UNIVERSITY OF DUBLIN
TRINITY COLLEGE
SCHOOL OF MATHEMATICS

Evariste Galois 1811-1832

Course 373
Finite Fields

Timothy Murphy

Contents

Chapter 1

The Prime Fields

YOU WILL BE FAMILIAR with *finite* or *modular* arithmetic—in which an integer $m > 0$ is chosen as *modulus*, and we perform the arithmetic operations (addition, subtraction and multiplication) *modulo* m.

These operations define the structure of a *commutative ring* on the set of remainders

$$\{0, 1, 2, \dots, m-1\}.$$

(Recall that a commutative ring is defined by 2 binary operations—addition and multiplication—satisfying the usual laws of arithmetic: addition and multiplication are both commutative and associative, and multiplication is distributive over addition.)

We denote this ring by $\mathbb{Z}/(m)$ (said: 'the ring Z modulo m'). We can think of $\mathbb{Z}/(m)$ either as the set $\{0, 1, \dots, m-1\}$ of remainders; or as the set of congruence classes

$$\bar{a} = \{\dots, a-2m, a-m, a, a+m, a+2m, \dots\} \quad (a = 0, 1, 2, \dots, m-1).$$

The latter is 'classier'; but the former is perfectly adequate, and probably preferable for our purposes.

Example 1. Let $m = 6$. Addition and multiplication in $\mathbb{Z}/(6)$ are given by

$+$	0	1	2	3	4	5
0	0	1	2	3	4	5
1	1	2	3	4	5	0
2	2	3	4	5	0	1
3	3	4	5	0	1	2
4	4	5	0	1	2	3
5	5	0	1	2	3	4

$\times$	0	1	2	3	4	5
0	0	0	0	0	0	0
1	0	1	2	3	4	5
2	0	2	4	0	2	4
3	0	3	0	3	0	3
4	0	4	2	0	4	2
5	0	5	4	3	2	1

Proposition 1. *Suppose p is* prime. *Then each non-zero element $a \in \mathbb{Z}/(p)$ is* invertible, *ie there exists an element $b \in \mathbb{Z}/(p)$ such that*

$$ab \equiv 1 \pmod{p}$$

Proof. Consider the p remainders

$$a \cdot 0 \bmod p,\ a \cdot 1 \bmod p,\ \dots,\ a \cdot (p-1) \bmod p.$$

These are distinct. For if

$$a \cdot r \equiv a \cdot s \pmod{p},$$

where $0 \le r < s \le p-1$, then

$$a \cdot (s-r) \equiv 0 \pmod{p}.$$

In other words,

$$p \mid a(s-r).$$

Since p is prime, this implies that

$$p \mid a \text{ or } p \mid s-r.$$

Both these are impossible, since $0 < a < p$ and $0 < s-r < p$.

Since the p remainders $a \cdot i \bmod p$ above are distinct, they must constitute the full set of remainders modulo p (by the Pigeon-Hole Principle). In particular, they must include the remainder 1, ie for some b

$$a \cdot b \equiv 1 \pmod{p}.$$

□

Recall that a *field* is a commutative ring with precisely this property, i.e. in which every non-zero element is invertible.

Corollary 1. *For each prime p, $\mathbb{Z}/(p)$ is a field.*

Definition 1. *We denote this field by $\mathbb{F}_p$.*

The reason for the double notation—$\mathbb{F}_p$ and $\mathbb{Z}/(p)$—is this. We shall show later that there exists a unique field $\mathbb{F}_{p^n}$ for each prime-power p^n. The fields $\mathbb{F}_p$ form so to speak the lowest layer in this hierarchy.

Nb: $\mathbb{F}_{p^n}$ is *not* the same as the ring $\mathbb{Z}/(p^n)$, unless $n = 1$. Indeed, it is easy to see that $\mathbb{Z}/(m)$ cannot be a field unless m is prime.

Finite fields are often called *Galois fields*, in honour of their discoverer, the French mathematician Évariste Galois. As you probably know, Galois died in a duel (not even over a woman!) at the age of 21.

The notation $\mathbf{GF}(q)$ is sometimes used in place of $\mathbb{F}_q$, although $\mathbb{F}_q$ seems to be becoming standard, presumably to emphasize that finite fields should be considered on a par with the familiar fields $\mathbb{Q}, \mathbb{R}, \mathbb{C}$.

Example 2. Addition and multiplication in $\mathbb{F}_7$ are given by

$+$	0	1	2	3	4	5	6
0	0	1	2	3	4	5	6
1	1	2	3	4	5	6	0
2	2	3	4	5	6	0	1
3	3	4	5	6	0	1	2
4	4	5	6	0	1	2	3
5	5	6	0	1	2	3	4
6	6	0	1	2	3	4	5

$\times$	0	1	2	3	4	5	6
0	0	0	0	0	0	0	0
1	0	1	2	3	4	5	6
2	0	2	4	6	1	3	5
3	0	3	6	2	5	1	4
4	0	4	1	5	2	6	3
5	0	5	3	1	6	4	2
6	0	6	5	4	3	2	1

Summary: For each prime p the remainders modulo p form a field $\mathbb{F}_p$ containing p elements.

Chapter 2

The Prime Subfield of a Finite Field

A SUBFIELD OF A FIELD F is a subset $K \subset F$ containing 0 and 1, and closed under the arithmetic operations—addition, subtraction, multiplication and division (by non-zero elements).

Proposition 2. *Suppose F is a field. Then F contains a smallest subfield P.*

Proof. Any intersection of subfields is evidently a subfield. In particular, the intersection of *all* subfields of F is a subfield P contained in every other subfield. ☐

Definition 2. *We call the smallest subfield P of a field F the* prime *(or rational) subfield of F.*

Definition 3. *The* characteristic *of a field F is defined to be the smallest integer $n > 0$ such that*

$$n \cdot 1 = \overbrace{1 + 1 + \cdots + 1}^{n \text{ times}} = 0,$$

if there is such an integer; or 0 otherwise.

Proposition 3. *The characteristic of a field is either a prime or 0. The characteristic of a finite field is always a prime.*

Proof. Suppose the characteristic n of the field F is a non-prime integer, say $n = rs$, where $1 < r, s$. Since $1 \cdot 1 = 1$, repeated application of the distributive law gives

$$(r \cdot 1)(s \cdot 1) = (\overbrace{1 + 1 + \cdots + 1}^{r \text{ times}})(\overbrace{1 + 1 + \cdots + 1}^{s \text{ times}}) = n \cdot 1 = 0.$$

Since F is a field, it follows that either $r \cdot 1 = 0$ or $s \cdot 1 = 0$; and in either case the characteristic of F is less than n, contrary to hypothesis.

Now suppose F is finite. Then the sequence

$$0,\ 1,\ 1+1,\ 1+1+1,\ \ldots$$

must have a repeat; say

$$r \cdot 1 = s \cdot 1$$

where $r < s$. Then

$$(s - r) \cdot 1 = 0,$$

and so F has finite characteristic. ☐

Proposition 4. *If F is a field of characteristic p, then its prime subfield $P \subset F$ is uniquely isomorphic to $\mathbb{F}_p$:*

$$\operatorname{char} F = p \Longrightarrow P = \mathbb{F}_p.$$

Proof. If F has characteristic p then we can define a map

$$\Theta : \mathbb{F}_p \to F$$

by

$$r \mapsto r \cdot 1 \quad (r = 0, 1, \ldots, p - 1).$$

It is readily verified that this map preserves addition and multiplication, and so is a *homomorphism*. (We always take 'homomorphism' to mean *unitary* homomorphism, i.e. we assume that $\Theta(1) = 1$.)

Now a homomorphism of *fields* is necessarily injective. For suppose $\Theta a = \Theta b$, where $a \neq b$. Let $c = b - a$. Then

$$\begin{aligned} \Theta a = \Theta b &\Longrightarrow \Theta c = 0 \\ &\Longrightarrow \Theta(1) = \Theta(cc^{-1}) = \Theta c \Theta c^{-1} = 0 \\ &\Longrightarrow \Theta(x) = \Theta(x \cdot 1) = \Theta(x)\Theta(1) = 0, \end{aligned}$$

for all x.

Thus Θ defines an isomorphism between $\mathbb{F}_p$ and $\operatorname{im}\Theta$.

But every subfield of F contains the element 1, and so also contains $r \cdot 1 = 1 + \cdots + 1$. Hence the field $\operatorname{im}\Theta$ is contained in every subfield of F, and so must be its prime subfield:

$$P = \operatorname{im}\Theta \cong \mathbb{F}_p.$$

Finally, the isomorphism Θ is unique, since

$$\Theta 1 = 1 \Longrightarrow \Theta r = \Theta(1 + \cdots + 1) = \Theta 1 + \ldots \Theta 1 = r \cdot 1.$$

□

Corollary 2. *$\mathbb{F}_p$ is the only field containing p elements.*

Much the same argument shows that the prime subfield of a field of characteristic 0—which as we have seen must be infinite—is uniquely isomorphic to the rational field $\mathbb{Q}$:

$$\operatorname{char} F = 0 \Longrightarrow P = \mathbb{Q}.$$

Summary: Every finite field F contains one of the prime fields $\mathbb{F}_p$ as its smallest (or prime) subfield.

Chapter 3

Finite Fields as Vector Spaces

SUPPOSE THAT F is a finite field of characteristic p, with prime subfield $P = \mathbb{F}_p$. Then we can regard F as *a vector space over P*.

You may be more familiar with vector spaces over $\mathbb{C}$ and $\mathbb{R}$. In fact the full panoply of linear algebra—the concepts of basis, dimension, linear transformation, etc—carry over unchanged to the case of vector spaces over a finite field.

Theorem 1. *Suppose F is a finite field of characteristic p. Then F contains p^n elements, for some n:*

$$\|F\| = p^n.$$

Proof. Suppose that F, as a vector space, has dimension n over P. Then we can find a basis $\{e_1, e_2, \ldots, e_n\}$ for F over P. Each element $a \in F$ is then uniquely expressible in the form

$$a = \lambda_1 e_1 + \lambda_2 e_2 + \cdots + \lambda_n e_n.$$

There are just p choices for each coordinate λ_i; so the total number of elements in F is

$$\overbrace{p \cdot p \cdots p}^{n \text{ times}} = p^n.$$

$\square$

By convention, we usually denote the number of elements in F by q. So we have shown that

$$q = p^n :$$

every finite field has prime-power order.

We are going to show—this is one of our main aims—that there is in fact exactly one finite field (up to isomorphism) of each prime order p^n, which we shall denote by $\mathbb{F}_{p^n}$.

Proposition 5. *Suppose the finite field F contains p^n elements; and suppose K is a subfield of F. Then K contains p^m elements, where $m \mid n$.*

Proof. In the proof of the Theorem above we considered F as a vector space over P, and we showed that if this space has dimension n then

$$\|F\| = \|P\|^n.$$

But we can equally well consider F as a vector space over K. Our argument now shows that if this space has dimension d then

$$\|F\| = \|K\|^d.$$

If $\|F\| = p^n$, it follows that $\|K\| = p^m$, where $n = md$. $\square$

Another way to prove this result is to consider the multiplicative groups

$$F^\times = F - \{0\},\ K^\times = K - \{0\},$$

formed by the non-zero elements of F and K. These groups have orders $p^n - 1$ and $p^m - 1$. Since $K^\times$ is a subgroup of $F^\times$, it follows by Lagrange's Theorem that

$$(p^m - 1) \mid (p^n - 1).$$

We leave it to the reader to show that this is true if and only if $m \mid n$.

We shall see later that in fact $\mathbb{F}_{p^n}$ contains exactly one subfield with p^m elements if $m \mid n$; as we may say,

$$\mathbb{F}_{p^m} \subset \mathbb{F}_{p^n} \iff m \mid n.$$

We can exploit the vector-space structure of F in other ways (apart from proving that $\|F\| = p^n$). Suppose $a \in F$. Then multiplication by a defines a map

$$\mu_a : F \to F : x \mapsto ax.$$

This map is evidently a *linear transformation* of F, regarded as a vector space over P. It follows that we can speak of its *trace* and *determinant*; and these will in turn define functions

$$T,\ D : F \to P$$

on F with values in P:

$$T(a) = \operatorname{tr} \mu_a, \quad D(a) = \det \mu_a.$$

We'll return to these functions later, when we have finite fields to hand in which to see them at work. At present the only finite fields we know about are the prime fields $\mathbb{F}_p$ and $T(a)$ and $D(a)$ both reduce trivially to a in this case.

Summary: The number of elements in a finite field is necessarily a prime-power:

$$\|F\| = p^n.$$

Chapter 4

Looking for $\mathbb{F}_4$

DOES THERE EXIST a field with 4 elements? (This is the first case in which there could exist a non-prime field.) A bull-headed approach—with a little help from the computer—will surely succeed in such a simple case.

Let's suppose, then, that the field F has $4 = 2^2$ elements. We know that F must have characteristic 2, so that

$$x + x = 0$$

for all $x \in F$.

Two of the elements of F are 0 and 1. Let the two others be called $\perp$ and $\top$ (said: *bottom* and *top*). Thus

$$F = \{0, 1, \perp, \top\}.$$

Consider the element $\perp + 1$. A little thought shows that it cannot be 0, 1 or $\perp$. For example,

$$\begin{aligned} \perp + 1 = 0 &\implies (\perp + 1) + 1 = 0 + 1 \\ &\implies \perp + (1 + 1) = 1 \\ &\implies \perp + 0 = 1 \\ &\implies \perp = 1, \end{aligned}$$

which contradicts our choice of $\perp$ as an element of F *different* from 0 and 1.

Now we can draw up the addition-table for F:

$+$	0	1	$\perp$	$\top$
0	0	1	$\perp$	$\top$
1	1	0	$\top$	$\perp$
$\perp$	$\perp$	$\top$	0	1
$\top$	$\top$	$\perp$	1	0

Turning to the multiplication table, let's see what we already know:

$\times$	0	1	$\perp$	$\top$
0	0	0	0	0
1	0	1	$\perp$	$\top$
$\perp$	0	$\perp$		
$\top$	0	$\top$		

Evidently it suffices to determine $\perp^2 = \perp \times \perp$, since the remaining products will then follow on applying the distributive law.

We have 4 choices:

$\perp^2 = 0$ Since $\perp$ is non-zero, it has an inverse $\perp^{-1}$. Thus

$$\begin{aligned}
\perp^2 = 0 &\implies \perp^{-1}(\perp^2) = 0 \\
&\implies (\perp^{-1}\perp)\perp = 0 \\
&\implies 1 \cdot \perp = 0 \\
&\implies \perp = 0,
\end{aligned}$$

contrary to our assumption that $\perp$ differs from 0 and 1.

$\perp^2 = 1$ This gives

$$\perp^2 - 1 = (\perp - 1)(\perp + 1) = 0.$$

Since F is a field, this implies that

$$\perp = 1 \text{ or } \perp = -1.$$

In fact since F has characteristic 2, $-1 = 1$ and so

$$\perp^2 = 1 \implies \perp = 1,$$

again contrary to assumption.

More simply, since F has characteristic 2,

$$\perp^2 - 1 = (\perp - 1)^2,$$

the middle term $-2\perp$ vanishing.

$\perp^2 = \perp$ This implies that

$$\perp(\perp - 1) = 0$$

and so either $\perp = 0$ or $\perp = 1$, both of which are excluded.

$\perp^2 = \top$ As Sherlock Holmes said, *When all other possibilities have been exhausted, the one remaining, however improbable, must be true.* So in this case we conclude that we must have

$$\perp^2 = \top.$$

Now we can complete our multiplication table

$$\begin{aligned}
\perp \times \top &= \perp(\perp + 1) = \perp^2 + \perp = \perp + \top = 1, \\
\top \times \top &= (\perp + 1)^2 = \perp^2 + 1 = \top + 1 = \perp.
\end{aligned}$$

$\times$	0	1	$\perp$	$\top$
0	0	0	0	0
1	0	1	$\perp$	$\top$
$\perp$	0	$\perp$	$\top$	1
$\top$	0	$\top$	1	$\perp$

So if there *is* a field with 4 elements, this must be it. But do these tables in fact define a field?

This is a convenient point to review exactly what we mean by a field, by listing the *Field Axioms*.

Definition 4. *A field F is defined by giving*

1. *A set F with 2 distinguished elements 0 and 1;*
2. *Two binary operations on F, ie 2 maps*

$$+ : F \times F \to F, \quad \times : F \times F \to F,$$

subject to the axioms:

(F1) *addition is associative: for all $a, b, c \in F$,*

$$a + (b + c) = (a + b) + c;$$

(F2) *addition is commutative: for all $a, b \in F$,*

$$b + a = a + b;$$

(F3) *for all $a \in F$,*

$$a + 0 = a;$$

(F4) *for each $a \in F$, there is a $b \in F$ such that*

$$a + b = 0;$$

(F5) *multiplication is associative: for all $a, b, c \in F$,*

$$a(bc) = (ab)c;$$

(F6) *multiplication is commutative: for all $a, b \in F$,*

$$ba = ab;$$

(F7) *for all $a \in F$,*

$$a \cdot 1 = a;$$

(F8) *multiplication is distributive over addition: for all $a, b, c \in F$,*

$$a(b + c) = ab + ac.$$

(F9) *for each $a \neq 0$ in F, there is a $b \in F$ such that*

$$ab = 1;$$

The rationals $\mathbb{Q}$, the reals $\mathbb{R}$ and the complex numbers $\mathbb{C}$ are examples of fields, as of course are the finite (or galois) fields $\mathbb{F}_p$.

Proposition 6. *Suppose F is a field. Then*

1. *for each $a, b \in F$, the equation*

$$a + x = b$$

has a unique solution;

2. *for each $a, b \in F$ with $a \neq 0$, the equation*

$$ay = b$$

has a unique solution.

Proof. By (F4) there exists a c such that

$$a + c = 0.$$

But then

$$\begin{aligned} a + (c + b) &= (a + c) + b && \text{by (F1)} \\ &= 0 + b \\ &= b && \text{by (F3)} \end{aligned}$$

Thus $x = c + b$ is a solution of the equation $a + x = b$. It is moreover the only solution, since

$$\begin{aligned} a + x = b = a + y &\implies c + (a + x) = c + (a + y) \\ &\implies (c + a) + x = (c + a) + y && by (F1) \\ &\implies (a + c) + x = (a + c) + y && by (F2) \\ &\implies 0 + x = 0 + y \\ &\implies x = y && by (F3). \end{aligned}$$

The second part of the Proposition is proved in an exactly analogous way. □

Returning to our prospective field F of 4 elements: to prove that this *is* a field we must verify that the axioms (F1–F9) hold.

This is a straighforward, if tedious, task. To verify (F1), for example, we must consider $4^3 = 64$ cases, since each of the 3 elements a, b, c can take any of the 4 values $0, 1, \perp, \top$.

Let's pass the task on to the computer, by giving a little C program to test the axioms.

```
#include <stdio.h>

typedef enum{zero, one, bottom, top} GF4;

char *el[4] = {"0", "1", "b", "t"};

GF4 add[4][4] = {
  {zero, one, bottom, top},
  {one, zero, top, bottom},
  {bottom, top, zero, one},
  {top, bottom, one, zero}
};

GF4 mul[4][4] = {
  {zero, zero, zero, zero},
  {zero, one, bottom, top},
  {zero, bottom, top, one},
  {zero, top, one, bottom}
};
```

```
main() {
  GF4 x, y, z;

  /* testing (F1) */

  for(x = zero; x <= top; x++)
    for (y = zero; y <= top; y++)
      for (z = zero; z <= top; z++)
        if (add[add[x][y]][z] != add[x][add[y][z]])
  printf("(%s + %s) + %s != %s + (%s + %s)\n",
    el[x], el[y], el[z], el[x], el[y], el[z]);

  /* testing (F2) */

  for(x = zero; x <= top; x++)
    for (y = zero; y <= top; y++)
      if (add[x][y] != add[y][x])
printf("%s + %s != %s + %s\n", el[x], el[y], el[y], el[x]);

  /* testing (F3) */

  for(x = zero; x <= top; x++)
    if (add[x][zero] != x)
printf("%s + 0 != %s\n", el[x], el[x]);

  /* testing (F4) */

  for(x = zero; x <= top; x++) {
    for (y = zero; y <= top; y++)
      if (add[x][y] == 0)
break;
    if (y > top)
      printf("%s + x = 0 has no solution in x\n", el[x]);
  }

  /* testing (F5) */

  for(x = zero; x <= top; x++)
    for (y = zero; y <= top; y++)
      for (z = zero; z <= top; z++)
        if (mul[mul[x][y]][z] != mul[x][mul[y][z]])
  printf("(%s * %s) * %s != %s * (%s * %s)\n",
    el[x], el[y], el[z], el[x], el[y], el[z]);

  /* testing (F6) */

  for(x = zero; x <= top; x++)
    for (y = zero; y <= top; y++)
```

```
      if (mul[x][y] != mul[y][x])
printf("%s * %s != %s * %s\n", el[x], el[y], el[y], el[x]);

  /* testing (F7) */

  for(x = zero; x <= top; x++)
    if (mul[x][one] != x)
printf("%s * 1 != %s\n", el[x], el[x]);

  /* testing (F8) */

  for(x = zero; x <= top; x++)
    for (y = zero; y <= top; y++)
      for (z = zero; z <= top; z++)
        if (mul[add[x][y]][z] != add[mul[x][z]][mul[y][z]])
  printf("(%s + %s) * %s != %s * %s + %s * %s\n",
    el[x], el[y], el[z], el[x], el[z], el[y], el[z]);

  /* testing (F9) */

  for(x = one; x <= top; x++) {
    for (y = one; y <= top; y++)
      if (mul[x][y] == 1)
break;
    if (y > top)
      printf("%s * x = 1 has no solution in x\n", el[x]);
  }

}
```

Not a very strenuous test for the computer, admittedly. But at least it shows who is boss.

Summary: There is just one field with 4 elements, as we expected.

Chapter 5

The Multiplicative Group of a Finite Field

SUPPOSE F is a field. The non-zero elements

$$F^\times = F - \{0\}$$

form a group under multiplication. (We could even take this as the definition of a field: a commutative ring whose non-zero elements form a multiplicative group.)

If F contains q elements, then $F^\times$ contains $q-1$ elements. It follows from Lagrange's Theorem for finite groups that

$$a^{q-1} = 1$$

for all $a \in F^\times$.

(There is a very simple proof of Lagrange's Theorem for a finite *abelian*—or commutative—group

$$A = \{a_1, a_2, \dots, a_n\}.$$

Suppose $a \in A$. Consider the n products

$$aa_1, aa_2, \dots, aa_n.$$

These are distinct, since

$$ax = ay \Longrightarrow x = y.$$

Hence they must be all the elements of A, in some order:

$$\{aa_1, aa_2, \dots, aa_n\} = \{a_1, a_2, \dots, a_n\}.$$

Multiplying together the elements on each side,

$$(aa_1)(aa_2)\dots(aa_n) = a_1a_2\dots a_n.$$

In other words,

$$a^n a_1a_2\dots a_n = a_1a_2\dots a_n.$$

Hence

$$a^n = 1,$$

on dividing both sides by $a_1a_2\dots a_n$.)

Theorem 2. *Suppose F is a finite field. Then the multiplicative group $F^\times$ is cyclic.*

Proof. Recall that a group G is said to be *of exponent* e (where e is a positive integer) if

$$g^e = 1$$

for all $g \in G$, and there is no smaller positive integer with this property. (Another way of expressing this is to say that e is the lcm of the orders of the elements of G.)

By Lagrange's Theorem, the exponent e of a finite group G divides its order:

$$e \mid \|G\|.$$

In general a group of exponent e need not contain an element of order e. For example, the symmetric group S_3 has exponent 6 (since it contains elements of orders 2 and 3); but it has no element of order 6 — otherwise it would be cyclic. However, an *abelian* group always has this property.

Lemma 1. *Suppose A is a finite abelian group, of exponent e. Then there exists an element $a \in A$ of order e.*

Proof of Lemma. Let

$$e = p_1^{e_1} \cdots p_r^{e_r}.$$

There must exist an element $a \in A$ of order $p_1^{e_1} m$ for some m, since otherwise p_1 would occur to a lower power in e. Then

$$a_1 = a^m$$

has order $p_1^{e_1}$. Similarly there exist elements $a_2, \ldots, a_r$ or orders $p_2^{e_2}, \ldots, p_r^{e_r}$.

Sublemma 1. *In an abelian group A, if a has order m and b has order n, and* $\gcd(a, b) = 1$*, then ab has order mn.*

Proof of Sublemma. Suppose ab has order d. Since

$$(ab)^{mn} = (a^m)^n (b^n)^m,$$

we have $d \mid mn$.

On the other hand,

$$(ab)^d = 1 \Longrightarrow (ab)^{nd} = 1 \Longrightarrow a^{nd} = 1,$$

since $b^{nd} = (b^n)^d = 1$. But a has order m; consequently

$$m \mid nd \Longrightarrow m \mid d,$$

since $\gcd(m, n) = 1$. Similarly $n \mid d$. But then

$$mn \mid d,$$

since $\gcd(m, n) = 1$.

We conclude that $d = mn$. ∎

The orders of the elements $a_1, \ldots, a_r$ are mutually co-prime. It follows from the Sublemma that their product

$$a_1 \cdots a_r$$

is of order

$$p_1^{e_1} \cdots p_r^{e_r} = e$$

□

Now suppose the multiplicative group F^* has exponent e. Then each of the $q-1$ elements $a \in F^*$ satisfies the polynomial equation

$$x^e - 1 = 0.$$

But a polynomial $p(x)$ of degree d has at most d roots. It follows that

$$q - 1 \leq e.$$

Since $e \mid q-1$ we conclude that

$$e = q - 1.$$

Hence, by our Lemma, F^* contains an element a of order $q-1$, which therefore generates F^* (since this group has $q-1$ elements). In particular, F^* is cyclic. □

Definition 5. *Suppose F is a finite field. A generator of $F^\times$ is called a* primitive element *(or* primitive root*) of F*

Our Theorem can thus be stated in the form: *Every finite field possesses at least one primitive element.*

Recall that Euler's function $\phi(n)$ (for positive integers n) is defined to be the number of numbers i in the range

$$\{0, 1, 2, \ldots, n-1\}$$

coprime to n (ie with $\gcd(i, n) = 1$). Thus

$$\phi(1) = 1,\ \phi(2) = 1,\ \phi(3) = 2,\ \phi(4) = 2,\ \phi(5) = 4,\ \phi(6) = 2,\ \phi(7) = 6,\ \phi(8) = 4,$$

and so on.

Proposition 7. *The number of primitive roots in F is $\phi(q-1)$.*

Proof. Since we know that

$$F^* = C_{q-1},$$

the result is a consequence of the following Lemma.

Lemma 2. *The cyclic group C_n has $\phi(n)$ generators*

Proof of Lemma. Suppose g is a generator of C_n. We have to determine how many of the elements g^r with $0 \leq r < n$ are also generators of C_n.

Sublemma 2. *The order of $g^r \in C_n$ is*

$$\frac{n}{\gcd(n, r)}$$

Proof of Sublemma. Let the order of g^r be d; and let $\gcd(n, r) = e$. Then

$$n = en',\ r = er' \qquad (\gcd(n', r') = 1).$$

Hence

$$(g^r)^{n'} = (g^{en'})^{r'} = (g^n)^{r'} = 1,$$

since $g^n = 1$. It follows that

$$d \mid n'.$$

On the other hand,

$$\begin{aligned} (g^r)^d = 1 &\implies g^{rd} = 1 \\ &\implies n \mid rd \\ &\implies n' \mid r'd \\ &\implies n' \mid d, \end{aligned}$$

since $\gcd(n', r') = 1$.

We conclude that

$$d = n' = \frac{n}{e} = \frac{n}{\gcd(n, r)}.$$

□

In particular, the number of elements of order n in C_n, ie the number of generators of C_n, is equal to the number of integers r in the range $0 \le r < n$ which are coprime to n. But that, by definition, is $\phi(n)$. □

□

Recall the explicit formula for $\phi(n)$: if

$$n = p_1^{e_1} p_2^{e_2} \dots p_s^{e_s}$$

then

$$\phi(n) = p_1^{e_1-1}(p_1 - 1) p_2^{e_2-1}(p_2 - 1) \dots p_s^{e_s-1}(p_s - 1).$$

This follows from the fact that the function $\phi(n)$ is *multiplicative* in the number-theoretic sense, ie

$$\phi(mn) = \phi(m)\phi(n) \text{ if } \gcd(m, n) = 1.$$

(This in turn is a simple consequence of the Chinese Remainder Theorem.) The result now follows from the particular case $n = p^e$. But the only numbers in $\{0, 1, 2, \dots, p^e - 1\}$ *not* coprime to p^e are the multiples of p; and there are just p^{e-1} of these. Hence

$$\phi(p^e) = p^e - p^{e-1} = p^{e-1}(p - 1).$$

So now it is easy to determine the number of primitive elements in a finite field. For example, $\mathbb{F}_{2^4}$ has $\phi(15) = 8$ primitive elements, while $\mathbb{F}_{2^5}$ has $\phi(31) = 30$ primitve elements.

Surprisingly, perhaps, it is just as difficult to prove our theorem for the elementary finite fields $\mathbb{F}_p$ as in the general case. Moreover, there is really no better way of finding *a primitive root modulo p* (ie a primitive element of $\mathbb{F}_p$) than testing the elements $2, 3, 5, 6, \dots$ successively. (We can at least omit powers like 4; for if 4 were primitive 2 would certainly be so.)

On the other hand, once we have found one primitive element $a \in F^\times$ it is easy to determine the others; they are just the powers

$$a^r \text{ where } \gcd(r, q-1) = 1.$$

As an illustration, consider the field $\mathbb{F}_7 = \mathbb{Z}/(7)$. We find that

$$2^3 = 8 \equiv 1.$$

Thus 2 has order 3, and is not primitive. But $3^2 \equiv 2,\ 3^3 \equiv 6$. Since the order of every non-zero element must divide $q - 1 = 6$, we conclude that 3 has order 6, and so is a primitive root modulo 7. There are just $\phi(6) = 2$ primitive elements; and these are the elements 3^r where $0 \le r < 6$ and $\gcd(r, 6) = 1$; in other words $r = 1$ and $r = 5$. Thus the full set of primitive roots modulo 7 is

$$3, 3^5 = 5.$$

(We may note that since $3^6 \equiv 1$,

$$3^5 = 3^{-1}.$$

And clearly, if a is a primitive element of $F^\times$ then so is its inverse a^{-1}.)

Summary: The multiplicative group $F^\times$ of a finite field F is cyclic. The generators of this group are called the *primitive* elements of the field.

Chapter 6

$\mathbb{F}_{16}$

WE CONSTRUCTED $\mathbb{F}_4$ knowing almost nothing of finite fields. Now, with a little more knowledge let's raise our sights a little.

Suppose the field F contains 16 elements. Let π be a primitive element of F. Then

$$F = \{0, 1, \pi, \pi^2, \ldots, \pi^{14}\}.$$

Since

$$\pi^{15} = 1,$$

the multiplication in F is completely determined:

$\times$	0	1	π	π^2	π^3	π^4	π^5	π^6	π^7	π^8	π^9	π^{10}	π^{11}	π^{12}	π^{13}	π^{14}
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	0	1	π	π^2	π^3	π^4	π^5	π^6	π^7	π^8	π^9	π^{10}	π^{11}	π^{12}	π^{13}	π^{14}
π	0	π	π^2	π^3	π^4	π^5	π^6	π^7	π^8	π^9	π^{10}	π^{11}	π^{12}	π^{13}	π^{14}	1
π^2	0	π^2	π^3	π^4	π^5	π^6	π^7	π^8	π^9	π^{10}	π^{11}	π^{12}	π^{13}	π^{14}	1	π
π^3	0	π^3	π^4	π^5	π^6	π^7	π^8	π^9	π^{10}	π^{11}	π^{12}	π^{13}	π^{14}	1	π	π^2
π^4	0	π^4	π^5	π^6	π^7	π^8	π^9	π^{10}	π^{11}	π^{12}	π^{13}	π^{14}	1	π	π^2	π^3
π^5	0	π^5	π^6	π^7	π^8	π^9	π^{10}	π^{11}	π^{12}	π^{13}	π^{14}	1	π	π^2	π^3	π^4
π^6	0	π^6	π^7	π^8	π^9	π^{10}	π^{11}	π^{12}	π^{13}	π^{14}	1	π	π^2	π^3	π^4	π^5
π^7	0	π^7	π^8	π^9	π^{10}	π^{11}	π^{12}	π^{13}	π^{14}	1	π	π^2	π^3	π^4	π^5	π^6
π^8	0	π^8	π^9	π^{10}	π^{11}	π^{12}	π^{13}	π^{14}	1	π	π^2	π^3	π^4	π^5	π^6	π^7
π^9	0	π^9	π^{10}	π^{11}	π^{12}	π^{13}	π^{14}	1	π	π^2	π^3	π^4	π^5	π^6	π^7	π^8
π^{10}	0	π^{10}	π^{11}	π^{12}	π^{13}	π^{14}	1	π	π^2	π^3	π^4	π^5	π^6	π^7	π^8	π^9
π^{11}	0	π^{11}	π^{12}	π^{13}	π^{14}	1	π	π^2	π^3	π^4	π^5	π^6	π^7	π^8	π^9	π^{10}
π^{12}	0	π^{12}	π^{13}	π^{14}	1	π	π^2	π^3	π^4	π^5	π^6	π^7	π^8	π^9	π^{10}	π^{11}
π^{13}	0	π^{13}	π^{14}	1	π	π^2	π^3	π^4	π^5	π^6	π^7	π^8	π^9	π^{10}	π^{11}	π^{12}
π^{14}	0	π^{14}	1	π	π^2	π^3	π^4	π^5	π^6	π^7	π^8	π^9	π^{10}	π^{11}	π^{12}	π^{13}

Turning to addition, consider the sum

$$\pi^i + \pi^j.$$

If $j > i$, we can write this is

$$\pi^i(1 + \pi^{j-i}).$$

Thus addition will be determined if we know $1 + \pi^i$ for each i $(0 \le i \le 14)$.

We have

$$1 + \pi^i = \pi^j,$$

for some $j = \sigma(i)$. Since we are working in characteristic 2, we can rewrite this as

$$\pi^i + \pi^{\sigma(i)} = 1.$$

Evidently σ is a permutation of $\{1, 2, \dots, 14\}$ of order 2,

$$\sigma^2 = 1,$$

ie $\sigma(\sigma(i)) = i$ for all i. Furthermore,

$$\sigma(i) \neq i,$$

since $\pi^i + \pi^i = 0$. Thus σ is a permutation of type 2^7, that is, it splits into 7 2-cycles.

There are

$$\frac{14!}{2^7 7!} = 13 \cdot 11 \cdot 9 \cdot 7 \cdot 5 \cdot 3 = 135135$$

such permutations. For we can write the numbers $1, 2, \dots, 14$ is $14!$ ways; and if we now bracket successive pairs, eg

$$(3, 7)(8, 2)(10, 13)(14, 4)(6, 1)(5, 9)(12, 11),$$

we have a permutation of type 2^7. Each such permutation arises in $2^7 7!$ ways; for firstly, we can write each of the 7 pairs in 2 ways, and secondly, we can order the 7 pairs in $7!$ ways.

Recall that $\pi^{15} = 1$. Thus we can regard the exponents of π as numbers modulo 15, and σ as a permutation of $\mathbb{Z}/(15)$. With this understanding, σ defines the addition of 2 different powers of π by the rule:

$$\pi^i + \pi^j = \pi^{i+\sigma(j-i)}.$$

The permutation σ determines the addition table, while the multiplication table is already known. In principle we could go through the 135135 cases, examining in each case if the 9 field axioms were satisfied.

For example, if

$$\sigma = (1, 11)(2, 8)(3, 12)(4, 5)(6, 13)(7, 9)(10, 14)$$

the addition table starts

$+$	0	1	π	π^2	π^3	π^4	π^5	π^6	π^7	π^8	π^9	π^{10}	π^{11}	π^{12}	π^{13}	π^{14}
0	0	1	π	π^2	π^3	π^4	π^5	π^6	π^7	π^8	π^9	π^{10}	π^{11}	π^{12}	π^{13}	π^{14}
1	1	0	π^{11}	π^8	π^{12}	π^5	π^4	π^{13}	π^9	π^2	π^7	π^{14}	π	π^3	π^6	π^{10}
π	π	π^{11}	0	π^{12}	π^9	π^{13}	π^6	π^5	π^{14}	π^{10}	π^3	π^8	1	π^2	π^4	π^7
π^2	π^2	π^8	π^{12}	0	π^{13}	π^{10}	π^{14}	π^7	π^6	1	π^{11}	π^4	π^9	π	π^3	π^5
π^3	π^3	π^6	π^9	π^{13}	0	π^{14}	π^{11}	1	π^8	π^7	π	π^{12}	π^5	π^{10}	π^2	π^4

On going through the axioms, it is clear that all except the first two are automatically satisfied for every permutation σ of type 2^7—we only need to consider the associativity and commutativity of addition.

Taking commutativity first,

$$\begin{aligned} 1 + \pi^i = \pi^i + 1 &\implies \pi^{\sigma(i)} = \pi^{i+\sigma(-i)} \\ &\implies \sigma(i) = i + \sigma(-i) \\ &\implies \sigma(i) - \sigma(-i) = i. \end{aligned}$$

It's readily verified that if this holds for all $i \in \{1, \ldots, 14\}$ then $\pi^i + \pi^j = \pi^j + \pi^i$ for all i, j.

Suppose this is so. Turning to associativity,

$$\begin{aligned} \pi^i + (1 + \pi^j) = (\pi^i + 1) + \pi^j &\implies \pi^i + \pi^{\sigma(j)} = \pi^{\sigma(i)} + \pi^j \\ &\implies \pi^{i+\sigma(\sigma(j)-i)} = \pi^{j+\sigma(\sigma(i)-j)} \\ &\implies \sigma(\sigma(j) - i) + i = \sigma(\sigma(i) - j) + j. \end{aligned}$$

Again, it is readily verified that if this holds for all $i, j \in \{1, \ldots, 14\}$ then $\pi^i + (\pi^j + \pi^k) = (\pi^i + \pi^j) + \pi^k$ for all i, j, k.

So all(!) we have to do is to run through the 135135 permutations, and determine in each case whether or not these 2 relations hold:

$$\sigma(i) - \sigma(-i) = i$$

for all i, and

$$\sigma(\sigma(j) - i) + i = \sigma(\sigma(i) - j) + j.$$

for all i, j. If these relations do hold then the permutation yields a field of 16 elements.

But it's more instructive to look at the question from another point of view.

Since $(a + b)^2 = a^2 + b^2$ in a field of characteristic 2,

$$\begin{aligned} \pi^i + \pi^{\sigma(i)} = 1 &\implies \pi^{2i} + \pi^{2\sigma(i)} = 1 \\ &\implies \sigma(2i) = 2\sigma(i). \end{aligned}$$

In other words, σ commutes with multiplication by 2. Formally, let

$$\mu : \mathbb{Z}/(15) \to \mathbb{Z}/(15)$$

be the map corresponding to multiplication by 2:

$$\mu(i) = 2i.$$

Then

$$\sigma\mu = \mu\sigma.$$

Now μ is a permutation of $\mathbb{Z}/(15)$; in cyclic notation

$$\mu = (1, 2, 4, 8)(3, 6, 12, 9)(5, 10)(7, 14, 13, 11).$$

(Under μ, in other words, $1 \mapsto 2$, $2 \mapsto 4$, etc.)

There is a neat way of expressing commutation among permutations. Note first that

$$\sigma\mu = \mu\sigma \iff \sigma\mu\sigma^{-1} = \mu.$$

But if we have a permutation expressed in cyclic form, say

$$\theta = (a_1, a_2, \ldots, a_r)(b_1, b_2, \ldots, b_s) \ldots$$

then it is readily verified that for any other permutation g,

$$g\theta g^{-1} = (ga_1, ga_2, \ldots, ga_r)(gb_1, gb_2, \ldots, gb_s) \ldots.$$

(Under $g\theta g^{-1}$,

$$ga_i \overset{g^{-1}}{\mapsto} a_i \overset{\theta}{\mapsto} a_{i+1} \overset{g}{\mapsto} ga_{i+1},$$

for example.) Applying this to the present case,

$$\sigma\mu\sigma^{-1} = (\sigma 1, \sigma 2, \sigma 4, \sigma 8)(\sigma 3, \sigma 6, \sigma 12, \sigma 9)(\sigma 5, \sigma 10)(\sigma 7, \sigma 14, \sigma 13, \sigma 11).$$

This must be the same as μ:

$$\sigma\mu\sigma^{-1} = \mu.$$

It follows that σ must permute the three 4-cycles in some way, and must send the 2-cycle to itself. Since $\sigma i \neq i$, the latter implies that

$$\sigma(5) = 10,\ \sigma(10) = 5.$$

In other words,

$$\pi^5 + \pi^{10} = 1.$$

As for the action of σ on the three 4-cycles, there are 2 possibilities. Either σ sends each 4-cycle into itself, or it interchanges two and sends the third into itself. In any case, it must send at least one 4-cycle into itself.

Suppose σ sends the 4-cycle (a, b, c, d) into itself:

$$(\sigma a, \sigma b, \sigma c, \sigma d) = (a, b, c, d),$$

Then we must have

$$\sigma a = c,\ \sigma b = d,\ \sigma c = a,\ \sigma d = b.$$

For we know that $\sigma a \neq a$; while for example

$$\sigma(a) = b \Longrightarrow \sigma(b) = \sigma(2a) = 2\sigma(a) = 2b = c,$$

whereas we know that

$$\sigma(a) = b \Longrightarrow \sigma(b) = a,$$

since $\sigma^2 = 1$.

We have dramatically reduced the number of possibilities; there is just one σ sending each 4-cycle into itself, while if σ swaps 2 4-cycles, we can choose the one that goes to itself in 3 ways, and there are then 4 ways of defining σ, eg if

$$\sigma(1, 2, 4, 8) = (3, 6, 12, 9)$$

we have 4 possibilities:

$$\left.\begin{array}{llll} \sigma 1 = 3, & \sigma 2 = 6, & \sigma 4 = 12, & \sigma 8 = 9; \\ \sigma 1 = 6, & \sigma 2 = 12, & \sigma 4 = 9, & \sigma 8 = 3; \\ \sigma 1 = 12, & \sigma 2 = 9, & \sigma 4 = 3, & \sigma 8 = 6; \\ \sigma 1 = 9, & \sigma 2 = 3, & \sigma 4 = 6, & \sigma 8 = 12. \end{array}\right\}$$

Note that in each of these cases, σ is completely determined. In the first case, for example,

$$\sigma = (1, 3)(2, 6)(4, 12)(5, 10)(7, 13)(8, 9)(11, 14).$$

Thus we have reduced the 135135 cases to just $1 + 3 \cdot 4 = 13$.

Observe that in this reduction we have not invoked the conditions that commutativity and associativity of addition impose on σ. Consider the commutativity condition:

$$\sigma(-i) = \sigma(i) - i.$$

Suppose σ sends the 4-cycle $(1, 2, 4, 8)$ into itself. As we have seen this implies that

$$\sigma 1 = 4.$$

But this in turn implies that

$$\sigma(-1) = \sigma(14) = \sigma(1) - 1 = 3.$$

Since 14 and 3 belong to different 4-cycles, the first of our 13 cases is out; σ must send one 3-cycle into itself, and interchange the other two.

There is one last reduction. We should perhaps have noted earlier that the three 4-cycles are not on the same footing. While $(1, 2, 4, 8)$ and $(7, 14, 13, 11)$ correspond to primitive elements (since the numbers are co-prime to 15), the 4-cycle $(3, 6, 12, 9)$ corresponds to elements of order 5.

Now we could have chosen any of the permutations π^i, with i co-prime to 15, in place of π. In particular we could have chosen π^7 in place of π. This would have interchanged the two primitive 4-cycles. Thus the case in which the third 4-cycle is sent into itself is effectively the same as that in which the first 4-cycle is sent into itself. We have reduced the 13 cases to 8.

In fact, if the first 4-cycle is sent into itself, σ is determined completely. For as we saw, we must have

$$\begin{aligned}&\sigma 1 = 4,\ \sigma 2 = 8,\ \sigma 3 = 14,\ \sigma 5 = 10,\\ &\sigma 13 = \sigma(-2) = \sigma 2 - 2 = 6, \sigma 12 = \sigma(-3) = \sigma 3 - 3 = 11.\end{aligned}$$

Thus

$$\sigma = (1, 4)(2, 8)(3, 14)(5, 10)(6, 13)(7, 9)(11, 12).$$

Suppose the second 4-cycle is sent into itself. Then

$$\sigma 3 = 12,\ \sigma 6 = 9, \sigma 12 = 3, \sigma 9 = 6.$$

But

$$\sigma 3 = 12 \Longrightarrow \sigma 12 = \sigma(-3) = \sigma 3 - 3 = 9.$$

We conclude that this case cannot occur.

In the end therefore we are left with only 1 case, (corresponding to the permutation above) leading to the addition table

$+$	0	1	π	π^2	π^3	π^4	π^5	π^6	π^7	π^8	π^9	π^{10}	π^{11}	π^{12}	π^{13}	π^{14}
0	0	1	π	π^2	π^3	π^4	π^5	π^6	π^7	π^8	π^9	π^{10}	π^{11}	π^{12}	π^{13}	π^{14}
1	1	0	π^4	π^8	π^{14}	π	π^{10}	π^{13}	π^9	π^2	π^7	π^5	π^{12}	π^{11}	π^6	π^3
π	π	π^4	0	π^5	π^9	1	π^2	π^{11}	π^{14}	π^{10}	π^3	π^8	π^6	π^{13}	π^{12}	π^7
π^2	π^2	π^8	π^5	0	π^6	π^{10}	π	π^3	π^{12}	1	π^{11}	π^4	π^9	π^7	π^{14}	π^{13}
π^3	π^3	π^{14}	π^9	π^6	0	π^7	π^{11}	π^2	π^4	π^{13}	π	π^{12}	π^5	π^{10}	π^8	1
π^4	π^4	π	1	π^{10}	π^7	0	π^8	π^{12}	π^3	π^5	π^{14}	π^2	π^{13}	π^6	π^{11}	π^9
π^5	π^5	π^{10}	π^2	π	π^{11}	π^8	0	π^9	π^{13}	π^4	π^6	1	π^3	π^{14}	π^7	π^{12}
π^6	π^6	p^{13}	π^{11}	π^3	π^2	π^{12}	π^9	0	π^{10}	π^{14}	π^5	π^7	π	π^4	1	π^8
π^7	π^7	π^9	p^{14}	π^{12}	π^4	π^3	π^{13}	π^{10}	0	π^{11}	1	π^6	π^8	π^2	π^5	π
π^8	π^8	π^2	π^{10}	1	π^{13}	π^5	π^4	π^{14}	π^{11}	0	π^{12}	π	π^7	π^9	π^3	π^6
π^9	π^9	π^7	π^3	π^{11}	π	π^{14}	π^6	π^5	1	π^{12}	0	π^{13}	π^2	π^8	π^{10}	π^4
π^{10}	π^{10}	π^5	π^8	π^4	π^{12}	π^2	1	π^7	π^6	π	π^{13}	0	π^{14}	π^3	π^9	π^{11}
π^{11}	π^{11}	π^{12}	π^6	π^9	π^5	π^{13}	π^3	π	π^8	π^7	π^2	π^{14}	0	1	π^4	π^{10}
π^{12}	π^{12}	π^{11}	π^{13}	π^7	π^{10}	π^6	π^{14}	π^4	π^2	π^9	π^8	π^3	1	0	π	π^5
π^{13}	π^{13}	π^6	π^{12}	π^{14}	π^8	π^{11}	π^7	1	π^5	π^3	π^{10}	π^9	π^4	π	0	π^2
π^{14}	π^{14}	π^3	π^7	π^{13}	1	π^9	π^{12}	π^8	π	π^6	π^4	π^{11}	π^{10}	π^5	π^2	0

It only remains to verify that this addition table (together with the ealier multiplication table) do indeed defined a field; that is, the identities enshrining the commutativity and associativity of addition hold.

We pass the task on to the computer.

```
#include <stdio.h>

int sigma[16] = {0, 4, 8, 14, 1, 10, 13, 9, 2, 7, 5, 12, 11, 6, 3};
main()
{
  int i, j;
  int err = 0;

  /* testing (F1) */

  for(i = 1; i <= 14; i++)
    for(j = 1; j <= 14; j++)
      if (sigma[i] != j)
if ((15 + sigma[(15 + sigma[i] -j) % 15] + j) % 15 !=
  (15 + sigma[(15 + sigma[j] -i) % 15] + i) % 15) {
    err++;
        printf("\\pi^%d + (1 + \\pi^%d) != (\\pi^%d + 1) + \\pi^%d\n",
  i, j, i, j);
  }
  if (!err) printf("Axiom F1 satisfied\n");

  /* testing (F2) */

  err = 0;
  for(i = 1; i <= 14; i++)
    if ( (15 + sigma[i] - sigma[15-i]) % 15 != i) {
      err++;
      printf("1 + \\pi^%d != \\pi^%d + 1\n", i, i);
    }
  if (!err) printf("Axiom F2 satisfied\n");

}
```

Summary: We have shown that there exists just 1 field $\mathbb{F}_{2^4}$ containing 16 elements. Moreover, our construction would apply in principle to any finite field $\mathbb{F}_{2^n}$ of characteristic 2.

Chapter 7

Polynomials over a Finite Field

A POLYNOMIAL over a field F is a formal expression

$$f(x) = c_n x^n + c_{n-1} x^{n-1} + \cdots + c_0 \quad (c_i \in F).$$

When F is finite we must distinguish between the polynomial $f(x)$ and the map

$$x \mapsto f(x) : F \to F$$

which it defines; for 2 different polynomials may define the same map, or what comes the the same thing, a polynomial may vanish for all elements of F, as for example the polynomial

$$f(x) = x^2 - x,$$

in the field $\mathbb{F}_2$.

The polynomials over F can be added and multiplied—we assume that the constructions are familiar—and so constitute a commutative ring (with 1) which we denote by $F[x]$.

Example 3. There are just 8 polynomials of degree ≤ 2 over $\mathbb{F}_2$, namely

$$0, 1, x, x+1, x^2, x^2+1, x^2+x, x^2+x+1.$$

We have

$$(x+1) + (x^2+x+1) = x^2, \quad (x+1)(x^2+x+1) = x^3+1.$$

We will be dealing almost exclusively with polynomials over a prime field P. Many of the questions concerning a finite field F can be expressed in terms of the polynomials over its prime subfield, which are generally much easier to get hold of, particularly with a computer.

At the same time, the study of the ring $P[x]$ of polynomials over the prime field P is a subject of great interest in its own right. There is a remarkable analogy between the ring $P[x]$ and the familiar ring of integers $\mathbb{Z}$. Almost every question that one can ask about $\mathbb{Z}$—for example, questions concerning the distribution of the primes—can equally well be asked of $P[x]$. To take an extreme example, the Riemann hypothesis (or more accurately, conjecture)—which has baffled generations of mathematicians—can be proved relatively easily in $P[x]$. (Usually it is simpler to establish a proposition in $P[x]$ than in $\mathbb{Z}$.)

Definition 6. *A polynomial $f(x)$ of degree ≥ 1 over the field F is said to be* prime *(or* indecomposable*) if it cannot be expressed as the product of 2 polynomials of lower degree over F.*

Example 4. There are just 5 prime polynomials of degree ≤ 3 over $\mathbb{F}_2$, namely

$$x, x+1, x^2+x+1, x^3+x+1, x^3+x^2+1.$$

Proposition 8. *(The Prime Factorisation Theorem) Every polynomial over the field F is expressible as a product of prime polynomials over F, unique up to order (and scalar multiples).*

Proof. This is almost identical with the usual proof in the classical case $\mathbb{Z}$.

Lemma 3. *Suppose $f(x), g(x) \in F[x]$; and suppose $g \neq 0$. Then we can divide f by g to obtain quotient $q(x)$ and remainder $r(x)$:*

$$f(x) = q(x)g(x) + r(x) \quad (\deg r < \deg g).$$

Lemma 4. *Suppose $f(x), g(x) \in F[x]$. Then f and g have a greatest common divisor*

$$d(x) = \gcd(f(x), g(x))$$

such that

$$d(x) \mid f(x), g(x);$$

and if $e(x) \in F(x)$ then

$$e(x) \mid f(x), g(x) \Longrightarrow d(x) \mid e(x).$$

Furthermore, we can find polynomials $u(x), v(x) \in F[x]$ such that

$$u(x)f(x) + v(x)g(x) = d(x).$$

Proof. We apply the Euclidean algorithm to $f(x)$ and $g(x)$:

$$\begin{aligned}
f(x) &= q_1(x)g(x) + r_1(x) \\
g(x) &= q_2(x)r_1(x) + r_2(x) \\
r_1(x) &= q_3(x)r_2(x) + r_3(x) \\
&\dots \\
r_{i-1}(x) &= q_{i+1}(x)r_i(x).
\end{aligned}$$

the process must end with an exact division, since the degrees of the remainders are strictly decreasing:

$$\deg g > \deg r_1 > \deg r_2 > \dots.$$

Now it is easy to see that the last non-zero remainder $r_i(x)$ is the required polynomial:

$$r_i(x) = \gcd(f(x), g(x)).$$

For on the one hand, going up the chain we see successively that

$$\begin{aligned}
r_i(x) &\mid r_{i-1}(x), \\
r_i(x) &\mid r_{i-2}(x), \\
&\dots \\
r_i(x) &\mid g(x), \\
r_i(x) &\mid f(x).
\end{aligned}$$

On the other hand, if $e(x) \mid f(x), g(x)$ then going down the chain we see successively that

$$
\begin{array}{rcl}
e(x) & \mid & r_1(x), \\
e(x) & \mid & r_2(x), \\
 & & \ldots \\
e(x) & \mid & r_i(x).
\end{array}
$$

Finally, going down the chain we can successively express $r_1(x), r_2(x), \ldots$ in the form

$$
\begin{aligned}
r_j(x) &= u_j(x)f(x) + v_j(x)g(x) \\
r_{j+1}(x) &= r_{j-1}(x) - q_{j+1}(x)r_j(x) \\
&= \left(u_{j-1}(x) - q_{i+1}(x)u_i(x)\right) f(x) + \left(v_{j-1}(x) - q_{i+1}(x)v_i(x)\right) g(x) \\
&= u_{j+1}(x)f(x) + v_{j+1}(x)g(x),
\end{aligned}
$$

where

$$
u_{j+1}(x) = \left(u_{j-1}(x) - q_{i+1}(x)u_i(x)\right), \; v_{j+1}(x) = \left(v_{j-1}(x) - q_{i+1}(x)v_i(x)\right);
$$

until finally we obtain an expression for $r_i(x) = \gcd(f, g)$ of the form

$$
\gcd(f(x), g(x) = u(x)f(x) + v(x)g(x),
$$

as required. □

Example 5. Working over $\mathbb{F}_2$, suppose

$$
f(x) = x^5 + x^2 + 1, \; g(x) = x^4 + x^3 + 1.
$$

We have

$$
\begin{aligned}
f(x) + xg(x) &= x^4 + x^2 + x + 1, \\
f(x) + xg(x) + g(x) &= x^3 + x^2 + x = r_1(x).
\end{aligned}
$$

This is the first step of the euclidean algorithm. Continuing,

$$
\begin{aligned}
g(x) + xr_1(x) &= x^3 + x + 1, \\
g(x) + xr_1(x) + r_1(x) &= x^2 + x + 1 = r_2(x), \\
r_1(x) + xr_2(x) &= 1 = r_3(x).
\end{aligned}
$$

Hence

$$
\gcd\left(f(x), g(x)\right) = 1,
$$

and working backwards we find that

$$
\begin{aligned}
1 &= r_1(x) + xr_2(x) \\
&= r_1(x) + g(x) + (x+1)r_1(x) \\
&= g(x) + xr_1(x) \\
&= g(x) + f(x) + (x+1)g(x) \\
&= f(x) + xg(x).
\end{aligned}
$$

Returning to the proof of the Prime Factorisation Theorem—sometimes call the Fundamental Theorem of Arithmetic—

Lemma 5. *Suppose $p(x), f(x), g(x) \in F[x]$; and suppose p is prime. Then*

$$p(x) \mid f(x)g(x) \Longrightarrow p(x) \mid f(x) \text{ or } p(x) \mid g(x).$$

Proof. Consider

$$d(x) = \gcd(p(x), f(x)).$$

Since $d(x)$ by definition divides $p(x)$; and since $p(x)$ by definition has only the factors 1 and itself, either $d(x) = 1$ or $d(x) = p(x)$.

If $d(x) = p(x)$ then

$$p(x) \mid f(x)$$

(since $d(x) \mid f(x)$) and we are done.

On the other hand if $d(x) = 1$, then by the Lemma above we can find $u(x), v(x) \in F[x]$ such that

$$u(x)p(x) + v(x)f(x) = 1.$$

Multiplying by $g(x)$,

$$u(x)p(x)g(x) + b(x)f(x)g(x) = g(x).$$

Now $p(x)$ divides both terms on the left (since $p(x) \mid f(x)g(x)$). Hence

$$p(x) \mid g(x).$$

□

Turning to the proof of the Proposition, if $f(x)$ is not a prime then we can factorise it

$$f(x) = u(x)v(x)$$

into 2 polynomials of lesser degree. If these are not prime, they can again be split; until finally we must attain an expression for $f(x)$ as a product of primes.

Finally, if we have 2 expressions for $f(x)$ as products of primes

$$p_1(x) \cdots p_m(x) = f(x) = q_1(x) \cdots q_n(x)$$

then the last Lemma shows that the p's and q's must be the same, up to order. □

Proposition 9. *Suppose F is a finite field, with prime subfield P. Each element $a \in F$ is a root of a unique prime polynomial $m(x)$ over P.*

If $\|F\| = p^n$ then the degree of $m(x)$ is $\leq n$.

For each polynomial $f(x)$ over P,

$$f(a) = 0 \Longleftrightarrow m(x) \mid f(x).$$

Proof. If $\|F\| = p^n$, then

$$\dim_P F = n.$$

Hence if $a \in F$, the $n+1$ elements

$$1, a, a^2, \ldots, a^n$$

must be linearly dependent, ie

$$c_0 + c_1 a + c_2 a^2 + \cdots + c_n a^n = 0$$

for some $c_i \in P$ (not all zero). In other words a is a root of the polynomial

$$f(x) = c_0 + c_1 x + c_2 x^2 + \cdots + c_n x^n = 0.$$

Now let $m(x)$ be the polynomial of smallest degree ≥ 1 satisfied by a. Then

$$\deg m(x) \leq \deg f(x) \leq n.$$

Also $m(x)$ must be prime. For if

$$m(x) = u(x)v(x)$$

then

$$0 = m(a) = u(a)v(a) \Longrightarrow u(a) = 0 \text{ or } v(a) = 0,$$

since F is a field. But that contradicts the minimality of $m(x)$.

Finally, suppose $f(a) = 0$. Divide $f(x)$ by $m(x)$:

$$f(x) = m(x)q(x) + r(x),$$

where $\deg r(x) < \deg m(x)$. Then

$$r(a) = f(a) - m(a)q(a) = 0,$$

and so $r(x) = 0$ by the minimality of $m(x)$, ie $m(x) \mid f(x)$.

This last result shows in particular that $m(x)$ is the *only* prime polynomial (up to a scalar multiple) satisfied by a. □

Remarks 1.

1. Another way of seeing that $a \in F$ satisfies an equation of degree $\leq n$ is to consider the linear map $\mu_a : F \to F$ defined by multiplication by a:

   $$\mu_a(t) = at.$$

   By the Cayley-Hamilton theorem, this linear transformation satisfies its own characteristic equation

   $$\chi_a(x) = \det\left(xI - \mu_a\right).$$

   It follows that a also satisfies this equation:

   $$\chi_a(a) = 0.$$

2. We shall see in Chapter 9 that if $a \in \mathbb{F}_{p^n}$ then the minimal polynomial of a must have degree $d \mid n$.

   Conversely—and more surprisingly—we shall find that all the roots of *any* prime polynomial of degree $d \mid n$ lie in $\mathbb{F}_{p^n}$.

Summary: Each element $a \in F$ is the root of a unique prime polynomial $m(x) \in P[x]$.

Chapter 8

The Universal Equation of a Finite Field

IN AN INFINITE FIELD, a polynomial $p(x)$ cannot vanish for all values of x unless it vanishes identically, ie all its coefficients vanish. For if p is of degree d it cannot have more than d roots, by the Remainder Theorem. In a finite field, however, the position is quite different.

Theorem 3. *Suppose F is a finite field of order q. Then every element $a \in F$ satisfies the equation*

$$U(x) \equiv x^q - x = 0.$$

Proof. By Lagrange's Theorem

$$a^{q-1} = 1$$

for all $a \in F^\times$. Multiplying by a,

$$a^q = a.$$

But this is also satisfied by $a = 0$. Thus it is satisfied by all $a \in F$. □

Corollary 2. *Suppose F is a finite field of order q. Then*

$$x^q - x \equiv \prod_{a \in F} (x - a)$$

over F.

Corollary 3. *Suppose F is a finite field of order q; and suppose $p(x) \in P[x]$, where P is the prime subfield of F. Then*

$$p(x) = 0 \text{ for all } x \in F \iff U(x) \mid p(x).$$

Corollary 4. *Suppose F is a finite field of order q; and suppose $a \in F$. Then the minimal polynomial $m(x)$ of a is a factor of the universal polynomial:*

$$m(x) \mid U(x).$$

Let

$$U_n(x) \equiv x^{p^n} - x.$$

We want to show that

$$U_m(x) \mid U_n(x) \iff m \mid n.$$

It turns out to be simpler to prove a more difficult result.

Proposition 10. *Let* $d = \gcd(m,n)$. *Then*

$$\gcd\left(U_m(x), U_n(x)\right) = U_d(x),$$

where

$$U_m(x) \equiv x^{p^m} - x,\ U_n(x) \equiv x^{p^n} - x,\ U_d(x) \equiv x^{p^d} - x.$$

Proof. Recall the recursive version of the euclidean algorithm (for calculating $\gcd(m,n)$), enshrined in the following C-code.

```
unsigned gcd( unsigned m, unsigned n )
{
  if( m == 0 ) return n;
  if( n == 0 ) return m;
  if ( m < n ) return gcd( m, n - m );
  return gcd( n, m - n );
}
```

Following this idea, we prove the result by induction on $\max(m,n)$. The result is trivial if $m = n$, or $m = 0$, or $n = 0$. We may therefore assume, without loss of generality, that $0 < m < n$. Let

$$n = m + r.$$

By the binomial theorem,

$$\left(x^{p^m} - x\right)^p = x^{p^{m+1}} - x^p,$$

all the terms except the first and last in the expansion vanishing. Repeating this r times,

$$\begin{aligned} U_m(x)^{p^r} &= \left(x^{p^m} - x\right)^{p^r} \\ &= x^{p^{m+r}} - x^{p^r} \\ &= x^{p^n} - x^{p^r} \\ &= U_n(x) - U_r(x). \end{aligned}$$

It follows from this that

$$\gcd\left(U_m(x), U_n(x)\right) = \gcd\left(U_r(x), U_m(x)\right).$$

But by the inductive hypothesis,

$$\begin{aligned} \gcd\left(U_r(x), U_m(x)\right) &= U_{\gcd(r,m)}(x) \\ &= U_{\gcd(m,n)}(x), \end{aligned}$$

since

$$\gcd(r,m) = \gcd(m,n).$$

∎

Corollary 5. *We have*

$$U_m(x) \mid U_n(x) \iff m \mid n.$$

> **Summary:** In a finite field, every element satisfies the universal equation
>
> $$x^q = x,$$
>
> where $q = \|F\|$.

Chapter 9

Uniqueness of the Finite Fields

IF WE ARE NOT YET in a position to show that the field $\mathbb{F}_{p^n}$ exists for each prime powers p^n, we can at least show that there is at most one such field.

Theorem 4. *Two finite fields with the same number of elements are necessarily isomorphic.*

Proof. Suppose F, F' are finite fields with

$$\|F\| = q = \|F'\|.$$

(Of course we know that q must be a prime-power: $q = p^n$.)

Choose a primitive root $\pi \in F$. Let its minimal polynomial be $m(x)$. Then

$$m(x) \mid x^q - x.$$

Now let us go across to F'. Since

$$x^q - x = \prod_{a' \in F'} (x - a'),$$

m(x) must factor completely in F', say

$$m(x) = (x - a'_1) \cdots (x - a'_d).$$

Choose any of these roots as π', say $\pi' = a'_1$. We are going to define an isomorphism

$$\Theta : F \to F'$$

under which

$$\pi \mapsto \pi'.$$

Observe first that π' must be a primitive root in F', ie it must have order $q-1$. For suppose its order were $d < q-1$. Then π' would satisfy the equation

$$x^d - 1.$$

Now $m(x)$, as a prime polynomial satisfied by π', must in fact be its minimal polynomial. Hence

$$m(x) \mid x^d - 1.$$

But then, going back to F, this implies that

$$\pi^d - 1 = 0,$$

ie π has order $< q-1$. We conclude that π' must be a primitive root in F'.

Thus π and π' each generates a cyclic group C_{q-1}. So we can certainly define a group isomorphism

$$\Theta : F^\times \to F'^\times : \pi^i \mapsto \pi'^i.$$

We can extend this to a bijection

$$\Theta : F \to F'$$

by adding the rule $0 \mapsto 0$.

This bijection Θ certainly preserves multiplication:

$$\Theta(ab) = \Theta(a)\Theta(b)$$

for all $a, b \in F$. It remains to show that it also preserves addition, ie

$$\Theta(a+b) = \Theta(a) + \Theta(b).$$

If one (or both) of a and b is 0 this holds trivially; so we may assume that $a, b \neq 0$. There are 2 cases to consider, according as $a+b=0$ or not.

Dealing first with the second (and general) case, let

$$a = \pi^i,\ b = \pi^j,\ a+b = \pi^k.$$

Thus

$$\pi^i + \pi^j = \pi^k$$

in F. In other words, π satisfies the equation

$$x^i + x^j - x^k = 0.$$

It follows that

$$m(x) \mid x^i + x^j - x^k.$$

Going across to F', we deduce that π' also satisfies the equation

$$x^i + x^j - x^k = 0.$$

In other words

$$\pi'^i + \pi'^j = \pi'^k$$

Thus

$$\Theta(a) + \Theta(b) = \Theta(a+b),$$

as required.

It remains to consider the trivial case

$$a + b = 0.$$

If the characteristic is 2 then this implies that $a = b$, in which case it is evident that $\Theta(a) = \Theta(b)$, and so

$$\Theta(a) + \Theta(b) = 0.$$

If the characteristic is odd, then we note that -1 is the only element in F of order 2; for the polynomial

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

has just the 2 roots ± 1. (This is a particular case of our earlier result that the number of elements in F of order $d \mid q-1$ is $\phi(d)$.) In fact we must have

$$-1 = \pi^{\frac{q-1}{2}}$$

since the element on the right certainly has order 2.

Thus if we suppose that $i > j$ (as we may without loss of generality)

$$\begin{aligned} \pi^i + \pi^j = 0 &\implies \pi^{i-j} = -1 \\ &\implies i - j = \frac{q-1}{2} \\ &\implies (\pi')^{i-j} = -1 \\ &\implies \pi'^i + \pi'^j = 0; \end{aligned}$$

so addition is preserved in this case also.

We have shown that the bijection $\Theta : F \to F'$ preserves addition and multiplication; in other words, it is an isomorphism. $\square$

Summary: There is at most 1 field $\mathbb{F}_{p^n}$ with p^n elements. (It remains to be shown that this field actually exists!)

Chapter 10

Automorphisms of a Finite Field

THE AUTOMORHPHISM GROUP G of a field F is usually called its *Galois group. Galois theory* establishes a correspondence between subfields of F and subgroups of G. To each subfield $K \subset F$ we associate the subgroup

$$\{g \in G : gx = x \text{ for all } x \in K\}.$$

Conversely, to each subgroup $H \subset G$ we associate the subfield

$$\{x \in F : gx = x \text{ for all } g \in H\}.$$

In the case of a finite field F, as we shall see, this establishes a one-one correspondence between the subfields of F and the subgroups of G.

Proposition 11. *Suppose F is a finite field of characteristic p. Then the map*

$$a \mapsto a^p$$

is an automorphism of F.

Proof. The map evidently preserves multiplication:

$$(ab)^p = a^p b^p.$$

Less obviously, it also preserves addition:

$$(a + b)^p = a^p + b^p.$$

For on expanding the left-hand side by the binomial theorem, all the terms except the first and last vanish. For

$$p \mid \binom{p}{i} \quad (i = 1, \dots, p - 1),$$

since p divides the numerator but not the denominator of

$$\binom{p}{i} = \frac{p(p-1)\cdots(p-i+1)}{1 \cdot 2 \cdots i}.$$

Finally, the map is injective since

$$a^p = 0 \Longrightarrow a = 0.$$

Since F is finite, this implies that the map is bijective, and so an automorphism of F. ☐

Remarks 2. 1. This is an astonishing result. In characteristic p, the map

$$x \mapsto x^p$$

is *linear*.

2. The map $a \mapsto a^p$ is an injective endomorphism for *any* field F of characteristic p. But it may not be bijective if F is infinite.

Definition 7. *We call the automorphism $a \mapsto a^p$ the* Frobenius automorphism *of F, and denote it by Φ.*

Theorem 5. *Suppose F is a finite field, with*

$$\|F\| = p^n.$$

Then the automorphism group of F is a cyclic group of order n, generated by the Frobenius automorphism:

$$\operatorname{Aut} \mathbb{F}_{p^n} = C_n = \{I, \Phi, \Phi^2, \ldots, \Phi^{n-1} : \Phi^n = I\}$$

Proof.

Lemma 6. *The Frobenius automorphism Φ of $\mathbb{F}_{p^n}$ has order n*

Proof of Lemma. We know that

$$a^{p^n} = a$$

for all $a \in F$. We can rewrite this as

$$\Phi^n(a) = a$$

for all a, ie

$$\Phi^n = I.$$

Suppose

$$\Phi^m = I$$

for some $m < n$. In other words

$$a^{p^m} = a$$

for all $a \in F$. This is an equation of degree p^m with $p^n > p^m$ roots: an impossibility.

We conclude that Φ has order n. □

We must show that

$$I, \Phi, \Phi^2, \ldots, \Phi^{n-1}$$

are the *only* automorphisms of $\mathbb{F}_{p^n}$.

Lemma 7. *Each automorphism Θ of a finite field F leaves invariant each element of its prime subfield P.*

Proof of Lemma. If $c \in P$, we have

$$c = 1 + \cdots + 1.$$

Hence

$$\Theta(c) = \Theta(1) + \cdots + \Theta(1) = 1 + \cdots + 1 = c.$$

□

Lemma 8. *The only elements of a finite field F left invariant by the Frobenius automorphism Φ are the elements of its prime subfield P.*

Proof of Lemma. By the last lemma, the p elements of P are all roots of the equation

$$\Phi(a) \equiv a^p = a.$$

Since this equation has degree p, they are *all* the roots. □

Lemma 9. *Suppose π is a primitive element of the finite field F. Then any automorphism Θ of F is completely determined by its action on π; that is, if Θ, Θ' are 2 such automorphisms then*

$$\Theta(\pi) = \Theta'(\pi) \Longrightarrow \Theta = \Theta'.$$

Proof of Lemma. Since every element $\alpha \neq 0$ in F is of the form $\alpha = \pi^i$ for some i, the result follows from the fact that

$$\Theta(\pi) = \Theta'(\pi) \Longrightarrow \Theta(\pi^i) = \Theta'(\pi^i).$$

□

Let π be a primitive element in F. Consider the product

$$f(x) = (x - \pi)(x - \Phi\pi) \cdots (x - \Phi^{n-1}\pi).$$

Applying the automorphism Φ to this product,

$$\begin{aligned} f^{\Phi}(x) &= (x - \Phi\pi)(x - \Phi^2\pi) \cdots (x - \pi) \\ &= f(x), \end{aligned}$$

the n factors simply being permuted cyclically. Thus $f(x)$ is left unchanged by Φ. From the Lemma above, this implies that the coefficients of $f(x)$ all lie in the prime subfield P:

$$f(x) \in P[x].$$

Now suppose Θ is an automorphism of F. Then

$$f^{\Theta}(x) = f(x),$$

since the coefficients of $f(x)$, being in P, are left unchanged by Θ. Thus

$$\begin{aligned} f^{\Theta}(x) &= (x - \Theta\pi)(x - \Theta\Phi\pi) \cdots (x - \Theta\Phi^{n-1}\pi) \\ &= (x - \pi)(x - \Phi\pi) \cdots (x - \Phi^{n-1}\pi). \end{aligned}$$

It follows that

$$\Theta\pi = \Phi^i\pi$$

for some i. But by the last lemma, this implies that

$$\Theta = \Phi^i.$$

□

Proposition 12. *Suppose $p(x) \in P[x]$ is a prime polynomial of degree d; and suppose $p(x)$ has a root α in the finite field F. Then all the roots of $p(x)$ lie in F; they are in fact the d elements*

$$\left\{\alpha, \Phi\alpha, \ldots, \Phi^{d-1}\alpha\right\}.$$

Proof. Since the automorphism Φ leaves the elements of the prime field P fixed,

$$p(\alpha) = 0 \Longrightarrow p(\Phi\alpha) = 0$$

Thus $\Phi\alpha = \alpha^p$ is also a root of $p(x)$. So by the same argument are $\Phi^2\alpha, \Phi^3\alpha, \ldots$.

On the other hand, we saw in the proof of the last Proposition that

$$f(x) \equiv \prod_{0 \le i < n} \left(x - \Phi^i \alpha\right) \in P[x].$$

Since $p(x)$ is the minimal polynomial of α, and α is a root of $f(x)$, it follows that

$$p(x) | f(x).$$

But $f(x)$ factorises completely in F. Hence the same is true of $p(x)$; and its roots must lie among the roots

$$\left\{\alpha, \Phi\alpha, \ldots, \Phi^{n-1}\alpha\right\}$$

of $f(x)$.

Let e be the least integer > 0 such that

$$\Phi^e \alpha = \alpha.$$

Then the elements

$$\left\{\alpha, \Phi\alpha, \ldots, \Phi^{e-1}\alpha\right\}$$

are all distinct. For if $0 \le i < j \le e$,

$$\Phi^i \alpha = \Phi^j \alpha \Longrightarrow \Phi^{j-i} \alpha = \alpha,$$

on applying the automorphism Φ^{-i}. But since $0 < j - i < e$ that contradicts the minimality of e.

On the other hand, we saw that the elements of this reduced set are all roots of $p(x)$. In fact they are all the roots. For we know that every root is of the form $\Phi^i\alpha$; and if

$$i = eq + r \qquad (0 \le r < e),$$

then

$$\Phi^i \alpha = \Phi^r \alpha.$$

Finally, since $p(x)$ is of degree d, it has just d roots. Hence $d = e$. □

Proposition 13. *The field $\mathbb{F}_{p^n}$ has exactly one subfield containing p^m elements for each $m \mid n$.*

Proof. We know from Chapter 3 that if $F \subset \mathbb{F}_{p^n}$ contains p^m elements then $m \mid n$; and we also know that in this case

$$F = \mathbb{F}_{p^m}.$$

It follows that all the elements of F satisfy the equation

$$x^{p^m} = x.$$

Since this equation has at most p^m roots in $\mathbb{F}_{p^n}$, it follows that

$$F = \left\{x \in \mathbb{F}_{p^n} : x^{p^m} = x\right\}.$$

Conversely, suppose $m \mid n$. Let

$$\begin{aligned} F &= \{x \in \mathbb{F}_{p^n} : x^{p^m} = x\} \\ &= \{x \in \mathbb{F}_{p^n} : \Phi^m x = x\}. \end{aligned}$$

Then F is a subfield of $\mathbb{F}_{p^n}$, since Φ^m is an automorphism of $\mathbb{F}_{p^n}$:

$$\begin{aligned} x, y \in F &\implies \Phi^m x = x, \Phi^m y = y \\ &\implies \Phi^m(x+y) = x+y, \Phi^m(xy) = xy \\ &\implies x+y, xy \in F. \end{aligned}$$

But we saw in Chapter 8 that

$$m \mid n \implies U_m(x) \mid U_n(x).$$

Since $U_n(x)$ factorises completely in F, the same must be true of $U_m(x)$. In other words, $U_m(x)$ has p^m roots in $\mathbb{F}_{p^n}$, ie

$$\|F\| = p^m.$$

□

In conclusion, let us see how this fits in with the general remarks on galois theory with which the chapter opened.

A cyclic group C_n has just 1 subgroup of order m for each $m \mid n$ (ie each m allowed by Lagrange's Theorem). These subgroups are all cyclic themselves. Suppose Φ generates C_n. If $n = md$ then the subgroup of order m is generated by Φ^d:

$$C_m = \left\{1, \Phi^d, \Phi^{2d}, \ldots, \Phi^{(m-1)d}\right\}.$$

According to the prescription of galois theory this corresponds to the subfield

$$\begin{aligned} K &= \left\{x \in \mathbb{F}_{p^n} : \Phi^d x = x\right\} \\ &= \left\{x \in \mathbb{F}_{p^n} : x^{p^d} = x\right\} \\ &= \mathbb{F}_{p^d}. \end{aligned}$$

Thus we have a one-one correspondence between subfields and subgroups:

$$\mathbb{F}_{p^m} \longleftrightarrow C_{n/m}.$$

Notice that under this correspondence, *the larger the subfield the smaller the subgroup*: if $K \longleftrightarrow S, K' \longleftrightarrow S'$,

$$K \subset K' \implies S \supset S'.$$

It follows from this that the Galois correspondence sends *intersections* into *joins*, and vice versa:

$$K \cap K' \longleftrightarrow \langle S, S' \rangle, \quad \langle K, K' \rangle \longleftrightarrow S \cap S'.$$

(The join $\langle K, K' \rangle$ of 2 subfields K, K' is the smallest subfield containing both K and K'; Similarly the join $\langle S, S' \rangle$ of 2 subgroups S, S' is the smallest subgroup containing both S and S'.)

Concretely, if $\mathbb{F}_{p^n}$ exists, and $d \mid n, e \mid n$ then we can regard $\mathbb{F}_{p^d}$ and $\mathbb{F}_{p^e}$ as subfields of $\mathbb{F}_{p^n}$:

$$\mathbb{F}_{p^d}, \mathbb{F}_{p^e} \subset \mathbb{F}_{p^n}.$$

It follows from the galois correspondence that

$$\mathbb{F}_{p^d} \cap \mathbb{F}_{p^e} = \mathbb{F}_{p^{\gcd(d,e)}}, \quad \langle \mathbb{F}_{p^d}, \mathbb{F}_{p^e} \rangle = \mathbb{F}_{p^{\mathrm{lcm}(d,e)}}.$$

Summary: A finite field has just one subfield of each allowed size:

$$\mathbb{F}_{p^m} \subset \mathbb{F}_{p^n} \iff m \mid n.$$

Chapter 11

Wedderburn's Theorem

IF WE RELAX THE CONDITION that multiplication should be commutative, but retain all the other laws of arithmetic, we are left with the axioms for a *skew-field* or *division-algebra*. (We shall use the term skew-field.) Note that with this definition, fields (ie commutative fields) are also skew-fields.

The most familiar example of a non-commutative skew-field is furnished by the *quaternions*

$$\mathbb{H} = \{t + xi + yj + zk : t, x, y, z \in \mathbb{R}\},$$

with multiplication defined by

$$i^2 = j^2 = k^2 = -1,\ ij = -ji = k,\ jk = -kj = i,\ ki = -ik = j.$$

In fact one can show that the only finite-dimensional skew-fields over $\mathbb{R}$ are: $\mathbb{R}$ itself, the complex numbers $\mathbb{C}$, and the quaternions $\mathbb{H}$.

Theorem 6. *Every finite skew-field is commutative.*

Proof. Suppose S is a finite skew-field. Let F be the *centre* of S, ie

$$F = \{z \in S : zs = sz \text{ for all } s \in S\}.$$

We have to prove in effect that $F = S$.

To this end we assume that $F \neq S$; we shall show that this leads to a contradiction. We do this by 'counting conjugates' in the multiplicative group

$$S^\times = S - \{0\}.$$

Let

$$\|F\| = q = p^m.$$

Just as in the commutative case, in Chapter 3, we can regard S as a vector space over F. As there, we deduce that

$$\|S\| = \|F\|^n,$$

where $n = \dim_F S$.

Recall that 2 elements h, k of a finite group G are said to be conjugate (and we write $h \sim k$) if there is an element $g \in G$ such that

$$k = ghg^{-1}.$$

Conjugacy is an equivalence relation; so G is partitioned into *conjugacy classes*.

Lemma 10. *Suppose G is a finite group; and suppose $g \in G$. Then the number of elements conjugate to g is*

$$\frac{\|G\|}{\|Z(g)\|},$$

where

$$Z(g) = \{z \in G : zg = gz\}$$

Proof of Lemma. Each element $x \in G$ defines a conjugate xgx^{-1} of g. We shall see that each conjugate arises just $\|Z(g)\|$ times in this way.

Suppose $h \sim g$, say

$$h = x_0 g x_0^{-1}.$$

Then

$$\begin{aligned} xgx^{-1} = h = x_0 g x_0^{-1} &\iff x_0^{-1} x g = g x_0^{-1} x \\ &\iff x_0^{-1} x \in Z(g) \\ &\iff x \in x_0 Z(g). \end{aligned}$$

Thus just $\|Z(s)\|$ elements $x \in G$ give rise to $h \sim g$. Since this holds for each conjugate of g, the number of conjugates is

$$\frac{\|G\|}{\|Z(g)\|}.$$

□

We apply this result with $G = S^\times$.

Lemma 11. *Suppose $s \in S$. Then*

$$Z(s) = \{z \in S : zs = sz\}$$

is a sub-skew-field of S.

Corollary 6. *With the same notation,*

$$\|Z(s)\| = q^d$$

for some $d \mid n$

Proof of Lemma. Regarding $Z(s)$ as a skew-field over F, we see that

$$\|Z(s)\| = q^d$$

If $s = 0$ the result is trivial. Suppose not; then $s \in S^\times$, and $Z(s)^\times$ is a subgroup of $S^\times$. Hence, by Lagrange's Theorem,

$$q^d - 1 \mid q^n - 1.$$

As we have already seen, this implies that

$$d \mid n.$$

(For on dividing n by d, say $n = md + r$ (where $0 \le r < d$), we have

$$q^n - 1 = q^{md+r} - 1 = q^r(q^m d - 1) + (q^r - 1).$$

Thus

$$q^d - 1 \mid q^n - 1,\ q^d - 1 \mid q^{md} - 1 \Longrightarrow q^d - 1 \mid q^r - 1.$$

But that is impossible unless $r = 0$, since $q^d - 1 > q^r - 1$.) □

Proof of Lemma. We can regard S as a vector space over the *skew-field* $Z(s)$. Usually we consider linear algebra over a *commutative* field; but the fundamental theory—the notions of dimension and basis—extends to vector spaces over a skew-field. In particular, if

$$\dim_{Z(s)} S = e$$

then

$$q^n = \|S\| = \|Z(s)\|^e = q^{de},$$

and so $n = de$, ie

$$d \mid n.$$

□

Lemma 12. *The number of elements conjugate to $s \in S^\times$ is*

$$\frac{q^n - 1}{q^d - 1}$$

for some $d \mid n$.

Proof of Lemma. The number of elements conjugate to s is

$$\frac{\|S^\times\|}{\|Z(s)^\times\|} = \frac{q^n - 1}{q^d - 1}$$

by our last result. □

An element $s \in S^\times$ lies in a conjugacy class by itself if and only if $s \in F^\times$. Thus there are just $q - 1$ such elements. Each of the remaining conjugacy classes contains

$$\frac{q^n - 1}{q^d - 1}$$

elements, for some $d \mid n$ ($d \neq n$).

So counting the elements in the various conjugacy classes gives an equation of the form

$$q^n - 1 = q - 1 + \frac{q^n - 1}{q^{d_1} - 1} + \frac{q^n - 1}{q^{d_2} - 1} + \cdots .$$

We are going to show that all the fractions

$$\frac{q^n - 1}{q^d - 1}$$

share a common factor $f > 1$, which also divides $q^n - 1$. It will follow that

$$f \mid q - 1.$$

But that, as we shall see, is impossible since $f > q$. We thus arrive at a contradiction.

Definition 8. *Suppose n is a positive integer. Let*

$$\omega = e^{\frac{2\pi i}{n}}.$$

Then the cyclotomic polynomial $C_n(x)$ *is defined to be*

$$C_n(x) = \prod_{0<i<n,\ \gcd(i,n)=1} (x - \omega^i).$$

Thus $C_n(x)$ is a polynomial of degree $\phi(n)$ (where $\phi(n)$ is Euler's function).

Lemma 13. *For each $n > 0$,*

$$x^n - 1 = \prod_{d|n} C_d(x).$$

Proof of Lemma. We know that

$$x^n - 1 = \prod_{0 \le i < n} (x - \omega^i).$$

We divide the factors $x - \omega^i$ according to the value of $\gcd(i, n)$.

Suppose $n = de$. Then

$$\gcd(i, n) = d \iff i = dj,\ \gcd(j, e) = 1,\ 0 \le j < e.$$

Thus

$$\prod_{\gcd(i,n)=d, 0 \le i < n} (x - \omega^i) = \prod_{\gcd(j,e)=1, 0 \le j < e} (x - \sigma^j),$$

where

$$\sigma = \omega^d = e^{\frac{2\pi i}{e}}.$$

In other words,

$$\prod_{\gcd(i,n)=d, 0 \le i < n} (x - \omega^i) = C_e(x).$$

We conclude that

$$x^n - 1 = \prod_{d|n} C_{\frac{n}{d}}(x).$$

Since $\frac{n}{d}$ runs over the factors of n as d does, we can rewrite our last result as

$$x^n - 1 = \prod_{d|n} C_d(x).$$

□

Corollary 7. *The cyclotomic polynomial $C_n(x)$ has integer coefficients.*

Proof of Lemma. We argue by induction on n. Suppose the result true of $C_m(x)$ for all $m < n$.

We have

$$C_n(x) = \frac{x^n - 1}{\prod_{d|n, d \ne n} C_d(x)}.$$

Each cyclotomic polynomial is evidently *monic*, ie has leading coefficient 1. But if we divide $f(x)$ by $g(x)$, where both $f(x)$ and $g(x)$ have integer coefficients and $g(x)$ is monic, say

$$f(x) = q(x)g(x) + r(x) \quad (\deg r(x) < \deg g(x)),$$

then both $q(x)$ and $r(x)$ have integer coefficients. (This is clear if we derive $q(x)$ and $r(x)$ by repeatedly reducing the degree of $f(x)$ by subtracting terms of the form $ax^r g(x)$.)

Since by our inductive hypothesis the factors $C_d(x)$ have integer coefficients, and each is monic, the same is true of their product. Hence $C_n(x)$, as the quotient of $x^n - 1$ by this product, also has integer coefficients. □

Example 6. We have

$$\begin{aligned}
C_1(x) &= x-1\\
C_2(x) &= \frac{x^2-1}{x-1} = x+1\\
C_3(x) &= \frac{x^3-1}{x-1} = x^2+x+1\\
C_4(x) &= \frac{x^4-1}{C_1(x)C_2(x)} = \frac{x^4-1}{(x-1)(x+1)} = x^2+1\\
C_5(x) &= \frac{x^5-1}{x-1} = x^4+x^3+x^2+x+1\\
C_6(x) &= \frac{x^6-1}{(x-1)C_2(x)C_3(x)} = \frac{x^6-1}{(x-1)(x+1)(x^2+x+1} = \frac{x^6-1}{(x+1)(x^3-1)} = x^2-x+1.
\end{aligned}$$

Lemma 14. *If $d \mid n$ $(d \neq n)$ then*

$$C_n(q) \mid \frac{q^n-1}{q^d-1}.$$

Proof of Lemma. Let

$$\frac{x^n-1}{x^d-1} = f(x).$$

Then $f(x)$ has integer coefficients.

We know that

$$C_n(x) \mid f(x).$$

It follows on substituting $x = q$ that

$$C_n(q) \mid f(q),$$

ie

$$C_n(q) \mid \frac{q^n-1}{q^d-1}$$

for all $d \mid n$ $(d \neq n)$. □

Thus we see that the number of elements in each conjugacy class in $S^\times - F^\times$ is divisible by

$$f = C_n(q).$$

Since $C_n(q) \mid q-1$, we conclude that

$$C_n(q) \mid q-1.$$

But

$$C_n(q) = \prod_{\gcd(i,n)=1} (q-\omega^i)$$

and so

$$|C_n(q)| = \prod |q-\omega^i| \geq (q-1)^{\phi(n)},$$

since

$$|q-\omega^i| \geq q-1.$$

Moreover there is equality only if each factor is $q-1$, which is the case only if $n = 1$. Thus if $n \neq 1$,

$$|C_n(q)| > q-1.$$

But this contradicts our assertion that

$$C_n(q) \mid q - 1.$$

We conclude that our original hypothesis is untenable, ie $F = S$, and so S is commutative. $\square$

Summary: There are no 'finite quaternions'; every finite skew-field is commutative.

Chapter 12

Existence of $\mathbb{F}_{p^n}$

WE CLAIMED at the outset that *finite fields abound in nature.* But to date, apart from the prime fields $\mathbb{F}_p$ we knew about anyway, we have only come up with a couple of trivial examples. It is surely time to 'put up or shut up'!

We shall see in the next two Sections how finite fields can arise in two contexts— in group representation theory, and in the theory of algebraic numbers.

Unfortunately, it does not seem possible in either case to prove that every finite field arises in this way, without delving deeply into one theory or the other.

We are therefore forced—against out natural inclination—to demonstrate the existence of $\mathbb{F}_{p^n}$ by *construction.* This we carry out in Sections 3 and 4.

So for the first 2 Sections that follow, we are in 'waffle mode'. The results we quote are not required for the theory that follows; and the discussion can be ignored without danger.

12.1 Looking for $\mathbb{F}_{p^n}$: 1. Among group representations

How do fields—or skew-fields— arise 'naturally'?

One way is as *the endomorphism ring of a simple object in an abelian category.* If that sounds highfalutin, a concrete example should make it clearer.

Suppose G is a finite group. Recall that a *representation* α of G in a finite-dimensional vector space V is defined by an action of G on V, ie a map

$$G \times V \to V : (g, v) \mapsto gv,$$

satisfying the conditions:

1. $(gh)v = g(hv)$;
2. $ev = v$;
3. $g(u + v) = gu + gv$;
4. $g(\lambda v) = \lambda(gv)$.

We say that the space V, with the action of G on it, constitutes a *G-space.* If U, V are 2 G-spaces, a map $t : U \to V$ is said to be a *G-map* if it preserves the action of G, ie

$$t(gv) = g(tv)$$

for all $v \in V$ and all $g \in G$. (The category of G-spaces and G-maps is an example of an abelian category—*abelian* because maps $u, v : U \to V$ can be *added.*)

The representation α of G in V is said to be *simple* (or *irreducible*) if no proper subspace of $U \subset V$ is stable under G, ie

$$gu \in U \text{ for all } g \in G, u \in U \Longrightarrow U = \{0\} \text{ or } V.$$

Suppose α is a representation of G in the vector space V over the field K. Then the G-maps $t : V \to V$ form a ring $E(\alpha)$, the *endomorphism-ring* of α (or V). This ring is not in general commutative, but it has an identity element 1.

Now suppose α is *simple*. In that case $E(\alpha)$ is a skew-field. For suppose $t \in E(\alpha)$, ie t is a G-map

$$t : V \to V.$$

It is readily verified that both

$$\ker t = \{v \in V : tv = 0\} \quad \text{and} \quad \operatorname{im} t = \{v \in V : v = tu \text{ for some } u \in V\}$$

are stable subspaces of V. Since α is simple, this implies that

$$\ker t = \{0\} \text{ or } V, \quad \operatorname{im} t = \{0\} \text{ or } V.$$

But $\ker t = V$ and $\operatorname{im} t = 0$ each imply that $t = 0$. Thus if $t \neq 0$,

$$\ker t = 0, \quad \operatorname{im} t = V.$$

In other words, t is both injective and surjective. This implies that t is invertible, ie there exists a G-map $u : V \to V$ such that $tu = 1$. Thus every element $t \neq 0$ in $E(\alpha)$ is invertible, ie $E(\alpha)$ is a skew-field.

This skew-field contains the scalar field:

$$K \subset E(\alpha).$$

Moreover, $E(\alpha)$ is a finite-dimensional vector space over K. For $E(\alpha)$ is a subset of the space $\hom(V, V)$ of *all* linear maps $t : V \to V$, and so

$$\dim_K E(\alpha) \leq \dim \hom(V, V) = (\dim V)^2.$$

In other words, $E(\alpha)$ *is a finite-dimensional division-algebra over* K.

In the familiar case, where $K = \mathbb{C}$ and we are dealing with representations over the complex numbers, it follows that

$$E(\alpha) = K.$$

For since $\mathbb{C}$ is algebraically-closed the only finite-dimensional division algebra over $\mathbb{C}$ is $\mathbb{C}$ itself. This is the well-known *Schur's Lemma*: In a simple representation, the only linear maps $t : V \to V$ commuting with all elements of G are the multiplies of the identity. In matrix terms,

$$TA(g) = A(g)T \text{ for all } g \in G \Longrightarrow T = \lambda I.$$

But this is far from the case if K is *finite*, say $K = P = \mathbb{F}_p$. In this case,

$$E(\alpha) = \mathbb{F}_{p^n},$$

where in general $n > 1$.

So we have a prescription for finding finite fields: Take a simple representation (of which there is an enormous choice), and determine $E(\alpha)$.

Unfortunately, this is not a practical way of showing that $\mathbb{F}_{p^n}$ exists for all prime powers p^n, as we would have to delve into representation theory too deeply. However, it is easy enough to see that *if* $\mathbb{F}_{p^n}$ exists then it arises in this way!

To see this, suppose $\mathbb{F}_q$ exists, where $q = p^n$. Consider the cyclic group

$$C_{q-1} = \mathbb{F}_q^{\times}.$$

This has a natural representation, ρ say, in $\mathbb{F}_q$, regarded as an n-dimensional vector space over $P = \mathbb{F}_p$:

$$(g, \alpha) \mapsto g\alpha \qquad (g \in \mathbb{F}_q^{\times}, \alpha \in \mathbb{F}_q).$$

It is easy to see that this representation ρ is simple. For suppose U is a proper subspace of $\mathbb{F}_q$. Choose any element $\alpha \neq 0$ in U. Then

$$\begin{aligned} g = \alpha^{-1} \in G &\implies g\alpha = 1 \in U \\ &\implies g \cdot 1 = g \in U \end{aligned}$$

for all $g \in \mathbb{F}_q^{\times}$. Thus $U = \mathbb{F}_q$.

On the other hand, suppose $\alpha \in \mathbb{F}_q$. Let

$$\mu_\alpha : \mathbb{F}_q \to \mathbb{F}_q$$

be the map defined by multiplication by α:

$$\mu_\alpha(x) = \alpha x.$$

This map evidently commutes with the action of C_{q-1}. Thus

$$\mu_\alpha \in E(\rho).$$

So we have a natural injection

$$\mathbb{F}_q \subset E(\alpha).$$

In fact it is easy to see that every endomorphism $\gamma \in E(\alpha)$ arises in this way. For

$$\gamma(1) = \alpha \implies \gamma = \mu_\alpha.$$

Thus we may say that

$$E(\alpha) = \mathbb{F}_q.$$

We have seen therefore that every finite field F does arise in this way—as the endomorphism ring of a simple representation over a prime field.

But as we have already pointed out, this is far from proving that $\mathbb{F}_{p^n}$ exists for all p^n!

12.2 Looking for $\mathbb{F}_{p^n}$: 2. In number theory

We know that

$$\mathbb{F}_p = \mathbb{Z}/(p).$$

It is natural to ask if the other non-prime finite fields can be constructed in the same way, taking some other ring in place of the integers $\mathbb{Z}$.

A concrete example will show what we mean. Let $\mathbb{Z}[i]$ denote the ring of gaussian integers:

$$\mathbb{Z}[i] = \{m + ni : m, n \in Z\}.$$

The number 2 is no longer prime in $\mathbb{Z}[i]$:

$$2 = (1+i)(1-i).$$

Each of these factors has remainder field $\mathbb{F}_2$, eg

$$Z[i]/(1+i) = \mathbb{F}_2.$$

For modulo $1+i$,

$$m + ni \equiv \begin{cases} 0 & \text{if } m+n \text{ is even,} \\ 1 & \text{if } m+n \text{ is odd.} \end{cases}$$

On the other hand, 3 remains prime in $\mathbb{Z}[i]$, and

$$\mathbb{Z}[i]/(3) = \mathbb{F}_{3^2}.$$

In fact, an odd prime p remains prime in $\mathbb{Z}[i]$ if $p \equiv 3 \bmod 4$, and splits if $p \equiv 1 \bmod 4$. Thus

$$\mathbb{Z}[i]/(p) = \mathbb{F}_{p^2} \text{ if } p \equiv 3 \bmod 4.$$

$\mathbb{Z}[i]$ is an example of an *algebraic integer ring*. These form the subject matter of *algebraic number theory*, and have been much studied—the initial impetus arising from attempts to prove Fermat's Last Theorem. It would take us too far out of our way to go into the theory here. But in fact every finite field $\mathbb{F}_{p^n}$ does arise in this way, with the proviso that instead of prime *elements* one really has to consider prime *ideals*. The 2 rings $\mathbb{Z}$ and $\mathbb{Z}[i]$ are both *principal ideal domains*, in which every ideal is of the form (θ), consisting of all multiples of some generator θ. In such a case there is no real distinction between elements and ideals. But in general the Fundamental Theorem of Arithmetic, according to which each $n \in \mathbb{N}$ is expressible in the form

$$n = p_1^{e_1} p_2^{e_2} \cdots p_r^{e_r},$$

and on which all depends, remains true only in the realm of ideals.

12.3 Extension fields

It is time to return to mathematical mode.

Proposition 14. *Suppose F is a finite field, with prime subfield P; and suppose $\alpha \in F$. Then there is a smallest subfield $K \subset F$ containing α; K consists of all elements expressible in the form $f(\alpha)$, where $f(x) \in P[x]$:*

$$K = \{f(\alpha) : f(x) \in P[x]\}.$$

Proof. Any intersection of subfields is itself a subfield. So there is certainly a smallest subfield containing α, namely the intersection K of all subfields containing α.

Clearly

$$\alpha \in K \Longrightarrow f(\alpha) \in K$$

for all polynomials $f(x) \in P[x]$.

The result will therefore follow if we can show that the set L of elements of the form $f(\alpha)$ forms a field. This set L is evidently closed under addition, subtraction and multiplication. It only remains to show that it is closed under division by non-zero elements.

Lemma 15. *A subset $S \subset G$ of a finite group G closed under multiplication is necessarily a subgroup.*

Proof of Lemma. We have to show that

$$s \in S \Longrightarrow s^{-1} \in S.$$

By Lagrange's Theorem,

$$s^n = 1,$$

where $n = \|G\|$. Thus

$$s^{-1} = s^{n-1} = s \cdot s \cdots s \in S.$$

$\square$

Applying this result to the subset $L^\times = L - \{0\}$, we conclude that $L^\times$ is a subgroup of $F^\times$. But this implies that L is a subfield of F, and so $L = K$. $\square$

Proposition 15. *If the minimal polynomial of $\alpha \in F$ has degree d, then the smallest subfield $K \subset F$ containing α contains p^d elements:*

$$K = \mathbb{F}_{p^d}.$$

Proof.

Lemma 16. *Suppose F is a finite field, with prime subfield P; and suppose $\alpha \in F$ has minimal polynomial $m(x) \in P[x]$. If $f(x), g(x) \in P[x]$ then*

$$f(\alpha) = g(\alpha) \Longleftrightarrow m(x) \mid f(x) - g(x).$$

Proof of Lemma. This follows at once from the fact that

$$f(\alpha) = 0 \Longleftrightarrow m(x) \mid f(x).$$

$\square$

Suppose $f(x) \in P[x]$. Dividing $f(x)$ by $m(x)$,

$$f(x) = m(x)q(x) + r(x) \qquad (\deg r(x) < d).$$

By the Lemma,

$$f(\alpha) = r(\alpha).$$

Thus each element $\beta \in K$ is expressible in the form

$$\beta = c_0 + c_1\alpha + c_2\alpha^2 + \cdots + c_{d-1}\alpha^{d-1} \qquad (c_i \in P).$$

Furthermore, the Lemma shows that this expression is unique.

It follows that the d elements

$$1, \alpha, \alpha^2, \ldots, \alpha^{d-1}$$

form a basis for K as a vector space over P. Thus

$$\dim_P K = d,$$

and so

$$\|K\| = p^d.$$

$\square$

Suppose $p(x) \in P[x]$ is a prime polynomial. If there *is* an extension field $K \supset P$ containing a root α of $p(x)$, the Proposition above shows us how it is constructed. Now we must try to turn analysis into synthesis.

Theorem 7. *Suppose $m(x) \in P[x]$ is a prime polynomial over the prime field $P = \mathbb{F}_p$. Then there exists a finite field F of characteristic p, and an element $\alpha \in F$, such that α has minimal polynomial $m(x)$;*

Proof. The last proposition tells us how to construct F. Each polynomial $f(x) \in P[x]$ defines an element $\bar{f} \in F$; and

$$\bar{f} = \bar{g} \iff m(x) \mid f(x) - g(x) \qquad (f(x), g(x) \in P[x]).$$

In other words

$$F = P[x]/\left(m(x)\right).$$

More precisely, the elements of F consist of the equivalence classes in $P[x]$ under the equivalence relation

$$f(x) \equiv g(x) \quad \text{if} \quad m(x) \mid f(x) - g(x).$$

Addition, subtraction and multiplication of elements of F, ie classes in $P[x]$, is defined by taking representatives of these classes, adding, subtracting or multiplying them, and returning the class of the result. The resulting classes are independent of the choice of representatives, since eg

$$m(x) \mid f_1(x) - f_2(x),\ m(x) \mid g_1(x) - g_2(x) \Longrightarrow m(x) \mid f_1(x)g_1(x) - f_2(x)g_2(x).$$

We have constructed a *ring* $F = P[x]/\left(m(x)\right)$. It remains to show that F is in fact a *field*. Suppose $\bar{f} \in F$; and suppose $\bar{f} \neq 0$, ie

$$m(x) \nmid f(x).$$

Then

$$\gcd\left(f(x), m(x)\right) = 1;$$

for the only factor of $m(x)$ apart from 1 is $m(x)$ itself.

It follows that we can find $a(x), b(x) \in P[x]$—for example, by the euclidean algorithm—such that

$$a(x)f(x) + b(x)m(x) = 1.$$

But this implies that

$$m(x) \mid a(x)f(x) - 1,$$

and so by definition

$$\bar{a}\bar{f} = 1.$$

Thus $\bar{f}$ is invertible in F. Therefore F is a field.

Consider the polynomial

$$i(x) = x,$$

and the corresponding element $\bar{i} \in F$. From the definition of multiplication in F,

$$\bar{i}^2 = \overline{x^2},\ \bar{i}^3 = \overline{x^3},$$

etc. More generally, for any polynomial $f(x) \in P[x]$,

$$f\left(\bar{i}\right) = \overline{f(x)}.$$

In particular

$$m\left(\bar{i}\right) = \overline{m(x)} = 0.$$

Thus $m(x)$ has a root in the field F, namely $\bar{i}$. $\square$

12.4 Constructing $\mathbb{F}_{p^n}$

Theorem 8. *For each prime power p^n there exists a field $\mathbb{F}_{p^n}$ containing p^n elements.*

Proof. We may assume by induction that there exist fields $\mathbb{F}_{p^m}$ containing p^m elements for all $m < n$.

Consider the prime factorisation of the universal polynomial

$$U_n(x) \equiv x^{p^n} - x = f_1(x)f_2(x)\dots f_r(x)$$

over the prime field $P = \mathbb{F}_p$.

Lemma 17. *At least one of the factors $f(x) = f_i(x)$ is not a factor of $U_m(x)$ for any $m \mid n$:*

$$f(x) \mid x^{p^n} - x, \quad f(x) \not| \, x^{p^m} - x$$

if $m \mid n$, $m < n$.

Proof of Lemma. Notice that if $\mathbb{F}_{p^n}$ exists, and α is any element of $\mathbb{F}_{p^n}$ not in any proper subfield—for example, α could be any primitive element of $\mathbb{F}_{p^n}$—then its minimal polynomial $m(x)$ will satisfy the lemma.

But without this assumption, a crude counting argument suffices. For each $m \mid n$, the sum of the degrees of all prime polynomials dividing $U_m(x)$ cannot exceed the degree of $U_m(x)$:

$$\sum_{f(x) \mid U_m(x)} \deg f(x) \le p^m.$$

Summing this over all factors m of n apart from n itself,

$$\begin{aligned} \sum_{f(x) \mid U_m(x),\, m \mid n,\, m<n} \deg f(x) &\le \sum_{m \mid n,\, m<n} p^m \\ &\le \sum_{m<n} p^m \\ &= \frac{p^n - 1}{p - 1} \\ &< p^n. \end{aligned}$$

since

$$\sum_{f(x) \mid U_n(x)} \deg f(x) = p^n,$$

it follows that at least one factor $f(x)$ of $U_n(x)$ divides no $U_m(x)$. □

Now let $F = P[x]/f(x)$ be the field extension corresponding to such a factor $f(x)$, as defined above. Then $F = \mathbb{F}_{p^m}$ for some m. We must show that $m = n$.

As we saw, F contains an element α satisfying $f(\alpha) = 0$. Since $f(x) \mid U_n(x)$ it follows that

$$U_n(\alpha) = 0.$$

On the other hand, since $\alpha \in \mathbb{F}_{p^m}$,

$$U_m(\alpha) = 0.$$

It follows that α satisfies

$$\gcd\left(U_n(x), U_m(x)\right) = U_d(x),$$

where $d = \gcd(n, m)$. Now $f(x)$ is the minimal polynomial of α. Hence

$$f(x) \mid U_d(x).$$

Since $d \mid n$, this contradicts the defining property of $f(x)$, unless $d = n$.

But now we have constructed a field $F = \mathbb{F}_{p^m}$, where $n \mid m$. It follows that F has a subfield $\mathbb{F}_{p^n}$ containing p^n elements. □

> **Summary:** There exists one and only one field $\mathbb{F}_{p^n}$ containing p^n elements, for each prime power p^n.

Chapter 13

Prime Polynomials over a Prime Field

AS WE HAVE SEEN (particularly in the last chapter), there is an intimate relation between the finite fields $\mathbb{F}_{p^n}$ of characteristic p and polynomials—in particular prime polynomials—over the prime field $P = \mathbb{F}_p$. The following result summarises the relation.

Proposition 16. *Suppose $\alpha \in \mathbb{F}_{p^n}$. Then the minimal polynomial of α over $P = \mathbb{F}_p$ is a prime polynomial of degree $d \mid n$.*

Conversely, if $p(x)$ is a prime polynomial of degree d in $P[x]$ then the roots of $p(x)$ lie in $\mathbb{F}_{p^n}$ if and only if $d \mid n$.

Proof. Let $m(x)$ be the minimal polynomial of $\alpha \in \mathbb{F}_{p^n}$. Suppose $\deg m(x) = d$. Let K be the smallest subfield containing α. Then

$$\dim_P K = d,$$

where $P = \mathbb{F}_p$. In other words,

$$K = \mathbb{F}_{p^d}.$$

But since $K \subset \mathbb{F}_{p^n}$ this implies that

$$d \mid n.$$

Conversely, suppose $p(x)$ is a prime polynomial of degree d over P. By the construction of an algebraic extension in the last Chapter, we can find a field $F \supset P$ in which $p(x)$ has a root α. (In fact, as we saw, this means that $p(x)$ factorises completely in F.)

Let K be the smallest subfield containing α. As we just saw

$$K = \mathbb{F}_{p^d}.$$

It follows that α satisfies the universal equation

$$U_d(x) = 0.$$

Hence

$$p(x) \mid U_d(x).$$

But if $d \mid n$,

$$U_d(x) \mid U_n(x)$$

Thus

$$p(x) \mid U_n(x),$$

and so $p(x)$ factorises completely in $\mathbb{F}_{p^n}$. $\square$

Corollary 8. *Let*

$$U(x) \equiv x^{p^n} - x$$

over $P = \mathbb{F}_p$. Then the prime factorisation of $U(x)$ takes the form

$$U(x) = \prod_{\deg m(x) | n} m(x),$$

where $m(x)$ runs over all prime polynomials of degree $d \mid n$ over P.

Corollary 9. *If $\Pi(n) = \Pi_p(n)$ denotes the number of prime polynomials of degree n over the prime field $P = \mathbb{F}_p$, then*

$$\sum_{d|n} d\Pi(d) = p^n.$$

Proof. This follows from the previous Corollary on comparing degrees. □

Corollary 10. *The number of prime polynomials of degree n over $P = \mathbb{F}_p$ is given by*

$$\Pi(n) = \frac{1}{n} \sum_{d|n} \mu\left(\frac{n}{d}\right) d^n,$$

where $\mu(n)$ is Möbius' function:

$$\mu(n) = \begin{cases} 0 & \text{if } n \text{ has a repeated prime factor} \\ (-1)^e & \text{if } n \text{ has } e \text{ distinct prime factors.} \end{cases}$$

Proof. This follows from the previous Corollary on applying Möbius' inversion formula:

$$F(n) = \sum_{d|n} f(d) \Longrightarrow f(n) = \sum_{d|n} \mu\left(\frac{n}{d}\right) F(d).$$

□

Example 7. The number of prime polynomials of degree 6 over $P = \mathbb{F}_2$ is

$$\begin{aligned} \Pi(6) = \Pi_2(6) &= \frac{1}{6}\left(\mu(1)2^6 + \mu(2)2^3 + \mu(3)2^2 + \mu(6)2^1\right) \\ &= \frac{1}{6}\left(2^6 - 2^3 - 2^2 + 2^1\right) \\ &= \frac{54}{6} \\ &= 9. \end{aligned}$$

In determining these 9 polynomials, note that if

$$p(x) = c_0 + c_1 x + c_2 x^2 + c_3 x^3 + c_4 x^4 + c_5 x^5 + c_6 x^6$$

is a prime polynomial of degree 6 then

1. The first and last coefficients c_0 and c_6 must not vanish:

$$c_0 = c_6 = 1.$$

2. The sum of the coefficients must be non-zero, or else $1+x$ would divide $p(x)$:

$$c_0 + c_1 + c_2 + c_3 + c_4 + c_5 + c_6 = 1.$$

This leaves just 16 possibilities.

Suppose $p(x)$ is prime. Then so is the polynomial obtained by taking the coefficients in reverse order

$$\tilde{p}(x) = x^6 p\left(\frac{1}{x}\right).$$

For it is easy to verify that if $p(x)$ factorises so does $\tilde{p}(x)$.

So the prime polynomials of degree 6 occur in pairs, except for those which are 'symmetrical', ie $\tilde{p}(x) = p(x)$. There are just 4 symmetrical polynomials among the 16 we are examining, namely

$$1+x^3+x^6, 1+x+x^3+x^5+x^6, 1+x^2+x^3+x^4+x^6, 1+x+x^2+x^3+x^4+x^5+x^6.$$

So 1 or 3 of these is prime; and then 4 or 3 of the remaining 6 pairs are prime.

If one of our 16 polynomials is *not* prime then it must have a prime factor of degree 2 or 3. (We have excluded the prime factors x and $1+x$ of degree 1.)

The number of prime polynomials of degree 2 is

$$\Pi(2) = \frac{1}{2}\left(2^2 - 2^1\right) = 1,$$

while

$$\Pi(3) = \frac{1}{3}\left(2^3 - 2^1\right) = 2.$$

Given that a prime polynomial of degree > 1 over $\mathbb{F}_2$ must have an odd number of non-zero coefficients, as we remarked above, we see that the prime of degree 2 must be

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

while the two primes of degree 3 are

$$r(x) = 1 + x + x^3,\ s(x) = 1 + x^2 + x^3.$$

If one of our 16 polynomials is not prime, then it is either divisible by $q(x)$ or else it is the product of 2 primes of degree 2, ie it is one of

$$\begin{aligned} r(x)^2 &= 1 + x^2 + x^6, \\ s(x)^2 &= 1 + x^4 + x^6, \\ r(x)s(x) &= 1 + x + x^2 + x^3 + x^4 + x^5 + x^6. \end{aligned}$$

It is easy to see if a polynomial $p(x)$ is divisible by $q(x)$, since

$$x^3 - 1 = (x-1)(x^2 + x + 1),$$

and therefore

$$x^3 \equiv 1 \bmod q(x).$$

So, for example,

$$p(x) = x^6 + x^5 + x^2 + x + 1 \equiv 1 + x^2 + x^2 + x + 1 \equiv x \bmod q(x),$$

and so

$$q(x) \not| p(x).$$

Dividing the 4 symmetrical polynomials of degree 6 by each of these in turn, we see that just 1 is prime, namely the first:

$$1 + x^3 + x^6.$$

Thus just 4 out of the 6 pairs of asymmetric polynomials are prime. We can exclude the pair

$$(1+x+x^3)^2 = 1+x^2+x^6,\ (1+x^2+x^3)^2 = 1+x^4+x^6.$$

Just one non-prime pair more to go!

It is evident that

$$(1+x+x^2)^3,\ (1+x+x^3)(1+x^2+x^3)$$

are both symmetric. It follows that the last non-prime of degree 6 must be the product of $1+x+x^2$ and a prime of degree 4.

Now

$$\Pi(4) = \frac{1}{4}\left(2^4 - 2^2\right) = 3.$$

The 3 prime polynomials of degree 4 are

$$1+x+x^4,\ 1+x^3+x^4,\ 1+x+x^2+x^3+x^4.$$

So our last non-prime pair is

$$(1+x+x^2)(1+x+x^4) = 1+x^3+x^4+x^5+x^6$$

and its 'conjugate'

$$(1+x+x^2)(1+x^3+x^4) = 1+x+x^2+x^3+x^6.$$

So if we represent the polynomial by its coefficients as a sequence of bits,

$$c_0 + c_1x + c_2x^2 + c_3x^3 + c_4x^4 + c_5x^5 + c_6x^6 \longleftrightarrow (c_0c_1c_2c_3c_4c_5c_6),$$

then our 9 prime polynomials of degree 6 are

$$\begin{array}{cc} (1000011) & (1100001) \\ (1000101) & (1010001) \\ \multicolumn{2}{c}{(1001001)} \\ (1001011) & (1101001) \\ (1001101) & (1011001) \end{array}$$

Definition 9. *Suppose $p(x)$ is a prime polynomial of degree d over $P = \mathbb{F}_p$. Let α be a root of $p(x)$ in $\mathbb{F}_{p^d}$. Then $p(x)$ is said to be* primitive *if α is primitive.*

Proposition 17. *Suppose $p(x)$ is a prime polynomial of degree d over $P = \mathbb{F}_p$; and suppose $\alpha \in \mathbb{F}_{p^d}$ is a root of $p(x)$. Then the order of α in $\mathbb{F}_{p^d}^\times$ is equal to the order of x modulo $p(x)$, ie the least integer $e > 0$ such that*

$$p(x) \mid x^e - 1.$$

Proof. Suppose

$$\alpha^e = 1.$$

Then α satisfies the equation

$$x^e - 1 = 0.$$

But $p(x)$ is the minimal polynomial of α. hence

$$p(x) \mid x^e - 1,$$

or in other words,

$$x^e \equiv 1 \bmod p(x).$$

Conversely,

$$\begin{aligned} x^e \equiv 1 \bmod p(x) &\implies p(x) \mid x^e - 1 \\ &\implies \alpha^e - 1 = 0 \\ &\implies \alpha^e = 1. \end{aligned}$$

□

Corollary 11. *With the same notation, the order of x modulo $p(x)$ divides $p^d - 1$.*

Corollary 12. *Suppose $p(x)$ is a prime polynomial of degree d over P. Then $p(x)$ is primitive if and only if x has order $p^d - 1$ modulo $p(x)$.*

Proposition 18. *The number of primitive polynomials of degree d is*

$$\frac{\phi(p^d - 1)}{d},$$

where $\phi(n)$ denotes Euler's function.

Proof.

Lemma 18. *If $\alpha \in \mathbb{F}_q$ is primitive, then so are all its conjugates*

$$\alpha, \Phi\alpha, \Phi^2\alpha, \ldots.$$

Proof of Lemma. Suppose $\Phi\alpha$ is not primitive. In other words $\Phi\alpha$ had degree $d < q - 1$. Then

$$\begin{aligned} (\Phi\alpha)^d = 1 &\implies \Phi\left(\alpha^d\right) = 1 \\ &\implies \alpha^d = 1, \end{aligned}$$

since Φ is an automorphism. □

There are $\phi(p^d - 1)$ primitive elements in $\mathbb{F}_{p^d}$. Each primitive polynomial $p(x)$ of degree d has d of these elements as roots. Thus the number of such polynomials is

$$\frac{\phi(p^d - 1)}{d}.$$

□

Example 8. The number of primitive polynomials of degree 6 over $\mathbb{F}_2$ is

$$\begin{aligned} \frac{\phi(2^6 - 1)}{6} &= = \frac{\phi(63)}{6} \\ &= \frac{\phi(3^2)\phi(7)}{6} \\ &= \frac{3 \cdot 2 \cdot 6}{6} \\ &= 6. \end{aligned}$$

So of our 9 prime polynomials of degree 6, just 6 are primitive and 3 non-primitive.

It is a straightforward matter to establish that if $p(x)$ is primitive then so is its 'conjugate' $\tilde{p}(x)$. (We leave the proof of this to the reader.) So it follows that our symmetric prime of degree 6 *cannot* be primitive (or there would be an odd number of primitive polynomials). Let us verify this.

It is sufficient, as we have seen, to determine the order of x modulo $p(x)$. If $p(x)$ is primitive this will be $2^6 - 1 = 63$. In any case, it will be a factor of 63.

Taking

$$p(x) = 1 + x^3 + x^6,$$

we have

$$x^6 \equiv x^3 + 1 \bmod p(x),$$

and so

$$\begin{aligned} x^9 &\equiv x^6 + x^3 \\ &\equiv 1. \end{aligned}$$

Thus x has order 9 modulo $p(x)$, and so $p(x)$ is not primitive. (We've actually shown that the order *divides* 9; but since the order of x modulo $p(x)$ is manifestly greater than the degree of $p(x)$, the order must in fact *be* 9.)

We leave it to the student to determine which of the 4 pairs of asymmetric primes is *not* primitive.

Summary: The prime polynomials over the $P = \mathbb{F}_p$ divide into 2 classes: primitive and non-primitive. We are able to compute both the number of prime polynomials, and the number of primitive polynomials, of a given degree.

Appendix A

Galois Theory

A.1 The Galois Correspondence

Definition 10. *Suppose G is finite group of automorphisms of the field K. Let k be the set of fixed elements under G:*

$$k = \{\theta \in K : g\theta = \theta \textit{ for all } g \in G\}.$$

Then we say that K is a galois extension *of k.*

We shall show that in this case

1. k is a subfield of K;
2. $\deg_k K$ is finite;
3. G is the full group of automorphisms of K over k:

$$G = \operatorname*{Aut}_k K.$$

It will follow in particular from this that if K is a galois extension of k then we can take $G = \mathrm{Aut}_k K$; so the property depends only on K and k (and not on G).

Examples 1.

1. The finite field

$$K = \mathcal{F}_(p^n)$$

is a galois extension of $\mathcal{F}_(p)$, with

$$G = \{I, \Phi, \Phi^2, \ldots, \Phi^{n-1}\},$$

where Φ is the Frobenius automorphism $x \mapsto x^p$.

2. The Gaussian rationals

$$K = \mathbb{Q}(i),$$

ie the field of complex numbers of the form $x + yi$, where $x, y \in \mathbb{Q}$, is a galois extension of $\mathbb{Q}$, with

$$G = \{I, C\},$$

where C is complex conjugation $x + yi \mapsto x - yi$.

3. The quadratic number field

$$K = \mathbb{Q}(\sqrt{2}),$$

ie the field of real numbers of the form $x + y\sqrt{2}$, where $x, y \in \mathbb{Q}$, is a galois extension of $\mathbb{Q}$, with

$$G = \{I, J\},$$

where J is the map $x + y\sqrt{2} \mapsto x - y\sqrt{2}$.

4. The cyclotomic field

$$K = \mathbb{Q}(\omega),$$

where $\omega = e^{2\pi i/n}$, is a galois extension of $\mathbb{Q}$; G is the group of $\phi(n)$ automorphisms of the form

$$\omega \mapsto \omega^i,$$

where $\gcd(i, n) = 1$.

Definition 11. *Suppose G is a finite group of automorphisms of K. Then For each subgroup $S \subset G$ we set*

$$\mathcal{F}(S) = \{\theta \in K : g\theta = \theta \text{ for all } g \in S\}.$$

2. *For each subfield $F \subset K$ we set*

$$\mathcal{S}(F) = \{g \in G : g\theta = \theta \text{ for all } \theta \in F\}.$$

As indicated above, we assume that

$$k = \mathcal{F}(G),$$

ie k denotes the set of elements left fixed by all the automorphisms in G,

Proposition 19. *Suppose G is a finite group of automorphisms of K. Then*

1. *For each subgroup $S \subset G$, $\mathcal{F}(S)$ is a subfield of K.*

2. *For each subfield $F \subset K$, $\mathcal{S}(F)$ is a subgroup of G.*

3. *If S is a subgroup of G then*

$$S \subset \mathcal{SF}(S);$$

4. *If F is a subfield of K then*

$$F \subset \mathcal{FS}(F);$$

5. *If S, T are subgroups of G then*

$$S \subset T \Longrightarrow \mathcal{F}(S) \supset \mathcal{F}(T);$$

6. *If E, F are subfields of K then*

$$E \subset F \Longrightarrow \mathcal{S}(E) \supset \mathcal{S}(F);$$

7. *For each subgroup $S \subset G$,*

$$\mathcal{FSF}(S) = \mathcal{F}(S).$$

In other words,

$$\mathcal{FSF} = \mathcal{F}.$$

8. *For each subfield $F \subset K$,*

$$\mathcal{SFS}(F) = \mathcal{S}(F).$$

In other words,

$$\mathcal{SFS} = \mathcal{S}.$$

Proof. All these results are immediate, except perhaps the last two.

For (7) we note that by (3)

$$S \subset \mathcal{SF}(S).$$

Hence

$$\mathcal{F}(S) \supset \mathcal{F}(\mathcal{SF}(S)),$$

by (5). On the other hand,

$$\mathcal{F}(S) \subset \mathcal{FS}(\mathcal{F}(S)),$$

on applying (4) with $\mathcal{F}(S)$ in place of F.

The last part (8) is proved similarly. □

It follows from the last 2 parts of this Proposition that if $F = \mathcal{F}(S)$, ie if F is the fixed field of some subgroup $S \subset G$, then

$$\mathcal{FS}(F) = F;$$

and similarly, if $S = \mathcal{S}(F)$, ie if S is the invariant subgroup of some subfield $F \subset K$ then

$$\mathcal{SF}(S) = S.$$

We shall show that *every* field F between k and K is the fixed field of some subgroup, and *every* subgroup $S \subset G$ is the invariant group of some subfield.

It will follow from this that the mappings

$$S \mapsto \mathcal{F}(S), \; F \mapsto \mathcal{S}(F)$$

establish a one-one correspondence between the subgroups of G and the subfields of K containing k. That is the Fundamental Theorem of Galois Theory.

A.2 Towers of Extensions

Proposition 20. *Suppose F is a subfield of K containing k,*

$$k \subset F \subset K;$$

and suppose $\deg_k F$ and $\deg_F K$ are both finite. Then $\deg_k K$ is finite, and

$$\deg_k K = \deg_k F \cdot \deg_F K.$$

Proof. Let $\{\epsilon_1, \ldots, \epsilon_r\}$ be a basis for F over k; and let $\{\eta_1, \ldots, \eta_s\}$ be a basis for K over F. Then the rs elements

$$\epsilon_i \eta_j \quad (1 \le i \le r,\ 1 \le j \le s)$$

form a basis for K over k.

For any $\theta \in K$ is uniquely expressible in the form

$$\theta = \sum_{1 \le j \le s} \xi_j \eta_j,$$

with $\xi_1, \ldots, \xi_s \in F$. But now each ξ_j is uniquely expressible in terms of the ϵ_i:

$$\xi_j = \sum_{1 \le i \le r} a_{ij} \epsilon_i,$$

where $a_{ij} \in k$, giving

$$\theta = \sum_{i,j} a_{ij} \epsilon_i \eta_j,$$

□

A.3 Algebraic Extensions

Recall that an element $\theta \in K$ is said to be *algebraic* over the subfield k if it satisfies a polynomial equation

$$x^n + c_1 x^{n-1} + \cdots + c_n = 0$$

with coefficients $c_i \in k$.

We say that K is an *algebraic extension* of k if every element $\theta \in K$ is algebraic over k. The algebraic extension K over k is said to be *simple* if

$$K = k(\alpha)$$

for some $\alpha \in K$. If this is so, and $m(x)$ is the minimal polynomial of α over k then

$$\deg_k K = \deg m(x),$$

with each element $\theta \in K$ uniquely expressible in the form

$$\theta = c_0 + c_1 \alpha + \cdots + c_{d-1} \alpha^{d-1},$$

where $d = \deg m(x)$.

Proposition 21. *An extension of finite degree is necessarily algebraic.*

Proof. Suppose $\deg_k K = d$; and suppose $\theta \in K$. The $d + 1$ elements

$$1, \theta, \theta^2, \ldots, \theta^d$$

must be linearly dependent over k, ie we can find $c_0, c_1, \ldots, c_d \in k$ such that

$$c_0 + c_1 \theta + \cdots + c_d \theta^d = 0.$$

In other words θ is a root of the polynomial

$$c_0 + c_1 x + \cdots + c_d x^d = 0.$$

□

Corollary 13. *If θ is algebraic over k then the extension $k(\theta)$ is algebraic.*

A.4 Conjugacy

We suppose in this Section that G is a finite group of automorphisms of the field K, and that $k = \mathcal{F}(G)$.

Definition 12. *Suppose $\theta \in K$. Then the elements $g\theta (g \in G)$ are called the* conjugates *of θ.*

The argument used in the proof of the following Proposition is frequently encountered in galois theory.

Proposition 22. *Suppose $\theta \in K$. Let the distinct conjugages of θ be*

$$\theta = \theta_1, \theta_2, \ldots, \theta_d;$$

Then the minimal polynomial of θ is

$$m(x) = (x - \theta_1) \cdots (x - \theta_d).$$

Proof. Consider the action of the automorphism $g \in G$ on $m(x)$. It is easy to see that g simply permutes the factors of $m(x)$:

$$\begin{aligned} m^g(x) &= (x - g\theta_1) \cdots (x - g\theta_d) \\ &= (x - \theta_1) \cdots (x - \theta_d) \\ &= m(x). \end{aligned}$$

It follows that the coefficients of $m(x)$ are invariant under all $g \in G$, and so lie in the groundfield k:

$$m(x) \in k[x].$$

Thus $m(x)$ is a polynomial over k satisfied by θ. If $M(x)$ is the minimal polynomial of θ, therefore,

$$M(x) \mid m(x).$$

But on applying the automorphism $g \in G$

$$M(\theta) = 0 \Longrightarrow M(g\theta) = 0,$$

since g leaves the coefficients of $M(x)$ fixed. Thus every conjugate θ_i of θ is a factor of $M(x)$, and so

$$m(x) \mid M(x).$$

Hence $M(x) = m(x)$, ie $m(x)$ is the minimal polynomial of θ. □

Corollary 14. *If $\theta \in K$ has d distinct conjugates then*

$$d = \deg_k k(\theta).$$

Recall that the polynomial $p(x)$ is said to be *separable* if it has distinct roots. We say that θ is separable over k if it is algebraic over k and its minimal polynomial $m(x)$ is separable; and we say that the algebraic extension F of k is separable if every element of F is separable over k.

In characteristic 0 (which is the case we are chiefly interested in), every algebraic element is separable; for if $g(x) = \gcd(m(x), m'(x))$ then $g(x) \mid m(x)$, and so $g(x) = 1$.

However, in finite characteristic p this argument may break down, since $m'(x)$ may vanish identically. This happens if (and only if) $m(x)$ contains only powers of x^p, say

$$m(x) = M(x^p).$$

In fact this cannot happen in our case; for we have seen that each element $\theta \in K$ satisfies an equation over k with *distinct* roots θ_i.

Corollary 15. *K is a separably algebraic extension of k.*

Proposition 23. *Suppose*

$$F = k(\theta),$$

where $\theta \in K$. Then

$$\deg_k F \cdot \|\mathcal{S}(F)\| = \|G\|.$$

Proof. Suppose θ has d conjugates. Then

$$\deg_k k(\theta) = d,$$

by the Corollary to the last Proposition.

On the other hand

$$\mathcal{S}(F) = \{g \in G : g\theta = \theta\};$$

for if g leaves θ fixed then it will leave every element of $k(\theta)$ fixed.

Let $S = \mathcal{S}(F)$. Then

$$\begin{aligned} g_1\theta = g_2\theta &\iff g_2^{-1}g_1\theta = \theta \\ &\iff g_2^{-1}g_1 \in S \\ &\iff g_1S = g_2S. \end{aligned}$$

This establishes a one-one correspondence between the conjugates of θ and the cosets of S. Hence the number d of conjugates is equal to the number of cosets, ie

$$d = \|G\|/\|S\|.$$

Thus

$$\deg_k k(\theta) \cdot \|S\| = d \cdot \|S| \ = \|G\|,$$

as required. □

A.5 The Correspondence Theorem

Theorem 9. *Suppose G is a finite group of automorphisms of the field K; and suppose $k = \mathcal{F}(G)$ is the field of fixed elements under G. Then*

1. *The maps*

   $$S \mapsto \mathcal{F}(S),\ F \mapsto \mathcal{S}(F)$$

   establish a one-one correspondence between subgroups $S \subset G$ and subfields $F \subset K$ containing k.

2. *If S and F correspond in this way then*

   $$\|S\| \cdot \deg_k F = \|G\|.$$

3. *In particular*

   $$\deg_k K = \|G\|.$$

4. *Each subfield F is a simple extension of k:*

   $$F = k(\theta).$$

5. *G is the full group of isomorphisms of K:*

   $$G = \operatorname*{Aut}_k K.$$

Proof. Let us assume that $\deg_k K < \infty$, as is implied by (3). We shall show at the end of the proof that this assumption is justified.

We argue by induction on G. Thus we may assume the result true for all proper subgroups $S \subset G$.

To establish the correspondence we have to show that $\mathcal{SF}(S) = S$ for every subgroup $S \subset G$, and $\mathcal{FS}(F) = F$ for every subfield $F \subset K$ containing k.

Lemma 19. *For each subgroup $S \subset G$ we have*

$$\mathcal{SF}(S) = S.$$

Proof of Lemma. This follows at once on applying our inductive hypothesis with S in place of G, and $k' = \mathcal{F}(S)$ in place of k. For the last part of the Theorem tells us that S is the *full* group of automorphisms of $\operatorname{Aut}_{k'} K$). $\square$

Lemma 20. *Suppose*

$$k \subset F, F' \subset K;$$

and suppose

$$\Theta : F \to F'$$

is an isomorphism over k. Then Θ can be extended to an automorphism of K over k.

Putting the matter the other way round, Θ is the restriction to F of some $g \in \operatorname{Aut}_k K$.

Proof of Lemma. Suppose $\theta \in K \setminus F$. Let

$$m(x) = (x - \theta_1) \cdots (x - \theta_d) = x^d + \gamma_1 x^{d-1} + \cdots + \gamma_d$$

be the minimal polynomial of θ over F.

We know that the minimal polynomial of θ over k is of the form

$$M(x) = (x - g_1\theta) \cdots (x - g_r\theta),$$

where $g_1\theta, \ldots, g_r\theta$ are the distinct conjugates of θ. Since $m(x) \mid M(X)$, we deduce that (1) the roots of $m(x)$ are distinct, and (2) these roots are all of the form $g\theta$.

Now consider the transform of $m(x)$ under Θ,

$$m^\Theta(x) \equiv x^d + (\Theta\gamma_1)x^{d-1} + \cdots + (\Theta\gamma_d).$$

Since

$$m^\Theta(x) \mid M^\Theta(x) = M(x),$$

we see that $m^\Theta(x)$ factorises completely in K.

Let θ' be any root of $m^\Theta(x)$. We extend Θ to a map

$$\Theta' : F(\theta) \to F'(\theta')$$

as follows. Suppose $\phi \in F(\theta$, say $\phi = p(\theta)$, where $p(x) \in F[x]$. Then

$$\phi = p(\theta) \mapsto \phi' = p^\Theta(\theta').$$

This is well-defined, since

$$p(\theta) = 0 \Longrightarrow m(x) \mid p(x) \Longrightarrow m^\Theta(x) \mid p^\Theta(x) \Longrightarrow p^\Theta(\theta') = 0.$$

Since Θ' clearly preserves addition and multiplication, it is an isomorphism extending Θ to $F(\theta)$.

We can extend the isomorphism repeatedly in this way to

$$F(\theta_1, \ldots, \theta_r$$

until finally we must reach K since we are assuming that $\deg_k K$ is finite.

As it stands, we only know that this extension is an endomorphism of K. However, a linear transformation $t : V \to V$ of a finite-dimensional vector space V is bijective if and only if it is injective (that is, if $\det t \neq 0$). Thus we have extended the isomorphism Θ to an automorphisms $g \in \mathrm{Aut}_k K$. □

Lemma 21. *Suppose*

$$k \subset F \subset K.$$

Then

$$\deg_k F \cdot \|\mathcal{S}(F)\| = \|G\|.$$

Proof of Lemma. We argue by induction on $\deg_k F$. Let us suppose the result holds for F; and suppose $\theta \in K \setminus F$. Let

$$m(x) = (x - \theta_1) \cdots (x - \theta_d)$$

be the minimal polynomial of θ over F. In the proof of the last Lemma we showed how to construct an isomorphism $F(\theta) \to F(\theta_i)$ for each root θ_i of $m(x)$. These isomorphisms extend — by the same Lemma — to automorphisms

$$g_1, \ldots, g_d \in \operatorname*{Aut}_F K = \mathcal{S}(F).$$

Let $S = \mathcal{S}(F)$; and suppose $g \in S$. Since g leaves $m(x)$ unchanged, $g\theta = \theta_i$ for some i. It follows that g restricts on $F(\theta)$ to one of our d isomorphisms, say the restriction of g_i. But then $g_i^{-1} g$ leaves θ fixed, and so leaves every element of $F' = F(\theta)$ fixed:

$$g_i^{-1} g \in \mathcal{S}(F') = S',$$

say. We deduce that

$$S = g_1 S' \cup \cdots \cup g_d S'.$$

Thus

$$\|S\| = \deg_F F' \cdot \|S'\|;$$

and so

$$\begin{aligned} \deg_k F' \cdot \|S'\| &= \deg_k F \cdot \deg_F F' \cdot \|S'\| \\ &= \deg_k F \cdot \|S\| \\ &= \|G\|, \end{aligned}$$

by the inductive hypothesis. □

Applying this Lemma with $F = K$,

$$\deg_k K = \|G\|,$$

since $\mathcal{S}(K) = \{e\}$.

Lemma 22. *For each subfield $F \subset K$ containing k we have*

$$\mathcal{F}\mathcal{S}(F) = F.$$

Proof of Lemma. We know that

$$F' = \mathcal{F}\mathcal{S}(F) \supset F,$$

and that

$$\mathcal{S}(F') = \mathcal{S}\mathcal{F}\mathcal{S}(F) = \mathcal{S}(F).$$

Thus from the last Lemma,

$$\deg_k F' = \frac{\|F\|}{\|\mathcal{S}(F')\|} = \frac{\|F\|}{\|\mathcal{S}(F)\|} = \deg_k F.$$

Hence

$$F' = F,$$

by Proposition 20. □

Lemma 23. *Suppose V is a vector space over an infinite field k; and suppose $U_1, \ldots, U_r$ are subspaces of V. Then*

$$V = \bigcup_{1 \le i \le r} U_i \Longrightarrow V = U_i$$

for some i.

Proof of Lemma. Suppose to the contrary that the U_i are all proper subspaces of V. We may suppose r minimal, so that

$$U_1 \cup \cdots \cup U_{r-1} \neq V.$$

Let

$$v \in V,\ v \notin U_1 \cup \cdots \cup U_{r-1};$$

and let

$$w \in V,\ w \notin U_r.$$

Consider the "line"

$$u = v + tw \quad (t \in k).$$

This cuts each U_i in at most one point; for if there were 2 such points then the whole line would lie in U_i. Thus if we choose t to avoid at most r values we can ensure that $u = v + tw$ does not lie in any of the subspaces, contrary to supposition. □

Lemma 24. *Suppose $k \subset F \subset K$. Then F is a simple extension of k:*

$$F = k(\theta).$$

Proof of Lemma. If k is finite, then so is F, and the result follows from the fact that a finite field F is a simple extension of every subfield $k \subset F$, eg $F = k(\pi)$, where π is a primitive root of F.

We may suppose therefore that k is infinite. By Lemma 22, each subfield $F \subset K$ containing k corresponds to the subgroup of $\mathcal{S}(F) \subset G$. Thus there can only be a finite number of such subfields.

It follows by the last Lemma that we can find $\theta \in F$ not belonging to any proper subfield of F containing k. But then $k(\theta)$ must be the whole of F:

$$k(\theta) = F.$$

□

Lemma 25. *G is the full group of automorphisms of K over k:*

$$G = \operatorname*{Aut}_k K.$$

Proof of Lemma. By the last Lemma,

$$K = k(\theta).$$

By Proposition 22 θ has minimal equation

$$m(x) = (x - g_1\theta) \cdots (x - g_n\theta),$$

where $g_1, \dots, g_n \in G$.

Every automorphism Θ of K over k must send θ into one of these conjugates $g\theta$. But this determines the automorphism completely. Hence $\Theta = g$. ☐

It only remains to show that $\deg_k K$ is finite. Suppose not. Then we can certainly find $\theta_1, \dots, \theta_n$ such that

$$\deg_k k(\theta_1, \dots, \theta_n) > \|G\|.$$

Now adjoin all the conjugates $m\theta_i$ of these elements; and let

$$F = k(g_1\theta_1, \dots, g_m\theta_n)$$

be the resulting subfield of K. Every automorphism $g \in G$ sends F into itself, since it merely permutes the elements $g_i\theta_j$. We can therefore apply the Theorem in this case, since $\deg_k F < \infty$. But then we conclude that

$$\deg_k F \le \|G\|,$$

contrary to construction. ☐

Corollary 16. *Suppose K is a galois extension of k; and suppose F is a subfield of K containing k:*

$$k \subset F \subset K.$$

Then K is a galois extension of F.

Corollary 17. *Suppose K is a finite extension of k. Then*

$$\operatorname*{Aut}_k K \le \deg_k K,$$

with equality if and only if the extension is galois.

Proof. First we must show that $G = \operatorname{Aut}_k K$ is finite. Suppose

$$K = k(\theta_1, \dots, \theta_n).$$

Let $m_i(x)$ be the minimal polynomial of θ_i. Then each automorphism $g \in G$ must send θ_i into another root $g\theta_i$ of $m_i(x)$. Thus there are only a finite number of choices for each $g\theta_i$; and since g is completely determined by the $g\theta_i$, there are only a finite number of choices for g.

Now we can apply the Theorem. Let

$$F = \mathcal{F}(G) = \{\theta \in K : g\theta = \theta \text{ for all } g \in G\}.$$

Then

$$\|G\| = \deg_F K \le \deg_k K,$$

with equality if and only if $F = k$, in which case the extension is galois, by definition. ☐

A.6 Normal Subgroups and Galois Extensions

Proposition 24. *Suppose F is a subfield of K containing k:*

$$k \subset F \subset K.$$

Then K is sent into itself by every $g \in G = \operatorname{Aut}_k K$ if and only if $\mathcal{S}(G)$ is a normal subgroup of G; and if this is so then

$$\operatorname*{Aut}_k F = \frac{G}{\mathcal{S}(F)}.$$

Proof. We know that

$$F = k(\theta)$$

for some $\theta \in K$, by Theorem 9(5). Let the conjugates of θ be

$$\theta_1 = \theta, \theta_2 = g_2\theta, \ldots, \theta_d = g_d\theta.$$

The automorphism $g \in G$ carries $k(\theta)$ into itself if and only if

$$g\theta \in k(\theta).$$

But θ and $g\theta$ have the same minimal polynomial, and so

$$\deg_k k(g\theta) = \deg_k k(\theta).$$

Thus

$$g\theta \in k(\theta) \Longleftrightarrow k(g\theta) = k(\theta).$$

Now

$$\begin{aligned}
\mathcal{S}\left(k(\theta)\right) &= \{h \in G : hg\theta = g\theta\} \\
&= \{h \in G : g^{-1}hg\theta = \theta\} \\
&= g\mathcal{S}\left(k(\theta)\right) g^{-1}.
\end{aligned}$$

Thus

$$k(g\theta) = k(\theta) \Longleftrightarrow \mathcal{S}\left(k(g\theta)\right) = \mathcal{S}\left(k(\theta)\right) \Longleftrightarrow g^{-1}Sg = S,$$

where $S = \mathcal{S}(k(\theta))$. In particular every $g \in G$ sends $k(\theta)$ into itself if and only if $g^{-1}Sg = S$ for all g, ie $S \lhd G$.

In this case, two automorphisms $g, h \in G$ induce the same automorphism of F if and only if they map θ into the same element. But

$$\begin{aligned}
g\theta = h\theta &\Longleftrightarrow h^{-1}g\theta = \theta \\
&\Longleftrightarrow h^{-1}g \in S \\
&\Longleftrightarrow hS = gS.
\end{aligned}$$

Thus the induced automorphisms of F are in one-one correspondence with the cosets of S, ie with the elements of the quotient-group G/S. It follows that

$$\operatorname*{Aut}_k F = G/S.$$

We note that these must be *all* the automorphisms of F over k, by Theorem 9(6). □

A.7 Splitting Fields

Definition 13. *The extension F of k is said to be a* splitting field *for the polynomial $p(x) \in k[x]$ if*

1. *$p(x)$ splits completely in F:*
$$p(x) = (x - \theta_1) \cdots (x - \theta_d) \quad (\theta_i \in F).$$
2. *F is generated by the roots of $p(x)$:*
$$F = k(\theta_1, \ldots, \theta_d).$$

Proposition 25. *Suppose K is a splitting field for the separable polynomial $p(x)$. Then K is a galois extension of k.*

Proof. Certainly
$$K = k(\theta_1, \ldots, \theta_d)$$
is of finite degree over k, by Proposition 20. Thus we may argue by induction on $\deg_k K$.

First let us dispose of the case in which k is finite. In this case K is a galois field
$$K = \mathcal{F}_(p^n);$$
and we know that $\mathcal{F}_(p^n)$ is a galois extension of all its subfields $\mathcal{F}_(p^m)$ (where $m \mid n$).

We may therefore assume that k is infinite. Let F be a minimal subfield of K containing k. Evidently K is the splitting field for $p(x)$ over F. Thus by our inductive hypothesis K is a galois extension of F.

There are 2 cases. Suppose first that there are two (or more) minimal subfields, F_1 and F_2. Then
$$\mathcal{F}(G) \subset F_1 \cap F_2 = k.$$
Hence K/k is galois.

Now suppose F is the unique minimal subfield. Since K/F is galois, K has only a finite number of subfields. By Lemma 23 we can choose $\phi \in K$ not in any of these subfields; and then
$$K = k(\phi).$$
Let $m(x)$ be the minimal polynomial of ϕ.

We can express ϕ as a polynomial in $\theta_1, \ldots, \theta_d$, say
$$\phi = f(\theta_1, \ldots, \theta_d).$$
For each permutation $\pi \in S_d$, let
$$\phi_\pi = f\left(\theta_{\pi(1)}, \ldots, \theta_{\pi(d)}\right) \quad (\pi \in S_d).$$
The coefficients of the product
$$P(x) = \prod_{\pi \in S_d} (x - \phi_\pi)$$
are all symmetric functions of $\theta_1, \ldots, \theta_d$, and so lie in k:
$$P(x) \in k[x].$$
It follows that all the roots of the minimal polynomial of θ, say
$$m(x) = (x - \theta_1) \ldots (x - \theta_d),$$
all lie in K.

Lemma 26. *Every element $\theta \in K$ is separable, ie θ is the root of a separable polynomial.*

Proof of Lemma. Let

$$g(x) = \gcd(m(x), m'(x)).$$

Then

$$g(x) \mid m(x).$$

Since $m(x)$ is irreducible, this implies that either $g(x)$ is constant, in which case $m(x)$ is separable, or else $m'(x)$ vanishes identically.

This is impossible in characteristic 0; so we need only consider the case of finite characteristic p.

In that case $m'(x) \equiv 0$ if and only if $m(x)$ contains only terms with powers x^{pr}; in other words,

$$m(x) = M(x^p) = x^{pr} + c_1 x^{p(r-1)} + \cdots + c_r.$$

It is easy to see that the pth powers form a subfield of K, say

$$K^p = \{\theta^p : \theta \in K\}.$$

Suppose $K^p \neq K$. If $K^p = k$ then

$$\theta_i^p \in k$$

for each of the roots θ_i of the generating polynomial $p(x)$. In other words, θ^i satisfies an equation

$$x^p - \theta_i^p \equiv (x - \theta_i)^p = 0$$

over k. But since $p(x)$ is separable, so is the minimal polynomial of θ_i. It follows that $\theta_i \in k$. Since this must hold for all the generators θ_i, $K = k$ and the result is trivial.

We may assume therefore that $F = K^p$ is a non-trivial subfield of K. Thus we can apply our inductive hypothesis, and deduce that the extension K/K^p is galois.

But if $\theta \in K$ then $\theta^p \in K^p$, and θ has minimal polynomial

$$x^p - \theta^p \equiv (x - \theta)^p.$$

It follows that every automorphism of K over K^p will leave θ fixed. Hence

$$\mathcal{G}(K/K^p) = \{e\},$$

and so the extension K/K^p is not galois, contrary to hypothesis. □

We have shown that $K = k(\theta)$, where the minimal polynomial $m(x)$ of θ splits completely in K into distinct factors:

$$m(x) = (x - \theta_1) \cdots (x - \theta_d).$$

For each root θ_i, the map

$$p(\theta) \mapsto p(\theta_i)$$

defines an automorphism of K over k. Thus

$$\deg_K k = d \leq \| \operatorname*{Aut}_k K \|.$$

It follows that K is a galois extension of k, by Corollary 2 to Theorem 9. □

Appendix A

The Normal Basis Theorem

As we have seen, we can regard a finite field F as a vector space over its prime subfield P. We often want to construct a *basis* for this vector space.

The simplest way to choose such a basis is to pick an element $\alpha \in F$ whose minimal polynomial has degree n—or equivalently, such that $F = P(\alpha)$. (For example, any primitive root of F will have this property.) For then the elements

$$\{1, \alpha, \alpha^2, \ldots, \alpha^{n-1}\}$$

are linearly independent, and so form a basis for F.

However, it is sometimes preferable to use a more specialized basis, namely one consisting of a complete family of conjugates

$$\left\{\gamma, \gamma^p, \ldots, \gamma^{p^{n-1}}\right\}$$

Such a basis is said to be *normal*; and the Normal Basis Theorem asserts the existence of normal bases in every finite field.

Theorem 10. *There exists an element $\alpha \in F = \mathcal{F}_(p^n)$ whose n conjugates*

$$\alpha, \pi\alpha, \pi^2\alpha, \ldots, \pi^{n-1}\alpha$$

form a basis for F over its prime subfield P.

Our proof of this theorem is based on a straightforward but perhaps unfamiliar result from linear algebra.

Suppose

$$T : V \to V$$

is a linear transformation of the finite-dimensional vector space over the scalar field k. Let $m(x)$ be the minimal polynomial of T.

(Recall that $m(x)$ is the polynomial of least degree satisfied by T, taken with leading coefficient 1. It has the property that

$$p(T) = 0 \Longleftrightarrow m(x) | p(x),$$

as is readily seen on dividing $p(x)$ by $m(x)$:

$$p(x) = m(x)q(x) + r(x) \qquad (\deg r(x) < \deg m(x))$$

(Incidentally, there certainly do exist polynomials $p(x)$ such that $p(T) = 0$. For the space $\hom(V, V)$ of all linear maps $T : V \to V$ has dimension n^2; and so the linear maps

$$I, T, T^2, \ldots, T^{n^2}$$

must be linearly independent, ie T satisfies an equation of degree $\leq n^2$. In fact, by the Cayley-Hamilton Theorem T satisfies its own characteristic equation

$$\chi_T(x) = \det(xI - T);$$

so the minimal polynomial of T actually has degree $\leq n$. But we don't need this.)

We can extend this notion of minimal polynomial as follows. Suppose $v \in V$. Consider the set of polynomials

$$I(v) = \{f(x) : f(T)v = 0\}.$$

This set is an *ideal* in the polynomial ring $k[x]$, ie it is closed under addition, and under multiplication by any polynomial in $k[x]$. It follows—since $k[x]$ is a *principle ideal doman*—that $I(v)$ consists of all the multiples of a polynomial $m_v(x)$. (It is easy to prove this result directly, taking $m_v(x)$ to be a polynomial of minimal degree in $I(v)$.) The main properties of this polynial are summarised in

Lemma 27. *1. $m_v(x) \| m(x)$ for all $v \in V$.*

2. $m(x) = \mathrm{lcm}_{v \in V}\, m_v(x)$.

3. If $u = f(T)v$ for some polynomial $f(x)$ then $m_u(x) \| m_v(x)$.

4. If $u, v \in V$ and $m_u(x), m_v(x)$ are co-prime then

$$m_{u+v}(x) = m_u(x) m_v(x).$$

Proof. 1. Since $m(T) = 0$, it follows that $m(T)v = 0$ for all v, and so

$$m_v(x) \| m(x).$$

2. It follows from the above that

$$f(x) = \operatorname*{lcm}_{v \in V} m_v(x)$$

is defined, with $f(x) \| m(x)$. But

$$f(T)v = 0$$

for *all* $v \in V$, and so

$$f(T) = 0.$$

Hence $f(x) = m(x)$.

3. We have

$$m_v(T)u = m_v(T)f(T)v = f(T)m_v(T)v = 0.$$

Hence $m_u(x) \| m_v(x)$.

4. Clearly

$$m_{u+v}(x) \| m_u(x) m_v(x).$$

Let

$$w = m_u(T)(u + v) = m_u(T)v;$$

and let $f(x) = m_w(x)$. Then

$$0 = f(T)w = f(T)m_u(T)v,$$

and so

$$m_v(x) \| f(x) m_u(x).$$

But since $m_u(x), m_v(x)$ are coprime, this implies that

$$m_v(x) \| f(x).$$

On the other hand, by part 3 of the Lemma,

$$f(x) \| m_{u+v}(x).$$

Hence

$$m_v(x) \| m_{u+v}(x),$$

and similarly

$$m_u(x) \| m_{u+v}(x).$$

Since $m_u(x), m_v(x)$ are coprime, this implies that

$$m_u(x) m_v(x) \| m_{u+v}(x),$$

from which the result follows.

□

Lemma 28. *There exists a vector v (sometimes called a* cylic *vector of T) such that $m_v(x) = m(x)$.*

Proof. Let

$$m(x) = p_1(x)^{e_1} p_2(x)^{e_2} \cdots p_r(x)^{e_r}$$

be the expression for the minimal polynomial $m(x)$ of T as a product of prime polynomials.

From part 2 of the Lemma above, for each $i (1 \le i \le r)$ we can find a vector u_i whose minimal polynomial is divisible by $p_i(x)^{e_i}$, say

$$m_{u_i}(x) = p_i(x)^{e_i} f_i(x).$$

But then

$$v_i = f_i(T) u_i$$

has minimal polynomial $p_i(x)^{e_i}$.

Now from part 4 of the Lemma above, if we set

$$v = v_1 + v_2 + \cdots + v_r$$

then

$$m_v(x) = m(x).$$

□

We shall apply this result to the fundamental automorphism π of $\mathcal{F}_(p^n)$.

Proof. Since $\pi : F \to F$ is a linear transformation, we can apply the Lemma above.

The minimal polynomial of π is

$$m(x) = x^n - 1.$$

For π satisfies $m(x) = 0$; and it cannot satisfy any equation of lower degree. For suppose

$$c_0 \pi^d + c_1 \pi^{d-1} + \ldots c_d = 0.$$

Then every element $\alpha \in F$ satisfies the equation

$$c_0 x^{p^d} + c_1 x^{p^{d-1}} + \dots c_d = 0.$$

But that is a contradiction, since the polynomial on the left has at most p^d roots.

By the Lemma, we can find a cyclic vector of π, ie an element $\alpha \in F$ whose minimal polynomial is $x^n - 1$. But this implies in particular that

$$\alpha, \pi\alpha, \pi^2\alpha, \dots, \pi^{n-1}\alpha$$

are linearly independent, and so form a basis for F over P. $\square$