A mini course on additive combinatorics

First draft. Dated Oct 24th, 2007

These are notes from a mini course on additive combinatorics given in Princeton University on August 23-24, 2007. The lectures were **Boaz Barak** (Princeton University), **Luca Trevisan** (University of California at Berkeley) and **Avi Wigderson** (Institute for Advanced Study, Princeton). The notes were taken by Aditya Bhaskara, Arnab Bhattacharyya, Moritz Hardt, Cathy Lennon, Kevin Matulef, Rajsekar Manokaran, Indraneel Mukherjee, Wolfgang Mulzer, Aaron Roth, Shubhangi Saraf, David Steurer, and Aravindan Vijayaraghavan.

The lectures were also videotaped, and the tapes appear on the course's website at `http://www.cs.princeton.edu/theory/index.php/Main/AdditiveCombinatoricsMinicourse`

The course organizers were Boaz Barak, Moses Charikar and Mitra Kelly.

Contents

2

Szemeredi's Regularity Lemma, and Szemeredi's Theorem for k=3
Luca Trevisan

Scribe(s): Kevin Matulef

In this lecture we give a sketch of Szemeredi's theorem for k=3. The proof consists of four steps. The first step, the Regularity Lemma, will be proven in a later lecture. In this lecture we explain how each subsequent step follows from the the previous one. The steps are:

1. The Regularity Lemma. Roughly, this says that every graph has a constant size approximate representation (the size of the representation depends only on the quality of the approximation).

2. The Triangle Removal Lemma. This says that if a graph has $o(n^3)$ triangles, then it is possible to remove $o(n^2)$ edges and break all the triangles.

3. Szemeredi's Theorem for k=3 for groups.

4. Szemeredi's Theorem for k=3 for the integers.

2.1 How (3) implies (4).

For convenience, we restate the two versions of Szemeredi's Theorem for $k = 3$ here:

Theorem 2.1.1 (Szemeredi's Theorem for $k = 3$ for Groups)**.** *$\forall\delta$, there exists $n(\delta)$ st $\forall N \geq n(\delta)$ and all subsets $A \subseteq \mathbb{Z}_N$ where $|A| \geq \delta N$, A contains a length-3 arithmetic progression (i.e. three points $a, b, c \in \mathbb{Z}_N$ such that $b - a \equiv c - b \bmod N$).*

Theorem 2.1.2 (Szemeredi's Theorem for $k = 3$ for the Integers)**.** *$\forall\delta$, there exists $n(\delta)$ st $\forall N \geq n(\delta)$ and all subsets $A \subseteq \{1, ..., N\}$ where $|A| \geq \delta N$, A contains a length-3 arithmetic progression (i.e. three points $a, b, c \in \{1, ..., N\}$ such that $b - a = c - b$).*

To see that Theorem 2.1.1 implies Theorem 2.1.2, fix δ and let $A \subseteq \{1, ..., N\}, |A| \geq \delta N$. Think of A as a subset of $\mathbb{Z}_{2N}$, so that $|A| \geq \frac{\delta}{2}|Z_{2N}|$. If N is large enough, Theorem 2.1.1 implies that there exists $a, b, c \in A$ s.t. $b - a \equiv c - b \pmod{2N}$. But since a, b, c are all in the range $\{1, ..., N\}$, then $b - a$ and $c - b$ are both in the range $\{1 - N, ..., N - 1\}$. So the only way that they can be equal modulo $2N$ is if in fact $b - a = c - b$ without the mod. This proves Theorem 2.1.2, contingent on Theorem 2.1.1.

2.2 How (2) implies (3).

We begin with a formal definition of the triangle removal lemma:

Lemma 2.2.1 (Triangle Removal Lemma)**.** *For all δ, there exists $\epsilon = \epsilon(\delta)$ st. for every graph G on n vertices, at least one of the following is true:*

1. *G can be made triangle-free by removing $< \delta n^2$ edges.*

2. *G has $\geq \epsilon n^3$ triangles.*

In this section we see how Lemma 2.2.1 implies Theorem 2.1.1.

To prove Theorem 2.1.1, start with a group H, and a subset $A \subseteq H$. Then construct a graph on $n = 3|H|$ vertices by making three vertex sets, call them X, Y, Z, each with $|H|$ vertices labeled according to H. Connect these vertices as follows:

- Connect a vertex $x \in X$ and $y \in Y$ if $\exists a \in A$ st $y = x + a$.

- Connect a vertex $y \in Y$ and $z \in Z$ if $\exists c \in A$ st $z = y + c$.

- Connect a vertex $x \in X$ and $z \in Z$ if $\exists b \in A$ st $z = x + b + b$.

Observe that for every choice of $x \in H$ and $a \in A$, the graph has triangles of the form $(x, x + a, x + a + a)$. These triangles are all edge disjoint, so the graph has at least $|H| \times |A|$ edge disjoint triangles. Since $|A| \geq \delta|H|$ and $|H| = n/3$, this means that to make the graph triangle free we must remove at least $|H| \times |A| = \delta n^2/9$ edges. Using this observation and Lemma 2.2.1 with $\epsilon = \epsilon(\delta/9)$, we conclude that the graph has at least $\epsilon n^3 = 9\epsilon|H|^3$ triangles.

There is a correspondence between triangles in the graph and length-3 arithmetic progression. If (x, y, z) is a triangle in the graph, then there exists $a, b, c \in A$ st $y = x+a$, $z = y+c$, and $z = x+b+b$. Putting these equations together and rearranging, we get $a - b = b - c$. Conversely, if (a, b, c) form an arithmetic progression in A, then it easy to see that for any x, the tuple $(x, x + a, x + b + b)$ forms a triangle in the graph. In fact, for each $x \in H$, there is a bijection between arithmetic progressions (a, b, c) and triangles of the form (x, y, z). Thus, the total number of triangles in the graph is $|H|\times$ (# of length-3 progressions in A).

Since the total number of triangles in the graph is at least $\epsilon n^3 = 9\epsilon|H|^3$, the total number of length-3 progression in A is at least $9\epsilon|H|^2$. This is counting the "trivial progressions" where $a = b = c$, but even when we exclude those progression we are left with at least $9\epsilon|H|^2 - \delta|H|$ length-3 progressions in A. This is positive when $|H|$ is large enough. Thus, Theorem 2.1.1 follows from Lemma 2.2.1.

2.3 How (1) implies (2).

We begin with some definitions:

Definition 2.3.1. For disjoint vertex sets A, B, the *density* between A and B is

$$d(A,B) = \frac{\#\text{ edges between } A \text{ and } B}{|A||B|}$$

Definition 2.3.2. Two disjoint vertex sets are *ϵ-regular* if $\forall S \subseteq A$ where $|S| \geq \epsilon|A|$, and $\forall T \subseteq B$ where $|T| \geq \epsilon|B|$, it holds that

$$|d(S,T) - d(A,B)| \leq \epsilon$$

Informally, a bipartite graph is *ϵ-regular* if its edges are dispersed like a random graph's.

Lemma 2.3.3 (Szemeredi's Regularity Lemma)**.** *$\forall \epsilon, t$, $\exists k = k(\epsilon,t), k \geq t$ st for every graph $G = (V,E)$ there exists a partition of G into $(V_1, ..., V_k)$ where $|V_1| = |V_2| = ... = |V_{k-1}| \geq |V_k|$ and at least $(1-\epsilon)\binom{k}{2}$ pairs (V_i, V_j) are ϵ-regular.*

Informally, Szemeredi's regularity lemma says that all graphs are mostly composed of random-looking bipartite graphs.

We would like to show that the Regularity Lemma (lemma 2.3.3) implies the Triangle Removal Lemma (lemma 2.2.1). To see this, start with an arbitrary graph G. The Regularity Lemma says we can find a $\frac{\delta}{10}$-regular partition with $t = \frac{10}{\delta}$ into $k = k(\frac{\delta}{10}, \frac{10}{\delta})$ blocks. Using this partition, we define a reduced graph G' as follows:

- Remove all edges between non-regular pairs.

  Since there are at most $\frac{\delta}{10}\binom{k}{2}$ non-regular pairs, and at most $(\frac{n}{k})^2$ edges between each pair, this step removes at most $\frac{\delta}{10}n^2$ edges.

- Remove all edges inside blocks.

  Since there are k blocks, and each block contains at most $\binom{n/k}{2}$ edges, this step removes at most $\frac{n^2}{k} \leq \frac{\delta}{10}n^2$ edges.

- Remove all edges between pairs of density less than $\frac{\delta}{2}$.

  There are at most $\frac{\delta}{2}(\frac{n}{k})^2$ edges between a pair of density less than $\frac{\delta}{2}$, and at most $\binom{k}{2}$ total such pairs, so this step removes at most $\frac{\delta}{2}n^2$ edges.

In total the reduced graph G' is at most $\frac{\delta}{10} + \frac{\delta}{10} + \frac{\delta}{2} < \delta$ far from the original graph G. Thus if G' contains no triangle, the first condition of the Triangle Removal Lemma is satisfied (since all the triangles can be removed from G by breaking $\leq \delta n^2$ edges). For the remainder of the proof then, we assume that G' does contain a triangle, and we wish to show that the second condition of the Triangle Removal Lemma is satisfied (i.e. there must by $\geq \epsilon n^3$ triangles in G).

If G' does contain a triangle, it must go between three different blocks, call them A,B, and C. Let $m = n/k$ be the size of the blocks. Since G' does not contain any edges between low-density pairs of blocks, we know that if there is an edge between A and B, then in fact there must be many edges.

More quantitatively, we claim that at most $m/4$ vertices in A can have $\leq \frac{\delta}{4}m$ neighbors in B. Otherwise, there would be a subset of A, call it A', where $|A'| \geq |A|/4$, such that $d(A', B) \leq \frac{\delta}{4}$. This violates the conditions of the Regularity Lemma, since $|d(A,B) - d(A',B)| \geq \frac{\delta}{4} \geq \frac{\delta}{10}$. Likewise, the same argument shows that at most $m/4$ vertices in A can have $\leq \frac{\delta}{4}m$ neighbors in C.

Since at most $m/4$ vertices in A can have $\leq \frac{\delta}{4}m$ neighbors in B, and at most $m/4$ vertices in A can have $\leq \frac{\delta}{4}m$ neighbors in C, there must be at least $m/2$ vertices in A that have both $\geq \frac{\delta}{4}m$ neighbors in B *and* $\geq \frac{\delta}{4}m$ neighbors in C.

Consider a single such vertex from A. Let S be its neighbor set in B and T be its neighbor set in C, where $|S| \geq \frac{\delta}{4}m$ and $|T| \geq \frac{\delta}{4}m$. How many edges go between S and T? Since $d(B,C) \geq \frac{\delta}{2}$, and B and C are $\frac{\delta}{10}$-regular, the number of edges going between S and T must be at least $(\frac{\delta}{2} - \frac{\delta}{10})|S||T| \geq \frac{\delta}{4}|S||T| \geq \frac{\delta^3}{64}m^2$. Thus a single vertex from A with high degree in both B and C accounts for at least $\frac{\delta^3}{64}m^2$ triangles in G'. Since there are at least $m/2$ such vertices from A, the total number of triangle must be at least $\frac{\delta^3}{128}m^3 = \frac{\delta^3}{128k^3}n^3$, which is what we wanted to show.

2.4 Brief Sketch of (1)

In this section we provide a brief sketch of the Regularity Lemma, the only remaining piece of the proof of Szemeredi's Theorem for $k = 3$. A more detailed proof will be given in a later lecture.

To prove the Regularity Lemma, we start with an arbitrary partition of the graph G into t subsets. If this partition is ϵ-regular, then we are done. Otherwise, we iteratively refine the partition until we get one that is ϵ-regular. We do this as follows:

If a partition is not ϵ-regular, that means there exist at least $\epsilon\binom{t}{2}$ non-regular pairs V_i, V_j. Let S_i and T_j be witnesses to the non-regularity of V_i, V_j; in other words, $S_i \subseteq V_i, |S_i| \geq \epsilon|V_i|$ and $T_j \subseteq V_j, |T_j| \geq \epsilon|V_j|$ where $|d(S_i, T_j) - d(V_i, V_j)| > \epsilon$. We refine the partition by subdividing all the subsets into intersections of $V_i, V_i \backslash S_i, V_j, V_j \backslash T_j$ over all non-regular pairs (i, j) (we actually even further refine the partition so that the resulting sets have equal size).

To analyze this process, we examine the quantity $\sum_{i,j} \frac{|V_i||V_j|}{n^2} d^2(V_i, V_j)$. This can be thought of as the variance of the densities between the sets containing two randomly chosen vertices. One can show that this quantity always increases upon refinement, and each iteration of the above process increases it by at least poly(ϵ). Since the quantity is at most 1, the refinement process must terminate after poly$(\frac{1}{\epsilon})$ iterations.

3

The Sum Product Theorem and its Applications

Avi Wigderson

Scribe(s): Aaron Roth and Cathy Lennon

Summary: This lecture contains the statement of the sum-product theorem and a broad survey of its applications to many areas of mathematics. Finally it sketches a proof; a full proof of the sum-product theorem is contained in the next lecture.

3.1 Preliminaries and Statement of the Sum-Product Theorem

Let $\mathbb{F}$ be a field, and let $A \subseteq \mathbb{F}$ be an arbitrary subset. In particular, A is not necessarily a subfield, and so is not necessarily closed under either field operation. We may therefore wish to study how close A is to a closed set under each field operation. The following definitions capture this:

Definition 3.1.1. For $A \subseteq \mathbb{F}$, the *sumset* $A + A$ is:

$$A + A = \{a + a' : a, a' \in A\}$$

Definition 3.1.2. For $A \subseteq \mathbb{F}$, the *product set* $A \times A$ is:

$$A \times A = \{a \times a' : a, a' \in A\}$$

If A is a subfield of $\mathbb{F}$, then $|A + A| = |A \times A| = |A|$. If A is not closed under some operation $\cdot$, then $|A \cdot A| > |A|$. In principle, $|A \cdot A|$ can be as large as $\Omega(|A|^2)$, but we may be interested in sets A such that $|A + A|$ or $|A \times A|$ is small.

3.1.1 Over the Reals

Example 3.1.3. For $\mathbb{F} = \mathbb{R}$ and $A = \{1, 2, 3, \ldots, k\}$ we have:

- $|A + A| < 2|A|$ (pairwise addition of elements in A yields integers in $\{2, ..., 2k\}$
- $|A \times A| = \Omega(|A|^2)$

For $\mathbb{F} = \mathbb{R}$ and $B = \{1, 2, 4, \ldots, 2^k\}$ we have:

- $|B \times B| \leq 2|B|$ (since this is $\{2^a : a \in |A + A|\}$)
- $|B + B| = \Omega(|B|^2)$ (Like the set of all length-k binary strings with norm 2)

In these examples, either the sumset or the product set is small, but not both. For $\mathbb{F} = \mathbb{R}$, does there exist an $A \subset \mathbb{R}$ for which $\max\{|A + A|, |A \times A|\}$ is 'small'? The sum-product theorem over the reals tells us that there is not:

Theorem 3.1.4 (Sum Product Theorem for $\mathbb{F} = \mathbb{R}$ [ES83a])**.** *For $\mathbb{F} = \mathbb{R}$, $\exists \epsilon > 0$ such that $\forall A \subset \mathbb{F}$ either:*

1. $|A + A| \geq |A|^{1+\epsilon}$

2. $|A \times A| \geq |A|^{1+\epsilon}$

The best ϵ known in the above theorem is $4/3$, but it is conjectured that it holds for $\epsilon = 2 - o(1)$.

3.1.2 Over a Finite Field

Suppose $\mathbb{F}$ is a finite field. Can Theorem 3.1.4 still hold? Not quite as stated:

- If $A \subset \mathbb{F}$ such that $|A| \geq C\mathbb{F}$ for some constant C, neither $|A + A|$ nor $|A \times A|$ can exceed $|A|$ by more than a constant factor. We must therefore have a condition that A is not 'too big.'
- If A is a subfield of $\mathbb{F}$ then it is closed under both field operations, and so $|A + A| = |A \times A| = |A|$. We must therefore have a condition that $\mathbb{F}$ contains on subfields.

The following version of the sum-product theorem holds for finite fields $\mathbb{F}_p$ for p prime (this restriction guarantees that there are no nontrivial subfields – variants of this theorem are possible for other fields with small subfields.)

Theorem 3.1.5 (Sum Product Theorem for $\mathbb{F} = \mathbb{F}_p$ [BT04][Kon03])**.** *For $\mathbb{F} = \mathbb{F}_p$ for p prime, $\exists \epsilon > 0$ such that $\forall A \subset \mathbb{F}$ with $|A| \leq |\mathbb{F}|^{.9}$, either:*

1. $|A + A| \geq |A|^{1+\epsilon}$

2. $|A \times A| \geq |A|^{1+\epsilon}$

The best known ϵ in the above theorem is .001, and it is known that ϵ cannot be larger than 3/2.

3.2 Applications

The sum-product theorem has wide ranging applications. In this section, we survey various fields of mathematics and give theorem statements that were state-of-the-art before application of the sum-product theorem, and then the corresponding improvement possible with the sum-product theorem.

3.2.1 Combinatorial Geometry

Consider the geometry of a plane $\mathbb{F}^2$. Let P be a set of points in $\mathbb{F}^2$ of cardinality n, and let L be a set of lines in $\mathbb{F}^2$ also of cardinality n. A natural quantity to study is

$$I = \{(l,p) : l \in L, p \in P \text{ such that point p lies on line l}\}$$

the set of incidences of points in P on lines in L. Using only simple combinatorial facts about lines and planes we may prove the following simple bound:

Theorem 3.2.1.

$$|I| \leq O(n^{3/2})$$

Proof. Let

$$\delta_{p,l} = \begin{cases} 1, & \text{if point } p \text{ lies on line } l; \\ 0, & \text{otherwise.} \end{cases}$$

Let $f(l) = \sum_{p \in P} \delta_{p,l}$ denote the number of points incident on line l and let $g(p) = \sum_{l \in L} \delta_{p,l}$ denote the number of lines incident on a point p. Node that $|I| = \sum_{p \in P} g(p) = \sum_{l \in L} f(l)$. then:

$$\begin{aligned} |I| &= \sum_{l \in L} f(l) \\ &= \sum_{l \in L} 1 \cdot f(l) \\ &\leq \sqrt{n} \cdot \sqrt{\sum_{l \in L} f(l)^2} \\ &= \sqrt{n} \cdot \sqrt{\sum_{l \in L} (\sum_{p \in P} \delta_{p,l})(\sum_{p \in P} \delta_{p,l})} \\ &= \sqrt{n} \cdot \sqrt{\sum_{p,p' \in P} \sum_{l \in L} \delta_{p,l}\delta_{p',l}} \\ &= \sqrt{n} \cdot \sqrt{\sum_{p \in P} g(p) + \sum_{p \neq p'} \sum_{l \in L} \delta_{p,l}\delta_{p',l}} \\ &\leq \sqrt{n} \cdot \sqrt{|I| + n^2} \end{aligned}$$

where the first inequality follows from Cauchy-Schwartz, and the second follows from the fact that there can be at most one line incident to any two distinct points, and there are fewer than n^2 pairs of points. Solving for $|I|$ we get the desired bound. □

For $\mathbb{F} = \mathbb{R}$ a tighter non-trivial bound is possible without the sum-product theorem:

Theorem 3.2.2 ([**?**],[Ele97])**.** *For* $\mathbb{F} = \mathbb{R}$ $|I| \leq O(n^{4/3})$

With the sum-product theorem, a big improvement is possible (this is a non-trivial consequence of a statistical-version of the sum-product theorem):

Theorem 3.2.3 ([BT04])**.** *For* $\mathbb{F} = \mathbb{F}_p$ *with* p *prime,* $|I| \leq n^{3/2-\epsilon}$

3.2.2 Analysis

Consider the following question which has many applications in analysis/PDE: what is the smallest area of a figure which contains a unit segment in every direction? [1]

In the Reals

Besicovitch proved:

Theorem 3.2.4 ([Bes63]). *One can construct $S \subseteq \mathbb{R}^2$ with arbitrarily small area such that S contains a unit segment in every direction.*

The construction is as follows: begin with a square of side length 2 with center O and add lines across the diagonals, partitioning it into four isosceles right triangles with hypotenuse of length 2 with common vertex O. Partition each hypotenuse into n equally sized segments, where the value of n depends on how small of an area is desired. For each endpoint, draw a line from O to the endpoint. This further partitions the square into $4n$ "elementary triangles". The resulting object looks like:

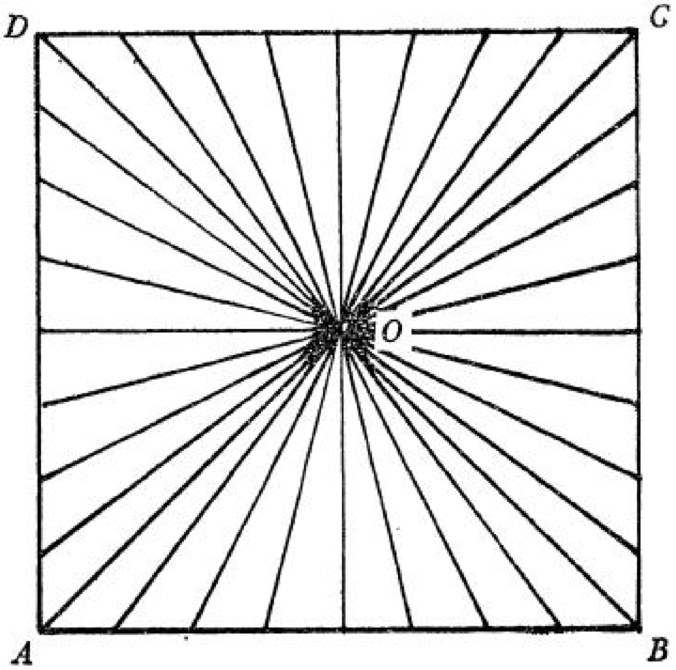

By constructing the figure like this, it is clear that considering each of the $4n$ elementary triangles and all possible segments with one endpoint at O and the other at a point in the base of the triangle, then the collection of all such segments for all triangles will have a range of 360° and each such segment will have a length of at least 1. Therefore,the collection of these $4n$ triangles contains a unit segment in all possible directions. Of course, the area of this square is quite large.

[1] Examples of such figures are a circle of radius 1 (area= $\pi/4 \approx .78$) and an equilateral triangle of height 1 (area= $\frac{1}{\sqrt{3}} \approx .58$). It was also (falsely) conjectured that the hypocycloid inscribed in a circle of diameter 3/2 was the figure of smallest area.

They key to his construction is following observation: arbitrary translations of the elementary triangles does not affect which unit segments are contained within the triangles. By translating them in such a way that they overlap, we can decrease the total area of the figure while maintaining the property of containing a unit segment in all directions.

We describe the translations for the bottom triangle, but the same procedure may be plied to the other three, resulting in a solution to the problem. For any integer $p \geq 2$ consider the sequence of triangles Δ_2, Δ_3,...,Δ_p, where each is a right isosceles triangle, and Δ_k has height k/p, base partitioned into 2^{k-2} segments, and a line joining each segment endpoint to the common vertex. Thus triangle Δ_k consists of 2^{k-2} elementary triangles. Each Δ_{k+1} can be constructed from Δ_k by bisecting each of the elementary triangles of Δ_k through their base, and scaling each so that it now has height $k+1$. This will create a series of overlapping triangles that can be translated to construct Δ_{k+1}. This process is known as "bisection and expansion". By starting with triangle Δ_2, and then using this technique repeatedly, after $p-2$ steps, we will have 2^{k-2} overlapping elementary triangles which can be translated to form Δ_p. Analysis of the process of bisection and expansion shows that each step increases the total area by at most $\frac{1}{p^2}$ and so at the final step the constructed object S_1 has area $\frac{2}{p}$. Choosing $p > 8/\epsilon$ will make the area of S_1 less than $\epsilon/4$. Doing this for each of the four original right triangle results in a figure of area at most ϵ. Since this figure consists of translations of the elementary triangles described above, it contains a unit segment in every direction.

One can also ask this question for $\mathbb{R}^d$ for $d > 2$ or for measures other than the Lebesque measure. A third variation is to consider this question for finite fields, and it is in this case that the sum-product theorem gives interesting results.

In Finite Fields

Call a set $S \subset (\mathbb{F}_p)^d$ a Kakeya set if S contains a line in all possible directions (for large p). By this we mean that for all $b \in \mathbb{F}_p^d \setminus \{0\}$ there is an $a \in \mathbb{F}_p^d$ such that $a + tb \in S$ for all $t \in \mathbb{F}_p$. Any particular line is of the form $a + tb$ and since $t \in \mathbb{F}_p$, each line contains p points. Define $B(d)$ to be the smallest r such that there exists a Kakeya set S with $|S| = \Omega(p^r)$. It has been conjectured that $B(d) = d$.

Before the sum-product theorem, the following was known:

Theorem 3.2.5. *(Trivial)* $B(d) \geq d/2$

Theorem 3.2.6 ([Wol99])**.** $B(d) \geq d/2 + 1$

However, one can use the sum-product theorem to show that:

Theorem 3.2.7 ([BT04])**.** $B(d) \geq d/2 + 1 + 10^{-10}$

We will show the trivial bound of $B(d) \geq d/2$. The proof relies on the following fact: if P is a collection of points in $\mathbb{F}_p^d$ and L is a collection of lines in $\mathbb{F}_p^d$, then

$$\{(p,l) \in P \times L : p \in L\}| \leq \min(|P|^{\frac{1}{2}}|L| + |P|, |P||L|^{\frac{1}{2}} + |L|)$$

To apply this, let P be any set of points in $\mathbb{F}_p^d$ which contains lines in every possible direction, and set L to be a set consisting of these lines. Since there are $\frac{p^d-1}{p-1}$ different directions, it follows that $|L| \geq \frac{p^d-1}{p-1}$. Since we assume that all of the points in each line lie in P, and since each line has p points, we have

$$p|L| = |\{(p,l) \in P \times L : p \in L\}| \leq |P||L|^{\frac{1}{2}} + |L|$$

$$\implies |P| \geq |L|^{\frac{1}{2}}(p-1) \geq (\frac{p^d-1}{p-1})^{\frac{1}{2}}(p-1) = ((p^d-1)(p-1))^{\frac{1}{2}} = (p^d(p-1-\frac{1}{p^{d-1}}+\frac{1}{p^d}))^{\frac{1}{2}} \geq p^{d/2}$$

So any Kakeya set S satisfies $|S| \geq p^{d/2}$ and hence $B(d) \geq d/2$.

3.2.3 Number Theory

For $\mathbb{F} = \mathbb{F}_p$ a finite field of prime order, let G be a multiplicative subgroup of $\mathbb{F}^*$. Define the fourier coefficient at a relative to G to be:

Definition 3.2.8 (Fourier coefficient at a)**.**

$$S(a,G) = \sum_{g \in G} \omega^{a \cdot g}$$

A natural quantity to study is the maximum fourier coefficient of any element in $\mathbb{F}^*$:

$$S(G) = \max_{a \in \mathbb{F}^*} |S(a,G)|$$

We may provide a trivial bound:

Theorem 3.2.9.

$$S(G) \leq |G|$$

Proof.

$$\begin{aligned} S(G) &= \max_{a \in \mathbb{F}^*} |\sum_{g \in G} \omega^{a \cdot g}| \\ &\leq \max_{a \in \mathbb{F}^*} \sum_{g \in G} |\omega^{a \cdot g}| \\ &= \sum_{g \in G} 1 \\ &= |G| \end{aligned}$$

□

We would instead like to show that $S(G) \leq |G|^{1-\epsilon}$.

Without the sum-product theorem, this was demonstrated for $|G| \geq p^{1/2}$ [Wol99], $|G| \geq p^{3/7}$ [HB96], and $|G| \geq p^{1/4}$ [KS99].

Using the sum-product theorem, more general results can be derived:

Theorem 3.2.10 ([BT04])**.** *For $\mathbb{F} = \mathbb{F}_p$ with p prime, and $|G| \geq p^\delta$, $S(G) \leq |G|^{1-\epsilon(\delta)}$*

Theorem 3.2.11 ([BK06])**.** *For $\mathbb{F} = \mathbb{F}_p$ with p prime, and $|G| \geq p^\delta$, $S(G^{k(\delta)}) \leq |G^k|^{1-\epsilon(\delta)}$*

Bourgain and Chung generalize these results to other fields.

3.2.4 Group Theory

Suppose that H is a finite group, and T is a set of generators for H. We can then study the structure of H with respect to T:

Definition 3.2.12 (Cayley Graph)**.** $\mathrm{Cay}(H;T)$, the Cayley Graph of H and T is the graph with vertex set $V = H$, and edge set $E = \{u, v : u \cdot v^{-1} \in T\}$. That is, the vertex set is the set of group elements, with edges corresponding to walks we can take among group elements by multiplying by the generators in T.

One interesting feature of a Cayley graph is its diameter $\mathrm{Diam}(H;T)$, the longest path between any two group elements in $\mathrm{Cay}(H;T)$. If $\mathrm{Cay}(H;T)$ is an expander graph, then the second largest eigenvalue over the matrix defining the markov process of a random walk over $\mathrm{Cay}(H;T)$, $\lambda(H;T) \leq 1 - \epsilon$, and so its diameter is $O(\log|H|)$.

For $H = \mathrm{SL}(2,p)$, the group of 2×2 invertible matrices over $\mathbb{F}_p$, Selberg, Lubotsky Phillips and Sarnak, and Margulis showed without the sum-product theorem a few sets of generators T for which $\mathrm{Cay}(H;T)$ is an expander [Sel65] [LS86] [Mar73]. Using the sum product theorem, Helfgott showed that for *all* T, the diameter of $\mathrm{Cay}(H;T)$ is $<$ polylog$(|H|)$ [Hel]. Also using the sum-product theorem, Borgain and Gamburd showed that for random T with $|T| = 2$, $\mathrm{Cay}(H;T)$ is an expander (with diameter $O(\log|H|)$) with high probability [BG06]. [BG06] also show that if $< T >$ is not cyclic in $H =\mathrm{SL}(2,\mathbb{Z})$, then $\mathrm{Cay}(H;T)$ is an expander.

3.2.5 Randomness Extractors and Dispersers

Randomness is a powerful tool used in many areas of computer science, and has practical applications because randomness appears to be prevalent in our physical experience of the world. However, theoretical applications often require uniform random bits, whereas when we observe randomness in nature, it is not of this clean form. We may have access to a random variable X with high entropy (in the sense that it would take many uniform random bits to sample from X), but that is far in statistical distance from the uniform distribution U_n – and perhaps the exact distribution of X is unknown. If S is a such a class of probability distributions over $\{0,1\}^n$, then $X \in S$ is often called a 'weak source of randomness'.

In order to make use of X in applications that require uniform random bits, we would like extractors and dispersers:

Definition 3.2.13 ((S,ϵ)-Disperser)**.** A function $f : \{0,1\}^n \to \{0,1\}^m$ for which all $X \in S$ satisfies:

$$|f(X)| \geq (1-\epsilon)2^m$$

is an (S,ϵ)-disperser. That is, $f(X)$ is a distribution with large support (although is not necessarily distributed close to uniform).

Definition 3.2.14 ((S,ϵ)-Extractor)**.** A function $f : \{0,1\}^n \to \{0,1\}^m$ for which all $X \in S$ satisfies:

$$\|f(x) - U_m\|_1 \leq \epsilon$$

is an (S,ϵ)-disperser. That is, $f(X)$ is ϵ-close to the uniform distribution (and in particular, must have large support).

The existence of extractors and dispersers is a Ramsey/Discrepancy theorem, but their explicit polynomial time construction is an important research area. Extractors and dispersers can be either randomized (seeded), or deterministic. It is easy to see that there cannot be deterministic constructions that work with all high entropy sources X, but we can avoid needing any uniform random bits if we make some assumption about the structure of S.

For example, let $S = L_k$, the set of affine subspaces of $\mathbb{F}_2^n$ of dimension $\geq k$. Say that f is optimal if $m = \Omega(k)$ and $\epsilon = 2^{-\Omega(k)}$ (information theoretic bounds tell us we cannot hope that $f(X)$ has higher entropy than X). Before the use of the sum-product theorem, it was known via the probabilistic method that there exist optimal affine extractors for all $k \geq 2\log n$. Explicit constructions for such extractors were known only for $k \geq n/2$. Using the sum-product theorem, better results are possible: [BW05] give an explicit construction for an affine disperser with $m = 1$ for all $k \geq \delta n$. [Bou07] gives an explicit construction for an optimal affine extractor for all $k \geq \delta n$. [GR05] give extractors for large finite fields of low dimension.

Alternately, we may consider $S = I_k = \{(X_1, X_2) : X_1, X_2 \in \{0,1\}^n$ independent, $H_\infty(X_i) \geq k\}$ where $H_\infty(X_i) \geq k$ implies that no element in X_i has probability greater than 2^{-k}. For simplicity, it is helpful to think about X_i as the uniform distribution over a support of k elements. Again, $f(X_1, X_2)$ cannot have higher entropy than (X_1, X_2), f is optimal if $m = \Omega(k)$ and $\epsilon = 2^{-\Omega(k)}$. Before the sum-product theorem, Erdos (using the probabilistic method) showed that there exists an optimal 2-source extractor for all $k \geq 2\log n$ in proving the existence of Ramsey Graphs[2]. [CG88] and [Vaz87] gave an explicit construction for an optimal 2-source extractor for all $k \geq n/2$. Using the sum-product theorem, (slightly) better results are possible: [Bou07] gives an explicit optimal 2-source extractor for $k \geq .4999n$. [BW05] give an explicit 2-source disperser for $m = 1$ and $k \geq \delta n$ (giving new constructions for bipartite Ramsey graphs), and [BW06] give an explicit 2-source disperser for $m = 1$ and $k \geq n^\delta$.

A Statistical Version of the Sum Product Theorem

The sum-product theorem tells us about the size of sets, and so is useful in constructing dispersers from a constant number of independent sources [3]. To construct extractors, however, which require not only that $f(A)$ have large support, but also that it be almost uniformly distributed, we would like a statistical analogue of the sum-product theorem. Intuitively, we would like a statement about the size of A in terms of its entropy. There are several measures we might consider: $H_0(A) = |\text{support}(A)|$ is simply the measure we have been using for the standard sum-product theorem. Shannon entropy $H_{\text{Shannon}} = \sum_{i \in A} -p_i \log p_i$ is one possibility. A stronger measure is $H_2(A) = -\log ||A||_2 \approx H_\infty(A) = \min_{i \in A} -\log(p_i)$. Note that we have: $H_2 \leq H_{Shannon} \leq H_0$.

We may rephrase our existing result over $\mathbb{F}_p$:

[2] A Ramsey graph has no clique and no independent set of size k. Note that a two-source disperser for I_k with $m = 1$ $f : \{0,1\}^{2n} \to \{0,1\}$ provides a construction for a bipartite Ramsey graph. Let the graph consist of two parts, each consisting of 2^n vertices, corresponding to the binary strings of length n. For two vertices x, y in different parts, let there be an edge (x, y) if $f(x, y) = 1$. In any two subsets S, T for $|S|, |T| \geq k$ in different parts, we are guaranteed by the definition of a disperser that $|f(S, T)| > 1 \Rightarrow f(S, T) = \{0, 1\}$ since we may consider S and T to be the uniform distribution over k elements. Therefore, between any two subsets $|S|, |T| \geq k$ in different parts, there are always edges, but never all of the possible edges, giving us a bipartite Ramsey graph.

[3] Let $f(x, y, z) = x \times y + z$. As a corollary of the sum-product theorem, we know that $|A \times A + A| \geq |A|^{1+\epsilon}$ for some constant ϵ. So long as $|A| \geq \mathbb{F}_p^\delta$ for some constant δ, by composing f with itself a constant number of times to obtain g, we have $g(A, \ldots, A) = \mathbb{F}_p$, and so we have a disperser.

Theorem 3.2.15 (Sum Product Theorem Over $\mathbb{F}_p$ [BT04] [Kon03])**.** *There exists an $\epsilon \geq 0$ such that for all $A \subseteq \mathbb{F}_p$ such that $H_0(A) \leq .9 \log p$, either:*

1. $H_0(A + A) > (1 + \epsilon)H_0(A)$ *or*
2. $H_0(A \times A) > (1 + \epsilon)H_0(A)$.

We would like an identical result in which we could replace H_0 with H_2 (or H_∞), but this is not possible. In particular, for every prime field $\mathbb{F}$, there is a distribution uniform over a subset of $\mathbb{F}$ of size 2^k (with $k < 0.9 \log p$) such that both $A + A$ and $A \times A$ put a constant probability over a set of size at most $O(2^k)$[4]. This gives $H_\infty(A) = k$, $H_\infty(A+A) \leq k + \log 1/c$ and $H_\infty(A \times A) \leq k + \log 1/c$ for some constant c, which violates our desired sum-product theorem. An analogous theorem is possible, however, if we replace sums and products with a convolution. Again, rephrasing our original theorem:

Theorem 3.2.16 ([BT04] [Kon03])**.** *There exists an $\epsilon \geq 0$ such that for all $A \subseteq \mathbb{F}_p$ such that $H_0(A) \leq .9 \log p$:*

$$H_0(A \times A + A) > (1 + \epsilon)H_0(A)$$

Now a statistical analogue is possible, as proven by Barak, Impagliazzo, and Wigderson:

Theorem 3.2.17 ([BW04])**.** *There exists an $\epsilon \geq 0$ such that for all $A \subseteq \mathbb{F}_p$ such that $H_2(A) \leq .9 \log p$:*

$$H_2(A \times A + A) > (1 + \epsilon)H_2(A)$$

Suppose that A, B, C are independent distributions over $\mathbb{F}_p$. Let the rate of distribution X be $r(X) = H_2(X)/\log p$, and let $r = \min(r(A), r(B), r(C))$. [BW04] show the following:

There exists a constant $\epsilon > 0$ such that:

- If $r \leq 0.9$, then $r(A, B, C) \geq (1 + \epsilon)r$
- If $r > 0.9$, then $r(A, B, C) = 1$.

Given this, and the statistical analogue of the sum-product theorem, we may construct extractors in the way that we were able to construct dispersers. We define:

- $f^1(A_1, A_2, A_3) = A_1 \times A_2 + A_3$
- $f^{t+1}(A_1, A_2, \ldots, A_{3^{t+1}}) = f^1(f^t(A_1, \ldots, A_{3^t}), f^t(A_{3^t+1}, \ldots, A_{2 \cdot 3^t}), f^t(A_{2 \cdot 3^t+1}, \ldots, A_{3^{t+1}}))$

By composing f^1 with itself t times, we are able to convert 3^t sources with entropy $H_2(A)$ to a single source with entropy at least $(1 + \epsilon)^t H_2(A)$. Using this construction, [BW04] give optimal explicit constructions for extractors over the set

$$S = \{(A_1, \ldots, A_c) : A_i \text{ independent over } \{0,1\}^n \text{ with } H_2(A_i) > k\}$$

for $k = \delta n$ and $c = \text{poly}(1/\delta)$ Note that for $\delta > 0$ a constant, this gives an extractor given only a *constant* number of independent input sources. Without using the sum-product theorem, Rao gets a stronger result, giving an explicit construction for optimal extractors over S with $k = n^\delta$ and $c = \text{poly}(1/\delta)$ [Rao06].

[4] Let A put probability 1/2 on an arithmetic progression, and probability 1/2 on a geometric progression. Then, as we saw above, $A + A$ will have probability 1/2 on a set of size $2|A|$, and $A \times A$ will have probability 1/2 on a set of size $2|A|$, which is what we need.

Definition 3.2.18 (Condenser)**.** For X a distribution on $\{0,1\}^n$, and $r(X) = H_2(X)/n$, then $f_c : \{0,1\}^n \to (\{0,1\}^m)^c$ is a condenser if there exists a constant $\epsilon > 0$ such that for all X with $r(X) \leq .9$:

$$\exists c > 0 : r(f_c(X)) \geq (1+\epsilon)r(X)$$

Intuitively, r is a measure of how close to uniform X is. f_c takes a weak random variable, creates a random variable over a possibly smaller space that is closer to uniform.

Using the sum product theorem, [BW05] give a condenser that iteratively boosts $r = \delta$ to $r = 0.9$ with $c = \text{poly}(1/\delta)$ for $m = \Omega(n)$. For constant δ therefore, c is constant.

3.3 Proof Sketch of the Sum-Product Theorem

Recall the statement of the sum-product theorem:

Theorem 3.3.1 ([BT04])**.** *Let $F = \mathbb{F}_p$, then for all $\delta < .9$ there exists an $\epsilon > 0$ such that for any $A \subset F$ satisfying $|A| = |F|^\delta$, either $|A + A| \geq |A|^{1+\epsilon}$ or $|A \times A| \geq |A|^{1+\epsilon}$.*

This theorem will be a consequence of the following two lemmas.

Lemma 3.3.2. *There exists a rational expression R_0 such that for all A, $|R_0(A)| > |A|^{1+\delta}$.*

where a rational expression $R(A)$ is a rational function in A, for example $R(A) = (A + A - A \times A)/(A \times A \times A)$.

Proof. Define the rational expression $R_0(A) := (A' - A')/(A' - A') = A''$, where $A' := (A - A)/(A - A)$. Also, define δ', δ'' to be the values such that $|A'| = |F|^{\delta'}$ and $|A''| = |F|^{\delta''}$. Then we claim that R_0 is the desired rational expression. This will follow from the following claim (to be proven subsequently): if $\delta \in (1/(k+1), 1/k)$ then $\delta' > \frac{1}{k}$. We know that there is some k such that δ is in the open interval $(1/(k+1), 1/k)$, (note here the strict inclusion since $|F|^{\frac{1}{k}}, |F|^{\frac{1}{k+1}} \notin \mathbb{N}$ but $|A| \in \mathbb{N}$). Applying the lemma twice gives that $\delta'' > 1/(k-1)$ and then

$$\delta(1+\delta) < \frac{1}{k} * \frac{k+1}{k} = \frac{k+1}{k^2} < \frac{k+1}{k^2-1} = \frac{1}{k-1} < \delta''$$

Putting this together gives

$$|F|^{\delta(1+\delta)} < |F|^{\delta''} = |R_0(A)|$$
$$\implies |R_0(A)| > (|F|^\delta)^{1+\delta} = |A|^{1+\delta}$$

which is what we wanted.

It is left to prove the claim: assume otherwise, ie that $\delta' < \frac{1}{k}$. Construct a sequence $s_0 = 1, s_1, ..., s_k \in F$ such that each s_j satisfies $s_j \notin s_0A' + s_1A' + ...s_{j-1}A'$. It is always possible to find such an s_j when $j \leq k$ because otherwise this would imply that $s_0A' + ...s_{j-1}A' = F \implies |A'| > |F|^{\frac{1}{j-1}}$, contradicting our assumption that $\delta' < \frac{1}{k}$.

Next, define a function $g : A^{k+1} \to F$ by $g(x_0, ..., x_k) = \sum s_i x_i$. By our choice of k, we have $|A|^{k+1} > |F|$, so g cannot be injective, and so there exists $x \neq y$ such that $\sum s_i x_i = g(x) = g(y) = \sum s_i y_i$. Choose j to be the largest index where x_j and y_j differ. Then $\sum s_i x_i = \sum s_i y_i \implies \sum s_i(x_i - y_i) = 0 \implies \sum_{i \leq j} s_i(x_i - y_i) = 0 \implies s_j = \sum_{i<j} s_i(x_i - y_i)/(y_j - x_j)$. Each $(x_i - y_i)/(y_j - x_j) \in A'$ and so $s_j = \sum_{i<j} s_i(x_i - y_i)(y_j - x_j) \in s_0A' + s_1A' + ...s_{j-1}A'$, a contradiction. □

Lemma 3.3.3. *If $|A+A| < |A|^{1+\epsilon}$ and $|A \times A| < |A|^{1+\epsilon}$ then for all R there exists a $c = c(R)$ such that there is some $B \subseteq A$ with $|B| > |A|^{1-c\epsilon}$ and $|R(B)| < B^{1+c\epsilon}$.*

We list only the ingredients for proving Lemma 2. A full proof will be given in the next lecture. Let G be an abelian group, $A \subseteq G$, $\epsilon > 0$ arbitrary, then the following theorems hold:

Theorem 3.3.4 ([Rao06])**.** $|A+A| < |A|^{1+\epsilon} \implies |A-A| < |A|^{1+2\epsilon}$

Theorem 3.3.5 ([Plu69],[Rao06])**.** $|A+A| < |A|^{1+\epsilon} \implies |A+kA| < |A|^{1+k\epsilon}$

Theorem 3.3.6 ([BS94], [Gow98])**.** $||A+A||^{-1} < |A|^{1+\epsilon} \implies \exists A' \subseteq A, |A'| > |A|^{1-5\epsilon}$ *but* $|A'+A'| < |A'|^{1+5\epsilon}$.

[The theorem now follows from these lemmas: assume for contradiction that the theorem is false, ie that there is some $\delta < .9$ such that for all $\epsilon > 0$ there is some set A, $|A| = |F|^{\delta}$ with both $|A+A| \leq |A|^{1+\epsilon}$ and $|A \times A| \leq |A|^{1+\epsilon}$. Let R_0 be as in LEMMA 1. Now A satisfies the hypotheses of LEMMA 2 and so considering $R = R_0$, there is a c, B with $|B| > |A|^{1-c\epsilon}$ such that $|B|^{1+\delta} < |R_0(B)| < |B|^{1+c\epsilon}$]

3.4 Conclusion

What we should take away from this survey lecture is that the sum-product theorem, despite having a simple statement, is fundamental, and has applications to many areas of mathematics (many more than have been touched upon here). It has proven to be useful in computer science, and presumably has further potential. We have touched upon the sum-product theorem over $\mathbb{R}$ and $\mathbb{F}_p$ for prime p, but it has other extensions: for example, over rings.

Unresolved questions about the sum-product theorem include determining the optimal value for ϵ. Over the reals, it is believed to be 1. It is known that over finite fields, we must have $\epsilon \leq 1/2$.

Open questions for which the sum-product theorem may prove useful include the construction of an extractor for entropy $< .4999n$, and the construction of a disperser for entropy $<< n^{o(1)}$.

4

Proof of the Sum-Product Theorem

Boaz Barak

Scribe(s): Arnab Bhattacharyya and Moritz Hardt

Summary: We give a proof of the sum-product theorem for prime fields. Along the way, we also establish two useful results in additive combinatorics: the Plünnecke-Ruzsa lemma and the Balog-Szemerédi-Gowers lemma.

4.1 Introduction

Given a finite set $A \subset \mathbb{Z}$, let $A + A \overset{def}{=} \{a + b : a, b \in A\}$ and $A \cdot A \overset{def}{=} \{a \cdot b : a, b \in A\}$. Then, the sum-product theorem states that either $A + A$ or $A \cdot A$ is a large set; more precisely: there exists an $\epsilon > 0$ such that $\max\{|A + A|, |A \cdot A|\} \geq |A|^{1+\epsilon}$ for any set A. The two contrasting situations of a set having a large doubling and a set having a large squaring are realized by a geometric progression and an arithmetic progression respectively. If A is a geometric progression of length n, then $|A \cdot A| \leq O(n)$ while $|A + A| \geq \Omega(n^2)$. On the other hand, if A is an arithmetic progression of length n, then $|A + A| \leq O(n)$ while $|A \cdot A| \geq \tilde{\Omega}(n^2)$. So, the sum-product theorem can be thought of as roughly saying that any set is either "close" to an arithmetic progression or a geometric progression.

The sum-product theorem for integers was formulated by Erdös and Szemerédi [ES83b] who conjectured that the correct ϵ is arbitrarily close to 1. The theorem for $\epsilon = \frac{1}{4}$ was proved by Elekes in [Ele97]. Here, though, we are going to examine the sum-product theorem for finite fields. That is, the case when A is a subset of a finite field $\mathbb{F}$. Clearly, the relationship between $|A + A|$ and $|A \cdot A|$ is uninteresting when A equals $\mathbb{F}$ or, in general, when A is any subfield of $\mathbb{F}$ because then, $|A + A| = |A \cdot A| = |A|$. Therefore, in order to avoid trivialities as well as technical complications, we are going to insist that $\mathbb{F}$ be a prime field (so that it does not contain *any* proper subfield) and that $|A|$ be significantly smaller than $|\mathbb{F}|$ (so that A is not the entire field). In such a setting, where $|A| < |\mathbb{F}|^{0.9}$, [BT04] and [Kon03] proved the sum-product theorem with $\epsilon \approx 0.001$. We base our exposition on [Gre05, TV06, BW04].

4.2 The Sum-Product Theorem for Prime Fields

Theorem 4.2.1 (Sum-Product Theorem)**.** *Given $\mathbb{F}$ a prime field, $A \subseteq \mathbb{F}$ with $|A| < |\mathbb{F}|^{0.9}$, then there exists $\epsilon > 0$ such that* $\max\{|A+A|, |A \cdot A|\} \geq |A|^{1+\epsilon}$.

Our proof for Theorem 4.2.1 is based on two key lemmas. Intuitively, the first lemma will say that if both $|A+A|$ and $|A \cdot A|$ are small, then so is $|r(A)|$ where $r(\cdot)$ is any rational expression. The second lemma will say that for a particular rational expression, $r^*(\cdot)$, if $|A+A|$ is small, it is true that $|r^*(A)|$ is large. So, if both $|A+A|$ and $|A \cdot A|$ are small, the two lemmas together will yield a contradiction.

In truth, the lemmas that we can actually prove are not as strong as those described in the above paragraph, but they suffice for the proof.

Lemma 4.2.2. *For any $0 < \rho < 1$, if $|A \cdot A| \leq |A|^{1+\rho}$ and $|A+A| \leq |A|^{1+\rho}$, then for every rational expression $r(\cdot)$, there exists $B \subseteq A$ such that $|B| \geq |A|^{1-O(\rho)}$ and $|r(B)| \leq |A|^{1+O(\rho)}$.*

Lemma 4.2.3. *For prime field $\mathbb{F}$ and $A \subseteq \mathbb{F}$ with $|A| \leq |\mathbb{F}|^{0.9}$ and $|A+A| \leq |A|^{1.1}$,*

$$\left| A\left(\frac{AA - AA}{A - A} + A\right)\right| \geq |A|^{1.1}$$

Theorem 4.2.1 now follows simply from the two lemmas.

Proof of Theorem 4.2.1 from Lemma 4.2.2 and Lemma 4.2.3. For the sake of contradiction, suppose there is an $A \subseteq \mathbb{F}$ such that $|A| \leq |\mathbb{F}|^{0.9}$ but $|A+A|$ and $|A \cdot A|$ are both less than $|A|^{1+\epsilon}$ for any arbitrarily small constant $\epsilon > 0$. By Lemma 4.2.2, there is an absolute constant c for which there exists $B \subseteq A$ with $|B| \geq |A|^{1-O(\epsilon)}$ such that $\left|B\left(\frac{BB-BB}{B-B} + B\right)\right| \leq |A|^{1+O(\epsilon)} \leq |B|^{1+c\epsilon}$. Also, note that $|B+B| \leq |A+A| < |A|^{1+\epsilon} \leq |B|^{1+O(\epsilon)}$. So, by Lemma 4.2.3, $\left|B\left(\frac{BB-BB}{B-B} + B\right)\right| \geq |B|^{1.1}$. Therefore, $c\epsilon \geq 0.1$ which means $\epsilon \geq 0.1/c$, a contradiction to the fact that ϵ can be arbitrarily close to 0. □

What remains now is to prove Lemma 4.2.2 and Lemma 4.2.3.

4.3 Two Useful Tools

In order to prove Lemma 4.2.2 and Lemma 4.2.3, we will establish two generally useful tools: the Plünnecke-Ruzsa lemma and the Balog-Szemerédi-Gowers lemma. Both of these results hold when the set elements are from an arbitrary abelian group. So, any true statement made in this section about sums and differences of elements of a set implies a corresponding true statement about products and quotients (simply by changing the name of the group operation).

4.3.1 The PR Lemma

The Plünnecke-Ruzsa (PR) lemma states that if A and B are of equal size and if $|A+B|$ is small, then $|n_1A - n_2A + n_3B - n_4B|$ is also small where n_1, n_2, n_3, n_4 are positive integers. So, it cannot be the case that $|A+B|$ is small but $|A+B+B|$ is large if A and B are of equal size. Note that such a statement has the flavor of Lemma 4.2.2 but it involves just one operation.

Lemma 4.3.1 (PR Lemma [Ruz96, Plü69] (*"Iterated Sums Lemma"*))**.** *For any abelian group G, if $A, B \subseteq G$ such that $|A+B| \leq K|A|$ and*[1] $|B| = |A|$, *then* $|A \pm A \pm \cdots \pm A \pm B \pm \cdots \pm B| \leq K^{O(1)}|A|$ *(where the $O(1)$ notation hides a constant depending on the number of $+$'s and $-$'s).*

The following corollary summarizes the take-home messages that will be useful later on.

Corollary 4.3.2. *For any abelian group G, if $A, B \subseteq G$ with $|A| = |B|$:*

- $|A + A| \leq K^{O(1)}|A| \Leftrightarrow |A - A| \leq K^{O(1)}|A|$
- $|A + B| \leq K|A| \Rightarrow |A + A| \leq K^{O(1)}|A|$

Also, if $|A \cdot B| \leq K|A|$ with $|B| = |A|$, then $|A^{n_1}A^{-n_2}B^{n_3}B^{-n_4}| \leq K^{O(1)}|A|$ for positive integers n_1, n_2, n_3, n_4 (where the $O(\cdot)$ notation hides a constant depending on n_1, n_2, n_3 and n_4).

Proof. For the first bulleted item, one direction of the implication follows from Lemma 4.3.1 by setting $B = A$ while the other direction follows from setting $B = -A$. The second bulleted item is a special case of Lemma 4.3.1. Finally, the last sentence in the corollary is the multiplicative version of Lemma 4.3.1, obtained by renaming the addition operation to be multiplication. □

Now we present a short ingenious proof of the PR Lemma due to Ruzsa.

Proof of Lemma 4.3.1. We first prove that the following two claims imply the lemma.

Claim 4.3.3. $|C - C| \leq \frac{|C-D|^2}{|D|}$. *In particular, if $|D| \geq K^{-O(1)}|C|$ and $|C - D| \leq K^{O(1)}|C|$, then $|C - C| \leq K^{O(1)}|C|$.*

Claim 4.3.4. *If $|A + B| \leq K|A|$ (for $|B| = |A|$), then there exists a set $S \subseteq A + B$ such that $|S| \geq |A|/2$ and $|A + B + S| \leq K^{O(1)}|A|$.*

To see that the PR Lemma follows from the two claims, note that the hypotheses of the lemma satisfy the hypotheses of Claim 4.3.4, and so there exists a set S such that $|S| \geq |A|/2$ and $|A + B + S| \leq K^{O(1)}|A|$. Now set $C = A + B$ and $D = -S$; then, $|D| = |S| \geq |A|/2 \geq |C|/2K$, and $|C - D| = |A + B + S| \leq K^{O(1)}|A| \leq K^{O(1)}|C|$. Application of the last sentence of Claim 4.3.3 then yields $|(A - A) + (B - B)| = |C - C| \leq K^{O(1)}|C| \leq K^{O(1)}|A|$. We can repeat this argument with $A - A$ and $B - B$ instead of A and B to get that $|\ell A - \ell A + \ell B - \ell B| \leq K^{O(1)}|A|$ for every constant ℓ. This clearly implies the conclusion of Lemma 4.3.1.

Proof of Claim 4.3.3. Consider the map $\phi : (C - D) \times (C - D) \to G$ that maps (x_1, x_2) to $x_1 - x_2$. For any element $x = c - c'$ in $C - C$, note that $\phi(c - d, c' - d) = x$ for every $d \in D$. So[2], $|C - C| \leq \frac{|C-D|^2}{|D|}$.

The second sentence of the claim comes from just plugging in the given bounds. □

[1] Although it is not needed in what follows, the conclusions of the lemma hold true even when $|B| = K^{\Theta(1)}|A|$ rather than $|B| = |A|$

[2] We will use this simple counting argument often. In general, whenever we have a map $f : X \to Y$ and we can say that for $Z \subseteq Y$, every element of Z has at least k elements in its preimage, then it is true that $|Z| \leq \frac{|X|}{k}$. This is easy to show.

Proof of Claim 4.3.4. Let $S \overset{def}{=} \{s \in A + B :$ there are at least $\frac{|A|}{2K}$ representations of s as $s = a' + b'\}$, the set of "popular sums." Let $N \overset{def}{=} |A| = |B|$. Consider the map $\lambda : (A+B)\times(A+B) \to G$ that takes (x_1, x_2) to $x_1 + x_2$. Now, for every element $x = a+b+s$ in $A+B+S$, each representation of s as $a' + b'$ where $a' \in A$ and $b' \in B$ provides a preimage of x; namely, $\lambda(a + b', a' + b) = x$. Since there are at least $N/2K$ representations of every element of S as the sum of an element of A and an element of B, $|A + B + S| \leq \frac{|A+B|^2}{N/2K} \leq K^{O(1)}N$.

The lower bound on the size of S comes from a Markov-style argument. Suppose there are fewer than $N/2$ elements in $A + B$ which have at least $N/2K$ representations as $a + b$ for $a \in A, b \in B$. Clearly, any element can have at most N representations. So, the total number of representations for all elements of $A + B$ is $< \frac{N}{2} \cdot N + |A + B| \cdot \frac{N}{2K} \leq \frac{N^2}{2} + \frac{N^2}{2} = N^2$; however, this is impossible because there are a total of exactly N^2 pairs (a, b) with $a \in A, b \in B$. □

□

4.3.2 The BSG Lemma

Suppose we have a set A such that $|A + A|$ is small. It necessarily follows that there are many elements in $A + A$ that can be represented in many ways as $a_1 + a_2$ where $a_1, a_2 \in A$. In fact, as shown in the proof of Claim 4.3.4, if $|A+A| \leq K|A|$, there must be at least $|A|/2$ elements in $A+A$ which have at least $\frac{|A|}{2K}$ representations. In this section, we answer a stronger question. If $A + A$ is small, does there necessarily exist a large subset $B \subseteq A$ such that each element b of $B+B$ has many representations of the form $b = a_1 + a_2$ where $a_1, a_2 \in A$? (Note that our previous observation does not guarantee this.) Roughly speaking, the Balog-Szemerédi-Gowers (BSG) lemma shows that such a B exists.

Lemma 4.3.5 (BSG Lemma [BS96, Gow98] (*"Many Representations Lemma"*))**.** *Suppose G is an abelian group and $A \subseteq G$. If $|A - A| \leq K|A|$, then $\exists B \subseteq A$ such that $|B| \geq K^{-O(1)}|A|$ and every element $b \in B - B$ has $K^{-O(1)}|A|^7$ representations as $b = a_1 - a_2 + a_3 - a_4 + a_5 - a_6 + a_7 - a_8$ with $a_1, \ldots, a_8 \in A$.*

Firstly, observe that the lemma immediately implies that $|B-B| \leq K^{O(1)}|A|$. This is so, because the map $f : A^8 \to G$ defined by $f(a_1, a_2, a_3, a_4, a_5, a_6, a_7, a_8) = a_1-a_2+a_3-a_4+a_5-a_6+a_7-a_8$ maps at least $K^{-O(1)}|A|^7$ elements to each element of $B - B$. Hence, $|B - B| \leq \frac{|A|^8}{K^{-O(1)}|A|^7} = K^{O(1)}|A|$. Secondly, note that the lemma refers to $B-B$ instead of $B+B$, contrary to our previous discussion. However, we can obtain an analogous result for $B + B$.

Lemma 4.3.6 (BSG Lemma for addition)**.** *If $|A - A| \leq K|A|$, then $\exists B \subseteq A$ such that $|B| \geq K^{-O(1)}|A|$ and every element $b \in B + B$ has $K^{-O(1)}|A|^{11}$ representations as a sum of* 12 *elements of $A \cup -A$.*

As discussed in the previous lecture, it is often useful to establish *statistical* variants of statements in arithmetic combinatorics. To do this, let us view $A - A$ as a distribution, specifically the one described by the experiment of picking a_1 uniformly at random from A, a_2 uniformly at random from A and then outputting $a_1 - a_2$. We will call this distribution $\mathcal{A} - \mathcal{A}$. Then, the hypothesis of Lemma 4.3.5 that $|A - A| \leq K|A|$ can be written as $H_0(\mathcal{A} - \mathcal{A}) \leq \log(K|A|)$, where

$H_0(X) \overset{def}{=} \log|supp(X)|$. The next lemma instead takes as hypothesis the condition $H_2(\mathcal{A}-\mathcal{A}) \leq \log(K|A|)$, where $H_2(X) \overset{def}{=} -\log\|X\|_2^2 = -\log\sum_u X(u)^2$. (Note that $\|X\|_2^2$ is also equal to the collision probability of X, defined as $\Pr_{u \leftarrow X, v \leftarrow X}[u = v]$.)

Corollary 4.3.7. *For G an abelian group, let A be a subset of G. Let $\mathcal{A}-\mathcal{A}$ denote the distribution obtained by picking a_1 uniformly at random from A, a_2 uniformly at random from A and outputting $a_1 - a_2$. If $H_2(\mathcal{A}-\mathcal{A}) \leq \log(K|A|)$, then there exists $B \subseteq A$ such that $|B| \geq K^{-O(1)}|A|$ and $|B-B| \leq K^{O(1)}|A|$.*

Note that using Corollary 4.3.2, we can also conclude that $|B+B| \leq K^{O(1)}|A|$. Corollary 4.3.7 provides a partial converse to the observation made in the beginning of this subsection that small $|A-A|$ implies existence of many elements in $A-A$ with many representations. In Corollary 4.3.7, the condition that $H_2(\mathcal{A}-\mathcal{A})$ is small is approximately the same as saying that there are many elements in $A-A$ with lots of representations as $a_1 - a_2$ with $a_1, a_2 \in A$. So, roughly, the statement of Corollary 4.3.7 is that if there are many elements in $A-A$ with many representations, then there exists a large subset B of A such that $|B-B|$ is small. We cannot necessarily assert that $|A-A|$ itself is small, because such an assertion would be incorrect. It turns out that there are sets A such that there many elements in $A-A$ with many representations and yet $|A-A|$ is large. For example, let A be the union of an arithmetic progression of length $N/2$ and a geometric progression of length $N/2$. In this case, $H_2(\mathcal{A}-\mathcal{A}) = \log\Theta(N)$, but $|A+A| \geq \Omega(N^2)$ due to the geometric progression.

We next turn to the proofs of the above lemmas and corollary. We will give the proofs of Lemma 4.3.5 and Lemma 4.3.6 and then, indicate how the proofs need to be changed in order to obtain Corollary 4.3.7.

Proof of the BSG lemma (Lemma 4.3.5). We begin with a graph-theoretic claim.

Claim 4.3.8 (Comb Lemma)**.** *If a graph G has N vertices and average degree at least ρN, then there exists a subset B of $\geq \rho^{-O(1)}N$ vertices such that for all $u, v \in B$, there are at least $\rho^{O(1)}N^3$ length-4 paths from u to v.*

Set $N = |A|$. Let us see how Claim 4.3.8 implies the BSG Lemma. Define the graph G with vertex set A. There is an edge between vertices $a, e \in A$ if $a-e$ has at least $N/(2K)$ representations in $A-A$. As worked out in the proof of Claim 4.3.4, there must be at least $N/2$ elements in $A-A$ having more than $N/(2K)$ representations, as $|A-A| \leq K|A|$. So, the average degree of the graph G must be at least $\frac{2\cdot(\#\text{ of edges in }G)}{N} \geq \frac{2\cdot\frac{N}{2}\cdot\frac{N}{2K}}{N} = \frac{N}{2K}$. Therefore, we can apply Claim 4.3.8 with $\rho = \frac{1}{2K}$ to get a set B of size such that there are at least $\rho^{O(1)}N^3$ length-4 paths between any two vertices in B. Consider two vertices $a, e \in B$ and let $\langle a, b, c, d, e\rangle$ be a length-4 path between them. Observe that we can write $a-e = (a-b)+(b-c)+(c-d)+(d-e)$. By definition of G, there must be at least $N/(2K)$ representations of $a-b, b-c, c-d$ and $d-e$. Furthermore, there are at least $\rho^{O(1)}N^3$ many length-4 paths between a and e. So, $a-e$ has a total of at least $K^{-O(1)}N^3\cdot(N/(2K))^4 \geq K^{-O(1)}N^7$ distinct representations of the form $a_1 - a_2 + a_3 - a_4 + a_5 - a_6 + a_7 - a_8$.

Proof of the Comb Lemma (Claim 4.3.8). The main idea of the proof is the following. Given a graph $G = (V, E)$ with average degree at least ρN where $N = |V|$, we will show the existence of a set $B \subseteq V$ such that for every vertex $u \in B$, there are many vertices in B that share many neighbors with u. So, given two vertices u and v in B, there will be many vertices that share many common neighbors with both u and v; this directly implies many length-4 paths between u and v. Let us work out the details.

First, we modify G so that the minimum degree of a vertex in V is $\rho N/10$. We do this by deleting from V any vertex with degree $< \rho N/10$. Clearly, we remove at most $\rho N^2/10$ edges; the number of vertices remaining in V is at least $\Omega(\rho N)$ by a Markov-style argument. For convenience, let us retain the names $G = (V, E)$ for the modified graph.

Next, pick a random vertex $x \in V$ and let $B' = \Gamma(x)$, the set of neighbors of x. Note that $\mathbb{E}_x\left[|B'|\right] \geq \frac{\frac{9}{10}\rho N^2}{N} = \frac{9}{10}\rho N$. We want to claim now that there are many length-2 paths between many pairs of vertices in B'. Say vertices a and b are *unfriendly* if $|\Gamma(a) \cap \Gamma(b)| < \rho^2 N/200$. Let X denote the number of unfriendly pairs of vertices in B'. For $a, b \in V$, let $X_{a,b}$ be the indicator variable that is 1 iff a and b are both in $B' = \Gamma(x)$ and are unfriendly with each other. $X = \sum_{a,b\in V} X_{a,b}$. For any $a, b \in V$,

$$\mathop{\mathbb{E}}_x\left[X_{a,b}\right] = \Pr_x[X_{a,b} = 1] \leq \Pr_x[a, b \in \Gamma(x)|\ a \text{ and } b \text{ are unfriendly}] \leq \rho^2/200$$

because for an unfriendly pair, only $\rho^2 N/200$ choices of x would lead to both being in $\Gamma(x)$; so, $\mathbb{E}_x\left[X\right] \leq \frac{\rho^2}{200}\binom{N}{2} \leq \frac{\rho^2}{400}N^2$.

Next, we want to extend this observation in order to show that for significantly many $a_1 \in B'$, there are many $a_2 \in B'$ such that a_1 and a_2 are friendly. For this purpose, let S denote the number of $a_1 \in B'$ such that a_1 is unfriendly with more than $\frac{\rho}{100}N$ vertices in B'. Then, it follows that $|S| \cdot \frac{\rho}{100}N \leq X$; therefore, $\mathbb{E}_x\left[|S|\right] \leq \frac{100}{\rho N}\frac{\rho^2}{400}N^2 = \frac{\rho}{4}N$. Therefore, $\mathbb{E}_x\left[|B'| - |S|\right] \geq (\frac{9}{10} - \frac{1}{4})\rho N \geq \frac{1}{2}\rho N$. So, by the averaging principle, there must exist an x such that $B \overset{def}{=} B' - S$ is of size at least $\frac{1}{2}\rho N$. Furthermore, because all the vertices in S have been removed, B has the property any vertex in B is unfriendly with at most $\frac{1}{100}\rho N$ other vertices in B.

Now suppose u and v are any two vertices in B. u has at most $\frac{1}{100}\rho N$ vertices in B unfriendly to it and v has at most $\frac{1}{100}\rho N$ vertices in B unfriendly to it; so, by the union bound, there can be at most $\frac{1}{50}\rho N$ vertices in B unfriendly to either u or v. Let w be one of the at least $(\frac{1}{2} - \frac{1}{50})\rho N = \frac{12}{25}\rho N$ vertices friendly to both u and v. Then, there are at least $\rho^2 N/200$ vertices x that are common neighbors of both u and w, and similarly at least $\rho^2 N/200$ vertices y that are common neighbors of both w and x. Each such choice of w, x, and y defines a path of length-4 between u and v. The number of such paths is at least $\frac{12}{25}\rho N \cdot (\rho^2 N/200)^2 \geq \rho^{O(1)}N^3$. □

□

Proof of the Additive BSG Lemma (Lemma 4.3.6). This proof is very similar to the proof of Lemma 4.3.5 above. For this reason, we skip some of the calculations that were already performed above. Again, we begin with a graph-theoretic claim.

Claim 4.3.9 (Bipartite Comb Lemma)**.** *If $G = (A, B, E)$ is a bipartite graph with $|A| = |B| = N$ and with the average degree at least ρN, then there exist subsets $A' \subseteq A$ and $B' \subseteq B$, each of size at least $K^{-O(1)}N$, such that for all $u \in A$ and $v \in B$, there are $\rho^{O(1)}N^2$ length-3 paths from u to v.*

Proof. The main idea of the proof is the following. We will show existence of a set $A' \subseteq A$ such that for every vertex $a_1 \in A'$, there are many vertices $a_2 \in A'$ that share many neighbors with

a_1. Also, we will show existence of another set $B' \in B$ such that every vertex in B' has lots of neighbors in A'. By our choice of parameters, for every pair of vertices $u \in A'$ and $v \in B'$, u will have a lot of length-2 paths to vertices in A' that are neighbors of v; so there will be many length-3 paths between u and v.

First, we modify G so that the minimum degree of a vertex in A is $\rho N/10$. We do this by deleting from A any vertex with degree $< \rho N/10$. Clearly, we remove at most $\rho N^2/10$ edges; the number of vertices remaining in A is at least $\frac{9}{10}\rho N$ by a Markov-style argument. The number of vertices remaining in B is still N. For convenience, let us retain the names A and B for the modified sets of vertices.

Next, we will define two vertices a and b to be *unfriendly* if $|\Gamma(a) \cap \Gamma(b)| < \rho^3 N/200$. Proceeding exactly like in the proof of Claim 4.3.8, we find that there exists a set $A' \subseteq A$ that is of size at least $\frac{1}{2}\rho N$ such that any vertex in A' is unfriendly with at most $\frac{1}{100}\rho^2 N$ other vertices in A'.

Next, let B' be the set of vertices in B that have more than $\frac{1}{50}\rho^2 N$ neighbors in A'. We want to lower-bound $|B'|$. To do so, first note that the number of edges between A' and B is at least $\frac{1}{10}\rho N \cdot |A'| \geq \frac{1}{20}\rho^2 N^2$. On the other hand, the number of edges between A' and B is at most $(N - |B'|)\frac{1}{50}\rho^2 N + |B'|N \leq \frac{1}{50}\rho^2 N^2 + |B'|N$. Combining the above two sentences, we get that $|B'| \geq (\frac{1}{20} - \frac{1}{50})\rho^2 N \geq \frac{1}{50}\rho^2 N$.

Finally, we show that the above A' and B' satisfy the claims made in the statement of the lemma. Consider any $a \in A'$ and $d \in B'$. Now, there must at least $\frac{1}{50}\rho^2 N$ neighbors of d in A'. At most $\frac{1}{100}\rho^2 N$ of these neighbors could be unfriendly with a; so, there must be at least $\frac{1}{100}\rho^2 N$ neighbors c such that $|\Gamma(a) \cap \Gamma(c)| \geq \frac{1}{200}\rho^3 N$. Each such $b \in \Gamma(a) \cap \Gamma(c)$ defines a path $\langle a, b, c, d\rangle$ between a and d. So, the number of these paths of length 3 between a and d is at least $\frac{1}{200}\rho^3 N \cdot \frac{1}{100}\rho^2 N \geq 10^{-4}\rho^5 N^2$. □

Claim 4.3.10. *If $|A + B| \leq K|A|$ with $|A| = |B| = N$, then there exist $A' \subseteq A$ and $B' \subseteq B$ with $|A'| \geq K^{-O(1)}N$ and $|B'| \geq K^{-O(1)}N$ such that each element $x \in A' + B'$ has at least $K^{-O(1)}N^5$ representations as $x = a_1 - a_2 + a_3 + b_1 - b_2 + b_3$ with $a_1, a_2, a_3 \in A$ and $b_1, b_2, b_3 \in B$.*

Proof. Define a bipartite graph G with the elements of A on one side and elements of B on the other. G has an edge between $a \in A$ and $b \in B$ iff $a + b$ has at least $KN/2$ representations. By a Markov argument, $|A + B| \leq K|A|$ implies that the degree of G is at least $N/(2K)$. Now, we apply Claim 4.3.9 with $\rho = \frac{1}{2K}$ to get subsets $A' \subseteq A$ and $B' \subseteq B$. Now, take any element $a + b \in A' + B'$. There exist at least $K^{-O(1)}N^2$ choices of $b' \in B$ and $a' \in A$ such that $\langle a, b', a', b\rangle$ is a path of length 3 in G. Then, we can write:

$$a + b = (a + b') - (a' + b') + (a' + b)$$

Each of the three terms above has at least $KN/2$ representations. So, in all, $a + b$ must have at least $K^{-O(1)}N^2 \cdot (KN/2)^3 \geq K^{-O(1)}N^5$ distinct representations. □

Let us show now that Claim 4.3.10 implies[3] Lemma 4.3.6. If $|A + A| \leq K|A|$, then using $B = A$ in Claim 4.3.10, there are subsets $A', B' \subseteq A$ such that $|A'|$ and $|B'|$ are at least $K^{-O(1)}N$ and each $x \in A' + B'$ has at least $K^{-O(1)}N^5$ representations as $x = a + b - c - d + e + f$ with

[3] In fact, Claim 4.3.10 also implies a weaker form of Lemma 4.3.5 with more terms in the representations. We chose to present the stronger proof because it was the one presented in class.

$a, b, c, d, e, f \in A$. Moreover, since $B' - A'$ is a subset of $A - A$ and so is also smaller than $K^{O(1)}|A|$ (by Corollary 4.3.2), we can apply Claim 4.3.10 once again (with $A = -A'$ and $B = B'$) to find subsets $A'' \subseteq A'$ and $B'' \subseteq B'$ such that every element of $y \in B'' - A''$ has $K^{-O(1)}N^5$ representations as $y = -a + b + c - d - e + f$ with $a, b, c, d, e, f \in A$.

We claim that every element in $A'' + A''$ has at least $K^{-O(1)}N^{11}$ representations as a sum of 12 elements from $A \cup -A$. Indeed, for every $x = a + c$ in $A'' + A''$, pick a $b \in B''$. Then, writing $a + c = (a + b) - (b - c)$, we see that x can be represented in $K^{-O(1)}N^5 \cdot K^{-O(1)}N^5 \cdot K^{-O(1)}N \geq K^{-O(1)}N^{11}$ ways as $x = (a_1 + b_1 - a_2 - b_2 + a_3 + b_3) - (-a_4 + b_4 + a_5 - b_5 - a_6 + b_6)$. □

We now sketch the proof of Corollary 4.3.7. In fact, it is very similar to the proof given above for Lemma 4.3.5. We define a graph G where there are edges between two vertices a and b iff there are $N/(2K)$ representations of $a - b$. Now, we just have to show that the average degree of G is $\Theta(N/K)$ and then the rest is exactly the same as in the proof of Lemma 4.3.5. This is not hard to show by a Markov-style argument.

4.4 Proof of Lemma 4.2.2

We now turn to the proof of our first main lemma. Intuitively, Lemma 4.2.2 says if $A \cdot A$ and $A + A$ are small, then we can also control the size of more complex rational expressions at least for a large subset $B \subseteq A$. We will actually prove a seemingly weaker claim.

Lemma 4.4.1. *If $A \subseteq \mathbb{F}$ satisfies $|A \cdot A| \leq K|A|$ and $|A + A| \leq K|A|$, then there exists a set $B \subseteq A$ with $|B| \geq \frac{1}{K^{O(1)}}|A|$ but $|B \cdot B + B \cdot B| \leq K^{O(1)}|A|$.*

However, it is not difficult to obtain Lemma 4.2.2 from this statement. First of all, our proof of the above statement extends straightforwardly to the expression $B^k + B^k$ for any $k > 2$. The reader finds this step done carefully in [BW04]. Once we have $B^k + B^k$, the PR Lemma allows us to iterate this sum and thus obtain any fixed length polynomial

$$p(B) = B^k + \cdots + B^k - B^k - \cdots - B^k.$$

Ultimately, we need rational expressions of the form $r(B) = p(B)/q(B)$ where p and q are polynomials as above. But the multiplicative version of PR tells us, if the product of polynomials $p(B)q(B)$ is small, then so is the rational $p(B)/q(B)$. On the other hand, we can bound the size of the product $p(B)q(B)$ by the size of a polynomial of the above form of higher degree and greater length.

4.4.1 A Proof Sketch With Unrealistic Assumptions

Before we come to the actual proof of Lemma 4.4.1, we will sketch the proof using a strongly idealized (possibly wrong) version of the BSG Lemma. We called this the "Dream BSG" Lemma during the lecture.

Lemma 4.4.2 ("Dream BSG")**.** *If $|A - A| \leq K|A|$ then there exists a subset $B \subseteq A$ with $|B| \geq \frac{1}{K^{O(1)}}|A|$ such that every $b \in B - B$ has $\frac{1}{K^{O(1)}}|A|$ representations as $b = a_1 - a_2$ with $a_1, a_2 \in A$. Moreover, every $b \in B + B$ has $\frac{1}{K^{O(1)}}|A|$ representations as $b = a_1 + a_2$ with $a_1, a_2 \in A$.*

Proof sketch of Lemma 4.4.1 using "Dream BSG". Apply the "Dream BSG" Lemma so as to obtain a subset $B \subseteq A$ with $|B| \geq \frac{1}{K^{O(1)}}|A|$ such that any $b \in B - B$ has $\frac{1}{K^{O(1)}}|A|$ representations as $a = x - y$ and every $b \in B \cdot B$ has $\frac{1}{K^{O(1)}}|A|$ representations as $b = x'y'$.

Claim 4.4.2.1. $|(B - B)B| \leq K^{O(1)}|A|$.

Proof. For every $a \in B - B$, we can write

$$a = x - y \tag{4.1}$$

where $x, y \in A$ in $\frac{1}{K^{O(1)}}|A|$ different ways. If we multiply Eq. 4.1 by $z \in B$, we get that az has $\frac{1}{K^{O(1)}}|A|$ different representations as $az = x' + y'$ with $x', y' \in A \cdot A$. Thus,

$$|(B - B)B| \leq \frac{|A \cdot A|^2}{\frac{1}{K^{O(1)}}|A|} = K^{O(1)}|A|.$$

□

Now, let $a, b \in B \cdot B$ where $b = y_1 y_2$. We have $\frac{1}{K^{O(1)}}N$ pairs (x_1, x_2) such that

$$a - b = x_1 x_2 - y_1 y_2 = (x_1 - y_1)x_2 + y_1(x_2 - y_2).$$

That is, for each pair (x_1, x_2) we get a unique representation of $a - b$ in terms of $z_1 + z_2$ with $z_i \in (B - B)B$. By our previous claim, this implies

$$|B \cdot B - B \cdot B| \leq \frac{1}{K^{O(1)}}N.$$

By the PR Lemma, the same is true for $B \cdot B + B \cdot B$. □

4.4.2 The Actual Proof

Our actual proof of Lemma 4.4.1 is conceptually very similar to the previous proof sketch. It is somewhat more cumbersome due to the weaker guarantees of the true BSG Lemmas.

Proof of Lemma 4.2.2. We will first strengthen our assumptions by proving the following claim.

Claim 4.4.2.2. *There exists a set $A' \subseteq A$ with $|A'| \geq \frac{1}{K^{O(1)}}N$ such that*

$$(A' - A')A'^{11}A'^{-11} \leq K^{O(1)}N.$$

Proof. By Lemma 4.3.5, there is a set $A' \subseteq A$ of size at least $\frac{1}{K^{O(1)}}N$ such that every element $b \in A' - A'$ has $\frac{1}{K^{O(1)}}N^7$ representations as

$$b = a_1 - a_2 + a_3 - a_4 + a_5 - a_6 + a_7 - a_8 \tag{4.2}$$

with $a_i \in A$. Multiply Equation 4.2 by an arbitrary $x \in A^{11}A^{-11}$. Thus every element $d \in (A' - A')A'^{11}A'^{-11}$ has $\frac{1}{K^{O(1)}}N^7$ representations as

$$d = c_1 - c_2 + c_3 - c_4 + c_5 - c_6 + c_7 - c_8$$

with $c_i \in A^{12}A^{-11}$. But, $|A^{12}A^{-11}| \leq K^{O(1)}N$ by the PR Lemma. □

Hence, we may assume without loss of generality

$$(A-A)A^{11}A^{-11} \leq K^{O(1)}N.$$

Now, apply Lemma 4.3.6 so as to obtain $B \subseteq A$ with $|B| \geq \frac{1}{K^{O(1)}}N$ such that every $b \in B \cdot B$ has $\frac{1}{K^{O(1)}}N^{11}$ representations as $b = x_1 x_2 \cdots x_{12}$ with $x_i \in A \cup A^{-1}$.

For every $a, b \in B \cdot B$. We fix $b = y_1 \cdots y_{12}$ and vary over the representations $a = x_1 \cdots x_{12}$. For each such representation of a we obtain a unique representation of $a - b$ as

$$\begin{aligned} a - b &= x_1 x_2 \cdots x_{12} - y_1 \cdots y_{12} \\ &= (x_1 - y_1) x_2 x_3 x_4 \cdots x_{12} \\ &+ y_1 (x_2 - y_2) x_3 x_4 \cdots x_{12} \\ &+ y_1 y_2 (x_3 - y_3) x_4 \cdots x_{12} \\ &\vdots \\ &+ y_1 y_2 \cdots y_{11} (x_{12} - y_{12}). \end{aligned}$$

That is, there are $\frac{1}{K^{O(1)}}N$ representations of $a - b$ as $z_1 + z_2 + \cdots + z_{12}$ where the each z_i is in a set of size at most $|(A-A)A^{11}A^{-11}| \leq K^{O(1)}N$. Thus,

$$|B \cdot B - B \cdot B| \leq K^{O(1)}N.$$

An application of the PR Lemma finishes the proof. □

4.5 Proof of Lemma 4.2.3

In order to conclude the proof of the Sum-Product Theorem, we need to exhibit a rational expression $r(A)$ which grows even if $A + A$ is small. In the previous lecture, we already saw such a rational expression, namely the one from [BW04]. In this lecture, we give an independent (possibly simpler) proof with a different rational expression. Recall the expression from Lemma 4.2.3,

$$A\left(A + \frac{AA - AA}{A - A}\right).$$

It is helpful to see expression as the composition of two parts. The tricky part is to show that set $A(A + x)$ grows compared to A at least for some field element $x \in \mathbb{F}$. But, $A(A + x)$ is not quite yet a rational expression in A. So, we would like to be able to replace x by some rational term $r(A)$. This is where the part $(AA - AA)/(A - A)$ comes into play.

Proposition 4.5.1. *Let $\mathbb{F}$ be a prime field. For any $A \subseteq \mathbb{F}$ and $x \in \mathbb{F}$, if $|A(A + x)| < |A|^2$, then $x \in \frac{AA-AA}{A-A}$.*

Proof. If $|A(A + x)| < |A|^2$, then we have $a, b, c, d \in A$ with $a \neq c$ such that $a(b + x) = c(d + x)$. Hence, $x = \frac{ab-cd}{a-c}$. □

Clearly, it only remains to prove the following claim.

Claim 4.5.2. *Under the assumptions of Lemma 4.2.3, there exists an $x \in \mathbb{F}$ such that*

$$|A|^{1.1} < |A(A+x)| < |A|^2$$

Proof. Suppose otherwise and let

$$\Lambda = \{x : |A(A+x)| = |A|^2\}.$$

We will prove a sequence of subclaims which turns out to contradict the fact that we are working over a field of prime order.

Claim 4.5.2.1.

$$\Lambda \neq \emptyset$$

Proof. For a random $x \in \mathbb{F}$, we are interested in the expected *collision probability* of $A(A+x)$. Recall,

$$\mathrm{cp}(A(A+x)) = \Pr_{a,b,c,d}[a(b+x) = c(b+x)].$$

We have,

$$\begin{aligned}
\mathbb{E}_{x\in\mathbb{F}}[\mathrm{cp}(A(A+x))] &= \frac{1}{|\mathbb{F}||A|^4}\sum_{x\in\mathbb{F}}\#\{a,b,c,d : a(b+x) = c(d+x)\} \\
&= \frac{1}{|\mathbb{F}||A|^4}\sum_{x\in\mathbb{F}}\#\{a,b,c,d : ab-cd = (c-a)x\} \qquad (4.3)
\end{aligned}$$

Observe, if a quadruple (a,b,c,d) has $a-c \neq 0$, then there is precisely one $x \in \mathbb{F}$ that satisfies $ab-cd = (c-a)x$. We have $|A|^4 - |A|^3$ such tuples. On the other hand, if $ab-cd = a-c = 0$, then every x satisfies this equation. But in this case, there are less than $2|A|^2$ such quadruples. Hence, we have the upper bound

$$\sum_{x\in\mathbb{F}}\#\{a,b,c,d : ab-cd = (c-a)x\} \leq |A|^4 - |A|^3 + 2|\mathbb{F}||A|^2 \qquad (4.4)$$

A simple calculation shows, if the field $\mathbb{F}$ is large enough, Equation 4.3 and Equation 4.4 imply

$$\mathbb{E}_{x\in\mathbb{F}}[\mathrm{cp}(A(A+x))] < |A|^{-1.1}.$$

But, by our assumption we ruled out any $x \in \mathbb{F}$ with

$$|A|^{-2} < \mathrm{cp}(A(A+x)) < |A|^{-1.1}.$$

Hence, there exists an x with $\mathrm{cp}(A(A+x)) = |A|^{-2}$ which is equivalent to

$$|A(A+x)| = |A|^2.$$

□

Claim 4.5.2.2. *There exists an element $d \neq 0$ such that*

$$\Lambda + d \subseteq \Lambda.$$

Proof. We want to prove there exists an element $d \neq 0$ such that for all $x \in \Lambda$, we have $d + x \in \Lambda$. Notice, by our assumption it is sufficient to show that $|A(A + d + x)| > |A|^{1.1}$ for all $x \in \Lambda$.

Let $x \in \Lambda$. Then, the mapping $(a, b) \mapsto a(b + x)$ is collision-free on $A \times A$. In particular, the mapping is collision free on the domain $A \times A'$ with $A' = A \cap (A + d)$, since $A' \subseteq A$. Thus for all d,

$$|A(A + d + x)| \geq |A((A \cap (A + d)) + x)| \geq |A| \cdot |A \cap (A + d)|.$$

So, this means we are done if there is an element $d \neq 0$ such that $|A \cap (A + d)| > |A|^{0.1}$. But,

$$A \cap (A + d) = \{a \mid \exists b : a - b = d\}.$$

Hence,

$$\sum_{d \in (A-A) \setminus \{0\}} |A \cap (A + d)| = |A|^2 - |A|.$$

Thus, there exists $d \neq 0$ with

$$|A \cap (A + d)| \geq \frac{|A|^2 - |A|}{|A - A| - 1} > \frac{|A| - 1}{|A|^{0.1}}.$$

If $|A| > 2$, we have $(|A| - 1)/|A|^{0.1} > |A|^{0.1}$ as desired. The remaining cases of $|A| \leq 2$ can be checked separately. □

Now, let $x \in \Lambda$ as given by Claim 4.5.2.1. By Claim 4.5.2.2, there exists $d \neq 0$ such that $x + d \in \Lambda$. Indeed, $x + kd \in \Lambda$ for any $k \geq 0$. On the other hand, $\mathbb{F}$ is of prime order. So, there will be a $k \geq 0$ such that $x + kd = 0$. But this is a contradiction, since

$$|A(A + 0)| = |A \cdot A| < |A|^2.$$

□

$=$

5

Proof of Szemerédi's Regularity Lemma

Luca Trevisan

Scribe(s): Wolfgang Mulzer

Summary: We give a proof of Szemerédi's Regularity Lemma [Sze78], which states essentially that any graph can be partitioned into a constant number of pieces such that the distribution of the edges between almost any pair of pieces is pseudo-random.

5.1 Szemerédi's Regularity Lemma

Previously we encountered Szemerédi's Regularity Lemma and saw how it can be used to prove Szemerédi's theorem for $k = 3$. Now we are going to prove the Regularity Lemma.

Recall that a pair (U, W) of subsets of vertices of a graph $G = (V, E)$ is ε-regular if the number of edges between any pair of large enough subsets of U and V is ε-close to what we would expect if the edges between them were chosen independently with the same probability. More precisely:

Definition 5.1.1. Let $\varepsilon > 0$, $G = (V, E)$ be a graph, and $U, W \subseteq V$. The pair (U, W) is called *ε-regular*, if for all $S \subseteq U$ and $T \subseteq W$ with $|S| \geq \varepsilon|U|$ and $|T| \geq \varepsilon|W|$ we have

$$|d(S,T) - d(U,W)| \leq \varepsilon,$$

where $d(A, B)$ denotes the *edge density*

$$d(A,B) \overset{def}{=} \frac{|E(A,B)|}{|A||B|}$$

between $A, B \subseteq V$.

With this definition in mind, we can give the precise statement of the Regularity Lemma:

Theorem 5.1.2. *For every $\varepsilon > 0$ and $t \in \mathbb{N}$ there is a constant $k(\varepsilon, t)$ such that for every graph $G = (V, E)$ with at least t vertices there is a partition $(V_1, V_2, \ldots, V_k)$ of the vertices with $t \leq k \leq k(\varepsilon, t)$ and $|V_1| = |V_2| = \cdots = |V_{k-1}| \geq |V_k|$ such that at least $(1 - \varepsilon)\binom{k}{2}$ of the pairs (V_i, V_j) are ε-regular.*

The Lemma states that any graph can essentially be partitioned into a constant number of pieces such that the distribution of the edges between almost any pair of pieces is pseudo-random, that is, ε-regular. Note that the Lemma is only meaningful for dense graphs, because for sparse graphs the density of the edges between the pieces of the partition tends to 0.

Note also that we make no statement about the edges inside each piece of the partition. However, we can specify the minimum size t of of the partition, which gives us control over the fraction of edges between the pieces.

Furthermore, choosing a random partition of k pieces will not do the job. Indeed, if we take a large bipartite graph and randomly partition its vertices into k pieces, then every piece will contain many vertices from both sides of the bipartition, and hence no pair of pieces will be ε-regular, because we can find large subsets with no edges between them.

Considering this counterexample and the many deep consequences of the Regularity Lemma, we expect the proof to be quite complicated and involved. However, a straightforward greedy method along with a clever choice of potential function yields the desired result, as we will now see.

Proof. The partition is generated by the following greedy method:

Algorithm 5.1.3 (Generating an ε-regular partition)**.**

1. Let $\mathcal{P} = (V_1, V_2, \ldots, V_k)$ be an arbitrary partition of V into k pieces with $|V_1| = \cdots = |V_{k-1}| \geq |V_k|$, where $k = \max(\frac{1}{\varepsilon}, t)$.
2. If $\mathcal{P}$ fulfills the requirements of the theorem, then STOP.

   Otherwise for at least $\varepsilon\binom{k}{2}$ pairs (V_i, V_j) we get subsets $S_{ij} \subseteq V_i$, $S_{ji} \subseteq V_j$ such that both $|S_{ij}| \geq \varepsilon|V_i|$, $|S_{ji}| \geq \varepsilon|V_j|$ and we have $|d(S_{ij}, S_{ji}) - d(V_i, V_j)| \geq \varepsilon$.
3. Subdivide each V_i into at most 2^{k-1} sets according to the σ-algebra generated by the sets $S_{ij} \subseteq V_i$ (see Figure 5.1). Call the refined partition $\mathcal{P}'$. We have $|\mathcal{P}'| \leq k2^{k-1}$.
4. Subdivide the pieces of $\mathcal{P}'$ into pieces of size $n/\left(k2^{k-1}\right)^2$ and (possibly) a remainder of smaller size. Recombine the remainder pieces arbitrarily into pieces of size $n/\left(k2^{k-1}\right)^2$ and (possibly) one piece of smaller size. Denote the resulting partition by $\mathcal{P}$, set $k := |\mathcal{P}|$ and go to Step 2.

We need to show that Algorithm 5.1.3 stops after a constant number of steps. In order to measure the progress we make in each step, we define a potential function that represents the regularity of the current partition.

More precisely, let $\mathcal{P} = (V_1, V_2, \ldots, V_k)$ be a partition of V. For each $A, B \subseteq V$, define a random variable $X[A, B]$ as follows: Let $v_1, v_2 \leftarrow_{\mathrm{R}} V$. Then $X[A, B] \overset{def}{=} d(A, B)^2$ if $v_1 \in A$ and $v_2 \in B$, and $X[A, B] \overset{def}{=} 0$, otherwise. Now we define the potential function $\Phi(\mathcal{P})$:

$$\Phi(\mathcal{P}) \overset{def}{=} \sum_{1 \leq i < j \leq |\mathcal{P}|} \mathbb{E}[X[V_i, V_j]] = \sum_{1 \leq i < j \leq |\mathcal{P}|} \frac{|V_i||V_j|}{|V|^2} d(V_i, V_j)^2.$$

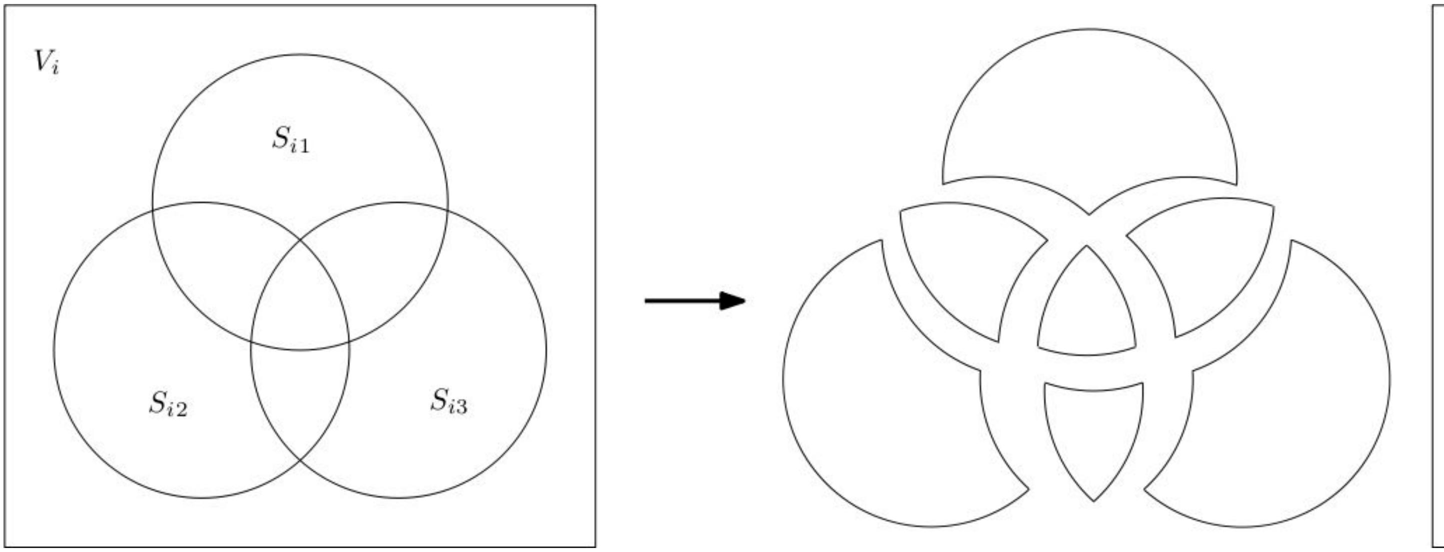

Figure 5.1: Splitting the piece V_i according to the σ-algebra generated by S_{i1}, S_{i2}, S_{i3}.

The potential function can be interpreted as the variance of the random variable Z which is defined by taking $v_1, v_2 \leftarrow_{\mathrm{R}} V$ and letting $Z \overset{def}{=} d(V_i, V_j)$ if $v_1 \in V_i$ and $v_2 \in V_j$ (and $Z \overset{def}{=} 0$ if v_1, v_2 are in the same piece of the partition). Since all the $d(V_i, V_j)$ are at most 1, it follows that $\Phi(\mathcal{P}) \leq 1$ for any partition $\mathcal{P}$.

Next, we show that $\Phi(\mathcal{P}') \geq \Phi(\mathcal{P})$ for any refinement $\mathcal{P}'$ of $\mathcal{P}$. It suffices to consider refinements that split a piece V_i into two pieces S, $V_i \backslash S$. In this case, we have

$$\Phi(\mathcal{P}') - \Phi(\mathcal{P}) = \mathbb{E}[X[S, V_i \backslash S]] + \sum_{j \neq i} \left(\mathbb{E}[X[S, V_j]] + \mathbb{E}[X[V_i \backslash S, V_j]] - \mathbb{E}[X[V_i, V_j]]\right).$$

Now if we define a random variable Y_j by picking $v \leftarrow_{\mathrm{R}} V_i$ and setting $Y_j \overset{def}{=} d(S, V_j)$ if $v \in S$ and $Y_j \overset{def}{=} (V_i \backslash S, V_j)$ if $v \in V_i \backslash S$, we have

$$\mathbb{E}[X[V_i, V_j]] = \frac{|V_i||V_j|}{|V|^2}\mathbb{E}[Y_j]^2 \leq \frac{|V_i||V_j|}{|V|^2}\mathbb{E}[Y_j^2] = \mathbb{E}[X[S, V_j]] + \mathbb{E}[X[V_i \backslash S, V_j]],$$

since for any random variable W we have $\mathbb{E}[W]^2 \leq \mathbb{E}[W^2]$. Hence, refining the partition can only increase the potential function.

Now we show that splitting a non-regular pair (V_i, V_j) increases the potential function significantly. Let $S \subseteq V_i$, $T \subseteq V_j$ be such that $|S| \geq \varepsilon |V_i|$, $|T| \geq \varepsilon |V_j|$ and $|d(S, T) - d(V_i, V_j)| \geq \varepsilon$. Let $\mathcal{P}'$ be the partition resulting from $\mathcal{P}$ by splitting (V_i, V_j) into $(S, V_i \backslash S, T, V_j \backslash T)$. We have

$$\Phi(\mathcal{P}') - \Phi(\mathcal{P}) \geq \mathbb{E}[X[S, T]] + \mathbb{E}[X[S, V_j \backslash T]] + \mathbb{E}[X[V_i \backslash S, T]] + \mathbb{E}[X[V_i \backslash S, V_j \backslash T]] - \mathbb{E}[X[V_i, V_j]],$$

since by the above calculation the contribution of all other pairs involving V_i or V_j is non-negative. Now define a random variable Y by picking $v_1 \leftarrow_{\mathrm{R}} V_i$, $v_2 \leftarrow_{\mathrm{R}} V_j$ and letting $Y \overset{def}{=} d(A, B)$ if $v_1 \in A$,

$v_2 \in B$, for $A \in (S, V_i \backslash S)$, $B \in (T, V_j \backslash T)$. We have $\mathbb{E}[Y] = d(V_i, V_j)$, and with probability at least ε^2, Y deviates from its mean by more than ε. Hence, we have

$$\mathrm{Var}[Y] = \mathbb{E}[(Y - \mathbb{E}[Y])^2] \geq \varepsilon^2 \cdot \varepsilon^2.$$

That means

$$\begin{aligned} \mathbb{E}[X[S,T]] + \mathbb{E}[X[S,V_j \backslash T]] + \mathbb{E}[X[V_i \backslash S, T]] + \mathbb{E}[X[V_i \backslash S, V_j \backslash T]] &= \frac{|V_i||V_j|}{|V|^2} \mathbb{E}[Y^2] \\ &\geq \frac{|V_i||V_j|}{|V|^2} \left(\mathbb{E}[Y]^2 + \varepsilon^4 \right) \\ &= \mathbb{E}[X[V_i, V_j]] + \frac{|V_i||V_j|}{|V|^2} \varepsilon^4, \end{aligned}$$

so the potential function increases by at least $\varepsilon^4 |V_i||V_j| / |V|^2$.

Now let us consider the partition $\mathcal{P}'$ obtained by refining $\mathcal{P}$ in Step 3. We have

$$\begin{aligned} \Phi(\mathcal{P}') - \Phi(\mathcal{P}) \geq & \sum_{\substack{1 \leq i < j \leq |\mathcal{P}| \\ (V_i, V_j) \text{ not regular}}} \left(\sum_{\substack{A,B \in \mathcal{P}' \\ A \subseteq V_i, B \subseteq V_j`}} \mathbb{E}[X[A,B]] - \mathbb{E}[V_i, V_j] \right) \\ \geq & \sum_{\substack{1 \leq i < j \leq |\mathcal{P}| \\ (V_i, V_j) \text{ not regular}}} \left(\sum_{\substack{A \in (S_{ij}, V_i \backslash S_{ij}) \\ B \in (S_{ji}, V_j \backslash S_{ji})}} \mathbb{E}[X[A,B]] - \mathbb{E}[V_i, V_j] \right) \\ \geq & \sum_{\substack{1 \leq i < j \leq |\mathcal{P}| \\ (V_i, V_j) \text{ not regular}}} \frac{|V_i||V_j|}{|V|^2} \varepsilon^4 \\ \geq & \, \Omega\left(\varepsilon^5\right), \end{aligned}$$

since $\mathcal{P}'$ refines every partition which splits a non-regular pair (V_i, V_j) according to the witness sets S_{ij}, S_{ji}, and since at least an ε-fraction of all pairs is non-regular.

In Step 4, we recombine some sets of the partition, and hence the potential function can decrease. However, the total number of elements in the remainder sets is at most $n/k2^{k-1}$. Therefore, the total contribution of the remainder sets to the potential function is at most $O(1/k2^{k-1})$, which is negligible if we choose the initial k large enough.

Hence, in each iteration, the potential function increases by $\Omega\left(\varepsilon^5\right)$, and since the potential is bounded by 1, after $O\left(1/\varepsilon^5\right)$ steps Algorithm 5.1.3 terminates. In each step the size of the partition grows exponentially, and hence letting $k(\varepsilon, t)$ be a tower of height $O(1/\varepsilon^5)$ proves the Theorem. Gowers proved that this tower dependence on ε is necessary [Gow97]. □

The proof of the Regularity Lemma given above is not quite constructive, because it does not tell us how to find the witness sets S_{ij} efficiently. However, there exists a polynomial time algorithm which computes the partition promised by the Regularity Lemma [ADL$^+$94].

6

Szemeredi's Theorem for $k = 3$: Roth's proof

Luca Trevisan

Scribe(s): Aditya Bhaskara, Aravindan Vijayaraghavan

Summary: We will look at Roth's proof of Szemeredi's theorem for the case $k = 3$. We first prove the theorem for subsets of $\mathbb{F}_p^n$ instead of $\{1, 2, \ldots, N\}$ as it is cleaner and illustrates the idea of the proof. Then we prove the theorem for integers, and finally look at a construction of Behrend of an AP-free subset of $[n]$ of size $\frac{n}{2^{\Theta(\sqrt{\log n})}}$.

6.1 Introduction

Szemeredi's theorem states that any subset of $\{1, 2, \ldots, N\}$ of positive *density* contains arbitrarily long arithmetic progressions. More formally, given a positive integer k and $\delta > 0$, there exists an N_0 such that for all $N > N_0$, any $S \subseteq [N]$ with $|S| > \delta N$ contains a k-term AP.

In this lecture, we will prove the result for $k = 3$, using techniques from Fourier analysis. The proof presented here is due to Roth. To illustrate the idea of the proof, we will first prove the result over $\mathbb{F}_p^n$ instead of $\{1, 2, \ldots, N\}$.

6.2 Roth's theorem over $\mathbb{F}_p^n$

Suppose p is a fixed prime. We will prove the following '$\mathbb{F}_p^n$ analogue' of the theorem.

Theorem 6.2.1. *Suppose $\delta > 0$. Given a set $A \subseteq \mathbb{F}_p^n$, with $|A| \geq \delta |\mathbb{F}_p^n|$, there exist $a, b \in \mathbb{F}_p^n$ such that $a, a + b, a + 2b \in A$.*

In fact, we prove something stronger – that there exists a constant c_p depending only on p such that the conclusion holds whenever $|A| > \frac{c_p p^n}{n}$. Consider the characteristic function of A. We abuse notation a little and denote it also by A (which maps $\mathbb{F}_p^n \to \mathbb{C}$). So $A(x) = 1$ if $x \in A$ and 0 otherwise.

6.2.1 Outline of the Proof

The idea now is to analyze the Fourier coefficients of the function A. The rough outline of the proof is the following. We prove that either the set A has 'lots' of APs (in the sense that given random $a, b \in \mathbb{F}_p^n$, the probability that $a, a+b, a+2b$ are in A is positive), or that A has a density at least $\delta + \frac{\delta^2}{4}$ on an *affine subspace* of $\mathbb{F}_p^n$ of dimension $(n-1)$.

Now if the second case happens, we look only at points on that subspace, and since it's a hyperplane, arithmetic progressions map back to arithmetic progressions so it suffices to prove that 3-APs exist in this set of points. So repeating the process, we either find several 3-APs or move to a still smaller space (with the given set of points having an even higher density), and so on. The density however obviously cannot keep growing (it is, afterall, bounded by 1) so this process has to stop.

In one 'step', the density increases from $\delta \to \delta + \frac{\delta^2}{4}$. This gives an obvious bound of $O\left(\frac{1}{\delta^2}\right)$ on the number of steps. However a slightly more careful argument shows that we need at most $O\left(\frac{1}{\delta}\right)$ steps. Note that in $\frac{4}{\delta}$ steps the density increases from δ to at least $\delta(1+\frac{\delta}{4})^{\frac{4}{\delta}} > 2\delta$. And hence in the next $\frac{4}{2\delta} = \frac{2}{\delta}$ steps the density increases from 2δ to 4δ, and so on. Thus in $\frac{4}{\delta}(1+\frac{1}{2}+\frac{1}{4}+\cdots+\frac{1}{2^t}) < \frac{8}{\delta}$ steps the density increases to 1 (for appropriate t).

6.2.2 Definitions

We start with a few definitions required for the proof. Given two functions $f, g : \mathbb{F}_p^n \to \mathbb{C}$, define the inner product as $\langle f, g\rangle = \mathbb{E}_x[f(x)\overline{g(x)}]$. We now construct an orthonormal basis for the space of functions $f : \mathbb{F}_p^n \to \mathbb{C}$. Note for now that this space has dimension p^n.

Given a point $t \in \mathbb{F}_p^n$, define $\chi_t : \mathbb{F}_p^n \to \mathbb{C}$ as $\chi_t(x) = \omega^{t_1x_1+t_2x_2+\cdots+t_nx_n} = \omega^{t.x}$, where $\omega = e^{2i\pi/p}$, a primitive pth root of unity. Note that $\langle \chi_t, \chi_t\rangle = 1$ for all $t \in \mathbb{F}_p^n$. Further if $s, t \in \mathbb{F}_p^n$ and $s \neq t$, then

$$\langle \chi_s, \chi_t\rangle = \mathbb{E}_x[\omega^{(s-t).x}] = \prod_i \mathbb{E}_{x_i}[\omega^{(s_i-t_i)x_i}] = 0$$

The last equality is because $(s_i - t_i) \neq 0$ for some i, and for this i we have $\mathbb{E}_{x_i}[\omega^{(s_i-t_i)x_i}] = 0$. Thus the set $\{\chi_t : t \in \mathbb{F}_p^n\}$ is an orthonormal set of functions. Since there are p^n of them it is an orthonormal basis for functions $f : \mathbb{F}_p^n \to \mathbb{C}$.

6.2.3 The Proof

We now formally state and prove what was stated in the outline above.

Theorem 6.2.2. *Suppose $A \subseteq \mathbb{F}_p^n$ has size $\delta|\mathbb{F}_p^n|$. Then one of the following holds*

1. *There are at least $\frac{\delta^3}{2}|\mathbb{F}_p^n|^2 - |A|$ arithmetic progressions in A.*
2. *There exists a subspace H of $\mathbb{F}_p^n$ of dimension $(n-1)$ such that the density of A on H is at least $\delta + \frac{\delta^2}{4}$.*

Proof. Writing the function A in the basis defined above, we have

$$A(x) = \sum_{t\in\mathbb{F}_p^n} \widehat{A}(t)\chi_t(x)$$

where $\widehat{A}(t) = \mathbb{E}_x[A(x)\overline{\chi_t(x)}]$. Thus by the assumption on the size of the set A, we have $\widehat{A}(\mathbf{0}) = \delta$. Also, we have $\sum_t |\widehat{A}(t)|^2 = \langle A, A\rangle = \mathbb{E}_x[|A(x)|^2] = \delta$, the last equality holding because A is a $0-1$ function.

We will now consider the quantity $\mathbb{E}_{x,y}[A(x)A(x+y)A(x+y+y)]$ (denoted E from now). Clearly, this denotes the fraction of the total possible 'arithmetic progressions' that are contained in A.

$$\begin{aligned} E &= \mathbb{E}_{x,y}\Big[\big(\sum_a \widehat{A}(a)\chi_a(x)\big)\big(\sum_b \widehat{A}(b)\chi_b(x+y)\big)\big(\sum_c \widehat{A}(c)\chi_c(x+y+y)\big)\Big] \\ &= \sum_{a,b,c} \widehat{A}(a)\widehat{A}(b)\widehat{A}(c)\mathbb{E}_{x,y}\big[\chi_a(x)\chi_b(x+y)\chi_c(x+2y)\big] \\ &= \sum_{a,b,c} \widehat{A}(a)\widehat{A}(b)\widehat{A}(c)\mathbb{E}_x\big[\chi_a(x)\chi_b(x)\chi_c(x)\big]\mathbb{E}_y\big[\chi_b(y)\chi_c(y)\chi_c(y)\big] \\ &= \sum_{a,b,c} \widehat{A}(a)\widehat{A}(b)\widehat{A}(c)\mathbb{E}_x\big[\chi_{a+b+c}(x)\big]\mathbb{E}_y\big[\chi_{b+2c}(y)\big] \end{aligned}$$

In between, we used the fact that the χ's are characters, i.e., $\chi_t(x+y) = \chi_t(x)\chi_t(y)$. The expectations over x, y are both non-zero iff $a+b+c=0$ and $b+2c=0$, i.e., $c=a$ and $b=-2a$. And in this case the expectations are 1. Thus

$$\begin{aligned} E &= \sum_a \widehat{A}(a)^2\widehat{A}(-2a) \\ &= \delta^3 + \sum_{a\neq\mathbf{0}} \widehat{A}(a)^2\widehat{A}(-2a) \end{aligned}$$

Now define $M = \max_{a\neq\mathbf{0}} |\widehat{A}(a)|$. By the above, we have $E = \delta^3 + \sum_{a\neq\mathbf{0}} \widehat{A}(a)^2\widehat{A}(-2a) \geq \delta^3 - M\sum_a |\widehat{A}(a)|^2 = \delta^3 - \delta M$.

If $M \leq \frac{\delta^2}{2}$, we have $E \geq \delta^3/2$, thus a constant fraction of all the possible AP's are actually in the set A. So the number of 3-APs in A is at least $\frac{\delta^3}{2}|\mathbb{F}_p^n|^2 - |A|$ (the $|A|$ is due to the trivial APs), and we are done.

Thus suppose $M > \frac{\delta^2}{2}$, so there exists an $a \neq \mathbf{0}$ such that $|\widehat{A}(a)| = |\mathbb{E}_x\big[A(x)\omega^{-a.x}\big]| > \frac{\delta^2}{2}$.The aim in this case is to show that there exists a $c \in \mathbb{F}_p$ such that $\mathbb{E}_{\{x|a.x=c\}}[A(x)] > \delta + \frac{\delta^2}{4}$. We have $|\mathbb{E}_x[A(x)\omega^{-a.x}]| \geq \frac{\delta^2}{2}$. Observe that if $a \neq \mathbf{0}$, $a.x$ is uniformly distributed in $\{0, 1, \ldots, (p-1)\}$. Thus we have

$$\Big|\frac{1}{p}\mathbb{E}_{\{x|a.x=0\}}[A(x)\omega^0] + \frac{1}{p}\mathbb{E}_{\{x|a.x=1\}}[A(x)\omega^1] + \cdots + \frac{1}{p}\mathbb{E}_{\{x|a.x=(p-1)\}}[A(x)\omega^{p-1}]\Big| > \frac{\delta^2}{2}$$

Also if we put $B(x) = A(x) - \delta$, so that $\mathbb{E}_x[B(x)] = 0$, we have $\widehat{B}(a) = \widehat{A}(a)$ for $a \neq 0$, thus the above equation holds if A is replaced by B. Thus, using triangle inequality, we get

$$\frac{1}{p}\Big|\mathbb{E}_{\{x|a.x=0\}}[B(x)]\Big| + \frac{1}{p}\Big|\mathbb{E}_{\{x|a.x=1\}}[B(x)]\Big| + \cdots + \frac{1}{p}\Big|\mathbb{E}_{\{x|a.x=(p-1)\}}[B(x)]\Big| > \frac{\delta^2}{2} \tag{6.1}$$

If we denote $\mathbb{E}_{\{x|a.x=i\}}[B(x)]$ by α_i, then since $\mathbb{E}_x[B(x)] = 0$, we have $\alpha_0 + \alpha_1 + \cdots + \alpha_{p-1} = 0$. Eqn.6.1 says that $\frac{1}{p}(|\alpha_0| + |\alpha_1| + \cdots + |\alpha_{p-1}|) > \frac{\delta^2}{2}$. These together imply that there exists an i with $(\alpha_i + |\alpha_i|) > \frac{\delta^2}{2}$. Thus for this i, α_i is positive and is $> \delta^2/4$.

Thus there exists a c such that $\mathbb{E}_{\{x|a.x=c\}}A(x) > \delta + \frac{\delta^2}{4}$. This completes the proof. □

Till now we had assumed that $|\mathbb{F}_p^n|$ is 'large enough' (say we denote this by N). One can see that for the quantity in condition (1) of Theorem 6.2.2 to be positive we need $N > \frac{2}{\delta^2}$. Further the dimension of the space must be at least $\frac{8}{\delta}$. Thus it suffices to have $N(= p^n) > p^{\frac{8}{\delta}}$, equivalently $|A| > \frac{8}{\log p}(p^n/n)$, as we claimed earlier.

6.3 Roth's theorem over $\mathbb{Z}$

In this section we will prove Roth's theorem. In particular,

Theorem 6.3.1. *Suppose $\delta > 0$. Then there is an absolute constant C such that for all $N > 2^{2^{\frac{C}{\delta}}}$, any $A \subseteq \{0, 1, \ldots, N-1\}$ of size $|A| = \delta N$ necessarily contains a non-trivial arithmetic progression of length* 3.

6.3.1 Outline and Definitions

The proof closely resembles the proof of Roth's theorem over $\mathbb{F}_p^n$. We show an analogous result – either the set A has lots of APs of length 3 or there exists an arithmetic progression P of size $\Omega(\sqrt{N})$ such that A has a density of at least $\delta + \frac{\delta^2}{64}$ on P. We study the problem initially over $\mathbb{Z}_N$, performing Fourier analysis on this group, and later see how we can move to $\mathbb{Z}$.

As before, if the first case above does not occur, we turn our attention to the subset $A' = A \cap P$ of A and iterate this argument. This can be done because as before, APs map back to APs in the original set. Thus we do this process until we either find several 3-APs or until the density of A' on an AP is > 1 (which is not possible). Notice that in this case arithmetic progressions in some sense play the role of hyperplanes.

For clarity, we will redefine some of the terms we will use for the $\mathbb{Z}_N$ case. Given two functions $f, g : \mathbb{Z}_N \to \mathbb{C}$, define the inner product as $\langle f, g\rangle = \mathbb{E}_x[f(x)\overline{g(x)}]$. The space of functions $f : \mathbb{Z}_N \to \mathbb{C}$ clearly has dimension N. For $k \in \mathbb{Z}_N$, consider the function $\chi_k : \mathbb{Z}_N \to \mathbb{C}$ defined by $\chi_k(x) = \omega^{kx}$, where $\omega = e^{\frac{2i\pi}{k}}$, a primitive Nth root of unity.

As before it's easy to see that these functions form an orthonormal basis for functions $f : \mathbb{Z}_N \to \mathbb{C}$. So any such f can be written as $f(x) = \sum_{k=0}^{N-1} \widehat{f}(k)\chi_k(x)$, where $\widehat{f}(k) = \mathbb{E}_x[f(x)\overline{\chi_k(x)}]$. We will call $\widehat{f}(k)$ the kth Fourier coefficient of f. Note also that clearly $|\widehat{f}(k)| \leq \mathbb{E}_x|f(x)|$

Another simple identity which will be quite useful is that $\widehat{f}(k)\overline{\widehat{g}(k)}$ is the Fourier transform of the function $G(x) = \mathbb{E}_y[f(y)\overline{g(y-x)}]$, and hence

$$|\widehat{f}(k)||\widehat{g}(k)| \leq \mathbb{E}_x|\mathbb{E}_y[f(y)\overline{g(y-x)}]| \tag{6.2}$$

This equality has the following crucial consequence. Suppose $A(x)$ and $B(x)$ are the characteristic functions of two sets $A, B \subseteq \mathbb{Z}_N$ (abusing notation as before). Then if $|\widehat{A}(k)|$ and $|\widehat{B}(k)|$ are both 'large', then there must exist an x such that $|\mathbb{E}_y[f(y)\overline{g(y-x)}]|$ is 'large'. But this quantity is precisely the cardinality of $A \cap (B + x)$. So this means that if A and B have a common 'large' Fourier coefficient, A has a 'large' intersection with a translate of B.

Later on we will prove that for any r, there exists an AP of size $\Omega(\sqrt{N})$ which has a 'large' rth Fourier coefficient. This along with the above (and with suitable meanings for 'large') will imply the main result.

6.3.2 The Proof

As outlined above, we will prove the following

Theorem 6.3.2. *Let $A \subseteq \{0, 1, \dots, N\}$ with $|A| = \delta N$ and $N \geq \frac{8}{\delta^2}$. Then one of the following holds*

1. *A contains at least $\delta^3 N^2 - |A|$ non-trivial APs of length 3.*
2. *There exists an arithmetic progression P of length $|P| > \frac{1}{256}\delta^2\sqrt{N}$ such that A has a density at least $\delta + \delta^2/64$ on P.*

We first consider this problem over $\mathbb{Z}_N$. As before, A is the characteristic function of set A and we see that $\widehat{A}(0) = \delta$ and $\sum_t |\widehat{A}(k)|^2 = \delta$. Again, we consider the quantity $E = \mathbb{E}_{x,y}[A(x)(x + y)A(x + y + y)]$, which signifies the fraction of total possible length-3 progressions (over $\mathbb{Z}_N$) in A. As in the $\mathbb{F}_p^n$ case,

$$\begin{aligned} E &= \mathbb{E}_{x,y}A(x)A(x+y)A(x+y+y) = \sum_{k=0}^{N-1} \widehat{A}(k)^2\widehat{A}(-2k) \\ &= \delta^3 + \sum_{k=1}^{N-1} \widehat{A}(k)^2\widehat{A}(-2k) \end{aligned}$$

Define $M = \max_{k\neq 0} |\widehat{A}(k)|$. We have,

$$\begin{aligned} E &= \delta^3 + \sum_{k\neq 0} \widehat{A}(k)^2\widehat{A}(-2k) \\ &\geq \delta^3 - M\sum_{k\neq 0} \widehat{A}(k)^2 = \delta^3 - \delta M \end{aligned}$$

Case 1: $M < \frac{\delta^2}{8}$

As before, if $M < \delta^2/2$, $E \geq \frac{\delta^3}{2}$ so a constant fraction of all the possible APs in $\mathbb{Z}_N$ are in set A (including the trivial ones). However we need to count the number of APs in $\mathbb{Z}$ instead of APs in $\mathbb{Z}_N$. If $x, x + y + y \in A \cap [N/3, 2N/3)$, then $x, x + y, x + y + y$ are also APs in $\mathbb{Z}$. Hence, the following proposition follows easily by the application of Cauchy-Schwarz inequality and Plancherel's identity.

Proposition 6.3.3. *If $M = \max_{k\neq 0} |\widehat{A}(k)| < \frac{\delta^2}{8}$, and if $|A \cap [\frac{N}{3}, \frac{2N}{3})| \geq \frac{\delta N}{4}$, then $E \geq \frac{\delta^3}{32}$.*

However, it is to be noted that this includes δN trivial length-3 APs, and since $N \geq \frac{8}{\delta^2}$, there exists non-trivial length-3 APs in this case.

Further, if $|A \cap [N/3, 2N/3)| < \frac{\delta N}{4}$, then either $|A \cap [0, N/3)|$ or $|A \cap [2N/3, N)| \geq \frac{9\delta}{8}(\frac{N}{3})$. Hence, in this case, there is an AP (over $\mathbb{Z}$) P of length $|P| \geq N/3$ such that $|A \cap P| \geq (\delta + \frac{\delta}{8})|P|$.

Case 2: $M \geq \delta^2/8$

We show in this case that A has increased density on an AP of large size. We prove this by finding a long arithmetic progression P whose fourier transform is large at the some k where the fourier transform of A is also large, so that A has increased density on a translate of P.

Lemma 6.3.4. *For any k such that $1 \leq k \leq N-1$, there exists an arithmetic progression P_1 of length $\geq \sqrt{N}/4$ and common difference d such that $|\widehat{P_1}(k)| \geq \frac{|P_1|}{2N}$ and $|P_1|d < N$.*

Proof. Given k, partition $[0, N-1]^2$ into $\lceil \sqrt{N} - 1 \rceil^2$ equal squares. Consider the points $\{(0,0), (1,k), \ldots, (N-1, (N-1)k)\}$. By pigeon hole principle, we see that there exist $l, m \in [0, N-1]$, with $l < m$, such that $d = m - l \leq \sqrt{N}$ and $kd \leq \sqrt{N}$ (mod N). We now claim that the required arithmetic progression $P_1 = \{\ldots, -2d, d, 0, d, 2d, \ldots\}$ with length $|P_1| = \lfloor \sqrt{N}/\pi \rfloor$.

Now $|\widehat{P_1}(k) - \frac{|P_1|}{N}| \leq |\mathbb{E}_x[P_1(x)(e^{-\frac{2i\pi}{N}xk} - 1)]$. Now for $x \in P_1$, we have $\frac{2\pi}{N}xk \leq \frac{2\pi}{\sqrt{N}}t \leq 1$. Bounding the segment length by the arc length, we get $|\widehat{P_1}(k) - \frac{|P_1|}{N}| \leq \frac{|P_1|}{2N}$, thus proving the result. □

We now show that the given set has a higher density over a translate of one of the APs constructed above.

Lemma 6.3.5. *Suppose $|\widehat{A}(k)| \geq \epsilon$ for some $k \in \mathbb{Z}_N$. Let P an arithmetic progression in $\mathbb{Z}_N$ such that $\widehat{P}(k) \geq \frac{|P|}{2N}$. Then there exists an x such that $|A \cap (P + x)| \geq (\delta + \frac{\epsilon}{4})|P|$ (i.e., it has a higher density over the progression).*

Proof. As before, define $B(x) = A(x) - \delta$. We have $\widehat{B}(0) = 0$ and $\widehat{B}(k) = \widehat{A}(k)$ for $k \neq 0$. Now note that

$$|A \cap (P + x)| \geq (\delta + \frac{\epsilon}{4})|P| \iff \mathbb{E}_y[B(y)P(y-x)] \geq \frac{\epsilon}{4N}|P| \tag{6.3}$$

Thus it suffices to prove that there exists x such that $\mathbb{E}_y[B(y)P(y-x)] \geq \frac{\epsilon}{4N}|P|$.
Let $G(x) = \mathbb{E}_y[B(y)P(y-x)]$. Using Eqn.6.2, we get

$$\mathbb{E}_x|G(x)| \geq |\widehat{G}(k)| \geq \frac{\epsilon}{2N}|P|$$

Since $G(x)$ has a mean value of 0,

$$\mathbb{E}_x[|G(x)| + G(x)] \geq \frac{\epsilon}{2N}|P|$$

Therefore there exists an x such that $\mathbb{E}_y[B(y)P(y-x)] \geq \frac{\epsilon}{4N}|P|$, as desired. □

From the above lemmas, and since $M \geq \delta^2/8$, we have an arithmetic progression P' on $\mathbb{Z}_N$ of length $\geq \sqrt{N}/4$ such that$|A \cap P'| \geq (\delta + \delta^2/32)|P'|$. We now need to construct a long arithmetic progression P over $\mathbb{Z}$ from P', with increased density of A (compared to $[0, N-1]$).
We know that P' has a common difference d such that $d|P'| < N$ and hence it can be split into two arithmetic progressions P_1 and P_2 in $\mathbb{Z}$. Let $|P_1| \leq |P_2|$. If $|P_1| \leq \frac{\delta^2}{64}|P'|$, then $|A \cap P_2| \geq (\delta + \delta^2/32)|P'| - |P_1| \geq (\delta + \delta^2/64)|P'|$ ($P = P_2$ here). Else, A has density $\geq (\delta + \delta^2/64)$ in at least one of P_1 or P_2 (this being the required P). This completes the proof of Theorem 6.3.2.

To complete the proof of Roth's theorem, we might have to consider $\frac{C}{\delta}$ 'steps' to get to a density bigger than 1 (thus obtaining the contradiction). Thus we need N to be such that $N^{(\frac{1}{2})^{C/\delta}}$ is at least a constant, or equivalently, $\delta > \frac{C'}{\log\log N}$.

6.4 Behrend's Construction

We now show the construction of an AP-free subset of $[n]$ of size $\Omega(n^{1-\epsilon})$ for all $\epsilon > 0$. The construction is due to Behrend (1940's and is still unbeaten).

Theorem 6.4.1. *(Behrend) There exists an AP-free subset of* $[n]$ *of size at least* $\frac{n}{2^{c\sqrt{\log n}}}$, *for some absolute constant* c.

Proof. Consider points in $\mathbb{R}^d$ such that $x_1^2 + \cdots + x_d^2 = m$ and $0 \leq x_i \leq k$. There exists some $m \leq dk^2$ for which there are at least k^d/dk^2 solutions. Fix this as the choice of m and call the set of points S.

Now we consider a natural encoding of a point $(x_1, x_2, \ldots, x_d)$ into $\mathbb{Z}$. Each of the coordinates is at most k. Thus look at $f : (x_1, \ldots, x_d) \to x_1 + (2k+1)x_2 + \cdots + (2k+1)^{d-1}x_d$. Now if $f(\mathbf{x}) + f(\mathbf{y}) = 2f(\mathbf{z})$, we must have $\mathbf{x} + \mathbf{y} = 2\mathbf{z}$, which does not happen for the chosen points as they lie on a sphere.

Thus $\{f(\mathbf{x}) : \mathbf{x} \in S\}$ is a 3-AP-free set of k^d/dk^2 integers each at most $(2k+1)^d$. Given n, pick $d = \sqrt{\log n}$, $(2k+1) = 2^{\sqrt{\log n}}$, then $k^d/dk^2 = n/2^{\Theta(\sqrt{\log n})}$ which is what we wanted. □

7

Gowers Uniformity Norms and Sketch of Gowers' proof of Szemerédi's Theorem

Luca Trevisan

Scribe(s): Indraneel Mukherjee, David Steurer

Summary: We give a sketch of Gowers' proof of Szemerédi's theorem using the Gowers Uniformity Norms for the case $k > 3$.

7.1 Introduction

Szemerédi's theorem states that any subset A of a cyclic group $\mathbb{Z}_N$ of prime order contains an arithmetic progression of length k, as long as the density of A in $\mathbb{Z}_N$ is at least $\delta_{N,k}$, where $\delta_{N,k} = o(1)$ for fixed k and $N \to \infty$. The theorem implies that a set $A \subseteq \mathbb{Z}$ of positive density in $\mathbb{Z}$ contains arbitrarily long arithmetic progressions. Szemerédi's theorem has also been generalized to arbitrary finite additive groups.

Randomness vs structure. A common theme in the proofs of Szemerédi's theorem is the dichotomy between randomness and structure. The main step in Gowers' proof, for example, is to show that every subset A of $\mathbb{Z}_N$ satisfies at least one of the following two conditions:

C1 The set A contains, up to constant factors, as many arithmetic progressions of length k as a randomly chosen set of the same density as A.

C2 There exists an arithmetic progression $P \subseteq \mathbb{Z}_N$ of length $N^{\Omega(1)} - 2^{\delta^{-O(1)}}$ such that A has significantly higher density in P than $\mathbb{Z}_N$, specifically $\delta' > \delta + \delta^{O(1)}$, where $\delta' = |A \cap P|/|P|$ and $\delta = |A|/N$.

Assuming that every set satisfies C1 or C2, one can show Szemerédi's theorem by a simple induction. The resulting quantitive bound on $\delta_{N,k}$ would be $\delta_{N,k} = (\log\log N)^{-\Omega_k(1)}$.

We say that a set is *k-pseudorandom* if it satisfies condition C1 for arithmetic progressions of length k.

Fourier-analytic Approach. Gowers' approach is inspired by Roth's Fourier-analytic proof of Szemerédi's theorem for $k = 3$ (Roth's theorem). For a subset A of a cyclic group $\mathbb{Z}_N$ of prime order, let $\hat{A}\colon \mathbb{Z}_N \to \mathbb{C}$ denote the Fourier transform of the characteristic function $\mathbf{1}_A$ of A. Further, let $\delta = \hat{A}(0)$ be the density of A in $\mathbb{Z}_N$. We say that a set A is *linearly uniform* if $|\hat{A}(x)| < c\delta^2$ for all $x \neq 0$. The two main steps of Roth's proof are as follows.

- Uniformity implies pseudorandomness: If the set A is linearly uniform then it is 3-pseudorandom.
- Non-uniformity implies density increment: If A is not linearly uniform, then A is significantly denser in an arithmetic progression of $\mathbb{Z}_N$ with parameters as in C2.

In Roth's proof, linear uniformity is used to distinguish between sets that are 3-pseudorandom and sets that allow a density increment argument.

In Section 7.2, we will see that the first step of Roth's proof cannot be carried out for $k > 3$, that is, we show that a linearly uniform set need not be k-pseudorandom for $k > 3$. This example shows that if we want to follow the structure of Roth's proof to show Szemerédi's theorem for $k > 3$, we need a new notion of uniformity that allows us to distinguish between k-pseudorandom sets and sets that allow a density increment argument.

Gowers Uniformity. One of the main innovations of Gowers' proof of Szemerédi's theorem are the Gowers uniformity norms $U^d(f)$ for functions $f\colon \mathbb{Z}_N \to \mathbb{C}$ and integers $d > 0$. We will define these norms in Section 7.3. We say that a set $A \subseteq \mathbb{Z}_N$ is k*-uniform* if $U^{k-1}(\mathbf{1}_A - \delta) < c_k\delta^k$. With this notion of uniformity, Gowers' proof of Szemerédi's theorem has the same structure as the Fourier-analytic proof of Roth's theorem.

- Uniformity implies pseudorandomness: If A is k-uniform then it is k-pseudorandom.
- Non-uniformity implies density increment: If A is not k-uniform, then A is significantly denser in an arithmetic progression P of $\mathbb{Z}_N$ with parameters as in C2.

Organization. In these notes, we present a sketch of Gowers' proof of Szemerédi's theorem. We will focus on the case of arithmetic progressions of length 4 and the ambient group $\mathbb{F}_p^n$ for a fixed prime $p > 3$.

In Section 7.4, we discuss the first step of Gowers' proof, that uniformity implies pseudorandomness. Most proofs in this section directly translate to the case of $k > 4$ and arbitrary ambient groups.

In Section 7.5, we show the second step of Gowers' proof, that non-uniformity implies density increment. This step is considerably more difficult than the first one. The presented proof assumes a strong inverse theorem for the Gowers uniformity norm. A corresponding theorem for $k > 4$ is not known. The presented proof also exploits the structure of the ambient group $\mathbb{F}_p^n$. In this sense, the discussion in Section 7.5 does not easily generalize to larger k or arbitrary ambient groups.

7.2 Linear Uniformity vs Pseudorandomness

In this section, we show that linearly uniform sets need not be k-pseudorandom for $k \geq 4$.

Lemma 7.2.1. *The quadratic surface*

$$A = \Big\{ x \in \mathbb{F}_p^n \mid \sum_{i \le n/2} x_i x_{n/2+i} = 0 \in \mathbb{F}_p \Big\}$$

is linearly uniform, but contains a $\Theta(\delta^3)$ fraction of all arithmetic progressions of length k (k-APs) of $\mathbb{F}_p^n$ for every $k \ge 3$, where $\delta \approx 1/p$ is the density of A. In particular, A is not k-pseudorandom for $k \ge 4$.

Proof. Suppose we can show that A is linearly uniform and has density $\delta \approx 1/p$. Then Roth's theorem shows that A contains $\Theta(\delta^3 |F_p^n|^2)$ arithmetic progressions of length 3. Since A is a quadratic surface, every line of $\mathbb{F}_p^n$ either intersects A in at most 2 points[1] or is completely contained in A. Hence every 3-AP contained in A extends to a k-AP for any k. Therefore, A contains a $\Theta(\delta^3)$ fraction of all k-APs of $\mathbb{F}_p^n$ for every $k \ge 3$.

We will write a vector $x \in \mathbb{F}_p^n$ as (u, v) where $u, v \in \mathbb{F}_p^{n/2}$ form the first and last $n/2$ coordinates, respectively of x.

We first compute the density of A as

$$\begin{aligned} \hat{A}(0) &= \Pr_{u,v}(\langle u, v\rangle = 0) \\ &= \Pr_u(u = 0) + \Pr_u(u \ne 0) \Pr_{u,v}(\langle u, v\rangle = 0 | u \ne 0) \\ &= p^{-n/2} + (1 - p^{-n/2})\tfrac{1}{p} = 1/p + 2^{-\Omega(n)}. \end{aligned}$$

We now estimate the non-zero Fourier coefficients of A. We will write $v \perp u$ to denote that $\langle v, u\rangle = 0$. Note that $A = \{(u, v) \mid u \perp v\}$.

The Fourier-transform of A at point $c = (c_1, c_2) \ne 0$ has value

$$\hat{A}(c) = \mathbb{E}_x[\mathbf{1}_A(x) \cdot \omega^{\langle c, x\rangle}] = \mathbb{E}_u\Big[\omega^{\langle c_1, u\rangle} \Pr_v(v \perp u) \mathbb{E}_v[\omega^{\langle c_2, v\rangle} \mid u \perp v]\Big] \tag{7.1}$$

where ω is a primitive p-th root of unity. If $c_2 = 0 \ne c_1$ then the above is equal to

$$\mathbb{E}_u \omega^{\langle c_1, u\rangle} \Pr(v \perp u) = \tfrac{1}{p}\mathbb{E}_u \omega^{\langle c_1, u\rangle} + (1 - \tfrac{1}{p}) \Pr(u = 0) = (1 - \tfrac{1}{p}) p^{-n/2} < p^{-n/2}$$

where we derived the second equality using the fact that ω is a primitive root and the expression $\langle c_1, u\rangle$ takes all values in $\mathbb{F}_p$ equally often as u varies over $\mathbb{F}_p^{n/2}$.

So assume $c_2 \ne 0$. In case $u \nparallel c_2$, the random variable $\langle c_2, v\rangle$ is uniformly distributed over $\mathbb{F}_p$ when v is chosen uniformly from the subspace orthogonal to u. Hence,

$$\forall u \nparallel c_2 : \quad \mathbb{E}_v[\omega^{\langle c_2, v\rangle} \mid u \perp v] = \tfrac{1}{p} \sum_{i=0}^{p-1} \omega^i = 0$$

Therefore, only the vectors u parallel to c_2 have non-zero contribution to the expectation in Equation 7.1 and we can upper-bound $|\hat{A}(c)| \le \Pr_u(u \parallel c_2) = 1/p^{n/2-1}$.

□

[1] Using a parametrization $\ell = \{x + \lambda y; \lambda \in \mathbb{F}_p\}$ of the line, the intersection of ℓ and a quadratic surface corresponds to the solution set of a polynomial equation $q(x + \lambda y) = 0$ that is at most quadratic in λ.

7.3 Gowers uniformity norm

The proof of Roth's theorem shows that the only subsets of $\mathbb{F}_p^n$ that are not 3-pseudorandom are sets that are correlated with a function $\omega^{a(x)}$ for a degree-1 polynomial a over $\mathbb{F}_p$, where the correlation is measured by the (absolute) value of the inner product

$$\mathbb{E}_{x \leftarrow_{\mathrm{R}} \mathbb{F}_p^n}[\mathbf{1}_A \cdot \omega^{a(x)}]$$

In the last section, we saw that linearly uniform quadratic surfaces are not 4-pseudorandom. This phenomenon suggests that a uniformity norm that measures distance from 4-pseudorandomness should be large if the set is strongly correlated with a polynomial of degree 2. A natural choice for such a norm would be the maximum correlation with a quadratic polynomial. It is however difficult, though not impossible, to relate the maximum correlation to the number of 4-APs in the set.

Instead the Gowers uniformity norm estimates the correlation with polynomials of a certain degree in a more indirect way, which then allows to relate that norm relatively easily to the number of arithmetic progressions of a certain length in the set.

For a function $f\colon \mathbb{F}_p^n \to \mathbb{C}$ and a vector $y \in \mathbb{F}_p^n$, the *derivative* in direction y is the function $D_y f\colon \mathbb{F}_p^n \to \mathbb{C}$ with

$$D_y f(x) = f(x)\overline{f(x+y)}.$$

Note that D_y acts as a difference operator in the exponent if we write $f(x) = \omega^{g(x)}$ for a p-th root of unity ω. Hence, if g is a degree-d polynomial in x over $\mathbb{F}_p$, then D_y reduces the degree of the polynomial in the exponent by at least one, e.g. $D_y(\omega^{x_1 x_2}) = \omega^{-y_1 x_2 - x_1 y_2 + y_1 y_2}$, where the exponent is now linear in x. So if we apply the difference operator three times on a quadratic polynomial, the resulting function is the constant function $\omega^0 = 1$.

The idea behind the Gowers uniformity norm is that instead of measuring the maximum correlation of f with an unknown degree-$(d-1)$ polynomial, we can as well measure the expected correlation of $D_{y_1 \cdots y_d} f$ with the constant function $\omega^0 = 1$, where the directions $y_1, \ldots, y_d$ are chosen uniformly at random from $\mathbb{F}_p^n$ and $D_{y_1 \cdots y_d}$ denotes the composition $D_{y_1} \circ D_{y_2} \circ \ldots \circ D_{y_d}$ of difference operators.

Definition 7.3.1. The degree-d Gowers uniformity norm $U^d(f)$ of a function $f\colon G \to \mathbb{C}$, where G is a finite group, is defined as

$$U^d(f) \overset{def}{=} (\mathbb{E}_{x, y_1, \ldots, y_d \leftarrow_{\mathrm{R}} G} D_{y_1 \cdots y_d} f(x))^{1/2^d}$$

Notice the expected value of a random derivative $\mathbb{E}_{x,y} D_y f(x) = |\mathbb{E}_x f(x)|^2$ is always a non-negative real number. So we also have,

$$\mathbb{E}_{x, y_1, \ldots, y_d} D_{y_1 \cdots y_d} f = \mathbb{E}_{y_2, \cdots, y_d} \mathbb{E}_{x, y_1} D_{y_1}(D_{y_2 \cdots y_d} f(x)) = \mathbb{E}_{y_2, \cdots, y_d} |\mathbb{E}_x D_{y_2 \cdots y_d} f(x)|^2 \geq 0$$

which shows that the Gowers norm is well-defined and non-negative.

The Gowers norm also satisfies the triangle inequality. However, since we do not use the triangle inequality here, we omit its proof.

We will need the following two simple facts, which can be verified easily.

Fact 7.3.2.

$$U^k(f)^{2^k} = \mathbb{E}_y\left[U^{k-1}(D_y f)^{2^{k-1}}\right]$$

Fact 7.3.3.

$$U^1(f) = |\mathbb{E}_x f(x)|$$

Inverse theorems. We argued heuristically that any set that is correlated with a degree-$(d-1)$ polynomial should have a large degree-d Gowers uniformity norm. In fact, the maximum correlation with a degree-$(d-1)$ polynomial is always a lower bound on the degree-d Gowers norm. Since we do not need it here, we omit the proof of this fact.

It is conjectured that correlation with a polynomial (on some subspace) is in fact the only obstruction to uniformity.

So far, such an "inverse theorem" is only known for the degree-3 Gowers norm.

Theorem 7.3.4 (Inverse theorem for U^3)**.** *Suppose $f : \mathbb{F}_p^n \to \mathbb{C}$ is a function that takes values whose magnitudes are bounded by 1 everywhere. Let ω denote a p-th root of unity. If $U^3(f) > \eta$, then there exists a subspace W of dimension at least $n - \eta^{-O(1)}$, and, for each coset $y + W$, a quadratic polynomial $q_y(x)$ over $\mathbb{F}_p$ defined on $y + W$, such that*

$$\mathbb{E}_{y \leftarrow_R \mathbb{F}_p^n} |\mathbb{E}_{x \leftarrow_R y+W} f(x)\omega^{q_y(x)}| = \Omega(\eta^{O(1)})$$

Assuming this theorem, it is relatively easy to carry out the second step of Gowers' strategy for the proof of Szemerédi's theorem for $k = 4$, that is, to show that non-uniformity implies that the set has higher density on some affine subspace of low codimension (cf. Section 7.5, Proposition 7.5.1).

7.4 Uniformity implies pseudorandomness

In order to study the pseudorandomness of a set with respect to arithmetic progressions of length k, we introduce the following degree-k form on functions $f \colon \mathbb{F}_p^n \to \mathbb{C}$,

$$\Lambda_k(f) = \mathbb{E}_{x,y \leftarrow_\mathrm{R} \mathbb{F}_p^n} f(x)f(x+y)f(x+2y)\cdots f(x+(k-1)y) \tag{7.2}$$

For a set $A \subseteq \mathbb{F}_p^n$, the value of $\Lambda_k(\mathbf{1}_A)$ gives the fraction of k-APs of $\mathbb{F}_p^n$ that are completely contained in A. For a random set A of density δ, we expect that it contains $\delta^k|\mathbb{F}_p^n|$ arithmetic progressions of length k and hence $\mathbb{E}_A \Lambda_k(\mathbf{1}_A) = \delta^k$. Hence, for a k-pseudorandom set we would require $\Lambda_k(\mathbf{1}_A) = \Theta(\delta^k)$.

It will be more convenient to consider Λ_k on the normalized function $f = \mathbf{1}_A - \delta$ instead of $\mathbf{1}_A$. The following lemma relates $\Lambda_k(f)$ and $\Lambda_k(\mathbf{1}_A)$ for the case $k = 4$ using the Fourier-spectrum of f.

Lemma 7.4.1.

$$\left|\Lambda_4(f) - (\Lambda_4(\mathbf{1}_A) - \delta^4)\right| \leq 4\delta^2 \|\hat{f}\|_\infty \tag{7.3}$$

Proof. Notice that $\hat{f}(0) = \mathbb{E}_x[f] = 0$.

For any x, y, z, w, we have

$$\begin{aligned}\mathbf{1}_A(x)\mathbf{1}_A(y)\mathbf{1}_A(z)\mathbf{1}_A(w) - \delta^4 &= (f(x)+\delta)(f(y)+\delta)(f(z)+\delta)(f(w)+\delta) - \delta^4 \\ &= f(x)f(y)f(z)f(w) + \delta \sum f(u)f(v)f(t) \\ &\quad + \delta^2 \sum f(u)f(v) + \delta^3 \sum f(u)\end{aligned}$$

where the summations are over distinct u, v, t belonging to $\{x, y, z, w\}$. If x, y, z, w are consecutive points of a random 4-AP in $\mathbb{F}_p^n$, then they are pairwise independent and uniformly distributed. Hence the third and fourth summations on the right side disappear in expectation. Further, the magnitude of the expectation of the second sum can be bounded, as in the proof of Roth's theorem, by $4\delta^2\|\hat{f}\|_\infty$. □

If $\|\hat{f}\|_\infty = \Omega(\delta^2)$ then, as in the proof of Roth's theorem, we can find an affine subspace of low codimension, in which the set A has significantly higher density than δ. To prove Szemerédi's theorem for $k = 4$, we could then apply the induction hypthesis on the restriction of A to that subspace.

So we may assume $\|\hat{f}\|_\infty \ll \delta^2$. In this case, the set A is 4-pseudorandom if and only if $|\Lambda_4(f)| \ll \delta^4$. The next lemma will show that $|\Lambda_4(f)|$ can be upperbounded by the degree-3 Gowers uniformity norm of f. Hence the set A is 4-pseudorandom if $U^3(f) \ll \delta^4$.

Lemma 7.4.2 (Generalized von Neumann theorem)**.** *For numbers $c_0, \ldots, c_{k-1} \in \mathbb{F}_p$ that are pairwise distinct and functions $g_0, g_1, \ldots, g_{k-1} \colon \mathbb{F}_p^n \to \mathbb{C}$ with $\|g_i\|_\infty \leq 1$, we have*

$$\left|\mathbb{E}_{x,d}\Big[g_0(x+c_0d)g_1(x+c_1d)\cdots g_{k-1}(x+c_{k-1}d)\Big]\right| \leq U^{k-1}(g_0)$$

where $x, d \leftarrow_R \mathbb{F}_p^n$.

In particular for $c_i = i$, $k \leq p$ and $f = g_0 = g_1 = \ldots = g_{k-1}$, we get

$$\Lambda_k(f) \leq U^{k-1}(f).$$

Using the previous lemma and the fact $U^{d-1}(f) \leq U^d(f)$ whose proof we omit here, one can show the following lemma which essentially generalizes Lemma 7.4.1 to larger k.

Lemma 7.4.3.

$$|\Lambda_k(\mathbf{1}_A) - \delta^k| = O_k(U^{k-1}(f))$$

The lemma implies that there exists a constant c_k such that a subset A of $\mathbb{F}_p^n$ is k-pseudorandom if $U^{k-1}(\mathbf{1}_A - \delta) \leq c_k\delta^k$. We say that sets A with $U^{k-1}(\mathbf{1}_A - \delta) \leq c_k\delta^k$ are *k-uniform.* This establishes the first step of Gowers' proof of Szemerédi's theorem, that uniformity implies pseudorandomness.

Proposition 7.4.4. *Every k-uniform subset of $\mathbb{F}_p^n$ is k-pseudorandom.*

Note that our notion of k-uniformity actually implies that there is no correlation of the set with a polynomial of degree $k-2$. Therefore, k-uniformity is usually refered to as uniformity of order $k-2$ [TV06]. This also explains the term linear uniformity used to refer to sets with $|\hat{f}| \ll \delta^2$, which turns out to be equivalent to $U^2(f) \ll \delta^3$, our condition for 3-uniformity.

In the remainder of the section, we present a proof of the generalized von Neumann theorem.

Proof of Lemma 7.4.2. Let $\Lambda_{k,\mathbf{c}}(g_0,\dots,g_{k-1})$ denote the k-linear form $\mathbb{E}_{x,d}[\prod_i g_i(x+c_i d)]$ that we try to relate to the Gowers uniformity norm.

We induce on k. For $k=2$, note that the random variables $x+c_0d$ and $x+c_1d$ are pairwise independent since $c_0 \neq c_1$ in $\mathbb{F}_p$. Hence, the bilinear form evaluates to $\Lambda_{2,\mathbf{c}}(g_0,g_1) = \mathbb{E}_x g_0(x)\mathbb{E}_{x'} g_1(x')$. By Fact 7.3.3, we can then verify

$$|\Lambda_{2,\mathbf{c}}(g_0,g_1)| = |\mathbb{E}_x g_0(x)\mathbb{E}_{x'} g_1(x')| \leq |\mathbb{E}_x g_0(x)| = U^1(g_0)$$

For $k > 1$, we will bound the form $\Lambda_{k,\mathbf{c}}$ by the expected value of a $(k-1)$-linear form $\Lambda_{k-1,\bar{\mathbf{c}}}$ over a certain distribution of arguments. Specifically, we have the following claim.

Claim 7.4.4.1.

$$|\Lambda_{k,\mathbf{c}}(g_0,\dots,g_{k-1})|^2 \leq |\mathbb{E}_y \Lambda_{k-1,\bar{\mathbf{c}}}(D_y g_0, \bar{g}_1,\dots,\bar{g}_{k-2})|$$

where $\bar{c}_0,\dots,\bar{c}_{k-2}$ are pairwise distinct numbers in $\mathbb{F}_p$ and $\bar{g}_1,\dots,\bar{g}_{k-2}\colon \mathbb{F}_p^n \to \mathbb{C}$ are random functions depending on the variable y such that $\Pr_y(\|\bar{g}_i\|_\infty \leq 1) = 1$.

Assuming Claim 7.4.4.1, we can end the proof of the current lemma as follows.

$$\begin{aligned}
|\Lambda_{k,\mathbf{c}}(g_0,\dots,g_{k-1})|^2 &\leq |\mathbb{E}_y \Lambda_{k-1,\bar{\mathbf{c}}}(D_y g_0, \bar{g}_1,\dots,\bar{g}_{k-2})| && \text{(Claim 7.4.4.1)}\\
&\leq (\mathbb{E}_y |\Lambda_{k-1,\bar{\mathbf{c}}}(D_y g_0, \bar{g}_1,\dots,\bar{g}_{k-2})|^{2^{k-2}})^{1/2^{k-2}} && \text{(Hölder)}\\
&\leq (\mathbb{E}_y U^{k-2}(D_y g_0)^{2^{k-2}})^{1/2^{k-2}} && \text{(induction hypothesis)}\\
&= (U^{k-1}(g_0)^{2^{k-1}})^{1/2^{k-2}} = U^{k-1}(g_0)^2 && \text{(Fact 7.3.2)}
\end{aligned}$$

Proof Sketch of Claim 7.4.4.1. We will prove the claim for the case $k=4$. The proof for general k is completely analogous.

We want to bound the following 4-linear form

$$\Lambda_{3,\mathbf{c}}(g_0,\dots,g_3) = \mathbb{E}_{x,d}\ g_0(x+c_0d)g_1(x+c_1d)g_2(x+c_2d)g_3(x+c_3d) \tag{7.4}$$

In order to eliminate g_3, we decouple the argument of g_3 from the variable d by substituting $\bar{x}$ for $x+c_3d$. Since the distribution of $\bar{x}$ is uniform, we can rewrite (7.4) as

$$\begin{aligned}
\Lambda_{3,\mathbf{c}}(g_0,\dots,g_3) &= \mathbb{E}_{\bar{x},d}\ g_0(\bar{x}+\bar{c}_0d)g_1(\bar{x}+\bar{c}_1d)g_2(\bar{x}+\bar{c}_2d)g_3(\bar{x})\\
&= \mathbb{E}_{\bar{x}} g_3(\bar{x})\ \mathbb{E}_d g_0(\bar{x}+\bar{c}_0d)g_1(\bar{x}+\bar{c}_1d)g_2(\bar{x}+\bar{c}_2d)
\end{aligned}$$

where $\bar{c}_i = c_i - c_3$.

Using $\|g_3\|_\infty \leq 1$, we eliminate g_3 and arrive at

$$\begin{aligned}
|\Lambda_{3,\mathbf{c}}(g_0,\dots,g_3)| &\leq \mathbb{E}_{\bar{x}}\ \left|\mathbb{E}_d\ g_0(\bar{x}+\bar{c}_0d)g_1(\bar{x}+\bar{c}_1d)g_2(\bar{x}+\bar{c}_2d)\right|\\
&\leq \left(\mathbb{E}_{\bar{x}}\ \left|\mathbb{E}_d\ g_0(\bar{x}+\bar{c}_0d)g_1(\bar{x}+\bar{c}_1d)g_2(\bar{x}+\bar{c}_2d)\right|^2\right)^{1/2} && \text{(Cauchy-Schwarz)}\\
&= \left(\mathbb{E}_{\bar{x}}\ \mathbb{E}_{d,d'} D_{\bar{c}_0(d'-d)}g_0(\bar{x}+\bar{c}_0d)\ D_{\bar{c}_1(d'-d)}g_1(\bar{x}+\bar{c}_1d)\ D_{\bar{c}_2(d'-d)}g_2(\bar{x}+\bar{c}_2d)\right)^{1/2}
\end{aligned}$$

Since $\bar{c}_0 \neq 0 \in \mathbb{F}_p$, the random variable $y = \bar{c}_0(d'-d)$ is uniformly distributed. Substituting y for $\bar{c}_0(d'-d)$, we obtain

$$|\Lambda_{3,\mathbf{c}}(g_0,\dots,g_3)| \leq \left(\mathbb{E}_y\ \mathbb{E}_{\bar{x},d}\ D_y g_0(\bar{x}+\bar{c}_0d)\ D_{y\bar{c}_1/\bar{c}_0}g_1(\bar{x}+\bar{c}_1d)\ D_{y\bar{c}_2/\bar{c}_0}g_2(\bar{x}+\bar{c}_2d)\right)^{1/2}$$

Since each $\bar{g}_i = D_{y\bar{c}_1/\bar{c}_0} g_1$ is bounded by 1 for all y and the numbers $\bar{c}_i = c_i - c_3 \in \mathbb{F}_p$ are pairwise distinct, we get as desired

$$|\Lambda_{3,\mathbf{c}}(g_0, \dots, g_3)| \leq \left(\mathbb{E}_y \; \mathbb{E}_{\bar{x},d} \; D_y g_0(\bar{x} + \bar{c}_0 d) \; \bar{g}_1(\bar{x} + \bar{c}_1 d) \; \bar{g}_2(\bar{x} + \bar{c}_2 d)\right)^{1/2}$$

□

7.5 Non-uniformity implies density increment

Assuming the inverse theorem for U^3, we can carry out the density increment argument for sets that are not 4-uniform.

Proposition 7.5.1 (Non-uniformity implies density increment). *Let $A \subseteq \mathbb{F}_p^n$ be set of density δ.*

If $U^3(\mathbf{1}_A - \delta) > \eta$ then there exists an affine subspace of dimension at least $n/2 - \eta^{-O(1)}$ on which A has density $\delta + \Omega(\eta^{O(1)})$.

Using Proposition 7.4.4 and Proposition 7.5.1, one can show Szemerédi's theorem for the groups $\mathbb{F}_p^n$ and $k = 4$ by a simple induction, similar to the induction in Roth's proof for $k = 3$.

There are generalizations of Proposition 7.5.1 known that allow to proof Szemerédi's theorem for groups $\mathbb{Z}_N$ and larger k. It is interesting that such a generalization for $k > 4$ could be proved, since a corresponding inverse theorem for the norm U^{k-1} is not know.

The proof of Proposition 7.5.1 goes in two steps. First, we show that for any non-uniform set A there exists an affine subspace W of large dimension such that A has significantly higher density in a quadratic surface of W. For this step we are using the inverse theorem for U^3.

Second, we show that any quadratic surface of a vector space over $\mathbb{F}_p$ can be paritioned into affine subspaces of large dimension. Hence in one of those affine subspaces, the set A must have density as least as large as in the quadratic surface.

Lemma 7.5.2. *If $U^3(\mathbf{1}_A - \delta) > \eta$, then there exists an affine subspace $W \subseteq \mathbb{F}_p^n$ of dimension $n - \eta^{-O(1)}$ and a quadratic polynomial q such that*

$$\mathbb{E}_{x \leftarrow_R S} \mathbf{1}_A(x) \geq \delta + \eta^{O(1)}$$

where S is the quadratic surface $\{x \in W \mid q(x) = 0\}$.

Proof. Using the inverse theorem, we know there exists a subspace W of dimension $n - 1/\epsilon$, where $\epsilon = \Omega(\eta^{O(1)})$, for which the following holds. Partition $\mathbb{F}_p^n$ into cosets $\{W_y\}$ of W. Then for each coset W_y, there exists a quadratic function q_y such that the average correlation $\mathbb{E}_y |\mathbb{E}_{W_y} f(x) \omega^{q_y(x)}|$ is at least ϵ.

Let $S_{yz} \subseteq W_y$ be the quadratic surface $\{x \in W_y \mid q_y(x) = z\}$. Note that the collection $\{S_{yz}\}_{yz}$

forms a partition of $\mathbb{F}_p^n$ that refines the partition $\{W_y\}_y$. We can thus write,

$$\begin{aligned}\epsilon \le \mathbb{E}_y|\mathbb{E}_{W_y} f(x)\omega^{q_y(x)}| &\le \mathbb{E}_y \sum_z \left|\mathbb{E}_{W_y} f \cdot \mathbf{1}_{S_yz}\right| && \text{(triangle inequality)}\\ &= \sum_y \Pr_{\mathbb{F}_p^n}(W_y) \sum_z \left|\mathbb{E}_{W_y} f \cdot \mathbf{1}_{S_yz}\right| \\ &= \sum_y \Pr_{\mathbb{F}_p^n}(W_y) \sum_z \Pr_{W_y}(S_{yz})|\mathbb{E}_{S_{yz}} f| \\ &= \sum_{yz} \Pr_{\mathbb{F}_p^n}(S_{yz})|\mathbb{E}_{S_{yz}} f|\end{aligned}$$

Since $\sum_{yz} \Pr_{\mathbb{F}_p^n}(S_{yz})\mathbb{E}_{S_{yz}} f = \mathbb{E}_{\mathbb{F}_p^n} f = 0$, we have

$$\epsilon \le \sum_{yz} \Pr_{\mathbb{F}_p^n}(S_{yz})|\mathbb{E}_{S_{yz}} f| = 2 \sum_{yz\in I^+} \Pr_{\mathbb{F}_p^n}(S_{yz})\mathbb{E}_{S_{yz}} f$$

where the second sum is only over the pairs y, z with $\mathbb{E}_{S_{yz}} f \ge 0$, denoted by I^+.

From the above, we conclude by an averaging argument that there exists a pair y, z such that

$$\mathbb{E}_{S_{yz}} \mathbf{1}_A = \delta + \mathbb{E}_{S_{yz}} f \ge \delta + \epsilon/2.$$

The affine subspace W_y and the quadratic polynomial $q = q_y - z$ are as desired by the lemma. □

Note that by translating the set A, we can assume that the subspace W from Lemma 7.5.2 is a linear subspace.

Lemma 7.5.3. *Any quadratic surface $S = \{x \in W \mid q(x) = 0\}$ of a vector space W over $\mathbb{F}_p$ can be partitioned into affine subspaces, each of dimension at least* $\dim(W)/2 - 5/2$.

Proof. We can write q as $q(x) = \langle x, Mx\rangle + \langle a, x\rangle + b$, where $M\colon W \to W$ is a symmetric linear operator, $a \in W$, and $b \in \mathbb{F}_p$. Let $Q(x) = \langle x, Mx\rangle$ denote the quadratic form given by M.

Let U be the linear subspace of W spanned by all $x \in W$ with $Q(x) = 0$.

Observe that q restricted to a coset $y + U$ is an affine linear function, as for every $u \in U$,

$$Q(y + u) = Q(y) + 2\langle u, My\rangle + Q(u) = Q(y) + 2\langle u, My\rangle.$$

Let $\ell_y\colon W \to \mathbb{F}_p$ be an affine linear function that agrees with q on $y + U$. Note that $S \cap (y + U)$ is equal to the intersection of $y + U$ and the affine hyperplane $\{x \in W \mid \ell_y(x) = 0\}$. Hence $S \cap (y + U)$ is either empty or an affine subspace of $y + U$ of dimension at least $\dim(U) - 1$.

In order to prove the lemma it remains to show that $\dim(U) \ge \dim(W)/2 - 3/2$.

Note that for every pair $x, y \in U$, the inner product $\langle y, Mx\rangle = 0$ vanishes, since over a field of characteristic $p > 2$,

$$0 = \tfrac{1}{2}Q(x + y) = \tfrac{1}{2}(Q(x) + 2\langle y, Mx\rangle + Q(y)) = \langle y, Mx\rangle.$$

Hence M maps U into the orthogonal complement of U.

Let $U' = \{y \in W \mid \forall x \in U\colon \langle y, Mx\rangle = 0\}$ be the orthogonal complement of $M(U) \subseteq W$. Note that U is a subspace of U'. In order to conclude $\dim(U) \ge \dim(W)/2 - 3/2$, we will show that the U' is not much larger than U.

Claim 7.5.3.1. $\dim(U') < \dim(U) + 3$

Assuming the claim, we can finish the proof as follows. The dimension of $M(U)$ can be at most the dimension of U. Also, the dimension of W is equal to $\dim(M(U)) + \dim(U') \leq \dim(U) + \dim(U')$. By Claim 7.5.3.1, we then conclude $\dim(W) < 2\dim(U) + 3$.

Proof of Claim 7.5.3.1. For sake of contradiction, assume that U' has dimension at least $\dim(U) + 3$. Then the quotient vector space U'/U contains three linearly independent vectors $x_1 + U$, $x_2 + U$, and $x_3 + U$.

Consider the quadratic polynomial $Q' \in \mathbb{F}_p[a_1, a_2, a_3]$ obtained by the substitution $Q' = Q(a_1x_1 + a_2x_2 + a_3x_3)$. By the Chevalley-Warning theorem, the number of solutions to $Q'(a) = 0$ is a multiple of the characteristic. Since $Q'(0) = 0$, the number of solutions is not zero, and hence there exists at least one other solution $a = (a_1, a_2, a_3) \neq 0$.

By construction, the vector $x' = a_1x_1 + a_2x_2 + a_3x_3$ is in the kernel of Q, that is, $Q(x') = 0$. Furthermore, since x_1, x_2, x_3 are linearly independent in U'/U, the vector x' is not contained in U. This contradicts the fact that U is the span of vectors in the kernel of Q.

□

8

Applications: Direct Product Theorems

Avi Wigderson

Scribe(s): Shubhangi Saraf

Summary: We show how Gowers Uniformity can be used to obtain XOR lemmas for correlation with low degree $GF(2)$ polynomials.

8.1 Introduction

A basic computer science question is to decide if a function f is in a given complexity class $\mathcal{C}$. This question in some sense deals with "worst-case" complexity. When we deal with average-case complexity, we consider the following question: for a function f that's not contained in $\mathcal{C}$, how well does $\mathcal{C}$ approximate f?

Consider a boolean valued function $f : \{0,1\}^n \to \{-1,+1\}$. Though we're restricting our attention to $\mathbb{F}_2^n$, many of the results proved in this lecture can be extended to $\mathbb{F}_p^n$. Throughout the lecture we'll also restrict our attention to the uniform distribution on inputs.

Consider $\max_{c\in\mathcal{C}} \Pr_{x\in\{0,1\}^n}[c(x) = f(x)]$. This quantity is a measure of how well f is approximated by $\mathcal{C}$. Note that if it equals 1, then f is exactly computed by some member of $\mathcal{C}$.

Another related measure of agreement is the correlation.

$$\mathrm{Corr}(f,\mathcal{C}) \doteq \max_{c\in\mathcal{C}} |E_x\left[f(x)c(x)\right]| = \max_{c\in\mathcal{C}} |\Pr_x\left[f(x) = c(x)\right] - \Pr_x\left[f(x) \neq c(x)\right]|.$$

Observe that $\mathrm{Corr}(f,\mathcal{C}) \in [0,1]$. Clearly when $\mathrm{Corr}(f,\mathcal{C}) = 1$ then f fully correlates with a member of $\mathcal{C}$, and when $\mathrm{Corr}(f,\mathcal{C}) = 0$ then f is completely uncorrelated. In fact, when $\mathrm{Corr}(f,\mathcal{C})$ is close to 0, one can use this information to build pseudorandom generators. Hence pseudorandomness is one reason to study correlation. Another reason is that if it is known that a function f has low correlation with a complexity class, then this information can be used to prove non-inclusion in a stronger complexity class. For instance, if for some complexity class $\mathcal{C}$, $\mathrm{Corr}(f,\mathcal{C}) < 1/t$, then $f \notin \mathcal{C}'$, where $\mathcal{C}'$ is class of functions that are obtained by taking the majority of t functions of $\mathcal{C}$.

If a function f has low correlation with a complexity class, it is in some sense "hard" to compute by that complexity class. Functions with extremely low correlation demonstrate certain pseudorandom properties and hence are of interest. It would be useful to have a method by which when given a function that is hard to compute, we can produce a function that is even harder to compute.

Hardness Amplification: Suppose $\mathrm{Corr}(f, \mathcal{C}) \leq \alpha < 1$ and we want to amplify the "hardness". One idea for this amplification which we hope might work is to take the exclusive OR of f on many independent copies.

$$f^{\oplus t}(y_1, y_2, \ldots, y_t) = \prod_{i=1}^{t} f(y_i)$$

where $y_i \in \{0,1\}^n$. If f took values in $\{0,1\}$ where $+1$ and -1 are identified 0 and 1 respectively, then this product is equivalent to taking the parity of f on all the x_i. The hope is that $\mathrm{Corr}(f^{\oplus t}, \mathcal{C}) \leq \alpha^t$. A big question is the following: for what $\mathcal{C}$ can this hope be materialized?

Let $\mathcal{C}_0$ denote the class of all constant functions. In this case there is perfect exponential decay, since for any function f,

$$\mathrm{Corr}(f^{\oplus t}, \mathcal{C}_0) = |E\left[f^{\oplus t}\right]| = |(E\left[f\right])|^t = \mathrm{Corr}(f, \mathcal{C}_0)^t.$$

Yao's XOR Lemma [Yao82](whose first proof appears in [Lev85]) gives a highly nontrivial complexity class for which something similar is true. Let $\mathcal{C} = P/poly$, which is the class of all polynomial sized circuits. Then,

$$\forall t, \mathrm{Corr}(f^{\oplus t}, P/poly) \leq \mathrm{Corr}(f, P/poly)^t + \frac{1}{n}.$$

The problem is that we don't have any explicit lower bounds for the above class or else we could use it to amplify the hardness. In the rest of the lecture we'll discuss correlation with low degree polynomials over $GF(2)$ and prove some hardness amplification results.

8.2 Low degree polynomials over $GF(2)$

Let P_d denote the class of all degree (at most) d polynomials in $GF(2)$ over any set of variables $x_1, x_2, \ldots, x_n$. We can assume that the polynomials take values in $\{-1, +1\}$ by raising -1 to the $\{0,1\}$ value of the polynomial. For example, $(-1)^{x_1 \oplus x_3} \in P_1$ and $(-1)^{x_1 x_2 \oplus x_3} \in P_2$.

P_d is an important and well studied complexity class. AC^0 denotes the class of functions computable by constant depth, polynomial sized circuits with unbounded fanin AND and OR gates. A number of AC^0 bounds were proved in a series of results. It was proved that PARITY $\notin AC^0$. Now, suppose we add PARITY to AC^0, i.e. in addition to AND and OR, we also allow $Mod\,2$ gates. Call the resulting complexity class $AC^0[2]$. Razborov [Raz87] proved that MAJORITY $\notin AC^0[2]$. The proof of the result consists of two parts and both parts deal with the correlation of functions to complexity classes of polynomials.

1. $\mathrm{Corr}(AC^0[2], P_{(\log n)^{O(1)}}) \geq 1 - \frac{1}{n}$.

2. $\mathrm{Corr}(MAJORITY, P_{\sqrt[3]{n}}) \leq \frac{1}{\sqrt{n}}$.

Both the above results are nontrivial. It easily follows that MAJORITY $\notin AC^0[2]$, since if it was, then by the first part $\mathrm{Corr}(MAJORITY, P_{(\log n)^{O(1)}}) \geq 1 - \frac{1}{n}$. However this contradicts the second part.

A big open question is to find an explicit function f for which $\mathrm{Corr}(f, P_{\log n}) \leq \frac{1}{n}$. Razborov showed that MAJORITY is an example of a function f for which $\mathrm{Corr}(f, P_{\log n}) \leq \frac{1}{\sqrt{n}}$, and there has been no improvement since then.

In the rest of the lecture we'll obtain hardness amplification results for the class P_d and use it to give explicit functions that have low correlation with P_d.

8.3 XOR lemma for correlation with low degree polynomials over $GF(2)$

Most of the results here appeared in [Vio06], and are extended and put in a more general context in [VW07]. The main result of this section is the following result: Let $f : \{0,1\}^n \to \{-1,+1\}$ be a function. If $\mathrm{Corr}(f, P_d) \leq 1 - 2^{-d}$, then $\forall t, \mathrm{Corr}(f^{\oplus t}, P_d) \leq \exp(-\frac{t}{4^d})$.

To prove XOR lemmas, the basic hope is to find a norm on boolean functions that in some sense captures its correlation with a given complexity class. Informally, we'd like the norm N to have the following properties:

1. $\forall f : \{0,1\}^n \to \mathbb{R}, N(f) \in \mathbb{R}$.
2. for every function f, $N(f) \approx \mathrm{Corr}(f, \mathcal{C})$.
3. If f and g are two functions on disjoint inputs, then $N(fg) = N(f)N(g)$.

Once we have such a norm, up to the "$\approx$", we get the XOR lemma in a trivial way, since

$$\mathrm{Corr}(f^{\oplus t}, \mathcal{C}) \approx N(f^{\oplus t}) = N(F)^t \approx \mathrm{Corr}(f, \mathcal{C})^t.$$

The question is - where do we get such norm? [VW07] show that when $\mathcal{C} = P_{d-1}$, then Gower's degree-d norm U_d satisfies all the previously mentioned properties. The norm was introduced by Gowers in [Gow98] and [Gow01], and also independently by [AKK$^+$03].

Let $\overline{Y} = y_1, y_2, \ldots y_d \in \{0,1\}^n$. We define the *cube* $C(\overline{Y})$ to be the following *multiset* of points spanned by $y_1, \ldots, y_d$.

$$C(\overline{Y}) = \left\{ \bigoplus_{j \in S} y_j | S \subseteq [d] \right\}.$$

Let $f : \{0,1\}^n \to \{-1,+1\}$ be a function. The degree-d norm of f is defined as

$$U_d(f) = E_{x,y_1,\ldots y_d}\left[\big(\prod_{z \in C(\overline{Y})} f(x \oplus z)\big) \right].$$

Note that since $C(\overline{Y})$ is actually a multiset, this may also be written as

$$U_d(f) = E_{x,y_1,\ldots y_d \in \{0,1\}^n}\left[\prod_{S \subseteq [d]} f\left(x \oplus \bigoplus_{j \in S} y_j \right) \right].$$

Note that $(\prod_{z\in C(\overline{Y})} f(x\oplus z))$ is shifting the cube by x and multiplying all the values on the cube, which is like taking the XOR of all the values if $1 \rightarrow 0$ and $-1 \rightarrow 1$.

The functional U_d can be defined more generally for complex valued functions. In the definition of U_d, all the terms in the product that correspond to the sum of an odd number of elements of $\overline{Y}$ are conjugated. This ensures that U_d is real valued.

For $z \in C(\overline{Y})$, let S_z denote the number of elements of $\overline{Y}$ that are summed to obtain z. For a complex number w, let $w^{\underline{i}}$ be w if i is even, and it's conjugate $\overline{w}$ if i is odd.

Let $f : \{0,1\}^n \rightarrow \mathbb{C}$.

$$U_d(f) = E_{x,y_1,\dots y_d\in\{0,1\}^n}\left[\prod_{z\in C(\overline{Y})} f(x\oplus z)^{\underline{S_z}}\right].$$

In most of the results that follow, it will suffice to consider boolean valued functions. However, for the result in Section 8.4 it will be convenient to consider the more general case of complex valued functions.

Proposition 8.3.1. *If $p \in P_{d-1}$, then for all x and $\overline{Y} = \langle y_1, y_2, \dots, y_d\rangle \in (2^n)^d$, $(\prod_{z\in C(\overline{Y})} p(x\oplus z)) = 1$.*

Proof. First note that if $\overline{Y}$ does not span a d dimensional space, then each member of $C(\overline{Y})$ appears in $C(\overline{Y})$ an even number of times and hence $(\prod_{z\in C(\overline{Y})} p(x\oplus z)) = 1$. If $\overline{Y}$ does span a d dimensional space, then there is a linear transformation that maps $\overline{Y} = \langle y_1, y_2, \dots, y_d\rangle$ to $\langle e_1, e_2, \dots e_d\rangle$, where $e_i \in \{0,1\}^n$ is the vector with a 1 only in the i^{th} coordinate. Such a linear transformation maps a degree d polynomial p to another degree d polynomial say p', and if the result is true for p' and $\langle e_1, e_2, \dots e_d\rangle$, it also holds for p and $\langle y_1, y_2, \dots, y_d\rangle$. Hence, we only need to consider $\overline{Y} = \langle e_1, e_2, \dots e_d\rangle$. We observe that it suffices to prove the result when p is a monomial. Let $p = \prod_{i\in I} x_i$, where $I \subseteq [n]$ with $|I| \leq d-1$. Pick $j \in [d] \setminus I$. Note that for all vectors y, $p(y + e_j) = p(y)$. Hence we can pair the terms in the product $(\prod_{z\in C(\overline{Y})} p(x\oplus z))$ such that p takes the same value on the members of each pair. This proves the result.

Another way of viewing this result is to observe that $U_d(f)$ is like taking the average of d derivatives of f, each of which reduces the degree of f by 1. First observe that if p is a polynomial of degree d in n variables, then $\forall y \in \{0,1\}^n$, $p(x\oplus y)\times p(x)$ is a polynomial of degree at most $d-1$ as any variable that occurs with even degree in a monomial is redundant. This is analogous to taking the first derivative of p in the direction specified by y. On the same lines, for $y_1, \dots, y_k \in \{0,1\}^n$, we could define the derivative of order k with respect to $y_1, \dots, y_k$ as $\left[\prod_{z\in C(\overline{Y})} p(x\oplus z)\right]$, where $C(\overline{Y})$ is the cube of $y_1, \dots y_k$. This is a polynomial of degree at most $d-k$, and the result easily follows. □

8.3.1 The property testing perspective

It is not difficult to show that the condition in Proposition 8.3.1 actually characterizes P_{d-1}, i.e. $p \in P_{d-1}$ if and only if for all x and $\overline{Y}$, $(\prod_{z\in C(\overline{Y})} p(x\oplus z)) = 1$. This characterization suggests a very natural way of testing whether a given function f is a degree d polynomial or very far from it. Consider the following property tester for P_{d-1}:

Pick $x, y_1, y_2, \dots, y_d$ at random. If $(\prod_{z\in C(\overline{Y})} f(x\oplus z)) = 1$, then accept. Otherwise reject.

Observe that $U_d(f)$ completely captures the probability that the tester accepts. $U_d(f) = \Pr[\text{tester accepts}] - \Pr[\text{tester rejects}]$. By Proposition 8.3.1, if $f \in P_{d-1}$ then the tester accepts with probability 1.

It was shown by [AKK$^+$03] that if $\mathrm{Corr}(f, P_{d-1}) \leq 1 - 2^{-d}$ then $U_d(f) \leq 1 - 4^{-d}$. This result begins to give some indication of how well the functional U_d captures correlation. The following is a high level sketch of their proof: By Proposition 8.3.1 we know that when f is a polynomial of degree $d-1$, if we restrict to any d dimensional subspace, the parity of f evaluated on the subspace is 0. This can be used to "predict" the value of f at a point of the d dimensional subspace by knowing the values of f at the rest of the points in the subspace. For any function f, define a function g such that for a point y, $g(y)$ is the majority over all subspaces containing y of the "predicted" value of y on that subspace. [AKK$^+$03] show that when the property tester mentioned above accepts with sufficiently high probability (more specifically if $U_d(f) \geq 1 - 4^{-d}$) then the function g is close to f ($\mathrm{Corr}(f, g) \geq 1 - 2^{-d}$). Moreover they also show that $g \in P_{d-1}$. Together, they imply that if $\mathrm{Corr}(f, P_{d-1}) \leq 1 - 2^{-d}$ then $U_d(f) \leq 1 - 4^{-d}$.

8.3.2 Properties of Gowers norm

Lemma 8.3.2. *Let $f : \{0,1\}^n \to \{-1,+1\}$ and $g : \{0,1\}^{n'} \to \{-1,+1\}$ be two functions. Define $(f \otimes g) : \{0,1\}^n \times \{0,1\}^{n'} \to \{-1,+1\}$ by $(f \otimes g)(x,y) = f(x) \cdot g(y)$. Then $U_d(f \otimes g) = U_d(f) \cdot U_d(g)$.*

Proof. The proof follows immediately from the definitions. □

Lemma 8.3.3. *Let $f : \{0,1\}^n \to \mathbb{C}$ be a function. For all functions p that belong to P_{d-1}, $U_d(f \cdot p) = U_d(f)$. Here $(f \cdot p)$ is defined to be the product of f and p on the same input.*

Proof. $U_d(f \cdot p) = E_{x,y_i,\ldots y_d}\left[(\prod_{z \in C(\overline{Y})} f(x \oplus z))\right]\left[(\prod_{z \in C(\overline{Y})} p(x \oplus z))\right]$. By Proposition 8.3.1, for all x and $\overline{Y}$ we have that $\left[(\prod_{z \in C(\overline{Y})} p(x \oplus z))\right] = 1$. Hence $U_d(f \cdot p) = E_{x,y_i,\ldots y_d}\left[(\prod_{z \in C(\overline{Y})} f(x \oplus z))\right] = U_d(f)$. □

Lemma 8.3.4. $\forall f : \{0,1\}^n \to \mathbb{C}$, *$Corr(f, P_{d-1}) \leq U_d(f)^{\frac{1}{2^d}}$.*

Proof. [1] We first show that $\forall g : \{0,1\}^n \to \mathbb{C}, E[g]^{2^d} \leq U_d(g)$. The proof is essentially the Cauchy–Schwarz inequality applied $d-1$ times.

[1] Taken from [VW07].

$$\begin{aligned}
U_d(f) &= E_{x,y_1,\dots y_d \in \{0,1\}^n}\left[\prod_{z\in C(\overline{Y})} f(x\oplus z)^{\underline{S_z}}\right]\\
&= E_{y_1,\dots y_{d-1} \in \{0,1\}^n}\left[E_{x,y_d}\left[\prod_{z\in C(\overline{Y}-y_d)} f(x\oplus z)^{\underline{S_z}}\cdot f(x\oplus z\oplus y_d)^{\underline{S_z+1}}\right]\right]\\
&= E_{y_1,\dots y_{d-1} \in \{0,1\}^n}\left[E_x\left[\prod_{z\in C(\overline{Y}-y_d)} f(x\oplus z)^{\underline{S_z}}\right]\overline{\cdot E_x\left[\prod_{z\in C(\overline{Y}-y_d)} f(x\oplus z)^{\underline{S_z}}\right]}\right]\\
&= E_{y_1,\dots y_{d-1} \in \{0,1\}^n}\left[|E_x\left[\prod_{z\in C(\overline{Y}-y_d)} f(x\oplus z)^{\underline{S_z}}\right]|^2\right]\\
&\geq |E_{y_1,\dots y_{d-1} \in \{0,1\}^n}\left[E_x\left[\prod_{z\in C(\overline{Y}-y_d)} f(x\oplus z)^{\underline{S_z}}\right]\right]|^2\\
&= U_{d-1}(f)^2
\end{aligned}$$

Now, $E\left[g\right] = |E_x\left[g(x)\right]| = \sqrt{U_1(g)} \leq U_2(g)^{1/2^2} \leq \ldots \leq U_d(g)^{1/2^d}$. Hence $E\left[g\right]^{2^d} \leq U_d(g)$.

For $p \in P_{d-1}$, $\mathrm{Corr}(f,p) = |E_x\left[(f\cdot p)(x)\right]|$. But $|E_x\left[(f\cdot p)(x)\right]| = U_1(f\cdot p)^{1/2} \leq U_2(f\cdot p)^{1/2^2} \leq \ldots \leq U_d(f\cdot p)^{1/2^d} = U_d(f)^{1/2^d}$. This completes the proof of the lemma. □

Conclusion of proof of XOR lemma

To conclude, if $\mathrm{Corr}(f, P_{d-1}) \leq 1 - 2^{-d}$, then

$$\mathrm{Corr}(f^{\oplus t}, P_{d-1}) \leq U_d(f^{\oplus t})^{1/2^d} = U_d(f)^{t/2^d} \leq (1-1/4^d)^{t/2^d} \leq exp(-t/8^d).$$

The first inequality follows from Lemma 8.3.4, the next equality follows from Lemma 8.3.2, and the next inequality follows from the result of [AKK+03].

8.4 Correlation of GF(2) polynomials with Mod_3^n

We use the results proved in the previous section to give an explicit example of a function that has exponentially decaying (in the number of variables) correlation with low degree $GF(2)$ polynomials.

Let $\omega \in \mathbb{C}$ be a primitive cube root of unity. Consider the function $Mod_3^n : \{0,1\}^n \to \mathbb{C}$ which is defined as $Mod_3^n(x_1, x_2, \ldots, x_n) = \omega^{\sum x_i} = \prod_{i=1}^n \omega^{x_i}$.

Bourgain [Bou05] showed that $\mathrm{Corr}(Mod_3^n, P_{d-1}) \leq \exp(\frac{-n}{8^d})$, where now the correlation between two functions f and g is defined as the norm of the complex number $E\left[fg\right]$.We show here how to derive a slightly better bound as a consequence of the results in [VW07].

By Lemma 8.3.4, $\mathrm{Corr}(Mod_3^n, P_{d-1}) \leq U_d(Mod_3^n)^{\frac{1}{2^d}}$. By Lemma 8.3.2, $U_d(Mod_3^n)^{\frac{1}{2^d}} = U_d(Mod_3^1)^{\frac{n}{2^d}}$. Consider the function $Mod_3^1 : \{0,1\} \to \mathbb{C}$, where $Mod_3^1(0) = 1$ and $Mod_3^1(1) = \omega$.

$$U_d(Mod_3^1) = E_{x,y_1,\dots y_d \in \{0,1\}} \left[\prod_{z \in C(\overline{Y})} (\omega^{(x \oplus z)})^{\underline{S_z}} \right].$$

Now, if any one of $y_1, \dots y_d$ is 0, then $\prod_{z \in C(\overline{Y})} (\omega^{(x \oplus z)})^{\underline{S_z}} = 1$, since for ever term in the product that is ω, there is one that is ω^2. Also, if $y_1 = y_2 = \dots = y_d = 1$, then $\prod_{z \in C(\overline{Y})} (\omega^{(x \oplus z)})^{\underline{S_z}} = \frac{\omega^{2^d} + \omega^{-2^d}}{2} = \frac{\omega + \omega^{-1}}{2} = \frac{-1}{2}$. Hence $U_d(Mod_3^1) = (1 - 2^{-d}) \cdot 1 + 2^{-d} \cdot (-1/2) \leq 1 - 2^{-d}$.

Thus we conclude that $\mathrm{Corr}(Mod_3^n, P_{d-1}) \leq (1 - 2^{-d})^{\frac{n}{2^d}} \leq \exp(\frac{-n}{4^d})$.

The above correlation bound for the complex-valued function Mod_3 also implies a bound for a Boolean analogue of the Mod_3 function for an appropriate definition of correlation [VW07].

9

Applications to PCPs
Luca Trevisan

Scribe(s): Rajsekar Manokaran

Summary: In this lecture, we will use Gowers uniformity to design a relaxed linearity test that optimizes on the query complexity. We will see how this construction can be extended to obtain a PCP verifier with similar query complexity, although assuming the Unique Games Conjecture.

9.1 PCPs and Query Complexity

In this section, we will see a construction of a PCP verifier for Unique-Games-Hard languages that makes q queries, has almost perfect completeness and soundness $\frac{q+1}{2^q} + \epsilon$ for arbitrarily small $\epsilon > 0$.

9.1.1 Linearity Testing

Linearity testing is central to constructing PCP verifiers and will hence serve as a good starting point to construct query efficient PCP verifiers.

Definition 9.1.1. A function $f : \{0,1\}^n \to \{-1,1\}$ is said to be linear if $\forall x, y \in \{0,1\}^n$:

$$f(x)f(y) = f(x+y)$$

equivalently, f is linear if $\exists a_1, a_2, \ldots a_n \in \{0,1\}$ such that:

$$f(x) = (-1)^{a_1x_1+a_2x_2+\ldots+a_nx_n}$$

In the linearity testing problem, we are given oracle access to a function f and we have to distinguish between the following two cases:

1. f is linear

2. f is "far" from being linear: For every linear function g, f agrees with g in at most $1/2+\epsilon$ fraction of the input.

Blum, Luby and Rubinfeld [BLR93] gave a very simple such linearity test:

Algorithm 9.1.2 (BLR Linearity Test)**.**
Input: Oracle access to a boolean function f

1. Choose $x, y \leftarrow_{\mathrm{R}} \{0,1\}^n$

2. *Accept* if and only if $f(x)f(y) = f(x+y)$

The above test can be analyzed quite elegantly by looking at the fourier spectrum of f. It is clear that the above test always accepts linear functions. The following theorem will show that the test rejects functions that agree in less than a $1/2+\epsilon$ fraction with probability at least $1/2+\epsilon$.

Theorem 9.1.3. *If Algorithm 9.1.2 accepts a (boolean) function f with probability greater than $1/2+\epsilon$, then there exists a set $S \subseteq [n]$ such that the (linear) function $\chi_S = (-1)^{+_{\{i\in S\}} x_i}$ agrees in at least $1/2+\epsilon$ fraction with f.*

Proof.

$$\begin{aligned}
\Pr_{x,y}[f(x)f(y) = f(x+y)] &= \tfrac{1}{2} + \tfrac{1}{2}\mathbb{E}_{x,y} f(x)f(y)f(x+y) \\
&= \tfrac{1}{2} + \tfrac{1}{2}\mathbb{E}_{x,y} \sum_{a,b,c} \hat{f}_a \hat{f}_b \hat{f}_c \chi_a(x)\chi_b(y)\chi_c(x+y) \\
&= \tfrac{1}{2} + \tfrac{1}{2} \sum_{a,b,c} \hat{f}_a \hat{f}_b \hat{f}_c \mathbb{E}_x \chi_a(x)\chi_c(x) \mathbb{E}_y \chi_b(y)\chi_c(y) && \begin{pmatrix}\text{by linearity of } \chi_S \text{ (and} \\ \text{of expectation)}\end{pmatrix} \\
&= \tfrac{1}{2} + \tfrac{1}{2} \sum_S \hat{f}_S^3 && (\chi_S \text{ form a orthonormal basis }) \\
&\le \tfrac{1}{2} + \tfrac{1}{2} \sum_S \hat{f}_S^2 \max_S \hat{f}_S \\
&\le \tfrac{1}{2} + \tfrac{1}{2} \max_S \hat{f}_S
\end{aligned}$$

Since $\hat{f}_S = \mathbb{E}_x[f(x)\chi_S(x)]$, the probability that f agrees with χ_S is:

$$\begin{aligned}
\Pr_x[f(x) = \chi_S(x)] &= \tfrac{1}{2} + \tfrac{1}{2}\mathbb{E}_x f(x)\chi_S(x) \\
&= \tfrac{1}{2} + \tfrac{1}{2}\hat{f}_S \\
&\ge \Pr_{x,y}[f(x)f(y) = f(x+y)] && (\text{choosing an } S \text{ that maximizes } \hat{f}_S)
\end{aligned}$$

■

Thus, this test rejects functions far from being linear with probability about $\frac{1}{2}$. The soundness can be improved simply by repeating this test independently. Repeating t times, we get a test that queries the function f at $3t$ points and with "soundness" $1/2^t$ (we are interested in cases where ϵ

is arbitrarily small). We are interested in improving the *query complexity* of the test. This test, queries f at q points and has soundness $1/2^{\frac{q}{3}}$.

Samorodnitsky and Trevisan [ST00] considered the following extension to test, called the complete graph test:

Algorithm 9.1.4 (Complete Graph Test)**.**
Input: Oracle access to a boolean function f

```
1. Choose $x_1, \ldots x_k \leftarrow_{\mathrm{R}} \{0,1\}^n$
2. For every pair $(i,j)$; $1 \le i < j \le k$, check if $f(x_i)f(x_j) = f(x_i + x_j)$
3. Accept only if every test accepts
```

We evaluate the function at $k + \binom{k}{2}$ points and perform $\binom{k}{2}$ tests. [ST00] show that in case the function is far from being linear, the tests behave almost independently thus accepting with probability around $(\frac{1}{2})^{q-\sqrt{2q}}$.

9.1.2 Hypergraph Tests and Gowers Uniformity

For further improving the query complexity, we would like to look at the so called complete hypergraph test [Sam05]:

Algorithm 9.1.5 (Complete Hypergraph Test)**.**
Input: Oracle access to a boolean function f

```
1. Choose $x_1, \ldots x_d \leftarrow_{\mathrm{R}} \{0,1\}^n$
2. For every pair $S \subseteq [d]$; $|S| \ge 2$, check if $\prod_{i \in S} f(x_i) = f(\sum_{i \in S} x_i)$
3. Accept only if every test accepts
```

Strictly speaking, the test needs to be parametrized by d. We will informally refer to the test as the hypergraph test and refer to it as the d-hypergraph test when the parameter d needs to be explicitly specified.

The test queries f at $q = 2^d - 1$ points and performs $2^d - d - 1$ tests. Again, if the tests behave independently when f is far from being linear, we will get a linearity test that attains a soundness value of $(q+1)/2^q$ This, however, turns out to be false: there exists functions which are far from being linear on which the tests are not all mutually independent. For example, the function $f(x) = (-1)^{x_1 x_2 + \ldots + x_{n-1} x_n}$ is far from begin linear. However, as is shown in [ST00], the test accepts f with probability at least $2^{-q+\Omega(\sqrt{q})}$. It is also known [ST06] that any linearity test that makes q queries and accepts linear functions with probability atleast c must accept f with probability at least $(1-c) + 2^{-q+\Omega(\sqrt{c})}$. Thus, the basic linearity testing problem does not have much room for improvement.

However, as we will see, if we are willing to relax the definition of the test to a "$(d-1)$ degree linearity test", we can use the d-hypergraph test (or in fact any hypergraph test where each edge in the hypergraph contains at most d vertices). We will begin by defining the Gowers inner product, Gowers uniformity [Gow98, Gow01] and the "generalized" linearity test.

Definition 9.1.6 (Gowers inner product)**.** The Gowers dimension-d inner product of a collection $\{f_S\}_{S\subseteq[d]}$ of functions $f_S : \{0,1\}^n \to \mathbb{R}$ is defined as:

$$\langle\{f_S\}\rangle_{U^d} \equiv \mathbb{E}_{x,x_1,\dots,x_d}\left[\prod_{S\subseteq[d]} f_S\left(x+\sum_{i\in S} x_i\right)\right]$$

Definition 9.1.7 (Gowers Uniformity)**.** The Gowers dimension-d uniformity of $f : \{0,1\}^n \to \mathbb{R}$ is defined as:

$$U^d(f) \equiv \mathbb{E}_{x,x_1,\dots,x_d}\left[\prod_{S\subseteq[d]} f\left(x+\sum_{i\in S} x_i\right)\right] = \langle\{f\}\rangle_{U^d}$$

Definition 9.1.8 ($(d-1)$-degree linearity test)**.** Given oracle access to a function f, a $(d-1)$-degree linearity test distinguishes between the following two cases:

1. f is linear
2. $U^d(f) \le \epsilon$

We are now in a position to prove that the d-hypergraph test tests $(d-1)$-degree linearity. Since the "completeness" case is trivial, we will be done if we prove the following theorem.

Theorem 9.1.9. *If f is a function such that $U^d(f) \le \epsilon^{2^{d+1}}$, then f passes the d-hypergraph test with probability at most $1/2^{2^d-d-1} + \epsilon = 1/2^k + \epsilon$, then*

We will need the following two simple claims about the gowers uniformity and inner product [Gow01, GT04].

Claim 9.1.9.1. $[\langle\{f_S\}\rangle_{U^d}]^{2^d} \le \prod_S U^d(f_S)$

Proof.

$$\begin{aligned}
|(\langle\{f_S\}\rangle_{U^d})| &= \left|\mathbb{E}_{x_1,\dots,x_{d-1}}\mathbb{E}_{x,x_d}\prod_{S\subseteq[d]} f_S\left(x+\sum_{i\in S}x_i\right)\right| \\
&= \left|\mathbb{E}_{x_1,\dots,x_{d-1}}\left[\mathbb{E}_x\prod_{S\subseteq[d-1]} f_S\left(x+\sum_{i\in S}x_i\right)\right]\left[\mathbb{E}_x\prod_{S\subseteq[d-1]} f_{S\cup\{d\}}\left(x+\sum_{i\in S}x_i\right)\right]\right| \\
&\le \left[\mathbb{E}_{x_1,\dots,x_{d-1}}\left|\mathbb{E}_x\prod_{S\subseteq[d-1]} f_S\left(x+\sum_{i\in S}x_i\right)\right|^2\right]^{\frac{1}{2}}\left[\dots\right]^{\frac{1}{2}} \\
&= \left[\mathbb{E}_{x_1,\dots,x_d,x}\prod_{S\subseteq[d-1]} f_S\left(x+\sum_{i\in S}x_i\right)\right]^{\frac{1}{2}}\left[\dots\right]^{\frac{1}{2}}
\end{aligned}$$

Doing this recursively gives the claim. ■

Claim 9.1.9.2. $U^{d-1}(f) \leq \sqrt{U^d(f)}$

Proof.

$$[U^{d-1}(f)]^2 \leq \mathbb{E}_{x_1,\ldots,x_{d-1}} \left[\mathbb{E}_x \prod_{S\subseteq[d-1]} f\left(x + \sum_{i\in S} x_i\right)\right]^2$$

$$= \mathbb{E}_{x,x_1,\ldots,x_d} \left[\prod_{S\subseteq[d]} f\left(x + \sum_{i\in S} x_i\right)\right] = \langle\{f\}\rangle_{U^d}$$

■

Proof of Theorem 9.1.9. If a function f passes the d-hypergraph test with probability at least $1/2^k + \epsilon$, then:

$$\tfrac{1}{2^k} + \epsilon \leq \mathbb{E}_{x_1,\ldots x_d} \left[\prod_{S\subseteq[d],|S|\geq 2} \left(\frac{1 + \prod_{i\in S} f(x_i) f(\sum_{i\in S} x_i)}{2}\right)\right]$$

$$\leq 1/2^k \left[\sum_{X\subseteq 2^{[d]}} \mathbb{E}_{x_1,\ldots x_d} \prod_{S\in X} \prod_{i\in S} f(x_i) f(\sum_{i\in S} x_i)\right]$$

Hence, there exists a non-empty set of "edges" $X \subseteq 2^{[d]}$, giving:

$$\mathbb{E}_{x_1,\ldots,x_t} \prod_{S\in X} \prod_{i\in S} f(x_i) f(\sum_{i\in S} x_i) \geq \epsilon$$

If t is the size of the largest edge in X, then assuming without loss of generality that $[t] \in X$, we can fix the variables $x_{t+1}, \ldots x_d$ such a way that the above expectation (over $x_1 \ldots x_t$) is at least ϵ. Further, we can group the fixed terms thus forming a collection of functions $\{f_S\}_{S\subseteq[t]}$ such that

$$\mathbb{E}_{x_1,\ldots,x_t} \prod_{S\subseteq[t]} \left[f_S(\sum_{i\in S} x_i)\right] \geq \epsilon$$

Note that since $[t]$ is the largest edge, $f_{[t]} = f$. Using this, we can lower bound the gowers inner product of a related set of functions as follows:

$$\left[\mathbb{E}_{x_1,\ldots,x_t} \prod_{S\subseteq[t]} f_S(\sum_{i\in S} x_i)\right]^2 \leq \mathbb{E}_{x_1,\ldots,x_{t-1}} \left[E_{x_t} \prod_{S\subseteq[t]} f_S(\sum_{i\in S} x_i)\right]^2$$

$$= \mathbb{E}_{x_1,\ldots,x_{t-1}} \prod_{S\subseteq[t],t\notin S} \left(f_S(\sum_{i\in S} x_i)\right)^2 \left[E_{x_t} \prod_{S\subseteq[t],t\in S} f_S(\sum_{i\in S} x_i)\right]^2$$

$$= \mathbb{E}_{x_1,\ldots,x_{t-1}} \left[E_{x_t} \prod_{S\subseteq[t-1]} f_{S\cup\{t\}}(x_t + \sum_{i\in S} x_i)\right]^2$$

$$= \langle\{f_{S\cup\{t\}}\}\rangle_{U^t}$$

Using Claim 9.1.9.1 and Claim 9.1.9.2, we get

$$\epsilon^2 \leq \left|\langle\{f_{S\cup\{t\}}\rangle_{U^t}\right| \leq min_S \left|U^t(g_S)\right|^{1/2^t} \leq U^d(f)^{1/2^d}$$

■

This bound on the soundness of the hypergraph test is also known to be tight [ST06].

9.1.3 Gowers Uniformity, Long Codes and Influence

In this section, we sketch the construction of the PCP verifier. For constructing the verifier, we will want to test long codes. This is done by extending the hypergraph test to the setting of several functions and by introducing noise in the test.

Algorithm 9.1.10 (γ-noisy hypergraph test)**.**

Fix a hypergraph $([t], E)$ and a $\gamma > 0$.

Let μ_γ denote the distribution over $\{0,1\}^n$ where each bit is independently chosen to be 1 with probability γ.

Input: Oracle access to a boolean function $\{g_a\}_{a\in[t]\cup E}$

1. Choose $x_1, \ldots x_t \leftarrow_{\mathrm{R}} \{0,1\}^n$

2. Choose $n_1, \ldots n_t \leftarrow_{\mathrm{R}} \mu_\gamma$

3. For every $e \in E$, choose $n_e \leftarrow_{\mathrm{R}} \mu_\gamma$

4. Accept if and only if
$$\forall e \in E, \prod_{i\in e} g_i(n_i + x_i) = g_e(n_e + \sum_{i\in e} x_i)$$

Given a collection of m balanced boolean functions $g_1, \ldots, g_m$, the long code test asks to distinguish between the following two cases:

1. The functions are all equal to the same long code

2. No variable is "influential" in more than a single function in the collection

This test is analysed by first proving that for a collecton of functions, if the t-dimensional gowers uniformity is large, then there exists a variable i which is influential in at least two of the functions (provided the functions are balanced). We can ensure that the functions (or long codes) the verifier has to deal with are always balanced by an elegant (and quite standard) technique called *folding*. Now, when the t-dimension gowers uniformity is small, by Theorem 9.1.9, the tests are almost independent and hence in the soundness case, the acceptance probability is very low. Hence, Algorithm 9.1.10 is a long code tester.

9.1.4 The PCP Verifier

In this section, we will see how we can compose an outer verifier for NP obtained assuming the unique games conjecture [Kho02] with an inner verifier to obtain a PCP verifier with the promised query complexity. The following claim is due to a result of Khot and Regev [KR03]. Assuming the UGC, for every integer q, and every $\gamma > 0$, there is a $\sigma(q, \gamma)$ such that there is a reduction from SAT to unique games on alphabets of size σ such that:

1. Every constraint involve q variables. A constraint is specified by the q variables, $v_1, \ldots, v_q$ along with q permutations $f_1, \ldots, f_q$ on the alphabet. A constraint is said to be satisfied by an assignment A to the variables if $f_1(A(v_1)) = f_2(A(v_2)) = \ldots = f_q(A(v_q))$.

2. If the formula is satisfiable, there is an assignment that satisfies at least a $1 - \gamma$ fraction of the constraints

3. A constraint is said to be satisfied weakly if $f_i(A(v_i))$ are not all different

4. If the formula is not satisfiable, then, no assignment satisfies more than a γ fraction even in the weak sense defined above

The inner verifier is as follows. Given a q-ary unique game, the verifier expects the long code of each constraint as the proof. Given a proof, the verifier picks a random constraint, say involving $v_1, \ldots, v_q$ and permutations $f_1, \ldots f_q$. Let $h_1, \ldots, h_q$ be the functions representing the supposed long codes (after folding) of the assignment to the corresponding variables. The verifier runs a γ-noisy hypergraph test on the functions $h_1 \circ f_1, \ldots, h_q \circ f_q$ (with the complete hypergraph on q vertices).

The completeness of this verifier is $1 - q\gamma$. We look at the soundness of the verifier. Suppose the verifier accepts with probability $1/2^{2^q - q - 1} + \epsilon = 1/2^k + \epsilon$, then a $\epsilon/2$ fraction of the edges accept with probability $1/2^k + \epsilon/2$. For every variable, we assign it randomly, one of the influential coordinates (or assign a random coordinate if there is no influential coordinate). Since a boolean function can not have too many influential variables, each edge is weakly satisfied with a good probability. Thus, taking expectation, we would expect a good fraction of the edges to be weakly satisfied. We fix γ and ϵ depending on the guarantees we obtain on the influence of variables in the "good" constraints. Wrapping up, we get a PCP verifier that has completeness arbitrarily close to 1 while having soundness arbitrarily close to $1/2^k$. Thus, assuming UGC, $NP = PCP_{1-\delta, \frac{w+1}{2^w}+\delta}(O(\log n), w)$ for every w of the form $2^t - 1$.

Bibliography

[AKK$^+$03] Alon, Kaufman, Krivelevich, Litsyn, and Ron. Testing Low-Degree Polynomials over GF(2). In *RANDOM: International Workshop on Randomization and Approximation Techniques in Computer Science*. LNCS, 2003.

[ADL$^+$94] N. Alon, R. A. Duke, H. Lefmann, V. Rödl, and R. Yuster. The algorithmic aspects of the regularity lemma. *J. Algorithms*, 16(1):80–109, 1994.

[BS94] A. Balog and E. Szemerédi. A statistical theorem of set addition. *Combinatorica*, 14(3):263, 1994.

[BS96] A. Balog and E. Szemerédi. A Statistical Theorem of set addition. *Combinatorica*, 14:263–268, 1996.

[BW04] I. R. Barak, B. and A. Wigderson. Extracting randomness using few independent sources. *FOCS*, pages 384–393, 2004.

[BW05] K. G. S. R. S. B. Barak, B. and A. Wigderson. Simulating independence: new constructions of condensers, ramsey grahs, dispersers, and extractors. *Proceedings of the thirty-seventh annual ACM symposium on theory of computing*, pages 1–10, 2005.

[BW06] R. A. S. R. Barak, B. and A. Wigderson. 2-source dispersers for sub-polynomial entropy and Ramsey graphs beating the Frankl-Wilson construction. *Proceedings from the thirty-eighth annual ACM symposium on theory of computing*, pages 671–680, 2006.

[Bes63] A. S. Besicovitch. The Kakeya Problem. *The American Mathematical Monthly*, 70(7):697–706, 1963.

[BLR93] M. Blum, M. Luby, and R. Rubinfeld. Self-testing/correcting with applications to numerical problems. *J. Comput. Syst. Sci.*, 47(3):549–595, 1993.

[BK06] G. A. Bourgain, J. and S. Konyagin. Estimates for the number of sums and products and for exponential sums in fields of prime order. *Journal of the London Mathematical Society*, 73(2):380–398, 2006.

[Bou05] J. Bourgain. Estimation of certain exponential sums arising in complexity theory. In *C. R. Math. Acad. Sci. Paris, 340(9)*, pages 627–631, 2005.

[Bou07] J. Bourgain. On the construction of affine extractors. *GAFA*, 17:22–57, 2007.

[BG06] J. Bourgain and A. Gamburd. Growth and generation in $SL_2(\mathbb{Z}/p\mathbb{Z})$. *C. R. Acad. Sci. Paris*, 342(10):717–721, 2006.

[BT04] K. N. Bourgain, J. and T. Tao. A Sum-Product Estimate in Finite Fields, and Applications. *GAFA*, 14(1):27–57, 2004.

[CG88] B. Chor and O. Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM Journal of Computing*, 17(2):230–261, 1988.

[Ele97] G. Elekes. On the Number of Sums and Products. *Acta Arithmetica*, 81:365–367, 1997.

[ES83a] P. Erdös and E. Szemerédi. On Sums and Products of Integers. *Studies in Pure Mathematics*, pages 213–218, 1983.

[ES83b] P. Erdös and E. Szemerédi. On Sums and Products of Integers. In L. Alpar, G. Halasz, and A. Sarközy, editors, *Studies in Pure Mathematics: To the memory of Paul Turán*, pages 213–218. Birkhaüser, Basel, 1983.

[GR05] A. Gabizon and R. Raz. Deterministic extractors for affine sources over large fields. *FOCS*, pages 407–416, 2005.

[Gow98] T. Gowers. A new proof of Szemrédi's theorem for arithmetic progressions of length four. *GAFA*, 8:529–551, 1998.

[Gow01] T. Gowers. A new proof of Szemerédi's theorem. *GAFA*, 11:465–588, 2001.

[Gow98] W. Gowers. Fourier Analysis and Szemerédi's Theorem. *Proceedings of the International Congress of Mathematicians*, 1:617–629, 1998.

[Gow97] W. T. Gowers. Lower bounds of tower type for Szemerédi's uniformity lemma. *Geom. Funct. Anal.*, 7(2):322–337, 1997.

[Gre05] B. Green. Some Notes On The Bourgain-Katz-Tao Sum-Product Theorem, 2005.

[GT04] B. Green and T. Tao. The primes contain arbitrarily long arithmetic progressions, 2004.

[GT05] B. Green and T. Tao. An inverse theorem for the Gowers U^3 norm, 2005.

[HB96] D. Heath-Brown. An Estimate for Heilbronn's Exponential Sum. *Analytic Number Theory, Proceedings of a Conference to Honor Heini Halberstam*, pages 451–463, 1996.

[Hel] H. Helfgott. Growth and generation in $SL_2(\mathbb{Z}/p\mathbb{Z})$. *ARXIV*.

[Kho02] S. Khot. On the power of unique 2-prover 1-round games. In *STOC '02: Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pages 767–775, New York, NY, USA, 2002. ACM Press.

[KR03] S. Khot and O. Regev. Vertex cover might be hard to approximate to within $2-\epsilon$, 2003.

[Kon03] S. Konyagin. A Sum-Product Estimate in Fields of Prime Order. *Arxiv*, 2003.

[KS99] S. Konyagin and I. Shparlinski. Character sums with exponential functions and their applications. *Cambridge Uniersity Press*, 1999.

[Lev85] L. A. Levin. One-way functions and pseudorandom generators. In *STOC '85: Proceedings of the seventeenth annual ACM symposium on Theory of computing*, pages 363–365, New York, NY, USA, 1985. ACM Press.

[LS86] P. R. Lubotzky, A. and P. Sarnak. Explicit expanders and the Ramanujan conjecture. *STOC*, page 240, 1986.

[Mar73] G. A. Margulis. Explicit constructions of expanders. *Problemy Peredači Informacii*, 9(4):71–80, 1973.

[Plü69] H. Plünnecke. Eigenschaften und Abschätzungen von Wirkingsfunktionen. *Berichte der Gesellschaft für Mathematik und Datenverarbeitung*, 22, 1969.

[Plu69] H. Plunneke. Eigenschaften und Abschätzungen von Wirkingsfunktionen. *Gesellschaft für Mathematik und Datenverarbeitung*, 1969.

[Rao06] A. Rao. Extractors for a constant number of polynomially small min-entropy independent sources. *Proceedings of the thirty-eighth annual ACM symposium on theory of computing*, pages 497–506, 2006.

[Raz87] A. Razborov. Lower bounds on the size of bounded-depth networks over a complete basis with logical addition (Russian). In *Mat. Zametki, 41(4)*, pages 598–607, 1987.

[Ruz96] I. Ruzsa. Sums of Finite Sets. In D. Chudnovsky, G. Chudnovsky, and M. Nathanson, editors, *Number Theory: New York Seminar*. Springer-Verlag, 1996.

[Sam05] A. Samorodnitsky. Hypergraph linearity and quadraticity tests for boolean functions. Manuscript, 2005.

[ST00] A. Samorodnitsky and L. Trevisan. A PCP characterization of NP with optimal amortized query complexity. In *STOC '00: Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 191–199, New York, NY, USA, 2000. ACM Press.

[ST06] A. Samorodnitsky and L. Trevisan. Gowers uniformity, influence of variables, and PCPs. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 11–20, New York, NY, USA, 2006. ACM Press.

[Sel65] A. Selberg. On the estimation of Fourier coefficients of modular forms. *Proceedings of the Symposium of Pure Math*, 8:1–15, 1965.

[Sze78] E. Szemerédi. Regular partitions of graphs. In *Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976)*, volume 260 of *Colloq. Internat. CNRS*, pages 399–401. CNRS, Paris, 1978.

[TV06] T. Tao and V. Vu. *Additive Combinatorics.* Cambridge University Press, 2006.

[Vaz87] U. Vazirani. Efficiency considerations in using semi-random sources. *Proceedings from the ninteenth annual ACM conference on theory of computing*, pages 160–168, 1987.

[Vio06] E. Viola. New correlation bounds for GF(2) polynomials using Gowers uniformity. *Electronic Colloquium on Computational Complexity (ECCC)*, (097), 2006.

[VW07] E. Viola and A. Wigderson. Norms, XOR Lemmas, and Lower Bounds for GF(2) Polynomials and Multiparty Protocols. In *IEEE Conference on Computational Complexity*, pages 141–154. IEEE Computer Society, 2007.

[Wol99] T. Wolff. Recent Work Connected with the Kakeya Problem. *Prospects in Mathematics, American Mathematical Society*, pages 129–162, 1999.

[Yao82] A. C. Yao. Theory and Applications of Trapdoor Functions (Extended Abstract). In *23rd Annual Symposium on Foundations of Computer Science*, pages 80–91, Chicago, Illinois, 3–5 Nov. 1982. IEEE.