HARMONIC ANALYSIS: LECTURE EIGHT

ASSOCIATE PROFESSOR PO LAM YUNG, TRANSCRIBED BY DOMINIC CONNORS
(U7289061)

Towards a Proof of the Local Smoothing Conjecture

In the last lecture, we surveyed the strategies used in modern attempts to prove the local smoothing conjecture: the reverse square function estimate strategy of Mockenhaupt, Seeger and Sogge (1992); the use of decoupling inequalities, which was pioneered by Wolff (2000) with subsequent contributions from Laba, Pramanik, Seeger, Bourgain and Demeter; and, the use of wave envelopes by Guth, Wang and Zhang (2020) with subsequent contributions from Gan and Wu (2025).

Local Smoothing Conjecture (Sogge, 1991). Suppose that u is a solution to the wave equation and $\hat{f}$ is supported in an annulus of size R. Then, for any $\epsilon > 0$

$$\|u\|_{L^p(\mathbb{R}^n \times [1,2])} \lesssim_\epsilon R^\epsilon \|f\|_{L^p(\mathbb{R}^n)} \tag{1}$$

We began an investigation of Wolff's decoupling strategy which we will continue in this lecture.

Decoupling

Our aim is to understand how we might find estimates for a solution to the half wave equation u on $\mathbb{R}^n_x \times \mathbb{R}_t$ where

$$u(x,t) = e^{it\sqrt{-\Delta}} f(x) = \sum_{\theta \in \Pi} u_\theta(x,t) \qquad \text{where} \qquad u_\theta(x,t) = e^{it\sqrt{-\Delta}} f_\theta(x) \tag{2}$$

If the Fourier transform of f is supported on an annulus of internal and external radius R and $2R$, respectively i.e. $\operatorname{Supp} \widehat{f} \subset \{|\xi| \sim R\}$. Let Π be a finite collection of $R \times \sqrt{R} \times \cdots \times \sqrt{R}$ rectangles θ that efficiently cover the annulus. Each θ creates a spherical cap of area roughly $R^{\frac{n-1}{2}}$ and the surface area of the outer sphere is of order R^{n-1} so $|\Pi| \approx R^{n-1}/R^{\frac{n-1}{2}} = R^{\frac{n-1}{2}}$. We decompose f so that $f = \sum_{\theta \in \Pi} f_\theta$ with $\operatorname{Supp} \widehat{f_\theta} \subset \theta \cap \operatorname{Supp} \widehat{f}$.

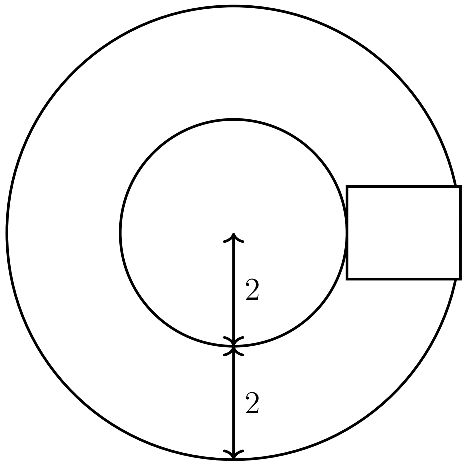

FIGURE 1. The annulus $\{|\xi| \sim 2\}$ in $\mathbb{R}^2$ and a θ in its covering Π.

In the last lecture, we considered the following result of Bourgain and Demeter which was published in *The Annals of Mathematics* in 2015. It had been conjectured by Wolff. Recall that $A \lessapprox B$ means that for every $\varepsilon > 0$ there is some $c_\varepsilon > 0$ such that $A \leq c_\varepsilon R^\varepsilon B$.

Theorem One (Bourgain and Demeter, 2015). If $p \geq 2$, $R \geq 1$, $u = \sum_{\theta \in \Pi} u_\theta$ as above and

$$\alpha(p) = \begin{cases} \frac{n-1}{2}\left(\frac{1}{2} - \frac{1}{p}\right) & \text{if } 2 \leq p \leq \frac{2(n+1)}{n-1} \\ (n-1)\left(\frac{1}{2} - \frac{1}{p}\right) - \frac{1}{p} & \text{if } p \geq \frac{2(n+1)}{n-1} \end{cases} \tag{3}$$

Then,

$$\|u\|_{L^p(\mathbb{R}^n \times [1,2])} \lessapprox R^{\alpha(p)} \left(\sum_{\theta \in \Pi} \|u_\theta\|^p_{L^p(\mathbb{R}^n \times [1,2])} \right)^{\frac{1}{p}} \tag{4}$$

Theorem one is implied by the following stronger result, theorem two, via an application of the Hölder inequality, as we shall see below. Bourgain and Demeter proved theorem one by first proving theorem two. It was Wolff who suggested that we look at inequalities of the form we see in the theorem below.

Theorem Two (Bourgain and Demeter). If $p \geq 2$, $R \geq 1$, $u = \sum_{\theta \in \Pi} u_\theta$ as above and

$$d(p) = \begin{cases} 0 & \text{if } 2 \leq p \leq \frac{2(n+1)}{n-1} \\ \frac{n-1}{2}\left(\frac{1}{2} - \frac{1}{p}\right) - \frac{1}{p} & \text{if } p \geq \frac{2(n+1)}{n-1} \end{cases} \tag{5}$$

Then,

$$\|u\|_{L^p(\mathbb{R}^n\times[1,2])} \lessapprox R^{d(p)} \left(\sum_{\theta\in\Pi} \|u_\theta\|^2_{L^p(\mathbb{R}^n\times[1,2])} \right)^{\frac{1}{2}} \tag{6}$$

Proving that theorem one follows from theorem two is straightforward. Let the sequence $(a_\theta)_{\theta\in\Pi}$ be given by $a_\theta = \|u_\theta\|_{L^p(\mathbb{R}^n\times[1,2])} \geq 0$ and suppose $\frac{1}{2} = \frac{1}{r} + \frac{1}{p}$. Using the Hölder inequality for $\ell^p(\Pi)$ spaces, we find that the bound from theorem two is controlled by the bound in theorem one:

$$\begin{aligned} R^{d(p)} \left(\sum_{\theta\in\Pi} \|u_\theta\|^2_{L^p(\mathbb{R}^n\times[1,2])} \right)^{\frac{1}{2}} &= R^{d(p)} \left(\sum_{\theta\in\Pi} a_\theta^2 \right)^{\frac{1}{2}} \\ &\leq R^{d(p)} \left(\sum_{\theta\in\Pi} 1^r \right)^{\frac{1}{r}} \left(\sum_{\theta\in\Pi} a_\theta^p \right)^{\frac{1}{p}} \text{ [by the } \ell^p \text{ Hölder inequality]} \\ &= R^{d(p)} |\Pi|^{\frac{1}{2}-\frac{1}{p}} \left(\sum_{\theta\in\Pi} a_\theta^p \right)^{\frac{1}{p}} \\ &\approx R^{d(p)} (R^{\frac{n-1}{2}})^{\frac{1}{2}-\frac{1}{p}} \left(\sum_{\theta\in\Pi} \|u_\theta\|^p_{L^p(\mathbb{R}^n\times[1,2])} \right)^{\frac{1}{p}} \\ &= R^{\alpha(p)} \left(\sum_{\theta\in\Pi} \|u_\theta\|^p_{L^p(\mathbb{R}^n\times[1,2])} \right)^{\frac{1}{p}} \text{ [by the definition of } \alpha \text{ and } d \text{]} \end{aligned}$$

The Tomas-Stein exponent $\frac{2(n+1)}{n-1}$ is derived by solving $d(p) = \frac{n-1}{2}\left(\frac{1}{2} - \frac{1}{p}\right) - \frac{1}{p} = 0$ for p.

As we shall see, Bourgain and Demeter proved theorem two using induction on scale. They opted to prove the stronger theorem two instead of theorem one directly because proofs by induction lend themselves to proving stronger statements as you then have a stronger induction hypothesis to work with. Theorem three will be more general than theorem two and not explicitly related to the wave equation, although the setup will be derived from it; we really only care about where the Fourier transform is mapped to.

Before we do this, we note that $\left\|f(x,t)\mathbb{1}_{[1,2]}(t)\right\|_{L^p(\mathbb{R}^n_x\times\mathbb{R}_t)} = \|f(x,t)\|_{L^p(\mathbb{R}^n\times[1,2])}$ by the properties of indicator functions on product spaces. So, we can reformulate theorem two.

(Equivalent formulation of) Theorem Two. If $p \geq 2$, $R \geq 1$, $u = \sum_{\theta \in \Pi} u_\theta$ as above and

$$d(p) = \begin{cases} 0 & \text{if } 2 \leq p \leq \frac{2(n+1)}{n-1} \\ \frac{n-1}{2}\left(\frac{1}{2} - \frac{1}{p}\right) - \frac{1}{p} & \text{if } p \geq \frac{2(n+1)}{n-1} \end{cases} \tag{7}$$

Then,

$$\left\| u(x,t)\mathbb{1}_{[1,2]}(t) \right\|_{L^p(\mathbb{R}^n_x \times \mathbb{R}_t)} \lessapprox R^{d(p)} \left(\sum_\theta \left\| u_\theta(x,t)\mathbb{1}_{[1,2]}(t) \right\|^2_{L^p(\mathbb{R}^n_x \times \mathbb{R}_t)} \right)^{\frac{1}{2}} \tag{8}$$

Waving Goodbye to the Wave Equation

We will now exploit this subtle reformulation to understand the Fourier support of the controlled term in theorem two, which will inspire theorem three. As we mentioned before, theorem three yields theorem two despite not explicitly referencing the wave equation or its solutions at all. Our solution to the half wave equation $u : \mathbb{R}^n_x \times \mathbb{R}_t \to \mathbb{C}$ is given by $u(x,t) = e^{it\sqrt{-\Delta}} f(x) = \int_{\mathbb{R}^n} e^{2\pi i (x\cdot\xi + t|\xi|)} \widehat{f}(\xi) d\xi$. The solution u is defined on $n+1$ dimensional spacetime but for each point in time t, $u(\cdot, t) : \mathbb{R}^n_x \to \mathbb{C}$ is defined on n dimensional space. For the sake of clarity, we will use the notation

$$\mathcal{F}_{x,t} v(\xi, \tau) = \int_{\mathbb{R}_t} \int_{\mathbb{R}^n_x} v(x,t) e^{-2\pi i (x\cdot\xi + t\tau)} dx dt \tag{9}$$

$$\mathcal{F}_x v(\xi, t) = \int_{\mathbb{R}^n_x} v(x,t) e^{-2\pi i x\cdot\xi} dx \tag{10}$$

to denote the $\mathbb{R}^n_x \times \mathbb{R}_t$ spacetime Fourier transform of $v : \mathbb{R}^n_x \times \mathbb{R}_t \to \mathbb{R}$ and the $\mathbb{R}^n_x$ spatial Fourier transform of $v(\cdot, t) : \mathbb{R}^n_x \to \mathbb{C}$, respectively. Note that $\mathcal{F}_{x,t} = \mathcal{F}_x \circ \mathcal{F}_t = \mathcal{F}_t \circ \mathcal{F}_x$. Using this notation, we see that our solution $u(x,t) = \mathcal{F}_x^{-1}\left(e^{2\pi i t|\xi|} \mathcal{F}_x f\right)$ and its spatial Fourier transform $\mathcal{F}_x u(\xi, t) = e^{2\pi i t|\xi|} \mathcal{F}_x f$. Although we will not prove it here, it can be easily verified that

$$\partial_t u(x,t) = \int_{\mathbb{R}^n} 2\pi i |\xi| e^{2\pi i (x\cdot\xi + t|\xi|)} \mathcal{F}_x f(\xi) d\xi = \mathcal{F}_x^{-1}(2\pi i |\xi| \mathcal{F}_x u) \tag{11}$$

$$\sqrt{-\Delta_x} u(x,t) = \int_{\mathbb{R}^n} 2\pi |\xi| e^{2\pi i (x\cdot\xi + t|\xi|)} \mathcal{F}_x f(\xi) d\xi = \mathcal{F}_x^{-1}(2\pi |\xi| \mathcal{F}_x u) \tag{12}$$

So, $\mathcal{F}_x(\partial_t u) = 2\pi i|\xi|\mathcal{F}_x u$ and $\mathcal{F}_x(\sqrt{-\Delta_x}u) = 2\pi|\xi|\mathcal{F}_x u$. Unsurprisingly, we can derive the half wave equation from the above and the linearity of the inverse Fourier transform:

$$\partial_t u - i\sqrt{-\Delta_x}u = \mathcal{F}_x^{-1}(2\pi i|\xi|\mathcal{F}_x u) - i\mathcal{F}_x^{-1}(2\pi|\xi|\mathcal{F}_x u) = 0 \tag{13}$$

We shall not prove it here, but it can be easily verified that $\mathcal{F}_{x,t}(\partial_t u)(\xi,\tau) = 2\pi i\tau\mathcal{F}_{x,t}u(\xi,\tau)$ and $\mathcal{F}_{x,t}(\sqrt{-\Delta_x}u)(\xi,\tau) = 2\pi|\xi|\mathcal{F}_{x,t}u(\xi,\tau)$. It follows immediately that

$$0 = \mathcal{F}_{x,t}(0) = \mathcal{F}_{x,t}(\partial_t u - i\sqrt{-\Delta_x}u) = 2\pi i(\tau - |\xi|)\mathcal{F}_{x,t}u \tag{14}$$

Hence, $\mathcal{F}_{x,t}u \neq 0$ implies that $\tau = |\xi|$ i.e. the support of the spacetime Fourier transform of u is contained within the light cone: $\operatorname{Supp}\mathcal{F}_{x,t}u \subseteq \{(\xi,\tau) \in \mathbb{R}^n_\xi \times \mathbb{R}_\tau : \tau = |\xi|\}$.

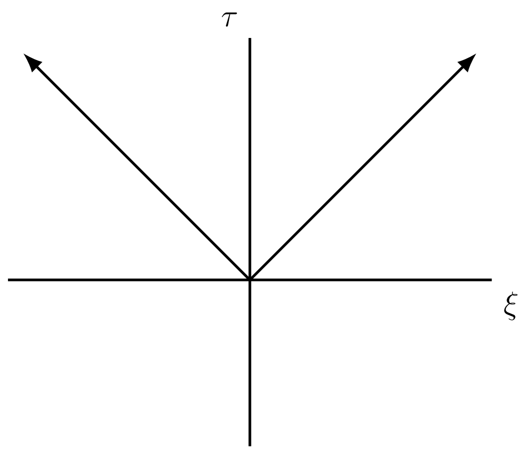

FIGURE 2. The light cone in $\mathbb{R}_\xi \times \mathbb{R}_\tau$.

Recall that $u_\theta = \mathcal{F}_x^{-1}\left(e^{2\pi it|\xi|}\mathcal{F}_x f_\theta\right)$. So,

$$\begin{aligned}
\mathcal{F}_{x,t}u_\theta &= \mathcal{F}_{x,t}\left(\mathcal{F}_x^{-1}\left(e^{2\pi it|\xi|}\mathcal{F}_x f_\theta\right)\right) \\
&= \mathcal{F}_t \circ \mathcal{F}_x \circ \mathcal{F}_x^{-1}\left(e^{2\pi it|\xi|}\mathcal{F}_x f_\theta\right) \\
&= \mathcal{F}_t\left(e^{2\pi it|\xi|}\mathcal{F}_x f_\theta\right) \\
&= \mathcal{F}_t\left(e^{2\pi it|\xi|}\right)\mathcal{F}_x f_\theta \\
&= \delta(\tau - |\xi|)\mathcal{F}_x f_\theta
\end{aligned} \tag{15}$$

Also recall that $\theta \subset \mathbb{R}^n_\xi$ contains the support of $\mathcal{F}_x f_\theta$. So, if $(\xi,\tau) \notin \theta \times \mathbb{R}_\tau \subset \mathbb{R}^n_\xi \times \mathbb{R}_\tau$, then $\mathcal{F}_x f_\theta(\xi) = 0$, which in turn implies that $\mathcal{F}_{x,t}u_\theta(\xi,\tau) = 0$ by the above. Hence, $\operatorname{Supp}\mathcal{F}_{x,t}u_\theta \subseteq \theta \times \mathbb{R}_\tau$. Combining this with the light cone description of the support of

$\mathcal{F}_{x,t}u$, we have that the support of $\mathcal{F}_{x,t}u_\theta$ is contained within the projection of $\theta \subset \mathbb{R}^n_\xi$ onto the light cone: $\operatorname{Supp}\mathcal{F}_{x,t}u_\theta \subseteq \{(\xi, |\xi|) \in \mathbb{R}^n_\xi \times \mathbb{R}_\tau : \xi \in \theta\}$.

Recall the heuristic from the third lecture that for an interval $I \subset \mathbb{R}$ and $a \in I$ that $\widehat{\mathbb{1}}_I(\xi) \overset{\text{Heuristic}}{=} e^{-2\pi i a\xi}|I|\mathbb{1}_{[0,\frac{1}{|I|}]}(\xi)$. With this, we have that

$$\begin{aligned}
\mathcal{F}_{x,t}\left(u_\theta \mathbb{1}_{[1,2]_t}\right)(\xi,\tau) &= \mathcal{F}_t\left(\mathcal{F}_x u_\theta \mathbb{1}_{[1,2]_t}\right) \\
&= \mathcal{F}_{x,t}u_\theta *_\tau \mathcal{F}_t\left(\mathbb{1}_{[1,2]_t}\right) \\
&\overset{\text{Heuristic}}{=} \mathcal{F}_{x,t}u_\theta *_\tau \left(e^{-2\pi i\xi}\mathbb{1}_{[0,1]_\tau}\right) \\
&= \left(\delta(\tau - |\xi|)\mathcal{F}_x f_\theta\right) *_\tau \left(e^{-2\pi i\xi}\mathbb{1}_{[0,1]_\tau}\right) \\
&= \int_{\mathbb{R}} \delta(\tau - s - |\xi|)\mathcal{F}_x f_\theta e^{-2\pi i\xi}\mathbb{1}_{[0,1]_\tau}(s)ds \\
&= \mathcal{F}_x f_\theta(\xi) e^{-2\pi i\xi}\mathbb{1}_{[0,1]_\tau}(\tau - |\xi|)
\end{aligned} \tag{16}$$

So, we have (heuristically) that the support of $\mathcal{F}_{x,t}\left(u_\theta\mathbb{1}_{[1,2]_t}\right)$ is contained within the one-neighbourhood of the projection of the rectangle θ onto the light cone, R_θ i.e. $\operatorname{Supp}\mathcal{F}_{x,t}\left(u_\theta\mathbb{1}_{[1,2]_t}\right) \overset{\text{Heuristic}}{\subseteq} (\theta \times \mathbb{R}_t) \cap \{(\xi,\tau) \in \mathbb{R}^n_\xi \times \mathbb{R}_\tau : 0 \le \tau - |\xi| \le 1\} \subseteq \bigcup_{\xi\in\theta} B_1(\xi, |\xi|) = R_\theta$. The forthcoming arguments involving R_θ would work just as well if we had defined R_θ as the thickened projection of the rectangle θ onto the light cone: $R_\theta = \{(\xi,\tau) \in \mathbb{R}^n_\xi \times \mathbb{R}_\tau : |\tau - |\xi|| \le 1, \xi \in \theta\}$.

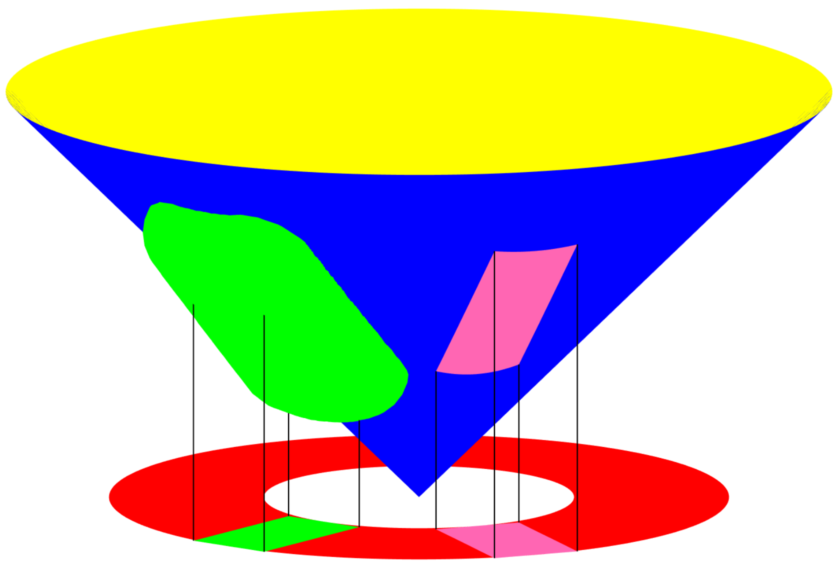

FIGURE 3. The light cone in $\mathbb{R}^2_\xi \times \mathbb{R}_\tau$, a green and a pink rectangle in Π, our covering of the annulus $\{|\xi| \sim 3\}$, the projection of the pink rectangle onto the light cone and $R_{\text{Green}} \supseteq \operatorname{Supp}\mathcal{F}_{x,t}(u_{Green}\mathbb{1}_{[1,2]})$.

Although theorems one and two as posed are specifically related to wave equation solutions, many of the properties of our solution u and the components of its decomposition u_θ are not specifically related to the wave equation; they are a consequence of the properties of their Fourier transformations, as we have seen in our investigation above. With this in mind, we will now state the more general theorem three that focuses on the Fourier related properties we have established. Theorems two and (consequently) one are implied by theorem three.

Theorem Three. If $p \geq 2$, $R \geq 1$, each $u_\theta \in \mathcal{S}(\mathbb{R}^{n+1})$ with $\operatorname{Supp} \widehat{u_\theta} \subseteq R_\theta$ defined as above and

$$d(p) = \begin{cases} 0 & \text{if } 2 \leq p \leq \frac{2(n+1)}{n-1} \\ \frac{n-1}{2}\left(\frac{1}{2} - \frac{1}{p}\right) - \frac{1}{p} & \text{if } p \geq \frac{2(n+1)}{n-1} \end{cases} \tag{17}$$

Then,

$$\left\| \sum_{\theta \in \Pi} u_\theta \right\|_{L^p(\mathbb{R}^{n+1})} \lessapprox R^{d(p)} \left(\sum_{\theta \in \Pi} \|u_\theta\|^2_{L^p(\mathbb{R}^{n+1})} \right)^{\frac{1}{2}} \tag{18}$$

This theorem is proved by interpolating on p. We immediately have that it is true when $p = 2$, by the approximate orthogonality of $(\widehat{u_\theta})_{\theta \in \Pi} \subset L^2(\mathbb{R}^{n+1})$:

$$\Big\| \sum_{\theta \in \Pi} u_\theta \Big\|^2_{L^2(\mathbb{R}^{n+1})} = \Big\| \sum_{\theta \in \Pi} \widehat{u_\theta} \Big\|^2_{L^2(\mathbb{R}^{n+1})} \approx R^0 \sum_{\theta \in \Pi} \Big\| \widehat{u_\theta} \Big\|^2_{L^2(\mathbb{R}^{n+1})} \tag{19}$$

where $d(2) = 0$. The $p = \infty$ case follows from an application of the $\ell^2(\Pi)$ Hölder inequality:

$$\begin{aligned} \Big| \sum_{\theta \in \Pi} u_\theta(z) \Big|^2 &\leq \Big(\sum_{\theta \in \Pi} |u_\theta(z)| \Big)^2 \\ &\leq |\Pi| \Big(\sum_{\theta \in \Pi} |u_\theta(z)|^2 \Big) \\ &\approx R^{\frac{n-1}{2}} \sum_{\theta \in \Pi} |u_\theta(z)|^2 \\ &\leq R^{\frac{n-1}{2}} \sum_{\theta \in \Pi} \|u_\theta\|^2_{L^\infty(\mathbb{R}^{n+1})} \end{aligned} \tag{20}$$

So,

$$\Big\| \sum_{\theta\in\Pi} u_\theta \Big\|_{L^\infty(\mathbb{R}^{n+1})} \lesssim R^{\frac{n-1}{4}} \left(\sum_{\theta\in\Pi} \|u_\theta\|^2_{L^\infty(\mathbb{R}^{n+1})} \right)^{\frac{1}{2}} \tag{21}$$

where $d(\infty) = \frac{n-1}{4}$. Thus to interpolate, we only need that for the Tomas-Stein exponent $p_0 = \frac{2(n+1)}{n-1}$

$$\left\| \sum_{\theta\in\Pi} u_\theta \right\|_{L^{p_0}(\mathbb{R}^{n+1})} \lessapprox \left(\sum_{\theta\in\Pi} \|u_\theta\|^2_{L^{p_0}(\mathbb{R}^{n+1})} \right)^{\frac{1}{2}} \tag{22}$$

It follows immediately from the Cauchy-Schwartz inequality for $\ell^2(\Pi)$ that

$$\left\| \sum_{\theta\in\Pi} u_\theta \right\|_{L^{p_0}(\mathbb{R}^{n+1})} \leq \sum_{\theta\in\Pi} \|u_\theta\|_{L^{p_0}(\mathbb{R}^{n+1})} \leq |\Pi|^{\frac{1}{2}} \left(\sum_{\theta\in\Pi} \|u_\theta\|^2_{L^{p_0}(\mathbb{R}^{n+1})} \right)^{\frac{1}{2}} \tag{23}$$

This estimate is insufficient; we need to drop the $|\Pi|^{\frac{1}{2}}$ from this trivial estimate. However, it makes our goal stated above seem reasonable. In the next section, we will investigate how inequality (22) holds for $p_0 = \frac{2(n+1)}{n-1}$ using Pramanik and Seeger's iteration method which yields cone decoupling estimates using parabolas.

Parabola Decoupling in Two Dimensions

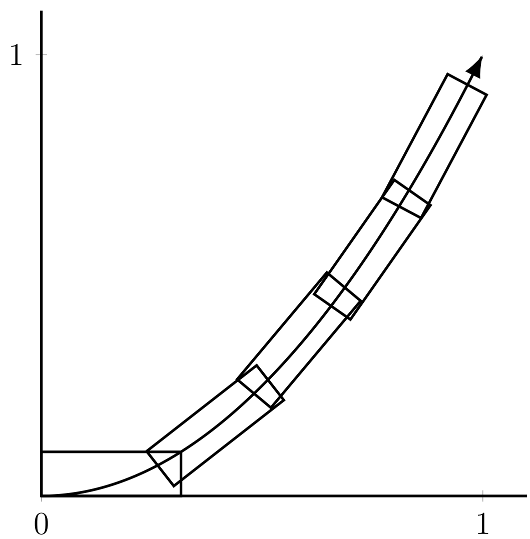

Figure 4. A covering of the parabola neighbourhood by Π.

We will state the following theorem without proof, but we will investigate its plausibility later. We will only consider the two dimensional argument as it is the most illustrative.

Theorem Four (Bourgain and Demeter). Suppose Π is a collection of $R^{-\frac{1}{2}} \times R^{-1}$ rectangles which cover an arc of the parabola $\{(\xi, \xi^2) : \xi \in (0,1)\}$, $p > 2$, $f_\theta \in \mathcal{S}(\mathbb{R}^2)$ with Supp $\widehat{f_\theta} \subset \theta$ for all $\theta \in \Pi$ and

$$d(p) = \begin{cases} 0 & \text{if } 2 \le p \le \frac{2(n+1)}{n-1} = 6 \\ \frac{n-1}{2}\left(\frac{1}{2} - \frac{1}{p}\right) - \frac{1}{p} = \frac{1}{2}\left(\frac{1}{2} - \frac{3}{p}\right) & \text{if } p \ge \frac{2(n+1)}{n-1} = 6 \end{cases} \tag{24}$$

Then,

$$\left\| \sum_{\theta \in \Pi} f_\theta \right\|_{L^p(\mathbb{R}^2)} \lessapprox R^{d(p)} \left(\sum_{\theta \in \Pi} \|f_\theta\|^2_{L^p(\mathbb{R}^2)} \right)^{\frac{1}{2}} \tag{25}$$

Referring to the bottom-right rectangle in the figure above, we see that the rectangles are maximal; if their proportions changed, then the parabola wouldn't go through the it's top-right corner. The bounding term is a multiple of the $\ell^2(\Pi)$ norm of $(\|f_\theta\|_{L^p(\mathbb{R}^2)})_{\theta \in \Pi}$.

It may not be obvious why we want to give a decoupling estimate for a thickened parabola; let us outline our plan of attack: instead of just tiling the light cone by thickened rectangles and then decoupling, similar to what we did in theorem three, we will cover the light cone by thickened parabolas then decouple and then cover those by thickened rectangles and decouple again. This will give us a tighter bound than just covering and decoupling with rectangles.

Cone Decoupling in Three Dimensions

Now, we turn our attention to the decoupling on a cone in three dimensions. The light cone in three dimensions is given by the equation $\xi_2^2 = \xi_3^2 - \xi_1^2 = (\xi_3 + \xi_1)(\xi_3 - \xi_1)$ where $\xi_3 \ge 0$. Lets investigate an approximation of the cone at a point. We can consider a point $(0, t, t)$ without loss of generality due to the rotational symmetry of the cone. Consider the linear change of variables $\eta_1 = \xi_3 + \xi_1$, $\eta_3 = \xi_3 - \xi_1$ and $\eta_2 = \xi_2$ where the equation for the light cone becomes $\eta_2^2 = \eta_1 \eta_3$. We can think of the light cone as the graph of the function $f : \mathbb{R}_{\eta_2} \times \mathbb{R}_{\eta_3} \to \mathbb{R}_{\eta_1}$ given by $f(x, y) = \frac{x^2}{y}$. The second order Taylor approximation of f about the point (t, t, t) is $f(x, y) = 2x - y + \frac{(x-y)^2}{t} + O\Big(\frac{\|(x-t, y-t)\|^3}{t^2}\Big)$. Up to terms of second order, f is linear along the span of $(1, 1)$ as $f(a, a) \approx a$ and quadratic along its orthogonal complement, the span of (1,-1), as $f(a, -a) \approx 3a + \frac{4a^2}{t}$; the graph of f is a parabolic cylinder. Hence, we can approximate a neighbourhood of the point (t, t, t) on the cone by a neighbourhood of (t, t, t) on the parabolic cylinder. The intersection of this approximating parabolic cylinder and the cone is a parabola (seen in red below). We will not show them here, but some tedious calculations show that points with the same $\mathbb{R}_{\eta_1}$ component (height) generate and sit on the same intersecting parabola despite having different Taylor expansions about them (which are equations for parabolic cylinders). The Taylor expansion of a light cone implies that we must use parabola, not another conic section, for our decoupling estimate. We state without proof that Pramanik-Seeger iteration, which uses our decoupling estimate for thickened parabolas and Fubini's theorem, yields the estimate for the cone.

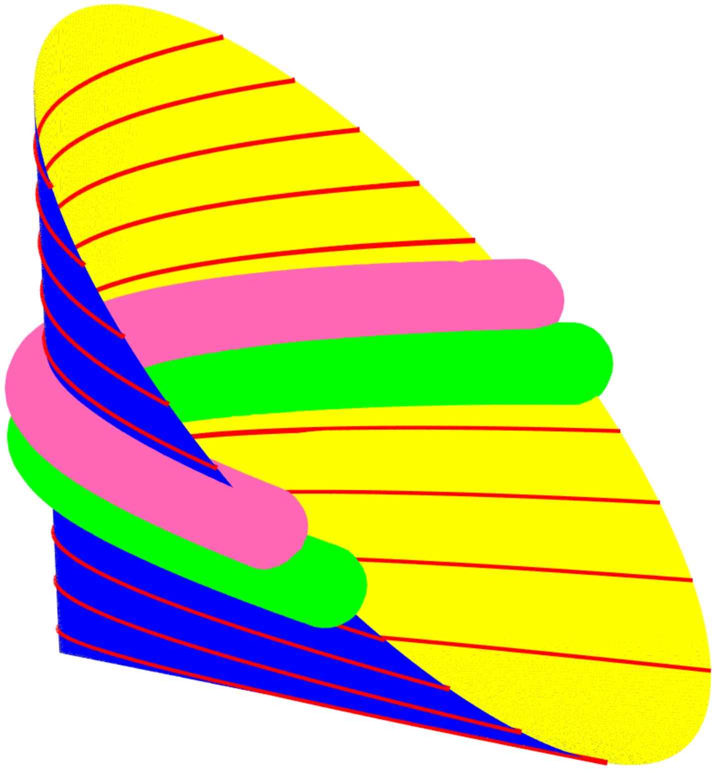

Figure 5. The transformed light cone in $\mathbb{R}_{\eta_2} \times \mathbb{R}_{\eta_3} \times \mathbb{R}_{\eta_1}$, red parabolas on the surface of the light cone and some (pink and green) thickened parabolas in the covering of the light cone.

Comparing Reverse Square Estimates and Decoupling Methods

We will now investigate how we might prove theorem four. Naturally, we ask why it is that ℓ^2 decoupling easier than reversed square function estimates? This is because we can use an induction on scale argument to prove the former that we hinted at earlier when we introduced theorem two. Bourgain and Demeter's original proof uses induction on scale along with multi-linear techniques and multi-linear Kakeya estimates. Simpler proofs (that we will investigate shortly) using connections to discrete phenomena were later found. We will now develop an intuition for how theorem four can be proved inducting on the scale of R. Recall the inequality of interest is

$$\left\|\sum_{\theta\in\Pi} f_\theta\right\|_{L^p(\mathbb{R}^2)} \lessapprox R^{d(p)}\left(\sum_{\theta\in\Pi}\|f_\theta\|^2_{L^p(\mathbb{R}^2)}\right)^{\frac{1}{2}} \tag{26}$$

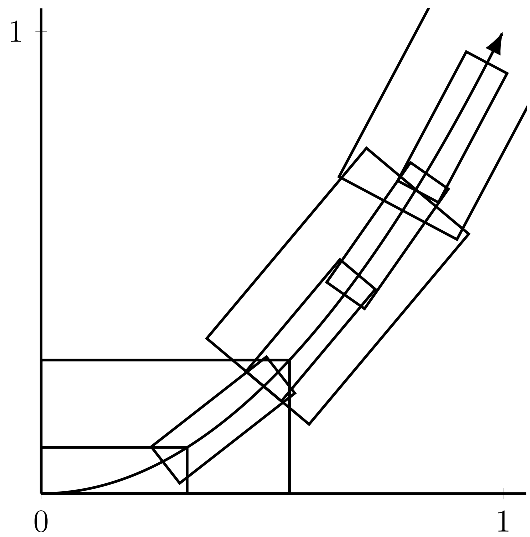

FIGURE 6. A covering of the parabola neighbourhood by Π and Λ.

For our induction argument, we will assume that the estimate holds for all $S < R$. That is to say, that for any S such that $S^{-1} > R^{-1}$, there is a collection of $R^{-\frac{1}{2}} \times R^{-1}$ (small) rectangles, Π, covering R^{-1} neighbourhood of the parabola for which the estimate holds and a collection of $S^{-\frac{1}{2}} \times S^{-1}$ (big) rectangles, Λ, for which we will prove the estimate. Our induction hypothesis is

$$\left\|\sum_{\theta\in\Pi} f_\theta\right\|_{L^p(\mathbb{R}^2)} \approx \left\|\sum_{\tau\in\Lambda}\left(\sum_{\theta\subset\tau} f_\theta\right)\right\|_{L^p(\mathbb{R}^2)} \overset{\text{Induction Hypothesis}}{\leq} S^{d(p)}\left(\sum_{\tau}\left\|\sum_{\theta\subset\tau} f_\theta\right\|^2_{L^p(\mathbb{R}^2)}\right)^{\frac{1}{2}} \tag{27}$$

Our induction hypothesis implies that for all $\tau \in \Lambda$,

$$\left\|\sum_{\theta\subset\tau} f_\theta\right\|^2_{L^p(\mathbb{R}^2)} \leq \left(\frac{R}{S}\right)^{2d(p)}\left(\sum_{\theta\subset\tau}\|f_\theta\|^2_{L^p(\mathbb{R}^2)}\right) \tag{28}$$

From which we get

$$\left\|\sum_{\theta} f_\theta\right\|_{L^p(\mathbb{R}^2)} \overset{\text{Almost}}{\leq} S^{d(p)}\cdot\left(\frac{R}{S}\right)^{d(p)}\left(\sum_{\tau\in\Lambda}\sum_{\theta\subset\tau}\|f_\theta\|^2_{L^p(\mathbb{R}^2)}\right)^{\frac{1}{2}} \tag{29}$$

Which is not bad at all; intuitively, it seems plausible we can use this line of reasoning to get the result we want. Decoupling works much better than the classical square estimates results. Using the reverse square estimate method, all we get is that

$$\left\|\sum_{\theta\in\Pi} f_\theta\right\|_{L^p} = \left\|\sum_{\tau\in\Lambda}\left(\sum_{\theta\subset\tau} f_\theta\right)\right\| \overset{\text{Reverse Square Function}}{\lesssim} \left\|\left(\sum_{\tau}\left|\sum_{\theta\subset\tau} f_\theta\right|^2\right)^{\frac{1}{2}}\right\|_{L^p} \tag{30}$$

and

$$\left\|\sum_{\theta\subset\tau} f_\theta\right\| \overset{\text{Reverse Square Function}}{\lesssim} \left\|\left(\sum_{\theta\subset\tau} |f_\theta|^2\right)^{\frac{1}{2}}\right\|_{L^p} \tag{31}$$

but this does not close the desired induction.

Discrete Restriction Estimates and Exponential Sum Estimates

Recall that the Tomas-Stein exponent was derived by setting $d(p) = 0$. This is no idle coincidence; ℓ^2-decoupling estimates for paraboloids have links to the discrete analogue of the Tomas-Stein paraboloid theorem.

Paraboloid Theorem (Tomas and Stein). Let $d\sigma$ be the surface measure of a paraboloid P in $\mathbb{R}^n$, $f \in L^2(P)$, $x \in \mathbb{R}^n$ and set $(fd\sigma)^\vee(x) = \int_P f(\xi)e^{ix\cdot\xi}d\sigma(\xi)$. Then, there is some constant C_n such that

$$\|(fd\sigma)^\vee\|_{L^{\frac{2(n+1)}{n-1}}(\mathbb{R}^n)} \leq C_n \|f\|_{L^2(P)} \tag{32}$$

$(fd\sigma)^\vee$ describes the dispersion of waves travelling in different directions. Now, consider a similar setup in a periodic setting.

Discrete Paraboloid Theorem. Let $D = \{(k_1, \dots, k_n) \in \mathbb{Z}^{n-1}\times\mathbb{Z} : \sum_{i=1}^{n-1} |k_i|^2 = k_n\}$ be a discrete paraboloid in $\mathbb{R}^n$, $d\sigma$ the discrete measure on it, $f \in \ell^2(D)$, $x \in \mathbb{Z}^n$ and set $(fd\sigma)^\vee(x) = \sum_{k\in D} f(k)e^{2\pi ix\cdot k}$. Then for every $\varepsilon > 0$, there is some constant C_ε such that

$$\left\|\sum_{k\in D, |k|\leq N} f(k)e^{2\pi ix\cdot k}\right\|_{L^{\frac{2(n+1)}{n-1}}([0,1]^n)} \leq C_\varepsilon N^\varepsilon \left(\sum_{k\in D,|k|\leq N} |f(k)|^2\right)^{\frac{1}{2}} \tag{33}$$

This can be proved using Bourgain and Demeter's decoupling theorem. Instead of this weak result, we might analogously have expected the stronger result that there is some constant C_n such that

$$\|(fd\sigma)^\vee\|_{L^{\frac{2(n+1)}{n-1}}([0,1]^n)} \leq C_n \left(\sum_{k\in D} |f(k)|^2\right)^{\frac{1}{2}} \tag{34}$$

This is in fact not true. Consider this counterexample: choose $N \in \mathbb{N}$ and consider the two dimensional case where $f_N(k) = 1$ for $1 \leq k \leq N$ and $f_N(k) = 0$ otherwise. As $N \to \infty$:

$$\|(f_N d\sigma)^\vee\|_{L^6([0,1]^2)} \gtrsim N^{\frac{1}{2}}(\log N)^{\frac{1}{6}} \tag{35}$$

Zane Li and Wooley used these exponential sum estimates to give a simple proof of ℓ^2 decoupling for the paraboloid in $\mathbb{R}^2$. Related ideas in other contexts have been used to study the zeros of the Riemann Zeta function.

Now, we will quickly look at how the difference between the normal and discrete paraboloid theorems affects the solutions to the Schrödinger equation (which is a modification of the wave equation). Consider a solution $u : \mathbb{R}^n_x \times \mathbb{R}_t \to \mathbb{C}$ to the Schrödinger equation $i\partial_t u(x,t) = \Delta_x u(x,t)$ on $\mathbb{R}^n$. It has the form $u(x,t) = \int_{\mathbb{R}^n} \widehat{f}(\xi) e^{2\pi i(x\cdot\xi + t|\xi|^2)} d\xi$. The operator disperses energy: wave packets at different frequencies move apart as t grows. The estimate given in the paraboloid theorem quantifies the dispersion.

Now consider a solution $v : (\mathbb{R}/\mathbb{Z})^n_x \times \mathbb{R}_t \to \mathbb{C}$ to the Schrödinger equation $i\partial_t v(x,t) = \Delta_x v(x,t)$ on the n-torus $(\mathbb{R}/\mathbb{Z})^n$. It has the form $v(x,t) = \sum_{k\in\mathbb{Z}^n} \widehat{f}(k) e^{2\pi i(x\cdot k + t|k|^2)}$. On the n-torus, frequencies are discrete and rational relations among the different values of $|k|^2$ allow waves to rephase and repeatedly overlap; the discrete paraboloid theorem does not give a good bound: there is less dispersion on the compact n-torus $(\mathbb{R}/\mathbb{Z})^n$ than on $\mathbb{R}^n$.