf{P}_1$, the role of f_1 in the conclusion can be replaced by any of the functions $f_0, \dots, f_\ell$ (with an appropriate modification of the formulas for (5.29)). For $f_2, \dots, f_\ell$, this is obvious by symmetry, whereas for f_0 this follows after shifting $\mathbf{x} \mapsto \mathbf{x} - \mathbf{P}_m$ for some arbitrary $m \in [\ell]$ and observing (as in the proof of Proposition 5.23 below) that the resulting collections of polynomials (5.30) are identical.

Proof. We will proceed under the assumption that $\mathbf{P}_1$ has the largest degree among the $\mathbf{P}_j$'s; at the end we will describe the necessary modifications in the general case.

Proposition 5.22 provides $r = O_{d,\ell}(1)$ and $m_1, \dots, m_r \in \mathbb{N}$ such the families $\mathcal{Q}_i = \partial_{m_i} \dots \partial_{m_1} \mathcal{P}$ all descend from $\mathcal{P}$ and $\mathcal{Q}_r$ is linear in z . Let

$$\mathbf{b}_1(\underline{h})z, \dots, \mathbf{b}_s(\underline{h})z$$

be the elements of $\mathcal{Q}_r$. Since the van der Corput operation at most doubles the length of a polynomial family, we have $s := |\mathcal{Q}_r| \leq 2^r |\mathcal{P}| \ll_{d,\ell} 1$. Applying Lemma 5.18 inductively to each of $\mathcal{Q}_0, \dots, \mathcal{Q}_{r-1}$ (with $\mathcal{Q}_0 := \mathcal{P}$), we deduce that

$$\sum_{\underline{h} \in \mathbb{Z}^r} \mu_H(\underline{h}) \sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{z \in [K]} f_0(\mathbf{x}) \cdot \prod_{j=1}^s f_j(\mathbf{x} + \mathbf{b}_j(\underline{h})z) \gg_{d,D,\ell} \delta^{O_{d,\ell}(1)} N^D.$$

The power of δ depends only on the number of iterations of Lemma 5.18; in particular, it is independent of D .

Since the family $\mathcal{Q}_r$ descends from $\mathcal{P}$, the polynomials $\mathbf{b}_1, \dots, \mathbf{b}_s$ are all distinct. Lemma 5.19 then implies that

$$\sum_{\underline{h} \in \mathbb{Z}^r} \mu_H(\underline{h}) \sum_{\underline{m}, \underline{m}' \in [\pm M]^s} \mathbb{E}_{\mathbf{x} \in \mathbb{Z}^D} \Delta'_{\mathbf{C}_1(\underline{h}), \dots, \mathbf{C}_s(\underline{h}); (\underline{m}, \underline{m}')} f_1(\mathbf{x}) \gg_{d,D,\ell} \delta^{O_{d,\ell}(1)} N^D,$$

where $\mathbf{C}_1, \dots, \mathbf{C}_s$ are the polynomials $\mathbf{b}_1, \mathbf{b}_1 - \mathbf{b}_2, \dots, \mathbf{b}_1 - \mathbf{b}_s$. Since $\mu_H(\underline{h}) \leq \frac{1_{[-2H, 2H]^r}(\underline{h})}{(4H+1)^r}$ pointwise and the inner average is nonnegative (by virtue of being an average of box norms), the inequality in the conclusion of the proposition follows after we replace H with $H/2$.

That the polynomials $\mathbf{C}_1, \dots, \mathbf{C}_s$ are nonzero is a consequence of the distinctness of $\mathbf{b}_1, \dots, \mathbf{b}_s$, which in

turn is a consequence of the tuple $\mathcal{Q}_r$ being normal. The form of $\mathbf{C}_1, \dots, \mathbf{C}_s$ is then a consequence of the properties (1)-(3) from Definition 5.20.

We now address the case where $\mathbf{P}_1$ does not have the largest degree. Then some other $\mathbf{P}_m$ has the largest degree, and we can translate $\mathbf{x} \mapsto \mathbf{x} - \mathbf{P}_m(z)$ so that

$$\begin{aligned} \sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{z \in [K]} f_0(\mathbf{x}) \cdot f_1(\mathbf{x} + \mathbf{P}_1(z)) \cdots f_\ell(\mathbf{x} + \mathbf{P}_\ell(z)) \\ = \sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{z \in [K]} f_0(\mathbf{x} - \mathbf{P}_m(z)) \cdot \prod_{j \in [\ell]} f_j(\mathbf{x} + \mathbf{P}_j(z) - \mathbf{P}_m(z)). \end{aligned}$$

We set

$$\tilde{\mathbf{P}}_j = \begin{cases} \mathbf{P}_j - \mathbf{P}_m, & j \neq m, \\ -\mathbf{P}_m, & j = m, \end{cases}$$

so that $\tilde{\mathbf{P}}_1$ has the largest degree among the $\tilde{\mathbf{P}}_j$'s and we can apply the argument above (after exchanging f_0 and f_m). The desired result now follows from the observation that

$$\{\mathbf{P}_1, \mathbf{P}_1 - \mathbf{P}_j : j \in [2, \ell]\} \quad \text{and} \quad \{\tilde{\mathbf{P}}_1, \tilde{\mathbf{P}}_1 - \tilde{\mathbf{P}}_j : j \in [2, \ell]\} \quad (5.30)$$

are identical as sets and hence these families of polynomials have the same leading coefficients. $\square$

5.5 Model example: concatenation of degree-1 box norms

To illustrate some of the ideas that go into the proof of Theorem 5.10, we start with a model concatenation result for averages of degree-1 box norms.

Proposition 5.24. *Let $d, D, r \in \mathbb{N}$ with $d \leq r$. Let $\delta > 0$ and $H, M, N \in \mathbb{N}$ satisfy*

$$\delta^{-Or(1)} \ll_{D,r} H \leq M \quad \text{and} \quad H^d M \leq N, \quad (5.31)$$

and let $\mathbf{C}(\underline{h}) = \sum_{\underline{u} \in \{0,1\}^r} \gamma_{\underline{u}} \underline{h}^{\underline{u}} \in \mathbb{Z}^D[h_1, \dots, h_r]$ be a multilinear polynomial of degree d with coefficients of size $O(1)$. Then for all 1-bounded functions $f : \mathbb{Z}^D \rightarrow \mathbb{C}$ supported on $[N]^D$, the bound

$$\mathbb{E}_{\underline{h} \in [\pm H]^r} \|f\|_{\mathbf{C}(\underline{h}) \cdot [\pm M]}^2 \geq \delta N^D \quad (5.32)$$

implies that

$$\|f\|_{[\pm N]^D, E}^4 \gg_{D,r} \delta^{O_r(1)} N^D,$$

where

$$E = \sum_{\underline{u} \in \{0,1\}^r} \gamma_{\underline{u}} \cdot [\pm H^{|\underline{u}|} M]. \quad (5.33)$$

We observe that the multiset E is a generalized arithmetic progression spanned by the coefficients $\gamma_{\underline{u}}$, and that the length of the interval corresponding to the coefficient $\gamma_{\underline{u}}$ is of the same order of magnitude as the maximum value of the monomial $\underline{h}^{\underline{u}} m$. Moreover, if $\gamma_{\underline{u}} = \mathbf{0}$, i.e., the monomial $\underline{h}^{\underline{u}}$ does not appear in the polynomial $\mathbf{C}$, then the corresponding term does not contribute to the set E .

Example 5.25. To illustrate the content of Proposition 5.24 with a specific example, we consider the polynomial

$$\mathbf{C}(h_1, h_2) = \mathbf{e}_1 h_1 h_2 + \mathbf{e}_2 h_1 = (h_1 h_2, h_1),$$

where $D = 2$, $H = M = N^{1/3}$; for simplicity, we take N to be a perfect cube. Proposition 5.24 tells us that if

$$\sum_{x_1, x_2, m} \mathbb{E}_{h_1, h_2 \in [\pm N^{1/3}]} \mu_{N^{1/3}}(m) f(x_1, x_2) \overline{f(x_1 + h_1 h_2 m, x_2 + h_1 m)} \geq \delta N^2,$$

or equivalently

$$\sum_{\mathbf{x}, m} \mathbb{E}_{h \in [\pm N^{1/3}]^2} \mu_{N^{1/3}}(m) f(\mathbf{x}) \overline{f(\mathbf{x} + (\mathbf{e}_1 h_1 h_2 + \mathbf{e}_2 h_1) m)} \geq \delta N^2,$$

then

$$\sum_{\mathbf{x}} \mathbb{E}_{\mathbf{n}_1, \mathbf{n}'_1 \in [\pm N]^2} \mathbb{E}_{\mathbf{n}_2, \mathbf{n}'_2 \in [\pm N] \times [\pm N^{2/3}]} \Delta'_{(\mathbf{n}_1, \mathbf{n}'_1), (\mathbf{n}_2, \mathbf{n}'_2)} f(\mathbf{x}) \gg \delta^{O(1)} N^2.$$

Writing $\mathbf{n}_2 = (n_{21}, n_{22})$ and $\mathbf{n}'_2 = (n'_{21}, n'_{22})$, we see that the variables n_{21}, n'_{21} correspond to the monomial $h_1 h_2$ (equivalently, to the coefficient $\underline{u} = (1, 1)$) and hence run along the coefficient $\mathbf{e}_1$ of $h_1 h_2$ in $\mathbf{C}$; the variables n_{22}, n'_{22} correspond to the monomial h_1 (equivalently, to $\underline{u} = (1, 0)$) and therefore run along the coefficient $\mathbf{e}_2$ of h_1 in $\mathbf{C}$.

In applications, we usually want to discard the directions that correspond to lower-degree terms in the

original polynomial, such as the direction $\mathbf{e}_1$ corresponding to the index $(1, 0)$ in the example above. We accomplish this using Lemma 5.15, and obtain the following streamlined version of Proposition 5.24.

Corollary 5.26. *Let $d, D, r \in \mathbb{N}$ with $d \leq r$. Let $N \in \mathbb{N}$ and $\delta > 0$ satisfy $N \gg_{D,r} \delta^{-O_r(1)}$, and let $C(\underline{h}) = \sum_{\underline{u} \in \{0,1\}^r} \gamma_{\underline{u}} \underline{h}^{\underline{u}} \in \mathbb{Z}^D[h_1, \dots, h_r]$ be a multilinear polynomial of degree d with coefficients of size $O(1)$. Then for all 1-bounded functions $f : \mathbb{Z}^D \rightarrow \mathbb{C}$ supported on $[N]^D$, the lower bound*

$$\mathbb{E}_{\underline{h} \in [\pm N^{1/(d+1)}]^r} \|f\|_{C(\underline{h}) \cdot [\pm N^{1/(d+1)}]}^2 \geq \delta N^D$$

implies that

$$\|f\|_{(\gamma_{\underline{u}} \cdot [\pm N])^2}^4 \gg_{D,r} \delta^{O_r(1)} N^D$$

for each $\underline{u} \in \{0,1\}^r$ with $|\underline{u}| = d$ and $\gamma_{\underline{u}} \neq 0$.

Proof. Proposition 5.24 applied with $H = M = N^{1/(d+1)}$ gives

$$\|f\|_{[\pm N]^D, E}^4 \gg_{D,r} \delta^{O_r(1)} N^D,$$

where E is as in (5.33). Since the coefficients $\gamma_{\underline{u}}$ have size $O(1)$, we can find vectors $\mathbf{v}_1, \dots, \mathbf{v}_k \in \mathbb{Z}^D$ (with $k < D$) of size $O(1)$ such that the set

$$\{\gamma_{\underline{u}} : \underline{u} \in \{0,1\}^r, |\underline{u}| = d, \gamma_{\underline{u}} \neq 0\} \cup \{\mathbf{v}_1, \dots, \mathbf{v}_k\}$$

generates a subgroup in $\mathbb{Z}^D$ of index $O(1)$. Hence there exists a positive integer such that

$$C \cdot [\pm N/C]^D \subseteq \sum_{\substack{\underline{u} \in \{0,1\}^r, \\ |\underline{u}| = d}} \gamma_{\underline{u}} \cdot [\pm CN] + \sum_{i=1}^k \mathbf{v}_i \cdot [\pm CN],$$

and so by successively applying properties (7), (4), (5) and (6) of Lemma 5.15, we get

$$\|f\|_{E', E}^4 \gg_{D,r} \delta^{O_r(1)} N^D$$

for

$$E' = \sum_{\substack{\underline{u} \in \{0,1\}^r, \\ |\underline{u}| = d}} \gamma_{\underline{u}} \cdot [\pm N].$$

To conclude from here, we pick a single $\underline{u}$ and apply Lemma 5.15(5). □

Corollary 5.26 allows us to upgrade the conclusion of Example 5.25 to

$$\sum_{\mathbf{x}} \mathbb{E}_{n_1, n'_1, n_2, n'_2 \in [\pm N]} \Delta'_{(n_1 \mathbf{e}_1, n'_1 \mathbf{e}_1), (n_2 \mathbf{e}_1, n'_2 \mathbf{e}_1)} f(\mathbf{x}) \gg \delta^{O(1)} N^2.$$

5.5.1 Sketch of the proof of Proposition 5.24

We now outline the proof of Proposition 5.24 in order to present the logic behind our concatenation argument in a relatively simple setting. This is an expanded version of the overview given in Section 5.2. The argument gets considerably more complicated in the setting of box norms of degree greater than 1, so the reader may find the simplified argument useful for grasping the main ideas. In the sketch, we use as a black box some technical results that will be derived in later sections. . Suppose that (5.32) holds. Expanding the definition of the box norm gives us

$$\mathbb{E}_{\underline{h} \in [\pm H]^r} \mathbb{E}_{m, m' \in [\pm M]} \sum_{\mathbf{x}} \Delta'_{(\mathbf{C}(\underline{h})m, \mathbf{C}(\underline{h})m')} f(\mathbf{x}) \geq \delta N^D.$$

Shifting $\mathbf{x} \mapsto \mathbf{x} - \mathbf{C}(\underline{h})m$, replacing $m' - m$ by m , and changing the order of summation, we get

$$\sum_{\mathbf{x}} f(\mathbf{x}) \mathbb{E}_{\underline{h} \in [\pm H]^r} \sum_m \mu_M(m) \overline{f(\mathbf{x} + \mathbf{C}(\underline{h})m)} \geq \delta N^D.$$

We will use several applications of the Cauchy–Schwarz inequality to replace the polynomial $\mathbf{C}(\underline{h})m$ by one with better equidistribution properties: Specifically, we want to replace it by a polynomial in a larger number of variables that does not assume any particular value more than a constant multiple of the “expected” number of times. Each application of the Cauchy–Schwarz inequality, which increases the number of variables, can be understood as a “smoothing” of the polynomial differencing parameter.

Step 1: Replacing $\mathbf{C}(\underline{h})m$ by a polynomial with better equidistribution properties.

Applying the Cauchy–Schwarz inequality in $\mathbf{x}$, we double the $\underline{h}, m$ variables and obtain

$$\sum_{\mathbf{x}} \mathbb{E}_{\underline{h}_1, \underline{h}_2 \in [\pm H]^r} \sum_{m_1, m_2} \mu_M(m_1, m_2) f(\mathbf{x} + \mathbf{C}(\underline{h}_1)m_1) \overline{f(\mathbf{x} + \mathbf{C}(\underline{h}_2)m_2)} \geq \delta^2 N^D,$$

then we shift $\mathbf{x} \mapsto \mathbf{x} - \mathbf{C}(\underline{h}_1)m_1$ to get

$$\sum_{\mathbf{x}} f(\mathbf{x}) \mathbb{E}_{\underline{h}_1, \underline{h}_2 \in [\pm H]^r} \sum_{m_1, m_2} \mu_M(m_1, m_2) \overline{f(\mathbf{x} + \mathbf{C}(\underline{h}_2)m_2 - \mathbf{C}(\underline{h}_1)m_1)} \geq \delta^2 N^D.$$

Applying the Cauchy–Schwarz inequality in $\mathbf{x}$ again gives

$$\sum_{\mathbf{x}} \mathbb{E}_{\substack{\underline{h}_1, \underline{h}_2, \\ \underline{h}_3, \underline{h}_4 \in [\pm H]^r}} \sum_{\substack{m_1, m_2, \\ m_3, m_4}} \mu_M(m_1, m_2, m_3, m_4) \\ f(\mathbf{x} + \mathbf{C}(\underline{h}_2)m_2 - \mathbf{C}(\underline{h}_1)m_1) \overline{f(\mathbf{x} + \mathbf{C}(\underline{h}_4)m_4 - \mathbf{C}(\underline{h}_3)m_3)} \geq \delta^4 N^D,$$

and, after changing variables $\mathbf{x} \mapsto \mathbf{x} - (\mathbf{C}(\underline{h}_2)m_2 - \mathbf{C}(\underline{h}_1)m_1)$, $m_2 \mapsto -m_2$, $m_3 \mapsto -m_3$, we obtain

$$\mathbb{E}_{\substack{\underline{h}_1, \underline{h}_2, \\ \underline{h}_3, \underline{h}_4 \in [\pm H]^r}} \sum_{\substack{m_1, m_2, \\ m_3, m_4}} \mu_M(m_1, m_2, m_3, m_4) \\ \sum_x f(\mathbf{x}) \overline{f(\mathbf{x} + \mathbf{C}(\underline{h}_1)m_1 + \mathbf{C}(\underline{h}_2)m_2 + \mathbf{C}(\underline{h}_3)m_3 + \mathbf{C}(\underline{h}_4)m_4)} \geq \delta^4 N^D.$$

Thus, at the cost of increasing the power of δ in our lower bound, we have replaced $\mathbf{C}(\underline{h})m$ by the polynomial

$$\sum_{i \in [4]} \mathbf{C}(\underline{h}_i)m_i = \mathbf{C}(\underline{h}_1)m_1 + \mathbf{C}(\underline{h}_2)m_2 + \mathbf{C}(\underline{h}_3)m_3 + \mathbf{C}(\underline{h}_4)m_4,$$

which is more uniformly distributed in its range than the original polynomial $\mathbf{C}(\underline{h})m$.

However, this new polynomial is still not quite as equidistributed as we want. For instance, if $\mathbf{C}(\underline{h}) = \mathbf{e}_1 h_1 h_2 + \mathbf{e}_2 h_1$ (as in Example 5.25), then the new polynomial

$$\begin{aligned} & \mathbf{e}_1(h_{11}h_{21}m_1 + h_{12}h_{22}m_2 + h_{13}h_{23}m_3 + h_{14}h_{24}m_4) \\ & + \mathbf{e}_2(h_{11}m_1 + h_{12}m_2 + h_{13}m_3 + h_{14}m_4) \end{aligned}$$

is still not equidistributed in its “natural” range $\mathbf{e}_2 \cdot [\pm 4H^2M] + \mathbf{e}_1 \cdot [\pm 4HM]$. The reason is that

$$g = \gcd(h_{11}h_{21}, h_{12}h_{22}, h_{13}h_{23}, h_{14}h_{24})$$

is quite large for many h ’s. For example, if $p \ll H^{\Omega(1)}$ is prime, then $p^2 | g$ with probability $\approx (3/p^2 - 2/p^3)^4$, which is much larger than the probability $\approx (1/p^2)^4$ that p^2 divides the gcd of four uniformly random integers of size $\asymp H^2$. This means that the $\mathbf{e}_2$ -coordinate of the polynomial will be a multiple of p^2 disproportionately often as the h ’s and m ’s range.

To eliminate this arithmetic obstruction in the general setting of Proposition 5.24, we use more applications of the Cauchy–Schwarz inequality to increase the number of the variables $h_{2i}, \dots, h_{ri}$ and uniformize the polynomial further; after enough duplications of variables using the Cauchy–Schwarz inequality, the

resulting polynomial will be equidistributed in its natural range. Writing $\underline{h}_i = (h_{1i}, \tilde{h}_i)$, we now apply the Cauchy–Schwarz inequality twice in $\mathbf{x}$ and h_{1i} to obtain

$$\mathbb{E}_{(h_{1i})_{i_1 \in [4]} \in [\pm H]^4} \mathbb{E}_{(\tilde{h}_{i_1 i_2})_{i_1, i_2 \in [4]} \in [\pm H]^{16(r-1)}} \sum_{(m_{i_1 i_2})_{i_1, i_2 \in [4]}} \mu_M((m_{i_1 i_2})_{i_1, i_2}) \overline{\sum_{\mathbf{x}} f(\mathbf{x}) f\left(\mathbf{x} + \sum_{i_1, i_2 \in [4]} \mathbf{C}(h_{1i_1}, \tilde{h}_{i_1 i_2}) m_{i_1 i_2}\right)} \geq \delta^{16} N^D.$$

We continue in this manner: Applying the Cauchy–Schwarz inequality twice in the variables

$$\mathbf{x}, h_{1i_1}, h_{2i_1 i_2}, \dots, h_{(l-1)i_1 \dots i_{l-1}}$$

for each $l = 3, 4, \dots, r$ (as in the $l = 2$ step just illustrated; the $l = 1$ step appeared at the very beginning of the proof), and bounding the Fejér kernels pointwise, we arrive at the bound

$$\mathbb{E}_{(h_{1i_1})_{i_1 \in [4]} \in [\pm H]^4} \cdots \mathbb{E}_{(h_{ri_1 \dots i_r})_{i_1, \dots, i_r \in [4]} \in [\pm H]^{4^r}} \mathbb{E}_{(m_{i_1 \dots i_r})_{i_1, \dots, i_r \in [4]} \in [\pm 2M]^{4^r}} \left| \sum_{\mathbf{x}} f(\mathbf{x}) f\left(\mathbf{x} + \sum_{i_1, \dots, i_r \in [4]} \mathbf{C}(h_{1i_1}, \dots, h_{ri_1 \dots i_r}) m_{i_1 \dots i_r}\right) \right| \gg_r \delta^{4^r} N^D.$$

Step 2: Applying the equidistribution properties of the new polynomial.

We have now replaced $\mathbf{C}(\underline{h})m$ by the massive polynomial

$$\sum_{i_1, \dots, i_r \in [4]} \mathbf{C}(h_{1i_1}, \dots, h_{ri_1 \dots i_r}) m_{i_1 \dots i_r} = \sum_{i_1, \dots, i_r \in [4]} \gamma_{\underline{u}} \sum_{i_1 \in [4]} h_{1i_1}^{u_1} \sum_{i_2 \in [4]} h_{2i_1 i_2}^{u_2} \cdots \sum_{i_r \in [4]} h_{ri_1 \dots i_r}^{u_r} m_{i_1 \dots i_r}. \quad (5.34)$$

For instance, if $\mathbf{C}(\underline{h}) = \mathbf{e}_1 h_1 h_2 + \mathbf{e}_2 h_1$ (as in Example 5.25), then the polynomial (5.34) equals

$$\mathbf{e}_1 \sum_{i_1, i_2 \in [4]} h_{1i_1} h_{2i_1 i_2} m_{i_1 i_2} + \mathbf{e}_2 \sum_{i_1, i_2 \in [4]} h_{1i_1} m_{i_1 i_2}.$$

The crucial point is that for a generic choice of h 's, the subpolynomials of (5.34) corresponding to different $\underline{u}$'s are simultaneously well distributed in their ranges. This result on simultaneous equidistribution is one of the new inputs of our work, and we outsource its exact statement and proof to Proposition 5.35. For now, however, let us briefly explain how this good distribution property of the polynomial (5.34) arises from its

“multiscale” structure. For a generic choice of $h_{ri_1 \dots i_r}$, the polynomial

$$\left(\sum_{i_r \in [4]} m_{i_1 \dots i_r}, \sum_{i_r \in [4]} h_{ri_1 \dots i_r} m_{i_1 \dots i_r} \right)$$

is well distributed in its range; this allows us to essentially replace both sums by new variables $t_{i_1 \dots i_{r-1}0}$ and $t_{i_1 \dots i_{r-1}1}$ ranging over $[\pm 4M], [\pm 4HM]$ respectively. Then for a generic choice of $h_{(r-1)i_1 \dots i_{r-1}}$, the polynomial

$$\left(\sum_{i_{r-1} \in [4]} t_{i_1 \dots i_{r-1}0}, \sum_{i_{r-1} \in [4]} t_{i_1 \dots i_{r-1}1}, \sum_{i_{r-1} \in [4]} h_{(r-1)i_1 \dots i_{r-1}} t_{i_1 \dots i_{r-1}0}, \sum_{i_{r-1} \in [4]} h_{(r-1)i_1 \dots i_{r-1}} t_{i_1 \dots i_{r-1}1} \right)$$

is well distributed in its range $[\pm 16M] \times [\pm 16HM] \times [\pm 16HM] \times [\pm 16H^2M]$, and so on.¹⁶ Because of this iterative scheme, it suffices to show that (5.34) is well distributed in its range for $r = 1$, and this can be accomplished using elementary counting arguments.

For $l \in [r]$ and $\eta > 0$, define the set

$$\begin{aligned} \mathcal{H}_{l,\eta} = \{ & (h_{li_1 \dots i_l})_{i_1, \dots, i_l \in [4]} \in [\pm H]^{4^l} : |h_{li_1 \dots i_l} - h_{li'_1 \dots i'_l}| \geq \eta H, \\ & \gcd(h_{li_1 \dots i_l} - h_{li'_1 \dots i'_l}, h_{li'_1 \dots i'_l} - h_{li''_1 \dots i''_l}) \leq \eta^{-1} \\ & \text{for all distinct } (i_1, \dots, i_l), (i'_1, \dots, i'_l), (i''_1, \dots, i''_l) \in [4]^l \}. \end{aligned}$$

This is a set of tuples $(h_{li_1 \dots i_l})_{i_1, \dots, i_l}$ whose coordinates are not too close to each other and do not have large common factors. It is not difficult to show that for a small $\eta > 0$, a generic tuple in $[\pm H]^{4^l}$ belongs to this set. Indeed, taking $\mathcal{H}_l := \mathcal{H}_{l, c\delta^{1/c}}$ for sufficiently small $c = c(r) > 0$, we have by Lemma 5.38 that all but a $O_r(\delta^{4^r})$ -proportion of the elements of $[\pm H]^{4^l}$ lie in $\mathcal{H}_l$, and hence

$$\begin{aligned} & \mathbb{E}_{(h_{1i_1})_{i_1 \in [4]} \in [\pm H]^4} \cdots \mathbb{E}_{(h_{ri_1 \dots i_r})_{i_1, \dots, i_r \in [4]} \in [\pm H]^{4^r}} \mathbb{E}_{(m_{i_1 \dots i_r})_{i_1, \dots, i_r \in [4]} \in [\pm 2M]^{4^r}} \\ & \prod_{l \in [r]} \mathbf{1}_{\mathcal{H}_l(h_{li_1 \dots i_l})_{i_1, \dots, i_l}} \cdot \left| \sum_{\mathbf{x}} f(\mathbf{x}) f\left(\mathbf{x} + \sum_{i_1, \dots, i_r \in [4]} \mathbf{C}(h_{1i_1}, \dots, h_{ri_1 \dots i_r}) m_{i_1 \dots i_r}\right) \right| \gg_r \delta^{4^r} N^D. \end{aligned}$$

We have thus excluded the h 's whose coefficients are close to each other or have large common factors.

For $(n_{\underline{u}})_{\underline{u}} \in \mathbb{Z}^{2^r}$, let

$$\mathcal{N}((n_{\underline{u}})_{\underline{u}}) = \mathbb{E}_{(h_{1i_1})_{i_1 \in [4]} \in [\pm H]^4} \cdots \mathbb{E}_{(h_{ri_1 \dots i_r})_{i_1, \dots, i_r \in [4]} \in [\pm H]^{4^r}} \mathbb{E}_{(m_{i_1 \dots i_r})_{i_1, \dots, i_r \in [4]} \in [\pm 2M]^{4^r}}$$

¹⁶For this to happen, we need for each $l \in [r]$ to sum over at least 3 indices i_l ; we sum over 4 since the Cauchy–Schwarz inequality doubles the number of variables each time.

$$\prod_{l \in [r]} \mathbf{1}_{\mathcal{H}_l((h_{li_1 \dots i_l})_{i_1, \dots, i_l})} \cdot \prod_{\underline{u} \in \{0,1\}^r} \mathbf{1}\left(\sum_{i_1 \in [4]} h_{1i_1}^{u_1} \cdots \sum_{i_r \in [4]} h_{ri_1 \dots i_r}^{u_r} m_{i_1 \dots i_r} = n_{\underline{u}}\right),$$

so that

$$\sum_{(n_{\underline{u}})_{\underline{u}}} \mathcal{N}((n_{\underline{u}})_{\underline{u}}) \cdot \left| \sum_{\mathbf{x}} f(\mathbf{x}) \overline{f\left(\mathbf{x} + \sum_{\underline{u} \in \{0,1\}^r} \gamma_{\underline{u}} n_{\underline{u}}\right)} \right| \gg_r \delta^{4^r} N^D. \quad (5.35)$$

The quantity $\mathcal{N}((n_{\underline{u}}))$ is the normalized count of $(\underline{h}_{1i_1}, \dots, \underline{h}_{ri_1 \dots i_r}, \underline{m}_{i_1 \dots i_r})$ such that for each $\underline{u} \in \{0,1\}^r$, the subpolynomial associated with $\gamma_{\underline{u}}$ in (5.34) equals $n_{\underline{u}}$. The key equidistribution result on the polynomial (5.34) is the assertion that the normalised count $\mathcal{N}$ satisfies the pointwise bound

$$\mathcal{N}((n_{\underline{u}})_{\underline{u}}) \ll_r \delta^{-O_r(1)} M^{-2^r} H^{-r2^{r-1}}. \quad (5.36)$$

This bound is a special case of Proposition 5.35 to be derived later on.

Let us explain why the bound (5.36) is of the right order of magnitude. Note first that $\mathcal{N}((n_{\underline{u}})_{\underline{u}}) = 0$ unless each $|n_{\underline{u}}| \leq C_{\underline{u}} H^{|\underline{u}|} M$ for some $C_{\underline{u}} > 0$. Defining

$$E_0 = \left\{ (n_{\underline{u}})_{\underline{u} \in \{0,1\}^r} : n_{\underline{u}} \in [-C_{\underline{u}} H^{|\underline{u}|} M, C_{\underline{u}} H^{|\underline{u}|} M], \underline{u} \in \{0,1\}^r \right\},$$

we observe that

$$|E_0| \asymp_r \prod_{\underline{u} \in \{0,1\}^r} (H^{|\underline{u}|} M) = M^{2^r} H^{r2^{r-1}}$$

and note that the exponents agree with those in (5.36). Combining these observations with (5.35) and (5.36), we get

$$\mathbb{E}_{(n_{\underline{u}})_{\underline{u}} \in E_0} \left| \sum_{\mathbf{x}} f(\mathbf{x}) \overline{f\left(\mathbf{x} + \sum_{\underline{u} \in \{0,1\}^r} \gamma_{\underline{u}} n_{\underline{u}}\right)} \right| \gg_r \delta^{O_r(1)} N^D.$$

Defining

$$E_1 = \sum_{\underline{u} \in \{0,1\}^r} \gamma_{\underline{u}} \cdot [\pm C_{\underline{u}} H^{|\underline{u}|} M],$$

we have

$$\mathbb{E}_{\mathbf{n} \in E_1} \left| \sum_{\mathbf{x}} f(\mathbf{x}) \overline{f(\mathbf{x} + \mathbf{n})} \right| \gg_r \delta^{O_r(1)} N^D.$$

Using the Cauchy–Schwarz inequality to remove the absolute value and recalling that f is supported on $[N]^D$, we obtain

$$\mathbb{E}_{\mathbf{n} \in E_1} \mathbb{E}_{\mathbf{s} \in [\pm N]^D} \sum_{\mathbf{x}} f(\mathbf{x}) \overline{f(\mathbf{x} + \mathbf{s})} \overline{f(\mathbf{x} + \mathbf{n})} f(\mathbf{x} + \mathbf{s} + \mathbf{n}) \gg_r \delta^{O_r(1)} N^D.$$

Two further applications of the Cauchy–Schwarz inequality to double $\mathbf{n}, \mathbf{s}$ let us introduce Fejér kernels, and we get

$$\|f\|_{E_1, [\pm N]^D}^4 \gg_r \delta^{O_r(1)} N^D.$$

To conclude the proof, we use Lemma 5.15(6) to replace E_1 by E .

5.6 Concatenation of averages of higher-degree seminorms

In the previous section, we showed how to concatenate an average of degree-1 box norms with polynomial directions. The argument becomes significantly more complicated when we deal with higher-degree box norms. In the degree-1 setting, applying the Cauchy–Schwarz inequality (as in the proof of Proposition 5.24) did not cause the degree of the box norm to increase. This pleasant property ceases to hold in the setting of higher-degree box norms; rather, each application of the Cauchy–Schwarz and Gowers–Cauchy–Schwarz Inequalities will involve passing to norms of much higher degree. The objective of the present section is to carry out this higher-degree argument. As noted in the proof outline, this stage of the concatenation arguments works for any box norms, not necessarily those that come out of PET and involve polynomials. To emphasize the robustness of this part of the concatenation argument, we phrase the results in this section in a more abstract setting of an average of box norms on an arbitrary finitely generated abelian group G . We will thus deal with expressions of the form

$$\mathbb{E}_{i \in I} \|f\|_{H_{1i}, \dots, H_{si}}^{2^s},$$

where I is a finite indexing set and $H_{1i}, \dots, H_{si} \subseteq G$ are finite nonempty multisets for each $i \in I$. In our applications, the set I will be $\underline{h} \in [\pm H]^r$, and the multisets $H_{1i}, \dots, H_{si}$ for $i \in I$ will be the sets

$$\mathbf{C}_1(\underline{h}) \cdot [\pm M], \dots, \mathbf{C}_s(\underline{h}) \cdot [\pm M]$$

for $\underline{h} \in [\pm H]^r$.

The following lemma will serve as the basis for the concatenation argument.

Lemma 5.27. *Let $A, \delta > 0$, let $s \in \mathbb{N}$, let G be a finitely generated abelian group, let I be a finite indexing set, and let $H_i, K_{1i}, \dots, K_{si} \subseteq G$ be finite multisets for each $i \in I$. Suppose that $B \subseteq G$ is a set satisfying*

$$|\text{supp}(B + B - B)| \leq A|B|.$$

Then for all 1-bounded functions f supported on B , the bound

$$\mathbb{E}_{i \in I} \|f\|_{H_i, K_{1i}, \dots, K_{si}}^{2^{s+1}} \geq \delta|B|$$

implies that

$$\mathbb{E}_{i_0, i_1 \in I} \|f\|_{K_{1i_0}, \dots, K_{si_0}, K_{1i_1}, \dots, K_{si_1}, H_{i_1} - H_{i_0}}^{2^{2s+1}} \gg_{A, s} \delta^{2^{2s+1}} |B|.$$

In applications, we will take $B = [N]^D$ and $A = 3^D$ since

$$\text{supp}(B + B - B) = [2 - N, 2N - 1]^D.$$

Proof. Let $K_i := K_{1i} \times \dots \times K_{si}$. The alternative formula (5.17) for box norms yields

$$\mathbb{E}_{i \in I} \sum_x \sum_{\underline{k} \in K_i} \sum_{h \in H_i} \mu_{K_i}(\underline{k}) \mu_{H_i}(h) \Delta_{\underline{k}, h} f(x) \geq \delta|B|.$$

We expand

$$\Delta_{\underline{k}, h} f(x) = f(x) \overline{f(x + h)} \Delta_{\underline{k}}^* \Delta_h f(x),$$

where $\Delta_{\underline{k}}^* f(x) = \prod_{\epsilon \in \{0,1\}^s \setminus \{0\}} \mathcal{C}^{|\epsilon|} f(x + \epsilon \cdot \underline{k})$, and change the order of summation to get

$$\sum_x f(x) \mathbb{E}_{i \in I} \sum_{\underline{k} \in K_i} \sum_{h \in H_i} \mu_{K_i}(\underline{k}) \mu_{H_i}(h) \overline{f(x + h)} \Delta_{\underline{k}}^* \Delta_h f(x) \geq \delta|B|.$$

An application of the Cauchy–Schwarz inequality in x gives

$$\sum_x \mathbb{E}_{i_0, i_1 \in I} \sum_{\substack{\underline{k}_0 \in K_{i_0}, h_0 \in H_{i_0}, \\ \underline{k}_1 \in K_{i_1}, h_1 \in H_{i_1}}} \mu_{K_{i_0}}(\underline{k}_0) \mu_{K_{i_1}}(\underline{k}_1) \mu_{H_{i_0}}(h_0) \mu_{H_{i_1}}(h_1) \overline{f(x + h_0)} f(x + h_1)$$

$$\Delta_{\underline{k}_0}^* \Delta_{h_0} f(x) \overline{\Delta_{\underline{k}_1}^* \Delta_{h_1} f(x)} \geq \delta^2 |B|.$$

We crucially observe that for each fixed $i_0, i_1 \in I$ and $h_0 \in H_{i_0}, h_1 \in H_{i_1}$, the average

$$\sum_x \sum_{\substack{\underline{k}_0 \in K_{i_0}, \\ \underline{k}_1 \in K_{i_1}}} \mu_{K_{i_0}}(\underline{k}_0) \mu_{K_{i_1}}(\underline{k}_1) f(x + h_0) \overline{f(x + h_1)} \Delta_{\underline{k}_0}^* \Delta_{h_0} f(x) \overline{\Delta_{\underline{k}_1}^* \Delta_{h_1} f(x)} \quad (5.37)$$

is a box inner product of 1-bounded functions along $K_{1i_0}, \dots, K_{si_0}, K_{1i_1}, \dots, K_{si_1}$, with $f(x + h_0) \overline{f(x + h_1)}$ being the term with multi-index $\underline{0}$. Moreover, since f is supported on B , the summand in (5.37) vanishes unless

$$x + h_j, x + \underline{\epsilon}_j \cdot \underline{k}_j, x + \underline{\epsilon}_j \cdot \underline{k}_j + h_j \in B \quad \text{for all } j \in \{0, 1\}, \underline{\epsilon}_j \neq \underline{0}.$$

In particular, if all of these elements lie in B , then we also have

$$x + \underline{\epsilon}_0 \cdot \underline{k}_0 + \underline{\epsilon}_1 \cdot \underline{k}_1 + h_0 = (x + \underline{\epsilon}_0 \cdot \underline{k}_0 + h_0) + (x + \underline{\epsilon}_1 \cdot \underline{k}_1 + h_1) - (x + h_1) \in B + B - B$$

for any $\underline{\epsilon}_0, \underline{\epsilon}_1 \in \{0, 1\}^s \setminus \{0\}$, and so we can introduce the term

$$\prod_{\underline{\epsilon}_0, \underline{\epsilon}_1 \in \{0, 1\}^s \setminus \{0\}} \mathbf{1}_{B+B-B-h_0}(x + \underline{\epsilon}_0 \cdot \underline{k}_0 + \underline{\epsilon}_1 \cdot \underline{k}_1) \quad (5.38)$$

into (5.37) without changing the value of the expression. The purpose of inserting this term is to ensure that if we view (5.37) as an inner product along $K_{1i_0}, \dots, K_{si_0}, K_{1i_1}, \dots, K_{si_1}$, then the functions evaluated at $x + \underline{\epsilon}_0 \cdot \underline{k}_0 + \underline{\epsilon}_1 \cdot \underline{k}_1$ with $\underline{\epsilon}_0, \underline{\epsilon}_1 \neq \underline{0}$ have support of size bounded by $|\text{supp}(B + B - B)|$, and hence by $A|B|$ by our assumption.

We now apply the Gowers-Cauchy-Schwarz inequality to (5.37) twisted by the product (5.38) to obtain

$$\mathbb{E}_{i_0, i_1 \in I} \sum_x \sum_{\substack{\underline{k}_0 \in K_{i_0}, \\ \underline{k}_1 \in K_{i_1}}} \mu_{K_{i_0}}(\underline{k}_0) \mu_{K_{i_1}}(\underline{k}_1) \mathbb{E}_{\substack{h_0 \in H_{i_0}, \\ h_1 \in H_{i_1}}} \Delta_{\underline{k}_0, \underline{k}_1}(f(x + h_0) \overline{f(x + h_1)}) \gg_{A,s} \delta^{2^{2s+1}} |B|,$$

or equivalently

$$\mathbb{E}_{i_0, i_1 \in I} \sum_x \mathbb{E}_{\substack{\underline{k}_0, \underline{k}'_0 \in K_{i_0}, \\ \underline{k}_1, \underline{k}'_1 \in K_{i_1}}} \mathbb{E}_{\substack{h_0 \in H_{i_0}, \\ h_1 \in H_{i_1}}} \Delta'_{(\underline{k}_0, \underline{k}'_0), (\underline{k}_1, \underline{k}'_1)}(f(x) \overline{f(x + h_1 - h_0)}) \gg_{A,s} \delta^{2^{2s+1}} |B|$$

after we expand the kernels $\mu_{K_{i_0}}, \mu_{K_{i_1}}$ and change variables $x \mapsto x - h_0$. Using the Cauchy-Schwarz

inequality to double the variables h_0, h_1 yields

$$\mathbb{E}_{i_0, i_1 \in I} \sum_x \mathbb{E}_{\substack{\underline{k}_0, \underline{k}'_0 \in K_{i_0}, \\ \underline{k}_1, \underline{k}'_1 \in K_{i_1}}} \mathbb{E}_{\substack{h_0, h'_0 \in H_{i_0}, \\ h_1, h'_1 \in H_{i_1}}} \Delta'_{(\underline{k}_0, \underline{k}'_0), (\underline{k}_1, \underline{k}'_1), (h_1 - h_0, h'_1 - h'_0)} f(x) \gg_{A,s} \delta^{2^{2s+2}} |B|.$$

This expression is now an average (indexed by i_0, i_1) of box norms along $K_{i_0}, K_{i_1}, H_{i_1} - H_{i_0}$. The inductive formula for box norms gives the claimed bound

$$\mathbb{E}_{i_0, i_1 \in I} \|f\|_{K_{1i_0}, \dots, K_{si_0}, K_{1i_1}, \dots, K_{si_1}, H_{i_1} - H_{i_0}}^{2^{2s+1}} \gg_{A,s} \delta^{2^{2s+2}} |B|.$$

We reiterate our convention that $H_{i_1} - H_{i_0}$ denotes the Minkowski sum counted with multiplicity and in particular is itself a multiset. $\square$

By repeatedly applying Lemma 5.27, we derive the following concatenation result, whose proof follows closely that of [68, Proposition 2.3].

Proposition 5.28 (Concatenation of box norms, version I). *Let $A, \delta > 0$, let $s \in \mathbb{N}$, let G be a finitely generated abelian group, let I be a finite indexing set, and let $H_{1i}, \dots, H_{si} \subseteq G$ be finite multisets for each $i \in I$. Suppose that $B \subseteq G$ is a set satisfying*

$$|\text{supp}(B + B - B)| \leq A|B|.$$

Then for all 1-bounded functions f supported on B , the bound

$$\mathbb{E}_{i \in I} \|f\|_{H_{1i}, \dots, H_{si}}^{2^s} \geq \delta |B| \tag{5.39}$$

implies that

$$\mathbb{E}_{\substack{i_{\underline{\epsilon}} \in I, \\ \underline{\epsilon} \in \{0,1\}^s}} \|f\|^{2^d} \{H_{ji_{\underline{\epsilon}}} - H_{ji_{\underline{\epsilon}'}} : j \in [s], \underline{\epsilon}, \underline{\epsilon}' \in \{0,1\}^s \\ \text{with } (\epsilon_1, \dots, \epsilon_{s-j}) = (\epsilon'_1, \dots, \epsilon'_{s-j}), \epsilon_{s+1-j} = 1, \epsilon'_{s+1-j} = 0\}} \gg_{A,s} \delta^{O_s(1)} |B|,$$

where d is the degree of the box norms on the left-hand side.

Since the indices $\underline{\epsilon}, \underline{\epsilon}'$ are always distinct, the multiset $H_{ji_{\underline{\epsilon}}} - H_{ji_{\underline{\epsilon}'}}$ is usually more uniformly distributed than both of the multisets $H_{ji_{\underline{\epsilon}}}, H_{ji_{\underline{\epsilon}'}}$.

Proof. The proof of Proposition 5.28 relies on a gradual concatenation of the “unconcatenated” multisets $H_{1i}, \dots, H_{si}$ using Lemma 5.27. We repeatedly use the inductive formula for box norms in order to reinterpret

the average in such a way that successive applications of Lemma 5.27 concatenate the multisets $H_{1i}, \dots, H_{si}$ one by one.

Starting with (5.39), we apply Lemma 5.27 to bound

$$\mathbb{E}_{i_0, i_1 \in I} \|f\|_{H_{1i_0}, \dots, H_{(s-1)i_0}, H_{1i_1}, \dots, H_{(s-1)i_1}, H_{si_1} - H_{si_0}}^{2^{s-1}} \gg_{A,s} \delta^{O_s(1)} |B|. \quad (5.40)$$

Instead of having s unconcatenated multisets indexed by i , we now have $s-1$ unconcatenated multisets indexed by i_0 and another $s-1$ unconcatenated multisets indexed by i_1 , together with the concatenated multiset $H_{si_1} - H_{si_0}$. Even though the total number of multisets has nearly doubled, we have made progress because the number of unconcatenated multisets for each index i_0, i_1 decreased by 1.

At the next stage of the argument, we will apply Lemma 5.27 twice to concatenate $H_{(s-1)i_0}$ and then $H_{(s-1)i_1}$. As a consequence, the two indices i_0, i_1 will be replaced by four indices $i_{00}, i_{01}, i_{10}, i_{11}$, and for each of them we will have exactly $s-2$ unconcatenated multisets. We will continue in this manner: At each stage, the number of indices $i_{\underline{e}}$ will double, but the number of unconcatenated multisets with each index $i_{\underline{e}}$ will decrease by 1. Eventually, after s steps, we will be left with 2^{s-1} unconcatenated multisets $H_{1i_{\underline{e}}}$, and 2^{s-1} applications of Lemma 5.27 will allow us to concatenate them all without producing any new unconcatenated multisets. This will finish the argument.

To illustrate this general strategy, let us see in detail what happens at the second stage, i.e., after we have obtained the bound (5.40). Using the induction formula for box norms, we can express (5.40) as

$$\begin{aligned} & \mathbb{E}_{i_1 \in I} \mathbb{E}_{h_{1i_1}, h'_{1i_1} \in H_{1i_1}} \dots \mathbb{E}_{h_{(s-1)i_1}, h'_{(s-1)i_1} \in H_{(s-1)i_1}} \\ & \mathbb{E}_{i_0 \in I} \left\| \Delta'_{(h_{1i_1}, h'_{1i_1}), \dots, (h_{(s-1)i_1}, h'_{(s-1)i_1})} f \right\|_{H_{1i_0}, \dots, H_{(s-1)i_0}, H_{si_1} - H_{si_0}}^{2^s} \gg_{A,s} \delta^{O_s(1)} |B|. \end{aligned}$$

For each fixed $i_1, h_{1i_1}, h'_{1i_1}, \dots, h_{(s-1)i_1}, h'_{(s-1)i_1}$, we apply Lemma 5.27 separately to each average over i_0 , obtaining

$$\begin{aligned} & \mathbb{E}_{i_1 \in I} \mathbb{E}_{h_{1i_1}, h'_{1i_1} \in H_{1i_1}} \dots \mathbb{E}_{h_{(s-1)i_1}, h'_{(s-1)i_1} \in H_{(s-1)i_1}} \mathbb{E}_{i_{00}, i_{01} \in I} \\ & \left\| \Delta'_{(h_{1i_1}, h'_{1i_1}), \dots, (h_{(s-1)i_1}, h'_{(s-1)i_1})} f \right\|_{H_{(s-1)i_{01}} - H_{(s-1)i_{00}}, H_{si_1} - H_{si_{00}}, H_{si_1} - H_{si_{01}}}^{2^{s-1}} \gg_{A,s} \delta^{O_s(1)} |B|. \end{aligned}$$

Using the inductive formula for box norms and swapping the indices of Δ' and the box norm, we can rearrange this to get

$$\begin{aligned}
& \mathbb{E}_{i_{00}, i_{01} \in I} \mathbb{E}_{h_{1i_{00}}, h'_{1i_{00}} \in H_{1i_{00}}, h_{1i_{01}}, h'_{1i_{01}} \in H_{1i_{01}}} \dots \mathbb{E}_{h_{(s-1)i_{00}}, h'_{(s-1)i_{00}} \in H_{(s-1)i_{00}}, h_{(s-1)i_{01}}, h'_{(s-1)i_{01}} \in H_{(s-1)i_{01}}} \\
& \left\| \mathbb{E}_{i_1 \in I} \Delta' \begin{pmatrix} (h_{1i_{00}}, h'_{1i_{00}}), \dots, (h_{(s-2)i_{00}}, h'_{(s-2)i_{00}}) \\ (h_{1i_{01}}, h'_{1i_{01}}), \dots, (h_{(s-2)i_{01}}, h'_{(s-2)i_{01}}), \\ (h_{(s-1)i_{01}} - h_{(s-1)i_{00}}, h'_{(s-1)i_{01}} - h'_{(s-1)i_{00}}) \end{pmatrix} f \right\|_{\substack{H_{1i_1}, \dots, H_{(s-1)i_1}, \\ H_{si_1} - H_{si_{00}}, H_{si_1} - H_{si_{01}}}}^{2^{s+1}} \gg_{A,s} \delta^{O_s(1)} |B|.
\end{aligned}$$

By Lemma 5.27 applied separately to each average over $i_1 \in I$, we have

$$\begin{aligned}
& \mathbb{E}_{i_{00}, i_{01}, i_{10}, i_{11} \in I} \mathbb{E}_{h_{1i_{00}}, h'_{1i_{00}} \in H_{1i_{00}}, h_{1i_{01}}, h'_{1i_{01}} \in H_{1i_{01}}} \dots \mathbb{E}_{h_{(s-1)i_{00}}, h'_{(s-1)i_{00}} \in H_{(s-1)i_{00}}, h_{(s-1)i_{01}}, h'_{(s-1)i_{01}} \in H_{(s-1)i_{01}}} \\
& \left\| \Delta' \begin{pmatrix} (h_{1i_{00}}, h'_{1i_{00}}), \dots, (h_{(s-2)i_{00}}, h'_{(s-2)i_{00}}) \\ (h_{1i_{01}}, h'_{1i_{01}}), \dots, (h_{(s-2)i_{01}}, h'_{(s-2)i_{01}}), \\ (h_{(s-1)i_{01}} - h_{(s-1)i_{00}}, h'_{(s-1)i_{01}} - h'_{(s-1)i_{00}}) \end{pmatrix} f \right\|_{\substack{H_{1i_{10}}, \dots, H_{(s-2)i_{10}}, H_{1i_{11}}, \dots, H_{(s-2)i_{11}}, \\ H_{(s-1)i_{11}} - H_{(s-1)i_{10}}, H_{si_{10}} - H_{si_{00}}, H_{si_{11}} - H_{si_{00}}, \\ H_{si_{10}} - H_{si_{01}}, H_{si_{11}} - H_{si_{01}}}}^{2^{2s+1}} \gg_{A,s} \delta^{O_s(1)} |B|.
\end{aligned}$$

Another application of the inductive formula for box norms lets us write this as

$$\mathbb{E}_{i_{00}, i_{01}, i_{10}, i_{11} \in I} \|f\|^{2^{4s-2}}_{\substack{H_{1i_{00}}, \dots, H_{(s-2)i_{00}}, H_{1i_{01}}, \dots, H_{(s-2)i_{01}}, \\ H_{1i_{10}}, \dots, H_{(s-2)i_{10}}, H_{1i_{11}}, \dots, H_{(s-2)i_{11}}, \\ H_{(s-1)i_{01}} - H_{(s-1)i_{00}}, H_{(s-1)i_{11}} - H_{(s-1)i_{10}}, \\ H_{si_{10}} - H_{si_{00}}, H_{si_{11}} - H_{si_{00}}, \\ H_{si_{10}} - H_{si_{01}}, H_{si_{11}} - H_{si_{01}}}} \gg_{A,s} \delta^{O_s(1)} |B|,$$

which can be expressed more compactly as

$$\mathbb{E}_{i_{00}, i_{01}, i_{10}, i_{11} \in I} \|f\|^{2^{4s-2}}_{\substack{\{H_{ji_{\underline{\epsilon}}}: \underline{\epsilon} \in \{0,1\}^2, j=1, \dots, s-2\}, \\ \{H_{(s-1)i_{\underline{\epsilon}}} - H_{(s-1)i_{\underline{\epsilon}'}}: \underline{\epsilon}, \underline{\epsilon}' \in \{0,1\}^2 \text{ with } \epsilon_1 = \epsilon'_1, \epsilon_2 = 1, \epsilon'_2 = 0\}, \\ \{H_{si_{\underline{\epsilon}}} - H_{si_{\underline{\epsilon}'}}: \underline{\epsilon}, \underline{\epsilon}' \in \{0,1\}^2 \text{ with } \epsilon_1 = 1, \epsilon'_1 = 0\}}} \gg_{A,s} \delta^{O_s(1)} |B|.$$

We have thus successfully concatenated all of the multisets $H_{(s-1)i_{\underline{\epsilon}}}$ and $H_{si_{\underline{\epsilon}}}$.

At the next stage, we concatenate the multisets $H_{(s-2)i_{\underline{\epsilon}}}$. Applying Lemma 5.27 and the induction formula for box norms four times (as in the previous stage), each time to an average over $i_{00}, i_{01}, i_{10}, i_{11}$ (respectively), we arrive at the inequality

$$\mathbb{E}_{\substack{i_{\underline{\epsilon}} \in I, \\ \underline{\epsilon} \in \{0,1\}^3}} \|f\|^{2^{d'}}_{\substack{\{H_{ji_{\underline{\epsilon}}}: \underline{\epsilon} \in \{0,1\}^3, j=1, \dots, s-3\}, \\ \{H_{ji_{\underline{\epsilon}}} - H_{ji_{\underline{\epsilon}'}}: j=s-2, s-1, s, \underline{\epsilon}, \underline{\epsilon}' \in \{0,1\}^3 \\ \text{with } (\epsilon_1, \dots, \epsilon_{s-j}) = (\epsilon'_1, \dots, \epsilon'_{s-j}), \epsilon_{s+1-j} = 1, \epsilon'_{s+1-j} = 0\}}} \gg_{A,s} \delta^{O_s(1)} |B|,$$

where d' is the degree of the box norms on the left-hand side. This time, we have successfully concatenated the multisets $H_{(s-2)i_{\underline{\epsilon}}}$. At the next step, eight applications of Lemma 5.27 and the induction formula for box norms allow us to concatenate multisets $H_{(s-3)i_{\underline{\epsilon}}}$. Continuing the argument in this manner, we arrive at the conclusion of the proposition after a total of

$$1 + 2 + 2^2 + \cdots + 2^{s-1} = 2^s - 1$$

applications of Lemma 5.27 and the induction formula for box norms. $\square$

We now present two corollaries of Proposition 5.28 which are better suited for applications in $\mathbb{Z}^D$. In what follows, we order the set $\{0, 1\}^s$ lexicographically. Identifying $\{0, 1\}^s$ with $[2^s]$, we obtain the following corollary.

Corollary 5.29 (Concatenation of box norms, version II). *Let $A, D, s, N \in \mathbb{N}$, let I be a finite indexing set, and let $H_{ji} \subseteq [\pm AN]^D$ be finite nonempty multisets for each $i \in I$ and $j \in [s]$. Then for all 1-bounded functions $f : \mathbb{Z}^D \rightarrow \mathbb{C}$ supported on $[N]^D$, the bound*

$$\mathbb{E}_{i \in I} \|f\|_{H_{1i}, \dots, H_{si}}^{2^s} \geq \delta N^D$$

implies that

$$\mathbb{E}_{k_1, \dots, k_t \in I} \|f\|_{\{H_{jk_{i_1}} - H_{jk_{i_2}} : j \in [s], 0 < i_1 < i_2 \leq t\}}^{2^d} \gg_{A, D, s} \delta^{O_s(1)} N^D, \quad (5.41)$$

where $t = 2^s$ and d is the degree of the box norms on the left-hand side.

Proof. Proposition 5.28 immediately implies that

$$\mathbb{E}_{\substack{i_{\underline{\epsilon}} \in I, \\ \underline{\epsilon} \in \{0, 1\}^s}} \|f\|_{\substack{\{H_{ji_{\underline{\epsilon}}} - H_{ji_{\underline{\epsilon}'}} : j \in [s], \underline{\epsilon}, \underline{\epsilon}' \in \{0, 1\}^s \\ \text{with } (\epsilon_1, \dots, \epsilon_{s-j}) = (\epsilon'_1, \dots, \epsilon'_{s-j}), \epsilon_{s+1-j} = 1, \epsilon'_{s+1-j} = 0\}}}^{2^{d'}} \gg_{A, D, s} \delta^{O_s(1)} N^D,$$

where d' is the degree of the box norms on the left-hand side. Ordering $\{0, 1\}^s$ lexicographically and identifying it with $[2^s] = [t]$ in an order-preserving way, we obtain a subset $\mathcal{K}_0 \subseteq [t] \times [t]$ such that

$$\mathbb{E}_{k_1, \dots, k_t \in I} \|f\|_{\{H_{jk_{i_1}} - H_{jk_{i_2}} : j \in [s], (i_1, i_2) \in \mathcal{K}_0\}}^{2^{d'}} \gg_{A, D, s} \delta^{O_s(1)} N^D.$$

Note that $\mathcal{K}_0 \subseteq \{(i_1, i_2) \in [t] \times [t] : i_1 < i_2\}$. Our next goal is to replace $\mathcal{K}_0$ by the full set $\mathcal{K} = \{(i_1, i_2) \in$

$[t] \times [t] : i_1 < i_2\}$. By the assumption on the multisets H_{ji} , we have

$$[N]^D - (H_{jk_1} - H_{jk_2}) \subseteq (-2AN, (2A+1)N]^D,$$

and so

$$|\text{supp}([N]^D - (H_{jk_1} - H_{jk_2}))| \leq (4A+1)^D N^D \ll_{A,D} N^D.$$

The corollary follows from this observation and repeated applications of Lemma 5.15(3), the monotonicity property for box norms. $\square$

The next corollary will let us take the directions in the concatenated box norms to be arbitrarily long sums of the original directions rather than just double sums (as in Corollary 5.29).

Corollary 5.30 (Iterated concatenation of box norms). *Let $A, D, s, N \in \mathbb{N}$, and let ℓ be a power of 2. Let I be a finite indexing set, and let $H_{ji} \subseteq [\pm AN]^D$ be symmetric multisets for each $i \in I$ and $j \in [s]$. Then for all 1-bounded functions $f : \mathbb{Z}^D \rightarrow \mathbb{C}$ supported on $[N]^D$, the bound*

$$\mathbb{E}_{i \in I} \|f\|_{H_{1i}, \dots, H_{si}}^{2^s} \geq \delta N^D$$

implies that

$$\mathbb{E}_{k_1, \dots, k_t \in I} \|f\|_{\{H_{jk_{i_1}} + \dots + H_{jk_{i_\ell}} : j \in [s], 1 \leq i_1 < \dots < i_\ell \leq t\}}^{2^d} \gg_{A,D,\ell,s} \delta^{O_{\ell,s}(1)} N^D, \quad (5.42)$$

for some $t = O_{\ell,s}(1)$, where d is the degree of the box norms on the left-hand side.

It will be useful later that the indices $i_1, \dots, i_\ell$ in the sum

$$H_{jk_{i_1}} + \dots + H_{jk_{i_\ell}}$$

in (5.42) are all distinct. The assumption of the symmetry of the multisets H_{ji} (meaning that $H_{ji} = -H_{ji}$) in the corollary is not necessary, and we include it only so that we do not have to keep track of the minus signs in the expressions $H_{jk_{i_1}} + \dots + H_{jk_{i_\ell}}$.

Proof. The proof proceeds by induction on ℓ , which we recall ranges over powers of 2. For $\ell = 1$, the claim is trivial, and for $\ell = 2$, the lower bound

$$\mathbb{E}_{k_1, \dots, k_{t_1} \in I} \|f\|_{\{H_{jk_{i_1}} + H_{jk_{i_2}} : j \in [s], 0 < i_1 < i_2 \leq t_1\}}^{2^{d_1}} \gg_{A,D,s} \delta^{O_s(1)} N^D, \quad (5.43)$$

where $t_1 = 2^s$ and $d_1 = \binom{t_1}{2}$ is the degree of the box norms appearing, follows from Corollary 5.29 and the fact that the sets H_{ji} are symmetric. We will now prove the case $\ell = 4$, and which will make the induction scheme for larger ℓ clear.

To prove the result for $\ell = 2$, we apply the case $\ell = 1$ to (5.43), taking $(k_1, \dots, k_{t_1}) \in I^{t_1}$ in place of $i \in I$. This gives

$$\mathbb{E}_{\substack{k_{il} \in I: \\ i \in [t_1], l \in [t'_1]}} \|f\|_{\{H_{jk_{i_1}l_1} + H_{jk_{i_2}l_1} + H_{jk_{i_1}l_2} + H_{jk_{i_2}l_2} : j \in [s], 0 < i_1 < i_2 \leq t_1, 0 < l_1 < l_2 \leq t'_1\}}^{2^{d'_1}} \gg_{A,D,s} \delta^{O_s(1)} N^D, \quad (5.44)$$

where $t'_1 = 2^{d_1}$ and d'_1 is the degree of the box norms appearing. Reindexing $(i, l) \in [t_1] \times [t'_1]$ as $i \in [t_2]$ and using the monotonicity property and the assumption on H_{ji} , we argue as in the proof of Corollary 5.29 that

$$\mathbb{E}_{k_1, \dots, k_{t_2} \in I} \|f\|_{\{H_{jk_{i_1}} + \dots + H_{jk_{i_4}} : j \in [s], 0 < i_1 < \dots < i_4 \leq t_2\}}^{2^{d_2}} \gg_{A,D,s} \delta^{O_s(1)} N^D,$$

where d_2 is the degree of the box norms appearing. We continue in the same manner for larger ℓ . $\square$

5.7 Equidistribution estimates for concatenation

The purpose of this section is to obtain upper bounds on the number of solutions to systems of multilinear equations that appear in connection with the manipulations in Section 5.8. These results will culminate in Proposition 5.35, which is precisely the ingredient needed to complete the proof of Theorem 5.10.

The starting point for our argument is the following simple lemma.

Lemma 5.31 (Bound on solutions to linear congruences). *Let $\ell \geq 2$, let $h_1, \dots, h_\ell$ be positive integers, and let $M \in \mathbb{N}$. Then*

$$\max_{t \in \mathbb{Z}} \sum_{m_1, \dots, m_{\ell-1} \in [\pm M]} \mathbf{1}_{h_\ell \mathbb{Z} + t}(h_1 m_1 + \dots + h_{\ell-1} m_{\ell-1}) \ll_\ell M^{\ell-1} \frac{\gcd(h_1, \dots, h_\ell)}{h_\ell} + M^{\ell-2}.$$

Proof. First, we observe that if $(m'_1, \dots, m'_{\ell-1}) \in [\pm M]^{\ell-1}$ is a solution to the congruence

$$h_1 m_1 + \dots + h_{\ell-1} m_{\ell-1} \equiv t \pmod{h_\ell},$$

then every other solution in $[\pm M]^{\ell-1}$ takes the form $(m_1, \dots, m_{\ell-1}) = (m'_1 + m''_1, \dots, m'_\ell + m''_\ell)$, where

$(m_1'', \dots, m_\ell'')$ is a solution to the homogeneous congruence

$$h_1 m_1 + \dots + h_{\ell-1} m_{\ell-1} \equiv 0 \pmod{h_\ell}$$

and $m_i'' \in [\pm M] - m_i' \subseteq [\pm 2M]$. Therefore, it suffices to prove the $t = 0$ case of the lemma (at the cost of changing the implicit constant in the conclusion of the lemma).

We proceed by induction on ℓ . For the base case $\ell = 2$, we have

$$h_2 | h_1 m_1 \iff \frac{h_2}{\gcd(h_1, h_2)} | m_1,$$

and hence

$$\sum_{m_1 \in [\pm M]} \mathbf{1}_{h_2 \mathbb{Z}}(h_1 m_1) \ll M \frac{\gcd(h_1, h_2)}{h_2} + 1$$

(where $+1$ is the boundary term).

For the induction step, write

$$\begin{aligned} \sum_{m_1, \dots, m_{\ell-1} \in [\pm M]} \mathbf{1}_{h_\ell \mathbb{Z}}(h_1 m_1 + \dots + h_{\ell-1} m_{\ell-1}) = \\ \sum_{m_1, \dots, m_{\ell-2} \in [\pm M]} \sum_{m_{\ell-1} \in [\pm M]} \mathbf{1}_{h_\ell \mathbb{Z} - (h_1 m_1 + \dots + h_{\ell-2} m_{\ell-2})}(h_{\ell-1} m_{\ell-1}). \end{aligned}$$

The crucial observation is that the inner summand vanishes if $h_1 m_1 + \dots + h_{\ell-2} m_{\ell-2}$ is not a multiple of $g := \gcd(h_{\ell-1}, h_\ell)$. Now suppose that g does divide $h_1 m_1 + \dots + h_{\ell-2} m_{\ell-2}$; the induction hypothesis tells us that this occurs for at most

$$O_\ell \left(M^{\ell-2} \frac{\gcd(h_1, \dots, h_\ell)}{g} + M^{\ell-3} \right)$$

choices of $m_1, \dots, m_{\ell-2}$, where we have used the identity

$$\gcd(h_1, \dots, h_{\ell-2}, \gcd(h_{\ell-1}, h_\ell)) = \gcd(h_1, \dots, h_\ell).$$

For each such choice of $m_1, \dots, m_{\ell-2}$, we have

$$\mathbf{1}_{h_\ell \mathbb{Z} - (h_1 m_1 + \dots + h_{\ell-2} m_{\ell-2})}(h_{\ell-1} m_{\ell-1}) = \mathbf{1}_{(h_\ell/g) \mathbb{Z} - (h_1 m_1 + \dots + h_{\ell-2} m_{\ell-2})/g}((h_{\ell-1}/g) m_{\ell-1}).$$

Now the $\ell = 2$ case of the proposition tells us that the sum over $m_{\ell-1}$ has size

$$O\left(M \frac{\gcd(h_{\ell-1}/g, h_{\ell}/g)}{h_{\ell}/g} + 1\right) = O\left(M \frac{g}{h_{\ell}} + 1\right)$$

since $\gcd(h_{\ell-1}/g, h_{\ell}/g) = 1$ by the definition of g . Combining these two observations and using $\gcd(h_1, \dots, h_{\ell}) \leq g \leq h_{\ell}$, we bound

$$\sum_{m_1, \dots, m_{\ell-1} \in [\pm M]} \mathbf{1}_{h_{\ell}\mathbb{Z}}(h_1 m_1 + \dots + h_{\ell-1} m_{\ell-1}) \ll_{\ell} M^{\ell-1} \frac{\gcd(h_1, \dots, h_{\ell})}{h_{\ell}} + M^{\ell-2},$$

as desired. $\square$

Corollary 5.32 (Bound on solutions to linear equations). *Let $\ell \geq 2$, let $h_1, \dots, h_{\ell}$ be positive integers, and let $M \in \mathbb{N}$. Then*

$$\max_{t \in \mathbb{Z}} \sum_{m_1, \dots, m_{\ell} \in [\pm M]} \mathbf{1}(h_1 m_1 + \dots + h_{\ell} m_{\ell} = t) \ll_{\ell} M^{\ell-1} \frac{\gcd(h_1, \dots, h_{\ell})}{h_{\ell}} + M^{\ell-2}.$$

We will later apply this corollary with

$$\max_{i \in [\ell]} |h_i| \leq H \leq M, \quad \gcd(h_1, \dots, h_{\ell}) \leq \eta^{-1}, \quad \text{and} \quad h_{\ell} \geq \eta H$$

(for small positive η), in which case the right-hand side is $\ll_{\ell} \eta^{-2} M^{\ell-1} H^{-1}$.

Proof. For each choice of $m_1, \dots, m_{\ell-1}$, there is at most one value of m_{ℓ} satisfying $h_1 m_1 + \dots + h_{\ell} m_{\ell} = t$, so we can bound

$$\sum_{m_1, \dots, m_{\ell} \in [\pm M]} \mathbf{1}_{h_1 m_1 + \dots + h_{\ell} m_{\ell} = t} \leq \sum_{m_1, \dots, m_{\ell-1} \in [\pm M]} \mathbf{1}_{h_{\ell}\mathbb{Z}+t}(h_1 m_1 + \dots + h_{\ell-1} m_{\ell-1});$$

here we have an inequality instead of an equality because on the right-hand side we have thrown out the constraint $m_{\ell} \in [\pm M]$. The result now follows from Lemma 5.31. $\square$

In order to state our more general results, we require a bit of notation. Define

$$\mathcal{K}_{t,\ell,r} := \{(k_{li})_{(l,i) \in [r] \times [\ell]} \in [t]^{r\ell} : 1 \leq k_{l1} < \dots < k_{l\ell} \leq t \text{ for all } l \in [r]\} \quad (5.45)$$

to be the collection of admissible tuples of indices. For $l \in [r]$ and $\eta > 0$, also define

$$\begin{aligned}
\mathcal{H}_{l,\eta} &:= \{(h_{lk_1 \dots k_l})_{k_1, \dots, k_l \in [t]} \in [\pm H]^{t^l} : |h_{lk_1 \dots k_l} - h_{lk'_1 \dots k'_l}| \geq \eta H, \\
&\quad \gcd(h_{lk_1 \dots k_l} - h_{lk'_1 \dots k'_l}, h_{lk'_1 \dots k'_l} - h_{lk''_1 \dots k''_l}) \leq \eta^{-1} \\
&\quad \text{for distinct } (k_1, \dots, k_l), (k'_1, \dots, k'_l), (k''_1, \dots, k''_l) \in [t]^l\} \quad (5.46)
\end{aligned}$$

to be the set of tuples of h 's whose pairwise differences are not too small and whose greatest common divisors are not too big. The following lemma, which lets us establish anti-concentration results on many pairs of “scales” simultaneously, will turn out to be the key component of the proof of Proposition 5.35 below.

Lemma 5.33. *Let $\eta > 0$ and $H, M, l, \ell, r, s, t \in \mathbb{N}$ satisfy $H \leq M$, $3 \leq \ell \leq t$, and $l \leq r$. Let $\mathcal{K} = \mathcal{K}_{t,\ell,r}$ be as in (5.45) and $\mathcal{H} = \mathcal{H}_{l,\eta}$ be as in (5.46). Then*

$$\begin{aligned}
&\mathbb{E}_{\substack{h_{lk_1 \dots k_l} \in [\pm H]: \\ k_1, \dots, k_l \in [t]}} \mathbf{1}_{\mathcal{H}}((h_{lk_1 \dots k_l})_{k_1, \dots, k_l}) \cdot \prod_{i_1, \dots, i_{l-1} \in [\ell]} \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{u_{l+1}, \dots, u_r \in \{0,1\}} \\
&\quad \max_{n_0, n_1 \in \mathbb{Z}} \mathbb{E}_{\substack{|m_{i_l}| \leq H^{u_{l+1} + \dots + u_r} M: \\ i_l \in [\ell]}} \mathbf{1}\left(\sum_{i_l \in [\ell]} m_{i_l} = n_0\right) \cdot \mathbf{1}\left(\sum_{i_l \in [\ell]} h_{lk_{1i_1} \dots k_{li_{i_l}}} m_{i_l} = n_1\right) \\
&\quad \ll_{l,\ell,r,s,t} \eta^{-2^{r-l+1} s^{\ell-1} |\mathcal{K}|} M^{-2^{r-l+1} s^{\ell-1} |\mathcal{K}|} H^{-(r-l+1) 2^{r-l} s^{\ell-1} |\mathcal{K}|}.
\end{aligned}$$

Example 5.34. We will illustrate the content of this lemma and the method of proof for $r = 2$, $l = 1$, $t = 4$, $\ell = 3$, $s = 1$ before supplying the proof in the general case. Letting

$$\mathcal{K}' = \{(1, 2, 3), (1, 2, 4), (1, 3, 4), (2, 3, 4)\} \quad \text{so that} \quad \mathcal{K} = \mathcal{K}' \times \mathcal{K}'$$

and $|\mathcal{K}| = 16$, and renaming $h_{1k_{1i_1}}$ as simply h_{k_i} , we can express the conclusion of the lemma as

$$\begin{aligned}
&\mathbb{E}_{h_1, h_2, h_3, h_4 \in [\pm H]} \mathbf{1}_{\mathcal{H}}(h_1, h_2, h_3, h_4) \cdot \left(\prod_{\underline{k} \in \mathcal{K}'} \prod_{u_2 \in \{0,1\}} \max_{n_0, n_1 \in \mathbb{Z}} \right. \\
&\quad \left. \mathbb{E}_{|m_1|, |m_2|, |m_3| \leq H^{u_2} M} \mathbf{1}(m_1 + m_2 + m_3 = n_0) \cdot \mathbf{1}(h_{k_1} m_1 + h_{k_2} m_2 + h_{k_3} m_3 = n_1) \right)^4 \\
&\quad \ll \eta^{-64} M^{-64} H^{-64},
\end{aligned}$$

where the fourth power comes from the fact that we are splitting

$$\prod_{(k_{1i_1}, k_{2i_2})_{i_1, i_2 \in [3]} \in \mathcal{K}} = \prod_{(k_{1i_1})_{i_1 \in [3]} \in \mathcal{K}'} \prod_{(k_{2i_2})_{i_2 \in [3]} \in \mathcal{K}'}$$

and the terms in the product are independent of k_{2i_2} (so the product over $\underline{k} \in \mathcal{K}'$ above corresponds to the

product over $(k_{1i_1})_{i_1 \in [3]} \in \mathcal{K}'$ while the fourth power corresponds to the product over $(k_{2i_2})_{i_2 \in [3]} \in \mathcal{K}'$. Our normalized count can then be expanded as

$$\mathbb{E}_{h_1, h_2, h_3, h_4 \in [\pm H]} \mathbf{1}_{\mathcal{H}}(h_1, h_2, h_3, h_4) \cdot \left(\prod_{\underline{k} \in \mathcal{K}'} \max_{n_{00}, n_{01}, n_{10}, n_{11} \in \mathbb{Z}} \right. \\ \left. \mathbb{E}_{\substack{|m_{01}|, |m_{02}|, |m_{03}| \leq M, \\ |m_{11}|, |m_{12}|, |m_{13}| \leq HM}} \mathbf{1} \left(\begin{array}{l} m_{01} + m_{02} + m_{03} = n_{00}, \\ m_{11} + m_{12} + m_{13} = n_{01} \end{array} \right) \cdot \mathbf{1} \left(\begin{array}{l} h_{k_1} m_{01} + h_{k_2} m_{02} + h_{k_3} m_{03} = n_{10}, \\ h_{k_1} m_{11} + h_{k_2} m_{12} + h_{k_3} m_{13} = n_{11} \end{array} \right) \right)^4.$$

Thus, for each $(h_1, h_2, h_3, h_4) \in \mathcal{H}$ and $\underline{k} \in \mathcal{K}'$, we are left with bounding the number of solutions to the system

$$\begin{aligned} m_{01} + m_{02} + m_{03} &= n_{00} & \text{and} & & h_{k_1} m_{01} + h_{k_2} m_{02} + h_{k_3} m_{03} &= n_{10}, \\ m_{11} + m_{12} + m_{13} &= n_{01} & \text{and} & & h_{k_1} m_{11} + h_{k_2} m_{12} + h_{k_3} m_{13} &= n_{11}. \end{aligned}$$

Using the equations on the left to solve for m_{03} and m_{13} and substituting the resulting expressions into the equations on the right reduces the system to

$$\begin{aligned} (h_{k_1} - h_{k_3})m_{01} + (h_{k_2} - h_{k_3})m_{02} &= n_0 \\ (h_{k_1} - h_{k_3})m_{11} + (h_{k_2} - h_{k_3})m_{12} &= n_1 \end{aligned}$$

for $n_0 = n_{10} - h_{k_3} n_{00}$ and $n_1 = n_{11} - h_{k_3} n_{01}$, and each of these two equations can then be handled separately using Corollary 5.32 and the assumptions on $(h_1, h_2, h_3, h_4) \in \mathcal{H}_1$.

The exponent 64 comes from the fact that the numbers $n_{00}, n_{01}, n_{10}, n_{11}$ lie in the intervals

$$n_{00} \in [\pm 3M], \quad n_{01}, n_{10} \in [\pm 3HM], \quad n_{11} \in [\pm 3H^2M]$$

and hence there are $\asymp H^4 M^4$ possibilities for the quadruple $(n_{00}, n_{01}, n_{10}, n_{11})$. Because of the assumption on (h_1, h_2, h_3, h_4) , we can use Corollary 5.32 to ensure that the quadruple $(n_{00}, n_{01}, n_{10}, n_{11})$ is roughly equidistributed in the appropriate box as the m 's range; in particular, the probability of hitting each individual quadruple $(n_{00}, n_{01}, n_{10}, n_{11})$ is

$$O((M^{-2}\eta^{-1}/(\eta H)) \cdot ((HM)^{-2}\eta^{-1}/(\eta H))) = O(\eta^{-4}H^{-4}M^{-4}).$$

Raising this quantity to the power 4 (because $|\mathcal{K}'| = 4$), we see that the expression in parentheses above is $O(\eta^{-16}H^{-16}M^{-16})$, and raising this to the power 4 gives the desired exponent 64.

Proof of Lemma 5.33. We let all implicit constants depend on l, ℓ, r, s, t . Using the equation $\sum_{i_l \in [\ell]} m_{i_l} = n_0$ we can substitute $m_\ell = n_0 - \sum_{i_l \in [\ell-1]} m_{i_l}$ into the equation

$$\sum_{i_l \in [\ell]} h_{lk_{1i_1} \dots k_{li_l}} m_{i_l} = n_1,$$

so that our normalized count is bounded by

$$\begin{aligned} & \mathbb{E}_{\substack{h_{lk_1 \dots k_l} \in [\pm H]: \\ k_1, \dots, k_l \in [t]}} \mathbf{1}_{\mathcal{H}}((h_{lk_1 \dots k_l})_{k_1, \dots, k_l}) \cdot \prod_{i_1, \dots, i_{l-1} \in [\ell]} \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{u_{l+1}, \dots, u_r \in \{0,1\}} H^{-(u_{l+1} + \dots + u_r)} M^{-1} \\ & \max_{n \in \mathbb{Z}} \mathbb{E}_{\substack{|m_{i_l}| \leq H^{u_{l+1} + \dots + u_r} M: \\ i_l \in [\ell-1]}} \mathbf{1} \left(\sum_{i_l \in [\ell-1]} (h_{lk_{1i_1} \dots k_{li_l}} - h_{lk_{1i_1} \dots k_{(l-1)i_{l-1}} k_{l\ell}}) m_{i_l} = n \right), \end{aligned}$$

where we have set $n = n_1 - n_0 h_{lk_{1i_1} \dots k_{(l-1)i_{l-1}} k_{l\ell}}$. By Corollary 5.32, the inner max is in turn bounded by

$$H^{-(u_{l+1} + \dots + u_r)} M^{-1} \frac{\gcd(h_{lk_{1i_1} \dots k_{(l-1)i_{l-1}} k_{li_l}} - h_{lk_{1i_1} \dots k_{(l-1)i_{l-1}} k_{l\ell}} : i_l \in [\ell-1])}{\max(h_{lk_{1i_1} \dots k_{(l-1)i_{l-1}} k_{li_l}} - h_{lk_{1i_1} \dots k_{(l-1)i_{l-1}} k_{l\ell}} : i_l \in [\ell-1])}$$

as long as $\ell \geq 3$ and $H \leq M$. Due to our assumption on the set $\mathcal{H}$, the gcd in the numerator is at most η^{-1} and the maximum in the denominator is at least ηH , so we can bound our normalized count by

$$\begin{aligned} & \ll \mathbb{E}_{\substack{h_{lk_1 \dots k_l} \in [\pm H]: \\ k_1, \dots, k_l \in [t]}} \prod_{i_1, \dots, i_{l-1} \in [\ell]} \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{u_{l+1}, \dots, u_r \in \{0,1\}} \eta^{-2} M^{-2} H^{-1-2(u_{l+1} + \dots + u_r)} \\ & \ll \eta^{-2^{r-l+1} s \ell^{l-1} |\mathcal{K}|} M^{-2^{r-l+1} s \ell^{l-1} |\mathcal{K}|} H^{-(r-l+1) 2^{r-l} s \ell^{l-1} |\mathcal{K}|}. \end{aligned}$$

□

The next result, in which we establish simultaneous anti-concentration for expressions on many scales, is the main estimate of this section. It gives an upper bound on the number of solutions to the systems of multilinear equations that naturally appear while concatenating box norms in the proof of Theorem 5.10.

Proposition 5.35 (Solutions to systems of multilinear equations). *Let $\eta > 0$, let $\ell, r, s, t \in \mathbb{N}$ with $3 \leq \ell \leq t$, and let $H, M \in \mathbb{N}$ with $H \leq M$. Let $\mathcal{K} = \mathcal{K}_{t,\ell,r}$ be as in (5.45), and for each $l \in [r]$, let $\mathcal{H}_l := \mathcal{H}_{l,\eta}$ be as in (5.46). Then*

$$\begin{aligned}
& \max_{\substack{n_{j\underline{k}\underline{u}} \in \mathbb{Z}: j \in [s], \\ \underline{k} \in \mathcal{K}, \underline{u} \in \{0,1\}^r}} \mathbb{E}_{\substack{m_{j\underline{k}i_1 \dots i_r} \in [\pm M]: \\ j \in [s], \underline{k} \in \mathcal{K}, i_1, \dots, i_r \in [\ell]}} \mathbb{E}_{\substack{h_{lk_1 \dots k_l} \in [\pm H]: \\ k_1, \dots, k_r \in [t], l \in [r]}} \left(\prod_{l \in [r]} \mathbf{1}_{\mathcal{H}_l}((h_{lk_1 \dots k_l})_{k_1, \dots, k_l \in [t]}) \right) \\
& \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{\underline{u} \in \{0,1\}^r} \mathbf{1} \left(\sum_{i_1, \dots, i_r \in [\ell]} h_{1k_1i_1}^{u_1} \dots h_{rk_ri_r}^{u_r} m_{j\underline{k}i_1 \dots i_r} = n_{j\underline{k}\underline{u}} \right) \\
& \ll_{\ell, r, s, t} \eta^{-O_{\ell, r, s, t}(1)} M^{-2^r s |\mathcal{K}|} H^{-r 2^{r-1} s |\mathcal{K}|}.
\end{aligned}$$

(Here, the superscript u 's are exponents, not indices.)

On a first read, the reader may wish to consider the case $t = \ell$, so that $\mathcal{K}$ consists of only one element, k_{li_i} can be replaced by k_l , and the product $\prod_{\underline{k} \in \mathcal{K}}$ disappears.

In the applications of Proposition 5.35 to Theorem 5.10, s is fixed and ℓ (the number of terms in each equation) can be taken as large as we want, whereas we have no control over t , which grows with ℓ and will typically be much larger than ℓ .

Before proving Proposition 5.35, we illustrate the statement with several examples.

Example 5.36. If $s = 1$, i.e., we are using Proposition 5.35 to concatenate an average of degree-1 box norms, then we will always have $t = \ell$, whence $|\mathcal{K}| = 1$. Suppose that $t = \ell = 3$ and $r = 2$, i.e., we are dealing with a system of equations coming from concatenating an average

$$\mathbb{E}_{h_1, h_2 \in [\pm H]} \|f\|_{(\beta_{00} + \beta_{10}h_1 + \beta_{01}h_2 + \beta_{11}h_1h_2) \cdot [\pm M]}^2.$$

The resulting system consists of equations of the form

$$\begin{aligned}
& h_{11}^{u_1} (h_{211}^{u_2} m_{11} + h_{212}^{u_2} m_{12} + h_{213}^{u_2} m_{13}) + h_{12}^{u_1} (h_{221}^{u_2} m_{21} + h_{222}^{u_2} m_{22} + h_{223}^{u_2} m_{23}) \\
& + h_{13}^{u_1} (h_{231}^{u_2} m_{31} + h_{232}^{u_2} m_{32} + h_{233}^{u_2} m_{33}) = n_{u_1 u_2}
\end{aligned}$$

for $(u_1, u_2) \in \{0,1\}^2$. Since $n_{u_1 u_2}$ lies in $[\pm 9H^{u_1+u_2}M]$, we expect to get a cancellation on the order of

$$\ll \prod_{u_1, u_2 \in \{0,1\}} H^{-(u_1+u_2)} M^{-1} = H^{-4} M^{-4},$$

and modulo the powers of η , this is indeed what Proposition 5.35 gives.

Example 5.37. Suppose now that $r = 1$, $s = 4$, $t = 6$, and $\ell = 4$, so that $|\mathcal{K}| = 15$. Then we have a system of 30 equations, which we can group into 15 pairs of equations

$$m_{\underline{k}1} + m_{\underline{k}2} + m_{\underline{k}3} + m_{\underline{k}4} = n_{\underline{k}0} \quad \text{and} \quad m_{\underline{k}1}h_{k_1} + m_{\underline{k}2}h_{k_2} + m_{\underline{k}3}h_{k_3} + m_{\underline{k}4}h_{k_4} = n_{\underline{k}1}$$

corresponding to the 15 tuples

$$\underline{k} \in \mathcal{K} = \{(k_1, k_2, k_3, k_4) : 1 \leq k_1 < k_2 < k_3 < k_4 \leq 6\}.$$

Each $n_{\underline{k}0}$, $n_{\underline{k}1}$ gives cancellation of order $\asymp M^{-1}$, $\asymp H^{-1}M^{-1}$ (respectively), so the total cancellation will be of order $\asymp H^{-15}M^{-30}$ (modulo powers of η , which we ignore here for the simplicity of exposition). We remark that this example does not figure in the proof of Theorem 5.10 (but nonetheless is covered by Proposition 5.35).

Let us see how the heuristics from the preceding two examples extend to the general case. For each $(j, \underline{k}, \underline{u})$, we expect the value $n_{j\underline{k}\underline{u}}$ to be distributed fairly uniformly in an interval $[\pm C_{j\underline{k}\underline{u}} H^{u_1 + \dots + u_r} M]$ for some $C_{j\underline{k}\underline{u}} > 0$; the key uniform distribution property follows from Corollary 5.32 and the fact that the tuples $(h_{lk_1 \dots k_l})_{k_1, \dots, k_l}$ lie in the sets $\mathcal{H}_l$. So for each $(j, \underline{k}, \underline{u})$ we get a cancellation of the order of

$$\asymp_r (H^{-(u_1 + \dots + u_r)} M^{-1}), \quad (5.47)$$

which is (up to constants) the probability that a uniformly random element of the interval $[\pm C_{j\underline{k}\underline{u}} H^{u_1 + \dots + u_r} M]$ assumes any given value. The powers of η in the bound come from restricting the variables h_l to the admissible ranges $\mathcal{H}_l$, and using Corollary 5.32 to bound solutions for h 's in this range. The final bound then follows upon taking the product of (5.47) over all $j \in [s]$, $\underline{u} \in \{0, 1\}^r$, and $\underline{k} \in \mathcal{K}$.

Proof of Proposition 5.35. We let all implicit constants depend on ℓ, r, s, t . The proof goes by induction on r . The $r = 1$ case is the content of Lemma 5.33. The induction step is in general quite notationally complicated, so, to make it more palatable, we first present the case $r = 2$.

The case $r = 2$.

In this case, our claimed bound takes the form

$$\begin{aligned} & \max_{\substack{n_{j\underline{k}00}, n_{j\underline{k}01}, \\ n_{j\underline{k}10}, n_{j\underline{k}11} \in \mathbb{Z}: \\ j \in [s], \underline{k} \in \mathcal{K}}} \mathbb{E}_{\substack{m_{j\underline{k}i_1 i_2} \in [\pm M]: \\ j \in [s], \underline{k} \in \mathcal{K}, i_1, i_2 \in [\ell]}} \mathbb{E}_{\substack{h_{1k_1}, h_{2k_1 k_2} \in [\pm H]: \\ k_1, k_2 \in [t]}} \mathbf{1}_{\mathcal{H}_1}((h_{1k_1})_{k_1}) \cdot \mathbf{1}_{\mathcal{H}_2}((h_{2k_1 k_2})_{k_1, k_2}) \\ & \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{\underline{u} \in \{0, 1\}^2} \mathbf{1} \left(\sum_{i_1, i_2 \in [\ell]} h_{1k_{1i_1}}^{u_1} h_{2k_{1i_1} k_{2i_2}}^{u_2} m_{j\underline{k}i_1 i_2} = n_{j\underline{k}\underline{u}} \right) \\ & \ll \eta^{-O(1)} M^{-4s|\mathcal{K}|} H^{-4s|\mathcal{K}|}. \end{aligned}$$

For each $(j, \underline{k}) \in [s] \times \mathcal{K}$, we have the system of four equations

$$\begin{aligned} \sum_{i_1, i_2 \in [\ell]} m_{j\underline{k}i_1i_2} &= n_{j\underline{k}00} \\ \sum_{i_1, i_2 \in [\ell]} h_{1k_{i_1}} m_{j\underline{k}i_1i_2} &= n_{j\underline{k}10} \\ \sum_{i_1, i_2 \in [\ell]} h_{2k_{i_1}k_{i_2}} m_{j\underline{k}i_1i_2} &= n_{j\underline{k}01} \\ \sum_{i_1, i_2 \in [\ell]} h_{1k_{i_1}} h_{2k_{i_1}k_{i_2}} m_{j\underline{k}i_1i_2} &= n_{j\underline{k}11}. \end{aligned}$$

This system is quadratic in the h variables, and our goal is to decouple it into two separate systems that are linear in h . We accomplish this by introducing two dummy variables

$$t_{i_1j\underline{k}0} = \sum_{i_2 \in [\ell]} m_{j\underline{k}i_1i_2} \quad \text{and} \quad t_{i_1j\underline{k}1} = \sum_{i_2 \in [\ell]} h_{2k_{i_1}k_{i_2}} m_{j\underline{k}i_1i_2}$$

and renaming $t_{i_1i_2j\underline{k}} = m_{j\underline{k}i_1i_2}$. Let us mention the motivation for choosing these variables. The above system of four equations has three different scales: The first is at scale $\asymp M$, the second and third are at scale $\asymp HM$, and the fourth is at scale $\asymp H^2M$. In order to apply the two-scale result from Lemma 5.33, we want to think of the fourth equation as having scale $\asymp H(HM)$, where the HM here “matches” the scale of the third equation. This is precisely the role played by the second set of dummy variables, which is at scale $\asymp HM$. The first set of dummy variables, at scale $\asymp M$, pertains to the first and second equations, and the two sets of dummy variables are chosen in such a way that they are related at consecutive scales as in Lemma 5.33.

Thus, solving the original system is equivalent to solving

$$\begin{aligned} \sum_{i_1 \in [\ell]} t_{i_1j\underline{k}0} &= n_{j\underline{k}00} \quad \text{and} \quad \sum_{i_1 \in [\ell]} h_{1k_{i_1}} t_{i_1j\underline{k}0} = n_{j\underline{k}10} \\ \sum_{i_1 \in [\ell]} t_{i_1j\underline{k}1} &= n_{j\underline{k}01} \quad \text{and} \quad \sum_{i_1 \in [\ell]} h_{1k_{i_1}} t_{i_1j\underline{k}1} = n_{j\underline{k}11}, \end{aligned}$$

together with

$$\sum_{i_2 \in [\ell]} t_{i_1i_2j\underline{k}} = t_{i_1j\underline{k}0} \quad \text{and} \quad \sum_{i_2 \in [\ell]} h_{2k_{i_1}k_{i_2}} t_{i_1i_2j\underline{k}} = t_{i_1j\underline{k}1} \quad \text{for each } i_1 \in [\ell],$$

where the variables $t_{i_1j\underline{k}0}$ and $t_{i_1j\underline{k}1}$ have ranges $[\pm \ell M]$ and $[\pm \ell HM]$, respectively. Converting the expecta-

tion over the m 's to a sum, we see that our normalized count is bounded by $O(M^{-s\ell^2|\mathcal{K}|})$ times

$$\begin{aligned}
& \max_{\substack{n_{j\underline{k}00}, n_{j\underline{k}01}, \\ n_{j\underline{k}10}, n_{j\underline{k}11} \in \mathbb{Z}: \\ j \in [s], \underline{k} \in \mathcal{K}}} \mathbb{E}_{\substack{h_{1k_1} \in [\pm H]: \\ k_1 \in [t]}} \mathbf{1}_{\mathcal{H}_1}((h_{1k_1})_{k_1}) \sum_{\substack{t_{i_1 j \underline{k} 0} \in [\pm \ell M], \\ t_{i_1 j \underline{k} 1} \in [\pm \ell H M]: \\ j \in [s], \underline{k} \in \mathcal{K}, i_1 \in [\ell]}} \\
& \left(\prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{\underline{u} \in \{0,1\}^2} \mathbf{1} \left(\sum_{i_1 \in [\ell]} h_{1k_1 i_1}^{u_1} t_{i_1 j \underline{k} u_2} = n_{j \underline{k} u} \right) \right) \\
& \times \left(\mathbb{E}_{\substack{h_{2k_1 k_2} \in [\pm H]: \\ k_1, k_2 \in [t]}} \mathbf{1}_{\mathcal{H}_2}((h_{2k_1 k_2})_{k_1, k_2}) \sum_{\substack{t_{i_1 i_2 j \underline{k}} \in [\pm M]: \\ j \in [s], \underline{k} \in \mathcal{K}, i_1, i_2 \in [\ell]}} \right. \\
& \left. \prod_{i_1 \in [\ell]} \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{u_2 \in \{0,1\}} \mathbf{1} \left(\sum_{i_2 \in [\ell]} h_{2k_1 i_1 k_2 i_2}^{u_2} t_{i_1 i_2 j \underline{k}} = t_{i_1 j \underline{k} u_2} \right) \right);
\end{aligned}$$

we will deal first with the system involving the variables $h_{2k_1 k_2}$ and then with the system involving the variables h_{1k_1} . Note that

$$\begin{aligned}
& \sum_{\substack{t_{i_1 i_2 j \underline{k}} \in [\pm M]: \\ j \in [s], \underline{k} \in \mathcal{K}, i_1, i_2 \in [\ell]}} \prod_{i_1 \in [\ell]} \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{u_2 \in \{0,1\}} \mathbf{1} \left(\sum_{i_2 \in [\ell]} h_{2k_1 i_1 k_2 i_2}^{u_2} t_{i_1 i_2 j \underline{k}} = t_{i_1 j \underline{k} u_2} \right) \\
& = \prod_{i_1 \in [\ell]} \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \sum_{\substack{t_{i_1 i_2 j \underline{k}} \in [\pm M]: \\ i_2 \in [\ell]}} \mathbf{1} \left(\sum_{i_2 \in [\ell]} t_{i_1 i_2 j \underline{k}} = t_{i_1 j \underline{k} 0} \right) \cdot \mathbf{1} \left(\sum_{i_2 \in [\ell]} h_{2k_1 i_1 k_2 i_2} t_{i_1 i_2 j \underline{k}} = t_{i_1 j \underline{k} 1} \right).
\end{aligned}$$

Thus, we have

$$\begin{aligned}
& \mathbb{E}_{\substack{h_{2k_1 k_2} \in [\pm H]: \\ k_1, k_2 \in [t]}} \mathbf{1}_{\mathcal{H}_2}((h_{2k_1 k_2})_{k_1, k_2}) \sum_{\substack{t_{i_1 i_2 j \underline{k}} \in [\pm M], \\ j \in [s], \underline{k} \in \mathcal{K}, i_1, i_2 \in [\ell]}} \\
& \prod_{i_1 \in [\ell]} \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{u_2 \in \{0,1\}} \mathbf{1} \left(\sum_{i_2 \in [\ell]} h_{2k_1 i_1 k_2 i_2}^{u_2} t_{i_1 i_2 j \underline{k}} = t_{i_1 j \underline{k} u_2} \right) \\
& = \mathbb{E}_{\substack{h_{2k_1 k_2} \in [\pm H]: \\ k_1, k_2 \in [t]}} \mathbf{1}_{\mathcal{H}_2}((h_{2k_1 k_2})_{k_1, k_2}) \prod_{i_1 \in [\ell]} \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \\
& \sum_{\substack{t_{i_1 i_2 j \underline{k}} \in [\pm M]: \\ i_2 \in [\ell]}} \mathbf{1} \left(\sum_{i_2 \in [\ell]} t_{i_1 i_2 j \underline{k}} = t_{i_1 j \underline{k} 0} \right) \cdot \mathbf{1} \left(\sum_{i_2 \in [\ell]} h_{2k_1 i_1 k_2 i_2} t_{i_1 i_2 j \underline{k}} = t_{i_1 j \underline{k} 1} \right).
\end{aligned}$$

Lemma 5.33 applied with $l = r = 2$ allows us to bound this expression by

$$O(\eta^{-O(1)} M^{s\ell^2|\mathcal{K}| - 2s\ell|\mathcal{K}|} H^{-s\ell|\mathcal{K}|}).$$

Similarly, the count of solutions to the system involving the variables h_{1k_1} is

$$\begin{aligned}
& \max_{\substack{n_{j\underline{k}00}, n_{j\underline{k}01}, \\ n_{j\underline{k}10}, n_{j\underline{k}11} \in \mathbb{Z}: \\ j \in [s], \underline{k} \in \mathcal{K}}} \mathbb{E}_{\substack{h_{1k_1} \in [\pm H]: \\ k_1 \in [t]}} \mathbf{1}_{\mathcal{H}_1}((h_{1k_1})_{k_1}) \sum_{\substack{t_{i_1 j \underline{k} 0} \in [\pm \ell M], \\ t_{i_1 j \underline{k} 1} \in [\pm \ell H M]: \\ j \in [s], \underline{k} \in \mathcal{K}, i_1 \in [\ell]}} \\
& \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{\underline{u} \in \{0,1\}^2} \mathbf{1} \left(\sum_{i_1 \in [\ell]} h_{1k_1 i_1}^{u_1} t_{i_1 j \underline{k} u_2} = n_{j \underline{k} u} \right) \\
& \leq \mathbb{E}_{\substack{h_{1k_1} \in [\pm H]: \\ k_1 \in [t]}} \mathbf{1}_{\mathcal{H}_1}((h_{1k_1})_{k_1}) \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{u_2 \in \{0,1\}} \\
& \max_{n_0, n_1 \in \mathbb{Z}} \sum_{\substack{t_{i_1 j \underline{k} u_2} \in [\pm \ell H^{u_2} M]: \\ i_1 \in [\ell]}} \mathbf{1} \left(\sum_{i_1 \in [\ell]} t_{i_1 j \underline{k} u_2} = n_0 \right) \cdot \mathbf{1} \left(\sum_{i_1 \in [\ell]} h_{1k_1 i_1} t_{i_1 j \underline{k} u_2} = n_1 \right),
\end{aligned}$$

and this can be bounded by

$$O(\eta^{-O(1)} M^{2s\ell|\mathcal{K}| - 4s|\mathcal{K}|} H^{s\ell|\mathcal{K}| - 4s|\mathcal{K}|})$$

using the case $l = 1, r = 2$ of Lemma 5.33. Note that the big centered equations had an *inequality* rather than an *equality* because we interchanged the max with the sums and products.

Combining these two bounds and the normalising factor $O(M^{-s\ell^2|\mathcal{K}|})$, we bound the overall normalised count by

$$\begin{aligned}
& \ll M^{-s\ell^2|\mathcal{K}|} \cdot \left(\eta^{-O(1)} M^{s\ell^2|\mathcal{K}| - 2s\ell|\mathcal{K}|} H^{-s\ell|\mathcal{K}|} \right) \cdot \left(\eta^{-O(1)} M^{2s\ell|\mathcal{K}| - 4s|\mathcal{K}|} H^{s\ell|\mathcal{K}| - 4s|\mathcal{K}|} \right) \\
& = \eta^{-O(1)} M^{-4s|\mathcal{K}|} H^{-4s|\mathcal{K}|},
\end{aligned}$$

as claimed.

The case $r > 2$.

We handle the case $r > 2$ similarly. For $l \in [r]$, $i_1, \dots, i_l \in [\ell]$, $j \in [s]$, $\underline{k} \in \mathcal{K}$, $\underline{u} \in \{0,1\}^r$, we define the auxiliary variable

$$t_{i_1 \dots i_l j \underline{k} \underline{u}} = \sum_{i_{l+1}, \dots, i_r \in [\ell]} h_{(l+1)k_{i_1} \dots k_{(l+1)i_{l+1}}}^{u_{l+1}} \cdots h_{rk_{i_1} \dots k_{ri_r}}^{u_r} m_{j \underline{k} i_1 \dots i_r}$$

at scale $H^{u_{l+1} + \dots + u_r} M$. Notice that $t_{i_1 \dots i_l j \underline{k} \underline{u}}$ depends on only the coordinates $u_{l+1}, \dots, u_r$ of $\underline{u}$ and that the t 's satisfy the relation

$$t_{i_1 \dots i_l j \underline{k} \underline{u}} = \sum_{i_{l+1} \in [\ell]} h_{(l+1)k_{i_1} \dots k_{(l+1)i_{l+1}}}^{u_{l+1}} t_{i_1 \dots i_{l+1} j \underline{k} \underline{u}}. \quad (5.48)$$

Iterating (5.48) through intermediate scales, we see that for each $j \in [s]$ and $\underline{k} \in \mathcal{K}$, we can rewrite

$$\begin{aligned}
& \sum_{\substack{m_{j\underline{k}i_1 \dots i_r} \in [\pm M]: \\ i_1, \dots, i_r \in [\ell]}} \mathbf{1} \left(\sum_{i_1, \dots, i_r \in [\ell]} h_{1k_{1i_1}}^{u_1} \cdots h_{rk_{1i_1} \dots k_{ri_r}}^{u_r} m_{j\underline{k}i_1 \dots i_r} = n_{j\underline{k}u} \right) \\
&= \sum_{\substack{|t_{i_1 j \underline{k} u}| \ll H^{u_2 + \dots + u_r} M: \\ i_1 \in [\ell]}} \mathbf{1} \left(\sum_{i_1 \in [\ell]} h_{1k_{1i_1}}^{u_1} t_{i_1 j \underline{k} u} = n_{j\underline{k}u} \right) \\
&\times \sum_{\substack{|t_{i_1 i_2 j \underline{k} u}| \ll H^{u_3 + \dots + u_r} M: \\ i_1, i_2 \in [\ell]}} \prod_{i_1 \in [\ell]} \mathbf{1} \left(\sum_{i_2 \in [\ell]} h_{2k_{1i_1} k_{2i_2}}^{u_2} t_{i_1 i_2 j \underline{k} u} = t_{i_1 j \underline{k} u} \right) \\
&\times \cdots \times \sum_{\substack{|t_{i_1 \dots i_r j \underline{k} u}| \leq M: \\ i_1, \dots, i_r \in [\ell]}} \prod_{i_r \in [\ell]} \mathbf{1} \left(\sum_{i_r \in [\ell]} h_{rk_{1i_1} \dots k_{ri_r}}^{u_r} t_{i_1 \dots i_r j \underline{k} u} = t_{i_1 \dots i_{r-1} j \underline{k} u} \right),
\end{aligned}$$

where we note that $t_{i_1 \dots i_r j \underline{k} u} = m_{j\underline{k}i_1 \dots i_r}$. Hence our normalized count is bounded by $O(M^{-s\ell^r |\mathcal{K}|})$ times

$$\begin{aligned}
& \max_{\substack{n_{j\underline{k}u} \in \mathbb{Z}: j \in [s], \\ \underline{k} \in \mathcal{K}, u \in \{0,1\}^r}} \mathbb{E}_{\substack{h_{1k_1} \in [\pm H]: \\ k_1 \in [t]}} \mathbf{1}_{\mathcal{H}_1}((h_{1k_1})_{k_1}) \sum_{\substack{|t_{i_1 j \underline{k} u}| \ll H^{u_2 + \dots + u_r} M: j \in [s], \\ \underline{k} \in \mathcal{K}, u_2, \dots, u_r \in \{0,1\}, i_1 \in [\ell]}} \\
& \left(\prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{u_1, \dots, u_r \in \{0,1\}} \mathbf{1} \left(\sum_{i_1 \in [\ell]} h_{1k_{1i_1}}^{u_1} t_{i_1 j \underline{k} u} = n_{j\underline{k}u} \right) \right) \\
& \times \mathbb{E}_{\substack{h_{2k_1 k_2} \in [\pm H]: \\ k_1, k_2 \in [t]}} \mathbf{1}_{\mathcal{H}_2}((h_{2k_1 k_2})_{k_1, k_2}) \sum_{\substack{|t_{i_1 i_2 j \underline{k} u}| \ll H^{u_3 + \dots + u_r} M: j \in [s], \\ \underline{k} \in \mathcal{K}, u_3, \dots, u_r \in \{0,1\}, i_1, i_2 \in [\ell]}} \\
& \left(\prod_{i_1 \in [\ell]} \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{u_2, \dots, u_r \in \{0,1\}} \mathbf{1} \left(\sum_{i_2 \in [\ell]} h_{2k_{1i_1} k_{2i_2}}^{u_2} t_{i_1 i_2 j \underline{k} u} = t_{i_1 j \underline{k} u} \right) \right) \\
& \times \cdots \times \mathbb{E}_{\substack{h_{rk_1 \dots k_r} \in [\pm H]: \\ k_1, \dots, k_r \in [t]}} \mathbf{1}_{\mathcal{H}_r}((h_{rk_1 \dots k_r})_{k_1, \dots, k_r}) \sum_{\substack{|t_{i_1 \dots i_r j \underline{k} u}| \leq M: j \in [s], \\ \underline{k} \in \mathcal{K}, i_1, \dots, i_r \in [\ell]}} \\
& \prod_{i_1, \dots, i_{r-1} \in [\ell]} \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{u_r \in \{0,1\}} \mathbf{1} \left(\sum_{i_r \in [\ell]} h_{rk_{1i_1} \dots k_{ri_r}}^{u_r} t_{i_1 \dots i_r j \underline{k} u} = t_{i_1 \dots i_{r-1} j \underline{k} u} \right).
\end{aligned}$$

As in the case $r = 2$, we have decoupled our original system of degree- r multilinear equations into r smaller systems, each consisting of a pair of multilinear equations at consecutive scales. We will next use Lemma 5.33 to bound the counts of solutions for these subsystems individually, starting from the system involving the variables $h_{rk_1 \dots k_r}$. The count of solutions to this system can be rephrased as

$$\mathbb{E}_{\substack{h_{rk_1 \dots k_r} \in [\pm H]: \\ k_1, \dots, k_r \in [t]}} \mathbf{1}_{\mathcal{H}_r}((h_{rk_1 \dots k_r})_{k_1, \dots, k_r}) \prod_{i_1, \dots, i_{r-1} \in [\ell]} \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}}$$

$$\sum_{\substack{|t_{i_1 \dots i_r j \underline{k} \underline{u}}| \leq M: \\ i_r \in [\ell]}} \prod_{u_r \in \{0,1\}} \mathbf{1} \left(\sum_{i_r \in [\ell]} h_{r k_{1 i_1} \dots k_{r i_r}}^{u_r} t_{i_1 \dots i_r j \underline{k} \underline{u}} = t_{i_1 \dots i_{r-1} j \underline{k} \underline{u}} \right),$$

and bounded by $\ll \eta^{-O(1)} M^{s\ell^r |\mathcal{K}| - 2s\ell^{r-1} |\mathcal{K}|} H^{-s\ell^{r-1} |\mathcal{K}|}$ using the case $l = r$ of Lemma 5.33. Similarly, the count corresponding to the system involving the variables $h_{(r-1)k_1 \dots k_{r-1}}$ can be rephrased as

$$\begin{aligned} & \mathbb{E}_{\substack{h_{(r-1)k_1 \dots k_{r-1}} \in [\pm H]: \\ k_1, \dots, k_{r-1} \in [t]}} \mathbf{1}_{\mathcal{H}_{r-1}}((h_{(r-1)k_1 \dots k_{r-1}})_{k_1, \dots, k_{r-1}}) \prod_{i_1, \dots, i_{r-2} \in [\ell]} \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{u_r \in \{0,1\}} \\ & \sum_{\substack{|t_{i_1 \dots i_{r-1} j \underline{k} \underline{u}}| \ll H^{u_r} M: \\ i_{r-1} \in [\ell]}} \prod_{u_{r-1} \in \{0,1\}} \mathbf{1} \left(\sum_{i_{r-1} \in [\ell]} h_{(r-1)k_{1 i_1} \dots k_{(r-1) i_{r-1}}}^{u_{r-1}} t_{i_1 \dots i_{r-1} j \underline{k} \underline{u}} = t_{i_1 \dots i_{r-2} j \underline{k} \underline{u}} \right) \end{aligned}$$

and bounded by

$$\ll \eta^{-O(1)} M^{2s\ell^{r-1} |\mathcal{K}| - s\ell^r |\mathcal{K}| - 4s\ell^{r-1} |\mathcal{K}|} H^{s\ell^{r-1} |\mathcal{K}| - 4s\ell^{r-1} |\mathcal{K}|}$$

with the help of the case $l = r - 1$ of Lemma 5.33. Continuing in this manner and combining the results, we can bound our initial normalized count by

$$\begin{aligned} & \ll M^{-s\ell^r |\mathcal{K}|} \cdot \prod_{l \in [r]} \eta^{-O(1)} M^{2^{r-l} s\ell^l |\mathcal{K}| - 2^{r-l+1} s\ell^{l-1} |\mathcal{K}|} \\ & H^{(r-l)2^{r-l-1} s\ell^l |\mathcal{K}| - (r-l+1)2^{r-l} s\ell^{l-1} |\mathcal{K}|}, \end{aligned}$$

where in each l -multiplicand, the negative exponents of η, M, H come from an application of Lemma 5.33 to bound the count of solutions to the system involving the variables $h_{l k_1 \dots k_l}$ and the positive exponents come from counting the admissible values of the variables $t_{i_1 \dots i_l j \underline{k} \underline{u}}$. The telescoping nature of the exponents leads to cancellations: The exponents of M and H cancel out as

$$\begin{aligned} & -s\ell^r |\mathcal{K}| + (s\ell^r |\mathcal{K}| - 2s\ell^{r-1} |\mathcal{K}|) + \dots + (2^{r-1} s\ell |\mathcal{K}| - 2^r s |\mathcal{K}|) = -2^r s |\mathcal{K}|, \\ & -s\ell^{r-1} |\mathcal{K}| + (s\ell^{r-1} |\mathcal{K}| - 4s\ell^{r-2} |\mathcal{K}|) + \dots + ((r-1)2^{r-2} s\ell |\mathcal{K}| - r2^{r-1} s |\mathcal{K}|) = -r2^{r-1} s |\mathcal{K}|, \end{aligned}$$

respectively, yielding the desired final bound

$$\ll \eta^{-O(1)} M^{-2^r s |\mathcal{K}|} H^{-r2^{r-1} s |\mathcal{K}|}.$$

□

The last result in this section ensures that the set $\mathcal{H}_{l,\eta}$ fills almost all of $[\pm H]^{t^l}$ when $\eta > 0$ is small.

Lemma 5.38. *Let $\eta > 0$, $H, l, t \in \mathbb{N}$, and let $\mathcal{H} = \mathcal{H}_{l,\eta}$ be as in (5.46). Then all but an $O_{l,t}(\eta)$ -proportion of the elements of $[\pm H]^{t^l}$ lie in $\mathcal{H}$.*

Proof. Setting $k = (k_1, \dots, k_l)$ and $h_k = h_{l k_1, \dots, k_l}$, we can rewrite

$$\mathcal{H} = \{(h_k)_{k \in [t^l]} \in [\pm H]^{t^l} : \gcd(h_k - h_{k'}, h_{k'} - h_{k''}) \leq \eta^{-1}, \\ |h_k| \geq \eta H \text{ for all distinct } k, k', k'' \in [t^l]\}.$$

We will show that there are very few elements of $[\pm H]^{t^l}$ failing each condition for membership in $\mathcal{H}$, and then we will use the union bound to conclude that $[\pm H]^{t^l} \setminus \mathcal{H}$ is small. For the gcd condition, we estimate

$$\begin{aligned} \mathbb{E}_{h, h', h'' \in [\pm H]} \mathbf{1}(\gcd(h - h'', h' - h'') \geq \eta^{-1}) &\ll \mathbb{E}_{h, h' \in [\pm 2H]} \mathbf{1}(\gcd(h, h') \geq \eta^{-1}) \\ &\ll \frac{1}{H^2} \sum_{g=\eta^{-1}}^{2H} \sum_{h, h' \in [\pm 2H]} \mathbf{1}(g|h, h') \\ &\ll \frac{1}{H^2} \sum_{g=\eta^{-1}}^{2H} (H/g)^2 \\ &= \sum_{g=\eta^{-1}}^{2H} \frac{1}{g^2} \ll \eta. \end{aligned}$$

For the size condition, we note that $|h| < \eta H$ for an $O(\eta)$ -fraction of $h \in [\pm H]$. The conclusion of the lemma now follows from the union bound. $\square$

5.8 Concatenation along polynomials

In this section, we combine the conclusions of the previous two sections in order to complete the proof of Theorem 5.10. First, we specialize the results of Section 5.6 to averages of polynomial box norms.

Proposition 5.39. *Let $d, r, s \in \mathbb{N}$ with $d \leq r$, and let ℓ be a power of 2. There exists a positive integer $t = O_{\ell, r, s}(1)$ such that the following holds. Let $\delta > 0$ and $D, H, M, N, V \in \mathbb{N}$ satisfy $H^d M V \leq N$, and let $C_1, \dots, C_s \in \mathbb{Z}^D[\underline{h}]$ be polynomials of degree at most d with coefficients of size at most V . Then for all 1-bounded functions $f : \mathbb{Z}^D \rightarrow \mathbb{C}$ supported on $[N]^D$, the bound*

$$\mathbb{E}_{\underline{h} \in [\pm H]^r} \|f\|_{C_1(\underline{h}) \cdot [\pm M], \dots, C_s(\underline{h}) \cdot [\pm M]}^{2^s} \geq \delta N^D \quad (5.49)$$

implies that

$$\mathbb{E}_{\substack{h_{lk_1 \dots k_l} \in [\pm H]: \\ l \in [r], k_1, \dots, k_r \in [t]}} \|f\|_{\{\sum_{i_1, \dots, i_r=1}^{\ell} \mathbf{C}_j(h_{1k_1 i_1}, \dots, h_{rk_1 i_1} \dots h_{r i_r}) \cdot [\pm M]: \\ j \in [s], \underline{k} \in \mathcal{K}\}}^{2^w} \gg_{D, \ell, r, s} \delta^{O_{\ell, r, s}(1)} N^D,$$

where $\mathcal{K} = \mathcal{K}_{t, \ell, r}$ and w is the degree of the box norms appearing.

Proof. The proof consists of several applications of Corollary 5.30. First, we apply it with the indexing set $\underline{h} \in [\pm H]^r$ to obtain a natural number $t_1 = O_{\ell, s}(1)$ such that

$$\mathbb{E}_{\underline{h}_1, \dots, \underline{h}_{t_1} \in [\pm H]^r} \|f\|_{\{\mathbf{C}_j(\underline{h}_{k_{11}}) \cdot [\pm M] + \dots + \mathbf{C}_j(\underline{h}_{k_{1\ell}}) \cdot [\pm M]: \\ j \in [s], 1 \leq k_{11} < \dots < k_{1\ell} \leq t_1\}}^{2^{w_1}} \gg_{D, \ell, s} \delta^{O_{\ell, s}(1)} N^D,$$

where w_1 is the degree of the box norms appearing.

For the second application of Corollary 5.30, we write $\underline{h}_k = (h_{1k}, \tilde{h}_{1k})$ and take the indexing set to be the set of tuples $(\tilde{h}_{11}, \dots, \tilde{h}_{1t_1}) \in [\pm H]^{(r-1)t_1}$. Then Corollary 5.30, applied separately for each fixed choice of $h_{11}, \dots, h_{1t_1}$, gives $t_2 = O_{\ell, s}(1)$ for which

$$\mathbb{E}_{\substack{h_{1k_1} \in [\pm H]: \\ k_1 \in [t_1]}} \mathbb{E}_{\substack{\tilde{h}_{1k_1 k_2} \in [\pm H]^{r-1}: \\ k_1 \in [t_1], k_2 \in [t_2]}} \|f\|_{\{\mathbf{C}_j(h_{1k_{11}}, \tilde{h}_{1k_{11} k_{21}}) \cdot [\pm M] + \dots + \mathbf{C}_j(h_{1k_{11}}, \tilde{h}_{1k_{11} k_{2\ell}}) \cdot [\pm M] \\ + \dots + \mathbf{C}_j(h_{1k_{1\ell}}, \tilde{h}_{1k_{1\ell} k_{21}}) \cdot [\pm M] + \dots + \mathbf{C}_j(h_{1k_{1\ell}}, \tilde{h}_{1k_{1\ell} k_{2\ell}}) \cdot [\pm M]: \\ j \in [s], 1 \leq k_{11} < \dots < k_{1\ell} \leq t_1, 1 \leq k_{21} < \dots < k_{2\ell} \leq t_2\}}^{2^{w_2}} \gg_{D, \ell, s} \delta^{O_{\ell, s}(1)} N^D,$$

where w_2 is the degree of the box norms appearing.

For the next step, we similarly write $\tilde{h}_{1k_1 k_2} = (h_{2k_1 k_2}, \tilde{h}_{2k_1 k_2})$ and take our indices to be the tuples $(\tilde{h}_{2k_1 k_2})_{k_1 \in [t_1], k_2 \in [t_2]} \in [\pm H]^{(r-2)t_1 t_2}$. Applying Corollary 5.30 separately for each $h_{1k_1}, h_{2k_1 k_2} \in [\pm H]$, with $k_1 \in [t_1], k_2 \in [t_2]$, we obtain $t_3 = O_{\ell, s}(1)$ for which

$$\mathbb{E}_{\substack{h_{1k_1} \in [\pm H]: \\ k_1 \in [t_1]}} \mathbb{E}_{\substack{h_{1k_1 k_2} \in [\pm H]: \\ k_1 \in [t_1], k_2 \in [t_2]}} \mathbb{E}_{\substack{\tilde{h}_{2k_1 k_2 k_3} \in [\pm H]^{r-2}: \\ k_1 \in [t_1], k_2 \in [t_2], k_3 \in [t_3]}} \|f\|_{\{\sum_{i_1, i_2, i_3=1}^{\ell} \mathbf{C}_j(h_{1k_{1i_1}}, h_{2k_{1i_1} k_{2i_2}}, \tilde{h}_{2k_{1i_1} k_{2i_2} k_{3i_3}}) \cdot [\pm M]: \\ j \in [s], 1 \leq k_{1i_1} < \dots < k_{1i_\ell} \leq t_1, i_\ell \in [3]\}}^{2^{w_3}} \gg_{D, \ell, s} \delta^{O_{\ell, s}(1)} N^D,$$

where w_3 is the degree of the box norms appearing. Continuing in this fashion $r - 2$ more times, we obtain

$t_3, \dots, t_r = O_{\ell, r, s}(1)$ such that

$$\mathbb{E}_{\substack{h_1 k_1 \dots k_\ell \in [\pm H]: \\ l \in [r], k_l \in [t_l]}} \|f\|_{\substack{\sum_{i_1, \dots, i_r=1}^\ell \mathbf{C}_j(h_1 k_{1i_1}, \dots, h_r k_{1i_1} \dots k_{ri_r}) \cdot [\pm M]: \\ j \in [s], 1 \leq k_{l1} < \dots < k_{l\ell} \leq t_l, l \in [r]}}^{2^{w_r}} \gg_{D, \ell, r, s} \delta^{O_{\ell, r, s}(1)} N^D,$$

where w_r is the degree of the box norms appearing.

Although this inequality is just as useful as the conclusion of the proposition, it will later be notationally convenient to replace the t_l 's with a single parameter. Setting $t = \max(t_1, \dots, t_r)$ and noting that the coordinates of the polynomial

$$\sum_{i_1, \dots, i_r=1}^\ell \mathbf{C}_j(h_1 k_{1i_1}, \dots, h_r k_{1i_1} \dots k_{ri_r})$$

are at most $O_{d, r, \ell}(VH^d)$ for all valid choices of the h 's, we deduce that

$$\sum_{i_1, \dots, i_r=1}^\ell \mathbf{C}_j(h_1 k_{1i_1}, \dots, h_r k_{1i_1} \dots k_{ri_r}) \cdot [\pm M] \subset [-AN, AN]^D$$

for some $A = O_{d, \ell, r, s}(1)$; it is here that we are using the assumption on the coefficients of $\mathbf{C}_1, \dots, \mathbf{C}_s$ and the bound $H^d MV \leq N$. We then deduce the result from the monotonicity property of box norms (Lemma 5.15(3)). $\square$

We will finally use Proposition 5.39, together with the equidistribution results from Section 5.7, to prove Theorem 5.10, which we restate for convenience.

Theorem 5.40 (Theorem 5.10). *Let $r, s \in \mathbb{N}$. There exists a positive integer $t = O_{r, s}(1)$ such that the following holds. Let $d, D, H, M, N, V \in \mathbb{N}$ and $\delta > 0$ satisfy*

$$d \leq r, \quad \delta^{-O_{r, s}(1)} \ll_{D, r, s} H \leq M \quad \text{and} \quad H^d MV \leq N,$$

and let $\mathbf{C}_1, \dots, \mathbf{C}_s \in \mathbb{Z}^D[h_1, \dots, h_r]$ be multilinear polynomials $\mathbf{C}_j(\underline{h}) = \sum_{\underline{u} \in \{0, 1\}^r} \gamma_{j\underline{u}} \underline{h}^{\underline{u}}$ of degree at most d and with coefficients at most V . Then for all 1-bounded functions $f : \mathbb{Z}^D \rightarrow \mathbb{C}$ supported on $[N]^D$, the bound

$$\mathbb{E}_{\underline{h} \in [\pm H]^r} \|f\|_{\mathbf{C}_1(\underline{h}) \cdot [\pm M], \dots, \mathbf{C}_s(\underline{h}) \cdot [\pm M]}^{2^{st}} \geq \delta N^D$$

implies that

$$\|f\|_{E_1^t, \dots, E_s^t}^{2^{st}} \gg_{D, r, s} \delta^{O_{r, s}(1)} N^D,$$

where

$$E_j = \sum_{\underline{u} \in \{0,1\}^d} \gamma_{j\underline{u}} \cdot [\pm H^{|\underline{u}|+1}].$$

Proof. The proof amounts to combining the conclusions of Propositions 5.39 and 5.35. We will apply these results with $\ell = 4$ since we need ℓ to be a power of 2 (as required by Proposition 5.39), and we want it to be at least 3 (in order to apply Proposition 5.35). We let all implicit constants depend on D, r, s^{17} and note that the powers of δ appearing in the proof will be independent of D . We also let $\mathcal{K} = \mathcal{K}_{t,4,r}$, which we recall from Section 5.7 is the set of admissible tuples of indices

$$\mathcal{K} = \{(k_{li})_{(l,i) \in [r] \times [4]} \in [t]^{4r} : 1 \leq k_{l1} < \dots < k_{l4} \leq t \text{ for all } l \in [r]\}.$$

By Proposition 5.39, there is some positive integer $t = O(1)$ (independent of D) for which

$$\mathbb{E}_{\substack{h_{lk_1 \dots k_l} \in [\pm H]: \\ l \in [r], k_1, \dots, k_r \in [t]}} \|f\|_{\sum_{i_1, \dots, i_r \in [\ell]} \mathbf{C}_j(h_{1k_{1i_1}}, \dots, h_{rk_{1i_1} \dots k_{ri_r}}) \cdot [\pm M]}^{2^w} \gg \delta^{O(1)} N^D, \quad (5.50)$$

where w is the degree of the box norms appearing. We will now make use of our equidistribution estimates. It is necessary first to restrict to “generic” tuples of the h ’s.

Recall from (5.46) the definition of the set of tuples $\mathcal{H}_{l,\eta} \subset [\pm H]^{t^l}$, which by Lemma 5.38 contains all but an $O(\eta)$ -proportion of the elements of $[\pm H]^{t^l}$. Write $\mathcal{H}_l = \mathcal{H}_{l,\eta}$ and take $\eta = c\delta^{1/c}$ for some $c > 0$ sufficiently small that the contribution of tuples $(h_{lk_1 \dots k_l})_{k_1, \dots, k_l} \in [\pm H]^{t^l} \setminus \mathcal{H}_l$ is at most half (say) of the lower bound in (5.50). Removing such “bad” tuples of h ’s and expanding (5.50), we obtain

$$\begin{aligned} \mathbb{E}_{\substack{h_{lk_1 \dots k_l} \in [\pm H]: \\ l \in [r], k_1, \dots, k_r \in [t]}} \left(\prod_{l \in [r]} \mathbf{1}_{\mathcal{H}_l}((h_{lk_1 \dots k_l})_{k_1, \dots, k_l}) \right) \sum_{\substack{m_{j\underline{k}i_1 \dots i_r} \in [\pm M]: \\ j \in [s], \underline{k} \in \mathcal{K}, i_1, \dots, i_r \in [4]}} \mu_M((m_{j\underline{k}i_1 \dots i_r})_{j, \underline{k}, i_1, \dots, i_r}) \\ \sum_x \Delta_{\{ \sum_{i_1, \dots, i_r \in [4]} \mathbf{C}_j(h_{1k_{1i_1}}, \dots, h_{rk_{1i_1} \dots k_{ri_r}}) m_{j\underline{k}i_1 \dots i_r} : f(x) \gg \delta^{O(1)} N^D \}} \quad (5.51) \end{aligned}$$

For each $j \in [s]$, $\underline{k} \in \mathcal{K}$, $\underline{u} \in \{0,1\}^r$, we introduce dummy variables

$$n_{j\underline{k}\underline{u}} = \sum_{i_1, \dots, i_r \in [4]} h_{1k_{1i_1}}^{u_1} \dots h_{rk_{1i_1} \dots k_{ri_r}}^{u_r} m_{j\underline{k}i_1 \dots i_r} \quad (5.52)$$

¹⁷and also on d , which, however, is bounded by r

at scale $H^{u_1+\dots+u_r}M$, and we define

$$\mathcal{N}(\underline{n}) = \mathbb{E}_{\substack{m_{j\underline{k}i_1\dots i_r} \in [\pm M]: \\ j \in [s], \underline{k} \in \mathcal{K}, i_1, \dots, i_r \in [4]}} \mathbb{E}_{\substack{h_{l\underline{k}_1\dots k_l} \in [\pm H]: \\ l \in [r], k_1, \dots, k_r \in [t]}} \left(\prod_{l \in [r]} \mathbf{1}_{\mathcal{H}_l}((h_{l\underline{k}_1\dots k_l})_{k_1, \dots, k_l}) \right) \\ \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{\underline{u} \in \{0,1\}^r} \mathbf{1} \left(\sum_{i_1, \dots, i_r \in [4]} h_{1\underline{k}i_1}^{u_1} \cdots h_{r\underline{k}i_r}^{u_r} m_{j\underline{k}i_1\dots i_r} = n_{j\underline{k}\underline{u}} \right).$$

In words, the kernel $\mathcal{N}(\underline{n})$ is the probability that (5.52) is satisfied when the m 's are chosen (independently, uniformly) at random from $[\pm M]$ and the tuples of h 's are chosen (independently, uniformly) at random from the sets $\mathcal{H}_l$. Inserting the definition of $\mathcal{N}(\underline{n})$ into (5.51) and applying the triangle inequality gives us the neat bound

$$\sum_{\substack{n_{j\underline{k}\underline{u}} \ll H^{|\underline{u}|}M: \\ j \in [s], \underline{k} \in \mathcal{K}, \underline{u} \in \{0,1\}^r}} \mathcal{N}(\underline{n}) \cdot \left| \sum_{\mathbf{x}} \Delta_{\{\sum \underline{u} \in \{0,1\}^r\}} \gamma_{j\underline{u}n_{j\underline{k}\underline{u}}: j \in [s], \underline{k} \in \mathcal{K}} f(\mathbf{x}) \right| \gg \delta^{O(1)} N^D.$$

The kernel $\mathcal{N}$ is supported on the tuples $(n_{j\underline{k}\underline{u}})_{j, \underline{k}, \underline{u}}$ satisfying $n_{j\underline{k}\underline{u}} \ll H^{|\underline{u}|}M$, of which there are

$$\ll \prod_{j \in [s]} \prod_{\underline{k} \in \mathcal{K}} \prod_{\underline{u} \in \{0,1\}^r} H^{|\underline{u}|}M \ll M^{2^r s |\mathcal{K}|} H^{r 2^{r-1} s |\mathcal{K}|}.$$

It is at this point that we crucially incorporate the bound

$$\mathcal{N}(\underline{n}) \ll \delta^{-O(1)} M^{-2^r s |\mathcal{K}|} H^{-r 2^{r-1} s |\mathcal{K}|} \quad (5.53)$$

provided by Proposition 5.35 to get the estimate

$$\mathbb{E}_{\substack{n_{j\underline{k}\underline{u}} \ll H^{|\underline{u}|}M: \\ j \in [s], \underline{k} \in \mathcal{K}, \underline{u} \in \{0,1\}^r}} \left| \sum_{\mathbf{x}} \Delta_{\{\sum \underline{u} \in \{0,1\}^r\}} \gamma_{j\underline{u}n_{j\underline{k}\underline{u}}: j \in [s], \underline{k} \in \mathcal{K}} f(\mathbf{x}) \right| \gg \delta^{O(1)} N^D.$$

From here, we use the Cauchy–Schwarz inequality to remove the absolute value and introduce Fejér kernels, and then use Lemma 5.15(6) to trim the boxes (as in the end of the proof sketch of Proposition 5.24). To conclude the proof, we use Lemma 5.15 to replace the differencing with respect to $[\pm N]^D$ by a differencing with respect to one of the boxes E_j ; by the monotonicity of box norms, we can also increase the multiplicity of every other E_i by 1, and we conclude the theorem with $t+1$ in place of t . $\square$

5.9 Proof of Theorem 5.6

We conclude by deriving Theorem 5.6 from Proposition 5.23 (the output of PET) and Theorem 5.10 (the concatenation result). Assume that

$$\left| \sum_{\mathbf{x} \in \mathbb{Z}^D} \mathbb{E}_{z \in [K]} f_0(\mathbf{x}) \cdot f_1(\mathbf{x} + \mathbf{P}_1(z)) \cdots f_\ell(\mathbf{x} + \mathbf{P}_\ell(z)) \right| \geq \delta N^D.$$

Proposition 5.23 gives us $r, s = O_{d,\ell}(1)$ and nonzero r -variable multilinear polynomials $\mathbf{C}_1, \dots, \mathbf{C}_s \in \mathbb{Z}^D[h]$ (depending only on $\mathbf{P}_1, \dots, \mathbf{P}_\ell$) of degree at most $d-1$, such that

$$\mathbf{C}_j(\underline{h}) = \sum_{\substack{\underline{u} \in \{0,1\}^r, \\ |\underline{u}| \leq d-1}} (|\underline{u}|+1)! \cdot (\beta_{1(|\underline{u}|+1)} - \beta_{w_{j\underline{u}}(|\underline{u}|+1)}) \underline{h}^{\underline{u}}$$

for some indices $w_{j\underline{u}} \in [0, \ell]$ (with $\beta_{0(|\underline{u}|+1)} := 0$), for which

$$\mathbb{E}_{\underline{h} \in [\pm H]^r} \|f_1\|_{\mathbf{C}_1(\underline{h}) \cdot [\pm M], \dots, \mathbf{C}_s(\underline{h}) \cdot [\pm M]}^{2^s} \gg_{d,D,\ell} \delta^{O_{d,\ell}(1)} N^D$$

holds for all $\delta^{-O_{d,\ell}(1)} \ll_{d,D,\ell} H \leq M \ll_{d,D,\ell} \delta^{O_{d,\ell}(1)} K$. Furthermore, by property (3) of Definition 5.20, $\beta_{1(|\underline{u}|+1)} - \beta_{w_{j\underline{u}}(|\underline{u}|+1)}$ is the leading coefficient of $\mathbf{P}_1 - \mathbf{P}_{w_{j\underline{u}}}$, which implies that

$$\deg(\mathbf{P}_1 - \mathbf{P}_{w_{j\underline{u}}}) = |\underline{u}| + 1.$$

By Theorem 5.10 (with $d-1$ in place of d), there exists $t_0 = O_{d,\ell}(1)$ such that

$$\|f\|_{E_{10}^{t_0}, \dots, E_{s0}^{t_0}}^{2^{st_0}} \gg_{d,D,\ell} \delta^{O_{d,\ell}(1)} N^D$$

for the boxes

$$E_{j0} = \sum_{\substack{\underline{u} \in \{0,1\}^r, \\ |\underline{u}| \leq d-1}} d_{1w_{j\underline{u}}}! \cdot (\beta_{1d_{1w_{j\underline{u}}}} - \beta_{w_{j\underline{u}}d_{1w_{j\underline{u}}}}) \cdot [\pm H^{d_{1w_{j\underline{u}}}}]$$

as long as $H^{d-1}MV \leq N$. We now set $H = M \asymp_{d,D,\ell} \delta^c K$ for a suitably large constant $c = c_{d,\ell} > 0$, which, together with the assumption on the size of K , ensures that the aforementioned conditions on H, M are satisfied. Using parts (4)-(6) of Lemma 5.15, we perform the following operations on each box E_{j0} : Remove all of the directions making up E_{j0} except for one direction corresponding to a (nonzero) highest-degree monomial; enlarge H to K (at the loss of the factor $O_{d,D,\ell}(\delta^{O_{d,\ell}(1)})$ in the lower bound); and remove the

factor $d_{1w_{j\underline{u}}}!$. In total, we replace each box E_{j0} by

$$(\beta_{1d_{1w_{j\underline{u}}}} - \beta_{w_{j\underline{u}}d_{1w_{j\underline{u}}}}) \cdot [\pm K^{d_{1w_{j\underline{u}}}}]$$

for some $w_{j\underline{u}} \neq 1$. The theorem now follows, with $t = t_0\ell$, from the monotonicity property of box norms, where we have implicitly used the bound

$$|(\beta_{1d_{1w_{j\underline{u}}}} - \beta_{w_{j\underline{u}}d_{1w_{j\underline{u}}}}) \cdot K^{d_{1w_{j\underline{u}}}}| \ll_D VK^{d_{1w_{j\underline{u}}}} \ll_{d,D,\ell} VK^d \leq N.$$

5.10 One more PET result

In the last section, we prove one more result required in the following chapter.

Proposition 5.41. *Let $d, r \in \mathbb{N}$ with $d \geq 2$, and let $P \in \mathbb{Z}[z]$ be a polynomial of degree d with coefficients of size at most V . There exists a positive integer $t = O_{d,r}(1)$ such that the following holds. Let $\delta > 0$ and $K, N \in \mathbb{N}$ satisfy $\delta^{-O_{d,r}(1)} \ll_{d,r} K \ll_{d,r} (N/V)^{1/d}$. Then for every 1-bounded function $f : \mathbb{Z} \rightarrow \mathbb{C}$ supported on $[N]$, the bound*

$$\sum_{h_1, \dots, h_r} \mu_K(\underline{h}) \sum_x \mathbb{E}_{z \in [K]} \prod_{\underline{\epsilon} \in \{0,1\}^r} \mathcal{C}^{|\underline{\epsilon}|} f(x + P(z + \underline{\epsilon} \cdot \underline{h})) \geq \delta N \quad (5.54)$$

implies that

$$\|f\|_{U^t(\beta_d \cdot [\pm N/\beta_d])}^{2^t} \gg_{d,r} \delta^{O_{d,r}(1)} N,$$

where β_d is the leading coefficient of P .

Proof. Write $P(z) = \sum_{i=0}^d \beta_i z^i$. Setting

$$P_{\underline{\epsilon}}(z, \underline{h}) := P(z + \underline{\epsilon} \cdot \underline{h}) - P(\underline{\epsilon} \cdot \underline{h}),$$

we first show that the family $\mathcal{Q} := \{P_{\underline{\epsilon}} : \underline{\epsilon} \in \{0,1\}^r\}$ (ordered so that $P_{\underline{1}}$ comes first) descends from $\mathcal{P} = \{P - P(0)\}$. Indeed, $\mathcal{Q}$ is normal, for which it is crucial that $d \geq 2$, since otherwise $P_{\underline{\epsilon}}(z, \underline{h}) = \beta_1 z$ for every $\underline{\epsilon}$, as the contributions of $\underline{h}$ are absorbed into the term not depending on z . Moreover, a simple

computation gives that

$$P_{\underline{\epsilon}}(z, \underline{h}) = \sum_{i=0}^d z^i \sum_{\substack{\underline{u} \in \mathbb{N}_0^r: |\underline{u}| \leq d-i, \\ \text{supp}(\underline{u}) \subseteq \text{supp}(\underline{\epsilon})}} \binom{|\underline{u}|+i}{\underline{u}} \beta_{|\underline{u}|+i} \underline{h}^{\underline{u}},$$

and hence the formula (5.24) holds with $w_{\underline{\epsilon}\underline{u}} = \mathbf{1}(\text{supp}(\underline{u}) \subseteq \text{supp}(\underline{\epsilon}))$ and $w_{\underline{u}} = 0$. This immediately implies property (1) from Definition 5.20. One also sees that for distinct $\underline{\epsilon}, \underline{\epsilon}' \in \{0, 1\}^r$, the leading coefficient of $P_{\underline{\epsilon}} - P_{\underline{\epsilon}'}$ as a polynomial in z is $d\beta_d(\underline{\epsilon} - \underline{\epsilon}') \cdot \underline{h}$. In particular, the leading coefficient of $P_{\underline{\epsilon}} - P_{\underline{\epsilon}'}$ is a multilinear homogeneous polynomial of degree 1. This gives property (2) of Definition 5.20, and property (3) follows since β_d is the leading coefficient of P . Hence $\mathcal{Q}$ indeed descends from $\mathcal{P}$. This allows us to use Proposition 5.22 to find positive integers $r' = O_{d,r}(1)$ and $m_1, \dots, m_{r'}$ such that $\mathcal{Q}' = \partial_{m_{r'}, \dots, m_1} \mathcal{Q}$ is linear in z and has size $s := |\mathcal{Q}'| \ll_{d,r} 1$.

Since Proposition 5.22 also gives that $\mathcal{Q}'$ descends from $\mathcal{P}$, the members of $\mathcal{Q}'$ are multilinear homogeneous polynomials of degree $d - 1$ (homogeneity follows from remark (4) below Definition 5.20). Letting $\underline{h} = (h_1, \dots, h_{r''})$ from now on with $r'' := r + r'$, we can enumerate the elements of $\mathcal{Q}'$ as

$$b_1(\underline{h})z, \dots, b_s(\underline{h})z,$$

where

$$b_j(\underline{h}) = \sum_{\underline{u} \in I_j} d! \beta_d \underline{h}^{\underline{u}}$$

and the sets $I_j \subseteq \{\underline{u} \in \{0, 1\}^{r''} : |\underline{u}| = d - 1\}$ for $j \in [s]$ are distinct and nonempty.

Applying the Cauchy-Schwarz and van der Corput Inequalities r' times to (5.54) much like¹⁸ in the proof of Proposition 5.23, we deduce the bound

$$\mathbb{E}_{\underline{h} \in [\pm H]^{r''}} \|f\|_{b_1(\underline{h}) \cdot [\pm H], (b_1(\underline{h}) - b_2(\underline{h})) \cdot [\pm H], \dots, (b_1(\underline{h}) - b_s(\underline{h})) \cdot [\pm H]}^{2^s} \gg_{d,D,\ell} \delta^{O_{d,\ell}(1)} N^D,$$

valid for all $H \ll_{d,\ell} \delta^{O_{d,\ell}(1)} K$. Theorem 5.10 combined with Lemma 5.15(5) gives a positive integer $t =$

¹⁸Technically, the expression (5.54) differs from the one in Proposition 5.23 in that in (5.54), we have an extra averaging over $h_1, \dots, h_r$ upfront. However, the argument for obtaining the bound in terms of an average of box norms for the averaged expression in (5.54) goes exactly the same way as in Proposition 5.23. Check e.g. the proof of [69, Corollary 4.3] starting from (14) to see why such averaged counting operators are no harder to analyze than ones without an additional average over $h_1, \dots, h_r$ upfront. The applications of the Cauchy-Schwarz and van der Corput Inequalities give us new shifting parameters $h_{r+1}, \dots, h_{r''}$; hence the polynomial parameters $\underline{h} = (h_1, \dots, h_{r''})$ come from combining the “old” $h_1, \dots, h_r$ with the “new” $h_{r+1}, \dots, h_{r''}$.

$O_{d,r}(1)$ for which

$$\|f\|_{U^t(d!\beta_d \cdot [\pm H^d])}^{2^t} \gg_{d,r} \delta^{O_{d,r}(1)} N.$$

Properties (3) and (7) of Lemma 5.15 allow us to replace the progression $d!\beta_d \cdot [\pm H^d]$ with $\beta_d \cdot [\pm N/\beta_d]$ upon taking $H \asymp_{d,\ell} \delta^{O_{d,\ell}(1)} K \asymp_{d,\ell} \delta^{O_{d,\ell}(1)} N/V$, giving the claimed result.

□

Chapter 6

Polynomial corners

6.1 Introduction

In recent years, there has been a surge of interest in obtaining “reasonable” upper bounds for the Polynomial Szemerédi Theorem of Bergelson and Leibman [9]. This theorem asserts that if $\mathbf{P}_1, \dots, \mathbf{P}_\ell \in \mathbb{Z}^D[z]$ are polynomials satisfying some mild local conditions, then every subset of $\mathbb{Z}^D$ with positive upper Banach density must contain a (nontrivial) occurrence of the polynomial progression

$$\mathbf{x}, \mathbf{x} + \mathbf{P}_1(z), \dots, \mathbf{x} + \mathbf{P}_\ell(z).$$

This result provided a unified proof for several earlier results in arithmetic Ramsey theory over $\mathbb{Z}$, including Roth’s pioneering theorem [93] on sets avoiding 3-term arithmetic progressions, Szemerédi’s later generalization [106] to sets avoiding ℓ -term arithmetic progressions for arbitrary fixed $\ell \geq 3$, and work of Sárközy [96, 97] and Furstenberg [40] (independently) on sets avoiding length-2 polynomial progressions.

The Polynomial Szemerédi Theorem, despite its useful generality, comes with a drawback: Because the proof uses methods from ergodic theory, the result is ineffective and provides no bounds on sets avoiding polynomial progressions. Roth and Sárközy proved their theorems with fairly good bounds. Szemerédi provided bounds of tower-type which, although strictly-speaking “effective”, are not particularly “reasonable”. The work of Gowers on obtaining improved bounds for Szemerédi’s Theorem [42, 43] gave birth to the field of higher-order Fourier analysis, and since then there has been substantial work on obtaining quantitative versions of various instances of the Polynomial Szemerédi Theorem.

Of particular interest is an early result due to Shkredov [102, 103], who obtained reasonable bounds for

subsets of $[N]^2$ avoiding the “corner” configuration

$$(x, y), (x + z, y), (x, y + z).$$

His argument involved a suitable notion of “pseudorandom sets” with which one can run a modification of Roth’s original density-increment strategy. In the context of corner-free subsets of $\mathbb{F}_p^D \times \mathbb{F}_p^D$, Green [46, 47] later gave a simplified exposition of Shkredov’s approach; see also work of Lacey and McClain [71] for an improvement in $\mathbb{F}_2^D \times \mathbb{F}_2^D$. Peluse’s recent work on subsets of $\mathbb{F}_p^D \times \mathbb{F}_p^D$ avoiding the “ L -shaped configuration” $(x, y), (x + z, y), (x, y + z), (x, y + 2z)$ built on some of the ideas from Shkredov’s approach.

Until recently, Shkredov’s work on corners was the only instance of a “reasonable bound” for a multidimensional configuration over the integers. Our main result is a polynomial extension of Shkredov’s theorem.

Theorem 6.1. *Let $P \in \mathbb{Z}[z]$ be a polynomial with an integer root of multiplicity 1. If $A \subseteq [N]^2$ does not contain a configuration of the form $(x, y), (x + P(z), y), (x, y + P(z))$ with $P(z) \neq 0$, then*

$$|A| \ll_P \frac{N^2}{(\log \log \log N)^c}$$

for some absolute constant $c > 0$.

Theorem 6.1 and the concurrently released work of Peluse, Prendiville and Shao [87] on the configuration $(x, y), (x + z, y), (x, y + z^2)$ provide the first examples of reasonable bounds for multidimensional polynomial (i.e., non-linear) progressions over the integers.

Using a calculation of Shkredov [102], we can take any $c < 1/73$. The reader may notice that our statement has an artificial-looking constraint on the polynomial P . This constraint, which is slightly stronger than intersectivity, arises from a technical point related to the use of the W -trick, as discussed later.

Note added in revision: Jaber, Liu, Lovett, Ostuni, and Swahney [59] have recently obtained a quantitative improvement of Shkredov’s result. In particular, their results imply that there is some constant $c > 0$ such that every corner-free subset of $[N]^2$ has size at most $N^2 / \exp((\log N)^c)$. By using this input in place of Shkredov’s bound [102], we can improve the estimate in Theorem 6.1 to $A \ll_P N^2 / \exp((\log \log N)^c)$ for some absolute constant $c > 0$. This change in the argument can be carried through quite locally: The (very easy) updates are necessary only in the statement and application of Proposition 6.60.

A key analytic step in our argument is the following inverse theorem for polynomial corners, which may be of independent interest.

Theorem 6.2 (Inverse theorem for polynomial corners). *Let $d \in \mathbb{N}$, let $P \in \mathbb{Z}[z]$ be a degree- d polynomial with leading coefficient β_d , let $\delta \in (0, 1/10)$, and let $f_0, f_1, f_2 : \mathbb{Z}^2 \rightarrow \mathbb{C}$ be 1-bounded functions supported on*

$[N]^2$ for some integer $N \geq \exp(\log(1/\delta)^{\Omega_P(1)})$. If

$$\left| \sum_{x,y} \mathbb{E}_{z \in [(N/\beta_d)^{1/d}]} f_0(x,y) f_1(x+P(z),y) f_2(x,y+P(z)) \right| \geq \delta N^2,$$

then we simultaneously have

$$\begin{aligned} \sum_{x,y} f_0(x,y) b_1(x) b_2(y) &\geq \exp(-\log(1/\delta)^{O_P(1)}) N^2 \quad \text{for some 1-bounded } b_1, b_2 : \mathbb{Z} \rightarrow \mathbb{C}, \\ \sum_{x,y} f_1(x,y) b_1(x+y) b_2(y) &\geq \exp(-\log(1/\delta)^{O_P(1)}) N^2 \quad \text{for some 1-bounded } b_1, b_2 : \mathbb{Z} \rightarrow \mathbb{C}, \\ \sum_{x,y} f_2(x,y) b_1(x) b_2(x+y) &\geq \exp(-\log(1/\delta)^{O_P(1)}) N^2 \quad \text{for some 1-bounded } b_1, b_2 : \mathbb{Z} \rightarrow \mathbb{C}. \end{aligned}$$

In the proof of Theorem 6.1, we will end up using a slightly different version of this inverse theorem that is better suited to W -tricks (Theorem 6.10); the proof simplifies in several ways for the W -trick-free version stated here. Theorem 6.10 leads to a “comparison theorem” with a “model” counting operator (Theorem 6.20), which in Section 6.7 we will use to deduce the following improvement on the result of Peluse, Sah, and Sawhney [88].

Corollary 6.3. *Let $P \in \mathbb{Z}[z]$ be a polynomial with an integer root of multiplicity 1. If $A \subseteq [N]$ does not contain a configuration of the form $x, x+P(z), x+2P(z)$ with $P(z) \neq 0$, then*

$$|A| \ll \frac{N}{\exp((\log \log N)^{c_P})}$$

for some constant $c_P > 0$.

Before we say more about our main theorem, let us briefly survey related results. Previous work has focused on 1-dimensional polynomial progressions over the integers and both 1-dimensional and multidimensional polynomial progressions in finite-field settings. One of the many important differences between the integers and finite fields is that for a fixed polynomial $P \in \mathbb{Z}[z]$ of degree at least 2, the image of P is an equidistributed set of positive density in (a sequence of) finite fields, while in the integers it is a sparse set with local biases; since there are more “opportunities” to find polynomial progressions in finite fields $\mathbb{F}_p$, it is generally easier to prove bounds on the sizes of sets avoiding polynomial progressions in finite-field settings, and much more is known here than over the integers.

Early results for 1-dimensional polynomial progressions in $\mathbb{F}_p$ (p a prime) are due to Bourgain and Chang [20], Dong, Li, and Sawin [31], and Peluse [84]. In a major breakthrough, Peluse [83] proved power-saving

upper bounds for subsets of $\mathbb{F}_p$ avoiding the polynomial progression

$$x, x + P_1(z), \dots, x + P_\ell(z)$$

for any given linearly independent polynomials $P_1, \dots, P_m \in \mathbb{Z}[z]$ with constant coefficients all zero. Her proof introduced the now-well-used technique of degree-lowering for Gowers norms. Peluse’s work has been generalized in several ways. For instance, Kuca [67] proved reasonable upper bounds for subsets of $\mathbb{F}_p$ avoiding a given progression of the form

$$x, x + z, \dots, x + (\ell - 1)z, x + z^\ell, \dots, x + z^{\ell+k-1}.$$

Later on, Leng [72] developed a U^3 -version of degree-lowering to obtain reasonable upper bounds for “complexity-1” configurations

$$x, x + P(z), x + Q(z), x + P(z) + Q(z).$$

Lastly, Bergelson and Best [8] generalized Peluse’s result to general finite commutative rings.

Less is known about multidimensional progressions in finite fields. Regarding linear configurations, we have already mentioned results on corners and L-shaped configurations. In the higher-degree case, Kuca [68, 69] proved a general Szemerédi-type theorem for patterns of the form

$$\mathbf{x}, \mathbf{x} + \mathbf{v}_1 P_1(z), \dots, \mathbf{x} + \mathbf{v}_\ell P_\ell(z),$$

where $P_1, \dots, P_\ell$ are linearly independent polynomials with constant coefficients all zero and $\mathbf{v}_1, \dots, \mathbf{v}_\ell$ are arbitrary vectors in $\mathbb{F}_p^D$. When $\ell = 2$, special cases of this result had been established earlier by Han, Lacey, and Yang [54].

Let us now turn to polynomial progressions in $\mathbb{Z}$. In a sequence of papers, Peluse [81] as well as Peluse and Prendiville [85, 86] adapted Peluse’s original degree-lowering argument for progressions of the form $x, x + P_1(z), \dots, x + P_m(z)$ where the polynomials P_i have distinct degrees. Prendiville [90] also proved reasonable bounds for sets avoiding arithmetic progressions with common difference of the form z^d for fixed d ; the homogeneity of the polynomial z^d allowed him to use a density-increment argument. Peluse, Sah, and Sawhney [88], adapting a U^3 -degree-lowering argument developed in [72], addressed the case of three-term arithmetic progressions with common difference $z^2 - 1$. They proved a “transference” result relating the number of 3-term arithmetic progressions to the number of 3-term arithmetic progressions with common

difference $z^2 - 1$.

By a simple projection argument (see the proof of Corollary 6.3 below), Theorem 6.1 contains the setting studied by Peluse, Sah, and Sawhney as a special case. Hence, our main theorem can be understood as a common generalization of Shkredov's corners theorem and the result of Peluse, Sah, and Sawhney. To prove Theorem 6.1, we will establish a transference result for corners, in the style of [88], and then use Shkredov's work as a black box. The starting point for our argument is a PET and quantitative concatenation result carried out in the previous chapter. The bulk of the present chapter is devoted to obtaining degree-lowering for the polynomial corners configuration. Our degree-lowering argument is the first to use box norms instead of Gowers norms, and its success attests to the robustness of degree-lowering as a tool in arithmetic combinatorics.

6.1.1 Future directions

Our results and methods suggest several avenues for future inquiry. The most natural open problem is extending Theorem 6.1 to a larger class of polynomials P . A particularly appealing example, recently highlighted by Peluse [82, Problem 3.17], is $P(z) = z^2$. (Theorem 6.1 fails to cover $P(z) = z^2$ because 0, the only integer root, has multiplicity $2 > 1$.) One strategy, which we attempted initially before turning to transference arguments, is obtaining $U^1 \times U^1$ -control for the P -polynomial corner counting operator and then running Shkredov's original density-increment argument for corners. Our degree-lowering argument goes through with $\tilde{P}$ replaced by an arbitrary (non- W -tricked) polynomial $P \in \mathbb{Z}[z]$, which gives the norm-control necessary for starting Shkredov's iterative argument. The pitfall of this approach, however, is that in later stages of the iteration one must deal with subsets of pseudorandom sets rather than subsets of $[N]^2$, and this creates complications in both the arguments of this chapter and the PET and concatenation results of the previous chapter. It seems possible that one could generalize our arguments to the setting of pseudorandom sets as in [102, 103], but we have not pursued this.

The polynomial progressions investigated in this chapter belong to the class of linear patterns (in our case, corners) with polynomial differences. Other works on obtaining bounds for such configurations have proceeded either by modifying arguments for the corresponding linear patterns (as in, e.g., [90]) or by using transference results to relate the counts of linear and polynomial patterns (as in the present work and [67, 88]). Reasonable bounds are known for L -shaped configurations and arithmetic progressions of length $\ell \geq 4$, but their polynomial analogues are currently unknown, so it would be interesting to extend our methods to establish transference results for these two particular configurations. It would also be interesting to use this machinery to obtain more general transference principles, even for cases where there are not currently

reasonable bounds for linear configurations.

6.2 Proof strategy

We now give an overview of the proof of Theorem 6.1. Our overarching goal is to relate the number of polynomial corners to the number of ordinary corners; once this is accomplished, we can “transfer” Shkredov’s result (in a standard supersaturation version) from ordinary corners to polynomial corners.

Let $A \subseteq [N]^2$ be a set of density δ , and let $P \in \mathbb{Z}[z]$ be a polynomial of degree d with an integer root ρ of multiplicity 1. Then

$$P(z - \rho) = \beta_d z^d + \cdots + \beta_1 z$$

for some $\beta_1, \dots, \beta_d \in \mathbb{Z}$ with $\beta_1, \beta_d \neq 0$; without loss of generality we may assume that $\beta_d > 0$. The number of corners in A is $N\Lambda^{\text{Corners}}(1_A, 1_A, 1_A)$, where Λ^{Corners} is the normalized corner-counting operator at scale N given by

$$\Lambda^{\text{Corners}}(f_0, f_1, f_2) := \sum_{x, y \in \mathbb{Z}} \sum_{z \in [N]} \mathbb{E} f_0(x, y) f_1(x + z, y), f_2(x, y + z).$$

The number of P -polynomial corners in A is

$$(N/\beta_d)^{1/d} \Lambda(1_A, 1_A, 1_A),$$

where Λ is the normalized P -polynomial corner-counting operator at scale N given by

$$\Lambda(f_0, f_1, f_2) := \sum_{x, y \in \mathbb{Z}} \sum_{z \in [(N/\beta_d)^{1/d}]} \mathbb{E} f_0(x, y) f_1(x + P(z), y) f_2(x, y + P(z)).$$

In order to show that $|\Lambda^{\text{Corners}} - \Lambda|$ is small, we must make both an Archimedean and a non-Archimedean modification.

First, since the set $\{P(z) : z \in \mathbb{N}\}$ becomes sparser at larger scales, we introduce a similarly-decaying weight

$$\nu(z) := d^{-1} 1_{[N]}(z) (N/(z+1))^{(d-1)/d}.$$

into Λ^{Corners} to obtain the model corner-counting operator

$$\Lambda^{\text{Model}}(f_0, f_1, f_2) := \sum_{x, y \in \mathbb{Z}} \sum_{z \in [N]} \mathbb{E} f_0(x, y) f_1(x + z, y) f_2(x, y + z) \nu(z).$$

Second, since the set $\{P(z) : z \in \mathbb{N}\}$ may be biased with respect to small moduli, we “pre-sieve” the set A before attempting to compare the counts of corners and of P -polynomial corners. More precisely, for W a very slowly-growing function of N , we restrict our attention to the elements of A lying in a single coset of $(\beta_1^2 W \mathbb{Z})^2$; by the pigeonhole principle, there is some such coset $(b_1 + \beta_1^2 W \mathbb{Z}) \times (b_2 + \beta_1^2 W \mathbb{Z})$ in which A has relative density at least δ (at scale N). Let

$$A' := \{(x', y') \in \mathbb{Z}^2 : (\beta_1^2 W x' + b_1, \beta_1^2 W y' + b_2) \in A\} \subseteq [N/(\beta_1^2 W)]^2$$

be this “slice” of A . We claim that in order to find a P -polynomial corner in A , it suffices to find a $\tilde{P}$ -polynomial corner in A' , where $\tilde{P}$ is the W -tricked auxiliary polynomial

$$\tilde{P}(z) := \frac{P(\beta_1 W z - \rho)}{\beta_1^2 W} = \beta_d \beta_1^{d-2} W^{d-1} z^d + \beta_{d-1} \beta_1^{d-3} W^{d-2} z^{d-1} + \cdots + z \in \mathbb{Z}[z]. \quad (6.1)$$

(The linear coefficient 1 of $\tilde{P}$ will be convenient later for technical reasons.) Indeed, if

$$(x', y'), (x' + \tilde{P}(z'), y'), (x' + \tilde{P}(z'), y' + \tilde{P}(z'))$$

form a $\tilde{P}$ -polynomial corner in A' , then A contains the P -polynomial corner

$$(x, y), (x + P(z), y), (x, y + P(z))$$

with

$$x := \beta_1^2 W x' + b_1, \quad y := \beta_1^2 W y' + b_2, \quad \text{and} \quad z := \beta_1 W z' - \rho.$$

Hence, we replace Λ with the W -tricked counting operator

$$\Lambda^W(f_0, f_1, f_2) := \sum_{x, y} \mathbb{E}_{z \in [(N/W_d)^{1/d}]} f_0(x, y) f_1(x + \tilde{P}(z), y) f_2(x, y + \tilde{P}(z)),$$

where $W_d := \beta_d \beta_1^{d-2} W^{d-1}$ is the leading coefficient of $\tilde{P}$. The main work of this chapter is showing that, with these modifications made, the counting operators Λ^{Model} and Λ^W are very close (Proposition 6.20).

This comparison theorem boils down to showing that both Λ^{Model} and Λ^W are “controlled” by appropriate $U^1 \times U^1$ -box norms of their arguments. The control for Λ^{Model} (Proposition 6.19) is fairly straightforward and follows from an argument similar to the proof of [88, Lemma 4.1], so we will focus on the control for Λ^W . We wish to show that for all 1-bounded functions f_0, f_1, f_2 supported on $[N]^2$, the largeness of $|\Lambda^W(f_0, f_1, f_2)|$

implies (with quantitative control) the largeness of a box norm of f_0 in the directions $\mathbf{e}_1 := (1, 0)$ and $\mathbf{e}_2 := (0, 1)$. The control in terms of f_1 (with the directions $\mathbf{e}_1, \mathbf{e}_1 - \mathbf{e}_2$) and f_2 (with the directions $\mathbf{e}_2, \mathbf{e}_2 - \mathbf{e}_1$) is completely symmetrical. The precise statement of this result is in Theorem 6.10, which constitutes a W -tricked version of Theorem 6.2.

The starting point of this endeavor is the PET induction scheme and a quantitative concatenation result, as carried out in the previous chapter. The output of these arguments is the existence of a natural number $t = O_d(1)$ such that the lower bound $|\Lambda^W(f_0, f_1, f_2)| \geq \delta N^2$ implies that

$$\|f_0\|_{(\mathbf{e}_1 W_d \cdot [\pm N/W_d])^t, (\mathbf{e}_2 W_d \cdot [\pm N/W_d])^t}^{2^{2t}} \gg_d \delta^{O_d(1)} N^2, \quad (6.2)$$

and likewise for f_1, f_2 (cf. Proposition 6.11). For simplicity of exposition, we will temporarily pretend that $W = 1$; at the end of this sketch, we will mention how to deal with the additional technical complications arising from the presence of W . For notational simplicity, we will also make the unimportant assumption that $\beta_d = \beta_1 = 1$ (so that $W_d = 1$ and the z -variable ranges over $[N^{1/d}]$ in the definition of Λ^W).

In order to replace t with 1, we will carry out an iterative degree-lowering procedure (see Proposition 6.12 below): We will show that if the largeness of $|\Lambda^W(f_0, f_1, f_2)|$ is already known to imply the largeness of $\|f_0\|_{(\mathbf{e}_1 \cdot [\pm N])^k, (\mathbf{e}_2 \cdot [\pm N])^\ell}^{2^{k+\ell}}$ for some k, ℓ with $k \geq 2$, then it also implies the largeness of $\|f_0\|_{(\mathbf{e}_1 \cdot [\pm N])^{k-1}, (\mathbf{e}_2 \cdot [\pm N])^\ell}^{2^{k+\ell-1}}$ (and likewise with the roles of k, ℓ reversed). We will arrive at our goal after $2t - 2$ iterations of this degree-lowering procedure. This is the first time that degree-lowering has been carried out using box norms; all previous degree-lowering arguments in the literature aim at lowering degrees of Gowers norms. Our argument thus testifies to the versatility of the degree-lowering approach pioneered by Peluse [83].

By the dual-difference interchange (a consequence of the Cauchy-Schwarz inequality; see Lemma 6.9), it suffices to obtain our degree-lowering result when $k = 2$ and $\ell = 1$. Suppose that $|\Lambda^W(f_0, f_1, f_2)| \geq \delta N^2$. We begin by applying the “stashing” trick first introduced by Manners [77]: Since f_0 is supported on $[N]^2$, using the Cauchy-Schwarz inequality to eliminate f_0 from Λ^W gives

$$\delta^2 N^4 \leq |\Lambda^W(f_0, f_1, f_2)|^2 \leq N^2 \Lambda^W(\mathcal{D}_0(f_1, f_2), f_1, f_2),$$

where $\mathcal{D}_0(f_1, f_2) : \mathbb{Z}^2 \rightarrow \mathbb{C}$, the *dual function* of the first argument of Λ^W , is given by

$$\mathcal{D}_0(f_1, f_2)(x, y) := \mathbb{E}_{z \in [N^{1/d}]} f_1(x + \tilde{P}(z), y) f_2(x, y + \tilde{P}(z)).$$

The hypothesis of the degree-lowering statement tells us that

$$\|\mathcal{D}_0(f_1, f_2)\|_{\mathbf{e}_1 \cdot [\pm N], \mathbf{e}_1 \cdot [\pm N], \mathbf{e}_2 \cdot [\pm N]}^8$$

is large (i.e., at least $c(\delta)N^2$ for some explicit $c(\delta) > 0$). Then the inverse theorem for the $U^2 \times U^1$ -norm (see Lemma 6.28 below) implies that $\mathcal{D}_0(f_1, f_2)$ correlates with a “ $U^2 \times U^1$ -obstruction” of the form $g(x)e(a(y)x + b(y))$, where $g : \mathbb{Z} \rightarrow \mathbb{C}$ is a 1-bounded function, $a : \mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}$ is a phase, and $b : \mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}$ is a shift.¹ Unraveling the definition of $\mathcal{D}_0(f_1, f_2)$ and ignoring the terms $g(x), b(y)$ (which are of secondary importance), we find that

$$\Lambda'(a, f_1, f_2) := \sum_{x, y} \mathbb{E}_{z \in [N^{1/d}]} e(a(y)x) f_1(x + \tilde{P}(z), y) f_2(x, y + \tilde{P}(z))$$

is large in absolute value.

We now aim to show that the phase a is close to some constant α for many values of y (see Proposition 6.13 below); this step “lowers” the degree-2 obstruction $e(a(y)x)$ to the degree-1 obstruction $e(\alpha x)$. Before sketching the proof of this structural result for a , let us see how to use it to complete the degree-lowering procedure. Substituting α for $a(y)$ and shifting $x \mapsto x - \tilde{P}(z)$ in the expression for $\Lambda'(a, f_1, f_2)$, we find that

$$\sum_{x, y} \mathbb{E}_{z \in [N^{1/d}]} e(\alpha(x - \tilde{P}(z))) f_1(x, y) f_2(x - \tilde{P}(z), y + \tilde{P}(z))$$

is large. Grouping the exponential with f_2 , we recognize this expression as an instance of the counting operator for the 2-point configuration $(x, y), (x - \tilde{P}(z), y + \tilde{P}(z))$. Adapting a classical Fourier-analytic argument of Sárközy [96, 97] for 2-point polynomial configurations, we conclude that

$$\|f_1\|_{(\mathbf{e}_2 - \mathbf{e}_1) \cdot [\pm N], \mathbf{e}_2 \cdot [\pm N]}^4$$

is large. A further sequence of manipulations in the flavor of stashing let us transfer this “degree-2 control” from f_1 back to f_0 .

It remains to analyze the phase function a . It turns out that it suffices to understand the case where f_2 is a “ $U^1 \times U^1$ -obstruction” of the form $b_1(y)b_2(x + y)$ for some 1-bounded functions $b_1, b_2 : \mathbb{Z} \rightarrow \mathbb{C}$, in which

¹To understand where this inverse theorem comes from, notice that differentiating once in the y -direction removes the $g(x)$ term and differentiating twice in the x -direction removes the $e(a(y)x + b(y))$ term.

case we are dealing with the expression

$$\sum_{x,y} \mathbb{E}_{z \in [N^{1/d}]} e(a(y)x) f_1(x + \tilde{P}(z), y) b_1(y + \tilde{P}(z)) b_2(x + y + \tilde{P}(z)).$$

Shifting $x \mapsto x - \tilde{P}(z)$, we reach an expression in which the functions f_1 and b_2 are independent of z . Then, by applying the triangle inequality in x, y and pigeonholing in x , we deduce that

$$\sum_{y \in [N]} \left| \mathbb{E}_{z \in [N^{1/d}]} e(-a(y)\tilde{P}(z)) b_1(y + \tilde{P}(z)) \right|$$

is large. A variant of the PET argument from the previous chapter lets us control this expression by a Gowers norm of b_1 , and the inverse theorem for the Gowers norms produces a nilsequence φ of bounded degree, dimension, and complexity such that

$$\sum_{y \in [N]} \left| \mathbb{E}_{z \in [N^{1/d}]} e(a(y)\tilde{P}(z)) \varphi(y + \tilde{P}(z)) \right| \tag{6.3}$$

is large. To analyze (6.3), we require input from the theory of nilsequences.

We use an iterative process, itself a microcosm of degree-lowering, to show that the nilsequence in (6.3) can be replaced by nilsequences of progressively smaller step (see Proposition 6.18). This argument makes use of recent quantitative equidistribution results for polynomial sequences on nilmanifolds due to Leng [73]. At the end of this iteration, we have replaced $\varphi(y)$ by a step-1 nilsequence, namely, a function of the form $e(p(y))$ for some polynomial p . We then use Weyl's inequality to reduce to the case where $p(y) = \alpha_1 y + \alpha_0$ is linear. Another application of Weyl's inequality shows that there are many values of y such that $a(y) - \alpha_1$ is major-arc, i.e., close to a rational of small height; pigeonholing in these rational shifts finally gives that $a(y)$ is nearly constant for many y 's.

Figure 6.1 provides a visual overview of the main steps of our argument.

Our degree-lowering arguments must be modified in several ways to account for the parameter W . Recall that the initial norm-control in (6.2) involves differencing with respect to boxes with common difference W (rather than 1). In order to manipulate such expressions and use inverse theorems for box norms, we must split all of the parameters in sight (e.g., x, y, z) into arithmetic progressions with common difference W . This maneuver, which often requires introducing iterated W -tricks, is quite a nuisance but does not change the essential nature of our arguments.

More importantly, the miniature PET induction in the lead-up to (6.3) yields Gowers-norm control where the differencing parameters are boxes with side lengths equal to the leading coefficient of $\tilde{P}$. Since this leading

coefficient is a constant multiple of a fixed positive integer power of W , this step of the argument requires us to replace W with some $W^{O(1)}$, and we must keep track of the accumulation of these exponents in all of the iterations of the degree-lowering. (We also have to perform $W^{O(1)}$ -tricks instead of W -tricks.)

We conclude with a short discussion of the shape of bounds in Theorem 6.1. Of the three logarithms in the bound in Theorem 6.1, two come from the supersaturation version of Shkredov's result [102, 103], and the third arises in the transference step of our argument. That we lose only one logarithm in comparing the counts of polynomial and linear corners is reliant on a recent equidistribution result for nilsequences due to Leng [73] and the quasipolynomial inverse theorem of Gowers norms [75] (which is itself reliant on [73]). If we instead used bounds for equidistribution of nilsequences provided by Tao and Teräväinen [109] (who quantified earlier equidistribution results of Green and Tao [51]), we would end up with a large number of iterated logarithms bounded in terms of the degree of the polynomial P .

6.2.1 Organization




Figure 6.1: The main steps and logical implications in the proof of Theorem 6.1.

After recalling definitions and preliminaries in Section 6.3, we embark on the proof of the main technical result of our work, Theorem 6.10, which asserts that the counting operator Λ^W is controlled by suitable

degree-2 box norms. In Section 6.4, we show how to use our degree-lowering result, Proposition 6.12, to deduce Theorem 6.10 from the PET and quantitative concatenation results of the previous chapter. We also reduce Proposition 6.12 to Proposition 6.13, an inverse theorem for the auxiliary operator Λ' . The lion's share of the proof of Proposition 6.13 is then carried out in Section 6.5. The key input in this argument is Proposition 6.18, a structural result on (6.3), which we prove in Section 6.6 using nilsequence machinery.

The arguments in Sections 6.4–6.6 make use of several standard results which we collect as appendices in the final six sections of this chapter: inverse theorems for various box norms in Section 6.8; a summary of nilsequence theory in Section 6.9; a fraction comparison argument in Section 6.10; and a Sárközy-style inverse theorem for a 2-point polynomial progressions (derived with circle-method techniques) in Section 6.11.

We conclude the main body of the chapter with Section 6.7, where we finally complete the proof of Theorem 6.1. Proposition 6.19 establishes the desired box-norm control for Λ^{Model} , and we compare Λ^W and Λ^{Model} in Proposition 6.20. The latter proposition requires two more ingredients from the appendiceal sections: a Fourier-uniformity result for certain polynomial weights from Section 6.12; and a supersaturation version of Shkredov's corners theorem from Section 6.13.

6.3 Definitions and preliminaries

We recall some basic notation and foundation results, along the lines of what we used in the previous chapter.

6.3.1 Basic notation

We write $\mathbb{N}, \mathbb{Z}, \mathbb{R}, \mathbb{C}$ for the sets of positive integers, integers, real numbers, and complex numbers (respectively). For reals $a < b$, we set $[a, b] := \{[a], a + 1, \dots, [b]\}$, and we abbreviate $[0, N - 1]$ as $[N]$ and $[-N + 1, N - 1]$ as $[\pm N]$. We remark that our somewhat unusual convention for $[N]$ (the norm would be $[N] = [1, N]$) behaves well under passing to arithmetic progressions: For $q|N$, we have the partition $[N] = \bigcup_{j \in [q]} (q \cdot [N/q] + j)$.

We write $\mathbf{x} = (x_1, \dots, x_D)$ for an element of $\mathbb{Z}^D$ (for fixed $D \in \mathbb{N}$) and x for an element of $\mathbb{Z}$. When the context is clear, we often write a sum over $\mathbb{Z}^D$ or $\mathbb{Z}$ as $\sum_{\mathbf{x}}$ or $\sum_x$ instead of $\sum_{\mathbf{x} \in \mathbb{Z}^D}$ or $\sum_{x \in \mathbb{Z}}$. We also write $\underline{h} = (h_1, \dots, h_s)$ for a tuple of s integers, and we denote its ℓ^1 -norm by $|\underline{h}| := |h_1| + \dots + |h_s|$. We will often combine these two pieces of notation; we use $\underline{h}\mathbf{v}$ to indicate the tuple $(h_1\mathbf{v}, \dots, h_s\mathbf{v})$.

We denote the indicator function of a set X by 1_X , and the indicator function of an event E by $1(E)$. If X is finite and nonempty, we let $\mathbb{E}_{x \in X} := \frac{1}{|X|} \sum_{x \in X}$ be the average over X .

We call a function $f : X \rightarrow \mathbb{C}$ *1-bounded* if $\sup_{x \in X} |f(x)| \leq 1$.

We use the following standard asymptotic notation. Let $f, g : X \rightarrow \mathbb{C}$, with g taking positive real values. We write $f = O(g)$, $f \ll g$, $g \gg f$, or $g = \Omega(f)$ if there exists $C > 0$ such that $|f(x)| \leq Cg(x)$ for all $x \in X$, and we write $f \asymp g$ or $f = \Theta(g)$ if $f \ll g$ and $f \gg g$. If the constant C depends on a parameter, we record this dependence with a subscript, e.g., $f = O_P(g)$, $f \ll_P g$.

We let $\mathcal{C} : \mathbb{C} \rightarrow \mathbb{C}$ denote the complex conjugation operator $\mathcal{C}z = \bar{z}$.

6.3.2 W -trick notation

Throughout this chapter, we let $w = w(N) > 10$ be a (fixed) slowly-growing function of N ; we will later take $w(N)$ to be a small constant (depending on P) times $\log N$. Let

$$W = W(N) := \prod_{\substack{p < w \\ p \text{ prime}}} p$$

be the product of the primes up to w . Standard results on the growth rate of Chebyshev's function give $\log W \asymp w$.

It will often be useful to restrict functions to arithmetic progressions, which we capture in the following definition.

Definition 6.4 (*V-tricked function*). For $V \in \mathbb{N}$ and $r \in \mathbb{Z}$, we can parameterize the arithmetic progression $V\mathbb{Z} + r \subseteq \mathbb{Z}$ by $\{Vx + r : x \in \mathbb{Z}\}$. If f is a function on $\mathbb{Z}$, then its restriction to $V\mathbb{Z} + r$, with this parametrization, is

$$f_{(r,V)}(x) := f(Vx + r).$$

Likewise, for $r_1, r_2 \in \mathbb{Z}$ and f a function on $\mathbb{Z}^2$, the restriction of f to $(V\mathbb{Z} + r_1) \times (V\mathbb{Z} + r_2)$ can be parameterized as

$$f_{(r_1, r_2, V)}(x, y) := f(Vx + r_1, Vy + r_2).$$

We refer to the operation of passing to such a subprogression as a *V-trick*, and we say that the resulting functions are *V-tricked*.

Throughout this chapter, we will usually take V to be a small (nonnegative integer) power of W and r, r_1, r_2 to be elements of $[V]$. We remark that a composition of two such *V-tricks* is itself a *V-trick*. To see this symbolically for $f : \mathbb{Z} \rightarrow \mathbb{C}$, one can calculate that

$$(f_{(r_1, V_1)})_{(r_2, V_2)}(x) = f_{(r_1, V_1)}(V_2x + r_2)$$

$$\begin{aligned}
&= f(V_1(V_2x + r_2) + r_1) \\
&= f_{(V_1r_2+r_1, V_1V_2)}(x),
\end{aligned}$$

where V_1V_2 is a power of W if V_1, V_2 are. The case of functions on $\mathbb{Z}^2$ is analogous.

Definition 6.5 (*V-tricked polynomial*). For V -tricks of polynomials, we will use slightly different notation that will turn out to be helpful for keeping track of divisibility properties of various coefficients. Let $Q \in \mathbb{Z}[z]$ be a polynomial. For $V \in \mathbb{N}$ and $r \in \mathbb{Z}$, define the V -tricked polynomial of Q to be

$$Q_{[r,V]}(z) := \frac{Q(Vz + r) - Q(r)}{V} \in \mathbb{Z}[z].$$

That $Q_{[r,V]}$ has integer coefficients can be easily verified using the Binomial Theorem.

As in the case of V -tricks for functions, the composition of two V -tricks for a polynomial is again a V -trick. In particular, we have

$$\begin{aligned}
(Q_{[r_1,V_1]})_{[r_2,V_2]}(z) &= \frac{Q_{[r_1,V_1]}(V_2z + r_2) - Q_{[r_1,V_1]}(r_2)}{V_2} \\
&= \frac{Q(V_1(V_2z + r_2) + r_1) - Q(r_1) - Q(V_1r_2 + r_1) + Q(r_1)}{V_1V_2} \\
&= Q_{[V_1r_2+r_1, V_1V_2]}(z).
\end{aligned}$$

Note that iterated V -tricks (in both the function setting and the polynomial setting) do not in general commute.

6.3.3 Van der Corput Inequalities

If $E \subseteq \mathbb{Z}^D$ is finite and nonempty, we define the *Fejér kernel*

$$\mu_E(\mathbf{h}) := \mathbb{E}_{\mathbf{h}_1, \mathbf{h}_2 \in E} 1_{\mathbf{h}_1 - \mathbf{h}_2}(\mathbf{h}).$$

Note that $\sum_{\mathbf{h}} \mu_E(\mathbf{h}) = 1$ and that $0 \leq \mu_E(\mathbf{h}) \leq \frac{1_{E-E}(\mathbf{h})}{|E|}$ pointwise. If $E = [\pm H]$, we set

$$\mu_H(h) := \mu_{[\pm H]}(h) = \mathbb{E}_{h_1, h_2 \in [\pm H]} 1_{h_1 - h_2}(h) = \frac{1}{2H-1} \left(1 - \frac{|h|}{2H-1} \right)_+.$$

This definition differs slightly from those in [81, 85, 86, 90] in that we take the ranges of h_1, h_2 to be $[\pm H]$ rather than $[H]$. Making the range of h_1, h_2 symmetric around the origin will be convenient for us later.

We will use the following form of the standard van der Corput inequality, as discussed in the previous chapter.

Lemma 6.6 (The van der Corput inequality). *Let $\delta \in (0, 1]$, and let $H, N \in \mathbb{N}$ with $H \leq \frac{\delta^2}{4}N$. If $f : \mathbb{Z} \rightarrow \mathbb{C}$ is a 1-bounded function satisfying*

$$\left| \mathbb{E}_{x \in [N]} f(x) \right| \geq \delta, \quad (6.4)$$

then

$$\mathbb{E}_{h, h' \in [\pm H]} \mathbb{E}_{x \in [N]} f(x+h) \overline{f(x+h')} \geq \frac{\delta^2}{4}$$

and

$$\sum_h \mu_H(h) \mathbb{E}_{x \in [N]} f(x) \overline{f(x+h)} \geq \frac{\delta^2}{4}.$$

6.3.4 Box norms

We now briefly review the theory of box norms on $\mathbb{Z}^D$; we will later work with the groups $\mathbb{Z}$ and $\mathbb{Z}^2$. Given $\mathbf{h}, \mathbf{h}' \in \mathbb{Z}^D$ and a finitely supported function $f : \mathbb{Z}^D \rightarrow \mathbb{C}$, we define the (multiplicative) discrete derivatives $\Delta_{\mathbf{h}} f(\mathbf{x}) := f(\mathbf{x}) \overline{f(\mathbf{x} + \mathbf{h})}$ and $\Delta'_{(\mathbf{h}, \mathbf{h}')} f(\mathbf{x}) := f(\mathbf{x} + \mathbf{h}) \overline{f(\mathbf{x} + \mathbf{h}')}$. Similarly, compositions of such discrete derivatives are given by

$$\begin{aligned} \Delta_{\mathbf{h}_1, \dots, \mathbf{h}_s} f(\mathbf{x}) &= \Delta_{\mathbf{h}_1} \cdots \Delta_{\mathbf{h}_s} f(\mathbf{x}) = \prod_{\underline{\epsilon} \in \{0,1\}^s} \mathcal{C}^{|\underline{\epsilon}|} f(\mathbf{x} + \underline{\epsilon} \cdot \underline{\mathbf{h}}), \\ \Delta'_{(\mathbf{h}_1, \mathbf{h}'_1), \dots, (\mathbf{h}_s, \mathbf{h}'_s)} f(\mathbf{x}) &= \Delta'_{(\mathbf{h}_1, \mathbf{h}'_1)} \cdots \Delta'_{(\mathbf{h}_s, \mathbf{h}'_s)} f(\mathbf{x}) = \prod_{\underline{\epsilon} \in \{0,1\}^s} \mathcal{C}^{|\underline{\epsilon}|} f(\mathbf{x} + (\underline{1} - \underline{\epsilon}) \cdot \underline{\mathbf{h}} + \underline{\epsilon} \cdot \underline{\mathbf{h}'}), \end{aligned}$$

where $\underline{\mathbf{h}} = (\mathbf{h}_1, \dots, \mathbf{h}_s)$ and again $\mathcal{C}z = \bar{z}$ for $z \in \mathbb{C}$.

For finite sets $E_1, \dots, E_s \subseteq \mathbb{Z}^D$, we define the *box norm* of f along $E_1, \dots, E_s$ to be

$$\|f\|_{E_1, \dots, E_s} := \left(\mathbb{E}_{\mathbf{h}_1, \mathbf{h}'_1 \in E_1} \cdots \mathbb{E}_{\mathbf{h}_s, \mathbf{h}'_s \in E_s} \sum_{\mathbf{x}} \Delta'_{(\mathbf{h}_1, \mathbf{h}'_1), \dots, (\mathbf{h}_s, \mathbf{h}'_s)} f(\mathbf{x}) \right)^{1/2^s}.$$

We can equivalently express this as

$$\|f\|_{E_1, \dots, E_s} = \left(\sum_{\mathbf{x}, \mathbf{h}_1, \dots, \mathbf{h}_s} \mu_{E_1}(\mathbf{h}_1) \cdots \mu_{E_s}(\mathbf{h}_s) \Delta_{\mathbf{h}_1, \dots, \mathbf{h}_s} f(\mathbf{x}) \right)^{1/2^s} \quad (6.5)$$

after shifting $\mathbf{x} \mapsto \mathbf{x} - \mathbf{h}'_1 - \dots - \mathbf{h}'_s$ and removing the extraneous averages over $\mathbf{h}'_1, \dots, \mathbf{h}'_s$. For the sake of completeness, we also define $\|f\|_\emptyset := \sum_{\mathbf{x}} f(\mathbf{x})$ (but note that this is not a seminorm). If a set E repeats t times, we often write E^t instead of $E, \dots, E$. We also write $\|f\|_{U^s(E)} := \|f\|_{E^s}$ for the degree- s Gowers norm of f along E . The following lemma lists several properties of box norms that we established in the previous chapter.

Lemma 6.7 (Properties of box norms,). *Let $\delta \in (0, 1)$, let $d, D, N, s \in \mathbb{N}$, and let $E_1, \dots, E_s \subseteq \mathbb{Z}^D$ be finite and nonempty sets. Let $f : \mathbb{Z}^D \rightarrow \mathbb{C}$ be a 1-bounded function supported on $[N]^D$. Then the following properties hold:*

1. (Inductive formula) *For any $i \in \{1, \dots, s\}$, we have*

$$\|f\|_{E_1, \dots, E_s}^{2^s} = \mathbb{E}_{\mathbf{h}_1, \mathbf{h}'_1 \in E_1} \cdots \mathbb{E}_{\mathbf{h}_i, \mathbf{h}'_i \in E_i} \left\| \Delta'_{(\mathbf{h}_1, \mathbf{h}'_1), \dots, (\mathbf{h}_i, \mathbf{h}'_i)} f \right\|_{E_{i+1}, \dots, E_s}^{2^{s-i}}.$$

2. (Permutation invariance) *For any permutation σ on $\{1, \dots, s\}$, we have*

$$\|f\|_{E_1, \dots, E_s} = \|f\|_{E_{\sigma(1)}, \dots, E_{\sigma(s)}}.$$

3. (Monotonicity) *If f is 1-bounded and $|\mathbb{Z}^D - E_s| \leq CN^D$ for some $C > 0$, then*

$$\|f\|_{E_1, \dots, E_{s-1}}^{2^{s-1}} \geq \delta N^D \implies \|f\|_{E_1, \dots, E_s}^{2^s} \geq \delta^2 N^D / C. \quad (6.6)$$

4. (Enlarging the sets) *If $s \geq 2$ and $E'_i \subseteq \mathbb{Z}^D$ is a finite set satisfying $E'_i \supseteq E_i$ for each $i \in \{1, \dots, s\}$, then*

$$\|f\|_{E_1, \dots, E_s} \leq \left(\frac{|E'_1| \cdots |E'_s|}{|E_1| \cdots |E_s|} \right)^{1/2^{s-1}} \|f\|_{E'_1, \dots, E'_s}.$$

5. (Trimming the lengths of boxes) *Let*

$$E_i = \beta_i \cdot [\pm H_i] \quad \text{and} \quad E'_i = \beta_i \cdot [\pm \kappa_i H_i]$$

for some $\beta_i \in \mathbb{Z}^D$, $\kappa_i > 0$, and $H_i \in \mathbb{N}$. If f is 1-bounded, then

$$\|f\|_{E_1, \dots, E_s}^{2^s} \geq \delta N^D \implies \|f\|_{E'_1, \dots, E'_s}^{2^s} \gg_s \delta^{2^s} \min(1, \delta^{2^s} / \kappa)^{2^s} N^D$$

for $\kappa = \max_i \kappa_i$.

6. (Passing to sub-arithmetic progressions) Suppose that f is 1-bounded and

$$E_i = \beta_i \cdot [\pm H_i] \quad \text{and} \quad E'_i = q_i \beta_i \cdot [\pm H_i / q_i]$$

for some $\beta_i \in \mathbb{Z}^D$ and $q_i, H_i \in \mathbb{N}$ satisfying $q_i \leq C\delta^{2^s} H_i$ for all i and some constant $C > 0$. Then

$$\|f\|_{E_1, \dots, E_s}^{2^s} \geq \delta N^D \implies \|f\|_{E'_1, \dots, E'_s}^{2^s} \gg_{s, C} \delta N^D.$$

6.3.5 Dual-difference interchange

On several occasions, we apply a standard dual-difference interchange trick which has been a linchpin in all previous degree-lowering arguments. Its proof rests on the following observation.

Lemma 6.8. *Let $D, t \in \mathbb{N}$, let $C > 0$, let $E \subseteq \mathbb{Z}^D$ be finite and nonempty, and for each $i \in \{1, \dots, t\}$, let $F_i(\mathbf{x}) = \mathbb{E}_{z_i \in Z_i} f_i(\mathbf{x}, z_i)$ be a function from $\mathbb{Z}^D$ to $\mathbb{C}$ supported on a set of size at most CN^D , where Z_i is a nonempty finite set and $f_i : \mathbb{Z}^D \times Z_i \rightarrow \mathbb{C}$ is 1-bounded. If*

$$\left| \sum_{\mathbf{x}, \mathbf{h}} \mu_E(\mathbf{h}) \Delta_{(\mathbf{h}, 0)} \prod_{i=1}^t \mathbb{E}_{z_i \in Z_i} f_i(\mathbf{x}, z_i) \right| \geq \delta N^D,$$

then

$$\sum_{\mathbf{x}, \mathbf{h}} \mu_E(\mathbf{h}) \prod_{i=1}^t \mathbb{E}_{z_i \in Z_i} \Delta_{(\mathbf{h}, 0)} f_i(\mathbf{x}, z_i) \geq (\delta^2 / C) N^D.$$

Proof. We first expand $\mu_E(\mathbf{h})$ so that

$$\left| \sum_{\mathbf{x}} \mathbb{E}_{\mathbf{h} \in E} \prod_{i=1}^t \mathbb{E}_{z_i \in Z_i} f_i(\mathbf{x} + \mathbf{h}, z_i) \cdot \mathbb{E}_{\mathbf{h}' \in E} \prod_{i=1}^t \mathbb{E}_{z'_i \in Z_i} \overline{f_i}(\mathbf{x} + \mathbf{h}', z'_i) \right| \geq \delta N^D.$$

Using the identity

$$\prod_{i=1}^t \mathbb{E}_{z_i \in Z_i} g(z_i) = \mathbb{E}_{z_1 \in Z_1, \dots, z_t \in Z_t} \prod_{i=1}^t g(z_i) \tag{6.7}$$

and the formula for F_i , we can further recast this as

$$\left| \sum_{\mathbf{x}} \mathbb{E}_{z_1 \in Z_1, \dots, z_t \in Z_t} \left(\mathbb{E}_{\mathbf{h}' \in E} \prod_{i=1}^t \overline{F_i}(\mathbf{x} + \mathbf{h}') \right) \mathbb{E}_{\mathbf{h} \in E} \prod_{i=1}^t f_i(\mathbf{x} + \mathbf{h}, z_i) \right| \geq \delta N^D.$$

We then apply the Cauchy-Schwarz inequality to double $\mathbf{h}$ and use the 1-boundedness of f_i 's together with

the bound on the support of F_i to get

$$\begin{aligned} (\delta^2/C)N^D &\leq \sum_{\mathbf{x}} \mathbb{E}_{z_1 \in Z_1, \dots, z_t \in Z_t} \left| \mathbb{E}_{\mathbf{h} \in E} \prod_{i=1}^t f_i(\mathbf{x} + \mathbf{h}, z_i) \right|^2 \\ &= \sum_{\mathbf{x}, \mathbf{h}} \mu_E(\mathbf{h}) \mathbb{E}_{z_1 \in Z_1, \dots, z_t \in Z_t} \prod_{i=1}^t \Delta_{(\mathbf{h}, 0)} f_i(\mathbf{x}, z_i). \end{aligned}$$

The result follows from applying (6.7) in the reverse direction. □

Iterating Lemma 6.8, we obtain the following.

Lemma 6.9 (Dual-difference interchange). *Let $C > 0$, let $D, r, s \in \mathbb{N}$ with $1 \leq r \leq s$, let $E_1, \dots, E_s \subseteq [\pm CN^D]$ be nonempty sets, and let $F(\mathbf{x}) = \mathbb{E}_{z \in Z} f(\mathbf{x}, z)$ be a function from $\mathbb{Z}^D$ to $\mathbb{C}$, where Z is a nonempty finite set and $f : \mathbb{Z}^D \times Z \rightarrow \mathbb{C}$ is a 1-bounded function supported on $[\pm CN^D] \times Z$. If*

$$\|F\|_{E_1, \dots, E_s}^{2^s} \geq \delta N^D,$$

then

$$\sum_{\mathbf{h}_1, \dots, \mathbf{h}_r} \mu_{E_1}(\mathbf{h}_1) \cdots \mu_{E_r}(\mathbf{h}_r) \|F^{\mathbf{h}_1, \dots, \mathbf{h}_r}\|_{E_{r+1}, \dots, E_s}^{2^{s-r}} \gg_{C,r} \delta^{2^r} N^D,$$

where

$$F^{\mathbf{h}_1, \dots, \mathbf{h}_r}(\mathbf{x}) := \mathbb{E}_{z \in Z} \Delta_{(\mathbf{h}_1, 0), \dots, (\mathbf{h}_r, 0)} f(\mathbf{x}, z).$$

Proof. Expanding the definition of the box norm, we get that

$$\sum_{\mathbf{h}_2, \dots, \mathbf{h}_s} \mu_{E_2}(\mathbf{h}_2) \cdots \mu_{E_s}(\mathbf{h}_s) \cdot \sum_{\mathbf{x}, \mathbf{h}_1} \mu_{E_1}(\mathbf{h}_1) \Delta_{(\mathbf{h}_1, 0), \dots, (\mathbf{h}_s, 0)} \mathbb{E}_{z \in Z} f(\mathbf{x}, z) \geq \delta N^D.$$

We apply Lemma 6.8 for a fixed tuple $(\mathbf{h}_2, \dots, \mathbf{h}_s) \in E_2 \times \cdots \times E_s$, viewing the iterated multiplicative derivative $\Delta_{(\mathbf{h}_2, 0), \dots, (\mathbf{h}_s, 0)}$ as a product of 2^{s-1} terms (hence, the multiplicative derivative plays the role of the product over $i \in \{1, \dots, t\}$ in the statement of Lemma 6.8). This lets us move the differencing over $\mathbf{h}_1$ inside the expectation, just like in the statement of Lemma 6.8), giving

$$\sum_{\mathbf{h}_1, \dots, \mathbf{h}_s} \mu_{E_1}(\mathbf{h}_1) \cdots \mu_{E_s}(\mathbf{h}_s) \cdot \sum_{\mathbf{x}} \Delta_{(\mathbf{h}_2, 0), \dots, (\mathbf{h}_s, 0)} \mathbb{E}_{z \in Z} \Delta_{(\mathbf{h}_1, 0)} f(\mathbf{x}, z) \gg_C \delta^2 N^D.$$

We continue in the same fashion to move the differencings over $\mathbf{h}_2, \dots, \mathbf{h}_r$ inside the expectation; here we are implicitly using the fact that the functions $\Delta_{(\mathbf{h}_1, 0), \dots, (\mathbf{h}_j, 0)} f$ for $\mathbf{h}_i \in E_i$, $1 \leq i \leq j \leq r$ are supported on $[\pm CN^D] \times Z$. $\square$

6.3.6 Fourier analysis

We conclude the preliminaries section with a brief summary of Fourier analysis on $\mathbb{Z}$, which we shall use on occasion. For a finitely supported function $f : \mathbb{Z} \rightarrow \mathbb{C}$ and $\xi \in \mathbb{R}/\mathbb{Z}$, we define the *Fourier transform* $\hat{f} : \mathbb{R}/\mathbb{Z} \rightarrow \mathbb{C}$ via

$$\hat{f}(\xi) := \sum_x f(x) e(\xi x).$$

The Fourier inversion formula gives $f(\xi) = \int_{\mathbb{R}/\mathbb{Z}} \hat{f}(\xi) e(-\xi x) d\xi$, and the Parseval Identity gives $\|f\|_2 = \|\hat{f}\|_2$. Here the L^p -norm of a function $f : \mathbb{Z} \rightarrow \mathbb{C}$ is defined by $\|f\|_p := \left(\sum_x |f(x)|^p \right)^{1/p}$ for $1 \leq p < \infty$ and $\|f\|_\infty := \sup_{x \in \mathbb{Z}} |f(x)|$, the L^p -norm of a (Lebesgue-measurable) function $F : \mathbb{R}/\mathbb{Z} \rightarrow \mathbb{C}$ is defined by $\|F\|_p := \left(\int_{\mathbb{R}/\mathbb{Z}} |F(\xi)|^p d\xi \right)^{1/p}$ for $1 \leq p < \infty$ and $\|F\|_\infty := \sup_{\xi \in \mathbb{R}/\mathbb{Z}} |F(\xi)|$.

6.4 Degree-lowering

For the remainder of the chapter, we fix the integer polynomial P together with its degree d and the auxiliary W -tricked polynomial $\tilde{P}$ defined in (6.1). Recall that for 1-bounded functions $f_0, f_1, f_2 : \mathbb{Z}^2 \rightarrow \mathbb{C}$ supported on $[N]^2$, we have the counting operator

$$\Lambda^W(f_0, f_1, f_2) = \sum_{x, y \in \mathbb{Z}} \mathbb{E}_{z \in [(N/W_d)^{1/d}]} f_0(x, y) f_1(x + \tilde{P}(z), y) f_2(x, y + \tilde{P}(z)),$$

where W_d is the leading coefficient of $\tilde{P}$. Recall also that $\mathbf{e}_1 := (1, 0)$ and $\mathbf{e}_2 := (0, 1)$ denote the usual basis vectors in $\mathbb{Z}^2$. As described in the proof outline above, the main work of this chapter is showing that Λ^W is “controlled” by appropriate degree-2 box norms of its arguments.

Theorem 6.10 (Inverse theorem for Λ^W). *There is a constant $C = C_P > 0$ such that the following holds. Let $\delta \in (0, 1/10)$, and let $f_0, f_1, f_2 : \mathbb{Z}^2 \rightarrow \mathbb{C}$ be 1-bounded functions supported on $[N]^2$ for some integer $N \geq CW^C \exp(\log(1/\delta)^C)$. If*

$$|\Lambda^W(f_0, f_1, f_2)| \geq \delta N^2,$$

then there exists some positive integer power $V = W^{O_P(1)}$ such that

$$\begin{aligned} \|f_0\|_{\mathbf{e}_1 V \cdot [N/V], \mathbf{e}_2 V \cdot [N/V]}^4 &\geq \exp(-\log(1/\delta)^{O_P(1)}) N^2, \\ \|f_1\|_{\mathbf{e}_1 V \cdot [N/V], (\mathbf{e}_2 - \mathbf{e}_1) V \cdot [N/V]}^4 &\geq \exp(-\log(1/\delta)^{O_P(1)}) N^2, \\ \text{and } \|f_2\|_{\mathbf{e}_2 V \cdot [N/V], (\mathbf{e}_2 - \mathbf{e}_1) V \cdot [N/V]}^4 &\geq \exp(-\log(1/\delta)^{O_P(1)}) N^2. \end{aligned}$$

Theorem 6.10 is a W -tricked version of Theorem 6.2 from the introduction. While we prove only the former, the latter follows from a simplified version of the proof of Theorem 6.10 in which we take $W = V = 1$ in all places and replace $\tilde{P}$ with P . We remark that while Theorem 6.10 has the same restriction on P as Theorem 6.1, no such restriction is present in the statement of Theorem 6.2. This is because the fraction-comparison argument Lemma 6.47, whose proof in the W -tricked setting uses the special structure of P , holds trivially for all non- W -tricked polynomials.

The purpose of this section is to reduce Theorem 6.10 to Proposition 6.13 (below); the corresponding part of the argument in the proof sketch of Section 6.2 is the paragraph after the definition of Λ' .

To prove Theorem 6.10, we will use a degree-lowering argument. The starting point for the argument is the following special case of Corollary 5.7 from the previous chapter.

Proposition 6.11. *There exist a constant $C = C_d > 0$ and a positive integer $t = O_d(1)$ such that for all $\delta \in (0, 1/10)$ and 1-bounded functions $f_0, f_1, f_2 : \mathbb{Z}^2 \rightarrow \mathbb{C}$ supported on $[N]^2$ for some integer $N \geq CW_d \delta^{-C}$, the bound*

$$|\Lambda^W(f_0, f_1, f_2)| \geq \delta N^2$$

implies that

$$\begin{aligned} \|f_0\|_{(\mathbf{e}_1 W_d \cdot [\pm N/W_d])^t, (\mathbf{e}_2 W_d \cdot [\pm N/W_d])^t}^{2^{2t}} &\gg_d \delta^{O_d(1)} N^2 \\ \|f_1\|_{(\mathbf{e}_1 W_d \cdot [\pm N/W_d])^t, ((\mathbf{e}_2 - \mathbf{e}_1) W_d \cdot [\pm N/W_d])^t}^{2^{2t}} &\gg_d \delta^{O_d(1)} N^2 \\ \text{and } \|f_2\|_{(\mathbf{e}_2 W_d \cdot [\pm N/W_d])^t, ((\mathbf{e}_2 - \mathbf{e}_1) W_d \cdot [\pm N/W_d])^t}^{2^{2t}} &\gg_d \delta^{O_d(1)} N^2. \end{aligned}$$

Our goal in the next few sections is to reduce t to 2 in the box norms appearing in Proposition 6.11. The following proposition captures the main iterative step of this process.

Proposition 6.12 (Degree-lowering). *Let $k, \ell \in \mathbb{N}$ with $\max(k, \ell) \geq 2$. For every $C_1 > 0$ there exists $C_2 = C_2(C_1, P, k, \ell) > 0$ such that the following holds. Let $V \leq W^{C_1}$ be a nonnegative integer power of W , and suppose that for all $\delta \in (0, 1/10)$, all positive integers $N \geq C_2 W^{C_2} \exp(\log(1/\delta)^{C_2})$, and all 1-bounded*

functions $f_0, f_1, f_2 : \mathbb{Z}^2 \rightarrow \mathbb{C}$ supported on $[N]^2$, the lower bound

$$|\Lambda^W(f_0, f_1, f_2)| \geq \delta N^2$$

implies that

$$\|f_0\|_{(\mathbf{e}_1 V \cdot [\pm N/V])^k, (\mathbf{e}_2 V \cdot [\pm N/V])^\ell}^{2^{k+\ell}} \geq \exp(-\log(1/\delta)^{C_1}) N^2.$$

Then there exists some positive integer power $\tilde{V} = V^{O_{C_1, P, k, \ell}(1)}$ such that

$$\|f_0\|_{(\mathbf{e}_1 \tilde{V} \cdot [N/\tilde{V}])^{k-1}, (\mathbf{e}_2 \tilde{V} \cdot [\pm N/\tilde{V}])^\ell}^{2^{k+\ell-1}} \geq \exp(-\log(1/\delta)^{O_{C_1, P, k, \ell}(1)}) N^2$$

if $k \geq 2$ and

$$\|f_0\|_{(\mathbf{e}_1 \tilde{V} \cdot [N/\tilde{V}])^k, (\mathbf{e}_2 \tilde{V} \cdot [\pm N/\tilde{V}])^{\ell-1}}^{2^{k+\ell-1}} \geq \exp(-\log(1/\delta)^{O_{C_1, P, k, \ell}(1)}) N^2$$

if $\ell \geq 2$.

Before proving Proposition 6.12, we show how it implies Theorem 6.10.

Proof of Theorem 6.10 assuming Propositions 6.11 and 6.12. By Proposition 6.11, there exists a positive integer $t = O_d(1)$ such that

$$\|f_0\|_{(\mathbf{e}_1 W_d \cdot [\pm N/W_d])^t, (\mathbf{e}_2 W_d \cdot [\pm N/W_d])^t}^{2^{2t}} \gg_d \delta^{O_d(1)} N^2.$$

Using Lemma 6.7(4), we can replace W_d by W^{d-1} at the cost of a factor of $O_P(1)$ in the lower bound. We now apply Proposition 6.12 $2t - 2$ times. Since N is chosen sufficiently large, we end up with a positive integer power $V = W^{O_d(1)}$ such that

$$\|f_0\|_{\mathbf{e}_1 V \cdot [\pm N/V], \mathbf{e}_2 V \cdot [\pm N/V]}^4 \geq \exp(-\log(1/\delta)^{O_P(1)}) N^2.$$

This gives the claimed control in terms of f_0 .

The control in terms of f_1, f_2 follows by symmetry. To see this, notice that the affine linear map $(x, y) \mapsto (2N - 1 - (x + y), y)$ sends $[N]^2$ into $[2N]^2$ and preserves the set of corners with each given side length, with the caveat that the two points on the horizontal side of each corner have been “exchanged”. More precisely, for each i , define the function $\tilde{f}_i : [2N]^2 \rightarrow \mathbb{C}$ by $\tilde{f}_i(x, y) := f_i(2N - 1 - (x + y), y)$. Then

the change of variables $x \mapsto 2N - 1 - x - y - \tilde{P}(z)$ gives the identity

$$\Lambda^W(f_0, f_1, f_2) = \Lambda^W(\tilde{f}_1, \tilde{f}_0, \tilde{f}_2).$$

The previous paragraph (with N, δ replaced by $2N, \delta/4$) tells us that

$$\|\tilde{f}_1\|_{\mathbf{e}_1 V \cdot [\pm 2N/V], \mathbf{e}_2 V \cdot [\pm 2N/V]}^4 \geq \exp(-\log(1/\delta)^{O_P(1)})N^2.$$

Using Lemma 6.7(5) to replace $[\pm 2N/V]$ by $[\pm N/V]$ and unwinding the definition of $\tilde{f}_1$, we conclude that

$$\|f_1\|_{\mathbf{e}_1 V \cdot [N/V], (\mathbf{e}_2 - \mathbf{e}_1) V \cdot [N/V]}^4 \geq \exp(-\log(1/\delta)^{O_P(1)})N^2.$$

The control for f_2 follows from an analogous change of variables. □

6.4.1 Proof of Proposition 6.12: preliminary maneuvers

We now begin the proof of Proposition 6.12. All implicit constants are allowed to depend on P, k, l . The initial assumption

$$|\Lambda^W(f_0, f_1, f_2)| \geq \delta N^2$$

implies that

$$\|f_0\|_{(\mathbf{e}_1 V \cdot [\pm N/V])^k, (\mathbf{e}_2 V \cdot [\pm N/V])^\ell}^{2^{k+\ell}} \geq \exp(-\log(1/\delta)^{O(1)})N^2 \tag{6.8}$$

for all 1-bounded functions $f_0, f_1, f_2 : \mathbb{Z}^2 \rightarrow \mathbb{C}$ supported on $[N]^2$. Our proof strategy is as follows. Applying several standard maneuvers such as stashing, the dual-difference interchange and an inverse theorem for the $U^2 \times U^1$ box norm, we reduce the problem to the study of the auxiliary operator

$$\Lambda'(a, \tilde{f}_1, \tilde{f}_2) := \sum_{x, y} \mathbb{E}_{z \in [(N/W_a)^{1/d} V^{-1}]} e(a(y)x) \tilde{f}_1(x + \tilde{P}_{[r, V]}(z), y) \tilde{f}_2(x, y + \tilde{P}_{[r, V]}(z)).$$

The phase $e(a(y)x)$ is the “main term” that arises from the $U^2 \times U^1$ inverse theorem. Our main technical advancement is a structural result on Λ' : We show that if $\Lambda'(a, \tilde{f}_1, \tilde{f}_2)$ is large, then a is “approximately constant” in the sense that there are many arithmetic progressions of difference $V' = V^{O(1)}$ on each of which a is constant a positive proportion of the time. Morally speaking, this means that $\Lambda'(a, \tilde{f}_1, \tilde{f}_2)$ can be large

only when $e(a(y)x)$ is actually a $U^1 \times U^1$ obstruction. We use this information to obtain the auxiliary norm control

$$\|f_1\|_{(\mathbf{e}_1 V V' \cdot [\pm N/\tilde{V}])^{k-1}, (\mathbf{e}_2 - \mathbf{e}_1) \tilde{V} \cdot [\pm N/\tilde{V}], (\mathbf{e}_2 \tilde{V} \cdot [\pm N/\tilde{V}])^{\ell-1}}^4 \geq \exp(-\log(1/\delta)^{O(1)}) N^2,$$

where $\tilde{V} = VV'$. We then transfer this norm-control from f_1 to f_0 .

Let us fill in some more details of this strategy. In doing so, we will also highlight several techniques used in various guises throughout the chapter.

We define the *dual function*

$$\mathcal{D}_0(f_1, f_2)(x, y) := \mathbb{E}_{z \in [(N/W_d)^{1/d}]} f_1(x + \tilde{P}(z), y) f_2(x, y + \tilde{P}(z))$$

with respect to the first argument of Λ^W . The Cauchy–Schwarz inequality gives

$$\begin{aligned} |\Lambda^W(f_0, f_1, f_2)|^2 &= \left| \sum_{x, y} f_0(x, y) \mathcal{D}_0(f_1, f_2)(x, y) \right|^2 \\ &\leq N^2 \cdot \sum_{x, y} |\mathcal{D}_0(f_1, f_2)(x, y)|^2 \\ &= N^2 \cdot \Lambda^W(\mathcal{D}_0(f_1, f_2), \overline{f_1}, \overline{f_2}), \end{aligned}$$

for all 1-bounded functions f_0 supported on $[N]^2$. It follows that if $\Lambda^W(f_0, f_1, f_2) \geq \delta N^2$, then we also have

$$\Lambda^W(\mathcal{D}_0(f_1, f_2), \overline{f_1}, \overline{f_2}) \geq \delta^2 N^2.$$

This trick whereby we replace an arbitrary function f_0 by the structured function $\mathcal{D}_0(f_1, f_2)$ is often called *stashing*, and it will be the bread and butter of our arguments. In particular, once we know that Λ^W is controlled by some norm $\|\cdot\|$ of its first argument, we see that the largeness of $\Lambda^W(f_0, f_1, f_2)$ implies not only the largeness of $\|f_0\|$ but also the largeness of $\|\mathcal{D}_0(f_1, f_2)\|$. In this particular case, stashing and (6.8) together give

$$\|\mathcal{D}_0(f_1, f_2)\|_{(\mathbf{e}_1 V \cdot [\pm N/V])^k, (\mathbf{e}_2 V \cdot [\pm N/V])^\ell}^{2^{k+\ell}} \geq \exp(-\log(1/\delta)^{O(1)}) N^2. \quad (6.9)$$

The next step is to apply the dual–difference interchange to (6.9) in order to arrive at an average of

$U^2 \times U^1$ -box norms to which we can apply an inverse theorem. To begin, applying Lemma 6.9 to (6.9) gives

$$\sum_{\substack{\underline{h}_1 \in \mathbb{Z}^{k-2}, \\ \underline{h}_2 \in \mathbb{Z}^{\ell-1}}} \mu_{N/V}(\underline{h}_1, \underline{h}_2) \|\mathcal{D}_0^{\underline{h}_1, \underline{h}_2}(f_1, f_2)\|_{(\mathbf{e}_1 V \cdot [\pm N/V])^2, \mathbf{e}_2 V \cdot [\pm N/V]}^8 \geq \exp(-\log(1/\delta)^{O(1)}) N^2,$$

where

$$\mathcal{D}_0^{\underline{h}_1, \underline{h}_2}(f_1, f_2)(x, y) := \mathbb{E}_{z \in [(N/W_d)^{1/d}]} \Delta_{\underline{h}_1 V \mathbf{e}_1, \underline{h}_2 V \mathbf{e}_2}(f_1(x + \tilde{P}(z), y) f_2(x, y + \tilde{P}(z)))$$

for $\underline{h}_1 \in \mathbb{Z}^{k-2}$ and $\underline{h}_2 \in \mathbb{Z}^{\ell-1}$. (We recall here that $\underline{h}_1 V \mathbf{e}_1 = (h_{11} V \mathbf{e}_1, \dots, h_{1(k-2)} V \mathbf{e}_1)$, and similarly for $\underline{h}_2 V \mathbf{e}_2$.) Since $\mu_{N/V}(\underline{h}_1, \underline{h}_2) \ll (N/V)^{-k-\ell+3}$ pointwise, the popularity principle provides at least $\exp(-\log(1/\delta)^{O(1)})(N/V)^{k+\ell-3}$ choices of $(\underline{h}_1, \underline{h}_2)$ such that

$$\|\mathcal{D}_0^{\underline{h}_1, \underline{h}_2}(f_1, f_2)\|_{(\mathbf{e}_1 V \cdot [\pm N/V])^2, \mathbf{e}_2 V \cdot [\pm N/V]}^8 \geq \exp(-\log(1/\delta)^{O(1)}) N^2. \quad (6.10)$$

For the next several steps, fix such a choice of $\underline{h}_1, \underline{h}_2$. (The reader can now see why all of the main ideas appear in the case $k = 2, \ell = 1$ described in the proof sketch.)

The box norm in (6.10) is not yet ripe for an application of an inverse theorem for the $U^2 \times U^1$ box norms because of the multiplicative factor V in the boxes $\mathbf{e}_1 V \cdot [\pm N/V]$ and $\mathbf{e}_2 V \cdot [\pm N/V]$. To bring the box norm into a form amenable to an application of the inverse theorem, we split the ranges of x, y into arithmetic progressions of common difference V and observe the identity

$$\|g\|_{(\mathbf{e}_1 V \cdot [\pm N/V])^2, \mathbf{e}_2 V \cdot [\pm N/V]}^8 = \sum_{j_1, j_2 \in [V]} \|g_{(j_1, j_2, V)}\|_{(\mathbf{e}_1 \cdot [\pm N/V])^2, \mathbf{e}_2 \cdot [\pm N/V]}^8. \quad (6.11)$$

for any compactly supported $g : \mathbb{Z} \rightarrow \mathbb{C}$; we recall the V -tricked function $g_{(j_1, j_2, V)}(x, y) = g(Vx + j_1, Vy + j_2)$ from Definition 6.4. Notice that after this change of variables, the boxes in the box norm have side length $[N/V]$ instead of $V \cdot [N/V]$.

We will of course apply the identity (6.11) to the functions $\mathcal{D}_0^{\underline{h}_1, \underline{h}_2}(f_1, f_2)$. This maneuver requires us to adapt the range of z accordingly: Recalling the definition of the V -trick for polynomials, we replace $\tilde{P}(z)$ by $\tilde{P}_{[r, V]}(z) = \frac{\tilde{P}(Vz+r) - \tilde{P}(r)}{V}$. On changing variables $z \mapsto Vz + r$, we can express the dual function as

$$\begin{aligned} & \mathcal{D}_0^{\underline{h}_1, \underline{h}_2}(f_1, f_2)(Vx + j_1, Vy + j_2) \\ &= \mathbb{E}_{z \in [(N/W_d)^{1/d}]} \Delta_{V \underline{h}_1 \mathbf{e}_1, V \underline{h}_2 \mathbf{e}_2}(f_1(Vx + j_1 + \tilde{P}(z), Vy + j_2) f_2(Vx + j_1, Vy + j_2 + \tilde{P}(z))) \\ &= \mathbb{E}_{r \in [V]} \mathbb{E}_{z \in [(N/W_d)^{1/d} V^{-1}]} \Delta_{V \underline{h}_1 \mathbf{e}_1, V \underline{h}_2 \mathbf{e}_2}(f_1(V(x + \tilde{P}_{[r, V]}(z)) + j_1 + \tilde{P}(r), Vy + j_2) \end{aligned}$$

$$\begin{aligned}
& \times f_2(Vx + j_1, V(y + \tilde{P}_{[r,V]}(z)) + j_2 + \tilde{P}(r)) \\
& = \mathbb{E}_{r \in [V]} \mathbb{E}_{z \in [(N/W_d)^{1/d}V^{-1}]} \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} (f_{1,j_1+\tilde{P}(r),j_2}(x + \tilde{P}_{[r,V]}(z), y) f_{2,j_1,j_2+\tilde{P}(r)}(x, y + \tilde{P}_{[r,V]}(z))),
\end{aligned}$$

where we also set

$$f_{l,j_1,j_2}(x, y) := (f_l)_{(j_1,j_2,V)}(x, y) = f_l(Vx + j_1, Vy + j_2)$$

for brevity. We observe that these V -tricked functions are supported on an interval of length N/V . In the formula above (and in similar expressions in the future), we assume without loss of generality that $(N/W_d)^{1/d}V^{-1} \in \mathbb{N}$ (as we may since N is large compared to V), so that splitting the range of z into arithmetic progressions incurs no error.

Invoking (6.11) with $g = \mathcal{D}_0^{\underline{h}_1, \underline{h}_2}(f_1, f_2)$ and applying the triangle inequality to (6.10) gives

$$\begin{aligned}
& \mathbb{E}_{j_1, j_2, r \in [V]} \left\| \mathbb{E}_{z \in [(N/W_d)^{1/d}V^{-1}]} \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{1,j_1+\tilde{P}(r),j_2}(\cdot + \tilde{P}_{[r,V]}(z), \cdot) \right. \\
& \quad \left. \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{2,j_1,j_2+\tilde{P}(r)}(\cdot, \cdot + \tilde{P}_{[r,V]}(z)) \right\|_{(\mathbf{e}_1 \cdot [\pm N/V])^2, \mathbf{e}_2 \cdot [\pm N/V]}^8 \\
& \geq \exp(-\log(1/\delta)^{O(1)})(N/V)^2.
\end{aligned}$$

Since each summand contributes at most $(N/V)^2$, the popularity principle provides at least $\exp(-\log(1/\delta)^{O(1)})V^3$ choices of $(j_1, j_2, r) \in [V]^3$ such that

$$\begin{aligned}
& \left\| \mathbb{E}_{z \in [(N/W_d)^{1/d}V^{-1}]} \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{1,j_1+\tilde{P}(r),j_2}(\cdot + \tilde{P}_{[r,V]}(z), \cdot) \right. \\
& \quad \left. \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{2,j_1,j_2+\tilde{P}(r)}(\cdot, \cdot + \tilde{P}_{[r,V]}(z)) \right\|_{(\mathbf{e}_1 \cdot [\pm N/V])^2, \mathbf{e}_2 \cdot [\pm N/V]}^8 \\
& \geq \exp(-\log(1/\delta)^{O(1)})(N/V)^2.
\end{aligned}$$

For the next several steps, fix such a choice of j_1, j_2, r . For each “good” choice of j_1, j_2, r from the previous step, the $U^2 \times U^1$ -inverse theorem (Lemma 6.28) tells us that there exist a 1-bounded function $g : \mathbb{Z} \rightarrow \mathbb{C}$ supported on $[N]$ and phase functions $a, b : \mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}$ (all depending on $\underline{h}_1, \underline{h}_2, j_1, j_2, r$) such that

$$\begin{aligned}
& \sum_{x, y} \mathbb{E}_{z \in [(N/W_d)^{1/d}V^{-1}]} \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{1,j_1+\tilde{P}(r),j_2}(x + \tilde{P}_{[r,V]}(z), y) \\
& \quad \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{2,j_1,j_2+\tilde{P}(r)}(x, y + \tilde{P}_{[r,V]}(z)) g(x) e(a(y)x + b(y)) \geq \exp(-\log(1/\delta)^{O(1)})(N/V)^2.
\end{aligned}$$

Hiding the dependence on $\underline{h}_1, \underline{h}_2, j_1, j_2, r$, we define

$$\begin{aligned}\tilde{f}_1(x, y) &:= e(b(y))\Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{1, j_1 + \tilde{P}(r), j_2}(x, y) \\ \text{and } \tilde{f}_2(x, y) &:= g(x)\Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{2, j_1, j_2 + \tilde{P}(r)}(x, y);\end{aligned}$$

then the previous inequality can be rewritten as

$$\begin{aligned}\sum_{x, y} \mathbb{E}_{z \in [(N/W_d)^{1/d} V^{-1}]} e(a(y)x) \tilde{f}_1(x + \tilde{P}_{[r, V]}(z), y) \tilde{f}_2(x, y + \tilde{P}_{[r, V]}(z)) \\ \geq \exp(-\log(1/\delta)^{O(1)})(N/V)^2.\end{aligned}$$

6.4.2 Examining the structure of a

The main task is showing that the phase function a possesses some extra structure; specifically, we will show that $a(y)$ is approximately constant much of the time. For future reference, define the counting operator

$$\Lambda'(a, \tilde{f}_1, \tilde{f}_2) := \sum_{x, y} \mathbb{E}_{z \in [(N/W_d)^{1/d} V^{-1}]} e(a(y)x) \tilde{f}_1(x + \tilde{P}_{[r, V]}(z), y) \tilde{f}_2(x, y + \tilde{P}_{[r, V]}(z)).$$

The key structural result for the function a is the following proposition.

Proposition 6.13. *For every $C_1 > 0$ there exists $C_2 = C_2(C_1, P) > 0$ such that the following holds. Let $\delta \in (0, 1/10)$, let $V \leq W^{C_1}$ be a nonnegative integer power of W , and let $r \in [V]$. Let $\tilde{f}_1, \tilde{f}_2 : \mathbb{Z}^2 \rightarrow \mathbb{C}$ be 1-bounded functions supported on $[N/V]^2$, for some integer $N \geq C_2 V^{C_2} \exp(\log(1/\delta)^{C_2})$, and let $a : \mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}$ be a phase function. If*

$$|\Lambda'(a, \tilde{f}_1, \tilde{f}_2)| \geq \delta(N/V)^2,$$

then there exist a positive integer power $V' = V^{O_{C_1, P}(1)}$ and phases $\alpha_0, \dots, \alpha_{V'-1} \in \mathbb{R}/\mathbb{Z}$ such that for at least $\exp(-\log(1/\delta)^{O_{C_1, P}(1)})V'$ values of $j \in [V']$, the bound

$$\|V' a_{(j, V')}(y) - \alpha_j\|_{\mathbb{R}/\mathbb{Z}} \leq \exp(\log(1/\delta)^{O_{C_1, P}(1)})VV'/N$$

holds for at least $\exp(-\log(1/\delta)^{O_{C_1, P}(1)})N/(VV')$ choices of $y \in [N/(VV')]$.

Unfortunately, this set of y 's where the phase can be well-approximated by constants may contribute little to the value of $\Lambda'(a, \tilde{f}_1, \tilde{f}_2)$. To obtain a consequence better suited to our purposes, we bootstrap Proposition 6.13 and produce some such set of y 's that does make a significant contribution to $\Lambda'(a, \tilde{f}_1, \tilde{f}_2)$.

The exact statement is as follows.

Proposition 6.14. *For every $C_1 > 0$ there exists $C_2 = C_2(C_1, P) > 0$ such that the following holds. Let $\delta \in (0, 1/10)$, let $V \leq W^{C_1}$ be a nonnegative integer power of W , and let $r \in [V]$. Let $\tilde{f}_1, \tilde{f}_2 : \mathbb{Z}^2 \rightarrow \mathbb{C}$ be 1-bounded functions supported on $[N/V]^2$, for some integer $N \geq C_2 V^{C_2} \exp(\log(1/\delta)^{C_2})$, and let $a : \mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}$ be any function. If*

$$|\Lambda'(a, \tilde{f}_1, \tilde{f}_2)| \geq \delta(N/V)^2,$$

then there exist a positive integer power $V' = V^{O_{C_1, P}(1)}$, an integer

$$\exp(-\log(1/\delta)^{O_{C_1, P}(1)})N \leq N' \leq N,$$

and phases $\alpha_0, \dots, \alpha_{V'-1} \in \mathbb{R}/\mathbb{Z}$ such that

$$\begin{aligned} \sum_{j \in [V']} \sum_{i \in [N/N']} \sum_y \left| \sum_{x \in [N'/(VV')]} \sum_{z \in [(N/W_d)^{1/d} V^{-1}]} \mathbb{E} e(\alpha_y \pmod{V'} x) \right. \\ \left. \tilde{f}_1(V'x + j + \frac{N'}{V}i + \tilde{P}_{[r, V]}(z), y) \tilde{f}_2(V'x + j + \frac{N'}{V}i, y + \tilde{P}_{[r, V]}(z)) \right| \\ \geq \exp(-\log(1/\delta)^{O_{C_1, P}(1)})(N/V)^2. \end{aligned}$$

Proposition 6.14 essentially says that we can replace $a(y)$ by a phase function $y \mapsto \alpha_y \pmod{V'}$ that is constant on arithmetic progressions of difference V' . The summation over $i \in [N/N']$ is a technical annoyance, necessitated by the proof, that will be easily removed in our later application of Proposition 6.14.

Proof. Before applying Proposition 6.13, we replace $a(y)$ by a modified phase function $\tilde{a}(y)$ with better properties. An important observation is that, by Plancherel's Identity, if $g : \mathbb{Z} \rightarrow \mathbb{C}$ is any 1-bounded function supported on $[M]$, then

$$\int_{\mathbb{R}/\mathbb{Z}} |\hat{g}(\xi)|^2 da = \sum_x |g(x)|^2 \leq M;$$

it follows that $|\hat{g}(\xi)| \leq 2M^{1/2}$ for all ξ outside of a set of measure at most $1/4$ (say). Write

$$\Lambda'(a, \tilde{f}_1, \tilde{f}_2) = \sum_y F(y),$$

where

$$F(y) := \sum_x \sum_{z \in [(N/W_d)^{1/d} V^{-1}]} \mathbb{E} e(a(y)x) \tilde{f}_1(x + \tilde{P}_{[r, V]}(z), y) \tilde{f}_2(x, y + \tilde{P}_{[r, V]}(z)).$$

Say that $y \in [N/V]$ is “good” if $|F(y)| \geq (\delta/2)(N/V)$ and that it is “bad” otherwise. The sum of $F(y)$ over all bad y ’s has absolute value at most $(\delta/2)(N/V)^2$, so

$$\left| \sum_{\text{good } y} F(y) \right| \geq \left| \Lambda'(a, \tilde{f}_1, \tilde{f}_2) \right| - (\delta/2)(N/V)^2 \geq (\delta/2)(N/V)^2. \quad (6.12)$$

We are now ready to define the modified phase function $\tilde{a}(y)$. On good y ’s, we define $\tilde{a}(y)$ to be equal to $a(y)$. Let $y_1, \dots, y_R$ be the bad elements of $[N/V]$. We iteratively choose the values of $\tilde{a}(y_1), \tilde{a}(y_2), \dots, \tilde{a}(y_R)$ such that each $\tilde{a}(y_\ell)$ is at least $(10N/V)^{-1}$ away from $\tilde{a}(y_1), \dots, \tilde{a}(y_{\ell-1})$ and satisfies

$$\left| \sum_x \sum_{z \in [(N/W_d)^{1/d} V^{-1}]} \mathbb{E} e(\tilde{a}(y_\ell)x) \tilde{f}_1(x + \tilde{P}_{[r,V]}(z), y_\ell) \tilde{f}_2(x, y_\ell + \tilde{P}_{[r,V]}(z)) \right| \leq 2(N/V)^{1/2};$$

that this is possible follows from the Plancherel observation from the beginning of this paragraph, since the already-chosen positions of $\tilde{a}(y_1), \dots, \tilde{a}(y_{\ell-1})$ “rule out” at most $1/5 < 3/4$ of $\mathbb{R}/\mathbb{Z}$ for the position of $\tilde{a}(y_\ell)$.

Notice that

$$\sum_{\text{bad } y} \left| \sum_x \sum_{z \in [(N/W_d)^{1/d} V^{-1}]} \mathbb{E} e(a(y)x) \tilde{f}_1(x + \tilde{P}_{[r,V]}(z), y) \tilde{f}_2(x, y + \tilde{P}_{[r,V]}(z)) \right| \leq 2(N/V)^{3/2}.$$

Combining this with (6.12), we conclude that

$$\left| \Lambda'(\tilde{a}, \tilde{f}_1, \tilde{f}_2) \right| \geq (\delta/2)(N/V)^2 - 2(N/V)^{3/2} \geq (\delta/3)(N/V)^2.$$

Fix $C_1 > 0$. We now apply Proposition 6.13 with $\delta/3$ in place of δ and with $\tilde{a}$ in place of a . This yields $C := C_2 = C_2(C_1, P) > 0$ and $V' = V^{O_P(1)}$, a set $\mathcal{V} \subset [V']$ of size $|\mathcal{V}| \geq \exp(-\log(1/\delta)^C)V'$, and for each $k \in \mathcal{V}$ a phase $\alpha_k \in \mathbb{R}/\mathbb{Z}$ and a set $S_k \subseteq [N/(VV')]$ of size $|S_k| \geq \exp(-\log(1/\delta)^C)N/(VV')$ such that

$$\|V' \tilde{a}_{(k,V')}(y') - \alpha_k\|_{\mathbb{R}/\mathbb{Z}} \leq \exp(\log(1/\delta)^C)VV'/N \quad (6.13)$$

for all $y' \in S_k$. The separation property of the phases $\tilde{a}(y_\ell)$ ensures that for each $k \in \mathcal{V}$, there are at most $\exp(\log(1/\delta)^C)V'$ bad y ’s of the form $y_i = V'y' + k$ with $y' \in S_k$. Thus, for each $k \in \mathcal{V}$, there are at least

$$|S_k| - \exp(\log(1/\delta)^C)V' \gg \exp(-\log(1/\delta)^C)N/(VV')$$

values of $y' \in S_k$ such that $V'y' + k$ is good. For each such $y = V'y' + k$, recall that $\tilde{a}(y) = a(y)$, and consider the quantity $F(y)$, which we know has absolute value at least $(\delta/2)N/V$. For some constant $C' > C$, set

$N' \sim (C')^{-1} \exp(-2 \log(1/\delta)^{2C'}) N$ and reparamaterize the x -variable as

$$V'x + j + \frac{N'}{V}i,$$

where $j \in [V']$, $i \in [N/N']$, and now $x \in [N'/(VV')]$; this amounts to splitting the x -variable into arithmetic progressions of common difference V' and length $N'/(VV')$. Applying the triangle inequality in j, i , we find that

$$\sum_{j \in [V']} \sum_{i \in [N/N']} \left| \sum_{x \in [N'/(VV')]} \mathbb{E}_{z \in [(N/W_d)^{1/d} V^{-1}]} e \left(a(y)(V'x + j + \frac{N'}{V}i) \right) \right. \\ \left. \tilde{f}_1(V'x + j + \frac{N'}{V}i + \tilde{P}_{[r,V]}(z), y) \tilde{f}_2(V'x + j + \frac{N'}{V}i, y + \tilde{P}_{[r,V]}(z)) \right| \geq (\delta/2)N/V.$$

The phase $e(a(y)(j + iN'/V))$ is independent of x and can be absorbed into the absolute value; the remaining phase is $e(V'a(y)x)$. The key point is that now replacing $V'a(y)$ by α_k incurs a small loss, since (6.13) lets us estimate

$$|e(V'a(y)x) - e(\alpha_k x)| \leq \exp(\log(1/\delta)^C) \frac{VV'}{N} \cdot \frac{N'}{VV'} \leq \exp(-\log(1/\delta)^{C'})/2.$$

(It was for this estimate that we introduced the extra parameter N' and the summation over i .) Plugging this into the previous bound and using the triangle inequality, we deduce that

$$\sum_{j \in [V']} \sum_{i \in [N/N']} \left| \sum_{x \in [N'/(VV')]} \mathbb{E}_{z \in [(N/W_d)^{1/d} V^{-1}]} e(\alpha_k x) \tilde{f}_1(V'x + j + \frac{N'}{V}i + \tilde{P}_{[r,V]}(z), y) \right. \\ \left. \tilde{f}_2(V'x + j + \frac{N'}{V}i, y + \tilde{P}_{[r,V]}(z)) \right| \geq (\delta/2 - \exp(-\log(1/\delta)^{C'})/2)N/V \geq (\delta/3)N/V.$$

The number of y 's for which we have this estimate at least

$$\gg \exp(-\log(1/\delta)^C) V' \cdot \exp(-\log(1/\delta)^C) N/(VV') \geq \exp(-\log(1/\delta)^C) N/V$$

(adjusting the constant C). This, together with the nonnegativity of the contribution from the remaining y 's, establishes the proposition. $\square$

6.4.3 Auxiliary norm control on f_1 .

Recall from before that we have the lower bound

$$\sum_{x,y} \mathbb{E}_{z \in [(N/W_d)^{1/d} V - 1]} \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{1,j_1 + \tilde{P}(r), j_2}(x + \tilde{P}_{[r,V]}(z), y) \\ \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{2,j_1, j_2 + \tilde{P}(r)}(x, y + \tilde{P}_{[r,V]}(z)) g(x) e(a(y)x + b(y)) \geq \exp(-\log(1/\delta)^{O(1)})(N/V)^2$$

for at least

$$\exp(-\log(1/\delta)^{O(1)})(N/V)^{k+\ell-3}$$

tuples $(\underline{h}_1, \underline{h}_2) \in \mathbb{Z}^{k+\ell-3}$ and at least $\exp(-\log(1/\delta)^{O(1)})V^3$ values $(j_1, j_2, r) \in [V]^3$, where g, a, b depend on $\underline{h}_1, \underline{h}_2, j_1, j_2, r$. Proposition 6.14 gives an integer

$$\exp(-\log(1/\delta)^{O(1)})N \leq N' \leq N,$$

a positive integer power $V' = V^{O(1)}$, and 1-bounded functions $b_{1,j,j',i}, b_{2,j,j',i} : \mathbb{Z} \rightarrow \mathbb{C}$ (the latter dependent on $\underline{h}_1, \underline{h}_2, j_1, j_2, r$) such that

$$\sum_{i \in [N/N']} \sum_{j, j' \in [V']} \sum_y \sum_{x \in [N'/(VV')]} \mathbb{E}_{z \in [(N/W_d)^{1/d} V - 1]} b_{1,j,j',i}(x) b_{2,j,j',i}(y) \\ \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{1,j_1 + \tilde{P}(r), j_2}(V'x + j + \frac{N'}{V}i + \tilde{P}_{[r,V]}(z), V'y + j') \\ \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{1,j_1, j_2 + \tilde{P}(r)}(V'x + j + \frac{N'}{V}i, V'y + j' + \tilde{P}_{[r,V]}(z)) \\ \geq \exp(-\log(1/\delta)^{O(1)})(N/V)^2$$

for at least $\exp(-\log(1/\delta)^{O(1)})V^3$ values $(j_1, j_2, r) \in [V]^3$. Specifically, $b_{1,j,j',i}(x) = e(\alpha_{j'}x)$ for the phase $\alpha_{j'}$ provided by Proposition 6.14 while $b_{2,j,j',i}(y)$ is the phase of the average in the conclusion of Proposition 6.14. By the pigeonhole principle, we can choose a single $i \in [N/N']$ for which the sum over j, j', y, x, z satisfies the lower bound, at the cost of losing an extra factor of $\exp(-\log(1/\delta)^{O(1)})$, modifying $b_{1,j,j',i}$ by a phase if necessary, and dropping the index i from $b_{1,j,j',i}, b_{2,j,j',i}$. We then V' -trick the polynomial $\tilde{P}_{[r,V]}$, simplify the notation by setting

$$f_{j_1, j_2, j, j'}^l(x, y) := \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{l, j_1, j_2}(V'x + j + \frac{N'}{V}i, V'y + j') \\ = \Delta_{V\underline{h}_1 \mathbf{e}_1, V\underline{h}_2 \mathbf{e}_2} f_l(VV'x + Vj + j_1 + N'i, VV'y + Vj' + j_2) \quad (6.14)$$

for $l = 1, 2$, and shift $y \mapsto y - \tilde{P}_{[Vr'+r, VV']}(z)$, obtaining

$$\sum_{j, j', r' \in [V']} \sum_y \sum_{x \in [N'/(VV')]} \mathbb{E}_{z \in [(N/W_d)^{1/d}(VV')^{-1}]} b_{1,j,j'}(x) b_{2,j,j'}(y - \tilde{P}_{[Vr'+r, VV']}(z))$$

$$f_{j_1+\tilde{P}(r),j_2,j+\tilde{P}_{[r,V]}(r'),j'}^1(x+\tilde{P}_{[Vr'+r,VV']}(z),y-\tilde{P}_{[Vr'+r,VV']}(z)) \\ f_{j_1,j_2+\tilde{P}(r),j,j'+\tilde{P}_{[r,V]}(r')}^2(x,y) \geq \exp(-\log(1/\delta)^{O(1)})(N/V)^2$$

for the aforementioned values $(j_1, j_2, r) \in [V]^3$. Importantly, this new expression is amenable to an application of a Sárközy-type estimate. By Corollary 6.57 applied with

$$b_{1,j,j'}(x)1_{[N'/(VV')]}(x)f_{j_1,j_2+\tilde{P}(r),j,j'+\tilde{P}_{[r,V]}(r')}^2(x,y), \quad b_{2,j,j'}(y)f_{j_1+\tilde{P}(r),j_2,j+\tilde{P}_{[r,V]}(r'),j'}^1(x,y)$$

in place of f_0, f_1 and with $\mathbf{v} = \mathbf{e}_1 - \mathbf{e}_2$, we obtain the lower bound

$$\sum_{j,j',r' \in [V']} \sum_{x,y} \left| \sum_{z \in [\pm N''/(VV')]} b_{2,j,j'}(y - qz) f_{j_1+\tilde{P}(r),j_2,j+\tilde{P}_{[r,V]}(r'),j'}^1(x + qz, y - qz) \right| \\ \geq \exp(-\log(1/\delta)^{O(1)})(N/V)^3$$

for some positive integers $q \leq \exp(\log(1/\delta)^{O(1)})$ and $\exp(-\log(1/\delta)^{O(1)})N \leq N'' \leq N$. We apply the Cauchy–Schwarz inequality to remove the absolute value and then change variables $(x, y) \mapsto (x - qz, y + qz)$ and $z \mapsto -z$, so that

$$\sum_{j,j',r' \in [V']} \sum_{x,y,z} \mu_{N''/(VV')}(z) \Delta_{qz} b_{2,j,j'}(y) \Delta_{qz(\mathbf{e}_2 - \mathbf{e}_1)} f_{j_1+\tilde{P}(r),j_2,j+\tilde{P}_{[r,V]}(r'),j'}^1(x,y) \\ \geq \exp(-\log(1/\delta)^{O(1)})V'(N/V)^2.$$

In order to get rid of $b_{2,j,j'}$, we introduce an additional averaging over $[\pm N/(VV')]$ in the x -direction and then apply the Cauchy–Schwarz inequality to double this new variable. Using Lemma 6.7(4) to remove q , and summing over the admissible values of $j_1, j_2, r \in [V]$, we thus obtain

$$\sum_{\substack{j_1,j_2,r \in [V], \\ j,j',r' \in [V']}} \|f_{j_1+\tilde{P}(r),j_2,j+\tilde{P}_{[r,V]}(r'),j'}^1\|_{(\mathbf{e}_2 - \mathbf{e}_1) \cdot [\pm N/(VV')], \mathbf{e}_1 \cdot [\pm N/(VV')]}^4 \\ \geq \exp(-\log(1/\delta)^{O(1)})VV'N^2.$$

We now want to remove the shifts by $\tilde{P}(r)$ and $\tilde{P}_{[r,V]}(r')$ from the indices j_1, j' . Notice first that the function $f_{j_1,j_2,j+s_3V',j'+s_4V'}^1$ is a translate of the function $f_{j_1,j_2,j,j'}^1$ and hence $\|f_{j_1,j_2,j,j'}^1\|_{(\mathbf{e}_2 - \mathbf{e}_1) \cdot [\pm N/(VV')], \mathbf{e}_1 \cdot [\pm N/(VV')]}^4$ depends only on the residue classes of j, j' modulo V' . We can therefore replace the summation over

$j, j' \in [V']$ by a summation over $j, j' \in \mathbb{Z}/V'\mathbb{Z}$ and then shift $j \mapsto j - \tilde{P}_{[r,V]}(r')$ to obtain

$$\sum_{j_1, j_2, r \in [V]} \sum_{j, j' \in \mathbb{Z}/V'\mathbb{Z}} \|f_{j_1 + \tilde{P}(r), j_2, j, j'}^1\|_{(\mathbf{e}_2 - \mathbf{e}_1) \cdot [\pm N/(VV')], \mathbf{e}_1 \cdot [\pm N/(VV')]}^4 \geq \exp(-\log(1/\delta)^{O(1)})VN^2.$$

Similarly, the identity

$$f_{j_1 + s_1 V, j_2 + s_2 V, j, j'}^1 = f_{j_1, j_2, j + s_1, j' + s_2}^1,$$

makes $\sum_{j, j' \in \mathbb{Z}/V'\mathbb{Z}} \|f_{j_1 + \tilde{P}(r), j_2, j, j'}^1\|_{(\mathbf{e}_2 - \mathbf{e}_1) \cdot [\pm N/(VV')], \mathbf{e}_1 \cdot [\pm N/(VV')]}^4$ depend only on the residue class of j_1, j_2 modulo V , and so we can replace the sum over $j_1, j_2 \in [V]$ by the sum over $j_1, j_2 \in \mathbb{Z}/V\mathbb{Z}$. Shifting $j_1 \mapsto j_1 - \tilde{P}(r)$, we then get

$$\sum_{j_1, j_2 \in \mathbb{Z}/V\mathbb{Z}} \sum_{j, j' \in \mathbb{Z}/V'\mathbb{Z}} \|f_{j_1, j_2, j, j'}^1\|_{(\mathbf{e}_2 - \mathbf{e}_1) \cdot [\pm N/(VV')], \mathbf{e}_1 \cdot [\pm N/(VV')]}^4 \geq \exp(-\log(1/\delta)^{O(1)})N^2,$$

and so

$$\|\Delta_{Vh_1\mathbf{e}_1, Vh_2\mathbf{e}_2} f_1\|_{(\mathbf{e}_2 - \mathbf{e}_1)VV' \cdot [\pm N/(VV')], \mathbf{e}_1 VV' \cdot [\pm N/(VV')]}^4 \geq \exp(-\log(1/\delta)^{O(1)})N^2.$$

Since this holds for at least $\exp(-\log(1/\delta)^{O(1)})(N/V)^{k+\ell-3}$ choices of $(h_1, h_2) \in \mathbb{Z}^{k+\ell-3}$, we get from properties (1) and (6) of Lemma 6.7 that

$$\|f_1\|_{(\mathbf{e}_1 VV' \cdot [\pm N/(VV')])^{k-1}, (\mathbf{e}_2 - \mathbf{e}_1)VV' \cdot [\pm N/(VV')], (\mathbf{e}_2 VV' \cdot [\pm N/(VV')])^{\ell-1}}^{2^{k+\ell-1}} \geq \exp(-\log(1/\delta)^{O(1)})N^2. \quad (6.15)$$

6.4.4 Completing the proof of Proposition 6.12

So far, we have proved that if

$$|\Lambda^W(f_0, f_1, f_2)| \geq \delta N^2$$

implies that

$$\|f_0\|_{(\mathbf{e}_1 V \cdot [\pm N/V])^k, (\mathbf{e}_2 V \cdot [\pm N/V])^\ell}^{2^{k+\ell}} \geq \exp(-\log(1/\delta)^{O(1)})N^2,$$

then the norm-control (6.15) also follows. This lower bound is not particularly useful in itself, but we will use it as an intermediate step towards the desired norm-control on f_0 . To this end, we first use stashing to replace f_1 in Λ^W with the dual function

$$\mathcal{D}_1(f_0, f_2)(x, y) := \mathbb{E}_{z \in [(N/W_d)^{1/d}]} f_0(x - \tilde{P}(z), y) f_2(x - \tilde{P}(z), y + \tilde{P}(z)),$$

so that (6.15) holds with $\mathcal{D}_1(f_0, f_2)$ in place of f_1 . We then apply the dual-difference interchange (Lemma 6.9) to this expression, so that

$$\begin{aligned} \sum_{\substack{\underline{h}_1 \in \mathbb{Z}^{k-1}, \\ \underline{h}_2 \in \mathbb{Z}^{\ell-1}}} \mu_{N/(VV')}(\underline{h}_1, \underline{h}_2) \cdot \|\mathcal{D}_1^{\underline{h}_1, \underline{h}_2}(f_0, f_2)\|_{(\mathbf{e}_2 - \mathbf{e}_1)VV' \cdot [\pm N/(VV')]}^2 \\ \geq \exp(-\log(1/\delta)^{O(1)})N^2, \end{aligned}$$

where

$$\mathcal{D}_1^{\underline{h}_1, \underline{h}_2}(f_0, f_2)(x, y) := \mathbb{E}_{z \in [(N/W_d)^{1/d}]} \Delta_{VV'\underline{h}_1\mathbf{e}_1, VV'\underline{h}_2\mathbf{e}_2}(f_0(x - \tilde{P}(z), y) f_2(x - \tilde{P}(z), y + \tilde{P}(z))).$$

Thus for at least $\exp(-\log(1/\delta)^{O(1)})(N/(VV'))^{k+\ell-2}$ values of $(\underline{h}_1, \underline{h}_2) \in \mathbb{Z}^{k+\ell-2}$, we have the lower bound

$$\|\mathcal{D}_1^{\underline{h}_1, \underline{h}_2}(f_0, f_2)\|_{(\mathbf{e}_2 - \mathbf{e}_1)VV' \cdot [\pm N/(VV')]}^2 \geq \exp(-\log(1/\delta)^{O(1)})N^2.$$

We temporarily fix such a choice of $(\underline{h}_1, \underline{h}_2) \in \mathbb{Z}^{k+\ell-2}$ and repeat the above procedure of splitting the domains of x, y into arithmetic progressions of difference VV' . This yields

$$\begin{aligned} & \mathcal{D}_1^{\underline{h}_1, \underline{h}_2}(f_0, f_2)(VV'x + j_1, VV'y + j_2) \\ &= \mathbb{E}_{z \in [(N/W_d)^{1/d}]} \Delta_{VV'\underline{h}_1\mathbf{e}_1, VV'\underline{h}_2\mathbf{e}_2}(f_0(VV'x + j_1 - \tilde{P}(z), VV'y + j_2) \\ & \quad \times f_2(VV'x + j_1 - \tilde{P}(z), VV'y + j_2 + \tilde{P}(z))) \\ &= \mathbb{E}_{r \in [VV']} \mathbb{E}_{z \in [(N/W_d)^{1/d}(VV')^{-1}]} \Delta_{\underline{h}_1\mathbf{e}_1, \underline{h}_2\mathbf{e}_2}(f_{0, j_1 - \tilde{P}(r), j_2}(x - \tilde{P}_{[r, VV']}(z), y) \\ & \quad \times f_{2, j_1 - \tilde{P}(r), j_2 + \tilde{P}(r)}(x - \tilde{P}_{[r, VV']}(z), y + \tilde{P}_{[r, VV']}(z))), \end{aligned}$$

where $f_{l, j_1, j_2}(x, y) = (f_l)_{(j_1, j_2, VV')}(x, y) = f_l(VV'x + j_1, VV'y + j_2)$ for $l = 0, 2$. Hence

$$\begin{aligned}
& \mathbb{E}_{j_1, j_2, r \in [VV']} \left\| \mathbb{E}_{z \in [(N/W_d)^{1/d}(VV')^{-1}]} \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{0, j_1 - \tilde{P}(r), j_2}(\cdot - \tilde{P}_{[r, VV']}(z), \cdot) \right. \\
& \quad \left. \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{2, j_1 - \tilde{P}(r), j_2 + \tilde{P}(r)}(\cdot - \tilde{P}_{[r, VV']}(z), \cdot + \tilde{P}_{[r, VV']}(z)) \right\|_{(\mathbf{e}_2 - \mathbf{e}_1) \cdot [\pm N/(VV')]}^2 \\
& \geq \exp(-\log(1/\delta)^{O(1)}(N/(VV'))^2).
\end{aligned}$$

Applying the inverse theorem for the U^1 norm (Lemma 6.22)² for the values $(j_1, j_2, r) \in [VV']^3$ where the above norm is at least $\exp(-\log(1/\delta)^{O(1)}(N/(VV'))^2)$, we obtain a 1-bounded function $b : \mathbb{Z} \rightarrow \mathbb{C}$ (depending on $\underline{h}_1, \underline{h}_2, j_1, j_2, r$) such that

$$\begin{aligned}
& \sum_{x, y} \mathbb{E}_{z \in [(N/W_d)^{1/d}(VV')^{-1}]} b(x+y) \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{0, j_1 - \tilde{P}(r), j_2}(x - \tilde{P}_{[r, VV']}(z), y) \\
& \quad \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{2, j_1 - \tilde{P}(r), j_2 + \tilde{P}(r)}(x - \tilde{P}_{[r, VV']}(z), y + \tilde{P}_{[r, VV']}(z)) \\
& \geq \exp(-\log(1/\delta)^{O(1)}(N/(VV'))^2).
\end{aligned}$$

Changing variables $(x, y) \mapsto (x + \tilde{P}_{[r, VV']}(z), y - \tilde{P}_{[r, VV']}(z))$ simplifies this expression to

$$\begin{aligned}
& \sum_{x, y} \mathbb{E}_{z \in [(N/W_d)^{1/d}(VV')^{-1}]} b(x+y) \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{0, j_1 - \tilde{P}(r), j_2}(x, y - \tilde{P}_{[r, VV']}(z)) \\
& \quad \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{2, j_1 - \tilde{P}(r), j_2 + \tilde{P}(r)}(x, y) \\
& \geq \exp(-\log(1/\delta)^{O(1)}(N/(VV'))^2).
\end{aligned}$$

After an application of Corollary 6.57, with

$$b(x+y) \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{2, j_1 - \tilde{P}(r), j_2 + \tilde{P}(r)}(x, y) \quad \text{and} \quad \Delta_{\underline{h}_1 \mathbf{e}_1, \underline{h}_2 \mathbf{e}_2} f_{0, j_1 - \tilde{P}(r), j_2}(x, y)$$

in place of $f_0(x, y), f_1(x, y)$ and with $\mathbf{v} = \mathbf{e}_2$, maneuvers similar to those in the proof of Theorem 6.10 give

$$\|f_0\|_{(\mathbf{e}_1 \tilde{V} \cdot [N/\tilde{V}])^{k-1}, (\mathbf{e}_2 \tilde{V} \cdot [N/\tilde{V}])^\ell}^{2^{k+\ell-1}} \geq \exp(-\log(1/\delta)^{O(1)}) N^2$$

with $\tilde{V} = VV'$. Exchanging the roles of x and y gives the last assertion of the proposition.

²Technically, to apply this result, the length of the boxes should be larger by a factor $\exp(\log(1/\delta)^{O(1)})$ than the support of the function whose norm we are evaluating. This can be ensured by replacing the box $(\mathbf{e}_2 - \mathbf{e}_1)VV' \cdot [\pm N/(VV')]$ by $(\mathbf{e}_2 - \mathbf{e}_1)VV' \cdot [\pm N'/(VV')]$ (with $N' \asymp \exp(\log(1/\delta)^C)N$ for an appropriate $C > 0$) before performing the Dual-Difference Interchange using Lemma 6.7(4).

6.5 Reducing Proposition 6.13 to a nilsequence problem

The next two sections are devoted to the proof of Proposition 6.13; this will complete the degree-lowering argument from the previous section. Recall the counting operator

$$\Lambda'(a, \tilde{f}_1, \tilde{f}_2) = \sum_{x,y} \mathbb{E}_{z \in [(N/W_d)^{1/d} V^{-1}]} e(a(y)x) \tilde{f}_1(x + \tilde{P}_{[r,V]}(z), y) \tilde{f}_2(x, y + \tilde{P}_{[r,V]}(z));$$

our goal is to show that if $\Lambda'(a, \tilde{f}_1, \tilde{f}_2)$ is large for some 1-bounded functions $\tilde{f}_1, \tilde{f}_2$, then there are many arithmetic progressions of difference $V' = V^{O_P(1)}$ on each of which the phase a is approximately constant. In this section we will use higher-order Fourier analysis manipulations to replace the functions $\tilde{f}_1, \tilde{f}_2$ with nilsequences; in the next section we will use sophisticated machinery from the theory of nilsequences to complete the proof of Proposition 6.13. Our starting point is the following norm-control.

Proposition 6.15. *For every $C_1 > 0$ there exists $C_2 = C_2(C_1, P) > 0$ such that the following holds. Let $\delta \in (0, 1/10)$, let $V \leq W^{C_1}$ be a nonnegative integer power of W , and let $r \in [V]$. Let $\tilde{f}_1, \tilde{f}_2 : \mathbb{Z}^2 \rightarrow \mathbb{C}$ be 1-bounded functions supported on $[N/V]^2$, for some integer $N \geq C_2 V^{C_2} \delta^{-C_2}$, and let $a : \mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}$ be a phase function. If*

$$|\Lambda'(a, \tilde{f}_1, \tilde{f}_2)| \geq \delta(N/V)^2,$$

then we have the norm-control

$$\|\tilde{f}_2\|_{(\mathbf{e}_2 - \mathbf{e}_1) \cdot [\pm N/V], \mathbf{e}_1 \cdot [\pm N/V]}^4 \gg_{C_1, P} \delta^{O_{C_1, P}(1)} (N/V)^2.$$

Naïvely, one would expect to have $\{\mathbf{e}_2 - \mathbf{e}_1, \mathbf{e}_2\}$ -control on $\tilde{f}_2$. Our norm-control looks slightly different because, as will become clear in the proof, taking discrete derivatives in the $\mathbf{e}_1$ -direction is very convenient for dealing with the $e(a(y)x)$ term in Λ' . The argument would give the same conclusion with $\|\tilde{f}_1\|_{(\mathbf{e}_2 - \mathbf{e}_1) \cdot [\pm N/V], (\mathbf{e}_1 \cdot [\pm N/V])^2}$, but we do not need this.

Proof. We fix $C_1 > 0$ and let all the constants depend on C_1, P . We introduce an extra averaging over $[N/V]$ in the x -direction³ and apply the Cauchy–Schwarz inequality in order to double the variable x ; this gives

$$\sum_{x,y,h} \mu_{N/V}(h) \mathbb{E}_{z \in [(N/W_d)^{1/d} V^{-1}]} e(a(y)h) \Delta_{h\mathbf{e}_1} \tilde{f}_1(x + \tilde{P}_{[r,V]}(z), y) \Delta_{h\mathbf{e}_1} \tilde{f}_2(x, y + \tilde{P}_{[r,V]}(z)) \geq \delta^2 (N/V)^2. \quad (6.16)$$

³If we did not want the range of the new variable h to be restricted, then we could apply the Cauchy–Schwarz inequality directly.

Shifting $x \mapsto x - \tilde{P}_{[r,V]}(z)$ gives

$$\sum_{x,y,h} \mu_{N/V}(h) \mathbb{E}_{z \in [(N/W_d)^{1/d} V^{-1}]} e(a(y)h)$$

$$\Delta_{h\mathbf{e}_1} \tilde{f}_1(x, y) \Delta_{h\mathbf{e}_1} \tilde{f}_2(x - \tilde{P}_{[r,V]}(z), y + \tilde{P}_{[r,V]}(z)) \geq \delta^2 (N/V)^2.$$

The crucial point is that now only $\tilde{f}_2$ has $\tilde{P}_{[r,V]}(z)$ in its argument and so we can apply the Sárközy-type estimate (Corollary 6.57) (after using the popularity principle to find many h 's for which the inner expression is large in absolute value). Hence there exist positive integers⁴ $q \ll \delta^{-O(1)}$ and $\delta^{O(1)}N \ll N' \leq N$ such that

$$\sum_{x,y,h} \mu_{N/V}(h) \left| \mathbb{E}_{z \in [N'/V]} \Delta_{h\mathbf{e}_1} \tilde{f}_2(x - qz, y + qz) \right| \gg \delta^{O(1)} (N/V)^2.$$

Applying the Cauchy-Schwarz inequality and then using Lemma 6.7(4) to replace N' with N and drop q , we get the desired bound on $\tilde{f}_2$. $\square$

The next result shows that the largeness of $\Lambda'(a, \tilde{f}_1, \tilde{f}_2)$ for some functions $\tilde{f}_1, \tilde{f}_2$ implies some correlation of the phase a with nilsequences on arithmetic progressions. We refer the reader to Definition 6.37 for definition of a nilsequence.

Proposition 6.16. *For every $C_1 > 0$ there exists $C_2 = C_2(C_1, P) > 0$ such that the following holds. Let $\delta \in (0, 1/10)$, let $V \leq W^{C_1}$ be a positive integer power of W , and let $r \in [V]$. Let $a : \mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}$ be a phase function, and let $\tilde{f}_1, \tilde{f}_2 : \mathbb{Z}^2 \rightarrow \mathbb{C}$ be 1-bounded functions supported on $[N/V]^2$, for some integer $N \geq C_2 V^{C_2} \exp(\log(1/\delta)^{C_2})$. If*

$$|\Lambda'(a, \tilde{f}_1, \tilde{f}_2)| \geq \delta (N/V)^2,$$

then there exist a positive integer $k = O_P(1)$, a positive integer power $V' = V^{O_{C_1, P}(1)}$, and a function $\varphi : \mathbb{Z} \rightarrow \mathbb{C}$ such that

$$\begin{aligned} \mathbb{E}_{j, r' \in [V']} \sum_{y \in [N/(VV')]} \left| \mathbb{E}_{z \in [(N/W_d)^{1/d} (VV')^{-1}]} e(-a_{(j, V')}(y) V' \cdot \tilde{P}_{[Vr' + r, VV']}(z)) \right. \\ \left. \varphi_{(j + \tilde{P}_{[r, V]}(r'), V')}(y + \tilde{P}_{[Vr' + r, VV']}(z)) \right| \geq \exp(-\log(1/\delta)^{O_{C_1, P}(1)}) N/(VV') \end{aligned}$$

and for each $j \in [V']$, the restriction $\varphi_{(j, V')}$ is a 1-bounded, degree- k nilsequence of complexity $\exp(\log(1/\delta)^{O_{C_1, P}(1)})$ on a $\log(1/\delta)^{O_{C_1, P}(1)}$ -dimensional nilmanifold.

⁴A direct application of Corollary 6.57 gives q_h that depends on h ; however, q can be chosen uniformly by the pigeonhole principle since there are only $O(\delta^{-O(1)})$ possibilities for q_h .

Proof. The proof consists of massaging the expression $\Lambda'(a, \tilde{f}_1, \tilde{f}_2)$, with the help of the same techniques as in the proof of Theorem 6.10, until the structure of a along arithmetic progressions becomes apparent. We fix $C_1 > 0$ and let all constants depend on C_1, P without notating the dependence explicitly.

Step 1: Replacing $\tilde{f}_2$ by a box norm obstruction.

Our first objective is to use Proposition 6.15 in conjunction with relevant box norm inverse theorems and stashing to replace $\tilde{f}_2$ by a “structured” function. By Proposition 6.15, we have the lower bound

$$\|\tilde{f}_2\|_{(\mathbf{e}_2 - \mathbf{e}_1) \cdot [\pm N/V], \mathbf{e}_1 \cdot [\pm N/V]}^4 \gg \delta^{O(1)}(N/V)^2.$$

The stashing argument gives

$$\|\mathcal{D}'_2(a, \tilde{f}_1)\|_{(\mathbf{e}_2 - \mathbf{e}_1) \cdot [\pm N/V], \mathbf{e}_1 \cdot [\pm N/V]}^4 \gg \delta^{O(1)}(N/V)^2,$$

for the dual function

$$\begin{aligned} \mathcal{D}'_2(a, \tilde{f}_1)(x, y) &:= \mathbb{E}_{z \in [(N/W_d)^{1/d} V^{-1}]} e(a(y - \tilde{P}_{[r, V]}(z))x) \\ &\quad \tilde{f}_1(x + \tilde{P}_{[r, V]}(z), y - \tilde{P}_{[r, V]}(z)) 1_{[N/V]}(y), \end{aligned}$$

where we were able to insert the term $1_{[N/V]}$ because $\tilde{f}_2$ is supported on $[N/V]^2$. By the inverse theorem for the $U^1 \times U^1$ -box norm (Lemma 6.27), there are 1-bounded functions $b_1, b_2 : \mathbb{Z} \rightarrow \mathbb{C}$ such that

$$\begin{aligned} \sum_{x, y} \mathbb{E}_{z \in [(N/W_d)^{1/d} V^{-1}]} b_1(y) b_2(x + y) e(a(y - \tilde{P}_{[r, V]}(z))x) \\ \tilde{f}_1(x + \tilde{P}_{[r, V]}(z), y - \tilde{P}_{[r, V]}(z)) 1_{[N/V]}(y) \gg \delta^{O(1)}(N/V)^2, \end{aligned}$$

and the presence of the $1_{[N/V]}(y)$ term lets us assume that b_1 is supported on $[N/V]$. The change of variables $y \mapsto y + \tilde{P}_{[r, V]}(z)$ gives

$$\begin{aligned} \sum_{x, y} \mathbb{E}_{z \in [(N/W_d)^{1/d} V^{-1}]} e(a(y)x) \tilde{f}_1(x + \tilde{P}_{[r, V]}(z), y) \\ b_1(y + \tilde{P}_{[r, V]}(z)) b_2(x + y + \tilde{P}_{[r, V]}(z)) \gg \delta^{O(1)}(N/V)^2. \end{aligned}$$

We observe that $b_2(x + y + \tilde{P}_{[r, V]}(z))$ is a function of $(x + \tilde{P}_{[r, V]}(z), y)$. Shifting $x \mapsto x - \tilde{P}_{[r, V]}(z)$ and setting

$\tilde{f}_1'(x, y) := e(a(y)x)\tilde{f}_1(x, y)b_2(x + y)$, we get

$$\sum_{x, y} \mathbb{E}_{z \in [(N/W_d)^{1/d}V^{-1}]} e(-a(y)\tilde{P}_{[r, V]}(z))\tilde{f}_1'(x, y)b_1(y + \tilde{P}_{[r, V]}(z)) \gg \delta^{O(1)}(N/V)^2.$$

Recall that $\tilde{f}_1'$ is 1-bounded and supported on $[N/V]^2$. Using the triangle inequality in the x - and y -variables, we can eliminate $\tilde{f}_1'$ entirely and obtain

$$\sum_{y \in [N/V]} \left| \mathbb{E}_{z \in [(N/W_d)^{1/d}V^{-1}]} e(-a(y)\tilde{P}_{[r, V]}(z))b_1(y + \tilde{P}_{[r, V]}(z)) \right| \gg \delta^{O(1)}N/V.$$

We have thus reduced matters to analyzing the auxiliary operator

$$\Lambda''(b_1) := \sum_{y \in [N/V]} \left| \mathbb{E}_{z \in [(N/W_d)^{1/d}V^{-1}]} e(-a(y)\tilde{P}_{[r, V]}(z))b_1(y + \tilde{P}_{[r, V]}(z)) \right|. \quad (6.17)$$

Step 2: Replacing b_1 by a nilsequence.

Our next goal is to show that Λ'' can be controlled by a Gowers norm of b_1 via a standard argument based on the Cauchy-Schwarz argument; this will let us replace b_1 by a nilsequence (at least on arithmetic progressions of not-too-small difference). Let $b_3 : \mathbb{Z} \rightarrow \mathbb{C}$ be a 1-bounded function supported on $[N/V]$ such that

$$\Lambda''(b_1) = \sum_y b_3(y) \mathbb{E}_{z \in [(N/W_d)^{1/d}V^{-1}]} e(-a(y)\tilde{P}_{[r, V]}(z))b_1(y + \tilde{P}_{[r, V]}(z)) \gg \delta^{O(1)}N/V.$$

Shifting $y \mapsto y - \tilde{P}_{[r, V]}(z)$, stashing, and shifting back $y \mapsto y + \tilde{P}_{[r, V]}(z)$, we find that

$$\sum_y \overline{b_3(y)} \mathbb{E}_{z \in [(N/W_d)^{1/d}V^{-1}]} e(a(y)\tilde{P}_{[r, V]}(z))\tilde{b}_1(y + \tilde{P}_{[r, V]}(z)) \gg \delta^{O(1)}N/V,$$

where

$$\tilde{b}_1(y) := \mathbb{E}_{z \in [(N/W_d)^{1/d}V^{-1}]} e(-a(y - \tilde{P}_{[r, V]}(z))\tilde{P}_{[r, V]}(z))b_3(y - \tilde{P}_{[r, V]}(z))1_{[N/V]}(y).$$

Note that the indicator function $1_{[N/V]}(y)$ can be included since b_1 is supported on this set. Now, using the Cauchy-Schwarz inequality to double the z -variable, we eliminate $b_3(y)$ so that

$$\begin{aligned} \sum_{y, h_1} \mu_{(N/W_d)^{1/d}V^{-1}}(h_1) \mathbb{E}_{z \in [(N/W_d)^{1/d}V^{-1}]} e(a(y)(\tilde{P}_{[r, V]}(z) - \tilde{P}_{[r, V]}(z + h_1))) \\ \tilde{b}_1(y + \tilde{P}_{[r, V]}(z))\overline{\tilde{b}_1(y + \tilde{P}_{[r, V]}(z + h_1))} \gg \delta^{O(1)}N/V. \end{aligned}$$

While doing so, we have also replaced the polynomial $\tilde{P}_{[r,V]}(z)$ of degree d in z in the exponential phase by the polynomial $\tilde{P}_{[r,V]}(z) - \tilde{P}_{[r,V]}(z + h_1)$ of degree $d - 1$. Using the Cauchy–Schwarz inequality to double the z -variable d more times, we can eliminate the exponential involving $a(y)$ altogether. This gives

$$\sum_{y, h_1, \dots, h_{d+1}} \mu_{(N/W_d)^{1/d} V^{-1}}(h_1, \dots, h_{d+1}) \mathbb{E}_{z \in [(N/W_d)^{1/d} V^{-1}]} \prod_{\underline{\epsilon} \in \{0,1\}^{d+1}} \mathcal{C}^{|\underline{\epsilon}|} \tilde{b}_1(y + \tilde{P}_{[r,V]}(z + \underline{\epsilon} \cdot \underline{h})) \gg \delta^{O(1)} N/V.$$

By Proposition 5.41 (from the previous chapter) and the formula $V_d = \beta_d \beta_1^{d-2} V^{d-1} W^{d-1}$, we have

$$\left\| \tilde{b}_1 \right\|_{U^{k+1}(V_d \cdot [\pm N/(V V_d)])}^{2^{k+1}} \gg \delta^{O(1)} N/V$$

for some integer $k = O(1)$. Take $V' = V^{O(1)}$ to be the smallest positive integer power of V such that $V_d | \beta_d \beta_1 V'$. Since $\beta_1, \beta_d = O(1)$, we can apply parts (4) and (6) of Lemma 6.7 to replace the boxes $V_d \cdot [\pm N/(V V_d)]$ by $V' \cdot [\pm N/(V V')]$, obtaining

$$\left\| \tilde{b}_1 \right\|_{U^{k+1}(V' \cdot [\pm N/(V V')])}^{2^{k+1}} \gg \delta^{O(1)} N/V.$$

Splitting the domain of $\tilde{b}_1$ into arithmetic progressions of difference V' , we find using the popularity principle that

$$\left\| (\tilde{b}_1)_{(j, V')} \right\|_{U^{k+1}([\pm N/(V V')])}^{2^{k+1}} \gg \delta^{O(1)} N/(V V')$$

for an $\Omega(\delta^{O(1)})$ -proportion of $j \in [V']$. The U^{k+1} -inverse theorem [75, Theorem 1.2] tells us that for each such $j \in [V']$, there is a 1-bounded⁵ degree- k nilsequence φ_j of complexity $\exp(\log(1/\delta)^{O(1)})$ on an $O(\log(1/\delta)^{O(1)})$ -dimensional nilmanifold (see Definition 6.37 for the definition of a nilsequence and its complexity) such that

$$\sum_y (\tilde{b}_1)_{(j, V')}(y) \varphi_j(y) \geq \exp(-\log(1/\delta)^{O(1)}) N/(V V').$$

For each remaining value of $j \in [V']$, set $\varphi_j := 0$. Thus, φ_j is a 1-bounded degree- k nilsequence of complexity $\exp(\log(1/\delta)^{O(1)})$ on an $O(\log(1/\delta)^{O(1)})$ -dimensional nilmanifold for *each* $j \in [V']$.

⁵1-boundedness is not mentioned in the statement of [75, Theorem 1.2]; however, the control on the Lipschitz norms of φ_j ensures that the nilsequences coming from [75, Theorem 1.2] can be rescaled to be 1-bounded at an acceptable loss of the factor $\exp(-\log(1/\delta)^{O(1)})$ in the lower bound.

Define $\varphi : \mathbb{Z} \rightarrow \mathbb{C}$ by setting $\varphi_{(j,V')}(y) = \varphi(V'y + j)$ to equal $\varphi_j(y)$ for each $j \in [V']$ and $y \in \mathbb{Z}$, so that

$$\sum_y \tilde{b}_1(y) \varphi(y) \geq \exp(-\log(1/\delta)^{O(1)})N/V.$$

Expanding the definition of $\tilde{b}_1$, shifting $y \mapsto y + \tilde{P}_{[r,V]}(z)$, and using the triangle inequality to eliminate b_3 , we obtain

$$\sum_{y \in [N/V]} \left| \mathbb{E}_{z \in [(N/W_d)^{1/d}V^{-1}]} e(-a(y)\tilde{P}_{[r,V]}(z)) \varphi(y + \tilde{P}_{[r,V]}(z)) \right| \geq \exp(-\log(1/\delta)^{O(1)})N/V.$$

Since φ is a nilsequence only when restricted to arithmetic progressions of common difference V' , we also split the ranges of y and z into arithmetic progressions of common difference V' . This gives

$$\begin{aligned} \mathbb{E}_{j,r' \in [V']} \sum_{y \in [N/(VV')]} \left| \mathbb{E}_{z \in [(N/W_d)^{1/d}(VV')^{-1}]} e(-a_{(j,V')}(y)V' \cdot \tilde{P}_{[Vr'+r, VV']}(z)) \right. \\ \left. \varphi_{(j+\tilde{P}_{[r,V]}(r'), V')}(y + \tilde{P}_{[Vr'+r, VV']}(z)) \right| \geq \exp(-\log(1/\delta)^{O(1)})N/(VV'), \end{aligned}$$

as desired. $\square$

Notice that the last step of splitting into progressions introduced a factor of V' in the first term; this arises essentially because we V' -tricked the exponential function.

6.6 Nilsequence application

In the previous section, we arrived at expressions of the form

$$\sum_{y \in [N/\tilde{V}]} \left| \mathbb{E}_{z \in [(N/W_d)^{1/d}\tilde{V}^{-1}]} e(\tilde{a}(y) \cdot \tilde{P}_{[\tilde{r}, \tilde{V}]}(z)) \tilde{\varphi}(y + \tilde{P}_{[\tilde{r}, \tilde{V}]}(z)) \right| \quad (6.18)$$

with

$$\tilde{V} = VV', \quad \tilde{r} = Vr' + r, \quad \tilde{\varphi} = \varphi_{(j+\tilde{P}_{[r,V]}(r'), V')}, \quad \text{and} \quad \tilde{a} = -a_{(j,V')}V',$$

where $\tilde{\varphi}$ is a nilsequence. This section is devoted to extracting information about the phase a using the largeness of (6.18). The main underlying machinery comes from the equidistribution results of [73]. We refer the reader to Section 6.9 for definitions and notation related to nilsequences.

Our main technical input is the following result, which is the special case $\ell = 1$ of [73, Theorem 3].

Theorem 6.17. *For every $k \in \mathbb{N}$ there is a constant $C = C(k) > 0$ such that the following holds. Let G/Γ be a nilmanifold of degree k , step s , complexity at most M , and dimension m , with associated filtration $G_\bullet$. Let $\delta \in (0, 1/10)$, and let $N > (M/\delta)^{Cm^C}$ be an integer. Let $F : G/\Gamma \rightarrow \mathbb{C}$ be a 1-bounded M -Lipschitz function with a vertical character ξ satisfying $0 < |\xi| \leq M/\delta$.*

Suppose $g \in \text{poly}(\mathbb{Z}, G)$ with

$$\left| \mathbb{E}_{n \in [N]} F(g(n)\Gamma) \right| \geq \delta.$$

Then there exist $0 < r \leq \dim(G/[G, G])$ and horizontal characters $\eta_1, \dots, \eta_r$ of size at most $(M/\delta)^{O_k(m)^{O_k(1)}}$ satisfying

$$\|\eta_i \circ g\|_{C^\infty[N]} \leq (M/\delta)^{O_k(m)^{O_k(1)}},$$

and ξ annihilates all $(s-1)$ -fold commutators of elements in the subgroup $G' := \bigcap_{i=1}^r \ker(\eta_i)$.

The use of this theorem is that it allows one to iteratively “reduce” the step of g .

In the following proposition, we let $\tilde{V}_d$ denote the leading coefficient of $\tilde{P}_{[\tilde{r}, \tilde{V}]}$.

Proposition 6.18. *Let $s, k \in \mathbb{N}$ with $s \leq k$. For every $C_1 > 0$ there exist constants $C_2 = C_2(C_1, P, k)$, $C_3 = C_3(C_1, P, k) > 1$ such that the following holds. Let G/Γ be a nilmanifold of degree k , step s , complexity at most $M \geq 2$, and dimension $m \in \mathbb{N}$, with associated filtration $G_\bullet$, and let $g \in \text{poly}(\mathbb{Z}, G)$. Let $\delta \in (0, 1/10)$, and let $F : G/\Gamma \rightarrow \mathbb{C}$ be a 1-bounded M -Lipschitz function with a vertical character ξ satisfying $0 < |\xi| \leq M/\delta$. Let $\tilde{V} \leq W^{C_1}$ be a nonnegative integer power of W , let $\tilde{r} \in [\tilde{V}]$, and let $N \geq \tilde{V}_d \tilde{V}^{C_2} (M/\delta)^{2C_2 dm^{C_2}}$ be an integer. Lastly, let $A \subseteq [N/\tilde{V}]$ and $B \subseteq \mathbb{Z}$ be nonempty arithmetic progressions of common difference at most M/δ satisfying*

$$|B| \geq C_2 \cdot (N/(\tilde{V} \tilde{V}_d))^{1/d} (M/\delta)^{-C_2 m^{C_2}}. \quad (6.19)$$

If

$$\sum_{y \in A} \left| \mathbb{E}_{z \in B} e(\tilde{a}(y) \cdot \tilde{P}_{[\tilde{r}, \tilde{V}]}(z)) F(g(y + \tilde{P}_{[\tilde{r}, \tilde{V}]}(z))\Gamma) \right| \geq \delta |A|, \quad (6.20)$$

then the following holds with $T := |B|^d \tilde{V}_d$.

- *When $s > 1$: There exists a factorization $g(n) = \varepsilon(n)g'(\tilde{n})\gamma(n)$, where ε is $((M/\delta)^{C_3 m^{C_3}}, T)$ -smooth, g' takes values in an $(M/\delta)^{C_3 m^{C_3}}$ -rational normal subgroup G' of G such that $G'/\ker(\xi)$ has step at most $s-1$, and γ is $(M/\delta)^{C_3 m^{C_3}}$ -rational.*
- *When $s = 1$ and $k > 1$: There is a polynomial $p(n) = \sum_{j=0}^k \alpha_j n^j$ such that $F(g(n)\Gamma) = e(p(n))$, and there is a positive integer $q \leq C_3 \delta^{-C_3}$ such that $\|q\alpha_j\|_{\mathbb{R}/\mathbb{Z}} \leq C_3 \delta^{-C_3} / T^j$ for all $j \geq 2$.*

- When $s = 1$ and $k = 1$: There is a phase $\alpha \in \mathbb{R}/\mathbb{Z}$ such that $\|\tilde{a}(y) - \alpha\|_{\mathbb{R}/\mathbb{Z}} \leq C_3 \delta^{-C_3}/T$ for at least $\delta^{C_3}|A|$ values of $y \in A$.

We remark that in the third case, the pigeonhole principle trivially provides some α such that $\|\tilde{a}(y) - \alpha\|_{\mathbb{R}/\mathbb{Z}} \leq C_3 \delta^{-C_3}/T$ for at least $2\delta^{-C_3}|A|/T$ values of y ; the proposition improves the quantity $2\delta^{-C_3}|A|/T$ by a factor of around $\delta^{O(1)}T$, which is quite substantial since $|B|$ is large in applications.

Proof. We fix $C_1 > 0$ and let all constants depend on C_1, P, k (and hence also on s since $s \leq k$). We address the three bullet-pointed cases one-by-one.

Case 1: $s > 1$.

Suppose first that $s > 1$. By the pigeonhole principle, there is some $y_0 \in A$ such that

$$\left| \mathbb{E}_{z \in B} e(\tilde{a}(y_0) \tilde{P}_{[\tilde{r}, \tilde{v}]}(z)) F(g(y_0 + \tilde{P}_{[\tilde{r}, \tilde{v}]}(z)) \Gamma) \right| \geq \delta. \quad (6.21)$$

Since ξ is nonzero, there is some $g_s \in G_{(s)}$ such that $\xi(g_s) = a(y_0)$. Now define the polynomial sequence $\tilde{g} : \mathbb{Z} \rightarrow G$ (with respect to the same filtration $G_\bullet$) via

$$\tilde{g}(n) := g_s^n g(y_0 + n).$$

Since $\text{poly}(\mathbb{Z}, G)$ is closed under pointwise multiplication, we have $\tilde{g} \in \text{poly}(\mathbb{Z}, G)$. Notice that

$$\pi_{\text{horiz}}(\tilde{g}(n)) = \pi_{\text{horiz}}(g(y_0 + n))$$

since $g_s^n \in G_{(s)}$ has image zero in the horizontal torus. The identity

$$F(\tilde{g}(n) \Gamma) = e(\tilde{a}(y_0)n) F(g(y_0 + n) \Gamma)$$

then lets us rewrite (6.21) as

$$\left| \mathbb{E}_{z \in B} F(\tilde{g}(\tilde{P}_{[\tilde{r}, \tilde{v}]}(z)) \Gamma) \right| \geq \delta.$$

Now, consider the filtration $(\mathbb{Z}_i)_{i=0}^\infty$ on $\mathbb{Z}$ given by

$$\mathbb{Z}_i = \begin{cases} \mathbb{Z} & \text{if } i \leq d, \\ 0 & \text{if } i > d; \end{cases}$$

with respect to this filtration, we have $\tilde{P}_{[\tilde{r}, \tilde{v}]} \in \text{poly}(\mathbb{Z}, \mathbb{Z})$. By [107, Theorem 1.6.9], the composition

$\tilde{g} \circ \tilde{P}_{[\tilde{r}, \tilde{V}]}(z)$ is a polynomial sequence on G with a slightly modified filtration (now of degree dk).

Since

$$|B| \geq C_2 \cdot (N/(\tilde{V}\tilde{V}_d))^{1/d} (M/\delta)^{-C_2 m^{C_2}} \geq C(M/\delta)^{C_2 m^{C_2}}$$

and $C_2 > 0$ can be chosen to be sufficiently large, Theorem 6.17 gives us horizontal characters $\eta_1, \dots, \eta_r$ of size $O((M/\delta)^{O(m)^{O(1)}})$ (for some $0 < r \leq \dim(G/[G, G])$) such that each η_i satisfies

$$\|\eta_i \circ \tilde{g} \circ \tilde{P}_{[\tilde{r}, \tilde{V}]} \|_{C^\infty[|B|]} \ll (M/\delta)^{O(m)^{O(1)}}, \quad (6.22)$$

and such that ξ annihilates all $(s-1)$ -fold commutators of elements of $G' := \bigcap_{i=1}^r \ker(\eta_i)$. Since

$$|B| \geq C_2 \cdot (N/(\tilde{V}\tilde{V}_d))^{1/d} (M/\delta)^{-C_2 m^{C_2}} \geq C_2 \tilde{V}^{(C_2-1)/d} \quad (6.23)$$

and $C_2 > 0$ is sufficiently large, Lemma 6.47 tells us that

$$\|\eta_i \circ \tilde{g}\|_{C^\infty[T]} \ll (M/\delta)^{O(m)^{O(1)}}$$

for each η_i (recall that $T = |B|^d \tilde{V}_d$). Since

$$|y_0| \leq N/\tilde{V} \ll (M/\delta)^{O(m)^{O(1)}} T \quad (6.24)$$

by (6.19), Lemma 6.46 further implies that each

$$\|\eta_i \circ g\|_{C^\infty[T]} \ll (M/\delta)^{O(m)^{O(1)}}$$

(where the $O(1)$ in this inequality may differ from the $O(1)$ in the preceding inequality). The conclusion of the proposition now follows from [73, Lemma A.1].

Case 2: $s = 1, k > 1$.

Now suppose that $s = 1$. Since $s = 1$, there is a polynomial $p(n) = \sum_{j=0}^k \alpha_j n^j$ such that $F(g(n)\Gamma) = e(p(n))$. Substituting this into (6.20) and pigeonholing in y gives that

$$\left| \mathbb{E}_{z \in [B]} e \left(\tilde{a}(y_0) \tilde{P}_{[\tilde{r}, \tilde{V}]}(z) + \sum_{j=0}^k \alpha_j (y_0 + \tilde{P}_{[\tilde{r}, \tilde{V}]}(z))^j \right) \right| \geq \delta$$

for some $y_0 \in A$. Weyl's inequality (see, e.g., [51, Lemma 4.4]) tells us that there is a positive integer

$q \ll \delta^{-O(1)}$ such that

$$\left\| q \cdot R(y_0 + \tilde{P}_{[\tilde{r}, \tilde{V}]}(z)) \right\|_{C^\infty[|B|]} \ll \delta^{-O(1)}.$$

for $R(n) := p(n) + \tilde{a}(y_0)(n - y_0)$. Applying Lemma 6.47 with (6.24) and then Lemma 6.46 with (6.23), we find that $\|q\alpha_j\|_{\mathbb{R}/\mathbb{Z}} \ll \delta^{-O(1)}/T^j$ for all $j \geq 2$. (We do not make use of the conclusion that $\|q(\tilde{a}(y_0) + \alpha_1)\|_{\mathbb{R}/\mathbb{Z}} \ll \delta^{-O(1)}/T$.) This concludes the proof of the proposition in the second case.

Case 3: $s = k = 1$.

Finally, suppose that $s = 1$ and moreover $k = 1$. Then there are $\alpha_0, \alpha_1 \in \mathbb{R}/\mathbb{Z}$ such that $F(g(n)\Gamma) = e(\alpha_1 n + \alpha_0)$. Now (6.20) reads

$$\sum_{y \in A} \left| \mathbb{E}_{z \in [B]} e(\tilde{a}(y) \tilde{P}_{[\tilde{r}, \tilde{V}]}(z) + \alpha_1(y + \tilde{P}_{[\tilde{r}, \tilde{V}]}(z))) \right| \geq \delta|A|.$$

The popularity principle provides at least $\delta|A|/2$ values of $y \in |A|$ such that

$$\left| \mathbb{E}_{z \in [B]} e(R_y \circ \tilde{P}_{[\tilde{r}, \tilde{V}]}(z)) \right| \geq \delta/2,$$

where $R_y(n) := (\tilde{a}(y) + \alpha_1)n$. Hence, by Weyl's inequality, we have

$$\left\| q_y R_y \circ \tilde{P}_{[\tilde{r}, \tilde{V}]} \right\|_{C^\infty(|B|)} \ll \delta^{-O(1)}$$

for some positive integer $q_y \ll \delta^{-O(1)}$. Lemma 6.47 then gives $\|q_y R_y\|_{C^\infty(T)} \ll \delta^{-O(1)}$, which is tantamount to

$$\|q_y(\tilde{a}(y) + \alpha_1)\|_{\mathbb{R}/\mathbb{Z}} \ll \delta^{-O(1)}/T.$$

Pigeonholing in y , we can find q such that this bound holds for $\gg \delta^{O(1)}|A|$ values of $y \in A$. Finally, pigeonholing in the shifts of α_1 by integer multiples of $1/q$, we can eliminate q . $\square$

6.6.1 Completing the proof of Proposition 6.13

We are now in a position to complete the proof of Proposition 6.13. We let all constants depend on P . By Proposition 6.16, there exist a positive integer $k = O(1)$, a positive integer power $V' = V^{O(1)}$, and a function $\varphi : \mathbb{Z} \rightarrow \mathbb{C}$ such that for each $j \in [V']$, the restriction $\varphi_{(j, V')}$ is a 1-bounded, degree- k nilsequence of complexity $\exp(\log(1/\delta)^{O(1)})$ on a $\log(1/\delta)^{O(1)}$ -dimensional nilmanifold. Moreover, we have

$$\mathbb{E}_{j,r' \in [V']} \sum_{y \in [N/(VV')]} \left| \sum_{z \in [(N/W_d)^{1/d}(VV')^{-1}]} \mathbb{E} e(-a_{(j,V')}(y)V' \cdot \tilde{P}_{[Vr'+r, VV']}(z)) \right. \\ \left. \varphi_{(j+\tilde{P}_{[r,V]}, V')}(y + \tilde{P}_{[Vr'+r, VV']}(z)) \right| \geq \exp(-\log(1/\delta)^{O(1)})N/(VV').$$

By the popularity principle, there exists $r' \in [V']$ and a set $\mathcal{V} \subseteq [V']$ of size $|\mathcal{V}| \geq \exp(-\log(1/\delta)^{O(1)})V'$ such that for each $j \in \mathcal{V}$, we have the bound

$$\sum_{y \in [N/\tilde{V}]} \left| \sum_{z \in [(N/W_d)^{1/d}\tilde{V}^{-1}]} \mathbb{E} e(\tilde{a}_{(j,V')}(y) \cdot \tilde{P}_{[\tilde{r}, \tilde{V}]}(z)) \tilde{\varphi}_j(y + \tilde{P}_{[\tilde{r}, \tilde{V}]}(z)) \right| \geq \eta N/\tilde{V}, \quad (6.25)$$

where $\eta = \exp(-\log(1/\delta)^{O(1)})$ and

$$\tilde{V} = VV', \quad \tilde{r} = Vr' + r, \quad \tilde{a}(y) = -a(y)V', \quad \text{and} \quad \tilde{\varphi}_j(y) = \varphi_{(j+\tilde{P}_{[r,V]}, V')}(y).$$

Fix $j \in \mathcal{V}$. We will iteratively apply Proposition 6.18. This will let us first reduce to the case where (6.25) holds with $\tilde{\varphi}_j$ replaced by a degree-1 nilsequence on a 1-dimensional nilmanifold, and then deduce that the phase function $\tilde{a}_{(j,V')}$ is nearly constant for a positive proportion of y 's. Our proof proceeds by iteratively reducing the step of the nilmanifold underlying $\tilde{\varphi}_j$. At each stage, we will have parameters $\eta, A, B, g, G/\Gamma$, and F . Since the step is initially at most $\log(1/\delta)^{O(1)}$ and we decrease the step at each stage, there are at most $\log(1/\delta)^{O(1)}$ stages. Keeping track of how our bounds degrade at each stage, we can guarantee that throughout the iterative process we always have the bounds

$$\eta \geq \exp(-\log(1/\delta)^{O(1)}), \quad N \gg \tilde{V}_d \tilde{V}^{\Omega(1)} \eta^{-O(m)^{O(1)}}, \quad |A| \gg \eta^{O(m)^{O(1)}} \frac{N}{\tilde{V}}, \\ |B| \gg \eta^{O(m)^{O(1)}} (N/(\tilde{V}\tilde{V}_d))^{1/d}, \quad T = |B|^d \tilde{V}_d \gg \eta^{O(m)^{O(1)}} \frac{N}{\tilde{V}}, \\ M \ll \eta^{-O(m)^{O(1)}}, \quad \text{and} \quad m \ll \log(1/\delta)^{O(1)}.$$

We remind the reader that m, M stand for dimension and complexity of the underlying nilmanifold, and we observe that the frequently appearing quantity $\eta^{O(m)^{O(1)}}$ is of the shape $\exp(-\log(1/\delta)^{O(1)})$. Our lower bound $N \gg V^{\Omega(1)} \exp(\log(1/\delta)^{\Omega(1)})$ guarantees that the iteration will not halt before the claimed result is proven.

We begin with $A = [N/\tilde{V}]$ and $B = [(N/(\tilde{V}\tilde{V}_d))^{1/d}] = [(N/W_d)^{1/d}\tilde{V}^{-1}]$. We also start with $\eta \asymp \exp(-\log(1/\delta)^{O(1)})$, for which the lower bound in (6.25) holds.

Step 1: Fourier-expanding F .

In order to apply Proposition 6.18, we will first Fourier-expand the function F . Suppose that $F : G/\Gamma \rightarrow \mathbb{C}$

is an M -Lipschitz function on a nilmanifold G/Γ of complexity M , and that

$$\sum_{y \in A} \left| \mathbb{E}_{z \in B} e(\tilde{a}_{(j, \tilde{V})}(y) \tilde{P}_{[\tilde{r}, \tilde{V}]}(z)) F(g(y + \tilde{P}_{[\tilde{r}, \tilde{V}]}(z))\Gamma) \right| \geq \eta |A|. \quad (6.26)$$

By [73, Lemma A.6], we can Fourier-expand F along the $G'_{(s)}$ -torus as

$$\tilde{F} = \sum_{|\xi| \leq (M/\eta)^{O(m)^{O(1)}}} F_\xi + O((\eta/M)^{O(m)^{O(1)}}),$$

where each F_ξ is an $(M/\eta)^{O(m)^{O(1)}}$ -Lipschitz vertical character of frequency ξ . By replacing η with $(\eta/M)^{O(m)^{O(1)}}$, we may assume that $F = F_\xi$ and that F is 1-Lipschitz. Thus, we have

$$\sum_{y \in A} \left| \mathbb{E}_{z \in B} e(\tilde{a}_{(j, \tilde{V})}(y) \tilde{P}_{[\tilde{r}, \tilde{V}]}(z)) F(g(y + \tilde{P}_{[\tilde{r}, \tilde{V}]}(z))\Gamma) \right| \gg (\eta/M)^{O(m)^{O(1)}} |A| \quad (6.27)$$

(The purpose of these manipulations is to replace the function F in (6.26) by a vertical character.) Finally, we relabel $(\eta/M)^{O(m)^{O(1)}}$ as simply η so our hypothesis becomes

$$\sum_{y \in A} \left| \mathbb{E}_{z \in B} e(\tilde{a}_{(j, \tilde{V})}(y) \tilde{P}_{[\tilde{r}, \tilde{V}]}(z)) F(g(y + \tilde{P}_{[\tilde{r}, \tilde{V}]}(z))\Gamma) \right| \gg \eta |A|. \quad (6.28)$$

Step 2: Applying Proposition 6.18.

Proposition 6.18 applied to (6.28) supplies a factorization $g = \varepsilon g' \gamma$, where γ is an $(M/\eta)^{O(m)^{O(1)}}$ -rational sequence, ε is an $((M/\eta)^{O(m)^{O(1)}}, T)$ -smooth sequence, and g' is a sequence lying in an $(M/\eta)^{O(m)^{O(1)}}$ -rational subnilmanifold G'/Γ' of step $\leq s-1$ for some normal subgroup $G' \subseteq G$ with lattice $\Gamma' := G' \cap \Gamma$. Since ε is $((M/\eta)^{O(m)^{O(1)}}, T)$ -smooth, it is also $(\sigma^{-1}, |A|)$ -smooth for some $\sigma = (\eta/M)^{O(m)^{O(1)}}$.

Step 3: Pigeonholing in a subprogression.

Let $R \ll (M/\eta)^{O(m)^{O(1)}}$ be the period⁶ of γ . We now pigeonhole in subprogressions A' of A and B' of B , of common difference R and lengths

$$(\eta/M)^{O(m)^{O(1)}} |A| \ll |A'| \leq \sigma |A| \eta^2 \quad \text{and} \quad (\eta/M)^{O(m)^{O(1)}} |B| \ll |B'| \leq \sigma |B| \eta^2,$$

to find some such A', B' such that

$$\sum_{y \in A'} \left| \mathbb{E}_{z \in B'} e(\tilde{a}_{(j, \tilde{V})}(y) \tilde{P}_{[\tilde{r}, \tilde{V}]}(z)) F(g(y + \tilde{P}_{[\tilde{r}, \tilde{V}]}(z))\Gamma) \right| \geq (\eta/2) |A'|.$$

⁶We use here [73, Lemma B.14], which says that if γ is Q -rational, then it is $O(Q^{O(1)})$ -periodic.

Recall that we have the factorization

$$g(y + \tilde{P}_{[\tilde{r}, \tilde{v}]}(z)) = \varepsilon(y + \tilde{P}_{[\tilde{r}, \tilde{v}]}(z))g'(y + \tilde{P}_{[\tilde{r}, \tilde{v}]}(z))\gamma(y + \tilde{P}_{[\tilde{r}, \tilde{v}]}(z)).$$

The term $\gamma(y + \tilde{P}_{[\tilde{r}, \tilde{v}]}(z))\Gamma$ is constant for $y \in A'$ and $z \in B'$ because A', B' have common difference R .

Thus, there is some $\gamma_0 \in G$ such that for such y, z , our factorization can be written as

$$\varepsilon(y + \tilde{P}_{[\tilde{r}, \tilde{v}]}(z))g'(y + \tilde{P}_{[\tilde{r}, \tilde{v}]}(z))\gamma_0,$$

up to right-multiplication by Γ . By the smoothness properties of ε , we have

$$\begin{aligned} d_{G/\Gamma}(\varepsilon(y + \tilde{P}_{[\tilde{r}, \tilde{v}]}(z)), \varepsilon_0) &\leq (\sigma^{-1}/|A|) \cdot |A'| \leq \eta^2 \leq \eta/4, \\ \text{and } d_{G/\Gamma}(\varepsilon_0, \text{id}_G) &\leq \sigma^{-1} \ll (M/\eta)^{O(m)^{O(1)}} \end{aligned}$$

for some $\varepsilon_0 \in G$. Then, by the triangle inequality (recall that F has Lipschitz constant at most 1), we have

$$\sum_{y \in A'} \left| \mathbb{E}_{z \in B'} e(\tilde{a}_{(j, \tilde{v})}(y) \tilde{P}_{[\tilde{r}, \tilde{v}]}(z)) F(\varepsilon_0 g'(y + \tilde{P}_{[\tilde{r}, \tilde{v}]}(z)) \gamma_0 \Gamma) \right| \geq (\eta/4) |A'|.$$

By [73, Lemma B.6], we can write $\gamma_0 = \{\gamma_0\}[\gamma_0]$ where $[\gamma_0] \in \Gamma$ and $|\psi(\{\gamma_0\})| \leq \frac{1}{2}$. Expanding

$$\varepsilon_0 g'(y + \tilde{P}_{[\tilde{r}, \tilde{v}]}(z)) \gamma_0 = \varepsilon_0 \{\gamma_0\} (\{\gamma_0\}^{-1} g'(y + \tilde{P}_{[\tilde{r}, \tilde{v}]}(z)) \{\gamma_0\}) [\gamma_0]$$

and setting $\tilde{F} := F|_{G'}(\varepsilon_0 \{\gamma_0\} \cdot)$ (this function is $O((M/\eta)^{O(m)^{O(1)}})$ -Lipschitz by [73, Lemma B.4]), we deduce from the normality of G' in G that

$$\{\gamma_0\}^{-1} g'(y + \tilde{P}_{[\tilde{r}, \tilde{v}]}(z)) \{\gamma_0\} \in \text{poly}(\mathbb{Z}, G')$$

and that

$$\sum_{y \in A'} \left| \mathbb{E}_{z \in B'} e(\tilde{a}_{(j, \tilde{v})}(y) \tilde{P}_{[\tilde{r}, \tilde{v}]}(z)) \tilde{F}(\{\gamma_0\}^{-1} g'(y + \tilde{P}_{[\tilde{r}, \tilde{v}]}(z)) \{\gamma_0\}) \Gamma' \right| \gg (\eta/M)^{O(m)^{O(1)}} |A'|.$$

Step 4: Completing the iteration.

We continue the iteration by redefining

$$\begin{aligned} A &\rightsquigarrow A', \quad B \rightsquigarrow B', \quad g(n) \rightsquigarrow \{\gamma_0\}^{-1} g'(n) \{\gamma_0\}, \\ G/\Gamma &\rightsquigarrow G'/\Gamma', \quad \eta \rightsquigarrow (\eta/M)^{O(m)^{O(1)}}, \quad \text{and} \quad F \rightsquigarrow \tilde{F}. \end{aligned}$$

Step 5: The special case $s = 1$.

When $s = 1$ and $k \geq 2$, we have $F(g(n)\Gamma) = e(p(n))$ for some $p(n) = \sum_{j=0}^k \alpha_j n^j \in \mathbb{R}[n]$. Applying Proposition 6.18, we see that there exists some $q \ll \eta^{-O(m)^{O(1)}}$ such that

$$\|q\alpha_j\|_{\mathbb{R}/\mathbb{Z}} \ll \frac{(M/\eta)^{O(m)^{O(1)}}}{T^j} \quad \text{for all } 2 \leq j \leq k.$$

Notice that when n is restricted to an arithmetic progression of difference q , the function $\sum_{j=2}^k \alpha_j n^j$ is slowly-varying and $p(n)$ is very close to a constant plus $\alpha_1 n$. Combining this with the pigeonholing argument from Step 2, which lets us replace A, B by arithmetic progressions of difference q , we reduce to the case where $k = 1$.

Finally, when $s = 1$ and $k = 1$, the last case of Proposition 6.18 gives for each $j \in \mathcal{V}$ a phase $\alpha_j \in \mathbb{R}/\mathbb{Z}$ such that

$$\begin{aligned} \|a_{(j, \tilde{V})}(y) + \alpha_j\|_{\mathbb{R}/\mathbb{Z}} &= \|a_{(j, \tilde{V})}(y)V' - \tilde{\alpha}_j\|_{\mathbb{R}/\mathbb{Z}} \ll \frac{(M/\eta)^{O(m)^{O(1)}}}{T} \\ &\ll \frac{(M/(\delta\eta))^{O(m)^{O(1)}}\tilde{V}}{N} \leq \exp(\log(1/\delta)^{O(1)}) \frac{\tilde{V}}{N} \end{aligned}$$

for at least $\exp(-(\log(1/\delta))^{O(1)})N/\tilde{V}$ values of $y \in [N/\tilde{V}]$. For $j \notin \mathcal{V}$, we define α_j arbitrarily. This completes the proof of Proposition 6.13.

6.7 Finishing the argument

Recall that

$$\Lambda^{\text{Model}}(f_0, f_1, f_2) = \sum_{x, y} \mathbb{E}_{z \in [N]} f_0(x, y) f_1(x + z, y) f_2(x, y + z) \nu(z)$$

where $\nu(z) = d^{-1} 1_{[N]}(z)(N/z)^{(d-1)/d}$. Our goal in this section is to compare Λ^W with Λ^{Model} . We already know from Theorem 6.10 that Λ^W is controlled by suitable degree-2 box norms of f_0, f_1, f_2 , and we now show that the same holds for Λ^{Model} .

Proposition 6.19. *Let $\delta \in (0, 1/10)$, let $N \in \mathbb{N}$, and let $f_0, f_1, f_2 : \mathbb{Z}^2 \rightarrow \mathbb{C}$ be 1-bounded functions supported*

on $[N]^2$. If

$$\Lambda^{\text{Model}}(f_0, f_1, f_2) \geq \delta N^3, \quad (6.29)$$

then $\|f_0\|_{\mathbf{e}_1 \cdot [\pm N], \mathbf{e}_2 \cdot [\pm N]}^4, \|f_1\|_{\mathbf{e}_1 \cdot [\pm N], (\mathbf{e}_2 - \mathbf{e}_1) \cdot [\pm N]}^4, \|f_2\|_{\mathbf{e}_2 \cdot [\pm N], (\mathbf{e}_2 - \mathbf{e}_1) \cdot [\pm N]}^4 \gg_d \delta^{O_d(1)} N^2$.

Proof. We will replace ν by a smoother weight function that is easier to work with. Let $\varepsilon > 0$ be a small parameter to be chosen later, and let $\varphi : \mathbb{Z} \rightarrow [0, 1]$ be a $4\varepsilon^{-1}$ -Lipschitz function that equals 1 on $[2\varepsilon N, (1 - \varepsilon)N - 1]$ and vanishes outside of $[\varepsilon N, N - 1]$. Define the cut-off weight function

$$\nu_\varepsilon(z) := d^{-1}(N/(z + 1))^{(d-1)/d} \varphi(z).$$

Then ν_ε equals ν away from the boundary of $[N]$, and we have the ℓ^1 -bound

$$\sum_{z \in \mathbb{Z}} |\nu_\varepsilon(z) - \nu(z)| \ll \varepsilon^{1/d} N.$$

Our construction of ν_ε ensures that it is $O(\varepsilon^{-3}N^{-1})$ -Lipschitz. Define the function $g : \mathbb{R}/\mathbb{Z} \rightarrow [0, 1]$ by setting $g(z) := \varepsilon^{1-1/d} \nu_\varepsilon(zN)$ for $0 \leq z < 1$; it follows that g is $O(\varepsilon^{-3})$ -Lipschitz. By [80, Lemma 4.1], there is a trigonometric polynomial $h : \mathbb{R}/\mathbb{Z} \rightarrow [0, 1]$ of the form

$$h(x) = \sum_{m \ll \varepsilon^{-10}} a_m e(mx),$$

with each $|a_m| \leq 1$, such that

$$\|g - h\|_\infty \leq \varepsilon.$$

Unwinding the various definitions, we see that

$$\left| \nu_\varepsilon(z) - \varepsilon^{1/d-1} \sum_{m \ll \varepsilon^{-10}} a_m e(mz/N) \right| \leq \varepsilon^{1/d}$$

for all $z \in [N]$. This, combined with our earlier ℓ^1 -bound, implies that $\Lambda^{\text{Model}}(f_0, f_1, f_2)$ differs by at most $O(\varepsilon^{1/d} N^3)$ from

$$\sum_{x, y} \sum_{z \in [N]} f_0(x, y) f_1(x + z, y) f_2(x, y + z) \cdot \varepsilon^{1/d-1} \sum_{m \ll \varepsilon^{-10}} a_m e(mz/N).$$

In particular, if we set ε to be a sufficiently small constant times δ^d , then the latter quantity has absolute

value at least $\delta N^3/2$. Applying the triangle inequality and distributing the phases $e(mz/N)$, we obtain

$$\sum_{m \ll \delta^{-10d}} \left| \sum_{x,y} \sum_{z \in [N]} f_0(x,y) f_1^{(m)}(x+z,y) f_2^{(m)}(x,y+z) \right| \gg \delta^{2d} N^3,$$

where

$$f_1^{(m)}(x,y) := f_1(x,y) e(mx/N) \quad \text{and} \quad f_2^{(m)}(x,y) := f_2(x,y) e(-mx/N)$$

are still 1-bounded functions supported on $[N]$. Pigeonholing, we find some m such that

$$\left| \sum_{x,y} \sum_{z \in [N]} f_0(x,y) f_1^{(m)}(x+z,y) f_2^{(m)}(x,y+z) \right| \gg \delta^{12d} N^3,$$

and then the desired norm-control on f_0 follows from the usual Cauchy–Schwarz maneuvers. The norm-control on f_1, f_2 follows analogously. $\square$

We are finally ready to show that Λ^W and Λ^{Model} are always close together.

Proposition 6.20. *There exists $C = C(P) > 0$ such that the following holds. Let $f_0, f_1, f_2 : \mathbb{Z}^2 \rightarrow \mathbb{C}$ be 1-bounded functions supported on $[N]^2$. If N, W satisfy $N \geq CW^C$ and $W \geq \exp \exp(\log(1/\delta)^C)$, then*

$$|\Lambda^W(f_0, f_1, f_2) - \Lambda^{\text{Model}}(f_0, f_1, f_2)| < \delta N^2.$$

Proof. We let all constants depend on P . Define the “difference” counting operator

$$\Lambda^*(f_0, f_1, f_2) := \Lambda^W(f_0, f_1, f_2) - \Lambda^{\text{Model}}(f_0, f_1, f_2)$$

and the dual function

$$\begin{aligned} \mathcal{D}_0^*(f_1, f_2)(x, y) &:= \mathbb{E}_{z \in [(N/W_d)^{1/d}]} f_1(x + \tilde{P}(z), y) f_2(x, y + \tilde{P}(z)) \\ &\quad - \mathbb{E}_{z \in [N]} f_1(x+z, y) f_2(x, y+z) \nu(z). \end{aligned}$$

Assume for the sake of contradiction that

$$|\Lambda^*(f_0, f_1, f_2)| \geq \delta N^2;$$

then stashing gives

$$|\Lambda^*(\mathcal{D}_0^*(f_1, f_2), f_1, f_2)| \geq \delta^2 N^2,$$

and we conclude by the triangle inequality that either

$$|\Lambda^W(\mathcal{D}_0^*(f_1, f_2), f_1, f_2)| \geq \delta^2 N^2/2 \quad \text{or} \quad |\Lambda^{\text{Model}}(\mathcal{D}_0^*(f_1, f_2), f_1, f_2)| \geq \delta^2 N^2/2.$$

According to which of these inequalities holds, we apply either Theorem 6.10 or Proposition 6.19 and Lemma 6.7(6) to find a positive integer power $V = W^{O(1)}$ for which

$$\|\mathcal{D}_0^*(f_1, f_2)\|_{\mathbf{e}_1 V \cdot [\pm N/V], \mathbf{e}_2 V \cdot [\pm N/V]}^4 \geq \exp(-\log(1/\delta)^{O(1)}) N^2.$$

As usual, we must pass to arithmetic progressions of difference V before using the $U^1 \times U^2$ -inverse theorem.

Write $g_{0,j_1,j_2} := (\mathcal{D}_0^*(f_1, f_2))_{(j_1,j_2,V)}$, so that

$$\mathbb{E}_{j_1,j_2 \in [V]} \|g_{0,j_1,j_2}\|_{\mathbf{e}_1 \cdot [\pm N/V], \mathbf{e}_2 \cdot [\pm N/V]}^4 \geq \exp(-\log(1/\delta)^{O(1)}) (N/V)^2.$$

The inverse theorem for the $U^1 \times U^1$ -box norm (Lemma 6.27) gives 1-bounded functions $b_{1,j_1,j_2}, b_{2,j_1,j_2} : \mathbb{Z} \rightarrow \mathbb{C}$, for each $j_1, j_2 \in [V]$, such that

$$\mathbb{E}_{j_1,j_2 \in [V]} \sum_{x,y} g_{0,j_1,j_2}(x,y) b_{1,j_1,j_2}(x) b_{2,j_1,j_2}(y) \geq \exp(-\log(1/\delta)^{O(1)}) (N/V)^2.$$

Define the weights

$$\tilde{\nu}(z) := \frac{N}{(N/W_d)^{1/d}} 1_{\tilde{P}([(N/W_d)^{1/d}])}(z) \quad \text{and} \quad \nu^* := \tilde{\nu} - \nu,$$

so that

$$\mathbb{E}_{z \in [(N/W_d)^{1/d}]} f_1(x + \tilde{P}(z), y) f_2(x, y + \tilde{P}(z)) = \mathbb{E}_{z \in [N]} f_1(x + z, y) f_2(x, y + z) \tilde{\nu}(z)$$

and

$$\mathcal{D}_0^*(f_1, f_2)(x, y) = \mathbb{E}_{z \in [N]} f_1(x + z, y) f_2(x, y + z) \nu^*(z).$$

Then

$$\mathbb{E}_{j_1, j_2, r \in [V]} \sum_{x, y} \mathbb{E}_{z \in [N/V]} b_{1, j_1, j_2}(x) b_{2, j_1, j_2}(y) f_{1, j_1+r, j_2}(x+z, y) f_{2, j_1, j_2+r}(x, y+z) \nu_r^*(z) \geq \exp(-\log(1/\delta)^{O(1)})(N/V)^2,$$

where $f_{l, j_1, j_2} := (f_l)_{(j_1, j_2, V)}$ and $\nu_r^* := \nu_{(r, V)}^*$.

We will now deduce the largeness of the norm $\|\widehat{\nu^*}\|_\infty$, in contradiction with the Fourier-uniformity estimate from Proposition 6.58. By pigeonholing in $j_1, j_2, r \in [V]$, we can find functions

$$f_l'(x, y) := \begin{cases} b_{2, j_1, j_2}(y) f_{1, j_1+r, j_2}(x, y), & l = 1 \\ b_{1, j_1, j_2}(x) f_{2, j_1+r, j_2}(x, y), & l = 2 \end{cases}$$

for which

$$\left| \sum_{x, y} \mathbb{E}_{z \in [N/V]} f_1'(x, y) f_2'(x-z, y+z) \nu_r^*(z) \right| \geq \exp(-\log(1/\delta)^{O(1)})(N/V)^2.$$

This new expression is similar to one in the proof of Roth's theorem and can be treated accordingly. Indeed, introducing a new shifting parameter h along the vector $(-1, 1)$ and pigeonholing in $x, y \in [N/V]^2$, we see that

$$\left| \sum_h \mathbb{E}_{z \in [N/V]} f_{1, x, y}'(h) f_{2, x, y}'(h+z) \nu_r^*(z) \right| \geq \exp(-\log(1/\delta)^{O(1)})(N/V),$$

where $f_{l, x, y}'(h) = f_l'(x-h, y+h)$. Fourier expanding all three functions, extracting $\|\widehat{\nu_r^*}\|_\infty$, and applying the Plancherel identity together with the bounds on $\|f_{l, x, y}'\|_2 \ll (N/V)^{1/2}$, we deduce that

$$\|\widehat{\nu_r^*}\|_\infty \geq \exp(-\log(1/\delta)^{O(1)})(N/V).$$

But Proposition 6.58 tells us that

$$\|\widehat{\nu_r^*}\|_\infty \ll N/(Vw^c)$$

for some constant $c > 0$ depending only on P . Hence $w \leq \exp(\log(1/\delta)^{O(1)})$ and so $W \leq \exp \exp(\log(1/\delta)^{O(1)})$, which contradicts the assumption on W as long as the $O(1)$ term is sufficiently large.

□

We are finally ready to conclude the proof of Theorem 6.1.

Proof of Theorem 6.1. It suffices to prove the result with N replaced by WN . Let $A \subseteq [WN]^2$ be a set with density γ . By the pigeonhole principle, there exist $j_1, j_2 \in [\beta_1^2 W]$ such that the set

$$B := \{(x, y) \in [N]^2 : \beta_1^2 W(x, y) + (j_1, j_2) \in A\}$$

has size at least γN^2 . Since $\nu(z) \geq d^{-1}$ pointwise for $z \in [N]$, the supersaturation result (Proposition 6.60) tells us that

$$\Lambda^{\text{Model}}(1_B, 1_B, 1_B) \geq d^{-1} \exp(-\exp(\gamma^{-c'})) N^2$$

for some absolute $c' > 0$. Setting $\delta := \exp(-\exp(\gamma^{-c'}))/(2d)$, we obtain from Proposition 6.20 that

$$\Lambda^W(1_B, 1_B, 1_B) \gg_P \exp(-\exp(\gamma^{-c'})) N^2$$

for any $W \geq \exp \exp(\log(1/\delta)^{\Omega_P(1)})$ satisfying

$$N \gg_P W^{\Omega_P(1)} \geq \exp \exp(\log(1/\delta)^{\Omega_P(1)}) \geq \exp \exp \exp(\Omega_P(\gamma^{-c'})).$$

In particular, the set B contains nontrivial configurations (x, y) , $(x + \tilde{P}(z), y)$, $(x, y + \tilde{P}(z))$, and hence the set A contains nontrivial configurations of (x, y) , $(x + P(z), y)$, $(x, y + P(z))$. This completes the proof. $\square$

We will now deduce Corollary 6.3.

Proof of Corollary 6.3. It suffices to prove the result with N replaced by WN . Let $A \subseteq [WN]$ be a set of density γ . By the pigeonhole principle, there exists $j \in [\beta_1^2 W]$ such that the set

$$B := \{x \in [N] : \beta_1^2 Wx + j \in A\}$$

has size at least γN . Since $\nu(z) \geq d^{-1}$ pointwise for $z \in [N]$, a standard Varnavides-style supersaturation version of the result of Kelley and Meka [61] tells us that

$$\sum_x \mathbb{E}_{z \in [N]} 1_B(x) 1_B(x+z) 1_B(x+2z) \geq d^{-1} \exp(-\log(1/\gamma)^{C'}) N$$

for some absolute constant $C' > 0$.

Consider the set

$$B' := \{(x, y) \in [2N]^2 : y - x \in B\},$$

and note that

$$N \sum_x 1_B(x) 1_B(x+z) 1_B(x+2z) \asymp \sum_{x,y} 1_{B'}(x,y) 1_{B'}(x+z,y) 1_{B'}(x,y+z)$$

for each $z \in [N]$. Using this correspondence and applying Proposition 6.20 with $\delta := \exp(-\log(1/\gamma)^{C'})/(10d)$, we conclude that

$$\sum_x \mathbb{E}_{z \in [(N/W_d)^{1/d}]} 1_B(x) 1_B(x + \tilde{P}(z)) 1_B(x + 2\tilde{P}(z)) \gg_P \exp(-\log(1/\gamma)^{O_P(1)}) N$$

as long as $W \geq \exp \exp(\log(1/\delta)^{\Omega_P(1)})$ satisfies

$$N \gg_P W^{\Omega_P(1)} \geq \exp \exp(\log(1/\delta)^{\Omega_P(1)}) \geq \exp \exp(\log(1/\gamma)^{\Omega_P(1)}).$$

In particular, the set B contains nontrivial configurations x , $x + \tilde{P}(z)$, $x + 2\tilde{P}(z)$, and hence the set A contains nontrivial configurations x , $x + P(z)$, $x + 2P(z)$. This completes the proof. $\square$

6.8 Inverse theorems for box norms

We require inverse theorems for several box norms and directional Gowers norms that arise in our arguments. Although these results are more or less standard, we include statements and proofs here for ease of reference.

6.8.1 U^1 -inverse theorems

Heuristically, a function has large degree-1 Gowers norm in the direction $\mathbf{v}$ if it correlates with a function that is periodic in the $\mathbf{v}$ -direction. Since we are working with norms in which the differencing parameters h are constrained to lie in intervals, making this heuristic precise is somewhat delicate. In particular, we need the “length” of the differencing parameter to be sufficiently large relative to the size of the support of the function in question, where “sufficiently large” depends on the quality of the lower bound on the Gowers norm. This issue does not appear for higher-degree box norms since Lemma 6.7(4) allows us to manipulate the lengths of boxes at will.

We start with the U^1 -inverse theorem for functions on $\mathbb{Z}$.

Lemma 6.21 (1-dimensional U^1 -inverse theorem). *Let $\delta \in (0, 1]$ and $N, N' \in \mathbb{N}$ with $N' \geq \delta^{-1/2} N \geq \delta^{-1}$.*

If $f : \mathbb{Z} \rightarrow \mathbb{C}$ is a 1-bounded function supported on $[N]$, then

$$\|f\|_{[\pm N']}^2 \geq \delta N \implies \left| \sum_x f(x) \right| \gg \delta^{1/4} N.$$

Proof. By definition, we have

$$\|f\|_{[\pm N']}^2 = \sum_{x, x' \in [N]} \mu_{N'}(x' - x) f(x) \overline{f(x')}.$$

The idea is that since N' is sufficiently large relative to N , we have $\mu_{N'}(x - x') \approx 1/(2N' - 1)$ for all $x, x' \in [N]$. More precisely, since f is 1-bounded and $N' \geq \delta^{-1/2} N \geq N$, we can make the crude comparison

$$\begin{aligned} \left| (2N' - 1) \|f\|_{[N']}^2 - \sum_{x, x'} f(x) \overline{f(x')} \right| &\leq \sum_{x, x' \in [N]} |(2N' - 1) \mu_{N'}(x' - x) - 1| \\ &= \frac{1}{2N' - 1} \sum_{x, x' \in [N]} |x - x'| \\ &\leq \frac{N^2(N - 1)}{2N' - 1} \leq \delta^{1/2} N^2/2. \end{aligned}$$

Hence

$$\begin{aligned} \left| \sum_x f(x) \right|^2 &= \sum_{x, x'} f(x) \overline{f(x')} \geq (2N' - 1) \delta N - \delta^{1/2} N^2/2 \\ &\geq (3/2) \delta^{1/2} N - \delta N \gg \delta^{1/2} N^2 \end{aligned}$$

since $N \geq \delta^{-1/2}$. □

Lemma 6.22 (2-dimensional U^1 -inverse theorem). *Let $\delta \in (0, 1]$ and $N, N' \in \mathbb{N}$ with $N' \geq (\delta/2)^{-1/2} N \geq 2\delta^{-1}$, and let $f : \mathbb{Z}^2 \rightarrow \mathbb{C}$ be a 1-bounded function supported on $[N]^2$. Then*

$$\|f\|_{\mathbf{e}_1, [\pm N']}^2 \geq \delta N^2 \implies \sum_{x, y} f(x, y) b(y) \gg \delta^{5/4} N^2 \text{ for some 1-bounded } b : \mathbb{Z} \rightarrow \mathbb{C};$$

$$\|f\|_{\mathbf{e}_2, [\pm N']}^2 \geq \delta N^2 \implies \sum_{x, y} f(x, y) b(x) \gg \delta^{5/4} N^2 \text{ for some 1-bounded } b : \mathbb{Z} \rightarrow \mathbb{C};$$

and

$$\|f\|_{(\mathbf{e}_2 - \mathbf{e}_1), [\pm N']}^2 \geq \delta N^2 \implies \sum_{x, y} f(x, y) b(x + y) \gg \delta^{5/4} N^2 \text{ for some 1-bounded } b : \mathbb{Z} \rightarrow \mathbb{C}.$$

Proof. We begin with the first statement; the second follows by symmetry. For each $y \in \mathbb{Z}$, define the

function $f_y : \mathbb{Z} \rightarrow \mathbb{Z}$ via $f_y(x) := f(x, y)$. Then we have the identity

$$\|f\|_{\mathbf{e}_1 \cdot [\pm N']}^2 = \sum_{y \in [N]} \|f_y\|_{[\pm N']}^2.$$

The popularity principle supplies at least $\delta N/2$ values of $y \in [N]$ for which $\|f_y\|_{[\pm N']}^2 \geq \delta N/2$. Applying Lemma 6.21 for each such y and using nonnegativity for the remaining y 's, we find that

$$\sum_{y \in [N]} \left| \sum_{x \in \mathbb{Z}} f(x, y) \right| \gg \delta N \cdot \delta^{1/4} N.$$

Adding 1-bounded fudge factors $b(y)$ to remove the absolute values, we obtain

$$\sum_{x, y} f(x, y) b(y) \gg \delta^{5/4} N^2,$$

as desired.

The proof of the third statement goes similarly. This time, we foliate $[N]^2$ using 1-dimensional arithmetic progressions with common difference $\mathbf{e}_2 - \mathbf{e}_1$; each such arithmetic progression intersects $[N]^2$ in at most N elements. Applying the popularity principle and Lemma 6.21, and then removing the absolute values as in the first statement, we arrive at the desired conclusion. $\square$

6.8.2 Unnormalized box norms

To prove inverse theorems for higher-degree box norms, it is often easier to start with unnormalized box norms. Given $\mathbf{v}_1, \dots, \mathbf{v}_s \in \mathbb{Z}^D$, we define the *unnormalized box norm* of a compactly supported function $f : \mathbb{Z}^D \rightarrow \mathbb{C}$ along $\mathbf{v}_1, \dots, \mathbf{v}_s$ to be

$$\|f\|_{\mathbf{v}_1, \dots, \mathbf{v}_s} := \left(\sum_{\mathbf{x}, h_1, \dots, h_s} \Delta_{\mathbf{v}_1 h_1, \dots, \mathbf{v}_s h_s} f(\mathbf{x}) \right)^{1/2^s}.$$

If $D = 1$ and $\mathbf{v}_1 = \dots = \mathbf{v}_s = 1$, we also set $\|f\|_{U^s} := \|f\|_{\mathbf{v}_1, \dots, \mathbf{v}_s}$. The following lemma relates unnormalized and normalized box norms.

Lemma 6.23 (Relating normalized and unnormalized box norms). *Let $s \geq 2$, let $D, H_1, \dots, H_s \in \mathbb{N}$, and let $\mathbf{v}_1, \dots, \mathbf{v}_s \in \mathbb{Z}^D$. For any compactly supported function $f : \mathbb{Z}^D \rightarrow \mathbb{C}$, we have*

$$\|f\|_{\mathbf{v}_1 \cdot [\pm H_1], \dots, \mathbf{v}_s \cdot [\pm H_s]}^{2^s} \leq \prod_{i=1}^s \frac{1}{2H_i - 1} \cdot \|f\|_{\mathbf{v}_1, \dots, \mathbf{v}_s}^{2^s}.$$

Proof. Using the inductive formula for box norms, the pointwise bound $\mu_H(h) \leq 1/(2H-1)$, and nonnegativity, we have

$$\begin{aligned}
\|f\|_{\mathbf{v}_1 \cdot [\pm H_1], \dots, \mathbf{v}_s \cdot [\pm H_s]}^{2^s} &= \sum_{h_2, \dots, h_s} \mu_{H_2}(h_2) \cdots \mu_{H_s}(h_s) \|\Delta_{\mathbf{v}_2 h_2, \dots, \mathbf{v}_s h_s} f\|_{\mathbf{v}_1 \cdot [\pm H_1]}^2 \\
&\leq \prod_{i=2}^s \frac{1}{2H_i - 1} \cdot \sum_{h_2, \dots, h_s} \|\Delta_{\mathbf{v}_2 h_2, \dots, \mathbf{v}_s h_s} f\|_{\mathbf{v}_1 \cdot [\pm H_1]}^2 \\
&= \prod_{i=2}^s \frac{1}{2H_i - 1} \cdot \sum_{h_1} \mu_{H_1}(h_1) \|\Delta_{\mathbf{v}_1 h_1} f\|_{\mathbf{v}_2, \dots, \mathbf{v}_s}^{2^{s-1}} \\
&\leq \prod_{i=1}^s \frac{1}{2H_i - 1} \cdot \|f\|_{\mathbf{v}_1, \dots, \mathbf{v}_s}^{2^s};
\end{aligned}$$

note that the applications of nonnegativity crucially used the assumption $s \geq 2$. $\square$

6.8.3 U^2 -inverse theorems

The starting point for U^2 -inverse theorems is the following well-known 1-dimensional result (see, e.g., [81, Lemma 2.4]).

Lemma 6.24 (1-dimensional U^2 -inverse theorem). *Let $\delta \in (0, 1]$ and $H, N \in \mathbb{N}$ with $\delta N \leq H \leq N$. If $f : \mathbb{Z} \rightarrow \mathbb{C}$ is a 1-bounded function supported on $[N]$, then*

$$\|f\|_{[\pm H], [\pm H]}^4 \geq \delta N \implies \sum_x f(x) e(ax + b) \gg \delta^{3/2} N \text{ for some } a, b \in \mathbb{R}/\mathbb{Z}.$$

Proof. By Lemma 6.23 and the assumption $H \geq \delta N$, we have $\|f\|_{U^2}^4 \gg \delta^3 N^3$. Fourier inversion and Plancherel's Identity give

$$\|f\|_{U^2}^4 = \|\hat{f}\|_4^4 \leq \|\hat{f}\|_2^2 \cdot \|\hat{f}\|_\infty^2 = \|f\|_2^2 \cdot \|\hat{f}\|_\infty^2 \leq N \cdot \|\hat{f}\|_\infty^2.$$

Hence $\|\hat{f}\|_\infty \gg \delta^{3/2} N$, which is equivalent to the conclusion of the lemma. $\square$

Lemma 6.25 (2-dimensional U^2 -inverse theorem). *Let $\delta \in (0, 1]$ and $H, N \in \mathbb{N}$ with $\delta^{-O(1)} \ll \delta N \leq H \leq N$, and let $f : \mathbb{Z}^2 \rightarrow \mathbb{C}$ be a 1-bounded function supported on $[N]^2$. Then*

$$\|f\|_{(\mathbf{e}_1 \cdot [\pm H])^2}^4 \geq \delta N^2 \implies \sum_{x, y} f(x, y) e(a(y)x + b(y)) \gg \delta^{3/2} N^2$$

for some 1-bounded $a, b : \mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}$;

$$\|f\|_{(\mathbf{e}_2 \cdot [\pm H])^2}^4 \geq \delta N^2 \implies \sum_{x,y} f(x,y) e(a(x)y + b(x)) \gg \delta^{3/2} N^2$$

for some 1-bounded $a, b : \mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}$.

This lemma follows from Lemma 6.24 just as Lemma 6.22 follows from Lemma 6.21.

6.8.4 $U^1 \times U^1$ -inverse theorems

Let $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{Z}^2$ be vectors such that $\langle \mathbf{v}_1, \mathbf{v}_2 \rangle_{\mathbb{Z}} = \mathbb{Z}^2$. Then the unnormalized box norm along $\mathbf{v}_1, \mathbf{v}_2$ can be written as

$$\|f\|_{\mathbf{v}_1, \mathbf{v}_2} = \left(\sum_{m, m', n, n'} f(\mathbf{v}_1 m + \mathbf{v}_2 n) \overline{f(\mathbf{v}_1 m' + \mathbf{v}_2 n)} f(\mathbf{v}_1 m + \mathbf{v}_2 n') \overline{f(\mathbf{v}_1 m' + \mathbf{v}_2 n')} \right)^{1/4}$$

and satisfies the following inverse theorem.

Lemma 6.26 (Inverse theorem for the unnormalized $U^1 \times U^1$ -norm). *Let $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{Z}^2$ be vectors of length $O(1)$ satisfying $\langle \mathbf{v}_1, \mathbf{v}_2 \rangle_{\mathbb{Z}} = \mathbb{Z}^2$, and let $N \in \mathbb{N}$. If $f : \mathbb{Z}^2 \rightarrow \mathbb{C}$ is a 1-bounded function supported on $[N]^2$, then there are 1-bounded functions $b_1, b_2 : \mathbb{Z} \rightarrow \mathbb{C}$ such that*

$$\|f\|_{\mathbf{v}_1, \mathbf{v}_2}^4 \ll N^2 \cdot \sum_{m, n} f(\mathbf{v}_1 m + \mathbf{v}_2 n) b_1(m) b_2(n).$$

Proof. The pigeonhole principle gives some $n' \in \mathbb{Z}$, necessarily of size $O(N)$, such that

$$\|f\|_{\mathbf{v}_1, \mathbf{v}_2}^4 \ll N \left| \sum_{m, m', n} f(\mathbf{v}_1 m + \mathbf{v}_2 n) \overline{f(\mathbf{v}_1 m' + \mathbf{v}_2 n)} f(\mathbf{v}_1 m + \mathbf{v}_2 n') \overline{f(\mathbf{v}_1 m' + \mathbf{v}_2 n')} \right|.$$

Let $S := \{m' \in \mathbb{Z} : \mathbf{v}_1 m' + \mathbf{v}_2 n' \in [N]^2\}$. Then $|S| \ll N$, and the result follows upon setting $b_1(m) := \overline{f(\mathbf{v}_1 m + \mathbf{v}_2 n')}$ and letting $b_2(n)$ be equal to $\mathbb{E}_{m' \in S} \overline{f(\mathbf{v}_1 m' + \mathbf{v}_2 n)} f(\mathbf{v}_1 m' + \mathbf{v}_2 n')$ times a suitable complex number of norm 1 that makes the expression in the absolute value positive. $\square$

Combining this lemma with Lemma 6.23 immediately gives the following result for unnormalized box norms.

Lemma 6.27 (Inverse theorem for the $U^1 \times U^1$ -norm). *Let $\delta \in (0, 1]$ and $H, N \in \mathbb{N}$ with $\delta N \leq H \leq N$. Let $f : \mathbb{Z}^2 \rightarrow \mathbb{C}$ be a 1-bounded function supported on $[N]^2$. Then*

$$\|f\|_{\mathbf{e}_1 \cdot [\pm H], \mathbf{e}_2 \cdot [\pm H]}^4 \geq \delta N^2 \implies \sum_{x,y} f(x,y) b_1(x) b_2(y) \gg \delta^3 N^2$$

for some 1-bounded $b_1, b_2 : \mathbb{Z} \rightarrow \mathbb{C}$;

$$\|f\|_{\mathbf{e}_1 \cdot [\pm H], (\mathbf{e}_2 - \mathbf{e}_1) \cdot [\pm H]}^4 \geq \delta N^2 \implies \sum_{x,y} f(x,y) b_1(x+y) b_2(y) \gg \delta^3 N^2$$

for some 1-bounded $b_1, b_2 : \mathbb{Z} \rightarrow \mathbb{C}$;

$$\|f\|_{\mathbf{e}_2 \cdot [\pm H], (\mathbf{e}_2 - \mathbf{e}_1) \cdot [\pm H]}^4 \geq \delta N^2 \implies \sum_{x,y} f(x,y) b_1(x) b_2(x+y) \gg \delta^3 N^2$$

for some 1-bounded $b_1, b_2 : \mathbb{Z} \rightarrow \mathbb{C}$.

6.8.5 $U^2 \times U^1$ -inverse theorem

Recall that our degree-lowering result Theorem 6.10 is based on replacing “ $U^2 \times U^1$ -obstructions” with “ $U^1 \times U^1$ -obstructions”. We are finally ready to give a precise characterization of the former obstructions.

Lemma 6.28. *Let $\delta \in (0, 1]$ and $H, N \in \mathbb{N}$ with $\delta^{O(1)} \ll \delta N \leq H \leq N$, and let $f : \mathbb{Z}^2 \rightarrow \mathbb{C}$ be a 1-bounded function supported on $[N]^2$. Then*

$$\|f\|_{\mathbf{e}_1 \cdot [\pm H], \mathbf{e}_1 \cdot [\pm H], \mathbf{e}_2 \cdot [\pm H]}^8 \geq \delta N^2 \implies \sum_{x,y} f(x,y) g(x) e(a(y)x + b(y)) \gg \delta^{O(1)} N^2$$

for some 1-bounded $g : \mathbb{Z} \rightarrow \mathbb{C}$ and phase functions $a, b : \mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}$;

$$\|f\|_{\mathbf{e}_1 \cdot [\pm H], \mathbf{e}_2 \cdot [\pm H], \mathbf{e}_2 \cdot [\pm H]}^8 \geq \delta N^2 \implies \sum_{x,y} f(x,y) g(y) e(a(x)y + b(x)) \gg \delta^{O(1)} N^2$$

for some 1-bounded $g : \mathbb{Z} \rightarrow \mathbb{C}$ and phase functions $a, b : \mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}$.

Proof. We only prove the first statement as the other can be proved by swapping the roles of x and y . The inductive formula for box norms gives

$$\delta^8 N^2 \leq \|f\|_{\mathbf{e}_1 \cdot [\pm H], \mathbf{e}_1 \cdot [\pm H], \mathbf{e}_2 \cdot [\pm H]}^8 = \sum_h \mu_H(h) \|\Delta_{h\mathbf{e}_2} f\|_{\mathbf{e}_1 \cdot [\pm H], \mathbf{e}_1 \cdot [\pm H]}^4.$$

By Lemma 6.25, there are phases $\alpha(y, y+h), \beta(y, y+h) \in \mathbb{R}/\mathbb{Z}$ such that

$$\sum_{x,y,h} \mu_H(h) \Delta_{h\mathbf{e}_2} f(x,y) e(\alpha(y, y+h)x + \beta(y, y+h)) \gg \delta^{O(1)} N^2$$

and the summand is real and nonnegative for each fixed h, y . In particular, we have

$$\sum_{x,y,h} \Delta_{he_2} f(x,y) e(\alpha(y,y+h)x + \beta(y,y+h)) \gg \delta^{O(1)} N^3.$$

Replacing the variable h with the variable $z := y + h$, we get

$$\sum_{x,y,z} f(x,y) \overline{f(x,z)} e(\alpha(y,z)x + \beta(y,z)) \gg \delta^{O(1)} N^3.$$

The pigeonhole principle gives an integer $|z_0| \ll N$ such that

$$\sum_{x,y} f(x,y) \overline{f(x,z_0)} e(\alpha(y,z_0)x + \beta(y,z_0)) \gg \delta^{O(1)} N^2.$$

The lemma now follows upon setting

$$g(x) := \overline{f(x,z_0)}, \quad a(y) := \alpha(y,z_0), \quad \text{and} \quad b(y) := \beta(y,z_0).$$

□

6.9 Nilsequence definitions

This section contains (mostly standard) definitions pertaining to nilsequences.

Definition 6.29 ($(t-1)$ -fold commutator). For elements $g_1, \dots, g_t$ of a group G , we define

$$[g_1, g_2, \dots, g_t] := [[g_1, g_2], g_3], \dots, g_t.$$

We will call any such commutator a $(t-1)$ -fold commutator of the elements $g_1, \dots, g_t$.

We now introduce the lower central series.

Definition 6.30. The *lower central series* of a nilpotent Lie group G is the sequence of nested subgroups $G = G_{(0)} = G_{(1)} \supseteq G_{(2)} \supseteq G_{(3)} \cdots$ where $G_{(i)} := [G, G_{(i-1)}]$ for $i \geq 1$. The *step* of G is the smallest integer s such that $G_{(s+1)} = \text{Id}_G$.

Recall that the *height* of the rational number a/b (for a, b coprime integers) is $\max(|a|, |b|)$. We will now define a nilmanifold.

Definition 6.31 (Nilmanifold). A *nilmanifold* of degree k , step s , complexity at most M , and dimension m consists of the following data:

1. an s -step connected and simply connected nilpotent real Lie group G and a discrete cocompact subgroup Γ of G ;
2. a *filtration* $G_\bullet = (G_i)_{i=0}^\infty$ of nested subgroups $G = G_0 = G_1 \supseteq G_2 \supseteq \cdots$ such that $G_i = \{\text{Id}_G\}$ for all $i > k$ and $[G_i, G_j] \subseteq G_{i+j}$ for all i, j ; and
3. a *Mal'cev basis*, i.e., a basis $\{X_1, \dots, X_m\}$ of the Lie algebra $\mathfrak{g} := \log(G)$.

We furthermore require the data to satisfy the following conditions:

1. For $1 \leq i, j \leq m$, we have

$$[X_i, X_j] = \sum_{l > \max(i, j)} a_{ijl} X_l$$

for some rational numbers a_{ijl} of height at most M .

2. For $m_i := \dim(G_i)$, we have $\log(G_i) = \text{span}\{X_j : m - m_i < j \leq m\}$.
3. Each element $g \in G$ can be expressed uniquely as $\exp(t_1 X_1) \exp(t_2 X_2) \cdots \exp(t_m X_m)$ with $t_1, \dots, t_m \in \mathbb{R}$; we let $\psi : g \mapsto (t_1, \dots, t_m)$ denote the associated *coordinate map* from G to $\mathbb{R}^m$.
4. Γ consists of precisely the elements g for which $\psi(g) \in \mathbb{Z}^m$.

A Mal'cev basis that satisfies the last three of these conditions is said to be *adapted* to the filtration $G_\bullet$.

Note that condition (ii) implies that each G_i is closed, connected, and simply connected.

The lower central series (as defined above) gives the so-called *standard filtration* on a nilpotent Lie group G , and we emphasize that there are filtrations other than the standard filtration. The standard filtration on G is “minimal” in the sense that any other filtration (G_i) satisfies $G_i \supseteq G_{(i)}$ for each $i \in \mathbb{N}$ [107, Exercise 1.6.2].

A nilmanifold can be endowed with a metric as follows.

Definition 6.32 (Metric on a nilmanifold). Given a nilmanifold G/Γ with coordinate map ψ , we define a metric d_G on G to be the right-invariant metric

$$d_G(x, y) := \inf \left\{ \sum_{i=1}^n \min\{|\psi(x_i x_{i-1}^{-1})|, |\psi(x_{i-1} x_i^{-1})|\} : x_0 = x, x_n = y \right\}.$$

Setting $d(x\Gamma, y\Gamma) := \inf\{d(x', y') : x'\Gamma = x\Gamma, y'\Gamma = y\Gamma\}$, we obtain a metric on G/Γ .

This metric makes G/Γ into a compact metric space and hence can be used to define the Lipschitz parameter of a function on G/Γ .

Definition 6.33 (Lipschitz norm). Let (X, d_X) be a metric space. The *Lipschitz constant* of a function $F: X \rightarrow \mathbb{C}$ is the quantity

$$\sup_{x \neq y \in X} \frac{|F(x) - F(y)|}{d_X(x, y)},$$

and the *Lipschitz norm* of F is

$$\|F\|_{\text{Lip}(X)} := \|F\|_{L^\infty(X)} + \sup_{x \neq y \in X} \frac{|F(x) - F(y)|}{d_X(x, y)}.$$

If $\|F\|_{\text{Lip}(X)} \leq M$, then we say that F is an *M-Lipschitz function*.

An important property of Lipschitz functions on a nilmanifold is that they can be approximated by finite linear combinations of vertical characters (as defined below).

We now turn to polynomial sequences.

Definition 6.34 (Polynomial sequences). Let G be an s -step connected and simply connected nilpotent real Lie group with a degree- k filtration $(G_i)_{i=0}^\infty$. The set $\text{poly}(\mathbb{Z}, G)$ of *polynomial sequences* to G consists of the maps $g: \mathbb{Z} \rightarrow G$ of the form

$$g(n) = \prod_{i=0}^k g_i^{(n)} \quad (n \in \mathbb{Z})$$

for some $g_i \in G_i$. It is known (see, e.g., [58, Chapter 14]) that $\text{poly}(\mathbb{Z}, G)$ forms a group under pointwise multiplication.

Our next definition is a smoothness norm on real polynomials that captures the extent to which their coefficients are close to integers.

Definition 6.35 (Smoothness norm). Let $p(x) = \sum_{j=0}^D \alpha_j x^j \in \mathbb{R}[x]$ be a polynomial of degree D . Its *smoothness norm* at scale $T > 0$ is defined to be

$$\|p\|_{C^\infty[T]} := \max_{1 \leq j \leq D} T^j \|\alpha_j\|_{\mathbb{R}/\mathbb{Z}}.$$

Note that this definition of smoothness is written using the monomial basis, which differs slightly from [73, Definition 2.2]. By [73, Lemma A.10], these notions are equivalent up to a multiplicative factor of $O_D(1)$ if one is willing to replace p by qp for some natural number $q = O_D(1)$.

Two important special classes of polynomial sequences on G are those whose coefficients are “rational” and those which grow slowly and remain “close” to the identity.

Definition 6.36 (Rational and smooth sequences). Given $Q \in \mathbb{N}$, we say an element g of G is *rational with height Q* (or *Q -rational*) if $g^Q \in \Gamma$. We say that a polynomial sequence γ on G is *Q -rational* if $\gamma(n)$ is Q -rational for each $n \in \mathbb{Z}$.

We say that a polynomial sequence ε on G is (K, T) -*smooth* if for each $n \in \mathbb{Z}$, we have

$$d_G(\varepsilon(n+1), \varepsilon(n)) \leq K/T \quad \text{and} \quad d_G(\varepsilon(0), \text{id}_G) \leq K.$$

We are finally ready to define a nilsequence.

Definition 6.37 (Nilsequence). A *nilsequence of complexity M* is a sequence $\varphi : \mathbb{Z} \rightarrow \mathbb{C}$ of the form $\varphi(n) = F(g(n)\Gamma)$, where g is a polynomial sequence on G and $F : G/\Gamma \rightarrow \mathbb{C}$ is an M -Lipschitz function on a nilmanifold G/Γ of complexity M .

We now shift gears and discuss horizontal and vertical characters on G/Γ .

Definition 6.38 (Horizontal and vertical components). The *horizontal component* of G is $G/G_{(2)} = G/[G, G]$, and the *vertical component* of G is $G_{(s)}$.⁷ We let $\pi_{\text{horiz}} : G \rightarrow G/G_{(2)}$ denote the quotient map. We use d_{horiz} and $d_{(s)}$ to denote the dimensions of $G/G_{(2)}$ and $G_{(s)}$, respectively.

We remark that $G/[G, G]\Gamma$, often called the *horizontal torus*, is isomorphic to the torus $(\mathbb{R}/\mathbb{Z})^{d_{\text{horiz}}}$; the Mal'cev basis adapted to G/Γ provides an isomorphism. Horizontal characters come from characters of the horizontal torus, as follows.

Definition 6.39 (Horizontal character). A *horizontal character* on G/Γ is a continuous homomorphism $\eta : G \rightarrow \mathbb{R}$ such that $\eta(\Gamma) \subseteq \mathbb{Z}$.

Note that since $\mathbb{R}$ is abelian, every horizontal character η annihilates $[G, G]$ and hence descends to a character on the horizontal torus. Since the characters of $(\mathbb{R}/\mathbb{Z})^{d_{\text{horiz}}}$ correspond to elements of $\mathbb{Z}^{d_{\text{horiz}}}$, our choice of Mal'cev basis provides a correspondence between horizontal characters of G/Γ and elements of $\mathbb{Z}^{d_{\text{horiz}}}$. More explicitly, each η is of the form $\eta(g) = k \cdot \psi(g)$ for some $k \in \mathbb{Z}^{d_{\text{horiz}}}$.

We will also require the notion of a vertical character.

Definition 6.40 (Vertical character). Let H be a connected, simply connected subgroup of the center $Z(G)$ which is *rational* in the sense that $H \cap \Gamma$ is cocompact in H , and let $\eta : H \rightarrow \mathbb{R}$ be a continuous homomorphism such that $\eta(H \cap \Gamma) \subseteq \mathbb{Z}$. We say that a function $F : G/\Gamma \rightarrow \mathbb{C}$ is an *H -vertical character* with *frequency η* if

$$F(gx) = e(\eta(g))F(x)$$

⁷In [51], the vertical component is defined to be G_ℓ where G_ℓ is the smallest nontrivial element of the filtration of G .

for all $g \in H$ and $x \in G/\Gamma$. We refer to a $G_{(s)}$ -vertical character as simply a *vertical character*.

Horizontal and vertical characters come with natural notions of size.

Definition 6.41 (Size of a character). If η is a horizontal character on G/Γ associated with the element $k \in (\mathbb{R}/\mathbb{Z})^{d_{\text{horiz}}}$, then the *size* (or *modulus*) of η , denoted $|k|$ or $\|k\|_\infty$, is the L^∞ norm of the vector k .

Let H be a rational subgroup of the center $Z(G)$, and let $\xi: H/(H \cap \Gamma) \rightarrow \mathbb{R}/\mathbb{Z}$ be a homomorphism. The *size* (or *modulus*) of ξ , denoted $|\xi|$, is the Lipschitz constant of the function $x \mapsto e(\xi(x))$ on $H/(H \cap \Gamma)$. Equivalently,

$$|\xi| := \sup_{x \neq y \in H/(H \cap \Gamma)} \frac{|e(\xi(x)) - e(\xi(y))|}{d_{G/\Gamma}(x, y)}.$$

Finally, we will need quantitative measures of rationality for subspaces and subgroups.

Definition 6.42 (Rational subspace). Let V be a vector space with a basis $\mathcal{B}$, and let W be a subspace of V . We say that W is *at most Q -rational* if W has a basis $\mathcal{B}'$ such that each $w_j \in \mathcal{B}'$ can be written as

$$w_j = \sum_{v_i \in \mathcal{B}} a_{ji} v_i$$

for some rationals a_{ji} of height at most Q .

Definition 6.43 (Rational subgroup and subnilmanifold). Let G/Γ be a nilmanifold with Mal'cev basis $\mathcal{X}$. A connected and simply connected subgroup $H \subseteq G$ is *Q -rational* if $\mathfrak{h} := \log(H)$ is Q -rational with respect to $\mathfrak{g}$ with the basis $\mathcal{X}$. Define $\Gamma' := H \cap \Gamma$. We then say that the nilmanifold H/Γ' , equipped with the filtration $H_\bullet = (H_i)_{i=0}^\infty$ given by $H_i := G_i \cap H$, is a *Q -rational subnilmanifold* of G/Γ .

Remark 6.44. If H is a Q -rational subgroup of G , then by [74, Lemma B.12] there exists a $Q^{O_k(m^{O(1)})}$ -rational Mal'cev basis for $H/(\Gamma \cap H)$.

6.10 Fraction-comparison estimates

In this section, we prove elementary “fraction-comparison estimates” in the style of [88] (see, e.g., the proof of Proposition 5.6 there). These estimates are useful in the nilsequences part of our argument. While the first two lemmas are straightforward, the third makes use of some important properties of V -tricked polynomials. Recall the smoothness norms from Definition 6.35.

Lemma 6.45. *Let $D, Q, T \in \mathbb{N}$, let $R \in \mathbb{R}[z]$ be a polynomial of degree D , and let $a, b \in \mathbb{Z}$ with $|a| \leq QT$ and $0 < |b| \leq Q$ nonzero. Then the polynomial $\tilde{R}(z) := R(a + bz)$ satisfies*

$$\|b^D R\|_{C^\infty[T]} \ll_D Q^{D-1} \|\tilde{R}\|_{C^\infty[T]}.$$

Proof. Assume that $\|\tilde{R}\|_{C^\infty[T]} \leq M$; we will show that $\|b^D R\|_{C^\infty[T]} \ll_D Q^{D-1} M$. Write $R(z) = \sum_{i=0}^D \alpha_i z^i$. Then

$$\tilde{R}(z) = R(a + bz) = \sum_{i=0}^D \alpha_i (a + bz)^i = \sum_{i=0}^D \alpha_i \sum_{j=0}^i \binom{i}{j} b^j z^j a^{i-j} = \sum_{j=0}^D b^j z^j \sum_{i=j}^D \alpha_i \binom{i}{j} a^{i-j}.$$

The assumption $\|\tilde{R}\|_{C^\infty[T]} \leq M$ means that

$$\left\| b^j \cdot \sum_{i=j}^D \alpha_i \binom{i}{j} a^{i-j} \right\|_{\mathbb{R}/\mathbb{Z}} \leq M/T^j \quad (6.30)$$

for each $1 \leq j \leq D$. In particular the $j = D$ instance of (6.30) gives

$$\|b^D \alpha_D\|_{\mathbb{R}/\mathbb{Z}} \leq M/T^D.$$

Next, the $j = D - 1$ instance of (6.30) gives

$$\|b^{D-1}(\alpha_{D-1} + Da\alpha_D)\|_{\mathbb{R}/\mathbb{Z}} \leq M/T^{D-1};$$

multiplying the expression in $\|\cdot\|_{\mathbb{R}/\mathbb{Z}}$ by b and using the triangle inequality (and the assumption on the sizes of a, b), we find that

$$\|b^D \alpha_{D-1}\|_{\mathbb{R}/\mathbb{Z}} \leq QM/T^{D-1} + D(QT)(M/T^D) \ll_D QM/T^{D-1}.$$

Continuing inductively in this manner, we find that

$$\|b^D \alpha_j\|_{\mathbb{R}/\mathbb{Z}} \ll_D Q^{D-j} M/T^j \leq Q^{D-1} M/T^j$$

for each $1 \leq j \leq D$, as desired. □

For ease of reference we record the following special case of this lemma.

Lemma 6.46. *Let $D, T \in \mathbb{N}$, and let $R \in \mathbb{R}[z]$ be a polynomial of degree D . If $|y_0| \leq QT$ and*

$$\|R(y_0 + \cdot)\|_{C^\infty[T]} \leq Q,$$

then

$$\|R\|_{C^\infty[T]} \ll_D Q^D.$$

Recall that we are interested in the V -tricked polynomial

$$\tilde{P}_{[r,V]}(z) := \frac{\tilde{P}(Vz + r) - \tilde{P}(r)}{V},$$

where $V = W^{O_P(1)}$ is a nonnegative integer power of W and $r \in [V]$; as mentioned previously, $\tilde{P}_{[r,V]}$ has integer coefficients, zero constant term and degree $d \geq 2$. Let us denote the coefficient of z^i in $\tilde{P}_{[r,V]}(z)$ by $V_i \in \mathbb{Z}[W, V, r, \beta_1, \dots, \beta_d]$. We record that

$$V_d = \beta_d \beta_1^{d-2} W^{d-1} V^{d-1}$$

and that V_1 is equal to 1 plus W times some integer-coefficient polynomial in $W, V, r, \beta_1, \dots, \beta_d$. Thus, if every prime factor of $\beta_1 \beta_d$ also divides W (which is the case for w larger than an absolute constant depending only on P), then V_1 is equivalent to 1 modulo every prime dividing V_d and we have

$$\gcd(V_d, V_1) = 1.$$

The following lemma tells us that precomposing a polynomial with $\tilde{P}_{[r,V]}$ cannot significantly decrease its (properly normalized) smoothness norm.

Lemma 6.47. *Let $D \in \mathbb{N}$. For every $C_1 > 0$ there exists a constant $C_2 = C_2(C_1, D, P) > 0$ such that the following holds. Let $V \leq W^{C_1}$ be a nonnegative integer power of W , let $r \in [V]$, and let $T, Q \in \mathbb{N}$. Let $R \in \mathbb{R}[z]$ be a polynomial of degree D . Assume that every prime factor of $\beta_1 \beta_d$ also divides W . If*

$$\|R \circ \tilde{P}_{[r,V]}\|_{C^\infty[(T/V_d)^{1/d}]} \leq Q$$

and $T > C_2 W^{C_2}$, then

$$\|R\|_{C^\infty[T]} \leq C_2 Q.$$

The scales of these smoothness norms come from the estimate $\tilde{P}_{[r,V]}((T/V_d)^{1/d}) \asymp T$.

Proof. We fix $C_1 > 0$ and let all the constants depend on P and D . The proof proceeds by induction on the degree D , where the base case $D = 0$ holds vacuously. For the induction step, write $R(z) = \sum_{i=0}^D \alpha_i z^i$. Recall that $\tilde{P}_{[r,V]}(z) = V_d z^d + \dots + V_1 z$. We compute

$$R(\tilde{P}_{[r,V]}(z)) = \sum_{i=0}^D \alpha_i (V_d z^d + \dots + V_1 z)^i$$

$$\begin{aligned}
&= \sum_{i=0}^D \alpha_i \sum_{j_1+\dots+j_d=i} \binom{i}{j_1, \dots, j_d} (V_1 z)^{j_1} \dots (V_d z^d)^{j_d} \\
&= \sum_{k=0}^{dD} z^k \sum_{i=0}^D \alpha_i \sum_{\underline{j} \in \mathcal{P}(i,k)} \binom{i}{j_1, \dots, j_d} V_1^{j_1} \dots V_d^{j_d},
\end{aligned}$$

where $\mathcal{P}(i, k)$ denotes the set of tuples $\underline{j} = (j_1, \dots, j_d)$ of nonnegative integers satisfying $j_1 + \dots + j_d = i$ and $j_1 + 2j_2 + \dots + dj_d = k$. The assumption $\|R \circ \tilde{P}_{[r, V]}\|_{C^\infty[(T/V_d)^{1/d}]} \leq Q$ tells us that

$$\left\| \sum_{i=0}^D \alpha_i \sum_{\underline{j} \in \mathcal{P}(i,k)} \binom{i}{j_1, \dots, j_d} V_1^{j_1} \dots V_d^{j_d} \right\|_{\mathbb{R}/\mathbb{Z}} \leq Q(V_d/T)^{k/d} \quad (6.31)$$

for each $1 \leq k \leq dD$. Let E_k denote the expression inside $\|\cdot\|_{\mathbb{R}/\mathbb{Z}}$ in (6.31).

We will show that $\|\alpha_D\|_{\mathbb{R}/\mathbb{Z}} \leq Q/T^D$. Since $E_{dD} = \alpha_D V_d^D$, Equation (6.31) for $k = dD$ tells us that

$$\|\alpha_D V_d^D\|_{\mathbb{R}/\mathbb{Z}} \leq Q(V_d/T)^D,$$

and hence there is some integer n such that

$$|\alpha_D - n/V_d^D| \leq Q/T^D. \quad (6.32)$$

We will aim to show that n is divisible by V_d^D . To accomplish this, we will find a second rational approximation for α_D using a denominator that is coprime to V_d^D . To begin, we have $E_1 = \alpha_1 V_1$. Next, $E_2 = \alpha_2 V_1^2 + \alpha_1 V_2$, and so

$$V_1 E_2 - V_2 E_1 = V_1^3 \alpha_2.$$

We continue with $E_3 = \alpha_3 V_1^3 + \alpha_2(2V_1 V_2) + \alpha_1 V_3$ and

$$V_1^2 E_3 - 2V_2(V_1 E_2 - V_2 E_1) - V_1 V_3 E_1 = V_1^5 \alpha_3.$$

Continuing inductively in this manner, we obtain a positive integer γ and polynomials $s_1, \dots, s_d \in \mathbb{Z}[V_1, \dots, V_d]$ such that

$$s_1 E_1 + \dots + s_d E_d = V_1^\gamma \alpha_D$$

holds (as an identity in the formal variables $V_1, \dots, V_d, \alpha_1, \dots, \alpha_D$); note that the s_i 's depend only on the parameters d, D . The crucial point for the induction is that E_k (for $1 \leq k \leq D$) equals $\alpha_k V_1^k$ plus a linear combination of $\alpha_{k-1}, \dots, \alpha_1$ with coefficients in $\mathbb{Z}[V_1, \dots, V_d]$.

Recalling that $V_i = O(W^{d-1}V^{d-1}) = O(W^{O(1)})$ and applying (6.31) for $1 \leq k \leq D$, we find that

$$\|V_1^\gamma \alpha_D\|_{\mathbb{R}/\mathbb{Z}} = O(QW^{O(1)}/T^{1/d}).$$

Thus there is an integer n' such that

$$|\alpha_D - n'/V_1^\gamma| = O(QW^{O(1)}/T^{1/d}).$$

Comparing this with (6.32) gives

$$|n/V_d^D - n'/V_1^\gamma| = O(QW^{O(1)}/T^{1/d}).$$

Now we must have

$$n/V_d^D = n'/V_1^\gamma$$

since otherwise we would have $|n/V_d^D - n'/V_1^\gamma| \geq 1/(V_d^D V_1^\gamma)$, which (due to the constraint on T) is not of size $O(QW^{O(1)}/T^{1/d})$. Since V_1, V_d are coprime, we conclude that n is divisible by V_d^D , and (6.32) tells us that

$$\|\alpha_D\|_{\mathbb{R}/\mathbb{Z}} \leq Q/T^D,$$

as desired.

We claim that

$$\left\| \alpha_D(\tilde{P}_{[r,V]})^D \right\|_{C^\infty[(T/V_d)^{1/d}]} = O(Q).$$

Indeed, since each $V_i \ll (WV)^{d-1} \ll V_d$, we see that the z^i -coefficient of $(\tilde{P}_{[r,V]})^D$ has size $\ll V_d^D$. Then for each $0 \leq i \leq dD$, the distance from the z^i -coefficient of $\alpha_D(\tilde{P}_{[r,V]})^D$ to the nearest integer is at most

$$\ll (Q/T^D) \cdot V_d^D \leq (V_d/T)^{i/d} Q;$$

this establishes the claim. Hence

$$\left\| (R \circ \tilde{P}_{[r,V]}) - \alpha_D(\tilde{P}_{[r,V]})^D \right\|_{C^\infty[(T/V_d)^{1/d}]} = O(Q),$$

and the lemma follows from the induction hypothesis. $\square$

6.11 Sárközy-type results

The goal of this section is to make precise the heuristic that

$$\sum_{\mathbf{x}} \mathbb{E}_{z \in [(N/V_d)^{1/d}]} f_0(\mathbf{x}) f_1(\mathbf{x} + \mathbf{v} \tilde{P}_{[r,V]}(z)) \approx \sum_{\mathbf{x}} \mathbb{E}_{z \in [N]} f_0(\mathbf{x}) f_1(\mathbf{x} + \mathbf{v} z), \quad (6.33)$$

where $V = W^{O_P(1)}$ is a nonnegative integer power of W , $r \in [V]$ (in particular, if $V = 1$, then $r = 0$ and hence $\tilde{P}_{[r,V]} = \tilde{P}$), and V_i 's are the coefficients of $\tilde{P}_{[r,V]}$. We will do so via circle-method arguments that originated in work of Sárközy [96, 97], and we will more closely follow a method due to Green [48].

Throughout this section, define $K := (N/V_d)^{1/d}$, so that $|\tilde{P}_{[r,V]}(z)| \ll_P N$ for $z \in [K]$. For $\xi \in \mathbb{R}/\mathbb{Z}$, we also let

$$S(\xi) := \sum_{z \in [K]} e(\xi \tilde{P}_{[r,V]}(z))$$

denote the Fourier transform of the (multi)set $\tilde{P}_{[r,V]}([K])$.

6.11.1 Circle method preliminaries

The first step towards (6.33) is showing that the counting operator for a 2-point polynomial progression is always controlled by the major arcs of its arguments. Let us make a few more definitions. Let $\varepsilon > 0$ be a constant to be determined later. For coprime $0 \leq a < q \leq K^\varepsilon$, we define the *major arc*

$$\mathfrak{M}_{a,q} := \left\{ \xi \in \mathbb{R}/\mathbb{Z} : \left| \xi - \frac{a}{q} \right| \leq K^{-(d-\varepsilon)} \right\}.$$

As long as $\varepsilon < 1/3$, the intervals $\mathfrak{M}_{a,q}$ are disjoint. So we partition $\mathbb{R}/\mathbb{Z}$ into the *major arcs* $\mathfrak{M}$ and the *minor arcs* $\mathfrak{m}$, where

$$\mathfrak{M} := \bigcup_{\substack{1 \leq a \leq q \leq K^\varepsilon, \\ \gcd(a,q)=1}} \mathfrak{M}_{a,q} \quad \text{and} \quad \mathfrak{m} := (\mathbb{R}/\mathbb{Z}) \setminus \mathfrak{M}.$$

A key input for controlling the minor arcs is the following upper bound on the even moments of F .

Lemma 6.48 (Crude bound on even moments of F). *There exists some positive integer $s_0(d) > d$ such that for all $s \geq s_0(d)$ and $\eta > 0$, we have*

$$\int_{\mathbb{R}/\mathbb{Z}} |S(\xi)|^{2s} d\xi \ll_{P,s,\eta} (WV)^{2s} K^{2s-d+\eta}.$$

The presence of both W and V in the bound is because V can be any nonnegative power of W , including 1. In applications, we will take W and V to be sufficiently small in terms of K so that the right-hand side can be replaced by $O_{P,s,\eta}(K^{2s-d+\eta})$.

Proof. We prove this result by reducing S , an exponential sum in $\tilde{P}_{[r,V]}$, to an exponential sum in P , which can then be treated with classical circle-method results. Recalling that

$$\tilde{P}_{[r,V]}(z) = \frac{\tilde{P}(Vz + r) - \tilde{P}(r)}{V} = \frac{P(\beta_1 W V z + \beta_1 W r - \rho)}{\beta_1^2 W V} - \frac{\tilde{P}(r)}{V},$$

we can recast $S(\xi)$ as

$$S(\xi) = \sum_{z \in [K]} e\left(\frac{\xi P(Az + B)}{\beta_1^2 W V} + \xi_0\right)$$

for

$$A = \beta_1 W V, \quad B = \beta_1 W r - \rho, \quad \text{and} \quad \xi_0 = -\frac{\xi \tilde{P}(r)}{V},$$

and so

$$\int_{\mathbb{R}/\mathbb{Z}} |S(\xi)|^{2s} d\xi = \int_0^1 \left| \sum_{z \in [K]} e\left(\frac{\xi P(Az + B)}{\beta_1^2 W V}\right) \right|^{2s} d\xi.$$

By changing variables $\xi \mapsto \beta_1^2 W V \xi$ and then using nonnegativity to extend the domain of integration from $[0, 1/(\beta_1^2 W V))$ to $[0, 1)$, we can upper bound this quantity by

$$\beta_1^2 W V \cdot \int_0^1 \left| \sum_{z \in [K]} e(\xi P(Az + B)) \right|^{2s} d\xi, \tag{6.34}$$

Expanding the $2s$ -th power and using orthogonality, we can recast (6.34) as $\beta_1^2 W V$ times

$$\sum_{\substack{z_1, \dots, z_s, \\ z'_1, \dots, z'_s \in [K]}} 1(P(Az_1 + B) + \dots + P(Az_s + B) = P(Az'_1 + B) - \dots - P(Az'_s + B)). \tag{6.35}$$

Since $Az + B \in [AK] - \rho$ for every $z \in [K]$, we can use nonnegativity to upper-bound (6.35) by

$$\sum_{\substack{z_1, \dots, z_s, \\ z'_1, \dots, z'_s \in [AK] - \rho}} 1(P(z_1) + \dots + P(z_s) = P(z'_1) - \dots - P(z'_s)),$$

which on Fourier expanding equals

$$\int_0^1 \left| \sum_{z \in [AK]} e(\xi P(z + \rho)) \right|^{2s} d\xi.$$

By e.g. [115, Theorem 14.1], this is at most

$$\ll_{P,\eta} (AK)^{2s-d+\eta} \ll_{P,s,\eta} (WV)^{2s-1} K^{2s-d+\eta},$$

and the result follows by combining this with the factor $\beta_1^2 WV$. $\square$

The following lemma is a generalization of [88, Lemma B.4], and it follows from a nearly identical proof, which we omit.

Lemma 6.49 (Minor arc control). *Let $s_0(d)$ be as in Lemma 6.48, and let $s > s_0(d)$ be a positive integer. There exists an absolute $C = C(P, s) > 0$ such that for all $\varepsilon \in (0, 1)$ and positive integers $K^\varepsilon \gg_{P,s,\varepsilon} V^C$, we have the pointwise bound*

$$\sup_{\xi \in \mathfrak{m}} |S(\xi)| \leq K^{1-\varepsilon/C}$$

as well as the L^{2s} -bound

$$\int_{\mathfrak{m}} |S(\xi)|^{2s} \ll_{P,s,\varepsilon} K^{2s-d-\varepsilon/C}.$$

In order to handle the major arcs, we start with the following classical decomposition result for F (see e.g. [114, Corollary 5.2]).

Lemma 6.50. *Let $0 < \varepsilon < 1/3$ and $0 \leq a < q \leq K^\varepsilon$ with a, q coprime. If $\xi \in \mathfrak{M}_{a,q}$, then*

$$S(\xi) = \mathbb{E}_{u \in [q]} e(a\tilde{P}_{[r,V]}(u)/q) \int_0^K e\left(\left(\xi - \frac{a}{q}\right) \tilde{P}_{[r,V]}(z)\right) dz + O(K^{2\varepsilon}).$$

We recall that the assumption $\varepsilon < 1/3$ ensures the disjointness of the major arcs.

In what follows, let

$$S(a, q) := \mathbb{E}_{u \in [q]} e(a\tilde{P}_{[r,V]}(u)/q) = \mathbb{E}_{u \in \mathbb{Z}/q\mathbb{Z}} e(a\tilde{P}_{[r,V]}(u)/q)$$

be the exponential sum appearing in Lemma 6.50. The following lemma gives bounds on $S(a, q)$ for various

values of a, q . It is inspired by [23, Lemma 5.3]; the proof is almost identical, so we omit it.

Lemma 6.51. *Let $1 \leq a < q$ be coprime positive integers. Then*

$$|S(a, q)| \ll_{P, \eta} q^{\eta-1/d} \quad \text{for every } \eta > 0.$$

Moreover, $S(a, q) = 0$ whenever $\gcd(q, W) > 1$.

Next, we upper-bound the integral appearing in Lemma 6.50. It essentially follows from the van der Corput Lemma (e.g., [105, pp. 332]) and we omit the proof.

Lemma 6.52. *We have*

$$\left| \int_0^K e(\xi \tilde{P}_{[r, V]}(z)) dz \right| \ll_P 1 + |\xi|^{-1/d} V_d^{-1/d}.$$

Combining the previous three lemmas, we can finally upper-bound the contribution of the major arcs.

Lemma 6.53 (Major arc control). *For $s > d$ and $\varepsilon \in (0, 1/10]$, we have*

$$\int_{\mathfrak{M}} |S(\xi)|^{2s} d\xi \ll_{P, s} \frac{K^{2s-d}}{V_d}$$

as long as $K \geq V_d^2$.

We recall that the parameter ε figures implicitly in the definition of the major arcs $\mathfrak{M}$; since the left-hand side of the inequality in the lemma is non-decreasing as ε increases, the important point is that we can take ε bounded away from 0.

Proof. We let all constants depend on P, s . Applying Lemma 6.50, integrating over all of the major arcs, and changing variables $\xi \mapsto \xi + a/q$ on each $\mathfrak{M}_{a, q}$, we have

$$\begin{aligned} \int_{\mathfrak{M}} |S(\xi)|^{2s} d\xi &= \left(1_{(a, q)=(0, 1)} + \sum_{\substack{1 \leq a < q \leq K^\varepsilon, \\ \gcd(a, q)=1}} |S(a, q)|^{2s} \right) \cdot \int_{-K^{-d+\varepsilon}}^{K^{-d+\varepsilon}} \left| \int_0^K e(\xi \tilde{P}_{[r, V]}(z)) dz \right|^{2s} d\xi \\ &\quad + O\left(1_{(a, q)=(0, 1)} + \sum_{\substack{1 \leq a < q \leq K^\varepsilon, \\ \gcd(a, q)=1}} 1 \right) (K^{2s-1-d+3\varepsilon}). \end{aligned}$$

We bound the terms in the expression above one by one, starting with the contribution of the sums

$S(a, q)$. For every $\eta > 0$, we can bound

$$1_{(a,q)=(0,1)} + \sum_{\substack{1 \leq a < q \leq K^\varepsilon, \\ \gcd(a,q)=1}} |S(a, q)|^{2s} \ll_\eta 1 + \sum_{\substack{1 \leq a < q \leq K^\varepsilon, \\ \gcd(a,q)=1}} q^{\eta-2s/d} \ll_\eta 1 + K^{\varepsilon(\eta+2-2s/d)},$$

where the first inequality follows from an application of Lemma 6.51 with $2\eta s$ in place of η , and the second bound follows from trivially estimating the number of a 's in $1 \leq a < q$ by q and then trivially bounding $q \leq K^\varepsilon$. Importantly, this expression is $O(1)$ as long as $s > d$ and $\eta < 2/d$. Since there are no further restrictions on the choice of η , we can take any η in the range above.

To bound the contribution of the integral, we split it over ξ at K^{-d}/V_d , obtaining

$$\begin{aligned} \int_{-K^{-d+\varepsilon}}^{K^{-d+\varepsilon}} \left| \int_0^K e\left(\xi \tilde{P}_{[r,V]}(z)\right) dz \right|^{2s} d\xi \\ \leq \int_{-K^{-d}/V_d}^{K^{-d}/V_d} K^{2s} d\xi + 2 \int_{K^{-d}/V_d}^1 \left| \int_0^K e\left(\xi \tilde{P}_{[r,V]}(z)\right) dz \right|^{2s} d\xi. \end{aligned}$$

The first term can be bounded trivially by K^{2s-d}/V_d . To bound the second term, we apply Lemma 6.52, which gives

$$\int_{K^{-d}/V_d}^1 \left| \int_0^K e\left(\xi \tilde{P}_{[r,V]}(z)\right) dz \right|^{2s} d\xi \ll \int_{K^{-d}/V_d}^1 |\xi|^{-2s/d} V_d^{-2s/d} d\xi \ll \frac{K^{2s-d}}{V_d},$$

where we use Lemma 6.52 to bound the integrand and the assumption $s > d$ to ensure that the exponent is of the right order.

Lastly, the error term is at most

$$\ll \left(1_{(a,q)=(0,1)} + \sum_{\substack{1 \leq a < q \leq K^\varepsilon, \\ \gcd(a,q)=1}} 1 \right) K^{2s-1-d+3\varepsilon} \ll K^{2s-1-d+5\varepsilon},$$

which can be bounded by $\ll K^{2s-d}/V_d$ since $K \geq V_d^2$ and $\varepsilon \leq 1/10$. □

We summarize the computations performed in this section in the result below.

Proposition 6.54 (Even moment bound). *There exists $s_1(d) \in \mathbb{N}$ and $C = C(d) > 0$ such that if $s \geq s_1(d)$ and $N \geq V_d^C$, then*

$$\int_{\mathbb{R}/\mathbb{Z}} |S(\xi)|^{2s} d\xi \ll_{P,s} \frac{K^{2s}}{N}.$$

Proof. Let $s_1(d) = \max(s_0(d), d+1)$, where s_0 comes from Lemma 6.49 and $d+1$ is the smallest possible s allowed by Lemma 6.53. We fix $s > s_1(d)$ and let all the implicit constants depend on P, s . We split $\mathbb{R}/\mathbb{Z}$ into major and minor arcs with $\varepsilon = 1/10$. Then we can use Lemma 6.49 to bound the minor arc contribution as

$$\int_{\mathfrak{m}} |S(\xi)|^{2s} d\xi \ll K^{2s-d-\varepsilon/C_1}$$

where $C_1 > 0$ is the constant from Lemma 6.49. Since $N = K^d V_d$, this can be bounded by $\ll \frac{K^{2s}}{N}$ as long as $V_d \leq K^{\varepsilon/C_1}$, which we can assume by taking the implicit constant in $N \geq V_d^{\Omega(1)}$ to be sufficiently large. We also implicitly use the assumption $N \geq V_d^{\Omega(1)}$ to ensure that $K^{\varepsilon/4} \geq V_d$ so that the assumption on the coefficients of $\tilde{P}_{[r,V]}$ in Lemma 6.49 is satisfied. To bound the major arc contribution by the same final expression, we simply apply Lemma 6.53. □

6.11.2 Applications to Sárközy-type estimates

We are now ready to show that counting operators for Sárközy-type configurations are controlled by the major arcs of their arguments.

Proposition 6.55. *For every $C_1 > 0$ there exists $C_2 = C_2(C_1, P) > 0$ such that the following holds. Let $\delta \in (0, 1/10)$, let $V \leq W^{C_1}$ be a nonnegative integer power of W , let $r \in [V]$, and let $N \geq C_2 V_d^{C_2} \delta^{-C_2}$ be an integer. Let $f_0, f_1 : \mathbb{Z} \rightarrow \mathbb{C}$ be 1-bounded functions supported on $[N]$. If*

$$\left| \sum_x \mathbb{E}_{z \in [K]} f_0(x) f_1(x + \tilde{P}_{[r,V]}(z)) \right| \geq \delta N, \quad (6.36)$$

then for every $i = 0, 1$ there exist phases $\xi_i \in \mathbb{R}/\mathbb{Z}$ and positive integers $q_i \ll_{C_1, P} \delta^{-O_{C_1, P}(1)}$ such that $\|q_i \xi_i\|_{\mathbb{R}/\mathbb{Z}} \ll_{C_1, P} \delta^{-O_{C_1, P}(1)}/N$ and $|\hat{f}_i(\xi_i)| \gg_{C_1, P} \delta^{O_{C_1, P}(1)} N$.

Proof. We fix $C_1 > 0$ and let all constants depend on P . We prove the bound only for $i = 0$, since one can then deduce the result for $i = 1$ by shifting $x \mapsto x - \tilde{P}_{[r,V]}(z)$ and exchanging the roles of f_0, f_1 . Our first goal is to show that the largeness of (6.36) implies the largeness of some Fourier coefficient of f_0 , and then we use the fraction comparison argument to deduce that this Fourier coefficient must be a major arc.

Define the multiset $E := \{\tilde{P}_{[r,V]}(z) : z \in [K]\}$, so that our assumption can be expressed as

$$\left| \sum_{x, z} f_0(x) f_1(x + z) 1_E(z) \right| \geq \delta N K.$$

Note that $F = \widehat{1_E}$. Now, Fourier-expanding, extracting the largest Fourier coefficient, and applying the triple Hölder inequality (with exponents $\frac{2s}{s-1}, 2, 2s$) followed by Parseval's Identity, we get

$$\begin{aligned} \delta NK &\leq \|\widehat{f_0}\|_\infty^{1/s} \cdot \int_{\mathbb{R}/\mathbb{Z}} |\widehat{f_0}(\xi)|^{1-1/s} \cdot |\widehat{f_1}(\xi)| \cdot |F(\xi)| \, d\xi \\ &\leq \|\widehat{f_0}\|_\infty^{1/s} \cdot \|f_0\|_2^{1-1/s} \cdot \|f_1\|_2 \cdot \left(\int_{\mathbb{R}/\mathbb{Z}} |F(\xi)|^{2s} \, d\xi \right)^{1/(2s)} \end{aligned}$$

for any $s \geq 2$. We can trivially bound $\|f_0\|_2^{1-1/s} \cdot \|f_1\|_2 \leq N^{(2s-1)/(2s)}$, and Proposition 6.54 gives

$$\int_{\mathbb{R}/\mathbb{Z}} |S(\xi)|^{2s} \, d\xi \ll \frac{K^{2s}}{N}$$

for any s sufficiently large in terms of d . Combining all these bounds together, we get the bound $\|\widehat{f_i}\|_\infty \gg \delta^{O(1)}N$, i.e. we have shown that the Sárközy configuration is controlled by Fourier analysis.

Now, let $\tilde{f}_0 = \mathbb{E}_{z \in K} \overline{f_1}(x + \tilde{P}_{[r,V]}(z))$. Stashing, we can replace f_0 by $\tilde{f}_0$ in (6.36) at the cost of squaring δ in the lower bound. Together with the just proved Fourier control, this gives us

$$\left| \sum_x \mathbb{E}_{z \in [K]} \overline{f_1}(x + \tilde{P}_{[r,V]}(z)) e(-\xi x) \right| \gg \delta^{O(1)}N$$

for some $\xi \in \mathbb{R}/\mathbb{Z}$. Shifting $x \mapsto x - \tilde{P}_{[r,V]}(z)$, splitting the sums over x and z , and using $\|f_1\|_1 \leq N$, we get that

$$\left| \sum_x f_1(x) e(\xi x) \right| \gg \delta^{O(1)}N \quad \text{and} \quad \left| \mathbb{E}_{z \in [K]} e(-\xi \tilde{P}_{[r,V]}(z)) \right| \gg \delta^{O(1)}.$$

It remains to show that ξ is a major arc. By Weyl's inequality (e.g. [51, Lemma 4.4]), the second inequality can only hold if there exists a positive integer $q \ll \delta^{-O(1)}$ for which $\|q\xi \tilde{P}_{[r,V]}\|_{C^\infty[K]} \ll \delta^{-O(1)}$. The result then follows from Lemma 6.47. $\square$

The following two corollaries make precise the heuristic (6.33).

Corollary 6.56. *For every $C_1 > 0$ there exists $C_2 = C_2(C_1, P) > 0$ such that the following holds. Let $\delta \in (0, 1/10)$, let $V \leq W^{C_1}$ be a nonnegative integer power of W , let $r \in [V]$, and let $N \geq C_2 V_d^{C_2} \delta^{-C_2}$ be an integer. If $f_0, f_1 : \mathbb{Z} \rightarrow \mathbb{C}$ are 1-bounded functions supported on $[N]$ satisfying*

$$\left| \sum_x \mathbb{E}_{z \in [K]} f_0(x) f_1(x + \tilde{P}_{[r,V]}(z)) \right| \geq \delta N,$$

then for each $i \in \{0, 1\}$ there exists a positive integer $q_i \ll_{C_1, P} \delta^{-O_{C_1, P}(1)}$ such that

$$\mathbb{E}_{u \in [\pm N]} \left| \mathbb{E}_{z \in [N']} f_i(u + q_i z) \right| \gg_{C_1, P} \delta^{O_{C_1, P}(1)} \quad \text{for all } 1 \leq N' \ll_{C_1, P} \delta^{O_{C_1, P}(1)} N.$$

Proof. We fix $C_1 > 0$ and let all constants depend on P . We just proof the claim for $i = 1$, as the other case follows by symmetry. By Proposition 6.55, there exist a phase $\xi \in \mathbb{R}/\mathbb{Z}$ and a positive integer $q \ll \delta^{-O(1)}$ such that $\|q\xi\|_{\mathbb{R}/\mathbb{Z}} \ll \delta^{-O(1)}/N$ and

$$\left| \sum_x f_1(x) e(\xi x) \right| \gg \delta^{O(1)} N.$$

Set $N'_0 \asymp \delta^C N$ for some $C > 1$ to be chosen later, and let $N' \in [N'_0]$ be arbitrary. Introducing an extra averaging over arithmetic progressions of length N' and common difference q gives

$$\left| \sum_u \mathbb{E}_{z \in [N']} f_1(u + qz) e(\xi(u + qz)) \right| \gg \delta^{O(1)} N.$$

The triangle inequality gives

$$\sum_u \left| \mathbb{E}_{z \in [N']} f_1(u + qz) e(\xi qz) \right| \gg_d \delta^{O_d(1)} N.$$

We have

$$|e(\xi qz) - 1| \leq |z| \cdot \|\xi q\|_{\mathbb{R}/\mathbb{Z}} \ll \delta^{-O(1)} N'/N$$

for all $z \in [N']$, so, with C sufficiently large, we obtain

$$\sum_u \left| \mathbb{E}_{z \in [N']} f_1(u + qz) \right| \gg \delta^{O(1)} N.$$

Finally, since f_1 is supported on $[N]$, there is no contribution from u 's outside of $[\pm N]$, and we get

$$\mathbb{E}_{u \in [\pm N]} \left| \mathbb{E}_{z \in [N']} f_1(u + qz) e(\xi u) \right| \gg \delta^{O(1)},$$

as desired. □

We now deduce a multidimensional version of this corollary.

Corollary 6.57. *For every $C_1 > 0$ there exists $C_2 = C_2(C_1, P) > 0$ such that the following holds. Let $\delta \in (0, 1/10)$, let $V \leq W^{C_1}$ be a nonnegative integer power of W , let $r \in [V]$, and let $N \geq C_2 V_d^{C_2} \delta^{-C_2}$ be*

an integer. Let $D \in \mathbb{N}$, and let $\mathbf{v} \in \mathbb{Z}^D$ have coordinates of size $O(1)$. If $f_0, f_1 : \mathbb{Z}^D \rightarrow \mathbb{C}$ are 1-bounded functions supported on $[N]^D$ satisfying

$$\left| \sum_{\mathbf{x}} \mathbb{E}_{z \in [K]} f_0(\mathbf{x}) f_1(\mathbf{x} + \mathbf{v} \tilde{P}_{[r, V]}(z)) \right| \geq \delta N^D,$$

then for each $i \in \{0, 1\}$ there exists a positive integer $q_i \ll_{C_1, D, P} \delta^{-O_{C_1, D, P}(1)}$ such that

$$\sum_{\mathbf{x}} \left| \mathbb{E}_{z \in [N']} f_i(\mathbf{x} + q_i \mathbf{v} z) \right| \gg_{C_1, D, P} \delta^{O_{C_1, D, P}(1)} N^D \quad \text{for all } 1 \leq N' \leq \delta^{O_{C_1, D, P}(1)} N.$$

Proof. We fix $C_1 > 0$ and let all constants depend on D and P . We just proof the claim for $i = 1$, as the other case follows by symmetry. Introducing an extra averaging over $\mathbf{v} \cdot [N]$ gives

$$\sum_{\mathbf{x}} \left| \sum_{y \in [N]} \mathbb{E}_{z \in [K]} f_0(\mathbf{x} + \mathbf{v} y) f_1(\mathbf{x} + \mathbf{v}(y + \tilde{P}_{[r, V]}(z))) \right| \geq \delta N^{D+1}.$$

Since the coordinates of $\mathbf{v}$ have size $O(1)$, the functions $\mathbf{x} \mapsto f_0(\mathbf{x} + \mathbf{v} y)$ vanish whenever $|\mathbf{x}| \gg N$ as y ranges over $[N]$. Pigeonholing in $\mathbf{x}$, we can therefore find $\Omega(\delta N^D)$ values $\mathbf{x} \in \mathbb{Z}^D$ for which

$$\left| \sum_{y \in [N]} \mathbb{E}_{z \in [K]} f_0(\mathbf{x} + \mathbf{v} y) f_1(\mathbf{x} + \mathbf{v}(y + \tilde{P}_{[r, V]}(z))) \right| \gg \delta N, \quad (6.37)$$

and so the previous corollary gives us positive integers $q_{\mathbf{x}} \ll \delta^{-O(1)}$ and $\delta^{O(1)} N \ll N'_{0, \mathbf{x}} \leq N$ such that

$$\mathbb{E}_{u \in [\pm N]} \left| \mathbb{E}_{z \in [N'_{\mathbf{x}}]} f_1(\mathbf{x} + \mathbf{v}(u + q_{\mathbf{x}} z)) \right| \gg \delta^{O(1)}$$

for all $N'_{\mathbf{x}} \in [N'_{0, \mathbf{x}}]$. We set $N'_0 := \min_{\mathbf{x}} N'_{0, \mathbf{x}}$ to eliminate the dependence on $\mathbf{x}$, where the minimum is taken over the $\Omega(\delta N^D)$ good values $\mathbf{x} \in \mathbb{Z}^D$. Since the lower bounds on $N'_{0, \mathbf{x}}$ are independent on $\mathbf{x}$, we deduce that $N'_0 \gg \delta^{O(1)} N$. Similarly, since there are only $O(\delta^{-O(1)})$ choices of $q_{\mathbf{x}}$, the pigeonhole principle produces some $q \ll \delta^{-O(1)}$ for which

$$\sum_{\mathbf{x}} \mathbb{E}_{u \in [\pm N]} \left| \mathbb{E}_{z \in [N']} f_1(\mathbf{x} + \mathbf{v}(u + qz)) \right| \gg \delta^{O(1)} N^D$$

for all $N' \leq N'_0$. The result follows upon shifting $\mathbf{x} \mapsto \mathbf{x} - \mathbf{v} u$. $\square$

6.12 Fourier uniformity of difference of weights

The purpose of this section is to show that the weight ν^* from the proof of Proposition 6.20 is highly Fourier-uniform on arithmetic progressions of difference V . Our argument is a slight modification of the proof of [88, Lemma B.8]. Recall that the leading term of $\tilde{P}(z)$ is $W_d z^d$, where $W_d = \beta_d \beta_1^{d-2} W^{d-1}$, and our weights are

$$\begin{aligned}\tilde{\nu}(z) &= \frac{N}{(N/W_d)^{1/d}} 1_{\tilde{P}([(N/W_d)^{1/d}])}(z) = N^{(d-1)/d} W_d^{1/d} \sum_{z' \in [(N/W_d)^{1/d}]} 1_{\tilde{P}(z')}(z), \\ \nu(z) &= d^{-1} 1_{[N]}(z) (N/(z+1))^{(d-1)/d}, \quad \text{and} \quad \nu^*(z) = \tilde{\nu}(z) - \nu(z).\end{aligned}$$

The following proposition captures the main estimate.

Proposition 6.58. *For every $C_1 > 0$ there exist $C_2 = C_2(C_1, P)$, $c = c(C_1, P) > 0$ such that the following holds. If $V \leq W^{C_1}$ is a nonnegative integer power of W and $N \geq W^{C_2}$, then*

$$\max_{r \in [V]} \max_{\theta \in \mathbb{R}/\mathbb{Z}} \left| \sum_{z \in \mathbb{Z}} e(\theta z) \nu^*(Vz + r) \right| \ll_{C_1, P} \frac{N}{V w^c}.$$

Notice that the trivial bound (from the triangle inequality) is $\ll N/V$; our gain is in the (relatively modest) w^{-c} term.

Proof. We fix $C_1 > 0$ and let all the constants depend on C_1 and P . Hensel's Lemma and the identity $\tilde{P}(z) \equiv z \pmod{W}$ ensure that for each $r \in [V]$ there is a unique $r' \in [V]$ such that $\tilde{P}(r') \equiv r \pmod{V}$. It follows that $\tilde{P}(z') \equiv r \pmod{V}$ if and only if $z' \equiv r' \pmod{V}$, so

$$\begin{aligned}\sum_z e(\theta z) \tilde{\nu}(Vz + r) &= N^{(d-1)/d} W_d^{1/d} \sum_z e(\theta z) \sum_{z' \in [(N/W_d)^{1/d}]} 1_{\tilde{P}(z')}(Vz + r) \\ &= N^{(d-1)/d} W_d^{1/d} \sum_{z' \in [(N/W_d)^{1/d} V^{-1}]} e\left(\theta \frac{\tilde{P}(Vz' + r') - r}{V}\right).\end{aligned}$$

Substituting z for z' and $V\theta$ for θ , and then multiplying both $\tilde{\nu}$ and ν by $e(\theta r)$, we find that it suffices to prove the inequality

$$\begin{aligned}\left| N^{(d-1)/d} W_d^{1/d} \sum_{z \in [(N/W_d)^{1/d} V^{-1}]} e(\theta \tilde{P}(Vz + r')) \right. \\ \left. - d^{-1} \sum_{z \in [N/V]} (N/(Vz + r))^{(d-1)/d} e(\theta(Vz + r + 1)) \right| \ll \frac{N}{V w^c}\end{aligned}$$

for all $r \in [V]$ and $\theta \in \mathbb{R}/\mathbb{Z}$.

We will require different arguments when $V\theta$ is in a minor arc, when it is in the major arc around 0, and when it is in a major arc not around 0.⁸

We first analyze when the contribution of ν can be large. By summation by parts, we have that

$$d^{-1} \left| \sum_{z \in [N/V]} (N/(Vz + r + 1))^{(d-1)/d} e(\theta(Vz + r)) \right|$$

equals

$$\begin{aligned} & d^{-1} \left| \sum_{z \in [N/V]} (N/(Vz + r + 1))^{(d-1)/d} e(\theta Vz) \right| \\ & \ll \sum_{t \in [N/V-1]} \left(\left(\frac{N}{Vt + r + 1} \right)^{(d-1)/d} - \left(\frac{N}{V(t+1) + r + 1} \right)^{(d-1)/d} \right) \left| \sum_{z \in [t]} e(V\theta z) \right| \\ & \quad + \left| \sum_{z \in [N/V]} e(V\theta z) \right| \\ & \ll \sum_{t \in [N/V-1]} (N/V)^{(d-1)/d} t^{(1-2d)/d} \min \left\{ t, \frac{1}{\|V\theta\|_{\mathbb{R}/\mathbb{Z}}} \right\} + \min \left\{ N/V, \frac{1}{\|V\theta\|_{\mathbb{R}/\mathbb{Z}}} \right\} \\ & \ll \frac{(N/V)^{(d-1)/d}}{\|V\theta\|_{\mathbb{R}/\mathbb{Z}}^{1/d}} + \min \left\{ N/V, \frac{1}{\|V\theta\|_{\mathbb{R}/\mathbb{Z}}} \right\}, \end{aligned}$$

where in the second-to-last line we split the sum at $t = 1/\|V\theta\|_{\mathbb{R}/\mathbb{Z}}$. This quantity is $\ll N/(Vw^c)$ unless $\|V\theta\|_{\mathbb{R}/\mathbb{Z}} < Vw^{cd}/N$.

We now analyze when the contribution of $\tilde{\nu}$ can be large. By [51, Proposition 4.3] and [73, Lemma A.10], we have

$$\left| N^{(d-1)/d} W_d^{1/d} \sum_{z \in [(N/W_d)^{1/d} V^{-1}]} e(\theta \tilde{P}(Vz + r')) \right| \ll \frac{N}{Vw^c}$$

unless there is some positive integer $q \ll w^{O(c)}$ such that

$$\left\| q\theta \tilde{P}(V \cdot + r') \right\|_{C^\infty[(N/W_d)^{1/d} V^{-1}]} \ll w^{O(c)}. \quad (6.38)$$

Let $R(z) := q\theta(Vz + \tilde{P}(r'))$, so that $R \circ \tilde{P}_{[r', V]}(z) = q\theta \tilde{P}(Vz + r')$. Setting $T = N/V$ and recalling that

⁸These major and minor arcs are different from the major and minor arcs in Section 6.11.

$V_d = W_d V^{d-1}$, we can write the estimate (6.38) as

$$\left\| R \circ \tilde{P}_{[r', V]} \right\|_{C^\infty[(T/V_d)^{1/d}]} \ll w^{O(c)}.$$

Lemma 6.47 then gives

$$\|qV\theta\|_{\mathbb{R}/\mathbb{Z}} \ll \frac{Vw^{O(c)}}{N}.$$

Hence the $\tilde{\nu}$ -contribution is $\ll N/(Vw^c)$ unless there is some positive integer $q \ll w^{O(c)}$ such that $\|qV\theta\|_{\mathbb{R}/\mathbb{Z}} \ll Vw^{O(c)}/N$.

Comparing the outputs of the previous two paragraphs, we see that the conclusion of the proposition holds in the “minor arc” case where there is no $q \ll w^{O(c)}$ such that $\|qV\theta\|_{\mathbb{R}/\mathbb{Z}} \ll Vw^{O(c)}/N$. Henceforth, we assume that we are in the “major arc” case where there is some $q \ll w^{O(c)}$ such that $\|qV\theta\|_{\mathbb{R}/\mathbb{Z}} \leq Vw^{O(c)}/N$. In this case, we write

$$V\theta = a/q + \theta^*,$$

where a, q are coprime integers with $1 \leq q \ll w^{O(c)}$, and

$$|\theta^*| \ll Vw^{O(c)}/N \leq (W_d/N)^{1-1/(4d)} V^{d-1/4}$$

(say) since N bounded from below be a sufficiently large power of W that is allowed to depend on d . Since the implicit constant in the term $w^{O(c)}$ depends only on d , choosing c sufficiently small (depending on d) guarantees that $q < w$.

First, assume that $q > 1$. The ν -contribution is $\ll N/(Vw^c)$. To handle the $\tilde{\nu}$ -contribution, we use Lemma 6.50 with $Q(z) = \tilde{P}_{[r', V]}(z)$, $K = (N/W_d)^{1/d} V^{-1}$, and $\varepsilon = 1/4$ to express

$$\begin{aligned} & \sum_{z \in [(N/W_d)^{1/d} V^{-1}]} e(\theta \tilde{P}(Vz + r')) \\ &= e(\theta \tilde{P}(r')) \mathbb{E}_{u \in [q]} e((a/q) \tilde{P}_{[r', V]}(u)) \int_0^{(N/W_d)^{1/d} V^{-1}} e(\theta^* \tilde{P}_{[r', V]}(x)) dx \\ & \quad + O((N/W_d)^{1/(2d)} V^{-1/2}). \end{aligned}$$

Since q is smaller than w , it shares a prime factor with W , and so Lemma 6.51 tells us that the average

$\mathbb{E}_{u \in [q]} e((a/q)\tilde{P}_{[r', V]}(u))$ vanishes. Hence

$$N^{(d-1)/d} W_d^{1/d} \sum_{z \in [(N/W_d)^{1/d} V^{-1}]} e(\theta \tilde{P}(Vz + r')) \ll N/(Vw^c)$$

(with room to spare) by the assumption that N is larger than a suitable power of V .

Finally, assume that $q = 1$, in which case $a = 0$ and $\|V\theta\|_{\mathbb{R}/\mathbb{Z}} \ll Vw^{O(c)}/N$. We will break the sum in the ν -contribution into intervals and show that each one approximately cancels out with a single summand from the $\tilde{\nu}$ -contribution. Since

$$\tilde{P}_{[r', V]}(z) - W_d V^{d-1} z^d \ll V^{O(1)} z^{d-1}$$

and

$$\tilde{P}(r') - r = Vu \quad \text{for some integer } |u| \ll W_d V^{d-1},$$

we have

$$\left| \sum_{z \in [(N/W_d)^{1/d} V^{-1}]} (e(\theta \tilde{P}(Vz + r')) - e(\theta(W_d V^d z^d + r))) \right| \ll V^{O(1)}.$$

This lets us simplify the $\tilde{\nu}$ -contribution by removing lower-order terms, and it suffices to show that

$$\left| N^{(d-1)/d} W_d^{1/d} e(\theta r) \sum_{z \in [(N/W_d)^{1/d} V^{-1}]} e(\theta W_d V^d z^d) - d^{-1} \sum_{z \in [N/V]} (N/(Vz + r + 1))^{(d-1)/d} e(\theta(Vz + r)) \right| \ll \frac{N}{Vw^c}.$$

We break the second sum into intervals of the form

$$I_z = [A_z, B_z] := [W_d V^{d-1} z^d, W_d V^{d-1} (z+1)^d - 1]$$

with length $|I_z| = W_d V^{d-1} dz^{d-1}(1 + O(z^{-1}))$ for $z \in [(N/W_d)^{1/d} V^{-1}]$. For $t \in I_z$, we bound

$$\begin{aligned} \left| \left(\frac{N}{Vt + r + 1} \right)^{(d-1)/d} - \left(\frac{N}{VA_z} \right)^{(d-1)/d} \right| &\ll \left| \frac{N}{Vt + r + 1} - \frac{N}{VA_z} \right| \cdot \left(\frac{N}{VA_z} \right)^{-1/d} \\ &\leq \left| \frac{VA_z - Vt - r - 1}{Vt + r + 1} \right| \cdot \left(\frac{N}{VA_z} \right)^{(d-1)/d} \\ &\ll \frac{|I_z|}{A_z} \cdot \left(\frac{N}{VA_z} \right)^{(d-1)/d} \ll z^{-1} \left(\frac{N}{VA_z} \right)^{(d-1)/d} \end{aligned}$$

and

$$|e(\theta \cdot Vt) - e(\theta \cdot VA_z)| \leq \|V\theta\|_{\mathbb{R}/\mathbb{Z}} \cdot |I_z| \ll \frac{V^{O(1)}}{N^{1/d}} \ll N^{-1/(2d)} \quad (6.39)$$

as long as $N \gg V^{\Omega(1)}$. Summing over $t \in I_z$, we obtain the key estimate

$$\begin{aligned} & d^{-1} \sum_{t \in I_z} (N/(Vt + r + 1))^{(d-1)/d} e(\theta(Vt + r)) \\ &= e(\theta r) d^{-1} \sum_{t \in I_z} (N/(VA_z))^{(d-1)/d} e(\theta VA_z) (1 + O(N^{-1/2d} + z^{-1})) \\ &= e(\theta r) W_d V^{d-1} z^{d-1} (N/(W_d V^d z^d))^{(d-1)/d} e(\theta W_d V^d z^d) (1 + O(N^{-1/2d} + z^{-1})) \\ &= e(\theta r) N^{(d-1)/d} W_d^{1/d} e(\theta W_d V^d z^d) (1 + O(N^{-1/2d} + z^{-1})). \end{aligned}$$

The lemma follows by summing over $z \in [(N/W_d)^{1/d} V^{-1}]$ and noting that

$$\left| \sum_{z \in [(N/W_d)^{1/d} V^{-1}]} N^{(d-1)/d} W_d^{1/d} (N^{-1/2d} + z^{-1}) \right| \ll \frac{N}{V w^c},$$

again with room to spare since W is sufficiently small compared to N . $\square$

6.13 Supersaturation

The purpose of this short section is to prove a supersaturation result for corners. We apply Varnavides' now-standard trick from [112] to Shkredov's result on corners [102].

Theorem 6.59 ([102]). *There exists a constant $c > 0$ such that the following holds: If $A \subseteq [N]^2$ has size δN^2 for some $\delta \in (0, 1/10)$ and $N \geq \exp \exp(\delta^{-c})$, then A contains a nontrivial corner.*

Proposition 6.60. *There exists a constant $c' > 0$ such that the following holds: If $A \subseteq [N]^2$ has size δN^2 for some $\delta \in (0, 1/10)$ and $N \geq \exp \exp(\delta^{-c'})$, then A contains at least $(\exp(-\exp(\delta^{-c'}))) N^3$ nontrivial corners.*

Proof. Let $M := \exp \exp((\delta/2)^c)$, where c is the constant from Theorem 6.59. For each $1 \leq d \leq N/(3M)$ and $\mathbf{u} \in [-d(M-1), N-1]^2$, define the set

$$\lambda(d, \mathbf{u}) := A \cap (\mathbf{u} + d \cdot [M]^2).$$

For fixed d , each element of A is contained in exactly M^2 of the sets $\lambda(d, \mathbf{u})$, so

$$\sum_{\mathbf{u}} |\lambda(d, \mathbf{u})| = M^2 |A| = \delta M^2 N^2.$$

Say that the pair $(d, \mathbf{u})$ is *good* if $|\lambda(d, \mathbf{u})| \geq (\delta/2)M^2$, and let $X(d)$ denote the number of $\mathbf{u}$'s such that $(d, \mathbf{u})$ is good. We can bound

$$\delta M^2 N^2 = \sum_{\mathbf{u}} |\lambda(d, \mathbf{u})| \leq X(d)M^2 + (N + d(M-1))^2 \cdot (\delta/2)M^2,$$

and so (for instance)

$$X(d) \geq \delta N^2 - (\delta/2)(N + N/3)^2 = (\delta/9)N^2.$$

Hence the total number of good pairs $(d, \mathbf{u})$ (for all d) is at least $(\delta/(27M))N^3$. Theorem 6.59 tells us that if $(d, \mathbf{u})$ is good, then $\lambda(d, \mathbf{u})$ contains a nontrivial corner. Finally, each nontrivial corner is contained in at most M^3 such sets $\lambda(d, \mathbf{u})$ (there are at most M choice of d and then at most M^2 choices of $\mathbf{u}$), so we conclude that the number of nontrivial corners in A is at least

$$\frac{\delta}{27M^4} N^3 \geq (\exp(-\exp(\delta^{-c'})))N^3$$

with a suitable choice of $c' > 0$. □

Bibliography

- [1] A. Adve and V. Giri, A spectral gap for spinors on hyperbolic surfaces. *Preprint*, available at <https://people.math.ethz.ch/~vgiri01/spinorgaps.pdf> (2024).
- [2] N. Alon, Large sets in finite fields are sumsets. *J. Number Theory*, **126** (2007), 110–118.
- [3] N. Alon and J. H. Spencer, *The Probabilistic Method, Fourth Edition*. Wiley, 2016.
- [4] J. Barajas and O. Serra, The lonely runner with seven runners. *Electr. J. Combin.*, **15.1** (2008), #R48.
- [5] M. Beck, S. Hoşten, and M. Schymura, Lonely Runner Polyhedra. *Integers*, **19** (2019), #A29.
- [6] M. Beck and M. Schymura, Deep lattice points in zonotopes, lonely runners, and lonely rabbits. Preprint arXiv:2301.12182v1 (2023).
- [7] V. Bergelson, Weakly mixing PET. *Ergodic Theory Dynam. Systems*, **7.3** (1987), 337–349.
- [8] V. Bergelson and A. Best, On the polynomial Szemerédi theorem in finite commutative rings. *Preprint* arXiv:arXiv:2303.17928v1 (2023).
- [9] Bergelson and A. Leibman, Polynomial extensions of van der Waerden’s and Szemerédi’s theorems. *J. Amer. Math. Soc.*, **9** (1996), 725–753.
- [10] U. Betke and J. M. Wills, Untere Schranken für zwei diophantische Approximations-Funktionen. *Monatsh. Math.*, **76.3** (1972), 214–217.
- [11] A. Bhardwaj, V. Narayanan, and H. Venkataraman, A few more Lonely Runners. Preprint arXiv:2211.08749v2 (2022).
- [12] W. Bienia, L. Goddyn, P. Gvozdjak, A. Sebő, and M. Tarsi, Flows, view obstructions, and the lonely runner. *J. Combin. Theory, Ser. B*, **72.1** (1998), 1–9.
- [13] A. Biswas and J. Saha, Minimal additive complements in finitely generated abelian groups. *Ramanujan J.*, **57** (2022), 215–238.

- [14] A. Biswas and J. Saha, On additive co-minimal pairs. *J. Number Theory*, **223** (2021), 350–370.
- [15] A. Biswas and J. Saha, On minimal complements in groups. *Ramanujan J.*, **55** (2021), 823–847.
- [16] J.-P. Bode and H. Harborth, Directed paths of diagonals within polygons. *Discrete Math.*, **299** (2005), 3–10.
- [17] T. Bohman, R. Holzman, and D. Kleitman, Six lonely runners. *Electr. J. Combin.*, **8.2** (2001), # R3.
- [18] T. Bohman and F. Peng, Coprime mappings and lonely runners. *Mathematika*, **68.3** (2022), 784–804.
- [19] J. Bourgain, On arithmetic progressions in sums of sets of integers. In *A Tribute to Paul Erdős* (Cambridge University Press, 1990), 105–109.
- [20] J. Bourgain and M.-C. Chang, Nonlinear Roth type theorems in finite fields. *Israel J. Math.*, **221** (2017), 853–867.
- [21] Matija Bucić, Bryce Frederickson, Alp Müyesser, Alexey Pokrovskiy, and Liana Yepremyan, Towards Graham’s rearrangement conjecture via rainbow paths. *Preprint* arXiv:2503.01825v1 (2025).
- [22] A. Burcroff and N. Luntzlara, Sets arising as minimal additive complements in the integers. *J. Number Theory*, **247** (2023), 15–34.
- [23] J. Chapman and S. Chow, Generalised Rado and Roth criteria. *Ann. Sc. Norm. Super. Pisa Cl. Sci.* (2023), 38 pp.
- [24] Y. Chen and Q. Yang, On a problem of Nathanson related to minimal additive complements. *SIAM J. Discrete Math.*, **26.4** (2012), 1532–1536.
- [25] Q. Chu, N. Frantzikinakis, and B. Host, Ergodic averages of commuting transformations with distinct degree polynomial iterates. *Proc. Lond. Math. Soc.*, **102** (2011), 801–842.
- [26] S. Costa and S. Della Fiore, Weak Freiman isomorphisms and sequencings of small sets. *Preprint* arXiv:2407.15785 (2024).
- [27] S. Costa and M. A. Pellegrini, Some new results about a conjecture by Brian Alspach. *Arch. Math. (Basel)*, **115** (2020), 479–488.
- [28] T. W. Cusick, View-obstruction problems. *Aequationes Math.*, **9.2** (1973), 165–170.
- [29] T.W. Cusick and M.E. Flahive, *The Markoff and Largange Spectra*. AMS: Mathematical Surveys and Monographs, **30**, 1989.

- [30] T. W. Cusick and C. Pomerance, View-obstruction problems, III. *J. Number Theory*, **19.2** (1984), 131–139.
- [31] D. Dong, X. Li, and W. Sawin, Improved estimates for polynomial Roth type theorems in finite fields. *J. Anal. Math.*, **141** (2020), 689–705.
- [32] S. Donoso, A. Ferré-Moragues, A. Koutsogiannis, and W. Sun, Decomposition of multicorrelation sequences and joint ergodicity. *Ergodic Theory Dynam. Systems*, **44.2** (2024), 423–480.
- [33] S. Donoso, A. Koutsogiannis, B. Kuca, W. Sun, and K. Tsinas, Seminorm estimates and joint ergodicity for pairwise independent Hardy sequences. *Preprint* arXiv:2410.15130v1 (2024).
- [34] S. Donoso, A. Koutsogiannis, and W. Sun, Seminorms for multiple averages along polynomials and applications to joint ergodicity. *J. Anal. Math.*, **146** (2022), 1–64.
- [35] M. Einsiedler, Applications of measure rigidity of diagonal actions. *Proceedings of the International Congress of Mathematicians* (2010)
- [36] M. Einsiedler, E. Lindenstrauss, P. Michel and A. Venkatesh, Distribution of periodic torus orbits on homogeneous spaces. *Duke Math. J.*, **148(1)** (2009), 119–174.
- [37] M. Einsiedler, E. Lindenstrauss, P. Michel and A. Venkatesh, Distribution of periodic torus orbits and Duke’s theorem for cubic fields. *Annals of Math.*, **173** (2011), 815–885.
- [38] P. Erdős and R. L. Graham, *Old and new problems and results in combinatorial number theory*. L’Enseignement mathématique (1980), Université de Genève.
- [39] H. T. Fan and A. Sun, Amending and continuing the Lonely Runner Spectrum Conjecture. Preprint ArXiv:2306.10417v1 (2023).
- [40] H. Furstenberg, Ergodic behavior of diagonal measures and a theorem of Szemerédi on arithmetic progressions. *J. Anal. Math.*, **31.1** (1977), 204–256.
- [41] S. M. Girvin and K. Yang, *Modern Condensed Matter Physics*. Cambridge University Press, 2019.
- [42] W. T. Gowers, A new proof of Szemerédi’s theorem. *Geom. Funct. Anal.*, **11.3** (2001), 465–588.
- [43] W. T. Gowers, A new proof of Szemerédi’s theorem for arithmetic progressions of length four. *Geom. Funct. Anal.*, **8.3** (1998), 529–551.
- [44] W. T. Gowers, Arithmetic progressions in sparse sets. In *Current developments in mathematics*, Int. Press (2000), 149–196.

- [45] R. Graham, On sums of integers taken from a fixed sequence. In *Proceedings, Washington State University Conference on Number Theory* (1971), 22–40.
- [46] B. Green, An argument of Shkredov in the finite field setting. Available at <https://people.maths.ox.ac.uk/greenbj/papers/corners.pdf> (2025).
- [47] B. Green, Finite field models in additive combinatorics. In *Surveys in combinatorics, 2005*, London Math. Soc. Lecture Note Ser. Vol. 237 (2005), 1–27.
- [48] B. Green, On arithmetic structures in dense sets of integers. *Duke Math. J.*, **114** (2002), 215–238.
- [49] B. Green and M. Sawhney, Primes of the form $p^2 + 4q^2$. *Preprint* arXiv:2410.04189v1 (2024).
- [50] B. Green and T. Tao, Linear equations in primes. *Ann. of Math.*, **171.3** (2010), 1753–1850.
- [51] B. Green and T. Tao, The quantitative behaviour of polynomial orbits on nilmanifolds. *Ann. of Math.*, **175** (2012), 465–540.
- [52] L. Habsieger and I. Ruzsa, Additive completion and disjoint translations. *Acta Math. Hungar.*, **69.4** (1995), 273–280.
- [53] A. Hajnal and E. Szemerédi, Proof of a conjecture of P. Erdős. In *Combinatorial theory and its applications, II (Proc. Colloq., Balatonfüred, 1969)* (1970), 601–623.
- [54] R. Han, M. T. Lacey, and F. Yang, A polynomial Roth theorem for corners in finite fields. *Mathematika*, **67.4** (2021), 885–896.
- [55] M. Henze and R. D. Malikiosis, On the covering radius of lattice zonotopes and its relation to view-obstructions and the lonely runner conjecture. *Aequationes Math.*, **91.2** (2017), 331–352.
- [56] J. Hicks, M. Ollis, and J. Schmitt, Distinct partial sums in cyclic groups: polynomial method and constructive approaches. *J. Combin. Des.*, **27.6** (2019), 369–385.
- [57] I.S. Honkala, Modified bounds for covering codes. *IEEE Trans. Inform. Theory*, **37.2** (1991), 351–365.
- [58] B. Host and B. Kra, *Nilpotent structures in ergodic theory*, AMS (2018).
- [59] M. Jaber, Y. Liu, S. Lovett, A. Ostuni, and M. Sawhney, Quasipolynomial bounds for the corners theorem. *Preprint* arXiv:2504.07006v1 (2025).
- [60] T. Kato, *Perturbation theory for linear operators*. Springer: Classics in Mathematics, **132**, 1995.

- [61] Z. Kelley and R. Meka, Strong bounds for 3-progressions. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)* (2023), 933–973.
- [62] S. Kiss, C. Sándor, and Q. Yang, On minimal additive complements of integers. *J. Combin. Theory Ser. A*, **162** (2019), 344–353.
- [63] A.J. Kollár and P. Sarnak, Gap sets for the spectra of cubic graphs. *Comm. Amer. Math. Soc.*, **1** (2021), 1–38.
- [64] D. Koz, M. Mirek, S. Peluse, and J. Wright, The multilinear circle method and a question of Bergelson. *Preprint* arXiv:2411.09478v2 (2024).
- [65] N. Kravitz, Barely lonely runners and very lonely runners: a refined approach to the Lonely Runner Problem. *Combin. Theory*, **1** (2021), #17.
- [66] N. Kravitz, Rearranging small sets for distinct partial sums. *Integers*, **24** (2024), #A113.
- [67] B. Kuca, Further bounds in the polynomial Szemerédi theorem over finite fields. *Acta Arith.*, **198** (2021), 77–108.
- [68] B. Kuca, Multidimensional polynomial patterns over finite fields: Bounds, counting estimates and Gowers norm control. *Adv. Math.*, **448** (2023), #109700.
- [69] B. Kuca, Multidimensional polynomial Szemerédi theorem in finite fields for polynomials of distinct degrees. *Israel J. Math.*, **259.2** (2022), 589–620.
- [70] A. Kwon, A note on minimal additive complements of integers. *Discrete Math.*, **342.7** (2019), 1912–1918.
- [71] M. T. Lacey and W. McClain, On an argument of Shkredov on two-dimensional corners. *Online J. Anal. Comb.*, **2** (2007).
- [72] J. Leng, A quantitative bound for Szemerédi’s theorem for a complexity one polynomial progression over $\mathbb{Z}/N\mathbb{Z}$. *Discrete Anal.*, **5** (2024), 35pp.
- [73] J. Leng, Efficient Equidistribution of Nilsequences. *Preprint* arXiv:2312.10772v1 (2023).
- [74] J. Leng, Efficient Equidistribution of Periodic Nilsequences and Applications. *Preprint* arXiv:2306.13820v1 (2023).
- [75] J. Leng, A. Sah, and M. Sawhney, Quasipolynomial bounds on the inverse theorem for the Gowers $U^{s+1}[N]$ -norm. *Preprint* arXiv:2402.17994v1 (2024).

- [76] R.D. Malikiosis, F. Santos, and M. Schymura, Linearly-exponential checking is enough for the Lonely Runner Conjecture and some of its variants. Preprint ArXiv:2411.06903v1 (2024).
- [77] F. Manners, True complexity and iterated Cauchy-Schwarz. *Preprint* arXiv:2109.05731v1 (2021).
- [78] A. Müyesser and A. Pokrovskiy, A random Hall–Paige conjecture. *Inventiones Math.*, to appear (arXiv:2204.09666v1, 2022).
- [79] M. Nathanson, Problems in additive number theory, IV: Nets in groups and shortest length g -adic representations. *Int. J. Number Theory*, **7.8** (2011), 1999–2017.
- [80] M. Ortega and S. Prendiville, Extremal Sidon sets are Fourier uniform, with applications to partition regularity. *J. Théor. Nombres Bordeaux*, **35.1** (2023), 115–134.
- [81] S. Peluse, Bounds for sets with no polynomial progressions. *Forum Math. Pi*, **8:e16** (2020), 1–55.
- [82] S. Peluse, Finite field models in arithmetic combinatorics – twenty years on. *Preprint* arXiv:2312.08100v1 (2023).
- [83] S. Peluse, On the polynomial Szemerédi theorem in finite fields. *Duke Math. J.*, **168.5** (2019), 749–774.
- [84] S. Peluse, Three-term polynomial progressions in subsets of finite fields. *Israel J. Math.*, **228** (2018), 379–405.
- [85] S. Peluse and S. Prendiville, A polylogarithmic bound in the nonlinear Roth theorem. *Int. Math. Res. Not.*, **2022.8** (2022), 5658–5684.
- [86] S. Peluse and S. Prendiville, Quantitative bounds in the non-linear Roth theorem. *Inv. Math.*, **238.3** (2024), 865–903.
- [87] S. Peluse, S. Prendiville, and X. Shao, Bounds in a popular multidimensional nonlinear Roth theorem. *Preprint* arXiv:2407.08338v1 (2024).
- [88] S. Peluse, A. Sah, and M. Sawhney, Effective bounds for Roth’s theorem with shifted square common difference. *Preprint* arXiv:2309.08359v1 (2023).
- [89] C. Pomerance, Coprime matchings. *Integers*, **22** (2022), #A83.
- [90] S. Prendiville, Quantitative bounds in the polynomial Szemerédi theorem: the homogeneous case. *Discrete Anal.*, **5** (2017), 34 pp.

- [91] S. Prendiville, The inverse theorem for the nonlinear Roth configuration: an exposition. *Preprint* arXiv:2003.04121v1 (2020).
- [92] J. Renault, View-obstruction: a shorter proof for 6 lonely runners. *Discrete Math.*, **287.1–3** (2004), 93–101.
- [93] K. F. Roth, On certain sets of integers. *J. Lond. Math. Soc.*, **28.1** (1953), 104–109.
- [94] I. Ruzsa, An infinite Sidon sequence. *J. Number Theory*, **68** (1998), 63–71.
- [95] A. Sah and M. Sawhney, Enumerating coprime permutations. *Mathematika*, **68.4** (2022), 1120–1134.
- [96] A. Sárközy, On difference sets of sequences of integers. I. *Acta Math. Hungar.*, **31.1–2** (1978), 125–149.
- [97] A. Sárközy, On difference sets of sequences of integers. III. *Acta Math. Hungar.*, **31.3–4** (1978), 355–386.
- [98] P. Sarnak, Special values of Selberg’s zeta function. *Number theory trace formulas and discrete groups*, ed. E. Bombieri and D. Goldfeld (1989), 457–465.
- [99] W. Sawin, comment on the post “Ordering subsets of the cyclic group to give distinct partial sums”. MathOverflow (2015), <https://mathoverflow.net/q/202857>.
- [100] W. Schmidt, Asymptotic formulae for point lattices of bounded determinant and subspaces of bounded height. *Duke Math. J.*, **35.2** (1968), 327–339.
- [101] C. Series, The geometry of Markoff numbers. *Math. Intelligencer*, **7.3** (1985), 20–29.
- [102] I. D. Shkredov, On a generalization of Szemerédi’s theorem. *Proc. Lond. Math. Soc.*, **93.3** (2006), 723–760.
- [103] I. D. Shkredov, On a problem of Gowers. *Izv. Math.*, **70.2** (2006), 385–425.
- [104] C. L. Siegel and K. Chandrasekharan, *Lectures on the Geometry of Numbers*. Springer, 1989.
- [105] E. Stein, *Harmonic Analysis: Real-variable Methods, Orthogonality and Oscillatory Integrals*. Princeton University Press (1993).
- [106] E. Szemerédi, On sets of integers containing k elements in arithmetic progression. *Acta Arith.*, **27.1** (1975), 199–245.
- [107] T. Tao, *Higher order Fourier analysis*. AMS, (2012).
- [108] T. Tao, Some remarks on the Lonely Runner Conjecture. *Contrib. Discrete Math.*, **13.2** (2018), 1–31.

- [109] T. Tao and J. Teräväinen, Quantitative bounds for Gowers uniformity of the Möbius and von Mangoldt functions. *J. Eur. Math. Soc.*, **27.4** (2023), 1321–1384.
- [110] T. Tao and T. Ziegler, Concatenation theorems for anti-Gowers uniform functions and Host- Kra characteristic factors. *Discrete Anal.*, **13** (2016), 61 pp.
- [111] K. Tsinas, Joint ergodicity of Hardy field sequences. *Trans. Amer. Math. Soc.*, **376** (2023), 3191–3263.
- [112] P. Varnavides, On certain sets of positive density. *J. Lond. Math. Soc.*, **s1-34.3** (1959), 358–360.
- [113] J. M. Wills, Zwei Sätze über inhomogene diophantische Approximation von Irrationalzahlen. *Monatsh. Math.*, **71.3** (1967), 263–269.
- [114] T. Wooley, *Analytic number theory: introduction to the circle method and its application, 2023*, lecture notes available at <https://www.math.purdue.edu/~twooley/2023ant/2023antnotes.pdf> (2025).
- [115] T. Wooley, Nested efficient congruencing and relatives of Vinogradov’s mean value theorem. *Proc. Lond. Math. Soc.*, **118.4** (2019), 942–1016.