Polynomial Method in Combinatorics

Youri Tamitegama [1]

Supervisor: Bogdan Nica
Mcgill University

Winter 2018

[1] Email: youri.tamitegama@mail.mcgill.ca

Abstract

These notes are an exploration of a surprisingly powerful perspective that can be used to solve combinatorial problems. This technique boils down to, given a question that is combinatorial in nature, reducing it to a question about the zero set of one or several polynomials.

This approach may seem a bit strange at first, especially since it is often hard to see the link between the statements of some of the problems that can be solved using this technique and polynomials. But in recent years, many important problems in combinatorics that were thought to be hard were solved quite simply using this method. Among these problems are the Kakeya and Nikodym conjectures in finite fields, the Dyson conjecture and the cap set problem.

Contents

1 Alon's Combinatorial Nullstellensatz **2**

1.1 Main theorems 2

1.2 Two classical applications 5

1.3 Graphs 6

1.4 The permanent lemma 7

1.5 Hypercube Coverings 9

1.6 Exercises 10

2 The Kakeya Set Problem **11**

2.1 Introduction 11

2.2 Dvir's proof for quantitative Kakeya sets 12

2.3 The Kakeya conjecture 16

2.4 Sharper bounds 17

2.5 An Explicit Construction 23

3 A reformulation using Lagrange Interpolation **25**

3.1 The Dyson Conjecture 26

4 The Cap Set Problem **30**

4.1 Introduction 30

4.2 The game of SET 30

4.3 Notation 31

4.4 Main Result 32

Chapter 1

Alon's Combinatorial Nullstellensatz

In this chapter we discuss Alon's paper on his combinatorial Nullstellensatz, which can be seen as the paper that roots the polynomial method. The main reference will be the paper that he published in 1999, [Alo99].

The term *Nullstellensatz* means "theorem about zeros". This is no fluke, the theorem is precisely a statement about the general shape and size of the sets of zeros of a polynomial depending on its highest degree terms.

The main idea comes from the basic fact that a polynomial in one variable of degree d cannot have more than d roots. We can also turn this around to see that given any set of $d+1$ elements, one of them will necessarily be a non-root of the polynomial. Of course, in one dimension, this observation is next to useless. But when generalized to multiple dimensions, this very simple fact can become quite powerful and as we will see, has many interesting applications.

1.1 Main theorems

Hilbert's Nullstellensatz states that given an arbitrary set of n-variate polynomials g_i over an algebraically closed field F, if some other n-variate polynomial f vanishes over the common zeros of the g_i's, then f raised to some power is contained in the ideal generated by the g_i's.

In other words, for such an f, $\exists k \in \mathbb{N}$ such that

$$f^k = \sum_{i=1}^{m} h_i g_i$$

where m is the number of g_i.

In Alon's Nullstellensatsz, not only do we restrict ourselves to the case when there are as many g_i's as there are dimensions, but we also make a very specific choice of g_i's that allows us to conclude that they form a basis for all polynomials vanishing on their common zeros. This property leads in particular to a corollary about existence of non vanishing elements in sets that are much larger than the degree of a polynomial.

We begin by presenting and proving the two main theorems of the Combinatorial Nullstellensatz. We will then see several interesting applications of these theorems.

Theorem 1.1.1. *Let F be an arbitrary field, $f \in F[x_1, \ldots, x_n]$. Let $S_1, \ldots, S_n$ be nonempty subsets of F and define $g_i(x_i) = \prod_{s \in S_i}(x_i - s)$. If f vanishes over all common zeros of $g_1, \ldots, g_n$, then there are polynomials $h_1, \ldots, h_n \in F[x_1, \ldots, x_n]$ satisfying $deg(h_i) \leq deg(f) - deg(g_i)$ such that*

$$f = \sum_{i=1}^{n} h_i g_i$$

Moreover, if $f, g_1, \ldots, g_n$ lie in $R[x_1, \ldots, x_n]$ for some subring R of F then there are polynomials $h_i \in R[x_1, \ldots, x_n]$

In a sense, it states that the g_i's form a basis for any polynomial vanishing on the entirety of $S_1 \times \ldots \times S_n$.

Theorem 1.1.2. *Let F be an arbitrary field, and let $f \in F[x_1, \ldots, x_n]$. Suppose the degree $deg(f)$ of f is $\sum_{i=1}^{n} t_i$, where each t_i is a nonnegative integer, and suppose the coefficient of the term $\prod_{i=1}^{n} x_i^{t_i}$ in f is nonzero. Then if $S_1, \ldots, S_n$ are subsets of F with $|S_i| > t_i$, there is a point $(s_1, \ldots, s_n) \in S_1 \times \ldots \times S_n$ so that*

$$f(s_1, \ldots, s_n) \neq 0$$

This second formulation of the theorem is the one that is most commonly used in applications. Its classical use is to find some polynomial that admits a non root in some product of subsets of $\mathbb{F}$ if and only if some property that we want holds. Then by the conditions of the problem, one shows that the desired coefficient is nonzero and that the polynomial has small degree, allowing the application of the Nullstellensatz.

Lemma 1.1.3. *Let $f = f(x_1, \ldots, x_n)$ be a polynomial in n variables over some field $\mathbb{F}$. Suppose that the degree of f in the ith variable is at most t_i, let $S_i \subset \mathbb{F}$ be a set of at least $t_i + 1$ distict elements of $\mathbb{F}$. If $f(x_1, \ldots, x_n) = 0$ for all n-tuples $(x_1, \ldots, x_n) \in S_1 \times \ldots \times s_n$, then f is the zero polynomial.*

Proof (Lemma). By induction on n.
$n = 1$. This is the well-known statement that a polynomial of degree at most d cannot have more than d roots.
$n > 1$. Rewrite f as a polynomial in x_n:

$$f(x_1, \ldots, x_{n-1}, x_n) = \sum_{i=1}^{t_n} f_i(x_1, \ldots, x_{n-1}) x_n^i$$

Fixing some tuple $(s_1, \ldots, s_{n-1}) \in S_1 \times \ldots \times S_{n-1}$, we see that as a single variable polynomial (in the nth coordinate), by hypothesis, f is identically zero on S_n. Thus $\forall\ 1 \leq i \leq n-1$, $f_i(s_1, \ldots, s_n) = 0$. Hence each f_i is a polynomial in $n-1$ variables such that its degree in the jth variable is at most t_j that vanishes on all points $(s_1, \ldots, s_{n-1}) \in S_1 \times \ldots \times S_{n-1}$, with $|S_j| \geq t_j + 1$.

We apply the induction hypothesis to get that f_i are zero polynomials, and conclude that f must also be the zero polynomial. □

As we will see in the next chapter, the idea of this proof has been recycled by Tao to give a bound on the size of Kakeya sets.

It seems tempting to try and derive the second theorem directly from this lemma, but the condition that the degree of f in the ith variable is at most t_i is not necessarily satisfied. Indeed, we could have a term that has degree in x_i which is $t_i + 1$ while the total degree of f is still exactly $\sum_{i=1}^{n} t_i$.

Proof (1st Nullstellensatz). Define $t_i = |S_i| - 1$ and $g_i(x_i) = \prod_{s \in S_i}(x_i - s)$. Isolating the highest degree term of g_i we write

$$g_i(x_i) = x_i^{t_i+1} - \sum_{j=0}^{t_i} g_{ij} x_i^j$$

When $x_i \in S_i$ we have $g_i(x_i) = 0$ and thus the equality $x_i^{t_i+1} = \sum_{j=0}^{t_i} g_{ij} x_i^j$.

Using this equality, we can modify f by repeatedly replacing every instance of x_i^k where $k > t_i$ by $x^{k-(t_i+1)} \sum_{j=0}^{t_i} g_{ij} x_i^j$.

Calling the newly obtained polynomial $\widetilde{f}$, we can check that we can rearrange $f - \widetilde{f}$ so that it is of the form $\sum_{i=1}^{n} h_i g_i$, where each polynomial h_i has degree at most $deg(f) - deg(g_i)$. Looking at a single substitution, we see that

$$\begin{aligned} f - \widetilde{f} &= h_i(x) x_i^{t_i+1} - h_i(x) \sum_{j=0}^{t_i} g_{ij} x_i^j \\ &= h_i g_i \end{aligned}$$

One can then generalize this after a notation heavy calculation. Note further that $\widetilde{f}$ now has degree at most t_i in the ith variable.

Moreover, as we replaced terms of f by term that evaluate to the same values on the cross product of the S_i's, we have equality between f and $\widetilde{f}$ for all $x \in S_1 \times \ldots \times S_n$. Thus $\widetilde{f}$ is zero on $S_1 \times \ldots \times S_n$. Applying the lemma, we conclude that f_i is the zero polynomial. Thus

$$f = \sum_{i=1}^{n} h_i g_i$$

□

Proof (2nd Nullstellensatz). We may assume that $|S_i| = t_i + 1$. Suppose by contradiction that f is zero on the whole $S_1 \times \ldots \times S_n$. Then we may write it as $\sum_{i=1}^{n} h_i g_i$ with g_i defined as previously. But then we clearly have a contradiction as if we look at the coefficient of the term $\prod_{j=1}^{n} x^{t_j}$, it must come from at least one of the summands $g_i h_i$. This means that h_i must have degree at least $\sum_{j \neq i} t_j$. But then the term $g_i h_i$ has degree $1 + \sum_{j=1}^{n} t_j$, which is larger than the degree of f. Contradiction. □

1.2 Two classical applications

From this theorem we can quickly derive two well-known results. First is the weaker version of an old theorem in finite fields:

Theorem 1.2.1 (Chevalley-Warning). *Let p be a prime, let $P_1, \ldots, P_m \in \mathbb{F}_p[x_1, \ldots, x_n]$. If $n > \sum_{i=1}^{m} deg(P_i)$ and the polynomials P_i have a common zero $(c_1, \ldots, c_n)$, then they have another common zero.*

Proof. Suppose by contradiction that the point $(c_1, \ldots, c_n)$ is the only common zero.

First, we note that for a fixed point $(s_1, \ldots, s_n) \in \mathbb{F}_p^n$, the following holds:

$$\prod_{i=1}^{m} (1 - P_i(s_1, \ldots, s_n)^{p-1}) = \begin{cases} 1, & \text{if } \forall i, P_i(s_1, \ldots, s_n) = 0 \\ 0, & \text{if } \exists i \ s.t. \ P_i(s_1, \ldots, s_n) \neq 0 \end{cases}$$

In other words, this product is an indicator of the common zeros of the P_i's. We can further note that its degree is at most $(p-1) \sum_{i=1}^{m} deg(P_i)$.

Next, we can define an indicator of the point $(c_1, \ldots, c_n)$:

$$\prod_{j=1}^{n} \prod_{c \in \mathbb{Z}_p, c \neq c_j} (s_j - c) = \begin{cases} 1/\delta, & \text{if}(s_1, \ldots, s_n) = (c_1, \ldots, c_n) \\ 0, & \text{otherwise} \end{cases}$$

Where δ is some non zero constant.
Now if we define the following polynomial:

$$f(x_1, \ldots, x_n) = \prod_{i=1}^{m} (1 - P_i(x_1, \ldots, x_n)^{p-1}) - \delta \prod_{j=1}^{n} \prod_{c \in \mathbb{Z}_p, c \neq c_j} (x_j - c)$$

We have by the above observations that f vanishes both at $(c_1, \ldots, c_n)$ and at every other point of $\mathbb{F}_p^n$ (since by assumption $(c_1, \ldots, c_n)$ is the only common zero). We can also see the coefficient of the term $\prod_{i=1}^{n} x_i^{p-1}$ comes only from the second sum since by an above remark, the degree of the first term is $(p-1) \sum_{i=1}^{m} deg(P_i) < n(p-1)$. This means that the coefficient must be δ, thus nonzero.

Applying the second Nullstellensatsz to the set $\mathbb{F}_p \times \ldots \times \mathbb{F}_p$, we conclude that f must have a non root in $\mathbb{F}_p^n$, which is a contradiction. □

There exists a stronger version of this theorem, which states that the number of common zeros needs in fact to be a multiple of the characteristic of the field. However, it seems hard to derive this result from the Nullstellensatsz.

One is tempted approach this stronger version using a modification of the above polynomial:

$$f(x_1, \dots, x_n) = \prod_{i=1}^{m}(1 - P_i(x_1, \dots, x_n)^{p-1}) - \sum_{i=1}^{k} \delta_k \prod_{j=1}^{n} \prod_{c \in \mathbb{Z}_p, c \neq c_{ij}} (x_j - c)$$

where k is the number of common roots and $c_i = (c_{i1}, \dots, c_{in})$ the common roots. But this approach fails as $\sum_{i=1}^{k} \delta_k$ might be zero.

Next is a classical result in additive combinatorics.

Theorem 1.2.2 (Cauchy-Davenport). *Given A, B non-empty subsets of $\mathbb{F}_p$, for some prime p, the following holds:*

$$|A + B| \geq min\{p, |A| + |B| - 1\}$$

Proof. When $|A| + |B| > p$ then taking any $g \in \mathbb{F}_p$, the sets A and $g - B$ must intersect, thus we can write $g = a + b$ for some $a \in A$, $b \in B$.

Otherwise, assume by contradiction that $|A + B| < |A| + |B| - 1$. Take a subset C of $\mathbb{F}_p$ such that $A + B \subset C$ and $|C| = |A| + |B| - 2$. Then if we define

$$f(x, y) = \prod_{c \in C}(x + y - c)$$

we have that $f(a, b) = 0$ for $a \in A$, $b \in B$, since $A + B \subset C$. Furthermore, the coefficient of the monomial $x^{|A|-1}y^{|B|-1}$ is $\binom{|C|}{|A|-1}$. Since $|C| = |A| + |B| - 2$ and $|A| + |B| \leq p$, this coefficient is nonzero in $\mathbb{F}_p$. Finally, this monomial is a maximum degree term in $\mathbb{F}$. We can thus apply the Nullstellensatz to $A \times B$ and derive a contradiction. □

This bound is tight. Indeed, if we pick A and B to be two singletons, $|A+B| = 1$, achieving the bound.

The statement does not hold for general finite fields, as the characteristic of the field can be a divisor of the coefficient of the $x^{|A|-1}y^{|B|-1}$.

1.3 Graphs

Proposition 1.3.1. *In any 4-regular simple graph, there exists a 3-regular subgraph.*

This proposition is almost a special case (when $p = 3$) of the following theorem. But it does not quite work with a perfectly 4-regular graph. That being said, if we allow ourselves to slightly raise the average degree (eg. by adding an extra edge), we get that the newly obtained graph has a 3-regular subgraph.

Theorem 1.3.2. *For any prime p, any loopless graph G=(V,E) with average degree bigger than 2p-2 and maximum degree at most 2p-1 contains a p-regular subgraph.*

Proof. For each edge $e \in E$ we define a variable x_e. The idea is to use these variables as selectors for the edges: $x_e = 1$ means e is in the subgraph, otherwise it is not. Let $a_{v,e} = 1$ if v is incident to e, 0 otherwise.

Since the maximum degree is $2p-1$, for a fixed $v \in V$, $\sum_{e \in E} a_{v,e} x_e = 0 \pmod p$ is equivalent to saying that either exactly p of the x_e are 1 or they are all 0 (i.e. p or no edges have been chosen, respectively).

Now define the following polynomial:

$$f = \prod_{v \in V} (1 - (\sum_{e \in E} a_{v,e} x_e)^{p-1}) - \prod_{e \in E} (1 - x_e)$$

By the above observation, the first product is nonzero if and only if each individual vertex has either degree 0 or p. But the second product is 1 if and only if all x_e are zero (i.e. no edges have been picked), and otherwise zero. Thus f is zero for any assignment of x_e unless at least one x_e is nonzero and all vertices have degree p or zero, i.e. the subgraph induced by the selected edges is p-regular.

The degree of the first product is $(p-1)|V|$, and by our initial hypothesis on the average degree we $2|E| > |V|(2p-2)$. Hence the coefficient of $\prod_{e \in E} x_e$ comes only from the second product. We can see that it is nonzero.

Applying the Nullstellensatsz to the set $\{0,1\}^{|E|}$, we obtain that f has a non root in this set, i.e. G has a p-regular subgraph.

□

1.4 The permanent lemma

Recall that the definition of the permanent of a matrix is the same as for a determinant, but without taking into account the sign of the permutations:

$$perm(A) = \sum_{\sigma \in S_n} \prod_{i=1}^{n} a_{i\sigma(i)}$$

The idea behind the permanent lemma is the following: suppose you're given an $n \times n$ matrix A, a vector $\mathbf{b} = (b_1, \ldots, b_n) \in \mathbb{F}^n$ and subsets $S_1, \ldots, S_n$ of $\mathbb{F}$ of size at least 2. Further suppose that we want to find some vector $\mathbf{x} \in S_1 \times \ldots \times S_n$ such that the vector $A\mathbf{x}$ differs in every coordinate from $\mathbf{b}$. Then the existence of such a vector is equivalent to saying that the polynomial

$$\prod_{i=1}^{n} \Big((\sum_{j=1}^{n} x_j a_{ij}) - b_i \Big)$$

admits a non root in the set $S_1 \times \ldots \times S_n$. The link with the Nullstellensatsz is now clear, and the permanent takes its place in the story as the coefficient of the term $\prod_{i=1}^{n} x_i$.

Lemma 1.4.1 (Permanent lemma)**.** *Let $A = (a_{ij})$ be an $n \times n$ matrix over a field $\mathbb{F}$ such that its permanent is nonzero over $\mathbb{F}$. Then for any vector $\mathbf{b} = (b_1, \dots, b_n) \in \mathbb{F}^n$ and for any family of subsets $S_1, \dots, S_n$ of $\mathbb{F}$, each of size at least 2, there is a vector $\mathbf{x} \in S_1 \times \dots \times S_n$ such that for every i the ith coordinate of $A\mathbf{x}$ differs from b_i.*

Proof. It follows from the above discussion. □

From this lemma we can derive another classical result in additive combinatorics.

Theorem 1.4.2 (Erdos-Ginzburg-Ziv)**.** *For any prime p, any sequence of 2p-1 members of Z_p contains a subsequence of cardinality p such that the sum of its members is 0 in Z_p*

This theorem can be proven in several ways. First we will see that it follows as an immediate corollary of the permanent lemma. It is also possible to derive it directly from the combinatorial Nullstellensatz as we will see later. The proof using the permanent is not very natural, but illustrates well the use of the permanent lemma.

Permanent proof. Given $a_1, \dots, a_{2p-1} \in Z_p$, WLOG $a_1 \leq a_2 \leq \dots \leq a_{2p-1}$ consider the $p-1 \times p-1$ all 1 matrix A.

If we define $p-1$ sets $S_i = \{a_i, a_{i+p-1}\}$, then given any vector $\mathbf{b} \in Z_p^n$, if we can apply the permanent lemma we get that there exists a selection of $p-1$ elements of the subsequence such that they do not sum to any of $\mathbf{b}$'s coordinates. Noticing that for any such selection, the element a_{2p-1} will never be picked, we let the coordinates of $\mathbf{b}$ be all elements of Z_p but $-a_{2p-1}$. Thus there exists $\mathbf{x} \in S_1 \times \dots \times S_n$ such that

$$a_{2p-1} + \sum_{i=1}^{p-1} x_i = 0 \pmod{p}$$

Which gives us what we wanted.

Only two things could prevent us from being able to apply the permanent lemma. First the permanent could be zero, but $perm(A) = (p-1)! \neq 0$. Indeed, each of the summands in the definition of the permanent are 1 and there are $|Sym_{p-1}| = (p-1)!$ of them, where by Sym_{p-1} we denote the symmetric group of $p-1$ elements. Second, the sets S_i could be in fact of size 1 (we could have $a_i = a_{i+p-1}$). But in that case, as the elements are ordered, we would have a subsequence of p elements all equal to each other: $a_i = a_{i+1} = \dots = a_{i+p-1}$. But then there is no need to apply the permanent lemma as this subsequence already sums to zero. □

Direct proof. Suppose not. Let $n = 2p-1$. For each a_i in our sequence $a_1, \dots a_n$ we will assign a variable x_i with again the idea that it is a selector for a_i. We would like to design a polynomial f with the property that $f(x_1, \dots, x_n) = 1$ is equivalent to having exactly p variables x_i taking value 1 and the sum of their associated elements is 0 (mod p). Once this is done, we would like to apply the second combinatorial

Nullstellensatsz to the cartesian product $\{0,1\}^n$. In other words, this will give us a satisfiable assignment of f.

To take care of the first part, we notice the following: for $x_i \in \{0,1\}$ we have $\sum x_i = 0$ iff exactly 0 or p of the x_i's are 1. Raising it to the power $p-1$ we get that:

$$(\sum x_i)^{p-1} = \begin{cases} 0, & \text{if exactly } p \text{ or } 0 \text{ elements have been picked} \\ 1, & \text{otherwise} \end{cases}$$

For the second part, notice that our assumption tells us that $\forall a_1, \ldots, a_p$ subsequence of p elements, we have $\sum a_i \neq 0$, i.e. $(\sum a_i)^{p-1} = 1$.

Joining these two gadgets together, we get an expression that takes value 1 iff exactly 0 or p x_i's have been picked and the sum of the corresponding a_i's sum to 0 (mod p):

$$(1 - (\sum x_i)^{p-1}) \cdot (1 - (\sum x_i a_i)^{p-1})$$

But this is not enough to be able to apply our theorem. We still need to take care of the degree of the polynomial, as well as the case when no x_i is picked. This can be done the following way:

$$f(x_1, \ldots, x_n) = (1 - (\sum_{i=1}^{n} x_i)^{p-1}) \cdot (1 - (\sum_{i=1}^{n} x_i a_i)^{p-1}) - \prod_{i=1}^{n}(1 - x_i)$$

By adding this last term, we get that the coefficient of the monomial $\prod_{i=1}^n x_i$ is nonzero, which allows us to apply the theorem on the cartesian product $\{0,1\}^n$. At the same time, we also ensure that when all x_i's are zero, f is also zero. This term is also zero everywhere else and thus does not change the value of the polynomial at any other point. □

1.5 Hypercube Coverings

Theorem 1.5.1. *Let $H_1, \ldots, H_m$ be a family of hyperplanes in $\mathbb{R}^n$ that cover all vertices of the unit cube $\{0,1\}^n$, but one. Then $m \geq n$.*

First we note that this bound is tight. Indeed, if we define the ith hyperplane to contain all points of the hypercube such that the ith coordinate is 1, then clearly the origin is not in any of these, and all points with a nonzero coordinate are contained in one of the hyperplanes. Also these are well defined hyperplanes: their equation is $x_i = 1$.

Proof. WLOG, we may assume the uncovered vertex is the origin. To each hyperplane H_i is associated an equation: $a_1x_1 + \ldots + a_nx_n = b_i$. Let $g_i(x_1, \ldots, x_n) = -b_i + \sum_{j=1}^n a_jx_j$. It is a polynomial of degree 1.

Suppose for contradiction that $m < n$. We know that for each point $(s_1, \ldots, s_n) \in \{0,1\}^n \backslash \{0\}$, the product $\prod_{i=1}^m g_i(s_1, \ldots, s_n)$ is zero as they are each contained in at least one of the hyperplanes. Further, we can see that this product will be some non-zero constant $1/\delta$ (the product of the constant terms of the g_i's) when evaluated at the origin.

Define now the following polynomial:

$$f = \prod_{i=1}^{m} g_i(x_1, \dots, x_n) - \delta \prod_{i=1}^{n}(1 - x_i)$$

By the above observations, f vanishes entirely on $\{0,1\}^n$.

However, by assumption $m < n$, which means that the degree of the first product is less the degree of the second, so we conclude that the coefficient of $\prod_{i=1}^{n} x_i$ is (in absolute value) δ, thus nonzero. Applying the Nullstellensatsz to $\{0,1\}^n$ we get a contradiction. □

Following is a variant of the previous statement that appeared in the 2007 IMO.

IMO 2007 Question 6. *Let n be a positive integer. Consider $S = \{(x,y,z)|x,y,z \in \{0,1,\dots,n\}, x+y+z > 0\}$ as a set of $(n+1)^3$ points in 3D space.*

Determine the smallest number of planes, the union of which contains S but not $(0,0,0)$.

One can easily find a collection of $3n$ such planes, but without knowing about the combinatorial Nullstellensatsz, showing that it's impossible to find a smaller such collection of planes seems quite hard.

Proof. The proof is virtually the same as for the previous theorem. Let m be the size of the collection of planes. We define the following polynomial:

$$\prod_{i=1}^{m} g_i(x,y,z) - \delta \prod_{i=1}^{n}(1 - x/i)(1 - y/i)(1 - z/i)$$

where by $g_i(x,y,z)$ we denote the equation of the ith plane. Then by the same arguments as above, we derive a contradiction when $m < 3n$. □

1.6 Exercises

Proposition 1.6.1. *Let $\mathcal{F}$ be a set system such that $|\mathcal{F}| > 2|X|$, where $X = \bigcup_{F \in \mathcal{F}} F$. There is a nonempty subcollection $\mathcal{F}' \subset \mathcal{F}$, such that for each element $x \in X$, the number of sets in $\mathcal{F}'$ containing x is divisible by 3.*

Chapter 2

The Kakeya Set Problem

2.1 Introduction

The Kakeya conjecture was first formaly stated by the japanese mathematician Soichi Kakeya in 1917: "What is the least area in the plane required to continuously rotate a needle of unit length and zero thickness around completely (i.e. by 360 degrees)?". We call such sets Kakeya needle sets.

Even though it may seem very unintuitive at first, it turns out, as Besicovitch showed in 1928 [Bes28], that one can achieve this using an arbitrarily small amount of area. His construction relied on the fact that translating the needle requires a set of zero measure.

This observation closed the question in $\mathbb{R}^2$, but there is a generalization of the statement to $\mathbb{R}^n$ that remains open. It has little to do with the polynomial method and thus won't be treated, but can be found in section 2 of [Fur08].

Here, we're interested by an adaptation of this problem to finite fields proposed by Wolf in 1999. We begin by defining the central notion of Kakeya sets in $\mathbb{R}^n$.

Definition 2.1.1. *A Kakeya set $K \subseteq \mathbb{R}^n$ is a set such that it contains a unit line segment in every direction.*

It is important to note that this notion is different from the notion of Kakeya needle sets as it is not required that we may continuously rotate a unit length segment within it. In fact, one can construct Kakeya sets of zero (Lebesgue) measure, while Kakeya needle sets may only be arbitrarily small.

This being said, if we try to adapt the notion of a Kakeya set to finite fields, we can drop the continuity condition as we work over finitely many elements. There are also other notions such as "unit length", or simply "length of a segment" in $\mathbb{F}^n$ that are hard to define. It seems to be simpler to drop this altogether and to consider sets which contain lines in every direction.

Definition 2.1.2. *A Kakeya set $K \subseteq \mathbb{F}^n$ is a set such that $\forall y \in \mathbb{F}^n$, there is a line with direction y that is contained in K.*

There is one last notion that needs to be adapted. That is the "area" of a set. It was agreed that it would correspond to the proportion of the field that the subset represents. With all these modified notions in mind, the adaptation of the Kakeya conjecture to finite fields is natural:

Theorem 2.1.3 (Finite field Kakeya conjecture)**.** *Let $\mathbb{F}_q$ be a finite field of size q, and let $K \subseteq \mathbb{F}^n$ be a Kakeya set. Then $|K| \geq c_n q^n$, where $c_n > 0$ does not depend on q.*

In other words, as the size of the underlying field grows, the size of the smallest Kakeya sets will remain the same.

The motivation behind this adaptation was to get a better understanding of Kakeya sets in $\mathbb{R}^n$ by looking at a similar problem in a simpler setting, which solution could then be generalized to the full $\mathbb{R}^n$ conjecture. For quite some time it was thought that the finite version was just as hard as the infinite version.

That was until using the polynomial method, Dvir was able to find the initial bound of $|K| \geq c_n q^{n-1}$ [Dvi09]. Immediately after, Alon and Tao managed to slightly change the arguments of the proof proposed by Dvir and obtain a better bound of $|K| \geq c_n q^n$ [Tao13], which closes the finite field Kakeya conjecture.

However, the lower bound obtained was still far from the best known upper bound of
$\frac{1}{2^{n-1}}q^n + o(q^{d-1})$ [SS+08]. After polishing once more the arguments of Dvir, Saraf and Sudan et al. [DKSS09] were able close this gap by providing a nearly matching lower bound of $\frac{1}{2^{n-1}}q^n$.

Here, we will present both Dvir's very first result, as well as some of the main improvements that have been made to it. In the very last section we will see the construction that gives the upper bound mentionned earlier.

2.2 Dvir's proof for quantitative Kakeya sets

The very first bound proposed by Dvir was presented as a corollary of a theorem for "quantitative" Kakeya sets.

Definition 2.2.1. *Let $\mathbb{F}_q$ be a finite field. A (γ,δ)-Kakeya set $K \subseteq \mathbb{F}^n$ is a set such that there exists some set $L \subset \mathbb{F}^n$ of size at least $\delta \cdot q^n$ such that $\forall y \in L$, there is a line with direction y that intersects K in at least $\gamma \cdot q$ points.*

In short, it is a set such that for a large proportion of directions, there is a line with that direction that has a large intersection with the set. The δ parameter being for the number of directions, and γ for the size of the intersections.

We can also notice that for $\gamma = \delta = 1$, K is a regular Kakeya set.

Theorem 2.2.2. *Let $K \subseteq \mathbb{F}^n$ be a (γ,δ)-Kakeya set. Then*

$$|K| \geq \binom{d+n-1}{n-1}$$

where

$$d = q \cdot min\{\delta, \gamma\} - 2$$

To prove this, we will first find a certain polynomial of small degree that vanishes on our Kakeya set $K \subseteq \mathbb{F}^n$. Then we will show that it needs to vanish on the whole field, which leads to a contradiction.

In Dvir's original paper, he uses the Schwartz-Zippel lemma to derive the contradiction. Since the lemma was in the original paper, and is an interesting statement on its own, we will use it as written in the paper. But it should be noted that the use of Schwartz-Zippel can be entirely avoided, and in hindsight it seems more appropriate to use the combinatorial nullstellensatz.

Lemma 2.2.3. *(Schwartz-Zippel) Let $\mathbb{F}$ be a finite field of size q, $f \in \mathbb{F}[x_1, ..., x_n]$ a polynomial of degree at most d. The set of points Z on which f vanishes satisfies:*

$$|Z| \leq dq^{n-1}$$

Proof. By induction on n. WLOG we can assume that $d < q$, otherwise, the statement is true.
The base case is the trivial one dimensional case.
When $n > 1$, the idea is to fix the value of one dimension and to count the zeros of the resulting $n-1$ variate polynomial by considering two cases: when it is entirely zero and when it is not. We define $f_t(x_1, \ldots, x_{n-1}) = f(x_1, \ldots, x_{n-1}, t)$ to be the $n-1$ variate polynomial obtained by assigning the fixed value t to the nth variable of f. If f_t vanishes entirely, by the induction hypothesis, it needs to be the zero polynomial as it has degree less than q. Hence, viewing f as a polynomial in the nth coordinate, it is zero for this value t, so we can write

$$f(x_1, \ldots, x_n) = (x_n - t)g_0(x_1, \ldots, x_n)$$

Indeed, If we denote by E the set of all such points t, we can repeatedly factor out such points until we get:

$$f(x_1, \ldots, x_n) = g(x_1, \ldots, x_n) \prod_{t \in E} (x_n - t)$$

where $g(x_1, \ldots, x_n)$ is now a polynomial of degree at most $d - |E|$ that does not vanish entirely when $x_n \in E$. This means that when we fix a value t along the nth axis we have two possibilities for the corresponding hyperplane: either $t \in E$ and the entire hyperplane vanishes, yielding q^{n-1} solutions, or $t \notin E$ and the number of solutions is precisely the number of solutions of $g(x_1, \ldots, x_{n-1}, t)$, which can now be seen as a $n-1$ variable polynomial and so we can apply the induction hypothesis. We thus get:

$$Z(f) = \bigcup_{t \in E} (F^{n-1} \times \{t\}) \cup \bigcup_{t' \notin E} (Z(g) \times \{t'\})$$

Which gives

$$\begin{aligned}
&|Z(f)| \leq |E|q^{n-1} + (d - |E|)q^{n-2}|E| \\
&|Z(f)| \leq dq^{n-1}
\end{aligned}$$

□

Now the proof of the main theorem.

Proof. Suppose by contradiction that we have

$$|K| < \binom{d+n-1}{n-1} \tag{2.1}$$

The number $\binom{d+n-1}{n-1}$ is precisely the dimension of the space W of homogeneous polynomials of degree d.

Indeed, the dimension of this space is equal to the number of n variable monomials of degree exactly d. And this number can be counted as follows: we're distributing d units of degree among n variables, with possibly variables with degree 0. This is the same as grouping d stars into n groups by listing them in line and inserting $n-1$ separators (bars) in between them. In other words, we have a total of $d+n-1$ elements, and we're picking $n-1$ elements among them to be separators. This counting argument is commonly referred to as the "stars and bars" technique.

So there must exist a nonzero homogeneous polynomial g of degree d which vanishes on K.

This can be justified by considering the evaluation map $e \colon W \to \mathbb{F}^{|K|}$ that associates to $f \in W$ the tuple of values that it takes on K, i.e. $e(f) = (f(x))_{x \in K}$. It is a linear isomorphism from one vector space to another with the second one being of strictly smaller dimension, which means that the kernel of the map is non trivial.

Recall that by definition of a (γ, δ)-Kakeya set, there exists some set $L \subset \mathbb{F}^n$ of size at least $\delta \cdot q^n$ such that $\forall y \in L$, there is a line with direction y that intersects K in at least $\gamma \cdot q$ points. We take such a set L and claim that g vanishes entirely on it as well.

Once this claim is proven, the theorem follows as we have

$$\begin{aligned}
d &= q \cdot min\{\delta, \gamma\} - 2 \\
\Rightarrow \frac{d}{q} &< \delta \\
\text{recall that by definition of L, } |L| &\geq \delta \cdot q^n \\
\Rightarrow |L| &> dq^{n-1}
\end{aligned}$$

which is a contradiction with the Schwartz-Zippel lemma.

Proof of the claim.

Let $y \in L$. By hypothesis, there is some line $L_{z,y} = \{z + ty \colon t \in \mathbb{F}\}$ which intersects

K in at least γq points. By choice of q, this means that $|L_{z,y} \cap K| \geq d+2$. In particular, since g vanishes on K, this means that g vanishes on at least $d+2$ points of the form $z + a_i y$. If one of them is zero, we still have at least $d+1$ non-zero points on which g vanishes. Since these points are non-zero, we can multiply by their inverse and get points of the form $a_i^{-1}z + y$ that are now on the line $L_{y,z} = \{y + tz \colon t \in \mathbb{F}\}$.

Moreover, because g is homogeneous, g vanishes on these d+1 points. Indeed, if $x \in \mathbb{F}^n$ is a root of g and $\lambda \in \mathbb{F}$ is a scalar, $g(\lambda x) = \lambda^d g(x) = 0$. But this implies that g must vanish entirely on $L_{y,z}$, in particular when $t = 0$ i.e. on y. Thus $g(y) = 0$.

□

From there, one can set $\gamma = \delta = 1$ to get the bound. Indeed,

$$\begin{aligned} |K| &\geq \binom{q-2+n-1}{n-1} \\ &\geq \frac{(q-1)^{n-1}}{(n-1)!} \\ \text{when q is large } &\sim \frac{q^{n-1}}{(n-1)!} \end{aligned}$$

We conclude:

Theorem 2.2.4. *Let $K \in \mathbb{F}^n$ be a Kakeya set. Then*

$$|K| \geq c_n \cdot q^{n-1}$$

where c_n is independent of q.

From there, Dvir observes that the product of Kakeya sets is a Kakeya set and uses this fact to obtain an even better bound.

We can check: if we have $K_1 \subset \mathbb{F}^k$ and $K_2 \subset \mathbb{F}^l$ Kakeya sets, and let $x \in \mathbb{F}^k \times \mathbb{F}^l$ be an arbitrary direction. Denote by x_1 it's coordinates in $\mathbb{F}^k$ and by x_2 it's coordinates in $\mathbb{F}^l$. Then by the Kakeya property of K_1 and K_2 there are elements z_1 and z_2 such that the lines L_{z_1,x_1} and L_{z_2,x_2} (same notation as before) are contained in K_1, K_2 respectively. Writing $z = (z_1, z_2)$, we see that the line $L_{z,x}$ is contained in $K_1 \times K_2$.

Corollary 2.2.4.1. *For every integer n and every $\epsilon > 0$, there exists a constant $c_{n,\epsilon}$ depending only on n and ϵ such that any Kakeya set $K \subset \mathbb{F}^n$ satisfies*

$$|K| \geq c_{n,\epsilon} \cdot q^{n-\epsilon}$$

With the above observation and the previously shown bound in mind, the proof is very straightforward.

Proof. Let K be a Kakeya set in $\mathbb{F}^n$ and consider its product with itself r times, where $r > 0$ is an integer. By the above observation it is still a Kakeya set in $\mathbb{F}^{nr}$. Applying the known bound to this newly obtained set, we get

$$\begin{aligned} |K|^r &\geq c_{nr} \cdot q^{nr-1} \\ |K| &\geq c_{n,r} \cdot q^{n-1/r} \end{aligned}$$

□

This concludes the original statements proposed by Dvir, but it is not the end of the story just yet.

2.3 The Kakeya conjecture

Shortly after Dvir published his paper, Alon and Tao managed to modify his arguments slightly and prove the Kakeya conjecture [Tao13]. The theorem is the following

Theorem 2.3.1. *Let $K \subseteq \mathbb{F}_q^n$ be a Kakeya set. Then*

$$|K| \geq \binom{q+n-1}{n}$$

Proof. First we recall that using the stars and bars technique, it is possible to count the dimension of the space of polynomials of degree at most $q-1$. Indeed, adding a dummy variable to make the sum of the degrees equal $q-1$, we're choosing n bars out of $q-1+n$ elements. Thus the dimension of this space is $\binom{q+n-1}{n}$.

Suppose by contradiction that

$$|K| < \binom{q+n-1}{n}$$

Then, as stated earlier, there exists a nonzero polynomial $f \in \mathbb{F}[x_1, \ldots, x_n]$ of degree at most $q-1$ that vanishes on the whole of K. Let d be its degree. We will prove that it needs to vanish on the whole field $\mathbb{F}^n$.

As K is a Kakeya set, we have that for each nonzero direction $\mathbf{v} \in \mathbb{F}^n$, there is some point $\mathbf{x} \in \mathbb{F}^n$ such that $\forall t \in \mathbb{F}$, $f(\mathbf{x} + t\mathbf{v}) = 0$. Restricting f to points of this line, we can view it as a one dimensional polynomial:

$$\begin{aligned} g(t) &= f(x_1 + tv_1, \ldots, x_n + tv_n) \\ &= f_d(\mathbf{v})t^d + [\text{terms of degree } < d] \end{aligned}$$

where f_d is the polynomial of degree at most d in $\mathbf{v}$ denoting the coefficient of the term t^d. It is important to note that it does not depend on $\mathbf{x}$ at all, so if we let $\mathbf{x}$ vary, this polynomial will remain the same. We note that for fixed $\mathbf{x}$ and $\mathbf{v}$ the polynomial g is zero for each $t \in \mathbb{F}$. As its degree is smaller than q, it must be the zero polynomial, which implies $f_d(\mathbf{v}) = 0$. Now repeating this process for each direction $\mathbf{v} \in \mathbb{F}^n$, we obtain that f_d is a polynomial of degree at most $q-1$ vanishing on the whole of $\mathbb{F}^n$.

There are two similar ways to conclude the proof.
The first is the Schwartz-Zippel lemma: the number of zeros of f_d should be at most $(q-1)q^{n-1}$.
The second is the combinatorial Nullstellensatz: letting $S_i = \mathbb{F}$, we get that there is a nonzero homogeneous polynomial of degree $d < q$ that vanishes on $S_1 \times \ldots \times S_n$. Taking any of it's terms with nonzero coefficient, we get a contradiction.

□

From there, using the same reasoning as for the result for quantitative Kakeya sets, we derive

$$|K| \geq \frac{q^n}{n!}$$

2.4 Sharper bounds

To obtain sharper bounds, Saraf and Sudan [SS$^+$08] use what they call the "extended method of multiplicities". The main idea is the same as in the previously seen proofs, but this time we will find a polynomial that vanishes on the Kakeya set with high multiplicity, and then show that its homogeneous part will need to also vanish on the whole field with high multiplicity. And this fact will force the nearly sharp upper bound of $\frac{1}{2^n}q^n$ on the size of Kakeya sets.

Throughout this section, given a vector of integers $\mathbf{i} = (i_1, \ldots, i_n) \in \mathbb{Z}^n$, we denote its weight $\sum_{j=1}^n i_j$ by $wt(\mathbf{i})$. We will also occasionally use the notation $\mathbf{x}^{\mathbf{i}}$ for $\prod_{j=1}^n x_j^{i_j}$.

A key notion in the method of multiplicities is the Hasse derivative. We begin by defining it, along with the notion of multiplicity.

Definition 2.4.1 (Hasse derivative). *Let $f \in \mathbb{F}[\mathbf{x}]$ and $\mathbf{i}$ be a nonnegative vector. The $\mathbf{i}$th Hasse derivative of f, denoted $f^{(\mathbf{i})}(\mathbf{x})$ is the coefficient of $\mathbf{z}^{\mathbf{i}}$ in the polynomial $\widetilde{f}(\mathbf{x}, \mathbf{z}) = f(\mathbf{x} + \mathbf{z}) \in F[\mathbf{z}]$. That is, $\widetilde{f}$ viewed as a function in $\mathbf{z}$ only.*

In other words,

$$f(\mathbf{x} + \mathbf{z}) = \sum_i f^{(\mathbf{i})}(\mathbf{x})\mathbf{z}^{\mathbf{i}} \tag{2.2}$$

Definition 2.4.2 (multiplicity). *Let $f \in \mathbb{F}[\mathbf{x}]$ and $\mathbf{a} \in \mathbb{F}^n$. The multiplicity of f at $\mathbf{a}$ is defined to be the largest integer m such that for all $\mathbf{i}$ with $wt(\mathbf{i}) < m$, $f^{(\mathbf{i})}(\mathbf{a}) = 0$.*

For example, if the multiplicity of f at $\mathbf{a}$ is precisely the degree of f, then the polynomial $\widetilde{f}(\mathbf{a}, \mathbf{z})$ is homogeneous of degree $deg(f)$.

We now list a few properties of multiplicities that will be usefull soon. In some of the proofs we will use properties of the Hasse derivative. We will not prove these here, but they can be found in the original paper by Saraf and Sudan [SS$^+$08].

Lemma 2.4.3. *Let $f \in \mathbb{F}[\mathbf{x}]$ and $\mathbf{i}$ a vector of nonnegative integers. Denote by H_g the homogeneous part of highest degree of a polynomial g. Then either $H_f^{(\mathbf{i})} = H_{f^{(\mathbf{i})}}$ or $H_f^{(\mathbf{i})}$ is the zero polynomial.*

We will take the above statement for granted.

Lemma 2.4.4. *If $f(\mathbf{x}) \in \mathbb{F}[\mathbf{x}]$ and $\mathbf{a} \in \mathbb{F}^n$ are such that $mult(f, \mathbf{a}) = m$, then $mult(f^{(i)}, \mathbf{a}) \geq m - wt(\mathbf{i})$.*

Put in words, the multiplicity of the $\mathbf{i}$th derivative of f at some point $\mathbf{a}$ cannot be less than the multiplicity of f at $\mathbf{a}$ minus the weight of the vector $\mathbf{i}$.

Proof. By hypothesis, for any vector $\mathbf{k}$ with $wt(\mathbf{k}) < m$, we have $f^{(\mathbf{k})}(\mathbf{a}) = 0$. Now, if we take any vector $\mathbf{j}$ such that $wt(\mathbf{j}) < m - wt(\mathbf{i})$, then by a basic property of the Hasse derivative, we have a relation between the $\mathbf{i} + \mathbf{j}$ derivative of f and the jth derivative of the ith derivative of f:

$$\binom{\mathbf{i}+\mathbf{j}}{\mathbf{j}} f^{(\mathbf{i}+\mathbf{j})}(x) = (f^{(\mathbf{i})})^{(\mathbf{j})}(x)$$

but since $wt(\mathbf{i} + \mathbf{j}) = wt(\mathbf{i}) + wt(\mathbf{j}) < m$, we get that $(f^{(\mathbf{i})})^{(\mathbf{j})}(\mathbf{a}) = f^{(\mathbf{i}+\mathbf{j})}(\mathbf{a}) = 0$. Since the $\mathbf{j}$th derivative of $f^{(\mathbf{i})}$ is zero at $\mathbf{a}$ for all $\mathbf{j}$ such that $wt(\mathbf{j}) < m - wt(\mathbf{i})$, we conclude that $mult(f^{(\mathbf{i})}, \mathbf{a}) \geq m - wt(\mathbf{i})$. □

In the original paper by Saraf and Sudan, the following lemma is derived as a corollary to a more general statement about the behavior of multiplicities under composition of polynomial tuples. Here we will only look at the case where we compose with a line.

Lemma 2.4.5 (multiplicities under when restricting to a line)**.** *Let $f(\mathbf{x}) \in \mathbb{F}[\mathbf{x}]$ and $\mathbf{x}$, $\mathbf{v} \in \mathbb{F}^n$. Let $f_{\mathbf{x},\mathbf{v}}(t)$ be the single variable polynomial formed by restricting f to the line $\mathbf{x} + t\mathbf{v}$, i.e. $f(\mathbf{x} + t\mathbf{v}) \in \mathbb{F}[t]$. Then for any fixed $\alpha \in \mathbb{F}$,*

$$mult(f_{\mathbf{x},\mathbf{v}}, \alpha) \geq mult(f, \mathbf{x} + \alpha\mathbf{v})$$

Said in words, the multiplicity at a point of $\mathbb{F}$ of the restriction of f to a line is at least the multiplicity of the original f at the corresponding point in $\mathbb{F}^n$.

The proof of this statement is left out, but can be found in the original paper.

Lemma 2.4.6 (Schwartz-Zippel lemma for multiplicities)**.** *Let $f \in \mathbb{F}^n$ be a non-zero polynomial of total degree at most d. Then,*

$$\sum_{\mathbf{v} \in \mathbb{F}^n} mult(f, \mathbf{v}) \leq d \cdot q^{n-1} \tag{2.3}$$

Proof. By induction on n.
When $n = 1$, we show that if $mult(f, \alpha) = m$, then $(x - \alpha)^m$ divides $P(x)$. This can be seen from the definition of multiplicity: $f^{(i)}(\alpha) = 0$ for all i such that $wt(i) < m$, which implies that z^m divides $f(\alpha + z) = \sum f^{(i)}(\alpha)z^i$. Substituting $x = z - v$ we see that $(x - \alpha)^m$ divides $f(x)$.

When $n > 1$, consider f as a one dimensional polynomial in terms of the last variable x_n:

$$f(x_1, \ldots, x_n) = \sum_{j=0}^{k} f_j(x_1, \ldots, x_{n-1})x_n^j \tag{2.4}$$

where $f_k(x_1, \ldots, x_{n-1}) \neq 0$ (i.e. k is the degree of f in the nth variable) and WLOG $k > 0$ (otherwise we are in the $n - 1$ case and thus we get the bound directly from the induction hypothesis).

For any tuple $(\alpha_1, \ldots, \alpha_{n-1}) \in \mathbb{F}^{n-1}$ denote by $m_{\alpha_1,\ldots,\alpha_{n-1}}$ the multiplicity of f_k at that point, i.e. $m_{\alpha_1,\ldots,\alpha_{n-1}} = mult(f_k, (\alpha_1, \ldots, \alpha_{n-1}))$. The claim is that for each such tuple,

$$\sum_{\beta \in \mathbb{F}} mult(f, (\alpha_1, \ldots, \alpha_{n-1}, \beta)) \leq m_{\alpha_1,\ldots,\alpha_{n-1}} \cdot q + k$$

Once this claim is proven, the result follows. Indeed, summing over all $(\alpha_1, \ldots, \alpha_{n-1}) \in \mathbb{F}^{n-1}$ we get the bound

$$\sum_{(\alpha_1,\ldots,\alpha_n) \in \mathbb{F}^n} mult(f, (\alpha_1, \ldots, \alpha_n)) \leq k \cdot q^{n-1} + \sum_{(\alpha_1,\ldots,\alpha_{n-1}) \in \mathbb{F}^{n-1}} m_{\alpha_1,\ldots,\alpha_{n-1}} \cdot q \tag{2.5}$$

But f was a polynomial of degree d, so the expression $f_k(x_1, \ldots, x_{n-1})x_n^k$ also has degree at most d. In particular, the polynomial f_k has degree at most $d-k$. Applying the induction hypothesis to f_k we are able to bound the rightmost term in the above inequality:

$$\begin{aligned}\sum_{(\alpha_1,\ldots,\alpha_{n-1}) \in \mathbb{F}^{n-1}} m_{\alpha_1,\ldots,\alpha_{n-1}} &\leq deg(f_k) \cdot q^{n-2} \\ &\leq (d-k) \cdot q^{n-2}\end{aligned}$$

Combining this with equation 2.5, we get the result.

We now prove the claim.

Fix $\alpha_1, \ldots, \alpha_{n-1} \in \mathbb{F}^{n-1}$ and take $\mathbf{i} = (i_1, \ldots, i_{n-1})$ such that $wt(\mathbf{i}) = m_{\alpha_1,\ldots,\alpha_{n-1}}$ and $f_k^{(\mathbf{i})}(x_1, \ldots, x_{n-1}) \neq 0$ (note that by definition the multiplicity is the smallest integer such that this is possible). Let $(\mathbf{i}, 0)$ denote the vector $(i_1, \ldots, i_{n-1}, 0)$; we note that

$$f^{(\mathbf{i},0)}(x_1, \ldots, x_n) = \sum_{j=0}^{k} f_j^{(\mathbf{i})}(x_1, \ldots, x_{n-1})x_n^j$$

Indeed, plugging in $\mathbf{x} + \mathbf{z}$ in the equation 2.4, one can see that the coefficient of each term $\mathbf{z}^{(\mathbf{i},0)} \cdot x_n^j$ is precisely the $\mathbf{i}$th derivative of f_j. Factoring out $\mathbf{z}^{(\mathbf{i},0)}$ we get the above formula.

In particular, this tells us that $f^{(\mathbf{i},0)}$ is not the zero polynomial as we picked $f_k^{\mathbf{i}}$ to be nonzero.

We use lemma 2.4.4 to get

$$mult(f, (\alpha_1, \ldots, \alpha_n)) \leq wt((\mathbf{i}, 0)) + mult(f^{(\mathbf{i},0)}, (\alpha_1, \ldots, \alpha_n))$$

Now recall that by choice of $\mathbf{i}$, we have $wt(\mathbf{i}, 0) = m_{\alpha_1,\ldots,\alpha_{n-1}}$. Moreover, applying 2.4.5 with $\mathbf{x} = (\alpha_1, \ldots, \alpha_{n-1}, 0)$ and $\mathbf{v} = (0, \ldots, 0, \alpha_n)$, we obtain the bound

$$mult(f, (\alpha_1, \ldots, \alpha_n)) \leq m_{(\alpha_1,\ldots,\alpha_{n-1})} + mult(f^{(\mathbf{i},0)}(\alpha_1, \ldots, \alpha_{n-1}, t), \alpha_n)$$

where $f^{(\mathbf{i},0)}(\alpha_1, \ldots, \alpha_{n-1}, t)$ is a single variable polynomial in t of degree k.

We sum over all $\alpha_n \in \mathbb{F}^n$, apply the base case to $f^{(\mathbf{i},0)}(\alpha_1, \dots, \alpha_{n-1}, t)$ and get

$$\sum_{\alpha_n \in \mathbb{F}} mult(f, (\alpha_1, \dots, \alpha_n)) \leq m_{\alpha_1, \dots, \alpha_{n-1}} \cdot q + k$$

as desired.

□

Lemma 2.4.7 (Interpolation lemma for multiplicities). *Given a set $K \subset \mathbb{F}^n$ and nonnegative integers m, d such that*

$$\binom{m+n-1}{n} \cdot |K| < \binom{d+n}{n} \tag{2.6}$$

there exists a nonzero polynomial $P = P_{m,K} \in \mathbb{F}[x]$ of total degree at most d such that $mult(P, a) \geq m$ for every $a \in K$.

Proof. The number of possible monomials for P is $\binom{d+n}{n}$.
Additionally, for each $a \in K$, the condition $mult(P, a) \geq m$ imposes $\binom{m+n-1}{n}$ constraints on the coefficients of P. Indeed, for each vector i such that we have $wt(i) < m$, there is one constraint $P^{(i)}(a) = 0$. Since the total number of linear constraints is $\binom{m+n-1}{n} \cdot |K|$, which is strictly less then the number of coefficients, there is a nontrivial polynomial that satisfies these constraints. □

We are now ready to present Saraf and Sudan's improved lower bound on the size of Kakeya sets.

The proof is basically line for line the same as the one exposed in the previous section, with the sole difference that the polynomial we're now considering vanishes with high multiplicity on K, and then using the above technical lemmas and the same trick as before we're able to show that its homogeneous part vanishes with high multiplicity on the whole field, showing the contradiction. The resulting bound is then much better.

Theorem 2.4.8 (Improved lower bound on the size of Kakeya sets). *If $K \subset \mathbb{F}^n$ is a Kakeya set, then $|K| \geq \frac{q^n}{(2-1/q)^n}$*

Proof. Let l be a large multiple of q and let

$$m = 2l - l/q \tag{2.7}$$

$$d = lq - 1 \tag{2.8}$$

d will be the bound on the degree of the polynomial vanishing on K.
m will be the multiplicity of the zeros of P on K.
l will be the multiplicity of the zeros of the homogeneous part of P.
Because of this choice of d, m and l we have the following list of facts that we will use throughout the proof:

$$d < lq \tag{2.9}$$

$$(m-l)q > d-l \tag{2.10}$$

If we have some d^* such that $d^* \geq m$, then by 2.7,

$$d^* \geq m \geq l \tag{2.11}$$

We will first prove that

$$|K| \geq \frac{\binom{d+n}{n}}{\binom{m+n-1}{n}}$$

Suppose not. Then by the interpolation lemma for multiplicities, there exists some polynomial $f \in \mathbb{F}[\mathbf{x}]$ such that f vanishes with multiplicity at least m at each point of K and that has degree at most d, say d^*. Because of this property we have $d^* \geq m$ and recalling 2.11, $d^* \geq l$. We isolate the homogeneous part of f:

$$f(\mathbf{x}) = f_{d^*}(\mathbf{x}) + [\text{terms of degree } < d^*]$$

We will now show that for each direction $\mathbf{v} \in \mathbb{F}^n$, $mult(f_{d^*}, \mathbf{v}) \geq l$.

Once this claim is proven, we get a contradiction with the Schwartz-Zippel lemma for multiplicities 2.4.6.

Indeed, it states that $d^* \cdot q^{n-1} \geq \sum_{\mathbf{v} \in \mathbb{F}^n} mult(f, \mathbf{v})$, but the above claim gives us

$$lq^n \leq \sum_{\mathbf{v} \in \mathbb{F}^n} mult(f, \mathbf{v})$$

And since we also know by 2.9 that $d^* \cdot q^{n-1} < lq^n$, we get a contradiction.

We now prove the claim.
Let $\mathbf{v} \in \mathbb{F}^n$.
We would like to show that for every vector $\mathbf{i}$ such that $wt(\mathbf{i}) < l$ the $\mathbf{i}$th Hasse derivative of the homogeneous part of f is zero at $\mathbf{v}$, i.e. $(f_{d^*})^{(\mathbf{i})}(\mathbf{v}) = 0$. Pick such a vector $\mathbf{i}$.

Since K is a Kakeya set, we can find an $\mathbf{x} \in \mathbb{F}^n$ such that the line $\mathbf{x} + t\mathbf{v} \colon t \in \mathbb{F}$ is entirely contained in K. Then by a previous lemma (2.4.4), the multiplicity of the $\mathbf{i}$th derivative of f at any point $\mathbf{a} \in \mathbb{F}^n$ is at most $m - wt(\mathbf{i})$. In particular this is true at each point of the line $\mathbf{x} + t\mathbf{v}$, i.e.

$$mult(f^{(\mathbf{i})}, (\mathbf{x} + t\mathbf{v})) \geq m - wt(\mathbf{i}) \tag{2.12}$$

Consider now the restriction of the $\mathbf{i}$th derivative of f to the line $\mathbf{x} + t\mathbf{v}$:

$$h_{\mathbf{x},\mathbf{v}}(t) = f^{(\mathbf{i})}(\mathbf{x} + t\mathbf{v})$$

Then if $h_{\mathbf{x},\mathbf{v}}$ is of degree d', we have $d' \leq d^* - wt(\mathbf{i})$.

Since the line $\mathbf{x} + t\mathbf{v}$ is contained in K, we know from a previous lemma on the behavior of multiplicities under composition (2.4.5) that $mult(h_{\mathbf{x},\mathbf{v}}, t) \geq mult(f^{(\mathbf{i})}, \mathbf{x}+t\mathbf{v})$, joining this with 2.12 we have $mult(h_{\mathbf{x},\mathbf{v}}, t) \geq m - wt(\mathbf{i})$ for every point $t \in \mathbb{F}$.

Now, since $wt(\mathbf{i}) \leq l - 1$ and

$$(m - l) \cdot q > d^* - l$$

(by choice of parameters, 2.10), we get the following:

$$\sum_{t \in \mathbb{F}} mult(h_{\mathbf{x},\mathbf{v}}, t) \cdot q \geq (m - wt(\mathbf{i})) \cdot q > d^* - wt(\mathbf{i}) \geq deg(h_{\mathbf{x},\mathbf{v}})$$

Which means $h_{\mathbf{x},\mathbf{v}}$ is the zero polynomial, by the Schwartz-Zippel lemma in the case $n = 1$. It implies in particular that the coefficient of the leading term $t^{d'}$ in $h_{\mathbf{x},\mathbf{v}}$ is zero. Notice now that the coefficient of $t^{d'}$ in $h_{\mathbf{x},\mathbf{v}}$ is precisely the homogeneous part of highest degree of the $\mathbf{i}$th derivative of f. Indeed, it becomes clear when writing it out in the same form as in Tao's proof:

$$f^{(\mathbf{i})}(x_1 + tv_1, \ldots, x_n + tv_n) = f^{(\mathbf{i})}_{d'}(\mathbf{v})t^{d'} + [\text{terms of degree } < d']$$

We conclude that $f^{(\mathbf{i})}_{d'}(\mathbf{v})t^{d'} = 0$.

But then by a property of the Hasse derivative 2.4.3, we get in either case that the $\mathbf{i}$th derivative of the highest degree homogeneous part of f is zero when evaluated at $\mathbf{v}$. As this is true for each vector with weight strictly less than l, we conclude that the multiplicity of f_{d^*} at $\mathbf{v}$ is at least l.

Now that this is done, we have the bound

$$|K| \geq \frac{\binom{d+n}{n}}{\binom{m+n-1}{n}}$$

After plugging in the values of m and d into this equation, we follow the same kind of reasoning as in the previous proofs to get the desired bound:

$$\begin{aligned}
\frac{\binom{d+n}{n}}{\binom{m+n-1}{n}} &= \frac{\binom{lq-1+n}{n}}{\binom{2l-l/q+n-1}{n}} \\
&= \frac{(lq-1+n)!(2l-l/q-1)!}{(lq-1)!(2l-l/q+n-1)!} \\
&= \frac{\prod_{i=1}^{n}(lq-1+i)}{\prod_{i=1}^{n}(2l-l/q-1+i)}
\end{aligned}$$

As this is true for all l such that l is a large multiple of q, we may let l tend to infinity:

$$|K| \geq \lim_{l\to\infty} \prod_{i=1}^{n} \left(\frac{q - 1/l + i/l}{2 - 1/q - 1/l + i/l} \right) = \left(\frac{q}{2-1/q} \right)^n$$

.

□

Which yields:

$$\begin{aligned}|K| &\geq \frac{q^{2n}}{(2q-1)^n} \\ &\geq \frac{1}{2^n}q^n\end{aligned}$$

2.5 An Explicit Construction

A construction of a Kakeya set of size at most $\frac{1}{2^{n-1}} \cdot q^n + O(q^{n-1})$ in fields of odd characteristic due to Kopparty [SS+08]. There also exists a similar construction for fields of characteristic 2.

Lemma 2.5.1 (Upper bound for Kakeya sets)**.** *Let $\mathbb{F}_q$ be a field of odd characteristic. Then there exists a Kakeya set K such that*

$$\frac{1}{2^{n-1}} \cdot q^n + O(q^{n-1})$$

Let $D = \{(\beta_1, \dots, \beta_{n-1}, \alpha) \colon \forall i, \beta_i + \alpha^2 \text{ is a square in } \mathbb{F}_q\}$.

The claim is that the following set is Kakeya and has the desired size :

$$K = D \cup (\mathbb{F}^{n-1} \times \{0\}) \tag{2.13}$$

First the size.
To get the size of D let's count the number of distinct squares in $\mathbb{F}$. Consider the polynomial $x^2 - c$, where c is a square. If c is nonzero, this polynomial has exactly two roots, which are $\sqrt{c}$ and $-\sqrt{c}$. But if $c = 0$, zero is the unique root. So if we look at the map that takes elements of $\mathbb{F}$ to their squares, the image has size $\frac{q-1}{2} + 1 = \frac{q+1}{2}$. This means that picking α first, we have $\frac{q+1}{2}$ choices left for each of the β_i's. Thus

$$\begin{aligned}|D| &= \left(\frac{q+1}{2}\right)^{n-1} \cdot q \\ &= \frac{q^n}{2^{n-1}} + O(q^{n-1})\end{aligned}$$

Finally, $|\mathbb{F}^{n-1} \times \{0\}| = q^{n-1}$. So taking the union of these two sets, we get the desired size.

We now show it indeed is Kakeya set.
Take a direction $\mathbf{b} = (b_1, \dots, b_n)$. If $b_n = 0$ then immediately the line $t\mathbf{b}$ is contained in $\mathbb{F}^{n-1} \times \{0\}$ and so it is contained in K.

Suppose now b_n is nonzero. Let $\mathbf{x} = \left(\left(\frac{b_1}{2b_n}\right)^2, \dots, \left(\frac{b_{n-1}}{2b_n}\right)^2, 0\right)$.

Then we have

$$\mathbf{x} + t\mathbf{b} = \left(\left(\frac{b_1}{2b_n}\right)^2 + tb_1, \dots, \left(\frac{b_{n-1}}{2b_n}\right)^2 + tb_{n-1}, tb_n\right)$$

In the notation above it means that $\alpha = tb_n$ and $\beta_i = \left(\frac{b_i}{2b_n}\right)^2 + tb_i$. But squaring the first and summing them it becomes apparent that it is a square:

$$\begin{aligned}\alpha^2 + \beta_i &= (tb_n)^2 + tb_i + \left(\frac{b_i}{2b_n}\right)^2 \\ &= \left(tb_n + \frac{b_i}{2b_n}\right)^2\end{aligned}$$

Thus the line is contained in D and so in K.

Chapter 3

A reformulation using Lagrange Interpolation

In this chapter, we mainly follow the work of Karasev and Petrov [KP12]. In the original theorem, we have a criterion for the existence of a non-root of a polynomial in a large set. There is however a different way to approach this relation between the structure of the polynomial and the sets on which it vanishes. Given a set of point on which the polynomial is zero, we can view this as a constraint on the polynomial. In a sense, we're "forcing" it to pass through certain points in space, and by doing so, we impose a certain structure on it (given of course that its degree is small). In one dimension, there is a well-known such statement, the Lagrange interpolation formula.

Theorem 3.0.1 (Lagrange Interpolation Formula)**.** *Given n points (x_i, y_i) in a field $\mathbb{F}$, there is a unique polynomial f over $\mathbb{F}$ of degree $n-1$ that passes through all the points. It can be found using the Lagrange Interpolation Formula:*

$$f(x) = \sum_{i=1}^{n} y_i \frac{\prod_{j \neq i}(x - x_j)}{\prod_{j \neq i}(x_i - x_j)}$$

When the polynomial is in this form, it is easy to see that the coefficient C of the leading term will be

$$C = \sum_{i=1}^{n} \frac{y_i}{\prod_{j \neq i}(x_i - x_j)}$$

Now, turning this fact around, if we have some polynomial of degree $n-1$, if we consider its values at n distinct points $\{\alpha_1, \ldots, \alpha_n\} =: A$, then it passes through $(\alpha_i, f(\alpha_i))$. Thus its leading coefficient C:

$$C = \sum_{\alpha \in A} \frac{f(\alpha)}{\prod_{\beta \in A, \beta \neq \alpha}(\alpha - \beta)}$$

In particular, if C is nonzero, then one of the $f(\alpha)$'s must be nonzero. Generalizing this statement, we get a version of Alon's Nullstellensatz that gives us information about some of the coefficients of the polynomial.

Theorem 3.0.2 (Reformulation of the Combinatorial Nullstellensatz). *Let $\mathbb{F}$ be a field, $f(x_1, \ldots, x_n)$ a polynomial over that field of degree at most $t_1 + \ldots + t_n$, where $t_1, \ldots, t_n$ are positive integers. Let $S_i \subset \mathbb{F}$ such that $|S_i| = t_i + 1$.*

Denote by C the coefficient of the term $\prod x_i^{ti}$ in f, and by $g_i(x) = \prod_{\alpha \in S_i}(x - \alpha)$. Then,

$$C = \sum_{(\alpha_1,\ldots,\alpha_n)\in S_1\times\ldots S_n} \frac{f(\alpha_1, \ldots, \alpha_n)}{g_1'(\alpha_1)\ldots g_n'(\alpha_n)} \tag{3.1}$$

Note that in the equation above, $g_i'(\alpha_i)$ is simply the product of $(x - \beta)$ over all $\beta \in S_i$, $\beta \neq \alpha_i$.

Unlike the previous applications that had to do with the existence (or non existence) of objects with certain properties, this new formulation of the theorem allows a proof of the Dyson conjecture, which provides a closed form for the constant coefficient of a certain type of polynomials.

Proof. The case $n = 1$ is precisely the above discussion.
$n > 1$. Isolating the term we're interested in, we can write $h = f - C \cdot \prod x_i^{ti}$, where h is a polynomial such that the coefficient of $\prod x_i^{t_i}$ is zero. As 3.1 is linear, it is equivalent to show that the relation holds for h.

Looking at a single term of h, we can see that it has degree in some variable, WLOG the nth, less than t_n (i.e. less than $|S_n|$). So if we fix some tuple $(\alpha_1, \ldots, \alpha_{n-1}) \in S_1 \times \ldots \times S_n$, applying the one dimension case to h restricted to a single variable, we get that

$$\sum_{\beta\in S_n} \frac{f(\alpha_1, \ldots, \alpha_{n-1}, \beta)}{g_1'(\alpha_1)\ldots g_{n-1}'(\alpha_{n-1})g_n'(\beta)} = 0$$

Repeating this over each tuple and every term of h, we obtain the desired relation. □

3.1 The Dyson Conjecture

We now show a first application of the theorem.

Theorem 3.1.1 (Dyson Conjecture). *Let $a_1, \ldots, a_n$ be integers and $a = \sum_{i=1}^n a_i$. In the Laurent polynomial*

$$\prod_{i\neq j}(1 - x_i/x_j)^{a_i}$$

the constant term C is of the form

$$C = \frac{a!}{a_1! \ldots a_n!}$$

This conjecture was first stated in 1962 by Dyson, then proven independently by Wilson and Gunson. The proof below is much shorter.

Proof. First, we observe that the constant coefficient in the above polynomial is in fact the same as the coefficient of the term $\prod x_i^{a-a_i}$ in the polynomial

$$f(x_1, \ldots, x_n) = \prod_{i<j} (-1)^{a_j} (x_j - x_i)^{a_i + a_j}$$

Indeed, if we group each pair $(1 - x_i/x_j)^{a_i}(1 - x_j/x_i)^{a_j}$ and multiply it by $x_j^{a_i} x_i^{a_j}$, we get the term $(-1)^{a_j}(x_j - x_i)^{a_i+a_j}$.

Now that we have a polynomial in a form that is easier to manipulate, here are the main lines of the proof.

The reformulation of the Nullstellensatz gives us the expression of C in function of sets S_i of size $a - a_i + 1$ We will modify the polynomial by adding low degree terms which will not change the coefficient C, but that will give us an expression for C that will be simpler to compute. Namely, we will choose sets S_i such that the modified polynomial $\widetilde{f}$ will have a unique nonzero value on $S_1 \times \ldots \times S_n$. This way, instead of having to compute a sum, we will have an expression of the form

$$C = \frac{f(\beta_1, \ldots, \beta_n)}{\prod g_i'(\beta_i)}$$

For each i, let $S_i = \{0, \ldots, a - a_i\}$. Note that the sets S_i have size $a - a_i + 1$ and so fit exactly in our reformulation of the Nullstellensatz. Another observation to make is that if $\alpha \in S_i$ then the segment $[\alpha, \alpha + a_i - 1]$ is contained in the segment $[0, a - 1]$.

Now, let's define our new function $\widetilde{f}$ from f by replacing the terms $(x_j - x_i)^{a_i + a_j}$ by

$$C_{i,j}(x_1, \ldots, x_n) = \prod_{s=-a_i+1}^{a_j} (x_j - x_i + s)$$

$$\widetilde{f} = \prod_{i<j} C_{i,j}(x_1, \ldots, x_n)$$

By doing so, we indeed preserve the degree, but we are adding low degree terms.

These $C_{i,j}$ terms can be interpreted as being almost the indicator of whether the two segments $\Delta_i = [\alpha_i, \alpha_i + a_i - 1]$ and $\Delta_j = [\alpha_j, \alpha_j + a_j - 1]$ intersect or not. More precisely,

$$C_{i,j} = \begin{cases} 0, & \text{if the segments } \Delta_i \text{ and } \Delta_j \text{ intersect,} \\ & \text{or if } x_i = x_j + a_j \\ 1, & \text{otherwise} \end{cases}$$

So for $C_{i,j}$ to be nonzero, not only do the segments need to be disjoint, x_i (recall that $i < j$) should not lie right after Δ_j. With this last condition, we are restricting

$\widetilde{f}$ to have a unique nonzero point in $S_1 \times \ldots \times S_n$, the point $(\beta_1, \ldots, \beta_n)$, where $\beta_i = a_1 + \ldots + a_{i-1}$, with $\beta_1 = 0$. In other words, all $C_{i,j}$'s are nonzero if and only if all segments are disjoint and that they are in the correct order (i.e. $\Delta_1, \Delta_2, \ldots, \Delta_n$).

We still need to check that if the segments are not in the correct order, one of the $C_{i,j}$'s must be zero. Suppose that all intervals are disjoint and that for some $j > i$, Δ_i comes after Δ_j. WLOG we can assume that they are adjacent (in the sense that there are no intervals between the two). By the above observation, in order for $C_{i,j}$ to be nonzero, we need $x_i > x_j + a_j$. Thus the point $x_j + a_j$ is contained in no interval. But as all intervals live in $[0, a-1]$ and that they contain a elements in total, there must be an intersection. Contradiction.

We conclude that the expression for C given by the Nullstellensatsz is

$$C = \frac{\widetilde{f}(\beta_1, \ldots, \beta_n)}{\prod g_i'(\beta_i)}$$

where $g_i(x) = \prod_{\alpha \in S_i}(x - \alpha) = \prod_{s=0}^{a-a_i}(x - s)$.

From here on, all that is left is to explicitely compute C:

$$\begin{aligned}
C_{i,j}(\beta_1, \ldots, \beta_n) &= \prod_{s=-a_i+1}^{a_j} [(a_1 + \ldots + a_{j-1}) - (a_1 + \ldots + a_{i-1}) + s] \\
&= \prod_{s=-a_i+1}^{a_j} (a_i + \ldots + a_{j-1} + s) \\
&= \prod_{s=1}^{a_j+a_i} (a_{i+1} + \ldots + a_{j-1} + s) \\
&= \frac{(a_i + \ldots + a_j)!}{(a_{i+1} + \ldots + a_{j-1})!}
\end{aligned}$$

$$\begin{aligned}
g_i'(\beta_i) &= \prod_{s=0}^{a-a_i} (\beta_i - s) \\
&= (\prod_{s=1}^{\beta_i} s)(\prod_{k=1}^{a-a_i-\beta_i} (-k)) \\
&= (-1)^{a_{i+1}+\ldots+a_n} \beta_i!(a_{i+1} + \ldots + a_n)!
\end{aligned}$$

In $C_{i,j}$, the numerator and denominator cancel out, leaving us with in the numerator:

$$(a_1 + \ldots + a_j)!, \forall 1 < j < n \tag{3.2}$$

$$(a_i + \ldots + a_n)!, \forall n > i > 1 \tag{3.3}$$

$$a! \tag{3.4}$$

And single terms $a_i, \forall 1 \leq i \leq n$ in the denominator.

The terms 3.2 and 3.3 cancel out with $g_i'(\beta_i)$; the signs cancel out as well, giving us

$$C = \frac{a!}{a_1! \dots a_n!} \tag{3.5}$$

as desired. □

Chapter 4

The Cap Set Problem

4.1 Introduction

In this chapter, we go over a result by Ellenberg and Giswijt on three-term arithmetic progressions. They give an significantly better upper bound for the size of sets containing no three-term arithmetic progressions in $\mathbb{Z}/3\mathbb{Z}^n$, and by doing so they provide an answer to what is known as the cap set problem. The main reference in this chapter is [EG16].

An arithmetic progression is a sequence of elements $a_0, a_0 + d, a_0 + 2d, etc \ldots$ As the name indicates, a three-term arithmetic progression is a set of three elements $a_0, a_0 + d, a_0 + 2d$. A long lasting problem in additive combinatorics is to find large subsets of abelian groups with no three-term arithmetic progression. The classical groups that have been examined are $\mathbb{Z}/N\mathbb{Z}$ and $\mathbb{Z}/3\mathbb{Z}^n$.

Here we will see an upper bound on the size of such sets discovered independently by Ellenberg and Giswijt in 2016. They were able to make a significant improvement on this bound: the best previous attempt by Bateman and Katz [BK12] yielded a bound of $O(3^n/n^{1+\epsilon})$, while the bound here is of $O(2.756^n)$. Also note that the best known lower bound is of about 2.2^n, due to a construction of Edel [Ede04]. So although the bound we will see here solves the cap set problem by telling us that the size of the max cap set grows at least exponentially slower than the size of the group, we are still far from having matching upper and lower bounds.

4.2 The game of SET

A motivation to the version of the problem where the group is $\mathbb{Z}/3\mathbb{Z}^n$ (commonly called the cap set problem) comes from the game of SET [Aus16].

In the game of SET, we have a deck of a large number of cards (usually 81), where each card has one out of three values for several features. These features are traditionally things like the color, shape, shadowing and number of the objects. The goal of the game is, given a small subset of the deck, to find a set of 3 cards such that for each individual feature, the values of the cards are either the same or all different. This task can be quite hard at times, and sometimes it is in fact impossible. When this happens, the players repeatedly add 3 cards to the subset of revealed cards until a set can be found. A natural question that arises is: how big should the subset of the deck be in order to guarantee that there is a set? This is the question that the result of Ellenberg and Gisjwit answers.

In the traditional deck of SET, we can view each card as being a point in $\mathbb{Z}/3\mathbb{Z}^4$ by viewing each feature as a coordinate and labelling the values for a feature $0, 1, 2$. In these terms, the goal of the game is precisely to find a three arithmetic progression. It turns out that for this traditional deck of 81 cards, the maximum subset with no three arithmetic progression has size 20 (higlighted in yellow below). This was proven in 1971 by Pellegrino [Hil83].

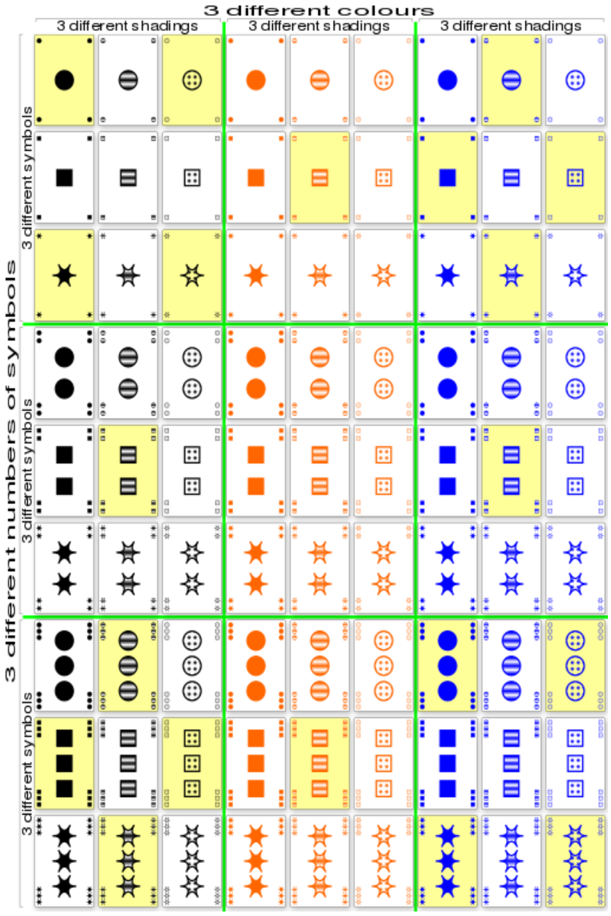

4.3 Notation

We begin by introducing the notation that will be used throughout this chapter.

$$\begin{aligned} M_n &: \text{monomials of n variables with degree in each variable at most } q-1 \\ M_n^d &: \text{subset of } M_n \text{ formed by monomials of total degree at most } d \\ S_n &: \mathbb{F}_q\text{-vector space spanned by } M_n \\ S_n^d &: \text{subspace of } S_n \text{ formed by polynomials of total degree at most } d \\ m_d &: dim\ S_n^d \text{ i.e. number of elements in } M_n^d \end{aligned}$$

We will also use the evaluation map

$$e\colon S_n \longrightarrow \mathbb{F}_q^{\mathbb{F}_q^n}, p \mapsto (p(a))_{a\in\mathbb{F}_q^n}$$

This map is a linear isomorphism and from it we can see that the indicator polynomials form a basis for S_n.

4.4 Main Result

Lemma 4.4.1. *Let $\mathbb{F}_q$ be a finite field, $A \subset \mathbb{F}_q^n$, $\alpha, \beta, \gamma \in \mathbb{F}_q$ summing to zero. If $f \in S_n^d$ is such that $f(\alpha a + \beta b) = 0$, $\forall a, b \in A$ distinct, then the number of $a \in A$ such that $f(-\gamma a) \neq 0$ is at most $2m_{d/2}$.*

Proof. To see this, we look at the square matrix B of size $|A|$ with entries $B_{ab} = f(\alpha a + \beta b)$ for $a, b \in A$. We then try to show that its rank is bounded by $2m_{d/2}$. Once we know this, combining this with the facts that B is diagonal (by hypothesis, $B_{ab} = f(\alpha a + \beta b) = 0$ when a, b distinct) and that $\alpha a + \beta a = -\gamma a$ (again by hypothesis) we conclude that the number of $a \in A$ such that $f(-\gamma a) \neq 0$ is at most $2m_{d/2}$.

The bound on the rank comes from a clever reordering of the terms of each entry of the matrix. We know that each entry is of the form

$$f(\alpha x + \beta y) = \sum_{m,m' \in M_n^d, deg(mm') \leq d} c_{m,m'} m(x) m'(y)$$

Now since the total degree is less than d, for each term $m(x)m'(y)$ either m or m' is of degree at most $d/2$. We can thus split the single sum above in two separate sums, each sum being indexed by monomials of degree at most $d/2$:

$$f(\alpha x + \beta y) = \sum_{m \in M_n^{d/2}} m(x) F_m(y) + \sum_{m' \in M_n^{d/2}} m'(y) G_m(x)$$

Now if we pick at a single monomial $m \in M_n^{d/2}$ and look at the $|A| \times |A|$ matrix with entries $m(a)F_m(b)$, we can see that it has rank 1. We conclude that B has rank at most $2m_{d/2}$ as it is the sum of $2m_{d/2}$ matrices of rank 1. □

We now look at the main theorem for this section, from which the solution to the cap set problem is derived. The theorem is stated for a field of size q, but in its application we really are interested only in the case where q is 3.

Theorem 4.4.2. *Let α, β, γ be elements of $\mathbb{F}_q$ such that they sum to zero and $\gamma \neq 0$. Let $A \subset \mathbb{F}_q^n$ such that*

$$\alpha a_1 + \beta a_2 + \gamma a_3 = 0$$

has no solutions $(a_1, a_2, a_3) \in A^3$ except when $a_1 = a_2 = a_3$.

Then $|A| \leq 3m_{(q-1)n/3}$.

The statement above seems a bit more general than our problem. We are interested in sets that contain no three-term arithmetic progressions. In other words, to use the same notation as in the above statement, we are interested in A for $\alpha = 1, \beta = 1, \gamma = 1$ (in $(\mathbb{Z}/3\mathbb{Z})^n$).

Indeed, having a three-term arithmetic progression a_1, a_2, a_3 means by definition that there exists a_0, $d \in (\mathbb{Z}/3\mathbb{Z})^n$ such that

$$\begin{aligned} a_1 &= a_0 \\ a_2 &= a_0 + d \\ a_3 &= a_0 + 2d \end{aligned}$$

So with the above settings of α, β, γ, we get $\alpha a_1 + \beta a_2 + \gamma a_3 = 0$ in $(\mathbb{Z}/3\mathbb{Z})^n$.

From this we see that if a set A is as in the statement of the theorem then there is no three-term arithmetic approximation. And if some set has no three-term approximation, it cannot have any solution, other than $a_1 = a_2 = a_3$.

Indeed, if $a_1 + a_2 + a_3 = 0$ with WLOG a_1 non trivial, then the sum of each coordinate is either 0 or divisible by 3. From this we can easily see they form an arithmetic progression.

For each coordinate i, if they all share a common value, set a_0 to be that value in coordinate i and d to be 0. Otherwise, we must have the values $0, 1, 2$ in some order. One can easily work out what the assignments of the ith coordinate of a_0 and d should be depending on the order:

$$\begin{aligned} \{0,1,2\} &\rightarrow a_0 = 0, d = 1 \\ \{0,2,1\} &\rightarrow a_0 = 0, d = 2 \\ \{1,2,0\} &\rightarrow a_0 = 1, d = 1 \\ \{1,0,2\} &\rightarrow a_0 = 1, d = 2 \\ \{2,0,1\} &\rightarrow a_0 = 2, d = 1 \\ \{2,1,0\} &\rightarrow a_0 = 2, d = 2 \end{aligned}$$

All this to say that the bound given by the theorem is also a bound for our sets.

Proof. Let $d \in [0, (q-1)n]$ be an integer. By V denote the space of polynomials in S_n^d vanishing on the complement of $-\gamma A$ (for simplicity, write $(-\gamma A)^c = X$). It has dimension at least $m_d - q^n + |A|$.

Indeed, if we denote by W the space of functions $X \to \mathbb{F}_q$, we have $dimW = |X|$. And as the restriction ϕ of polynomials over $\mathbb{F}_q^n$ to X is a linear homomorphism, we get that

$$dimKer(\phi) + dimImg(\phi) = dimS_n^d$$

Of course, the kernel of the map ϕ is the subspace V, and its image is contained in W, thus $|X| \geq dimImg(\phi)$. Hence the inequality

$$\begin{aligned} dimV + |X| &\geq m_d \\ dimV &\geq m_d - q^n + |A| \end{aligned}$$

Let $S(A) = \{g \in \mathbb{F}^n \colon g = \alpha a_1 + \beta a_2, a_1 \neq a_2 \in A\}$. Notice that this set is precisely all the elements on which the polynomial was vanishing in the previous lemma. And in fact we are in a position where we can apply the lemma. Indeed, by hypothesis, for no distinct $a_1, a_2 \in A$ is there $a_3 \in A$ such that $\alpha a_1 + \beta a_2 = -\gamma a_3$. Thus $S(A)$ and $-\gamma A$ do not intersect, i.e. $S(A) \subset (-\gamma A)^c$, so any polynomial $f \in V$ vanishes on $S(A)$. We conclude that if Σ is the support of such an f, $|\Sigma| \leq 2m_{d/2}$ (note that the entire support is contained in $-\gamma A$).

Now, if we pick $f \in V$ such that f has maximal support, then we also get the bound $|\Sigma| \geq dimV$.

Indeed, suppose for contradiction that $|\Sigma| < dimV$. Then if we consider the evaluation map $e \colon V \longrightarrow \mathbb{F}^{|\Sigma|}$, it necessarily has non trivial kernel. We can then take $g \in V$ nonzero such that g vanishes on Σ. But looking at $f + g$, we see that it is nonzero on Σ, and that as g is nonzero, there is some point $s \notin \Sigma$ at which it is nonzero, and so $f + g$ is also nonzero at s. We have thus constructed an element of V with strictly larger support than f which contradicts the choice of f.

Now putting all the inequalities together we get

$$\begin{aligned} dimV &\leq |\Sigma| \leq 2m_{d/2} \\ m_d - q^n + |A| &\leq 2m_{d/2} \\ |A| &\leq 2m_{d/2} + (q^n - m_d) \end{aligned}$$

Let $d = 2(q-1)n/3$. Then $|A| \leq 2m_{(q-1)n/3} + (q^n - m_{2(q-1)n/3})$.

Recall $|M_n| = q^n$ and $m_d = |M_n^d|$. So $q^n - m_d$ corresponds to the number of monomials of M_n such that they are of degree strictly larger than d. But as there is a bijection between these and monomials of degree stricly less than $(q-1)n - d$, that number is at most $m_{(q-1)n-d}$. Note that the bijection is simply the complement with respect to the degrees: $x_1^{d_1} \dots x_n^{d_n} \mapsto x_1^{(q-1)n-d_1} \dots x_n^{(q-1)n-d_n}$.

We thus have

$$\begin{aligned} |A| &\leq 2m_{(q-1)n/3} + m_{(q-1)n/3} \\ &= 3m_{(q-1)n/3} \end{aligned}$$

□

What is left to do is to bound the number $m_(q-1)n/3$. Ellenberg and Giswijt gave a probabilistic interpretation to this number in order to bound it.

If we take $X_1, \dots, X_n$ to be i.i.d. discrete random variable taking values $\{0, 1, \dots, q-1\}$ with uniform probability, we can see that

$$\mathbb{P}[\sum \frac{X_i}{n} \leq (q-1)/3] = m_{(q-1)n/3}/q^n$$

This can then be seen as a large deviation problem and solved using Cramer's theorem.

But there is another way to obtain the bound that was proposed by Tao [Tao16] that does not require these advanced probability tools.

Let $q = 3$. We can count the number of monomials of total degree at most d and degree in each variable at most 2 in the following way. With the following notation

$$\begin{aligned} a&\colon \text{number of variable of degree 0} \\ b&\colon \text{number of variable of degree 1} \\ c&\colon \text{number of variable of degree 2} \end{aligned}$$

we can view m_d as the sum over all a, b, c such that $a + b + c = n$ and $b + 2c \leq d$ of the number of possible choices of a, b, c variables. Thus,

$$m_d = \sum_{a+b+c=n;\ a,b,c\geq 0;\ b+2c\leq d} \frac{n!}{a!b!c!}$$

We can rewrite $a + b + c = n$ as $\alpha n + \beta n + \gamma n = n$ for some α, β, γ summing to 1. From here we can use Stirling's formula:

$$n! = (1 + o(1))\sqrt{2\pi n} \cdot n^n e^{-n}$$

Which essentially tells us that $n! = n^n e^{-n}$ up to a polynomial factor.

We get

$$\begin{aligned} \frac{n!}{a!b!c!} &\sim \frac{n^n e^{-n}}{(\alpha n)^{\alpha n}(\beta n)^{\beta n}(\gamma n)^{\gamma n} e^{-(\alpha+\beta+\gamma)n}} \\ &= \frac{1}{(\alpha^\alpha \beta^\beta \gamma^\gamma)^n} \\ &= exp(n \cdot h(\alpha, \beta, \gamma)) \end{aligned}$$

Where $h(\alpha, \beta, \gamma) = \alpha log\frac{1}{\alpha} + \beta log\frac{1}{\beta} + \gamma log\frac{1}{\gamma}$, is commonly called the entropy function.

The number of triples (a, b, c) in the expression of m_d is upperbounded by the number of triples (a, b, c) of non negative numbers summing to n. But by the stars and bars counting technique, this number is $\binom{n+3}{3}$ which is $O(n^3)$. We can thus make the observation that m_d is a sum of a polynomial number of terms, and hence can write that, up to a polynomial factor, we have

$$m_d = exp(n \cdot \max h(\alpha, \beta, \gamma))$$

where the maximum is taken over all α, β, γ summing to 1 such that $\alpha, \beta, \gamma \geq 0$ and $\beta + 2\gamma \leq d/n$. Setting $d = (3-1)n/3$ as in the bound given by the theorem, this is a maximization problem with constraints as above, and the last constraint becomes $\beta + 2\gamma \leq 2/3$.

Using the method of Lagrange multipliers, we obtain

$$\begin{aligned} \alpha &= \frac{32}{3(15 + \sqrt{32})} \\ \beta &= \frac{4(\sqrt{33} - 1)}{3(15 + \sqrt{32})} \\ \gamma &= \frac{(\sqrt{33} - 1)^2}{6(15 + \sqrt{33})} \end{aligned}$$

This yields $h(\alpha, \beta, \gamma) = 1.013455$ and thus the bound of $O(2.756^n)$.

Bibliography

[Alo99] Noga Alon. Combinatorial nullstellensatz. *Combinatorics, Probability and Computing*, 8(1-2):7–29, 1999.

[Aus16] David Austin. Game. set. polynomial, 2016.

[Bes28] AS Besicovitch. On kakeya's problem and a similar one. *Mathematische Zeitschrift*, 27(1):312–320, 1928.

[BK12] Michael Bateman and Nets Katz. New bounds on cap sets. *Journal of the American Mathematical Society*, 25(2):585–613, 2012.

[DKSS09] Zeev Dvir, Swastik Kopparty, Shubhangi Saraf, and Madhu Sudan. Extensions to the method of multiplicities, with applications to kakeya sets and mergers. In *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on*, pages 181–190. IEEE, 2009.

[Dvi09] Zeev Dvir. On the size of kakeya sets in finite fields. *Journal of the American Mathematical Society*, 22(4):1093–1097, 2009.

[Ede04] Yves Edel. Extensions of generalized product caps. *Designs, Codes and Cryptography*, 31(1):5–14, 2004.

[EG16] Jordan S Ellenberg and Dion Gijswijt. On large subsets of $\mathbb{F}_q^n$ with no three-term arithmetic progression. *arXiv preprint arXiv:1605.09223*, 2016.

[Fur08] Markus Furtner. The kakeya problem, 2008.

[Hil83] R. Hill. On Pellegrino's 20-Caps in S4, 3. In A. Barlotti, P.V. Ceccherini, and G. Tallini, editors, *Combinatorics '81 in honour of Beniamino Segre*, volume 78 of *North-Holland Mathematics Studies*, pages 433–447. North-Holland, 1983.

[KP12] RN Karasev and FV Petrov. Partitions of nonzero elements of a finite field into pairs. *Israel Journal of Mathematics*, 192(1):143–156, 2012.

[SS$^+$08] Shubhangi Saraf, Madhu Sudan, et al. An improved lower bound on the size of kakeya sets over finite fields. *Analysis & PDE*, 1(3):375–379, 2008.

[Tao13] Terence Tao. Algebraic combinatorial geometry: the polynomial method in arithmetic combinatorics, incidence combinatorics, and number theory. *arXiv preprint arXiv:1310.6482*, 2013.

[Tao16] Terence Tao. A symmetric formulation of the croot-lev-pach-ellenberg-gijswijt capset bound, 2016.