What is the Polynomial Method?

Niranjan Balachandran*

Abstract

The polynomial method is a relatively new algebraic tool (and philosophy) that has over the past ten years enabled researchers to settle several long-standing open problems arising from diverse areas such as Combinatorial and Finite Geometry, Additive Combinatorics, Number theory, and so on. I shall in this note, provide an introduction to what this method is all about, and as an attempt to expound on this new technique, go over four problems in which substantial progress happened from a prior state of virtual hopelessness. In particular, we shall consider the following:

1. Dvir's solution of the Finite Kakeya Conjecture,
2. Guth-Katz' solution of the Joints' Problem,
3. The Cap-set problem and the work of Ellenberg-Gisjwijt, and
4. A function field analogue of Sárközy's theorem, due to Green.

1 Introduction: 4 problems

The Polynomial method is a relatively new algebraic technique (and philosophy!) that has been used to great effect to settle several long-standing open problems. often with a strong combinatorial flavour, arising from several areas - Combinatorial and Finite Geometry, Additive Combinatorics, Number theory, to name a few. The basic premise of the method is the following. Given a combinatorial configuration/set (usually with strong rigidity properties - we will return to what this means, later) in a vector space, one associates a polynomial of 'somewhat' low degree, such that the polynomial vanishes on the set in question. And since polynomials typically demonstrate very strong robustness properties, the bearing of these properties on the set is studied, in order to shed further light on the original combinatorial problem. This description sounds very vague, and too generic to be of any specific use, so in this note, we shall expound on this method by deconstructing the proofs of four recent results. These results describe what I believe form the basic setting for the Polynomial method. For problems and solutions that involve more sophisticated techniques, the reader is referred to the recent book by Guth [12] or the survey paper by Tao [18] for instance.

To get started, let us first take a look at the following problems.

*Department of Mathematics, Indian Institute of Technology Bombay, Mumbai, India. email: niranj (at) math.iitb.ac.in.

1. **The Finite Kakeya Problem:** Let q be a prime power. A set $K \subset \mathbb{F}_q^n$ is called a *Kakeya Set* if for every direction v, there is some line ℓ_v along the direction v, that is contained in K, i.e., $\ell_v \subset K$. If K is a Kakeya set in $\mathbb{F}_q^n$, then is it true that $|K| \geq c_n q^n$ for some constant $c_n > 0$? (that does NOT depend on q).

   The Finite Kakeya problem was proposed by T. Wolff as a 'toy' version of the more general Kakeya Conjecture about the Hausdorff dimension of a Kakeya set in $\mathbb{R}^n$, which is still unsettled. For a while, it appeared that the Finite Kakeya problem was as hard as the 'real' version, till the finite version was settled in the affirmative by Dvir in 2008 ([6]).

2. **The Joints' problem:** This problem in Combinatorial Geometry was proposed in a paper by Chazelle *et al* ([4]) in the early 90s. Suppose $\mathcal{L}$ is a collection of L lines in $\mathbb{R}^3$. By a *Joint*, we shall mean a point $v \in \mathbb{R}^3$ such that three non-coplanar lines of $\mathcal{L}$ pass through v. The Joints' problem asks: Given a collection $\mathcal{L}$ with $|\mathcal{L}| = L$, what is the maximum possible number of joints they determine?

   It is clear that the number of joints determined by $\mathcal{L}$, which we shall denote $j(\mathcal{L})$, satisfies $j(\mathcal{L}) \leq \binom{L}{2}$ since a joint is necessarily an intersection point. If one tries to construct a collection of lines with 'many' joints, then the following 'grid' example suggests itself; if we designate that all the points of the $N \times N \times N$ grid turn out to be joints, then the collection of lines $\mathcal{L}$ is simply all the axis-parallel lines that intersect this grid. In particular, this sets $L = 6N^2$ and $j(\mathcal{L}) = N^3$, so this achieves $j(\mathcal{L}) \geq \Omega(L^{3/2})$. The proposers of the problem proved that $j(\mathcal{L}) \lesssim L^{7/4}$ ([4]). The correct order of the number of joints remained in question till it was settled by Guth and Katz in 2010 ($j(\mathcal{L}) \lesssim L^{3/2}$) in [13].

3. **The Cap-set problem:** This problem is motivated by an old problem in additive combinatorics. An old conjecture of Erdős-Turán (which was settled in its fullness by Szemerédi) states the following: Given and integer $k \geq 3$, and $\varepsilon > 0$ there exists $n_0(k, \varepsilon)$ such that the following holds. For any $N \geq n_0$, and any subset $A \subset [N]$ with $|A| \geq \varepsilon N$ A contains a k term AP. For the case $k = 3$, the current best known result states the following: For $N \gg 0$ and any $A \subset [N]$ with $|A| \geq C\frac{N(\log\log N)^4}{\log N}$ has a 3-term AP ([15]). On the other hand, a classic (and beautiful) construction of Behrend ([3]) shows that there exist sets A with $|A| \geq Ne^{-C\sqrt{\log N}}$ with no 3-AP. In particular, for any $\varepsilon > 0$ the size of a maximal subset of $[N]$ without a 3-AP is larger than $N^{1-\varepsilon}$. In order to see if one could improve upon Behrend's construction significantly (there have been minor improvements, the latest due to Elkin, [9]) the corresponding 'toy' version frames the same question over finite fields. Since the smallest field in which a 3-AP makes sense is $\mathbb{F}_3$, the following question - the Cap-Set problem[1] the first interesting case: What is the size of a largest set $A \subset \mathbb{F}_3^n$ without any 3-term APs? The first interesting result in the regard was due to Meshulam who shows that if $|A| \geq \Omega(3^n/n)$ then A contains a 3-AP. This was improved a little by Bateman and Katz ($3^n/n^{1+\varepsilon}$). The best known lower bound is about 2.2^n (Edel, [8]). It came as a bit of a shock then, when Ellenberg, and independently, Gisjwijt, following a recent brilliant idea in a paper of Croot-Lev-Pach (see [5]) showed an upper bound of $o(2.756^n)$ ([10]). Thus, the toy version of the problem does not shed any light on the original problem.

4. **Function Field analogue of Sárközy's theorem:** A. Sárközy proved the following theorem ([16]) in 1978: Given a polynomial $f \in \mathbb{Z}[T]$ with $f(0) = 0$, there exists a constant $c_f > 0$

[1]The terminology borrows from the card game '**SET**' which is played with a set of 81 cards, each bearing a picture with 4 attributes, each of which takes one of 3 different values. A win or a 'Set' is a set of 3 cards such that for each of the 4 attributes, either the set witnesses all possible values, or is a constant across the cards. Modelling this over $\mathbb{F}_3^4$, a set is precisely a 3-AP.

such that the following holds: For $N \gg 0$ and any set $A \subset [N]$ with $|A| \geq N(\log N)^{-c_f}$ there exist $a \neq b$ with $a, b \in A$ such that $a - b = f(u)$ for some $u \in \mathbb{Z}$. Green proved ([11]) in 2016, the following function field analogue of Sárközy's theorem, over finite fields: Given a polynomial $F \in \mathbb{F}_q[T]$ of degree k with $F(0) = 0$, there exists $0 < c := c(k, q) < 1$ such that for any subset A of polynomials (in T) of degree less than n with $|A| > q^{(1-c)n}$ there exist $\alpha(T) \neq \beta(T)$ in A such that $\alpha(T) - \beta(T) = F(\gamma(T))$ for some $\gamma(T) \in \mathbb{F}_q[T]$.

This theorem proves a stronger version (somewhat similar to the Cap-set problem case) of Sárközy's theorem than the integer case.

In the next section, I shall expound on the basic principle of the Polynomial method, and then deconstruct each of these proofs in the following two sections. But before we embark on that task, I make one (in my opinion) important point and distinction, which I shall explain in a little greater detail in the next section. The Polynomial method is very reminiscent of the Linear Algebra method (see [2] or [14]), or the technique known as the Combinatorial Nullstellensatz (CN) of Alon ([1]). Both these techniques also work on the premise of encoding the (usually combinatorial) problem in terms of polynomials. While some researchers include the CN within the ambit of the Polynomial method, most people do think of the Linear Algebra method as distinct from the Polynomial method. I maintain that *neither of these very interesting techniques can be regarded as expressing the basic philosophy of the Polynomial method.* I will explain my stance in more detail in the next section.

2 The Premise of the Polynomial Method

The basic premise of the Polynomial method goes as follows. To attempt to solve a combinatorial problem about the cardinality of a set with some rigidity properties, one picks a polynomial f (often in several variables), of 'low degree', such that the set in question is contained in the set of zeroes of f (which we shall denote by $Z(f)$). As mentioned in the last paragraph of the preceding section, this bears a great deal of resemblance with the Linear algebra method, or the technique of CN. However, what distinguishes the Polynomial method (both technically, as well as philosophically) are the following:

- In either the Linear Algebra method or the usage of CN, the *power of the technique depends very critically on how cleverly/effectively/robustly the polynomials are chosen.* In other words, the method rests greatly on knowing the polynomial(s) *explicitly.* The Polynomial method, on the other hand, derives its power in the *relative anonymity of the polynomial is what makes the method work* (when it works). Not being able to explicitly write down such a polynomial is underscored by the stronger fact that the polynomial's degree is under greater control here.
- In the use of the Linear Algebra method, there are many instances, where the corresponding vector spaces are vector spaces of polynomials over a field, and the polynomials are only treated as vectors in the corresponding vector spaces. But the polynomial method relies on two other aspects:
  - Polynomials have some very strong rigidity properties. For example, a polynomial (of one variable) of degree k has at most k zeroes (over a field). In other words, knowing the value of a polynomial of degree less than k at k distinct points determines the polynomial

uniquely. Thus getting polynomials of 'low degree' are crucial because they demonstrate a greater sense of rigidity.

- Each $f \in \mathbb{F}[X_1, X_2, \ldots, X_n]$ also naturally defines a function $\phi_f : \mathbb{F}^n \to \mathbb{F}$. In fact if $\mathbb{F} = \mathbb{F}_q$ then *every function* $\phi : \mathbb{F}_q^n \to \mathbb{F}_q$ is in fact realizable as a polynomial function.

The power of the polynomial method relies very crucially on how well one can combine/juggle this duality between polynomials as vectors and polynomials as robust functions.

To emphasize the first point made, we start with a simple question. Suppose we wish to find a polynomial $f(X, Y) \in \mathbb{R}[X, Y]$ such that the points $(i, 2^i)$ for $i = 1, \ldots, 2018$ lie in $Z(f)$. While explicit choices such as $f(X, Y) = \prod_{i=1}^{2018}(X - i)$ do accomplish this, one can get the same with polynomials of much smaller degree.

Lemma 1. *Suppose $S \subset \mathbb{F}^n$ is a finite set and suppose $|S| < \binom{n+d}{n}$. Then there exists a non-trivial polynomial $f(X_1, \ldots, X_n) \in \mathbb{F}_q[X_1, \ldots, X_n]$ of degree at most d such that $S \subset Z(f)$. In particular, given a finite set S, there exists a polynomial f of degree at most $n|S|^{1/n}$ such that $S \subset Z(f)$*

Proof. The vector space of polynomials in n variables, of degree at most d has dimension $\binom{n+d}{n}$ (Exercise!). To get a polynomial f in n variables that vanishes at each point of S, this imposes $|S|$ linear equations in $\binom{n+d}{n}$ variables. Under the hypothesis, it is a simple linear algebra result that there exists a non-trivial solution to this system of linear equations, which translates to the existence of a non-trivial polynomial f. The latter statement is a simple computation that is an instructive exercise. □

In particular, to answer the question posed above, one can in fact find a polynomial $f(X, Y)$ of degree at most 89 which is a much smaller than the 'nice', explicit polynomials we mentioned earlier.

The robustness of polynomials is described in the following lemma.

Lemma 2. *Suppose $f \in \mathbb{F}[X_1, \ldots, X_n]$ is a polynomial of degree at most d. Then for any line ℓ, either $\ell \subset Z(f)$ or $|\ell \cap Z(f)| \leq d$.*

Proof. Let ℓ be parametrized as $\ell(t) = \mathbf{a}+\mathbf{v}t$ as $t \in \mathbb{F}$. Restrict f to the line; in other words, consider the polynomial $f_\ell(T) = f(\mathbf{a}+\mathbf{v}T) = f(a_1+v_1T, \ldots, a_n+v_nT)$. Since this is a univariate polynomial of degree at most d, the result follows from the basic properties of univariate polynomials. □

I shall now deconstruct the proofs of the first and second problems in the next section, and the proofs of problems 3 and 4 in the final section. I have grouped them this way keeping with the thematic similarities in these proofs.

3 Problems 1 & 2

3.1 Dvir's solution to the Finite Kakeya Problem

Suppose K is a Kakeya set in $\mathbb{F}_q^n$. We wish to prove that K is 'large'. Now, if not, i.e., suppose $|K| \leq cq^n$ for some c. Then, by lemma 1 there is a non-trivial polynomial $f \in \mathbb{F}_q[X_1, \ldots, X_n]$ of degree at most $nc^{1/n}q$ such that $K \subset Z(f)$. Now if $nc^{1/n} \leq 1/2$ (say), then in particular we have $deg(f) \leq q/2$.

The main point here is this: Since $deg(f) \leq q/2$ by the observation above, several lines of $\mathbb{F}_q^n$ are contained in $Z(f)$. To make this more precise, write $f = f_H + f_R$, where f_H denotes the homogeneous part of the highest degree terms of f. Fix $\mathbf{v}$, and the corresponding line $\ell(\mathbf{v}) := \mathbf{a} + \mathbf{v}t$ which is contained in K by definition. Then, we have $f(a_1 + v_1 t, \ldots, a_n + v_n t) = 0$ for all $t \in \mathbb{F}_q$. Since this is a polynomial of degree at most $q/2$ which vanishes at q elements, by lemma 2, this is identically the zero polynomial, so in particular, the coefficient of t^k, viz., $f_H(\mathbf{v})$ (check!) equals 0. Since v is arbitrary, it follows that f_H vanishes at all $\mathbf{v}$, which implies that $f_H \equiv 0$. But this contradicts that this is the leading term of f; this completes the proof.

This proof in particular, shows that any Kakeya set in $\mathbb{F}_q^n$ has size at least $q^n/(2n)^n$. But one can by the usage of a slightly more sophisticated technique improve this bound to $q^n/2^n$ (see [7]). This is almost best possible, as shown by Saraf and Sudan (see [17]) paper gives an example of a Kakeya set of size $q^n/2^{n-1} + O(q^{n-1})$.

3.2 The Joints' Problem

Again, suppose $\mathcal{L}$ is a collection of L lines in $\mathbb{R}^3$ that determine the maximum number of joints. Let us try an induction proof. Denote the maximum number of joints determined by L lines by $j(L)$. Suppose there is a line $\ell \in \mathcal{L}$ such that the number of joints lying on ℓ is somewhat 'small', say $f(L)$. Then we have

$$j(L) \leq f(L) + j(L-1).$$

Now, what determines 'small' in this context? Let us now turn on to our new-found philosophy. Our main object of interest here is the set of joints $J = J(\mathcal{L})$, so let F be a minimum degree polynomial such that $J \subset Z(F)$. In particular, $deg(F) \leq 3|J|^{1/3}$, by lemma 1. Now, with the thinking tool that lemma 2 offers, note that if the line ℓ has strictly more than $3|J|^{1/3}$ joints, then ℓ is contained in $Z(F)$. Here's the dichotomy: For every line $\ell \in \mathcal{L}$, either

- $|\ell \cap Z(F)| \leq 3j(L)^{1/3}$ **or**
- $\ell \subset Z(F)$.

So, suppose now, that some line contains at most $3j(L)^{1/3}$ joints. This gives us,

$$\begin{aligned} j(L) &\leq 3j(1)^{1/3} + \cdots + 3j(L)^{1/3} \\ &\leq 3Lj(L)^{1/3} \end{aligned}$$

and that gives

$$j(L)^{2/3} \leq 3L \;\Rightarrow\; j(L) \leq 3^{3/2}L^{3/2}.$$

Now, let us take stock of the situation. We have seen that if there is a line with at most $3j(L)^{1/3}$ joints on it, and *this happens in every configuration of lines*, then we are home, by induction. So, suppose there is no such line. In other words, *every line* of $\mathcal{L}$ has more than $3|J|^{1/3}$ joints on it; by the dichotomy observation of the preceding paragraph, it follows that every line of $\mathcal{L}$ is contained in $Z(F)$.In particular, F, when restricted along the lines of $\mathcal{L}$, is identically zero, so the directional derivative of F along any line of $\mathcal{L}$ is also zero. Hence, for each $\ell \in \mathcal{L}$ we have $\langle \nabla F, \mathbf{v}_\ell \rangle = 0$, where $\mathbf{v}_\ell$ denotes a unit vector along the line ℓ.

Now, since a joint has three non-coplanar lines through it, it follows that for *each joint* p,

$$\nabla F(p) = \mathbf{0}.$$

Consequently,

$$\frac{\partial F}{\partial x}(\mathbf{p}) = \frac{\partial F}{\partial y}(\mathbf{p}) = \frac{\partial F}{\partial z}(\mathbf{p}) = 0.$$

But then since F is a non-trivial polynomial of minimum degree with $F(\mathbf{p}) = 0$, each of the partial derivatives above must be zero, but that forces that $F \equiv 0$, and that is a contradiction!

4 Problems 3 & 4

4.1 The Cap-set problem

So, let's take a look at the cap-set problem. While Ellenberg and Gisjwijt's proof actually works over all $\mathbb{F}_p$, we shall restrict our attention to the case $p = 3$ since this already demonstrates all main ideas of the proof. One of the advantages of this case is that a 3 term AP a, b, c is simply one with the property $a + b + c = \mathbf{0}$.

We seek an upper bound for the size of a cap-set A, i.e., a set A with no 3-APs, so we look to prove a statement of the form,

$$A \text{ is not very large.}$$

If this is not the case, then one equivalently makes the statement

$$\overline{A} \subset \mathbb{F}_3^n \text{ is large.}$$

So, our newfound philosophy suggests: Start with $f \in \mathbb{F}_3[X_1, \ldots, X_n]$, a 'low-degree' polynomial such that f vanishes on $\overline{A}$.

Now, as mentioned earlier, the payoff of the method depend on how well we utilize their identity as functions. As described earlier, let $\mathcal{W}_n$ denote the vector space of polynomials in $X_1, \ldots, X_n$ over $\mathbb{F}_3$ generated by monomials $X_1^{a_1} \cdots X_n^{a_n}$ with $0 \leq a_i \leq 2$ for each i, and let $\mathcal{W}_{n,d}$ denote

the subspace of $\mathcal{W}_n$ comprising of polynomials of degree at most d. In keeping with the general principle, we wish to choose a polynomial in $\mathcal{W}_{n,d}$ for a suitable d that vanishes on $\overline{A}$. But how do we know there is a non-zero f here that we may pick, and what would be an optimal d?

Let $\mathcal{F}$ denote the vector space of all functions $f : \mathbb{F}_3^n \to \mathbb{F}_3$ that vanish at every point of $\overline{A}$, then we seek a non-zero $f \in \mathcal{W}_{n,d} \cap \mathcal{F}$, for an optimal value of d. Writing $m_{n,d} = \dim(\mathcal{W}_{n,d})$ we have

$$\begin{aligned}
\dim(\mathcal{W}_{n,d} \cap \mathcal{F}) &\geq m_{n,d} - |\overline{A}| \\
&= m_{n,d} - 3^n + |A| \\
&= 3^n - m_{n,2n-d-1} - 3^n + |A| \text{ (since } m_{n,d} = 3^n - m_{n,2n-d-1}. \text{ Why?!)} \\
&= |A| - m_{n,2n-d-1}.
\end{aligned}$$

Now, if the RHS above is non-positive then we have $|A| \leq m_{n,2n-d-1}$. So, let us assume that the RHS> 0. Since functions that are maximally non-zero on A are likely to be more useful, let g be an element of $\mathcal{W}_{n,d} \cap \mathcal{F}$ of maximum support $S \subset A$. I claim that $|S| \geq \dim(\mathcal{W}_{n,d} \cap \mathcal{F})$. Indeed, suppose not. Then the subspace of $\mathcal{W}_{n,d} \cap \mathcal{F}$ consisting of functions that disappear on S has dimension at least one. Pick a non-zero h from this subspace, and now consider the element $g + h \in \mathcal{W}_{n,d} \cap \mathcal{F}$; this has strictly larger support than g contradicting the assumption about g.

Now, (and this is a natural combinatorial idea that arises naturally in many applications of the Linear Algebra method) consider the $|A| \times |A|$ matrix $\mathcal{M}$ with rows and columns indexed by the elements of A whose $(a, b)^{th}$ element is $f(-a - b)$.

This is natural because, since if $-a - b = c \in A$ for distinct $a \neq b$, then we have $a + b + c = \mathbf{0}$ which is not possible by definition. By the choice of f, it follows that $\mathcal{M}$ is a diagonal matrix, and so $r(\mathcal{M}) \geq |S| \geq |A| - m_{n,2n-d-1}$, where S is the set described in the preceding paragraph.

The last part of this proof follows a trick of Croot-Lev-Pach (see [5]). Denote $\mathbf{X} := (X_1, \ldots, X_n)$ and $\mathbf{Y} := (Y_1, \ldots, Y_n)$ and consider the polynomial $F(\mathbf{X}, \mathbf{Y}) := f(\mathbf{X} + \mathbf{Y})$, so that the $(a, b)^{th}$ entry of $\mathcal{M}$ is $F(a, b)$. Since f has degree at most d, so does F. Then we can write

$$F(\mathbf{X}, \mathbf{Y}) = \sum_{\mathfrak{m} \in \mathcal{W}_{n,d/2}} \mathfrak{m}(\mathbf{X}) M_{\mathfrak{m}}(\mathbf{Y}) + \sum_{\mathfrak{m} \in \mathcal{W}_{n,d/2}} \mathfrak{m}(\mathbf{Y}) M_{\mathfrak{m}}(\mathbf{X}).$$

But now, (and this is the Croot-Lev-Pach trick) observe that this identity allows us to realize the matrix $\mathcal{M}$ as the sum of at most $2m_{n,d/2}$ rank one matrices, (since each summand of either sum describes a rank one matrix) so $r(\mathcal{M}) \leq 2m_{n,d/2}$ (!).

This coupled with the previous observations gives us

$$|A| \leq m_{n,2n-d-1} + 2m_{n,d/2}.$$

To complete the proof now, we look to optimize over d. First, setting (ignoring divisibility issues) $2n - d - 1 = d/2$ gives $d = 2(2n - 1)/3$, so we have $|A| \leq 3m_{n,(2n-1)/3}$.

Now $m_{n,d}$ is precisely the number of n-tuples $(a_1, \ldots, a_n)$ such that $0 \leq a_i \leq 2$ and $\sum a_i \leq d$. This does not admit an explicit simple analytic form answer, so we turn to estimates.

Suppose a_i $(1 \leq i \leq n)$ are independent random variables taking values in $\{0, 1, 2\}$ uniformly; the well-known Hoeffding inequality for sums of iid random variables gives us tight concentration inequalities:

$$\mathbb{P}(\big|\sum_i a_i - n\big| > t) \leq 2\exp\left(\frac{-t^2}{8n}\right),$$

so taking $t = n/3$ gives in particular, that the number of $(a_1, \ldots, a_n)$ that count towards $m_{n,d}$ for $d = (2n-1)/3$ is at most $3^n \exp(-n/72) = (2.971\cdots)^n$.

To get the better bound of $o(2.756^n)$ (á la Ellenberg-Gisjwijt) one needs sharper probabilistic estimates. In fact they use a sharper result of Crámer on large deviations to get their result. I refer the interested reader to their paper [10] for those details. As mentioned in the introduction, the best known lower bound gives a set of size about 2.2^n without a 3-AP in $\mathbb{F}_3^n$, so there is still ample room to obtain the best possible base for this exponential order.

4.2 Sárközy's theorem for function fields

This proof (due to Green) again is very similar to the Ellenberg-Gisjwijt proof, so I shall mainly demonstrate the case $F(T) = T^k$; the general case involves a slightly worse value for c, and I shall skip that part here.

We wish to show a suitable $c = c(k, q)$ such that if a set $A \subset (\mathbb{F}_q[T])_{<n}$ of polynomials of degree less than n has size at least $q^{(1-c)n}$ then there exist $\alpha(T) \neq \beta(T)$ in A such that $\alpha(T) - \beta(T) = (\gamma(T))^k$ for some $\gamma(T) \in \mathbb{F}_q[T]$.

Let us frame this another way. Write

$$(a_0 + a_1T + \cdots + a_{m-1}T^{m-1})^k = g_0(a_0, \ldots, a_{m-1}) + g_1(a_0, \ldots, a_{m-1})T + \cdots + g_{(m-1)k}(a_0, \ldots, a_{m-1})T^{(m-1)k}$$

for polynomial functions $g_0, \ldots, g_{mk}$. Ignoring floor and ceiling signs, if we set $m = (n-1)/k + 1$, so that $k(m-1) = n-1$, this defines the map $\Phi : \mathbb{F}_q^m \to \mathbb{F}_q^n$ defined as

$$\Phi(a_0, \ldots, a_{m-1}) = (g_0(a_0, \ldots, a_{m-1}), \ldots, g_{n-1}(a_0, \ldots, a_{m-1}))$$

which essentially encodes taking the k^{th} power map for polynomials. The result we seek asks for the following: Suppose $A \subset (\mathbb{F}_q[T])_{<n}$ is a set of polynomials of degree less than n such that $(A - A) \cap \mathrm{Im}(\Phi) = (0)$ (here 0 denotes the zero polynomial), then how large can A be? We wish to show that $|A| \leq q^{(1-c)n}$ for some $c = c(k, q)$, that we shall determine in a short while now.

Taking a combinatorial cue from the Ellenberg-Gisjwijt proof, let us consider an $|A| \times |A|$ matrix $\mathcal{M}$ whose $(a, b)^{th}$ entry equals $|\Phi^{-1}(a - b)|$. Again, the choice is natural since this choice of entries coupled with the observation of the preceding paragraph implies that $\mathcal{M}$ is in fact the identity matrix of order $|A|$. Now, to get a bound on the rank, we could use the same Croot-Lev-Pach trick, provided we can represent these matrix entries arising from some polynomial $h(X_1, \ldots, X_n, Y_1, \ldots, Y_n)$.

If we have such a polynomial of degree at most d, say, then again, writing

$$h(\mathbf{x}-\mathbf{y}) = \sum_{\mathfrak{m}\in\mathcal{W}_{n,d/2}} \mathfrak{m}(\mathbf{x})M_{\mathfrak{m}}(\mathbf{y}) + \sum_{\mathfrak{m}\in\mathcal{W}_{n,d/2}} \mathfrak{m}(\mathbf{y})M_{\mathfrak{m}}(\mathbf{x}),$$

so that again, the rank of $|A| = \mathcal{M} \leq 2m_{n,d/2}$. Thus, the problem has now shifted to showing that the function $\psi(\mathbf{x}) = |\Phi^{-1}(\mathbf{x})|$ for $\Phi = (g_0, \ldots, g_{n-1}) : \mathbb{F}_q^m \to \mathbb{F}_q^n$ arising as outlined earlier can be described by a polynomial of 'somewhat small' degree.

Now, all functions defined over finite fields and taking values in $\mathbb{F}_q$ come from polynomials. The relevant function here (that defined for us the matrix entries) is integer valued, so we modify the $(a,b)^{th}$ entry of $\mathcal{M}$ to be $|\Phi^{-1}(a-b)| \pmod p$ where $q = p^r$ for some r.

Now, some general observations. As remarked earlier, every function $\psi : \mathbb{F}_q^n \to \mathbb{F}_q$ can be realized as a polynomial. We write

$$F_\psi(\mathbf{x}) = \sum_{\mathbf{a}\in\{0,1,\ldots,q-1\}^n} \widehat{\psi(\mathbf{a})}\mathbf{x}^{\mathbf{a}}, \tag{1}$$

its 'Fourier' expansion, as we may call it (following Ben Green). The main point to note is this: Suppose the degree of F_ψ is small, say d (which we shall estimate in a moment). Then we turn again, to the Croot-Lev-Pach trick: The entries of $\mathcal{M}$ are $F_\psi(\mathbf{x}-\mathbf{y})$, and since we have, as before,

$$F_\psi(\mathbf{X}-\mathbf{Y}) = \sum_{\mathfrak{m}\in\mathcal{W}_{n,d/2}} \mathfrak{m}(\mathbf{X})M_{\mathfrak{m}}(\mathbf{Y}) + \sum_{\mathfrak{m}\in\mathcal{W}_{n,d/2}} \mathfrak{m}(\mathbf{Y})M_{\mathfrak{m}}(\mathbf{X})$$

, it follows again, for the same reasons, that $|A| = r(\mathcal{M}) \leq 2m_{n,d/2}$. Again, we shall then use Hoeffding's inequality to bound the RHS here.

So, the problem reduces to: If F_ψ is as described above, with $\psi(x) = |\Phi^{-1}(x)|$ where $\Phi = (g_0, \ldots, g_{n-1}) : \mathbb{F}_q^m \to \mathbb{F}_q^n$, then what is the degree (or a bound on the degree) of F_ψ? Note that since the map Φ captures the map $a(T) \to a(T)^k$ in $\mathbb{F}_q[T]$, it follows that each g_i has degree at most k.

Note that stating that the degree of F_ψ is at most d is equivalent to verifying that $\widehat{\psi(\mathbf{a})} = 0$ whenever $|\mathbf{a}| := a_1 + \cdots + a_n > D$. So, if we have an explicit formula for $\widehat{\psi(\mathbf{a})}$, then we could use that to some avail.

This, reduces to more standard fare, especially, if you think of F_ψ as a Fourier expansion. Let us write F instead of F_ψ; we claim

$$\widehat{\psi(\mathbf{a})} = \sum_{\mathbf{x}\in\mathbb{F}_q^n} F(\mathbf{x}) \prod_{i=1}^{n} \sigma_{a_i}(x_i) \tag{2}$$

where $\sigma_a : \mathbb{F}_q \to \mathbb{F}_q$ is defined as

$$\begin{aligned} \sigma_a(x) &= -x^{q-1-a} \qquad \text{if } a > 0 \\ &= 1 - x^{q-1} \qquad \text{if } a = 0. \end{aligned}$$

For the moment, let us assume this claim. Then, note that

$$\widehat{\psi(\mathbf{a})} = \sum_{\mathbf{x}\in\mathbb{F}_q^n} \left|\Phi^{-1}(\mathbf{x})\right| \prod_1^n \sigma_{a_i}(x_i) = \sum_{\mathbf{y}\in\mathbb{F}_q^m} \prod_1^n \sigma_{a_i}(g_{i-1}(\mathbf{y})).$$

The main point to note here is that for all a_i, $\sigma_{a_i}(T)$ is a polynomial (in T) of degree $q-1-a_i$, so since each g_i has degree at most k, it follows that $\widehat{\psi(\mathbf{a})}$ is a sum (over all $\mathbf{y}\in\mathbb{F}_q^m$) of monomials (in the y_i) of degree $\sum_i (q-1-a_i)k = (q-1)k - k|\mathbf{a}|$. But here is a useful fact about summing polynomials over finite fields; if $h(x_1,\ldots,x_m)$ is a monomial of degree less than $q-1$, then

$$\sum_{\mathbf{x}\in\mathbb{F}_q^m} h(x_1,\ldots,x_m) = 0.$$

Hence, if $|\mathbf{a}| > (q-1)(n-m/k)$ then it follows that $\widehat{\psi(\mathbf{a})} = 0$; that proves that degree of F_ψ is at most $d = (q-1)(n-m/k)$. Then, coupling this observation with the consequence from the Croot-Lev-Pach trick, and plugging $m = \frac{n+k-1}{k}$, and using Hoeffding's inequality, gives us that if $A-A$ does not admit a non-zero element which is a k^{th} power, then

$$|A| \leq 2m_{n,d/2} \leq 2q^n \exp\left(-\frac{(n+k-1)^2}{8nk^2d^2}\right) \leq 2q^{(1-c)n} \text{ with } c = c(n,q) = O(k^2d^2\log q)^{-1}. \tag{3}$$

Finally, the claim (regarding the coefficient $\widehat{\psi(\mathbf{a})}$). This is more or less standard fare, so I shall only give a sketch. Substituting for $\mathcal{F}(\mathbf{x})$ using (1) in the RHS of (2) we have

$$\sum_{\mathbf{x}\in\mathbb{F}_q^n} F(\mathbf{x}) \prod_{i=1}^n \sigma_{a_i}(x_i) = \sum_{\mathbf{b}} \widehat{\psi(\mathbf{b})} \left(\sum_{\mathbf{x}\in\mathbb{F}_q^n} \mathbf{x}^{\mathbf{b}} \prod_{i=1}^n \sigma_{a_i}(x_i) \right). \tag{4}$$

And now, we argue that if $\mathbf{b} = \mathbf{a}$ then the inner sum equals 1 and is zero otherwise. Indeed, suppose $\mathbf{b} = \mathbf{a} = (a_1,\ldots,a_r,0,\ldots,0)$. Then the inner sum in (4) equals

$$\sum_{\mathbf{x}\in\mathbb{F}_q^n} x_1^{a_1}\cdots x_r^{a_r}(-x_1^{q-1-a_1})\cdots(-x_r^{q-1-a_r})\mathbf{1}_{\{x_i=0\ \forall r<i\leq n} = \sum_{x_1,\ldots,x_r\in\mathbb{F}_q^*} (-1)^r x_1^{q-1}\cdots x_r^{q-1} = 1$$

as claimed. The other part is similar, and relies on the observation made earlier, that any monomial of degree less than $q-1$ sums to zero when summed over all $x\in\mathbb{F}_q$. I shall omit the details.

Green actually gives a better bound than the one we have up in 3 by doing the following simple trick. Write $k = k_0 + k_1 q + \cdots + k_r q^r$ (writing k in bases q), and then note that for any polynomial $g(T)$,

$$(g(T))^k = (g(T))^{k_0}\cdot(g(T))^{k_1}\cdots(g(T))^{k_r} = (g(T))^{k_0+\cdots+k_r},$$

so if we set $D(k,q)$ to denote the sum of the digits of k in q-ary notation, then the degrees of the g_i can be improved to $D(k,q)$ instead of d we had earlier. This (with some other calculations) yields a better bound of $|A| \leq 2q^{(1-c)n}$ with $c = c(k,q) = (2k^2D^2(k,q)\log q)^{-1}$.

5 Conclusion

We have seen four problems whose solutions invoke the polynomial method, and this is in fact only the tip of the iceberg. I believe that over the next few years, we are bound to see many ore, quite sophisticated applications of this technique which fuse more advanced ideas from algebraic geometry. Some more sophisticated tools appear in the monograph of Guth [12], but even those, in my opinion, are only the beginning. The future probably holds more sophisticated tools that may lead to solutions of several other open problems.

References

[1] N. Alon, Combinatorial Nullstellensatz, *Combin. Probab. Comput.* **8**(1999), 7-29.

[2] L. Babai, P. Frankl, *Linear Algebra Methods in Combinatorics*, lecture notes, University of Chicago, 1988.

[3] F. A. Behrend, On sets of integers which contain no three terms in arithmetical progression. Proc. Nat. Acad. Sci. U. S. A. 32, (1946). 331-332.

[4] B. Chazelle, H. Edelsbrunner, L. Guibas, R. Pollack, R. Seidel, M. Sharir, and J. Snoeyink, Counting and cutting cycles of lines and rods in space. *Comput. Geom.* **1** (1992), no. 6, 305-323.

[5] E. Croot, V. Lev, and P. Pach, Progression-free sets in $\mathbb{Z}_4^n$ are exponentially small, *Ann. of Math. (2)* **185**(2017), no. 1, 331337.

[6] Z. Dvir, On the size of Kakeya sets in finite fields. *J. Amer. Math. Soc.*, **22**(2009), no. 4, 1093-1097.

[7] Z. Dvir, S. Kopparty, S. Saraf, and M. Sudan, Extensions to the method of multiplicities, with applications to Kakeya sets, and mergers, 50^{th} *Annual IEEE Symp. Found. Comp. Sci. (FOCS)*, 2009, 181-190.

[8] Y. Edel, Extensions of generalized product caps, *Des. Codes. Crypt.*, **31**(2004), no. 1, 5-14.

[9] M. Elkin, An improved construction of progression-free sets, *Israel J. Math.*. **184** (2011), Vol 1, 93-128.

[10] J. Ellenberg, D. Gisjwijt, On large subsets of $\mathbb{F}_q^n$ with no three-term arithmetic progression. *Ann. of Math.*(2), **185** (2017), no. 1, 339-343.

[11] B. Green, Sárközy's theorem in Function fields, *Q. J. Math.* **68** (2017), no. 1, 237-242.

[12] L. Guth, *Polynomial Methods in Combinatorics*, **64** University Lecture Series, American Math Society, 2016.

[13] L. Guth, and N. Katz, Algebraic methods in discrete analogs of the Kakeya problem. *Adv. Math.* **225**(2010), no. 5, 2828-2839.

[14] J. Matoušek, *33 Miniatures in Linear Algebra: Mathematical and algorithmic applications of linear algebra*, Student Mathematical Library, Vol. 53, American Mathematical Society, 2010.

[15] T. Sanders, On Roth's theorem on progressions. *Ann. of Math.*, **174**(2011), Vol. 1, 619-636.

[16] A. Sárközy, On difference sets of sequences of integers III, *Acta Arith. Acad. Scient. Hungar.*, **3**(1978), 355-386.

[17] S. Saraf, M. Sudan, Improved lower bound on the size of Kakeya sets over finite fields, *Anal. PDE*, **1**(2008), No. 3, 375-379.

[18] T. Tao, Algebraic combinatorial geometry: the polynomial method in arithmetic combinatorics, incidence combinatorics, and number theory, *EMS Surv. Math. Sci.* **1** (2014), 1-46.