Chapter 3

Field Fundamentals

3.1 Field Extensions

If F is a field and $F[X]$ is the set of all *polynomials over* F, that is, polynomials with coefficients in F, we know that $F[X]$ is a Euclidean domain, and therefore a principal ideal domain and a unique factorization domain (see Sections 2.6 and 2.7). Thus any nonzero polynomial f in $F[X]$ can be factored uniquely as a product of irreducible polynomials. Any root of f must be a root of one of the irreducible factors, but at this point we have no concrete information about the existence of roots and how they might be found. For example, X^2+1 has no real roots, but if we consider the larger field of complex numbers, we get two roots, $+i$ and $-i$. It appears that the process of passing to a larger field may help produce roots, and this turns out to be correct.

3.1.1 Definitions

If F and E are fields and $F \subseteq E$, we say that E is an *extension* of F, and we write $F \leq E$, or sometimes E/F.

If E is an extension of F, then in particular E is an abelian group under addition, and we may multiply the "vector" $x \in E$ by the "scalar" $\lambda \in F$, and the axioms of a vector space are satisfied. Thus if $F \leq E$, then E is a vector space over F. The dimension of this vector space is called the *degree* of the extension, written $[E : F]$. If $[E : F] = n < \infty$, we say that E is a *finite extension* of F, or that *the extension* E/F *is finite*, or that E is *of degree* n *over* F.

If f is a nonconstant polynomial over the field F, and f has no roots in F, we can always produce a root of f in an extension field of F. We do this after a preliminary result.

3.1.2 Lemma

Let $f\colon F \to E$ be a homomorphism of fields, i.e., $f(a+b) = f(a)+f(b)$, $f(ab) = f(a)f(b)$ (all $a, b \in F$), and $f(1_F) = 1_E$. Then f is a monomorphism.

Proof. First note that *a field F has no ideals except $\{0\}$ and F*. For if a is a nonzero member of the ideal I, then $ab = 1$ for some $b \in F$, hence $1 \in I$, and therefore $I = F$. Taking I to be the kernel of f, we see that I cannot be all of F because $f(1) \neq 0$. Thus I must be $\{0\}$, so that f is injective. ♣

3.1.3 Theorem

Let f be a nonconstant polynomial over the field F. Then there is an extension E/F and an element $\alpha \in E$ such that $f(\alpha) = 0$.

Proof. Since f can be factored into irreducibles, we may assume without loss of generality that f itself is irreducible. The ideal $I = \langle f(X) \rangle$ in $F[X]$ is prime (see (2.6.1)), in fact maximal (see (2.6.9)). Thus $E = F[X]/I$ is a field by (2.4.3). We have a problem at this point because F need not be a subset of E, but we can place an isomorphic copy of F inside E via the homomorphism $h\colon a \to a + I$; by (3.1.2), h is a monomorphism, so we may identify F with a subfield of E. Now let $\alpha = X + I$; if $f(X) = a_0 + a_1X + \cdots + a_nX^n$, then

$$\begin{aligned} f(\alpha) &= (a_0 + I) + a_1(X + I) + \cdots + a_n(X + I)^n \\ &= (a_0 + a_1X + \cdots + a_nX^n) + I \\ &= f(X) + I \end{aligned}$$

which is zero in E. ♣

The extension E is sometimes said to be obtained from F by *adjoining a root* α of f. Here is a further connection between roots and extensions.

3.1.4 Proposition

Let f and g be polynomials over the field F. Then f and g are relatively prime if and only if f and g have no common root in any extension of F.

Proof. If f and g are relatively prime, their greatest common divisor is 1, so there are polynomials $a(X)$ and $b(X)$ over F such that $a(X)f(X) + b(X)g(X) = 1$. If α is a common root of f and g, then the substitution of α for X yields $0 = 1$, a contradiction. Conversely, if the greatest common divisor $d(X)$ of $f(X)$ and $g(X)$ is nonconstant, let E be an extension of F in which $d(X)$ has a root α (E exists by (3.1.3)). Since $d(X)$ divides both $f(X)$ and $g(X)$, α is a common root of f and g in E. ♣

3.1.5 Corollary

If f and g are distinct monic irreducible polynomials over F, then f and g have no common roots in any extension of F.

Proof. If h is a nonconstant divisor of the irreducible polynomials f and g, then up to multiplication by constants, h coincides with both f and g, so that f is a constant multiple of g. This is impossible because f and g are monic and distinct. Thus f and g are relatively prime, and the result follows from (3.1.4). ♣

If E is an extension of F and $\alpha \in E$ is a root of a polynomial $f \in F[X]$, it is often of interest to examine the field $F(\alpha)$ generated by F and α, in other words the smallest subfield of E containing F and α (more precisely, containing all elements of F along with α). The field $F(\alpha)$ can be described abstractly as the intersection of all subfields of E containing F and α, and more concretely as the collection of all rational functions

$$\frac{a_0 + a_1\alpha + \cdots + a_m\alpha^m}{b_0 + b_1\alpha + \cdots + b_n\alpha^n}$$

with $a_i, b_j \in F, m, n = 0, 1, \ldots,$ and $b_0 + b_1\alpha + \cdots + b_n\alpha^n \neq 0$. In fact there is a much less complicated description of $F(\alpha)$, as we will see shortly.

3.1.6 Definitions and Comments

If E is an extension of F, the element $\alpha \in E$ is said to be *algebraic* over F is there is a nonconstant polynomial $f \in F[X]$ such that $f(\alpha) = 0$; if α is not algebraic over F, it is said to be *transcendental* over F. If every element of E is algebraic over F, then E is said to be an *algebraic extension* of F.

Suppose that $\alpha \in E$ is algebraic over F, and let I be the set of all polynomials g over F such that $g(\alpha) = 0$. If g_1 and g_2 belong to I, so does $g_1 \pm g_2$, and if $g \in I$ and $c \in F[X]$, then $cg \in I$. Thus I is an ideal of $F[X]$, and since $F[X]$ is a PID, I consists of all multiples of some $m(X) \in F[X]$. Any two such generators must be multiples of each other, so if we require that $m(X)$ be monic, then $m(X)$ is unique. The polynomial $m(X)$ has the following properties:

(1) If $g \in F[X]$, then $g(\alpha) = 0$ if and only if $m(X)$ divides $g(X)$.

(2) $m(X)$ is the monic polynomial of least degree such that $m(\alpha) = 0$.

(3) $m(X)$ is the unique monic irreducible polynomial such that $m(\alpha) = 0$.

Property (1) follows because $g(\alpha) = 0$ iff $g(X) \in I$, and $I = \langle m(X) \rangle$, the ideal generated by $m(X)$. Property (2) follows from (1). To prove (3), note that if $m(X) = h(X)k(X)$ with $\deg h$ and $\deg k$ less than $\deg m$, then either $h(\alpha) = 0$ or $k(\alpha) = 0$, so that by (1), either $h(X)$ or $k(X)$ is a multiple of $m(X)$, which is impossible. Thus $m(X)$ is irreducible, and uniqueness of $m(X)$ follows from (3.1.5).

The polynomial $m(X)$ is called the *minimal polynomial* of α over F, sometimes written as $\min(\alpha, F)$.

3.1.7 Theorem

If $\alpha \in E$ is algebraic over F and the minimal polynomial $m(X)$ of α over F has degree n, then $F(\alpha) = F[\alpha]$, the set of polynomials in α with coefficients in F. In fact, $F[\alpha]$ is the set $F_{n-1}[\alpha]$ of all polynomials *of degree at most* $n - 1$ with coefficients in F, and $1, \alpha, \ldots, \alpha^{n-1}$ form a basis for the vector space $F[\alpha]$ over the field F. Consequently, $[F(\alpha) : F] = n$.

Proof. Let $f(X)$ be any nonzero polynomial over F of degree $n - 1$ or less. Then since $m(X)$ is irreducible and $\deg f < \deg m$, $f(X)$ and $m(X)$ are relatively prime, and there are polynomials $a(X)$ and $b(X)$ over F such that $a(X)f(X) + b(X)m(X) = 1$. But then $a(\alpha)f(\alpha) = 1$, so that any nonzero element of $F_{n-1}[\alpha]$ has a multiplicative inverse. It follows that $F_{n-1}[\alpha]$ is a field. (This may not be obvious, since the product of two polynomials of degree $n-1$ or less can have degree greater than $n-1$, but if $\deg g > n-1$, then divide g by m to get $g(X) = q(X)m(X) + r(X)$ where $\deg r(X) < \deg m(X) = n$. Replace X by α to get $g(\alpha) = r(\alpha) \in F_{n-1}[\alpha]$. Less abstractly, if $m(\alpha) = \alpha^3 + \alpha + 1 = 0$, then $\alpha^3 = -\alpha - 1$, $\alpha^4 = -\alpha^2 - \alpha$, and so on.)

Now any field containing F and α must contain all polynomials in α, in particular all polynomials of degree at most $n-1$. Therefore $F_{n-1}[\alpha] \subseteq F[\alpha] \subseteq F(\alpha)$. But $F(\alpha)$ is the smallest field containing F and α, so $F(\alpha) \subseteq F_{n-1}[\alpha]$, and we conclude that $F(\alpha) = F[\alpha] = F_{n-1}[\alpha]$. Finally, the elements $1, \alpha, \dots, \alpha^{n-1}$ certainly span $F_{n-1}[\alpha]$, and they are linearly independent because if a nontrivial linear combination of these elements were zero, we would have a nonzero polynomial of degree less than that of $m(X)$ with α as a root, contradicting (2) of (3.1.6). ♣

We now prove a basic multiplicativity result for extensions, after a preliminary discussion.

3.1.8 Lemma

Suppose that $F \le K \le E$, the elements $\alpha_i, i \in I$, form a basis for E over K, and the elements $\beta_j, j \in J$, form a basis for K over F. (I and J need not be finite.) Then the products $\alpha_i\beta_j, i \in I, j \in J$, form a basis for E over F.

Proof. If $\gamma \in E$, then γ is a linear combination of the α_i with coefficients $a_i \in K$, and each a_i is a linear combination of the β_j with coefficients $b_{ij} \in F$. It follows that the $\alpha_i\beta_j$ span E over F. Now if $\sum_{i,j} \lambda_{ij}\alpha_i\beta_j = 0$, then $\sum_i \lambda_{ij}\alpha_i = 0$ for all j, and consequently $\lambda_{ij} = 0$ for all i, j, and the $\alpha_i\beta_j$ are linearly independent. ♣

3.1.9 The Degree is Multiplicative

If $F \le K \le E$, then $[E : F] = [E : K][K : F]$. In particular, $[E : F]$ is finite if and only if $[E : K]$ and $[K : F]$ are both finite.

Proof. In (3.1.8), we have $[E : K] = |I|$, $[K : F] = |J|$, and $[E : F] = |I||J|$. ♣

We close this section by proving that every finite extension is algebraic.

3.1.10 Theorem

If E is a finite extension of F, then E is an algebraic extension of F.

Proof. Let $\alpha \in E$, and let $n = [E : F]$. Then $1, \alpha, \alpha^2, \dots, \alpha^n$ are $n + 1$ vectors in an n-dimensional vector space, so they must be linearly dependent. Thus α is a root of a nonzero polynomial with coefficients in F, which means that α is algebraic over F. ♣

Problems For Section 3.1

1. Let E be an extension of F, and let S be a subset of E. If $F(S)$ is the subfield of E generated by S over F, in other words, the smallest subfield of E containing F and S, describe $F(S)$ explicitly, and justify your characterization.
2. If for each $i \in I$, K_i is a subfield of the field E, the *composite* of the K_i (notation $\bigvee_i K_i$) is the smallest subfield of E containing every K_i. As in Problem 1, describe the composite explicitly.
3. Assume that α is algebraic over F, with $[F[\alpha] : F] = n$. If $\beta \in F[\alpha]$, show that $[F[\beta] : F] \leq n$, in fact $[F[\beta] : F]$ divides n.
4. The minimal polynomial of $\sqrt{2}$ over the rationals $\mathbb{Q}$ is $X^2 - 2$, by (3) of (3.1.6). Thus $\mathbb{Q}[\sqrt{2}]$ consists of all numbers of the form $a_0 + a_1\sqrt{2}$, where a_0 and a_1 are rational. By Problem 3, we know that $-1 + \sqrt{2}$ has a minimal polynomial over $\mathbb{Q}$ of degree at most 2. Find this minimal polynomial.
5. If α is algebraic over F and β belongs to $F[\alpha]$, describe a systematic procedure for finding the minimal polynomial of β over F.
6. If E/F and the element $\alpha \in E$ is transcendental over F, show that $F(\alpha)$ is isomorphic to $F(X)$, the field of rational functions with coefficients in F.
7. Theorem 3.1.3 gives one method of adjoining a root of a polynomial, and in fact there is essentially only one way to do this. If E is an extension of F and $\alpha \in E$ is algebraic over F with minimal polynomial $m(X)$, let I be the ideal $\langle m(X) \rangle \subseteq F[X]$. Show that $F(\alpha)$ is isomorphic to $F[X]/I$. [Define $\varphi\colon F[X] \to E$ by $\varphi(f(X)) = f(\alpha)$, and use the first isomorphism theorem for rings.]
8. In the proof of (3.1.3), we showed that if f is irreducible in $F[X]$, then $I = \langle f \rangle$ is a maximal ideal. Show that conversely, if I is a maximal ideal, then f is irreducible.
9. Suppose that $F \leq E \leq L$, with $\alpha \in L$. What is the relation between the minimal polynomial of α over F and the minimal polynomial of α over E?
10. If $\alpha_1, \dots, \alpha_n$ are algebraic over F, we can successively adjoin the α_i to F to obtain the field $F[\alpha_1, \dots, \alpha_n]$ consisting of all polynomials over F in the α_i. Show that

$$[F[\alpha_1, \dots, \alpha_n] : F] \leq \prod_{i=1}^{n} [F(\alpha_i) : F] < \infty$$

3.2 Splitting Fields

If f is a polynomial over the field F, then by (3.1.3) we can find an extension E_1 of F containing a root α_1 of f. If not all roots of f lie in E_1, we can find an extension E_2 of E_1 containing another root α_2 of f. If we continue the process, eventually we reach a complete factorization of f. In this section we examine this idea in detail.

If E is an extension of F and $\alpha_1, \dots, \alpha_k \in E$, we will use the notation $F(\alpha_1, \dots, \alpha_k)$ for the subfield of E generated by F and the α_i. Thus $F(\alpha_1, \dots, \alpha_k)$ is the smallest subfield of E containing all elements of F along with the α_i. ("Smallest" means that $F(\alpha_1, \dots, \alpha_k)$ is the intersection of all such subfields.) Explicitly, $F(\alpha_1, \dots, \alpha_k)$ is the collection of all rational functions in the α_i with nonzero denominators.

3.2.1 Definitions and Comments

If E is an extension of F and $f \in F[X]$, we say that f *splits* over E if f can be written as $\lambda(X-\alpha_1)\cdots(X-\alpha_k)$ for some $\alpha_1,\dots,\alpha_k \in E$ and $\lambda \in F$.

(There is a subtle point that should be mentioned. We would like to refer to the α_i as "the" roots of f, but in doing so we are implicitly assuming that if β is an element of some extension E' of E and $f(\beta)=0$, then β must be one of the α_i. This follows upon substituting β into the equation $f(X)=\lambda(X-\alpha_1)\cdots(X-\alpha_k)=0$.)

If K is an extension of F and $f \in F[X]$, we say that K is a *splitting field* for f over F if f splits over K but not over any proper subfield of K containing F.

Equivalently, K is a splitting field for f over F if f splits over K and K is generated over F by the roots $\alpha_1,\dots,\alpha_k$ of f, in other words, $F(\alpha_1,\dots,\alpha_k)=K$. For if K is a splitting field for f, then since f splits over K we have all $\alpha_j \in K$, so $F(\alpha_1,\dots,\alpha_k)\subseteq K$. But f splits over $F(\alpha_1,\dots,\alpha_k)$, and it follows that $F(\alpha_1,\dots,\alpha_k)$ cannot be a proper subfield; it must coincide with K. Conversely, if f splits over K and $F(\alpha_1,\dots,\alpha_k)=K$, let L be a subfield of K containing F. If f splits over L then all α_i belong to L, so $K=F(\alpha_1,\dots,\alpha_k)\subseteq L\subseteq K$, so $L=K$.

If $f \in F[X]$ and f splits over the extension E of F, then E contains a unique splitting field for f, namely $F(\alpha_1,\dots,\alpha_k)$.

3.2.2 Proposition

If $f \in F[X]$ and $\deg f = n$, then f has a splitting field K over F with $[K:F]\le n!$.

Proof. We may assume that $n \ge 1$. (If f is constant, take $K=F$.) By (3.1.3), F has an extension E_1 containing a root α_1 of f, and the extension $F(\alpha_1)/F$ has degree at most n. (Since $f(\alpha_1)=0$, the minimal polynomial of α_1 divides f; see (3.1.6) and (3.1.7).) We may then write $f(X)=(X-\alpha_1)^{r_1}g(X)$, where α_1 is not a root of g and $\deg g \le n-1$. If g is nonconstant, we can find an extension of $F(\alpha_1)$ containing a root α_2 of g, and the extension $F(\alpha_1,\alpha_2)$ will have degree at most $n-1$ over $F(\alpha_1)$. Continue inductively and use (3.1.9) to reach an extension of degree at most $n!$ containing all the roots of f. ♣

If $f \in F[X]$ and f splits over E, then we may pick any root α of f and adjoin it to F to obtain the extension $F(\alpha)$. Roots of the same irreducible factor of f yield essentially the same extension, as the next result shows.

3.2.3 Theorem

If α and β are roots of the irreducible polynomial $f \in F[X]$ in an extension E of F, then $F(\alpha)$ is isomorphic to $F(\beta)$ via an isomorphism that carries α into β and is the identity on F.

Proof. Without loss of generality we may assume f monic (if not, divide f by its leading coefficient). By (3.1.6), part (3), f is the minimal polynomial of both α and β. By (3.1.7), the elements of $F(\alpha)$ can be expressed uniquely as $a_0+a_1\alpha+\cdots+a_{n-1}\alpha^{n-1}$, where the a_i belong to F and n is the degree of f. The desired isomorphism is given by

$$a_0+a_1\alpha+\cdots+a_{n-1}\alpha^{n-1}\to a_0+a_1\beta+\cdots+a_{n-1}\beta^{n-1}. \quad ♣$$

If f is a polynomial in $F[X]$ and F is isomorphic to the field F' via the isomorphism i, we may regard f as a polynomial over F'. We simply use i to transfer f. Thus if $f = a_0 + a_1X + \cdots a_nX^n$, then $f' = i(f) = i(a_0) + i(a_1)X + \cdots + i(a_n)X^n$. There is only a notational difference between f and f', and we expect that splitting fields for f and f' should also be essentially the same. We prove this after the following definition.

3.2.4 Definition

If E and E' are extensions of F and i is an isomorphism of E and E', we say that i is an F-*isomorphism* if i fixes F, that is, if $i(a) = a$ for every $a \in F$. F-homomorphisms, F-monomorphisms, etc., are defined similarly.

3.2.5 Isomorphism Extension Theorem

Suppose that F and F' are isomorphic, and the isomorphism i carries the polynomial $f \in F[X]$ to $f' \in F'[X]$. If K is a splitting field for f over F and K' is a splitting field for f' over F', then i can be extended to an isomorphism of K and K'. In particular, if $F = F'$ and i is the identity function, we conclude that any two splitting fields of f are F-isomorphic.

Proof. Carry out the construction of a splitting field for f over F as in (3.2.2), and perform exactly the same steps to construct a splitting field for f' over F'. At every stage, there is only a notational difference between the fields obtained. Furthermore, we can do the first construction inside K and the second inside K'. But the comment at the end of (3.2.1) shows that the splitting fields that we have constructed coincide with K and K'. ♣

3.2.6 Example

We will find a splitting field for $f(X) = X^3 - 2$ over the rationals $\mathbb{Q}$.

If α is the positive cube root of 2, then the roots of f are $\alpha, \alpha(-\frac{1}{2} + i\frac{1}{2}\sqrt{3})$ and $\alpha(-\frac{1}{2} - i\frac{1}{2}\sqrt{3})$. The polynomial f is irreducible, either by Eisenstein's criterion or by the observation that if f were factorable, it would have a linear factor, and there is no rational number whose cube is 2. Thus f is the minimal polynomial of α, so $[\mathbb{Q}(\alpha) : \mathbb{Q}] = 3$. Now since α and $i\sqrt{3}$ generate all the roots of f, the splitting field is $K = \mathbb{Q}(\alpha, i\sqrt{3})$. (We regard all fields in this example as subfields of the complex numbers $\mathbb{C}$.) Since $i\sqrt{3} \notin \mathbb{Q}(\alpha)$ (because $\mathbb{Q}(\alpha)$ is a subfield of the reals), $[\mathbb{Q}(\alpha, i\sqrt{3}) : \mathbb{Q}(\alpha)]$ is at least 2. But $i\sqrt{3}$ is a root of $X^2 + 3 \in \mathbb{Q}(\alpha)[X]$, so the degree of $\mathbb{Q}(\alpha, i\sqrt{3})$ over $\mathbb{Q}(\alpha)$ is a most 2, and therefore is exactly 2. Thus

$$[K : \mathbb{Q}] = [Q(\alpha, i\sqrt{3}) : \mathbb{Q}] = [Q(\alpha, i\sqrt{3}) : \mathbb{Q}(\alpha)][\mathbb{Q}(\alpha) : \mathbb{Q}] = 2 \times 3 = 6.$$

Problems For Section 3.2

1. Find a splitting field for $f(X) = X^2 - 4X + 4$ over $\mathbb{Q}$.
2. Find a splitting field K for $f(X) = X^2 - 2X + 4$ over $\mathbb{Q}$, and determine the degree of K over $\mathbb{Q}$.

3. Find a splitting field K for $f(X) = X^4 - 2$ over $\mathbb{Q}$, and determine $[K : \mathbb{Q}]$.
4. Let $\mathcal{C}$ be a *family* of polynomials over F, and let K be an extension of F. Show that the following two conditions are equivalent:

   (a) Each $f \in \mathcal{C}$ splits over K, but if $F \leq K' < K$, then it is not true that each $f \in \mathcal{C}$ splits over K'.

   (b) Each $f \in \mathcal{C}$ splits over K, and K is generated over F by the roots of all the polynomials in $\mathcal{C}$.

   If one, and hence both, of these conditions are satisfied, we say that K is a *splitting field for* $\mathcal{C}$ over F.
5. Suppose that K is a splitting field for the finite set of polynomials $\{f_1, \dots, f_r\}$ over F. Express K as a splitting field for a single polynomial f over F.
6. If m and n are distinct square-free positive integers greater than 1, show that the splitting field $\mathbb{Q}(\sqrt{m}, \sqrt{n})$ of $(X^2 - m)(X^2 - n)$ has degree 4 over $\mathbb{Q}$.

3.3 Algebraic Closures

If f is a polynomial of degree n over the rationals or the reals, or more generally over the complex numbers, then f need not have any rational roots, or even real roots, but we know that f always has n complex roots, counting multiplicity. This favorable situation can be duplicated for any field F, that is, we can construct an algebraic extension C of F with the property that any polynomial in $C[X]$ splits over C. There are many ways to express this idea.

3.3.1 Proposition

If C is a field, the following conditions are equivalent:

(1) Every nonconstant polynomial $f \in C[X]$ has at least one root in C.

(2) Every nonconstant polynomial $f \in C[X]$ splits over C.

(3) Every irreducible polynomial $f \in C[X]$ is linear.

(4) C has no proper algebraic extensions.

If any (and hence all) of these conditions are satisfied, we say that C is *algebraically closed*.

Proof. (1) implies (2): By (1) we may write $f = (X - \alpha_1)g$. Proceed inductively to show that any nonconstant polynomial is a product of linear factors.

(2) implies (3): If f is an irreducible polynomial in $C[X]$, then by (2.9.1), f is nonconstant. By (2), f is a product of linear factors. But f is irreducible, so there can be only one such factor.

(3) implies (4): Let E be an algebraic extension of C. If $\alpha \in E$, let f be the minimal polynomial of α over C. Then f is irreducible and by (3), f is of the form $X - \alpha$. But then $\alpha \in C$, so $E = C$.

(4) implies (1): Let f be a nonconstant polynomial in $C[X]$, and adjoin a root α of f to obtain $C(\alpha)$, as in (3.1.3). But then $C(\alpha)$ is an algebraic extension of C, so by (4), $\alpha \in C$. ♣

It will be useful to embed an arbitrary field F in an algebraically closed field.

3.3.2 Definitions and Comments

An extension C of F is an *algebraic closure* of F if C is algebraic over F and C is algebraically closed.

Note that C is minimal among algebraically closed extensions of F. For if $F \le K \le C$ and $\alpha \in C, \alpha \notin K$, then since α is algebraic over F it is algebraic over K. But since $\alpha \notin K$, the minimal polynomial of α over K is a nonlinear irreducible polynomial in $K[X]$. By (3) of (3.3.1), K cannot be algebraically closed.

If C is an algebraic extension of F, then in order for C to be an algebraic closure of F it is sufficient that every polynomial in $F[X]$ (rather than $C[X]$) splits over C. To prove this, we will need the following result.

3.3.3 Proposition

If E is generated over F by finitely many elements $\alpha_1, \dots, \alpha_n$ algebraic over F (so that $E = F(\alpha_1, \dots, \alpha_n)$), then E is a finite extension of F.

Proof. Set $E_0 = F$ and $E_k = F(\alpha_1, \dots, \alpha_k)$, $1 \le k \le n$ (so $E_n = E$). Then $E_k = E_{k-1}(\alpha_k)$, where α_k is algebraic over F and hence over E_{k-1}. But by (3.1.7), $[E_k : E_{k-1}]$ is the degree of the minimal polynomial of α_k over E_{k-1}, which is finite. By (3.1.9), $[E : F] = \prod_{k=1}^{n} [E_k : E_{k-1}] < \infty$. ♣

3.3.4 Corollary

If E is an extension of F and A is the set of all elements in E that are algebraic over F (the *algebraic closure of F in E*), then A is a subfield of E.

Proof. If $\alpha, \beta \in A$, then the sum, difference, product and quotient (if $\beta \neq 0$) of α and β belong to $F(\alpha, \beta)$, which is a finite extension of F by (3.3.3), and therefore an algebraic extension of F by (3.1.10). But then $\alpha + \beta, \alpha - \beta, \alpha\beta$ and α/β belong to A, proving that A is a field. ♣

3.3.5 Corollary (Transitivity of Algebraic Extensions)

If E is algebraic over K (in other words, every element of E is algebraic over K), and K is algebraic over F, then E is algebraic over F.

Proof. Let $\alpha \in E$, and let $m(X) = b_0 + b_1 X + \cdots + b_{n-1} X^{n-1} + X^n$ be the minimal polynomial of α over K. The b_i belong to K and are therefore algebraic over F. If $L = F(b_0, b_1, \dots, b_{n-1})$, then by (3.3.3), L is a finite extension of F. Since the coefficients of $m(X)$ belong to L, α is algebraic over L, so by (3.1.7), $L(\alpha)$ is a finite extension of L. By (3.1.9), $L(\alpha)$ is a finite extension of F. By (3.1.10), α is algebraic over F. ♣

Now we can add another condition to (3.3.1).

3.3.6 Proposition

Let C be an algebraic extension of F. Then C is an algebraic closure of F if and only if every nonconstant polynomial in $F[X]$ splits over C.

Proof. The "only if" part follows from (2) of (3.3.1), since $F \subseteq C$. Thus assume that every nonconstant polynomial in $F[X]$ splits over C. If f is a nonconstant polynomial in $C[X]$, we will show that f has at least one root in C, and it will follow from (1) of (3.3.1) that C is algebraically closed. Adjoin a root α of f to obtain the extension $C(\alpha)$. Then $C(\alpha)$ is algebraic over C by (3.1.7), and C is algebraic over F by hypothesis. By (3.3.5), $C(\alpha)$ is algebraic over F, so α is algebraic over F. But then α is a root of some polynomial $g \in F[X]$, and by hypothesis, g splits over C. By definition of "splits" (see (3.2.1)), all roots of g lie in C, in particular $\alpha \in C$. Thus f has at least one root in C. ♣

To avoid a lengthy excursion into formal set theory, we argue intuitively to establish the following three results. (For complete proofs, see the appendix to Chapter 3.)

3.3.7 Theorem

Every field F has an algebraic closure.

Informal argument. Well-order $F[X]$ and use transfinite induction, beginning with the field $F_0 = F$. At stage f we adjoin all roots of the polynomial f by constructing a splitting field for f over the field $F_{<f}$ that has been generated so far by the recursive procedure. When we reach the end of the process, we will have a field C such that every polynomial f in $F[X]$ splits over C. By (3.3.6), C is an algebraic closure of F. ♣

3.3.8 Theorem

Any two algebraic closures C and C' of F are F-isomorphic.

Informal argument. Carry out the recursive procedure described in (3.3.7) in both C and C'. At each stage we may use the fact that any two splitting fields of the same polynomial are F-isomorphic; see (3.2.5). When we finish, we have F-isomorphic algebraic closures of F, say $D \subseteq C$ and $D' \subseteq C'$. But an algebraic closure is a minimal algebraically closed extension by (3.3.2), and therefore $D = C$ and $D' = C'$. ♣

3.3.9 Theorem

If E is an algebraic extension of F, C is an algebraic closure of F, and i is an embedding (that is, a monomorphism) of F into C, then i can be extended to an embedding of E into C.

Informal argument. Each $\alpha \in E$ is a root of some polynomial in $F[X]$, so if we allow α to range over all of E, we get a collection S of polynomials in $F[X]$. Within C, carry out the recursive procedure of (3.3.7) on the polynomials in S. The resulting field lies inside C and contains an F-isomorphic copy of E. ♣

Problems For Section 3.3

1. Show that the converse of (3.3.3) holds, that is, if E is a finite extension of F, then E is generated over F by finitely many elements that are algebraic over F.
2. An *algebraic number* is a complex number that is algebraic over the rational field $\mathbb{Q}$. A *transcendental number* is a complex number that is not algebraic over $\mathbb{Q}$. Show that there only countably many algebraic numbers, and consequently there are uncountably many transcendental numbers.
3. Give an example of an extension C/F such that C is algebraically closed but C is not an algebraic extension of F.
4. Give an example of an extension E/F such that E is an algebraic but not a finite extension of F.
5. In the proof of (3.3.7), why is C algebraic over F?
6. Show that the set A of algebraic numbers is an algebraic closure of $\mathbb{Q}$.
7. If E is an algebraic extension of the infinite field F, show that $|E| = |F|$.
8. Show that any set S of nonconstant polynomials in $F[X]$ has a splitting field over F.
9. Show that an algebraically closed field must be infinite.

3.4 Separability

If f is a polynomial in $F[X]$, we can construct a splitting field K for f over F, and all roots of f must lie in K. In this section we investigate the multiplicity of the roots.

3.4.1 Definitions and Comments

An irreducible polynomial $f \in F[X]$ is *separable* if f has no repeated roots in a splitting field; otherwise f is *inseparable*. If f is an arbitrary polynomial, not necessarily irreducible, then we call f separable if each of its irreducible factors is separable.

Thus if $f(X) = (X-1)^2(X-3)$ over $\mathbb{Q}$, then f is separable, because the irreducible factors $(X-1)$ and $(X-3)$ do not have repeated roots. We will see shortly that over a field of characteristic 0 (for example, the rationals), every polynomial is separable. Here is a method for testing for multiple roots.

3.4.2 Proposition

If

$$f(X) = a_0 + a_1X + \cdots + a_nX^n \in F[X],$$

let f' be the *derivative* of f, defined by

$$f'(X) = a_1 + 2a_2X + \cdots + na_nX^{n-1}.$$

[Note that the derivative is a purely formal expression; we completely ignore questions about existence of limits. One can check by brute force that the usual rules for differentiating a sum and product apply].

If g is the greatest common divisor of f and f', then f has a repeated root in a splitting field if and only if the degree of g is at least 1.

Proof. If f has a repeated root, we can write $f(X) = (X-\alpha)^r h(X)$ where $r \geq 2$. Applying the product rule for derivatives, we see that $(X - \alpha)$ is a factor of both f and f', and consequently $\deg g \geq 1$. Conversely, if $\deg g \geq 1$, let α be a root of g in some splitting field. Then $(X - \alpha)$ is a factor of both f and f'. We will show that α is a repeated root of f. If not, we may write $f(X) = (X - \alpha)h(X)$ where $h(\alpha) \neq 0$. Differentiate to obtain $f'(X) = (X - \alpha)h'(X) + h(X)$, hence $f'(\alpha) = h(\alpha) \neq 0$. This contradicts the fact that $(X - \alpha)$ is a factor of f'. ♣

3.4.3 Corollary

(1) Over a field of characteristic zero, every polynomial is separable.

(2) Over a field F of prime characteristic p, the irreducible polynomial f is inseparable if and only if f' is the zero polynomial. Equivalently, f is a polynomial in X^p; we abbreviate this as $f \in F[X^p]$.

Proof. (1) Without loss of generality, we can assume that we are testing an irreducible polynomial f. The derivative of X^n is nX^{n-1}, and in a field of characteristic 0, n cannot be 0. Thus f' is a nonzero polynomial whose degree is less than that of f. Since f is irreducible, the gcd of f and f' is either 1 or f, and the latter is excluded because f cannot possibly divide f'. By (3.4.2), f is separable.

(2) If $f' \neq 0$, the argument of (1) shows that f is separable. If $f' = 0$, then $\gcd(f, f') = f$, so by (3.4.2), f is inseparable. In characteristic p, an integer n is zero if and only if n is a multiple of p, and it follows that $f' = 0$ iff $f \in F[X^p]$. ♣

By (3.4.3), part (1), every polynomial over the rationals (or the reals or the complex numbers) is separable. This pleasant property is shared by finite fields as well. First note that a finite field F cannot have characteristic 0, since a field of characteristic 0 must contain a copy of the integers (and the rationals as well), and we cannot squeeze infinitely many integers into a finite set. Now recall the binomial expansion modulo p, which is simply $(a + b)^p = a^p + b^p$, since p divides $\binom{p}{k}$ for $1 \leq k \leq p - 1$. [By induction, $(a + b)^{p^n} = a^{p^n} + b^{p^n}$ for every positive integer n.] Here is the key step in the analysis.

3.4.4 The Frobenius Automorphism

Let F be a finite field of characteristic p, and define $f\colon F \to F$ by $f(\alpha) = \alpha^p$. Then f is an automorphism. In particular, if $\alpha \in F$ then $\alpha = \beta^p$ for some $\beta \in F$.

Proof. We have $f(1) = 1$ and

$$f(\alpha + \beta) = (\alpha + \beta)^p = \alpha^p + \beta^p = f(\alpha) + f(\beta),$$
$$f(\alpha\beta) = (\alpha\beta)^p = \alpha^p\beta^p = f(\alpha)f(\beta)$$

so f is a monomorphism. But an injective function from a finite set to itself is automatically surjective, and the result follows. ♣

3.4.5 Proposition

Over a finite field, every polynomial is separable.

Proof. Suppose that f is an irreducible polynomial over the finite field F with repeated roots in a splitting field. By (3.4.3), part (2), $f(X)$ has the form $a_0 + a_1X^p + \cdots + a_nX^{np}$ with the $a_i \in F$. By (3.4.4), for each i there is an element $b_i \in F$ such that $b_i^p = a_i$. But then

$$(b_0 + b_1X + \cdots + b_nX^n)^p = b_0^p + b_1^pX^p + \cdots + b_n^pX^{np} = f(X)$$

which contradicts the irreducibility of f. ♣

Separability of an element can be defined in terms of its minimal polynomial.

3.4.6 Definitions and Comments

If E is an extension of F and $\alpha \in E$, then α is *separable over* F if α is algebraic over F and $\min(\alpha, F)$ is a separable polynomial. If every element of E is separable over F, we say that E *is a separable extension of* F or *the extension* E/F *is separable* or E *is separable over* F. By (3.4.3) and (3.4.5), every algebraic extension of a field of characteristic zero or a finite field is separable.

3.4.7 Lemma

If $F \le K \le E$ and E is separable over F, then K is separable over F and E is separable over K.

Proof. Since K is a subfield of E, K/F is separable. If $\alpha \in E$, then since α is a root of $\min(\alpha, F)$, it follows from (1) of (3.1.6) that $\min(\alpha, K)$ divides $\min(\alpha, F)$. By hypothesis, $\min(\alpha, F)$ has no repeated roots in a splitting field, so neither does $\min(\alpha, K)$. Thus E/K is separable. ♣

The converse of (3.4.7) is also true: If K/F and E/K are separable, then E/F is separable. Thus we have *transitivity of separable extensions.* We will prove this (for finite extensions) in the exercises.

In view of (3.4.6), we can produce many examples of separable extensions. Inseparable extensions are less common, but here is one way to construct them.

3.4.8 Example

Let $F = \mathbb{F}_p(t)$ be the set of rational functions (in the indeterminate t) with coefficients in the field with p elements (the integers mod p). Thus an element of F looks like

$$\frac{a_0 + a_1 t + \cdots + a_m t^m}{b_0 + b_1 t + \cdots + b_n t^n}.$$

with the a_i and b_j in $\mathbb{F}_p$. Adjoin $\sqrt[p]{t}$, that is, a root of $X^p - t$, to create the extension E. Note that $X^p - t$ is irreducible by Eisenstein, because t is irreducible in $\mathbb{F}_p[t]$. (The product of two nonconstant polynomials in t cannot possibly be t.) The extension E/F is inseparable, since

$$X^p - t = X^p - (\sqrt[p]{t})^p = (X - \sqrt[p]{t})^p,$$

which has multiple roots.

Problems For Section 3.4

1. Give an example of a separable polynomial f whose derivative is zero. (In view of (3.4.3), f cannot be irreducible.)
2. Let $\alpha \in E$, where E is an algebraic extension of a field F of prime characteristic p. Let $m(X)$ be the minimal polynomial of α over the field $F(\alpha^p)$. Show that $m(X)$ splits over E, and in fact α is the only root, so that $m(X)$ is a power of $(X - \alpha)$.
3. Continuing Problem 2, if α is separable over the field $F(\alpha^p)$, show that $\alpha \in F(\alpha^p)$.
4. A field F is said to be *perfect* if every polynomial over F is separable. Equivalently, every algebraic extension of F is separable. Thus fields of characteristic zero and finite fields are perfect. Show that if F has prime characteristic p, then F is perfect if and only if every element of F is the p^{th} power of some element of F. For short we write $F = F^p$.

   In Problems 5-8, we turn to transitivity of separable extensions.
5. Let E be a finite extension of a field F of prime characteristic p, and let $K = F(E^p)$ be the subfield of E obtained from F by adjoining the p^{th} powers of all elements of E. Show that $F(E^p)$ consists of all finite linear combinations of elements in E^p with coefficients in F.
6. Let E be a finite extension of the field F of prime characteristic p, and assume that $E = F(E^p)$. If the elements $y_1, \dots, y_r \in E$ are linearly independent over F, show that $y_1^p, \dots, y_r^p$ are linearly independent over F.
7. Let E be a finite extension of the field F of prime characteristic p. Show that the extension is separable if and only if $E = F(E^p)$.
8. If $F \le K \le E$ with $[E : F] < \infty$, with E separable over K and K separable over F, show that E is separable over F.
9. Let f be an irreducible polynomial in $F[X]$, where F has characteristic $p > 0$. Express $f(X)$ as $g(X^{p^m})$, where the nonnegative integer m is a large as possible. (This makes sense because $X^{p^0} = X$, so $m = 0$ always works, and f has finite degree, so m is bounded above.) Show that g is irreducible and separable.

10. Continuing Problem 9, if f has only one distinct root α, show that $\alpha^{p^m} \in F$.
11. If E/F, where char $F = p > 0$, and the element $\alpha \in E$ is algebraic over F, show that the minimal polynomial of α over F has only one distinct root if and only if $\alpha^{p^n} \in F$ for some nonnegative integer n. (In this case we say that α is *purely inseparable* over F.)

3.5 Normal Extensions

Let E/F be a field extension. In preparation for Galois theory, we are going to look at monomorphisms defined on E, especially those which fix F. First we examine what an F-monomorphism does to the roots of a polynomial in $F[X]$.

3.5.1 Lemma

Let $\sigma\colon E \to E$ be an F-monomorphism, and assume that the polynomial $f \in F[X]$ splits over E. If α is a root of f in E, then so is $\sigma(\alpha)$. Thus σ permutes the roots of f.

Proof. If $b_0 + b_1\alpha + \cdots + b_n\alpha^n = 0$, with the $b_i \in F$, apply σ and note that since σ is an F-monomorphism, $\sigma(b_i) = b_i$ and $\sigma(\alpha^i) = (\sigma(\alpha))^i$. Thus

$$b_0 + b_1\sigma(\alpha) + \cdots + b_n(\sigma(\alpha))^n = 0. \quad ♣$$

Now let C be an algebraic closure of E. It is convenient to have C available because it will contain all the roots of a polynomial $f \in E[X]$, even if f does not split over E. We are going to count the number of embeddings of E in C that fix F, that is, the number of F-monomorphisms of E into C. Here is the key result.

3.5.2 Theorem

Let E/F be a finite separable extension of degree n, and let σ be an embedding of F in C. Then σ extends to exactly n embeddings of E in C; in other words, there are exactly n embeddings τ of E in C such that the restriction $\tau|_F$ of τ to F coincides with σ. In particular, taking σ to be the identity function on F, there are exactly n F-monomorphisms of E into C.

Proof. An induction argument works well. If $n = 1$ then $E = F$ and there is nothing to prove, so assume $n > 1$ and choose an element α that belongs to E but not to F. If f is the minimal polynomial of α over F, let $g = \sigma(f)$. (This is a useful shorthand notation, indicating that if a_i is one of the coefficients of f, the corresponding coefficient of g is $\sigma(a_i)$.) Any factorization of g can be translated via the inverse of σ to a factorization of f, so g is separable and irreducible over the field $\sigma(F)$. If β is any root of g, then there is a unique isomorphism of $F(\alpha)$ and $(\sigma(F))(\beta)$ that carries α into β and coincides with σ on F. Explicitly,

$$b_0 + b_1\alpha + \cdots + b_r\alpha^r \to \sigma(b_0) + \sigma(b_1)\beta + \cdots + \sigma(b_r)\beta^r.$$

Now if $\deg g = r$, then $[F(\alpha) : F] = \deg f = \deg g = r$ as well, so by (3.1.9), $[E : F(\alpha)] = n/r < n$. By separability, g has exactly r distinct roots in C, so there are exactly r possible choices of β. In each case, by the induction hypothesis,the resulting embedding of $F(\alpha)$ in C has exactly n/r extensions to embeddings of E in C. This produces n distinct embeddings of E in C extending σ. But if τ is any embedding of F in C that extends σ, then just as in (3.5.1), τ must take α to a root of g, i.e., to one of the β's. If there were more than n possible τ's, there would have to be more than n/r possible extensions of at least one of the embeddings of $F(\alpha)$ in C. This would contradict the induction hypothesis. ♣

3.5.3 Example

Adjoin the positive cube root of 2 to the rationals to get $E = \mathbb{Q}(\sqrt[3]{2})$. The roots of the irreducible polynomial $f(X) = X^3 - 2$ are $\sqrt[3]{2}$, $\omega\sqrt[3]{2}$ and $\omega^2\sqrt[3]{2}$, where $\omega = e^{i2\pi/3} = -\frac{1}{2} + \frac{1}{2}i\sqrt{3}$ and $\omega^2 = e^{i4\pi/3} = -\frac{1}{2} - \frac{1}{2}i\sqrt{3}$.

Notice that the polynomial f has a root in E but does not split in E (because the other two roots are complex and E consists entirely of real numbers). We give a special name to extensions that do not have this annoying drawback.

3.5.4 Definition

The algebraic extension E/F is *normal* (we also say that E is *normal over* F) if every irreducible polynomial over F that has at least one root in E splits over E. In other words, if $\alpha \in E$, then all *conjugates* of α over F (i.e., all roots of the minimal polynomial of α over F) belong to E.

Here is an equivalent condition.

3.5.5 Theorem

The finite extension E/F is normal if and only if every F-monomorphism of E into an algebraic closure C is actually an F-automorphism of E. (The hypothesis that E/F is finite rather than simply algebraic can be removed, but we will not need the more general result.)

Proof. If E/F is normal, then as in (3.5.1), an F-monomorphism τ of E into C must map each element of E to one of its conjugates. Thus by hypothesis, $\tau(E) \subseteq E$. But $\tau(E)$ is an isomorphic copy of E , so it must have the same degree as E over F. Since the degree is assumed finite, we have $\tau(E) = E$. (All we are saying here is that an m-dimensional subspace of an m-dimensional vector space is the entire space.) Conversely, let $\alpha \in E$, and let β be any conjugate of α over F. As in the proof of (3.5.2), there is an F-monomorphism of E into C that carries α to β. If all such embeddings are F-automorphisms of E, we must have $\beta \in E$, and we conclude that E is normal over F. ♣

3.5.6 Remarks

In (3.5.2) and (3.5.5), the algebraic closure can be replaced by any fixed normal extension of F containing E; the proof is the same. Also, the implication $\tau(E) \subseteq E \Rightarrow \tau(E) = E$

holds for any F-monomorphism τ and any finite extension E/F; normality is not involved.

The next result yields many explicit examples of normal extensions.

3.5.7 Theorem

The finite extension E/F is normal if and only if E is a splitting field for some polynomial $f \in F[X]$.

Proof. Assume that E is normal over F. Let $\alpha_1, \dots, \alpha_n$ be a basis for E over F, and let f_i be the minimal polynomial of α_i over $F, i = 1, \dots, n$. Since f_i has a root α_i in E, f_i splits over E, hence so does $f = f_1 \cdots f_n$. If f splits over a field K with $F \subseteq K \subseteq E$, then each α_i belongs to K, and therefore K must coincide with E. Thus E is a splitting field for f over F. Conversely, let E be a splitting field for f over F, where the roots of f are α_i, $i = 1, \dots, n$. Let τ be an F-monomorphism of E into an algebraic closure. As in (3.5.1), τ takes each α_i into another root of f, and therefore τ takes a polynomial in the α_i to another polynomial in the α_i. But $F(\alpha_1, \dots, \alpha_n) = E$, so $\tau(E) \subseteq E$. By (3.5.6), τ is an automorphism of E, so by (3.5.5), E/F is normal. ♣

3.5.8 Corollary

Let $F \le K \le E$, where E is a finite extension of F. If E/F is normal, so is E/K.

Proof. By (3.5.7), E is a splitting field for some polynomial $f \in F[X]$, so that E is generated over F by the roots of f. But then $f \in K[X]$ and E is generated over K by the roots of f. Again by (3.5.7), E/K is normal. ♣

3.5.9 Definitions and Comments

If E/F is normal and separable, it is said to be a *Galois extension*; we also say that E is *Galois over* F. It follows from (3.5.2) and (3.5.5) that if E/F is a finite Galois extension, then there are exactly $[E : F]$ F-automorphisms of E. If E/F is finite and separable but not normal, then at least one F-embedding of E into an algebraic closure must fail to be an automorphism of E. Thus in this case, the number of F-automorphisms of E is less than the degree of the extension.

If E/F is an arbitrary extension, the *Galois group* of the extension, denoted by $\mathrm{Gal}(E/F)$, is the set of F-automorphisms of E. (The set is a group under composition of functions.)

3.5.10 Example

Let $E = \mathbb{Q}(\sqrt[3]{2})$, as in (3.5.3). The Galois group of the extension consists of the identity automorphism alone. For any $\mathbb{Q}$-monomorphism σ of E must take $\sqrt[3]{2}$ into a root of $X^3 - 2$. Since the other two roots are complex and do not belong to E, $\sqrt[3]{2}$ must map to itself. But σ is completely determined by its action on $\sqrt[3]{2}$, and the result follows.

If E/F is not normal, we can always enlarge E to produce a normal extension of F. If C is an algebraic closure of E, then C contains all the roots of every polynomial in $F[X]$, so C/F is normal. Let us try to look for a smaller normal extension.

3.5.11 The Normal Closure

Let E be a finite extension of F, say $E = F(\alpha_1, \dots, \alpha_n)$. If $N \supseteq E$ is any normal extension of F, then N must contain the α_i along with all conjugates of the α_i, that is, all roots of $\min(\alpha_i, F), i = 1, \dots, n$. Thus if f is the product of these minimal polynomials, then N must contain the splitting field K for f over F. But K/F is normal by (3.5.7), so K must be the smallest normal extension of F that contains E. It is called the *normal closure* of E over F.

We close the section with an important result on the structure of finite separable extensions.

3.5.12 Theorem of the Primitive Element

If E/F is a finite separable extension, then $E = F(\alpha)$ for some $\alpha \in E$. We say that α is a *primitive element* of E over F.

Proof. We will argue by induction on $n = [E : F]$. If $n = 1$ then $E = F$ and we can take α to be any member of F. If $n > 1$, choose $\alpha \in E \setminus F$. By the induction hypothesis, there is a primitive element β for E over $F(\alpha)$, so that $E = F(\alpha, \beta)$. We are going to show that if $c \in F$ is properly chosen, then $E = F(\alpha + c\beta)$. Now by (3.5.2), there are exactly n F-monomorphisms of E into an algebraic closure C, and each of these maps restricts to an F-monomorphism of $F(\alpha + c\beta)$ into C. If $F(\alpha + c\beta) \neq E$, then $[F(\alpha + c\beta) : F] < n$, and it follows from (3.5.2) that at least two embeddings of E, say σ and τ, must coincide when restricted. Therefore

$$\sigma(\alpha) + c\sigma(\beta) = \tau(\alpha) + c\tau(\beta),$$

hence

$$c = \frac{\sigma(\alpha) - \tau(\alpha)}{\tau(\beta) - \sigma(\beta)}. \tag{1}$$

(If $\tau(\beta) = \sigma(\beta)$ then by the previous equation, $\tau(\alpha) = \sigma(\alpha)$. But an F-embedding of E is determined by what it does to α and β, hence $\sigma = \tau$, a contradiction.) Now an F-monomorphism must map α to one of its conjugates over F, and similarly for β. Thus there are only finitely many possible values for the ratio in (1). If we select c to be different from each of these values, we reach a contradiction of our assumption that $F(\alpha + c\beta) \neq E$. The proof is complete if F is an infinite field. We must leave a gap here, to be filled later (see (6.4.4)). If F is finite, then so is E (since E is a finite-dimensional vector space over F). We will show that the multiplicative group of nonzero elements of a finite field E is cyclic, so if α is a generator of this group, then $E = F(\alpha)$. ♣

Problems For Section 3.5

1. Give an example of fields $F \leq K \leq E$ such that E/F is normal but K/F is not.
2. Let $E = \mathbb{Q}(\sqrt{a})$, where a is an integer that is not a perfect square. Show that $E/\mathbb{Q}$ is normal.

3. Give an example of fields $F \le K \le E$ such that E/K and K/F are normal, but E/F is not. Thus transitivity fails for normal extensions.
4. Suppose that in (3.5.2), the hypothesis of separability is dropped. State and prove an appropriate conclusion.
5. Show that $E = \mathbb{Q}(\sqrt{2}, \sqrt{3})$ is a Galois extension of $\mathbb{Q}$.
6. In Problem 5, find the Galois group of $E/\mathbb{Q}$.
7. Let E be a finite extension of F, and let K be a normal closure (= minimal normal extension) of E over F, as in (3.5.11). Is K unique?
8. If E_1 and E_2 are normal extensions of F, show that $E_1 \cap E_2$ is normal over F.

Appendix To Chapter 3

In this appendix, we give a precise development of the results on algebraic closure treated informally in the text.

A3.1 Lemma

Let E be an algebraic extension of F, and let $\sigma\colon E \to E$ be an F-monomorphism. Then σ is surjective, hence σ is an automorphism of E.

Proof. Let $\alpha \in E$, and let $f(X)$ be the minimal polynomial of α over F. We consider the subfield L of E generated over F by the roots of f that lie in E. Then L is an extension of F that is finitely generated by algebraic elements, so by (3.3.3), L/F is finite. As in (3.5.1), σ takes a root of f to a root of f, so $\sigma(L) \subseteq L$. But $[L : F] = [\sigma(L) : F] < \infty$ (σ maps a basis to a basis), and consequently $\sigma(L) = L$. But $\alpha \in L$, so $\alpha \in \sigma(L)$. ♣

The following result, due to Artin, is crucial.

A3.2 Theorem

If F is any field, there is an algebraically closed field E containing F.

Proof. For each nonconstant polynomial f in $F[X]$, we create a variable $X(f)$. If T is the collection of all such variables, we can form the ring $F[T]$ of all polynomials in all possible finite sets of variables in T, with coefficients in F. Let I be the ideal of $F[T]$ generated by the polynomials $f(X(f))$, $f \in F[X]$. We claim that I is a proper ideal. If not, then $1 \in I$, so there are finitely many polynomials $f_1, \dots, f_n$ in $F[X]$ and polynomials $h_1, \dots, h_n$ in $F[T]$ such that $\sum_{i=1}^n h_i f_i(X(f_i)) = 1$. Now only finitely many variables $X_i = X(f_i), i = 1, \dots, m$, can possibly appear in the h_i, so we have an equation of the form

$$\sum_{i=1}^{n} h_i(X_1, \dots, X_m) f_i(X_i) = 1 \tag{1}$$

where $m \geq n$. Let L be the extension of F formed by successively adjoining the roots of $f_1, \dots, f_n$. Then each f_i has a root $\alpha_i \in L$. If we set $\alpha_i = 0$ for $n \leq i < m$ and then set $X_i = \alpha_i$ for each i in (1), we get $0 = 1$, a contradiction.

Thus the ideal I is proper, and is therefore contained in a maximal ideal $\mathcal{M}$. Let E_1 be the field $F[T]/\mathcal{M}$. Then E_1 contains an isomorphic copy of F, via the map taking $a \in F$ to $a + \mathcal{M} \in E_1$. (Note that if $a \in \mathcal{M}, a \neq 0$, then $1 = a^{-1}a \in \mathcal{M}$, a contradiction.) Consequently, we can assume that $F \leq E_1$. If f is any nonconstant polynomial in $F[X]$, then $X(f) + \mathcal{M} \in E_1$ and $f(X(f) + \mathcal{M}) = f(X(f)) + \mathcal{M} = 0$ because $f(X(f)) \in I \subseteq \mathcal{M}$.

Iterating the above procedure, we construct a chain of fields $F \leq E_1 \leq E_2 \leq \cdots$ such that every polynomial of degree at least 1 in $E_n[X]$ has a root in E_{n+1}. The union E of all the E_n is a field, and every nonconstant polynomial f in $E[X]$ has all its coefficients in some E_n. Therefore f has a root in $E_{n+1} \subseteq E$. ♣

A3.3 Theorem

Every field F has an algebraic closure.

Proof. By (3.5.12), F has an algebraically closed extension L. If E is the algebraic closure of F in L (see 3.3.4), then E/F is algebraic. Let f be a nonconstant polynomial in $E[X]$. Then f has a root α in L (because L is algebraically closed). We now have α algebraic over E (because $f \in E[X]$), and E algebraic over F. As in (3.3.5), α is algebraic over F, hence $\alpha \in E$. By (3.3.1), E is algebraically closed. ♣

A3.4 Problem

Suppose that σ is a monomorphism of F into the algebraically closed field L. Let E be an algebraic extension of F, and α an element of E with minimal polynomial f over F. We wish to extend σ to a monomorphism from $F(\alpha)$ to L. In how many ways can this be done?

Let σf be the polynomial in $(\sigma F)[X]$ obtained from f by applying σ to the coefficients of f. Any extension of f is determined by what it does to α, and as in (3.5.1), the image of α is a root of σf. Now the number of distinct roots of f in an algebraic closure of F, call it t, is the same as the number of distinct roots of σf in L; this follows from the isomorphism extension theorem (3.2.5). Thus the number of extensions is at most t. But if β is any root of σf, we can construct an extension of σ by mapping the element $h(\alpha) \in F(\alpha)$ to $(\sigma h)(\beta)$; in particular, α is mapped to β. To show that the definition makes sense, suppose that $h_1(\alpha) = h_2(\alpha)$. Then $(h_1 - h_2)(\alpha) = 0$, so f divides $h_1 - h_2$ in $F[X]$. Consequently, σf divides $\sigma h_1 - \sigma h_2$ in $(\sigma F)[X]$, so $(\sigma h_1)(\beta) = (\sigma h_2)(\beta)$.

We conclude that the number of extensions of σ is the number of distinct roots of f in an algebraic closure of F.

Rather than extend σ one element at a time, we now attempt an extension to all of E.

A3.5 Theorem

Let $\sigma \colon F \to L$ be a monomorphism, with L algebraically closed. If E is an algebraic extension of F, then σ has an extension to a monomorphism $\tau \colon E \to L$.

Proof. Let $\mathcal{G}$ be the collection of all pairs (K, μ) where K is an intermediate field between F and E and μ is an extension of σ to a monomorphism from K to L. We partially order $\mathcal{G}$ by $(K_1, \mu) \leq (K_2, \rho)$ iff $K_1 \subseteq K_2$ and ρ restricted to K_1 coincides with μ. Since $(F, \sigma) \in \mathcal{G}$, we have $\mathcal{G} \neq \emptyset$. If the pairs $(K_i, \mu_i), i \in I$, form a chain, there is an upper bound (K, μ) for the chain, where K is the union of the K_i and μ coincides with μ_i on each K_i. By Zorn's lemma, $\mathcal{G}$ has a maximal element (K_0, τ). If $K_0 \subset E$, let $\alpha \in E \setminus K_0$. By (3.5.12), τ has an extension to $K_0(\alpha)$, contradicting maximality of (K_0, τ). ♣

A3.6 Corollary

In (3.5.12), if E is algebraically closed and L is algebraic over $\sigma(F)$, then τ is an isomorphism.

Proof. Since E is algebraically closed, so is $\tau(E)$. Since L is algebraic over $\sigma(F)$, it is algebraic over the larger field $\tau(E)$. By (1) $\iff$ (4) in (3.3.1), $L = \tau(E)$. ♣

A3.7 Theorem

Any two algebraic closures L and E of a field F are F-isomorphic.

Proof. We can assume that F is a subfield of L and $\sigma\colon F \to L$ is the inclusion map. By (3.5.12), σ extends to an isomorphism τ of E and L, and since τ is an extension of σ, it is an F-monomorphism. ♣

A3.8 Theorem (=Theorem 3.3.9)

If E is an algebraic extension of F and C is an algebraic closure of F, then any embedding of F into C can be extended to an embedding of E into C.

Proof. Repeat the proof of (3.5.12), with the mapping μ required to be an embedding. ♣

A3.9 Remark

The argument just given assumes that E is a subfield of C. This can be assumed without loss of generality, by (3.5.12), (3.5.12) and (3.5.12). In other words, we can assume that an algebraic closure of F contains a specified algebraic extension of F.

A3.10 Theorem

Let E be an algebraic extension of F, and let L be the algebraic closure of F containing E (see 3.5.12). If σ is an F-monomorphism from E to L, then σ can be extended to an automorphism of L.

Proof. We have L algebraically closed and L/E algebraic, so by (3.5.12) with E replaced by L and F by E, σ extends to a monomorphism from L to L, an F-monomorphism by hypothesis. The result follows from (3.5.12). ♣