Chapter 10, Field Extensions

You are assumed to know Section 10.1. Everything you have learned in linear algebra applies regardless of what the field of scalars is. In particular, the definitions of **vector space, linear independence, basis** and **dimension** are unchanged. And the main theorems still hold, such as existence of a basis for any vector space and the fact that the number of elements in any basis is the same (called the dimension). Our main interest in vector spaces will be to apply them to situations where we have one field contained in another.

Assume that $F \subseteq K$ are fields. Then K is an additive abelian group and we can multiply elements of K by scalars (elements) from F satisfying the usual distributive and associative laws. Thus K is a vector space over F. In particular, there exists a basis of elements of K such that every element of K is a linear combination of elements of the basis with scalars from F. We have already worked with some examples.

$\mathbb{Q} \subseteq K = \mathbb{Q}(\sqrt{d})$ for any squarefree integer d. Then $K = \{\, a + b\sqrt{d} \mid a, b \in \mathbb{Q} \,\}$. But this just says that K has dimension 2 over $\mathbb{Q}$ with a basis $\{1, \sqrt{d}\}$. Another example we did long ago was $\mathbb{Q}(\sqrt[3]{2})$, a vector space over $\mathbb{Q}$ with basis $\{1, \sqrt[3]{2}, \sqrt[3]{4}\}$.

On the other hand, there are examples like $\mathbb{Q} \subseteq \mathbb{R}$ where the extension is infinite dimensional. We shall limit ourselves to finite dimensional extensions in this course. When K is a finite dimensional extension of F, we write $[K : F]$ for the dimension $\dim_F K$. We get two immediate results:

(1) $[K : F] = 1$ iff $K = F$.
This is a consequence of the fact that a one-dimensional vector space is the same as the field of scalars.

(2) (Theorem 10.5) Let K, L be finite dimensional extension fields of F and assume they are **isomorphic over F** (that is, there is an isomorphism $f\colon K \to L$ with $f(a) = a$ for all $a \in F$). It follows that $[K : F] = [L : F]$.
This is a consequence of the fact that isomorphic vector spaces have the same dimension.

There is one new result in section 1 and we shall have great need for it.

Theorem 10.4. *Let $F \subseteq K \subseteq L$ be fields. If $[K : F]$ and $[L : K]$ are finite, then $[L : F]$ is finite and $[L : F] = [L : K][K : F]$.*

Proof. We shall prove this by actually constructing a basis for L over F. Assume that we have bases $\{u_1, \dots u_m\}$ for K over F and $\{v_1, \dots v_n\}$ for L over K. We shall show that the set $\mathcal{B} = \{\, u_i v_j \mid 1 \le i \le m,\ 1 \le j \le n \,\}$ is a basis for L over F. Since $u_i v_j = u_k v_l$ implies $u_i v_j - u_k v_l = 0$ is a dependence relation in L with coefficients $u_i, -u_k \in K$, this cannot happen and the set has mn distinct elements. Therefore, if we can show $\mathcal{B}$ is a basis, we will have $[L : F] = mn = [L : K][K : F]$.

We must show two things: $\mathcal{B}$ spans L with coefficients from F and the elements are linearly independent. For the former, assume that $w \in L$. Then we can write w as a linear combination of the v_j's with coefficients from K, say $w = a_1v_1 + \cdots + a_nv_n$. Each $a_j \in K$, so it can be written as a linear combination of the u_i's with coefficients from F, say $a_j = b_{1j}u_1 + \cdots b_{mj}u_j$, $j = 1, \ldots, n$. Substituting these expressions into the equation for w gives w as a linear combination of elements u_iv_j with coefficients $b_{ij} \in F$.

Now assume we have a linear combination equal to zero: $\sum a_{ij}u_iv_j = 0$, $a_{ij} \in F$. Rearrange the terms to obtain

$$\sum_{j=1}^{n}\left(\sum_{i=1}^{m} a_{ij}u_i\right) v_j = 0.$$

Since the v_j's are linearly independent, each coefficient $\sum_{i=1}^{m} a_{ij}u_i \in K$ must be zero. But the u_i's are linearly independent over F, hence all $a_{ij} = 0$. Therefore $\mathcal{B}$ is a linearly independent set. □

The converse of this theorem also holds: if $[L : F]$ is finite, then both $[K : F]$ and $[L : K]$ are finite.

Proof. $[L : K]$ is finite: if $u_1, \ldots, u_n$ is a basis for L over F, then the set also spans L over K since $F \subseteq K$. The elements may no longer be linearly independent, but from linear algebra we know that any spanning set contains a basis, hence $[L : K] \leq [L : F]$.

$[K : F]$ is finite: K is a subspace of L (as vector spaces over F), so $\dim_F K \leq \dim_F L$, again from linear algebra facts (the proof is to start with 1_K and add elements of K as long as you can find one which is linearly independent of the set already obtained; this process must stop in at most $[L : F]$ steps or you would have a linearly independent subset of L with more elements than a basis). □

Simple Extensions.

We saw in Chapter 5 that we can always build extensions of a field F by forming the polynomial ring in a variable x and then factoring out by the principal ideal generated by an irreducible polynomial $p(x)$. In fact, we saw that $K = F[x]/(p(x))$ is a field which contains a root of $p(x)$, namely the image of x in K. There is also another point of view we can take, starting with a larger field and looking for its subfields: assume $F \subseteq K$ are fields and $u \in K$. We define $F(u)$ to be the intersection of all subfields of K containing both F and u. It is easy to see that any intersection of fields is again a field, so $F(u)$ is a field extension of F called a **simple extension** since it is generated by a single element. There are two possibilities:

(1) u satisfies some nonzero polynomial with coefficients in F, in which case we say u is **algebraic** over F and $F(u)$ is an **algebraic extension** of F.

(2) u is not the root of any nonzero polynomial over F, in which case we say u is **transcendental** over F and $F(u)$ is an **transcendental extension** of F.

Examples: $\sqrt{2}$ is algebraic over $\mathbb{Q}$. π is transcendental over $\mathbb{Q}$. πi is algebraic over $\mathbb{R}$, but transcendental over $\mathbb{Q}$.

If u is transcendental over F, then the homomorphism $F[x] \to F(u)$ defined by $x \mapsto u$ has kernel zero. Therefore the field of quotients $F(x)$ is isomorphic to $F(u)$. We will not pursue this case in this course. Our next two theorems show that the algebraic case is exactly the case mentioned above with a quotient ring of $F[x]$.

Theorem 10.6. *Let K be an extension field of F and $u \in K$ an algebraic element over F. There exists a unique irreducible monic polynomial $p(x) \in F[x]$ with u as a root. For any polynomial $g(x) \in F[x]$, if $g(u) = 0$, then $p(x)$ divides $g(x)$. We call $p(x)$ the* **minimal polynomial** *of u over F.*

Proof. Consider the homomorphism $F[x] \to K$ defined by evaluation of a polynomial at u. Since the image is a subring of a field, the kernel is a prime ideal in the PID $F[x]$, say $(p(x))$. We know that p is not the zero polynomial since there is some polynomial which u satisfies. Multiplying by a constant, we may assume that $p(x)$ is monic. It is irreducible because the image is an integral domain (being a subring of a field; by Theorem 5.11, the image is actually a field). Any polynomial over F with u as a root is in the ideal, hence is a multiple of $p(x)$. The only irreducible polynomials in the ideal are the associates of $p(x)$, and only one of those, namely $p(x)$ itself, is monic. Thus we have the uniqueness of $p(x)$. □

Theorem 10.7. *Let K be an extension field of F and $u \in K$ an algebraic element over F with minimal polynomial $p(x)$ of degree n. Then*

(1) *$F(u) \cong F[x]/(p(x))$;*

(2) *$\{1, u, u^2, \ldots, u^{n-1}\}$ is a basis of the vector space $F(u)$ over F; and therefore*

(3) *$[F(u) : F] = n$.*

Proof. (1) was proved in the proof of Theorem 10.6 since $F(u)$ is the image of the homomorphism in that proof; indeed, it clearly maps onto $F[u]$, but since the image is a field, it must actually be $F(u)$ (in fact, this shows they are equal). (2) was discussed at the end of Chapter 5. By the division algorithm we can write any $f(x) \in F[x]$ in the form $f(x) = p(x)q(x) + r(x)$ where $r(x) = 0$ or has degree less than $\deg p(x)$. Thus $f(x) \equiv r(x) \pmod{p(x)}$, and is thus written in terms of the powers of $x + (p(x))$; the isomorphic image of this is u, so all elements of $F(u)$ are written as linear combinations of powers of u. Since $\deg r(x) < n = \deg p(x)$, we only need the powers up to $n-1$. $\square$

Example 1. We consider $K = \mathbb{Q}(\sqrt{2}, \sqrt{3})$. We wish to use Theorem 10.4 to show that it has dimension 4 over $\mathbb{Q}$; then it is easily seen that $\{1, \sqrt{2}, \sqrt{3}, \sqrt{6}\}$ is a basis. $[\mathbb{Q}(\sqrt{2}) : \mathbb{Q}] = 2$; since $x^2 - 3$ has degree 2 over $\mathbb{Q}$, either $[K : \mathbb{Q}(\sqrt{2})] = 2$ (and we get the desired 4 over $\mathbb{Q}$, or $[K : \mathbb{Q}(\sqrt{2})] = 1$. This latter can only happen if $x^2 - 3$ is reducible over $\mathbb{Q}(\sqrt{2})$. We know its factors over $\mathbb{R}$, so it is only reducible if $\pm\sqrt{3} \in \mathbb{Q}(\sqrt{2})$. Assume $\sqrt{3} = a + b\sqrt{2}$, square both sides and you get $\sqrt{2} \in \mathbb{Q}$, a contradiction. Therefore K has dimension 4 over $\mathbb{Q}$.

We next look at K from the standpoint of the previous theorem. Let $u = \sqrt{2} + \sqrt{3}$. Then $\mathbb{Q}(u)$ is a subfield of K, and therefore has dimension 2 or 4 over $\mathbb{Q}$ (it isn't 1 and it divides 4). Find its minimal polynomial: $u^2 = 5 + 2\sqrt{6}$; $u^4 - 10u^2 + 25 = (u^2 - 5)^2 = 24$, so u satisfies $x^4 - 10x^2 + 1$. Is it irreducible? Eisenstein's criterion does not apply. It is reducible modulo 2 ($(x+1)^4$) and modulo 3 ($(x^2+1)^2$) and modulo 5 ($(x^2+2)(x^2-2)$)... It is irreducible modulo 13, but that would be hard to show by hand. From the mod 3 case, we see that it has no linear factors. Thus if it is reducible, we have $x^4 - 10x^2 + 1 = (x^2 + ax + b)(x^2 + cx + d) = x^4 + (a+c)x^3 + (ac + b + d)x^2 + (ad + bc)x + bd$, which implies

$$\begin{aligned} a &= -c \\ -10 &= -a^2 + b + d \\ a(d-b) &= 0 \\ bd &= 1 \end{aligned}$$

Since $bd = 1$, we have $b = d = \pm 1$ so $b + d = \pm 2 \neq a^2 - 10$. Therefore the polynomial is irreducible and is thus the minimal polynomial of u. Therefore, by Theorem 10.7(2), another basis for K over $\mathbb{Q}$ is $\{1, u, u^2, u^3\} = \{1, \sqrt{2} + \sqrt{3}, 5 + 2\sqrt{6}, 11\sqrt{2} + 9\sqrt{3}\}$.

Theorem 10.7 has a corollary that we shall make great use of in the future as we deal with isomorphisms of algebraic extensions of a field F. An immediate consequence is that if two elements u, v in some extension field of F satisfy the same irreducible polynomial over F, then $F(u) \cong F(v)$, as they are both isomorphic to the same quotient of $F[x]$. For example $x^4 - 2$ has roots $\sqrt[4]{2}$ and $i\sqrt[4]{2}$, so these generate isomorphic extensions of $\mathbb{Q}$, even though we normally think of one as being a subfield of $\mathbb{R}$ and the other as being a subfield of $\mathbb{C}$ that is not contained in $\mathbb{R}$. We can generalize this idea to isomorphic base fields, rather than a fixed base field F.

Corollary 10.8. *Let $\sigma\colon F \to E$ be an isomorphism of fields. Let u be algebraic over F with minimal polynomial $p(x) \in F[x]$. Let v be algebraic over E with minimal polynomial*

$\sigma(p(x)) \in E[x]$ [in the sense that $\sigma\colon F \to E$ has a unique extension to an isomorphism $\sigma\colon F[x] \to E[x]$ defined by applying the isomorphism to the coefficients of polynomials]. Then σ extends to an isomorphism of fields $\bar{\sigma}\colon F(u) \to E(v)$ such that $\bar{\sigma}(u) = v$.

Proof. Consider the composition $F[x] \cong E[x] \to E[x]/(\sigma p(x)) \cong E(v)$. It is surjective since each mapping is. The kernel contains $p(x)$ since $p(x) \mapsto \sigma p(x) \mapsto 0$. On the other hand any element of the kernel is in $\sigma^{-1}(\ker(E[x] \to E[x]/(\sigma p(x)))) = \sigma^{-1}(\sigma p(x)) = (p(x))$. Thus we obtain an isomorphism of $F[x]/(p(x)) \cong E(v)$. But $F[x]/(p(x)) \cong F(u)$, so we are done. □

Algebraic extensions.

We generalize the idea of simple algebraic extensions as follows:

Definition, page 347. An extension field K of F is called an **algebraic extension** of F if every element of K is algebraic over F.

This now allows algebraic extensions of infinite dimension such as $\mathbb{Q}(\sqrt{2}, \sqrt{3}, \sqrt{5}, \dots)$ over $\mathbb{Q}$. (At least Example 1 suggests that it should be infinite dimensional because one would expect that none of the square roots would interact, just as they didn't for $\sqrt{2}$ and $\sqrt{3}$.)

Theorem 10.9. *If $n = [K : F] < \infty$, then K is an algebraic extension of F.*

Proof. Let $u \in K$. The set $\{1, u, u^2, \dots, u^n\}$ has $n + 1 > [K : F]$ elements, so must be linearly dependent over F. The dependence relation gives a polynomial satisfied by u. □

The contrapositive says that if K contains any transcendental element v, then the extension is infinite dimensional. We also know this is true, since then the homomorphism $F[x] \to K$ carrying $x \mapsto v$ has kernel zero.

An extension of a field F is called **finitely generated** if it has the form $F(u_1, u_2, \dots, u_n)$, defined as the intersection of all subfields of a field K which contains $F \cup \{u_1, u_2, \dots, u_n\}$. This is really just an iterated form of our earlier definition since $F(u_1, u_2, \dots, u_n) = F(u_1, u_2, \dots, u_{n-1})(u_n)$, and thus the elements can be added to F one at a time. It turns out that for algebraic extensions, being finitely generated is actually the same as being finite dimensional.

Theorem 10.10 and converse. *Let K be an algebraic extension of F. $[K : F] < \infty$ if and only if $K = F(u_1, u_2, \dots, u_n)$ for some elements $u_1, u_2, \dots, u_n \in K$.*

Proof. ($\Longrightarrow$) Let u_1 be any element of K not in F. Then $F \subsetneq F(u_1)$. If $F(u_1) = K$, we are done; otherwise, choose $u_2 \in K$, not in $F(u_1)$, so that $F(u_1) \subsetneq F(u_1, u_2)$. Continue this process. At each step the dimension over F increases. Since $[K : F]$ is finite, the process must terminate in a finite number of steps, say n, yielding $K = F(u_1, u_2, \ldots, u_n)$.

($\Longleftarrow$) Now assume that $K = F(u_1, u_2, \ldots, u_n)$. For each k, the extension $F(u_1, u_2, \ldots, u_k)$ of $F(u_1, u_2, \ldots, u_{k-1})$ is a simple extension, hence has finite dimension by Theorem 10.7. Iterating Theorem 10.4 yields $[K : F] = [K : F(u_1, u_2, \ldots, u_{n-1})][F(u_1, u_2, \ldots, u_{n-1}) : F(u_1, u_2, \ldots, u_{n-2})] \cdots [F(u_1) : F]$, which is finite. □

It is common to refer to extensions satisfying Theorem 10.10 simply as **finite extensions**.

A very important example of an infinite dimensional algebraic extension is the set of all elements of $\mathbb{C}$ which are algebraic over $\mathbb{Q}$. This is called the set of **algebraic numbers**. Another example is the set of all elements of $\mathbb{R}$ which are algebraic over $\mathbb{Q}$. This is called the set of **real algebraic numbers**. But why are these fields? That is, why is the sum and product of algebraic elements again algebraic? It is certainly not easy to see what the minimal polynomials might be like. But our theory using vector spaces now makes it clear that it is so: for indeed, if $u, v \in K$ are algebraic over F, then $F(u, v)$ is a finite dimensional vector space over F; and thus the subspace $F(u - v)$ (or $F(uv^{-1})$) must also be finite dimensional over F. By Theorem 10.9, it is an algebraic extension, so its element $u - v$ is algebraic over F. This is the essence of Corollary 10.12 in the book. Note that this gives no information about the converse: we have no idea whether $e + \pi$ is algebraic or transcendental over $\mathbb{Q}$.

Exercise 11, page 351. Let $u, v \in K$ be algebraic over a subfield F with minimal polynomials $p(x)$ and $q(x)$ of degrees m, n, respectively. Assume first that $\gcd(m, n) = 1$. We claim that $[F(u, v) : F] = mn$. We know that $[F(u, v) : F] = [F(u, v) : F(u)][F(u) : F]$ by Theorem 10.4, hence $m \mid [F(u, v) : F]$. Similarly, $n \mid [F(u, v) : F]$, so $mn \mid [F(u, v) : F]$ because m and n are relatively prime. On the other hand, $[F(u, v) : F(u)] \leq n$ since the minimal polynomial for v over $F(u)$ can have no greater degree than the minimal polynomial $q(x)$ over F. (In fact, the minimal polynomial over $F(u)$ must divide $q(x)$ by Theorem 10.6.) Therefore $[F(u, v) : F] = [F(u, v) : F(u)][F(u) : F] \leq mn$. Since $mn \mid [F(u, v) : F]$, they must be equal.

This sometimes holds if $\gcd(m, n) \neq 1$ as in our Example 1 with $\sqrt{2}$ and $\sqrt{3}$. But it may also fail. A trivial example is given by taking $u = v$. A nontrivial example comes from $\mathbb{Q}(\sqrt{2}, \sqrt[4]{2})$ whose degree over $\mathbb{Q}$ is only 4 since it equals $\mathbb{Q}(\sqrt[4]{2})$.

An example using this result is that $[\mathbb{Q}(\sqrt{2}, \sqrt[3]{2}) : \mathbb{Q}] = 6$ since $\gcd(2, 3) = 1$.

Exercise 13, page 351, is closely related to Exercise 11. Let $m = \deg p(x)$ and $n =$

$\deg q(x)$. We saw above that $[F(u,v):F] = mn = [F(u,v):F(u)][F(u):F]$. This implies that $[F(u,v):F(u)] = n$, which equals the degree of the minimal polynomial of v over $F(u)$ by Theorem 10.7. By Theorem 10.6, this polynomial divides $q(x)$ (since $q(v) = 0$ in $F(u)$), and therefore must be $q(x)$ since they have the same degree. That is, $q(x)$ remains irreducible over $F(u)$. In our specific example, this says $x^3 - 2$ is irreducible over $\mathbb{Q}(\sqrt{2})$.

Splitting fields.

Let $F \subseteq K$ be fields and let $f(x) \in F[x]$. We say that $\mathbf{f(x)}$ **splits over K** if it factors into linear factors in $K[x]$. If $\deg f(x) = n$, this means that $f(x)$ has n roots in K counting multiplicities (i.e., for $f(x) = (x+1)^2$, we count the root 1 twice). Given a polynomial $f(x) \in F[x]$, we are interested in constructing the smallest field K containing F and all the roots of $f(x)$. We call this field the **splitting field** of $f(x)$ over F. Note that if $f(x) = c(x-u_1)\cdots(x-u_n)$ in $K[x]$ and K is the splitting field, then $K = F(u_1, \ldots, u_n)$ since it is generated by F together with the roots of $f(x)$.

Examples: F is the splitting field for every linear polynomial over F.
$\mathbb{C}$ is the splitting field for any irreducible quadratic polynomial over $\mathbb{R}$. In particular, this is true for the polynomial x^2+1.
$K = \mathbb{Q}(\sqrt{2}, \sqrt{3})$ is the splitting field for $x^4 - 10x^2 + 1$ over $\mathbb{Q}$ by Example 1. K is also the splitting field for $x^2 - 3$ over $\mathbb{Q}(\sqrt{2})$. The polynomial does not need to be irreducible. K is the splitting field for $(x^2-2)(x^2-3)(x+1)$ over $\mathbb{Q}$ as well. However, our main interest is in irreducible polynomials since then we know more about the dimension of K over F.

Our immediate goal is to show that splitting fields always exist and are unique up to isomorphism. We will then go on to see that they have much stronger and more surprising properties.

Theorem 10.13. *Let F be a field and let $f(x) \in F[x]$ be a polynomial of degree $n > 0$. Then there exists a splitting field K of $f(x)$ over F with $[K:F] \leq n!$.*

Proof. We construct K using induction on n. If $n = 1$, then $K = F$ works since the root of $f(x)$ already lies in F. Assume that the theorem holds for polynomials of degree $n-1$ and $n \geq 2$. Since $F[x]$ is a UFD, we can find a monic, irreducible factor $p(x)$ of $f(x)$ in $F[x]$. We can construct the field $F[x]/(p(x))$ which has a root u of $p(x)$ (Theorem 5.11). Call this field $F(u)$. Its dimension over F is equal to $\deg p(x) \leq \deg f(x) = n$. Over the field $F(u)$, we factor $f(x) = (x-u)g(x)$ for some $g(x) \in F(u)[x]$. By the induction hypothesis, there exists a splitting field K for $g(x)$ over $F(u)$ with $[K:F(u)] \leq (n-1)!$. But then K has all the roots of $f(x)$ and is generated by those roots, hence is a splitting field of $f(x)$ over F. Furthermore, $[K:F] = [K:F(u)][F(u):F] \leq n!$. □

As you might guess, splitting fields for a given polynomial are all isomorphic. In fact, they are unique if you are working inside some large enough field to begin with; that is, if you want the splitting field of $f(x) \in \mathbb{Q}[x]$, there is a unique choice inside $\mathbb{C}$. But we do not always have a big field available, as was the situation when we constructed a splitting field in the previous theorem. So in general, the best we can ask for is isomorphism. In an abstract sense, with the base field $\mathbb{Q}$, we can't tell $\sqrt{2}$ from $-\sqrt{2}$; they are simply the two roots of the irreducible polynomial x^2-2. Of course, they have precise meanings inside a bigger field like $\mathbb{R}$.

Theorem 10.14. *Let $\sigma\colon F \to E$ be a field isomorphism, $f(x) \in F[x]$ with $\deg f = n > 0$, and $\sigma(f(x))$ the corresponding polynomial in $E[x]$. If K is a splitting field of $f(x)$ over F and L is a splitting field of $\sigma(f(x))$ over E, then σ extends to an isomorphism $K \cong L$.*

Proof. Again we induct on n. If $n = 1$, then $K = F$ and $L = E$ so σ is the required isomorphism. Now assume that $n \geq 2$ and the theorem holds for polynomials of degree $n-1$. Let $p(x)$ be a monic irreducible polynomial which divides $f(x)$ in $F[x]$. From the isomorphism $F[x] \cong E[x]$, the polynomial $\sigma p(x)$ is a monic irreducible divisor of $\sigma f(x)$. Let u be a root of $p(x)$ in K and v a root of $\sigma p(x)$ in L (since these polynomials must split in the respective fields). Corollary 10.8 says σ extends to an isomorphism $F(u) \cong E(v)$ carrying u to v. Now factor $f(x) = (x-u)g(x)$ over $F(u)$, and correspondingly, $\sigma f(x) = (x-v)\sigma g(x)$ over $E(v)$. Apply the induction hypothesis to the polynomial $g(x) \in F(u)[x]$ of degree $n-1$ which splits in K. We obtain an isomorphism $K \cong L$ extending the isomorphism $F(u) \cong E(v)$ which, in turn, extends σ. □

Note that an explicit isomorphism can be constructed step by step as in the proof. For example, if $f(x) = (x^2-2)(x^2-3) \in \mathbb{Q}[x]$, with $F = E = \mathbb{Q}$, then we first take an irreducible factor $p(x) = x^2-2$ and construct an isomorphism $F(u) = \mathbb{Q}(\sqrt{2}) \to E(v) = \mathbb{Q}(-\sqrt{2})$ taking $u = \sqrt{2}$ to $v = -\sqrt{2}$. Note that I have $\deg p$ choices for where to send u. Then I factor $f(x) = (x-\sqrt{2})(x^3+\sqrt{2}x^2-3x-3\sqrt{2})$ and choose an irreducible factor of the second factor, like x^2-3, and repeat the process once more to get an isomorphism $\mathbb{Q}(\sqrt{2},\sqrt{3}) \cong \mathbb{Q}(-\sqrt{2},-\sqrt{3})$ (if I happen to make that choice for u and v the second time). Note in particular that it did not require $n = 4$ steps in this case—sometimes you get extra roots when you adjoin one of them, but not always. Of course, you always do for quadratic polynomials.

Example 2. Let's find the splitting field of x^3-2 over $\mathbb{Q}$. It certainly contains $u_1 = \sqrt[3]{2}$. It also contains the other two cube roots of 2 in $\mathbb{C}$, namely $u_2 = \sqrt[3]{2}\frac{-1+\sqrt{3}i}{2}$ and $u_3 = \sqrt[3]{2}\frac{-1-\sqrt{3}i}{2}$. So the splitting field is $\mathbb{Q}(u_1,u_2,u_3)$. What is its dimension over $\mathbb{Q}$? Notice that the two complex cube roots of 1, $\frac{-1\pm\sqrt{3}i}{2}$ are roots of the equation $x^3-1 = (x-1)(x^2+x+1)$, so both have the minimal polynomial x^2+x+1. Thus the splitting field can also be written as $\mathbb{Q}(\sqrt[3]{2},\omega)$, where $\omega = \frac{-1+\sqrt{3}i}{2}$. Since these two elements have minimal polynomials of the relatively prime degrees 3 and 2, respectively, the dimension over $\mathbb{Q}$ is the product, 6, by Exercise 13 done earlier.

The most important definition for studying Galois theory is the following:

Definition. An algebraic extension K of F is **normal** if whenever an irreducible polynomial over F has one root in K, then it splits in K.

This would seem to be an extremely strong condition since it refers to all polynomials, but in fact it applies to all splitting fields! That is, as soon as a field K is a splitting field for one polynomial over F, it is also a splitting field for any irreducible polynomial over F with a single root in K.

Theorem 10.15. *K is a splitting field for some polynomial over F if and only if K is a finite dimensional normal extension of F.*

Proof. $(\implies)$ We have $K = F(u_1, \dots u_n)$ where $u_1, \dots, u_n$ are the roots of some polynomial $f(x) \in F[x]$. By Theorem 10.10, $[K : F] < \infty$. To show that K is a normal extension of F, let $p(x)$ be an irreducible polynomial over F with a root $v \in K$ and let L be the splitting field of $p(x)$ over K. Our goal is to show that $L \subseteq K$, or more specifically, every root w of $p(x)$ lies in K. Since $p(x)$ is irreducible, Corollary 10.8 implies that $F(v) \cong F(w)$ via an isomorphism which fixes F. Now $K(w) = F(w)(u_1, \dots, u_n)$, hence is a splitting field for $f(x)$ over $F(w)$. Since K is a splitting field for $f(x)$ over F, it is also a splitting field for $f(x)$ over $F(v) \subseteq K$. By Theorem 10.14, the isomorphism $F(v) \cong F(w)$ extends to an isomorphism of splitting fields for $f(x)$: $K \xrightarrow{\sim} K(w)$ taking $v \mapsto w$ and fixing the subfield F. But then these two vector spaces over F have the same dimension, and since one contains the other, they are equal. Therefore $w \in K$ for every root w of $p(x)$ as desired.

$(\impliedby)$ Since K is a finite dimensional extension of F, it can be written $K = F(u_1, \dots, u_n)$ with each u_i satisfying some minimal polynomial $p_i(x)$. Each $p_i(x)$ splits in K since it is normal, and therefore $f(x) = p_1(x) \cdots p_n(x)$ splits over K. It follows that K is the splitting field for $f(x)$ over F. □

Some of this theory can be extended to infinite algebraic extensions. One can construct an algebraic extension $\bar{F}$ of a field F in which every polynomial over F splits. Theorem 10.14 can be extended to show that such a field is unique up to isomorphism; $\bar{F}$ is called the **algebraic closure** of F. An example is $\bar{\mathbb{Q}}$, the field of all algebraic numbers, which is the algebraic closure of $\mathbb{Q}$. Any field, for which every polynomial over it splits over the field itself, is called **algebraically closed**. Another algebraically closed field is $\mathbb{C}$, a fact usually proved using complex analysis. The algebra proofs require the intermediate value theorem for $\mathbb{R}$ plus some Galois theory (Chapter 11).

Exercise 13, page 358: A splitting field for x^6+x^3+1 over $\mathbb{Q}$ is given by $\mathbb{Q}(\rho, \sqrt[3]{\rho}, \sqrt[3]{\bar\rho})$ because $x^6+x^3+1=(x^3-\rho)(x^3-\bar\rho)$, where $\rho=\frac{-1+\sqrt{3}i}{2}$ is a cube root of 1.

Exercise 16, page 358: A splitting field for x^3+x+1 over $\mathbb{Z}_2$ is given by $\mathbb{Z}_2(\rho)$ where ρ is a root of x^3+x+1 because this gives a field of 8 elements $\{0,1,\rho,\rho+1,\rho^2,\rho^2+1,\rho^2+\rho,\rho^2+\rho+1\}$ over which $x^3+x+1=(x+\rho)(x+\rho^2)(x+\rho^2+\rho)$. This field is commonly denoted by $\mathbb{F}_8$.

Separability.

We say a polynomial of degree n is **separable** if it has n *distinct* roots in some splitting field. Thus x^2-1 is separable over $\mathbb{R}$, but $(x^2+1)^2$ is not since i is a multiple root in the splitting field $\mathbb{C}$. Our real concern is when irreducible polynomials are separable, and the answer is pretty much always in the cases we will encounter in this course. Our interest in this is that it is a needed condition for our later work; inseparability leads to considerably different results. Consequently, we define an element of an extension field K of F to be **separable over F** if it is algebraic and its minimal polynomial is separable. And we define an extension field K to be **separable over F** if every element of K is separable over F.

The main test for separability uses the derivative. Of course, we do not have the usual limits of calculus available, so we define the derivative formally:

$$\text{For } f(x)=a_0+a_1x+\cdots+a_nx^n, \text{ the derivative is } f'(x)=a_1+2a_2x+\cdots+na_nx^{n-1}.$$

The usual sum, product and chain rule formulas hold for polynomials; this can be proved using this definition, but the underlying reason is that they hold over $\mathbb{R}$ and they are purely formal formulas, so they must also hold over any field F. Note that there are some differences from the usual theorems over $\mathbb{R}$. For example, over $\mathbb{Z}_p$, the derivative of x^p+1 is zero!

Lemma 10.16 and converse. *The polynomial $f(x)\in F[x]$ is separable iff* $\gcd(f,f')=1$.

Proof. Let K be a splitting field for $f(x)$ and let $a\in K$ be any root of $f(x)$. Then $f(x)=(x-a)^m g(x)$ with $m\geq 1$, $g(a)\neq 0$. Now $f'(x)=m(x-a)^{m-1}g(x)+(x-a)^m g'(x)=(x-a)^{m-1}[mg(x)+(x-a)g'(x)]$, where the second factor, evaluated at a, gives $mg(a)+0=0 \iff m=0\in K$. Now if $f(x)$ is separable then every power m is 1, so $f'(a)\neq 0$ and $x-a$ is not a factor of $f'(x)$. Thus none of the factors of $f(x)$ divides $f'(x)$ and they must be relatively prime. Conversely, if $f(x)$ is not separable, then some $m>1$ and $f'(a)=0$, hence $x-a$ is a common factor of $f(x)$ and $f'(x)$, so they are not relatively prime. □

Recall that we discussed the **characteristic** of a ring R last semester. For a ring with 1_R, it was seen that it is the number $n\geq 0$ such that the homomorphism $\mathbb{Z}\to R$ defined by $f(k) = k1_R$ has kernel (n). When $R = K$ is a field, the image of $\mathbb{Z}$ must be an integral domain, so n is a prime or 0. Thus every field has characteristic either 0 (in which case it contains a copy of the rational numbers) or a prime p (in which case it contains a copy of the finite field $\mathbb{Z}_p$. The field $\mathbb{Q}$ or $\mathbb{Z}_p$ is called the **prime subfield** of K.

Theorem 10.17. *Every irreducible polynomial over a field of characteristic zero is separable, and hence every algebraic extension is separable.*

Proof. Let $p(x) = a_n x^n + \cdots + a_0$, $a_n \neq 0$ be irreducible. Then $p'(x) = na_n x^{n-1} = \cdots + a_1$ has degree $n - 1$, so is nonzero and has no factor in common with $p(x)$ since $p(x)$ is irreducible. Thus $\gcd(p, p') = 1$ and $p(x)$ is separable. □

What goes wrong in characteristic p? For finite fields nothing goes wrong. The simplest example of a nonseparable extension is given in Exercise 15, page 363. This works for any p as well as 2. Let $F = \mathbb{Z}_p(t)$ be the field of quotients of the ring of polynomials in one variable over $\mathbb{Z}_p$. Consider the polynomial $f(x) = x^p - t$. t is an irreducible element of $\mathbb{Z}_p[t]$, so this polynomial is irreducible by Eisenstein's criterion. But $f'(x) = 0$, so $\gcd(f, f') = f$ and $f(x)$ is not separable. In fact, what happens in a splitting field, is that if u is any root, then $f(x) = (x - u)^p$; that is, u is the only root and it has multiplicity p.

Finite Fields.

Like finite groups, these have lots of applications and are a major part of a course in applied algebra. They are used in combinatorics, coding theory, cryptography, projective geometry, etc. In order to make this chapter independent of ring theory, the author has done a lot of ring theory at the beginning of the section—which we skip. Note that a finite field K is also a finite abelian group (under addition) and we again use the word **order** for the number of elements in it. Also, K can be thought of as a vector space over its prime subfield $\mathbb{Z}_p$ of some dimension n. Thus, as a vector space, $K \cong \mathbb{Z}_p^n$ and has p^n elements. This proves

Theorem 10.23. *If K is a finite field of characteristic p, then $|K| = p^n$ where $[K : \mathbb{Z}_p] = n$.* □

Arithmetic in characteristic p

$$(a + b)^{p^n} = a^{p^n} + b^{p^n}.$$

Proof. This is just the binomial theorem plus a fact about binomial coefficients. Use induction on n. If $n = 1$, then it follows from the fact that all the other coefficients in the binomial expansion have the form $\binom{p}{k} = \dfrac{p!}{k!(p-k)!}$ for $0 < k < p$; thus the denominator is not divisible by p, but the numerator is, and so the coefficient is 0 in $\mathbb{Z}_p$. Now do the obvious inductive step:

$$(a+b)^{p^{n+1}} = ((a+b)^{p^n})^p = (a^{p^n} + b^{p^n})^p = a^{p^{n+1}} + b^{p^{n+1}}. \qquad \square$$

Theorem 10.25 (Characterization of Finite Fields). *For each prime p and integer $n \geq 1$, there is a unique (up to isomorphism) field of order p^n. It is the splitting field of $x^{p^n} - x$ over $\mathbb{Z}_p$.*

Proof. We know that $x^{p^n} - x$ has some splitting field K over $\mathbb{Z}_p$ and it is unique up to isomorphism. Since $\frac{d}{dx}x^{p^n} - x = -1$, it is relatively prime to $x^{p^n} - x$, which thus has no repeated roots; so $x^{p^n} - x$ has precisely p^n different roots c and they all satisfy $c^{p^n} = c$. Let S be the set of roots; for $a, b \in S$, we have $(ab)^{p^n} = a^{p^n} b^{p^n} = ab$, so $ab \in S$, $(a^{-1})^{p^n} = (a^{p^n})^{-1} = a^{-1}$, so $a^{-1} \in S$, and $(a-b)^{p^n} = a^{p^n} - b^{p^n} = a - b$, so $a - b \in S$. It follows that S is a field, and is certainly generated by the roots of $x^{p^n} - x$, so $S = K$. Thus there exists a field with p^n elements and it is the splitting field of $x^{p^n} - x$ over $\mathbb{Z}_p$. Now let L be any other field with p^n elements. It's multiplicative group of nonzero elements L^* has $p^n - 1$ elements. By Lagrange's Theorem, any nonzero element of $c \in L$ satisfies $c^{p^n - 1} = 1$, so c is a root of $x^{p^n - 1} - 1$. Therefore every element of L satisfies $x(x^{p^n-1} - 1) = x^{p^n} - x$; it follows that L is also a splitting field for $x^{p^n} - x$, so is isomorphic to K. $\square$

We denote the field of order p^n by $\mathbb{F}_{p^n}$ or $GF(p^n)$ and call it the **Galois field of order $\mathbf{p^n}$**. Note that we have actually shown that

$$x^{p^n} - x = \prod_{a \in \mathbb{F}_{p^n}} (x - a).$$

While $x^{p^n} - x$ is certainly not irreducible, we can use group theory to show that there always are irreducible polynomials of order p^n over $\mathbb{Z}_p$.

Theorem 10.28. *$\mathbb{F}_{p^n}$ is a simple extension of $\mathbb{Z}_p$. There exists an irreducible polynomial of degree n over $\mathbb{Z}_p$.*

Proof. By Theorem 7.15, the multiplicative group of nonzero elements of $\mathbb{F}_{p^n}$ is cyclic. If u is a generator of this group, then $\mathbb{F}_{p^n} = \{0, u, u^2, \ldots, u^{p^n - 1}\} = \mathbb{Z}_p(u)$. The minimal polynomial of u is irreducible over $\mathbb{Z}_p$ and has degree $[\mathbb{F}_{p^n} : \mathbb{Z}_p] = n$. $\square$

When is $K = \mathbb{F}_{p^m}$ contained in $L = \mathbb{F}_{p^n}$? If $K \subseteq L$, then $[L : K]m = [L : K][K : \mathbb{Z}_p] = [L : \mathbb{Z}_p] = n$, so $m \mid n$. On the other hand, if $m \mid n$, say $n = mr$, then $c^{p^m} = c \implies c^{p^n} =$

$c^{p^{mr}} = (c^{p^m})^{p^{m(r-1)}} = c^{p^{m(r-1)}} = \cdots = c$, so any element of K is in L because it satisfies the appropriate polynomial.