

1. PREPARATORY REMARKS

Fix $n \geq 1$ an integer. Let K_n/\mathbf{Q} be a splitting field of $X^n - 1$, so the group of n th roots of unity in K has order n (as \mathbf{Q} has characteristic not dividing n) and is cyclic (as is any finite subgroup of the multiplicative group of a field). By Exercise 3 in HW4, we have a natural injective homomorphism of groups

$$\text{Gal}(K_n/\mathbf{Q}) \hookrightarrow (\mathbf{Z}/n\mathbf{Z})^\times,$$

and it is an isomorphism when n is a prime power due to an argument with Eisenstein's criterion to prove that the polynomial $\Phi_{p^e}(X) = (X^{p^e} - 1)/(X^{p^{e-1}} - 1) = \Phi_p(X^{p^{e-1}})$ for $e \geq 1$ is *irreducible* over \mathbf{Q} . Concretely, the roots of Φ_{p^e} (in a splitting field over \mathbf{Q}) are clearly the full set of *primitive* p^e th roots of unity. Hence, the statement that the extension K_{p^e}/\mathbf{Q} is “as big as is possible” really says that all primitive p^e th roots of unity are Galois conjugates of each other over \mathbf{Q} (i.e., have the same minimal polynomial over \mathbf{Q}). This is a very *special* property of \mathbf{Q} .

For example, over \mathbf{R} it is clear for odd $n > 1$ that the $\varphi(n)$ primitive n th roots of unity break up into $\varphi(n)/2$ $\text{Gal}(\mathbf{C}/\mathbf{R})$ -conjugate pairs. Likewise, over \mathbf{F}_p (as well as many other interesting fields) one cannot say that “all primitive n th roots of unity are created equal”: they might have *different* minimal polynomials over the ground field. For example, if $n|(p-1)$ then \mathbf{F}_p^\times contains a full set of n th roots of unity, so these are most definitely not “created equal”. As one instance, \mathbf{F}_7 contains 6 distinct 6th roots of unity; the primitive cube roots of unity in \mathbf{F}_7 are 2 and 4 (i.e., $X^2 + X + 1 = (X-2)(X-4)$ is a factorization of Φ_3 in $\mathbf{F}_7[X]$).

The aim of this handout is to show that $\text{Gal}(K_n/\mathbf{Q}) \rightarrow (\mathbf{Z}/n\mathbf{Z})^\times$ is an isomorphism even when n is not a prime power. We cannot hope to use Eisenstein's criterion, so we shall require another method (due to Gauss) which proceeds by a totally different approach and in particular proves the irreducibility of Φ_{p^e} in $\mathbf{Q}[X]$ by methods unrelated to Eisenstein's criterion. The key to Gauss' idea is to exploit the magical properties of the p th power map in $\mathbf{F}_p[X]$ for an auxiliary prime p . Since the number of primitive n th roots of unity in K_n is exactly $\varphi(n) = |(\mathbf{Z}/n\mathbf{Z})^\times|$ (as for the number of generators of any cyclic group of order n), and a Galois conjugate of a primitive n th root of unity is again a primitive n th root of unity (why?), to say they are Galois conjugates of each other over \mathbf{Q} is to say that one of them (and hence all of them) has minimal polynomial over \mathbf{Q} of degree $|(\mathbf{Z}/n\mathbf{Z})^\times|$. Any primitive n th root of unity ζ in K_n actually generates K_n over \mathbf{Q} , so $[K_n : \mathbf{Q}]$ is the degree of the minimal polynomial of ζ over \mathbf{Q} . Since $[K_n : \mathbf{Q}] = |\text{Gal}(K_n/\mathbf{Q})|$, we conclude that $\text{Gal}(K_n/\mathbf{Q}) \hookrightarrow (\mathbf{Z}/n\mathbf{Z})^\times$ is an isomorphism *if and only if* all primitive n th roots of unity in K_n are Galois conjugate over \mathbf{Q} (i.e., have the same minimal polynomial). This is what allows one to say that “all primitive n th roots of unity are created equal over \mathbf{Q} ”; there is no *algebraic* way to distinguish them from each other once it is seen that $\text{Gal}(K_n/\mathbf{Q})$ acts transitively on this set (i.e., they're all roots of the same minimal polynomial over \mathbf{Q}). The examples in the preceding paragraph show that this fails over other ground fields (such as \mathbf{R} and finite fields).

As an important special case of the role of transitive Galois group actions to eliminate the possibility of algebraic distinguishability, when $n = 4$, one *cannot* say in any intrinsic sense that there is a “preferred” solution to $T^2 + 1 = 0$ in \mathbf{C} (view \mathbf{C} as simply an abstract choice of splitting field of $T^2 + 1$ over \mathbf{R}). It can be shown that there is only one absolute value on \mathbf{C} extending the usual one on \mathbf{R} , and this is how we give \mathbf{C} a topological structure of normed field. Since the topology of \mathbf{C} determines the subfield \mathbf{R} as the closure of \mathbf{Q} , this subfield is intrinsic to \mathbf{C} with its topology, and $\text{Gal}(\mathbf{C}/\mathbf{R})$ permutes the two solutions to $T^2 + 1 = 0$. No canonical mathematical construction should depend on a choice of solution to this equation (though constructions may use such a choice if the end result is independent of the choice, such as in the statement of the residue formula in complex analysis).

2. MAIN RESULT

With the motivation and background set up, it is time to prove the main result. We begin with some notation. Fix $n \geq 1$ and K_n/\mathbf{Q} a splitting field of $X^n - 1$. Define

$$\Phi_n(X) = \prod (X - \zeta) \in K_n[X],$$

where ζ runs over all *primitive* n th roots of unity in K_n (i.e., all generators of the intrinsic order n *cyclic* group of solutions to $T^n - 1 = 0$ in K_n). It is clear from the intrinsic nature of primitive n th roots of unity that the action of $\text{Gal}(K_n/\mathbf{Q})$ permutes these around. Hence, even without knowing if $\text{Gal}(K_n/\mathbf{Q})$ is “big”, it is clear that the monic polynomial $\Phi_n(X)$ is *invariant* under the action of $\text{Gal}(K_n/\mathbf{Q})$ (which just shuffles its linear factors under the natural action on $K_n[X]$). Hence, *by Galois theory* (!) the coefficients of Φ_n must lie in \mathbf{Q} ! Its degree is clearly $|(\mathbf{Z}/n\mathbf{Z})^\times|$. The main aim is therefore to prove:

Theorem 2.1. (Gauss) *The polynomial $\Phi_n \in \mathbf{Q}[X]$ is irreducible.*

Proof. By construction, $\Phi_n \in \mathbf{Q}[X]$ is monic, and over the extension field K_n we see that Φ_n divides $X^n - 1$. Since formation of gcd commutes with extension of the ground field, the divisibility $\Phi_n | (X^n - 1)$ in $\mathbf{Q}[X]$ must hold because it is true in $K_n[X]$ (i.e., Φ_n serves as a gcd of Φ_n and $X^n - 1$). By Gauss’ Lemma, since $X^n - 1 \in \mathbf{Q}[X]$ has integral coefficients, any *monic* factorization in $\mathbf{Q}[X]$ is necessarily in $\mathbf{Z}[X]$. That is, if we write $X^n - 1 = \Phi_n h$ with $h \in \mathbf{Q}[X]$, then since h is visibly monic (as $X^n - 1$ and Φ_n are monic) it follows that both Φ_n and h must lie in $\mathbf{Z}[X]$.

Now suppose that Φ_n is *not* irreducible in $\mathbf{Q}[X]$, so there is a factorization $\Phi_n = fg$ in $\mathbf{Q}[X]$ with monic f and g of positive degree. We may also suppose f is *irreducible*. By Gauss’ Lemma applied to the monic factorization $fg = \Phi_n$ with $\Phi_n \in \mathbf{Z}[X]$, we must have $f, g \in \mathbf{Z}[X]$. We seek to derive a contradiction. In $K_n[X]$ we have the monic factorization $\Phi_n = \prod (X - \zeta)$ where the product runs over all *primitive* n th roots of unity in K_n . Since f and g both have positive degree, there must exist *distinct* primitive n th roots of unity ζ and ζ' in K_n such that $X - \zeta$ is a factor of f and $X - \zeta'$ is a factor of g ; that is, $f(\zeta) = 0$ and $g(\zeta') = 0$ in K_n .

We can write $\zeta' = \zeta^r$ for some unique $r \in (\mathbf{Z}/n\mathbf{Z})^\times$ since ζ and ζ' are both primitive n th roots of unity. Since $\zeta \neq \zeta'$, we must have $r \neq 1$. Choose a positive integer representing this residue class r , and again denote it by r , so $r > 1$ and $\text{gcd}(r, n) = 1$. Consider the prime factorization $r = \prod p_j$ with primes p_j not necessarily pairwise distinct. To go from ζ to $\zeta' = \zeta^r$ we successively raise to exponents p_1 , then p_2 , etc. Since $f(\zeta) = 0$ and $g(\zeta') = 0$, so $f(\zeta') \neq 0$ and $g(\zeta) \neq 0$ (as the factorization $\Phi_n = fg$ and separability of Φ_n forces f and g to have no common roots), there must exist a least j for which $\zeta^{p_1 \cdots p_{j-1}}$ is a root of f and its p_j th power is a root of g . In other words, we can find a primitive n th root of unity ζ_0 and a prime p not dividing n such that $f(\zeta_0) = 0$ and $g(\zeta_0^p) = 0$. From this fact we shall deduce a contradiction.

Since f is irreducible over \mathbf{Q} , it must be the minimal polynomial of ζ_0 . But $g(\zeta_0^p) = 0$, so $g(X^p) \in \mathbf{Q}[X]$ has ζ_0 as a root. Thus, $f|g(X^p)$ is $\mathbf{Q}[X]$. We can therefore write $g(X^p) = fq$ in $\mathbf{Q}[X]$, with q necessarily monic (as $g(X^p)$ and f are monic). Since $g(X^p)$ has coefficients in \mathbf{Z} , Gauss’ Lemma once again ensures that $q \in \mathbf{Z}[X]$. Thus, the identity $g(X^p) = fq$ takes place in $\mathbf{Z}[X]$. Now reduce mod p ! In $\mathbf{F}_p[X]$, we get

$$\bar{f}\bar{q} = \bar{g}(X^p) = \bar{g}(X)^p,$$

the final equality using the fact that $a^p = a$ for all $a \in \mathbf{F}_p$. The monicity of f and g , coupled with the condition that both have positive degree, ensures that $\bar{f}, \bar{g} \in \mathbf{F}_p[X]$ have *positive degree* (though may well be reducible). From the divisibility relation $\bar{f}|\bar{g}^p$ we conclude that \bar{f} and \bar{g} must have a non-trivial irreducible factor in common. Hence, the product $\bar{f}\bar{g}$ has a non-trivial irreducible factor appearing with multiplicity more than 1. But in $\mathbf{Q}[X]$ we have $fg = \Phi_n | (X^n - 1)$, so by Gauss’ Lemma we even get $fg | (X^n - 1)$ in $\mathbf{Z}[X]$, so we may reduce mod p to get $\bar{f}\bar{g} | (X^n - 1)$ in $\mathbf{F}_p[X]$. It follows that $X^n - 1 \in \mathbf{F}_p[X]$ has a non-trivial square factor and hence is not separable. But this is absurd, since p doesn’t divide n and hence the derivative test ensures that $X^n - 1 \in \mathbf{F}_p[X]$ is separable! Contradiction. ■

Note that although the Φ_n ’s were constructed abstractly, they’re easy to compute explicitly in $\mathbf{Z}[X]$. In fact, since each n th root of unity is a primitive d th root of unity for a unique $d|n$, we obviously have $X^n - 1 = \prod_{d|n} \Phi_d(X)$. This allows a computer to recursively compute $\Phi_n(X)$ by long division using $X^n - 1$ and $\Phi_d(X)$ ’s for proper divisors d of n ; the abstract theory is needed to prove that when Φ_d is known for all $d|n$ with $d < n$, then $\prod_{d|n, d < n} \Phi_d$ actually divides $X^n - 1$ in $\mathbf{Q}[X]$; by thinking about the root structure in a splitting field, such a divisibility claim is obvious (without thinking about root structure in a big extension, such issues become quite tedious to handle).