

CDM

Finite Fields

KLAUS SUTNER

CARNEGIE MELLON UNIVERSITY

SPRING 2023



1 Ideals

2 The Structure Theorem

We know that every finite field carries two apparently separate structures: additive and multiplicative.

addition $\mathbb{F} \cong (\mathbb{Z}_p)^k \quad (a_1, \dots, a_k)$

multiplication $\mathbb{F}^\times \cong \mathbb{Z}_{p^k-1} \quad g^i$

The problem is that we have absolutely no idea how to unify the two.

Time to get serious about building a finite field.

We would like to follow the construction of $\mathbb{Q}(\sqrt{2})$ from above, adjoining a root of $x^2 - 2 = 0$ to the rationals. But this time, we won't rely on intuition and prior knowledge of the reals. For example, consider the polynomial

$$f = x^2 + x + 1 \in \mathbb{F}_2[x]$$

We can easily check that f has no root over \mathbb{F}_2 .

So how do we expand \mathbb{F}_2 to a field \mathbb{F} where f has a root?

This time:

- We do not know a convenient big field like \mathbb{R} that we can use as a safe sandbox, and
- we have no intuitive idea what a root of f looks like.

So, we can't just do

$$\mathbb{Q}(\sqrt{2}) = \{a + b\sqrt{2} \mid a, b \in \mathbb{Q}\} \subseteq \mathbb{R}$$

But: we can try to lift this construction to our new setting. To wit, we applied the simplification rule

$$x^2 \rightsquigarrow 2$$

to all polynomials over \mathbb{Q} . This produces expressions $a + bx$, $a, b \in \mathbb{Q}$, that turn out to form a field (the “unknown” x works just like the root we are after).

We want $x^2 + x + 1 = 0$, so we use the simplification rule

$$x^2 \rightsquigarrow x + 1$$

and apply it to all polynomials in $\mathbb{F}_2[x]$. We are in characteristic 2, so plus is minus.

With luck, we might wind up with a finite field that has a root for f .

Here is one of the occasions where it is useful to think of a polynomial as an expression, a term in some formal language.

On that view, we can apply the rewrite rule $x^2 \rightsquigarrow x + 1$ to try to simplify the expression. More precisely, we use this rule plus all the standard simplifications we can apply to our terms (associativity, commutativity, cancellation, ...).

For those concerned about the StringWorld approach to life, not to worry, we will unearth the actual algebraic meaning behind this rewrite process in a moment.

So what happens to an arbitrary polynomial $p(x) \in \mathbb{F}_2[x]$ if we apply this rule systematically? Essentially, we can smash all the higher powers of x . Here is an example.

$$\begin{aligned}x^6 + x^3 + x + 1 &\rightsquigarrow (x + 1)^3 + x(x + 1) + x + 1 \\ &\rightsquigarrow (x^3 + x^2 + x + 1) + (x^2 + x) + x + 1 \\ &\rightsquigarrow x(x + 1) + (x + 1) + 1 \\ &\rightsquigarrow x + 1\end{aligned}$$

Proposition

x^k reduces to $1, x, x + 1$, depending on $k \bmod 3$.

So $x^6 + x^3 + x + 1 \rightsquigarrow 1 + 1 + x + 1 = x + 1$.

The simplification process is highly nondeterministic, there are many choices along the way.

This might cause a huge headache: if we apply the rules in one particular way, we get a different result from when we apply the rules in another way.

One really needs to make sure the process is **confluent**: application order does not matter, the final result is always the same. More later.

In general, if we start with a polynomial $f \in \mathbb{F}[x]$ of degree d , we get a simplification rule

$$x^k \rightsquigarrow a_{k-1}x^{k-1} + a_{k-2}x^{k-2} + \dots + a_1x + a_0$$

But then we can reduce all polynomials down to polynomials of degree at most $d - 1$. If the coefficient field has size q , the collection of polynomials of degree less than d , $\mathbb{F}_{<d}[x]$, has size q^d .

In particular if $\mathbb{F} = \mathbb{Z}_p$ for some prime p we get p^d reduced polynomials.

We want to use $\mathbb{F}_{<d}[x]$ as the carrier set for our extension field $\mathbb{F} \subseteq \mathbb{K}$. What are the operations?

- Addition is simply addition of polynomials in $\mathbb{F}[x]$.
- Multiplication is multiplication of polynomials in $\mathbb{F}[x]$ followed by a reduction: we have to apply the simplification rule until we get back to a polynomial of degree less than d .

We have an algorithm, but we need to work out the algebraic meaning of all of this.

Our simplification process induces an equivalence relation on $\mathbb{F}[x]$: two polynomials are equivalent if they reduce to the same polynomial in $\mathbb{F}_{<d}[x]$.

In fact, we get a **congruence** \approx : our simplification is compatible with the field operations.

So we can form a quotient ring, which turns out to be exactly the field we are looking for:

$$\mathbb{K} = \mathbb{F}[x] / \approx$$

Definition

Let R be a commutative ring. An **ideal** $I \subseteq R$ is a subset that is closed under addition and under multiplication by arbitrary ring elements: $a \in I, b \in R$ implies $ab \in I$.

So an ideal is much more constrained than a subring: it has to be closed by multiplication from the outside. Ideals are hugely important since they produce congruences and thus allow us to form a quotient structure:

$$a = b \pmod{I} \quad \text{iff} \quad a - b \in I.$$

As a consequence, arithmetic in this quotient structure is well-behaved: E.g.

$$a = a', b = b' \pmod{I} \quad \Rightarrow \quad a + b = a' + b', ab = a'b' \pmod{I}$$

Suppose \mathbb{F} is a field and consider an irreducible polynomial $f(x)$ and the principal ideal $(f(x)) = f(x)\mathbb{F}[x]$ that it generates.

We identify two polynomials when their difference is divisible by f :

$$h(x) = g(x) \pmod{f(x)} \iff f(x) \mid h(x) - g(x)$$

Let d be the degree of f . Then any polynomial h is equivalent to a polynomial g of degree less than d : write $h(x) = q(x)f(x) + g(x)$ by polynomial division.

What is the smallest ideal containing elements $a_1, \dots, a_k \in R$?

All we need is linear combinations: the ideal **generated** by a_1, \dots, a_k is

$$(a_1, \dots, a_k) = \{ r_1 a_1 + \dots + r_k a_k \mid r_i \in R \}$$

In particular for $k = 1$ we have

$$(a) = \{ ra \mid r \in R \}$$

This is the **principal ideal** generated by a .

The ideals $\{0\}$ and R are called trivial, all others are proper.

Note that a field is a commutative ring that has no proper ideals.

Definition

A **principal ideal domain (PID)** is an integral domain, all of whose ideals are principal.

Important examples of PIDs are

- the integers \mathbb{Z} (think GCD)
- the Gaussian integers $\mathbb{Z}[i]$
- a polynomial ring $\mathbb{F}[x]$ where \mathbb{F} is a field

Counterexamples: $\mathbb{Z}[x]$ and $\mathbb{F}[x, y]$ both fail to be PIDs.

Suppose we have an extension $\mathbb{F} \subseteq \mathbb{K}$ with $\alpha \in \mathbb{K}$ algebraic over \mathbb{F} . Let

$$I = \{ f \in \mathbb{F}[x] \mid f(\alpha) = 0 \}$$

Then I is an ideal and we must have $I = (g)$.

The polynomial g has minimal degree among all the annihilators of α , and we may safely assume that g is monic.

Definition

This polynomial g is the **minimal polynomial** of α over \mathbb{F} .

In algebra it is important to come up with the right notion of substructure: just picking a subset that is closed under the algebraic operations is often not very interesting.

- For groups, normal subgroups are arguably more important than plain subgroups.
- For rings, ideals are arguably more important than subrings.
- But for vector spaces, sub-vector-spaces are just the right notion.

Ideals provide the right type of equivalence relation for the construction of a finite field from a polynomial ring. Alas, the ideals cannot be chosen arbitrarily, we need to start from special polynomials, in analogy to the modulus being prime in the integer case.

Definition

A polynomial is **irreducible** if it is not the product of polynomials of smaller degree.

Irreducibility is necessary when we try to construct a field $\mathbb{F}[x]/(f)$: otherwise we do not even get an integral domain.

For suppose $f(x) = f_1(x)f_2(x)$ where both f_1 and f_2 have degree at least 1. Then $1 \leq \deg(f_i) < \deg(f)$, so neither f_1 or f_2 can be simplified in $\mathbb{F}[x]/(f)$. In particular both elements in $\mathbb{F}[x]/(f)$ are non-zero, but their product is zero.

Fix some prime p .

Question:

How many irreducible polynomials of degree m are there in $\mathbb{F}_p[x]$?

Let's write I_m^p for this number, so trivially $I_m^p \leq p^m$.

Lemma (Gauss)

$$I_m^p = \frac{1}{m} \sum_{d|m} \mu(m/d) p^d$$

Recall the Möbius function μ :

$$\mu(n) = \begin{cases} +1 & \text{if } n \text{ square-free, even number of prime factors} \\ -1 & \text{if } n \text{ square-free, odd number of prime factors} \\ 0 & \text{otherwise.} \end{cases}$$

One can show that

$$(p^m - 2p^{m/2})/m \leq I_m^p \leq p^m/m$$

E.g., $I_{50}^2 = 22517997465744$, about 2 percent.

Here are some numerical values for characteristic 2.

1-5	2	1	2	3	6
6-10	9	18	30	56	99
11-15	186	335	630	1161	2182
16-20	4080	7710	14532	27594	52377

$$x, 1 + x$$

$$1 + x + x^2$$

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

$$1 + x^3 + x^4, 1 + x + x^4, 1 + x + x^2 + x^3 + x^4$$

$$1 + x^3 + x^5, 1 + x^2 + x^5, 1 + x^2 + x^3 + x^4 + x^5, 1 + x + x^3 + x^4 + x^5$$

$$1 + x + x^2 + x^4 + x^5, 1 + x + x^2 + x^3 + x^5$$

All irreducibles in $\mathbb{F}_2[x]$ up to degree 5.

Lemma

$$x^{p^k} - x = \prod (f \mid f \text{ monic, irreducible, } \deg(f) \mid k)$$

There is a fairly good test for irreducibility that assumes we have access to the prime factors of m (a reasonable assumption).

Theorem (Rabin)

Suppose $f \in \mathbb{F}_p[x]$ is a monic polynomial of degree m . Then f is irreducible iff f divides $x^{p^d} - x$ but f and $x^{p^{d/q}} - x$ are coprime for all prime divisors q of m .

Over \mathbb{F}_2 , the polynomial

$$f(x) = x^3 + x + 1$$

is irreducible. Let $I = (f(x))$ be the ideal generated by f .

The first few powers of x modulo I are:

$$1, x, x^2, x + 1, x^2 + x, x^2 + x + 1, x^2 + 1$$

These are actually all polynomials of degree less-than 3, except 0.

So $\mathbb{F}_{<3}[x]$ forms an integral domain, and hence a field, if multiplication is understood modulo I .

OK, but where is the root of f ?

We write α for (the equivalence class of) x for emphasis, $\alpha = x \bmod f(x)$.

Then $\alpha \in \mathbb{K}$ is a root of f in the extension field \mathbb{K} .

Why? We have by brute force

$$f(\alpha) = \alpha^3 + \alpha + 1 = 0 \pmod{I}$$

Yes, this is a bit lame. One would have hoped for some kind of fireworks, some clever way of writing down the root in terms of some fancy polynomial.

But, it's really no different from the $\sqrt{2}$ example, just less familiar.

Again, algebraically, it is best to think of the extension field $\mathbb{F}_2 \subseteq \mathbb{K}$ as a quotient structure, as the polynomials modulo f :

$$\mathbb{K} = \mathbb{F}_2[x]/(f(x))$$

With a view towards algorithms, we can make things more combinatorial by keeping track of coefficient vectors, in this case

$$c_2x^2 + c_1x + c_0 \rightsquigarrow (c_2, c_1, c_0)$$

where $c_i \in \mathbb{F}_2$ is just a single bit.

In this setting the additive structure is trivial: it's just componentwise addition of these triples mod 2.

$$(c_2, c_1, c_0) + (c'_2, c'_1, c'_0) = (c_2 + c'_2, c_1 + c'_1, c_0 + c'_0)$$

As observed before, the additive group of these fields is just a Boolean group. Note that this operation is trivial to implement (xor on bit-vectors, can even be done in 32 or 64 bit blocks).

For other characteristics, though, we have to use modular numbers.

How about multiplication? Since multiplication increases the degree, we can't just multiply out, but we have to simplify using our rule $x^3 \rightarrow x + 1$ afterwards.

The product

$$(c_2, c_1, c_0) \cdot (c'_2, c'_1, c'_0) = (d_2, d_1, d_0)$$

is given by the coefficient triple

$$\begin{aligned}d_2 &= c_2 c'_0 + c_1 c'_1 + c_0 c'_2 + c_2 c'_2 \\d_1 &= c_1 c'_0 + c_0 c'_1 + c_2 c'_1 + c_1 c'_2 + c_2 c'_2 \\d_0 &= c_0 c'_0 + c_2 c'_1 + c_1 c'_2\end{aligned}$$

This is a bit messy, and it gets more messy when we deal with larger degree polynomials. Still, we could hard-wire a circuit.

Recall that α is the equivalence class of x . We have already checked that α is the generator of \mathbb{F}^\times . Here are the corresponding vector representations.

$$\alpha^0 = 1 \qquad = (0, 0, 1)$$

$$\alpha^1 = \alpha \qquad = (0, 1, 0)$$

$$\alpha^2 = \alpha^2 \qquad = (1, 0, 0)$$

$$\alpha^3 = \alpha + 1 \qquad = (0, 1, 1)$$

$$\alpha^4 = \alpha^2 + \alpha \qquad = (1, 1, 0)$$

$$\alpha^5 = \alpha^2 + \alpha + 1 \qquad = (1, 1, 1)$$

$$\alpha^6 = \alpha^2 + 1 \qquad = (1, 0, 1)$$

Careful, though, it is in general **not** the case that α generates the whole multiplicative group.

For this to work, we need to choose particular irreducible polynomials in our construction, so-called **primitive polynomials**.

For example, there are 9 monic irreducibles of degree 6 in $\mathbb{F}_2[x]$:

$$1+x^5+x^6, 1+x^3+x^6, 1+x^2+x^4+x^5+x^6, 1+x^2+x^3+x^5+x^6, 1+x+x^6, \\ 1+x+x^4+x^5+x^6, 1+x+x^3+x^4+x^6, 1+x+x^2+x^5+x^6, 1+x+x^2+x^4+x^6$$

But 3 of them fail to be primitive:

$$1+x^3+x^6, 1+x+x^2+x^4+x^6, 1+x^2+x^4+x^5+x^6$$

We really obtain a field this way, not just some ring.

	h	h^{-1}
1	1	1
2	α	$1 + \alpha^2$
3	α^2	$1 + \alpha + \alpha^2$
4	$1 + \alpha$	$\alpha + \alpha^2$
5	$1 + \alpha^2$	α
6	$\alpha + \alpha^2$	$1 + \alpha$
7	$1 + \alpha + \alpha^2$	α^2

This table duly defines an involution: $(h^{-1})^{-1} = h$.

1 Ideals

2 **The Structure Theorem**

Recall the big theorem we announced some time ago:

Theorem

Every finite field \mathbb{F} has cardinality p^k where p is prime and the characteristic of \mathbb{F} , and $k \geq 1$.

Moreover, for every p prime and $k \geq 1$ there is a finite field of cardinality p^k .

Lastly, all fields of cardinality p^k are isomorphic.

We have already taken care of parts 1 and 2:

- Since \mathbb{F} is finite vector space over \mathbb{Z}_p where p is the characteristic of \mathbb{F} it must have size p^k , p prime, $k \geq 1$.
- Since there are irreducible polynomials over \mathbb{Z}_p of degree k for any k we can always construct a finite field of the form $\mathbb{Z}_p[x]/(f)$ of size p^k .

The Problem:

It is absolutely unclear that all these quotient rings are isomorphic.

Issue 1 Suppose we use some irreducible polynomial f . Say f has roots α and β . Why should we have

$$\mathbb{F}_p(\alpha) \cong \mathbb{F}_p(\beta)$$

Issue 2 Suppose f and g are two irreducible polynomials of the same degree. Why should we have

$$\mathbb{F}_p[x]/(f) \cong \mathbb{F}_p[x]/(g)$$

We will first deal with issue 1, and then extend the method to handle issue 2.

First a few handy tools.

Let's collect some tools to compare rings and fields.

Definition

Let R and S be two rings and $f : R \rightarrow S$. f is a **ring homomorphism** if

$$f(g + h) = f(g) + f(h) \quad \text{and} \quad f(gh) = f(g)f(h).$$

If f is in addition injective/surjective/bijective we speak about monomorphisms, epimorphism and isomorphisms, respectively. The **kernel** of a ring homomorphism is the set of elements that map to 0.

Notation: **$\ker(f)$** .

Note that $f(0) = 0$. Since $f(x) = f(y)$ iff $x - y \in \ker(f)$ a ring homomorphism is a monomorphism iff its kernel is trivial: $\ker(f) = \{0\}$.

It is easy to see that the kernel of any ring homomorphism $f : R \rightarrow S$ is an ideal in R .

Recall that our rings always have a multiplicative unit (as opposed to abominable rngs). So one requires

$$f(1) = 1.$$

These are sometimes called **unital ring homomorphisms**. In particular field homomorphisms are unital.

Lemma

If $f : \mathbb{F} \rightarrow \mathbb{K}$ is a field homomorphism, then f is injective.

Proof.

$\ker(f) \subseteq \mathbb{F}$ is an ideal. But in a field there are only two ideals: $\{0\}$ and the whole field. Since $f(1) = 1$, 1 is not in the kernel, so the kernel must be $\{0\}$ and f is injective.

□

Here is a somewhat surprising example of a homomorphism.

Definition

Let R be a ring of characteristic $p > 0$.

The **Frobenius homomorphism** is defined by the map $R \rightarrow R, x \mapsto x^p$.

The Frobenius map is indeed a ring homomorphism since R has characteristic p :

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

Over a finite field we even get an automorphism. The orbits of a non-zero element look like

$$a, a^p, a^{p^2}, \dots, a^{p^{k-1}}$$

Exercise

Use the binomial theorem to prove that the Frobenius map is a homomorphism.

Back to our uniqueness problem. As a first step, consider two roots of the same polynomial. More precisely, let $f(x) \in \mathbb{F}[x]$ irreducible, $\alpha_1 \neq \alpha_2$ two roots of f . Consider two corresponding simple field extensions $\mathbb{F} \subseteq \mathbb{K}_i = \mathbb{F}(\alpha_i)$ where $\alpha_i \in \mathbb{K}_i$.

Theorem (Extension Isomorphism Theorem)

There is a unique isomorphism $\varphi : \mathbb{K}_1 \rightarrow \mathbb{K}_2$ such that $\varphi(\alpha_1) = \alpha_2$ and φ is the identity on \mathbb{F} .

Proof.

We exploit the fact that $\mathbb{K}_i = \mathbb{F}[\alpha_i]$ and define

$$\varphi(p(\alpha_1)) = p(\alpha_2)$$

for any $p \in \mathbb{F}[x]$.

It is straightforward but tedious to check that φ has all the right properties.

The interesting part is to verify well-definedness.

To this end, suppose $p(\alpha_1) = p'(\alpha_1)$.

Then f divides $p - p'$, say $p - p' = f \cdot q$.

But then $(p - p')(\alpha_2) = f(\alpha_2) \cdot q(\alpha_2) = 0$.

□

Exercise

Work out the rest of the proof.

The following fact is often useful to establish an isomorphism. Suppose $f : R \rightarrow S$ is an epimorphism (no major constraint, otherwise replace S by the range of f). Then $R/\ker(f)$ is isomorphic to S .

For example, we can use this technique to prove our old theorem about field extensions by adjoining roots.

More precisely, let $\mathbb{F}(\alpha)$ be the smallest field $\mathbb{F} \subseteq \mathbb{F}(\alpha) \subseteq \mathbb{K}$ that contains a root $\alpha \in \mathbb{K}$ of some polynomial $f \in \mathbb{F}[x]$. Then

$$\mathbb{F}(\alpha) = \{g(\alpha) \mid g \in \mathbb{F}[x]\} = \mathbb{F}[\alpha]$$

rather than, say, the collection of rational functions over \mathbb{F} evaluated at α .

To see why, note that the right hand side is the range of the evaluation map

$$\begin{aligned}\nu: \mathbb{F}[x] &\longrightarrow \mathbb{K} \\ g &\mapsto g(\alpha)\end{aligned}$$

that evaluates g at α , producing a value in \mathbb{K} . It is easy to check that ν is a ring homomorphism and clearly $(f) \subseteq \ker(\nu)$.

We may safely assume that f is monic and has minimal degree in $\mathbb{F}[x]$ of all polynomials with root α . Then f is irreducible and we have

$$\ker(\nu) = \{p \in \mathbb{F}[x] \mid f \text{ divides } p\} = (f)$$

This shows that the range of ν is isomorphic to $\mathbb{F}[x]/(f)$ and hence a field.

Irreducibility is essential here, otherwise $f(x) = (x^2 - 2)(x^2 - 3) = x^4 - 5x^2 + 6$ with $\alpha = \sqrt{2}$ over $\mathbb{F} = \mathbb{Q} \subseteq \mathbb{C} = \mathbb{K}$ would produce a non-integral domain.

Note that this is the third time we encounter kernels.

- For a general function $f : A \rightarrow B$ the kernel relation is given by $f(x) = f(y)$.
- For a group homomorphism $f : A \rightarrow B$ the kernel is given by $\{x \in A \mid f(x) = 1\}$.
- For a ring homomorphism $f : A \rightarrow B$ the kernel is given by $\{x \in A \mid f(x) = 0\}$.

In the last two cases we can easily recover the classical kernel relation and the definition as stated turns out to be more useful.

Still, there is really just one idea.

Back to the problem of showing that there is only “one” finite field \mathbb{F}_{p^k} of size p^k . To understand finite fields completely we need just one more idea.

Definition

Let $f \in \mathbb{F}[x]$ monic, $\mathbb{F} \subseteq \mathbb{K}$. Field \mathbb{K} is a **splitting field** of f if

- $f(x) = (x - \alpha_1) \dots (x - \alpha_d)$ in $\mathbb{K}[x]$, and
- $\mathbb{K} = \mathbb{F}(\alpha_1, \dots, \alpha_d)$.

Needless to say, the $\alpha_i \in \mathbb{K}$ are exactly the roots of f . Thus, in a splitting field we can decompose the polynomial into linear factors.

In other words, \mathbb{K} is the smallest field where f splits into linear factors; by adjoining all the roots of f we get all of \mathbb{K} .

Example

\mathbb{C} is the splitting field of $x^2 + 1 \in \mathbb{R}[x]$.

It is hugely surprising that over \mathbb{C} any non-constant real polynomial can already be decomposed into linear factors, everybody splits already.

Example

Consider $f(x) = x^8 + x \in \mathbb{F}_2[x]$. Then

$$f(x) = x(x+1)(x^3+x^2+1)(x^3+x+1)$$

Adjoining one root of $g(x) = x^3 + x + 1$ already produces the splitting field of f : the other irreducible factor of degree 3 also splits.

$$x^8 + x = x(x+1)(x^3 + x^2 + 1)(x^3 + x + 1)$$

element	root of
0	x
α^0	$x + 1$
α^1	$x^3 + x + 1$
α^2	$x^3 + x + 1$
α^3	$x^3 + x^2 + 1$
α^4	$x^3 + x + 1$
α^5	$x^3 + x^2 + 1$
α^6	$x^3 + x^2 + 1$

Our next goal is to establish the following result.

Theorem (Splitting Field Theorem)

*For any irreducible polynomial there exists a splitting field.
Moreover, any two such splitting fields are isomorphic.*

Note that we have all the tools to construct a splitting field: we just keep adjoining roots of irreducible factors of the given polynomial.

But for the uniqueness part we need a bit more machinery.

Basic problem: what would happen in the last example if we had chosen $x^3 + x + 1$ rather than $x^3 + x^2 + 1$? We get isomorphic vector spaces, but why should the multiplicative structure be the same?

Suppose we have an isomorphism $\theta : \mathbb{F}_1 \rightarrow \mathbb{F}_2$ and $f_1 \in \mathbb{F}_1[x]$. Set $f_2 = \theta(f_1)$ and let $\mathbb{F}_i \subseteq \mathbb{K}_i$ be splitting fields for f_i .

Lemma

There exists a isomorphism $\varphi : \mathbb{K}_1 \rightarrow \mathbb{K}_2$ such that $\varphi \upharpoonright \mathbb{F}_1 = \theta$.

Proof.

We construct φ by repeated application of the Extension Isomorphism theorem.

Technically, we perform induction on the degree $d = [\mathbb{K}_1 : \mathbb{F}_1]$ of the splitting extension.

The case $d = 1$ is trivial, since then $\mathbb{F}_i = \mathbb{K}_i$.

So suppose $d > 1$. We may safely assume that f is not irreducible; otherwise the Extension Isomorphism theorem suffices.

Since $\mathbb{F}_1 \neq \mathbb{K}_1$ there is some irreducible factor $g_1 \in \mathbb{F}[x]$ of f_1 of degree larger than 1.

g_1 splits in \mathbb{K}_1 , say, it has a root α_1 .

Set $g_2 = \theta(g_1)$ with root α_2 in \mathbb{K}_2 .

By the Extension Isomorphism theorem from above, there is a unique isomorphism $\theta' : \mathbb{F}_1[\alpha_1] \rightarrow \mathbb{F}_2[\alpha_2]$ such that $\theta'(\alpha_1) = \alpha_2$ and $\theta' \upharpoonright \mathbb{F}_1 = \theta$.

But $[\mathbb{K}_i[\alpha_i] : \mathbb{F}_i] = \deg g_i < \deg f_i$, so by the induction hypothesis we are done.

□

Corollary

Splitting fields are unique up to isomorphism.

Now we can pin down the structure of all finite fields: they are splitting fields (and hence uniquely determined).

Theorem

There is a unique (up to isomorphism) finite field of size p^k .

Proof.

Let $q = p^k$ and consider $f = x^q - x \in \mathbb{F}_p[x]$.

f has q roots, which form a field. For let a and b two roots, then:

$$f(a + b) = (a + b)^q - (a + b) = a^q - a + b^q - b = 0$$

$$f(ab) = (ab)^q - (ab) = a^q b^q - ab = 0$$

Hence the roots of f form the whole splitting field of f . By the Splitting Field theorem, this field is unique up to isomorphism. \square

Note that the prime subfield of \mathbb{F}_q , $q = p^k$, consists of all fixed points of the Frobenius morphism: $x = x^p$, or $x^p - x = 0$.

Similarly, the whole field can be construed as the fixed points of the map $x \mapsto x^q$, an iterated version of the Frobenius morphism:

$$x^q - x = 0$$

Also note that \mathbb{F}_q is far removed from being algebraically closed: $x^k - 1$ has the obvious root 1, but no others. So

$$(x^k - 1)/(x - 1) = x^{k-1} + \dots + x + 1$$

has no roots whatsoever.

We can exploit knowledge of the splitting field to study roots of irreducible polynomials.

Consider the irreducible $f = x^3 + x^2 + 1 \in \mathbb{F}_2[x]$.

Let α be a root of f in the splitting field \mathbb{K} . By long division (exploiting $\alpha^3 + \alpha^2 + 1 = 0$) we find all the roots of f :

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

Hence

$$\mathbb{K} = \mathbb{F}_2(\alpha) = \{ a\alpha^2 + b\alpha + c \mid a, b, c \in \mathbb{F}_2 \}$$

Consider characteristic $p = 5$ and $k = 2$.

$$\begin{aligned}x^{25} - x &= x(1+x)(2+x)(3+x)(4+x) \\ &\quad (2+x^2)(3+x^2)(1+x+x^2)(2+x+x^2)(3+2x+x^2)(4+2x+x^2) \\ &\quad (3+3x+x^2)(4+3x+x^2)(1+4x+x^2)(2+4x+x^2)\end{aligned}$$

The factorization of $x^{25} - x$ was done by an algorithm, one of the great success stories of computer algebra dating back to the 1960s.


At any rate, there are 10 irreducible quadratic polynomials to choose from.

Which one should we pick?

The factorization of $X^{11^2} - x$.

$$\begin{aligned}
 & x(1+x)(2+x)(3+x)(4+x)(5+x)(6+x)(7+x)(8+x)(9+x)(10+x) \\
 & (1+x^2)(3+x^2)(4+x^2)(5+x^2)(9+x^2)(1+x+x^2)(4+x+x^2)(6+x+x^2) \\
 & (7+x+x^2)(8+x+x^2)(2+2x+x^2)(4+2x+x^2)(5+2x+x^2)(6+2x+x^2) \\
 & (10+2x+x^2)(3+3x+x^2)(6+3x+x^2)(8+3x+x^2)(9+3x+x^2)(10+3x+x^2)(2+4x+x^2) \\
 & (5+4x+x^2)(7+4x+x^2)(8+4x+x^2)(9+4x+x^2)(1+5x+x^2)(2+5x+x^2)(3+5x+x^2) \\
 & (7+5x+x^2)(10+5x+x^2)(1+6x+x^2)(2+6x+x^2)(3+6x+x^2)(7+6x+x^2) \\
 & (10+6x+x^2)(2+7x+x^2)(5+7x+x^2)(7+7x+x^2)(8+7x+x^2)(9+7x+x^2)(3+8x+x^2) \\
 & (6+8x+x^2)(8+8x+x^2)(9+8x+x^2)(10+8x+x^2)(2+9x+x^2)(4+9x+x^2) \\
 & (5+9x+x^2)(6+9x+x^2)(10+9x+x^2)(1+10x+x^2)(4+10x+x^2)(6+10x+x^2) \\
 & \qquad (7+10x+x^2)(8+10x+x^2)
 \end{aligned}$$

This time, there are 55 quadratic irreducible polynomials to pick from.



Number of terms in the factorization of $x^n - x$ modulo 2, for $n \leq 200$.

$$\Phi(n) = |\mathbb{Z}_n^*| = |\{x < n \mid \gcd(x, n) = 1\}|$$

Euler proved the following product formula

$$\Phi(n) = n \prod_{p|n} (1 - 1/p)$$

where the product is over all primes dividing n .

This is easy to compute given the prime factorization of n , not so easy otherwise.

Definition

Let \mathbb{F} be a finite field and $f \in \mathbb{F}[x]$ irreducible. f is **primitive** if $x \bmod f$ is a generator of the multiplicative subgroup in the extension field $\mathbb{F}[x]/(f)$. The roots of a primitive polynomial are also called primitive.

The size of the multiplicative subgroup \mathbb{F}_q^\times is $q - 1$, $q = p^k$, and we know that the group is cyclic.

Hence there must be $\phi(q - 1)$ generators in this subgroup, corresponding to the number of primitive polynomials.

Since any of the roots of a corresponding primitive polynomial is a generator, the number of primitive polynomials of degree k is

$$\frac{\Phi(q-1)}{k}$$

For example, in the case $p = 5$, $k = 2$ there are 8 primitive elements and 4 polynomials.

There is an alternative way to describe primitive polynomials that avoids references to the extension field construction.

Definition

Let $f \in \mathbb{F}[x]$ such that $f(0) \neq 0$. The **order** or **exponent** of f is the least $e \geq 1$ such that f divides $x^e - 1$.

In other words, $x^e = 1 \pmod{f}$.

So an irreducible f is primitive iff it has order $p^k - 1$ where p is the characteristic and k the degree of f .

For example, $f = 2 + 4x + x^2$ is primitive.

α	x	α^{13}	$4x$
α^2	$3 + x$	α^{14}	$2 + 4x$
α^3	$3 + 4x$	α^{15}	$2 + x$
α^4	$2 + 2x$	α^{16}	$3 + 3x$
α^5	$1 + 4x$	α^{17}	$4 + x$
α^6	2	α^{18}	3
α^7	$2x$	α^{19}	$3x$
α^8	$1 + 2x$	α^{20}	$4 + 3x$
α^9	$1 + 3x$	α^{21}	$4 + 2x$
α^{10}	$4 + 4x$	α^{22}	$1 + x$
α^{11}	$2 + 3x$	α^{23}	$3 + 2x$
α^{12}	4	α^{24}	1

So $\mathbb{F}_{5^2}^*$ is indeed cyclic with generator α , and \mathbb{F}_{5^2} has dimension 2 as a vector space over \mathbb{F}_5 , as required.

A lot is known about primitive polynomials: 904,000 hits on google.

There are tables [Hansen](#).

There are well-analyzed algorithms:

Nirmal R. Saxena & Edward J. McCluskey

Primitive Polynomial Generation Algorithms—Implementation
and Performance Analysis

CRC 2004

A lot is known about special-form polynomials:

Richard P. Brent, Paul Zimmermann

Twelve new primitive binary trinomials

arXiv 2016

New primitive polynomials over \mathbb{F}_2 of degree 42, 643, 801; 43, 112, 609; and 74, 207, 281.

E. R. Berlekamp

Algebraic Coding Theory

McGraw-Hill, 1968.

R. Lidl, H. Niederreiter

Introduction to Finite Fields and their Applications

Cambridge University Press, 1986.