LECTURE NOTES 1 FOR 254A

TERENCE TAO

1. Introduction

The aim of this course is to tour the highlights of *arithmetic combinatorics* - the combinatorial estimates relating to the sums, differences, and products of finite sets, or to related objects such as arithmetic progressions. The material here is of course mostly combinatorial, but we will also exploit the Fourier transform at times. We will also discuss the recent applications of this theory to geometric combinatorics problems, and in particular the Kakeya problem.

The setup is as follows. Let Z be an abelian additive group: typical examples of Z are the integers $\mathbf{Z}$, a cyclic group $Z = \mathbf{Z}/N\mathbf{Z}$, or a lattice $\mathbf{Z}^n$. These are all discrete examples; one could of course consider continuous groups such as the real line $\mathbf{R}$ or the circle $\mathbf{T}$ but this will not add any new features to the theory, because we will always be dealing with finite subsets of Z.

Let A, B be two finite subsets of Z. We can then form their sum set

$$A + B := \{a + b : a \in A, b \in B\}$$

or their difference set

$$A - B := \{a - b : a \in A, b \in B\}.$$

If Z is also a ring, then we could form their product set also:

$$A \cdot B := \{ab : a \in A, b \in B\},$$

although we will not deal with this set for several weeks.

The type of questions we will be concerned with are the following: what are the relationships between the relative sizes of the sets A, B, $A + B$, $A - B$, $A + A$, $A - A$, $A + B + B$, etc.? A typical question is the following: if A is "essentially closed under addition", in the sense that $|A + A| \sim |A|$, does this mean that A is also essentially closed under subtraction, in the sense that $|A - A| \sim |A|$? And is A essentially closed under iterated addition, in the sense that $|A + A + A| \sim |A|$, etc.? Of course, the number of questions one could ask here is endless, and we do not yet have a completely satisfactory theory, but there are a number of very useful tools developed which can attack these problems.

Some of what we do has the flavor of "approximate group theory". A subgroup G of Z has the property of being closed under addition, or more precisely that $G+G = G$. In particular, $|G + G| = |G|$. We do not directly deal with subgroups here, but

instead with approximate subgroups in which (for instance) $A + A$ is only slightly larger than A. The question is then to what extent does the machinery and intuition from group theory (e.g. the concepts of cosets, quotient spaces, homomorphisms, etc.) carry over to this approximate setting.

Another variant we will consider is when we no longer consider complete sum sets $A + B$, but rather partial sum sets of the form

$$\{a + b : (a, b) \in G\}$$

for some "large" subset G of $A \times B$, as well as the associated partial difference set

$$\{a - b : (a, b) \in G\}.$$

Knowing the relationship between the size of these two sets has immediate application to the Kakeya problem, which we will discuss later in this course. (I'll just give one small hint about the relationship here: the Kakeya problem concerns how well various line segments, pointing in different directions, can overlap each other. If a line segment connects two points a and b, then the point $(a + b)/2$ must also lie in the line segment, whereas the difference $(a - b)$ is essentially the direction of the line segment. Thus if one can compress line segments in lots of directions into a very small set, one should end up with a situation in which a certain partial sum set is small but the corresponding partial difference set is large. Thus any estimate we have connecting partial sum sets with partial difference sets should lead to some bound on the size of Kakeya sets.)

The above questions were all concerned with sums and differences, and the methods we will use to deal with them are mostly combinatorial, and based on basic arithmetic facts, such as the commutativity and associativity of addition, or of such equivalences as "$a + b = c + d$ if and only if $a - d = c - b$". However, we will also rely sometimes on the Fourier transform. The reason this transform comes in is the following: if f is a function supported on a set A, and g is supported on B, then the convolution $f * g$ is supported on $A + B$, while $f * \tilde{g}$ is supported on $A - B$, where $\tilde{g}(x) := g(-x)$ is the reflection of g. Thus, it is plausible that knowledge about how $f * g$ and $f * \tilde{g}$ are related will lead to information about how $A + B$ and $A - B$ are related. To analyze these convolutions, the most obvious tool to use is the Fourier transform.

(A side note: there are basically three aspects to the Fourier transform: the *analytic* side, which has to do with estimates such as Plancherel, Hausdorff-Young, Littlewood-Paley, Sobolev, and the uncertainty principle; the *algebraic* side, which has to do with group actions (symmetries), characters, representation theory, and so forth; and the *arithmetic* side, which has to do with how the Fourier transform measures how well a set is closed under addition or subtraction. The three aspects are of course related, but of course the emphasis in this course will mostly be on the arithmetic side of the Fourier transform).

Closely related to the theory of sums and differences is that of *arithmetic progressions*. Arithmetic progressions are the most obvious example of "approximate groups": if $A := \{a + jr : 1 \leq j \leq N\}$ is an arithmetic progression of size N, then $A + A$ is another arithmetic progression of almost the same size ($2N - 1$, to be

precise). More generally, we can consider an arithmetic progression of dimension d, which is something of the form

$$A := \{a + j_1 r_1 + \ldots + j_d r_d : 1 \leq j_s \leq N_s \text{ for } j = 1, \ldots, d\};$$

for generic choices of spacing $r_1, \ldots, r_d$, this has size $N_1 \ldots N_d$, while $A + A$ is only slightly larger with size $(2N_1 - 1) \ldots (2N_d - 1) \sim 2^d |A|$. Thus arithmetic progressions are good examples of sets where $|A + A|$ is comparable to A. One could also take a large subset $A' \subset A$ of an arithmetic progression to obtain another set for which $|A' + A'|$ is comparable to $|A'|$. (For instance, A' could be $N/2$ numbers from $\{1, \ldots, N\}$, selected at random). There is a very deep theorem known as *Freiman's theorem* which gives a converse to this statement: if $A + A$ has size comparable to A, then A is a large subset of an arithmetic progression of small dimension. We'll prove this theorem later in this course.

Arithmetic progressions are closely related to the Fourier transform. The key connection is the following: (infinite) arithmetic progressions are nothing more than the level sets of characters (such as $\exp(2\pi i k x)$)! (Another, closely related connection is the Poisson summation formula).

Later on in this course we will study arithmetic progressions in more detail. An old conjecture of Erdös - still unsolved - is whether the prime numbers have arithmetic progressions of arbitrary length. Even in the simplest non-trivial case of arithmetic progressions of length 3, we do not know the answer (though in one sense we are very close, only off by a square root of a logarithm - more on this later). It may be that this question will be resolved using some deep facts from number theory, but it is also quite possible that one can use a far cruder argument, using only the fact that the set of primes in $\{1, \ldots, N\}$ has density $\sim 1/\log N$. This leads to the following problem: given any $k \geq 3$, what density do we need of subsets of $\{1, \ldots, N\}$ to guarantee a non-trivial arithmetic progression of length k? (By non-trivial we mean that the spacing of the progression is non-zero). There is a famous result of Szemeredi in this direction: given any $\delta > 0$ and $k \geq 3$, it is true that every subset of $\{1, \ldots, N\}$ of density at least δ (i.e. cardinality at least δN) has a non-trivial arithmetic progression of length k, provided N is large enough depending on δ and k. We will give several proofs of this result for small values of k, and give a beautiful (but rather unusual) proof by Furstenburg using ergodic theory in the general case.

2. Bounds on $A + B$

We now begin with one of the most basic questions: given the cardinalities $|A|$, $|B|$ of two non-empty finite sets A, B in a group Z, what can one say about $|A + B|$? (This is of course a vacuous question if A or B is empty).

We have a trivial upper bound

$$|A + B| \leq |A||B|$$

and this is sharp, as can be seen when A and B are generic (and Z is large). The more interesting question is what lower bound one can get on $|A + B|$.

We first make a basic observation: if we translate A or B by any amount, this does not affect the cardinalities of A, B, or $A+B$ (or of $A-B$, etc.).

Lemma 2.1. *If A and B are non-empty finite subsets of Z, then we have $|A+B| \geq |A| + |B| - 1$.*

Proof We exploit the fact that the integers are ordered. We may translate A so that $\sup(A) = 0$, and translate B so that $\inf(B) = 0$. Since $0 \in A, B$, we thus see that $A+B$ contains $A \cup B$. But A consists of non-positive integers and B consists of non-negative integers, so $|A \cup B| = |A| + |B| - 1$, and the claim follows. ■

These bounds are sharp; see Exercises 1 and 2. Note from Exercise 2 that the lower bound is only attained of A and B are arithmetic progressions.

This settles the problem when the group Z is the integers. What about other groups? We first observe that all torsion-free groups are "equivalent" to Z, in the following sense.

Definition 2.2. An abelian group Z is *torsion-free* if one has $nx \neq 0$ for all non-zero x and all $n \geq 1$, where $nx = x + \ldots + x$ is the summation of n copies of x.

Recall that an isomorphism $\phi : Z \to Z'$ between two abelian groups is a bijection such that $\phi(x+y) = \phi(x) + \phi(y)$ for all $x, y \in Z$. Unfortunately, isomorphisms are very rare (e.g. there is no isomorphism between $\mathbf{Z}$ and $\mathbf{Z}^2$). Thus we will use a more relaxed notion of isomorphism.

Definition 2.3. Let $k \geq 2$, let $A \subset Z$ be the subset of one abelian group, and let $B \subset Z'$ be the subset of another abelian group. A *Freiman isomorphism of order k* $\phi : A \to B$ is a bijection from A to B such that for any $x_1, \ldots, x_k, y_1, \ldots, y_k \in A$, we have

$$\phi(x_1) + \ldots + \phi(x_k) = \phi(y_1) + \ldots + \phi(y_k)$$

if and only if

$$x_1 + \ldots + x_k = y_1 + \ldots + y_k.$$

The idea is that a Freiman isomorphism looks exactly like an actual isomorphism as long as you are only allowed to perform at most k additions. Observe that if ϕ is a Freiman isomorphism of order k, then it is automatically a Freiman isomorphism of order k' for all $k' \leq k$. Also, the composition of two Freiman isomorphisms of order k is another Freiman isomorphism of order k, as is the inverse of a Freiman isomorphism of order k. Note that every genuine isomorphism is a Freiman isomorphism of every order, as is any translation map $\phi(x) := x + x_0$. Finally, we see that if there is a Freiman isomorphism on $A \cup B \subset Z$ of order at least 2 which maps A to A' and B to B', then $A+B$ will have the same cardinality as $A'+B'$, in fact ϕ induces an explicit bijection between the two sets. (Note also that the same is true for $A-B$, because of the simple identity $x_1 + x_2 = y_1 + y_2 \iff x_1 - y_2 = y_1 - x_2$. If one also wants to preserve the cardinality of more complicated expressions such as $A+B-B$ then one needs higher order isomorphisms, of course.).

We can now make precise the statement that torsion-free abelian groups are no richer than the integers, for the purposes of understanding sums and differences of finite sets.

Lemma 2.4. *Let A be a finite subset of a torsion-free abelian group. Then for any integer k, there is a Freiman isomorphism $\phi : A \to \phi(A)$ to some finite subset $\phi(A)$ of the integers.*

Note that the converse is trivial: one can always embed the integers in any other torsion-free abelian group.

Proof We may extend Z to be a vector space over the rationals $\mathbf{Q}$ (by replacing Z with $Z \otimes_{\mathbf{Z}} \mathbf{Q}$). Without loss of generality we may translate A so that it contains 0. Now look at $\mathrm{span}(A)$, the span of A over the rationals. This is a finite-dimensional vector space over $\mathbf{Q}$ and is thus isomorphic to $\mathbf{Q}^n$ for some n. Since A is finite, we thus see that A is contained in some lattice which is isomorphic to $\mathbf{Z}^n$. Thus without loss of generality we may assume that $Z = \mathbf{Z}^n$.

Now let M be a large integer, and define the map $\phi : \mathbf{Z}^n \to \mathbf{Z}$ by

$$\phi(a_1, \ldots, a_n) := a_1 + a_2 M + a_3 M^2 + \ldots + a_n M^{n-1}.$$

(i.e. we view elements of $\mathbf{Z}^n$ as digit strings of integers base M. If M is large enough depending on A and k, we see that this is a Freiman isomorphism (because if M is large enough we never have to "carry" a digit). ■

As a corollary, we see that Lemma 2.1 extends to all torsion-free abelian groups (and indeed, any inequality involving a finite number of sums and differences which works for $\mathbf{Z}$, will work for all torsion-free abelian groups). Note that something slightly non-trivial is going on here, because Lemma 2.1 relied crucially on the ordering of $\mathbf{Z}$, which is not an arithmetic property and is not preserved under Freiman isomorphisms.

Now we look at what happens when there is torsion, e.g. if $Z = \mathbf{Z}/N\mathbf{Z}$ for some N. The first thing we see is that we have the trivial bound $|A + B| \geq \max(|A|, |B|)$, since $A + B$ must contain at least one translate of A and at least one translate of B. This is sharp: if A, B are equal to the same finite subgroup G of Z, then $|A + B| = |A| + |B|$. (See Exercise 3 for a more precise formulation of when $|A + B| \geq \max(|A|, |B|)$ is sharp.)

Now let's look at what happens when there aren't any proper subgroups, i.e. when $Z = \mathbf{Z}/p\mathbf{Z}$ for some prime p.

Theorem 2.5 (Cauchy-Davenport inequality)**.** *If A, B are any two non-empty subsets of $\mathbf{Z}/p\mathbf{Z}$ for some prime p, then*

$$|A + B| \geq \min(|A| + |B| - 1, p).$$

This theorem is superficially similar to the torsion-free inequality in Lemma 2.1, but is more non-trivial to prove - we will in fact give two and a half proofs of this

inequality. Note that we need the p on the right-hand side, since clearly we cannot make $|A+B|$ exceed p.

We first deal with a trivial case, when $|A|+|B|-1 \geq p$. Then $|A|+|B|$ is larger than $|\mathbf{Z}/p\mathbf{Z}|$, and by the pigeonhole principle A and $x-B$ will always intersect for any $x \in \mathbf{Z}/p\mathbf{Z}$. Thus $A+B = \mathbf{Z}/p\mathbf{Z}$, and so $|A+B| = p$ as desired. So we only need to prove the inequality $|A+B| \geq |A|+|B|-1$ when $|A|+|B|-1 < p$.

We now begin the first proof of the Cauchy-Davenport inequality, based on contradiction. Suppose we have a counterexample to the inequality, so that $|A+B| < |A|+|B|-1 \leq p$ for some A and B. Let's assume the counterexample is *minimal* in the sense that $|A|$ is as small as possible. (Note that the inequality is trivial when $|A| = 1$, so $|A| > 1$). By translating A and B we may assume that A, B have a point in common.

Now we do a little trick, known as the *Dyson e-transform*[1]: we replace A and B by $A' := A \cap B$ and $B' := A \cup B$. From inclusion-exclusion we see that $|A'|+|B'| = |A|+|B|$. Also we observe that $A'+B'$ is a subset of $A+B$ (because

$$A'+B' = A'+(A\cup B) \subseteq (A'+B)\cup(A'+A) \subseteq (A+B)\cup(B+A) = A+B$$

). Thus A', B' will be another counterexample to the Cauchy-Davenport inequality (we've kept $|A|+|B|-1$ the same size, but made $A+B$ equal or smaller cardinality; note that A' is non-empty by assumption). This will contradict minimality of A unless $A' = A$,i.e. if A is contained in B.

Thus when we have a minimal counterexample A, B to the Cauchy-Davenport inequality, we know that $A \subseteq B$ whenever A and B intersect. More generally, by translation invariance we see that $A+x \subseteq B$ whenever $A+x$ and B intersect. In particular, this means that $B-A+A \subseteq B$ (since $A+x$ intersects B precisely when $x \in B-A$). Thus $|B-A+A| \leq |B|$. By Exercise 3 this means that B is a translate of some subgroup of $\mathbf{Z}/p\mathbf{Z}$, which means that either $|B| = 1$ or $B = \mathbf{Z}/p\mathbf{Z}$. But in either case we can easily check that we do not have a counterexample to Cauchy-Davenport. Thus there is no minimal counterexample to Cauchy-Davenport.

Now we give a completely different proof of Cauchy-Davenport, due to Elon, Nathanson, and Ruzsa. Write $m := |A|$, $n := |B|$; we may assume that $m+n-1 < p$. Write $F := \mathbf{Z}/p\mathbf{Z}$, and consider the set F^A, the set of F-valued functions on A. This is an m-dimensional vector space over the field F. Now consider the polynomials $1, x, x^2, \ldots, x^{m-1}$, thought of as functions from A to F (i.e. elements of F^A). These m functions are linearly independent (over F) in F^A, because polynomials of degree $m-1$ have at most $m-1$ zeroes. Thus these functions form a basis for F^A. Similarly the functions $1, y, y^2, \ldots, y^{m-1}$ form a basis for F^B.

Now consider $F^A \otimes F^B = F^{A\times B}$, the space of F-valued functions $f(x,y)$ for $x \in A$ and $y \in B$. Then the functions $x^j y^k$, for $0 \leq j \leq m-1$ and $0 \leq k \leq n-1$, form a basis for $F^{A\times B}$. In particular, if we let S be the $mn-1$ monomials $\{x^j y^k : 0 \leq j \leq m-1, 0 \leq k \leq n-1, (j,k) \neq (m-1, n-1)\}$, then $x^{m-1}y^{n-1}$ does not lie in

[1]Actually, this is the Dyson 0-transform. The e-transform is the 0-transform composed with a translation by e.

the span of S (over F). On the other hand, observe that $x^j y^k$ *does* lie in the span of S if $0 \leq j < m-1$ (because y^k is a linear combination of $1, y, \ldots, y^{n-1}$) or if $0 \leq k < n-1$ (because x^j is a linear combination of $1, x, \ldots, x^{m-1}$).

Now suppose that the Cauchy-Davenport inequality failed, so that $|A+B| < m+n-1$. Then there exists a set $C \subseteq F$ of cardinality $|C| = m+n-2$ such that $A+B \subseteq C$. Thus if we define the function $f \in F^{A \times B}$ by

$$f(x,y) := \prod_{z \in C} (x+y-z)$$

then f vanishes on $A \times B$, i.e. $f \equiv 0$. However, if we expand $f(x,y)$ out completely, we get a monomial term of the form

$$\left(\begin{array}{c} m+n-2 \\ m-1 \end{array} \right) x^{m-1} y^{n-1},$$

plus a lot of other monomial terms involve $x^j y^k$ where either $j < m-1$ or $k < n-1$ or both, and which thus lie in the span of S. But since $m_n - 1 < p$, we see that $\left(\begin{array}{c} m+n-2 \\ m-1 \end{array} \right)$ is not a multiple of p, and so it is invertible in F. This contradicts the fact that $x^{m-1} y^{n-1}$ does not lie in the span of S.

Now we give half of a proof of the Cauchy-Davenport inequality. Note that the Cauchy-Davenport inequality is equivalent to the more symmetrical:

Proposition 2.6. *Let A, B, C be non-empty subsets of F such that $|A|+|B|+|C| \geq p+2$. Then $A+B+C = \mathbf{Z}_p$.*

Indeed, to see how this proposition implies Cauchy-Davenport, just set $C := -(\mathbf{Z}/p\mathbf{Z} \backslash (A+B))$ and take contrapositives.

We will use the Fourier transform. Let $l^2(F)$ be the vector space of complex-valued functions on F. Given any $f \in l^2(F)$, we can define the Fourier transform $\hat{f} \in l^2(F)$ by

$$\hat{f}(\xi) := \frac{1}{\sqrt{p}} \sum_{x \in F} e^{2\pi i x \xi / p} f(x).$$

This is an isometry from $l^2(F)$ to $l^2(F)$. Now look at the subspace $l^2(A)$ of $l^2(F)$; this is the $|A|$-dimensional space of complex-valued functions on A. The Fourier transform $\widehat{l^2(A)}$ of this space is thus also $|A|$ dimensional. Heuristically, this means that given any $|A|$ distinct frequencies $\xi_1, \ldots, \xi_{|A|}$, we should be able to find a function $f \in l^2(A)$ whose Fourier coefficients $\hat{f}(\xi_1), \ldots, \hat{f}(\xi_{|A|})$ are equal to anything we specify. (This is true if the characters $e^{2\pi i x \xi_j / p}$ for $j = 1, \ldots, |A|$ are linearly independent on A). If we believe this, then in particular we can make one Fourier coefficient equal to 1 and $|A|-1$ other coefficients equal to 0.

Since $|A|+|B|+|C| \geq p+2$, we can find subsets X, Y, Z covering $F - \{0\}$ with cardinality $|A|-1$, $|B|-1$, and $|C|-1$ respectively. Then if we believe the above heuristic, we can find functions f, g, h in $l^2(A)$, $l^2(B)$, $l^2(C)$ such that f has a non-zero Fourier coefficient at 0 but zero coefficients on X, and similarly for g, h and

Y, Z. But by Parseval, this implies that the Fourier coefficients of $f * g * h$ are zero everywhere except at 0 where it is non-zero; this implies that $f * g * h$ is a non-zero constant. Since $f * g * h$ is supported on $A + B + C$, the claim follows.

Unfortunately I haven't been able to prove the linear independence of the characters necessary to make this proof work (even though one has a vast amount of freedom in selecting the sets X, Y, Z), however it does show how the Fourier transform can (in principle) be used to control additive information of sets.

3. On A and $A + B$

Now let's go back to a general abelian group Z, and suppose that we have two finite non-empty sets A, B. We have the bound $|A + B| \geq |A|$. Call a set A *B-invariant* if we have $|A + B| = |A|$; of course this can only happen when A has at least as many elements as B. From Exercise 3 we know that B-invariant sets are unions of cosets of some subgroup G, generated by some translate of B. Note that $|G|$ must be larger than or equal to $|B|$, but less than or equal to $|A|$.

Thus we have a satisfactory description of B-invariant sets. Now let us call A *essentially B-invariant* if $|A + B| \sim |A|$. Is there an analogous structure theorem for essentially B-invariant sets - that they are essentially cosets of a group generated by B (or one of its translates)?

It turns out the answer is yes - this is a variant of Freiman's theorem - but it is not easy to prove. However, we can begin to approach this fact with a number of partial results of this nature.

First of all, suppose that A and A' are two genuinely B-invariant sets. Then A and A' are unions of cosets of the same sub-group G. This gives us some additional information on the sum set $A+A'$. First of all, it must also be unions of cosets of G, but secondly, it cannot be as large as $|A||A'|$, because $G + G = G$. Indeed, we now have the upper bound $|A + A'| \leq |A||A'|/|G|$, and similarly $|A - A'| \leq |A||A'|/|G|$ (this is basically because we can pass to the quotient group Z/G and use the bound $|(A + A')/G| \leq |A/G||A'/G|$, etc.). In particular, if we use the crude bound $|G| \geq |B|$, we obtain $|A \pm A'| \leq |A||A'|/|B|$.

Let's look at the dual situation. Suppose A is both genuinely B-invariant and genuinely B'-invariant. Then it is not hard to see that A must consist of cosets of some group G which contains both a translate of B and a translate of B'. In particular, this means that $B \pm B'$ is contained in a coset of G. Using the trivial bound $|G| \leq |A|$, we thus obtain the bound $|B \pm B'| \leq |A|$.

The first question we can ask is whether these crude bounds continue to hold for essentially B-invariant sets instead of genuinely B-invariant sets. One can partially answer this question by means of the following elementary lemma of Imre Ruzsa.

Lemma 3.1. [4] *If U, V, W are three non-empty finite subsets of an abelian group Z, then $|V - W| \leq \frac{|U+V||U+W|}{|U|}$.*

Note that this is a generalization of the trivial bound $|V - W| \leq |V||W|$. In the special case where U is a subgroup, this bound in fact comes from the trivial bound applied to the abelian group Z/U, and then pulled back to Z. (In that particular case one can improve the left-hand side to $|U+V-W|$, and more generally one can do so when U is essentially closed under addition, see Q10. However for general U one cannot hope for such an estimate, see Q9). One can think of this lemma as sort of a triangle inequality: if one can control the arithmetic interaction between U and V, and between U and W, then one can also control the arithmetic interaction between V and W.

Proof Consider the linear map $\pi : V \times W \to V - W$ defined by $\pi(x, y) := x - y$. This map is clearly surjective, and so we can find a partial inverse $f : V - W \to V \times W$ such that $\pi(f(w)) = w$ for all $w \in V - W$. In particular, the points $f(w)$ all have different values of π as w varies in $V - W$.

Let $U^\Delta \in Z \times Z$ be the diagonal $U^\Delta := \{(u, u) : u \in U\}$. Observe that $(V \times W) + U^\Delta = (U+V) \times (U+W)$. In particular, for any $w \in V - W$, the sets $f(w)+U^\Delta$ lie in $(U+V) \times (U+W)$. Furthermore, since U^Δ is in the null space of π, we see that the sets $f(w) + U^\Delta$ are all disjoint. Since each set $f(w) + U^\Delta$ has cardinality $|U|$ and we have $|V - W|$ such sets, we have $|U||V - W| \leq |U + V||U + W|$ as desired. ■

In this proof we see a number of interesting tricks, notably the trick of selecting a section (or partial inverse) f to an injective function π, which is then used to obtain a certain disjointness property. We will meet this trick again later in this course.

From this lemma we see that if A and A' are two essentially B-invariant sets, then

$$|A - A'| \leq \frac{|A + B||A' + B|}{|B|} \lesssim \frac{|A||A'|}{|B|}$$

which is consistent with our previous discussion. Or if A is both essentially B-invariant and B'-invariant, then

$$|B - B'| \leq \frac{|A + B||A + B'|}{|A|} \lesssim |A|$$

which is also consistent with our previous discussion. Of course, we would also like to control $|A + A'|$ and $|B + B'|$, but we do not yet have the technology to do so; the problem is that if $|A + B| \sim |A|$ we do not yet know if $|A - B| \sim |A|$ (this seems true, based on analogy with genuinely B-invariant sets, and we will eventually prove something like this, but we certainly don't know it yet).

Another thing we see from the analogy between essentially B-invariant sets and genuinely B-invariant sets is that if A is essentially B-invariant, then it should also be essentially $B + B$-invariant, essentially $B - B$-invariant, essentially $B + B + B$-invariant, etc. We will prove something like this in the next three sections.

4. Plünnecke's theorem

The purpose of this section is to prove

Theorem 4.1 (Plünnecke's theorem)**.** [3] *Let A, B be two finite non-empty subsets of an abelian group Z, and suppose that $|A+B| \leq K|A|$ for some real number $K \geq 1$. Then there is some non-empty subset A' of A such that $|A' + B + B| \leq K^2|A'|$.*

In particular, every essentially B-invariant set contains an essentially $B+B$-invariant set. Of course this theorem can be iterated to construct $B + B + B + B$-invariant sets, etc.

The idea behind this theorem is that if the passage from A to $A + B$ "magnifies" cardinality by K, then the passage from A to $A+B+B$ should magnify cardinality by at most K^2. In order to quantify this notion of magnification, we shall need some machinery from graph theory.

Set $V_0 := A$, $V_1 := A + B$, $V_2 := A + B + B$. We can define a directed graph $G = G[A, B]$ connecting V_0 to V_1 and V_1 to V_2 as follows: for every $a \in A$ and $b \in B$, we connect $a \in V_0$ to $a + b \in V_1$; we let $E_{0\to 1}$ be the set of all such edges. For every $a + b \in A + B$ and $c \in B$, we connect $a + b \in V_1$ to $a + b + c \in V_2$.

This graph $G[A, B]$ is an example of a *commutative graph*, which we now define.

Definition 4.2. Let Z be an abelian group. A *commutative graph* (or *Plünnecke graph*) of depth 2 is a graph G with three (possibly overlapping) finite sets of vertices $V_0, V_1, V_2 \subset Z$, and two sets of directed edges $E_{0\to 1}$, $E_{1\to 2}$, such that each edge e in $E_{0\to 1}$ connects V_0 to V_1, and each edge in $E_{1\to 2}$ connects V_1 to V_2. Furthermore, we have the following *commuting square property*: if the edge $a \to a + b$ lies in $E_{0\to 1}$, and the edge $a + b \to a + b + c$ lies in $E_{1\to 2}$, then the edge $a \to a + c$ must also lie in $E_{0\to 1}$, and the edge $a + c \to a + b + c$ must also lie in $E_{1\to 2}$.

The commuting graph property is so named because it reflects a basic feature of addition, namely that the translation maps $a \mapsto a + b$ and $a \mapsto a + c$ commute. There are commuting graphs of higher order depths, but we will not need them here (but see Exercise 6).

Observe from the commuting square property that every edge $a \to a + b$ in $E_{0\to 1}$ induces an injection from edges $a + b \to a + b + c$ emanating from $a + b$, and edges $a \to a + c$ emanating from a; we call this the *pullback map* induced by the edge $a \to a + b$. Similarly, any edge $d \to d + c$ in $E_{1\to 2}$ induces a *pushforward map* from edges $d - b \to d$ terminating at d, to edges $d + c - b \to d + c$ terminating at $d + c$. Intuitively, these maps show us that the behavior of the graph from V_0 to V_1 must somehow be similar to that from V_1 to V_2.

We do allow V_0, V_1, V_2 to overlap, but this can be easily fixed by replacing Z with the product space $Z \times \mathbf{Z}$, and replacing V_0, V_1, V_2 by their respective lifts $V_0 \times \{0\}$, $V_1 \times \{1\}$, $V_2 \times \{2\}$, and adjusting the edges accordingly. Thus we will take V_0, V_1, V_2 to be disjoint. (Note how the freedom to take Cartesian products of our

abelian group gives us additional room to maneuver; we will see more examples of this product trick later on).

Let $G = (V_0, V_1, V_2, E_{0\to 1}, E_{1\to 2})$ be a commuting graph. If A' is a subset of V_0, we define $G(A') \subseteq V_1$ to be the set of all vertices in V_1 which are connected via an edge in $E_{0\to 1}$ to some vertex in A'. Similarly we define $G^2(A') \subseteq V_2$ to be the set of all vertices in V_2 which are connected via an edge in $E_{1\to 2}$ to some vertex in $G(A')$, or equivalently the set of all vertices in V_2 which are connected via a path of length 2 to a vertex in A'. In the special case $G = G[A, B]$, observe that $G(A')$ is just $A' + B$ and $G^2(A')$ is just $A' + B + B$.

Plünnecke's theorem will then follow from the following proposition.

Proposition 4.3. *Let G be a commutative graph, and suppose that $|V_1| < K|V_0|$. Then $|G^2(A')| < K^2|A'|$ for some $A' \subseteq V_0$.*

To prove this proposition, let us first prove it in the special case when $C = 1$:

Proposition 4.4. *Let G be a commutative graph such that $|V_1| < |V_0|$. Then $|G^2(A')| < |A'|$ for some $A' \subseteq V_0$.*

This Proposition asserts, informally, that if G is not an expanding map, then neither is G^2. We prove this theorem in the next section.

5. Some graph theory

To prove Proposition 4.4, we recall *Menger's theorem* from graph theory. Let G be a directed graph, and let A and B be two subsets of G. We define $MAXFLOW(A \to B; G)$ to be the maximum number of disjoint paths in G which connect a vertex in A with a vertex in B. We define $MINCUT(A \to B; G)$ to be the minimum number of vertices one needs to remove from G in order to disconnect A from B.

Theorem 5.1 (Menger's theorem)**.** $MAXFLOW(A \to B; G) = MINCUT(A \to B; G)$.

This theorem has a nice intuitive interpretation; the only obstructions to flowing from A to B are cuts which disconnect A from B. In particular, if one cannot disconnect A from B using $s-1$ cuts, then one must be able to find s disjoint paths from A to B.

Proof It is obvious that $MAXFLOW$ is less than or equal to $MINCUT$, since if S is a set disconnecting A from B then every path from A to B must pass through S, and so one can have at most $|S|$ disjoint paths. The other direction, however, is not as obvious.

We induct on the number of edges in G. If there are no edges in G then $MAXFLOW$ and $MINCUT$ are both equal to $|A \cap B|$. Now we assume that there is at least one edge $(a \to b)$ in G, and the claim has already been proven for all fewer edges.

Let $s := MINCUT(A \to B; G)$. We have to construct s disjoint paths from A to B. Let us first look at the quotient graph $G/(a = b)$, with a and b identified into a single vertex $\{a, b\}$, and the edge $(a \to b)$ deleted. Let $S/(a = b)$ be a minimal subset of $G/(a = b)$ disconnecting $A/(a = b)$ from $B/(a = b)$. If $S/(a = b)$ has at least s elements, then we can apply the induction hypothesis to $G/(a = b)$ (which has fewer edges than G) to construct s disjoint paths in $G/(a = b)$ from $A/(a = b)$ to $B/(a = b)$, which clearly lifts to s disjoint paths in G from A to B as desired. Thus we can assume $S/(a = b)$ has fewer than s elements. But S disconnects A from B, and thus has at least s elements. The only way this can happen is if S contains both a and b, and $|S| = s$.

Now consider $MINCUT(A \to S; G)$. This must be at least s, because if we could disconnect A from S using fewer than s elements, then we could disconnect A from B using the same elements (because every path from A to B must pass through something in S). In fact we see that $MINCUT(A \to S; G - (a \to b))$ is at least s, because the edge from a to b plays no role in the connectivity between A and S. By the induction hypothesis we can thus find s disjoint paths from A to S; note each path must have a distinct endpoint since $|S| = s$. Similarly we can find s disjoint paths from S to B, with each path having a distinct initial point. The paths in the first family must be disjoint from the paths in the second family, except at S, since otherwise we could construct a path from A to B which avoids S entirely. Now all we do is concatenate the s paths from A to S with the s paths from S to B to create the s disjoint paths from A to B. ∎

Now we prove Proposition 4.4. Write $s := MAXFLOW(V_0 \to V_2; G)$. Since V_1 disconnects V_0 from V_2, we have $s \leq |V_1| < |V_0|$. On the other hand, by Menger's theorem there is a set $S \subseteq V_0 \cup V_1 \cup V_2$ of cardinality $|S| = s < |V_0|$ which disconnects V_0 from V_2.

Write $S = S_0 \cup S_1 \cup S_2$, where $S_j := S \cap V_j$. The plan now is to push all the vertices in the S_1 component of the disconnecting set over to V_0.

Let G' be the subgraph of G whose edges $E'_{0\to1} \cup E'_{1\to2}$ are the union of all the paths from $V_0 \backslash S_0$ to $V_2 \backslash S_2$. Since G has the commutative property, we see easily that G' does also. Also, S_1 disconnects V_0 from V_2 in G'. In particular, we have $MINCUT(V_0 \to V_2; G') = |S_1|$ (if the mincut was any smaller, we could contradict the minimality of $|S| = s$). Thus there are $|S_1|$ disjoint paths in G', which of course pass through distinct points in S_1. Let $W_0 \subseteq V_0 \backslash S_0$ denote the initial points of these paths and $W_1 \subseteq V_2 \backslash S_2$ denote the final points, thus $|W_0| = |S_1| = |W_2|$.

Now we use the commutative property of G'. By pulling back on these $|S_1|$ disjoint paths, we see that we have an injection from the edges in $E'_{1\to2}$ to the edges in $E'_{0\to1}$ emanating from W_0, since every edge in $E'_{1\to2}$ emanates from some vertex in S_1. Similarly, by pushing forward on these paths, we see that we have an injection from the edges in $E'_{0\to1}$ to the edges in $E'_{1\to2}$ entering W_2. The only way both of these statements can be true is if every edge in $E'_{0\to1}$ emanates from W_0, and every edge in $E'_{1\to2}$ terminates at W_2. But then we can replace $S_0 \cup S_1 \cup S_2$ by $S_0 \cup W_0 \cup S_2$ to obtain a set of cardinality s disconnecting V_0 from V_2 in G. Since

$S_0 \cup W_0 \cup W_2$ disconnects S, the set $G^2(V_0 \backslash (S_0 \cup W_0))$ must be contained in W_2. But since $s < |V_0|$, the set $V_0 \backslash (S_0 \cup W_0)$ has larger cardinality than $|W_2|$, and we are done.

6. The Cartesian product trick

We have proven Proposition 4.3 in the special case $C = 1$. Now we leverage this special case to the general case by a surprisingly powerful trick, that of taking Cartesian products. We have already seen a very tiny hint of this when we used a lifting trick to make V_0, V_1, and V_2 disjoint.

Let G and $\tilde{G}$ be two commutative graphs in two abelian groups Z and $\tilde{Z}$, with vertex sets V_0, V_1, V_2 and $\tilde{V}_0$, $\tilde{V}_1$, $\tilde{V}_2$ respectively, and edge sets $E_{0\to 1}, E_{1\to 2}$ and $\tilde{E}_{0\to 1}$, $\tilde{E}_{2\to 2}$. Then we can form a Cartesian product $G \times \tilde{G}$ in $Z \times \tilde{Z}$ with vertex sets $V_0 \times \tilde{V}_0$, $V_1 \times \tilde{V}_1$, $V_2 \times \tilde{V}_2$ and edge sets $E_{0\times 1} \times \tilde{E}_{0\times 1}$ and $E_{1\times 2} \times \tilde{E}_{1\times 2}$, where the direct sum of two edges $(a \to b)$ and $(\tilde{a} \to \tilde{b})$ is understood to be the edge $((a, \tilde{a}) \to (b, \tilde{b}))$.

One can easily verify that the Cartesian product of two commutative graphs is still commutative. For instance, the Cartesian product of $G[A, B]$ and $G[\tilde{A}, \tilde{B}]$ is $G[A \times \tilde{A}, B \times \tilde{B}]$.

Let us also define the concept of *magnification ratio*. If G is a commutative graph, we define the magnification ratio $D(G)$ to be the infimum

$$D(G) := \inf_{A' \subseteq V_0, \text{ non-empty}} \frac{|G^2(A')|}{|A'|};$$

thus $D(G)$ is the minimum amount by which G^2 expands sets. Proposition 4.4 thus asserts that if $\frac{|V_1|}{|V_0|}$ is less than 1, then so is the ratio $D(G)$. We wish to prove Proposition 4.3, which asserts that if $\frac{|V_1|}{|V_0|}$ is less than K, then the ratio $D(G)$ is less than K^2.

Two examples should be kept in mind. Let $Z = \mathbf{Z}^k$, let $A := \{0\}$, and let B be the standard basis for $\mathbf{Z}^k$. Then in the commutative graph $G_k := G[A, B]$, we have $|V_0| = 1$, $|V_1| = k$, and $D(G) = \frac{k(k-1)}{2}$ (why?). Dually, we have the reflected commutative graph $G_k^\dagger$, in which the sets V_0 and V_2 are swapped, in which $|V_0| = \frac{k(k-1)}{2}$, $|V_1| = k$, and $D(G) = \frac{2}{k(k-1)}$. Note that in both cases we are pretty close to the bound in Proposition 4.3, up to a factor of 2 or so.

Now we make a key observation connecting magnification ratios with Cartesian products:

Lemma 6.1. *Let G and $\tilde{G}$ be commutative graphs. Then $D(G \times \tilde{G}) = D(G) \times D(\tilde{G})$.*

Proof Let $d := D(G)$ and $\tilde{d} := D(\tilde{G})$. Then there exists $A' \subseteq V_0$ such that $|G^2(A')| = d|A'|$, while there exists $\tilde{A}' \subseteq \tilde{V}_0$ such that $|\tilde{G}^2(\tilde{A}')| = \tilde{d}|\tilde{A}'|$. Multiplying together, we obtain

$$|(G \times \tilde{G})^2(A' \times \tilde{A}')| = d\tilde{d}|A' \times \tilde{A}'|$$

which implies that $D(G \times \tilde{G}) \geq d\tilde{d}$.

Now we show the reverse inequality, i.e. for every $\Omega \subseteq V_0 \times V_0'$, that

$$|(G \times \tilde{G})^2(\Omega)| \geq d\tilde{d}|\Omega|.$$

The trick is to factorize the left-hand side. Let $I_{V_0} = G[V_0, \{0\}]$ be the trivial commutative graph whose vertex sets V_0, V_1, V_2 are all equal to V_0, and the edges are just the loops on V_0. Similarly define $I_{\tilde{V}_2}$. Then observe that

$$(G \times \tilde{G})^2(\Omega) = (G \times I_{\tilde{V}_2})^2(I_{V_0} \times \tilde{G})^2(\Omega);$$

this basically arises because we can think of paths in $G \times \tilde{G}$ as a path in $I_{V_0} \times \tilde{G}$, followed by a path in $G \times I_{\tilde{V}_2}$. Also, it is easy to see that $D(I_{V_0} \times \tilde{G}) = D(\tilde{G})$ and $D(G \times I_{\tilde{V}_2}) = D(G)$. The claim follows. ∎

Now we have enough tools to finish the proof of Proposition 4.3. Let G be a commutative graph such that $|V_1|/|V_0| < K$ for some $K > 1$. Let k be an integer between $2K+1$ and $2K+2$, so that $|V_1|/|V_0|^{\frac{2}{k-1}} < 1$. Then if we take the Cartesian product of G with $G_k^\dagger$, the new graph $G \times G_k^\dagger$ has a ratio $|V_1|/|V_0|$ less than 1. Thus Proposition 4.4 applies, and $D(G \times G_k^\dagger) < 1$. By Lemma 6.1, we thus have

$$D(G) < 1/D(G_k^\dagger) = \frac{k(k-1)}{2} \leq 10K^2.$$

This is almost what we want, except that we have lost a constant of 10. To remove this loss we resort to Cartesian products again. Let M be a large integer, and consider $G \times \ldots \times G$, the Cartesian product of M copies of G. Applying the bound we have with G replaced by $G \times \ldots \times G$ (and K replaced by K^M) we obtain

$$D(G \times \ldots \times G) \leq 10K^{2M}.$$

But then applying Lemma 6.1 repeatedly and then taking M^{th} roots, we obtain

$$D(G) \leq K^2$$

as desired. (This weird phenomenon occurs everywhere in this subject - all constants can be magically converted to 1 just by raising everything to sufficiently high powers. See, constants really don't matter after all!) This completes the proof of Proposition 4.3, and hence of Plünnecke's theorem.

7. Boosting the size of A'

We have just proven Plünnecke's theorem, which roughly speaking says that if A is essentially B-invariant, then some subset A' of A is essentially $B+B$-invariant. While this is a nice theorem to have, it seems a little unsatisfactory in one respect:

the set A' could be much smaller than A. Fortunately, we can remedy this quite easily.

Corollary 7.1. *Let A and B be non-empty subsets of an abelian group Z such that $|A+B| \leq K|A|$, and let $0 < \delta < 1$ be a parameter. Then there exists a subset A' of A of cardinality $|A'| \geq (1-\delta)|A|$ such that $|A'+B+B| \leq \frac{2K^2}{\delta}|A|$.*

Thus one can make A' nearly as large as A, although we begin to lose in the constants as one tries to push A' closer and closer to being 100% of A. One cannot improve this factor K^2/δ by much; see Q11.

Proof We shall iterate Plünnecke's theorem by performing the algorithm. Let $A_0 := A$. By Plünnecke's theorem we may find a non-empty subset A'_0 of A_0 such that

$$|A'_0 + B + B| \leq \frac{|A_0+B|^2}{|A_0|^2}|A'_0| \leq \frac{K^2|A|^2}{|A_0|^2}|A'_0|.$$

Now set $A_1 := A_0 \backslash A'_0$. If $|A_1| < \delta|A|$, we terminate the algorithm. Otherwise we apply Plünnecke's theorem again to find a non-empty subset A'_1 of A_1 such that

$$|A'_1 + B + B| \leq \frac{|A_1+B|^2}{|A_1|^2}|A'_1| \leq \frac{K^2|A|^2}{|A_1|^2}|A'_1|.$$

Now we set $A_2 := A_1 \backslash A'_1$. If $|A_2| < \delta|A|$ we terminate the algorithm. We repeat this procedure until it terminates (which it must, as each new A_j is strictly smaller than the previous one) at some A_k with $|A_k| < \delta|A|$. We then set $A' := A \backslash A_k$; clearly we have $|A'| \geq (1-\delta)|A|$. Also, by construction

$$\begin{aligned}
|A'+B+B| &\leq \sum_{j=0}^{k-1} |A'_j + B + B| \\
&\leq \sum_{j=0}^{k-1} \frac{K^2|A|^2}{|A_j|^2}|A'_j| \\
&= K^2|A|^2 \sum_{j=0}^{k-1} \frac{|A_j| - |A_{j+1}|}{|A_j|^2} \\
&\leq K^2|A|^2 (\frac{1}{|A_{k-1}|} + \sum_{j=0}^{k-2} \frac{|A_j| - |A_{j+1}|}{|A_j||A_{j+1}|}) \\
&= K^2|A|^2 (\frac{1}{|A_{k-1}|} + \sum_{j=0}^{k-2} \frac{1}{|A_{j+1}|} - \frac{1}{|A_j|} \\
&\leq 2K^2|A|^2/|A_{k-1}| \\
&\leq 2K^2|A|/\theta
\end{aligned}$$

as desired. ∎

Now one could ask why we couldn't boost A' up all the way to A and obtain a result of the form "If A is essentially B-invariant, then it is also essentially $B+B$-invariant". The following example (due to Ruzsa) shows that this is too naive.

Proposition 7.2. [4] *Let n be a large integer. Then there exists sets A of size $|A| \sim n^2$ and B of size $|B| \sim n$, such that $|A+B| \sim n^2$ but $|A+B+B| \sim n^3$.*

Proof We shall use the abelian group $\mathbf{Z}^2$. We shall need a set B for which $B+B$ is much larger than B; a typical choice is the set

$$B := \{(i,0) : i = 1, \ldots, n\} \cup \{(0,j) : j = 1, \ldots, n\};$$

observe that $|B| \sim n$ but that $|B+B| \sim n^2$.

Now let A_0 be the $n \times n$ square $\{(i,j) : i,j = 1, \ldots, n\}$; observe that $|A_0| \sim n^2$, $|A_0+B| \sim n^2$, and $|A_0+B+B| \sim n^2$. Thus A_0 is not a counterexample to our claim; it is both essentially B-invariant and essentially $B+B$-invariant. However, we can worsen things by adding a small number of maverick elements to A_0.

Specifically, let A_1 be any collection of n points, sufficiently separated in space. Then $|A_1| \sim n$, $|A_1+B| \sim n^2$, and $|A_1+B+B| \sim n^3$. If we then set $A := A_0 \cup A_1$ then we obtain our desired counterexample. ∎

Note how a key fact used here was that $B+B$ was much larger than B. If $B+B$ was the same size as B one could do much better; see Q10.

Thus we really do need to remove some exceptional elements before passing from essentially B-invariant sets to $B+B$-invariant sets. We will see more examples like this later on.

8. Some consequences of Plünnecke's theorem

From Plünnecke's theorem we have some immediate corollaries. To state them, we need some notation. Let us write $2A$ for $A+A$, $3A$ for $A+A+A$, etc, with the convention that $0A = \{0\}$. (This should be kept separate from $2 \cdot A := \{2a : a \in A\}$, $3 \cdot A := \{3a : a \in A\}$, etc.). Plünnecke's theorem thus says that if you can control $A+B$, then you can also control $A'+2B$ for some refinement A' of A. Iterating this, we can get

Corollary 8.1 (Iterated Plünnecke's theorem)**.** *Let A, B be two finite non-empty subsets of an abelian group Z, and suppose that $|A+B| \leq K|A|$ for some real number $K \geq 1$. Then for every $n = 1, 2, 3, \ldots$ there is some non-empty subset A_n of A such that $|A_n + nB| \leq K^{C(n)}|A_n|$ for some absolute constant $C(n)$.*

Indeed, this claim when m is a power of two just follows by iterating Plünnecke's theorem $\log_2 m$ times, and then one makes the simple observation that once the theorem is proven for any large n, it automatically holds for smaller n (because if $n < n'$ then nA is contained in a translate of $n'A$).

One can in fact take $C(n) = n$; see Exercise 6.

From the above Corollary and Lemma 3.1 we see that

$$|nB - nB| \leq \frac{|A_n + nB|^2}{|A_n|} \leq K^{2C(n)}|A_n| \leq K^{2C(n)}|A|.$$

Since nB is contained in a translate of mB when $n \geq m$, we thus easily obtain

Corollary 8.2 (Sumset estimates.)**.** *Let A, B be two finite non-empty subsets of an abelian group Z, and suppose that $|A+B| \leq K|A|$ for some real number $K \geq 1$. Then for every $n, m \geq 0$ we have the bound*

$$|nB - mB| \leq K^{C(n,m)}|A|.$$

Thus if A is essentially B-invariant, then the sum or difference of an arbitrary number of copies of B cannot get much larger than A. This is consistent with our intuition that if A is essentially B-invariant, then B is essentially contained inside a coset of a subgroup G, and A is essentially the union of such cosets.

Let us call a set A *essentially closed under addition* if $|A + A| \lesssim |A|$. The sumset estimates, specialized to the case $A = B$, then say that if A is essentially closed under addition, then it is also essentially closed under subtraction, and more generally any expression of the form $A \pm A \pm A \ldots \pm A$ also has size comparable to $|A|$. This means that A behaves essentially like an additive subgroup of Z (especially if we translate A so that A contains the origin) - with one basic caveat: the more arithmetic operations we perform on A, the larger the exponent $C(n, m)$ can get, and so eventually we might lose our closure properties if we insist on doing too much arithmetic. (A basic example to keep in mind here is when Z is the integers and $A := \{1, \ldots, N\}$. This set is more or less closed under arithmetic operations as long as you don't take too many of them.)

The exponents $C(n)$ and $C(n, m)$ can be computed, and the values that were worked out here are not the best exponents known, but we will not bother to make these exponents explicit and sharp in these lectures.

One can swap addition and subtraction here; if A is essentially closed under subtraction, then it is also essentially closed under addition (one just uses $B := -A$ instead of $B := +A$ in the sumset estimates. See also Exercise 7).

One should compare this situation with that of subgroups and submonoids. A submonoid G is closed under sums of two elements: $G + G \subseteq G$, and hence by iteration it is closed under sums of arbitrary elements: $nG \subseteq G$. If G is a finite monoid, then it must have torsion, and thus $n \cdot G = -G$ for some large n, and thus G is also closed under negation. Hence G is a subgroup and thus closed under all arithmetic operations. The above sumset estimates can thus be viewed as a perturbation of this basic principle: closure under one arithmetic operation (+ or -) implies closure under the other operation, and under iterates of these operations.

9. Covering one set by another

Compare the following two notions:

- A set G is closed under addition if $G + G = G$.
- A set A is essentially closed under addition if $|A + A| \sim |A|$.

Up until now I have been trying to convince you that the second notion is basically a mild generalization of the first, although it is much more robust (it is stable under translations, finite unions of translations, and also refinements - passing from A to a subset A' of comparable cardinality) But the observant reader should have noted that there seems to be a large gap between the two, because the first notion is about equating two *sets*, whereas the second notion is only about comparing two *cardinalities*. Saying that $|A + A|$ and $|A|$ have comparable cardinalities does not necessarily mean that they have similar structure... does it?

On the other hand, we do know that if $|A+A| = |A|$, then A is indeed a translate of a subgroup G, and so $A = G+x$ for some x (see Exercise 3). So, up to translations (which, it should be clear by now, are quite harmless), the extreme limiting case of the second notion does indeed correspond to the first. So this helps close the gap between the two notions, at least in the extreme limiting case when $|A + A| = |A|$. However, this doesn't help us when, say, $|A + A| = |A| + 1$; now the two notions seem quite different again.

To close this gap we use yet another clever elementary lemma of Imre Ruzsa.

Lemma 9.1 (Ruzsa's quotient lemma)**.** [5] *Let A and B be finite sets. Then there exists a set X of cardinality $|X| \leq \frac{|A+B|}{|A|}$ such that $B \subseteq X + A - A$. In other words we can cover B by at most $\frac{|A+B|}{|A|}$ translates of $A - A$.*

Proof Consider the sets $x + A$, where x ranges over B. These are a collection of sets contained in $A + B$, each with cardinality $|A|$. Let us consider a maximal disjoint family of such sets, i.e. a collection of the form $\{x + A : x \in X\}$ such that the $x + A$ are all disjoint. Such a collection can easily be obtained by the greedy algorithm. By the disjointness, we see that there can be at most $\frac{|A+B|}{|A|}$ such sets. Now let b be any element of B. By maximality, $b + A$ must intersect one of the $x + A$, i.e. $b \in x + A - A$ for some $x \in X$, i.e. $b \in X + A - A$, thus $B \subseteq X + A - A$ as desired. ∎

Note how powerful this "maximal disjoint family" trick is. Note what is going on in the case when A is a subgroup, and B is a union of cosets of A (so that $|A + B| = |B|$). Then what X is doing is simply selecting a single element from each coset inside B, and of course X will end up having cardinality $|B|/|A|$. Thus this lemma is a generalization of the familiar algebraic notion of taking a quotient by a subgroup, which of course turns each coset into a single element of the quotient space.

To apply this quotient lemma, suppose that A is essentially closed under addition: $|A+A| \sim |A|$. Then by sumset estimates we also have $|nA - mA + A| \sim |A|$. By the above lemma, we can thus cover $nA - mA$ by $O(1)$ translates of $A - A$. In particular, if we set $G := A - A$, then $G + G$ and $G - G$, and more generally $nG - mG$, is contained in $O(1)$ translates of G.

Thus: if A is essentially closed under addition, then there is a slightly larger set $G := A - A$, which is "even more closed under addition" in the sense that $G + G$ or $G - G$ is contained in $O(1)$ translates of G. Note that A is itself contained in a translate of G, and has comparable size, $|A| \sim |G|$. This provides a satisfying connection between the approximate notion of "essentially closed under addition" and the more precise notion of "genuinely closed under addition".

One may ask why we need this intermediate set G - why couldn't we just use A directly? In other words, could we say that if A is essentially closed under addition, that $A + A$ is contained in $O(1)$ translates of A, etc? Looking at our first model example - that of $A = \{1, \ldots, N\}$ in the integers $Z = \mathbf{Z}$, this seems plausible. However, this statement is false. Let us set A to be a random subset of $\{1, \ldots, N\}$ by flipping an independent coin for each integer from 1 to N and placing that integer in A if the coin turns up heads. This is almost surely a set of cardinality $\approx N/2$. (By "almost surely" I mean that the probability of this statement being false is exponentially small in N). The set $A + A$ almost surely contains a large interval, e.g. $\{N/2, \ldots, 3N/2\}$ (why?). However, it is difficult to cover such an interval with only a small number of translates of A. For instance, if you were to use two translates of A to cover $\{N/2, \ldots, 3N/2\}$, you would almost surely only be able to cover $3/4$ of this interval (why?). With three translates you could cover $7/8$ at most, almost surely, and so forth. Indeed, if you are good with probability you will see that you will almost surely need $\sim \log N$ translates to cover all of $A + A$.

This illustrates an important principle: even if a set A contains a lot of "holes" - in this case, it is missing half the elements of the interval $\{1, \ldots, N\}$, the set $A - A$ is much better behaved and will have almost no holes.

It turns out that this logarithmic loss is more or less sharp: if you do insist on trying to cover things using A instead of $A - A$, you only lose that logarithmic factor.

Theorem 9.2. *Let A be a finite subset of an abelian group Z of cardinality $|A| = N > 1$ which is essentially closed under addition. Then for any fixed $n, m \geq 0$ there is a set $X = X(n,m)$ with cardinality $|X| \lesssim \log N$ such that $mA - nA \subseteq X + A$.*

In other words, a set like $A - A$ (or more generally $mA - nA$) can be covered by $O(\log N)$ translates of A.

To illustrate the basic idea, let us first show a baby version of this theorem, which is already somewhat interesting:

Lemma 9.3. *Let G be a finite group of cardinality $|G| = N$, and let A be a non-empty subset of G. Then one can cover G using at most $O(\frac{N}{|A|} \log N)$ translates*

of A; in other words, there exists a subset X of G with cardinality $|X| \lesssim \frac{N}{|A|} \log N$ such that $X + A = G$.

Proof We shall use a random argument. We may assume that $N \gg 1$ is large. Let $1 \leq r \ll N$ be a number to be chosen later. We construct X by taking each element of G and placing it in X with probability r/N, with each element being chosen independently of all the others. Clearly $|X| = \sum_{x \in G} \chi_X(x)$ has expected value of r, since $\chi_X(x)$ has expectation r/N for each $x \in G$. Furthermore, it is not hard to see that

$$|X|^2 = \sum_{x \in G} \sum_{y \in G} \chi_X(x)\chi_X(y) = |X| + \sum_{x,y \in G: x \neq y} \chi_X(x)\chi_X(y)$$

has expectation $r + N(N-1)\frac{r}{N}\frac{r}{N} \sim r^2$, for similar reasons. By Markov's inequality we thus see that $|X| = O(r)$ with probability at least 99%.

Now we compute the expected value of $|X + A|$. First we take any element $g \in G$ and ask what is the probability that g lies in $X + A$. This is the same as saying that X has some intersection with $g - A$. But $g - A$ has $|A|$ elements and each element has independently an r/N chance of lying in X. Thus the probability that X intersects $g - A$ is

$$1 - (1 - \frac{r}{N})^{|A|}.$$

Summing over all g, we see that the expected value of $|X + A| = \sum_{g \in G} \chi_{X+A}(g)$ is

$$N - N(1 - \frac{r}{N})^{|A|} \geq N - Ne^{-r|A|/N},$$

since $e^{-x} \geq 1 - x$ for all $0 \leq x \leq 1$ by convexity of e^{-x}. Thus we can choose $r = O(\frac{N}{|A|} \log N)$ such that $|X + A|$ has expected value at least $N - 0.01$, which means that it has to be equal to N at least 99% of the time. But then we see that $X + A = G$ and $|X| = O(r)$ at least 98% of the time, and we are done (since we only need to find at least one such X. ∎

Another way to prove this lemma is sketched in Exercise 8. Note that this argument also shows that if you only want $X + A$ to fill up, say, 90% of G, rather than all of G, then you only need $O(N/|A|)$ translates rather than $O(\frac{N}{|A|} \log N)$. (Thus, fighting for that last 10% of G takes a lot of effort! Compare with the discussion on exceptional sets after Corollary 7.1).

Now we prove Theorem 9.2. Fix m, n, and write $B := mA - nA$. The idea is now to run the same argument, but using $A - A$ as a proxy for the group G; we already saw earlier that in many ways $A - A$ functions like a group. However, for technical reasons we first need to improve Lemma 9.1 a little bit:

Lemma 9.4 (Improved quotient lemma)**.** *Let A and B be finite sets. Then there exists a set X in Z of cardinality at most $\frac{2|A+B|}{|A|}$ such that $X + A - A$ covers B, and moreover for every $y \in B$ there are at least $|A|/2$ triplets $(x, a, a') \in X \times A \times A$ such that $x + a - a' = y$.*

Thus not only does $X + A - A$ cover B; it covers B multiple times. The only loss here compared to Ruzsa's quotient lemma is that X could be twice as big.

Proof We perform the following algorithm. Initialize X to be the empty set, so that $X + A - A$ is also the empty set. We now run the following loop. If we cannot find any element y in B which is "sufficiently disjoint from $X + A - A$" in the sense that $|(y + A) \cap (X + A)| < |A|/2$, we terminate the algorithm. Otherwise, if there is such an element y, we add it to X, and then repeat the algorithm.

Every time we add an element to X, the size of $|X + A|$ increases by at least $|A|/2$, by construction. However, $X + A$ must always lie within the set $B + A$. Thus this algorithm terminates after at most $\frac{2|A+B|}{|A|}$ steps.

Now let y be any element of B. By construction, we have $|(y+A)\cap(X+A)| \geq |A|/2$, and hence y has at least $|A|/2$ representations of the form $x + a - a'$ for some $(x, a, a') \in X \times A \times A'$, as desired. ∎

We are now ready to cover B by translates of A. We apply Lemma 9.4 to create our set X; note that $|A| = N$ and $|B + A| = O(N)$ by sumset estimates, so $|X| = O(1)$. Let $1 \leq r \ll N$ be a number to be chosen later (we may of course assume that N is large since the claim is trivial for N small). Let Y be a subset of $X \times A$ chosen randomly, so that each pair (x, a') of $X \times A$ has an independent probability of $\frac{r}{N|X|}$ of being selected. By the argument in Lemma 9.3 we see that $|Y| = O(r)$ with probability at least 99%. Now pick any element $y \in B$ and let us ask what is the probability that y lies in $Y + A$. From Lemma 9.4 we know that there are at least $N/2$ pairs $(x, a') \in X \times A'$ such that y lies in $x - a' + A$. Each of these pairs has independently a $r/N|X|$ chance of lying in Y. Thus the probability that y lies in $Y + A$ is at least

$$1 - (1 - \frac{r}{N|X|})^{N/2} \geq 1 - e^{-r/2|X|}.$$

Thus the expected value of $|(Y + A) \cap B$ is at least

$$|B|(1 - e^{-r/2|X|}).$$

If we choose r to be a sufficiently large multiple of $\log N$, this is greater than $|B| - 0.01$, and so we have $B \subseteq Y + A$ with probability 99%. Thus we have $B \subseteq Y + A$ and $|Y| \lesssim \log N$ for 98% of the choices of Y, and we are done. ∎

We now have quite a satisfactory theory of the structure of sets A which are essentially closed under addition or subtraction; we now know that they are also closed under more complicated sums and differences, and we can cover these more complicated sums and differences by a small number of translates of either A or $A - A$, thus keeping close to our heuristic that A should basically be a translate of a subgroup, or a large subset thereof. Even this is not the best thing we can say; we can make this heuristic even more precise, thanks to a deep theorem known as *Freiman's theorem*, but we won't tackle it until the next set of notes.

10. $A+B$ and $A-B$

In the previous section we saw that if A was essentially closed under addition - so that $|A+A| \sim |A|$, then it was also essentially closed under subtraction $|A-A|$, and indeed we were able to say a lot more than this. Now one could ask whether the same thing is true when one has two sets instead of one. For instance, if we know that $|A+B| \sim |A|$, does this mean that $|A-B| \sim |A|$?

On the positive side, we have Q3, which implies that if $|A+B| = |A|$, then B is contained in some subgroup G, and A is the union of cosets of G, and so $|A-B| = |A|$ as well. On the other hand, this argument also gives that $|A+B+B| = |A|$, and we already know from Proposition 7.2 that $A+B+B$ is a problem. As it turns out, $A-B$ is also not well behaved, and for similar reasons:

Proposition 10.1. [4] *For any integer $K \geq 1$, there exist finite subsets A, B of an abelian group Z such that $|A+B| \sim |A|$ but $|A-B| \gtrsim |A|^{2-\log 6/\log 7}$.*

Proof One of the key ingredients of Proposition 7.2 was to use a B such that $B+B$ was a very different size from B. Here, the analogous trick is to use a B such that $B-B$ is a very different size from $B+B$. To do this, first observe that we clearly cannot use a B which is symmetric, e.g. $B = -B$, or more generally $B = x_0 - B$. The simplest non-symmetric set is $B := \{0,1,3\}$ in the group $\mathbf{Z}/7\mathbf{Z}$; for this set, $B-B = \{-3,-2,-1,0,1,2,3\}$ has cardinality 7, while $B+B = \{0,1,2,3,4,6\}$ has cardinality 6. This difference may not look like much, but we can use the Cartesian product trick to boost this. Let N be a large integer, and let $B := \{0,1,3\}^N$ in the group $Z := (\mathbf{Z}/7\mathbf{Z})^N$. Then $|B| = 3^N$, $|B-B| = 7^N$, and $|B+B| = 6^N$.

Thus if we were to set $A = B$, then $A-B$ is significantly larger than $A+B$. This however is not enough, because $A+B$ is much larger than A.

Another option is to set $A = (\mathbf{Z}/7\mathbf{Z})^N$. Then $A+B$ is the same size as A, but unfortunately so is $A-B$.

To get a genuine counterexample we shall take a combination of the two previous examples. Specifically, we set $Z := (\mathbf{Z}/7\mathbf{Z})^N \times \mathbf{Z}$, so that Z contains an infinite number of copies of $(\mathbf{Z}/7\mathbf{Z})$. We let B be as before (identifying $(\mathbf{Z}/7\mathbf{Z})^N$ with $(\mathbf{Z}/7\mathbf{Z})^N \times \{0\}$), and let A consist of one copy of $(\mathbf{Z}/7\mathbf{Z})^N$ and roughly $(7/6)^N$ copies of B. Then A has cardinality comparable to $7^N + (7/6)^N 3^N \sim 7^N$, $A+B$ has cardinality $7^N + (7/6)^N 6^N \sim 7^N$, but $A-B$ has cardinality $7^N + (7/6)^N 7^N \sim (49/6)^N$. The proof follows. ∎

Thus essentially B-invariant sets aren't necessary $(-B)$-invariant. It is possible to make an example in the torsion-free group $\mathbf{Z}$ instead of the group $(\mathbf{Z}/7\mathbf{Z})^N \times \mathbf{Z}$; the idea is to use the base 7 representation of the integers. We leave the details to the reader (or see [4]).

In the case of $A+B$ and $A+B+B$, Plünnecke's theorem (and Corollary 7.1) allow us to salvage a satisfactory result if we pass from A to a subset A'; we know that

we can make A' nearly the same size as A, though we begin to pay if we try to make A' too close to A. We could hope to pull off a similar trick with $A+B$ and $A-B$, but now we will lose a logarithm.

Proposition 10.2. [4] *For any integer $n \geq 1$, there exist finite non-empty subsets A, B of an abelian group Z such that $|A| \sim |A+B| \sim C^n$, and $|B| \sim n$, but $|A'-B| \gtrsim n|A'|$ for all non-empty subsets A' of A.*

Proof The trick here is to make $A+B$ "small" (so many sums of the form $a+b$ collide) but $A-B$ "large" (so many differences of the form $a-b$ are distinct; in fact, since $|A'-B| \sim |A'||B|$, we need a very large fraction of the differences to be distinct). Furthermore, this largeness of $A-B$ has to be fairly "uniform"; unlike Proposition 10.2, we cannot get away with using a small exceptional set to generate this largeness of $A-B$.

We shall work in the group $\mathbf{Z}^{2n}$. Let A be the set

$$A := \{(x_1, x_2, \ldots, x_{2n}) \in \mathbf{Z}^{2n} : x_1 + \ldots + x_n = n; x_1, \ldots, x_n \geq 0\}.$$

In other words, A is the set of partitions of n consecutive objects into $2n$ groups. The number of such partitions is $\left(\begin{array}{c} 3n-1 \\ 2n-1 \end{array} \right)$, because every time you choose $2n-1$ objects out of an ordered sequence of $3n-1$ objects, the remaining n objects are partitioned into $2n$ groups, and this map from choices to partitions is bijective. From Stirling's formula $n! \sim n^n e^{-n} n^{-1/2}$ we thus see that $|A| \sim C^n$ for some C (actually $C = 27/4$, if you must know).

Now let $B := \{e_1, \ldots, e_{2n}\}$ be the basis elements of $\mathbf{Z}^{2n}$, thus $|B| = 2n$. Clearly $A+B$ is equal to

$$A+B = \{(x_1, x_2, \ldots, x_{2n}) \in \mathbf{Z}^{2n} : x_1 + \ldots + x_n = n+1; x_1, \ldots, x_n \geq 0\}$$

and thus has cardinality

$$|A+B| = \left(\begin{array}{c} 3n \\ 2n-1 \end{array} \right) = \frac{3n}{n+1} \left(\begin{array}{c} 3n-1 \\ 2n-1 \end{array} \right) \sim |A|.$$

Now look at $A-B$. The problem here is that every element $(x_1, x_2, \ldots, x_{2n})$ in A has at least n co-efficients which are zero, and thus generates n elements in $A-B$ for which exactly one of the coefficients is equal to -1. Conversely, given such an element of $A-B$ one can reconstruct the elements of A and B which generated it; i.e. these elements of $A-B$ are all distinct. Thus $|A-B| \geq n|A|$, and more generally $|A'-B| \geq n|A'|$ for every subset A' of A. ∎

The factor n is pretty small compared to A, it is comparable to $\log |A|$. We now show that this logarithmic loss is about as bad as it can get:

Proposition 10.3. *Let A, B be finite non-empty subsets of an abelian group Z such that $|A| \sim |A+B|$. Then for any $\varepsilon > 0$ we can find some subset A' of A such that $|A'| \sim |A|$ and $|A'-B| \leq C_\varepsilon |A|^{1+\varepsilon}$.*

Proof We let N be a large number to be chosen later. By the Iterated Plünnecke theorem we can find a subset A' of A such that $|A' + 2^N B| \leq C(N)|A'|$; by using

either the statement or proof of Corollary 7.1 one can also make $|A'| \sim |A|$. Now consider the sequence of numbers

$$1, |B|, |2B|, |4B|, \ldots, |2^N B|.$$

These are an increasing sequence of numbers bounded above by $C(N)|A|$, so by the pigeonhole principle there exists $0 \leq j < N$ such that

$$|2^{j+1}B| \leq C(N)|A|^{1/N}|2^j B|.$$

Now let $B_* := 2^j B$, thus

$$|B_* + B_*| \leq C(N)|A|^{1/N}|B_*|.$$

By Lemma 3.1 we have

$$|A' - B_*| \leq \frac{|A' + B_*||B_* + B_*|}{|B_*|} \leq C(N)|A|^{1/N}|A' + B*| \leq C(N)|A|^{1/N}|A|$$

as desired, if N is chosen large enough. ∎

This exploiting of the pigeonhole principle in a long monotone sequence is a useful trick: the point is that a long monotone sequence must be close to constant at some point along the sequence.

This above discussion concludes the limit of what we will say about the *cardinality* of sets such as A, $A+B$, $A+A$, $A-B$, etc. In the next set of notes we will analyze the *structure* of these sets also.

11. Exercises

- Q1. Let $m \geq 1$, $n \geq 1$, and $m + n - 1 \leq s \leq mn$. Show that there exists sets A, B of integers such that $|A| = m$, $|B| = n$, and $|A + B| = s$. (Thus there are no further relationships on $|A|, |B|, |A+B|$ beyond the bounds $|A| + |B| - 1 \leq |A + B| \leq |A||B|$ already proven.)
- Q2. Let A, B be non-empty finite sets of integers with $|A|, |B| \geq 2$. Show that $|A+B| = |A|+|B|-1$ if and only if A and B are arithmetic progressions with exactly the same spacing.
- Q3. Let A, B be non-empty finite subsets of an abelian group Z. Show that $|A + B| = |A|$ if and only if A is the union of cosets of some finite subgroup G of Z, and B is contained in a translate of the same subgroup G. (As a corollary, note that $|A+A| = |A|$ if and only if A is a translate of a subgroup G).
- Q4*. Let A, B be non-empty finite subsets of the field $\mathbf{Z}/p\mathbf{Z}$, and suppose that $|A + B| = |A| + |B| - 1 < p$. Show that A and B are arithmetic progressions with exactly the same spacing. (Hint: Use the Dyson e-transform proof, and induct on the size of A. The key is to show that if A' and B' are progressions with the same spacing, and $A' + B' = A + B$, then A and B are also progressions with the same spacing. You may find it convenient to use translations and isomorphisms to transform arithmetic progressions to an interval $\{1, 2, \ldots, m\}$. Note that if the interval is small enough, then one can perform a Freiman isomorphism to transform $\mathbf{Z}/p\mathbf{Z}$ back to $\mathbf{Z}$.)

- Q5. Suppose that A is a finite collection of women, and B is a finite collection of men, such that some of the women are compatible with some of the men (i.e. there is a graph G from A to B describing the compatibility). For every set $A' \subseteq A$ of women, let $G(A')$ denote the set of men who are compatible with at least one woman in A'. Show that it is possible for each woman in A to marry a separate man with whom she is compatible, if and only if one has $|G(A')| \geq |A'|$ for all $|A'| \subseteq A$. (This is known as *Hall's marriage theorem*, and can be deduced from Menger's theorem. One could say that this theorem lists all the possible (graph-theoretical) obstructions to marriage.)
- Q6**. Extend Plünnecke's theorem to higher sums. More precisely, show that if A, B are finite non-empty subsets of an abelian group Z such that $|A+B| \leq C|A|$, and $n \geq 1$, then there is a non-empty subset A' of A such that $|A'+nB| \leq C^n|A'|$. (Note that iterating Plünnecke's theorem only gives this for n which are powers of two. For general n, the main trick is to start with a separating set of the obvious graph on $V_0 \cup \ldots \cup V_n$, and move all the intermediate separating points down to V_0.)
- Q7. Suppose that A, B are finite sets such that $|A+B| \sim |A| \sim |B|$. Show that A, B, and $A \cup B$ are all essentially closed under addition (Hint: Use the iterated Plünnecke theorem). Conclude that any expression of sums and differences of A and B, e.g. $A+A-A+B-B-B$, has cardinality comparable to $|A|$.
- Q8. Let A and B be any two non-empty subsets of a finite abelian group G. Use the pigeonhole principle to show that there exists a translate $B+x$ of B in G such that $|A \cap (B+x)| \geq \frac{|A||B|}{|G|}$, and hence that

$$\frac{|G| - |A \cup (B+x)|}{|G|} \leq \frac{|G|-|A|}{|G|}\frac{|G|-|B|}{|G|}.$$

  Iterate this fact to give a different proof of Lemma 9.3.
- Q9. Show that for any $n \geq 1$, there exists an abelian group Z and a finite set U such that $|U| = 4^n$, $|U+U| = 10^n$, and $|U+U-U| = 28^n$. (Hint: first do the case $n = 1$). Conclude that one cannot hope to replace the left-hand side of Lemma 3.1 by $|U+V-W|$, even if we are prepared to lose a constant or a logarithm in the estimate.
- Q10. Let A, A', B be finite non-empty subsets of an abelian group Z such that B is essentially closed under addition. The purpose of this question is to show that B behaves much like a subgroup of Z for purposes of quotienting out by B.
- (a) Show that $|A+B| \sim |A|$ if and only if A is contained in $O(|A|/|B|)$ translates of B.
- (b) Show that if $|A+B| \sim |A|$, then $|A+mB-nB| \sim |A|$ for all m, n (with the constants depending on m and n, of course.
- (c) Show that $|A+A'+B| \lesssim \frac{|A+B||A'+B|}{|B|}$ (compare this with Lemma 3.1).
- Q11. Let $0 < \delta < 1$. Show that there exists finite non-empty subsets A,B of an abelian group $\mathbf{Z}$ such that $|A+B| \sim |A|$ such that for every subset A' of A for which $|A'| \geq (1-\delta)|A|$, we have $|A'+B+B| \gtrsim |A|/\delta$. (Hint: adapt the proof of Proposition 7.2). Thus one cannot get rid of the $1/\delta$ factor in Corollary 7.1.

References

[1] B. Green, *Edinburgh lecture notes on Freiman's theorem*, preprint.
[2] M. Nathanson, *Additive number theory. Inverse problems and the geometry of sumsets*, Graduate Texts in Mathematics 165, Springer-Verlag, New York, 1996.
[3] H. Plünnecke, *Eigenschaften un Abschätzungen von Wirkingsfunktionen*, BMwF-GMD-22 Gesellschaft für Mathematik und Datenverarbeitung, Bonn 1969.
[4] I. Ruzsa, *Sums of finite sets*, Number Theory: New York Seminar; Springer-Verlag (1996), D.V. Chudnovsky, G.V. Chudnovsky and M.B. Nathanson editors.
[5] I. Ruzsa, *An analog of Freiman's theorem in groups*, Structure theory of set addition, Astérisque No. 258 (1999), 323–326.

Department of Mathematics, UCLA, Los Angeles CA 90095-1555

E-mail address: tao@math.ucla.edu

LECTURE NOTES 2 FOR 254A

TERENCE TAO

1. The structure of sets essentially closed under addition

Let Z be an abelian group, and let A be a non-empty finite subset of Z which is essentially closed under addition in the sense that $|A+A| \sim |A|$. In the previous week's notes, we were able to deduce quite a bit of information about A. For instance, from sumset estimates we know that any additive combination $A \pm A \pm A \ldots \pm A$ of A also has cardinality comparable to A. In fact by combining this with Rusza's quotient lemma we have a stronger statement, that any such additive combination is contained in the union of $O(1)$ translates of $A-A$. In particular the set $G := A-A$ is very close to being a subgroup of Z, in the sense that $G+G \subseteq X+G$ for some small set X. Throughout these notes we call a set X *small* if we have $|X| = O(1)$.

Let us make some notation. We say that a set A' is a *refinement* of another set A if $A' \subseteq A$ and $|A'| \sim |A|$. We say that a set A' is a *small convolution* of another set A if $A' = X+A$ for some small set A. We say that A' is a *refined small convolution* of A if it is the refinement of a small convolution of A. Thus the set $G = A-A$ mentioned above is a refined small convolution of itself. Observe that the relation of being a refined small convolution is transitive (but not quite symmetric, recall that to cover a set by its refinement could require a logarithmic number of covers).

This is already a fairly satisfactory state of affairs. For comparison with the exactly closed under addition case, if we know that $|A+A| = |A|$ then we know that A is a translate $A = x+G$ of some subgroup G of Z, so in particular $G = A-A$ and $G+G = G$.

However, we can still ask for more. We know a lot about the structure of finite abelian groups G; they are the direct product of a finite number of cyclic groups (we won't prove this fact here, but it is standard in any graduate algebra course). This is a very explicit description of these sets. We could ask for a similarly explicit description of sets essentially closed under addition.

To get an idea of what one could hope for, note first that if A is essentially closed under addition, and A' is a refinement, then A' is also essentially closed under addition, since

$$|A'| \leq |A'+A'| \leq |A+A| \sim |A| \sim |A'|.$$

Also, if A is essentially closed under addition, and A' is a small convolution of A, i.e. $A' = X+A$ for some small set X, then A' is also essentially closed under

addition, since

$$|A'| \leq |A' + A'| = |X + X + A + A| \leq |X + X||A + A| \lesssim |A| \leq |A'|.$$

Thus any characterization of sets essentially closed under addition must be stable under both refinement and small convolutions, and hence under refined small convolutions. It must also be stable under Cartesian products and bijective Freiman homomorphisms of order at least 2; to remind you of the terminology, let us recall

Definition 1.1. Let $k \geq 2$, let $A \subset Z$ be the subset of one abelian group, and let $B \subset Z'$ be the subset of another abelian group.

A *Freiman homomorphism of order* k $\phi : A \to B$ is a map such that for any $x_1, \ldots, x_k, y_1, \ldots, y_k \in A$, we have

$$\phi(x_1) + \ldots + \phi(x_k) = \phi(y_1) + \ldots + \phi(y_k)$$

whenever

$$x_1 + \ldots + x_k = y_1 + \ldots + y_k.$$

If moreover ϕ is bijective and the inverse is also a Freiman homomorphism of order k, we say that ϕ is a *Freiman isomorphism of order* k.

Thus, for instance, bijective Freiman homomorphisms of order at least 2 preserve the size of A and cannot increase the size of $A + A$. Note that bijective Freiman homomorphisms are not necessarily Freiman isomorphisms, which would in fact *preserve* the size of $A + A$.

In light of these invariances, the best structure theorem one can hope for is something of the form

Conjecture 1.2. *Let A be essentially closed under addition. Then A is a refined small convolution of a set P, where P is some explicit set which is very similar to a product of finite cyclic groups, which is manifestly essentially closed under addition, and is also stable under Cartesian products and bijective Freiman homomorphisms of order at least 2.*

It turns out that this conjecture is more or less true, for very satisfactory classes of objects P. When the ambient group Z is torsion free or has very large torsion, one can take P to be something called a *generalized arithmetic progression* and the result is known as Freiman's theorem. When the ambient group Z has small torsion, then one can take P to be a genuine subgroup of Z and the result is due to Rusza (and is in fact quite short, given all the machinery we've already developed). There are intermediate cases where Z has intermediate or mixed torsion where one can state some theorems of a similar flavor, but we won't do so here.

2. The bounded torsion case

We begin with the case where Z has bounded torsion, i.e. there is some small positive integer $r = O(1)$ such that $rx = 0$ for all $x \in Z$. For instance, Z could be the direct sum of a large number of copies of $\mathbf{Z}/r\mathbf{Z}$. In this case we have

Theorem 2.1. [?] *Let Z have bounded torsion, and let $A \subseteq Z$ be essentially closed under addition. Then A is a refinement of a subgroup G of Z.*

Clearly the converse is true: subgroups are essentially closed under addition, and hence so are refinements. One may wonder what happed to the small convolutions, but in the bounded torsion case a small convolution $X + G$ of a subgroup G is a refinement of the slightly larger subgroup $\langle X \rangle + G$, where $\langle X \rangle$ is the subgroup generated by X (and clearly has cardinality at most $r^{|X|} = O(1)$. Clearly the subgroup property is stable under Cartesian products; it takes a little more effort to see that a Freiman homomorphic image of order 2 of a subgroup is a translated subgroup (Exercise 1), which is thus a refinement of a slightly larger subgroup.

Proof Without loss of generality we may assume that $0 \in A$ (since we can just add it in if it isn't already there). Then A is a refinement of $A-A$, and so it suffices to show that the set $G_0 := A - A$ is a refinement of some subgroup. If we let $\langle G_0 \rangle$ be the group generated by G_0, it thus suffices to show that $|\langle G_0 \rangle| \lesssim |G_0|$.

As mentioned in the previous section, we already know that

$$G_0 + G_0 \subseteq X + G_0$$

for some small set X. Iterating this we see that

$$G_0 + G_0 + G_0 \subseteq X + X + G_0$$

$$G_0 + G_0 + G_0 + G_0 \subseteq X + X + X + G_0$$

etc. Doing this r times and taking unions, we thus obtain

$$\langle G_0 \rangle \subset \langle X \rangle + G_0,$$

and hence

$$|\langle G_0 \rangle| \leq |\langle X \rangle||G_0| \leq r^{|X|}|G_0| \lesssim |G_0|$$

as desired. ∎

Doing this argument a little more carefully and explicitly one can get a more precise statement: if $|A+A| \leq K|A|$ and Z has torsion r, then A is contained in a subgroup G of cardinality at most $K^2 r^{K^4}|A|$. Thus this theorem is only effective when r and K are fairly small, otherwise the passage from X to $\langle X \rangle$, in particular, is rather lossy.

3. The torsion-free case

We now turn to the torsion-free case (when $nx \neq 0$ for all $x \in Z \backslash \{0\}$ and $n = 1, 2, 3, \ldots$), which is substantially more difficult. As far as I know, nobody has managed to adapt the simple argument in the previous section to this case (of course, one would have to replace "subgroup" by another notion, e.g. "generalized arithmetic progression", since in the torsion-free case all non-trivial subgroups are infinite). It is still possible to proceed, but we must now supplement our combinatorial techniques with a powerful new weapon - the Fourier transform.

Before we do so, though, let us try to figure out what "subgroup" should be replaced with in the torsion-free case. To put it another way, what explicit finite sets P are there which are manifestly essentially closed under addition?

A simple example is the interval $[0,N] := \{n \in \mathbf{Z} : 0 \leq n \leq N\}$ in the torsion-free group $\mathbf{Z}$. Since $|[0,N]| = N+1$ and $[0,N]+[0,N] = [0,2N]$, we see that $[0,N]$ is essentially closed under addition. More generally, (and keeping in mind that our sets P should be stable under Cartesian products), if we have any multi-index $N := (N_1, \ldots, N_d)$, we can define the box

$$[0,N] := \{(n_1, \ldots, n_d) \in \mathbf{Z}^d : 0 \leq n_i \leq N_i \text{ for } i = 1, \ldots, d\}.$$

Since $|[0,N]| = \prod_{i=1}^d (N_i+1)$ and $[0,N]+[0,N] = [0,2N]$, we see that

$$|[0,N]+[0,N]| \leq 2^d |[0,N]|$$

and so $[0,N]$ is essentially closed under addition if d is bounded. We call d the *dimension* or *rank* of the box $[0,N]$, and call the number $|[0,N]| = \prod_{i=1}^d (N_i+1)$ the *volume* of the box.

We also need to be stable under bijective Freiman homomorphisms of order 2. To this end, we make the following definitions. If $\phi : \mathbf{Z}^d \to Z$ is an affine homomorphism (i.e. the translation of a genuine homomorphism) and $[0,N]$ is a box of dimension d, we call $\phi([0,N])$ an *generalized arithmetic progression* of dimension d, length N and volume $|[0,N]|$. If ϕ is in fact injective on $[0,N]$, then we call this generalized arithmetic progression *proper*. It is easy to see that a generalized arithmetic progression takes the form

$$P := \{a + \sum_{i=1}^d n_i v_i : 0 \leq n_i \leq N_i \text{ for } i = 1, \ldots, d\} = \{a + n \cdot v : n \in [0,N]\}$$

for some fixed group elements $a, v_1, \ldots, v_d \in Z$, where $v := (v_1, \ldots, v_d) \in Z^d$, and the dot product $n \cdot v$ is defined in the obvious manner. Also, it is clear that this progression is proper if all the sums $a + \sum_{i=1}^d n_i v_i$. When $d = 1$ this is clearly just the familiar notion of arithmetic progression. We call $a = \phi(0)$ the *base point* of the progression P, and we call $v_1, \ldots, v_d$ the *basis vectors* of P.

If ϕ is injective on $[0,N]$, then it is definitely a bijective Freiman homomorphism of order 2 from $[0,N]$ to $\phi([0,N])$, and hence all proper generalized arithmetic progressions are essentially closed under addition if the dimension is bounded. (Conversely, these are the only images of $[0,N]$ under bijective Freiman homomorphisms; see Exercise 2). It is not as obvious, but in fact even the improper generalized arithmetic progressions are essentially closed under addition; we will see this in the next section.

We can now state Freiman's theorem, which asserts that up to refined small convolution, proper generalized arithmetic progressions are the only sets essentially closed under addition.

Theorem 3.1 (Freiman's theorem)**.** [2] *Let Z be a torsion-free abelian group, and let $A \subseteq Z$ be essentially closed under addition. Then A is a refined small convolution of a proper generalized arithmetic progression P of bounded rank.*

It turns out we can drop the "small convolution" bit of this theorem; more on this later.

We now sketch how this theorem is proven. Firstly, we observe using Rusza's quotient lemma that it suffices to show that some arithmetic combination of A, such as $2A-2A$, can be refined to a proper generalized arithmetic progression. Next, we use some Freiman isomorphisms to set Z to be a cyclic group, of order not much larger than $|A|$. The task is now to find a large arithmetic progression in $2A-2A$. It is here that the Fourier transform comes in handy; $2A-2A$ is the Fourier support of the non-negative function $|\hat{\chi_A}|^4$. The fact that A is essentially closed under addition and that A is a large subset of Z will mean that $|\hat{\chi_A}|^4$ is concentrated on a small set, which will imply that $2A-2A$ contains the "dual" of that set. This dual is something known as a *Bohr neighbourhood*. The last step is to show that Bohr neighbourhoods contain large proper generalized arithmetic progressions.

4. Reduction to finding a large arithmetic progression inside $2A-2A$

Let A be essentially closed under addition. Our task is to place A inside (a small convolution) of a proper generalized arithmetic progression P of comparable size. Clearly we may assume that A is large, since the task is trivial if $A = O(1)$.

It may seem strange, but up to small convolutions, the task of placing A inside P is almost the same as placing P inside A, or in some enlargement of A such as $2A-2A$. This is analogous to the statement in group theory that if a group contains a group H of small index. then G is contained in a small extension of H. To extend this to our situation of sets essentially closed under addition, we recall (a consequence of) Rusza's quotient lemma from last week's notes, which we rephrase here in our new language:

Lemma 4.1 (Rusza's quotient lemma)**.** *If $|A+B| \sim |A| \sim |B|$, then A is a refined small convolution of $B-B$.*

Thus if we can find a proper generalized arithmetic progression P of bounded rank such that $|A+P| \sim |A| \sim |P|$, Rusza's quotient lemma will then give that A is a refined small convolution of $P-P$. Since P is a generalized arithmetic progression of bounded rank, $P-P$ is a refined small convolution of P, and the claim then follows. (To see what is going on here, pretend that the progression P is actually a subgroup. Then the condition $|A+P| \sim |P|$ asserts that the quotient set A/P is small, and thus that A is contained in a small convolution of P).

So, we need to find a large proper progression P which is essentially A-invariant. But from sumset estimates we already know that all additive combinations of A, such as $2A-2A$, are already essentially A-invariant. Thus if we can find a proper generalized arithmetic progression P which refines $2A-2A$ and is of bounded rank, then we are done.

It remains to show that $2A - 2A$ contains a large subset (of cardinality $\sim |A|$) which is also a proper generalized arithmetic progression of bounded rank. To do this we first use Freiman homomorphisms to restrict A to an abelian group Z of comparable cardinality.

5. Restricting A to a cyclic group

Right now A is a subset of an arbitrary torsion-free abelian group Z. This group, being infinite is far too large to do any accurate Fourier analysis on, given how finite A is. (The uncertainty principle tells us that Fourier analysis is more useful when applied to large objects rather than small ones). On the other hand, we don't really use all of Z - in fact, we are only looking for progressions inside $2A - 2A$. This set is reasonably small (it has size $\sim |A|$) and behaves sort of like a group, so it seems reasonable that we can somehow quotient out the rest of Z and just work in a group of order comparable to A. This we can do, thanks to the technology of Freiman homomorphisms.

We will need Freiman isomorphisms of order 8 on A. These will induce Freiman isomorphisms of order 2 on $2A-2A$, which is enough to preserve proper generalized arithmetic progressions (Exercise 2).

First note that since Z is torsion free, A can be mapped onto the integers via a Freiman isomorphism of order 8 (Lemma 2 of last week's notes). So we can assume that Z is the integers $\mathbf{Z}$. This is better than Z being arbitrary, however the infinite nature of $\mathbf{Z}$ will still cause problems (because we will soon want to talk about a "random dilation" of A, and this only makes sense for finite groups). So now we truncate the integers to be finite.

By translating A we may assume that A lives on the positive integers, and in particular lies inside the interval $[1, p/8]$ for some very large prime $p \gg |A|$. Since $[1, p/8]$ can be mapped via a Freiman isomorphism of order 8 into the field $\mathbf{Z}/p\mathbf{Z}$, we can thus assume now that A lives inside the field $\mathbf{Z}/p\mathbf{Z}$. (This is no longer torsion free, but we will not need the torsion-free assumption any more).

We have now localized A to live in a finite group, but this is not particularly useful right now because p could be so much larger than $|A|$ that A only will occupy a tiny fraction of $\mathbf{Z}/p\mathbf{Z}$, which suggests via the uncertainty principle that Fourier analysis at this stage would not be helpful. To fix this we will perform a random projection of A onto a much smaller group $\mathbf{Z}/N\mathbf{Z}$, for some N which is comparable in size to A (we will choose N later).

But first, we use a random dilation trick, to give us some freedom to avoid some obstacles later on. Let λ be a randomly chosen invertible element of the field $\mathbf{Z}/p\mathbf{Z}$; this induces an additive isomorphism $x \to \lambda x$ on $\mathbf{Z}/p\mathbf{Z}$. In particular, we observe that the set λA is Freiman isomorphic to A of order 8 (in fact, of any order). So we have the freedom to randomly dilate A. (Of course, A would then no longer live in $[0, p/8]$, but we soon use a pigeonholing trick to compensate for this fact).

Now let $N \ll p$ be chosen later, and define the projection $\pi : \mathbf{Z}/p\mathbf{Z} \to \mathbf{Z}/N\mathbf{Z}$ by setting

$$\pi(n) := (n \mod N) \text{ for all } N = 0, 1, 2, \ldots, p-1.$$

This is not quite an additive homomorphism, however note that for $j = 0, 1, \ldots, 7$, π is a Freiman homomorphism of order 8 when restricted to the set $Z_j := (jp/8, (j+1)p/8]$, which is a set which occupies roughly $\frac{1}{8}$ of the original field $\mathbf{Z}/p\mathbf{Z}$. By the pigeonhole principle, for each λ there exists a $0 \leq j < 8$ such that the set $A' := \lambda A \cap Z_j$ refines λA. Thus if we set $B := \pi(A') \subseteq \mathbf{Z}/N\mathbf{Z}$, then the map $\pi : A' \to B$ is a surjective Freiman homomorphism of order 8.

This is beginning to look pretty good, but there is a problem: this map might not be injective (so that there might be collisions $\pi(x) = \phi(x')$ for two distinct $x, x' \in A'$). More generally, there may be collisions in $8A'$, so that

$$\pi(x_1) + \ldots + \pi(x_8) = \pi(x'_1) + \ldots + \pi(x'_8)$$

while $x_1 + \ldots + x_8 \neq x'_1 + \ldots + x'_8$, where all sixteen elements are in A'; this is the only obstruction to π being a Freiman *isomorphism* of order 8. Fortunately, this type of collision rarely occurs, if N is large enough and λ is chosen randomly. Indeed, if we do have the above collision, then the integer $(x_1 + \ldots + x_8) - (x'_1 + \ldots + x'_8)$ must be a non-zero multiple of N. Thus we will have no collisions if the set $8A' - 8A' \subseteq \lambda(8A - 8A)$ contains no multiples of N, where we think of elements of $\mathbf{Z}/p\mathbf{Z}$ as being integers from 0 to $p-1$. However, for each non-zero element y in $8A - 8A$, the probability that the randomly dilated element λy will be a multiple of N is $O(1/N)$, since p is prime. Thus, if we choose $N \gg |8A - 8A|$, there will be a non-zero probability that there are no collisions. Since $|8A - 8A| \sim |A|$, we can do this by choosing N to be a sufficiently large but bounded multiple of $|A|$.

To summarize, if we choose N large but still comparable to A we can find a refinement A' of λA which is Freiman-isomorphic of order 8 to a subset B of $\mathbf{Z}/N\mathbf{Z}$. If we can show that $2B - 2B$ contains a proper bounded-rank generalized arithmetic progression of size comparable to $|B|$, this implies (by pulling back the Freiman isomorphisms) that $2A' - 2A'$, and hence $2A - 2A$, similarly contain a proper bounded-rank generalized arithmetic progression of size comparable to $|B| = |A'| \sim |A|$, and we will be done.

Thus to conclude the proof of Freiman's theorem, we have to prove

Proposition 5.1. *Let N be a large integer, and let A be a subset of $\mathbf{Z}/N\mathbf{Z}$ with cardinality $|A| \sim N$ and such that $|A+A| \sim |A|$. (The latter condition is superfluous from the former, but there is a reason why we still keep it around, which we will get to later). Then $2A - 2A$ contains a proper bounded-rank generalized arithmetic progression P of cardinality $|P| \sim N$.*

Now we have achieved the goal of making A a large subset of the group, which allows for Fourier analysis to enter the picture. This we will do in the next section.

(A remark: The observant reader may ask why we aren't done already proving Freiman's theorem, because $\mathbf{Z}/N\mathbf{Z}$ is certainly a proper arithmetic progression, and by our assumptions it isn't much larger than A itself. The problem is that the

projection π is not an isomorphism on all of $\mathbf{Z}/N\mathbf{Z}$; at best it is an isomorphism on smaller sets such as $8A$ or $2A-2A$, and you can't pull the progression $\mathbf{Z}/N\mathbf{Z}$ to get any sort of meaningful progression in the original group. So Freiman isomorphisms are somewhat subtle maps - once you stray too far from the original set A (where "too far" depends on the order of the isomorphism), the mapping properties become terrible.)

Note, by the way, that since the magnitude of N is determined up to a constant, one can easily arrange matters so that N is prime (e.g. by using the prime number theorem; much more elementary proofs are available that there is a prime between N and CN for large enough N and some large absolute constant C).

6. Fourier analysis on finite abelian groups

Let Z be an additive finite abelian group. Then there exists a map $e : Z \times Z \to S^1$, where $S^1 := \{z \in \mathbf{C} : |z| = 1\}$ is the unit circle, which is multiplicative in the sense that

$$e(x+x',\xi) = e(x,\xi)e(x',\xi); \quad e(x,\xi+\xi') = e(x,\xi)e(x,\xi') \text{ for all } x,x',\xi,\xi' \in Z$$

and which is non-degenerate in the sense that for every non-zero x there exists a ξ such that $e(x,\xi) \neq 1$, and similarly for every non-zero ξ there exists a x such that $e(x,\xi) \neq 1$. We call e a *bi-character* of Z. Note in particular that $e(0,\xi) = e(x,0) = 1$, and $e(x,-\xi) = e(-x,\xi) = \overline{e(x,\xi)}$ for all $x,\xi \in \mathbf{Z}$. We usually refer to x as the *position variable* and ξ as the *frequency variable*.

The existence of such a bi-character is easiest to see in the cyclic case $Z = \mathbf{Z}/N\mathbf{Z}$, in which case one can just take

$$e(x,\xi) := \exp(2\pi i x\xi/N).$$

Also, if Z is the direct sum $Z = Z_1 \times Z_2$ of two smaller abelian groups, with bi-characters e_1 and e_2 respectively, then the tensor product

$$e_1 \otimes e_2((x_1,x_2),(\xi_1,\xi_2)) := e_1(x_1,\xi_1)e_2(x_2,\xi_2)$$

is a bi-character on Z. Since every finite abelian group is the direct sum of cyclic groups, we thus see that every abelian group has at least one such bi-character. (There are actually many bi-characters, basically because Z contains lots of automorphisms, but it never makes much of a difference in Fourier analysis which one you pick. One can make things more canonical by letting ξ range not in Z, but rather in the dual group Z^*, which consists of the characters on Z., in which case $e(x,\xi)$ is just the tautological map of ξ applied to x.)

Once one selects and fixes a bi-character e, one can then do Fourier analysis. Let dx denote normalized counting measure on Z, thus

$$\int_Z f(x)\ dx := \frac{1}{|Z|}\sum_{x\in Z} f(x),$$

and let $d\xi$ denote discrete measure on Z, thus

$$\int_Z g(\xi)\ d\xi := \sum_{\xi \in Z} g(\xi).$$

The functions $e(\cdot,\xi)$ then form an orthonormal basis of $L^2(Z,dx)$ (Exercise 3). Because of this, if we define define the *Fourier transform* $\hat{f}(\xi)$ of a complex-valued function $f(x)$ on Z by the formula

$$\hat{f}(\xi) := \langle f, e(\cdot,\xi)\rangle_{L^2(\mathbf{Z},dx)} := \int_Z f(x)\overline{e(x,\xi)}\ dx$$

then we have the *Fourier inversion formula*

$$f(x) = \int_Z \hat{f}(\xi)e(x,\xi)\ d\xi = \sum_{\xi \in Z} \hat{f}(\xi)e(x,\xi)$$

(Exercise 3). Furthermore, we have the *Parseval relation*

$$\langle f, g\rangle_{L^2(Z,dx)} = \langle \hat{f}, \hat{g}\rangle_{L^2(Z,d\xi)}$$

and hence the *Plancherel formula*

$$\|f\|_{L^2(Z,dx)} = \|\hat{f}\|_{L^2(Z,d\xi)}$$

(Exercise 3). Furthermore, if we define convolution

$$f * g(x) := \int_Z f(y)g(x-y)\ dy$$

then we have

$$\widehat{f * g}(\xi) = \hat{f}(\xi)\hat{g}(\xi)$$

(Exercise 3), and if we define reflection $\tilde{f}$ by

$$\tilde{f}(x) := \overline{f(-x)}$$

then we have

$$\hat{\tilde{f}}(\xi) = \overline{\hat{f}}(\xi)$$

(Exercise 3).

We now apply these Fourier identities to the characteristic function $f(x) = \chi_A(x)$ of some subset $A \subseteq Z$ of cardinality $|A| = c|Z|$ for some $0 < c \leq 1$; we shall think of c as being reasonably close to 1. Then from Plancherel we have

$$\sum_{\xi \in Z} |\hat{\chi}_A(\xi)|^2 = \int_Z |\chi_A(x)|\ dx = c \tag{1}$$

while from the triangle inequality we have the crude pointwise estimate

$$|\hat{f}(\xi)| \leq \int_Z |\chi_A(x)|\ dx \leq c. \tag{2}$$

Together, these bounds imply that the number of "large" Fourier coefficients is rather small: for any $0 < \varepsilon \leq 1$, we have

$$|\{\xi \in Z : |\hat{\chi}_A(\xi)| \geq \varepsilon c\}| \leq \varepsilon^{-2}c^{-1}. \tag{3}$$

This is a basic consequence of orthogonality; the only way that $|\hat{\chi}_A(\xi)| = |\langle \chi_A, e(\cdot,\xi)\rangle|$ can come close to the theoretical maximum of c is if χ_A is close to parallel to the character $e(\cdot,\xi)$, but because the characters are all orthogonal, it is difficult for χ_A

to be close to parallel to too many of these characters at once. Later on we will obtain a more refined estimate for these "resonant frequencies" ξ where the Fourier transform is large.

Now suppose that A is essentially closed under addition, so that $|A+A| \leq K|A|$ for some $K = O(1)$. Then $\chi_A * \chi_A$ is supported on $A+A$ which has normalized measure at most cK. On the other hand, we have

$$\|\chi_A * \chi_A\|_{L^1(Z,dx)} = \|\chi_A\|_{L^1(Z,dx)}\|\chi_A\|_{L^1(Z,dx)} = c^2,$$

and hence by Hölder's inequality (or Cauchy-Schwarz)

$$\|\chi_A * \chi_A\|_{L^2(Z,dx)} \geq c^{3/2}K^{-1/2}.$$

By Plancherel this implies that

$$\sum_{\xi \in Z} |\hat{\chi}_A(\xi)|^4 \geq c^3 K^{-1}. \tag{4}$$

In contrast, notice from (1), (2) that

$$\sum_{\xi \in Z} |\hat{\chi}_A(\xi)|^4 \leq c^2 \sum_{\xi \in Z} |\hat{\chi}_A(\xi)|^2 = c^3 \tag{5}$$

so we have a very precise control on the l^4 norm of $\hat{\chi}_A$ here, especially if K is small. In fact, we can even get a concentration estimate on $\hat{\chi}_A$. A slight modification of (5) gives

$$\sum_{\xi \in Z : |\hat{\chi}_A(\xi)| \leq \varepsilon c} |\hat{\chi}_A(\xi)|^4 \leq \varepsilon^2 c^2 \sum_{\xi \in Z} |\hat{\chi}_A(\xi)|^2 = \varepsilon^2 c^3,$$

so if we choose $\varepsilon := \frac{1}{2\sqrt{K}}$, then by (4) we have

$$\sum_{\xi \in \Lambda} |\hat{\chi}_A(\xi)|^4 \geq \frac{3}{4} \sum_{\xi \in Z} |\hat{\chi}_A(\xi)|^4 \tag{6}$$

where Λ is the set of *resonant frequencies*

$$\Lambda := \{\xi \in Z : |\hat{\chi}_A(\xi)| \geq \frac{c}{2\sqrt{K}}\}.$$

On the other hand, from (3) we have

$$|\Lambda| \leq 4Kc^{-1}. \tag{7}$$

Thus, if A is essentially closed under addition, *a small set Λ of frequencies will dominate the L^4 norm of $\hat{\chi}_A$*. (This should be contrasted with the case when A is exactly closed under addition: see Exercise 4).

To access this structural information about $|\hat{\chi}_A(\xi)|^4$, we introduce the function $f(x)$ defined by

$$f := \chi_A * \chi_A * \tilde{\chi}_A * \tilde{\chi}_A.$$

Observe from our Fourier identities that f is supported on $2A-2A$ and has Fourier transform

$$\hat{f}(\xi) = |\hat{\chi}_A(\xi)|^4.$$

In particular, f has non-negative Fourier transform which is in some sense highly concentrated in Λ. Remember that we wanted to show that $2A-2A$ contained a large arithmetic progression; it thus suffices to show that the support of f contains

a large proper arithmetic progression. By the Fourier inversion formula, it thus suffices to find a large proper bounded-rank generalized arithmetic progression P for which the Fourier series

$$f(x) = \sum_{\xi \in Z} |\hat{\chi}_A(\xi)|^4 e(x,\xi) \tag{8}$$

is non-zero for all $x \in P$.

Now we use the fact that $|\hat{\chi}_A|^4$ is concentrated in Λ; by the uncertainty principle, this should mean that f should be concentrated on some sort of dual set of Λ. To make this precise, let $X \subseteq Z$ denote the set

$$X := \{x \in Z : |e(x,\xi) - 1| < 1/4 \text{ for all } x \in \Lambda\}; \tag{9}$$

this is an example of a *Bohr set* and we will have more to say about this set later. Then it is clear that

$$\mathrm{Re} \sum_{\xi \in \Lambda} |\hat{\chi}_A(\xi)|^4 e(x,\xi) \geq \frac{3}{4} \sum_{\xi \in Z} |\hat{\chi}_A(\xi)|^4,$$

so in particular

$$|\sum_{\xi \in \Lambda} |\hat{\chi}_A(\xi)|^4 e(x,\xi)| \geq \frac{3}{4} \sum_{\xi \in Z} |\hat{\chi}_A(\xi)|^4.$$

On the other hand, from (6) and the triangle inequality we have

$$|\sum_{\xi \in Z \backslash \Lambda} |\hat{\chi}_A(\xi)|^4 e(x,\xi)| \leq \frac{1}{3} \sum_{\xi \in \Lambda} |\hat{\chi}_A(\xi)|^4.$$

Adding this together, we see thus see that the expression (8) is non-zero for all $x \in X$. In other words, *the Bohr set X is contained in $2A-2A$.* (It is not yet clear how useful this is, because we haven't shown any lower bounds on the size of X, but we will come to this shortly).

Let us now apply the above discussion to the situation inherited from the previous section, in which we have a refinement A of Z_N such that $|A+A| \sim |A|$; thus both c and K are comparable to one, and so by (7) Λ is a small set. We have shown that $2A-2A$ contains a Bohr set X, which is sort of a dual to Λ. Thus all we need to do is show that the Bohr set contains a proper bounded-rank generalized arithmetic progression P of size comparable to N, and we will have proven Freiman's theorem. This is the purpose of the next section.

7. Some lattice theory

Let us remind ourselves of the current situation. We have a large integer N, and a small set Λ in $\mathbf{Z}/N\mathbf{Z}$, so that $|\Lambda| = O(1)$. We wish to show that the Bohr set (9) contains a proper bounded-rank generalized arithmetic progression of volume $\sim N$.

We haven't fixed the bicharacter e yet; let us use the standard one $e(x,\xi) := \exp(2\pi i x\xi/N)$. Then our Bohr set takes the form

$$X = \{x \in Z : \|\frac{x\xi}{N}\| < \delta N \text{ for all } x \in \Lambda\}$$

for some small absolute constant δ (corresponding to the 1/4 factor in (9); one can shrink X a bit and take δ to be 1/20 for concreteness). Here $\|x\|$ denotes the distance from the real number x to the nearest integer.

In contrast, note that the set

$$\{x \in \mathbf{R} : \|\frac{x\xi}{N}\| = 0\}$$

is an infinite arithmetic progression for each $x \in \Lambda$. So X is in some sense the intersection of $|\Lambda| = O(1)$ many neighbourhoods of arithmetic progressions, and we have to somehow show that X thus itself contains a large generalized arithmetic progression of bounded rank. To do this we use some lattice theory.

First note that we can erase 0 from Λ since it contributes nothing to the Bohr set. Let us now enumerate Λ as $\Lambda = \{\xi_1, \ldots, \xi_k\}$ for some $k = O(1)$. We work on the torus $T^k := \mathbf{R}^k/\mathbf{Z}^k$, and isolate the vector

$$w := (\frac{\xi_1}{N}, \ldots, \frac{\xi_k}{N}) + \mathbf{Z}^k$$

in this torus. Since $Nw = 0$, the expression $xw \in T^k$ makes sense for every $x \in \mathbf{Z}/N\mathbf{Z}$. Observe that the set

$$X' := \{x \in \mathbf{Z}/N\mathbf{Z} : xw \in B(0,\delta)\},$$

where $B(0,\delta)$ is the ball of radius δ around the origin in the torus T^k, is contained in the Bohr set X. Thus it will suffice to show that X' contains a proper bounded-rank generalized arithmetic progression P of cardinality $\sim N$.

As a warm up let us first show the much weaker statement that $|X'| \sim N$ (so that we at least have enough elements to support such a large arithmetic progression). We can cover the torus T^k by about $O((C/\delta)^k)$ balls of radius $\delta/2$. By the pigeonhole principle, one of these balls, say $B(x_0, \delta/2)$, contains at least $\gtrsim (\delta/C)^k N$ multiples of w. Subtracting, this means that $B(x_0,\delta/2) - B(x_0,\delta/2) = B(0,\delta)$ contains $\gtrsim (\delta/C)^k N$ multiples of w. Since δ is an absolute constant and $k = O(1)$, we see that $|X'| \gtrsim N$ and we are done.

We still have to construct P. The idea is to use basis vectors $v_i \in \mathbf{Z}/N\mathbf{Z}$ such that $v_i w$ is very close to the origin; this will allow us to use many multiples of v_i in our progression and still have the property that $xw \in B(0,\delta)$ for all $x \in P$. Of course we need the progression to be proper, which we will achieve by ensuring that the $v_i w$ are "linearly independent".

We turn to the details. We will work locally in the ball $B(0,\delta)$, and observe that the torus $\mathbf{T}^k$ is indistinguishable from Euclidean space $\mathbf{R}^k$ in this region. This allows us to use tools from linear algebra.

We use the greedy algorithm. Let $X'w := \{xw : x \in X'\}$; these are the multiples of w which lie in $B(0,\delta)$, and we have already seen that this set is quite large, $|X'w| = |X'| \sim N$. Let $v_1 w$ be the non-zero element of $X'w$ which is closest to the origin (if there is a tie, choose arbitrarily). Let $v_2 w$ be the non-zero element of $X'w$ which is linearly independent from $v_1 w$ and is closest to the origin. Let $v_3 w$

be the non-zero element of $X'w$ which is linearly independent from v_1w and v_2w and which is closest to the origin. We continue in this manner until we can go no further; we thus obtain linearly independent vectors $v_1w, \ldots, v_{k'}w$ in $X'w$ for some $1 \leq k' \leq k$ which span $X'w$. For $j = 1, \ldots, k$ let $r_j = |v_jw|$ denote the distance from v_jw to the origin, thus $0 \leq r_1 \leq r_2 \leq \ldots \leq r_{k'} \leq \delta$.

We now use a volume packing argument to get upper bounds on the distances r_j. For $j = 0, \ldots, k'$, let V_j denote the span of $v_1w, \ldots, v_jw$, thus $\{0\} = V_0 \subset V_1 \subset V_2 \subset \ldots \subset V_{k'}$ form a k'-dimensional complete flag of subspaces. Let π_j denote the orthogonal projection onto V_j, and consider the "box"

$$B = \{x \in V_{k'} : |\pi_{j+1}(x) - \pi_j(x)| < r_j/2k \text{ for all } 0 \leq j < k'\}.$$

Note from the triangle inequality that this box is contained in $B(0, \delta/2)$. Observe that B cannot contain any multiples of w other than the origin. For, if we had some non-zero multiple vw of w in B, let V_j be the first subspace in the flag which contains v_j, thus vw is independent of $v_1, \ldots, v_{j-1}$. By the triangle inequality, we would then have

$$|vw| = |\pi_j(vw)| < r_1/2k + \ldots + r_j/2k \leq r_j,$$

contradicting the minimality of r_j.

If we now let $B/2$ denote the set $\{x/2 : x \in B\}$, then by convexity and symmetry we have $B/2 - B/2 = B$. Thus by the previous discussion, the sets $xw + B/2$ are disjoint as x varies over X'. Since the sets $xw + B/2$ lie in the k'-dimensional disk $B(0, 2\delta) \cap V_{k'}$, and $|X'| \sim N$, a volume packing argument thus yields that the k'-dimensional volume of $B/2$ is $O(1/N)$ (allowing constants to depend in δ and k), so the same applies to B. But a simple computation shows that the k'-dimensional measure of B is comparable to $r_1 \ldots r_{k'}$, so we have shown that

$$r_1 \ldots r_{k'} \lesssim 1/N. \tag{10}$$

Now let $N_1, \ldots, N_{k'} \geq 1$ be positive integers such that N_j is the integer part of $1 + \frac{\delta}{kr_j}$. By the previous (and the assumption $k = O(1)$) we have

$$N_1 \ldots N_{k'} \gtrsim N.$$

If we let $P \subseteq \mathbf{Z}/N\mathbf{Z}$ be the generalized arithmetic progression

$$P := \{\sum_{j=1}^{k'} n_j v_j : 0 < n_j \leq N_j \text{ for all } j = 1, \ldots, k'\}$$

then P has volume $\gtrsim N$, and by construction we also see that $Pw \subseteq B(0, \delta)$. Thus P is contained inside X' and is thus contained inside the Bohr set. Also from linear independence we see that P is proper (note that the linear combinations never leave $B(0, \delta)$, so we don't encounter any wraparound issues arising from the torus, and we are done. This completes the proof of Freiman's theorem.

(Remark: one can get more precise bounds on $r_1 \ldots r_k$, by a classical theorem known as *Minkowski's second theorem*, but we do not need to do so here.)

8. Some more lattice theory: Proper and improper generalized arithmetic progressions

The main result of this section is the following result of Gowers and Walters:

Theorem 8.1. *Let Z be a torsion-free abelian group. Then every generalized arithmetic progression in Z with bounded rank is the refinement of a proper generalized arithmetic progression with equal or lesser rank.*

Now observe that any small convolution $X + P$ of a generalized arithmetic progression P of bounded rank d is a refinement of a generalized arithmetic progression of rank $d+|X|$; indeed, X is clearly contained in a rather small arithmetic progression Q of rank $|X|$ and volume $2^{|X|}$, and then $X + P$ is a refinement of the generalized arithmetic progression $Q + P$. Thus by Lemma 8.1, the small convolution of a proper bounded-rank progression is the refinement of another proper bounded-rank progression. This allows us to remove the "small convolution" caveat in Freiman's theorem.

The proof I give here is not Gowers and Walters' original proof (which I haven't seen), and is probably much clumsier, but it does illustrate some interesting techniques from lattice theory.

Before we prove this theorem, let us recall an algebraic analogue which has a similar flavor.

Lemma 8.2. *Let Γ be a lattice in $\mathbf{R}^d$, i.e. a discrete additive subgroup. Then there is an invertible linear transformation which transforms Γ to the standard lattice $\mathbf{Z}^k$ for some $0 \leq k \leq d$. We call k the rank of the lattice Γ.*

If one thinks of finitely generated lattices as the analogue of generalized progressions, and invertible images of the standard lattice as the analogue of proper generalized progressions, then we see that this lemma is in some sense an algebraic analogue of Theorem 8.1.

Proof This result can be proven algebraically (using the fact that every finitely generated torsion-free group is isomorphic to some standard lattice), but we will use an analytic proof, similar in spirit to the Euclidean algorithm (which in fact corresponds to the $d = 1$ case of this Lemma).

We first observe that we may assume that the vectors in Γ span $\mathbf{R}^d$, else we could pass from $\mathbf{R}^d$ to a smaller vector space and continue the argument. With this assumption, we will show that Γ is in fact isomorphic to $\mathbf{Z}^d$.

Since Γ is discrete, there exists some ball $B(0, \varepsilon)$ which contains no lattice point in Γ other than the origin. We fix this $\varepsilon > 0$.

By the spanning assumption, we can find d linearly independent vectors $v_1, \ldots, v_d$ in Γ, so in particular the volume $|v_1 \wedge \ldots \wedge v_d|$ is strictly positive. On the other

hand, by Minkowski's theorem (Exercise 5) we know that this volume is bounded from below by some constant $c = c(\varepsilon) > 0$ depending only on ε.

We now use the method of descent. Let $\langle v_1, \ldots, v_d \rangle$ be the lattice consisting of integer combinations of the $v_1, \ldots, v_d$. This is a sublattice of Γ. If it is in fact equal to Γ, then Γ is clearly isomorphic to $\mathbf{Z}^d$ since we can use an invertible transformation to map the linearly independent vectors $v_1, \ldots, v_d$ to the generators of $\mathbf{Z}^d$. Now suppose that the vectors $\langle v_1, \ldots, v_d \rangle$ do not generate Γ. Then the half-open paralleopiped $\{\sum_{i=1}^d t_i v_i : 0 \leq t_i < 1\}$ generated by the vectors $v_1, \ldots, v_d$, being a fundamental domain of $\langle v_1, \ldots, v_d \rangle$, must contain a lattice point w in Γ which is non-zero. Write $w = \sum_{i=1}^d t_i v_i$; without loss of generality we may assume that $t_d > 0$. We may also assume that $t_d \leq 1/2$ since we could replace w by $v_1 + \ldots + v_d - w$ otherwise. Then the volume $|v_1 \wedge \ldots \wedge v_{d-1} \wedge w|$ is at most half that of $|v_1 \wedge \ldots \wedge v_d|$. We thus replace v_d by w and repeat the above argument. Because of our absolute lower bound on the volume of paralleopipeds, this argument must eventually terminate, at which point we have found the desired presentation for Γ. ■

Call a lattice vector v in Γ *irreducible* if it is not of the form nw for some integer $n > 1$ and some $w \in \Gamma$. Clearly every lattice vector is the integer multiple of an irreducible lattice vector.

One corollary of the above Lemma is that irreducible vectors can always be factored out:

Corollary 8.3. *Let Γ be a lattice of rank k, and let w be a non-zero irreducible vector in Γ. Then we have a factorization $\Gamma = \langle w \rangle + \Gamma'$ where $\langle w \rangle$ is the space of integer multiples of w, and Γ' is a lattice of rank $k-1$.*

Proof By Lemma 8.2 we may take Γ equal to $\mathbf{Z}^k$, at which point w takes the form $(a_1, \ldots, a_k)$ for some integers $a_1, \ldots, a_k$. Since v is irreducible, the $a_1, \ldots, a_k$ have no common factor, and thus we have $n_1 a_1 + \ldots + n_k a_k = 1$ for some integers $n_1, \ldots, n_k$. By use of row operations (which leaves $\mathbf{Z}^k$ invariant) we may thus assume that w is equal to the basis vector e_k. But then the claim follows by setting $\Gamma' := \mathbf{Z}^{k-1}$. ■

We now prove Theorem 8.1. We shall actually prove the following variant, which is clearly equivalent to Theorem 8.1 but is easier to prove for inductive purposes:

Proposition 8.4. *Let Z be a torsion-free abelian group, let P be a generalized arithmetic progression in Z of rank $k = O(1)$, and let Q be a generalized arithmetic progression in Z of rank $d = O(1)$ and cardinality $|Q| = O(1)$. Then $Q + P$ is a refinement of a proper generalized arithmetic progression of bounded rank.*

Thus we have split our progression into a large progression P and a small progression Q. The reason we do this is that our inductive step will sometimes reduce the rank k of the large progression at the cost of increasing the rank d of the small progression,

and then use a different argument to reduce the small rank d while keeping the large rank k constant.

Proof As mentioned above, we will induct on both k and d. We call a pair P', Q' of progressions with ranks k', d' *simpler* than P, Q with ranks k, d if $k' < k$, or if $k' = k$ and $d' < d$, or if $k' = k$ and $d' = d$ and P' has a smaller volume than P. We may assume inductively that the Proposition has already been proven for all pairs P', Q' of progressions which are simpler than P, Q.

The basic idea is to use the fact that an arithmetic progression can be improper only when there is at least one collision between two different sums in the progression. One can then use this collision to reduce the rank of either P or Q without increasing the size of P and Q too much.

First let us look at the case where P is improper; we will aim to use this improperness to drop the rank k of P by one; the small progression Q will make its appearance at the end. We may assume that P has base point 0, i.e.

$$P := \{n \cdot v : n \in [0, N]\}$$

for some multivector $v \in Z^k$ and multinumber $N := (N_1, \ldots, N_k)$. It will be convenient to rescale by N. If $x \in \mathbf{R}^k$, let x/N denote the vector

$$x/N := (x_1/N_1, \ldots, x_k/N_k),$$

and let v' denote the multivector

$$v' := (N_1 v_1, \ldots, N_k v_k),$$

thus we have

$$P = \{x \cdot v' : x \in [0, N]/N\}.$$

Since P is improper, we have some $x, x' \in [0, N]/N$ such that $x \cdot v' = x' \cdot v'$. Let $\ker(v') \subseteq \mathbf{Z}^k/N$ be the lattice

$$\ker(v') := \{x \in \mathbf{Z}^k/N : x \cdot v' = 0\};$$

we thus see that $\ker(v')$ contains a non-zero vector $w \in [-N, N]/N$, namely $w := x - x'$. We may of course assume that w is irreducible in $\ker(v')$.

Morally, Corollary 8.3 should now let us replace the rank k lattice $\mathbf{Z}^k/N$ by one of rank $k-1$, but we are not dealing with the infinite lattice here, only a bounded subset of it, and so we have to do a little more work.

Write $w = (w_1, \ldots, w_k)$. Without loss of generality we may assume that the k^{th} coefficient dominates in the sense that

$$|w_j| \leq |w_k| \text{ for } j = 1, \ldots, k. \tag{11}$$

One should think of this as saying that w is essentially aligned with the e_k basis vector. In particular w_k is non-zero.

Now suppose that w is small, $|w| \ll 1$. Then we can replace P by a progression of the same rank but smaller volume, allowing us to use the induction hypothesis.

Here's how. Let N' denote the multinumber

$$N' := (kN_1, \ldots, kN_{k-1}, N_k|w_k|).$$

Observe that for every $x \in [0,N]/N$ there exists some integer n such that $x - nw$ lies in $[-N', N']/N$, just by the Euclidean algorithm in the k^{th} co-ordinate. Since $x \cdot v' = (x - nw) \cdot v'$, we can thus cover the arithmetic progression P by another progression P' of the same rank k, but of volume $O(|[-N', N']|) = O(|w_k||[0,N]|)$. This is smaller if $|w|$ is sufficiently small. Thus in this case one can use the induction hypothesis to close the argument. (Note that P' can easily be verified to be a small convolution of P, so that P is thus a refinement of P').

It remains to consider the case when w is never small; i.e. there is a ball $B^k(0,\varepsilon)$ in $\mathbf{R}^k$ which contains no lattice point of $\ker(v')$ other than the origin, where $\varepsilon = \varepsilon(k) > 0$ depends only on k. This tells us (by volume-packing arguments) that the lattice $\ker(v')$ is fairly sparse; for instance, it can only have a bounded number of elements in the bounded set $[-N,N]/N$. In particular this implies that

$$|P| \sim |[0,N]|.$$

Define a linear projection π_w from $\mathbf{R}^k$ to $\mathbf{R}^{k-1}$ by

$$\pi_w(x) = x - \frac{x_k}{w_k} w.$$

Let $\Gamma \subseteq \mathbf{R}^{k-1}$ denote the lattice $\Gamma := \pi_w(\mathbf{Z}^k/N)$; this lattice Γ is discrete (indeed, it lives in the discrete lattice $\frac{1}{|w_k|N_1 \ldots N_k}\mathbf{Z}^{k-1}$) and has rank $k-1$ (since it contains $\mathbf{Z}^{k-1}/N$). Now, let $B(0,C)$ be a large ball in $\mathbf{R}^{k-1}$. We claim that the set $\Gamma \cap B(0,C)$ has $O(|[0,N]|)$ elements and that

$$P \subseteq (\Gamma \cap B(0,C)) \cdot v'.$$

To verify the latter claim, let $x \cdot v'$ be any element of P. Since $w \cdot v' = 0$, we have $x \cdot v' = \pi_w(x) \cdot v'$. But since $x \in [0,N]/N$, a simple computation using (11) shows that $\pi_w(x) \in \Gamma \cap B(0,C)$ if C is large enough depending on k, and we are done. To prove the former claim, note that Γ consists of elements of the form $\pi_w(x)$ for some $x = (x_1, \ldots, x_k) \in \mathbf{Z}^k$. Since $\pi_w(w) = 0$, we may assume that $|x_k| < |w_k|$. But then if $\pi_w(x) \in B(0,C)$, one easily sees that $|x| \lesssim 1$. Since x also lies in $\mathbf{Z}^k/N$, the claim follows.

Since $\Gamma \cap B(0,C)$ has $O(|[0,N]|)$ elements, we see by volume packing arguments (noting that Γ contains $\mathbf{Z}^{k-1}$ and so we can start tiling all of $\mathbf{R}^{k-1}$ with integer translates of $B(0,C)$) that

$$|\mathbf{R}^{k-1}/\Gamma| \gtrsim 1/|[0,N]|. \tag{12}$$

The plan is to replace the progression P by the slightly larger object $(\Gamma \cap B(0,C)) \cdot v'$, which has rank $k-1$. The trouble is that we have not yet demonstrated that $\Gamma \cap B(0,C)$ is an arithmetic progression. We already know that Γ is a lattice of rank $k-1$, and hence by Lemma 8.2 can be generated by $k-1$ linearly independent lattice vectors. In principle, this allows us to cover $\Gamma \cap B(0,C)$ by an arithmetic progression of rank $k-1$, but if the parallelopiped generated by these vectors is too

degenerate then this progression will be very inefficient (it will require many more points than $|[0,N]|$, and so we will no longer be able to use the induction hypothesis since the old progression P will not be a refinement of the new one). So we need a way to obtain a non-degenerate parallelopiped in Γ. Fortunately we can re-use the algorithm from the previous section. Let u^1 be a non-zero vector in Γ of minimal length; let u^2 be a non-zero vector linearly independent of u^1 with minimal length subject to these constraints; and so forth up to u^{k-1} (we must get $k-1$ vectors this way because Γ spans $\mathbf{R}^{k-1}$. By repeating the volume packing argument used to prove (10) in the last section, we have that

$$|u^1| \dots |u^{k-1}| \sim |\mathbf{R}^{k-1}/\Gamma|. \tag{13}$$

In particular, this implies that the vectors $u^1, \dots, u^{k-1}$ are not very degenerate:

$$|\mathbf{R}^{k-1}/\Gamma| \leq |u^1 \wedge \dots \wedge u^{k-1}| \leq |u^1| \dots |u^{k-1}| \lesssim |\mathbf{R}^{k-1}/\Gamma|. \tag{14}$$

Also by the greedy nature of the construction, and the fact that Γ contains $\mathbf{Z}^{k-1}/N$, we see that the vectors u^i are all bounded.

Let U be the sublattice Γ generated by the linearly independent vectors $u^1, \dots, u^{k-1}$. By (14) we see that U has bounded index:

$$|\Gamma/U| \lesssim 1.$$

In particular Γ is a small convolution of U, and $\Gamma \cap B(0,C)$ is a small convolution of $U \cap B(0,2C)$.

Let x be any element of $U \cap B(0,2C)$. Since $x \in U$, we have

$$x = \sum_{i=1}^{k-1} n_i u^i$$

for some integers n_i. These integers can be computed by Cramer's rule:

$$n_i = \frac{u^1 \wedge \dots \wedge u^{i-1} \wedge x \wedge u^{i+1} \wedge \dots \wedge u^{k-1}}{u^1 \wedge \dots \wedge u^{k-1}}$$

where the quotient of two $(k-1)$-forms in $\mathbf{R}^{k-1}$ is defined in the obvious manner. By (14) we thus have $n_i = O(1/|u^i|)$. Thus $U \cap B(0,2C)$, and hence $(U \cap B(0,2C)) \cdot v'$, is contained in an arithmetic progression of rank $k-1$ and size

$$[\lfloor C/|w_1| \rfloor, \dots, \lfloor C/|w_{k-1}| \rfloor]$$

which thus has volume

$$\lesssim \frac{1}{|w_1| \dots |w_{k-1}|} \lesssim \frac{1}{|\mathbf{R}^{k-1}/\Gamma|} \lesssim |[0,N]|$$

by (14) and (12).

To summarize, we have covered the arithmetic progression P, which has rank k and cardinality $|P| \sim |[0,N]|$, by a small convolution of another arithmetic progression - call it P' - which has rank $k-1$ and volume $O(|[0,N]|)$. Since every small set can be contained in a small arithmetic progression, we thus see that P is a refinement of $Q' + P'$ for some small arithmetic progression Q'. We now apply the induction hypothesis, replacing P by P' (thus reducing the rank of P) and adding Q' to Q

(this increases the rank of Q, but we still have a simpler pair of progressions by definition) to get what we want. This concludes the induction when P is improper.

Now we look at the case when P is proper; the idea is now to work with Q in much the same way that P was worked on in the previous argument.

If $Q+P$ is already proper, then we are done. So suppose $Q+P$ is improper; we thus have $q, q' \in Q$ and $p, p' \in P$ such that $q+p = q'+p'$. Since P is already proper, we have $q \neq q'$.

By translation symmetry we may assume Q has base point 0, so

$$Q := \{a + m \cdot y : m \in [0, M]\}$$

for some multivector $y \in Z^d$, and some multinumber M of rank d; since Q is small, M is bounded. Since we have $q - q' \in P - P$, we thus have $m \cdot y \in P - P$ for some $m \in [-M, M]$. We can write $m = cm'$ where m' is irreducible in $\mathbf{Z}^d$, and $c = O(1)$ is non-zero. Then $m' \cdot v \in (P-P)/c$.

By Corollary 8.3, $\mathbf{Z}^d$ is the direct sum of a lattice of rank $d-1$, and the integer multiples $\langle m' \rangle$ of m'. Since M is small, we can thus cover Q by $Q' + [-C, C]m' \cdot v$ for some small arithmetic progression Q' of rank $d-1$, and some bounded constant C. Thus we have

$$Q + P \subseteq Q' + (\bigcup_{|r| \leq C} \frac{r}{c}(P-P) + P).$$

But one can easily verify that $(\bigcup_{|r|\leq C} \frac{r}{c}(P-P)+P)$ is contained in a progression P' of the same rank as P and of comparable volume. Thus $Q+P$ is a refinement of the simpler pair of progressions $Q' + P'$, and we can use the induction hypothesis again. This closes the induction. ∎

9. Disassociated sets, and Chang's refinement of Freiman's theorem

In previous sections we have proven Freiman's theorem, which states that if $|A+A| \leq K|A|$ then A is a refined small convolution of an arithmetic progression P of dimension d at most $d \leq C_K$, where the constants of the refinement and small convolution are at most C'_K. However, the estimates we have for the dimension d are pretty bad; our bound for C_K is essentially the same as $|\Lambda|$, which is as large as K^C, while the bound for C'_K is exponential in the dimension and thus looks like $\exp(CK^C)$.

Recently, Chang [1] has refined the dimension bound from K^C to $CK \log K$ (which is close to the optimal conjecture of CK), with a corresponding improvement of the C'_K constant to $\exp(CK(\log K)^2)$. The idea is to use more information about the resonant set Λ than just merely bounding the size $|\Lambda|$ as in (7).

Let us recall our situation. We are working in a cyclic group $\mathbf{Z}/N\mathbf{Z}$, and have a subset A of $\mathbf{Z}/N\mathbf{Z}$ of size $|A| = cN$, and we consider the resonant set

$$\Lambda := \{\xi \in \mathbf{Z}/N\mathbf{Z} : |\hat{\chi}_A(\xi)| \geq \alpha c\}$$

for some $0 < \alpha < 1$ (in the previous applications, we had $\alpha := 1/(2\sqrt{K})$). Then the Plancherel identity

$$\sum_{\xi \in \mathbf{Z}/N\mathbf{Z}} |\hat{\chi}_A|^2 = c$$

gives the cardinality bound

$$|\Lambda| \leq \alpha^{-2} c^{-1}.$$

This bound is fairly sharp (see for instance Q4). However, we can do a bit better.

Definition 9.1. A *cube* of dimension d is a generalized arithmetic progression of dimension d, base point equal to 0, and all lengths $N_1, \ldots, N_d$ equal to 1; in other words, a cube is a set of the form

$$\{\sum_{i=1}^{d} \epsilon_i v_i : \epsilon_i = 0, 1\}$$

for some basis vectors $v_1, \ldots, v_d$. Alternatively, a cube is any linear image of $[0,1]^d$.

Theorem 9.2. [1] *With the above notation, Λ is contained in a cube of dimension $O(\alpha^{-2} \log(1/c))$.*

This result is especially impressive in the case $\alpha \sim 1$ and $c \ll 1$. Then the set Λ has cardinality $O(1/c)$, but here we are saying a much stronger statement, that Λ is in fact contained in a cube of dimension $O(\log(1/c))$. (In the context of Q6, this implies that any subgroup of cardinality c is contained in a cube of dimension $O(\log(1/c))$, which is a true but not entirely obvious statement. To get some idea of what's going on, observe that the set $\{1, 2, \ldots, N\}$ is contained in the cube with basis vectors 2^0, 2^1, $\ldots$, $2^{\lfloor \log_2 N \rfloor}$).

To prove this theorem we need a definition.

Definition 9.3. A set $\xi_1, \xi_2, \ldots, \xi_n$ of frequencies is said to be *dissociated* if the cube

$$\{\sum_{i=1}^{d} \epsilon_i \xi_i : \epsilon_i = 0, 1\}$$

is proper. Equivalently, the only solution to

$$\sum_{i=1}^{d} \epsilon_i \xi_i = 0; \quad \epsilon_i \in \{-1, 0, 1\}$$

occurs when all the ϵ_i are equal to zero.

A key example to keep in mind here is when $\xi_j = 2^j$. The intuition then is that dissociated frequencies behave much like lacunary frequencies.

To prove Theorem 9.2, it will suffice to show

Proposition 9.4. *With the above notation, any dissociated set $\xi_1, \xi_2, \ldots, \xi_n$ in Λ can have cardinality at most $O(\alpha^{-2} \log(1/c))$.*

Let us see why Proposition 9.4 implies Theorem 9.2. Let $\xi_1, \ldots, \xi_n$ be a *maximal* dissociated set in Λ. Then any other frequency ξ in Λ must be of the form

$$\xi = \sum_{i=1}^n \epsilon_i \xi_i$$

for some $\epsilon_i \in \{-1, 0, 1\}$, since otherwise we could add ξ to our dissociated set and contradict maximality. But this means that all vectors ξ in Λ are contained in the cube with basis vectors $\pm\xi_1, \ldots, \pm\xi_n$, and we are done.

We now prove the Proposition. First we observe that we can replace Λ by the slight variant

$$\Lambda_\theta := \{\xi \in \mathbf{Z}/N\mathbf{Z} : \mathrm{Re} e^{i\theta} \hat{\chi}_A(\xi) \geq \frac{1}{2}\alpha c\}.$$

Since one can cover Λ by a finite number of the Λ_θ, it suffices to prove the claim for a single Λ_θ.

Fix θ, and consider the normalized exponential sum

$$f(x) := \frac{e^{-i\theta}}{\sqrt{n}} \sum_{j=1}^n e^{2\pi i \xi_j x/n}.$$

From Plancherel's theorem, we see that this function is normalized in L^2:

$$\|f\|_{L^2(\mathbf{Z}/N\mathbf{Z}; dx)} = 1.$$

Also from Parseval's inequality, we know that f has a large inner product with χ_A:

$$\langle f, \chi_A \rangle_{L^2(\mathbf{Z}/N\mathbf{Z}; dx)} = \frac{1}{\sqrt{n}} \sum_{j=1}^n e^{-i\theta} \overline{\hat{\chi}_A(\xi_j)},$$

and hence by construction of Λ_θ

$$\mathrm{Re} \frac{1}{N} \sum_{x \in A} f(x) \gtrsim \frac{1}{\sqrt{n}} n\alpha c,$$

or equivalently

$$\frac{1}{|A|} \sum_{x \in A} \mathrm{Re} f(x) \gtrsim \sqrt{n}\alpha.$$

Our aim is to show that $n = O(\alpha^{-2} \log(1/c))$. If this is not true, then $\sqrt{n}\alpha \gg \log(1/c)^{1/2}$, and hence we have

$$\frac{1}{|A|} \sum_{x \in A} \mathrm{Re} f(x) \gg \log(1/c)^{1/2}. \tag{15}$$

Thus while f is L^2 normalized, its average on the subset A of $\mathbf{Z}/N\mathbf{Z}$ is somewhat large. If we just use the L^2 bound, we cannot get a contradiction, because the hypothesis $|A| = cN$ combined with the L^2 normalization only allows us to bound the left-hand side of (15) by $O(c^{-1/2})$ (how? use Cauchy-Schwarz). Fortunately, we can take advantage of the dissociativity to improve our bounds on f. The point is that f is composed of dissociated frequencies, which should act in a largely

independent manner, and so we expect a "large deviation estimate" or "law of large numbers" that says that the magnitude of f does not fluctuate very much, and indeed we expect exponential decay estimates for the set where f is very large (this is where the log is eventually going to come from). Specifically, we have

Lemma 9.5. *In the above notation, we have*

$$\int_{\mathbf{Z}/N\mathbf{Z}} e^{\mathrm{Re} f(x)}\ dx \lesssim 1.$$

Proof The presence of the exponential is a familiar one when dealing with a sum of "independent" quantities, since the exponential converts this to a *product* of independent quantities, at which point one can exploit the independence.

Since

$$\mathrm{Re} f(x) = n^{-1/2} \sum_{j=1}^{n} \cos(2\pi \xi_j x/n - \theta)$$

we have

$$e^{\mathrm{Re} f(x)} = \prod_{j=1}^{n} e^{n^{-1/2}\cos(2\pi i \xi_j x/n - \theta)}.$$

Now we estimate the exponential by a Taylor expansion:

$$e^{n^{-1/2}\cos(2\pi i \xi_j x/n - \theta)} \leq 1 + n^{-1/2}\cos(2\pi i \xi_j x/n - \theta) + O(n^{-1}).$$

Since $(1 + O(n^{-1}))^n = O(1)$, we thus have

$$e^{\mathrm{Re} f(x)} \lesssim \prod_{j=1}^{n} (1 + n^{-1/2}\cos(2\pi i \xi_j x/n - \theta)),$$

which we split using exponentials as

$$e^{\mathrm{Re} f(x)} \lesssim \prod_{j=1}^{n} (1 + \frac{1}{2} n^{-1/2} e^{-i\theta} e^{2\pi i \xi_j x/n} + \frac{1}{2} n^{-1/2} e^{i\theta} e^{-2\pi i \xi_j x/n}).$$

Now we multiply this out. The right-hand side contains a 1, plus a whole bunch of multiples of plane waves $e^{2\pi i \xi x/n}$ for various values of ξ. However, since the ξ_j are dissociated, none of those plane waves are constant. Thus if we sum in x, they all cancel, and we obtain

$$\int_{\mathbf{Z}/N\mathbf{Z}} e^{\mathrm{Re} f(x)}\ dx \lesssim 1$$

as desired. □

From this Lemma and the Jensen inequality

$$\exp(\frac{1}{|A|} \sum_{x \in A} \mathrm{Re} f(x)) \leq \frac{1}{|A|} \sum_{x \in A} e^{\mathrm{Re} f(x)}$$

we obtain the desired contradiction to (15), since $|A| = cN$, and Proposition 9.4 follows. ■

10. Exercises

- Q1. Let G be a subgroup of an abelian group Z, and let $\phi : G \to G'$ be a Freiman homomorphism of order 2 from G to some subset G' of another abelian group Z'. Show that G' is the translate of some subgroup of Z'. (Hint: first normalize so that $\phi(0) = 0$).
- Q2. Let P be a proper generalized arithmetic progression, and let $\phi : P \to \phi(P)$ be a Freiman homomorphism of order 2 from P to some subset $\phi(P)$ of an abelian group Z. Show that $\phi(P)$ is a generalized arithmetic progression of the same length as P, and furthermore that $\phi(P)$ is proper if ϕ is injective. (Hint: This is a very similar argument to Q1. First normalize so that P is a box and $\phi(0) = 0$).
- Q3. Verify all the unproved claims in Section 6.
- Q4. Let A be a subset of a finite abelian group Z with $|A| = c|Z|$ and $|A + A| = |A|$, and let $e(x, \xi)$ be a fixed bi-character used to define the Fourier transform. Show that there is a set $\Lambda \subseteq Z$ of size $|\Lambda| = c^{-1}$ such that $|\hat{\chi}_A| = c\chi_\Lambda$; thus one has perfect concentration of the Fourier transform in a small set. Contrast this with the computations in Section 6. Show in addition that Λ is a subgroup of Z.
- Q5. Let Λ be a lattice (i.e. a discrete additive subgroup) of $\mathbf{R}^d$ whose quotient space $\mathbf{R}^d/\Lambda$ has finite volume. (A typical example is $\mathbf{Z}^d$). Let K be a convex bounded subset of $\mathbf{R}^d$ which is symmetric around the origin, and such that $|K| > 2^d \mathbf{R}^d/\Lambda$. Show that K must contain a lattice point in Λ other than the origin. (This is *Minkowski's first theorem*. Hint: adapt some of the arguments in Section 7). Give an example to show that the constant 2^d in the above theorem cannot be improved.
- Q6*. Prove (13) by doing a volume packing argument on the quotient space $\mathbf{R}^d/\Gamma$. (It is possible to compute a sharp constant in (13) - in fact, it is $2^d/|B(0,1)|$ - and the claim in fact holds for all possible norms on $\mathbf{R}^d$, not just the Euclidean one. This statement is known as *Minkowski's second theorem*, but we will not prove it here).
- Q7 (a). Let $a_1, \ldots, a_N$ be a square-summable sequence of positive reals:

$$\sum_{j=1}^N a_j^2 \leq 1.$$

For any $\lambda > 0$, prove the bound

$$P(|\sum_{j=1}^N \varepsilon_j a_j| > \lambda) \lesssim e^{-c\lambda^2}$$

for some absolute constant $c > 0$, where the $\varepsilon_j = \pm 1$ are independent, identically distributed unbiased random signs, and P denotes the probability of an event. (Hint: compute the expectation of $\exp(t \sum_{j=1}^N \varepsilon_j a_j)$ for any parameter $t > 0$, in the spirit of Lemma 9.5.) This is of course a special case of the *law of large numbers*.

- Q7(b) Under the same assumptions on a_i, prove *Zygmund's inequality*

$$|\{\theta \in [0,2\pi] : |\sum_{j=1}^{N} a_j e^{2\pi i 2^j \theta}| > \lambda\}| \lesssim e^{-c\lambda^2}$$

for any $\lambda > 0$. (In fact, one can replace the frequencies 2^j by any other dissociated set).

References

[1] M. Chang, *A polynomial bound in Freiman's theorem*, Duke Math. J. **113** (2002), no. 3, 399–419.
[2] G. Freiman, *Structure theory of set addition*, Asterisque **258** (1999), 1–33.
[3] B. Green, *Edinburgh lecture notes on Freiman's theorem*, preprint.
[4] I. Ruzsa, *An analog of Freiman's theorem in groups*, Structure theory of set addition, Astérisque No. 258 (1999), 323–326.

Department of Mathematics, UCLA, Los Angeles CA 90095-1555

E-mail address: tao@math.ucla.edu

LECTURE NOTES 3 FOR 254A

TERENCE TAO

1. Introduction

In the previous two sets of notes we have been concerned with the size of the sumset

$$\{a+b : (a,b) \in A \times B\}$$

and difference set

$$\{a-b : (a,b) \in A \times B\}$$

and how the sizes of these sets were related to each other, and whether they revealed anything about the inner structure of A and B.

However, in many applications we do not have control over the full sumset or full difference set: we merely have control over a "large portion" of the sumset or difference set. In other words, we may only have good control on a set such as

$$\{a+b : (a,b) \in G\}$$

or

$$\{a-b : (a,b) \in G\},$$

where G is a "large" subset of $A \times B$, and then we are interested in asking the same sorts of questions (how does a partial sumset control a partial difference set? What does this say about the structure of A and B?)

As one can imagine, the results here are more incomplete than in the case of total sumsets and total difference sets. Nevertheless, we still have some important results, notably the *Balog-Szemeredi theorem*, which asserts that control of partial sumsets can be used to imply control of total sumsets if one passes to a refinement.

To give one indication of how partial sum sets or partial difference sets come up, suppose we have two sets A and B of very different cardinalities, say $|A| \gg |B|$, but we know that A is essentially B-invariant, so that $|A+B| \sim |A|$. Now we want to say something about $A-B$. In week 1 notes we know that we can say something if we lose a small power of A and pass to a refinement. Here is another approach. We have a projection $\pi_{+1} : A \times B \to A+B$ defined by the addition map $\pi_{+1}(a,b) := a+b$. This is a map from a large set (of cardinality $|A||B|$) to a small set (of cardinality $\sim |A|$). This means that there must be substantial failure of injectivity in π_{+1}; more precisely, (see Q1), we must have

$$|\{(a,b,a',b') \in A \times B \times A \times B : a+b = a'+b'\}| \geq \frac{|A|^2|B|^2}{|A+B|} \gtrsim |A||B|^2.$$

Now we use a basic relationship between addition and subtraction: if $a+b = a'+b'$, then $a - b' = a' - b$. Thus if one has many sum collisions, then one must also have many difference collisions:

$$|\{(a,b',a',b) \in A \times B \times A \times B : a - b' = a' - b\}| \gtrsim |A||B|^2.$$

Thus, if we define $\pi_{-1} : A \times B \to A - B$ to be the subtraction map $\pi_{-1}(a,b) := a - b$, then we have

$$|\{(x,x') \in (A \times B)^2 : \pi_{-1}(x) = \pi_{-1}(x')\}| \gtrsim |A||B|^2.$$

To use this, we need a definition. Let us call a difference $d \in A - B$ *popular* if it can be written as $a - b$ in $\gtrsim |B|$ different ways, i.e.

$$|\{x \in A \times B : \pi_{-1}(x) = d\}| \gtrsim |B|.$$

Note that no difference x can be written in more than $|B|$ ways, since once you choose b there is no choice for a (since $a = x + b$). So popular differences are ones which can be written in a near-maximal number of ways. Let G be the set of pairs $x \in A \times B$ for which $\pi_{-1}(x)$ is popular. If $x \notin G$, then there are $\ll |G|$ elements $x' \in A \times B$ for which $\pi_{-1}(x') = \pi_{-1}(x)$. Thus we have

$$|\{(x,x') \in (A \times B)^2 : \pi_{-1}(x) = \pi_{-1}(x'), x \notin G\}| \ll |A \times B||B| = |A||B|^2.$$

Thus we must have

$$|\{(x,x') \in (A \times B)^2 : \pi_{-1}(x) = \pi_{-1}(x'), x \in G\}| \gtrsim |A||B|^2.$$

Since each $x \in G$ contributes at most $|B|$ elements to this set, we thus have

$$|G| \gtrsim |A||B|,$$

i.e. G is a refinement of $A \times B$. The set

$$\{a - b : (a,b) \in G\}$$

is the set of all popular differences. Since each difference requires $\sim |B|$ elements from G, we have

$$|\{a - b : (a,b) \in G\}| \lesssim |G|/|B| \leq |A \times B|/|B| = |A|.$$

To summarize, starting from the hypothesis that all of the sum set $A + B$ had size $\sim |A|$, we can conclude that a large portion of the difference set $A - B$ also has size $\sim |A|$. (One can weaken the hypothesis, assuming instead that a large portion of $\{a + b : (a,b) \in G'\}$ the sumset has size $\sim |A|$, although because at some point we swap (a,b,a',b') with (a,b',a',b) the set G and G' are not directly related.)

We would like to then reverse these types of implications, so that if one assumes that a large part of $A+B$ (for instance) is small, then perhaps all of $A+B$ is small. As stated, such a claim is clearly false. For instance, let N be a large number, and let $A = B$ consist of the arithmetic progression $\{1, 2, \ldots, N\}$, together with N generic integers. Then $|A| = |B| = 2N$, and a large portion of $A + B$ is contained inside the small set $\{2, \ldots, 2N\}$ (indeed, one quarter of $A \times B$ projects onto this set), but $A + B$ is huge, of the order of N^2. The problem here is that A and B contain a large number of "junk" elements which do not contribute at all to the portion G of $A \times B$ which is projecting onto the small set, but which do bloat the size of $A + B$ unnecessarily.

The *Balog-Szemerédi theorem* [1] asserts that in such situations, one can always trim away the junk and make $A+B$ small again:

Theorem 1.1 (Balog-Szemerédi theorem)**.** *Let N be a large integer, and let A and B be finite subsets of an abelian group Z such that $|A| \sim |B| \sim N$. Suppose also that there is a refinement $G \subseteq A \times B$ of $A \times B$ such that the differences of $a-b$ are small:*

$$|\{a-b : (a,b) \in G\}| \lesssim N.$$

Then we can find refinements A', B' of A and B respectively such that

$$|A'-B'| \lesssim N.$$

Recall that a set A' is a *refinement* of A if $A' \subseteq A$ and $|A'| \sim |A|$. From sumset estimates (see Week 1 notes) we know that once $|A'-B'| \lesssim N$, we know that A' and B' (and even $A' \cup B'$) are essentially closed under addition, subtraction, and any other finite combination of these operations; if we are in a torsion-free group we also know that A', B' are a refinement of a generalized arithmetic progression of small rank. IN short, we know just about everything about A' and B'. So the Balog-Szemerédi theorem says that if most of $A-B$ is small, then A and B both contain large components which are very highly structured. (As mentioned before, the rest of A and B may well be random noise).

Note that the claim would be trivial if the set G contained a large Cartesian product $A' \times B'$. However, all we know is that G is a large subset of $A \times B$, and this is not enough to obtain a large Cartesian product (see Q3).

We now present a proof by Gowers [4] of the Balog-Szemerédi theorem which gives quite good bounds (the final constants are basically just polynomials of the original constants, as opposed to the exponential (or worse) constants that can come out of things like Freiman's theorem).

Note: throughout these notes, we shall be using the Cauchy-Schwarz estimate (Q1) and the popularity argument (Q2) quite often.

2. Proof of the Balog-Szeméredi theorem

The basic idea of the proof is as follows. Fix N, A, B, G; clearly we may assume that A and B are disjoint (by translating one set if necessary). We need to find refinements A' and B' such that $A'-B'$ is very small, of size $O(N)$. Since $|A'| \sim N$ and $|B'| \sim N$, the trivial bound for $A'-B'$ is $O(N^2)$ - because the differences $a-b$ of $A'-B'$ could all be distinct. This is far too weak a bound. *However*, if we knew that all the differences of $A'-B'$ were *popular* - i.e. for every $a-b$ in $A'-B'$, there existed $\gtrsim N$ other pairs $(a',b') \in A \times B$ for which $a-b = a'-b'$ - then this would ensure that $A'-B'$ only has cardinality $O(N)$ (because each distinct element $A'-B'$ is associated to $\gtrsim N$ pairs from $A \times B$, and there are only $\sim N^2$ pairs in $A \times B$ to go around). So, a naive way to try to prove this theorem is to try to choose A' and B' so that all the differences in $A'-B'$ are popular. (This is a manifestation of a basic principle: if one wants to make a set (in this case $A'-B'$)

small, try to make it so that each element of the set consumes a large number of resources (in this case, a large number of differences), and ensure that the total number of resources (in this case, the pairs $A \times B$) is bounded. The main problem is to prevent a proliferation of a large number of elements, each of which consumes very little resources).

Define a (symmetric) relation $\sim$ on $A \times B \cup B \times A$ by setting $a \sim b$ or $b \sim a$ if $a - b$ is a popular difference. From the hypothesis on G and the popularity argument in Q2 we see that this relation is quite dense; the set of pairs $(a, b) \in A \times B$ such that $a \sim b$ is $\gtrsim N^2$. We can think of this as a bipartite graph; for sake of analogy, we will think of A, B as distinct sets of people, and think of the relation $a \sim b$ as saying that person a *knows* person b. By the resource consumption argument given above, the set

$$D := \{a - b : a \sim b\}$$

of popular differences is small:

$$|D| \lesssim N.$$

Unfortunately, even though we have lots of connections (lots of pairs of people know each other), we don't necessarily have large sets A', B' such that *all* of the people in A' know *all* the people in B' - i.e. we don't know that we can make A' and B' *totally connected* via the relation $\sim$. (See Q3; also if you have some experience with Ramsey theory, you know that this is a rather tall order).

The trick is to weaken the relation $\sim$ by using degrees of separation (as in the infamous "six degrees of separation"). We first set $0 < \varepsilon \leq 1$ to be a small parameter (comparable to 1) to be chosen later, and define a (symmetric) relation $\sim\sim$ on $B \times B$ by setting $b \sim\sim b'$ if

$$|\{a' \in A : b \sim a' \sim b'\}| \geq \varepsilon N.$$

i.e. two people b, b' are related by $\sim\sim$ if they are connected by one degree of separation in multiple ($\gtrsim N$) ways. We will say that b and b' *communicate* if they are related by $\sim\sim$ (they don't directly know each other, but they have a lot of mutual friends). We then define a relation $\sim\sim\sim$ on $A \times B$ by setting $a \sim\sim\sim b$ if

$$|\{b' \in B : a \sim b' \sim\sim b\}| \gtrsim N;$$

roughly speaking, we have $a \sim\sim\sim b$ if a and b are connected by three degrees of separation in multiple ($\gtrsim N^2$) ways. We say that a is *aware of* b if $a \sim\sim\sim b$ (they may not directly know each other, but a knows many people who communicate with b). The intuition is that this relation is much weaker than than $\sim$, and thus it should be easier to find large groups A', B' of people which are totally connected by this relation. (There is an analogy here with some arguments in the previous notes, where a set A did not contain a large arithmetic progression, but the variant $2A - 2A$ did).

First, we check that even though the awareness relation $\sim\sim\sim$ is weak, it still consumes enough resources that the difference set is still small:

Lemma 2.1. *The set* $\{a - b : a \sim\sim\sim b\}$ *has cardinality* $O(N)$.

Proof Let $a \sim\sim\sim b$. By definition, we can find $\gtrsim \varepsilon N^2$ pairs (a', b') such that $a \sim b' \sim a' \sim b$. From the identity

$$a - b = (a - b') - (a' - b') + (a' - b)$$

we thus see that there are $\gtrsim N^2$ solutions to the equation

$$a - b = d_1 - d_2 + d_3; \quad d_1, d_2, d_3 \in D.$$

(Note that different pairs (a', b') will lead to different triples of popular differences (d_1, d_2, d_3)). Thus each difference $a - b$ with $a \sim\sim\sim b$ is associated with $\gtrsim N^2$ triples (d_1, d_2, d_3) in D^3. But as discussed before D only has cardinality $O(N)$, hence D^3 has cardinality $O(N^3)$; thus there are only at most $O(N)$ possible values for $a - b$, and the lemma follows. ∎

So the whole problem now is to find large sets A', B' which are totally connected by the awareness relation $\sim\sim\sim$. An obvious sub-problem here is to find a large set B' such that the elements of B' are related to each other by the communication relation $\sim\sim$ quite often.

The first thing to do is to refine B by kicking out all the unfriendly people, as they clearly are not going to be part of the final set B'. Let $B_1 \subset B$ be the set of all the people b in B who know $\gtrsim N$ people in A:

$$|\{a \in A : a \sim b\}| \gtrsim N \text{ for all } b \in B_1.$$

By the popularity argument we know that

$$|\{(a, b) \in A \times B_1 : a \sim b\}| \gtrsim N^2$$

and in particular that $|B_1| \sim N$. So even after we remove all the unfriendly people from B, we still have a large group of people B_1, and a lot of pairs of people between A and B_1 who know each other.

Now, we pick a random person $a_0 \in A$ and set $B_2 = B_2(a_0) \subseteq B_1$ to be all the people that a_0 knows directly:

$$B_2 := \{b \in B_1 : b \sim a_0\}.$$

Let us call $B_2(a_0)$ the *friends* of a_0. The point is that a pair of communicating people b, b' are much more likely to both be the friends of a_0, than a pair of non-communicating people b, b' (since they just don't have that many mutual friends). Inverting this, this should mean that most friends of a_0 communicate a lot.

Also, while the set B_2 may occasionally be small (because a_0 might be particularly unfriendly), it will be large on the average. Indeed, since $|\{(a, b) \in A \times B_1 : a \sim b\}| \sim N^2$, we see from the popularity argument (Q2) that $|B_2|$ has expectation $\sim N$, and in particular we have $|B_2| \sim N$ with probability ~ 1.

Now let's condition on the event that $|B_2| \sim N$; this event has probability ~ 1 so this conditioning does not significantly distort probabilities. Take two uncommunicative people $b, b' \in B_1$, so that $b \not\sim\sim b'$. By definition, b and b' communicate with $O(\varepsilon N)$ people in A. Thus the probability that these people both lie in B_2

is $O(\varepsilon)$. Summing over all uncommunicative pairs (there are $O(N^2)$ of them) and using linearity of expectation, we obtain

$$E(|\{(b,b') \in B_2 : b \not\sim b'\}|) \lesssim \varepsilon N^2 \lesssim \varepsilon |B_2|^2.$$

Thus on the average, at most $O(\varepsilon)$ of the pairs in B_1 are uncommunicative. By the pigeonhole principle, we may thus find an $a_0 \in A$ such that $B_2 := B_2(a_0)$ has size $|B_2| \sim N$ and we have the bound

$$|\{(b,b') \in B_2 : b \not\sim b'\}| \lesssim \varepsilon |B_2|^2.$$

Call an element $b \in B_2$ *uncommunicative* if

$$|\{b' \in B_2 : b \not\sim b'\}| \gtrsim \varepsilon^{1/2} |B_2|.$$

Clearly, there are at most $O(\varepsilon^{1/2}|B_2|)$ uncommunicative elements of B_2 (this is the popularity argument again!). Thus if we let $B_* \subseteq B_2$ be the uncommunicative elements of B_2, then B_* is fairly small:

$$|B_*| \lesssim \varepsilon^{1/2} |B_2|.$$

By construction of B_1, everybody in B_1 knows $\gtrsim N$ people in A. Since B_2 is a subset of B_1, we thus have

$$|\{(a,b) \in A \times B_2 : a \sim b\}| \gtrsim N|B_2| \gtrsim N^2;$$

thus a lot of people in A know a lot of people in B_2. By the popularity argument, there thus exists a refinement A' of A such that everybody in A' knows many people from B_2:

$$|\{b' \in B_2 : a \sim b'\}| \gtrsim N \text{ for all } a \in A'.$$

We are now ready to define B': we simply set $B' := B_2 \backslash B_*$, i.e. we choose all the communicative people from B_2. Since B_* is so small (if ε is small enough), we have $|B'| \sim N$, so that B' is indeed a refinement of B. We now claim that $a \sim\sim\sim b$ for every $a \in A'$, $b \in B'$.

To prove this, fix $a \in A'$ and $b \in B'$. By construction, a knows $\gtrsim N$ people b' in B_2. Since B_* is so small (if ε is small enough), a must therefore know $\gtrsim N$ people b' in B'. Since b is not uncommunicative, it will communicate with all but $O(\varepsilon^{1/2}N)$ people in B'. Thus, if ε is small enough, it will communicate with $\gtrsim N$ of the people b' mentioned above - and hence $a \sim\sim\sim b$. This, combined with the previous Lemma, shows that $|A' - B'| \lesssim N$, and the result follows. □

Note how every time we make the relation weaker, we can make it denser. Regarding the original knowledge relation $a \sim b$, we could only say things like "people in B_1 know many ($\gtrsim N$) people in A". For the slightly weaker communication relation $b \sim\sim b'$, we could say things like "people in B' communicate all but a few ($O(\varepsilon^{1/2}N)$) people in B_2". And finally regarding the weakest notion of awareness $a \sim\sim\sim b$, we can say that "people in A' are aware of *all* the people in B'".

The Balog-Szemerédi theorem can be made a little more quantitative: an inspection of the above proof reveals that the constants only grow polynomially rather than exponentially. More precisely, we have

Theorem 2.2 (Quantitative Balog-Szemerédi theorem)**.** *Let N be a large integer, and let A and B be finite subsets of an abelian group Z such that $|A| = |B| = N$. Suppose also that there is a set $G \subseteq A \times B$ of $A \times B$ with $|G| \geq N^2/K$ such that the differences of $a - b$ are small:*

$$|\{a - b : (a, b) \in G\}| \leq KN.$$

Then we can find subsets A', B' of A and B respectively such that

$$|A'|, |B'| \gtrsim K^{-C} N$$

and

$$|A' - B'| \lesssim K^C N.$$

One could even work out these constants C explicitly (they are of the order of 10 or so, depending on how efficiently you run the above argument), but we will not do so here. This quantitative version has applications to arithmetic progressions, as we will see in a later set of notes; for now, however, we shall focus on its application to the Kakeya problem.

3. A variant of the Balog-Szemerédi theorem, and applications to the Kakeya problem

Before we get to the Kakeya problem, we need a variant of the Balog-Szemerédi theorem. To do so, let us first digress on popular rows and columns.

Suppose that A and B are finite sets with cardinality $|A| \sim |B| \sim N$, and let $H \subseteq A \times B$ be a large subset of $A \times B$, so that $|H| \sim N^2$. One can think of $A \times B$ as looking roughly like a square, either considered as the union of $\sim N$ rows of the form $A \times \{b\}$, or as the union of $\sim N$ columns of the form $\{a\} \times B$. Each of these rows and columns contain a certain number of elements from H.

Since there are $\sim N$ rows, and the total number of elements of H is $\sim N^2$, we see that each row contains $\sim N$ *on the average*. However, it is not necessarily the case that *every* row contains $\sim N$ elements of H; there might be a few rows which contain far fewer. However, such rows, by definition, do not contain many elements of H. Indeed, by the popularity argument (Q2), if we throw away all the *unpopular rows* (those rows which only have $\ll N$ elements of H), then one still retains a large fraction of H (and in particular, we still have $\sim N^2$ elements of H). Since each row can contain at most $O(N)$ elements from H, the number of rows remaining must still be comparable to N. Thus one can always refine things a little bit so that all the rows become popular.

One can similarly throw away the unpopular columns and ensure that all the columns are popular, while still retaining a large fraction of H. Unfortunately, it is not so easy to *simultaneously* ensure that the rows and columns are popular, because the operation of throwing away unpopular rows may turn previously popular columns unpopular, and vice versa. However, we always have the freedom to make *either* the rows popular, or the columns popular, depending on which is more convenient at the time.

By doing this, one can modify the previous Balog-Szemerédi argument to retain a large portion of a reference set H:

Theorem 3.1 (Modified Balog-Szemerédi theorem)**.** *Let N be a large integer, and let A and B be finite subsets of an abelian group Z such that $|A| \sim |B| \sim N$. Suppose also that there is a refinement $G \subseteq A \times B$ of $A \times B$ such that the differences of $a - b$ are small:*

$$|\{a - b : (a, b) \in G\}| \lesssim N.$$

Suppose also we have another refinement $H \subseteq A \times B$ (so that $|H| \sim N^2$). Then we can find refinements A', B' of A and B respectively such that

$$|A' - B'| \lesssim N$$

and also

$$|H \cap (A' \times B')| \sim N^2.$$

Note that if one sets $H := A \times B$ then one gets the ordinary Balog-Szemerédi theorem. The proof of this theorem is almost the same as that of the usual Balog-Szemerédi, however every time one prepares to pass from A or B to a refinement of A or B, one must make sure that either all the columns are popular or all the rows are popular (so that the refinement captures most of H). We leave the actual verification of this modified version to Q4.

A more quantitative version is available. The following bound was proven by Bourgain [2], using the above arguments:

Theorem 3.2. [2] *Let N be a large integer, and let A and B be finite subsets of an abelian group Z such that $|A| \sim |B| \sim N$. Suppose also that there is a set $H \subseteq A \times B$ of cardinality*

$$|H| \geq N^2/K$$

whose sum set is small:

$$|\{a + b : (a, b) \in H\}| \leq N.$$

Then we can find subsets A', B' of A and B respectively such that

$$|H \cap (A' \times B')| \gtrsim K^{-9} N^2.$$

and also

$$|A' - B'| \lesssim K^{13} N^{-1} |H \cap (A' \times B')|.$$

Note that the existence of this set H with small sumset implies the existence of the set G with small difference set, by the argument in the introduction.

Now we apply this Theorem to Kakeya sets. These sets can be studied both in Euclidean space or in finite fields; for simplicity we treat the finite field case (which has fewer technicalities).

Let F be a finite field (e.g. $F := \mathbf{Z}/p\mathbf{Z}$ for some prime p), and let $n \geq 2$ be an integer, thus F^n is vector space over F. We think of n as being fixed, but $|F|$ as

being large; our implicit constants may depend on n but not on F. A *line* in F^n is any set l of the form

$$l = \{x + tv : t \in F\}$$

where $x \in F^n$ and $v \in F^n \backslash \{0\}$; we refer to v as the *direction* of l (though this direction is only defined up to scalar multiplication by invertible elements of F, thus there are about $|F|^{n-1}$ possible distinct directions). A *Besicovitch set* in F^n is said to be any set $E \subseteq F^n$ which contains a line in every direction.

The *finite field Kakeya set conjecture* asserts that one has the bound

$$|E| \geq c(n, \varepsilon)|F|^{n-\varepsilon}$$

for every $\varepsilon > 0$ and some $c(n, \varepsilon) > 0$ independent of $|F|$; thus Besicovitch sets are conjectured to fill up practically all of F^n. This conjecture is known to be true in two dimensions, but is not known in higher dimensions; however, partial bounds exist. For instance, here is an easy estimate:

Proposition 3.3. *Every Besicovitch set E in F^n has cardinality*

$$|E| \gtrsim |F|^{(n+1)/2}.$$

Proof Fix the Besicovitch set E. We may assume that F is large (actually we just need $|F| \geq 4$) since the claim is trivial otherwise.

For every $t \in F$, consider the "horizontal plane" $F^{n-1} \times \{t\}$, and let E_t denote the "horizontal slice"

$$E_t := E \cap (F^{n-1} \times \{t\}).$$

Clearly we have

$$\sum_{t \in F} |E_t| = |E|$$

so the average size of $|E_t|$ is $|E|/|F|$. Indeed, from Chebyshev's inequality we have

$$|\{t \in F : |E_t| \geq 2|E|/|F|\}| \leq |F|/2$$

and hence we have at least $|F|/2$ values of t for which $|E_t| \leq 2|E|/|F|$. In particular, we have at least two such values of t. By an affine transformation, we may assume those values are 0 and 1, thus

$$|E_0|, |E_1| \leq 2|E|/|F|.$$

Now consider the lines in E. For every "velocity" $w \in F^{n-1}$, there must be a line l_w with direction $(1, w)$ that lies in E. This line must intersect E_0 at some point $a(w) \in E_0$, and intersect E_1 at some other point $b(w) \in E_1$, so that $b(w) - a(w) = (1, w)$. Thus to each $w \in F^{n-1}$ we can associate a pair $(a(w), b(w)) \in E_0 \times E_1$ of points, so that the line between them lies in E and has direction $(1, w)$. But since two points determine at most one line, all these pairs are distinct. Thus if we let H be the set of all such pairs, we have $|H| = |F|^{n-1}$. But we also have the trivial bound

$$|H| \leq |E_0 \times E_1| = |E_0||E_1| \lesssim (|E|/|F|)^2.$$

Combining the two bounds we get $|E| \gtrsim |F|^{(n+1)/2}$ as desired. ∎

We now use Bourgain's version of the Balog-Szemerédi theorem to improve this a bit for large dimension n, though we will have to assume F has characteristic greater than 2. (The *characteristic* of a finite field F is the prime p such that $px = 0$ for all $x \in F$).

Proposition 3.4. *Let F have characteristic greater than 2. Every Besicovitch set E in F^n has cardinality*

$$|E| \gtrsim |F|^{(13n+12)/25}.$$

Proof Again, we take $|F|$ large, and consider the slices E_t. By Chebyshev we have

$$|\{t \in F : |E_t| \geq 10|E|/|F|\}| \leq |F|/10.$$

In particular, we can find a non-trivial arithmetic progression a, $a+r$, $a+2r$ of length 3 such that $|E_a|, |E_{a+r}|, |E_{a+2r}| \leq 10|E|/|F|$. (Indeed, we just pick a and r randomly, and we will have this happening with at least a 70% chance). By an affine transformation (using the assumption that the characteristic is greater than 2) we may set $a := 0$ and $r := 1/2$. Thus we have $|E_0|, |E_{1/2}|, |E_1| \leq N$, where we set $N := 10|E|/|F|$.

As before, for every velocity $w \in F^{n-1}$ we have a pair $(a(w), b(w))$ in $E_0 \times E_1$, such that $b(w) - a(w) = (1, w)$ and the line between $a(w)$ and $b(w)$ lies in E. In particular, the midpoint $(a(w) + b(w))/2$ lies in $E_{1/2}$. Thus, if we let $G \subseteq E_0 \times E_1$ be the set of pairs of the form $(a(w), b(w))$, then $|G| = |F|^{n-1}$, and in fact

$$|\{b - a : (a, b) \in G\}| = |F|^{n-1}$$

(since all the differences in G are distinct), while

$$|\{b + a : (a, b) \in G\}| = |E_{1/2}| \leq N.$$

Write $K := N^2/|F|^{n-1}$, so that $|G| = N^2/K$. Now we apply Bourgain's version of the Balog-Szemerédi theorem, with H replaced by G, to obtain sets $A' \subseteq E_0$ and $B' \subseteq E_1$ such that

$$|A' - B'| \lesssim K^{13} N^{-1} |G \cap (A' \times B')|.$$

However, since all the differences in G are distinct, we have

$$|G \cap (A' \times B')| \leq |A' - B'|$$

which implies that

$$K^{13} N^{-1} \lesssim 1$$

which (after some algebra) implies that $N^{25/13} \gtrsim |F|^{n-1}$, and hence that $|E| \gtrsim |F|^{(13n+12)/25}$. ∎

4. Another approach to the Kakeya problem

The material here is only a brief taste of what goes into the Kakeya problem. For a more thorough discussion of the Kakeya problem and its connection to arithmetic combinatorics, see [7].

While the Balog-Szemerédi theorem has many uses (in particular it is used in Gowers proof of Szemerédi's theorem, to be covered in a later set of notes), it

is not the most efficient tool to attack the Kakeya problem. One can improve upon the above result by dealing with the slices E_t and the set G directly.

To simplify the discussion, let us assume for the moment that the slices E_t are evenly distributed, so each slice has cardinality comparable to its average value $N := |E|/|F|$. Also write $Z := F^n$. Then, as in the proof of Proposition 3.4, we have a set $G \subseteq Z \times Z$ of cardinality $|G| = |F|^{n-1}$ whose differences $b - a$ are all distinct, but whose sums are all small

$$|\{a + b : (a, b) \in G\}| \leq |E_{1/2}| \lesssim N.$$

In other words, we have

$$|\pi_+(G)| \lesssim N.$$

More generally, given any slice E_t with $t \neq 0$, we have

$$|\{(1 - t)a + tb : (a, b) \in G\}| \leq |E_t| \lesssim N$$

and so if we define $\pi_r(a, b) := a + rb$ for any $r \in F$, we have

$$|\pi_{t/(1-t)}(G)| \lesssim N.$$

Thus almost all of the arithmetic projections of G are small (of size $O(N)$), except for the projection π_{-1}, which is large, because of the distinct differences property:

$$|\pi_{-1}(G)| = |G| = |F|^{n-1}.$$

Our task is to get a lower bound for $|E|$ in terms of $|F|$, or - what amounts to the same thing - getting an upper bound for $|F|$ in terms of N. This leads to the following question: if we have upper bounds on various projections $\pi_r(G)$ of a set G, does this imply upper bounds on the size of $\pi_{-1}(G)$? Well, have the obvious bound

$$|\pi_{-1}(G)| \leq |G| \leq |\pi_0(G) \times \pi_\infty(G)| = |\pi_0(G)||\pi_\infty(G)|$$

where we adopt the convention $\pi_\infty(a, b) := b$. So if we have $|\pi_0(G)|, |\pi_\infty(G)| \lesssim N$, then $|G| \lesssim N^2$; combining this with the formula $N = |E|/|F|$ we get the bound $|E| \gtrsim N^{(n+1)/2}$, which we obtained before.

The argument in Proposition 3.4 implies the following improvement to this trivial bound:

Proposition 4.1. *Let $G \subset Z \times Z$ be a set such that $|\pi_0(G)|, |\pi_1(G)|, |\pi_\infty(G)| \lesssim N$. Then $|\pi_{-1}(G)| \lesssim N^{2-1/13}$.*

This corresponds to the numerology $|E| \gtrsim N^{(13n+12)/25}$, as discussed earlier. The proof of this Proposition is almost identical to that in Proposition 3.4 (note that we can assume that π_{-1} is injective on G, by removing redundant elements of G if necessary).

We can improve this bound a bit (at the cost of requiring one more projection), by using a somewhat different argument than that used to prove the Balog-Szemerédi theorem.

Proposition 4.2. [5] *Let $G \subset Z \times Z$ be a set such that $|\pi_0(G)|, |\pi_1(G)|, |\pi_2(G)|, |\pi_\infty(G)| \lesssim N$. Then $|\pi_{-1}(G)| \lesssim N^{2-1/4}$.*

This bound then leads to the improved bound $|E| \gtrsim N^{(4n+3)/7}$ for Besicovitch sets when the characteristic is greater than 3.

Proof We think of G as living on a two-dimensional space $Z \times Z$. We shall construct various geometric objects in this space. We may of course assume that π_{-1} is injective on G, so that we just need to bound $|G|$.

Define a *vertical line segment* to be a pair $(g, g') \in G \times G$ such that $\pi_0(g) = \pi_0(g')$; we abuse notation and write write $\pi_0((g, g'))$ for $\pi_0(g) = \pi_0(g')$. Let V be the set of all vertical line segments. By Cauchy-Schwarz (Q1) we have

$$|V| \gtrsim |G|^2/N. \tag{1}$$

We define a new relation $\sim_0$ by setting $v_1 \sim_4 v_2$ if v_1 and v_2 have the same value of $\pi_2 \otimes \pi_\infty$. By Cauchy-Schwarz (Q1) we have

$$|\{(v_1, v_2) \in V \times V : v_1 \sim_0 v_2\}| \gtrsim |V|^2/N^2 \gtrsim |G|^4/N^4$$

by (1).

We will complement this lower bound with an upper bound

$$|\{(v_1, v_2) \in V \times V : v_1 \sim_0 v_2\}| \lesssim N^3$$

which will give $|G| \lesssim N^{2-1/4}$. This is better than the trivial bound of N^4 (which can be obtained, e.g. by specifying four projections of the four vertices of the trapezoid formed by v_1, v_2. The key is to use the injectivity of π_{-1}, and more precisely the identity

$$\pi_{-1}(g_2') = -\pi_1(g_1') + 2\pi_1(g_2) - 2\pi_\infty(g_1).$$

Thus if one specifies $\pi_1(g_1')$, $\pi_1(g_2)$, and $\pi_\infty(g_1)$, then one also specifies $\pi_{-1}(g_2')$, hence g_2' (by injectivity of π_{-1}), which then together with $\pi_1(g_2)$ gives g_2, which together with $\pi_\infty(g_1)$ gives g_1, which together with $\pi_1(g_1')$ gives g_1'. Thus (v_1, v_2) is completely determined by the three projections $\pi_1(g_1')$, $\pi_1(g_2)$, and $\pi_\infty(g_1)$, and the upper bound of $O(N^3)$ follows. ■

If we are allowed to use more projections than π_0, π_1, π_2, π_∞, then we can improve the gain $1/4$ in $N^{2-1/4}$ somewhat; the record to date is $N^{2-0.325\ldots}$, using a very large number of slices (and requiring the characteristic of F to be large) [6]. If we could get down to $N^{1+\varepsilon}$ for any ε then we would have proven the Kakeya conjecture (not just in finite fields, but also in Euclidean space; see Q5).

If one doesn't wish to involve the projection $\pi_2(G)$, the results we have are weaker:

Proposition 4.3. [5] *Let $G \subset Z \times Z$ be a set such that $|\pi_0(G)|, |\pi_1(G)|, |\pi_\infty(G)| \lesssim N$. Then $|\pi_{-1}(G)| \lesssim N^{2-1/6}$.*

This bound then leads to the improved bound $|E| \gtrsim N^{(6n+5)/11}$ for Besicovitch sets when the characteristic is greater than 2.

Proof We use the same vertical line segments V as before. Again we assume π_{-1} is injective on G.

Now we define some relations between vertical line segments. If $v_1 = (g_1, g_1')$ and $v_2 = (g_2, g_2')$ are vertical line segments, we set $v_1 \sim_1 v_2$ if $\pi_\infty(g_1) = \pi_\infty(g_2)$ and $\pi_\infty(g_1') = \pi_\infty(g_2')$ (i.e. v_1 and v_2 form a rectangle). Equivalently, $v_1 \sim_1 v_2$ if v_1 and v_2 have the same image under $\pi_\infty \oplus \pi_\infty : V \to \pi_\infty(G) \times \pi_\infty(G)$.

Similarly, we define a relation $\sim_2$ on vertical line segments by setting $v_1 \sim_2 v_2$ if v_1 and v_2 have the same image under $\pi_1 \oplus \pi_1$ (so v_1 and v_2 form a parallelogram), and finally define $v_1 \sim_3 v_2$ if v_1 and v_2 have the same inage under $\pi_1 \oplus \pi_\infty$ (so v_1 and v_2 form a certain type of trapezoid).

Heuristically speaking, given two randomly selected vertical line segments v_1, v_2, the probability that $v_1 \sim_1 v_2$ should be about $1/N^2$, since the image of the map $\pi_1 \oplus \pi_1$ has cardinality $O(N^2)$. Similarly for $\sim_2$ and $\sim_3$. Thus, if we pick four vertical line segments at random, v_1, v_2, v_3, v_4, the probability that

$$v_1 \sim_1 v_2 \sim_2 v_3 \sim_3 v_4$$

should be about $1/N^6$, and so the number of such quadruplets should be about $|V|^4/N^6$. In general, we have this as a lower bound:

$$|\{(v_1, v_2, v_3, v_4) \in V^4 : v_1 \sim_1 v_2 \sim_2 v_3 \sim_3 v_4\}| \gtrsim |V|^4/N^6;$$

this is part of an abstract combinatorial lemma, see Q6.

We now claim an upper bound

$$|\{(v_1, v_2, v_3, v_4) \in V^4 : v_1 \sim_1 v_2 \sim_2 v_3 \sim_3 v_4\}| \lesssim |V|N^2;$$

combining this with the previous bound and (1) we obtain the desired bound on $|G|$ after some algebra. The point of this upper bound $|V|N^2$ is that it improves somewhat on the trivial bound on $|V|N^3$ (since one can completely specify a quadruple (v_1, v_2, v_3, v_4) in the above set by first specifying v_1, then (say) specifying $\pi_0(v_2)$, $\pi_0(v_3)$, and $\pi_0(v_4)$). Of course, to do so we have to use our hypothesis that G has distinct differences (i.e. π_{-1} is injective on G). Basically, the idea is to use this injectivity to cut down the number of degrees of freedom inherent in this system (v_1, v_2, v_3, v_4) of vertical line segments.

Write $v_j = (g_j, g_j')$ for $j = 1, 2, 3, 4$. There are two observations. The first is that v_1, v_2, v_3 all have the same length:

$$g_1 - g_1' = g_2 - g_2' = g_3 - g_3'.$$

In particular, if one specifies v_1, then this already fixes the length of v_2 and v_3.

The second observation is that the length of v_3, combined with the projections $\pi_0(v_3)$ and $\pi_\infty(g_4)$, determines the projection $\pi_{-1}(g_4')$ by the identity

$$\pi_{-1}(g_4') = \pi_\infty(g_4) - \pi_0(v_3) - (g_3 - g_3').$$

Thus if one specifies v_1, $\pi_0(v_3)$, and $\pi_\infty(g_4)$, this determines $\pi_{-1}(g_4')$. But since π_{-1} is injective, this determines g_4'; this, together with $\pi_\infty(g_4)$ determines v_4. This, together with $\pi_0(v_3)$, determines v_3; and this together with v_1 determines v_2. Thus (v_1, v_2, v_3, v_4) is completely determined by the three parameters v_1, $\pi_0(v_3)$, $\pi_\infty(g_4)$, which explains the upper bound of $|V|N^2$. ∎

For future application we note an extension of this Proposition, in which we impose some precise control on the multiplicity of π_{-1}.

Proposition 4.4. [5] *Let $G \subset Z \times Z$ be a set such that $|\pi_0(G)|, |\pi_1(G)|, |\pi_\infty(G)| \lesssim N$. Suppose also that π_{-1} has multiplicity exactly K on G, in the sense that*

$$|\{x \in G : \pi_{-1}(G) = y\}| = K$$

for all $y \in \pi_1(G)$. Then $|\pi_{-1}(G)| \lesssim N^{2-1/6}K^{-5/6}$.

We leave the proof of this Proposition (which is a mild generalization of the previous one) to Q9.

There is however a limit as to how far these methods can go if we refuse to admit too many slices. For instance, the following argument of Ruzsa shows that the $1/6$ in the above estimate cannot be improved to beyond $0.20824\ldots$:

Proposition 4.5. *Let m be a large number, and let $N := (27/4)^m$. There exists a torsion-free abelian group Z and a $G \subseteq Z \times Z$ such that*

$$|\pi_0(G)|, |\pi_1(G)|, |\pi_2(G)|, |\pi_\infty(G)| \lesssim N$$

but

$$|\pi_{-1}(G)| \gtrsim 27^m = N^{\log(27)/\log(27/4)} = N^{2-0.20824\ldots}.$$

Proof Let m be a large integer. We take $Z = \mathbf{Z}^{3m}$, the space of $3m$-tuples of integers. Let $A \subset Z$ denote the space of $3m$-tuples t which consist of $2m$ zeroes and m ones; clearly

$$N := |A| = \frac{(3m)!}{m!(2m)!} \approx (27/4)^m = N,$$

where we have used Stirling's formula $n! \sim n^{1/2}(n/e)^n$ (see Q8).

Now let $G \subset A \times A$ denote the space of pairs (t_1, t_2) of $3m$-tuples, such that the m ones in t_1 occupy a completely disjoint set of co-ordinates than the m ones in t_2. Thus

$$|G| = \frac{(3m)!}{m!m!m!} \approx (27)^m = N^{\log(27)/\log(27/4)}.$$

One can easily check that π_0, π_1, π_∞ all map G to a space isomorphic to A, and so we are done. ∎

5. Application to multilinear convolution-type operators

This section is set in the continuous case $\mathbf{R}$, and so we now use $|E|$ to denote Lebesgue measure rather than cardinality of E.

We now present an argument of Michael Christ [3] which uses Proposition 4.4 to give a non-trivial estimate for the multilinear convolution-type operator T, defined by

$$T(f,g,h)(x) := \int_{\mathbf{R}} f(x+t)g(x-t)h(t)\ dt,$$

where f, g, h are functions on the real line. (If we set $h(t) := 1/t$ then this is the famous *bilinear Hilbert transform*, but the results here will not address this important operator).

To simplify the discussion, let us assume that f, g, h are characteristic functions

$$f = \chi_F, g = \chi_G, h = \chi_H$$

and F, G, H have Lebesgue measure $O(1)$. Then the function $T(f,g,h)(x)$ is just

$$T(f,g,h)(x) = \int_{\mathbf{R}} \chi_F(x+t)\chi_G(x-t)\chi_H(t)\ dt = |H \cap (x-G) \cap (F-x)|.$$

In particular, $T(f,g,h)$ is non-negative and uniformly bounded. Furthermore, since $T(f,g,h)$ is pointwise dominated by $\chi_F * \chi_G(2x)$, it is in L^1. In particular, we have the distributional bounds

$$|\{T(f,g,h) \geq \lambda\}| \lesssim \lambda^{-1}$$

for $0 < \lambda \lesssim 1$.

One can ask the question as to whether these bounds can be improved. Interestingly, one can do so by means of the argument in Proposition 4.4, with the same gain of $1/6$:

Proposition 5.1. [3] *For $0 < \lambda \lesssim 1$, we have*

$$|\{T(f,g,h) \geq \lambda\}| \lesssim \lambda^{-1+1/6}$$

Fix λ. First of all, we observe by a standard limiting argument that we may assume that F, G, H are all unions of intervals of some very small length $0 < \delta \ll \lambda$. Next, we observe the following continuous version of Proposition 4.4:

Proposition 5.2. [5] *Let $G \subset \mathbf{R} \times \mathbf{R}$ be a set consisting of unions of $\delta \times \delta$ squares such that $|\pi_0(G)|, |\pi_1(G)|, |\pi_\infty(G)| \lesssim N$. Suppose also that π_{-1} has fibers of measure $\sim K$ on G, in the sense that*

$$|\{x \in G : \pi_{-1}(G) = y\}| \sim K$$

for all $y \in \pi_{-1}(G)$. Then $|\pi_{-1}(G)| \lesssim N^{2-1/6}K^{-5/6}$.

Indeed, one can rescale δ to be 1 (replacing N by N/δ and K by K/δ, replace the set G with a discretized variant on $\mathbf{Z} \times \mathbf{Z}$, and apply Proposition 5.2. (We leave this as an exercise).

To apply this lemma, what we do is that we let E be the set

$$E := \{x \in \mathbf{R} : T(f,g,h) \geq \lambda\} + O(\delta),$$

and for each $x \in E$ we choose a subset T_x of $H \cap (x-G) \cap (F-x) + O(\delta)$ of measure $\sim \lambda$. We then set

$$\mathbf{G} := \{(t+x, t-x) : x \in E, t \in T_x\}.$$

Then $\mathbf{G}$ is essentially the union of $\delta \times \delta$-squares, and one has the estimates

$$|\pi_0(\mathbf{G})|, |\pi_1(\mathbf{G})|, |\pi_\infty(\mathbf{G})| \lesssim 1$$

by the hypotheses on F, G, H. Furthermore we have

$$|\{x \in G : \pi_{-1}(G) = y\}| \sim \lambda$$

for all $y \in \pi_{-1}(G)$. Thus by Proposition 5.2 we have

$$|\pi_{-1}(G)| \lesssim \lambda^{-5/6}$$

and the claim follows.

One can modify Proposition 4.5 to show that one cannot improve the 1/6 in this estimate to beyond $0.20824\ldots$. Also, if one replaces T by the superficially similar operator

$$T(f,g,h)(x) := \int_{\mathbf{R}} f(x+t)g(x-\sqrt{2}t)h(t)\ dt,$$

(or indeed replacing $\sqrt{2}$ by any other irrational number) then there is no improvement over the trivial bound of λ^{-1}). Intuitively, this is because irrational projections cannot be used in arithmetic ways to control the size of rational projections; the actual construction however is a bit complex, and can be found in [3].

6. Exercises

- Q1. Show that if there is a map $f : X \to Y$ from one finite non-empty set X to another Y, then we have

  $$|\{(x,x') \in X : f(x) = f(x')\}| \geq \frac{|X|^2}{|Y|}.$$

  If $|X| \gg |Y|$, strengthen this to

  $$|\{(x,x') \in X : f(x) = f(x') \text{ and } x \neq x'\}| \gtrsim \frac{|X|^2}{|Y|}.$$

  We refer to this estimate as the *Cauchy-Schwarz estimate*.
- Q2. Let $f : X \to Y$ be a map from one finite non-empty set to another. Call an element $y \in Y$ *popular* if

  $$|\{x \in X : f(x) = y\}| \geq \frac{1}{2}\frac{|X|}{|Y|},$$

  i.e. y receives at least half of the expected number of points from X. Then show that

  $$|\{x \in X : f(x) \text{ is popular }\}| \geq \frac{1}{2}|X|.$$

  In other words, over half of the points in X have popular images under f. (We refer to this result as the *popularity argument*). Can you see how Q2 can be used to imply (a slightly weaker form of) the result in Q1?
- Q3. Let N be a large integer, and let A and B be two sets of cardinality N. Show that there exists a set $G \subset A \times B$ of cardinality $\sim N^2$ (i.e. a refinement of $A \times B$) which does not contain any Cartesian product of the form $A' \times B'$ for any refinements A', B' of N. (Hint: There are deterministic constructions, but a random construction is simplest: let G be constructed randomly by permitting each pair (a,b) to lie in G with probability 1/2, with all the probabilities being independent. Show that for each A', B', the

probability that G contains $A' \times B'$ is extremely small - of the order of e^{-CN^2} or so, while the number of sets A', B' are only e^{CN} or so).

- Q4*. Prove the (non-quantitative) modified Balog-Szemerédi theorem.
- Q5*. Let $n \geq 2$. A *Besicovitch set* E in $\mathbf{R}^n$ is a compact set which contains a unit line segment in every direction; thus for every unit vector ω there exists a line segment l_ω in E of length 1 and direction ω. For any $\delta > 0$, let $N_\delta(E)$ be the δ-neighborhood of E, and let $|N_\delta(E)|$ be the n-dimensional volume of this δ-neighborhood.

  (a) Prove the bound

$$|N_\delta(E)| \geq c(E,n)\delta^{(n-1)/2},$$

  where the constant $c(E,n)$ depends on E and n but not δ. (This implies that the *Minkowski dimension* of E is at least $(n+1)/2$).

  (b) Use Bourgain's argument to improve this bound to

$$|N_\delta(E)| \geq c(E,n)\delta^{12(n-1)/25}.$$

  (This implies that the Minkowski dimension of E is at least $(13n+12)/25$).

  (c) Use the Katz-Tao argument to improve this to

$$|N_\delta(E)| \geq c(E,n)\delta^{4(n-1)/7}.$$

  (This implies that the Minkowski dimension of E is at least $(4n+3)/7$).

Q6. Let $f_1 : V \to X_1$, $f_2 : V \to X_2$, $f_3 : V \to X_3$ be functions from one finite nonempty set to another. Show that

$$|\{(v_1,v_2,v_3,v_4) \in V^4 : f_1(v_1) = f_1(v_2); f_2(v_2) = f_2(v_3); f_3(v_3) = f_3(v_4)\}| \gtrsim \frac{|V|^4}{|X_1||X_2||X_3|}.$$

Hint: Apply the popularity argument in Q2 three times, each time refining V slightly. Note that the right-hand side is consistent with the probabilistic intuition that two random elements of X_1 should be equal with probability $1/|X_1|$, etc.

Q7. Use the Cartesian product trick (cf. the proof of Plünnecke's theorem in Week 1 notes) to show that the $\lesssim$'s in the Katz-Tao estimate can be replaced by $\leq$'s. More precisely, show that if $|\pi_0(G)|, |\pi_1(G)|, |\pi_2(G)|, |\pi_\infty(G)| \leq N$, then $|\pi_{-1}(G)| \leq N^{2-1/4}$.

Q8. Prove Stirling's formula $n! \sim n^{1/2}(n/e)^n$ by comparing the sum $\sum_{j=1}^n \log j$ with the integral $\int_1^n \log x \; dx$ (using for instance the trapezoid rule).

Q9. Prove Proposition 4.4, and then deduce Proposition 5.2.

References

[1] A. Balog, E. Szemerédi, *A statistical theorem of set addition*, Combinatorica, **14** (1994), 263–268.

[2] J. Bourgain, *On the dimension of Kakeya sets and related maximal inequalities*, GAFA **9** (1999), 256–282.

[3] M. Christ, *On certain elementary trilinear operators*, Math. Research Letters, **8** (2001), 43–56.

[4] T. Gowers, *A new proof of Szemerédi's theorem for arithmetic progressions of length four*, GAFA **8** (1998), 529–551.

[5] N. Katz, T. Tao, *Bounds on arithmetic progressions, and applications to the Kakeya conjecture*, Math. Res. Letters **6** (1999), 625–630.
[6] N. Katz, T. Tao, *New bounds for Kakeya problems*, to appear, J. d'Analyse Jerusalem.
[7] T. Tao, *Edinburgh lecture notes on the Kakeya problem*, www.math.ucla.edu/∼tao/preprints/Kakeya.html

Department of Mathematics, UCLA, Los Angeles CA 90095-1555

E-mail address: tao@math.ucla.edu

LECTURE NOTES 4 FOR 254A

TERENCE TAO

1. Arithmetic progressions

In this weeks notes we begin a new topic - that of finding arithmetic progressions inside a reasonably dense set of integers. The basic question is the following: let k be a small integer (in practice, we only take $k = 3$ or $k = 4$), and let N be a very large integer. If A is a subset of $\{1, \ldots, N\}$, how large does A have to be in order to ensure that it contains at least one proper arithmetic progression of length k? (In this set of notes all our arithmetic progressions shall be of rank 1; i.e. no generalized arithmetic progressions will be considered here). Of course we take $k \geq 3$ since the $k \leq 2$ cases are quite trivial.

One famous conjecture in this direction is Erdös's conjecture that the primes contain infinitely many arithmetic progressions of arbitrary length. More generally, he conjectured that any subset A of the natural numbers with $\sum_{n \in A} 1/n = \infty$ should contain infinitely many progressions of arbitrary length. In the above language of subsets of $\{1, \ldots, N\}$, this roughly corresponds to sets A of density $\sim 1/\log N$ (and in particular this is what happens to the primes, thanks to the prime number theorem). However, it is not yet known whether subsets of $\{1, \ldots, N\}$ of density $\sim 1/\log N$ must contain progressions of length k for large enough N, even when $k = 3$; the best result in this direction is by Bourgain, who showed that we have progressions of length 3 when the density is at least $\sim \sqrt{\log \log N}/\sqrt{\log N}$. We will prove this result in next week's results. (Remark: In the special case of the primes, Van der Corput in 1935 proved the infinitude of arithmetic progressions of length 3, but this is because the primes have better distributional properties than a generic set of density $1/logN$. The corresponding question about progressions of primes of length 4 is still open).

In this set of notes we prove a basic result in this area, known as *Szemerédi's theorem*:

Theorem 1.1. *Let $0 < \delta \leq 1$ and $k \geq 3$. Then there exists an $N(\delta, k)$ such that for any $N > N(\delta, k)$ and any subset A of $\{1, \ldots, N\}$ of cardinality $|A| \geq \delta N$, there exists an proper arithmetic progression in A of length k.*

Actually, we shall just prove this for $k = 3$ (this is due to Roth) and $k = 4$ (using an argument by Gowers). The $k = 4$ argument does eventually extend to general k, but the details are quite technical and will not be pursued here. The methods used will incorporate almost everything we have done in previous lecture notes.

2. Roth's argument

We now prove Theorem 1.1 when $k = 3$. The idea is to induct downwardly on δ. Of course when $\delta = 1$ the claim is trivial. Indeed, even when $\delta > 2/3$ the claim is easy, for the set $\{a \in \{1, \ldots, N-2\} : \text{ at least one of } a, a+1, a+2 \text{ is not in } A\}$ has cardinality at least $3(1-\delta)N$, and so as soon as N is large enough that $3(1-\delta)N < N - 2$ we can find $a \in \{1, \ldots, N\}$ such that $a, a+1, a+2$ are all in A.

Of course, δ is a continuous parameter, not a discrete one, and so the induction has to be slightly careful - in particular, the step size of the induction should not be allowed to vanish to zero unexpectedly. Let $P(\delta)$ be the property that Theorem 1.1 holds for δ (and $k = 3$), thus we already have $P(\delta)$ for $\delta > 2/3$. Our task is to show $P(\delta)$ is true for all $\delta > 0$. To show this, the main inductive step is the following:

Proposition 2.1. *For each $0 < \delta < 1$ there exists an $\varepsilon = \varepsilon(\delta) > 0$ such that ε depends continuously on δ, and such that $P(\delta + \varepsilon)$ implies $P(\delta)$.*

We now indicate why Proposition 2.1 implies Theorem 1.1 for $k = 3$. To see this, suppose for contradiction that $P(\delta)$ failed for some δ (and thus for all smaller δ). Let δ_* be the supremum of all the δ for which $P(\delta)$ failed, so that $P(\delta)$ is true for $\delta_* > \delta$ and false for $\delta_* < \delta$. But then by continuity and positivity of $\varepsilon(\delta)$ we can find a $\delta < \delta_*$ for which $\delta + \varepsilon(\delta) > \delta_*$, contradicting Proposition 2.1.

So it remains to prove Proposition 2.1. Fix δ: we let $\varepsilon = \varepsilon(\delta) > 0$ be chosen later. By hypothesis, there exists an integer M_0 such that for any $M > M_0$ subset B of $\{1, \ldots, M\}$ of density $|B| \geq (\delta + \varepsilon)M$, that B contains a orioer arithmetic progression of length 3. Clearly, this also implies that for any proper arithmetic progression P of cardinality $|P| \geq M_0$, and any subset $B \subseteq P$ with $|B| \geq (\delta+\varepsilon)|P|$, that B must also contain an arithmetic progression of length 3. Our task is thus to show that for N large enough (depending on M, δ, ε) and any A in $\mathbf{Z}/N\mathbf{Z}$ of density $|A| \geq \delta N$, that A also contains a progression of length 3.

The idea is to introduce a dichotomy - either A contains a "compact factor" - a subprogression of $\mathbf{Z}/N\mathbf{Z}$ of size at least M_0 where the density of A increases from δ to $\delta + \varepsilon$ - in which case we can apply the inductive hypothesis. Otherwise, A is "mixing" - it has uniformly density $\delta+O(\varepsilon)$ even when restricted to subprogressions. In such cases we can use some theory of exponential sums to show that the Fourier coefficients of A are small, which by some Fourier analysis will easily deduce the existence of many arithmetic progressions. The intuition here is that in the mixing case, A behaves like a random subset of $\{1, \ldots, N\}$ of density δ (i.e. it is as if each element of $\{1, \ldots, N\}$ would be in A with probability δ independently), and for random sets it is clear that one has $\sim \delta^3 N^2$ arithmetic progressions (since there are $\sim N^2$ progressions overall, and each progression has probability $\sim \delta^3$ of being contained in A).

(This terminology of compact factor and mixing is not completely accurate, but I use it by analogy with another, ergodic theory proof of Szemeredi's theorem given by Furstenburg).

If there exists an arithmetic progression $P \subseteq \{1, \dots, N\}$ of length at least M for which $|A \cap P| \geq (\delta + \varepsilon)|P|$, then by the inductive hypothesis we know that $A \cap P$, and hence A, contains a proper arithmetic progression of length 3. So, we may assume without loss of generality that

Assumption. We have

$$|A \cap P| \leq (\delta + \varepsilon)|P| \tag{1}$$

for every proper arithmetic progression $P \subseteq \{1, \dots, N\}$ of size $|P| \geq M$. In particular, we know that $|A| = (\delta + O(\varepsilon))N$.

This assumption means that A is very evenly distributed; for instance, if we partition $\{1, \dots, N\}$ into $\sim k$ progressions of length $\sim N/k$, then the δN elements of A are almost perfectly divided among those progressions. This will eventually show that the Fourier coefficients of χ_A are very small. But to do that, we must first place A inside a cyclic group, similarly to our treatment of Freiman's theorem.

Let p be a prime between N and $(1-\varepsilon)N$ (which we can always choose for N large enough, thanks to the prime number theorem), and let A' be the set $A \cap \{1, \dots, p\}$, thought of as a subset of $\mathbf{Z}/p\mathbf{Z}$; observe that $|A'| = (\delta + O(\varepsilon))p$. We know that $|A' \cap P| \leq (\delta + \varepsilon)|P|$ for all sufficiently long progressions P in $\{1, \dots, p\}$. But this is not the same as all sufficiently long progressions in $\mathbf{Z}/p\mathbf{Z}$, because of wraparound issues; to make this distinction let us call a *genuine arithmetic progression* in $\mathbf{Z}/p\mathbf{Z}$, one which is still a genuine progression when $\mathbf{Z}/p\mathbf{Z}$ is replaced with $\{1, \dots, p\}$ in the obvious manner.

Fortunately, our density bound for genuine arithmetic progressions can easily be extended to (sufficiently long) $\mathbf{Z}/p\mathbf{Z}$ arithmetic progressions, because very such progression can be decomposed into genuine arithmetic progressions:

Proposition 2.2. *There exists an $M_1 = M_1(M, \varepsilon)$ such that one has*

$$|A' \cap P| \leq (\delta + O(\varepsilon))|P|$$

for any proper arithmetic progression P in $\mathbf{Z}/p\mathbf{Z}$ (not necessarily genuine) of length $|P| \geq M_1$.

Proof Let $r \neq 0$ be the spacing of a proper progression P in $\mathbf{Z}/p\mathbf{Z}$. Let us first consider the case when $0 < r \ll p/M$. Then we can partition P into the disjoint union of *genuine* arithmetic progressions of the same step size r, simply by making a cut every time the progression wraps around the end of $\{1, \dots, p\}$. Except possibly for the first and last progression, each of the genuine progressions has length at least M, and so the density of A in those genuine progressions is at most $\delta + \varepsilon$ by assumption. Adding up all these density estimates for those genuine progressions of length at least M, we obtain the bound

$$|A' \cap P| \leq (\delta + \varepsilon)|P| + 2M.$$

A similar bound obtains when $0 < -r \ll p/M$.

Now suppose that r is arbitrary. Consider the first $O(M)$ multiples of r in $\mathbf{Z}/p\mathbf{Z}$. By the pigeonhole principle, two of them must be $\ll p/M$ apart, thus we can find

a non-zero $j = O(M)$ such that $|jr \mod p| \ll p/M$. We can then (if P is long enough) partition P into $|j|$ disjoint progressions, all of length $\sim |P|/|j|$, with step size $jr \mod p$. Applying the previous estimate to all of these progressions and adding up, we obtain

$$|A' \cap P| \leq (\delta + \varepsilon)|P| + 2|j|M \leq (\delta + \varepsilon)|P| + O(M^2),$$

and the claim follows if $|P|$ is large enough. ∎

Let $f(x)$ denote the function

$$f(x) := \chi_{A'}(x) - |A'|/p$$

on $\mathbf{Z}/p/Z$; in other words, f is the characteristic function of A', normalized to have mean zero. Then from the previous proposition we see that

$$\sum_{x \in P} f(x) \lesssim \varepsilon |P|$$

for any arithmetic progression P of length $\geq M_1$. In other words, f has mean at most $O(\varepsilon)$ on any sufficiently long progression. To put it another way, we have an upper bound

$$f * \frac{\chi_P}{|P|}(x) \lesssim \varepsilon$$

for any x (since we just replace P by $x - P$ in the previous. It would be nice if we also could have a lower bound too:

$$|f * \frac{\chi_P}{|P|}(x)| \lesssim \varepsilon$$

since this would be a very strong statement about f. However, we cannot quite deduce this from what we know; it asserts that on every sub-progression, that not only is $|A' \cap P|$ have density at most $\delta + O(\varepsilon)$, it also has density at *least* $\delta - O(\varepsilon)$. Nevertheless, despite the fact that we cannot obtain this bound directly, we can obtain a usable substitute, namely

$$\|f * \frac{\chi_P}{|P|}(x)\|_1 \lesssim \varepsilon;$$

in other words, we don't have *uniform* control on how negative $f * \chi_P|P|$ can be, but we do have L^1 control. The reason for this is very simple: since f has mean zero, we know that $f * \frac{\chi_P}{|P|}$ also has mean zero, and in particular the positive and negative parts have equal L^1 norm. But the positive part is uniformly $O(\varepsilon)$, and hence $O(\varepsilon)$ in L^1 norm as well, thus the negative part is also.

This has some consequences for the Fourier transform of f:

Lemma 2.3. *For any $\xi \in \mathbf{Z}/p\mathbf{Z}$, we have $|\hat{f}(\xi)| \lesssim \varepsilon$.*

This statement is sometimes referred to as *linear uniformity* of f; $f(x)$ is somewhat orthogonal to all functions of the form $e^{2\pi i x\xi/p}$, which have phases linear in x. For comparision, if A is chosen randomly, then one almost surely has $|\hat{f}(\xi)| \lesssim p^{-1/2+}$ (see Q1). Thus this lemma asserts that sets which do not concentrate on arithmetic progressions have linear uniformity; for a partial converse, see Q3.

If we pretend that f is itself a phase function, $f(x) = e^{2\pi i\phi(x)}$, then the above lemma asserts that ϕ does not behave like a linear phase, i.e. $\phi(x) \not\approx x\xi/p + c$ for any ξ and constant c, where we are vague as to what the symbol $\approx$ means.

Proof The case $\xi = 0$ is trivial, so let us assume $\xi \neq 0$. Since $\mathbf{Z}/p\mathbf{Z}$ is a field, we can find $r \neq 0$ such that $r\xi = 1 \mod p$. Let P be the progression $\{jr : 0 \leq j < M_1\}$; by the preceding discussion we have

$$|\hat{f}(\xi)|\frac{p}{|P|}|\hat{\chi}_P(\xi)| \lesssim \varepsilon.$$

On the other hand, by construction we easily see that

$$\hat{\chi}_P(\xi) = \frac{1}{p}\sum_{j=0}^{M_1-1} e^{2\pi ij/p} \sim |P|/p$$

if p is sufficiently large depending on M_1. The claim follows. ■

Now armed with this Fourier information, we can finish the proof. We introduce the trilinear form $T(f, g, h)$ for three functions on $\mathbf{Z}/p\mathbf{Z}$ defined by

$$T(f, g, h) := \int_{\mathbf{Z}/p\mathbf{Z}}\int_{\mathbf{Z}/p\mathbf{Z}} f(x)g(x + r)h(x + 2r)\, dxdr.$$

Observe that the number of proper (but non-genuine) arithmetic progressions in A' of length 3 is precisely

$$p^2T(\chi_{A'}, \chi_{A'}, \chi_{A'}) - |A'|$$

(the $-|A'|$ being there just to remove the trivial progressions of step 0; we count $a, a+r, a+2r$ as being a different progression from $a+2r, a+r, a$). So we can find (non-genuine) progressions as soon as we prove an estimate of the form

$$T(\chi_{A'}, \chi_{A'}, \chi_{A'}) \gg \delta/p. \tag{2}$$

Later on we will see how we need to modify this to obtain genuine progressions.

The basic estimate we need here is that $T(f, g, h)$ is small if one of f, g, h has small Fourier coefficients:

Proposition 2.4. *We have*

$$|T(f, g, h)| \leq \|f\|_2\|g\|_2\|\hat{h}\|_\infty$$

and similarly for permutations.

To get some intuition for this proposition, suppose that $f(x) = e^{2\pi i\phi(x)}$, $g(x) = e^{2\pi i\psi(x)}$, and $h(x) = e^{2\pi i\eta(x)}$. Then the quantity $T(f, g, h)$ is large if $\phi(a) + \psi(a + r) + \eta(a + 2r)$ behaves like a constant, i.e.

$$\phi(a) + \psi(a + r) + \eta(a + 2r) \approx c \text{ for most } a, r$$

where we are vague about what $\approx$ and "for most" mean. This Proposition says that this type of phase cancellation can happen only when ϕ, ψ, and η all behave somewhat linearly - each of f, g, and h need to have one large Fourier coefficient in order to obtain the above type of estimate. In fact, the proof will show that the only way $T(f, g, h)$ can be large is if there exists a ξ such that $\phi(a) \approx a\xi/p + c_1$, $\psi(a) \approx -2a\xi/p + c_2$, and $\eta(a) \approx a\xi/p + c_3$. (See Q6.)

Proof We use the Fourier inversion formula to write

$$f(x) = \sum_{\xi_1 \in \mathbf{Z}/p\mathbf{Z}} \hat{f}(\xi_1) e^{2\pi i x \xi_1/p}$$

$$g(x+r) = \sum_{\xi_2 \in \mathbf{Z}/p\mathbf{Z}} \hat{g}(\xi_2) e^{2\pi i (x+r) \xi_2/p}$$

$$h(x+2r) = \sum_{\xi_3 \in \mathbf{Z}/p\mathbf{Z}} \hat{h}(\xi_3) e^{2\pi i (x+2r) \xi_3/p}$$

which allows us to write $T(f,g,h)$ as

$$\sum_{\xi_1,\xi_2,\xi_3 \in \mathbf{Z}/p\mathbf{Z}} \hat{f}(\xi_1)\hat{g}(\xi_2)\hat{h}(\xi_3) \int_{\mathbf{Z}/p\mathbf{Z}} \int_{\mathbf{Z}/p\mathbf{Z}} e^{2\pi i (x\xi_1 + (x+r)\xi_2 + (x+2r)\xi_3)/p}\, dx dr.$$

The inner r integral vanishes unless $\xi_2 + 2\xi_3 = 0 \mod p$; the inner x integral vanishes unless $\xi_1 + \xi_2 + \xi_3 = 0 \mod p$. Thus the only non-zero contribution arises when $\xi_3 = \xi_1$ and $\xi_2 = -2\xi_1$, thus

$$T(f,g,h) = \sum_{\xi \in \mathbf{Z}/p\mathbf{Z}} \hat{f}(\xi)\hat{g}(-2\xi)\hat{h}(\xi).$$

The claim then follows by taking $\hat{h}$ out in supremum, and applying Cauchy-Schwarz and Plancherel to what remains. ■

To apply this proposition, we write $\chi_{A'} = |A'|/p + f$, and split $T(\chi_{A'}, \chi_{A'}, \chi_{A'})$ accordingly into 8 terms. 7 of these can be estimated using the above proposition and lemma as $O(\delta^2\varepsilon) + O(\delta\varepsilon^2) + O(\varepsilon^3)$, while the 8th is $T(|A'|/p, |A'|/p, |A'|/p) = |A'|^3/p^3$, which is roughly δ^3. Thus if we choose $\varepsilon \ll \delta$ we obtain (2) as desired. This gives us a non-genuine arithmetic progression of length 3; in fact, it gives us $\sim \delta^3 p^2$ such progressions, which is consistent with the random case.

To upgrade this to a genuine arithmetic progression, we use a cutoff function. Let ψ be a bump function adapted to $[p/3, 2p/3]$. We will show that

$$T(\psi\chi_{A'}, \psi\chi_{A'}, \psi\chi_{A'}) \gg \delta/p;$$

this will show that $A' \cap [p/3, 2p/3]$ contains an arithmetic progression, which must then necessarily be genuine. We split $\psi\chi_{A'} = \psi|A'|/p + \psi f$. Observing that the Fourier transform of ψ is bounded and decays rapidly away from the origin, we see that ψf is also linearly uniform (its Fourier coefficients are also $O(\varepsilon)$). Also, $T(\psi,\psi,\psi) \sim c$ for some $c > 0$ independent of p, depending only on ψ. The claim follows. This concludes the proof of Roth's theorem. □

We now remark on what the exact dependence of $N(\delta)$ is on δ, as predicted by the above proof. Let $N(\delta)$ be the smallest number for which the conclusion of Szemeredi's theorem (with $k = 3$) holds for that valaue of δ. An inspection of the above argument shows that we can take $\varepsilon = c\delta$ for some absolute constants c, C, and can take $N_0 = CM^C\delta^{-C}\varepsilon^{-C}$ for some other absolute constants C. If we assume that $M \geq \delta^{-1}$ (which is quite reasonable, given it is difficult to find many

sets in $\{1,\dots,M\}$ of density δ if $M < \delta^{-1}$, we can thus take $N_0 = CM^C$. Thus the above argument gives a recursive bound of the form

$$N(\delta) \leq CN(\delta + c\delta)^C.$$

Also, we have $N(\delta) < 10$ when $\delta > 0.9$ (for instance), by the argument given for $\delta > 2/3$. Putting these two together, we can obtain a bound of the form

$$N(\delta) \leq \exp(\exp(C/\delta))$$

for some other absolute constant C; see Q2. Inverting this, we thus see that for any large N that every subset of $\{1,\dots,N\}$ of density at least $C/\log\log N$ will contain a non-trivial arithmetic progression of length 3. Note that this is far too weak to address the question of what happens for the primes, which have a density of about $1/\log N$; in the next week's notes we will give some progress toward closing this gap.

Call a set $A' \subseteq \mathbf{Z}/p\mathbf{Z}$ *linearly uniform* if the function $f = \chi_{A'} - |A'|/p$ has small Fourier coefficients, say $\|\hat{f}\|_\infty \lesssim \varepsilon$. The above argument says that if A' is linearly uniform and has density $\sim \delta$, then A' contains $\sim \delta^3 p^2$ arithmetic progressions; roughly speaking, this is because the linear uniformity guarantees that the events $a \in A'$, $a + r \in A'$, and $a + 2r \in A'$ behave more or less "independently". Note that if A' had a high concentration in an arithmetic progression then we would not have this independence.

3. Gowers proof of Szemeredi's theorem when $k = 4$.

Roth's argument above appeared all the way back in 1953. One might think that it would be an easy manner to extend that $k = 3$ argument to, say, $k = 4$, but the first such proof had to wait until 1969 by Szemeredi, and required some very deep combinatorial results (in particular, Szemeredi's uniformity lemma). Only in 1998 did Gowers finally extend the Fourier-analytic stile of Roth's argument to $k = 4$, and then in 2001 to all k.

The main difficulty when $k = 4$ is the following. We can once again give the dichotomy between "compact factor" (where $A \cap P$ is denser than A) or "mixing" (when A is uniformly of density δ). The compact factor case is again easy by induction; the problem is how to use the mixing hypothesis. In theory, in this case one expects about $\delta^4 p^2$ arithmetic progressions of length 4 in A', which is of course much larger than the $\sim \delta p$ trivial arithmetic progressions when p is large. Unfortunately, in this case the Fourier argument is not as useful. If we introduce the quadrilinear form

$$Q(f,g,h,k) := \int_{\mathbf{Z}/p\mathbf{Z}} \int_{\mathbf{Z}/p\mathbf{Z}} f(x)g(x+r)h(x+2r)k(x+3r)\,dxdr$$

then the number of $\mathbf{Z}/p\mathbf{Z}$ arithmetic progressions of length 4 in A' is

$$p^2 Q(\chi_{A'}, \chi_{A'}, \chi_{A'}, \chi_{A'}) - |A'|.$$

The main term is

$$p^2 Q(|A'|/p, |A'|/p, |A'|/p, |A'|/p) \sim \delta^4 p^2$$

as expected, and most of the other terms can be treated using Proposition 2.4 (because if there is at least one constant factor $|A'|/p$ then one can collapse the quadrilinear form to a trilinear one, which can then be treated using some variant of Proposition 2.4). However, the term $p^2Q(f,f,f,f)$ is more resistant to analysis. A Fourier argument shows that $p^2Q(f,f,f,f)$ is equal to

$$p^2 \sum_{\xi_1,\xi_2,\xi_3,\xi_4 \in \mathbf{Z}/p\mathbf{Z}: \xi_1+\xi_2+\xi_3+\xi_4=\xi_2+2\xi_3+3\xi_4=0} \hat{f}(\xi_1)\hat{f}(\xi_2)\hat{f}(\xi_3)\hat{f}(\xi_4),$$

or equivalently

$$p^2 \sum_{\xi_3,\xi_4 \in \mathbf{Z}/p\mathbf{Z}} \hat{f}(\xi_3+2\xi_4)\hat{f}(-2\xi_3-3\xi_4)\hat{f}(\xi_3)\hat{f}(\xi_4).$$

Using the bounds $\|\hat{f}\|_{l^2} \sim \delta^{1/2}$ and $\|\hat{f}\|_\infty \lesssim \varepsilon$, the best bound that one can obtain on expression term is $O(\delta^2 p^2)$ (how? Use Cauchy-Schwarz on $\mathbf{Z}/p\mathbf{Z} \times \mathbf{Z}/p\mathbf{Z}$, splitting the four factors into pairs of two), which dominates the main term.

The basic problem is that linear uniformity is not strong enough to control Q. However, if we strengthen linear uniformity to something called *quadratic uniformity*, then we will be in good shape. At this point it is not quite clear what quadratic uniformity should mean; in fact, the difficulty in even coming up with the correct notion here is one of the key stumbling blocks in the argument. (An obvious guess for quadratic uniformity is that the function f should be somewhat orthogonal, not just to linear phase functions, but also to quadratic phase functions as well; it turns out that this WILL imply the notion of quadratic uniformity which we will use, but only if we localize it to relatively short sub-progressions, and the implication is highly nontrivial, using both Freiman's theorem and the Balog-Szemeredi theorem).

Let us first make some notation. Let us use $o(1)$ to denote any quantity depending on ε which goes to zero as ε goes to 0, so in particular we can make $o(1)$ much smaller than δ by choosing ε sufficiently small. We say that a statement $P(k)$ is true for *most* $k \in \mathbf{Z}/p\mathbf{Z}$ if the number of exceptional k for which $P(k)$ false is only $o(p)$.

We say that a function f on $\mathbf{Z}/p\mathbf{Z}$ is *linearly uniform* if $\|f\|_\infty \lesssim 1$, but $\|\hat{f}\|_\infty = o(1)$; as we saw earlier, this basically meant that f did not concentrate on any arithmetic progression. We say that f is *quadratically uniform* if $\|f\|_\infty \lesssim 1$, and $f(x+k)\overline{f}(x)$ is linearly uniform for most $k \in \mathbf{Z}/p\mathbf{Z}$. In other words (see Q5), f is quadratically uniform if

$$\sum_{k \in \mathbf{Z}/p\mathbf{Z}} \|\widehat{f(\cdot+k)\overline{f}}\|_\infty = o(p).$$

To understand this concept, let us first suppose that f is a phase function, i.e. $f(x) = e^{2\pi i \phi(x)}$ for some real-valued ϕ; thus the condition $\|f\|_\infty \lesssim 1$ is automatic. Saying that f is linearly uniform is the same as saying that the function $e^{2\pi i(\phi(x)-x\xi/p)}$ has mean $o(1)$ for all ξ; thus the function $\phi(x)$ is never "aligned" with any linear function $x\xi/p$. Saying that f is quadratically uniform is the same as saying that for most k, the function $\phi(x+k)-\phi(x)$ is linearly uniform; this is somewhat like saying that the "derivative" of ϕ never "behaves linearly", and is

thus a statement that ϕ "never behaves quadratically". For instance, the function $f(x) = e^{2\pi i x^2/p}$ can easily be shown to be linearly uniform but not quadratically uniform (see Q4, Q7(a)).

Let's say that a subset A' of $\mathbf{Z}/p\mathbf{Z}$ is quadratically uniform if the mean-zero function $F = \chi_{A'} - |A'|/p$ is quadratically uniform. We now claim that if A' is quadratically uniform and has density $\sim \delta$, then it will contain $\sim \delta^4 p^2$ arithmetic progressions (i.e. the number predicted by a random argument). Roughly speaking, the reason why we need quadratic uniformity is that this is what it takes to make the events $a \in A$, $a+r \in A$, $a+2r \in A$, and $a+3r \in A$ behave "independently". For instance, take A to be the set

$$A = \{x \in \mathbf{Z}/p\mathbf{Z} : e^{2\pi i x^2/p} = 1 + O(\delta)\}.$$

Then A has density $\sim \delta$, and it turns out to be linearly uniform (for much the same reason that $e^{2\pi i x^2/p}$ is linearly uniform), but is not quadratically uniform; if $a, a+r, a+2r$ all lie in A, then there is a very high probability that $a+3r$ also lies in A, basically because of the identity

$$(a+3r)^2 = a^2 - 3(a+r)^2 + 3(a+2r)^2.$$

Now we prove that quadratic uniformity implies $\sim \delta^4 p^2$ arithmetic progressions of length 4:

Lemma 3.1. *Suppose $\chi_{A'} - |A'|/p$ is both linearly uniform and quadratically uniform. Then A' contains $\sim \delta^4 p^2$ $\mathbf{Z}/p\mathbf{Z}$-progressions of length 4.*

Proof We first recall the analysis from Roth's argument: if f, g, h are bounded functions, and at least one of f, g, h are linearly uniform, then

$$\int_{\mathbf{Z}/p\mathbf{Z}} \int_{\mathbf{Z}/p\mathbf{Z}} f(a)g(a+r)h(a+2r)\, da dr = T(f,g,h) = o(1). \tag{3}$$

Also, since F is quadratically uniform, we know the function $F(x+k)F(x)$ linearly uniform for most k. Applying (3) with $f(x) = F(x+k)F(x)$, $g(x) = F(x+2k)F(x)$, $h(x) = F(x+3k)F(x)$, we thus have

$$\int_{\mathbf{Z}/p\mathbf{Z}} \int_{\mathbf{Z}/p\mathbf{Z}} F(a)F(a+k)F(a+r)F(a+2k+r)F(a+2r)F(a+3k+2r)\, da dr = o(1)$$

for most k. For those exceptional values of k, this integral is still bounded, of course. Thus we can integrate over all k and obtain

$$\int_{\mathbf{Z}/p\mathbf{Z}} \int_{\mathbf{Z}/p\mathbf{Z}} \int_{\mathbf{Z}/p\mathbf{Z}} F(a)F(a+k)F(a+r)F(a+2k+r)F(a+2r)F(a+3k+2r)\, da dr dk = o(1).$$

Replacing k by the variable $s := r+k$, and a by the variable $b := a-r$, this becomes

$$\int_{\mathbf{Z}/p\mathbf{Z}} \int_{\mathbf{Z}/p\mathbf{Z}} \int_{\mathbf{Z}/p\mathbf{Z}} F(b+r)F(b+s)F(b+2r)F(b+2s)F(b+3r)F(b+3r)\, db dr ds = o(1),$$

thus

$$\int_{\mathbf{Z}/p\mathbf{Z}} |\int_{\mathbf{Z}/p\mathbf{Z}} F(b+r)F(b+2r)F(b+3r)\, dr|^2 db = o(1),$$

In particular, by Cauchy-Schwarz we have $Q(F,F,F,F) = o(1)$, which implies that $Q(\chi_{A'}, \chi_{A'}, \chi_{A'}, \chi_{A'}) \sim \delta^4$ as desired (since we can split $\chi_{A'} = |A'|/p + F$ and use Proposition 2.4 and its variants (using the linear uniformity of A') to control crossterms. ∎

In light of this lemma, and the linear uniformity already obtained from Roth's argument, it thus suffices to show

Proposition 3.2. *Let A' be a subset of $\mathbf{Z}/p\mathbf{Z}$ of density $\delta + O(\varepsilon)$, which obeys the conclusions of Proposition 2.2 (i.e. A' has uniform density on reasonably long arithmetic progressions). Then, if p is large enough, then A' is quadratically uniform.*

To prove this proposition, we will argue by contradiction. Let $\sigma = o(1)$ be a quantity to be chosen later. Suppose that A' is not quadratically uniform; this means that for at least σp values of k, we have a frequency $\xi(k)$ such that

$$|\widehat{\overline{F(x)}F(x+k)}(\xi(k))| \geq \sigma. \tag{4}$$

Let B be the set of all such k, and let $\Omega \in \mathbf{Z}/p\mathbf{Z} \times \mathbf{Z}/p\mathbf{Z}$ be the graph

$$\Omega := \{(k, \xi(k))\}.$$

One can think of Ω as the set of "resonant frequencies" of the two-variable function $\overline{F(x)}F(x+k)$ (if we take Fourier transforms in x but not in k).

We now perform a number of somewhat mysterious Fourier manipulations to eventually discover some additive information about Ω. Before we proceed, let us give an intuitive argument. Suppose that $F(x) = e^{2\pi i \phi(x)}$. Then the estimate (4) says, essentially, that $\phi(x+k) - \phi(x) \approx \xi(k)x/p + c$ for "most" x and k. It turns out the only way this can work out is if ξ is roughly additive, i.e. Ω is in some sense somewhat closed under addition. Indeed, we heuristically expect ϕ to be linear, and $\xi(k)$ should be something like the slope of ϕ, multiplied by k. (See Q7(b)).

For all $k \in B$, we expand (4), using the fact that F is real, to obtain

$$|\int F(x)F(x+k)e^{2\pi i x\xi(k)/p}\, dx| \geq \sigma.$$

We square this to get

$$\int\int F(x)F(y)F(x+k)F(y+k)e^{2\pi i(x-y)\xi(k)/p}\, dxdy \geq \sigma^2.$$

Writing $y = x+u$, and writing $G_u(x) := F(x)F(x+u)$, this becomes

$$\int\int G_u(x)G_u(x+k)e^{-2\pi i u\xi(k)/p}\, dxdu \geq \sigma^2 p.$$

Integrating this over all $k \in B$, we obtain

$$\int (\int\int G_u(x)G_u(x+k)H_u(k)\, dxdk)\, du \geq \sigma^3$$

where $H_u(k) := \chi_B(k) e^{-2\pi i u \xi(k)/p}$. But by Proposition 2.4, we have that

$$|\int \int G_u(x) G_u(x+k) H_u(k)\ dxdk| \lesssim \|\hat{H}_u\|_\infty$$

since G_u is bounded. Thus we have

$$\int \|\hat{H}_u\|_\infty du \geq \sigma^3.$$

Thus we must have $\|\hat{H}_u\|_\infty \gtrsim \sigma^C$ for at least $O(\sigma^C p)$ values of u. For all such u, we have $\sum_{\eta \in \mathbf{Z}/p\mathbf{Z}} |\hat{H}_u(\eta)|^4 \gtrsim \sigma^C$; thus we have

$$\int \sum_{\eta \in \mathbf{Z}/p\mathbf{Z}} |\hat{H}_u(\eta)|^4 du \geq \sigma^C.$$

Note however that

$$\hat{H}_u(\eta) = \int_B e^{-2\pi i(\eta k + u\xi(k))}\ dk = \int_\Omega e^{-2\pi i (k,\xi)\cdot(\eta,u)}\ dkd\xi = \hat{\chi}_\Omega(\eta, u)$$

where we give Ω the product of normalized measure dk and counting measure $d\xi$. Thus by Plancherel

$$\int \sum_{\eta \in \mathbf{Z}/p\mathbf{Z}} |\hat{H}_u(\eta)|^4 du = \|\chi_\Omega * \chi_\Omega\|_2^2 = p^{-3} |\{a,b,c,d \in \Omega : a+b=c+d\}|.$$

Thus we have

$$|\{a,b,c,d \in \Omega : a+b=c+d\}| \gtrsim \sigma^C p^3 \gtrsim \sigma^C |\Omega|^3.$$

In other words, Ω is partly closed under addition, in the language of last week's notes. Thus by the Balog-Szemeredi theorem, there exists $\Omega' \subseteq \Omega$ with $|\Omega'| \gtrsim \sigma^C |\Omega|$ such that Ω' is totally closed under addition, $|\Omega' + \Omega'| \lesssim \sigma^C |\Omega'|$. By Freiman's theorem (adapted to $\mathbf{Z}/p\mathbf{Z} \times \mathbf{Z}/p\mathbf{Z}$; this is not a torsion-free abelian group, but the argument still essentially extends), Ω' is thus contained inside a proper generalized arithmetic progression P of rank $C(\sigma)$ and size at most $C(\sigma)p$, thus in particular $|\Omega'| \gtrsim C(\sigma)|P|$. But this progression P can be partitioned into one-dimensional progressions of size at least $\gtrsim p^{c(\sigma)}$; thus by the pigeonhole principle we can thus find a one-dimensional arithmetic progression Q of size at least $\gtrsim p^{c(\sigma)}$ such that

$$|\Omega \cap Q| \geq |\Omega' \cap Q| \gtrsim c(\sigma)|Q|.$$

This progression Q takes the form

$$Q := \{(k, ak+b : k \in R\}$$

where R is an arithmetic progression in $\mathbf{Z}/p\mathbf{Z}$ of size $\gtrsim p^{c(\sigma)}$, and $a, b \in \mathbf{Z}/p\mathbf{Z}$. Thus we have

$$|\widehat{\overline{F(x)} F(x+k)}(ak+b)| \geq \sigma$$

for at least $c(\sigma)|R|$ values of $k \in \mathbf{R}$. In particular we have

$$\sum_{k \in R} |\widehat{\overline{F(x)} F(x+k)}(ak+b)|^2 \geq c(\sigma)|R|.$$

To obtain a contradiction, it will thus suffice to show the following "localized, linearized" form of quadratic uniformity:

Proposition 3.3. *Let A' be a subset of $\mathbf{Z}/p\mathbf{Z}$ of density $\delta + O(\varepsilon)$, which obeys the conclusions of Proposition 2.2 (i.e. A' has uniform density on reasonably long arithmetic progressions). Then, if p is large enough, then*

$$\sum_{k \in R} |\widehat{\overline{F(x)}F(x+k)}(ak+b)|^2 = o(|R|)$$

for all $a, b \in \mathbf{Z}/p\mathbf{Z}$ and all arithmetic progressions R of size at least $|R| \geq M_2 = M_2(M_1, \delta, \varepsilon)$.

We now prove this in the next section to finish the proof of Gowers' theorem.

4. Linearized quadratic uniformity

Let R be a reasonably long arithmetic progression, and let a, b be in $\mathbf{Z}/p\mathbf{Z}$. We now have to bound the expression

$$\sum_{k \in R} |\widehat{\overline{F(x)}F(x+k)}(ak+b)|^2. \tag{5}$$

Since

$$\widehat{\overline{F(x)}F(x+k)}(ak+b) = \int \overline{F(x)}F(x+k)e^{-2\pi i(ak+b)x/p}\, dx$$

we can write

$$|\widehat{\overline{F(x)}F(x+k)}(ak+b)|^2 = \int\int \overline{F(x)}F(y)F(x+k)\overline{F(y+k)}e^{-2\pi i(ak+b)(x-y)/p}\, dxdy.$$

Making the substitution $y = x + u$, we can thus rewrite (5) as

$$\int \overline{F(x)} \sum_{k \in R} \int F(x+u)F(x+k)\overline{F(x+u+k)}e^{2\pi i(ak+b)u/p}\, du\, dx.$$

We need to show that this is $o(|R|)$. It suffices to prove this uniformly in x, i.e. it suffices to show that

$$|\sum_{k \in R} \int F(x+u)F(x+k)\overline{F(x+u+k)}e^{2\pi i(ak+b)u/p}\, du| = o(|R|).$$

By translation invariance we may take $x = 0$:

$$|\sum_{k \in R} \int F(u)F(k)\overline{F(u+k)}e^{2\pi i(ak+b)u/p}\, du| = o(|R|).$$

Now for the key trick. The phase function $(ak+b)u$ is bilinear in k and u. Thus we can write it (mod p) as

$$(ak+b)u = Q_1(u) + Q_2(k) + Q_3(u+k)$$

for some (inhomogeneous) quadratic polynomials Q_1, Q_2, Q_3 with integer coefficients. Thus it will suffice to show

$$|\int\int F_1(u)F_2(k)F_3(u+k)\, dudk| = o(|R|/p), \tag{6}$$

where

$$F_1(u) := F(u)e^{2\pi i Q_1(u)/p}$$
$$F_2(k) := F(k)e^{2\pi i Q_2(k)/p}\chi_R(k)$$
$$F_3(x) := \overline{F(x)}e^{2\pi i Q_3(x)/p}.$$

At this point we shall oversimplify the argument for sake of exposition, and clean up the oversimplification later.

The left-hand side of (6) is $T(F_1, F_2, F_3)$, and F_1 and F_3 are clearly bounded in L^2. Thus by Proposition 2.4 it thus suffices to show that

$$\|\hat{F}_2\|_\infty = o(|R|/p),$$

or in other words

$$\sum_{k\in R} F(k)e^{2\pi i Q_2(k)/p}e^{-2\pi i\xi k/p} = o(|R|).$$

The linear phase $-\xi k$ can be absorbed into the arbitrary quadratic phase Q_2, and so it suffices to show

Proposition 4.1. *For any progression R of length at least M_2, and any quadratic polynomial $Q(k)$ with integer coefficients, we have*

$$\sum_{k\in R} F(k)e^{2\pi i Q(k)/p} = o(|R|).$$

This is a much stronger version of Lemma 2.3, which dealt with linear phase functions instead of quadratic ones, and worked on all of $\mathbf{Z}/p\mathbf{Z}$ instead of localizing to the relatively short interval R. As it turns out, this Proposition is not quite true (and so this proof is not quite rigorous), however we will indicate the fix needed for this proposition a bit later.

Proof (Informal) As mentioned above, we cannot quite prove this proposition using only the hypothesis we have, which is the upper bound

$$F * \frac{\chi_P}{|P|}(x) \lesssim \varepsilon$$

for all x and all progressions P of length at least M_1. However, we can derive it if we assume that we can have the lower bound as well as the upper bound:

$$|F * \frac{\chi_P}{|P|}(x)| \lesssim \varepsilon = o(1). \tag{7}$$

As mentioned earlier, we do not actually have this estimate, except in an L^1-averaged sense; so as it turns out we only get an L^1-averaged version of this proposition, but this turns out to be sufficient to close the argument. More on this later.

Anyway, suppose we have (7). The idea is to split R into sub-progressions, on which $Q(k)/p$ is essentially constant (mod 1), more precisely

Lemma 4.2. *If M_2 is sufficiently large, then one can partition R as the disjoint union of progressions $P_1, \ldots, P_K$, such that each progression P_j is of length between M_1 and $2M_1$, and such that $e^{2\pi i Q(k)/p}$ is within $o(1)$ of a constant c_j on each progression P_j.*

Assuming this lemma and (7), we have

$$\sum_{k\in P_j} F(k)e^{2\pi iQ(k)/p} = c\sum_{k\in P_j} F(k) + o(|P_j|) = co(|P_j|) + o(|P_j|) = o(|P_j|)$$

and the claim follows by summing over j.

It remains to prove the lemma. By applying a linear transformation we may assume that $R = \{1, 2, \dots, r\}$ for some $r \geq M_2$.

Now we write $Q(k) = ak^2 + bk + c$. We need to partition R into progressions P_j for which Q/p is essentially constant mod p, or more precisely

$$\|\frac{ak^2 + bk + c - C_j}{p}\| = o(1)$$

for all $k \in P_j$, where $\|x\|$ denotes the distance from x to the nearest integer. Another way of saying this is that

$$\|\frac{a(k^2 - (k')^2) + b(k - k')}{p}\| = o(1)$$

for all $k, k' \in P_j$.

We first observe from Weyl's theorem on the distribution of squares[1] that we can find $1 \leq d \leq C(M_1, \varepsilon)$ such that

$$\|\frac{ad^2}{p}\| = o(M_1^{-100}).$$

We can then partition R into progressions Q_i of length $\sim M_1^{10}$ and spacing d. Let $Q_i = \{k_0 + jd : j = O(M_1^{10})\}$ be one of these progressions. Then if $k = k_0 + jd$ and $k' = k_0 + j'd$, ten

$$\|\frac{a(k^2 - (k')^2) + b(k - k')}{p}\| = \|\frac{a(j^2 - (j')^2)d^2}{p} + \frac{(2ak_0 + bd)(j - j')}{p}\| = \|\frac{(2ak_0 + bd)(j - j')}{p}\| + o(1)$$

by choice of d. So now we use Kronecker's theorem[2] to find a number $J = O(M_1^5)$ such that

$$\|\frac{(2ak_0 + bd)}{p}J\| = O(M_1^{-5}).$$

If we then divide Q_i into progressions P_s of length between M_1 and M_1, and of spacing Jd, we thus have

$$\|\frac{(2ak_0 + bd)(j - j')}{p}\| = O(M_1^{-4}) = o(1)$$

[1]This theorem states that the integer parts of $\|\alpha n^2)_{n=1}^{\infty}$ are uniformly distributed if α is irrational. There is a quantitative version of this which says that these integer parts are "approximately uniformly distributed" if α is not too close to a rational of small denominator. On the other hand, if α is close to a rational of small denominator, then αn^2 is close to an integer quite often. We omit the precise treatment of Weyl's theorem here due to lack of time; however Gowers [2] has shown that given any $a \in \mathbf{Z}/p\mathbf{Z}$ and any integer $D > 1$, we can find $1 \leq d \leq D$ such that $\|\frac{ad^2}{p}\| \lesssim D^{-1/8}$.

[2]This is the same as Weyl's theorem, but for αn instead of αn^2: specifically, given any α and any N, there exists $1 \leq n \leq N$ such that $\|\alpha n\| \lesssim 1/N$. Actually the argument here is quite simple, based on the pigeonhole principle; at least two of the multiples of α must be within $1/N$ of each other, so the difference will be within $O(1/N)$ of the nearest integer.

for all $k_0 + jd, k_0 + j'd \in P_s$ (so that $(j - j')/J = O(M_1)$) and the claim follows. ■

The problem with this Proposition, of course, is that assumes (7), when in fact we only have this bound on the average in an L^1 sense. If one works one's way through the above argument, this implies that we have the following weaker averaged version of the Proposition: if R is as above, and for each $y \in \mathbf{Z}/p\mathbf{Z}$ we assign a quadratic polynomial Q_y, then we have

$$\int |\sum_{k \in R+y} F(k) e^{2\pi i Q_y(k)/p}|\, dy = o(|R|).$$

It turns out that this is still enough to prove (6). The trick is to partition the u variable into translates of R, apply the preceding argument to each translate, and then average up. We omit the details.

5. Exercises

- Q1. Let p be a large prime and $0 < \delta < 1$. Let A be a random subset of $\mathbf{Z}/p\mathbf{Z}$, formed by letting each element of $\mathbf{Z}/p\mathbf{Z}$ lie in A with an independent probability of δ.

  (a) Show that the expected number of proper arithmetic progressions in A is exactly $\delta^3 p(p-1)$.

  (b) Let $f := \chi_A - \delta$. Show that for any frequency ξ, we have

  $$\mathbf{E}(\exp(\sqrt{p}\mathrm{Re}\hat{f}(\xi)) \lesssim 1$$

  where the implicit constant is allowed to depend on δ. Hint: write

  $$\sqrt{p}\mathrm{Re}\hat{f}(\xi) = \sum_{j=1}^{p} p^{-1/2} \cos(2\pi j\xi/p) X_j$$

  where the X_j are iid random variables which equal $1-\delta$ with probability δ, and $-\delta$ with probability $1-\delta$ (so in particular X_j, and all products of distinct X_j, have mean zero). Now argue as in the dissociated set estimates in the last set of notes.

  A similar estimate obtains if $\mathrm{Re}\hat{f}$ is replaced by $-\mathrm{Re}\hat{f}$ or $\pm\mathrm{Im}\hat{f}$. Conclude that one has $\|\hat{f}\|_\infty \lesssim p^{-1/2}\log p$ with probability at least $1 - p^{-100}$.
- Q2. Given the bounds

  $$N(\delta) \leq CN(\delta + c\delta)^C.$$

  Also, we have $N(\delta) < 10$ when $\delta > 0.9$, deduce a bound of the form

  $$N(\delta) \leq \exp(\exp(C\delta))$$

  for all $0 < \delta \leq 1$. (Use induction. It may help to make the constants C more explicit to distinguish them from each other (i.e. use C_1, C_2, etc.)
- Q3. Let A be a subset of $\mathbf{Z}/p\mathbf{Z}$ of density $|A| = \delta p$, and suppose that $\chi_A - \delta$ is linearly uniform, thus $|\widehat{\chi_A}(\xi)| = o(1)$ for all $\xi \neq 0$. Let P be a proper arithmetic progression in $\mathbf{Z}/p\mathbf{Z}$ of size $|P| \sim p$. Show that $|A \cap P| = (\delta + o(1))|P|$. Thus we have a partial converse to Roth's argument (which derives linear uniformity from uniform density on moderately long progressions); given linear uniformity, we can derive uniform density on *very* long progressions. (Hint:

by an affine transformation we may assume that P is an interval. Use Parseval to estimate $\int (\chi_A - \delta)\psi$, where ψ is a suitable bump function which is supported on P and equals 1 on most of P.

- Q4. Let p be an odd prime. Let f be the function $f(x) := \exp(2\pi i x^2/p)$ on $\mathbf{Z}/p\mathbf{Z}$. Show that $\|f\|_\infty = 1$, that $\|\hat{f}\|_\infty = p^{-1/2}$, but that for every k, $\|\widehat{f(\cdot+k)\overline{f}}\|_\infty = 1$. Thus f is linearly uniform but not quadratically uniform.
- Q5. Let $f(x)$ be any bounded function on $\mathbf{Z}/p\mathbf{Z}$, thus $\|f\|_\infty \lesssim 1$. Show that the statement $\int_{\mathbf{Z}/p\mathbf{Z}} |f(x)| = o(1)$ is true if and only if $f(x) = o(1)$ for most $x \in \mathbf{Z}/p\mathbf{Z}$.
- Q6. Let $f(x)$, $g(x)$, $h(x)$ be maps from one abelian group Z to another Z' such that we have the identity $f(a) + g(a+r) + h(a+2r) = c$ for all $a, r \in Z$ and some fixed $c \in Z'$. Show that f and g and h are affine homomorphisms from Z to Z' (e.g. $f(x) = F(x) + k$, where $k \in Z'$ and F is a genuine homomorphism from Z to Z'). Hint: first normalize so that $f(0) = g(0) = h(0) = 0$. (Why is this result analogous to Proposition 2.4?)
- Q7* (a). Let $f(x)$, $g(x)$, $h(x)$, $k(x)$ be maps from one abelian group Z to another Z' such that we have the identity $f(a) + g(a+r) + h(a+2r) + k(a+3r) = c$ for all $a, r \in Z$ and some fixed $c \in Z'$. Show that for every $k \in \mathbf{Z}$, the function $f(x+k) - f(x)$ is an affine homomorphism from Z to Z', thus $f(x+k) - f(x) = \phi(k)(x) + c(k)$, where $\phi(k)$ is a homomorphism from Z to Z' and $c(k)$ is an element of Z'. (Why is this result analogous to the statement that quadratic uniformity implies the expected number of arithmetic progressions of length 4?)

  (b) Show that the map $k \to \phi(k)$ is itself a homomorphism (from Z to $Hom(Z, Z')$). Conclude that f is an affine quadratic form, i.e. $f(x) = Q(x,x) + L(x) + C$ for some bilinear form $Q : Z \times Z \to Z'$, some linear form (i.e. homomorphism) $L : Z \to Z'$, and some constant $C \in Z'$. Similarly for g, h, k. (Why is this result analogous to the claim that the set Ω of resonant frequencies is partly closed under addition?)

References

[1] T. Gowers, *A new proof of Szemeredi's theorem for arithmetic progressions of length 4*, GAFA **8** (1998), 529-551.

[2] T. Gowers, *A new proof of Szemeredi's theorem*, GAFA **11** (2001), 465-588.

[3] K.F. Roth, *On certain sets of integers*, J. London Math. Soc. **28** (1953), 245-252.

[4] E. Szemeredi, *On sets of integers containing no four elements in arithmetic progression*, Acta Math. Acad. Sci. Hungar. **20** (1969), 89–104.

[5] E. Szemeredi, *On sets of integers containing no k elements in arithmetic progression*, Acta Arith. **27** (1975), 299–345.

Department of Mathematics, UCLA, Los Angeles CA 90095-1555

E-mail address: tao@math.ucla.edu

LECTURE NOTES 5 FOR 254A

TERENCE TAO

1. More on arithmetic progressions of length 3

In last week's notes we proved (among other things) *Roth's theorem*, which is the following:

Theorem 1.1. *Let $0 < \delta \leq 1$. Then there exists an $N(\delta)$ such that for any $N > N(\delta, k)$ and any subset A of $\{1, \dots, N\}$ of cardinality $|A| \geq \delta N$, there exists an proper arithmetic progression in A of length 3.*

Roth's argument is ultimately based on proving the recursive inequality

$$N(\delta) \lesssim N(\delta + c\delta^2)^2$$

which ultimately leads to an estimate for $N(\delta)$ of the form

$$N(\delta) \lesssim e^{e^{C/\delta}}.$$

To put this another way, to ensure that a subset A of $\{1, \dots, N\}$ contains a non-trivial arithmetic progression of length 3, Roth's argument requires that A have density at least $\gtrsim 1/\log\log N$.

One can ask whether this can be improved. Heath-Brown and Szemeredi [4] were able to improve this bound from $1/\log\log N$ to $1/(\log N)^{1/20}$; the main new idea over Roth's argument is to try to exploit data from several Fourier coefficients simultaneously, rather than Roth's approach of locating one large Fourier coefficient and using that to refine the set. Bourgain [2] then improved this further, to $(\log\log N)^{1/2}/(\log N)^{1/2}$, which is the best known result to date.

It is not known what the sharp relationship between N and δ should be. An old conjecture of Erdös and Turan [3] suggests that the correct density should be roughly $1/\log N$ (more precisely, one should have infinitely many progressions whenever A is a subset of the natural numbers with $\sum_{n \in A} 1/n$ divergent; this corresponds to a density somewhere between $1/\log N$ and $1/(\log N)^{1+\varepsilon}$, roughly speaking; see Q1).

The best known counterexample of a set with *no* arithmetic progressions of length 3 is almost sixty years old, and is due to Behrend [1] (see also [8]). This example has density $\exp(-C\sqrt{\log N})$, which is substantially less dense than the positive results (which require a density of at least $\exp(-C\log\log N)$ or so). It is quite remarkable that this simple example has not been improved in so long a span of time. (The best that is known is that there are slightly larger sets, of density

$\exp(-C(\log N)^{1/(k-1)})$, which contain no arithmetic progressions of length k; see [6], [5]).

We begin with Behrend's example.

2. Behrend's example

The objective of this section is to prove

Theorem 2.1. *Let $N \gg 1$ be a large number. Then there exists a subset $A \subset \{1, \ldots, N\}$ of cardinality $|A| \gtrsim N \exp(-C\sqrt{\log N})$ which does not contain any proper arithmetic progressions of length 3.*

Behrend's key observation is the following: in any Euclidean space $\mathbf{R}^d$, no matter how large the dimension d, a sphere $\{x \in \mathbf{R}^d : |x| = r\}$ cannot contain any proper arithmetic progressions of length 3. This is basically because the sphere is the boundary of a strictly convex body.

Moving from continuous space back to discrete, we see that any set of the form

$$\{(x_1, \ldots, x_d) \in \mathbf{Z}^d : x_1^2 + \ldots + x_d^2 = R\}$$

for any integers $d \geq 1$ and $R > 0$, cannot contain any proper arithmetic progressions of length 3.

Now we have to map this example back to $\{1, \ldots, N\}$. Let n, d, and R be large integers to be chosen later, and consider the set

$$S_{n,d,R} := \{(x_1, \ldots, x_d) \in \{1, 2, \ldots, n\}^d : x_1^2 + \ldots + x_d^2 = R\};$$

thus $S_{n,d,R}$ is free of proper arithmetic progressions of length 3 (or higher). Note that as R ranges from n to nd^2, the sets $S_{n,d,R}$ cover the cube $\{1, 2, \ldots, n\}^d$, which has cardinality n^d. Thus by the pigeonhole principle there exists an R with

$$|S_{n,d,R}| \gtrsim n^d/(nd^2).$$

Now we map $S_{n,d,R}$ to $\{1, \ldots, N\}$. Recall that the map $\phi : \{1, \ldots, n\}^d \to \phi(\{1, \ldots, n\})^d$ defined by

$$\phi(x_1, \ldots, x_d) := \sum_{j=1}^{d} x_j (2n)^{j-1}$$

is a Freiman isomorphism of order 2. In particular, $\phi(S_{n,d,R})$ also has cardinality $\gtrsim n^d/(nd^2)$ and has no arithmetic progressions of length 3. It is also contained in $\{1, \ldots, N\}$ if $N \gg C^d n^d$. Thus if we set $n := cN^{1/d}$, then we have found a subset A of $\{1, \ldots, N\}$ without arithmetic progressions of length 3, which has cardinality

$$|A| \gtrsim c^d N/(N^{1/d} d^2) \gtrsim N \exp(-Cd - \frac{\log N}{d} - C \log d).$$

Thus if we set $d \sim \log N$, we obtain the result.

3. Bourgain's argument

We now present Bourgain's argument, which states that any subset of $\{1,\dots,N\}$ with density $\gg \sqrt{\log\log N}/\sqrt{\log N}$ will contain an arithmetic progression. Equivalently, this asserts that $N(\delta) \lesssim (\frac{1}{\delta})^{C/\delta^2}$; see Q2.

By embedding $\{1,\dots,N\}$ in a cyclic group of order $p \sim 3N$, it suffices to prove the following:

Theorem 3.1. *Let $0 < \delta_0 \leq 1$, and let $p \gg (\frac{1}{\delta_0})^{C/\delta_0^2}$ be a prime. Then every subset A of $\mathbf{Z}/p\mathbf{Z}$ of density at least δ_0 contains a proper arithmetic progression (in $\mathbf{Z}/p\mathbf{Z}$) of length 3.*

We now begin the proof of Theorem 3.1. Assume for contradiction that A contains no proper arithmetic progressions of length 3. We will also assume that $|A| = \delta_0 p$ (since one can just adjust δ_0 if necessary if $|A| > \delta_0 p$).

We recall the trilinear form $T(f,g,h)$ defined by

$$T(f,g,h) := \int_{\mathbf{Z}/p\mathbf{Z}} \int_{\mathbf{Z}/p\mathbf{Z}} f(a)g(a+r)h(a+2r)\,da\,dr = p^{-2} \sum_{a,r\in\mathbf{Z}/p\mathbf{Z}} f(a)g(a+r)h(a+2r)$$

where integration is over normalized counting measure on $\mathbf{Z}/p\mathbf{Z}$. Then, since A has no proper arithmetic progressions, we see that

$$T(\chi_A,\chi_A,\chi_A) = p^{-2} \sum_{a\in\mathbf{Z}/p\mathbf{Z}} \chi_A(a)\chi_A(a)\chi_A(a) = |A|/p^2 = \delta_0/p.$$

More generally, we see that for any functions ψ_1,ψ_2,ψ_3 on $\mathbf{Z}/p\mathbf{Z}$, we have

$$T(\chi_A\psi_1,\chi_A\psi_2,\chi_A\psi_3) = p^{-2}\sum_{a\in A} \psi_1(x)\psi_2(x)\psi_3(x). \tag{1}$$

On the other hand, we have the identity

$$T(f,g,h) = \sum_{\xi\in\mathbf{Z}/p\mathbf{Z}} \hat{f}(\xi)\hat{g}(-2\xi)\hat{h}(\xi). \tag{2}$$

As before, the idea is to use both (1) and (2) to force $\chi_A - \delta$ to have a large Fourier coefficient, say $\widehat{\chi_A - \delta}(\xi)$ is large. This rather directly shows that A is fairly dense (density at least $\delta + c\delta^2$) on a *Bohr neighborhood* - some translate of the set $\{x \in \mathbf{Z}/p\mathbf{Z} : \|x\xi/p\| \ll \delta^2\}$. Roughly speaking, the idea is to let ψ be a probability distribution on this Bohr neighborhood, and then note that $\hat{\psi}(\xi)$ is large, thus $(\chi_A - \delta) * \psi$ is large, thus $\chi_A - \delta$ is large on some translate of the support of ψ.

In our previous proof of Roth's theorem, we then passed from the Bohr neighborhood to an arithmetic progression, since that was Freiman equivalent to an interval $\{1,\dots,M\}$ and we could continue the induction. However, this step of passing from neighborhoods to progressions is quite inefficient (in particular, it causes the square in the recursion $N(\delta) \lesssim N(\delta + c\delta^2)^2$). Bourgain's trick is to avoid this step

by staying with the Bohr neighborhood. At the first iteration we only restrict to a Bohr neighborhood of dimension 1, but each stage introduces another dimension to the Bohr neighborhood, because we add another frequency restriction to the previous neighborhood at each stage. Also while Bohr neighborhoods are very similar in spirit to arithmetic progressions and intervals, they are not quite as well-behaved under sums; for instance the sumset of $\{x : \|x\omega/p\| \leq 1/K\}$ with itself can have significantly larger cardinality (about 2^d as large, if d is the dimension of the Bohr neighborhood). However, things are much better if one takes the sumset of $\{x : \|x\omega/p\| \leq 1/K\}$ with $\{x : \|x\omega/p\| \leq 1/K'\}$ if $K' \gg K$ (e.g. $K' \geq d^2K$), so we will be using tricks like this to keep the Bohr neighbourhoods behaving as much like intervals or arithmetic progressions as possible. Actually to make this work neatly we will not quite use Bohr neighborhoods, but rather exponential cutoff functions which are localized near Bohr neighborhoods. Each new Bohr neighborhood has a density which is δ^2 better than the previous one, but has one higher dimension, and the width K has to narrow by a factor of δ^C (mostly because of this K versus K' issue). Thus the procedure has to halt after $O(1/\delta)$ iterations, in which case we have a Bohr neighborhood of dimension $O(1/\delta)$ and width $O(\delta^{C/\delta})$, so a net density of about $O(\delta^{C/\delta^2})$. This is only meaningful if $p \gg \delta^{-C/\delta^2}$, whence the restriction on p.

We now turn to the details.

4. Bohr cutoff functions

We shall need the following cutoff functions.

Given any $x \in \mathbf{Z}/p\mathbf{Z}$, write $\|x/p\|$ for the distance from x/p to the nearest integer, thus $\|x/p\|$ ranges between 0 and 1/2. Given any $d \geq 0$ and any vector $\omega = (\omega_1, \ldots, \omega_d) \in (\mathbf{Z}/p\mathbf{Z})^d$, write

$$\|x\omega/p\| := \sum_{j=1}^{d} \|x\omega_j/p\|,$$

and let ϕ_ω be the function

$$\phi_\omega(x) := \exp(-\|x\omega_j/p\|).$$

We shall be using cutoff functions of the form ϕ_ω^K for various large numbers $K \gg 1$. Roughly speaking, this is a smooth projection to the Bohr neighborhood

$$\{x \in \mathbf{Z}/p\mathbf{Z} : \sum_{j=1}^{d} \|\frac{x\omega_j}{p}\| \lesssim 1/K\}, \tag{3}$$

and is somewhat similar to the characteristic function of a subgroup or of an arithmetic progression (since the above set is somewhat closed under addition, especially if d is not too large.

Note that ϕ_ω is always even and positive, and in particular equals 1 at the origin 0. We observe the following crude lower bound on ϕ_ω^K.

Lemma 4.1. *We have*

$$\int \phi_\omega^K \gtrsim (CdK)^{-d}.$$

Proof This is Minkowski's theorem again. We cover the torus $\mathbf{R}^d/\mathbf{Z}^d$ into about $(CdK)^d$ cubes of side length $1/CdK$. By the pigeonhole principle, one of these cubes must contain at least $(CdK)^{-d}p$ points of the form $x\omega$, where $x \in \mathbf{Z}/p\mathbf{Z}$. Subtracting, we thus see that the cube centered at the origin of sidelength $2/CdK$ also contains at least $(CdK)^{-d}p$ points of the form $x\omega$. Since $\phi_\omega^K \sim 1$ on this cube, the claim follows. ∎

We now introduce the normalized functions $\psi_{\omega,K}$ by

$$\psi_{\omega,K} := \frac{\phi_\omega^K}{\int \phi_\omega^K};$$

thus $\psi_{\omega,K}$ is the L^1-normalized version of ψ_ω^K, and can be thought of as a probability measure concentrated on the Bohr neighborhood (3).

One measure of this is the following localization estimate.

Lemma 4.2. *For any $\lambda \gg d^2$, we have*

$$\int_{\|x\omega/p\| \geq \lambda/K} \psi_{\omega,K}(x)\ dx \lesssim e^{-\lambda/2}.$$

Proof It suffices to show that

$$\int_{(\lambda+1)/K \geq \|x\omega/p\| \geq \lambda/K} \psi_{\omega,K}(x)\ dx \lesssim e^{-\lambda/2},$$

since the claim then follows by summing in λ. Since $\phi_\omega^K(x) \sim e^{-\lambda}$ on this set, it thus suffices to show that

$$|\{x \in \mathbf{Z}/p\mathbf{Z} : (\lambda+1)/K \geq \|x\omega/p\| \geq \lambda/K\}| \lesssim e^{\lambda/2} \int \phi_\omega^K.$$

But we have

$$\int \phi_\omega^K \lesssim |\{x \in \mathbf{Z}/p\mathbf{Z} : \|x\omega/p\| \leq 1/K\}|.$$

Now we can cover the region

$$\{y \in \mathbf{R}^d/\mathbf{Z}^d : (\lambda+1)/K \geq \|y\| \geq \lambda/K\}$$

by $O((C\lambda)^d) = O(e^{\lambda/2})$ l^1-balls of radius $1/2K$. Each one of these balls can contain at most $|\{x \in \mathbf{Z}/p\mathbf{Z} : \|x\omega/p\| \leq 1/K\}|$ points of the form $x\omega/p$ by the Minkowski's theorem argument, and the claim follows. ∎

We now prove a basic estimate, that narrower projections are essentially convolution idempotents for wider projections.

Lemma 4.3. *Let $\omega \in (\mathbf{Z}/p\mathbf{Z})^d$ and $K' \gg d^2K$. Then we have the pointwise estimate*

$$\int \psi_{\omega,K}(x-jy)\psi_{\omega/k,K'}(y)\ dy = (1+O(\frac{K}{K'}))\psi_{\omega,K}(x)$$

for $j = 1, 2$ and $k = 1, 2$. More generally, for any subset A' of $\mathbf{Z}/p\mathbf{Z}$ and any $\omega_{d+1} \in \mathbf{Z}/p\mathbf{Z}$, we have

$$\int \psi_{\omega,K}(x-jy)\chi_{A'}(x-jy)\psi_{(\omega/k,\omega_{d+1}),K'}(y)\ dy = (\int \chi_{A'}(x-jy)\psi_{\omega,K'}(y)\ dy + O(d^2\frac{K}{K'}))\psi_{\omega,K}(x). \tag{4}$$

In particular we have

$$\|\psi_{\omega,K} * \psi_{\omega/k,K'} - \psi_{\omega,K}\|_1 \lesssim d^2\frac{K}{K'}.$$

Proof It suffices to prove (4). Fix x. We have to show that

$$\int (\psi_{\omega,K}(x-jy) - \psi_{\omega,K}(x))\chi_{A'}(x-jy)\psi_{(\omega/k,\omega_{d+1}),K'}(y)\ dy = O(d^2\frac{K}{K'})\psi_{\omega,K}(x).$$

Dividing out by $\psi_{\omega,K}(x)$ and using the triangle inequality, it will suffice to show

$$\int (e^{K\|ky\omega'/p\|} - 1)\psi_{\omega',K'}(y)\ dy = O(d^2\frac{K}{K'})$$

where $\omega' := (\omega/k, \omega_{d+1}) \in (\mathbf{Z}/p\mathbf{Z})^{d+1}$. We split this into the region where $\|y\omega'/p\| \gg d^2/K'$ and $\|y\omega'/p\| \lesssim d^2/K'$. In the first part we use Lemma 4.2 and decomposition into shells $\lambda \leq \|y\omega'/p\| \leq \lambda + 1$; for the second part we use the Taylor approximation

$$e^{K\|ky\omega'/p\|} - 1 = O(d^2\frac{K}{K'})$$

and the L^1-normalization of $\psi_{\omega',K'}$. ∎

This estimate has two important consequences for our application. First, the functional T is large when applied to these measures $\psi_{\omega,K}$:

Corollary 4.4. *Let $\omega \in (\mathbf{Z}/p\mathbf{Z})^d$ and $K' \gg d^2K$. Then we have*

$$T(\psi_{\omega,K}^{1/2}, \psi_{\omega,K'}, \psi_{\omega,K}^{1/2}) = 1 + O(d^2K/K').$$

More generally, for any subset $A' \subseteq \mathbf{Z}/p\mathbf{Z}$ we have

$$T(\psi_{\omega,K}^{1/2}, \psi_{\omega,K'}, \chi_{A'}\psi_{\omega,K}^{1/2}) = (1 + O(d^2K/K'))\int \chi_{A'}\psi_{\omega,K}. \tag{5}$$

Proof It suffices to prove (5). The left-hand side is

$$\int \chi_{A'}(x)\psi_{\omega,K}(x)^{1/2}(\int \psi_{\omega,K}(x-2y)^{1/2}\psi_{\omega,K'}(y)\ dy)\ dx,$$

where we have used the even-ness of $\psi_{\omega,K}$. Since $\psi_{\omega,K}^{1/2}$ is a scalar multiple of $\psi_{\omega,K/2}$, we may use Lemma 4.3 to write this as

$$\int \chi_{A'}(x)\psi_{\omega,K}(x)^{1/2}\psi_{\omega,K}(x)^{1/2}(1 + O(d^2K/K'))\ dx$$

and the claim follows from the L^1 normalization of $\psi_{\omega,K}$. ∎

Secondly, if the Fourier coefficient of a wide projection is large, then so is the coefficient of the narrow projection:

Lemma 4.5. *Let $\omega \in (\mathbf{Z}/p\mathbf{Z})^d$ and $K' \gg d^2 K$. Suppose that $\xi \in \mathbf{Z}/p\mathbf{Z}$ is such that*

$$|\hat{\psi}_{\omega,K}(-2\xi)| \gg d^2 \frac{K}{K'}.$$

Then we have

$$|\hat{\psi}_{\omega/2,K'}(\xi)| \sim 1.$$

Proof Write $\omega' := \omega/2$. Let $y \in \mathbf{Z}/p\mathbf{Z}$ be such that $\|y\omega'/p\| \leq d^2/K'$. Then if we write $y' := y/2$ in $\mathbf{Z}/p\mathbf{Z}$, then $\|y'\omega/p\| \leq d^2/K'$. In particular, we have

$$\psi_{\omega,K}(x+y') - \psi_{\omega,K}(x) = O(d^2 \frac{K}{K'} \psi_{\omega,K}).$$

Integrating this against $e^{-2\pi i(-2\xi)x/p}$, we obtain

$$(e^{2\pi i i(2\xi)y'/p} - 1)\hat{\psi}_{\omega,K}(-2\xi) = O(d^2 \frac{K}{K'}).$$

By hypothesis we thus have

$$|e^{2\pi i i(2\xi)y'/p} - 1| \ll 1,$$

so in particular

$$|e^{-2\pi i i\xi y/p} - 1| \ll 1.$$

Integrating this against $\psi_{\omega',K'}$ we obtain

$$\int_{\|y\omega'/p\| \leq d^2/K'} (e^{-2\pi i i\xi y/p} - 1)\psi_{\omega',K'}(y)\, dy| \ll 1.$$

On the other hand, from Lema 4.2 we have

$$\int_{\|y\omega'/p\| \geq d^2/K'} (e^{-2\pi i i\xi y/p} - 1)\psi_{\omega',K'}(y)\, dy| \ll 1.$$

Adding the two, we obtain the claim. ■

Finally, we present one of the inductive steps, reminiscent of Roth's argument: if $\chi_{A'}$ has density δ on $\psi_{\omega,K}$, and $(\chi_{A'} - \delta)\psi_{\omega,K}$ has a large Fourier coefficient, then some translate of $\chi_{A'}$ has density strictly greater than δ for some refined cutoff $\psi_{\omega',K'}$.

Lemma 4.6. *Let $\omega \in (\mathbf{Z}/p\mathbf{Z})^d$ and $K' \gg d^2 K$. Suppose that A' is a subset of $\mathbf{Z}/p\mathbf{Z}$, and set $\delta := \int \chi_{A'}\psi_{\omega,K}$. Write $f := (\chi_{A'} - \delta)\psi_{\omega,K}$. Suppose we have a frequency $\xi \in \mathbf{Z}/p\mathbf{Z}$ such that $|\hat{f}(\xi)| \gg \frac{d^2 K}{K'}$. Then there exists an $\omega' \in (\mathbf{Z}/p\mathbf{Z})^{d+1}$ and a translate A'' of A' such that*

$$\int \chi_{A''}\psi_{\omega',K'} \geq \delta + c|\hat{f}(\xi)|.$$

Proof Write $\omega' := (\omega, \xi)$. By construction (and Lemma 4.2), we see that

$$\hat{\psi}_{\omega',K'}(\xi) \sim 1$$

since the phase $e^{-2\pi i x\xi/p}$ is close to 1 on the effective support of $\psi_{\omega',K'}$. Thus we have

$$|\hat{f}(\xi)\hat{\psi}_{\omega',K'}(\xi)| \sim |\hat{f}(\xi)|$$

and in particular that

$$\|f * \psi_{\omega',K'}\|_1 \gtrsim |\hat{f}(\xi)|.$$

Since f has mean zero, so does $f * \psi_{\omega',K'}$. Thus we have

$$\int (f * \psi_{\omega',K'})_+ \gtrsim |\hat{f}(\xi)|,$$

and by the pigeonhole principle so there exists an x such that

$$f * \psi_{\omega',K'}(x) \gtrsim |\hat{f}(\xi)|\psi_{\omega,K}(x).$$

But by Lemma 4.3, we have

$$\psi_{\omega,K} * \psi_{\omega',K'}(x) = (1 + O(\frac{d^2K}{K'}))\psi_{\omega,K}(x)$$

and

$$\chi_{A'}\psi_{\omega,K} * \psi_{\omega',K'}(x) = (1 + O(\frac{d^2K}{K'}))(\chi_{A'} * \psi_{\omega',K'})(x)\psi_{\omega,K}(x),$$

so we conclude that

$$\chi_{A'} * \psi_{\omega',K'}(x) - \delta \gtrsim |\hat{f}(\xi)|$$

and the claim follows. ∎

5. Putting it all together

Using all the above tools, we can now conclude Bourgain's argument. Recall that we are assuming A to have density δ_0 and no proper arithmetic progressions of length 3. Let C_0 be a large absolute constant (e.g. $C_0 = 1000$); assume that $p \gg (1/\delta)^{C_0^5/\delta^2}$. The key inductive step is the following:

Proposition 5.1. *Let $d \leq C_0/\delta_0$, let $\omega \in (\mathbf{Z}/p\mathbf{Z})^d$, and let $1 \leq K \leq \delta^{C_0/\delta_0}$. Suppose we have a translate A' of A such that*

$$\int \chi_{A'}\psi_{\omega,K} = \delta$$

for some $\delta_0 \leq \delta \leq 1$. Then there exists $\omega' \in (\mathbf{Z}/p\mathbf{Z})^{d+1}$, $K' \leq \delta^{-C}K$ and a translate A'' of A' such that

$$\int \chi_{A''}\psi_{\omega,K'} = \delta'$$

where $\delta' \geq \delta + c\delta^2$.

Proof The first step is to refine $\psi_{\omega,K}$ a little bit in such a way that the density of A' remains close to δ.

Set $K_j := (cd^2\delta^2)^j K$ for $j = 1, 2, 3$. From Lemma 4.3 we have

$$\int \chi_{A'}(\psi_{\omega,K} * \psi_{\omega,K_j} - \psi_{\omega,K}) = O(c\delta^2)$$

for $j = 1, 2$, and hence

$$\int \chi_{A'}(\psi_{\omega,K} * \psi_{\omega,K_j}) = \delta + O(c\delta^2)$$

which can be rewritten (using the L^1 normalization of $\psi_{\omega,K}$) as

$$\int \chi_{A'}(\psi_{\omega,K} * \psi_{\omega,K_j}) = \delta + O(c\delta^2)$$

$$\int (\chi_{A'} * \psi_{\omega,K_j} - \delta)\psi_{\omega,K} = O(c\delta^2).$$

Also we can assume that

$$(\chi_{A'} * \psi_{\omega,K_j} - \delta)(x) < c\delta^2$$

for all $x \in \mathbf{Z}/p\mathbf{Z}$, since we are done otherwise (by setting $K' := K$, $\omega' = (\omega, 0)$, and $A'' := A' - x$). In particular we have

$$\int (\chi_{A'} * \psi_{\omega,K_j} - \delta)_+ \psi_{\omega,K} = O(c\delta^2).$$

Combining this with the previous we see that

$$\int \sum_{j=1}^{2} |\chi_{A'} * \psi_{\omega,K_j} - \delta| \psi_{\omega,K} = O(c\delta^2).$$

By the pigeonhole principle we can thus find an x_0 such that

$$\sum_{j=1}^{2} |\chi_{A'} * \psi_{\omega,K_j} - \delta|(x_0) = O(c\delta^2).$$

Without loss of generality we can take $x_0 = 0$. Thus we have

$$\int \chi_{A'}(x)\psi_{\omega,K_j}(x) = \delta + O(c\delta^2) \tag{6}$$

for $j = 1, 2$.

Now we compute the quantity

$$T(\chi_{A'}\psi_{\omega,K_1}^{1/2}, \chi_{A'}\psi_{\omega,K_2}, \chi_{A'}\psi_{\omega,K_1}^{1/2}) - T(\delta\psi_{\omega,K_1}^{1/2}, \delta\psi_{\omega,K_2}, \delta\psi_{\omega,K_1}^{1/2}). \tag{7}$$

Since A has no proper arithmetic progressions of length 3, we have (1). In particular from Lemma 4.1 we have the very crude bound

$$T(\chi_{A'}\psi_{\omega,K_1}^{1/2}, \chi_{A'}\psi_{\omega,K_2}, \chi_{A'}\psi_{\omega,K_1}^{1/2}) \leq (CdK_2)^{Cd}/p.$$

On the other hand, from Corollary 4.4 we have

$$T(\delta\psi_{\omega,K_1}^{1/2}, \delta\psi_{\omega,K_2}, \delta\psi_{\omega,K_1}^{1/2}) \sim \delta^3.$$

From our lower bounds on p and upper bounds on d, K_2 we thus have

$$|(7)| \sim \delta^3.$$

Now we split (7) into pieces. We first investigate the quantity

$$T(\chi_{A'}\psi_{\omega,K_1}^{1/2}, \chi_{A'}\psi_{\omega,K_2}, \chi_{A'}\psi_{\omega,K_1}^{1/2}) - T(\chi_{A'}\psi_{\omega,K_1}^{1/2}, \delta\psi_{\omega,K_2}, \chi_{A'}\psi_{\omega,K_1}^{1/2}). \tag{8}$$

Suppose first that $|(8)| \gtrsim \delta^3$. By (2) we thus have

$$\sum_{\xi} |\widehat{\chi_{A'}\psi_{\omega,K_1}^{1/2}}(\xi)|^2 |\widehat{(\chi_{A'} - \delta)\psi_{\omega,K_2}}(-2\xi)| \gtrsim \delta^3.$$

But by Plancherel we have

$$\sum_\xi |\widehat{\chi_{A'}\psi^{1/2}_{\omega,K_1}}(\xi)|^2 = \int \chi_{A'}\psi_{\omega,K_1} \sim \delta$$

and hence we have a frequency ξ such that

$$|\widehat{(\chi_{A'} - \delta)\psi}_{\omega,K_2}(-2\xi)| \gtrsim \delta^2.$$

By Lemma 4.6 we can thus find a translate A'' of A, an $\omega' \in (\mathbf{Z}/p\mathbf{Z})^{d+1}$ such that

$$\int \chi_{A''}\psi_{\omega',K_3} \geq \delta + c\delta^2$$

as desired.

We may now suppose that $|(8)| \ll \delta^3$, which implies that

$$T(\chi_{A'}\psi^{1/2}_{\omega,K_1}, \delta\psi_{\omega,K_2}, \chi_{A'}\psi^{1/2}_{\omega,K_1}) - T(\delta\psi^{1/2}_{\omega,K_1}, \delta\psi_{\omega,K_2}, \delta\psi^{1/2}_{\omega,K_1})$$

has magnitude $\sim \delta^3$. We split this as a linear combination of

$$T(\delta\psi^{1/2}_{\omega,K_1}, \delta\psi_{\omega,K_2}, (\chi_{A'} - \delta)\psi^{1/2}_{\omega,K_1}) \tag{9}$$

and

$$T((\chi_{A'} - \delta)\psi^{1/2}_{\omega,K_1}, \delta\psi_{\omega,K_2}, (\chi_{A'} - \delta)\psi^{1/2}_{\omega,K_1}) \tag{10}$$

where we use the symmetry $T(f,g,h) = T(h,f,g)$. To estimate (9), we observe from Corollary 4.4 that

$$T(\delta\psi^{1/2}_{\omega,K_1}, \delta\psi_{\omega,K_2}, \delta\psi^{1/2}_{\omega,K_1}) = \delta^3 + O(c\delta^4)$$

and

$$T(\delta\psi^{1/2}_{\omega,K_1}, \delta\psi_{\omega,K_2}, \chi_{A'}\psi^{1/2}_{\omega,K_1}) = \delta^2(1 + O(c\delta))(\int \chi_{A'}\psi_{\omega,K_1}) = \delta^3 + O(c\delta^4)$$

and so (9) is $O(c\delta^4)$. Thus we must have $|(10)| \sim \delta^3$, thus by (2)

$$\sum_\xi |\widehat{(\chi_{A'} - \delta)\psi}^{1/2}_{\omega,K_1}(\xi)|^2 |\widehat{\psi_{\omega,K_2}}(-2\xi)| \gtrsim \delta^3.$$

On the other hand, from Plancherel we have

$$\sum_\xi |\widehat{(\chi_{A'} - \delta)\psi}^{1/2}_{\omega,K_1}(\xi)|^2 = \int (\chi_{A'} - \delta)^2\psi_{\omega,K_1} = O(\delta),$$

and thus we have

$$\sum_{\xi : |\widehat{\psi_{\omega,K_2}}(-2\xi)| \gtrsim c\delta^2} |\widehat{(\chi_{A'} - \delta)\psi}^{1/2}_{\omega,K_1}(\xi)|^2 |\widehat{\psi_{\omega,K_2}}(-2\xi)| \gtrsim \delta^3.$$

By Lemma 4.5 we thus have

$$\sum_\xi |\widehat{(\chi_{A'} - \delta)\psi}^{1/2}_{\omega,K_1}(\xi)|^2 |\widehat{\psi_{\omega/2,K_3}}(\xi)|^2 \gtrsim \delta^3,$$

which by Plancherel implies that

$$\int |((\chi_{A'} - \delta)\psi^{1/2}_{\omega,K_1}) * \psi_{\omega/2,K_3}|^2 \gtrsim \delta^3. \tag{11}$$

So now let us investigate the quantity

$$|((\chi_{A'} - \delta)\psi_{\omega,K_1}^{1/2}) * \psi_{\omega/2,K_3}|(x).$$

We can rewrite this as

$$\int (\chi_{A'}(x-y) - \delta)\psi_{\omega,K_1}^{1/2}(x-y)\psi_{\omega/2,K_3}(y)\ dy.$$

By Lemma 4.3, this is equal to

$$\int (\chi_{A'}(x-y) - \delta)\psi_{\omega/2,K_3}(y)\ dy\psi_{\omega,K_1}^{1/2}(x) + O(c\delta^2\psi_{\omega,K_1}^{1/2}(x)).$$

The error term contributes $O(c\delta^4)$ to (11), and hence we have

$$\int |\int (\chi_{A'}(x-y) - \delta)\psi_{\omega/2,K_3}(y)\ dy|^2\psi_{\omega,K_1}(x)\ dx \gtrsim \delta^3. \tag{12}$$

Note that we may assume as before that

$$\int (\chi_{A'}(x-y) - \delta)\psi_{\omega/2,K_3}(y)\ dy \leq c\delta^2 \tag{13}$$

for all x, since we are done otherwise. In particular, the expression in absolute values in (12) is bounded by δ, and we thus have

$$\int |\int (\chi_{A'}(x-y) - \delta)\psi_{\omega/2,K_3}(y)\ dy|\psi_{\omega,K_1}(x)\ dx \gtrsim \delta^2.$$

From this and (13) we have

$$|\int \int (\chi_{A'}(x-y) - \delta)\psi_{\omega/2,K_3}(y)\ dy\psi_{\omega,K_1}(x)\ dx| \gtrsim \delta^2.$$

or equivalently,

$$|\int (\chi_{A'} - \delta)(\psi_{\omega/2,K_3} * \psi_{\omega,K_1})| \gtrsim \delta^2.$$

From Lemma 4.3 we may replace $\psi_{\omega/2,K_3} * \psi_{\omega,K_1}$ by ψ_{ω,K_1} with an acceptable error, thus

$$|\int (\chi_{A'} - \delta)\psi_{\omega,K_1}| \gtrsim \delta^2.$$

But this contradicts (6). This proves the Proposition. ∎

Observe that the map $\delta \to \delta + c\delta^2$ will reach size > 1 after most C/δ_0 steps, starting from δ_0. Thus iterating Proposition 5.1 this many times (starting with $d = 0$, $K = 1$, and $A' = A$), we can eventually get the density greater than 1, a contradiction. Thus A must contain a progression of length 3.

6. Exercises

- Q1. Let $\varepsilon > 0$ be a small number. Suppose we knew that for every large $N \gg 1$ and every subset $A \subseteq \{1, \ldots, N\}$ of density at least $1/(\log N)^{1+\varepsilon}$ contained a non-trivial arithmetic progression of length 3. Deduce that for any subset $B \subset \mathbf{Z}^+$ with $\sum_{n\in A} 1/n = \infty$ contains an infinite number of arithmetic progressions of length 3.

- Q2. Show that the following two statements are equivalent: (a) For every $N \gg 1$ and every subset A of $\{1, \ldots, N\}$ of density $\gg \sqrt{\log\log N}/\sqrt{\log N}$, A contains a proper arithmetic progression of order 3. (b) For every δ and every $N \gg (\frac{1}{\delta})^{C/\delta^2}$, any subset A of $\{1, \ldots, N\}$ of density at least δ will contain an arithmetic progression of order 3. (Note: this equivalence has nothing to do with arithmetic progressions, and is basically a matter of algebra.)

References

[1] F. A. Behrend, *On sets of integers which contain no three terms in arithmetic progression*, Proc. Nat. Acad. Sci. 32 (1946), 331–332

[2] J. Bourgain, *On triples in arithmetic progression*, GAFA **9** (1999), 968–984.

[3] P. Erdös, P. Turan, *On some sequences of integers*, J. London Math. Soc. **11** (1936), , 261–264.

[4] D.R. Heath-Brown, *Integer sets containing no arithmetic progressions*, J. London Math. Soc. **35** (1987), 385–394.

[5] I. Laba, M. Lacey, *On sets of integers not containing long arithmetic progressions*, preprint.

[6] R.A. Rankin, *Sets of integers containing not more than a given number of terms in arithmetic progression*, Proc. Roy. Soc. Edinburgh Sect. A **65** (1960/1961), 332–344.

[7] K.F. Roth, *On certain sets of integers*, J. London Math. Soc. **28** (1953), 245–252.

[8] R. Salem, D.C. Spencer, *On sets of integers which contain no three terms in arithmetic progression*, Proc. Nat. Acad. Sci. 32 (1942), 561-563.

Department of Mathematics, UCLA, Los Angeles CA 90095-1555

E-mail address: tao@math.ucla.edu

LECTURE NOTES 6 FOR 254A

TERENCE TAO

1. Sums and products

We now leave the topic of arithmetic progressions, and return to topics closer to those studied upon at the beginning of the course, where we studied the relationship between such sets as $A + A$ and $A - A$. This time, however, we shall look at the sets $A + A$ and $A \cdot A$, where

$$A \cdot A = \{ab : a, b \in A\}.$$

Note that if one wanted to study $A \cdot A$ and A/A (where A does not contain 0, of course), one could just take logarithms and reduce things to a problem and sums and differences. But to study $A + A$ and $A \cdot A$ simultaneously is a different matter, and the theory here is much less well understood.

From Freiman's theorem we expect $|A + A| \sim |A|$ only when A resembles an arithmetic progression; similarly we expect $|A \cdot A| \sim |A|$ only when A resembles a geometric progression. Since arithmetic progressions and geometric progressions do not resemble each other, it thus seems reasonable to expect at least one of $|A + A|$ and $|A \cdot A|$ to be large, at least when A is a subset of $\mathbf{R}$. This is indeed the case; the best known bound on this is Elekes's theorem $\max(|A + A|, |A \cdot A|) \gtrsim |A|^{5/4}$, which we shall prove shortly.

Interestingly, the geometry of the plane now begins to play a role. The reason can already be seen in the familiar equation $y = mx + b$ for a line in the plane; this is an equation which involves both multiplication and addition, which is exactly the problem encountered above.

To address the geometry of the plane, we first need to understand the topology of the plane - and in particular, the concept of *crossing numbers*. This will be handled next.

2. Crossing numbers

Let $G = (V, E)$ be a finite undirected, multiplicity-free graph with a collection V of vertices, and E of edges. A *drawing* $\phi(G)$ of G is a map ϕ on (V, E) which maps each vertex v in V to a point $\phi(v)$ in $\mathbf{R}^2$, while mapping each edge e in E, with vertices v_1 and v_2, to an (open) curve $\phi(e)$ in $\mathbf{R}^2$ with vertices $\phi(v_1)$ and $\phi(v_2)$. We say that a drawing of G is *proper* if (a) the points $\phi(v)$ are all distinct, (b) the

(open) curves $\phi(e)$ are simple and do not pass through any points in $\phi(V)$, (c) any two distinct open curves $\phi(e)$, $\phi(e')$ only intersect in a finite number of points, (d) any two distinct open curves with a shared endpoint are disjoint. It is clear that every finite graph G has at least one drawing which is proper (map the vertices to randomly selected points, and the edges to line segments). Given any proper drawing $\phi(G)$ of a graph G, define the *crossing number* $cr(\phi(G))$ to be the number of times the edges $\phi(e)$, $\phi(e')$ cross each other, i.e.

$$cr(\phi(G)) = \sum_{\{e,e'\} \in E \times E : e \neq e'} |\phi(e) \cap \phi(e')|.$$

Thus $cr(\phi(G))$ is a non-negative integer. We define the *crossing number* $cr(G)$ of a graph G to be

$$cr(G) := \inf_{\phi} cr(\phi(G)),$$

i.e. the minimal crossing number over all proper drawings ϕ of G.

We now pursue the question of how to estimate $cr(G)$, especially from below. Clearly we have $cr(G') \leq cr(G)$ if G' is a subgraph of G; we can convert this monotonicity property into a strict monotonicity property if $cr(G)$ is positive.

Lemma 2.1. *Let $G = (V, E)$ be a graph such that $cr(G)$ is positive. Then there exists an edge e in E such that $cr(G) \geq cr(G - \{e\}) + 1$.*

Proof Let $M := \inf_{e \in E} cr(G - \{e\})$, and let ϕ be a proper drawing of G. We know that $cr(\phi(G))$ is positive; we need to show that it is at least $M + 1$.

Since $cr(\phi(G))$ is positive, we have distinct edges $\phi(e)$ and $\phi(e')$ which intersect at least 1. But we know that $cr(\phi(G - \{e\})$ intersects at least M times. The claim follows. ■

Next, we give a simple condition to show when $cr(G)$ is positive.

Lemma 2.2. *If $|E| > 3|V| - 6$, then $cr(G)$ is positive.*

Proof Suppose for contradiction that we had a graph G and a drawing ϕ with no crossings, but for which $|E| > 3|V| - 6$. The drawing $\phi(G)$ subdivides the plane into a collection F of open regions (including one infinite region). By Euler's formula we have $|V| - |E| + |F| = 2$. But also, since each edge is adjacent to two faces, and each face is adjacent to at least three edges, we have $3|F| \leq 2|E|$. Combining the two equations we obtain $|E| \leq 3|V| - 6$, contradiction. ■

Combining these two Lemmas together we immediately obtain the bound

$$cr(G) \geq \max(|E| - 3|V| + 6, 0).$$

Now we improve this estimate using randomization arguments. Let $G = (V, E)$ be a graph, and let $0 < p < 1$ be a parameter to be chosen later. Let $G' = (V', E')$ be the (random) sub-graph of G formed by allowing each vertex $v \in V$ to lie in V' with independent probability p, and let $e \in E$ lie in E' if both its vertices lie

in V' (so each edge has a probability p^2 of occurring in E', in particular we have $\mathbf{E}(|E'|) = p^2|E|$. Let $\phi(G)$ be a drawing of G, which of course induces a drawing $\phi(G')$ of G'. Each crossing of $\phi(G)$ requires two distinct edges arising from four distinct points, and thus has a probability p^4 of also occurring in $\phi(G')$. These events are not independent, but expectation is still linear, so we have

$$\mathbf{E}(cr(\phi(G')) = p^4 cr(G).$$

On the other hand, using the above bound we have

$$\mathbf{E}(cr(\phi(G')) \geq \mathbf{E}(|E'| - 3|V| + 6) = \mathbf{E}(|E'|) - 3\mathbf{E}(|V|) + 6 = p^2|E| - 3p|V| + 6.$$

Thus we have the estimate

$$cr(G) \geq \frac{p^2|E| - 3p|V| + 6}{p^4}.$$

If $|E| \geq 5|V|$, then we can pick $p = 4|V|/|E|$, and obtain the *crossing number bound*

$$cr(G) \gtrsim |E|^3/|V|^2 \text{ when } |E| \geq 5|V|.$$

This bound is quite sharp (see Q1). It can be rewritten as

$$|E| \lesssim cr(G)^{1/3}|V|^{2/3} + |V|. \tag{1}$$

3. The Szemerédi-Trotter theorem

We now use the crossing number estimate to obtain a basic estimate on incidences between points and lines.

Let P be a finite collection of points in the plane, and let L be a finite collection of lines. We consider the quantity $I(P, L) := |\{(p, l) \in P \times L : p \in l\}|$, the set of incidences between P and L.

Form the graph G whose vertices V are the set P of points, and whose edges E are the bounded line segments in L formed by any two points in P. Clearly $|V| = |P|$. As for $|E|$, note that

$$|E| = \sum_{l \in L} (|P \cap l| - 1)_+ \geq \sum_{l \in L} |P \cap l| - |L| = I(P, L) - |L|.$$

On the other hand, since each two lines in L intersect in at most one point, we have $cr(G) \leq |L|^2$. Applying (1), we thus see that

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

This is the Szemerédi-Trotter theorem (which initially had a much longer proof [16]!). This bound is sharp (Q2).

4. Elekes's theorem

Now we can prove a bound on $A + A$ and $A \cdot A$, when A is a finite subset of the reals $\mathbf{R}$. Let P be the set

$$(A + A) \times (A \cdot A),$$

and let L consist of all the lines $\{(x,y) : y = a(x-a')\}$ where $a, a' \in A$. Then $|P| = |A+A||A \cdot A|$ and $L = |A|^2$. Also, each line L is incident to at least $|A|$ points in P (why?), so $I(P,L) \geq |L||A| = |A|^3$. From the Szemerédi-Trotter theorem we conclude that

$$|A|^3 \lesssim |A|^{4/3}|P|^{2/3} + |P| + |A|^2,$$

from which we conclude (if $|A| \gg 1$; the case $|A| \lesssim 1$ can be dealt with manually) that

$$|P| \gtrsim |A|^{5/2},$$

and thus we obtain Elekes's theorem [7]

$$\max(|A+A|, |A \cdot A|) \gtrsim |A|^{5/4}. \tag{3}$$

It is not known what the best bound on this quantity is; (3) has not been improved despite the argument being remarkably simple. An old conjecture of Erdös and Szemerédi [9] claims that

$$\max(|A+A|, |A \cdot A|) \gtrsim |A|^{2-\varepsilon},$$

i.e. either A is almost totally non-closed under addition or totally non-closed under multiplication.

5. The finite field case

Elekes's theorem showed that when A was a finite subset of $\mathbf{R}$, then there was some significant failure of A to be closed under either addition or multiplication. We could ask whether the same is true in finite fields F. Of course, if A is a subfield of F, then $A+A$ and $A \cdot A$ are both equal to A, so in that case there is no failure of closure under either multiplication or addition. (Of course the crossing number argument relied on Euler's formula, and thus on the topology of the plane, and so does not work in finite fields).

However, suppose F was a prime finite field $F := \mathbf{Z}/p\mathbf{Z}$. Then F has no non-trivial subfields, and so the above counterexample does not apply. In this case we can get a small improvement, though not yet on the level of Elekes's theorem:

Theorem 5.1. [4] *Let F be a finite field of prime order, and let A be a subset of F such that $|F|^\delta < |A| < |F|^{1-\delta}$ for some $\delta > 0$. Then we have $\max(|A+A|, |A \cdot A|) \gtrsim |A|^{1+\varepsilon}$ for some $\varepsilon = \varepsilon(\delta) > 0$.*

This theorem requires a certain amount of machinery, though fortunately we have already developed most of it. The idea is to prove by contradiction, assuming that $|A+A|$ and $|A \cdot A|$ are close to $|A|$, and conclude that A is behaving very much like a subfield of F - enough so that one can do some basic linear algebra, and eventually conclude that $|F|$ is a power of $|A|$ (or more precisely $|A'|$ for some variant A' of A), a contradiction since $|F|$ is prime.

6. First ingredient: Sum-product estimates

Recall the sum-set estimates from week 1 notes, which said (among other things) that if $|A+A| \sim |A|$, then in fact $|A \pm A \pm \ldots \pm A| \sim |A|$. Exponentiating this, we see that if A does not contain 0 and $|A \cdot A| \sim |A|$, then in fact we have the product-set estimates $|A \cdot A \cdot \ldots \cdot A/(A \cdot \ldots \cdot A)| \sim |A|$.

Now suppose that (to prove Theorem 5.1 by contradiction) we knew that we had both $|A+A| \sim |A|$ and $|A \cdot A| \sim |A|$. Then sum estimates and product estimates allow us to control any expression of A which involves just sums and differences, or just products and quotients. But we do not yet know how to control expressions which involve both, e.g. $A \cdot A + A \cdot A$. However, it turns out that we can do so if we refine A a little:

Lemma 6.1. [11] *Let A be a non-empty subset of F such that*

$$|A+A|, |A \cdot A| \leq K|A|.$$

We shall use $X \lessapprox Y$ to denote the estimate $X \leq CK^C Y$. Then there is a subset A' of A with $|A'| \gtrapprox |A|$ such that

$$|A' \cdot A' - A' \cdot A'| \lessapprox |A'|.$$

This only controls one joint expression of A'; we shall control all the others as well in the next section.

To prove Lemma 6.1, we recall Gowers' quantitative formulation [10] of the Balog-Szemeredi lemma [1]:

Theorem 6.2. [10], [3] *Let A, B be finite subsets of an additive group with cardinality $|A| = |B|$, and let G be a subset of $A \times B$ with cardinality*

$$|G| \approx |A||B|$$

such that we have the bound

$$|\{a+b : (a,b) \in G\}| \lessapprox |A|.$$

Then there exists subsets A', B' of A and B respectively with $|A'| \geq cK^{-C}|A|$, $|B'| \geq cK^{-C}|B|$ such that

$$|A' - B'| \lessapprox |A|.$$

Indeed, we have the stronger statement that for every $a' \in A$ and $b' \in B$, there are $\gtrapprox |A|^5$ solutions to the problem

$$a' - b' = (a_1 - b_1) - (a_2 - b_2) + (a_3 - b_3); a_1, a_2, a_3 \in A; b_1, b_2, b_3 \in B.$$

Proof (Sketch) Call an element d a *popular difference* if it can be written in the form $a-b$ for $\gtrapprox |A|$ pairs $(a,b) \in A \times B$. The hypotheses on G ensure that there are $\gtrapprox$ popular differences. Thus by the construction in previous notes, we can find large subsets A' and B' of A and B respectively with size $\gtrapprox |A|$ such that every element a' of A' is aware of every element b' of B', thus there are $\gtrapprox |A|^2$ solutions to the problem

$$a' - b' = d_1 - d_2 + d_3$$

where d_1, d_2, d_3 are popular differences. The claim follows. ∎

Now we can prove Lemma 6.1. Without loss of generality we may assume that $|A| \gg 1$ is large, and that $0 \neq A$ (since removing 0 from A does not significantly affect any of the hypotheses).

We first observe from Theorem 6.2 that we can find subsets C, D of A with $|C|, |D| \approx |A|$ such that every element in $C - D$ has $\gtrapprox |A|^5$ representations of the form

$$a_1 - a_2 + a_3 - a_4 + a_5 - a_6; \quad a_1, \ldots, a_6 \in A.$$

Multiplying this by an arbitrary element of $A \cdot A \cdot A/A \cdot A$, we see that every element of $(C - D) \cdot A \cdot A \cdot A/A \cdot A$ has $\gtrapprox |A|^5$ representations of the form

$$b_1 - b_2 + b_3 - b_4 + b_5 - b_6; \quad b_1, \ldots, b_6 \in A \cdot A \cdot A \cdot A/A \cdot A.$$

However, by the multiplicative form of sumset estimates, the set $A \cdot A \cdot A \cdot A/A \cdot A$ has cardinality $\approx |A|$. Thus by Fubini's theorem we have

$$|(C - D) \cdot A \cdot A \cdot A/A \cdot A| \lessapprox |A|. \tag{4}$$

Now we refine C and D. Since $|C|, |D| \approx |A|$ and $|A \cdot A| \approx |A|$, we have $|CD| \approx |C|, |D|$, and hence by the multiplicative form of Theorem 6.2, we can find subsets C', D' of C, D with $|C'|, |D'| \approx |A|$ such that every element in $C'D'$ has $\gtrapprox |A|^5$ representations in the form

$$\frac{c_1 d_1 c_3 d_3}{c_2 d_2}; \quad c_1, c_2, c_3 \in C; \quad d_1, d_2, d_3 \in D.$$

Now let $c, c' \in C'$ and $d, d' \in D$ be arbitrary. By the pigeonhole principle there thus exist $c_2 \in C$, $d_2 \in D$ such that we have $\gtrapprox |A|^3$ solutions to the problem

$$cd = \frac{c_1 d_1 c_3 d_3}{c_2 d_2}; \quad c_1, c_3 \in C; \quad d_1, d_3 \in D.$$

We can rewrite this as

$$cd - c'd' = x_1 - x_2 + x_3 + x_4$$

where

$$\begin{aligned} x_1 &= \frac{(c_1 - d')d_1 c_3 d_3}{c_2 d_2} \\ x_2 &= \frac{d'(c' - d_1)c_3 d_3}{c_2 d_2} \\ x_3 &= \frac{d'c'(c_3 - d_2)d_3}{c_2 d_2} \\ x_4 &= \frac{d'c'd_2(c_2 - d_3)}{c_2 d_2}. \end{aligned}$$

For fixed c, d, c', d', c_2, d_2, it is easy to see that the map from (c_1, c_3, d_1, d_3) to (x_1, x_2, x_3, x_4) is a bijection. Since all the x_j lie in $(C - D) \cdot A \cdot A \cdot A/A \cdot A$, we thus have $\gtrapprox |A|^3$ ways to represent $cd - c'd'$ in the form $x_1 - x_2 + x_3 - x_4$, where x_1, x_2, x_3, x_4 all lie in $(C - D) \cdot A \cdot A \cdot A/A \cdot A$. By (4) and Fubini's theorem we thus have

$$|C'D' - C'D'| \lessapprox |A|.$$

In particular we have $|C'D'| \lessapprox |A| \lessapprox |C'|$, which by the multiplicative form of sumset estimates implies $|C'/D'| \approx |C'|$. By considering the fibers of the quotient map $(x,y) \to x/y$ on $C' \times D'$ and using the pigeonhole principle, we thus see that there must be a non-zero field element x such that $|C' \cap D'x| \approx |A|$. If we then set $A' := C' \cap D'x$ we have $|A'A' - A'A'| \lessapprox |A|$ as desired.

7. Second ingredient: Iterated sum and product set estimates

We now prove the following lemma, which is in the spirit of sum-set estimates:

Lemma 7.1. *Let A be a non-empty subset of a finite field F, and suppose that we have the bound*

$$|A \cdot A - A \cdot A| \leq K|A|$$

for some $K \geq 1$. We adopt the normalization that $1 \in A$. Then for any polynomial P of several variables and integer coefficients, we have

$$|P(A, A, \ldots, A)| \leq CK^C|A|$$

where the constants C depend of course on P.

Proof We need some notation. We say that a set A is *essentially contained* in B, and write $A \Subset B$, if we have $A \subseteq X + B$ for some set X of cardinality $|X| \leq CK^C$.

We have the following simple lemma of Ruzsa [14]:

Lemma 7.2. *Let A and B be subsets of F such that $|A+B| \leq CK^C|A|$ or $|A-B| \leq CK^C|A|$. Then $B \Subset A - A$.*

Proof By symmetry we may assume that $|A+B| \leq CK^C|A|$. Let X be a maximal subset of B with the property that the sets $\{x + A : x \in X\}$ are all disjoint. Since the sets $x + A$ all have cardinality $|A|$ and are all contained in $A + B$, we see from disjointness that $|X||A| \leq |A + B|$, and hence $|X| \leq CK^C$. Since the set X is maximal, we see that for every $b \in B$, the set $b + A$ must intersect $x + A$ for some $x \in X$. Thus $b \in x + A - A$, and hence $B \subseteq X + A - A$ as desired. ∎

Call an element $x \in F$ *good* if we have $x \cdot A \Subset A - A$.

Proposition 7.3. *The following three statements are true.*

- *Every element of A is good.*
- *If x and y are good, then $x + y$ and $x - y$ is good.*
- *If x and y are good, then xy is good.*

(Of course, the implicit constants in "good" vary at each occurence).

Proof Let us first show that every element of A is good. Since $1 \in A$, we have

$$|A.A - A| \leq |A.A - A.A| \leq K|A|$$

and hence by Lemma 7.2

$$A.A \Subset A - A \tag{5}$$

which implies in particular that every element of x is good.

Now suppose that x and y are good, thus $x \cdot A \Subset A - A$ and $y \cdot A \Subset A - A$. Then

$$(x+y) \cdot A \subseteq x \cdot A + y \cdot A \Subset A - A + A - A.$$

On the other hand, since $|A - A| \leq |A.A - A.A| \leq K|A|$, we have from sumset estimates that

$$|A - A + A - A + A| \leq CK^C|A|$$

and hence by Lemma 7.2

$$A - A + A - A \Subset A - A. \tag{6}$$

Thus by transitivity of $\Subset$ we have $(x+y) \cdot A \Subset A - A$ and hence $x + y$ is good. A similar argument shows that $x - y$ is good.

Now we need to show that xy is good. Since $x \cdot A \Subset A - A$ we have

$$xy \cdot A \Subset y \cdot A - y \cdot A.$$

But since $y \cdot A \Subset A - A$, we have

$$xy \cdot A \Subset A - A - A + A.$$

By (6) we conclude that xy is good. ∎

By iterating this proposition we see that for any integer polynomial P, every element of $P(A, \ldots, A)$ is good[1].

Write $A^2 := A \cdot A$, $A^3 := A \cdot A \cdot A$, etc. We now claim inductively that $A^k \Subset A - A$ for all $k = 0, 1, 2, 3, \ldots$. The case $k = 0, 1$ are trivial, and $k = 2$ has already been covered by (5). Now suppose inductively that $k > 2$, and that we have already proven that $A^{k-1} \Subset A - A$. Thus

$$A^{k-1} \subseteq X + A - A$$

for some set X of cardinality $|X| \leq CK^C$. Clearly we may restrict X to the set $A^{k-1} - (A - A)$. In particular, every element of X is good. We now multiply by A to obtain

$$A^k \subseteq X \cdot A + A \cdot A - A \cdot A.$$

Since every element of X is good, and $|X| \leq CK^C$, we see that $X \cdot A \Subset A - A$. By (5) we thus have

$$A^k \Subset A - A + A - A - (A - A).$$

But by arguing as in the proof of (6) we have

$$A - A + A - A - (A - A) \Subset A - A,$$

and thus we can close the induction.

[1] An alternate way to proceed at this point is to show that the number of good points is at most $\lessapprox N$; indeed, it is easy to show that any good point is contained inside $(A - A + A - A)/(A - A)$ if N is sufficiently large, where we exclude 0 from the denominator $A - A$ of course. We omit the details.

Since $A^k \Subset A - A$ for every k, and $A - A \pm (A - A) \Subset A - A$ by (6), we thus see that every integer combination of A^k is essentially contained in $A - A$. In particular $P(A, \dots, A) \Subset A - A$ for every integer polynomial A, and the claim follows. ■

8. Third ingredient: Expanding sum estimates

The third and final ingredient is a statement of the form that $|A + B|$ can become much larger than $|A|$ or $|B|$ separately. We first recall the Cauchy-Davenport inequality

$$|A + B| \geq \min(|A| + |B| - 1, |F|) \tag{7}$$

for any non-empty subsets A, B of F. If we are allowed to arbitrarily dilate one of the sets A, B then we can improve subtantially on this inequality:

Lemma 8.1. *Let A, B be finite non-empty subsets of a finite field F, and lte $F^* := F - \{0\}$ denote the invertible elements of F. Then there exists $\xi \in F^*$ such that*

$$|A + B\xi| \geq \min(\frac{1}{2}|A||B|, \frac{1}{10}|F|). \tag{8}$$

Proof We may assume without loss of generality that $|A||B| \leq \frac{1}{2}|F|$, since if $|A||B| > \frac{1}{2}|F|$ we may remove some elements from A and B without affecting the right-hand side of (8). Let ξ be an element of F^*. We use the inclusion-exclusion principle[2] and the invertibility of ξ to compute

$$\begin{aligned}
|A + B\xi| &= |\bigcup_{a \in A} a + B\xi| \\
&\geq \sum_{a \in A} |a + B\xi| - \frac{1}{2} \sum_{a,a' \in A: a \neq a'} |(a + B\xi) \cap (a' + B\xi)| \\
&\geq \sum_{a \in A} |B| - \frac{1}{2} \sum_{a,a' \in A: a \neq a'} \sum_{b,b' \in B} \delta_{a+b\xi, a'+b'\xi} \\
&= |A||B| - \frac{1}{2} \sum_{a,a' \in A: a \neq a'} \sum_{b,b' \in B: b \neq b'} \delta_{\xi, (a-a')/(b-b')},
\end{aligned}$$

[2]To verify our use of the principle, suppose an element x lies in N of the sets $a+B\xi$ for some $N \geq 1$. Then the sum $\sum_{a \in A} |a+B\xi|$ counts x N times, while the sum $\sum_{a,a' \in A: a \neq a'} |(a+B\xi) \cap (a'+B\xi)|$ counts x $N(N-1)$ times. Since $N - \frac{N(N-1)}{2}$ is always less than or equal to 1, the claim follows. An alternate way to obtain this lemma (which gives slightly worse bounds when $|A||B| \ll |F|$, but somewhat better bounds when $|A||B| \gg |F|$) is by using the Cauchy-Schwarz inequality $\|\chi_A * \chi_{B\xi}\|_{l^1}^2 \leq \|\chi_A * \chi_{B\xi}\|_{l^2} |A + B\xi|$ and again randomizing over ξ.

where $\delta_{i,j}$ is the Kronecker delta function. If we average this over all $\xi \in F_*$ we obtain

$$\begin{aligned}\frac{1}{|F_*|}\sum_{\xi\in F_*}|A+B\xi| &\geq |A||B| - \frac{1}{2}\sum_{a,a'\in A: a\neq a'}\sum_{b,b'\in B: b\neq b'}\frac{1}{|F|-1}\\ &\geq |A||B| - \frac{1}{2}\frac{|A|^2|B|^2}{|F|-1}\\ &\geq \frac{1}{2}|A||B|\end{aligned}$$

by our hypothesis $|A||B| \leq \frac{1}{2}|F|$. The claim (8) then follows by the pigeonhole principle. ∎

9. Proof of Theorem 5.1

We now have all the machinery needed to prove Theorem 5.1. We basically follow the Edgar-Miller approach, see [6]. We write F for $\mathbf{Z}/q\mathbf{Z}$, and let $F^* := F - \{0\}$ be the invertible elements of F. Let $\delta > 0$, and let A be a subset of F such that $|F|^\delta < |A| < |F|^{1-\delta}$.

Let $0 < \varepsilon \ll 1$ be a small number depending on δ to be chosen later. In this section we use $X \lesssim Y$ to denote the estimate $X \leq C(\delta,\varepsilon)Y$ for some $C(\delta,\varepsilon) > 0$. Suppose for contradiction that

$$|A+A|, |A\cdot A| \lesssim |A|^{1+\varepsilon};$$

Then by Lemma 6.1, and passing to a refinement of A if necessary, we may assume that

$$|A\cdot A - A\cdot A| \lesssim |A|^{1+C\varepsilon}.$$

We may normalize $1 \in A$. By Lemma 7.1 we thus have

$$|P(A,\ldots,A)| \lesssim |A|^{1+C\varepsilon} \tag{9}$$

for any polynomial P with integer coefficients, where the constants C depend of course on P.

Our first objective is to obtain a linear surjection from A^k to F for sufficiently large k:

Lemma 9.1. *There exists a positive integer $k \sim 1/\delta$, and invertible field elements $\xi_1,\ldots,\xi_k \in F^*$, such that*

$$F = A\xi_1 + \cdots + A\xi_k.$$

In other words, we have a linear surjection from A^k to F.

Proof Iterating Lemma 8.1 about $O(1/\delta)$ times, we obtain $\xi_1,\ldots,\xi_k \in F^*$ such that

$$|A\xi_1 + \cdots + A\xi_k| \geq \frac{|F|}{10}.$$

The lemma then obtains after $O(1)$ applications of the Cauchy-Davenport inequality (7), increasing k as necessary. ∎

Next, we reduce the rank k of this surjection, at the cost of replacing A by a polynomial expression of A.

Lemma 9.2. *Let B be a non-empty subset of F, and suppose $k > 1$ is such that there is a linear surjection from B^k to F. Then there is a linear surjection from $\tilde{B}^{k-1}$ to F, where $\tilde{B} := B \cdot (B - B) + B \cdot (B - B)$.*

Proof By hypothesis, we have a surjection

$$B^k \to F : (a_1, \ldots, a_k) \mapsto \sum_{j \leq k} a_j \xi_j$$

for some $\xi_1, \ldots, \xi_k \in F$. Our map cannot be one-to-one, since otherwise

$$|B|^k = |F| \text{ (contradicting primarily of } |F|).$$

Thus there are $(b_1, \ldots, b_k) \neq (b'_1, \ldots b'_k) \in B^k$ with

$$(b_1 - b'_1)\xi_1 + \cdots + (b_k - b'_k)\xi_k = 0. \tag{10}$$

Let $b_k \neq b'_k$. By the surjection property

$$F = B\xi_1 + \ldots + B\xi_k;$$

since F is a field, we thus have

$$F = B\xi_1(b_k - b'_k) + \cdots + B\xi_k(b_k - b'_k)$$

and substituting $(b_k - b'_k)\xi_k$ from (10)

$$\begin{aligned} F &= B\xi_1(b_k - b'_k) + \cdots + B\xi_{k-1}(b_k - b'_k) - B(b_1 - b'_1)\xi_1 - \cdots - B(b_{k-1} - b'_{k-1})\xi_{k-1} \\ &\subset \tilde{B}\xi_1 + \cdots + \tilde{B}\xi_{k-1} \end{aligned}$$

and the claim follows. ∎

Starting with Lemma 9.1 and then iterating Lemma 9.2 k times, we eventually get a linear surjection from a polynomial expression $P(A, \ldots, A)$ of A to F, and thus

$$|P(A, \ldots, A)| \geq |F|.$$

But this contradicts (9), if ε is sufficiently small depending on δ. This contradiction proves Theorem 5.1. ∎

Remark. Suppose the finite field F did not have prime order. Then the analogue of Theorem 5.1 fails, since one can take A to be a subfield G of F, or a large subset of such a subfield G. It turns out that one can adapt the above argument to show that these are in fact the only ways in which Theorem 5.1 can fail (up to dilations, of course):

Theorem 9.3. *Let A be a subset of a finite field F such that $|A| > |F|^\delta$ for some $0 < \delta < 1$, and suppose that $|A + A|, |A \cdot A| \leq K|A|$ for some $K \gg 1$. Then there exists a subfield G of F of cardinality $|G| \leq K^{C(\delta)}|A|$, a non-zero field element $\xi \in F - \{0\}$, and a set $X \subseteq F$ of cardinality $|X| \leq K^{C(\delta)}$ such that $A \subseteq \xi G \cup X$.*

It is interesting to compare the above theorem to Freiman's theorem which does not assume control on $|A \cdot A|$ but has a dependence on constants which is significantly worse than polynomial. It seems possible that the constant $C(\delta)$ can be made independent of δ, but we do not know how to do so.

Proof (Sketch) Of course, we may assume that $|F| \geq K^{C(\delta)}$ for some large $C(\delta)$. We repeat the argument used to prove Theorem 5.1. This argument allows us to find a refinement A' of A with $|A'| \geq K^{-C}A$ such that $|A' \cdot A' - A' \cdot A'| \leq K^C|A|$. By dilating A and A' if necessary we may assume as before that $1 \in A'$ (as we shall see, this normalization allows us to take $\xi = 1$ in the conclusion of this theorem). By Lemma 7.1 we thus have $|P(A', \ldots, A')| \leq K^C|A'|$ for all integer polynomials P, with the constant C depending on P of course. We may assume $0 \in A'$ since adding 0 to A' and A do not significantly affect the above polynomial bounds.

We now claim that A' is contained in some subfield G of F of cardinality $|G| \leq K^{C(\delta)}|A|$. The argument in Lemma 9.1 still gives a surjection from $(A')^k$ to F for some $k \sim 1/\delta$. We then attempt to use Lemma 9.2 to drop the rank of this surjection down to 1. If we can reduce the rank all the way to one, then we have by arguing as before that $|F| \leq K^{C(k)}|A|$, so the claim follows by setting $G := F$. The only time we run into difficulty in this iteration is if we discover a linear surjection from some $\tilde{A}^{k'}$ to F with $k' > 1$ which is also injective, where $\tilde{A}$ is some polynomial expression of $P(A', \ldots, A')$. An inspection of the proof of Lemma 9.2, combined with the normalizations $0, 1 \in A'$, reveals that $\tilde{A}$ must contain A'. If we have $|\tilde{A} + \tilde{A}| > |\tilde{A}|$, then the linear map from $(\tilde{A} + \tilde{A})^{k'} \to F$ is surjective but not injective, which allows us to continue the iteration of Lemma 9.2. Similarly if $|\tilde{A} \cdot \tilde{A}| > |\tilde{A}|$. Thus the only remaining case is when $|\tilde{A}| = |\tilde{A} + \tilde{A}| = |\tilde{A} \cdot \tilde{A}|$. But this, combined with the fact that $0, 1 \in \tilde{A}$, implies that $\tilde{A} = \tilde{A} + \tilde{A} = \tilde{A} \cdot \tilde{A}$, and hence that $\tilde{A}$ is a subfield of F. Since $|\tilde{A}| \leq K^{C(k)}|A'|$, the claim follows. This shows that A' is a subset of G. Since $|A + A'| \leq K|A|$, we see from Lemma 7.2 that $A \Subset A' - A'$, and hence $A \Subset G$. Thus there exists a set Y of cardinality $|Y| \leq K^{C(\delta)}$ such that $A \subseteq G + Y$.

Let $y \in Y - G$. To finish the proof it will suffice to show that $|A \cap (G + y)| \leq K^{C(\delta)}$ for all such y. But observe that for any two distinct $x, x' \in G + y$, the sets xG and $x'G$ do not intersect except at the origin (for if $xg = x'g'$, then $g \neq g'$, and hence $x = (x' - x)\frac{g'}{g - g'} \in G$, contradicting the hypotheses that $x \in G + y$ and $y \notin G$). In particular, the sets $x(A' - \{0\})$ and $x'(A' - \{0\})$ are disjoint. Thus

$$K|A| \geq |A \cdot A| \geq |A \cap (G + y)||A' - \{0\}| \geq |A \cap (G + y)|K^{-C}|A|$$

and the claim follows. ∎

10. Some basic combinatorics

In later sections we shall use the sum-product estimate in Theorem 5.1 to various combinatorial problems in finite geometries. In doing so we will repeatedly use a number of basic combinatorial tools, which we collect here for reference.

We shall frequently use the following elementary observation: If B is a finite set, and $\mu : B \to \mathbf{R}^+$ is a function such that

$$\sum_{b \in B} \mu(b) \geq X,$$

then we have

$$\sum_{b \in B: \mu(b) \geq X/2|B|} \mu(b) \geq X/2.$$

We refer to this as a "popularity" argument, since we are restricting B to the values b which are "popular" in the sense that μ is large.

We shall frequently use the following version of the Cauchy-Schwarz inequality.

Lemma 10.1. *Let A, B be finite sets, and let $\sim$ be a relation connecting pairs $(a,b) \in A \times B$ such that*

$$|\{(a,b) \in A \times B : a \sim b\}| \gtrsim X$$

for some $X \gg |B|$. Then

$$|\{(a,a',b) \in A \times A \times B : a \neq a'; a, a' \sim b\}| \gtrsim \frac{X^2}{|B|}.$$

Proof Define for each $b \in B$, define $\mu(b) := |\{a \in A : a \sim b\}|$. Then by hypothesis we have

$$\sum_{b \in B} \mu(b) \gtrsim X.$$

In particular, by the popularity argument we have

$$\sum_{b \in B: \mu(b) \gtrsim X/|B|} \mu(b) \gtrsim X.$$

By hypothesis, we have $X/|B| \gg 1$. From this and the previous, we obtain

$$\sum_{b \in B: \mu(b) \gtrsim X/|B|} \mu(b)(\mu(b)-1) \gtrsim X(X/|B|)$$

and the claim follows. ∎

A typical application of the above Lemma is the standard incidence bound on lines in a plane F^2, where F is a finite field.

Corollary 10.2. *Let F^2 be a finite plane. For an arbitrarily collection $P \subseteq F^2$ of points and L of lines in F^2, we have*

$$|\{(p,l) \in P \times L : p \in l\}| \lesssim |P|^{1/2}|L| + |P| \tag{11}$$

Proof We may of course assume that the left-hand side of (11) is $\gg |P|$, since the claim is trivial otherwise. From Lemma 10.1 we have

$$|\{(p,l,l') \in P \times L \times L : p \in l \cap l'; l \neq l'\}| \gtrsim |P|^{-1}|\{(p,l) \in P \times L : p \in l\}|^2.$$

On the other hand, $|l \cap l'|$ has cardinality $O(1)$ if $l \neq l'$, thus

$$|\{(p,l,l') \in P \times L \times L : p \in l \cap l'; l \neq l'\}| \lesssim |L|^2.$$

Combining the two estimates we obtain the result. ∎

Note that this is markedly inferior to the Szemerédi-Trotter theorem.

11. Application: A Szemerédi-Trotter type theorem in finite fields

We now use the one-dimensional sum-product estimate to obtain a key two-dimensional estimate, namely an incidence bound of Szemerédi-Trotter type.

Let F be a finite field, and consider the projective finite plane PF^3, which is the set $F^3 - \{(0,0,0)\}$ quotiented by dilations. We embed the ordinary plane F^2 into PF^3 by identifying (x,y) with the equivalence class of $(x,y,1)$; PF^3 is thus F^2 union the line at infinity. Let $1 \leq N \leq |F|^2$ be an integer, and let P be a collection of points and L be a collection of lines in F^2. We consider the problem of obtaining an upper bound on the number of incidences

$$|\{(p,l) \in P \times L : p \in l\}|.$$

From Corollary 10.2 and the duality between points and lines in two dimensions we have the easy bounds

$$|\{(p,l) \in P \times L : p \in l\}| \leq \min(|P||L|^{1/2} + |L|, |L||P|^{1/2} + |P|), \tag{12}$$

see e.g. [2]. In a sense, this is sharp: if we set $N = |F|^2$, and let P be all the points in $F^2 \subset PF^3$ and L be most of the lines in F^2, then we have roughly $|F|^3 \sim N^{3/2}$ incidences. More generally if G is any subfield of F then one can construct a similar example with $N = |G|^2$, P being all the points in G^2, and L being the lines with slope and intercept in G.

We can use the sum-product estimate (Theorem 5.1) to obtain a non-trivial improvement to this:

Theorem 11.1. [4] *Let F be the finite field $F := \mathbf{Z}/q\mathbf{Z}$ for some prime q, and let P and L be points and lines in PF^3 with cardinality $|P|, |L| \leq N = |F|^\alpha$ for some $0 < \alpha < 2$. Then we have*

$$|\{(p,l) \in P \times L : p \in l\}| \leq CN^{3/2-\varepsilon}$$

for some $\varepsilon = \varepsilon(\alpha) > 0$ depending only on the exponent α.

Proof We may assume that $N \gg 1$ is large. By adding dummy points and lines we may assume that $|P| = |L| = N$.

Fix $N = |F|^\alpha$, and let $0 < \varepsilon \ll 1$, be chosen later. Suppose for contradiction that we can find points P and lines L with $|P| = |L| = N$ such that

$$|\{(p,l) \in P \times L : p \in l\}| \gtrsim N^{3/2-\varepsilon};$$

we shall use the sum-product estimates to obtain a contradiction if ε is sufficiently small. Our arguments follow those in [11], [17].

We first use the popularity argument to control how many points are incident to a line and vice versa. For each $p \in P$, define the multiplicity $\mu(p)$ at p by

$$\mu(p) := |\{l \in L : p \in l\}|.$$

Then by hypothesis

$$\sum_{p \in P} \mu(p) \gtrsim N^{3/2-\varepsilon}$$

and hence by the popularity argument and the hypothesis $|P| = N$

$$\sum_{p \in P : \mu(p) \gtrsim N^{1/2-\varepsilon}} \mu(p) \gtrsim N^{3/2-\varepsilon}.$$

On the other hand, we observe that

$$\begin{aligned}
\sum_{p \in P : \mu(p) \gg N^{1/2+\varepsilon}} \mu(p)^2 &\ll N^{-1/2-\varepsilon} \sum_{p \in P} \mu(p)(\mu(p)-1) \\
&\lesssim N^{-1/2-\varepsilon} \sum_{p \in P} |\{(l,l') \in L \times L : p \in l, l'; l \neq l'\}| \\
&= N^{-1/2-\varepsilon} \sum_{l,l' \in L : l \neq l'} |\{p \in P : p \in l, l'\}| \\
&\leq N^{-1/2-\varepsilon} \sum_{l,l' \in L : l \neq l'} 1 \\
&\leq N^{1/2+\varepsilon}.
\end{aligned}$$

Thus if we set $P' \subseteq P$ to be the set of all points p in P such that

$$N^{1/2-\varepsilon} \lesssim \mu(p) \lesssim N^{1/2+\varepsilon}$$

then we have

$$\sum_{p \in P'} \mu(p) \gtrsim N^{3/2-\varepsilon}.$$

For each $l \in L$, define the multiplicity $\lambda(l)$ by

$$\lambda(l) := |\{p \in P' : p \in l\}|,$$

then we can rewrite the previous as

$$\sum_{l \in L} \lambda(l) \gtrsim N^{3/2-\varepsilon}.$$

By the popularity argument we thus have

$$\sum_{l \in L : \lambda(l) \gtrsim N^{1/2-\varepsilon}} \lambda(l) \gtrsim N^{3/2-\varepsilon}.$$

On the other hand, we have

$$\begin{aligned}\sum_{l\in L:\lambda(l)\gg N^{1/2+\varepsilon}} \lambda(l) &\lesssim N^{-1/2-\varepsilon} \sum_{l\in L} \lambda(l)(\lambda(l)-1)\\ &\lesssim N^{-1/2-\varepsilon} \sum_{l\in L} |\{(p,p')\in P'\times P' : p,p'\in l; p\neq p'\}|\\ &= N^{-1/2-\varepsilon} \sum_{p,p'\in P':p\neq p'} |\{l\in L : p,p'\in l\}|\\ &\leq N^{-1/2-\varepsilon} \sum_{p,p'\in P':p\neq p'} 1.\end{aligned}$$

Thus if we set $L'\subset L$ to be the set of all lines l in L such that

$$N^{1/2-\varepsilon} \lesssim \lambda(p) \lesssim N^{1/2+\varepsilon}$$

then we have

$$\sum_{l\in L'} \lambda(l) \gtrsim N^{3/2-\varepsilon}.$$

For each $p\in P'$, let $\mu'(p)$ denote the multiplicity

$$\mu'(p) := |\{l\in L' : p\in l\}|;$$

clearly $\mu'(p)\leq \mu(p)$. We can then rewrite the previous estimate as

$$\sum_{p\in P'} \mu'(p) \gtrsim N^{3/2-\varepsilon}.$$

Thus by the popularity argument, if we set $P''\subseteq P'$ to be the set of all points p in P' such that

$$\mu'(p) \gtrsim N^{1/2-\varepsilon}$$

then we have

$$\sum_{p\in P''} \mu'(p) \gtrsim N^{3/2-\varepsilon}.$$

or equivalently

$$|\{(p,l)\in P''\times L' : p\in l\}| \gtrsim C_0 N^{3/2-\varepsilon}.$$

Since $|L'|\leq N$, we have in particular that

$$|P''| \gtrsim N^{1/2-\varepsilon}. \tag{13}$$

The next step is to capture a large portion of the popular point set P' inside a Cartesian product $A\times B$, possibly after a projective transformation. The key observation is that such a product arises, modulo projective transformations, whenever one intersects two "bushes" of lines.

Let p_0 be any point in P''. Then by construction there are $\gtrsim N^{1/2-\varepsilon}$ lines l in L' containing p_0. Each of these lines l contains $\gtrsim N^{1/2-\varepsilon}$ points p in P'; of course, all but one of these are distinct from p_0. Thus we have

$$|\{(p,l)\in P'\times L' : p,p_0\in l, p\neq p_0\}| \gtrsim N^{1-2\varepsilon}.$$

Let us define a relation $\sim$ on P by defining $p \sim p'$ if $p \neq p'$ and there is a line in L' containing both p and p'. Since two distinct points determine at most one line, we thus have

$$|\{p \in P' : p \sim p_0\}| \gtrsim N^{1-2\varepsilon} \text{ for all } p_0 \in P''.$$

Summing this over all p_0 in P'', we obtain

$$|\{(p_0, p) \in P'' \times P' : p \sim p'\}| \gtrsim N^{1-2\varepsilon}|P''|.$$

Since $|P'| \leq N$, we thus see by Lemma 10.1 that

$$|\{(p_0, p_1, p) \in P'' \times P'' \times P' : p \sim p_0, p_1; p_0 \neq p_1\}| \gtrsim N^{1-C\varepsilon}|P''|^2.$$

By the pigeonhole principle, there thus exist distinct points $p_0, p_1 \in P''$ such that

$$|\{p \in P' : p \sim p_0, p_1\}| \gtrsim N^{1-C\varepsilon}. \tag{14}$$

Fix these p_0, p_1. By applying a projective linear transformation (which maps lines to lines and preserves incidence) we may assume that p_0, p_1 are both on the line at infinity. Indeed, we may assume that $p_0 = [(1,0,0)]$ and $p_1 = [(0,1,0)]$, where $[(x,y,z)]$ is the equivalence class of (x,y,z) in PF^3.

We first eliminate those points p in (14) on the line at infinity. Such points can only occur if the line at infinity is in L'. But then that line contains at most $O(N^{1/2+\varepsilon})$ points in P', by the definition of L'. Thus if ε is sufficiently small we have

$$|\{p \in P' \cap F^2 : p \sim [(1,0,0)], [(0,1,0)]\}| \gtrsim N^{1-C\varepsilon}.$$

Consider the lines in L' which pass through $[(1,0,0)]$. In the plane F^2, these lines be horizontal, i.e. they are of the form $\{(x,y) \in F^2 : y = b\}$ for some $b \in F$. Let $B \subseteq F$ denote the set of all such b. Since each line contains at least $cN^{1/2-\varepsilon}$ points in P', and $|P'| \leq N$, we know that $|B| \lesssim N^{1/2+\varepsilon}$. Similarly the lines in L' which pass through $[(0,1,0)]$ must in F^2 be vertical lines of the form $\{(x,y) \in F^2 : x = a\}$ for $a \in A$, where $|A| \leq CN^{1/2+\varepsilon}$. We thus have

$$|P' \cap (A \times B)| \gtrsim N^{1-C\varepsilon} \tag{15}$$

and

$$|A|, |B| \leq CN^{1/2+\varepsilon} \tag{16}$$

Now that we have placed P' in a Cartesian grid, the next step is to exploit the form $y = mx + b$ of lines in F^2 to obtain some additive and multiplicative information on A and B.

Define $P_0 := P' \cap (A \times B)$. By definition of P' we have

$$|\{l \in L : p \in l\}| \gtrsim N^{1/2-\varepsilon} \text{ for all } p \in P_0;$$

summing over P_0 using (15) and rearranging, we obtain

$$|\{(p,l) \in P_0 \times L : p \in l\}| \gtrsim N^{3/2-C\varepsilon}.$$

Let L_0 be those lines in L which are not horizontal. Since horizontal lines can contribute at most $|P_0| \leq N$ incidences to the above expression, we have (if ε is sufficiently large)

$$|\{(p,l) \in P_0 \times L_0 : p \in l\}| \gtrsim N^{3/2-C\varepsilon}.$$

By the popularity argument, if we let L_1 denote those lines in L_0 such that

$$|\{p \in P_0 : p \in l\}| \gtrsim N^{1/2-C\varepsilon}$$

we thus have

$$|\{(p,l) \in P_0 \times L_1 : p \in l\}| \gtrsim N^{3/2-C\varepsilon}$$

if the implicit constants are chosen appropriately.

Define a relation $\sim$ between B and L_1 by defining $b \sim l$ if there is a point p in the row $P_0 \cap (A \times \{b\})$ such that $p \in l$. Note that such a point p is unique since l is not horizontal, and thus

$$|\{(b,l) \in B \times L_1 : b \sim l\}| \gtrsim N^{3/2-C\varepsilon}.$$

By Lemma 10.1, we thus have

$$|\{(b,b',l) \in B \times B \times L_1 : b, b' \sim l\}| \gtrsim N^{2-C\varepsilon}.$$

By (16) and the pigeonhole principle, we thus conclude that there exists distinct heights $b, b' \in B$ such that

$$|\{l \in L_1 : b, b' \sim l\}| \gtrsim N^{1-C\varepsilon}.$$

Fix this b, b'. By an affine transformation of the vertical variable (which does not affect the line at infinity) we may assume that $b = 0$ and $b' = 1$. Since each line $l \in L_1$ contains $\gtrsim N^{1/2-C\varepsilon}$ points (x,t) in P_0, and hence in $A \times B$, and most of these have $t \neq 0, 1$ since l is not horizontal, we have

$$|\{(x,t,l) \in A \times B \times L_1 : 0, 1 \sim l; (x,t) \in l; t \neq 0, 1\}| \gtrsim N^{3/2-C\varepsilon}.$$

By definition of the relation $a \sim l$, we thus have

$$|\{(x,t,l,x_0,x_1) \in A \times B \times L_1 \times A \times A : (x_0,0), (x,t), (x_1,1) \in l; t \neq 0,1\}| \gtrsim N^{3/2-C\varepsilon}.$$

Since the three points $(x_0, 0)$, (x,t), $(x_1, 1)$ determine l, and

$$x = x_0 + (x_1 - x_0)t,$$

we thus have

$$|\{(t,x_0,x_1) \in B \times A \times A : (1-t)x_0 + tx_1 \in A; t \neq 0,1\}| \gtrsim N^{3/2-C\varepsilon}. \tag{17}$$

Note that this is somewhat similar to saying that $(1-B).A + B.A \subseteq A$, so we are getting close to being able to apply our sum-product estimate. But first we must perform some Balog-Szemerédi type refinements.

Let $A' \subseteq A$ denote those x_1 in A for which

$$|\{(t,x_0) \in B \times A \times A : (1-t)x_0 + tx_1 \in A; t \neq 0,1\}| \gtrsim N^{1-C\varepsilon}.$$

From (16), (17) and the popularity argument we have

$$|\{(t,x_0,x_1) \in B \times A \times A' : (1-t)x_0 + tx_1 \in A; t \neq 0,1\}| \gtrsim N^{3/2-C\varepsilon} \tag{18}$$

if the implicit constants are chosen correctly.

In particular, from (16) again we have

$$|A'| \gtrsim N^{3/2-C\varepsilon}/|A||B| \gtrsim N^{1/2-C\varepsilon}. \tag{19}$$

Also, by (18), the pigeonhole principle and (16) we may find $t_0 \in B$ such that $t_0 \neq 0,1$ and

$$|\{(x_0,x_1) \in A \times A' : (1-t_0)x_0 + t_0x_1 \in A\}| \gtrsim N^{1/2-C\varepsilon}$$

By (16) we have

$$|\{(x_0,x_1) \in A \times A' : (1-t_0)x_0 + t_0x_1 \in A\}| \gtrsim N^{-C\varepsilon}|A||A'|.$$

By (19), (16) and Theorem 6.2 applied to the sets $(1-t_0)A$ and t_0A', we thus have a subsets $(1-t_0)\tilde{A}$ of $(1-t_0)A$ and t_0A'' of t_0A' with cardinalities at least $\gtrsim N^{1/2-C\varepsilon}$ such that

$$|(1-t_0)\tilde{A} + t_0A'| \lesssim N^{1/2+C\varepsilon}.$$

By (19), (16) and sumset estimates, this implies in particular that

$$|t_0A' + t_0A'| \lesssim N^{1/2+C\varepsilon}$$

and hence

$$|A' + A'| \lesssim N^{1/2+C\varepsilon}. \tag{20}$$

Now we return to (18). From (16) and the pigeonhole principle we may find an $x_0 \in A$ such that

$$|\{(t,x_1) \in B \times A' : (1-t)x_0 + tx_1 \in A; t \neq 0,1\}| \gtrsim N^{1-C\varepsilon}.$$

By a translation in the horizontal variables x_0, x_1, A, A' we may assume that $x_0 = 0$. Thus

$$|\{(t,x_1) \in (B\backslash\{0\}) \times (A'\backslash\{0\}) : tx_1 \in A\}| \gtrsim N^{1-C\varepsilon},$$

since the contribution of 0 is easily controlled by (16). By (16) and the multiplicative form of Theorem 6.2, we can thus find a subset A'' of $A'\backslash\{0\}$ with $|A''| \gtrsim N^{1/2-C\varepsilon}$ and

$$|A'' \cdot A''| \lesssim N^{1/2+C\varepsilon}.$$

On the other hand, from (20) we have

$$|A'' + A''| \lesssim N^{1/2+C\varepsilon}.$$

But this gives a contradiction to the sum product estimate (Theorem 5.1) if ε is sufficiently small. ∎

12. Application to the distance set problem

We now work in the finite field plane F^2. Given any two points (x_1, y_1), (x_2, y_2), we define the distance $d((x_1, y_1), (x_2, y_2)) \in F$ by

$$d((x_1, y_1), (x_2, y_2)) = (x_1 - x_2)^2 + (y_1 - y_2)^2$$

(we omit the square root to avoid some distracting technicalities). Given any collection P of points in F^2, we define the distance set $\Delta(P) \subseteq F$ by

$$\Delta(P) := \{d(p, p') : p, p' \in P\}.$$

The *Erdös distance problem* is to obtain the best possible lower bound for $|\Delta(P)|$ in terms of $|P|$. If -1 is a square, thus $i^2 = -1$ for some $i \in F$, then the set $P := \{(x, ix) : x \in F\}$ has cardinality $|P| = |F|$ but $\Delta(P) = \{0\}$ has cardinality 1. To avoid this degenerate case[3] we assume that -1 is not a square, so any two distinct points have a non-zero distance. From the fact that any two "circles" intersect in at most two points, it is possible to use extremal graph theory to obtain the bound

$$|\Delta(P)| \geq c|P|^{1/2};$$

see also [8]. This bound is sharp if one takes $P = F^2$, so that $\Delta(P)$ is essentially all of F. Similarly if one takes $P = G^2$ for any subfield G of F. However, as in the previous section one can hope to improve this bound when no subfields are available.

From the obvious identity

$$\Delta(A \times A) = (A - A)^2 + (A - A)^2$$

it is clear that this problem has some connection to the sum-product estimate. Indeed, any improvement to the trivial bound on $|\Delta(P)|$ can be used (in combination with Lemma 6.1 and Lemma 7.1) to obtain a bound of the form in Lemma 5.1. We now present the converse implication, using the sum-product bounds already obtained to derive a new bound on the distance problem.

Theorem 12.1. [4] *Let $F = \mathbf{Z}/p\mathbf{Z}$ for some prime $p = 3 \mod 4$ (so -1 is not a square), and let P be a subset of F^2 of cardinality $|P| = N = |F|^\alpha$ for some $0 < \alpha < 2$. Then we have*

$$|\Delta(P)| \gtrsim N^{1/2+\varepsilon}$$

for some $\varepsilon = \varepsilon(\alpha) > 0$.

Remark. In the Euclidean analogue to this problem, with N points in $\mathbf{R}^2$, it is conjectured [8] that the above estimate is true for all $\varepsilon < 1/2$. Currently, this is known for all $\varepsilon < \frac{4e}{5e-1} - \frac{1}{2} \approx 0.364$ [15]. However, the Euclidean results depend (among other things) on crossing number technology and thus do not seem to obviously extend to the finite field case.

Proof We shall exploit the Szemerédi-Trotter-type estimate in Theorem 11.1 in much the same way that the actual Szemerédi-Trotter theorem [16] was exploited in [5] for the Euclidean version of Erdös's distance problem, or how a Furstenburg

[3] We thank Alex Iosevich for pointing out the need to exclude this case.

set estimate was used in [11], [17] to imply a Falconer distance set problem result. The key geometric observation is that the set of points which are equidistant from two fixed points lie on a line (the perpendicular bisector of the two fixed points).

We may assume that $|F|$ and $|P|$ are large; in particular, we may assume that F has characteristic greater than 2. Fix N, and suppose for contradiction that

$$|\Delta(P)| \lesssim N^{1/2+\varepsilon}$$

for some small $0 < \varepsilon \ll 1$ to be chosen later. For any point $p \in P$, we clearly have the identity

$$\{(p', r) \in P \times \Delta(P) : d(p, p') = r\}| = |P| = N$$

so by Lemma 10.1

$$\{(p', p'', r) \in P \times P \times \Delta(P) : d(p, p') = d(p, p'') = r; p' \neq p''\}| \gtrsim N^{3/2-\varepsilon}.$$

We can of course eliminate the r variable:

$$\{(p', p'') \in P \times P : d(p, p') = d(p, p''); p' \neq p''\}| \gtrsim N^{3/2-\varepsilon}.$$

Summing this over all $p \in P$ and rearranging, we obtain

$$\sum_{p', p'' \in P : p' \neq p''} |\{p \in P : d(p, p') = d(p, p'')\}| \geq cN^{5/2-\varepsilon}.$$

By the pigeonhole principle, there thus exists $p_0 \in P$ such that

$$\sum_{p' \in P : p' \neq p_0} |\{p \in P : d(p, p') = d(p, p_0)\}| \geq cN^{3/2-\varepsilon}.$$

By translation invariance we may take $p_0 = (0, 0)$. Writing $p' = (a, b)$ and $p = (x, y)$, this becomes

$$\sum_{(a,b) \in P : (a,b) \neq (0,0)} |\{(x, y) \in P : (x - a)^2 + (y - b)^2 = x^2 + y^2\}| \geq cN^{3/2-\varepsilon}.$$

Thus if we let $l(a, b)$ denote the perpendicular bisector of $(0, 0)$ and (a, b):

$$l(a, b) := \{(x, y) \in F^2 : (x-a)^2+(y-b)^2 = x^2+y^2\} = \{(x, y) \in F^2 : 2ax+2by = a^2+b^2\}$$

and let L be the collection of lines $\{l(a, b) : (a, b) \in P \backslash \{0, 0\}\}$, then we have

$$\{(p, l) \in P \times L : p \in l\} \geq cN^{3/2-\varepsilon}.$$

But since all the lines $l(a, b)$ are distinct, we have $|L| = N - 1$, while $|P| = N$. Thus this clearly contradicts Theorem 11.1, and we are done. ∎

The sum-product estimate has a number of other applications, for instance it gives the best known bound on the Kakeya problem for finite fields in three dimensions. We will not detail this here, but refer the interested reader to [4].

13. Exercises

- Q1. Show that (1) is sharp in the following sense: given any $|E| \geq 5|V|$, one can find a graph G with $cr(G) \sim |E|^3/|V|^2$. (Hint: consider n equally spaced points in a circle, and consider the drawing formed by connecting only those points which have at most k points between them using straight line segments).
- Q2. Find a set of N^3 points and N^3 lines which have N^4 incidences between them. (Hint: use the grid $\{1, \ldots, N\} \times \{1, \ldots, N^2\}$). Conclude that the bound (2) is sharp in the case $|L| = |P|$. (For an additional challenge: can you show it is sharp in the general case?)

References

[1] A. Balog, E. Szemerédi, *A statistical theorem of set addition*, Combinatorica, **14** (1994), 263–268.

[2] B. Bollobas, *Extremal Graph Theory*, Academic Press, London 1978.

[3] J. Bourgain, *On the dimension of Kakeya sets and related maximal inequalities*, Geom. Funct. Anal. **9** (1999), no. 2, 256–282.

[4] J. Bourgain, N. Katz, T. Tao, *A sum-product estimate in finite fields, and applications*, to appear, GAFA.

[5] F. Chung, E. Szemerédi, W. T. Trotter Jr., *On the number of different distances*, Discrete and Computational Geometry **7** (1992), 1–11.

[6] G. A. Edgar, C. Miller, *Borel subrings of the reals*, Proc. Amer. Math. Soc. 131 (2003), 1121–1129.

[7] G. Elekes, *On the number of sums and products*, Acta Arith. **81** (1997), 365–367.

[8] P. Erdös, *On sets of distances of n points*, American Mathematical Monthly, **53** (1946), 248–250.

[9] P. Erdös, E. Szemerédi, *On sums and products of integers*, Studies in Pure Mathematics, 213–218. Birkhaüser, Basel, 1983.

[10] T. Gowers, *A new proof of Szemerédi's theorem for arithmetic progressions of length four*, Geom. Funct. Anal. **8** (1998), no. 3, 529–551.

[11] N. Katz, T. Tao, *Some connections between the Falconer and Furstenburg conjectures*, New York J. Math., **7** (2001), 148–187.

[12] M. Nathanson, *Additive number theory. Inverse problems and the geometry of sumsets.* Graduate Texts in Mathematics, 165. Springer-Verlag, New York, 1996.

[13] I. Ruzsa, *Sums of finite sets*, Number Theory: New York Seminar; Springer-Verlag (1996), D.V. Chudnovsky, G.V. Chudnovsky and M.B. Nathanson editors.

[14] I. Ruzsa, *An analog of Freiman's theorem in groups*, Structure theory of set addition, Astérisque No. 258 (1999), 323–326.

[15] J. Solymosi, G. Tardos, C. D. Tóth, *Distinct distances in the plane*, Discrete Comput. Geom. **25** (4) (2001), 629–634.

[16] E. Szemerédi, W. T. Trotter Jr., *Extremal problems in discrete geometry*, Combinatorica **3** (1983), 381–392.

[17] T. Tao, *Finite field analogues of the Erdös, Falconer, and Furstenburg problems*, unpublished.

Department of Mathematics, UCLA, Los Angeles CA 90095-1555

E-mail address: tao@math.ucla.edu