Chapter V. Algebraic Extensions.

Re §V.1. Finite and algebraic extensions.

P.223, first paragraph of §V.1 [>]: In addition to referring to a field E as an *extension* of a field F, Lang will frequently refer to the pair $F \subseteq E$ as ''an extension''.

Note that ''E is a finite extension of F'' does *not* mean E is a finite field.

P.224, top paragraph [~]: Here $F[\alpha]$ means the ring generated over F by α; cf. the notation $A[S]$ introduced on p.90.

Here and subsequently Lang will speak of polynomials ''with leading coefficient 1''. There is a standard term for such polynomials, which he introduced on p.175, but unfortunately rarely uses: *monic polynomials*.

What Lang calls ''the irreducible polynomial of α over F'' and writes $\mathrm{Irr}(\alpha, F, X)$ is generally called the *minimal polynomial* of α over F. Note the three arguments of ''Irr''. The significance of α is clear; the argument F is needed because a given element α will have different minimal polynomials over different fields (e.g., i has minimal polynomial X^2+1 over the real numbers, but $X-i$ over the complexes). Finally, if we are dealing with more than one indeterminate, then simply saying ''the minimal polynomial of i over the reals'' would not show whether we meant X^2+1 or Y^2+1, so the third argument is also necessary.

The degree of $\mathrm{Irr}(\alpha, F, X)$ is called the **degree** of α over F. Likewise, the number $[E:F]$ that Lang associates with the extension E of F later on this page is called the **degree** of the extension. (The index refers to this page for the term ''degree'', though the word does not appear on the page.)

P.225, Corollary V.1.3 [=]: This should begin with the same sentence as the preceding proposition.

P.225, proof of Proposition V.1.4 [<]: The first statement of the proposition is equivalent to saying that $k[\alpha]$ is a field. There is an easier proof than Lang gives, using the first display on the preceding page. For since $(p(X))$ is an ideal of a principal ideal domain generated by an irreducible element, it is a *maximal* ideal, so $k[X]/(p(X))$ will be a field.

P.226, second paragraph [>]: Note that where Lang defines the compositum EF, it is not asserted that this consists of products of elements of E and elements of F, or even of sums of such products. (What form *will* its elements have?) In the last line of that paragraph, for ''we cannot define the compositum'' read ''the compositum is not defined''; i.e., the above definition is not applicable.

You might find it interesting to prove that given subfields E, F of a field L, if we write $k = E \cap F$, then if either E or F is algebraic over k, their compositum EF will be generated by E and F as a ring, i.e., will consist of sums of products of elements of E and elements of F. On the other hand, for an example of a compositum *not* having this form, start with a field k, let $L = k(s,t)$ where s and t are algebraically independent over k, and let $E = k(s)$, $F = k(t)$. You will find that in the ring generated by E and F, $s-t$ is not invertible; hence the compositum EF is bigger than that ring.

(The word ''compositum'' is sometimes used in a slightly different way, though not in Lang. Given extensions E and F of a field k, ''*a* compositum of E and F'' means any extension L of k given with homomorphisms, i.e., embeddings, of E and F into it over k, such that L is generated by the images of those embeddings. Under that definition, one can prove that any two extensions of k *do* have a compositum, though it is not, in general, unique, even up to isomorphism.)

P.226, middle paragraph [>]: Compare this with the notation introduced in the third paragraph of p.110, where $K(X_1, \ldots, X_n)$ is defined as the field of fractions of the polynomial ring $K[X_1, \ldots, X_n]$. The convention of which these are both cases is that when we list some elements in parentheses after the name of a field, we are naming a field extension generated by those elements. Whether these elements are independent indeterminates, as on p.110, or elements of a given larger field, as on this page, must be deduced from context (sometimes from the symbols used). If there is danger of confusion, the writer should say explicitly what is meant.

In the next paragraph Lang will define a field E to be *finitely generated* over a subfield k if and only if it can be written $k(\alpha_1, \dots, \alpha_n)$ for some finite family of elements $\alpha_1, \dots, \alpha_n$ of E. However, such an extension will not necessarily be finitely generated over k as a *ring*. This is because an extension $k(\alpha_1, \dots, \alpha_n)$ may or may not coincide with its subring $k[\alpha_1, \dots, \alpha_n]$.

P.227, last paragraph [~]: Lang's term ''distinguished class'' is essentially a mnemonic: Various classes of field extensions share certain properties, so if we call a class of extensions having these properties ''distinguished'', then instead of having to memorize this list of properties for each such class, we simply have to recall that we have proved that the class is distinguished. Note, however, that this definition differs from the commonest mathematical use of the word ''distinguished'', where, for instance, if X is a set and x an element of X, then the pair (X, x) may be referred to as ''a set with a distinguished element''.

(I would suggest replacing Lang's phrase ''$\mathcal{C}$ is a distinguished class'' with the phrase ''$\mathcal{C}$ has the three standard properties''.)

P.228, third paragraph [~]: Here ''E/F'' means ''E, viewed as an extension of F''. This allows us to abbreviate statements like ''E is algebraic over F'' and ''E is finite over F'' to ''E/F is algebraic'' and ''E/F is finite''. This notation is commonly used, though its similarity to the notation for factor-groups, factor-rings etc. is unfortunate. (Stewart's little book *Galois Theory* introduces instead the symbol $E:F$, which does not have this disadvantage.)

P.229, end of §V.1 [>]: Here are some further notes on the material in this section.

Hints of what is to come.

Algebraic extensions of fields are curious in that their structure is neither ''very tight'' nor ''very loose''. If we map an algebraic extension field E of a field k into another extension F of k, then for each $\alpha \in E$, the element of F to which it is mapped must be a root of the minimal polynomial of α over k. A polynomial of degree n has at most n roots in a field; thus, though there may be more than one element of F to which α can be sent, there cannot be more than $\deg_k(\alpha)$. This ''semi-rigid'' structure on homomorphisms among algebraic extension fields can be best investigated by examining the group of *automorphisms* of a single field extension F of k. We can see from the above that if F is finite over k, then $\mathrm{Aut}_k(F)$ is finite, and we shall find that its structure is intimately connected with the structure of the extension. This is the subject of Galois Theory.

In the above paragraph we used the fact that an irreducible polynomial of degree n over k has *at most* n roots in an extension field of k. Such a polynomial usually has a full complement of n distinct roots in a big enough extension of k. Usually – but not always! The situation where it does not is called ''inseparability'', and will be studied in §V.6. Let us note one example here. Suppose k_0 is a field of finite characteristic p, let t be an indeterminate over k_0, let $L = k_0(t)$, and define $k = k_0(t^p) \subseteq L$. Then as an extension of k, L is generated by the element t, which is a root of the polynomial $X^p - t^p \in k[X]$. In $L[X]$, this polynomial splits into linear factors, $X^p - t^p = (X-t)^p$; but we see from this factorization that it has only one root, t. Galois theory is limited to the study of extensions in which this phenomenon does not occur. These are called *separable* extensions, and introduced in §V.4.

On fields, and universal constructions.

We encountered the concepts of ''product'' and ''coproduct'' in the study of groups and rings, and it is natural to ask whether such constructions exist for fields. In general, the answer is no. Given fields E and F, if we form their product as rings, this will have zero-divisors $(0,1)\cdot(1,0) = 0$, and so will not be a field; and, in fact, there is in general no field with the required universal property. Likewise, if we try to construct a coproduct by looking at what elements we can name in a field with homomorphisms of E and F into it, we find that we can name all the *ring-theoretic* expressions in such elements, but to get a field, we must also have inverses of all *nonzero* elements. However, a given ring-theoretic expression in elements of E and F might have the value zero in one field L_1 containing copies of E and F, but be nonzero in another such field L_2. Hence if we introduce an inverse to this expression, we lose universality (we can't map to L_1), while if we don't, we don't get inverses of all nonzero elements, and so don't have a field.

Why this difference between fields and other classes of algebraic objects? The other classes we have seen can each be defined as consisting of all sets given with certain everywhere-defined operations, subject to certain fixed identities. (To put the definitions in this form sometimes requires a reformulation. For instance, the condition that every element of a group have an inverse is not an identity. But if we define a group in terms of *three* operations, composition (a binary operation), inverse (a unary operation) and neutral element (a ''zeroary'' operation), we can then state *all* the group axioms as identities.) Such a class of objects is known in universal algebra as a *variety of algebras*, and one can prove that in any variety, one has free objects, products, coproducts, etc.. The class of fields fails to be a variety in that the multiplicative inverse operation is not everywhere defined – it is defined only on nonzero elements. Another important class of algebraic objects that fails to be a variety is the class of *integral domains*, since the condition ''$xy = 0 \Rightarrow x = 0$ or $y = 0$'' is not an identity.

It is interesting to note, however, that if E and F are extensions of a field k, and we form their coproduct $E \amalg F$ as commutative k-algebras, then all field extensions of k generated by an image of E and an image of F can be constructed, up to isomorphism, as the *fields of fractions* of the *factor-rings* $E \amalg F/P$, as P ranges over all prime ideals of $E \amalg F$.

Quaternions.

Though the present chapter concerns fields, I would like to briefly introduce you to one important example of a noncommutative division ring: the ring of quaternions.

This is an algebra over the real numbers, generally denoted $\mathbf{H}$ (for its discoverer, William Rowan Hamilton), which is 4-dimensional as an $\mathbf{R}$-vector-space, with basis denoted $\{1, i, j, k\}$, and with multiplication determined by the conditions that 1 be the multiplicative neutral element, and that the other basis elements satisfy

$$\text{(c30)} \qquad \begin{gathered} i^2 = j^2 = k^2 = -1, \\ ij = -ji = k, \quad jk = -kj = i, \quad ki = -ik = j. \end{gathered}$$

Thus, the general element of $\mathbf{H}$ has the form $a + bi + cj + dk$ $(a, b, c, d \in \mathbf{R})$, and a product of two such elements is calculated by applying bilinearity over $\mathbf{R}$, and evaluating each of the resulting summands using (c30). To verify that this multiplication is associative, it suffices to check the associative law for each of the $4\cdot 4\cdot 4 = 64$ products in which each term is taken from $\{1, i, j, k\}$. Actually, it is easy to see that associativity will be automatic whenever one of the factors is 1; so we are reduced to the 27 cases where each term is taken from $\{i, j, k\}$; and in fact, since (c30) is invariant under cyclic permutations of $\{i, j, k\}$, the problem reduces to checking the nine cases in which the first factor is i; a not unreasonable amount of work.

Once we have shown that $\mathbf{H}$ is an associative algebra, how do we prove that every element is invertible? By a trick very much like that used for the complex numbers. Given a quaternion $\alpha = a + bi + cj + dk$, define $\bar\alpha = a - bi - cj - dk$. Then it is easy to check that $\alpha\bar\alpha = a^2 + b^2 + c^2 + d^2$. (Do you see why the ''mixed terms'', e.g., those with coefficient ab, or bc, all cancel?) Since this real number $\alpha\bar\alpha$ is invertible, an inverse to α is given by $\bar\alpha/\alpha\bar\alpha$.

Some miscellaneous observations on the quaternions: It follows from (c30) that the eight elements ± 1, $\pm i$, $\pm j$, $\pm k \in \mathbf{H}$ form a group; this is the group Lang calls the ''quaternion group'' on p.9, last paragraph (though he writes m in place of -1, since ''-1'' is not meaningful in an abstract group). Note also that the second line of (c30) agrees with the ''cross product'' operation on $\mathbf{R}^3$, although, of course, the first line does not. The division ring $\mathbf{H}$ clearly contains three copies of the complex numbers, $\mathbf{R}(i)$, $\mathbf{R}(j)$

and $\mathbf{R}(k)$. In fact, it contains uncountably many such copies, since it is not hard to show that any quaternion $\alpha = bi + cj + dk$ such that $b^2 + c^2 + d^2 = 1$ satisfies $\alpha^2 = -1$, and generates a commutative subfield of $\mathbf{H}$ isomorphic to $\mathbf{C}$.

Clearly $\{a + bi + cj + dk \mid a, b, c, d \in \mathbf{Z}\}$ is a subring of $\mathbf{H}$; but in fact, it is easy to verify that the larger subset consisting of all elements $(a + bi + cj + dk)/2$ such that a, b, c, d are integers all congruent to one another mod 2 also forms a subring, and this turns out to have nicer arithmetical properties. It is therefore the latter ring that is called the ring of *integer quaternions*.

There are other important examples of noncommutative division rings than $\mathbf{H}$, but none as easily described or as well-known, and it was in fact the first to be discovered.

Re §V.2. Algebraic closure.

P.229, beginning of §V.2 [<]: Be prepared for some messiness here. We will be dealing with situations that would look neater if we identified isomorphic objects, but instead, we will explicitly name isomorphisms between them.

Given an isomorphism between two objects, isn't it always safe to make an identification, so as to regard the two objects as the same, and the isomorphism as the identity? It frequently is; but a big exception is when the objects are already the same, or lie within a common structure being considered, since then making such an identification would amount to considering distinct elements of that structure as the same. And, in fact, we will soon be using the results of this section in the case where the isomorphisms referred to are automorphisms (or partial automorphisms) of one field.

Nevertheless, until such assumptions are brought in, in understanding results that Lang states for a given embedding σ of one field in another, it can be useful to first picture the case where the embedding is an inclusion.

The need to avoid treating isomorphic objects as the same leads Lang to make such distinctions in this section even in cases when the two objects will not in fact be part of a common object; for instance in not identifying a field k with its canonical image in a field $k[X]/I$ (I an ideal of $k[X]/I$) at the bottom of p.230. However, thinking how to handle such distinctions rigorously is good practice for us.

P.229, next-to-last paragraph [>]: Note that Lang not only writes a^σ for $\sigma(a)$, but also f^σ for the polynomial gotten by applying σ to every coefficient of f, while leaving the indeterminate X unchanged, so that one is still free to substitute an arbitrary element for that indeterminate. (In the notation introduced on p.99, next-to-last display, this polynomial would be written (σf). Another way of formulating the final display on this page would be $(\sigma f)(\tau\alpha) = (\tau f)(\tau\alpha) = \tau(f(\alpha)) = \tau(0) = 0$.)

P.230, proof of Lemma V.2.2 [~]: This can be done without any computation. We know that $E_1 E_2$ can be characterized as the least subfield of E containing E_1 and E_2, and that $\sigma(E_1)\sigma(E_2)$ is likewise the least subfield of L containing $\sigma(E_1)$ and $\sigma(E_2)$. From this, and the fact that σ gives an isomorphism of E with $\sigma(E) \subseteq L$, it immediately follows that it takes $E_1 E_2$ to $\sigma(E_1)\sigma(E_2)$.

P.230, 5 lines from bottom [~]: Here Lang pauses to prove something he should have noted (and almost did) on p.93: That any homomorphism from a field into a nonzero ring has kernel 0.

Setting aside that distraction, note the point of this paragraph in Lang: he shows that, given a polynomial f over a field k, one can *construct* an extension of k in which f has a root. You have probably seen this argument in an undergraduate algebra course; if not, take careful note of the method. There is a minor technical complication, mentioned earlier, that the field one gets does not contain k itself, but an isomorphic copy σk thereof; this will be patched up in Proposition V.2.3, but meanwhile, it causes the result proved here to take the less transparent form ''$p^\sigma(\xi) = 0$'' instead of ''$p(\xi) = 0$''. Recall that $p^\sigma(X)$ is the polynomial obtained by applying σ (equivalently, the restriction of σ to k) to all the coefficients of p, but not to the indeterminate X.

It may seem confusing to have a polynomial in which the coefficients are elements of $k[X]/(p(X))$, and yet in which X appears as the indeterminate. If so, you might find the reasoning easier to follow if you give the indeterminate in p^σ a different name, writing something like $p^\sigma(Y)$. (Lang doesn't actually write ''$p^{\sigma}(X)$'' in the proof, but he writes p^{σ}; and in thinking about the argument, one should be able to refer to the indeterminate in that polynomial.)

P.231, end of proof of Proposition V.2.3 [>]: It would actually have been better to have proved a general lemma that if $\sigma: k \to F$ is an embedding of fields, then there is an extension E of k such that σ extends to an isomorphism $E \cong F$, since this is what the proof really shows. Lang will use this fact again on the next page, where he speaks of ''using the same type of set-theoretic argument as in Proposition 2.3''.

P.231, proof of Theorem V.2.5 [>].

Remarks on the construction of algebraically closed field extensions

To construct an algebraically closed extension field of k, Lang first adjoins a root of every polynomial over k, then does the same for every polynomial over the resulting field, and so on, iterating this countably many times, and finally takes the union. This is a useful general technique to be aware of. Actually, in *this* case, the successive iterations are not needed; one can prove that an algebraic extension of k in which every polynomial over k has a root is already algebraically closed. (This is Exercise VI:L28, p.325.) But at this point in the development, it is easier to invoke the countable chain of extensions than to prove that one iteration suffices!

One might ask whether one can't prove the existence of an algebraically closed extension of any field k by Zorn's Lemma. Clearly the union of any chain of algebraic extensions of k is again an algebraic extension. Doesn't Zorn's Lemma therefore give a maximal such extension? The trouble with this approach is that the algebraic extensions of k do not form a *set*. (Zorn's Lemma was only proved for partially ordered *sets*; and, indeed, if the restriction to sets could be removed, one could use the resulting ''lemma'' to prove such statements as that there was a largest set, which is absurd.)

But with some additional work, one can indeed get a Zorn's Lemma proof based on this idea. Given k, the first step is to get a bound on the cardinalities of algebraic extensions of k. I claim that such a bound is given by the cardinality of the polynomial ring $k[X]$. Indeed, any algebraic extension E of k can be mapped set-theoretically to $k[X]$ by sending each element to its minimal polynomial, and this map is finitely-many-to-one, since each polynomial has only finitely many roots in E. But a set A which can be mapped in a finitely-many-to-one fashion to an infinite set B has cardinality no greater than that of B, giving the asserted bound. (Lang will sketch approximately the same argument very briefly on p.235, though not for the purpose of giving an alternative construction of algebraically closed fields.) In view of this result, we can find a set S whose cardinality is *greater* than that of any algebraic extension of k. Without loss of generality, let us assume that S contains k. We may now let P be the set consisting of all field extensions of k whose underlying sets are contained in S, and partially order P by the relation ''is an extension of''. It is easy to verify that this set is inductive; thus P has a maximal element, K.

It remains to prove that K is in fact algebraically closed. Assuming the contrary, let L be a proper algebraic extension of K. This will still have cardinality less than that of S. Moreover, since K has cardinality less than that of S, the cardinality of $S - K$ is the same as that of S; in particular, we can map the set $L - K$ bijectively to some subset $U \subseteq S - K$. Using the identity map on K and the above bijection on $L - K$, we get a bijection between L and the subset $K \cup U$ of S, which is the identity on K. We can use this bijection to transport the field structure of L onto $K \cup U$. Since L is a proper algebraic extension of K, $K \cup U$ with this field structure will be a proper algebraic extension of K whose underlying set lies in S. But this contradicts the maximality of K, so we conclude that K is indeed algebraically closed.

(In setting up the above proof, we might have thought it would be enough to take S to have cardinality greater than *or equal to* that of any algebraic extension of k; e.g., denumerable if k was so. You should look carefully at the above proof, and see why this would not have worked.)

I will make some remarks at the end of this section on the idea behind Lang's proof of the existence of an algebraically closed extension of k, in particular, the first part, where it is shown that one can simultaneously adjoin roots to all polynomials of positive degree over k.

P.233, Proposition V.2.7 [~]: In the preceding paragraphs, Lang has proved more than he records in the proposition, and in fact, he will use the stronger result repeatedly, so you should note it: *If* $E = k(\alpha)$, *where* α *is algebraic over* k, *with minimal polynomial* p, *and* $\sigma: k \to L$ *is an embedding of* k *in a field* L, *then the extensions of* σ *to* E *correspond bijectively to the distinct roots of* p^{σ} *in* L; *namely, for each such root* β, *there is a unique extension of* σ *carrying* α *to* β. (In particular, L does not have to be algebraically closed.)

In many applications of this result, L will also be an extension of k, and σ the inclusion map. Note that in that case, the conclusion simplifies to: *Homomorphisms* $E \to L$ *over* k *correspond bijectively to roots of* p *in* L; *namely, for each such root* β, *there is a unique such homomorphism carrying* α *to* β.

P.234, Corollary V.2.9 [=]: In the proof (though not in the statement) of the corollary, the ''identity mapping on k'' should be ''the inclusion $k \to E'$''.

P.235, second and third paragraphs [<]: These are confusingly written, and can be skipped.

P.235, end of §V.2 [>]:

''Compactness'' results (sketch, optional).

Let us look again at Lang's proof that a field k admits an extension in which every polynomial $f \in k[X]$ of positive degree has a root. He begins by showing that for every *finite* set of such polynomials f, one can get an extension of k in which these all have roots. He then adjoins to k an infinite family of indeterminates X_f, one for each f, and forms the ideal I generated by the elements $f(X_f)$. To complete the proof by dividing out by a maximal ideal containing I, we need to know that I is a *proper* ideal of the polynomial ring.

To get this, Lang observes that if it contained 1, then the subideal generated by some finite family $f_1(X_{f_1}), \dots, f_n(X_{f_n})$ would also. But one can construct a homomorphism from our polynomial ring into a field in which $f_1, \dots, f_n$ all have roots, which has $f_1(X_{f_1}), \dots, f_n(X_{f_n})$ in its kernel; hence the ideal generated by these elements does not contain 1, giving the result needed for the proof.

In general, an argument showing that if a set S of conditions expressible in an appropriate language has the property that every *finite* subset of S can be satisfied simultaneously, then all of S can be satisfied simultaneously, is called by logicians a *compactness* theorem; and they have compactness theorems for various sorts of languages. To see the reason for the term, let L be a language, and M the class of all ''models'' of L (structures in which we can interpret the statements of the language. E.g., in our case, the language would be one with symbols for addition, multiplication, and the elements of k, and with variable-symbols, quantifiers, etc., sufficient to express the condition that a structure forms *a field containing* k, and such that we can write, for any $a_0, \dots, a_{n-1} \in k$, the condition

$$(\exists\, x)\ \ x^n + a_{n-1}x^{n-1} + \dots + a_0 = 0.$$

The class M of models would consist of all sets given with operations and elements to be interpreted as the addition, multiplication, and the scalars from k.) Let us topologize M with the least topology such that for each proposition P in the language L, the set of models satisfying P is a closed set. Then the statement that any set S of propositions all of whose finite subsets admit models itself admits a model is equivalent to saying that M is *compact* under this topology! (Well, almost. The trouble is that, as we have described it, M is not a *set*. This can be gotten around in various ways, e.g., by limiting M to consist of models whose underlying set is a subset of a fixed set of large enough cardinality. In any case, logicians generally do not think of the topological background when they use the term ''compactness theorem''.)

Since beginning graduate students are not assumed to have seen any general compactness theorems, Lang goes through an ad hoc ''compactness argument'' for this case.

We remark that if a language is *too* powerful, compactness will generally fail. For instance, suppose a language can express the condition (*P*): ''the set X is finite'', as well as the more elementary conditions (*Q1*): ''X has more than 1 element'', (*Q2*): ''X has more than 2 elements'', (*Q3*): ''X has more than 3 elements'', etc.. Then we see that any finite subset of $\{(P), (Q1), (Q2), (Q3), \dots\}$ can be satisfied simultaneously, but not all of them together; so we do not have compactness for this language. Note that if a language allows one to quantify over functions, then (assuming other more elementary properties), one can say ''there exists a function from X to X which is one-to-one but not onto'', i.e., that X is infinite; and negating this, one can say that X is finite. Hence languages for which one expects compactness results should in general not allow quantification over functions. You can verify, similarly, that compactness fails for a ring-theoretic language in which one can say, ''x is transcendental'' along with more elementary statements about elements of a field.

Re §V.3. Splitting fields and normal extensions.

P.236, line 3 [>]: Lang has defined the **splitting field** of a polynomial. Here are some examples. A splitting field for the polynomial $X^2 - 2$ over $\mathbf{Q}$ is $\mathbf{Q}(\sqrt{2})$, since over this field we have the factorization $X^2 - 2 = (X - \sqrt{2})(X + \sqrt{2})$. On the other hand, a splitting field of $X^4 - 2$ cannot be constructed simply by adjoining to $\mathbf{Q}$ a fourth root of 2. For instance, if we let $\alpha \in \mathbf{C}$ denote the positive real fourth root of 2, then over $\mathbf{Q}(\alpha)$, the polynomial $X^4 - 2$ factors as $(X - \alpha)(X + \alpha)(X^2 + \alpha^2)$, with a quadratic factor that is irreducible over the real numbers, hence over its subfield $\mathbf{Q}(\alpha)$; so $\mathbf{Q}(\alpha)$ is not a splitting field of this polynomial. On the other hand, you should verify that a splitting field of $X^4 - 2$ *is* given by $\mathbf{Q}(\alpha, i)$. (Remember to check the condition that the field in question is generated by roots of the polynomial!)

Another example: Let $f(X) = X^5 - 1$. If we write $\zeta_5 = e^{2\pi i/5}$, then the roots of f are 1, ζ_5, ${\zeta_5}^2$, ${\zeta_5}^3$, ${\zeta_5}^4$, and it follows immediately that $\mathbf{Q}(\zeta_5)$ is a splitting field for f.

Digression: The existence of splitting fields (or of algebraic closures, obtained in the preceding section) gives an interesting alternative proof of the italicized statement on p.192, that if elements $t_1, \dots, t_n$ are algebraically independent, then the symmetric polynomials $s_1, \dots, s_n$ in these elements are also algebraically independent, at least in the case where the base-ring A is an integral domain. To see this, let K be the field of fractions of A, let $K(t_1, \dots, t_n)$ be the field of fractions of $A[t_1, \dots, t_n]$, and let $K[t_1, \dots, t_n]$ be the subring thereof generated by the indicated elements.

Because $t_1, \dots, t_n$ are assumed algebraically independent, $K[t_1, \dots, t_n]$ can be identified with the polynomial ring in n indeterminates over K; so if L is any field extension of K and $\alpha_1, \dots, \alpha_n$ any elements of L, we get a substitution homomorphism $K[t_1, \dots, t_n] \to L$ that fixes elements of K and sends each t_i to α_i. This homomorphism will send the symmetric polynomials in the ts to the symmetric polynomials in the αs, that is, to the coefficients of the polynomial $\prod_i (X - \alpha_i)$; so if the symmetric polynomials in the ts satisfied a nontrivial polynomial relation $f(s_1, \dots, s_n) = 0$, the coefficients c_i of every polynomial of the form $\prod (X - \alpha_i)$ for any n-tuple of elements α_i of any extension field of K would also satisfy $f(c_1, \dots, c_n) = 0$.

But by the existence of splitting fields, every monic polynomial over an extension field of K does have the form $\prod (X - \alpha_i)$ in some larger extension field; and since we can take any n-tuple of elements as the coefficients of $X^{n-1}, \dots, X^0$ in a monic polynomial, this would mean that every n-tuple of elements c_i satisfied $f(c_1, \dots, c_n) = 0$. Taking an L containing algebraically independent elements $c_1, \dots, c_n$, we get an immediate contradiction.

P.237, statement and proof of Theorem V.3.3 [=]: A clearer statement of **NOR 1** is:

NOR 1. *For every embedding σ of K in k^{a} over k, one has $K^\sigma = K$.*

The idea of the equivalent conditions **NOR 1 - NOR 3** is to avoid the deficiency we saw above in the field gotten by adjoining to $\mathbf{Q}$ a single fourth root of 2, namely, that of containing some, but not all roots of an irreducible polynomial.

In the proof of the theorem, the first paragraph actually shows that **NOR 1** $\Rightarrow$ **NOR 3** $\Rightarrow$ **NOR 2**, and the second paragraph shows that **NOR 2** $\Rightarrow$ **NOR 1**, so the third paragraph, showing **NOR 3** $\Rightarrow$ **NOR 1**, is unnecessary.

Re §V.4. Separable extension.

P.239, beginning of §V.4 [<]: Suppose F is a field, E an algebraic extension of F, and L an algebraically closed field containing F. Then it is natural to study the embeddings of E in L over F.

As a slightly more general situation, we might assume that L does not necessarily contain F itself, but an isomorphic copy $\sigma(F)$, where $\sigma\colon F \to L$ is an embedding of fields, in which case we should understand ''embeddings of E in L *over* F'' to mean embeddings that make commuting diagrams with the inclusion of F in E, and the map $\sigma\colon F \to L$. In a sense, this is no more general than the first situation, since we can reduce it to that situation by mentally ''identifying'' elements of F with their images in L, and thus regarding L as an extension of F. But, as I discussed at the beginning of §V.2, we will want to avoid formally making such identifications, so that we will be able to use these results in the case where F *is* a subfield of L but σ is *not* the inclusion map. Nevertheless, in reading the material on this page (up to Theorem V.4.1), you might first think of L and L' as algebraically closed *extensions* of F, and σ and τ as the inclusions, then reread those paragraphs without making that assumption.

The point Lang is leading up to here is that in looking at the embeddings of E over F into an algebraically closed extension L of F, one gets the same number of maps, independent of the L one uses.

P.240, paragraph after Corollary V.4.2 [<]: Skip this paragraph. (Lang will make this all clear two sections later.)

P.240, last two sentences [~]: These refer to ''multiple roots''. This means multiple roots in an *algebraic closure* of k (or in any field in which the polynomial in question splits). These polynomials do not, in general, have roots in k itself. The same applies to later uses of the phrase.

Incidentally, note that for any finite extension E of a field k within k^{a}, we have

$$\#\text{ of automorphisms of } E \text{ over } k \;\le\; \#\text{ of embeddings of } E \text{ in } k^{\mathrm{a}} \;\le\; [E:k].$$

By the results of the last section, E is *normal* over k if and only if one has equality at the first $\le$ sign, while by the definition just given, it is *separable* over k if and only if one has equality at the second. So normality plus separability means that it has as many automorphisms as one could hope for. We will see in the next chapter that this pair of conditions defines an important class of extensions.

P.241, proof of Theorem V.4.3 [>]: Some further basic facts that you should be able to verify at this point, but which Lang doesn't state formally, are recorded in

Lemma V.4.c1. *Let k be a field and f an irreducible polynomial over k. Let α be a root of f in an extension field of k, and let E be a splitting field of f over k. Then the following conditions are equivalent.*

(i) f *is separable.*

(ii) $k(\alpha)$ *is separable over* k.

(iii) E *is separable over* k. $\square$

P.241, last display [=]: An ''a'' appears in place of what should be the second ''α''.

P.243, line below first displayed equation [=]: For ''the same field'' read ''the above field''.

P.243, last line [=]: After ''polynomials'' add ''over an algebraic closure of E''.

The idea of this argument is striking: different subextensions of $k(\alpha)$ can be told apart by the different minimal polynomials satisfied by the single element α over each of them.

P.244, end of §V.4 [>]:

Notes on the proof of the Theorem of the Primitive Element.

The preceding proof becomes more understandable if we extract from it the following general result:

Lemma V.4.c2. *A vector space V over an infinite field k cannot be written as the union of finitely many proper subspaces $V_1, \ldots, V_n$.*

Proof. Let x be a point of V chosen to *minimize* the number of V_i containing it. We claim that x lies in none of the V_i; this will establish the lemma. Assuming the contrary, let us choose $y \in V$ lying outside of some V_i that contains x, and consider the family of points $x + cy$ $(c \in k)$, which we can picture as a line in V. Any two distinct points of this line clearly span the subspace spanned by x and y; so any subspace of V that does not contain *both* x and y contains *at most* one point of the line. Applying this to each of the finitely many spaces V_i, and noting that because k is infinite, there are infinitely many points on our line, we conclude that some point of our line lies in only those V_i which contain *both* x and y. In particular, it lies in fewer of the V_i than x does, contradicting our choice of x, and completing the proof of the lemma. □

Now in the proof of the Theorem of the Primitive Element, since each subextension of E is in particular a k-subspace, we can replace the second and third paragraphs of the proof by a quick application of the above lemma.

This lemma likewise simplifies the proof of the last statement of the theorem, that every finite separable extension has a primitive element: In the second paragraph of p.244, for each pair of distinct embeddings σ, τ of E in k^{a} let us write $V_{\sigma,\tau}$ for $\{\alpha \in E \mid \sigma\alpha = \tau\alpha\}$. Then by the above lemma, we can find $\alpha \in E$ lying in none of these subspaces. Hence α has degree equal to at least the number of distinct σ, which, since E is separable, is $[E:k]$; so α generates E over k.

Re §V.5. Finite fields.

P.245, long paragraph starting in middle of page [>]: The fact that the roots of $X^{p^n} - X$ form a subring of $\mathbf{F}_p^{\mathrm{a}}$, which Lang proves by detailed calculation, is a consequence of an easily verified and important general fact:

Lemma V.5.c1. *Let $f, g: A \to B$ be homomorphisms of rings. Then $\{a \in A \mid f(a) = g(a)\}$ forms a subring $C \subseteq A$.*

The same result is true with ''ring'' replaced by ''A-algebra'', ''A-module'', or ''monoid''.

Moreover, in the ring and monoid cases, if an element a of C has an inverse $a^{-1} \in A$, then $a^{-1} \in C$. Hence in these cases, if A is a field, respectively a group, so is C. □

The subring, subgroup, etc., characterized above is often called the *difference kernel* of f and g, because in the case of modules (or abelian groups), it can be described as the *kernel* of the *difference*-map $f - g$. It is called by categorists the *equalizer* of f and g. If $A = B$, the equalizer of the identity endomorphism and an endomorphism g is called the *fixed subring* (or *fixed subgroup* etc.) of g.

Now the set of roots of $X^{p^n} - X$ in $\mathbf{F}_p^{\mathrm{a}}$ can be described as the fixed subring of the nth power of the Frobenius endomorphism; hence by the above lemma it forms a subfield.

Re §V.6. Inseparable extensions.

P.247, first sentence of §V.6 [>]: Before ''can be omitted'', insert ''except for Proposition 6.1''.

(Also, I disagree with the opinion that this section is more technical, and by implication less interesting than the rest of this chapter.)

P.247, statement of Proposition V.6.1 [=]: Note that ''root of f'' means ''root of f in k^{a}''. This usage occurs throughout this section; thus, you will find frequent references to the multiplicities of the roots of an irreducible polynomial. Since the roots are in k^{a}, while the polynomials are irreducible over k, there is no contradiction.

Even if one is only interested in separable extensions, this proposition is important, since it shows that every algebraic extension of a field of characteristic 0 – or more generally, every algebraic extension whose degree is not divisible by the characteristic – *is* separable. So Lang's invitation, earlier on this page, to omit §V.6 should have referred only to the material after this proposition.

P.248, end of proof of Proposition V.6.1 [>]: This proof can be made clearer, beginning with the second paragraph on p.248 (''Consider ...''). In the second line, for ''a polynomial'', read ''the polynomial f'''. The result Lang then deduces should in fact be made a lemma: ''*If an irreducible polynomial f is inseparable* (*has multiple roots in its splitting field*), *then the characteristic of the base-field is a prime. An irreducible polynomial f over a field of prime characteristic p is inseparable if and only if it has the form $g(X^p)$.*'' Now consider a general irreducible polynomial f over a field k of characteristic p. Let μ be the largest integer such that we can write $f(X) = h(X^{p^\mu})$. Any factorization of h over our base-field would yield a factorization of f, so h is also irreducible; hence, by the above lemma and our choice of μ, h is separable. Thus, its factorization over k^{a} has the form $h(X) = \Pi\,(X-\beta_i)$, where the β_i are distinct and separable. For each i let us choose $\alpha_i \in k^{\mathrm{a}}$ satisfying ${\alpha_i}^{p^\mu} = \beta_i$. We can now factor f:

$$f(X) \;=\; h(X^{p^\mu}) \;=\; \Pi\,(X^{p^\mu} - \beta_i) \;=\; \Pi\,(X^{p^\mu} - {\alpha_i}^{p^\mu}) \;=\; \Pi\,(X-\alpha_i)^{p^\mu}.$$

This shows that every root has multiplicity p^μ, and has separable p^μth power, completing the proof of the proposition.

We see that, loosely speaking, inseparability is a phenomenon of pth roots in characteristic p. We gave an example of this phenomenon earlier, in our comments on §V.1 (second paragraph of *Hints of what is to come*).

P.250, third line [=]: The argument beginning 8 lines from the bottom of the preceding page (at the word ''Then'') and ending here can be done more simply: by Theorem VI.6.1, every root of the minimal polynomial of α has multiplicity p^μ for some μ; but since the separable degree is 1, there is only one root. So the polynomial has the form $(X-\alpha)^{p^\mu}$. But this equals $X^{p^\mu} - \alpha^{p^\mu}$, hence particular, $\alpha^{p^\mu} \in k$.

P.250, first line of proof of Proposition V.6.5 [=]: ''tower theorem'' should be ''tower property''.

P.250, Corollary V.6.8 and its proof [~]: The same argument shows, more generally, that if K is normal over k, and K_0 is any subextension which is carried into itself by all automorphisms of K, then K_0 is also normal over k.

P.251, line below the diagram [=]: For ''the corollaries of Proposition 6.1'' read ''Theorem 4.1 and Corollary 6.4''.

P.251, first line of Corollary V.6.10 [=]: For ''$x^p,\ x \in E$'' read ''x^p *where* $x \in E$''.

P.251, first line of Proof of Corollary V.6.10 [=]: Note that from the assumption $E = E^p k$, we get $E = E^p k = (E^p k)^p k = E^{p^2} k^p k = E^{p^2} k$ (the last step because k ''absorbs'' its subfield k^p). Iterating this argument, we get $E = E^{p^m} k$ for all nonnegative m, a fact which will be used two lines later.

P.251, Proposition V.6.11, second line [=]: Rather than turning to the next chapter for the definition of *fixed field*, you can recall Lemma V.5.c1, this Companion, comment to p.245, and note that in that situation if A is a field then C will be a subfield.

P.251, third-from-last line [=]: After ''By definition'' add ''of K^G''.

P.252, first full paragraph [<]: This result can be gotten without reducing to the finite case. Given $\alpha \in K$, let $\alpha_1, \dots, \alpha_r$ be the distinct roots of the minimal polynomial of α, let $f(X) = \Pi\,(X-\alpha_i)$, and observe that the coefficients of f will be G-invariant.

P.252, Proof of Corollary V.6.12 [~]: Suppose E is an algebraic extension of k. How do we apply Proposition V.6.11, as Lang asks us to? The point is to first take the normal closure F of E in k^{a}. Then that proposition says that F is a separable extension of a purely inseparable extension; but being perfect, k has no purely inseparable extension, so F is separable, so its subextension E is also separable, as claimed.

Having finished this section, let us note a couple of interesting

Examples involving inseparable extensions.

Both our examples will be based on the following observation:

Let k_0 be a field, p be any prime, and $k_0(u_1, u_2)$ an extension gotten by adjoining independent transcendentals u_1 and u_2 to k_0. Then we see that $u_1{}^p$ and $u_2{}^p$ will also be independent transcendentals over k_0 (since any polynomial relation that they might satisfy could be written as a relation satisfied by u_1 and u_2, and no such relations hold). Moreover, we claim that

(c31) $$[k_0(u_1, u_2) : k_0(u_1, u_2{}^p)] \;=\; [k_0(u_1, u_2{}^p) : k_0(u_1{}^p, u_2{}^p)] \;=\; p,$$
$$\text{whence} \qquad [k_0(u_1, u_2) : k_0(u_1{}^p, u_2{}^p)] \;=\; p^2.$$

Indeed, the degrees in the top line are bounded above by p because each u_i satisfies the polynomial $X^p - u_i{}^p$ over the subfield in question. To also get a lower bound of p, suppose that $1, u_i, \ldots, u_i{}^{p-1}$ were linearly dependent over that field. Taking a linear dependence relation among them, and clearing denominators, we would get an equation $f = 0$ holding in $k_0[u_1, u_2]$. By looking at the residues of the exponents modulo p, we can see that f would be a nonzero polynomial, a contradiction.

The above observations are valid without restriction on the characteristic of k, but in our applications, k_0 will be a field of characteristic p. Our first application will be the construction of

A nonprimitive finite field extension.

Theorem V.4.6 gave a necessary and sufficient condition for a finite extension of a field k to be primitive (i.e., generated over k by a single element), and it was shown that every separable extension satisfied this condition. Moreover, the easiest examples of inseparable extensions, gotten by adjoining to a field a single inseparable element, are ipso facto also primitive.

But suppose we adjoin to a field k of characteristic p pth roots of two elements of k, getting a field $k(a^{1/p}, b^{1/p})$. The Frobenius endomorphism (the pth-power map) carries the elements of k and the two elements $a^{1/p}$ and $b^{1/p}$ into k, hence it carries the whole field they generate into k. This shows that every element of that extension satisfies a polynomial of the form $X^p - c$ $(c \in k)$. But $k(a^{1/p}, b^{1/p})$ will generally have degree p^2 over k, so that a primitive element would have to have degree p^2. Hence there is no primitive element.

To complete the example, we need an explicit case of the situation we have just claimed holds ''in general'', i.e., two elements a and b such that $[k(a^{1/p}, b^{1/p}) : k] = p^2$. This is gotten from (c31) by taking any field k_0 of characteristic p, and letting $k = k_0(u_1{}^p, u_2{}^p)$, $a = u_1{}^p$, and $b = u_2{}^p$.

According to the Theorem of the Primitive Element, an extension of the sort we have constructed will have infinitely many subextensions. Can we describe an infinite family of such subextensions explicitly? Yes. In general, given an extension $k(a^{1/p}, b^{1/p})$ of degree p^2 in characteristic p, I claim that the subextensions $k(a^{1/p} + c\, b^{1/p})$ as c ranges over the distinct elements of k are all distinct. For each of them has degree $\leq p$ over k, but we see by linear algebra that the compositum of the extensions we get using any two values of c contains both $a^{1/p}$ and $b^{1/p}$, hence is the whole field $k(a^{1/p}, b^{1/p})$. Since this has degree p^2 over k, no two of these extensions of degree p can be equal. (But are we sure that whenever we have such an extension $k(a^{1/p}, b^{1/p})$, there must be infinitely many distinct elements $c \in k$? Yes; because on a finite field k of characteristic p, the pth-power map, being a one-to-one map of a finite set into itself, is a permutation, so in the finite case, there can be no elements a, b which do not have pth roots in k.)

What I have done above is essentially Lang's Exercise V:L24.

Our second construction will involve

Mixing separable and inseparable extensions.

If E is an algebraic extension of a field k, and F its largest *separable* subextension, then Proposition V.6.6 shows that E is *purely inseparable* over F. Thus, E can be obtained by a separable extension followed by a purely inseparable extension. There is also a largest purely *inseparable* subextension of E, and we might expect that E would be *separable* over this field, giving a decomposition of the reverse sort.

Lang shows in Proposition V.6.11 that this is so if E is normal over k, but warns, in the preceding paragraph, that it is not true in general. Let me show how it can fail. The idea will be to adjoin to k all the roots of a separable polynomial, then adjoin a pth root of just *one* of these. Then the inseparable stage of the extension will have been done in an ''asymmetric'' way, the effect of which *cannot* be obtained by starting with a purely inseparable extension of the base field k.

So let us again look at the situation of (c31) with k_0 a field of characteristic p, and recall that $u_1{}^p$ and $u_2{}^p$ are algebraically independent over k_0. Let us call these two elements t_1 and t_2, let s_1, s_2 denote the *elementary symmetric functions* in this pair of elements, and let us write

$$k \;=\; k_0(s_1, s_2).$$

Thus, k is the field of symmetric rational functions in t_1 and t_2, and $k_0(t_1, t_2)$ is generated over k by the roots t_1, t_2 of the quadratic polynomial $X^2 - s_1X + s_2$; equivalently, by *one* of the roots, since each can be obtained from the other by subtracting from s_1. We now look at the larger fields gotten by adjoining the pth roots of one or both of t_1 and t_2:

$$E \;=\; k_0(u_1, t_2) \;=\; k(u_1), \qquad L \;=\; k_0(u_1, u_2) \;=\; k(u_1, u_2).$$

Suppose that an element $x \in E$ satisfies $x^p \in k$. From the fact that k consists of elements symmetric in t_1 and t_2, one can see that x will be symmetric in u_1 and u_2. But since x lies in E, u_2 will appear everywhere with exponents divisible by p, hence by symmetry the same will be true of u_1. Thus, x will be a symmetric rational function in t_1 and t_2, hence will lie in k. This means that the largest purely inseparable subextension of E over k is k itself. Since E is not separable over this, we have the desired example.

This construction is delicate. To see this, suppose that, more generally, we adjoin to any field k of characteristic p the roots of an arbitrary quadratic polynomial $X^2 - bX + c$, calling these t_1 and t_2, and then adjoin $t_1{}^{1/p}$, and call the resulting field E. In $k(t_1, t_2)$, we have $t_1 + t_2 = b$, hence in k^{a} we have $t_1{}^{1/p} + t_2{}^{1/p} = b^{1/p}$. (In k^{a} every element has a pth root, and in fields of characteristic p, pth roots are unique when they exist, so the Frobenius endomorphism $x \mapsto x^p$ of k^{a} is bijective, hence is an automorphism. The inverse of this automorphism may be written $x \mapsto x^{1/p}$, and the above equation follows.) Now suppose b and c happened to be chosen so that b is a pth power in k, but c is not; i.e., $b^{1/p} \in k$ but $c^{1/p} \notin k$. When we adjoin $t_1{}^{1/p}$ to $k(t_1, t_2)$, we can get from this the element $t_2{}^{1/p} = b^{1/p} - t_1{}^{1/p}$, and multiplying $t_1{}^{1/p}$ and $t_2{}^{1/p}$ together, we get $c^{1/p} \in E$. Thus, there exists an element $c^{1/p} \in E - k$ purely inseparable over k, which is what the preceding example was aimed at avoiding. In this new example, one finds in fact that E is separable over $k(c^{1/p})$, i.e., this E *is* a separable extension of a purely inseparable extension of k. Essentially the same thing happens if we reverse the hypotheses, and assume $c^{1/p} \in k$ but $b^{1/p} \notin k$.

Chapter VI. Galois Theory.

Re §VI.1. Galois extensions.

P.261, last sentence (continued on next page), ''It is a field because ... '' [=]: This is a special case of Lemma V.5.c1 (this Companion, comment to Lang's p.245).

P.262, line 7 [>]: Lang here mentions several notations for the Galois group of K over k. In the next few pages he will write $G(K/k)$, but later on, he will also use the other notations, rather at random.

P.262, first paragraph of the proof of Theorem VI.1.2 [~]: This repeats the argument used in Proposition V.6.11, and you can easily get the result being proved (that $k = K^G$) from that proposition. Lang doesn't call on the proposition because he made §V.6 optional.

Incidentally, where Lang organizes this proof as a demonstration that the correspondence $H \mapsto K^H$ is one-to-one and onto, I think it is more natural to show that that map, and the map $F \mapsto \mathrm{Gal}(K/F)$ are inverse to one another. Indeed, that is what his proof of one-one-ness and onto-ness reduces to.

P.263, first paragraph [>]: Note that the statement ''H belongs to F'' does *not* mean ''H is the Galois group of F''. Indeed, as an extreme case, the Galois group of our whole extension K is $G(K/F)$, but the subgroup thereof ''belonging to K'' is $\{1\}$.

Incidentally, one also uses the words ''belong to'' in the converse sense, saying ''F is the subfield belonging to the subgroup H'' to mean F is the fixed field of H. (Lang uses it this way at the bottom of p.268.)

P.263, second line of proof of Corollary VI.1.6 [=]: Delete the first occurrence of the word ''finite''.

P.264, first sentence of proof of Theorem VI.1.8 [=]: An easier way to say this would be ''Let $\alpha \in K$ and choose $\sigma_1, \ldots, \sigma_k \in G$ so that $\sigma_1 \alpha, \ldots, \sigma_k \alpha$ are the distinct members of the orbit of α under G''. Since the orbit of an element of a G-set is permuted by every member of G, we can skip the second sentence of the proof and jump directly to the third. (Alternatively, we could drop mention of $\sigma_1, \ldots, \sigma_k$ entirely, and begin the proof by writing $\alpha_1, \ldots, \alpha_k$ for the orbit of α under G, and using these elements in defining f.)

Four lines after the display in this proof, ''splits in'' should be ''splits into''.

P.264, proof of the Theorem of Artin (VI.1.8) [>]: The first part of this proof again repeats the argument used in Proposition V.6.11.

Note that Theorem VI.1.2 showed that under certain hypotheses, every intermediate field ''belongs to'' some group, while the Theorem of Artin shows that under other hypotheses, every subgroup belongs to some subfield. The limitations of the two results are different: Theorem VI.1.2 requires that we be given a separable normal extension, but needs no assumption of finiteness; the Theorem of Artin requires no separability or normality assumptions, (since it does not start with a field *extension*, but with a single field and a group of automorphisms, and *gives* separability and normality in the conclusion), but finiteness is essential. When we put these results together we get the first assertion of Theorem VI.1.1, a beautiful and powerful bijective correspondence, holding for *finite separable normal* extensions.

P.264, proof of Corollary VI.1.9 [>]: This completes the proof of the first sentence of Theorem VI.1.1, which is what is generally called the Fundamental Theorem of Galois Theory. (The second sentence, to be proved in Theorem VI.1.10, is among the many important auxiliary facts about this correspondence.)

Let us note that although the formulation of that theorem looks about the same from one text to another, there are subtle differences, resulting from the fact that different texts define ''Galois extension'' in different ways. Some authors, like Lang, define this to mean a separable normal algebraic extension; others to mean an algebraic extension E/k such that k is the fixed field of $G(E/k)$. Moreover, one may define ''separable'' either, like Lang, in terms of the number of embeddings in an algebraic closure, or by the condition that the minimal polynomial of every element of the extension have no multiple roots.

Whatever definitions a text gives, it generally supplies the additional results that establish that these definitions are equivalent to the other formulations just mentioned. Thus, Lang's Theorem V.4.3 does this for the concept of separable extension, and Theorem VI.1.8 (Artin's Theorem) for that of normal extension.

So if one takes the conjunction of the Fundamental Theorem of Galois Theory with these supplementary results, this is effectively the same from text to text, even though the statements of the Theorem alone may not be equivalent.

P.264, Remark [>]: I develop the bijective correspondence between subextensions of an *infinite* Galois extension and appropriate subgroups of its Galois group in the comments to Lang's §VI.14 below (though that section is not covered in Math 250A). You can read those comments now (without even looking at that section of Lang), although in the *examples* given, you will have to temporarily take on faith a couple of plausible facts about particular finite extensions, which follow from results in §§VI.3 and VI.8.

Returning to p.264, suppose F is a finite separable extension of a field k, but *not* normal. We know that F can be embedded in a Galois extension K, and, being a field between k and K, it will then ''belong to'' some subgroup $H < G = G(K/k)$. Since the inclusion relations among subfields of K containing k are the reverse of those among subgroups of G, we can say that *subextensions* of F correspond to *subgroups of* G *that contain* H. Thus, even though the extension field F is not Galois, its subextensions can be described using Galois theory. What about the automorphism group of F over k – can you figure out how to characterize it in terms of G and H? Since it is a finite group of automorphisms of F, the Theorem of Artin says that it is the Galois group of F over a subfield k'. By the Fundamental Theorem of Galois Theory, this k' is the subfield of K belonging to some subgroup of G. Can you characterize that subgroup in terms of H? (If you can't answer these questions now, you might think about them again after finishing this section.)

P.265, slightly below middle, ''There is no avoiding the contravariance ... '' [~]: I am not sure what Lang means by this. I think the point he is making is that if we want to write conjugation by λ in exponential notation, and have it compose correctly for an operation written to the right of its argument, we must define it to take σ to $\lambda^{-1}\sigma\lambda$ rather than $\lambda\sigma\lambda^{-1}$; but then it will take $G(\lambda K/\lambda k)$ to $G(K/k)$, rather than vice versa.

Something group-theorists sometimes do to avoid this sort of problem is to use a left-superscript notation; e.g., writing $\lambda\sigma\lambda^{-1} = {}^{\lambda}\sigma$. (Lang does this on p.69, second and third lines of first Example.)

P.266, statement and proof of Theorem VI.1.12, and diagram on next page [=]: When we adjoin the elements of F to K, not all the automorphisms of K over k may extend to automorphisms of the resulting field over F. Those that do extend form a subgroup, which we know must correspond to some field between k and K. Which field? This theorem says it is $K \cap F$!

If you stop to think about it, you will see that if F has some elements in common with K other than the elements of k, this will indeed prevent automorphisms of K that move those elements from extending to automorphisms of KF over F. The theorem tells us that conversely, F must contain enough of K to lead to whatever failures of automorphisms to extend occur.

The result is sometimes called the *Theorem on Natural Irrationalities*. I think that ''irrationalities'' is an old way of saying ''elements not in the base field'' (generalizing the term from the case where the base field is $\mathbf{Q}$), and the theorem says that F contains those ''irrationalities'' from K that it ''naturally'' ought to.

In the proof, Lang jumps right to proving the last sentence of the theorem. Of the two earlier assertions, the statement that KF is Galois over F holds because normality and separability go over to this extension, while the statement that K is Galois over $K \cap F$ is immediate, since $K \cap F$ is an intermediate field of a Galois extension (Theorem VI.1.2).

P.266, last parenthetical paragraph [~]: Lang has proved the last assertion of the theorem only in the case where K is finite over k (or KF is finite over F), and he now says that topological methods are used in the infinite case. In fact, the infinite case can be deduced from the finite case. Let σ be an automorphism of K over $K \cap F$ which we wish to extend to an automorphism σ_{KF} of KF over F. We note that K is the union of its subextensions K_0 that are finite and normal over k, that each of these subextensions is carried into itself by σ, and that by the finite case, which Lang has proved, the restriction of σ to each such K_0 extends in a unique way to an automorphism $\sigma_{K_0 F}$ of $K_0 F$ over F. Since every element $x \in KF$ may be obtained by field operations from finitely many elements of K and of F, every such x lies in a subextension $K_0 F$, and we may thus define

(c32) $$\sigma_{KF}(x) = \sigma_{K_0 F}(x), \qquad \text{where } x \in K_0 F \text{ and } K_0 \text{ is finite over } k.$$

We claim that this definition is independent of our choice of a K_0 for each x, and does indeed give an automorphism. For note that if K_0 and K_1 are two finite normal subextensions of K, then $K_0 K_1$ is also one. The map $\sigma_{K_0 K_1 F}$ will act on K_0 and K_1 via σ, and will carry $K_0 F$ and $K_1 F$ into themselves by normality, hence by the uniqueness of $\sigma_{K_0 F}$ and $\sigma_{K_1 F}$, it will agree with those extensions on their domains. Hence if an element x lies in both $K_0 F$ and $K_1 F$, we see that $\sigma_{K_0 F}(x) = \sigma_{K_0 K_1 F}(x) = \sigma_{K_1 F}(x)$, so that (c32) gives a well-defined map. That this map is an automorphism can be deduced similarly: Given elements $x \in K_0 F$, $y \in K_1 F$, these will both lie in $K_0 K_1 F$, on which σ_{KF} acts by the automorphism $\sigma_{K_0 K_1 F}$; thus, we get the required equations $\sigma_{KF}(x+y) = \sigma_{KF}(x) + \sigma_{KF}(y)$, etc..

P.267, diagram, and following line [>]: Lang says, ''It is suggestive to think of opposite sides of a parallelogram being equal''. Note, however, that this only applies under the particular hypotheses of the preceding theorem, which essentially says that *if* one of the lower edges of such a parallelogram is a Galois extension, then the parallel upper edge is also Galois, with the same Galois group. The example marked **Warning** that Lang gives after the next corollary shows that even the equality of degrees between opposite sides can fail if the lower edge is not assumed normal. Moreover, in that example, each of the upper edges of the parallelogram is an extension of degree 2, and every quadratic extension is normal. (Once we adjoin one root ρ of a quadratic polynomial f to a field, the other root is a root of the linear polynomial $f(X)/(X-\rho)$, hence lies in the extension, so the extension is normal.) So in this diagram, the *top* two edges are normal, yet opposite edges are still not in any sense ''equal''. Incidentally, in that ''Warning'', all occurrences of the symbol ''E'' should be changed to ''K'' (six occurrences in the last six lines) for conformity with the preceding material.

P.269, last phrase of §VI.1 [=]: ''Notation should be functorial with respect to the ideas!'' was a slogan of Lang's. It is not clear to me exactly how he was using the word ''functorial''. One way of looking at it is that when one has a functor between categories, not only does each object of the domain category yield an object of the codomain category, but morphisms yield morphisms, in a way that elegantly matches the structure of the domain category, and that Lang was implying that notation should similarly elegantly match ideas.

Note, also, that in pre-category-theoretic usage, when one had a construction of a new object $F(X)$ from an old object X, and one (eventually) recognized that every map $a: X \to Y$ induced a map $F(X) \to F(Y)$, one generally denoted that map by an ad hoc symbol such as a^*; but with the realization that the constructions on objects and on morphisms form one entity, a ''functor'', we now write $F(a): F(X) \to F(Y)$. Lang may have been saying that notation should reflect ideas in something like the way that functorial notation does. His use of k^{a}, k^{s} and k^{ab} rather than arbitrary bars and tildas does this.

P.269, end of §VI.1 [>]:

Galois connections.

The Fundamental Theorem of Galois Theory is an example of a very general type of mathematical situation: One has two sets S and T, and a certain binary *relation* on them, i.e., a subset $R \subseteq S \times T$. (In our case, S is an extension field K of a field k, T is the set $G(K/k)$, and R is the ''fixed-element'' relation, $\{(x, \sigma) \in K \times G(K/k) \mid \sigma(x) = x\}$.) For every subset $X \subseteq S$ one defines $X^* \subseteq T$ to be $\{t \in T \mid (\forall s \in X)\ (s, t) \in R\}$, and likewise for every subset $Y \subseteq T$ one defines $Y^* \subseteq S$ to be $\{s \in S \mid$

$(\forall t \in Y)\ (s,t) \in R\}$. These operators constitute what is called a *Galois connection* between S and T, the general concept being named after the important special case we have been studying.

There are a number of basic results about Galois connections that are trivial to prove, but worth knowing, to keep from having to rediscover them in case after case! In particular, (i) $X \subseteq X' \Rightarrow X* \supseteq X'*$; (ii) the operators $**$ from subsets of S to subsets of S, and from subsets of T to subsets of T, are *closure operators* on S and T respectively (as defined in this Companion, comments on §III.1); and (iii) the operators $*$ give a bijective inclusion-reversing correspondence between closed subsets of S and closed subsets of T with respect to these operators. (These and other results are proved, and numerous examples discussed, in §5.5 of the 245 notes.)

What *cannot* be gotten from the general theory of Galois connections are *characterizations* of the closed sets in a given case – for example, the result that when K is a finite separable normal extension of k and R is the ''fixed element'' relation between elements of K and elements of $G(K/k)$, then the closed subsets of K are precisely all intermediate fields, and the closed subsets of $G(K/k)$ are all its subgroups. For each Galois connection, this problem must be investigated using the methods of the area of mathematics in which it arises.

Re §VI.2. Examples and applications.

P.269, beginning of §VI.2 [<]: There is a fundamental class of examples which Lang has given us the background for, but which he fails to note here! Since it is of a slightly different sort from the examples he gives, let us note it before you start this section.

We have seen that if K is a field of characteristic p, then the pth power operation is an endomorphism of K, called the Frobenius endomorphism, whose fixed field is precisely the subfield of p elements (since $X^p - X$ can have at most p roots). If K is a finite field, say $\mathbf{F}_{p^n}$ (in the notation of §V.5), then this endomorphism will necessarily be an automorphism, and be of finite order; hence by the Theorem of Artin, the group it generates will be the full automorphism group of $\mathbf{F}_{p^n}$ over its fixed field $\mathbf{F}_p$. Thus, each field $\mathbf{F}_{p^n}$ is a Galois extension of $\mathbf{F}_p$, with cyclic Galois group, which must have order $[\mathbf{F}_{p^n} : \mathbf{F}_p] = n$.

One can also see (in more than one way) that for every m dividing n, the subgroup generated by the mth power of the Frobenius map corresponds to the subfield $\mathbf{F}_{p^m} \subseteq \mathbf{F}_{p^n}$.

P.269, first paragraph of §VI.2 [>]: There are some general observations about this situation that we should make here for use in the examples to follow. First, f will be irreducible if and only if its Galois group G is *transitive* on the set of roots. Indeed, if we have a factor g of f in $k[X]$, then g is fixed under G, hence the set of its roots must be carried into itself by this group; and, conversely, given a subset A of the set of roots of f that is carried into itself by G, we see that $\prod_{\alpha \in A} (X - \alpha)$ is a factor of f in $K[X]$ which is invariant under G, hence lies in $k[X]$, i.e., a factor of f in the latter ring. Thus, f has no proper factor if and only if G acts transitively on the set of its roots. Finally, note that if we adjoin to k one of the roots of f, say α_1, then the subgroup of G belonging to the resulting subextension $k(\alpha_1) \subseteq K$ will be the group of elements of G fixing α_1. (Can you state a precise proof of this?) If, with Lang, we regard G as a subgroup of S_n, i.e., as acting formally on $\{1, \ldots, n\}$ rather than on $\{\alpha_1, \ldots, \alpha_n\}$, then this becomes the isotropy subgroup of the element 1.

P.269, third from last line [=]: For ''This comes from completing the square and'' read ''This may be seen by taking the irreducible polynomial of any generator of K over k, and completing the square, or equivalently,''.

We remark that in the preceding paragraph, Lang introduced the assumption char $k \neq 2$ because for the characteristic 2 case, a different method is needed, not taught in ''elementary school'', which we will see at the end of §VI.7. For similar reasons he will assume char $k \neq 2, 3$ in Example 2.

P.270, line 10 [=]: ''In the second case $k(\alpha)$ is not normal over k'': This is because the corresponding group is the stabilizer of $1 \in \{1, 2, 3\}$ in S_3, which is not a normal subgroup. In contrast, the stabilizer of 1 in A_3 is normal therein; equivalently, it fixes all members of the orbit of 1 under that group.

P.270, line below display ''Δ = ...'' [>]: Lang writes ''The set of σ in G which leave δ fixed is precisely the set of even permutations''. You can probably convince yourself that this is true, but let us develop formally the corresponding result for equations of arbitrary degree. Note that D and δ are the polynomials that Lang introduced on pp.192-193 (though he is writing D here for what he called D_f there, and Δ here for what he called D there, i.e., the result of substituting the coefficients of a particular polynomial for the arguments of D. I will instead distinguish these explicitly as $D(s_1, \ldots, s_n)$ and $D(a_1, \ldots, a_n)$.)

Lemma VI.2.c1. *Let n be a positive integer, and let $D \in \mathbf{Z}[X_1, \ldots, X_n]$ be the* discriminant polynomial, *that is, the unique polynomial such that in the polynomial ring $\mathbf{Z}[t_1, \ldots, t_n]$, if one writes s_i for the ith elementary symmetric polynomial in the t_i, one has $\prod_{i<j}(t_i - t_j)^2 = D(s_1, \ldots, s_n)$.*

Then for any field k not of characteristic 2 and any separable monic polynomial $f = X^n + a_{n-1}X^{n-1} + \ldots + a_0$ of degree n over k, if we denote by E the splitting field of f, and by $\alpha_1, \ldots, \alpha_n$ the roots of f in E, then the following three conditions are equivalent:

(a) *$G(E/k)$, regarded as a group of permutations of $\alpha_1, \ldots, \alpha_n$, lies in the alternating group A_n (i.e., acts by even permutations on these roots).*

(b) *The element $\delta(\alpha) = \prod_{i<j}(\alpha_i - \alpha_j)$ of E lies in k.*

(c) *The element $D(a_{n-1}, \ldots, a_0) \in k$ is a square in k.*

Hence, the fixed field of the group of elements of $G(E/k)$ which act by even permutations on the roots of f is $k(\delta(\alpha)) = k(D(a_{n-1}, \ldots, a_0)^{1/2})$.

Proof. It is easy to see that the automorphism of $\mathbf{Z}[t_1, \ldots, t_n]$ that acts on the subscripts of the t's by a permutation of the form $(i, i+1)$ takes $\delta(t) = \prod_{i<j}(t_i - t_j)$ to $-\delta(t)$. Since an *odd* (respectively *even*) permutation can be characterized as a permutation which can be written as the product of an odd (respectively even) number of permutations of the form $(i, i+1)$, we see that the action of any odd permutation on $\mathbf{Z}[t_1, \ldots, t_n]$ will send $\delta(t)$ to its negative, and any even permutation will send it to itself.

Now if an automorphism θ of E acts by a permutation π_θ on the roots $\alpha_1, \ldots, \alpha_n$ of f, then it is easy to see that the map $\mathbf{Z}[t_1, \ldots, t_n] \to E$ carrying t_i to α_i makes a commuting square with the automorphism θ of the field, and the automorphism of that polynomial ring acting by π_θ on the subscripts of the indeterminates. It follows that θ will send $\delta(\alpha)$ to itself if π_θ is even, and to its negative if π_θ is odd. Hence, if all members of $G(E/k)$ act by even permutations on $\alpha_1, \ldots, \alpha_n$, then $\delta(\alpha)$ is fixed under that group, hence belongs to k, proving (a)$\Rightarrow$(b); and conversely, if this element belongs to k, then all members of the Galois group fix it, hence act by even permutations; so (b)$\Rightarrow$(a).

It is also clear that $D(a_{n-1}, \ldots, a_0) = \delta(\alpha)^2$; hence (b)$\Rightarrow$(c). Since the equation $X^2 - D(a_{n-1}, \ldots, a_0) = 0$ can have at most two roots in E, its only roots are $\pm\,\delta(\alpha)$, so (c)$\Rightarrow$(b).

The last assertion is easily deduced from the above results. □

Thus Lang is essentially applying the above lemma with $n = 3$, using the formula for $\Delta = D(0, a, b)$ obtained in Chapter IV (p.193). Another important application of the above lemma will be made in Exercise VI.3:1.

Let us record one other important fact about the discriminant. You should find it straightforward to prove, by the same approach as above,

Lemma VI.2.c2. *Let n and D be as in the preceding lemma, let k be* any *field, and let $f = X^n + a_{n-1}X^{n-1} + \ldots + a_0$ be any monic polynomial of degree n over k. Then f is inseparable (has at least one multiple root in k^{a}) if and only if $D(a_{n-1}, \ldots, a_0) = 0$.* □

P.271, the two lines after second display (''The field ...'') [=]: Replace these two lines by, ''Since α is real, $i \notin \mathbf{Q}(\alpha)$. But i satisfies the polynomial $X^2 + 1$, hence it has''. And the next sentence should begin with ''So''.

P.271, next-to-last line of text, ''This gives us the structure of G'' [>]: It is easy to see that the four roots of $X^4 - 2$ in the complex plane form the vertices of a square, and that the group Lang shows to be the Galois group of this equation acts on these roots as the full symmetry group of that square! You should be able to determine the field extension corresponding to each subgroup in the diagram at the bottom of that page. In that diagram, the symbols $\langle\ ...\ \rangle$ should more logically be $\{\ ...\ \}$, since Lang is listing all the elements of each subgroup, rather than just a generating set. (A wide class of extensions having this same Galois group, or subgroups thereof, is examined in Exercises VI.2:1 and VI.2:2.)

P.272, the two lines after second display [>]: Since F is the fixed field of a group of order $n!$, K has degree $n!$ over it, hence it has degree at least $n!$ over $k(s_1, \ldots, s_n)$, with strict inequality if this field is strictly smaller that F. But since K is generated over that subfield by the roots of the polynomial f which has coefficients in that subfield, the degree of K over that subfield is also at most $n!$. Hence that subfield is exactly F.

Here, as in the proof of Artin's Theorem (p.264), and of the Primitive Element Theorem (p.243, argument at the bottom of the page), the fact that a subfield contains the coefficients of one or more irreducible polynomials helps us ''pin down'' which subfield it is.

P.272, last paragraph of Example 4 [=]: The theorem referred to in the first sentence of the paragraph is one we have not had, but the consequence that there exists a Galois extension E of $\mathbf{Q}$ such that $G(E/\mathbf{Q}) \cong S_n$ is proved for the special case of n prime in Example 6 on the next page, and for the general case in Exercise VI:L14, p.322. Concerning the open question Lang then discusses, note that if G is a finite group, we can embed it in a symmetric group S_n. If we let E be an extension of $\mathbf{Q}$ having S_n as Galois group, as above, and F the intermediate field corresponding to the given group $G \subseteq S_n$, then $G(E/F) \cong G$. Thus, every finite group is the Galois group of a Galois extension E/F such that E and F are both finite extensions of $\mathbf{Q}$. The open question is whether one can always do this with $F = \mathbf{Q}$. The final two sentences, about ''specializing parameters'' is a technical point which you may ignore.

P.272, beginning of Example 5 [~]: Lang mentions that $\mathbf{R}$ is ''an ordered field''. This means a field with a total ordering that satisfies certain of the properties of the ordering of the real numbers. Here all you need to know is that you will be using in the proof some of the familiar properties of the order relation on the reals. The general definition of an ordered field may be found in §XI.1.

Toward the bottom of the page, Lang shows that an arbitrary complex number $a + bi$ has a square root. To see how one could discover the formula he gives, consider the problem of solving for c and d the equation $a + bi = (c + di)^2$. Expanding the right-hand side and equating real and imaginary parts, one gets a value for $c^2 - d^2$ and a value for cd. Squaring the latter and multiplying by -1, we see that we have values for the sum and product of c^2 and $-d^2$, i.e., the elementary symmetric functions in these two real numbers. Hence we can find c^2 and $-d^2$ by the quadratic formula. (Obviously, c^2 will be the larger root, and $-d^2$ the smaller.) The fact that the reals are an ordered field is used to show that the number whose square root one has to take in using this formula is nonnegative, and that the resulting values for c^2 and $-d^2$ are nonnegative and nonpositive respectively, hence that the values for c^2 and d^2 are both nonnegative. If one arbitrarily chooses a square root of the value obtained for c^2, and calls this c, then, depending on one's choice of a square root d of d^2, one finds that $(c + di)^2$ comes out as $a \pm bi$. In particular, one choice gives the desired value, $a + bi$.

P.273, beginning of second paragraph [=]: Change ''We now see'' to ''Recall''.

P.273, end of Example 5 [>]: The argument Lang has given may be summarized as follows. The fact that every finite group has a p-Sylow subgroup tells us that a group whose order is not a power of p has a subgroup whose index is not divisible by p; while we also know that every nontrivial group whose order *is* a power of p has a subgroup of index exactly p. Using the Fundamental Theorem of Galois

Theory, these statements translate to say that if a field has an extension of degree not a power of p, it has a nontrivial extension of degree relatively prime to p, while if it has a nontrivial extension of degree a power of p, it has an extension of degree exactly p. Now taking $p = 2$, we observe that if the complex numbers had an extension of degree not a power of 2, this would also be an extension of $\mathbf{R}$ of degree not a power of 2, hence by the above observations, $\mathbf{R}$ would have a nontrivial extension of degree not divisible by 2, and the minimal polynomial of some element thereof would be a real polynomial of odd degree having no real root, a contradiction. On the other hand, if $\mathbf{C}$ had a proper extension of degree a power of 2, then by the above observations, it would have an extension of degree exactly 2. But using the quadratic formula, and the existence of square roots of all complex numbers, we know that every quadratic polynomial over $\mathbf{C}$ splits, so this is also impossible. Thus $\mathbf{C}$ has *no* nontrivial finite extensions.

P.273, Example 6, sentence beginning ''When-'' [=]: For ''this group'' read ''a subgroup of S_n'', and for ''this representation of G'' read ''the Galois group of f, regarded as a group of permutations of the roots of f''.

P.274, through end of §VI.2 [<]: You can skip these last three examples, which assume material in later chapters, or not in the book.

If you do read Example 7, trusting the truth of the result Lang quotes in italics, you might see Exercise VI.2:4 for a clarification of an argument Lang gives later in that example regarding the Galois group of $X^2 - X - 1$ (mod 2).

If you have seen the topological concept of covering space, and had an undergraduate course in complex analysis, you can read in place of (or before) Example 9 the following more elementary notes on the relation between

Galois theory and Riemann surfaces (sketch, optional).

The Riemann sphere $\mathbf{C} \cup \{\infty\}$ (which Lang calls the Gauss sphere) is defined to consist of the complex plane with one ''point at infinity'' adjoined. (A picture can be found in most complex analysis texts.) Every member of $\mathbf{C}(t)$ can be regarded as a meromorphic function on this object; e.g., t is a function with a ''pole'' at infinity, but holomorphic everywhere else; $1/(t-1)$ has a pole at 1, but is holomorphic everywhere else (including ∞).

Now consider the extension field $\mathbf{C}(t^{1/2})$. At each point of the Riemann sphere other than 0 and ∞, there are two values that can be given to $t^{1/2}$, hence a general rational function of this element will be double-valued. To get around this problem, one regards elements of this field as meromorphic functions, not on the Riemann sphere, but on a ''two-sheeted branched covering of that sphere, with branch-points at 0 and ∞''. This means, essentially, a space X with a map into the Riemann sphere, such that if we remove those two points from the sphere, and their inverse images from X, what we are left with is a two-to-one covering map. (Certain conditions are also assumed on how the map behaves at the branch points, which I won't go into.)

This particular branched covering of the Riemann sphere can itself be identified with the Riemann sphere itself, by using the parameter $t^{1/2}$ in place of t; i.e., it is the Riemann sphere ''wrapped twice around itself''. But if we take an algebraic extension of $\mathbf{C}(t)$ that cannot be generated over $\mathbf{C}$ by a single element, for example, $\mathbf{C}(t,\ (t^3 - t)^{1/2})$, then the corresponding branched covering, when ''unwrapped'', is not a sphere, but a torus, or a more complicated surface. The objects one gets in this way are called ''compact Riemann surfaces''. Given a branched covering of the Riemann sphere by such a surface (or of one such surface by another), it can be shown that the Galois group of the corresponding field extension acts by self-homeomorphisms of the covering space, that respect the covering map, and hence permute the ''sheets''. In other words, the Galois group becomes the covering group of our topological covering!

(The Berkeley courses that consider Riemann surfaces are Math 205, 241 and 255.)

(Lang's Example 9 skips past the finite extensions of $\mathbf{C}(t)$ discussed above, and looks at what is, in effect, the *compositum* (union) of all such finite extensions whose branch points lie in a given finite set $\{P_1, \ldots, P_{n+1}\}$. (On line 5 he has P_{n-1} instead of P_{n+1} by mistake.) As discussed in the note to §VI.14 below, the Galois group of such an infinite extension can be described in terms the Galois groups of the finite extensions of which it is a union.)

Re §VI.3. Roots of unity.

P.277, first paragraph [>]: In any field, a *primitive* n*th root of unity* means an element of multiplicative order n. A field may or may not have such elements; for example, $\mathbf{Q}$ has a primitive square root of unity but no primitive nth root of unity for higher n, though its extension $\mathbf{C}$ has primitive nth roots of unity for all n. $\mathbf{F}_7$ has primitive square, cube and sixth roots of unity (hence so does any field of characteristic 7); it has no primitive twelfth root of unity, but $\mathbf{F}_{7^2}$ does. (Do you see why?) No field of characteristic p has a primitive pth root of unity (do you see why not?), hence no such field has a primitive nth root of unity for any n divisible by p.

In this paragraph, Lang assumes n not divisible by the characteristic of k, and works in k^{a} (so that what he is looking for could be described as nth roots of unity *over* k, rather than *in* k). In this context, he shows that the group $\boldsymbol{\mu}_n$ is cyclic of order n (do you see why its order will be exactly n?), hence, a generator of this group will indeed be a primitive nth root of unity.

P.277, line after third display, ''Hence $\sigma\zeta = \zeta^i$'' [>]: Note that since every nth root of unity has the form ζ^r for some r, σ will act on these roots by $\sigma(\zeta^r) = (\sigma\zeta)^r = \zeta^{ir} = (\zeta^r)^i$. In particular, the same integer i for which $\sigma\zeta = \zeta^i$ also satisfies $\sigma\zeta' = \zeta'^{\,i}$ for all primitive nth roots of unity ζ'. This is why Lang can write $i = i(\sigma)$, rather than $i = i(\sigma, \zeta)$.

P.277, fifth line from bottom [=]: The function $\varphi(n)$, called the *Euler phi function*, was defined on p.94.

P.277, third line from bottom [=]: Here is where Lang starts using variant notations for the Galois group (mentioned on p.262, end of second paragraph). Also, on this and the next line there are three occurrences of κ that should be k as in the beginning of this line.

P.278, statement of Theorem VI.3.1 [>]: To see what is needed to prove this result, consider any field k, and any positive integer n not divisible by the characteristic of k. Suppose we write $X^n - 1$ as a product of irreducibles $g_1 \ldots g_r$ in $k[X]$, and look at the roots of each factor in a splitting field of $X^n - 1$. Since all the roots of an irreducible polynomial over k satisfy precisely the same equations over k, we can say that for each i, all roots of g_i have the same multiplicative order. In particular, the $\varphi(n)$ primitive nth roots of unity, which are the roots of order n, will constitute the set of all roots of some subfamily of these factors. From these observations, we can see that Theorem VI.3.1 is equivalent to the statement that when $k = \mathbf{Q}$, all the primitive nth roots of unity are roots of a *single* irreducible factor of $X^n - 1$. (Lang will make this translation at the bottom of p.279; but in fact, it is implicit in his proof.)

P.278, proof of Theorem VI.3.1 [~]: The idea of this proof can be best understood if we ask ourselves how much of it works if we replace $\mathbf{Q}$ by a more general field k. So let k be any field of characteristic not dividing n, and let us look at the primitive nth roots of unity in a splitting field of $X^n - 1$. These roots are permuted transitively by the operations of raising to various positive integer powers m relatively prime to n. Let us ask which of these values of m have the property that exponentiation by m preserves (permutes) the set of roots of every irreducible factor of $X^n - 1$. If you examine the method of proof of the theorem, you will see that it shows that if k is the field of fractions of some UFD A which can be mapped homomorphically to a field $\mathbf{Z}/p\mathbf{Z}$ (p relatively prime to n), then exponentiation by p permutes the roots of each irreducible factor of $X^n - 1$. The key facts from which this follows are: (i) exponentiation by p is an automorphism of any algebraic extension of $\mathbf{Z}/p\mathbf{Z}$, hence permutes the roots of any polynomial over $\mathbf{Z}/p\mathbf{Z}$, (ii) the property that raising to the pth power sends roots of an irreducible polynomial g_i to roots of a polynomial g_j can be written as a divisibility relation, $g_i(X) \mid g_j(X^p)$, and so is preserved on going to homomorphic images (even if the irreducibility of g_i is lost), and (iii) $X^n - 1$ is separable over $\mathbf{Z}/p\mathbf{Z}$, so the distinct factors g_j of $X^n - 1$ over $k[X]$, when taken in $A[X]$ and then mapped into $\mathbf{Z}/p\mathbf{Z}[X]$, remain relatively prime, hence the various $g_j(X^p)$ are also relatively prime, hence $g_i(X)$ cannot divide more than one $g_j(X^p)$.

Returning to the special case $k=\mathbf{Q}$, note that the UFD $\mathbf{Z}$ has homomorphisms into all fields $\mathbf{Z}/p\mathbf{Z}$. Hence exponentiation by every positive integer prime to n permutes the roots of each factor of X^n-1 over $\mathbf{Q}$. But every primitive nth root of unity can be obtained from any other by raising it to a positive integer power; thus, all the primitive nth roots of unity are roots of a single factor of X^n-1.

P.278, proof of Corollary VI.3.2 [<]: In preparation for this proof, let us make some easy observations on the function φ. If m and n are relatively prime, we know that $Z_{mn} \cong Z_m \times Z_n$ (Proposition I.4.3(v)) and that the cyclic generators of this group correspond to the pairs (x,y) where x and y are generators of the respective factors. It follows that

$$\text{For } m \text{ and } n \text{ relatively prime, } \varphi(mn) = \varphi(m)\varphi(n).$$

A number-theoretic function with this property is called *multiplicative*. (Note that the word does *not* mean, as one might imagine, that the above equation holds for *all* m and n!) To evaluate such a function on all positive integers, it suffices to know its values on powers of primes. As noted on p.96, $\varphi(p^r) = (p-1)p^{r-1}$.

P.278, line following next-to-last display (''Our assertion ...'') [=]: Add at the end of the line, ''and Corollary 1.13''.

P.279, line following second display [=]: Note that the periods (i.e., orders) of nth roots of unity will be the divisors of n, and that if d is a divisor of n, then the dth roots of unity are a subset of the nth roots of unity. Hence the product defining $\Phi_d(X)$ is indeed a function of d alone, and not of n, even though we are looking at it as the product of a subset of the factors of X^n-1.

P.279, 5th line from bottom [=]: ''Cyclotomic'' means ''circle-cutting'', the idea being that the nth roots of unity divide the unit circle of the complex plane into n equal arcs.

P.280, assertions **1-5** [<]: You should be able to get assertions 3 and 4 as immediate consequences of the definition of Φ_n. Assertion 2 and the second part of assertion 1 are applications of assertion 4, while the first part of assertion 1 has already been proved.

Assertion 5, as well as some subsequent discussion on this page, involve *Möbius inversion*. You will not be responsible for this (or any of the material from assertion 5 through the first paragraph on p.281), but the method is a general and important one, so I will sketch the idea.

Suppose one has a function α (not assumed a homomorphism) from the positive integers to an abelian group A (which we will write *multiplicatively*), and another such function β related to α by the formula $\beta(n) = \prod_{d|n} \alpha(d)$, the product being taken over all positive integer divisors of n, including 1 and n. It is not hard to see that the values of α can be recovered from those of β. Indeed, assuming inductively that one has found $\alpha(m)$ for all proper divisors m of n, one can get $\alpha(n)$ by dividing $\beta(n)$ by the product of those values. So in fact, $\alpha(n)$ is expressible, using the group operation of A, in terms of the values of $\beta(d)$ as d ranges over the divisors of n. The exact formula can be shown to be $\alpha(n) = \prod_{de=n} \beta(d)^{\mu(e)}$, where μ, the *Möbius function*, is defined by the formula Lang gives on this page. He applies this result with $A = Q(X)^*$, $\alpha(n) = \Phi_n(X)$ and $\beta(n) = X^n - 1$. (Lang refers to this technique, in the last paragraph of the page, as ''inverting by the general formalism of convolutions''. This is the same ''convolution'' as in the definition of the multiplication operation of a monoid algebra, but to explain the connection would require a lengthy digression, in which we would turn each of the functions α and β above into a formal power series in infinitely many indeterminates; so I will not go into it here. Some of the ideas are developed in Exercises II:L12, p.116, and V:L21-22, p.254.)

P.281, 2nd display [=]: To make this calculation clearer, I would move the fraction to between the n-term sum and the final 0. In any case, the argument is that the sum equals that fraction by the general formula for $(X^n-1)/(X-1)$, and the fraction is zero because its numerator is zero. If ζ is a *primitive* nth root of unity, another way to see this same result is by noting that the sum is the sum of the roots of X^n-1, hence equals the coefficient of X^{n-1} in that polynomial, which is zero. In the field of complex numbers, we can also see the result geometrically: The powers of ζ are the vertices of a regular n-gon, which has center of gravity at the origin, hence 0 is the average of the vertices as complex numbers, hence they sum to 0. Finally, if ζ is a not necessarily primitive nth root of unity, it will still be a primitive dth root for some $d|n$, where $d \neq 1$ by assumption. We see that the sum in question will be the sum of the distinct powers of ζ each repeated n/d times. Hence by the ''primitive root'' case, this sum is zero.

P.281, statement of Theorem VI.3.3 [=]: Here, as in the preceding discussion, p is an *odd* prime (a prime $\neq 2$). Note that from the way S is defined, if we apply an automorphism that sends ζ to ζ^{μ}, then if μ is a square mod p, S will be unchanged, while if μ is a non-square, S will change to $-S$. Thus S^2 is invariant under the whole Galois group, and taking our base-field to be $\mathbf{Q}$, we conclude that S^2 must be a rational number, whose square root belongs to the indicated cyclotomic extension. The point of the theorem is to establish which rational number it is. The answer, we are told, is $\pm p$, with sign depending on whether -1 is a square mod p.

Let me sketch how one might come up with this answer and its proof. Squaring the formula defining S, we get the expression in the first display on the next page. What would be an enlightening way of grouping the terms? We might try to group them by the sum $\mu+\nu$. This looks nice in terms of the exponent of ζ; but it is not clear how the quadratic symbol will behave as μ and ν vary with constant sum. If we tried to group them according to μ or ν alone, we would just be undoing our squaring operation. Suppose, then, that we group them by the *ratio* of μ to ν, calling that ratio λ, so that ν becomes $\lambda\mu$. Then the ''numerator'' in the quadratic symbol becomes $\lambda\mu^2$, and since the quadratic symbol is unchanged under multiplying by a square, it will be the same in each group of terms. This gives us the right-hand side of the top line of Lang's second display; except that he has recycled the symbol ν with a new meaning, where I have introduced a new symbol λ; and I would write double summation, the outer Σ being over λ, which ranges over all nonzero residues modulo p. The quadratic symbol would be written occurring right after that Σ, and the inner summation would be $\Sigma_{\mu}\, \zeta^{\mu(\lambda+1)}$, where μ also ranges over nonzero residues modulo p.

Now evaluation the inner sums is complicated slightly by the exclusion of the residue $\mu = 0$. Note that the sum of the excluded terms is just the sum of all the quadratic symbols, multiplied by $\zeta^0 = 1$. Since the quadratic symbol takes on the values $+1$ and -1 equally often, the sum of those excluded terms is 0, and we can put them in. We then see that the summation $\Sigma_{\mu}\, \zeta^{\mu(\lambda+1)}$ is *almost* always just the sum of all powers of ζ, which is 0 (since it is unchanged under multiplying by ζ). The one exception is when $\lambda = 1$, where we are summing $\zeta^0 = 1$ p times, and the sum is p. Bringing in the coefficient $\left(\frac{\lambda}{p}\right)$, that term comes to $\left(\frac{-1}{p}\right) p$, as claimed.

(Incidentally, note that ''μ'' in this proof is just an index representing an integer modulo p; not to be confused with the group ''$\boldsymbol{\mu}_n$'' of p.277 nor the Möbius function $\mu(n)$ of p.280.)

Re §VI.4. Linear independence of characters.

P.283, top paragraph [>]: Given any field K and any set S, we know that the set K^S of K-valued functions on S becomes a K-vector-space if addition and scalar multiplication are defined componentwise. Lang is here regarding characters $G \to K$ as forming a *subset* of this vector space, and considering whether this subset is linearly dependent.

In the simplest case, where $G = \mathbf{N}$, the monoid of nonnegative integers, characters $G \to K$ are sequences of powers $(1, \alpha, \alpha^2, \dots)$. You may have been aware that distinct such sequences are always linearly independent. (In fact, the formula for the Vandermonde determinant shows that given n such sequences, the length-n vectors formed from their first n terms are linearly independent. Though this is not essential to the present reading, it is a fact of general importance in mathematics, and I will review it below.) The result of this section shows that the same is true for homomorphisms $G \to K$ for an arbitrary monoid G.

The Vandermonde determinant (sketch, optional).

The Vandermonde determinant is defined as the polynomial in n indeterminates $X_1, \ldots, X_n$ given by the determinant of the $n \times n$ matrix whose jth column is composed of the $j-1$st powers of the indeterminates:

(c33) $$\det((X_i^{\,j-1}))_{i,j=1,\ldots,n},$$

(or more generally, the ring element resulting from the substitution of some n values x_i for the indeterminates X_i in this expression). I claim that (c33) equals

(c34) $$\Pi_{i>j}\,(X_i - X_j).$$

The first step is to show that (c33) is divisible by all the factors shown in (c34). Intuitively, it is divisible by each term $X_i - X_j$ because if one sets X_i and X_j equal, then (c33) becomes the determinant of a matrix whose ith and jth rows are equal, and thus goes to zero. To formalize the concept of setting X_i and X_j equal, let us map the polynomial ring in $X_1, \ldots, X_n$ to its factor-ring by the ideal generated by any difference $X_i - X_j$. Then the image of (c33) in that factor-ring is the determinant of a matrix over that ring with two equal rows, and is thus zero; hence (c33) must lie in the ideal we have factored out. Since this is true for all $i < j$, and since the distinct terms $X_i - X_j$ $(i > j)$ are relatively prime, (c33) is divisible by (c34). Moreover, (c33) and (c34) are of the same degree, $0+1+\ldots+(n-1)$, hence they can only differ by a scalar factor. To determine this scalar, we look at the coefficient of the monomial

$$X_2\, X_3^2\, \ldots\, X_{n-1}^{n-2}\, X_n^{n-1}$$

in (c33) and in (c34). This monomial arises exactly once in the expansion of each (as the diagonal product in (c33), and as the product of the first term of each factor in (c34)), hence it has coefficient 1 in each, hence the scalar in question is 1.

Now note that whenever we map our polynomial ring into a field so as to take $X_1, \ldots, X_n$ to *distinct* values, $x_1, \ldots, x_n$, (c34) goes to a nonzero field element; hence the matrix whose determinant is taken in (c33) must have linearly independent rows. This shows, as we claimed earlier, that any n distinct sequences of the form $(1, \alpha, \alpha^2, \ldots, \alpha^{n-1})$ in a field are linearly independent.

P.284, before the line preceding Corollary VI.4.2 [<]: At this point read Corollary VI.5.4, p.286, which for some reason was put in the wrong section. (That corollary is a slight generalization of the observation you have just read on p.284 of Lang)

P.284, Corollary VI.4.2 [~]: This is the result we obtained in a more precise form using the Vandermonde determinant.

P.284, end of §VI.4 [>]. Here are two tangential notes related to the material of this section.

Further observations on linear independence of characters (optional).

The theorem on Linear Independence of Characters was proved in preparation for an application that will be made in the next section; but let us note that it can also be used (and is, in some textbooks) in place of one of the steps in Lang's development of the Fundamental Theorem of Galois Theory; namely, to show that the order of the automorphism group of a finite extension is at most the degree of that extension, or more generally, that the number of embeddings over k of an extension field E in an algebraic closure k^{a} is at most $[E:k]$ (p.240, top line). This was done in Lang by counting roots of polynomials; the alternative argument is as follows: An embedding of E in k^{a} is, in particular, a k^{a}-valued character on the multiplicative monoid of E; so the set of such embeddings is linearly independent. In addition to being a character, such an embedding is a k-linear map $E \to k^{\mathrm{a}}$, and such a map is determined by what it does on elements of a k-basis of E. Thus, the k^{a}-dimension of the space of all k-linear maps $E \to k^{\mathrm{a}}$ is $[E:k]$. Combining these two observations, we see that this number is an upper bound on the number of embeddings of E in k^{a} over k.

Turning to the way linear independence of characters is proved, it is interesting that essentially the same the argument is used in getting several important mathematical results: one assumes the existence of a linear relation among certain elements, takes such a relation that involves the fewest terms; then transforms it in some way, and shows that by taking a linear combination of the original relation and the transformed relation, one can eliminate one term and get a smaller relation, leading to a contradiction (unless that smaller relation is trivial, in which case some desired condition is satisfied). You have doubtless seen this method used to prove that eigenvectors of a linear map corresponding to distinct eigenvalues are linearly independent. Actually, as noted in the comments to p.562 in this Companion, linear independence of eigenvectors and linear independence of characters are essentially the same result. But I sketch below another application of the above technique which I do not see how to regard as giving ‘‘the same result’’. In fact, it gives an alternative proof (also used in some texts) of the step in the development of the Fundamental Theorem of Galois Theory complementary to the one mentioned above; namely, that if k is the fixed field of a finite group G of automorphisms of a field E, then $[E:k] \leq (G:1)$ (cf. Lang, Theorem VI.1.8, p.264).

To get this, let $G = \{g_1, \dots, g_n\}$. Then we want to show that any $r > n$ elements $\alpha_1, \dots, \alpha_r \in E$ are k-linearly dependent. Consider the r n-tuples $(g_1(\alpha_1), \dots, g_n(\alpha_1))$, $\dots$, $(g_1(\alpha_r), \dots, g_n(\alpha_r)) \in E^n$. As $r > n$ elements of an n-dimensional E-vector space, these must be E-linearly dependent. Let $c_1, \dots, c_r \in E$ be coefficients in a nontrivial linear relation that they satisfy.

This condition on $c_1, \dots, c_r$ says that for every $g \in G$, $c_1 g(\alpha_1) + \dots + c_r g(\alpha_r) = 0$. Applying g^{-1} to this relation, and writing $g^{-1} = h$, we see that this is equivalent to saying that for every $h \in G$, $h(c_1)\alpha_1 + \dots + h(c_r)\alpha_r = 0$. Clearly, the set of r-tuples $(c_1, \dots, c_r)$ satisfying this condition is an E-subspace of E^r closed under the action of G.

Let us take a nonzero element of this subspace which minimizes the number of nonzero entries, normalized to have some entry equal to 1. By applying any member of G, and subtracting the result from the given element, we get an element with fewer nonzero entries; hence by the minimality assumption, this must be zero. This means that the given element is invariant under the action of G, i.e., that it has entries in k. Hence it gives the coefficients of a linear dependence relation among $\alpha_1, \dots, \alpha_r$ over k, as desired.

Another result that can be gotten by the same approach is that of Exercise XIII:L32, p.550 (though Lang's hint leads to a different proof, using the preceding exercise on that page. For still another method, *using* linear independence of characters, see VI.5:1.)

On characters into division rings (sketch, optional).

If G is a monoid and E a division ring, we can again regard the set of maps $G \to E$ as a vector space E^G over E (let us say a *left* vector space, since there are two choices); and we may again ask whether the subset consisting of the *characters* $G \to E$ is E-linearly independent.

Let us look at a simple case. In the division ring $\mathbf{H}$ of quaternions (described in the last comment to §V.1, above), the elements ± 1, $\pm i$, $\pm j$, $\pm k$ are fourth roots of unity (all but the first pair being primitive fourth roots), hence each of these elements determines a character $Z_4 \to \mathbf{H}$, sending the generator $[1]$ of Z_4 to the chosen element. This gives us *eight* characters in the *four*-dimensional $\mathbf{H}$-vector-space $\mathbf{H}^{Z_4}$. Hence these characters must be linearly dependent! (There are in fact more than eight such characters: an easy calculation shows that a quarternion $a + bi + cj + dk$ is a primitive fourth root of unity if and only if $a = 0$ and $b^2 + c^2 + d^2 = 1$, so the 4-dimensional space $\mathbf{H}^{Z_4}$ contains *uncountably* many characters.)

Surprisingly, however, there is a version of linear independence of characters that works in division rings. In the above example, one can show that all the primitive fourth roots of unity in $\mathbf{H}$ are *conjugate* to one another. (For example, note that $(1+i)j = k(1+i)$, so $j = (1+i)^{-1}k(1+i)$.) And conjugacy turns out to be the key idea. Since the set of functions f from a monoid G to a division ring E can be multiplied on the right as well as on the left by elements of E, one can conjugate such a function f by an element $d\in E$, getting a function $x\mapsto d^{-1}f(x)d$, and it is easy to see that if f is a character, so are all its conjugates. Now if $f_1, ..., f_n$ are characters $G\to E$, and we take any left linear relation $a_1 f_1 + ... + a_n f_n = 0$ $(a_i\in E)$ that they satisfy, and collect terms according to the conjugacy classes of the characters f_i, one can prove that the sum over *each* conjugacy class will be zero; that is, every case of linear dependence among characters results from linear dependence *within* conjugacy classes. It remains to determine when a family of mutually conjugate characters, $d_1^{-1} f d_1, \ ..., \ d_n^{-1} f d_n$, is linearly dependent. This turns out to happen if and only if the elements $d_1, \ ..., \ d_n\in E$ are left linearly dependent over the *centralizer* C in E of the submonoid $f(G)$. Indeed, given a C-linear dependence relation among the d's, you should not find it hard to write down an E-linear dependence relation among the above mutually conjugate characters. The proof of the converse result, and of the fact that linear dependence relations among arbitrary families of characters all come from such relations among families of conjugate characters, are obtained by a fairly natural adaptation of the argument given by Lang in the commutative case; you might try to discover the details yourself. Note that when E is a field, every conjugacy class consists of a single character, so in this case the above assertion that all linear dependence relations are consequences of such relations within conjugacy classes is equivalent to the linear independence of characters.

Re §VI.5. The norm and trace.

P.284, beginning of §VI.5 [<]: Suppose α is an element of a finite separable extension E of a field k, and we take an algebraic closure k^{a} of k containing E, and look at the set of distinct *conjugates* of α in k^{a}; i.e., the images of α under all embeddings $E\to k^{\mathrm{a}}$. Then we see that both the *sum* and the *product* of these conjugates are invariant under $G(k^{\mathrm{a}}/k)$; hence both lie in the base-field k.

These two elements are useful in studying α, but there are problems when we consider more than one element at a time. For instance, we would like to be able to say that for $\alpha, \beta\in E$, the sum of the conjugates of $\alpha+\beta$ equals the sum of the conjugates of α plus the sum of the conjugates of β. But if α, β and $\alpha+\beta$ do not all have the same stabilizer in the Galois group, it may not be possible to pair off ''corresponding'' conjugates of these elements, as we would need to do to prove equality; and, indeed, the desired equality will not in general hold. What will work, however, is to associate to each $\alpha\in E$ the sum, respectively the product, of $\sigma(\alpha)$ as σ ranges over all embeddings $E\to k^{\mathrm{a}}$. These may be thought of as the sum and product of the conjugates of α ''counting multiplicities''. The multiplicities are such that we always have a sum or product of $[E:k]$ terms, regardless of the degree of the particular element α. This sum and product are called the *trace* and *norm* of α with respect to the extension E/k.

If we want to extend these concepts to possibly inseparable extensions, there are two plausible ways to go: We could again take the sum and product of $\sigma(\alpha)$ over all embeddings $\sigma: E\to k^{\mathrm{a}}$, or we could take the sum and product of all conjugates of α using a common multiplicity that makes their total number come to $[E:k]$. These give different values, because an inseparable extension E of k has fewer than $[E:k]$ distinct embeddings in k^{a}. Of these two approaches, the second turns out to have the better algebraic properties (the first may not even give elements of k), though the resulting functions are not as useful in the inseparable as in the separable case. (In particular, the trace is identically zero when E is inseparable over k.)

Thus, when you read the definitions of these functions in §VI.5, you should recognize the terms $p^{\mu} = [E:k]_i$ as supplying the extra multiplicity needed in the inseparable case. (You might find it easiest to understand this section if you first work through it assuming E/k is separable, so that $p^{\mu} = 1$, then reread it with the general case in mind.)

P.284, last two displays [~]: Lang writes (starting with the 4th printing) $N_{E/k}(\alpha) = N_k^E(\alpha) = ...$, and a similar formula for the trace, meaning that he will sometimes write these one way, and sometimes the other. (In earlier printings, he only had the second definition, but still used both notations.)

There is a general convention of mathematical typesetting, that one-letter symbols, when not in a special font such as boldface or script, are set in italics, while a symbol of more than one letter, if not in a special font, is set in roman. This is the reason the symbol N for ''norm'' is italic, while the symbol Tr for ''trace'' is roman. (Other examples of multi-letter symbols set in roman, used in Lang or elsewhere, are Irr, Hom, det, lim, sin, and log.)

The value of this convention is that if one sees an italic multi-letter expression such as Re, one can recognize it as a combination (e.g., denoting the left ideal of the ring R generated by an element e) rather than a single symbol (e.g., ''Re'' might denote ''the real part of –''). (Some French and German mathematical works use a variant convention, in which symbols consisting of a single *capital* letter are also set in roman, so that only single lower-case symbols are italicized; thus, in such works the above left ideal would be written $\mathrm{R}e$. There are still other conventions used by particular authors – for instance, some authors feel one should use roman letters for symbols with standard meanings, so that they write the base of the natural logarithms as e, and the derivative of y as $\mathrm{d}y/\mathrm{d}x$; or that one should interchange the roles of roman and italic symbols in italicized passages, so that what is ''$\sin x$'' in ordinary text becomes ''*sin* x'' in the statement of a theorem. But I don't recommend either of those conventions.)

P.286, proof of Theorem VI.5.2 [>]: Here Lang has in effect repeated the proof of Theorem III.6.4.

After that proof, read the three-line **Remark** on p.287, then come back to Corollary VI.5.3 and its proof (which in effect repeats the proof of Theorem III.6.1), then skip Corollary VI.5.4, which, as we noted earlier, belongs in the preceding section, and go on to Proposition VI.5.5.

P.287, proof of Proposition VI.5.5, first display [=]: Lang is very terse about the derivation of this formula. It can be looked at as an example of *Lagrange interpolation*. Recall that that is a technique for finding a polynomial of degree $<n$ that assumes n specified values at n specified points. One does this by first writing down n polynomials, each of which is zero at $n-1$ of the points and nonzero at the remaining one, then taking the unique linear combination of these that gives the right values at all the points. In this case, the n points are $\alpha_1, \dots, \alpha_n$, and a polynomial which is zero at all of these except α_i is $f(X)/(X-\alpha_i)$. To get the correct linear combination, we need to know the values of these polynomials at the α_i. If we write $f(X)/(X-\alpha_i) = g_i(X)$, then by differentiating the equation $(X-\alpha_i)g_i(X) = f(X)$ with respect to X and substituting $X = \alpha_i$, we find that $g_i(\alpha_i) = f'(\alpha_i)$; now the desired formula is easily deduced.

P.287, last display [~]: Here Lang is writing i for what he has called r so far.

It is curious that this proposition has established a dual basis for the powers of α in the field $k(\alpha)$ by performing a computation in the larger field $k(\alpha_1, \dots, \alpha_n)$.

A particular consequence of this result is that none of the β_i are zero; in fact, that they are linearly independent over k. Can you see a quick proof of either of these facts not using the proposition?

P.288, end of top paragraph [>]: Suppose E/k is a finite extension, and α an element of E. As noted in the above paragraph of Lang, multiplication by α can be looked at as a k-linear map $m_\alpha : E \to E$. Since we shall soon be looking at α as a member of two different fields, let us denote this operator on E more precisely as $m_{\alpha,E}$. What is the relation between the *minimal polynomial* of α and the *characteristic polynomial* of $m_{\alpha,E}$? (Recall from undergraduate linear algebra that the latter is the determinant of the matrix $XI - M_{\alpha,E}$, where I is the identity matrix, and $M_{\alpha,E}$ is the matrix representing $m_{\alpha,E}$ with respect to any k-basis of E.)

Let us first consider the case $E = k(\alpha)$. In that case, a basis for E over k is given by $1, \alpha, \dots, \alpha^{d-1}$, where $d = [E:k]$. It is easy to write down the matrix of $m_{\alpha,E}$ in terms of this basis, since that map sends the basis element α^i to the basis element α^{i+1} for each $i < d-1$, while it sends α^{d-1} to α^d, which is not one of the basis elements, but can be written as a linear combination of these elements using coefficients which are negatives of the corresponding coefficients of $\mathrm{Irr}(\alpha, k, X)$. The resulting matrix $M_{\alpha,E}$ is called the *companion matrix* of $\mathrm{Irr}(\alpha, k, X)$, and it is easy to compute its characteristic polynomial, which turns out to be precisely $\mathrm{Irr}(\alpha, k, X)$.

(Another way to show that the characteristic polynomial of $m_{\alpha,E}$ is $\mathrm{Irr}(\alpha, k, X)$, without computation, is to call on a result of linear algebra which you may have seen as an undergraduate, the *Cayley-Hamilton Theorem*, Theorem XIV.3.1 in Lang. This says that if m is a linear transformation on a finite-dimensional vector space, and f its characteristic polynomial, then $f(m) = 0$. Now since the map sending each $\beta \in E$ to $m_{\beta,E}$ is a k-algebra homomorphism from E to the algebra of k-vector-space endomorphisms of E, and $\mathrm{Irr}(\alpha, k, X)$ generates the ideal of polynomials satisfied by α in E, it also generates the ideal of polynomials satisfied by $m_{\alpha,E}$ in that endomorphism ring. But by the Cayley-Hamilton theorem, $m_{\alpha,E}$ satisfies its characteristic polynomial, hence the latter polynomial must be a multiple of $\mathrm{Irr}(\alpha, k, X)$. Since both these polynomials are monic of degree $[k(\alpha):k]$, they are equal, as claimed.)

In the case where E is not necessarily generated by α, let $w_1, \ldots, w_r$ be any basis of E over $k(\alpha)$, and let us take as a basis of E over k the set of products of the w's with the elements of some basis for $k(\alpha)$ over k. We find that the matrix representing $m_{\alpha,E}$ in terms of such a basis consists of a diagonal array of r copies of the matrix of $m_{\alpha,k(\alpha)}$. The characteristic polynomial of such a diagonal array is the product of the characteristic polynomials of the diagonal blocks, hence we find that in this general case, the characteristic polynomial of $m_{\alpha,E}$ is $\mathrm{Irr}(\alpha, k, X)^r$. i.e., $\mathrm{Irr}(\alpha, k, X)^{[E\,:\,k(\alpha)]}$.

Now it is easy to check that the linear map $m_{\alpha,k(\alpha)}$ has trace equal to the negative of the term-after-the-leading-term of $\mathrm{Irr}(\alpha, k, X)^{[E\,:\,k(\alpha)]}$. Indeed, of the diagonal entries of the companion matrix, only one is not automatically zero, and that one has the desired value. Similarly, the determinant of that map is $(-1)^{[k(\alpha):k]}$ times the constant term of $\mathrm{Irr}(\alpha, k, X)^{[E\,:\,k(\alpha)]}$. (Key to checking this: multiply the companion matrix by a matrix P that permutes the $[k(\alpha):k]$ rows cyclically so that the string of 1's moves onto the main diagonal. The determinant of the resulting matrix is easy to evaluate; note also that the matrix P we have multiplied by has determinant $(-1)^{[k(\alpha):k]+1}$.) But these coefficients, with the indicated adjustments of sign, are precisely $\mathrm{Tr}_{k(\alpha)/k}(\alpha)$ and $N_{k(\alpha)/k}(\alpha)$, by the definition of those trace and norm. Moreover, from the preceding paragraph we see that $m_{\alpha,E}$ has trace equal to $[E : k(\alpha)]$ times the trace of $m_{\alpha,k(\alpha)}$, and norm the $[E : k(\alpha)]$th power of the norm of $m_{\alpha,k(\alpha)}$. Now when $\mathrm{Irr}(\alpha, k, X)$ is raised to the power $[E : k(\alpha)]$, its term-after-the-leading-term and its constant term also get multiplied by $[E : k(\alpha)]$, respectively raised to that power. It follows that the statements that the trace and norm of the element α are the trace and determinant of the operation of multiplication by α, which we established for those operations on $k(\alpha)/k$, are also true on the arbitrary extension E/k. This is Proposition VI.5.6; but I thought it preferable that you see that these results on the trace and the norm are particular cases of a general observation relating the minimal and characteristic polynomials.

(I have often wondered whether one could get further interesting results by studying the other coefficients of the characteristic polynomial of $m_{\alpha,F}$.)

Incidentally, for a general (not necessarily irreducible) polynomial f, one defines its companion matrix in terms of the coefficients of f in exactly the same way. To interpret this matrix, note that the ring $k[X]/(f)$, even if it is not a field, still has a vector space basis $1, \alpha, \ldots, \alpha^{\deg(f)-1}$, where α is the image of X; so the companion matrix still represents the action of multiplication by α with respect to that basis.

Re §VI.6. Cyclic extensions.

P.288, statement of Theorem VI.6.1 (Hilbert's ‘‘Theorem 90’’) [>]: Let me give you

A brief motivation for Hilbert's Theorem 90 and its proof.

We know that the norm function of a field extension is invariant under the Galois group: $N(\sigma(\alpha)) = N(\alpha)$, and is multiplicative: $N(\alpha\beta) = N(\alpha)N(\beta)$. Hence if we have an extension with cyclic Galois group, $G(K/k) = \langle\sigma\rangle$, we see that we can get elements of norm 1 by taking quotients $\alpha/\sigma\alpha$ for $\alpha \in K-\{0\}$. Can we think of any other way to get such elements? Well, $\gamma/\sigma^m\gamma$ will similarly have norm 1 for any $m \in \mathbf{Z}$, $\gamma \in K-\{0\}$. But if we set $\alpha = \gamma(\sigma\gamma) \ldots (\sigma^{m-1}\gamma)$, we see that $\gamma/\sigma^m\gamma = \alpha/\sigma\alpha$, so this is still an example of the sort we first noted. Thus it seems plausible that expressions of the form $\alpha/\sigma\alpha$ give all elements of norm 1. Theorem VI.6.1 establishes this result. (The result first appeared as ''Theorem 90'' in Hilbert's book *Bericht über die Theorie der algebraischen Zahlkörper*, and has come to be called ''Hilbert's Theorem 90''.)

The proof Lang will give may be motivated as follows. Given β with norm 1, we want an α such that $\beta = \alpha/\sigma(\alpha)$. We can translate this to say we want to find an element α fixed by the k-linear map $\alpha \mapsto \beta\,\sigma(\alpha)$. Denoting this k-linear map φ, a quick computation shows that for all α, $\varphi^n(\alpha) = N(\beta)\,\sigma^n(\alpha)$, which by our hypotheses is α. Thus $0 = \varphi^n - 1 = (\varphi-1)(\varphi^{n-1} + \ldots + \varphi + 1)$. The second factor in this product is nonzero, by the linear independence of characters, hence has a nonzero element α in its image; this will lie in the kernel of $\varphi - 1$, i.e., will be an element of the form we want. You can verify that this operator $\varphi^{n-1} + \ldots + \varphi + 1$ is in fact the map used in the proof in Lang (first display on p.289).

P.289, end of proof of Theorem VI.6.1 [>]: Let us note

Another motivation for Hilbert's Theorem 90.

Hilbert's Theorem 90 can also be looked at in terms of a goal we are moving toward, and will achieve in the next section. This is to find necessary and sufficient conditions on a finite algebraic extension E/k for the elements of E to be *expressible in radicals*; i.e., for E to lie within a finite extension which is the top of a tower of fields, having k at the bottom, and such that each field in the tower is obtained from the field just below it by adjoining an element which has some power in that lower field. Now the properties of adjoining to a field k an nth root α of an element $a \in k$ (where n is not divisible by the characteristic of k) are simplest if our base field contains a primitive nth root of unity ζ_n. In this case, the extension is automatically Galois, since $X^n - a$ splits as $\prod\,(X - \zeta_n^i \alpha)$. I claim that if a is not an mth power in k for any $m > 1$ dividing n, then the Galois group is generated by an automorphism σ carrying α to $\zeta_n \alpha$. Indeed, the set of nth roots of unity (powers of ζ_n) such that our extension has an automorphism carrying α to itself times that root of unity is clearly a subgroup of $\langle \zeta_n \rangle$; thus it will be generated by ζ_n^m for some divisor m of n. We see that $\alpha^{n/m}$ is invariant under the Galois group, hence lies in k, and a is the mth power of this element, hence by our assumption on a, $m = 1$; this says that the Galois group has an element σ of the desired form. Further, each element of the Galois group is determined by which element $\zeta_n^i \alpha$ it sends α to, and all possibilities are achieved by powers of our σ, hence the Galois group is precisely $\langle \sigma \rangle$.

What about the converse? If k is a field containing a primitive nth root of unity ζ_n, and E a Galois extension of k with cyclic Galois group $\langle \sigma \rangle$ of order n, is E generated over k by an nth root of an element of k? It is not hard to see that this is equivalent to asking whether E has a nonzero element α such that $\sigma(\alpha) = \zeta_n \alpha$. Note that such a relation at least passes the test of consistency, in that applying σ successively n times would bring α back to α.

Rather than looking at this question in isolation, let us look generally at elements $\beta \in E$ (not necessarily in k) such that it would it be ''consistent'' for a nonzero element $\alpha \in k$ to satisfy

$$\sigma(\alpha) \;=\; \beta\alpha. \tag{c35}$$

To find the ''consistency'' condition needed, we compute

$$\begin{aligned}
\sigma^0(\alpha) &= \alpha,\\
\sigma^1(\alpha) &= \beta\alpha,\\
\sigma^2(\alpha) &= \sigma(\beta)\beta\alpha,\\
&\;\;\ldots\\
\sigma^n(\alpha) &= \sigma^{n-1}(\beta)\ldots\sigma(\beta)\beta\alpha,
\end{aligned}$$

Since $\sigma^n = 1$, the last equation simplifies to $\alpha = N(\beta)\alpha$, so the desired consistency condition is

$N(\beta) = 1$. Hilbert's Theorem 90 is the statement that this necessary condition is also sufficient for the existence of an element α satisfying (c35).

Returning to the particular case where $\beta = \zeta_n$ is a primitive nth root of unity in the base field k, we conclude that E will indeed contain an element α satisfying $\sigma\alpha = \zeta_n\alpha$, equivalently, an element whose nth power, but no lower power, satisfies $a \in k$; so $E = k(a^{1/n})$. The details are given in the proof of the next result in Lang, Theorem VI.6.2.

P.290, first sentence [~]:

Motivation for the additive form of Hilbert's Theorem 90.

We have seen that Hilbert's Theorem 90 is useful in studying cyclic Galois extensions where the base field contains a primitive nth root of unity. If we start with a cyclic extension K/k where k does not have such a root of unity, we may still be able to get information by adjoining such a root, and looking at what Lang calls the ''lifted'' extension $K(\zeta_n)/k(\zeta_n)$. But a situation where even this is not possible is if n is divisible by the characteristic of k, since no field of characteristic dividing n can contain a primitive nth root of unity.

What will a cyclic field extension look like in this case? The simplest situation to consider is that of an extension K/k of fields of characteristic p with Galois group $\langle\sigma\rangle$ cyclic of order p. Since we cannot find elements of K on which σ acts by multiplication by a primitive pth root of unity, what other plausible behavior can we come up with that, when applied p times, would bring us back to the original element?

Well, the behavior $\sigma(\alpha) = \alpha+1$ would have this property; so we may ask whether every cyclic extension of degree p in characteristic p has an element α satisfying this equation. Again, we can generalize: The condition on an element $\beta \in K$, where K/k is a cyclic extension of *any* order, and of *any* characteristic, that is needed for the equation $\sigma(\alpha) = \alpha+\beta$ to be consistent with $\sigma^n = 1$, is easily seen to be $\mathrm{Tr}(\beta) = 0$. The next result that Lang will prove, the ''additive form of Hilbert's Theorem 90'', says that this, too, is sufficient as well as necessary.

There is another way to look at that result, analogous to the ''brief motivation for Hilbert's Theorem 90'' that I gave in my comment on Lang's p.288. In that paragraph, if you replace ''norm'' by ''trace'', ''multiplicative'' by ''additive'', and ''quotient'' by ''difference'', you will get a brief motivation for Theorem VI.6.3.

P.290, first display [=]: As in the proof of the multiplicative form of Hilbert's Theorem 90, Lang uses exponential notation for the action of σ when writing out the formula for the desired element α; but this time he is inconsistent, mixing both notations in the same formula.

P.290, statement of Theorem VI.6.4 [<]: On the preceding page of Lang, when the multiplicative form of Hilbert's Theorem 90 was used to get an element α satisfying $\sigma(\alpha) = \zeta_n\alpha$, we noted that this would have minimal polynomial of the form $X^n - a$. If we similarly apply the additive form of Hilbert's Theorem 90 to get an element α satisfying $\sigma(\alpha) = \alpha+1$, what form will its minimal polynomial take? Well, the conjugates of α will be $\alpha, \alpha+1, \ldots, \alpha+p-1$, hence that polynomial can be written $\prod_{i=0,\ldots,p-1}(X-\alpha+i)$. To simplify this, let $Y = X-\alpha$; then it becomes $\prod(Y+i)$, i.e., the polynomial of degree p that has all elements of $\mathbf{Z}_p$ as roots. We recall that this polynomial is $Y^p - Y$. Hence $\prod(X-\alpha+i) = (X-\alpha)^p - (X-\alpha) = X^p - X - (\alpha^p - \alpha)$. This leads to the statement of Theorem VI.6.4.

P.290, next to last line [=]: ''$1, \ldots, p$'' should be ''$1, \ldots, p-1$''.

P.291, top paragraph [<]: The argument can be completed more briefly. After the first three words, continue: ''Indeed, σ must take α to $\alpha+i$ for some $i \in \{1, \ldots, p-1\}$, and under addition, i generates the group $\{0, \ldots, p-1\}$, hence $\langle\sigma\rangle$ is transitive on the set of roots $\{\alpha+i\}$ of f, proving irreducibility.''

P.291, end of §VI.6 [>]. Here are some further notes on the subject of this section.

Some exact sequences (optional).

Let K/k be a Galois extension with cyclic Galois group $\langle\sigma\rangle$. Hilbert's Theorem 90 can be looked at as saying that if we regard the norm function as a group homomorphism $K^* \to k^*$, then its kernel is the image of the endomorphism $\alpha \mapsto \sigma(\alpha)/\alpha$ of the group K^*. Now we know the kernel of that endomorphism; it is the set of elements $\alpha \in K^*$ fixed under σ, i.e., k^*. These results can be summarized by saying that the sequence of group homomorphisms

(c36) $$1 \to k^* \xrightarrow{\subseteq} K^* \xrightarrow{\alpha \mapsto \sigma(\alpha)/\alpha} K^* \xrightarrow{N} k^*$$

is exact.

This suggests the question: what can one say about the *cokernel* of the rightmost arrow above?

One case where we can answer this easily is if k and K are finite fields. Note that in an exact sequence of finite groups, each group is an extension of the image of the map coming into it by the image of the map going out of it. From this one can deduce that in a finite exact sequence beginning and ending with trivial groups, the product of the orders of the groups in even-numbered positions equals the product of the orders of the groups in odd-numbered positions. Thus, if we let C denote the cokernel of the last map of (c36), and tack $\to C \to 1$ onto that sequence, we get $(k^*:1)\cdot(K^*:1)\cdot(C:1) = (K^*:1)\cdot(k^*:1)$; hence $(C:1) = 1$. So if our fields are finite, the last map of (c36) is surjective.

If the fields are infinite, that map may or may not be surjective. For a negative example, let $k = \mathbf{R}$, $K = \mathbf{C}$. Then since the range of the norm map on nonzero complex numbers consists of the *positive* real numbers, the cokernel of the right-hand map of (c36) is isomorphic to Z_2.

The observations we made about exact sequences of finite groups also apply to exact sequences of *finite-dimensional* vector spaces, with sums of *dimensions* replacing products of group-orders. Since both fields involved in the additive form of Hilbert's Theorem 90 are finite-dimensional over k, we can conclude that the sequence

(c37) $$0 \to k \xrightarrow{\subseteq} K \xrightarrow{\alpha \mapsto \sigma(\alpha) - \alpha} K \xrightarrow{\mathrm{Tr}} k \to 0$$

is always exact. (Note that we generally use ''1'' to denote the trivial group when groups are written multiplicatively, ''0'' when they are written additively.)

Some linear algebra.

The special cases of Hilbert's Theorem 90 used in proving Theorems VI.6.2 and VI.6.4, namely the cases where $\beta = \zeta_n$ in the multiplicative version, and where $\beta = 1$ in the additive version, can be looked at in yet another way. If K is a cyclic extension of k, with Galois group $\langle\sigma\rangle$ of order n, then $\sigma: K \to K$, regarded as a linear transformation on the k-vector space K, satisfies the polynomial $t^n - 1$. If k has a primitive nth root of unity, this polynomial splits into *distinct* linear factors over k, hence the k-vector-space K has a basis B of eigenvectors of σ. I claim that no two of these eigenvectors have the same eigenvalue. For if distinct basis elements $\alpha, \beta \in B$ satisfied $\sigma(\alpha) = \zeta_n^i\alpha$ and $\sigma(\beta) = \zeta_n^i\beta$ with the same i, then α/β would be an element not in k which was fixed under σ, contradicting the Fundamental Theorem of Galois Theory. In particular, the n eigenvalues must comprise all n roots of $t^n - 1$, so some eigenvector α will have eigenvalue ζ_n, as desired. (We can see that up to scalars in k, the other eigenvectors will be the powers of α.)

In the additive case, where k has characteristic p, the polynomial $t^p - 1$ satisfied by σ factors as $(t-1)^p$. A linear transformation satisfying this polynomial has Jordan canonical form in which the nonzero entries consist of 1's down the main diagonal, and possibly some 1's in the next diagonal up. Since $\sigma \neq 1$, there must be at least one 1 above the main diagonal; looking at the ''highest'' of these, we see that the corresponding basis vector will be sent by σ to itself plus the preceding vector in our basis, which will be a fixed element. Thus, we get an α such that $\sigma(\alpha) = \alpha + c$ for some $c \in k$, and dividing by c, we get an element which adds 1 on applying σ, as in Theorem VI.6.4.

In §VI.10, some results are developed which can be regarded as analogs of Hilbert's Theorem 90 for noncyclic Galois extensions. Although there is not time to cover that section in 250A, I have put some notes on it in this Companion which bring out this relationship. That section and the notes to it may be read once one has read this section.

Re §VI.7. Solvable and radical extensions.

P.293, end of proof of Theorem VI.7.2 [>]: The result just proved answers the famous question of when a polynomial equation can be ''solved in radicals''! Galois Theory in fact arose out of the effort to answer this question. Note that if an equation has degree ≤ 4, its Galois group can be identified with a subgroup of S_4, a solvable group; so every polynomial equation of degree ≤ 4 is solvable in radicals. On the other hand, S_5 is not solvable (Theorem I.5.5), and there exist polynomials of degree 5 over $\mathbf{Q}$ whose Galois group is this group (Example 4 on p. 272; Example 6 on p. 273). Hence polynomials of degree 5 (and higher) are not in general solvable in radicals.

Another classical question which can be answered by Galois Theory is, ''Which lengths can be constructed from a given set of lengths by ruler and compass; and in particular for which n can one construct by ruler and compass a regular n-gon?'' For developments of these results, see, e.g., Hungerford's *Algebra*, or Stewart's *Galois Theory*.

Re §VI.8. Abelian Kummer theory.

This subject is named for Ernst Kummer (1810-1893).

P.293, first sentence of the section [=]: The phrase ''the theorem concerning cyclic extensions'' refers to Theorem VI.6.2 and its characteristic-p analog, Theorem VI.6.4.

P.294, 4th paragraph [~]: Note that if $a' = b^m a$ $(a, a', b \in k - \{0\})$, then to adjoin to k an mth root of a' is equivalent to adjoining an mth root of a. Thus, in classifying fields obtained by adjoining mth roots, we should use some system of elements which contains $b^m a$ if and only if it contains a. This is the idea behind looking at subgroups $B \subseteq k - \{0\}$ which contain all mth powers.

P.294, 5th paragraph [=]: In developing the properties of the map $\sigma \mapsto \omega_\sigma$, Lang fixes α and does not show it in the symbol ω_σ, though it appears in the definition thereof. Later in the paragraph, he switches to considering the dependence of this root of unity on both σ and α. He notes that this does not change if α is replaced by a different mth root of a, hence it can be regarded as a function of σ and a, which he writes $\langle \sigma, a \rangle$. Intuitively, this definition says $\langle \sigma, a \rangle = \sigma(a^{1/m})/a^{1/m} \in \boldsymbol{\mu}_m$, but because of the ambiguity of the symbol $a^{1/m}$, the actual definition must be a bit wordier. This operation turns out to be a bilinear map into a cyclic group of order m, so the duality of §I.9 becomes applicable; however, note that the groups here are written multiplicatively, in contrast to the notation of that section.

P.295, first display [=]: For ''$B/k^{*m} \approx G^\wedge$'' read ''$G \approx (B/k^{*m})^\wedge$''. Both are true, but the latter is the isomorphism he proves, and also the one that goes over to the infinite case.

P.295, fourth line of proof of Theorem VI.8.2 [=]: Lang says, ''Thus we may assume without loss of generality that B_2/k^{*m} is finitely generated, hence finite''. What he is asserting is that if we knew the desired result ''$k(B_1^{1/m}) \subseteq k(B_2^{1/m}) \Rightarrow B_1 \subseteq B_2$'' in the case where B_2/k^{*m} is finite, then by using the observation in the preceding sentence, we could deduce the same result in the general case. (Do you see how?) Since to prove that result it suffices to prove it under the assumption that B_2/k^{*m} is finite, that assumption will be made for the remainder of the paragraph.

Later in the proof, at the first display, change ''or'' to ''and also''.

P.295, sentence beginning in third line of last paragraph, ''Any finite subextension...'' [=]: The reasoning here is: Any finite subextension of K will still be Galois (Corollary VI.1.11), and its Galois group, being a finite homomorphic image of $\mathrm{Gal}(K/k)$, will be a finite abelian group of exponent p. Such a group is a direct product of cyclic abelian groups of exponent p (Theorem I.8.2), so by Corollary VI.1.6, our subextension will be a compositum of subfields with these Galois groups.

P.296, first display [~]: Why are we so interested in the operation $x \mapsto x^p - x$ that we invent a symbol for it? Because Theorem VI.6.4 showed us that every cyclic extension of degree p in characteristic p is generated by an element x satisfying $x^p - x = a$ for some a in the base field. Thus, the operation $x \mapsto x^p - x$ in characteristic p is analogous to the pth power operation in characteristic $\neq p$.

P.296, Theorem VI.8.3, definition of $\langle \sigma, a \rangle$ [~]: Just as, in the characteristic-not-dividing-m case, we made a definition which intuitively said $\langle \sigma, a \rangle = \sigma(a^{1/m})/a^{1/m} \in \boldsymbol{\mu}_m$, so here the idea is $\langle \sigma, a \rangle = \sigma(\wp^{-1}a) - (\wp^{-1}a) \in \mathbf{Z}_p$. Again this map turns out to be bilinear, so that we will be able to apply the duality of §I.9. (And this time, the groups in question are written additively.)

P.296, last paragraph [>]: In the characteristic-not-dividing-m case, there was no extra difficulty when m was divisible by the square of a prime, since one can still have primitive mth roots of unity for such m. But in characteristic p, while the map $\alpha \mapsto \alpha+1$ has order p, there is no obvious similar construction of order p^2. Hence more elaborate machinery is needed, for which Lang refers us to the exercises. (I have not had a chance to go through those and see how workable they are.)

P.296, end of §VI.8 [>]:

Overview of the above section, and one further result.

Suppose k is a field containing a primitive mth root of unity, and one adjoins to k the mth roots of a family of elements of k. Then the behavior of any element σ of the resulting Galois group is determined by specifying which mth roots of unity each of these adjoined elements gets multiplied by. Obviously, there are certain consistency conditions that must be satisfied. For instance, in $\mathbf{Q}(2^{1/2}, 3^{1/2}, 6^{1/2})$, what an automorphism σ does to $6^{1/2}$ will be determined by what it does to $2^{1/2}$ and $3^{1/2}$. Generally speaking, we can say that if we have adjoined mth roots $\alpha_1, \ldots, \alpha_r$ of elements $a_1, \ldots, a_r \in k$ (where this list may involve repetitions), and if the product $a_1 \ldots a_r$ is already an mth power in k, then $\alpha_1 \ldots \alpha_r$ must lie in k; hence the roots of unity by which σ multiplies these factors must have product 1.

Kummer theory tells us that these consistency conditions are the only restrictions on such automorphisms. To express this, one lets B denote the multiplicative group generated by all elements whose mth roots one wishes to adjoin to k, together with all nonzero elements which already have mth roots in k. For each element σ of the Galois group, one defines a homomorphism $\langle \sigma, - \rangle : B \to \boldsymbol{\mu}_m$, sending each $a \in B$ to the root of unity by which σ multiplies the mth roots of a. Theorem VI.8.1 says that, letting G denote $G(k(B^{1/m})/k)$, the resulting bilinear map

(c38) $$\langle -, - \rangle : \ G \times B \ \to \ \boldsymbol{\mu}_m$$

has kernels 1 and k^{*m} respectively on the left and right. Hence, if B/k^{*m} is finitely generated, equivalently, finite, we can conclude by Theorem I.9.2 that (c38) induces an isomorphism between G and the dual group to B/k^{*m} with respect to $\boldsymbol{\mu}_m$. Elements of this dual group represent possible ''consistent'' behaviors for maps $\langle \sigma, - \rangle$; hence every such behavior is realized by some σ.

Recall that the relation between subextensions of a Galois field extension and subgroups of its Galois group is order-*reversing*. The relation between subgroups of a finite abelian group and subgroups of its dual group is also order-reversing; hence Kummer theory yields an order-*preserving* correspondence between subextension of $k(B^{1/m})$ and subgroups of B containing k^{*m}.

What if B/k^{*m} is not finitely generated? In general, when we have a bilinear map from two abelian groups into a cyclic group, and we pass to the factor-groups of these groups by the appropriate kernels of the bilinear map, then each factor-group can be identified with a *subgroup* of the dual of the other. But I claim that in the present situation, we get an identification of the Galois group G with the whole of $\mathrm{Hom}(B/k^{*m}, \boldsymbol{\mu}_m)$. I will sketch the argument, which is another example of a ''compactness proof'' (cf. comments in this Companion on p.235, end of §V.2 of Lang), i.e., a proof that a certain infinite set of conditions will have a model if all its finite subsets do.

Given $h \in \mathrm{Hom}(B/k^{*m}, \boldsymbol{\mu}_m)$, we want to find a $\sigma \in G$ which induces h. Now for every subgroup $B_0 \subseteq B$ finitely generated over k^{*m}, the restriction of h to B_0 is a member of $\mathrm{Hom}(B_0/k^{*m}, \boldsymbol{\mu}_m)$, and so by the results proved in Lang, there is a unique $\sigma_{B_0} \in G_{B_0} = G(k(B_0{}^{1/m})/k)$ which induces this map. We see that given B_0 and B_1 each finitely generated over k^{*m}, the automorphisms σ_{B_0} and

σ_{B_1} agree on $k((B_0 \cap B_1)^{1/m})$. From the fact that B is the union of all its subgroups B_0 finitely generated over k^{*m}, it is not hard to deduce that these automorphisms fit together to give a single automorphism of $k(B^{1/m})$, which induces h, as claimed. (We ''put them together'' by taking their union as sets.)

(The analogs of the above considerations likewise apply to the case of exponent-p abelian extensions of a field of characteristic p.)

Re §VI.9. The equation $X^n - a = 0$.

P.297, statement of Theorem VI.9.1 [~]: This theorem gives necessary and sufficient conditions for a polynomial of the form $X^n - a$ to be irreducible. To motivate it, let us consider how such a polynomial might factor.

Clearly, if $X^n - a$ is reducible in $k[X]$, then its irreducible factors all have degrees $< n$, and conversely. (Incidentally, not all these factors need have the same degree, an extreme case being when $a = 1$.) So what we are asking is how an nth root (in an extension field) of an element $a \in k$ can have degree $< n$ over k. An obvious case where this happens is when $a = b^m$ where $b \in k$, $m \mid n$, and $m > 1$; then we can get a root of $X^n - a$ of degree $< n$ by taking an n/mth root of b. To say a is an mth power for some $m \mid n$ with $m > 1$ is equivalent to saying it is a pth power for some prime $p \mid n$, and this is the formulation used in the theorem.

There is another less obvious way $X^n - a$ can be reducible. It arises from the fact that $(1+i)^4 = -4$, so that the polynomial $X^4 - (-4)$ has the root $1+i$, which only has degree 2 over $\mathbf{Q}$, though -4 is not a square in $\mathbf{Q}$. The resulting factorization of $X^4 - (-4)$ over $\mathbf{Q}$ can by Gauss's Lemma be taken to be a factorization over $\mathbf{Z}$ (the factors are $(X^2 - 2X+2)(X^2+2X+2)$, cf. Remark 2 on p.183), and since $\mathbf{Z}$ can be mapped homomorphically into any field, this gives a factorization of $X^4 - (-4)$ over any field k. This in turn gives a factorization of $X^4 - (-4)b^4$ for any $b \in k$, and hence of $X^{4m} - (-4)b^4$ for any m. Thus, another necessary condition for $X^n - a$ to be irreducible is that if n has the form $4m$, a should not have the form $-4b^4$. Theorem VI.9.1 asserts that the two necessary conditions just noted are also sufficient for $X^n - a$ to be irreducible.

P.297, first two paragraphs of the proof of Theorem VI.9.1 [~]: This argument, showing that the statement of the theorem is true *if* it is true whenever n is a prime power, can be replaced by the following simpler argument, pointed out to me by Lenstra. Let α be a root of $X^n - a$, and let the prime factorization of n be $p_1^{r_1} ... p_t^{r_t}$. For each i, $\alpha^{n/p_i^{r_i}}$ is a $p_i^{r_i}$th root of a lying in $k(\alpha)$. Assuming the theorem is true for prime power exponents, this element has degree $p_i^{r_i}$, hence $p_i^{r_i} \mid [k(\alpha):k]$. Since this is true for all i, we get $n \mid [k(\alpha):k]$. But $[k(\alpha):k] \leq n$, so $[k(\alpha):k] = n$.

After this argument, continue from ''We now suppose that $n = p^r$ ...'' (11 lines from the bottom). However, here is a note about the next paragraph:

P.297, third from last paragraph [~]: Here Lang considers the purely inseparable case. The ''even more trivial argument'' that he refers to uses the Frobenius map. (Can you see how to do this?) To deduce the final sentence, note that the preceding sentence gives $[k(a^{1/p^r}) : k(\alpha)] = p^{r-1}$. It follows that $[k(a^{1/p^r}) : k] = p^r$.

Lang's induction also omits the base case, $r = 1$. To see that case, note that if $\alpha = a^{1/p} \notin k$, then the irreducible polynomial of α, being a divisor of $X^p - a$ of degree > 1, must have multiple roots. Then Proposition V.6.1 tells us that the degree is divisible by p; so it is not a proper divisor of $X^p - a$.

P.297, last paragraph [=]: The first word, ''Suppose'', should be ''By assumption''. On the second line, ''a root'' should be ''the root''. On the third line, ''and $\alpha^p = a$. Taking the norm from $k(\alpha)$ to k'' should be ''. Taking the norm from $k(\alpha)$ to k of the equation $\alpha^p = a$''. On the last line of the page, the justification for ''it follows'' is: the set of integers i such that a^i is a pth power in k clearly forms an additive group. It contains p by definition, and the preceding line shows that it contains d. As these integers are relatively prime, it contains 1.

P.298, second paragraph (''Assume ...'') [~]: Although one of our hypotheses is that a is not a pth power in k, this does not say whether or not, after we adjoin a pth root α of a to k, this α will be a pth power in the resulting field $k(\alpha)$, and the question becomes crucial to the proof. Hence Lang considers two cases. First he assumes it does not, and quickly completes the inductive argument. (Note that if p is odd, the assumption that α is not a pth power gives him the hypothesis for an inductive application of the theorem, while if $p=2$, he must in addition show that α is not of the form $-4\beta^4$ in this field. I would simplify the argument for the latter case as follows: After ''If $p=2$'' continue, ''... then our hypotheses says that α is not a square in $k(\alpha)$; hence applying the nontrivial automorphism of this field, we see that $-\alpha$ is also not a square, and so in particular cannot be written $4\beta^4$. Hence α cannot be written $-4\beta^4$ in $k(\alpha)$, and again by induction ...''.) The remainder of the proof of the theorem, starting with the last sentence of that paragraph, is devoted to showing that the case where α is a pth power in $k(\alpha)$ cannot occur.

P.298, before third-from-last display, ''In either case ...'' [=]: This is because -4 is a square in $k(i)$. The symbol $\pm$ is much used in the remainder of the proof. In each occurrence, it will simply mean ''either + or $-$'', with no implicit ''respectively''. That is, the ''+'' case of one occurrence need not correspond to the ''+'' case of another.

P.298, next-to-last display [>]: This contradiction has shown that the two factors of $X^{2^r}-a$ in $k(i)[X]$ are each irreducible in that ring, so by unique factorization in this ring, these are the only proper monic nonunit factors of $X^{2^a}-\alpha$ in $k(i)[X]$. Since these factors do not lie in $k[X]$, we conclude that the polynomial is irreducible in the latter ring, as desired.

P.299, end of line 2 [>]: To strengthen this result, after ''*or if* p *is odd*'' add a third case, ''*or if* -1 *is a square in* k''. (To see that the conclusion holds in this case, note that by the preceding cases, this case need only be considered when $p=2$ and 2 is not the characteristic. In that case, the condition ''*a is not a p-th power*'' together with -1 being a square implies $a\notin -4k^4$, as required to apply the theorem.)

P.299, statement of Corollary VI.9.3 [<]: Theorem VI.9.1, whose proof has taken up most of this section, is moderately interesting; but the real reason we took the time to prove it is for the sake of Corollary VI.9.3, a surprising result saying that any field k such that $1<[k^{\mathrm{a}}:k]<\infty$ must look very much like the single example that we know: the real numbers! Incidentally, that result is due to the same two people as Theorem VI.6.4, and the name ''the Artin-Schreier Theorem'' is applied to both statements.

P.299, second line of proof of Corollary VI.9.3 [=]: By ''inseparable over some subfield of degree > 1'' Lang means ''inseparable of degree > 1 over some subfield''.

P.299, middle of second paragraph of proof of Corollary VI.9.3 [>]: Lang refers to Exercise VI:L29; you should look over that exercise, and think about how to do it.

P.299, last four lines of second paragraph of proof of Corollary VI.9.3 [<]: This argument can be shortened with the help of the strengthened version of Corollary VI.9.2 mentioned above: Since k^{a} is the splitting field of X^p-a, and $i\in k_1\subseteq F$, that corollary shows that $X^{p^2}-a$ is irreducible, contradicting the fact that $[k^{\mathrm{a}}:F]=p$.

P.299, third paragraph of proof of Corollary VI.9.3 [<]: Let me give the content of this paragraph in a way that yields more explicit information about k. First note that for any $a\in k-\{0\}$, either a is a square, or $-a$ is a square, but not both. (Indeed, let σ be the nontrivial automorphism of $k^{\mathrm{a}}=k(i)$. If we take a square root of a and a square root of $-a$ in this field, then each is either fixed, or sent to its negative by σ. But their product is $\pm ia$, which is sent to its negative by σ, so *one* of these elements must be fixed, and the *other* sent to its negative. Thus, one lies in k and the other does not.) Let us call $a\in k-\{0\}$ ''positive'' if a is a square, and ''negative'' if $-a$ is a square. We claim that a sum of two positive elements, say u^2 and v^2, is positive. Indeed, u^2+v^2 can be written $(u+iv)(u-iv)$ in $k(i)$; note that neither of the factors is zero, so this element is nonzero. Since $k(i)$ is algebraically closed, we can take a square root of $u+iv$; say $u+iv = (x+iy)^2$ $(x,y\in k)$. Applying σ, we have $u-iv = (x-iy)^2$; multiplying these equations together, we see that $u^2+v^2 = (x^2+y^2)^2$, which is the desired result. Since 1 is positive, it now follows by induction that the sum of any number of 1's is positive, hence nonzero; hence k has characteristic 0, completing the proof of Corollary VI.9.3.

In fact, if for $a,b\in k$, we let ''$a > b$'' mean that $a-b$ is positive, then it is easy to verify that this is a total ordering on k, which obeys the usual laws relating inequalities and arithmetic operations. A field k with such an ordering is called an ''ordered field''; these are considered in Chapter XI of Lang. The above paragraph is essentially the proof of Proposition XI.2.4 (p.452).

Corollary VI.9.3 shows that a field k which has ''few'' algebraic extensions, in the sense that $[k^{\mathrm{a}}:k]$ is finite, must have very restrictive properties. We remark, however, if we weaken the sense of ''few extensions'', there are wider classes of examples. In particular, Exercises VI:L26 and VI:L27 (p.325) show two ways of getting fields with the property that every finite extension is cyclic, and it is not hard to show that such a field has at most one extension of each degree.

P.299, third from last line [=]: This result, that any nontrivial torsion element must have order 2, is true not only for $G(\mathbf{Q}^{\mathrm{a}}/\mathbf{Q})$ but for $G(k^{\mathrm{a}}/k)$ for any field k. This follows from Corollary VI.9.3: If σ is an automorphism of finite order $n>1$, then by Artin's Theorem, $[k^{\mathrm{a}}:(k^{\mathrm{a}})^{\langle\sigma\rangle}] = n$; but by the corollary referred to, this can only happen for $n = 2$. From that corollary we also see that if k has positive characteristic *or* contains a square root of -1, then $G(k^{\mathrm{a}}/k)$ has *no* nontrivial elements of finite order.

P.300, top [<]: We began this section by determining when a polynomial of the form $X^n - a$ was irreducible. It is natural to ask, in such a case, what its Galois group is. Lang introduces this subject as Example 2, and then obtains an answer, under certain assumptions, in Theorem VI.9.4. The idea is: first adjoin a primitive nth root of unity ζ_n, *then* adjoin an nth root α of a. The resulting Galois group will be an extension of its subgroup $G(k(\zeta_n,\alpha)/k(\zeta_n))$ by its factor-group $G(k(\zeta_n)/k)$. Under certain conditions, the first of these will be isomorphic to the additive group of $\mathbf{Z}/n\mathbf{Z}$, and the second to the multiplicative group of units of this ring, and the extension will be the natural semidirect product of these two groups. (You might look at the Galois group of the splitting field of X^3-3 over $\mathbf{Q}$ from this point of view.) Lang will describe this semidirect product as a group of matrices over the ring $\mathbf{Z}/n\mathbf{Z}$.

P.301, first two lines [=]: I don't understand the reasoning Lang gives after ''because''. The assertion that the commutator is the subgroup shown at the bottom of the preceding page is only true for n odd – you can easily check that for $n=2$, the commutator subgroup is instead the trivial subgroup – but Lang's words don't seem to use the assumption that n is odd. In any case, for n odd, consider the commutator of an element as on the bottom of the preceding page with the element as in the third display of that page having $b = 0$ and $d = 2$. (When n is odd, $2\in(\mathbf{Z}/n\mathbf{Z})^*$.)

P.301, 2nd through 8th lines after the diagram [=]: Where Lang says ''We apply the first part of the proof to $X^p-\beta$ over $k(\beta)$'', you should verify that the hypotheses are satisfied here, in particular that $[k(\beta,\boldsymbol{\mu}_p):k(\beta)] = \varphi(p)$. (Hint: $[k(\boldsymbol{\mu}_p):k]$ and $[k(\beta):k]$ are relatively prime.) In the next sentence, the point is that $k(\beta,\boldsymbol{\mu}_n)$ is an abelian subextension of $k(\beta)$, hence so is the left-hand side of the displayed equation; but by ''the first part of the proof, applied to $X^p-\beta$ over $k(\beta)$'', the intersection of $k(\alpha)$ with the maximal abelian subextension of the splitting field of $X^p-\beta$ over $k(\beta)$ is just $k(\beta)$.

Re §VI.10. Galois cohomology.

P.302, start of §VI.10 [<]:

Motivation of the ideas of this section.

Hilbert's Theorem 90, which we studied in §VI.6, concerned *cyclic* Galois extensions. Is an analogous result true in the noncyclic case?

The answer depends on how one tries to generalize the theorem. If one focuses on the element β, and regards the theorem as a way of constructing the most general element of norm 1, then the obvious analog in a noncyclic Galois extension would say that every element β of norm 1 can be written as a product, over various $\sigma \in G$ and $\alpha \in K^*$, of the elements $\sigma(\alpha)/\alpha$. This statement turns out to be false. (Here is the idea of a counterexample, shown to me by H. Lenstra. Start with a quadratic extension E/k such that E contains no square root of -1, but has an element β with norm $N_k^E(\beta) = -1$; the latter condition can be achieved by letting E be generated over k by a root β of an irreducible polynomial $X^2 + aX - 1$. Now take another quadratic extension F of k, and take their compositum EF within a common extension. Then $N_k^{EF}(\beta) = (-1)^2 = 1$; but if E and F are chosen properly, one can prove from the fact that $N_k^E(\beta) \neq 1$ that β is not a product of elements $\sigma(\alpha)/\alpha$ in the Galois extension EF.)

On the other hand, if we focus on the α in Hilbert's Theorem 90, we can look at that theorem as saying that if we prescribe in a consistent way how we want an element $\alpha \in K$ to behave under the action of our Galois group – specifically, what element of K it should be multiplied by on applying each element of that group – then we can in fact find an α having this behavior. Now if the Galois group is a cyclic group $\langle\sigma\rangle$, then to describe the behavior of α under this group, it suffices to say what element β it is multiplied by when the generator σ is applied, since its images under elements σ^i can be determined recursively from this datum. We have seen that the consistency condition is then $N(\beta) = 1$. For a general Galois group G, one might similarly specify the behavior of α under G by indicating how it is to be affected by the elements of some generating set for G; but it would be complicated to describe the general form the consistency conditions for these data would take. It is formally simpler just to think of the desired data as specifying for *every* $\sigma \in G$ the element β_σ by which α is multiplied on applying σ. The consistency condition is then easily shown to be

(c39) $$\beta_{\sigma\tau} = \beta_\sigma\,\sigma(\beta_\tau) \qquad (\sigma,\ \tau \in G).$$

In this section, Lang will prove a generalization of Hilbert's Theorem 90 saying that for every system of β's satisfying (c39), there exists an $\alpha \neq 0$ having the prescribed behavior under G. He will use homological language: a system of elements $(\beta_\sigma)_{\sigma \in G}$ $(\beta_\sigma \in K^*)$ is called a 1-*cocycle* in K^* if it satisfies (c39), and a 1-*coboundary* if it has the form $(\sigma(\alpha)/\alpha)_{\sigma \in G}$ for some $\alpha \in K^*$. The 1-coboundaries form a subgroup of the 1-cocycles, and the first assertion of Theorem VI.10.1 says that the factor-group of the 1-cocycles by the 1-coboundaries is trivial. This is the desired existence result. The second assertion gives the analogous generalization of the *additive* version of Hilbert's Theorem 90.

Some notational points to be aware of: Lang defines the concepts of 1-cocycle and 1-coboundary in a general abelian group on which G acts, which he writes additively, but in the first application, the group will be K^*, which one writes multiplicatively. I have used α and β above in such a way as to match the use of these symbols in the section on Hilbert's Theorem 90; but in this section Lang writes α where I have β.

P.303, second line from bottom [=]: ''so f is a coboundary'' should be ''so $(\tau - 1)f$ is a coboundary''.

Re §VI.11. Non-abelian Kummer extensions. We generally don't have time for these sections in Math 250A; but here are a few errata pointed out by Harry Gindi:

P.305, last line [=]: $\mathbf{Z}(p^{n(p)})$ should be $\mathbf{Z}/p^{n(p)}\mathbf{Z}$.

P.307, 4th-from-last display [=]: The final cy should be ct. On the next line, $\frac{1}{N}\Gamma$ should be Γ'.

P.308, line 5 [=]: $c_M(\Gamma)$ should be $e_M(\Gamma)$.

Re §VI.14. Infinite Galois extensions.

P.313, start of §VI.14 [<]: The first part of this section characterizes the Galois group of an infinite Galois extension K/k in terms of the Galois groups of its finite subextensions F/k. However, the description uses inverse limits (§I.10), which we generally don't look at in Math 250A. The latter part of the section consists mainly of references to advanced work in the literature. One of the basic results of the subject, the characterization of those subgroups of $G(K/k)$ that correspond to intermediate fields, is not mentioned, because the usual formulation is in topological terms.

Below I give, not a commentary on the section, but ''What I would have put in such section'', namely, a few motivating examples followed by a sketch of these two main results and their proofs, in a form that does not require concepts of topology or of inverse limits. I then point out for the student familiar with those tools how the results proved can be naturally reformulated in terms of them.

(I have a more extensive write-up of the subject, which I generally hand out in Math 250B, which assumes familiarity with the above two concepts, and also with the ring of p-adic integers, and which then develops a characterization of totally disconnected compact topological groups. It can be found at math.berkeley.edu/~gbergman/grad.hndts/infGal+profin.ps .)

Examples of infinite Galois extensions, and two theorems.

To get a feel for the subject, let us begin by determining the automorphism groups of three such extensions K/k. I will first list the examples, then sketch their properties, leaving most of the verification to you. (a) Let $k = \mathbf{Q}$, and adjoin the square roots of all positive rational numbers, equivalently, of all primes, getting $K = \mathbf{Q}(2^{1/2}, 3^{1/2}, 5^{1/2}, ...)$. (b) Let F be a field (of characteristic not 2) containing primitive 2^nth roots of unity for all n, let k be a simple transcendental extension field $k = F(t)$, and look at the Galois extension gotten by adjoining a 2^nth root of t for each n: $K = F(t, t^{1/2}, t^{1/4}, \ldots)$. (c) Fix a prime p, let $k = \mathbf{Q}$, and let $K = \mathbf{Q}(\zeta_p, \zeta_{p^2}, \ldots, \zeta_{p^m}, \ldots)$, where ζ_n denotes a primitive nth root of unity.

In case (a), we see that any $\sigma \in G(K/k)$ is determined by what it does to the elements $p^{1/2}$: some of these elements will be fixed, and others sent to their negatives. With the help of the results in §VI.8 (Kummer Theory), and my comments on that section, you can verify that for any set S of primes, there exists a σ such that $\{p \mid \sigma p^{1/2} = -p^{1/2}\}$ is precisely S. It follows that the Galois group of our extension is isomorphic to the group of all ± 1-valued functions on the set of primes, under pointwise multiplication; that is, to a direct product of countably many copies of Z_2.

In case (b), an element of the automorphism group $G(K/k)$ is similarly determined by what it does to $t^{1/2}$, $t^{1/4}$, etc.; but here the actions are not independent: It can send $t^{1/2}$ to $t^{1/2}$ or $-t^{1/2}$; if it does the former, it must send $t^{1/4}$ to $t^{1/4}$ or $-t^{1/4}$, while if it sends $t^{1/2}$ to $-t^{1/2}$, it must send $t^{1/4}$ to $it^{1/4}$ or $-it^{1/4}$. Likewise, for each behavior of an automorphism on $t^{1/4}$, there are two possible behaviors on $t^{1/8}$, etc.. Associating to each σ the sequence of roots of unity by which the successive elements $t^{2^{-n}}$ are modified, what we have is, first an element of $\boldsymbol{\mu}_2$, then an inverse image of this element in $\boldsymbol{\mu}_4$ under the squaring map $\boldsymbol{\mu}_4 \to \boldsymbol{\mu}_2$, then an inverse image of that element under the squaring map $\boldsymbol{\mu}_8 \to \boldsymbol{\mu}_4$, etc.. Replacing these groups $\boldsymbol{\mu}_{2^n}$ by the isomorphic groups $\mathbf{Z}/2^n\mathbf{Z}$ (e.g., if $F = \mathbf{C}$, identifying $[1] \in \mathbf{Z}/2^n\mathbf{Z}$ with $e^{2\pi i/2^n} \in \boldsymbol{\mu}_{2^n}$), we can say that any element of $G(K/k)$ is determined by a sequence $(\ldots, c_8, c_4, c_2)$ of elements taken from the successive terms of the diagram of groups and surjective homomorphisms

(c40) $$\ldots \to \mathbf{Z}/8\mathbf{Z} \to \mathbf{Z}/4\mathbf{Z} \to \mathbf{Z}/2\mathbf{Z},$$

such that each c_{2^i} is the image of the preceding element, $c_{2^{i+1}}$, under the corresponding arrow; and that the group operation of $G(K/k)$ corresponds to componentwise addition of such sequences. (In the language Lang introduces in §I.10, this group is the *inverse limit* of the system of groups and homomorphisms (c40). This particular inverse limit is the additive group of what is called the *ring of* 2*-adic integers*.) Note that every nonidentity element of this group of sequences has infinite order; for an element whose component at one stage has order 2^i $(i > 0)$ will have component of order 2^{i+1} at the stage immediately to the left, so there is no finite exponent common to all its components. Nevertheless, *locally*, i.e., on each finitely generated subextension $F(t, t^{1/2}, t^{1/4}, \ldots, t^{2^{-n}})$, each element of this group acts as an automorphism of finite order. Note also that this Galois group, like the preceding, is uncountable.

In case (c), the interested reader can verify, with the help of Theorem VI.3.1, that we have an analogous characterization in terms of the sequence of abelian groups of units in finite rings:

$$\ldots \ \to\ (\mathbf{Z}/p^3\mathbf{Z})^* \ \to\ (\mathbf{Z}/p^2\mathbf{Z})^* \ \to\ (\mathbf{Z}/p\mathbf{Z})^*.$$

One can also show that if p is odd, these groups are isomorphic to the additive groups

$$\ldots \ \to\ \mathbf{Z}/p^2(p-1)\mathbf{Z} \ \to\ \mathbf{Z}/p(p-1)\mathbf{Z} \ \to\ \mathbf{Z}/(p-1)\mathbf{Z}.$$

Using Proposition I.4.3(v), one can break this chain up as the direct product of a constant sequence of cyclic groups, $\mathbf{Z}/(p-1)\mathbf{Z}$, and the sequence

$$\ldots \ \to\ \mathbf{Z}/p^2\mathbf{Z} \ \to\ \mathbf{Z}/p\mathbf{Z} \to 0,$$

which, like (c40), yields an uncountable torsion-free abelian group. So in this case, our Galois group is the product of a finite group and an uncountable torsion-free group. (A similar result is true when $p = 2$, except that the torsion part does not have order $p-1$.)

I have given commutative examples for simplicity and brevity, but you should not find it hard to cook up noncommutative examples as well.

A description of how the Galois group of an arbitrary infinite Galois field extension is related to those of its finite subextensions, generalizing the above examples, is given in the following result, which is straightforward to prove:

Theorem VI.14.c1. *Let K/k be a (not necessarily finite) Galois extension of fields. Let X denote the set of all subfields $F \subseteq K$ that are* finite *and Galois over k, or, more generally, any subset of that set having the property that every subfield of K finite over k is contained in some member of X.*

Then any automorphism σ of K over k is determined by its restrictions to the members of X. Given a family $(\sigma_F)_{F\in X} \in \prod_{F\in X} G(F/k)$, the necessary and sufficient condition for (σ_F) to arise in the above manner from an element of $G(K/k)$ is the ''consistency condition'', that for every pair of fields $E \subseteq F$ both in X, the restriction homomorphism $G(F/k) \to G(E/k)$ carry σ_F to σ_E.

Thus, $G(K/k)$ can be identified with the subgroup of $\prod_{F\in X} G(F/k)$ consisting of elements $(\sigma_F)_{F\in X}$ satisfying this consistency condition. (This subgroup of the product is known as the inverse limit *of the groups $G(F/k)$ with respect to the system of restriction maps.)*

Idea of Proof. The ''consistency condition'' is clearly necessary, and conversely, if it holds for a family (σ_F), then the σ_F, regarded as partial automorphisms of K, fit together to give a total automorphism. □

We now come to a more meaty question: We know from Theorem VI.1.2 that every field L between k and K ''belongs to'' (i.e., is the fixed field of) some subgroup of $G(K/k)$; namely, $G(K/L)$. But which subgroups of $G(K/k)$ ''belong to'' subextensions? The answer is usually given in topological terms, but it can be stated using, instead, the following definition:

Given a set H of automorphisms of K, let us call an automorphism σ of K *H-approximable* if for every finite Galois subextension F of K, there is an element $\tau\in H$ such that $\tau|_F = \sigma|_F$. (Examples: Let $k = \mathbf{Q}$ and $K = \mathbf{Q}(2^{1/2}, 3^{1/2}, 5^{1/2}, \ldots)$, and let H consist of those automorphisms that change the sign of $p^{1/2}$ for only finitely many primes. You should find it easy to verify that *every* automorphism is H-approximable. The same is true for the set H of those automorphisms that change the sign of $p^{1/2}$ for an *even* finite number of primes. But if we let H be the set of automorphisms that change the change the sign of $p^{1/2}$ for a finite number of primes and such that the number of members of $\{2^{1/2}, 3^{1/2}, 5^{1/2}\}$ whose signs are changed is even, then the H-approximable automorphism are just the automorphisms that have the latter property; equivalently, the automorphisms that fix $30^{1/2}$.)

It is straightforward to check that the operation carrying each set H of automorphisms to the set of all

H-approximable automorphisms is a *closure operator* (as defined above in a comment to p.121 of Lang). We shall call that set the *approximation closure* of H, and call H *approximation-closed* if it is its own approximation closure. It is also easy to verify that when H is a subgroup of $G(K/k)$, its approximate-closure is again a subgroup. Let us now prove

Theorem VI.14.c2. *Let K/k be a Galois extension of fields. If H is any subgroup of $G(K/k)$, and K^H its fixed field, then $G(K/K^H)$ i.e., the subgroup belonging to K^H, is the* approximation closure *of H. Hence, the subgroups of $G(K/k)$ belonging to subextensions of K are precisely the approximation-closed subgroups.*

Proof. The last sentence is easily deduced from the preceding sentence, which we shall prove.

First we note that every automorphism σ in the approximation closure of H fixes every element $\alpha \in K^H$. Indeed, given $\alpha \in K^H$, we can find a finite Galois subextension F of K which contains it, and by definition of approximation-closure, there is an element $\tau \in H$ which agrees with σ on F. Since $\alpha \in K^H$, τ fixes α, hence so does σ.

Conversely, suppose $\sigma \in G(K/k)$ fixes all elements of K^H, and let F be any finite Galois subextension of K. We wish to find an element of H which agrees with σ on F. On restricting elements of H to F, we get a group H_F of automorphisms of F over k whose fixed field is $F \cap K^H$. By assumption, the restriction of σ to F fixes all members of this subfield, hence by Artin's Theorem, it must lie in H_F, i.e., there must exist some $\tau \in H$ such that $\tau|_F = \sigma|_F$. Thus, σ is H-approximable, as required. □

Remarks for students familiar with point-set topology. It is not hard to verify that what I have called ''approximation closure'' is the operation of closure under a *topology* on $G(K/k)$. (A basis N of open neighborhoods of the identity in this topology is given by the set of subgroups $G(K/F)$ belonging to the finite Galois subextensions $F \subseteq K$; a basis of open neighborhoods of a general point σ is given by the translate of this family, σN.) If you have seen the concept of an inverse limit of topological spaces, you can verify that the same topology arises by looking at $G(K/k)$ as the inverse limit of the finite *discrete* groups $G(F/k) \cong G(K/k)/G(K/F)$. A finite discrete set is compact, and an inverse limit of compact spaces is compact, so it is a compact topology. Lang refers to it in §VI.14 as the ''Krull topology''.

One may ask *which* topological groups occur as inverse limits of systems of finite groups. These turn out to be those which are compact Hausdorff and totally disconnected, a result proved in the handout mentioned earlier.

There is an interesting open question related to such groups. Note that for any positive integer n, the inverse limit of a system of finite groups all of which have exponent n will again have exponent n. (So, for instance, the Galois group of $\mathbf{Q}(2^{1/2}, 3^{1/2}, 5^{1/2}, \ldots)$ over $\mathbf{Q}$ has exponent 2.) In particular, an inverse limit of finite groups all having a common exponent will be a group whose elements all have finite order. The question is whether the converse holds. I will pose this question below in two equivalent forms; note that form (b) does not require familiarity either with topological groups or with inverse limits.

Open Question VI.14.c3. (a) *If G is a compact topological group all of whose elements have finite order, must G be of finite exponent? Equivalently,*

(b) *Given a sequence of* finite *groups and surjective homomorphisms, $\ldots \to G_3 \to G_2 \to G_1$, such that the groups G_i have no common finite exponent, can one always find a sequence of elements $(\ldots, g_3, g_2, g_1) \in \prod G_i$ such that for each i, the map $G_{i+1} \to G_i$ takes g_{i+1} to g_i, and such that the g_i have unbounded orders?*

The proof of equivalence of (a) and (b) requires nontrivial results that I will not go into here. (Note that (a) does not assume the group is totally disconnected.) I am told that there is in the literature at least one *incorrect* proof of an affirmative answer to the above open question.

Chapter VIII. Transcendental Extensions.

So far, I only have notes on one section of this chapter:

Re §VIII.1. Transcendence bases.

P.355, beginning of §VIII.1 [<]: If k is a field, $X_1, X_2, \ldots$ are independent indeterminates over k, and $k(X_1, \ldots, X_n) \cong k(X_1, \ldots, X_m)$ as k-algebras, must $m = n$? The results of this short section will show that the answer is yes. In fact, they will show that if some algebraic extension of $k(X_1, \ldots, X_n)$ is isomorphic as a k-algebra to some algebraic extension of $k(X_1, \ldots, X_m)$, then $m = n$. Thus, the number of independent indeterminates involved in this way in an extension K of a field k is an invariant of K.

P.355, display [=]: The notation $M_{(\nu)}(S)$ is as defined on p.101.

P.355, last four lines [>]: We really want to define the transcendence degree to be the common cardinality of all maximal algebraically independent subsets, but we have not yet proved that these are the same; hence Lang defines it to be the greatest of the cardinalities of such subsets. However, this hedge does not really work: we don't know at this point that there *is* such a greatest cardinality. E.g., it is conceivable that K might have maximal algebraically independent subsets of all finite cardinalities, but none of infinite cardinality, or that there might exist such subsets with cardinalities $\aleph_0$, $\aleph_1$, ... but none with cardinality $\aleph_\omega$. So we really ought to wait till we have proved that all such maximal subsets have the same cardinality before defining the transcendence degree.

The alternative to the term ''transcendence degree'' that Lang gives, ''dimension'', I do not recommend, because it invites confusion with dimension as a vector space over k. (The reason ''dimension'' may sometimes be used is that $k(X_1, \ldots, X_n)$ can be thought of as the field of rational functions on ''n-dimensional space'' over k; and, more generally, the transcendence degree of what algebraic geometers call ''the function field associated with an algebraic variety'' is the dimension of that variety.)

Ignore the sentence on the last two lines (''Actually ...''); Lang simply doesn't want to bother with proving the infinite case of the next theorem.

P.356, third paragraph [=]: where Lang says ''**transcendence base**'', I recommend ''transcendence **basis**''.

P.356, statement and proof of Theorem VIII.1.1 [<]: If you have an Addison Wesley printing before the 6th, the proof of this theorem there is incorrect, so be sure to read the proof given in the Errata section of this Companion. That incorporates the five comments marked [=] given immediately below, so in that case you can ignore these.

P.356, sentence containing the first display [=]: This is true because by the sentence before the theorem, w_1 is algebraic over $k(x_1, \ldots, x_m)$. Thus it satisfies a polynomial with coefficients in that field, and by clearing denominators we can rewrite this as a polynomial in $m+1$ indeterminates satisfied by $w_1, x_1, \ldots, x_m$.

P.356, sentence after first display, ''After renumbering ...'' [=]: This is because the preceding equation must actually involve at least one of the x_i, otherwise it would show w_1 algebraic over k, contradicting the assumption of algebraic independence. By renumbering the x's, we may assume that it involves x_1.

P.356, two lines later [=]: ''Otherwise w_1 would be a root of two irreducible polynomials ...''. Namely, f_1 and the irreducible factor of g_N in question. Note that if this factor vanished on the indicated n-tuple, it would indeed have to involve w_1, so as not to contradict algebraic independence of the x's; also, it cannot be the same irreducible polynomial as f_1, since f_1 involves x_1 and it does not.

P.356, three lines before last display [=]: change ''have found $w_1, \ldots, w_r$ $(r < n)$ such that'' to ''have found for some $r < n$ that''.

P.356, last line [=]: The final sentence holds because if we had $n > m$, then the preceding sentence would imply that w_{m+1} was algebraic over $k(w_1, \ldots, w_m)$, contradicting algebraic independence of $w_1, \ldots, w_n$.

Notes on the above proof. The idea of the above proof is to take a polynomial relation $f_1 = 0$ expressing the algebraicity of w_1 over $k(x_1, \ldots, x_m)$, choose any x_i that this relation involves, and then use the same relation to show that x_i is algebraic over the extension generated by w_1 and the other x's, and in this way ''trade off'' x's for w's one by one. However, there is a dangerous trap: If one merely takes any nonzero polynomial f_1 satisfied by w_1 over $k[x_1, \ldots, x_m]$, and any x_i that it ''involves'' (i.e., that occurs as a factor of some monomial in f_1), it will not necessarily be true that $f_1 = 0$ constitutes a nontrivial relation satisfied by x_i over the subfield in question. For instance, if $w_1 = x_2 + x_3$, and we take $f_1 = (x_1 - 1)(w_1 - x_2 - x_3)$, then we see that $f_1 = 0$ holds, and f_1 involves x_1, yet as a polynomial in x_1 over $k(w_1, x_2, x_3)$, it is trivial, because of the relation $w_1 = x_2 + x_3$ holding in the latter field. (And in fact, in this situation x_1 is *not* algebraic over $k(w_1, x_2, x_3)$.) So one needs some condition to keep this from happening. The condition Lang uses is that f_1 be irreducible; one that I think gives the nicest proof is to let f_1 be chosen so that the set of x's that it involves is minimal. Another way to get around the problem is to start the proof differently: Where the present proof ''threw in a w'' and showed that we could then ''eject an x'', the proof in Hungerford's text ''pulls an x out'' and shows that one of the w's can ''take its place''. That argument doesn't seem to contain a trap like the above.

More than one text, including printings of Lang before the 6th Addison Wesley printing, give proofs of this theorem that fall into the above trap. I taught from Lang for years and believed the argument; it was only in preparing to teach the same material in an undergraduate course, from another text that used the same argument, and thinking about exactly how to justify it step-by-step for the class, that I saw there was a problem!

P.357, lines 3-4 [>]: Lang says, ''The cardinality statement in the infinite case will be left as an exercise''. That proof uses the same idea as the proof of unique cardinality for bases of non-finitely-generated free abelian groups, given in Lemma I.7.c2 (this *Companion*, comments on Theorem I.7.3). Though there we were concerned with the construction ''the subgroup generated by'', and we here are dealing with the relation ''algebraic over'', the method still works, because if an element w is algebraic over an extension $k(X)$, then there is some finite subset of $X_0 \subseteq X$ such that w is algebraic over $k(X_0)$. You should be able to supply the remaining details.

Now that we know that the transcendence degree of K over k is well defined, we should have a symbol for it. The one I am used to, and will use in class, is $\mathrm{tr.deg.}_k\, K$.

The other point Lang leaves as an exercise is that one can build up a transcendence basis containing S and contained in Γ. To do this, take a maximal algebraically independent subset $\mathcal{B}$ of Γ containing S, and note that all elements of Γ, hence all elements of $k(\Gamma)$, hence all elements of K, must be algebraic over $k(\mathcal{B})$. So $\mathcal{B}$ is the desired transcendence basis.

P.357, end of §VIII.1 [>]:

Some notes and examples.

If K/k is any field extension, and we take a transcendence basis S for K over k, so that

$$k \subseteq k(S) \subseteq K,$$

then we see that the first of these inclusions is a field-of-rational-functions construction, also called a *pure transcendental extension*, while the second is an *algebraic* extension. Thus, every field extension can be written as an algebraic extension of a pure transcendental extension. However, though the pure transcendental extension $k(S)/k$ is unique *up to isomorphism* (since we have shown that the cardinality of the algebraic independent generating set is unique), it is far from unique. E.g., if $S = \{x\}$, a singleton, then $\{x^2\}$ is another transcendence basis. For another example, if S is any finite set $\{t_1, \ldots, t_n\}$, then the set $\{s_1, \ldots, s_n\}$ of elementary symmetric polynomials in these generators is also a transcendence basis. (Proof?)

If k is countable, it is easy to show that an extension gotten by adjoining to k countably many independent indeterminates will again be countable, and we also noted when discussing the construction of algebraic closure that an algebraic extension of a countable field is countable. Hence an extension of a countable field having countable transcendence degree is countable. It follows that the fields of real and complex numbers have *uncountable* transcendence degree over the rationals.

You should find it easy to verify that the field of ''trigonometric functions'' on the real line, i.e., the field of fractions of $\mathbf{R}[\sin x, \cos x]$, has transcendence degree 1 over $\mathbf{R}$. (In fact, it is pure transcendental, i.e., it has the form $\mathbf{R}(f(x))$ for some trigonometric function $f(x)$; but we cannot take this function to be either $\sin x$ or $\cos x$. Can you see why not, and can you find such an f?)

It is not hard to prove that a family of real or complex functions $e^{a_1 x}, e^{a_2 x}, \ldots, e^{a_n x}$ is algebraically independent over $\mathbf{R}$, respectively, $\mathbf{C}$, if and only if $a_1, \ldots, a_n$ are *linearly* independent over $\mathbf{Q}$. The key step is to verify first that a family of functions $e^{c_1 x}, \ldots, e^{c_m x}$ is *linearly* independent over $\mathbf{R}$ or $\mathbf{C}$ if and only if $c_1, \ldots, c_m$ are distinct.

You may have noticed a parallelism between the proofs of existence and unique cardinality for bases of vector spaces, and for transcendence bases of field extensions. There is, in fact, a general concept of which both are instances:

General ''dependence structures'', or ''matroids'' (sketch, optional).

Recall that if V is a vector space over a field, then the operator taking each subset $X \subseteq V$ to the subspace that it spans is a *closure operator*, i.e., that if we write the subspace spanned by X as $\mathrm{cl}(X)$, then the operator cl satisfies (c21)-(c23) (this Companion, comment to p.121 of Lang). You should likewise verify that if K is an extension of a field k, then the operator taking a subset $X \subseteq K$ to the field of all elements of K *algebraic* over $k(X)$ is a closure operator. Like most closure operators arising in algebra, these two satisfy the strengthening of (c22) mentioned (but not named) in our earlier discussion,

(c41) $$\mathrm{cl}(X) = \bigcup_{\text{finite } X_0 \subseteq X} \mathrm{cl}(X_0).$$

A closure operator satisfying (c41) is generally called *finitary*, or *algebraic*,

The above two examples also satisfy a much rarer condition, called the ''exchange property''. This says that for all $X \subseteq S$ and $s, t \in S$,

(c42) $$s \in \mathrm{cl}(X \cup \{t\}) - \mathrm{cl}(X) \implies t \in \mathrm{cl}(X \cup \{s\}).$$

(A more colorful term might be the ''You scratch my back, I'll scratch yours'' condition.) You should have no trouble writing down the proof of (c42) in the vector-space case; while in the field-extension case, you should not find it hard to prove (c42) by a modification of the inductive step of the argument I gave to fix Lang's incorrect proof of Theorem I.7.3.

If cl is *any* closure operator on a set S that satisfies (c41) and (c42), then let us say that a subset $X \subseteq S$ *spans* S if $\mathrm{cl}(X) = S$, and that X is *independent* if no $x \in X$ lies in $\mathrm{cl}(X - \{x\})$. Zorn's Lemma, together with (c41), shows that S has a maximal independent subset X. Assuming (c42), one can prove that any such subset spans S; so, calling an independent subset which spans S a *basis*, one has a proof of the existence of bases which includes both the vector-space and the field-extension cases. Moreover, one can prove in this context that any two bases of S have the same cardinality.

A set S with a closure operator cl satisfying (c41) and (c42) is often called a ''dependence system'' by algebraists. The same type of structure was discovered independently by people thinking about the relation of linear dependence among rows of a matrix, and has been extensively studied under the name *matroid*. A good work on the subject is *Matroid Theory* by D. J. A. Welsh, Academic Press (1976). I will use this term below (despite the fact that matroid theorists often restrict their attention to finite matroids, a restriction that I will not make).

Note that if S is a matroid and T a subset of S, then T becomes a matroid under the closure operator $\mathrm{cl}_T(X) = \mathrm{cl}(X) \cap T$; we may call (T, cl_T) a *submatroid* of (S, cl). One of the games in matroid theory is to see whether various sorts of matroids are isomorphic to submatroids of vector spaces.

For example, given a graph Γ (in the sense of graph theory), there is a classical matroid structure on its edges, under which the closure of a set X of edges consists of those edges e such that the two vertices of e are connected by a path in X. If one maps the set of edges of Γ into a vector space V having the vertices of Γ as a basis, by sending each edge to the difference of its vertices (in either order), then the above matroid structure is induced by that of the vector space V.

For another example, given a vector space V, if we let P be the *projective space* consisting of all 1-dimensional subspaces of V, then P has a natural matroid structure under which the closure of a set of points is the projective subspace that they span. This matroid can be embedded in the matroid structure of V by mapping each point of P to an arbitrary nonzero representative in V. There is also an axiomatic definition of a projective geometry P, which is equivalent to saying that P is a matroid with certain additional properties. It turns out that if $\dim(P)$ (which for geometric reasons is defined in this case to be one less than the number of elements in a basis of P) is > 2, then P will be isomorphic to the projective space associated with a vector space over some division ring. On the other hand, there are known to exist projective planes (2-dimensional projective geometries) which cannot be so represented.

Even the matroid structure associated with algebraic dependence can be embedded in that of a vector space when the characteristic is 0, via the theory of *derivations* (Lang, §VIII.5). However, in positive characteristic it cannot, so far as I know.

Note that the definition of algebraic dependence is based on that of linear dependence: An element α is algebraic over a field k if and only its powers $1, \alpha, \alpha^2, \ldots$ are linearly dependent over k. Now differential equations can be studied algebraically, using the concept of a *differential field* k, that is, a field given with a derivation $x \mapsto x'$. Here the introduction of solutions to differential equations becomes the construction of appropriate differential extension fields. An element f of a differential extension field K/k satisfies some differential equation over k if and only if its derivatives $f, f', f'', \ldots$ are *algebraically dependent* over k. This leads to a theory of *differential dependence* and *differential closure*, which is built upon algebraic dependence in the same way that the latter is built upon linear dependence, and which can again be shown to be a closure operator satisfying (c41) and (c42).

The above are all the examples I know of matroids. The concept, though elegant where it is applicable, only seems to apply to a limited class of natural cases.

(I make use of matroids and their relation with vector spaces in a paper which I consider one of my most interesting, *Constructing division rings as module-theoretic direct limits,* Trans. A. M. S. **354** (2002) 2079-2114.)

Chapter X. Noetherian Rings and Modules.

For this chapter, like the preceding, I have notes on only one section.

Re §X.4. Nakayama's Lemma.

Observe that despite the title of the chapter, rings in this section are *not* assumed Noetherian (first sentence of the section).

P.425, line after first display (or in the Addison Wesley printings, top line of page) [=]: For ‘‘namely a_s’’ read ‘‘namely $-a_s$’’.

P.425, end of proof of Lemma X.4.1 [>]: More generally, this proof works if we assume E is a left module over a not necessarily commutative ring A, and $\mathbf{a}$ is any 2-sided ideal contained in all maximal *left* ideals of A. For then, at the top of this page, we can say that the *left* ideal generated by $1+a$ is not proper, hence $1+a$ is left invertible, allowing us to complete the proof.

Remark: The intersection of all maximal left ideals of a ring coincides with the intersection of all maximal right ideals, and is thus a 2-sided ideal, called the (*Jacobson*) *radical* of A. So the assumption on $\mathbf{a}$ is really that it is contained in this radical. (Some properties of the Jacobson radical are developed in Exercises XVII:L1-5, p.661.)

Here is an alternative proof of the lemma, which I like. Obviously, it suffices to prove that a *nonzero* finitely generated module E satisfies $\mathbf{a}E \neq E$. Now because E is finitely generated, it has a maximal proper submodule F. (Do you recall how this is proved?) Thus, E/F is a simple module; hence the annihilator of every element of this module is a maximal left ideal of A. Therefore, $\mathbf{a}(E/F) = 0$, so $\mathbf{a}E \subseteq F$, so $\mathbf{a}E \neq E$.

The lemma is false without the finite generation hypothesis. For instance, if we take for A any local integral domain that is not a field, for $\mathbf{a}$ its maximal ideal, and for E its field of fractions, then $\mathbf{a}E = E$, but $E \neq 0$.

The next two lemmas in Lang, which both say virtually the same thing, modify the statement of the preceding lemma simultaneously in two ways: First, they apply it to a factor-module E/F, to get a statement that under certain conditions, a submodule $F \subseteq E$ must be all of E; and second, they assume A local, so that without loss of generality $\mathbf{a}$ can be taken to be the maximal ideal. These two modifications can, of course, each be made separately.

These and the remaining results of this section are also true for noncommutative A. (For the definition of a noncommutative local ring, see Exercise II.4:3.) Unfortunately, in noncommutative ring theory one cannot in general get local rings from arbitrary rings by localization, so these results, though still important, are not as powerful a tool in studying general noncommutative rings as in the commutative case.

P.425, proof of Theorem X.4.4 [=]: Where Lang says ''The first statement'', he means the two sentences, ''*Then* E *is free. In fact, ...*''. Likewise, by ''the second statement'' (bottom line of this page) he means the last sentence of the theorem.

Chapter XII. Absolute Values.

We don't cover this in Math 250, but Johann Schuster has sent me one erratum:

Re §XII.1. Definitions, Dependence, and Independence.

P.466, third display [=]: The inequality should be ''<''.

Also, the final conclusion of the proof would be clearer if, after defining λ earlier, one had noted that this is the real number such that $|x_0|_2 = |x_0|_1^{\lambda}$. At the end, one combines this with the definition of α as the real number such that $|x|_1 = |x_0|_1^{\alpha}$, and the equality one has just proved.