ADDITIVE COMBINATORICS: LECTURE 13

LEO GOLDMAKHER

Today we are going to develop the final tool required to prove the Freiman-Ruzsa theorem. First, recall what we're trying to prove:

Theorem 1 (Freiman-Ruzsa)**.** *Given a finite set $A \subset \mathbb{Z}$ such that $|A+A| \leq K|A|$. Then there exists a gAP $Q \supseteq A$ such that $|Q| \ll_K |A|$ and $\dim Q \ll_K 1$.*

Proof Strategy. First, we embed A into some finite group $\mathbb{Z}_N$; let $B \subseteq \mathbb{Z}_N$ denote the image of A. Applying our Bohr set machinery (Bogolyubov's Lemma combined with the Bohr-to-gAP proposition) produces a large, low-dimensional proper gAP $Q \subseteq 2B - 2B$. Now pull back Q to a proper gAP inside $2A - 2A$. Finally, Ruzsa's Reduction Lemma (from Lecture 8) allows us to construct out of this a small low-dimensional gAP containing A.

'QED'

Let us consider this approach more carefully. The argument begins with an injective map $\phi : A \hookrightarrow \mathbb{Z}_N$. Set $B := \phi(A)$, so that ϕ is a bijection between A and B. The Bohr set machinery yields a proper gAP $Q \subseteq 2B - 2B$ of dimension $\leq \frac{1}{\delta^2}$, where $\delta = \frac{|B|}{N}$ is the density of B inside $\mathbb{Z}_N$; to guarantee that the dimension is small, we need δ to be fairly large. Equivalently:

$$\textit{We need an injection } \phi : A \hookrightarrow \mathbb{Z}_N \textit{ such that } N \textit{ not much larger than } |A|. \tag{1}$$

The next step of the argument is to pull back the gAP from $2B-2B$ to a gAP in $2A-2A$. In other words:

$$\begin{aligned}&\textit{Given a bijection } \phi : A \to B\textit{, we need to construct } \psi : 2B - 2B \to 2A - 2A \textit{ which}\\ &\textit{preserves proper gAPs.}\end{aligned} \tag{2}$$

Making ψ an injection guarantees that it preserves properness, but ensuring that ψ preserves gAPs is more demanding. The most natural injective map which preserves gAPs is a group isomorphism, but this is too rigid – there are no isomorphisms between $\mathbb{Z}$ and $\mathbb{Z}_N$, for example. Instead we will use a weaker notion, that of a *Freiman k-isomorphism*. Before defining this, we recall from last lecture the concept of a Freiman k-homomorphism:

Definition. *Given abelian groups G and H and subsets $A \subseteq G$ and $B \subseteq H$, a* Freiman k-homomorphism *from A to B is a map $\phi : A \to B$ which satisfies the property*

$$\begin{gathered}a_1 + a_2 + \cdots + a_k = a_1' + a_2' + \cdots + a_k'\\ \Downarrow\\ \phi(a_1) + \phi(a_2) + \cdots + \phi(a_k) = \phi(a_1') + \phi(a_2') + \cdots + \phi(a_k')\end{gathered}$$

for any $a_i, a_i' \in A$.

Date: April 11, 2014.

Last lecture we made several observations:

- any genuine homomorphism is a k-homomorphism for any k;
- a k-homomorphism is an ℓ-homomorphism for any $\ell \leq k$;
- a translation is a k-homomorphism for every k; and
- 2-homomorphisms map APs to APs; injective 2-homomorphisms map proper APs to proper APs.

Exercise 1. *Show that an injective 2-homomorphism maps proper gAPs to proper gAPs of the same size and dimension.*

We can now define a Freiman k-isomorphism:

Definition. *Given abelian groups G and H and subsets $A \subseteq G$ and $B \subseteq H$, a* Freiman k-isomorphism *from A to B is a bijective map $\phi : A \to B$ such that both ϕ and ϕ^{-1} are k-homomorphisms. If this is the case, we write $A \simeq_k B$.*

Equivalently, $\phi : A \to B$ is a k-isomorphism if

$$\begin{gathered} a_1 + a_2 + \cdots + a_k = a'_1 + a'_2 + \cdots + a'_k \\ \Updownarrow \\ \phi(a_1) + \phi(a_2) + \cdots + \phi(a_k) = \phi(a'_1) + \phi(a'_2) + \cdots + \phi(a'_k) \end{gathered}$$

Exercise 2. *Show that there exist bijective k-homomorphisms which are* not *k-isomorphisms.*

It's clear that k-isomorphisms satisfy one of the requirements from (2) – they preserve proper gAPs – but it's not obvious how to extract a k-isomorphism $\psi : 2B - 2B \to 2A - 2A$ from an ℓ-isomorphism $\phi : A \to B$. Some thought suggests the following.

Exercise 3. *If $A \simeq_8 B$, then $2A - 2A \simeq_2 2B - 2B$.*

This completely settles our requirement (2), leaving us to focus on (1): how to k-isomorphically embed an arbitrary set $A \subseteq \mathbb{Z}$ into $\mathbb{Z}_N$ for some N which isn't too large. For example, the natural embedding of $A = \{0, 2, 4000\}$ is into $\mathbb{Z}_{4001}$, which is huge compared to A and thus causes problems when we apply our Bohr set machinery. A more potent example is the set $\{p \leq n\}$ of all primes up to n; this can certainly be embedded into $\mathbb{Z}_n$, but the primes occupy an arbitrarily small proportion of $\mathbb{Z}_n$ as $n \to \infty$, so we would not be able to get any bound on the dimension of the gAP produced by the Bohr set machinery.

In the above two examples we saw that the most obvious embedding is not useful, because it embeds the set into a group which is so large that the Bohr set machinery gives poor bounds on the dimension of the gAP. Perhaps there's a clever way to k-isomorphically embed these sets into a smaller group? Unfortunately, it turns out that some sets cannot be k-isomorphically embedded into *any* small groups.

Example 2. *Let $A = \{1, 2, 4, \ldots, 2^{n-1}\}$. I claim that A isn't 2-isomorphic to any subset of $\mathbb{Z}_N$ unless N is large. For, suppose $A \simeq_2 B \subseteq \mathbb{Z}_N$. Then*

$$\frac{n(n+1)}{2} = |A + A| = |B + B| \leq N.$$

In particular, if $N \leq \frac{1}{2}|A|^2$ then A doesn't embed 2-isomorphically into $\mathbb{Z}_N$ (and hence, doesn't embed k-isomorphically for any $k \geq 2$). We could still use this to produce a gAP using our Bohr

set technology, but the dimension would be bounded by $\frac{|A|^2}{4}$ (as opposed to being bounded by a constant, as in Freiman-Ruzsa).

This shows that there's no hope of k-isomorphically embedding A directly into $\mathbb{Z}_N$. However, if $mA - mA$ isn't too big, it turns out to be possible to embed a large chunk of A into $\mathbb{Z}_N$:

Lemma 3 (Ruzsa's Modelling Lemma)**.** *Fix $m \in \mathbb{N}$ and a finite $A \subseteq \mathbb{Z}$. Then for any prime $N > 2|mA - mA|$, $\exists A' \subseteq A$ of size $|A'| \geq \frac{|A|}{m}$ such that A' embeds m-isomorphically into $\mathbb{Z}_N$.*

We will prove this result below, but first we show how to use it to prove the Freiman-Ruzsa theorem.

Proof of Freiman-Ruzsa. Given $A \subseteq \mathbb{Z}$ of small doubling, say $|A + A| \leq K|A|$. We wish to show that A is contained in a low-dimensional gAP which is not much larger than A itself.

STEP 1: *Embed a large chunk of A into $\mathbb{Z}_N$, for some $N \asymp_K |A|$; denote the image by $B \subseteq \mathbb{Z}_N$.*

Plünnecke-Ruzsa implies $|8A - 8A| \leq K^{16}|A|$. Applying the Ruzsa Modelling Lemma, we deduce that for any prime $N > 2K^{16}|A|$ we can 8-isomorphically embed most of A into $\mathbb{Z}_N$. More precisely, fix a prime N satisfying

$$2K^{16}|A| < N < 4K^{16}|A|$$

(such a prime exists by Bertrand's Postulate). Ruzsa's Modelling Lemma implies the existence of $A' \subseteq A$ and $B \subseteq \mathbb{Z}_N$ such that $A' \simeq_8 B$ and $|A'| \asymp |A|$.

STEP 2: *Find a large, low-dimensional proper gAP inside $2B - 2B$.*

It follows from Step 1 that

$$|B| = |A'| \asymp |A| \asymp_K N,$$

so that $|B| = \delta N$ with $\delta \asymp_K 1$. Bogolyubov's Lemma (Lecture 10) produces a Bohr set $B(R, 1/4) \subseteq 2B - 2B$ such that $|R| \leq \frac{1}{\delta^2} \ll_K 1$. The Bohr-to-gAP proposition (Lecture 11) implies the existence of a proper gAP inside this Bohr set, of dimension $\ll_K 1$ and size $\gg_K N$.

STEP 3: *Pull back the above to a large, low-dimensional proper gAP inside $2A - 2A$.*

Since $A' \simeq_8 B$, we have $2A' - 2A' \simeq_2 2B - 2B$. Now 2-isomorphisms preserve proper gAPs, whence $2A' - 2A'$ contains a proper gAP of dimension $\ll_K 1$ and size $\gg_K N$. Observe that $2A - 2A \supseteq 2A' - 2A'$, and recall from Step 1 that $N \asymp_K |A|$. It follows that $2A - 2A$ contains a proper gAP of dimension $\ll_K 1$ and size $\gg_K |A|$.

STEP 4: *Produce a small, low-dimensional gAP containing A.*

By Ruzsa's Reduction lemma (Lecture 8), since we were able to find a proper gAP of dimension $\ll_K 1$ and size $\gg_K |A|$ inside $2A - 2A$, there must exist a gAP containing A of size $\ll_K |A|$ and dimension $\ll_K 1$.

This concludes the proof. □

All that remains for us to do is to prove the Ruzsa Modelling Lemma. The strategy of the proof is as follows. First, there is an obvious m-isomomorphism of A into $\mathbb{Z}_p$ for huge primes (reduction mod p); we will abuse notation and refer to the image of A inside $\mathbb{Z}_p$ as A as well. It now suffices to find an m-isomorphism from a large piece of A into a small finite field $\mathbb{Z}_N$. A first attempt at this is to find the most popular interval $I \subseteq \mathbb{Z}_p$ of length p/m, and reduce all the elements of $I \cap A$ modulo N. This gives an m-homomorphism from A to $\mathbb{Z}_N$, but it doesn't give an m-isomorphism. Ruzsa's ingenious idea is to tweak the above construction to produce a huge family of m-homomorphisms from large subsets of A to $\mathbb{Z}_N$. He then employs a counting argument to show that some of these homomorphisms *must* be m-isomorphisms.

Proof of Ruzsa Modelling Lemma. Recall that we are given a finite set $A \subseteq \mathbb{Z}$, a positive integer m, and a prime $N > 2|mA - mA|$. We wish to find a large chunk of A which m-isomorphically embeds into $\mathbb{Z}_N$. We begin by considering the most natural candidate for an m-homomorphic embedding of a large chunk of A. This attempt won't work, unfortunately, but it will serve as a foundation on which to build an isomorphism.

STEP 1: *m-isomorphically embed A into a finite (but potentially large) group $\mathbb{Z}_p$.*

Consider the map

$$\begin{aligned} \phi : A &\to \mathbb{Z}_p \\ a &\mapsto a \pmod p \end{aligned}$$

This is a genuine group homomorphism, and hence an m-homomorphism. So long as p is large enough, it is also injective, thus giving an m-homomorphic embedding of A into $\mathbb{Z}_p$. Moreover, if p is really huge, ϕ is an m-isomorphism. To see this, define the *diameter* of A to be $\operatorname{diam} A := (\max A - \min A)$, and pick any prime

$$p > m \cdot \operatorname{diam} A.$$

It is clear that ϕ is a bijective m-homomorphism from A onto $\phi(A)$; to see that it's an m-isomorphism, it suffices to verify that ϕ^{-1} is an m-homomorphism from $\phi(A)$ onto A. To this end, suppose

$$\phi(a_1) + \phi(a_2) + \cdots + \phi(a_m) = \phi(a_1') + \phi(a_2') + \cdots + \phi(a_m').$$

This means

$$a_1 + a_2 + \cdots + a_m \equiv a_1' + a_2' + \cdots + a_m' \pmod p,$$

or in other words

$$(a_1 - a_1') + \cdots + (a_m - a_m') \equiv 0 \pmod p. \tag{3}$$

Now each difference on the left hand side is smaller than $\operatorname{diam} A$, whence

$$\left|(a_1 - a_1') + \cdots + (a_m - a_m')\right| \leq m \cdot \operatorname{diam} A < p.$$

Combining this with (3) shows that

$$(a_1 - a_1') + \cdots + (a_m - a_m') = 0,$$

whence ϕ^{-1} is an m-homomorphism from $\phi(A)$ onto A. Thus, ϕ is an m-isomorphism.

STEP 2: *Construct an m-homomorphism from a large subset of $\mathbb{Z}_p$ to $\mathbb{Z}_N$.*

Since we can pick an arbitrarily large value of p in Step 1, we may assume that $p > N$. The most natural map from $\mathbb{Z}_p$ to $\mathbb{Z}_N$ is the projection $\pi : \mathbb{Z}_p \to \mathbb{Z}_N$, where $\pi(x) := x \pmod N$. An immediate problem is that π isn't well-defined!

Exercise 4. *Show (by example) that the projection $\pi : \mathbb{Z}_7 \to \mathbb{Z}_3$ as described above is not well-defined.* [Hint: the elements of $\mathbb{Z}_7$ aren't integers!]

Fortunately, there's a cheap fix: project via the integers. First, given $x \in \mathbb{Z}_p$ there exists a unique $x_p \in \{0, 1, 2, \ldots, p-1\}$ such that $x \equiv x_p \pmod p$. Now define the projection π by

$$\begin{aligned} \pi : \mathbb{Z}_p &\to \mathbb{Z}_N \\ x &\mapsto x_p \pmod N \end{aligned}$$

We now have a well-defined projection map, but it doesn't do what we want:

Exercise 5. *Show (by example) that π might not be an m-homomorphism.* [Hint: find an example in which it's not a 2-homomorphism.]

This dashes our immediate hopes of finding an m-homomorphism from $\mathbb{Z}_p$ to $\mathbb{Z}_N$. Still, we can use π to construct an m-homomorphism from a decent-sized chunk of $\mathbb{Z}_p$ to $\mathbb{Z}_N$. Pick any interval $I \subseteq \{0, 1, \ldots, p-1\}$ of length $\lfloor p/m \rfloor$; we may view I as a subset of $\mathbb{Z}_p$.

Exercise 6. *Prove that π is an m-homomorphism from I to $\mathbb{Z}_N$.*

STEP 3: *Construct an m-homomorphism from a large subset $A' \subseteq A$ to $\mathbb{Z}_N$.*

Partition $\{0, 1, \ldots, p-1\}$ into m intervals of length $\lfloor p/m \rfloor$ (one of the intervals is allowed to be 'incomplete', i.e. shorter). Let I be the interval among these which is most popular with respect to $\phi(A)$, i.e. such that $|I \cap \phi(A)|$ is maximal. Now set

$$A' := \phi^{-1}\Big(I \cap \phi(A)\Big) = \{a \in A : \phi(a) \in I\}.$$

Applying the pigeonhole principle and recalling that ϕ is a bijection, we see that

$$|A'| = |I \cap \phi(A)| \geq \frac{|\phi(A)|}{m} = \frac{|A|}{m}.$$

Since the composition of two m-homomorphisms is an m-homomorphism, we see that

$$\pi \circ \phi : A' \to \mathbb{Z}_N \quad \text{is an } m\text{-homomorphism.}$$

However, this isn't an m-isomorphism; π might not even be injective! Our goal is now to tweak the above construction to obtain a map which is an m-isomorphism from a large subset of A into $\mathbb{Z}_N$.

STEP 4: *Construct a large family of m-homomorphisms from large subsets of A to $\mathbb{Z}_N$.*

For any $\lambda \in \mathbb{Z}_p^\times$, consider the map

$$\begin{aligned} \phi_\lambda : A &\to \mathbb{Z}_p \\ a &\mapsto \lambda a \pmod p \end{aligned}$$

Exercise 7. *Prove that for each $\lambda \in \mathbb{Z}_p^\times$, there exists $A'_\lambda \subseteq A$ with $|A'_\lambda| \geq \frac{|A|}{m}$ such that $\pi \circ \phi_\lambda$ is an m-homomorphism from A'_λ to $\mathbb{Z}_N$.*

We thus have a large family of m-homomorphisms from large subsets of A to $\mathbb{Z}_N$. I claim that one of these is an m-isomorphic embedding of A into $\mathbb{Z}_N$, so long as N is large enough.

STEP 5: *Count the number of non-isomorphisms.*

Let's call $\lambda \in \mathbb{Z}_p^\times$ *good* if $\pi \circ \phi_\lambda$ is an m-isomorphism from A'_λ to $\mathbb{Z}_N$; otherwise, if it is merely an m-homomorphism, we say λ is *bad*. I claim there are at most $\frac{p}{N}|mA - mA|$ bad λ's; it follows that so long as $N > 2|mA - mA|$, there exists a good λ, and the theorem is proved!

Our plan of attack is to show that any bad λ must satisfy a congruence modulo p; then we will trivially count the number solutions to this congruence, thus giving a bound on the number of bad λ's. If λ is bad, there must exist $a_i, a'_i \in A'_\lambda$ such that

$$a_1 + a_2 + \cdots + a_m \neq a'_1 + a'_2 + \cdots + a'_m$$

but

$$\pi \circ \phi_\lambda(a_1) + \cdots + \pi \circ \phi_\lambda(a_m) = \pi \circ \phi_\lambda(a'_1) + \cdots + \pi \circ \phi_\lambda(a'_m). \tag{4}$$

Set $b_i := \phi_\lambda(a_i)$ and $b'_i := \phi_\lambda(a'_i)$, with the understanding that $b_i, b'_i \in \{0, 1, 2, \ldots, p-1\}$. Then (4) can be rewritten

$$b_1 + \cdots + b_m \equiv b'_1 + \cdots + b'_m \pmod{N}.$$

Since ϕ_λ is an isomorphism and λ is bad, we know that

$$b_1 + \cdots + b_m \neq b'_1 + \cdots + b'_m.$$

Without loss of generality, we may assume that the left hand side is larger than the right; it follows from the above two displays that

$$(b_1 + \cdots + b_m) - (b'_1 + \cdots + b'_m) = \ell N \tag{5}$$

for some integer $\ell \geq 1$.

Exercise 8. *Prove that $\ell \leq \frac{p}{N}$.* [Hint: where do the b_i's live?]

It follows from (5) that

$$(b_1 + \cdots + b_m) - (b'_1 + \cdots + b'_m) \equiv \ell N \pmod{p}$$

whence (by the definition of b_i and b'_i) we deduce

$$\lambda^{-1} \equiv \ell^{-1} N^{-1} \Big((a_1 + \cdots + a_m) - (a'_1 + \cdots + a'_m) \Big) \pmod{p}.$$

How many possible values can the right hand side take? From above, there are at most p/N choices of ℓ. The quantity $(a_1 + \cdots + a_m) - (a'_1 + \cdots + a'_m)$ takes at most $|mA - mA|$ distinct values modulo p. Thus, there are at most $\frac{p}{N}|mA - mA|$ bad λ's, as claimed at the start of this step; the theorem follows. □

WWW.MATH.TORONTO.EDU/LGOLDMAK/APM461/
E-mail address: lgoldmak@math.toronto.edu