Now we come to the second main issue in the real case. We have Lemma 14.8. Following the finite field case, we would like to iterate this lemma. However, there is an issue with this iteration, which is that we do not know whether $Q(A)$ is a $(\delta, s+\epsilon, \delta^{-\epsilon})$-set, and in fact this is likely not true in general. Instead, we will use that $Q(A)$ contains a $(\delta, s+\epsilon, \delta^{-\epsilon})$-set. It takes significant extra work to prove this fact. We will discuss the issues more next time.

15. Bourgain's Projection Theorem II

April 10

Definition 15.1. *A $(\delta, s, C)_d$-set is a set $X \subset B^d(0,1)$ such that*

$$|X \cap B(x,r)|_\delta \leq Cr^s|x|_\delta.$$

Remark 15.2. *We think of a (δ, s, C) set as a set which is 'non-concentrated' on the scale δ with degree s.*

Using this language we can rewrite the Bourgain projection theorem as.

Theorem 15.3. *Given $0 < t < 2$, $0 < s \leq 1$, there exist $\varepsilon, \eta > 0$ such that*

If $X \subset B^2(0,1)$ is a $(\delta, t, \delta^{-\eta})_2$-set with $|X|_\delta = \delta^{-t}$ and $D \subset [0,1]$ is a $(\delta, s, \delta^{-\eta})_1$-set. Then there exists some $\theta \in D$ such that

$$\min_{\substack{X' \subset X \\ |X'|_\delta \geq \delta^\eta |X|_\delta}} |\pi_\theta X'| \geq \delta^{-\frac{t}{2}-\varepsilon}$$

Last time we saw that there exists a polynomial Q such that for every $0 < s < 1$ there exists $\varepsilon, \eta > 0$ such that if A is a $(\delta, s, \delta^{-\eta})_1$-subset of $[0,1]$ and $|A|_\delta = \delta^{-s}$ then $|Q(A)|_\delta \geq \delta^{-s-\varepsilon}$. Now we cannot yet iterate this because we do not know that $Q(A)$ is a non-concentrated, in fact this is not true, but we can ask for $Q(A)$ to contain a $(\delta, s+\varepsilon, \delta^{-\eta})$ set (though with different ε, η).

In these notes, we discuss some of the ideas to deal with this technical issue, although we don't give a complete proof.

Let us quickly confirm some properties of non-concentrated sets.

Lemma 15.4. *If X is a $(\delta, s, C)_d$-set then:*

1. *$|X|_\rho \geq C^{-1}\rho^{-s}$ for all $\rho \in [\delta, 1]$.*
2. *If $Y \subset X$ and $|Y|_\delta \geq \frac{1}{K}|X|_\delta$ then Y is a $(\delta, s, CK)_1$-set.*

Intuitively (i) tells us that if X is non-concentrated on scale δ then it is large on all scales at least δ, (ii) tells us that this concept is preserved under taking 'dense' subsets.

Proof. (1) If $X \subset \bigcup_{i=1}^m B(x_i,\rho)$ then

$$|X|_\delta \le \sum_{i=1}^m |X \cap B(x_i,\rho)|_\delta$$

but we know that $|X \cap B(x_i,\rho)|_\delta \le C\rho^s|X|_\delta$ so

$$|X|_\delta \le mC\rho^s|X|_\delta \implies m \ge C^{-1}\rho^{-s}.$$

(2) This is even simpler since

$$|Y \cap B(x,\rho)|_\delta \le |X \cap B(x,\rho)|_\delta \le C\rho^s|X|_\delta \le (CK)\rho^s|Y|_\delta$$

□

Now due to this lemma if we want $Q(A)$ to contains a $(\delta, s+\varepsilon, \delta^{-\eta})$ set then it must be that $|Q(A)|_\rho \ge \rho^{-s-\varepsilon}$ $\forall\delta \in [\delta, 1]$.

Now we notice two important things about the above property.

- We don't get this for free because A need not be a $(\rho, s, \delta^{-\eta})$-set for $\rho \in [\delta, 1]$.
- This property is necessary but not sufficient.

We can fix both of these problems with one framework, that of the 'uniform set', which is very useful even outside of this theory.

Assume that δ is some negative power of 2, we will denote by $\mathcal{D}_\delta$ the set of δ-mesh cubes tiling $\mathbb{R}^d$. For any given set X we denote by $\mathcal{D}_\delta(X)$ the set of those cubes that intersect X. We then define $|X|^*_\delta := |\mathcal{D}_\delta(X)|$ and notice that $|X|^*_\delta \sim |X|_\delta$ as we saw in the last lecture.

Definition 15.5. *Given $\Delta \in 1/\mathbb{N}$ and $m \in \mathbb{N}$, A set $X \subset [0,1]^d$ is (Δ, m)-uniform if for any $j \in \{0, \dots, m-1\}$ and for any cube $Q \in \mathcal{D}_\delta(X)$ we have*

$$|Q \cap X|^*_{\Delta^{j+1}} = R_j$$

where R_j is independent of Q.

The numbers R_j are called 'branching factors' of X.

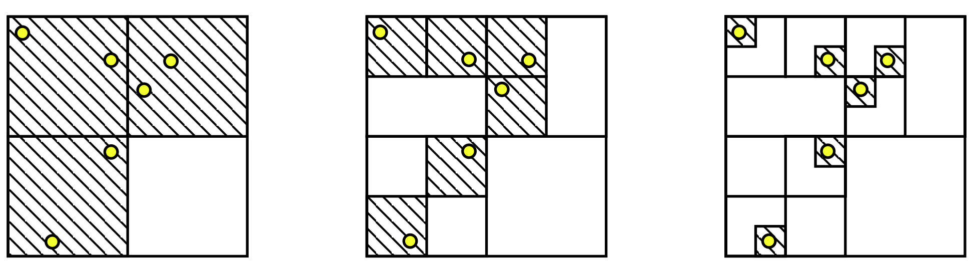

$R_1 = 3$ $R_2 = 2$ $R_3 = 1$

FIGURE 14. A $(1/2, 2)$-uniform set with its 3 branching factors

We can see why these uniform sets are useful with the following lemma.

Lemma 15.6. *Let $X \subset [0,1]^d$ be a (Δ, m)-uniform set and let $\delta = \Delta^m$.*

(a) If $|X|_\rho \geq C^{-1}\rho^{-s}$ for all $\rho \in \{1, \Delta, \Delta^2, \ldots, \Delta^m\}$, then X is a $(\delta, s, O_\Delta(C))_\delta$ set.

(b) If X is a (δ, s, C) set then X is also a $(\rho, s, O_\Delta(C))$ for all $\rho \in [\delta, 1]$.

If we believe this, and we know that A and $Q(A)$ are both uniform, then that immediately solves both our problems and lets us continue the proof. Before we explain how to make A and $Q(A)$ uniform let us prove this lemma.

Proof. (a) Let $\rho = \Delta^j$ and Q some cube in $\mathcal{D}_{\Delta^j}(X)$. We clearly have the recursive relation $|X \cap Q|_{\Delta^{i+1}} = R_i |X \cap Q|^*_{\Delta^i}$ which when iterated gives us

$$|X \cap Q|^*_{\Delta^m} = R_j R_{j+1} \cdots R_{m-1} |X \cap Q|^*_{\Delta^j}$$

but we know that $|X \cap Q|^*_{\Delta^j} = 1$ precisely because Q is a Δ^j cube. We thus have

$$|X \cap Q|^*_{\Delta^m} = \frac{R_0 R_1 \cdots R_{m-1}}{R_0 R_1 \cdots R_{j-1}},$$

Now the numerator here is precisely $|X|^*_{\Delta^m}$ and the denominator is $|X|^*_{\Delta^j}$ so by assumption we have $|X|^*_{\Delta^j} \gtrsim C^{-1}\rho^{-s}$ which gives us

$$|X \cap Q|^*_{\Delta^m} \lesssim C\rho^s |X|^*_{\Delta^m}.$$

This shows that X is a $(\delta, s, C)_\delta$ set at scales $1, \Delta, \ldots, \Delta^m$. For the scales in between we can sandwich them between two powers of Δ, this loses us an extra factor of at most $O_\Delta(C)$.

(b) Again let $\rho = \Delta^{j_0}$, then for any j with $0 \leq j \leq j_0$, let Q be some square in $\mathcal{D}_{\Delta^j}(X)$ then we have again

$$|X \cap Q|_{\Delta^{j_0}} = \frac{R_0 \cdots R_{j_0-1}}{R_0 \cdots R_{j-1}} = \frac{|X|^*_\rho}{|X|^*_{\Delta^j}} \lesssim C\rho^s |X|_\rho$$

where in the last step we applied the previous lemma for (δ, s, C) sets. Again the sandwiching gives us an extra factor of $O_\Delta(C)$. □

Now we learn an important tool, which is the method to make any set uniform.

Lemma 15.7 (Uniformization)**.** *Let $\delta = \Delta^m$, $X \subset [0,1]^d$, and let μ be an arbitrary sub-additive set function (eg. $\mu(B) = |B|_\delta$). Then there exists a subset $Y \subset X$ such that Y is (Δ, m)-uniform and*

$$\mu(Y) \geq \left[2d \ln\left(\frac{1}{\Delta}\right)\right]^{-m} \mu(X) = \delta^{-\sigma}\mu(X)$$

where $\sigma = \frac{\ln(2\ln\frac{1}{\Delta})}{\ln\frac{1}{\Delta}}$.

Note that as $\Delta \to 0$ we have $\sigma \to 0$ so we can make this power arbitrarily small by picking Δ at the end.

Proof. We will construct a uniform subset by thinking of X as a tree, and pruning it from the leaves to make it uniform. We will do this step by step, first we set $X_m = X$, then at each step, from X_j we construct X_{j-1} by removing enough mass from level j to make it uniform.

To do this let $X_{j-1,\ell} = \{Q \in \mathcal{D}_{\Delta^{j-1}}(X) : |Q \cap X|^*_{\Delta^j} \in [2^\ell, 2^{\ell+1}]\}$ where ℓ ranges between 0 and $d \ln \frac{1}{\Delta}$. This splits X_j into $d \ln \frac{1}{\Delta}$ different pieces across which we have similar magnitude branching on level j. Then because we have a sub-additive function

$$\mu(X_j) \leq \sum_{\ell=0}^{d \ln \frac{1}{\Delta}} \mu(X_{j-1,\ell}).$$

so we can pick the 'largest' piece and lost at most a factor of $d \ln 1\Delta$. Assume that $X_{j-1,\ell}$ is that piece, we set X_{j-1} to be the $X_{j-1,\ell}$ where at level j we removed enough of the set to get the branching factor to be exactly 2^ℓ. Since the branching factors are all within a factor of 2 away from 2^ℓ this loses us at most half of the 'measure' of $X_{j-1,\ell}$ so that

$$2d \ln \frac{1}{d} \mu(X_{j-1}) \geq d \ln \frac{1}{d} \mu(X_{j-1,\ell}) \geq \mu(X_j)$$

iterating this process m times gives us exactly the lemma. □

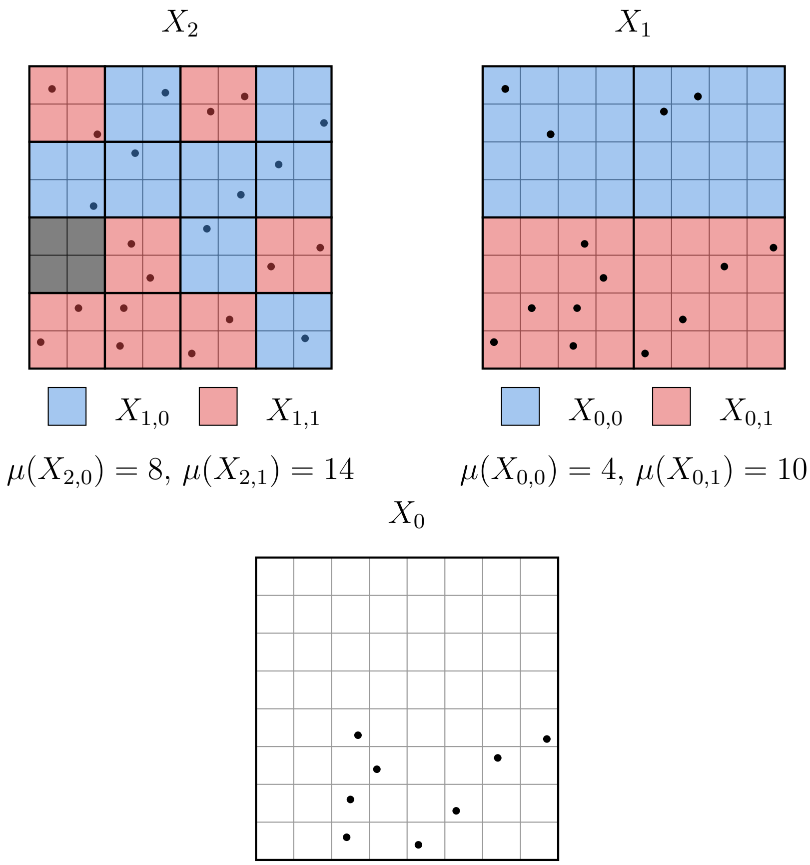

FIGURE 15. Applying uniformization with $\Delta = 1/2$ and $m = 2$ to a set.

Now we return to our original goal. We recall that A is a (δ, s, C)-set. By Lemma 14.8 from last lecture, we know that $|Q(A)|_\delta \geq \delta^{-s-\varepsilon}$, where Q is a fixed polynomial. However, we don't yet know whether $Q(A)$ contains a $(\delta, s+\epsilon, C')$ set, and so we cannot iterate.

Using the uniformization lemma, we can reduce to the case that A is uniform. In this case, we know that A is a (ρ, s, C) set for all $\rho \geq \delta$. Now, by Lemma 14.8, we know that $|Q(A)|_\rho \geq \rho^{-s-\varepsilon}$ for all $\rho \geq \delta$. If we knew that $Q(A)$ was uniform, then it would follow that $Q(A)$ is a $(\delta, s+\epsilon, C')$ set with a reasonable C'. However, just because A is uniform, it does not tells us that $Q(A)$ is uniform.

The main enemy here is that $|Q(A)|_\rho$ may be large, and $|Q(A)|_\delta$ may be large, but it could still happen that there is a subset $B \subset Q(A)$ so that $|B|_\rho \ll |Q(A)|_\rho$ and yet $|Q(A) \setminus B|_\delta \ll |Q(A)|_\delta$. (It's a good exercise to draw a picture of this scenario.)

This enemy scenario sounds somewhat bizarre and even unlikely, but it takes a fair amount of work to rule it out. And it involves somewhat changing the outline of the proof. We will discuss these somewhat technical but yet important issues next time.