ADDITIVE COMBINATORICS: WINTER 2007

K. Soundararajan

Introduction

The aim of this course is to study additive problems in number theory. Broadly, given a sufficiently large set of integers A (or more generally a subset of some abelian group) we are interested in understanding additive patterns that appear in A. An important example is whether A contains non-trivial arithmetic progressions of some given length k. One reason for considering arithmetic progressions is that they are quite indestructible structures: they are preserved under translations and dilations of A, and they cannot be excluded for trivial congruence reasons. For example the pattern a, b and $a+b$ all being in the set seems quite close the arithmetic progression case a, b, $(a+b)/2$, but the former case can never occur in any subset of the odd integers (and such subsets can be very large). Another class of questions we can ask is whether all numbers can be written as a sum of s elements from a given set A. For example, all numbers are sums of four squares, nine cubes etc. Waring's problem and the Goldbach conjectures are two classical examples. In the same spirit, given a set A of N integers we may ask for information about the sumset $A+A := \{a+b : a, b \in A\}$. If there are not too many coincidences, then we may expect $|A+A| \gg N^2$. But when A is an AP note that $|A+A| \leq 2|A|-1$. One of our goals for the class will be Freiman's theorem that if the sumset is small then A looks like a "generalized arithmetic progression."

The subject may be said to begin with a beautiful result of van der Waerden (1927).

van der Waerden's Theorem. *Let k and r be given. There exists a number $N = N(k,r)$ such that if the integers in $[1,N]$ are colored using r colors, then there is a non-trivial monochromatic k term arithmetic progression.*

van der Waerden's proof was by an ingenious elementary induction argument on k and r. The proof does not give any good bound on how large $N(k,r)$ needs to be. A more general result was subsequently found by Hales and Jewett (1963), with a nice refinement of Shelah (1988), but again the bounds for the van der Waerden numbers are quite poor.

The Hales-Jewett Theorem. *Let k and r be given. There exists a number $N = N(k,r)$ such that if the points in $[1,k]^N$ are colored using r colors then there is a monochromatic "combinatorial line". Here a combinatorial line is a collection of k points of the following type: certain of the coordinates are fixed, and a certain non-empty set of coordinates are designated as "wildcards" taking all the values from 1 to k.*

A picturesque way of describing the Hales-Jewett theorem is that a "tic-tac-toe" game of getting k in a row, played by r players, always has a result in sufficiently high dimensions. Since there is obviously no disadvantage to going first, the first player wins; but no constructive strategy solving the game is known. One can recover van der Waerden's theorem by thinking of $[1,k]^N$ as giving the base k digits (shifted by 1) of numbers in $[0, k^N - 1]$.

Erdős and Turan proposed a stronger form of the van der Waerden, partly in the hope that the solution to the stronger problem would lead to a better version of van der Waerden's theorem.

The Erdős-Turán conjecture. *Let δ and k be given. There is a number $N = N(k,\delta)$ such that any set $A \subset [1,N]$ with $|A| \geq \delta N$ contains a non-trivial arithmetic progression*

of length k.

In 1953, Roth proved the Erdős-Turán conjecture in the case $k = 3$.

Roth's Theorem. *There exists a positive constant C such that if $A \subset [1, N]$ with $|A| \geq CN/\log\log N$ then A has a non-trivial three term AP.*

In other words, $N(\delta, 3) \leq \exp(\exp(C/\delta))$ for some positive constant C. This stronger result does in fact give a good bound on the van der Waerden numbers for $k = 3$. We know now thanks to Bourgain that $|A| \gg N(\log\log N/\log N)^{1/2}$ suffices. Thus the double exponential bound can be replaced by a single exponential.

Let $r_3(N)$ denote the size of the largest subset of $[1, N]$ having no non-trivial three term APs. Then as mentioned above, $r_3(N) \ll N\sqrt{\log\log N/\log N}$. What is the true nature of $r_3(N)$? If we pick a random set A in $[1, N]$ we may expect that it has about $|A|^3/N$ three term APs. This suggests that $r_3(N)$ is perhaps of size $N^{1/3}$. However, in 1946 Behrend found an ingenious construction that does much much better.

Behrend's Theorem. *There exists a set $A \subset [1, N]$ with $|A| \gg B\exp(-c\sqrt{\log N})$ containing no non-trivial three term arithmetic progressions. In other words $r_3(N) \gg N\exp(-c\sqrt{\log N})$.*

Roth's proof is based on Fourier analysis. It falls naturally into two parts: either the set A looks random in which case we may easily count the number of three term progressions, or the set has some structure which can be exploited to find a subset with increased density. The crucial point is that the idea of randomness here can be made precise in terms of the size of the Fourier coefficients of the set. This argument is quite hard to generalize to four term progressions (or longer), and was only extended recently with the spectacular work of Gowers.

Returning to the Erdős-Turán conjecture, the next big breakthrough was made by Szemerédi who in 1969 established the case $k = 4$, and in 1975 dealt with the general case $k \geq 5$. His proof was a tour-de-force of extremely ingenious and difficult combinatorics. One of his ingredients was van der Waerden's theorem, and so this did not lead to a good bound there.

Szemerédi's Theorem. *Given k and $\delta > 0$, there exists $N = N(k, \delta)$ such that any set $A \subset [1, N]$ with $|A| \geq \delta N$ contains a non-trivial k term arithmetic progression.*

An entirely different approach was opened by the work of Furstenberg (1977) who used ergodic theoretic methods to obtain a new proof of Szemerédi's theorem. The ergodic theoretic approach also did not lead to any good bounds, but was useful in proving other results previously inaccessible. For example, it led to a multi-dimensional version of Szemerédi's theorem, also a density version of the Hales-Jewett theorem (due to Katznelson and Ornstein), and also allowed for the common difference of the APs to have special shapes (e.g. squares).

In 1998-2001 Gowers made a major breakthrough by extending Roth's harmonic analysis techniques to prove Szemerédi's theorem. This approach finally gave good bounds for the van der Waerden numbers.

Gowers's Theorem. *There exists a positive constant c_k such that any subset A in $[1, N]$ with $|A| \gg N/(\log\log N)^{c_k}$ contains a non-trivial k term arithmetic progression.*

In this course, we hope to give an account of Gowers's proof in the case $k = 4$. One of the major insights of Gowers is the development of a "quadratic theory of Fourier analysis" which substitutes for the "linear Fourier analysis" used in Roth's theorem. Gowers's ideas have transformed the field, opening the door to many spectacular results, most notably the work of Green and Tao.

The Green-Tao Theorem (2003). *The primes contain arbitrarily long non-trivial arithmetic progressions.*

Note that up to N there are about $N/\log N$ primes. This density is much smaller than what would be covered by Gowers's theorem; even in the case $k = 3$ it is not covered by the best known results on $r_3(N)$. We will not be able to cover the Green-Tao theorem, but will give some of the ideas in the simple case $k = 3$. Another result along those lines is the celebrated three primes theorem.

Vinogradov's theorem (1937). *Every large odd number is the sum of three primes.*

Another brilliant result of Green and Tao, developing Gowers's ideas, is that $r_4(N) \ll N(\log N)^{-c}$ where $r_4(N)$ denotes the largest cardinality of a set in $[1, N]$ containing no four term progressions.

Another theme that we shall explore, and which also plays an important role in Gowers's proof, is Freiman's theorem on sumsets. If A is a set of N integers then $A + A$ is bounded above by $N(N + 1)/2$, and below by $2N - 1$. The lower bound is attained only when A is highly structured, and is an arithmetic progression of length N. Clearly if A is a subset of an arithmetic progression of length CN then $|A + A| \le 2C|A|$. More generally suppose $d_1, \ldots, d_k$ are given numbers, and consider the set

$$\{a_0 + a_1 d_1 + \ldots + a_k d_k : \quad 1 \le a_i \le N_i \text{ for } 1 \le i \le k\}.$$

We may think of this as a generalized arithmetic progression of dimension d. Note that this generalized AP has cardinality at most $N_1 \cdots N_k$. If these sums are all distinct (so that the cardinality equals $N_1 \cdots N_k$) we call the GAP *proper*. Note that if A is contained in a gAP of dimension k and size $\le CN$ then $|A+A| \le 2^k CN$. Freiman's theorem provides a converse to this showing that all sets with small sumsets must arise in this fashion.

Freiman's theorem. *If A is a set with $|A + A| \le C|A|$ then there exists a proper GAP of dimension k (bounded in terms of C) and size $\le C_1|A|$ for some constant C_1 depending only on C.*

Qualitatively Freiman's theorem says that any set with a small sumset looks like an arithmetic progression. Similarly we may expect that a set with a small product set should look like a geometric progression. But of course no set looks simultaneously like an arithmetic and a geometric progression! Thus we may surmise, as did Erdős and Szemerédi that either the sumset or the product set must be large.

Erdős-Szemerédi Conjecture. *If A is a set of N integers then*

$$|A + A| + |A \cdot A| \gg N^{2-\epsilon},$$

for any $\epsilon > 0$.

This is currently known for $\epsilon > 3/4$ (indeed a little better) thanks to results of Erdős-Szemerédi, Solymosi, Elekes ... . The sum-product theory (and its generalizations) is another very active problem in additive combinatorics, and has led to many important applications (bounding exponential sums etc).

We end this introduction by giving a brief description of how ergodic theory connects up with these combinatorial problems. The subject begins with a simple recurrence theorem of Poincaré.

Poincaré recurrence. *Let X be a probability space with measure μ, and let T be a measure preserving transformation (so $\mu(T^{-1}A) = \mu(A)$). For any set V with positive measure there exists a point $x \in V$ such that for some natural number n, $T^n x$ also is in V.*

Proof. This is very simple: note that the sets V, $T^{-1}V$, $T^{-2}V$, ... cannot all be disjoint. Therefore $T^{-m}V \cap T^{-m-n}V \neq \emptyset$ for some natural numbers m and n. But this gives readily that $V \cap T^n V \neq \emptyset$ as needed.

It is clear from the proof that the number n in Poincaré's result may be found below $1/\mu(V)$. As an example, we may take X to be the circle $\mathbb{R}/\mathbb{Z}$, and take V to be the interval $[-1/2Q, 1/2Q]$, and T to be the map $x \to x + \theta$ for some fixed number θ. We thus obtain:

Dirichlet's Theorem. *For any real number θ, and any $Q \geq 1$ there exists $1 \leq q \leq Q$ such that $\|q\theta\| \leq 1/Q$. Here $\|x\|$ denotes the distance between x and its nearest integer.*

If X happens also to be a separable (covered by countably many open sets) metric space, then we can divide X into countably many balls of radius $\epsilon/2$. Then it follows that almost every points of X returns to within ϵ of itself. That is, almost every point is *recurrent*.

We don't really need a probability space to find recurrence. Birkhoff realized that this can be achieved purely topologically and holds for compact metric spaces.

Birkhoff's Recurrence Theorem. *Let X be a compact metric space, and T be a continuous map. Then there exists a recurrent point in X; namely, a point x such that there is a sequence $n_k \to \infty$ with $T^{n_k} x \to x$.*

Proof. Since X is compact, any nested sequence of non-empty closed sets $Y_1 \supset Y_2 \supset Y_3 \ldots$ has a non-empty intersection. Consider T-invariant closed subsets of X; that is, Y with $TY \subset Y$. By Zorn's lemma and our observation above, there exists a non-empty minimal closed invariant set Y. Let y be any point in Y and consider the closure of y, Ty, $T^2 y$, ... This set is plainly a closed invariant subset of Y, and by minimality equals Y. Therefore y is recurrent.

These are some basic simple results, of the same depth as Dirichlet's pigeonhole principle and its application to Diophantine approximation. In the example of Diophantine approximation, we see that if $\|n\theta\|$ is small then so are $\|2n\theta\|$, $\|3n\theta\|$ etc. This suggests the possibility of *multiple recurrence*.

Topological Multiple Recurrence. *Let X be a compact metric space, and T be a continuous map. For any integer $k \geq 1$ there exists a point $x \in X$ and a sequence $n_\ell \to \infty$ with $T^{jn_\ell}x \to x$ for each $1 \leq j \leq k$.*

This theorem is analogous to van der Waerden's theorem, and indeed implies it. To see this, let $\Lambda = \{1, \dots, r\}$ represent r colors, and consider $\Omega = \Lambda^{\mathbb{Z}}$. Thus Ω is the space of all r colorings of the integers, and by $x \in \Omega$ we understand a particular r coloring of the integers. We make Ω into a compact metric space (check using sequential compactness), by taking as the metric $d(x, y) = 0$ if $x = y$ and $d(x, y) = 2^{-\ell}$ where ℓ is the least magnitude for which either $x(\ell) \neq y(\ell)$ or $x(-\ell) \neq y(-\ell)$. We define the shift map T by $Tx(n) = x(n + 1)$. Now suppose we are given a coloring ξ of the integers. Take X to be the closure of $T^n\xi$ where n ranges over all integers. By definition this is a closed invariant compact metric space, and so by the Topological Multiple Recurrence Theorem there is a $x \in X$ and some $n \in \mathbb{Z}$ with $x(0) = x(n) = x(2n) = \ldots = x(kn)$. But from the definition of the space X we may find an $m \in \mathbb{Z}$ such that $T^m\xi$ and x agree on the interval $[-kn, kn]$. Then it follows that $\xi(m) = \xi(m + n) = \ldots = \xi(m + kn)$ producing a $k + 1$ term AP.

The above argument gives an infinitary version of the van der Waerden theorem where we color all the integers. But from it we may deduce the finite version. Suppose not, and there are r colorings of $[-N, N]$ with no monochromatic k-APs for each natural number N. Extend each of these colorings arbitrarily to $\mathbb{Z}$, obtaining an element in Ω. By compactness we may find a limit point in Ω of these elements. That limit point defines a coloring of $\mathbb{Z}$ containing no monochromatic k-APs, and this is a contradiction.

The ergodic theoretic analog of Szemerédi's theorem is Furstenberg's multiple recurrence theorem for measure preserving transformations, and this implies Szemerédi by an argument similar to the one above.

Furstenberg's Theorem. *Let X be a probability measure space and let T be a measure preserving transformation. If V is a set of positive measure, then there exists a natural number n such that $V \cap T^{-n}V \cap T^{-2n}V \cap \ldots \cap T^{-kn}V$ has positive measure.*

The Hales-Jewett Theorem

We begin with a warm-up result, which although unrelated may help set the mood.

Schur's Theorem. *Given any positive number r, if $N \geq N(r)$ and the integers in $[1, N]$ are colored using r colors then there is a monochromatic solution to $x + y = z$.*

First we need a special case of Ramsey's theorem.

Lemma. *Suppose that the edges of the complete graph K_N are colored using r colors. If $N \geq N(r)$ then there is a monochromatic triangle.*

Proof. We will use induction on r. It is very well known that if $r = 2$ and $N \geq 6$ then there is a monochromatic triangle. Suppose we know the result for $r-1$ colorings, and we need $N \geq N(r-1)$ for that result. Pick a vertex. There are $N-1$ edges coming out of it. So for some color there are $\geq \lceil (N-1)/r \rceil$ edges starting from this vertex having the same color. Now the complete graph on the other vertices of these edges must be colored using only $r-1$ colors. Thus if $N \geq rN(r-1) - r + 2$ we are done.

Proof of Schur's Theorem. Consider the complete graph on N vertices labeled 1 through N. Color the edge joining a to b using the color of $|a - b|$. By our lemma, if N is large then there is a monochromatic triangle. Suppose its vertices are $a < b < c$ then $(c - a) = (c - b) + (b - a)$ is a solution proving Schur's theorem.

Let k and r be given natural numbers. Consider the cube $[1, k]^N$, and color each point in it using r colors. The Hales-Jewett theorem says that if N is sufficiently large then there will be a monochromatic line having k points. Here a (combinatorial) line means the following: Let $\mathbf{x} = (x_1, \ldots, x_N)$ be a point, and let A be a non-empty subset of $[1, N]$. By $\mathbf{x} \oplus jA$ (where $1 \leq j \leq k$) we denote the point $\mathbf{y}(j)$ whose coordinates are given by $y_i(j) = x_i$ if $i \notin A$ and $y_i(j) = j$ if $i \in A$. The line $\mathbf{x} \oplus A$ consists of the points $\mathbf{x} \oplus jA$ for $1 \leq j \leq k$. In other words, A describes a set of coordinates whose entries are wildcards taking all the values from 1 to k.

As a special case consider $k = 3$ and $r = 2$ which corresponds (essentially) to a game of tic-tac-toe. The Hales-Jewett theorem guarantees that in high dimension a game of tic-tac-toe never ends in a draw. Since the first person has a free move, and can steal any winning strategy that the second person devises, it follows that the first player should win such games.

We will now give two proofs of the Hales-Jewett theorem; the second, due to Shelah, being a small but very important modification of the first. The proofs both proceed by induction on k and r. Let $HJ(k, r)$ denote the least N for which the theorem holds; we wish to show that this is finite, and also derive some bounds for it. Note that if $k = 1$ there is nothing to prove and we may take $HJ(1, r) = 1$. Consider next the case that $k = 2$. Take $N = r$ and note that two of the $r+1$ points $(1, 1, \ldots, 1)$, $(1, 1, \ldots, 1, 2)$, $(1, \ldots, 2, 2)$, $\ldots$, $(1, 2, 2, \ldots, 2)$, $(2, 2, \ldots, 2)$ must have the same color. Thus $HJ(2, r) \leq r$. Exercise: show that $HJ(2, r) = r$.

First proof. We assume that $HJ(1, r)$, $\ldots$, $HJ(k-1, r)$ all exist, for all values of r. We now want to show that $HJ(k, r)$ exists. Let N_1, $\ldots$, N_r denote a very rapidly increasing sequence which we will specify below, and set $N = N_1 + \ldots + N_r$.

Consider $[1,k]^N$ as $[1,k]^{N_1} \times \ldots \times [1,k]^{N_r}$. Given a point $\mathbf{x}_r$ in $[1,k]^{N_r}$ we may define an r-coloring of $[1,k]^{N_1} \times \ldots \times [1,k]^{N_{r-1}}$ by setting the color of $(\mathbf{x}_1, \ldots, \mathbf{x}_{r-1})$ to be our original color for $(\mathbf{x}_1, \ldots, \mathbf{x}_r)$. Call this r-coloring of $[1,k]^{N_1} \times \ldots \times [1,k]^{N_{r-1}}$ as $\phi_{\mathbf{x}_r}$ say. The number of possibilities for $\phi_{\mathbf{x}_r}$ is naturally $r^{k^{N_1+\ldots+N_{r-1}}}$. We now view these possibilities as a palette for coloring $\mathbf{x}_r$ using $r^{k^{N_1+\ldots+N_{r-1}}}$ colors. If N_r is sufficiently large — precisely, $N_r \geq HJ(k-1, r^{k^{N_1+\ldots+N_{r-1}}})$ — then by induction hypothesis we may find a point $\mathbf{y}_r$ and a non-empty $A_r \subset [1, N_r]$ such that $\phi_{\mathbf{y}_r \oplus jA_r}$ is the same coloring for each $1 \leq j \leq k-1$. In other words, for any given choice of $\mathbf{x}_1$, ..., $\mathbf{x}_{r-1}$, the color of $(\mathbf{x}_1, \ldots, \mathbf{x}_{r-1}, \mathbf{y}_r \oplus jA_r)$ does not change as j varies from 1 to $k-1$.

We now want to repeat the same argument for $\mathbf{x}_{r-1}$. Given $\mathbf{x}_{r-1}$ we have an r-coloring of $[1,k]^{N_1} \times \ldots [1,k]^{N_{r-2}} \times [1,k]$ by setting the color of $(\mathbf{x}_1, \ldots, \mathbf{x}_{r-2}, \mathbf{y}_r \oplus jA_r)$ to be the original color of $(\mathbf{x}_1, \ldots, \mathbf{x}_{r-2}, \mathbf{x}_{r-1}, \mathbf{y}_r \oplus jA_r)$. There are $r^{2k^{N_1+\ldots+N_{r-2}}}$ such possible colorings — since the colorings for $\mathbf{y}_r \oplus jA_r$ are the same for $1 \leq j \leq k-1$ we have a 2 in place of the more obvious k. Again we view each of these coloring possibilities as a palette of colors for $\mathbf{x}_{r-1}$. Thus if N_{r-1} is sufficiently large — precisely, $N_{r-1} \geq HJ(k-1, r^{2k^{N_1+\ldots+N_{r-2}}})$ — then we may find $\mathbf{y}_{r-1}$ and a non-empty subset $A_{r-1} \subset [1, N_{r-1}]$ such that, given $\mathbf{x}_1$, ..., $\mathbf{x}_{r-2}$, and $1 \leq j_r \leq k$, the color of $(\mathbf{x}_1, \ldots, \mathbf{x}_{r-2}, \mathbf{y}_{r-1} \oplus j_{r-1}A_{r-1}, \mathbf{y}_r \oplus j_r A_r)$ does not change as j_{r-1} varies from 1 to $k-1$.

We continue in this manner. In stage ℓ we require that $N_{r-\ell+1} \geq HJ(k-1, r^{2^{\ell-1}k^{N_1+\ldots+N_{r-\ell}}})$, and produce $\mathbf{y}_{r-\ell+1}$ and a non-empty subset $A_{r-\ell+1} \subset [1, N_{r-\ell+1}]$. After r stages we will have produced points $\mathbf{y}_1$, ..., $\mathbf{y}_r$ and non-empty sets A_1, ..., A_r such that the points $(\mathbf{y}_1 \oplus j_1 A_1, \mathbf{y}_2 + j_2 A_2, \ldots, \mathbf{y}_r \oplus j_r A_r)$ and $(\mathbf{y}_1 \oplus j_1' A_1, \mathbf{y}_2 + j_2' A_2, \ldots, \mathbf{y}_r \oplus j_r' A_r)$ have the same color if for each $1 \leq i \leq r$ either $j_i = j_i' = k$ or $j_i, j_i' < k$. But this is tantamount to having an alphabet with just two elements: either $j_i < k$ or $j_i = k$, and having an r coloring on an r-dimensional cube on this 2 element alphabet. Thus we have reduced to $HJ(2, r)$! The proof follows.

The bounds produced by this proof are obviously not even astronomical.

Shelah's proof. The proof given above reduces $HJ(k, r)$ to $HJ(2, \ )$ using $HJ(k-1, \ )$. Shelah instead uses $HJ(2, \ )$ to reduce to $HJ(k-1, \ )$ and the effect of this simple change on bounds is dramatic!

Let R be a large parameter to be chosen, and let N_1, ..., N_R be a rapidly growing sequence, and set $N = N_1 + \ldots + N_R$. Consider $[1,k]^N$ as $[1,k]^{N_1} \times \ldots \times [1,k]^{N_R}$. Consider $\mathbf{x}_R$ in $[1,k]^{N_R}$ and associate to it an r-coloring $\phi_{\mathbf{x}_R}$ of $[1,k]^{N_1} \times \ldots \times [1,k]^{N_{R-1}}$ by giving $(\mathbf{x}_1, \ldots, \mathbf{x}_{R-1})$ the original color of $(\mathbf{x}_1, \ldots, \mathbf{x}_R)$. Note that there are $r^{k^{N_1+\ldots+N_{R-1}}}$ such colorings. As before, we think of these as a palette of colors for the $\mathbf{x}_R$. Thus if N_R is sufficiently large — precisely $N_R \geq HJ(2, r^{k^{N_1+\ldots+N_{R-1}}}) = r^{k^{N_1+\ldots+N_{R-1}}}$ — then we may find $\mathbf{y}_R$ and a non-empty set $A_R \subset [1, N_R]$ such that $\phi_{\mathbf{y}_R \oplus jA_R}$ determine the same coloring for $j = 1$ and $j = 2$. In other words, given any $\mathbf{x}_1$, ..., $\mathbf{x}_{R-1}$, the colors of $(\mathbf{x}_1, \ldots, x_{R-1}, \mathbf{y}_R \oplus jA_R)$ are the same for $j = 1$ and $j = 2$.

Now we repeat the same for $\mathbf{x}_{R-1}$. Consider the coloring on $[1,k]^{N_1} \times \cdots [1,k]^{N_{R-2}} \times [1,k]$ by setting the color of $(\mathbf{x}_1, \ldots, \mathbf{x}_{R-2}, \mathbf{y}_r \oplus jA_R)$ to be the original color of $(\mathbf{x}_1, \ldots, \mathbf{x}_{R-2}, \mathbf{x}_{R-1}, \mathbf{y}_r \oplus jA_R)$. There are $r^{(k-1)k^{N_1+\ldots+N_{R-2}}}$ such colorings — the $(k-1)$ arises because the colors for $j = 1$ and $j = 2$ are the same. Therefore, viewing these as a palette of colors for $\mathbf{x}_{R-1}$ we see

that if N_{R-1} is large — precisely, $N_{R-1} \geq HJ(2, r^{(k-1)k^{N_1+\ldots+N_{R-2}}}) = r^{(k-1)k^{N_1+\ldots+N_{R-2}}}$ — then there exists $\mathbf{y}_{R-1}$ and a non-empty $A_{R-1} \subset [1, N_{R-1}]$ such that given $\mathbf{x}_1, \ldots, \mathbf{x}_{R-2}$, and $1 \leq j_R \leq k$, the color of $(\mathbf{x}_1, \ldots, \mathbf{x}_{R-2}, \mathbf{y}_{R-1} \oplus j_{R-1}A_{R-1}, \mathbf{y}_R \oplus j_R A_R)$ is the same for $j_{R-1} = 1$ or 2.

We continue in this manner. In stage ℓ we require that $N_{R-\ell+1} \geq HJ(2, r^{(k-1)^{\ell-1}k^{N_1+\ldots+N_{R-\ell}}}) = r^{(k-1)^{\ell-1}k^{N_1+\ldots+N_{R-\ell}}}$ and produce $\mathbf{y}_{R-\ell+1}$ and a non-empty $A_{R-\ell+1} \subset [1, N_{R-\ell+1}]$. After R stages we find that $(\mathbf{y}_1 \oplus j_1 A_1, \ldots, \mathbf{y}_R \oplus j_R A_R)$ and $(\mathbf{y}_1 \oplus j_1' A_1, \ldots, \mathbf{y}_R \oplus j_R' A_R)$ have the same color if for each i either $1 \leq j_i, j_i' \leq 2$ or $j_i = j_i'$. This is tantamount to having a $k-1$ alphabet (1 and 2 are identified) and r-coloring an R dimensional cube on this alphabet. So if $R \geq HJ(k-1, r)$ then we are done.

Deduction of Van der Waerden's theorem. Van der Waerden's theorem states that if the numbers $[1, N]$ are r-colored then there is a mono-chromatic k-AP provided N is large in terms of k and r. One way to deduce this from Hales-Jewett is to identify $[1, k]^N$ with the numbers from $[1, k^N]$ by means of the base k expansion: thus $W(k, r) \leq k^{HJ(k,r)}$. Alternatively, consider an r-coloring of $[0, (k-1)N]$ and use this to get a coloring of $[1, k]^N$ by setting the color of $(x_1, \ldots, x_N)$ to be the color of $\sum(x_i - 1)$. A Hales-Jewett line on $[1, k]^N$ then gives a k-AP. Thus $W(k, r) \leq (k-1)HJ(k, r) + 1$.

Roth's Theorem

1. Basic Fourier Analysis.

Throughout we will write $\mathbb{Z}_N$ to denote $\mathbb{Z}/N\mathbb{Z}$. Often it may be convenient to assume that N is prime; since most of our results are insensitive to the exact value of N, this assumption is largely harmless. Given $f : \mathbb{Z}_N \to \mathbb{C}$, we define the Fourier transform $\hat{f} : \mathbb{Z}_N \to \mathbb{C}$ by setting

$$\hat{f}(r) = \sum_n f(n) e\Big(-\frac{rn}{N}\Big).$$

Here and throughout, $e(\theta)$ denotes $e^{2\pi i\theta}$. We may easily check the Fourier inversion formula,

$$f(n) = \frac{1}{N}\sum_r \hat{f}(r) e\Big(\frac{rn}{N}\Big).$$

Similarly we may easily check Parseval's formula

$$N\sum_n f(n)\overline{g(n)} = \sum_k \hat{f}(k)\overline{\hat{g}(k)},$$

and in the special case $f = g$

$$N\sum_n |f(n)|^2 = \sum_k |\hat{f}(k)|^2.$$

Notation. Given a set $A \in \mathbb{Z}_N$ we will use $A(n)$ to denote the characteristic function of A; that is, $A(n) = 1$ if $n \in A$ and 0 otherwise. Also, we will let f_A denote the *balanced function* $f_A(n) = A(n) - |A|/N$. Note that $\hat{A}(0) = |A|$ and that $\hat{f}_A(0) = 0$.

2. The Proof of Roth's Theorem.

Roth's Theorem says that if a positive density subset of the integers contains a 3-AP. We now give a streamlined proof of this due to Gowers.

Let $A \subset [1, N]$ with $|A| = \delta N$. We assume that N is odd, and let B be the set of even or odd numbers in A whichever is larger. Consider

$$\frac{1}{N}\sum_{r \pmod N} \hat{B}(r)^2\hat{A}(-2r) = \#\{x + y \equiv 2z \pmod N : \quad x, y \in B, z \in A\}.$$

Here $\hat{B}(r) = \sum_{b\in B} e(br/N)$, and similarly for $\hat{A}(r)$. By size and parity considerations, we find that $x + y \equiv 2z \pmod N$ in fact implies that $x + y = 2z$ so that we have a three term AP. There are $|B|$ such trivial 3-APs. So the number of non-trivial 3-APs is

$$\frac{1}{N}\sum_{r \pmod N} \hat{B}(r)^2\hat{A}(-2r) - |B| = \frac{|B|^2|A|}{N} - |B| + \frac{1}{N}\sum_{r\neq 0} \hat{B}(r)^2\hat{A}(-2r). \tag{1}$$

The proof now splits into two parts: when A has no large Fourier coefficients (A is random), and when A has a large Fourier coefficient (A has a structure – linear bias). The first case is easy, while in the second case we have to work a little to uncover the structure.

First suppose that for all $r \neq 0$ we have $|\hat{A}(r)| \leq \delta^2 N/4$. In this case

$$\frac{1}{N}\Big|\sum_{r\neq 0}\hat{B}(r)^2\hat{A}(-2r)\Big| \leq \frac{\delta^2}{4}\sum_r |\hat{B}(r)|^2 = \frac{\delta^2}{4}N|B| \leq \frac{|A||B|^2}{2N}.$$

From (1) we deduce that there are many 3-APs in this case.

So we may suppose that there exists $r \neq 0$ such that $|\hat{A}(r)| \geq \delta^2 N/4$. Equivalently we have

$$\Big|\sum_{a=1}^{N}(A(a)-\delta)e(ar/N)\Big| \geq \frac{\delta^2}{4}N, \tag{2}$$

where $A(a)$ denotes the characteristic function of A.

Let $1 \leq Q \leq N$ be a parameter to be chosen shortly. We use Dirichlet's theorem to find b/q where $q \leq Q$ and $(b,q) = 1$ such that $|r/N - b/q| \leq \frac{1}{qQ}$. We then divide $[1,N]$ into progressions $\pmod q$. There are q such progressions each with $N/q + O(1)$ elements. We now subdivide these progressions into M intervals each. Thus there are qM such intervals in all, and each interval contains about $N/(qM) + O(1)$ elements. Let I denote a typical such interval. We claim that on I, $e(ar/N)$ is more or less constant. Indeed it is $e(ab/q + a\theta)$ for some $|\theta| \leq 1/qQ$. Now $e(ab/q)$ is constant as all elements of I are in the same progression $\pmod q$. The variation in $e(a\theta)$ is at most $O(N|\theta|/M) = O(N/(qQM))$. Thus from (2) we find that

$$\begin{aligned}\frac{\delta^2 N}{4} &\leq \sum_I \Big|\sum_{a\in I}(A(a)-\delta)e(ar/N)\Big| = \sum_I \Big(\Big|\sum_{a\in I}(A(a)-\delta)\Big| + O\Big(\frac{N|I|}{qQM}\Big)\Big|\\ &= \sum_I \Big|\sum_{a\in I}(A(a)-\delta)\Big| + O\Big(\frac{N^2}{qQM}\Big).\end{aligned}$$

We will choose $Q = \sqrt{N}$ and $M = C\sqrt{N}/(q\delta^2)$ for a suitably large constant C. Then we obtain

$$\frac{\delta^2 N}{8} \leq \sum_I \Big|\sum_{a\in I}(A(a)-\delta)\Big|. \tag{3}$$

Plainly

$$0 = \sum_I \sum_{a\in I}(A(a)-\delta),$$

so that from (3) we may deduce the existence of an I with

$$\sum_I (A(a)-\delta) \geq \frac{\delta^2 N}{16qM}.$$

Recall that I contains about $N/(qM)$ elements, and so the relative density of A within I is at least $\delta + \delta^2/16$ — this is referred to as a density increment argument. Now we

take I and translate and dilate it so that it corresponds to the set $[1, N/qM]$. Note that APs are preserved under translation and dilation. We have thus extracted a set of density $\delta + \delta^2/16$ lying in $[1, N/(qM)] = [1, \delta^2\sqrt{N}/C]$, and it suffices to exhibit 3-APs in this set. Now we may repeat the entire argument for this set.

If we iterate the argument D/δ times for an appropriate constant D then we can get a density > 0.9 when it is easy to exhibit 3APs. After these iterations, the initial value of N would have been reduced to $\delta^4 N^{1/2^L}/C^2$. We wish this last quantity to be relatively large; say $\geq 10^3$. Thus we would like $N \geq (10^3 C^2/\delta^4)^{2^{D/\delta}}$. Equivalently if $\delta > c/\log\log N$, the argument works.

3. Behrend's Example.

Behrend constructed a surprisingly large set in $[1, N]$ with no 3-APs.

Behrend's Theorem A. *There is a set A in $[1, N]$ which is free of 3 APs and satisfies $|A| \gg N\exp(-c\sqrt{\log N})$. Here c is an absolute positive constant.*

Behrend's Theorem B. *There exists a set A in $[1, N]$ with $|A| \geq \delta N$ which has $\ll \delta^{c\log(1/\delta)} N^2$ three term progressions. Here c is an absolute positive constant, and $\delta > 0$.*

Proof of Theorem A. Consider points $(x_1, \dots, x_K)$ with $x_i \in [0, d]$. Thus there are $(d+1)^K$ such points. Consider $\sum_{i=1}^K x_i^2$. This is an integer in $[0, Kd^2]$ and so there exists $n \leq Kd^2$ such that $n = \sum x_i^2$ for more than $(d+1)^K/(Kd^2)$ tuples. That is, there is a sphere with many points. The argument rests on the fact that any line can intersect the sphere in at most two points.

Consider the set $A = \{\sum_{i=1}^K x_i(2d+1)^{i-1}\}$ where $(x_1, \dots, x_K)$ is a point on our sphere. Note that all elements of A are below $(2d+1)^K$ and that $|A| \geq (d+1)^K/(Kd^2)$. We claim that A has no 3-APs. For, if

$$\sum x_i(2d+1)^{i-1} + \sum z_i(2d+1)^{i-1} = \sum 2y_i(2d+1)^{i-1},$$

then $x_i + z_i = 2y_i$, as x_i, y_i, and z_i are all $\leq d$. In other words the points $(x_1, \dots, x_K)$, $(y_1, \dots, y_K)$, and $(z_1, \dots, z_K)$ all lie on a line, which is impossible.

Now take K about size $\sqrt{\log N}$, and d about size $e^{\sqrt{\log N}}$. Then A is a set in $[1, N]$ with $|A| \geq N\exp(-c\sqrt{\log N})$ with no 3-APs.

Proof of Theorem B. Let A_0 be a set of $2M\delta$ elements in $[1, M]$ such that A_0 is free of three term progressions. Consider $A \in [1, 2MK]$ which consists of all elements $a \equiv a_0 \pmod{2M}$ for some $a_0 \in A_0$. Thus $|A| = 2KM\delta$. If three elements in A lie in an AP, then, since A_0 has no 3-APs, they must all be congruent $\pmod{2M}$. The number of three term progressions in A is therefore $\ll (2M\delta)K^2 \ll (\delta/M)(2MK)^2$. Theorem A furnishes an example where $M = \exp(c(\log 1/\delta)^2)$, and so the result follows.

4. Fourier coefficients alone do not control four APs: A quadratic example.

Let $A \subset \mathbb{Z}_N$ be defined as follows: $a \in A$ if and only if there exists $|b| \leq N\delta/2$ with $a^2 \equiv b \pmod{N}$. We claim that $\hat{A}(k)$ is small for every $k \neq 0$, but there are substantially more four term arithmetic progressions in A than for a random set. It's clear from the definition that $|A| = \delta N + O(1)$.

For $k \neq 0$ we have that

$$\hat{A}(k) = \sum_{n \in A} e(-nk/N) = \sum_{n \pmod N} \Big(\frac{1}{N} \sum_{|b| \leq \delta N/2} \sum_{r \pmod N} e\Big(\frac{r(n^2-b)}{N}\Big)\Big) e\Big(-\frac{kn}{N}\Big).$$

Regrouping, this is

$$\frac{1}{N} \sum_{r \pmod N} \Big(\sum_{n \pmod N} e\Big(\frac{rn^2 - kn}{N}\Big)\Big)\Big(\sum_{|b| \leq \delta N/2} e\Big(-\frac{br}{N}\Big)\Big).$$

When $k \neq 0$ (and suppose N is odd) the inner sum over n above is a Gauss sum which is $\leq \sqrt{N}$ in magnitude. The sum over b is bounded by $\ll \min(\delta N, 1/\|r/N\|)$. It follows that

$$\hat{A}(k) \ll \frac{1}{N}\sqrt{N} \sum_{r \pmod N} \min\Big(\delta N, \frac{1}{\|r/N\|}\Big) \ll \sqrt{N} \log N.$$

This proves our first claim that $\hat{A}(k)$ is small for non-zero k.

To prove our second claim consider $B \subset A$ defined by $b \in B$ if and only if $b^2 \equiv c \pmod N$ for some $|c| \leq \delta N/14$. Plainly $|B| = \delta N/7 + O(1)$. The above argument shows readily that B has no non-trivial large Fourier coefficients, and therefore we may find $\gg \delta^3 N^2$ three term progressions in B. Consider one of those progressions a, $a+d$, $a+2d$ all in B (and hence in A). We claim that $a+3d$ belongs to A. To see this, note the identity

$$a^2 - 3(a+d)^2 + 3(a+2d)^2 - (a+3d)^2 = 0.$$

It follows that $(a+3d)^2$ has a representative $\pmod N$ of size less than $7(\delta N/14) = \delta N/2$, proving our claim. Thus A has many more four APs than we would expect for a random set (viz. $\delta^4 N^2$).

5. Varnavides's Theorem.

Varnavides's Theorem is a stronger form of Roth's theorem which counts the number of three term progressions.

Varnavides's Theorem. *For every $\epsilon > 0$ there exists $C(\delta) > 0$ such that if $A \subset [1, N]$ with $|A| \geq \delta N$ then A contains at least $C(\delta)N^2$ three term progressions.*

Proof. By Roth's theorem we know that there exists $M = M(\delta)$ such that any set of $\delta M/2$ elements in $[1, M]$ has a non-trivial three term progression. Now consider progressions $P(a, d) = a + [1, M]d$ in $[1, N]$ where we allow $d \leq \delta N/M^2$ and $a \leq N(1 - \delta/M)$.

We claim that for many choices of a and d one has $|A \cap P(a, d)| \geq \delta M/2$. Indeed we have that for any given d

$$\sum_{a \leq N(1-\delta/M)} |A \cap P(a,d)| \geq M \sum_{\substack{a \in A \\ Md \leq a \leq N(1-\delta/M)}} 1 \geq M\Big(\delta N - 2N\delta/M\Big).$$

It follows that for each d there are $\gg \delta N$ values of a with $|A \cap P(a,d)| \geq \delta M/2$, so that in total there are $\gg \delta^2 N^2/M^2$ good progressions $P(a,d)$.

By Roth's theorem each good progression contributes at least one three term progression in A. But of course some of these progressions could get over counted. Suppose we are given a progression x, $x+y$, $x+2y$ in A. To how many $P(a,d)$'s could this belong? Clearly d must be a divisor of y, and moreover $y/d \leq M$. Therefore there are at most M choices for d. Each choice of d fixes a in at most M ways. Therefore each progression is over counted at most M^2 times.

Thus we have exhibited $\gg \delta^2 N^2/M^4$ distinct three term progressions, and this proves Varnavides's Theorem.

6. The large spectrum of a set.

Suppose $A \subset \mathbb{Z}_N$ with $|A| = \delta N$. In our proof of Roth's theorem a crucial role was played by the large Fourier coefficients of A. Namely the set $R = R(\rho)$ of values r with $|\hat{A}(r)| \geq \rho |A|$. From Parseval we see easily that $|R| \leq \rho^{-2}\delta^{-1}$. A result of M. C. Chang (which we may explore later) says that the set of large Fourier coefficients has a very rigid structure.

Chang's Theorem. *The set $R(\rho)$ is contained in a cube of dimension at most $2\rho^{-2}\log(1/\delta)$. That is, there exist numbers $r_1, \ldots, r_k$ with $k \leq 2\rho^{-2}\log(1/\delta)$ such that each $r \in R$ may be written as $\sum \epsilon_j r_j$ where the ϵ_j take values -1, 0, or 1.*

Weyl's equidistribution theorem and differencing method

The idea of using Fourier analysis to study equidistribution goes back to Weyl. Here is his celebrated equidistribution criterion:

Weyl's criterion. *Let $u_1, u_2, \ldots,$ be a sequence of real numbers. We say that this sequence is uniformly distributed* $\pmod 1$ *if either of the following three equivalent statements holds:*

(1) For every interval $(\alpha, \beta) \in \mathbb{T}$ we have

$$\lim_{N\to\infty} \frac{1}{N} \#\{n \le N : u_n \pmod 1 \in (\alpha, \beta)\} = \beta - \alpha.$$

(2) For every non zero integer k we have

$$\sum_{n\le N} e(ku_n) = o(N),$$

as $N \to \infty$.

(3) For every Riemann integrable function f on $\mathbb{T}$ we have

$$\frac{1}{N} \sum_{n\le N} f(u_n) \to \int_{\mathbb{T}} f(u)du$$

as $N \to \infty$.

Sketch of proof. We check easily that (3) $\implies$ (2) and that (1) $\implies$ (3). To complete the proof one shows (2) $\implies$ (1). This follows by using Weierstrass's approximation theorem to approximate the characteristic function of $[\alpha, \beta]$ by trigonometric polynomials.

As an immediate application of Weyl's criterion we obtain that $\{n\theta\}$ is uniformly distributed $\pmod 1$ when θ is irrational. Weyl generalized this substantially by showing that the fractional parts of a polynomial $p(n)$ which has an irrational coefficient (and which is not the constant term) become equidistributed $\pmod 1$. To achieve this he added a crucial new idea which has come to be known as Weyl differencing.

To illustrate this, we first show that for odd N the Gauss sum $\sum_n e(n^2/N)$ has size $\sqrt{N}$; we used this previously in our quadratic example on four term progressions. We square the Gauss sum

$$\Big|\sum_n e(n^2/N)\Big|^2 = \sum_{n_1,n_2 \pmod N} e\Big(\frac{n_1^2 - n_2^2}{N}\Big).$$

Now if we write $n_1 = n_2 + h$ then $n_1^2 - n_2^2 = 2hn_2 + h^2$ is a linear polynomial in n_2. Thus the above equals

$$\sum_{h \pmod N} \sum_{n_2 \pmod N} e\Big(\frac{2hn_2 + h^2}{N}\Big) = N,$$

since only the term $h = 0$ survives. This proves our claim. The heart of the proof is that a degree two polynomial is reduced by differencing to a degree one polynomial, and this last sum is easy to evaluate.

Lemma 1. *Suppose* $|\alpha - a/q| \le 1/q^2$ *where* $(a,q)=1$, *and* $q \ge 2$. *Then for* $N \ge 1$ *we have*

$$\sum_{n\le N} e(n^2\alpha) \ll \frac{N}{\sqrt{q}} + \sqrt{(q+N)\log q}.$$

Proof. Call the sum in question S. Then

$$|S|^2 = \sum_{n_1,n_2\le N} e(\alpha(n_1^2-n_2^2)) = \sum_{|h|\le N} \sum_{\substack{n_2\ge \max(-h,0)\\ n_2\le \min(N,N-h)}} e(\alpha(2hn_2+h^2)).$$

Recalling that $\sum_{a\le n\le b} e(n\theta) \ll \min(b-a, 1/\|\theta\|)$, the above is

$$\ll \sum_{|h|\le N} \min(N, 1/\|2h\alpha\|).$$

Divide the terms $|h|\le N$ into intervals of length $q/2$. We may check easily that for each interval, the sum over h is $\ll N + q\log q$. Since there are $\ll N/q+1$ such intervals, our claimed estimate follows.

From Lemma 1 and Weyl's criterion it follows that if α is irrational then the fractional parts of αn^2 are equidistributed $\pmod 1$. In particular it follows that $\|n^2\alpha\|$ can be made less than ϵ for any given ϵ; we observed this as a consequence of van der Waerden's theorem. The argument extends inductively to cover all polynomials having at least one irrational coefficient (which is not simply the constant term). As an exercise, the reader may try to bound $\sum_{n\le N} e(n^k\alpha)$ for $k\ge 3$.

For our later applications we will need a more quantitative version of finding small fractional parts of $n^2\alpha$ (just as we needed Dirichlet's theorem in the proof of Roth's theorem).

Lemma 2. *Given a rational number* a/q *with* $(a,q)=1$, *there exists* $m \le M$ *with* $\|am^2/q\| \ll \sqrt{q}(\log q)^{\frac{3}{2}}/M$.

Proof. We may assume that $M \le q$. We want to find solutions to $am^2 \equiv b \pmod q$ with $|b|$ being small, say $\le L$. That is we want to estimate

$$\frac{1}{q}\sum_{|b|\le L}\sum_{r \pmod q}\sum_{m\le M} e\Big(\frac{(am^2-b)r}{q}\Big).$$

The term $r=0$ gives a main term of $(2L+1)M/q$. The terms $r\ne 0$ give, using Lemma 1 to estimate the sum over m,

$$\ll \frac{1}{q}\sum_{r\ne 0}\sqrt{q\log q}\min\Big(L, \frac{1}{\|r/q\|}\Big) \ll \sqrt{q}(\log q)^{\frac{3}{2}}.$$

It follows that if $L \gg q\sqrt{q}(\log q)^{\frac{3}{2}}/M$ we will have such solutions, proving the Lemma.

Corollary 3. *Let α be any real number. For every $M \geq 1$ there exists a natural number $m \leq M$ with $\|m^2\alpha\| \ll (\log M)/M^{\frac{1}{3}}$.*

Proof. First we find a rational number a/q with $q \leq Q$, $(a,q) = 1$ and $|\alpha - a/q| \leq 1/(qQ)$. If $q \leq M$ then by choosing $m = q$ we find that $\|\alpha m^2\| \leq q/Q \leq M/Q$. Suppose now that $q > M$. Using Lemma 2 we may find $m \leq M$ with $\|am^2/q\| \leq \sqrt{q}(\log q)^{\frac{3}{2}}/M$, and therefore

$$\|\alpha m^2\| \leq \frac{\sqrt{q}(\log q)^{\frac{3}{2}}}{M} + \frac{M^2}{qQ} \leq \frac{\sqrt{Q}(\log Q)^{\frac{3}{2}}}{M} + \frac{M}{Q}.$$

Choosing now $Q = M^{\frac{4}{3}}/\log M$ we obtain the desired conclusion.

It is expected that one can find $m \leq M$ with $\|m^2\alpha\| \ll M^{-1+\epsilon}$ for any $\epsilon > 0$. Heilbronn showed that one can achieve $\ll M^{-\frac{1}{2}+\epsilon}$, and Zaharescu (about ten years ago) obtained $\ll M^{-\frac{4}{7}+\epsilon}$.

Vinogradov's Three Primes Theorem

The still unresolved Goldbach conjecture states that every even number is the sum of two primes. In the 1930s Vinogradov proved the striking result that every large odd number is the sum of three primes. His method applies also to 3-APs in the primes, and this was done by van der Corput. We will now discuss these results; to illustrate the ideas, we will assume the truth of the Generalized Riemann Hypothesis and argue in the spirit of Hardy and Littlewood, but the assumption of GRH can be removed with more effort. Below, $\Lambda(n)$ denotes the von Mangoldt function $\Lambda(n) = \log p$ if n is a power of the prime p, and $\Lambda(n) = 0$ otherwise.

Vinogradov's Theorem. *Let N be a large natural number. Then*

$$\sum_{n_1+n_2+n_3=N} \Lambda(n_1)\Lambda(n_2)\Lambda(n_3) \sim \frac{N^2}{2}\mathfrak{S}(N),$$

where

$$\mathfrak{S}(N) = \prod_{p|N}\Big(1 - \frac{1}{(p-1)^2}\Big)\prod_{p\nmid N}\Big(1 + \frac{1}{(p-1)^3}\Big).$$

When N is odd we have $\mathfrak{S}(N) \gg 1$, and since the number of prime squares, cubes, etc. is small, we conclude that there are many ways of writing a large odd number as a sum of three primes.

Let

$$f(\alpha) = \sum_{n\le N} \Lambda(n)e(n\alpha).$$

Note that

$$\int_0^1 f(\alpha)^3 e(-N\alpha)d\alpha = \sum_{n_1+n_2+n_3=N} \Lambda(n_1)\Lambda(n_2)\Lambda(n_3),$$

and our aim is to obtain an asymptotic formula for this quantity. The insight of Hardy and Littlewood was that the exponential sum $f(\alpha)$ is large only when α is close to a rational number with small denominator, and such values of α give the dominant contribution to our integral. We already saw this feature for $\sum_{n\le N} e(n^2\alpha)$ from Weyl's method.

To gain an understanding of $f(\alpha)$ we will invoke the Generalized Riemann Hypothesis. What we need is the consequence of GRH for the distribution of primes in progressions. Recall that there are $\phi(q)$ Dirichlet characters $\chi \pmod q$.[1] If $\chi \pmod q$ is a non-trivial Dirichlet character, then on GRH for $x \ge q$ we have

$$\psi(x,\chi) = \sum_{n\le x} \Lambda(n)\chi(n) \ll x^{\frac12}\log^2 x. \tag{1}$$

[1] These are homomorphisms from $(\mathbb{Z}/q\mathbb{Z})^* \to \mathbb{T}$. The functions are extended to all of $\mathbb{Z}$ by setting $\chi(n) = 0$ if $(n,q) > 1$.

If $\chi_0 \pmod q$ is the trivial character then we have (on RH)

$$\psi(x,\chi_0) = \sum_{n\le x}\Lambda(n) + O(\log^2 qx) = x + O(x^{\frac12}\log^2(qx)). \tag{2}$$

Given $(a,q)=1$, the orthogonality relation

$$\frac{1}{\phi(q)}\sum_{\chi \pmod q}\chi(n)\overline{\chi(a)} = \begin{cases} 1 & \text{if } n\equiv a \pmod q \\ 0 & \text{otherwise}\end{cases}$$

shows that, for $x\ge q$,

$$\psi(x;q,a) = \sum_{\substack{n\le x\\ n\equiv a \pmod q}}\Lambda(n) = \frac{x}{\phi(q)} + O(x^{\frac12}\log^2 x). \tag{3}$$

Lemma 1. *Let a/q be a rational number with $(a,q)=1$. Then we have, assuming GRH,*

$$\sum_{n\le x}\Lambda(n)e(na/q) = \frac{\mu(q)}{\phi(q)}x + O(\sqrt{qx}\log^2 x).$$

Proof. We have

$$\sum_{n\le x}\Lambda(n)e(an/q) = O(\log^2 x) + \sum_{\substack{n\le x\\(n,q)=1}}\Lambda(n)e(an/q). \tag{4}$$

We could now split n into progressions $\pmod q$, and invoke (3). But there is some loss in doing this, and we would obtain an error term of $\ll q\sqrt{x}\log^2 x$ which is not sufficient for our purposes. A better way to proceed is to express $e(an/q)$ (for $(an,q)=1$) in terms of the multiplicative Dirichlet characters $\pmod q$:

$$e(an/q) = \frac{1}{\phi(q)}\sum_{b \pmod q}\sum_{\chi \pmod q}\chi(b)\overline{\chi(an)}e(b/q) = \frac{1}{\phi(q)}\sum_{\chi \pmod q}\overline{\chi(an)}\tau(\chi).$$

Here $\tau(\chi)$ denotes the Gauss sum

$$\tau(\chi) = \sum_{a \pmod q}\chi(a)e(a/q).$$

It is well known that $|\tau(\chi)|\le\sqrt{q}$, and equality holds there when χ is a *primitive* character (see for example, Davenport's *Multiplicative number theory*).

Using this in (4) we obtain that

$$\sum_{n\le x}\Lambda(n)e(an/q) = \frac{1}{\phi(q)}\sum_{\chi \pmod q}\overline{\chi(a)}\tau(\chi)\psi(x,\overline{\chi}) + O(\log^2 x).$$

Using the GRH bound (1) for all non-trivial χ we see that such terms contribute $\ll \sqrt{qx}\log^2 x$. It remains to handle the principal character χ_0. By (2) this term contributes

$$\frac{1}{\phi(q)}\tau(\chi_0)(x+O(\sqrt{x}\log^2 x) = \frac{\mu(q)}{\phi(q)}(x+O(\sqrt{x}\log^2 x)),$$

since $\tau(\chi_0)$ is a *Ramanujan sum* which is easily evaluated as $\mu(q)$. This completes our proof.

Lemma 1 allows us to handle $f(a/q)$, and to pass to $f(\alpha)$ for nearby α we use a standard technique known as partial summation.

Lemma 2. *Assume GRH. Let $\alpha = a/q+\beta$ where $(a,q)=1$. Then*

$$f(\alpha) = \frac{\mu(q)}{\phi(q)}\int_0^N e(\beta x)dx + O\Big((1+|\beta|N)\sqrt{qN}\log^2 N\Big).$$

Proof. Note that,

$$f(\alpha) = \int_0^N e(x\beta)d\Big(\sum_{n\le x}\Lambda(n)e(an/q)\Big) = \int_0^N e(x\beta)d\Big(\frac{\mu(q)}{\phi(q)}x + E(x,a/q)\Big),$$

say, for some error term $E(x,a/q)$. The first term gives the main term of the Lemma. As for the second term, integration by parts gives that it is

$$E(N,a/q)e(N\beta) - \int_0^N 2\pi i\beta e(x\beta)E(x,a/q)dx.$$

Using the bound of Lemma 1, we conclude the desired estimate.

Corollary 3. *Select $Q = N^{\frac{2}{3}}$, and let $|\alpha - a/q| \le 1/(qQ)$ with $(a,q)=1$ and $q\le Q$. Then, assuming GRH,*

$$f(\alpha) \ll \frac{N}{\phi(q)} + N^{\frac{5}{6}+\epsilon}.$$

Proof. Lemma 2 reveals (with any Q and $|\alpha - a/q| \le 1/(qQ)$ and $q\le Q$) that

$$f(\alpha) \ll \frac{N}{\phi(q)} + \Big(1+\frac{N}{qQ}\Big)\sqrt{qN}\log^2 N \ll \frac{N}{\phi(q)} + \Big(\sqrt{QN} + \frac{N^{\frac{3}{2}}}{Q}\Big)\log^2 N.$$

The optimal choice is $Q = N^{\frac{2}{3}}$, as in the Corollary, and the result follows.

Following Hardy and Littlewood, we say that points close to rational numbers with small denominators lie on *major arcs* while the *minor arcs* form the complementary set. Concretely, let us say that α lies on a major arc if $|\alpha - a/q| \le 1/(qQ)$ with $Q = N^{\frac{2}{3}}$, $(a,q)=1$ and $q \le (\log N)^{10}$.

Corollary 4. *If $\mathfrak{m}$ denotes the set of minor arcs, we have (on GRH)*

$$\int_{\mathfrak{m}} |f(\alpha)|^3 d\alpha \ll \frac{N^2}{(\log N)^9}.$$

Proof. From Corollary 3, it follows that if α lies on a minor arc then

$$f(\alpha) \ll \frac{N}{(\log N)^8},$$

since $\phi(q) \gg q/\log\log q$ and so if $q > (\log N)^{10}$ we have $\phi(q) \geq (\log N)^9$. Therefore

$$\int_{\mathfrak{m}} |f(\alpha)|^3 d\alpha \ll \frac{N}{(\log N)^9} \int_0^1 |f(\alpha)|^2 d\alpha = \frac{N}{(\log N)^9} \sum_{n \leq N} \Lambda(n)^2 \ll \frac{N}{(\log N)^8}.$$

It remains now to evaluate the major arc contribution. Precisely, we are interested in α with $|\alpha - a/q| \leq 1/(qQ)$ and $q \leq (\log N)^{10}$. Notice that the intervals for different rational numbers are disjoint. This major arc contribution equals

$$\int_{\mathfrak{M}} f(\alpha)^3 e(-N\alpha) d\alpha = \sum_{q \leq (\log N)^{10}} \sum_{\substack{1 \leq a \leq q \\ (a,q)=1}} \int_{-1/(qQ)}^{1/(qQ)} f(a/q+\beta)^3 e(-N(a/q+\beta)) d\beta. \tag{5}$$

From Lemma 2 we see that

$$f(a/q+\beta)^3 = \frac{\mu(q)^3}{\phi(q)^3}\Big(\int_0^N e(\beta x)dx\Big)^3 + O\Big(\frac{1}{\phi(q)^2}\min\Big(N^2, \frac{1}{|\beta|^2}\Big)(1+|\beta|N)\sqrt{qN}\log^2(qN)\Big) + O\Big((1+|\beta|N)^3(qN)^{\frac{3}{2}}\log^2(qN)\Big).$$

With a little calculation we see that the contribution of the error terms above to (5) is $\ll N^{\frac{11}{6}+\epsilon}$. Therefore, the major arc contribution is

$$\sum_{q \leq (\log N)^{10}} \frac{\mu(q)^3}{\phi(q)^3}\Big(\sum_{\substack{1 \leq a \leq q \\ (a,q)=1}} e(-Na/q)\Big)\Big(\int_{-1/(qQ)}^{1/(qQ)} \Big(\int_0^N e(\beta x)dx\Big)^3 e(-N\beta)d\beta\Big) + O(N^{\frac{11}{6}+\epsilon}). \tag{6}$$

The above factorizes nicely as a series and an integral. Let us tackle the integral first. By making a substitution $x = Ny$ and $N\beta = \xi$ this is

$$N^2 \int_{-N/(qQ)}^{N/(qQ)} \Big(\int_0^1 e(y\xi)dy\Big)^3 e(-\xi)d\xi = N^2\Big(\int_{-\infty}^{\infty} \Big(\int_0^1 e(y\xi)dy\Big)^3 e(-\xi)d\xi + O\Big(\frac{(qQ)^2}{N^2}\Big)\Big) = \frac{N^2}{2} + O((qQ)^2).$$

Using this in (6), our major arc contribution becomes

$$\frac{N^2}{2}\sum_{q\le(\log N)^{10}}\frac{\mu(q)^3}{\phi(q)^3}\Big(\sum_{\substack{1\le a\le q\\(a,q)=1}} e(-Na/q)\Big)+O(N^{\frac{11}{6}+\epsilon}).$$

The sum over a is a Ramanujan sum. We are only interested in it for square-free q. It is multiplicative, and equals -1 on primes not dividing N, and $p-1$ on the primes dividing N. At any rate, for a given q the sum over a is no more than $\phi(q)$. Thus the sum over q may be extended to infinity, with an error at most $\sum_{q>(\log N)^{10}}\mu(q)^2/\phi(q)^2\ll(\log N)^{-10}$. Therefore, our major arc contribution is

$$\sim\frac{N^2}{2}\sum_{q=1}^{\infty}\frac{\mu(q)^3}{\phi(q)^3}\Big(\sum_{\substack{1\le a\le q\\(a,q)=1}} e(-aN/q)\Big),$$

which by multiplicativity equals

$$\sim\frac{N^2}{2}\prod_p\Big(1-\frac{1}{(p-1)^3}\sum_{1\le a\le p-1}e(-Na/p)\Big)=\frac{N^2}{2}\mathfrak{S}(N).$$

This completes our conditional proof of Vinogradov's Theorem.

Sum-product estimates: A proof due to Solymosi

Let A denote a set of N real numbers. Erdős and Szemerédi had the insight that if $A+B$ is small for some set B (say of cardinality N) then A must have some additive structure, while if $A\cdot C$ (the set of products ac) is small for some set C then A has some multiplicative structure. Since we expect that multiplicative structures and additive structures should be independent of each other, they conjectured that one of the two sets must always be large. This circle of ideas has proved very fruitful in recent years. We illustrate this sum-product theory by giving a beautiful proof of Solymosi.

Theorem. *Let A, B and C be finite sets of real numbers, each having at least two elements. Then*

$$|A+B|\times|A\cdot C|\gg(|A|^3|B||C|)^{\frac{1}{2}}.$$

In particular, if A, B and C all have cardinality N then either $A+B$ or $A\cdot C$ has cardinality $\gg N^{\frac{5}{4}}$.

Proof. We remove 0 (if present) from the sets A and C; since our sets have at least two elements, the modified sets are non-empty. Suppose now that 0 is not in A or C.

For each element $a\in A$ we let $a'\in A$ denote the nearest neighbor of a. If there are two choices for a' pick the right hand neighbor. We wish to consider quadruples (a,a',b,c) where $a\in A$, $b\in B$, and $c\in C$. We call such a quadruple *good* if the following two conditions hold:

(i). The number of $u\in A+B$ with $|a+b-u|\le|a-a'|$ is at most $10|A+B|/|A|$.

(ii). The number of $v\in A\cdot C$ with $|ac-v|\le|ac-a'c|$ is at most $10|A\cdot C|/|A|$.

We will obtain upper and lower bounds for the number of good quadruples, and this will yield the Theorem.

First we consider the upper bound. Each quadruple (a,a',b,c) may be recovered uniquely from knowing $s_1=a+b$, $s_2=a'+b$, $p_1=ac$ and $p_2=a'c$. How many choices for these four numbers can lead to good quadruples? The first number s_1 can be chosen in $|A+B|$ ways. But the second number s_2 must then be one of the $10|A+B|/|A|$ elements of $A+B$ nearest to s_1; else there would be $>10|A+B|/|A|$ elements u of $A+B$ with $|s_1-u|\le|s_1-s_2|=|a-a'|$ which would contradict (a,a',b,c) being good. Similarly the third number can be chosen in $|A\cdot C|$ ways, but the fourth is then fixed in at most $10|A\cdot C|/|A|$ ways. Therefore the number of good quadruples is

$$\le 100\frac{|A+B|^2|A\cdot C|^2}{|A|^2}.$$

Now we consider the lower bound. Let b and c be given. We claim that there are at least $|A|/2$ values of a leading to a good quadruple. The claim shows that the number of good quadruples is at least $|A||B||C|/2$, which when combined with our upper bound above establishes the Theorem. To prove the claim consider

$$\sum_{a\in A}\#\{u\in A+B:\ \ |a+b-u|\le|a-a'|\}=\sum_{u\in A+B}\#\{a\in A:\ \ |(u-b)-a|\le|a-a'|\}.$$

The RHS counts for a given $(b-u)$ the number of a such that $(u-b)$ is as close as the nearest neighbor of a; clearly there are at most two such a, the ones enclosing $b-u$. Therefore the above is $\leq 2|A+B|$. It follows that at most $|A|/5$ values of $a \in A$ can satisfy $\#\{u \in A+B : \ |a+b-u| \leq |a-a'|\} \geq 10|A+B|$; or equivalently 80% of the elements in A satisfy criterion (i). Similarly 80% satisfy criterion (ii). Thus at least 60% satisfy both criteria, proving our claim.

A small modification to this proof (exercise) shows that the Theorem holds for sets of complex numbers as well. Solymosi has also shown that

$$|A+A|^8 \times |A \cdot A|^3 \gg |A|^{14-\epsilon}.$$

This inequality establishes firstly that either $A+A$ or $A \cdot A$ has cardinality $|A|^{\frac{14}{11}-\epsilon}$ which is the best currently known bound of this type. More remarkably it shows that if $|A+A| \leq C|A|$ (so that A looks like a generalized arithmetic progression) then $|A \cdot A| \gg |A|^{2-\epsilon}$ which is essentially best possible.

The sum-product phenomenon is quite general. A particularly useful version is due to the work of Bourgain, Katz, Tao, and Konyagin. Let p be a prime and $A \subset \mathbb{F}_p^*$ be a set with $|A| \leq p^{1-\delta}$. Then there is a constant $c = c(\delta) > 0$ such that $|A+A|+|A \cdot A| \geq c|A|^{1+c}$. In other words, there are no approximate subrings of $\mathbb{F}_p$.

There are other interesting proofs of the sum-product theorem. An elegant proof of Elekes uses as its main ingredient a beautiful result of Szemerédi and Trotter in incidence geometry.

Szemerédi-Trotter Theorem. *Suppose we are given m distinct lines in the plane, and n distinct points. Then the number of pairs (P, ℓ) where P is a point lying on a line ℓ is $\ll m+n+(mn)^{\frac{2}{3}}$.*

Deduction of sum-product estimates. For points take the set $(A+B) \times (A \cdot C)$ and for lines take $y = c(x-b)$ for each $b \in B$ and $c \in C$. Thus $n = |A+B||A \cdot C|$ and $m = |B||C|$. Each line contains $|A|$ points and so we obtain that

$$|A||B||C| \ll |B||C| + |A+B||A \cdot C| + (|B||C||A+B||A \cdot C|)^{\frac{2}{3}},$$

and the result follows.

There are also interesting non-abelian analogs of these results. For example, Helfgott has shown that if A is not contained in any proper subgroup of $SL_2(\mathbb{Z}/p\mathbb{Z})$ and $|A| < p^{3-\delta}$ for some $\delta > 0$ then $|A \cdot A \cdot A| > c|A|^{1+c}$ for some $c > 0$ depending only on δ.

We end by highlighting two striking propositions on sumsets (due to Solymosi) which follow easily from his method of proof above.

Proposition 1. *Let A be a set of N real numbers $a_1 < a_2 < \ldots < a_N$ such that the successive differences $a_{i+1} - a_i$ are all distinct. If B is any non-empty set of real numbers then $|A+B| \gg |A||B|^{\frac{1}{2}}$.*

First Proof. Given $a \in A$ let a' denote its nearest neighbor as before. Consider triples (a, a', b), and call such a triple good if the number of $u \in A+B$ with $|a+b-u| \leq |a-a'|$

is at most $10|A+B|/|A|$. Arguing as in the Theorem we obtain $\gg |A||B|$ good triples. To obtain an upper bound note that knowing $s_1 = a+b$ and $s_2 = a'+b$ we may recover (a, a', b); here we use that the consecutive differences are distinct! The number of choices of s_1, s_2 that lead to good triples is $\ll |A+B|^2/|A|$, and our result follows.

Second Proof. Consider the numbers in $A+B$. Divide them into about $N/2$ disjoint intervals with each interval containing about $|A+B|/(N/2)$ elements. We now count the number of pairs (i,b) where $1 \le i \le N$ and $b \in B$ such that $a_i + b$ and $a_{i+1} + b$ both lie in the same interval.

First we get a lower bound on the number of such pairs. Consider b as fixed. If (i,b) does not form a pair then there must be an interval jump between $a_i + b$ and $a_{i+1} + b$. Since there are only $N/2$ intervals, the number of such jumps is $\le N/2$. Therefore for $\ge N/2$ elements i we will get a pair (i,b). Thus the total number of pairs is $\gg N|B|$.

Now let us find an upper bound for this number. Choose two elements lying in the same interval. We claim that these can arise as $a_i + b$ and $a_{i+1} + b$ for at most one pair (i,b). Indeed the difference of the two elements must be $a_{i+1} - a_i$. This fixes i because of the assumption that the successive differences in A are distinct. Once i is fixed, naturally there is at most one choice for b. Thus the number of pairs is $\ll N\binom{|A+B|/(N/2)}{2} \ll |A+B|^2/N$.

Combining the upper and lower bounds we see the Proposition.

There is a generalization of Proposition 1 to two dimensions, which the reader may like to think about.

Proposition 2. *Let $A = \{a_1 < \ldots < a_N\}$ and $C = \{c_1 < \ldots < c_N\}$ be two N-element sets of real numbers such that the pairs of successive differences $(a_{i+1} - a_i, c_{i+1} - c_i)$ are all distinct. Then for any non-empty sets B and D we have*

$$|A+B||C+D| \gg (N^3|B||D|)^{\frac{1}{2}}.$$

Deduction of the Sum-Product estimate. Without loss of generality suppose that $A = \{a_1 < \ldots < a_N\}$ is a set of positive real numbers. Set $C = \{\log a\}$. Note that A and C satisfy the hypothesis of Proposition 2. Apply that Proposition with $B = A$ and $D = C$.

Erdős and Szemerédi have conjectured that $|A+A| + |A \cdot A| \gg |A|^{2-\epsilon}$. Solymosi has conjectured even more: if $|A| = |B| = |C|$ then $|A+B| + |A \cdot C| \gg |A|^{2-\epsilon}$. These remain open. We record that the ϵ in these conjectures is necessary. For example, if we take A to be the numbers from 1 to N, then $|A \cdot A| = o(N^2)$: this is the delightful multiplication table problem of Erdős and the order of magnitude of $|A \cdot A|$ has been recently determined by Kevin Ford. One can do better by choosing A to be the set of integers below M having exactly K distinct prime factors, and then choosing M and K appropriately. In this manner one obtains that the ϵ in Solymosi's conjecture must be $\gg 1/\log\log|A|$. Perhaps this is the right answer?

VII. Freiman's Theorem

1. Statement of the result, and the ingredients of the proof.

Let A be a subset of the integers with $|A| = N$. There are at most $N(N+1)/2$ possible distinct sums $a+b$ with a and b in A. Thus $|A+A| \le N(N+1)/2$. This upper bound can be attained when the numbers in A are spread apart: e.g. when they are powers of 2. At the other extreme we must always have $2N-1$ elements in the sumset $A+A$. For, if we arrange the elements in A as $a_1 < a_2 < \ldots < a_N$ then we have $a_1 + a_1 < a_1 + a_2 < \ldots < a_1 + a_N < a_2 + a_N < a_3 + a_N < \ldots < a_N + a_N$. It is easily checked that $|A+A| = 2N-1$ if and only if A consists of an arithmetic progression of length N.

Recall that a generalized arithmetic progression of dimension d is a set of numbers of the form $x_0 + \sum_{j=1}^{d} a_j x_j$ where the x_j are given integers, and $0 \le a_j < r_j$ for some numbers r_j. The size of a GAP is defined to be the number of (distinct) elements in this set. We say that a GAP is *proper* if all the sums above are distinct; thus its size is $\prod_j r_j$. If Q is a proper GAP of dimension d, then clearly $|2Q| \le 2^d |Q|$. Therefore a big subset of a proper GAP has small doubling. Freiman's remarkable theorem offers a converse to this statement.

Freiman's Theorem. *Let A be a set of integers for which there exists a set B with $|B| = |A|$ and $|A+B| \le C|A|$; in particular, A could be a set with $|A+A| \le C|A|$, or $|A-A| \le C|A|$. Then there exist constants d and S depending only on C such that A is contained in a GAP of dimension at most d, and size at most $S|A|$. In fact, we may choose $d \ll C^{34}$, and $S \ll e^{C^{34}}$.*

Above we have given A as a subset of a GAP. With more effort, one can ensure that this GAP is *proper*. We will describe below a beautiful proof of Freiman's theorem due to Ruzsa. There are three main ingredients in this proof:

• Plünnecke's inequalities. If A is a proper GAP then we see that not only does A have small doubling, but also the sets $kA - \ell A$ are small. Plünnecke showed that if A is a set with small doubling then $kA - \ell A$ is automatically also small. In fact, as in our statement of Freiman's theorem it is enough for A to have a set B with $|A| = |B|$ and $|A+B| \le C|A|$, and then it follows that $kA - \ell A$ is small. This is quite powerful, and by itself yields an elegant analog of Freiman's theorem for groups with bounded torsion (see §3 below).

• Freiman homomorhpisms and Ruzsa's embedding lemma. A Freiman k-homomorphism between sets is a map that preserves relations of the form $x_1 + \ldots + x_k = y_1 + \ldots + y_k$. A 2-homomorphism preserves arithmetic progressions. Ruzsa's embedding lemma allows one to pass from a set A (with small sumset $2A$) of integers to a large subset A' which is k-isomorphic to a subset of $\mathbb{Z}/N$. Here N is of size comparable to $|A|$. In other words, the problem for $\mathbb{Z}$ may be replaced with a problem for $\mathbb{Z}/N$ with N not too large.

• Bogolyubov's lemma and the structure of $2A - 2A$. Although a set A might lack additive structure and appear rough, the sumsets kA for large k start looking smoother. Bogolyubov found a simple Fourier argument that takes a large set A in $\mathbb{Z}/N$ and finds a highly structured set in $2A - 2A$. This highly structured set is a Bohr set, described by means of certain Diophantine inequalities. The second part of this argument uses the

geometry of numbers (in particular, Minkowski's second theorem) to extract from the Bohr set a large GAP.

Putting all these ingredients together one obtains Freiman's theorem.

2. Plünnecke's inequalities.

One of the key steps in the proof of Freiman's theorem is the following:

Plünnecke's Theorem A. *If A is a set of integers for which there exists a set B with $|A+B| \leq C|B|$, then*

$$|kA - \ell A| \leq C^{k+\ell}|B|,$$

for all integers k and ℓ. In particular, if $|B| = |A|$ then $|kA - \ell A| \leq C^{k+\ell}|A|$, and this conclusion holds for sets A satisfying $|A+A| \leq C|A|$ or $|A-A| \leq C|A|$.

The proof of this Theorem relies on some results of Plünnecke in graph theory. A *layered graph of level n* is a directed graph whose vertex set may be written as the disjoint union of $n+1$ sets V_0, V_1, ..., V_n, and whose edge set consists solely of edges from V_i to V_{i+1} for some $0 \leq i \leq n-1$. Given a layered graph G we define the i-th *magnification ratio* by

$$D_i(G) = \inf_{\substack{X \subset V_0 \\ X \neq \emptyset}} \frac{|\mathrm{Im}_i(X)|}{|X|},$$

where $\mathrm{Im}_i(X)$ denotes the set of vertices in V_i that may be reached by a path starting from some vertex in X.

Example 1. Consider a layered graph of level 3 with V_0 having three vertices, V_1 having one vertex, and V_2 having five vertices. Suppose all three points in V_0 are connected to the point in V_1 and that the point in V_1 is connected to four of the vertices in V_2. Then D_1 equals $1/3$ and D_2 equals $4/3$.

A *Plünnecke graph* is a special type of layered graph satisfying the following two properties:

Forward Splitting: Suppose that $0 \leq i \leq n-2$ and that u, v, w_1, ..., w_k are such that $u \in V_i$, $v \in V_{i+1}$, $w_j \in V_{i+2}$ and (u,v), (v,w_j) are all edges in our graph. Then there exist distinct v_1, ..., v_k in V_{i+1} such that (u,v_j), (v_j,w_j) are all edges in the graph.

Backward Splitting: Suppose that $0 \leq i \leq n-2$ and that there are distinct u_1, ..., u_k (in V_i), v in V_{i+1} and w in V_{i+2} so that (u_j,v), (v,w) are all edges in our graph. Then there exist distinct v_1, ..., v_k in V_{i+1} such that (u_j,v_j) and (v_j,w) are all edges in our graph.

It is clear that Example 1 is not a Plünnecke graph.

Example 2. This is the key example. Suppose A and B are any two sets of integers and take $V_i = A + iB$. Connect a vertex v in V_i to a vertex w in V_{i+1} exactly when $w - v \in B$. We can check that this graph is Plünnecke thanks to the commutativity of addition.

Example 3. This is a special case of Example 2 called the *independence graph*. We take A to be $\{0\}$ and B to be a set of h numbers such that all the possible n-fold sums of elements in B are distinct (for example, take B to be $\{1, (2n), (2n)^2, \ldots, (2n)^{h-1}\}$. Note that the i-th vertex set contains $\binom{h+i-1}{i}$ elements (exercise). Therefore $D_i(G)$ equals $\binom{h+i-1}{i}$ in this example. We will denote this graph by $I_n(h)$.

Example 4. Suppose that G and H are two layered graphs of level n with vertex sets $V_0, \dots V_n$ and $W_0, \dots, W_n$. The product graph $G \times H$ has a vertex set given by $V_j \times W_j$ for $j = 0, \dots, n$. There is an edge from (v_j, w_j) to (v_{j+1}, w_{j+1}) precisely if there (v_j, v_{j+1}) is an edge in G and (w_j, w_{j+1}) is an edge in H. With this definition the product of two layered graphs is a layered graph. Moreover, check that the product of two Plünnecke graphs is a Plünnecke graph.

Example 5. This is more of an observation on Plünnecke graphs which will be useful later on. Given a vertex $v_i \in V_i$ let $d_+(v_i)$ denote the number of edges coming out of v_i, and let $d_-(v_i)$ denote the number of edges coming into v_i. Suppose $v_{i+1} \in V_{i+1}$ is joined to v_i by an edge. Then the first Plünnecke condition shows that $d_+(v_i) \geq d_+(v_{i+1})$. The second Plünnecke condition shows that $d_-(v_i) \leq d_-(v_{i+1})$.

With these preliminaries in place, we can now state Plünnecke's result in graph theory.

Plünnecke's Theorem B. *Let G be a Plünnecke graph of level n. Then the magnification ratios satisfy the inequalities*

$$D_1 \geq D_2^{\frac{1}{2}} \geq D_3^{\frac{1}{3}} \geq \dots \geq D_n^{\frac{1}{n}}.$$

Before discussing the proof of Theorem B let us first deduce Theorem A.

Deduction of Theorem A. Consider the natural additive Plünnecke graph with $V_i = B + iA$. By assumption we know that $D_1 \leq |A + B|/|B| \leq C$. By Theorem B, it follows that for each $k \geq 1$ we have $D_k \leq C^k$. In other words, there exists a non-empty subset B' of B with $|B' + kA| \leq C^k |B'|$. It follows immediately that $|kA| \leq C^k |B|$ proving Theorem A in the case $\ell = 0$.

To deduce the full strength of the Theorem we need the following simple, but very useful Lemma of Ruzsa.

Ruzsa's Lemma. *For any three sets U, V, W we have*

$$|U||V - W| \leq |U + V||U + W|.$$

Proof. To each difference $d \in V - W$ associate a pair $(v(d), w(d)) \in V \times W$ with $v(d) - w(d) = d$. (Of course there may be many solutions to $v - w = d$; we just pick one.) Then given $(u, d) \in U \times (V - W)$ we may map it to $(u + v(d), u + w(d)) \in (U + V) \times (u + W)$. This map is injective, and therefore the inequality follows.

Returning now to the deduction of Theorem A, we see easily from Rusza's lemma that

$$|A||kA - \ell A| \leq |(k + 1)A||(\ell + 1)A| \leq C^{k+\ell+2}|B|^2,$$

which is certainly good enough for any applications, but with a little more effort we can recover the stated version. If $\ell \geq k$ then applying Theorem B twice we may find $B'' \subset B' \subset B$ with $|B' + kA| \leq C^k |B'|$ and $|B'' + \ell A| \leq C^\ell |B''|$ (why?). Now use Ruzsa's Lemma with $U = B''$, $V = kA$, and $W = \ell A$.

We now turn to the proof of Theorem B. There are two main steps in the argument.

Lemma 1. *Let G and H be layered graphs of level n. Then for each $i = 1, \ldots, n$ we have $D_i(G \times H) = D_i(G)D_i(H)$.*

Lemma 2. *Let G be a Plünnecke graph of level n with $D_n \geq 1$. Then there exist at least $|V_0|$ paths from V_0 to V_n such that the vertices of these paths are all disjoint. Consequently, $D_i \geq 1$ for all $i = 1, \ldots, n$.*

Assuming these lemmas let us now deduce Theorem B.

Deduction of Theorem B. When $D_n = 1$ the result follows by Lemma 2. Therefore we may assume that $0 < D_n < 1$ or that $D_n > 1$. Let's start with the former case. Let r and h be some natural numbers to be chosen, and consider the graph $G^r \times I_n(h)$. By Lemma 1 we know that

$$D_n(G^r \times I_n(n)) = D_n(G)^r D_n(I_n(h)) = D_n(G)^r \binom{h+n-1}{n} \geq D_n(G)^r \frac{h^n}{n!}.$$

Given r we will choose h to be the least number so that the above is at least 1; therefore $h \leq n!^{\frac{1}{n}} D_n(G)^{-\frac{r}{n}} + 1$. Then it follows from Lemma 2 that $D_i(G^r I_n(h)) \geq 1$ so that

$$D_i(G)^r \geq \frac{1}{D_i(I_n(h))} \geq h^{-i} \geq (n!^{\frac{1}{n}} D_n(G)^{-\frac{r}{n}} + 1)^{-i}.$$

Extract r-th roots above, and let r tend to infinity; it follows that $D_i(G) \geq D_n(G)^{i/n}$, proving Plünnecke's bound.

The case when $D_n > 1$ is similar, except that we reverse the independence graph $I_n(h)$ so that the magnification numbers are now small in size. We leave the details to the reader.

Lastly, we turn to the proof of the Lemmas. The first lemma is straightforward, but the second requires more careful thought, and relies on Menger's theorem from graph theory.

Proof of Lemma 1. Let V_i denote the vertex sets of G and W_i the vertex sets of H. Let A be a subset of V_0 with $|\mathrm{Im}_i(A)| = D_i(G)|A|$ and B a subset of W_0 with $|\mathrm{Im}_i(B)| = D_i(H)|B|$. Then $\mathrm{Im}_i(A \times B) = \mathrm{Im}_i(A) \times \mathrm{Im}_i(B)$ from which it follows that $|\mathrm{Im}_i(A \times B)| = D_i(G)D_i(H)|A \times B|$. This shows that $D_i(G \times H) \leq D_i(G)D_i(H)$.

We must now show that $D_i(G \times H) \geq D_i(G)D_i(H)$. Let X be a non-empty subset of $V_0 \times W_0$. We write X as the disjoint union of $\{v\} \times H_v$ where H_v denotes the set of all $w \in W_0$ with $(v, w) \in X$. Of course we need consider only those v for which H_v is non-empty. Let $Y \subset V_0 \times W_i$ denote the union of $\{v\} \times \mathrm{Im}_i(H_v)$. Then note that

$$|Y| = \sum_{v \in V_0} |\mathrm{Im}_i(H_v)| \geq D_i(H) \sum_{v \in V_0} |H_v| = D_i(H)|X|.$$

Now write Y as the union of sets $G_w \times \{w\}$ where w runs over elements in W_i. If $Z \subset V_i \times W_i$ denotes the union of $\mathrm{Im}_i(G_w) \times \{w\}$ then it is plain that $Z = \mathrm{Im}_i(X)$. Moreover

$$|Z| = \sum_{w \in W_i} |\mathrm{Im}_i(G_w)| \geq D_i(G) \sum_{w \in W_i} |G_w| = D_i(G)|Y| \geq D_i(G)G_i(H)|X|,$$

which completes our proof.

For the proof of Lemma 2 we require Menger's theorem from graph theory (see for example, Bollabas's *Modern Graph Theory*).

Menger's Theorem. *Let G be any graph, and let a and b be two distinct vertices with $(a.b)$ not an edge. Then the maximum number of vertex disjoint paths from a to b equals the minimum number of vertices separating a from b. Here two paths from a to b are called vertex disjoint if they share no vertices in common apart from a and b. A set of vertices is said to separate a and b if every path from a to b must contain one of these vertices.*

Proof of Lemma 2. Introduce two vertices a and b at the ends of our Plünnecke graph, connecting a to all vertices in V_0 and connecting all vertices in V_n to b. We wish to show that there are $|V_0|$ vertex disjoint paths from a to b. Suppose the maximum number of such paths is m, and our goal is to show that $m \geq |V_0|$. Let $\pi_1, \ldots, \pi_m$ be m such vertex disjoint paths. By Menger's theorem there exists a set S of m vertices $s_1, \ldots, s_m$ separating a and b. Suppose that we have labeled these vertices so that s_i lies on path π_i. In addition we choose our separating set S so that

$$\sum_{i=0}^{n} i|S \cap V_i|$$

is minimal.

Claim: Such a minimal set S is contained in $V_0 \cup V_n$.

Assuming the claim, we can finish the proof of Lemma 2. Consider the vertices in V_0 that are not in S. Since $D_n \geq 1$ these vertices must have an image in V_n of at least the same size. All those image vertices must necessarily be in S, completing our proof. Actually, this argument shows that (by the minimality property of S) our separating set is in fact V_0.

It remains lastly to prove the Claim. Suppose not. By rearranging if needed, we may assume that $s_1, \ldots, s_q$ are the elements of S lying in some V_i with $i \neq 0, n$. Let s_j^- denote the predecessor of s_j on path π_j, and s_j^+ denote its successor on that path. By the minimality condition imposed on S, we know that the set $\{s_1^-, \ldots, s_q^-, s_{q+1}, \ldots, s_m\}$ is *not* a separating set. This means that there exists some path π which does not contain any element of that set. Therefore there must exist some s (which is in $\{s_1, \ldots, s_q\}$ and thus in V_i) which lies on the path π. Let $r \in V_{i-1}$ denote the predecessor of s on π. Note that r is not equal to s_j^- for $j = 1, \ldots, q$.

Consider three sets of vertices $A = \{s_1^-, \ldots, s_q^-, r\} \subset V_{i-1}$, $B = \{s_1, \ldots, s_q\}$ and $C = \{s_1^+, \ldots, s_q^+\}$. Consider the graph induced by our original Plünnecke graph on these sets A, B and C: that is, consider the graph H on three layers A, B, and C with an edge connecting vertices in A and B (or B and C) precisely if that edge belonged to the graph G. Observe that every path (in G) from A to C must pass through B (that is, it remains a path in H); else, one cound extend that path into a path from a to b avoiding the set S. It follows that our subgraph H is Plünnecke.

Consider now our Plünnecke subgraph H. We use the notation of Example 5 above. The number of edges emanating from A is equal to the number of edges entering B

$$d_+(r) + \sum_{j=1}^{q} d_+(s_j^-) = \sum_{j=1}^{q} d_-(s_j).$$

Since H is Plünnecke we know that $d_-(s_j) \leq d_-(s_j^+)$ and that $d_+(s_j) \leq d_+(s_j^-)$. Thus the above is

$$\leq \sum_{j=1}^{q} d_-(s_j^+) = \sum_{j=1}^{q} d_+(s_j) \leq \sum_{j=1}^{q} d_+(s_j^-),$$

which is a contradiction. This settles our Claim, and the Lemma.

3. Freiman's theorem in a bounded torsion group.

Plünnecke's results apply not just to addition in $\mathbb{Z}$, but within any abelian group (check). If we consider a *torsion group* where each element has order at most r, then we can deduce the following elegant quantitative analog of Freiman's theorem (due to Ruzsa).

Ruzsa's version of Freiman's theorem in a bounded torsion group. *Let G be an abelian group, such that every element of G has order at most r. Let A be a finite subset of G such that $|A+A| \leq C|A|$. Then A is contained in a subgroup H with $|H| \leq C^2 r^{C^4}|A|$.*

Proof. It is helpful first to consider the case when A is symmetric: that is $a \in A$ implies $-a \in A$. We will find a set $X \subset 3A$ of size $\leq C^4$ such that the subgroup generated by A is contained in $2A + \langle X \rangle$ where $\langle X \rangle$ denotes the subgroup generated by X. Clearly $|\langle X \rangle| \leq r^{|X|} \leq r^{C^4}$ and $|2A| \leq C|A|$ by assumption, and a stronger inequality than claimed in the Theorem follows. We choose X to be a maximal subset of $3A$ such that the translates $A + x$ for $x \in X$ are all disjoint. Since the elements of $A + x$ all lie in $4A$ which has size $\leq C^4|A|$, it follows that $|X| \leq C^4$ as desired. By maximality, we see that if $t \in 3A$ then $A + t$ intersects $A + x$ for some $x \in X$. This implies that $t \in A - A + X = 2A + X$. Therefore $3A \subset 2A + X$, and hence $4A \subset 3A + X \subset 2A + 2X$, and so on. Therefore

$$\langle A \rangle = \cup_{j \geq 1} jA \subset 2A + \langle X \rangle,$$

as we wanted.

When A is not symmetric, the same argument works with a tiny modification. We consider maximal $X \subset 2A - A$ with $A + x$ disjoint. Then it follows that $2A - A \subset A - A + X$, and iterating this we get $jA - A \subset A - A + \langle X \rangle$. Hence $\langle A \rangle \subset A - A + \langle X \rangle$, and since $|A - A| \leq C^2$ we recover precisely the estimate of the Theorem.

4. Freiman Homomorphisms and Ruzsa's embedding lemma.

Definition. *Let A and B be subsets of some additive groups G and H. A Freiman k-homomorphism from A to B is a map ϕ such that if $x_1 + \ldots + x_k = y_1 + \ldots + y_k$ is a relation among elements in A then $\phi(x_1) + \ldots + \phi(x_k) = \phi(y_1) + \ldots + \phi(y_k)$ holds as a relation in B. If ϕ is invertible, and gives a k-homomorphism from B to A, then we say that ϕ is a k-isomorphism.*

Remark 1. Note that if ϕ is a k-homomorphism then it induces a map from kA to kB by sending $x_1 + \ldots + x_k$ to $\phi(x_1) + \ldots + \phi(x_k)$. If ϕ is a k-isomorphism, this map is a bijection and so $|kA| = |kB|$.

Remark 2. A 2-isomorphism preserves arithmetic progressions. For, if $a + c = b + b$ are three consecutive terms of an AP, then $\phi(a) + \phi(c) = \phi(b) + \phi(b)$ as well. Moreover,

a 2-isomorphism preserves GAPs. Precisely, suppose we have a GAP (in A) given by $a_0 + a_1x_1 + \ldots + a_kx_k$ with $0 \le x_j \le n_j$, then its image is the GAP $\phi(a_0) + (\phi(a_0 + a_1) - \phi(a_0))x_1 + (\phi(a_0 + a_2) - \phi(a_0))x_2 + \ldots + (\phi(a_0 + a_k) - \phi(a_0))x_k$.

Remark 3. If $\phi_1 : A \to B$ and $\phi_2 : B \to C$ are k-homomorphisms then $\phi_2 \circ \phi_1$ gives a k-homomorphism from A to C. If ϕ_1 and ϕ_2 are isomorphisms then so is the composition $\phi_2 \circ \phi_1$.

Example 1. Let $A \subset \mathbb{Z}$ and let N be any natural number, and consider the reduction map $\pmod N$. This obviously is a k-homomorphism for any k. But usually it is not an isomorphism. If N is very large then this map can be inverted. Precisely, if A lies in an interval of length I then for $N > I$ we can invert the map, and for $N > kI$ the inverse map is a k-homomorphism.

Example 2. Let q be a number coprime to N and consider the map $\phi : \mathbb{Z}/N \to \mathbb{Z}/N$ given by $x \to qx$. This is a k-isomorphism for any k.

Example 3. This is closely connected to example 1. Consider $\phi : \mathbb{Z}/N \to \mathbb{Z}$ by choosing a representative for a congruence class lying in $[1, N]$ and identifying that class with this integer. This is not a k-homomorphism as it stands, but if we restrict ϕ to a subset of $\mathbb{Z}/N$ such as $\{jN/k < x \le (j + 1)N/k\}$ then it is a k-homomorphism.

Example 4. Let A be a finite set in $\mathbb{Z}^d$. By translating (which is an isomorphism of all orders) we may assume that the elements of A all have positive coordinates. Pick a base b which is very large and exceeds k times the maximum coordinate of any element of A. Map $(a_1, \ldots, a_d) \in A$ to $a_1b + a_2b^2 + \ldots + a_db^d \in \mathbb{Z}$. This map gives a k-isomorphism, because the base has been chosen so large that there are no carries. This example can be used to extend Freiman's theorem from $\mathbb{Z}$ to $\mathbb{Z}^d$ for any d.

Example 5. Let A be a sparse set; for example, $A = \{1, 2, 4, \ldots, 2^{n-1}\}$. We claim that A is not 2-isomorphic to a subset B of $\mathbb{Z}/N$ for $N < n(n + 1)/2$. This follows since $|2A| = n(n + 1)/2 = |2B| \le N$.

Example 5 shows that not all sets can be embedded into $\mathbb{Z}/N$ for a small value of N. Our aim in this section is to show that if A is a set of integers with $|kA - kA| \le C|A|$ then we may extract a large subset A' of A such that A' is k-isomorphic to a subset of $\mathbb{Z}/N$ where N is a relatively small prime.

Ruzsa's Embedding Lemma. *Let A be a set of integers with $|kA - kA| \le C|A|$. Then for any prime $N > 2C|A|$ we may find a subset A' of A with $|A'| \ge |A|/k$ such that A' is k-isomorphic to a subset of $\mathbb{Z}/N$.*

Before we do this, we present a warm-up problem of Erdős which may illuminate the construction.

Proposition 1. *If A is any set of N integers, then we may extract a sum-free subset of A with cardinality $\ge N/3$.*

Proof. Pick a very large prime p. For each $1 \le q \le p - 1$ define A_q to be the elements in

A such that qa is congruent to a number in $(p/3, 2p/3)$. The probability that this happens is $1/3$, and so for some q we must have $|A_q| \geq |A|/3$. This is the desired subset of A: for, if $a+b=c$ in A_q then $qa+qb=qc$, but this is impossible as a congruence $\pmod p$.

Note that the proof implicitly uses the constructions of Examples 2 and 3, together with a simple probabilistic argument.

Proof of Ruzsa's Embedding Lemma. Let p be a very large prime. For each $1 \leq q \leq p-1$ we find a subset A_q of A such that $|A_q| \geq |A|/k$ and that the dilates qa for each $a \in A_q$ all lie in an interval $(jp/k, (j+1)p/k) \pmod p$ for some $0 \leq j \leq k-1$. Clearly such A_q exists by the pigeon-hole principle.

Now consider the map from A_q to $[1,p]$ by sending a to qa and then reducing that $\pmod p$ to get a representative in $[1,p]$. By construction the image of this map is in $[jp/k, (j+1)p/k]$ for some j, and this map is a k-isomorphism.

Take the image of the previous map, and view it $\pmod N$. Thus we have a map from A_q to $\mathbb{Z}/N$ and this is plainly a k-homomorphism. The question is if it can be inverted and is a k-isomorphism. We will show that if N is suitably large, then for some q it is indeed a k-isomoprhism.

If the map (call it ϕ_q) is not a k-isomorphism, then there must be a coincidence $\phi_q(a_1)+\ldots+\phi_q(a_k) = \phi_q(b_1)+\ldots+\phi_q(b_k)$ where the a_i and b_i are in A_q with $\sum a_i \neq \sum b_i$. This means that $(qa_1)_p + \ldots (qa_k)_p \neq (qb_1)_p + \ldots + (qb_k)_p$ (where $(n)_p$ denotes the reduction of $n \pmod p$ taken in $[1,p]$), but that the two sides are congruent $\pmod N$. That is for some non-zero number ℓ with $|\ell| \leq p/N$ we have

$$\ell N = (qa_1)_p + \ldots + (qa_k)_p - (qb_1)_p - \ldots - (qb_k)_p.$$

Note that we can get away with $|\ell| \leq p/N$ because by construction all $(qa)_p$ lie in an interval of size p/k. Viewing this relation $\pmod p$ we obtain that

$$\ell N \equiv q(a_1 + \ldots + a_k - b_1 - \ldots - b_k) \pmod p.$$

Given $a_1+\ldots+a_k-b_1-\ldots-b_k$ there are at most $2p/N$ bad values of q for which this can happen. Further, the number of choices for $a_1+\ldots+a_k-b_1-\ldots-b_k$ is at most $|kA-kA| \leq C|A|$. Therefore the total number of bad values of q for which a coincidence can occur is $\leq C|A|(2p/N)$, so that if $N > 2C|A|$ then there will be some value of q with no coincidences, giving us the desired isomorphism.

5. Bogolyubov's Lemma and Bohr sets.

Bogolyubov's Lemma says that if A is a subset of $\mathbb{Z}/N\mathbb{Z}$ with $|A| = \delta N$ then $2A-2A$ contains a highly structured set known as a Bohr set.

Definition. *Let N be a large prime and $r_1, \ldots, r_k$ be k distinct reduced residue classes* $\pmod N$. *We define the Bohr set $\mathcal{B}(r_1, \ldots, r_k; \delta_1, \ldots, \delta_k)$ to be the set of all residue classes s such that $\|sr_j/N\| \leq \delta_j$ for each $j = 1, \ldots, k$. If $\delta_1 = \delta_2 = \ldots = \delta_k$ and $K = \{r_1, \ldots, r_k\}$ then we will abbreviate the Bohr set $\mathcal{B}(r_1, \ldots, r_k; \delta_1, \ldots, \delta_k)$ as $\mathcal{B}(K; \delta)$. We will refer to k as the* dimension *of our Bohr set.*

We postpone to the next section a discussion of the structure of Bohr sets: the main input there will come from Minkowski's *geometry of numbers*.

Bogolyubov's Lemma. *Let A be a subset of $\mathbb{Z}/N$ with $|A| = \delta N$. Then $2A-2A$ contains a Bohr set $B(K;1/4)$ of dimension $k \le 1/\delta^2$.*

Proof. Let as before $\hat{A}(r)$ denote the Fourier transform of our set A. Note that

$$\sum_r |\hat{A}(r)|^2 = N|A| = \delta N^2,$$

by Parseval, and so there are at most $1/(\delta\lambda^2)$ large Fourier coefficients $|\hat{A}(r)| \ge \lambda|A|$. Of course $r = 0$ is one of these large Fourier coefficients, and let $r_1, \dots, r_k$ be the non-zero large Fourier coefficients. We claim that if $\lambda^2 = \delta$ is small enough, then the Bohr set $\mathcal{B}(r_1, \dots, r_k; 1/4)$ is contained in $2A - 2A$, which establishes Bogolyubov's Lemma.

Let b be an element of this Bohr set. Consider

$$\begin{aligned}\sum_{r \pmod N} |\hat{A}(r)|^4 e(br/N) &= \sum_{r \pmod N} |\hat{A}(r)|^r \cos(2\pi br/N) \\ &= N\#\{b \equiv a_1 + a_2 - a_3 - a_4 \pmod N\},\end{aligned}$$

and we must show that this is positive. The contribution of $r = 0$ is $|A|^4$. The contribution of r_j for all the large Fourier coefficients is positive, because $\cos(2\pi br_j/N) \ge \cos(\pi/2) = 0$ for elements in our Bohr set. Finally the contribution of the small Fourier coefficients (those less than $\lambda|A|$ in size) is in magnitude

$$< \lambda^2|A|^2 \sum_{r \pmod N} |\hat{A}(r)|^2 = \lambda^2|A|^3 N = |A|^4,$$

by our choice of λ. It follows that our sum above is indeed positive, and therefore $b \in 2A - 2A$ as desired.

6. The structure of Bohr sets: Input from the geometry of numbers.
We will show here that the Bohr sets have a lot of structure, and contain big generalized arithmetic progressions.

Remark 1. First we note that $\mathcal{B}(K;\delta)$ is pretty big. This is Dirichlet's theorem on Diophantine approximation. Divide $[0,1)^k$ into cubes of size δ^k. There are about $(1/\delta)^k$ such cubes. Now consider the N points $(nr_1/N, nr_2/N, \dots, nr_k/N)$ viewed (mod 1). By the pigeonhole principle one cube contains about $\delta^k N$ such points. The difference set of those points is contained in the desired Bohr set. Thus $|\mathcal{B}| \gg \delta^k N$.

Remark 2. Take an element in $\mathcal{B}(K; 2N^{-\frac{1}{k}})$. By our previous remark there is a non-zero $n \in \mathbb{Z}/N$ in this Bohr set. Clearly $n\ell$ will belong to $\mathcal{B}(K;\delta)$ for each $1 \le \ell \le \delta N^{\frac{1}{k}}/2$. Thus $\mathcal{B}(K;\delta)$ contains a 1-dimensional arithmetic progression of size $\ge \delta N^{\frac{1}{k}}/2$.

Our goal in this section is to show the following Proposition.

Proposition. *The Bohr set $\mathcal{B}(K;\delta)$ contains a proper GAP of dimension k and size at least $(\delta/k)^k N$.*

To achieve this we require some input from the *geometry of numbers*. We review briefly the facts that we will need; for further information the reader may consult Siegel's beautiful

Lectures on the geometry of numbers, or Cassels's *An Introduction to the Geometry of Numbers*.

A *lattice* of $\mathbb{R}^k$ is a discrete subgroup of $\mathbb{R}^k$. Equivalently, a lattice may be described as the subgroup generated by linearly independent vectors $v_1, \ldots, v_\ell$. The lattice is said to be of full rank if $\ell = k$.

Example 1. $\mathbb{Z}^k$ is a full rank lattice.

Example 2. Obviously $N\mathbb{Z}^k$ is a sublattice of the lattice $\mathbb{Z}^k$. Of particular interest to us is the lattice generated by $N\mathbb{Z}^k$ and the vector $(r_1, \ldots, r_k)$. For example, if $k = 2$, $N = 5$ and $(r_1, r_2) = (1, 1)$ then this lattice is generated by $(1, 1)$ and $(5, 0)$.

From now on, assume that our lattices have full rank. Given a lattice Λ with basis v_1, $\ldots$, v_k, a fundamental domain for $\mathbb{R}^k/\Lambda$ is the parallelopiped $\{\sum_{j=1}^k x_j v_j : \ \ 0 \le x_j < 1\}$. Note that every element of $\mathbb{R}^k$ can be written uniquely as a lattice vector plus an element from this fundamental parallelopiped. The volume of a fundamental parallelopiped is an invariant of the lattice, independent of the choice of the basis. We denote this volume by $\mathrm{Vol}(\Lambda)$.

Example 3. The lattice $\mathbb{Z}^k$ has volume 1. The lattice $N\mathbb{Z}^k$ has volume N^k. The lattice generated by $N\mathbb{Z}^k$ and $(r_1, \ldots, r_k)$ has volume N^{k-1} (assuming one of the r_j's is coprime to N).

An open convex set C in $\mathbb{R}^k$ will be called a convex body. The body is centrally symmetric (about the origin) if $x \in C$ implies $-x \in C$.

Blichtfeld's Lemma. *Let C be an open set and Λ a lattice. If the volume of C exceeds the volume of Λ then there are two points $x \neq y$ in C with $x - y \in \Lambda$.*

Proof. For each point $\lambda \in \Lambda$ consider the translate $\lambda + C$. If these translates intersect, then we are done. Suppose they are all disjoint. Consider a big box B. The box contains about $\mathrm{Vol}(B)/\mathrm{Vol}(\Lambda)$ lattice points. Thus the set of translates $\lambda + C$ for $\lambda \in B$ has volume at least about $\mathrm{Vol}(C)\mathrm{Vol}(B)/\mathrm{Vol}(\Lambda)$, but this should still be contained in a box just slightly larger than B. This gives a contradiction.

Minkowski's First Theorem. *Let C be a convex, centrally symmetric body such that $Vol(C) > 2^k\, Vol(\Lambda)$. Then C contains a non-zero lattice point.*

Proof. Use Blichtfeld's lemma for $\frac{1}{2}C$, and note that $C = \frac{1}{2}C - \frac{1}{2}C$.

Given a convex body C we let λC denote the dilated body $\{\lambda x : x \in C\}$. Let λ_1 be the infimum of values λ such that λC contains a non-zero lattice point. We may express Minkowski's first theorem as saying that $\lambda_1^k \le 2^k \mathrm{Vol}(\Lambda)/\mathrm{Vol}(C)$. In general this is best possible; consider $\Lambda = \mathbb{Z}^k$ and $C = (-1, 1)^k$. But in the above example we see that the closure of C contains not just one, but k linearly independent lattice points. This suggests a refinement of the bound in Minkowski's first theorem. Define the *successive minima* $\lambda_1 \le \lambda_2 \le \lambda_3 \le \ldots \le \lambda_k$ by letting λ_j denote the infimum of λ such that λC contains j linearly independent lattice vectors.

Minkowski's Second Theorem. *With the above notation we have*

$$\lambda_1 \cdots \lambda_k \le 2^k\, Vol(\Lambda)/\, Vol(C).$$

This remains a beautiful and deep result, for whose proof we refer to the books of Cassels or Siegel mentioned earlier. We are now in a position to establish the structure of Bohr sets.

Proof of the Proposition. We consider the full rank lattice Λ generated by $N\mathbb{Z}^k$ and $(r_1, \dots, r_k)$ (see examples 2 and 3). Its volume is N^{k-1}. Let C be the convex body $(x_1, \dots, x_k)$ with $|x_j| \le 1$. The volume of C is plainly 2^k. By Minkowski's second theorem there exist k linearly independent lattice points $\mathbf{b}_1, \dots, \mathbf{b}_k$ with $\mathbf{b}_j \in \lambda_j C$, and $\lambda_1 \cdots \lambda_k \le N^{k-1}$. Consider linear combinations $\sum_{j=1}^k \mathbf{b}_j n_j$ where $|n_j| \le \delta N/(k\lambda_i)$. It is clear that the coordinates of these vectors are bounded by $\sum_{j=1}^k \lambda_j(\delta N/(k\lambda_j)) = \delta N$. Furthermore, if we write each $\mathbf{b}_i$ as $(b_i r_1, \dots, b_i r_k) \pmod N$ then we have just shown that $\sum_{j=1}^r b_j n_j$ (with $|n_j| \le \delta N/\lambda_j$) lies in our Bohr set $\mathcal{B}(\{r_1, \dots, r_k\}; \delta)$. This is the sought-for GAP of dimension k.

It remains to show that our GAP is proper, and then evaluate its size. If $\sum_{j=1}^k b_j n_j = \sum_{j=1}^k b_j n_j'$ then it follows that $\sum_{j=1}^k \mathbf{b}_j n_j \equiv \sum_{j=1}^k \mathbf{b}_j n_j' \pmod N$. However the coordinates of both these vectors are all $\le \delta N$ in size; thus congruence $\pmod N$ implies equality. Linear independence of the $\mathbf{b}_i$'s now shows that $n_j = n_j'$. Thus our GAP is proper. Its size is plainly

$$\prod_{j=1}^k \Big(1 + 2\Big[\frac{\delta N}{k\lambda_j}\Big]\Big) \ge \prod_{j=1}^k \frac{\delta N}{k\lambda_j} \ge \Big(\frac{\delta}{k}\Big)^k N,$$

using the Minkowski bound.

7. The proof of Freiman's Theorem.

Let A be a subset of the integers with $|2A| \le C|A|$.

Step 1. By Plünnecke's inequality we then know that

$$|8A - 8A| \le C^{16}|A|.$$

Step 2. Applying now Ruzsa's embedding lemma, we may find $A_1 \subset A$ with $|A_1| \ge |A|/8$ such that A_1 is 8-isomorphic to a subset A_2 of $\mathbb{Z}/N$ with N being a prime between $2C^{16}|A|$ and $4C^{16}|A|$.

Step 3. By Bogolyubov's Lemma $2A_2 - 2A_2$ contains a Bohr set $\mathcal{B}(K; 1/4)$ of dimension $\le 1024C^{32}$.

Step 4. From the geometry of numbers we conclude that $2A_2 - 2A_2$ contains a proper generalized arithmetic progression of dimension $1024C^{32}$ and size $\gg e^{-C^{33}}|A_2|$. Now $2A - 2A$ is 2-isomorphic to $2A_2 - 2A_2$ and since arithmetic progressions are preserved under 2-isomorphisms, $2A - 2A$ contains a proper generalized arithmetic progression of dimension $1024C^{32}$ and size $\gg e^{-C^{33}}|A|$. Call this progression Q, and throw away excess elements if necessary so that Q is of size about $e^{-C^{33}}|A|$.

Take X to be a maximal set in A such that the translates $Q+x$ for $x \in X$ are all disjoint. Plainly $|X| = |Q + X|/|Q| \le C^5|A|/|Q|$ since $Q + X \subset 3A - 2A$ and using Plünnecke's inequality. Furthermore, by the maximality of x we see that for any $a \in A$ we must have $Q + a \cap Q + X \ne \emptyset$ which implies that $A \subset X + Q - Q$. Now $Q - Q$ is a generalized arithmetic progression of the same dimension as Q, and size $\le 2^{1024C^{32}}|Q| \ll |A|$. Viewing every element of X as contributing an extra dimension, we obtain Freiman's theorem.

8. Chang's refinement: getting a polynomial bound for the dimension.

We now give Chang's refinement which allows for a polynomial bound for the dimension. Recall that after step 4 (of §7) we produced a proper GAP Q in $2A-2A$ with dimension $1024C^{32}$ and size about $\exp(-C^{33})|A|$. Set $Q_0 = Q$ and let Y_0 denote a maximal set in A such that $Q_0 + y$ are disjoint for $y \in Y_0$. If $|Y_0| < 2C$ then we declare ourselves done, else we pick X_0 to be a subset of Y_0 with cardinality $2C$, and then set $Q_1 = Q_0 + X_0$. Now we repeat the same process, selecting a maximal Y_1 in A such that $Q_1 + y$ are disjoint for $y \in Y_1$, and picking a subset X_1 of Y_1 with cardinality $2C$. Proceed in this manner until we terminate in a set Y_t of cardinality $< 2C$. By the maximality of Y_t we get that

$$A \subset Y_t + Q_t - Q_t = Y_t + \sum_{j=0}^{t-1}(X_j - X_j) + (Q - Q).$$

The set on the RHS above is a GAP of dimension $\leq 2C(t+1) + 1024C^{32}$ and size $\leq 3^{2C(t+1)}2^{1024C^{32}}|Q| \ll 3^{2C(t+1)}|A|$. It remains now to estimate t.

Note that

$$|Q_t| = |Q_{t-1}||X_{t-1}| = 2C|Q_{t-1}| = \ldots = (2C)^t|Q|.$$

On the other hand,

$$Q_t \subset A + Q_{t-1} \subset \ldots \subset tA + Q \subset (t+2)A - 2A,$$

so that $|Q_t| \leq C^{t+4}|A|$ by Plünnecke. Since $|Q|$ is about $e^{-C^{33}}|A|$, estimates it follows that $t \ll C^{33}$. Thus the dimension of the GAP in Freiman's theorem is $\ll C^{34}$ and the size is $\ll \exp(C^{34})|A|$.

By adding some extra ingredients, including Chang's structure theorem on large Fourier coefficients, one can further reduce the bound on the dimension of the GAP is Freiman's theorem to $\ll C^2(\log C)^2$, and size factor S may be made $\leq \exp(2^{20}C^2(\log C)^2)$. The dimension of the GAP in Freiman's theorem cannot be made smaller than C: to see this, take A to be the sumset of the interval $[1, N/C]$ and a set S of cardinality C which is very sparse. Then the sumset $A + A$ is about size $C|A|$, but one cannot hope to realize A as a large subset of a GAP with dimension significantly smaller than C.

We end by giving a small variant of Freiman's result which will be useful in Gowers's proof.

Freiman Variant. *Let A be a subset of $\mathbb{Z}$ with $|A+A$ or $|A-A| \leq C|A|$. There exists a proper progression Q_0 of dimension $\leq 2048C^{32}$ and size $\gg e^{-C^{33}}|A|$ such that*

$$|A \cap Q_0| \gg C^{-16}|Q_0| \gg e^{-C^{34}}|A|.$$

Proof. We follow the argument of §7, making a few small changes. Most notably, in Bogolyubov's Lemma if we allow the dimension to go up to $2/\delta^2$ then we can ensure that every element in that Bohr set $\mathcal{B}(K; 1/4)$ has at least $|A|^4/(2N)$ representations as $a_1 + a_2 - a_3 - a_4$. Following now the argument earlier up to step 4, we find that $2A - 2A$ contains a proper GAP Q of dimension $2048C^{32}$ and size about $\exp(-C^{33})|A|$,

and moreover each element in Q has $\gg C^{-16}|A|^3$ representations as $a_1+a_2-a_3-a_4$. Thus

$$\sum_{a_1,a_2,a_3,a_4\in A} Q(a_1+a_2-a_3-a_4)=\sum_{q\in Q}\#\{q=a_1+a_2-a_3-a_4\}\gg |Q|C^{-16}|A|^3,$$

and so there exists some choice of a_2, a_3, a_4 such that

$$\sum_{a_1} Q(a_1+a_2-a_3-a_4)=A\cap(Q-a_2+a_3+a_4)\gg C^{-16}|Q|.$$

Taking $Q_0=Q-a_2+a_3+a_4$ we have proved our result.

VIII. Gowers's proof of Szemerédi's theorem for four term progressions

1. Roth's theorem revisited: The Gowers U^2 norm.
Let $f : \mathbb{Z}/N \to \mathbb{C}$ be given. We define its Gowers U^2 norm by

$$\|f\|_{U^2}^4 = \frac{1}{N^3} \sum_{a,b,c} f(a)\overline{f(a+b)f(a+c)}f(a+b+c).$$

Thus, this is an average of f over *parallelograms* a, $a+b$, $a+c$, $a+b+c$. Note that a three term progression a, $a+b$, $a+2b$ arises as a special parallelogram: namely a square with $b = c$. We will now see that the U^2 norm controls the distribution of functions in three term progressions.

Lemma 1. *Let f_1, f_2 and f_3 be three functions from $\mathbb{Z}/N$ to $\mathbb{C}$. Then*

$$\Big| \sum_{a,d} f_1(a)f_2(a+d)f_3(a+2d) \Big| \leq N^2 \min_{i=1,2,3} \|f_i\|_{U^2}.$$

Proof. Call the LHS as S. Cauchy's inequality gives that

$$S^2 \leq N \sum_a \Big| \sum_d f_2(a+d)f_3(a+2d) \Big|^2 = N \sum_a \sum_{d,e} f_2(a+d)\overline{f_2(a+e)}f_3(a+2d)\overline{f_3(a+2e)}.$$

We reparametrize this, setting $A = a+d$ and $A+k = a+e$, so that $a+2d = A+d$ and $a+2e = A+d+2k$. So, we get

$$S^2 \leq N \sum_{A,k} \Big| \sum_d f_3(A+d)\overline{f_3(A+d+2k)} \Big|.$$

Replace above $2k$ by k (assuming that N is odd), and use Cauchy's inequality again. Thus

$$S^4 \leq N^4 \sum_{A,k} \sum_{d,e} f_3(A+d)\overline{f_3(A+e)f_3(A+d+k)}f_3(A+e+k).$$

Writing $a = A+d$, $a+b = A+e$ (so that $b = e-d$), $c = k$, we obtain that the above equals

$$N^5 \sum_{a,b,c} f_3(a)\overline{f_3(a+b)f_3(a+c)}f_3(a+b+c) = N^8 \|f_3\|_{U^2}^4.$$

We have shown that $S \leq N^2 \|f_3\|_{U^2}$. By symmetry, the Lemma follows.

Lemma 2. *Let A be a subset of $[1, N]$ with $|A| = \delta N$. View A in $\mathbb{Z}/N$, and let f denote the balanced function of A: $f(n) = 1-\delta$ if $n \in A$, and $f(n) = -\delta$ otherwise. If $\|f\|_{U^2} \leq \delta^3/32$, then either there exist $N^2\delta^3/32$ three term progressions in A, or there exists a sub-progression of length $N/3$ on which A has density $\geq 9\delta/8$.*

Proof. Let A_1 and A_2 denote $A \cap [N/3, 2N/3]$. If this set has cardinality $\leq \delta N/4$ then one of the two sets $A \cap [1, N/3]$ or $A \cap [2N/3, N]$ must have cardinality $\geq 3\delta N/8$ and we would be done. If $|A \cap [N/3, 2N/3]| \geq \delta N/4$ then we see that the number of (genuine) three term progressions in A is at least

$$\sum_{a,d} A_1(a)A_2(a+d)A(a+2d) = \delta \sum_{a,d} A_1(a)A_2(a+d) + \sum_{a,d} A_1(a)A_2(a+d)f(a+2d)$$
$$\geq \frac{\delta^3}{16}N^2 - \|f\|_{U^2}N^2,$$

using Lemma 1. This proves the Lemma.

What does it mean for the U^2 norm to be large? Here we must fall back on Fourier coefficients. Notice that

$$\|f\|_{U^2}^4 = \frac{1}{N^4}\sum_k |\hat{f}(k)|^4.$$

Therefore, using Parseval we see that if $\|f\|_{U^2} \geq \delta^3/32$ (as in Lemma 2) then $\delta^{12}N^2 \ll \max_k |\hat{f}(k)|^2$ so that $\max_k |\hat{f}(k)| \gg \delta^6 N$. This is similar to, but weaker than, the criterion we obtained directly in Roth's proof. What has been gained is that the definition of the U^2 norm lends itself readily to generalization, unlike Roth's proof.

2. The Gowers U^3 norm and four term progressions.
Analogously to the U^2 norm, we define the U^3 norm of $f : \mathbb{Z}/N \to \mathbb{C}$ by

$$\|f\|_{U^3}^8 = \frac{1}{N^4}\sum_{a,b,c,d} f(a)\overline{f(a+b)f(a+c)f(a+d)}f(a+b+c)f(a+b+d)f(a+c+d)\overline{f(a+b+c+d)}.$$

This is an average of f over *parallelopipeds*. Now define

$$\Delta(f,k)(n) = f(n)\overline{f(n+k)}.$$

Then we see that

$$\|f\|_{U^3}^8 = \frac{1}{N}\sum_b \|\Delta(f,b)\|_{U^2}^4.$$

Just as the U^2 norm controls three term progressions, the U^3 norm controls four term progressions.

Lemma 1. *Let f_1, f_2, f_3, f_4 be functions from $\mathbb{Z}/N$ to $\mathbb{C}$. Then*

$$\Big|\sum_{a,d} f_1(a)f_2(a+d)f_3(a+2d)f_4(a+3d)\Big| \leq N^2 \min_{i=1,2,3,4}\|f_i\|_{U^3}.$$

Proof. Call the sum in question S. Cauchy's inequality gives that

$$S^2 \leq N\sum_a \Big|\sum_d f_2(a+d)f_3(a+2d)f_4(a+3d)\Big|^2$$
$$= N\sum_a\sum_{d,e} f_2(a+d)\overline{f_2(a+e)}f_3(a+2d)\overline{f_3(a+2e)}f_4(a+3d)\overline{f_4(a+3e)}.$$

Rewrite $a+d$ as a, and write $e=d+k$. Then the above is

$$N\sum_k\sum_{a,d}\Delta(f_2,k)(a)\Delta(f_3,2k)(a+d)\Delta(f_4,3k)(a+2d).$$

The inner sums over a and d are now over various functions evaluated on three term progressions. Therefore, by Lemma 1, the above is

$$\leq N^3\sum_k\|\Delta(f_4,3k)\|_{U^2}.$$

Replacing $3k$ by k (assuming N is not divisible by 3), Hölder's inequality now gives that this is

$$\leq N^4\|f_4\|_{U^3}^2.$$

The Lemma follows.

Lemma 2. *Let A be a subset of $[1,N]$ with $|A|=\delta N$. We view A in $\mathbb{Z}/N$ and let f denote the balanced function of A. If $\|f\|_{U^3}\leq\delta^4/144$ then either there exist $N^2\delta^4/72$ four term progressions in A, or there exists a sub-progression of length $2N/5$ on which A has density $\geq 25\delta/24$.*

Proof. Let $A_1=A_2=A\cap[2N/5,3N/5]$. If $|A_1|\leq\delta N/6$ then either $A\cap[1,2N/5]$ or $A\cap[3N/5,N]$ has size at least $5\delta N/12$, so that there would be a sub-progression of length $2N/5$ where A has density $25\delta/24$. Suppose now that $|A_1|=|A_2|\geq\delta N/6$. The number of genuine four term progressions in A exceeds

$$\sum_{a,d}A_1(a)A_2(a+d)A(a+2d)A(a+3d)=\delta\sum_{a,d}A_1(a)A_2(a+d)A(a+2d)$$
$$+\sum_{a,d}A_1(a)A_2(a+d)A(a+2d)f(a+3d).$$

The second term above is bounded in size by $N^2\|f\|_{U^3}$, by Lemma 1. The first term above equals

$$\delta^2\sum_{a,d}A_1(a)A_2(a+d)+\delta\sum_{a,d}A_1(a)A_2(a+d)f(a+2d)\geq\delta^2|A_1|^2-\delta N^2\|f\|_{U^2},$$

by Lemma 1 of the previous section. Therefore the number of genuine four term progressions in A is at least

$$\delta^4N^2/36-\delta N^2\|f\|_{U^2}-N^2\|f\|_{U^3}.$$

A homework problem shows that $\|f\|_{U^2}\leq\|f\|_{U^3}\leq\|f\|_{U^4}\ldots$, and therefore the Lemma follows.

The argument generalizes in an obvious manner, using $(k-1)$-dimensional parellelopipeds to define a Gowers U^{k-1} norm which controls the distribution of k term arithmetic progressions. Precisely, we define inductively

$$\|f\|_{U^k}^{2^k}=\frac{1}{N}\sum_b\|\Delta(f,b)\|_{U^{k-1}}^{2^{k-1}},$$

and then

$$\Big|\sum_{a,d}f_1(a)f_2(a+d)\cdots f_k(a+(k-1)d)\Big|\leq N^2\min_{i=1,\ldots,k}\|f_i\|_{U^k}.$$

3. Extracting structure out of a large U^3 norm.
Suppose now that we have a function $f : \mathbb{Z}/N \to [-1,1]$ with $\|f\|_{U^3} \geq \alpha$. We wish to extract some additive structure that f must possess. Recall that

$$\|f\|_{U^3}^8 = \frac{1}{N} \sum_{k \pmod N} \|\Delta(f,k)\|_{U^2}^4 \geq \alpha^8.$$

Since each $\|\Delta(f,k)\|_{U^2}$ is ≤ 1, we see that there are at least $\alpha^8 N/2$ values of k for which $\|\Delta(f,k)\|_{U^2}^4 \geq \alpha^8/2$. Let B denote the set of such values k.

For $k \in B$ we know that

$$\begin{aligned}\frac{\alpha^8}{2} \leq \|\Delta(f,k)\|_{U_2}^4 &= \frac{1}{N^4}\sum_{\ell} |\widehat{\Delta(f,k)}(\ell)|^4 \\ &\leq \Big(\max_{\ell} |\widehat{\Delta(f,k)}(\ell)|^2\Big)\frac{1}{N^4}\sum_{\ell} |\widehat{\Delta(f,k)}(\ell)|^2 \leq \frac{1}{N^2}\max_{\ell} |\widehat{\Delta(f,k)}(\ell)|^2,\end{aligned}$$

where we used Parseval at the last step. Therefore there exists some large Fourier coefficient of $\Delta(f,k)$. That is for each $k \in B$ we may find a $\phi(k)$ (making some choice among the large coefficients) such that

$$|\widehat{\Delta(f,k)}(\phi(k))| \geq N\alpha^4/2. \tag{1}$$

Summarizing the argument so far, we have produced a large set B (with at least $\alpha^8 N/2$ elements) such that for $k \in B$ there exists $\phi(k)$ satisfying (1).

The crucial observation in Gowers's proof is that the map $\phi(k)$ is far from arbitrary, and behaves "linearly" for many values of k.

Proposition. *Let $f : \mathbb{Z}/N \to [-1,1]$, and suppose that B and ϕ are as above. Then there are at least $\alpha^{64}2^{-12}N^3$ quadruples $(b_1,b_2,b_3,b_4) \in B^4$ with $b_1+b_2 = b_3+b_4$, and $\phi(b_1)+\phi(b_2) = \phi(b_3)+\phi(b_4)$.*

Proof. From (1) and our definition of B we know that

$$\sum_{k\in B} |\widehat{\Delta(f,k)}(\phi(k))|^2 \geq \frac{\alpha^{16}}{8}N^3. \tag{2}$$

Expanding out $\widehat{\Delta(f,k)}(\phi(k))$ we see that the LHS above equals

$$\begin{aligned}&\sum_{k\in B}\sum_{x,y} f(x)f(x+k)f(y)f(y+k)e\Big(-\frac{(x-y)\phi(k)}{N}\Big) \\ =&\sum_{x,u} f(x)f(x+u)\sum_{k\in B} f(x+k)f(x+u+k)e\Big(-\frac{u\phi(k)}{N}\Big) \\ \leq&\sum_{x,u}\Big|\sum_{k}\Delta(f,u)(x+k)h_u(k)\Big|, \qquad (3)\end{aligned}$$

where we have written $h_u(k) = B(k)e(-u\phi(k)/N)$ with $B(k)$ denoting the characteristic function of B. Using Cauchy's inequality above we conclude that

$$\frac{\alpha^{32}}{64}N^4 \leq \sum_u \sum_x \Big|\sum_k \Delta(f,u)(x+k)h_u(k)\Big|^2. \tag{4}$$

Let us consider the sums over x and k, treating u as fixed. Let $F_u(x) = \sum_k \Delta(f,u)(x+k)h_u(k)$. Then

$$\begin{aligned}\hat{F}_u(r) &= \sum_x F_u(x)e(-xr/N) = \sum_x \sum_k \Delta(f,u)(x+k)e(-(x+k)r/N)h_u(k)e(kr/N)\\ &= \widehat{\Delta(f,u)}(r)\hat{h}_u(-r).\end{aligned}$$

Thus Parseval gives

$$\sum_x |F_u(x)|^2 = \frac{1}{N}\sum_r |\widehat{\Delta(f,u)}(r)|^2|\hat{h}_u(-r)|^2 \leq \frac{1}{N}\Big(\sum_r |\widehat{\Delta(f,u)}(r)|^4\Big)^{\frac{1}{2}}\Big(\sum_r |\hat{h}_u(-r)|^4\Big)^{\frac{1}{2}}.$$

Trivially we see that $|\widehat{\Delta(f,u)}(r)| \leq N$, and by Parseval $\sum_r |\widehat{\Delta(f,u)}(r)|^2 \leq N^2$, and so the above is

$$\leq N\Big(\sum |\hat{h}_u(-r)|^4\Big)^{\frac{1}{2}}.$$

Inputing this into (4) we conclude that

$$\frac{\alpha^{32}}{64}N^3 \leq \sum_u \Big(\sum_r |\hat{h}_u(-r)|^4\Big)^{\frac{1}{2}},$$

which by Cauchy's inequality gives

$$\frac{\alpha^{64}}{2^{12}}N^5 \leq \sum_u \sum_r |\hat{h}_u(-r)|^4.$$

Upon recalling the definition of $h_u(k)$ we see that the RHS above precisely equals N^2 times the number of desired additive quadruples. Therefore, the Proposition holds.

4. Additive Quadruples and The Balog-Szemerédi-GowersTheorem.

Freiman's theorem finds structure in A provided $A+B$ is small for some set B of the same size as A. Suppose instead that we have a set B (with $|B| = |A|$) and we know that for many choices of $(a,b) \in A \times B$ we get $a+b$ lying in a small set, then can we still form any conclusions about A? More precisely, we assume that we are given a subset G of $A \times B$ with $|G| \geq \alpha|A|^2$ and such that $S = \{a+b : \quad (a,b) \in G\}$ is small. Then what can we conclude about A? The answer is that A contains a big subset A', and B contains a big subset B' such that $A'+B'$ is small; and now we can find structure in A' as in Freiman's theorem. This is the Balog-Szemerédi theorem.

We will need a variant of the Balog-Szemerédi theorem due to Gowers. Let A and B be sets with $|A| = |B|$, and $|A + A| \le C|A|$, and let $r_{A+B}(n)$ denote the number of the number of ways of writing n as $a + b$ with $a \in A$ and $b \in B$. Note that

$$\sum_n r_{A+B}(n) = |A||B|,$$

and that $|A + B|$ equals the number of n with $r_{A+B}(n) \neq 0$. Using Cauchy's inequality it follows that

$$\sum_n r_{A+B}(n)^2 \ge |A|^4/|A + B| \ge |A|^3/C.$$

The LHS counts the number of *additive quadruples* (a_1, b_1, a_2, b_2) with $a_1 + b_1 = a_2 + b_2$. Thus having small $A + B$ implies the existence of many additive quadruples. Gowers's variant assumes the existence of many additive quadruples, and finds large subsets of A and B whose sumset is small.

The Balog-Szemerédi-Gowers Theorem A. *Let A and B be subsets of an abelian group, with $|A| = |B|$. Suppose there are at least $\alpha|A|^3$ additive quadruples $(a_1, b_1, a_2, b_2) \in A \times B \times A \times B$ with $a_1 + b_1 = a_2 + b_2$. Then there are subsets A' of A and B' of B with*

$$|A'| \ge \alpha^2|A|/(16\sqrt{2}), \qquad |B'| \ge \alpha^2|B|/16, \qquad \text{and} \qquad |A' + B'| \le 2^{28}\alpha^{-13}|A|.$$

The Balog-Szemerédi-Gowers Theorem B. *Let A and B be two subsets of an abelian group with $|A| = |B|$. Let G be a subgraph of the complete bipartite graph between A and B, with G having at least $|A||B|/K$ edges. Suppose that $A +_G B = \{a + b : \quad (a, b) \in G\}$ has cardinality $|A +_G B| \le K_1|A|$. Then there exist subsets A' of A and B' of B with*

$$|A'| \ge |A|/(4\sqrt{2}K), \qquad |B'| \ge |B|/(4K), \qquad \text{and} \qquad |A' + B'| \le 2^{15}K^5K_1^3|A|.$$

Equivalence of the two versions. Suppose we are given A and B with many additive quadruples. That is $\sum_n r_{A+B}(n)^2 \ge \alpha|A|^3$. Then there are at least $\alpha|A|/2$ popular sums n with $r_{A+B}(n) \ge \alpha|A|/2$ (why?). Define the graph G by letting (a, b) be an edge in G precisely when $a + b$ is a popular sum. The number of edges in G is at least $\alpha^2|A|^2/4$, and since there can be at most $2|A|/\alpha$ popular sums, we also have $|A +_G B| \le 2|A|/\alpha$. Therefore version A follows from version B.

Conversely, suppose we are given a subgraph G as in version B. Then $\sum_n r_{A+_G B}(n) = |G|$, and so $\sum_n r_{A+_G B}(n)^2 \ge |G|^2/|A +_G B|$ by Cauchy's inequality. Therefore there are at least $|G|^2/|A +_G B|$ additive quadruples, and we can deduce version B from version A (at least, up to constants).

It is version B that we will focus on proving. The idea is that since G contains a proportion of all edges from A to B, there will likely be many paths of length 2 from A to A, and many paths of length 3 from A to B. Once we quantify these paths of lengths two and three, the result follows easily. The situation is analogous to the sumsets $2A$, $3A$ etc becoming more and more smooth.

Lemma 1. *Let G be an undirected bipartite graph having two vertex sets A and B (that is, the edges connect points in A to points in B). Suppose that the edge set has cardinality $|A||B|/K$ for some $K \geq 1$. Given $\epsilon \in (0,1)$, there exists a subset A' of A with $|A'| \geq |A|/(\sqrt{2}K)$ such that for at least a proportion $(1-\epsilon)$ of the pairs $(a_1, a_2) \in A' \times A'$ we have at least $\epsilon|B|/(2K^2)$ paths of length 2 in G connecting a_1 and a_2.*

Proof. For $a \in A$ let $B(a)$ denote the points in B connected to a, and similarly for $b \in B$ let $A(b)$ denote the points in A connected to b. Let Ω denote the subset of $A \times A$ consisting of pairs (a_1, a_2) for which there exits fewer than $\epsilon|B|/(2K^2)$ elements in $B(a_1) \cap B(a_2)$.

Clearly $\sum_{b \in B} |A(b)|$ equals the total number of edges $|A||B|/K$. By Cauchy's inequality it then follows that

$$\sum_{b \in B} \sum_{a_1, a_2 \in A(b)} 1 = \sum_{b \in B} |A(b)|^2 \geq |A|^2|B|/K^2.$$

Further

$$\sum_{b \in B} \sum_{\substack{a_1, a_2 \in A(b) \\ (a_1, a_2) \in \Omega}} 1 = \sum_{(a_1, a_2) \in \Omega} \sum_{b \in B(a_1) \cap B(a_2)} 1 \leq |\Omega|\epsilon|B|/(2K^2) \leq \epsilon|A|^2|B|/(2K^2).$$

Combining the above two relations we find that

$$\sum_{b \in B} \Big(|A(b)|^2 - \frac{1}{\epsilon}(|A(b)^2 \cap \Omega|) \Big) \geq \frac{|A|^2|B|}{2K^2},$$

so that for some $b \in B$ one has

$$|A(b)|^2 - \frac{1}{\epsilon}(|A(b)^2 \cap \Omega|) \geq |A|^2/(2K^2).$$

The Lemma follows upon taking A' to be this set $A(b)$.

Lemma 2. *Let G be a bipartite graph as above, having an edge set of size $|A||B|/K$. We may extract a set A'' of A such that $|A''| \geq |A|/(4\sqrt{2}K)$, each vertex in A'' has degree at least $|B|/(2K)$, and for each $a_1 \in A''$ there exist at least $(1 - 1/(16K))|A''|$ vertices $a_2 \in A''$ such that a_1 and a_2 are joined by at least $|B|/(256K^3)$ paths of length 2.*

Proof. We remove from A all vertices with degree $\leq |B|/(2K)$. Let $\tilde{A}$ denote the set of remaining vertices, and consider the induced subgraph on vertex sets $\tilde{A}$ and B. Since at most $|A||B|/(2K)$ edges are removed from our original graph, our new graph has at least $|A||B|/(2K)$ edges. Furthermore, $|\tilde{A}| \geq |A|/(2K)$.

We take $\epsilon = 1/(32K)$ in Lemma 1, and thus find a subset A' of $\tilde{A}$ with $|A'| \geq |A|/(2\sqrt{2}K)$ (why?) such that for a proportion $1 - 1/(32K)$ of the pairs $(a_1, a_2) \in A' \times A'$ we have at least $|B|/(256K^3)$ paths of length 2 connecting a_1 and a_2. It follows that for at most half of values $a_1 \in A'$ can there exist more than $1/(16K)$ of values $a_2 \in A'$ with (a_1, a_2) not connected by many paths of length 2. Take A'' to be the good half of A'. This proves our Lemma.

Lemma 3. *Let G be a bipartite graph as above having an edge set of size $|A||B|/K$. We may find subsets A' and B' of A and B with $|A'| \geq |A|/(4\sqrt{2}K)$ and $|B'| \geq |B|/(4K)$ such that for any $a \in A'$ and $b \in B'$ there exist $\geq |A||B|/(2^{15}K^5)$ paths of length three joining a and b.*

Proof. Take A' to be the set A'' extracted in Lemma 2. We must now find the set B'. We will take B' to be the set of vertices adjacent to at least $|A'|/(8K)$ elements from A'. Note that the number of edges connecting A' to B is at least $|A'||B|/(2K)$. Therefore at least $|B|/4K$ of the vertices in B must be connected to $|A'|/(8K)$ vertices in A'; in other words, $|B'| \geq |B|/(4K)$. If $a \in A'$ and $b \in B'$ then we have at least $|A'|/(8K)$ vertices in A' that are adjacent to b, and at most $|A'|/(16K)$ of these can have the property that there are few paths of length two connecting them to a. Thus there are at least $|A'|/(16K)$ vertices a_2 that are both adjacent to b, and have at least $|B|/(256K^3)$ paths of length two connecting a and a_2. Thus there are at least $|A||B|/(2^{15}K^5)$ paths of length three connecting a and b.

Proof of Balog-Szemerédi-Gowers B. By Lemma 3 we may extract large sets A' and B' with at least $|A||B|/(2^{15}K^5)$ paths of length 3 connecting any two elements in these sets. Thus given a and b in A' and B', we can find more than $|A||B|/(2^{15}K^5)$ pairs b_1 and a_2 in B and A with (a, b_1), (a_2, b_1), (a_2, b) all being edges in our graph G. That is $a + b_1 = x$, $a_2 + b_1 = y$ and $a_2 + b = z$ are all elements of $A +_G B$. Now note the identity

$$a + b = a + b_1 - (b_1 + a_2) + a_2 + b = x - y + z.$$

We know lots of solutions to this equation with x, y and z in $A +_G B$. But the total number of choices for x, y and z is at most $|A +_G B|^3 \leq K_1^3|A|^3$. Therefore the number of distinct possibilities for $a + b$ is at most

$$\frac{K_1^3|A|^3}{|A|^2/(2^{15}K^5)} = 2^{15}K^5K_1^3|A|,$$

which completes our proof.

5. Linearity of the map ϕ on a big set.

We now resume the argument from §3. Let $\Gamma = \{(b, \phi(b)) : b \in B\}$. So Γ is a subset of $(\mathbb{Z}/N)^2$ possessing many additive quadruples. By version A of the Balog-Szemerédi-Gowers Theorem we know that there exist subsets Γ_1 and Γ_2 of Γ with

$$|\Gamma_1| = |\Gamma_2| = 2^{-30}\alpha^{128}|\Gamma|, \qquad \text{and} \qquad |\Gamma_1 + \Gamma_2| \leq 2^{200}\alpha^{-840}|\Gamma|.$$

Now we would like to use Freiman's Theorem. However, our version of Freiman's theorem was for subsets of $\mathbb{Z}$, and to apply it here we need to make a couple of two isomorphisms. Identify $(\mathbb{Z}/N)^2$ with $[1, N]^2$ and divide it into four squares. Then we may pick subsets $\Gamma_3 \subset \Gamma_1$ and $\Gamma_4 \subset \Gamma_2$ with each being a quarter of the size of these sets, each contained in a square. The sets Γ_3 and Γ_4 are naturally two isomorphic to subsets of $\mathbb{Z}^2$. By picking a large base, these sets Γ_3 and Γ_4 are seen to be two isomorphic to two subsets

of the integers Γ_5 and Γ_6. Now apply Freiman's Theorem. More precisely, we will use the Freiman variant given in VII.8.

In this way, we conclude that there exists a proper progression Q in $(\mathbb{Z}/N)^2$ whose dimension is $\ll \alpha^{-2^{16}}$ and size $\gg \exp(-\alpha^{-2^{17}})N$ such that

$$|\Gamma \cap Q| \gg \alpha^{2^{16}}|Q|.$$

Since Q is proper it contains a one dimensional progression of length $\gg \exp(-\alpha^{-2^{16}})N^{\alpha^{2^{16}}}$. Cover Q by translates of this progression. It follows that there exists a one dimensional progression P_0 (in $(\mathbb{Z}/N)^2$) having size $\gg \exp(-\alpha^{-2^{16}})N^{\alpha^{2^{16}}}$ and satisfying $|\Gamma \cap P_0| \gg \alpha^{2^{16}}|P_0|$.

Summarizing, we have shown the following Proposition.

Proposition. *Keep the notations of §3. There is a progression P in $\mathbb{Z}/N$ of size $\gg \exp(-\alpha^{-2^{16}})N^{\alpha^{2^{16}}}$, and a linear function $n \to 2\lambda n + \mu$ such that $|B \cap P| \geq \eta|P|$ (with $\eta \gg \alpha^{2^{16}}$), and for $k \in B \cap P$ we have $\phi(k) = 2\lambda k + \mu$.*

6. Extracting quadratic bias.

Our aim is now to show that f correlates locally with a quadratic phase function. We start with the following simple version.

Proposition 1. *Let $f : \mathbb{Z}/N \to [-1, 1]$, and suppose that for some $\lambda \in \mathbb{Z}/N$ we have*

$$\sum_{k \in \mathbb{Z}/N} |\widehat{\Delta(f,k)}(2\lambda k)|^2 \geq \zeta N^3.$$

Then for some $r \in \mathbb{Z}/N$ we have

$$\Big|\sum_n f(n)e\Big(\frac{\lambda n^2 + rn}{N}\Big)\Big| \geq \sqrt{\zeta}N.$$

Proof. Expanding out the hypothesis we obtain that

$$\begin{aligned}\zeta N^3 &\leq \sum_k \sum_{x,y} f(x)f(x+k)f(y)f(y+k)e(2\lambda k(y-x)) \\ &= \sum_k \sum_{x,u} f(x)f(x+k)f(x+u)f(x+u+k)e(2\lambda ku).\end{aligned}$$

Now observe that $x^2 - (x+k)^2 - (x+u)^2 + (x+k+u)^2 = 2ku$, and so the above equals

$$\sum_{x,k,u} f(x)e(\lambda x^2)f(x+k)e(-\lambda(x+k)^2)f(x+u)e(-\lambda(x+u)^2)f(x+k+u)e(\lambda(x+k+u)^2).$$

If we set $g(x) = f(x)e(\lambda x^2)$, then the above equals $N^3\|g\|_{U^2}^4$. Therefore

$$\zeta N^3 \leq N^3\|g\|_{U^2}^4 = \frac{1}{N}\sum_r |\hat{g}(r)|^4 \leq N \max_r |\hat{g}(r)|^2$$

using Parseval. This yields the Proposition.

With a little more effort, the same argument extends to cover the information given in the Proposition of §5.

Proposition 2. *Keep the notations of §5. For every $x \in \mathbb{Z}/N$ there exists $r_x \in \mathbb{Z}/N$ such that*

$$\sum_x \Big| \sum_{k\in P+x} f(k) e\Big(-\frac{\lambda k^2 + r_x k}{N}\Big)\Big| \ge \frac{\eta\alpha^8}{4\sqrt{2}} N|P|.$$

Proof. In the notation of §5 we have that

$$\sum_{k\in P} |\widehat{\Delta(f,k)}(2\lambda k+\mu)|^2 \ge \frac{\eta\alpha^8}{4}|P|N^2.$$

Expanding out the LHS we obtain

$$\sum_{k\in P}\sum_{x,u} f(x)f(x+k)f(x+u)f(x+u+k)e\Big(\frac{(2\lambda k+\mu)u}{N}\Big).$$

Now each u may be written as $\ell + y$ where $\ell \in P$ and $y \in \mathbb{Z}/N$ in exactly $|P|$ ways. Therefore the above is

$$\frac{1}{|P|}\sum_{x,y}\sum_{k,\ell\in P} f(x)f(x+k)f(x+y+\ell)f(x+y+k+\ell)e\Big(\frac{(2\lambda k+\mu)(\ell+y)}{N}\Big).$$

For some $y \in \mathbb{Z}/N$ it follows that

$$(1)\quad \sum_x \Big|\sum_{k,\ell\in P} f(x+k)f(x+y+\ell)f(x+y+k+\ell)e\Big(\frac{2\lambda k\ell + 2\lambda k y + \mu\ell}{N}\Big)\Big| \ge \frac{\eta\alpha^8}{4}|P|^2 N.$$

Now we separate the variables k and ℓ by writing $2\lambda k\ell = \lambda((k+\ell)^2 - k^2 - \ell^2)$. We think of x as being fixed, and focus on the sums over k and ℓ. We write $g_1(k) = g_{1,x}(k) = f(x+k)e((-\lambda k^2 + 2\lambda k y)/N)$, $g_2(\ell) = f(x+y+\ell)e((-\lambda\ell^2+\mu\ell)/N)$, and $g_3(k+\ell) = f(x+y+k+\ell)e(\lambda(k+\ell)^2/N)$ provided k and ℓ are in P and $k+\ell$ is in $P+P$. For other values of k, ℓ or $k+\ell$ we set these functions equal to zero. Thus the inner sums in (1) give, by Parseval

$$\sum_{k,\ell} g_1(k)g_2(\ell)g_3(k+\ell) = \frac{1}{N}\sum_r \hat{g}_1(r)\hat{g}_2(r)\hat{g}_3(-r).$$

By Cauchy-Schwarz and Parseval, this is

$$\le (\max_r |\hat{g}_1(r)|)\Big(\frac{1}{N}\sum_r |\hat{g}_2(r)|^2\Big)^{\frac12}\Big(\frac{1}{N}\sum_r |\hat{g}_3(r)|^2\Big)^{\frac12} \le (\max_r |\hat{g}_1(r)|)|P|^{\frac12}(2|P|)^{\frac12}.$$

Using this in (1) we have shown that

$$\sum_x \max_r \Big|\sum_{k\in P} f(x+k)e\Big(\frac{-\lambda k^2 + 2\lambda k y - rk}{N}\Big)\Big| \ge \frac{\eta\alpha^8}{4\sqrt{2}}|P|N.$$

That is, for every x there exists r_x such that

$$\sum_x \Big|\sum_{k\in P+x} f(k)e\Big(-\frac{\lambda k^2 + r_x k}{N}\Big)\Big| \ge \frac{\eta\alpha^8}{4\sqrt{2}}|P|N,$$

which proves our Proposition.

7. Application of Weyl's inequality leading to density increment.
Let us recall the proof of Roth's theorem in the case when there is a large Fourier coefficient of f. Say $\hat{f}(r)$ is large. In this case, the idea was to dissect the interval $[1, N]$ into sub-progressions on which $e(rn/N)$ is roughly constant. This is achieved by using Dirichlet's theorem to approximate r/N by a/q for some small q, then splitting n into progressions $\pmod q$, and subdividing those progressions into small intervals.

Lemma 1. *Let P be a progression $\pmod N$ of length R, and let $\psi_1(x) = \alpha x$ be a linear function. There exists a partition of P into $\pmod N$ progressions $P_1, \dots, P_M$ each of length about $R^{\frac{1}{4}}$ (so M is about $R^{\frac{3}{4}}$), such that for $x, y \in P_j$ we have*

$$|e(\psi_1(x)/N) - e(\psi_1(y)/N)| \ll R^{-\frac{1}{4}}.$$

Proof. Without loss of generality we may suppose that $P = [1, R]$. By Dirichlet's Theorem we may find $q \leq \sqrt{R}$ such that $\|q\alpha/N\| \leq 1/\sqrt{R}$. Divide $[1, R]$ into the progressions $\pmod q$. If one of those progressions is $a + jq$ with $1 \leq j \leq R/q$, then divide that progression into sub-intervals for j each of length about $R^{\frac{1}{4}}$. In this way we arrive at the $R^{\frac{3}{4}}$ desired sub-progressions, and for any x, y in one of these sub-progressions we plainly have $\|(\psi_1(x) - \psi_1(y))/N\| \leq R^{\frac{1}{4}}\|q\alpha/N\| \leq R^{-\frac{1}{4}}$. This proves the Lemma.

We now extend the argument above to the case when f correlates locally with some quadratic polynomial, as in §6 above. The aim is to dissect that progression $P + x$ into sub-progressions on which the quadratic phase is roughly constant.

Lemma 2. *Let P be a progression $\pmod N$ of size R, and let $\psi_2(x) = \alpha x^2 + \beta x$ be some quadratic function. We may split P into $\pmod N$ progressions $P_1, \dots, P_M$ each of length about $R^{\frac{1}{128}}$ (so that M is about $R^{\frac{127}{128}}$), such that for $x, y \in P_j$ we have*

$$|e(\psi_2(x)/N) - e(\psi_2(y)/N)| \ll R^{-\frac{1}{128}}.$$

Proof. Without loss of generality we may suppose that $P = [1, R]$. By Weyl's Theorem (see Corollary 3 of Chapter IV) we may find $q \leq R^{\frac{1}{2}}$ such that $\|q^2\alpha/N\| \leq R^{-\frac{1}{8}}$. Divide P into the progressions $\pmod q$, and divide each of those progressions into intervals of length $R^{\frac{1}{32}}$. In this manner we obtain about $R^{\frac{31}{32}}$ sub-progressions each of length about $R^{\frac{1}{32}}$.

Consider one of the above sub-progressions. It looks like $a + jq$ for some a, and $1 \leq j \leq R^{\frac{1}{32}}$. Now

$$\frac{\psi_2(a+jq)}{N} = \frac{\alpha}{N}(a+jq)^2 + \frac{\beta}{N}(a+jq) = \frac{\alpha a^2 + \beta a}{N} + j^2\frac{q^2\alpha}{N} + j\frac{2\alpha aq + \beta q}{N}.$$

Here the first term is constant and doesn't vary on this sub-progression. The second term viewed $\pmod 1$ changes by at most $R^{\frac{1}{16}}\|q^2\alpha/N\| \leq R^{-\frac{1}{16}}$. The last term gives a linear polynomial in j, and we may employ Lemma 1 to make that term locally constant. That is, using Lemma 1 we may divide our values for j into sub-progressions of length $R^{\frac{1}{128}}$ on which the last term varies $\pmod 1$ by at most $R^{-\frac{1}{128}}$. This establishes our Lemma.

The progressions above are in $\mathbb{Z}/N$. We now show that progressions in $\mathbb{Z}/N$ can be broken up into not too many genuine (over $\mathbb{Z}$) progressions.

Lemma 3. *Let P be a $\mathbb{Z}/N$ progression of length R. Then we may partition P into $4\sqrt{R}$ genuine arithmetic progressions.*

Proof. Let the progression P be $a + jq$ with $1 \le j \le R$. We find, by Dirichlet's Theorem $\ell \le \sqrt{R}$ with $\|\ell q/N\| \le R^{-\frac{1}{2}}$. Use this to divide our progression P into sub-progressions of $j \pmod{\ell}$. Sub-divide those progressions into intervals, as needed. The whole argument is quite similar to Lemma 1.

Proposition 4. *Let P be a progression as in §6, and denote its size by R. For each $x \in \mathbb{Z}/N$ we may partition $P + x$ into about $4R^{\frac{255}{256}}$ genuine arithmetic progressions $P_{x,1}$, ..., $P_{x,M}$ (with M about $4R^{\frac{255}{256}}$) such that*

$$\sum_{x\in\mathbb{Z}/N}\sum_{j=1}^{M}\Big|\sum_{k\in P_{x,M}} f(k)\Big| \ge \frac{\eta\alpha^8}{8}NR.$$

Proof. We start with the estimate of Proposition 2 of §6. Then use Lemma 2 to split $P+x$ into about $R^{\frac{127}{128}}$ progressions $\pmod N$ on which the exponential factor there is essentially constant. Then subdivide these progressions using Lemma 3 to obtain genuine arithmetic progressions. Thus we obtain the Proposition.

Corollary 5 (Density Increment). *Keep the notations of Proposition 4. There is a genuine arithmetic progression Q of size at least $(\eta\alpha^8/128)R^{\frac{1}{256}}$ such that*

$$|A \cap Q| \ge |Q|(\delta + \eta\alpha^8/128).$$

Proof. Note that

$$\sum_{x\in\mathbb{Z}/N}\sum_{j=1}^{M}\sum_{k\in P_{x,M}} f(k) = 0.$$

Therefore, adding this to the estimate of Proposition 4 we obtain that

$$\sum_{x}\sum_{j=1}^{M}\max\Big(0, \sum_{k\in P_{x,M}} f(k)\Big) \ge \frac{\eta\alpha^8}{16}NR.$$

The contribution of terms with $|P_{x,M}| \le (\eta\alpha^8/128)R^{\frac{1}{256}}$ to the above is $\le (\eta\alpha^8/32)NR$. Hence the contribution of long intervals to the above sum is big, and the Corollary follows.

8. Gowers's Theorem for four term progressions.

We recapitulate the argument. We have a set $A \in \mathbb{Z}_N$ with $|A| = \delta N$, and f is the balanced function of A. We may suppose that $\|f\|_{U^3} \ge \delta^4/144$, else we have density increment on a subprogression of length $2N/5$ (see §2). This implies the existence of large Fourier coefficients $\widehat{\Delta}(f,k)(\phi(k))$ for k lying in a big set B. The map ϕ was seen to be weakly linear (possessing many additive quadruples) in §3. This weak linearity led, via Balog-Szemerédi-Gowers to linearity on a largish progression P (see §4 and 5). From the

linearity of the function ϕ we then obtained long progressions on which f had quadratic bias (Proposition 2 of §6). Finally in §7 we dissected those long progressions into sub-progressions where the quadratic phase function was roughly constant. This led finally to the increase in density of A on a sub-progression.

Precisely, the above argument leads to a (genuine) sub-progression of $[1, N]$ of size at least $\exp(-\delta^{-2^{20}})N^{\delta^{2^{17}}}$ on which the relative density of A is at least $\delta + 2^{-2^{20}}\delta^{2^{20}}$. Now we iterate this argument. The argument can be iterated at most $2^{2^{20}}\delta^{-2^{20}}$ times. After so many iterations, we obtain a very large set of numbers below $\exp(-2\delta^{-2^{20}})N^{\delta^{2^{40}}\delta^{-2^{20}}}$. If this is large, then we obtain the desired four term progressions. This is satisfied if $\delta \geq (\log\log N)^{-2^{-40}}$. Equivalently, if

$$N \geq \exp(\exp((1/\delta)^{2^{40}})),$$

then a set of density δ in $[1, N]$ contains a four term progression.