POLYNOMIAL METHOD IN COMBINATORICS

NEIL KOLEKAR

Abstract. The goal of this paper is to prove the combinatorial Nullstellensatz, an algebraic technique that has several applications in combinatorics, and can be applied through a process coined the Polynomial Method in Combinatorics, with applications specifically in Additive Combinatorics, Graph Theory, the Kakeya Set problem, and Cap Set problem.

1. What is the Polynomial Method?

The Polynomial Method in Combinatorics is a modern method used to solve a certain class of combinatorial problems by rephrasing the combinatorial nature of the problem into algebraic terms by encoding it in a polynomial. A basic outline of the method is the following:

(1) Associate the problem with some points of a vector space.
(2) Find the polynomial of minimal degree for which it vanishes on these points.
(3) Use the polynomial to settle the problem.

Notice that (3) is rather nontrivial – this is where we use Nullstellensatz, which we discuss in the following section.

2. Combinatorial Nullstellensatz

A well-known theorem of Hilbert, called Hilbert's Nullstellenatz, states that if F is an algebraically closed field, and $f, g_1, \ldots, g_m$ are polynomials in the ring of polynomials $F[x_1, \ldots, x_n]$, where f vanishes over all common zeros of $g_1, \ldots, g_m$, then there is an integer k and polynomials $h_1, \ldots, h_m$ in $F[x_1, \ldots, x_n]$ so that

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

In a new Nullstellensatz devised by Alon in 1999 (see [Alo99]), a more specific case where $m = n$ and each g_i takes the form $\prod_{s \in S_i}(x_i - s)$ is looked at, in which a stronger conclusion is claimed.

Theorem 2.1 (Combinatorial Nullstellensatz)**.** *Let F be an arbitrary field, and let $f = f(x_1, \ldots, x_n)$ be a polynomial in $F[x_1, \ldots, x_n]$. Let $S_1, S_2, \ldots, S_n$ be nonempty subsets of F and define $g_i(x_i) = \prod_{s \in S_i}(x_i - s)$. If $f(s_1, ..., s_n) = 0$ for all $s_i \in S_i$), then there are polynomials $h_1, \ldots, h_n \in F[x_1, \ldots, x_n]$ satisfying $\deg(h_i) \leq \deg(f) - \deg(g_i)$ so that*

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

A second formulation of Nullstellensatz can be seen as follows.

Theorem 2.2. *Let F be an arbitrary field, and let $f = f(x_1, ..., x_n)$ be a polynomial in $F[x_1, ..., x_n]$. 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 $\prod_{i=1} {x_i}^{t_i}$ in f is nonzero. Then, if $S_1, \ldots, S_n$ are subsets of F with $|S_i| > t_i$, there exist $s_1 \in S_1, s_2 \in S_2, \ldots, s_n \in S_n$ so that $f(s_1, \ldots, s_n) \neq 0$.*

The formulation in Theorem 2.2 is more well-used than in Theorem 2.1 in its applications. In particular, we will always be using Theorem 2.2 in the problems that we solve in this paper.

Date: February 2023.

2.1. **Introductory Consequences of Nullstellensatz.** In this section, we will present two classical applications of the combinatorial Nullstellensatz: the Chevalley-Warning theorem (initially proved by Chevalley in 1935) and the Cauchy-Davenport inequality (initially proved by Cauchy in 1813).

Theorem 2.3 (Chevalley-Warning)**.** *Let p be a prime, and 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. Assume for contradiction that $(c_1, \ldots, c_n)$ is the only common zero of the polynomials $P_1, \ldots, P_m$. Then, define the polynomial of n variables f by

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

where δ is chosen so that $f(c_1, \ldots, c_n) = 0$. Furthermore, it follows that $\delta \neq 0$ and that f vanishes on $(\mathbb{Z}_p)^n$.

Finally, for all $1 \leq i \leq n$, set $t_i = p-1$ and $S_i = \mathbb{Z}_p$; note that as the coefficient of $\prod_{i=1}^{n} x_i^{t_i}$ in f is $-\delta \neq 0$, and f vanishes on $S_1 \times \cdots \times S_n$, Theorem 2.2 implies that $f(s_1, \ldots, s_n) \neq 0$ for some $(s_1, \ldots, s_n) \in (\mathbb{Z}_p)^n$, which contradicts the earlier obtained fact that f vanishes on $(\mathbb{Z}_p)^n$, and we are done. ∎

Theorem 2.4 (Cauchy-Davenport)**.** *Given non-empty subsets A, B of $\mathbb{F}_p$, for some prime p, the following holds:*

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

Proof. When $|A|+|B| > p$, the result is trivial, so it suffices to show that the result holds when $|A|+|B| \leq p$; assume that $|A| + |B| \leq p$ for the remainder of the proof.

Assume for contradiction that $|A + B| < |A| + |B| - 1$. Choose some subset C of $\mathbb{Z}_p$ with $A + B \subset C$ and $|C| = |A| + |B| - 2$. Now, define the bivariate polynomial f by

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

Notice that f vanishes over $A \times B$ as $A + B \subset C$. Additionally, the coefficient of $x^{|A|-1}y^{|B|-1}$ in $f(x, y)$ is $\binom{|C|-1}{|A|-1}$, which is nonzero in $\mathbb{Z}_p$. Hence, by setting $t_1 = |A| - 1$, $t_2 = |B| - 1$, $S_1 = A$, and $S_2 = B$, Theorem 2.2 yields a contradiction, which completes the proof. ∎

3. Permanents

The Permanent Lemma is a reformulation of Theorem 2.2 in the context of permanents. The motivation for doing so is to (with more ease) encounter problems that involve expressions that easily arise from an element of a product of a matrix and a vector.

Lemma 3.1 (Permanent Lemma)**.** *Let $A = (a_{ij})$ be an $n \times n$ matrix over a field F such that its permanent is nonzero over F. Then for any vector $b = (b_1, \ldots, b_n) \in \mathbb{F}_n$ and for any family of subsets $S_1, \ldots, S_n$ of F, each of size at least 2, there is a vector $\boldsymbol{x} \in S_1 \times \ldots \times S_n$ such that for every i the ith coordinate of $A\boldsymbol{x}$ differs from b_i.*

The permanent-based proof of the following theorem due to Erdos, Ginzburg, and Ziv demonstrates the utility of permanents in solving certain combinatorial problems. An alternative approach would be considerably longer, including the proof that utilizes only the formulation in Theorem 2.2.

Theorem 3.2 (Erdos-Ginzburg-Ziv)**.** *For any prime p, any sequence of $2p - 1$ members of $\mathbb{Z}_p$ contains a subsequence of cardinality p such that the sum of its members is 0 in $\mathbb{Z}_p$.*

Proof. Consider the sequence $a_1, \ldots, a_{2p-1}$ with all elements in $\mathbb{Z}_p$, and without loss of generality suppose that $0 \leq a_1 \leq a_2 \leq \ldots a_{2p-1}$. Consider the $(p-1) \times (p-1)$ matrix A consisting entirely of 1s. Set $S_i = \{a_i, a_{i+p-1}\}$ and $t_i = 1$ for $1 \leq i \leq p-1$; given any vector $\mathbf{b} \in (\mathbb{Z}_p)^n$, the permanent lemma ensures that there exists a collection of $p-1$ elements of the subsequence such that they do not sum up to any of $\mathbf{b}$'s coordinates. Furthermore, no such collection ever has a_{2p-1} in it, and hence by letting all the coordinates of $\mathbf{b}$

be elements of $\mathbb{Z}_p \backslash \{-a_{2p-1}\}$, the permanent lemma applies to obtain that there exists some $\mathbf{x} \in S_1 \times \cdots \times S_n$ with

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

as required. ■

4. Graph Theoretic Applications

In this chapter, we survey two graph theoretic applications of the Polynomial Method, the first being a generalization of a theorem of Taśkinov on subgraphs, and the second being a theorem of Alon on list colorings.

4.1. Subgraphs.

A well-known conjecture of Berge and Sauer, proved by Taśkinov in 1982, states that any simple 4–regular graph contains a 3–regular subgraph. Though this false for graphs with multiple edges, adding one extra edge suffices to ensure the existence of a 3–regular subgraph. A more general version of this theorem can be seen to be proved rather quickly with the Polynomial Method.

Theorem 4.1. *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 any edge $e \in E$, define its indicator x_e as 1 when e is in the subgraph and 0 otherwise. Furthermore, for $v \in V$ and $e \in E$, let $a_{v,e}$ be 1 when v is incident to e and 0 otherwise. Note that, for fixed $v \in V$, $\sum_{e \in E} a_{v,e} x_e \equiv 0 \pmod{p}$ is equivalent to stating that exactly p or no edges have been chosen. Define

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

Notice that the first product is nonzero (1) if and only if for each $v \in V$, $\deg(v) \in \{0, p\}$; the second product is 1 if and only if no edges are picked, and 0 otherwise. All in all, f is zero unless the subgraph induced by the selected edges is p–regular.

The degree of the first product is $(p-1)|V|$, although the degree of f is $|E|$, which is bigger than $(p-1)|V|$ by the average degree hypothesis[1] so the coefficient of $\prod_{e \in E} x_e$ in f is nonzero. Hence Theorem 2.2 applies by setting $S_i = \{0, 1\}$ and $t_i = 1$ for $1 \leq i \leq |E|$. ■

Notice that the aforementioned theorem of Taśkinov is the $p = 3$ case of Theorem 4.1.

4.2. List Colorings.

We first recall the following definition concerning list colorings in Graph Theory.

Definition 4.2. A simple graph G is *k-choosable* if its possible to properly color its vertices given a list of k colors at each vertex.

In this section we prove the following theorem (due to Alon) about choosability.

Theorem 4.3 (Alon)**.** *A bipartite graph G is $\lfloor L(G) \rfloor + 1$ choosable, where*

$$L(G) \stackrel{\text{def}}{=} \max_{H \subseteq G} |E(H)|/|V(H)|.$$

(Notice that $L(G)$ is also half the maximum of the average degree of the subgraphs H of G.)

It is highly obscure how the Polynomial Method can be used to prove a result on list colorings; the key to using it is to show that there exists an *directed orientation* of any nonempty graph G with certain conditions we desire, and consider the *graph polynomial* of a directed orientation of G, whence we may exemplify the use of combinatorial Nullstellensatz.

Definition 4.4 (Graph Polynomial)**.** For a graph G, its graph polynomial is the polynomial f_G defined by

$$f_G(x_1, \ldots, x_n) = \prod_{\substack{(i,j) \in E(G) \\ i<j}} (x_i - x_j).$$

[1]Since the average degree is bigger than $2p-2$ as assumed, it follows that $|E| > \frac{1}{2}|V|(2p-2) = (p-1)|V|$.

We consider orientations of G as well; in this we simply direct each edge to obtain a directed graph, to which Definition 4.4 applies. In an orientation, an edge $v \to w$ is *ascending* if $v \leq w$. The parity of a graph is determined as *even* if it has an even number of edges and *odd* if it has an odd number of edges.

Definition 4.5. Define $\mathrm{DE_G}(d_1, \ldots, d_n)$ to be the set of all even orientations of G in which vertex i has indegree d_i, and define $\mathrm{DO_G}(d_1, \ldots, d_n)$ in a similar manner. Finally, we define $\mathrm{D_G} = \mathrm{DE_G}(d_1, \ldots, d_n) \cup \mathrm{DO_G}(d_1, \ldots, d_n)$.

We summarize the above terminology in the following lemma.

Lemma 4.6. *The coefficient of $\prod_{i=1}^{n} x_i^{d_i}$ in the graph polynomial of G is*

$$|DE_G(d_1, \ldots, d_n)| - |DO_G(d_1, \ldots, d_n)|.$$

Proof. In the expansion of f_G, each term corresponds to a choice of x_i and x_j for each (i, j) where $i < j$. Furthermore, a term has coefficient $+1$ if its corresponding orientation is even, and -1 if its corresponding orientation is odd. The conclusion follows, and we are done. ■

We now introduce Eulerian suborientations, which is a class of specific suborientations.

Definition 4.7. For some $D \in \mathrm{D_G}(d_1, \ldots, d_n)$, an *Eulerian suborientation* of D is a subgraph of D in which every vertex has equal indegree and outdegree. We say that such an Eulerian suborientation is *even* if it has an even number of edges, and *odd* if it has an odd number of edges. Furthermore, denote the sets of even and odd Eulerian suborientations of D by $\mathrm{EE}(D)$ and $\mathrm{EO}(D)$, respectively.

The following lemma establishes the interrelation of $\mathrm{DE_G}$ and $\mathrm{DO_G}$ with $\mathrm{EE}(D)$ and $\mathrm{EO}(D)$.

Lemma 4.8. *If $D \in DE_G(d_1, \ldots, d_n)$, then there are bijections*

$$DE_G(d_1, \ldots, d_n) \to EE(D)$$
$$DO_G(d_1, \ldots, d_n) \to EO(D).$$

Similarly, if $D \in DO_G(d_1, \ldots, d_n)$ then there are bijections

$$DE_G(d_1, \ldots, d_n) \to EO(D)$$
$$DO_G(d_1, \ldots, d_n) \to EE(D).$$

Proof. Consider any orientation $D' \in \mathrm{D}_G(d_1, \ldots, d_n)$. Let $D \star D'$ be the suborientation of D by including exactly the edges of D whose orientation in D' is in the opposite direction. This establishes a bijection $f : \mathrm{D}_G(d_1, \ldots, d_n) \to \mathrm{EE_G}(D) \cup \mathrm{EO_G}(D)$. Finally, realize that $D \star D'$ is even if D and D' are both even or both odd, and is odd otherwise. The desired result follows. ■

But why is Lemma 4.8 useful? Because when substituted in Lemma 4.6, we have the following:

Lemma 4.9. *The coefficient of $\prod_{i=1}^{n} x_i^{d_i}$ in the graph polynomial of G is*

$$\pm(|EE(d_1, \ldots, d_n)| - |EO(d_1, \ldots, d_n)|).$$

Proof. Note that the bijections in Lemma 4.8 preserve cardinality, and then apply Lemma 4.6 to finish. ■

Lemma 4.9 is sufficient for Theorem 2.2 to be applied, which justifies the introduction of Eulerian suborientations in the first place.

Theorem 4.10. *Let G be a graph on $\{1, \ldots, n\}$, and let $D \in D_G(d_1, \ldots, d_n)$. If $|EE(D)| \neq |EO(D)|$, then G is $(1 + \max_i d_i)$-choosable.*

Proof. We wish to, given $(1 + \max_i d_i)$ colors, properly color G with the properties assumed in the theorem statement. Then, because $|\mathrm{EE}(D)| \neq |\mathrm{EO}(D)|$, the coefficient of $\prod_{i=1}^{n} x_i^{d_i}$ in f_G is nonzero. Regarding the colors as real numbers, let S_i be the set of colors at vertex i so that $|S_i| = 1 + d_i$ for $1 \leq i \leq n$. Hence, by Theorem 2.2, one can select a color from each S_i so that f_G does not vanish, from which the conclusion follows. ■

It remains to find an orientation in which each indegree is at most $\lceil L(G) \rceil$. We prove this as a technical fact below.

Theorem 4.11. *Let G be a nonempty graph. Then, G has an orientation in which every indegree is at most $\lceil L(G) \rceil$.*

Proof. Let $E = E(G)$ and define

$$X = \underbrace{V(G) \sqcup V(G) \sqcup \ldots \sqcup V(G)}_{\lceil L(G) \rceil \text{ times}}.$$

Construct the bipartite graph $E \cup X$, where we connect $e \in E$ to $v \in X$ if v is an endpoint of e. By Hall's marriage theorem, we can match each edge in E to some vertex in X; as there are exactly $\lceil L(G) \rceil$ copies of each vertex in X, the statement follows. ■

Using the machinery of Theorems 4.10 and 4.11, we prove the final theorem.

Proof of Theorem 4.3. By Theorem 4.11, we can select $D \in \mathrm{D}_G(d_1, \ldots, d_n)$ where $\max_i d_i \leq \lceil L(G) \rceil$. Since G is bipartite, $\mathrm{EO}(D) = \emptyset$, so Theorem 4.10 applies and completes the proof. ■

5. Finite Field Kakeya Conjecture

A well-known problem in geometric measure theory is the *Kakeya conjecture*. The motivation behind its proposal began with the Kakeya needle problem.

Question 5.1 (Kakeya needle problem)**.** *What is the least area in which a line segment of width* 1 *can be freely rotated?*

In 1928, Besicovitch showed that there is no minimum area, or that the needle can be rotated using an arbitrarily small amount of positive area. His proof relied on the fact that translating the needle requires a set of zero measure. While Besicovitch's observation closed the question in $\mathbb{R}^2$, the question of whether a construction exists in $\mathbb{R}^n$, as of the time of this writing, is still open. At the center of this generalization is the notion of a Kakeya set.

Definition 5.2. A Kakeya set $K \subseteq \mathbb{R}^n$ is a set such that for all $y \in \mathbb{R}^n$, there is a line with direction y that is contained in K.

Now, the conjecture is as follows:

Conjecture 5.3 (Kakeya conjecture)**.** *Every Kakeya set K of $\mathbb{R}^n$ has Minkowski dimension and Hausdorff dimension equal to n.*

In 1999, Wolff proposed a analogous finite field version of the above conjecture (see Corollary 5.6) intended to be simpler, which was later proved by Dvir in 2008. We prove a more general version of it in Theorem 5.5, due to Tao and Alon (see [Tao13]), who modified Dvir's argument to arrive at its statement. First, we must define Kakeya sets in $\mathbb{F}^n$.

Definition 5.4. A Kakeya set $K \subseteq \mathbb{F}^n$ is a set such that for all $y \in \mathbb{F}^n$, there is a line with direction y that is contained in K.

Theorem 5.5 (Finite Field Kakeya Conjecture, Tao and Alon)**.** *Let $K \subseteq \mathbb{F}_q^n$ be a Kakeya set. Then $|K| \geq \binom{q+n-1}{n}$.*

Proof. Assume to the contrary that $|K| < \binom{q+n-1}{n}$. A "stars and bars" argument implies that 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.

Note that K is a Kakeya set, and so for each nonzero direction $\mathbf{v} \in {\mathbb{F}_q}^n$, there exists some point $\mathbf{x}$ such that for all $t \in \mathbb{F}$, $f(\mathbf{x} + t\mathbf{v}) = 0$. We can define $g(t)$ by $g(t) = f(x_1 + tv_1, \ldots, x_n + tv_n)$. Now let $t^d f_d(\mathbf{v})$ be the polynomial with maximal number of monomials, all of which have degree greater than or equal to d. Note that for fixed $\mathbf{x}$ and $\mathbf{v}$, g is zero for all t. As g has degree less than q, $f_d(\mathbf{v}) = 0$. Thus, f_d is a polynomial of degree at most $q-1$ vanishing on the whole of $\mathbb{F}_q^n$.

Now we finish with combinatorial Nullstellensatz: set $t_i = q-1$ and $S_i = \mathbb{F}_q$ for $1 \leq i \leq n$, so that there is a nonzero homogenous polynomial of degree less than q that vanishes on $S_1 \times \cdots \times S_n$, however, taking into consideration of any of its terms with a nonzero coefficient yields a contradiction to Theorem 2.2, as desired. ■

We may prove the initial proposal of the conjecture, as a consequence of Theorem 5.5.

Corollary 5.6. *Let $K \subseteq \mathbb{F}_q^n$ be a Kakeya set. Then $|K| \geq c_n q^n$, where $c_n > 0$ does not depend on q.*

In particular, by letting $c_n = \frac{1}{n!}$ and noting that $\binom{q+n-1}{n} \geq \frac{q^n}{n!}$, we see that Corollary 5.6, as initially proven by Dvir, is indeed true.

6. The Cap Set Problem

We give an introduction with the game of SET ([Aus16]). In the game of SET, we have a deck of 81 cards, and each card has one out of three values for several features. (These features can be things such as color, shape, shadowing, and number of the objects.) The objective 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. But in reality, SET is just a collection of all elements in $(\mathbb{Z}/3\mathbb{Z})^4$, and we want to find a three-term arithmetic progression. Furthermore, we can generalize this by considering elements of $(\mathbb{F}_q)^n$, to arrive at a game called *Cap Set*. The Cap Set problem concerns the sizes of subsets of $(\mathbb{F}_q)^n$ that contain no three-term arithmetic progressions. First, we define the following:

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

The main theorem is the following.

Theorem 6.1 (Cap Set Problem)**.** *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*

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

has no solutions $(a_1, a_2, a_3) \in A^3$ except when $a_1 = a_2 = a_3$. Then $|A| \leq 3m_{(q-1)n/3}$.

As this theorem has a highly involved proof, we present a sketch of it below. As opposed to the other problems studied in this paper, this one does not require the use of combinatorial Nullstellensatz, and instead uses the idea of the support of a polynomial and the Extremal Principle.

Proof. Let d be an integer between 0 and $(q-1)n$, inclusive, and let V be the set of polynomials in S_n^d vanishing on $(-\gamma A)^c = X$; using basic linear algebra, one can show that V has dimension at least $m_d - q^n + |A|$. Let $S(A) = \{g \in \mathbb{F}_q^n : g = \alpha a_1 + \beta a_2\}$. The technical fact that $S(A)$ and γA do not intersect holds, so all $f \in V$ vanish on $S(A)$. Therefore, if Σ is the support of $f \in V$, then $|\Sigma| \leq 2m_{d/2}$.

Furthermore, by picking $f \in V$ with maximal support, we claim $|\Sigma| \geq \dim(V)$. To prove this, assume to the contrary that $|\Sigma| < \dim(V)$. Then, take nonzero g such that g vanishes on Σ. Note that $f + g$ is nonzero on Σ, and thus there is some point $s \notin \Sigma$ at which it is nonzero, so $f + g$ is also nonzero at s. This constructs an element of V with larger support than f, which contradicts the maximality of the support of f.

Since $|\Sigma| \geq \dim(V)$ and $|\Sigma| \leq 2m_{d/2}$, so $|\Sigma| \leq 2m_{d/2}$. Choosing $d = 2(q-1)n/3$, we have $|A| \leq 2m_{(q-1)n/3} + q^n - m_{2(q-1)n/3}$. Recall from the previously defined definitions, $q^n - m_d$ is the number of monomials of M_n such that they are of degree strictly larger than d. Note that the bijection $x_1^{d_1} \cdots x_n^{d_n} \mapsto x_1^{(q-1)n-d_1} \cdots x_n^{(q-1)n-d_n}$ allows us to write $q^n - m_d \leq m_{(q-1)n-d}$. Hence,

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

as desired. ∎

In order to achieve a stronger bound on $|A|$, we must bound $m_{(q-1)n/3}$. One way to do so involves a probabilistic approach – let $X_1 \ldots, X_n$ be iid discrete random variables taking values on $\{0, 1, \ldots, q-1\}$ with uniform probability. Then,

$$\mathbb{P}\left(\sum_{i=1}^{n} X_i \leq \frac{n(q-1)}{3}\right) = m_{(q-1)n/3}/q^n$$

from which it follows that

$$\mathbb{P}\left(\sum_{i=1}^{n} \frac{X_i}{n} \leq \frac{(q-1)}{3}\right) = m_{(q-1)n/3}/q^n.$$

Now, one can view this as a large deviations problem and apply Cramer's theorem to obtain a bound of $O(2.756^n)$.

References

[Alo99] Noga Alon. "Combinatorial Nullstellensatz". In: (1999).

[Tao13] Terence Tao. "Algebraic combinatorial geometry: the polynomial method in arithmetic combinatorics, incidence combinatorics, and number theory". In: (2013).

[Aus16] David Austin. "Game. SET. Polynomial". In: (2016).