Larry Guth

Larry Guth spoke on applications of Fourier analysis to the problem of counting solutions to Diophantine equations. His initial focus was on sums of cubes: how many solutions are there to the equation

$$a_1^3 + a_2^3 + a_3^3 = b_1^3 + b_2^3 + b_3^3 \tag{1}$$

with $a_i, b_i \in \mathbb{Z}$ and $1 \le a_i, b_i \le N$? A heuristic argument suggests that the number of solutions, $S(N)$, satisfies

Conjecture 1. *For all $\varepsilon > 0$, $S(N) \le C_\varepsilon N^{3+\varepsilon}$ for some constant C_ε.*

The proof remains out of reach, but although the Fourier approach has not led to progress in the three-cubes problem, it does suggest a way to tackle other Diophantine problems of the same type.

In the three-cubes case, the starting point is the introduction of the trigonometric polynomial

$$f(X) = \sum_{a=1}^{N} \mathrm{e}^{2\pi \mathrm{i} a^3 X} .$$

By evaluating the integral on the right-hand side, we have the following.

Lemma 1. $S(N) = \int_0^1 |f(X)|^6 \, \mathrm{d}X$.

Starting from this formula for $S(N)$, the idea is to use estimates from Fourier analysis to obtain bounds on the number of solutions. For example, from the triangle inequality, $|f(X)| \le N$ while from orthogonality,

$$\int_0^1 |f|^2 = N .$$

By combining these, one has

Proposition 1. $S(N) \le N^5$

This is not much of an advance because the bound can be obtained directly without using Fourier analysis, but it does suggest that one might do better than the trivial bound by exploring the properties of f in more depth.

Suppose that we are given a function g with $|g(X)| \le N \; \forall X$ and $\int_0^1 |g|^2 = N$. There are two extreme behaviours. If $|g(X)|$ is 'spread out', with values close to $\sqrt{N}$ throughout the unit interval, then

$$\int_0^1 |g|^6 \sim N^3 .$$

If f had this behaviour, then the conjecture would follow. On the other hand, if $|g(X)|$ is 'focused', taking low values except on a set of total length $1/N$ on which it takes values close to N, then

$$\int_0^1 |g|^6 \sim N^5 .$$

In the Fourier approach, progress towards establishing the conjecture would involve showing that f exhibits behaviour closer to the first case than to the second.

This is hard, but there is another case in which the approach has recently proved fruitful, namely that of Vinogradov's conjecture concerning the number $J_{s,k}(N)$ of solutions to the system of Diophantine equations

$$\begin{aligned} a_1 + \cdots + a_s &= b_1 \cdots + b_s \\ a_1^2 + \cdots + a_s^2 &= b_1^2 \cdots + b_s^2 \\ &\vdots \\ a_1^k + \cdots + a_s^k &= b_1^k \cdots + b_s^k \end{aligned} \qquad (2)$$

with $a_i, b_i \in \mathbb{Z}$ and $1 \le a_i, b_i \le N$.

Conjecture 2. $J_{s,k}(N) \le C(s,k,\varepsilon)N^\varepsilon(N^s + N^{2s-(1+\cdots+k)})$.

On the right, N^s counts the number of diagonal solutions ($a_i = b_i$) while the second term comes from a heuristic guess for the number of non-diagonal solutions.

Vinogradov proved the conjecture for $s \ge Ck^2 \log k$. Its truth generally has been established by Bourgain, Demeter, and Guth.[1] Trevor Wooley had earlier proved the conjecture for $k = 3$, and more recently has given a different proof for the general case, by using his method of efficient congruencing.

At first sight it is puzzling that the second conjecture should be more tractable than the first when if anything it looks more complex. The answer lies in the additional symmetry in the second case: both equations (1) and (2) are invariant under dilation of the variables; but the system (2) has additional symmetry under translation. This is exploited more or less explicitly in all the approaches to Vinogradov's conjecture.

Guth outlined the idea behind his proof (with Bourgain and Demeter), starting with the special case $k = 2$. Here there is an older and simpler proof using unique factorization in $\mathbb{Z}[\mathrm{i}]$, but it is instructive to take a different route that uses Fourier analysis because the method then generalizes to give a proof of the full conjecture. When $k = 2$, the key object is the two-variable function

$$f(X) = \sum_{a=1}^{N} \mathrm{e}^{2\pi \mathrm{i}\omega_a . X}$$

where the 'frequencies' ω_a all lie on the parabola $\{(a, a^2)\}$ in the plane. They are defined in such a way that

$$J_{s,k}(N) = \int_{[0,1]^2} |f(X)|^{2s}\, \mathrm{d}X.$$

As with the first conjecture, the problem is to show that $|f(X)|$ is 'spread-out' rather than 'focused'. This is achieved by a multiscale approach. First the parabola is partitioned into the union of small arcs θ_j. Correspondingly, the sum defining f decomposes into sums of the functions

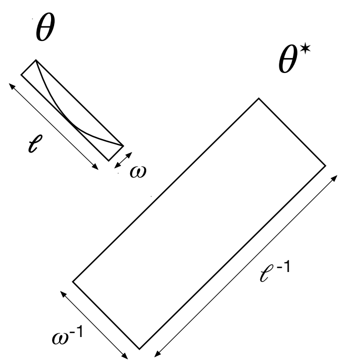

$$f_{\theta_j} = \sum_{\omega_a \in \theta_j} \mathrm{e}^{2\pi \mathrm{i}\omega_a . X}.$$

Each arc determines a rectangle θ^* in the X-plane, 'dual' to the smallest rectangle θ containing the arc itself. The dual rectangle is orthogonal to θ, with inverse width and length. Its translates tile the X-plane. A key lemma is that $|f_\theta|$ is roughly constant on each tile.

[1] Ann. Math. **184** (2016), 633–682.

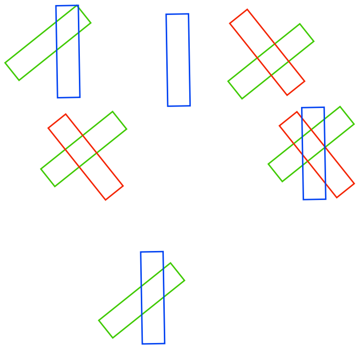

The idea of the proof is to show that f must be 'spread out'. The individual functions f_θ could be either spread out or concentrated. If each is spread out, then f must be as well. On the other hand, suppose that each f_θ is focused in a sparse set of translates of θ. Because the tilings by the different θs (shown on the left in different colors) are slanted in different directions, one can argue from the lemma that the places where the different $|f_\theta|$s are large cannot coincide, and so f is forced to be more spread out than the individual f_θs. This idea goes back to the 1970s, under the heading *restriction theory*, but it was not clear that, on its own, it would lead to estimates sharp enough to prove the conjecture, not least because in higher dimensions one comes up against Kakeya-type questions about the extent to which tubes pointing in different directions can overlap. The new advance comes from combining restriction theory with a multi-scale approach, in which one considers nested sequences of partitions, with the small arcs θ_i grouped into larger arcs τ_i. This gives much more detailed constraints on the extent to which it is possible for f to be focused.

Guth concluded his lecture by working through one particular example that illustrated how the information at different scales combines to tie down the behaviour of f itself in enough detail to prove the conjecture.