6. Fields I

1. *Adjoining things*

The general *intention* of *adjoining* a new element α to a field k is arguably clear: k itself does *not* contain a root of an equation, and we want to enlarge k so that it *does* include such a root. The possibility or legitimacy of doing so may seem to depend upon one's philosophical outlook, but the situation is more robust than that. [1]

Let k be a field. Let $k \subset K$ where K is a bigger field. For $\alpha \in K$, define the **field extension (in K) over k generated by α** [2]

$$k(\alpha) = \bigcap_{\text{fields } E \subset K,\ E \supset k,\ \alpha \in E} E$$

It is easy to check that the intersection of subfields of a common field is a field, so this intersection is a field. Rather than a single element, one could as well adjoin any subset of the over-field K. [3]

Before studying $k(\alpha)$ in more detail, consider a different procedure of *adjoining* something: for a commutative

[1] In the 19^{th} century there was widespread confusion or at least concern over issues of existence of *quantities* having various properties. Widespread belief in the legitimacy of the complex numbers was not in place until well into that century, and ironically was abetted by pictorial emphasis on complex numbers as *two-dimensional* things. The advent of the Hamiltonian quaternions in mid-century made the complex numbers seem innocent by comparison.

[2] The notation here is uncomfortably fragile: exchanging the parentheses for any other delimiters alters the meaning.

[3] This definition does *not* give a good computational handle on such field extensions. On the other hand, it is unambiguous and well-defined.

ring R with 1 which is a subring of an R-algebra A, for $\alpha \in A$, one might *attempt* to define [4]

$$R[\alpha] = \{ \text{ polynomials in } \alpha\}$$

One probably understands the intent, that this is

$$R[\alpha] = \{c_0 + c_1\alpha + \ldots + c_n\alpha^n : c_i \in R\}$$

More precisely, a proper definition would be

$$R[\alpha] = \text{the image in } A \text{ of the unique } R\text{-algebra homomorphism sending } x \text{ to } \alpha$$

where we invoke the universal mapping property of $R[x]$.

Specialize R again to be a field k, and let A be a (not necessarily commutative) k-algebra, $\alpha \in A$. Then the natural homomorphism

$$\varphi : k[x] \longrightarrow k[\alpha] \qquad (\text{by } x \longrightarrow \alpha)$$

has a kernel which is a principal ideal $\langle f\rangle$. [5] By the usual Isomorphism Theorem the map φ *descends* to the quotient by the kernel, giving an isomorphism

$$\overline{\varphi} : k[x]/\langle f\rangle \approx k[\alpha]$$

If $f = 0$, that is, if the kernel is trivial, then $k[\alpha]$ of course inherits properties of the polynomial ring. [6]

At this point we need to begin using the fact that a k-algebra A is a k-vectorspace. [7] The **degree** of A over k is

$$[A : k] = \text{degree of } A \text{ over } k = \text{ dimension of } A \text{ as } k\text{-vectorspace}$$

If $k[\alpha] \approx k[x]$, then, for example, the various powers of α are linearly independent over k, and $k[\alpha]$ is infinite-dimensional as a k-vectorspace. And there is *no* polynomial $P(x) \in k[x]$ such that $P(\alpha) = 0$. Especially in the simple situation that the k-algebra A is a *field*, such elements α with $k[\alpha] \approx k[x]$ are **transcendental** over k. [8]

On the other hand, a perhaps more interesting situation is that in which the kernel of the natural

$$k[x] \longrightarrow k[\alpha]$$

has non-zero kernel $\langle f\rangle$, with f monic without loss of generality. This f is the **minimal polynomial** of α (in A) over k.

Although our immediate concern is field extensions, there is at least one other useful application of this viewpoint, as follows. Let V be a k-vectorspace, and let A be the k-algebra

$$A = \text{End}_k V$$

[4] Note again the fragility of the notation: $k(\alpha)$ is generally quite different from $k[\alpha]$, although in some useful cases (as below) the two can coincide.

[5] ... since $k[x]$ is a principal ideal domain for k a field. For more general commutative rings R the corresponding discussion is more complicated, though not impossible.

[6] ... to which it is *isomorphic* by the just-demonstrated isomorphism!

[7] By *forgetting* the multiplication in A, if one insists.

[8] This is an essentially *negative* definition: there are no relations.

of k-linear maps (i.e., **endomorphisms**) of V to itself. For $T : V \longrightarrow V$ a k-linear map, we can consider the natural k-algebra map

$$k[x] \longrightarrow \mathrm{End}_k V \quad (\text{by } x \longrightarrow T)$$

We give $\mathrm{End}_k V$ a k-vectorspace structure value-wise by

$$(\alpha \cdot T)(v) = \alpha \cdot (Tv)$$

for $v \in V$ and $\alpha \in k$. If V is finite-dimensional, then $\mathrm{End}_k V$ is also finite-dimensional. [9] In particular, the kernel of the natural map from $k[x]$ cannot be just 0. Let f be the non-zero monic generator for the kernel. Again, [10] this monic is the **minimal polynomial** for T. The general construction shows that for any $P(x) \in k[x]$,

$$P(T) = 0 \in \mathrm{End}_k V \quad \text{if and only if } f \text{ divides } P$$

In particular, if the polynomial equation $f(x) = 0$ has a *root* λ in k, then [11] we can prove that T has **eigenvalue** λ. That is, there is a non-zero vector $v \in V$ (the λ-**eigenvector**) such that

$$Tv = \lambda \cdot v$$

Indeed, let $f(x) = (x - \lambda) \cdot g(x)$ for some $g(x) \in k[x]$. Since g is not the minimal polynomial for T, then there is a vector $w \in V$ such that $g(T) \cdot w \neq 0$. We claim that $v = g(T)w$ is a λ-eigenvector. Indeed,

$$0 = f(T) \cdot w = (T - \lambda) \cdot g(T)w = (T - \lambda) \cdot v$$

and by the previous comment $v = g(T)w$ is not 0. [12]

Returning to field extensions: let K be a field containing a smaller field k, $\alpha \in K$, and let f be the generator for the kernel of the natural map $k[x] \longrightarrow k[\alpha]$. We do assume that f is non-zero, so we can make f monic, without loss of generality. Since f is non-zero, we do call it the **minimal polynomial** of α over k, and, since α *has* a minimal polynomial over k, we say that α is **algebraic** over k. [13] If every element α of a field extension K of k is algebraic over k, then say that the field extension K itself is **algebraic** over k.

Once again, given any polynomial $P(x)$, there are *unique* $Q(x)$ and $R(x)$ with $\deg R < \deg f$ such that

$$P = Q \cdot f + R$$

and

$$P(\alpha) = Q(\alpha) \cdot f(\alpha) + R(\alpha) = Q(\alpha) \cdot 0 + R(\alpha) = R(\alpha)$$

[9] This is not hard to prove: let $e_1, \ldots, e_n$ be a k-basis for V. Then the k-linearity $T(\sum_i c_i e_i) = \sum_i c_i T(e_i)$ shows that T is determined completely by the collection of images Te_i. And $Te_i = \sum_j T_{ij} e_j$ for some collection of n^2 elements T_{ij} of k. Thus, if V is n-dimensional then its endomorphism algebra is n^2-dimensional.

[10] This is terminology completely consistent with linear algebra usage.

[11] From the fact that roots correspond perfectly to linear factors, for polynomials in one variable with coefficients in a field.

[12] Even this brief discussion of minimal polynomials and linear operators should suggest, and correctly so, that use of determinants and invocation of the Cayley-Hamilton theorem, concerning the *characteristic polynomial* of a linear operator, is not exactly to the point.

[13] Again, this situation, where $f(\alpha) = 0$ with a non-zero polynomial f, is in contrast to the case where α satisfies *no* algebraic equation with coefficients in k.

Letting $n = \deg f$, this implies that $1, \alpha, \alpha^2, \ldots, \alpha^{n-1}$ are a k-basis for $k[\alpha]$. [14]

[1.0.1] **Proposition:** For α algebraic over k (all inside K), the ring $k[\alpha]$ is a *field.* [15] That is, for α *algebraic*, $k(\alpha) = k[\alpha]$. The minimal polynomial f of α over k is *irreducible* in $k[x]$. And the degree (dimension) of $k(\alpha)$ over k is

$$[k(\alpha) : k] = \dim_k\, k(\alpha) = \deg f$$

Proof: First, from above,

$$k[\alpha] \approx k[x]/\langle f\rangle$$

To prove irreducibility, suppose we can write $f = g \cdot h$ with $g, h \in k[x]$ with proper factors. By minimality of f, neither $g(\alpha)$ nor $h(\alpha)$ is 0. But $f(\alpha) = 0$, so $g(\alpha$ and $h(\alpha)$ are zero-divisors, contradiction. [16]

Since $k(\alpha)$ is the smallest field inside the ambient field K containing α and k, certainly $k[\alpha] \subset k(\alpha)$. To prove equality, it would suffice to show that non-zero elements of $k[\alpha]$ have multiplicative inverses in $k[\alpha]$. For polynomial $g(x) \in k[x]$, $g(\alpha) \neq 0$ if and only if the minimal polynomial $f(x)$ of α over k does not divide $g(x)$. Since f is irreducible and does not divide g, there are polynomials r, s in $k[x]$ such that

$$1 = \gcd(f, g) = r(x) \cdot f(x) + s(x) \cdot g(x)$$

so, mapping x to α,

$$1 = r(\alpha) \cdot f(\alpha) + s(\alpha) \cdot g(\alpha) = r(\alpha) \cdot 0 + s(\alpha) \cdot g(\alpha) = s(\alpha) \cdot g(\alpha)$$

That is, $s(\alpha)$ is a multiplicative inverse to $g(\alpha)$, and $k[\alpha]$ is a field. The degree is as asserted, since the polynomials of degree $< \deg f$ are irredundant representatives for the equivalence classes of $k[x]/\langle f\rangle$.
///

2. *Fields of fractions, fields of rational functions*

For $k \subset K$ fields and $\alpha \in K$ transcendental over k, it is not true that $k[\alpha] \approx k(\alpha)$, in complete contrast to the case that α is algebraic, discussed just above. [17]

But from elementary mathematics we have the idea that for *indeterminate* [sic] x

$$k(x) = \text{ field of rational functions in } x \ = \{\frac{g(x)}{h(x)} : g, h \in k[x],\ h \neq 0\}$$

We can reconcile this primitive idea with our present viewpoint.

Let R be an *integral domain* (with unit 1) [18] and define the **field of fractions** Q of R to be the collection of ordered pairs (r, s) with $r, s \in R$ and $s \neq 0$, modulo the equivalence relation [19]

$$(r, s) \sim (r', s') \qquad \text{if} \qquad rs' = sr'$$

[14] Indeed, the identity $P = Qf + R$ shows that any polynomial in α is expressible as a polynomial of degree $< n$. This proves spanning. On the other hand, a linear dependence relation $\sum_i c_i \alpha^i = 0$ with coefficient c_i in k is nothing other than a polynomial relation, and our hypothesis is that any such is a (polynomial) multiple of f. Thus, the monomials of degrees less than $\deg f$ are linearly independent.

[15] This should be a little surprising.

[16] Everything is taking place inside the larger field K.

[17] In particular, since $k[\alpha] \approx k[x]$, $k[\alpha]$ is not a field at all.

[18] A definition can be made for more general commutative rings, but the more general definition has more complicated features which are not of interest at the moment.

[19] This would be the usual requirement that two fractions r/s and r'/s' be equal.

The addition is suggested by the usual addition of fractions, namely that

$$(r,s)+(r',s')=(rs'+r's,ss')$$

and the multiplication is the more obvious

$$(r,s)\cdot(r',s')=(rr',ss')$$

One should verify that these operations are well-defined on the quotient Q by that equivalence relation, that Q is a commutative ring with unit (the equivalence class of)$(1,1)$, that $r \longrightarrow (r,1)$ *injects* R to Q. This constructs the field of fractions.

The latter construction is *internal*, in the sense that it constructs a concrete thing in set-theoretic terms, given the original ring R. On the other hand, we can characterize the field of fractions *externally*, by properties of its mappings to other rings or fields. In particular, we have

[2.0.1] **Proposition:** For an integral domain R with unit 1, its field of fractions Q, with the natural inclusion $i : R \longrightarrow Q$, is the unique field (and inclusion of R into it) such that, for any *injective* ring homomorphism $\varphi : R \longrightarrow K$ with a field K, there is a unique $\tilde{\varphi} : Q \longrightarrow K$ such that

$$\varphi \circ i = \tilde{\varphi}$$

Specifically, $\tilde{\varphi}(r,s)=\varphi(r)/\varphi(s)$. [20]

[20] Implicitly we must claim that this is well-defined.

Proof: Indeed, try to define [21]

$$\tilde{\varphi}(r,s) = \varphi(r)/\varphi(s)$$

where the quotient on the right-hand side is in the field K, and the injectivity of φ assure that $s \neq 0$ implies that $\varphi(s) \neq 0$. This is certainly compatible with φ on R, since

$$\tilde{\varphi}(r,1) = \varphi(r)/\varphi(1) = \varphi(r)$$

and the smallest subfield of K containing R certainly must contain all such quotients. The main thing to check is that this definition really is well-defined, namely that if $(r,s) \sim (r',s')$, then

$$\tilde{\varphi}(r,s) = \tilde{\varphi}(r',s')$$

Do this as follows. The equivalence relation is that $rs' = r's$. Applying φ on R gives

$$\varphi(r)\varphi(s') = \varphi(r')\varphi(s)$$

Since φ is injective, for s, s' nonzero in R their images are nonzero in K, so we can divide, to obtain

$$\varphi(r)/\varphi(s) = \varphi(r')/\varphi(s')$$

This proves the well-definedness. That multiplication is preserved is easy, and that addition is preserved is straightforward. ///

To practice categorical arguments, we can also prove, without using formulas or explicit constructions:

[2.0.2] Proposition: Let Q' be a field with inclusion $i' : R \longrightarrow Q'$ such that, for every injective homomorphism $\varphi : R \longrightarrow K$ with a field K, there is a unique $\tilde{\varphi} : Q' \longrightarrow K$ such that

$$\varphi \circ i' = \tilde{\varphi}$$

Then there is a unique isomorphism $j : Q \longrightarrow Q'$ of the field of fractions Q of R (with inclusion $i : R \longrightarrow Q$) to Q' such that

$$i' = j \circ i$$

That is, up to unique isomorphism, there is only one field of fractions of an integral domain.

Proof: First prove that any field map $f : Q \longrightarrow Q$ such that $f \circ i = i$ must be the identity on Q. Indeed, taking $K = Q$ and $f = i : R \longrightarrow K$ in the defining property, we see that the identity map id_Q on Q has the property $\mathrm{id}_K \circ i = i$. The *uniqueness* property assures that any other f with this property must be id_K.

Then let Q' and $i' : R \longrightarrow Q'$ be another pair satisfying the universal mapping condition. Taking $\varphi = i' : R \longrightarrow Q'$ yields $\tilde{\varphi} : Q \longrightarrow Q'$ with $\varphi = \tilde{\varphi} \circ i$. Reversing the roles, taking $\varphi' = i : R \longrightarrow Q$ yields $\tilde{\varphi'} : Q' \longrightarrow Q$ with $\varphi' = \tilde{\varphi'} \circ i'$. Then (by the previous paragraph) $\tilde{\varphi} \circ \tilde{\varphi'} : Q \longrightarrow Q$ must be the identity on Q, and, similarly, $\tilde{\varphi'} \circ \tilde{\varphi} : Q' \longrightarrow Q$; must be the identity on Q'. Thus, $\tilde{\varphi}$ and $\tilde{\varphi'}$ are mutual inverses. This proves the isomorphism of the two objects. [22] ///

Thus, without having a larger field in which the polynomial ring $k[x]$ sits, we simply form the field of fractions of this integral domain, and denote it [23]

$$k(x) = \text{ field of fractions of } k[x] = \text{ rational functions in } x$$

[21] What else could it be?

[22] The *uniqueness* of the isomorphism also follows from discussion, since if there were two isomorphisms h and h' from Q to Q', then $h' \circ h^{-1} : Q \longrightarrow Q$ would be a non-identity map with the desired property, but only the identity on Q has the universal mapping property.

[23] To say that these are rational *functions* is a bit of a misnomer, but no worse than to refer to polynomial *functions*, which is also misleading but popular.

Despite having this construction available, it still may be true that for fields $k \subset K$, there is α in K *transcendental* over k, in the sense (above) that α satisfies no polynomial relation with coefficients in k. [24] In that case, we have the more general definition of $k(\alpha)$ as the intersection of all subfields of K containing k and containing α.

For notational consistency, we should check that $k(\alpha)$ is isomorphic to the field of fractions of $k[\alpha]$. And, indeed, since $k[x]$ injects to $k[\alpha]$ (taking x to α), by the mapping property characterization the field of fractions $k(x)$ of $k[x]$ has a unique injection j to the field $k(\alpha)$ extending the given map. Certainly $k(\alpha) \subset j(k(x))$, since $k(\alpha)$ is the intersection of all subfields of K containin k and α. Thus, the image of the injective map j is exactly $k(\alpha)$, and j is an isomorphism of $k(x)$ to $k(\alpha)$.

3. *Characteristics, finite fields*

The linear algebra viewpoint is decisive in understanding many elementary features of fields, for example, the result below on possible cardinalities of finite fields.

First, observe that any ring R is a $\mathbb{Z}$-algebra in a canonical [25] manner, with the action

$$n \cdot r = \begin{cases} \underbrace{r + \ldots + r}_{n} & (n > 0) \\ 0_R & (n = 0) \\ -(\underbrace{r + \ldots + r}_{|n|}) & (n < 0) \end{cases}$$

An easy but tedious induction proves that this $\mathbb{Z}$-algebra structure deserves the name. [26] As evidence for the naturality of this $\mathbb{Z}$-structure, notice that if $f : R \longrightarrow S$ is any ring homomorphism, then f is a $\mathbb{Z}$-algebra homomorphism when the above $\mathbb{Z}$-algebra structures are put on R and S.

When a ring R has an identity 1_R, there is a canonical $\mathbb{Z}$-algebra homomorphism $i : \mathbb{Z} \longrightarrow R$ by

$$i : n \longrightarrow n \cdot 1_R$$

Granting that the $\mathbb{Z}$-algebra structure on R works as claimed, the proof that this is a homomorphism is nearly trivial:

$$i(m + n) = (m + n) \cdot 1_R = m \cdot 1_R + n \cdot 1_R = i(m) + i(n)$$

$$i(m \cdot n) = (m \cdot n) \cdot 1_R = m \cdot (n \cdot 1_R) = m \cdot (1_R \cdot (n \cdot 1_R)) = (m \cdot 1_R) \cdot (n \cdot 1_R) = i(m) \cdot i(n)$$

Now consider the canonical $\mathbb{Z}$-algebra homomorphism $i : \mathbb{Z} \longrightarrow k$ for a field k. [27] If i is injective, then it extends to an injection of the field of fractions $\mathbb{Q}$ of $\mathbb{Z}$ into k. In this case, say k is of **characteristic zero**, and this canonical copy of $\mathbb{Q}$ inside k is the **prime field** inside k. If i is not injective, its kernel is a

[24] Again, more precisely, the condition that α be transcendental is that the natural map $k[x] \longrightarrow k[\alpha]$ by $x \longrightarrow \alpha$ has trivial kernel.

[25] This sort of use of *canonical* is meant for the moment to insinuate that there is no whimsical choice involved. A more precise formulation of what *canonical* could mean would require a category-theoretical set-up. We may do this later.

[26] The arguments to prove this are of the same genre as those proving the so-called Laws of Exponents. Here, one must show that $(m + n)r = mr + nr$ and $(mn)r = m(nr)$ for $m, n \in \mathbb{Z}$ and $r \in R$, and $m(rs) = (mr)s$ for $s \in R$.

[27] It is no coincidence that we begin our study of fields by considering homomorphisms of the two simplest interesting rings, $k[x]$ for a field k, and $\mathbb{Z}$, into rings and fields.

principal ideal in $\mathbb{Z}$, say $p\mathbb{Z}$ with $p > 0$. Since the image $i(\mathbb{Z})$ is inside a field, it is an integral domain, so $p\mathbb{Z}$ is a (non-zero) prime ideal, which implies that p is prime. This integer p is the **characteristic** of k. We know that $\mathbb{Z}/\langle p\rangle$ is a *field*. Then we see that (by the Isomorphism Theorem for rings) the homomorphism $i : \mathbb{Z} \longrightarrow k$ with kernel $p\mathbb{Z}$ induces an isomorphism

$$\mathbb{Z}/p \approx i(\mathbb{Z}) \subset k$$

This canonical copy of $\mathbb{Z}/p$ inside k is the **prime field** inside k.

A finite field with q elements is often denoted $\mathbb{F}_q$ or $GF(q)$. [28]

[3.0.1] **Theorem:** A finite field K has p^n elements for some prime p and integer n. [29] In particular, let $n = [K : \mathbb{F}_p]$ be the degree of K over its prime field $\mathbb{F}_p \approx \mathbb{Z}/p$ with prime p. Then

$$|K| = p^n$$

Proof: Let $\mathbb{F}_p$ be the prime field in K. Let $e_1, \ldots, e_n$ be a $\mathbb{F}_p$-basis for the $\mathbb{F}_p$-vectorspace K. Then there are p^n choices of coefficients $c_i \in \mathbb{F}_p$ to form linear combinations

$$\alpha = \sum_{i=1}^{n} c_i\, e_i \in K$$

so K has p^n elements. ///

[28] This notation begs the question of *uniqueness* (up to isomorphism) of a finite field once its cardinality is specified. We address this shortly.

[29] We will prove *existence* and uniqueness results for finite fields a bit later.

4. *Algebraic field extensions*

The first of the following two examples is amenable to *ad hoc* manipulation, but the second is designed to frustrate naive explicit computation.

[4.0.1] **Example:** Let γ be a root (in some field k of characteristic 0, thus containing the prime field $\mathbb{Q}$) of the equation

$$x^2 - \sqrt{2}x + \sqrt{3} = 0$$

Is γ a root of a polynomial equation with rational coefficients?

In the same spirit as *completing the square*, we can manipulate the equation $x^2 - \sqrt{2}x + \sqrt{3} = 0$ to make the square roots disappear, as follows. Move the x^2 to the opposite side and square both sides, to obtain

$$2x^2 - 2\sqrt{6}x + 3 = x^4$$

Then move everything but the remaining square root to the right-hand side

$$-2\sqrt{6}\,x = x^4 - 2x^2 - 3$$

and square again

$$24x^2 = x^8 - 4x^6 - 2x^4 + 6x^2 + 9$$

and then we find that γ is a root of

$$0 = x^8 - 4x^6 - 2x^4 - 18x^2 + 9$$

It is not so obvious that the original [30]

$$\gamma = \frac{\sqrt{2} \pm \sqrt{2 - 4\sqrt{3}}}{2}$$

are roots. [31]

[4.0.2] **Example:** Let α be a root of the equation

$$x^5 - x + 1 = 0$$

and let β be a root of the equation

$$x^7 - x + 1 = 0$$

Then let γ be a root of the equation

$$x^6 - \alpha x + \beta = 0$$

Is γ a root of a polynomial equation with rational coefficients?

In this second example manipulations at the level of the first example fail. [32] But one might speculate that in answering an existential question it might be possible to avoid explicit computations entirely, as in the proofs of the following results.

[30] Solving the original quadratic equation directly, by completing the square, for example.

[31] For that matter, it appears that the original equation has exactly two roots, while a degree 8 equation might have 8. Thus, we seem to have introduced 6 *spurious* roots in this process. Of course, an explanation for this is that there are two different square roots of 2 and two different square roots of 3 in k, so really $2 \cdot 2 = 4$ versions of the original quadratic equation, each with perhaps 2 roots in k.

[32] The provable limitations of familiar algebraic operations are packaged up in *Galois theory*, a bit later.

[4.0.3] **Proposition:** Let $k \subset K \subset L$ be fields, with $[K : k] < \infty$ and $[L : K] < \infty$. Then

$$[L : k] = [L : K] \cdot [K : k] < \infty$$

In particular, for a K-basis $\{E_i\}$ of L, and for a k-basis e_j of K, the set $\{E_i e_j\}$ is a k-basis for L. [33]

Proof: On one hand, any linear relation

$$\sum_{ij} A_{ij}\, E_i e_j = 0$$

with $A_{ij} \in k$ gives

$$\sum_i (\sum_j A_{ij}\, e_j)\, E_i = 0$$

so for each i we have $\sum_j A_{ij}\, e_j = 0$, by the linear independence of the E_i. And by the linear independence of the e_j we find that $A_{ij} = 0$ for all indices. On the other hand, given

$$\beta = \sum_i b_i E_i \in L$$

with $b_i \in K$, write $b_i = \sum_j a_{ij} e_j$ with $a_{ij} \in k$, and then

$$\beta = \sum_i (\sum_j a_{ij} e_j)\, E_i = \sum_{ij} a_{ij}\, E_i e_j$$

which proves the spanning property. Thus, the elements $E_i e_j$ are a k-basis for L. ///

A field extension K of a field k is **finite** if the degree $[K : k]$ is finite. *Finite* field extensions can be built up by adjoining elements. To economize on parentheses and brackets, [34] write

$$k(\alpha_1, \ldots, \alpha_n) \quad \text{for} \quad k(\alpha_1)(\alpha_2)\ldots(\alpha_n)$$

and

$$k[\alpha_1, \ldots, \alpha_n] \quad \text{for} \quad k[\alpha_1][\alpha_2]\ldots[\alpha_n]$$

[4.0.4] **Proposition:** Let K be a field containing k, and suppose that $[K : k] < \infty$. Then any element α in K is algebraic over k, and there are finitely-many $\alpha_1, \ldots, \alpha_n$ such that

$$K = k(\alpha_1, \ldots, \alpha_n) = k[\alpha_1, \ldots, \alpha_n]$$

In particular, *finite* extensions K are necessarily *algebraic.* [35]

Proof: Given $\alpha \in K$, the countably many powers $1, \alpha, \alpha^2, \ldots$ cannot be linearly independent over k, since the whole K is finite-dimensional over k. A linear dependence relation among these powers is a polynomial

[33] The first assertion is merely a qualitative version of the last. Note that this proposition does not mention field *elements* explicitly, but rather emphasizes the vector space structures.

[34] On might worry that this notation glosses over potential issues. But, for example, one can prove that a polynomial ring in two variables really is *naturally* isomorphic to a polynomial ring in one variable over a polynomial ring in one variable.

[35] The converse is not true. That is, some fields k admit extensions K with the property that every element in K is algebraic over k, but K is infinite-dimensional over k. The rational numbers $\mathbb{Q}$ can be proven to have this property, as do the p-adic numbers $\mathbb{Q}_p$ discussed later. It is not completely trivial to prove this.

equation satisfied by α. [36] If K is strictly larger than k, take $\alpha_1 \in K$ but not in k. Then $[k(\alpha_1) : k] > 1$, and the multiplicativity

$$[K : k] = [K : k(\alpha_1)] \cdot [k(\alpha_1) : k]$$

with $[K : k] < \infty$ implies that

$$[K : k(\alpha_1)] < [K : k]$$

If K is still larger than $k(\alpha_1)$, take α_2 in K not in $k(\alpha_1)$. Again,

$$[K : k(\alpha_1, \alpha_2)] < [K : k(\alpha_1)] < [K : k]$$

These degrees are positive integers, so a decreasing sequence must reach 1 in finitely-many steps (by Well-Ordering). The fact [37] that $k(\alpha) = k[\alpha]$ for α algebraic over k was proven earlier. ///

Let K and L be subfields of a larger field E. The **compositum** $K \cdot L$ of K and L is the smallest subfield of E containing both K and L. [38]

[4.0.5] Proposition: Let $k \subset E$ be fields. Let K, L be subfields of K containing k. Suppose that $[K : k] < \infty$ and $[L : k] < \infty$. Then

$$[K \cdot L \,:\, k] \le [K : k] \cdot [L : k] < \infty$$

In particular, if

$$K = k(\alpha_1, \ldots, \alpha_m) = k[\alpha_1, \ldots, \alpha_m]$$

$$L = k(\beta_1, \ldots, \beta_n) = k[\beta_1, \ldots, \beta_n]$$

then

$$K \cdot L = k(\alpha_1, \ldots, \alpha_m, \beta_1, \ldots, \beta_n) = k[\alpha_1, \ldots, \alpha_m, \beta_1, \ldots, \beta_n]$$

Proof: From the previous proposition, there do exist the α_i and β_j expressing K and L as k with finitely many elements adjoined as in the statement of the proposition. Recall that these mean that

$$K = \text{ intersection of subfields of } E \text{ containing } k \text{ and all } \alpha_i$$

$$L = \text{ intersection of subfields of } E \text{ containing } k \text{ and all } \beta_i$$

On one hand, $K \cdot L$ contains all the α_i and β_j. On the other hand, since these elements are algebraic over k, we do have

$$k(\alpha_1, \ldots, \alpha_m, \beta_1, \ldots, \beta_n) = k[\alpha_1, \ldots, \alpha_m, \beta_1, \ldots, \beta_n]$$

The left-hand side is a field, by definition, namely the smallest subfield [39] of E containing all the α_i and β_j. Thus, it contains K, and contains L. Thus, we have equality. ///

[36] A more elegant argument is to map $k[x]$ to K by $x \longrightarrow \alpha$, and note that the kernel must be non-zero, since otherwise the image would be infinite-dimensional over k.

[37] Potentially disorienting and quite substantial.

[38] As with many of these constructions, the notion of *compositum* does not make sense, or at least is not well-defined, unless the two fields lie in a common larger field.

[39] This discussion would appear to depend perhaps too much upon the larger ambient field E. In one sense, this is true, in that *some* larger ambient field is necessary. On the other hand, if K and L are both contained in a smaller subfield E' of E, we can replace E by E' for this discussion. One may reflect upon the degree to which the outcome genuinely depends upon any difference between E' and E, and how to avoid this concern.

[4.0.6] **Proposition:** Let $k \subset E$ be fields, and K, L fields between k and E. Let $\alpha \in L$ be algebraic over k. Then

$$[k(\alpha) : k] \geq [K(\alpha) : K]$$

Proof: Since α is algebraic over k, $k(\alpha) = k[\alpha]$, and the degree $[k(\alpha) : k]$ is the degree of the minimal polynomial of α over k. This degree cannot increase when we replace k by the larger field K, and we obtain the indicated inequality. ///

[4.0.7] **Proposition:** Let k be a field, K a field algebraic over k, and L a field containing K. Let $\beta \in L$ be algebraic over K. Then β is algebraic over k.

Proof: Let $M(x)$ be the monic irreducible in $K[x]$ which is the minimal polynomial for β over K. Let $\{\alpha_0, \ldots, \alpha_{n-1}\}$ be the *finite* set (inside K) of coefficients of $M(x)$. Each field $k(\alpha_i)$ is of finite degree over k, so by the previous proposition their compositum $k(\alpha_1, \ldots, \alpha_n)$ is finite over k. The polynomial $M(x)$ is in $k(\alpha_1, \ldots, \alpha_n)[x]$, so β is algebraic over $k(\alpha_1, \ldots, \alpha_n)$. From above, the degree of $k(\alpha_1, \ldots, \alpha_n)(\beta)$ over k is the product

$$[k(\alpha_1, \ldots, \alpha_n)(\beta) : k] = [k(\alpha_1, \ldots, \alpha_n)(\beta) : k(\alpha_1, \ldots, \alpha_n)] \cdot [k(\alpha_1, \ldots, \alpha_n) : k] < \infty$$

Thus, $k(\alpha_1, \ldots, \alpha_n)(\beta)$ is finite over k, and in particular β is algebraic over k.

///

[4.0.8] **Corollary:** Let $k \subset K \subset L$ be fields, with K algebraic over k and L algebraic over K. Then L is algebraic over k.

Proof: This is an element-wise assertion, and for each β in L the previous proposition proves the algebraicity. ///

[4.0.9] **Remark:** An arrangement of fields of the form $k \subset K \subset L$ is sometimes called a **tower** of fields, with a corresponding picture

$$\begin{tikzcd} L \arrow[d, dash] \\ K \arrow[d, dash] \\ k \end{tikzcd}$$

The situation that K and L are intermediate fields between k and E, with compositum KL, is depicted as

$$\begin{tikzcd} & E \arrow[d, dash] & \\ & KL \arrow[dl, dash] \arrow[dr, dash] & \\ K \arrow[dr, dash] & & L \arrow[dl, dash] \\ & k & \end{tikzcd}$$

5. *Algebraic closures*

A field K is **algebraically closed** if every non-constant polynomial $f(x) \in k[x]$ has at least one root $\alpha \in k$, that is,

$$f(\alpha) = 0$$

Upon division, this algebraic closure property implies that any polynomial in $K[x]$ factors into linear factors in $K[x]$.

Given a field k, a larger field K which is algebraically closed [40] *and* such that every element of K is algebraic over k, is an **algebraic closure of** k. [41]

[5.0.1] Theorem: Any field k has an algebraic closure $\overline{k}$, unique up to isomorphism. Any algebraic field extension E of k has at least one injection to $\overline{k}$ (which restricts to the identity on k).

Proof: (Artin) Let S be the set of monic irreducibles in $k[x]$, for each $s \in S$ let x_s be an indeterminate, and consider the polynomial ring

$$R = k[\ldots, x_s, \ldots] \quad (s \in S)$$

in S-many variables. [42] We claim that there is at least one maximal proper ideal M in R containing every $f(x_f)$ for $f \in S$. First, one must be sure that the ideal F generated by all $f(x_f)$ is *proper* in R. If F were not proper, there would be elements $r_i \in R$ and irreducibles f_i such that (a finite sum)

$$\sum_{i=1}^{n} r_i\, f_i(x_{f_i}) = 1$$

Make a finite field extension E of k such that all the finitely-many f_i have roots α_i in E, inductively, as follows. First, let $k_1 = k[x]/\langle f_1\rangle$. Then let F_2 be an *irreducible* factor of f_2 in k_1, and let $k_2 = k_1[x]/\langle F_2\rangle$. And so on, obtaining $E = k_n$. Using the universal mapping property of polynomial rings, we can send x_{f_i} to $\alpha_i \in E$, thus sending $f_i(x_{f_i})$ to 0. [43] Then the relation becomes

$$0 = 1$$

Thus, there is no such relation, and the ideal F is proper.

Next, we claim that F lies in a *maximal* proper ideal M in R. This needs an equivalent of the Axiom of Choice, such as Hausdorff Maximality or Zorn's Lemma. In particular, among all chains of *proper* ideals containing F

$$F \subset \ldots \subset I \subset \ldots$$

there exists a maximal chain. [44] The union of an ascending chain of proper ideals cannot contain 1, or else one of the ideals in the chain would contain 1, and would not be proper. Thus, the union of the ideals in a maximal chain is still proper. If it were not a *maximal* proper ideal then there would be a further (proper)

[40] Note that not only polynomials with coefficients in k must have roots in K, but polynomials with coefficients in K. Thus, one can perhaps imagine a different universe in which one makes a large enough field K such that all polynomials with coefficients in k have roots, but polynomials with coefficients in K need a larger field for their roots. That this does not happen, and that the process of constructing algebraic closures terminates, is the content of the theorem below.

[41] The second requirement is desirable, since we do not want to have algebraic closures be needlessly large. That is, an algebraic closure of k should not contain elements transcendental over k.

[42] This ostentatiously extravagant construction would not have been taken seriously prior to Bourbaki's influence on mathematics. It turns out that once one sacrifices a *little* finiteness, one may as well fill things out symmetrically and accept a *lot* of non-finiteness. Such extravagance will reappear in our modern treatment of tensor products, for example.

[43] No, we have no idea what happens to the r_i, but we don't care.

[44] Maximal in the sense that there is no other proper ideal J containing F that either contains or is contained in every element of the (maximal) chain.

ideal that could be added to the chain, contrary to assumption. Thus, we have a maximal ideal M in R. Thus, $K = R/M$ is a field.

By construction, for monic irreducible (non-constant) f the equation $f(Y) = 0$ has a root in K, namely the image of x_f under the quotient map, since $f(x_f) \in M$ for all irreducibles f. This proves that all non-constant polynomials in $k[x]$ have roots in K.

Now we prove that every element in $\overline{k}$ is algebraic over k. Let α_f be the image of x_f in $kbar$. Since α_f is a zero of f it is algebraic over k. An element β of $\overline{k}$ is a polynomial in finitely-many of the α_fs, say $\alpha_{f_1}, \ldots, \alpha f_n$. That is, $\beta \in k[\alpha_1, \ldots, \alpha_n]$, which is a field since each α_i is algebraic over k. Since (for example) the compositum (inside $\overline{k}$) of the algebraic extensions $k(\alpha_{f_i}) = k[\alpha_{f_i}]$ is algebraic, β is algebraic over k.

Next, we prove that non-constant $F(x) \in \overline{k}[x]$ has a zero in $\overline{k}$ (hence, it has *all* zeros in $\overline{k}$). The coefficients of F involve some finite list $\alpha_{f_1}, \ldots, \alpha_{f_n}$ out of all α_f, and $F(x)$ has a zero in $\overline{k}(\alpha_{f_1}, \ldots, \alpha_{f_n})[x]/\langle F \rangle$. Thus, since β is algebraic over an algebraic extension of k, it is algebraic over k, and, thus, is a root of a polynomial in $k[x]$.

Now consider an algebraic extension E of k, and show that it admits an imbedding into $\overline{k}$. First, if $\alpha \in E$, let f be the minimal polynomial of α over k, and let β be a zero of f in $\overline{k}$. Map $k[x] \longrightarrow \overline{k}$ by sending $x \longrightarrow \beta$. The kernel is exactly the ideal generated by f, so (by an isomorphism theorem) the homomorphism $k[x] \longrightarrow \overline{k}$ descends to an injection $k[\alpha] \longrightarrow \overline{k}$. This argument can be repeated to extend the inclusion $k \subset \overline{k}$ to any extension $E = k(\alpha_1, \ldots, \alpha_n)$ with α_i algebraic over k. We use an equivalent of the Axiom of Choice to complete the argument: consider the collection of ascending chains of fields E_i (containing k) inside E admitting families of injections $\psi_i : E_i \longrightarrow \overline{k}$ with the compatibility condition that

$$\psi_j \,|_{E_i}| = \psi_i \quad \text{for} \quad E_i \subset E_j$$

We can conclude that there is a *maximal* chain. Let E' be the union of the fields in this maximal chain. The field E' imbeds in $\overline{k}$ by ψ_i on E_i, and the compatibility condition assures us that this is well-defined. We claim that $E' = E$. Indeed, if not, there is $\alpha \in E$ that is not in E'. But then the first argument shows that $E'(\alpha)$ does admit an imbedding to $\overline{k}$ extending the given one of E', contradiction. Thus, $E' = E$ and we have, in fact, imbedded the whole algebraic extension E to $\overline{k}$.

Last, we prove that any other algebraic closure K of k is isomorphic to $\overline{k}$. [45] Indeed, since K and $\overline{k}$ are algebraic over k, we have at least one injection $K \longrightarrow \overline{k}$, and at least one injection $\overline{k} \longrightarrow K$, but there is no reason to think that our capricious construction assures that these are mutual inverses. A different mechanism comes into play. Consider K imbedded into $\overline{k}$. Our claim is that K is necessarily all of $\overline{k}$. Indeed, any element of $\overline{k}$ is algebraic over k, so is the zero of a polynomial f in $k[x]$, say of degree n, which has all n roots in the subfield K of $\overline{k}$ because K is algebraically closed. That is, every element of the overfield $\overline{k}$ is actually in the subfield K, so the two are equal. ///

Exercises

6.[5.0.1] Let γ be a root of the equation $x^2 + \sqrt{5}x + \sqrt{2} = 0$ in an algebraic closure of $\mathbb{Q}$. Find an equation with *rational* coefficients having root γ.

6.[5.0.2] Let γ be a root of the equation $x^2 + \sqrt{5}x + \sqrt[3]{2} = 0$ in an algebraic closure of $\mathbb{Q}$. Find an equation with *rational* coefficients having root γ.

6.[5.0.3] Find a polynomial with rational coefficients having a root $\sqrt{2} + \sqrt{3}$.

[45] Note that we do not claim uniqueness of the isomorphism. Indeed, typically there are many different maps of a given algebraic closure $\overline{k}$ to itself that fix the underlying field k.

6.[5.0.4] Find a polynomial with rational coefficients having a root $\sqrt{2}+\sqrt[3]{5}$.

6.[5.0.5] Let γ be a root of $x^5 - x + 1 = 0$ in an algebraic closure of $\mathbb{Q}$. Find a polynomial with rational coefficients of which $\gamma + \sqrt{2}$ is a root.

6.[5.0.6] Show that the field obtained by adjoining $\sqrt{2}$, $\sqrt[4]{2}$, $\sqrt[8]{2}$, $\sqrt[16]{2}$, $\ldots$, $\sqrt[2^n]{2}$, $\ldots$, to $\mathbb{Q}$ is *not* of finite degree over $\mathbb{Q}$.