

7 Galois extensions, Frobenius elements, and the Artin map

In our standard *AKLB* setup, A is a Dedekind domain with fraction field K , L/K is a finite separable extension, and B is the integral closure of A in L , a Dedekind domain with fraction field L . We now consider the case where L/K is also normal, hence Galois, and let $G := \text{Gal}(L/K)$ denote the Galois group; we will use *AKLBG* to denote this setup.

7.1 Splitting primes in Galois extensions

We begin by showing that the Galois group G acts on the ideal group \mathcal{I}_B (the invertible, equivalently, nonzero, fractional ideals of B) and that this action is compatible with the group structure of \mathcal{I}_B . More precisely, \mathcal{I}_B is a left G -module.

Definition 7.1. Let G be a group. A *left G -module* is an abelian group M equipped with a left G -action that commutes with its group operation; in additive notation we have $\sigma(a + b) = \sigma(a) + \sigma(b)$ for all $\sigma \in G$ and $a, b \in M$. One similarly defines a *right G -module* as an abelian group with a right G -action that commutes with the group operation.

Theorem 7.2. *Assume AKLBG. For each fractional ideal I of B and $\sigma \in G$ define*

$$\sigma(I) := \{\sigma(x) : x \in I\}.$$

The set $\sigma(I)$ is a fractional ideal of B , and this defines a group action on \mathcal{I}_B that makes it a left G -module. Moreover, the restriction of this action to $\text{Spec } B$ makes it a G -set.

Proof. We first show that $\sigma(B) = B$ for all $\sigma \in G$. Each $b \in B$ is integral over A , hence $f(b) = 0$ for some monic polynomial $f \in A[x]$, and we have

$$0 = \sigma(0) = \sigma(f(b)) = f(\sigma(b)),$$

so $\sigma(b)$ is also integral over A , hence an element of B , since B is the integral closure of A in L . This proves $\sigma(B) \subseteq B$, and the same argument shows $\sigma^{-1}(B) \subseteq B$, hence $B \subseteq \sigma(B)$ and therefore $\sigma(B) = B$ as claimed.

Each $\sigma \in G = \text{Gal}(L/K)$ is a field automorphism of L and thus commutes with addition and multiplication. It follows that if $I \subseteq L$ is a finitely generated B -module (a fractional ideal) then $\sigma(I)$ is a finitely generated $\sigma(B)$ -module, and $\sigma(B) = B$, so $\sigma(I)$ is a finitely generated B -module, hence a fractional ideal as claimed. We clearly have $\sigma((0)) = (0)$ for all $\sigma \in G$, so G permutes \mathcal{I}_B , the group of nonzero fractional ideals. We also have

$$(\sigma\tau)(I) = \{(\sigma\tau)(x) : x \in I\} = \{\sigma(\tau(x)) : x \in I\} = \{\sigma(y) : y \in \tau(I)\} = \sigma(\tau(I)),$$

and the identity clearly acts trivially, so we have a left G -action on \mathcal{I}_B .

Now let $I, J \in \mathcal{I}_B$ and $\sigma \in G$. Each $x \in IJ$ has the form $x = a_1b_1 + \cdots + a_nb_n$ with $a_i \in I$ and $b_i \in J$, and $\sigma(x) = \sigma(a_1)\sigma(b_1) + \cdots + \sigma(a_n)\sigma(b_n) \in \sigma(I)\sigma(J)$. Thus $\sigma(IJ) \subseteq \sigma(I)\sigma(J)$, and applying the same argument to $\sigma(I), \sigma(J)$, and σ^{-1} implies $\sigma^{-1}(\sigma(I)\sigma(J)) \subseteq IJ$ and therefore $\sigma(I)\sigma(J) \subseteq \sigma(IJ)$. Thus $\sigma(IJ) = \sigma(I)\sigma(J)$ for all $I, J \in \mathcal{I}_B$, implying that \mathcal{I}_B is a left G -module.

Let \mathfrak{p} be a prime of B and let $\sigma(\mathfrak{p}) = \mathfrak{q}_1^{e_1} \cdots \mathfrak{q}_n^{e_n}$ be the unique factorization of $\sigma(\mathfrak{p})$ in B . Applying σ^{-1} to both sides yields $\mathfrak{p} = \sigma^{-1}(\mathfrak{q}_1)^{e_1} \cdots \sigma^{-1}(\mathfrak{q}_n)^{e_n}$, and therefore $n = 1$ and $e_1 = 1$, since \mathfrak{p} is prime, thus $\sigma(\mathfrak{p}) = \mathfrak{q}_1$ is prime and the G -action on \mathcal{I}_B restricts to a G -action on $\text{MaxSpec } B$, and on $\text{Spec } B$, since G fixes $\{(0)\} = \text{Spec } B - \text{MaxSpec } B$. \square

Recall that by a prime of A (or K) we mean a nonzero prime ideal of A , and similarly for B (and L), and for any prime \mathfrak{p} of A we use $\{\mathfrak{q}|\mathfrak{p}\}$ to denote the set of primes \mathfrak{q} that lie above \mathfrak{p} (equivalently, for which $\mathfrak{p} = A \cap \mathfrak{q}$); in other words, $\{\mathfrak{q}|\mathfrak{p}\}$ is the fiber of the contraction map $\text{MaxSpec } B \rightarrow \text{MaxSpec } A$ above \mathfrak{p} .

Corollary 7.3. *Assume AKLBG. For each prime \mathfrak{p} of A the group G acts transitively on the set $\{\mathfrak{q}|\mathfrak{p}\}$; in other words, the orbits of the G -action on $\text{Spec } B$ are the fibers of the contraction map $\text{Spec } B \rightarrow \text{Spec } A$.*

Proof. Consider any $\sigma \in G$. For $\mathfrak{q}|\mathfrak{p}$ we have $\mathfrak{p}B \subseteq \mathfrak{q}$ and $\sigma(\mathfrak{p}B) \subseteq \sigma(\mathfrak{q})$, so $\sigma(\mathfrak{q})|\mathfrak{p}$ (note $\sigma(\mathfrak{p}B) = \mathfrak{p}B$ and in a Dedekind domain, to contain is to divide). Thus $\{\mathfrak{q}|\mathfrak{p}\}$ is closed under the action of G , we just need to show that it consists of a single orbit.

Let $\{\mathfrak{q}|\mathfrak{p}\} = \{\mathfrak{q}_1, \dots, \mathfrak{q}_n\}$ and suppose that \mathfrak{q}_1 and \mathfrak{q}_2 lie in distinct G -orbits. The primes $\mathfrak{q}_1, \dots, \mathfrak{q}_n$ are maximal ideals, hence pairwise coprime, so by the CRT we have a ring isomorphism

$$\frac{B}{\mathfrak{q}_1 \cdots \mathfrak{q}_n} \simeq \frac{B}{\mathfrak{q}_1} \times \cdots \times \frac{B}{\mathfrak{q}_n},$$

and we may choose $b \in B$ such that $b \equiv 0 \pmod{\mathfrak{q}_2}$ and $b \equiv 1 \pmod{\sigma^{-1}(\mathfrak{q}_1)}$ for all $\sigma \in G$ (by hypothesis, $\sigma(\mathfrak{q}_2) \neq \mathfrak{q}_1$ for all $\sigma \in G$, since $\mathfrak{q}_1, \mathfrak{q}_2$ lie in different G -orbits). Then $b \in \mathfrak{q}_2$ and

$$N_{L/K}(b) = \prod_{\sigma \in G} \sigma(b) \equiv 1 \pmod{\mathfrak{q}_1},$$

hence $N_{L/K}(b) \notin A \cap \mathfrak{q}_1 = \mathfrak{p}$. But $N_{L/K}(b) \in N_{L/K}(\mathfrak{q}_2) = \mathfrak{p}^{f_{\mathfrak{q}_2}} \subseteq \mathfrak{p}$, a contradiction. \square

As shown in the proof of Theorem 7.2, we have $\sigma(B) = B$ for all $\sigma \in G = \text{Gal}(L/K)$, thus each $\sigma \in G$ restricts to a ring automorphism of B that fixes every element of the subring $A = B \cap K$, and thus every element of any prime \mathfrak{p} of A . It follows that σ induces an isomorphism of residue field extensions $\bar{\sigma} \in \text{Hom}_{A/\mathfrak{p}}(B/\mathfrak{q}, B/\sigma(\mathfrak{q}))$ defined by $\bar{\sigma}(x + \mathfrak{q}) := \sigma(x) + \sigma(\mathfrak{q})$ for $x \in B$, which we may more compactly write as $\bar{\sigma}(\bar{x}) := \overline{\sigma(x)}$ (but note that the \bar{x} and $\sigma(x)$ are typically elements of different residue fields).

Corollary 7.4. *Assume AKLBG and let \mathfrak{p} be a prime of A . The residue field degrees $f_{\mathfrak{q}} := [B/\mathfrak{q} : A/\mathfrak{p}]$ are the same for every $\mathfrak{q}|\mathfrak{p}$, as are the ramification indices $e_{\mathfrak{q}} := v_{\mathfrak{q}}(\mathfrak{p}B)$.*

Proof. For each $\sigma \in G$ we have an isomorphism of the residue fields B/\mathfrak{q} and $B/\sigma(\mathfrak{q})$ that fixes A/\mathfrak{p} , so they clearly have the same degree $f_{\mathfrak{q}} = f_{\sigma(\mathfrak{q})}$, and G acts transitively on $\{\mathfrak{q}|\mathfrak{p}\}$, by Corollary 7.3, so the function $\mathfrak{q} \mapsto f_{\mathfrak{q}}$ must be constant on $\{\mathfrak{q}|\mathfrak{p}\}$.

For each $\sigma \in G$ we also have $\sigma(\mathfrak{p}) = \mathfrak{p}$ and $\sigma(B) = B$, so $\sigma(\mathfrak{p}B) = \mathfrak{p}B$, and for each $\mathfrak{q}|\mathfrak{p}$,

$$e_{\mathfrak{q}} = v_{\mathfrak{q}}(\mathfrak{p}B) = v_{\mathfrak{q}}(\sigma(\mathfrak{p}B)) = v_{\mathfrak{q}}\left(\sigma\left(\prod_{\mathfrak{r}|\mathfrak{p}} \mathfrak{r}^{e_{\mathfrak{r}}}\right)\right) = v_{\mathfrak{q}}\left(\prod_{\mathfrak{r}|\mathfrak{p}} \sigma(\mathfrak{r})^{e_{\mathfrak{r}}}\right) = v_{\mathfrak{q}}\left(\prod_{\mathfrak{r}|\mathfrak{p}} \mathfrak{r}^{e_{\sigma^{-1}(\mathfrak{r})}}\right) = e_{\sigma^{-1}(\mathfrak{q})}.$$

The transitivity of the G -action on $\{\mathfrak{q}|\mathfrak{p}\}$ again implies that $\mathfrak{q} \mapsto e_{\mathfrak{q}}$ is constant on $\{\mathfrak{q}|\mathfrak{p}\}$. \square

Corollary 7.4 implies that whenever L/K is Galois, we may unambiguously write $e_{\mathfrak{p}}$ and $f_{\mathfrak{p}}$ instead of $e_{\mathfrak{q}}$ and $f_{\mathfrak{q}}$; recall that we previously defined $g_{\mathfrak{p}} := \#\{\mathfrak{q}|\mathfrak{p}\}$.

Corollary 7.5. *Assume AKLBG. For each prime \mathfrak{p} of A we have $e_{\mathfrak{p}}f_{\mathfrak{p}}g_{\mathfrak{p}} = [L:K]$.*

Proof. This follows immediately from Theorem 5.43 and Corollary 7.4. \square

Example 7.6. Assume *AKLBG*. When $n := [L : K]$ is prime there are just three ways a prime \mathfrak{p} of A can split in B :

- $e_{\mathfrak{p}} = n$ and $f_{\mathfrak{p}} = g_{\mathfrak{p}} = 1$, in which case \mathfrak{p} is totally ramified in L ;
- $f_{\mathfrak{p}} = n$ and $e_{\mathfrak{p}} = g_{\mathfrak{p}} = 1$, in which case \mathfrak{p} remains inert in L if $B/\mathfrak{p}B$ is finite étale;
- $g_{\mathfrak{p}} = n$ and $e_{\mathfrak{p}} = f_{\mathfrak{p}} = 1$, in which case \mathfrak{p} splits completely in L if $B/\mathfrak{p}B$ is finite étale.

Recall from Definition 5.45 that we only defined the terms “remains inert” and “splits completely” for unramified primes, which includes the condition that all the residue field extensions B/\mathfrak{q} of A/\mathfrak{p} are separable, equivalently, that $B/\mathfrak{p}B$ is finite étale over A/\mathfrak{p} . This will automatically hold in the primary case of interest to us, where the residue field A/\mathfrak{p} is finite, hence perfect, and all residue field extensions are separable.

7.2 Decomposition and inertia groups

Definition 7.7. Assume *AKLBG*. For each prime \mathfrak{q} of B the *decomposition group* $D_{\mathfrak{q}}$ (also denoted $D_{\mathfrak{q}}(L/K)$) is the stabilizer of \mathfrak{q} in G .

Lemma 7.8. Assume *AKLBG* and let \mathfrak{p} be a prime of A . The decomposition groups $D_{\mathfrak{q}}$ for $\mathfrak{q}|\mathfrak{p}$ are all conjugate in G , with $\#D_{\mathfrak{q}} = e_{\mathfrak{p}}f_{\mathfrak{p}}$ and $[G : D_{\mathfrak{q}}] = g_{\mathfrak{p}}$.

Proof. Points in an orbit of group action have conjugate stabilizers, so the $D_{\mathfrak{q}}$ for $\mathfrak{q}|\mathfrak{p}$ are all conjugate, by Corollary 7.3. The orbit-stabilizer theorem implies $[G : D_{\mathfrak{q}}] = \#\{\mathfrak{q}|\mathfrak{p}\} = g_{\mathfrak{p}}$. We have $\#G = [L : K] = e_{\mathfrak{p}}f_{\mathfrak{p}}g_{\mathfrak{p}}$, by Corollary 7.5, so $\#D_{\mathfrak{q}} = \#G/[G : D_{\mathfrak{q}}] = e_{\mathfrak{p}}f_{\mathfrak{p}}$. \square

Let us now consider a particular prime $\mathfrak{q}|\mathfrak{p}$ of B (by writing $\mathfrak{q}|\mathfrak{p}$ we define \mathfrak{p} as $\mathfrak{q} \cap A$). As noted above, each $\sigma \in G$ induces a residue field isomorphism $\bar{\sigma} \in \text{Hom}_{A/\mathfrak{p}}(B/\mathfrak{q}, B/\sigma(\mathfrak{q}))$. For $\sigma \in D_{\mathfrak{q}}$, we have $\sigma(\mathfrak{q}) = \mathfrak{q}$, in which case $\bar{\sigma} \in \text{Aut}_{A/\mathfrak{p}}(B/\mathfrak{q})$. Moreover, the map $\sigma \mapsto \bar{\sigma}$ defines a group homomorphism $\pi_{\mathfrak{q}} : D_{\mathfrak{q}} \rightarrow \text{Aut}_{A/\mathfrak{p}}(B/\mathfrak{q})$, since for any $x \in B$ we have

$$\overline{\sigma\tau}(\bar{x}) = \overline{\sigma\tau(x)} = \overline{\sigma(\tau(x))} = \bar{\sigma}(\overline{\tau(x)}) = \bar{\sigma}(\bar{\tau}(\bar{x})).$$

Note that B/\mathfrak{q} need not be a Galois extension of A/\mathfrak{p} even when L is a Galois extension of K , which is why we write $\text{Aut}_{A/\mathfrak{p}}(B/\mathfrak{q})$ and not $\text{Gal}((B/\mathfrak{q})/(A/\mathfrak{p}))$.

Proposition 7.9. Assume *AKLBG* and let $\mathfrak{q}|\mathfrak{p}$ be a prime of B . The group homomorphism $\pi_{\mathfrak{q}} : D_{\mathfrak{q}} \rightarrow \text{Aut}_{A/\mathfrak{p}}(B/\mathfrak{q})$ defined by $\sigma \mapsto \bar{\sigma}$ is surjective and B/\mathfrak{q} is normal over A/\mathfrak{p} .

Proof. Let F be the separable closure of A/\mathfrak{p} in B/\mathfrak{q} and for $\bar{b} \in F$, pick $b \in B$ such that $b \equiv \bar{b} \pmod{\mathfrak{q}}$ and $b \equiv 0 \pmod{\sigma^{-1}(\mathfrak{q})}$ (so $\sigma(b) \equiv 0 \pmod{\mathfrak{q}}$) for all $\sigma \in G - D_{\mathfrak{q}}$; the CRT implies that such an b exists, since for $\sigma \in G - D_{\mathfrak{q}}$ the ideals \mathfrak{q} and $\sigma(\mathfrak{q})$ are distinct and therefore coprime (since they are maximal ideals). Now define

$$g(x) := \prod_{\sigma \in G} (x - \sigma(b)) \in A[x],$$

and let \bar{g} denote the image of g in $(A/\mathfrak{p})[x]$. Observe that \bar{b} is a root of a polynomial $\bar{g} \in (A/\mathfrak{p})[x]$ that splits completely in $(B/\mathfrak{q})[x]$, and our choice of \bar{b} was arbitrary, so this applies to every $\bar{b} \in F^{\times}$. It follows that F is a normal (hence Galois) extension of A/\mathfrak{p} , and we have $\text{Gal}(F/(A/\mathfrak{p})) \simeq \text{Aut}_{A/\mathfrak{p}}(B/\mathfrak{q})$, since F is the separable closure of A/\mathfrak{p} in B/\mathfrak{q} .

For each $\sigma \in G - D_{\mathfrak{q}}$ we have $\bar{\sigma}(\bar{b}) = 0$, so 0 is a root of $\bar{g}(x)$ with multiplicity at least $m = \#(G - D_{\mathfrak{q}})$, and the remaining roots are $\bar{\sigma}(\bar{b})$ for $\sigma \in D_{\mathfrak{q}}$, all of which are $\text{Gal}(F/(A/\mathfrak{p}))$ -conjugates of \bar{b} . It follows that $\bar{g}(x)/x^m$ divides a power of the minimal polynomial $f(x)$ of \bar{b} , but $f(x)$ is irreducible in $(A/\mathfrak{p})[x]$, so $\bar{g}(x)/x^m$ is a power of $f(x)$ and every $\text{Gal}(F/(A/\mathfrak{p}))$ -conjugate of \bar{b} has the form $\bar{\sigma}(\bar{b})$ for some $\sigma \in D_{\mathfrak{q}}$. Applying this to \bar{b} chosen so that $F = (A/\mathfrak{p})(\bar{b})$ (by the primitive element theorem) shows that the map $\pi_{\mathfrak{q}}: D_{\mathfrak{q}} \rightarrow \text{Aut}_{A/\mathfrak{p}}(B/\mathfrak{q}) \simeq \text{Gal}(F/(A/\mathfrak{p}))$ is surjective.

To show that B/\mathfrak{q} is a normal extension of A/\mathfrak{p} we proceed as we did for F : for each $b \in B$ define $g \in A[x]$ and $\bar{g} \in (A/\mathfrak{p})[x]$ as above to show that every $\bar{b} \in B/\mathfrak{q}$ is the root of a polynomial in $(A/\mathfrak{p})[x]$ that splits completely in $(B/\mathfrak{q})[x]$. \square

Definition 7.10. Assume *AKLBG*, and let $\mathfrak{q}|\mathfrak{p}$ be a prime of B . The kernel of the surjective homomorphism $\pi_{\mathfrak{q}}: D_{\mathfrak{q}} \rightarrow \text{Aut}_{A/\mathfrak{p}}(B/\mathfrak{q})$ is the *inertia group* $I_{\mathfrak{q}}$ of \mathfrak{q} .

Corollary 7.11. Assume *AKLBG* and let $\mathfrak{q}|\mathfrak{p}$ be a prime of B . We have an exact sequence

$$1 \longrightarrow I_{\mathfrak{q}} \longrightarrow D_{\mathfrak{q}} \longrightarrow \text{Aut}_{A/\mathfrak{p}}(B/\mathfrak{q}) \longrightarrow 1,$$

and $\#I_{\mathfrak{q}} = e_{\mathfrak{p}}[B/\mathfrak{q} : A/\mathfrak{p}]_i$.

We have shown that the residue field B/\mathfrak{q} is always a normal extension of the residue field A/\mathfrak{p} . Let us now suppose that it is also separable, hence Galois; this holds, for example, if A/\mathfrak{p} is a perfect field, and in particular, whenever A/\mathfrak{p} is a finite field. We then have

$$D_{\mathfrak{q}}/I_{\mathfrak{q}} \simeq \text{Aut}_{A/\mathfrak{p}}(B/\mathfrak{q}) = \text{Gal}((B/\mathfrak{q})/(A/\mathfrak{p})).$$

Proposition 7.12. Assume *AKLBG*, let $\mathfrak{q}|\mathfrak{p}$ be a prime of B , and suppose B/\mathfrak{q} is a separable extension of A/\mathfrak{p} . We have a tower of field extensions $K \subseteq L^{D_{\mathfrak{q}}} \subseteq L^{I_{\mathfrak{q}}} \subseteq L$ with

$$\begin{aligned} e_{\mathfrak{p}} &= [L : L^{I_{\mathfrak{q}}}] = \#I_{\mathfrak{q}}; \\ f_{\mathfrak{p}} &= [L^{I_{\mathfrak{q}}} : L^{D_{\mathfrak{q}}}] = \#D_{\mathfrak{q}}/\#I_{\mathfrak{q}}; \\ g_{\mathfrak{p}} &= [L^{D_{\mathfrak{q}}} : K] = \#\{\mathfrak{q}|\mathfrak{p}\}. \end{aligned}$$

The fields $L^{D_{\mathfrak{q}}}$ and $L^{I_{\mathfrak{q}}}$ are the *decomposition field* and *inertia field* associated to \mathfrak{q} .

Proof. The third equality follows immediately from Lemma 7.8. The second follows from Proposition 7.9 and the separability of $(B/\mathfrak{q})/(A/\mathfrak{p})$, since $D_{\mathfrak{q}}/I_{\mathfrak{q}} \simeq \text{Gal}((B/\mathfrak{q})/(A/\mathfrak{p}))$ has cardinality $f_{\mathfrak{p}} = [B/\mathfrak{q} : A/\mathfrak{p}]$. We then have $[L : L^{D_{\mathfrak{q}}}] = \#D_{\mathfrak{q}} = e_{\mathfrak{p}}f_{\mathfrak{p}}$ and $\#D_{\mathfrak{q}}/\#I_{\mathfrak{q}} = f_{\mathfrak{p}}$, so $\#I_{\mathfrak{q}} = e_{\mathfrak{p}}$, so the first equality also holds. \square

We now consider an intermediate field E lying between K and L . Let us fix a prime $\mathfrak{q}|\mathfrak{p}$ of B , and let $\mathfrak{q}_E := \mathfrak{q} \cap E$, so that $\mathfrak{q}|\mathfrak{q}_E$ and $\mathfrak{q}_E|\mathfrak{p}$, and let us use $\overline{G}_{\mathfrak{q}}(L/K) := \text{Aut}_{A/\mathfrak{p}}(B/\mathfrak{q})$, $\overline{G}_{\mathfrak{q}}(L/E) := \text{Aut}_{(B \cap E)/\mathfrak{q}_E}(B/\mathfrak{q})$, $\overline{G}_{\mathfrak{q}_E}(E/K) := \text{Aut}_{A/\mathfrak{p}}((B \cap E)/\mathfrak{q}_E)$ to denote the automorphism groups of the residue field extensions associated to the tower $K \subseteq E \subseteq L$. We use the notation $D_{\mathfrak{q}}(L/E)$ to denote the decomposition group of \mathfrak{q} relative to the extension L/E (note that L/E is Galois since L/K is), and similarly define $D_{\mathfrak{q}}(L/K)$, as well as $I_{\mathfrak{q}}(L/E)$ and $I_{\mathfrak{q}}(L/K)$. In the case that E/K is also Galois, we similarly use $D_{\mathfrak{q}_E}(E/K)$ and $I_{\mathfrak{q}_E}(E/K)$ to denote the decomposition and inertia group of \mathfrak{q}_E (subgroups of $\text{Gal}(E/K)$).

Proposition 7.13. *Assume AKLBG, let E be an intermediate field between K and L . Let \mathfrak{q} be a prime of B and let $\mathfrak{q}_E = \mathfrak{q} \cap E$ and $\mathfrak{p} = \mathfrak{q} \cap K$. Then*

$$I_{\mathfrak{q}}(L/E) = I_{\mathfrak{q}}(L/K) \cap \text{Gal}(L/E) \quad \text{and} \quad D_{\mathfrak{q}}(L/E) = D_{\mathfrak{q}}(L/K) \cap \text{Gal}(L/E).$$

If E/K is Galois, then we have the following commutative diagram of exact sequences:

$$\begin{array}{ccccccc}
 & & 1 & & 1 & & 1 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 1 & \longrightarrow & I_{\mathfrak{q}}(L/E) & \longrightarrow & I_{\mathfrak{q}}(L/K) & \longrightarrow & I_{\mathfrak{q}_E}(E/K) \longrightarrow 1 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 1 & \longrightarrow & D_{\mathfrak{q}}(L/E) & \longrightarrow & D_{\mathfrak{q}}(L/K) & \longrightarrow & D_{\mathfrak{q}_E}(E/K) \longrightarrow 1 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 1 & \longrightarrow & \overline{G}_{\mathfrak{q}}(L/E) & \longrightarrow & \overline{G}_{\mathfrak{q}}(L/K) & \longrightarrow & \overline{G}_{\mathfrak{q}_E}(E/K) \longrightarrow 1 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 & & 1 & & 1 & & 1
 \end{array}$$

Proof. Note that $D_{\mathfrak{q}}(L/E) \subseteq \text{Gal}(L/E) \subseteq \text{Gal}(L/K)$. An element σ of $\text{Gal}(L/K)$ lies in $D_{\mathfrak{q}}(L/E)$ if and only if it fixes E (hence lies in $\text{Gal}(L/E)$) and satisfies $\sigma(\mathfrak{q}) = \mathfrak{q}$ (hence lies in $D_{\mathfrak{q}}(L/K)$), which proves the first claim. For the second claim, the restriction of $\pi_{\mathfrak{q}}(L/K): D_{\mathfrak{q}}(L/K) \rightarrow \overline{G}_{\mathfrak{q}}(L/K)$ to $D_{\mathfrak{q}}(L/E)$ is the map $\pi_{\mathfrak{q}}(L/E): D_{\mathfrak{q}}(L/E) \rightarrow \overline{G}_{\mathfrak{q}}(L/E)$, hence the kernels agree after intersecting with $\text{Gal}(L/E)$.

The exactness of the columns follows from Corollary 7.11; we now argue exactness of the rows. Each row corresponds to an inclusion followed by a restriction in which the inclusion is precisely the kernel of the restriction (for the first two rows this follows from the two claims proved above and for the third row it follows from Proposition 7.9, since normal subextension of an algebraic extension are stable under automorphisms); exactness at the first two groups in each row follows. Surjectivity of the restriction maps follows from the bijection used in the proof of Lemma 4.10. We have a bijection $\text{Hom}_K(L, \Omega) \rightarrow \text{Hom}_E(L, \Omega) \times \text{Hom}_K(E, \Omega)$ whose second factor is restriction, and we may view this as a bijection $\phi: \text{Gal}(L/K) \rightarrow \text{Gal}(L/E) \times \text{Gal}(E/K)$. If $\sigma \in \text{Gal}(E/K)$ stabilizes \mathfrak{q}_E then $\phi^{-1}(1, \sigma) \in \text{Gal}(L/K)$ stabilizes \mathfrak{q} and restricts to σ ; this implies surjectivity of the restriction maps in the first two rows, and for the third we replace $L/E/K$ with the corresponding tower of residue field extensions (and forget about stabilizing \mathfrak{q}_E).

We now argue commutativity of the four corner squares which suffices to prove the commutativity of the entire diagram. The upper left square commutes because all the maps are inclusions. The upper right square commutes because inclusion and restriction commute. The lower left square commutes because the horizontal maps are inclusions and the vertical maps coincide on $D_{\mathfrak{q}}(L/E)$. The lower right square commutes because the horizontal maps are restrictions and the vertical maps agree after restriction to E . \square

Corollary 7.14. *Assume AKLBG, let E be an intermediate field between K and L . Let \mathfrak{q} be a prime of B and let $\mathfrak{q}_E = \mathfrak{q} \cap E$ and $\mathfrak{p} = \mathfrak{q} \cap K$. Then*

- $e_{\mathfrak{q}_E/\mathfrak{p}} = 1$ if and only if $E \subseteq L^{I_{\mathfrak{q}}}$, and
- $e_{\mathfrak{q}_E/\mathfrak{p}} = f_{\mathfrak{q}_E/\mathfrak{p}} = 1$ if and only if $E \subseteq L^{D_{\mathfrak{q}}}$,

where $I_{\mathfrak{q}}$ and $D_{\mathfrak{q}}$ are the inertia and decomposition groups of \mathfrak{q} .

Proof. Proposition 7.13 implies $I_{\mathfrak{q}}(L/E) = I_{\mathfrak{q}}(L/K) \cap \text{Gal}(L/E)$, and for $F = L^{I_{\mathfrak{q}}}$, we have $I_{\mathfrak{q}}(L/F) = I_{\mathfrak{q}}(L/K) = \text{Gal}(L/F)$. We also have $\text{Gal}(L/EF) = \text{Gal}(L/E) \cap \text{Gal}(L/F)$, so

$$I_{\mathfrak{q}}(L/E) = I_{\mathfrak{q}}(L/K) \cap \text{Gal}(L/E) = \text{Gal}(L/F) \cap \text{Gal}(L/EF) = \text{Gal}(L/EF) = I_{\mathfrak{q}}(L/EF).$$

Now $e_{\mathfrak{q}/\mathfrak{q}_E} = \#I_{\mathfrak{q}}(L/E) = \#I_{\mathfrak{q}}(L/EF) = e_{\mathfrak{q}/\mathfrak{q}_{EF}}$ is equal to $e_{\mathfrak{q}/\mathfrak{q}_F} = \#I_{\mathfrak{q}}(L/F)$ if and only if $E \subseteq F$. The first claim in the corollary follows, since $I_{\mathfrak{q}}(L/F) = I_{\mathfrak{q}}(L/K)$ implies $e_{\mathfrak{q}/\mathfrak{q}_F} = e_{\mathfrak{q}/\mathfrak{p}}$ which implies $e_{\mathfrak{q}_F/\mathfrak{p}} = 1$, since $e_{\mathfrak{q}/\mathfrak{q}_F} e_{\mathfrak{q}_F/\mathfrak{p}} = e_{\mathfrak{q}/\mathfrak{p}}$, by Lemma 5.38.

The proof of the second claim follows *mutatis mutandis*: replace $I_{\mathfrak{q}}$ by $D_{\mathfrak{q}}$ and $e_{\mathfrak{q}/x}$ by $e_{\mathfrak{q}/x} f_{\mathfrak{q}/x}$ throughout. \square

In our AKLBG setup, for any prime \mathfrak{p} of K we let $I_{\mathfrak{p}}$ and $D_{\mathfrak{p}}$ denote the subgroups of $G = \text{Gal}(L/K)$ generated by the inertia groups $I_{\mathfrak{q}}$ and the decomposition groups $D_{\mathfrak{q}}$ of the primes $\mathfrak{q}|\mathfrak{p}$, respectively, which we call the *inertia group* and *decomposition group* of \mathfrak{p} . The corresponding inertia field $L^{I_{\mathfrak{p}}}$ and decomposition field $L^{D_{\mathfrak{p}}}$ are Galois extensions of K that are characterized by the following corollary.

Corollary 7.15. *Assume AKLBG and let \mathfrak{p} be a prime of K . The fields $L^{I_{\mathfrak{p}}}$ and $L^{D_{\mathfrak{p}}}$ are Galois extensions of K , and for any intermediate field E we have $e_{\mathfrak{q}_E/\mathfrak{p}} = 1$ for all $\mathfrak{q}_E|\mathfrak{p}$ if and only if $E \subseteq L^{I_{\mathfrak{p}}}$, and $e_{\mathfrak{q}_E/\mathfrak{p}} = f_{\mathfrak{q}_E/\mathfrak{p}} = 1$ for all $\mathfrak{q}_E|\mathfrak{p}$ if and only if $E \subseteq L^{D_{\mathfrak{p}}}$.*

When A/\mathfrak{p} is a perfect field, the inertia field is the largest subfield of L in which \mathfrak{p} is unramified, and the decomposition field is the largest subfield in which \mathfrak{p} splits completely.

Proof. The fact that G acts transitively on $\{\mathfrak{q}|\mathfrak{p}\}$ means that $I_{\mathfrak{p}}$ is generated by a complete set of conjugate subgroups $I_{\mathfrak{q}}$ and is therefore stable under conjugation, hence normal, and similarly for $D_{\mathfrak{p}}$. It follows that the corresponding fixed fields are Galois extensions of K . The rest of the corollary follows immediately from Corollary 7.14. \square

7.3 Frobenius elements

We now add the further assumption that the residue fields A/\mathfrak{p} (and therefore B/\mathfrak{q}) are finite for all primes \mathfrak{p} of K .¹ This holds, for example, whenever K is a global field (a finite extension of \mathbb{Q} or $\mathbb{F}_q(t)$). In this situation B/\mathfrak{q} is necessarily a Galois extension of A/\mathfrak{p} (we don't need Proposition 7.9 for this, finite extensions of finite fields are always Galois). Indeed, recall that every finite extension of a finite field \mathbb{F} has a cyclic Galois group generated by the $\#\mathbb{F}$ -power Frobenius automorphism $x \mapsto x^{\#\mathbb{F}}$.

In order to simplify the notation, when working with finite residue fields we may write $\mathbb{F}_{\mathfrak{q}} := B/\mathfrak{q}$ and $\mathbb{F}_{\mathfrak{p}} := A/\mathfrak{p}$; these are finite fields of p -power order, where p is the characteristic of $\mathbb{F}_{\mathfrak{p}}$ (and of $\mathbb{F}_{\mathfrak{q}}$). Note that the field K (and L) need not have characteristic p (consider the case of number fields), but if the characteristic of K is positive then it must be p (consider the homomorphism $A \rightarrow A/\mathfrak{p}$ from the integral domain A to the field A/\mathfrak{p}).

Let $\mathfrak{q}|\mathfrak{p}$ be a prime of B . Corollary 7.11 gives us an exact sequence

$$1 \longrightarrow I_{\mathfrak{q}} \longrightarrow D_{\mathfrak{q}} \xrightarrow{\pi_{\mathfrak{q}}} \text{Gal}(\mathbb{F}_{\mathfrak{q}}/\mathbb{F}_{\mathfrak{p}}) \longrightarrow 1.$$

If \mathfrak{p} (equivalently, \mathfrak{q}) is unramified, then $e_{\mathfrak{p}} = e_{\mathfrak{q}} = 1$ and $I_{\mathfrak{q}}$ is trivial. In this case we have an isomorphism

$$\pi_{\mathfrak{q}}: D_{\mathfrak{q}} \xrightarrow{\sim} \text{Gal}(\mathbb{F}_{\mathfrak{q}}/\mathbb{F}_{\mathfrak{p}}).$$

¹There exist Dedekind domains A (PIDs even) with a mixture of finite and infinite residue fields; see [1].

The Galois group $\text{Gal}(\mathbb{F}_q/\mathbb{F}_p)$ is the cyclic group of order $f_p = [\mathbb{F}_q : \mathbb{F}_p]$ generated by the Frobenius automorphism

$$x \mapsto x^{\#\mathbb{F}_p}.$$

Note that the cardinality of the finite field \mathbb{F}_p is necessarily a power of its characteristic p . If $K = \mathbb{Q}$ and $\mathfrak{p} = (p)$ is a prime of \mathbb{Z} , then $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$ is the field with p elements, but in general the field \mathbb{F}_p need not be a prime field (consider $K = \mathbb{Q}(i)$ and $\mathfrak{p} = (7)$).

Definition 7.16. Assume *AKLBG* with finite residue fields and $\mathfrak{q}|\mathfrak{p}$ unramified. The inverse image of the Frobenius automorphism of $\text{Gal}(\mathbb{F}_q/\mathbb{F}_p)$ under $\pi_{\mathfrak{q}} : D_{\mathfrak{q}} \xrightarrow{\sim} \text{Gal}(\mathbb{F}_q/\mathbb{F}_p)$ is the *Frobenius element* $\sigma_{\mathfrak{q}} \in D_{\mathfrak{q}} \subseteq G$ (also called the *Frobenius substitution* [2, §8]).

Proposition 7.17. Assume *AKLBG* with finite residue fields and $\mathfrak{q}|\mathfrak{p}$ unramified. The Frobenius element $\sigma_{\mathfrak{q}}$ is the unique $\sigma \in G$ such that for all $x \in B$ we have

$$\sigma(x) \equiv x^{\#\mathbb{F}_p} \pmod{\mathfrak{q}}.$$

Proof. Clearly $\sigma_{\mathfrak{q}}$ has this property, we just need to show uniqueness. Suppose $\sigma \in G$ has the desired property. For any $x \in \mathfrak{q}$ we have $x \equiv 0 \pmod{\mathfrak{q}}$, and $\sigma(x) \equiv x^{\#\mathbb{F}_p} \pmod{\mathfrak{q}}$ implies $\sigma(x) \equiv 0 \pmod{\mathfrak{q}}$, so $\sigma(x) \in \mathfrak{q}$; it follows that $\sigma(\mathfrak{q}) = \mathfrak{q}$, and therefore $\sigma \in D_{\mathfrak{q}}$. The isomorphism $\pi_{\mathfrak{q}} : D_{\mathfrak{q}} \rightarrow \text{Gal}(\mathbb{F}_q/\mathbb{F}_p)$ maps both σ and $\sigma_{\mathfrak{q}}$ to the Frobenius automorphism $x \mapsto x^{\#\mathbb{F}_p}$, so we must have $\sigma = \sigma_{\mathfrak{q}}$. \square

Proposition 7.18. Assume *AKLBG* with finite residue fields and let \mathfrak{p} be an unramified prime of A . The set of Frobenius elements $\{\sigma_{\mathfrak{q}} : \mathfrak{q}|\mathfrak{p}\}$ is a conjugacy class of G .

Proof. Let $\sigma_{\mathfrak{q}}$ be a Frobenius element, let C be its conjugacy class, let $\mathfrak{q}'|\mathfrak{p}$, and let $\tau \in G$ satisfy $\mathfrak{q}' = \tau(\mathfrak{q})$ (the existence of τ is guaranteed by Corollary 7.3). For any $x \in B$ we have

$$\begin{aligned} \sigma_{\mathfrak{q}}(x) &\equiv x^{\#\mathbb{F}_p} \pmod{\mathfrak{q}} \\ \tau(\sigma_{\mathfrak{q}}(x)) &\equiv \tau\left(x^{\#\mathbb{F}_p}\right) \pmod{\tau(\mathfrak{q})} \\ (\tau\sigma_{\mathfrak{q}})(x) &\equiv \tau(x)^{\#\mathbb{F}_p} \pmod{\mathfrak{q}'} \\ (\tau\sigma_{\mathfrak{q}})(\tau^{-1}(x)) &\equiv \tau(\tau^{-1}(x))^{\#\mathbb{F}_p} \pmod{\mathfrak{q}'} \\ (\tau\sigma_{\mathfrak{q}}\tau^{-1})(x) &\equiv x^{\#\mathbb{F}_p} \pmod{\mathfrak{q}'}, \end{aligned}$$

where we applied τ to both sides in the second line and replaced x by $\tau^{-1}(x)$ in the fourth line. Proposition 7.17 implies $\sigma_{\mathfrak{q}'} = \tau\sigma_{\mathfrak{q}}\tau^{-1} \in C$.

Now consider any $g \in C$, pick $\tau \in G$ so that $\tau\sigma_{\mathfrak{q}}\tau^{-1} = g$, and let $\mathfrak{q}' := \tau(\mathfrak{q})$. We have $g(x) = (\tau\sigma_{\mathfrak{q}}\tau^{-1})(x) \equiv x^{\#\mathbb{F}_p} \pmod{\mathfrak{q}'}$ for all $x \in B$ by the argument above, and Proposition 7.17 implies that $g = \sigma_{\mathfrak{q}'}$ is a Frobenius element. \square

Definition 7.19. Assume *AKLBG* with finite residue fields and let \mathfrak{p} be an unramified prime of A . The G -conjugacy class $\{\sigma_{\mathfrak{q}} : \mathfrak{q}|\mathfrak{p}\}$ is the *Frobenius class* of \mathfrak{p} , denoted $\text{Frob}_{\mathfrak{p}}$.

It is common to abuse terminology and refer to $\text{Frob}_{\mathfrak{p}}$ as a Frobenius element $\sigma_{\mathfrak{p}} \in G$ representing its conjugacy class (so $\sigma_{\mathfrak{p}} = \sigma_{\mathfrak{q}}$ for some $\mathfrak{q}|\mathfrak{p}$); there is little risk of confusion so long as we remember that $\sigma_{\mathfrak{p}}$ is only determined up to conjugacy (which usually governs all the properties we care about). But there is one situation where this terminology is entirely correct. If G is abelian then each conjugacy class consists of a single element, in which case $\text{Frob}_{\mathfrak{p}} = \{\sigma_{\mathfrak{q}} : \mathfrak{q}|\mathfrak{p}\}$ is a singleton set and there is a unique choice for $\sigma_{\mathfrak{p}}$ (note that $\#\{\sigma_{\mathfrak{q}} : \mathfrak{q}|\mathfrak{p}\} = 1$ does not imply $\#\{\mathfrak{q}|\mathfrak{p}\} = 1$ because the map $\mathfrak{q} \rightarrow \sigma_{\mathfrak{q}}$ need not be injective).

7.4 Artin symbols

There is another notation commonly used to denote Frobenius elements that includes the field extension in the notation.

Definition 7.20. Assume *AKLBG* with finite residue fields. For each unramified prime \mathfrak{q} of L we define the *Artin symbol*

$$\left(\frac{L/K}{\mathfrak{q}}\right) := \sigma_{\mathfrak{q}}.$$

Proposition 7.21. Assume *AKLBG* with finite residue fields and $\mathfrak{q}|\mathfrak{p}$ unramified. Then \mathfrak{p} splits completely if and only if $\left(\frac{L/K}{\mathfrak{q}}\right) = 1$.

Proof. This follows directly from the definitions: if \mathfrak{p} splits completely then $e_{\mathfrak{p}}f_{\mathfrak{p}} = 1$ and $D_{\mathfrak{q}} = \langle \sigma_{\mathfrak{q}} \rangle = \{1\}$. Conversely, if $D_{\mathfrak{q}} = \langle \sigma_{\mathfrak{q}} \rangle = \{1\}$ then $e_{\mathfrak{p}}f_{\mathfrak{p}} = 1$ and \mathfrak{p} splits completely. \square

We will see later in the course that the extension L/K is completely determined by the set of primes \mathfrak{p} that split completely in L . Thus in some sense the Artin symbol captures the essential structure of L/K .

Proposition 7.22. Assume *AKLBG* with finite residue fields and let $\mathfrak{q}|\mathfrak{p}$ be unramified. Let E be an intermediate field between K and L , and define $\mathfrak{q}_E := \mathfrak{q} \cap E$. Then

$$\left(\frac{L/E}{\mathfrak{q}}\right) = \left(\frac{L/K}{\mathfrak{q}}\right)^{[\mathbb{F}_{\mathfrak{q}_E}:\mathbb{F}_{\mathfrak{p}}]},$$

and if E/K is Galois then $\left(\frac{E/K}{\mathfrak{q}_E}\right)$ is the restriction of $\left(\frac{L/K}{\mathfrak{q}}\right)$ to E .

Proof. For the first claim, note that $\#\mathbb{F}_{\mathfrak{q}_E} = (\#\mathbb{F}_{\mathfrak{p}})^{[\mathbb{F}_{\mathfrak{q}_E}:\mathbb{F}_{\mathfrak{p}}]}$. The second claim follows from the commutativity of the lower right square in the commutative diagram of Proposition 7.13: the Frobenius automorphism $x \mapsto x^{\#\mathbb{F}_{\mathfrak{p}}}$ of $\text{Gal}(\mathbb{F}_{\mathfrak{q}_E}/\mathbb{F}_{\mathfrak{p}})$ is the restriction of the Frobenius automorphism $x \mapsto x^{\#\mathbb{F}_{\mathfrak{p}}}$ of $\text{Gal}(\mathbb{F}_{\mathfrak{q}}/\mathbb{F}_{\mathfrak{p}})$ to $\mathbb{F}_{\mathfrak{q}_E}$. \square

When L/K is abelian, the Artin symbol takes the same value for all $\mathfrak{q}|\mathfrak{p}$ and we may instead write

$$\left(\frac{L/K}{\mathfrak{p}}\right) := \sigma_{\mathfrak{p}}.$$

In this setting we now view the Artin symbol as a function mapping unramified primes \mathfrak{p} to Frobenius elements $\sigma_{\mathfrak{p}} \in G$. We wish to extend this map to a multiplicative homomorphism from the ideal group \mathcal{I}_A to the Galois group $G = \text{Gal}(L/K)$, but ramified primes $\mathfrak{q}|\mathfrak{p}$ cause problems: the homomorphism $\pi_{\mathfrak{q}}: D_{\mathfrak{q}} \rightarrow \text{Gal}(\mathbb{F}_{\mathfrak{q}}/\mathbb{F}_{\mathfrak{p}})$ is not a bijection when \mathfrak{p} is ramified (it has nontrivial kernel $I_{\mathfrak{q}}$ of order $e_{\mathfrak{q}} = e_{\mathfrak{p}} > 1$).

For any set S of primes of A , let \mathcal{I}_A^S denote the subgroup of \mathcal{I}_A generated by the primes of A that do not lie in S (a free abelian group).

Definition 7.23. Let A be a Dedekind domain with finite residue fields. Let L be a finite abelian extension of $K = \text{Frac } A$, and let S be the set of primes of A that ramify in L . The *Artin map* is the homomorphism

$$\begin{aligned} \left(\frac{L/K}{\cdot}\right) : \mathcal{I}_A^S &\rightarrow \text{Gal}(L/K) \\ \prod_{i=1}^m \mathfrak{p}_i^{e_i} &\mapsto \prod_{i=1}^m \left(\frac{L/K}{\mathfrak{p}_i}\right)^{e_i}. \end{aligned}$$

Remark 7.24. We will prove in later lectures that the set S of ramified primes is finite, but the definition makes sense in any case, and more generally, for any set S that contains all ramified primes, a fact that we will use later.

One of the main results of class field theory is that the Artin map is surjective (this is part of what is known as *Artin reciprocity*). This is a deep theorem that we are not yet ready to prove, but we can verify that it holds in some simple examples.

Example 7.25 (Quadratic fields). Let $K = \mathbb{Q}$ and $L = \mathbb{Q}(\sqrt{d})$ for some square-free integer $d \neq 1$. Then $\text{Gal}(L/K)$ has order 2 and is certainly abelian. As you proved on Problem Set 2, the only ramified primes $\mathfrak{p} = (p)$ of $A = \mathbb{Z}$ are those that divide the *discriminant*

$$D := \text{disc}(L/K) = \begin{cases} d & \text{if } d \equiv 1 \pmod{4}, \\ 4d & \text{if } d \not\equiv 1 \pmod{4}. \end{cases}$$

If we identify $\text{Gal}(L/K)$ with the multiplicative group $\{\pm 1\}$, then

$$\left(\frac{L/K}{\mathfrak{p}}\right) = \left(\frac{\mathbb{Q}(\sqrt{d})/\mathbb{Q}}{(p)}\right) = \left(\frac{D}{p}\right) = \pm 1,$$

where $\left(\frac{D}{p}\right)$ is the *Kronecker symbol*. For odd primes $p \nmid D$ we have

$$\left(\frac{D}{p}\right) = \begin{cases} +1 & \text{if } D \text{ is a nonzero square modulo } p, \\ -1 & \text{if } D \text{ is not a square modulo } p, \end{cases}$$

and for $p = 2$ not dividing D (in which case $D = d \equiv 1 \pmod{4}$) we have

$$\left(\frac{D}{2}\right) = \begin{cases} +1 & \text{if } D \equiv 1 \pmod{8}, \\ -1 & \text{if } D \equiv 5 \pmod{8}. \end{cases}$$

The cyclotomic extensions $\mathbb{Q}(\zeta_n)/\mathbb{Q}$ provide another interesting example that you will have an opportunity to explore on Problem Set 4.

References

- [1] Raymond C. Heitmann, *PID's with specified residue fields*, Duke Math. J. **41** (1974), 565–582.
- [2] Jean-Pierre Serre, *Local fields*, Springer, 1979.