SUMMARY OF GALOIS THEORY (PT. 1) (NIS-B 1995).

25/10/95.

§1 Field extensions

(Much of the material in §1 and 2 was covered in the IB Rings and Modules course.)

Recall that a *field* is something in which the elements can be added, subtracted, multiplied and divided (except that division by zero is prohibited) and all the usual rules of arithmetic are true. In particular, addition and multiplication are commutative. Examples are $\mathbb{Q}$ (the rationals), $\mathbb{R}$ (the reals), $\mathbb{C}$ (the complexes), $\mathbb{F}_p := \mathbb{Z}/p\mathbb{Z}$ (the integers modulo p, where p is prime). A further very important class of examples is given as follows. Start with a PID R and an irreducible element f of R. Then the quotient ring $R/(f)$ is a field. For example, if K is a field, then $K[X]$ is a PID.

If K is a subfield of the field L, then L is an *extension* of K. We also say that "L/K is a (field) extension".

Assume that L is an extension of K and that $\alpha \in L$. Then $K(\alpha)$ is the subfield of L generated by K and α; concretely, this is the set of elements of L that can be written (not necessarily uniquely) as quotients of polynomials in α with coefficients in K. $K[\alpha]$ is the set of elements of L that can be written (not necessarily uniquely) as polynomials in α with coefficients in K. So there is a surjective ring homomorphism $\rho : K[X] \to K[\alpha]$ such that $\rho(\sum a_n X^n) = \sum a_n \alpha^n$. (N.B.: X will always denote an indeterminate.)

Definition/Proposition 1.1. *(1.1.1) α is algebraic if it is a root of a nonzero polynomial f in $K[X]$, that is, if $\ker\rho$ is non-zero. In this case $\ker\rho$ is the ideal generated by a unique monic irreducible polynomial f (recall that $K[X]$ is a PID, so that every ideal in $K[X]$ is generated by a single element). In this case f is the minimal polynomial of α. Moreover, $K[\alpha]$ is then isomorphic to $K[X]/(f)$, and so is a field. It is then equal to $K(\alpha)$. Note that $f(\alpha) = 0$, so that although f is irreducible over K, it picks up a root (= zero, by abuse of language) in the bigger field $K(\alpha)$.*

Conversely, given an irreducible polynomial $f \in K[X]$, we can construct a bigger field L in which f has a root α, by taking $L = K[X]/(f)$ and α to be the residue class of X in L. In fact, this particular L has the following important property (which is the first step in proving the uniqueness of splitting fields): if M/K is any extension in which f has a root β, then the given inclusion $K \hookrightarrow M$ extends to an inclusion $L \hookrightarrow M$, and this extension is unique subject to the requirement that $\alpha \mapsto \beta$.
(1.1.2) α is transcendental if it is not algebraic. In this case ρ is an isomorphism.
(1.1.3) α is separable (over K) if either it is transcendental or if it is algebraic and is a root of a polynomial with coefficients in K having no multiple roots.

Typeset by $\mathcal{A}\mathcal{M}\mathcal{S}$-TEX

(1.1.4) A field extension $K \subset L$ is called finite if L is finite-dimensional as a vector space over K. (N.B. The use of the word "finite" here does **NOT** *imply that K or L is a finite set.) This dimension is denoted $[L : K]$ and is called the degree of the extension.*

Lemma 1.2. *Suppose that α is algebraic over K and that its minimal polynomial $f \in K[X]$ has degree n. Then $[K(\alpha) : K] = n$, and $\{1, \alpha, \dots, \alpha^{n-1}\}$ is a K-basis of $K(\alpha)$.*

Proof. Exercise. □

Theorem 1.3. *("The tower law for field extensions.") If L/K and M/L are finite extensions, then so is M/K, and $[M : K] = [M : L].[L : K]$. Conversely, if M/K is finite, then so are M/L and L/K, and again $[M : K] = [M : L].[L : K]$.*

Proof. If $\{a_1, \dots, a_m\}$ is a K-basis of L and $\{b_1, \dots, b_n\}$ is an L-basis of M, then $\{a_i b_j\}$ is a K-basis of M. □

§2 Splitting fields

Given $f \in K[X]$, a splitting field for f over K is a field extension L/K such that f splits completely (that is, factors as a product of linear terms in $L[X]$,) and if M/K is any extension in which f splits completely, then we can embed L in M. Concretely, L is the field extension of K generated by the roots of f. One obvious but important remark is that if $K \hookrightarrow L \hookrightarrow M$ are fields and M/K is a splitting field for $f \in K[X]$, then M/L is a splitting field for $f \in L[X]$.

Theorem 2.1. *(1) Splitting fields exist.*
(2) Splitting fields are unique. More precisely, if $\phi : K \longrightarrow K_1$ is an homomorphism, if $\phi(f) = f_1$, if L is a splitting field for f over K and L_1 is a splitting field for f_1 over K_1, then ϕ can be extended to a homomorphism $\Phi : L \longrightarrow L_1$. Moreover, if ϕ is an isomorphism, then so is Φ.

This was proved in the Rings and Modules course last year.

§3 Separable extensions

Remember that if K is a field, then $K[X]$ is a PID, so a UFD. Hence any two elements of $K[X]$ have an HCF. Moreover, if H is the HCF of F and G, then $H = AF + BG$ for some $A, B \in K[X]$. That is, in terms of ideals, $(F, G) = (H)$.

Lemma 3.1. *Assume that L/K is a field extension and that $F, G \in K[X]$. Then the HCF of F and G in $K[X]$ is equal to their HCF in $L[X]$.*

Proof. Suppose that H is the HCF of F, G in $K[X]$. Then $\frac{F}{H}, \frac{G}{H} \in K[X]$ and there are $A, B \in K[X]$ such that $1 = A.\frac{F}{H} + B.\frac{G}{H}$. Suppose now that $P \in L[X]$ divides $\frac{F}{H}$ and $\frac{G}{H}$; then P divides 1, and so is in L. This means that H is the HCF of F, G in $L[X]$, as stated. □

A finite extension L/K is *separable* if every element of L is separable over K.

Theorem 3.2. *("The theorem of the primitive element".) (3.2.1) If L/K is finite and $L = K(\alpha, \beta_1, \ldots, \beta_r)$, where each β_i is separable over K, then there is an element θ of L such that $L = K(\theta)$.*
(3.2.2) If L/K is finite and separable, then there is an element θ of L such that $L = K(\theta)$.

Proof. We shall assume that K is infinite. The finite case will be covered later in the course.

It is clearly enough to prove (3.2.1). Induction on the number of generators of L easily reduces us to the case where $L = K(\alpha, \beta)$ and β is separable over K. Let $f, g \in K[X]$ be the minimal polynomials of α, β. Take a splitting field M/K of fg such that M contains L. Say $\alpha = \alpha_1, \ldots, \alpha_r$ are the zeroes of f and $\beta = \beta_1, \ldots, \beta_s$ are those of g Then choose $c \in K$ such that the elements $\alpha_i + c\beta_j$ are all distinct. Put $\theta = \alpha + c\beta$.

Define F by $F(X) = f(\theta - cX)$. Note that $F \in K(\theta)[X]$. We have $g(\beta) = 0$ and $F(\beta) = f(\alpha) = 0$. So β is a zero of F and g. In fact, β is their only common zero. For the other zeroes of g are $\beta_2, \ldots, \beta_s$ and for $i \neq 1$, $F(\beta_i) = f(\alpha + c(\beta - \beta_i))$, while by construction $\alpha + c(\beta - \beta_i)$ is never equal to an α_j.

Now consider the HCF H of F, g in the ring $K(\theta)[X]$. By Lemma 3.1, H is also their HCF in $M[X]$. Since β is the only common zero of F, g, β is the only zero of H in M. Since g is a product of linear terms in $M[X]$, so is H. Now g has no repeated factors, since it is separable, so the same holds for H. Hence $H = X - \beta$. But $H \in K(\theta)[X]$, and so $\beta \in K(\theta)$. Hence $\alpha \in K(\theta)$, so that $K(\alpha, \beta) \subset K(\theta) \subset K(\alpha, \beta)$. □

Note that if $f \in K[X]$, then the derivative of f with respect to X exists, although in a purely formal fashion. We denote it by df/dX or f'.

Proposition 3.3. *If* $\operatorname{char} K = 0$ *then any algebraic extension is separable.*

Proof. Suppose that $\alpha \in L$ with minimal polynomial $f \in K[X]$. Suppose that f has a repeated root β in a splitting field M, say. Then f and f' have a common zero, namely β, so that in $M[X]$ they are not coprime. Hence they are not coprime in $K[X]$, by Lemma 3.1. But f is irreducible in $K[X]$, so that $f' = 0$ identically. Then $f = \sum_n a_n X^{pn}$ where $a_n \in K$ and $p = \operatorname{char} K$. □

Definition. Suppose that L/K and M/K are extensions. Then a *K-homomorphism* $L \to M$ is a homomorphism $\phi : L \to M$ of fields such that $\phi(x) = x$ for all $x \in K$.

Lemma 3.4. *Suppose that L/K is algebraic, that $\alpha \in L$ with minimal polynomial f and that M/L contains a splitting field of f. Say* $\deg f = n$*. Then there are at most n distinct K-homomorphisms $K(\alpha) \to M$, and α is separable over K if and only if there are exactly n such homomorphisms.*

Proof. The K-homomorphisms $K(\alpha) \to M$ correspond precisely to the elements of M whose minimal polynomial is f. The result is now immediate. □

Proposition 3.5. *Suppose that α is algebraic and separable over K and that $\beta \in K(\alpha)$. Then β is separable over K.*

Proof. Suppose that f is the minimal polynomial of α and g that of β. Let M/K be a splitting field of fg. Say $[K(\beta) : K] = m$ and $[K(\alpha) : K(\beta)] = n$, so that

$[K(\alpha) : K] = mn$. By 3.4, there are mn distinct K-homomorphisms $K(\alpha) \to M$. On the other hand, given a K-homomorphism $\phi : K(\beta) \to M$, there are exactly n extensions of ϕ to $K(\alpha)$, again by 3.4. Since there are at most m such ϕ, by 3.4, there must be exactly m such, and we are done.

Proposition 3.6. *If α is algebraic and separable over K and β is algebraic and separable over $K(\alpha)$, then $L = K(\alpha, \beta)$ is separable over K.*

Proof. By (3.2.1) we have $L = K(\theta)$. Say $[K(\alpha) : K] = m$ and $[L : K(\alpha)] = n$. Then there are m distinct K-homomorphisms $\phi : K(\alpha) \to M$, where M is some fixed sufficiently large splitting field, and for each such ϕ there are n extensions of ϕ to a homomorphism $L \to M$. Hence there are mn distinct K-homomorphisms $L \to M$, so that θ is separable over K. Now 3.5 shows that every element of L is separable over K. □

Corollary 3.7. *If L/K and M/L are finite and separable, then so is M/K.*

Proof. Immediate from 3.6. □

§4 Galois extensions: first properties

Suppose that G is a *finite* group of automorphisms of a field L. Define the *field of invariants* $K = L^G = \{x \in L \mid g(x) = x \forall\, g \in G\}$. (It really is a field; this is easy to check.)

Definition 4.1. An extension L/K is *Galois* if there is a finite group G of automorphisms of L such that $K = L^G$.

Lemma 4.2. *Every element of L is algebraic over K, of degree at most $\#G$.*

Proof. Suppose $\alpha \in L$. Then α is a zero of the polynomial $\prod_{g \in G}(X - g(\alpha)) = f(X)$, say. Note that f is G-invariant, so that $f \in K[X]$. Since α is a zero of f, we are done. □

Lemma 4.3. *L/K is separable.*

Proof. Suppose that $\alpha \in L$. We must show that its minimal polynomial f is separable. Consider the set $\{s(\alpha) \mid s \in G\}$. Suppose that its distinct elements are $\{\alpha = \alpha_1, \dots, \alpha_r\}$ and put $g = \prod_i (X - \alpha_i)$. Then g is separable, and since its linear factors are permuted by G, it is G-invariant. Hence $g \in K[X]$. Also, $g(\alpha) = 0$, so that $f \mid g$. Then f is separable. □

Lemma 4.4. *L/K is finite.*

Proof. By Lemma 4.2, we can find $\alpha \in L$ such that $[K(\alpha) : K]$ is maximal. Assume that $\beta \in L - K(\alpha)$ (for else we are done). Since $[K(\alpha, \beta) : K(\alpha)] \leq [K(\beta) : K]$, it is finite, as is $[K(\alpha) : K]$. Hence, by the tower law, $K(\alpha, \beta)/K$ is finite. Since $K(\alpha, \beta) \subset L$, $K(\alpha, \beta)/K$ is separable. So by the theorem of the primitive element, we can write $K(\alpha, \beta) = K(\gamma)$. Then $K(\alpha) \subset K(\gamma)$, and $[K(\gamma) : K] \leq [K(\alpha) : K]$ by maximality. Hence $K(\gamma) = K(\alpha)$, so that $\beta \in K(\alpha)$, which is absurd. □

Theorem 4.5. $[L : K] = \#G$.

Proof. By Lemmas 4.3. and 4.4 and the theorem of the primitive element, we have $L = K(\alpha)$. As before, put $f(X) = \prod_{s\in G}(X - s(\alpha))$; then $f \in K[X]$. Let $g \in K[X]$ be the minimal polynomial of α over K. Then g divides f, since α is a zero of f. So $\deg g \leq \deg f = \#G$, while $[L : K] = \deg g$. So it is enough to show that f is irreducible. Then suppose that $f = f_1 f_2$, $f_i \in K[X]$. Since $L[X]$ is a UFD, there is a decomposition $G = G_1 \cup G_2$ of G into disjoint subsets such that $f_i = \prod_{s\in G_i}(X - s(\alpha))$. Without loss of generality, $1 \in G_1$; choose $t \in G_2$. Since $f_i \in K[X]$, $t(f_i) = f_i$. However, $X - t(\alpha)$ is a factor of $t(f_1)$ but not of f_1. □

Note that so far in this § we have regarded the field L and the group G as fundamental, and then constructed the subfield K afterwards. However, we shall in fact spend most of this course going in the other direction; that is, we start with an extension L/K and then consider the group $\mathrm{Aut}(L/K) = \{s \in \mathrm{Aut}(L) \mid s(x) = x \;\; \forall \;\; x \in K\}$.

Lemma 4.6. *If $s_1, \dots, s_n \in \mathrm{Aut}(L)$ are distinct, then they are linearly independent over L. That is, if $l_1, \dots, l_n \in L$ such that $\sum l_i s_i = 0$ (that is, $\sum l_i s_i(x) = 0$ for all $x \in L$), then $l_i = 0$ for all i.*

Proof. Without loss of generality we can assume that $\sum l_i s_i = 0$ is a shortest linear relation. Then all $l_i \neq 0$ and $n \geq 2$.

Since $s_1 \neq s_2$, there exists $y \in L$ with $s_1(y) \neq s_2(y)$. Now $\sum l_i s_i(yx) = 0$ for all x, so that $\sum l_i s_i(y).s_i(x) = 0$ for all x. On the other hand, multiply the equation $\sum l_i s_i(x) = 0$ by $s_1(y)$ and subtract from the previous equation; the result is $\sum_{i\geq 2} l_i(s_i(y) - s_1(y))s_i(x) = 0$, which is a shorter relation. □

Proposition 4.7. *Suppose that L/K is a finite extension. Then* $\mathrm{Aut}(L/K)$ *is finite.*

Proof. Note that the earlier results of this § show that every finite subgroup of $\mathrm{Aut}(L/K)$ has order at most $[L : K]$. However, this is not enough to prove 4.7. Instead, we shall use Lemma 4.6.

Say $[L : K] = n$, and pick a K-basis $\{x_i\}$ of L. If the result is false, then we can find distinct elements $s_1, \dots, s_{n+1}$ of $\mathrm{Aut}(L/K)$. Consider the $n \times n$ matrix $A = (s_j(x_i))$. If $\det A = 0$, then its columns are linearly dependent, so that there exist $m_1, \dots, m_n \in L$ such that $\sum_j m_j s_i(x_j) = 0$. However, this contradicts Lemma 4.6. So $\det A \neq 0$, so that there exist $l_1, \dots, l_n \in L$ such that $\sum l_j s_j(x_i) = s_{n+1}(x_i)$ for all i. Then $s_{n+1} - \sum l_j s_j = 0$, contradicting Lemma 4.6. □

Lemma 4.8. *If G is a finite group of automorphisms of a field L, then $G =$* $\mathrm{Aut}(L/L^G)$.

Proof. Put $K = L^G$ and $H = \mathrm{Aut}(L/K)$. Note that $G \subset H$. By 4.6 H is finite. Then by Theorem 4.5 $\#G = [L : K] = \#H$, so that $G = H$. □

§5 Galois extensions and separable splitting fields

The aim here is to show that these two classes of extensions are the same.

Theorem 5.1. *(5.1.1) A finite extension L/K is Galois if and only if it is separable and is the splitting field of some polynomial in $K[X]$.*
(5.1.2) L/K is Galois if it is the splitting field of a separable polynomial in $K[X]$.

Proof. (5.1.1) Assume that L/K is separable and is the splitting field of $f = f_1.\dots.f_r$, where $f_i \in K[X]$ is irreducible. Then the f_i are separable, and we can assume that they are distinct. We shall argue by induction on $\deg f$.

Suppose that α, β are roots in L of f_1. There is an isomorphism $\psi : K(\alpha) \to K(\beta)$ such that $\alpha \mapsto \beta$ and $\lambda \mapsto \lambda$ for all $\lambda \in K$. Then $L/K(\alpha)$ is a splitting field for $g = \frac{f}{(X-\alpha)}$, say, and $L/K(\beta)$ is a splitting field for $h = \frac{f}{(X-\beta)}$. Since $\psi(g) = h$, we can extend ψ to an isomorphism $\Psi : L \to L$, by Theorem 2.1(2). Note that $L/K(\alpha)$ is a splitting field for the separable polynomial $\frac{f}{(X-\alpha)}$, so is Galois by the induction hypothesis. Put $H = \mathrm{Aut}(L/K(\alpha))$ and $\mathrm{Aut}(L/K) = G$; by Proposition 4.7 both groups are finite, while clearly $H \subset G$. Put $K_1 = L^G$. Since $K(\alpha) = L^H$, $K_1 \subset K(\alpha)$. Similarly, $K_1 \subset K(\beta)$.

Assume that $K_1 \neq K$. Then $[K(\alpha) : K_1] < [K(\alpha) : K]$, so that f_1 factors over K_1. Say $f_1 = p.q.r$, where p, q are irreducible. Since f_1 is separable, p, q, r are paiwise coprime. So we can choose α to be a root of p and β to be a root of q.

Note that $K_1(\alpha) = K(\alpha)$ and $K_1(\beta) = K(\beta)$. Then Ψ induces an isomorphism $K_1(\alpha) \to K_1(\beta)$ such that $\alpha \mapsto \beta$, while $\Psi(x) = x$ for all $x \in K_1$, since $\Psi \in \mathrm{Aut}(L/K)$. Hence α and β have the same minimal polynomial over K_1, which is absurd. Hence $K_1 = K$.

Conversely, suppose that L/K is Galois. Then by Lemma 4.3 it is separable, so that by Theorem 3.2 $L = K(\alpha)$, say. Put $f(X) = \prod_{s \in G}(X - s(\alpha))$. Then L is a splitting field for f over K.
(5.1.2) Suppose that L/K is a splitting field for the separable polynomial $f \in K[X]$. If the roots of f are $\{\alpha_1, \dots, \alpha_r\}$, then $L = K(\alpha_1, \dots, \alpha_r)$, so that L is generated over K by elements whose minimal polynomials are separable (because they are factors of f). Then L is obtained from K by successively adjoining separable elements, and so is separable over K. □

§6 The fundamental theorem of Galois theory

Theorem 6.1. *(6.1.1) If L/K is a finite Galois extension with $\mathrm{Aut}(L/K) = G$, then there is a one-to-one correspondence between the subgroups of G and the fields between K and L given as follows:*
(i) Given a subgroup H of G, the corresponding field is $H' = L^H = \{x \in L \mid s(x) = x \ \forall\ s \in H\}$.
(ii) Given a field M between L and K, the corresponding subgroup is $M' = \{s \in G \mid s(x) = x \ \forall\ x \in M\}$.

We have $H'' = H$ and $M'' = M$.
(6.1.2) This correspondence reverses inclusions. That is, if $H_1 \subset H_2$ are subgroups of G, then $H_1' \supset H_2'$, while if $M_1 \subset M_2$ are intermediate fields, then $M_1' \supset M_2'$.
(6.1.3) If $H_1 \subset H_2$ are subgroups of G and $M_i = H_i'$, then $[M_1 : M_2] = [H_2 : H_1]$.
(6.1.4) L/H' is Galois, and $\mathrm{Gal}(L/H') = H$.

(6.1.5) H is normal in G if and only if H'/K is Galois. In this case $\mathrm{Gal}(H'/K) \cong G/H$. *More generally, if M is the Galois closure in L of H'/K, then* $\mathrm{Gal}(M/K) \cong G/H_1$, *where $H_1 = \cap_{s\in G} s^{-1}Hs$.*

Proof. (6.1.1) Suppose first that $K \subset M \subset L$. Put $H = M'$. Note that $H = \mathrm{Aut}(L/M)$. By 5.1 L/K is a separable splitting field, say for $f \in K[X]$. Then L/M is a splitting field for f, where f is regarded as an element of $M[X]$. Since f is separable, L/M is Galois, by 5.1. So, by the definition of "Galois", $H' = M$. That is, $M'' = M$.

Now suppose that $H \subset G$. Put $M = H'$. Then L/M is Galois, with group H, by 4.7. So $M' = H$, so that $H'' = H$.

Thus the maps $H \mapsto H'$ and $M \mapsto M'$ are inverse to each other. This proves (6.1.1).
(6.1.2) This is obvious.
(6.1.3) Since L/M_i is Galois with group H_i, we have $[L : M_i] = \#H_i$. Then $[H_2 : H_1] = \#H_2/\#H_1 = [L : M_2]/[L : M_1] = [M_1 : M_2]$, where we have used 4.4 and the tower law.
(6.1.4) This is obvious.
(6.1.5) Suppose that $H \subset G$ and $M = H'$. Let $s \in G$, $x \in M$ and $h \in H$. Then $s^{-1}hs(s^{-1}(x)) = s^{-1}(h(x)) = s^{-1}(x)$. Hence $s^{-1}Hs$ acts trivially on the subfield $s^{-1}(M)$ of L, so that $s^{-1}Hs \subset (s^{-1}(M))'$. Since $[(s^{-1}(M)) : K] = [M : K]$, it follows that $s^{-1}Hs = (s^{-1}(M))'$. Now assume that H is normal; then $s^{-1}Hs = H$, by definition, so that $s^{-1}(M) = M$, by (6.1.1). So there is a map $\phi : G \to \mathrm{Aut}(M/K)$ given by $\phi(s)(x) = s(x)$. It is easy to check that ϕ is a homomorphism and that $\ker\phi = H$. Since $\#\mathrm{Aut}(M/K) \le [M : K] = [L : K]/[L : M] = \#G/\#H = \#(G/H)$, it follows that ϕ is an isomorphism.

Conversely, suppose that M/K is Galois. Then M/K is the splitting field of some separable $f \in K[X]$. Suppose that α is a root of f and that $s \in G$. Then $0 = s(0) = s(f(\alpha)) = f(s(\alpha))$, so that G permutes the roots of f. But M is generated over K by the roots of f, so that G preserves M. That is, there is a homomorphism $\phi : G \to \mathrm{Aut}(M/K)$ given by $\phi(s)(x) = s(x)$. By definition, $\ker\phi = M' = H$, say, so that $\#G/\#H = [L : M] = \#\mathrm{Aut}(M/K)$. Hence ϕ is an isomorphism. □

Theorem 6.2. *Suppose that M/K is a finite separable extension, that $M = K(\theta)$, that $f \in K[X]$ is the minimal polynomial of θ and that L/K is a splitting field for f. Then L/K is the minimal Galois extension containing M. That is, L/K is Galois and if L_1/K is any Galois extension containing M, then there is a homomorphism $\phi : L \to L_1$ such that $\phi(x) = x$ for all $x \in M$.*

Proof. By the uniqueness of splitting fields, it is enough to show that f splits completely in L_1.

Suppose that $G = \mathrm{Gal}(L_1/K)$, and put $g(X) = \prod_{s\in G}(X - s(\theta))$. Then g is invariant under G, and so lies in $K[X]$. Also, $g(\theta) = 0$, so that $f \mid g$. Since g splits completely over L_1, so does f. □

Definition. This extension L/K is a (or the, by abuse of notation) *Galois closure* of M/K.

Theorem 6.3. *Suppose that L/K is a Galois extension and that $f \in K[X]$ is irreducible. Then f splits completely in L if it has a root in L.*

Proof. Suppose that $\alpha \in L$ is a root of f. Note that $f(s(\alpha)) = 0$ for all $s \in G$. Suppose that $\beta_1, \dots, \beta_r$ are the distinct elements $s(\alpha)$, where s runs over the elements of G. Put $g = \prod_i (X - \beta_i)$. Then g is G-invariant, and so lies in $K[X]$. Moreover, g divides f in $L[X]$, since every root of g is a root of f and g has no repeated roots. Hence, by Lemma 3.1, g divides f in $K[X]$. Since f is irreducible, $f = g$ and we are done. □

(To be continued.)

Some facts about finite groups.

This is a list of some definitions and theorems concerning finite groups that we need for the Galois theory course, but whose proof requires more time than we have.

(1) Symmetric groups S_n.

S_n is generated by the set of transpositions, and by $\{(12),\ (23),\ (34),\dots,\ (n-1,n)\}$. S_n is also generated by any pair consisting of a transposition and an n–cycle.

A subgroup H of S_n is *transitive* if H permutes $\{1,\dots,n\}$ transitively. The only transitive subgroup of S_3 is A_3. The transitive subgroups of S_4 are A_4, D_8, C_4, V. Here V is the group whose non–identity elements are the things whose cycle–type is (2)(2) and is isomorphic to $C_2 \times C_2$, C_4 is a cyclic group generated by a 4–cycle, D_8 is of order 8 and is the symmetry group of the square. S_4 permutes the elements of $V - \{1\}$ transitively, so that there is a homomorphism $\pi : S_4 \to S_3$ with $\ker \pi = V$ defined by this permutation action. Counting orders shows that π is surjective (so that, in the language of normal subgroups and quotient groups that is described below, $S_4/V \cong S_3$). The transitive subgroups of S_4 listed above (that is, those subgroups that permute 4 objects transitively) map to the subgroups $S_3, A_3, C_2, C_2, 1$ respectively, so that except for ambiguity between D_8 and C_4 they are distinguished by their images in S_3.

(2) Groups of order ≤ 8.

The Abelian groups of order ≤ 7 are either cyclic or $C_2 \times C_2$. The only non-Abelian group of order ≤ 7 is S_3. There are five groups of order 8: the Abelian groups C_8, $C_4 \times C_2$ and $C_2 \times C_2 \times C_2$ and the non-abelian groups D_8 and Q_8. Q_8 is the subset $\{\pm 1,\ \pm i,\ \pm j,\ \pm k\}$ of the quaternions. Recall that $i^2 = j^2 = k^2 = -1$ and $ij = k = -ji, jk = i = -kj, ki = j = -ik$. They can be distinguished e.g. by counting elements of order 2; D_8 has five and Q_8 has one. Also, D_8 is a subgroup of S_4 while Q_8 is not.

(3) Orbits and stabilizers.

If a group G acts on a set X, then the length of any orbit is the index of the stabilizer of any element of that orbit. Two elements of the same orbit have conjugate stabilizers.

(4) If H is a subgroup of G, then the order of H divides that of G, and the ratio is the index of H in G (the number of left cosets, or equivalently the number of right cosets). If p^n is a prime power dividing the order of G, then G has a subgroup of order p^n.

(5) Abelian groups.

An abelian group G is uniquely a direct product $C_1 \times \cdots \times C_r$ of cyclic groups such that the order of C_i divides that of C_{i+1}, and G is also uniquely a direct product $C'_1 \times \cdots \times C'_s$ of cyclic groups of prime power order. There is an integer n such that every element has order dividing n, and some element has order n.

Typeset by $\mathcal{A}\mathcal{M}\mathcal{S}$-TEX

(6) Normal subgroups and factor groups (= quotient groups).

If H is a subgroup of a group G,then H is *normal* if $sH = Hs$ for all $s \in G$, or equivalently $s^{-1}Hs = H$ for all s. If H is normal, then the set of left cosets $\{Hs \mid s \in G\}$ of H in G naturally forms a group, the *quotient group* G/H, with multiplication given by $(Hs)(Ht) = Hst$. There is a surjective homomorphism $\pi : G \to G/H$ with $\pi(s) = Hs$, and $\ker \pi = H$. Note that (assuming G to be finite) $\#(G/H) = \#(G)/\#(H)$.

For example, suppose that G acts on an object (geometric or algebraic), say X. By definition, this means that we are given a group homomorphism $\rho : G \to \mathrm{Aut}(X)$. Suppose that H acts trivially on X, i.e. that $\rho(h) = 1$ for all $h \in H$. Then the action of G on X naturally induces an action of G/H on X. That is, there is a homomorphism $\sigma : G/H \to \mathrm{Aut}(X)$ such that $\rho = \sigma \circ \pi$.

Crudely, the idea is that, given that H acts trivially, we can throw it away by passing to the smaller, and so less complicated, group G/H.

(7) Simple and soluble groups and composition series.

A group G is *simple* if every normal subgroup is either 1 or G. So an Abelian group is simple if and only if it is of prime order. See (7) below for other examples.

A *composition series* for G is a chain $1 = G_0 \subset G_1 \subset \cdots \subset G_m = G$ of subgroups of G such that each G_i is a normal subgroup of G_{i+1} (but not necessarily of G) and each composition factor G_{i+1}/G_i is simple. The Jordan-Hölder theorem states that if $1 = H_0 \subset H_1 \subset \cdots H_n = G$ is another composition series for G, then $m = n$ and the quotients G_{i+1}/G_i are equal to the quotients H_{j+1}/H_j in some order, but not necessarily the given one. That is, the group G determines its composition factors. Note that the composition factors do not determine the group; consider, for example, $C_2 \times C_2$ and C_4 or C_6 and S_3.

G is *soluble* or *solvable* if it has a composition series in which every factor is Abelian. Important examples are:

S_3. This has a composition series $1 \subset A_3 \subset S_3$.

S_4. This has a chain of normal subgroups $1 \subset V \subset A_4 \subset S_4$. The qotients are respectively V, C_3 and C_2.

The subgroup B of upper triangular matrices in the group $GL_2(R)$, where R is any finite ring (e.g. $R = \mathbb{Z}/n\mathbb{Z}$ or a finite field). The subgroup U consisting of those matrices with 1's on the diagonal is normal in B, $U \cong R$ (where R is the additive group) and $B/U \cong (R^*)^2$, where R^* is the group of multiplicative units units in R.

Any subgroup or quotient group of a soluble group. Conversely, if N is normal in G and N and G/N are both soluble, then so is G.

(8) Alternating groups.

The alternating group A_n is generated for example by the 3-cycles $(123), (124)$, $\ldots, (12n)$, and is simple. In consequence, S_5 is not soluble if $n \geq 5$.

Here is a proof of simplicity for $n = 5$ (for the general case see van der Waerden):

Recall from the IA Algebra and Geometry course that A_5 is the group of rotations of a regular icosahedron. Suppose that $g \in A_5$ with $g \neq 1$. Then g is determined by its axis A and its angle θ, where $0 < \theta \leq \pi$. If $\theta < \pi$ then we must also specify the direction of A. There are four possibilities:

(i) A passes through the midpoint of an edge. Then $\theta = \pi$ and g has order 2. There are 30 edges, so 15 such g and they are conjugate.
(ii) A passes through the midpoint of a face. Then $\theta = 2\pi/3$ and g has order 3. There are 20 faces, so 20 such g and they are conjugate.
(iii) A passes through two opposite vertices. Then $\theta = 2\pi/5$ or $4\pi/5$ and g has order 5. There are 12 vertices, so there is one conjugacy class consisting of the 12 such g for which $\theta = 2\pi/5$ and another conjugacy class consisting of the 12 for which $\theta = 4\pi/5$.

So A_5 has 5 conjugacy classes, with respectively $1, 15, 20, 12$ and 12 members.

Now suppose that H is a normal subgroup of A_5 of order n. Then by definition H is a union of conjugacy classes in A_5, so that $n = 1 + a.15 + b.20 + c.12 + d.12$, where $a, \ldots, d$ are 0 or 1. Also, n divides 60, by Lagrange's theorem. It is easy to see that the only solutions are $a = \cdots = d = 0$ and $a = \cdots = d = 1$, corresponding to $H = 1$ and $H = A_5$.