Quantitative results in arithmetic combinatorics

Thomas F. Bloom

A dissertation submitted to the University of Bristol in accordance with the requirements for award of the degree of Doctor of Philosophy in the Faculty of Science.

School of Mathematics

July 2014

ABSTRACT

In this thesis we study the generalisation of Roth's theorem on three term arithmetic progressions to arbitrary discrete abelian groups and translation invariant linear equations. We prove a new structural result concerning sets of large Fourier coefficients and use this to prove new quantitative bounds on the size of finite sets which contain only trivial solutions to a given translation invariant linear equation. In particular, we obtain a quantitative improvement for Roth's theorem on three term arithmetic progressions and its analogue over $\mathbb{F}_q[t]$, the ring of polynomials in t with coefficients in a finite field $\mathbb{F}_q$.

We prove arithmetic inverse results for $\mathbb{F}_q[t]$ which characterise finite sets A such that $|A + t \cdot A| / |A|$ is small. In particular, when $|A + A| \ll |A|$ we prove a quantitatively optimal result, which is the $\mathbb{F}_q[t]$-analogue of the Polynomial Freiman-Ruzsa conjecture in the integers.

In joint work with Timothy G. F. Jones we prove new sum-product estimates for finite subsets of $\mathbb{F}_q[t]$, and more generally for any local fields, such as $\mathbb{Q}_p$. We give an application of these estimates to exponential sums.

To Hobbes, who had better things to think about.

ACKNOWLEDGEMENTS

I would like to thank my supervisor Professor Trevor Wooley, who has been a constant source of support, advice, and inspiration.

Over the past four years I have benefitted greatly from stimulating discussions with many mathematicians; thanks in particular to Tom Sanders, Julia Wolf, Ernie Croot, Olof Sisask, Giorgis Petridis, Kevin Henriot, and Thomas Hudson. The JACKET seminars were a fun and productive forum for discussion; thanks to everyone involved, and especially Tim Jones for providing a non-trivial upper bound for my Erdős number.

My family has been unwavering in their love and support, for which I am immensely grateful. I would like to thank my friends for their ceaseless efforts to distract me from maths, and their partial success. Finally, I would like to thank my wife Julia, for keeping me sane and adding the commas.

AUTHOR'S DECLARATION

I declare that the work in this dissertation was carried out in accordance with the requirements of the University's Regulations and Code of Practice for Research Degree Programmes and that it has not been submitted for any other academic award. Except where indicated by specific reference in the text, the work is the candidate's own work. Work done in collaboration with, or with the assistance of, others, is indicated as such. Any views expressed in the dissertation are those of the author.

SIGNED: ..... DATE:.....

CONTENTS

CHAPTER 1

INTRODUCTION

This thesis is concerned with arithmetic combinatorics, which is a field of relatively recent vintage. It is difficult to define precisely, but can be characterised as the study of approximate structure in finite subsets of abelian groups (or, more generally, any ring).

In this thesis we prove new quantitative results related to three classical problems of arithmetic combinatorics: locating three term arithmetic progressions, characterising those sets A with small doubling constant $|A+A|/|A|$, and showing that any finite set grows under at least one of addition and multiplication (commonly known as the sum-product phenomenon).

A common thread between our results is the emphasis on studying these problems over $\mathbb{F}_q[t]$, the ring of polynomials in t with coefficients in $\mathbb{F}_q$. The traditional setting for most problems in arithmetic combinatorics is the integers $\mathbb{Z}$, but their analogues in $\mathbb{F}_p^\infty$, the infinite-dimensional vector space over $\mathbb{F}_p$, are also increasingly studied. Due to its rigid additive structure many arguments used for $\mathbb{Z}$ are considerably simplified over $\mathbb{F}_p^\infty$, while most of their salient features are preserved. This also makes $\mathbb{F}_p^\infty$ a very useful setting for exploring the potential of new methods in arithmetic combinatorics. Often, having proved a result over $\mathbb{F}_p^\infty$, the extension to $\mathbb{Z}$ poses purely technical (though still challenging) difficulties. For a survey of arithmetic combinatorics in $\mathbb{F}_p^\infty$ and its relation to the classical integer problems we refer the reader to the survey by Green [28].

There is a well-known analogy in number theory between $\mathbb{Z}$ and $\mathbb{F}_q[t]$, and the latter is often used as an arithmetic model case for the former, although the model is not always perfect. This is perhaps best exemplified by the fact that Weil solved the $\mathbb{F}_q[t]$-analogue of the Riemann hypothesis in the 1940s, while the integer version remains intractable. There has been much interest in problems of additive number theory translated to the $\mathbb{F}_q[t]$-setting; we mention, for example, the work of Kubota [37] and Liu and Wooley [41] on Waring's problem over $\mathbb{F}_q[t]$.

It is perhaps surprising, then, that $\mathbb{F}_q[t]$ has received comparatively little attention in arithmetic combinatorics. It is useful both as a way to study the arithmetic combinatorics

of $\mathbb{F}_p^\infty$ which involves behaviour more exotic than simple addition (for example, the 'shift map' which in $\mathbb{F}_q[t]$ corresponds to multiplication by t), and also as a model for the arithmetic combinatorics of $\mathbb{Z}$. We have found that $\mathbb{F}_q[t]$ can act as an intermediary between $\mathbb{F}_p^\infty$ and $\mathbb{Z}$, possessing the rigid structure of the former and the rich arithmetic behaviour of the latter. We hope that this thesis gives some demonstration of its potential.

1.1 Translation invariant equations

1.1.1 Three term arithmetic progressions

How large can a set $A \subset \{1, \dots, N\}$ be while not containing any non-trivial three term arithmetic progressions; that is, a sequence of the shape $a, a+d, a+2d$ with $d \neq 0$? This is one of the core problems of arithmetic combinatorics.

The question was originally inspired by a theorem of van der Waerden [76], one of the earliest results in Ramsey theory: if the integers are partitioned into finitely many parts then one of the parts must contain arbitrarily long arithmetic progressions. Erdős and Turán [21] conjectured that this was true simply because one of the sets in such a partition must contain a positive proportion of the integers, and that any set of similar density must also contain arbitrarily long arithmetic progressions.

We will restrict our attention to three term arithmetic progressions, which are the first non-trivial case. Let $R(N)$ denote the size of the largest subset of $\{1, \dots, N\}$ that contains no non-trivial three term arithmetic progression (we will henceforth implicitly assume that a three term progression is non-trivial), and let $r(N) = R(N)N^{-1}$. The problem of estimating $r(N)$ was first considered by Erdős and Turán [21], inspired by the result of van der Waerden. Their best estimate was that, for any $\epsilon > 0$, if N is sufficiently large depending on ϵ then

$$r(N) < \frac{3}{8} + \epsilon.$$

The arguments of [21] were based on the trivial inequality $R(M+N) \leq R(M) + R(N)$, and explicit calculations of $R(N)$ for small N. For example, since $R(8) = 4$ it follows that $r(N) \leq 1/2 + 3/N$ for all $N \geq 1$. As mentioned above, Erdős and Turán conjectured in [21] that

$$r(N) = o(1),$$

which can be viewed as a quantitatively stronger form of van der Waerden's theorem for three term arithmetic progressions. This conjecture was resolved by Roth [50] with an inspired elaboration of the circle method. Roth's method was powerful enough to give the explicit quantitative bound

$$r(N) \ll \frac{1}{\log\log N}.$$

There has since been much effort spent on improving this quantitative bound. Since the original work of Roth [50] other, quite distinct, proofs that $r(N) = o(N)$ have been given; perhaps most strikingly, Furstenberg [23] gave a proof using ergodic theory. Such alternative methods have, however, never managed to improve upon the original bound of Roth; subsequent quantitative improvements have always used the analytic framework provided by Roth, now commonly referred to as the density increment method.

The idea is simple: if a subset $A \subset \{1, \dots, N\}$ fails to have three term arithmetic progressions and N is sufficiently large then it must have increased density on some smaller arithmetic progression $P \subset \{1, \dots, N\}$. Most crucially the property of containing no three term arithmetic progressions is preserved under dilations and translations, and hence we have some $M \leq N$ and $A' \subset \{1, \dots, M\}$ with density strictly larger than the density of A within $\{1, \dots, N\}$. We can now iterate this argument; since the density is always at most 1, however, we can only perform this argument a bounded number of times. The only reason we must be forced to halt is if we are working within $\{1, \dots, N'\}$ where N' is too small for this argument to be carried out. Backtracking to the original density α gives an upper bound for α in terms of N as required, provided at each step M is not too small compared to N.

The quantitative improvements of upper bounds on $r(N)$ since the work of Roth [50] have all followed this basic strategy, although the details have become quite sophisticated. It is worth mentioning why such quantitative improvements, above the estimate $r(N) = o(N)$, are useful. It has long been observed that if one could show that

$$r(N) \ll \frac{(\log\log N)^{1-\delta}}{\log N}$$

for some $\delta > 0$ then one would obtain as a corollary that the primes contain infinitely many three term arithmetic progressions. This already follows from the work of Vinogradov on the ternary Goldbach problem, but it should be true simply because the primes are

sufficiently dense, without using any deeper properties of their distribution. Furthermore, if one could show that

$$r(N) \ll \frac{1}{(\log N)(\log\log N)^2}, \tag{1.1}$$

say, then by partial summation it follows that if $\sum_{a\in A} 1/a$ diverges then A contains infinitely many three term arithmetic progressions. This would confirm a conjecture of Erdős, who also conjectured that the same condition should be sufficient for containing infinitely many k-term arithmetic progressions for any $k \geq 3$.

The first quantitative improvement on the result of Roth came from unpublished work of Szemerédi, who showed that

$$r(N) \ll \exp\left(-O\left(\sqrt{\log\log N}\right)\right).$$

It was subsequently observed by Heath-Brown and Szemerédi that the ideas that led to this bound could be coupled with the large sieve to yield further progress. Heath-Brown [33] showed that

$$r(N) \ll \frac{1}{(\log N)^c}$$

for some absolute constant $c > 0$, and a more direct approach by Szemerédi [73] showed explicitly that

$$r(N) \ll_\epsilon \frac{1}{(\log N)^{1/4-\epsilon}} \text{ for all } \epsilon > 0.$$

The next leap forward was achieved by Bourgain [6], who showed that

$$r(N) \ll \left(\frac{\log\log N}{\log N}\right)^{1/2},$$

which he later improved [9] to

$$r(N) \ll \frac{(\log\log N)^2}{(\log N)^{2/3}}.$$

By building upon the method of Bourgain [9], Sanders [61] obtained the further improvement

$$r(N) \ll_\epsilon \frac{1}{(\log N)^{3/4-\epsilon}} \text{ for all } \epsilon > 0.$$

Soon afterwards, by combining a new probabilistic method of Croot and Sisask [17] with combinatorial techniques Sanders [60] improved this to

$$r(N) \ll \frac{(\log\log N)^6}{\log N}.$$

This result comes, for the first time, within a whisker of the conjectured upper bound of (1.1) which would prove the conjecture of Erdős.

In this thesis we will prove the modest improvement

$$r(N) \ll \frac{(\log\log N)^4}{\log N},$$

which has appeared in [4], and where the record now stands. The significance of this bound is largely due to the fact that it is proved by quite a different method than that used by Sanders [60]; it avoids the probabilistic method of Croot and Sisask and instead operates almost entirely within Fourier space in the tradition of Roth. It is inspired by recent related work by Bateman and Katz [1] which we will discuss further below.

1.1.2 Generalisations

In this thesis we will widen our scope from the original problem of studying the density of subsets of $\{1, \dots, N\}$ without three term arithmetic progressions, to a much more general class of problems which we are able to tackle via the same method. We first observe that a three term arithmetic progression is precisely a solution to the equation

$$x_1 + x_2 - 2x_3 = 0,$$

with a trivial progression corresponding to a trivial solution where $x_1 = x_2 = x_3$. With this observation it is natural to speculate about similar results concerning subsets of $\{1, \dots, N\}$ with no such trivial solutions to any linear equation of the shape

$$c_1x_1 + \cdots + c_sx_s = 0 \tag{1.2}$$

where $c_i \in \mathbb{Z}$. The case $s \leq 2$ is trivial, and hence we shall restrict ourselves to the case $s \geq 3$ and $c_i \neq 0$ for $1 \leq i \leq s$. It is also extremely important that the solutions are translation invariant; that is, that if $(x_1, \dots, x_s)$ is a solution then so too is $(x_1 + x, \dots, x_s + x)$. If this condition fails then it is easy to construct sets of positive density having no solution to (1.2). Translation invariance is equivalent to the condition that

$$c_1 + \cdots + c_s = 0. \tag{1.3}$$

With such translation invariance we always have the trivial solutions $x_1 = \cdots = x_s$, but when $s \geq 4$ there can be other degeneracies within the coefficients that create more trivial

solutions. A study of such translation invariant equations was carried out by Ruzsa [55]. Following [55] we decompose $\{1,\dots,s\}$ as $I_1 \sqcup \cdots \sqcup I_\ell$, where for all $1 \le j \le \ell$

$$\sum_{i\in I_j} c_i = 0 \text{ and } \sum_{i\in I'} c_i \neq 0 \text{ if } I' \subsetneq I_j,$$

and define a solution $(x_1,\dots,x_s)$ to (1.2) to be trivial if for all $1 \le m \le \ell$ we have $x_i = x_j$ if $i,j \in I_m$. Ruzsa referred to the maximum such ℓ as the genus of the equation (1.2), and when $\ell \ge 2$ Ruzsa proved that we can obtain very strong bounds on the analogous density problem by simple combinatorial arguments. We will make no further mention of the genus, but we stress that, while our methods are valid for any translation invariant equation over the integers of the shape (1.2), the methods of Ruzsa deliver far superior bounds for any such equation with genus larger than 1.

Before going further, we make another immediate generalisation of the problem: the equation (1.2) makes sense with the variables taken from any abelian group G and the coefficients taken from R, the endomorphism ring of G (see below for some examples). We will attempt to tackle this generalisation of the original problem, although always preserving the crucial translation invariance condition (1.3).

In order to discuss the wide variety of such problems in a cohesive manner we will use the following definition. Let G be some abelian group and $B \subset G$ be a finite subset. Let R be the endomorphism ring of G and R^* the multiplicative subgroup of injective endomorphisms. Let $s \ge 3$ and $\mathbf{c} \in (R^*)^s$ satisfy (1.3). The generalised Roth problem for $(G,B,\mathbf{c})$ asks for the size of the largest $A \subset B$ that contains only trivial solutions to (1.2). We denote the size of this A by $R_{\mathbf{c}}(B;G)$, and let $r_{\mathbf{c}}(B;G) = R_{\mathbf{c}}(B;G)\,|B|^{-1}$. The coefficients $\mathbf{c}$ will be considered to be fixed, and the implied constants in the bounds which follow will depend on $\mathbf{c}$ (and, of course, also on G, although not on B).

This abstract definition is able to capture a wide variety of interesting problems. We will now give some examples of the kind of problem that we will consider.

1. When $G = \mathbb{Z}$ the endomorphism ring is $R = \mathbb{Z}$, and so $r_{\mathbf{c}}(\{1,\dots,N\};\mathbb{Z})$ is the density of the largest subset of $\{1,\dots,N\}$ without solutions to (1.2), where the coefficients $c_i \in \mathbb{Z}\backslash\{0\}$. For example, the density of the largest subset of $\{1,\dots,N\}$ without three term arithmetic progressions, is precisely $r_{(1,1,-2)}(\{1,\dots,N\};\mathbb{Z})$ (which we denoted by $r(N)$ in the previous section).

2. When $G = \mathbb{Z}^d$ the endomorphism ring is $R = M_d(\mathbb{Z})$, the ring of $d \times d$ matrices with integer entries. This setting contains many interesting higher-dimensional problems; for example, the density of the largest subset of $\{1, \dots, N\}^2$ without right-angled isosceles triangles is $r_{(c_1,c_2,c_3)}(\{1, \dots, N\}^2; \mathbb{Z}^2)$ where

$$c_1 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \; c_2 = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \text{ and } c_3 = \begin{pmatrix} -1 & -1 \\ 1 & -1 \end{pmatrix}.$$

3. When $G = \mathbb{F}_p^\infty$ the endomorphism ring R is rather large. It certainly contains $\mathbb{F}_p$; thus, for example, when $p \geq 3$ the quantity $r_{(1,1,-2)}(\mathbb{F}_p^n; \mathbb{F}_p^\infty)$ is concerned with sets without three term arithmetic progressions, as in the integer case.

4. There are other interesting endomorphisms of $\mathbb{F}_p^\infty$ to consider – for example, the 'shift' operator which is equivalent to multiplication by t if we view G as the additive group of $\mathbb{F}_p[t]$. For example, $r_{(1,t,-1-t)}(\mathbb{F}_p[t]_{\deg<n}; \mathbb{F}_p[t])$ is the density of the largest set of polynomials over $\mathbb{F}_p$ with degree less than n which contains no non-trivial solution to

$$x_1 + tx_2 = (1 + t)x_3.$$

5. Similarly, we can introduce the Frobenius morphism $\phi : x \mapsto x^p$ and consider sets of polynomials of degree less than n without solutions to

$$x_1 + tx_2 + x_2^p = (1 + t)x_3 + x_3^p;$$

we observe that this equation is indeed translation invariant, and the corresponding density is $r_{(1,t+\phi,-1-t-\phi)}(\mathbb{F}_p[t]_{\deg<n}; \mathbb{F}_p[t])$.

6. Our formulation also includes the problem of bounding the density of subsets of $\{1, \dots, N\}$ without k-term arithmetic progressions for $k > 3$. When $k = 4$, for example, this is equivalent to bounding $r_{\mathbf{c}}(\{1, \dots, N\}^2 \cap D; \mathbb{Z}^2)$ where

$$c_1 = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \; c_2 = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}, \; c_3 = \begin{pmatrix} -1 & 0 \\ 0 & -2 \end{pmatrix}, \text{ and } c_4 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

and D is the plane $\{(x_1, x_2) \in \mathbb{R}^2 : x_1 - x_2 = 0\}$. We observe, however, that c_1 is not injective on $\mathbb{Z}^2$; this is not, however, an insurmountable problem. Far more significant is the fact that the plane D is not closed under dilation by the

coefficients c_i, which is a significant obstacle for the approach given in this thesis. For longer arithmetic progressions one needs to use tools from the nascent field of higher order Fourier analysis, as pioneered by Gowers [26] in his groundbreaking work on Szemerédi's theorem.

7. Another intriguing problem, considered by Shkredov [67], is that of bounding the density of subsets of $\{1,\dots,N\}^2$ without 'corners': right-angled triangles whose sides are parallel to the axes. This is equivalent to bounding $r_{\mathbf{c}}(\{1,\dots,N\}^3 \cap D'; \mathbb{Z}^3)$ where

$$c_1 = \begin{pmatrix} -1 & 1 & 0 \\ -1 & -1 & 0 \\ -2 & 0 & 2 \end{pmatrix}, \; c_2 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 0 & -1 \end{pmatrix}, \text{ and } c_3 = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & -1 \end{pmatrix},$$

and D' is the plane $\{(x_1,x_2,x_3) \in \mathbb{R}^2 : x_2 - x_3 = 0\}$. In this case, all the coefficients are indeed injective, but once again the fact that the plane D' is not closed under dilation by the coefficients c_i prevents the methods outlined in this thesis from applying to this problem.

When $G = \mathbb{Z}$ this problem was first considered by Roth [51], who extended his earlier method to show that

$$r_{\mathbf{c}}(\{1,\dots,N\};\mathbb{Z}) \ll \frac{1}{\log\log N}.$$

This generalisation was not taken up in the subsequent literature concerning Roth's theorem on three term arithmetic progressions, although many of the methods in the literature are straightforward to generalise. What is less straightforward is obtaining a bound that improves as s increases. We showed in [3], generalising the bound on $r(N)$ obtained by Sanders [60], that

$$r_{\mathbf{c}}(\{1,\dots,N\};\mathbb{Z}) \ll \left(\frac{(\log\log N)^6}{\log N}\right)^{s-2}.$$

In this thesis we will generalise our improvement for the traditional case of Roth's theorem to obtain the bound

$$r_{s,\mathbf{c}}(\{1,\dots,N\};\mathbb{Z}) \ll \left(\frac{(\log\log N)^4}{\log N}\right)^{s-2}. \tag{1.4}$$

When $s \geq 6$ far greater savings can be made by alternative methods, as shown by Schoen and Shkredov [66], who showed that

$$r_{\mathbf{c}}(\{1,\dots,N\};\mathbb{Z}) \ll \exp\left(-O\left(\left(\frac{\log N}{\log\log N}\right)^{1/6}\right)\right) \text{ when } s \geq 6.$$

1.1.3 Vector spaces over a finite field

As mentioned above, it is now a well-known heuristic that the finite field setting $G = \mathbb{F}_p^\infty$ is a useful model of the case $G = \mathbb{Z}$, where the finite vector space $\mathbb{F}_p^n$ corresponds to the arithmetic progression $\{1,\dots,N\}$. In particular, this model is useful because arguments over $\mathbb{Z}/N\mathbb{Z}$ often become greatly simplified over $\mathbb{F}_p^n$.

This observation was first made in the context of Roth's theorem by Meshulam [45]. He observed that the density increment argument of Roth runs far more efficiently when the group has finite characteristic, and showed that

$$r_{\mathbf{c}}(\mathbb{F}_p^n;\mathbb{F}_p^\infty) \ll \frac{1}{n} \text{ when } \mathbf{c} \in \mathbb{F}_p^3.$$

Since $\left|\mathbb{F}_p^n\right| = p^n$ this is comparable to a bound of $r(N) \ll 1/\log N$ in the integer case, which is still out of reach. This was generalised by Liu and Spencer [39] to

$$r_{\mathbf{c}}(\mathbb{F}_p^n;\mathbb{F}_p^\infty) \ll \frac{1}{n^{s-2}} \text{ when } \mathbf{c} \in \mathbb{F}_p^s.$$

Again, Schoen and Shkredov [66] have obtained far superior bounds when $s \geq 6$, namely

$$r_{\mathbf{c}}(\mathbb{F}_p^n;\mathbb{F}_p^\infty) \ll \exp(-O(n^{1/5})) \text{ when } \mathbf{c} \in \mathbb{F}_p^s.$$

In a spectacular recent development, the 'log barrier' was finally broken for $\mathbb{F}_p^n$ by Bateman and Katz [1], who showed that

$$r_{\mathbf{c}}(\mathbb{F}_p^n;\mathbb{F}_p^\infty) \ll \frac{1}{n^{1+\delta}} \text{ for some absolute } \delta > 0 \text{ when } \mathbf{c} \in \mathbb{F}_p^3.$$

1.1.4 Polynomial rings over a finite field

Upon examination of these results a question immediately presents itself: since the endomorphism ring of $\mathbb{F}_p^\infty$ is much larger than $\mathbb{F}_p$ can we obtain similar bounds when the coefficients are endomorphisms more exotic than those generated from simple addition?

In fact, one can interpret $\mathbb{F}_p^\infty$ as the additive group of the polynomial ring $\mathbb{F}_q[t] = \mathbb{F}_q[t]$, where q is some power of p and $\mathbb{F}_q$ is the finite field of order q. Any endomorphism of the additive group of the polynomial ring $\mathbb{F}_q[t]$ thus corresponds to an endomorphism of $\mathbb{F}_p^\infty$, and we are led to the study of Roth's problem over $\mathbb{F}_q[t]$. Not only is this an interesting problem in its own right, but it is also a way to study translation invariant equations over $\mathbb{F}_p^\infty$ with the coefficients being quite exotic endomorphisms, and acts as a model for studying the problem over $\mathbb{Z}$ – one that is, in many respects, far more useful than the conventional $\mathbb{F}_p^\infty$ model.

When considering the coefficients of an equation of the shape (1.2) over $\mathbb{F}_q[t]$, there are obvious candidates – namely, coefficients drawn from $\mathbb{F}_q[t]$, as each $a \in \mathbb{F}_q[t]$ generates a natural endomorphism of $\mathbb{F}_q[t]$. Furthermore, just as the obvious analogue of an arithmetic progression $\{1, \dots, N\}$ in $\mathbb{F}_p^\infty$ is $\mathbb{F}_p^n$, the obvious set to consider in a Roth type problem is the $\mathbb{F}_q[t]$-arithmetic progression $\mathbb{F}_q[t]_n = \{x \in \mathbb{F}_q[t] : \deg x < n\}$. The study of how large a subset of $\mathbb{F}_q[t]_n$ can be while containing only trivial solutions to the equation (1.2), with the coefficients drawn from $\mathbb{F}_q[t]$, is the $\mathbb{F}_q[t]$-analogue of the traditional setting of Roth's problem over $\mathbb{Z}$.

Given that $\mathbb{F}_q[t]$ has much in common with $\mathbb{F}_p^\infty$ one might hope that Meshulam's bound can be generalised for this setting. This, sadly, is not achievable. Liu and Spencer [40] have, however, managed this in the case when all the coefficients have degree 0; that is, when all the coefficients are from $\mathbb{F}_q$,

$$r_{\mathbf{c}}(\mathbb{F}_q[t]_n; \mathbb{F}_q[t]) \ll \frac{1}{n^{s-2}} \text{ when } \mathbf{c} \in \mathbb{F}_q^s.$$

We observe that the left hand side is equal to $r_{\mathbf{c}}(\mathbb{F}_q^n; \mathbb{F}_q^\infty)$, and thus this result is a direct generalisation of Meshulam's bound for $\mathbb{F}_p^\infty$.

The general problem for arbitrary $\mathbf{c} \in \mathbb{F}_q[t]^s$, allowing the coefficients to have positive degree, is more difficult. With such coefficients the problem is closer to that over $\mathbb{Z}$ than that over $\mathbb{F}_p^\infty$. This problem was first addressed independently by Liu and Zhao [42] and the author [3]. Liu and Zhao [42] adapted the method of Bourgain [9] to deliver the bound

$$r_{s,\mathbf{c}}(\mathbb{F}_q[t]_n; \mathbb{F}_q[t]) \ll \left(\frac{(\log n)^2}{n}\right)^{\frac{2}{3}(s-2)\left(1-\frac{s-3}{4s-9}\right)} \text{ when } \mathbf{c} \in \mathbb{F}_q[t]^s.$$

In [3] we similarly adapted the method of Sanders [60] to prove the superior bound

$$r_{s,\mathbf{c}}(\mathbb{F}_q[t]_n; \mathbb{F}_q[t]) \ll \left(\frac{(\log n)^5}{n}\right)^{s-2} \text{ when } \mathbf{c} \in \mathbb{F}_q[t]^s.$$

In [4] we improved this to

$$r_{s,\mathbf{c}}(\mathbb{F}_q[t]_n;\mathbb{F}_q[t]) \ll \left(\frac{(\log n)^2}{n}\right)^{s-2} \text{ when } \mathbf{c}\in\mathbb{F}_q[t]^s. \tag{1.5}$$

In this thesis we will prove this bound; indeed, we will present our arguments in some generality so that we can deduce both the bound (1.4) for the integer problem and also the bound (1.5) for the $\mathbb{F}_q[t]$ setting.

We finally mention one more generalisation. While $\mathbb{F}_q[t]$ includes far more endomorphisms of $\mathbb{F}_p^\infty$ than $\mathbb{F}_p$ there are yet more; in particular, there are also those generated from the Frobenius endomorphism $x\mapsto x^p$ of $\mathbb{F}_q[t]$. The study of such endomorphisms leads naturally to Drinfeld modules. We defer the definitions and precise statement of results to Section 2.2.2. As a demonstration, however, we mention the following simple corollary of our results. There exists an absolute constant $\delta>0$ such that if $A\subset\mathbb{F}_q[t]_n$ has only trivial solutions to the translation invariant equation

$$x_1+tx_2=(1+t)x_3+(x_3-x_2)^p$$

then

$$|A|\ll_q \frac{q^n}{n^\delta}.$$

1.1.5 Lower bounds

We conclude our survey of Roth type problems with a brief discussion on the dual problem of proving lower bounds for the functions $r_{\mathbf{c}}$. This thesis makes no attempt to improve the lower bounds for $r(N)$ or related quantities, but it is interesting to observe how wide the divide still is, despite recent quantitative progress, between the best known upper and lower bounds.

At the time of the paper of Erdős and Turán [21] the best known construction for a set of integers with no three term arithmetic progressions was the set given by those integers whose ternary expansion only contains the digits 0 and 2, which leads to the lower bound

$$r(N)\gg N^{\log 2/\log 3-1}.$$

This was conjectured by some to be the best possible; or more generally that some bound of the shape $r(N)\ll N^{-\delta}$ should be achievable, for some absolute constant $\delta>0$. This

was shown to be false by Salem and Spencer [58], who showed that

$$r(N) \gg \exp\left(-O\left(\frac{\log N}{\log\log N}\right)\right).$$

This construction was soon after improved by Behrend [2], who gave a lower bound of the shape

$$r(N) \gg \exp\left(-O\left(\sqrt{\log N}\right)\right).$$

Remarkably, almost 70 years later this bound has not been substantially improved, which has led many to conjecture that this is the correct order of magnitude of $r(N)$. We, however, would not be surprised if $\exp(-O\left((\log N)^{1/3}\right))$ is closer to the truth; we will give some justification and other conjectures in Section 2.1.1. A slight improvement has been obtained by Elkin [19], and an alternative simpler proof was given by Green and Wolf [32]. More precisely, Behrend's argument leads to

$$r(N) \gg (\log N)^{-1/4} \exp\left(-\sqrt{8\log 2 \log N}\right),$$

while Elkin [19] obtains

$$r(N) \gg (\log N)^{1/4} \exp\left(-\sqrt{8\log 2 \log N}\right).$$

Thus, even an improvement of the constant in the exponential would be a significant achievement at this time. The only other setting where lower bounds have received attention is the finite field model $\mathbb{F}_p^n$. In this setting the constructions of [58] and [2] fail, and it may still be the case that $r(\mathbb{F}_p^n) \ll p^{-\delta n}$ for some absolute $\delta > 0$. Lower bound constructions in this setting rely on finding explicit constructions for small n and lifting them using the trivial inequality $r(\mathbb{F}_p^{mn}) \gg r(\mathbb{F}_p^n)^m$. Thus, for example, Edel [18] constructs a set $A \subset \mathbb{F}_3^{480}$ such that A contains no solutions to $x + y + z = 0$ and

$$|A| = 2^{327}\left(2^{73} + 3^7 7^{76}\right) \geq \left(3^{480}\right)^{0.72485},$$

which implies that for large n we have

$$r(\mathbb{F}_3^n) \gg 3^{-0.27515n},$$

which remains the best known lower bound for the problem over $\mathbb{F}_3^n$.

1.2 Inverse sumset theorems

Another fundamental problem of arithmetic combinatorics is the behaviour of the doubling constant $|A+A|/|A|$ where A is a finite set of integers. More precisely, we seek a full description of sets with a small doubling constant. It is not hard to find sets of integers with a small doubling constant. It is easy to check, for example, that if $A = \{1, \dots, N\}$ then $|A+A| = 2|A| - 1$; moreover, since the doubling constant is preserved under arbitrary translations and dilations the same equality holds for any arithmetic progression $A = a \cdot \{1, \dots, N\} + b$ (where $a, b \in \mathbb{Z}$).

More generally, we define a generalised arithmetic progression of dimension d to be a set of the shape

$$A = \{a_1 x_1 + \cdots + a_d x_d + b : 1 \leq x_i \leq N_i\}$$

for some integers $a_1, \dots, a_d \in \mathbb{Z}\backslash\{0\}$ and $N_1, \dots, N_d \geq 1$. It is easy to check that if A is a generalised progression of dimension d then $|A+A| \leq 2^d |A|$.

Thus we see that generalised progressions of small dimension provide many examples of sets which have a small doubling constant. It is possible to go yet further, however. We observe that if $B \subset A$ and $|B| \geq K^{-1}|A|$ then $|B+B| \leq |A+A| \leq K(|A+A|/|A|)|B|$. In particular, if A has a small doubling constant then any dense subset also has a small doubling constant. More generally, we say that B is d-covered by A if there exists a set X of size $|X| \leq d$ such that $B \subset A + X$. In such a case we have

$$|B+B| \leq |A+A+X+X| \leq |X|^2 |A+A| \leq \left(d^2 \frac{|A+A|}{|A|}\frac{|A|}{|B|}\right)|B|.$$

Combining our discussion thus far we obtain the following lemma.

Lemma 1.1. *Let A be a set which is K-covered by a generalised arithmetic progression P of dimension $O(\log K)$ such that $|P| \leq K|A|$. Then $|A+A| \leq K^{O(1)}|A|$.*

Our attempts to construct sets with a small doubling constant have been quite rudimentary so far, and it is natural to speculate whether one can find other, more exotic, examples. It is one of the most striking results of arithmetic combinatorics that this is not the case – in other words, one can prove that if A has doubling constant K then there exists some $d(K)$ such that A is $\exp(d(K))$-covered by a generalised progression P of dimension $d(K)$ such that $|P| \leq \exp(d(K))|A|$. We refer to such a theorem as an

inverse sumset theorem; qualitatively, such theorems offer a complete characterisation of sets with a small doubling constant.

The first general such inverse theorem was proved by Freiman [22], who showed that there exists some such function $d(K)$, although without any explicit bounds. Some time later Ruzsa [56] reawakened interest in such problems by providing a much simpler proof with explicit quantitative bounds; namely, $d(K) \ll K^{O(1)}$. He conjectured that it should be possible to take $d(K) \ll \log K$, a conjecture now commonly known as the Polynomial Freiman-Ruzsa conjecture. The example of a generalised arithmetic progression shows that this would be the best possible.

It is worth pointing out that Ruzsa's paper [56] does not, in fact, state an inverse theorem in the form above. There are several different types of inverse theorem, all of which say that if $|A+A|/|A|$ is small then A is 'close' in some sense to a small arithmetic progression. There are, of course, various measures of closeness we could use here. In this discussion we have chosen the one that we find most natural. Other forms of inverse theorem are essentially equivalent via simple combinatorial arguments.

Chang [13] later proved the bound

$$d(K) \ll K^2(\log K)^3,$$

and Schoen [64] improved this to the sub-polynomial bound

$$d(K) \ll \exp(O(\sqrt{\log K})).$$

Using a remarkable new probabilistic method of Croot and Sisask [17], Sanders [62] achieved, for the first time, poly-logarithmic bounds and showed that

$$d(K) \ll (\log K)^4 \log\log K.$$

In a subsequent paper Sanders [63] both simplified his proof and incorporated an argument of Konyagin to improve this to

$$d(K) \ll (\log K)^3(\log\log K)^{O(1)},$$

where the record now stands.

As with Roth's theorem we may ask the analogous question with $\mathbb{Z}$ replaced by $\mathbb{F}_p^\infty$. In other words, we seek to characterise sets $A \subset \mathbb{F}_p^n$ such that the doubling constant

$|A+A|/|A|$ is small. Here, of course, there are plentiful examples when the doubling constant is 1 – any subspace will do. As above, any set which is efficiently covered by a relatively small subspace will also have a small doubling constant. As with the integers, one can show that if $A \subset \mathbb{F}_p^n$ has $|A+A| \leq K|A|$ then there is a constant $f(K)$ and a finite subspace $P \leq \mathbb{F}_p^\infty$ such that A is $\exp(f(K))$-covered by P and $|P| \leq \exp(f(K))|A|$. This was first proved by Ruzsa [57], who also gives a conjecture of Marton that $f(K) \ll \log K$ should be achievable.

As above, the best known quantitative bounds are proved by Sanders [63], who established that

$$f(K) \ll (\log K)^3 (\log\log K)^{O(1)}.$$

In fact, Sanders [63] proves a more general result for arbitrary abelian groups of which the integer and $\mathbb{F}_p^n$ versions are immediate corollaries. The statement for general abelian groups, which was first addressed by Green and Ruzsa [30], is a little more technical to state, and we will not discuss it further; for more details see the paper of Sanders [63].

Instead, we return to our main theme and establish an inverse theorem for $\mathbb{F}_q[t]$, the ring of polynomials over the finite field $\mathbb{F}_q$. A sumset inverse theorem of the type considered above reduces to the case of $\mathbb{F}_p^\infty$, and so the results above hold. There is, however, an alternative question – we may instead seek to characterise the finite subsets of $\mathbb{F}_q[t]$ such that the t-doubling constant $|A+t\cdot A|/|A|$ is small, where $\mathbb{F}_q[t]$ is the ring of polynomials in t with coefficients in $\mathbb{F}_q$.

Aside from being an interesting problem in its own right, such a characterisation has applications to a variety of other problems in arithmetic combinatorics over $\mathbb{F}_q[t]$. Indeed, the traditional inverse theorems over $\mathbb{Z}$ have many important applications; we mention, in particular, the connection to the problem of obtaining quantitative bounds for the density of subsets of $\{1,\dots,N\}$ with no non-trivial four term arithmetic progressions.

An inverse theorem allows one to pass from a weak hypothesis to a strong structural conclusion, which is extremely useful for arithmetic applications. When considering the integers, which are generated by addition, it is sufficient to study the behaviour of the sumset $A+A$; if that can be controlled then this leads to a full understanding of the arithmetic nature of A. In $\mathbb{F}_q[t]$, however, addition is only half of the story – to capture the arithmetic structure of A one needs to understand both how A behaves under addition and how it interacts with the transcendental t.

As with $\mathbb{Z}$ it is easy to construct sets with a small t-doubling constant by using the notion of an $\mathbb{F}_q[t]$-generalised arithmetic progression, which is a set of the shape

$$\{a_1x_1 + \cdots + a_dx_d + b : \deg x_i < n_i\}$$

for some $a_1, \ldots, a_d, b \in \mathbb{F}_q[t]$ and integers $n_1, \ldots, n_d \geq 1$. By combining the arguments used for inverse theorems in other settings, especially those in the work of Sanders, with the rigid arithmetic structure of $\mathbb{F}_q[t]$ we are able to prove the following characterisation.

Theorem 1.2. *Let $A \subset \mathbb{F}_q[t]$ and $K \geq 4$ be such that $|A + t \cdot A| \leq K|A|$ and $|A + \alpha \cdot A| \leq K|A|$ for all $\alpha \in \mathbb{F}_q$. There exists a $\mathbb{F}_q[t]$-generalised arithmetic progression P of dimension $h(K)$ such that A is $\exp(h(K))$-covered by P and $|P| \leq \exp(h(K))|A|$ where*

$$h(K) \ll (\log K)^3(\log\log K)^5.$$

Furthermore, if we introduce the additional hypothesis that $|A + A| \ll |A|$ then we are able to prove the best possible quantitative result, which delivers the same quantitative bound as the Polynomial Freiman-Ruzsa conjecture (under the strong assumption that $|A + A| / |A| = O(1)$).

Theorem 1.3. *Let $A \subset \mathbb{F}_q[t]$ and $K \geq 4$ be such that $|A + t \cdot A| \leq K|A|$ and $|A + \alpha \cdot A| \leq K|A|$ for all $\alpha \in \mathbb{F}_q$. Furthermore, suppose that $|A + A| \ll |A|$. There exists a $\mathbb{F}_q[t]$-generalised arithmetic progression P of dimension $h(K)$ such that A is $\exp(h(K))$-covered by P and $|P| \leq \exp(h(K))|A|$ where*

$$h(K) \ll \log K,$$

where the implied constant depends on $|A + A| / |A|$.

In other words, from a quantitative perspective the difficult part of understanding the arithmetic structure of subsets of $\mathbb{F}_q[t]$ lies entirely with the sumset; we are able to prove quantitatively optimal results concerning any additional structure of $\mathbb{F}_q[t]$ over $\mathbb{F}_p^n$.

1.3 The sum-product phenomenon

The third problem of arithmetic combinatorics that we will consider in this thesis is a manifestation of the sum-product phenomenon. Roughly speaking, this is the idea that no

set can be both additively and multiplicatively structured at the same time; in particular, for any ring R and any reasonable finite set $A \subset R$ either the sumset $A+A$ or the product set AA must be almost as large as possible, barring trivial obstacles such as the presence of zero-divisors.

To be more precise in what follows, we will say that δ is permissible for a collection of finite sets $\mathcal{A}$ if for all $\epsilon > 0$ there exists a constant $C_\epsilon > 0$ such that for all $A \in \mathcal{A}$

$$\max(|A+A|, |AA|) \geq C_\epsilon |A|^{1+\delta-\epsilon}. \tag{1.6}$$

The sum-product heuristic then says that if $\mathcal{A}$ is some reasonable collection of finite subsets of any ring then 1 is permissible for $\mathcal{A}$; this is clearly the best possible. For example, such a conjecture is plausible for the collection of all finite subsets of an infinite field.

This problem was first considered by Erdős and Szemerédi for subsets of $\mathbb{R}$. In [20] they proved that there exists some absolute constant $c > 0$ which is permissible for the collection of finite subsets of $\mathbb{R}$, and conjectured that 1 is permissible. There has been steady progress on the problem over the real numbers, culminating with Solymosi [72] who showed that $1/3$ is permissible.

It is also reasonable to conjecture that 1 is permissible for finite subsets of $\mathbb{C}$, although this problem is more subtle. Solymosi [71] showed that $1/4$ is permissible, and the method of Solymosi [72] was generalised by Konyagin and Rudnev [36] to show that $1/3$ is permissible for all finite sets of complex numbers.

When one considers subsets of a finite field $\mathbb{F}_q$ there is an obvious obstacle to the sum-product phenomenon; namely, the fact that $\max(|A+A|, |AA|) \leq q$, which prevents (1.6) from holding if $|A|$ is too large as a function of q. If one only considers small subsets of $\mathbb{F}_q$, however, there are no such trivial obstacles, and hence 1 should be permissible here also.

This problem was first considered by Bourgain, Katz, and Tao [12], though only for q prime. They proved the following analogue of the Erdős-Szemerédi result: for all $\delta > 0$ there exists some $c(\delta) > 0$ which is permissible for all sets $A \subset \mathbb{F}_p$ such that $p^\delta < |A| < p^{1-\delta}$. It is, of course, crucial here that all implicit constants are absolute and do not depend on p.

When $|A| < p^{1/2}$ this estimate was made explicit by Garaev [24], who showed that $1/14$ is permissible; this was successively improved by Katz and Shen [34] and Bourgain

and Garaev [10]. The current best known result is due to Rudnev [53] who showed that 1/11 is permissible.

For general finite fields $\mathbb{F}_q$ the situation becomes more tangled, as there are more obstacles to a sum-product result; if A is a subfield of $\mathbb{F}_q$, for example, then clearly both $A+A$ and AA are as small as possible. Nonetheless, Li and Roche-Newton [38] have extended the method of Rudnev [53] to show that 1/11 is permissible for all subsets of $\mathbb{F}_q$ that are not 'too close' to a subfield.

Another interesting ring is the ring of polynomials over a given field; Croot and Hart [15] considered the case $\mathbb{C}[t]$, and showed that there exists some absolute constant $c > 0$ which is permissible for all finite sets of monic polynomials in $\mathbb{C}[t]$.

Finally, Tao [74] has considered the sum-product problem for arbitrary rings (not necessarily commutative), although the possibility of zero-divisors makes the theorems quite technical.

In this thesis we will prove, in joint work with Timothy G. F. Jones, new sum-product estimates for non-archimedean local fields, which have not thus far been considered in the sum-product literature. We recall that a non-archimedean local field is a locally compact topological field F equipped with a non-archimedean absolute value; that is, an absolute value $|\cdot| : F \to \mathbb{R}$ such that $|x+y| \leq \max(|x|, |y|)$ for all $x, y \in F$. A more concrete definition is that non-archimedean local fields are all finite extensions of $\mathbb{Q}_p$ for some prime p and all fields of Laurent series $\mathbb{F}_q(\!(t^{-1})\!)$ for some finite field $\mathbb{F}_q$.

We will show that for finite subsets of such fields 1/5 is permissible. Since the only archimedean local fields are $\mathbb{R}$ and $\mathbb{C}$ this result, combined with the above results of [72] and [36], implies the following.

Theorem 1.4. *Let F be any local field and let $\epsilon > 0$. For any finite $A \subset F$ we have*

$$\max(|A+A|, |AA|) \gg_{F,\epsilon} |A|^{6/5-\epsilon}.$$

Of course, since $\mathbb{F}_q[t] \subset \mathbb{F}_q(\!(t^{-1})\!)$ we obtain, in particular, the following sum-product result for $\mathbb{F}_q[t]$.

Theorem 1.5. *For any $\epsilon > 0$ and finite $A \subset \mathbb{F}_q[t]$ we have*

$$\max(|A+A|, |AA|) \gg_{q,\epsilon} |A|^{6/5-\epsilon}.$$

1.4 Technical background and definitions

We conclude by summarising some background material which we will use frequently.

1.4.1 Notation

For any group G and functions $f, g : G \to \mathbb{C}$ we write $f = O(g)$ or $f \ll g$ if there exists an absolute constant $C > 0$ such that $|f(x)| \leq C\,|g(x)|$ for all $x \in G$. Such a constant will be, in every instance, explicitly computable, but we often choose to suppress the constants to make the arguments more readable and transparent. At times the constant C may depend on other variables, which we indicate with the use of subscripts; thus, for example, $f \ll_s g$ indicates that the implied constant may depend on the parameter s.

When $\delta \in (0, 1]$ we write $\mathcal{L}(\delta) = 2 + \lceil \log(1/\delta) \rceil$. In particular, $\mathcal{L}(\delta) \geq 2$ is an integer and $e^{\mathcal{L}(\delta)} \geq \delta^{-1}$. Whenever we use logarithms we will always implicitly assume that the argument is sufficiently large to ensure positivity; when the term $\log\log K$ occurs, for example, we implicitly assume that $K > e$.

We write $x \approx y$ when $y \leq x < 2y$.

When B is a finite subset of an abelian group G we say that $A \subset B$ has density α, or that A has density α in B, if $\alpha = |A| \,/\, |B|$. It will sometimes be convenient to use the normalised counting measure on such a set B, which we will denote by β, so that $\beta(A) = |A \cap B| \,/\, |B|$.

1.4.2 Plünnecke-Ruzsa estimates

We will require various results from arithmetic combinatorics in our arguments, and in general we will state these explicitly as and when they are required. There are, however, sumset estimates that we will state now and use frequently in what follows. These are often collectively known as Plünnecke-Ruzsa estimates, and allow one to control higher order sumsets with a bound on a single sumset.

The first estimates of this type are due to Plünnecke [47]. Ruzsa [54] later both simplified the proof and proved a more general version (for a comprehensive discussion of the history and proofs we refer the reader to [25]). Furthermore, a wonderfully simple proof of Plünnecke's inequality was recently found by Petridis [46]. The following general form of such an estimate will suffice for our needs.

Theorem 1.6 (Plünnecke-Ruzsa estimate)**.** *Let $A, B_1, \dots, B_m$ be finite subsets of an abelian group G. If $|A + B_i| \leq K_i |A|$ for $1 \leq i \leq m$ then*

$$|B_1 + \cdots + B_m| \leq K_1 \cdots K_m |A| .$$

Thus, for example, if $|A + A| \leq K |A|$ then $|nA| \leq K^n |A|$ for all $n \geq 1$.

1.4.3 Fourier analysis on discrete abelian groups

Chapters 2 and 3 will make heavy use of Fourier analysis on an arbitrary discrete abelian group G, and we summarise the necessary background here. The standard reference is Rudin [52], where proofs of all the following facts can be found.

Let G be an abelian group with the discrete topology. A character is a homomorphism $\gamma : G \to \mathbb{C}$ such that $|\gamma(x)| = 1$ for all $x \in G$. The set of all characters on G forms an abelian group $\widehat{G}$. There is a natural topology on $\widehat{G}$, which is the coarsest topology such that the map $\gamma \mapsto \gamma(x)$ is continuous for every $x \in G$. This topology makes $\widehat{G}$ a compact group.

As locally compact abelian groups both G and $\widehat{G}$ have Haar measures – translation invariant regular measures – which are unique up to a multiplicative constant and defined on all Borel sets. We normalise these to be the counting measure on G and the probability measure on $\widehat{G}$. We denote the Haar measure on $\widehat{G}$ by μ.

For any $p \geq 1$ we define the L^p norm of a function $f : G \to \mathbb{C}$ as

$$\|f\|_p = \left(\sum_x |f(x)|^p \right)^{1/p}$$

and define the L^∞ norm to be $\sup_{x \in G} |f(x)|$. By $L^p(G)$ we denote the space of all functions $f : G \to \mathbb{C}$ such that $\|f\|_p < \infty$. Furthermore, if $X \subset G$ then by $L^p(X)$ we denote the space of all functions $f \in L^p(G)$ which are supported on X, i.e. $f(x) = 0$ if $x \notin X$.

For any $f \in L^1(G)$ we define the Fourier transform $\widehat{f} : \widehat{G} \to \mathbb{C}$ by

$$\widehat{f}(\gamma) = \sum_x f(x)\gamma(-x).$$

For any $f, g : G \to \mathbb{C}$ we have the inner product

$$\langle f, g \rangle = \sum_x f(x)\overline{g(x)}.$$

For $f, g \in L^1(G)$ this satisfies the Parseval formula

$$\langle f, g\rangle = \int \widehat{f}(\gamma)\overline{\widehat{g}(\gamma)}\,\mathrm{d}\gamma.$$

For $f, g : G \to \mathbb{C}$ we define the convolution as

$$f * g(x) = \sum_y f(y)g(x-y)$$

and have the identity $\widehat{f * g} = \widehat{f} \cdot \widehat{g}$. We write $f^{(s)}$ for the s-fold convolution of f with itself.

1.4.4 Polynomial rings

We denote the ring of polynomials with coefficients in some finite field $\mathbb{F}_q$ by $\mathbb{F}_q[t]$. When we use this notation the letter q will always denote the size of the coefficient field $\mathbb{F}_q$ and t the indeterminate.

There are many analogues between $\mathbb{F}_q[t]$ and $\mathbb{Z}$; we refer the reader to Rosen [49] for a thorough discussion. For our purposes, all that we will need is the fact that $\mathbb{F}_q[t]$ is a commutative ring of characteristic p where $q = p^r$. Furthermore, it is equipped with a valuation map $\deg : \mathbb{F}_q[t]\backslash\{0\} \to \mathbb{N}$ where if $x = \sum_{0\le n\le N} a_n t^n$ and $a_N \neq 0$ then $\deg x = N$. We further set $\deg 0 = -\infty$.

By $\mathbb{F}_q[t]_n$ we denote the finite subgroup of polynomials of degree strictly less than n. We observe that $\mathbb{F}_q[t]_n$ is a finite $\mathbb{F}_q$-vector space of size q^n. In arithmetic combinatorics it plays the same role for $\mathbb{F}_q[t]$ as $\{1, \dots, N\}$ does for $\mathbb{Z}$.

Finally, we observe that the dual group $\widehat{\mathbb{F}_q[t]}$ can be identified with the torus group

$$\left\{\sum_{n<0} a_n t^n : a_n \in \mathbb{F}_q\right\},$$

where $\gamma = \sum_{n<0} a_n t^n$ corresponds to the character on $\mathbb{F}_q[t]$ defined by

$$f_\gamma(x) = \exp(2\pi i \mathrm{Tr}(\mathrm{res}_{-1}(\gamma x))/p)$$

and res_{-1} is the residue map which picks out the coefficient of t^{-1} and $\mathrm{Tr} : \mathbb{F}_q \to \mathbb{F}_p$ is the trace map.

1.4.5 Other definitions

There are quite a few non-standard definitions that we will employ. These will be defined before they are used, but we also collect the most important here for ease of reference.

When $X, X' \subset G$ are finite sets of an abelian group G we say that X' is δ-*sheltered* by X (or that X δ-shelters X') if $|(X'+X)\backslash X| \leq \delta |X|$. If $\Gamma \subset \widehat{G}$ and $X \subset G$ we say that X has δ-*control* of Γ if, for all $\gamma \in \Gamma$ and $x \in X$, we have $|1-\gamma(x)| \leq \delta$. For any $f \in L^1(G)$ and $\eta \in [0,1]$ we define the η-spectrum to be

$$\Delta_\eta(f) = \left\{\gamma \in \widehat{G} : \left|\widehat{f}(\gamma)\right| \geq \eta \|f\|_1\right\}.$$

Furthermore, we define the η-level spectrum to be

$$\tilde{\Delta}_\eta(f) = \left\{\gamma \in \widehat{G} : \left|\widehat{f}(\gamma)\right| \approx \eta \|f\|_1\right\}.$$

We finally flag up some notation which will be used in different ways at different points in this thesis. Firstly, the use of $\langle X \rangle$ where X is some finite subset of a group G. This notation is used in three quite distinct (though related) ways in this thesis. In Chapter 2 it is used to denote the smallest closed subgroup which contains X. In Chapter 3 it is used to denote the set

$$\left\{\sum_{x \in X} \epsilon_x x : \epsilon_x \in \{-1,0,1\}\right\}, \tag{1.7}$$

while in Chapter 4, when $G = \mathbb{F}_q[t]$, it is used to denote the $\mathbb{F}_q$-subspace generated by X. Which definition is in use will be clear from the context.

Secondly, the notion of covering. In Chapter 4 it will be used in the traditional sense; thus we say that X is r-covered by Y if there exists a set Z of size at most r such that $X \subset Y + Z$. In Chapter 3, however, we say that X is r-covered by Y if there exists a set Z of size at most r such that

$$X \subset Y - Y + \langle Z \rangle,$$

with $\langle Z \rangle$ used in the sense of (1.7).

CHAPTER 2

TRANSLATION INVARIANT EQUATIONS

The objects of study in this and the following chapter are translation invariant equations; more precisely, we will study how dense a subset of a finite structured subset of an abelian group can be while containing only trivial solutions to a given translation invariant equation.

We first fix some notation that will hold throughout the following two chapters. We fix some abelian group G equipped with the discrete topology, and let $\widehat{G}$ denote its compact dual group. As usual, we equip G with the counting measure and $\widehat{G}$ with the Haar probability measure, which we denote by μ. We let R be the endomorphism ring of G and let R^* be the multiplicative subgroup of injective endomorphisms. Finally, if $\mathbf{c} \in R^s$ then we say that $\mathbf{c}$ is commutative (or that $\mathbf{c}$ commutes) if $c_i c_j = c_j c_i$ for all $1 \leq i \leq j \leq s$.

Let $s \geq 3$ and $\mathbf{c} = (c_1, \dots, c_s) \in (R^*)^s$ be a commutative s-tuple such that $c_1 + \cdots + c_s = 0$. For any finite set A we will study the solutions to

$$c_1 x_1 + \cdots + c_s x_s = 0 \tag{2.1}$$

with $x_i \in A$. We say that such a solution is trivial if all the x_i are identical. Given some fixed finite set $X \subset G$ possessing a certain amount of structure we will study how large a set $A \subset X$ can be if it contains no non-trivial solutions to (2.1). For example, if we take $G = \mathbb{Z}$ with $X = \{1, \dots, N\}$ and $\mathbf{c} = (1, 1, -2)$ then this problem becomes that of bounding the density of sets of integers which contain no non-trivial three term arithmetic progressions.

In this chapter we undertake a general study of such problems and give new quantitative bounds that offer an improvement even in the traditional case $G = \mathbb{Z}$. Our main tool is Theorem 2.3 which encapsulates an efficient density increment method and the statement of which is quite general. We postpone the proof of Theorem 2.3 to the following chapter. In this chapter we will use it as a black box and examine in detail how to apply it to various cases of interest.

Our results take a particularly strong form when X is a finite subgroup. In particular,

we will prove the following theorem, which is a quantitative version of the analogue of Roth's theorem for $\mathbb{F}_q[t]$, the ring of polynomials with coefficients in $\mathbb{F}_q$.

Theorem 2.1. *Let $A \subset \mathbb{F}_q[t]_N$ with density α and $\mathbf{c} \in (\mathbb{F}_q[t]\backslash\{0\})^s$ be such that $c_1 + \cdots + c_s = 0$. If A has only trivial solutions to (2.1) then*

$$\alpha \ll_{\mathbf{c},s,q} \left(\frac{(\log N)^2}{N}\right)^{s-2}.$$

In fact, we will prove this theorem while keeping the dependence on $\mathbf{c}$ and q explicit. This upper bound is an improvement of our earlier upper bound $((\log N)^5)/N)^{s-2}$ which appeared in [3].

The application to cases where X is not a subgroup is more technically involved. We will prove the following theorem, which is currently the best known quantitative bound for Roth's theorem on arithmetic progressions.

Theorem 2.2. *Let $A \subset \{1, \ldots, N\}$ with density α and $\mathbf{c} \in (\mathbb{Z}\backslash\{0\})^s$ be such that $c_1 + \cdots + c_s = 0$. If A has only trivial solutions to (2.1) then*

$$\alpha \ll_{\mathbf{c},s} \left(\frac{(\log\log N)^4}{\log N}\right)^{s-2}.$$

Before delving into the technical framework it is worth emphasising that our methods have certain limitations in which equations of the shape (2.1) they can handle. The coefficient tuple $\mathbf{c}$ must satisfy three conditions:

1. translation invariance, i.e. $c_1 + \cdots + c_s = 0$,
2. commutativity, i.e. $c_i c_j = c_j c_i$ for $1 \leq i \leq j \leq s$, and
3. injectivity, i.e. $c_i \in R^*$ for $1 \leq i \leq s$.

The first condition, translation invariance, is to be expected, and is forced upon us by the iterative nature of the argument. Indeed, even in the case $G = \mathbb{Z}$ it is easy to see that without translation invariance one can have sets of constant density within $\{1, \ldots, N\}$ for arbitrarily large N which do not contain any solutions to (2.1). Consider, for example, the equation $x_1 + x_2 - x_3 = 0$ and the set of odd numbers.

The second condition, commutativity, we hope can be weakened somewhat. It is not an issue when the endomorphism ring is commutative, as with $G = \mathbb{Z}$, but it becomes a significant barrier for higher-dimensional settings such as $G = \mathbb{Z}^d$.

The final condition, injectivity, we also hope could be weakend; indeed, we believe that it is possible to adapt our methods to give results of similar strength while allowing for one of the coefficients to be non-injective, although we have not been able to achieve this in a rigorous fashion.

Finally, we remark that our methods are limited not only by the properties of the equation (2.1), but also by the form of the initial set X. Thus, as discussed in the introduction, we say nothing new about the density of subsets of $\{1, \dots, N\}$ without k-term progressions for $k > 3$ or of subsets of $\{1, \dots, N\}^2$ without 'corners', right-angled triangles whose sides are parallel to the axes. These problems can be formulated in terms of a single translation invariant equation of the shape (2.1), but the correspoding initial sets X are not amenable to our methods – in short because they are not even approximately closed under dilation by the coefficients c_i.

2.1 A general framework

Let A be any finite subset of G and $\mathbf{c} \in R^s$ for some $s \geq 3$. We wish to count solutions to (2.1), and so we define

$$\Upsilon_{\mathbf{c}}(A) = \langle (c_1 \cdot A) * \cdots * (c_{s-1} \cdot A), ((-c_s) \cdot A) \rangle.$$

We observe that $\Upsilon_{\mathbf{c}}(A)$ is a count of all the solutions to (2.1) with the variables lying in the set A, including the trivial solutions where all of the variables are identical. The number of such solutions can be trivially bounded by $|A|$, however, and hence if A contains only trivial solutions to (2.1) then we must have $\Upsilon_{\mathbf{c}}(A) \ll |A|$. In particular, a suitable lower bound on $\Upsilon_{\mathbf{c}}(A)$ for arbitrary sets A (which will depend on $|A|$ itself) will lead to an upper bound on the size of sets which have only trivial solutions to (2.1).

In broad strokes, our approach follows the density increment strategy first exploited by Roth [50]. The strategy is as follows. Suppose that $A \subset H$, where H is some finite subset of G which has a fair degree of structure. If $\Upsilon_{\mathbf{c}}(A)$ is small, and in particular smaller than the expected count $|A|^s / |H|$ which would hold if A were randomly distributed throughout H, then A does not behave randomly and in particular is not equally distributed over

large subsets of H. In particular, one can find a large subset $H' \subset H$, which is also very structured, and $x \in G$ such that

$$\frac{|(A-x) \cap H'|}{|H'|} \geq (1+\nu)\frac{|A \cap H|}{|H|},$$

for some $\nu > 0$. The key observation is that the number of solutions to (2.1) with the variables chosen from A is at least the number of solutions with the variables chosen from $(A-x) \cap H'$; this is where the fact that (2.1) is translation invariant is crucial.

We now repeat this argument; hence either $\Upsilon_{\mathbf{c}}((A-x) \cap H')$ is large or there is some structured $H'' \subset H'$ on which a translate of A has increased density, and so on. Since the density is always bounded above by 1, however, this iteration must halt after some bounded number of steps. It halts with some $A^\sharp \subset A - x$ and a structured subset $H^\sharp$ such that $\Upsilon_{\mathbf{c}}(A^\sharp)$ is at least the expected count $\left|A^\sharp\right|^s / \left|H^\sharp\right|$. Since $A^\sharp$ contains only trivial solutions to (2.1), if $\alpha = |A| / |H|$,

$$\left|A^\sharp\right| \geq \Upsilon_{\mathbf{c}}(A^\sharp) \gg \frac{\left|A^\sharp\right|^s}{|H^\sharp|} \geq \left|H^\sharp\right|^{s-2} \alpha^{s-1} \left|A^\sharp\right|.$$

It follows that $\alpha \leq \left|H^\sharp\right|^{-(s-2)/(s-1)}$, which will give a non-trivial upper bound for the density of A within H as required, provided $\left|H^\sharp\right|$ is not too small.

The strength of the result obtained will thus depend on two parameters – the size of the density increment ν, which controls how many steps of the iteration are required, and the relative size of the next structured set, $|H'| / |H|$, which is a factor lost in each step of the iteration. These two factors will give a lower bound for the size of the final structured set $H^\sharp$, and hence an upper bound for α.

Of course, to turn this idea into a rigorous argument we need to make precise what kind of structured subset is required to run the argument. In particular, it needs to be structured enough to be able to carry out the argument, but not so structured that we cannot find a relatively large similarly structured subset for the next stage. The subtlety therefore lies in carrying out the argument with the bare minimum of structural requirements.

The structural requirements can be roughly summarised by the necessity of being able to study solutions to (2.1) locally within H, which implies that H must be mostly closed under addition and dilations by the coefficients $\mathbf{c}$. To insist that H be itself closed under

such operations is far too restrictive for our applications; it turns out that it suffices for H to have enough subsets which do not 'escape' too far from it under these operations. To show more precisely what we mean, we recall the following crucial definition, which will be used extensively in the definitions to follow: we say that X' is δ-*sheltered* by X if

$$|(X + X')\backslash X| \leq \delta |X|.$$

Given that H is structured enough for us to be able to locally study solutions to (2.1), we also need to know what kind of subset $H' \subset H$ we must pass to to be able to obtain the requisite density increment. We shall see in Chapter 3 that the important property is that H' is suitably controlled on some finite subset of $\widehat{G}$ – we recall that X has δ-*control* of Γ if $|1 - \gamma(x)| \leq \delta$ for all $x \in X$ and $\gamma \in \Gamma$.

We now state precisely the kind of structures that our method requires to produce a density increment, and that are produced at the next stage. These definitions are to a certain degree quite artificial, as they package together the various technical hypotheses that will be required in the proof of our main theorem. The core idea is quite natural, however – given a set X we will pass to a subset X' which is both sheltered by X and is also well controlled on some bounded number of characters.

Let $X \subset G$ be a finite set, $\delta \in [0, 1]$ and $\mathbf{c} \in R^s$. Furthermore, let $(\tilde{B}, \tilde{B}')$ be a pair of finite symmetric subsets of G, each containing 0. We call such a pair *good* for $(X, \mathbf{c}; \delta)$ if

1. the sets $c_1 \cdot \tilde{B}$, $c_s \cdot \tilde{B}$ and $c_1 c_s \cdot \tilde{B}'$ are all δ-sheltered by X, and

2. $c_2 \cdot \tilde{B}' + \cdots + c_{s-1} \cdot \tilde{B}'$ is δ-sheltered by $\tilde{B}$.

Let $\mathfrak{B}$ be a collection of finite symmetric subsets of G, each containing 0 and $C \geq 1$. We say that $\mathfrak{B}$ is *permissible* for $(X, \mathbf{c}; \delta, d)$ if there exists a good pair for $(X, \mathbf{c}; \delta)$, say $(\tilde{B}, \tilde{B}')$, such that

1. $c_1 \cdot \tilde{B}$, $c_s \cdot \tilde{B}$, and $c_1 c_s \cdot \tilde{B}'$ are all members of $\mathfrak{B}$, and

2. for each finite $\Gamma \subset \widehat{G}$ of size at most d and each $1 < i < s$ the collection $\mathfrak{B}$ contains $c_1 c_s c_i \cdot X'$ for some $X' \subset \tilde{B}'$ which has $(4|\Gamma|)^{-1}$-control of Γ and is δ-sheltered by $\tilde{B}'$.

A *chain* from X of length K, denoted by $\mathfrak{X} = (\mathfrak{X}_0, \dots, \mathfrak{X}_K)$, is a sequence of collections of finite symmetric subsets of G, where $\mathfrak{X}_0 = \{X\}$ and $\mathfrak{X}_{i+1} \subset \cup_{X \in \mathfrak{X}_i} \mathcal{P}(X)$. By $X \in \mathfrak{X}$ we mean that there exists some $0 \leq i \leq K$ such that $X \in \mathfrak{X}_i$.

Finally, let $\delta \in [0,1]$ and $d : \mathbb{N} \to \mathbb{R}_+$. We say that a chain $\mathfrak{X}_0, \dots, \mathfrak{X}_K$ satisfies $DI(\delta, d; \mathbf{c})$ if, for $i \geq 1$ and every $X' \in \mathfrak{X}_{i-1}$, the collection $\mathfrak{X}_{i+1}$ contains some collection permissible for $(X', \mathbf{c}; \delta, d(i))$. This property, with an appropriate choice for δ and d, encapsulates precisely the necessary hypotheses to allow the previously outlined density increment argument to succeed; thus, at the ith stage of the argument, we will be working within some $X' \in \mathfrak{X}_i$, and we will show that A has increased density on some $X'' \in \mathfrak{X}_{i+1}$.

Another abstraction of the kind of set system required to run a density increment argument is the notion of a 'Bourgain system' outlined by Green and Sanders [31]. This concept, however, lies in a somewhat orthogonal direction; roughly speaking, a Bourgain system is a source for sets with which to create the next link of a chain that satisfies $DI(\delta; r; \mathbf{c})$ for some suitable parameters δ and r.

We are finally ready to state our main theorem concerning translation invariant equations, which will be proved in Chapter 3. We recall that a trivial solution to (2.1) is one where $x_1 = \dots = x_s$.

Theorem 2.3. *There exists a constant $C(s) > 1$, depending only on s, such that the following holds. Let $\mathbf{c} \in (R^*)^s$ be a commutative s-tuple such that $c_1 + \dots + c_s = 0$. Let $X \subset G$ be a finite symmetric set and $A \subset X$ with density α. Let*

$$\delta = \exp(-C\mathcal{L}(\alpha)^2) \text{ and } r(i) = C(1 + C^{-1})^{-i/(s-2)} \alpha^{-1/(s-2)} \mathcal{L}(\alpha) \text{ for } i \geq 0.$$

If A contains only trivial solutions to (2.1) then there exists an integer $0 \leq K \ll_s \mathcal{L}(\alpha)$ such that if $\mathfrak{X}$ is a chain from X of length K which satisfies $DI(\delta, r; \mathbf{c})$ then there exists $X' \in \mathfrak{X}$ such that

$$|X'|^{-1} \gg \alpha^2. \tag{2.2}$$

The appearance of α^2 may give the impression that bounds of the strength $\alpha \ll |X|^{-1/2}$ can follow from an application of Theorem 2.3. This is misleading, however – passing from one link in the chain to the next will always incur a substantial loss in the size of $|X'|$, which will depend on α, and hence it is the size of $|X'| / |X|$ which will dominate in (2.2).

We remind the reader that Theorem 2.3 is valid for any abelian group G with endomorphism ring R. For it to have a non-trivial conclusion, however, we need to be able to control the cardinality of the sets making up each link in a chain – this is a tricky matter, and places strict demands on the structure of the original set X. In the rest of this chapter we discuss this matter in some detail. As one might expect, it is far simpler when X is a finite subgroup of G, largely due to the fact that subgroups completely shelter all of their subsets.

2.1.1 Conjectures and comparison with previous results

Theorem 2.3 may be viewed as a packaging of the density increment argument, as carried out in full generality, without specialising it to any particular special cases such as the integers. Many of the previous methods used for Roth-type theorems can be reinterpreted in this fashion. There are three important quantitative parameters: K, δ, and d.

The parameter K essentially measures how many times we can run the density increment argument before we must halt due to the trivial bound $\alpha \leq 1$; thus if at each stage we can pass from a set of density α to a set of density $(1+\Omega(1))\alpha$ then $K \approx \mathcal{L}(\alpha)$. On the other hand, if we can only achieve a density increment of $\alpha \mapsto (1+\Omega(\alpha))\alpha$ then $K \approx \alpha^{-1}$.

The parameter δ measures how much shelter is required at each stage; essentially, how additively closed our approximately structured sets must be in order to have enough structure for the kind of relative 'local' Fourier analysis we will perform at each step.

Finally, the function d is how many characters we need to control at each step. That is, to find set $X' \subset X$ on which A has increased density, we take some subset of X which is sufficiently closed under addition and then passing to a subset on which roughly d characters are trivial.

We first consider the consequences of Theorem 2.3 for the classical problem of bounding the density of subsets of $\{1,\dots,N\}$ without three term arithmetic progressions. By constructing suitable chains explicitly for this setting, as we will do in Section 2.2, one can show that if Theorem 2.3 holds with parameters $K(\alpha)$, $\delta(\alpha)$ and $d(i) = d(i;\alpha)$ then, whenever $A \subset \{1,\dots,N\}$, with $|A| = \alpha N$, has no three term arithmetic progressions we have

$$\mathcal{L}(\delta)\sum_{i=0}^{K}(K-i)d(i) \gg \log N.$$

For example, the bounds

$$K \ll \mathcal{L}(\alpha),\ \delta \gg \exp(-O(\mathcal{L}(\alpha)^2) \text{ and } d(i) \ll c^i \alpha^{-1}\mathcal{L}(\alpha) \text{ for some } c < 1$$

of Theorem 2.3 lead to the bound

$$\alpha \ll \frac{(\log\log N)^4}{\log N}.$$

We summarise below how some previous work, when reinterpreted in this manner, compares to the bounds in Theorem 2.3 (we omit mention of absolute multiplicative constants in the table below; the constant $0 < c < 1$ denotes some fixed quantity).

A notable omission from the table below are the results of Bourgain [9] and Sanders [61], which give the bounds $\alpha \ll (\log N)^{-2/3+o(1)}$ and $\alpha \ll (\log N)^{-3/4+o(1)}$ respectively. While these methods fall very much within the density increment framework, they use a delicate case analysis which is not easy to interpret in the framework of Theorem 2.3. We further remark that we credit both Roth [50] and Bourgain [6] for the same approach since the core mechanism of achieving that kind of density increment was already present in Roth [50]; the achievement of Bourgain was to refine the technical machinery to most efficiently exploit that density increment, refinements which we will make full use of in the next chapter. Similarly, the first row credits Heath-Brown [33], Szemerédi [73] and Chang [13] since the described density increment can be obtained by combining the ideas present in those papers, although this approach and the stated bound on α have not been published.

Similarly, one may construct explicit chains for $\mathbb{F}_p^n$ in such a manner as to show that if Theorem 2.3 holds with parameters $K(\alpha)$, $\delta(\alpha)$ and $d(i) = d(i;\alpha)$ then, whenever $A \subset \mathbb{F}_p^n$, with $|A| = \alpha\mathbb{F}_p^n$, has no three term arithmetic progressions we have

$$\sum_{i=0}^{K} d(i) \gg_p n.$$

We observe in particular that there is no longer any dependence on δ (essentially because δ was the parameter which controlled how sheltered subsets were, and a subspace in $\mathbb{F}_p^n$ is closed under addition and hence automatically shelters all its subsets completely). For a similar reason, a factor of K has also been lost.

As above, we can interpret some existing results in this fashion.

Method of	K	δ	$d(i)$	$\alpha \ll$
Heath-Brown [33] Szemerédi [73] Chang [13]	$\mathcal{L}(\alpha)$	$\alpha^{O(1)}$	$c^i\alpha^{-2}\mathcal{L}(\alpha)$	$\frac{(\log\log N)^{3/2}}{(\log N)^{1/2}}$
Roth [50] Bourgain [9]	α^{-1}	$\alpha^{O(1)}$	1	$\frac{(\log\log N)^{1/2}}{(\log N)^{1/2}}$
Sanders [60]	$\mathcal{L}(\alpha)$	$\alpha^{O(1)}$	$c^i\alpha^{-1}\mathcal{L}(\alpha)^4$	$\frac{(\log\log N)^6}{\log N}$
Theorem 2.3	$\mathcal{L}(\alpha)$	$\exp(-O(\mathcal{L}(\alpha)^2))$	$c^i\alpha^{-1}\mathcal{L}(\alpha)$	$\frac{(\log\log N)^4}{\log N}$

Table 2.1: Density increment results over $\{1, \dots, N\}$

	K	$d(i)$	$\alpha \ll$
Meshulam [45]	α^{-1}	1	$1/n$
Sanders [60]	$\mathcal{L}(\alpha)$	$c^i\alpha^{-1}\mathcal{L}(\alpha)^4$	$(\log n)^4/n$
Bateman and Katz [1]	α^{-c}	1	$1/n^{1/c}$
Theorem 2.3	$\mathcal{L}(\alpha)$	$c^i\alpha^{-1}\mathcal{L}(\alpha)$	$\log n/n$

Table 2.2: Density increment results over $\mathbb{F}_p^n$

We recall that c denotes some fixed quantity $0 < c < 1$, and hence the result of Bateman and Katz is the strongest.

It is an interesting question to ask how strong a density increment theorem such as Theorem 2.3 could be; any improvement in the parameters K, d, or δ would immediately yield quantitative progress in problems such as Roth's theorem. In Table 2.3 we make three possible conjectures as to where the truth lies, and the consequences for three

different Roth-type problems.

	'One large character; α-shelter' conjecture	'One large character; constant shelter' conjecture	'Many large characters; α-shelter' conjecture
K	$\mathcal{L}(\alpha)$	$\mathcal{L}(\alpha)$	1
δ	$\alpha^{O(1)}$	1	$\alpha^{O(1)}$
$d(i)$	1	1	$\mathcal{L}(\alpha)$
$A \subset \{1,\dots,N\}$ $x+y=2z$ solution free $\alpha \ll$	$\exp(-O((\log N)^{1/3}))$	$\exp(-O((\log N)^{1/2}))$	$\exp(-O((\log N)^{1/2}))$
$A \subset \mathbb{F}_q[t]_{\deg<n}$ $x+ty=(1+t)z$ solution free $\alpha \ll$	$\exp(-O(n^{1/2}))$	$\exp(-O(n^{1/2}))$	q^{-cn}
$A \subset \mathbb{F}_p^n$ $x+y=2z$ solution free $\alpha \ll$	p^{-cn}	p^{-cn}	p^{-cn}

Table 2.3: Possible conjectures for the best-possible density increment approach

The first is the situation where, if A had solutions to the given equation, then there is a character $\gamma \in \widehat{G}$ such that $\left|\widehat{A}(\gamma)\right| \gg \alpha$, and that the surrounding machinery would need a degree of additive shelter comparable to α to be successfully carried out.

The second is the situation where, as before, there exists some large character, but also that only some constant degree of shelter would be required. We believe that this is too ambitious, however; roughly speaking, having $\delta \approx \alpha$ ensures that, when performing

the density increment argument relative to some $A \subset X$, we are passing to some

$$|(X' + X) \backslash X| \leq |A| / 100,$$

say, and hence when considering elements of $((X' \cap A) + A) \cap (-2 \cdot A)$ (say, which is roughly what is considered when counting three term arithmetic progressions) only a small fraction lies outside the set X, where our methods cannot reach it (since all our arguments are performed relative to X). Thus imposing a shelter parameter of $\delta \approx \alpha$ seems quite natural.

If one could only have a constant shelter parameter, then for reasonably sparse sets A it is possible that all of $((X' \cap A) + A) \cap (-2 \cdot A)$ lies outside X which would destroy any hope of an argument using techniques relative to X succeeding.

Finally, the third possibility is that, as before, some degree of shelter comparable to α is necessary, but that there exists not just one large character, but almost as many as possible, roughly α^{-3}.

We believe that the first conjecture is the most plausible, and hence the correct bound for Roth's theorem over the integers is of the shape $\exp(-(\log N)^{1/3})$, rather than the Behrend-type bound $\exp(-(\log N)^{1/2})$. The model setting of $\mathbb{F}_q[t]$ suggests a way to determine where the truth lies, by examining in detail whether it is possible to achieve a lower bound of Behrend-type strength for this problem.

2.2 Strong structure

In this section we address the situation when X is a finite subgroup of G. Although our arguments are valid for any abelian group G they are only of interest when G has many such finite subgroups. When $G = \mathbb{Z}$, for example, the only finite subgroup is $\{0\}$, and we understand subsets of this group well enough without recourse to the arguments in this section.

If $X \subset G$ then we define

$$[X] = \left\{\gamma \in \widehat{G} : \gamma(x) = 1 \text{ for all } x \in X\right\},$$

and if $\Gamma \subset \widehat{G}$ then

$$[\![\Gamma]\!] = \{x \in G : \gamma(x) = 1 \text{ for all } \gamma \in \Gamma\}.$$

We observe that by continuity of the characters $[X]$ is a closed (and, in particular, measurable) subset of $\widehat{G}$. We first record some basic properties of the operators $[\cdot]$ and $[\![\cdot]\!]$. In this chapter $\langle X\rangle$ denotes the subgroup of G generated by X, and similarly $\langle\Gamma\rangle$ denotes the closed subgroup of $\widehat{G}$ generated by $\Gamma\subset\widehat{G}$.

These constructions are known as annihilators, and the following basic facts can be generalised to arbitrary locally compact abelian groups; we refer the reader to Chapter 2 of Rudin [52] for more details.

Lemma 2.4. *Let $\Gamma\subset\widehat{G}$ and $X\subset G$. We have the following properties:*

1. *if $\Gamma'\subset\Gamma$ then $[\![\Gamma']\!]\supset[\![\Gamma]\!]$,*
2. *if $X'\subset X$ then $[X']\supset[X]$,*
3. *$[\![[X]]\!]\supset X$ and $[[\![\Gamma]\!]]\supset\Gamma$,*
4. *$[\![\Gamma]\!]=[\![\langle\Gamma\rangle]\!]$ and $[X]=[\langle X\rangle]$,*
5. *if $X\leq G$ is a finite subgroup then $[X]$ is a subgroup of $\widehat{G}$ and $\mu([X])=|X|^{-1}$,*
6. *if $X\leq G$ is a finite subgroup then $X=[\![[X]]\!]$,*
7. *if $\Gamma\subset\widehat{G}$ has positive measure then $[\![\Gamma]\!]$ is a finite subgroup of G of size $\mu(\langle\Gamma\rangle)^{-1}$.*

Proof. The first three properties follow immediately from the definitions. By property 3 it follows that $[[\![\Gamma]\!]]$ is a closed subgroup of $\widehat{G}$ which contains Γ and hence $[[\![\Gamma]\!]]\supset\langle\Gamma\rangle\supset\Gamma$. By the first three properties it follows that

$$[\![\Gamma]\!]\subset[\![[[\![\Gamma]\!]]\!]]\subset[\![\langle\Gamma\rangle]\!]\subset[\![\Gamma]\!]\,,$$

so that $[\![\Gamma]\!]=[\![\langle\Gamma\rangle]\!]$. The argument for $[X]=[\langle X\rangle]$ is similar, which proves property 4.

Let $X\leq G$ be a finite subgroup. Since X is closed under addition it is easy to check that if $\gamma\notin[X]$ then $\widehat{X}(\gamma)=0$, and if $\gamma\in[X]$ then $\widehat{X}(\gamma)=|X|$, and hence by Parseval's identity,

$$|X|=\int_\gamma\left|\widehat{X}(\gamma)\right|^2\,\mathrm{d}\gamma=\int_{\gamma\in[X]}\left|\widehat{X}(\gamma)\right|^2\,\mathrm{d}\gamma=\mu([X])\,|X|^2\,,$$

and property 5 follows. We verify property 7 in a similar fashion. Let $\Gamma\subset\widehat{G}$ with positive measure. We observe as above that $\widehat{[\![\Gamma]\!]}(\gamma)=|[\![\Gamma]\!]|$ for all $\gamma\in\Gamma$. Furthermore, if $\gamma\notin\langle\Gamma\rangle$

then there exists some $x \in [\![\Gamma]\!]$ such that $\gamma(x) \neq 1$ and hence $\widehat{[\![\Gamma]\!]}(\gamma) = 0$. It follows that

$$|X| = \int \left|\widehat{[\![\Gamma]\!]}(\gamma)\right|^2 \mathrm{d}\gamma = \int_{\langle\Gamma\rangle} \left|\widehat{[\![\Gamma]\!]}(\gamma)\right|^2 \mathrm{d}\gamma \geq \mu(\Gamma)\,|[\![\Gamma]\!]|^2\,,$$

and property 7 follows.

It remains to prove property 6. As above, since $[\![[X]]\!]$ is a subgroup which contains X we have $\langle X\rangle \subset [\![[X]]\!]$. Furthermore, by property 7 we know that $[\![[X]]\!]$ is a finite subgroup. It is thus sufficient to show that

$$|[\![[X]]\!]| = \mu([[\![[X]]\!]])^{-1} \leq |\langle X\rangle| = \mu([X])^{-1},$$

which follows from property 3. This proves property 6. □

We recall that the left action of R on G induces a right action of R on $\widehat{G}$, defined by $(\gamma c)(x) = \gamma(cx)$ for any $\gamma \in \widehat{G}$ and $x \in G$.

Lemma 2.5. *Let $\Lambda_0, \Lambda \subset \widehat{G}$ and $c \in R$. If $\Lambda_0 \cdot c \subset \Lambda$ then $c \cdot [\![\Lambda]\!] \subset [\![\Lambda_0]\!]$.*

Proof. We need to show that for all $\gamma \in \Lambda_0$ and all $y \in c\cdot[\![\Lambda]\!]$ we have $\gamma(y) = 1$. Unpacking the definitions, this is simply saying that if $\lambda(z) = 1$ for all $\lambda \in \Lambda$ then $\gamma(cz) = (\gamma c)(z) = 1$, which follows since $\gamma c \in \Lambda_0 \cdot c \subset \Lambda$. □

We will now use the operators $[\cdot]$ and $[\![\cdot]\!]$ to present an explicit construction for a permissible collection for any finite subgroup. As mentioned above, because a finite subgroup completely shelters all of its subsets we are able to construct a permissible collection with the shelter parameter $\delta = 0$.

Lemma 2.6. *Suppose that $\mathbf{c} \in (R^*)^s$ is a commutative s-tuple. Let $X = a \cdot [\![\Lambda]\!] \leq G$ be a finite subgroup where $a \in R^*$ commutes with $\mathbf{c}$, and let $\mathfrak{B}$ be the collection of all finite subgroups of the form $a' \cdot [\![(\Lambda \cdot T) \cup \Gamma)]\!]$ for some $a' \in R^*$ which commutes with $\mathbf{c}$, finite set $\Gamma \subset \widehat{G}$ of size $|\Gamma| \leq d$ and $T \subset \{c_1, \ldots, c_s\}^2$. Then $\mathfrak{B}$ is permissible for $(X, \mathbf{c}; 0, d)$.*

Proof. We shall first construct a good pair for $(X, \mathbf{c}; 0)$. Since X is a subgroup it suffices to find a pair $(\tilde{B}, \tilde{B}')$ of finite subgroups of G such that $c_1 \cdot \tilde{B}, c_s \cdot \tilde{B}, c_1c_s \cdot \tilde{B}' \subset X$ and $c_i \cdot \tilde{B}' \subset \tilde{B}$ for $1 < i < s$. Let

$$\tilde{B} = a \cdot [\![\Lambda \cdot \{c_1, c_s\}]\!]\,,$$

so that $c_1 \cdot \tilde{B}, c_s \cdot \tilde{B} \subset X$ by Lemma 2.5. Similarly, let

$$\tilde{B}' = a \cdot [\![\Lambda \cdot \{c_1, c_s\} \cdot \{c_2, \dots, c_{s-1}\} \cup \Lambda \cdot c_1 c_s]\!] .$$

It follows that $c_1 c_s \cdot \tilde{B}' \subset X$. Furthermore, we also have by Lemma 2.5 that $c_i \cdot \tilde{B}' \subset \tilde{B}$ for $1 < i < s$. This proves that $(\tilde{B}, \tilde{B}')$ is a good pair for $(X, \mathbf{c}; 0)$. For $\mathfrak{B}$ to be permissible we first require that it contain the subgroups $c_1 \cdot \tilde{B}$, $c_s \cdot \tilde{B}$ and $c_1 c_s \cdot \tilde{B}'$. For this it suffices to note that

$$c_1 \cdot \tilde{B} = c_1 a \cdot [\![\Lambda \cdot \{c_1, c_s\}]\!] \in \mathfrak{B},$$

for example. Furthermore, for each finite $\Gamma \subset \widehat{G}$ of size at most C and $1 < i < s$ we require $\mathfrak{B}$ to contain $c_1 c_s c_i \cdot X'$ where X' is some subgroup of $\tilde{B}'$ with control of Γ. For this it suffices to take

$$X' = a \cdot [\![\Lambda \cdot \{c_1, c_s\} \cdot \{c_2, \dots, c_{s-1}\} \cup \Lambda \cdot c_1 c_s \cup \Gamma \cdot a]\!] .$$

□

This construction can be iterated to give an explicit construction of chain from X with the $DI(0, d(i); \mathbf{c})$ property, where

$$d(i) = C(1+C)^{-i/(s-2)} \alpha^{-1/(s-2)} \mathcal{L}(\alpha).$$

Namely, we let $\mathfrak{X}_0 = \{X\}$ and for $j > 0$ every member of $\mathfrak{X}_j$ is of the shape $a \cdot [\![\Lambda_j]\!]$ for some $a \in R^*$ which commutes with $\mathbf{c}$ and some

$$\Lambda_j \subset \left(\Lambda \cup \bigcup_{i=1}^{j} \Gamma_i \right) \cdot \{1, c_1, \dots, c_s\}^{2j}$$

for some finite sets $\Gamma_i \subset \widehat{G}$ such that $|\Gamma_i| \leq d(i)$ for $1 \leq i \leq j$. Lemma 2.6 ensures that for every $X \in \mathfrak{X}_{i-1}$ the link $\mathfrak{X}_i$ contains some collection of sets which is permissible for $(X, \mathbf{c}; 0, d(i))$ as required. Summing the geometric series we can include all such Γ_i in a single finite $\Gamma \subset \widehat{G}$ of size $O_C(\alpha^{-1/(s-2)} \mathcal{L}(\alpha))$. In particular, the following is an immediate corollary of Theorem 2.3.

Theorem 2.7. *Let $X = [\![\Lambda]\!] \leq G$ be a finite subgroup and $A \subset X$ with density α. Let $\mathbf{c} \in (R^*)^s$ be a commutative s-tuple such that $c_1 + \dots + c_s = 0$. Finally, suppose that A contains only trivial solutions to* (2.1).

There exists an integer $0 \leq K \ll_s \mathcal{L}(\alpha)$ *and a finite set* $\Gamma \subset \widehat{G}$ *such that* $|\Gamma| \ll_s \alpha^{-1/(s-2)}\mathcal{L}(\alpha)$ *and*

$$\mu(\langle(\Lambda \cup \Gamma) \cdot \{1, c_1, \dots, c_s\}^{2K}\rangle) \gg_s \alpha^2. \tag{2.3}$$

Proof. This follows immediately from our chain construction above, Theorem 2.3 and part 5 of Lemma 2.4. □

2.2.1 Function fields

As a first demonstration of the theory in the subgroup case we discuss the problem where $G = \mathbb{F}_q[t]$, the ring of polynomials in the indeterminate t over $\mathbb{F}_q$, the field with q elements. Since this is a vector space over $\mathbb{F}_q$ there are certainly many subgroups in both G and $\widehat{G}$. Furthermore, $\mathbb{F}_q[t]$ has a natural structure as an $\mathbb{F}_q[t]$-module, and hence there is a natural embedding of $\mathbb{F}_q[t]\backslash\{0\} \subset R^*$, the ring of injective endomorphisms on $\mathbb{F}_q[t]$. In particular, we fix some $\mathbf{c} \in (\mathbb{F}_q[t]\backslash\{0\})^s$ such that $c_1 + \cdots + c_s = 0$, and observe that the set $\{1, c_1, \dots, c_s\}^{2K}$ is a subset of $\mathbb{F}_q[t]_{2\ell K+1}$, where $\ell = \max_{1\leq i\leq s} \deg c_i$ and $\mathbb{F}_q[t]_n$ denotes the set of polynomials of degree less than n.

We desire an upper bound for the density of subsets of some structured set X with only trivial solutions to (2.1). Such a structured set does need to be a subgroup for Theorem 2.7 to hold, but the conclusion of Theorem 2.7 shows that we also need to have a reasonable amount of control on how $\langle[X] \cdot \mathbb{F}_q[t]_{2\ell K+1}\rangle$ grows as K increases. Consider, for example, the case

$$X = \left\{\sum_{i=0}^{N} a_i t^{2i} : a_i \in \mathbb{F}_q\right\}.$$

It is easy to check that

$$[X] = \left\{x \in \widehat{\mathbb{F}_q[t]} : \deg x < -2N\right\} + \left\{\sum_{i=0}^{2N} a_i t^{-2i} : a_i \in \mathbb{F}_q\right\}.$$

In particular, $\mu(\langle[X] \cdot \mathbb{F}_q[t]_2\rangle) = \mu(\widehat{\mathbb{F}_q[t]}) = 1$, and hence Theorem 2.7 will deliver only the trivial bound $\alpha \leq 1$ when $\ell > 0$, even though X is a finite subgroup of $\mathbb{F}_q[t]$. When $\ell = 0$ all the coefficients lie in $\mathbb{F}_q$, and here one can use simpler methods to obtain good quantitative results for all $\mathbb{F}_q$-subspaces $X \leq G$, as done by Liu and Spencer [39].

Hence we shall limit our attention to the case $\ell \geq 1$, which entails restricting our focus from general finite subgroups of $\mathbb{F}_q[t]$ to those with suitably regular arithmetic behaviour.

In particular, we shall consider the case when

$$X = \mathcal{A}_N = \{x \in \mathbb{F}_q[t] : \deg x < N\},$$

which is the $\mathbb{F}_q[t]$-analogue of the interval $\{1, \dots, N\}$ in the integers. In this case we have

$$[X] = \{x \in \widehat{\mathbb{F}_q[t]} : \deg x < -N\}.$$

It follows that the subgroup $[X] \cdot \mathbb{F}_q[t]_{2\ell K+1} \leq \widehat{\mathbb{F}_q[t]}$ has measure of at most $q^{2\ell K+1-N}$. Furthermore, we may trivially contain any finite set $\Gamma \subset \widehat{\mathbb{F}_q[t]}$ in a finite subgroup of $\widehat{\mathbb{F}_q[t]}$ of size at most $q^{|\Gamma|}$, and $\Gamma \cdot \mathbb{F}_q[t]_n$ in a subgroup of size at most $q^{n|\Gamma|}$. It follows that

$$\mu(\langle([X] \cup \Gamma) \cdot \{1, c_1, \dots, c_s\}^{2K}\rangle) \ll q^{O(\ell K|\Gamma|)-N},$$

where the implied constants are absolute. Thus the equation (2.3), on recalling the bounds for K and $|\Gamma|$ present in Theorem 2.7, becomes

$$\mathcal{L}(\alpha) \gg_s \log q \left(N - \ell\alpha^{-1/(s-2)}\mathcal{L}(\alpha)^2\right).$$

Rearranging this inequality yields the following corollary of Theorem 2.7, which gives a strong quantitative result for Roth's problem for linear equations with coefficients from $\mathbb{F}_q[t]$.

Corollary 2.8. *Let $A \subset \mathbb{F}_q[t]_N$ with density α and $\mathbf{c} \in (\mathbb{F}_q[t]\backslash\{0\})^s$, with $\ell = \max \deg c_i$, be such that $c_1 + \cdots + c_s = 0$. Then if A has only trivial solutions to* (2.1) *we have*

$$\alpha \ll_s \left(\ell \frac{(\log N)^2}{N}\right)^{s-2},$$

where $\delta > 0$ is some absolute constant depending only on s.

We observe in particular that the implied constant hidden in the $\ll$ depends only on s, and not on the coefficients $\mathbf{c}$ or even on q. Theorem 2.1 is an immediate corollary.

2.2.2 Drinfeld modules

An interesting feature of $\mathbb{F}_q[t]$ is that there are other natural ways to give it the structure of an $\mathbb{F}_q[t]$-module, which arise from the linearity of the Frobenius map $x \mapsto x^p$. In

particular, for any $r \geq 0$ and $a_1, \dots, a_r \in \mathbb{F}_q[t]$ with $a_r \neq 0$ we may define a linear map $\rho_t : \mathbb{F}_q[t] \to \mathbb{F}_q[t]$ by

$$\rho_t(x) = tx + a_1 x^p + a_2 x^{p^2} + \cdots + a_r x^{p^r}.$$

If we also define $\rho_1(x) = x$ then these combine to give, for any $a \in \mathbb{F}_q[t]$, an $\mathbb{F}_q$-linear map $\rho_a : \mathbb{F}_q[t] \to \mathbb{F}_q[t]$ such that $\rho_{ab} = \rho_a \rho_b$. This induces an action of $\mathbb{F}_q[t]$ on $\mathbb{F}_q[t]$ – namely, by letting $c \cdot a = \rho_c(a)$. Explicitly,

$$\rho_c(a) = \sum_{i=0}^{d} a_i \rho_{t^i}(a) \text{ where } c = \sum_{i=0}^{d} a_i t^i,$$

and ρ_{t^i} is defined as the i-fold composition of ρ_t. The map ρ is known as a Drinfeld module of rank r; the case $r = 0$ gives the trivial action of $\mathbb{F}_q[t]$ on $\mathbb{F}_q[t]$ discussed in the previous section. The concept of a Drinfeld module can be vastly generalised; see Chapter 13 of [49] for more details.

By arguing exactly as in the previous section a similar result to Corollary 2.8 could be obtained, with $\mathbb{F}_q[t]_N$ replaced by $\mathbb{F}_q[t]_{N,\rho} = \{\rho_x(1) : x \in \mathbb{F}_q[t]_n\}$. This is to be expected, as everything is isomorphic to the rank 0 case. What is more interesting, however, is that one can control the behaviour of $\mathbb{F}_q[t]_N$ itself relative to a Drinfeld module, leading to the following result.

Corollary 2.9. *Let ρ be any Drinfeld module. There exist constants $C, \delta > 0$, depending only on ρ and s, such that the following holds. Let $A \subset \mathbb{F}_q[t]_N$ with density α and $\mathbf{c} \in (\mathbb{F}_q[t] \backslash \{0\})^s$, with $\ell = \max \deg c_i$, be such that $c_1 + \cdots + c_s = 0$. Then, if there are only trivial solutions to*

$$\rho_{c_1}(x_1) + \cdots + \rho_{c_s}(x_s) = 0$$

with $x_i \in A$, we have

$$\alpha \ll_s \frac{1}{N^{C/\ell}},$$

where the implied constant depends only on s.

Proof. For any finite $\Gamma \subset \widehat{G}$ the set $\Gamma \cdot \{1, c_1, \dots, c_s\}^{2K}$ is trivially contained in a finite subgroup of $\widehat{G}$ of size at most $q^{O(\ell K |\Gamma|)}$. Furthermore, if $\Lambda = [\mathbb{F}_q[t]_N] \cdot \mathbb{F}_q[t]_{8K\ell}$ then $[\![\Lambda]\!]$ contains every finite subgroup $X \leq \mathbb{F}_q[t]$ such that

$$\bigcap_{a \in \mathbb{F}_q[t]_{8\ell K}} \rho_a(X) \subset \mathbb{F}_q[t]_N,$$

and so $\mu(\Lambda) \leq |X|^{-1}$. In general, we claim that $\cap_{a \in \mathbb{F}_q[t]_M} \rho_a(\mathbb{F}_q[t]_{N/C^M}) \subset \mathbb{F}_q[t]_N$ for some constant $C > 0$ depending only on ρ. This follows since for any $x \in \mathbb{F}_q[t]$ with degree d there is a constant C such that $\deg \rho_t(x) = Cd$, and hence in general if a has degree M then $\deg \rho_a(x) \leq C^M d$.

In particular, $\mu(\Lambda) \leq q^{-N/C^{\ell K}}$ for some $C > 0$ depending only on ρ. Theorem 2.7 then implies that

$$\exp_q \left(\ell \alpha^{-1/(s-2)} \mathcal{L}(\alpha)^2 - N \alpha^{O_{\rho,s}(\ell)} \right) \geq \alpha,$$

and the claim follows. □

As a concrete example, we consider the simplest non-trivial example of a Drinfeld module; namely, the Carlitz module defined by $\rho_t(x) = tx + x^p$. Corollary 2.9 implies that there exists an absolute constant $C > 0$ such that if $A \subset \mathbb{F}_q[t]_N$ has density $\alpha \gg N^{-C}$ then it contains non-trivial solutions to both

$$x_1 + tx_2 - (1+t)x_3 = 0 \text{ and } x_1 + tx_2 - (1+t)x_3 = (x_3 - x_2)^p.$$

In general, we obtain the following version of the Erdős-Turán conjecture for $\mathbb{F}_q[t]$, which is particularly interesting in that it applies simultaneously to all Drinfeld modules.

Theorem 2.10. *Suppose that $A \subset \mathbb{F}_q[t]$ has positive lower density in the sense that*

$$\liminf_{N \to \infty} |A \cap \mathbb{F}_q[t]_N| \, q^{-N} > 0.$$

Then for any $s \geq 3$, coefficients $\mathbf{c} \in (\mathbb{F}_q[t] \backslash \{0\})^s$ such that $c_1 + \cdots + c_s = 0$, and Drinfeld module ρ, the set A contains infinitely many solutions to

$$\rho_{c_1}(x_1) + \cdots + \rho_{c_s}(x_s) = 0$$

such that the variables x_i are all distinct.

2.3 Weak structure

We have seen in the previous section that when every link in the chain consists of finite subgroups then the construction of the next link in the chain is quite straightforward. In the general case, however, the construction is more delicate. We still require our initial set X to possess a fair degree of additive structure, but it is not necessary that

it be a subgroup. In particular, we will be able to apply our constructions in the case $X = \{1, \dots, N\}$, which is the classical setting of Roth's theorem.

We first give some preliminary definitions. We say that X has additive growth of K if for all $n \geq 1$ we have

$$|2nX| \leq K\,|nX|\,.$$

The sets for which we are able to construct a suitable chain will be restricted to those which are of the shape nX where n is large and X is a symmetric set with additive growth $O(1)$. The primary example is when $X = \{-1, 0, 1\}$ so that $nX = \{-n, \dots, n\}$. In fact, in the integers, this example is essentially captures all the sets that we are able to construct chains for. Indeed, the condition of bounded additive growth certainly implies that X has polynomial growth in the sense that there exists some $d \geq 1$ with $|nX| \leq n^d\,|X|$ for all $n \geq 1$. It follows from inverse sumset results (see, for example, the survey of Sanders [63]) that X is well-approximated in some sense by a generalised arithmetic progression. More generally, for an arbitrary abelian group our methods are limited to sets which closely resemble coset progressions, which are the sum of a subgroup and an arithmetic progression. Thus, for the case $G = \mathbb{Z}$ at least, the reader would lose little in taking $X = \{-1, 0, 1\}$ in what follows, although we state our results and proofs more generally to make plain precisely what structure is required.

We say that $\rho : \widehat{G} \to [0, 1]$ is cofinite if $\rho(\gamma) = 1$ for all but finitely many $\gamma \in \widehat{G}$. For any cofinite ρ and $X \subset G$ we define the Bohr set with width ρ and ground set X as

$$X(\rho) = \{x \in X : |1 - \gamma(x)| \leq 2\rho(\gamma) \text{ for all } \gamma \in \widehat{G}\}.$$

The importance of such sets in the study of translation invariant equations was first recognised by Bourgain [6], although the definition we give here is more general, in that we allow arbitrary ground sets X, and our discussion is valid for any abelian group G. For any cofinite $\rho_1, \rho_2 : \widehat{G} \to (0, 1]$ let

$$\Phi(\rho_1, \rho_2) = \prod_{\gamma \in \widehat{G}} \frac{\rho_1(\gamma)}{\rho_2(\gamma)},$$

which is a finite product by cofiniteness, and let $\Phi(\rho) = \Phi(\rho, 1)$. If $\lambda \in [0, 1]$ and ρ is cofinite then we will abuse notation by writing $\lambda\rho$ for the cofinite dilate defined by

$$(\lambda\rho)(\gamma) = \begin{cases} \lambda(\rho(\gamma)) & \text{if } \rho(\gamma) < 1 \text{ and} \\ 1 & \text{otherwise.} \end{cases}$$

Furthermore, by $\rho_1 \wedge \rho_2$ we denote the cofinite function defined by

$$(\rho_1 \wedge \rho_2)(\gamma) = \min(\rho_1(\gamma), \rho_2(\gamma)).$$

Finally, if ρ is cofinite then $\mathrm{rk}(\rho)$ denotes the number of $\gamma \in \widehat{G}$ such that $\rho(\gamma) < 1$.

Each link in our chain will be composed of sets of form $X(\rho)$ for suitably chosen X and ρ, and so it is crucial to have a good lower bound on the size of such sets. For general finite X such a bound cannot be obtained, but if X has a reasonable amount of additive structure then we can use a probabilistic argument. We will require the following covering lemma due to Ruzsa [57]. We will give a proof of this lemma as Lemma 4.15 in Chapter 4.

Lemma 2.11 (Ruzsa [57])**.** *Let X and Y be finite subsets of G. If $|X+Y| \leq K|Y|$ then $X \subset Y - Y + T$ for some $T \subset G$ such that $|T| \leq K$.*

In the original work by Bourgain [6] a lower bound on the size of Bohr sets was obtained by Fourier analysis, and in particular the properties of the Fejer kernel. This argument is special to the integers, however. In the book of Tao and Vu [75] they give (as their Lemma 4.20) an alternative argument for the case when X is a finite group, which is robust enough to adapt for our purposes.

Lemma 2.12. *For all finite $X \subset G$ and cofinite $\rho_1, \rho_2 : \widehat{G} \to (0,1]$ such that $\rho_2 \leq \rho_1$ we have*

$$|X(\rho_1)| \leq 4^{\mathrm{rk}(\rho_2)} \Phi(\rho_1, \rho_2) \, |(X - X)(\rho_2)| \, .$$

In particular if ρ is cofinite then

$$|(X - X)(\rho)| \geq 4^{-\mathrm{rk}(\rho)} \Phi(\rho) \, |X| \, .$$

Furthermore, if Y is a finite set such that $|X+Y| \leq K|Y|$ then

$$|X(\rho_1)| \leq K 4^{\mathrm{rk}(\rho_2)} \Phi(\rho_1, \rho_2) \, |(2Y - 2Y)(\rho_2)| \, .$$

Proof. Let Γ denote the set of $\gamma \in \widehat{G}$ such that $\rho_2(\gamma) < 1$ and for $\gamma \in \Gamma$ consider the set of $\nu \in \mathbb{C}$ such that $|\nu| = 1$ and $|1 - \nu| \leq 2\rho_1(\gamma)$. It is clear that there exist at most $k_\gamma \leq \lceil 2\rho_1(\gamma)/\rho_2(\gamma) \rceil$ circles of radius $\rho_2(\gamma)$ which completely cover such a set. In particular $X(\rho_1)$ is covered by at most $\prod_{\gamma \in \Gamma} k_\gamma$ many sets of the shape

$$\{x \in X : \gamma(x) \in D_\gamma \text{ for all } \gamma \in \Gamma\},$$

where each D_γ is a circle of radius at most $\rho_2(\gamma)$. If such a set is denoted by X' then it is clear that $X' - X' \subset (X - X)(\rho_2)$. Hence

$$|X(\rho_1)| \leq \prod_{\gamma \in \Gamma} k_\gamma \, |(X - X)(\rho_2)| \leq 4^{\mathrm{rk}(\rho_2)} \Phi(\rho_1, \rho_2) \, |(X - X)(\rho_2)|$$

as required. If $|X + Y| \leq K |X|$ then Lemma 2.11 implies that X is covered by at most K translates of $Y - Y$, and hence each such X' is covered by at most K sets of the shape X'' such that $X'' - X'' \subset (2Y - 2Y)(\rho_2)$, and the lemma follows as before. □

Lemma 2.12 is particularly useful when applied to sets with bounded additive growth.

Corollary 2.13. *Suppose that a finite symmetric set $X \subset G$ has additive growth of K. For all $n \geq m \geq 2^{-j} n \geq 4$,*

$$|(mX)(\rho/2)| \geq K^{-j-4} 8^{-\mathrm{rk}(\rho)} \, |(nX)(\rho)| \, .$$

Proof. Let $m' = \lfloor m/4 \rfloor \geq 1$ and choose k such that $2^{k-1} < n \leq 2^k$, so that $m' \geq 2^{k-j-3}$. Since X has additive growth of K,

$$|nX + m'X| \leq |2nX| \leq \left|2^{k+1} X\right| \leq K^{j+4} \left|2^{k-j-3} X\right| \leq K^{j+4} \, |m'X| \, .$$

It follows from Lemma 2.12 that

$$|(4m'X)(\rho/2)| \geq K^{-j-4} 8^{-\mathrm{rk}(\rho)} \, |(nX)(\rho)| \, ,$$

and the claim follows since $4m'X \subset mX$. □

To construct a chain we need to produce, for any given set in some link of the chain, sets for the next chain which are sheltered by it, and are sufficiently structured in their own right. One might first hope that, if our original set is a Bohr set, then all subsets which are also Bohr sets whose width functions are suitably chosen dilates of the original width function are suitable for this purpose. In general, however, it is impossible to produce the required amount of shelter from such sets – the best one can hope for is the trivial relationship $X(\rho) + X(\lambda\rho) \subset (2X)((1 + \lambda)\rho)$, and hence the sumset can grow by a factor exponential in the rank of ρ, which is far too costly for our purposes.

It was Bourgain [6] who first saw how to avoid this difficulty, and hence enabled the use of Bohr sets of the type $X(\rho)$ in proving Roth's theorem. In particular, by a simple

covering argument he showed that, while an arbitrary Bohr set may not shelter its subsets very well, there exist many 'regular' Bohr sets which do shelter smaller Bohr sets very effectively. In particular, these regular Bohr sets are so plentiful that we may restrict the chain to include only regular Bohr sets without any real quantitative cost.

In general, we say that (X, n, ρ) is (δ, κ)-*regular* if, letting $\lambda = \kappa\delta \mathrm{rk}(\rho)^{-1}$, for all $1 \leq s \leq \lambda n$ and $\rho' \leq \lambda\rho$ the set $(sX)(\rho')$ is δ-sheltered by $(nX)(\rho)$. Bourgain's covering argument from [6] allows one to always 'regularise' any given Bohr set without much cost.

We remark that this problem of regularity is the reason that our results are limited to ground sets of the shape nX for some structured X and large n, as can be seen from examining the proof below. In turn, as mentioned above, this essentially limits us, in the integers, to arithmetic progressions. If a suitable alternative route to regularity is found it may be possible to extend our results to other cases of interest, such as when X is the set of the first n prime numbers or square numbers.

Lemma 2.14. *Let $\delta \in (0, 1]$ and $\rho : \widehat{G} \to (0, 1]$ be a cofinite function. If X has additive growth of $K \geq 4$ and $n \geq 4$ then there exists $\rho' \approx \rho$ and $n' \approx n$ such that (X, n', ρ') is $(\delta, 2^{-5}(\log K)^{-1})$-regular.*

Proof. Let $r = \mathrm{rk}(\rho)$ and choose some integer m such that $2^4\delta^{-1}r \log K > m \geq 2^3\delta^{-1}r \log K$. Let $t = \lceil n/2 \rceil$ and observe that if $s \leq 2^{-5}(\log K)^{-1}\delta r^{-1}n$ then $ms < n/2$, so that $t + ms \leq n$. Let $X_j = (t + js)X$ for $0 \leq j \leq m$ and $\lambda_i = 1/2 + i/2m$ for $0 \leq i \leq m$, so that

$$1/2 = \lambda_0 < \lambda_2 < \cdots < \lambda_m = 1.$$

Hence

$$\frac{|(tX)(\rho/2)|}{|(nX)(\rho)|} \leq \frac{|X_0(\lambda_0\rho)|}{|X_m(\lambda_m\rho)|} = \prod_{i=0}^{m-1} \frac{|X_i(\lambda_i\rho)|}{|X_{i+1}(\lambda_{i+1}\rho)|}.$$

Suppose that each of the factors on the right hand side is at most $(1 + \delta)^{-1}$. Using the inequality $1 + x \geq e^{0.567x}$, valid for all $0 \leq x \leq 1$,

$$\frac{|(tX)(\rho/2)|}{|(nX)(\rho)|} \leq (1 + \delta)^{-m} \leq \exp(-0.567\delta m) < K^{-5}8^{-r},$$

since $\delta m/r > 5 \log K + \log 8/0.567$. By Corollary 2.13, however, the left hand side is at least $K^{-5}8^{-r}$ and we have a contradiction. It follows that for some $0 \leq i < m$

$$|X_{i+1}(\lambda_{i+1}\rho)| \leq (1 + \delta) |X_i(\lambda_i\rho)|,$$

that is,

$$|((t+is)X)(\lambda_i\rho) + (sX)(\rho/2m)| \leq (1+\delta)\,|((t+is)X)(\lambda_i\rho)|\,.$$

The proof is completed by letting $n' = t + is$ and $\rho' = \lambda_i\rho$. □

Before presenting our chain construction we need to make a couple of technical definitions that will control how dilations by the coefficients $c_i \in R$ affect Bohr sets. Given finite sets $A \subset R$ and $X \subset G$ we say that A is *m-constrained* over X if $a \cdot X \subset mX$ for all $a \in A$. We will also abuse this notation for tuples $\mathbf{c} \in R^s$ by saying that $\mathbf{c}$ is m-constrained over X if the set $\{c_1, \dots, c_s\}$ is.

If $a \in R$ and ρ is cofinite then we define

$$(\rho \circ a)(\gamma) = \begin{cases} \inf_{\gamma' a = \gamma} \rho(\gamma') & \text{if } \gamma = \gamma' a \text{ for some } \gamma' \in \widehat{G}, \text{ and} \\ 1 & \text{otherwise.} \end{cases}$$

If $A \subset R^*$ then we let $\rho \circ A = \min_{a \in A} \rho \circ a$. We observe that this operation commutes with dilation, so that $\lambda(\rho \circ a) = (\lambda\rho) \circ a$. Furthermore, $\mathrm{rk}(\rho \circ A) \leq |A|\,\mathrm{rk}(\rho)$.

These definitions will be important because, as can be seen from the definition of a permissible collection, to construct a permissible collection for a Bohr set X we will need to be able to construct a sheltered Bohr set X' such that $a \cdot X' \subset X$. The following lemma demonstrates how our definitions will allow this.

Lemma 2.15. *Let $X \subset G$ be a finite set and let $A \subset R$ be a finite set such that A is m-constrained over X. If $n' \leq n/m$ and $X' = (n'X)(\rho \circ A)$ then $a \cdot X' \subset (nX)(\rho)$ for all $a \in A$.*

Proof. By the definition of m-constraint it is clear that $a \cdot (n'X) \subset nX$ for all $a \in A$. Let $x \in (n'X)(\rho \circ A)$ and suppose that $\gamma \in \widehat{G}$ is such that $\rho(\gamma) < 1$. We have

$$|1 - \gamma(ax)| = |1 - (\gamma a)(x)| \leq 2(\rho \circ A)(\gamma a) \leq 2\rho(\gamma)$$

by definition, and hence $ax \in (nX)(\rho)$ as required. □

We now at last combine our technical lemmata to perform a chain construction which is valid even for sets which are not subgroups, although we will still require a fairly strict amount of additive control; as noted above, in the case $G = \mathbb{Z}$ this is restrictive enough to essentially limit us to arithmetic progressions.

Lemma 2.16. *Let $X \subset G$ be a finite symmetric set with additive control $K \geq 4$ and $\rho : \widehat{G} \to [0,1]$ be cofinite with rank r. Let $\mathbf{c} \in (R^*)^s$ be commutative and m-constrained over X. Let $\kappa = 2^{-5}(\log K)^{-1}$. If (X, n, ρ) is (δ, κ)-regular and $a \in R^*$ commutes with $\mathbf{c}$ then the following collection $\mathfrak{B}$ is permissible for $(a \cdot ((nX)(\rho)), \mathbf{c}; \delta, d)$.*

The collection $\mathfrak{B}$ comprises every set of the shape $a' \cdot ((n'X)(\rho'))$ where

1. *(X, n', ρ') is (δ, κ)-regular,*
2. *$a' \in R^*$ commutes with $\mathbf{c}$,*
3. *$n' \gg_s \lfloor m^{-2}\lambda n \rfloor$, and*
4. *$\rho' \geq \lambda(\rho \circ \{1, c_1, \dots, c_s\}^2) \wedge \rho_\Gamma$, for some cofinite ρ_Γ with rank at most d satisfying $\rho_\Gamma \gg 1/d$,*

where $\lambda \gg (\delta/r \log K)^3$.

Loosely speaking, our chain construction allows us to proceed from a set of the shape $a \cdot ((nX)(\rho))$ to one of the shape $a' \cdot ((n'X)(\rho'))$ where n' and ρ' are suitably controlled as above.

Proof. We observe that, without loss of generality, we may suppose that $a = 1$, for all the sets in our construction can be dilated by a without harming any of the necessary shelter or control properties.

We first need to construct a good pair $(\tilde{B}, \tilde{B}')$ for $((nX)(\rho), \mathbf{c}; \delta)$. The hypothesis of regularity guarantees that if $\lambda_1 = \kappa\delta/r$ and $n_1 = \lfloor \lambda_1 n \rfloor$ then $(n_1X)(\lambda_1\rho)$ is δ-sheltered by $(nX)(\rho)$. For $(\tilde{B}, \tilde{B}')$ to satisfy the first property of a good pair it thus suffices that $c_1 \cdot \tilde{B}, c_s \cdot \tilde{B}, c_1c_s \cdot \tilde{B}' \subset (n_1X)(\lambda_1\rho)$. Let $\tilde{n} \approx \lfloor n_1/m \rfloor$ and $\tilde{\rho} \approx \lambda_1\rho \circ \{1, c_1, c_s\}$ be chosen such that $(X, \tilde{n}, \tilde{\rho})$ is (δ, κ)-regular. By Lemma 2.15 if $\tilde{B} = (\tilde{n}X)(\tilde{\rho})$ then $c_1 \cdot \tilde{B}, c_s \cdot \tilde{B} \subset (n_1X)(\lambda_1\rho)$ as required.

We further require that $c_2 \cdot \tilde{B}' + \cdots + c_{s-1} \cdot \tilde{B}'$ be δ-sheltered by $\tilde{B}$ and $c_1c_s \cdot \tilde{B}' \subset (n_1X)(\lambda_1\rho)$. By the regularity of $\tilde{B}$ imposed above if $\lambda_2 = \kappa\delta/3rs$ and $n_3 \leq \lambda_2\tilde{n}$ then (on noting that $\tilde{\rho}$ has rank at most $3r$ by construction) $s((n_3X)(\lambda_2\tilde{\rho}))$ is δ-sheltered by $\tilde{B}$. In particular, if we choose $\tilde{n}' \approx \lfloor \lambda_2\tilde{n}/m \rfloor$ and $\tilde{\rho}' \approx \lambda_2\tilde{\rho} \circ \{1, c_2, \dots, c_{s-1}, c_1c_s\}$ such that $(X, \tilde{n}', \tilde{\rho}')$ is (δ, κ)-regular and let $\tilde{B}' = (\tilde{n}'X)(\tilde{\rho}')$, then by Lemma 2.15 once again $\{1, c_2, \dots, c_{s-1}, c_1c_s\} \cdot \tilde{B}' \subset (\lambda_2\tilde{n}X)(\lambda_2\tilde{\rho})$, which implies the required amount of shelter.

This completes the construction of a suitably good pair $(\tilde{B}, \tilde{B}')$. We pause to observe the bounds

$$\tilde{n}' \gg_s \left(\frac{\kappa\delta}{rm}\right)^2 n \text{ and } \tilde{\rho}' \gg_s \left(\frac{\kappa\delta}{r}\right)^2 \rho \circ \{1, c_1, \dots, c_s\}^2.$$

It remains to extend this to give a construction of a collection $\mathfrak{B}$ which is permissible for $((nX)(\rho), \mathbf{c}; \delta, d)$. We first require that $c_1c_s \cdot \tilde{B}'$, $c_1 \cdot \tilde{B}$ and $c_s \cdot \tilde{B}$ are all members of $\mathfrak{B}$, which is clearly the case with the collection given in the statement of the lemma. We now fix some finite $\Gamma \subset \widehat{G}$ of size at most d and some $1 < i < s$. It remains to include in $\mathfrak{B}$ some $c_1c_sc_i \cdot X'$ where $X' \subset \tilde{B}'$ has $(4d)^{-1}$-control of Γ and is δ-sheltered by $\tilde{B}'$.

With $\lambda_3 = \kappa\delta/4sr$, we choose an integer $n' \approx \lfloor \lambda_3 \tilde{n}' \rfloor$ and $\rho' \approx \lambda_3 \tilde{\rho}'$ so that (X, n', ρ') is (δ, κ)-regular and $(n'X)(\rho')$ is δ-sheltered by $\tilde{B}'$. Furthermore, we certainly have $X' \subset \tilde{B}'$ as required. Finally, we let $X' = (n'X)(\rho' \wedge \rho_\Gamma)$, where $\rho_\Gamma(\gamma) \approx 1/4d$ if $\gamma \in \Gamma$ and is 1 otherwise, and is chosen such that $(X, N', \rho' \wedge \rho_\Gamma)$ is (δ, κ)-regular. This completes the construction of X', and it is clear that it has the form specified in the statement of the lemma. □

We may now iterate this construction to create a chain suitable for an application of Theorem 2.3, which yields the following.

Theorem 2.17. *Let $X \subset G$ be a finite symmetric set with additive control $K \geq 4$ and $n \geq 4$, and let $A \subset nX$ with density α. Let $\mathbf{c} = (c_1, \dots, c_s) \in (R^*)^s$ be a commutative s-tuple which is m-constrained by X such that $c_1 + \cdots + c_s = 0$. Finally, suppose that A contains only trivial solutions to* (2.1)*.*

There exist integers $0 \leq M \ll_s \mathcal{L}(\alpha)$ and

$$n' \gg_s \left\lfloor \left(\frac{\exp(-\mathcal{L}(\alpha)^2)}{m \log K}\right)^M n \right\rfloor,$$

and a cofinite ρ' of rank at most $O_s(\alpha^{-1/(s-2)}\mathcal{L}(\alpha))$ such that

$$\rho' \gg \left(\frac{\exp(-\mathcal{L}(\alpha)^2)}{\log K}\right)^M$$

so that

$$\left|(n'X)(\rho' \circ \{1, c_1, \dots, c_s\}^M)\right|^{-1} \gg_s \alpha^2.$$

Finally, the lower bound on the size of Bohr sets provided by Lemma 2.12 leads to the desired lower bound for α. It only remains to introduce one final definition to convert the awkward width function $\rho \circ \{1, \dots, c_s\}^M$ into something more amenable. Let $F_d : \widehat{G} \to \mathcal{P}(\widehat{G})$ be such that $|F_d(\gamma)| \leq d$ for all $\gamma \in \widehat{G}$ and, furthermore, $F_d(F_d(\Gamma)) = F_d(\Gamma)$ for all finite $\Gamma \subset \widehat{G}$. We say that (F_d, m) *dominates* $A \subset R$ if, for any cofinite $\rho : \widehat{G} \to [0,1]$, we have, for all $\gamma \in \widehat{G}$ and $x \in G$,

$$\text{if } |1 - \gamma'(x)| \leq \rho(\gamma)/m \text{ for all } \gamma' \in F_d(\gamma) \text{ then } |1 - \gamma(x)| \leq 2(\rho \circ A)(\gamma).$$

In our applications, the function F_d will in fact depend only on G. The point is that we will need to be able to control the rank and size of width functions of the shape $\rho \circ A^M$ for some A depending on the coefficients $\mathbf{c}$, which the parameters F_d and m offer. We admit that this definiton is a little obtuse, and refer the reader to the applications to follow for clarification.

Theorem 2.18. *Let $X \subset G$ be a finite symmetric set with additive control $K \geq 4$ and $n \geq 4$, and let $A \subset nX$ with density α. Let $\mathbf{c} = (c_1, \dots, c_s) \in (R^*)^s$ be a commutative s-tuple which is m-constrained by X such that $c_1 + \cdots + c_s = 0$. Finally, suppose that (F_d, m') dominates $\{1, c_1, \dots, c_s\}$.*

If A contains only trivial solutions to (2.1) then there exists an integer

$$n' \gg_s \lfloor \exp\left(-O_s\left(\mathcal{L}(\alpha)\left(\mathcal{L}(\alpha)^2 + \log\log K + \log m\right)\right)\right) n \rfloor$$

such that

$$\exp\left(O_s\left(d\alpha^{-1/(s-2)}\mathcal{L}(\alpha)^2\left(\mathcal{L}(\alpha)^2 + \log\log K + \log m'\right)\right)\right) \gg_s |n'X|. \tag{2.4}$$

Proof. By Theorem 2.17 there exist integers $0 \leq M \ll_s \mathcal{L}(\alpha)$ and

$$n' \gg_s \lfloor \exp\left(-\mathcal{L}(\alpha)O_s\left(\mathcal{L}(\alpha)^2 + \log\log K + \log m\right)\right) n \rfloor$$

and a cofinite ρ' of rank at most $O_s(\alpha^{-1/(s-2)}\mathcal{L}(\alpha))$ such that

$$\rho' \gg \exp(-\mathcal{L}(\alpha)O_s(\mathcal{L}(\alpha)^2 + \log\log K))$$

and

$$\left|(n'X)(\rho' \circ \{1, \dots, c_s\}^M)\right|^{-1} \gg_s \alpha^2.$$

We now use the hypothesis that (F_d, m') dominates $\{1, \dots, c_s\}$. If we let ρ'' be the width function defined by

$$\rho''(\gamma) = \begin{cases} (m')^{-M} \inf_{\gamma' \in F_d^{-1}(\gamma)} \rho'(\gamma') & \text{where } F_d^{-1}(\gamma) = \{\gamma' \in \widehat{G} : \gamma \in F_d(\gamma')\} \neq \emptyset, \text{ and} \\ 1 & \text{otherwise} \end{cases}$$

then it follows from the definitions that $\mathrm{rk}(\rho'') \leq d\mathrm{rk}(\rho')$,

$$\rho'' \gg \exp(-\mathcal{L}(\alpha) O_s(\mathcal{L}(\alpha)^2 + \log\log K + \log m')),$$

and that for any Y we have $Y(\rho'') \subset Y(\rho' \circ \{1, \dots, c_s\}^M)$. In particular, by Lemma 2.12 it follows that

$$\exp\left(O_s\left(d\alpha^{-1/(s-2)}\mathcal{L}(\alpha)^2\left(\mathcal{L}(\alpha)^2 + \log\log K + \log m'\right)\right)\right) \gg_s |n'X|$$

as required. □

2.3.1 The integers

The most important application of Theorem 2.18 is to the classical setting of Roth's problem; namely, where $G = \mathbb{Z}$ (we recall that the endomorphism ring of $\mathbb{Z}$ is identical to $\mathbb{Z}$). On recalling the definitions it is clear from the triangle inequality that, if $\mathbf{c} \in (\mathbb{Z}\backslash\{0\})^s$ with $\ell = \max_{1 \leq i \leq s} |c_i|$, then $\mathbf{c}$ is ℓ-constrained by any finite symmetric $X \subset \mathbb{Z}$. Furthermore, by the triangle inequality (ι_1, ℓ) dominates $\{1, c_1, \dots, c_s\}$, where ι_1 is the identity function. Thus the technical burdens of the previous section are immediately lightened.

It remains to choose a suitable ground set $X \subset \mathbb{Z}$ which has bounded additive control. The obvious choice is $X = \{-1, 0, 1\}$ which has additive control of 2. As discussed in the previous section, for the integers this is, in some sense, the only reasonable choice to make. With this choice the bound (2.4) thus becomes

$$\exp\left(O_s\left(\alpha^{-1/(s-2)}\mathcal{L}(\alpha)^2\left(\mathcal{L}(\alpha)^2 + \log \ell\right)\right)\right) \gg_s n.$$

After some simple algebra the following theorem is an immediate corollary of Theorem 2.18.

Theorem 2.19. *Let $A \subset \{1, \dots, N\}$ with density α and $\mathbf{c} = (\mathbb{Z}\backslash\{0\})^s$ be such that $c_1 + \cdots + c_s = 0$. Let $\ell = \max_{1 \leq i \leq s} |c_i|$. If A contains only trivial solutions to (2.1) then*

$$\alpha \ll_s \left(\frac{(\log\log N)^2\left((\log\log N)^2 + \log \ell\right)}{\log N}\right)^{s-2},$$

where the implied constant depends only on s.

In particular, taking $\mathbf{c} = (1, 1, -2)$ we obtain a new quantitative bound for Roth's theorem on three term arithmetic progressions: if $A \subset \{1, \dots, N\}$ with density α contains only trivial three term arithmetic progressions then

$$\alpha \ll \frac{(\log\log N)^4}{\log N}.$$

It is interesting to observe the dependence on ℓ in this result. For example, we can obtain the following result, which generalises our bound for three term arithmetic progressions to a host of other equations simultaneously.

Corollary 2.20. *If $A \subset \{1, \dots, N\}$ has density*

$$\alpha \gg \frac{(\log\log N)^4}{\log N}$$

then A contains non-trivial solutions to the equation

$$sx_1 + (t - s)x_2 = tx_3$$

for every $1 \leq s \leq t \leq \exp((\log\log N)^2)$.

2.3.2 Multi-dimensional integers

A benefit of developing our theory in such general terms is that it is straightforward to obtain multi-dimensional generalisations. In this section we consider the problem where $G = \mathbb{Z}^d$; we recall that $\widehat{G} = \mathbb{T}^d$. Analogous to the previous section, a suitable ground set is given by $X = \{-1, 0, 1\}^d$, which has additive control $K = 2^d$. If we fix some $\mathbf{c} \in \mathrm{GL}_d(\mathbb{Z})^s$ and $\ell = \max_{1\leq i\leq s} \max_{1\leq j,k\leq d} |c_i(jk)|$ then it follows from the triangle inequality that $\mathbf{c}$ is $d\ell$-constrained over X.

Furthermore, we will define F_{d^2} to be the function which takes the character $\gamma = (\gamma_1, \dots, \gamma_d)$ to the set $\{e_{ij}\gamma\}_{1\leq i,j\leq d}$, where e_{ij} is the matrix which has 1 in the (i, j)th entry and 0 everywhere else. It follows from the triangle inequality that $(F_{d^2}, d\ell)$ dominates $\{1, c_1, \dots, c_s\}$.

As in the previous section, we thus have the following immediate corollary of Theorem 2.18.

Theorem 2.21. *Let $A \subset \{1,\dots,N\}^d$ with density α and $\mathbf{c} = (c_1,\dots,c_s) \in \mathrm{GL}_d(\mathbb{Z})^s$ be such that $c_1 + \cdots + c_s = 0$ and $c_i c_j = c_j c_i$ for $1 \le i \le j \le s$. If A contains only trivial solutions to* (2.1) *then*

$$\alpha \ll_s \left(d \frac{(\log\log N)^2 \left((\log\log N)^2 + \log d + \log \ell\right)}{\log N} \right)^{s-2},$$

where the implied constant depends only on s.

In some cases this gives a quantitative improvement on the work of Prendiville [48], which established a bound of $\alpha \ll_{d,\ell} 1/\log\log N$ in the case $s = 3$, though Prendiville's work is more general in that it is able to handle one of the coefficients having rank less than d. We mention the following example from [48] to which our Theorem 2.21 does apply: finding right-angled isosceles triangles in $A \subset \{1,\dots,N\}^2$, which are solutions to the equation

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} x_1 + \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} x_2 + \begin{pmatrix} -1 & 1 \\ -1 & -1 \end{pmatrix} x_3 = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

A similar problem to which our theory does not apply is the problem of sets without 'corners', right-angled triangles with sides parallel to the axes, for which quantitative results have been achieved by Shkredov [67]. We recall from the introduction that this is equivalent to studying solutions of the equation

$$\begin{pmatrix} -1 & 1 & 0 \\ -1 & -1 & 0 \\ -2 & 0 & 2 \end{pmatrix} \mathbf{x}_1 + \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 0 & -1 \end{pmatrix} \mathbf{x}_2 + \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & -1 \end{pmatrix} \mathbf{x}_3 = \begin{pmatrix} 0 \\ 0 \end{pmatrix},$$

with variables $\mathbf{x}_i \in A \subset \{-N,\dots,N\}^3 \cap D$, where D is the plane $\{(x_1,x_2,x_3) : x_2 = x_3\}$. We observe that

$$\{-N,\dots,N\}^3 \cap D = N\left(\{-1,0,1\}^3 \cap D\right) = NX,$$

say. The set X certainly has bounded additive control. The problem is the hypothesis of constraint; we recall that for our theory to be applicable we required there to be some m such that $c_i \cdot X \subset mX$ for $1 \le i \le 3$. This condition fails completely for this problem, however, since the plane D is not closed under dilation by the coefficients c_i. It is for precisely the same reason that our methods cannot handle k-term arithmetic progressions for $k > 3$.

CHAPTER 3

AN IMPROVED DENSITY INCREMENT ARGUMENT

In this chapter we prove Theorem 2.3, which provides a lower bound for the density of sets without solutions to a given translation invariant equation, which we applied in the previous chapter. In broad strokes our proof follows the traditional density increment strategy first used in the original paper of Roth [50]. It incorporates improvements to this original method by Szemerédi [73], Heath-Brown [33], Bourgain [6, 9] and Sanders [60, 61]. The main new ingredient of the approach in this chapter, and the primary reason for the quantitative improvements on previous work, is a new structural result on the set of large Fourier coefficients inspired by related results of Bateman and Katz [1] and Shkredov [69].

The method presented here has appeared, in the case of finite groups, in [4]. An alternative method, which generalised the work of Sanders [60] on the traditional case of Roth's theorem to more general translation invariant equations, appeared earlier in [3]; this alternative method is both more technically challenging and delivers weaker bounds, but we include a proof of the main lemma involved as it includes a generalisation of the combinatorial methods of [60] which may be useful for other problems.

3.1 A heuristic discussion

We begin with a heuristic outline of the proof. We recall that we are working within an abelian group G with endomorphism ring R; for simplicity we will restrict ourselves in this discussion to the case when G is finite, our equation has only three variables, and R is a commutative integral domain. In particular, if $|G| = N$ then the Haar probability measure on $\widehat{G}$ is the measure assigning the weight $1/N$ to each $\gamma \in \widehat{G}$. We will simplify matters still further by not keeping track of some logarithmic factors; to this end we write $X \lesssim_\alpha Y$ when $X \ll \mathcal{L}(\alpha)^C Y$ for some absolute constant C, and similarly for $\tilde{O}(\cdot)$ and $\tilde{\Omega}(\cdot)$.

The problem is this: having fixed some $c_1, c_2, c_3 \in R$ such that $c_1 + c_2 + c_3 = 0$, and

some finite set $X \subset G$ with a certain amount of structure, what is the size of the largest subset of X which contains only trivial solutions to the equation

$$c_1x_1 + c_2x_2 + c_3x_3 = 0, \tag{3.1}$$

where a trivial solution is one which satisfies $x_1 = x_2 = x_3$? The natural approach is to find some lower bound for the inner product

$$\Upsilon_{\mathbf{c}}(A) = \langle (c_1 \cdot A) * (c_2 \cdot A), (-c_3 \cdot A) \rangle, \tag{3.2}$$

which counts the number of solutions (including trivial ones) to (3.1). The lower bound will depend only on the size of A; thus, suppose we have the lower bound $\Upsilon_{\mathbf{c}}(A) \geq L(|A|)$. If A contains only trivial solutions then $L(|A|) \leq \Upsilon_{\mathbf{c}}(A) \leq |A|$, which after rearranging will give an upper bound on $|A|$. We find a lower bound for (3.2) using the traditional density increment strategy, which is an iterative argument due originally to Roth [50].

We first observe that, if $A \subset X$ is a random set with density α and X is roughly closed under addition and dilation by the coefficients c_i, then the expected value of $\Upsilon_{\mathbf{c}}(A)$ is $\alpha |A|^2$; indeed, having fixed any $x_1, x_2 \in A$ the probability that the value of x_3 which satisfies (3.1) belongs to A is α. It follows that if $\Upsilon_{\mathbf{c}}(A)$ is much smaller than this expected value then A does not behave like a random set, and hence should *not* be evenly distributed within X. In particular, if we partition X into a small number of structured parts $X_1 \sqcup \cdots \sqcup X_k$ then A is relatively concentrated in some part of this partition. By iterating this argument we may increase the degree of concentration until it exceeds some trivial upper bound, and so we must halt with a lower bound for $\Upsilon_{\mathbf{c}}(A)$.

As usual with linear problems, we measure how 'structured' the set A is by its Fourier coefficients. Thus the iterative step of the density increment strategy breaks down into two stages:

1. show that if (3.2) is smaller than the expected value of $\alpha |A|^2$ then A (or some dilate of A) has many large non-trivial Fourier coefficients, and

2. show that if A has many non-trivial large Fourier coefficients then it has greater than expected density on some large structured subset $X' \subset X$; that is, for some $\nu > 0$ we have
$$\alpha' = \frac{|A \cap X'|}{|X'|} \geq (1 + \nu)\alpha'.$$

Given such information, it is clear how to repeatedly apply this argument to yield a lower bound for $\Upsilon_{\mathbf{c}}(A)$. Since $A \cap X'$ also contains no non-trivial solutions to (3.1) we can repeat this argument for $A \cap X' \subset X'$, and so on, until we arrive at some set $A_0 \subset A$ and some structured set $X_0 \subset X$ such that $A_0 \subset X_0$ has at least the expected number of solutions to (3.1), so that $\Upsilon_{\mathbf{c}}(A) \geq \Upsilon_{\mathbf{c}}(A_0) \geq \alpha_0 |A_0|^2 \geq \alpha |A|^2 (|X_0| |X|^{-1})^2$. Provided $|X_0|$ is not too small relative to $|X|$ this gives us a non-trivial lower bound for $\Upsilon_{\mathbf{c}}(A)$.

There are two fundamental parameters which control how small X_0 is: how many times we may have to run the density increment argument until we have the desired lower bound on Υ, and how small X' is in relation to X at each stage. The former is controlled by ν, for after K iterations of the argument we have density $\alpha_K \geq (1+\nu)^K \alpha$, and hence the trivial bound $\alpha_K \leq 1$ forces the argument to halt after at most $K \lesssim_\alpha \nu^{-1}$ steps.

The latter is more delicate to control, for how large X' can be taken depends not only on the size of X but also on its finer structural properties. We recall that X needs to be roughly closed under addition and dilation by the coefficients c_i; for the purposes of this discussion we shall capture this property by supposing that there is some constant $K(\mathbf{c}) \geq 2$ such that $|c_1 \cdot X + c_2 \cdot X + c_3 \cdot X| \leq K^{d_X} |X|$ for some integer $d_X \geq 1$, which we shall refer to as the dimension of X. In our argument X' will be a subset of X which is controlled on some set $\Gamma \subset \widehat{G}$ (i.e. $\gamma(x) \approx 1$ for all $x \in X'$ and $\gamma \in \Gamma$), and subject to some additional structural requirements which we will ignore for this discussion. If $|\Gamma| = r$ then we are able to obtain the lower bound $|X'| \geq \exp(-O(d_X + r)) |X|$ and $d_{X'} \leq d_X + O(r)$. It thus suffices to control the parameter r, which we will call the rank increment.

Suppose, then, that we can prove the argument above with density increment ν and rank increment r. We begin with some $A \subset G$ with density α. Iterating the argument above leads to a sequence of sets $G \supset X_1 \supset \cdots \supset X_K$ for some $K \lesssim_\alpha \nu^{-1}$, with $d_{X_i} \ll ir$,

$$|X_K| \geq \exp(-O(d_{X_1} + \cdots + d_{X_K} + Kr))N \geq \exp(-O(K^2 r))N,$$

and

$$\Upsilon_{\mathbf{c}}(A) \geq \alpha^3 |X_K|^2 \geq \alpha^3 \exp(-O(K^2 r))N^2.$$

In particular, if A contains only trivial solutions to (3.1) then $\Upsilon_{\mathbf{c}}(A) \ll \alpha N$, and hence

$$r\nu^{-2} \gtrsim_\alpha \log N,$$

which leads to an upper bound on α, the quality of which depends on how the parameters ν and r depend on the density α.

Thus, for example, when $G = \mathbb{Z}/N\mathbb{Z}$ the method of Bourgain [6] gives a density increment of $\nu \gg \alpha$ and a rank increase of $r \ll 1$, which implies that

$$\alpha \ll_\epsilon \frac{1}{(\log N)^{1/2-\epsilon}} \text{ for all } \epsilon > 0.$$

On the other hand, the method of Sanders [60] achieves a density increment of $\nu \gg 1$ and a rank increase of $r \lesssim_\alpha \alpha^{-1}$, which implies that

$$\alpha \ll_\epsilon \frac{1}{(\log N)^{1-\epsilon}} \text{ for all } \epsilon > 0.$$

We now discuss our new method, which also achieves a density increment of $\nu \gg 1$ and a rank increase of $r \lesssim_\alpha \alpha^{-1}$. This is done in a more direct fashion than the method of Sanders, which leads to some marginal quantitative improvements in the hidden factors of $\mathcal{L}(\alpha)$.

For simplicity, we let $A \subset G$ be a set of density α. It is, of course, crucial that our argument below is equally valid when G is replaced with some suitably structured $X \subset G$, and in particular it is important that the sets X produced in each stage of the density increment argument are structured enough for a 'local' Fourier analysis to hold.

We first observe that by taking the Fourier transform we have

$$\begin{aligned}\Upsilon_{\mathbf{c}}(A) &= \frac{1}{N}\sum_{\gamma\in\widehat{G}} \widehat{(c_1\cdot A)}(\gamma)\widehat{(c_2\cdot A)}(\gamma)\widehat{(-c_3\cdot A)}(\gamma)\\ &= \alpha\,|A|^2 + \frac{1}{N}\sum_{\gamma\neq 0} \widehat{(c_1\cdot A)}(\gamma)\widehat{(c_2\cdot A)}(\gamma)\widehat{(-c_3\cdot A)}(\gamma).\end{aligned}$$

In particular, if $\Upsilon_{\mathbf{c}}(A)$ is much smaller than $\alpha\,|A|^2$ then

$$\left|\frac{1}{N}\sum_{\gamma\neq 0} \widehat{(c_1\cdot A)}(\gamma)\widehat{(c_2\cdot A)}(\gamma)\widehat{(-c_3\cdot A)}(\gamma)\right| \gg \alpha\,|A|^2\,.$$

By Hölder's inequality there exists some $1 \leq i \leq 3$ such that

$$\frac{1}{N}\sum_{\gamma\neq 0}\left|\widehat{(c_i\cdot A)}(\gamma)\right|^3 \gg \alpha\,|A|^2\,.$$

For convenience, we shall suppose henceforth that $c_i = 1$; the passage to more general coefficients is a purely technical difficulty, as we can pass from A to arbitrary dilations of A without altering the fundamental property of containing only trivial solutions to (3.1).

We now observe that many characters in $\widehat{G}$ can be discarded without affecting this lower bound; namely, if $\Delta_\eta(A) = \{\gamma \in \widehat{G} : \left|\widehat{A}(\gamma)\right| \geq \eta |A|\}$ then, by Parseval's identity,

$$\frac{1}{N} \sum_{\gamma \notin \Delta_\eta(A)} \left|\widehat{A}(\gamma)\right|^3 \leq \eta |A| \frac{1}{N} \int \sum_\gamma \left|\widehat{A}(\gamma)\right|^2 = \eta |A|^2 .$$

In particular, there exists some absolute constant $c > 0$ such that

$$\frac{1}{N} \sum_{\gamma \in \Delta_{c\alpha(A)} \setminus \{0\}} \left|\widehat{A}(\gamma)\right|^3 \gg \alpha |A|^2 . \tag{3.3}$$

Using the trivial upper bound $\left|\widehat{A}(\gamma)\right| \leq |A|$ leads to the lower bound

$$\frac{1}{N} \sum_{\gamma \in \Delta_{c\alpha(A)} \setminus \{0\}} \left|\widehat{A}(\gamma)\right|^2 \gg \alpha |A| . \tag{3.4}$$

More generally, we say that A has correlation of strength τ with Δ if

$$\frac{1}{N} \sum_{\gamma \in \Delta \setminus \{0\}} \left|\widehat{A}(\gamma)\right|^2 \gg \tau |A| .$$

To see how correlation can be converted into density increment, we first observe that if a finite set $X \subset G$ is controlled on γ then $\left|\widehat{X}(\gamma)\right| \gg |X|$. In particular, if A has correlation of strength τ with Δ then by Parseval's identity, since $\widehat{(A - \alpha)}(\gamma) = \widehat{A}(\gamma)$ if $\gamma \neq 0$,

$$\|(A - \alpha) * X\|_2^2 = \frac{1}{N} \sum_{\gamma \neq 0} \left|\widehat{A}(\gamma)\right|^2 \left|\widehat{X}(\gamma)\right|^2 \gg \tau |A| |X|^2 .$$

Assuming X has a reasonable amount of additive structure the left hand side is approximately $\|A * X\|_2^2 - \alpha |A| |X|^2$, and hence by the trivial bound $\|f\|_2^2 \leq \|f\|_\infty \|f\|_1$ we have

$$\|A * X\|_\infty \|A * X\|_1 \gg |A| |X|^2 (\alpha + \tau) .$$

Since $\|A * X\|_1 = |A| |X|$ we have thus shown that correlation of strength τ with Δ implies a density increment of $\tau\alpha^{-1}$ on any set X which is controlled on Δ and has a reasonable amount of additive structure.

We showed above that if A has only trivial solutions to (3.1) then it has correlation of strength $\Omega(\alpha)$ with $\Delta_{c\alpha}(A)$. By Parseval's identity we have the immediate upper bound $|\Delta_\eta(A)| \ll \eta^{-2}\alpha^{-1}$, and so our discussion thus far leads to a density increment of $\nu \gg 1$ and rank increment of $r \ll \alpha^{-3}$.

We can do better, however, and start with the simple observation that if there is a set Λ such that

$$\Delta \subset \langle \Lambda \rangle = \left\{ \sum_{\lambda \in \Lambda} \epsilon_\lambda \lambda : \epsilon_\lambda \in \{-1, 0, 1\} \right\}, \tag{3.5}$$

then control on Λ implies control on Δ. In particular, we say that Δ is d-covered if (3.5) is possible with some Λ of size at most d. Thus, by the above, if we have correlation of strength $\Omega(\tau)$ with a d-covered set Δ then we can prove a density increment of $\nu \gg 1$ and a rank increment of $r \ll d$.

In the course of improving the quantitative bounds for the inverse sumset problem Chang [13] showed that $\Delta_\eta(A)$ is $\tilde{O}_\alpha(\eta^{-2})$-covered. This fact, combined with the above, yields a density increment of $\nu \gg 1$ and a rank increment of $r \ll \alpha^{-2}$. Chang's lemma is quantitatively the best possible result, as shown by Green [27], and hence an alternative approach is needed for any improvement.

We recall that in passing from (3.3) to (3.4) we used the trivial bound $\left|\widehat{A}(\gamma)\right| \leq |A|$. Instead, we now use something a little more refined. We can partition $\Delta_{c\alpha}(A)$ into at most $\tilde{O}_\alpha(1)$ sets of the shape

$$\tilde{\Delta}_\eta(A) = \{\gamma \in \widehat{G} : \left|\widehat{A}(\gamma)\right| \approx \eta\,|A|\}$$

for $\eta \gg \alpha$. It follows from (3.3) by the pigeonhole principle that there exists some $\eta \gg \alpha$ such that

$$\frac{1}{N} \sum_{\gamma \in \tilde{\Delta}_\eta(A)} \left|\widehat{A}(\gamma)\right|^3 \gtrsim_\alpha \alpha\,|A|^2\,,$$

and hence

$$|\Delta_\eta(A)| \geq \left|\tilde{\Delta}_\eta(A)\right| \gtrsim_\alpha \eta^{-3}.$$

This should be compared to the trivial bound $|\Delta_\eta(A)| \ll \eta^{-2}\alpha^{-1}$. In particular, let $\Delta' \subset \Delta_\eta(A)$ be any set such that $|\Delta'| \gg \eta\,|\Delta_\eta(A)|$. We have

$$\frac{1}{N} \sum_{\gamma \in \Delta' \setminus \{0\}} \left|\widehat{A}(\gamma)\right|^2 \gg \eta^2 \alpha\,|A|\,|\Delta'| \gtrsim_\alpha \alpha\,|A|\,.$$

Thus we have shown that there exists $\eta \gg \alpha$ such that A has correlation of strength $\tilde{\Omega}_\alpha(\alpha)$ with *any* subset $\Delta' \subset \Delta_\eta(A)$ such that $|\Delta'| \gg \eta\,|\Delta_\eta(A)|$.

In our new method we exploit this fact using a new structural result which will act as a quantitatively superior substitute for the previously mentioned lemma of Chang. This was inspired by similar structural results in the recent progress made by Bateman and Katz [1] for Roth's theorem over $\mathbb{F}_3^n$. We recall that Chang's lemma says that $\Delta_\eta(A)$ is $\tilde{O}_\alpha(\eta^{-2})$-covered. Our new structural lemma says that there exists a set $\Delta' \subset \Delta_\eta(A)$ such that $|\Delta'| \gg \eta\,|\Delta_\eta(A)|$ which is $\tilde{O}_\alpha(\eta^{-1})$-covered. Thus, if we are willing to pass to an η-dense subset of the spectrum we can achieve a covering which is more efficient by a factor of η than that given by Chang's lemma.

As shown above, in the present argument we can pass to an η-dense subset while still preserving the necessary amount of correlation, and so our new structural lemma can be applied. Thus our argument leads to the conclusion that A has correlation of strength $\tilde{\Omega}_\alpha(\alpha)$ with some set Δ' which is $\tilde{O}_\alpha(\alpha^{-1})$-covered, giving a density increment of $\nu \gtrsim_\alpha 1$ and rank increment of $r \lesssim_\alpha \alpha^{-1}$, as required.

This concludes our heuristic discussion of how to obtain our improved quantitative bounds for Roth type problems; of course, there is a fair amount of technical difficulty in making these ideas rigorous: we must work over a general abelian group, precisely define what kind of 'structure' is required for local Fourier analysis, keep track of the factors of $\mathcal{L}(\alpha)$, and so on.

3.2 Creating correlation

In this section we show how a bound for a convolution in physical space can be converted to more useful Fourier information, which we shall use in the next section to generate a density increment suitable for an iterative argument. The key observation is that, by taking the Fourier transform, we have the identity

$$\langle f_1 * \cdots * f_{s-1}, f_s \rangle = \int \widehat{f_1}(\gamma) \cdots \overline{\widehat{f_s}(\gamma)}\,\mathrm{d}\gamma.$$

In particular, a lower bound for the left hand side leads, via the triangle inequality and Hölder's inequality, to a lower bound on $\int \left|\widehat{f}_i(\gamma)\right|^s \mathrm{d}\gamma$ for some $1 \leq i \leq s$. Provided $s \geq 3$ we can exploit the cancellation inherent in Parseval's identity to bound the contribution to this integral from those γ such that $\left|\widehat{f}_i(\gamma)\right|$ is small. In particular, if Δ is the set of

large Fourier coefficients of $\widehat{f}_i$ (for a suitable notion of large) then this will imply that $\int_\Delta \left|\widehat{f}_i(\gamma)\right|^s \, \mathrm{d}\gamma$ is large.

Rather than speak directly of the set of large Fourier coefficients in this section we will use the more flexible notion of correlation, which replaces the characteristic function of the set Δ with the Fourier transform of an arbitrary $g \in L^1(G)$. For our application we will be concerned with functions f that are balanced functions of sets; that is, $f = A - \alpha B$ for some $A \subset B$ with density α. This is particularly useful because a suitably strong upper bound for the convolution of sets leads to a lower bound for the convolution of their balanced functions.

Let B be a finite subset of G with some subset $A \subset B$ with density α, and let $g \in L^1(G)$. We say that A in B has correlation of strength τ, or τ-correlates, with g if, letting $\mathbf{A} = A - \alpha B$,

$$\int \left|\widehat{\mathbf{A}}(\gamma)\right|^2 |\widehat{g}(\gamma)| \, \mathrm{d}\gamma \geq \tau \, |A| \, \|g\|_1 .$$

By Parseval's identity the left hand side is at most $\|\mathbf{A}\|_2^2 \, \|g\|_1 = (1-\alpha) \, |A| \, \|g\|_1$, and hence we may always assume that $\tau \in [0, 1-\alpha]$. The importance of correlation in the study of translation invariant equations can be traced to the work of Szemerédi [73] and Heath-Brown [33], where the value of working with the spectrum as a whole, rather than a single large Fourier coefficient, was first recognised.

We will prove two lemmas which convert an upper bound on the convolution of sets to some non-trivial correlation of one of the sets. The first, which is all that will be required to prove Theorem 2.3, is an elementary combinatorial argument. We present a simple proof which allows us to create correlation from only a bare minimum of structural hypotheses. We recall that X' is δ-sheltered by X if $|(X + X')\backslash X| \leq \delta \, |X|$.

Lemma 3.1. *Let B be a finite symmetric subset of G. Let A_1 and A_s be subsets of B with relative densities α_1 and α_s respectively, and let $A_2, \ldots, A_{s-1}$ be any finite subsets of G such that $A_2 + \cdots + A_{s-1}$ is $2^{-2}\alpha$-sheltered by B, where $\alpha = \min(\alpha_1, \alpha_s)$. Then either*

$$\langle A_1 * A_2 * \cdots * A_{s-1}, A_s \rangle \geq 2^{-2} \, |A_1| \cdots |A_s| \, |B|^{-1} , \tag{3.6}$$

or for some $i \in \{1, s\}$ the set A_i in B has correlation of strength $2^{-2}\alpha_i$ with $A_j^{(s-2)}$ for some $1 < j < s$.

Proof. Let $\mathbf{A}_i = A_i - \alpha_i B$ for $i \in \{1, s\}$, and let $f = A_2 * \cdots * A_{s-1}$. By linearity of the inner product

$$\langle \mathbf{A}_1 * f, \mathbf{A}_s \rangle = \langle A_1 * f, A_s \rangle - \alpha_1 \langle B * f, A_s \rangle - \alpha_s \langle A_1 * f, B \rangle + \alpha_1 \alpha_s \langle B * f, B \rangle. \quad (3.7)$$

We observe that f is supported on $A_2 + \cdots + A_{s-1}$, and hence

$$\begin{aligned} |\langle A_1 * f, B \rangle - \|f\|_1 |A_1|| &= \sum_{x \notin B} A_1 * f(x) \\ &\leq \|f\|_1 |(A_2 + \cdots + A_{s-1} + B) \backslash B| \\ &\leq 2^{-2} \|f\|_1 |A_1|, \end{aligned}$$

since $A_2 + \cdots + A_{s-1}$ is $2^{-2}\alpha_1$-sheltered by B. In particular, $\langle A_1 * f, B \rangle \geq (1 - 2^{-2}) |A_1| \|f\|_1$. By a similar argument $\langle B * f, A_s \rangle \geq (1 - 2^{-2}) |A_s| \|f\|_1$. Furthermore, we have the trivial bound $\langle B * f, B \rangle \leq \|f\|_1 |B|$. Combining these bounds with (3.7) we obtain the upper bound

$$\langle \mathbf{A}_1 * f, \mathbf{A}_s \rangle \leq \langle A_1 * f, A_s \rangle - 2^{-1} \alpha_1 \alpha_s \|f\|_1 |B|.$$

In particular, either (3.6) is true or

$$\langle \mathbf{A}_1 * f, \mathbf{A}_s \rangle \leq -2^{-2} \alpha_1 \alpha_s \|f\|_1 |B|.$$

Taking the Fourier transform of the left hand side and applying the triangle inequality implies that

$$\int \left| \widehat{f}(\gamma) \right| \left| \widehat{\mathbf{A}_1}(\gamma) \right| \left| \widehat{\mathbf{A}_s}(\gamma) \right| \, \mathrm{d}\gamma \geq 2^{-2} \alpha_1 \alpha_s \|f\|_1 |B|.$$

Applying the Cauchy-Schwarz inequality followed by Hölder's inequality it follows that for some $i \in \{1, s\}$ and some $2 \leq j \leq s - 1$ we have

$$\int \left| \widehat{A_j}(\gamma) \right|^{s-2} \left| \widehat{\mathbf{A}_i}(\gamma) \right|^2 \, \mathrm{d}\gamma \geq 2^{-2} |A_j|^{s-2} \alpha_i^2 |B|,$$

and the lemma follows from the definition of correlation. □

Lemma 3.1 is the most direct way to convert an upper bound of an inner product of a convolution to some non-trivial correlation, and is all that will be required for the proof of Theorem 2.3. Using a new probabilistic technique due to Croot and Sisask [17], however, Sanders [60] recently found an ingenious alternative argument, which is quantitatively stronger for some applications. We generalised the argument of [60] to function fields and $s > 3$ variables in [3]. In this thesis we will generalise the techniques further still and prove the following, more involved, correlation lemma.

Lemma 3.2. *Let B be a finite symmetric subset of G. Let A_1 and A_s be subsets of B with relative densities α_1 and α_s respectively, and let $A_2, \dots, A_{s-1}$ be any finite subsets of G. Let $\alpha = \min(\alpha_1, \alpha_s)$.*

Furthermore, suppose that there are finite symmetric sets $B_2', \dots, B_{s-1}', B_2'', \dots, B_{s-1}'', B'''$ such that for each $2 \le i \le s-1$

1. *$A_i \subset B_i'$ with density at least α',*
2. *$2\ell B''' + B_2'' + \cdots + B_{s-1}''$ is $2^{-1}\alpha$-sheltered by B,*
3. *B_i' is $2^{-2s}\alpha$-sheltered by B,*
4. *B_i'' is $2^{-2}\alpha'$-sheltered by B_i', and*
5. *$\left|B_{s-1}'' + B'''\right| \le 2\left|B_{s-1}''\right|$.*

Then either

1. *$\langle A_1 * \cdots * A_{s-1}, A_s\rangle \ge \exp\left(-2^5 s\alpha_1^{-1/(s-2)}\mathcal{L}(\alpha')\right)\prod_{i=2}^{s-1}|B_i''|\,|A_s|$, or*
2. *there is a finite set D with $\beta'(D) \ge 2^{-1}\alpha'$ such that A_1 in B has correlation of strength $2^{-2s}\alpha_1$ with $D * D$, or*
3. *for any integer $\ell \ge 1$ there is a finite set D with*

$$\beta'''(D) \ge \exp\left(-2^{26+2s}\ell^2\alpha_1^{-1/(s-2)}\mathcal{L}(\alpha')\mathcal{L}(\alpha_s)\right)$$

such that A_s in B has correlation of strength $2^{-4-s}\alpha_s$ with $D^{(\ell)}$.

It is immediately apparent that Lemma 3.2 is more complex, both in the hypotheses required and the conclusion delivered, than the simple Lemma 3.1. Most of the differences are largely cosmetic, however, with regards to the quantitative bounds delivered for the problems considered in Chapter 2. The important difference lies in the objects with which correlation of strength $\approx \alpha$ is obtained. When using the customary structural lemma of Chang, Lemma 3.2 is more effective; it is unnecessary for our purposes, however, as we prove a new structural lemma for which Lemma 3.1 is sufficient.

We illustrate this by considering, for simplicity, the case $s = 3$. The structural lemma of Chang (which we will state explicitly in the following section) can be used to convert

correlation of strength $\Omega(\alpha)$ with $D^{(r)}$, where D is contained inside some structured set with density δ, into a density increment of strength $\Omega(1)$ on a set with relative density $|X'|/|X| \approx \exp(-\alpha^{-2/r}\mathcal{L}(\delta))$.

Lemma 3.1 delivers correlation with a set D of density $\Omega(\alpha)$, and hence using Chang's lemma would deliver a density increment on a structured subset with relative density $\exp(-\alpha^{-2}\mathcal{L}(\alpha))$. Lemma 3.2, however, delivers correlation either with $D^{(2)}$, where D is a set of density $\Omega(\alpha)$, or with $D^{(\ell)}$ where D is a set of density

$$\delta \gg \exp(-O(\ell^2\alpha^{-1}\mathcal{L}(\alpha)^2))$$

and $\ell \geq 1$ is any integer. In particular, using Chang's lemma this yields a density increment on a structured subset with relative density either $\exp(-\alpha^{-1}\mathcal{L}(\alpha))$ or

$$\exp(-\alpha^{-1-2/\ell}\ell^2\mathcal{L}(\alpha)^2) \approx \exp(-\alpha^{-1}\mathcal{L}(\alpha)^4),$$

if we choose $\ell \approx \mathcal{L}(\alpha)$. Thus, when using Chang's lemma, Lemma 3.2 is superior to Lemma 3.1 by a factor of almost α^{-1}. This improvement led to the quantitative improvements of Sanders [60], which were generalised in [3].

We include a proof of Lemma 3.2 at the end of this chapter for this reason, and because the tools required offer an alternative point of view. Before that, however, we prove a new structural lemma which is an alternative to Chang's lemma, and which means that the much simpler Lemma 3.1 will suffice. This is particularly beneficial in that fewer technical hypotheses are required, which makes the task of constructing chains in Chapter 2 far simpler.

Roughly speaking, our new structural result, stated below as Theorem 3.4, can be used to convert correlation of strength $\Omega(\alpha)$ with $D^{(\ell)}$, where D is contained inside some structured set with density δ, into a density increment of strength $\Omega(1)$ on a set with relative density $|X'|/|X| \approx \exp(-\alpha^{-1/\ell}\mathcal{L}(\delta))$. Combining this with Lemma 3.1, for example, results in a density increment on a set with relative density $\exp(-\alpha^{-1}\mathcal{L}(\alpha))$. An application of Lemma 3.2, in contrast, would yield a density increment on a set with relative density either $\exp(-\alpha^{-1/2}\mathcal{L}(\alpha))$ or

$$\exp(-\alpha^{-1-1/\ell}\ell^2\mathcal{L}(\alpha)^2) \approx \exp(-\alpha^{-1}\mathcal{L}(\alpha)^4),$$

making the optimal choice of $\ell \approx \mathcal{L}(\alpha)$. Thus a direct application of Lemma 3.1 is quantitatively better than the worst result of applying Lemma 3.2.

3.3 From correlation to density increment

We first state our new structural result for spectra and show how this can be combined with Lemma 3.1 to prove Theorem 2.3. We say that Δ is d-covered by Λ if there exists a finite set $\Gamma \subset \widehat{G}$ of size at most d such that $\Delta \subset \Lambda - \Lambda + \langle \Gamma \rangle$, where

$$\langle \Gamma \rangle = \left\{ \sum_{\gamma \in \Gamma} \epsilon_\gamma \gamma : \epsilon_\gamma \in \{-1, 0, 1\} \right\}.$$

We first recall, for comparison, the structural result of Chang. This is proved by Chang [13] in the special case when G is a finite group and $B = G$. The following more general version can be proved using the technique of Shkredov [70]. We will not require this result in our method.

Theorem 3.3. *Let B be a finite subset of G. Let $f \in L^1(B)$ and $\delta = \|f\|_1 \|f\|_\infty^{-1} |B|^{-1}$. The set $\Delta_\eta(f)$ is d-covered by $2\Delta_{\exp(-8\mathcal{L}(\delta)\mathcal{L}(\eta))}(B)$ for some $d \ll \eta^{-2}\mathcal{L}(\delta)$.*

Our new structural result shows that any function supported on the spectrum of large Fourier coefficients has a large proportion of its weight supported on a subset which is more efficiently covered than the full spectrum; Theorem 3.3, on the other hand, manages to cover the entire spectrum, but less efficiently.

Theorem 3.4. *Let B be a finite subset of G. Let $f \in L^1(B)$ and $\delta = \|f\|_1 \|f\|_\infty^{-1} |B|^{-1}$. For any function $\omega : \widehat{G} \to \mathbb{R}_+$ supported on $\Delta_\eta(f)$ there exists $\Delta' \subset \Delta_\eta(f)$ such that*

$$\int_{\Delta'} \omega(\gamma)\, \mathrm{d}\gamma \geq 2^{-10}\eta \int \omega(\gamma)\, \mathrm{d}\gamma,$$

and Δ' is d-covered by $2\Delta_{\exp(-8\mathcal{L}(\delta)\mathcal{L}(\eta))}(B)$ for some $d \leq 2^{11}\eta^{-1}\mathcal{L}(\delta)$.

For a more easily digestible version one should take $B = G$ as a finite group, so that $\Delta_{\exp(-8\mathcal{L}(\delta)\mathcal{L}(\eta))}(B) = \{0\}$. For many applications, however, one must work with sets which are not finite groups and have a much weaker degree of additive structure, for which the generality of Theorem 3.4 as we have stated it will be necessary.

We shall also need the following technical lemma which gives a convenient relationship between shelter and control on a spectrum. This simple argument was used by Green and Konyagin [29], in the proof of their Lemma 3.6.

Lemma 3.5. *Let B and B' be finite subsets of G and $\epsilon, \delta \in (0,1]$. If B' is a symmetric set which is $\delta\epsilon$-sheltered by B then B' has 2δ-control of $\Delta_\epsilon(B)$.*

Proof. For any $\gamma \in \widehat{G}$ and $x \in G$

$$\begin{aligned}(1-\gamma(x))\widehat{B}(\gamma) &= \sum_{y\in B}\gamma(y) - \sum_{y\in B+x}\gamma(y)\\ &= \sum_{y\in B\setminus(B+x)}\gamma(y) - \gamma(x)\sum_{y\in B\setminus(B-x)}\gamma(y).\end{aligned}$$

By the triangle inequality,

$$|1-\gamma(x)|\left|\widehat{B}(\gamma)\right| \leq |B\setminus(B+x)| + |B\setminus(B-x)|\,.$$

In particular, if $\gamma \in \Delta_\epsilon(B)$ and $x \in B'$ then

$$|1-\gamma(x)|\,\epsilon\,|B| \leq 2\,|(B'+B)\setminus B| \leq 2\delta\epsilon\,|B|\,,$$

so that $|1-\gamma(x)| \leq 2\delta$ and the lemma follows. □

We now turn to the task of converting a hypothesis of correlation of a given set to the conclusion that this set has increased density on some structured subset. As promised we will use Theorem 3.4 to do this efficiently. If Theorem 3.3 is used instead then a similar but quantitatively weaker statement can be proved, with $\tau^{-1/r}$ below replaced by $\tau^{-2/r}$.

Lemma 3.6. *Let $r \geq 1$ be any integer. Let B and B' be finite symmetric subsets of G and $A \subset B$ with density α. Let $f \in L^1(B')$ and $\delta = \|f\|_1 \|f\|_\infty^{-1} |B'|^{-1}$.*

Suppose that A in B has correlation of strength τ with $f^{(r)}$. Then there exists a finite set $\Gamma \subset \widehat{G}$ of size $|\Gamma| \leq 2^{13}\tau^{-1/r}\mathcal{L}(\delta)$ such that if B'' is a finite symmetric subset of G such that

1. *B'' has 2^{-2}-control of $\langle\Gamma\rangle$,*
2. *B'' is $2^{-30r^2}\alpha$-sheltered by B, and*
3. *B'' is $\exp(-2^5 r\mathcal{L}(\tau)\mathcal{L}(\delta))$-sheltered by B',*

then $\|A * B''\|_\infty \geq (\alpha + 2^{-15r^2}\tau)\,|B''|$.

Proof. Without loss of generality we may suppose that $\|f\|_1 = 1$. If the conclusion holds for τ then it also holds for any $\tau' \leq \tau$ so we may, without loss of generality, assume that

$$\int \left|\widehat{\mathbf{A}}(\gamma)\right|^2 \left|\widehat{f}(\gamma)\right|^r \, \mathrm{d}\gamma = \tau \, |A| \, .$$

Let $\eta = (\tau/2)^{1/r}$. By Parseval's identity

$$\int_{\widehat{G}\setminus\Delta_\eta(f)} \left|\widehat{\mathbf{A}}(\gamma)\right|^2 \left|\widehat{f}(\gamma)\right|^r \, \mathrm{d}\gamma \leq \eta^r \, \|\mathbf{A}\|_2^2 \leq 2^{-1}\tau \, |A| \, ,$$

and hence

$$\int_{\Delta_\eta(f)} \left|\widehat{\mathbf{A}}(\gamma)\right|^2 \left|\widehat{f}(\gamma)\right|^r \, \mathrm{d}\gamma \geq 2^{-1}\tau \, |A| \, .$$

For all $i \geq 0$ let $\Delta_i = \tilde{\Delta}_{2^i\eta}(f)$, so that

$$\sum_{i\geq 0} \left(2^{i+1}\eta\right)^r \int_{\Delta_i} \left|\widehat{\mathbf{A}}(\gamma)\right|^2 \, \mathrm{d}\gamma \geq 2^{-1}\tau \, |A| \, .$$

We now apply Theorem 3.4 individually to each $i \geq 0$ with the weight functions ω_i defined by $\omega_i(\gamma) = \left|\widehat{\mathbf{A}}(\gamma)\right|^2$ for $\gamma \in \Delta_i$ and $\omega_i(\gamma) = 0$ otherwise. Let Δ_i' be the sets given by Theorem 3.4, so that

$$(2^i\eta)^{r-1} \int_{\Delta_i'} \left|\widehat{\mathbf{A}}(\gamma)\right|^2 \, \mathrm{d}\gamma \geq 2^{-10}(2^i\eta)^r \int_{\Delta_i} \left|\widehat{\mathbf{A}}(\gamma)\right|^2 \, \mathrm{d}\gamma.$$

It follows that

$$\sum_{i\geq 0} (2^i\eta)^{r-1} \int_{\Delta_i'} \left|\widehat{\mathbf{A}}(\gamma)\right|^2 \, \mathrm{d}\gamma \geq 2^{-r-11}\tau \, |A| \, .$$

By Hölder's inequality

$$\sum_{i\geq 0} (2^i\eta)^{r-1} \int_{\Delta_i'} \left|\widehat{\mathbf{A}}(\gamma)\right|^2 \, \mathrm{d}\gamma \leq \left(\sum_{i\geq 0} (2^i\eta)^r \int_{\Delta_i} \left|\widehat{\mathbf{A}}(\gamma)\right|^2 \, \mathrm{d}\gamma\right)^{1-1/r} \left(\sum_{i\geq 0} \int_{\Delta_i'} \left|\widehat{\mathbf{A}}(\gamma)\right|^2 \, \mathrm{d}\gamma\right)^{1/r} .$$

Furthermore,

$$\sum_{i\geq 0} (2^i\eta)^r \int_{\Delta_i} \left|\widehat{\mathbf{A}}(\gamma)\right|^2 \, \mathrm{d}\gamma \leq \int \left|\widehat{\mathbf{A}}(\gamma)\right|^2 \left|\widehat{f}(\gamma)\right|^r \, \mathrm{d}\gamma = \tau \, |A| \, .$$

It follows that, if $\Delta' = \cup_{i\geq 0}\Delta_i'$, then

$$\int_{\Delta'} \left|\widehat{\mathbf{A}}(\gamma)\right|^2 \, \mathrm{d}\gamma \geq 2^{-r^2-11r}\tau \, |A| \, .$$

Furthermore, Theorem 3.4 ensures that for each $i \geq 0$ there is a finite set $\Gamma_i \subset \widehat{G}$ of size $|\Gamma_i| \leq 2^{11-i}\eta^{-1}\mathcal{L}(\delta)$ such that, if $\Lambda = \Delta_{\exp(-2^4 r\mathcal{L}(\tau)\mathcal{L}(\delta))}(B)$, then

$$\Delta_i' \subset \langle \Gamma_i \rangle + 4\Lambda.$$

Hence if $\Gamma = \cup_{i\geq 0}\Gamma_i$ then $|\Gamma| \leq 2^{11}\eta^{-1}\mathcal{L}(\delta)\sum_{i\geq 0} 2^{-i} \leq 2^{12}\eta^{-1}\mathcal{L}(\delta)$ and

$$\Delta' \subset \langle \Gamma \rangle + 4\Lambda.$$

Suppose that B'' has 2^{-2}-control of $\langle\Gamma\rangle$ and 2^{-4}-control of Λ. It follows that it has 2^{-1}-control of Δ', so that $\left|\widehat{B''}(\gamma)\right| \geq |B''|/2$ for all $\gamma \in \Delta'$. Hence

$$\|\mathbf{A} * B''\|_2^2 = \int \left|\widehat{\mathbf{A}}(\gamma)\right|^2 \left|\widehat{B''}(\gamma)\right|^2 \mathrm{d}\gamma \geq 2^{-r^2-11r-2}\tau\,|A|\,|B''|^2\,.$$

Recalling that $\mathbf{A} = A - \alpha B$ and expanding out the left hand side yields

$$\|A * B''\|_2^2 + \alpha^2\|B * B''\|_2^2 - 2\alpha\langle B * B'', A * B''\rangle \geq 2^{-r^2-11r-2}\tau\,|A|\,|B''|^2\,.$$

By the Cauchy-Schwarz inequality, since $A * B''$ is supported on $B + B''$,

$$\begin{aligned}\left|\langle B * B'', A * B''\rangle - |A|\,|B''|^2\right| &\leq \|A * B''\|_2\,\|B * B'' - |B''|\,(B + B'')\|_2 \\ &\leq |B''|\,|A|^{1/2}\,\|B * B'' - |B''|\,(B + B'')\|_2\,.\end{aligned}$$

Furthermore, since B'' is $2^{-30r^2}\alpha$-sheltered by B we have

$$\begin{aligned}\|B * B'' - |B''|\,(B + B'')\|_2^2 &= \|B * B''\|_2^2 - 2\,|B|\,|B''|^2 + |B''|^2\,|B + B''| \\ &\leq |B''|^2\,|(B + B'')\backslash B| \\ &\leq 2^{-30r^2}\,|A|\,|B''|\,.\end{aligned}$$

It follows that

$$\langle B * B'', A * B''\rangle \geq (1 - 2^{-15r^2})\,|A|\,|B''|^2$$

and hence, using the trivial bound $\|B * B''\|_2^2 \leq |B|\,|B''|^2$, it follows that

$$\|A * B''\|_2^2 \geq (\alpha\,|A| + 2^{-15r^2}\tau\,|A|)\,|B''|^2\,.$$

The left hand side is at most $|A|\,|B''|\,\|A * B''\|_\infty$ and the lemma follows since the hypothesis of $\exp(-2^5 r\mathcal{L}(\tau)\mathcal{L}(\delta))$-shelter implies 2^{-4}-control of Λ by Lemma 3.5. □

Lemma 3.6 is the engine by which we will convert correlation to a density increment which can then be iterated. Coupling it with the correlation-producing Lemma 3.1 yields the following corollary.

Corollary 3.7. *Let B and $B'_2, \dots, B'_{s-1}$ be finite symmetric subsets of G such that $0 \in B'_i$ for $1 < i < s$. Let A_1 and A_s be subsets of B with densities α_1 and α_s respectively, and for $1 < i < s$ let $A_i \subset B'_i$ with density at least α'. Suppose that $B'_2 + \cdots + B'_{s-1}$ is $2^{-30s^2}\alpha$-sheltered by B, where $\alpha = \min(\alpha_1, \alpha_s)$. Then either*

$$\langle A_1 * A_2 * \cdots * A_{s-1}, A_s \rangle \geq 2^{-2} |A_1| \cdots |A_s| |B|^{-1}, \tag{3.8}$$

or there exists a finite set $\Gamma \subset \widehat{G}$ of size $|\Gamma| \leq 2^{15}\alpha^{-1/(s-2)}\mathcal{L}(\alpha')$ and $1 < i < s$ such that if $B'' \subset B'_i$ is a symmetric set such that

1. *B'' has 2^{-2}-control of $\langle \Gamma \rangle$ and*
2. *B'' is $\exp(-2^6 s\mathcal{L}(\alpha)\mathcal{L}(\alpha'))$-sheltered by B'_i,*

then $\|A_j * B''\|_\infty \geq (1 + 2^{-15s^2})\alpha_j |B''|$ for some $j \in \{1, s\}$.

Proof. This is an immediate consequence of Lemmata 3.1 and 3.6. In particular, it is simple to check that the hypotheses are strong enough to provide the required amount of shelter; thus, for instance, since $B'' \subset B'_i$ and B'_i is $2^{-26s^2}\alpha$-sheltered by B it certainly follows that B'' is $2^{-30s^2}\alpha$-sheltered by B. □

For the proof of Theorem 2.3 it remains to apply Corollary 3.7 to the case when $A_i = c_i \cdot A$ (or rather, some subset of $c_i \cdot A$ chosen such that the hypotheses are met) and frame the hypotheses using the language of chains from Chapter 2. We recall the key definitions from Chapter 2: a good pair for $(X, \mathbf{c}; \delta)$ is a pair $(\tilde{B}, \tilde{B}')$ of finite symmetric subsets of G, each containing 0, such that

1. the sets $c_1 \cdot \tilde{B}$, $c_s \cdot \tilde{B}$, and $c_1 c_s \cdot \tilde{B}'$ are all δ-sheltered by X, and
2. $c_2 \cdot \tilde{B}' + \cdots + c_{s-1} \cdot \tilde{B}'$ is δ-sheltered by $\tilde{B}$.

If $\mathfrak{B}$ is a collection of finite symmetric subsets of G, each containing 0, then $\mathfrak{B}$ is permissible for $(X, \mathbf{c}; \delta, d)$ if there exists some pair of sets $(\tilde{B}, \tilde{B}')$ which is a good pair for $(X, \mathbf{c}; \delta)$ such that

1. $c_1c_s \cdot \tilde{B}'$, $c_1 \cdot \tilde{B}$, $c_s \cdot \tilde{B}$ are all members of $\mathfrak{B}$ and

2. for each finite $\Gamma \subset \widehat{G}$ of size at most d and each $2 \leq i \leq s-1$, the collection $\mathfrak{B}$ contains $c_1c_sc_i \cdot X'$ for some $X' \subset \tilde{B}'$ which has $(4|\Gamma|)^{-1}$-control of Γ and is δ-sheltered by $\tilde{B}'$.

Finally, we recall that a trivial solution to the equation

$$c_1x_1 + \cdots + c_sx_s = 0 \tag{3.9}$$

is one where $x_1 = \cdots = x_s$.

Lemma 3.8. *Let $\mathbf{c} \in (R^*)^s$ be a commutative s-tuple such that $c_1 + \cdots + c_s = 0$. Let X be a finite symmetric set and let $A \subset X$ with density α. Suppose that A has only trivial solutions to* (3.9) *and let*

$$\delta = \exp(-2^{14}s^2\mathcal{L}(\alpha)^2) \text{ and } d = 2^{15}\alpha^{-1/(s-2)}\mathcal{L}(\alpha).$$

Then if $\mathfrak{B}$ is permissible for $(X, \mathbf{c}; \delta, d)$ either

1. *there exists $X' \in \mathfrak{B}$ such that $\alpha^2 \leq 2^4|X'|^{-1}$, or*

2. *there exist $X' \in \mathfrak{B}$ and $A' \subset X'$ with density α' such that $A' \subset a \cdot A - x$ for some $a \in R^*$ which commutes with $\mathbf{c}$ and $x \in G$, and*

$$\alpha' \geq \alpha(1 + 2^{-20s^2}).$$

Proof. Let $\mathfrak{B}$ be a fixed collection of finite symmetric sets which is permissible for $(X, \mathbf{c}; \delta, d)$, and let $(\tilde{B}, \tilde{B}')$ be some associated good pair for $(X, \mathbf{c}; \delta)$. Let $B = c_1c_s \cdot \tilde{B}$ and $B' = c_1c_s \cdot \tilde{B}'$. Let $\epsilon > 0$ be some constant to be chosen later. We observe that if Y is $\epsilon\alpha$-sheltered by X then

$$|A||Y| - \sum_{x \in X} A * Y(x) = \sum_{x \notin X} A * Y(x) \leq |Y|\,|(X+Y)\backslash X| \leq \epsilon|A||Y|.$$

In particular, provided $\delta \leq \epsilon\alpha$,

$$\sum_{x \in X} \Big(A * (c_1 \cdot \tilde{\beta})(x) + A * (c_s \cdot \tilde{\beta})(x) + A * \beta'(x)\Big) \geq 3(1-\epsilon)|A|,$$

whence there exists some $x \in X$ such that

$$A * (c_1 \cdot \tilde{\beta})(x) + A * (c_s \cdot \tilde{\beta})(x) + A * \beta'(x) \geq 3(1-\epsilon)\alpha.$$

An easy calculation shows that either one summand has size at least $(1+\epsilon)\alpha$, and we are in the second case of the lemma provided $\epsilon \geq 2^{-20s^2}$, or all summands are at least $(1-5\epsilon)\alpha$, so we henceforth assume that we have fixed an $x \in X$ such that

$$(1-5\epsilon)\alpha \leq A * (c_1 \cdot \tilde{\beta})(x), A * (c_s \cdot \tilde{\beta})(x), A * \beta'(x) \leq (1+\epsilon)\alpha.$$

Let

$$A_1 = c_1 \cdot (A - x) \cap B \text{ and } A_s = -c_s \cdot (A - x) \cap B.$$

Since c_1 and $c_1 c_s$ belong to R^* they preserve the cardinalities of sets of G, and hence

$$\alpha_1 = \frac{|A_1|}{|B|} = \frac{\left|c_1 \cdot (A-x) \cap c_1 c_s \cdot \tilde{B}\right|}{\left|\tilde{B}\right|} = A * (c_s \cdot \tilde{\beta})(x).$$

Similarly, $\alpha_s = A * (c_1 \cdot \tilde{\beta})(x)$, and hence $(1-5\epsilon)\alpha \leq \alpha_1, \alpha_s \leq (1+\epsilon)\alpha$. Furthermore, if we let

$$A_i = c_i \cdot (A - x) \cap c_i \cdot B' \text{ for } 2 \leq i \leq s-1$$

then $\alpha_i' = |A_i| / |c_i \cdot B'| \geq (1-5\epsilon)\alpha \geq \alpha/2$, say.

We are now in a position to apply Corollary 3.7. We certainly have that $\delta \leq 2^{-30s^2}(1-5\epsilon)\alpha$, so that the second condition of a good pair ensures that $c_2 \cdot B' + \cdots + c_{s-1} \cdot B'$ is $2^{-30s^2}\min(\alpha_1, \alpha_s)$-sheltered by B, as required. If (3.8) holds then

$$\langle A_1 * \cdots * A_{s-1}, A_s \rangle \geq 2^{-2} |A_1| \cdots |A_s| |B|^{-1}.$$

By hypothesis, however, A has only trivial solutions to $c_1 x_1 + \cdots + c_s x_s = 0$, and hence, since $c_1 + \cdots + c_s = 0$, a solution to $a_1 + \cdots + a_{s-1} = a_s$ with $a_i \in A_i$ is fixed whenever we have fixed $a_2, \ldots, a_{s-1}$, and so

$$\langle A_1 * \cdots * A_{s-1}, A_s \rangle \leq |A_2| \cdots |A_{s-1}|.$$

In particular, $2^2 \geq |A_1| |A_s| |B|^{-1} \geq 2^2 \alpha^2 |B|$ and the first alternative holds.

Otherwise, we are in the second case of Corollary 3.7, and hence there is some $\Gamma \subset \widehat{G}$ of size $|\Gamma| \leq d$ and $1 < i < s$ such that if $B'' \subset c_i \cdot B'$ is a symmetric set which has 2^{-2}-control of $\langle \Gamma \rangle$ and is δ-sheltered by $c_i \cdot B'$, then

$$\|A_j * \beta''\|_\infty \geq (1 + 2^{-15s^2})(1-5\epsilon)\alpha \geq (1 + 2^{-16s^2})\alpha,$$

provided $\epsilon \leq 2^{-15s^2-3}$, say, for some $j \in \{1, s\}$. We let $A' = (A_i - y) \cap B''$ for an appropriate $y \in G$; it remains to verify that we can choose some $B'' \in \mathfrak{B}$ to satisfy the previous conditions.

By definition there is some $X' \subset \tilde{B}'$ such that $c_1c_sc_i \cdot X' \in \mathfrak{B}$, X' has 2^{-2}-control of $\Gamma \cdot c_1c_sc_i$, and X' is δ-sheltered by $\tilde{B}'$. Since $c_1c_sc_i \cdot \tilde{B}' = c_i \cdot B'$ it suffices to take $B'' = c_1c_sc_i \cdot X'$ and the proof is complete. □

We now come at last to the proof of Theorem 2.3. We recall that a chain $\mathfrak{X}_0, \ldots, \mathfrak{X}_K$ satisfies $DI(\delta, d; \mathbf{c})$ if for every $X \in \mathfrak{X}_{i-1}$ the collection $\mathfrak{X}_i$ contains some collection permissible $(X, \mathbf{c}; \delta, d(i))$. Of course, the definition of a chain has been chosen to precisely fit the hypotheses, and so Theorem 2.3 will be an immediate corollary of Lemma 3.8. For convenience, we restate Theorem 2.3 below.

Theorem 2.3. *There exists a constant $C(s) > 1$, depending only on s, such that the following holds. Let $\mathbf{c} \in (R^*)^s$ be a commutative s-tuple such that $c_1 + \cdots + c_s = 0$. Let $X \subset G$ be a finite symmetric set and $A \subset X$ with density α. Let*

$$\delta = \exp(-C\mathcal{L}(\alpha)^2) \text{ and } r(i) = C(1 + C^{-1})^{-i/(s-2)}\alpha^{-1/(s-2)}\mathcal{L}(\alpha) \text{ for } i \geq 0.$$

If A contains only trivial solutions to (2.1) then there exists an integer $0 \leq K \ll_s \mathcal{L}(\alpha)$ such that if $\mathfrak{X}$ is a chain from X of length K which satisfies $DI(\delta, r; \mathbf{c})$ then there exists $X' \in \mathfrak{X}$ such that

$$|X'|^{-1} \gg \alpha^2.$$

Proof. We let $C > 1$ be some constant to be determined later. Let $K \geq 0$ be maximal such that there exists at least one chain from X of length K which satisfies $DI(\delta, r; \mathbf{c})$, and for every such chain $\mathfrak{X}_0, \ldots, \mathfrak{X}_K$ and $0 \leq i \leq K$ there exists some $X_i \in \mathfrak{X}_i$ and $A_i \subset X_i$ with density

$$\alpha_i = |A_i| / |X_i| \geq (1 + C^{-1})^i\alpha,$$

where $A_i \subset a \cdot A - x$ for some $a \in R^*$ which commutes with $\mathbf{c}$ and $x \in G$. Since $\alpha_i \leq 1$ for $0 \leq i < K$ it follows that $K \ll_C \mathcal{L}(\alpha)$.

If there is no chain of length $K+1$ which satisfies $DI(\delta, r; \mathbf{c})$ then the conclusion of the theorem holds vacuously, so we may suppose that such a chain exists. Let $\mathfrak{X}_0, \ldots, \mathfrak{X}_{K+1}$ be some such chain. By construction there exists some $X' \in \mathfrak{X}_K$ and $A' \subset X'$ such that $\alpha' \geq (1 + C^{-1})^K\alpha$ and $A' \subset a \cdot A - x$ for some $a \in R^*$ which commutes with $\mathbf{c}$ and $x \in G$.

Crucially, since $c_1 + \cdots + c_s = 0$, this implies that any non-trivial solution to (3.9) with the variables in A' can be lifted to a non-trivial solution with the variables in A. In particular, A' has only trivial solutions to (3.9), and so we can apply Lemma 3.8.

Since the chain satisfies $DI(\delta, r; \mathbf{c})$ the collection $\mathfrak{X}_{K+1}$ contains some $\mathfrak{B}$ which is permissible for

$$(X', \exp(-2^{14}s^2\mathcal{L}(\alpha)^2), \mathbf{c}, 2^{-15}\alpha_i^{-1/(s-2)}\mathcal{L}(\alpha)),$$

provided we choose C sufficiently large. By Lemma 3.8 either there exists some $X'' \in \mathfrak{X}_{K+1}$ such that

$$\alpha^2 \leq \alpha'^2 \leq 2^4 |X'|^{-1},$$

and the theorem holds, or there exist $X'' \in \mathfrak{X}_{K+1}$ and $A'' \subset X''$ with density α'' such that $A'' \subset a \cdot A' - x$ for some $a \in R^*$ which commutes with $\mathbf{c}$ and $x \in G$, and

$$\alpha'' \geq \alpha'(1 + 2^{-20s^2}) \geq (1 + C^{-1})^{K+1}\alpha,$$

provided $C \geq 2^{20s^2}$. Since $\mathfrak{X}_0, \ldots, \mathfrak{X}_{K+1}$ was an arbitrary chain from X of length $K + 1$ which satisfies $DI(\delta, r; \mathbf{c})$ this contradicts the maximality of K, and the proof is complete. □

3.4 Structure of spectra

In this section we prove the result which is at the heart of our quantitative improvements to Roth's theorem: a new result about the additive structure of spectra, that is, sets of large Fourier coefficients of a given function. This has already been stated as Theorem 3.4. We will first prove a finitary version and then show how Theorem 3.4 follows.

The proof is as follows. We first obtain an upper bound for the additive energy of any set without the required additive structure using an argument of Bateman and Katz [1], and then combine this with a lower bound for the additive energy of spectra due to Shkredov [69] to deduce that spectra possess the required structure.

The structure in question is that of being efficiently covered by another set; namely, we say that a finite set X is *d-covered* if there exists a set X' of size at most d such that $X \subset \langle X' \rangle$, where

$$\langle X' \rangle = \left\{ \sum_{x \in X'} \epsilon_x x : \epsilon_x \in \{-1, 0, 1\} \text{ for all } x \in X' \right\}$$

and $\langle\emptyset\rangle = \{0\}$. The fact that the spectra of sets can be efficiently covered was first observed by Chang [13], who proved that if G is finite and $A \subset G$ with density α then $\Delta_\eta(A)$ is d-covered for some $d \ll \eta^{-2}\mathcal{L}(\alpha)$. This should be compared with the trivial bound $|\Delta_\eta(A)| \leq \eta^{-2}\alpha^{-1}$, which follows from Parseval's identity. In other words, in applications where one can pass to a covering of a set rather than the set itself one can save a factor of $\alpha^{-1}\mathcal{L}(\alpha)^{-1}$ over the trivial bound.

The conclusion of Chang's lemma cannot be quantitatively improved, as demonstrated by Green [27], who showed that for a wide range of α and η one can construct $A \subset \mathbb{Z}/N\mathbb{Z}$ of density α such that $\Delta_\eta(A)$ cannot be d-covered by any $d \ll \eta^{-2}\mathcal{L}(\alpha)$; these constructions were generalised and extended by Shkredov [68].

The new structural result proved in this section shows that these quantitative bounds can be improved if we are prepared to pass to some dense subset of $\Delta_\eta(A)$; namely, there exists some $\Delta' \subset \Delta_\eta(A)$ such that $|\Delta'| \gg \eta\,|\Delta_\eta(A)|$ and Δ' is d-covered for some $d \ll \eta^{-1}\mathcal{L}(\alpha)$. For the applications to Roth's theorem, as shown in the previous section, one can pass to a dense subset without cost, and hence this result can be used as a quantitatively superior version of Chang's lemma.

Due to the iterative nature of the density increment strategy we shall in fact require the more general form of Theorem 3.4, which applies to any function $f \in L^1(G)$ and has a more flexible notion of density. On a first read-through, however, it is simplest to take f in what follows to be the characteristic function of some set, G to be a finite abelian group and $B = G$.

We first require some technical definitions. For any $\Gamma \subset \widehat{G}$, weight function $\omega : \widehat{G} \to \mathbb{R}_+$ with finite support, and integer $m \geq 1$ we define the additive energy as

$$E_{2m}(\omega, \Gamma) = \sum_{\gamma_1,\ldots,\gamma'_m} \omega(\gamma_1)\cdots\omega(\gamma'_m)\Gamma\left(\sum_{i=1}^m \gamma_i - \sum_{j=1}^m \gamma'_j\right).$$

Similarly, for any integers $t_1, t_2 \geq 1$ we define the restricted energy as

$$E^\sharp_{t_1,t_2}(\omega, \Gamma) = \sum_{\substack{\Delta_1 \in \binom{\widehat{G}}{t_1}, \Delta_2 \in \binom{\widehat{G}}{t_2} \\ \Delta_1 \cap \Delta_2 = \emptyset}} \prod_{\gamma \in \Delta_1 \cup \Delta_2} \omega(\gamma)\Gamma\left(\sum_{\gamma \in \Delta_1} \gamma - \sum_{\gamma' \in \Delta_2} \gamma'\right);$$

since ω has finite support both of these sums are well-defined, and there are no issues with convergence. For any ω and Γ we define $E_0(\omega, \Gamma) = E^\sharp_0(\omega, \Gamma) = 1$. Observe that E and

$E^\sharp$ differ, not only in the restriction on repeating elements, but also in that the former is sensitive to permutations of γ_i. We will provide a connection between the energy and restricted energy in Lemma 3.10 below.

We say that Δ is Γ-*dissociated* if for all $k \geq 1$ and $\lambda \in \widehat{G}$ there are at most 2^k many pairs Δ_1, Δ_2 of disjoint subsets of Δ such that $|\Delta_1 \cup \Delta_2| = k$ and

$$\sum_{\gamma \in \Delta_1} \gamma - \sum_{\gamma' \in \Delta_2} \gamma' \in \Gamma + \lambda.$$

This should be compared with the usual notion of dissociativity in arithmetic combinatorics, which is equivalent to the property of having no pair of disjoint subsets $\Delta_1, \Delta_2 \subset \Delta$ with $\sum_{\gamma \in \Delta_1} \gamma - \sum_{\gamma' \in \Delta_2} \gamma' = 0$.

Finally, we say that S has Γ-*dimension* d if d is the size of the largest Γ-dissociated subset of S. We observe that the dimension of a non-empty set is always at least 1 since any singleton set is trivially Γ-dissociated for all $\Gamma \subset \widehat{G}$. Furthermore, if Δ is Γ-dissociated then it is also Γ'-dissociated for any translate Γ' of Γ.

If a set has a very large dimension then almost all of the set is dissociated, and hence one would expect few additive relations between its elements, so it should have small additive energy. The following lemma verifies this intuition.

Lemma 3.9. *Let $\Gamma \subset \widehat{G}$. If a finite set $S \subset \widehat{G}$ has Γ-dimension $|S| - k$ then for all $t_1, t_2 \geq 0$,*

$$E^\sharp_{t_1,t_2}(S, \Gamma) \leq 4^k 2^{t_1+t_2}.$$

Proof. Let $S = S_0 \sqcup S_1$ where S_0 is Γ-dissociated and $|S_1| = k$. By separating the contribution from the subsets of S_1 we obtain the estimate

$$\begin{aligned} E^\sharp_{t_1,t_2}(S, \Gamma) &= \sum_{\substack{\Delta_1 \in \binom{S}{t_1}, \Delta_2 \in \binom{S}{t_2} \\ \Delta_1 \cap \Delta_2 = \emptyset}} \Gamma\left(\sum_{\gamma \in \Delta_1} \gamma - \sum_{\gamma' \in \Delta_2} \gamma'\right) \\ &\leq \sum_{\substack{0 \leq r_1 \leq t_1 \\ 0 \leq r_2 \leq t_2}} \binom{k}{r_1}\binom{k}{r_2} \sup_\lambda \sum_{\substack{\Delta_1 \in \binom{S_0}{t_1 - r_1}, \Delta_2 \in \binom{S_0}{t_2 - r_2} \\ \Delta_1 \cap \Delta_2 = \emptyset}} \Gamma\left(\sum_{\gamma \in \Delta_1} \gamma - \sum_{\gamma' \in \Delta_2} \gamma' + \lambda\right). \end{aligned}$$

Since S_0 is Γ-dissociated, however, the inner summand is bounded above by $2^{t_1+t_2}$ and the lemma follows. □

For the main result of this section we need to convert a conclusion about the restricted additive energy to the full additive energy, for which the following lemma will suffice.

Lemma 3.10. *Let $\Gamma \subset \widehat{G}$ and let $\omega : \widehat{G} \to \mathbb{R}_+$ be some weight function with finite support and $\omega_2 = \left(\sum_\gamma \omega(\gamma)^2\right)^{1/2}$. Then for any $m \geq 2$,*

$$E_{2m}(\omega, \Gamma) \leq 2^{4m}(m!)^2\omega_2^{2m} \sum_{0\leq t_1,t_2\leq m} \frac{\omega_2^{-t_1-t_2}}{((m-t_1)!(m-t_2)!)^{1/2}} \sup_\lambda E^\sharp_{t_1,t_2}(\omega, \Gamma+\lambda).$$

Proof. We divide the range of summation of E_{2m} according to the size of the subsets of $\{\gamma_1, \dots, \gamma_m\}$ and $\{\gamma'_1, \dots, \gamma'_m\}$ consisting of elements that occur with multiplicity 1. This leads to the upper bound

$$E_{2m}(\omega, \Gamma) \leq \sum_{0\leq l_1,l_2\leq m} G_{m-l_1}(\omega)G_{m-l_2}(\omega)\binom{m}{l_1}\binom{m}{l_2} \sup_\lambda F_\lambda(l_1, l_2), \tag{3.10}$$

where

$$F_\lambda(l_1, l_2) = \sum_{\substack{\gamma_1,\dots,\gamma'_{l_2} \\ \gamma_i\neq\gamma_j\ \gamma'_i\neq\gamma'_j\ i\neq j}} \omega(\gamma_1)\cdots\omega(\gamma'_{l_2})\Gamma\left(\sum_{i=1}^{l_1}\gamma_i - \sum_{j=1}^{l_2}\gamma'_j - \lambda\right)$$

and

$$G_k(\omega) = \sum_\Delta^* \prod_{\gamma\in\Delta} \omega(\gamma) \leq \frac{k!}{(\lfloor k/2\rfloor)!}\left(\sum_\gamma \omega(\gamma)^2\right)^{k/2}, \tag{3.11}$$

the first sum being restricted to those ordered k-tuples $\Delta \in \widehat{G}^k$ where each element occurs with multiplicity at least 2. The sum $F_\lambda(l_1, l_2)$ is almost a dilated copy of the restricted energy $E^\sharp_{l_1,l_2}$ except that it lacks the restriction $\gamma_i \neq \gamma'_j$ for all $1 \leq i \leq l_1$ and $1 \leq j \leq l_2$. To introduce this we partition $F_\lambda(l_1, l_2)$ according to the number of common elements between the γ_i and γ'_i; thus

$$F_\lambda(l_1, l_2) \leq \sum_{i=0}^{\min(l_1,l_2)} \binom{l_1}{i}\binom{l_2}{i} i!\omega_2^{2i}(l_1-i)!(l_2-i)!E^\sharp_{l_1-i,l_2-i}(\omega, \Gamma+\lambda). \tag{3.12}$$

Combining (3.10), (3.11) and (3.12) and simplifying the expression shows that $E_{2m}(\omega, \Gamma)$ is at most

$$(m!)^2\omega_2^{2m} \sum_{0\leq l_1,l_2\leq m} \sum_{i=0}^{\min(l_1,l_2)} \frac{\omega_2^{2i-l_1-l_2}}{i!(\lfloor (m-l_1)/2\rfloor)!(\lfloor (m-l_2)/2\rfloor)!} \sup_\lambda E^\sharp_{l_1-i,l_2-i}(\omega, \Gamma+\lambda).$$

Relabelling $t_1 = l_1 - i$ and $t_2 = l_2 - i$ this is at most

$$(m!)^2 \omega_2^{2m} \sum_{0 \leq t_1, t_2 \leq m} \omega_2^{-t_1 - t_2} \sup_{\lambda} E^{\sharp}_{t_1, t_2}(\omega, \Gamma + \lambda) f(m, t_1, t_2),$$

where

$$f(m, t_1, t_2) = \sum_{i=0}^{m - \max(t_1, t_2)} \frac{1}{i!(\lfloor (m - t_1 - i)/2 \rfloor)!(\lfloor (m - t_2 - i)/2 \rfloor)!}.$$

Finally, a tedious calculation using the elementary inequality $n!/(\lfloor n/2 \rfloor!)^2 \leq 2(n+1)^{1/2} 2^n$, valid for all $n \geq 0$, shows that the inner sum is at most $2^{4m}((m - t_1)!(m - t_2)!)^{-1/2}$ and the lemma follows. □

The final technical lemma of this section provides a relationship between covering and dimension that will be important in the proof of Lemma 3.12.

Lemma 3.11. *Let $\Gamma \subset \widehat{G}$ be a symmetric set. If $\Delta \subset \widehat{G}$ has Γ-dimension r then there is a partition $\widehat{G} = \Lambda_0 \sqcup \Lambda_1$ where Λ_0 is $2r$-covered by Γ and for all $\gamma \in \Lambda_1$ the set $\Delta \cup \{\gamma\}$ has Γ-dimension $r + 1$.*

Proof. By hypothesis we can decompose Δ as $\Delta_0 \sqcup \Delta_1$ where Δ_0 is Γ-dissociated and $|\Delta_0| = r$. Let Δ' be the set of all $\gamma \in \widehat{G}$ such that $\Delta_0 \cup \{\gamma\}$ is not Γ-dissociated. We claim that a suitable decomposition is provided by $\Lambda_0 = \Delta' \cup \Delta_0$ and $\Lambda_1 = \widehat{G} \backslash \Lambda_0$.

Firstly, let $\gamma \in \Lambda_1$. By construction the set $\Delta_0 \cup \{\gamma\}$ is Γ-dissociated, and hence $\Delta \cup \{\gamma\}$ has, by definition, Γ-dimension of at least $r + 1$, since $|\Delta_0 \cup \{\gamma\}| = r + 1$. It remains to show that Λ_0 is $2r$-covered by Γ; for this, it suffices to show that

$$\Delta_0 \cup \Delta' \subset \Gamma - \Gamma + \langle \Delta_0 \rangle + \langle \Delta_0 \rangle.$$

This is obvious for Δ_0. Let $\gamma \in \Delta'$. By construction $\Delta_0 \cup \{\gamma\}$ is not Γ-dissociated, and hence there exists $k \geq 1$ and $\lambda \in \widehat{G}$ such that there are more than 2^k many triples $(\epsilon, \Delta_1', \Delta_2')$ such that $\epsilon \in \{-1, 0, 1\}$, the sets Δ_1' and Δ_2' are disjoint subsets of Δ_0 with $|\Delta_1' \cup \Delta_2'| + |\epsilon| = k$, and

$$\epsilon \gamma + \sum_{\gamma_1' \in \Delta_1'} \gamma_1' - \sum_{\gamma_2' \in \Delta_2'} \gamma_2' \in \Gamma + \lambda.$$

If there exists at least one such triple with $\epsilon = 0$ and at least one with $\epsilon \neq 0$ then it is easy to check that this implies that $\gamma \in \Gamma - \Gamma + \langle \Delta_0 \rangle - \langle \Delta_0 \rangle$ as required. If $\epsilon \equiv 0$ for

all such triples then this contradicts the Γ-dissociativity of Δ_0. Hence we can assume that $\epsilon \in \{-1, 1\}$ for all such triples; this is clearly impossible for $k = 1$, and for $k > 1$ we observe that by the pigeonhole principle there are strictly more than 2^{k-1} such triples with identical ϵ. This, however, is another contradiction to the Γ-dissociativity of Δ_0, considering the translate $\Gamma + \lambda - \epsilon\gamma$. Thus $\gamma \in \Gamma - \Gamma + \langle\Delta_0\rangle - \langle\Delta_0\rangle$ as required, and the proof is complete. □

The following lemma is crucial, and uses random sampling to prove a hereditary version of our earlier intuition: namely, if a set is such that every large subset is not efficiently covered then we must have particularly small additive energy. The argument is a variant on that used by Bateman and Katz in [1, Section 5]. There, however, they only wish to bound the 8-fold additive energy, whereas for our purposes we shall need to deal with the $2m$-fold additive energy where $m \to \infty$, and hence we have taken care to make the dependence on m explicit.

We treat the constants in this argument, as in the rest of this chapter, quite crudely; it is certainly possible to improve them, but such improvements would have a negligible effect on the main results.

Lemma 3.12. *Let $\Gamma \subset \widehat{G}$ be a symmetric set and $\omega : \widehat{G} \to \mathbb{R}_+$ be a function with finite support. Let $\omega_2 = \left(\sum_\gamma \omega(\gamma)^2\right)^{1/2}$. Let $d \geq n \geq 2$ and $d/4 \geq m \geq 2$ be integers chosen such that $\omega_2 \leq m^{1/2}d^{-1}\omega_1$. Then either there is a finite set $\Delta \subset \widehat{G}$ such that*

$$\sum_{\gamma\in\Delta} \omega(\gamma) \geq \frac{n}{d}\sum_\gamma \omega(\gamma)$$

and Δ is $2d$-covered by Γ, or

$$E_{2m}(\omega, \Gamma) \leq 2^{13m+6n} m^{2m} d^{-2m} \left(\sum_\gamma \omega(\gamma)\right)^{2m}.$$

Proof. Without loss of generality we may suppose that $\sum \omega(\gamma) = 1$. We first observe that either we are in the first case, or every subset $\Delta \subset \widehat{G}$ which is $2d$-covered by Γ satisfies $\sum_{\gamma\in\Delta} \omega(\gamma) \leq nd^{-1}$, which we shall assume henceforth.

Let $S \subset \widehat{G}$ be a random set of size at most d chosen by selecting d elements of $\widehat{G}$ at random, where we choose $\gamma \in \widehat{G}$ with probability $\omega(\gamma)$. We claim that for $k \geq 0$ the set S has Γ-dimension $d - k$ with probability at most $n^k/k!$.

Suppose that we have selected $d' \leq d$ elements of S, say S', and that S' has Γ-dimension r. By Lemma 3.11 we can partition $\widehat{G} = \Lambda_0 \sqcup \Lambda_1$ such that Λ_0 is $2d$-covered by Γ and for all $\gamma \in \Lambda_1$ the set $S' \cup \{\gamma\}$ has Γ-dimension of at least $r+1$. Thus, in our model,

$$\mathbb{P}(\dim(S' \cup \{\gamma\}) \leq \dim(S')) \leq \sum_{\gamma \in \Lambda_0} \omega(\gamma) \leq \frac{n}{d},$$

since Λ_0 is $2d$-covered by Γ. From this estimate, combined with the trivial observations that the empty set has Γ-dimension 0 and that Γ-dimension is non-decreasing, it follows that the probability that S has Γ-dimension $d-k$ is at most the probability that k events with probability at most n/d occur in d independent trials, which is at most

$$\binom{d}{k} n^k d^{-k} \leq n^k/k!$$

as required. By Lemma 3.9 it follows that for all $\lambda \in \widehat{G}$ and integers $t_1, t_2 \leq m$ we have

$$\mathbb{E}E^\sharp_{t_1,t_2}(S, \Gamma + \lambda) \leq 4^m \sum_{k=0}^{\infty} \frac{(4n)^k}{k!} = 4^m e^{4n}.$$

Let $1 \leq k \leq 2m$. For any distinct $\gamma_1, \ldots, \gamma_k \in \widehat{G}$ the probability that $\gamma_1, \ldots, \gamma_k \in S$ is at least

$$k!\binom{d}{k}\omega(\gamma_1)\cdots\omega(\gamma_k)\left(1 - \sum_{i=1}^{k}\omega(\gamma_i)\right)^{d-k}.$$

Since $k \leq d/2$ we have $k!\binom{d}{k} \geq (d/2)^k$. Furthermore, by the Cauchy-Schwarz inequality

$$\sum_{i=1}^{k}\omega(\gamma_i) \leq (2m)^{1/2}\omega_2 \leq 2md^{-1} \leq 1/2,$$

so that the second factor is at least

$$\exp\left(-d\sum_{i=1}^{k}\omega(\gamma_i)\right) \geq e^{-2m}.$$

It follows that the probability that $\gamma_1, \ldots, \gamma_k \in S$ is at least $2^{-5m}d^k\omega(\gamma_1)\cdots\omega(\gamma_k)$. By linearity of expectation, assuming $t_1 + t_2 \leq m$,

$$\mathbb{E}E^\sharp_{t_1,t_2}(S, \Gamma + \lambda) \geq 2^{-5m}d^{t_1+t_2}E^\sharp_{t_1,t_2}(\omega, \Gamma + \lambda),$$

and so, for all $\lambda \in \widehat{G}$ and $0 \leq t_1, t_2 \leq m$,

$$E^{\sharp}_{t_1,t_2}(\omega, \Gamma + \lambda) \leq 2^{7m} e^{4n} d^{-t_1-t_2}.$$

From Lemma 3.10 it follows that

$$E_{2m}(\omega, \Gamma) \leq 2^{11m} e^{4n} m! \omega_2^{2m} \left(\sum_{0 \leq t \leq m} \frac{(m!)^{1/2} (\omega_2^{-1} d^{-1})^t}{((m-t)!)^{1/2}} \right)^2.$$

For brevity let $r = \omega_2^{-2} d^{-2} \geq m^{-1}$. By the Cauchy-Schwarz inequality the inner factor is at most

$$(m+1) \sum_{0 \leq t \leq m} \frac{m!}{(m-t)!} r^t \leq (m+1) \sum_{0 \leq t \leq m} (erm)^t \leq (m+1)^2 (erm)^m,$$

say. In particular

$$E_{2m}(\omega, \Gamma) \leq 2^{13m} e^{4n} m^{2m} \omega_2^{2m} r^m = 2^{13m} e^{4n} m^{2m} d^{-2m},$$

and the lemma follows. □

Lemma 3.12 is quite general, and although we have stated it for the dual group $\widehat{G}$ the proof is valid for any abelian group. We shall apply it when ω is supported on the spectrum of some function, which is useful because there is a powerful lower bound on the additive energy of such sets, due originally to Shkredov [69]. In that paper it was shown that if G is a finite group and $A \subset G$ with density α and $\Delta \subset \Delta_\eta(A)$ then $E_{2m}(\Delta, \{0\}) \gg \eta^{2m} \alpha |\Delta|^{2m}$. A simpler proof of this result was given by Shkredov in [70]; it is straightforward to generalise this proof to provide a lower bound for the more general version of additive energy we are considering in this section, which we shall do now.

Lemma 3.13. *Let $\epsilon \in [0,1]$ and $B \subset G$ be any finite set. Let $f \in L^1(B)$ and $\omega : \widehat{G} \to \mathbb{R}_+$ be a function with finite support contained in $\Delta_\eta(f)$. For all integers $m \geq 1$,*

$$E_{2m}(\omega, \Delta_\epsilon(B)) \geq \left(\sum_\gamma \omega(\gamma) \right)^{2m} \left(\left(\eta \frac{\|f\|_1}{\|f\|_{2m/(2m-1)} |B|^{1/2m}} \right)^{2m} - \epsilon \right).$$

Proof. Without loss of generality we may suppose that $\sum_\gamma \omega(\gamma) = 1$. Let χ be defined by $\chi(x) = \sum_\gamma \omega(\gamma) \overline{c_\gamma} \gamma(-x)$, where $c_\gamma \widehat{f}(\gamma) = \left|\widehat{f}(\gamma)\right|$. By construction whenever $\omega(\gamma) \neq 0$ we have $\left|\widehat{f}(\gamma)\right| \geq \eta \|f\|_1$, whence

$$\sum_x f(x) \overline{\chi(x)} = \sum_\gamma \omega(\gamma) \left|\widehat{f}(\gamma)\right| \geq \eta \|f\|_1.$$

By Hölder's inequality, however,

$$\left(\sum_x f(x)\overline{\chi(x)}\right)^{2m} \leq \left(\sum_x |f(x)|^{2m/(2m-1)}\right)^{2m-1} \left(\sum_x B(x)\,|\chi(x)|^{2m}\right).$$

It remains to note that, by the triangle inequality,

$$\begin{aligned}\sum_x B(x)\,|\chi(x)|^{2m} &= \sum_x B(x)\left|\sum_\gamma c_\gamma \omega(\gamma)\gamma(x)\right|^{2m} \\ &\leq \sum_{\gamma_1,\ldots,\gamma'_m} \omega(\gamma_1)\cdots\omega(\gamma'_m)\left|\widehat{B}(\gamma_1+\cdots+\gamma_m-\gamma'_1-\cdots-\gamma'_m)\right|.\end{aligned}$$

It follows that

$$\begin{aligned}\eta^{2m}\,\|f\|_1^{2m} &\leq \|f\|_{2m/(2m-1)}^{2m} \sum_{\gamma_1,\ldots,\gamma'_m} \omega(\gamma_1)\cdots\omega(\gamma'_m)\left|\widehat{B}(\gamma_1+\cdots-\gamma'_m)\right| \\ &\leq \|f\|_{2m/(2m-1)}^{2m}\left(|B|\,E_{2m}(\omega,\Delta_\epsilon(B)) + \epsilon\,|B|\right),\end{aligned}$$

and the proof is complete. □

Finally, we can prove the technical heart of our argument, the finitary version of the aforementioned alternative to Chang's lemma. Again, for our application we need a fairly general statement, but at first glance the reader should take $B = G$ and $\epsilon \to 0$, so that $\Delta_\epsilon(B) = \{0\}$.

Theorem 3.14. *Suppose that $f : B \to \mathbb{C}$ and let $\alpha = \|f\|_1 / \|f\|_\infty |B|$. Let $\omega : \widehat{G} \to \mathbb{R}_+$ be a function with finite support contained in $\Delta_\eta(f)$. Let $0 \leq \epsilon \leq \exp(-8\mathcal{L}(\eta)\mathcal{L}(\alpha))$. There is a set $\Delta' \subset \Delta_\eta(f)$ such that*

$$\sum_{\gamma\in\Delta'} \omega(\gamma) \geq 2^{-12}\eta \sum_\gamma \omega(\gamma)$$

and Δ' is $2^{14}\mathcal{L}(\alpha)\eta^{-1}$-covered by $\Delta_\epsilon(B)$.

Proof. Without loss of generality we may suppose that $\sum\omega(\gamma) = 1$. Let $\omega_2 = \left(\sum_\gamma \omega(\gamma)^2\right)^{1/2}$. Suppose first that $\omega_2 \geq 2^{-12}\mathcal{L}(\alpha)^{-1/2}\eta$ and let Δ' be a random set selected by including $\gamma \in \widehat{G}$ independently with probability $2^{13}\eta^{-1}\mathcal{L}(\alpha)\omega(\gamma)$. Then, if Δ' is this randomly chosen set we have, by Chernoff's inequality, that $|\Delta'| \leq 2^{14}\eta^{-1}\mathcal{L}(\alpha)$ with probability at least $7/8$, say, and

$$\mathbb{E}\sum_{\gamma\in\Delta'} \omega(\gamma) \geq 2^{13}\eta^{-1}\mathcal{L}(\alpha)\omega_2^2 \geq 2^{-11}\eta,$$

and hence by Markov's inequality we have $\sum_{\gamma\in\Delta'}\omega(\gamma)\geq 2^{-12}\eta$ with probability at least $1/2$, and the lemma follows.

Otherwise, we let $n=m=\mathcal{L}(\alpha)$ and $d=\lfloor 2^{12}\eta^{-1}m\rfloor$ and apply Lemmata 3.13 and 3.12. We can suppose, by the above, that $\omega_2\leq 2^{-12}\mathcal{L}(\alpha)^{-1/2}\eta\leq m^{1/2}d^{-1}$, as is necessary for the application of Lemma 3.12. Lemma 3.13 implies that

$$E_{2m}(\omega,\Delta_\epsilon(B))\geq\left(\eta\frac{\|f\|_1}{\|f\|_{2m/(2m-1)}|B|^{1/2m}}\right)^{2m}-\epsilon.$$

We have the trivial bound $\|f\|_{2m/(2m-1)}\leq\|f\|_\infty^{1/2m}\|f\|_1^{1-1/2m}$, and hence if $\epsilon\leq\eta^{2m}\alpha/2$ then

$$E_{2m}(\omega,\Delta_\epsilon(B))\geq\eta^{2m}\alpha/2.$$

By Lemma 3.12 either there is a set Δ' such that

$$\sum_{\gamma\in\Delta'}\omega(\gamma)\geq\frac{m}{d}$$

and Δ' is $2d$-covered by $\Delta_\epsilon(B)$, or

$$\eta^{2m}\alpha\leq 2^{19m+1}m^{2m}d^{-2m}.$$

In particular,

$$d\leq 2^{10}m\eta^{-1}\alpha^{-1/2m},$$

which contradicts our initial choice of d and m, and the proof is complete. □

This concludes the proof of Theorem 3.4 in the special case of ω having finite support. The extension to any $\omega:\widehat{G}\to\mathbb{R}_+$ is a routine exercise in compactness (we recall that $\widehat{G}$ is a compact group).

Theorem 3.4. *Let B be a finite subset of G. Let $f\in L^1(B)$ and $\delta=\|f\|_1\|f\|_\infty^{-1}|B|^{-1}$. For any function $\omega:\widehat{G}\to\mathbb{R}_+$ supported on $\Delta_\eta(f)$ there exists $\Delta'\subset\Delta_\eta(f)$ such that*

$$\int_{\Delta'}\omega(\gamma)\,\mathrm{d}\gamma\geq 2^{-10}\eta\int\omega(\gamma)\,\mathrm{d}\gamma,$$

and Δ' is d-covered by $2\Delta_{\exp(-8\mathcal{L}(\delta)\mathcal{L}(\eta))}(B)$ for some $d\leq 2^{11}\eta^{-1}\mathcal{L}(\delta)$.

Proof. Let $\Delta = \Delta_{\exp(-8\mathcal{L}(\alpha)\mathcal{L}(\eta))}(B)$ for brevity. We observe that, since $\left|\widehat{B}\right|$ is a continuous function from $\widehat{G}$ to $\mathbb{C}$ and

$$\Delta \supset \{\gamma : \left|\widehat{B}(\gamma)\right| > \exp(-8\mathcal{L}(\delta)\mathcal{L}(\eta)\},$$

the set Δ contains a non-empty open subset of $\widehat{G}$. Since $\widehat{G}$ is compact it follows that there exists some finite $\Lambda \subset \widehat{G}$ such that $\widehat{G} = \Delta + \Lambda$. Hence we can decompose $\widehat{G}$ into finitely many disjoint sets P_λ, where $P_\lambda \subset \Delta + \lambda$ for all $\lambda \in \Lambda$. For each $\lambda \in \Lambda$ let γ_λ be some element of $P_\lambda \cap \Delta_\eta(f)$, or $\gamma_\lambda = \lambda$ if this intersection is empty. We let

$$\Lambda' = \{\gamma_\lambda : \lambda \in \Lambda \text{ such that } P_\lambda \cap \Delta_\eta(f) \neq \emptyset\}$$

and consider the weight function ω' supported on $\Lambda' \subset \Delta_\eta(f)$ defined by

$$\omega'(\gamma_\lambda) = \int_{P_\lambda} \omega(\gamma)\,\mathrm{d}\gamma.$$

We observe that, since ω is supported on $\Delta_\eta(f)$,

$$\sum_\gamma \omega'(\gamma) = \sum_{\lambda\in\Lambda} \int_{P_\lambda} \omega(\gamma)\,\mathrm{d}\gamma = \int \omega(\gamma)\,\mathrm{d}\gamma.$$

By Theorem 3.14 there exists some $\Lambda'' \subset \Lambda'$ such that

$$\int_{\Delta+\Lambda''} \omega(\gamma)\,\mathrm{d}\gamma \geq \sum_{\lambda\in\Lambda''} \omega'(\lambda) \geq 2^{-10}\eta \int \omega(\gamma)\,\mathrm{d}\gamma.$$

Furthermore, Λ'' is d-covered by Δ for some $d \leq 2^{11}\eta^{-1}\mathcal{L}(\alpha)$, and hence $\Delta+\Lambda''$ is d-covered by 2Δ as required. □

3.5 An alternative method

In this final section we prove Lemma 3.2, which encapsulates the method of Sanders [60] and its subsequent generalisation in [3]. As discussed in Section 3.2 the new structural results of the previous section mean that this combinatorial approach is no longer the most efficient route, as it leads to a quantitatively poorer result while increasing the complexity of the hypotheses. Nonetheless, we include a proof here as a demonstration of how these combinatorial techniques can be used.

The two main ideas are a combinatorial 'thickening' transformation coupled with a probabilistic sampling technique developed by Croot and Sisask [17]. It is important that

both arguments take place entirely within physical space, as opposed to the usual heavy dependence on Fourier analysis. The combinatorial transformation essentially converts a sumset $L+S$ into another sumset $L'+S' \subset L+S$ where L is larger than L' and S' is not much smaller than S. By iterating this construction we may pass to the situation where L is very large, and in particular is very dense inside some structured set. By applying this to the sumset $c_1 \cdot A + c_2 \cdot A$ we obtain the lower bound

$$\langle (c_1 \cdot A) * (c_2 \cdot A), (-c_3 \cdot A) \rangle \gg \langle L * S, (-c_3 \cdot A) \rangle$$

where L is now very dense inside a structured set. We could apply traditional Fourier techniques to show that either the latter inner product is large or there is a density increment, but this leads to a quantitatively very poor result; although we have gained in the density of L, the cost that must be paid is a dramatic reduction in the density of S. Sanders [60], however, applied instead the the probabilistic sampling method of Croot and Sisask which asymmetrically weights the sizes of L and S, allowing the increased density of L created by the combinatorial procedure to be exploited efficiently.

We first discuss the combinatorial thickening procedure. The key observation is that for any sets L and S and any $x \in G$ we have the inclusion

$$(L \cup (S-x)) + (L \cap (S+x)) \subset L+S.$$

Since this holds for arbitrary $x \in G$ we have the freedom to choose x such that $L \cup (S-x)$ is very large and $L \cap (S+x)$ is not too small. This idea was first used by Mann [43], and has resurfaced in various forms since; in [75] it is referred to as the e-transform. Sanders [60] observed that if L is contained in a finite group with density λ, say, and this construction is repeated λ^{-1} many times then it produces sets L' and S' such that $L'+S' \subset L+S$ and $L' \subset G$ with constant density. A set with constant density inside a group is, in many ways, essentially as structured as the group itself, and hence analysis of $L'+S'$ will be much easier than $L+S$. As usual, this argument is robust enough to apply even when $L \subset X$ for some X with a modicum of structure, which we shall do below.

The essence of the technique is contained in the following technical lemma. We recall that for a finite set $B \subset G$ the function β denotes the relative density of a set within B, so that $\beta(K) = |B \cap K| / |B|$.

Lemma 3.15. *Let B and B' be symmetric finite subsets of G, and K, L, and S be any finite subsets of G. Let $T \subset B'$ with density τ. Suppose that S is $2^{-2}\tau$-sheltered by B' and B' is $2^{-2}\beta(K)$-sheltered by B. Then either*

1. *there are $y \in B'$ and $S_0 \subset S$ such that, if $L_0 = L \cup (K + y)$, then $\beta(L_0) \geq \beta(L) + 2^{-1}\beta(K)$ and $|S_0| \geq 2^{-2}\tau |S|$ and*

$$L_0 * S_0(x) \leq L * S(x) + K * T(x)$$

   for all $x \in G$, or

2. *there is a set $D \subset B'$ such that $|D| \geq 2^{-1} |T|$ and*

$$\langle K, (L \cap B) * D \rangle \geq 2^{-2} |D| \, |K \cap B| .$$

Proof. Let $\mathcal{S} = \{x \in B' : T * (-S)(x) \geq 2^{-2}\tau |S|\}$. Whenever $x \in \mathcal{S}$ by definition $|S \cap (T - x)| \geq 2^{-2}\tau |S|$. Furthermore, since S is $2^{-2}\tau$-sheltered by B' and $T \subset B'$,

$$|\langle T * (-S), B' \rangle - |S| \, |T|| \leq \sum_{x \notin B'} T * (-S)(x) \leq |S| \, |(B' - S) \backslash B'| \leq 2^{-2} |T| \, |S|$$

and hence

$$(1 - 2^{-2}) |S| \, |T| \leq \langle T * (-S), B' \rangle \leq |\mathcal{S}| \, |S| + 2^{-2} |S| \, |T| .$$

In particular, it follows that $|\mathcal{S}| \geq 2^{-1} |T|$. Let $L' = L \cap B$, $K' = K \cap B$ and define

$$D = \left\{ x \in B' : L' * (-K)(x) \geq 2^{-2} |K'| \right\} .$$

Suppose first that $|D| < 2^{-1} |T|$, so that $\mathcal{S} \backslash D \neq \emptyset$. Since B' is $2^{-2}\beta(K)$-sheltered by B we have, as above,

$$|\langle \mathcal{S} \backslash D, K' * B \rangle - |K'| \, |\mathcal{S} \backslash D|| \leq 2^{-2} |K'| \, |\mathcal{S} \backslash D| ,$$

and hence

$$\sum_{y \in \mathcal{S} \backslash D} |K' \cap (B - y)| = \langle \mathcal{S} \backslash D, K' * B \rangle \geq \left(1 - 2^{-2}\right) |K'| \, |\mathcal{S} \backslash D| .$$

By the pigeonhole principle there exists $y \in \mathcal{S} \backslash D$ such that $|K' \cap (B - y)| \geq (1 - 2^{-2}) |K'|$ and hence

$$|(L \cup (K + y)) \cap B| \geq |L'| + |K \cap (B - y)| - L' * (-K)(y) \geq |L'| + 2^{-1} |K'| .$$

Thus we are in the first case, letting $S_0 = S \cap (T - y)$, since for any $x \in G$

$$(L \cup (K + y)) * (S \cap (T - y))(x) \leq L * S(x) + (K + y) * (T - y)(x) = L * S(x) + K * T(x).$$

Otherwise, $|D| \geq 2^{-1} |T|$ and, by the definition of D,

$$\langle K, L' * (-D) \rangle \geq 2^{-2} |D| \, |K \cap B| ,$$

and the proof is complete. □

Lemma 3.15 can be applied repeatedly as follows.

Lemma 3.16. *Let $c \in [0, 1/2]$ be an arbitrary constant, and let B and B' be finite symmetric subsets of G. Let $T \subset B'$ with density τ, and K be any finite subset of G. Suppose that T' is a finite set which is $2^{-2}\tau$-sheltered by B' and B' is $2^{-2}\beta(K)$-sheltered by B. Finally, suppose that there is some $z \in B'$ such that $T' \subset T - z$ and $|(K + z) \cap B| \geq |K \cap B| /2$. Then for any $m \geq 1$ either*

1. *there are sets $S \subset T'$ and $X \subset B'$ with $1 \leq |X| \leq m$ such that, if $L = K + X$, then*
   a) *$L * S(x) \leq m(K * T(x))$ for all $x \in G$,*
   b) *$\beta(L) \geq \min(c, 2^{-1} m\beta(K))$, and*
   c) *$|S| \geq (2^{-2}\tau)^{m-1} |T'|$,*
2. *or there is a set $D \subset B'$ with $|D| \geq 2^{-1} |T|$ and $C \subset B$ such that*

$$\beta(C) \leq \min(c, m\beta(K)) \text{ and } \langle K, C * D \rangle \geq 2^{-2} |D| \, |K \cap B| .$$

Proof. We prove this by induction on m. When $m = 1$ we set $S = T'$ and $X = \{z\}$. We clearly have, for all $x \in G$, the required bound $(K + z) * T'(x) \leq K * T(x)$ since $T' \subset T - z$. Furthermore, by hypothesis $\beta(L) \geq \beta(K)/2$, and hence we are in the first case of the lemma.

We shall henceforth suppose that the lemma is true for some fixed $m \geq 1$ and prove the claim for $m + 1$. If we are in the second case then we are done immediately; hence suppose otherwise, and we have X_m and S_m as given by the inductive hypothesis. Let $L_m = K + X_m$. If $\beta(L_m) \geq \min(c, 2^{-1}(m + 1)\beta(K))$ then we are in the first case of the

lemma, letting $X_{m+1} = X_m$ and $S_{m+1} = S_m$. Otherwise we are in a position to apply Lemma 3.15 with $L = L_m$ and $S = S_m$.

Since $\beta(L_m) \leq \min(c, 2^{-1}(m+1)\beta(K))$ this implies that either we are in the second case of the lemma (relabelling $L_m \cap B = C$), or there are $x \in B'$ and a set $S_{m+1} \subset S_m$ such that, if $X_{m+1} = X_m \cup \{x\}$ then, letting $L_{m+1} = L_m \cup (K + x) = K + X_{m+1}$,

1. $L_{m+1} * S_{m+1}(x) \leq L_m * S_m(x) + K * T(x) \leq (m+1)(K * T(x))$ for all $x \in G$,
2. $\beta(L_{m+1}) \geq \beta(L_m) + 2^{-1}\beta(K) \geq 2^{-1}(m+1)\beta(K)$, and
3. $|S_{m+1}| \geq 2^{-2}\tau\,|S_m| \geq (2^{-2}\tau)^m\,|T'|$,

and the proof is complete. □

Lemma 3.16 will suffice for creating lower bounds for the convolution of two sets $A_1 * A_2$, which is sufficient for the traditional $s = 3$ case of Roth's theorem. It first appeared in the work of Sanders [60] in the case $s = 3$ and $G = \mathbb{Z}/N\mathbb{Z}$. In [3] we observed that this idea could be applied iteratively to multiple convolutions $A_1 * \cdots * A_k$, and thus extended the method of Sanders to translation invariant equations in $s \geq 4$ variables.

Lemma 3.17. *Let $m \geq 2$ be any integer. Let B, $B'_2, \ldots, B'_k$ and $B''_2, \ldots, B''_k$ be finite symmetric subsets of G. Let $A_1 \subset B$ with density α_1 and for $2 \leq i \leq k$ let $A_i \subset B'_i$ with density α_i. Finally, suppose that B''_i is $2^{-2}\alpha_i$-sheltered by B'_i and that B'_i is $(2^{-2}\alpha_1, 2^{-k-1})$-sheltered by B for $2 \leq i \leq k$.*

Then either

1. *there are sets $S_i \subset B''_i$ for $2 \leq i \leq k$ and $X_i \subset B'_i$ such that $1 \leq |X_i| \leq m$ for $2 \leq i \leq k$, and with $L = A_1 + X_2 + \cdots + X_k$ and $\lambda = |L \cap B| / |B|$ we have*
   a) *$L * S_2 * \cdots * S_k(x) \leq m^{k-1} A_1 * \cdots * A_k(x)$ for all $x \in G$,*
   b) *$\lambda \geq \min(2^{-k-1}, (2^{-1}m)^{k-1}\,\alpha_1)$, and*
   c) *$|S_i| \geq 2^{-2m+1}\alpha_i^m\,|B''_i|$ for $2 \leq i \leq k$,*
2. *or for some $2 \leq i \leq k$ there is a set $D \subset B'_i$ with $|D| \geq 2^{-1}\alpha'\,|B'_i|$ such that A_1 in B has correlation of strength $2^{-2k-2}m^{1-k}$ with $D * D$.*

Proof. We use strong induction on k; thus we shall fix $k \geq 2$ and assume that the theorem is known to hold for all $k' < k$. In particular, we shall assume that there are $X_1 = \{0\}$ and $X_i \subset B_i'$ and $S_i \subset B_i''$ for $2 \leq i \leq k-1$ such that, if $L = A_1 + X_1 + \cdots + X_{k-1}$ and $\lambda = |L \cap B| / |B|$, then

1. $L * S_2 * \cdots * S_{k-1}(x) \leq m^{k-2} A_1 * \cdots * A_{k-1}(x)$ for all $x \in G$,
2. $\lambda \geq \min(2^{-k}, (2^{-1}m)^{k-2} \alpha_1)$, and
3. $|S_i| \geq 2^{-2m+1} \alpha_i^m |B_i''|$ for $2 \leq i \leq k-1$.

This is obvious for $k = 2$, and for $k > 2$ this follows from the induction hypothesis, for if the second case of the theorem holds for $k-1$ then it already holds for k and we are done.

Since B_k'' is $2^{-1}\alpha_k$-sheltered by B' and $A_k \subset B_k'$ we have

$$|\langle B_k'' * A_k, B_k' \rangle - |A_k| |B_k''|| \leq \sum_{x \notin B_k'} B_k'' * A_k(x) \leq 2^{-1} |A_k| |B_k''|,$$

so that

$$\sum_{z \in B_k'} B_k'' * A_k(z) \geq 2^{-1} |A_k| |B_k''|,$$

and hence if $B_0 = \{z \in B_k' : B'' * A_k(z) \geq 2^{-1}\alpha_k |B_k''|\}$ then $B_0 \neq \emptyset$. Furthermore, since B_k' is $2^{-1}\lambda$-sheltered by B it follows similarly that

$$\sum_{z \in B_0} B * L(z) = \sum_{x \in B} (-B_0) * L(x) \geq 2^{-1} |L \cap B| |B_0|,$$

and hence there exists $z \in B_0$ such that $B * L(z) \geq 2^{-1} |L \cap B|$. We fix such z and let $T' = B_k'' \cap (A_k - z)$, so that $\beta_k''(T') \geq 2^{-1}\alpha_k$.

We first observe that if $\lambda \geq 2^{-k}$ then we are done immediately, letting $S_k = T'$ and $X_k = \{z\}$, and hence we can suppose that $\lambda \geq 2^{-k+2}\alpha_1 |X_1| \cdots |X_{k-1}|$. We now apply Lemma 3.16 with $K = L$ and $T = A_k$, and $c = 2^{-k-1}$. If the first case of Lemma 3.16 holds then there are sets $X_k \subset B_k'$ and $S_k \subset B_k''$ such that if $L' = L + X_k = A_1 + X_1 + \cdots + X_k$ and $\lambda' = |L' \cap B| / |B|$, then

1. $L' * S_k(x) \leq m(L * A_k(x))$ for all $x \in G$,
2. $\lambda' \geq \min\left(2^{-k-1}, 2^{-1}m\lambda\right) \geq \min\left(2^{-k-1}, (2^{-1}m)^{k-1}\alpha_1\right)$, and

3. $|S_k| \geq (2^{-2}\alpha_k)^{m-1} |T'| \geq 2^{-2m+1}\alpha_k^m |B_k''|$,

and we are in the first case of the lemma. Otherwise, there is a set $D \subset B_k'$ with $|D| \geq 2^{-1}|A_k|$ and C such that $|C| \leq \min(2^{-k-1}, m\lambda)|B|$ and, if $L = A_1 + X_1 + \cdots + X_{k-1}$, then

$$\begin{aligned}\langle L, C * D\rangle &\geq 2^{-2}|D|\,\lambda\,|B| \\ &\geq 2^{-k}|D|\,\alpha_1\,|X_1|\cdots|X_{k-1}|\,|B|\,.\end{aligned}$$

By the triangle inequality there exists some $x \in G$ such that

$$\langle (A_1 + x), C * D\rangle \geq 2^{-k}|D|\,\alpha_1\,|B|\,.$$

Replacing C by $C - x$ we may assume that $x = 0$. Furthermore, we have the trivial bound $\langle B, C * D\rangle \leq |C|\,|D| \leq 2^{-k-1}|D|\,|B|$. In particular, if we let $\mathbf{A_1} = (A_1 - \alpha_1)B$, then by the Cauchy-Schwarz inequality

$$\begin{aligned}\|\mathbf{A_1} * D\|_2\,|C|^{1/2} &\geq \langle \mathbf{A_1}, C * D\rangle \\ &\geq 2^{-k-1}|D|\,\alpha_1\,|B|\,.\end{aligned}$$

Hence

$$\|\mathbf{A_1} * D\|_2^2 \geq 2^{-2k-2}|C|^{-1}\alpha_1^2|B|^2|D|^2 \geq 2^{-2k-2}m^{-1}\lambda^{-1}\alpha_1^2|B|\,|D|^2\,.$$

It follows that A_1 in B has correlation of strength $2^{-2k-2}m^{-1}\lambda^{-1}\alpha_1$ with $D * D$ and the claim follows, since we trivially have $\lambda \leq \alpha_1|X_1|\cdots|X_{k-1}| \leq m^{k-2}\alpha_1$. □

It is helpful to view Lemma 3.17 with most of the technicalities stripped away; the point is that we can bound the convolution $A_1 * \cdots * A_k$ from below by another convolution $L * S_2 * \cdots * S_k$ such that L has density roughly $m^{k-1}\alpha_1$ and each S_i has density roughly α_i^m. In particular, if we choose $m \approx \alpha_1^{-1/(k-1)}$ then L has constant density within some structured set, and hence must be very structured itself. This fact is only useful, however, if we can work with the convolution $L * S_2 * \cdots * S_k$ in some asymmetric way to lessen the impact of the density of the other sets decreasing from $\exp(-\mathcal{L}(\alpha_i))$ to $\exp(-\alpha_1^{-1/(k-1)}\mathcal{L}(\alpha_i))$.

Again, traditional Fourier analytic techniques fail here, as they treat each part of a convolution with equal weighting and so the gains from the increased density of L are

swamped by the losses resulting from the sparsity of the sets S_i. The new method of Croot and Sisask [17] allows for one of the sets in a convolution to be given special weighting, which enables the combinatorial techniques above to be exploited.

This allowed Sanders [60] to achieve a bound of $1/(\log N)^{1+o(1)}$ for the traditional case of Roth's theorem. The novelty of [3] lies in generalising this technique from $A_1 * A_2$ to $A_1 * \cdots * A_k$ to prove good quantitative bounds for translation invariant equations in $s > 3$ variables. The generalisation lies entirely in the combinatorial arguments presented above; the application of the method of Croot and Sisask is the same as that in [60], but we include it here for a complete proof of the correlation-producing Lemma 3.2.

We first state the sampling theorem of Croot and Sisask; there have been several proofs of this theorem, presented in [17], [60] and [16]. We recall the translation operator τ_x defined by $\tau_x(f(y)) = f(y-x)$.

Theorem 3.18 (Croot-Sisask [17])**.** *Let $\eta \in (0,1]$ and $p \geq 2$. Let S and T be finite subsets of G such that $|S+T| \leq K|S|$ for some $K \geq 2$. For all $g \in L^p(G)$ there exist $u \in T$ and $X \subset T - u$ with*

$$|X| \geq \exp(-2^{12} p\eta^{-2} \log K)\,|T|\,,$$

such that

$$\|\tau_x(S * g) - S * g\|_p \leq \eta\,|S|\,\|g\|_p$$

for all $x \in X$.

We now use Theorem 3.18 to prove another correlation-producing lemma, following the method of [60].

Lemma 3.19. *Let $\ell \geq 1$ be any integer. Let B, B', and B'' be finite symmetric subsets of G such that $|B' + B''| \leq 2|B'|$. Let A and L be subsets of B with densities α and λ respectively, and suppose that $S_k \subset B'$ with density σ_k. Finally, suppose that $2\ell B'' + S_2 + \cdots + S_k$ is $2^{-1}\lambda$-sheltered by B. Then either*

$$\langle L * S_2 * \cdots * S_k, A\rangle \geq 2^{-4}\lambda\,|S_2| \cdots |S_k|\,|A|\,,$$

or there is $u \in B''$ and $D \subset B'' - u$ with

$$|D| \geq \exp\left(-2^{22}\ell^2\lambda^{-2}\mathcal{L}(\alpha)\log(2\sigma_k^{-1})\right)|B''|$$

such that A in B has correlation of strength $2^{-4}\lambda\alpha$ with $D^{(\ell)}$.

Proof. We first observe that since $S_k \subset B'$ we have that

$$|S_k + B''| \leq |B' + B''| \leq 2\,|B'| = 2\sigma_k^{-1}\,|S_k|\,.$$

By Theorem 3.18 there are $u \in B''$ and $D \subset B'' - u$ with

$$|D| \geq \exp\left(-2^{12}\eta^{-2}\ell^2 p \log(2\sigma_k^{-1})\right)|B''|\,,$$

such that for all $d \in D$

$$\|\tau_d f - f\|_p \leq \eta\ell^{-1}\,|L|^{1/p}\,|S_2|\cdots|S_k|\,,$$

where $f = L * S_2 * \cdots * S_k$. By the triangle inequality it follows that

$$\|\delta * f - f\|_p \leq \eta\,|L|^{1/p}\,|S_2|\cdots|S_k|\,,$$

where $|D|^{\ell}\,\delta = D^{(\ell)}$. By Hölder's inequality

$$\begin{aligned}|\langle \delta * f, A\rangle - \langle f, A\rangle| &\leq \eta\,|L|^{1/p}\,|S_2|\cdots|S_k|\,|A|^{1-1/p}\\ &\leq 2^{-3}\lambda\,|S_2|\cdots|S_k|\,|A|\end{aligned}$$

taking $p = \mathcal{L}(\alpha)$ and $\eta = 2^{-5}\lambda$, say. Hence either we are in the first case or

$$\langle \delta * f, A\rangle \leq 2^{-4}\lambda\,|S_2|\cdots|S_k|\,|A|\,. \tag{3.13}$$

By hypothesis we know that $\ell D + S_2 + \cdots + S_k$ is $2^{-1}\lambda$-sheltered by B, whence

$$\langle \delta * f, B\rangle \geq 2^{-1}\,|L|\,|S_2|\cdots|S_k|\,. \tag{3.14}$$

Combining the inequalities (3.13), (3.14) and using the Cauchy-Schwarz inequality yields

$$2^{-2}\lambda\,|S_2|\cdots|S_k|\,|A| \leq |\langle \delta * f, \mathbf{A}\rangle| \leq \|\mathbf{A} * (-\delta)\|_2\,|L|^{1/2}\,|S_2|\cdots|S_k|$$

and the conclusion follows. □

It remains to combine Lemmata 3.17 and 3.19, both of which produce some correlation, to prove the more efficient correlation-producing Lemma 3.2. As discussed above, in the case $s = 3$ and $G = \mathbb{Z}$ this is essentially already present in [60]; the novelty here is in the generalisation of this technique to general abelian groups and $s > 3$.

Lemma 3.2. *Let B be a finite symmetric subset of G. Let A_1 and A_s be subsets of B with relative densities α_1 and α_s respectively, and let $A_2, \dots, A_{s-1}$ be any finite subsets of G. Let $\alpha = \min(\alpha_1, \alpha_s)$.*

Furthermore, suppose that there are finite symmetric sets $B_2', \dots, B_{s-1}', B_2'', \dots, B_{s-1}'', B'''$ such that for each $2 \le i \le s-1$

1. *$A_i \subset B_i'$ with density at least α',*
2. *$2\ell B''' + B_2'' + \cdots + B_{s-1}''$ is $2^{-1}\alpha$-sheltered by B,*
3. *B_i' is $2^{-2s}\alpha$-sheltered by B,*
4. *B_i'' is $2^{-2}\alpha'$-sheltered by B_i', and*
5. *$\left|B_{s-1}'' + B'''\right| \le 2\left|B_{s-1}''\right|$.*

Then either

1. *$\langle A_1 * \cdots * A_{s-1}, A_s\rangle \ge \exp\left(-2^5 s\alpha_1^{-1/(s-2)}\mathcal{L}(\alpha')\right)\prod_{i=2}^{s-1}|B_i''|\,|A_s|$, or*
2. *there is a finite set D with $\beta'(D) \ge 2^{-1}\alpha'$ such that A_1 in B has correlation of strength $2^{-2s}\alpha_1$ with $D * D$, or*
3. *for any integer $\ell \ge 1$ there is a finite set D with*

$$\beta'''(D) \ge \exp\left(-2^{26+2s}\ell^2\alpha_1^{-1/(s-2)}\mathcal{L}(\alpha')\mathcal{L}(\alpha_s)\right)$$

such that A_s in B has correlation of strength $2^{-4-s}\alpha_s$ with $D^{(\ell)}$.

Proof. We first apply Lemma 3.17 with $k = s-1$ and $m = \lceil 2\alpha_1^{-1/(s-2)}\rceil$. This implies that either there is, for some $2 \le i < s$, a set $D \subset B_i'$ with $|D| \ge 2^{-1}\alpha'|B_i'|$ such that A_1 in B has correlation of strength $2^{-2s}\alpha_1$ with $D * D$, and we are in the first case of the lemma, or there are sets $S_i \subset B_i''$ for $1 < i < s$ and $L \subset B$ such that

1. $L * S_2 * \cdots * S_{s-1}(x) \le 2^s\alpha_1^{-1}A_1 * \cdots * A_{s-1}(x)$ for all $x \in G$,
2. $\beta(L) \ge 2^{-s}$, and
3. $|S_i| \ge \exp(-2^3\mathcal{L}(\alpha_i)\alpha_1^{-1/(s-2)})\,|B_i''|$ for $1 < i < s$.

By Lemma 3.19 either

$$\langle L * S_2 * \cdots * S_{s-1}, A_s \rangle \geq \exp(-2^4 s\mathcal{L}(\alpha')\alpha_1^{-1/(s-2)}) \prod_{i=2}^{s-1} |B_i''| \, |A_s| \,,$$

and hence

$$\langle A_1 * \cdots * A_{s-1}, A_s \rangle \geq \exp(-2^5 s\mathcal{L}(\alpha')\alpha_1^{-1/(s-2)}) \prod_{i=2}^{s-1} |B_i''| \, |A_s|$$

as required, or there are $u \in B'''$ and $D \subset B''' - u$ with

$$|D| \geq \exp\left(-2^{26+2s}\ell^2\alpha_1^{-1/(s-2)}\mathcal{L}(\alpha_s)\mathcal{L}(\alpha')\right)|B'''|$$

such that A_s in B has correlation of strength $2^{-4-s}\alpha_s$ with $D^{(\ell)}$, and the proof is complete. □

CHAPTER 4

ARITHMETIC INVERSE THEOREMS IN FUNCTION FIELDS

We recall that $\mathbb{F}_q[t]$ denotes the ring of polynomials over the finite field $\mathbb{F}_q$, where q is some fixed prime power p^r. In this chapter we prove a collection of related arithmetic inverse theorems for finite subsets of $\mathbb{F}_q[t]$. In this introductory section we will suppose that we are in the simplest case when $q = p$.

Let A be a finite subset of $\mathbb{F}_p[t]$. If one is just concerned with the sumset $A + A$ then one may as well view A as a finite subset of $\mathbb{F}_p^n$ for some $n \geq 0$. Questions about the behaviour of $A + A$ therefore belong in the realm of arithmetic combinatorics over $\mathbb{F}_p^n$, which has received a great deal of attention.

In $\mathbb{F}_p[t]$, however, we have more arithmetic available to us than just addition; in particular, we may also ask about the behaviour of $A + t \cdot A$, and especially how large such a set can be as a function of $|A|$. It is easy to verify the trivial inequalities

$$2|A| - 1 \leq |A + t \cdot A| \leq |A|^2,$$

valid for any finite $A \subset \mathbb{F}_p[t]$. Furthermore, both bounds are sharp, so there is not much more that we can say in general. In the spirit of the inverse results for the size of the sumset discussed in the introduction, however, one can hope for some characterisation of the cases when $A + t \cdot A$ is small – that is, a result stating that $|A + t \cdot A|$ is small if and only if A is 'close' to a certain kind of structure.

From Plünnecke's inequality we have that $|A + A| \leq |A + t \cdot A|^2 / |A|$, and hence if $|A + t \cdot A|$ is small then certainly $|A + A|$ is small, and so, viewing A as a subset of some $\mathbb{F}_p^n$, the inverse results discussed in the introduction imply that A is close to being an $\mathbb{F}_p$-vector space. This tells us a great deal about the deeper additive structure of A, but it fails as a characterisation of the cases when $|A + t \cdot A| / |A|$ is small.

That is, while it is true that if $|A + t \cdot A| / |A|$ is small then A is close to being an $\mathbb{F}_p$-vector space the converse fails dramatically. Indeed, it is possible for A to be a finite $\mathbb{F}_p$-vector space, and yet $|A + t \cdot A| = |A|^2$, the maximum possible. Such a set is given by,

for example,

$$A = \left\{ \sum_{i=0}^{n} a_i t^{2i} : a_i \in \mathbb{F}_p \right\}$$

for any $n \geq 0$.

To get an idea of what kind of characterisation is appropriate we should first ask for examples of sets where $|A + t \cdot A|$ is small. By considering the analogy with $\mathbb{Z}$ we quickly arrive at the concept of an $\mathbb{F}_p[t]$-arithmetic progression – a set of the form $a \cdot \mathbb{F}_p[t]_N + x$ for some $a, x \in \mathbb{F}_p[t]$, where $\mathbb{F}_p[t]_N$ denotes the set of polynomials in $\mathbb{F}_p[t]$ of degree less than N. In general, we define an arithmetic space of dimension d to be a set of the shape

$$\mathbb{F}_p[t]_{n_1} \cdot x_1 \oplus \cdots \oplus \mathbb{F}_p[t]_{n_d} \cdot x_d$$

for some $n_1, \ldots, n_d \geq 0$ and $x_1, \ldots, x_d \in \mathbb{F}_p[t]$, and where the sum is direct, considering each component as an $\mathbb{F}_q$-vector space. This is the $\mathbb{F}_p[t]$-analogue of a generalised arithmetic progression in the integers. It is easy to see that if A is an arithmetic space of dimension d then $A + A = A$ and $|A + t \cdot A| \leq p^d |A|$. We say that A is r-covered by V if there exists a set X of size $|X| \leq r$ such that $A \subset V + X$. The proof of the following lemma is immediate.

Lemma 4.1. *Let V be an arithmetic space of dimension d and let A be a finite set which is $\exp(d)$-covered by V such that $|A| \geq \exp(-d)|V|$. Then $|A + t \cdot A| \leq \exp(O_p(d))|A|$.*

The goal, then, is to prove the converse result which will complete our characterisation of sets A such that $|A + t \cdot A|$ is small. By adapting the proof of the sharpest quantitative bounds available for the inverse results over the integers, we are able to prove the following.

Theorem 4.2. *Let $A \subset \mathbb{F}_p[t]$ be a finite set and let $K \geq 4$ be such that $|A + t \cdot A| \leq K|A|$. There exists some*

$$d \ll (\log K)^3 (\log \log K)^{3 \log_2 3}$$

and an arithmetic space of dimension d such that $|V| \leq \exp(d)|A|$ and A is $\exp(d)$-covered by V.

Lemma 4.1 shows that the best bound one could hope for here is $d \ll \log K$. That such a bound in the analogous result over the integers is attainable was conjectured by Ruzsa [57], who credits the analogous conjecture over $\mathbb{F}_p^n$ to Marton. In general, such

a conjecture is now commonly referred to as the Polynomial Freiman-Ruzsa conjecture. There is no reason not to expect the same for the analogous problem in $\mathbb{F}_p[t]$. One attractive aspect of the theory over $\mathbb{F}_p[t]$ is that we are able to prove such a sharp result with the additional hypothesis that $|A+A| \ll |A|$. In other words, for sets which are approximately closed under addition we are able to prove the quantitatively optimal result which characterises sets A such that $|A + t \cdot A| / |A|$ is small.

Theorem 4.3. *Let $A \subset \mathbb{F}_q[t]$ be a finite set and let $K_1, K_2 \geq 2$ be such that $|A+A| \leq K_1 |A|$ and $|A + t \cdot A| \leq K_2 |A|$. Then there is some $d \ll_{p,K_1} \log K_2$ and an arithmetic space of dimension d such that $|V| \leq \exp(d) |A|$ and A is $\exp(d)$-covered by V.*

The above discussion assumes that $q = p$ for some prime p. For the general case when q is a prime power one must slightly strengthen the hypothesis; in particular, one needs to control not only $|A+A|$ but also $|A + a \cdot A|$ for all $a \in \mathbb{F}_q$. By Plünnecke's inequality and the trivial inclusion $n \cdot A \subset nA$ it is easy to see that for $q = p$ control of $|A+A|$ suffices.

This chapter will be structured as follows. In the first section we address the case when A is a finite $\mathbb{F}_q$-vector space, and hence has the strongest possible amount of additive structure. In this situation we are able to exploit the rigid structure of $\mathbb{F}_q[t]$ to prove a very sharp arithmetic inverse theorem.

In the second section we prove some results from the literature that will be needed in the following section; the statements that we require are slightly more general than those present in the literature, although the proofs are largely trivial adaptations.

In the third section we adapt the latest methods used for additive inverse results, as presented by Sanders [63], to weaken the rigid hypotheses of the results from the first section. We have taken some care with our arguments to prove Theorem 4.3 with good dependency on both K_1 and K_2.

Finally, in the fourth section we give an application of such inverse results to an $\mathbb{F}_q[t]$-analogue of a problem of Konyagin and Łaba [35].

4.1 Decomposition of vector spaces

In this section we will prove an arithmetic inverse result with a strong conclusion but with a correspondingly strong hypothesis – in particular, we will be concerned only with

finite $\mathbb{F}_q$-vector spaces V and prove structural results for such sets under the assumption that $|V + t \cdot V|$ is small.

We first introduce some notation. We recall that $\mathbb{F}_q[t]_N = \{x \in \mathbb{F}_q[t] : \deg x < N\}$. A $\mathbb{F}_q$-vector space V has arithmetic dimension k if k is minimal such that there exist $d_1, \ldots, d_k \geq 1$ and $x_1, \ldots, x_k \in \mathbb{F}_q[t]$ such that

$$V = \mathbb{F}_q[t]_{d_1} \cdot x_1 \oplus \cdots \oplus \mathbb{F}_q[t]_{d_k} \cdot x_k.$$

Observe in particular that, if we consider $d_1 = \cdots = d_k = 1$, the arithmetic dimension of V is at most the dimension of V considered as a vector space over $\mathbb{F}_q$. We further say that V has strong arithmetic dimension k if k is minimal such that there exist $d_1, \ldots, d_k \geq 1$ and $x_1, \ldots, x_k \in \mathbb{F}_q[t]$ such that

$$V = \mathbb{F}_q[t]_{d_1} \cdot x_1 \oplus \cdots \oplus \mathbb{F}_q[t]_{d_k} \cdot x_k$$

and $d_1 + \deg x_1 < \cdots < d_k + \deg x_k$. It is clear that the strong arithmetic dimension is always at least the arithmetic dimension; in fact, we will shortly see that they are always identical, and hence this definition is redundant. It is convenient, however, for the inductive proof of Theorem 4.5.

For the main structural theorem of this section we shall require the following technical lemma. This is essentially the univariate case of the important theorem that a minimal Gröbner basis always exists for any finite set in a polynomial ring; in the univariate case, the proof is particularly simple.

Lemma 4.4. *If $V \subset \mathbb{F}_q[t]$ is a finite $\mathbb{F}_q$-vector space then there exists a decomposition of the form*

$$V = \mathbb{F}_q[t]_1 \cdot x_1 \oplus \cdots \oplus \mathbb{F}_q[t]_1 \cdot x_d,$$

where $\deg x_1 < \deg x_2 < \cdots < \deg x_d$.

Proof. We use induction on $\dim V$. If $\dim V = 0$ then the result is trivial. Otherwise, let $x \in V \backslash \{0\}$ be a monic polynomial of minimal degree, and let $V = \mathbb{F}_q[t]_1 \cdot x \oplus W$. The result follows from the inductive hypothesis and the fact that if $w \in W \backslash \{0\}$ then $\deg w > \deg x$. This latter fact is true because otherwise we must have a monic $w \in W \backslash \{0\}$ such that $\deg w = \deg x$, and hence $w - x \in V \backslash \{0\}$ has degree strictly less than $\deg x$, which contradicts the minimality of $\deg x$. □

The following theorem is a very strong arithmetic inverse theorem for $\mathbb{F}_q$-vector spaces V. In particular, if the additive structure of V is as strong as possible then we obtain the best possible inverse result concerning what structure can be deduced from $|V + t \cdot V| / |V|$. The rigid arithmetic structure of $\mathbb{F}_q[t]$ allows for a constructive algebraic proof. We emphasise that the following theorem offers a complete characterisation of subspaces V such that $|V + t \cdot V| / |V|$ is small.

Theorem 4.5. *Let $V \subset \mathbb{F}_q[t]$ be a finite $\mathbb{F}_q$-vector space. We have $|V + t \cdot V| = q^r |V|$ if and only if V has arithmetic dimension r.*

In the proof the following simple fact will be used frequently: if $x, y \in \mathbb{F}_q[t]$ are such that $\deg x < \deg y$, then $\deg(x + y) = \deg y$.

Proof. It follows immediately from the definitions that if V has arithmetic dimension r then $|V + t \cdot V| \leq q^r |V|$. Using the fact that the strong arithmetic dimension is at least the arithmetic dimension, it thus suffices to show that if $|V + t \cdot V| \leq q^r |V|$ then V has strong arithmetic dimension of at most r. We show this by induction on r. The case $r = 0$ is trivial, since the trivial lower bound $|V + t \cdot V| \geq 2 |V| - 1$ forces $V = \{0\}$. We shall hence assume that $r \geq 1$ and that the claim has been proved for $r' < r$. Let

$$V = \mathbb{F}_q[t]_1 \cdot x_1 \oplus \cdots \oplus \mathbb{F}_q[t]_1 \cdot x_\ell$$

be a decomposition of the type provided by Lemma 4.4, and for $1 \leq s \leq \ell$ let

$$V_{\leq s} = \mathbb{F}_q[t]_1 \cdot x_1 \oplus \cdots \oplus \mathbb{F}_q[t]_1 \cdot x_s.$$

We observe that if $x \in V_{\leq s} \backslash \{0\}$ then $\deg x_1 \leq \deg x \leq \deg x_s$. Furthermore, if $x \in V \backslash V_{\leq s}$ then $\deg x > \deg y$ for all $y \in V_{\leq s}$. Let $1 \leq s \leq \ell$ be maximal such that $V_{\leq s}$ has strong arithmetic dimension of at most r. If $s = \ell$ then the claim follows immediately, so suppose that $1 \leq s < \ell$ and that $V_{\leq s}$ has strong arithmetic dimension of $1 \leq r' \leq r$. We must have $r' = r$, or this contradicts the maximality of s by considering the decomposition

$$V_{\leq s+1} = V_{\leq s} \oplus \mathbb{F}_q[t]_1 \cdot x_{s+1} = \mathbb{F}_q[t]_{d_1} \cdot y_1 \oplus \cdots \oplus \mathbb{F}_q[t]_{d_{r'}} \cdot y_{r'} \oplus \mathbb{F}_q[t]_1 \cdot x_{s+1}.$$

Hence $V_{\leq s}$ has strong arithmetic dimension r, whence we have some decomposition

$$V_{\leq s} = \mathbb{F}_q[t]_{d_1} \cdot y_1 \oplus \cdots \oplus \mathbb{F}_q[t]_{d_r} \cdot y_r,$$

such that $d_1 + \deg y_1 < \cdots < d_r + \deg y_r$; furthermore, since $\deg x_s$ is maximal over all $x \in V_{\leq s}$, and so is $t^{d_r-1}y_r$, we have that $d_r + \deg y_r = \deg x_s + 1 < \deg tx_\ell$. If $|V_{\leq s} + t \cdot V_{\leq s}| < q^r |V|$ then by induction $V_{\leq s}$ has strong arithmetic dimension of less than r, which is a contradiction as noted above. It follows that $|V_{\leq s} + t \cdot V_{\leq s}| = q^r |V_{\leq s}|$ and hence, since the $\dim V_{\leq s} + r$ many elements

$$\{y_1, \ldots, t^{d_1}y_1, y_2, \ldots, y_r, \ldots, t^{d_r}y_r\}$$

span the $\mathbb{F}_q$-vector space $V_{\leq s} + t \cdot V_{\leq s}$, they are also linearly independent over $\mathbb{F}_q$. In particular, if

$$\alpha_1 t^{d_1}y_1 + \cdots + \alpha_r t^{d_r}y_r \in V_{\leq s}$$

with $\alpha_i \in \mathbb{F}_q$ then we must have $\alpha_i = 0$ for $1 \leq i \leq r$.

We now use the hypothesis $|V + t \cdot V| \leq q^r |V|$ to observe that the $\dim V + r + 1$ elements

$$\{y_1, \ldots, t^{d_1}y_1, \ldots, t^{d_r}y_r\} \cup \{x_{s+1}, \ldots, x_\ell\} \cup \{tx_\ell\} \subset V + t \cdot V$$

are linearly dependent over $\mathbb{F}_q$. Since the degree of tx_ℓ is strictly larger than that of all elements of $V \cup t \cdot V_{\leq s}$ there must exist $\alpha_i \in \mathbb{F}_q$ for $1 \leq i \leq r$, not identically zero, such that

$$z = \alpha_1 t^{d_1}y_1 + \cdots + \alpha_r t^{d_r}y_r \in V.$$

If $\alpha_r = 0$ then $\deg z \leq d_{r-1} + \deg y_{r-1} < d_r + \deg y_r = \deg x_s + 1$, and hence $z \in V_{\leq s}$, which contradicts the above. Hence we must have $\deg z = d_r + \deg y_r = \deg x_s + 1$, and hence $z \in V_{\leq s+1}$. Let $1 \leq i \leq r$ be such that $\alpha_i \neq 0$ and d_i is minimal, and let $z = t^{d_i}y$, say. In particular we have that $t^j y \in V_{\leq s}$ for $0 \leq j < d_i$ and $t^{d_i}y \in V_{\leq s+1}$. We claim that $V_{\leq s+1}$ has strong arithmetic dimension r, with a suitable decomposition provided by

$$\mathbb{F}_q[t]_{d_1} \cdot y_1 \oplus \cdots \oplus \mathbb{F}_q[t]_{d_{i-1}} \cdot y_{i-1} \oplus \mathbb{F}_q[t]_{d_{i+1}} \cdot y_{i+1} \oplus \cdots \oplus \mathbb{F}_q[t]_{d_r} \cdot y_r \oplus \mathbb{F}_q[t]_{d_i+1} \cdot y. \quad (4.1)$$

This contradicts the maximality of s and completes the proof. The vector space (4.1) is contained in $V_{\leq s+1}$, and since $\dim V_{\leq s+1} = \dim V_{\leq s} + 1$, comparing dimensions shows that the vector spaces are equal, provided only that this sum is indeed direct. If the sum is not direct then we have $a_j \in \mathbb{F}_q[t]_{d_j}$, not identically zero, and $\beta \in \mathbb{F}_q$ such that

$$a_1 y_1 + \cdots + a_{i-1}y_{i-1} + a_{i+1}y_{i+1} + \cdots + a_r y_r + a_i y + \beta t^{d_i} y = 0.$$

If $\beta = 0$ then this contradicts the orthogonality of the original decomposition of $V_{\leq s}$, and so $\beta \neq 0$. Since the degree of the final summand is $d_i + \deg y = \deg z = d_r + \deg y_r$ which is strictly larger than the degree of all the other summands, the left hand side cannot be zero, which is a contradiction. Finally, the fact that this decomposition is a witness to $V_{\leq s+1}$ having strong arithmetic dimension r follows from the fact that $d_i + 1 + \deg y = \deg z + 1 > d_r + \deg y_r$. □

We now prove a similar statement, weakening the condition that V be a vector space to merely having very tight control on the growth of its additive sets.

Lemma 4.6. *Let $X \subset \mathbb{F}_q[t]$ be a finite set such that $\mathbb{F}_q \cdot X = X$. If $|nX| \leq \ell_1 |X|$ and $|n(X + t \cdot X)| \leq \ell_2 |X|$ for all $n \geq 0$ then there is a finite $\mathbb{F}_q$-vector space $V \subset \langle X \rangle$ such that $X \subset V$, and furthermore $|V| \leq \ell_1 |X|$ and $|V + t \cdot V| \leq \ell_2 |X|$.*

Proof. Since $(n+1)X \supset nX$ for all $n \geq 0$, and $|nX|$ is bounded above by a constant independent of n there exists some $n_0 \geq 0$ such that $mX = n_0 X$ for all $m \geq n_0$. We claim that $V = n_0 X$ is a $\mathbb{F}_q$-vector space; it is clearly closed under dilations from $\mathbb{F}_q$, and if $v, w \in V$ then $v + w \in 2n_0 X = n_0 X = V$, so that V is closed under addition. The lemma now follows trivially. □

The following theorem is an immediate consequence of Theorem 4.5 and Lemma 4.6.

Theorem 4.7. *Let $X \subset \mathbb{F}_q[t]$ be a finite set such that $\mathbb{F}_q \cdot X = X$. If $|nX| \leq K_1 |X|$ and $|n(X + t \cdot X)| \leq K_2 |X|$ for all $n \geq 0$ then there is an arithmetic space $V \subset \langle X \rangle$ of arithmetic dimension of at most $\log_q(K_2)$ such that $X \subset V$ and $|V| \leq K_1 |X|$.*

4.2 Random sampling and covering lemmata

4.2.1 Random sampling

The power of random sampling in arithmetic combinatorics was demonstrated by Croot and Sisask [17], who used it to create large sets of L^p almost-periods for convolutions; most crucially, their results have only a polynomial dependence on p where traditional Fourier analytic techniques have an exponential dependence. Their ideas have seen many applications, and played a crucial part in Sanders' quantitative improvement of Roth's theorem, as discussed in Chapter 3.

Sanders also used these ideas to give a dramatic quantitative improvement of Freiman's theorem in [62]. An attractive feature of the Croot-Sisask method is that since it samples from physical space, rather than Fourier space, it is very robust and hence we can control arithmetic behaviour more exotic than simple addition. In this section we will adapt the proof of the Croot-Sisask theorem to prove a version suitable for our purposes, and construct a set X such that $X + t \cdot X$ forms a large set of almost-periods.

In this section we keep our results general, and unless otherwise stated G is an arbitrary abelian group with the discrete topology.

We will use the random sampling argument of Croot, Łaba and Sisask [16], and give an almost self-contained proof, appealing only to the classical Marcinkiewicz-Zygmund inequality [44].

Lemma 4.8 (Marcinkiewicz-Zygmund inequality)**.** *Let $p \geq 2$ and X be any random variable such that $\mathbb{E}|X - \mathbb{E}X|^p < \infty$. If $X_1, \dots, X_k$ are independently sampled from X then*

$$\mathbb{E}\left|\frac{1}{k}\sum_{j=1}^{k} X_j - \mathbb{E}X\right|^p \leq \left(\frac{Cp}{k}\right)^{p/2} \mathbb{E}|X - \mathbb{E}X|^p,$$

where $C > 0$ is an absolute constant.

We now use this to prove the following general sampling lemma of [16].

Lemma 4.9. *Let $p \geq 2$ and $g_1, \dots, g_n \in L^p(G)$. There is some $k \ll \eta^{-2}p$ such that*

$$\left\|\frac{1}{n}\sum_{i=1}^{n} g_i - \frac{1}{k}\sum_{j=1}^{k} g_{\sigma_j}\right\|_p < \eta \max_{1\leq i\leq n} \|g_i\|_p$$

for at least $n^k/2$ many $\sigma \in [n]^k$.

Proof. Let h be chosen uniformly at random from $g_1, \dots, g_n$. For any $x \in G$

$$\mathbb{E}h(x) = \frac{1}{n}\sum_{j=1}^{n} g_j(x) = f(x),$$

say. Let $h_1, \dots, h_k$ be independently chosen copies of h, where k is some integer to be chosen later. By Lemma 4.8, for any $x \in G$,

$$\mathbb{E}\left|\frac{1}{k}\sum_{j=1}^{k} h_j(x) - f(x)\right|^p \leq \left(\frac{Cp}{k}\right)^{p/2} \mathbb{E}|h(x) - f(x)|^p.$$

In particular, summing over all $x \in G$ yields

$$\sum_{x\in G} \mathbb{E}\left|\frac{1}{k}\sum_{j=1}^{k} h_j(x) - f(x)\right|^p \le \left(\frac{Cp}{k}\right)^{p/2} \mathbb{E}\,\|h-f\|_p^p .$$

By the triangle inequality, if $M = \max_{1\le i\le n} \|g_i\|_p$ then $\|h-f\|_p^p \le (2M)^p$, whence

$$\mathbb{E}\left\|\frac{1}{k}\sum_{j=1}^{k} h_j - f\right\|_p^p \le \left(\frac{CpM^2}{k}\right)^{p/2}.$$

In particular we can choose $k \ll p\eta^{-2}$ such that

$$\mathbb{E}\left\|\frac{1}{k}\sum_{j=1}^{k} h_j - f\right\|_p^p \le (\eta M)^p/2.$$

By Markov's inequality it follows that with probability at least $1/2$

$$\left\|\frac{1}{k}\sum_{j=1}^{k} h_j - f\right\|_p^p < (\eta M)^p,$$

and the proof is complete. □

We now apply Lemma 4.9 to a convolution to prove the following version of the Croot-Sisask theorem, the essence of which first appeared in [17]. It is convenient here to use the translation operator τ_x, defined as $\tau_x f(y) = f(y-x)$.

Theorem 4.10. *Let $p \ge 2$, $S \subset G$ be some finite set, and $g \in L^p(G)$. For all $\eta \in (0,1]$ there are $k \ll p\eta^{-2}$ and $\mathcal{L} \subset S^k$ such that $|\mathcal{L}| \ge |S|^k/2$, and if $(x,\dots,x) \in \mathcal{L}-\mathcal{L}$ then*

$$\|\tau_x(S*g) - S*g\|_p \le \eta\,|S|\,\|g\|_p .$$

Proof. We apply Lemma 4.9 to the decomposition $S*g = \sum_{y\in S}\tau_y g$. This yields an integer $k \ll \eta^{-2}p$ and a set $\mathcal{L} \subset S^k$ such that $|\mathcal{L}| \ge |S|^k/2$, and if $\mathbf{y} \in \mathcal{L}$ then

$$\left\|S*g - \frac{|S|}{k}\sum_{i=1}^{k}\tau_{y_i} g\right\|_p < \frac{\eta}{2}\,|S|\,\|g\|_p .$$

It follows that if $(z,\dots,z) \in \mathcal{L}-\mathcal{L}$, say $z = y_i - y_i'$ for $1 \le i \le k$, then by the triangle inequality

$$\|\tau_z(S*g) - S*g\|_p \le \left\|\tau_z(S*g) - \frac{|S|}{k}\sum_{i=1}^{k}\tau_{y_i} g\right\|_p + \left\|S*g - \frac{|S|}{k}\sum_{i=1}^{k}\tau_{y_i} g\right\|_p .$$

The theorem follows since the translation operator τ is an isometry and $\tau_{-z}\tau_{y_i} = \tau_{y_i'}$ for $1 \leq i \leq k$. □

To generate a set of almost-periods from the conclusion of Theorem 4.10 we need to be able to produce some set X such that $(x, \ldots, x) \in \mathcal{L} - \mathcal{L}$ for all $x \in X$; this can be an awkward matter since although $\mathcal{L}$ is dense inside S^k it need not resemble a product set itself. Fortunately, if S has a reasonable amount of additive structure then this does force $\mathcal{L} - \mathcal{L}$ to contain a fairly large diagonal set.

Lemma 4.11. *Let S and T be finite subsets of G such that $|S + T| \leq K|S|$. Let $\mathcal{L} \subset S^k$ be such that $|\mathcal{L}| \geq |S|^k/2$. Then there exists $u \in T$ and a set $X \subset T - u$ such that*

$$|X| \geq (2K)^{-k}|T|$$

and $(x, \ldots, x) \in \mathcal{L} - \mathcal{L}$ for all $x \in X$.

Proof. Let $\tilde{T} = \{(t, \cdots, t) : t \in T\} \subset T^k$, and observe that $\left|\mathcal{L} + \tilde{T}\right| \leq |S + T|^k \leq K^k |S|^k$. In particular, by the Cauchy-Schwarz inequality,

$$\left\|\mathcal{L} * \tilde{T}\right\|_2^2 K^k |S|^k \geq \left(\sum_x \mathcal{L} * \tilde{T}(x)\right)^2 = |T|^2 |\mathcal{L}|^2 .$$

Furthermore,

$$\left\|\mathcal{L} * \tilde{T}\right\|_2^2 = \langle \tilde{T} * (-\tilde{T}), \mathcal{L} * (-\mathcal{L})\rangle = \sum_{t_1, t_2 \in \tilde{T}} \mathcal{L} * (-\mathcal{L})(t_1 - t_2),$$

and hence by averaging there are $t \in T$ and $X \subset T - t$ such that

$$|X| \geq K^{-k}|\mathcal{L}|\,|S|^{-k}|T| \geq K^{-k}|T|/2$$

and for all $x \in X$ we have $\mathcal{L} * (-\mathcal{L})(\tilde{x}) > 0$ as required. □

Lemma 4.11 is robust enough to be iterated to handle more general situations. For the following lemma we suppose that G is an R-module for some ring R.

Lemma 4.12. *Let S and T be finite subsets of G and let $A \subset R^*$ be any finite set. Suppose that $K \geq 1$ is such that $|S + a \cdot T| \leq K|S|$ for all $a \in A$. Let $\mathcal{L} \subset S^k$ be such that $|\mathcal{L}| \geq |S|^k/2$. There exists some $X \subset \langle T \rangle$ such that*

$$|X| \geq (2K)^{-k|A|}|T|$$

and $(ax, \ldots, ax) \in 2^{|A|}(\mathcal{L} - \mathcal{L})$ for all $x \in X$ and $a \in A$.

Proof. We use induction on $|A|$; the case $A = \emptyset$ is trivial. Suppose then that $|A| \geq 1$, and fix any $b \in A$. By induction there is some $u \in G$ and $X \subset \langle T \rangle$ such that $|X| \geq (2K)^{-(|A|-1)k} |T|$, and for all $a \in A \backslash \{b\}$ we have $(ax, \dots, ax) \in 2^{|A|-1}(\mathcal{L} - \mathcal{L})$. We now apply Lemma 4.11 with T replaced by $b \cdot X$. It follows that there exists some $bu' \in b \cdot X$ and $b \cdot X' \subset b \cdot (X - u')$ such that

$$|X'| = |b \cdot X'| \geq (2K)^{-k} |X| \geq (2K)^{-|A|k} |T| ,$$

and furthermore $(bx, \dots, bx) \in \mathcal{L} - \mathcal{L}$ for all $x \in X'$.

It remains to observe that if $x \in X'$ then $x \in X - X$, and hence $(ax, \dots, ax) \in 2^{|A|}(\mathcal{L} - \mathcal{L})$ for all $a \in A'$. □

Combining Lemma 4.12 with the set $\mathcal{L}$ produced by Lemma 4.10 immediately yields the following.

Theorem 4.13. *Let $S \subset G$ be a finite set and $g \in L^p(G)$. Let $T \subset G$ and $A \subset R^*$ be finite sets, and $K \geq 2$ be such that $|S + a \cdot T| \leq K |S|$ for all $a \in A$. Then there exists some $X \subset \langle T \rangle$ such that*

$$|X| \geq \exp\left(-O_{|A|}\left(p\eta^{-2} \log K\right)\right) |T| ,$$

and for all $x \in \sum_{a \in A} a \cdot X$

$$\|\tau_x(S * g) - S * g\|_p \leq \eta |S| \, \|g\|_p .$$

Proof. By Theorem 3.18 there are $k \ll_{|A|} p\eta^{-2}$ and $\mathcal{L} \subset S^k$ such that $|\mathcal{L}| \geq |S|^k / 2$, and if $(x, \dots, x) \in \mathcal{L} - \mathcal{L}$ then

$$\|\tau_x(S * g) - S * g\|_p \leq |A|^{-1} 2^{-|A|} \eta |S| \, \|g\|_p .$$

By the triangle inequality if $(z, \dots, z) \in 2^{|A|}(\mathcal{L} - \mathcal{L})$ then

$$\|\tau_z(S * g) - S * g\|_p \leq |A|^{-1} \eta |S| \, \|g\|_p .$$

By Lemma 4.12 there exist $u \in G$ and $X \subset T - u$ such that

$$|X| \geq \exp(-O_{|A|}\left(p\eta^{-2} \log K\right)) |T| ,$$

and for any $a \in A$ and all $x \in a \cdot X$

$$\|\tau_x(S * g) - S * g\|_p \leq |A|^{-1} \eta |S| \, \|g\|_p .$$

The result follows from another application of the triangle inequality. □

4.2.2 Covering lemmata

We shall also need some well-known covering lemmata due to Ruzsa and Chang. We will here prove a more abstract covering lemma of which both shall be simple corollaries; the generality of Lemma 4.14 is more than we require for the main results of this chapter, but we will indulge ourselves in a small digression.

Let

$$\Sigma_k A = \left\{ \sum_{x \in B} x : B \subset A \text{ and } 1 \leq |B| \leq k \right\}.$$

We say that A is k-dissociated if $\Sigma_k A$ has the maximum possible size (note that this differs from the concept of Γ-dissociativity used in Chapter 3). In other words, A is k-dissociated if whenever we have $B_1, B_2 \subset A$ with $1 \leq |B_1| \leq |B_2| \leq k$ and $\sum_{x \in B_1} x = \sum_{y \in B_2} y$, then $B_1 = B_2$. Finally, for any $K \geq 1$ let $f_k(K)$ be the minimal n such that

$$\sum_{i=1}^{k} \binom{n}{k} > K.$$

We observe that the left hand side is precisely the cardinality of a k-dissociated set of size n.

The proof of the following general covering lemma is a generalisation of the proofs of the covering lemmata of Ruzsa [57] and Chang [13].

Lemma 4.14. *Let S and A be finite subsets of G and let $k_1, \ldots, k_r, K_1, \ldots, K_r \geq 1$ be parameters such that*

$$|S + \Sigma_{k_1} A + \cdots + \Sigma_{k_r} A| \leq K_1 \cdots K_r |S|.$$

Then there exist $1 \leq s \leq r$ and sets $T_1, \ldots, T_s \subset A$ such that T_i is k_i-dissociated for $1 \leq i \leq s$,

$$A \subset S - S + \Sigma_{k_1} T_1 - \Sigma_{k_1} T_1 + \cdots + \Sigma_{k_s} T_s - \Sigma_{k_{s-1}} T_s,$$

$|\Sigma_{k_i} T_i| = f_{k_i}(K_i)$ for $1 \leq i < s$ and $|\Sigma_{k_s} T_s| \leq K_j$.

Proof. We say that $1 \leq j \leq r$ is good if there exist $T_1, \ldots, T_j \subset A$ such that for $1 \leq i \leq j$ the set T_i is k_i-dissociated and $|\Sigma_{k_i} T_i| > K_i$, and furthermore if $S_0 = S$ and $S_i = S_{i-1} + \Sigma_{k_i} T_i$ for $1 \leq i \leq j$ then $|S_i| = |S_{i-1}| \, |\Sigma_{k_i} T_i|$. In particular, for $1 \leq i \leq j$ we have $|S_i| > K_i |S_{i-1}|$, and hence if j is good then

$$K_1 \cdots K_j |S| < |S_j| \leq \left| S + \Sigma_{k_1} A + \cdots + \Sigma_{k_j} A \right|.$$

It follows that r is not good; let $1 \leq s \leq r$ be minimal such that s is not good. Since $s-1$ is good there exist $T_1, \dots, T_{s-1} \subset A$ such that T_i is k_i-dissociated for $1 \leq i < s$ and if $S_0 = S$ and $S_i = S_{i-1} + \Sigma_{k_i} T_i$ then $|S_i| = |S_{i-1}|\,|\Sigma_{k_i} T_i|$ for $1 \leq i < s$. We further have that $|\Sigma_{k_i} T_i| > K_i$; by removing elements from each T_i if necessary, we may suppose that $|\Sigma_{k_i} T_i| = f_{k_i}(K_i)$. Let $T_s \subset A$ be a maximal set which is both k_s-dissociated and such that $|S_{s-1} + \Sigma_{k_s} T_s| = |S_{s-1}|\,|\Sigma_{k_s} T_s|$.

Let $x \in A \backslash T_s$ and let $T' = T_s \cup \{x\}$. By maximality either T' is k_s-dissociated and $|S_{s-1} + \Sigma_{k_s} T'| < |S_{s-1}|\,|\Sigma_{k_s} T'|$, or there are distinct subsets $U, V \subset T'$ of size at most k_s such that

$$\sum_{u \in U} u = \sum_{v \in V} v.$$

In the former case there must be some $s_1, s_2 \in S_{s-1}$ and $r_1, r_2 \in \Sigma_{k_s} T'$ such that $s_1 + r_1 = s_2 + r_2$. Since $|S_{s-1} + \Sigma_{k_s} T_s| = |S_{s-1}|\,|\Sigma_{k_s} T_s|$, however, at least one of the elements in the sum forming either r_1 or r_2 must be x, and hence $x \in S_{s-1} - S_{s-1} + \Sigma_{k_s} T_s - \Sigma_{k_s - 1} T_s$. A similar conclusion follows from the second situation, using the fact that T_s is k_s-dissociated. □

The above is rather abstract and general. For our purposes there are two simple corollaries that will be useful, both of which have already appeared in the literature. The simplest demonstration of Lemma 4.14 is to take $r = 1 = k_1$ when we recover the well-known covering lemma of Ruzsa [57].

Lemma 4.15. *Suppose that $|S + A| \leq K\,|S|$. Then there exists $T \subset A$ with $|T| \leq K$ such that*

$$A \subset S - S + T.$$

Alternatively we may choose $r = 1$ and $k_1 = k$ to deduce the following lemma due to Chang [13].

Lemma 4.16. *Suppose that $|S + kA| < 2^k\,|S|$ and $0 \in A$. Then there exists a dissociated set T with $|T| < k$ such that*

$$A \subset S - S + \langle T \rangle - \langle T \rangle.$$

Proof. This follows from Lemma 4.14 and the trivial observation that $\Sigma_k A \subset kA$. It only remains to observe that if T is dissociated and $|T| \geq k$ then $2^k \leq |\Sigma_k T|$. □

4.3 Covering structured sets

In this section we adapt the arguments of Sanders [63] to show that if $|A+A|$ and $|A+t\cdot A|$ are both small relative to $|A|$ then A can be efficiently covered by a set suitable for an application of Theorem 4.7. The arguments of [63] represent the best-known approach to inverse sumset theorems at the time of writing, and are a synthesis of arguments by Sanders, Schoen and Konyagin. The adaptation to a more general arithmetic inverse theorem presented here is straightforward, but we present the proofs in full detail. The main point of interest is that there are now two parameters, $|A+A|/|A|$ and $|A+t\cdot A|/|A|$, and it emerges that the method is asymmetric in its treatment of these parameters; the real cost is in the doubling parameter $|A+A|/|A|$, and with this fixed we are able to obtain polynomial type bounds in terms of the parameter $|A+t\cdot A|/|A|$.

We use the notation $A^\circ = A-A$ for any finite $A\subset \mathbb{F}_q[t]$. The following lemma is a generalisation of Proposition 4.2 from [62]. The field $\mathbb{F}_q$ is fixed throughout this section; in particular, all implicit constants may depend on q, the size of the finite field.

Lemma 4.17. *Let A, S and T be any finite subsets of $\mathbb{F}_q[t]$ such that $\mathbb{F}_q\cdot T = T$. Let $K_1, K_2 \geq 4$ be such that*

$$|A+S| \leq K_1|A| \text{ and } \max(|S+T|, |S+t\cdot T|) \leq K_2|S|.$$

Then for any $m\geq 1$ there is a set $X\subset\langle T\rangle$ such that $\mathbb{F}_q\cdot X = X$,

$$|X| \geq \exp(-O(m^2\log K_1\log K_2))|T|$$

and we have $m(X+t\cdot X)\subset A^\circ + S^\circ$.

Proof. Let $\eta>0$ and $p\geq 2$ be parameters to be chosen later. By Theorem 4.13 there exists a set X' such that

$$|X'| \geq \exp(-O(p\eta^{-2}\log K_2))|T|,$$

and, if $X=\sum_{a\in\mathbb{F}_q} a\cdot X'$, then for all $x\in m(X+t\cdot X)$,

$$\|\tau_x(A+S)*(-S) - (A+S)*(-S)\|_p \leq m\eta|A+S|^{1/p}|S|.$$

In particular, by Hölder's inequality, for all $x\in m(X+t\cdot X)$,

$$|\langle \tau_x(A+S)*(-S), A\rangle - \langle (A+S)*(-S), A\rangle| \leq m\eta|A+S|^{1/p}|S||A|^{1-1/p}.$$

Since $\langle (A+S)*(-S), A\rangle = |A|\,|S|$ this implies that for all $x \in m(X+t\cdot X)$

$$|(A+S)*(-S)*(-A)(x) - |A|\,|S|| \leq m\eta K_1^{1/p}\,|A|\,|S|\,.$$

Choosing the parameters $p = \log K_1$ and $\eta = 1/2^3 m$, say, implies that $m(X+t\cdot X) \subset A^\circ + S^\circ$ as required. □

Lemma 4.17 is already powerful enough to prove a strong inverse sumset theorem, and the sumset analogue was a key component of the quantitative breakthrough of Sanders [62]. To demonstrate how we first show how containment of a large sumset can be combined with the strong structural result Theorem 4.7 to prove an arithmetic inverse result.

Lemma 4.18. *There exists an absolute constant $C > 0$ such that the following holds. Let A and S be any finite subsets of $\mathbb{F}_q[t]$ such that $|A+S| \leq K\,|A|$ for some $K \geq 4$. Let $X \subset \mathbb{F}_q[t]$ be a finite set such that $\mathbb{F}_q \cdot X = X$ and $|X| \geq L^{-1}\,|A|$ for some $L \geq K^{8C}$. Finally, suppose that*

$$\lceil C\log K\rceil (X + t\cdot X) \subset A^\circ + S^\circ.$$

Then there is an arithmetic space $V \subset \langle X\rangle$ of arithmetic dimension $d \ll \log L$ such that A is $L^{O(1)}$-covered by V and $|V| \leq L^{O(1)}\,|A|$.

Proof. Let $C > 0$ be some constant to be chosen later, and $m = \lceil C\log K\rceil$. We first observe that by the Plünnecke-Ruzsa estimates for any $n \geq 1$ we have $|n(A^\circ + S^\circ)| \leq K^{O(n)}\,|A|$. In particular, for any $n \geq 1$, since $3mn(X+t\cdot X) + X \subset 4n(A^\circ + S^\circ)$ we have

$$\begin{aligned} |3mn(X+t\cdot X) + X| &\leq \exp(O(n\log K))\,|A| \\ &\leq \exp(O(n\log K))L\,|X| \\ &< 2^{mn}\,|X|\,, \end{aligned}$$

provided C is sufficiently large and $\log L \leq 2^{-3}C(\log K)n$. By the hypothesis on L we may choose some suitable $1 \leq n \ll \log L/\log K$, so that $mn \ll \log L$. It follows from Lemma 4.16 that there is a $\mathbb{F}_q$-vector space T of dimension $O(\log L)$ such that

$$3(X+t\cdot X) \subset T + 2X.$$

We now consider the set $X' = 2X$. For any $n \geq 1$ we have, by induction, $n(X'+t\cdot X') \subset T + X'$, and hence

$$|n(X'+t\cdot X')| \leq L^{O(1)}\,|X'|\,.$$

Furthermore,

$$|X + A| \leq |A + A^\circ + S^\circ| \leq K^{O(1)} |A| \leq \exp(O(\log K + \log L)) |X|,$$

and hence by Lemma 4.15 the set A is $L^{O(1)}$-covered by X'. By Theorem 4.7 there is an arithmetic space $V \subset \mathbb{F}_q[t]$ of arithmetic dimension $O(\log L)$ such that $X' \subset V$ and $|V| \leq L^{O(1)} |X'|$, and the proof is complete. □

Lemma 4.18 shows that, to obtain good quantitative bounds in our desired arithmetic inverse result, the important thing to control is how large a set X we can take such that the $O(\log K)$-fold sumset of $X + t \cdot X$ is contained in $A^\circ + S^\circ$. Lemma 4.17 immediately provides a suitable set X, and indeed, does so with very good bounds. Using this we can prove the following result; the bounds match those obtained by Sanders [62] for the analogous result over $\mathbb{Z}$, and we again stress that the covering arguments we use in converting the strong inverse result Theorem 4.7 to a more general inverse result are those developed by Sanders.

The following result will be improved immediately afterwards, but we include it here as a demonstration of the power of Lemma 4.17.

Theorem 4.19. *Let $A \subset \mathbb{F}_q[t]$ be a finite set and $K_1, K_2 \geq 4$ be such that*

$$|A + A| \leq K_1 |A| \text{ and } \sum_{\alpha \in \mathbb{F}_q} |A + \alpha \cdot A| + |A + t \cdot A| \leq K_2 |A|.$$

Then there is a $d \ll (\log K_1)^3 \log K_2$ and an arithmetic space of arithmetic dimension at most d such that A is $\exp(O(d))$-covered by V and $|V| \leq \exp(O(d)) |A|$.

Proof. Let $T = \sum_{a \in \mathbb{F}_q} a \cdot A$. It is clear that $a \cdot T = T$ for all $a \in \mathbb{F}_q$, and by the Plünnecke-Ruzsa estimates we certainly have that

$$\max(|A + T|, |A + t \cdot T|) \leq K_2^{O(1)} |A|.$$

By Lemma 4.17 for any $m \ll \log K_1$ there is a set $X \subset \mathbb{F}_q[t]$ such that $\mathbb{F}_q \cdot X = X$,

$$|X| \geq \exp(-O((\log K_1)^3 \log K_2)) |A|$$

and $m(X + t \cdot X) \subset 2A^\circ$. The theorem follows from Lemma 4.18 with $L \ll \exp(O((\log K_1)^3 \log K_2))$. □

Theorem 4.3 follows immediately. It is possible, however, to improve the dependence on K_1, by a stunning application of the pigeonhole principle given by Konyagin and expounded by Sanders [63], as we shall do now.

We first prove the following technical lemma.

Lemma 4.20. *Let A, S and T be any finite subsets of $\mathbb{F}_q[t]$ such that $\mathbb{F}_q \cdot T = T$. Let $K_1, K_2 \geq 4$ be such that*

$$|A+S| \leq K_1 \min(|A|, |S|) \text{ and } \max(|S+T|, |S+t\cdot T|) \leq K_2 |S|.$$

Then for all integers $m \geq 1$ there are sets S' and $T' \subset \langle T \rangle$ such that $\mathbb{F}_q \cdot T' = T'$,

$$|T'| \geq \exp(-O(m^2 \log K_1 \log K_2))\,|T|$$

such that $S \subset S' \subset A^\circ + S^\circ + S$ and

$$|S' + T' + t\cdot T'| \leq K_1^{O(1/m)} |S'|.$$

Proof. An application of Lemma 4.17 yields some T' with $\mathbb{F}_q \cdot T' = T'$,

$$|T'| \geq \exp(-O(m^2 \log K_1 \log K_2))\,|T|$$

and $m(T' + t\cdot T') \subset A^\circ + S^\circ$. In particular, by the Plünnecke-Ruzsa estimates,

$$|S + m(T' + t\cdot T')| \leq |A^\circ + S^\circ + S| \leq K_1^{O(1)} |S|.$$

By the pigeonhole principle there exists $0 \leq l < m$ with

$$|S + l(T' + t\cdot T') + T' + t\cdot T'| \leq K_1^{O(1/m)} |S + l(T' + t\cdot T')|.$$

The proof is complete, letting $S' = S + l(T' + t\cdot T')$. □

A single application of Lemma 4.20, with $m = 1$, replaces the parameter K_2 by K_1, while reducing the size of T by a factor of $\exp(-O(\log K_1 \log K_2))$, which already allows the $(\log K_1)^3 \log K_2$ factor in Theorem 4.19 to be improved to $(\log K_1)^4 + (\log K_1)(\log K_2)$. The real power of Lemma 4.20, however, lies in iteration, which will allow us to reduce the $(\log K_1)^4$ here to just shy of $(\log K_1)^3$. We now prove the key lemma, which first replaces K_2 by K_1 and then repeatedly applies Lemma 4.20 to reduce the size of K_1 by an exponential factor at each stage, until we halt in the best possible situation with $K_2 \ll 1$.

Lemma 4.21. *Let A and B be finite subsets of $\mathbb{F}_q[t]$ such that $\mathbb{F}_q \cdot B = B$. Let $K_1, K_2 \geq 4$ be such that*

$$|A+A| \leq K_1 |A| \text{ and } \max(|A+B|, |A+t\cdot B|) \leq K_2 |A|.$$

Then there is a positive integer $r \ll (\log\log K_1)^{\log_2 3}$, a set $A-A \subset S \subset r(A-A)$, and a set $T \subset \langle B\rangle$ such that $\mathbb{F}_q \cdot T = T$,

$$|S+T+t\cdot T| \ll |S|,$$

and

$$|T| \geq \exp\left(-O\left((\log K_1)^3 (\log\log K_1)^{3\log_2 3} + \log K_1 \log K_2\right)\right)|B|.$$

Proof. Let m_i be some sequence of positive integers to be chosen later, c be some absolute constant (though it may depend on q) to be chosen later, and $n_i = (3^{i+1}-1)/2$. Define the sequence ρ_i by letting $\rho_0 = c\log K_2$ and $\rho_i = cn_i \log K_1 m_i^{-1}$ for $i \geq 1$. We shall prove by induction that for all $i \geq 0$ we can find sets S_i and T_i such that $\mathbb{F}_q \cdot T_i = T_i$,

1. $A^\circ \subset S_i \subset n_i A^\circ$,
2. $|S_i + T_i + t\cdot T_i| \leq \exp(\rho_i)|S_i|$, and
3. $|T_i| \geq \exp\left(-O\left(\sum_{j=0}^{i-1} 3^i m_{i+1}^2 \rho_i\right)\log K_1\right)|B|$.

We begin the induction by letting $S_0 = A^\circ$ and $T_0 = B$, and the required estimates follow from the Plünnecke-Ruzsa estimates, which imply that

$$|A - A + B + t\cdot B| \leq K_2^{O(1)}|A|,$$

and hence the second condition follows provided c is sufficiently large. Suppose now that $i \geq 0$ and we have S_i and T_i as above. We apply Lemma 4.20 with the parameters m_{i+1}, S_i and T_i. By the Plünnecke-Ruzsa estimates we have

$$|A + S_i| \leq |(n_i+1)A^\circ| \leq K_1^{O(3^i)}|A|,$$

and hence Lemma 4.20 yields S_{i+1} and T_{i+1} such that $\mathbb{F}_q \cdot T_{i+1} = T_{i+1}$ and

$$|T_{i+1}| \geq \exp\left(-O\left(3^i m_{i+1}^2 \rho_i\right)\log K_1\right)|T_i|$$

as required. Furthermore, $S_{i+1} \subset (1+3n_i)A^\circ = n_{i+1}A^\circ$ and

$$|S_{i+1} + T_{i+1} + t \cdot T_{i+1}| \leq \exp(O(3^{i+1}m_{i+1}^{-1}) \log K_1) \, |S_{i+1}| \leq \exp(\rho_{i+1}) \, |S_{i+1}| \, ,$$

again provided c is chosen sufficiently large.

We now choose

$$m_i = 3^i \lceil (\log K_1)^{1-1/2^{i-1}} \rceil \text{ for all } i \geq 1.$$

In particular, $m_{i+1}^2 \rho_i \ll m_{i+1}^2 m_i^{-1} 3^i \log K_1 \ll 3^{2i} (\log K_1)^2$ for $i \geq 1$, and hence

$$|T_i| \geq \exp\left(-O\left(3^{3i}(\log K_1)^3 + \log K_1 \log K_2\right)\right) |B| \, .$$

Furthermore, if we choose $n \geq 1$ such that $2^{n-1} \geq \log\log K_1$ then $m_n \gg 3^n \log K_1$. It follows that $\rho_n \ll 1$ and hence

$$|S_n + T_n + t \cdot T_n| \ll |S_n|$$

as required. We may choose such an n such that $3^{3n} \ll (\log\log K_1)^{3\log_2 3}$ and the proof is complete. □

We now combine Lemmata 4.17, 4.18, and 4.21 to prove our strongest inverse result.

Theorem 4.22. *Let A and B be finite subsets of $\mathbb{F}_q[t]$ such that $\mathbb{F}_q \cdot B = B$. Let $K_1, K_2 \geq 4$ be such that*

$$|A + A| \leq K_1 \, |A| \, , \ |A + B + t \cdot B| \leq K_2 \, |A| \ \textit{and} \ |B| \geq K_3^{-1} \, |A| \, .$$

Then there is an arithmetic space $V \subset \langle B \rangle$ of arithmetic dimension K' such that A is $\exp(K')$-covered by V and $|V| \leq \exp(K') \, |A|$ for some

$$K' \ll (\log K_1)^3 (\log\log K_1)^{3\log_2 3} + \log K_1 \log K_2 + \log K_3.$$

Proof. By Lemma 4.21 there is an integer $r \ll (\log\log K_1)^2$, a set $A^\circ \subset S \subset rA^\circ$, a set $T \subset \langle B \rangle$ such that $\mathbb{F}_q \cdot T = T$ and

$$|S + T + t \cdot T| \ll |S|$$

and

$$|T| \geq \exp(-O((\log K_1)^3 (\log\log K_1)^{3\log_2 3} + \log K_1 \log K_2)) \, |B| \, .$$

We now apply Lemma 4.17, with m to be chosen later, to give a set $X \subset \langle B \rangle$ such that $\mathbb{F}_q \cdot X = X$,

$$\begin{aligned} |X| &\geq \exp(-O(m^2 \log K_1 (\log\log K_1)^{\log_2 3}))\, |T| \\ &\geq \exp\Big(- O\Big(m^2 \log K_1 (\log\log K_1)^{\log_2 3} \\ &\qquad + (\log K_1)^3 (\log\log K_1)^{3\log_2 3} + \log K_1 \log K_2\Big)\Big)\, |B| \end{aligned}$$

and $m(X + t\cdot X) \subset A^\circ + S^\circ$. In particular we have satisfied the hypotheses of Lemma 4.18 with $K = \exp(O(\log K_1 (\log\log K_1)^{\log_2 3}))$ and

$$\log L \ll m^2 \log K_1 (\log\log K_1)^{\log_2 3} + (\log K_1)^3 (\log\log K_1)^{3\log_2 3} + \log K_1 \log K_2 + \log K_3,$$

provided

$$m \gg \log K_1 (\log\log K_1)^{\log_2 3} + (\log L)^{1/3}.$$

A simple calculation shows that this is satisfied with our L for some

$$m \ll \log K_1 (\log\log K_1)^{\log_2 3} + (\log K_1 \log K_2)^{1/3} + (\log K_3)^{1/3}.$$

The conclusion now follows from the conclusion of Lemma 4.18, and the observation that by the Plünnecke-Ruzsa estimates we can, without loss of generality, assume that $K_1 \leq K_2^{O(1)}$. □

Theorem 4.22 is already a quantitatively strong arithmetic inverse theorem for $\mathbb{F}_q[t]$. The following corollary is, however, a simpler version which will be sufficient for many applications.

Corollary 4.23. *Let $A \subset \mathbb{F}_q[t]$ be a finite set and $K_1, K_2 \geq 4$ are such that*

$$|A + A| \leq K_1 |A|\,,\ |A + \alpha\cdot A| \leq K_2 |A| \ \text{ for all } \alpha \in \mathbb{F}_q \text{ and } |A + t\cdot A| \leq K_2 |A|\,.$$

Then there is an arithmetic space $V \subset \langle A \rangle$ of arithmetic dimension at most K' such that A is $\exp(K')$-covered by V and $|V| \leq \exp(K')\,|A|$ for some

$$K' \ll (\log K_1)^3 (\log\log K_1)^{3\log_2 3} + \log K_1 \log K_2.$$

Proof. We apply Theorem 4.22 with $B = \sum_{\alpha\in\mathbb{F}_q} \alpha\cdot A$ and observe that the necessary control on the additive growth follows from the Plünnecke-Ruzsa estimates. □

We remark that when $\mathbb{F}_q = \mathbb{F}_p$ for some prime p the hypothesis $|A + \alpha \cdot A| \leq K_2 |A|$ can be discarded by the Plünnecke-Ruzsa estimates, since $\alpha \cdot A \subset pA$ and we can replace K_2 by $\max(K_2, K_1^{O(p)})$ without affecting the strength of the conclusion.

For the application in the following section it is more convenient to use an alternative form, which follows by a standard covering argument. We recall that an $\mathbb{F}_q[t]$-arithmetic progression is a set of the form $a \cdot \mathbb{F}_q[t]_n + x$ for some $n \geq 0$ and $a, x \in \mathbb{F}_q[t]$.

Corollary 4.24. *Let $A \subset \mathbb{F}_q[t]$ be a finite set and $K_1, K_2 \geq 4$ are such that*

$$|A + A| \leq K_1 |A|\,, \ |A + \alpha \cdot A| \leq K_2 |A| \ \text{ for all } \alpha \in \mathbb{F}_q \text{ and } \ |A + t \cdot A| \leq K_2 |A|\,.$$

Then there exists a $\mathbb{F}_q$-vector space $T \subset \langle A \rangle$ of dimension at most K' such that $A - A + T$ contains an $\mathbb{F}_q[t]$-arithmetic progression P of size $|P| \gg |A|^{1/K'}$ for some

$$K' \ll (\log K_1)^3 (\log\log K_1)^{3\log_2 3} + \log K_1 \log K_2.$$

Proof. Let V be the arithmetic space given by Corollary 4.23, and K' be the given parameter. Since A is $\exp(K')$-covered by V, we have for any $k \geq 1$ the bound $|A + kV| \leq \exp(K')|V|$. By Lemma 4.16 there exists a finite $\mathbb{F}_q$-vector space $T \subset \langle A \rangle$ of dimension $O(K')$ such that $V \subset A - A + T$. The proof is complete after observing that since V has arithmetic dimension at most K' it contains an arithmetic progression of size at least $|V|^{1/K'}$. □

4.4 Transcendence and adding

We now present an application of the inverse theorems we have proven, which is an $\mathbb{F}_q[t]$-analogue of a problem first considered by Konyagin and Łaba [35]. Let A be a finite subset of $\mathbb{R}$ and ξ be any transcendental element. When $A \subset \mathbb{Q}$ it is clear that the only additive relations of the shape $a_1 + \xi a_2 = a_3 + \xi a_4$ are the trivial ones, and hence $|A + \xi \cdot A| = |A|^2$; indeed, this holds even if ξ is irrational. For general $A \subset \mathbb{R}$, we may have some non-trivial relations, but if ξ is transcendental then there should be relatively few, and hence we should be able to provide some non-trivial lower bound for $|A + \xi \cdot A|$.

Konyagin and Łaba proved that $|A + \xi \cdot A| \gg (\log |A|)^{1-o(1)} |A|$. Sanders [59] later observed that such lower bounds can be obtained by combining simple modelling arguments

with an inverse sumset result. Using such an argument with the sharpest known form of such an inverse result, Sanders [62] improved this lower bound to

$$|A + \xi \cdot A| \gg \exp(O((\log |A|)^c)) \, |A|$$

for some absolute constant $c > 0$. An example by Green, given in [35], shows that this is almost the best possible result. Namely, if one takes $A = \{\sum_{i=1}^{m} a_i \xi^i : 1 \leq a_i \leq n\}$ for suitable choices of n and m then one can show that

$$|A + \xi \cdot A| \ll \exp(O((\log |A|)^{1/2})) \, |A| \, .$$

We now consider this problem in the non-archimedean setting, with $\mathbb{R}$ replaced by $\mathfrak{k}$, which is the completion of the rational function field $\mathbb{F}_q(t)$. Since transcendence over $\mathbb{F}_q[t]$ is the obvious analogue of transcendence over $\mathbb{Z}$, one might hope for similar lower bounds to the above to hold for $|A + \xi \cdot A|$ when A is any finite subset of $\mathfrak{k}$ and $\xi \in \mathfrak{k}$ is any element transcendental over $\mathbb{F}_q[t]$.

A moment's thought shows that this is too ambitious; indeed, the analogue of the example outlined above already dashes our hopes. In particular, if $A = \{\sum_{i=0}^{n} a_i \xi^i : a_i \in \mathbb{F}_q\}$ then it is easy to show that $|A + \xi \cdot A| \leq q \, |A|$. Recalling that we take $\mathbb{F}_q$ to be fixed, this is essentially a constant upper bound, and hence no non-trivial lower bound can be given.

On examination of this example, however, some hope returns – for since such a set is contained in $\mathbb{F}_q[\xi]$ and ξ is transcendental over $\mathbb{F}_q[t]$ we have $|A + t \cdot A| \gg |A|^2$. Thus one might hope that if $A \subset \mathfrak{k}$ does not grow when added to its dilation by some transcendental element, then this forces growth when added to its dilation by t.

The construction above is easily adapted to such a situation. Consider the set $A = \{\sum_{i=1}^{n} a_i \xi^i : a_i \in \mathbb{F}_q[t]_m\}$. It is easy to show that $|A| = q^{nm}$ and $|A + t \cdot A| = q^n \, |A|$. Furthermore, $|A + \xi A| = q^m \, |A|$. It follows that if $|A + t \cdot A| = K_1 \, |A|$ and $|A + \xi \cdot A| = K_2 \, |A|$ then $(\log K_1)(\log K_2) \approx \log |A|$.

This should be compared to the case $A \subset \mathbb{R}$, when we study the single parameter K given by $|A + \xi \cdot A| = K \, |A|$ and our construction gives a set A with $(\log K)^2 \approx \log |A|$. Thus we see that in the analogous situation in $\mathbb{F}_q[t]$ there is a 'splitting' of the parameter K into two distinct parameters, and we may now ask for non-trivial lower bounds on the size of such parameters.

By combining our inverse results for $\mathbb{F}_q[t]$ with the argument of Sanders [59] we are able to prove such a result. Namely, we show that if $\xi \in \mathfrak{k}$ is transcendental over $\mathbb{F}_q[t]$ and $A \subset \mathfrak{k}$ is a finite set with $|A + t \cdot A| = K_1 |A|$ and $|A + \xi \cdot A| = K_2 |A|$ then either

$$\min(\log K_1, \log K_2) \gg (\log |A|)^{1/6-o(1)} \text{ or } (\log K_1)(\log K_2) \gg (\log |A|)^{1/2}.$$

In particular, we always have $\max(\log K_1, \log K_2) \gg (\log |A|)^{1/6-o(1)}$, which is a similar result to that obtained in the problem over $\mathbb{R}$.

Following [59], we first prove the following simple consequence of the Plünnecke-Ruzsa estimates.

Lemma 4.25. *Let $A \subset \mathfrak{k}$ be any finite set. If $\xi \in \mathfrak{k}\backslash\{0\}$ and $|A + \xi \cdot A| \leq K |A|$ then for all $\ell \geq 1$*

$$\left|A - A + \xi(A - A) + \cdots + \xi^{\ell-1}(A - A)\right| \leq K^{O(\ell)} |A| .$$

Proof. Let $A' = A - A$ and $B' = \xi(A - A)$. By the Plünnecke-Ruzsa estimates we have, for all $k \geq 1$, the upper bounds $|kB'| = |kA'| \leq K^{O(k)} |A|$. In particular, $|3A' - 3A'| \leq K^{O(1)} |A|$. By Corollary 4.15 there is some S with $|S| \leq K^{O(1)}$ such that

$$3A' - 2A' \subset A' - A' + S.$$

Similarly, there is some T with $|T| \leq K^{O(1)}$ such that $\xi \cdot A \subset A - A + T$, and hence if $T' = 2T - T$ then $|T'| \leq K^{O(1)}$ and $B' \subset A' - A' + T'$. We define a sequence of sets T_ℓ by $T_1 = \{a\}$ for some $a \in A'$ and $T_{\ell+1} = S + T' - T' + \xi \cdot T_\ell$. We claim that, for all $\ell \geq 1$,

$$A' + \xi \cdot A' + \cdots + \xi^{\ell-1} \cdot A' \subset A' - A' + T_\ell.$$

This certainly holds for $\ell = 1$. Suppose that this inclusion holds for $\ell \geq 1$; then

$$\begin{aligned} A' + \cdots + \xi^{\ell} \cdot A' &\subset A' + \xi \cdot (A' - A' + T_\ell) \\ &= A' + B' - B' + \xi \cdot T_\ell \\ &\subset 5A' + T' - T' + \xi \cdot T_\ell \\ &\subset A' - A' + S + T' - T' + \xi \cdot T_\ell \end{aligned}$$

as required. The result follows from the Plünnecke-Ruzsa estimates and the trivial bound $|T_\ell| \ll K^{O(\ell)}$. □

We now use Lemma 4.25 and Corollary 4.24 to prove the main result.

Theorem 4.26. *Let $A \subset \mathfrak{k}$ be a finite set. Suppose that ξ_1 and ξ_2 are algebraically independent over $\mathbb{F}_q$, and that $K_1, K_2 \geq 2$ are such that*

$$|A + \xi_1 \cdot A| \leq K_1 |A| \text{ and } |A + \xi_2 \cdot A| \leq K_2 |A|.$$

Then either

$$\min(\log K_1, \log K_2) \gg_q \frac{(\log |A|)^{1/6}}{(\log \log |A|)^{\log_2 3}}$$

or

$$(\log K_1)(\log K_2) \gg_q (\log |A|)^{1/2}.$$

Proof. Without loss of generality, we may suppose that $K_2 \leq K_1$ and that $0 \in A$, so that $A \cup \xi_2 \cdot A \subset A + \xi_2 \cdot A$. By the Plünnecke-Ruzsa estimates we have that $|A + A| \leq K_2^2 |A|$. Since $A \cup \xi_2 \cdot A$ is finite it is contained in a finite-dimensional $\mathbb{F}_q[\xi_1]$-module, say $\mathbb{F}_q[\xi_1]^d \cdot v$. For some large integer N we consider the map $f : \mathbb{F}_q[\xi_1]^d \cdot v \to \mathbb{F}_q[\xi_1]$ defined by

$$f(x_1 v_1 + \cdots + x_d v_d) = x_1 + x_2 \xi_1^N + \cdots + x_d \xi_1^{N(d-1)}.$$

Assuming N is sufficiently large, depending on $A \cup \xi_2 \cdot A$, we have that, for all $x, y \in A \cup \xi_2 \cdot A$,

$$f(x) + f(y) = f(x') + f(y') \text{ implies } x + y = x' + y'$$

and

$$f(x) + \xi_1 \cdot f(y) = f(x') + \xi_1 \cdot f(y') \text{ implies } x + \xi_1 y = x' + \xi_1 y'.$$

In particular, if $A' = f(A \cup \xi_2 \cdot A)$ then, by the Plünnecke-Ruzsa estimates,

$$|A' + A'| \leq |A + A + \xi_2 \cdot A + \xi_2 \cdot A| \leq K_2^{O(1)} |A|$$

and

$$|A' + \xi_1 \cdot A'| \leq |A + \xi_2 \cdot A + \xi_1 \cdot A + \xi_1 \xi_2 \cdot A| \leq K_1^{O(1)} |A|.$$

We observe that in the proof leading to Corollary 4.24 all that was used was that t was transcendental over $\mathbb{F}_q$; in particular, the conclusion is equally valid replacing $\mathbb{F}_q[t] = \mathbb{F}_q[t]$ by $\mathbb{F}_q[\xi_1]$. Hence there exists a $\mathbb{F}_q$-vector space $T \subset \mathbb{F}_q[\xi_1]$ of dimension at most K' such that $A' - A' + T$ contains a $\mathbb{F}_q[\xi_1]$-arithmetic progression P of size $|P| \gg |A'|^{1/K'}$ for some

$$K' \ll (\log K_2)^3 (\log \log K_2)^{3 \log_2 3} + (\log K_1)(\log K_2).$$

We observe that, since f is an $\mathbb{F}_q$-isomorphism on $\langle A' \rangle$ the set $f^{-1}(T)$ is also a $\mathbb{F}_q$-vector space, and furthermore $f^{-1}(P)$ is also a $\mathbb{F}_q[\xi_1]$-arithmetic progression. Applying the inverse map f^{-1} this implies that there is a progression $P' \subset \mathfrak{k}$ of size at least $|P| \gg |A|^{1/K'}$ such that

$$P' = f^{-1}(P) \subset f^{-1}(A' - A' + T) \subset A - A + \xi_2 \cdot A - \xi_2 \cdot A + f^{-1}(T).$$

In particular, for any $l \geq 1$,

$$\left|P' + \xi_2^2 \cdot P' + \cdots + \xi_2^{2(l-1)} \cdot P'\right| \leq q^{lK'} \left|(A - A) + \xi_2 \cdot (A - A) + \cdots + \xi_2^{2l-1} \cdot (A - A)\right|.$$

Since ξ_2 is transcendental over $\mathbb{F}_q[\xi_1]$, however, the left hand side is at least $|P'|^l$. Furthermore, by Lemma 4.25 the right hand side is at most $(q^{K'} K_2^{O(1)})^l \, |A|$. It follows that

$$|A|^{O(l/K')-1} \leq q^{lK'}.$$

Hence we can choose $l \asymp K'$ such that $|A| \leq q^{O(K'^2)}$, and hence

$$\log |A| \ll (\log K_2)^6 (\log \log K_2)^{6 \log_2 3} + (\log K_1)^2 (\log K_2)^2,$$

and the theorem follows. □

CHAPTER 5

SUM-PRODUCT ESTIMATES FOR NON-ARCHIMEDEAN FIELDS

All work in this chapter is joint work with Timothy G. F. Jones and most has been published in [5].

We recall that δ is *permissible* for a collection of finite sets $\mathcal{A}$ if for all $\epsilon > 0$ there exists a constant $C_\epsilon > 0$ such that for all $A \in \mathcal{A}$

$$\max(|A+A|, |AA|) \geq C_\epsilon |A|^{1+\delta-\epsilon}.$$

The sum-product heuristic then says that if $\mathcal{A}$ is some reasonable collection of finite subsets of any ring then 1 is permissible for $\mathcal{A}$; this is clearly the best possible. For a more thorough discussion of the sum-product phenomenon we refer to the introduction. We recall, however, that it is known that $1/3$ is permissible for all finite subsets of $\mathbb{C}$ [72, 36] and that for any finite field of prime order $\mathbb{F}_p$ the constant $1/11$ is permissible for all $A \subset \mathbb{F}_p$ such that $|A| < p^{1/2}$ [53].

In this chapter we will prove new sum-product estimates for non-archimedean local fields, which have not thus far been considered in the sum-product literature. We recall that a non-archimedean local field is a locally compact topological field F equipped with a non-archimedean absolute value; that is, an absolute value $|\cdot| : F \to \mathbb{R}$ such that $|x+y| \leq \max(|x|, |y|)$ for all $x, y \in F$. More concretely any non-archimedean local field is either a finite extension of $\mathbb{Q}_p$ for some prime p or a field of Laurent series $\mathbb{F}_q(\!(t^{-1})\!)$ for some finite field $\mathbb{F}_q$.

We will show that for finite subsets of such fields $1/5$ is permissible. Since the only archimedean local fields are $\mathbb{R}$ and $\mathbb{C}$ this result, combined with the fact that $1/3$ is permissible for $\mathbb{C}$, implies the following.

Theorem 5.1. *Let F be any local field and let $\epsilon > 0$. For any finite $A \subset F$ we have*

$$\max(|A+A|, |AA|) \gg_{F,\epsilon} |A|^{6/5-\epsilon}.$$

Since $\mathbb{F}_q[t]$, a polynomial ring over a finite field, is contained in a non-archimedean local field our result is also valid for any finite set of such polynomials, with an implied

constant dependent on q. This should be compared to the result of Croot and Hart [15], that there exists an absolute constant $\delta > 0$ which is permissible for all finite subsets of $\mathbb{C}[t]$; our methods are not robust enough to apply to $\mathbb{C}[t]$ since they rely crucially on the finiteness of the residue field. On the other hand, we are able to give the fairly large and explicit exponent of $1/5$ for $\mathbb{F}_q[t]$.

The exponent of $1/5$ for non-archimedean local fields lies between the $1/11$ known for finite fields and the $1/3$ known for the archimedean local fields $\mathbb{R}$ and $\mathbb{C}$. It is natural to conjecture, as for the archimedean local fields $\mathbb{R}$ and $\mathbb{C}$, that the correct answer for non-archimedean local fields is 1.

Aside from their intrinsic interest sum-product results over $\mathbb{F}_q[t]$ may have applications to constructions in theoretical computer science. Sum-product results over finite fields have been used to construct efficient randomness extractors; see for example the work of Bourgain in [7] and [8]. A key idea in these constructions is the observation that a string of n bits can be interpreted as an element of the field $\mathbb{F}_{2^n}$ so that the full power of the sum-product machinery can be brought to bear. An alternative is to interpret it as an element of $\mathbb{F}_2[t]$. Given that the sum-product results now available in $\mathbb{F}_2[t]$ are better than those in $\mathbb{F}_{2^n}$ (with an exponent of $1/5$ rather than $1/11$) we expect that constructions along a similar line to those in [7] will be quantitatively stronger over $\mathbb{F}_2[t]$ than over $\mathbb{F}_{2^n}$.

For number theoretic and geometric applications it is worth pointing out that any global function field, that is, a field of transcendence degree 1 over a finite field, can be embedded into $\mathbb{F}_q(\!(t^{-1})\!)$ for some q, where q depends only on the field of constants and genus of the function field. In particular the conclusion of Theorem 5.1 will also hold for any finite subset of a global function field.

We now state our main results and some immediate corollaries. Let F be a non-archimedean local field with a non-archimedean absolute value $|\cdot|$. We define the ring of integers to be $\mathcal{O} = \{a \in F : |a| \leq 1\}$ and let $\mathfrak{m} = \{a \in F : |a| < 1\}$. We then define the residue field of F to be $\mathcal{O}/\mathfrak{m}$ – crucially, because this field is both compact and discrete, it is finite.

Theorem 5.2. *Let F be a non-archimedean local field with a residue field of size q. Let A, B and D be any finite subsets of F. Then*

$$q\,|A+B|^3\,|AD|^2 \gg \frac{|A|^3\,|B|^2\,|D|}{(\log|B|)^2(\log|A|)^4},$$

where the implied constant is absolute.

Theorem 5.1 is an immediate corollary. We observe that the dependence on q here is the best possible, which follows from considering the set $A = \mathbb{F}_q \subset \mathbb{F}_q(\!(t^{-1})\!)$. This contrasts with the finite field setting, where we think of q as being large. Clearly any finite field sum-product result for which the constants depended on q would be meaningless, since we could write everything as $O_q(1)$.

We also have the following corollary.

Corollary 5.3. *Let F be a non-archimedean local field with a residue field of size q. Let A be any finite subset of F. If $|A+A| \leq K|A|$, then for any finite $D \subset F$ and any $\epsilon > 0$*

$$|AD| \gg_{q,K,\epsilon} |A|^{1-\epsilon} |D|^{1/2},$$

and if $|AA| \leq K|A|$, then for any finite $B \subset F$ and any $\epsilon > 0$

$$|A+B| \gg_{q,K,\epsilon} |B|^{2/3-\epsilon} |A|^{2/3-\epsilon}.$$

When A, B, and D are sets of integers a similar result was proved with the best possible exponents by Chang [14].

For many applications it is more convenient to work with a more flexible measure of multiplicative structure than $|AD|$. In particular, let $E_\times(A,D)$ denote the number of $(a,a',d,d') \in A^2 \times D^2$ such that $ad = a'd'$. Our methods allow us to prove the following energy sum-product result.

Theorem 5.4. *Let F be a non-archimedean local field with a residue field of size q. Let A, B and D be any finite subsets of F. Then*

$$q|D|^9 |A|^7 |B|^{-2} |A+B|^3 \gg \frac{E_\times(A,D)^6}{(\log|B|)^2(\log|A|)^4},$$

where the implied constant is absolute.

This should be compared to the main result of Solymosi [71], who showed that for any finite $A, D \subset \mathbb{Z}$ we have $|A+A||D+D| \gg E_\times(A,D)/\log|D|$.

We note that by the Cauchy-Schwarz inequality $|A|^2 |D|^2 \leq |AD| E_\times(A,D)$, and hence Theorem 5.4 delivers a sum-product exponent of $1/9$, which is weaker than the $1/5$ given by Theorem 5.2. For some applications to exponential sums (such as that in Section 5.3),

however, the ability to directly bound the multiplicative energy can be very useful, since in general one cannot deduce upper bounds for $E_\times(A,D)$ from lower bounds for $|AD|$, and hence Theorem 5.4 is, in a sense, qualitatively stronger than Theorem 5.2.

The approach used here is based on a geometric argument used by Solymosi [71] to show that 1/4 is permissible for finite subsets of the complex numbers, coupled with some unique structural properties of non-archimedean geometry.

The rest of the chapter is structured as follows. The first section collects some necessary background on non-archimedean geometry. The second section introduces the crucial concept of a separable set and uses it with a generalisation of the argument of [71] to deduce our results. We conclude by demonstrating how the quantitative strength of such results can be exploited by generalising to local fields a result on exponential sums due to Bourgain, Glibichuk, and Konyagin [11].

5.1 Preliminaries

Let F be a non-archimedean local field equipped with a non-archimedean norm $|\cdot|$. In particular, for all $x, y \in F$ we have $|x+y| \leq \max(|x|, |y|)$, and furthermore $|x| = |-x|$. Both of these properties will be used frequently without further mention. Since F is non-archimedean it has a very rigid geometry, which we will be able to exploit when proving our sum-product estimates. A particular concern will be the behaviour of balls, which are sets of the shape

$$B(x,r) = \{y \in F : |x-y| \leq r\} \text{ for some } x \in F \text{ and } r \in \mathbb{R}_+.$$

We will call r the radius of the ball $B(x,r)$. The fact that F is non-archimedean implies the following standard result.

Lemma 5.5. *If B_1 and B_2 are balls in F then either they are disjoint, or $B_1 \subseteq B_2$, or $B_2 \subseteq B_1$. If in addition B_1 and B_2 have the same radius then either they are disjoint or $B_1 = B_2$.*

Proof. Let $B_1 = B(x,r)$ and $B_2 = B(y,s)$. If there exists $a \in B(x,r) \cap B(y,s)$ then

$$|x-y| \leq \max\{|a-x|, |a-y|\} \leq \max\{r,s\}.$$

If $r \leq s$ this implies that $B(x,r) \subseteq B(y,s)$ since if $b \in B(x,r)$ then

$$|y-b| \leq \max\{|y-x|, |b-x|\} \leq \max\{r,s\} = s.$$

Conversely if $s \leq r$ then $B(y,s) \subseteq B(x,r)$. In particular, if $r = s$ then $B(x,r) = B(y,s)$. □

We shall prove a general result concerning partial product sets, of which the results in the introduction are immediate corollaries. If $G \subset A \times D$ then the partial product set is defined by $A \overset{G}{\cdot} D = \{ad : (a,d) \in G\}$. We will prove the following theorem.

Theorem 5.6. *Let F be a non-archimedean local field with a residue field of size q. Let A, B and D be any finite subsets of F and $G \subset A \times D$. Then*

$$q\,|D|^3\,|A|\,|B|^{-2}\,|A+B|^3\left|A \overset{G}{\cdot} D\right|^2 \gg \frac{|G|^4}{(\log|B|)^2(\log|A|)^4},$$

where the implied constant is absolute.

Theorem 5.2 follows immediately upon taking $G = A \times D$. For Theorem 5.4 we invoke the following simple consequence of the pigeonhole principle; see, for example, Lemma 2.30 of [75].

Lemma 5.7. *Let A and D be any finite subsets of F. There exists a graph $G \subset A \times D$ such that*

$$|G| \gg \frac{E_\times(A,D)}{|A|^{1/2}\,|D|^{1/2}}$$

and

$$\left|A \overset{G}{\cdot} D\right| \ll \frac{|A|^2\,|D|^2}{E_\times(A,D)},$$

where the implied constant is absolute.

The proof of Theorem 5.6 builds upon an approach of Solymosi [71] for sum-products in $\mathbb{C}$, which shows that $1/4$ is permissible for all finite subsets of $\mathbb{C}$. When adapting this method to our problem the non-archimedean geometry is a mixed blessing.

First, the bad news. Solymosi's argument fails at a critical point in the non-archimedean setting, for the following reason. For each $a \in A$, let $a' \in A \setminus \{a\}$ be such that $|a-a'|$ is minimal, and let B_a be the ball of radius $|a-a'|$ centred on a. Solymosi's argument

uses the crucial fact that a single complex number can be contained in at most $O(1)$ of the B_a. This fails spectacularly in F, where an element could be contained in as many as $|A|$ of the B_a, as demonstrated by the following example: let $F = \mathbb{F}_q(\!(t^{-1})\!)$ and $A = \{t^j : 0 \leq j \leq n\}$, so that

$$B_{t^j} = \left\{x \in F : |x| \leq q^j\right\}$$

for $j \geq 1$ and $B_1 = B_t$, meaning that every one of the $|A|$ balls contains 0 as an element.

But all is not lost. In the example above we actually have $|A + A| \approx |A|^2$, and so a strong sum-product estimate holds despite the failure of Solymosi's argument. In fact we will be able to show that something like this is possible whenever Solymosi's argument fails, by considering a special type of structure to be defined in the following section: separable sets.

Separable sets have a paucity of additive structure – so much so, in fact, that the sumset of a separable set has almost maximal size. The idea is to show that a large separable set must exist whenever Solymosi's argument fails. Combining this with an analysis of separable sets as having large sumsets will lead to a proof of Theorem 5.6.

The following section analyses separable sets and shows that their sumsets have maximal growth. We will then adapt Solymosi's proof from [71] to establish that if the sumset and product set are both small then there must exist a large separable set, and use this to prove Theorem 5.6.

5.2 Separable sets and chains

A finite set $A \subset F$ is *separable* if its elements can be indexed as $A = \{a_1, \ldots, a_{|A|}\}$ in such a way that for each $1 \leq j \leq |A|$ there is a ball B_j with $A \cap B_j = \{a_1, \ldots, a_j\}$.

Lemma 5.8. *If $A \subset F$ is a finite separable set then $|kA| \geq (k!)^{-2} |A|^k$ for any $k \geq 1$.*

Proof. Let $E_{2k}(A)$ denote the k-fold additive energy of A, i.e. the number of solutions to

$$a_1 + \ldots + a_k = b_1 + \ldots + b_k \tag{5.1}$$

with $a_i, b_i \in A$. By the Cauchy-Schwarz inequality $|A|^{2k} \leq |kA| \, E_{2k}(A)$, and so it suffices to show that $E_{2k}(A) \leq (k!)^2 |A|^k$.

Let $\{a_1, \dots, a_{|A|}\}$ be an ordering of A such that for $1 \leq j \leq |A|$ there is a ball B_j such that $B_j \cap A = \{a_1, \dots, a_j\}$. It suffices to show that there are at most $|A|^k$ many solutions to (5.1) such that $a_1 \leq \cdots \leq a_k$ and $b_1 \leq \cdots \leq b_k$ with respect to this ordering of A, for $E_{2k}(A)$ is at most $(k!)^2$ multiplied by the number of such solutions. Suppose that there exists some $1 \leq i \leq k$ such that $a_i \neq b_i$, and let such i be maximal. It follows that

$$a_1 + \cdots + a_i = b_1 + \cdots + b_i.$$

We may suppose, without loss of generality, that $a_i < b_i$. Let $B = B(x, r)$ be a ball such that $a_j \in B$ for $1 \leq j \leq i$ and $b_j \in B$ for $1 \leq j < i$, but $b_i \notin B$. It follows that

$$\begin{aligned} |b_i - x| &= |a_i + \cdots + a_1 - b_{i-1} - \cdots - b_1 - x| \\ &= |(a_i - x) + \cdots + (a_1 - x) - \cdots - (b_1 - x)| \\ &\leq \max(|a_i - x|, \dots, |b_1 - x|) \\ &\leq r, \end{aligned}$$

and hence $b_i \in B$, which is a contradiction. We must therefore have $a_i = b_i$ for all $1 \leq i \leq k$, and hence $E_{2k}(A) \leq (k!)^2 |A|^k$ and the proof is complete. □

The driving force of our argument is the following lemma, which shows that if the sum and product set of A are both small then A must contain a large separable set. For this we adapt the argument of [71] (used there for the archimedean field $\mathbb{C}$) to the non-archimedean setting. We remark that all of the analysis in the proof below is non-archimedean; indeed some of the facts of non-archimedean geometry deployed here are manifestly false in $\mathbb{C}$.

A couple of new definitions are required. For a finite set $A \subset F$ and an element $a \in A$, define $r_A(a) = \min_{a' \in A \setminus \{a\}} |a - a'|$ and $B_A(a) = B(a, r_A(a))$. Additionally, for any $n \geq 1$ we say that $C = (c_1, \dots, c_n) \in A^n$ is an A-chain of length n if $c_i \neq c_j$ for $1 \leq i < j \leq n$ and $B_A(c_1) \subseteq \cdots \subseteq B_A(c_n)$.

The following argument, a strengthened form of that found in [71], finds a large chain in A as long as the sumset and partial product set are both small. If this condition were to fail then a suitable sum-product result would follow immediately.

We recall the notation $N \approx k$ for $k \leq N < 2k$.

Lemma 5.9. *Let A, B and D be any finite subsets of F and $G \subset A \times D$. Then A contains an A-chain of cardinality*

$$N \gg \frac{|G|^4 |B|}{|A| |D|^3 |A+B|^2 \left|A \overset{G}{\cdot} D\right|^2 (\log |A|)^4}.$$

Proof. For any $d \in D$ let $G_d = \{a \in A : (a,d) \in G\}$. For each $a \in A$ and $d \in D$ write $N_d(a)$ for the maximal length N of an A-chain $C = (c_1, \ldots, c_N)$ for which $c_N = a$ and $c_i \in G_d$ for $1 \leq i \leq N$. We observe that $N_d(a) \leq |B_A(a) \cap G_d|$, since if $(c_1, \ldots, c_N)$ is such a maximal A-chain then $c_i \in G_d$ for $1 \leq i \leq N$ by definition and for each $c \in C$ we have $c \in B_A(c) \subseteq B_A(a)$.

We begin with some preliminary pigeonholing. Let D' be the set of $d \in D$ such that $|G_d| \geq |G|/2|D|$. We observe that

$$|G| = \sum_{d \in D} |G_d| \leq |D| \frac{|G|}{2|D|} + |D'| |A|$$

and hence $|D'| \geq |G|/2|A|$.

Fix some $d \in D'$ and for $0 \leq j \leq \log_2 |A|$ define $A_d(j)$ to be the set of $a \in G_d$ for which $N_d(a) \approx 2^j$. Since the sets $A_d(j)$ form a partition of G_d there exists some j_d for which $|A_d(j_d)| \gg |G| / |D| \log |A|$. Let $2^{j_d} = k_d$ and $A_d(j_d) = A_d$, so that for all $a \in A_d$ we have $N_d(a) \approx k_d$.

We now perform another dyadic pigeonholing over D' itself. For $0 \leq i \leq \log_2 |A|$ let D'_i be the set of $d \in D'$ such that $k_d \approx 2^i$. Once again, since the sets D'_i form a partition of D' there exists some D'' such that $|D''| \gg |G| / |A| \log |A|$ and an integer k such that $k_d \approx k$ for all $d \in D''$. It suffices to show that

$$k \gg \frac{|G|^4 |B|}{|A| |D|^3 |A+B|^2 \left|A \overset{G}{\cdot} D\right|^2 (\log |A|)^4}.$$

To this end, we say that a pair $(a,c) \in A \times B$ is d-additively good if

$$|(A+B) \cap (B_A(a) + c)| \leq \frac{16k |A+B|}{|A_d|},$$

and that $(a,d) \in G$ is multiplicatively good if

$$\left|(A \overset{G}{\cdot} D) \cap (B_A(a) \cdot d)\right| \leq \frac{16k \left|A \overset{G}{\cdot} D\right|}{|A_d|}.$$

We say that a quadruple $(a,b,c,d) \in A \times A \times B \times D''$ is good if $a \in A_d$, $b \in B_A(a) \cap G_d$, (a,c) is d-additively good and (a,d) is multiplicatively good. Let Q be the number of good quadruples. We shall bound Q from below to obtain

$$Q \gg \frac{k\,|G|\,|D''|\,|B|}{|D|\log|A|} \tag{5.2}$$

and bound it from above to obtain

$$Q \ll \frac{k^2\,|A+B|^2\left|A \overset{G}{\cdot} D\right|^2 |D|^2 (\log|A|)^2}{|G|^2}. \tag{5.3}$$

Comparing these bounds yields the the required bound on k since $|D''| \gg |G| \,/\, |A| \log|A|$.

We shall first establish (5.2). For any $d \in D''\backslash\{0\}$ we have, by rearranging the summation, that

$$\sum_{a\in A_d}\left|(A \overset{G}{\cdot} D) \cap (B_A(a)\cdot d)\right| = \sum_{v\in(A\overset{G}{\cdot}D)\cdot d^{-1}} |C_d(v)|$$

where $C_d(v)$ is the set of $a \in A_d$ with $v \in B_A(a)$. We observe that, for any v, the elements of $C_d(v)$ may be ordered to form an A-chain. This follows from Lemma 5.5 since for any $x,y \in C_d(v)$ we have $v \in B_A(x) \cap B_A(y)$ and so either $B_A(x) \subseteq B_A(y)$ or $B_A(y) \subseteq B_A(x)$. In particular, since $c_n \in A_d$ and $c_i \in G_d$ for $1 \le i \le n$ we have $|C_d(v)| \le N_d(c_n) < 4k$. We therefore have

$$\sum_{a\in A_d}\left|(A \overset{G}{\cdot} D) \cap (B_A(a)\cdot d)\right| < 4k\left|A \overset{G}{\cdot} D\right|$$

and hence there are at least $3\,|A_d|\,/4$ many $a \in A_d$ for which (a,d) is multiplicatively good. By an analogous argument for fixed $c \in A$ there are at least $3\,|A_d|\,/4$ elements $a \in A_d$ for which (a,c) is d-additively good.

It follows that for any fixed $d \in D''$ and $c \in B$ there are at least $|A_d|\,/2$ elements $a \in A_d$ such that (a,c) is d-additively good and (a,d) is multiplicatively good. Furthermore, for each $a \in A_d$ there are at least k many $b \in G_d$ such that $b \in B_A(a) \cap G_d$ since $k \le N_d(a) \le |B_A(a) \cap G_d|$. It follows that

$$Q \gg \sum_{d\in D''}\sum_{c\in B} k\,|A_d| \gg \frac{k\,|G|\,|D''|\,|B|}{|D|\log|A|}$$

as required.

We now prove (5.3). We observe, and this is the key observation of [71], that the map

$$(a, b, c, d) \mapsto (a + c, b + c, ad, bd)$$

is injective and so it suffices to bound the number of possibilities for this latter expression, subject to the constraint that (a, b, c, d) is good. There are certainly at most $|A + B|$ possibilities for $a + c$ and at most $\left|A \overset{G}{\cdot} D\right|$ for ad, so it suffices to show that if these are fixed then there are at most $\ll k\,|A + B|\,|D| \log |A| / |G|$ possibilities for $b + c$ and at most $\ll k\left|A \overset{G}{\cdot} D\right| |D| \log |A| / |G|$ for bd. We shall prove this for $b + c$; the argument for bounding the number of possible bd is similar.

We first observe that if $a + c = a' + c'$ then either

$$B_A(a) + c \subseteq B_A(a') + c' \text{ or } B_A(a') + c' \subseteq B_A(a) + c,$$

since both sets are balls with the same centre $a + c$. As a consequence, if $G' \subseteq A \times B$ is the set of d-additively good pairs (a, c) for any $d \in D''$, then for any $x \in A \overset{G'}{+} B$ there is a fixed pair (a_x, c_x) which is d-additively good for some $d \in D''$ such that

$$B_A(a) + c \subseteq B_A(a_x) + c_x$$

whenever $a + c = x$ and (a, c) is d-additively good for any $d \in D''$. Thus if $a + c = x$ is the fixed first co-ordinate and $b + c$ is a possible second co-ordinate then since $b \in B_A(a) \cap A$ and $c \in B$ we have

$$\begin{aligned} b + c &\in (A + B) \cap (B_A(a) + c) \\ &\subseteq (A + B) \cap (B_A(a_x) + c_x)\,. \end{aligned}$$

Since (a_x, c_x) is d-additively good for some $d \in D''$, there are, as required, at most $\ll k\,|A + B|\,|D| \log |A| / |G|$ possibilities for $b + c$, which concludes the proof. □

The following result shows that any chain contains a large separable subset, allowing Lemma 5.8 to be applied to the chain found by Lemma 5.9.

Lemma 5.10. *Let q be the size of the residue field of F and let A be any finite subset of F. If C is the set of elements of an A-chain then C contains a separable set of cardinality at least $|C| / q$.*

Proof. It is clear that any subset $\{c_1, \ldots, c_n\} \subseteq C$ with

$$B_A(c_1) \subsetneq \ldots \subsetneq B_A(c_n)$$

is separable. Define an equivalence relation on elements of A by $a \sim b$ if and only if $B_A(a) = B_A(b)$. To prove the lemma it suffices to show that each equivalence class contains at most q elements of A.

We first observe that if $a \sim b$ and $a \neq b$ then $|a - b| = r_A(a) = r_A(b)$. Indeed, since $B_A(a) = B_A(b)$ it follows that $b \in B_A(a)$ and so $|a - b| \leq r_A(a)$. By minimality, however, $|a - b| \geq r_A(a)$ and so $|a - b| = r_A(a)$.

Suppose, for a contradiction, that there is an equivalence class containing distinct elements $a_1, \ldots, a_{q+1}$. It suffices to show that there are distinct i, j, k in $\{1, \ldots, q+1\}$ such that

$$|a_k - a_j| \neq |a_k - a_j|,$$

for this contradicts the previous paragraph, as both must be equal to $r_A(a_k)$. Considering the differences $b_i = a_1 - a_{i+1}$ for $1 < i \leq q+1$ this follows from the fact that for any $b_1, \ldots, b_q \in F^*$ such that $|b_1| = \cdots = |b_q|$ there exist $1 \leq i < j \leq q$ such that $|b_i - b_j| < |b_j|$.

This fact, in turn, follows from the assumption that the residue field has cardinality q. We recall that the residue field is defined as $\mathcal{O}/\mathfrak{m}$, where

$$\mathcal{O} = \{x \in F : |x| \leq 1\} \text{ and } \mathfrak{m} = \{x \in F : |x| < 1\}.$$

Without loss of generality, we may suppose that $|b_i| = 1$ for $1 \leq i \leq k$. Since $|\mathcal{O}/\mathfrak{m}| = q$ by the pigeonhole principle there must exist $1 \leq i < j \leq q$ such that $b_i - b_j \in \mathfrak{m}$, and the proof is complete. □

Theorem 5.6 follows by combining Lemma 5.8 with Lemmata 5.9 and 5.10 and Plünnecke's inequality as follows.

Proof of Theorem 5.6. Let

$$L = \frac{|G|^4 |B|}{q |A| |D|^3 |A + B|^2 \left|A \overset{G}{\cdot} D\right|^2 (\log |A|)^4}.$$

By Lemmata 5.9 and 5.10 the set A contains a separable subset U of cardinality $\Omega(L)$, where the implied constant is absolute. For any $k \geq 1$, Lemma 5.8 implies that

$$|kA| \geq |kU| \gg k^{-2k} L^k.$$

By Plünnecke's inequality $|kA| \leq |A+B|^k / |B|^{k-1}$ and hence

$$|A+B| \gg k^{-2} L \, |B|^{1-1/k} \gg \frac{|G|^4 \, |B|^{2-1/k}}{qk^2 \, |A| \, |D|^3 \, |A+B|^2 \left| A \overset{G}{\cdot} D \right|^2 (\log |A|)^4}.$$

The proof is completed by taking $k = \lceil \log |B| \rceil$. □

5.3 An exponential sum estimate

In this section we demonstrate how our sum-product estimates can be used to give strong estimates for exponential sums by inserting them into arguments due to Bourgain, Glibichuk, and Konyagin in [11], where they prove a sum-product estimate for subsets of $\mathbb{F}_p$ and then use it to give an upper bound for an exponential sum over $\mathbb{F}_p$.

Their arguments are fairly robust and we are able to generalise them to handle exponential sums over finite subsets of $\mathbb{Z}$ and $\mathbb{F}_q[t]$. We also keep track of the constants to give an explicit result. The constants in these results could certainly be slightly improved with a little more effort.

We shall use the following strong form of the quantitative Balog-Szemerédi-Gowers theorem, due to Schoen [65].

Lemma 5.11 (Schoen)**.** *Let G be an abelian group and A be a finite subset such that*

$$E(A) = \left| \{(a,b,c,d) \in A^4 : a+b = c+d\} \right| = \kappa \, |A|^3 .$$

There exists $A' \subset A$ such that $|A'| \gg \kappa \, |A|$ and $|A'-A'| \ll \kappa^{-4} \, |A'|$.

We remark that if A is a finite subset of a field F with $0 \notin A$ then an identical lemma holds, replacing additive by multiplicative energy and the difference set $A'-A'$ by the ratio set A'/A'.

We will first use Theorem 5.2 to prove the following explicit generalisation of Theorem 7 from [11].

Theorem 5.12. *Let F be a non-archimedean local field, with a residue field of size q, equipped with the counting measure and let $f, g : F \to \mathbb{C}$ be any functions with finite support. For any $\epsilon > 0$*

$$\sum_{y\neq 0} g(y) \sum_{\substack{x_1,x_2,x_3,x_4 \\ x_1+yx_2=x_3+yx_4}} f(x_1)f(x_2)f(x_3)f(x_4)$$

$$\ll_{q,\epsilon} \|f\|_1^2 \|g\|_1 \|f\|_2^2 \max\left(\left(\frac{\|f\|_2}{\|f\|_1}\right)^2, \left(\frac{\|f\|_1}{\|f\|_2}\right)^{23/90+\epsilon}\left(\frac{\|g\|_2}{\|g\|_1}\right)^{24/90-\epsilon}\right).$$

Proof. Without loss of generality we may suppose that $\|f\|_1 = \|g\|_1 = 1$, and by the triangle inequality we may assume that f and g take only non-negative real values. For this proof we will use the notation $f \circ f = f * (-f)$. With this definition the sum to be estimated is

$$\sum_{y\neq 0} g(y) \sum_x f \circ f(x) f \circ f(yx) = \eta \|f\|_2^2,$$

say. Using the trivial estimates $\|f \circ f\|_\infty \leq \|f\|_2^2$ and $\|f \circ f\|_1 = 1$ it follows immediately that $\eta \leq 1$. Furthermore, we may assume that $\eta \geq 2^{-3}\|f\|_2^2$, or else the theorem follows immediately. Let S be the set of $x \neq 0$ such that $f \circ f(x) \geq 2^{-3}\eta \|f\|_2^2$. The strategy is to give lower bounds for the multiplicative and additive energy of S in terms of η; combined with Lemma 5.11 and Theorem 5.2 this will give a suitable upper bound on η.

By our assumptions we have

$$\sum_{y\neq 0} g(y) \sum_{x=0 \text{ or } yx=0} f \circ f(x) f \circ f(yx) \leq 2\|f\|_2^4 \leq 2^{-2}\eta \|f\|_2^2,$$

and

$$\sum_{y\neq 0} g(y) \sum_{x\notin S\cup\{0\}} f \circ f(x) f \circ f(yx) \leq 2^{-3}\eta \|f\|_2^2$$

and similarly for $yx \notin S \cup \{0\}$. It follows that

$$\sum_{y\neq 0} |S \cap yS|\, g(y) \geq \|f\|_2^{-4} \sum_{y\neq 0} g(y) \sum_{\substack{x\in S \\ yx\in S}} f \circ f(x) f \circ f(yx) \geq 2^{-1}\eta \|f\|_2^{-2},$$

and hence in particular $|S| \geq 2^{-1}\eta \|f\|_2^{-2}$. Furthermore, we have

$$|S| \leq 2^3 \eta^{-1} \|f\|_2^{-2} \sum_x f \circ f(x) = 2^2 \eta^{-1} \|f\|_2^{-2}.$$

Let Λ be the set of $y \neq 0$ such that $|S \cap yS| \geq 2^{-2}\eta \|f\|_2^{-2}$. Estimating as above we have

$$\sum_{y \in \Lambda} |S \cap yS| \, g(y) \geq 2^{-2}\eta \|f\|_2^{-2} .$$

By the pigeonhole principle there exists some $\Lambda' \subset \Lambda$ and $\eta^{-1} \gg \eta' \gg \eta$ such that for all $y \in \Lambda'$ we have $|S \cap yS| \approx \eta' \|f\|_2^{-2}$ and

$$\eta' \|f\|_2^{-2} \sum_{y \in \Lambda'} g(y) \gg \sum_{y \in \Lambda'} |S \cap yS| \, g(y) \gg \mathcal{L}(\eta)^{-1} \eta \|f\|_2^{-2} .$$

By the Cauchy-Schwarz inequality the left hand side is bounded above by $\eta' \|f\|_2^{-2} |\Lambda|^{1/2} \|g\|_2$, and hence, if $\delta = \|g\|_2^{-2} \|f\|_2^2$, then

$$|\Lambda'| \gg_\epsilon (\eta')^{-2} \eta^{2+\epsilon} \delta \|f\|_2^{-2} \gg (\eta')^{-2} \eta^{3+\epsilon} \delta \, |S| .$$

It follows that

$$E_\times(S) = \sum_y |S \cap yS|^2 \gg (\eta')^2 \|f\|_2^{-4} |\Lambda'| \gg_\epsilon \eta^{5+\epsilon} \delta \, |S|^3 .$$

By Lemma 5.11 there is some $S' \subset S$ such that $|S'| \gg \delta \eta^{5+\epsilon} |S|$ and $|S'/S'| \ll \delta^{-4} \eta^{-20-\epsilon} |S'|$. Furthermore,

$$\|f\|_2 \left(\sum_{x,y \in S'} f \circ f(x-y) \right)^{1/2} \geq \sum_{x \in S'} f \circ f(x) \geq 2^{-2} \eta \|f\|_2^2 \, |S'| .$$

It follows that if T is the set of z such that $f \circ f(z) \geq 2^{-5} \eta^2 \|f\|_2^2$ then

$$\sum_{x,y \in S'} 1_{x-y \in T} f \circ f(x-y) \gg \|f\|_2^2 \, \eta^2 \, |S'|^2 .$$

By the pigeonhole principle there is some $T' \subset T$ and $1 \gg \eta'' \gg \eta^2$ such that for all $z \in T''$ we have $f \circ f(z) \approx \eta'' \|f\|_2^2$ and

$$\sum_{x,y \in S'} 1_{x-y \in T} \gg (\eta'')^{-1} \|f\|_2^{-2} \sum_{x,y \in S'} 1_{x-y \in T} f \circ f(x-y) \gg (\eta'')^{-1} \eta^{2+\epsilon} \, |S'|^2 .$$

By the Cauchy-Schwarz inequality the left hand side is at most $|T'|^{1/2} E_+(S')^{1/2}$. Using the trivial bound $|T'| \ll 2^5 (\eta'')^{-1} \|f\|_2^{-2}$ this implies

$$E_+(S') \gg (\eta'')^{-1} \eta^{4+\epsilon} \|f\|_2^2 \, |S'|^4 \gg \delta \eta^{9+\epsilon} \, |S'|^3 .$$

Thus by Lemma 5.11 once again we have some S'' such that $|S''| \gg \delta\eta^{9+\epsilon}|S'|$ and

$$|S'' - S''| \ll \delta^{-4}\eta^{-36-\epsilon}|S''| \text{ and } |S''/S''| \ll \delta^{-5}\eta^{-29-\epsilon}|S''|.$$

By Theorem 5.2 then we must have, for any $\epsilon > 0$,

$$\delta^{-22}\eta^{-166} \gg |S''|^{1-\epsilon} \gg \left(\delta^2\eta^{14}\|f\|_2^{-2}\right)^{1-\epsilon},$$

and hence $\eta^{180}\delta^{24-\epsilon} \ll_\epsilon \|f\|_2^{2-\epsilon}$ as required. □

We will shortly apply this estimate to exponential sums, as in [11]. Before that, however, we take a brief detour to prove a corollary of a sum-product flavour.

Corollary 5.13. *Let F be any non-archimedean local field with a residue field of size q. Let A and B be any finite subsets of F such that $0 \notin B$. For any $\epsilon > 0$ there exists some $y \in B$ such that*

$$|A + y \cdot A| \gg_{q,\epsilon} |A| \min\left(|A|, \left(\frac{|B|}{|A|^{23/24}}\right)^{2/15-\epsilon}\right).$$

In particular, if $|B| \geq |A|^{203/24+\epsilon}$ then there exists some $y \in B$ such that $|A + y \cdot A| \gg_{q,\epsilon} |A|^2$.

Proof. By the Cauchy-Schwarz inequality for any $y \in B$

$$|A|^4 = \left(\sum_z \sum_{x_1,x_2 \in A} 1_{x_1+yx_2=z}\right)^2 \leq |A + y \cdot A| \sum_{x_1,x_2,x_3,x_4 \in A} 1_{x_1+yx_2=x_3+yx_4}.$$

Summing over all $y \in B$ implies that

$$|A|^4|B| \leq \left(\max_{y\in B}|A + y \cdot A|\right) \sum_{y\in B} \sum_{x_1,x_2,x_3,x_4\in A} 1_{x_1+yx_2=x_3+yx_4}.$$

Theorem 5.12 implies that, for any $\epsilon > 0$,

$$|A|^4|B| \ll_{\epsilon,q} \left(\max_{y\in B}|A + y \cdot A|\right)|A|^3|B| \max\left(|A|^{-1}, |A|^{23/180+\epsilon}|B|^{-24/180+\epsilon}\right)$$

and the result follows immediately. □

We now come to our estimate for exponential sums. The $\mathbb{F}_p$ analogue of the following result, although not explicitly stated in [11], is implicitly proven in the course of the proof of their Theorem 5. We follow their proof, taking care to keep track of the constants.

Theorem 5.14. *Let F be a non-archimedean local field, with residue field of cardinality q, and let G be a subring of F. Let $A \subset G$ be any finite set such that $0 \notin A$. Then for all $k \geq 1$ if $r \geq 2^{19k}$ then*

$$\int_{\widehat{G}} \left| \sum_{x_1,\dots,x_k \in A} \xi(x_1 \cdots x_k) \right|^{2r} \mathrm{d}\xi \ll_{k,q} |A|^{2rk - k^{1/253}},$$

where the integral is taken over the dual group of the discrete additive group of G.

Thus, for example, this theorem is applicable for $G = \mathbb{Z}$ or $G = \mathbb{F}_q[t]$. For $\mathbb{Z}$, of course, one could certainly improve the constants by using the better sum-product estimates of Solymosi [72]. We observe that by orthogonality this integral is a count of the number of solutions to

$$\sum_{i=1}^{r} x_{i1} \cdots x_{ik} = \sum_{i=1}^{r} y_{i1} \cdots y_{ik}$$

with $x_{ij}, y_{ij} \in A$. In particular the integral is always a positive integer, and furthermore we have the trivial lower bound $|A|^{rk}$ and the trivial upper bound $|A|^{2rk-1}$.

Finally, we remark that if the strongest possible quantitative estimates were available for both the sum-product problem and the Balog-Szemerédi-Gowers theorem then we would be able to replace $k^{1/253}$ here with $k^{1/21}$; it may be the case that $k^{1-o(1)}$ is the correct bound, for which it seems a different argument would be needed.

Proof. For any $s \geq 0$ and $r \geq 1$ let

$$\mu_s(z) = |A|^{-2^s} \sum_{x_1,\dots,x_{2^s} \in A} 1_{x_1 \cdots x_{2^s} = z} \text{ and } f_{s,r} = \mu_s^{(r)} \circ \mu_s^{(r)}$$

and

$$\delta_{s,r} = \|f_{s,r}\|_2^2 = \int |\widehat{\mu_s}(\xi)|^{4r} \, \mathrm{d}\xi.$$

We claim that $\delta_{s,r}$ decreases as either r or s increases. This is clear for increasing r since $|\widehat{\mu_s}(\xi)| \leq 1$, and for increasing s follows from the identity

$$\widehat{\mu_{s+1}}(\xi) = \sum_{x,z} \xi(xz) \mu_s(z) \mu_s(x) \tag{5.4}$$

and Hölder's inequality. It suffices to prove that for any $s \geq 0$ and $r = 2^{18(2^s-1)}$ we have the estimate

$$\delta_{s,r} \ll_{s,q} |A|^{-2^{2s/253}}. \tag{5.5}$$

We shall use induction on s. We first observe that the case $s = 0$ follows immediately from Parseval's identity.

For now, we fix some $s \geq 0$ and let $r \geq 1$ be any integer. Let $\eta \in (0, 1]$ be some parameter to be chosen later, and let $\Delta_s = \Delta_\eta(\mu_s)$. By (5.4) and the triangle inequality if $\xi \in \Delta_{s+1}$ then

$$\sum_x \left| \sum_z \xi(xz)\mu_s(z) \right| \mu_s(x) \geq \eta.$$

By Hölder's inequality

$$\sum_{y,x} \mu_s(x)\xi(yx)f_{s,r}(y) = \sum_x \left| \sum_z \xi(xz)\mu_s(z) \right|^{2r} \mu_s(x) \geq \eta^{2r}.$$

By Hölder's inequality once again we have

$$\sum_y \left| \sum_x \mu_s(x)\xi(yx) \right|^{4r} f_{s,r}(y) \geq \eta^{8r^2}.$$

Integrating over all $\xi \in \Delta_{s+1}$ this implies

$$\sum_y f_{s,r}(y) \int |\widehat{\mu_{s+1}}(\xi)|^{4r} \left| \sum_x \mu_s(x)\xi(yx) \right|^{4r} \mathrm{d}\xi \geq \eta^{8r^2+4r} m(\Delta_{s+1}),$$

where m is the Haar measure on $\widehat{G}$, normalised so that $m(\widehat{G}) = 1$. By orthogonality

$$\int |\widehat{\mu_{s+1}}(\xi)|^{4r} \left| \sum_x \mu_s(x)\xi(yx) \right|^{4r} \mathrm{d}\xi = \sum_{\substack{x_1,x_2,x_3,x_4 \\ x_1+yx_2=x_3+yx_4}} f_{s+1,r}(x_1)f_{s,r}(x_2)f_{s+1,r}(x_3)f_{s,r}(x_4),$$

and hence

$$\sum_y \sum_{\substack{x_1,x_2,x_3,x_4 \\ x_1+yx_2=x_3+yx_4}} f_{s+1,r}(x_1)f_{s,r}(x_2)f_{s+1,r}(x_3)f_{s,r}(x_4)f_{s,r}(y) \geq \eta^{8r^2+4r} m(\Delta_{s+1}).$$

The contribution from $y = 0$ is

$$f_{s,r}(0) \left\| f_{s+1,r} \right\|_2^2 = \left\| \mu_s^{(r)} \right\|_2^2 \delta_{s+1,r} \leq \delta_{s,r}^{1/2} \delta_{s+1,r} \leq \delta_{s,r}^{3/2}.$$

In particular, if $f = f_{s+1,r} + f_{s,r}$ then

$$\delta_{s,r}^{3/2} + \sum_{y \neq 0} \sum_{\substack{x_1,x_2,x_3,x_4 \\ x_1+yx_2=x_3+yx_4}} f(x_1)f(x_2)f(x_3)f(x_4)f_{s,r}(y) \gg \eta^{8r^2+4r} m(\Delta_{s+1}).$$

We now observe that $\|f\|_2 \ll \delta_{s,r}^{1/2}$ and apply Theorem 5.12. It follows that

$$\delta_{s,r}^{1+1/181} \gg \eta^{8r^2+4r} m(\Delta_{s+1,\eta}).$$

Furthermore, for any r' we have

$$\delta_{s+1,r'} \ll m(\Delta_{s+1,\eta}) + \eta^{4r'} \ll \eta^{-8r^2-4r}\delta_{s,r}^{1+1/181} + \eta^{4r'}.$$

We now choose

$$\eta^{-8r^2-4r} = \delta_{s,r}^{1/182-1/181} \text{ and } r' \geq \frac{181 \cdot 183}{4}(8r^2+4r),$$

so that $\eta^{4r'} \leq \delta_{s,r}^{1+1/182}$ and hence

$$\delta_{s+1,r'} \ll \delta_{s,r}^{1+1/182}.$$

In particular, if we have (5.5) for some $s \geq 0$ and $r = 2^{18(2^s-1)}$ then, if $r' = 2^{18(2^{s+1}-1)}$, we have

$$\delta_{s+1,r'} \ll \left(|A|^{-2^{2s/253}}\right)^{183/182} \ll |A|^{-2^{2(s+1)/253}},$$

and the proof is complete. □

BIBLIOGRAPHY

[1] M. Bateman and N. H. Katz. New bounds on cap sets. *J. Amer. Math. Soc.*, 25(2):585–613, 2012.

[2] F. A. Behrend. On sets of integers which contain no three terms in arithmetical progression. *Proc. Nat. Acad. Sci. U. S. A.*, 32:331–332, 1946.

[3] T. F. Bloom. Translation invariant equations and the method of Sanders. *Bull. Lond. Math. Soc.*, 44(5):1050–1067, 2012.

[4] T. F. Bloom. A quantitative improvement on Roth's theorem on arithmetic progressions. *arXiv:1405.5800*, 2014.

[5] T. F. Bloom and T. G. F. Jones. A sum-product estimate in function fields. *Int. Math. Res. Not. IMRN*, 44(5):1050–1067, 2012.

[6] J. Bourgain. On triples in arithmetic progression. *Geom. Funct. Anal.*, 9(5):968–984, 1999.

[7] J. Bourgain. More on the sum-product phenomenon in prime fields and its applications. *Int. J. Number Theory*, 1(1):1–32, 2005.

[8] J. Bourgain. On the construction of affine extractors. *Geom. Funct. Anal.*, 17(1):33–57, 2007.

[9] J. Bourgain. Roth's theorem on progressions revisited. *J. Anal. Math.*, 104:155–192, 2008.

[10] J. Bourgain and M. Z. Garaev. On a variant of sum-product estimates and explicit exponential sum bounds in prime fields. *Math. Proc. Cambridge Philos. Soc.*, 146(1):1–21, 2009.

[11] J. Bourgain, A. A. Glibichuk, and S. V. Konyagin. Estimates for the number of sums and products and for exponential sums in fields of prime order. *J. London Math. Soc. (2)*, 73(2):380–398, 2006.

[12] J. Bourgain, N. Katz, and T. Tao. A sum-product estimate in finite fields, and applications. *Geom. Funct. Anal.*, 14(1):27–57, 2004.

[13] M.-C. Chang. A polynomial bound in Freiman's theorem. *Duke Math. J.*, 113(3):399–419, 2002.

[14] M.-C. Chang. Sum and product of different sets. *Contributions to Discrete Math.*, 1(1):57–67, 2006.

[15] E. Croot and D. Hart. On sums and products in $\mathbb{C}[x]$. *Ramanujan J.*, 22(1):33–54, 2010.

[16] E. Croot, I. Łaba, and O. Sisask. Arithmetic progressions in sumsets and L^p-almost-periodicity. *Combin. Probab. Comput.*, 22(3):351–365, 2013.

[17] E. Croot and O. Sisask. A probabilistic technique for finding almost-periods of convolutions. *Geom. Funct. Anal.*, 20(6):1367–1396, 2010.

[18] Y. Edel. Extensions of generalized product caps. *Des. Codes Cryptogr.*, 31(1):5–14, 2004.

[19] M. Elkin. An improved construction of progression-free sets. *Israel J. Math.*, 184:93–128, 2011.

[20] P. Erdős and E. Szemerédi. On sums and products of integers. In *Studies in pure mathematics*, pages 213–218. Birkhäuser, Basel, 1983.

[21] P. Erdős and P. Turán. On Some Sequences of Integers. *J. London Math. Soc.*, S1-11(4):261–264, 1936.

[22] G. A. Freiman. Nachala strukturnoi teorii slozheniya mnozhestv. *Kazan. Gosudarstv. Ped. Inst.*, 1966.

[23] H. Furstenberg. Ergodic behaviours of diagonal measures and a theorem of Szemerédi on arithmetic progressions. *J. Analyse Math.*, 31:204–256, 1977.

[24] M. Z. Garaev. An explicit sum-product estimate in $\mathbb{F}_p$. *Int. Math. Res. Not. IMRN*, 11:Art. ID rnm035, 2007.

[25] A. Geroldinger and I. Z. Ruzsa. *Combinatorial Number Theory and Additive Group Theory*. Advanced Courses in Mathematics - CRM Barcelona. Birkhäuser, 2009.

[26] W. T. Gowers. A new proof of Szemerédi's theorem. *Geom. Funct. Anal.*, 11(3):465–588, 2001.

[27] B. Green. Some constructions in the inverse spectral theory of cyclic groups. *Comb. Prob. Comp.*, 2:127–138, 2003.

[28] B. Green. Finite field models in additive combinatorics. In *Surveys in combinatorics 2005*, volume 327 of *London Math. Soc. Lecture Note Ser.*, pages 1–27. Cambridge Univ. Press, Cambridge, 2005.

[29] B. Green and S. Konyagin. On the Littlewood problem modulo a prime. *Canad. J. Math.*, 61:141–164, 2009.

[30] B. Green and I. Z. Ruzsa. Freiman's theorem in an arbitrary abelian group. *J. Lond. Math. Soc. (2)*, 75(1):163–175, 2007.

[31] B. Green and T. Sanders. A quantitative version of the idempotent theorem in harmonic analysis. *Ann. of Math.*, 168(3):1025–1054, 2008.

[32] B. Green and J. Wolf. A note on Elkin's improvement of Behrend's construction. In *Additive number theory*, pages 141–144. Springer, New York, 2010.

[33] D. R. Heath-Brown. Integer sets containing no arithmetic progressions. *J. London Math. Soc. (2)*, 35(3):385–394, 1987.

[34] N. H. Katz and C.-Y. Shen. A slight improvement to Garaev's sum product estimate. *Proc. Amer. Math. Soc.*, 136(7):2499–2504, 2008.

[35] S. Konyagin and I. Łaba. Distance sets of well-distributed planar sets for polygonal norms. *Israel J. Math.*, 152:157–179, 2006.

[36] S. V. Konyagin and M. Rudnev. On new sum-product-type estimates. *SIAM J. Discrete Math.*, 27(2):973–990, 2013.

[37] R. M. Kubota. Waring's problem for $\mathbb{F}_q[x]$. *Dissertationes Math. (Rozprawy Mat.)*, 117:60pp., 1974.

[38] L. Li and O. Roche-Newton. An improved sum-product estimate for general finite fields. *SIAM J. Discrete Math.*, 25(3):1285–1296, 2011.

[39] Y.-R. Liu and C. V. Spencer. A generalization of Meshulam's theorem on subsets of finite abelian groups with no 3-term arithmetic progression. *Des. Codes Cryptogr.*, 52(1):83–91, 2009.

[40] Y.-R. Liu and C. V. Spencer. A generalization of Roth's theorem in function fields. *Int. J. Number Theory*, 5(7):1149–1154, 2009.

[41] Y.-R. Liu and T. D. Wooley. Waring's problem in function fields. *J. Reine Angew. Math.*, 638:1–67, 2010.

[42] Y.-R. Liu and X. Zhao. A generalization of Roth's theorem in function fields. *Michigan Math. J.*, 61(4):839–866, 2012.

[43] H. B. Mann. A proof of the fundamental theorem on the density of sets of positive integers. *Ann. Math.*, 43:523–527, 1942.

[44] J. Marcinkiewicz and A. Zygmund. Sur les foncions independantes. *Fund. Math.*, 28:60–90, 1937.

[45] R. Meshulam. On subsets of finite abelian groups with no 3-term arithmetic progressions. *J. Combin. Theory Ser. A*, 71(1):168–172, 1995.

[46] G. Petridis. New proofs of Plünnecke-type estimates for product sets in groups. *Combinatorica*, 32(6):721–733, 2012.

[47] H. Plünnecke. *Eigenschaften und Abschätzungen von Wirkungsfunktionen.* Gesellschaft für Mathematik und Datenverarbeitung, 1969.

[48] S. Prendiville. Matrix progressions in multidimensional sets of integers. available at http://www.personal.reading.ac.uk/ pr905055/MatrixProgressions.pdf.

[49] M. Rosen. *Number theory in function fields*, volume 210 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002.

[50] K. F. Roth. On certain sets of integers. *J. London Math. Soc.*, 28:104–109, 1953.

[51] K. F. Roth. On certain sets of integers. II. *J. London Math. Soc.*, 29:20–26, 1954.

[52] W. Rudin. *Fourier analysis on groups.* A Wiley-Interscience Publication, 1962.

[53] M. Rudnev. An improved sum-product inequality in fields of prime order. *Int. Math. Res. Not. IMRN*, 16:3693–3705, 2012.

[54] I. Z. Ruzsa. An application of graph theory to additive number theory. *Scientia, Ser. A.*, 3:97–109, 1989.

[55] I. Z. Ruzsa. Solving a linear equation in a set of integers. I. *Acta Arith.*, 65(3):259–282, 1993.

[56] I. Z. Ruzsa. Generalized arithmetical progressions and sumsets. *Acta Math. Hungar.*, 65(4):379–388, 1994.

[57] I. Z. Ruzsa. An analog of Freiman's theorem in groups. *Astérisque*, 258:xv, 323–326, 1999. Structure theory of set addition.

[58] R. Salem and D. C. Spencer. On sets of integers which contain no three terms in arithmetical progression. *Proc. Nat. Acad. Sci. U. S. A.*, 28:561–563, 1942.

[59] T. Sanders. Appendix to: "Roth's theorem on progressions revisited" by J. Bourgain. *J. Anal. Math.*, 104:193–206, 2008.

[60] T. Sanders. On Roth's theorem on progressions. *Ann. of Math. (2)*, 174(1):619–636, 2011.

[61] T. Sanders. On certain other sets of integers. *J. Anal. Math.*, 116:53–82, 2012.

[62] T. Sanders. On the Bogolyubov-Ruzsa lemma. *Anal. PDE*, 5(3):627–655, 2012.

[63] T. Sanders. The structure theory of set addition revisited. *Bull. Amer. Math. Soc. (N.S.)*, 50(1):93–127, 2013.

[64] T. Schoen. Near optimal bounds in Freiman's theorem. *Duke Math. J.*, 158(1):1–12, 2011.

[65] T. Schoen. New bounds in Balog-Szemerédi-Gowers theorem. *available at www.staff.amu.edu.pl/ schoen/remark-B-S-G.pdf*, 2013.

[66] T. Schoen and I. D. Shkredov. Roth's theorem in many variables. *Israel J. Math.*, 199(1):287–308, 2014.

[67] I. D. Shkredov. On a generalization of Szemerédi's theorem. *Proc. London Math. Soc.*, 93(3):723–760, 2006.

[68] I. D. Shkredov. Some examples of sets of large trigonometric sums. *Mat. Sb.*, 198(12):105–140, 2007.

[69] I. D. Shkredov. On sets of large trigonometric sums. *Izv. Ros. Akad. Nauk, Ser. Mat.*, 72(1):161–182, 2008.

[70] I. D. Shkredov. On sumsets of dissociated sets. *Online J. Anal. Combinatorics*, 4:1–27, 2009.

[71] J. Solymosi. On sum-sets and product-sets of complex numbers. *J. Théor. Nombres Bordeaux*, 17(3):921–924, 2005.

[72] J. Solymosi. Bounding multiplicative energy by the sumset. *Adv. Math.*, 222(2):402–408, 2009.

[73] E. Szemerédi. Integer sets containing no arithmetic progressions. *Acta Math. Hungar.*, 56(1-2):155–158, 1990.

[74] T. Tao. The sum-product phenomenon in arbitrary rings. *Contrib. Discrete Math.*, 4(2):59–82, 2009.

[75] T. Tao and V. Vu. *Additive combinatorics*, volume 105 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2006.

[76] B. L. van der Waerden. Beweis einer baudetsche vermutung. *Nieuw Arch. Wisk.*, 15:212–216, 1927.