A Refinement of Multivariate Value Set Bounds

Luke Smith*, Daqing Wan

340 Rowland Hall (Bldg. # 400), University of California, Irvine, Irvine, CA 92697-3875.

Abstract

Over finite fields, if the image of a polynomial map is not the entire field, then its cardinality can be bounded above by a significantly smaller value. Earlier results bound the cardinality of the value set using the degree of the polynomial, but more recent results make use of the powers of all monomials.

In this paper, we explore the geometric properties of the Newton polytope and show how they allow for tighter upper bounds on the cardinality of the multivariate value set. We then explore a method which allows for even stronger upper bounds, regardless of whether one uses the multivariate degree or the Newton polytope to bound the value set. Effectively, this provides improvement of a degree matrix-based result given by Zan and Cao, making our new bound the strongest upper bound thus far.

Keywords: Value Set, polynomial image set, multivariate polynomials, Newton polytopes, p-adic liftings
2010 MSC: 11T06, 11T55, 11H06

1. Recent Multivariate Value Set Theorems

For a given polynomial $f(x)$ over a finite field $\mathbb{F}_q$, let $V_f := \mathrm{Im}(f)$ denote the value set of f. Determining the cardinality and structure of the value set

*Corresponding author
Email addresses: smithla@uci.edu (Luke Smith), dwan@math.uci.edu (Daqing Wan)

is a problem with a rich history and wide variety of uses in number theory, algebraic geometry, coding theory and cryptography.

Relevant to this paper are theorems which provide upper bounds on the cardinality of our value set when $f(x)$ is not a permutation polynomial.[1] These upper bounds have been extensively studied in the case of univariate polynomials (see [10] for more information), but results on multivariate polynomial maps have gained more recent attention.

A result published by Mullen, Wan, and Wang in 2012 [8] gives a bound on the value set of polynomial maps, one with no error terms:

Theorem 1.1. *Let $f(x_1, ..., x_n) = (f_1(x_1, ..., x_n), ..., f_n(x_1, ..., x_n))$ be a polynomial map over the vector space $\mathbb{F}_q^n$, and let* $\deg f = \max_i \deg f_i$.

$$\text{If } |V_f| < q^n, \text{ then } |V_f| \leq q^n - \min\left\{q, \frac{n(q-1)}{\deg f}\right\}.$$

Since the time their paper was published, multiple refinements have been made to this theorem.

One approach towards improving Theorem 1.1 is to replace the term $\frac{n(q-1)}{\deg f}$ by using different properties of the polynomial map f. Note that the degree only takes one monomial of f into account, so it is reasonable to expect tighter bounds on $|V_f|$ if we account for every monomial. Smith [10] improved upon Theorem 1.1 by generalizing Mullen, Wan, and Wang's p-adic lifting approach and utilizing the Newton polytope $\Delta(f)$ of the polynomial map f. The Newton polytope is constructed using all monomials of f using discrete geometry, meaning it encodes more information than deg f and allows for a stronger statement to be made:

Theorem 1.2 (Smith [10], 2014)**.** *Let $f(x_1, ..., x_n) = (f_1(x_1, ..., x_n), ..., f_n(x_1, ..., x_n))$ be a polynomial map over the vector space $\mathbb{F}_q^n$, let $\Delta(f)$ be the Newton polytope of f, and let μ_f be a certain constant (defined explicitly later) dependent on $\Delta(f)$.*

$$\text{If } |V_f| < q^n, \text{ then } |V_f| \leq q^n - \min\{q, \mu_f \cdot (q-1)\},$$

[1] Permutation polynomials have also been studied extensively in literature, in view of their application to cryptography and combinatorics. For more information about other ways value sets have been studied historically, please refer to [6].

Zan and Cao also refine Theorem 1.1 by using the degree matrix D_f of the polynomial map f in order to account for all of the monomials of f. Their approach generalizes the p-adic lifting technique as well and improves upon Smith's statement in [10]:

Theorem 1.3 (Zan, Cao [15], 2014)**.** *Let $f(x_1, ..., x_n) = (f_1(x_1, ..., x_n), ..., f_n(x_1, ..., x_n))$ be a polynomial map over the vector space $\mathbb{F}_q^n$ and let D_f be the degree matrix of f.*

$$\textit{If } |V_f| < q^n, \textit{ then } |V_f| \leq q^n - \min\{q, \omega_f\},$$

where the constant ω_f (defined explicitly later) depends on D_f.

Overall, each new refinement gives us stronger bounds, i.e. $\omega_f \geq \mu_f \cdot (q-1) \geq \frac{n}{\deg f}(q-1)$ (see [1] and [15]). In addition, in the univariate case, it has been shown that there are instances when ω_f is strictly larger than $\frac{q-1}{\deg f}$ (as opposed to μ_f always being equal to $\frac{1}{\deg f}$ when $n = 1$). However, since each of these bounds are of the form $|V_f| \leq q^n - \min\{C_f, q\}$ with C_f dependent on the theorem, we are limited to removing at most q elements from these cardinality bounds.

Another type of improvement on Theorem 1.1 removes this dependence on subtracting the minimum of two constants. Though still dependent on the polynomial map degree, a theorem by Kosters allows for a stronger bound whenever $n > \deg f$:

Theorem 1.4 (Kosters [7], 2014)**.** *Let $f(x_1, ..., x_n) = (f_1(x_1, ..., x_n), ..., f_n(x_1, ..., x_n))$ be a polynomial map over the vector space $\mathbb{F}_q^n$, and let $\deg f = \max_i \deg f_i$.*

$$\textit{If } |V_f| < q^n, \textit{ then } |V_f| \leq q^n - \frac{n(q-1)}{\deg f}.$$

In order to achieve this result, Kosters completely averted the use of p-adic liftings, instead using a method more akin to Turnwald's univariate proof in [11].

2. Main Result

In this paper, we will refine these multivariate value set bounds even further, removing the minimum condition from Theorems 1.2 and 1.3, ultimately proving the following theorem:

Theorem 2.1. *Let $f(x_1, ..., x_n) = (f_1(x_1, ..., x_n), ..., f_n(x_1, ..., x_n))$ be a polynomial map over the vector space $\mathbb{F}_q^n$ and let D_f be the degree matrix of f.*

$$\text{If } |V_f| < q^n, \text{ then } |V_f| \leq q^n - \omega_f,$$

where the constant ω_f depends on D_f.

To properly convey the significance of this bound in relation to prior bounds, we will describe the Newton polytope in Section 3 and the degree matrix in Section 4. We will also define the constants associated with these objects and connections between the two.

3. The Newton Polytope

Let F be an arbitrary field and let $h \in F[x_1, ..., x_n]$. If we write h in the form

$$h(x_1, ..., x_n) = \sum_{j=1}^{m} a_j X^{D_j}, \quad a_j \in F^* \tag{1}$$

where

$$D_j = (d_{1j}, ..., d_{nj})^T \in \mathbb{Z}_{\geq 0}^n, \quad X^{D_j} = x_1^{d_{1j}} \cdots x_n^{d_{nj}}, \tag{2}$$

then we have the following definition:

Definition 3.1 (Newton polytope)**.** The Newton polytope of polynomial $h \in F[x_1, ..., x_n]$, $\Delta(h)$, is the convex closure of the set $\{D_1, ..., D_m\} \cup \{(0, ..., 0)\}$ in $\mathbb{R}^n$.

Geometric properties of the Newton polytope, such as its dilation by $k \in \mathbb{R}$, its volume or its decomposition into other polytopes via Minkowski Sum, are useful tools in discerning properties of their associated polynomials. For more information, see [3], [13], and [12].

The significance of the Newton polytope to the multivariate value set problem comes from the definition of the following quantity:

Definition 3.2 (Minimal dilation factor μ_h)**.** Let F be a field, let $h \in F[x_1, ..., x_n]$, and let $\Delta(h)$ be the Newton polytope of h.

$$\mu_h := \inf\{k \in \mathbb{R}_{>0} \mid k\Delta(h) \cap \mathbb{N}^n \neq \varnothing\}.$$

In other words, μ_h is the infimum of all positive real numbers k such that the dilation of $\Delta(h)$ by k contains a lattice point with strictly positive coordinates, and we define $\mu_h = \infty$ if such a dilation does not exist. For our purposes, since the vertices of our polytopes have integer coordinates, μ_h will always be finite and rational so long as we consider h which is not a polynomial in some proper subset of $\{x_1, ..., x_n\}$. If h is polynomial in a proper subset of $\{x_1, ..., x_n\}$, then we may make a linear change of variables $\{z_1, ..., z_\nu\}$, $\nu < n$, which allows us to consider $\Delta(h(z_1, ..., z_\nu)) \subset \mathbb{R}^\nu$, where μ_h will be finite.

The quantity μ_f is used by Adolphson and Sperber [1] to put a lower bound on the q-adic valuation ord_q of the number of $\mathbb{F}_q$-rational points on a variety V, $N(V)$, over $\mathbb{F}_q$. Namely, let $V = Z(f_1, ..., f_m)$ be the vanishing set of $f_1, ..., f_n$, where $f_i \in \mathbb{F}_q[x_1, ..., x_n]$. If the collection of polynomials $f_1, ..., f_m$ is not polynomial in some proper subset of $x_1, ..., x_n$, then we have for $f(x_1, ..., x_n, y_1, ..., y_m) = f_1(x_1, ..., x_n)y_1 + \cdots + f_m(x_1, ..., x_n)y_m$,

$$\mathrm{ord}_q(N(V)) \geq \mu_f - m.$$

Note that in the above definitions, the multivariate polynomial h maps the vector space F^n into its base field F. However, for the value set problem, we are interested in studying the polynomial vector $f : \mathbb{F}_q^n \longrightarrow \mathbb{F}_q^n$. Fortunately, the definitions we have developed in this section can be extended to polynomial vectors. If we denote the support of h by $\Gamma(h) \coloneqq \{D_1, ..., D_m\}$, then we define $\Delta(f)$ to be the convex closure of $\Gamma(f_1) \cup \cdots \cup \Gamma(f_n) \cup \{(0, ..., 0)\}$ in $\mathbb{R}^n$.

4. The Degree Matrix and Comparison of Constants

For our multivariate polynomial h as in Section 3, we define the $n \times m$ degree matrix of h, $D_h \coloneqq (D_1, ..., D_m) \in \mathbb{Z}_{\geq 0}^{n \times m}$. The degree matrix has been used by Cao and his collaborators in [2], [4], and [5] in rational point counting and p-adic estimates. In relation to the value set problem, Zan and Cao use the degree matrix in [15] as a succinct way of keeping track of the exponent vectors D_j that does not explicitly rely on a geometry. Using this, they define the following invariant of h.

Definition 4.1 (Integral dilation factor ω_h)**.** Let F be a field, let $h \in$

$F[x_1, ..., x_n]$ be as in equation (1).

$$\omega_h := \min \left\{ \sum_{j=1}^{m} k_j \,\middle|\, k_j \in \{0, 1, ..., q-1\}, \sum_{j=1}^{m} k_j D_j \in (q-1)\mathbb{N}^n \right\}. \tag{3}$$

This constant can be thought of as the minimal number of exponent vectors (up to $q-1$ duplicates of each) needed to be summed together to reach a lattice point where all coordinates are positive multiples of $q-1$. Again, so long as h is not polynomial in some proper subset of $\{x_1, ..., x_n\}$, ω_h will always exist.

Though ω_f and μ_f may seem different by their definitions, a lemma in [3] gives us that

$$\mu_f = \min \left\{ \sum_{j=1}^{m} \alpha_j \,\middle|\, \alpha_j \in \mathbb{Q}_{\geq 0}, \sum_{j=1}^{m} \alpha_j D_j \in \mathbb{N}^n \right\}. \tag{4}$$

Intuitively, studying the dilation of $\Delta(f)$ is equivalent to studying linear combinations of the exponent vectors geometrically. Because of similarity, we can use both D_f and $\Delta(f)$ to study μ_f and ω_f.[2]

In fact, because of this similarity, we have a direct comparison of the two terms proven by [15]. This, alongside a result of Adolphson and Sperber [1], gives us the following inequalities:

Lemma 4.2. $\omega_f \geq \mu_f(q-1) \geq \frac{n(q-1)}{d}$.

Not only does ω_f provide a better value set bound for nonpermutation polynomials, but [15] gives sharp examples which improve previously known univariate bounds.

5. Single Variable Value Set

To provide insight towards the proof of our main result, we will investigate upper bounds of $|V_f|$ for the case when f is a single variable polynomial. Parts of this proof will generalize to the multivariate case.

[2]Our main theorem which is dependent on ω_f uses $\Delta(f)$ in the proof given.

Theorem 5.1. *Let $f(x) \in \mathbb{F}_q[x]$ be a single variable polynomial of degree $d > 0$. If $|V_f| < q$, then*

$$|V_f| \leq q - \frac{q-1}{d}.$$

The proof of this theorem relies on the following definition:

Definition 5.2 (The quantity $U(f)$)**.** Let $\mathbb{Z}_q$ denote the ring of p-adic integers with uniformizer p and residue field $\mathbb{F}_q$. Also let $\tilde{f}(x) \in \mathbb{Z}_q[x]$ be the lifting of f taking coefficients from the Teichmüller lifting $L_q \subset \mathbb{Z}_q$ of $\mathbb{F}_q$. Then we define $U(f)$ to be the smallest positive integer k such that the sum

$$S_k(f) \coloneqq \sum_{x \in L_q} \tilde{f}(x)^k \not\equiv 0 \pmod{pk}.$$

By taking into account the following sum,

$$\sum_{x \in L_q} x^k = \begin{cases} 0, & q-1 \nmid k, \\ q-1, & q-1 \mid k, k \neq 0, \\ q, & k = 0, \end{cases} \tag{5}$$

and remembering that we are only summing over a finite number of terms, we have that, for f not identically zero, $\frac{q-1}{d} \leq U(f)$. We also have that if f is a permutation polynomial, then $S_k(f) = S_k(x) = \sum_{x \in L_q} x^k$, implying $U(f) = q-1$. The fact that $U(f)$ exists for all nonpermutation polynomials as well is a corollary of lemma 5.3. Overall, the lemma and the above argument give us that

$$\frac{q-1}{d} \leq U(f) \leq q-1.$$

Theorem 5.1 also follows directly from the lemma 5.3:

Lemma 5.3. *If $|V_f| < q$, then*

$$|V_f| \leq q - U(f).$$

The proof of this result is given by Wan, Shiue, and Chen in [14], and their paper also includes more details regarding this lemma. Mullen, Wan, and Wang [8] also describe an alternate proof of this lemma presented to them by Lenstra through private communication.

6. From Single Variable to Multivariable

Let $f(x_1, ..., x_n) = (f_1(x_1, ..., x_n), ..., f_n(x_1, ..., x_n))$ be a polynomial vector, and note deg $f = \max_i\{\deg f_i\}$. This maps the vector space $\mathbb{F}_q^n$ to itself. Now, take a basis $e_1, ..., e_n$ of $\mathbb{F}_{q^n}$ over $\mathbb{F}_q$. Denote $x = x_1e_1 + \cdots + x_ne_n$ and define

$$g(x) \coloneqq f_1(x_1, ..., x_n)e_1 + \cdots + f_n(x_1, ..., x_n)e_n.$$

In this way, we can think of the function g as a non-constant univariate polynomial map from the finite field $\mathbb{F}_{q^n}$ to itself. Even better, we have the equality $|V_f| = |g(\mathbb{F}_{q^n})|$. Therefore, using Lemma 5.3, we know

$$\text{if } |V_f| < q^n, \text{then } |V_f| \leq q^n - U(g),$$

where g is viewed as a univariate polynomial.

Unfortunately, as a univariate polynomial, we do not have good control of the univariate degree of g in relation to the multivariate degree of f. Even if one were to construct a closed form for $g(x)$ using methods such as Lagrange Interpolation, the degree of g would likely be high enough as to make the resulting upper bound on $|V_f|$ trivial. Because of these issues with the degree of g, we cannot use the bounds from the previous section directly, and must rely on another method to bound $U(g)$.

Previously, we introduced $g(x)$ as a univariate polynomial. However, using a basis $e_1, ..., e_n$ of $\mathbb{F}_{q^n}$ over $\mathbb{F}_q$ as before, we can also define a multivariate polynomial

$$g(x_1, ..., x_n) \coloneqq f_1(x_1, ..., x_n)e_1 + \cdots + f_n(x_1, ..., x_n)e_n$$

mapping the vector space $\mathbb{F}_q^n$ into the field $\mathbb{F}_{q^n}$. In this sense, g as a multivariate polynomial shares some important properties with f as a polynomial vector, such as the fact that $\deg(g) = \max_i\{\deg f_i\}$. Whereas the paper by Mullen, Wan, and Wang determine a bound for $U(g)$ relying on the multivariate degree of f, in this paper we will use the Newton polytope of the multivariate polynomial $g(x_1, ..., x_n)$ to improve upon these bounds. With this in mind, we define $\Delta(f) \coloneqq \Delta(g(x_1, ..., x_n))$, $\mu_f \coloneqq \mu_{g(x_1,...,x_n)}$, $\omega_f \coloneqq \omega_{g(x_1,...,x_n)}$ and prove our main theorem.

7. A Method to Improve Prior Proofs

A major limitation of the p-adic liftings methods in the proofs given in [8], [10], and [15] is how they limit the p-divisibility we can discern from the sum $S_k(f)$. Indeed, if we immediately split $S_k(f)$ amongst the monomials of the multivariate polynomial $g(x_1, ..., x_n)^k$ (as is done in [8] and [10], and as is equivalent to the methods of [15]), we lose much of the structure and divisibility of each term. Therefore, we will manipulate our summand to leverage a larger p-adic valuation before splitting it into monomials. To do this, we need the following lemma:

Lemma 7.1. *Let $x_1, ..., x_n$ be in a commutative ring R, and let $e \in \mathbb{N}$. Then*

$$(x_1 + \cdots + x_n)^{p^e} = x_1^{p^e} + \cdots + x_n^{p^e} + ph_1(x_1^{p^{e-1}}, ..., x_n^{p^{e-1}}) + p^2h_2(x_1^{p^{e-2}}, ..., x_n^{p^{e-2}}) + \cdots + p^e h_e(x_1, ..., x_n)$$

where $h_t(x_1^{p^{e-t}}, ..., x_n^{p^{e-t}}) \in R[x_1, ..., x_n]$ is such that $\deg\, h_t(x_1, ..., x_n) = p^t$.

Proof. We use the multinomial theorem on the left hand side of the above equation.

$$(x_1 + \cdots + x_n)^{p^e} = x_1^{p^e} + \cdots + x_n^{p^e} + \sum_{\substack{a_1+\cdots+a_n=p^e \\ a_1 \neq p^e, ..., a_n \neq p^e}} \binom{p^e}{a_1, ..., a_n} x_1^{a_1} \cdots x_n^{a_n}. \quad (6)$$

For simplicity of notation, let

$$A = \binom{p^e}{a_1, ..., a_n} x_1^{a_1} \cdots x_n^{a_n}.$$

Then the sum in (6) can be split as follows:

$$\sum_{a_1+\cdots+a_n=p^e} A = x_1^{p^e} + \cdots + x_n^{p^e} + \sum_{\substack{a_1+\cdots+a_n=p^e \\ p^{e-1}||(a_1,...,a_n)}} A + \sum_{\substack{a_1+\cdots+a_n=p^e \\ p^{e-2}||(a_1,...,a_n)}} A + \cdots + \sum_{\substack{a_1+\cdots+a_n=p^e \\ p||(a_1,...,a_n)}} A + \sum_{\substack{a_1+\cdots+a_n=p^e \\ p \nmid a_\epsilon \text{ for some } \epsilon}} A.$$

Now, let

$$\sigma_t = \sum_{\substack{a_1+\cdots+a_n=p^e \\ p^{e-t}||(a_1,\ldots,a_n)}} A, \;\; 1 \leq t \leq e.$$

If we can show for $1 \leq t \leq e$ that σ_t has the form $p^t h_t(x_1^{p^{e-t}}, ..., x_n^{p^{e-t}})$ with deg $h_t(x_1, ..., x_n) = p^t$, then the proof is done.

Notice that the summand A always has degree $a_1 + \cdots + a_n = p^e$, which means deg $\sigma_t = p^e$. Since $p^{e-t}|a_\epsilon$ for all ϵ between 1 and n, we know that σ_t has the form $\tau_t(x_1^{p^{e-t}}, ..., x_n^{p^{e-t}}) \in R[x_1, ..., x_n]$ and deg $\tau_t(x_1, ..., x_n) = \frac{p^e}{p^{e-t}} = p^t$.

The fact that $p^t | \binom{p^e}{a_1,\ldots,a_n}$ under the conditions that $p^{e-t}||(a_1, ..., a_n)$ has an elegant proof by Singmaster in [9]. Therefore, we have $p^t|\tau_t(x_1^{p^{e-t}}, ..., x_n^{p^{e-t}})$. This tells us σ_t has the form $p^t h_t(x_1^{p^{e-t}}, ..., x_n^{p^{e-t}})$ with deg $h_t(x_1, ..., x_n) = p^t$, and thus the lemma is proved.

□

Let f be a polynomial map over $\mathbb{F}_q^n$, char $\mathbb{F}_q = p$. Also let $e_1, ..., e_n$ be a basis of the field $\mathbb{F}_{q^n}$ over $\mathbb{F}_q$, and let $x = x_1e_1 + \cdots + x_ne_n$ as before in Section 6, allowing us to identify the polynomial map $f(x_1, .., x_n) = (f_1(x_1, ..., x_n), ..., f_n(x_1, ..., x_n))$ with the univariate polynomial $f(x)$ or multivariate polynomial $f(x_1, ..., x_n) = f_1(x_1, ..., x_n)e_1 + \cdots + f_n(x_1, ..., x_n)e_n$. Also let $S_k(f)$ and $U(f)$ be as in Section 5.1. To improve upon the p-adic lifting method, we will apply Lemma 7.1 to $f(x_1, ..., x_n)^k$, split $S_k(f)$ amongst these polynomials, and then split the summand polynomials further into monomials.

Write $k = p^e k_1$ with $p \nmid k_1$. For simplicity of notation, assume f has already been lifted with coefficients in L_{q^n}. Then by Lemma 7.1, there exist polynomials $F_0, ..., F_e \in \mathbb{F}_{q^n}[x_1, ..., x_n]$ such that

$$\begin{aligned}(f_1(x_1, ..., x_n)e_1 + \cdots + f_n(x_1, ..., x_n)e_n)^{p^e} &= F_0(x_1^{p^e}, ..., x_n^{p^e}) + pF_1(x_1^{p^{e-1}}, ..., x_n^{p^{e-1}}) \\ &\quad + \cdots + p^e F_e(x_1, ..., x_n),\end{aligned}$$

where deg $F_t(x_1, ..., x_n) \leq dp^t$. This means that

$$(f_1(x_1,...,x_n)e_1+\cdots+f_n(x_1,...,x_n))^k = \Big(F_0(x_1^{p^e},...,x_n^{p^e})+pF_1(x_1^{p^{e-1}},...,x_n^{p^{e-1}}) + \cdots + p^eF_e(x_1,...,x_n)\Big)^{k_1}$$

$$= \sum_{b_0+\cdots+b_e=k_1} \binom{k_1}{b_0,...,b_e} p^{b_1+2b_2+\cdots+eb_e} F_0(x_1^{p^e},...,x_n^{p^e})^{b_0}\cdots F_e(x_1,...,x_n)^{b_e}.$$

Now for fixed $b_0,...,b_e$, let λ be the positive integer such that $b_\lambda \neq 0, b_{\lambda+1} = \cdots = b_e = 0$, and let $y_i = x_i^{p^{e-\lambda}}$. This means we can reduce the power and degree of our summand polynomials in the following way:

$$F_0(x_1^{p^e},...,x_n^{p^e})^{b_0}\cdots F_\lambda(x_1^{p^{e-\lambda}},...,x_n^{p^{e-\lambda}})^{b_\lambda} = F_0(y_1^{p^\lambda},...,y_n^{p^\lambda})^{b_0}\cdots F_\lambda(y_1,...,y_n)^{b_\lambda}. \quad (7)$$

Note that each term may have a different substitution, but we may split $S_k(f)$ amongst each summand to bound the p-divisibility of the entire sum. Using the reduction of $f(x_1,..,x_n)^k$ to (7), we are given sums of the form

$$p^{b_1+2b_2+\cdots+\lambda b_\lambda} \sum_{y_1,...,y_n\in L_q} F_0(y_1^{p^\lambda},...,y_n^{p^\lambda})^{b_0}\cdots F_\lambda(y_1,...,y_n)^{b_\lambda}. \quad (8)$$

Now the fact that $b_\lambda \neq 0$ tells us this sum is divisible by p^λ, i.e. $S_k(f) \equiv 0 \pmod{p^\lambda}$. From here, we must further split this summand product into monomials and determine the p-divisibility of the smaller sums. Let

$$F_0(y_1^{p^\lambda},...,y_n^{p^\lambda})^{b_0}\cdots F_\lambda(y_1,...,y_n)^{b_\lambda} = \sum_{j=1}^{m} c_j Y^{W_j}, \quad c_j \in \mathbb{F}_{q^n}^*,$$

where

$$W_j = (w_{1j},...,w_{nj})^T \in \mathbb{Z}_{\geq 0}^n, \quad Y^{W_j} = y_1^{w_{1j}}\cdots y_n^{w_{nj}}. \quad (9)$$

This allows the sum in (8), and ultimately $S_k(f)$, to be split among the monomials in (9) into sums of the form

$$p^{b_1+2b_2+\cdots+\lambda b_\lambda} \sum_{y_1,...,y_n\in L_q} c_j Y^{W_j} = c_j p^{b_1+2b_2+\cdots+\lambda b_\lambda} \prod_{i=1}^{n} \sum_{y_i\in L_q} y_i^{w_{ij}}. \quad (10)$$

Let C be an upper bound on k, i.e. $k < C$ (C will be made explicit in the next section). Our goal is to show $C \leq U(f)$ and therefore come up with a nicer bound on $|V_f|$ (thanks to Lemma 5.3). Let v_p be the p-adic valuation with $v_p(p) = 1$, and let ℓ_j be the number of nonzero entries of W_j. We can accomplish our goal by showing the sum in (10) is congruent to zero mod $p^\lambda q^{n-\ell_j}$, and that $v_p(p^\lambda q^{n-\ell_j}) \geq v_p(pk)$, i.e.

$$\lambda + (n - \ell_j)v_p(q) \geq e + 1.$$

If this holds true for all monomials, then $S_k(f) \equiv 0 \bmod pk$ and $C \leq U(f)$.

Now the sum in (10) equals zero if one of the w_{ij}'s is not divisible by $q-1$, so all that is left is to consider the case when $q-1|w_{ij}$ for all i. In this case, since $b_\lambda \neq 0$, and since ℓ_j is the number of nonzero w_{ij}, we have $n - \ell_j$ zero terms, which tells us

$$p^{b_1+2b_2+\cdots+\lambda b_\lambda} \sum_{y_1,\ldots,y_n \in L_q} c_j Y^{W_j} \equiv 0 \,(\text{mod } p^\lambda q^{n-\ell_j}).$$

The above substitution method allows us to refine the recently published results mentioned in Section 1, whose proofs simply used the monomials of $f(x_1, ..., x_n)^k$ directly. These proofs required that $k < q$ to bound the value set, but our proofs do not. The next section will show how the added structure our method provides a tighter upper bound on the cardinality of the value set.

8. Improved Integral Dilation Bound

Theorem 8.1. *Let $f(x_1, ..., x_n) = (f_1(x_1, ..., x_n), ..., f_n(x_1, ..., x_n))$ be a polynomial vector over the vector space $\mathbb{F}_q^n$. Without loss of generality, suppose f is not polynomial in some subset of $\{x_1, ..., x_n\}$. Let ω_f be the integral dilation factor associated with $\Delta(f)$. If $|V_f| < q^n$, then*

$$|V_f| \leq q^n - \omega_f.$$

Proof. If we can show that, for $1 \leq k < \omega_f$ and $k = p^e k_1$ with $p \nmid k_1$,

$$S_k(f) \coloneqq \sum_{x \in L_{q^n}} \tilde{f}(x)^k = \sum_{x_1,\ldots,x_n \in L_q} \left(\tilde{f}_1(x_1, ..., x_n)\tilde{e}_1 + \cdots + \tilde{f}_n(x_1, ..., x_n)\tilde{e}_n\right)^k \equiv 0 \,(\text{mod } pk),$$

then $U(f) \geq \omega_f$ and we are done by Lemma 5.3.

For simplicity of notation, assume f is already lifted to characteristic zero over L_{q^n}. Split $S_k(f)$ into sums of the form (8). Notice that, by our substitution and Lemma 7.1, the exponent vectors of the monomials of the product $F_0(y_1^{p^\lambda}, ..., y_n^{p^\lambda})^{b_0} \cdots F_\lambda(y_1, ..., y_n)^{b_\lambda}$ are contained in $\frac{k}{p^{e-\lambda}}\Delta(f)$. When we further split these sums into sums of the form (10), we have that

$$p^{b_1+2b_2+\cdots+\lambda b_\lambda} \sum_{y_1,...,y_n \in L_q} c_j Y^{W_j} \equiv 0 \,(\text{mod } p^\lambda q^{n-\ell_j}).$$

Since this sum equals 0 if one of the w_{ij}'s is not divisible by $q-1$, assume $q-1|w_{ij}$ for all i. By this assumption, we have that

$$W_j \in \frac{k}{p^{e-\lambda}}\Delta(f) \cap (q-1)\mathbb{Z}^n_{\geq 0}. \tag{11}$$

To further develop this proof, we require additional terminology.

Definition 8.2 (The quantity γ)**.**

$$\gamma := \min \left\{ |S| \,\middle|\, S \subseteq \{W_1, ..., W_m\}, \sum_{W_j \in S} W_j \in \mathbb{N}^n \right\}.$$

In other words, γ is the size of smallest subset of the exponent vectors, $\{W_1, ..., W_m\}$, such that the sum of its elements lie in $\mathbb{N}^n$. Since $f(x_1, ..., x_n)$ is not polynomial in some proper subset of $\{x_1, ..., x_n\}$, we have that the polynomials $F_0(y_1^{p^\lambda}, ..., y_n^{p^\lambda}), ..., F_\lambda(y_1, ..., y_n)$ are not either. This means γ will exist. Also, assume without loss of generality that $W_1, ..., W_\gamma$ satisfy the sum property of γ, i.e.

$$W_1 + \cdots + W_\gamma \in \mathbb{N}^n.$$

Using this and (11), we have that

$$W_1 + \cdots + W_\gamma \in \frac{\gamma k}{p^{e-\lambda}}\Delta(f) \cap (q-1)\mathbb{N}^n, \tag{12}$$

which means that $\omega_f \leq \frac{\gamma k}{p^{e-\lambda}}$. Reorganizing this, and using our assumption on k at the beginning of the proof, we have $\frac{p^{e-\lambda}}{\gamma}\omega_f \leq k < \omega_f$, or $p^{e-\lambda} < \gamma$. To make use of this inequality, we have the following lemma:

Lemma 8.3. *For all integers $1 \leq j \leq m$, we have*

$$\gamma - 1 \leq n - \ell_j.$$

Proof. Let W_u be such that $\ell_u \geq \ell_j$ for all $1 \leq j \leq m$. If $\ell_u = n$, then $\gamma = 1$ and we are done. If not, W_u has $n - \ell_u$ components which are zero and we can pick elements from $\{W_1, ..., W_{u-1}, W_{u+1}, ..., W_m\}$ to add to W_u until the resulting sum is an element of $\mathbb{N}^n$. This implies it is possible to pick $n-\ell_u+1$ vectors from $\{W_1, ..., W_m\}$ whose sum will lie in $\mathbb{N}^n$. By the definition of γ, we must have $\gamma \leq n - \ell_u + 1$. But by our assumption on W_u, this means that $\gamma - 1 \leq n - \ell_j$ for all j. □

With the help of Lemma 8.3 and (12), we have that $p^{e-\lambda} \leq \gamma - 1 \leq n - \ell_j$. If we can show that

$$v_p\left(p^{\lambda} q^{p^{e-\lambda}}\right) \geq v_p(pk),$$

then $S_k(f) \equiv 0 \bmod (pk)$ and we are done. In other words, if $r = e - \lambda$ we must show

$$p^r v_p(q) \geq r + 1. \tag{13}$$

Fortunately, this is true for all primes p and all positive integers r.

□

9. Analysis of Cardinality Bounds

Our main bound proven in Section 8 is sharp. Let $N(x_1, ..., x_{n-1})$ be the field norm of $\mathbb{F}_{q^{n-1}}$ over $\mathbb{F}_q$. Kosters [7] illustrates that Theorem 1.4 is sharp using the map $f(x_1, ..., x_n) = (x_1, x_2, ..., N(x_1, ..., x_{n-1})x_n)$. Based on this example, we give the following sharp example for Theorem 8.1. Let $h(x_1, x_2, ..., x_n) = (x_1, x_2, ..., N(x_1, ..., x_{n-1})^a x_n)$ with a in $\mathbb{N}$. Because $N(x_1, ..., x_{n-1})$ is a polynomial containing the monomials $x_1^{n-1}, ..., x_{n-1}^{n-1}$ with nonzero coefficients, we have that $(a, a, ..., a, 1) \in \Delta(h)$. This explicitly tells us $\Delta(h) \cap \mathbb{N}^n \neq \varnothing$. We also have for all $V = (v_1, ..., v_n) \in \Delta(h) \cap \mathbb{N}^n$, $v_n = 1$. This implies $\omega_h = q-1$. In addition, since the preimage $N^{-1}(0) = \{(0, ..., 0)\}$, we are given $|V_h| = q^n - (q-1)$. This example highlights the flexibility granted by the use of constants derived from the Newton polytope, since deg $h = a(n-1)+1$ does not allow for a sharp cardinality bound. This flexibility also

gives us more freedom to make substitutions when generating more sharp examples. If $z_1(x), ..., z_{n-1}(x)$ are univariate permutation polynomials in $\mathbb{F}_q[x]$, then the maps $g(x_1, ..., x_n) = (z_1(x_1), ..., z_{n-1}(x_{n-1}), N(x_1, ..., x_{n-1})^a x_n)$ and $h(z_1(x_1), ..., z_{n-1}(x_{n-1}), x_n)$ will share the same constants and value set cardinality as $h(x_1, ..., x_n)$.

Using the constant ω_f also has an advantage when determining bounds on univariate value sets. In this case, since $n = 1$, we have that $\mu_f = \frac{1}{\deg f}$ for all $f \in \mathbb{F}_q[x]$, but Zan and Cao [15] give a sharp example which improves upon this for ω_f. If $f(x) = x^7 + ax \in \mathbb{F}_{19}[x]$ with $a \neq 0, 4, 5, 8, 16, 17$, then it is easy to check that $\omega_f = 6$, $|V_f| = 13 = 19 - \omega_f < 19 - \left\lceil \frac{1}{7}(18) \right\rceil = 16$.

Note that, in general, it is not immediately clear how large of an improvement the strongest bound in Theorem 8.1 provides over our bounds in Theorem 1.4. The first author of the present paper has addressed in [10] that an effective method for calculating μ_f is not directly clear from the definitions given. However, calculation of ω_f should be much more efficient complexity-wise, since only a finite amount of values need to be checked to determine the minimum value. This quantity of values to check by brute force grows with complexity $O(q^n)$ and is therefore polynomial in q (though exponential in n). Therefore, there is much value in the use of ω_f even when it is equal to $\mu_f \cdot (q - 1)$.

10. Future Work

It is important to consider whether results such as in Section 8 apply in more general settings. For instance, there are cases when it is more convenient to use rational interpolated form of a map than its polynomial form, especially when the monomials of the rational interpolation have much smaller degree. Even if we strictly considered Laurent polynomials, where we have $f(x) \in \mathbb{F}_q[x, x^{-1}]$ or the Laurent polynomial map $f(x_1, ..., x_n) = (f_1(x_1, ..., x_n), ..., f_n(x_1, ..., x_n))$ with multivariate polynomial $f_i(x_1, ..., x_n) \in \mathbb{F}_q[x_1, x_2, ..., x_n, x_1^{-1}, x_2^{-1}, ..., x_n^{-1}]$, can we apply the geometry of the Newton polytope to bound their cardinalities? Would such bounds be any stronger than those obtained by using a polynomial-interpolated form of the map?

References

[1] A. Adolphson and S. Sperber. p-adic estimates for exponential sums and the theorem of Chevalley-Warning. *Ann. Sci. Ecole Norm. S.*, 20(4):545–556, 1987.

[2] W. Cao. Smith normal form of augmented degree matrix and its applications. *Linear Algebra Appl.*, 431(10):1778 – 1784, 2009.

[3] W. Cao. Dilation of Newton polytope and p-adic estimate. *Discrete Comput. Geom.*, 45(3):522–528, 2011.

[4] W. Cao and Q. Sun. On a class of equations with special degrees over finite fields. *Acta Arith.*, 130:195–202, 2007.

[5] J. Chen and W. Cao. Degree matrices and divisibility of exponential sums over finite fields. *Arch. Math.*, 94(5):435–441, 2010.

[6] J. Hill. Weil image sums (and some related problems), http://untruth.org/s/p9.html. 2011.

[7] M. Kosters. Polynomial maps on vector spaces over a finite field. *Finite Fields Th. App.*, 31(0):1 – 7, 2015.

[8] G. L. Mullen, D. Wan, and Q. Wang. Value sets of polynomial maps over finite fields. *Q. J. Math.*, pages 1191 – 1196, 2012.

[9] D. Singmaster. Divisibility of binomial and multinomial coefficients by primes and prime powers. In J. V. E. Hoggatt and M. Bicknell-Johnson, editors, *A Collection of Manuscripts Related to the Fibonacci Sequence, 18th Anniversary Volume*, pages 98–114. 1980.

[10] L. Smith. Polytope bounds on multivariate value sets. *Finite Fields Th. App.*, 28(0):132 – 139, 2014.

[11] G. Turnwald. A new criterion for permutation polynomials. *Finite Fields Th. App.*, 1(1):64 – 82, 1995.

[12] D. Wan. Variation of p-adic Newton polygons for L-functions of exponential sums. *Asian J. Math.*, 8(3):427–472, 09 2004.

[13] D. Wan. Lectures on zeta functions over finite fields. In D. Kaledin and Y. Tschinkel, editors, *Higher-Dimensional Geometry over Finite Fields*, pages 244–268. IOS Press, 2008.

[14] D. Wan, P. J.-S. Shiue, and C. S. Chen. Value sets of polynomials over finite fields. *P. Am. Math. Soc.*, 119(3):711–717, 1993.

[15] H. Zan and W. Cao. Powers of polynomials and bounds of value sets. *J. Number Theory*, 143:286 – 292, 2014.