Roth's Theorem: Overview

Roth's Theorem

Let $\delta > 0$, and let $N = N(\delta) = e^{e^{1000/\delta}}$. Then a set $A \subset [N]$ with $|A| \geq \delta N$ contains a 3-term AP.

I Fourier Analysis on $\mathbb{Z}_N$

Let N be a prime, let $\omega = e^{2\pi i/N}$, and let $f : \mathbb{Z}_N \to \mathbb{C}$. The Fourier transform $\hat{f}$ of f is defined by $\hat{f}(r) = \sum_s f(s)\omega^{-rs}$. If also $g : \mathbb{Z}_N \to \mathbb{C}$, then the convolution $f * g$ is defined by the formula $(f * g)(s) = \sum_t f(t)\overline{g(t-s)}$. We have the following identities:

$$
\begin{aligned}
(f * g)^\wedge(r) &= \hat{f}(r)\overline{\hat{g}(r)} \\
\sum_r \hat{f}(r)\overline{\hat{g}(r)} &= N \sum_s f(s)\overline{g(s)} \\
\sum_r |\hat{f}(r)|^2 &= N \sum_s |f(s)|^2 \\
\sum_r \hat{f}(r)\omega^{rs} &= N f(s)
\end{aligned}
$$

II The fundamental formula and its consequences

Our strategy will be to show that either A contains plenty of 3-term APs, or there exists a (long) subprogression $P \subset [N]$ such that $|A \cap P| \geq (\delta + c\delta^2)|P|$. We will call such a P a high density subprogression, or **hdsp**. Repeated use of this fact will give the result.

Suppose that $A, B, C \subset \mathbb{Z}_N$. Then the number of $(s, t, u) \in A \times B \times C$ with $s + u = 2t$ is given by

$$
\begin{aligned}
N^{-1} \sum_r \hat{A}(r)\hat{B}(-2r)\hat{C}(r) &= N^{-1}|A||B||C| + N^{-1} \sum_{r \neq 0} \hat{A}(r)\hat{B}(-2r)\hat{C}(r) \\
&\geq N^{-1}|A||B||C| - \max_{r \neq 0} |\hat{A}(r)||B|^{1/2}|C|^{1/2}
\end{aligned}
$$

Applying this with $B = C = A \cap (N/3, 2N/3)$ gives that *either* A contains at least $\delta^3 N^2/50 - N$ 3-term APs, *or* $|B| \leq \delta N/5$ (whence $A \setminus B$ lies in a **hdsp**), *or* there is some $r \neq 0$ such that $|\hat{A}(r)| \geq \delta^2 N/10$. We must show that this third possibility also enables us to find a **hdsp**.

III A high density subprogression mod N

At this stage, the proof is morally over – a large Fourier coefficient should guarantee some sort of "periodicity" in A, which should easily yield a **hdsp**. However, we will have to unravel it. We divide the unit circle into M equal arcs $I_1, \ldots, I_M$, where $M \approx 40\pi\delta^{-2}$. Each arc contains about $N/M = \delta^2 N/40\pi$ consecutive powers of ω. For $1 \le j \le M$, set

$$P_j = \{s \in \mathbb{Z}_N : \omega^{-rs} \in I_j\}$$

and pick $s_j \in P_j$. Now

$$\hat{A}(r) = \sum_s A(s)\omega^{-rs} = \sum_{j=1}^{M} \sum_{s \in P_j} A(s)\omega^{-rs} \approx \sum_{j=1}^{M} \sum_{s \in P_j} A(s)\omega^{-rs_j} = \sum_{j=1}^{M} |A \cap P_j| \omega^{-rs_j}$$

Since the numbers ω^{-rs_j} are spread almost evenly around the circle, $\sum_{j=1}^{M} \omega^{-rs_j}$ is very small. Consequently, the sizes of the intersections $|A \cap P_j|$ must differ from each other by at least a certain amount, and indeed the above reasoning can be made to show that, for some j,

$$|A \cap P_j| \ge (\delta + \delta^2/40)|P_j| = \delta' |P_j|$$

IV A high density subprogression

We are almost done, except that P_j is only an arithmetic progression (of common difference $-r^{-1}$) mod N, rather than a genuine AP. In fact, three issues present themselves:

- P_j might "overlap" several times, e.g. we could have $N = 101$ and $P_j = \{30, 60, 90, 19, 49, 79, 8, 38, 68, 98, 27\}$
- P_j might not overlap, but it might "pass through" 0, e.g. we could have $N = 11$ and $P_j = \{6, 8, 10, 1, 3\}$
- N might not be prime, e.g. 10 isn't prime

Roughly speaking, these are dealt with as follows. If $P_j = \{s_0, s_1, \ldots, s_{l-1}\}$ has length l, we look at the first $m \approx \sqrt{l}$ terms. Two of these must be within N/m of each other, say s_a and s_b, with $b > a$. But then s_{b-a} is within N/m of s_0. Writing $u = b - a$, we consider the sequences $Q_j^0 = \{s_0, s_u, s_{2u}, \ldots\}, Q_j^1 = \{s_1, s_{u+1}, s_{2u+1}, \ldots\}, Q_j^2 = \{s_2, s_{u+2}, s_{2u+2}, \ldots\}, \ldots$. These are still mod N progressions, but since the common difference in each one is less than N/m is magnitude, we can split each Q_j^i into genuine APs (call them R_j^k), all but two of which have length at least $\sqrt{l}$.

The argument now proceeds as follows. If the density of A in P_j is at least δ' (as above), then the density of A must be at least δ' in one of the subprogressions R_j^k. But what if the specific high density R_j^k turns out to be one of the "ends" of a Q_j^i, which potentially is very short? This is really to do with the second issue above, which we handle using the following lemma. Suppose P_1 and P_2 are disjoint APs (e.g. $P_1 = \{6, 8, 10\}$ and $P_2 = \{1, 3\}$). Suppose also that $|A \cap (P_1 \cup P_2)| \ge (\delta + \delta^2/40)|P_1 \cup P_2|$. Then either both P_1 and P_2 have length at least $\delta^2/80|P_1 \cup P_2|$, or at most one of them, say P_1, has length less than $\delta^2/80|P_1 \cup P_2|$, and then $|A \cap P_2| \ge (\delta + \delta^2/80)|P_2|$. The proof of this lemma is just a simple calculation.

Finally, the third issue is easily resolved using *Bertrand's Postulate*, proved by Chebyshev: for any $n \ge 1$, there is always a prime p satisfying $n \le p \le 2n$. This has an elementary proof.

AS November 2009