Visibly irreducible polynomials over finite fields

Evan M. O'Dorney

August 31, 2018

Abstract

Lenstra, in this MONTHLY, has pointed out that a cubic over $\mathbb{F}_5 = \mathbb{Z}/5\mathbb{Z}$ of the form $(x-a)(x-b)(x-c) + \lambda(x-d)(x-e)$, where $\{a,b,c,d,e\}$ is some permutation of $\{0,1,2,3,4\}$, is irreducible because every element of $\mathbb{F}_5$ is a root of one summand but not the other. We classify polynomials over finite fields that admit an irreducibility proof with this structure.

1 Introduction.

In a past note in this MONTHLY [5], Lenstra relates how he was trying to set an examination problem of a standard genre—namely, factoring a polynomial over a finite field—whose answer could be verified by a quick, humanly comprehensible argument. He chose the following polynomial:

$$f(x) = x^3 - 3x^2 - x - 3 \in \mathbb{F}_5[x].$$

(Here and throughout this article, $\mathbb{F}_q$ denotes the field with q elements.) Built in was the following solution:

$$f(x) = (x^3 - x) - (3x^2 + 3) = x(x+1)(x-1) - 3(x+2)(x-2), \tag{1}$$

which shows that f is in fact irreducible: for if it factored, it would have to have a linear factor, and each of the five possible linear factors over $\mathbb{F}_5$ divides one but not the other of the two summands of (1). The same proof applies to any polynomial over $\mathbb{F}_5$ of the form

$$(x-a)(x-b)(x-c) + \lambda(x-d)(x-e), \tag{2}$$

where $\{a,b,c,d,e\}$ are the elements $\{0,1,2,3,4\} = \mathbb{F}_5$ in some order and $\lambda \in \mathbb{F}_5^\times$ is a nonzero constant. Lenstra proves that every monic irreducible cubic over $\mathbb{F}_5$ has the form (2) in a unique way, up to permuting the factors in each term, and gives a pleasant algorithm for finding a, b, c, d, e, and λ.

In this article we address the natural question ([5], p. 818) of the extent to which this phenomenon extends to other degrees of polynomials and other fields.

We can say at once that the phenomenon is not restricted to $\mathbb{F}_5$. Taking $\mathbb{F}_2$, the simplest of all fields, and writing the quadratic polynomial

$$f(x) = x^2 + x + 1 = (x)^2 + (x+1),$$

we see that f is irreducible, because the two possible linear factors $x, x+1$ each divide one but not the other of the two terms of the decomposition. This is not the only such irreducibility proof for this polynomial: equally effective are

$$f(x) = (x^2 + 1) + x = (x+1)^2 + x$$

and

$$f(x) = (x^2 + x) + 1 = x(x+1) + 1.$$

The same argument applies to cubics over $\mathbb{F}_2$ such as

$$f(x) = x^3 + x + 1 = (x)^3 + (x+1) = x(x+1)^2 + 1 = \cdots.$$

But if we write a quartic in a form like

$$f(x) = x^4 + x + 1 = (x)^4 + (x + 1),$$

the irreducibility is no longer clear. We see that f has no linear factor, but a quartic polynomial could still factor as the product of two quadratics. If we know (somehow) that $x^2 + x + 1$ is the only irreducible quadratic over $\mathbb{F}_2$, then we can write

$$f(x) = x^4 + x + 1 = x(x + 1)(x^2 + x + 1) + 1,$$

and now it *is* visible that f is not divisible by either of the linear factors x, $x + 1$ or the quadratic factor $x^2 + x + 1$, and hence f must be irreducible.

Motivated by the foregoing examples, we make the following definition:

Definition. A *visibly irreducible decomposition (VID)* of degree $d \geq 2$ over the finite field $\mathbb{F}_q$ is a sum $f_1(x) + f_2(x) + \cdots + f_r(x) = f(x)$ of $r \geq 2$ nonzero polynomials $f_i(x) \in \mathbb{F}_q[x]$ of degree at most d with the following properties:

(VID-1) Every irreducible polynomial $p(x)$ of degree not exceeding $d/2$ (known as an *operative factor*) divides all but exactly one of the $f_i(x)$. This makes it visible that $p(x) \nmid f(x)$.

(VID-2) Exactly one of the f_i actually has degree d, the others having degree less than d. This makes it visible that f has degree exactly d.

Condition (VID-2) may seem a bit arbitrary, but it ensures that the sum $f(x)$ has degree exactly d, without the need to check the sum of the leading coefficients. Without it, $f(x)$ could be an irreducible polynomial of any degree from $\lfloor d/2 \rfloor + 1$ to d inclusive—or could be a constant! Condition (VID-2) is also motivated by symmetry considerations, as will be explained in Section 3. At the end of this article we will briefly note what happens if it is removed.

Our main result is the following determination of which polynomials admit a VID.

Theorem 1.

(a) *For the following pairs (q, d),* EVERY *irreducible polynomial of degree d over $\mathbb{F}_q$ admits a VID:*

- $(2, 2)$, $(3, 2)$
- $(2, 3)$, $(3, 3)$, $(4, 3)$, $(5, 3)$
- $(2, 4)$
- $(2, 5)$
- $(2, 6)$
- $(2, 7)$.

(b) *For $(q, d) = (3, 5)$, exactly* HALF *of all irreducible quintics over $\mathbb{F}_3$ admit a VID.*

(c) *For all other q and d,* NO *irreducible polynomial admits a VID.*

2 No VID's for large fields or high degrees.

We begin with the proof of Theorem 1(c), which restricts the (q, d) pairs to be considered to a finite list. The method is quite straightforward.

Lemma 1. *If a VID $f_1 + \cdots + f_r$ of degree d exists over $\mathbb{F}_q$, then*

$$dr \geq (r - 1) \left[1 + \left| \bigcup_{n=1}^{\lfloor d/2 \rfloor} \mathbb{F}_{q^n} \right| \right], \tag{3}$$

where the union is taken within the algebraic closure $\overline{\mathbb{F}_q}$ (which contains a unique isomorphic copy of $\mathbb{F}_{q^n}$ for each n).

Proof. Let $\xi \in \mathbb{F}_{q^n}$, $1 \leq n \leq \lfloor d/2 \rfloor$. The minimal polynomial $p(x)$ of ξ is irreducible of degree $n \leq \lfloor d/2 \rfloor$ and thus must divide all but one of the f_i; therefore ξ is a root of the product $f_1 f_2 \cdots f_r$ of multiplicity at least $r - 1$. But this is a product of one factor of degree d and $r - 1$ factors of degree at most $d - 1$, so

$$d + (d-1)(r-1) \geq (r-1) \left| \bigcup_{n=1}^{\lfloor d/2 \rfloor} \mathbb{F}_{q^n} \right|,$$

which simplifies to (3). $\qquad\square$

Lemma 2. *The bound* (3) *can hold only for the pairs* (q, d) *mentioned in Theorem 1*(a),(b).

Proof. The bound is weakest when $r = 2$, so it suffices to determine when it can hold in this case. We have

$$2d \geq 1 + \left| \bigcup_{n=1}^{\lfloor d/2 \rfloor} \mathbb{F}_{q^n} \right| \geq 1 + \left| \mathbb{F}_{q^{\lfloor d/2 \rfloor}} \right| = 1 + q^{\lfloor d/2 \rfloor}. \tag{4}$$

In particular,

$$2d \geq 1 + 2^{\lfloor d/2 \rfloor}$$

which is seen to hold only when $d \leq 7$ or $d = 9$. But the $d = 9$ case, upon substituting back into (3), yields

$$18 = 2d \geq 1 + \left| \mathbb{F}_{q^3} \cup \mathbb{F}_{q^4} \right| = 1 + q^3 + q^4 - q \geq 23,$$

which is untrue. So we have $2 \leq d \leq 7$. For each d, (4) bounds the value of q by

$$q \leq \left\lfloor (2d - 1)^{1/\lfloor d/2 \rfloor} \right\rfloor.$$

Plugging $d = 2, 3, 4, 5, 6, 7$ into this fanciful-looking expression yields the bounds of 3, 5, 2, 3, 2, and 2 respectively, precisely as desired. $\qquad\square$

3   Symmetry.

Proving Theorem 1(a) is a finite problem: for each (q, d), there are a finite number of irreducibles, and we simply need to write a VID for each one! However, throughout this article, we will strive to prove results conceptually rather than resorting to computation. In this section, we describe a family of symmetries that allow us to consider only a small number of irreducibles per (q, d) pair.

The symmetries are best described in terms of *homogeneous forms* of degree d in two variables X, Y. These are in bijection with one-variable polynomials of degree at most d, via the standard operations of *homogenization*

$$f(x) \mapsto F(X, Y) = Y^d f(X/Y)$$

and *dehomogenization*

$$F(X, Y) \mapsto f(x) = F(x, 1),$$

and we will frequently identify one-variable polynomials with their homogenizations.

In the homogeneous context, we have the following attractive notion of VID:

Definition. A *(homogeneous) VID* of degree d over a finite field $\mathbb{F}_q$ is a sum $F_1(X, Y) + F_2(X, Y) + \cdots + F_r(X, Y) = F(X, Y)$ of $r \geq 2$ nonzero homogeneous forms of degree d over $\mathbb{F}_q$, satisfying a single property:

(HVID) Every irreducible homogeneous form $P(X, Y)$ of degree not exceeding $d/2$ (called an *operative factor*) divides all but exactly one of the $F_i(x)$. This makes it visible that $P(X, Y) \nmid F(X, Y)$.

We see that (HVID) for each operative factor $P(X, Y)$ corresponds to (VID-1) for the corresponding inhomogeneous operative factor $p(x)$, *except* for the special operative factor $P(X, Y) = Y$, for which (HVID) corresponds to (VID-2). Thus this notion of VID is entirely compatible with the one above.

Now, the group $\mathrm{GL}_2(\mathbb{F}_q)$ acts on homogeneous forms of degree d by linear change of variables

$$\begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix} \cdot F(X,Y) = F(\alpha X + \gamma Y, \beta X + \delta Y).$$

The scalar matrix αI acts by multiplication by α^d, and thus the quotient $\Gamma = \mathrm{PGL}_2(\mathbb{F}_q)$ acts on the set of forms of degree d up to scaling. Moreover, the Γ-action preserves irreducibility and acts on VID's of each degree up to scaling (where *scaling* a VID means scaling all its summands F_i by a single scalar $\alpha \in \mathbb{F}_q^\times$).

Though we will not need it, the Γ-action can be described directly on inhomogeneous polynomials as

$$\begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix} \cdot f(x) = (\beta x + \delta)^d f\left(\frac{\alpha x + \gamma}{\beta x + \delta} \right);$$

in such form it was studied in [4].

Let $\mathcal{I}(q,d)$ be the set of irreducible homogeneous forms of degree d over $\mathbb{F}_q$, up to scaling. These are in bijection with the monic irreducible one-variable polynomials of degree d if $d \geq 2$. The size of $\mathcal{I}(q,d)$ is given by the classical formula (due to Gauss in the case q prime; see [3] for a simple proof in the general case):

$$|\mathcal{I}(q,d)| = \begin{cases} \dfrac{1}{d} \sum_{k \mid d} \mu(k) q^{d/k}, & d \geq 2, \\ q+1, & d = 1 \end{cases} \tag{5}$$

where $\mu(k)$ is the Möbius function. If one $F \in \mathcal{I}(q,d)$ admits a VID, then so do all irreducibles in the Γ-orbit of F. Therefore we will begin by counting the Γ-orbits on $\mathcal{I}(q,d)$. We begin with a pair of simple results.

Lemma 3.

(a) *The group $\mathrm{GA}_1(\mathbb{F}_q)$ of affine transformations of the line (which is also the subgroup of transformations in Γ fixing one linear form $Y \in \mathcal{I}(q,1)$) acts simply transitively on $\mathbb{F}_{q^2} \backslash \mathbb{F}_q$.*

(b) *Γ acts simply transitively (by linear fractional transformations) on $\mathbb{F}_{q^3} \backslash \mathbb{F}_q$.*

Proof. The proof method in each case is similar:

(a) Since $|\mathrm{GA}_1(\mathbb{F}_q)| = q(q-1) = |\mathbb{F}_{q^2} \backslash \mathbb{F}_q|$, it is enough to prove that the stabilizer of each point is trivial. Suppose $\gamma \in \mathrm{GA}_1(\mathbb{F}_q)$ fixes $\xi \in \mathbb{F}_{q^2} \backslash \mathbb{F}_q$. Then γ also fixes $\tau(\xi)$, where $\tau \in \mathrm{Gal}(\mathbb{F}_{q^2}/\mathbb{F}_q)$ is a generator. Since $\xi \notin \mathbb{F}_q$, we have $\xi \neq \tau(\xi)$. An affine transformation that fixes two points must be the identity.

(b) Since $|\Gamma| = q(q-1)(q+1) = |\mathbb{F}_{q^3} \backslash \mathbb{F}_q|$, it is enough to prove that the stabilizer of each point is trivial. Suppose $\gamma \in \Gamma$ fixes $\xi \in \mathbb{F}_{q^3} \backslash \mathbb{F}_q$. Then γ also fixes $\tau(\xi)$ and $\tau^2(\xi)$, where $\tau \in \mathrm{Gal}(\mathbb{F}_{q^3}/\mathbb{F}_q)$ is a generator. Since $\xi \notin \mathbb{F}_q$, the three conjugates $\xi, \tau(\xi), \tau^2(\xi)$ are distinct. A linear fractional transformation that fixes three points must be the identity. $\square$

Lemma 4. *The values of (q,d) for which $\mathcal{I}(q,d)$ consists of a single Γ-orbit are as follows: all $d \leq 3$, and $(2,4)$ and $(2,5)$.*

The Γ-action on $\mathcal{I}(q,d)$ is well studied: formulas have been published for the number of fixed points of various elements and subgroups of Γ [1, 2, 4]. Nevertheless, no one in the literature seems to have posed before the simple question of when $\mathcal{I}(q,d)$ is a single Γ-orbit.

Proof of Lemma 4. By Lemma 3, Γ transitively permutes the elements of $\mathbb{F}_{q^2} \backslash \mathbb{F}_q$ (respectively, $\mathbb{F}_{q^3} \backslash \mathbb{F}_q$) and thus also transitively permutes their minimal polynomials, which comprise $\mathcal{I}(q,2)$ (respectively, $\mathcal{I}(q,3)$). Here we are using that the minimal polynomial of $\xi \in \mathbb{F}_{q^d}$ is, upon homogenization, the lowest-degree form defined over $\mathbb{F}_q$ divisible by $X - \xi Y$; and Γ acts on these linear forms up to scaling as it does on the elements $\xi \in \mathbb{F}_{q^d} \cup \{\infty\}$ via linear fractional transformations. (The matrix

$$g = \begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix}$$

4

does not act by $\xi \mapsto \frac{\alpha\xi+\beta}{\gamma\xi+\delta}$, as is natural, but by $\xi \mapsto \frac{\delta\xi-\gamma}{-\beta\xi+\alpha}$; but this is of no significance for the orbits.) This takes care of the cases where $d \leq 3$.

In the cases $(q, d) = (2, 4)$ and $(2, 5)$, we can prove the lemma by bounding the point stabilizers in the following standard way, which will also be important later:

Lemma 5. *If $d \geq 3$, then for every $F(X, Y) \in \mathcal{I}(q, d)$, the stabilizer Γ_F is a cyclic group of order dividing d.*

Proof. Let $f(x)$ be the dehomogenization of F. The stabilizer $\Gamma_F = \Gamma_f$ permutes the roots of f in $\mathbb{F}_{q^d}$ and thus maps naturally into $\mathrm{Gal}(\mathbb{F}_{q^d}/\mathbb{F}_q) \cong C_d$, a cyclic group. Any $\gamma \in \Gamma_f$ that fixes one root must fix all the roots, and since there are at least three roots, we must in fact have $\gamma = I$. So Γ_f maps isomorphically to a subgroup of C_d. $\qquad\square$

In the cases $(q, d) = (2, 4)$ and $(2, 5)$, we get that all orbits have size

$$\frac{|\Gamma|}{|\Gamma_f|} \geq \frac{|\Gamma|}{\gcd(|\Gamma|, d)} = \begin{cases} 3, & d = 4 \\ 6, & d = 5. \end{cases}$$

By Gauss's formula (5), there are exactly 3 irreducible quartics and 6 irreducible quintics over $\mathbb{F}_2$, implying that there is only a single orbit in these cases.

For all other (q, d), there is *more* than one orbit. For $d \geq 5$ this can be seen simply by showing that

$$|\mathcal{I}(q, d)| > |\Gamma|,$$

an exercise in bounding. For $d = 4$, we claim that the point stabilizer of any irreducible $f \in \mathcal{I}(q, 4)$ has order at least 2. Consider the permutation of the roots of f given by the square of Frobenius: $\tau(\xi) = \xi^{q^2}$. As this permutation is in the Klein four group, it preserves the cross ratio of the four roots (the unique invariant

$$\frac{(\xi_1 - \xi_2)(\xi_3 - \xi_4)}{(\xi_1 - \xi_3)(\xi_2 - \xi_4)}$$

of quadruples of distinct points on the projective line) and thus is given by a linear fractional transformation, which, being characterized in a Galois-invariant way, is defined over $\mathbb{F}_q$. So after verifying the weaker inequality

$$|\mathcal{I}(q, 4)| > \frac{|\Gamma|}{2}$$

for $q \geq 3$, it follows that there is more than one orbit. $\qquad\square$

4 Construction of VID's.

Writing a VID of a given degree is a quite intuitive matter: we place each operative factor in the appropriate summands and repeat factors (or, in rare cases, add higher-degree factors) to bring the total degree of each term up to d. For instance, in the case $(q, d) = (4, 3)$, $r = 2$, the operative factors are five linear forms $L_1, \ldots, L_5$. Because the degree of each term cannot exceed 3, they must appear in the distribution

$$L_1 L_2 + L_3 L_4 L_5;$$

then, to bring the degree of the first term up to 3, we add another factor of L_1 or L_2 (not L_3, L_4, or L_5!) to get the VID shape

$$L_1^2 L_2 + \alpha L_3 L_4 L_5,$$

where the relative scaling α as well as the ordering of the forms L_i can freely vary. For clarity's sake we include a formal exposition.

Definition. A *shape* of degree d is a sum

$$\mathfrak{S} = \sum_{i=1}^{r} \left(\alpha_i \prod_{j=1}^{s_i} \mathfrak{P}_{ij} \right)$$

of products of formal factors $\mathfrak{P}_{ij}$, with two attached pieces of data:

(a) a positive integer $\deg \mathfrak{P}_{ij}$ for each factor, to be thought of as a degree, with each summand having total degree $\sum_j \deg \mathfrak{P}_{ij} = d$;

(b) an equivalence relation $\mathfrak{P}_{ij} \equiv \mathfrak{P}_{k\ell}$ among the factors, respecting degree (that is, such that two equivalent factors have the same degree).

Definition. An *instance* of a shape over a field $\mathbb{F}$ is an actual sum $\sum_i F_i$ given by replacing each formal factor $\mathfrak{P}_{ij}$ with an actual homogeneous polynomial of the specified degree over $\mathbb{F}$, so that equivalent factors get replaced by the same (or proportional) forms and inequivalent factors by nonproportional forms, and specifying the relative scalings α_i of the terms. Rescaling the entire sum, or fiddling with the scalings of each factor without changing the overall relative scalings of each term, will be considered to yield the same instance. Permuting the terms F_i will also be considered to yield the same instance, if the shape happens to be invariant under some such permutation.

Definition. A *visibly irreducible shape (VIS)* of degree d over $\mathbb{F}_q$ is a shape of degree d in which, for $1 \leq n \leq d/2$, there are exactly $|\mathcal{I}(q,d)|$ inequivalent factors of degree n and each appears in all but one summand of the shape.

These definitions have been arranged to make it obvious that every VID is an instance of a unique VIS, and every instance of a VIS is a VID. We now proceed with the construction.

4.1 The one-orbit cases.

The cases (q,d) where $\mathcal{I}(q,d)$ is a single Γ-orbit are the simplest to analyze. One simply has to write a single VIS $\mathfrak{S}$; then its instances (being a Γ-invariant set) represent all irreducibles in $\mathcal{I}(q,d)$. Moreover, each irreducible is represented the same number of times, which may readily be computed by dividing the number of instances of $\mathfrak{S}$ by $|\mathcal{I}(q,d)|$.

Considerations of space prevent us from classifying *all* VIS's, though such a classification is certainly within reach; we limit ourselves to listing one VIS per (q,d) pair. We write shapes as follows: L, Q, and C with possible subscripts denote linear, quadratic, and cubic factors respectively, those with different subscripts being inequivalent. The formal coefficient α_i of one term can be suppressed, and all the α_i can be suppressed if $q = 2$.

q	d	r	Example VIS	$\lvert\mathcal{I}(q,d)\rvert$	# of VID's of this shape per irred
2	2	$2,3$	$L_1L_2 + L_2L_3 + L_3L_1$	1	1
2	3	any	$L_1^2L_2 + L_2^2L_3 + L_3^2L_1$	2	1
2	4	$2,3,4$	$L_1^2L_2L_3 + Q^2$	3	1
2	5	any	$L_1^4L_2 + L_3Q^2$	6	1
3	2	2	$L_1L_2 + \alpha L_3L_4$	3	2
3	3	$2,3,4$	$L_1^3 + \alpha L_2L_3L_4$	8	1
4	3	2	$L_1^2L_2 + \alpha L_3L_4L_5$	20	3
5	3	2	$L_1L_2L_3 + \alpha L_4L_5L_6$	40	1

Table 1: VID's of irreducibles in the cases where there is a single Γ-orbit.

The value of r, the number of terms, is constrained by Lemma 1. For two (q,d) pairs, namely $(2,3)$ and $(2,5)$, all values of $r \geq 2$ are admissible, and for a striking reason: there is a single term $T = L_1L_2L_3$

(respectively, $T = L_1L_2L_3Q$) that is divisible by all the operative factors, and hence T can be tacked on to a VIS any number of times without affecting visible irreducibility! Using T, we can also concoct VIS's such as

$$L_1L_2L_3 + C, \tag{6}$$

expressing translation symmetries of the sets $\mathcal{I}(2,3)$ and $\mathcal{I}(2,5)$. It is a matter of taste whether an expression like (6) is truly *visibly* irreducible, insomuch as the irreducibility of the sum rests on the irreducibility of a term C of the same degree! Fortunately, this point is of little consequence for us, since every polynomial admitting a VID will turn out to have one like those in Table 1, with each summand involving powers of the operative factors only.

In all other cases, $r \leq 4$. In Table 1, we have chosen neither the longest nor the shortest VIS but rather the most symmetrical, minimizing the number of distinct instances and thus minimizing the number of VID's of that shape per irreducible, shown in the last column of the table. In six cases, indeed, we can make the VID unique.

The VIS $L_1L_2 + L_2L_3 + L_3L_1$ tabulated for quadratics over $\mathbb{F}_2$ is more symmetric than the shape $L_1^2 + L_2L_3$ discovered above and is the first instance of a *visibly rootless Lagrange interpolation.* Recall that Lagrange interpolation is a general method for computing a polynomial of minimal degree attaining specified values at an arbitrary finite list of points by summing polynomials that vanish at all but one of the given points. In the present context, it is easy to see that the products

$$L_1 \cdots \hat{L}_i \cdots L_{q+1},$$

consisting of all but one linear form, form a basis for the homogeneous forms of degree $d = q$ over $\mathbb{F}_q$. If a form F of degree q has no roots over $\mathbb{F}_q$, then each basis element has nonzero coefficient, and the sum is a *visibly rootless* expansion of F. After $q = 2$, the next case $q = 3$ yields the VIS

$$L_1L_2L_3 + \alpha L_1L_2L_4 + \beta L_1L_3L_4 + \gamma L_2L_3L_4,$$

of maximal length $r = 4$, which represents the 8 irreducible cubics over $\mathbb{F}_3$ just by varying the signs $\alpha, \beta, \gamma \in \mathbb{F}_3^\times = \{\pm 1\}$. For degree $d = 4$ onward, rootless polynomials are no longer necessarily irreducible. Rootlessness is a less deep notion than irreducibility, and visibly rootless expansions are easily shown to exist for all polynomials provided that d is large compared to q.

4.2 Corollaries using selected VID's.

The cases $(q, d) = (2, 4)$ and $(q, d) = (3, 3)$ of Table 1 are also noteworthy in yielding the following novel relations among irreducible polynomials in those degrees.

Corollary 1. *The three nonzero linear forms L_i and the three irreducible quartic forms D_i over $\mathbb{F}_2$ are in a canonical bijection respecting the Γ-action, given by*

$$L_1 \mapsto D_1 = L_1^2L_2L_3 + Q^2 = L_1^4 + L_2L_3Q = L_1^2Q + L_2^2L_3^2.$$

Proof. It is easy to see that there is only one way for the symmetric group $\Gamma = S_3$ to act transitively on a set of size 3, up to isomorphism, namely its natural action on three letters. The three letters can be distinguished by their stabilizers, which are the three order-2 subgroups of S_3.

Consequently, the three L_i and the three D_i comprise isomorphic Γ-sets. Each of the three VID's listed is invariant under swapping L_2 with L_3 and thus must represent the unique irreducible quartic D_1 fixed by this transposition. $\square$

Corollary 2. *Let C be an irreducible cubic form over $\mathbb{F}_3$. There is a unique irreducible cubic form C' over $\mathbb{F}_3$ such that*

- $C + C'$ *is a cube,*

- $C - C'$ *is the product of three distinct linear factors.*

Proof. By Table 1, C can be uniquely decomposed as $L_1^3 + L_2L_3L_4$. (Here we are noting that every element of $\mathbb{F}_3$ is a cube, so the term L_1^3, which was a priori only a cube up to scaling, is in fact the cube of a linear form L_1; and we scale L_2, L_3, and L_4 so that the equality holds.) We see that $C' = L_1^3 - L_2L_3L_4$ satisfies the conditions. Conversely, for any C' satisfying the conditions,

$$C = \frac{C + C'}{2} + \frac{C - C'}{2}$$

is a VID of C of the shape $L_1^3 + L_2L_3L_4$, the two summands clearly being coprime. $\qquad\square$

4.3 Sextics over $\mathbb{F}_2$.

There remain two cases of Theorem 1(a) for which there are multiple Γ-orbits. These require more work, for instead of merely displaying one VIS, we must find a VIS for each orbit and carefully verify that they indeed cover all the orbits.

The $(2^6 - 2^3 - 2^2 + 2)/6 = 9$ irreducible sextics over $\mathbb{F}_2$ are a case in point. Three of these form a *special orbit* with point stabilizer of size 2, represented by the self-reciprocal polynomial $x^6 + x^3 + 1$. The other six form a *generic orbit* with trivial point stabilizer. (There are many ways to verify the sizes and stabilizers of these orbits.) The VIS

$$F_{L_1} = L_1^2 L_2 L_3 Q + C_1 C_2$$

is symmetric under swapping L_2 with L_3 or C_1 with C_2; indeed, it has exactly three instances, which must represent the irreducibles in the special orbit. One finds that there is just one other VIS in this degree,

$$F_{L_1, L_2, L_3, C_1, C_2} = L_1^2 L_2 C_1 + L_3 Q C_2. \tag{7}$$

It is completely asymmetric and yields 12 instances, some of which must necessarily have the same sum. But who is to say, except by explicit computation, that they are not just additional VID's for the special orbit?

To shed more light on this question, recall Lenstra's proof [5] that the unique VIS for $(q, d) = (5, 3)$—appearing in the last row of Table 1—represents all irreducible cubics over $\mathbb{F}_5$. His method is quite different from ours: after observing that there are the same number of irreducible cubics as VID's of this shape, he shows directly that no two of the VID's have the same value. Suppose that

$$L_1 L_2 L_3 + \alpha L_4 L_5 L_6 = \beta L_1' L_2' L_3' + \gamma L_4' L_5' L_6' \tag{8}$$

for some constants $\alpha, \beta, \gamma \in \mathbb{F}_5^\times$ and permutation $\{L_1', \ldots, L_6'\}$ of $\{L_1, \ldots, L_6\}$. Then observe that some pair of terms, one on each side of (8), must share at least two linear factors. Assume they share exactly two (the other case is trivial): we can reindex so that the relation (8) takes the form

$$L_1 L_2 L_3 + \alpha L_4 L_5 L_6 = \beta L_1 L_2 L_4 + \gamma L_3 L_5 L_6$$
$$L_1 L_2 (L_3 - \beta L_4) = L_5 L_6 (\gamma L_3 - \alpha L_4). \tag{9}$$

Now the common value of the two sides is a cubic form divisible by four distinct linear forms L_1, L_2, L_5, L_6, which is impossible. At the heart of the proof is a "compare and factor" technique (9), by which two similar sums are subtracted term by term and proved to be unequal. This "compare and factor" method will enable us to avoid brute-force computation of orbit representatives and their VID's.

We are now ready to prove the $(q, d) = (2, 6)$ case of Theorem 1. In fact, we have the following:

Theorem 2. *The asymmetric shape* (7) *represents every irreducible sextic over* $\mathbb{F}_2$.

Proof. We ask whether a VID of the asymmetric shape $F_{L_1, L_2, L_3, C_1, C_2}$ can equal one of the symmetric shape F_{L_1}. In fact, it can: in an application of the "compare and factor" method, the potential equality

$$F_{L_1, L_2, L_3, C_1, C_2} = F_{L_3}$$
$$L_1^2 L_2 C_1 + L_3 Q C_2 = L_1 L_2 L_3^2 Q + C_1 C_2$$

can be written as

$$L_3 Q (C_2 + L_1 L_2 L_3) = C_1 (L_1^2 L_2 + C_2),$$

which holds if and only if

8

(a) $C_1 = C_2 + L_1 L_2 L_3$, and

(b) $L_3 Q = L_1^2 L_2 + C_2$.

Equation (a) *always* holds (this was noted above in (6)). Equation (b) may or may not hold: $L_1^2 L_2 + L_3 Q$ is a visibly irreducible cubic that may or may not be C_2. The six instances of the asymmetric shape (7) that satisfy (b) form a Γ-orbit that represents the special orbit, each sextic therein occurring twice. As for the six instances that do *not* satisfy (b), we leave it to the reader to apply the "compare and factor" method to eliminate the other possibilities

$$F_{L_1, L_2, L_3, C_1, C_2} = F_{L_1} \quad \text{and} \quad F_{L_1, L_2, L_3, C_1, C_2} = F_{L_2},$$

and to conclude that these instances necessarily represent the generic orbit. As these six instances form a Γ-orbit, we obtain that each irreducible in the generic orbit actually has a unique VID. $\square$

This completes the construction of VID's for the two orbits and the solution of the $(q, d) = (2, 6)$ case of Theorem 1. This may seem like a lot of fuss considering the small number of polynomials involved. But we will now apply the same method to the septimic (7th-degree) case.

4.4   Septimics over $\mathbb{F}_2$.

There are $(2^7 - 2)/7 = 18$ irreducible septimics over $\mathbb{F}_2$. By Lemma 5, all point stabilizers are trivial and there are 3 orbits, each of size 6. Ideally, we would seek a VIS that represents all 18 septimics. There are several VIS's in this degree, but unfortunately, none has more than $3! \cdot 1! \cdot 2! = 12$ instances, due to the limited permutations of the L_i's, Q, and the C_i's. (No VIS includes a factor of degree 4 or greater, as then, even for $r = 2$, the total degree of the two terms would be at least $3(1) + 1(2) + 2(3) + 1(4) = 15 > 7 + 7$.) However, the following VID schema is a union of two shapes that are sufficiently similar to allow the "compare and factor" method to work both within and between them.

Theorem 3. *Any irreducible septimic over $\mathbb{F}_2$ is uniquely of the visibly irreducible schema*

$$L_1^i L_2^{4-i} C_1 + L_3^2 Q C_2 \tag{10}$$

for some orderings L_1, L_2, L_3 and C_1, C_2 of the operative forms and some integer i, $0 < i < 4$, up to the symmetry that takes $i \mapsto 2 - i$ and swaps L_1 with L_2.

Proof. Because there are $\frac{3! \cdot 2! \cdot 3}{2} = 18$ VID's within the schema, it is enough to show that no two have equal sum. Using the "compare and factor" method, we make the following observation: If $F_1, \ldots, F_4$ are quartics (not necessarily irreducible or even distinct) such that

$$F_1 C_1 + F_2 C_2 = F_3 C_1 + F_4 C_2$$

but $F_1 \neq F_3$, then from the factorization

$$(F_1 - F_3) C_1 = (F_4 - F_2) C_2,$$

we get that $F_1 - F_3 = L C_2$ and $F_4 - F_2 = L C_1$ for some $L \in \{L_1, L_2, L_3\}$. In particular, $F_1 - F_3$ is divisible by exactly one L_i and not by Q.

Assume that

$$L_1^i L_2^{4-i} C_1 + L_3^2 Q C_2 = L_1'^j L_2'^{4-j} C_1' + L_3'^2 Q C_2'$$

are two distinct VID's for the same irreducible septimic within the schema (10), where $\{L_1', L_2', L_3'\}$ and $\{C_1', C_2'\}$ are permutations of $\{L_1, L_2, L_3\}$ and $\{C_1, C_2\}$, respectively. If $C_2' = C_2$, then the coefficients F_2, F_4 of C_2 on each side are both divisible by Q, which is impossible by the observation above. So $C_2' = C_1$ and $C_1' = C_2$. But now the difference of the coefficients of C_1 on each side is

$$F_1 - F_3 = L_1^i L_2^{4-i} + L_3'^2 Q.$$

If $L_3' = L_3$, then $F_1 - F_3$ is divisible by none of the L_i. (Indeed, it is a visibly irreducible quartic.) But if L' is one of the other L_i, say L_1, then $F_1 - F_3$ is divisible by both L_1 and L_3 since L_1 divides both terms and L_3 divides neither term. So in no case can $F_1 - F_3$ be divisible by exactly one L_i, completing the proof of the theorem and of Theorem 1(a). $\square$

5   Quintics over $\mathbb{F}_3$.

As promised in Theorem 1, this is the unique case in which we get VID's for some but not all of the irreducibles of one degree over a field. There are $(3^5 - 3)/5 = 48$ irreducible quintics over $\mathbb{F}_3$, up to scaling. The group Γ has size 24 (indeed, it is isomorphic to S_4, permuting the four L_i freely). The point stabilizer of any irreducible quintic is trivial by Lemma 5, so there are two orbits, each of size 24. In writing a VID, we note that equality holds in Lemma 1 with $r = 2$, so we must have just two terms f_1, f_2 such that

$$f_1 f_2 = \alpha L_1 L_2 L_3 L_4 Q_1 Q_2 Q_3$$

for some $\alpha \in \mathbb{F}_3^{\times}$. There is but a single way, up to reindexing, to split the degree-10 polynomial on the right into the product of two quintics: thus there is only a single VIS

$$F_{L_1,Q_1,\alpha} = L_1 Q_2 Q_3 + \alpha L_2 L_3 L_4 Q_1.$$

It has 24 instances (there are 4 choices for L_1, 3 for Q_1, and 2 for α). We conclude that they are the 24 quintics in one orbit. As there are no other VIS's, the 24 quintics in the other orbit do not admit a VID. This completes the proof of Theorem 1.

6   VID's without visible degree.

To return to our starting point, our notion of VID of a one-variable polynomial included a condition on the degrees of the summands (VID-2), at first seemingly arbitrary, but ultimately explained in terms of the corresponding condition on homogeneous polynomials (HVID) respecting their richer Γ-symmetry. This article would be incomplete without a few remarks on what would go differently if (VID-2) were removed.

The proof of Lemma 1 remains unchanged and yields the weaker bound

$$dr \geq (r-1) \left| \bigcup_{n=1}^{\lfloor d/2 \rfloor} \mathbb{F}_{q^n} \right|$$

in which a summand of $r - 1$ is omitted from the right-hand side. Continuing in the manner of Lemma 2 yields the same possible (q, d) pairs, with one addition: $(q, d) = (4, 2)$. A quadratic over $\mathbb{F}_4$ cannot have a VID in the sense used throughout this article, but a sum of the shape

$$L_1 L_2 + \alpha L_3 L_4, \tag{11}$$

omitting the exceptional linear form $L_5(X, Y) = Y$, can be irreducible. Sums of this shape do not have a Γ-action, but there is an action by the stabilizer Γ_∞ of L_5, which is none other than the group $\mathrm{GA}_1(\mathbb{F}_4)$ of affine transformations of $\mathbb{F}_4$. By Lemma 3(a), Γ_∞ permutes $\mathcal{I}(4, 2)$ transitively. There are $|\mathcal{I}(4, 2)| = (4^2 - 4)/2 = 6$ irreducibles. The shape (11) has 9 instances (fixing L_5), of which 3 have the value L_5^2 up to scaling (one choice of α for each choice of L_1, L_2, L_3, L_4). So we have the following uniqueness theorem.

Theorem 4. *An irreducible quadratic $f(x)$ over $\mathbb{F}_4$ can be expressed uniquely in the form*

$$\alpha(x - a)(x - b) + \beta(x - c)(x - d)$$

up to commutativity, where $\{a, b, c, d\} = \mathbb{F}_4$ and $\alpha, \beta \in \mathbb{F}_4^{\times}$ are distinct scalars.

7   Conclusion.

Lenstra's "compare and factor" method, coupled with an awareness of the symmetry of the situation, demonstrate for us that VID's of the same degree tend to "repel" each other and fill out all irreducibles of a given degree. However, the obtainable families of irreducibles peter out after a finite list, and no case comes close to exceeding the $4 \cdot |\mathcal{I}(5, 3)| = 160$ cubics in Lenstra's $\mathbb{F}_5$ example. So the question remains: Is the "compare and factor" method, for all its beauty, applicable only to a finite total number of objects? Or are there structures of higher degree, perhaps even in more dimensions, that can be handled in a subtly analogous way?

8 Acknowledgments.

References

[1] Ahmadi, O. (2011). Generalization of a theorem of Carlitz. *Finite Fields Appl.* 17(5): 473–480.

[2] Carlitz, L. (1967). Some theorems on irreducible reciprocal polynomials over a finite field. *J. reine angew. Math.* 227: 212–220.

[3] Chebolu, S. K., Mináč, J. (2011). Counting irreducible polynomials over finite fields using the inclusion-exclusion principle. *Math. Mag.* 84(5): 369–371.

[4] Garefalakis, T. (2011). On the action of $\mathrm{GL}_2(\mathbb{F}_q)$ on irreducible polynomials over $\mathbb{F}_q$. *J. Pure Appl. Algebra.* 215(8): 1835–1843.

[5] Lenstra, H. (2010). Irreducible cubics modulo five. *Amer. Math. Monthly.* 117(9): 817–821.