Topics in additive combinatorics

Frederick Manners

Balliol College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity term 2016

This thesis is dedicated to my family,
for their constant support.

Acknowledgements

First and foremost, the author is grateful to his supervisor Ben Green for his advice, support and collaboration.

The author was funded during this degree by the Engineering and Physical Sciences Research Council, and received additional support from the University of Cambridge and Trinity College, Cambridge (2012–3), and the University of Oxford and Balliol College, Oxford (2013–6).

The author also owes a great debt to all those others who have collaborated on, discussed or otherwise stimulated this and other research. In particular, but not exclusively, this includes: Sean Eberhard, Rudi Mrazović, Péter Varjú, Yonatan Gutman, Jonathan Lee, Jakub Konieczny, Przemysław Mazur, Aled Walker, Sofia Lindqvist, Fernando Shao, Olaf Sisask, Thomas Bloom, Bryna Kra, Elon Lindenstrauss and Terence Tao.

He would also like to thank the following institutions for providing a stimulating research environment, and for their hospitality, at various stages during the preparation of this D.Phil.: the Mathematisches Forschungsinstitut Oberwolfach, the Alfréd Rényi Institute of Mathematics, the Institute for Mathematics and its Applications in Minneapolis, the Centre de recherches mathématiques in Montreal, the Banff International Research Station; the University of Bristol, the London School of Economics, Queen Mary University of London, the Univeristy of York, the Univeristy of Cambridge and the Clay Mathematics Institute.

Regarding Chapter 2, the author is grateful to Przemysław Mazur for pointing out the generalization alluded to in Remark 2.2.2. Related to Chapter 5, he would like to thank Robin Pemantle for a discussion in connection with the method used to prove Theorem 5.4.1, and Fernando Shao for useful conversations. Similarly regarding Chapter 6, the author is particularly grateful to Jonathan Lee for extensive discussions of these ideas; and to Sean Eberhard, for pointing out the improvement to Lemma 6.5.3.

Abstract

We present a number of results in both pure and applied additive combinatorics.

Contents

Chapter 1

Introduction

This thesis comprises a collection of results concerned variously with the tools of additive combinatorics that have been developed over recent years, or with applications of those tools to proving theorems in related areas.

By these tools, we mean anything from very classical Fourier analysis and analytic number theory; to the classical results of additive combinatorics, concerning sumsets, small doubling, Freĭman's theorem, additive energy, and so on; to more recent structural developments, notably the inverse theory for the Gowers U^k-norms, pioneered by Gowers [30], Green and Tao [33], and Green, Tao and Ziegler [35, 36], and others, sometimes referred to informally as *higher-order Fourier analysis*.

The first three chapters of this work are attempts at modest extensions to this state of knowledge.

In Chapter 2, we prove a weak analogue of Freĭman's theorem, with better bounds than can be achieved by other methods. More precisely: if a set X of integers has $|X + X| \leq K|X|$, then the Freĭman dimension lemma states that the Freĭman dimension of $\dim X$ of X satisfies (essentially) $\dim X \leq K - 1$. This is a reasonably weak conclusion that plays an early role in some proofs of the full Freĭman theorem.

The dimension lemma is best possible as stated. However, we obtain the following refinement: there is a subset $X' \subseteq X$ with $|X|/|X'| \leq \exp\left(O(\log^2 K)\right)$ such that $\dim X' = O(\log K)$.

Chapter 3 is concerned with the study of *approximate polynomials*. There is a close connection (although not one that has been written down) between the inverse theory of the Gowers U^{s+2}-norm, and the problem of classifying *1% polynomial* functions of degree s on finite groups. Roughly, if a polynomial of degree s is a function whose $(s+1)$-th derivatives vanish, a 1% polynomial is a function whose derivatives vanish a positive proportion of the time.

Here, we consider a variant of this notion, namely functions whose $(s+1)$-th derivatives take few distinct values, and obtain a strong structure theorem under this stronger hypothesis.

Finally, Chapter 4 derives a variant of the usual Green–Tao–Ziegler inverse theorem for the U^k-norms, working over cyclic groups $\mathbb{Z}/N\mathbb{Z}$ rather than the interval $\{1,\dots,N\}$, from the original result. This result had an application in work of Candela and Sisask [11].

By applications, we mean results which bring some of this machinery to bear on questions not directly concerned with the machinery itself.

In this category, Chapter 5 applies Fourier and circle method techniques to the problem of counting triples of bijections $\pi_1,\pi_2,\pi_3\colon \{1,\dots,n\} \to \mathbb{Z}/n\mathbb{Z}$ such that $\pi_1+\pi_2=\pi_3$, which, aside from any intrinsic interest, is closely connected to the open so-called "Queens problem".

Finally, Chapter 6 answers a question (the "Pyjama problem") in elementary metric geometry, using input from Fourier theory and topological dynamics.

These chapters are to some degree orthogonal, and each may be read independently of the others. However, they are far from being unrelated. Indeed, we have already stated that there are close connections between the study of approximate polynomials (Chapter 3) and the inverse theory of the Gowers norms (Chapter 4). Furthermore, there is a better-understood connection between the inverse theory of the U^3-norm and inverse sumset questions related to Freĭman's theorem, meaning Chapter 2 is in the same spirit as these other two.

As another example, the main obstruction to applying the approach of Chapter 5 to the full Queens problem, is the somewhat incomplete status of the fundamentals of higher-order Fourier analysis in general and the inverse theory of the U^3-norm in particular.

Indeed, it is our intention to stress throughout that all of these pieces of work are rooted in the same circle of ideas.

Since each chapter contains a substantial introduction to the problem it discusses, we will defer to those and not give more detailed outlines of each chapter at this point.

Some of this work has been produced in collaboration. In particular, Chapter 5 is based on the paper [21] written together with Sean Eberhard and Rudi Mrazović. In other cases where no explicit co-authorship took place, the author has benefited from innumerable fruitful conversations with a large number of people, whom he has attempted to acknowledge above or within each chapter.

Chapters 2, 4 and 6 are derived from the author's previous papers or preprints [43], [44], [45] respectively.

1.1 Notation and conventions

Any notation particular to each chapter will be introduced towards the start of the chapter in question. We mention here those conventions that have global scope.

As is standard, for quantities X, Y we write $X = O(Y)$ to mean $|X| \leq CY$ for some constant $C > 0$, and write $X = O_{r_1,\ldots,r_k}(Y)$ to mean $|X| \leq C(r_1, \ldots, r_k)Y$ for some quantity $C(r_1, \ldots, r_k) > 0$ depending only on the parameters $r_1, \ldots, r_k$. If we write $X \ll_{r_1,\ldots,r_k} Y$ this is a synonym for $X = O_{r_1,\ldots,r_k}(Y)$.

If f and g are functions, we write $f = o_{r_1,\ldots,r_k;x\to x_0}(g)$ to mean that $f(x)/g(x) \to 0$ as $x \to x_0$, where the rate of convergence depends only on the quantities $r_1, \ldots, r_k$. The expression $f = o_{n\to\infty}(g)$ is usually abbreviated to $f = o(g)$, etc..

For a finite set S and function f, we will denote the average of f on the set S by

$$\mathbb{E}_{s\in S} f(s) = \frac{1}{|S|} \sum_{s\in S} f(s) \ .$$

Unless otherwise specified, finite groups G are equipped with uniform probability measure, and their dual groups $\widehat{G}$ are equipped with counting measure. This convention extends to the definition of convolutions, Fourier transforms, L^p norms and so on.

We also use the standard notation $e(x) = e^{2\pi i x}$.

Chapter 2

Finding a low-dimensional piece of a set

In this chapter, we show that a finite set of integers $A \subseteq \mathbb{Z}$ with $|A+A| \leq K|A|$ contains a large piece $X \subseteq A$ with Freĭman dimension $O(\log K)$, where large means $|A|/|X| \ll \exp(O(\log^2 K))$. This can be thought of as a major quantitative improvement on Freĭman's dimension lemma; or as a weak Freĭman–Ruzsa theorem with almost polynomial bounds.

The methods used, centered around an increment strategy for additive energy, differ from the usual tools in this area and may have further potential.

Most of our argument takes place over $\mathbb{F}_2^n$, which is itself curious. There is a possibility that the above bounds could be improved, assuming sufficiently strong results in the spirit of the Polynomial Freĭman–Ruzsa Conjecture over finite fields.

This chapter is based on the preprint [43].

2.1 Introduction

2.1.1 Motivation and statements

Our primary motivation comes from considering Freĭman-type theorems for subsets of the integers, of the following sort.

Theorem 2.1.1. *Suppose $A \subseteq \mathbb{Z}$ and $|A+A| \leq K|A|$. Then there exists a generalized arithmetic progression[1] P of rank at most $C_1(K)$, such that $|P| \leq C_2(K)|A|$ and $|A \cap P| \geq |A|/C_3(K)$, for some quantities C_1, C_2, C_3 depending only on K.*

If one insists that $A \subseteq P$ then one cannot do better than $C_1(K) = K^{O(1)}$ and $C_2(K), C_3(K) = \exp\left(K^{O(1)}\right)$; such a result is due to Freĭman [25, 26]. If, as stated, we allow just a large subset of A to be contained in P, the current overall best bounds

[1] By a *generalized arithmetic progression of rank r*, we mean a set of the form $\{a_0 + a_1 n_1 + \cdots + a_r n_r : 0 \leq n_i < N_i\}$ where $a_0, \ldots, a_r$ are integers and $N_1, \ldots, N_R$ are positive integers.

in Theorem 2.1.1 are due to Sanders [50], and are of the form $C_1(K) = O(\log^4 K)$, $C_2(K), C_3(K) = \exp(O(\log^4 K))$.

It is also possible to control the rank $C_1(K)$ optimally by $O(\log K)$, at the expense of poorer bounds on C_2 and C_3; see e.g. [4, 5, 25, 31].

It has been widely conjectured that one can take $C_1(K) = O(\log K)$ while also insisting that $\log(C_2(K))$ and $\log(C_3(K))$ are both $O\big(\log^{1+o(1)} K\big)$ (and indeed sharper conjectures could be made).

This chapter has nothing to add to the state of the art in Theorem 2.1.1. Instead, we consider a model problem that asks for significantly less structural information than Theorem 2.1.1. To state this, we briefly recall the notion of Freĭman dimension.

Definition 2.1.2. Let G, G' be abelian groups and $A \subseteq G$, $A' \subseteq G'$ finite subsets. We say a function $\phi\colon A \to A'$ is a *Freĭman homomorphism* if $\phi(x) + \phi(w) = \phi(y) + \phi(z)$ whenever $x, y, z, w \in A$ and $x + w = y + z$. We say it is a *Freĭman isomorphism* if it has an inverse that is a Freĭman homomorphism.

Now let $A \subseteq \mathbb{Z}$. The *Freĭman dimension* of A, denoted $\dim(A)$, is the largest positive integer d such there exists $A' \subseteq \mathbb{Z}^d$ which affinely generates $\mathbb{Z}^d$ (that is, A' is not contained in any proper affine subspace of $\mathbb{Z}^d$), such that A is Freĭman isomorphic to A'.

A key classical result about Freĭman dimension, due to Freĭman [26] (or see e.g. [59, Lemma 5.13]), is the following.

Theorem 2.1.3 (Freĭman dimension lemma)**.** *Let $A \subseteq \mathbb{Z}$ be such that $|A+A| \leq K|A|$. Then* $\dim(A) \leq K - 1 + O(K^2/|A|)$.

In fact a slightly sharper result is proven, but the details of this need not concern us currently.

There is at least some connection between Freĭman dimension and bounds in Theorem 2.1.1. Specifically, Chang's proof of Freĭman's theorem [14] shows that one may take the rank of the progression P to be *precisely* the Freĭman dimension $\dim(A)$; the same bound is obtained by the proof in [59, Chapter 5.6].

The bound in Theorem 2.1.3 is essentially sharp: for instance, consider a set such as $\{a + 100^b N\colon 1 \leq a \leq N, 1 \leq b \leq k\}$ which has doubling about $(k+1)$ and Freĭman dimension about k. However, examples with very large Freĭman dimension appear all to look like a union of a few structured pieces as in this case, each of which individually has much smaller Freĭman dimension. So, it is reasonable to expect a

much stronger bound to hold if we are allowed to pass to a large *subset* of A. Note this is entirely analogous to the situation in Theorem 2.1.1.

Hence, the following is – at least in some respects – a model problem on the road towards polynomial bounds in Theorem 2.1.1.

Conjecture 2.1.4. *Suppose $A \subseteq \mathbb{Z}$ is a finite set and $|A+A| \leq K|A|$. Then we can find a subset $X \subseteq A$ with $|A|/|X| \ll K^{O(1)}$ and $\dim(X) = O(\log K)$.*

This is equivalent to what is called the "Weak Polynomial Freĭman–Ruzsa Conjecture" in [15]. It is a consequence of any sufficiently strong analogue of Theorem 2.1.1 over $\mathbb{Z}^d$ (for all d).

In this chapter, we fail to prove Conjecture 2.1.4, but obtain the following near miss as our main result.

Theorem 2.1.5. *Let $A \subseteq \mathbb{Z}$ be a finite set with $E(A) \geq |A|^3/K$. Then we can find a subset $X \subseteq A$ such that $\dim(X) = O(\log K)$ and $|A|/|X| \ll \exp(O(\log^2 K))$.*

Note this is stated in terms of the additive energy

$$E(A) = \#\{x, y, z, w \in A \colon x - y = z - w\}$$

of the set rather than its doubling constant. Since a set with small doubling has large energy, the energy formulation is at least as strong. In fact two formulations are essentially equivalent, by invoking the Balog–Szemerédi–Gowers theorem. However, the proof of Theorem 2.1.5 does not use this result, and works with energy directly.

As we have stated above, a large part of the interest in Theorem 2.1.5 from our perspective is contained in some of the techniques that are used in the proof, which appear to be novel and not terribly closely related to the existing literature in this area.[2] In a sense, Theorem 2.1.5 could be thought of as an application of these techniques rather than the primary motivation behind them.

In particular, the tools used are essentially disjoint from those appearing in Sanders' work [50]. Also, we spend most of our time working over the group $\mathbb{F}_2^n$, which is unconventional given the result itself concerns subsets of $\mathbb{Z}$.

Although it is not our real goal, we state one application of our main result. By combining Theorem 2.1.5 with the argument in [59, Chapter 5.6] used as a black box, we obtain the following version of Theorem 2.1.1 as an easy consequence.

[2]We have since been informed [Ben Green, personal communication] of previous unpublished work of Gowers which includes some closely related ideas.

Corollary 2.1.6. *In Theorem 2.1.1, we may take* $C_1(K) = O(\log K)$ *and also* $C_2(K), C_3(K) \ll \exp(\exp(O(\log^2 K)))$.

The bound on C_2, C_3 is not competitive compared to e.g. [31] (although there is perhaps scope for optimizing this argument); we mention this corollary only because our derivation differs substantially from those arguments.

2.1.2 Outline of the proof

Our general strategy is one of *additive energy increment*. That is, we argue that if $\dim(A)$ is substantially larger than $\log K$, we can find a large subset $A_1 \subseteq A$ whose (normalized) additive energy is noticeably larger than that of A. Since normalized energy (which we have not yet defined) is bounded above by 1, this process must terminate after relatively few steps.

To achieve this, we follow the general philosophy of [31] or [52], which both argue along the following lines. We may assume that, after a Freĭman isomorphism, we have identified A with a subset of $\mathbb{Z}^d$.

- If A were in fact a measurable subset of $\mathbb{R}^d$, it would have to have doubling exponential in d (by the Brunn–Minkowski inequality) or additive energy exponentially small in d (by the sharp form of Young's inequality due to Beckner [1] and Brascamp and Lieb [9]).
- Since it does not, we deduce that A is not properly d-dimensional in some sense.
- We use this information to make progress by passing to a large subset of A.

In place of the Brunn–Minkowski or sharp Young inequalities, at the heart of our argument is a variant of the sharp Young inequality for "U^3 energy" (which we will define). This has the advantage of having a dramatically simpler proof than the sharp Young inequality itself. We can then transfer this argument from $\mathbb{R}^d$ to $\mathbb{Z}^d$, and while the same conclusion cannot always hold in that setting, if it fails, the simplicity of the original proof translates into very strong structural information about our set.

Specifically, we will be able to show that if $A \subseteq \mathbb{Z}^d$ has normalized U^3 energy much bigger than 2^{-d}, then the set must be very unevenly distributed modulo 2.

We then seek to use this knowledge to obtain an energy increment. The analysis required for this takes place largely over the finite vector space $\mathbb{F}_2^d$. This rather unexpected feature suggests the intriguing possibility that strong conjectures over $\mathbb{F}_2^n$ (in the spirit of the Polynomial Freĭman–Ruzsa conjecture over $\mathbb{F}_2^n$) could imply strong results over the integers, and specifically further strengthenings of Theorem 2.1.5.

2.1.3 Layout of the chapter

In Section 2.2, we introduce the notion of U^3 energy, and explore the variant of the sharp Young inequality described above. We also prove a result stating that if $A \subseteq \mathbb{Z}^d$ fails to have small energy then it has a very uneven distribution modulo 2.

In Section 2.3, we deal with the additive energy increment argument and complete the proof of Theorem 2.1.5.

At some point, we will need a useful inequality relating U^3 energy to usual additive energy, due to Shkredov [53, Proposition 35]. For interest and completeness, we give a self-contained alternative proof using entropy in Section 2.4. This discussion is more or less orthogonal to the rest of the chapter.

Finally, in Section 2.5, we describe – for general interest only – a closely related result to the material of Section 2.2, showing that sets of small doubling in $\mathbb{Z}^d$ occupy few residue classes modulo 2. Although we do not make use of this result in the main argument, we include it mainly because its proof is very different to the methods of Section 2.2, and it is curious that such dissimilar arguments should obtain such similar and yet – to the author at least – unfamiliar conclusions.

2.1.4 Notation

In this chapter, we have frequently adopted the slightly non-standard habit of using integral signs to denote sums or averages over discrete or even finite groups. Hence for instance given a function $f\colon \mathbb{Z}/N\mathbb{Z} \to \mathbb{R}$ we might write

$$\|f\|_2 = \left(\int_x |f(x)|^2\right)^{1/2}$$

even though the domain is finite. This allows us to be agnostic in our notation concerning the choice of normalization of the measure (e.g. counting measure or probability measure); i.e. the above could denote either a sum or an average, depending on context. Often, any *consistent* choice will be valid, in that we frequently manipulate quantities or equalities that are unaffected by the choice of normalization.

On a few occasions, though, it will be genuinely important that counting measure is used on the groups $\mathbb{Z}^d$ and $\mathbb{F}_2^d$, and that L^p- and U^k-norms on these groups are defined accordingly. Note this runs contrary to the usual convention of the document. We therefore declare that, for the rest of this chapter, counting measure is used throughout for these groups, and will flag up when this actually matters.

Various symbols, notably A', η and K, will retain their meanings throughout large parts of the argument. The original definitions can be found in the discussion at the start of Section 2.3 (but before Section 2.3.1).

For a set X, we will sometimes denote its indicator function also by X rather than 1_X, to aid legibility.

2.2 Reduction modulo 2 in $\mathbb{Z}^d$ and U^3 energy

For reasons that are not easy to justify at this point, our arguments will necessarily be concerned not with the usual notion of additive energy

$$E(A) = \#\{x, y, z, w \in A \colon x - y = z - w\} = \|1_A\|_{U^2}^4$$

but with the much less standard notion of U^3 *energy*, which we denote

$$E_{U^3}(A) = \|1_A\|_{U^3}^8$$

where in both cases we assume counting measure in the definition of our U^k-norms. Concretely we therefore have

$$E_{U^3}(A) = \sum_{x,h_1,h_2,h_3 \in \mathbb{Z}} A(x)A(x+h_1)A(x+h_2)A(x+h_1+h_2)A(x+h_3) \\ A(x+h_1+h_3)A(x+h_2+h_3)A(x+h_1+h_2+h_3) \ .$$

As usual, there is a multilinear variant of the U^3 energy, which we denote by

$$\langle\!\langle A_1, \dots, A_8 \rangle\!\rangle_{U^3} = \sum_{x,h_1,h_2,h_3 \in \mathbb{Z}} A_1(x)A_2(x+h_1)A_3(x+h_2)A_4(x+h_1+h_2)A_5(x+h_3) \\ A_6(x+h_1+h_3)A_7(x+h_2+h_3)A_8(x+h_1+h_2+h_3)$$

for sets $A_1, \dots, A_8$, and which is the same as the U^3 Gowers inner product.

The reader may reasonably be concerned that the U^3-norm of a set might be a genuinely different quantity than its energy, or U^2-norm: certainly these are radically different notions when applied to complex-valued functions. However, it turns out that in the setting of indicator functions, or more generally non-negative real functions, these are comparable measures of additive structure.

More precisely, we state the following.

Proposition 2.2.1. *If Z is any abelian group and $f \colon Z \to \mathbb{R}_{\geq 0}$ a non-negative finitely supported function, then*

$$\|f\|_{U^3} \geq \|f\|_{U^2}^2 / \|f\|_1 \ .$$

In the other direction,

$$\|f\|_{U^3} \leq \|f\|_{U^2}^{1/2} \|f\|_4^{1/2} .$$

We remark that these inequalities are independent of the choice of measure on Z used to define these norms.

Proof. The second inequality is straightforward: in the expression

$$\|f\|_{U^3}^8 = \int_{x,y,h_1,h_2} f(x)f(x+h_1)f(x+h_2)f(x+h_1+h_2)f(y)f(y+h_1)f(y+h_2)f(y+h_1+h_2)$$

we apply the AM–GM inequality $abcd \leq (a^4 + b^4 + c^4 + d^4)/4$ pointwise to the last four terms.

The first is significantly less trivial. In the case that f is the indicator function of a set (which is all we need, and from which the general case could be deduced), this result is due to Shkredov [53, Proposition 35]; his proof uses a Loomis–Whitney-type result attributed to Bollobás and Thomason [6], and the tensor power trick.

In Section 2.4, for interest, we give a self-contained alternative (but morally similar) proof using entropy. □

Remark 2.2.2. A natural analogue of Proposition 2.2.1 for higher U^k norms does exist: in general for $k \geq 2$ and a non-negative function f, one can show using very similar methods that

$$\|f\|_{U^k}^{3k-2} \leq \|f\|_{U^{k+1}}^{2k-2} \|f\|_{U^{k-1}}^k ,$$

and further results can be obtained by chaining these together. The exponents are the unique choice that makes sense under scaling the measure. This generalization was pointed out to the author by Przemysław Mazur.

We recall that Young's inequality implies the bound $\|f\|_{U^2} \leq \|f\|_{4/3}$. We begin by recalling a version of this for the U^3-norm.

Proposition 2.2.3. *Let N be an odd prime and suppose $f\colon \mathbb{Z}/N\mathbb{Z} \to \mathbb{C}$. We have $\|f\|_{U^3} \leq \|f\|_2$.*

This is a special case of a general result for all U^k-norms; see [22]. The usual proof is iterated application of Young's inequality for convolutions. We note that in this case there is an even easier proof.

Proof. We consider

$$\|f\|_{U^3}^8 = \int_{x,h_1,h_2,h_3} (f(x)f(x+h_1+h_2)f(x+h_1+h_3)f(x+h_2+h_3)) \\ \left(\overline{f(x+h_1)f(x+h_2)f(x+h_3)f(x+h_1+h_2+h_3)}\right) .$$

and writing the bracketed terms as

$$F(x,h_1,h_2,h_3) = f(x)f(x+h_1+h_2)f(x+h_1+h_3)f(x+h_2+h_3)$$
$$G(x,h_1,h_2,h_3) = f(x+h_1)f(x+h_2)f(x+h_3)f(x+h_1+h_2+h_3)$$

we apply the Cauchy–Schwarz inequality to obtain

$$\begin{aligned}\|f\|_{U^3}^8 &\leq \left(\int_{x,h_1,h_2,h_3} |F(x,h_1,h_2,h_3)|^2\right)^{1/2} \left(\int_{x,h_1,h_2,h_3} |G(x,h_1,h_2,h_3)|^2\right)^{1/2} \\ &= \left(\int_{x,h_1,h_2,h_3} |f(x)|^2|f(x+h_1+h_2)|^2|f(x+h_1+h_3)|^2|f(x+h_2+h_3)|^2\right)^{1/2} \\ &\quad \left(\int_{x,h_1,h_2,h_3} |f(x+h_1)|^2|f(x+h_2)|^2|f(x+h_3)|^2|f(x+h_1+h_2+h_3)|^2\right)^{1/2} \\ &= \int_{x,h_1,h_2,h_3} |f(x)|^2|f(x+h_1+h_2)|^2|f(x+h_1+h_3)|^2|f(x+h_2+h_3)|^2\end{aligned}$$

where the last line follows by a straightforward change of variables. But by another change of variables, this integral is precisely

$$\left(\int_x |f(x)|^2\right)^4$$

as required. □

Over the domain $\mathbb{R}^d$ rather than $\mathbb{Z}/N\mathbb{Z}$, there is a sharp form of Young's inequality due to Beckner [1] and Brascamp–Lieb [9], which in particular implies that $\|f\|_{U^2} \leq C^d\|f\|_{4/3}$ for some explicit constant $C < 1$. Again, we consider an analogue for the U^3-norm.

Proposition 2.2.4. *Let $f\colon \mathbb{R}^d \to \mathbb{C}$ be a well-behaved function (say, continuous and compactly supported). We have $\|f\|_{U^3} \leq (1/2)^{d/8}\|f\|_2$.*

As before, this can be deduced by iterated application of the *sharp* Young inequality – again see [22] for such a proof. Alternatively, it is a special case of the more general Brascamp–Lieb inequality [9].

However, as remarked in their original paper [9], some special cases of the Brascamp–Lieb inequality have unusually easy proofs. It turns out this is one such instance, and that we have essentially already seen the proof.

Proof. We apply the Cauchy–Schwarz argument as above to deduce that

$$\|f\|_{U^3}^8 \leq \int_{x,h_1,h_2,h_3\in\mathbb{R}^d} |f(x)|^2|f(x+h_1+h_2)|^2|f(x+h_1+h_3)|^2|f(x+h_2+h_3)|^2 \ .$$

We now apply the linear change of variables

$$\begin{aligned} r &= x \\ s &= x+h_1+h_2 \\ t &= x+h_1+h_3 \\ u &= x+h_2+h_3 \end{aligned} \tag{2.2.1}$$

and note that the corresponding matrix

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 \\ 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 \end{pmatrix}^{\otimes d} \tag{2.2.2}$$

has determinant $(-2)^d$. It follows that

$$\|f\|_{U^3}^8 \leq 2^{-d}\|f\|_2^8$$

as required. □

Of course, we are ultimately interested not in functions on $\mathbb{R}^d$ but in subsets of $\mathbb{Z}^d$. Since the argument above is so elementary, it is straightforward to transfer it to this setting and see what we obtain.

Proposition 2.2.5. *If $f\colon \mathbb{Z}^d \to \mathbb{C}$ then*

$$\|f\|_{U^3}^8 \leq \int_{r,s,t,u} [r+s+t+u \equiv 0 \pmod 2] \ |f(r)|^2|f(s)|^2|f(t)|^2|f(u)|^2 \ .$$

Proof. We apply the same Cauchy–Schwarz argument as above, and again make the change of variables (2.2.1). This time, we observe that the image of the matrix (2.2.2) applied to $\mathbb{Z}^{4d}$ consists of precisely those tuples $(r,s,t,u) \in \mathbb{Z}^{4d}$ such that $r+s+t+u \in 2\mathbb{Z}^d$, as required. Indeed, one inclusion is clear, and equality of these two sets then follows from the determinant calculation above. □

The right hand side expression is a function only of what we shall define as

$$\begin{aligned} \eta\colon (\mathbb{Z}/2\mathbb{Z})^d &\to \mathbb{R}_{\geq 0} \\ \omega &\mapsto \left(\int_x [x \bmod 2 = \omega] \ |f(x)|^2\right)^{1/2} \end{aligned}$$

i.e. the distribution of the L^2 mass of f modulo 2. Moreover, we see that if η is a constant function, i.e., the mass of the original function f was equidistributed over residue classes modulo 2, then this expression is precisely $\|f\|_2^8/2^d$, just as we saw over $\mathbb{R}^d$.

In other words, applying this to $f = 1_A$ for some finite set $A \subseteq \mathbb{Z}^d$, we deduce the following.

Slogan 2.2.6. *If $A \subseteq \mathbb{Z}^d$ is a finite set and the energy $E_{U^3}(A)$ is much bigger than $2^{-d}|A|^4$, then A must be highly non-equidistributed modulo 2.*

Remark 2.2.7. One possible interpretation for this rather bizarre-looking statement is that if A were properly d-dimensional then its energy should be exponentially small in d; so the non-flat distribution modulo 2 is a kind of witness to the fact that A is not properly d-dimensional.

We now make this observation more precise.

First, we note that – since positive and negative signs are irrelevant modulo 2 – the expression appearing on the right hand side of Proposition 2.2.5 is precisely $\|\eta^2\|_{U^2}^4$ where η is defined as above. Here, we stress the importance of using counting measure on $(\mathbb{Z}/2\mathbb{Z})^d$ to define this norm.

Using the Young's inequality bound $\|h\|_{U^2} \leq \|h\|_{4/3}$, we obtain the following.

Proposition 2.2.8. *Let $f\colon \mathbb{Z}^d \to \mathbb{C}$ and let $\eta\colon (\mathbb{Z}/2\mathbb{Z})^d \to \mathbb{R}_{\geq 0}$ be defined as above. Then*

$$\|f\|_{U^3} \leq \|\eta\|_{8/3} \ .$$

Proof. We know that $\|f\|_{U^3}^8 \leq \|\eta^2\|_{U^2}^4$ and $\|\eta^2\|_{U^2} \leq \|\eta^2\|_{4/3} = \|\eta\|_{8/3}^2$. The result follows. □

It follows that if $\|f\|_{U^3}/\|f\|_2 = 1/K$ is large then so is $\|\eta\|_{8/3}/\|\eta\|_2$ (since $\|f\|_2 = \|\eta\|_2$). Since 8/3 is bigger than 2, this tells us that η must have a positive proportion of its L^2 mass on a just few points of $(\mathbb{Z}/N\mathbb{Z})^d$. Specifically, there is an $x \in (\mathbb{Z}/2\mathbb{Z})^d$ such that $\eta(x) \geq K^{-4}\|\eta\|_2$.

In the case that $f = 1_A$ is the indicator function of a set, a result of the same flavour stated in terms of doubling constants is provided (for interest only) as Proposition 2.5.1 below.

2.3 Energy increment arguments

The key quantity in the remainder of the argument is what we call the *normalized* U^3 *energy* of a set:

$$\mathscr{E}(A) = \big(E_{U^3}(A)/|A|^4\big)^{1/8} = \|A\|_{U^3}/\|A\|_2 \ .$$

Clearly the same quantity also makes sense for any (non-negative) finitely supported function f; i.e. we define

$$\mathscr{E}(f) = \|f\|_{U^3}/\|f\|_2$$

for any discrete abelian group Z and any finitely supported $f\colon Z \to \mathbb{R}_{\geq 0}$. We remark that this quantity does not depend on the normalization of Haar measure on Z.

Our detailed strategy to use the results of the previous section to prove Theorem 2.1.5 is as follows.

- Let $A \subseteq \mathbb{Z}$ be a finite set with large normalized U^3 energy, $\mathscr{E}(A) = 1/K$.
- Write $d = \dim(A)$. If $d = O(\log K)$, we are done. If not, we proceed as follows.
- By the definition of Freĭman dimension, there is a Freĭman isomorphic copy $A' \subseteq \mathbb{Z}^d$ of A, which affinely generates $\mathbb{Z}^d$. Note (see Proposition 2.3.11) that Freĭman isomorphism preserves quantities such as the U^3 energy of a set.
- We now inspect the function $\eta\colon (\mathbb{Z}/2\mathbb{Z})^d \to \mathbb{R}_{\geq 0}$ defined in the previous section for $f = 1_{A'}$, i.e., the L^2 distribution of A' modulo 2. By Proposition 2.2.8, this has a positive proportion of its L^2 mass concentrated on only a few values.
- We will now use this information to obtain an additive energy increment upon passing to a subset of A'. Specifically, we will find a subset $S \subseteq (\mathbb{Z}/2\mathbb{Z})^d$, define

  $$A'_S := \{a \in A'\colon a \bmod 2 \in S\} \subseteq A'$$

  and show that

  $$\mathscr{E}(A'_S) \geq \alpha\, \mathscr{E}(A')$$

  and

  $$|A'|/|A'_S| \leq \beta$$

  for suitable parameters $\alpha > 1$ and β.

  Finally, we pass to the corresponding set $A_S \subseteq A$, where the exact same conclusions hold. We then iterate this whole argument from the beginning on A_S.

The aim is of course to make the parameter α as large as possible while keeping β fairly small. We might hope to have $\alpha \geq 1.01$ and $\beta = O(K^{O(1)})$. In this case, since the normalized energy increases by a factor of α at each stage and is bounded above by 1, the algorithm can proceed for at most $O(\log K)$ iterations, meaning the size of the final set relative to A is about $K^{-O(\log K)} = \exp(-O(\log^2 K))$ as required.

Our procedure for finding this energy increment is slightly involved, and splits into several cases. We stress that there is no one obvious way to proceed in this phase of the argument, and there may well be more elegant ways to combine the various options than we have managed to find.

The three cases we will consider are as follows.

(A) We know that a positive proportion of the L^2 mass of η is clustered on a few points where η is large. Suppose that a significant amount of L^2 mass can also be found on points where η is *small*; currently this cannot be ruled out.

Then, by discarding these points – which by Proposition 2.2.8 has very little effect on $\|A'\|_{U^3}$ – we can obtain a fairly respectable energy increment at very little cost. In the above notation, we set

$$S = \{x \in (\mathbb{Z}/2\mathbb{Z})^d \colon \eta(x) \text{ is large}\} .$$

So, we may henceforth assume that almost all of the L^2 mass of $\eta(x)$ is found at points where $\eta(x) \gg K^{-C}$.

(B) The mainstay of the strategy is to restrict to a single fiber modulo 2 (i.e., to a single residue class modulo 2 in $\mathbb{Z}^d$). In the above notation we take $S \subseteq (\mathbb{Z}/2\mathbb{Z})^d$ to be a single point. By the previous case, this involves a density penalty of $O(K^C)$.

We exploit the fact that the energy of these fibers is, on average (in an appropriate sense), at least the energy of the whole set. Moreover, it will usually be slightly larger, *unless* the function η has normalized energy close to the maximum, i.e., $\mathscr{E}(\eta) \geq 0.99$.

(C) The final case handles this possibility that η has close to maximal energy.

Fortunately, in this case there is a structure theorem, that asserts that η is close to being the indicator function of a coset of a subgroup of $(\mathbb{Z}/2\mathbb{Z})^d$, correctly normalized. Since η has small support compared to the whole group $(\mathbb{Z}/2\mathbb{Z})^d$, this has to be a proper subgroup.

However, η cannot be entirely supported on this coset, since by assumption A' affinely generates $\mathbb{Z}^d$. So, there are a few stray points of η lying outside.

The strategy is then (usually) to delete these extraneous points. One can argue that this obtains an energy increment: possibly a very modest one, but at a similarly modest cost in density.

In the main case, Case B, we increase the normalized energy by a factor of at least about 1.001 (say) and lose $O(K^{O(1)})$ in the density. As we remark above, it follows this case can occur only $O(\log K)$ times, and so the density loss is of the form $\exp(O(\log^2 K))$.

The analysis of Cases A and C is rather more delicate: there is no lower bound on the energy increase, and so each might occur an unbounded number of times. However, in the final evaluation we will see that these cases are in fact very efficient compared to Case B, achieving a polynomial trade-off between energy and density.

Remark 2.3.1. The unwelcome $\log^2 K$ comes from the inherently wasteful practice of restricting to precisely *one* fiber modulo 2 in Case B, which loses a factor of $K^{O(1)}$ density, in exchange for relatively small energy increment.

One might reasonably conjecture that a better energy increment can always be obtained using the more refined practice of passing to a coset as in Case C. However, our analysis does not show this. Moreover, such a conjecture seems likely to be roughly as hard as the Polynomial Freĭman–Ruzsa Conjecture over $\mathbb{F}_2^n$.

As mentioned above, this does leave open the possibility that a sufficiently good understanding of quantitative structural theory over $\mathbb{F}_2^n$ could imply significantly improved results over $\mathbb{Z}$ using these methods.

The remainder of this section will be concerned with making the above sketch precise, and in particular with the analysis of these three cases. This argument is rather technical (although rarely actually difficult) and has little conceptual content beyond what we have already said.

Finally before we proceed, we shall fix some further notation. As well as

$$A'_S = \{a \in A' \colon a \bmod 2 \in S\}$$

for $S \subseteq (\mathbb{Z}/2\mathbb{Z})^d$, we will write

$$A'_\omega = \{a \in A' \colon a \bmod 2 = \omega\}$$

for the degenerate case where $S = \{\omega\}$ is a single point $\omega \in (\mathbb{Z}/2\mathbb{Z})^d$. So, for instance, $\eta(\omega) = \|A'_\omega\|_2$, and we continue to use this symbol in what follows.

As above, the constant $K \geq 1$ will always be defined by $\mathscr{E}(A') = 1/K$.

2.3.1 Case A: discarding points where η is small

Recall we are trying to show that either we can obtain a reasonable energy increment, or we may assume that, up to a negligible error, η is supported on large values.

We formalize this as follows, continuing to use the notation introduced above.

Proposition 2.3.2. *Let $0 < R < 1$ be a parameter. The following dichotomy holds: either η can be decomposed as*

$$\eta = \eta_{\text{big}} + \eta_{\text{sml}}$$

where η_{big}, η_{sml} have disjoint supports, $\eta_{\text{big}}(x)/\|\eta\|_2 \geq 2^{-8}R^4K^{-4}$ whenever it is non-zero and $\|\eta_{\text{sml}}\|_2 \leq R\|\eta\|_2$; or we can find a set $S \subseteq (\mathbb{Z}/2\mathbb{Z})^d$ such that

$$\frac{\mathscr{E}(A'_S)}{\mathscr{E}(A')} \geq (|A'|/|A'_S|)^{1/4} \ .$$

In the latter case, we may further assume that $A'_S \neq A'$.

Observe that the energy increment we obtain here is a power of the density loss, which is acceptable and indeed much more efficient than we hope to achieve overall.

We first isolate the following lemma.

Lemma 2.3.3. *Let*

$$S := \{x \in (\mathbb{Z}/2\mathbb{Z})^d \colon \eta(x)/\|\eta\|_2 \geq W\}$$

for some parameter W. Define $t := |A'_S|/|A'|$. Then

$$\frac{\mathscr{E}(A'_S)}{\mathscr{E}(A')} \geq t^{-1/2}\left(1 - W^{1/4}K\sqrt{1-t}\right) \ .$$

Proof. Write $f = 1_{A'}$, $f_{\text{big}} = 1_{A'_S}$ and $f_{\text{sml}} = f - f_{\text{big}}$. Similarly define $\eta_{\text{big}} = \eta \cdot 1_S$ and $\eta_{\text{sml}} = \eta - \eta_{\text{big}}$.

By Proposition 2.2.8 we have that

$$\begin{aligned}\|f_{\text{sml}}\|_{U^3} &\leq \|\eta_{\text{sml}}\|_{8/3} \\ &\leq \|\eta_{\text{sml}}\|_\infty^{1/4}\|\eta_{\text{sml}}\|_2^{3/4} \\ &\leq W^{1/4}\|\eta_{\text{sml}}\|_2 \ .\end{aligned}$$

By the triangle inequality, $\|f_{\text{big}}\|_{U^3} \geq \|f\|_{U^3} - \|f_{\text{sml}}\|_{U^3}$ and so

$$\begin{aligned}\frac{\mathscr{E}(f_{\text{big}})}{\mathscr{E}(f)} &= \frac{\|f\|_2}{\|f_{\text{big}}\|_2}\frac{\|f_{\text{big}}\|_{U^3}}{\|f\|_{U^3}} \\ &\geq t^{-1/2}\left(1 - \frac{W^{1/4}\|\eta_{\text{sml}}\|_2}{\|f\|_{U^3}}\right) \\ &= t^{-1/2}\left(1 - W^{1/4}K\sqrt{1-t}\right)\end{aligned}$$

as required. □

Proof of Proposition 2.3.2. We set $W = 2^{-8}R^4K^{-4}$ in Lemma 2.3.3 and take S to be the set from the statement of that lemma, i.e., S is the set of points x where $\eta(x) \geq W$. Similarly, define η_{big}, η_{sml} as above.

If the first case of the dichotomy does not hold with these choices, it must be that

$$\|\eta_{\text{sml}}\|_2 \geq R\|\eta\|_2$$

since the other conditions hold by definition. Using the notation $t = |A'_S|/|A'|$ as above and noting that

$$(\|\eta_{\text{big}}\|_2/\|\eta\|_2)^2 = t, \quad (\|\eta_{\text{sml}}\|_2/\|\eta\|_2)^2 = 1 - t$$

this says precisely that $(1-t) \geq R^2$.

In that case, we apply Lemma 2.3.3 to deduce that

$$\begin{aligned}\frac{\mathscr{E}(A'_S)}{\mathscr{E}(A')} &\geq t^{-1/2}\left(1 - W^{1/4}K\sqrt{1-t}\right) \\ &= t^{-1/2}\left(1 - R\sqrt{1-t}/4\right) \\ &\geq t^{-1/2}\left(1 - (1-t)/4\right) \\ &\geq t^{-1/4}\end{aligned}$$

as required (since $3/4 + t/4 \geq t^{1/4}$ by the AM–GM inequality). □

2.3.2 Case B: the small energy regime and passing to a fiber

Our goal in this section is the following result.

Proposition 2.3.4. *Again we continue with previous notation. Fix a parameter $\varepsilon > 0$. Suppose that*

$$\mathscr{E}(\eta) \leq 1 - \varepsilon$$

Then either we may find a set $S \subseteq (\mathbb{Z}/2\mathbb{Z})^d$ *such that*

$$\frac{\mathscr{E}(A'_S)}{\mathscr{E}(A')} \geq (|A'|/|A'_S|)^{1/4}$$

and $A'_S \neq A'$*, as in Proposition 2.3.2; or there is some* $\omega \in (\mathbb{Z}/2\mathbb{Z})^d$ *such that*

$$|A'|/|A'_\omega| \leq 2^{24}K^{16}\varepsilon^{-8}$$

and

$$\frac{\mathscr{E}(A'_\omega)}{\mathscr{E}(A')} \geq 1 + \varepsilon/2 \ .$$

First, we note the following inequality concerning the fibering properties of U^3 energy, which follows easily from the Gowers–Cauchy–Schwarz inequality for the U^3-norm.

Lemma 2.3.5. *Let* G*,* H *be two abelian groups and* $\phi\colon G \to H$ *a surjective group homomorphism. Let* $f\colon G \to \mathbb{R}_{\geq 0}$ *be a finitely supported function. For* $h \in H$*, write*

$$f_h(x) = f(x) \cdot 1_{\phi^{-1}(h)}(x) \ .$$

Finally, define

$$\begin{aligned} \gamma\colon H &\to \mathbb{R}_{\geq 0} \\ h &\mapsto \|f_h\|_{U^3} \ . \end{aligned}$$

Then $\|f\|_{U^3} \leq \|\gamma\|_{U^3}$*.*

Proof. Clearly

$$\|f\|_{U^3}^8 = \left\| \sum_{h \in H} f_h \right\|_{U^3}^8$$

and expanding this out we obtain

$$\|f\|_{U^3}^8 = \sum_{h_1,\ldots,h_8 \in H} \langle\!\langle f_{h_1}, \ldots, f_{h_8} \rangle\!\rangle_{U^3} \ .$$

If $(h_1, \ldots, h_8)$ do *not* form a 3-dimensional parallelepiped in H, the corresponding Gowers inner product will vanish. Hence

$$\|f\|_{U^3}^8 = \sum_{h,r_1,r_2,r_3 \in H} \langle\!\langle f_h, f_{h+r_1}, \ldots, f_{h+r_1+r_2+r_3} \rangle\!\rangle_{U^3} \ .$$

By the Gowers–Cauchy–Schwarz inequality, we deduce

$$\|f\|_{U^3}^8 \leq \sum_{h,r_1,r_2,r_3 \in H} \|f_h\|_{U^3} \|f_{h+r_1}\|_{U^3} \ldots \|f_{h+r_1+r_2+r_3}\|_{U^3}$$

and the right hand side is precisely $\|\gamma\|_{U^3}^8$, as required. □

For our current application, we will apply this when $G = \mathbb{Z}^d$, $H = (\mathbb{Z}/2\mathbb{Z})^d$, ϕ is the obvious projection map and $f = 1_{A'}$. In particular, we will now write

$$\begin{aligned}\gamma\colon (\mathbb{Z}/2\mathbb{Z})^d &\to \mathbb{R}_{\geq 0}\\ h &\mapsto \|A'_h\|_{U^3}\ .\end{aligned}$$

From this point, the proof of Proposition 2.3.4 is essentially straightforward, but the details are rather involved.

Proof of Proposition 2.3.4. Let us first sketch the case where η only takes large values on its support. We have that

$$\begin{aligned}\|A'\|_{U^3} &\leq \|\gamma\|_{U^3}\\ &\leq \sup_{x\colon \eta(x)\neq 0}\{\gamma(x)/\eta(x)\}\|\eta\|_{U^3}\\ &\leq \sup_{x\colon \eta(x)\neq 0}\{\gamma(x)/\eta(x)\}(1-\varepsilon)\|\eta\|_2\end{aligned}$$

where we have applied Lemma 2.3.5 and noted that if $\eta(x) = 0$ then necessarily $\gamma(x) = 0$.

So this states precisely that there exists some $x \in (\mathbb{Z}/2\mathbb{Z})^d$ with $\eta(x) \neq 0$ such that

$$\gamma(x)/\eta(x) \geq (1-\varepsilon)^{-1}\|A'\|_{U^3}/\|A'\|_2$$

and since $\gamma(x) = \|A'_x\|_{U^3}$, $\eta(x) = \|A'_x\|_2$ this is what we want – provided we can be sure that $|A'|/|A'_x| = O(K^{O(1)})$ for any x with $\eta(x) \neq 0$.

By applying Proposition 2.3.2, we are free to assume this, up to a small error in the L^2 norm in some sense. The technical work is now to rerun this argument, tolerating this further error term.

Setting $R = \varepsilon/2K$ in Proposition 2.3.2, either we have the energy increment conclusion as required, or we may decompose

$$\eta = \eta_{\text{big}} + \eta_{\text{sml}}$$

where $\eta_{\text{big}}(x) \geq 2^{-12}\varepsilon^4 K^{-8}\|\eta\|_2$ on its support $S \subseteq (\mathbb{Z}/2\mathbb{Z})^d$, and $\|\eta_{\text{sml}}\|_2/\|\eta_2\| \leq \varepsilon/2K$. We similarly decompose $\gamma = \gamma_{\text{big}} + \gamma_{\text{sml}}$ where $\gamma_{\text{big}} = 1_S \cdot \gamma$. Then arguing

much as before,

$$\begin{aligned}\|A'\|_{U^3} &\leq \|\gamma\|_{U^3} \\ &\leq \|\gamma_{\mathrm{big}}\|_{U^3} + \|\gamma_{\mathrm{sml}}\|_{U^3} \\ &\leq \|\gamma_{\mathrm{big}}\|_{U^3} + \|\gamma_{\mathrm{sml}}\|_2 \\ &\leq \sup_{x\in S}\{\gamma(x)/\eta(x)\}\|\eta_{\mathrm{big}}\|_{U^3} + \|\eta_{\mathrm{sml}}\|_2 \\ &\leq \sup_{x\in S}\{\gamma(x)/\eta(x)\}(1-\varepsilon)\|\eta\|_2 + \|\eta_{\mathrm{sml}}\|_2\end{aligned}$$

where we have made use of the pointwise bound $\gamma(x) \leq \eta(x)$ and the trivial bound $\|\eta_{\mathrm{big}}\|_{U^3} \leq \|\eta\|_{U^3}$. Hence, there exists an $x \in S$ such that

$$\begin{aligned}\frac{\gamma(x)/\eta(x)}{\mathscr{E}(A')} &\geq (1-\varepsilon)^{-1}\left(1 - \frac{\|\eta_{\mathrm{sml}}\|_2/\|\eta\|_2}{\|A'\|_{U^3}/\|A'\|_2}\right) \\ &\geq (1-\varepsilon)^{-1}(1-(\varepsilon/2K)\cdot K) \geq 1+\varepsilon/2\end{aligned}$$

as required. Also, the condition

$$|A'|/|A'_x| \leq 2^{24}K^{16}\varepsilon^{-8}$$

is satisfied since $x \in S$ (and by the definition of S). □

2.3.3 Case C: large energy and passing to a coset

Recall that our final case handles the possibility that η has very close to maximal normalized U^3 energy. Specifically, we will prove the following.

Proposition 2.3.6. *There exists some absolute constant $\delta > 0$ such that the following holds. Suppose $\mathscr{E}(\eta) \geq 1-\delta$, and $K \leq 2^{d/8-4}$. Then there exists a set $S \subseteq (\mathbb{Z}/2\mathbb{Z})^d$ such that*

$$\frac{\mathscr{E}(A'_S)}{\mathscr{E}(A')} \geq (|A'|/|A'_S|)^{1/8}$$

and furthermore $A'_S \neq A'$.

Again we note that in this phase of the argument we get an energy increment which is polynomially large compared to the density penalty.

Our key ingredient is the following result.

Proposition 2.3.7. *For all $\varepsilon > 0$ there exists $\delta > 0$ such that the following holds. Let G be any discrete abelian group and $f\colon G \to \mathbb{R}_{\geq 0}$ a finitely supported function with $\|f\|_2 = 1$. Suppose that $\|f\|_{U^3} \geq 1-\delta$.*

Then there exists a finite subgroup $H \leq G$ and an element $x \in G$ such that

$$\left\| f - \frac{1}{|H|^{1/2}} 1_{x+H} \right\|_2 \leq \varepsilon \ .$$

Proof. This is essentially a special case of [22, Theorem 1.4]. There, the authors consider functions which need not be real-valued and non-negative, and deduce that they are close in L^2 to a function $\chi(y) = \frac{1}{|H|^{1/2}} 1_{x+H}(y) e(P(y))$ where P is some quadratic polynomial (in fact, they also prove a result for all U^k-norms and any LCA group). When f is real-valued and non-negative, it is clear that replacing χ by $|\chi|$ can only decrease the L^2 distance from f, so we deduce the above as an easy consequence. □

So, we may assume that η is roughly the indicator function of a coset. As we remarked earlier, this will be a proper coset, since η has small support. In particular, this means that most, but not all, of the mass of η lies in some co-dimension one coset of $(\mathbb{Z}/2\mathbb{Z})^d$. Under these circumstances, it is always possible to obtain a modest energy increment, as the following lemma shows.

Lemma 2.3.8. *Let $\phi\colon \mathbb{Z}^d \to \mathbb{Z}/2\mathbb{Z}$ be a surjective homomorphism. For $i \in \{0,1\}$ write $A'_i = \phi^{-1}(\{i\})$, and write*

$$\alpha_i = (|A'_i|/|A'|)^{1/2} = \|A'_i\|_2/\|A'\|_2$$

so that $\alpha_0^2 + \alpha_1^2 = 1$. Suppose that $\min\{\alpha_0, \alpha_1\} \leq 0.1$. Then for some $i \in \{0,1\}$ we have

$$\frac{\mathscr{E}(A'_i)}{\mathscr{E}(A')} \geq (|A'|/|A'_i|)^{1/8} \ .$$

Proof. The analysis is very similar to that of the previous section. We define

$$\gamma_i = \|A'_i\|_{U^3}$$

for $i \in \mathbb{Z}/2\mathbb{Z}$. By Lemma 2.3.5 we have that

$$\|A'\|_{U^3} \leq \|\gamma\|_{U^3} = \left(\gamma_0^8 + 14\gamma_0^4\gamma_1^4 + \gamma_1^8\right)^{1/8} \ .$$

Write $\gamma'_i = \gamma_i/\|A'\|_{U^3}$. Then we may assume for contradiction that

$$\gamma'_i/\alpha_i < \alpha_i^{-1/4}$$

for each i, since otherwise the conclusion of the lemma is satisfied. That is, we may assume $\gamma'_i < \alpha_i^{3/4}$. Combining this with

$$(\gamma'_0)^8 + 14(\gamma'_0)^4(\gamma'_1)^4 + (\gamma'_1)^8 \geq 1$$

from above, we find that

$$\alpha_0^6 + 14\alpha_0^3\alpha_1^3 + \alpha_1^6 > 1 \ .$$

From this fact, together with $\alpha_0^2 + \alpha_1^2 = 1$, we wish to deduce that $\min\{\alpha_0, \alpha_1\} > 0.1$. This is now a routine algebraic exercise. We set $t = \alpha_0^2 - \alpha_1^2$, which lies in the range $[-1, 1]$, so that

$$\left(\frac{1+t}{2}\right)^3 + 14\left(\frac{1+t}{2}\right)^{3/2}\left(\frac{1-t}{2}\right)^{3/2} + \left(\frac{1-t}{2}\right)^3 > 1$$

which, after expansion and rearrangement, yields

$$2 + 6t^2 + 14(1-t^2)^{3/2} > 8$$

which further rearranges to

$$(1-t^2)^{3/2} > \frac{3}{7}(1-t^2)$$

and hence we deduce that

$$|t| < \frac{2\sqrt{10}}{7}$$

and so

$$\min\{\alpha_0, \alpha_1\} = \left(\frac{1-|t|}{2}\right)^{1/2} > 0.21$$

as required. □

With this unenlightening computation complete, we formalize the discussion above to deduce Proposition 2.3.6.

Proof of Proposition 2.3.6. By Proposition 2.3.7 we may choose $\delta > 0$ such that for η satisfying the condition of the statement, there is a subgroup $H \leq (\mathbb{Z}/2\mathbb{Z})^d$ and $x \in (\mathbb{Z}/2\mathbb{Z})^d$ such that

$$\left\| \eta/\|\eta\|_2 - \frac{1}{|H|^{1/2}} 1_{x+H} \right\|_2 \leq 0.1 \ .$$

We now argue that (WLOG) H is proper. Suppose for contradiction that it is not. Once again, it is convenient to apply Proposition 2.3.2, now with parameter $R = 1/2$ (say); so that, either we have an energy increment at least as large as the one we require and are done, or we may decompose $\eta = \eta_{\text{big}} + \eta_{\text{sml}}$ as in that statement. By the triangle inequality, we deduce that

$$\left\| \eta_{\text{big}}/\|\eta\|_2 - 2^{-d/2} \right\|_2 \leq 0.1 + \|\eta_{\text{sml}}\|_2/\|\eta\|_2 \leq 0.6 \ .$$

However, since $\eta_{\text{big}}(y)/\|\eta\|_2$ is either 0 or at least $2^{-12}K^{-4} \geq 2^4 \cdot 2^{-d/2}$, we have that

$$\sum_{y \in (\mathbb{Z}/2\mathbb{Z})^d} \left| \eta_{\text{big}}(y)/\|\eta\|_2 - 2^{-d/2} \right|^2 \geq \sum_{y \in (\mathbb{Z}/2\mathbb{Z})^d} 2^{-d} = 1$$

which is a contradiction.

Hence, we may assume that H is proper, and so there exists an index two subgroup $H \leq H' < (\mathbb{Z}/2\mathbb{Z})^d$. We know that almost all of the L^2 mass of η (or equivalently A') is contained in just one of the cosets of H'. Specifically, if $(y + H')$ is the coset *not* containing $(x + H)$, then

$$\|(\eta/\|\eta\|_2) \cdot 1_{y+H'}\|_2 \leq \left\| \eta/\|\eta\|_2 - \frac{1}{|H|^{1/2}} 1_{x+H} \right\|_2 \leq 0.1 \ .$$

Hence, quotienting by H' we obtain a map $\mathbb{Z}^d \to \mathbb{Z}/2\mathbb{Z}$ such that the hypotheses of Lemma 2.3.8 apply. This completes the proof. □

2.3.4 Completing the proof of Theorem 2.1.5

It remains only to put all these pieces together to deduce Theorem 2.1.5.

Initially, we are given a set $A \subseteq \mathbb{Z}$ such that $E(A)/|A|^3 \geq 1/L$. By Proposition 2.2.1, we deduce that the U^3 energy is also large, in that $\mathscr{E}(A) \geq L^{-1/2}$.

We now formally set up a recursive procedure for passing to subsets of A with larger U^3 energy, stopping when the resulting set has small Freĭman dimension.

Write $A_0 = A$, and write $K_i = \mathscr{E}(A_i)$. All our previous discussion will establish the following recursive fact.

Lemma 2.3.9. *Given a set $A_n \subseteq \mathbb{Z}$, either $\dim(A_n) \leq 8 \log_2 K_n + 32$, or there exists a proper subset $A_{n+1} \subsetneq A_n$ such that $K_{n+1} \leq K_n$ and*

$$\log_2 \left(\frac{|A_n|}{|A_{n+1}|} \right) \ll \big(\log_2 K_0 + O(1) \big) \log_2 \left(\frac{K_n}{K_{n+1}} \right) \ .$$

Remark 2.3.10. Reviewing the preceding arguments, it is fairly clear that the constant 32 appearing here could be replaced by 0.001 (or indeed any fixed positive constant), at the expense of worsening the implied $O(1)$ constants elsewhere. This means that the bound on the Freĭman dimension of $8 \log_2(K) + 0.001$ we obtain this way is essentially best possible in some sense, at least if we are allowed to replace the indicator function A with a (Gaussian) density.

The reason we do not stress this optimality elsewhere in this chapter is that it only holds with reference to the U^3 energy. When we attempt to transfer this result to other measures of structure – such as U^2 energy or doubling – even the leading constant in $O(\log K)$ is no longer optimal.

Proof of Theorem 2.1.5 assuming Lemma 2.3.9. Since each inclusion is proper, the process must terminate. That is, there is some N such that A_N is defined and

$$\dim(A_N) \leq 8 \log_2 K_N + 32 \leq 8 \log_2 K_0 + 32 \leq 4 \log_2 L + 32 \ .$$

By telescoping the recursive size inequality, we deduce that

$$\log_2 \left(\frac{|A_0|}{|A_N|} \right) \ll \big(\log_2 K_0 + O(1) \big) \log_2 \left(\frac{K_0}{K_N} \right)$$

and bounding $K_N \geq 1$ trivially we deduce

$$|A_0|/|A_N| \ll \exp\big(O(\log^2 K_0)\big) \ll \exp\big(O(\log^2 L)\big)$$

as required. □

Proof of Lemma 2.3.9. Let $d = \dim(A_n)$. By the definition of Freĭman dimension, we may find a subset $A'_n \subseteq \mathbb{Z}^d$ and a Freĭman (2-)isomorphism $\phi\colon A_n \leftrightarrow A'_n$.

We remark the following.

Proposition 2.3.11. *The U^3 energy is preserved under Freĭman isomorphisms. In the current setting, this means that $\|A_n\|_{U^3} = \|A'_n\|_{U^3}$.*

Proof. It suffices to show that a Freĭman homomorphism sends 3-dimensional parallelepipeds to 3-dimensional parallelepipeds. Certainly this holds for 2-dimensional parallelograms, also known as additive quadruples: indeed, this is exactly the definition of a Freĭman homomorphism.

We observe that for an abelian group G, a configuration $\{0,1\}^3 \to G$ is a 3-dimensional parallelepiped (that is, has the form $\omega \mapsto g_0 + \omega_1 g_1 + \omega_2 g_2 + \omega_3 g_3$) if and only if every codimension one face of $\{0,1\}^3$ (that is, sets of the form $\{\omega \in \{0,1\}^d\colon \omega_i = r\}$ for fixed $i \in \{1,2,3\}$ and $r \in \{0,1\}$) maps to an additive quadruple. The proof of this is an easy exercise, and it clearly implies the result. □

We now invoke one of Proposition 2.3.4 and Proposition 2.3.6: which depends on whether the quantity $\mathscr{E}(\eta)$ is bigger or smaller than $(1-\delta)$, where δ is the positive absolute constant appearing in Proposition 2.3.6. We set A'_{n+1} to be the set A'_S produced by either of those results as appropriate, and define $A_{n+1} = \phi^{-1}(A'_{n+1})$.

We conclude that either

$$K_n/K_{n+1} \geq (|A_n|/|A_{n+1}|)^{1/8}$$

as implied by either the first case of Proposition 2.3.4 or by Proposition 2.3.6; or,

$$K_n/K_{n+1} \geq 1 + \delta/2$$

and

$$|A_n|/|A_{n+1}| \leq 2^{24} K_n^{16} \delta^{-8} \leq 2^{24} K_0^{16} \delta^{-8}$$

as returned by the second case of Proposition 2.3.4. Taking logarithms, we find that either

$$\log_2 \left(\frac{|A_n|}{|A_{n+1}|} \right) \leq 8 \log_2 \left(\frac{K_n}{K_{n+1}} \right)$$

or

$$\begin{aligned} \log_2 \left(\frac{|A_n|}{|A_{n+1}|} \right) &\leq 24 + 8 \log(1/\delta) + 16 \log K_0 \\ &\leq \frac{24 + 8 \log(1/\delta) + 16 \log K_0}{\log(1 + \delta/2)} \log \left(\frac{K_n}{K_{n+1}} \right) \end{aligned}$$

either of which is acceptable for Lemma 2.3.9. □

Remark 2.3.12. The only bar to obtaining completely explicit constants in Theorem 2.1.5 is currently the mysterious constant δ obtained from the structure theorem for functions with almost maximal U^3 energy. Doubtless this constant is moderately sensible; however, no explicit dependencies are given in [22] and one would have to examine their proofs, which in turn rely on the classification of almost extremizers in Young's inequality due to Fournier [24].

Alternatively, one might reasonably seek a self-contained proof of Proposition 2.3.7, exploiting the fact that we only require it for non-negative functions, rather than going via [22].

Given we have made no attempt to optimize constants in any part of this argument, actually pursuing this would perhaps be in poor taste. The purpose of this comment is to reassure the reader that the implied constants in Theorem 2.1.5 are not too ridiculous.

2.4 An alternative proof of an inequality of Shkredov

We now give, for interest only, an alternative proof of the harder direction of Proposition 2.2.1, a result due to Shkredov [53]. Since this argument is orthogonal to the main theme of the chapter, we place it separately here.

Recall we are trying to show the following.

Proposition 2.4.1. *If Z is any abelian group and $f\colon Z \to \mathbb{R}_{\geq 0}$ a non-negative finitely supported function, then*

$$\|f\|_{U^3} \geq \|f\|_{U^2}^2 / \|f\|_1 \ .$$

When $f = 1_X$ is the indicator function of a set, this can be interpreted as follows: if X contains many 2-dimensional parallelepipeds – say, $\delta|X|^3$ of them – then it contains many 3-dimensional parallelepipeds, i.e. at least $\delta^4|X|^4$. A result of this flavour is implicit in [34], although the implied proof is quite complicated (involving the Balog–Szemeredi–Gowers theorem and Freĭman modelling) and obtains a worse exponent.

We note that Proposition 2.4.1 does *not* follow from the elementary fact of monotonicity of the Gowers norms, usually stated as $\|f\|_{U^3} \geq \|f\|_{U^2}$. Indeed, this formula assumes probability measure on Z is being used (and so is false with our conventions); to correct this requires a factor of the measure $\mu(Z)$ of the whole group, which renders the result useless unless f is dense. In our case, Z could even be infinite and so f far from dense.

Proof of Proposition 2.4.1. Define a function

$$\begin{aligned} R\colon Z \times Z &\to \mathbb{R}_{\geq 0} \\ (a,b) &\mapsto \sum_{x \in Z} f(x)f(x+a)f(x+b)f(x+a+b) \end{aligned}$$

or in other words, the number of parallelepipeds with sides a and b, counted weighted by f. We note that

$$\sum_{a,b \in Z} R(a,b) = \|f\|_{U^2}^4 \tag{2.4.1}$$

and

$$\sum_{a,b \in Z} R(a,b)^2 = \|f\|_{U^3}^8 \ . \tag{2.4.2}$$

We also have that

$$\sum_{b \in Z} R(a,b) = \left(f * f_-(a)\right)^2$$

where we write f_- for the function $f_-(x) = f(-x)$; and since $\sum_a f * f_-(a) = \|f\|_1^2$, we deduce that

$$\sum_{a \in Z} \left(\sum_{b \in Z} R(a,b) \right)^{1/2} = \|f\|_1^2 \ . \tag{2.4.3}$$

One might hope that in fact (2.4.1), (2.4.2) and (2.4.3) are the only facts we need to bear in mind to prove the result: that is, that we can forget where R came from

and just treat it as an arbitrary non-negative function on $Z \times Z$. In other words, it would suffice to show the following purely analytic result.

Lemma 2.4.2. *Let $R\colon Z \times Z \to \mathbb{R}_{\geq 0}$ be an arbitrary function. Then*

$$\sum_{a,b\in Z} R(a,b)^2 \geq \frac{\left(\sum_{a,b\in Z} R(a,b)\right)^4}{\left[\sum_{a\in Z}\left(\sum_{b\in Z} R(a,b)\right)^{1/2}\right]^2 \left[\sum_{b\in Z}\left(\sum_{a\in Z} R(a,b)\right)^{1/2}\right]^2} \,.$$

Fortunately this is true. Although Lemma 2.4.2 is essentially completely elementary, we are not aware of a completely elementary proof. Instead, we rely on some standard entropy inequalities.

Proof of Lemma 2.4.2. Define

$$W(a,b) := R(a,b)\left(\sum_{a,b\in Z} R(a,b)\right)^{-1}$$

i.e. normalizing so that $\sum_{a,b} W(a,b) = 1$. We may think of $W(a,b)$ as the joint probability distribution of a pair of random variables X, Y on Z, so that $\mathbb{P}[X = x, Y = y] = W(x,y)$. Then, the quantities

$$W_X(a) = \sum_{b\in Z} W(a,b)$$

and

$$W_Y(b) = \sum_{a\in Z} W(a,b)$$

are precisely the marginal distributions of X and Y.

What we want to prove is then that

$$\sum_{a,b\in Z} W(a,b)^2 \geq \left(\sum_{a\in Z} W_X(a)^{1/2}\right)^{-2} \left(\sum_{b\in Z} W_Y(b)^{1/2}\right)^{-2} \,.$$

This statement is of the following flavour: "if X and Y are concentrated on a few values then so is (X,Y)". A more usual way to express the same sentiment might be in terms of the entropy inequality $H(X,Y) \leq H(X) + H(Y)$. But noting the inequalities (which are standard under the title "monotonicity of Rényi entropy", or are easy convexity arguments)

$$2\log\left[\sum_a W_X(a)^{1/2}\right] \geq H(X)$$

and similarly for Y; and

$$\log\left[\sum_{a,b} W(a,b)^2\right] \geq -H(X,Y)$$

this implies the result.

□

This completes the proof of Proposition 2.4.1. □

2.5 A result about doubling and reduction modulo 2

The following result is an analogue of Proposition 2.2.8 for sets, in terms of an hypothesis on the doubling constant rather than the U^3 energy. This fact was not needed in the main argument, and is stated for interest only.

Proposition 2.5.1. *Suppose $A \subseteq \mathbb{Z}^d$ is a finite set, and let $|A+A| = K|A|$.*

Consider the image of A under reduction modulo 2; that is,

$$\tilde{A} := \{x \bmod 2 \colon x \in A\} \subseteq (\mathbb{Z}/2\mathbb{Z})^d \ .$$

Then $|\tilde{A}| \leq K^6$.

Our key ingredient will be a well-known covering argument due to Ruzsa [49], which we state in its most general form.

Lemma 2.5.2. *Let X, Y be finite subsets of the same abelian group Z.*

Then there exists an integer k, $k \leq |X+Y|/|X|$, and elements $\{y_1, \ldots, y_k\} \subseteq Y$ such that

$$Y \subseteq \bigcup_{i=1}^{k} (y_i + X - X) \ .$$

We briefly recall the proof.

Proof. The trick is to select a maximal set $\{y_1, \ldots, y_k\} \subseteq Y$ such that the sets $(y_i + X)$ are disjoint. By maximality, for every $y \in Y$ there exists $i \in [k]$ such that $(y+X) \cap (y_i + X) \neq \emptyset$. Equivalently, $y \in y_i + X - X$. Furthermore, note that $(y_i + X)$ are disjoint subsets of $X+Y$, each of size $|X|$, so clearly $k|X| \leq |X+Y|$ as required. □

We deduce the following, which is a slight variant on the well-known result that if A has small doubling then $A - A$ is an approximate group.

Corollary 2.5.3. *Suppose A is a finite subset of an abelian group Z, and $|A+A| = K|A|$. Then $(2A-2A)$ is covered by at most K^6 translates of $2\cdot(A-A)$.*

Note by $2\cdot X$ we mean the *dilate* $\{2x\colon x\in X\}$, as opposed to the sumset $2X = X+X$.

Proof. We apply the Ruzsa covering lemma with $X = 2\cdot A$ and $Y = 2A-2A$. We deduce that Y is covered by k translates of $2\cdot A - 2\cdot A = 2\cdot(A-A)$, where

$$k \leq |2\cdot A + 2A - 2A|/|A| \leq |2A+2A-2A|/|A| \leq K^6$$

where we have used the obvious inclusion $2\cdot A \subseteq 2A$ and Plünnecke's inequalities. □

Remark 2.5.4. Our preferred way of thinking about this result is that $A-A$ satisfies some kind of rough convexity. Indeed, for a convex subset $X\subseteq\mathbb{R}^d$ we have $2\cdot X = 2X$. It is fairly clear this cannot happen for non-trivial subsets of $\mathbb{Z}^d$, since there is a parity obstruction to writing $x_1+x_2 = 2x_3$ for general $x_1,x_2,x_3\in X$. However, if $X\subseteq\mathbb{Z}^d$ were the intersection of $\mathbb{Z}^d$ with a convex subset of $\mathbb{R}^d$, we might expect the slight weakening $2X\subseteq 2\cdot X + \{-1,0,1\}^d$ to hold.

It is this property that the previous corollary seeks to emulate.

Our original proposition is a trivial consequence of Corollary 2.5.3.

Proof of Proposition 2.5.1. By Corollary 2.5.3 we have that $2A-2A\subseteq 2\cdot(A-A)+S$ for some set S, $|S|\leq K^6$. Therefore $2A-2A$ takes at most K^6 distinct values modulo 2. Equivalently, $|2\tilde{A}-2\tilde{A}|\leq K^6$. It follows trivially that $|\tilde{A}|\leq K^6$. □

Remark 2.5.5. The same argument shows that the number of residues modulo N is bounded by K^{4+N} for any positive integer N.

Chapter 3

Functions whose derivatives take boundedly many values

A function $f\colon \mathbb{Z} \to \mathbb{R}$ is a polynomial of degree s, in the usual sense, if and only if all its discrete derivatives of order $(s+1)$ are identically zero.

Here, we ask what one can say about functions $f\colon \mathbb{Z} \to \mathbb{R}$ satisfying the weaker hypothesis that their $(s+1)$-th discrete derivatives take only K distinct values, for some bounded positive integer K. This is by no means the only possible approximate variant of the original hypothesis, and in fact turns out to be a fairly strong one.

Examples of such functions include: genuine polynomials; any function that takes only a bounded number of distinct values; and (less trivially) functions of the form $a \cdot \{p(x)\}$, where $p\colon \mathbb{Z} \to (\mathbb{R}/\mathbb{Z})^d$ is a genuine polynomial, $\{\cdot\}\colon (\mathbb{R}/\mathbb{Z})^d \to [0,1)^d$ is the fractional part map, and $a \in \mathbb{R}^d$.

In this chapter, we show that all such functions f are built out of these examples, in a strong sense. Such a result might be of interest (i) for its own sake, (ii) for the techniques used in its proof, or (iii) as a possible stepping stone to a stronger result applicable to weaker notions of approximate polynomials.

3.1 Introduction

Throughout, let G and A be arbitrary abelian groups, both notated additively.

Given a function $f\colon G \to A$ and an element $a \in G$, we define the *discrete derivative* of f at a to be

$$\begin{aligned} \partial_a f\colon G &\to A \\ x &\mapsto f(x) - f(x+a)\ . \end{aligned}$$

We say f is a *polynomial of degree* s if for any $a_1, \dots, a_{s+1} \in G$ the function

$$\partial_{a_1} \dots \partial_{a_{s+1}} f$$

is identically zero. In familiar cases – e.g., when $G = \mathbb{Z}$ and $A = \mathbb{R}$ – this coincides with the usual notion of a polynomial of degree s.

We can relax this definition in several ways. One natural one is to consider functions f whose $(s+1)$-th derivatives are *often* zero. If G is finite and $f\colon G \to A$ is a function such that

$$\#\left\{(a_1, \dots, a_{s+1}, x) \in G^{s+2} \colon \partial_{a_1} \dots \partial_{a_{s+1}} f(x) = 0\right\} \geq \varepsilon |G|^{s+2}$$

we say that f is an *ε-approximate polynomial*, or informally a 1% polynomial, of degree s.

Alternatively, suppose (again when G is finite) that a dense subset $X \subseteq G$ and a function $f\colon X \to A$ are given. The derivative $\partial_{a_1} \dots \partial_{a_{s+1}} f(x)$ is only defined when all the points

$$\left\{x + \omega \cdot a \colon \omega \in \{0,1\}^{s+1}\right\}$$

lie in X, where $a = (a_1, \dots, a_{s+1})$ and $\omega \cdot a$ denotes $\sum_{\omega_i = 1} a_i$. If $\partial_{a_1} \dots \partial_{a_{s+1}} f(x) = 0$ whenever this condition holds, then we say f is an *X-local polynomial*.

When G is infinite, e.g. $G = \mathbb{Z}$, similar notions might be formalized by restricting to finite intervals, or by considering other density notions, but we will not consider such statements here.

Some interesting examples of 1% or local polynomials come from *bracket polynomials*; e.g. in degree 2, functions such as

$$\begin{aligned} &\mathbb{Z} \to \mathbb{R} \\ &x \mapsto \alpha\{\beta x\{\gamma x\}\} \end{aligned}$$

where here $\{\cdot\}\colon \mathbb{R} \to [0,1)$ is the fractional part map and $\alpha, \beta, \gamma \in \mathbb{R}$ are parameters. More generally, any function arising from addition, multiplication and application of the fractional part map will have these properties for appropriate parameters s and ε or X respectively.

It seems likely that these are in some sense the only examples of approximate or local polynomials. However, to our knowledge, no proof of such a result is known; and there is good reason to believe that any such proof must be at least as hard as, or make use of, the inverse theorem for the Gowers U^{s+2}-norm.[1]

In this chapter, we instead work with the following stronger variant of these hypotheses.

[1]The link between approximate polynomials and Gowers norm statements is in fact the author's original motivation for thinking about such questions.

Definition 3.1.1. Let $f\colon G \to A$ be a function and $S \subseteq A$ a finite set. We say f is a *polynomial of degree s up to S-errors*, if for all $a_1, \ldots, a_{s+1}, x \in G$ we have

$$\partial_{a_1} \ldots \partial_{a_{s+1}} f(x) \in S \ . \tag{3.1.1}$$

Clearly the case $S = \{0\}$ corresponds to genuine polynomials. In general, it is easy to see that $0 \in S$ and (WLOG) $S = -S$.

When G is infinite, (3.1.1) is always a non-trivial condition. For finite G, the reader should imagine S having some fixed size $|S| = K$ as $|G|$ becomes large.

It is not clear that this is in fact a stronger notion than the two we have already discussed, but this turns out to be the case in an appropriate sense.

Our main result is a strong classification of polynomials up to S-errors. We first state this very concretely when $G = \mathbb{Z}$.

Theorem 3.1.2. *Let $s \geq 0$ be given, let $S \subseteq \mathbb{R}$ be a finite set, $|S| = K$, and let $f\colon \mathbb{Z} \to \mathbb{R}$ be a polynomial of degree s up to S-errors.*

Then there exist:

- *a non-negative integer $d \leq K$;*
- *real numbers r_i for $1 \leq i \leq d$ with $|r_i| \leq K \max\{|s|\colon s \in S\}$;*
- *real numbers $\alpha_i^{(0)}, \ldots, \alpha_i^{(s)}$ for $1 \leq i \leq d$, as well as $\beta^{(0)}, \ldots, \beta^{(s)}$;*
- *functions $\tau_i\colon \mathbb{Z} \to \mathbb{Z}$ for $1 \leq i \leq d$ such that $|\tau_i(x)| \leq K^2 5^K$ for all $1 \leq i \leq d$ and $x \in \mathbb{Z}$;*

such that

$$f(x) = \beta^{(0)} + \beta^{(1)}x + \cdots + \beta^{(s)}x^s + \sum_{i=1}^{d} r_i \left(\tau_i(x) + \left\{\alpha_i^{(0)} + \alpha_i^{(1)}x + \cdots + \alpha_i^{(s)}x^s\right\}\right)$$

for all $x \in \mathbb{Z}$, where $\{\cdot\}\colon \mathbb{R} \to [0,1)$ is the fractional part map.

The statement in the general case is as follows.

Theorem 3.1.3. *Let $s \geq 0$ be given, let $S \subseteq \mathbb{R}$ be a finite set, $|S| = K$, and let $f\colon G \to \mathbb{R}$ be a polynomial of degree s up to S-errors.*

Then there exist:

- *a non-negative integer $d \leq K$;*
- *a polynomial (in the above sense) $p\colon G \to (\mathbb{R}/\mathbb{Z})^d$ of degree s;*
- *a function $\tau\colon G \to \mathbb{Z}^d$ with $|\tau(x)_i| \leq K^2 5^K$ for all x and $1 \leq i \leq d$;*
- *a vector $r \in \mathbb{R}^d$; and*

- *a polynomial (again in the above sense) $F\colon G \to \mathbb{R}$ of degree s;*

such that

$$f(x) = F(x) + r \cdot \big(\tau(x) + \{p(x)\}\big) \tag{3.1.2}$$

where this time $\{\cdot\}\colon (\mathbb{R}/\mathbb{Z})^d \to [0,1)^d$ denotes the obvious bijective lift.

Finally, we have the bound $\|r\|_\infty \leq K \max\{|s|\colon s \in S\}$.

Remark 3.1.4. In other words, this states that polynomials up to S-errors are (essentially) bracket polynomials with at most one top-level bracket. Hence there is a notable contrast with the local or 1% polynomial cases when $s > 1$, which both allow more exotic examples with up to s brackets.

One might then hope that a classification of 1% or local polynomials might proceed by some kind of iteration of this approach to introduce more and more brackets. However, this remains speculative.

Remark 3.1.5. The case $s = 1$ of Theorem 3.1.3 is already non-trivial: it asserts that functions such that $f(x) - f(x+a) - f(x+b) + f(x+a+b)$ takes only boundedly many distinct values are essentially bracket-linear. This statement has roughly the same power as Freĭman's theorem, insofar as each can be deduced from the other using relatively elementary tools. As such, the proof of Theorem 3.1.3 must include an implicit proof of Freĭman's theorem. To our knowledge, this is not quite the same as any of the existing proofs, but is closely related to *universal ambient group* approaches, such as that appearing in [59, Section 5.6].

Remark 3.1.6. The proof of Theorem 3.1.3 is most natural in the case of finite groups G, which is still non-trivial. The generalization to infinite G is straightforward, provided one is happy to work in a infinitary style, i.e. exploiting translation-invariant means and their properties.

We will run our exposition nominally in the finite setting, and later offer some remarks concerning the infinite case. In the first instance, the reader will lose nothing by assuming that $G = \mathbb{Z}/N\mathbb{Z}$ for some large prime N.

The rough outline of the proof is to build the torus $(\mathbb{R}/\mathbb{Z})^d$ and the maps p and τ out of a suitable vector space of functions with the same properties as f. This is very similar in spirit to applications of the universal ambient group, in particular to a proof of Freĭman's theorem [59, Chapter 5.6], where the vector space of Freĭman homomorphisms on a set of small doubling is used to build a generalized progression.

3.1.1 Layout of the chapter

Section 3.2 gives a more detailed summary of the approach. The construction of $(\mathbb{R}/\mathbb{Z})^d$, p, F etc. is described in Section 3.3; and the required quantitative information, as well as some loose ends, are derived in Section 3.4. Finally, Section 3.5 gives some remarks concerning the case of infinite groups.

3.2 High-level aspects of the proof

We first fix some further notation. It will be convenient to treat the iterated derivative $\partial_{a_1} \dots \partial_{a_k} f(x)$ as a function of all of the variables $a_1, \dots, a_k, x$, not just the last. So, given a function $f\colon G \to A$ and an integer $k \geq 0$, we define

$$\begin{aligned} \partial^k f\colon G^{k+1} &\to A \\ (a_1, \dots, a_k; x) &\mapsto \partial_{a_1} \dots \partial_{a_k} f(x) \\ &= \sum_{\omega \in \{0,1\}^k} (-1)^{|\omega|} f(x + \omega \cdot a) \end{aligned}$$

where $|\omega| = \sum_{i=1}^k \omega_i$ and $\omega \cdot a$ denotes the element $\sum_{\omega_i=1} a_i$ of G.

Given a finite set $S \subseteq \mathbb{R}$ and $f\colon G \to \mathbb{R}$ such that f is a polynomial of degree s up to S-errors, it makes sense to partition G^{s+2} into $K = |S|$ pieces, according to the value of $\partial^{s+1} f$ at each point. That is, if we enumerate $S = \{s_1, \dots, s_K\}$ arbitrarily, we may define sets

$$X_i = \left\{ z \in G^{s+2} \colon \partial^{s+1}(z) = s_i \right\}$$

for $1 \leq i \leq K$, which by assumption partition G^{s+2}.

Throughout, we use the notation $j(z)$ for $z \in G^{s+2}$ to denote the unique $j \in \{1, \dots, K\}$ such that $z \in X_j$.

The key idea of the proof is the following. Rather than considering the function f in isolation, we will consider simultaneously all functions $g\colon G \to \mathbb{R}$ with the property that $\partial^{s+1} g$ is constant on each cell X_i. It is clear that the space of such functions forms a real vector space, which we denote by V; i.e.:

$$V = \left\{ g\colon G \to \mathbb{R} \colon \forall i, \forall z, z' \in X_i \colon \partial^{s+1} g(z) = \partial^{s+1} g(z') \right\} . \tag{3.2.1}$$

The point is that we can use this vector space V to build the torus $(\mathbb{R}/\mathbb{Z})^d$ and the polynomial map $p\colon G \to (\mathbb{R}/\mathbb{Z})^d$ appearing in the conclusion of Theorem 3.1.3.

Slightly more accurately, we note that there is a subspace of V consisting of all genuinely polynomial maps $G \to \mathbb{R}$ of degree s, which we denote by $\mathrm{poly}_s(G, \mathbb{R}) \subseteq V$.

For Theorem 3.1.3, we need a procedure for decomposing $f \in V$ into $f = F + f'$ where F is a true polynomial and f' will have the form we require. So, what we want is a canonical *complementation* $V = \mathrm{poly}_s(G, \mathbb{R}) \oplus V'$ of the subspace $\mathrm{poly}_s(G, \mathbb{R})$. Correctly, it is this space V', isomorphic to the quotient $V/\mathrm{poly}_s(G, \mathbb{R})$, that gives rise to the torus $(\mathbb{R}/\mathbb{Z})^d$ and the map p.

Note that we have an exact sequence

$$0 \to \mathrm{poly}_s(G, \mathbb{R}) \to V \xrightarrow{\partial^{s+1}} \mathbb{R}^K \tag{3.2.2}$$

where $\mathbb{R}^K$ records the value of the derivative $\partial^{s+1} g$ on each cell X_i. So, the quotient $V/\mathrm{poly}_s(G, \mathbb{R})$ is in turn isomorphic to the image $\partial^{s+1}(V)$ in $\mathbb{R}^K$. Note this already shows that $\dim V' \leq K$.

What we need then is precisely a *section* of ∂^{s+1} on its image, i.e. a map $\sigma \colon \partial^{s+1}(V) \to V$ such that $\partial^{s+1} \circ \sigma$ is the identity. We then reconstruct f' from the derivative $\partial^{s+1} f$ by setting $f' = \sigma(\partial^{s+1} f)$.

In fact this process is easy to describe: explicitly,

$$f'(x) = \int_{a_1, \dots, a_{s+1} \in G} \partial^{s+1} f(a_1, \dots, a_{s+1}; x) \ . \tag{3.2.3}$$

When G is finite, the integral sign simply denotes the uniform measure on G, i.e. $\frac{1}{|G|^k} \sum_{a_1, \dots, a_k \in G} [\dots]$. If G is infinite, we have to interpret the integral in (3.2.3) as an *invariant mean*, or finitely-additive translation-invariant probability measure on G. In that setting, this construction of f' is used for a different but closely related purpose in [57], which gives an elegant proof of a generalization of the Hyers–Ulam stability theorem.

It is not hard to verify that $F := f - f'$ as defined is a polynomial $G \to \mathbb{R}$. It remains, then, to explain precisely how $(\mathbb{R}/\mathbb{Z})^d$ and p arise from V' as we have claimed, and to verify that this f' then has the form required by Theorem 3.1.3.

Unfortunately, formally the proof differs from the picture we have just given, in that it turns out to be more natural to work with the *dual* (or more accurately, pre-dual) of these vector spaces of functions, than with the spaces themselves. So, in our actual construction, all the arrows will be reversed relative to this discussion, although the underlying ideas are the same.

The remainder of the proof splits into two pieces: an algebraic part, describing in detail how the various objects $(\mathbb{R}/\mathbb{Z})^d$, p, τ, r provided by Theorem 3.1.3 arise in this way; and an analytic part, which establishes the necessary quantitative control on d, τ, r etc..

3.3 The algebraic construction

The algebraic construction is notationally somewhat opaque but essentially completely formal. As mentioned above, it is closely related to the *universal ambient group* construction in the case of Freĭman homomorphisms (which are the same thing as local polynomials of degree 1). For an exposition on this subject, see e.g. [59, Chapter 5.5].

We define W to be the real vector space with basis $\{e_x \colon x \in G\}$, quotiented by the relations

$$\sum_{\omega \in \{0,1\}^{s+1}} (-1)^{|\omega|} e_{x+\omega \cdot a} - \sum_{\omega \in \{0,1\}^{s+1}} (-1)^{|\omega|} e_{x'+\omega \cdot a'} = 0 \tag{3.3.1}$$

for all $i \in \{1, \dots, K\}$ and $(a; x), (a'; x') \in X_i$. There is an obvious map $\iota \colon G \to W$ given by $\iota(x) = e_x$.

The motivation for this definition is that the space of functions V defined above corresponds precisely to the space of linear functionals $W \to \mathbb{R}$. Indeed, given a function $g \colon G \to \mathbb{R}$ that lies in V, we may define a linear map on W by sending e_x to $g(x)$ (which is well-defined on the quotient by (3.2.1)); and conversely a linear functional $\phi \colon W \to \mathbb{R}$ yields a function $\phi \circ \iota$ on G, which lies in V by (3.3.1).

Now let Γ be the subgroup of W generated (over $\mathbb{Z}$, not $\mathbb{R}$) by the elements

$$\gamma(a; x) = \sum_{\omega \in \{0,1\}^k} (-1)^{|\omega|} e_{x+\omega \cdot a} \tag{3.3.2}$$

for $(a; x) \in G^{s+2}$. These are somehow *formal derivatives*: evaluating a functional at $\gamma(a; x)$ is the same as evaluating $\partial^{s+1} f(a; x)$.

Note that the relations (3.3.1) state precisely that $\gamma(z) = \gamma(z')$ whenever $z, z' \in X_j$, and hence $\gamma(z)$ takes at most K distinct values and Γ is generated by a set of size at most K. We also write γ_j to denote the element $\gamma(z)$ for any $z \in X_j$.

Let $U = \langle \Gamma \rangle_{\mathbb{R}}$ denote the real span of Γ, i.e. a vector subspace of W, and set $d = \dim U \leq K$.

Remark 3.3.1. The subspace U is dual to the image $\partial^{s+1}(V)$ in the heuristic discussion above.

We will need the following additional facts about this set-up, whose proofs are deferred until Section 3.4 as they arise naturally in that context.

Lemma 3.3.2. *With notation as above:*

(*i*) *the group* Γ *is discrete, i.e. a lattice in* U*, and hence* U/Γ *is isomorphic to a torus* $(\mathbb{R}/\mathbb{Z})^d$*;*

(*ii*) *there is a canonical linear projection* $\pi\colon W \to U$ *(i.e.,* $\pi|_U$ *is the identity on* U*).*

Remark 3.3.3. Under the duality we have been working with, the projection $\pi\colon W \to U$ is dual to the section σ of ∂^{s+1} mentioned in Section 3.2. For now, we can be agnostic about how π is defined, but it will be important to use the averaged construction (3.2.3) (or rather its dual) in order to obtain quantitative control later.

We now have all the tools in place to describe the construction of the objects required by Theorem 3.1.3.

Because f is an element of V, we may associate to it a linear functional $\phi\colon W \to \mathbb{R}$. The projection π gives us a decomposition $W = U \oplus \ker \pi$ via $w \mapsto \big(\pi(w), w - \pi(w)\big)$, and we apply this to get

$$\begin{aligned} f(x) &= \phi(\iota(x)) \\ &= \phi\big(\pi(\iota(x))\big) + \phi\big(\iota(x) - \pi(\iota(x))\big)\ . \end{aligned}$$

We first consider the right-hand term,

$$\begin{aligned} F\colon G &\to \mathbb{R} \\ x &\mapsto \phi\big(\iota(x) - \pi(\iota(x))\big) \end{aligned}$$

and claim that it is actually a genuine polynomial on G of degree s. Indeed,

$$\begin{aligned} \partial^{s+1} F(a;x) &= \sum_{\omega \in \{0,1\}^{s+1}} (-1)^{|\omega|} \phi\big(\iota(x + \omega \cdot a) - \pi(\iota(x + \omega \cdot a))\big) \\ &= \phi\big(\gamma(a;x) - \pi(\gamma(a;x)\big) \end{aligned}$$

by exploiting linearity of ϕ and π, and where $\gamma(a;x)$ is defined as in (3.3.2). However, $\pi(\gamma(a;x)) = \gamma(a;x)$ by assumption on π, and so the right hand side is zero as required.

It remains to deal with the left-hand term $x \mapsto \phi\big(\pi(\iota(x))\big)$. Because Γ is a lattice in U, we may (arbitrarily for now) choose an integral basis for Γ, so that U is identified with $\mathbb{R}^d$, the lattice Γ with $\mathbb{Z}^d$ and the quotient U/Γ with the torus $(\mathbb{R}/\mathbb{Z})^d$. Writing $\{\cdot\}\colon (\mathbb{R}/\mathbb{Z})^d \to [0,1)^d$ for the canonical lift, and likewise $\lfloor x \rfloor \in \mathbb{Z}^d$ for the element $x - \{x \bmod 1\}$ given $x \in \mathbb{R}^d$, we may define

$$\begin{aligned} \tau\colon G &\to \mathbb{Z}^d \\ x &\mapsto \lfloor \pi(\iota(x)) \rfloor \end{aligned}$$

and

$$\begin{aligned} p\colon G &\to (\mathbb{R}/\mathbb{Z})^d \\ x &\mapsto \pi(\iota(x)) \bmod 1 \end{aligned}$$

so that $\phi\big(\pi(\iota(x))\big) = \phi\big(\tau(x) + \{p(x)\}\big)$.

Finally, using the standard inner product with respect to our chosen basis for Γ, we may choose the unique $r \in \mathbb{R}^d$ so that $\phi(u) = r \cdot u$ for all $u \in U$. Combining all these definitions, we find that

$$f(x) = F(x) + r \cdot \big(\tau(x) + \{p(x)\}\big) \ ,$$

exactly as required for Theorem 3.1.3.

We now claim that p is a polynomial map of degree s. Essentially this is because all formal derivatives lie in Γ, and we have already quotiented out by that. Formally, we compute

$$\begin{aligned} \partial^{s+1} p(a;x) &= \sum_{\omega \in \{0,1\}^{s+1}} (-1)^{|\omega|} \pi\big(\iota(x + \omega \cdot a)\big) \bmod \Gamma \\ &= \pi(\gamma(a;x)) \bmod \Gamma \\ &= \gamma(a;x) \bmod \Gamma = 0 \end{aligned}$$

where we have used linearity of π, together with the fact that π fixes elements of Γ.

Putting all this together, we have constructed all the objects $(\mathbb{R}/\mathbb{Z})^d \cong U/\Gamma$, p, τ, r, F appearing in the conclusion of Theorem 3.1.3, satisfying (3.1.2) and their respective algebraic properties. It remains only to:

- prove Lemma 3.3.2;
- specify a particular integral basis for Γ;
- show that the function τ, as defined, is bounded coordinate-wise by $K^2 5^K$; and
- verify the claimed bound on $\|r\|_\infty$;

where we note that the latter two depend on the correct choice of integral basis. We turn to these points in the next section.

3.4 Quantitative bounds

We return briefly to the discussion in Section 3.2. Recall we were interested in the exact sequence

$$0 \to \mathrm{poly}_s(G, \mathbb{R}) \to V \xrightarrow{\partial^{s+1}} \mathbb{R}^K$$

and wished, among other things, to obtain a simple description of the image $\partial^{s+1}(V)$. In the dual language of the previous section, this is equivalent to investigating what relations are satisfied by the generating set $\gamma(a;x)$ of U.

It is certainly true that functions of the form $h = \partial^k g$ satisfy some linear *cocycle relations*, which in this context are

$$\begin{aligned} h(a_1, \ldots, a_i + a_i', \ldots, a_k; x) = \quad & h(a_1, \ldots, a_i, \ldots, a_k; x) \\ & + h(a_1, \ldots, a_i', \ldots, a_k; x + a_i) \end{aligned} \tag{3.4.1}$$

for all $1 \leq i \leq k$ and $a \in G^k$, $x, a_i' \in G$. We might call functions h satisfying these conditions *cocycles*. So, certainly the image of $\partial^{s+1}(V)$ in $\mathbb{R}^K$ must be contained in the subspace cut out by such relations.

It turns out that:

(i) in fact, equality holds; i.e. all functions h satisfying (3.4.1) have the form $\partial^{s+1} g$ for some g (a kind of cohomological triviality statement);

(ii) indeed, one shows this by exhibiting a linear map σ from the space of cocycles in $\mathbb{R}^K$ to V, such that $\partial^{s+1} \circ \sigma$ is the identity on that space.

As a consequence, we get a bounded complexity description of the image $\partial^{s+1}(V)$ as a subspace of $\mathbb{R}^K$ by taking a suitable subset of the relations (3.4.1).

The map σ is constructed exactly as in (3.2.3), i.e.:

$$(\sigma h)(x) = \int_{a_1, \ldots, a_k \in G} h(a_1, \ldots, a_k; x)$$

but it will be important to verify the properties of this construction when applied to an arbitrary cocycle h, not just the case $h = \partial^{s+1} g$, in order to show this cohomological triviality.

We have already mentioned the similarity of this approach to methods in [57]. It was also inspired by arguments in [10], although this case is much more straightforward than the analogous ones treated there.

As before, all these remarks are identical in spirit but formally dual to what we actually want to prove, which we now state formally.

Lemma 3.4.1. *We fix notation as in previous sections, and assume temporarily that G is finite. Then the following hold.*

(*i*) *Consider the linear map* $\mu\colon \mathbb{R}^K \to U$ *that sends the standard basis vectors* $(\varepsilon_j)_{j=1}^K$ *in* $\mathbb{R}^K$ *to* γ_j *respectively.*

Again we abuse notation to write $\varepsilon(z) = \varepsilon_j$ *whenever* $z \in X_j$*; i.e.* $\varepsilon(z)$ *is defined to be* $\varepsilon_{j(z)}$ *for any* $z \in G^{s+2}$*.*

Then $\ker \mu$ *is generated by the elements*

$$\begin{aligned}\varepsilon(a_1, \dots, a_i + a_i', \dots, a_{s+1}; x) - \varepsilon(a_1, \dots, a_i, \dots, a_{s+1}; x) \\ - \varepsilon(a_1, \dots, a_i', \dots, a_{s+1}; x + a_i)\end{aligned} \tag{3.4.2}$$

for any $a_1, \dots, a_{s+1}, a_i', x \in G$ *and* $1 \le i \le s+1$*. Equivalently, the relations*

$$\begin{aligned}\gamma(a_1, \dots, a_i + a_i', \dots, a_{s+1}; x) = \quad & \gamma(a_1, \dots, a_i, \dots, a_{s+1}; x) \\ & + \gamma(a_1, \dots, a_i', \dots, a_{s+1}; x + a_i)\end{aligned}$$

hold for any choice of parameters, and generate all the relations satisfied by $(\gamma_j)_{j=1}^K$*.*

(*ii*) *Define a map* $\pi\colon W \to U$ *by*

$$\sum \alpha_x e_x \mapsto \sum \alpha_x \left(\int_{a_1, \dots, a_k \in G} \gamma(a; x) \right)$$

(*where* $\{e_x\colon x \in G\}$ *are the basis used to define* W*). Then* $\pi|_U$ *is the identity.*

(*iii*) *There exist vectors* $(\beta_1, \dots, \beta_d) \in \Gamma$ *such that for every* $1 \le i \le K$ *we may write*

$$\gamma_i = \sum_{j=1}^{d} \varphi_{ij} \beta_j$$

for some $\varphi_{ij} \in \mathbb{Z}$ *with* $|\varphi_{ij}| \le K^2 5^K$*. Conversely, each* β_i *may be expressed as*

$$\beta_i = \sum_{j=1}^{K} t_{ij} \gamma_j$$

for some $t_{ij} \in \mathbb{R}$ *with* $|t_{ij}| \le 1$*.*

Note that it is intended that parts (i) and (ii) of this lemma are dual to the remarks (i) and (ii) above, respectively.

We first show that this is enough to complete the proof of Theorem 3.1.3.

Proof of the remainder of Theorem 3.1.3 given Lemma 3.4.1. It is clear from part (iii) of the lemma that Γ is a lattice in U with $(\beta_i)_{i=1}^d$ as an integral basis; and part (ii) asserts the existence of the projection π. Hence Lemma 3.3.2 is proven, and it suffices to prove the other statements listed at the end of Section 3.3.

For every $x \in G$, the element $\pi(\iota(x))$ is a convex combination of elements $\gamma(a;x)$, and hence by (iii) may be expressed as

$$\pi\big(\iota(x)\big) = \sum_{i=1}^{d} y_i \beta_i$$

for some $y_i \in \mathbb{R}$ with $|y_i| \leq K^2 5^K$. By definition, $\tau(x) \in \mathbb{Z}^d$ is the vector $(\lfloor y_i \rfloor)_{i=1}^d$, and so $|\tau(x)_i| \leq K^2 5^K$ for each i as claimed.

Finally, note that, for $1 \leq j \leq K$ and any $(a;x) \in X_j$,

$$\begin{aligned} s_j &= \partial^{s+1} f(a;x) \\ &= \sum_{\omega \in \{0,1\}^{s+1}} (-1)^{|\omega|} \phi\big(\pi(\iota(x + \omega \cdot a))\big) \\ &= \phi(\gamma(a;x)) = \phi(\gamma_j) \ , \end{aligned}$$

and also, $r_i = \phi(\beta_i)$ for each $1 \leq i \leq d$ by definition. It follows that

$$r_i = \sum_{j=1}^{K} t_{ij} s_j$$

where t_{ij} are the coefficients from (iii), and hence $|r_i| \leq K \|s\|_\infty$ as required. □

We spend the remainder of this section proving Lemma 3.4.1.

At the risk of running out of letters, we let $Z \subseteq \mathbb{R}^K$ denote the subspace generated by elements of the form (3.4.2); i.e. (i) states that $Z = \ker \mu$.

We first consider the easy direction, i.e. that $Z \subseteq \ker \mu$. This is equivalent to saying that the equation

$$\begin{aligned} \gamma(a_1, \ldots, a_i + a_i', \ldots, a_{s+1}, x) = \quad & \gamma(a_1, \ldots, a_i, \ldots, a_{s+1}, x) \\ & + \gamma(a_1, \ldots, a_i', \ldots, a_{s+1}, x + a_i) \end{aligned} \tag{3.4.3}$$

does actually hold every choice of $1 \leq i \leq s+1$ and $a_1, \ldots, a_{s+1}, a_i', x \in G$. This follows immediately from expanding out the definition of $\gamma(a;x)$ and noting the obvious cancellation.

So, from now on we may write μ^* for the well-defined map $\mathbb{R}^K/Z \to W$, and to show (i) it suffices to show that μ^* is injective.

We will prove the rest of (i) and (ii) in one go, by verifying the following.

Lemma 3.4.2. *There is a canonical linear map $\pi^*\colon W \to \mathbb{R}^K/Z$ such that $\pi^* \circ \mu^*$ is the identity on $\mathbb{R}^K/Z$ and $\mu^* \circ \pi^* = \pi$, with π defined as in Lemma 3.4.1(ii).*

Under the formal duality, this states exactly the properties of the map σ we claimed above.

The lemma clearly implies that μ^* is injective, proving (i). But μ is clearly surjective by construction, so μ^* is an isomorphism, and it follows entirely formally π is a projection. Indeed,

$$\begin{aligned}\pi(u) &= \mu^* \circ \pi^*(u)\\ &= \mu^* \circ \pi^* \circ \mu^*\big((\mu^*)^{-1}u\big)\\ &= \mu^* \circ \big((\mu^*)^{-1}u\big) = u\end{aligned}$$

for all $u \in U$, giving (ii).

Proof of Lemma 3.4.2. We define π^* by

$$\pi^*\colon \sum \alpha_x e_x \mapsto \sum \alpha_x \left(\int_{a_1,\dots,a_{s+1}\in G} \varepsilon(a;x)\right)$$

where $\varepsilon(z) = \varepsilon_{j(z)}$ are the basis vectors of $\mathbb{R}^K$ as defined in Lemma 3.4.1(i) (and we recall that for $z \in G^{s+2}$, $j(z)$ denotes the unique $j \in \{1,\dots,K\}$ such that $z \in X_j$). The identity $\pi = \mu^* \circ \pi^*$ is immediate from the definitions of π and π^*.

We now verify that $\pi^* \circ \mu^*(\varepsilon(a;x)) = \varepsilon(a;x) \pmod Z$ for any $(a;x) \in G^{s+2}$, which suffices. By the definition of Z we have

$$\varepsilon(a;x) = \varepsilon\big(a_1 + a_1', a_2, \dots, a_{s+1}; x\big) - \varepsilon\big(a_1', a_2, \dots, a_{s+1}; x + a_1\big) \pmod Z$$

for any $a_1' \in G$. If we now integrate over this parameter a_1', we deduce that

$$\varepsilon(a;x) = \int_{a_1'\in G} \varepsilon\big(a_1', a_2, \dots, a_{s+1}; x\big) - \varepsilon\big(a_1', a_2, \dots, a_{s+1}; x + a_1\big) \pmod Z$$

where we have used translation-invariance of the integral over a_1' (i.e. $\int_x h(x+y) = \int_x h(x)$). Applying this argument iteratively to the integrands for each coordinate $2 \le i \le s+1$ it turn, we obtain

$$\varepsilon(a;x) = \int_{a_1',\dots,a_{s+1}'\in G} \sum_{\omega\in\{0,1\}^{s+1}} (-1)^{|\omega|}\varepsilon\big(a_1', \dots, a_{s+1}'; x + \omega\cdot a\big) \pmod Z .$$

On the other hand, we can expand out

$$\begin{aligned}\pi^* \circ \mu^*\big(\varepsilon(a;x)\big) &= \pi^* \left(\sum_{\omega \in \{0,1\}^{s+1}} (-1)^{|\omega|} e_{x+\omega\cdot a} \right) \\ &= \sum_{\omega \in \{0,1\}^{s+1}} (-1)^{|\omega|} \int_{a'_1,\dots,a'_{s+1} \in G} \varepsilon\big(a'_1,\dots,a'_{s+1};x+\omega\cdot a\big)\end{aligned}$$

which gives the result. □

It now remains only to verify part (iii) of Lemma 3.4.1, concerning the existence of a suitable integral basis $(\beta_i)_{i=1}^d$ of Γ.

Recall that we have by assumption an over-determined generating set $\{\gamma_j : 1 \leq j \leq K\}$ for Γ, and wish to find an integral basis whose basis elements are bounded complexity linear combinations of these generators (and vice versa).

The key fact at our disposal is that, by virtue of (i), we have a complete list of relations satisfied by the γ_j (which are the same thing as generators of $\ker \mu$), all of which have small integer coefficients.

From this, one can immediately deduce *some* quantitative control of the kind required. Indeed, given this information, for fixed K there are only finitely many different possible sets of relations that the γ_i can satisfy. By choosing an fixed integral basis $(\beta_i)_{i=1}^d$ arbitrarily in each case, we can then trivially ensure that $|\varphi_{ij}|$ and $|t_{ij}|$ are bounded only in terms of K.

However, we will do a little more work to achieve more reasonable quantitative bounds. The arguments required are reasonably standard, somewhat technical, and largely uninteresting.

Reordering if necessary, we will assume that $(\gamma_j)_{j=1}^d$ are linearly independent, where we recall $d = \dim U$. So, the $\mathbb{Z}$-span $\Lambda = \langle \gamma_1, \dots, \gamma_d \rangle_{\mathbb{Z}}$ is certainly a lattice in U.

Also, we choose a basis $(k_i)_{i=1}^{K-d}$ for $\ker \mu$ consisting of elements of the form given by (3.4.2); so in particular each k_i has the form $\varepsilon(z) - \varepsilon(z') - \varepsilon(z'')$ for some $z, z', z'' \in G^{s+2}$.

Claim 3.4.3. *The index* $[\Gamma : \Lambda]$ *is bounded by* $5^{(K-d)/2}$. *In particular,* Γ *is a lattice in* U.

Proof of claim. The surjective group homomorphism

$$\begin{aligned} \mathbb{Z}^K &\to \Gamma/\Lambda \\ (a_j)_{j=1}^K &\mapsto \sum_{j=1}^K a_j\gamma_j \pmod{\Lambda} \end{aligned}$$

has kernel containing the $(K-d)$ vectors $(k_i)_{i=1}^{K-d}$, as well as the d vectors $(\varepsilon_j)_{j=1}^d$.

It follows that the image has cardinality at most the absolute value of the determinant of the $K \times K$ matrix with these K vectors as its rows. By the Hadamard inequality, this determinant is bounded by the product of the ℓ^2 norms of the rows, which is at most $5^{(K-d)/2}$ as required.[2] □

We recall the following standard fact, which follows from [54, Theorem 18] (or its proof).

Proposition 3.4.4 ([54, Theorem 18])**.** *Suppose Λ is a lattice of rank r and $(y_1, \dots, y_r)$ are elements of Λ. Then there is an integral basis $(x_1, \dots, x_r)$ for Λ and rationals $(c_i)_{i=1}^r, (c_{ij})_{1\le j<i\le r}$ such that*

$$\begin{aligned} x_1 &= c_1 y_1 \\ x_2 &= c_{21} y_1 + c_2 y_2 \\ &\vdots \\ x_r &= c_{r1} y_1 + \cdots + c_{r(r-1)} y_{r-1} + c_r y_r \end{aligned}$$

where furthermore $c_i = 1/a_i$ for positive integers a_i with $\prod_{i=1}^r a_i$ equal to the index of $\langle y_1, \dots, y_r\rangle_{\mathbb{Z}}$ in Λ, and $0 \le c_{ij} < c_j \le 1$ for all $1 \le j < i \le r$.

We apply the proposition to Γ with the y's supplied by $(\gamma_j)_{j=1}^d$, and let $(\beta_i)_{i=1}^d$ be the resulting x's. The second condition in (iii) is then immediate.

Conversely, by inverting the lower triangular matrix c_{ij} we obtain expressions

$$\gamma_i = \sum_{j=1}^d \varphi_{ij}\beta_j$$

for $1 \le i \le d$ and some integers $(\varphi_{ij})_{i,j=1}^d$. So, our next task is to bound the coefficients that can appear when inverting this lower-triangular matrix.

[2] The row $\varepsilon(z) - \varepsilon(z') - \varepsilon(z'')$ might have ℓ^2 norm $\sqrt{5}$, not $\sqrt{3}$ as one might first think, if e.g. $\varepsilon(z') = \varepsilon(z'')$.

This task is simplified by defining $y'_i = c_i y_i$ and writing $x_i = \sum_j c'_{ij} y'_j$, since the conditions on c_{ij} translate into $c'_i = 1$ and $0 \le c'_{ij} < 1$ for $0 \le j < i \le r$. So, we have

$$y'_i = x_i - \sum_{j<i} c'_{ij} y_j$$

and an easy inductive argument shows that

$$y'_i = \sum_{1 \le j \le i} \psi_{ij} x_j$$

for some coefficients ψ_{ij} with $\psi_{ii} = 1$ and $|\psi_{ij}| \le 2^{i-j-1}$ for every $1 \le j < i \le d$.

Since $\varphi_{ij} = a_i \psi_{ij}$, we deduce

$$|\varphi_{ij}| \le a_i 2^{\max(i-j-1,0)} \le 5^{(K-d)/2} 2^d \tag{3.4.4}$$

for $1 \le i, j \le d$.

To handle the remaining vectors γ_j for $d < j \le K$, we consider a $(K-d) \times K$ matrix M whose rows are the vectors $(k_i)_{i=1}^{K-d} \in \ker \mu$ defined above, which we recall form a basis for all the relations satisfied by the γ_j. Hence in particular $\sum_{j=1}^{K} M_{ij} \gamma_j = 0$ for all $1 \le i \le K - d$.

Writing $M = (L|T)$ where T consists of the latter $(K-d)$ columns, it follows from our assumptions (that k_i and $(\gamma_j)_{j=1}^d$ are linearly independent) that T is invertible. Hence

$$\sum_{j=1}^{K} (T^{-1}L|\,\mathrm{id})_{ij} \gamma_j = 0$$

or equivalently

$$\gamma_{d+i} = -\sum_{j=1}^{d} (T^{-1}L)_{ij} \gamma_j$$

for each $1 \le i \le K - d$.

Since T is an integral matrix, with each row having at most one $+1$ and two (-1)'s and the rest zero, by Cramer's rule and the Hadamard inequality we have $|(T^{-1})_{ij}| \le 5^{(K-d)/2}$ for all $1 \le i, j \le K - d$ (see the proof of Claim 3.4.3), and hence

$$|(T^{-1}L)_{ij}| \le (K-d) 5^{(K-d)/2}$$

for each i, j. Combining this with (3.4.4), we get

$$\begin{aligned} |\varphi_{ij}| &\le d(K-d) 5^{(K-d)/2)} \cdot 5^{(K-d)/2} 2^d \\ &\le K^2 5^K \end{aligned}$$

for all $1 \le i \le K$ and $1 \le j \le d$, which completes the proof of Lemma 3.4.1.

3.5 Infinitary considerations

The only place we have used finiteness of the group G in this argument is in the statement and proof of Lemma 3.4.1. There, we used this property to write down integrals over G of the form

$$\int_{a_1,\ldots,a_k \in G} [\ldots]$$

where the integral denoted an average over the group, i.e.

$$\frac{1}{|G|^k} \sum_{a_1,\ldots,a_k \in G} [\ldots] .$$

However, we now remark that the only properties of this integral operation that we actually used were:

- finite additivity, i.e. $\int (f+g) = \int f + \int g$;
- the probability normalization $\int 1 = 1$;
- convexity properties, i.e. if $C \subseteq \mathbb{R}^d$ is a closed convex set and $f(x) \in C$ for all x then $\int f \in C$; and
- translation-invariance, i.e. $\int_{x \in G} f(x+y) = \int_{x \in G} f(x)$ for any $y \in G$.

Moreover, the integrand is always a vector in a finite dimensional real vector space, irrespective of the cardinality of G.

Hence, if the group G is equipped with a translation-invariant, finitely additive probability measure ν, then one can verify that the whole proof goes through unchanged, interpreting every integral symbol $\int_{x \in G} f$ as the *translation-invariant mean* $\int f(x) d\nu(x)$ associated to the measure ν.

The existence of such a ν is equivalent to the amenability of G, and it is a standard fact that all (discrete) abelian groups are amenable. However, even for (say) $G = \mathbb{Z}$ such a mean cannot be explicitly constructed, meaning the proof we obtain this way is strongly infinitary in flavour. It should be possible, but not necessarily rewarding, to finitize the proof at least in simple cases such as $G = \mathbb{Z}$.

We note that Fubini's theorem fails in general for finitely additive measures; but, although we are occasionally notationally lax concerning the order of integration in expressions such as $\int_{a_1,\ldots,a_k \in G}$, a close reading of the proofs shows that this is not actually an issue.

Chapter 4

Periodic nilsequences and inverse theorems on cyclic groups

The inverse theorem for the Gowers norms, in the form proved by Green, Tao and Ziegler, applies to functions on an interval $[M]$. A recent paper of Candela and Sisask requires a stronger conclusion when applied to N-periodic functions; specifically, that the corresponding nilsequence should also be N-periodic in a strong sense.

In most cases, this result is implied by work of Szegedy (and Anatolín Camarena and Szegedy) on the inverse theorem. This deduction is given in Candela and Sisask's paper. In this chapter, we give an alternative proof, which uses only the Green–Tao–Ziegler inverse theorem as a black box. The result is also marginally stronger, removing a technical condition from the statement.

The proof centers around a general construction in the category of nilsequences and nilmanifolds, which is possibly of some independent interest.

This chapter is based on the preprint [44].

4.1 Introduction

The inverse theorem for the Gowers norms has a central role in many recent developments in additive combinatorics and related fields. We briefly recall the statement, as it appears in [36, Conjecture 1.2].

Theorem 4.1.1. *Fix an integer $s \geq 0$ and $\delta > 0$. Then there exists a finite collection $\mathscr{M}_{s,\delta}$ of tuples (G, Γ, d_G), where G is an s-step filtered, nilpotent Lie group, Γ a lattice in G (i.e. a discrete co-compact subgroup), and d_G a left-invariant Riemannian metric on G, such that the following holds: for any positive integer N, and any function $f\colon [N] \to \mathbb{C}$ with $|f| \leq 1$ and $\|f\|_{U^{s+1}[N]} \geq \delta$, there exists*

(*i*) *a tuple $(G, \Gamma, d_G) \in \mathscr{M}_{s,\delta}$;*

(*ii*) *a polynomial map* $p\colon \mathbb{Z} \to G$*;*

(*iii*) *a function* $F\colon G \to \mathbb{C}$ *that is automorphic* [1] *with respect to* Γ*, bounded in magnitude by* 1 *and is* $O_{s,\delta}(1)$*-Lipschitz with respect to* d_G*;*

such that the nilsequence $F \circ p$ *correlates with* f*, i.e.* $|\mathbb{E}_{x\in[N]} f(x) \cdot \overline{F \circ p(x)}| \gg_{s,\delta} 1$.

We will not reproduce definitions of the various terms here, referring the reader to [36] itself, [58] or the forthcoming [32].

This result only applies to functions whose domain is an interval in $\mathbb{Z}$. However, in the case $s = 1$, which is covered by classical Fourier analysis, similar statements can be made in much more general domains; in particular cyclic groups $\mathbb{Z}/N\mathbb{Z}$. Moreover, in this setting the structured functions $F \circ p$ can be taken to be characters on $\mathbb{Z}/N\mathbb{Z}$, and hence naturally respect the algebraic structure of $\mathbb{Z}/N\mathbb{Z}$.

For general s, in some sense the algebraic content is reflected in the polynomial map p. So, in formulating a version of Theorem 4.1.1 for cyclic groups, it makes some sense to require that p respects the algebraic structure of $\mathbb{Z}/N\mathbb{Z}$.

Definition 4.1.2. Let G be a (filtered) nilpotent Lie group and Γ a lattice in G. We say a polynomial map $p\colon \mathbb{Z} \to G$ is *N-periodic* if $p(x+N)p(x)^{-1} \in \Gamma$ for all x.

Informally, we refer to a function of the form $F \circ p\colon \mathbb{Z} \to \mathbb{C}$, for an N-periodic polynomial map $p\colon \mathbb{Z} \to G$ and some automorphic function $F\colon G \to \mathbb{C}$, as an N-periodic nilsequence.

Remark 4.1.3. Note that an N-periodic nilsequence in this sense should not be confused with the weaker property that $F \circ p$ is periodic *as a function*, i.e. that $F \circ p(x+N) = F \circ p(x)$ for all x.

We can now state the main result, the N-periodic analogue of Theorem 4.1.1.

Theorem 4.1.4 (Main theorem)**.** *Fix an integer* $s \geq 1$ *and* $\delta > 0$*. Then there exists a finite collection* $\mathscr{M}'_{s,\delta}$ *of tuples* (G, Γ, d_G) *as in Theorem 4.1.1, such that the following holds: for any positive integer* N*, and any function* $f\colon \mathbb{Z}/N\mathbb{Z} \to \mathbb{C}$ *with* $|f| \leq 1$ *and* $\|f\|_{U^{s+1}(\mathbb{Z}/N\mathbb{Z})} \geq \delta$*, there exists*

(*i*) *a tuple* $(G, \Gamma, d_G) \in \mathscr{M}'_{s,\delta}$*;*

(*ii*) *an* N*-periodic polynomial map* $p\colon \mathbb{Z} \to G$*;*

[1] This simply means that $F(\gamma x) = F(x)$ for all $x \in G$ and $\gamma \in \Gamma$, i.e. F descends to a function $\Gamma\backslash G \to \mathbb{C}$. One could also use the term "Γ-periodic (on the left)".

(*iii*) *a function* $F\colon G \to \mathbb{C}$ *that is automorphic with respect to* Γ*, bounded in magnitude by* 1 *and is* $O_{s,\delta}(1)$*-Lipschitz with respect to* d_G*;*

such that $F \circ p$ *correlates with* f*, i.e.* $|\mathbb{E}_{x \in \mathbb{Z}/N\mathbb{Z}} f(x) \cdot \overline{F \circ p(x)}| \gg_{s,\delta} 1$.

Such a result finds an application in recent work of Candela and Sisask [11], on convergence (over prime $N \to \infty$) of the minimum number of solutions to a fixed system of linear equations in a dense subset $A \subseteq \mathbb{Z}/N\mathbb{Z}$. The full strength of Definition 4.1.2 is required for this application to succeed.

In that paper, the authors establish Theorem 4.1.4 assuming a technical condition on N, as a consequence of Szegedy's approach to the inverse theorem for the Gowers norms (see in particular [56] as well as joint work with Anatolín Camarena [10]).

Our approach is to prove the following result, which constructs periodic nilsequences from non-periodic ones. Assuming this, it is straightforward to deduce Theorem 4.1.4.

Theorem 4.1.5. *Suppose we are given an (*s*-step filtered) nilpotent Lie group* G*, a lattice* Γ *in* G *and a left-invariant Riemannian metric* d_G*, as well as a polynomial map* $p\colon \mathbb{Z} \to G$*. Also suppose* $F\colon G \to \mathbb{C}$ *is an automorphic,* K*-Lipschitz function bounded in magnitude by* 1*. Further, suppose* $\phi\colon \mathbb{R}/\mathbb{Z} \to [0,1]$ *is a smooth function supported on* $[0,1/2]$*, and* N *is a positive integer.*

Then there exists:

(*i*) *a (canonical)* s*-step filtered nilpotent Lie group* $\tilde{G}$ *and lattice* $\tilde{\Gamma}$ *in* $\tilde{G}$*, which depend only on* G *and* Γ*;*

(*ii*) *a (canonical)* N*-periodic polynomial map* $\tilde{p}\colon \mathbb{Z} \to \tilde{G}$ *which depends only on* p *(and* G, Γ*); and*

(*iii*) *an* $O_{K,\Gamma,G,\phi}(1)$*-Lipschitz function* $\tilde{F}\colon \tilde{G} \to \mathbb{C}$*, automorphic with respect to* $\tilde{\Gamma}$ *and bounded in magnitude by one;*

such that

$$\tilde{F}(\tilde{p}(x)) = \phi(x/N) \cdot F(p(x \bmod N))$$

for all x*, where* $x \bmod N$ *refers to the representative in* $\{0, \dots, N-1\}$.

4.2 Deduction of Theorem 4.1.4

We record the following easy lemma.

Lemma 4.2.1. *Suppose $f\colon \mathbb{Z}/N\mathbb{Z} \to \mathbb{C}$ is supported on an interval J in $\mathbb{Z}/N\mathbb{Z}$ satisfying $N/20 \leq |J| \leq N/2$. Write $f'\colon J \to \mathbb{C}$ for the restriction of f to that interval. Then for any $s \geq 1$,*

$$\|f'\|_{U^{s+1}(J)} = C\|f\|_{U^{s+1}(\mathbb{Z}/N\mathbb{Z})}$$

for some constant C independent of f, satisfying $1 \ll_s C \ll_s 1$.

Proof. By translation-invariance we may assume $J = \{0, \dots, |J| - 1\}$. It suffices to prove that any parallelepiped $(x + h \cdot \omega)_{\omega \in \{0,1\}^{s+1}}$ in $\mathbb{Z}/N\mathbb{Z}$ that is entirely contained in J is also a parallelepiped with respect to $\mathbb{Z}$ (after embedding $\{0, \dots, |J| - 1\}$ in $\mathbb{Z}$ in the obvious way). The constant C then arises purely from the normalization constant in the definition of a Gowers norm on an interval and can be ignored.

It suffices to check the case $s = 1$, since any configuration (x_ω) is a parallelepiped if and only if all its 2-dimensional faces are parallelepipeds. In other words, we want to know that if $x, y, z, w \in \{0, \dots, |J| - 1\}$ and $x - y - z + w \equiv 0 \pmod{N}$ then $x - y - z + w = 0$; but this is clear as $-2(|J| - 1) \leq x - y - z + w \leq 2(|J| - 1)$. □

We proceed to the proof.

Proof of Theorem 4.1.4 assuming Theorem 4.1.5. We are free to assume, to avoid tedious issues, that N is not too small (say $N \geq 100$) since the result is trivial for small N (using, say, the U^2 case).

We choose 20 closed intervals $I_0, \dots, I_{19}$ in $\mathbb{R}/\mathbb{Z}$, each of width exactly $1/10$, whose interiors cover $\mathbb{R}/\mathbb{Z}$. For sake of argument we can fix $I_m = [m/20, m/20 + 1/10] \pmod 1$. Now choose a smooth partition of unity (ρ_m) on $\mathbb{R}/\mathbb{Z}$ adapted to I_m. The functions ρ_m descend to a partition of unity (ϕ_m) on $\mathbb{Z}/N\mathbb{Z}$ in the obvious way, i.e. by $\phi_m\colon x \mapsto \rho_m(x/N)$. Then, each such ϕ_m is supported on an interval of length either $\lfloor N/10 \rfloor$ or $\lceil N/10 \rceil$; call this interval J_m.

We deduce that

$$\begin{aligned}
\|f\|_{U^{s+1}(\mathbb{Z}/N\mathbb{Z})} &= \left\| \sum_m \phi_m \cdot f \right\|_{U^{s+1}(\mathbb{Z}/N\mathbb{Z})} \\
&\leq \sum_m \|\phi_m \cdot f\|_{U^{s+1}(\mathbb{Z}/N\mathbb{Z})} \\
&= C \sum_m \|\phi_m \cdot f\|_{U^{s+1}(J_m)}
\end{aligned}$$

where we have used the triangle inequality for U^{s+1} and Lemma 4.2.1.

Hence, for some m we have $\|\phi_m \cdot f\|_{U^{s+1}(J_m)} \gg_s \delta$. Translating everything if necessary, we can assume that ρ_m is supported on $[0, 1/2]$ and $J_m \subseteq \{0, \dots, N/2\}$.

We apply Theorem 4.1.1 to the function $\phi_m \cdot f$ on J_m to obtain $(G, \Gamma, d_G) \in \mathscr{M}_{s,\Omega_s(\delta)}$ and a nilsequence $\psi = F \circ p$, where $p\colon \mathbb{Z} \to G$ is a polynomial map, such that ψ correlates with $\phi_m \cdot f$, i.e. $|\mathbb{E}_{x \in J_m}(\phi_m \cdot f)(x)\overline{\psi(x)}| \gg_{s,\delta} (1)$.

Now apply Theorem 4.1.5 to the nilsequence ψ and the smooth cut-off ρ_m. We find that $\psi'\colon x \mapsto \rho_m(x/N)\psi(x \bmod N)$ is an N-periodic nilsequence with respect to some $\tilde{G}, \tilde{\Gamma}$ (depending only on G, Γ).

Defining $\mathscr{M}'_{s,\delta}$ to be those $(\tilde{G}, \tilde{\Gamma}, d_{\tilde{G}})$ obtained by applying Theorem 4.1.5 to elements of $\mathscr{M}_{s,\Omega_s(\delta)}$, and observing that

$$\begin{aligned}\left|\mathbb{E}_{x \in \mathbb{Z}/N\mathbb{Z}} f(x)\overline{\psi'(x)}\right| &= \left|\mathbb{E}_{x \in \{0,\dots,N-1\}} f(x)\overline{\phi_m(x)\psi(x)}\right| \\ &= \left|\mathbb{E}_{x \in J_m}(\phi_m \cdot f)(x)\overline{\psi(x)}\right| \gg_{s,\delta} 1 \ ,\end{aligned}$$

this completes the proof. □

4.3 The proof of Theorem 4.1.5

The construction we will use is very closely related to that described in [36, Appendix C], which in turn is based on an argument of Furstenberg [27, page 31]. The goal there was rather different to ours; namely, to show that any polynomial nilsequence can be realized as a linear one, at the expense of augmenting G, Γ to some $\tilde{G}$, $\tilde{\Gamma}$. Both problems require one to "invent" the Heisenberg group, or a variant of it, in the special case of the input $G = \mathbb{R}$, $\Gamma = \mathbb{Z}$; and it turns out that their general cases are also closely related.

Our proof therefore recaps that construction, with only a few (albeit significant) modifications. We quote a number of facts from [36, Appendix C] without reproducing the proofs.

Proof of Theorem 4.1.5. We consider the nilpotent Lie group of polynomial maps $\mathrm{poly}(\mathbb{Z}, G)$ with pointwise multiplication. There is an isomorphism $\mathrm{poly}(\mathbb{Z}, G) \cong \mathrm{poly}(\mathbb{R}, G)$ given by restriction (see [36, Appendix C]). We also consider the subgroup $\mathrm{poly}(\mathbb{Z}, \Gamma)$. By [36, Lemma C.1] the latter is discrete and co-compact.

We also define the shift action

$$\begin{aligned}T\colon \mathbb{R} \times \mathrm{poly}(\mathbb{R}, G) &\to \mathrm{poly}(\mathbb{R}, G) \\ (t, p) &\mapsto p(\cdot + t)\end{aligned}$$

and thereby the semi-direct product $\tilde{G} = \mathrm{poly}(\mathbb{R}, G) \rtimes_T \mathbb{R}$. To avoid confusion we note that by this ordering we mean to imply a group operation

$$(g, t) * (g', t') = (g \cdot t(g'), t + t') \ .$$

This has a subgroup $\tilde{\Gamma} = \mathrm{poly}(\mathbb{Z}, \Gamma) \rtimes_T \mathbb{Z}$ (since the action of T respects $\mathrm{poly}(\mathbb{Z}, \Gamma)$). The latter is discrete and co-compact in $\tilde{G}$. Indeed, every right coset $\tilde{\Gamma} x$ has a (unique) representative $(\mathrm{poly}(\mathbb{Z}, \Gamma) g, t)$ where $t \in [0, 1)$.

We can also be explicit about the filtration on $\tilde{G}$, which is $\tilde{G}_0 = \tilde{G}$, $\tilde{G}_1 = \mathrm{poly}(\mathbb{R}, G^{+1}) \rtimes_T \mathbb{R}$ and $\tilde{G}_i = \mathrm{poly}(\mathbb{R}, G^{+i}) \rtimes_T \{0\}$ for $i \geq 2$. Here G^{+i} denotes the group G_i with the shifted filtration $G_j^{+i} = G_{i+j}$. Note in particular that $\tilde{G}_{s+1}$ is trivial (noting $s \geq 1$).

Let $q \in \mathrm{poly}(\mathbb{R}, G)$ be a rescaled version of p, by $q(x) = p(xN)$. Now we define

$$\begin{aligned} \tilde{p} \colon \mathbb{Z} &\to \tilde{G} \\ x &\mapsto (\mathrm{id}_G, x/N) * (q, 0) \end{aligned}$$

which is certainly a polynomial map, and moreover is N-periodic since

$$\begin{aligned} \tilde{p}(x + N) &= (\mathrm{id}_G, x/N + 1) * (q, 0) \\ &= (\mathrm{id}_G, 1) * (\mathrm{id}_G, x/N) * (q, 0) \\ &\in \tilde{\Gamma}\, \tilde{p}(x) \ . \end{aligned}$$

Finally, we define $\tilde{F}$ on the fundamental domain $(\mathrm{poly}(\mathbb{Z}, \Gamma) \backslash \mathrm{poly}(\mathbb{R}, G)) \rtimes_T [0, 1)$ by

$$\tilde{F}(\mathrm{poly}(\mathbb{Z}, \Gamma)\, x, t) = \phi(t) \cdot F(\Gamma\, x(0))$$

and extend periodically by $\tilde{\Gamma}$. In other words,

$$\tilde{F}(\mathrm{poly}(\mathbb{Z}, \Gamma)\, x, t) = \phi(\{t\}) \cdot F(\Gamma\, x(-\lfloor t \rfloor))$$

for general t.

It is now straightforward to check that

$$\begin{aligned} \tilde{F}(\tilde{p}(x)) &= \tilde{F}(q(\cdot + x/N), x/N) \\ &= \phi(\{x/N\}) \cdot F(q(x/N - \lfloor x/N \rfloor)) \\ &= \phi(\{x/N\}) \cdot F(p(x \bmod N)) \end{aligned}$$

as required.

There remains only the statement that $\tilde{F}$ is $O_{K,G,\Gamma,\phi}(1)$-Lipschitz. For this to make sense, we need to equip $\mathrm{poly}(\mathbb{R}, G)$ and $\tilde{G}$ with left-invariant Riemannian metrics $d_{\mathrm{poly}(\mathbb{R},G)}$, $d_{\tilde{G}}$, which define corresponding metrics $d_{\mathrm{poly}(\mathbb{Z},\Gamma)\backslash \mathrm{poly}(\mathbb{R},G)}$, $d_{\tilde{\Gamma}\backslash\tilde{G}}$. There are various ways one could make these choices canonical; ultimately it doesn't matter as any two choices will be bi-Lipschitz on the respective compact manifolds.

It is certainly clear that $\tilde{F}$ is continuous: this is immediate except at points (t, x) where $t \in \mathbb{Z}$, but the support properties of ϕ guarantee continuity there also. Note also that the definition of $\tilde{F}$ depends only on F, G and Γ, not p or N.

Using compactness of $\tilde{\Gamma}\backslash\tilde{G}$, it is routine to adapt these observations to show $\tilde{F}$ is Lipschitz and to obtain the quantitative statement about the Lipschitz constant.

□

4.4 An example

We consider the case where $s = 2$ and the original nilsequence f is the function

$$\begin{aligned} f\colon \mathbb{Z} &\to \mathbb{C} \\ x &\mapsto e(\alpha x^2) \end{aligned}$$

for some real number α. Note that this is trivially a (polynomial) nilsequence by composing the polynomial map $p\colon \mathbb{Z} \to \mathbb{R}$, $p(x) = \alpha x^2$ with the automorphic function $F(t) = e(t)$.

In general, p fails to be N-periodic unless $\alpha \in \frac{1}{N}\mathbb{Z}$. Moreover, f can fail to correlate significantly with any N-periodic phase quadratic: i.e. (to prove Theorem 4.1.4) it would not suffice just to alter the value of α.

Theorem 4.1.5 requires us (essentially) to consider the function f' obtained by repeating f N-periodically, i.e. $f'(x) = e(\alpha N^2\{x/N\}^2)$. This is a bracket quadratic, and hence general machinery due to Bergelson and Leibman [3] tells us that it is – essentially – a nilsequence. With such a simple example we can make this very explicit. To simplify matters, we consider the case where $\alpha = a/N^2$ for some integer a; these examples *are* fairly dense in the whole space.

Let

$$G = \begin{pmatrix} 1 & \mathbb{R} & \mathbb{R} \\ 0 & 1 & \mathbb{R} \\ 0 & 0 & 1 \end{pmatrix}; \quad \Gamma = \begin{pmatrix} 1 & \mathbb{Z} & \mathbb{Z} \\ 0 & 1 & \mathbb{Z} \\ 0 & 0 & 1 \end{pmatrix}$$

so G is the usual Heisenberg group and Γ the standard lattice. We abbreviate the matrix $\left(\begin{smallmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{smallmatrix}\right)$ to (x, y, z). Now define

$$\begin{aligned} p\colon \mathbb{Z} &\to G \\ t &\mapsto (t/N, 2at/N, at^2/N^2) \end{aligned}$$

which is a polynomial map. Moreover, it is N-periodic, since

$$p(t+N) = \begin{pmatrix} 1 & t/N+1 & at^2/N^2+2at/N+a \\ 0 & 1 & 2at/N+2a \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 1 & a \\ 0 & 1 & 2a \\ 0 & 0 & 1 \end{pmatrix} p(t) \in \Gamma p(t) \ .$$

A fundamental domain for $\Gamma\backslash G$ is given by elements $(x, y, z) \in G$ with $x, y, z \in [0, 1)$, and the map sending a point in G to the corresponding point in the fundamental domain is $(x, y, z) \mapsto (\{x\}, \{y\}, \{z - y\lfloor x \rfloor\})$, so we define $F\colon G \to \mathbb{C}$ by $F(x, y, z) = \phi(x)e(z)$ for points in the fundamental domain, and extend Γ-periodically. Here $\phi\colon \mathbb{R}/\mathbb{Z} \to \mathbb{R}_{\geq 0}$ is our smooth function supported on $[0, 1/2]$. In other words, $F(x, y, z) = \phi(\{x\})e(z - y\lfloor x \rfloor)$. It is now straightforward to check that $F \circ p(t) = \phi(t/N)e(at^2/N^2)$ if $0 \leq t < N$, which together with the N-periodicity above gives the result.

It turns out that this construction is *not* exactly what comes out of the general machinery of Section 4.3. However, the two are closely, albeit slightly subtly, related. We will not run through the construction in detail, but sketch some of the features.

In the Heisenberg example above, the group G is 2-step nilpotent with a 2-step filtration. The group $\tilde{G}$ from the general construction turns out to be 3-step nilpotent *as a group*, even though we were careful to give it a 2-step filtration.

This apparent paradox is a consequence of the fact that the filtration of $\tilde{G}$ is not *proper*, i.e. $\tilde{G}_0 \neq \tilde{G}_1$. Hence all the 3-step behaviour of $\tilde{G}$ is hidden in $\tilde{G}_0$ where it does not significantly affect the filtration.[2]

Although it may seem illegal to use a 3-step nilpotent group when $s = 2$, this is in fact compatible with the definitions of [36]. Indeed, the same phenomenon arises when applying the construction in Appendix C of that paper.

In the special case that the image of $\tilde{p}$ is contained in $\tilde{G}_1$, we can restrict everything to $\tilde{G}_1$ equipped with the proper filtration $\tilde{G}_1 = \tilde{G}_1 \supseteq \tilde{G}_2 \supseteq \{\mathrm{id}\}$. Hence $\tilde{G}_1$ is a 2-step nilpotent group. Moreover, it turns out that this special case occurs whenever $\alpha = a/N^2$ as required above. We also find that $\tilde{G}_1$ is exactly the Heisenberg group, and indeed this process yields exactly the explicit construction we have just described.

For general α we do not have that the image of $\tilde{p}$ is contained in $\tilde{G}_1$; but it is contained in some *coset* of $\tilde{G}_1$. So, by shifting everything by a small element of $\tilde{G}$ we can move to the previous case and again identify $\tilde{F} \circ \tilde{p}$ with a Heisenberg nilsequence.

In fact, it is generally true that any non-proper nilsequence in the above sense can be identified with a proper one, but we will not prove such a result.

[2] A group with an improper s-step filtration need not even be nilpotent in general.

Chapter 5

Summing triples of bijections

In this chapter, we prove an asymptotic for the number of additive triples of bijections $\{1, \dots, n\} \to \mathbb{Z}/n\mathbb{Z}$, that is, the number of pairs of bijections $\pi_1, \pi_2 \colon \{1, \dots, n\} \to \mathbb{Z}/n\mathbb{Z}$ such that the pointwise sum $\pi_1 + \pi_2$ is also a bijection.

This problem is equivalent to counting the number of orthomorphisms or complete mappings of $\mathbb{Z}/n\mathbb{Z}$, to counting the number of arrangements of n mutually nonattacking semiqueens on an $n \times n$ toroidal chessboard, and to counting the number of transversals in a cyclic Latin square.

The method of proof is a version of the Hardy-Littlewood circle method from analytic number theory, adapted to the group $(\mathbb{Z}/n\mathbb{Z})^n$.

This chapter is based on joint work with Sean Eberhard and Rudi Mrazović [21].

5.1 Introduction

Let S be the set of bijections $\{1, \dots, n\} \to \mathbb{Z}/n\mathbb{Z}$, thought of as a subset of the group $(\mathbb{Z}/n\mathbb{Z})^n$. We are interested in counting the number s_n of additive triples in S, that is, the number of pairs of bijections $\pi_1, \pi_2 \colon \{1, \dots, n\} \to \mathbb{Z}/n\mathbb{Z}$ such that the pointwise sum $\pi_1 + \pi_2$ is also a bijection.

The number s_n has been studied somewhat extensively, but under a different guise. Since S is invariant under precomposition with permutations of $\{1, \dots, n\}$ it is easy to see by arbitrarily identifying $\{1, \dots, n\}$ with $\mathbb{Z}/n\mathbb{Z}$ that $s_n/n!$ is the number of permutations π of $\mathbb{Z}/n\mathbb{Z}$ such that $x \mapsto \pi(x) - x$ is also a permutation. Such maps are called orthomorphisms or complete mappings, and afford the following fun interpretation. Define a semiqueen to be a chess piece which can move any distance horizontally, vertically, or diagonally in the northeast-southwest direction. Then orthomorphisms represent ways of arranging n mutually nonattacking semiqueens on an $n \times n$ toroidal chessboard. Thus $s_n/n!$ is the number of such arrangements.

In another guise this problem is that of counting the number of transversals of the cyclic Latin square. Here a transversal of an $n \times n$ Latin square is a set of n squares with no two sharing the same row, column, or symbol, and the cyclic Latin square is the Latin square with (i, j) entry given by $i + j \pmod n$. Then as above the number of such transversals is $s_n/n!$.

If n is even then $s_n = 0$. Indeed, in this case for any bijection $\pi\colon \{1, \dots, n\} \to \mathbb{Z}/n\mathbb{Z}$ we have $\sum_{i=1}^{n} \pi(i) = n/2$, so the sum of two bijections is never again a bijection. If n is odd then it is easy to see that $s_n > 0$, but estimating s_n is not easy.

In 1991, Vardi [62] made a conjecture equivalent to the following.

Conjecture 5.1.1. *There are constants $c_1 > 0$ and $c_2 < 1$ such that for any large enough odd number n we have*

$$c_1^n n!^2 \leq s_n \leq c_2^n n!^2.$$

The upper bound in this conjecture is known, and various authors have made incremental improvements to the constant c_2. Cooper and Kovalenko [20] showed that $c_2 = e^{-0.08854}$ is acceptable, and this was later improved by Kovalenko [38] to $c_2 = 1/\sqrt{2}$ and by McKay, McLeod and Wanless [46] to $c_2 = 0.614$. Recently, Taranenko [60] considered the more general problem of bounding the number of transversals in an arbitrary Latin square, and a consequence of his bound is that one can take $c_2 = 1/e + o(1)$. Glebov and Luria [29] proved the same bound using a somewhat simpler method based on entropy.

There has been much less progress on the lower bound. Nontrivial lower bounds for s_n have been proved under various arithmetic assumptions about n. For example, Cooper [18] proved that $s_n \geq e^{\frac{1}{2}\sqrt{n}\log n} n!$ under the hypothesis that n has a divisor of size roughly $\sqrt{n}$, while if, say, n is prime then a lower bound of Rivin, Vardi, and Zimmermann [48] for the torodial queens problem gives $s_n \geq 2^{\sqrt{(n-1)/2}} n!$. These lower bounds were recently superseded by work of Cavenagh and Wanless [13], which showed that $s_n \geq 3.246^n n!$ for all odd n. However, this lower bound is still a long way from the one in Conjecture 5.1.1.

Some researchers have also tried investigating the growth rate of s_n numerically: see Cooper, Gilchrist, Kovalenko, and Novaković [19] and Kuznetsov [39], [40], [41]. Finally, we should also mention that $s_n/(n \cdot n!)$ is sequence A003111 in [55].

In this chapter we prove Vardi's conjecture with the optimal values $c_1, c_2 = 1/e + o(1)$. In fact our estimate is much more precise – we compute s_n up to a factor of $1 + o(1)$.

Theorem 5.1.2. *Let n be an odd integer. Then*

$$s_n = (e^{-1/2} + o(1))n!^3/n^{n-1}.$$

The most surprising feature of this estimate is the appearance of the constant $e^{-1/2}$. This constant arises as the sum of the singular series in an argument resembling the circle method from analytic number theory, but it can be rationalized heuristically as follows. If $\pi_1, \pi_2 \colon \{1, \dots, n\} \to \mathbb{Z}/n\mathbb{Z}$ are random bijections then the sum $f = \pi_1 + \pi_2$ is something like a random function subject to $\sum_{x \in \mathbb{Z}/n\mathbb{Z}} f(x) = 0$, and so we might guess that the probability $\pi_1 + \pi_2$ is a bijection is about $n \cdot n!/n^n$. However if we fix two elements $x, y \in \mathbb{Z}/n\mathbb{Z}$ then the difference

$$f(x) - f(y) = (\pi_1(x) - \pi_1(y)) + (\pi_2(x) - \pi_2(y))$$

is the sum of two uniformly random nonzero elements, so $f(x) = f(y)$ with probability $1/(n-1)$, not $1/n$. Thus the probability that $f(x) \neq f(y)$ is smaller than we previously suggested by a factor of

$$\frac{1 - 1/(n-1)}{1 - 1/n} = 1 - 1/n^2 + O(1/n^3).$$

There are a total of $n^2/2 + O(n)$ pairs x, y, so we might guess that the probability $\pi_1 + \pi_2$ is a bijection is more like

$$\left(1 - 1/n^2\right)^{n^2/2} n \cdot n!/n^n \approx e^{-1/2} n!/n^{n-1}.$$

Our proof applies with only notational modifications when $\mathbb{Z}/n\mathbb{Z}$ is replaced by any abelian group G of odd order. There are additional complications however when G is allowed to have even order, say when $G = (\mathbb{Z}/2\mathbb{Z})^d$, and we do not know whether the same method can be made to work in this case.

Like the toroidal semiqueens problem, the toroidal queens problem – the problem of counting the number of ways of arranging n mutually nonattacking queens on an $n \times n$ toroidal chessboard – is equivalent to counting the number of solutions to a particular linear system in the set of bijections, namely the system

$$\pi_1 + \pi_2 = \pi_3,$$
$$\pi_1 - \pi_2 = \pi_4.$$

Linear systems of complexity 2 like this one are known not to be controlled by Fourier analysis alone, but there is some hope that one could use the higher-order theory pioneered by Gowers. This is an interesting avenue which has not yet been fully explored.

5.1.1 Notation

In addition to the standard notation discussed above, we use a primed sum $\sum'_{s_1,\dots,s_k \in S}$ to denote the sum over all k-tuples of distinct elements $s_1, \dots, s_k \in S$.

5.2 Outline of the proof

Our approach to proving Theorem 5.1.2 is Fourier-analytic. Write $G = \mathbb{Z}/n\mathbb{Z}$ for n an odd integer, and write $S \subset G^n$ for the set of bijections $\{1, \dots, n\} \to G$. Our goal is to compute the quantity

$$\mathbb{E}_{x,y \in G^n} 1_S(x) 1_S(y) 1_S(x+y) = \langle 1_S * 1_S, 1_S \rangle.$$

A standard application of Parseval's identity shows that this quantity can be expressed in terms of the Fourier transform of 1_S as

$$\sum_{\chi \in \widehat{G^n}} |\widehat{1_S}(\chi)|^2 \widehat{1_S}(\chi). \tag{5.2.1}$$

Here we identify $\widehat{G}$ with $\frac{1}{n}\mathbb{Z}/\mathbb{Z}$ in the usual way, so that for $\chi = (r_1, \dots, r_n) \in (\frac{1}{n}\mathbb{Z}/\mathbb{Z})^n$ we have explicitly

$$\widehat{1_S}(r_1, \dots, r_n) = \frac{1}{n^n} \sum_{x_1,\dots,x_n}{}' e(-r_1 x_1 - \cdots - r_n x_n), \tag{5.2.2}$$

where, as mentioned in the previous section, we use $\sum'$ to denote the sum over distinct $x_1, \dots, x_n \in G$, i.e., the sum over S. In fact it is clear that $\widehat{1_S}$ is real-valued, since $S = -S$, and hence we may drop the absolute value signs in (5.2.1).

The form of the proof can then be viewed as an analogue of the Hardy-Littlewood circle method in analytic number theory, adapted to the group G^n rather than $\mathbb{Z}$. We borrow some nomenclature from the classical setting.

- There are a small number of characters $\chi \in \widehat{G^n}$, namely the characters $\chi = (r_1, \dots, r_n)$ for which almost all of the r_i are equal, that make a substantial contribution to the sum (5.2.1). We call the totality of these χ the *major arcs*. We compute explicitly the contribution of these χ to (5.2.1), up to small errors: this is an analogue of the singular series, which accounts for the main term in Theorem 5.1.2, including the unexpected constant $e^{-1/2}$. For all this see Section 5.3.

- We bound the contribution from all other characters using the triangle inequality; that is, we obtain an upper bound for

$$\sum_{\text{all other } \chi} |\widehat{1_S}(\chi)|^3$$

that is smaller than the main term. We call such χ collectively the *minor arcs*.

A large part of this chapter is therefore devoted to obtaining good bounds for Fourier coefficients $\widehat{1_S}(\chi)$, either pointwise or on average in a suitable sense. In fact we will need to combine several different arguments which are effective in different regimes.

5.2.1 Preliminaries on $\widehat{1_S}$

In order to describe how the characters $\widehat{G}^n$ are divided up into different pieces in which different approaches will be effective, we will need some preliminary remarks.

First, note that the function $\widehat{1_S}(r_1, \dots, r_n)$ is invariant under permutation of the r_i:

$$\widehat{1_S}(r_1, \dots, r_n) = \widehat{1_S}(r_{\sigma(1)}, \dots, r_{\sigma(n)})$$

for any bijection $\sigma\colon \{1, \dots, n\} \to \{1, \dots, n\}$. This is immediate from the definition (5.2.2). Hence it makes sense to specify a Fourier coefficient of interest in the form $\widehat{1_S}(r_1^{a_1}, \dots, r_k^{a_k})$, where $r_i \in \widehat{G}$ are distinct, the a_i are positive integers such that $\sum a_i = n$, and the notation r^a means r repeated a times.

Second, observe that $\widehat{1_S}(r_1, \dots, r_n) = 0$ unless $\sum r_i = 0$. This follows from the fact that S is invariant under global shifts $(x_i) \mapsto (x_i + t)$, so one can compute

$$\begin{aligned}
\widehat{1_S}(r_1, \dots, r_n) &= \frac{1}{n^n} \sum_{x_1, \dots, x_n}{}' e(-r_1 x_1 - \cdots - r_n x_n) \\
&= \frac{1}{n^n} \sum_{x_1, \dots, x_n}{}' e(-r_1(x_1 + t) - \cdots - r_n(x_n + t)) \\
&= \widehat{1_S}(r_1, \dots, r_n) e(-(r_1 + \cdots + r_n)t).
\end{aligned}$$

Dually, note that $\widehat{1_S}$ is invariant under global shifts. This follows from the fact that $\sum x_i = 0$ for any $(x_i) \in S$. Hence,

$$\begin{aligned}
\widehat{1_S}(r_1, \dots, r_n) &= \frac{1}{n^n} \sum_{x_1, \dots, x_n}{}' e(-r_1 x_1 - \cdots - r_n x_n) \\
&= \frac{1}{n^n} \sum_{x_1, \dots, x_n}{}' e(-(r_1 + t)x_1 - \cdots - (r_n + t)x_n) \\
&= \widehat{1_S}(r_1 + t, \dots, r_n + t).
\end{aligned}$$

Finally, note the trivial bound

$$|\widehat{1_S}(\chi)| \leq \widehat{1_S}(0) = \frac{n!}{n^n}.$$

5.2.2 Entropy ranges for minor arcs

We now explain a straightforward but still fairly powerful bound on $|\widehat{1_S}(\chi)|$ that follows directly from these elementary considerations.

Proposition 5.2.1 (Entropy bound)**.** *We have*

$$|\widehat{1_S}(r_1^{a_1}, \dots, r_k^{a_k})| \leq \binom{n}{a_1, \dots, a_k}^{-1/2} \left(\frac{n!}{n^n}\right)^{1/2}.$$

Proof. Let $\mathcal{O} \subset \widehat{G}^n$ denote the set of characters obtained by permuting the elements of $(r_1^{a_1}, \dots, r_k^{a_k})$. Hence, $|\mathcal{O}| = \binom{n}{a_1,\dots,a_k}$ and by permutation invariance, $\widehat{1_S}$ takes a constant value on $\mathcal{O}$. Thus by Parseval,

$$\begin{aligned}\frac{n!}{n^n} &= \sum_{\chi \in \widehat{G}^n} |\widehat{1_S}(\chi)|^2 \\ &\geq \sum_{\chi \in \mathcal{O}} |\widehat{1_S}(\chi)|^2 \\ &= |\mathcal{O}||\widehat{1_S}(r_1^{a_1}, \dots, r_k^{a_k})|^2,\end{aligned}$$

and the result follows. □

The term "entropy bound" refers to the fact that the quantity $H(\chi) = \frac{1}{n}\log\binom{n}{a_1,\dots,a_k}$ is roughly the entropy $\sum_{i=1}^k (a_i/n)\log(n/a_i) \in [0, \log n]$ of a random variable taking the value r_i with probability a_i/n. In a slight abuse of nomenclature, we refer to this first quantity $H(\chi)$ as the *entropy of* χ. In this language, the bound of Proposition 5.2.1 is precisely $\exp(-Hn/2)(n!/n^n)^{1/2}$ or roughly $\exp(-Hn/2 - n/2)$.

This is already sufficient to control the contribution to (5.2.1) for most characters χ.

Corollary 5.2.2. *We have*

$$\sum_{\chi:\, H(\chi) \geq R} |\widehat{1_S}(\chi)|^3 \leq \exp((3-R)n/2)\left(\frac{n!}{n^n}\right)^3.$$

Proof. By Parseval and Proposition 5.2.1,

$$\begin{aligned}\sum_{\chi\colon H(\chi)\geq R}|\widehat{1_S}(\chi)|^3 &\leq \left(\sum_{\chi\colon H(\chi)\geq R}|\widehat{1_S}(\chi)|^2\right)\sup_{\chi\colon H(\chi)\geq R}|\widehat{1_S}(\chi)| \\ &\leq \left(\sum_{\chi\in\widehat{G^n}}|\widehat{1_S}(\chi)|^2\right)\sup_{\chi\colon H(\chi)\geq R}\exp(-H(\chi)n/2)\left(\frac{n!}{n^n}\right)^{1/2} \\ &\leq \exp(-Rn/2)\left(\frac{n!}{n^n}\right)^{3/2} \\ &\leq \exp((3-R)n/2)\left(\frac{n!}{n^n}\right)^{3}\end{aligned}$$

as required, where we have used the estimate $n! > (n/e)^n$. □

This is smaller than the main term for any fixed $R > 3$, so from now on such high-entropy characters need not concern us. We refer to such characters as the *high-entropy minor arcs*. Since a typical character χ has entropy comparable to $\log n$, the high-entropy case comprises almost all characters in some sense, so this is a good first step.

On the opposite extreme we have characters χ with entropy $o(1)$, such as $\chi = \left(r^k, -r^k, 0^{n-2k}\right)$ for $k = o(n)$. The characters with entropy $O\left(\frac{\log n}{n}\right)$ are the major arcs, but between $O\left(\frac{\log n}{n}\right)$ and $o(1)$ we have the *low-entropy minor arcs*. Such characters are necessarily of the form $(\dots, r^{n-o(n)})$, i.e., they take a single value almost all the time. By the remarks above we may shift without loss of generality to assume that this value r is zero. Thus the low-entropy minor arcs are closely related to the set of *sparse* characters, characters χ comprised almost entirely of zeros.

That leaves the characters χ with $o(1) \leq H(\chi) \leq 3$, which form the *medium-entropy minor arcs*. A good model case are characters such as $\chi = (r^{n/3}, -r^{n/3}, 0^{n/3})$, which are particularly troublesome.

There are two fundamental areas of inefficiency in the arguments of Proposition 5.2.1 and Corollary 5.2.2 that prevent them from giving good bounds for low- or medium-entropy minor arcs. First, the bound on $|\widehat{1_S}|$ given by Proposition 5.2.1 is less effective for smaller H, and in particular worse than trivial when $H < 1$. Clearly there is no hope for this programme unless we can find a bound on $|\widehat{1_S}|$ that is nontrivial throughout the range $0 < H < 1$, and moreover obtains a substantial exponential saving for most of that range.

Second, in the proof of Corollary 5.2.2 we made use of the convenient bound

$$\sum_{\chi \in X} |\widehat{1_S}(\chi)|^3 \leq \left(\sum_{\chi \in \widehat{G}} |\widehat{1_S}(\chi)|^2 \right) \sup_{\chi \in X} |\widehat{1_S}(\chi)|$$

for some appropriate set $X \subset \widehat{G}^n$. It is fairly clear that we cannot afford this luxury for small entropies: since the main term has order $(n!/n^n)^3$, we need $|\widehat{1_S}(\chi)| = o((n!/n^n)^2)$ for all $\chi \in X$ for this to be effective. This is a saving of about $\exp(-n)$ over the trivial bound of $n!/n^n$, and it is simply not reasonable to expect such a bound to hold if, say, $H(\chi) = 1/1000$. One way around this is to pigeonhole the characters χ into various sets X_i, and apply such an $L^2 \cdot L^\infty$ bound on each set. To make this work, we would need a bound on

$$\sum_{\chi \in X_i} |\widehat{1_S}(\chi)|^2$$

which makes a significant saving over the crude estimate

$$\sum_{\chi \in X_i} |\widehat{1_S}(\chi)|^2 \leq \sum_{\chi \in \widehat{G}^n} |\widehat{1_S}(\chi)|^2$$

employed in Corollary 5.2.2. Alternatively, one could try to bound the L^3 sum

$$\sum_{\chi \in X_i} |\widehat{1_S}(\chi)|^3$$

directly, using an L^∞ bound and an estimate for the size of X_i.

We address both of these inefficiencies in order to reach our final bound. Specifically, we will do each of the following.

- Improving on Proposition 5.2.1, we obtain a good general-purpose bound for $|\widehat{1_S}(\chi)|$ that is nontrivial for values of $H(\chi)$ approaching zero. Roughly speaking, where Proposition 5.2.1 gives the bound $e^{-Hn/2-n/2}$ we will prove the bound $e^{-Hn/2-n}$. This appears in Section 5.4.
- We use this bound to estimate the contribution to (5.2.1) from χ with $10^{-10} < H(\chi) < 10$, say, by using a dyadic decomposition and a simple estimate for the number of characters of a given entropy. Thus we dispatch the medium-entropy minor arcs. This is covered in Section 5.7.
- Separately, we obtain a good estimate for

$$\sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^2,$$

i.e., an improvement to Parseval when summing only over those characters χ having exactly m nonzero terms. Here we might imagine $m \leq n/1000$. We call this a *sparseval* bound, and prove this in Section 5.5.

- Finally, we obtain a slightly different L^∞ bound specialized to the case of sparse characters χ. This appears in Section 5.6. The total contribution from the low-entropy minor arcs is controlled by combining this with the L^2 bound above.

In summary, we split the universe of all χ into several slightly overlapping regions – major arcs, low-entropy minor arcs, medium-entropy minor arcs, and large-entropy minor arcs – and apply a cocktail of different bounds and explicit computations adapted to each region.

5.2.3 Interpretations of minor arc bounds

As we have said, a significant part of our effort will be spent proving nontrivial bounds on Fourier coefficients $|\widehat{1_S}(\chi)|$ for χ in the minor arcs. Although such bounds are at first glance fairly esoteric, in fact they have natural interpretations. For instance, they are intimately related to questions of the following flavour.

Question 5.2.3. *Suppose the elements of $G = \mathbb{Z}/n\mathbb{Z}$ are divided into k sets $A_1, \dots, A_k$, each of a pre-determined size $a_i \approx n/k$, at random. Consider the random variables*

$$T_i = \sum_{r \in A_i} r \in G$$

for $i = 1, \dots, k$. The vector $T = (T_1, \dots, T_k)$ is a random variable taking values in the subgroup $H = \{(x_i) \in G^k : \sum x_i = 0\}$ of G^k. How close is T to being equidistributed on H, in a quantitative sense? In other words how close is the law of T to the uniform distribution on H?

Because of the connection of this question to the size of the Fourier coefficients $\widehat{1_S}(r_1^{a_1}, \dots, r_k^{a_k})$, our bounds give a strong answer to this question in many regimes. It is not inconceivable that these kind of bounds may find applications elsewhere.

5.3 Major arcs

In this section we compute $\widehat{1_S}(r_1, \dots, r_m, 0, \dots, 0)$ explicitly whenever m is bounded. To this end observe that

$$\widehat{1_S}(r_1, \dots, r_m, 0, \dots, 0) = \frac{(n-m)!}{n^n} \sum_{x_1, \dots, x_m}{}' e(-r_1 x_1 - \cdots - r_m x_m). \tag{5.3.1}$$

The sum over distinct $x_1, \ldots, x_m$ can be related to a sum over partitions of $\{1, \ldots, m\}$ using a type of Möbius inversion: for any function $F(x_1, \ldots, x_m)$ we have

$$\sum_{x_1,\ldots,x_m}{}' F(x_1, \ldots, x_m) = \sum_{\mathcal{P}} \mu(\mathcal{P}) \sum_{\substack{x_1,\ldots,x_m \\ x_i = x_j \text{ whenever } i \overset{\mathcal{P}}{\sim} j}} F(x_1, \ldots, x_m),$$

where the outer sum runs over partitions $\mathcal{P}$ of $\{1, \ldots, m\}$, and

$$\mu(\mathcal{P}) = (-1)^{m-|\mathcal{P}|} \prod_{P \in \mathcal{P}} (|P| - 1)!.$$

See, e.g., [7] for more information. To apply this to $\widehat{1_S}$, say that a partition $\mathcal{P}$ of $\{1, \ldots, m\}$ *kills* $(r_1, \ldots, r_m)$ if $\sum_{i \in P} r_i = 0$ for each $P \in \mathcal{P}$, and observe that

$$\sum_{\substack{x_1,\ldots,x_m \\ x_i = x_j \text{ whenever } i \overset{\mathcal{P}}{\sim} j}} e(-r_1 x_1 - \cdots - r_m x_m) = \begin{cases} n^{|\mathcal{P}|} & \text{if } \mathcal{P} \text{ kills } (r_1, \ldots, r_m), \\ 0 & \text{otherwise,} \end{cases}$$

so

$$\widehat{1_S}(r_1, \ldots, r_m, 0, \ldots, 0) = \frac{(n-m)!}{n^n} \sum_{\mathcal{P} \text{ killing } (r_1,\ldots,r_m)} \mu(\mathcal{P}) n^{|\mathcal{P}|}. \tag{5.3.2}$$

The following two observations are immediate from this formula:

- Suppose every killing partition of $(r_1, \ldots, r_m)$ has at most k parts. Then

$$|\widehat{1_S}(r_1, \ldots, r_m, 0, \ldots, 0)| \leq O_m\left(\frac{1}{n^{m-k}} \frac{n!}{n^n}\right). \tag{5.3.3}$$

- Suppose that m is even, that $r_1, \ldots, r_m \in \widehat{G}$ are nonzero, and that $(r_1, \ldots, r_m)$ is killed by a unique partition with $m/2$ parts. In other words suppose that $r_1, \ldots, r_m \in \widehat{G}$ are distinct and nonzero, and that, up to a permutation, $r_{2j} = -r_{2j-1}$ for each $j = 1, \ldots, m/2$. Then

$$\widehat{1_S}(r_1, \ldots, r_m, 0, \ldots, 0) = \frac{(-1)^{m/2}}{n^{m/2}} \frac{n!}{n^n} + O_m\left(\frac{1}{n^{m/2+1}} \frac{n!}{n^n}\right). \tag{5.3.4}$$

Proposition 5.3.1. *If m is even then*

$$\sum_{m\text{-sparse } \chi} \widehat{1_S}(\chi)^3 = \frac{(-1)^{m/2}}{2^{m/2}(m/2)!} \frac{n!^3}{n^{3n}} + O_m\left(\frac{1}{n} \frac{n!^3}{n^{3n}}\right),$$

while if m is odd then

$$\sum_{m\text{-sparse } \chi} \widehat{1_S}(\chi)^3 = O_m\left(\frac{1}{n} \frac{n!^3}{n^{3n}}\right).$$

Proof. First of all note by permutation-invariance that

$$\sum_{m\text{-sparse }\chi} \widehat{1_S}(\chi)^3 = \binom{n}{m} \sum_{r_1,\dots,r_m\neq 0} \widehat{1_S}(r_1,\dots,r_m,0,\dots,0)^3. \tag{5.3.5}$$

For each $(r_1,\dots,r_m)$ choose a maximal-size partition $\mathcal{P}$ which kills $(r_1,\dots,r_m)$, and split the sum up according to $\mathcal{P}$. Suppose $\mathcal{P}$ has k parts. Since $\mathcal{P}$ cannot have singletons we have $k \le m/2$, and by (5.3.3) we have

$$|\widehat{1_S}(r_1,\dots,r_m,0,\dots,0)| \le O_m\left(\frac{1}{n^{m-k}}\frac{n!}{n^n}\right).$$

Since the number of $(r_1,\dots,r_m)$ killed by $\mathcal{P}$ is bounded by n^{m-k}, the contribution to (5.3.5) from these $(r_1,\dots,r_m)$ is bounded by

$$\binom{n}{m} n^{m-k}\left(\frac{O_m(1)}{n^{m-k}}\frac{n!}{n^n}\right)^3 = O_m\left(\frac{1}{n^{m-2k}}\frac{n!^3}{n^{3n}}\right),$$

which is satisfactory unless $k = m/2$. Moreover the number of $(r_1,\dots,r_m)$ killed by at least two partitions $\mathcal{P}$ with $m/2$ parts is bounded by $n^{m/2-1}$, so the contribution from these $(r_1,\dots,r_m)$ is bounded by

$$\binom{n}{m} n^{m/2-1} O_m\left(\frac{1}{n^{m/2}}\frac{n!}{n^n}\right)^3 = O_m\left(\frac{1}{n}\frac{n!^3}{n^{3n}}\right),$$

which is again satisfactory. Thus we may restrict our attention to those $(r_1,\dots,r_m)$ which are killed by a unique partition $\mathcal{P}$ with $m/2$ parts, and for these (5.3.4) gives

$$\widehat{1_S}(r_1,\dots,r_m,0,\dots,0) = \frac{(-1)^{m/2}}{n^{m/2}}\frac{n!}{n^n} + O_m\left(\frac{1}{n^{m/2+1}}\frac{n!}{n^n}\right).$$

For a fixed such $\mathcal{P}$ the number of such $(r_1,\dots,r_m)$ is $n^{m/2} + O_m(n^{m/2-1})$, and the number of such $\mathcal{P}$ is

$$(m-1)(m-3)\cdots 1 = \frac{m!}{2^{m/2}(m/2)!},$$

so the total contribution from these $(r_1,\dots,r_m)$ is

$$\begin{aligned}\binom{n}{m}\frac{m!}{2^{m/2}(m/2)!}(n^{m/2}+O_m(n^{m/2-1}))&\left(\frac{(-1)^{m/2}}{n^{m/2}}\frac{n!}{n^n} + O_m\left(\frac{1}{n^{m/2+1}}\frac{n!}{n^n}\right)\right)^3 \\ &= \frac{(-1)^{m/2}}{2^{m/2}(m/2)!}\frac{n!^3}{n^{3n}} + O_m\left(\frac{1}{n}\frac{n!^3}{n^{3n}}\right).\end{aligned}$$

This proves the proposition. □

5.4 Square-root cancellation for general Fourier coefficients

The aim of this section is to prove the following refinement of Proposition 5.2.1.

Theorem 5.4.1. *Suppose $\chi = (r_1^{a_1}, \dots, r_k^{a_k})$, where $\sum_{i=1}^k a_i = n$. Then*

$$|\widehat{1_S}(\chi)| \leq \binom{n+k-1}{k-1}^{1/2} \binom{n}{a_1, \dots, a_k}^{-1/2} \frac{n!}{n^n}.$$

Remark 5.4.2.

- The terms depending on k are easily seen to be a small error in the regimes we care about. Indeed, if $H(\chi) = O(1)$ then χ can take at most $O(n/\log n) = o(n)$ distinct values, and the additional term is then $\exp(o(n))$, and thus small compared to the exponential entropy saving.
- The saving over Proposition 5.2.1 is an apparently modest factor of around $e^{-n/2+o(n)}$. This may be uninspiring if H is large, but if H is small it is decisive.
- In general this bound is not sharp, often by a substantial amount. However, it is sometimes attained asymptotically for very special, arithmetically structured classes of characters. For instance, if we temporarily allow n to be even, then one can compute directly that

  $$\log \left[|\widehat{1_S}((1/2)^a, 0^{n-a})|/(n!/n^n) \right] \sim -\frac{1}{2} \log \binom{n}{a}$$

  for large n. Similarly, numerical evidence strongly suggests that

  $$\log \left[|\widehat{1_S}((1/3)^a, (-1/3)^a, 0^{n-2a})|/(n!/n^n) \right] \sim -\frac{1}{2} \log \binom{n}{a, a, n-2a}$$

  if n is divisible by 3; for example, when $n = 3003$,

  $$\log \left[|\widehat{1_S}((1/3)^{1001}, (-1/3)^{1001}, 0^{1001})|/(n!/n^n) \right] = -1649.01782245...,$$

  while

  $$-\frac{1}{2} \log \binom{3003}{1001, 1001, 1001} = -1645.46757758....$$

  But, for example, the evidence also suggests that

  $$\log \left[|\widehat{1_S}((1/3)^a, 0^{n-a})|/(n!/n^n) \right] \sim -\frac{2}{3} \log \binom{n}{a}.$$

From the above remarks, we see that one might hope to improve on Theorem 5.4.1, but only by ruling out the particular characters discussed above and others similar to them. For instance, a strengthening is almost certainly true under the assumption that n is prime. Thus one might think of Theorem 5.4.1 as the best general-purpose bound available.

We should also comment briefly on the term "square-root cancellation". If $\chi = (r_1^{a_1}, \dots, r_k^{a_k})$ then

$$\widehat{1_S}(\chi) = \frac{n!}{n^n} \mathop{\mathbb{E}}_{\text{partitions } \mathcal{P}} e\left(-r_1\left(\sum_{x\in P_1} x\right) - \cdots - r_k\left(\sum_{x\in P_k} x\right)\right)$$

where the expectation is over all ordered partitions $\mathcal{P} = (P_1, \dots, P_k)$ of G into k pieces with $|P_i| = a_i$. The number of such partitions is precisely $\binom{n}{a_1,\dots,a_k}$; assuming that the phases in the average behave randomly then we should expect square-root cancellation, and we heuristically recover the bound in Theorem 5.4.1 up to lower-order terms.

Proof of Theorem 5.4.1. Let γ be the Gaussian measure on $\mathbb{C}^k$ defined with respect to Lebesgue measure λ by

$$\frac{d\gamma}{d\lambda} = \frac{1}{\pi^k} \exp\left(-\sum_{i=1}^k |z_i|^2\right).$$

We claim that

$$n^n \widehat{1_S}(\chi) = \int_{\mathbb{C}^k} z_1^{a_1} \cdots z_k^{a_k} \prod_{x\in G} \left(\sum_{i=1}^k e(-r_i x)\overline{z_i}\right) d\gamma. \tag{5.4.1}$$

One can check this by expanding the product and using the identity

$$\int_{\mathbb{C}^k} \prod_{i=1}^k z_i^{a_i}\overline{z_i}^{b_i}\, d\gamma = \begin{cases} \prod_{i=1}^k a_i! & \text{if } a_i = b_i \text{ for each } i, \\ 0 & \text{otherwise.} \end{cases} \tag{5.4.2}$$

We can now proceed by bounding the integral in (5.4.1) using various techniques. Applying the Cauchy–Schwarz inequality we bound the right hand side by

$$\left(\int_{\mathbb{C}^k} |z_1|^{2a_1} \cdots |z_k|^{2a_k}\, d\gamma\right)^{1/2} \left(\int_{\mathbb{C}^k} \prod_{x\in G} \left|\sum_{i=1}^k e(-r_i x)\overline{z_i}\right|^2 d\gamma\right)^{1/2}. \tag{5.4.3}$$

The second factor here is exactly $\left(\prod_{i=1}^k a_i!\right)^{1/2}$ by (5.4.2). If we now apply the AM–GM inequality to the last term and evaluate the resulting integral, we get

$$\begin{aligned}\int_{\mathbb{C}^k} \prod_{x\in G}\left|\sum_{i=1}^k e(-r_i x)\overline{z_i}\right|^2 d\gamma &\le \int_{\mathbb{C}^k}\left(\frac{1}{n}\sum_{x\in G}\left|\sum_{i=1}^k e(-r_i x)\overline{z_i}\right|^2\right)^n d\gamma\\ &= \int_{\mathbb{C}^k}\left(\sum_{i=1}^k |z_i|^2\right)^n d\gamma\\ &= \sum_{s_1+\cdots+s_k=n}\binom{n}{s_1,\dots,s_k}\int_{\mathbb{C}^k}|z_1|^{2s_1}\cdots|z_k|^{2s_k}\,d\gamma\\ &= n!\sum_{s_1+\cdots+s_k=n} 1\\ &= n!\binom{n+k-1}{k-1}.\end{aligned}$$

Thus the theorem follows by dividing (5.4.3) by n^n. □

Remark 5.4.3. Though the identity (5.4.1) is easily verified directly, we briefly sketch here a possible route by which one might be motivated to write down such a formula in the first place.

Let $\{X_x : x \in G\}$ be indeterminates, and observe that the left hand side of (5.4.1) is precisely the coefficient of $\prod_{x\in G} X_x$ in the multivariate polynomial

$$\prod_{i=1}^k\left(\sum_{x\in G} X_x e(-r_i x)\right)^{a_i},$$

and so equivalently the value of the derivative

$$\left(\prod_{x\in G}\frac{\partial}{\partial X_x}\right)\prod_{i=1}^k\left(\sum_{x\in G} X_x e(-r_i x)\right)^{a_i} \tag{5.4.4}$$

evaluated at zero (although in fact this evaluation is redundant as this expression is a constant function).

In some generality, it is possible to equate a quantity of the form

$$\frac{\partial^{b_1}}{\partial Y_1^{b_1}}\cdots\frac{\partial^{b_k}}{\partial Y_k^{b_k}} f\big|_{Y_1=\cdots=Y_k=0},$$

where $f(Y_1,\dots,Y_k)$ is a suitable holomorphic function of k complex variables, with the corresponding integral expression

$$\int_{\mathbb{C}^k}\overline{Y_1^{b_1}}\dots\overline{Y_k^{b_k}}\, f(Y_1,\dots,Y_k)\,d\gamma,$$

where again γ denotes Gaussian measure. This can be verified by applying the case $k = 1$ iteratively, which in turn follows by averaging the usual Cauchy integral formula over different radii. This formula can be considered a particular higher-dimensional variant of the Cauchy integral formula.

The identity (5.4.1) follows by applying this identity to (5.4.4) and making an orthogonal change of variables.

5.5 Sparseval

We recall from Section 5.2 that one of our goals is to obtain good bounds on

$$\sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^2, \tag{5.5.1}$$

where the sum is over all characters $\chi = (r_1, \dots, r_n)$ with precisely m nonzero entries.

First consider the related sum

$$\sum_{\leq m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^2, \tag{5.5.2}$$

i.e., the sum over characters with *at most* m non-zero entries. This can be bounded by applying Parseval to the set S_m of injections $\{1, \dots, m\} \to G$ as a subset of the group G^m. Indeed, letting $N(\chi)$ be the set of indices i for which r_i is nonzero, we have

$$\begin{aligned}
\sum_{\leq m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^2 &\leq \sum_{|N|=m} \sum_{N(\chi)\subset N} |\widehat{1_S}(\chi)|^2 \\
&= \binom{n}{m} \sum_{r_1,\dots,r_m} |\widehat{1_S}(r_1, \dots, r_m, 0, \dots, 0)|^2 \\
&= \binom{n}{m} \sum_{r_1,\dots,r_m} \frac{(n-m)!^2}{n^{2n-2m}} |\widehat{1_{S_m}}(r_1, \dots, r_m)|^2 \\
&= \binom{n}{m} \frac{(n-m)!^2}{n^{2n-2m}} \frac{n!}{n^m (n-m)!} \\
&= \frac{n^m}{m!} \frac{n!^2}{n^{2n}}.
\end{aligned} \tag{5.5.3}$$

Here we used (5.3.1).

This is our first nontrivial *sparseval* bound. It turns out that overcounting as above is too inefficient in the context of the wider argument, when, say, m/n is a small constant. The purpose of the rest of this section therefore is to improve the bound on (5.5.1) as much as possible.

Theorem 5.5.1. *If $m \le n/2$ then*

$$\sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^2 \le O(m^{1/4}) e^{O(m^{3/2}/n^{1/2})} \binom{n}{m}^{1/2} \frac{n!^2}{n^{2n}}$$

Proof. The starting point of our strategy is to apply inclusion–exclusion to obtain an expression for (5.5.1).

For $m \le n$ let

$$Q(m,n) = \frac{n^{2n}}{n!^2} \sum_{m\text{-sparse } \chi} |\widehat{1_S}(r)|^2,$$

and observe from (5.5.3) that

$$\begin{aligned}
\frac{n^m}{m!} &= \frac{n^{2n}}{n!^2} \sum_{|N|=m} \sum_{N(\chi)\subset N} |\widehat{1_S}(\chi)|^2 \\
&= \frac{n^{2n}}{n!^2} \sum_{k=0}^{m} \sum_{|N(\chi)|=k} |\widehat{1_S}(\chi)|^2 \binom{n-k}{m-k} \\
&= \sum_{k=0}^{m} \binom{n-k}{m-k} Q(k,n).
\end{aligned}$$

By inverting this relation we have

$$\begin{aligned}
Q(m,n) &= \sum_{k=0}^{m} (-1)^{m-k} \binom{n-k}{m-k} \frac{n^k}{k!} \\
&= \frac{1}{m!} \sum_{k=0}^{m} \binom{m}{k} (-1)^k (n-m+1)\cdots(n-k) n^k. \qquad (5.5.4)
\end{aligned}$$

This expression exhibits a vast amount of cancellation. To exploit this, we employ a generating function argument, followed by some complex analysis in the spirit of the saddle-point method (see [23, Chapter VIII] for background); see also Remark 5.5.2 for some further intuition.

Observe from (5.5.4) that $Q(m,n)$ is exactly the coefficient of X^m in

$$f(X) = \frac{n^{n+1} e^X}{(n+X)^{n-m+1}}. \qquad (5.5.5)$$

Thus by Cauchy's formula

$$Q(m,n) = \frac{n^{n+1}}{i2\pi} \oint_{|z|=r} \frac{e^z}{(n+z)^{n-m+1} z^{m+1}}\, dz,$$

for any r in the range $0 < r < n$, so

$$Q(m,n) \le \frac{n^{n+1}}{r^m} \max_{|z|=r} \left| \frac{e^z}{(n+z)^{n-m+1}} \right|.$$

We claim that $|e^z/(n+z)^{n-m+1}|$ has only two local maxima on $|z|=r$: one at $z=+r$ and the other at $z=-r$. To see this note that if $|z|=r$ and $\Re z = t$ then

$$\left|\frac{e^z}{(n+z)^{n-m+1}}\right|^2 = \frac{e^{2t}}{(n^2+r^2+2nt)^{n-m+1}},$$

so

$$\frac{d}{dt}\log\left|\frac{e^z}{(n+z)^{n-m+1}}\right|^2 = 2 - \frac{2n(n-m+1)}{n^2+r^2+2nt}.$$

This function has a unique pole at $t=-(n+r^2/n)/2$, which is to the left of the allowed region $-r \le t \le r$, and it has a unique zero somewhere corresponding to a minimum, not a maximum, so the claim holds. Thus $|e^z/(n+z)^{n-m+1}|$ is bounded on $|z|=r$ by

$$\frac{e^{\pm r}}{(n\pm r)^{n-m+1}},$$

so

$$Q(m,n) \le \frac{n^{n+1}e^{\pm r}}{r^m(n\pm r)^{n-m+1}} = \frac{e^{\pm r}}{(1\pm r/n)^{n-m+1}}\frac{n^m}{r^m}.$$

Here

$$\begin{aligned}\log\left(\frac{e^{\pm r}}{(1\pm r/n)^{n-m+1}}\right) &= \pm r - (n-m+1)\big(\pm r/n - r^2/2n^2 + O(r^3/n^3)\big)\\ &= r^2/2n + O\big(r^3/n^2 + rm/n\big).\end{aligned}$$

Taking $r=(mn)^{1/2}$ we thus have, by Stirling's formula,

$$Q(m,n) \le e^{(1/2+O((m/n)^{1/2}))m}\frac{n^{m/2}}{m^{m/2}} \le O(m^{1/4})e^{O(m^{3/2}/n^{1/2})}\binom{n}{m}^{1/2}. \qquad \square$$

Remark 5.5.2. Although it is straightforward to verify from (5.5.4) that $Q(m,n)$ is the coefficient of X^m in (5.5.5), this proof is unsatisfying in that one needs to know the formula for f in advance. Here is an alternative proof of this fact which does not have this fault.

Let ∂ be the forward difference operator defined on functions $g:\mathbb{N}\to\mathbb{R}$ by

$$\partial g(k) = g(k+1) - g(k),$$

and observe that

$$Q(m,n) = \frac{1}{m!}\partial^m g_{m,n}(0),$$

where $g_{m,n}$ is the polynomial defined by

$$g_{m,n}(k) = n^k(n-k)(n-k-1)\cdots(n-m+1).$$

Intuitively, the sum (5.5.4) exhibits enormous cancellation, precisely because it is an expression for a high derivative of a very "smooth" function $g_{m,n}$. To take advantage of this smoothness we investigate the first derivative $\partial g_{m,n}$. We compute

$$\begin{aligned}\partial g_{m,n}(k) &= n^{k+1}(n-k-1)(n-k-2)\cdots(n-m+1) - n^k(n-k)(n-k-1)\cdots(n-m+1)\\ &= kn^k(n-k-1)(n-k-2)\cdots(n-m+1)\\ &= \frac{k}{n}g_{m,n}(k+1).\end{aligned}$$

From this it is possible to control the higher derivatives: using the Leibniz-type formula

$$\partial(gh)(k) = \partial g(k)h(k+1) + g(k)\partial h(k)$$

we derive the recurrence

$$\begin{aligned}\partial^\ell g_{m,n}(k) &= \partial^{\ell-1}\left(\frac{k}{n}g_{m,n}(k+1)\right)\\ &= \frac{k}{n}\partial^{\ell-1}g_{m,n}(k+1) + \frac{\ell-1}{n}\partial^{\ell-2}g_{m,n}(k+2),\end{aligned}$$

which holds for $\ell \geq 2$. In particular, letting

$$a_\ell = \frac{1}{\ell!}\partial^\ell g_{m,n}(m-\ell),$$

we have the recurrence

$$\ell a_\ell = \frac{m-\ell}{n}a_{\ell-1} + \frac{1}{n}a_{\ell-2} \tag{5.5.6}$$

for $(a_\ell)_{\ell\geq 0}$. Observe that $a_0 = n^m$, $a_1 = (m-1)n^{m-1}$, and that $Q(m,n) = a_m$.

In Section 5.6 we will bound solutions to recurrences like (5.5.6) in a direct fashion. For (5.5.6) itself, it is convenient to employ a generating function technique. For an indeterminate X put

$$f(X) = \sum_{\ell=0}^{\infty} a_\ell X^\ell,$$

and observe by multiplying (5.5.6) by $X^{\ell-1}$ and summing for $\ell \geq 2$ that

$$f'(X) - a_1 = \frac{m-1}{n}(f(X) - a_0) - \frac{1}{n}Xf'(X) + \frac{1}{n}Xf(X)\ .$$

By inputting $a_0 = n^m$ and $a_1 = (m-1)n^{m-1}$ and rearranging we arrive at

$$(n-X)f'(X) = (m-1+X)f(X).$$

The formula (5.5.5) is obtained by solving this differential equation.

5.6 An L^∞ bound for low-entropy minor arcs

The main result of this section is the following proposition.

Proposition 5.6.1. *Let χ be a character with exactly m nonzero coordinates, where $m \leq n/3$. Then*

$$|\widehat{1_S}(\chi)| \leq e^{O\left(m^{3/2}/n^{1/2}+m^{1/2}\right)} \cdot 2^{-m/2} \binom{n}{m}^{-1/2} \cdot \frac{n!}{n^n}.$$

Remark 5.6.2. It is instructive to compare this bound to what one can deduce from Theorem 5.4.1. If χ has exactly m nonzero coordinates and takes precisely k distinct nonzero values, i.e., $\chi = (r_1^{a_1}, \ldots, r_k^{a_k}, 0^{n-m})$ where $\sum a_i = m$, then Theorem 5.4.1 gives us a bound of

$$|\widehat{1_S}(\chi)| \leq \binom{n+k}{k}^{1/2} \binom{n}{a_1, \ldots, a_k, n-m}^{-1/2} \frac{n!}{n^n}$$

and the worst case of this bound over all choices of $k \leq m$ and $a_1, \ldots, a_k$ can be computed approximately as

$$|\widehat{1_S}(\chi)| \leq \binom{n}{m}^{-1/2} e^{o(m)} \frac{n!}{n^n},$$

provided that m is significantly greater than $\sqrt{n}$. In other words a character similar to $\chi = (r^m, 0^{n-m})$ is asymptotically about as bad as the worst case. Moreover if we allow n to be even and set $r = 1/2$ then the bound from Theorem 5.4.1 is essentially sharp, as remarked in that section.

However, in this setting, Proposition 5.6.1 improves on this bound by a significant factor of $2^{-m/2}$. Clearly then such an improvement is only possible under the assumption that n is odd, and the proof of Proposition 5.6.1 must exploit this assumption in a fundamental way. Since we do not know how to adapt the proof of Theorem 5.4.1 to exploit the absence of 2-torsion, we employ a more specialized approach that suffices in the sparse regime.

We observe that this is the unique stage of the proof in which we need the full strength of the hypothesis that G has odd order. Elsewhere we only need to assume that $\sum_{x \in G} x = 0$.

Proof of Proposition 5.6.1. Let $\chi = (r_1, \ldots, r_m, 0^{n-m})$. The key tool in our proof is an exact recursive formula for Fourier coefficients $\widehat{1_S}(\chi)$, in terms of related Fourier

coefficients for which $m' < m$. Specifically, as long as r_m is nonzero we have

$$\begin{aligned}\widehat{1_S}(\chi) &= \frac{(n-m)!}{n^n} \sum_{x_1,\ldots,x_m}{}' e(-r_1x_1 - \cdots - r_mx_m) \\ &= \frac{(n-m)!}{n^n} \sum_{x_1,\ldots,x_{m-1}}{}' \sum_{x_m \neq x_1,\ldots,x_{m-1}} e(-r_1x_1 - \cdots - r_mx_m) \\ &= -\frac{(n-m)!}{n^n} \sum_{x_1,\ldots,x_{m-1}}{}' \sum_{x_m = x_1,\ldots,x_{n-1}} e(-r_1x_1 - \cdots - r_mx_m) \\ &= -\frac{1}{n-m+1} \sum_{i=1}^{m-1} \widehat{1_S}(r_1, \ldots, r_i + r_m, \ldots, r_{m-1}, 0, \ldots, 0) \ . \end{aligned} \tag{5.6.1}$$

As an aside, we remark that this formula can be used for efficient explicit computation of certain kinds of Fourier coefficient.

Let U_m be the maximal value of the $|\widehat{1_S}|$ taken over all characters with exactly m nonzero coordinates. If we apply the triangle inequality to the right hand side of (5.6.1) and bound each $|\widehat{1_S}(\chi')|$ appearing there by the appropriate value $U_{m'}$ (where $m' < m$), we obtain recursive bounds on the U_m.

However, there is a subtlety in this process: the number of nonzero coordinates of

$$(r_1, \ldots, r_i + r_m, \ldots, r_{m-1}, 0, \ldots, 0)$$

might be either $(m-1)$ (if $r_i + r_m \neq 0$) or $(m-2)$ (if $r_i + r_m = 0$), and we do not have much control over which occurs. Since we expect $U_{m-1} < U_{m-2}$, it is in our interests to land in the former case as much as possible. We have the freedom to reorder the r_i before applying (5.6.1), so our goal is to optimize the recursive bound over all choices of r_m.

For each $i \leq m$ let $\mathcal{N}_i = \{j \leq m \colon r_j = -r_i\}$. By reordering, we may assume without loss of generality that $|\mathcal{N}_i|$ is minimized when $i = m$. Suppose $j \in \mathcal{N}_m$. Then $r_j = -r_m$, so since $|G|$ is odd we have $r_j \neq r_m$, so $\mathcal{N}_m \cap \mathcal{N}_j = \emptyset$. Thus by minimality $|\mathcal{N}_m| \leq m/2$.

Hence, applying (5.6.1) we deduce the bound

$$\begin{aligned}|\widehat{1_S}(\chi)| &\leq \frac{|\mathcal{N}_m| U_{m-2} + (m - 1 - |\mathcal{N}_m|) U_{m-1}}{n-m+1} \\ &\leq \frac{1}{n-m+1} \max\left\{(m-1)U_{m-1},\ (m/2)U_{m-2} + (m/2-1)U_{m-1}\right\}\end{aligned}$$

where the former term in the maximum covers the case $U_{m-1} > U_{m-2}$ and the latter the more probable case $U_{m-1} \leq U_{m-2}$. Hence

$$U_m \leq \frac{1}{n-m+1} \max\left\{(m-1)U_{m-1},\ (m/2)U_{m-2} + (m/2-1)U_{m-1}\right\}. \tag{5.6.2}$$

Our task is now to obtain bounds on U_m by solving this recurrence, which, ignoring the maximum, resembles a kind of time-dependent Fibonacci sequence.

In Section 5.5, Remark 5.5.2, we dealt with a similar recurrence using generating function methods. Here, the presence of the maximum of two alternatives, and possibly other reasons, make this approach less attractive. Instead, we will bound (5.6.2) by more hands-on methods.

Specifically, we will phrase (5.6.2) in terms of a product of 2×2 matrices (as one might for the Fibonacci sequence), and control the L^2 norm of the result by bounding the $L^2 \to L^2$ operator norms of each matrix in the product.

Write $\alpha_m = \frac{m}{2(n-m+1)}$, $\beta_m = \frac{m-2}{2(n-m+1)}$, and $\gamma_m = \frac{m-1}{n-m+1}$. Additionally, we work with a rescaled version $V_m = (\alpha_1 \alpha_2 \dots \alpha_m)^{-1/2} U_m$ of U_m, for $m \geq 1$. From (5.6.2) we deduce that

$$V_m \leq \max \left\{ \gamma_m \alpha_m^{-1/2} \, V_{m-1}, \ \alpha_m^{1/2} \alpha_{m-1}^{-1/2} \, V_{m-2} + \beta_m \alpha_m^{-1/2} \, V_{m-1} \right\} \tag{5.6.3}$$

holds for any $m \geq 3$.

If we define matrices

$$M_m = \begin{pmatrix} \gamma_m \alpha_m^{-1/2} & 0 \\ 1 & 0 \end{pmatrix} \quad \text{and} \quad N_m = \begin{pmatrix} \beta_m \alpha_m^{-1/2} & \alpha_m^{1/2} \alpha_{m-1}^{-1/2} \\ 1 & 0 \end{pmatrix},$$

and vectors $v_m = (V_m, V_{m-1})^T$, we can write (5.6.3) equivalently as

$$v_m \leq \max \left\{ M_m v_{m-1}, N_m v_{m-1} \right\} \ .$$

Using the easily obtainable bounds $\alpha_m \alpha_{m-1}^{-1} = 1 + O(1/m)$ and $\beta_m^2 \alpha_m^{-1} \leq m/n$, we get that

$$\operatorname{Tr} \left(N_m^T N_m \right) \leq 2 + m/n + O(1/m)$$

and

$$\det \left(N_m^T N_m \right) = 1 + O(1/m),$$

and so by considering the singular values of N_m, we get the operator norm bound

$$\| N_m \|_{L^2 \to L^2} \leq 1 + O((m/n)^{1/2} + m^{-1/2}) \ .$$

Completely analogously, using the easy bound $\gamma_m^2 \alpha_m^{-1} \leq 4m/n$, we deduce

$$\| M_m \|_{L^2 \to L^2} \leq 1 + O(m/n) \ .$$

Thus we have

$$\begin{aligned} V_m &\le \sqrt{V_m^2 + V_{m-1}^2} \\ &\le \left(\prod_{r=3}^{m} \max\left\{ \|M_r\|_{L^2 \to L^2}, \|N_r\|_{L^2 \to L^2} \right\} \right) \cdot \sqrt{V_2^2 + V_1^2} \\ &\le \left(\prod_{r=3}^{m} \left(1 + O\left((r/n)^{1/2} + r^{-1/2} \right)\right) \right) \cdot \sqrt{V_2^2 + V_1^2} \ . \end{aligned}$$

We already know from Section 5.2 that $U_1 = 0$ and it is evident from (5.6.1) that $U_2 = O(n!/n^{n+1})$. Hence, we can bound the term $\sqrt{V_2^2 + V_1^2}$ in the above inequality by $O(n!/n^n)$, which gives

$$V_m \le e^{O\left(m^{3/2}/n^{1/2} + m^{1/2}\right)} \cdot \frac{n!}{n^n} \ ,$$

and the claim stated in the proposition follows easily from this. □

5.7 End of the argument

To estimate the sum of cubes

$$\sum_{\chi \in \widehat{G}^n} \widehat{1_S}(\chi)^3$$

we divide the set of all $\chi \in \widehat{G}^n$ into three regions depending on the value of $H(\chi)$: a high-entropy range $H \ge 10$, a medium-entropy range $\varepsilon \le H \le 10$, and a low-entropy range $H \le \varepsilon$. Here ε is a small positive constant which will be chosen at the end of the argument.

In the high-entropy range $H \ge 10$ it is enough to use the bound from Corollary 5.2.2,

$$\sum_{\chi:\ H(\chi) \ge 10} |\widehat{1_S}(\chi)|^3 \le e^{-3.5n} n!^3 / n^{3n}.$$

In the medium-entropy range $\varepsilon \le H \le 10$, we first need a bound for the number of characters of a given entropy.

Lemma 5.7.1. *If $H \le 10$ then the number of characters of entropy at most H is bounded by $e^{Hn + o(n)}$.*

Proof. Every character of entropy at most H has an orbit under the permutation action of size at most e^{Hn}, so it suffices to show that there are only $e^{o(n)}$ such orbits of characters of entropy at most 10. Every orbit is uniquely specified by giving the

number a_r of appearances of r for each $r \in \widehat{G}$, so we must show that the number of nonnegative integer vectors $(a_r)_{r\in\widehat{G}}$ with $\sum_r a_r = n$ and satisfying

$$\frac{n!}{\prod_{r\in\widehat{G}} a_r!} \leq e^{10n}$$

is $e^{o(n)}$. Let $\delta > 0$ be a small constant, and let t be the sum of the a_r for which $a_r \leq e^{-11/\delta}n$. Then

$$e^{-11n}n^n \leq e^{-10n}n! \leq \prod_{r\in\widehat{G}} a_r! \leq \prod_{r\in\widehat{G}} a_r^{a_r} \leq (e^{-11/\delta}n)^t n^{n-t} = e^{-11t/\delta}n^n,$$

so $t \leq \delta n$. Since at most $e^{11/\delta}$ of the a_r are bigger than $e^{-11/\delta}n$, the number of $(a_r)_{r\in\widehat{G}}$ is bounded by the number of ways of choosing the set B of at most $e^{11/\delta}$ indices r for which a_r is big, times the number of ways of choosing these a_r, times the number of ways of choosing $(a_r)_{r\notin B}$ so that $\sum_{r\notin B} a_r = n - \sum_{r\in B} a_r \leq \delta n$. This is bounded by

$$n^{e^{11/\delta}} n^{e^{11/\delta}} \binom{n+\delta n-1}{\delta n-1}.$$

The claimed estimate now follows by applying Stirling's formula and choosing δ appropriately. □

We conclude by combining Lemma 5.7.1 with Theorem 5.4.1. Note that if the entropy of χ is $O(1)$ then the number of distinct coordinates of χ is $O(n/\log n)$, so Theorem 5.4.1 implies

$$|\widehat{1_S}(\chi)| \leq e^{-H(\chi)n/2+o(n)}n!/n^n.$$

Thus for any $H \leq 10$ we have

$$\begin{aligned}\sum_{\chi\colon 9H/10<H(\chi)\leq H} |\widehat{1_S}(\chi)|^3 &\leq |\{\chi\colon H(\chi)\leq H\}| \max_{\chi\colon H(\chi)\geq 9H/10} |\widehat{1_S}(\chi)|^3 \\ &\leq e^{-7Hn/20+o(n)}n!^3/n^{3n}.\end{aligned}$$

Decomposing the range $[\varepsilon, 10]$ into ranges of this type and bounding crudely, we obtain

$$\sum_{\chi\colon \varepsilon\leq H(\chi)\leq 10} |\widehat{1_S}(\chi)|^3 \leq O(\log(1/\varepsilon))e^{-7\varepsilon n/20+o(n)}n!^3/n^{3n}.$$

In the low-entropy range $H \leq \varepsilon$, one can easily verify that $\chi = (r_1, \ldots, r_n)$ must repeat some coordinate $r \in \widehat{G}$ at least $(1-\varepsilon)n$ times. Thus by global shift-invariance of $\widehat{1_S}$ we have

$$\sum_{\chi\colon H(\chi)\leq\varepsilon} \widehat{1_S}(\chi)^3 = n\sum_{m=0}^{\varepsilon n} \sum_{\substack{m\text{-sparse } \chi \\ H(\chi)\leq\varepsilon}} \widehat{1_S}(\chi)^3. \tag{5.7.1}$$

By combining Proposition 5.6.1 with Theorem 5.5.1 we have that for any $m \leq \varepsilon n$,

$$\begin{aligned}\sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^3 &\leq \left(\max_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|\right) \sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^2 \\ &\leq e^{O\left(m^{3/2}/n^{1/2}+m^{1/2}\right)} 2^{-m/2} \frac{n!^3}{n^{3n}} \\ &\leq e^{O\left(\varepsilon^{1/2}m+m^{1/2}\right)} 2^{-m/2} \frac{n!^3}{n^{3n}}.\end{aligned}$$

As long as ε is sufficiently small depending on the constant implicit in the $O(\varepsilon^{1/2}m)$ term, this is negligible except when m has size $O(1)$. In this range, we can apply Proposition 5.3.1. We introduce a further parameter M and split the sum (5.7.1) into the ranges $1 \leq m \leq 2M$ and $2M < m \leq \varepsilon n$, to obtain

$$\begin{aligned}\sum_{\chi\colon H(\chi)\leq\varepsilon} \widehat{1_S}(\chi)^3 = n \sum_{m=0}^{M} \frac{(-1)^m}{2^m m!} \frac{n!^3}{n^{3n}} + nO_M\left(\frac{1}{n}\frac{n!^3}{n^{3n}}\right) \\ + O\left(n \sum_{m=2M}^{\varepsilon n} e^{O\left(\varepsilon^{1/2}m+m^{1/2}\right)} 2^{-m/2} \frac{n!^3}{n^{3n}}\right).\end{aligned}$$

Finally, by combining the three bounds we have just proved for each range $H \leq \varepsilon$, $\varepsilon \leq H \leq 10$ and $H \geq 10$, we deduce

$$\begin{aligned}\sum_{\chi\in\widehat{G^n}} \widehat{1_S}(\chi)^3 = n \sum_{m=0}^{M} \frac{(-1)^m}{2^m m!} \frac{n!^3}{n^{3n}} + nO_M\left(\frac{1}{n}\frac{n!^3}{n^{3n}}\right) \\ + O\left(n \sum_{m=2M}^{\varepsilon n} e^{O\left(\varepsilon^{1/2}m+m^{1/2}\right)} 2^{-m/2} \frac{n!^3}{n^{3n}}\right) \\ + O\left(\log(1/\varepsilon) e^{-7\varepsilon n/20+o(n)} \frac{n!^3}{n^{3n}}\right) \\ + O\left(e^{-3.5n} \frac{n!^3}{n^{3n}}\right).\end{aligned}$$

Theorem 5.1.2 now follows by choosing ε and M appropriately.

Chapter 6

A solution to the Pyjama problem

The *pyjama stripe* is the subset of $\mathbb{R}^2$ consisting of a vertical strip of width 2ε around every integer x-coordinate. The *Pyjama problem* asks whether finitely many rotations of the pyjama stripe around the origin can cover the plane.

The purpose of this chapter is to answer this question in the affirmative, for all positive ε. The problem is reduced to a statement closely related to Furstenberg's $\times 2$, $\times 3$ Theorem from topological dynamics, and is proved by analogy with that result.

This chapter is based on the paper [45].

6.1 Introduction

Fix some $\varepsilon > 0$. Let $E = E(\varepsilon)$ denote a *pyjama stripe* viewed as a subset of $\mathbb{C}$:

$$E := \{z \in \mathbb{C} : \Re(z) \,(\mathrm{mod}\, 1) \in (-\varepsilon, \varepsilon)\}$$

i.e. a vertical strip of width 2ε around each integral x-coordinate; see Figure 6.1. (The term "pyjama stripe" is a reference to the resemblance of E to the pattern on a pair of stripy pyjamas.)

For θ a unit complex number, write $E_\theta := \theta^{-1}E$ for the corresponding rotated pyjama stripe, i.e. the set obtained by rotating E by θ^{-1} about the origin.

Problem 6.1.1 (Pyjama problem)**.** *Do finitely many rotations of E cover the plane? I.e., does there exist a finite collection $\{\phi_1, \ldots, \phi_k\}$ of unit complex numbers such that $\bigcup_i E_{\phi_i} = \mathbb{C}$?*

Again see Figure 6.1. The goal of this chapter is to answer this in the affirmative, for all positive ε.

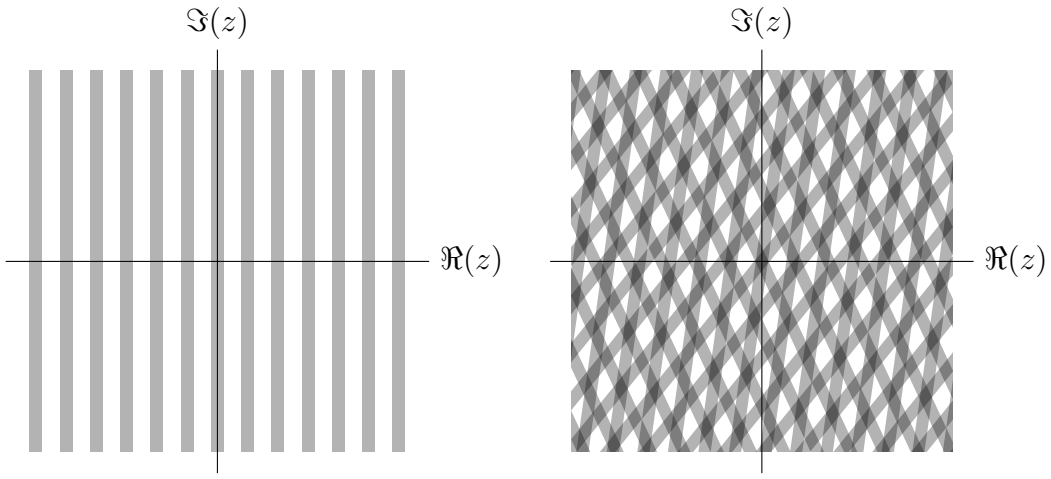

Figure 6.1: A pyjama stripe (left); and a finite set of rotated stripes that fail to cover the plane (right)

The problem was first stated in [37], and was the subject of the more recent note [42], which inspired many of the ideas of this chapter. For a brief history and background of the problem, the reader is referred to [42, Section 1].

Our proof proceeds by reducing the Pyjama problem to a statement (Lemma 6.5.3) that resembles Furstenberg's $\times 2$, $\times 3$ Theorem [28, Part IV], but which – as far as the author is aware – does not appear in the literature. The proof of this result follows that of the $\times 2$, $\times 3$ Theorem very closely, but faces significant technical complications, which bring into play a certain amount of p-adic analysis. Where possible, we have tried to ensure the chapter is nonetheless readable by someone without an extensive background in that field.

We note that this exposition differs slightly from the published version [45]: following a suggestion from Sean Eberhard, we have amended Lemma 6.5.3 to obtain a stronger conclusion and bring it into closer analogy with the $\times 2$, $\times 3$ Theorem. This change prompts other minor ones to other statements.

6.2 Overview of the proof

6.2.1 Statement

For concreteness, we will work from the outset with an explicit finite set of rotations. At this stage, this choice is completely unmotivated; the motivation is spread over the next few sections.

We fix some notation. Take $\mathcal{P}_5 = 1 + 2i$, $\mathcal{P}_{13} = 2 + 3i$; these are primes of $\mathbb{Z}[i]$ dividing 5 and 13 respectively. Write $\theta_5 = \mathcal{P}_5/\overline{\mathcal{P}_5}$ and $\theta_{13} = \mathcal{P}_{13}/\overline{\mathcal{P}_{13}}$, which in

particular are unit complex numbers with rational coefficients (i.e. elements of $\mathbb{Q}(i)$).

We restate the result in this setting.

Theorem 6.2.1. *Let $\varepsilon > 0$ be arbitrary. Define*

$$\Theta_N = \{{\theta_5}^r {\theta_{13}}^s \,:\, r, s \in \mathbb{Z},\, 0 \le r,\, s \le N\}\ .$$

Also, for $n \ge 1$ define $\zeta_1^{(n)}$, $\zeta_2^{(n)}$, $\zeta_3^{(n)}$ to be any triple of unit complex numbers satisfying

$$n\left(\zeta_1^{(n)} + \zeta_2^{(n)}\right) = \zeta_3^{(n)}$$

and let $\Theta' = \Theta'(n, N) = \zeta_1^{(n)}\,\Theta_N \cup \zeta_2^{(n)}\,\Theta_N \cup \zeta_3^{(n)}\,\Theta_N$.

Then there exist n, N (depending on ε) such that $\bigcup_{\theta\in\Theta'} E_\theta = \mathbb{C}$.

Remark 6.2.2.

(i) Finding a triple $\zeta_1^{(n)}, \zeta_2^{(n)}, \zeta_3^{(n)}$ as above is equivalent to finding a triangle (in $\mathbb{C}$) with side lengths $(n, n, 1)$; indeed, $\zeta_i^{(n)}$ correspond to the unit vectors in the direction of the edges of such a triangle. By the converse to the triangle inequality, such a triangle exists, and it is straightforward to find one explicitly.

(ii) Due to the infinitary nature of some of the methods used, the dependence of n and N (and thereby of $|\Theta'| = 3(N+1)^2$) on ε given by our proof is completely ineffective.

(iii) The significance of the numbers 5 and 13 is merely that they are primes of the form $4k + 1$. Any pair of such numbers would work, but again we fix these for concreteness.

6.2.2 Rational obstructions and periodic coverings

It is noted in [42, Section 2] that *periodic coverings* – that is, choices of rotations $\phi_1, \ldots, \phi_k$ such that all the sets E_{ϕ_i} have a common period lattice Λ (of rank 2) – are attractive when approaching the Pyjama problem. Indeed, to check that $\bigcup_i E_{\phi_i} = \mathbb{C}$ for such a covering, it would suffice to check that some fundamental domain of Λ was covered, which is essentially a finite task.

In particular, in the special case that all the rotations ϕ_i have rational coefficients – i.e. are elements of $\mathbb{Q}(i)$ – we can choose some $D \in \mathbb{Z}[i]$ such that $D\phi_i \in \mathbb{Z}[i]$ for all i, and then deduce that $D\,\mathbb{Z}[i]$ is a period lattice. Indeed, for any $z \in E_{\phi_i}$ and $r \in \mathbb{Z}[i]$ we have that $\phi_i\, z \in E$ (by definition), hence $\phi_i(z + D\,r) \in E$ (as E is $\mathbb{Z}[i]$-periodic)

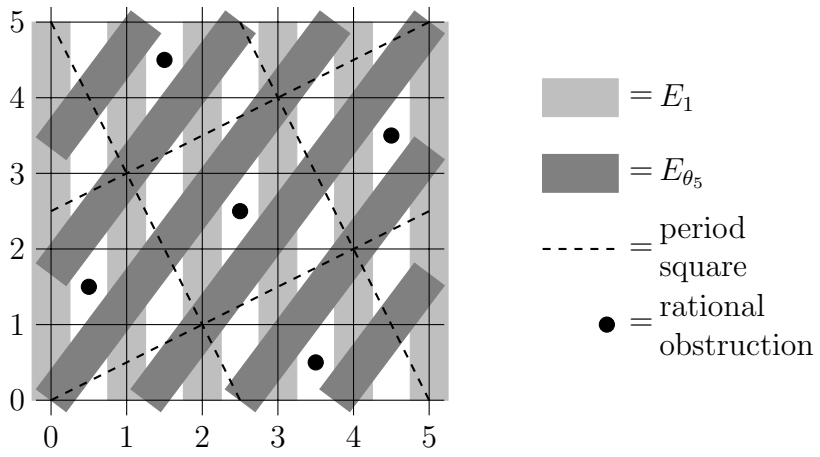

Figure 6.2: A periodic configuration $\{1, \theta_5\}$ showing the rational obstructions $\frac{1+i}{2}D$.

and so we deduce $z + D\,r \in E_{\phi_i}$. Note the sets Θ_N of rotations appearing in the statement of Theorem 6.2.1 fall within this special case.

However, it is also shown in [42, Theorem 2.3] that these periodic coverings cannot possibly cover all of $\mathbb{C}$ when ε is small (specifically, when $\varepsilon < 1/3$). The authors locate certain *rational obstructions*, i.e. points with a specific rational form which are guaranteed not to be covered by $\bigcup_i E_{\phi_i}$. For example, if the ϕ_i are in $\mathbb{Q}(i)$ and D is as above, then the point $\frac{1+i}{2}D$ is not covered (for $\varepsilon < 1/2$); this follows from the fact that if $\frac{r+s\,i}{t}$ is a unit vector (where r, s, t are integers and the ratio is in its lowest terms) then necessarily t is odd.

More generally, for any fixed $\varepsilon > 0$ one can identify a finite set of tuples (a, b, m) of integers such that, for any choice of $\phi_i \in \mathbb{Q}(i)$ and D defined as above, the points $\frac{a+b\,i}{m}D$ are not covered. (Translates of these by $D\,\mathbb{Z}[i]$ are therefore also not covered.) We term all of these *rational obstructions*. Their number increases without bound as $\varepsilon \to 0$.

A simple example of a rational configuration with rotations $\{1, \theta_5\}$ is given in Figure 6.2, with the period square and the rational obstruction corresponding to $\frac{1+i}{2}D$ indicated.

Our strategy to overcome these obstructions is as follows.

- We show that this is in some sense all that can go wrong. That is, for suitably chosen $\phi_i \in \mathbb{Q}(i)$, the *only* points in $\mathbb{C} \setminus \bigcup_i E_{\phi_i}$ are rational obstructions of the above form, or points reasonably close to them. This result is termed the *rationality lemma* and is stated formally in Lemma 6.3.1. Its proof will occupy

the majority of the chapter. The rational rotations used correspond to the Θ_N in the statement of Theorem 6.2.1.

- We then use a trick to deal with the missing points. As this trick necessarily involves adding some irrational rotations to the set, we term it the *irrational trick*; it corresponds to the $\zeta_i^{(n)}$ in the statement of Theorem 6.2.1. We note that this trick is essentially a variation of a technique used in [42, Theorem 3.1] to construct a covering with $\varepsilon = 1/3 - 1/48$. Section 6.3 describes this part of the argument.

The proof of the rationality lemma is by far the lengthier and harder part of the argument; we now turn to the strategy for proving that.

6.2.3 The rationality lemma and topological dynamics

The first stage in the proof of the rationality lemma is to deduce it from an analogous infinitary statement. This process is very similar to (and inspired by) the infinitary reformulation of the Pyjama problem itself in [42, Theorem 4.1], with a few notable differences: most significantly, that we are reformulating the rationality lemma rather than the Pyjama problem. The new version is Lemma 6.5.3, the *infinitary irrationality lemma*.

Our main reason for making this reformulation is that the new version bears a striking resemblance to Furstenberg's $\times 2$, $\times 3$ Theorem [28, Section IV]. Various extensions of Furstenberg's result are even closer, both formally and in spirit, to what we need: e.g. work of Berend [2], generalizing Furstenberg's result to arbitrary finite-dimensional connected compact abelian groups; and generalizations that allow an isometric direction pursued by Muchnik [47] and more recently by Wang [63]. However, Lemma 6.5.3 does not follow from any of these results, or (as far as the author is aware) from any others in the published literature, although it could be seen as a case of a common generalization of all of them. Here, we give a self-contained proof, by adapting Furstenberg's original argument – or more accurately, a variant of it due to Boshernitzan [8] – to our context.

6.2.4 Layout of the chapter

Section 6.3 states the rationality lemma and deduces Theorem 6.2.1 from it, using the irrational trick.

Section 6.4 reproduces (for reference) a sketch proof of Furstenberg's $\times 2$, $\times 3$ Theorem, that will be used as a model for the proof of Lemma 6.5.3.

Section 6.5 deals with the jump to the infinitary setting, and proves the relevant reductions.

Section 6.6 provides some background on the limit object $\widehat{A}$ appearing in the infinitary rationality lemma, and states some standard properties of it. We then state a number of less standard but essentially straightforward auxiliary results that we will need to run the proof of the $\times 2$, $\times 3$ Theorem; the lengthier and more technical proofs are consigned to Section 6.8.

Finally, Section 6.7 reruns the proof from Section 6.4 in the new context, thereby proving the infinitary rationality lemma.

6.2.5 Notation

We use $\Re(z)$ to denote the real part of a complex number. In a metric space, $B_R(x)$ will denote the open ball of radius R about x, and $\overline{B}_R(x)$ will similarly denote the closed ball.

6.3 The irrational trick

We state the rationality lemma formally.

Lemma 6.3.1 (Rationality lemma)**.** *Let $\Theta_N = \{{\theta_5}^r {\theta_{13}}^s : 0 \leq r, s \leq N\}$ be as above, and let $\varepsilon > 0$ be fixed. There is an absolute constant $T \ll 1$ and a parameter $n \ll_\varepsilon 1$, such that for every $R > 0$ there exists $D \in \mathbb{Z}[i]$ with $|D| \geq R$, and a positive integer $N \ll_{\varepsilon,R} 1$, so that the following holds: any point of $\mathbb{C}$ not contained in $\bigcup_{\theta\in\Theta_N} E_\theta$ is at distance at most T from a very rational point $r \in \mathbb{C}$, in the sense that $r \in \frac{1}{n} D\,\mathbb{Z}[i] \setminus D\mathbb{Z}[i]$.*

Phrased another way, let $W = \mathbb{C} \setminus \bigcup_{\theta\in\Theta_N} E_\theta$ be the set of missing points; then W is contained in the set

$$W' = \bigcup \left\{ \overline{B}_T(r) \,:\, r \in \frac{1}{n} D\,\mathbb{Z}[i],\, r \notin D\,\mathbb{Z}[i] \right\} \subseteq \mathbb{C}\ .$$

Remark 6.3.2.

(i) We see that W' is periodic with period $D\,\mathbb{Z}[i]$, as expected.

(ii) It will be important that $n = n(\varepsilon)$ and T are bounded independently of R. This implies that the set W' becomes very sparse as D becomes large. (In reality, T is always bounded by 0.51 or so, but we do not need to know this.)

(iii) In fact, D will be some large power of $\overline{\mathcal{P}_5\mathcal{P}_{13}}$.

Once we know that the set of missing points is both very sparse and very structured, it is possible to combine several copies of Θ_N in such a way as to cover the plane. There are likely to be many ways to do this; we give one below.

Proof of Theorem 6.2.1 given Lemma 6.3.1. Take $n \ll_\varepsilon 1$ as in Lemma 6.3.1. We choose $|D|$ in that lemma to be sufficiently large that the ball $\overline{B}_{2nT}(0)$, and its translates by $D\,\mathbb{Z}[i]$, are contained in $\bigcup_{\theta\in\Theta_N} E_\theta$. Explicitly we can take any $|D| > 2n^2T + nT$, as the nearest point of W' to 0 is at least a distance $|D|/n - T$ away.

Write $\Theta' = \Theta'(n, N)$ as in the statement, and abbreviate $\zeta_i^{(n)}$ to ζ_i. Suppose $z \in \mathbb{C}$ does not lie in $\bigcup_{\theta\in\Theta'} E_\theta$. Equivalently, none of $\zeta_1 z$, $\zeta_2 z$, $\zeta_3 z$ lie in $\bigcup_{\theta\in\Theta_N} E_\theta$. By Lemma 6.3.1, we have:

$$\zeta_i z \in \overline{B}_T(r_i)$$

for $i \in \{1, 2, 3\}$, for some $r_i \in \frac{1}{n} D\mathbb{Z}[i]$ (but $r_i \notin D\mathbb{Z}[i]$).

In particular, $n\,(r_1 + r_2)$ lies in $D\,\mathbb{Z}[i]$, so $n(\zeta_1 z + \zeta_2 z)$ lies in some ball of radius $2nT$ around a point of $D\,\mathbb{Z}[i]$. By the choice of D, this point lies in $\cup_{\theta\in\Theta_N} E_\theta$. But this point is also precisely $\zeta_3 z$, a contradiction. □

6.4 Furstenberg's $\times 2$, $\times 3$ Theorem

Since it will be so central to our proof of the rationality lemma, we briefly review Furstenberg's $\times 2$, $\times 3$ Theorem and its proof here.

Recall that, for a topological dynamical system (Φ, X) (i.e. Φ is a semigroup acting continuously on a compact metric space X) we say $Y \subseteq X$ is invariant if $\Phi(Y) \subseteq Y$.

So, let $\Phi = \{2^r 3^s : r, s \in \mathbb{Z}_{\geq 0}\} \subseteq \mathbb{N}$ be a semigroup under multiplication, acting on $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ by multiplication. The theorem states:

Theorem 6.4.1 ([28, Section IV])**.** *The only closed, invariant subsets $Y \subseteq \mathbb{T}$ with respect to this action are $\mathbb{T}$ itself, and certain finite sets of rationals in $\mathbb{T}$.*

This in fact holds whenever Φ is a multiplicative subsemigroup of $\mathbb{N}$ that is *non-lacunary*, in the sense that Φ is not contained in any set of the form $\{a^r : r \in \mathbb{Z}_{\geq 0}\}$ for some $a \in \mathbb{N}$. However, the case when Φ is the semigroup generated by 2 and 3 is essentially as hard as the general case.

Both Furstenberg's and Boshernitzan's proofs make use of the following lemma and its corollaries. This is the point in the argument where the non-lacunary nature of Φ is used.

Lemma 6.4.2 (Related to [28, Lemma IV.1]). *Fix $\delta > 0$. For sufficiently small $\eta > 0$ we have that $\eta\Phi \cap [0,1)$ is δ-dense in $[0,1)$.*

Proof. It suffices to show that $\eta\Phi$ is $\delta/2$-dense on $[\delta/2, 1)$. Taking logarithms, it is then also sufficient to show that $\log(\eta\Phi) \cap [\log(\delta/2), 0]$ is δ'-dense on $[\log(\delta/2), 0]$, for some δ' sufficiently small in terms of δ (in fact δ' can be taken to be $\delta/2$).

Equivalently, this states that $\log\Phi = \{r\log 2 + s\log 3 : r, s \in \mathbb{Z}_{\geq 0}\}$ is δ'-dense on $[\log(\delta/2) + C, C]$ for $C = -\log\eta$ large enough. Crucially, $\log 3/\log 2$ is irrational, so we can choose N such that $\{s\log 3/\log 2 \pmod 1 : s \in \mathbb{Z}, 0 \leq s \leq N\}$ is $(\delta'/\log 2)$-dense on $\mathbb{R}/\mathbb{Z}$. Then $\{r + s\log 3/\log 2 : r, s \in \mathbb{Z}_{\geq 0}\}$ is $(\delta'/\log 2)$-dense on $[N\log 3/\log 2, \infty)$ and so $\log\Phi$ is δ'-dense on $[N\log 3, \infty)$, which gives the result. □

Corollary 6.4.3 ([28, Lemma IV.2]). *Suppose $Y \subseteq \mathbb{T}$ is a closed invariant subset which has 0 as a limit point. Then $Y = \mathbb{T}$. The same conclusion holds if Y has just a rational limit point.*

Proof sketch. For any $\delta > 0$, we can choose a non-zero $t \in Y$ small enough that $|t|\Phi \cap [0,1)$ is δ-dense in $[0,1)$, so *a fortiori* the orbit $\Phi(t) \subseteq Y$ is δ-dense in $\mathbb{T}$. But δ was arbitrary and Y was closed, so $Y = \mathbb{T}$. The case of a rational limit point is similar, or can be deduced from this. □

The arguments in Furstenberg [28] and Boshernitzan [8] now diverge; we broadly follow Boshernitzan.

Sketch proof of Theorem 6.4.1. Define $\Phi^{(m)} = \{2^{mr}3^{ms} : r, s \in \mathbb{Z}_{\geq 0}\} \subseteq \Phi$, i.e. the subsemigroup generated by 2^m, 3^m. The above results hold also for $\Phi^{(m)}$ for any positive integer m (as it is still non-lacunary).

If Y is finite it is straightforward to see it consists of rationals. Suppose Y is infinite. Let Y' be the set of limit points of Y; it is another closed invariant subset, and non-empty as Y is infinite. If Y' contains any rationals then $Y = \mathbb{T}$ by Corollary 6.4.3, so assume it doesn't.

Fix $\delta > 0$. Choose a finite, δ-dense set of rationals in $\mathbb{T}$ whose denominators are coprime to 2 and 3; call them $\{a_0, \ldots, a_{\ell-1}\}$. For ease of notation assume $a_0 = 0$. Then for some m, $\Phi^{(m)}$ fixes all the a_i.

For $0 \leq k < \ell$, let $X_k = (Y' - a_0) \cap \cdots \cap (Y' - a_k)$. We claim each such X_k is closed, $\Phi^{(m)}$-invariant and non-empty. This will complete the proof: if we know that $X_{\ell-1}$ is non-empty, then there exists some $x \in \mathbb{T}$ such that $\{x + a_0, x + a_1, \ldots, x + a_{\ell-1}\} \subseteq Y'$. In particular, Y' is δ-dense in $\mathbb{T}$, and hence so is Y; but $\delta > 0$ was arbitrary so $Y = \mathbb{T}$ (as it is closed).

The proof of this claim is by induction on k; since $X_0 = Y'$ the case $k = 0$ is clear. Suppose X_k is closed, invariant and non-empty. We have $X_{k+1} = X_k \cap (Y' - a_{k+1})$ so this is certainly closed; and since a_{k+1} fixed by $\Phi^{(m)}$, it is also $\Phi^{(m)}$-invariant.

Since $X_k \subseteq Y'$ it contains only irrational points; so for any $x \in X_k$, the orbit $\Phi^{(m)}(x)$ is infinite and contained in X_k. So X_k is infinite. Hence $(X_k - X_k)$ has 0 as a limit point, and is closed and $\Phi^{(m)}$-invariant. So $X_k - X_k = \mathbb{T}$ (by the corollary), and in particular there exist $x, y \in X_k$ such that $y - x = a_{k+1}$. So $x = y - a_{k+1} \in Y' - a_{k+1}$ and $x \in X_k$, so $x \in X_{k+1}$ as required. □

6.5 An infinitary reformulation of the rationality lemma

6.5.1 Introductory remarks

We have two main motivations for considering infinitary versions of the rationality lemma.

- The statement of Lemma 6.3.1 involves four interdependent parameters (ε, D, n, N), all of which can be successfully eliminated by passing to a suitable limit structure. As is typical, this makes the statements and proofs significantly cleaner at the expense of sacrificing quantitative control.
- More importantly, passing to the limit exposes the connection with topological dynamics and Furstenberg's proof of the $\times 2$, $\times 3$ Theorem. This link is much more obscure in the finitary setting.

We are not aware of any reason in principle preventing the whole argument from being finitized to obtain effective bounds on N and n in terms of ε. However, the bounds obtained in this way would likely be of very poor quality (i.e. at least of tower-type).

6.5.2 Dual groups and a limit space

In [42], the authors observe that the sets E_θ have a natural interpretation in terms of characters on $\mathbb{C}$.

The continuous characters on $\mathbb{C}$ all have the form:

$$\begin{aligned} \chi_z \,:\, \mathbb{C} &\to \mathbb{T} \\ w &\mapsto \Re\,(z\,w) \pmod 1 \end{aligned}$$

where $z \in \mathbb{C}$ is some complex number. So, simply by unravelling the definitions we can equivalently characterize the pyjama stripe E_θ by the formula:

$$E_\theta = \{z \in \mathbb{C} \,:\, \chi_z(\theta) \in (-\varepsilon, \varepsilon)\}\ .$$

The interesting feature of this definition is that it makes sense for *any* character on $\mathbb{C}$, continuous or not. Writing $\widehat{\mathbb{C}}$ for the group of all characters (not necessarily continuous) on $\mathbb{C}$, i.e. the Bohr compactification, we can similarly define:

$$E_\theta^* = \{\chi \in \widehat{\mathbb{C}} \,:\, \chi(\theta) \in (-\varepsilon, \varepsilon)\} \subseteq \widehat{\mathbb{C}} \,.$$

The notable properties of this situation are that $\widehat{\mathbb{C}}$, being the Pontryagin dual of $\mathbb{C}$ endowed with the discrete topology, is compact with respect to its natural topology (the topology of pointwise evaluation); and moreover E_θ^* is an open set in this topology. So, in [42] the authors conclude that the Pyjama problem is equivalent to the following infinitary question:

Problem 6.5.1 (Infinitary Pyjama problem; [42, Lemma 4.1])**.** *Let $\varepsilon > 0$. Is there a set Φ of unit complex numbers, which need not be finite and without loss of generality may be taken to be the whole unit circle, such that $\bigcup_{\phi\in\Phi} E_\phi^*$ is all of $\widehat{\mathbb{C}}$?*

Deducing the Pyjama problem from this is straightforward. If such a set Φ exists, then by compactness there is actually a finite set $\Phi = \{\phi_1, \ldots, \phi_k\}$ with the same property, i.e.:

$$\bigcup_i E_{\phi_i}^* = \widehat{\mathbb{C}} \,.$$

Now, intersecting both sides of this equation with the set $\{\chi_z \,:\, z \in \mathbb{C}\}$ of continuous characters gives

$$\bigcup_i E_{\phi_i} = \mathbb{C}$$

as required. The converse direction is less straightforward, and essentially follows from the fact that the continuous characters are dense in $\widehat{\mathbb{C}}$.

We refer to the space $\widehat{\mathbb{C}}$ as the *limit space* being used in the arguments of [42].

6.5.3 The smaller limit space $\widehat{A}$

One downside of this approach, in our opinion, is that this limit space $\widehat{\mathbb{C}}$ is very large (for instance, non-metrizable) and generally difficult to work with.

However, we observe that we only ever need to know the value of a given $\chi \in \widehat{\mathbb{C}}$ *at points of* Φ. So, the infinitary Pyjama problem is well-defined on the quotient of $\widehat{\mathbb{C}}$ obtained by restricting to Φ. If Φ is very large, this doesn't help us very much; but if Φ is chosen very conservatively, this quotient might be comparatively manageable.

Specifically, suppose we work with the set of rotations

$$\Theta = \Theta_\infty = \{{\theta_5}^r {\theta_{13}}^s \,:\, r,\, s \in \mathbb{Z}_{\geq 0}\} = \bigcup_{N\geq 1} \Theta_N \,.$$

We write A for the subring of $\mathbb{Q}(i)$ generated by $\mathbb{Z}[i]$, $1/\overline{\mathcal{P}_5}$ and $1/\overline{\mathcal{P}_{13}}$, denoted $A = \mathbb{Z}[i][1/\overline{\mathcal{P}_5}, 1/\overline{\mathcal{P}_{13}}]$. To put it another way, A consists of those elements of $\mathbb{Q}(i)$ whose denominators are of the form $\overline{\mathcal{P}_5}^r\overline{\mathcal{P}_{13}}^s$.

In particular $\Theta \subseteq A$ and so it suffices to consider the quotient of $\widehat{\mathbb{C}}$ obtained by restricting the characters to A. Equivalently, that space is precisely the Pontryagin dual $\widehat{A}$, where A is given the discrete topology.

By contrast to $\widehat{\mathbb{C}}$, the space $\widehat{A}$ (again with the topology of pointwise evaluation) is reasonably tame.[1] In general terms, it is metrizable and second-countable; but more importantly its structure can be understood entirely concretely as a certain quotient of a finite product of $\mathbb{C}$ and some p-adic fields. We will return to this viewpoint in Section 6.6.

Of course, since we know that rotations in $\mathbb{Q}(i)$ cannot cover $\mathbb{C}$, the analogue of Problem 6.5.1 for $\widehat{A}$ is false. Our aim is instead to transfer the rationality lemma (Lemma 6.3.1) to a statement on $\widehat{A}$. A first attempt is given below.

Note that continuous characters on $\mathbb{C}$ are certainly characters of A; that is, there is a map

$$\begin{aligned} \jmath_{\mathbb{C}} \,:\, \mathbb{C} &\to \widehat{A} \\ z &\mapsto \chi_z|_A \,. \end{aligned}$$

We write $\overline{B}_R^{\mathbb{C}}(0) \subseteq \widehat{A}$ for the image $\jmath_{\mathbb{C}}(\overline{B}_R(0))$ of the closed ball of radius R about 0 in $\mathbb{C}$, under $\jmath_{\mathbb{C}}$. For $x \in \widehat{A}$, similarly define $\overline{B}_R^{\mathbb{C}}(x) = x + \overline{B}_R^{\mathbb{C}}(0)$.

Lemma 6.5.2 (Infinitary rationality lemma, first formulation)**.** *Pick $\varepsilon > 0$, and let Θ, A be defined as above. Treat*

$$E_\theta^* = \{x \in \widehat{A} \,:\, x(\theta) \in (-\varepsilon, \varepsilon)\}$$

now as a subset of $\widehat{A}$. Define $W^* = \widehat{A} \setminus \bigcup_{\theta\in\Theta} E_\theta^*$ (the missing points). Then there exists $n \ll_\varepsilon 1$ and $R \ll 1$ such that W^* is contained in the set:

$$W'^* = \left\{ w \in \widehat{A} \,:\, n\,w \in \overline{B}_{nR}^{\mathbb{C}}(0) \right\}$$

In other words, every point w of W^* is close to an n-torsion point of $\widehat{A}$ in the sense that they differ by the image under $\jmath_{\mathbb{C}}$ of a small complex number.

[1] It is perhaps worth noting that the heuristic statements "we work with rational rotations so that the state space (i.e. fundamental domain in $\mathbb{C}$) is finite" and "we work with rotations in A so that the state space (i.e. $\widehat{A}$) is reasonably tame", are very closely related.

Note that, in this first pass reformulation, the parameters N and D have been eliminated but ε, n and T (analogous to R here) have not.

The deduction of Lemma 6.3.1 from this is somewhat less straightforward than in the case of Problem 6.5.1, but is nonetheless essentially formal.

Deduction of Lemma 6.3.1 from Lemma 6.5.2. Let $m \in \mathbb{N}$ and $\delta > 0$ be parameters to be specified later. Then the set

$$U_0 = \left\{ x \in \widehat{A} \,:\, x(i^\ell {\theta_5}^r {\theta_{13}}^s) \in (-\delta, \delta) \;\forall\; 0 \leq r, s \leq m,\; \ell \in \{0,1\} \right\}$$

is a standard open neighbourhood of 0 in $\widehat{A}$, as the latter has the topology of pointwise evaluation.

Define

$$U = \bigcup_{w \in W'^*} (w + U_0) \subseteq \widehat{A}$$

which is therefore an open neighbourhood of W'^*. Assuming Lemma 6.5.2, we have that

$$U \cup \bigcup_{\theta \in \Theta} E_\theta^* = \widehat{A}\,.$$

By compactness, and since all these sets are open, there exists N such that

$$U \cup \bigcup_{\theta \in \Theta_N} E_\theta^* = \widehat{A}\,.$$

Taking the preimage of both sides under $\jmath_{\mathbb{C}}$, we obtain

$$\jmath_{\mathbb{C}}^{-1}(U) \cup \bigcup_{\theta \in \Theta_N} E_\theta = \mathbb{C}$$

and rearranging this we get

$$\mathbb{C} \setminus \bigcup_{\theta \in \Theta_N} E_\theta \subseteq \jmath_{\mathbb{C}}^{-1}(U)\,.$$

Hence it suffices to verify that this mysterious set $\jmath_{\mathbb{C}}^{-1}(U)$ is contained in the set W' from the statement of Lemma 6.3.1, for appropriate values of the parameters. This is essentially just a case of unravelling the definitions.

Suppose $z \in \jmath_{\mathbb{C}}^{-1}(U)$. That means there exists $w \in W'^*$ and $u \in U_0$ such that $\chi_z = w + u$. So, $n\,\chi_z = nw + nu$; but we know that $nw = \chi_{nt}$ for some $t \in \mathbb{C}$ with $|t| \leq T$. So, $nu = \chi_{nz-nt}$.

By the definition of U_0, we have $u(i^\ell {\theta_5}^r {\theta_{13}}^s) \in (-\delta, \delta)$ for all $0 \leq r, s \leq m$ and $\ell \in \{0, 1\}$. So, $nu(i^\ell {\theta_5}^r {\theta_{13}}^s) \in (-n\delta, n\delta)$ for all such r, s, ℓ and hence

$$\Re\left(n(z-t)\, i^\ell {\theta_5}^r {\theta_{13}}^s\right) \pmod 1 \in (-n\delta, n\delta)$$

for all such r, s, ℓ, which implies that $n(z-t)\, {\theta_5}^r {\theta_{13}}^s$ is within distance $2n\delta$ of a point of $\mathbb{Z}[i]$ for all $0 \leq r, s \leq m$.

We isolate the following elementary claim.

Claim. *Write $D = \overline{\mathcal{P}_5}^a \overline{\mathcal{P}_{13}}^b$ for some $a, b \geq 0$. Fix some $x \in \mathbb{C}$, $\eta \leq 1/100$ and suppose for all integers r, s with $0 \leq r \leq a$, $0 \leq s \leq b$ we have that ${\theta_5}^r {\theta_{13}}^s x$ is within distance η of a point of $\mathbb{Z}[i]$. Then x is within distance η of a point of $D\mathbb{Z}[i]$.*

Proof of claim. We proceed by induction on a and b. The base case $a = 0$, $b = 0$ is by assumption. If $a > 0$, we deduce by inductive hypothesis that x and $\theta_5 x$ both lie within η of points of $\overline{\mathcal{P}_5}^{a-1} \overline{\mathcal{P}_{13}}^b \mathbb{Z}[i]$; call these points y_1, y_2 respectively. Now, $|\theta_5 y_1 - y_2| \leq 1/50$ and both $\theta_5 y_1$ and y_2 lie in $\overline{\mathcal{P}_5}^{a-2} \overline{\mathcal{P}_{13}}^b \mathbb{Z}[i]$, so $y_2 = \theta_5 y_1$ (as distinct points of $\overline{\mathcal{P}_5}^{-1} \mathbb{Z}[i]$ are a distance at least $1/\sqrt{5}$ apart). Hence $y_1 = {\theta_5}^{-1} y_2 \in \overline{\mathcal{P}_5}^a \overline{\mathcal{P}_{13}}^b \mathbb{Z}[i]$ as required.

The induction step on b is analogous. □

Applying the claim to $n(z-t)$ and setting $\delta = 1/(200n)$, we deduce that $n(z-t)$ is within $1/100$ of a point of $D\mathbb{Z}[i]$ where $D = \overline{\mathcal{P}_5}^m \overline{\mathcal{P}_{13}}^m$; so z is certainly within $(R + 1/100)$ of a point $k \in \frac{1}{n} D\mathbb{Z}[i]$ (recalling $|t| \leq R$). Setting $T = R + 1/100$, we deduce that z lies in W' – *unless* the point k lies in $D\mathbb{Z}[i]$ itself, which was specifically prohibited in the definition of W'.

We check directly that this last case doesn't happen. Indeed, suppose $z = k + t'$ for $k \in D\,\mathbb{Z}[i]$ and $|t'| \leq T$. Since ${\theta_5}^r \neq 1$ for all $r \geq 0$, the map $\tau \mapsto \theta_5 \tau$ is an aperiodic rotation on the unit circle, and so by standard arguments we have that $\{{\theta_5}^r\, t' : 0 \leq r \leq M\}$ is $\varepsilon/2$-dense on the circle $\{z \in \mathbb{C} : |z| = |t'|\}$ for some $M \ll_\varepsilon 1$. In particular there is some $0 \leq r \leq M$ such that ${\theta_5}^r\, t' \in E$. Assuming $m \geq M$ (which we can), we get ${\theta_5}^r\, k \in \mathbb{Z}[i]$ and hence ${\theta_5}^r\, z \in E$.

Noting that m was allowed to be arbitrarily large, we get $|D|$ arbitrarily large (independently of n, R and ε), thereby completing the proof. □

6.5.4 A dynamical reformulation

Although Lemma 6.5.2 puts the rationality lemma in a compact setting, the link with topological dynamics and the $\times 2$, $\times 3$ Theorem is not yet clear. Our second (and final!) reformulation requires a simple but important shift in perspective.

Rather than rotating the pyjama stripe E_1^* to E_θ^* and asking whether the rotated copies cover most of $\widehat{A}$, we keep the stripe fixed and rotate the *points of* $\widehat{A}$ *itself*. We then ask which points have rotations landing in the standard stripe E_1^*.

Formally, we define an action of Θ on $\widehat{A}$ by

$$\begin{aligned} \rho \,:\, \Theta \times \widehat{A} &\to \widehat{A} \\ (\theta, \chi) &\mapsto (a \mapsto \chi(\theta a)) \ . \end{aligned}$$

(We normally omit the ρ and write $\theta(\chi)$ to refer to this action.)

This is a continuous, multiplicative semigroup action; i.e. $(\Theta, \widehat{A})$ can be viewed as a topological dynamical system. We can therefore use the terminology of invariant subsets etc. in this context.

Note that

$$\begin{aligned} E_\theta^* &= \left\{ \chi \in \widehat{A} \,:\, \theta(\chi) \in E_1^* \right\} \\ &= \rho(\theta)^{-1}(E_1^*) \end{aligned}$$

where this once we use $\rho(\theta)^{-1}$ explicitly to denote the preimage, thereby avoiding confusion of θ^{-1} and $1/\theta$. (Since $1/\theta \notin A$, its action on $\widehat{A}$ is not well-defined.)

We are now in a position to state the full infinitary rationality lemma.

Lemma 6.5.3 (Infinitary rationality lemma)**.** *As above, treat* $(\Theta, \widehat{A})$ *as a topological dynamical system. If* $Y \subseteq \widehat{A}$ *is a closed,* Θ*-invariant set, then either* $Y = \widehat{A}$ *or* Y *is contained in a set of the form*

$$U = \bigcup_{i=1}^{k} \overline{B}_R^{\mathbb{C}}(u_i)$$

where $u_1, \ldots, u_k$ *is a finite collection of torsion points of* $\widehat{A}$*, and* $R > 0$ *is some real number.*

Remark 6.5.4.

(i) The approximate dictionary with the $\times 2$, $\times 3$ Theorem is then:

$$\begin{gathered}\widehat{A} \longleftrightarrow \mathbb{T} \\ \{\theta_5{}^r\theta_{13}{}^s \,:\, r,s \in \mathbb{Z}_{\geq 0}\} = \Theta \longleftrightarrow \Phi = \{2^r 3^s \,:\, r,s \in \mathbb{Z}_{\geq 0}\} \\ Y = \widehat{A} \longleftrightarrow Y = \mathbb{T} \\ Y \text{ is contained in some } U \longleftrightarrow Y \text{ is a finite set of rationals}\end{gathered}$$

We note that the last of these is weaker than the obvious analogue, namely "Y is a finite set of torsion points". However, such a statement would be false, as can be seen by taking, say, $Y = \overline{B}^{\mathbb{C}}_{1/10}(0) \subseteq \widehat{A}$.

(ii) For all this to make sense, it was crucial that Θ was a multiplicative semigroup. This forms part of the motivation for the choice of Θ and by extension of the Θ_N.

(iii) Notice that all the parameters (i.e. n, ε etc.) appearing in Lemma 6.3.1 have now been eliminated, with the exception of the constant R.

Having seen how Lemma 6.5.3 relates to the $\times 2$, $\times 3$ Theorem, we now see how it relates to Lemma 6.5.2 by deducing the latter from it. This is very straightforward.

Deduction of Lemma 6.5.2 from Lemma 6.5.3. Consider the closed, invariant subset of $\widehat{A}$ given by

$$\begin{aligned}W^* &= \bigcap_{\theta\in\Theta} \rho(\theta)^{-1}\left(\widehat{A} \setminus E_1^*\right) \\ &= \widehat{A} \setminus \bigcup_{\theta\in\Theta} \rho(\theta)^{-1} E_1^* \\ &= \widehat{A} \setminus \bigcup_{\theta\in\Theta} E_\theta^*\end{aligned}$$

(closed because E_1^* is open). Clearly $W \neq \widehat{A}$, as each E_θ^* is non-empty, and hence $W^* \subseteq U$ for some U of the specified form. Taking n such that $n\,u_i = 0\ \forall i$, this implies the conclusion of Lemma 6.5.2. □

To recap: in order to prove Theorem 6.2.1 (the Pyjama problem) it now suffices to prove Lemma 6.5.3 using the analogy with the proof of Theorem 6.4.1.

Remark 6.5.5. Dynamical systems of the form $(\Theta, \widehat{A})$ are by no means new. Such objects are referred to as "S-integer dynamical systems" in [16], and similar spaces are studied in [2], as mentioned in the introduction.

6.5.5 Remarks on the choice of Θ

The motivation for the choice of the rotations Θ, Θ_N and ultimately $\Theta'(n, N)$ (see Theorem 6.2.1) should now be complete, and we offer a few remarks.

We noted above that Θ needs to be a semigroup. All we really needed about the sets Θ_N was that their union was Θ; but taking Θ_N to be a long multidimensional progression on the generators is the most natural choice. Once we have chosen Θ_N, the choice of $\Theta'(n, N)$ is fixed by the use of the irrational trick (see Section 6.3).

We could in principle have used a larger semigroup in place of Θ, such as the set of all unit norm elements of $\mathbb{Q}(i)$, which clearly contains Θ. Although this cannot hurt in some sense – adding more rotations doesn't make the problem harder – we would be forced to pass to a larger limit structure, e.g. $\widehat{\mathbb{Q}(i)}$. This is problematic for two reasons.

(i) Working with this larger, more complicated space introduces yet more technical hurdles than already exist in Section 6.6.

(ii) Trying to get any benefit out of the extra rotations brings into play non-trivial questions about the distribution of rational points on the unit circle, or about the distribution of primes.

In the other direction, we could not have used a much smaller semigroup. Indeed, the appropriate results are false for the smaller set $\{{\theta_5}^r : r \geq 0\}$; by analogy with Theorem 6.4.1, we might say this set is lacunary in some sense.

In other words, Θ is the simplest example that works, in much the same way that $\{2^r 3^s : r, s \in \mathbb{Z}_{\geq 0}\}$ is the simplest case that works in the $\times 2$, $\times 3$ Theorem.

6.6 Technical results about $\widehat{A}$

6.6.1 Introductory remarks

The remainder of the chapter is dedicated to the proof of Lemma 6.5.3 by analogy with the $\times 2$, $\times 3$ Theorem. However, there are two sources of technical complication.

- The proof of the $\times 2$, $\times 3$ Theorem required a certain amount of detailed knowledge of the compact group $\mathbb{T}$: e.g. what its torsion points are; what its finite orbits under Φ are, and so forth. Transferring the proof to $\widehat{A}$ requires a similar degree of knowledge of $\widehat{A}$. Mostly these facts are not hard, but require slightly more justification in this less familiar setting.

This is most pronounced when it comes to proving an analogue of the density estimate of Lemma 6.4.2. Recall this was fairly straightforward after taking logarithms. Essentially the same proof works; but the act of "taking logarithms" on $\widehat{A}$ requires a very detailed knowledge of the structure of that space in terms of p-adic fields.

Similar difficulties are overcome in [2], although we have not followed that author's approach very closely.

- The statement of Lemma 6.5.3 necessarily involves the small complex balls $\overline{B}_R^{\mathbb{C}}$ that appear in the definition of U. This means that small complex errors have to be carried through the entire argument, making some of the statements much less transparent.

In this section, we quote a number of facts about the detailed structure of $\widehat{A}$, and go on to state some auxiliary results that address these technical difficulties. These are typically translations of trivial results on $\mathbb{T}$, and where possible we will make the connection explicit.

6.6.2 A description of $\widehat{A}$

Recall that A is the subring of $\mathbb{Q}(i)$ generated by $\mathbb{Z}[i]$, $1/\overline{\mathcal{P}_5}$ and $1/\overline{\mathcal{P}_{13}}$. That is, it consists of all elements of $\mathbb{Q}(i)$ whose denominators are of the form $\overline{\mathcal{P}_5}^r\overline{\mathcal{P}_{13}}^s$.

So far we have used nothing about $\widehat{A}$ other than its definition and topology. We now describe the characters on A explicitly.

This material is fairly standard and can be found in many places in the literature. The application of harmonic analysis to number fields is perhaps most famously associated with [61], and much of what follows can be extracted from that paper. The exact form of the results we need appears in [16, Section 3]. The author has found [17] to be a very good introduction to these ideas.

However, to keep things as approachable as possible to those unfamiliar with this material, this subsection takes on an expository flavour, while leaving rigorous details to the references.

6.6.2.1 A model case: $\widehat{\mathbb{Z}[1/2]}$

A slightly simpler case to consider is the dual group of $\mathbb{Z}[1/2]$, the ring of dyadic rationals treated as a discrete additive group. The dual $\widehat{\mathbb{Z}[1/2]}$ is often referred to as a *solenoid* or *the* 2*-solenoid*.

Clearly any continuous character on $\mathbb{R}$ restricts to one on $\mathbb{Z}[1/2]$. The continuous characters on $\mathbb{R}$ can be written

$$\begin{aligned} \chi_x \,:\, \mathbb{R} &\to \mathbb{T} \\ y &\mapsto \{x\,y\} \end{aligned}$$

where $\{\cdot\}$ denotes the fractional part; so, analogously to $\jmath_{\mathbb{C}} : \mathbb{C} \to \widehat{A}$ above, we get a map

$$\begin{aligned} \jmath_{\mathbb{R}} \,:\, \mathbb{R} &\to \widehat{\mathbb{Z}[1/2]} \\ x &\mapsto \chi_x|_{\mathbb{Z}[1/2]} \ . \end{aligned}$$

(It is not hard to verify that in fact $\jmath_{\mathbb{R}}$ is injective.)

As well as being contained in the reals, $\mathbb{Z}[1/2]$ is also contained in the 2-adic rationals $\mathbb{Q}_2$; indeed, $\mathbb{Q}_2$ is the completion of $\mathbb{Z}[1/2]$ with respect to the 2-adic metric. So analogously, any continuous character of $\mathbb{Q}_2$ restricts to a character of $\mathbb{Z}[1/2]$. The continuous characters of $\mathbb{Q}_2$ are all of the form:

$$\begin{aligned} \chi_a \,:\, \mathbb{Q}_2 &\to \mathbb{T} \\ b &\mapsto \{a\,b\}_2 \end{aligned}$$

where $a \in \mathbb{Q}_2$ and $\{x\}_2$ is the 2-*adic fractional part*, defined as the unique dyadic rational in $[0, 1)$ such that $x - \{x\}_2 \in \mathbb{Z}_2$ (the 2-adic integers). Hence we get another map

$$\begin{aligned} \jmath_{\mathbb{Q}_2} \,:\, \mathbb{Q}_2 &\to \widehat{\mathbb{Z}[1/2]} \\ a &\mapsto \chi_a|_{\mathbb{Z}[1/2]} \ . \end{aligned}$$

which is also injective. By combining the real and 2-adic characters on $\mathbb{Z}[1/2]$, we get a map

$$\begin{aligned} \jmath \,:\, \mathbb{R} \times \mathbb{Q}_2 &\to \widehat{\mathbb{Z}[1/2]} \\ (x, a) &\mapsto -\jmath_{\mathbb{R}}(x) + \jmath_{\mathbb{Q}_2}(a) \ . \end{aligned}$$

Now, $\jmath$ is not injective, and it is not too difficult to see that $\ker \jmath$ consists of pairs (r, r), where $r \in \mathbb{Z}[1/2]$ is treated as a real number and a 2-adic number respectively. Less straightforward is the fact that $\jmath$ is actually *surjective*. Assuming that fact, we get the following conclusion.

Proposition 6.6.1 (Structure of $\widehat{\mathbb{Z}[1/2]}$)**.**

(i) Consider the diagonal embedding

$$\begin{aligned} \imath^{\Delta} : \mathbb{Z}[1/2] &\to \mathbb{R} \times \mathbb{Q}_2 \\ r &\mapsto (r, r) \ . \end{aligned}$$

Then $\imath^{\Delta}(\mathbb{Z}[1/2])$ is a discrete and co-compact subgroup of $\mathbb{R} \times \mathbb{Q}_2$, and $\jmath$ (as defined above) gives rise to an isomorphism of topological groups

$$\tilde{\jmath} : (\mathbb{R} \times \mathbb{Q}_2)/\imath^{\Delta}(\mathbb{Z}[1/2]) \to \widehat{\mathbb{Z}[1/2]}$$

where the left hand side is given the natural (product, quotient) topology coming from the usual topologies on $\mathbb{R}$ and $\mathbb{Q}_2$. In particular, $\widehat{\mathbb{Z}[1/2]}$ is naturally a compact metric space.

(ii) A fundamental domain for $\imath^{\Delta}(\mathbb{Z}[1/2])$ is given by $[0, 1) \times \mathbb{Z}_2$. This does not give a rise to a topological isomorphism or an isomorphism of groups. It is, however, possible to simplify the above quotient to $(\mathbb{R} \times \mathbb{Z}_2)/\imath^{\Delta}(\mathbb{Z})$.

6.6.2.2 Applying to $\widehat{A}$

Much the same analysis carries over to $\widehat{A}$. (We will overload notation from the $\mathbb{Z}[1/2]$ case. Since we will never need to refer to $\mathbb{Z}[1/2]$ in the argument this should not cause too much confusion.)

We have already seen the complex characters in $\widehat{A}$ given by $\jmath_{\mathbb{C}}$. To define the appropriate non-Archimedean characters, we have to consider the non-Archimedean fields obtained by completing $\mathbb{Q}(i)$ with respect to the $\overline{\mathcal{P}_5}$- and $\overline{\mathcal{P}_{13}}$-adic metrics.

In fact, these completions are canonically (and topologically) isomorphic to the more usual non-Archimedean fields $\mathbb{Q}_5$ and $\mathbb{Q}_{13}$ respectively. The reason is that $\mathbb{Q}_5$ and $\mathbb{Q}_{13}$ already contain square roots of -1; sending i to one of these gives the isomorphism, and which root to choose is specified by the choice of e.g. $\overline{\mathcal{P}_5}$ rather than $\mathcal{P}_5$. In the sequel, we will talk in terms of $\mathbb{Q}_5$ and $\mathbb{Q}_{13}$ rather than the completions of $\mathbb{Q}(i)$ by $\overline{\mathcal{P}_5}$ and $\overline{\mathcal{P}_{13}}$; though we will continue to refer to the absolute values $|\cdot|_{\overline{\mathcal{P}_5}}$ and $|\cdot|_{\overline{\mathcal{P}_{13}}}$ on $\mathbb{Q}(i)$ from time to time.

Running the analysis as before gives the following result.

Proposition 6.6.2 (The structure of $\widehat{A}$)**.**

(i) There are canonically specified roots of -1, denoted $i_5 \in \mathbb{Q}_5$ and $i_{13} \in \mathbb{Q}_{13}$, such that the following holds. The field $\mathbb{Q}(i)$ is embedded in $\mathbb{Q}_5$ and $\mathbb{Q}_{13}$ by the unique field homomorphisms

$$\begin{aligned} \imath_{\mathbb{Q}_5} &: \mathbb{Q}(i) \to \mathbb{Q}_5 \\ \imath_{\mathbb{Q}_{13}} &: \mathbb{Q}(i) \to \mathbb{Q}_{13} \end{aligned}$$

given by taking i to i_p in each case. Under these embeddings, the absolute values $|\cdot|_{\overline{\mathcal{P}_p}}$ on $\mathbb{Q}(i)$ and $|\cdot|_p$ on $\mathbb{Q}_p$ agree (for $p \in \{5, 13\}$); e.g. $|\imath_{\mathbb{Q}_5}(q)|_5 = |q|_{\overline{\mathcal{P}_5}}$ for $q \in \mathbb{Q}(i)$, and similarly for $p = 13$.[2]

Also, $\mathbb{Q}(i)$ is embedded in $\mathbb{C}$ in the usual way, which we denote $\imath_{\mathbb{C}}$ for consistency.

(ii) Under the diagonal embedding

$$\begin{aligned} \imath^{\Delta} &: \mathbb{Q}(i) \to \mathbb{C} \times \mathbb{Q}_5 \times \mathbb{Q}_{13} \\ &\qquad q \mapsto (\imath_{\mathbb{C}}(q), \imath_{\mathbb{Q}_5}(q), \imath_{\mathbb{Q}_{13}}(q)) \end{aligned}$$

the image $\imath^{\Delta}(A)$ of A is discrete and co-compact with respect to the usual metric.

(iii) In addition to $\jmath_{\mathbb{C}} : \mathbb{C} \to \widehat{A}$ defined above, we define

$$\begin{aligned} \jmath_{\mathbb{Q}_5} &: \mathbb{Q}_5 \to \widehat{A} \\ &\qquad a \mapsto (r \mapsto \{a \cdot \imath_{\mathbb{Q}_5}(r)\}_5) \end{aligned}$$

where $\{\cdot\}_5$ is the 5-adic fractional part as above, and similarly

$$\begin{aligned} \jmath_{\mathbb{Q}_{13}} &: \mathbb{Q}_{13} \to \widehat{A} \\ &\qquad a \mapsto (r \mapsto \{a \cdot \imath_{\mathbb{Q}_{13}}(r)\}_{13}) \ . \end{aligned}$$

Combining these, we get a map

$$\begin{aligned} \jmath &: \mathbb{C} \times \mathbb{Q}_5 \times \mathbb{Q}_{13} \to \widehat{A} \\ &\qquad (z, a, b) \mapsto -\jmath_{\mathbb{C}}(z) + \jmath_{\mathbb{Q}_5}(a) + \jmath_{\mathbb{Q}_{13}}(b) \ . \end{aligned}$$

[2]These absolute values certainly agree up to an arbitrary exponent, and we choose normalizations to make them agree exactly.

(iv) The kernel of $\jmath$ is precisely $\imath^\Delta(A)$, and $\jmath$ is surjective, so $\jmath$ gives rise to a natural isomorphism of topological groups

$$\tilde{\jmath} : (\mathbb{C} \times \mathbb{Q}_5 \times \mathbb{Q}_{13})/\imath^\Delta(A) \to \widehat{A} .$$

(v) Consequently, $\widehat{A}$ is naturally a connected compact metric space, where the metric

$$d_{\widehat{A}}(x, y) = \inf \{|z|_{\mathbb{C}} + |a|_5 + |b|_{13} : (z, a, b) \in \mathbb{C} \times \mathbb{Q}_5 \times \mathbb{Q}_{13},\ \jmath(z, a, b) = x - y\}$$

is also translation-invariant.

(vi) A fundamental domain for $\imath^\Delta(A)$ is given by $[0, 1)^2 \times \mathbb{Z}_5 \times \mathbb{Z}_{13}$. This does not give rise to a topological isomorphism or an isomorphism of groups. However, it is possible to simplify the above quotient to $(\mathbb{C} \times \mathbb{Z}_5 \times \mathbb{Z}_{13})/\imath^\Delta(\mathbb{Z}[i])$, though we shall not use this fact.

(vii) The action ρ of Θ on $\widehat{A}$ defined above, is equivalent (under $\jmath$) to

$$\rho(\theta)(z, a, b) = (\imath_{\mathbb{C}}(\theta)z,\ \imath_{\mathbb{Q}_5}(\theta)a,\ \imath_{\mathbb{Q}_{13}}(\theta)b)$$

i.e. to multiplying by θ on each factor, combined with the appropriate embeddings. This action in fact makes sense on the whole of A, not just Θ.

Remarks on the proof. We have already given an overview of the approach to proving this. The content is in parts (ii) and (iv), for which we refer the reader to [16, Theorem 3.1] for a discussion of the rigorous details. The remaining parts are easy consequences of these together with well-known facts from the theory of local fields.

With this as background, we move on to some auxiliary results.

6.6.3 Torsion points and periodic points

Definition 6.6.3. For a positive integer m, we define $\Theta^{(m)} = \{{\theta_5}^{rm}{\theta_{13}}^{sm} : r, s \in \mathbb{Z}_{\geq 0}\}$, a subsemigroup of Θ.

We say a point $x \in \widehat{A}$ is *periodic* if $\Theta^{(m)}$ fixes x for some m.

Recall that, in the proof of the $\times 2$, $\times 3$ Theorem, we made an analogous definition of $\Phi^{(m)}$, and points fixed by $\Phi^{(m)}$ played an important role in the argument. We implicitly used the fact that the periodic points in the context of $\mathbb{T}$ are precisely rationals whose denominator is coprime to 2 and 3.

A related and completely trivial fact is that the torsion points of $\mathbb{T}$ are precisely the rationals. So, if $x \in \mathbb{T}$ is a torsion point then $2^r 3^s x$ is periodic for some $r, s \geq 0$.

In this subsection we transfer some of these results to $\widehat{A}$. First we classify the periodic and torsion points. Those proofs not given here can be found in Section 6.8.

Proposition 6.6.4. *Let $x \in \widehat{A}$, and take $(z, a, b) \in \mathbb{C} \times \mathbb{Q}_5 \times \mathbb{Q}_{13}$ any representative of x (i.e. $\jmath(z, a, b) = x$).*

(i) The following are equivalent:

(a) x is torsion;

(b) there exists $q \in \mathbb{Q}(i)$ such that $z = \imath_{\mathbb{C}}(q)$, $a = \imath_{\mathbb{Q}_5}(q)$ and $b = \imath_{\mathbb{Q}_{13}}(q)$; equivalently, $(z, a, b) \in \imath^{\Delta}(\mathbb{Q}(i))$;

(c) the orbit $\Theta(x)$ is finite.

(ii) The following are equivalent:

(a) x is periodic;

(b) there is some $q \in \mathbb{Q}(i)$ such that $z = \imath_{\mathbb{C}}(q)$, $a = \imath_{\mathbb{Q}_5}(q)$, $b = \imath_{\mathbb{Q}_{13}}(q)$, and additionally $|q|_{\mathcal{P}_5}, |q|_{\mathcal{P}_{13}} \leq 1$.

Note the appearance of $|\cdot|_{\mathcal{P}_5}$ and $|\cdot|_{\mathcal{P}_{13}}$ in (ii) (b); *not* $|\cdot|_{\overline{\mathcal{P}_5}}$ and $|\cdot|_{\overline{\mathcal{P}_{13}}}$.

One useful application is the following corollary.

Corollary 6.6.5. *If $x \in \widehat{A}$ is torsion then for some $\theta \in \Theta$, $\theta(x)$ is periodic.*

Proof. Given a $q \in \mathbb{Q}(i)$ satisfying (i)(b), we can pick a $\theta \in \Theta$ to remove powers of $\mathcal{P}_5$ and $\mathcal{P}_{13}$ from the denominator. Then $\theta\, q$ satisfies (ii)(b). □

6.6.4 Unions of complex balls

By a *complex ball* in $\widehat{A}$ we mean a set of the form $\overline{B}_R^{\mathbb{C}}(x)$, where we recall $\overline{B}_R^{\mathbb{C}}(0)$ is the image of the ball of radius R about 0 in the complex plane under the map $\jmath_{\mathbb{C}}$, and $\overline{B}_R^{\mathbb{C}}(x)$ is its translate by $x \in \widehat{A}$.

The proof of the $\times 2$, $\times 3$ Theorem made use of the following trivial facts.

(i) If $Y \subseteq \mathbb{T}$ is a finite $\Phi^{(m)}$-invariant set, then Y consists entirely of rationals.

(ii) If $x \in \mathbb{T}$ is irrational then its orbit $\Phi^{(m)}(x)$ is infinite.

In adapting these to our context, we not only need to talk about torsion points of $\widehat{A}$ in place of rationals, but also introduce small complex errors into the hypotheses and conclusions, as discussed in the introduction to this section. Here is the resulting analogous statement.

Proposition 6.6.6.

(i) Suppose $Y \subseteq \widehat{A}$ is closed, Θ-invariant, and contained in a finite union of complex balls (equivalently, of unit complex balls). Then it is contained in a finite union of complex balls with centers at torsion points of $\widehat{A}$.

(ii) Let $x \in \widehat{A}$. Then the orbit closure $\overline{\Theta(x)}$ is contained in a finite union of complex balls, if and only if x has the form $x = y + \jmath_{\mathbb{C}}(z)$ where y is torsion and $z \in \mathbb{C}$.

Moreover, the above holds replacing Θ by $\Theta^{(m)}$.

6.6.5 Non-Archimedean limits

Recall the crucial result for proof of the $\times 2$, $\times 3$ Theorem (Corollary 6.4.3): if $Y \subseteq \mathbb{T}$ is closed, invariant and has 0 as a limit point then $Y = \mathbb{T}$.

The naive generalization to $\widehat{A}$ – i.e. that if $Y \subseteq \widehat{A}$ is closed, invariant and has 0 as a limit point then $Y = \widehat{A}$ – is false, as again can be seen by taking $Y = \overline{B}^{\mathbb{C}}_{1/10}(0)$.

However, this counterexample is in some sense the only one. The view taken is as follows: in the same way that the singleton $\{0\} \subseteq \mathbb{T}$ does not have 0 as a limit point – the constant 0 sequence not being a valid way to approach 0 – similarly, $\overline{B}^{\mathbb{C}}_{1/10}(0)$ should not have 0 as a "proper" limit point, because sequences approaching 0 along a complex ball around 0 should be likewise invalid.

This is formalized in the following definition.

Definition 6.6.7. We say x is a *non-Archimedean limit point* of $Y \subseteq \widehat{A}$ if x is a limit point of $Y \setminus B^{\mathbb{C}}_{1/10}(x)$; equivalently, if there is a sequence $x_n \to x$ of points of Y such that $x_n - x \notin B^{\mathbb{C}}_{1/10}(0)$ (in which case we say x is a non-Archimedean limit of x_n). Denote by Y' the set of non-Archimedean limit points of Y.

Remark 6.6.8. The constant $1/10$ appearing in this definition is completely unimportant, and altering it does not affect the definition.

The analogue of the above-mentioned result will have to wait (see Lemma 6.7.1); for now we record some basic facts about these non-Archimedean limits.

Proposition 6.6.9. *For any $Y \subseteq \widehat{A}$, Y' is closed. If Y is $\Theta^{(m)}$-invariant, then so is Y'.*

This is, of course, analogous to the corresponding properties of the set of usual limit points of a set.

Proof. We can write Y' as:

$$Y' = \bigcap_{x \in \widehat{A}} \left(\overline{Y \setminus B^{\mathbb{C}}_{1/10}(x)} \right)$$

which is clearly closed. Invariance is clear. □

Also recall that we used the fact that an infinite set in a compact space has a limit point. Below is an analogue for non-Archimedean limits.

Proposition 6.6.10. *Suppose $Y \subseteq \widehat{A}$ is closed and not contained in a finite union of complex balls. Then Y' is non-empty, and $0 \in (Y - Y)'$.*

6.6.6 Density of $\mathbb{C}$, $\mathbb{Q}_5$ and $\mathbb{Q}_{13}$ in $\widehat{A}$

It is a standard fact that the embedded copies $\jmath_{\mathbb{C}}(\mathbb{C})$, $\jmath_{\mathbb{Q}_5}(\mathbb{Q}_5)$ and $\jmath_{\mathbb{Q}_{13}}(\mathbb{Q}_{13})$ of $\mathbb{C}$, $\mathbb{Q}_5$ and $\mathbb{Q}_{13}$ respectively in $\widehat{A}$, are all dense; equivalently, any $x \in \widehat{A}$ can be arbitrarily well approximated by complex characters, or by 5-adic or 13-adic characters.

We will only need this fact in the case of $\mathbb{Q}_5$ and $\mathbb{Q}_{13}$, and there we will actually need the following slight strengthening.

Proposition 6.6.11. *Suppose $J \leq \mathbb{Q}_5^\times$ is a finite index multiplicative subgroup, and H is some coset of J. Then $\jmath_{\mathbb{Q}_5}(H) = \jmath(0, H, 0)$ is dense in $\widehat{A}$. The same holds symmetrically for $\mathbb{Q}_{13}$.*

The proof is, as usual, in Section 6.8, and on this occasion is slightly lengthy. We mention one key step, namely the following related result about simultaneously approximating elements of $\mathbb{C}$, $\mathbb{Q}_5$ and $\mathbb{Q}_{13}$ by elements of A.

Proposition 6.6.12. *Let $z \in \mathbb{C}$ and $b \in \mathbb{Q}_5$ be given. For any $\delta > 0$, we can find an element $q \in A$ such that $|\imath_{\mathbb{C}}(q) - z|_{\mathbb{C}}$, $|\imath_{\mathbb{Q}_5}(q) - b|_5 \leq \delta$.*

The same holds swapping 5 and 13.

Proof. This is a special case of what is sometimes called the Strong Approximation Theorem (which is no more than the natural generalization of this statement); see [12, Chapter 10, Theorem 4.1]. □

6.6.7 Density of periodic points of $\widehat{A}$

In the $\times 2$, $\times 3$ Theorem, we noted that for all $\delta > 0$ there exists an m and a finite, δ-dense subset of $\mathbb{T}$ consisting of points fixed by $\Phi^{(m)}$. We now turn to the (completely unaltered) statement for $\widehat{A}$.

Proposition 6.6.13. *For all $\delta > 0$ there exists an $m \in \mathbb{Z}_{>0}$ and a finite set $S \subseteq \widehat{A}$ which is pointwise fixed by $\Theta^{(m)}$, such that S is δ-dense in $\widehat{A}$.*

We need the following result, very similar in flavour to Proposition 6.6.12 above. The number 7 is essentially arbitrary here.

Proposition 6.6.14. *Let $z \in \mathbb{C}$, $a \in \mathbb{Q}_5$ and $b \in \mathbb{Q}_{13}$ be given. For any $\delta > 0$ we can find $q \in A[1/7]$ (say) such that $|\imath_{\mathbb{C}}(q) - z|_{\mathbb{C}}$, $|\imath_{\mathbb{Q}_5}(q) - a|_5$, $|\imath_{\mathbb{Q}_{13}}(q) - b|_{13} \leq \delta$. Equivalently, $\imath^{\Delta}(A[1/7])$ is dense in $\mathbb{C} \times \mathbb{Q}_5 \times \mathbb{Q}_{13}$, or if you prefer $\jmath(\imath^{\Delta}(A[1/7]))$ is dense in $\widehat{A}$.*

Proof. This is also a special case of Strong Approximation; again see [12, Chapter 10, Theorem 4.1]. □

Proof of Proposition 6.6.13. We consider the sets $R_n = 7^{-n}A \subseteq \mathbb{Q}(i)$, and their images $S_n = \jmath\left(\imath^{\Delta}(R_n)\right) \subseteq \widehat{A}$ under the diagonal map. We note $S_n \cong A/7^n A$ and so is finite. By Proposition 6.6.14, we have that $\bigcup_n S_n$ is dense in $\widehat{A}$.

The fact that some S_n is δ-dense in $\widehat{A}$ follows from this and a routine compactness argument. Since the action of Θ permutes the finite set S_n, we may choose m such that θ_5^m and θ_{13}^m fix S_n pointwise, as required. □

6.7 Proof of the rationality lemma

6.7.1 The growth estimate, i.e. Lemma 6.4.2

Most of the hard work in running the proof of the $\times 2$, $\times 3$ Theorem is in proving the analogue of Lemma 6.4.2, which was used to show Corollary 6.4.3 which in turn was crucial for the inductive step of the main argument.

This proof is still highly technical. However, as we regard it as the core of the argument, we will place it here rather than in a later section. We begin with the following version.

Lemma 6.7.1. *Suppose $Y \subseteq \widehat{A}$ is a closed, Θ-invariant set with $\jmath_{\mathbb{C}}(w)$ as a non-Archimedean limit point for some $w \in \mathbb{C}$. Then $Y = \widehat{A}$.*

The same holds replacing Θ by $\Theta^{(m)}$ for any $m > 0$.

We need one or two more standard facts, the proofs of which again appear in Section 6.8.

Proposition 6.7.2. *Suppose $p > 2$ and $u \in \mathbb{Z}_p^\times$ is not a root of unity. Then $J_0 = \overline{\{u^r : r \geq 0\}}$ is a finite index subgroup of $\mathbb{Z}_p^\times$.*

Corollary 6.7.3. *Suppose $p > 2$, $u \in \mathbb{Z}_p^\times$ is not a root of unity, and $\alpha \in S^1$ is irrational (i.e. not torsion). Let $J_0 = \overline{\{u^r : r \geq 0\}}$ be as in Proposition 6.7.2. Then the closure*

$$\overline{\{(u^r, \alpha^r) \colon r \geq 0\}} \subseteq \mathbb{Z}_p \times S^1$$

is just $J_0 \times S^1$.

Finally, before embarking on the proof of Lemma 6.7.1, we introduce an abuse of notation: for the sake of visual clarity, where it is unambiguous to do so, we suppress the explicit use of $\imath_{\mathbb{C}}$, $\imath_{\mathbb{Q}_5}$ and $\imath_{\mathbb{Q}_{13}}$, implicitly treating $\mathbb{Q}(i)$ as a subset of $\mathbb{C}, \mathbb{Q}_5, \mathbb{Q}_{13}$ respectively.

Proof of Lemma 6.7.1. If $\jmath_{\mathbb{C}}(w) \in Y'$, we can choose a sequence x_n in Y such that $x_n \to \jmath_{\mathbb{C}}(w)$ is a non-Archimedean limit. We may identify $x_n - \jmath_{\mathbb{C}}(w)$ with points (z_n, a_n, b_n) of the fundamental domain $[-1/2, -1/2)^2 \times \mathbb{Z}_5 \times \mathbb{Z}_{13}$ of $\widehat{A}$. Then, z_n, a_n, b_n converge individually to 0.

Since the limit is non-Archimedean, we know that a_n and b_n cannot both be eventually zero. Suppose a_n is not eventually zero – the other case is symmetric. Passing to an appropriate subsequence we may assume $a_n \neq 0$ for all n.

Let J_0 denote the closure of $\{{\theta_{13}}^{mr} : r \in \mathbb{Z}_{\geq 0}\} \subseteq \mathbb{Z}_5^\times$. Since θ_{13} is not a root of unity in $\mathbb{Z}_5^\times$, we have that J_0 is a (closed) finite index subgroup of $\mathbb{Z}_5^\times$, by Proposition 6.7.2. Let J denote the subgroup of $\mathbb{Q}_5^\times$ spanned by J_0 and ${\theta_5}^m$; so J has finite index in $\mathbb{Q}_5^\times$ (this can be seen directly). Finally, pick H to be some coset of J such that $a_n \in H$ infinitely often; passing to an appropriate subsequence, assume $a_n \in H$ for all n. Set $H_0 = H \cap \mathbb{Z}_5^\times$ (a coset of J_0).

By Proposition 6.6.11, it suffices to show $\jmath(w, H, 0) \subseteq Y$, as $\jmath(0, H, 0)$ is dense in $\widehat{A}$ (and hence so is $\jmath(w, H, 0)$). So, fix any $\alpha \in H$; we will show that $\jmath(w, \alpha, 0) \in Y$. If $\alpha = 0$ this is trivial, so suppose it is not.

Fix a $\delta > 0$. Our aim is to find an n and integers $r, s \geq 0$ such that

$$d_{\widehat{A}}({\theta_{13}}^{mr}{\theta_5}^{ms}(x_n), \jmath(w, \alpha, 0)) \leq \delta \ .$$

Since ${\theta_{13}}^{mr}{\theta_5}^{ms}(x_n)$ lies in Y (by $\Theta^{(m)}$-invariance) and since Y is closed and δ is arbitrary, this again suffices.

The distance $d_{\widehat{A}}(\theta_{13}{}^{mr}\theta_5{}^{ms}(x_n), \jmath(w,\alpha,0))$ is bounded by the sum of pointwise distances:

$$|\theta_5{}^{ms}\theta_{13}{}^{mr}(z_n+w)-w|_{\mathbb{C}} + |\theta_5{}^{ms}\theta_{13}{}^{mr}a_n-\alpha|_5 + |\theta_5{}^{ms}\theta_{13}{}^{mr}b_n|_{13}\ .$$

We can immediately bound the first term by $|z_n|_{\mathbb{C}} + |\theta_5{}^{ms}\theta_{13}{}^{mr}w - w|_{\mathbb{C}}$.

The hard work is in approximating α by $\theta_5{}^{ms}\theta_{13}{}^{mr}$ while also guaranteeing that the $\mathbb{Q}_{13}$ term stays small as we do so. Of course, by choosing n large we can make b_n as small as we like; but some consideration is needed of the fact that r and s may themselves depend on n. As an added complication, $\theta_5{}^{ms}\theta_{13}{}^{mr}$ must simultaneously approximately fix w.

We first note that we may choose n sufficiently large that $|z_n|_{\mathbb{C}} \leq \delta/4$.

Next, we invoke Corollary 6.7.3 to pick N such that the following holds: for any $(x,t) \in J_0 \times S^1$, there is some r, $0 \leq r \leq N$, such that $|x - \theta_{13}^{mr}|_5 \leq |\alpha|_5^{-1}\,\delta/10$ and $|t-\theta_{13}^{mr}|_{\mathbb{C}} \leq |w|^{-1}\delta/4$. (Note this statement is again superficially stronger than Corollary 6.7.3 but follows from it easily by compactness.) If $w=0$ the latter condition is considered vacuous.

We can now consider the $\mathbb{Q}_{13}$ contribution. Again taking n sufficiently large, we can assume that $|b_n|_{13} \leq 13^{-Nm}\,\delta/4$; equivalently, that $|\theta_{13}{}^{mr}b_n|_{13} \leq \delta/4$ for all $0 \leq r \leq N$.

The remaining terms are the heart of the matter. Again taking n large enough, we may assume $|a_n|_5 \leq |\alpha|_5$. At this stage we fix some particular n, large enough in all the ways described above. It remains only to fix r and s.

For $\theta_5{}^{ms}\theta_{13}{}^{mr}a_n$ to approximate α, it should have the same valuation, which is controlled by s. So, we are forced to take $s = (v_5(a_n) - v_5(\alpha))/m \geq 0$ (recalling a_n is non-zero). As a_n and α lie in the same coset of J by assumption, we have that s is an integer and $u := \alpha/(\theta_5{}^{ms}a_n) \in J_0$.

Now we must choose r so that $\theta_{13}{}^{mr}$ approximates $(u, \theta_5{}^{-ms}) \in J_0 \times S^1$. Indeed, by our choice of N there is an r in the range $0 \leq r \leq N$ such that $|\theta_{13}{}^{mr} - u|_5 \leq |\alpha|_5^{-1}\,\delta/4$, whence

$$|\theta_{13}{}^{mr}\theta_5{}^{ms}a_n - \alpha|_5 = |\theta_{13}{}^{mr}-u|_5\,|\alpha|_5 \leq \delta/4\ ;$$

and also $|\theta_{13}{}^{mr} - \theta_5{}^{-ms}|_{\mathbb{C}} \leq |w|^{-1}\delta/4$, whence $|\theta_{13}{}^{mr}\theta_5{}^{ms}w - w|_{\mathbb{C}} \leq \delta/4$.

Putting everything together, we finally get:

$$\begin{aligned} &d_{\widehat{A}}(\theta_{13}{}^{mr}\theta_5{}^{ms}(x_n),\ \jmath(w,\alpha,0)) \\ &\qquad \leq |z_n|_{\mathbb{C}} + |\theta_{13}{}^{mr}\theta_5{}^{ms}w - w|_{\mathbb{C}} + |\theta_{13}{}^{mr}\theta_5{}^{ms}a_n - \alpha|_5 + |\theta_{13}{}^{mr}\theta_5{}^{ms}b_n|_{13} \\ &\qquad \leq \delta/4 + \delta/4 + \delta/4 + \delta/4 \leq \delta \end{aligned}$$

as required. □

In fact, Lemma 6.7.1 easily implies a slightly stronger version of itself.

Corollary 6.7.4. *Let $Y \subseteq \widehat{A}$ be a closed, Θ-invariant set, and $y = \jmath_{\mathbb{C}}(w) + t \in \widehat{A}$, where $w \in \mathbb{C}$ and t is a torsion point. Suppose y is a non-Archimedean limit point of Y. Then $Y = \widehat{A}$.*

Again, the same holds replacing Θ by $\Theta^{(m)}$.

Proof. By Corollary 6.6.5 we can choose $\theta_0 \in \Theta$ such that $\theta_0(t)$ is periodic. If x is periodic then so is $\theta(x)$ for any $\theta \in \Theta$, and so we can adjust θ_0 to lie in $\Theta^{(m)}$. Choose $m' > 0$ such that $m | m'$ and $\Theta^{(m')}$ fixes $\theta_0(t)$.

Since Y is a $\Theta^{(m)}$-invariant set which has $\jmath_{\mathbb{C}}(w) + t$ as a non-Archimedean limit point, it also has $\theta_0(\jmath_{\mathbb{C}}(w) + t)$ as a non-Archimedian limit point. Hence, $Y - \theta_0(t)$ is a $\Theta^{(m')}$-invariant set with $\jmath_{\mathbb{C}}(\theta_0 w)$ as a non-Archimedian limit point.

But by Lemma 6.7.1, this implies that $Y - \theta_0(t) = \widehat{A}$ and hence $Y = \widehat{A}$. □

6.7.2 The main argument

Finally, we finish the proof of the infinitary rationality lemma (Lemma 6.5.3). This part of the argument follows the proof of Theorem 6.4.1 almost line for line, which is possible only by repeatedly invoking results from Section 6.6.

Proof of Lemma 6.5.3. Let $Y \subseteq \widehat{A}$ be a closed invariant subset, and assume that Y is not contained in a finite union of complex balls with centers at torsion points.

By Proposition 6.6.6 (i) and Proposition 6.6.10, Y' is non-empty, and it is closed and invariant by Proposition 6.6.9. Also, we may assume Y' does not contain any points of the form $x = \jmath_{\mathbb{C}}(w) + t$ with $w \in \mathbb{C}$ and t torsion, as this would imply $Y = \widehat{A}$ by Corollary 6.7.4.

Fix an arbitrary $\delta > 0$ and choose m and $S \subseteq \widehat{A}$ according to Proposition 6.6.13; that is, S consists of fixed points of $\Theta^{(m)}$ and is δ-dense in $\widehat{A}$. Assume for convenience that $0 \in S$, and write $S = \{a_0, a_1, \ldots, a_{\ell-1}\}$ where $a_0 = 0$. For $0 \leq k < \ell$, define $X_k := (Y' - a_0) \cap \cdots \cap (Y' - a_k)$. Note that $X_0 = Y'$, X_k are nested, and furthermore $X_{k+1} = X_k \cap (Y' - a_{k+1})$ which implies (inductively) that X_k is closed and $\Theta^{(m)}$-invariant for all k (recalling that $\Theta^{(m)}$ fixes the a_i).

We claim inductively that X_k is non-empty; since $X_k \subseteq Y'$ and Y' contains no points $x = \jmath_{\mathbb{C}}(w) + t$ where $w \in \mathbb{C}$ and t is torsion, this immediately implies X_k is not contained in a finite union of complex balls (Proposition 6.6.6(ii)). The base case $k = 0$ is by assumption; suppose this holds for X_k. By Proposition 6.6.10, $X_k - X_k$

has 0 as a non-Archimedean limit point, and hence so does $Y' - X_k$. By Lemma 6.7.1, $Y' - X_k = \widehat{A}$ and in particular we can find $x \in X_k$, $y \in Y'$ such that $y - x = a_{k+1}$. I.e. $x \in (Y' - a_{k+1})$ and as $x \in X_k$ we deduce $x \in X_{k+1}$, completing the induction.

Hence for any $x \in X_{\ell-1}$, we have $\{x + a_0, x + a_1, \ldots, x + a_{\ell-1}\} \subseteq Y' \subseteq Y$. So Y contains a translate of S and hence is δ-dense in $\widehat{A}$. But $\delta > 0$ was arbitrary and Y is closed; hence $Y = \widehat{A}$ as required. □

6.8 Proofs of auxiliary results

Here we provide proofs of some of the less standard and more technical results of Section 6.6.

6.8.1 Torsion points and periodic points

Proof of Proposition 6.6.4. Consider (i). Unwrapping the definitions, we have that $n\,x = 0$ for $n \in \mathbb{Z}$ if and only if there is some $r \in A$ such that $n\,z = \imath_{\mathbb{C}}(r)$, $n\,a = \imath_{\mathbb{Q}_5}(r)$, $n\,b = \imath_{\mathbb{Q}_{13}}(r)$. Setting $q = r/n$ gives (a ⇒ b) and choosing $n \in \mathbb{Z}$ such that $n\,q \in A$ gives (b ⇒ a).

For (c ⇒ a), observe that by pigeonhole there exist $\phi \neq \psi \in \Theta$ such that $\phi(x) = \psi(x)$, whence $(\phi - \psi)(x) = 0$ and so, multiplying by $a \in \mathbb{Z}[i]$ such that $a(\phi - \psi) \in \mathbb{Z}$, we deduce x is torsion.

Finally, for (a, b ⇒ c) we note that $n\,x = 0 \Rightarrow n\theta(x) = 0$ for $\theta \in \Theta$, so it suffices to check that, for given n, the set $\{x \in \widehat{A} : n\,x = 0\}$ is finite. This is actually a group, and by (b) is isomorphic under $\imath^{\Delta}$ to the quotient $\{q \in \mathbb{Q}(i) : n\,q \in A\}/A$, which in turn is isomorphic (under $q \mapsto nq$) to A/nA. That this last group is finite is a standard fact.

We now turn to (ii). For (a ⇒ b), we write $\phi = (\theta_5\,\theta_{13})^m$ and suppose $\phi(x) = x$. That means there exists $r \in A$ such that $\imath_{\mathbb{C}}(\phi - 1)z = \imath_{\mathbb{C}}(r)$, $\imath_{\mathbb{Q}_5}(\phi - 1)a = \imath_{\mathbb{Q}_5}(r)$, $\imath_{\mathbb{Q}_{13}}(\phi - 1)b = \imath_{\mathbb{Q}_{13}}(r)$, so taking $q = r/(\phi - 1)$ it suffices to check that $|q|_{\mathcal{P}_5}$, $|q|_{\mathcal{P}_{13}} \leq 1$. Note $|r|_{\mathcal{P}_5}$, $|r|_{\mathcal{P}_{13}} \leq 1$ (since $r \in A$) and $|\phi|_{\mathcal{P}_5}$, $|\phi|_{\mathcal{P}_{13}} < 1$ since $m > 0$, so $|\phi - 1|_{\mathcal{P}_5}$, $|\phi - 1|_{\mathcal{P}_{13}} = 1$.

For (b ⇒ a), we observe as above that the orbit $\Theta(x)$ is contained in

$$\{x \in \widehat{A} : n\,x = 0\} \cong A/n\,A$$

for some $n \in \mathbb{Z}$, which crucially can now be taken to be coprime to 65. It follows that the corresponding multiplicative action of Θ on the ring $A/n\,A$ is now invertible, i.e. the image of Θ lies in $(A/n\,A)^{\times}$. Hence we can choose m to be $|(A/n\,A)^{\times}|$. □

6.8.2 Unions of complex balls

Proof of Proposition 6.6.6. We prove (i). Let $x_1, \ldots, x_k \in Y$ be such that $\bigcup_i \overline{B}_1^{\mathbb{C}}(x_i)$ covers Y. It suffices to show each x_i has the form $x_i = y_i + \jmath_{\mathbb{C}}(z_i)$ where y_i is torsion and $z_i \in \mathbb{C}$.

By pigeonhole, there exist distinct $\theta_1, \theta_2 \in \Theta^{(m)}$ such that $\theta_1(x_i)$ and $\theta_2(x_i)$ lie in the same ball $\overline{B}_1^{\mathbb{C}}(x_j)$. Hence $\theta_1(x_i) - \theta_2(x_i) \in \jmath_{\mathbb{C}}(\mathbb{C})$. Choosing $r \in \mathbb{Z}[i]$ such that $r(\theta_1 - \theta_2) = n \in \mathbb{Z}$ we obtain $n\, x_i \in \jmath_{\mathbb{C}}(\mathbb{C})$ as required.

Now consider (ii). $\Rightarrow$ follows trivially from (i): $\overline{\Theta^{(m)}(x)}$ is contained in a finite union of complex balls, is closed and invariant, and hence is contained in a finite union of complex balls with torsion centers, which *a fortiori* implies x itself has the specified form.

For $\Leftarrow$, we invoke Proposition 6.6.4 (i)(a $\Rightarrow$ c) and note the trivial fact that the orbit closure of a point $\jmath_{\mathbb{C}}(w)$, $w \in \mathbb{C}$ is contained in (in fact, is equal to) $\jmath_{\mathbb{C}}\left(\{z \in \mathbb{C} : |z| = |w|\}\right)$, which is certainly contained in some $\overline{B}_R^{\mathbb{C}}(0)$. □

6.8.3 Non-Archimedean limits

Proof of Proposition 6.6.10. By the hypothesis on Y, we can pick a sequence $x_i \in Y$ such that $x_i - x_j \notin \overline{B}_1^{\mathbb{C}}(0)$ for all $i \neq j$. Since Y is compact metric, passing to some subsequence we have $x_i \to x$ for some $x \in Y$.

It would suffice to show that $x_i - x \notin \overline{B}_{1/2}^{\mathbb{C}}(0)$ for all i, as then $x_i \to x$ and $x_i - x \to 0$ would be non-Archimedean limits in Y and $Y - Y$ respectively.

In fact this holds for all i with at most one exception: if $x_i - x \in \overline{B}_{1/2}^{\mathbb{C}}(0)$ and $x_j - x \in \overline{B}_{1/2}^{\mathbb{C}}(0)$ for $i \neq j$ then $x_i - x_j \in \overline{B}_1^{\mathbb{C}}(0)$, contradicting the choice of the x_i. Deleting the exceptional index (if there is one) gives the result. □

6.8.4 Density of $\mathbb{C}$, $\mathbb{Q}_5$ and $\mathbb{Q}_{13}$ in $\widehat{A}$

We need a standard fact for the proof of Proposition 6.6.11.

Proposition 6.8.1. *Any finite index multiplicative subgroup of $\mathbb{Z}_p^\times$ contains a subgroup $(1 + p^n\mathbb{Z}_p, \times)$ for some n.*

Proof. There is an isomorphism $(1 + p\mathbb{Z}_p, \times) \longleftrightarrow (\mathbb{Z}_p, +)$ given by the p-adic log and exp maps (for $p > 2$); see e.g. [51, Chapter II, Section 3, Proposition 8].

So if $J \leq \mathbb{Z}_p^\times$ has finite index then $J \cap (1 + p\mathbb{Z}_p)$ has finite index in $1 + p\mathbb{Z}_p$, which corresponds to a finite index subgroup of $(\mathbb{Z}_p, +)$ under the isomorphism; those are

all $(p^n\mathbb{Z}_p, +)$ for some n, which correspond to $(1+p^{n+1}\mathbb{Z}_p, \times)$ under the isomorphism the other way. □

Corollary 6.8.2. *Let J be a finite index subgroup of $\mathbb{Q}_p^\times$. There exists some $\eta > 0$ depending only on J, such that if $x, y \in \mathbb{Q}_p^\times$ and $|x-y|_p/|x|_p \leq \eta$ then x and y are in the same coset of J.*

Proof. We require $y/x \in J$. By Proposition 6.8.1 applied to the finite index subgroup $J_0 = J \cap \mathbb{Z}_p^\times \leq \mathbb{Z}_p^\times$, there is some $\eta > 0$ such that if $a \in \mathbb{Z}_p^\times$ and $|1-a|_p \leq \eta$ then $a \in J_0$. So, provided $\eta \leq 1$, we have $|1-a|_p \leq \eta \Rightarrow a \in J_0$ for all $a \in \mathbb{Q}_p^\times$. But $|1-y/x|_p = |x-y|_p/|x|_p$ and the result follows. □

We isolate one further result from the proof of Proposition 6.6.11.

Lemma 6.8.3. *Let J, H be as in the statement of Proposition 6.6.11. For any $\mu, \nu > 0$ we can find an $r \in A$ such that $|r|_{\mathbb{C}}$, $|r|_{\overline{\mathcal{P}_{13}}} \leq \mu$, $|r|_{\overline{\mathcal{P}_5}} \geq \nu$ and $\imath_{\mathbb{Q}_5}(r) \in H$. The same holds with 5 and 13 exchanged.*

Proof. Let $x_1, \ldots, x_k \in A$ be coset representatives for J; i.e. $\bigcup_i x_i\, J = \mathbb{Q}_5^\times$. (This is possible by Corollary 6.8.2 and the fact that A is dense in $\mathbb{Q}_p^\times$.) Let

$$C = \max\left\{|x_i|_{\mathbb{C}},\, |x_i|_{\overline{\mathcal{P}_{13}}},\, 1/|x_i|_{\overline{\mathcal{P}_5}}\right\}\ .$$

Pick any $y \in A$ such that $|y|_{\mathbb{C}}$, $|y|_{\overline{\mathcal{P}_{13}}} < 1$ and $|y|_{\overline{\mathcal{P}_5}} > 1$ (say, $y = \overline{\mathcal{P}_{13}}/\overline{\mathcal{P}_5}^{10}$). Then choose n large enough that $|y^n|_{\mathbb{C}}$, $|y^n|_{\overline{\mathcal{P}_{13}}} \leq \mu/C$ and $|y^n|_{\overline{\mathcal{P}_5}} \geq C\,\nu$. Now take i such that $x_i\, y^n \in H$, and observe $r = x_i\, y^n$ has the desired properties. □

Proof of Proposition 6.6.11. Let $x = \jmath(z, a, b) \in \widehat{A}$, with (z, a, b) in the fundamental domain (say). We take $\delta > 0$ arbitrary, and choose $q \in A$ such that $|\imath_{\mathbb{C}}(q) - z|_{\mathbb{C}}$, $|\imath_{\mathbb{Q}_{13}}(q) - b|_{13} \leq \delta/4$ (by Proposition 6.6.12).

If $a - \imath_{\mathbb{Q}_5}(q) \in H$ we are happy, as then

$$d_{\widehat{A}}((z,a,b),\, (0,\, a-\imath_{\mathbb{Q}_5}(q),\, 0)) = d_{\widehat{A}}((z,a,b),\, (\imath_{\mathbb{C}}(q),\, a,\, \imath_{\mathbb{Q}_{13}}(q))) \leq \delta/2\ .$$

If not, the above results allow us to perturb q to some $q - r$ so that this holds (i.e. $a - \imath_{\mathbb{Q}_5}(q-r) \in H$), without affecting the other properties of q too much.

Specifically, take $\eta > 0$ as in Corollary 6.8.2, and r as in Lemma 6.8.3 with parameters $\mu = \delta/4$ and $\nu = 10|a - \imath_{\mathbb{Q}_5}(q)|_5/\eta$. We claim $a - \imath_{\mathbb{Q}_5}(q-r) \in H$. Indeed,

as $\imath_{\mathbb{Q}_5}(r) \in H$ by construction and $|a - \imath_{\mathbb{Q}_5}(q)|_5 / |\imath_{\mathbb{Q}_5}(r)|_5 \leq \eta/10$, this follows by Corollary 6.8.2. Finally,

$$\begin{aligned} d_{\widehat{A}}((z,a,b),\,(0,a-\imath_{\mathbb{Q}_5}(q-r),0)) &= d_{\widehat{A}}((z,a,b),\,(\imath_{\mathbb{C}}(q-r),\,a,\,\imath_{\mathbb{Q}_{13}}(q-r)) \\ &\leq |z-\imath_{\mathbb{C}}(q)|_{\mathbb{C}} + |r|_{\mathbb{C}} + |\imath_{\mathbb{Q}_{13}}(q)-b|_{\overline{\mathcal{P}_{13}}} + |r|_{\overline{\mathcal{P}_{13}}} \\ &\leq \delta \end{aligned}$$

which, as δ was arbitrary, completes the proof. □

6.8.5 Proof of Proposition 6.7.2 and Corollary 6.7.3

Proof of Proposition 6.7.2. Using [51, Chapter II, Section 3, Proposition 7], we write $u = \xi\, v$ where ξ is a $(p-1)^{\text{st}}$ root of unity and $v \in 1 + p\mathbb{Z}_p$. By assumption $v \neq 1$. It suffices to show $\overline{\{v^{(p-1)r} : r \geq 0\}}$ is a finite index subgroup of $(1+p\mathbb{Z}_p, \times)$. Using the p-adic log/exp isomorphism ([51, Chapter II, Section 3, Proposition 8]) this reduces to showing that for $x \in \mathbb{Z}_p$ non-zero, the closure $\overline{\{(p-1)r\,x : r \geq 0\}}$ is a finite index subgroup of $(\mathbb{Z}_p, +)$; and in fact it is precisely $p^{v_p(x)}\mathbb{Z}_p$. □

Proof of Corollary 6.7.3. Fix $(x,t) \in J_0 \times S^1$, and let $\varepsilon > 0$ be given. Choose an integer $k \geq 0$ such that $p^{-k} \leq \varepsilon/2$.

We may choose an integer $\ell > 0$ such that $u^\ell \equiv 1 \pmod{p^k}$. Also, by definition of J_0, we may pick $a \geq 0$ such that $u^a \equiv x \pmod{p^k}$; and so the same holds replacing a with $a + n\ell$ for any n.

But since α is irrational so is α^ℓ, and hence we may choose n such that $d(\alpha^{\ell n}, t \cdot \alpha^{-a}) \leq \varepsilon/2$. Setting $r = a + \ell n$, it follows that $d((x,t),(u^r,\alpha^r)) \leq \varepsilon$ as required. □

References

[1] William Beckner, *Inequalities in Fourier analysis*, Ann. of Math. (2) **102** (1975), no. 1, 159–182. MR 0385456 (52 #6317)

[2] Daniel Berend, *Multi-invariant sets on compact abelian groups*, Trans. Amer. Math. Soc. **286** (1984), no. 2, 505–535. MR 760973 (86e:22009)

[3] Vitaly Bergelson and Alexander Leibman, *Distribution of values of bounded generalized polynomials*, Acta Math. **198** (2007), no. 2, 155–230. MR 2318563 (2008m:11149)

[4] Yuri Bilu, *Addition of sets of integers of positive density*, J. Number Theory **64** (1997), no. 2, 233–275. MR 1453212 (98e:11013)

[5] ———, *Structure of sets with small sumset*, Astérisque (1999), no. 258, xi, 77–108, Structure theory of set addition. MR 1701189 (2000h:11109)

[6] Béla Bollobás and Andrew Thomason, *Projections of bodies and hereditary properties of hypergraphs*, Bull. London Math. Soc. **27** (1995), no. 5, 417–424. MR 1338683

[7] M. Bóna, *Handbook of enumerative combinatorics*, Chapman and Hall/CRC, 2015.

[8] Michael D. Boshernitzan, *Elementary proof of Furstenberg's Diophantine result*, Proc. Amer. Math. Soc. **122** (1994), no. 1, 67–70. MR 1195714 (94k:11085)

[9] Herm Jan Brascamp and Elliott H. Lieb, *Best constants in Young's inequality, its converse, and its generalization to more than three functions*, Advances in Math. **20** (1976), no. 2, 151–173. MR 0412366 (54 #492)

[10] Omar Antolin Camarena and Balazs Szegedy, *Nilspaces, nilmanifolds and their morphisms*, Preprint, arXiv:1009.3825.

[11] Pablo Candela and Olof Sisask, *Convergence results for systems of linear forms on cyclic groups and periodic nilsequences*, SIAM J. Discrete Math. **28** (2014), no. 2, 786–810. MR 3211031

[12] J. W. S. Cassels, *Local fields*, London Mathematical Society Student Texts, vol. 3, Cambridge University Press, Cambridge, 1986. MR 861410 (87i:11172)

[13] N. J. Cavenagh and I. M. Wanless, *On the number of transversals in Cayley tables of cyclic groups*, Discrete Appl. Math. **158** (2010), no. 2, 136–146. MR 2563279 (2011a:05052)

[14] Mei-Chu Chang, *A polynomial bound in Freiman's theorem*, Duke Math. J. **113** (2002), no. 3, 399–419. MR 1909605 (2003d:11151)

[15] ———, *Some consequences of the polynomial Freiman-Ruzsa conjecture*, C. R. Math. Acad. Sci. Paris **347** (2009), no. 11-12, 583–588. MR 2532910 (2010e:11005)

[16] V. Chothi, G. Everest, and T. Ward, *S-integer dynamical systems: periodic points*, J. Reine Angew. Math. **489** (1997), 99–132. MR 1461206 (99b:11089)

[17] K. Conrad, *The character group of* $\mathbb{Q}$, Unpublished; available at `http://www.math.uconn.edu/~kconrad/blurbs/gradnumthy/characterQ.pdf`.

[18] C. Cooper, *A lower bound for the number of good permutations*, Data Recording, Storage and Processing (Nat. Acad. Sci. Ukraine) **213** (2000), 15–25.

[19] C. Cooper, R. Gilchrist, I. N. Kovalenko, and D. Novaković, *Deriving the number of "good" permutations, with applications to cryptography*, Kibernet. Sistem. Anal. (1999), no. 5, 10–16, 187. MR 1741073 (2000k:94029)

[20] C. Cooper and I. N. Kovalenko, *The upper bound for the number of complete mappings*, Theory Probab. Math. Statist. **53** (1996), 77–83.

[21] Sean Eberhard, Freddie Manners, and Rudi Mrazović, *Additive triples of bijections, or the toroidal semiqueens problem*, Preprint, arXiv:1510.05987.

[22] Tanja Eisner and Terence Tao, *Large values of the Gowers-Host-Kra seminorms*, J. Anal. Math. **117** (2012), 133–186. MR 2944094

[23] P. Flajolet and R. Sedgewick, *Analytic combinatorics*, Cambridge University Press, Cambridge, 2009. MR 2483235 (2010h:05005)

[24] John J. F. Fournier, *Sharpness in Young's inequality for convolution*, Pacific J. Math. **72** (1977), no. 2, 383–397. MR 0461034 (57 #1021)

[25] G. A. Freĭman, *Foundations of a structural theory of set addition*, American Mathematical Society, Providence, R. I., 1973, Translated from the Russian, Translations of Mathematical Monographs, Vol 37. MR 0360496 (50 #12944)

[26] Gregory A. Freĭman, *Structure theory of set addition*, Astérisque (1999), no. 258, xi, 1–33, Structure theory of set addition. MR 1701187 (2000j:11147)

[27] H. Furstenberg, *Recurrence in ergodic theory and combinatorial number theory*, Princeton University Press, Princeton, N.J., 1981, M. B. Porter Lectures. MR 603625 (82j:28010)

[28] Harry Furstenberg, *Disjointness in ergodic theory, minimal sets, and a problem in Diophantine approximation*, Math. Systems Theory **1** (1967), 1–49. MR 0213508 (35 #4369)

[29] R. Glebov and Z. Luria, *On the maximum number of latin transversals*, (2015), Preprint available at http://arxiv.org/pdf/1506.00983.pdf.

[30] W. T. Gowers, *A new proof of Szemerédi's theorem*, Geom. Funct. Anal. **11** (2001), no. 3, 465–588. MR 1844079

[31] B. Green and T. Tao, *Compressions, convex geometry and the Freiman-Bilu theorem*, Q. J. Math. **57** (2006), no. 4, 495–504. MR 2277597 (2007g:11013)

[32] Ben Green, *Hadamard lectures on nilsequences*, In preparation.

[33] Ben Green and Terence Tao, *An inverse theorem for the Gowers $U^3(G)$ norm*, Proc. Edinb. Math. Soc. (2) **51** (2008), no. 1, 73–153. MR 2391635

[34] ———, *An equivalence between inverse sumset theorems and inverse conjectures for the U^3 norm*, Math. Proc. Cambridge Philos. Soc. **149** (2010), no. 1, 1–19. MR 2651575 (2011g:11019)

[35] Ben Green, Terence Tao, and Tamar Ziegler, *An inverse theorem for the Gowers U^4-norm*, Glasg. Math. J. **53** (2011), no. 1, 1–50. MR 2747135

[36] ———, *An inverse theorem for the Gowers $U^{s+1}[N]$-norm.*, Ann. Math. (2) **176** (2012), no. 2, 1231–1372 (English).

[37] A. Iosevich, M. N. Kolountzakis, and M. Matolcsi, *Covering the plane by rotations of a lattice arrangement of disks*, ArXiv Mathematics e-prints (2006).

[38] I. N. Kovalenko, *On an upper bound for the number of complete mappings*, Kibernet. Sistem. Anal. (1996), no. 1, 81–85, 188. MR 1409532 (97g:05002)

[39] N. Yu. Kuznetsov, *Applying fast simulation to find the number of good permutations*, Cybernet. Systems Anal. **43** (2007), 830–837.

[40] ———, *Estimating the number of good permutations by a modified fast simulation method*, Cybernet. Systems Anal. **44** (2008), 547–554.

[41] ———, *Estimating the number of latin rectangles by the fast simulation method*, Cybernet. Systems Anal. **45** (2009), 69–75.

[42] R. D. Malikiosis, M. Matolcsi, and I. Z. Ruzsa, *A note on the pyjama problem*, ArXiv e-prints (2012).

[43] Freddie Manners, *Finding a low-dimensional piece of a set of integers*, Preprint, arXiv:1512.06272.

[44] ———, *Periodic nilsequences and inverse theorems on cyclic groups*, Preprint, arXiv:1404.7742.

[45] ———, *A solution to the pyjama problem*, Invent. Math. **202** (2015), no. 1, 239–270. MR 3402799

[46] B. D. McKay, J. C. McLeod, and I. M. Wanless, *The number of transversals in a Latin square*, Des. Codes Cryptogr. **40** (2006), no. 3, 269–284. MR 2251320 (2007f:05030)

[47] Roman Muchnik, *Semigroup actions on* $\mathbb{T}^n$, Geom. Dedicata **110** (2005), 1–47. MR 2136018 (2006i:37022)

[48] I. Rivin, I. Vardi, and P. Zimmermann, *The n-queens problem*, Amer. Math. Monthly **101** (1994), no. 7, 629–639. MR 1289272 (95d:05009)

[49] Imre Z. Ruzsa, *An analog of Freiman's theorem in groups*, Astérisque (1999), no. 258, xv, 323–326, Structure theory of set addition. MR 1701207 (2000h:11111)

[50] Tom Sanders, *On the Bogolyubov-Ruzsa lemma*, Anal. PDE **5** (2012), no. 3, 627–655. MR 2994508

[51] J.-P. Serre, *A course in arithmetic*, Springer-Verlag, New York, 1973, Translated from the French, Graduate Texts in Mathematics, No. 7. MR 0344216 (49 #8956)

[52] Xuancheng Shao, *Large values of the additive energy in* $\mathbb{R}^d$ *and* $\mathbb{Z}^d$, Math. Proc. Cambridge Philos. Soc. **156** (2014), no. 2, 327–341. MR 3177873

[53] Ilya Shkredov, *Energies and structure of additive sets*, Electron. J. Combin. **21** (2014), no. 3, Paper 3.44, 53. MR 3262281

[54] Carl Ludwig Siegel, *Lectures on the geometry of numbers*, Springer-Verlag, Berlin, 1989, Notes by B. Friedman, Rewritten by Komaravolu Chandrasekharan with the assistance of Rudolf Suter, With a preface by Chandrasekharan. MR 1020761

[55] N. J. A. Sloane, *The on-line encyclopedia of integer sequences*, `https://oeis.org/`.

[56] Balász Szegedy, *On higher order Fourier analysis*, Preprint, arXiv:1203.2260.

[57] László Székelyhidi, *Note on a stability theorem*, Canad. Math. Bull. **25** (1982), no. 4, 500–501. MR 674571

[58] Terence Tao, *Higher order Fourier analysis*, Graduate Studies in Mathematics, vol. 142, American Mathematical Society, Providence, RI, 2012. MR 2931680

[59] Terence Tao and Van H. Vu, *Additive combinatorics*, Cambridge Studies in Advanced Mathematics, vol. 105, Cambridge University Press, Cambridge, 2010, Paperback edition [of MR2289012]. MR 2573797

[60] A. A. Taranenko, *Multidimensional permanents and an upper bound on the number of transversals in Latin squares*, J. Combin. Des. **23** (2015), no. 7, 305–320. MR 3350096

[61] J. T. Tate, *Fourier analysis in number fields, and Hecke's zeta-functions*, Algebraic Number Theory (Proc. Instructional Conf., Brighton, 1965), Thompson, Washington, D.C., 1967, pp. 305–347. MR 0217026 (36 #121)

[62] I. Vardi, *Computational recreations in Mathematica*, Addison-Wesley Publishing Company, Advanced Book Program, Redwood City, CA, 1991. MR 1150054 (93e:00002)

[63] Zhiren Wang, *Rigidity of commutative non-hyperbolic actions by toral automorphisms*, Ergodic Theory Dynam. Systems **32** (2012), no. 5, 1752–1782. MR 2974218