ROOTS AND IRREDUCIBLE POLYNOMIALS

KEITH CONRAD

This handout, which accompanies the course on analogies between $\mathbf{Z}$ and $F[T]$, discusses some properties of polynomials in $F[T]$. The results in Sections 1 and 2 work with any F, but the results in Section 3 and 4 are (somewhat) special to the field $F = \mathbf{F}_p$. The main result in these notes in Theorem 3.7.

The notation $F[T]_{h(T)}$, for the ring of polynomials in $F[T]$ considered modulo $h(T)$, is used on the first-year number theory sets. We will instead use a more common notation from abstract algebra, and write $F[T]/h(T)$ for $F[T]_{h(T)}$.

As a matter of terminology, when $f(T) \in F[T]$, we will say $f(T)$ is a polynomial "over" F. For example, if we are thinking about $T^3 + 2T + 1$ as a polynomial in $\mathbf{F}_5[T]$ (as opposed to $\mathbf{F}_2[T]$, $\mathbf{R}[T]$, and so on), we might say "consider $T^3 + 2T + 1$ over $\mathbf{F}_5$."

1. Roots in larger fields

A polynomial in $F[T]$ may not have a root in F. If we are willing to enlarge the field F, then we can discover some roots.

Theorem 1.1. *Let F be a field and $\pi(T)$ be irreducible in $F[T]$. There is a field $E \supset F$ such that $\pi(T)$ has a root in E.*

Proof. Use $E = F[x]/\pi(x)$. It is left to the reader to check the details. □

Example 1.2. Consider $T^2 + 1 \in \mathbf{F}_3[T]$, which has no root in $\mathbf{F}_3$. The ring $\mathbf{F}_3[x]/(x^2 + 1)$ is a field containing $\mathbf{F}_3$. In this field $\overline{x}^2 = -1$, so the polynomial $T^2 + 1$ picks up a root $\overline{x}$ in $\mathbf{F}_3[x]/(x^2 + 1)$. The root $-\overline{x} = 2\overline{x}$ is also in this field.

When an irreducible polynomial over F picks up a root in a larger field E, more roots do not have to be in E. A simple example is $T^3 - 2$ in $\mathbf{Q}[T]$, which has only one root in $\mathbf{R}$.

By repeating the construction of the previous proof several times, we can always create a field with a full set of roots for our polynomial. We state this as a corollary.

Corollary 1.3. *Let F be a field and $f(T) = a_m T^m + \cdots + a_0$ be in $F[T]$ with degree $m \geq 1$. There is a field $K \supset F$ such that in $K[T]$,*

$$f(T) = a_m(T - \alpha_1) \cdots (T - \alpha_m). \tag{1.1}$$

Proof. Exercise. □

The situation in $\mathbf{F}_p[T]$ is much simpler than in $\mathbf{Q}[T]$. We will see later (Theorem 3.7) that for an irreducible in $\mathbf{F}_p[T]$, a larger field which contains one root must contain all the roots. This will not be needed until Section 3, but we give two examples now so the idea is clear.

Example 1.4. In $\mathbf{F}_7[T]$, $T^3 - 2$ is irreducible. It has a root in the field $\mathbf{F}_7[x]/(x^3 - 2)$, namely $\overline{x}$. It also has two other roots in this field, $2\overline{x}$ and $4\overline{x}$.

Example 1.5. In $\mathbf{F}_5[T]$, $T^3 + T^2 + 1$ is irreducible. In the field $\mathbf{F}_5[x]/(x^3 + x^2 + 1)$, the polynomial has the root $\overline{x}$ and also the roots $2\overline{x}^2 + 3\overline{x}$ and $3\overline{x}^2 + \overline{x} + 4$.

2. Divisibility and roots in $F[T]$

An important result on the first-year problem sets is the connection between roots and divisibility by linear polynomials. For $f(T) \in F[T]$ and $\alpha \in F$, $f(\alpha) = 0 \iff (T-\alpha)|f(T)$. The next result is an analogue for divisibility by higher degree polynomials in $F[T]$, provided they are irreducible. (All linear polynomials are irreducible.)

Theorem 2.1. *Let $\pi(T)$ be irreducible in $F[T]$ and let α be of a root of $\pi(T)$ in some larger field. For $h(T)$ in $F[T]$, $h(\alpha) = 0 \iff \pi(T)|h(T)$.*

Proof. If $h(T) = \pi(T)g(T)$, then $h(\alpha) = \pi(\alpha)g(\alpha) = 0$.

Now assume $h(\alpha) = 0$. We want to show $\pi|h$. By the division algorithm in $F[T]$, $h(T) = \pi(T)q(T) + r(T)$, where $q(T)$ and $r(T)$ are in $F[T]$ and $\deg r < \deg \pi$ or $r = 0$. Since α is a root of both $h(T)$ and $\pi(T)$, we get $r(\alpha) = 0$.

Suppose $r(T) \neq 0$. Because $\pi(T)$ is irreducible in $F[T]$ and $\deg r < \deg \pi$, $r(T)$ and $\pi(T)$ are relatively prime. Therefore we can write

$$1 = a(T)\pi(T) + b(T)r(T)$$

for some $a(T), b(T) \in F[T]$. Substitute α for T, and the right side vanishes. This is absurd, so $r(T) = 0$. □

Example 2.2. Consider $\pi(T) = T^2 - 2$ in $\mathbf{Q}[T]$. It has a root $\sqrt{2} \in \mathbf{R}$. For any $h(T) \in \mathbf{Q}[T]$, $h(\sqrt{2}) = 0 \iff (T^2-2)|h(T)$. This equivalence breaks down if we allow $h(T)$ to come from $\mathbf{R}[T]$: try $h(T) = T - \sqrt{2}$.

3. Roots of irreducibles in $\mathbf{F}_p[T]$

This section makes explicit the relations among the roots of an irreducible polynomial in $\mathbf{F}_p[T]$. In short, we can obtain all roots from any one root by repeatedly taking p-th powers. The precise statement is in Theorem 3.7.

The ring $\mathbf{F}_p[T]$, or more generally the ring $\mathbf{F}_p[T]/h$, contains $\mathbf{F}_p$. Therefore $p = 0$ in these rings.

Lemma 3.1. *Let A be a ring in which $p = 0$. Pick any x and y in A.*

a) $(x+y)^p = x^p + y^p$.

b) When A is a field, $x^p = y^p \implies x = y$.

Proof. a) By the binomial theorem,

$$(x+y)^p = x^p + \sum_{k=1}^{p-1} \binom{p}{k} x^{p-k}y^k + y^p.$$

For $1 \leq k \leq p-1$, the integer $\binom{p}{k}$ is a multiple of p, so the intermediate terms are 0 in A.

b) Now assume A is a field and $x^p = y^p$. Then $0 = x^p - y^p = (x-y)^p$. (Note $(-1)^p = -1$ for $p \neq 2$, and also for $p = 2$ since $2 = 0 \implies -1 = 1$ in A.) Since A is a field, from $(x-y)^p = 0$ we get $x - y = 0$, so $x = y$. □

Theorem 3.2. *For any $f(T) \in \mathbf{F}_p[T]$, $f(T)^{p^m} = f(T^{p^m})$ for $m \geq 0$.*

Proof. The case $m = 1$ was on the first-year number theory sets. Induct. □

Example 3.3. In $\mathbf{F}_5[T]$, $(2T^4 + T^2 + 3)^5 = 2T^{20} + T^{10} + 3$.

Lemma 3.4. *For $h(T)$ in $\mathbf{F}_p[T]$ with degree d, $\mathbf{F}_p[T]/h$ has size p^d.*

Proof. Exercise. □

Lemma 3.5. *When F is a finite field with size q, $a^q = a$ for all a in F.*

Proof. The equation is clear for $a = 0$. For $a \neq 0$ in F, $a^{q-1} = 1$ (similar to proof that $u^{\varphi(m)} = 1$ in U_m), so $a^q = a$. □

Theorem 3.6. *Let $\pi(T)$ be irreducible of degree d in $\mathbf{F}_p[T]$.*

a) In $\mathbf{F}_p[T]$, $\pi(T)|(T^{p^d} - T)$.

b) For $n \geq 0$, $\pi(T)|(T^{p^n} - T) \iff d|n$.

Proof. The divisibility in (a) in the same as the congruence $T^{p^d} \equiv T \bmod \pi(T)$, or equivalently the equation $\overline{T}^{p^d} = \overline{T}$ in $\mathbf{F}_p[T]/\pi$. Such an equation follows immediately from Lemmas 3.4 and 3.5, using the field $\mathbf{F}_p[T]/\pi$.

To prove ($\Longleftarrow$) in (b), write $n = kd$. Starting with $T^{p^d} \equiv T \bmod \pi$ (from (a)) and applying the p^d-th power to both sides k times, we obtain

$$T \equiv T^{p^d} \equiv T^{p^{2d}} \equiv \cdots \equiv T^{p^{kd}} \bmod \pi.$$

Thus $\pi(T)|(T^{p^n} - T)$.

Now we prove ($\Longrightarrow$) in (b). We assume

$$T^{p^n} \equiv T \bmod \pi \tag{3.1}$$

and want to show $d|n$. Write $n = dq + r$ with $0 \leq r < d$. We will show $r = 0$.

We have $T^{p^n} = T^{p^{dq}p^r} = (T^{p^{dq}})^{p^r}$. By ($\Longleftarrow$), $T^{p^{dq}} \equiv T \bmod \pi$, so $T^{p^n} \equiv T^{p^r} \bmod \pi$. Thus, by (3.1),

$$T^{p^r} \equiv T \bmod \pi. \tag{3.2}$$

This tells us that one particular element of $\mathbf{F}_p[T]/\pi$, the class of T, is equal to its own p^r-th power. Let's extend this property to all elements of $\mathbf{F}_p[T]/\pi$. For any $f(T) \in \mathbf{F}_p[T]$, $f(T)^{p^r} = f(T^{p^r})$ by Theorem 3.2. Combining with (3.2),

$$f(T)^{p^r} \equiv f(T) \bmod \pi.$$

Therefore, in $\mathbf{F}_p[T]/\pi$, the class of $f(T)$ is equal to its own p^r-th power. As $f(T)$ is a general polynomial in $\mathbf{F}_p[T]$, we have proved every $a \in \mathbf{F}_p[T]/\pi$ satisfies $a^{p^r} = a$ (in $\mathbf{F}_p[T]/\pi$). Recall r is the remainder when n is divided by d.

Consider now the polynomial $X^{p^r} - X$. When $r > 0$, this is a nonzero polynomial, with degree p^r. We have found p^d different roots of this polynomial in the field $\mathbf{F}_p[T]/\pi$, namely every element. Therefore $p^d \leq p^r$, so $d \leq r$. But, recalling where r came from, $r < d$. This is a contradiction, so $r = 0$. That proves $d|n$. □

Theorem 3.7. *Let $\pi(T)$ be irreducible in $\mathbf{F}_p[T]$ with degree d and $E \supset \mathbf{F}_p$ be a field in which $\pi(T)$ has a root, say α. Then $\pi(T)$ has roots $\alpha, \alpha^p, \alpha^{p^2}, \cdots, \alpha^{p^{d-1}}$. These d roots are distinct; more precisely, when i and j are nonnegative, $\alpha^{p^i} = \alpha^{p^j} \iff i \equiv j \bmod d$.*

Proof. Since $\pi(T)^p = \pi(T^p)$ by Theorem 3.2, we see α^p is also a root of $\pi(T)$, and likewise $\alpha^{p^2}, \alpha^{p^3}$, and so on by iteration. Once we reach α^{p^d} we have cycled back to the start: $\alpha^{p^d} = \alpha$ by Theorem 3.6a. (Write the divisibility in Theorem 3.6a as an equation in $\mathbf{F}_p[T]$ and then substitute α for T.)

Now we will show $\alpha^{p^i} = \alpha^{p^j} \iff i \equiv j \bmod d$, where $i, j \geq 0$. Since $\alpha^{p^d} = \alpha$, the implication ($\Longleftarrow$) is straightforward. To argue in the other direction, we may suppose without loss of generality that $i \leq j$, say $j = i + k$ with $k \geq 0$. Then

$$\alpha^{p^i} = \alpha^{p^j} \Longrightarrow \alpha^{p^i} = (\alpha^{p^k})^{p^i}.$$

We conclude $\alpha = \alpha^{p^k}$ by Lemma 3.1b. Therefore $\pi(T)|(T^{p^k} - T)$ in $\mathbf{F}_p[T]$ by Theorem 2.1, so $d|k$ by Theorem 3.6b. Thus, $i \equiv j \bmod d$. □

Since $\pi(T)$ has at most $d = \deg \pi$ roots in any field, Theorem 3.7 tells us $\alpha, \alpha^p, \cdots, \alpha^{p^{d-1}}$ are a complete set of roots of $\pi(T)$: in $E[T]$, $\pi(T)$ decomposes into (distinct!) linear factors.

Example 3.8. The polynomial T^3+T^2+1 is irreducible in $\mathbf{F}_2[T]$. In $E = \mathbf{F}_2[x]/(x^3+x^2+1)$, one root of the polynomial is $\overline{x}$. The other two roots are $\overline{x}^2$ and $\overline{x}^4$.

If we wish to write the third root without going beyond the second power of $\overline{x}$, note $x^4 \equiv x^2 + x + 1 \bmod x^3 + x^2 + 1$. Therefore, the roots of $T^3 + T^2 + 1$ in E are $\overline{x}$, $\overline{x}^2$, and $\overline{x}^2 + \overline{x} + 1$.

Now we can remove the mystery behind the listing of the roots in Example 1.5. There was no guessing or brute-force searching involved. The roots are $\overline{x}$, $\overline{x}^5$, and $\overline{x}^{25}$. Then remainders modulo $x^3 + x^2 + 1$ (in $\mathbf{F}_5[x]$) were computed for x^5 and x^{25}.

4. Counting irreducibles in $\mathbf{F}_p[T]$

A nice application of Theorem 3.6 is the next result, which goes back to Gauss. It provides a method to locate the irreducible polynomials of a given degree in $\mathbf{F}_p[T]$, by factoring a certain polynomial.

Theorem 4.1. *Let $n \geq 1$. In $\mathbf{F}_p[T]$,*

$$T^{p^n} - T = \prod_{d|n} \prod_{\substack{\pi \text{ monic} \\ \deg \pi = d}} \pi(T), \tag{4.1}$$

where $\pi(T)$ is irreducible.

Proof. From Theorem 3.6, the irreducible factors of $T^{p^n} - T$ in $\mathbf{F}_p[T]$ are the irreducibles with degree dividing n. Moreover, since $T^{p^n} - T$ is monic, we can write its prime factorization with only monic irreducible factors. What remains is to show that each irreducible factor of $T^{p^n} - T$ appears only once in the factorization. Let $\pi(T)$ be an irreducible factor of $T^{p^n} - T$ in $\mathbf{F}_p[T]$. We want to show $\pi(T)^2$ does not divide $T^{p^n} - T$.

There is a field E in which $\pi(T)$ has a root, say α. We will work in $E[T]$. Observe that

$$\begin{aligned} T^{p^n} - T &= T^{p^n} - T - (\alpha^{p^n} - \alpha) \\ &= (T - \alpha)^{p^n} - (T - \alpha) \text{ by Lemma 3.1a} \\ &= (T - \alpha)((T - \alpha)^{p^n - 1} - 1). \end{aligned}$$

The second factor in this last expression does not vanish at α, so it is not divisible by $T - \alpha$. Therefore $(T - \alpha)^2$ does not divide $T^{p^n} - T$. Since $(T - \alpha)|\pi(T)$, $\pi(T)^2$ does not divide $T^{p^n} - T$. □

Example 4.2. We factor $T^{2^n} - T$ in $\mathbf{F}_2[T]$ for $n = 1, 2, 3, 4$. We have

$$T^2 - T = T(T + 1),$$

$$T^4 - T = T(T+1)(T^2+T+1),$$
$$T^8 - T = T(T+1)(T^3+T+1)(T^3+T^2+1),$$
$$T^{16} - T = T(T+1)(T^2+T+1)(T^4+T+1)(T^4+T^3+1)(T^4+T^3+T^2+T+1).$$

Therefore we have a table listing all the irreducibles of each small degree in $\mathbf{F}_2[T]$:

n	Irreducibles of degree n in $\mathbf{F}_2[T]$
1	T, $T+1$
2	T^2+T+1
3	T^3+T+1, T^3+T^2+1
4	T^4+T+1, T^4+T^3+1, $T^4+T^3+T^2+T+1$

Factoring $T^{p^n} - T$ is, in practice, not the way to find all the monic irreducibles of degree n. For instance, with a computer that can handle finite field arithmetic, it is much easier to find all monic irreducibles of degree 6 in $\mathbf{F}_5[T]$ by running through all monics of degree 6 individually, and seeing which don't factor, than asking the computer to factor $T^{5^6} - T$ in $\mathbf{F}_5[T]$. This polynomial has degree $5^6 = 15625$ and 2635 irreducible factors (of which 2580 have degree 6). That is a large factorization for a computer to find.

The following theorem makes explicit an idea used in the proof of Theorem 4.1: divisibility relations in $F[T]$ can be checked by working over any larger field.

Theorem 4.3. *Let F be a field and K be a larger field. For $f(T)$ and $g(T)$ in $F[T]$, $f(T)|g(T)$ in $F[T]$ if and only if $f(T)|g(T)$ in $K[T]$.*

Proof. It is clear that divisibility in $F[T]$ implies divisibility in the larger $K[T]$. Conversely, suppose $f(T)|g(T)$ in $K[T]$. Then $g(T) = f(T)h(T)$ for some $h(T) \in K[T]$. By the division algorithm in $F[T]$, $g(T) = f(T)q(T) + r(T)$, where $q(T)$ and $r(T)$ are in $F[T]$ and $r(T) = 0$ or $\deg r < \deg f$. Comparing these two formulas for $g(T)$, the uniqueness of the division algorithm in $K[T]$ implies $q(T) = h(T)$ and $r(T) = 0$. Therefore $g(T) = f(T)q(T)$, so $f(T)|g(T)$ in $K[T]$. □

Let $\mathrm{N}_p(n)$ be the number of monic irreducibles of degree n in $\mathbf{F}_p[T]$. For instance, $\mathrm{N}_p(1) = p$. We will use Theorem 4.1 to give a formula for $\mathrm{N}_p(n)$, using Möbius inversion.

On the right side of (4.1), for each d dividing n there are $\mathrm{N}_p(d)$ different monic irreducible factors of degree d (each appearing just once). Taking degrees of both sides of (4.1),

$$p^n = \sum_{d|n} d\mathrm{N}_p(d).$$

This is an identity for all $n \geq 1$, so by Möbius inversion

$$n\mathrm{N}_p(n) = \sum_{d|n} \mu\left(\frac{n}{d}\right) p^d.$$

Therefore

$$\mathrm{N}_p(n) = \frac{1}{n}\sum_{d|n} \mu\left(\frac{n}{d}\right) p^d = \frac{p^n}{n} + \text{lower degree terms in } p. \tag{4.2}$$

Example 4.4. $\mathrm{N}_p(2) = \dfrac{p^2-p}{2}$, $\mathrm{N}_p(9) = \dfrac{p^9-p^3}{9}$, $\mathrm{N}_p(12) = \dfrac{p^{12}-p^6-p^4+p^2}{12}$.

Since any factor of n which is less than n is at most $n/2$, it is not hard to check from (4.2) that $\mathrm{N}_p(n) \sim p^n/n$ as $n \to \infty$. Since there are p^n monics of degree n in total, we interpret this asymptotic to mean the probability a random monic of degree n in $\mathbf{F}_p[T]$ is irreducible is (around) $1/n$. Thus, when sampling monics of degree 6 in $\mathbf{F}_5[T]$, around $1/6$ of them are irreducible.

Here is an analogue of the prime number theorem in $\mathbf{F}_p[T]$. The proof is left to the reader.

Theorem 4.5. *As $n \to \infty$,*

$$\#\{\text{monic irred. } \pi \in \mathbf{F}_p[T] : \deg \pi \leq n\} \sim \frac{(p-1)}{p} \cdot \frac{p^n}{n}.$$