Algebra
Prof. Fabien Morel
WS2025/2026

Written Notes by Leon Man

Contents

Fehler gerne bei leon.man@campus.lmu.de melden :)

1 Groups

1.1 Generalities

Definition 1.1.1. Let X be a set. An **internal law of composition** on a set X is a map:

$$\varphi : X \times X \to X$$

Definition 1.1.2. A **Magma** $(X, \circ)$ is a set X with an internal law of composition $\circ$ on X.

Definition 1.1.3. A **Category** $\mathcal{C}$ is a triple consisting of:

1. $Ob(\mathcal{C})$: A class of objects of $\mathcal{C}$.
2. For $B, C \in Ob(\mathcal{C})$, $\mathrm{Hom}_{\mathcal{C}}(B, C)$ is the set of morphisms from B to C.
3. Composition law: $\circ : \mathrm{Hom}_{\mathcal{C}}(A, B) \times \mathrm{Hom}_{\mathcal{C}}(B, C) \to \mathrm{Hom}_{\mathcal{C}}(A, C)$ is an associative map.
4. Identity: $\mathrm{id}_B \in \mathrm{Hom}_{\mathcal{C}}(B, B)$.

Let $A, B \in Ob(\mathcal{C}), f \in \mathrm{Hom}_{\mathcal{C}}(A, B)$.

1. f is called a **monomorphism** if for all objects Z and morphisms $g_1, g_2 : Z \to A$:
$$f \circ g_1 = f \circ g_2 \implies g_1 = g_2.$$
2. f is called an **epimorphism** if for all objects Z and morphisms $h_1, h_2 : B \to Z$:
$$h_1 \circ f = h_2 \circ f \implies h_1 = h_2.$$
3. f is called an **isomorphism** if there exists a morphism $g : B \to A$ such that:
$$g \circ f = \mathrm{id}_A \quad \text{and} \quad f \circ g = \mathrm{id}_B.$$
4. We denote the set of morphisms from an object to itself, $\mathrm{Hom}_{\mathcal{C}}(A, A)$, by $\mathrm{End}_{\mathcal{C}}(A)$. Its elements are called **endomorphisms**.
5. The subset of isomorphisms in $\mathrm{End}_{\mathcal{C}}(A)$ is denoted by $\mathrm{Aut}_{\mathcal{C}}(A)$. Its elements are called **automorphisms**.

Definition 1.1.4. A **monoid** is an associative magma $(M, \cdot)$ with a neutral element e.

Remark 1.1.5. Let M be an associative magma. If there exists a neutral element e, it is unique.

Proof. If f is another, then $e = ef = f$. □

Definition 1.1.6. Let M be a monoid. An element $m \in M$ is said to be **invertible** if there exists $m^{-1} \in M$ such that $m \cdot m^{-1} = m^{-1} \cdot m = e$.

Lemma 1.1.7. Let M be a monoid, $m \in M$. If m is invertible, its inverse is unique.

Proof. If $\tilde{m}, m'$ are inverses of m, $\tilde{m} = \tilde{m}e_M = \tilde{m}(mm') = (\tilde{m}m)m' = e_M m' = m'$. □

Definition 1.1.8. A **group** G is a monoid in which every element is invertible.

In the following, let G denote a group.

Example 1.1.9. 1. If M is a monoid, let $M^\times = \{m \in M \mid m \text{ is invertible}\}$. Then $M^\times$ is a group. For instance, if $M = M_n(K)$ (matrices), then $M^\times = GL_n(K)$.

2. Let X be a set. The group $\mathrm{Aut}_{\mathbf{Set}}(X)$ with composition is denoted by $S(X)$, the symmetric group.

3. In a category $\mathcal{C}$, the group of automorphisms of an object C is $\mathrm{Aut}(C)$.

Definition 1.1.10. If G, H are groups, a map $\psi : G \to H$ is a **group homomorphism** if:

$$\psi(xy) = \psi(x)\psi(y) \quad \forall x, y \in G.$$

The class of groups with group homomorphisms as morphisms forms a category we denote by **Grp**.

Remark 1.1.11. Note that in groups, monomorphisms are exactly injective homomorphisms and epimorphisms are exactly surjective homomorphisms.

Lemma 1.1.12. Let $\psi : G \to H$ be a group homomorphism, $g \in G$. Then:

1. $\psi(e_G) = e_H$
2. $\psi(g^{-1}) = \psi(g)^{-1}$

Proof. 1. $e_G \cdot e_G = e_G \implies \psi(e_G) \cdot \psi(e_G) = \psi(e_G) \cdot e_H$.

2. $\forall g \in G : \psi(g) \cdot \psi(g^{-1}) = \psi(e_H) = e_G$.

□

1.2 Subgroups, Quotients

Definition 1.2.1. A subset H of a group G is called a **subgroup** if:

1. H is stable by multiplication: $\forall a, b \in H, ab \in H$.
2. $e \in H$.
3. $\forall h \in H, h^{-1} \in H$.

The induced law of composition $\cdot$ from G gives H a group structure.
There is always a canonical embedding monomorphism $H \hookrightarrow G$.
If H is a subgroup of G, we sometimes write $H \leq G$.

Lemma 1.2.2. Let $H \subseteq G$ be a nonempty subset.

$$H \text{ is a subgroup } \Leftrightarrow \forall a, b \in H : a(b^{-1}) \in H.$$

Proof. ($\implies$): trivial.
($\Leftarrow$): Let $x \in H$. Then, $e_G = xx^{-1} \in H$,
$\forall x \in H : e_G \cdot x^{-1} \in H$, and
$\forall x, y \in H : y^{-1} \in H \implies x \cdot (y^{-1})^{-1} = xy \in H$, so H is a subgroup. □

Proposition 1.2.3. Let $\psi : G \to H$ be a homomorphism.

1. $\mathrm{Im}(\psi) = \{\psi(g) \in H \mid g \in G\}$ is a subgroup of H.
2. $\ker(\psi) = \{g \in G \mid \psi(g) = e_H\}$ is a subgroup of G.
3. ψ is a monomorphism if and only if $\ker(\psi) = \{e_G\}$.

Proof. 1. Let $h_1, h_2 \in \mathrm{im}(\psi)$ with $h_i = \psi(g_i)$. Then $h_1 h_2^{-1} = \psi(g_1)\psi(g_2)^{-1} = \psi(g_1 g_2^{-1}) \in \mathrm{im}(\psi)$. Thus, $\mathrm{im}(\psi) \leq H$.

2. Let $x, y \in \ker(\psi)$. Then $\psi(xy^{-1}) = \psi(x)\psi(y)^{-1} = e_H e_H^{-1} = e_H$, implying $xy^{-1} \in \ker(\psi)$. Thus, $\ker(\psi) \leq G$.

3. ($\Longrightarrow$) trivial.

   ($\Longleftarrow$) $\psi(x) = \psi(y) \iff \psi(x)\psi(y)^{-1} = e_H \iff \psi(xy^{-1}) = e_H$. If $\ker(\psi) = \{e_G\}$, this implies $xy^{-1} = e_G$, so $x = y$.

□

Example 1.2.4.
1. $(\mathbb{Z}, +) \subset (\mathbb{Q}, +) \subset (\mathbb{R}, +) \subset (\mathbb{C}, +)$ are strictly speaking only monomorphisms, but we treat them as subgroups.

2. $sgn : S_n \to \{\pm 1\}$ is an epimorphism.

3. Any monomorphism $\varphi : G \hookrightarrow H$ induces an isomorphism $G \xrightarrow{\cong} \mathrm{im}(\varphi) \subseteq H$.

4. If $\varphi : H \to G$ is an isomorphism of groups, clearly, $\varphi^{-1} : G \to H$ is a group homomorphism, so φ is an isomorphism in the category of groups.

1.2.1 Subgroups Generated by a Family

Lemma 1.2.5. Let $(H_i)_{i \in I}$ be a family of subgroups of a group G. Then $H := \bigcap_{i \in I} H_i$ is a subgroup of G.

Proof. Since $e_G \in H_i$ for all $i \in I$, we have $e_G \in H$, so $H \neq \emptyset$. Let $x, y \in H$. Then $x, y \in H_i$ for every $i \in I$. Since each $H_i \leq G$, it follows that $xy^{-1} \in H_i$ for all $i \in I$, and thus $xy^{-1} \in H$. Therefore, $H \leq G$. □

Definition 1.2.6. Let $(g_i)_{i \in I}$ be a family of elements in G. We denote by $\langle (g_i) \rangle$ the intersection of all subgroups of G containing the elements (g_i). It is called the subgroup **generated by** (g_i). We say that (g_i) **generates** G if $\langle (g_i) \rangle = G$.

Remark 1.2.7. It is equivalent to say that $\langle (g_i)_{i \in I} \rangle$ is the subset of all elements in G that can be written as a product of elements (g_i) or (g_i^{-1})

Definition 1.2.8.
1. A group G is said to be of **finite type** if there exists a finite family generating G.

2. If G is generated by one element g, G is called **cyclic**.

3. If G is finite, $|G|$ is called the **order** of G. If $g \in G$, the order of g is $|\langle g \rangle|$.

4. For a group G, $g, h \in G$, we denote $ghg^{-1}h^{-1}$ by $[g, h]$.
   $\langle [g, h]_{g,h \in G} \rangle := [G, G]$ is called the **commutator subgroup**.

5. The **symmetric group** S_n is the group of permutations on the set $\{1, \ldots, n\}$ ($\mathrm{Aut}_{\mathbf{Set}}(\{1, \ldots, n\})$) with composition. It has order $n!$.
   It can be shown that $S_n = \langle (\tau_{i,i+1})_{i \in \{1,\ldots,n-1\}} \rangle$, with $\tau_{i,j}$ denoting the transposition of the i-th and j-th element.

1.2.2 Quotient of a Group by a Subgroup

Definition 1.2.9. Let H be a subgroup of G. Let $g \in G$. The set $gH = \{gh \mid h \in H\}$ is called the **left coset** associated with g.
We say that $(x, y) \in G$ are **right congruent modulo** H if

$$\exists h \in H : y = x \cdot h (\Leftrightarrow x^{-1}y \in H).$$

Lemma 1.2.10. For a subgroup $H \leq G$, right congruence modulo H defines an equivalence relation on G with left cosets as equivalence classes.

Proof. Reflexivity: For any $x \in G$, $x^{-1}x = e_G \in H$, so $x \sim x$.
Symmetry: If $x \sim y$, then $x^{-1}y \in H$. Since H is a subgroup, $(x^{-1}y)^{-1} = y^{-1}x \in H$, implying $y \sim x$.
Transitivity: If $x \sim y$ and $y \sim z$, then $x^{-1}y \in H$ and $y^{-1}z \in H$. By closure, $(x^{-1}y)(y^{-1}z) = x^{-1}z \in H$, so $x \sim z$.
The equivalence class of $g \in G$ is $[g] = \{x \in G \mid g \sim x\} = \{x \in G \mid g^{-1}x = h \in H\} = \{gh \mid h \in H\} = gH$. □

Definition 1.2.11. Let $H \leq G$. We denote by $G/H := \{gH \mid g \in G\}$ the quotient of G by the equivalence "right congruence modulo H".
We get a canonical surjective map $G \twoheadrightarrow G/H$, $g \mapsto gH$.

Definition 1.2.12. Let $H \leq G$. If G/H is finite, its cardinality is denoted by $[G : H]$, called the **index** of H in G.

Remark 1.2.13. We observe that for all $g \in G$, the map $\sigma_g : H \to gH, h \mapsto gh$ has the inverse $\sigma_{g^{-1}}$, so it is a bijection. It follows that all left cosets have the same cardinality.

Thus, we have shown:

Corollary 1.2.14 (Lagrange theorem)**.** Let G be a finite group with subgroup H. Then,

$$|G| = |H|[G : H].$$

Example 1.2.15.

1. Consider the subgroup $n\mathbb{Z} \subset \mathbb{Z}$. The quotient is $\mathbb{Z}/n\mathbb{Z}$, so the index is:

$$[\mathbb{Z} : n\mathbb{Z}] = n.$$

2. The index of the alternating group in the symmetric group is 2:

$$[S_n : A_n] = 2.$$

Proof: Consider the sign homomorphism $\text{sgn} : S_n \to \{\pm 1\}$. We know that $A_n = \ker(\text{sgn})$. We will later see that this induces a bijection

$$S_n/A_n \xrightarrow{\cong} \{\pm 1\}.$$

3. If $H \leq G$ is a subgroup of index 2 (i.e., $[G : H] = 2$), then H is a normal subgroup ($H \trianglelefteq G$). *Proof: exercises.*

4. Let $H \leq G$. Consider the inversion map:

$$\chi : G \xrightarrow{\cong} G, \quad x \mapsto x^{-1}.$$

We have $\chi(H) = H$. This map sends a left coset to a right coset.

G^{op}, the **opposite group** is defined by the set G equipped with the multiplication $(x, y) \mapsto yx$. The map χ defines an isomorphism $G \cong G^{op}$. Consequently, χ induces a bijection between the set of left cosets and the set of right cosets:

$$G/H \xrightarrow{\cong} H\backslash G \quad (\text{or } G^{op}/H^{op}).$$

1.3 Normal Subgroups, Quotient Groups

We begin by analyzing the structure of a group "collapse." Let $\pi : G \twoheadrightarrow H$ be an epimorphism of groups (a surjective homomorphism). This map "collapses" elements of G onto H. To understand the internal structure of this collapse, we examine the kernel:

$$K := \ker(\pi) = \{g \in G \mid \pi(g) = e_H\}.$$

Lemma 1.3.1. For every $g \in G$, the preimage of $\pi(g)$ satisfies:

$$\pi^{-1}(\pi(g)) = gK = Kg.$$

Proof. Let $x \in G$. We have:

$$\begin{aligned} x \in \pi^{-1}(\{\pi(g)\}) &\iff \pi(x) = \pi(g) \\ &\iff \pi(g^{-1}x) = e_H \\ &\iff g^{-1}x \in K \\ &\iff x \in gK. \end{aligned}$$

Thus $\pi^{-1}(\{\pi(g)\}) = gK$. The equality $\pi^{-1}(\{\pi(g)\}) = Kg$ follows analogously. Since the preimage is equal to both the left and right cosets, we conclude $gK = Kg$. □

Corollary 1.3.2. For every $h \in H$, $\pi^{-1}(h)$ is a unique equivalence class in the quotient set G/K (or $K\backslash G$). Consequently, the map π induces a natural bijection:

$$\bar{\pi} : H \xrightarrow{\sim} G/K.$$

Since H is a group, this bijection allows us to impose the group structure of H onto the set of cosets G/K. Thus, G/K becomes a group, and we obtain an isomorphism:

$$G/K \cong H.$$

More generally, since any homomorphism $\psi : G \to H$ is surjective onto its image, there is always a canonical isomorphism $G/\ker\psi \xrightarrow{\cong} \operatorname{im}\psi$. This result is known as the first isomorphism theorem.

Remark 1.3.3. If $\pi : G \twoheadrightarrow H$ is an epimorphism, then G is partitioned into disjoint subsets, each isomorphic to the kernel (as sets). We can write G as a disjoint union:

$$G = \coprod_{h \in H} \pi^{-1}(h) = \coprod_{\alpha \in G/K} \alpha.$$

Having observed that kernels always satisfy the property $gK = Kg$, we abstract this property to define a special class of subgroups.

Definition 1.3.4. A subgroup $N \subseteq G$ is said to be **normal** (denoted $N \trianglelefteq G$) if its left and right cosets coincide for all elements of G:

$$\forall g \in G : gN = Ng.$$

Remark 1.3.5. We can conclude that if N is the kernel of any homomorphism, then N is normal. Conversely, we will later see that every normal subgroup is the kernel of the natural projection map $G \to G/N$. Thus, normal subgroups are exactly the kernels of homomorphisms.

Lemma 1.3.6. Let $N \subseteq G$ be a subgroup. It is easy to see that the following are equivalent:

1. N is normal ($gN = Ng$ for all $g \in G$).

2. N is *invariant under conjugation:*

$$\forall g \in G : gNg^{-1} = N.$$

3. For all $g \in G$ and all $n \in N$:

$$gng^{-1} \in N.$$

Example 1.3.7. 1. $A_n \trianglelefteq S_n$, since it is the kernel of the sign homomorphism.
Similarly, over some field K, $\mathrm{SL}_n(K) \trianglelefteq \mathrm{GL}_n(K)$ is the kernel of the determinant.

2. Any subgroup of an abelian group is normal.

3. Commutator subgroups are normal.

   Proof. To prove normality, it suffices to show that the subgroup is invariant under conjugation. Let $g \in G$ and let $c = [x, y] = xyx^{-1}y^{-1}$ be a generator of $[G, G]$. Conjugating c by g, we obtain:

$$\begin{aligned} g[x,y]g^{-1} &= g(xyx^{-1}y^{-1})g^{-1} \\ &= (gxg^{-1})(gyg^{-1})(gx^{-1}g^{-1})(gy^{-1}g^{-1}) \\ &= (gxg^{-1})(gyg^{-1})(gxg^{-1})^{-1}(gyg^{-1})^{-1} \\ &= [gxg^{-1}, gyg^{-1}]. \end{aligned}$$

   Since gxg^{-1} and gyg^{-1} are elements of G, the result $[gxg^{-1}, gyg^{-1}]$ is itself a commutator and thus belongs to $[G, G]$.

   Since the conjugate of any generator is in $[G, G]$, and conjugation is a homomorphism, the conjugate of any product of commutators is also in $[G, G]$. Therefore, $g[G, G]g^{-1} \subseteq [G, G]$ for all $g \in G$, proving that $[G, G]$ is normal. □

4. Consider the subgroup of S_n that fixes the element n, defined as $H := \{\sigma \in S_n \mid \sigma(n) = n\}$. This subgroup is naturally isomorphic to S_{n-1}. For $n \geq 3$, H is not normal in S_n. To see this, choose $h = \tau_{n-2,n-1} \in H$ and a conjugator $g = \tau_{n-1,n} \in S_n \setminus H$. The conjugate is:

$$ghg^{-1} = \tau_{n-1,n} \circ \tau_{n-2,n-1} \circ \tau_{n-1,n} = \tau_{n-2,n}.$$

   Since the resulting transposition moves n (mapping $n \mapsto n - 2$), it is not in H. Thus, $gHg^{-1} \not\subseteq H$.

5. If $H \subseteq G$ is not necessarily normal, we may define

$$N_G(H) =: \{g \in G \mid gHg^{-1} = H\},$$

   the **normaliser** of H in G. It is the largest (not necessarily normal) subgroup of G in which H is normal.

1.3.1 Quotient Groups

Lemma 1.3.8. Let $H \subseteq G$ be a normal subgroup. Then for all $x, y \in G$:

$$(xH) \cdot (yH) = (xy)H.$$

Proof. We compute the set product $(xH)(yH) = \{xh_1yh_2 \mid h_1, h_2 \in H\}$. Since H is normal, $Hy = yH$, meaning for any $h_1 \in H$, there exists $h' \in H$ such that $h_1y = yh'$. Thus, an element xh_1yh_2 becomes $xyh'h_2$, which lies in $(xy)H$. Conversely, any element $(xy)h \in (xy)H$ can be written as $x \cdot 1 \cdot y \cdot h$, which is in $(xH)(yH)$. □

Theorem 1.3.9. Let $H \subseteq G$ be normal. The operation defined by $(xH)\cdot(yH) = (xy)H$ induces a group structure on the set of cosets G/H.

The canonical map $\pi : G \to G/H$ defined by $g \mapsto gH$ is a group epimorphism with kernel H. The group G/H is called the **quotient group** of G by H.

Proof. The previous lemma confirms the operation is well-defined on cosets. Associativity is inherited directly from G. The identity element is the coset H (since $(xH)(H) = xH$), and the inverse of xH is $x^{-1}H$ (since $(xH)(x^{-1}H) = H$).

The map π is a homomorphism by construction: $\pi(xy) = (xy)H = (xH)(yH) = \pi(x)\pi(y)$. Finally, $\ker \pi = \{x \in G \mid xH = H\} = H$. □

1.3.2 Conjugation and Automorphisms

Lemma 1.3.10. Let $g \in G$. The map $c_g : G \to G$ defined by $x \mapsto gxg^{-1}$ is an automorphism of G.

This defines a map $c : G \to \mathrm{Aut}_{\mathbf{Grp}}(G)$, which is a group homomorphism.

Proof. For $x, y \in G$, $c_g(xy) = c_g(x)c_g(y) = gxg^{-1}gyg^{-1} = gxyg^{-1} = c_g(xy)$, so c_g is a homomorphism. It clearly has the inverse $c_{g^{-1}}$, so it is bijective.

For $g, h, x \in G$, $c(gh)(x) = c_{gh}(x) = ghxh^{-1}g^{-1} = gc_h(x)g^{-1} = c_g \circ c_h(x)$, so c is a homomorphism. □

Definition 1.3.11. Let c be defined as above. We denote its image by $\mathrm{Inn}(G)$, the group of **inner automorphisms**.

We denote the quotient group $\mathrm{Aut}_{\mathbf{Grp}}(G)/\mathrm{Inn}(G)$ by $\mathrm{Out}(G)$, the group of **outer automorphisms**.

$\ker c \trianglelefteq G =: Z(G)$ is called the **center** of G. It is the set of elements that commute with all elements in G.

Remark 1.3.12. $\mathrm{Inn}(G)$ is a normal subgroup of G.

Proof. Let $\varphi \in \mathrm{Aut}(G)$ be an automorphism and let $c_g \in \mathrm{Inn}(G)$ be an inner automorphism. Then, for any $x \in G$:

$$\begin{aligned}(\varphi \circ c_g \circ \varphi^{-1})(x) &= \varphi(c_g(\varphi^{-1}(x))) \\ &= \varphi(g \cdot \varphi^{-1}(x) \cdot g^{-1}) \\ &= \varphi(g) \cdot \varphi(\varphi^{-1}(x)) \cdot \varphi(g^{-1}) \\ &= \varphi(g) \cdot x \cdot \varphi(g)^{-1} \\ &= c_{\varphi(g)}(x).\end{aligned}$$

Thus, $\varphi \circ c_g \circ \varphi^{-1} = c_{\varphi(g)}$. Since $c_{\varphi(g)} \in \mathrm{Inn}(G)$, the subgroup is normal. □

Example 1.3.13. 1. Given a group G with subgroup H, we define

$$C_G(H) := \{g \in G \mid \forall h \in H : gh = hg\} \supseteq Z(H).$$

It is the largest subgroup of G whose elements commute with every element of H and is called the **centraliser** of H in G.

2. In $S_n (n \geq 3)$ The center is trivial, so $S_n \cong \mathrm{Inn}(S_n)$. For $n \neq 6$, every automorphism is inner, so $S_n \cong \mathrm{Aut}(S_n)$. For $n = 6$, there exists an exceptional outer automorphism, so $S_n \subsetneq \mathrm{Aut}(S_6)$.

3. In A_n, The center is trivial, so $A_n \cong \mathrm{Inn}(A_n)$. However, conjugation by an odd permutation in S_n (e.g., $(1\,2)$) induces an automorphism of A_n that is not inner to A_n. Thus $\mathrm{Out}(A_n) \cong \mathbb{Z}/2$.

4. For some field K, in $GL_n(K)$, The center consists of scalar matrices $K^\times \cdot I$. The inner automorphism group is the **projective general linear group:**

$$\mathrm{Inn}(GL_n(K)) \cong GL_n(K)/Z(GL_n(K)) =: PGL_n(K).$$

Since there are automorphisms that are not conjugations (e.g., $A \mapsto (A^T)^{-1}$), we have $\mathrm{Inn}(GL_n(K)) \subsetneq \mathrm{Aut}(GL_n(K))$.

1.3.3 Universal Property of the Quotient

Theorem 1.3.14. Let $H \trianglelefteq G$. Let $\varphi : G \to M$ be a group homomorphism such that $H \subseteq \ker(\varphi)$. There exists a unique homomorphism $\bar{\varphi} : G/H \to M$ such that the diagram commutes:

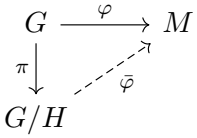

Remark 1.3.15. In other words for $H \trianglelefteq G$ and any group M, the map

$$\mathrm{Hom}_{\mathbf{Grp}}(G/H, M) \to \mathrm{Hom}_{\mathbf{Grp}}(G, M),\ \varphi \mapsto \varphi \circ \pi$$

is injective with image $\{\psi \in \mathrm{Hom}_{\mathbf{Grp}}(G, M) \mid \psi|_H = *\}$, where $*$ denotes the trivial homomorphism.

Proof. Let $\psi : G \to M$, $H \subseteq \ker \psi$. We have to show that $\psi = \varphi \circ \pi$ for some $\varphi : G/H \to M$. We define $\varphi(gH) = \psi(g)$. Then, $\forall g \in G : \varphi(\pi(g)) = \varphi(gH) = \psi(g)$, so $\psi = \phi \circ \pi$. Left to show is that φ is well-defined: Let $g_1 H = g_2 H$ in G/H. This implies that $g_1 \in g_2 H$ and since $H \subseteq \ker \varphi$, $\psi(g_1) = \psi(g_2)$, which is equivalent to $\varphi(g_1 H) = \varphi(g_2 H)$. It is easy to see that φ is a homomorphism and unique. □

Example 1.3.16.

1. The universal property of the kernel: Let $\varphi : G \to H$ be a homomorphism with kernel $K = \ker(\varphi)$ and inclusion $i : K \hookrightarrow G$. For any group M, composition with i yields a map:

$$i_* : \mathrm{Hom}_{\mathbf{Grp}}(M, K) \longrightarrow \mathrm{Hom}_{\mathbf{Grp}}(M, G).$$

This map is injective, and its image consists exactly of those homomorphisms $\psi : M \to G$ such that $\varphi \circ \psi$ is trivial.

$$\mathrm{Hom}_{\mathbf{Grp}}(M, K) \cong \{\psi \in \mathrm{Hom}_{\mathbf{Grp}}(M, G) \mid \varphi \circ \psi = 1\}.$$

2. The first isomorphism theorem is an immediate consequence of the universal property. Let $\varphi : G \to H$ be a group homomorphism.

Since $G/\ker\varphi$ is a quotient group with canonical projection $\pi : G \to G/\ker\varphi$, the universal property guarantees a unique homomorphism:

$$\bar{\varphi} : G/\ker\varphi \to H, \quad g\ker\varphi \mapsto \varphi(g).$$

Clearly, the kernel of $\bar{\varphi}$ is trivial, so the map is an isomorphism onto its image.

This yields the **canonical factorisation** of φ:

$$G \xrightarrow{\pi} G/\ker(\varphi) \xrightarrow[\sim]{\tilde{\varphi}} \mathrm{Im}(\varphi) \xhookrightarrow{\iota} H$$

such that $\varphi = \iota \circ \tilde{\varphi} \circ \pi$. Here, π is the canonical projection, ι is the inclusion map, and $\tilde{\varphi}$ is the isomorphism induced by the universal property.

Applying this result yields the following:

(a) $S_n \xrightarrow{\text{sgn}} \{\pm 1\} \implies S_n/A_n \cong \{\pm 1\}$.

(b) $GL_n(K) \xrightarrow{\det} K^\times \implies GL_n(K)/SL_n(K) \cong K^\times$.

3. Let $\varphi : G_1 \to G_2$ be a homomorphism. If $N_1 \trianglelefteq G_1$ and $N_2 \trianglelefteq G_2$ are normal subgroups such that $\varphi(N_1) \subseteq N_2$, there exists a unique induced map $\bar{\varphi}$ making the diagram commute:

$$\begin{array}{ccc} G_1 & \xrightarrow{\varphi} & G_2 \\ \downarrow{\scriptstyle \pi_1} & & \downarrow{\scriptstyle \pi_2} \\ G_1/N_1 & \xrightarrow{\exists!\,\bar{\varphi}} & G_2/N_2 \end{array}$$

We obtain this result by applying the universal property to the composition $\pi_2 \circ \varphi$.

1.3.4 Group Extensions

Definition 1.3.17. A sequence of groups and homomorphisms

$$\ldots \xrightarrow{f_{i-1}} G_i \xrightarrow{f_i} G_{i+1} \xrightarrow{f_{i+1}} \ldots$$

is called **exact at** G_i if the image of the incoming map equals the kernel of the outgoing map:

$$\mathrm{im}(f_{i-1}) = \ker(f_i).$$

The sequence is called **exact** if it is exact at every group G_i in the sequence.

Definition 1.3.18. A **group extension** of a group H by a group K is a short exact sequence:

$$* \longrightarrow K \xrightarrow{\iota} G \xrightarrow{\pi} H \longrightarrow *$$

Let $* \to K \to G \to H \to *$ be an extension.

Consider another extension $* \to K \to G' \to H \to *$. We say they are **isomorphic** if there exists an isomorphism $\varphi : G \xrightarrow{\cong} G'$ such that the following diagram commutes:

$$\begin{array}{ccccccccc} * & \longrightarrow & K & \longrightarrow & G & \longrightarrow & H & \longrightarrow & * \\ & & \| & & \cong\downarrow{\scriptstyle \varphi} & & \| & & \\ * & \longrightarrow & K & \longrightarrow & G' & \longrightarrow & H & \longrightarrow & * \end{array}$$

Remark 1.3.19. Let $* \longrightarrow K \xrightarrow{\iota} G \xrightarrow{\pi} H \longrightarrow *$ be a group extension. Assume K is abelian.

Since there is an embedding $K \hookrightarrow G$, we can treat K as a proper subgroup of G.
Consider the conjugation homomorphism $c : G \to \mathrm{Aut}_{\mathbf{Grp}}(K)$.
Since K is abelian, the conjugation by any element $h \in K$ is trivial (i.e., $hkh^{-1} = k$ for all $k \in K$). This means $K \subseteq \ker(c)$.

By the universal property of the quotient, the map c induces a unique homomorphism from the quotient group $H \cong G/K$ to the automorphism group of K:

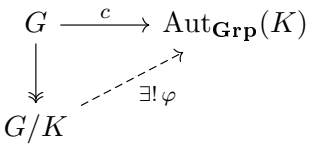

Thus, we obtain a homomorphism $\varphi : H \to \mathrm{Aut}_{\mathbf{Grp}}(K)$.

Example 1.3.20. Consider the extension:

$$K^\times \cong Z(GL_n(K)) \hookrightarrow GL_n(K) \twoheadrightarrow PGL_n(K).$$

In this specific case, the induced map $\varphi : PGL_n(K) \to K^\times$ is trivial, because $K^\times$ consists of scalar matrices λI, which lie in the center of $GL_n(K)$. φ is induced by the universal property of the quotient, so it has the mapping rule $MK^\times \mapsto c(M)$, with

$$c(M) : Z(GL_n(K)) \to Z(GL_n(K)),\ \lambda I \mapsto M\lambda I M^{-1} = \lambda I,$$

so all $MK^\times \in PGL_n(K)$ get mapped to the trivial automorphism.

This leads to the general definition:

Definition 1.3.21. An extension $* \to K \to G \to H \to *$ is called **central** if the kernel is contained in the center of the group, i.e., $K \subseteq Z(G)$. (In such extensions, the action of H on K is always trivial).

Definition 1.3.22. Let

$$* \longrightarrow K \longrightarrow G \xrightarrow{\pi} H \longrightarrow *$$

be an extension. If there is a $\sigma \in \mathrm{Hom}_{\mathbf{Grp}}(H, G)$ such that $\pi \circ \sigma = \mathrm{id}_H$, σ is called a **section** and the extension is called **split**.

Example 1.3.23.

1. Consider the extension given by the sign homomorphism:

$$* \longrightarrow A_n \hookrightarrow S_n \xrightarrow{\mathrm{sgn}} \{\pm 1\} \longrightarrow *$$

We can define a section $\sigma : \{\pm 1\} \to S_n$ by mapping -1 to any fixed transposition (e.g., $\tau = (1\,2)$):

$$\sigma(1) = \mathrm{id}, \quad \sigma(-1) = (1\,2).$$

Since $\mathrm{sgn}((1\,2)) = -1$, we have $\mathrm{sgn} \circ \sigma = \mathrm{id}$. The image of this section is a subgroup $H = \{\mathrm{id}, (1\,2)\} \cong \mathbb{Z}/2\mathbb{Z}$ that complements A_n.

2. Consider the extension given by the determinant:

$$* \longrightarrow SL_n(K) \hookrightarrow GL_n(K) \xrightarrow{\det} K^\times \longrightarrow *$$

We can define a section $s : K^\times \to GL_n(K)$ by embedding $K^\times$ into some diagonal entry of a matrix:

$$\sigma(\lambda) = \begin{pmatrix} \lambda & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{pmatrix}.$$

Clearly, $\det(\sigma(\lambda)) = \lambda \cdot 1 \cdots 1 = \lambda$, so $\det \circ \sigma = \mathrm{id}_{K^\times}$.

Lemma 1.3.24. Let $* \longrightarrow K \xrightarrow{i} G \xrightarrow{\pi} H \longrightarrow *$ be a split extension wth section σ. Then, the map

$$\Phi : K \times H \to G,\ (k,h) \mapsto k\sigma(h)$$

is a bijection.

Proof. Suppose $\Phi(k_1,h_1) = \Phi(k_2,h_2)$. Then:

$$k_1\sigma(h_1) = k_2\sigma(h_2).$$

Apply the projection π to both sides. Since $k_1, k_2 \in \ker(\pi)$, we have $\pi(k_1) = \pi(k_2) = e$. Also, since σ is a section, $\pi(\sigma(h)) = h$.

$$\pi(k_1\sigma(h_1)) = \pi(k_2 s(h_2)) \implies e \cdot h_1 = e \cdot h_2 \implies h_1 = h_2.$$

Substituting $h_1 = h_2$ back into the original equation, we can cancel $\sigma(h_1)$ from the right to obtain $k_1 = k_2$. Thus, the map is injective.

Let $g \in G$. We want to find $k \in K$ and $h \in H$ such that $g = k\sigma(h)$. Set $h = \pi(g)$. Consider the element $k := g \cdot \sigma(h)^{-1}$. We check if $k \in K$ by applying π:

$$\pi(k) = \pi(g\sigma(h)^{-1}) = \pi(g) \cdot \pi(\sigma(h))^{-1} = h \cdot h^{-1} = e.$$

Since $\pi(k) = e$, we have $k \in \ker(\pi) = K$. Thus, $g = k \cdot \sigma(h) = \Phi(k,h)$, so the map is surjective. □

Remark 1.3.25. Let $* \to K \to G \to H \to *$ be a split extension with section σ. Since K is normal in G, for every $h \in H$ and $k \in K$, the conjugate $\sigma(h)k\sigma(h)^{-1}$ lies in K. This defines a homomorphism:

$$\rho_\sigma : H \to \mathrm{Aut}(K), \quad h \mapsto (k \mapsto \sigma(h)k\sigma(h)^{-1}).$$

We observe how the multiplication in G behaves under this bijection. Let $g_1 = k_1\sigma(h_1)$ and $g_2 = k_2\sigma(h_2)$.

$$\begin{aligned} g_1 \cdot g_2 &= k_1\sigma(h_1) \cdot k_2\sigma(h_2) \\ &= k_1 \cdot \underbrace{\sigma(h_1)k_2\sigma(h_1)^{-1}}_{\rho_\sigma(h_1)(k_2)} \cdot \underbrace{\sigma(h_1)\sigma(h_2)}_{\sigma(h_1h_2)} \\ &= (k_1 \cdot \rho_\sigma(h_1)(k_2)) \cdot \sigma(h_1h_2). \end{aligned}$$

This calculation shows that the group operation on the pairs (k,h) involves the action ρ.

Definition 1.3.26. Let H, K be groups and let $\rho : H \to \mathrm{Aut}(K)$ be a homomorphism. The **semidirect product** of K and H with respect to ρ, denoted $K \rtimes_\rho H$, is the set $K \times H$ equipped with the multiplication:

$$(k_1,h_1) \cdot (k_2,h_2) := (k_1 \cdot \rho(h_1)(k_2),\ h_1h_2).$$

The group axioms are easy to verify.

Lemma 1.3.27. Let $1 \to K \to G \xrightarrow{\pi} H \to 1$ be a split extension with section $\sigma : H \to G$. Let $\rho : H \to \mathrm{Aut}(K)$ be the homomorphism defined by $h \mapsto c_{\sigma(h)}|_K$.

Then the map $\varphi : K \rtimes_\rho H \to G$ defined by $(k,h) \mapsto k \cdot \sigma(h)$ is an isomorphism of groups. Moreover, this isomorphism makes the following diagram commute (i.e., the extensions are isomorphic):

$$\begin{array}{ccccccccc} 1 & \longrightarrow & K & \xrightarrow{i_1} & K \rtimes_\rho H & \xrightarrow{\pi_2} & H & \longrightarrow & 1 \\ & & \| & & \cong\downarrow\varphi & & \| & & \\ 1 & \longrightarrow & K & \xrightarrow{i} & G & \xrightarrow{\pi} & H & \longrightarrow & 1 \end{array}$$

where $i_1(k) = (k,e)$ and $\pi_2(k,h) = h$.

Proof. Since the map $(k, h) \mapsto k \cdot \sigma(h)$ is a bijection from $K \times H$ to G, φ is a bijection.

Furthermore, we previously calculated that the multiplication in G satisfies:

$$(k_1\sigma(h_1)) \cdot (k_2\sigma(h_2)) = (k_1 \cdot \rho(h_1)(k_2)) \cdot \sigma(h_1h_2)$$

. Since this is exactly the definition of the multiplication in the semidirect product $K \rtimes_\rho H$, the map φ preserves the group operation. Thus, it is an isomorphism.

We verify the commutativity of the diagram by checking the left and right squares separately.

For the left square, we must show that $\varphi \circ i_1 = i$. Let $k \in K$. Following the top-right path, we compute $i_1(k) = (k, e_H)$, and applying φ yields $k \cdot \sigma(e_H)$. Since σ is a homomorphism, $\sigma(e_H) = e_G$, so the result is k. This matches the bottom path where $i(k) = k$.

For the right square, we must show that $\pi \circ \varphi = \pi_2$. Let $(k, h) \in K \rtimes_\rho H$. The top-right path yields $\pi_2(k, h) = h$. For the bottom-left path, we first apply φ to obtain $k \cdot \sigma(h)$. Applying π gives $\pi(k \cdot \sigma(h)) = \pi(k) \cdot \pi(\sigma(h))$. Since $k \in \ker(\pi)$ and σ is a section, this simplifies to $1 \cdot h = h$.

Since both squares commute, the extensions are isomorphic. □

Example 1.3.28. This lemma justifies the semidirect product structure of our previous examples:

1. $S_n \cong A_n \rtimes_\rho (\mathbb{Z}/2\mathbb{Z})$.

   Here, ρ describes how the transposition $(1\,2)$ conjugates the even permutations.

2. $GL_n(K) \cong SL_n(K) \rtimes_\rho K^\times$.

   Here, $\rho(\lambda)$ is the conjugation of $SL_n(K)$ by the diagonal matrix $\mathrm{diag}(\lambda, 1, \ldots, 1)$.

1.3.5 Nilpotent and Solvable Groups

Definition 1.3.29. We define a sequence of groups inductively by $G^{[0]} := G$ and

$$G^{[i+1]} := G^{[i]}/Z(G^{[i]}).$$

A group G is called **nilpotent** if there exists some $n \in \mathbb{N}$ such that $G^{[n]} = *$.

Definition 1.3.30. The derived subgroup of G is defined as $D(G) := [G, G]$, the subgroup generated by all commutators $ghg^{-1}h^{-1}$. The derived series is defined inductively by $D^{(0)}(G) := G$ and

$$D^{(i)}(G) := D(D^{(i-1)}(G)) \quad \text{for } i \geq 1.$$

A group G is called **solvable** if there exists some $n \in \mathbb{N}$ such that $D^{(n)}(G) = *$.

Lemma 1.3.31. Let G be a group. G is solvable if and only if there exists a **subnormal series**

$$* = G_n \trianglelefteq G_{n-1} \trianglelefteq \cdots \trianglelefteq G_1 \trianglelefteq G_0 = G$$

such that the quotients G_{i-1}/G_i are abelian for all $1 \leq i \leq n$.

Proof. ($\Rightarrow$) If G is solvable, then by definition the derived series terminates: $D^{(n)}(G) = \rfloor\langle\dashv\nabla\in$ for some n. Set $G_i := D^{(i)}(G)$. Recall that $D^{(i+1)}(G) = [D^{(i)}(G), D^{(i)}(G)]$. The quotient $D^{(i)}(G)/D^{(i+1)}(G)$ is abelian (since we are quotienting by the commutator subgroup). Thus, the derived series itself serves as the required subnormal series with abelian factors.
($\Leftarrow$) Suppose such a series exists: $= G_n \trianglelefteq \cdots \trianglelefteq G_0 = G$ with G_{i-1}/G_i abelian. We prove by induction that $D^{(k)}(G) \subseteq G_k$ for all k.

If $k = 1$, since $G_0/G_1 = G/G_1$ is abelian, the commutator subgroup $D(G) = [G, G]$ must be contained in G_1. Thus $D^{(1)}(G) \subseteq G_1$.

Let $k \geq 2$. Assume $D^{(k)}(G) \subseteq G_k$. We compute the next derived subgroup:

$$D^{(k+1)}(G) = [D^{(k)}(G), D^{(k)}(G)].$$

By the induction hypothesis, this is a subgroup of $[G_k, G_k]$. Since the quotient G_k/G_{k+1} is abelian, the commutator subgroup of G_k must lie inside G_{k+1} (i.e., $[G_k, G_k] \subseteq G_{k+1}$). Therefore:

$$D^{(k+1)}(G) \subseteq [G_k, G_k] \subseteq G_{k+1}.$$

By induction, $D^{(n)}(G) \subseteq G_n =$. Hence $D^{(n)}(G) =$, so G is solvable. □

Corollary 1.3.32. If G is a finite group, then G is solvable if and only if there exists a chain of subgroups

$$* = G_r \subset G_{r-1} \subset \cdots \subset G_0 = G$$

such that

$$G_{i-1}/G_i \cong \mathbb{Z}/p_i\mathbb{Z} \quad \text{where } p_i \text{ is prime.}$$

Proof. ($\Leftarrow$) If such a series exists, the factors $G_{i-1}/G_i \cong \mathbb{Z}/p_i\mathbb{Z}$ are cyclic and therefore abelian. By the Lemma, the existence of a series with abelian factors implies G is solvable.

($\Rightarrow$) Let G be a finite solvable group. By the Lemma, there exists a subnormal series where the quotients $A_i = G_{i-1}/G_i$ are abelian.

We rely on the fact that any finite abelian group can be written as a direct product (a trivial semidirect product) of cyclic subgroups:

$$A_i \cong \mathbb{Z}/n_1\mathbb{Z} \times \cdots \times \mathbb{Z}/n_k\mathbb{Z}$$

(This is a special case of the structure theorem for finitely generated R-modules over PIDs, which we will prove later.)

Using this decomposition, we can insert subgroups between G_i and G_{i-1} corresponding to these cyclic factors. Furthermore, each cyclic group $\mathbb{Z}/n_j\mathbb{Z}$ has a series where the factors are of prime order $\mathbb{Z}/p\mathbb{Z}$ (derived from the prime factorisation of n_j).

Applying this to every interval in the original series yields a composition series for G where every factor is isomorphic to $\mathbb{Z}/p\mathbb{Z}$. □

Lemma 1.3.33. Every nilpotent group is solvable.

Proof. Let G be a nilpotent group. By definition, the sequence of groups $G^{[0]} = G$ and $G^{[i+1]} = G^{[i]}/Z(G^{[i]})$ terminates at the trivial group for some n (i.e., $G^{[n]} = \{e\}$).

We define a sequence of subgroups in G inductively. Let $\pi_i : G \to G^{[i]}$ be the natural projection map composed of all the previous quotient steps. Let $Z_i = \ker(\pi_i)$.

For $i = 0$, $G^{[0]} = G$, so π_0 is the identity map. Thus $Z_0 = \{e\}$.

For $i = n$, $G^{[n]} = \{e\}$, so $Z_n = G$.

Since $G^{[i+1]}$ is a quotient of $G^{[i]}$, the map π_{i+1} is formed by composing π_i with the projection $p : G^{[i]} \to G^{[i]}/Z(G^{[i]})$. If $g \in Z_i$, then $\pi_i(g) = e$. The projection of e is e, so $\pi_{i+1}(g) = e$. Thus $g \in Z_{i+1}$. This gives us the chain:

$$* = Z_0 \trianglelefteq Z_1 \trianglelefteq \cdots \trianglelefteq Z_n = G.$$

Consider the quotient Z_{i+1}/Z_i. By definition of the sequence, the kernel of the map from $G^{[i]}$ to $G^{[i+1]}$ is exactly the center $Z(G^{[i]})$.

The elements of Z_{i+1} are exactly the elements in G that map into this center under π_i:

$$\pi_i(Z_{i+1}) = Z(G^{[i]}).$$

Since the kernel of π_i is Z_i, the map π_i induces an isomorphism between the factor group Z_{i+1}/Z_i and the image $Z(G^{[i]})$.

Since $Z(G^{[i]})$ is the center of a group, it is abelian. Therefore, Z_{i+1}/Z_i is abelian.

We have shown that the chain $Z_0 \trianglelefteq \cdots \trianglelefteq Z_n$ is a subnormal series with abelian factors. Thus, G is solvable. □

1.4 Operations of Groups on a Set

Definition 1.4.1. Let G be a group and X a set. A **left operation** of G on X is a homomorphism $\rho : G \to S(X)$ **(1)**,
or equivalently, a map $G \times X \to X$ **(2)**, $(g, x) \mapsto g \cdot x$, satisfying:

1. $\forall x \in X : e_G \cdot x = x$
2. $\forall g, h \in G, x \in X : g(hx) = (gh)x$

The set that is acted on by G is called a **G-set**.

Proposition 1.4.2. The definitions are ideed equivalent.

Proof. **(1)** $\to$ **(2)**: ρ maps $g \in G$ to permutations $\sigma_x : X \to X$, so we obtain a map

$$G \times X \to X,\ (g, x) \mapsto \rho(g)(x) =: \sigma_g(x).$$

Since ρ is a homomorphism,

$$\forall x \in X : e_G \cdot x = \sigma_{e_G}(x) = \mathrm{id}_X(x) = x.$$

as well as

$$\forall g, h \in G : g \cdot (h \cdot x) = \sigma_g(\sigma_h(x)) = (\sigma_g \circ \sigma_h)(x) = \sigma_{gh}(x) = (gh) \cdot x.$$

(2) $\to$ **(1)**: For each $g \in G$, we obtain a map $\sigma_g : X \to X$, $x \mapsto g \cdot x$.
It is a bijection since

$$\forall x \in X : \sigma_{g^{-1}}(\sigma_g(x)) = g^{-1} \cdot (g \cdot x) = (g^{-1}g) \cdot x = e_G \cdot x = x = g \cdot (g^{-1}) \cdot x = \sigma_g(\sigma_{g^{-1}}(x)),$$

so $\sigma_{g^{-1}}$ is its inverse. □

Example 1.4.3. Let G be a group.

1. Given any set S, there is the trivial operation $G \times S \to S$, $(g, s) \mapsto s$.
2. S_n operates on $\{1, \ldots, n\}$.
3. G operates on itself by left multiplication (translation), so we get a homomorphism

   $$\varphi : G \to S_G \cong S_{|G|}.$$

   Since

   $$k \in \ker \varphi \Leftrightarrow \varphi(k) = \mathrm{id}_G \Leftrightarrow \forall g \in G : k \cdot g = kg = g \Leftrightarrow k = e_G \implies \ker \varphi = *,$$

   φ is an embedding, so G is isomorphic to a subgroup of $S_{|G|}$. This result is known as Cayley's theorem.
4. G also operates on itself by conjugation ($g \cdot x = gxg^{-1}$).
5. G operates on the set of its subgroups by conjugation and on sets of left cosets by translation.
6. Let K be a field. $\mathrm{GL}_n(K)$ operates on the left on K^n by matrix multiplication.

Definition 1.4.4. Let X be a G-set, $x \in X$.

1. $\mathrm{Stab}_g(x) = \mathcal{I}_x = \{g \in G \mid g \cdot x = x\} \subseteq G$ is called the **stabiliser** or isotropy subgroup of x.

2. $G \cdot x = \{g \cdot x \mid g \in G\} \subseteq X$ is called the **orbit** of x.

3. x is called a **fixed point** of the operation if $\mathrm{Stab}_G(x) = G$. We denote the set of fixed points by $S^G = \{s \in S \mid g \cdot s = s, \forall g \in G\} \subseteq X$.

4. An operation is called **transitive** if $X \neq \varnothing$ and $\forall x, y \in X \, \exists g \in G : g \cdot x = y$.

5. Let $Y \subseteq X$. We say that Y is G**-invariant** if $\forall g \in G : g \cdot Y = Y$

Definition 1.4.5. Let X be a G-set. An equivalence relation is defined by:

$$x \sim y \iff \exists g \in G, gx = y$$

The equivalence classes are the orbits $G \cdot x = \{gx \mid g \in G\}$. The set of equivalence classes is denoted $G \setminus X$. It is in particular the disjoint union of all orbits.

Proof. We verify the axioms of an equivalence relation:
Reflexivity: For all $x \in X$, $e \cdot x = x$, so $x \sim x$.
Symmetry: If $x \sim y$, then $gx = y$ for some $g \in G$. Thus $x = g^{-1}y$, implying $y \sim x$.
Transitivity: If $x \sim y$ and $y \sim z$, then $y = gx$ and $z = hy$ for some $g, h \in G$. Then $z = h(gx) = (hg)x$, so $x \sim z$.

The equivalence class of x is $[x] = \{y \in X \mid x \sim y\} = \{y \in X \mid \exists g \in G, y = gx\} = \{gx \mid g \in G\} = G \cdot x$. □

Remark 1.4.6. Using the orbit decomposition, it is easy to see that a subset is G-invariant if and only if it is a union of orbits.

Note that unlike in the set of left cosets, the equivalence classes in $G \setminus X$ are not in bijection.

Theorem 1.4.7. Let X be a G-set, $x \in X$. The map $\varphi_x : G/\mathcal{I}_x \to G \cdot x$, $g\mathcal{I}_x \mapsto g \cdot x$ is well-defined and a bijection. In particular, if G is finite:

$$|G \cdot x| = \frac{|G|}{|\mathcal{I}_x|}.$$

Proof. Let $\varphi(g\mathcal{I}_x) = g \cdot x$. We verify well-definedness and injectivity simultaneously via the following chain of equivalences:

$$g\mathcal{I}_x = h\mathcal{I}_x \iff h^{-1}g \in \mathcal{I}_x \iff (h^{-1}g) \cdot x = x \iff g \cdot x = h \cdot x.$$

Reading left to right shows φ is well-defined; reading right to left shows φ is injective.
Surjectivity is immediate by the definition of the orbit $G \cdot x$. Thus, φ is a bijection. □

Definition 1.4.8. A map $f : X \to Y$ between G-sets is called G**-equivalent** if

$$\forall g \in G, x \in X : f(gx) = gf(x).$$

We may now define G-**Set**, the category of G-sets.

Example 1.4.9. Let X be a G-set.

1. Let $Y \subset X$ be G-invariant. Y is a G-set with left multiplication and the inclusion $Y \hookrightarrow G$ is G-equivalent.

   For instance, $\forall x \in X$, the orbit $G \cdot x \subset X$ is G-invariant.

2. $\forall x \in X$, the bijection $G/\mathcal{I}_x \xrightarrow{\cong} G \cdot x$ is G-equivalent.

   Indeed: let $g \in G$. $\varphi(g\mathcal{I}_x) = g \cdot x = g(e_G \cdot x) = g\varphi(I_x)$.

3. Let X be a G-set. We can decompose X into its orbits. Let I be a set of indices indexing the orbits, i.e., $G\backslash X = \{\mathcal{O}_i\}_{i\in I}$. We choose a system of representatives $(x_i)_{i\in I}$ such that $x_i \in \mathcal{O}_i$ for each i. Then the map defined by the disjoint union of the orbit bijections is a G-equivalent isomorphism:
$$\coprod_{i\in I} G/\mathcal{I}_{x_i} \xrightarrow{\cong} X.$$

4. Let $f : X \to Y$ be a G-equivalent map. The map f induces well-defined maps on the fixed points $f^G : X^G \to Y^G$
   and the quotient spaces $\bar{f} : G\backslash X \to G\backslash Y$, $G \cdot x \mapsto G \cdot f(x)$

5. Consider the symmetric group S_n acting on the set $N = \{1, \ldots, n\}$. This action is transitive. The stabiliser of the element n is the subgroup of permutations fixing n, which is isomorphic to S_{n-1} (permuting the first $n-1$ elements). Thus, we have the canonical S_n-equivalent bijection:
$$S_n/S_{n-1} \cong \{1, \ldots, n\}.$$

6. Let K be a field and $n \geq 1$. The general linear group $GL_n(K)$ acts on the vector space K^n. The orbits of this action are exactly $\{0\}$ and $K^n \setminus \{0\}$. In particular, the set of non-zero vectors $K^n \setminus \{0\}$ is a transitive $GL_n(K)$-set.

7. Let $P \in \mathbb{R}[X]$ be a polynomial with real coefficients, $P = \sum a_i X^i$. Let $R_P(\mathbb{C}) = \{z \in \mathbb{C} \mid P(z) = 0\}$ be the set of complex roots. Complex conjugation $\tau : \mathbb{C} \to \mathbb{C}, z \mapsto \bar{z}$ is a field automorphism. This defines an action of the group $G = \mathbb{Z}/2\mathbb{Z} = \{\mathrm{id}, \tau\}$ on $\mathbb{C}$.

   The set of roots $R_P(\mathbb{C})$ is a G-invariant subset.

   Proof. Since the coefficients a_i are real, $P(\bar{z}) = \overline{P(z)}$. Thus, if $P(z) = 0$, then $P(\bar{z}) = 0$. □

   We can decompose $R_P(\mathbb{C})$ into orbits under this action. The fixed points are the real roots, which form singleton orbits $\{x_i\}$. The non-real roots come in conjugate pairs, forming orbits of size 2, $\{z_j, \bar{z}_j\}$.
$$R_P(\mathbb{C}) = \{x_1, \ldots, x_r\} \sqcup \{z_1, \bar{z}_1\} \sqcup \cdots \sqcup \{z_s, \bar{z}_s\}.$$

8. Generalisation: Let L be a field and let $G \subset \mathrm{Aut}(L)$ be a finite subgroup of automorphisms. Let $K = L^G$ be the fixed subfield. For any polynomial $P \in K[X]$, the set of roots $R_P(L)$ is a finite G-invariant subset of L. (This is because for any $\sigma \in G$ and root y, $P(\sigma(y)) = \sigma(P(y)) = \sigma(0) = 0$).

9. Let X be a G-set. If x, y are two elements in the same orbit, then their stabilisers G_x and G_y are conjugate subgroups in G.

   Proof. Since x and y are in the same orbit, there exists $g \in G$ such that $y = g \cdot x$. We claim that $G_y = gG_xg^{-1}$.

   Let $h \in G_y$. By definition, $h \cdot y = y$. Substituting $y = g \cdot x$, we get $h \cdot (g \cdot x) = g \cdot x$. Multiplying by g^{-1} on the left (and using the group action axiom):
$$g^{-1}hg \cdot x = x.$$
   This implies $g^{-1}hg \in G_x$, or equivalently, $h \in gG_xg^{-1}$. Thus, $G_y \subseteq gG_xg^{-1}$. A similar argument shows the reverse inclusion. □

10. Let G operate on itself by conjugation. This induces a left operation of G on the set $U(G)$ of all subgroups of G. For a subgroup $H \in U(G)$, the normaliser $N_G(H)$ is the stabiliser of H under the conjugation action.

    It follows that if $N \subseteq G$ is a subgroup containing H such that $H \trianglelefteq N$, $N \subset I_H = N_G(H)$.

1.4.1 Cycle Decomposition in S_n

We analyze the action of a permutation $\sigma \in S_n$ on the set $\{1, \ldots, n\}$. The subgroup $\langle\sigma\rangle$ generated by σ acts on $\{1, \ldots, n\}$, decomposing it into disjoint orbits:

$$\{1, \ldots, n\} = \coprod_{\alpha \in \langle\sigma\rangle \backslash \{1,\ldots,n\}} \alpha = \{i \mid \sigma(i) = i\} \sqcup \coprod_{|\alpha| \geq 2} \alpha$$

The set of elements moved by σ, $\{i \mid \sigma(i) \neq i\}$, is called the **support** of σ, denoted $\mathrm{supp}(\sigma)$.

For each orbit α with $|\alpha| \geq 2$, we choose an element $i \in \alpha$. The orbit has the form $\alpha = \{i, \sigma(i), \ldots, \sigma^{l-1}(i)\}$ where $l = |\alpha|$. We define the **cycle** γ_α corresponding to this orbit as the permutation that cycles these elements and fixes everything else.

Lemma 1.4.10. Let γ_1 and γ_2 be cycles with disjoint support (i.e., $\mathrm{supp}(\gamma_1) \cap \mathrm{supp}(\gamma_2) = \emptyset$). Then

$$\gamma_1 \circ \gamma_2 = \gamma_2 \circ \gamma_1.$$

Proof. If $x \in \mathrm{supp}(\gamma_1)$, then $x \notin \mathrm{supp}(\gamma_2)$, so $\gamma_2(x) = x$. Thus $\gamma_2(\gamma_1(x)) = \gamma_1(x)$. Similarly $\gamma_1(\gamma_2(x)) = \gamma_1(x)$. If x is not in either support, both fix x. Thus the permutations are identical. □

Theorem 1.4.11. Every permutation $\sigma \in S_n$ can be written as a product of cycles with disjoint support:

$$\sigma = \gamma_1 \circ \cdots \circ \gamma_s$$

where γ_j corresponds to the distinct orbits of σ with size ≥ 2. This decomposition is unique up to the order of the factors (which commute together).

Proof. The decomposition follows directly from the partition of $\{1, \ldots, n\}$ into disjoint orbits. For any $x \in \{1, \ldots, n\}$:
If x is a fixed point of σ, it is not in the support of any γ_j, so both sides map x to x.
If x is moved by σ, it belongs to exactly one orbit α_j. Thus $x \in \mathrm{supp}(\gamma_j)$ and $x \notin \mathrm{supp}(\gamma_k)$ for $k \neq j$. The product on the right acts as γ_j on x, which matches $\sigma(x)$ by definition. □

Corollary 1.4.12. The order of σ is the least common multiple of the lengths of its disjoint cycles.

$$\mathrm{ord}(\sigma) = \mathrm{lcm}(l_1, \ldots, l_s).$$

Proof. Let $l_j = \mathrm{ord}(\gamma_j)$ be the length of the j-th cycle. If $\sigma^t = \mathrm{id}$, then $(\gamma_1 \circ \cdots \circ \gamma_s)^t = \mathrm{id}$. Since the cycles have disjoint support and thus commute, this expands to $\gamma_1^t \circ \cdots \circ \gamma_s^t = \mathrm{id}$. Since the supports are disjoint, each factor must individually be the identity: $\gamma_j^t = \mathrm{id}$ for all j. This holds if and only if $l_j \mid t$ for all j. The smallest such positive t is $\mathrm{lcm}(l_1, \ldots, l_s)$. □

1.4.2 Application to finite groups

Let G be a finite group operating on a finite set X. We distinguish between orbits of size 1 (fixed points) and larger orbits. The set of fixed points is $X^G = \{x \in X \mid \forall g \in G, g \cdot x = x\}$. If $x \notin X^G$, then its orbit has size $|\alpha| > 1$.

Lemma 1.4.13 (Class Equation)**.** Let R be a set of representatives for the non-trivial orbits of X (i.e., orbits with size > 1). Then:

$$|X| = |X^G| + \sum_{x \in R} [G : I_x].$$

Corollary 1.4.14. Let G be a group of order p^r, p prime, acting on a finite set X. Then:

$$|X| \equiv |X^G| \pmod{p}.$$

Proof. Consider the class equation from the Lemma above:

$$|X| - |X^G| = \sum_{x \in R} [G : I_x].$$

For every representative $x \in R$, the orbit is non-trivial, so the index $[G : I_x] > 1$. Since G is a p-group, the index $[G : I_x]$ must be a divisor of $|G| = p^r$. The only divisors of p^r greater than 1 are multiples of p.

Therefore, every term in the sum is divisible by p.

$$\sum_{x \in R} [G : I_x] \equiv 0 \pmod{p}.$$

Consequently, $|X| - |X^G| \equiv 0 \pmod{p}$. □

1.5 p-Groups and the Sylow Theorems

Definition 1.5.1. Let p be a prime number. A (finite) p-group G is a group of order p^r, $r \geq 0$.

Lemma 1.5.2. Any subgroup of a p-group is a p-group.

Proof. Clear. □

Theorem 1.5.3. Let G be a nontrivial finite p-group. Then, $Z(G) \neq *$.

Proof. Note that if we let G act on itself by conjugation,

$$G^G = \{h \in G \mid \forall g \in G : ghg^{-1} = h (\iff gh = hg)\} = Z(G).$$

We know that $|G| \equiv 0 \bmod p$ and, from the corollary from the previous section, that $|G| \equiv |Z(G)| \bmod p$.

It follows that $0 \equiv |Z(G)| \bmod p$, and since $|Z(G)| \geq 1$, $|Z(G)| \geq p$, so $Z(G)$ is in particular nontrivial. □

Corollary 1.5.4. Every finite p-group is nilpotent.

Proof. Let G be a finite p-group. We consider the sequence $G^{[0]} = G$ and $G^{[i+1]} = G^{[i]}/Z(G^{[i]})$.

If $G^{[i]}$ is non-trivial, it is a p-group, so its center is nontrivial. Since the order of a quotient group is given by $|G/N| = |G|/|N| = p^r/p^s = p^{r-s}$, the quotient $G^{[i+1]}$ remains a p-group with strictly smaller order. Thus, the sequence of groups must eventually terminate at $\{e\}$. □

Theorem 1.5.5. Let G be a finite group with g be a prime divisor of its order. Then, there is a subgroup of order p (which must be cyclic).

Proof. Let $X = \{(g_1, g_2, \ldots, g_p) \in G^p \mid g_1 g_2 \ldots g_p = e\}$. We can choose the first $p-1$ elements freely, and the last element is then uniquely determined (specifically $g_p = (g_1 \ldots g_{p-1})^{-1}$). This implies $|X| = |G|^{p-1}$. Since $p \mid |G|$, it follows that $p \mid |X|$.

Let $C \cong \mathbb{Z}/p\mathbb{Z}$, and let σ be a generator of C. Let C act on X by:

$$\sigma(g_1, \ldots, g_p) := (g_2, g_3, \ldots, g_p, g_1).$$

The shifted tuple is still in X. Indeed:

$$g_1 \ldots g_p = e \implies e = g_1^{-1}(g_1 g_2 \ldots g_p)g_1 = (g_1^{-1} g_1)(g_2 \ldots g_p g_1) = g_2 \ldots g_p g_1.$$

Since $\langle \sigma \rangle = C$, the entire action is well-defined.

Because $|C| = p$ and the size of any orbit must divide $|C|$, any orbit has size 1 or p. Let N_1 be the number of orbits of size 1 (the fixed points), and N_p the number of orbits of size p. Then the total size is the sum of the orbit sizes:

$$|X| = (N_1 \cdot 1) + (N_p \cdot p).$$

The fixed points are exactly the tuples of the form $(g, \ldots, g)$ with $g \in G$. For this tuple to be in X, the product must be identity: $g^p = e$. Thus, $N_1 = |\{g \in G \mid g^p = e\}|$.

We know $p \mid |X|$. From the equation $|X| = N_1 + pN_p$, we can write $N_1 = |X| - pN_p$. Since $p \mid |X|$ and $p \mid pN_p$, it follows that $p \mid N_1$.

We know $N_1 > 0$, because the tuple $(e, \ldots, e)$ is always in X (as $e^p = e$). Since $p \mid N_1$ and $N_1 \neq 0$, we must have $N_1 \geq p$.

Thus, there are at least $p-1$ non-identity elements g such that $g^p = e$ (which is equivalent to $\operatorname{ord}(g) = p$, because p is prime). □

Definition 1.5.6. Let G be a finite group with $|G| = n$. Let p be a prime number. We can write $n = p^r \cdot m$ where $\gcd(p, m) = 1$. Here, r is the p-adic valuation of n. Assume $p \mid n$. A **p-Sylow subgroup** (or Sylow-p-subgroup) of G is a subgroup $H \subseteq G$ such that $|H| = p^r$. (i.e., it is a p-subgroup of maximal possible order).

Example 1.5.7.

1. Let $G = \mathbb{Z}/n\mathbb{Z}$. If $n = p^r m$ as in the definition, then:

   $$\mathbb{Z}/n\mathbb{Z} \cong \mathbb{Z}/p^r\mathbb{Z} \times \mathbb{Z}/m\mathbb{Z}.$$

   The subgroup isomorphic to $\mathbb{Z}/p^r\mathbb{Z}$ is the unique p-Sylow subgroup. (Generally, if $n = p_1^{n_1} \ldots p_s^{n_s}$ is the prime decomposition, then $\mathbb{Z}/n\mathbb{Z} \cong \mathbb{Z}/p_1^{n_1}\mathbb{Z} \times \cdots \times \mathbb{Z}/p_s^{n_s}\mathbb{Z}$).

2. Let $G = S_3$. Order $|S_3| = 6 = 3 \times 2$. The subgroup $\langle (123) \rangle = \{\mathrm{id}, (123), (132)\}$ has order 3. It is a 3-Sylow subgroup.

3. Let $G = S_4$. Order $|S_4| = 24 = 3 \times 2^3$.

   The subgroup $\langle (123) \rangle \cong \mathbb{Z}/3\mathbb{Z}$ is a 3-Sylow subgroup.
   The Dihedral group D_4 (of order 8) embedded in S_4 is a 2-Sylow subgroup.

Theorem 1.5.8 (Sylow Theorems)**.** Let G be a finite group, $|G| = n$, with $n = p^r \cdot m$ and $\gcd(p, m) = 1$, $r \geq 1$.

1. For every $k \in \{1, \ldots, r\}$, there exists a subgroup $H \subseteq G$ such that $|H| = p^k$. (In particular, p-Sylow subgroups always exist).

2. If $H \subseteq G$ is a p-group and $P \subseteq G$ is a p-Sylow subgroup, then there exists $g \in G$ such that:

   $$H \subseteq gPg^{-1}.$$

   (In particular, all p-Sylow subgroups are conjugate to each other).

3. Let n_p be the number of p-Sylow subgroups in G. Then, $n_p \mid r$ and $n_p \equiv 1 \pmod{p}$.

Proof of the first theorem. Let $k \in \{0, \ldots, r\}$. Let $\Omega := \{S \subseteq G \mid |S| = p^k\}$. Since $\binom{|G|}{p^k}$ is the number of subsets of G with cardinality p^k, $|\Omega| = \binom{|G|}{p^k}$.

Claim: The highest exponent e such that $p^e \mid |\Omega|$ is $e = r - k$. Notation: $v_p(|\Omega|) = r - k$.

Indeed:

$$|\Omega| = \binom{p^r m}{p^k} = \frac{p^r m(p^r m - 1)\ldots(p^r m - p^k + 1)}{p^k(p^k - 1)\ldots 1} = \prod_{i=0}^{p^k-1} \frac{p^r m - i}{p^k - i}.$$

Using the property $v_p(xy) = v_p(x) + v_p(y)$, we have:

$$v_p(|\Omega|) = \sum_{i=0}^{p^k-1} \Big(v_p(p^r m - i) - v_p(p^k - i) \Big).$$

For the first term ($i = 0$): $v_p(p^r m) - v_p(p^k) = r - k$ (because $p \nmid m$).

We can now assume $1 \leq i < p^k$. We consider the p-adic valuation of i: $i = p^j \cdot y$, where $p \nmid y$. Since $i < p^k$, we have $j < k \leq r$. Then, $p^k - i = p^k - (p^j y) = p^j(p^{k-j} - y)$. Because $p \mid p^{k-j}$ and $p \nmid y$, we have $p \nmid (p^{k-j} - y)$. Thus, $v_p(p^k - i) = v_p(p^j(p^{k-j} - y)) = j$.

And $p^r m - i = p^r m - (p^j y) = p^j(p^{r-j} m - y)$. Because $p \mid p^{r-j} m$ and $p \nmid y$, we have $p \nmid (p^{r-j} m - y)$. Thus, $v_p(p^r m - i) = v_p(p^j(p^{r-j} m - y)) = j$.

Therefore, for all other terms ($i > 0$), $v_p(p^r m - i) - v_p(p^k - i) = j - j = 0$. Consequently, the sum collapses to the first term: $v_p(|\Omega|) = r - k$. □

Define a group action of G on Ω by $g \cdot S := \{gs \mid s \in S\}$. Note that $g \cdot S$ is indeed in Ω, because $h \mapsto gh$ maps bijectively.

Consider the orbit decomposition $|\Omega| = \sum_{\mathcal{O} \in G\backslash\Omega} |\mathcal{O}|$. By the claim above, $p^{r-k+1} \nmid |\Omega|$, so there must be an $S \in \Omega$ such that $p^{r-k+1} \nmid |G \cdot S|$. This implies $v_p(|G \cdot S|) \leq r - k$.

By the Orbit-stabiliser Theorem: $|G \cdot S| = \frac{|G|}{|\mathcal{I}_S|} \iff |G| = |\mathcal{I}_S| \cdot |G \cdot S|$. We analyze the valuations:

$$v_p(|G|) = r, \quad v_p(|G \cdot S|) \leq r - k.$$

Thus, $v_p(|\mathcal{I}_S|) = v_p(|G|) - v_p(|G \cdot S|) \geq r - (r - k) = k$. This implies $|\mathcal{I}_S| \geq p^k$.

Now we must show $|\mathcal{I}_S| \leq p^k$. Let $s_0 \in S$. Define $f : \mathcal{I}_S \to S$ by $h \mapsto hs_0$. Then $g : S \to \mathcal{I}_S$, $h \mapsto hs_0^{-1}$ acts as a left inverse on the image. Explicitly, let $h \in \mathcal{I}_S$:

$$g \circ f(h) = g(hs_0) = hs_0 s_0^{-1} = h \implies g \circ f = \mathrm{id}_{\mathcal{I}_S}.$$

Thus, f is injective, so $|\mathcal{I}_S| \leq |S| = p^k$.

To conclude, we have $|\mathcal{I}_S| \geq p^k$ and $|\mathcal{I}_S| \leq p^k \implies |\mathcal{I}_S| = p^k$. So we have found a subgroup of G of order p^k. □

Proof of the second theorem. Let P be a p-Sylow subgroup of G and let $\Omega := G/P$. Since $|G| = p^r m$ with $p \nmid m$ and $|P| = p^r$, it follows immediately from Lagrange's Theorem that $|\Omega| = m \not\equiv 0 \pmod{p}$.

Now, let H be any p-subgroup of G acting on Ω by left multiplication.

We have already shown that this must hold: $|\Omega| \equiv |\Omega^H| \pmod{p}$ (class equation).

Combining this observation with $|\Omega| \not\equiv 0 \pmod{p}$, we see that $|\Omega^H| \not\equiv 0 \pmod{p}$, so Ω^H is non-empty. Let $gP \in \Omega^H$. Then for all $h \in H$,

$$h(gP) = gP \implies g^{-1}hgP = P \implies g^{-1}Hg \subseteq P \implies H \subseteq gPg^{-1}.$$

If H is a Sylow subgroup, then $|H| = |P|$, forcing $H = gPg^{-1}$. □

Proof of the third theorem. Recall that $G = p^k \cdot m$.

Let $\Sigma = \mathrm{Syl}_p(G)$ be the set of all p-Sylow subgroups of G. By definition, $n_p = |\Sigma|$.

1. The divisibility condition ($n_p \mid m$)

The group G acts on Σ by conjugation. By the second theorem, this action is transitive. Fix a Sylow subgroup $P \in \Sigma$. We determine the stabiliser of P under this action:

$$\mathrm{Stab}_G(P) = \{g \in G \mid g \cdot P = P\} = \{g \in G \mid gPg^{-1} = P\}.$$

By definition, the set of elements that conjugate a subgroup to itself is the normaliser of that subgroup. Thus, $\mathrm{Stab}_G(P) = N_G(P)$.

The size of the orbit is given by the Orbit-stabiliser Theorem:

$$n_p = |\Sigma| = \frac{|G|}{|\mathrm{Stab}_G(P)|} = \frac{|G|}{|N_G(P)|}.$$

We consider the fraction representing the index of P in G:

$$\frac{|G|}{|P|} = \frac{|G|}{|N_G(P)|} \cdot \frac{|N_G(P)|}{|P|}.$$

We know that $|G| = p^r m$ and $|P| = p^r$. Substituting these values into the left side:

$$m = n_p \cdot \frac{|N_G(P)|}{|P|}.$$

Since P is a subgroup of $N_G(P)$, Lagrange's Theorem guarantees that the fraction $\frac{|N_G(P)|}{|P|}$ is an integer. Therefore, n_p must be a divisor of m.

2. The congruence condition ($n_p \equiv 1 \pmod p$)

We restrict the action of G on Σ to the subgroup P itself. P acts on Σ by conjugation. We know that

$$|\Sigma| \equiv |\Sigma^P| \pmod p.$$

The set of fixed points is $\Sigma^P = \{Q \in \Sigma \mid \forall x \in P, xQx^{-1} = Q\} = \{Q \in \Sigma \mid P \subseteq N_G(Q)\}$.

Clearly $P \in \Sigma^P$ because P is a subgroup of its own normaliser ($P \subseteq N_G(P)$).

Let $Q \in \Sigma^P$ be any fixed point. Then $P \subseteq N_G(Q)$. Inside the group $N_G(Q)$, both P and Q are subgroups of order p^r, so they are both p-Sylow subgroups of $N_G(Q)$.

By definition of the normaliser, Q is a normal subgroup of $N_G(Q)$. By the second theorem applied to the group $N_G(Q)$, all Sylow subgroups must be conjugate. Since Q is normal, it is the only conjugate of itself. Thus, P must equal Q.

The only fixed point is P itself, so $|\Sigma^P| = 1$.

$$n_p = |\Sigma| \equiv 1 \pmod p.$$

□

Example 1.5.9.

1. $n_p = 1 \iff$ the p-Sylow subgroup is normal.
2. Let A be a finite abelian group. If $p \mid |A|$, there exists a unique p-Sylow subgroup $A_p \subseteq A$.

   For example, consider $\mathbb{Z}/n\mathbb{Z}$. If $p \mid n$, write $n = p^r \cdot m$ with $p \nmid m, r \geq 1$. Then $\langle \bar{m} \rangle \subset \mathbb{Z}/n\mathbb{Z} \cong \mathbb{Z}/p^r\mathbb{Z} \times \mathbb{Z}/m\mathbb{Z}$ is the unique p-Sylow subgroup.
3. Consider S_3. $|S_3| = 6 = 2 \cdot 3$. $A_3 \trianglelefteq S_3$ is a 3-Sylow subgroup $\implies n_3 = 1$.

   For n_2: we know $n_2 \mid 6 \implies n_2 \in \{1, 2, 3, 6\}$. Also $n_2 \equiv 1 \pmod 2 \implies n_2 \in \{1, 3\}$. For any distinct i, j, the subgroups $\langle (ij) \rangle$ are 3 distinct 2-Sylow subgroups. Thus $n_2 = 3$.
4. Consider S_4. $|S_4| = 24 = 3 \cdot 2^3$. We know that $n_2 \mid 3$ and $n_2 \equiv 1 \pmod 2 \implies n_2 \in \{1, 3\}$. D_4 (of order 8) is a 2-Sylow subgroup, but it is not normal in $S_4 \implies n_2 \neq 1 \implies n_2 = 3$.

   For n_3: we know $n_3 \mid 8$ and $n_3 \equiv 1 \pmod 3 \implies n_3 \in \{1, 4\}$. Notice that $\langle (123) \rangle \neq \langle (124) \rangle$, but both are 3-Sylow subgroups $\implies n_3 = 4$.

2 Rings

2.1 Generalities

Definition 2.1.1. A **ring** R is a tuple $(R, +\cdot)$ with internal laws of compositon $+$, $\cdot$, such that:

1. $(R, +)$ is an abelian group with neutral element 0
2. $(R, \cdot)$ is an associative monoid with neutral element 1
3. $\forall a, b, c \in R : (a + b) \cdot c = a \cdot c + b \cdot c$ *(Distributive law)*

We say that R is **commutative** if $\forall a, b \in R : a \cdot b = b \cdot a$.

Definition 2.1.2. Let R, S be rings. A **ring homomorphism** $\varphi : R \to S$ is a map such that for all $r, s \in R$:

1. $\varphi(r + s) = \varphi(r) + \varphi(s)$
2. $\varphi(r \cdot s) = \varphi(r) \cdot \varphi(s)$

Now, one may define **Ring**, the category of rings.

Lemma 2.1.3. Let $r, s \in R$.

1. $r \cdot 0 = 0 \cdot r = 0$
2. $(-r) \cdot s = r \cdot (-s) = -(r \cdot s)$

Proof. 1. $r \cdot 0 = r \cdot (0 + 0) = r \cdot 0 + r \cdot 0 \implies r \cdot 0 = 0$.

2. $r \cdot (-s) + r \cdot s = r \cdot (s + (-s)) = r \cdot 0 = 0 \implies r \cdot (-s) = -(r \cdot s)$

□

Example 2.1.4. 1. $(\mathbb{Z}, +, \cdot)$, the ring of integers, is the initial object in **Ring**.

2. $(\mathbb{Q}, +, \cdot) \subset (\mathbb{R}, +, \cdot) \subset (\mathbb{C}, +, \cdot)$ are subrings.
3. For $n \geq 1$, $\mathbb{Z}/n\mathbb{Z}$ is a ring, the projection $\mathbb{Z} \to \mathbb{Z}/n\mathbb{Z}$ is a ring homomorphism.
4. In the above example, for $n = 0$, $\mathbb{Z}/\mathbb{Z} =: 0$ is the **zero ring**. A ring R is the zero ring if and only if $0_R = 1_R$.
5. Let K be a field, $n \geq 1$. $M_n(K)$ is a ring. More generally, let V be a vector space. Then, $\mathrm{End}_K(V)$ is a ring with multiplication defined as composition.
6. For rings R, S, $R \times S$ is a ring called the **product ring** of R and S with pointwise operations. More generally, if $(R_i)_{i \in \mathcal{I}}$ is a family of rings, $\prod_i R_i$ is a ring.
7. Let G be a group. Let $\mathbb{Z}[G]$ be the free abelian group on the set G.

   The elements of $\mathbb{Z}[G]$ are formal sums with finite support:

   $$\alpha = \sum_{g \in G} a_g [g]$$

   where $a_g \in \mathbb{Z}$ are coefficients, and $a_g = 0$ almost all g.

   Let $[g]$ denote the function that is 1 at g and 0 elsewhere. $\{[g] \mid g \in G\}$ generates the free abelian group $\mathbb{Z}[G]$.

We define multiplication on the generators using the group operation of G:

$$[g] \cdot [h] := [gh].$$

We extend this linearly to the whole set. If $\alpha = \sum a_g[g]$ and $\beta = \sum b_h[h]$, their product is:

$$\alpha \cdot \beta = \sum_{g,h \in G} (a_g b_h)[gh].$$

This gives $\mathbb{Z}[G]$ the structure of a ring, called the **group ring**. It is commutative if and only if G is abelian.

Definition 2.1.5. Let $S \subseteq R$ be a subset. S is called a **subring** of R if:

1. $(S, +)$ is a subgroup of $(R, +)$
2. $1_R \in S$, $\forall x, y \in S : x \cdot_R y \in S$. *($S$ is stable by product in R)*

Definition 2.1.6. Let R be a ring. The **multiplicative group**, denoted $R^\times$, is defined as:

$$R^\times := \{u \in R \mid \exists v \in R : uv = vu = 1_R\}.$$

This forms a group under the ring multiplication. Note that if R is non-commutative, $R^\times$ is generally a non-abelian group.

Example 2.1.7.

1. A ring R is called a **division algebra** if $R^\times = R \setminus \{0\}$. If R is commutative and a division algebra, R is a **field**.
2. Let K be a field. The group of units of the matrix ring $M_n(K)$ is the general linear group:
$$(M_n(K))^\times =: GL_n(K).$$
3. Let $n \in \mathbb{N}$ with $n \geq 2$. Consider the commutative ring $(\mathbb{Z}/n\mathbb{Z}, +, \cdot)$. The group of units is given by:
$$(\mathbb{Z}/n\mathbb{Z})^\times = \{\bar{\lambda} \mid \lambda \in \mathbb{Z}, \gcd(\lambda, n) = 1\}.$$

   Proof. Let $\lambda \in \mathbb{Z}$.

$$\begin{aligned} \bar{\lambda} \text{ is invertible} &\iff \exists \mu \in \mathbb{Z} : \bar{\lambda} \cdot \bar{\mu} = \bar{1} \\ &\iff \exists \mu \in \mathbb{Z} : n \mid (1 - \lambda\mu) \\ &\iff \exists \mu, m \in \mathbb{Z} : 1 = \lambda\mu + nm. \end{aligned}$$

   This implies that $\gcd(\lambda, n) = 1$. Indeed, if there were a prime p such that $p \mid n$ and $p \mid \lambda$, then p must divide the linear combination $\lambda\mu + nm = 1$, which is a contradiction. □

   Remark: It follows that if n is not prime, $\mathbb{Z}/n\mathbb{Z}$ is not a field.
4. Let $\varphi : R \to S$ be a ring homomorphism. Then φ maps units to units, i.e., $\varphi(R^\times) \subseteq S^\times$.

   Proof: Since $\varphi(1_R) = 1_S$, for any $x \in R^\times$ we have:
$$\varphi(xx^{-1}) = \varphi(1_R) = 1_S \implies \varphi(x)\varphi(x^{-1}) = 1_S.$$
   Similarly $\varphi(x^{-1})\varphi(x) = 1_S$. Thus $\varphi(x)$ is invertible in S with inverse $\varphi(x^{-1})$.

   Consequently, φ induces a group homomorphism:
$$\varphi^\times : R^\times \to S^\times.$$
   In categorical terms, this defines a functor from the category of rings to the category of groups:
$$\mathcal{F} : \mathbf{Ring} \to \mathbf{Grp}, \quad R \mapsto R^\times.$$

5. Any ring homomorphism mapping out of a field is injective.

   Proof. Consider a nonzero element in the domain. Since the domain is a field, this element is a unit. Because homomorphisms map units to units, its image must be a unit in the codomain, so it cannot be zero. Thus, the kernel is trivial. □

Definition 2.1.8. A commutative ring R is called an **integral domain** if $R \neq 0$ and

$$\forall x, y \in R : xy = 0 \implies x = 0 \text{ or } y = 0$$

In other words: the product of nonzero elements are nonzero. (1 is the empty product, so $R \neq 0$ is also covered)

Example 2.1.9. 1. $\mathbb{Z}$ is an integral domain.

2. Any field is an integral domain.

3. Finite integral domains are fields. *(proof: exercises)*

4. In integral domains, the following simplification is possible:

$$\forall a, b, c \in R, a \neq 0 : ab = ac \implies b = c.$$

   Proof. $ab - ac = 0 \iff a(b-c) = 0 \iff b - c = 0 \iff b = c.$ □

2.1.1 Formal Power Series and Polynomials

Definition 2.1.10. Let R be a commutative ring. We consider the set of sequences $R^{\mathbb{N}} = \mathrm{Hom}_{\mathbf{Set}}(\mathbb{N}, R)$. This set becomes a ring $(R[[x]], +, \cdot)$ under componentwise addition and the following product:

$$(a_n)_n \cdot (b_n)_n = \left(\sum_{i+j=n} a_i b_j \right)_n .$$

This ring is called the **ring of formal power series** over R.

We identify R as a subring of $R[[x]]$ via the injection $r \mapsto (r, 0, 0, \dots)$. Furthermore, we define the indeterminate X as the specific sequence $X := (0, 1, 0, 0, \dots)$. Since X^k corresponds to the sequence with 1 at index k and 0 elsewhere, every element $a = (a_n)_{n\in\mathbb{N}} \in R[[x]]$ can be uniquely written as:

$$a = \sum_{n=0}^{\infty} a_n X^n.$$

Remark 2.1.11. The definition of the product on $R[[X]]$ is motivated by the formal distribution of series multiplication. Since X is central, we can expand the product of two series as follows:

$$\left(\sum_{n=0}^{\infty} a_n X^n \right) \cdot \left(\sum_{m=0}^{\infty} b_m X^m \right) = \sum_{n,m} a_n b_m X^{n+m} = \sum_{k=0}^{\infty} \left(\sum_{n+m=k} a_n b_m \right) X^k.$$

Thus, the formal product is the unique multiplication compatible with the distribution of terms and the laws of exponents.

Definition 2.1.12. We denote by $R[X]$ the subring of $R[[X]]$ consisting of sequences $(a_n)_{n\in\mathbb{N}}$ with $a_n = 0$ for almost all n. It is called the **polynomial ring** over R.

Remark 2.1.13. We define the ring of formal power series in n variables inductively by

$$R[[X_1, \ldots, X_n]] := (R[[X_1, \ldots, X_{n-1}]])[[X_n]].$$

Analogously, the polynomial ring in n variables is defined inductively by:

$$R[X_1, \ldots, X_n] := (R[X_1, \ldots, X_{n-1}])[X_n].$$

Example 2.1.14. 1. If R is commutative, the **evaluation map**

$$\mathrm{ev} : R[x] \to \mathrm{Hom}_{\mathbf{Set}}(R, R), \quad P \mapsto f_P$$

is a **ring homomorphism**.

Proof. The preservation of addition is immediate. For multiplication, let $P = \sum a_n x^n$ and $Q = \sum b_m x^m$. For any $y \in R$, we have:

$$\begin{aligned} f_P(y) \cdot f_Q(y) &= \left(\sum_{n=0}^{\infty} a_n y^n\right) \cdot \left(\sum_{m=0}^{\infty} b_m y^m\right) \\ &= \sum_{s=0}^{\infty} \left(\sum_{n+m=s} a_n b_m\right) y^s \quad \text{(since } R \text{ is commutative)} \\ &= f_{P \cdot Q}(y). \end{aligned}$$

□

2. If R is an infinite field, the evaluation map is injective.

3. The polynomial ring $\mathbb{Z}[x_1, \ldots, x_n]$ satisfies the following universal property: Let R be a commutative ring. There is a natural bijection:

$$\mathrm{Hom}_{\mathbf{Ring}}(\mathbb{Z}[x_1, \ldots, x_n], R) \xrightarrow{\sim} R^n$$

given by evaluating the homomorphism at the generators:

$$\varphi \longmapsto (\varphi(x_1), \ldots, \varphi(x_n)).$$

Definition 2.1.15. Let R be a ring and let $P = \sum_{n=0}^{\infty} a_n X^n \in R[X]$ be a polynomial.

1. The **degree** of P, denoted $\deg P$, is defined as:

$$\deg P := \begin{cases} \max\{n \in \mathbb{N} \mid a_n \neq 0\} & \text{if } P \neq 0, \\ -\infty & \text{if } P = 0. \end{cases}$$

2. If $P \neq 0$ and $d = \deg P$, the coefficient a_d is called the **leading coefficient** of P.

3. The polynomial P is called **monic** if its leading coefficient is 1.

Remark 2.1.16. Let R be an integral domain. Then, for $P, Q \in R[X]$: $\deg(P \cdot Q) = \deg P + \deg Q$. *(proof: exercises)*

In particular, $R[X]$ is also an integral domain.

2.1.2 Field of Fractions of an Integral Domain

Definition 2.1.17. Let R be an integral domain. Consider the set of pairs $M = R \times (R \setminus \{0\})$. We define an equivalence relation $\sim$ on M by:

$$(a, b) \sim (c, d) \iff ad = bc.$$

The **field of fractions** of R, denoted $\mathrm{Frac}(R)$ is the set of equivalence classes $M/\sim$. We denote the class of (a, b) by the fraction $\frac{a}{b}$.

We define addition and multiplication on $\mathrm{Frac}(R)$ as follows:

$$\frac{a}{b} + \frac{c}{d} := \frac{ad + bc}{bd}, \qquad \frac{a}{b} \cdot \frac{c}{d} := \frac{ac}{bd}.$$

Proposition 2.1.18. Let R be an integral domain and let $\sim$ and $\mathrm{Frac}(R)$ be defined as above.

1. The relation $\sim$ is indeed an equivalence relation.
2. With the defined operations, $\mathrm{Frac}(R)$ is a field.
3. The map $\iota : R \to \mathrm{Frac}(R)$ defined by $\iota(r) = \frac{r}{1}$ is an injective ring homomorphism.

Proof. 1. Reflexivity and symmetry are immediate from the commutativity of R. For transitivity, suppose $(a, b) \sim (c, d)$ and $(c, d) \sim (e, f)$. Then $ad = bc$ and $cf = de$. Multiplying the first equation by f and the second by b, we get:

$$adf = bcf \quad \text{and} \quad bcf = bde \implies adf = bde.$$

Since R is an integral domain, we can cancel d to obtain $af = be$, so $(a, b) \sim (e, f)$.

2. It is easy to verify that addition and multiplication are well-defined and satisfy the ring axioms. The additive neutral element is $\frac{0}{1}$ and the multiplicative neutral element is $\frac{1}{1}$. For any non-zero element $\frac{a}{b} \neq \frac{0}{1}$, we have $a \neq 0$ (since $a \cdot 1 \neq b \cdot 0$). Thus, (b, a) is a valid pair in $R \times (R \setminus \{0\})$.
$$\frac{a}{b} \cdot \frac{b}{a} = \frac{ab}{ba} = \frac{1}{1}.$$
Thus, every non-zero element has a multiplicative inverse.

3. The map ι preserves addition and multiplication directly from the definitions. For injectivity, let $r \in \ker(\iota)$. Then:
$$\frac{r}{1} = \frac{0}{1} \iff r \cdot 1 = 1 \cdot 0 \iff r = 0.$$
Since the kernel is trivial, ι is injective.

□

Proposition 2.1.19. Let R be an integral domain, L a field and $\varphi : R \hookrightarrow L$ an injective homomorphism.

Proposition 2.1.20 (Universal Property of the Field of Fractions)**.** Let R be an integral domain and let L be a field. Let $\varphi : R \to L$ be an injective ring homomorphism. Then there exists a unique ring homomorphism $\psi : \mathrm{Frac}(R) \to L$ such that the following diagram commutes (i.e., $\psi \circ \iota = \varphi$):

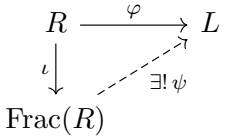

Proof. For any $x \in R$, we must have $\psi(\frac{x}{1}) = \varphi(x)$. For any $y \in R \setminus \{0\}$, since φ is injective and L is a field, $\varphi(y) \neq 0$, so $\varphi(y)$ is invertible in L. We define ψ by:

$$\psi\left(\frac{x}{y}\right) := \varphi(x) \cdot \varphi(y)^{-1}.$$

ψ is well-defined: suppose $\frac{a}{b} = \frac{c}{d}$. Then $ad = bc$. Applying φ:

$$\varphi(a)\varphi(d) = \varphi(b)\varphi(c).$$

Multiplying both sides by $\varphi(b)^{-1}\varphi(d)^{-1}$ yields $\varphi(a)\varphi(b)^{-1} = \varphi(c)\varphi(d)^{-1}$. Thus, the map is independent of the representative.

It is straightforward to verify that ψ is a homomorphism and is the unique solution. □

Remark 2.1.21. Intuitively, this states: "If a field L contains R, it must contain $\mathrm{Frac}(R)$."

Corollary 2.1.22. Let $\varphi : R \to S$ be an injective homomorphism between two integral domains. Then there exists a unique homomorphism $\mathrm{Frac}(\varphi) : \mathrm{Frac}(R) \to \mathrm{Frac}(S)$ making the following diagram commute:

$$\begin{array}{ccc} R & \xrightarrow{\quad\varphi\quad} & S \\ \big\downarrow & & \big\downarrow \\ \mathrm{Frac}(R) & \overset{\exists!\ \mathrm{Frac}(\varphi)}{\dashrightarrow} & \mathrm{Frac}(S) \end{array}$$

This is obtained by applying the universal property of $\mathrm{Frac}(R)$ to the composite map $R \to S \hookrightarrow \mathrm{Frac}(S)$.

Example 2.1.23.

1. The field of fractions of the integers is the field of rational numbers:

$$\mathrm{Frac}(\mathbb{Z}) = \mathbb{Q}.$$

2. Let K be a field. The field of fractions of the polynomial ring is called the **rational function field**:

$$K(X) := \mathrm{Frac}(K[X]).$$

Similarly, for n variables:

$$K(X_1, \ldots, X_n) := \mathrm{Frac}(K[X_1, \ldots, X_n]).$$

Note that the inclusion of rings induces an inclusion of fields. Since $K[X_1, \ldots, X_{n-1}] \subset K[X_1, \ldots, X_n]$, we have:

$$K(X_1, \ldots, X_{n-1}) \subset K(X_1, \ldots, X_n).$$

2.2 Ideals, Quotient Rings

In the following, let R be a commutative Ring.

Definition 2.2.1. Let $\mathfrak{a} \subseteq R$. $\mathfrak{a}$ is called an **ideal** if

1. $(\mathfrak{a}, +)$ is a subgroup of $(R, +)$.
2. For all $r \in R$ and $x \in \mathfrak{a}$, we have $rx \in \mathfrak{a}$.

Remark 2.2.2.

1. Let $\mathfrak{a} \subseteq R$ be an ideal. If $1 \in \mathfrak{a}$, $\mathfrak{a} = R$. Ideals that do not contain 1_R are called proper ideals.
2. Let K be a field, $\mathfrak{a} \subseteq K$ an ideal. Then, $\mathfrak{a} = 0$ or $\mathfrak{a} = K$.

Proof. 1. Let $r \in R$. $r = 1 \cdot r \in \mathfrak{a}$, so $\mathfrak{a} = R$.

2. Obviously, 0 is an ideal of K. If there is $x \neq 0 \in \mathfrak{a}$, because K is a field, $x^{-1} \in K$, $1 = xx^{-1} \in K$.

□

Example 2.2.3. 1. $\forall n \in \mathbb{Z} : n\mathbb{Z}$ is an ideal.

2. Let $a \in R$. $(a) := \{b \in R \mid \exists c \in R : b = ca\}$ is an ideal, called **principal ideal** generated by a.

3. If $\mathfrak{a}, \mathfrak{b}$ are ideals in R, so are $\mathfrak{a} + \mathfrak{b}$ and $\mathfrak{a} \cap \mathfrak{b}$.

4. Let G be a group. The **augmentation map** is the ring homomorphism $\varepsilon : \mathbb{Z}[G] \to \mathbb{Z}$ defined by summing the coefficients:

$$\varepsilon\left(\sum_{g \in G} a_g[g]\right) = \sum_{g \in G} a_g.$$

On basis elements, this simply looks like $\varepsilon([g]) = 1$ for all $g \in G$.

The **augmentation ideal**, denoted $I(G)$, is the kernel of this map:

$$I(G) := \ker(\varepsilon) = \left\{\sum a_g[g] \in \mathbb{Z}[G] \;\middle|\; \sum a_g = 0\right\}.$$

Definition 2.2.4. Let $(a_k)_{k \in A}$ be a family of elements in R.

$$((a_k)_{k \in A}) := \{x \in R \mid \exists (\lambda_k)_{k \in A} \in R^A : \lambda_k = 0 \text{ for almost all k } \in A,\ x = \sum_{k \in A} \lambda_k \cdot a_k\}$$

is an ideal, called the ideal generated by $(a_k)_{k \in A}$.

Definition 2.2.5. Let $\mathfrak{a}, \mathfrak{b}$ be ideals in R. We define:

$$\mathfrak{a} + \mathfrak{b} = \{a + b \mid a \in \mathfrak{a}, b \in \mathfrak{b}\},\ \mathfrak{a} \cdot \mathfrak{b} = \{\sum_{k=1}^{n} a_k \cdot b_k \mid n \in \mathbb{N}, a_k \in \mathfrak{a}, b_k \in \mathfrak{b}\}$$

Remark 2.2.6. Note that $\mathfrak{a}+\mathfrak{b}$ is "larger" than $\mathfrak{a}$ and $\mathfrak{b}$, since both are contained in their sum, but $\mathfrak{a} \cdot \mathfrak{b}$ is "smaller"; it is in fact easy to see that $\mathfrak{a} \cdot \mathfrak{b} \subseteq \mathfrak{a} \cap \mathfrak{b}$.

Lemma 2.2.7. Let $\varphi : R \to S$ be a ring homomorphism.
Then, $\ker \varphi = \{r \in R \mid \varphi(r) = 0_S\}$ is an ideal.

Proof. Since $\ker\varphi$ is the kernel of the underlying group homomorphism $(R, +) \to (S, +)$, $\ker\varphi$ is already an additive subgroup.
Let $r \in R, x \in \ker \varphi$. $\varphi(rx) = \varphi(r) \cdot 0_S = 0_S \implies rx \in \ker \varphi$. □

Definition 2.2.8. An ideal $\mathfrak{a} \subseteq R$ is said to be of **finite type** if

$$\exists (a_1, \ldots, a_n) \in R^n : (a_1, \ldots, a_r) = \mathfrak{a}.$$

Note that $(a_1, \ldots, a_n)$ refers to an n-tuple on the left, and a generated ideal on the right.
R is called **noetherian** if all of its ideals are of finite type.

Definition 2.2.9. R is called a **principal ideal domain** (PID) if R is an integral domain and any ideal is principal (generated by one element).

Theorem 2.2.10 (Definition of the quotient ring)**.** Let $\mathfrak{a} \subseteq R$ be an ideal. On the quotient abelian group $(R/\mathfrak{a}, +)$ ($(\mathfrak{a}, +)$ is normal in $(R, +)$ since both groups are abelian), there is a unique ring structure which makes the group canonical projection $\pi : R \to R/\mathfrak{a}$ a ring homomorphism.

Proof. Let $\pi : R \to R/\mathfrak{a}$ be the canonical projection.
We require $\cdot_{R/\mathfrak{a}} : ((R/\mathfrak{a}) \setminus \{0\})^2 \to R/\mathfrak{a},\ (a + \mathfrak{a}) \cdot (b + \mathfrak{a}) = (a \cdot_R b) + \mathfrak{a}$ for all $a,\ b \in R$ for π to be a ring homomorphism. Thus, our desired ring structure is unique.
We have to show that the operation is well-defined:
Let $a,\ \tilde{a} \in a + \mathfrak{a},\ b,\ \tilde{b} \in b + \mathfrak{a}$, so there are $x,\ y \in \mathfrak{a}$ such that $a + x = \tilde{a}$, $b + y = \tilde{b}$. To show: $\tilde{a}\tilde{b} \in ab + \mathfrak{a}$.

$$\tilde{a}\tilde{b} + \mathfrak{a} = (a + x)(b + y)\mathfrak{a} = ab + ay + bx + \underbrace{xy}_{\in \mathfrak{a}}$$

Because $\mathfrak{a}$ is an ideal, we get $ay,\ bx \in \mathfrak{a}$, so the entire expression is indeed element of $ab + \mathfrak{a}$. □

Proposition 2.2.11 (Universal property of the quotient ring)**.** Let $\mathfrak{a} \subseteq R$ be an ideal, S a commutative ring, $\pi : R \to R/\mathfrak{a}$ the canonical projection. Then,

$$\psi : \mathrm{Hom}_{\mathbf{Ring}}(R/\mathfrak{a}) \to \mathrm{Hom}_{\mathbf{Ring}}(R, S),\ \bar{\varphi} \mapsto \bar{\varphi} \circ \pi$$

is injective with image $\{\varphi \in \mathrm{Hom}_{\mathbf{Ring}}(R, S) \mid \varphi|_{\mathfrak{a}} = 0\}$.
In other words: If $\varphi \in \mathrm{Hom}_{\mathbf{Ring}}(R, S)$ with $\mathfrak{a} \subseteq \ker \varphi$, there is a unique $\bar{\varphi} \in \mathrm{Hom}_{\mathbf{Ring}}(R/\mathfrak{a}, S)$ such that the following diagram commutes:

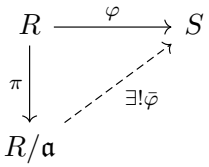

Proof. We want $\bar{\varphi}$ to map $r + \mathfrak{a} \in R/\mathfrak{a}$ to $\varphi(r) \in S$, then we would have

$$\bar{\varphi} \circ \pi(r) = \bar{\varphi}(r + \mathfrak{a}) = \varphi(r).$$

for all $r \in R$. Since we are mapping out of a quotient, we have to check whether this is well-defined: Let $\tilde{r} \in r + \mathfrak{a}$, so $r + \mathfrak{a} = \tilde{r} + \mathfrak{a}$ in $R/\mathfrak{a}$. It follows that there is an $a \in \mathfrak{a}$ with $\tilde{r} = r + a$, so

$$\bar{\varphi}(\tilde{r} + \mathfrak{a}) = \bar{\varphi}((r + a) + \mathfrak{a}) = \bar{\varphi}(r + \mathfrak{a}).$$

It is also clear that $\bar{\varphi}$ is a ring homomorphism.
Since π is surjective, the image of each element in $R/\mathfrak{a}$ is already defined by φ, making $\bar{\varphi}$ unique. □

Corollary 2.2.12 (Isomorphism theorem for rings)**.** Let $\varphi : R \to S$ be a ring homomorphism. Then,

$$R/\ker \varphi \cong \mathrm{im}\, \varphi.$$

Proof. Let $K = \ker \varphi$, $I = \mathrm{im}\, \varphi$.
The proof of the universal property already gives us a well-defined homomorphism $\bar{\varphi} : R/K \to I$, $r + K \mapsto \varphi(r)$.
Because

$$k + K \in \ker \bar{\varphi} \Leftrightarrow \varphi(k) = \bar{\varphi}(k + K) = 0_S \Leftrightarrow k \in K \Leftrightarrow k + K = K = 0_{R/K},$$

$\bar{\varphi}$ is injective. Because for all $\varphi(r) \in I$, $\bar{\varphi}(r + K) = \varphi(r)$, $\bar{\varphi}$ is an isomorphism between $R/\ker \varphi$ and $\mathrm{im}\, \varphi$. □

Definition 2.2.13. An ideal $\mathfrak{p}$ is called a **prime ideal** if $R/\mathfrak{p}$ is an integral domain. $\mathrm{Spec}(R)$ denotes the set of all prime ideals in R.
A proper ideal $\mathfrak{m}$ is called a **maximal ideal** if for all ideals $\mathfrak{a}$:

$$\mathfrak{m} \subseteq \mathfrak{a} \subseteq R \implies \mathfrak{a} = \mathfrak{m} \vee \mathfrak{a} = R.$$

Remark 2.2.14.

1. If $\mathfrak{p}$ is a prime ideal in R, $R/\mathfrak{p}$ is an integral domain, so in particular, $\mathfrak{p} = 0_{R/\mathfrak{p}} \neq 1_{R/\mathfrak{p}} \ni 1_R$, so $1 \notin \mathfrak{p}$, so $\mathfrak{p}$ is a proper ideal.
2. Let $a, b \in R$. If $a \cdot b \in \mathfrak{p}$,

   $$(a + \mathfrak{p}) \cdot (b + \mathfrak{p}) = a \cdot b + \mathfrak{p} = \mathfrak{p} = 0_{R/\mathfrak{p}},$$

   and since $R/\mathfrak{p}$ is an integral domain, either $a + \mathfrak{p} = 0_{R/\mathfrak{p}}$ or $b + \mathfrak{p} = 0_{R/\mathfrak{p}}$, so $a \in \mathfrak{p} \vee b \in \mathfrak{p}$.
3. Let $\mathcal{I}(R)$ denote the set of all proper ideals in R. It is partially ordered by set inclusion $\subseteq$. Maximal ideals are the maximal elements in $\mathcal{I}(R)$. One can prove the existence of maximal ideals in $R \neq 0$ using Zorns Lemma.

Lemma 2.2.15. Let $\mathfrak{m} \subset R$ be an ideal. Then, $\mathfrak{m}$ is maximal in R if and only if $R/\mathfrak{m}$ is a field.

Proof. In the exercises, it was proven that surjective homomorphisms preserve ideals.
Also, if we have ideals $\mathfrak{a} \subseteq R,\ \mathfrak{b} \subseteq R/\mathfrak{a}$, the pullback $\pi^{-1}(\mathfrak{b})$ is an ideal in R containing $\mathfrak{a}$.

It is then easy to see that the map

$$\Phi : \mathcal{I}(R/\mathfrak{a}) \to \mathcal{I}(R),\ \mathfrak{b} \mapsto \pi^{-1}(\mathfrak{b})$$

is injective,
inducing a bijection on its image $\{\tilde{\mathfrak{a}} \subseteq R \mid \mathfrak{a} \subseteq \tilde{\mathfrak{a}}\}$.

($\Rightarrow$): Let $\mathfrak{m}$ be an ideal in R. If $R/\mathfrak{m}$ is a field, any ideal $\mathfrak{a}$ in $R/\mathfrak{m}$ is either 0 or $R/\mathfrak{m}$,
so $\Phi(\mathfrak{a}) = \pi^{-1}(\mathfrak{a})$ is either $\pi^{-1}(0) = \mathfrak{m}$ or $\pi^{-1}(R/\mathfrak{m}) = R$, which are precisely the ideals in R containing $\mathfrak{m}$.
Thus, $\mathfrak{m}$ is maximal in R.

($\Leftarrow$): Let $r + \mathfrak{a}$ be a nonzero element in the quotient $R/\mathfrak{m}$.
This means that $r \in R$ and $r \notin \mathfrak{m}$. Consider $\mathfrak{b} := (r) + \mathfrak{m}$.
Since $r \in \mathfrak{b},\ r \notin \mathfrak{m} \implies \mathfrak{m} \subsetneq \mathfrak{b}$ and $\mathfrak{m}$ maximal, we can conclude that $\mathfrak{b} = R$ and in particular $1_R \in \mathfrak{b}$.
Hence, we can write $1 = u \cdot r + m,\ m \in \mathfrak{m},\ u \in R$, which becomes $\overline{1} = \overline{u} \cdot \overline{r} + \overline{0}$ in $R/\mathfrak{m}$, so we have found an inverse u for all $\overline{r} \in R/\mathfrak{m}$. □

2.2.1 Euclidean division in $R[X]$

Theorem 2.2.16. Let R be a commutative ring, $A \in R[x] \setminus \{0\}, d := \deg(A) \in \mathbb{N}$. Assume the leading coefficient a_d of A is invertible. Then,

$$\forall B \in R[x]\ \exists! Q, R \in R[x] : B = Q \cdot A + R,\ \deg(R) < d.$$

Proof. Existence:
We proceed by strong induction on n = $\deg(B)$.
Base Case: If $\deg(B) < d$, we can simply choose $Q = 0,\ R = B$.
Inductive Step: Let $\deg(B) = n,\ b_n$ being the leading coefficient of B.

We want to eliminate the leading term $b_n X^n$ of B by subtracting a multiple of A. We define the monomial term

$$T(X) = b_n \cdot a_d^{-1} \cdot X^{n-d}.$$

Now, consider the polynomial

$$B'(X) = B(X) - T(X) \cdot A(X).$$

Notice that the leading term of $T(X) \cdot A(X)$ is $(b_n \cdot a_d^{-1} \cdot a_d)X^{(n-d)+d} = b_n X^n$, which is the leading term of $B(X)$ as well.
We can conclude that $\deg(B'(X)) < n$, so we apply our induction hypothesis to B': There exist $Q', R' \in R[x]$, such that

$$B' = Q' \cdot A + R', \ \deg(R') < d.$$

If we substitute the definition of B', we get

$$B - T \cdot A = Q' \cdot A + R' \ \Leftrightarrow B = (Q' + T) \cdot A + R', \deg(R) < d.$$

Uniqueness:
Suppose (Q, R) as well as $(\tilde{Q}, \tilde{R})$ satisfy the conditions above, so

$$B = Q \cdot A + R = \tilde{Q} \cdot A + \tilde{R}, \deg(R), \deg(\tilde{R}) < d.$$

This is equivalent to

$$0 = (Q - \tilde{Q}) \cdot A + (R - \tilde{R}) \Leftrightarrow (Q - \tilde{Q})A = \tilde{R} - R.$$

Let $\mathcal{Q} = Q - \tilde{Q}$, $\mathcal{R} = \tilde{R} - R$, so

$$\mathcal{Q}A = \mathcal{R}.$$

Suppose for the sake of contradiction that $Q \neq 0$. Since the leading coefficient a_d is invertible and thus not a zero divisor, we get

$$\deg(\mathcal{Q}A) = \deg(\mathcal{Q}) + \deg(A) = \deg(\mathcal{Q}) + d.$$

Since $\deg(\mathcal{Q}) \geq 0$, we get $deg(\mathcal{Q}A) \geq d$. But since obviously $\deg(\mathcal{R}) < d$, we have a contradiction. Thus, $\mathcal{Q} = 0$ and consequently $\mathcal{R} = 0 \cdot A = 0$. □

Corollary 2.2.17. Let K be a field, $P \in K[X] \setminus \{0\}$.
Then, $K[X]/(P)$ is a K-vector space with basis $B = \{\overline{1}, \overline{x}, \ldots, \overline{x}^{\deg(P)-1}\}$.

Proof. We will prove a more general statement using the Euclidean division later. □

Definition 2.2.18. Let I be an integral domain. If there is a norm function $N : R \setminus \{0\} \to \mathbb{N}_0$ such that for all $a, b \in R$, $b \neq 0$ there are $q, r \in R$ such that $a = bq + r$ with $(r = 0) \vee (N(r) < N(b))$, I is called a **Euclidean domain**.

Corollary 2.2.19. Every Euclidean domain is a principal ideal domain.

Proof. Let R be a Euclidean domain with norm N and an ideal $\mathfrak{a}$. We must show that $\mathfrak{a}$ is principal.
If $\mathfrak{a} = 0$, $\mathfrak{a}$ is obviously principal. Suppose $\mathfrak{a} \neq 0$. Consider the set of the norms of all nonzero elements in $\mathfrak{a}$:

$$S = \{N(x) \,|\, x \in \mathfrak{a} \setminus \{0\}\}.$$

S is a nonempty subset of $\mathbb{N}$, so by the well-ordering principle, it must have a least element $N(d)$, $d \in \mathfrak{a}$. We show that $(d) = \mathfrak{a}$. Clearly, $(d) \subseteq \mathfrak{a}$. Let $a \in \mathfrak{a}$. Since R is a Euclidean domain, there are $q, r \in R$ such that

$$a = d \cdot q + r \Leftrightarrow r = a - d \cdot q,$$

where $r = 0$ or $r < d$. We know that $a \in \mathfrak{a}$ and $d \in \mathfrak{a}$, so r must be in $\mathfrak{a}$.
If $r \neq 0$, we would have $r < d$, contradicting our assumption that $N(d)$ is minimal. Thus, r must be 0, so $a = d \cdot q$ and we get $\mathfrak{a} \subseteq (d)$. □

Corollary 2.2.20. Let $P \in R[X]$, $k \in R$. Then,

$$P(k) = 0 \Leftrightarrow (X - k) \mid P.$$

Proof. Euclidean division of P by $(X - k)$, which is allowed because the leading coefficient of $X - k$ is invertible. We get

$$P = Q(X - k) + r.$$

Since $\deg(X - k) = 1$, $\deg(r) < 1$, so $r \in R$, and $P(k) = 0 \Leftrightarrow r = 0 \Leftrightarrow (X - k) \mid P$. □

Definition 2.2.21. Let A be a commutative ring, $P \in A[X]$. We define

$$R_P(A) := \{x \in A \mid P(x) = 0\}$$

to be set of roots of P in A.

Corollary 2.2.22. Let K be a field, $P \in K[X] \setminus \{0\}, \deg(P) = n(\geq 0)$. Then,

$$|R_P(K)| \leq n.$$

Proof. Induction on n. If $n = 0$, R is a nonzero constant and thus has no roots.
If $n = 1$, we can assume $P = aX + b$, where $a \neq 0$, so a invertible.
Assume that for all $P' \in K[X]$ with $\deg(P') = n' < n$, $|R_P(K)| \leq n'$. Let $\deg(P) = n$. If $R_P(K) = \varnothing$, there is nothing left to show. If not, let $x_1 \in P_R(K)$. From the previous corollary, we know that $(X - x_1) \mid P$, so there is a $P_1 \in K[X]$ such that

$$P = (X - x_1) \cdot P_1.$$

It is easy to see that $\deg(P_1) = \deg(P) - 1$, and

$$P(x) = 0 \Leftrightarrow x = x_1 \vee x \in R_{P_1}(K) \implies R_P(K) = \{x_1\} \cup R_{P_1}(K).$$

By our inductive hypothesis, we have

$$|R_P(K)| = |R_{P_1}(K)| + 1 \leq (n - 1) + 1 = n.$$

□

Remark 2.2.23. This holds for polynomial rings over an integral domain A as well, since for $P \in A[X] \setminus \{0\}$, $\deg(P) = n$:

$$|R_P(A)| \leq |R_P(\mathrm{Frac}(A))| \leq n.$$

Corollary 2.2.24. Let K be a field, $P \in K[X] \setminus \{0\}$, $|R_P(K)| = n$. Then,

$$P = a_n \prod_{r \in R_P(K)} (X - r),$$

where a_n is the leading coefficient of P.

Proof. Consider the polynomial $N = P - a_n \prod_{r \in R_P(K)}(X - r)$. Since the n-th terms of both P and $a_n \prod_{r \in R_P(K)}(X - r)$ are $(a_n X^n)$, it follows that
$\deg(N) \leq n - 1$. But clearly,

$$\forall y \in R_P(K) : N(y) = 0, |R_P(K)| = n > n - 1,$$

so N must be 0, otherwise this would contradict the previous corollary. □

Example 2.2.25. Consider $X^n - 1 \in \mathbb{C}[X]$. Let $\zeta = e^{\frac{2\pi i}{n}}$.
Since for all $i < n$, $(\zeta^i)^n = (\zeta^n)^i = 1^i = 1$, $R_{X^n-1}(\mathbb{C}) = \{1, \zeta, \ldots, \zeta^{n-1}\}$. By the previous corollary, $X^n - 1 = \prod_{k=0}^{n-1}(X - \zeta^k)$ in $\mathbb{C}[X]$.

Definition 2.2.26. Let K be a field.

$$\mu_n(K) = R_{X^n-1}(K) = \{x \in K \mid x^n = 1\}$$

is the set of the n-th roots of unity in K.

2.2.2 Relations Between Roots and Coefficients of Polynomials

Let K be a field. Consider a monic polynomial $P \in K[X]$ of degree $n \geq 1$. We can express P in its expanded coefficient form

$$P(X) = X^n + a_{n-1}X^{n-1} + \cdots + a_1X + a_0 = \sum_{k=0}^{n} a_k X^k$$

and its factored form (assuming it splits completely in K)

$$P(X) = \prod_{i=1}^{n}(X - \alpha_i)$$

Definition 2.2.27. Let $n \geq 1$, $i \in \{1, \ldots, n\}$. The **i-th elementary symmetric function** Σ_i is defined as:

$$\Sigma_i = \sum_{\substack{J \subseteq \{1,\ldots,n\} \\ |J|=i}} X^J \in \mathbb{Z}[X_1, \ldots, X_n], \text{ where } X^J := \prod_{j \in J} X_j.$$

It is clear that this is equivalent:

$$\Sigma_i = \sum_{1 \leq j_1 < j_2 < \cdots < j_i \leq n} X_{j_1} X_{j_2} \ldots X_{j_k}$$

Theorem 2.2.28. For the polynomial $P(X) = \sum_{i=0}^{n} a_i X^i$ with roots $\alpha_1, \ldots, \alpha_n$, the coefficients are given by:

$$a_i = (-1)^{n-i}\Sigma_{n-i}(\alpha_1, \ldots, \alpha_n)$$

for all $0 \leq i < n$.

Proof. We proceed by expanding the factored form of the polynomial and comparing coefficients with the standard form.

Consider the product:

$$P(X) = (X - \alpha_1)(X - \alpha_2)\ldots(X - \alpha_n)$$

To determine the coefficient of a specific term X^k in this expansion, we must choose X from exactly k of the factors,

and choose the constant term $(-\alpha_j)$ from the remaining $n - k$ factors.

Let $m = n - k$ be the number of roots chosen. The term involving X^k is formed by summing over all possible combinations of choosing m distinct roots.

For a specific choice of indices $J = \{j_1, \ldots, j_m\}$ with $|J| = m$, the contribution to the product is:

$$(-\alpha_{j_1})(-\alpha_{j_2})\ldots(-\alpha_{j_m})X^{n-m} = (-1)^m(\alpha_{j_1}\ldots\alpha_{j_m})X^{n-m}$$

Summing over all such subsets $J \subseteq \{1, \ldots, n\}$ of size m, the total term is:

$$\left((-1)^m \sum_{|J|=m} \prod_{j \in J} \alpha_j\right) X^{n-m}$$

Recognizing the inner sum as the elementary symmetric polynomial Σ_m, the coefficient of X^{n-m} is:

$$(-1)^m\Sigma_m$$

To match the index notation of the theorem, let i be the power of X, so $i = n - m$, which implies $m = n - i$.

Substituting this back into our expression for the coefficient a_i:

$$a_i = (-1)^{n-i}\Sigma_{n-i}.$$

□

Remark 2.2.29. Let $\sigma \in S_n$. A map

$$K[X_1, \ldots, X_n] \to K[X_1, \ldots, X_n],\ P(X_1, \ldots, X_n) \mapsto P(X_{\sigma(1)}, \ldots, P(X_{\sigma(n)})) := P^\sigma$$

is an automorphism of the polynomial ring. $\forall i \leq n : \Sigma_i^\sigma = \Sigma_i$.

2.3 R-Modules

Definition 2.3.1. Let R be a commutative ring. An **R-module** M is a triple $(M, +, \cdot)$ with

$$+ : M^2 \to M,\ \cdot : R \times M \to M,$$

such that: $\forall r, s \in R, m, n \in M$:

1. $(M, +)$ is an abelian group with neutral element 0
2. $(r + s) \cdot m = r \cdot m + s \cdot m$
3. $r \cdot (m + n) = r \cdot m + r \cdot n$
4. $(r \times s) \cdot m = r \cdot (s \cdot m)$, where $\times$ refers to the multiplication in R
5. $1_R \cdot m = m$

Example 2.3.2.

1. If $R = K$ is a field, any R-module is a K-vector space.
2. Any abelian group $(A, +)$ admits a unique $\mathbb{Z}$-module structure determined by $1 \cdot a = a$.
3. Let M be a module over the polynomial ring $K[X]$. This single object packages the data of a vector space and a linear map into one structure:

   Since $K \subset K[X]$, the module M is automatically a vector space V over K.

   $X \in K[X]$ must act on the vectors in V. We define this action as a map $T : V \to V$:

   $$T(v) := X \cdot v$$

   Crucially, because X commutes with scalars in the polynomial ring ($X\lambda = \lambda X$), the map T preserves scalar multiplication:

   $$T(\lambda v) = X \cdot (\lambda v) = (X\lambda) \cdot v = (\lambda X) \cdot v = \lambda \cdot (Xv) = \lambda T(v)$$

   Thus, T is a K-linear endomorphism.

Definition 2.3.3. Let M and N be R-modules. A map $\varphi : M \to N$ is called a homomorphism of R-modules or an **R-linear map** if it is a morphism of the additive groups and

$$\forall \lambda \in R \forall x \in M : \varphi(\lambda \cdot_M x) = \lambda \cdot_N \varphi(x)$$

Remark 2.3.4. One may define the category of R-modules R-**Mod**.

Example 2.3.5. Let $\varphi : R \to S$ be a commutative ring homomorphism, M be an S-module. We may consider the R-module M^φ with multiplication defined as:

$$\lambda \cdot_{M^\varphi} x = \varphi(\lambda) \cdot_M .$$

Definition 2.3.6. Let M be an R-module. A **sub-R-module** N is a subgroup of the additive group with the following property:

$$\forall \lambda \in R \forall x \in N : \lambda \cdot x \in N.$$

Example 2.3.7. R is trivially an R-module itself. Submodules are precisely the ideals in R.

Theorem 2.3.8. Let M be an R-module with sub-R-module N. On the quotient of the underlying abelian groups $(M,+)/(N,+)$, there exactly one map
$\cdot : R \times (M,+)/(N,+) \to (M,+)/(N,+)$ that makes the canonical projection $\pi : M \to M/N$ R-linear. The resulting module is called the **quotient module**.

Proof. Let $a+N, b+N \in M/N$. We define $\cdot : ((\lambda, m+N)) \mapsto (\lambda \cdot m)+N$. It is easy to check that it is well-defined, makes π R-linear and imposes a module structure on N/M. □

Example 2.3.9. If $f : M \to N$ is R-linear, $\ker(f)$ is a sub-R-module of M, $\mathrm{im}(f)$ is a sub-R-module of N. The quotient

$$N/\mathrm{im}(f) := \mathrm{coker}(f)$$

is called the **cokernel** of f. The kernel of the projection $N \to \mathrm{coker}(f)$ is $\mathrm{im}(f)$.

Definition 2.3.10. Let $(M_i)_{i\in\mathcal{I}}$ be a family of R-modules. We define the R-modules

$$\prod_{i\in\mathcal{I}} = \{(m_i)_{i\in\mathcal{I}} \,|\, \forall i \in \mathcal{I} : m_i \in M\},$$

called the **direct product** of $(M_i)_{i\in\mathcal{I}}$ and

$$\bigoplus_{i\in\mathcal{I}} M_i = \{(m_i)_{i\in\mathcal{I}} \,|\, \forall i \in \mathcal{I} : m_i \in M,\ m_i = 0 \text{ for almost all } i\},$$

called the **direct sum**.

Remark 2.3.11. It is clear that the direct sum is a submodule of the direct product. They are equal if and only if the index set $\mathcal{I}$ is finite.

Lemma 2.3.12 (Universal properties of the direct sum and direct product)**.** Let N be an R-module, $(M_i)_{i\in\mathcal{I}}$ a family of R-modules. There are bijections between the sets

$$\mathrm{Hom}_R(N, \prod_{i\in\mathcal{I}} M_i) \text{ and } \prod_{i\in\mathcal{I}} \mathrm{Hom}_R(N, M_i)$$

as well as

$$\mathrm{Hom}_R(\bigoplus_{i\in\mathcal{I}} M_i, N) \text{ and } \prod_{i\in\mathcal{I}} \mathrm{Hom}_R(M_i, N).$$

Remark 2.3.13. Categorically, The direct product is the universal object designed to be easy to map into, whereas the direct sum is the universal object designed to be easy to map out of.

Proof of the lemma. For the direct product, Let $\pi_k : \prod M_i \to M_k$ denote the canonical projection onto the k-th factor. We construct the bijection by defining maps in both directions.

Define $\Phi : \mathrm{Hom}_R(N, \prod M_i) \to \prod \mathrm{Hom}_R(N, M_i)$ by composing with the projections:

$$\Phi(f) = (\pi_i \circ f)_{i\in\mathcal{I}}.$$

Define $\Psi : \prod \mathrm{Hom}_R(N, M_i) \to \mathrm{Hom}_R(N, \prod M_i)$ as follows. Given a family of maps $(f_i)_{i\in I}$, we define the map $\Psi((f_i))$ by its action on an element $n \in N$:

$$\Psi((f_i)_{i\in\mathcal{I}})(n) = (f_i(n))_{i\in\mathcal{I}}.$$

This map is well-defined because the product allows arbitrary tuples.

The maps are inverse to each other:

$$(\Phi \circ \Psi)((f_i)_i) = \Phi(n \mapsto (f_i(n))_i) = (\pi_k \circ (n \mapsto (f_i(n))_i))_k = (f_k)_k.$$

$$(\Psi \circ \Phi)(f) = \Psi((\pi_i \circ f)_i) = (n \mapsto (\pi_i(f(n)))_i).$$

Since a tuple is determined by its components, this recovers f. l. Thus, Φ and Ψ are inverses.

For the direct sum, Let $\iota_k : M_k \to \bigoplus M_i$ denote the canonical injection of the k-th summand. Define $\Phi : \mathrm{Hom}_R(\bigoplus M_i, N) \to \prod \mathrm{Hom}_R(M_i, N)$ by composing with the injections:

$$\Phi(F) = (F \circ \iota_i)_{i \in \mathcal{I}}.$$

Define $\Psi : \prod \mathrm{Hom}_R(M_i, N) \to \mathrm{Hom}_R(\bigoplus M_i, N)$ as follows. Given a family of maps $(g_i)_{i \in \mathcal{I}}$, we define the map $\Psi((g_i))$ by its action on a tuple $(m_i)_{i \in \mathcal{I}} \in \bigoplus M_i$:

$$\Psi((g_i)_{i \in I})((m_i)_{i \in \mathcal{I}}) = \sum_{i \in \mathcal{I}} g_i(m_i).$$

This map is well-defined because, by the definition of the direct sum, $m_i = 0$ for almost all i, making the sum finite. $(\Phi \circ \Psi)((g_i)_i)_k = \Psi((g_i)_i) \circ \iota_k$. For any $x \in M_k$, $\iota_k(x)$ has x at index k and 0 elsewhere. Thus, the sum collapses to $g_k(x)$. So the result is g_k.
$(\Psi \circ \Phi)(F)$ acts on $(m_i)_i$ as $\sum_i (F \circ \iota_i)(m_i)$. Since F is linear, this is $F(\sum_i \iota_i(m_i))$. Since $(m_i)_i = \sum_i \iota_i(m_i)$, this recovers $F((m_i)_i)$.
Thus, Φ and Ψ are inverses. □

2.3.1 Sub-R-modules generated by a Family

Definition 2.3.14. Let M be an R-module, $(M_i)_{i \in \mathcal{I}}$ a family of its submodules. Consider the direct sum $\bigoplus_{i \in \mathcal{I}} M_i$. For each $i \in \mathcal{I}$, we have a canonical inclusion $\iota_i : M_i \hookrightarrow M$, so by the universal property, there is a unique map

$$\bigoplus_{i \in \mathcal{I}} M_i \to M,\ (m_i)_{i \in \mathcal{I}} \mapsto \sum_{i \in \mathcal{I}} \iota_i(m_i).$$

We call its image the **sum** of $(M_i)_{i \in \mathcal{I}}$, denoted by $\sum_{i \in \mathcal{I}} M_i$.
In other words, $\sum_{i \in \mathcal{I}} M_i = \{\sum_{i \in I} m_i \,|\, m_i \in M_i,\ m_i = 0 \text{ for almost all } i\}$.
Now, let $(a_i)_{i \in \mathcal{I}}$ be a family of elements of M. It is easy to see that

$$\Phi_{(a_i)_{i \in \mathcal{I}}} : \bigoplus_{i \in \mathcal{I}} R \to M,\ (\lambda_i)_{i \in \mathcal{I}} \mapsto \sum_{i \in \mathcal{I}} \lambda_i a_i$$

is R-linear. We denote its image by $((a_i)_{i \in \mathcal{I}})$ and call it the **sub-R-module of M generated by** $(a_i)_{i \in \mathcal{I}}$.
If $(a_i)_{i \in \mathcal{I}} \subset M$, $((a_i)_{i \in \mathcal{I}}) = \sum_{i \in \mathcal{I}} (a_i)$.

Definition 2.3.15. Let M be an R-module, $(e_i)_{i \in \mathcal{I}}$ a family of elements, $\Phi_{(e_i)_{i \in \mathcal{I}}} : \bigoplus_{i \in \mathcal{I}} \to M$ defined as above. We say that:

1. M is **generated by** $(e_i)_{i \in \mathcal{I}}$ if $\Phi_{(e_i)_{i \in \mathcal{I}}}$ is an epimorphism.
2. M is a **free R-module with basis** $(e_i)_{i \in \mathcal{I}}$ if $\Phi_{(e_i)_{i \in \mathcal{I}}}$ is an isomorphism.
3. M is a **finite type** R-module if it has a finite family of generators.

Remark 2.3.16. It is easy to see that the injectivity of Φ_B is equivalent to the linear independence of B.

Example 2.3.17. 1. Let $P \in R[X]$ be monic, $\deg(P) = n$. $R[X]/(P)$ is a free R-module with basis $B := \{\bar{1}, \bar{X}, \ldots, \bar{X}^{n-1}\}$.

Proof. We show that B is a linearly independent generating family.
Let $\overline{f} \in R[X]/(P)$. Since P is monic, we can perform Euclidean division to write $f = qP+r$ with $\deg(r) < n$ (or $r = 0$).
In the quotient, $\overline{P} = \overline{0}$, so $\overline{f} = \overline{r}$. Since $\deg(r) \leq n-1$, we can write $r = \sum_{i=0}^{n-1} a_i X^i$ for some $a_i \in R$. Thus, $\overline{f} = \sum_{i=0}^{n-1} a_i \overline{X}^i$, showing that B generates M.
Suppose $\sum_{i=0}^{n-1} a_i \overline{X}^i = \overline{0}$, $a_i \in R$. Let $g = \sum_{i=0}^{n-1} a_i X^i$. The condition $\overline{g} = \overline{0}$ implies $g \in (P)$, so $g = h \cdot P$ for some $h \in R[X]$. If $h \neq 0$, then $\deg(g) = \deg(h) + \deg(P) \geq n$ (since P is monic). However, by construction, $\deg(g) \leq n-1$. This contradiction implies $h = 0$, and thus $g = 0$. Therefore, all coefficients a_i are 0. □

2. $P[X]$ is a free R-module with basis $\{X^n \mid n \in \mathbb{N}\}$.

3. Let R be a noetherian ring. Then, if M is a finite type R-module, any submodule of M is of finite type. *proof: exercises*

4. $R[X][Y]/(XY-1)$ is a free R-module on $\{X^i \mid i \in \mathbb{Z}\}$. *proof: exercises*

2.3.2 Structure of Modules over Principal Ideal Domains

Definition 2.3.18. Let R be an integral domain, let M be an R-module.

$$M_{\text{tor}} := \{m \in M \mid \exists \lambda \in R \setminus \{0\} : \lambda \cdot m = 0\} \subseteq M$$

is a sub-R-module, called the **torsion submodule**. If $M_{\text{tor}} = \varnothing$, M is called **torsion free**.

Theorem 2.3.19 (Structure Theorem for Finitely Generated Modules over a PID)**.** Let R be a principal ideal domain and let M be a finite type R-module. There exists a unique integer $n \geq 0$ and a unique decreasing sequence of ideals:

$$R \supsetneq I_1 \supseteq I_2 \supseteq \cdots \supseteq I_n$$

such that there is an isomorphism of R-modules:

$$M \cong \bigoplus_{i=1}^{n} R/I_i.$$

The ideals $I_1, \ldots, I_n$ are called the **invariant factors** of M.

Remark 2.3.20. Since R is a PID, for each I_i there is an $a_i \in R$ such that $I_i = (a_i)$. Since

$$I_i \supseteq I_{i+1} \implies a_{i+1} \in I_i \implies \exists r \in R : r \cdot a_i = a_{i+1} \Leftrightarrow a_i \mid a_{i+1},$$

the chain of ideals implies the divisibility condition

$$a_1 \mid a_2 \mid \cdots \mid a_n.$$

Note that while the sequence of ideals I_i is unique, the generators a_i are not (they are determined only up to multiplication by a unit).

Remark 2.3.21. Let $s \in \{1, \ldots, n\}$ be the smallest integer such that $I_s \neq 0$ and $I_{s+1} = 0$. We get the decomposition

$$M \cong \underbrace{\bigoplus_{i=1}^{s} R/I_i}_{\mathcal{T}} \oplus \underbrace{\bigoplus_{i=s+1}^{s+r} R/I_i (\cong R^r)}_{\mathcal{F}}.$$

Note that because $\mathcal{F}$ is a finite direct sum, it is just the product R^r, which has a canonical basis, so $\mathcal{F}$ is free. Since R is an integral domain, $\mathcal{F}$ is in particular torsion free.

Now, let x be any non-zero element of the smallest ideal, i.e., $x \in I_s \setminus \{0\}$. Then, $x \in I_k$ for all $k \in \{1, \ldots, s\}$.

Let $t = (\bar{r}_1, \ldots, \bar{r}_s) \in \mathcal{T}$. Then:

$$x \cdot t = (x \cdot \bar{r}_1, \ldots, x \cdot \bar{r}_s) = (\overline{xr_1}, \ldots, \overline{xr_s}).$$

Since $x \in I_k$ for every k, the product xr_k lies in I_k, so $\overline{xr_k} = \bar{0}$ in R/I_k.

Thus, $x \cdot t = 0$. Since there exists a non-zero scalar x that kills every element in $\mathcal{T}$, every element in $\mathcal{T}$ is a torsion element. Therefore, $\mathcal{T}$ is isomorphic to the torsion submodule M_{tor}.

Example 2.3.22. 1. Let A be a finitely generated abelian group (a $\mathbb{Z}$-module of finite type).

There exists a unique integer $r \geq 0$ and a unique sequence of integers $a_1, \ldots, a_n$ such that:

$$A \cong \mathbb{Z}^r \oplus \bigoplus_{i=1}^{n} \mathbb{Z}/(a_i) (= \mathbb{Z}/a_i\mathbb{Z})$$

subject to the condition that the a_i are positive, non-invertible integers satisfying the divisibility chain $a_1 \mid a_2 \mid \cdots \mid a_n$.

If A is a finite abelian group, there cannot be a free part since its order would then no longer be finite. Thus, a finite abelian group is purely torsion and isomorphic to a direct product of cyclic groups:

$$A \cong \mathbb{Z}/a_1\mathbb{Z} \times \mathbb{Z}/a_2\mathbb{Z} \times \cdots \times \mathbb{Z}/a_n\mathbb{Z}$$

subject to the unique sequence of integers a_i being positive, non-invertible, and satisfying $a_1 \mid a_2 \mid \cdots \mid a_n$.

We first prove the following theorem which will be helpful when proving the structure theorem:

Theorem 2.3.23 (Adapted Basis Theorem)**.** Let R be a principal ideal domain. Let F be a free R-module of rank n, and let $M \subseteq F$ be a submodule.

Then:

1. M is a free R-module of rank $m \leq n$.
2. There exists a basis $\mathcal{B} = \{e_1, \ldots, e_n\}$ of F and elements $a_1, \ldots, a_m \in R \setminus \{0\}$ such that $\{a_1e_1, \ldots, a_me_m\}$ is a basis of M with $a_1 \mid a_2 \mid \cdots \mid a_m$.

Remark 2.3.24. Let R be a nonzero commutative ring, $n, m \in \mathbb{N}_{>1}$ such that $\varphi : R^n \to R^m$ is an isomorphism of R-modules. Then, $n = m$. This means that any finite type free R-module has a well-defined rank.

Proof. Let $\varphi : R^n \xrightarrow{\cong} R^m$ be an isomorphism. Since $R \neq 0$, there exists a maximal ideal $\mathfrak{m} \subset R$. Let $k = R/\mathfrak{m}$ be the residue field (here, we use commutativity of R, otherwise k would not be a field).

Reducing coefficients modulo $\mathfrak{m}$ induces a linear map between the k-vector spaces:

$$\bar{\varphi} : k^n \to k^m$$

Since φ is an isomorphism with inverse ψ, the reduction $\bar{\varphi}$ has inverse $\bar{\psi}$, making $\bar{\varphi}$ an isomorphism of vector spaces.

Since isomorphic vector spaces have the same dimension, we conclude:

$$n = \dim_k(k^n) = \dim_k(k^m) = m.$$

□

Proof of the adapted basis theorem. The previous remark allows us to proceed by induction on the rank n of F.

We proceed by induction on the rank n of F. The case $n = 0$ is trivial.

Let $n \geq 1$. Consider the set of all scalars obtainable by evaluating linear forms of F on vectors in M. We define the set $\Gamma(M) \subseteq R$:

$$\Gamma(M) = \{\phi(x) \mid \phi \in F^*, x \in M\}$$

where $F^* = \mathrm{Hom}_R(F, R)$. It is easily verified that $\Gamma(M)$ is an ideal of R. Since R is a PID, this ideal is principal, generated by some element $a_1 \in R$.

$$\Gamma(M) = (a_1)$$

If $a_1 = 0$, since F is free, there are particularly projections onto single coordinates which must all be zero, implying that M is zero, so we can assume $a_1 \neq 0$.

Then, by definition, there exists a form $\phi_1 \in F^*$ and a vector $x_1 \in M$ such that $\phi_1(x_1) = a_1$. We claim that a_1 divides x_1 in F. Let $\{\epsilon_i\}$ be an arbitrary basis of F and $\{\epsilon_i^*\}$ the dual basis. For any coordinate j, the value $\epsilon_j^*(x_1)$ lies in $\Gamma(M) = (a_1)$, so a_1 divides every coordinate of x_1. Thus, there exists a unique $e_1 \in F$ such that:

$$x_1 = a_1 e_1$$

Evaluating our chosen form ϕ_1 on this equation:

$$a_1 = \phi_1(x_1) = \phi_1(a_1 e_1) = a_1 \phi_1(e_1)$$

Since R is a domain and $a_1 \neq 0$, we conclude $\phi_1(e_1) = 1$.

The condition $\phi_1(e_1) = 1$ implies that the map $\pi(y) = \phi_1(y)e_1$ is a projection onto the submodule generated by e_1. This yields a direct sum decomposition of F:

$$F = Re_1 \oplus \ker(\phi_1)$$

We obtain a compatible decomposition for M. Let $y \in M$. Then $\phi_1(y) \in \Gamma(M) = (a_1)$, so $\phi_1(y)$ is a multiple of a_1.

$$y = \underbrace{\phi_1(y)e_1}_{\in R(a_1 e_1)} + \underbrace{(y - \phi_1(y)e_1)}_{\in \ker(\phi_1)}$$

Thus, $M = R(a_1 e_1) \oplus (M \cap \ker(\phi_1))$.

Let $F' = \ker(\phi_1)$ and $M' = M \cap \ker(\phi_1)$. F' is free of rank $n-1$. By the induction hypothesis, there exists a basis $\{e_2, \ldots, e_n\}$ of F' and scalars $a_2 \mid \cdots \mid a_k$ such that $\{a_2 e_2, \ldots, a_k e_k\}$ is a basis of M'. Combining these, $\{e_1, \ldots, e_n\}$ is a basis of F, and the basis for M is $\{a_1 e_1, a_2 e_2, \ldots, a_k e_k\}$. Finally, to see that $a_1 \mid a_2$, observe that $\Gamma(M')$ is generated by restrictions of forms to M', so $\Gamma(M') \subseteq \Gamma(M)$. This implies $(a_2) \subseteq (a_1)$, or $a_1 \mid a_2$. □

Now we have all the tools needed to prove the structure theorem.

Existence proof of the structure theorem. Let M be a finite type R-module. By definition, there is an epimorphism

$$\pi : R^n \twoheadrightarrow M$$

.

Let $K = \ker(\pi) \subset R^n$. By the first isomorphism theorem, we have $M \cong R^n/K$.

Since R is a PID and R^n is free, we can apply the adapted basis theorem to the submodule K, so K is free of rank $m \leq n$ and there exists a basis $(e_1, \ldots, e_n)$ of R^n and scalars $a_1, \ldots, a_m \in R \setminus \{0\}$ such that:

1. $(a_1e_1, \ldots, a_me_m)$ is a basis of K.
2. The divisibility condition holds: $a_1 \mid a_2 \mid \cdots \mid a_m$.

We can visualize the situation with the following commutative diagram:

$$\begin{array}{ccccc} K & \xhookrightarrow{\iota} & R^n & \xrightarrow{\pi} & M \cong R^n/\ker\pi = R^n/\mathrm{im}(\iota) = \mathrm{coker}(\iota) \\ \cong\uparrow & & =\uparrow & & \uparrow\cong \\ R^m & \xrightarrow[\psi]{} & R^n & \longrightarrow & \mathrm{coker}(\psi) \end{array}$$

Here, the map $\psi : R^m \to R^n$ is defined by mapping the canonical basis of R^m to the adapted basis elements of K. With respect to the basis $(e_1, \ldots, e_n)$, the map ψ is given by the diagonal matrix:

$$A := \begin{pmatrix} a_1 & 0 & \cdots & 0 \\ 0 & a_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & a_m \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & 0 \end{pmatrix} \in R^{n\times m}$$

From the diagram it is clear that to understand the structure of M we only have to understand the cokernel of the map ψ given by A. To do this, we decompose both R^n and $\mathrm{im}(\psi)$ into direct sums.

Since $(e_1, \ldots, e_n)$ is a basis, R^n is the direct sum of free modules generated by each basis vector:

$$R^n = \bigoplus_{i=1}^{n} Re_i = Re_1 \oplus Re_2 \oplus \cdots \oplus Re_n$$

The image of ψ is the submodule generated by the images of the basis vectors of the domain. By the definition of the matrix representation, the i-th column of A contains the coefficients of $\psi(\epsilon_i)$ (with ϵ_i being the i-th canonical basis vector) with respect to the basis $(e_1, \ldots, e_n)$.

Since A is diagonal with entries a_i, we have $\psi(\epsilon_i) = a_ie_i$. Consequently, the image submodule decomposes into a direct sum of the submodules generated by these elements:

$$\mathrm{im}(\psi) = \bigoplus_{i=1}^{n} R(a_ie_i) = Ra_1e_1 \oplus \cdots \oplus Ra_me_m \oplus 0 \oplus \cdots \oplus 0$$

We can conclude:

$$M \cong \mathrm{coker}(\psi) = \frac{\bigoplus_{i=1}^{n} Re_i}{\bigoplus_{i=1}^{n} Ra_ie_i}$$

By the universal property of the direct sum, the family of projection maps $\pi_i : Re_i \twoheadrightarrow Re_i/Ra_ie_i$ induces a unique surjective homomorphism:

$$\Phi : \bigoplus_{i=1}^{n} Re_i \longrightarrow \bigoplus_{i=1}^{n} \left(\frac{Re_i}{Ra_ie_i}\right)$$

The kernel of Φ is the direct sum of the kernels of the π_i, which is $\bigoplus_{i=1}^{n} Ra_ie_i = \mathrm{im}(\psi)$. Finally, we get:

$$M \cong R^n/\mathrm{im}(\psi) \cong (\bigoplus_{i=1}^{n} Re_i)/\ker\Phi \cong \mathrm{im}(\Phi) = \bigoplus_{i=1}^{n} \left(\frac{Re_i}{Ra_ie_i}\right)$$

We analyze the components Re_i/Ra_ie_i in two groups:

the torsion part ($1 \leq i \leq m$):

Consider the map

$$\varphi_i : Re_i \to R/(a_i),\ re_i \mapsto r + (a_i).$$

Since e_i is a basis element, the coefficient r is unique, making φ_i well-defined. The map is clearly surjective with kernel Ra_ie_i. Thus, by the first isomorphism theorem:

$$Re_i/Ra_ie_i \cong R/(a_i).$$

the free part ($m < i \leq n$)

In these coordinates, the image is zero ($a_i = 0$). We are quotienting by the zero submodule, which leaves the component unchanged:

$$Re_i/0 \cong Re_i \cong R.$$

Combining these parts, we obtain the decomposition:

$$M \cong R/(a_1) \oplus \cdots \oplus R/(a_m) \oplus R \oplus \cdots \oplus R$$

Remark: Finally, we remove any trivial terms where a_i is a unit (since $R/1 \cong 0$) to obtain the invariant factors. □

We introduce some tools for the uniqueness proof:

Definition 2.3.25. The **annihilator** of an R-module M is the set of scalars that kill every element in M:

$$\mathrm{Ann}(M) = \{r \in R \mid \forall x \in M, r \cdot x = 0\}.$$

It is easy to see that it is an ideal in R.

Lemma 2.3.26. Let $M \cong \bigoplus_{i=1}^{n} R/I_i$ be a decomposition with a descending chain of ideals $R \supsetneq I_1 \supseteq I_2 \supseteq \cdots \supseteq I_n$. Then the annihilator recovers the last ideal:

$$\mathrm{Ann}(M) = I_n$$

Proof. This is a special case of the lemma that follows. □

Definition 2.3.27. Let M and N be R-modules and let $k \geq 1$ be an integer.

Definition: A map $\phi : M^k \to N$ is called **R-multilinear** if it is R-linear in each variable separately. That is, for every index i, every $\lambda \in R$, and all $x_j, y \in M$:

$$\phi(x_1, \ldots, x_i + \lambda y, \ldots, x_k) = \phi(x_1, \ldots, x_i, \ldots, x_k) + \lambda\phi(x_1, \ldots, y, \ldots, x_k)$$

The map ϕ is called **alternating** if it vanishes whenever two arguments are equal:

$$\exists i \neq j \text{ such that } x_i = x_j \implies \phi(x_1, \ldots, x_k) = 0$$

Remark: The alternating property implies:

$$\phi(\ldots, x_i, \ldots, x_j, \ldots) = -\phi(\ldots, x_j, \ldots, x_i, \ldots)$$

Lemma 2.3.28. Let $M \cong \bigoplus_{i=1}^{n} R/\mathfrak{a}_i$

with the ideals now indexed in ascending order:

$$\mathfrak{a}_1 \subseteq \mathfrak{a}_2 \subseteq \cdots \subseteq \mathfrak{a}_n \subsetneq R$$

Let T_k be the ideal of scalars defined by the property that they annihilate all alternating k-linear forms on M:

$$T_k = \{\lambda \in R \mid \forall N, \forall \varphi : M^k \to N \ R\text{-multilinear and alternating}\ , \lambda \cdot \varphi = 0\}$$

Then, for every $1 \leq k \leq n$, we have the equality:

$$T_k = \mathfrak{a}_k$$

Proof. ($T_k \subseteq \mathfrak{a}_k$) We construct a specific form to test λ. Let $x \in M$. We denote the component of x in the j-th summand $R/\mathfrak{a}_j$ by $x^{(j)}$.

To define a determinant, we must view these components in a common ring. We choose the target $R/\mathfrak{a}_k$. For any $j \leq k$, we have $\mathfrak{a}_j \subseteq \mathfrak{a}_k$. This inclusion of ideals induces a natural map $\pi : R/\mathfrak{a}_j \to R/\mathfrak{a}_k$ defined by $r + \mathfrak{a}_j \mapsto r + \mathfrak{a}_k$. This is well-defined because any element in the smaller ideal $\mathfrak{a}_j$ is automatically in $\mathfrak{a}_k$ (i.e., the zero of the source maps to the zero of the target).

We define $\psi : M^k \to R/\mathfrak{a}_k$ by taking the first k components of the inputs, projecting them all into $R/\mathfrak{a}_k$, and computing the determinant:

$$\psi(x_1, \dots, x_k) = \det\left(\pi(x_i^{(j)})\right)_{1 \leq i,j \leq k}$$

We evaluate this on the generators of the first k summands. Let $e_j \in M$ be the element with the identity coset $\overline{1}$ in the j-th position and zero elsewhere. Since $\pi(\overline{1}) = \overline{1}$, the matrix becomes the identity:

$$\psi(e_1, \dots, e_k) = \overline{1} \in R/\mathfrak{a}_k$$

If $\lambda \in T_k$, then $\lambda \cdot \psi = 0$, so $\lambda \cdot \overline{1} = \overline{0}$ in $R/\mathfrak{a}_k$, implying $\lambda \in \mathfrak{a}_k$.

($\mathfrak{a}_k \subseteq T_k$) Let $\lambda \in \mathfrak{a}_k$. We show that λ annihilates any alternating k-linear map $\phi : M^k \to N$.

By multilinearity, it suffices to check the value of ϕ on tuples of the form $(x_1, \dots, x_k)$ where each x_m belongs to a specific summand $R/\mathfrak{a}_{j_m}$. Suppose two inputs, say x_1 and x_2, come from the same summand $R/\mathfrak{a}_j$. Since $R/\mathfrak{a}_j$ is cyclic, it is generated by a single element e. Thus, we can write $x_1 = r \cdot e$ and $x_2 = s \cdot e$ for some $r, s \in R$. Substituting this into the map:

$$\phi(x_1, x_2, \dots) = \phi(re, se, \dots) = rs \cdot \phi(e, e, \dots)$$

Since ϕ is alternating, $\phi(e, e, \dots) = 0$, so the entire term vanishes.

Conclusion: The map ϕ is non-zero only if the inputs $x_1, \dots, x_k$ come from k distinct summands. Let the indices of these distinct summands be $j_1 < j_2 < \cdots < j_k$. Since there are k distinct integers chosen from $\{1, \dots, n\}$, the largest index must satisfy $j_k \geq k$. Using the ascending chain of ideals:

$$\lambda \in \mathfrak{a}_k \subseteq \mathfrak{a}_{j_k}$$

The input x_k belongs to $R/\mathfrak{a}_{j_k}$, so it is annihilated by $\mathfrak{a}_{j_k}$ (and thus by λ).

$$\lambda \cdot \phi(\dots, x_k) = \phi(\dots, \lambda x_k) = \phi(\dots, 0) = 0$$

Thus, λ kills every non-zero term of the map, so $\lambda \in T_k$. □

Uniqueness proof of the structure theorem. Suppose we have two decompositions of a finitely generated R-module M:

$$M \cong R^s \oplus \bigoplus_{i=1}^{n} R/I_i \quad \text{and} \quad M \cong R^{s'} \oplus \bigoplus_{j=1}^{m} R/J_j$$

where the ideals form chains $R \supsetneq I_1 \supseteq I_2 \supseteq \cdots \supseteq I_n$ and $R \supsetneq J_1 \supseteq J_2 \supseteq \cdots \supseteq J_m$.

Uniqueness of the free part.

First, we separate the torsion and free parts. We have already concluded that the torsion submodule M_{tor} corresponds to the torsion part of the decomposition $\bigoplus_{i=1}^{n} R/I_i =: T$. Consider the quotient M/M_{tor}:

$$M/M_{\text{tor}} \cong (R^s \oplus T)/T \cong R^s$$

(It is easy to see that the quotient is isomorphic to the free part R^s by considering the canonical projection $\pi : R^s \oplus T \to R^s$; its kernel is precisely $T = M_{\text{tor}}$ because R is an integral domain.)

Analogously, we get $M/M_{\text{tor}} \cong R^{s'}$. Thus, M/M_{tor} is a free module over the commutative ring R. We have shown that in this case, s must equal s', so the free parts are equal. We now assume $s = 0$ and focus on the torsion part.

Uniqueness of the torsion part. If $s = 0$, M is a torsion module. Suppose we have two decompositions of M into cyclic factors:

$$M \cong \bigoplus_{i=1}^{n} R/\mathfrak{a}_i \quad \text{and} \quad M \cong \bigoplus_{j=1}^{m} R/\mathfrak{b}_j$$

To apply our Lemma, we order the ideals in ascending chains (renaming indices if necessary so that $\mathfrak{a}_1$ is the smallest ideal, i.e., the annihilator of M):

$$\mathfrak{a}_1 \subseteq \mathfrak{a}_2 \subseteq \cdots \subseteq \mathfrak{a}_n \subsetneq R$$

$$\mathfrak{b}_1 \subseteq \mathfrak{b}_2 \subseteq \cdots \subseteq \mathfrak{b}_m \subsetneq R$$

By the Lemma proved previously, the ideals in such a decomposition are entirely determined by the intrinsic properties of M. Specifically, the k-th ideal in the ascending chain is equal to $T_k(M)$, the ideal of scalars that annihilate all alternating k-linear forms on M. Since $T_k(M)$ is defined independently of the decomposition, we must have $\mathfrak{a}_k = \mathfrak{b}_k$ for all k where both are defined.

Suppose for contradiction that the lengths differ, for instance $m > n$. Consider the intrinsic ideal $T_{n+1}(M)$. Using the first decomposition ($M \cong \bigoplus_{i=1}^{n} R/\mathfrak{a}_i$), any alternating map with $n+1$ inputs must vanish, because it is impossible to choose $n+1$ distinct summands from a set of n. Thus, the annihilator is the whole ring:

$$T_{n+1}(M) = R$$

However, using the second decomposition, the Lemma implies:

$$T_{n+1}(M) = \mathfrak{b}_{n+1}$$

Since the summands in the decomposition are non-zero, the ideals are proper ($\mathfrak{b}_{n+1} \subsetneq R$). This leads to the contradiction $R = \mathfrak{b}_{n+1}$. Thus, we must have $n = m$.

We conclude that since $n = m$ and $\mathfrak{a}_k = \mathfrak{b}_k$ for all k, the invariant factors are unique. □

2.4 Divisibility and Factorisation in Integral Domains

In the following, let R be an integral domain.

Definition 2.4.1. 1. Let $a, b \in R$. One says that a **divides** b if there is a $q \in R$ such that $b = q \cdot a$, denoted $a \mid b$.

Remark 2.4.2. It is easy to see that $a \mid b \Leftrightarrow (b) \subset (a)$.

2. We say that $a, b \in R$ are **associated** if there is an invertible $c \in R^\times$ such that $b = c \cdot a$, denoted $a \sim_{\text{ass}} b$.

Remark 2.4.3. (a) $\sim_{\text{ass}}$ is an equivalence relation. In fact, $R^\times$ acts on R by multiplication. The orbits are the equivalence classes of $\sim_{\text{ass}}$.

(b) $a \sim_{\text{ass}} b \Leftrightarrow (a) = (b)$.

Proof. ($\Rightarrow$) is clear. For ($\Leftarrow$), let $(a) = (b)$. We can assume that both are nonzero, otherwise the proof is trivial. Let $b = qa$, $a = rb$ for $q, r \in R$, so $b = qrb \implies qrb - b = 0 \implies b(qr - 1) = 0 \implies qr = 1 \implies q, r \in R^\times$. □

Definition 2.4.4. Let $\pi \in R$, $\pi \neq 0$, $\pi \notin R^\times$.

1. π is called a **prime element** if $\forall a, b \in R : \pi \mid a \cdot b \implies \pi \mid a$ or $\pi \mid b$. $\mathcal{P}(R)$ denotes the set of prime elements.

2. π is called **irreducible** if $\forall a, b \in R : \pi = a \cdot b \implies a \in R^\times$ or $b \in R^\times$. $\mathrm{Irr}(R)$ denotes the set of irreducible elements.
(in other words, all divisors of π are associated to π).

Remark 2.4.5. Let $p \in R \setminus \{0\}$. p is a prime element if and only if (p) is a prime ideal.

Proof. ($\Rightarrow$) Assume p is a prime element. By definition, p is not a unit, so $(p) \subsetneq R$. Let $xy \in (p)$. Then $p \mid xy$. Since p is prime, $p \mid x$ or $p \mid y$, which implies $x \in (p)$ or $y \in (p)$. Thus, (p) is a prime ideal.

($\Leftarrow$) Assume (p) is a prime ideal. By definition of prime ideals, $(p) \subsetneq R$, so p is not a unit. Suppose $p \mid xy$. Then $xy \in (p)$. Since (p) is prime, $x \in (p)$ or $y \in (p)$, which implies $p \mid x$ or $p \mid y$. Thus, p is a prime element. □

Lemma 2.4.6. Let $\pi \in R$ be a prime element. Then, π is irreducible.

Proof. Let $a, b \in R$ such that $\pi = a \cdot b$. Since π is prime, assume without loss of generality that $\pi \mid a$, so $a = q\pi$, $q \in R$. Then,

$$\pi = a \cdot b = q\pi \cdot b = (qb)\pi \implies qb = 1 \implies b \in R^\times,$$

so π is irreducible. □

Example 2.4.7. In fields, there are no irreducible or prime elements.

2.4.1 Divisibility in Principal Ideal Domains

In the following, let R be a principal ideal domain.

Lemma 2.4.8. Let $(a_n)_{n\in\mathbb{N}}$ a sequence in R such that for all $n \geq 1$, $a_n \mid a_{n-1}$. Then, $\exists n_0 \in \mathbb{N} \, \forall n \geq n_0 : a_n \sim_{\text{ass}} a_{n_0}$.

Proof. We can reformulate the condition by

$$(a_0) \subseteq (a_1) \subseteq (a_2) \subseteq \dots.$$

Consider $I = \bigcup_{n\in\mathbb{N}}(a_n)$. It is easy to see that I is an ideal. Since R is a principal ideal domain, there must be $a \in R$ such that $I = (a)$. Since $a \in I = \bigcup_{n\in\mathbb{N}}(a_n)$, there must be an index $n_0 \in \mathbb{N}$ such that $a \in (a_{n_0})$ (which implies $(a) \subseteq (a_{n_0})$), so for all $k \geq n_0$:

$$(a_{n_0}) \subseteq (a_k) \subseteq \bigcup_{n\in\mathbb{N}} (a_n) = I = (a) \subseteq (a_{n_0}),$$

so $\forall k \geq n_0 : (a_k) = (a_{n_0})$, which is equivalent to $a_k \sim_{\text{ass}} a_{n_0}$. □

Corollary 2.4.9. Every non-zero, non-unit element $a \in R$ has at least one irreducible divisor.

Proof. We construct a sequence of divisors (a_n) with $a_0 = a$. As long as a_n is reducible, we write $a_n = a_{n+1}b_{n+1}$ with non-units a_{n+1}, b_{n+1}. This yields a sequence where $a_{n+1} \mid a_n$ for all n. By the lemma, there exists n_0 such that for all $n \geq n_0$, $a_n \sim_{\text{ass}} a_{n_0}$. In particular, $a_{n_0+1} \sim_{\text{ass}} a_{n_0}$, which implies $a_{n_0+1} = u \cdot a_{n_0}$ for a unit u. This means the factorisation step $a_{n_0} = a_{n_0+1}b_{n_0+1}$ did not split a_{n_0} into non-units, so the process must have terminated. Thus, the element a_{n_0} is irreducible. □

Corollary 2.4.10. Let $a \in R \setminus \{0\}$, $a \notin R^\times$. Then, there are irreducible $\pi_1, \ldots, \pi_n (n \geq 1)$ such that $a = \prod_{i=1}^{n} \pi_i$.

Proof. Let $a_0 = a$. We recursively construct a sequence of divisors. If a_k is a unit, the process terminates. If a_k is not a unit, by Corollary 1, there exists an irreducible π_{k+1} such that $a_k = \pi_{k+1} a_{k+1}$. This yields a sequence where $a_{k+1} \mid a_k$ for all k. By the previous lemma, there exists n_0 such that for all $k \geq n_0$, $a_{k+1} \sim_{\text{ass}} a_k$. The condition $a_{k+1} \sim_{\text{ass}} a_k$ implies $a_k = u \cdot a_{k+1}$ for a unit u. Substituting this into $a_k = \pi_{k+1} a_{k+1}$, we see that π_{k+1} must be a unit. Since π_{k+1} is irreducible (and thus a non-unit), the sequence cannot continue beyond index n_0. Therefore, the sequence terminates at some a_n which is a unit, yielding the factorisation $a = \pi_1 \cdots \pi_n \cdot u$. □

Remark 2.4.11. If R is an euclidean domain such as $\mathbb{Z}$ or $K[x]$ (where K is a field), the existence of factorisation can be proven more directly using the Euclidean norm, without relying on the abstract lemma.

Proof. Let R be a Euclidean domain equipped with a norm $N : R \setminus \{0\} \to \mathbb{N}$.

For $R = \mathbb{Z}$, let $N(a) = |a|$.

For $R = K[X]$, let $N(f) = \deg f$.

These norms satisfy the property: if $b \mid a$ and $\frac{a}{b}$ is not a unit, then $N(b) < N(a)$.

Let $a \in R$ be a non-zero, non-unit element. If a is reducible, we can write $a = a_1 b_1$ where a_1, b_1 are non-units. By the norm property, $N(a_1) < N(a)$. If a_1 is reducible, we write $a_1 = a_2 b_2$ with $N(a_2) < N(a_1)$. We repeat this process to obtain a sequence of divisors $a, a_1, a_2, \ldots$ with strictly decreasing norms:

$$N(a) > N(a_1) > N(a_2) > \ldots$$

Since $N(x) \in \mathbb{N}$, such a strictly decreasing sequence of natural numbers cannot be infinite. The process must terminate at some element a_k which is irreducible. □

Corollary 2.4.12. Let $a \in R \setminus \{0\}$, $a \notin R^\times$. Then, there are irreducible $\pi_1, \ldots, \pi_n (n \geq 1)$ such that $a = \prod_{i=1}^{n} \pi_i$.

Definition 2.4.13. Let R be a general integral domain, $a, b \in R \setminus \{0\}$. $d \in R$ is called the **greatest common divisor** (gcd) of a and b if

1. $d \mid a$ and $d \mid b$
2. $\forall x \in R : x \mid a$ and $x \mid b \implies x \mid d$

We denote d by $\gcd(a, b)$.

$m \in R$ is called the **least common multiple** (lcm) of a and b if

1. $a \mid m$ and $b \mid m$
2. $\forall x \in R : a \mid x$ and $b \mid x \implies m \mid x$

We denote m by $\operatorname{lcm}(a, b)$.

Remark 2.4.14.

1. Note that in general, a gcd may not exist.
2. If a gcd exists, it is unique up to multiplication by a unit of R.

   Proof. If d, d' are both gcd's of a and b, we have that $d \mid d'$ and $d' \mid d$, so $(a) = (b)$, which means that a and b are associated. □

Theorem 2.4.15. Let $a, b \in R \setminus \{0\}$. Then, $d = \gcd(a, b)$ exists and $(d) = (a, b)$.

Proof. Since R is a PID, we can write $(a,b) = (d)$ for some $d \in R$. It follows that $a \in (d)$, $b \in (d)$, so $d \mid a$, $d \mid b$, so d is a common divisor of a and b. Now, let d' be another common divisor of a and b. Then, $a \in (d')$ and $b \in (d')$, so

$$(a,b) \subseteq (d') \iff (d) \subseteq (d') \iff d' \mid d,$$

so d is the greatest common divisor of a and b. □

Because in particular $d \in (a,b)$, d must be a linear combination of a and b with coefficients in R. We obtain the following result:

Corollary 2.4.16 (Bézout formula)**.** For $a, b \in R \setminus \{0\}$ there are $u, v \in R$ such that $\gcd(a,b) = ua + vb$.

Remark 2.4.17. Let $a, b \in R$.

1. If $d \in R$ is a linear combination and common divisor of a and b, it must be the gcd.
   Indeed: If d' is a common divisor of a and b, it must divide any linear combination, so $d' \mid d$.

2. One says that a and b are **coprime** if $\gcd(a,b) = 1$. (This is equivalent to $\exists u, v \in R : ua + vb = 1$)

3. R_{ass} denotes the set of equivalence classes of the relation $\sim_{\text{ass}}$. The relation on R_{ass} defined by $\bar{a} \leq \bar{b} \iff a \mid b$ is a partial order.

   Proof. First, we verify that the relation is well-defined. Let $a' \in \bar{a}$ and $b' \in \bar{b}$. Then $a' = ua$ and $b' = vb$ for some units $u, v \in R^\times$.

   $$a \mid b \iff b = ka \iff v^{-1}b' = ku^{-1}a' \iff b' = (vku^{-1})a' \iff a' \mid b'.$$

   (a) *Reflexivity:* For any $a \in R$, we have $a = 1 \cdot a$, so $a \mid a$. Thus $\bar{a} \leq \bar{a}$.

   (b) *Antisymmetry:* Suppose $\bar{a} \leq \bar{b}$ and $\bar{b} \leq \bar{a}$. Then $a \mid b$ and $b \mid a$, so $a \sim_{\text{ass}} b \iff \bar{a} = \bar{b}$.

   (c) *Transitivity:* Suppose $\bar{a} \leq \bar{b}$ and $\bar{b} \leq \bar{c}$. Then $a \mid b$ and $b \mid c$, meaning $b = xa$ and $c = yb$ for some $x, y \in R$. Substituting, we get $c = y(xa) = (yx)a$, so $a \mid c$. Thus $\bar{a} \leq \bar{c}$.

   □

Theorem 2.4.18. Let R be a PID, $\pi \in R \setminus \{0\}, \pi \notin R^\times$. Then, the following conditions are equivalent:

1. π is a prime element
2. π is irreducible
3. $R/(\pi)$ is an integral domain
4. $R/(\pi)$ is a field

Proof. We have already shown 1. $\implies$ 2. and 4. $\implies$ 3. $\implies$ 1. is clear, so only 2. $\implies$ 4. is left to show.
Assume π is irreducible. Let $\overline{\alpha} \in (R/(\pi)) \setminus \{0\}$, which implies $\alpha \in R$ and $\alpha \notin (\pi)$. Consider the ideal generated by π and α:

$$(\pi) \subset (\pi, \alpha) \subseteq R$$

Since R is a PID, there exists $\delta \in R$ such that $(\delta) = (\pi, \alpha)$. Note that δ is a gcd of π and α. Since $(\pi) \subset (\delta)$, we have $\delta \mid \pi$. Because π is irreducible, δ must be either a unit or associated to π.

Since $\alpha \in (\delta)$ but $\alpha \notin (\pi)$, we have $(\delta) \neq (\pi)$, so δ cannot be associated to π. Therefore, δ is a unit, and $(\pi, \alpha) = (\delta) = R$. By Bézout's lemma, there exist $\mu, \nu \in R$ such that:

$$1 = \mu\pi + \nu\alpha$$

Projecting this equation into the quotient $R/(\pi)$, we get:

$$\overline{1} = \overline{\mu\pi} + \overline{\nu\alpha} = \overline{\nu\alpha}$$

Thus, every non-zero element $\overline{\alpha}$ has an inverse $\overline{\nu}$, so $R/(\pi)$ is a field. □

Remark 2.4.19. In particular, nonzero prime ideals are always maximal in PIDs.

2.4.2 Unique Factorisation Domains

Definition 2.4.20. An integral domain R is called a **factorial ring** or **unique factorisation domain** (UFD) if for all $a \in R \setminus \{0\}$ and $a \notin R^\times$:

$$\exists r \geq 1, \exists(\pi_1, \ldots, \pi_r) \in \mathbb{P}(R)^r \quad \text{such that} \quad a = \prod_{i=1}^{r} \pi_i$$

Remark 2.4.21. We have proven that in principal ideal domains, all nonzero elements can be factored into irreducible elements and that all irreducible elements are prime, so all PIDs are factorial rings.

Lemma 2.4.22. Let R be a factorial ring, $\pi \in R$. Then, $\pi \in \mathcal{P}(R) \iff \pi \in \mathrm{Irr}(R)$.

Proof. We have already shown ($\implies$). Let $\pi \in \mathrm{Irr}(R)$. Since R is a factorial ring, we can write

$$\pi = \pi_1 \ldots \pi_r, \forall i \in \{1, \ldots, r\} : \pi_i \in \mathcal{P}(R),$$

so in particular, for all i, $\pi_i \notin R^\times$. Since π is irreducible, r must be 1, so

$$\pi = \pi_1 \in \mathcal{P}(R)$$

□

Remark 2.4.23. We denote by $\mathcal{P}(R)_{\mathrm{ass}} := \{\bar{r} \in R_{\mathrm{ass}} \mid r \in \mathcal{P}(R)\}$ and by $(\pi_\alpha)_{\alpha \in \mathcal{P}(R)_{\mathrm{ass}}}$ **a set of representatives of prime elements** in R. We can canonically define such a set as e.g. the prime numbers in $\mathbb{Z}$ or the monic irreducible polynomials in $K[x]$, but for general rings, this requires the axiom of choice.

Theorem 2.4.24. Let R be a UFD and let $(\pi_\alpha)_\alpha$ be a set of representatives of prime elements. For all $a \in R \setminus \{0\}$, there exists a unique family $(n_\alpha)_\alpha$, with $n_\alpha \in \mathbb{N}$ and $n_\alpha = 0$ for almost all α, and a unique unit $u \in R^\times$ such that

$$a = u \cdot \prod_\alpha \pi_\alpha^{n_\alpha}.$$

Remark 2.4.25. We can reformulate the result as follows: the set of non-zero association classes $R_{\mathrm{ass}} \setminus \{0\}$ is a commutative monoid under multiplication. Specifically, $R_{\mathrm{ass}} \setminus \{0\}$ is a free commutative monoid on the set of prime association classes $P(R)_{\mathrm{ass}}$.
(Monoid $\implies$ Existence of factorisation, Free $\implies$ Uniqueness of factorisation)

(Note: The free monoid on a set of n elements is isomorphic to $\mathbb{N}^n$).

If $K = \mathrm{Frac}(R)$ is the fraction field of R, then the quotient group $K^\times/R^\times$ is the free abelian group on the set $P(R)_{\mathrm{ass}}$ via the inclusion:

$$P(R)_{\mathrm{ass}} \subset K^\times/R^\times, \quad \bar{\pi} \longmapsto (\bar{\pi}).$$

Lemma 2.4.26. Let R be an integral domain. Let $r \geq 1$, $s \geq 1$ and let $p_1, \ldots, p_r \in P(R)$ and $q_1, \ldots, q_s \in P(R)$ be prime elements.

If

$$p_1 \cdots p_r = q_1 \cdots q_s \quad \text{in } R,$$

then $r = s$, and there exists a permutation $\sigma \in S_r$ such that for all i, p_i and $q_{\sigma(i)}$ are associated.

Proof. We proceed by induction on r.
If $r = 1$, we have $p_1 = q_1 \ldots q_s$. Since p_1 is in particular irreducible, we can assume that $q_2, \ldots, q_s$ are invertible, so p_1 and q_1 are associated.
If $r \geq 2$, Consider the equality:

$$p_1 \cdots p_r = q_1 \cdots q_s.$$

Since p_r divides the left of the equation, we have $p_r \mid q_1 \cdots q_s$. By the definition of a prime element, p_r must divide at least one factor on the right. Thus, there exists an index $j \in \{1, \ldots, s\}$ such that $p_r \mid q_j$.

Since q_j is irreducible, p_r and q_j must be associated. Thus, $q_j = u \cdot p_r$ for some unit $u \in R^\times$.

Substituting this into the original equation and cancelling p_r (valid since R is an integral domain), we obtain:

$$p_1 \cdots p_{r-1} = u \cdot q_1 \cdots \widehat{q_j} \cdots q_s.$$

Absorbing the unit u into one of the remaining factors on the right (e.g., let $q_1' = uq_1$), we have a product of $r-1$ primes on the left equal to a product of $s-1$ primes on the right, so we can apply the inductive hypothesis, completing the proof. □

Proof of the theorem. Existence: Follows from the fact that R is a UFD, allowing any non-zero non-unit to be written as a finite product of irreducibles. We collect associated primes into the representatives π_α and combine the units into u.

Uniqueness: Suppose a has two such representations:

$$a = u \cdot \prod_\alpha \pi_\alpha^{n_\alpha} = v \cdot \prod_\alpha \pi_\alpha^{m_\alpha}.$$

Expanding these powers into linear products of primes, we apply the Lemma. The Lemma guarantees that the prime factors on the left and right are identical up to permutation and association.

Since $(\pi_\alpha)_\alpha$ is a set of distinct representatives (i.e., no two distinct π_α are associated), the association of factors implies strictly that $n_\alpha = m_\alpha$ for all α. Cancelling the prime powers from both sides yields $u = v$. □

Remark 2.4.27. 1. Let R be a UFD and let $(\pi_\alpha)_\alpha$ be a set of representatives of prime elements. Consider $a, b \in R \setminus \{0\}$ with factorisations:

$$a = u \cdot \prod_\alpha \pi_\alpha^{n_\alpha}, \quad b = v \cdot \prod_\alpha \pi_\alpha^{m_\alpha}, \quad (u, v \in R^\times).$$

It follows that:

(a) $a \mid b \iff \forall \alpha,\ n_\alpha \leq m_\alpha$.

(b) The greatest common divisor exists and is given (up to association) by:

$$\gcd(a, b) \sim \prod_\alpha \pi_\alpha^{\min(n_\alpha, m_\alpha)}.$$

(c) The least common multiple exists and is given by:

$$\operatorname{lcm}(a, b) \sim \prod_\alpha \pi_\alpha^{\max(n_\alpha, m_\alpha)}.$$

(d) From the above, we obtain the identity:

$$\gcd(a,b)\cdot \operatorname{lcm}(a,b) \sim a\cdot b.$$

(e) a and b are coprime if and only if for all α, $\min(n_\alpha, m_\alpha) = 0$.

Note: One may define the gcd of multiple elements $(a_1,\dots,a_n) \in (R\setminus\{0\})^n$ recursively:

$$\gcd(a_1,\dots,a_n) = \gcd(a_1, \gcd(a_2,\dots,a_n)).$$

2. *Gauss's Lemma I:* Let R be a factorial ring, $a,b,c \in R\setminus\{0\}$. If $a \mid bc$ and a,c are coprime, then $a \mid b$.

3. Let K be a field and let $P \in K[x]$. If $\deg P \le 3$, then:

$$P \text{ is irreducible } \iff R_P(K) = \emptyset,$$

Proof. ($\Rightarrow$) If $R_P(K) \neq \emptyset$, P has a root α and factorizes as $(X-\alpha)Q$, so it is reducible.

($\Leftarrow$) If P is reducible, $P = A\cdot B$ with $\deg A, \deg B \ge 1$. Since $\deg P \le 3$, necessarily one factor must have degree 1 (as $2+2=4>3$). A degree 1 factor implies the existence of a root. □

Consider $(X^2+1)^2$ in $\mathbb{R}[x]$. It has no roots in $\mathbb{R}$ ($R_P(\mathbb{R}) = \emptyset$) but is clearly reducible.

2.4.3 Polynomial Rings over UFDs

To conclude this chapter, we prove that the polynomial ring over a UFD is a UFD itself. We begin with the following

Observation: If A is an integral domain, $A_{\text{ass}} \cong \operatorname{Prin}(A)$ via the map $\bar{a} \mapsto (a)$, where $\operatorname{Prin}(A)$ denotes the set of principal ideals in A. Also note that the product of equivalence classes maps to the ideal product.

Proof. Let Φ be the map defined above. We construct the inverse map $\Psi : \operatorname{Prin}(A) \to A_{\text{ass}}$. Let $I \in \operatorname{Prin}(A)$ be a principal ideal. Choose a generator x such that $I = (x)$, and define $\Psi(I) = \bar{x}$.

Suppose I has two different generators, $I = (x) = (y)$. Then $x \mid y$ and $y \mid x$. Since A is a domain, this implies $x = uy$ for some unit $u \in A^\times$ (i.e., $x \sim_{\text{ass}} y$). Thus $\bar{x} = \bar{y}$ in A_{ass}, so the map is well-defined and is clearly the left and right inverse to Φ. □

Recall that For all $a,b \in A$:

$$(a) = (b) \iff a \sim_{\text{ass}} b$$
$$a \mid b \iff (b) \subseteq (a)$$

. Thus, we can easily see:

Lemma 2.4.28. Let $a,b \in R$, where R is a UFD. Then $\gcd(a,b) \in R_{\text{ass}}$. It is $\sup((a),(b)) \in (\operatorname{Prin}(A), \subseteq)$, where in a poset $(S,\le)$, for $a,b \in S$:

$$\sup(a,b) = \min\{c \in S \mid a \le c \text{ and } b \le c\}$$

Remark 2.4.29.

1. In general, $(a)+(b) \subsetneq (\gcd(a,b))$ (unless R is a PID).

2. For $(a_1,\dots,a_n) \in R^n$, we may define $\gcd(a_1,\dots,a_n) = \sup((a_1),\dots,(a_n))$ (if R is a UFD).

Definition 2.4.30. Let R be a UFD, $P \in R[X]\setminus\{0\}$. $P = \sum_{i=0}^n a_i X^i$, $a_n \neq 0$ is called **primitive** if $\gcd(a_1,\dots,a_n) = 1$.

Remark 2.4.31. This means that no prime element in R divides all coefficients of P.

Example 2.4.32. 1. Monic polynomials are primitive.

2. In $\mathbb{Z}[X]$, $2X+3$ is primitive.

Lemma 2.4.33 (Gauss' Lemma II)**.** Let R be a UFD and let $P, Q \in R[X]$ be primitive polynomials. Then their product $P \cdot Q$ is also primitive.

Proof. Suppose $P \cdot Q$ is not primitive. Then there exists an irreducible element $\pi \in \mathrm{Irr}(R)$ such that π divides all the coefficients of $P \cdot Q$. (This is equivalent to saying $\pi \mid (P \cdot Q)$ in $R[X]$.)

Consider the reduction homomorphism modulo π:

$$\Phi : R[X] \to (R/(\pi))[X], \quad A \mapsto \overline{A}$$

Recall that for any $A \in R[X]$, $\pi \mid A \iff \overline{A} = 0$ in $(R/(\pi))[X]$.

Our assumption $\pi \mid (P \cdot Q)$ implies:

$$\overline{P \cdot Q} = \overline{P} \cdot \overline{Q} = 0 \quad \text{in } (R/(\pi))[X]$$

Since π is irreducible in a UFD, the ideal (π) is prime. Therefore, the quotient ring $R/(\pi)$ is an integral domain, so the polynomial ring $R/(\pi)[X]$ is also an integral domain.

Since $\overline{P} \cdot \overline{Q} = 0$ in an integral domain, we must have:

$$\overline{P} = 0 \quad \vee \quad \overline{Q} = 0,$$

So either P or Q is not primitive. □

Remark 2.4.34. In a UFD R, generally, $R/(\pi)$, $\pi \in \mathcal{P}(R)$ is an integral domain but not a field.

Lemma 2.4.35. Let R be a UFD, $K = \mathrm{Frac}(R)$. Recall that we consider $R \subset K$. Let $P \in K[X] \setminus \{0\}$.

Then there exist $c_0 \in K \setminus \{0\}$ and $P_0 \in R[X]$ primitive, such that:

$$P = c_0 P_0$$

Furthermore, this decomposition is unique up to a unit in R. That is, if $P = c_1 P_1$ is another such decomposition, then there exists $u \in R^\times$ such that:

$$c_1 = u c_0 \quad \text{and} \quad P_1 = u^{-1} P_0$$

Proof. Existence. Write P using a common denominator:

$$P = \sum_{i=0}^{n} \frac{a_i}{b_i} X^i, \quad a_i \in R,\ b_i \in R \setminus \{0\}$$

$$= \frac{1}{\prod_j b_j} \left(\sum_{i=0}^{n} a_i' X^i \right), \quad \text{where } a_i' = \left(\prod_{j \neq i} b_j \right) a_i \in R$$

Let $d = \gcd(a_0', \ldots, a_n')$ inside R. Then we can write:

$$P = \frac{d}{\prod_j b_j} \cdot P_0, \quad \text{where } P_0 = \sum_{i=0}^{n} \frac{a_i'}{d} X^i \in R[X]$$

Clearly, $\gcd(\frac{a_0'}{d}, \ldots, \frac{a_n'}{d}) = 1$, so P_0 is primitive.

Uniqueness. Suppose $c_0 P_0 = c_1 P_1 = P$ is another decomposition. (In particular, $\frac{c_0}{c_1} P_0 = P_1$). Let (π_α) be a family of representatives of primes in R. Since R is a UFD, we can factor the fraction in K:

$$\frac{c_0}{c_1} = u \cdot \prod \pi_\alpha^{n_\alpha}, \quad n_\alpha \in \mathbb{Z},\ u \in R^\times$$

The equality $\frac{c_0}{c_1} P_0 = P_1$ becomes:

$$u \cdot \prod_{n_\alpha > 0} \pi_\alpha^{n_\alpha} \cdot P_0 = \prod_{n_\alpha < 0} \pi_\alpha^{-n_\alpha} \cdot P_1$$

If there is some $n_\alpha < 0$ (for some fixed α), then π_α divides the right hand side. By Gauss' Lemma I, since π_α is prime, it must divide the product on the left. However, π_α does not divide the primes with positive exponents (since $\pi_\alpha \not\sim_{\text{ass}} \pi_\beta$ for distinct indices), so π_α must divide P_0.

This implies π_α divides all coefficients of P_0, which is a contradiction because P_0 is primitive. Thus, for all α, we must have $n_\alpha = 0$ (and symmetrically $n_\alpha \not> 0$) $\implies \frac{c_0}{c_1} = u \in R^\times$. □

Remark 2.4.36.

1. Similarly, $\frac{c_1}{c_0} \in R^\times$, and $P_1 = \left(\frac{c_1}{c_0}\right) \cdot P_0$.
2. c_0, up to multiplication by a unit, is called the **content** of P. P_0 is called the **primitive part**.
3. Let $P \in K[X] \setminus \{0\}$, with $P = c_0 P_0$ where P_0 is primitive and $c_0 \in K \setminus \{0\}$. Of course, $P \in R[X] \iff c_0 \in R \setminus \{0\}$.
4. Gauss' Lemma II implies that the content of a product $P \cdot Q$ is the product of the contents.

   Proof. Write $P = c_0 P_0$ and $Q = d_0 Q_0$. Then $P \cdot Q = (c_0 d_0) \cdot (P_0 Q_0)$. Since P_0, Q_0 are primitive, their product $P_0 Q_0$ is primitive (by the lemma). Thus, the content of PQ is indeed $c_0 d_0$ (up to units). □

Theorem 2.4.37. Let R be a UFD. Then, the polynomial ring $R[X]$ is a UFD.

Proof. Let $K = \text{Frac}(R)$. We define two sets of elements in $R[X]$ and show they act as the prime elements.

1. Let $\mathcal{S}_1 := \mathcal{P}(R)$ be the set of prime elements of R. Let $\pi \in \mathcal{S}_1$. Then π is a prime element in $R[X]$.

   Indeed: If $\pi \mid PQ$ in $R[X]$, we must show $\pi \mid P$ or $\pi \mid Q$. Recalling the reduction map modulo π, we have:

   $$\pi \mid PQ \iff \overline{PQ} = \overline{P} \cdot \overline{Q} = \overline{0} \quad \text{in } (R/(\pi))[X]$$

   Since π is prime in R, the quotient $R/(\pi)$ is an integral domain. This implies $(R/(\pi))[X]$ is also an integral domain. Therefore, $\overline{P} \cdot \overline{Q} = \overline{0} \implies \overline{P} = \overline{0}$ or $\overline{Q} = \overline{0}$. This corresponds to $\pi \mid P$ or $\pi \mid Q$ in $R[X]$.

2. Let $\mathcal{S}_2 := \{P \in R[X] \mid P \text{ primitive and irreducible in } K[X]\}$. (Note that irreducibility in $K[X]$ implies $\deg(P) \geq 1$). Let $P \in \mathcal{S}_2$. Then P is a prime element in $R[X]$.

   Indeed: Let $P \mid QS$ where $Q, S \in R[X]$. Then clearly $P \mid QS$ in $K[X]$. Since K is a field, $K[X]$ is a Euclidean domain (hence a UFD), so irreducible elements are prime. Thus $P \in \mathcal{P}(K[X])$, so $P \mid Q$ or $P \mid S$ in $K[X]$. Assume without loss of generality that $P \mid Q$

in $K[X]$. Then $Q = PQ_1$ for some $Q_1 \in K[X]$. By the previous Lemma, we can write $Q_1 = cQ_2$ with $c \in K \setminus \{0\}$ and $Q_2 \in R[X]$ primitive. Substituting this back:

$$Q = P \cdot c \cdot Q_2 = c(PQ_2)$$

Since P and Q_2 are primitive, their product PQ_2 is primitive (by Gauss' Lemma). Comparing contents, we must have $c \in R$ (up to a unit). Therefore $Q_1 = cQ_2 \in R[X]$, which implies $P \mid Q$ in $R[X]$.

We can now show the existence of the factorisation. Let $P \in R[X] \setminus \{0\}$. We show P is a product of elements from $\mathcal{S}_1$ and $\mathcal{S}_2$. View P as an element of $K[X]$. Since $K[X]$ is a UFD, we have a factorisation into irreducibles in $K[X]$:

$$P = Q_1 \cdots Q_r, \quad \text{where } Q_i \in K[X] \text{ are irreducible.}$$

By the previous Lemma, for each i, we can write $Q_i = c_i\tilde{Q}_i$, where $c_i \in K$ and $\tilde{Q}_i \in R[X]$ is primitive. Note that $\tilde{Q}_i$ is associated to Q_i in $K[X]$, so it is still irreducible in $K[X]$. Thus $\tilde{Q}_i \in \mathcal{S}_2$. Now we have:

$$P = \left(\prod_{i=1}^{r} c_i\right) \tilde{Q}_1 \cdots \tilde{Q}_r$$

Let $C = \prod c_i$. Since $P \in R[X]$ and the product $\prod \tilde{Q}_i$ is primitive (Gauss' Lemma), the scalar C must actually be in $R \setminus \{0\}$. Since R is a UFD, we can factor C into primes of R:

$$C = u \cdot \pi_1 \cdots \pi_k, \quad \text{where } \pi_j \in \mathcal{S}_1$$

Thus, P is a product of primes from $\mathcal{S}_1$ and $\mathcal{S}_2$.

□

Remark 2.4.38. The prime elements in $R[X]$ are exactly the elements of $\mathcal{S}_1 \cup \mathcal{S}_2$ (up to multiplication by units in R), since we proved that any non-zero polynomial P admits a factorisation

$$P = u \cdot (\pi_1 \cdots \pi_k) \cdot (Q_1 \cdots Q_m)$$

where u is a unit, $\pi_i \in \mathcal{S}_1$, and $Q_j \in \mathcal{S}_2$. If P is itself an irreducible (prime) element, it cannot be decomposed into a product of two or more non-units. Therefore, exactly one factor in the list above must be non-unit, and all others must be units. This forces P to be associated either to a single π_1 (case $m = 0$) or to a single Q_1 (case $k = 0$).

Theorem 2.4.39 (Eisenstein Criterion)**.** Let R be a UFD and $P \in R[X]$ be given by $P = \sum_{i=0}^{n} a_i X^i$ with $n \geq 1$ and $a_n \neq 0$. If there is a prime $\pi \in \mathcal{P}(R)$ such that

1. $\pi \nmid a_n$.
2. For all $i \in \{0, \ldots, n-1\}$, $\pi \mid a_i$. (This is equivalent to saying $\overline{P} = \overline{a_n}X^n$ in $(R/(\pi))[X]$).
3. $\pi^2 \nmid a_0$.

Then P is irreducible in $\mathrm{Frac}(R)[X]$ (and irreducible in $R[X]$ if P is also primitive).

Proof. Suppose P is reducible in $K[X]$ (where $K = \mathrm{Frac}(R)$). That is, $P = QS$ with $Q, S \in K[X]$ and $\deg Q, \deg S \geq 1$.

Using the contents (or the previous Lemma), we can clear denominators to find polynomials in $R[X]$. Specifically, writing $Q = c(Q)Q_0$ and $S = c(S)S_0$ where Q_0, S_0 are primitive in $R[X]$, we get:

$$P = (c(Q)c(S))Q_0S_0 \implies P = \lambda Q_0 S_0$$

for some scalar $\lambda \in R$ (since $P \in R[X]$). We can absorb the scalar into one of the factors, say $\tilde{Q} = \lambda Q_0$. Thus we may assume $P = \tilde{Q}S_0$ in $R[X]$ with $\deg\tilde{Q}, \deg S_0 \geq 1$.

Now, consider the reduction modulo π:

$$\overline{a_n}X^n = \overline{P} = \overline{\tilde{Q}} \cdot \overline{S_0} \quad \in (R/(\pi))[X]$$

Since $\pi \nmid a_n$, we have $\overline{a_n} \neq 0$. Since π is prime, $R/(\pi)$ is an integral domain, so factorisation of monomials is unique. Therefore, the factors $\overline{\tilde{Q}}$ and $\overline{S_0}$ must be monomials of the form:

$$\overline{\tilde{Q}} = uX^m, \quad \overline{S_0} = vX^l$$

where $m + l = n$, and u, v are units in $R/(\pi)$.

Because $\deg(\tilde{Q}) + \deg(S_0) = n$, the degrees of the reductions must match the degrees of the original polynomials (i.e., the leading coefficients were not killed by π). If we assume P is reducible, then $m = \deg\tilde{Q} \geq 1$ and $l = \deg S_0 \geq 1$.

Consequently, the constant terms of $\overline{\tilde{Q}}$ and $\overline{S_0}$ must be 0 (since $m, l \geq 1$). This means π divides the constant coefficient of $\tilde{Q}$ and π divides the constant coefficient of S_0.

Let q_0 and s_0 be these constant coefficients. Since $a_0 = q_0 s_0$, and $\pi \mid q_0$ and $\pi \mid s_0$, it follows that:

$$\pi^2 \mid a_0$$

This contradicts assumption 3. Thus, P must be irreducible. □

3 Field Extensions

In the following, let K and L be fields.

3.1 Generalities

Definition 3.1.1. Let A be a commutative ring. An **A-algebra** is a pair (R, φ) consisting of a ring R and a ring homomorphism $\varphi : A \to R$ such that the image of A is contained in the center of R:

$$\varphi(A) \subseteq Z(R)$$

(i.e., $\forall a \in A, x \in R : \varphi(a) \cdot x = x \cdot \varphi(a)$).

A **homomorphism** of A-algebras between (R, φ) and (S, ψ) is a ring homomorphism $f : R \to S$ such that $f \circ \varphi = \psi$.

We may define the category of commutative A-algebras **A-cAlg**.

Remark 3.1.2.

1. The homomorphism φ defines a scalar multiplication $A \times R \to R$ given by $(a, x) \mapsto \varphi(a)x$. This gives R the structure of an A-module.
2. The condition $f \circ \varphi = \psi$ is equivalent to requiring that f is A-linear.

   Note that in the case where $A \subseteq R$ and $A \subseteq S$ are subrings, this is precisely the requirement that $f|_A = \mathrm{id}_A$.
3. Let A be a commutative K-algebra and let $u \in A$ be any element. We can formally justify evaluating a polynomial at an element $u \in A$ by considering the **evaluation map**

   $$\mathrm{ev}_u : K[X] \to A, \quad \sum_{i=0}^{n} a_i X^i \mapsto \sum_{i=0}^{n} a_i u^i.$$

   And writing $\mathrm{ev}_u(P) =: P(u)$. It is easy to see that ev_u is a morphism of K-algebas.

Proposition 3.1.3. Let K be a field and R a commutative K-algebra. The map

$$\Phi : \mathrm{Hom}_{K\text{-}\mathbf{cAlg}}(K[X], R) \to R, \quad \varphi \mapsto \varphi(X)$$

is a bijection.

Proof. We define $\Psi : R \to \mathrm{Hom}_{K\text{-}\mathbf{cAlg}}(K[X], R)$ by $z \mapsto \mathrm{ev}_z$.

First, for any $z \in R$, we have $(\Phi \circ \Psi)(z) = \Phi(\mathrm{ev}_z) = \mathrm{ev}_z(X) = z$. Thus $\Phi \circ \Psi = \mathrm{id}_R$.

Conversely, let $\varphi \in \mathrm{Hom}_{K\text{-}\mathbf{cAlg}}(K[X], R)$ and let $z = \varphi(X)$. Since φ is a K-algebra homomorphism, for any polynomial $P = \sum a_i X^i$ we have

$$\varphi(P) = \sum a_i \varphi(X)^i = \sum a_i z^i = \mathrm{ev}_z(P).$$

Thus $\varphi = \mathrm{ev}_z = \Psi(\Phi(\varphi))$, proving that $\Psi \circ \Phi = \mathrm{id}$. Since Φ has a two-sided inverse, it is a bijection. □

Definition 3.1.4. A **field extension** is a ring homomorphism $\iota : K \to L$.

Note that we have shown that ι must be injective since we are mapping out of a field. Thus, we usually consider K as a subfield of L and denote ι by $K \subset L$.

Clearly L is a K-algebra, so L is a K-vector space. $[L : K] =: \dim_K L$ is called the **degree** of the extension.

An extension $K \to L$ is called **finite** if $[L : K] < \infty$, and **infinite** otherwise.

Example 3.1.5.

1. Given a field M, if $K \subset L$ and $L \subset M$ are field extensions, then the composite inclusion $K \subset M$ is also a field extension.

2. We have the chain of inclusions $\mathbb{Q} \subset \mathbb{R} \subset \mathbb{C}$.

   The extension $\mathbb{C}/\mathbb{R}$ is finite. Since $\{1, i\}$ is a basis, we have $[\mathbb{C} : \mathbb{R}] = 2$.
   The extension $\mathbb{R}/\mathbb{Q}$ is infinite. This can be seen as $\{\sqrt{p} \mid p \in \mathcal{P}(\mathbb{Z})\}$ is an infinite linearly independent set over $\mathbb{Q}$.

3. Let K be a field. Consider the polynomial ring $K[X]$ and its field of fractions (rational functions) $K(X) = \mathrm{Frac}(K[X])$. By 1., the inclusion $K \subset K(X)$ defines a field extension.

   This extension is infinite, as the elements $\{1, X, X^2, \dots\}$ form a linearly independent set over K.

4. Let $P \in K[X]$ be a monic irreducible polynomial. (It follows that $\deg P \geq 1$.)

   The quotient ring $L = K[X]/(P)$ is a field (since (P) is a prime ideal and $K[X]$ is a PID).
   $\{1, \alpha, \dots, \alpha^{\deg(P)-1}\}$ is a basis of L over K, so:

$$[K[X]/(P) : K] = \deg(P)$$

Example 3.1.6. Let $P(X) = \sum_{i=0}^{n} a_i X^i$ be a polynomial in $K[X]$. Consider the quotient ring $L = K[X]/(P)$ and let $\alpha = \bar{X}$ be the equivalence class of X in L.

We calculate $P(\alpha)$ by substituting α into the expression for P:

$$P(\alpha) = \sum_{i=0}^{n} a_i(\alpha)^i = \sum_{i=0}^{n} a_i(\bar{X})^i = \sum_{i=0}^{n} a_i\overline{X^i} = \overline{\sum_{i=0}^{n} a_i X^i} = \overline{P(X)}.$$

Inside the quotient L, the element $\overline{P(X)}$ is 0 (since $P(X) \in (P)$).

Definition 3.1.7. Observe that for any ring R, there is a canonical morphism of rings

$$\psi_R : \mathbb{Z} \to R,\ n \mapsto \sum_{i=1}^{n} 1_R.$$

We define the **characteristic** of a field K, denoted $\mathrm{char}\, K$ using the map ψ_K as follows:

1. If ψ_K is injective, $\mathrm{char}\, K := 0$.

2. If ψ_K is not injective, its kernel is a proper ideal of $\mathbb{Z}$. Since $\mathbb{Z}$ is a PID, it is generated by a unique positive integer $p \in \mathbb{Z}$. We set $\mathrm{char}\, K := p$.

Remark 3.1.8.

1. If $\psi_K : \mathbb{Z} \to K$ is injective, by the universal property of the fraction field, ψ_K factors through $\mathrm{Frac}(\mathbb{Z}) = \mathbb{Q}$. Since $\mathbb{Q}$ is a field, we can consider $\mathbb{Q} \subseteq K$.

2. If ψ_K is not injective, we have $\mathbb{Z}/\ker\psi_K \cong \mathrm{im}\,\psi_K$. Since K is particularly an integral domain, the subring $\mathrm{im}\,\psi_K$ is an integral domain, so $ker\,\psi_K$ is a prime ideal, so its generator p is in $\mathcal{P}(\mathbb{Z})$. Since $p > 0$, p is a prime number.

Lemma 3.1.9. Let $\iota : K \to L$ be a field extension. Then, $\mathrm{char}\, K = \mathrm{char}\, L$.

Proof. Consider the canonical ring homomorphisms $\psi_K : \mathbb{Z} \to K$ and $\psi_L : \mathbb{Z} \to L$. Note that

$$\psi_L(n) = n \cdot 1_L = n \cdot \iota(1_K) = \iota(n \cdot 1_K) = \iota(\psi_K(n)).$$

Since K is a field, ι is injective. Thus:

$$n \in \ker(\psi_L) \iff \iota(\psi_K(n)) = 0_L \iff \psi_K(n) = 0_K \iff n \in \ker(\psi_K)$$

Therefore, $\ker(\psi_L) = \ker(\psi_K)$ and thus $\mathrm{char}\, K = \mathrm{char}\, L$. □

Example 3.1.10. 1. $\operatorname{char}\mathbb{Q} = \operatorname{char}\mathbb{R} = \operatorname{char}\mathbb{C} = \operatorname{char}\mathbb{C}[X] = 0$.

2. Consider the finite field $\mathbb{F}_p$ (which has p elements). Let $P \in \mathbb{F}_p[X]$ be a monic, irreducible polynomial. Then the quotient $L = \mathbb{F}_p[X]/(P)$ is a field extension of $\mathbb{F}_p$.

   We can conclude:
   $$\operatorname{char}(L) = \operatorname{char}(\mathbb{F}_p) = p$$
   Note: This extension L is a finite field with $p^{\deg(P)}$ elements, since L is an $\mathbb{F}_p$-vector space and thus isomorphic to $\mathbb{F}_p^{\deg P}$.

3.1.1 Transcendental and Algebraic Elements

Definition 3.1.11. Let A be a commutative K-algebra and let $x \in A$. Consider the evaluation homomorphism:
$$\mathrm{ev}_x : K[X] \to A, \quad P \mapsto P(x)$$

1. If ev_x is injective, x is called **transcendental** over K.

2. If ev_x is not injective, then $\ker(\mathrm{ev}_x) \neq \{0\}$. Since $K[X]$ is a PID, this kernel is generated by a unique monic polynomial $P_x \in K[X]$:
   $$\ker(\mathrm{ev}_x) = (P_x)$$
   In this case, x is called **algebraic** over K, and P_x is called the **minimal polynomial** of x over K.

Example 3.1.12. 1. If A is a field, it is now easy to see that P_x is irreducible in $K[X]$.

2. Consider $\mathbb{Q} \to \mathbb{C}$. Let $P = X^2 - 1 \in \mathbb{Q}[X]$ Then, $\sqrt{2} \in \mathbb{C}$ is a root of P, so $P \in \ker \mathrm{ev}_{\sqrt{2}}$. Thus, $\sqrt{2}$ is algebraic.

   More generally, if there is a monic $P \in \mathbb{Q}[X]$ with root $z \in \mathbb{C}$, z is algebraic over Q.

3. Let $K \to L$ be a finite extension. Then, all $x \in L$ are algebraic.
   Indeed: $\mathrm{ev}_x : K[X] \to L$ is particularly K-linear and $\dim_K(K[X]) = \infty > \dim_K(L)$, so ev_x cannot be injective.

Lemma 3.1.13. Let $x \in L$ and let $\varphi_x \in \mathrm{End}_K(L)$ be the multiplication map $y \mapsto xy$. For any polynomial $P \in K[X]$, we have $P(\varphi_x) = \varphi_{P(x)}$.

Proof. Since $\mathrm{End}_K(L)$ is a K-algebra, the expression $P(\varphi_x)$ is well-defined. Let $P(X) = \sum a_i X^i$. Using the linearity of the map assignment $z \mapsto \varphi_z$ and the fact that $(\varphi_x)^i = \varphi_{x^i}$:
$$P(\varphi_x) = \sum a_i \varphi_x^i = \sum a_i \varphi_{x^i} = \varphi_{\sum a_i x^i} = \varphi_{P(x)}$$
Consequence: $P(\varphi_x) = 0 \iff P(x) = 0$. Thus, the minimal polynomial of the element x is exactly the minimal polynomial of the linear operator φ_x. □

Definition 3.1.14. Let $K \to L$ be a field extension, $(x_i)_{i\in I}$ a family of elements in L.

The **sub-K-algebra** generated by $(x_i)_{i\in I}$, denoted $K[(x_i)_{i\in I}]$, is the smallest subring of L containing both K and all elements x_i.

Explicitly, it consists of all polynomial expressions in the x_i's with coefficients in K:
$$K[(x_i)_{i\in I}] = \{P(x_{i_1}, \ldots, x_{i_r}) \mid \{i_1, \ldots, i_r\} \subseteq I,\ P \in K[X_1, \ldots, X_r]\}$$

Note that $K[(x_i)_{i\in I}]$ is generally only an integral domain. We denote its field of fractions by $K((x_i)_{i\in I})$.

Clearly, there is an embedding $K[(x_i)_{i\in I}] \hookrightarrow L$, so the embedding factors through $K((x_i)_{i\in I})$. Thus, $K((x_i)_{i\in I}) \subseteq L$.

Example 3.1.15. Let $K \subset L$ be a field extension and let $x \in L$.

The K-algebra $K[x]$ is precisely the image of the evaluation homomorphism:

$$\mathrm{im}(ev_x) = \{P(x) \mid P \in K[X]\} = K[x] \subset L$$

If x is algebraic over K with minimal polynomial P_x, there is an isomorphism:

$$K[X]/(P_x) \xrightarrow{\cong} \mathrm{im}(\mathrm{ev}_x) = K[x]$$

Since P_x is irreducible, the quotient $K[X]/(P_x)$ is a field. Consequently, the subring $K[x]$ is a subfield of L. Moreover, $K \to K[X]$ is a finite extension since $K[X]/(P_x)$ is clearly finite-dimensional and isomorphic to $K[x]$.

3.2 Algebraic Extensions

Definition 3.2.1. Let $K \to L$ be a field extension. It is called **algebraic** if all $x \in L$ are algebraic over K.

Example 3.2.2.

1. Finite field extensions are algebraic.
2. $\mathbb{Q} \subset \mathbb{R}$ is not algebraic: for instance $e, \pi \in \mathbb{R}$ are transcendental over $\mathbb{Q}$. *(Lindemann-Weierstrass-Theorem)*

Lemma 3.2.3. Let $K \to L$ be a field extension, $x_1, \ldots, x_n \in L$.
If for all i, x_i is algebraic over K, the sub-K-algebra $K[x_1, \ldots, x_n] \subseteq L$ is a subfield and $\dim_K(K[x_1, \ldots, x_n]) \leq \infty$.

Proof. We proceed by induction on n. We have already shown the claim for $n = 1$.

If $n \geq 2$, let $F = K[x_1, \ldots, x_{n-1}]$. We can assume that F is a field and $[F : K] < \infty$.

Now consider $K[x_1, \ldots, x_n] = F[x_n]$. Since x_n is algebraic over K, there is a nonzero minimal polynomial $P_{x_n} \in K[X]$. Since $K \subset F$, we have $P_{x_n} \in F[X]$, so x_n is also algebraic over F.

By the base case logic (applied to the extension over F), $F[x_n]$ is a field and $[F[x_n] : F] < \infty$. Using the fact that for finite extensions $A \to B \to C$, $[C : A] = [C : B][B : A]$ *(proof: exercise)*,

$$[(K[x_1, \ldots, x_n] =)F[x_n] : K] = [F[x_n] : F] \cdot [F : K].$$

Since both factors on the right are finite, the total degree is finite. □

Remark 3.2.4. Let $K \to L$ be an extension, $x \in L$. The following are equivalent:

1. x is algebraic over K
2. $[K[x] : K] < \infty$
3. $K[x] = K(x)$
4. There is a field M with $K \subseteq M \subseteq L$ such that $x \in M$ and $\dim_K M \leq \infty$

In that case, we say that x is **of finite degree** over K.

Notation: $\deg_K x = \deg P_x = \dim_K K[x]$.

Proof. The previous lemma proves (1.) $\implies$ (2.) and

$$(1.) \implies (3.) \text{ (since } \mathrm{Frac}(A) = A \iff A \text{ is a field)}$$

as well as

$$(1.) \implies (4.) \text{ (we can pick } M = K[x])$$

so we only have to show that each point implies (1.).

$\neg(1) \implies \neg(2), \neg(3), \neg(4)$ Suppose x is transcendental over K. By definition, the evaluation map $ev_x : K[X] \to K[x]$ is injective. Since it is clearly surjective, we have an isomorphism of K-algebras:

$$K[x] \cong K[X]$$

We check the contrapositives:

$\neg(2.)$: The polynomial ring $K[X]$ has infinite dimension (basis $\{1, X, X^2, \dots\}$). Thus $[K[x] : K] = \infty$.

$\neg(3.)$: The polynomial ring $K[X]$ is not a field, so $K[x] \subsetneq \text{Frac } K[x] = K(x)$.

$\neg(4.)$: Suppose there was a field M with $x \in M$ and $[M : K] < \infty$. Since M is a subspace containing x and is closed under multiplication, it must contain the algebra generated by x, i.e., $K[x] \subseteq M$. This would imply:

$$\infty = \dim_K(K[x]) \leq \dim_K M < \infty$$

Thus, no such field M exists. □

Corollary 3.2.5. Let $K \to L$ be a field extension, with $x_1, \dots, x_n \in L$ algebraic over K. Given any polynomial $P \in K[X_1, \dots, X_n]$, $P(x_1, \dots, x_n) \in L$ is algebraic over K.

Proof. We have already shown that $K[x_1, \dots, x_n]$ is a finite extension and thus algebraic. Since $P(x_1, \dots x_n) \in K[x_1, \dots, x_n] \subseteq L$, $P(x_1, \dots, x_n)$ is algebraic over K. □

Remark 3.2.6. Particularly, if $x, y \in L$ are algebraic, so are $x + y, x - y, xy, \dots$.
However, there is no general formula to compute the minimal polynomials of these elements.

Corollary 3.2.7. Given a field extension $K \to L$, let $(x_i)_{i \in I}$ be a family of elements of L that are algebraic over K. Then, $K[(x_i)_{i \in I}]$ is a field.

Proof. Follows immediately from the fact that

$$K[(x_i)_{i \in I}] = \bigcup_{J \subseteq I, \text{ finite}} K[(x_j)_{j \in J}].$$

□

Corollary 3.2.8. Let $K \to L$, $L \to M$ be algebraic extensions. Then, the extension $K \to M$ is algebraic.

Proof. Let $x \in M$. Since M is algebraic over L, there exists a nonzero polynomial $Q \in L[X]$ such that $Q(x) = 0$. Let $Q = \sum_{i=0}^{n} a_i X^i$ with coefficients $a_i \in L$ and $n \geq 1$.

Consider the subfield generated by these coefficients:

$$L' := K[a_0, \dots, a_n] \subset L$$

As the elements a_i are algebraic over K (since $K \to L$ is algebraic), by the previous Lemma, $K[a_0, \dots, a_n]$ is a field that is finite-dimensional over K (and thus algebraic over K).

Now observe that $Q \in L'[X]$. Since $Q(x) = 0$, this implies that x is algebraic over L'. Consequently, the extension $L' \subset L'[x]$ is finite-dimensional.

We obtain the chain of fields:

$$\underbrace{K \subset L'}_{\text{finite}} \subset \underbrace{L'[x]}_{\text{finite}} \subset M$$

The total extension $L'[x]$ over K is finite-dimensional:

$$[L'[x] : K] = [L'[x] : L'] \cdot [L' : K] < \infty$$

Since x is contained in a finite extension $L'[x]$ of K, x is algebraic over K. Since x was an arbitrary element of M, the entire extension $K \subset M$ is algebraic. □

Definition 3.2.9. Let $K \to L$ be an extension. Consider the set

$$\bar{K}_L := \{x \in L \mid x \text{ is algebraic over } K\} \subseteq L.$$

$\bar{K}_L$ is a subfield of L and algebraic over K, called the **algebraic closure** of K in L.

A field Ω is called **algebraically closed** if every non-constant polynomial in $\Omega[X]$ admits a zero in Ω. It follows inductively that polynomials of degree $n \geq 1$ admit n zeroes.

Remark 3.2.10. If Ω is algebraically closed and $K \to \Omega$ is an extension, $\bar{K}_\Omega$ is algebraically closed. *(proof: exercises)*

Corollary 3.2.11. Given an extension $K \to L$, $\bar{K}_L$ is indeed a subfield and algebraic over K.

Proof. Let $x, y \in \bar{K}_L$. The extension $K[x, y]$ is algebraic over K. Since $K[x, y]$ contains $x - y$ and xy^{-1} (for $y \neq 0$), $\bar{K}_L$ is a subfield. The extension $K \to \bar{K}_L$ is algebraic by definition, as $\bar{K}_L$ consists precisely of elements algebraic over K. □

Corollary 3.2.12. Let $K \to L$ be an extension and $x_1, \ldots, x_n \in L$ algebraic elements over K. We denote $d_i := \deg_K x_i \geq 1$. Then, the following hold:

1. $[K[x_1, \ldots, x_n] : K] \leq \prod_i d_i$
2. $\forall i : d_i \mid [K[x_1, \ldots, x_n] : K]$ (particularly, $\operatorname{lcm}(d_1, \ldots, d_n) \mid [K[x_1, \ldots, x_n] : K]$)

Proof. Let $E = K[x_1, \ldots, x_n]$.

1. Define a chain of fields $K = E_0 \subset E_1 \subset \cdots \subset E_n = E$, where $E_i := K[x_1, \ldots, x_i]$.

   The element x_i is algebraic over K with degree $d_i = \deg(P_{x_i})$. Since $K \subseteq E_{i-1}$, the minimal polynomial of x_i over E_{i-1} must divide the original minimal polynomial P_{x_i} (viewed in $E_{i-1}[X]$). Therefore, the degree of the step is bounded:

   $$[E_i : E_{i-1}] \leq d_i$$

   The total degree is the product of the steps:

   $$[E : K] = \prod_{i=1}^{n} [E_i : E_{i-1}] \leq \prod_{i=1}^{n} d_i$$

2. Let $j \in \{1, \ldots, n\}$. We have $K \subseteq K[x_j] \subseteq E$, so

   $$[E : K] = [E : K[x_j]] \cdot [K[x_j] : K]$$

   We know that $[K[x_j] : K] = d_j$. Thus, d_j divides $[E : K]$ for every j.

□

Example 3.2.13. Consider the field extension $\mathbb{Q} \subset \mathbb{C}$ with elements $\alpha = \sqrt{2}$ and $\beta = \sqrt{3}$. We study the extension $\mathbb{Q}[\alpha, \beta]$.

We first determine the degree of the extension. Since $[\mathbb{Q}[\alpha] : \mathbb{Q}] = 2$ and $[\mathbb{Q}[\beta] : \mathbb{Q}] = 2$, the total degree satisfies:

$$[\mathbb{Q}[\alpha, \beta] : \mathbb{Q}] \leq [\mathbb{Q}[\alpha] : \mathbb{Q}] \cdot [\mathbb{Q}[\beta] : \mathbb{Q}] = 4.$$

Furthermore, since $\mathbb{Q}[\alpha] \subseteq \mathbb{Q}[\alpha, \beta]$, the degree must be a multiple of 2, leaving 2 or 4 as possibilities.

Assume for the sake of contradiction that $[\mathbb{Q}[\alpha, \beta] : \mathbb{Q}] = 2$. Then $\mathbb{Q}[\alpha] = \mathbb{Q}[\alpha, \beta]$, implying $\beta \in \mathbb{Q}[\alpha]$. Since $\{1, \alpha\}$ is a basis of $\mathbb{Q}[\alpha]$, there exist $u, v \in \mathbb{Q}$ such that $\beta = u + v\alpha$. Squaring both sides yields:

$$3 = u^2 + 4v^2 + 2uv\sqrt{2}.$$

Since $\sqrt{2} \notin \mathbb{Q}$, we must have $2uv = 0$.

1. If $v = 0$, then $\beta = u \in \mathbb{Q}$, contradicting $\beta \notin \mathbb{Q}$.

2. If $u = 0$, then $\beta = v\sqrt{2} \implies 3 = 2v^2 \implies v^2 = 3/2$, which is impossible for $v \in \mathbb{Q}$.

Thus, the assumption is false, and $[\mathbb{Q}[\alpha, \beta] : \mathbb{Q}] = 4$. Consequently, the set $\{1, \alpha, \beta, \alpha\beta\}$ forms a basis for the extension over $\mathbb{Q}$ *(proof: exercises)*.

Next, we show that $\gamma = \alpha + \beta$ generates the field. We check the degree $d = [\mathbb{Q}[\gamma] : \mathbb{Q}] \in \{1, 2, 4\}$.

1. If $d = 1$, then $\gamma \in \mathbb{Q}$. However, $\gamma^2 = 5 + 2\sqrt{6} \notin \mathbb{Q}$, a contradiction.

2. If $d = 2$, then γ would satisfy a quadratic equation $X^2 + uX + v = 0$ with coefficients $u, v \in \mathbb{Q}$. Substituting $\gamma = \alpha + \beta$:

$$\begin{aligned}(\alpha + \beta)^2 + u(\alpha + \beta) + v &= 0\\ (2 + 3 + 2\alpha\beta) + u\alpha + u\beta + v &= 0\\ (5 + v) \cdot 1 + u\alpha + u\beta + 2\alpha\beta &= 0.\end{aligned}$$

   This is a linear combination of the basis elements $\{1, \alpha, \beta, \alpha\beta\}$. Since these elements are linearly independent, all coefficients must be zero. However, the coefficient of $\alpha\beta$ is 2, which is non-zero. This is a contradiction.

Therefore, $[\mathbb{Q}[\gamma] : \mathbb{Q}] = 4$, so $\mathbb{Q}[\gamma] = \mathbb{Q}[\alpha, \beta]$.

Finally, to find the minimal polynomial of γ, we compute:

$$\begin{aligned}\gamma &= \sqrt{2} + \sqrt{3}\\ \gamma - \sqrt{2} &= \sqrt{3}\\ (\gamma - \sqrt{2})^2 &= 3\\ \gamma^2 - 2\sqrt{2}\gamma + 2 &= 3\\ \gamma^2 - 1 &= 2\sqrt{2}\gamma.\end{aligned}$$

Squaring both sides again eliminates the remaining root:

$$(\gamma^2 - 1)^2 = 8\gamma^2 \implies \gamma^4 - 2\gamma^2 + 1 = 8\gamma^2.$$

Rearranging terms yields the minimal polynomial $P = X^4 - 10X^2 + 1$.

3.2.1 The Category of Field Extensions

Remark 3.2.14. Let K be a field. We define the category of field extensions of K, denoted $\mathbb{F}_K$. The objects are field extensions L of K (formally, a pair (L, ι) of a field and an embedding). For two extensions L and M, the set of morphisms is defined as:

$$\mathrm{Hom}_{\mathbb{F}_K}(L, M) = \{f : L \to M \mid f \text{ is a field homomorphism s.t. } f|_K = \mathrm{Id}_K\}.$$

In other words, its objects are K-algebras that are fields, and its morphisms are all K-algebra homomorphisms.

Example 3.2.15. 0. If $f \in \mathrm{Hom}_K(L, M)$, then f preserves algebraicity. That is, for all $x \in L$ that are algebraic over K, the image $f(x) \in M$ is also algebraic over K. *(proof: exercise)*

1. $\mathrm{End}_{\mathbb{F}_{\mathbb{Q}}}(\mathbb{R}) = \{\mathrm{id}_{\mathbb{R}}\}$:
   Let $\sigma \in \mathrm{End}_{\mathbb{Q}}(\mathbb{R})$. Since σ is a field homomorphism fixing $\mathbb{Q}$, we have $\sigma(q) = q$ for all $q \in \mathbb{Q}$.

   First, we show that σ preserves non-negativity. Let $x \in \mathbb{R}$ with $x \geq 0$. Then there exists $y \in \mathbb{R}$ such that $x = y^2$. It follows that:

   $$\sigma(x) = \sigma(y^2) = (\sigma(y))^2 \geq 0.$$

   Consequently, σ is order-preserving. Let $a, b \in \mathbb{R}$ with $a \leq b$. Then $b - a \geq 0$, so:

   $$\sigma(b) - \sigma(a) = \sigma(b - a) \geq 0 \implies \sigma(a) \leq \sigma(b).$$

   For any rationals $q_1, q_2 \in \mathbb{Q}$ such that $q_1 \leq x \leq q_2$, the order-preserving property implies:

   $$\sigma(q_1) \leq \sigma(x) \leq \sigma(q_2).$$

   Since σ fixes $\mathbb{Q}$, this becomes:

   $$q_1 \leq \sigma(x) \leq q_2.$$

   As q_1 and q_2 can be chosen arbitrarily close to x, we conclude that $\sigma(x) = x$. Thus, $\sigma = \mathrm{id}_{\mathbb{R}}$.

2. $\mathrm{End}_{\mathbb{F}_{\mathbb{R}}}(\mathbb{C}) = \{\mathrm{id}_{\mathbb{C}}, \sigma\}$, with σ being the conjugation automorphism:
   Let $f \in \mathrm{End}_{\mathbb{R}}(\mathbb{C})$. By definition, $f(r) = r$ for all $r \in \mathbb{R}$. Since $i^2 = -1$, we have $f(i)^2 = f(i^2) = f(-1) = -1$. The only roots of $X^2 + 1$ in $\mathbb{C}$ are $\pm i$, so $f(i) \in \{i, -i\}$.

   For any $z = x + iy \in \mathbb{C}$ (with $x, y \in \mathbb{R}$):

   $$f(z) = f(x) + f(i)f(y) = x + f(i)y.$$

   If $f(i) = i$, then $f(z) = x + iy = z$, so $f = \mathrm{id}$.
   If $f(i) = -i$, then $f(z) = x - iy = \bar{z}$, so $f = \sigma$.

3. It can be shown that there are uncountably many morphisms in $\mathrm{End}_{\mathbb{Q}}(\mathbb{C})$

4. Let K be a field of characteristic 0. Recall that this means that $\mathbb{Q}$ can be embedded into K. Then, $\mathrm{End}_{\mathbf{Field}}(K) = \mathrm{End}_{\mathbb{F}_{\mathbb{Q}}}(K)$:

   Let $\sigma : K \to K$ be a field homomorphism. By definition, $\sigma(1_K) = 1_K$. Since σ respects addition, it fixes the subring generated by 1_K, which is isomorphic to $\mathbb{Z}$. Furthermore, since σ respects division, for any rational number $x = \frac{p}{q}$, we have $\sigma(x) = \sigma(p)\sigma(q)^{-1} = pq^{-1} = x$. Thus σ fixes the prime subfield $\mathbb{Q}$ pointwise.

The following lemma is fundamental to Galois theory.

Lemma 3.2.16. Let K be a field, let $P \in K[X]$ be a monic irreducible polynomial (so particularly, $\deg P \geq 1$).
Consider the quotient $K[X]/(P)$. Recall that it is a field extension of K of degree $\deg P$. Let $K \to M$ be another extension. Then, the map

$$\Psi : \mathrm{Hom}_{\mathbb{F}_K}(K[X]/(P), M) \to M,\ \varphi \mapsto \varphi(\alpha), \text{ where } \alpha = \bar{X} \in K[X]/(P)$$

is a bijection onto $R_P(M)$, the set of roots of P in M.

Proof. We can write $P = \sum_{i=0}^{n} a_i X^i$. First note that any K-extension morphism $\varphi : K[X]/(P) \to M$ maps $\alpha = \bar{X}$ to a root of P. Indeed:

$$P(\varphi(\alpha)) = \sum_{i=0}^{n} a_i \varphi(\alpha)^i = \varphi\left(\sum_{i=0}^{n} a_i \alpha^i\right) = \varphi(P(\alpha)) = \varphi(0) = 0.$$

Recall that there is a bijection $\Phi : \mathrm{Hom}_{K\text{-}\mathbf{cAlg}}(K[X], M) \to M$ given by $\phi \mapsto \phi(X)$. We visualize the restriction of this property to the quotient ring L using the following commutative diagram:

$$\begin{array}{ccc} \mathrm{Hom}_{K\text{-}\mathbf{cAlg}}(K[X], M) & \xrightarrow[\sim]{\Phi} & M \\ \uparrow & & \uparrow \\ \mathrm{Hom}_{\mathbb{F}_K}(K[X]/(P), M) & \overset{\Psi}{\dashrightarrow} & R_P(M) \end{array}$$

By the universal property of quotients, we can identify $\mathrm{Hom}_{\mathbb{F}_K}(K[X]/(P), M)$ with the subset of homomorphisms ϕ from $K[X]$ such that $\phi(P) = 0$. For any $\phi \in \mathrm{Hom}_{K\text{-}\mathbf{cAlg}}(K[X], M)$ and its corresponding element $z = \Phi(\phi)$, we have:

$$\phi(P) = 0 \iff P(z) = 0.$$

This shows that Φ restricts to a bijection between the subset of maps vanishing on (P) and the subset of roots $R_P(M)$. Therefore, the dashed map Ψ is an isomorphism. □

Remark 3.2.17.

1. It follows directly that given a root $\beta \in R_P(M)$, there is a unique K-extension morphism $\varphi : K[X]/(P) \to M$ with $\varphi(\bar{X}) = \beta$.

2. Let $a \in K \setminus K^2$ (meaning that a is not a second power in K). Then, $X^2 - a \in \mathrm{Irr}(K[X])$. Consider the extension $K \to K[X]/(X^2 - a) =: L$, sometimes written as $K[\sqrt{a}]$.
   If $K \to M$ is another field extension, there is a bijection

   $$\mathrm{Hom}_{\mathbb{F}_K}(K[X]/(X^2 - a),\, M) \cong \{x \in M \mid x^2 = a\}.$$

   For instance, if $\mathrm{char}(K) \neq 2$, $\forall x \in L^\times, x \neq -x$. $\alpha = \bar{X} \in L$, so $X^2 - a = (X + \alpha)(X - \alpha)$. Since morphisms in $\mathrm{End}_{\mathbb{F}_K}(L)$ are K-vector space endomorphisms, they are bijective, so

   $$\{\alpha, -\alpha\} \cong \mathrm{End}_{\mathbb{F}_K}(L) = \mathrm{Aut}_{\mathbb{F}_K}(L).$$

Lemma 3.2.18. Let $K \to L$ be a finite (and thus algebraic) extension. Then, there is an $r \geq 0$ such that there are $x_1, \ldots, x_r \in L$ with $L = K[x_1, \ldots, x_r]$.

Proof. Strong induction on the degree $[L : K]$.
If $K = L$, we are finished. If not, let $x \in L \setminus K$. Then, $[K[x] : K] > 1$, so the degree $[L : K[x]]$ is strictly smaller than $[L : K]$ and we are finished. □

Theorem 3.2.19. Let L, Ω be in $\mathbb{F}_K$, with $K \to L$ being finite and Ω being algebraically closed. Then,

$$\mathrm{Hom}_{\mathbb{F}_K}(L, \Omega) \neq \varnothing.$$

Proof. Because of the previous lemma, we can assume $L = K[x_1, \ldots, x_r]$, $x_1, \ldots, x_r \in L$, allowing us to proceed by induction on r. If $r = 1$, we can write

$$L = K[x] \cong K[X]/(P_x),\ x \in L.$$

Since Ω is algebraically closed, there is a $y \in R_{P_x}(\Omega)$, and because P_x is monic and irreducible, $P_y = P_x$, so we get a morphism

$$K[X]/(P_x) = K[X]/(P_y) \to \Omega.$$

If $n \geq 2$, let $K_1 = K[x_1]$ and consider the extensions $K \subset K_1 \subset L$. By the base case, there exists an embedding $\sigma : K_1 \to \Omega$. We identify K_1 with its isomorphic image $\sigma(K_1) \subset \Omega$. We observe that $L = K_1[x_2, \ldots, x_r]$ is an extension of K_1 generated by $r-1$ elements. Since Ω remains algebraically closed over the subfield $\sigma(K_1)$, we apply the induction hypothesis to extend σ to a homomorphism $\tau : L \to \Omega$. Thus, $\mathrm{Hom}_K(L, \Omega) \neq \varnothing$. □

Remark 3.2.20. It is possible to show that this holds for infinite algebraic extensions $K \to L$ as well using Zorns lemma.

3.3 Splitting Fields

Definition 3.3.1. Let $P \in K[X]$ be a monic polynomial. A field $L \supseteq K$ is called a **splitting field** of P over K if:

1. P splits into degree 1 (linear) factors in $L[X]$: $P = \prod_{i=1}^{\deg P}(X - \alpha_i)$
2. L is generated over K by the roots of P: $L = K[\alpha_i, \ldots, \alpha_{\deg P}]$

Remark 3.3.2. The extension to a splitting field $K \to L$ is finite. If $\deg(P) = n$, we will prove that $[L = K[x_1, \ldots, x_n] : K] \leq n!$.

Example 3.3.3.
1. $\mathbb{C}$ is a splitting field of $X^2 + 1 \in \mathbb{R}[X]$.
   More generally, if $a \in K \setminus K^2$, $K[X]/(X^2 - a)$ is a splitting field of $X^2 - a$.
2. $\mathbb{Q}[\sqrt[3]{2}]$ is *not* a splitting field of $X^3 - 2 \in \mathbb{Q}[X]$, since there are complex roots.

Lemma 3.3.4 (Extension Lemma)**.** Let K and K' be fields and let $\varphi : K \to K'$ be a field isomorphism. Let $P \in K[X]$ be non-constant with splitting field L over K, and let L' be a splitting field of P^φ over K' (where P^φ is obtained by applying φ to the coefficients of P). Then, φ can be extended to a field isomorphism $\sigma : L \xrightarrow{\cong} L'$.

Proof. Induction on $[L : K]$.

If $[L : K] = 1$, then $L = K$, so P splits into linear factors in $K[X]$:

$$P(X) = c\prod_{i=1}^{n}(X - a_i), \quad c, a_i \in K$$

Applying the coefficient-wise ring isomorphism $\tilde{\varphi} : K[X] \to K'[X]$ yields:

$$P^\varphi(X) = \varphi(c)\prod_{i=1}^{n}(X - \varphi(a_i))$$

Thus, P^φ splits into linear factors in $K'[X]$. Since L' is the splitting field of P^φ over K', and all roots $\varphi(a_i)$ are already in K', we have $L' = K'$, so φ is already an isomorphism of the splitting fields.

Assume $[L : K] = n > 1$ and the theorem holds for extensions of degree $< n$. Since $[L : K] > 1$, P has an irreducible factor $f \in K[X]$ with $\deg(f) \geq 2$.

Let $\alpha \in L$ be a root of f. Since $f \mid P \implies f^\varphi \mid P^\varphi$, and L' splits P^φ, there exists a root $\beta \in L'$ of f^φ.

The map φ induces an isomorphism of polynomial rings $\tilde{\varphi} : K[X] \to K'[X]$, mapping (f) to (f^{φ}). This descends to the quotient fields:

$$\bar{\varphi} : K[X]/(f) \xrightarrow{\cong} K'[X]/(f^{\varphi})$$

|Since $K[X]/(f) \cong K(\alpha)$ and $K'[X]/(f^{\varphi}) \cong K'(\beta)$, we obtain:

$$\hat{\varphi} : K(\alpha) \xrightarrow{\cong} K'(\beta)$$

such that $\hat{\varphi}|_K = \varphi$ and $\hat{\varphi}(\alpha) = \beta$.

We now treat L as a splitting field of P over $K(\alpha)$ and L' as a splitting field of P^{φ} over $K'(\beta)$. Since $[L : K(\alpha)] < n$, the induction hypothesis ensures an isomorphism $\sigma : L \to L'$ extending $\hat{\varphi}$, and thus extending φ. □

Theorem 3.3.5. Let $P \in K[X]$ be non-constant and monic. Then, splitting fields of P over K always exist and are unique up to K-isomorphism.

Proof. Existence:
Induction on $n := \deg P$.
If $n = 1$, P is already of the form $X - \alpha \in K[X]$, so there is nothing to show.
If $n \geq 2$, we can pick an irreducible factor $Q \in K[X]$ of P and consider the extension field

$$E := K[X]/(Q),\ \alpha_n := \bar{X} \in E.$$

Since $Q(\alpha_n) = 0$ and $Q \mid P$, $P(\alpha_n) = 0$, so we can factor

$$P = R \cdot (X - \alpha_n),$$

with R monic and of degree $n - 1$.

By induction hypothesis, there is a splitting field L of R, so

$$R = \prod_{i=1}^{n-1} (X - \alpha_i) \in L[X],$$

where $L = E[\alpha_1, \ldots, \alpha_{n-1}]$. We conclude that

$$P = \prod_{i=1}^{n} (X - \alpha_i) \in L[X] = E[\alpha_1, \ldots, \alpha_{n-1}] = K[\alpha_1, \ldots, \alpha_n].$$

Uniqueness up to K-isomorphism:
Note that the claim is simply the special case of the Extension Lemma where $K = K'$ and the base isomorphism φ is the identity map id_K. □

Remark 3.3.6. Let L be a splitting field of $P \in K[X]$ over K, $n := \deg P \geq 2$. Let

$$A = \{\alpha_1, \ldots, \alpha_n\} \subset L,\ P = \prod_{i=1}^{n} (X - \alpha_i) \in L[X].$$

Note that the α_i may repeat, so $m := |A| \leq n$.
We may assume that (after permutation), $A = \{\alpha_1, \ldots, \alpha_m\}$.
Consider $\sigma \in \mathrm{Hom}_{\mathbb{F}_K}(L, L) = \mathrm{Aut}_K(L)$ (because the degree of the splitting field is finite).
Note that $\sigma(R_P(L) \subseteq R_P(L))$, so the restriction

$$\sigma|_{R_P(L)} : (R_P(L) =) A \to A$$

is an isomorphism because $\sigma^{-1}|_{R_P(L)}$ is the inverse. Thus, we obtain a map

$$\Phi : \mathrm{Aut}_K(L) \to S_m,\ \sigma \mapsto \sigma|_{R_P(L)}$$

.

Lemma 3.3.7. The above defined map Φ is a monomorphism of groups.

Proof. Let $\sigma, \tau \in \mathrm{Aut}_K(L)$. For any root $\alpha \in A$, we have:

$$\Phi(\sigma \circ \tau)(\alpha) = (\sigma \circ \tau)|_{R_P(K)}(\alpha) = \sigma(\tau(\alpha)) = \Phi(\sigma)(\Phi(\tau)(\alpha)).$$

It follows that $\Phi(\sigma \circ \tau) = \Phi(\sigma) \circ \Phi(\tau)$, so Φ is indeed a group homomorphism.

Let $\sigma \in \ker(\Phi)$. By definition, this means $\Phi(\sigma) = \mathrm{id}_A$, so $\sigma(\alpha) = \alpha$ for all $\alpha \in A$. Since L is a splitting field of P over K, it is generated by the roots, i.e., $L = K[A]$. Since σ fixes K and the generating set A, it must fix the entire generated K-algebra L, so $\ker(\Phi) = *$. □

Corollary 3.3.8. Given a splitting field L of a polynomial $P \in K[X]$, we can conclude that there is a subgroup of S_m isomorphic to $\mathrm{Aut}_K(L)$, so

$$|\mathrm{Aut}_K(L)| \mid |S_m| = m! \mid n!,$$

where $n = \deg P \geq [L : K]$.

3.3.1 Application to Finite Fields

In the following, K denotes a finite field of characteristic $p > 0$ and cardinality q.

Lemma 3.3.9. K is a splitting field of $P = X^q - X \in \mathbb{F}_p$ over $\mathbb{F}_p$.

Proof. We show first that P splits into linear factors in K: Consider the multiplicative group $K^\times = K \setminus \{0\}$. Its order is $q - 1$, so by Lagrange's theorem, for any $x \in K \setminus \{0\}$,

$$x^{q-1} = x^{|K^\times|} = e_{K^\times} = 1 \implies x^q - x = 0.$$

Thus, $K = R_P(K)$, so

$$\forall x \in K : (X - x) \mid P \implies K[X] \ni Q := \prod_{x \in K} X - x \mid P,$$

but $\deg Q = |K| = q = \deg P$ and both Q and P are monic, so we can conclude $P = Q$.

Since clearly $K = \mathbb{F}_p[K]$, K is the splitting field of P over $\mathbb{F}_p$. □

Remark 3.3.10. Since we have shown that splitting fields are unique up to isomorphism, this shows that finite fields with given cardinality q are unique up to isomorphism.

Lemma 3.3.11. $|K| = p^{[K:\mathbb{F}_p]}$, which is particularly a prime power.

Proof. Since $\mathrm{char}\, K = p$, by definition, there is a canonical map $\psi_K : \mathbb{Z} \to K$ with kernel (p), so

$$\mathrm{im}\, \psi_k \cong \mathbb{Z}/(p) = \mathbb{F}_p.$$

Thus, there is the extension $\mathbb{F}_p \to K$, allowing us to treat K as an $\mathbb{F}_p$-vector space. We conclude:

$$|K| = |\mathbb{F}_p^{\dim_{\mathbb{F}_p}(K)}| = |\mathbb{F}_p|^{[K:\mathbb{F}_p]} = p^{[K:\mathbb{F}_p]}.$$

□

Definition 3.3.12. Let L be a field. A polynomial $P \in L[X]$ is called **separable** if it splits into distinct linear factors in a splitting field M of P over L. In other words, $|R_P(M)| = \deg P$.

Lemma 3.3.13. Let F be a field, $P \in F[X]$. If P and its formal derivative P' are coprime in $F[X]$, P is separable.

Proof. We prove the contrapositive. Suppose P has a multiple root α in some extension field. Then we can write $P = (X-\alpha)^2 Q$ for some polynomial Q. Using the formal product rule for derivatives:

$$P' = ((X-\alpha)^2)' \cdot Q + (X-\alpha)^2 \cdot Q' = 2(X-\alpha)Q + (X-\alpha)^2 Q'.$$

Notice that $(X-\alpha)$ divides both terms on the right-hand side. Therefore, $(X-\alpha) \mid P'$, so $(X-\alpha)$ is a common factor of P and P', contradicting the assumption. □

Remark 3.3.14. It follows that all irreducible polynomials P with coefficients in a field of characteristic 0 are separable since the derivative P' has strictly smaller degree and thus cannot be a non-trivial divisor. Note that characteristic 0 guarantees that the derivative of non-constant polynomials are not the zero polynomial.

Lemma 3.3.15. Let $F : K \to K$, $x \to x^p$, the Frobenius map. It is an automorphism of K.

Proof. Clearly, F is a homomorphism of the multiplicative group.
Let $x, y \in K$. Then, by the binomial theorem,

$$F(x+y) = (x+y)^p = \sum_{i=0}^{p} \binom{p}{i} x^i y^{p-i}.$$

It is easy to see that for $i \in \{1, \ldots, p-1\}$, $p \mid \binom{p}{i}$, so all summands but the first and last vanish and the term becomes $x^p + y^p = F(x) + F(y)$. (Note that this holds in any commutative $\mathbb{F}_p$-algebra.)

Thus, the map is automatically injective and since K is finite, it is bijective. □

Proposition 3.3.16. For every prime p and integer $n \geq 1$, there is a field of cardinality $q = p^n$.

Proof. Consider the polynomial $P = X^q - X \in \mathbb{F}_p[X]$ and let L be a splitting field of P over $\mathbb{F}_p$. We define the set of roots of P in L as

$$R := \{\alpha \in L \mid \alpha^q = \alpha\}.$$

We first show that R is a subfield of L. It is evident that $0, 1 \in R$ and that R is closed under multiplication and inversion. Regarding addition, we recall that since L has characteristic p, the map $\varphi : L \to L, x \mapsto x^{p^n}$ is the n-th iterate of the Frobenius map. Consequently, φ is a field homomorphism, so

$$(\alpha+\beta)^q = \alpha^q + \beta^q = \alpha + \beta$$

for all $\alpha, \beta \in R$. Thus, R is closed under addition and forms a subfield of L.

Consider the formal derivative $P' = qX^{q-1} - 1$. Since $q = p^n = 0$ in $\mathbb{F}_p$, we obtain $P' = -1$. It follows that $\gcd(P, P') = 1$, which implies that P is separable. Since P splits completely in L, it has exactly $\deg P = q$ distinct roots. Therefore, $|R| = q$, proving the existence of a field of cardinality q. □

Remark 3.3.17. We have thus established a complete classification: finite fields exist if and only if their cardinality is a prime power $q = p^n$, in which case the field is unique up to isomorphism.

3.4 Galois Extensions

Let K, L be arbitrary fields again.

Definition 3.4.1. An algebraically closed field Ω is called an **algebraic closure** of K if there is an algebraic extension $K \to \Omega$.

Theorem 3.4.2. Algebraic closures of K always exist and are unique up to K-isomorphism.

Proof. Existence:
Let F be any field. Let $\mathcal{S}$ be the set of all monic irreducible polynomials in $F[x]$ of degree ≥ 1. We define the polynomial ring in infinitely many variables:

$$R = F[X_f]_{f \in \mathcal{S}}$$

where we associate a distinct variable X_f to each $f \in \mathcal{S}$. Let I be the ideal of R generated by the polynomials $f(X_f)$ for all $f \in \mathcal{S}$.

We claim $I \neq R$. If $1 \in I$, there would exist a finite sum $1 = \sum_{i=1}^{n} g_i \cdot f_i(X_{f_i})$. However, we can choose a finite extension of F containing roots α_i for this finite set of polynomials f_i (for instance, a splitting field of $\prod_{i=1}^{n} f_i$). Evaluating the equation at $X_{f_i} = \alpha_i$ yields $1 = 0$, a contradiction.

By Zorn's Lemma, I is contained in a maximal ideal $\mathfrak{m}$. The residue field $F_1 = R/\mathfrak{m}$ is an extension of F in which every irreducible polynomial $f \in F[x]$ has at least one root (namely, the equivalence class of X_f).

We iterate this construction. Let $K_0 = K$. For each $n \geq 0$, let K_{n+1} be the field extension of K_n obtained by applying the above construction to K_n. This yields a chain of fields:

$$K_0 \subset K_1 \subset K_2 \subset \dots$$

Define $\Omega = \bigcup_{n=0}^{\infty} K_n$.

We show Ω is algebraically closed. Let $P \in \Omega[x]$ be a non-constant polynomial. Since P has finitely many coefficients, there exists some N such that $P \in K_N[x]$. By construction, P has a root in $K_{N+1} \subset \Omega$. Since Ω contains a root for every polynomial, it splits completely; hence Ω is algebraically closed.

Finally, since each extension K_{n+1}/K_n is algebraic (generated by roots of polynomials), the union Ω is algebraic over K. Thus, Ω is an algebraic closure of K.

Uniqueness up to K-isomorphism:
Let $\bar{K}$ and $\bar{K}'$ be two algebraic closures of K. We want a K-isomorphism $\sigma : \bar{K} \to \bar{K}'$.

Consider the set

$$\mathcal{P} := \{(E, \tau) \mid K \subseteq E \subseteq \bar{K},\ \tau \in \mathrm{Hom}_{\mathbb{F}_K}(E, \bar{K}')\}.$$

With the partial ordering

$$(E_1, \tau_1) \leq (E_2, \tau_2) \iff E_1 \subseteq E_2 \text{ and } \tau_2|_{E_1} = \tau_1.$$

Given a chain $(E_i, \tau_i)_{i \in I}$, it is easy to see that $(\bigcup_{i \in I} E_i, \bigcup_{i \in I} \tau_i)$, where we treat the τ_i as set-theoretic maps, i.e. subsets of $E_i \times K$, is an upper bound.

Zorn's Lemma implies the existence of a maximal pair (M, μ). Suppose for contradiction that $M \subsetneq \bar{K}$ and choose $\alpha \in \bar{K} \setminus M$. Let $f \in M[X]$ be the minimal polynomial of α over M. Since $\bar{K}'$ is algebraically closed, the polynomial f^{μ} has a root $\beta \in \bar{K}'$. [1] Applying the Extension Lemma to the extension $M \to M(\alpha)$ yields an embedding $\hat{\mu} : M(\alpha) \to \bar{K}'$ that restricts to μ on M and maps α to β. This construction $(M(\alpha), \hat{\mu})$ strictly exceeds the maximal element (M, μ), a contradiction. Thus $M = \bar{K}$.

[1] This is in some sense the infinite case of the extension lemma, the lecture didn't really cover this case in detail. To avoid awkward repetitions of this argument later, we work through the details here once. We will later simply reference this result.

A more general approach would be to first define splitting fields for *infinite* sets of polynomials; then, we could simply define the algebraic closure to be the splitting field of "all monic irreducible polynomials". Since we skipped that definition, we have to use this manual Zorn's Lemma construction instead. It feels a bit out of place compared to the rest of the chapter, but I included it nonetheless for the sake of completeness.

Finally, consider the image $\mu(\bar{K}) \subseteq \bar{K}'$. Since $\bar{K}$ is algebraically closed, its isomorphic image $\mu(\bar{K})$ is also algebraically closed. Because $\bar{K}'$ is an algebraic extension of K, it is algebraic over the intermediate field $\mu(\bar{K})$. However, the only algebraic extension of an algebraically closed field is the field itself. Therefore, $\mu(\bar{K}) = \bar{K}'$, and μ is the required K-isomorphism. □

Remark 3.4.3. It may be convenient to fix a single algebraic closure $\overline{K}$ of K. Then, for a polynomial $P \in K[x]$, the subfield $K[R_P(\bar{K})] \subseteq \overline{K}$ is clearly a splitting field of P over K.

Under this convention, we may refer to $K[R_P(\bar{K})]$ as *the* splitting field of P, understanding that it is the unique splitting field contained in $\overline{K}$.

For instance, we can fix the algebraic closure $\bar{\mathbb{Q}}$ of $\mathbb{Q}$ by treating it as a subfield of $\mathbb{C}$. Then, given $P \in \mathbb{Q}[X]$, we refer to $\mathbb{Q}[R_P(\bar{\mathbb{Q}})]$ as the splitting field of P over $\mathbb{Q}$.

3.4.1 Conjugate Elements

We fix an algebraic closure $\bar{K}$ of K.

Definition 3.4.4. $x, y \in \bar{K}$ are called **conjugate** over K if there is a $\sigma \in \mathrm{Aut}_K(\bar{K})$ with $\sigma(x) = y$.

Theorem 3.4.5. Let $x, y \in \bar{K}$. x and y are conjugate if and only if their minimal polynomials are the same in $K[X]$.

Proof. ($\implies$):
We have already concluded that given $P \in K[X]$ and $\sigma \in \mathrm{Aut}_K(\bar{K})$, $\sigma(P(x)) = P(\sigma(x))$.
This means that if $y = \sigma(x)$, because σ is an automorphism,

$$P(x) = 0 \iff \sigma(P(x)) = 0 \iff P(\sigma(x)) = P(y) = 0,$$

so their minimal polynomials must be identical.

($\impliedby$):
Let $P \in K[X]$ be the minimal polynomial of x and y.

Since $K(x) \cong K[X]/(P) \cong K(y)$, there exists a unique K-isomorphism $\phi : K(x) \to K(y)$ with $\phi(x) = y$. We now extend ϕ to the closure. Let $\mathcal{S}$ be the set of embeddings $\tau : E \to \bar{K}$ defined on subfields $K(x) \subseteq E \subseteq \bar{K}$ such that $\tau|_{K(x)} = \phi$. Ordering $\mathcal{S}$ by compatible extension allows us to apply Zorn's Lemma, yielding a maximal embedding (M, μ).

By the same maximality argument used for algebraic closures, we must have $M = \bar{K}$. If M were a proper subfield, we could pick $\alpha \in \bar{K} \setminus M$, find a root of its minimal polynomial in $\bar{K}$, and use the Extension Lemma to extend μ to $M(\alpha)$, contradicting maximality. Thus μ is defined on all of $\bar{K}$. Since $\mu(\bar{K})$ is an algebraically closed subfield of $\bar{K}$ containing K, and $\bar{K}$ is algebraic over K, we conclude $\mu(\bar{K}) = \bar{K}$. Thus μ is the required automorphism mapping x to y. □

Remark 3.4.6. Let $\bar{K}$ be an algebraic closure of K and set $G := \mathrm{Aut}_K(\bar{K})$. The group G operates on the set $\bar{K}$, which allows us to reframe the arithmetic of polynomials in terms of G-orbits.

Consider a polynomial $P \in K[X]$. Its set of roots $R_P(\bar{K}) \subseteq \bar{K}$ is G-invariant. Indeed, for any root $\alpha \in R_P(\bar{K})$ and any automorphism $\sigma \in G$,

$$P(\sigma(\alpha)) = \sigma(P(\alpha)) = \sigma(0) = 0,$$

so $\sigma(\alpha) \in R_P(\bar{K})$. Thus, G permutes the roots of P.

We show that the orbit $G \cdot \alpha$ of an element α is exactly the set of roots of its minimal polynomial μ_α. The inclusion $G \cdot \alpha \subseteq R_{\mu_\alpha}(\bar{K})$ follows from invariance. To see the reverse inclusion (transitivity), let β be any other root of the irreducible polynomial μ_α. We have shown that there must exist a unique K-isomorphism $\varphi : K(\alpha) \xrightarrow{\sim} K(\beta)$ such that $\varphi(\alpha) = \beta$.

As demonstrated in the proof of the previous theorem, this local isomorphism extends to a global automorphism $\sigma \in G$. Consequently, $\sigma(\alpha) = \beta$, proving that G acts transitively on the roots of any irreducible polynomial.

Therefore, if we take a square-free polynomial $P \in K[X]$ and factor it into distinct monic irreducibles $P = \pi_1 \cdots \pi_r$, the set of roots partitions into disjoint orbits:

$$R_P(\bar{K}) = O_1 \sqcup \cdots \sqcup O_r.$$

Here, each orbit O_i corresponds uniquely to the roots of the factor π_i.

3.5 Normal and separable Extensions

Definition 3.5.1. Let $K \to L$ be a finite field extension.

1. The extension is said to be **normal** if for all $x \in L$, the minimal polynomial P_x of x over K splits into linear factors in $L[X]$.
2. The extension is said to be **separable** if for all $x \in L$, P_x is separable.
3. The extension is said to be a **Galois extension** if it is both normal and separable.

If $K \to L$ is Galois, we denote its K-automorphism group $\mathrm{Aut}_K(L)$ as $\mathrm{Gal}(L/K)$ and call it the **Galois group**.

Remark 3.5.2.
1. If $K \subset L$ is Galois and $K \subset M \subset L$ is an intermediate field, then the extension $M \subset L$ is Galois.
2. If K is a finite field, any finite extension $K \to L$ is Galois.

Theorem 3.5.3. Let $K \to L$ be a finite extension. Let $\varphi_0 : L \to \bar{K}$ be a fixed K-embedding. The following conditions are equivalent:

1. $K \subset L$ is normal.
2. L is the splitting field of a polynomial $P \in K[X]$.
3. For any K-morphism $\varphi : L \to \bar{K}$, the images coincide: $\varphi(L) = \varphi_0(L)$.
4. The image is invariant under the automorphism group of the closure:

$$\forall \sigma \in \mathrm{Aut}_K(\bar{K}), \sigma(\varphi_0(L)) = \varphi_0(L).$$

Proof. (1) $\Rightarrow$ (2): Assume $K \subset L$ is normal. Since the extension is finite, we can find algebraic elements $\alpha_1, \ldots, \alpha_n$ such that $L = K(\alpha_1, \ldots, \alpha_n)$. Let P_i be the minimal polynomial of α_i over K. By normality, each P_i splits completely in L. Defining $P = \prod P_i$, we see that L is generated by the roots of P and contains all of them. Thus, L is the splitting field of P.

(2) $\Rightarrow$ (3): Assume L is the splitting field of $P \in K[X]$. Let R_P be the set of roots of P in $\bar{K}$. Any K-embedding $\varphi : L \to \bar{K}$ must permute these roots. Since $L = K(R_P)$, the image is determined solely by the roots:

$$\varphi(L) = \varphi(K(R_P)) = K(\varphi(R_P)) = K(R_P)$$

Since the set of roots R_P is fixed, this image is independent of φ. Thus $\varphi(L) = \varphi_0(L)$.

(3) $\Rightarrow$ (1): Identify L with $\varphi_0(L)$. Assume $\varphi(L) = L$ for any embedding φ. Let $\alpha \in L$ and let $\beta \in \bar{K}$ be any root of its minimal polynomial P_α. We must show $\beta \in L$. There exists

a K-isomorphism $\psi : K(\alpha) \to K(\beta)$. By the extension property, ψ extends to an embedding $\Phi : L \to \bar{K}$. By assumption (3), the image of this embedding must be L itself:

$$\Phi(L) = L$$

Since $\beta = \Phi(\alpha)$ and $\alpha \in L$, it follows that $\beta \in L$. Thus P_α splits in L, so L is normal.

(3) $\iff$ (4): This is immediate from the definition of the automorphism group. Any automorphism $\sigma \in \mathrm{Aut}_K(\bar{K})$ restricts to an embedding $\sigma|_L : L \to \bar{K}$. Conversely, any embedding extends to an automorphism of the closure. Thus, requiring the image to be invariant under all embeddings (3) is identical to requiring it to be invariant under all automorphisms (4). □

3.5.1 Galois Correspondence

Definition 3.5.4. Let $G \subset \mathrm{Aut}(L)$ be a subgroup. The **fixed field** of G is

$$L^G := \{x \in L \mid \sigma(x) = x \text{ for all } \sigma \in G\}.$$

This is a subfield of L.

Theorem 3.5.5 (Primitive Element Theorem)**.** Let L/K be a finite separable extension. Then $L = K(\alpha)$ for some $\alpha \in L$.

Proof. Case 1: K is infinite.

It suffices to show that any two-generator separable extension $K(\alpha, \beta)/K$ is generated by one element, the claim then follows by induction on the number of generators.

Let $P_\alpha, P_\beta \in K[x]$ be the minimal polynomials of α and β over K. Since L/K is separable, P_α and P_β are separable; let $\alpha = \alpha_1, \ldots, \alpha_r$ be the roots of P_α and $\beta = \beta_1, \ldots, \beta_s$ be the roots of P_β in some algebraic closure $\overline{K}$.

For each i and each $j \geq 2$, the equation $\alpha_i + c\beta_j = \alpha + c\beta$ has at most one solution $c \in K$. Since there are only finitely many such pairs (i, j) and K is infinite, we may choose $c \in K$ avoiding all of them. Set $\gamma := \alpha + c\beta$.

We claim $K[\gamma] = K[\alpha, \beta]$. It suffices to show $\beta \in K[\gamma]$. Consider $P := P_\alpha(\gamma - cX) \in K[\gamma][X]$. Then $P(\beta) = P_\alpha(\alpha) = 0$, so $\beta \in R_P(\bar{K})$. For $j \geq 2$ however, $P(\beta_j) = P_\alpha(\gamma - c\beta_j) \neq 0$ by our choice of c. Thus P and P_β only share the root β in $\overline{K}$, so $\gcd(P, P_\beta) = X - \beta$ in $\overline{K}[X]$. Since $P, P_\beta \in K[\gamma][X]$ and we can obtain the gcd through iterated Euclidean division (Euclidean algorithm) so that the coefficients stay within $K[\gamma]$, $X - \beta \in K[\gamma][X]$, hence $\beta \in K[\gamma]$.

Case 2: K is finite.

Then L is also finite, say $|L| = q$. It was shown in the exercises that the multiplicative group $L^\times$ is cyclic of order $q-1$; let α be a generator. Then α satisfies no polynomial over K of degree less than $[L : K]$, so $[K[\alpha] : K] = [L : K]$, giving $L = K[\alpha]$. □

Lemma 3.5.6 (Automorphism Bound)**.** Let $K \to L$ be a finite separable extension. Then,

$$|\mathrm{Aut}_K(L)| \leq [L : K].$$

Proof. By the primitive element theorem, there is $\alpha \in L$ with $L = K[\alpha]$. Any $\sigma \in \mathrm{Aut}_K(L)$ is determined entirely by $\sigma(\alpha)$, since α generates L over K and σ is K-linear. Moreover $\sigma(\alpha)$ must be a root of P_α: applying σ to $P_\alpha(\alpha) = 0$ and using that σ fixes K pointwise gives $P_\alpha(\sigma(\alpha)) = \sigma(P_\alpha(\alpha)) = 0$. Since P_α has at most $\deg P_\alpha = [L : K]$ roots in L, there are at most $[L : K]$ choices for $\sigma(\alpha)$, hence $|\mathrm{Aut}_K(L)| \leq [L : K]$. □

Theorem 3.5.7 (Artin)**.** Let L be a field and $G \subset \mathrm{Aut}(L)$ a finite subgroup. Then $L^G \to L$ is a Galois extension with $\mathrm{Gal}(L/L^G) = G$ and $[L : L^G] = |G|$.

Proof. Let $K := L^G$. We first show that $K \to L$ is Galois. Let $\alpha \in L$. The group G acts on L by evaluation, and the orbit of α is $G \cdot \alpha = \{\sigma(\alpha) \mid \sigma \in G\}$. We define

$$f_\alpha(x) := \prod_{\beta \in G\cdot\alpha} (x - \beta) \in L[x].$$

For any $\tau \in G$, the action of τ permutes the orbit $G \cdot \alpha$ (since $\tau(\sigma(\alpha)) = (\tau\sigma)(\alpha) \in G \cdot \alpha$), so when considering the obvious operation of G on $L[X]$, τ permutes the factors of f_α, leaving f_α and thus its coefficients invariant, meaning that they are in the fixed field K. Hence $f_\alpha \in K[x]$. Since f_α has distinct roots by construction, it is separable. Its roots all lie in L and α is one of them, so α is separable and algebraic over K, and f_α splits completely in L.

Since $\alpha \in L$ was arbitrary, every element of L is separable and algebraic over K and every minimal polynomial splits in L. Thus L/K is normal and separable, hence Galois.

We now show that $[L : K] = n$. Since L/K is Galois, there is a primitive element $\alpha \in L$ (such that $L = K[\alpha]$). The minimal polynomial P_α of α over K divides f_α, hence

$$[L : K] = \deg P_\alpha \leq \deg f_\alpha = |G \cdot \alpha| \leq |G|.$$

The reverse inequality follows from the automorphism bound lemma, so we get $[L : K] = |G|$.

It remains to show that $G = \mathrm{Gal}(L/K)$. $G \subseteq \mathrm{Gal}(L/K)$ is immediate from the definition (and particularly $|G| \leq |\mathrm{Gal}(L/K)|$). Because K-automorphisms σ of $L = K[\alpha]$ are uniquely determined by $\sigma(\alpha)$ which must be in $R_{P_\alpha}(L)$, we get

$$|G| \leq |\mathrm{Gal}(L/K)| \leq |R_{P_\alpha}| \leq \deg P_\alpha = [L : K] = |G|,$$

so $G = \mathrm{Gal}(L/K)$. □

Theorem 3.5.8 (Galois Correspondence)**.** Let $K \to L$ be a finite Galois extension, $G = \mathrm{Gal}(L/K)$. Then $|G| = [L : K]$. Moreover, the maps

$$\Psi : \mathcal{H} \to \mathcal{F}, \quad H \mapsto L^H, \qquad \Phi : \mathcal{F} \to \mathcal{H}, \quad M \mapsto \mathrm{Aut}_M(L),$$

between subgroups $\mathcal{H} = \{H \mid H \leq G\}$ and intermediate fields $\mathcal{F} = \{M \mid K \subseteq M \subseteq L\}$ are mutual inverses.

Proof. We first show that $|G| = [L : K]$. Since L/K is separable, the automorphism bound lemma gives $|G| \leq [L : K]$. For the reverse inequality, apply the primitive element theorem to obtain $\alpha \in L$ with $L = K(\alpha)$, and let $P_\alpha \in K[X]$ be its minimal polynomial, so $\deg P_\alpha = [L : K]$. Since L/K is normal, P_α splits completely in $L[X]$, and since L/K is separable, it has exactly $[L : K]$ distinct roots in L. Since each root $\beta \in R_{P_\alpha}(L)$ determines a unique K-automorphism $\sigma_\beta \in \mathrm{Aut}_K(L)$ with $\sigma_\beta(\alpha) = \beta$, $|G| \geq [L : K]$, and thus $|G| = [L : K]$.

Since every $\sigma \in G$ fixes K pointwise, we have $K \subseteq L^G$. Applying Artin's theorem to G acting on L gives $[L : L^G] = |G| = [L : K]$. Since $K \subseteq L^G \subseteq L$,

$$[L : K] = [L : L^G] \cdot [L^G : K] = [L : K] \cdot [L^G : K],$$

so $[L^G : K] = 1$, giving $K = L^G$.

Now we show that Φ and Ψ are inverse bijections.

Let $H \leq G$. We want to show that $(\Phi \circ \Psi)(H) = H$, i.e. $\mathrm{Aut}_{L^H}(L) = H$. Applying Artin's theorem directly to H acting on L gives $\mathrm{Aut}_{L^H}(L) = \mathrm{Gal}(L/L^H) = H$, so $\Phi \circ \Psi = \mathrm{id}_{\mathcal{H}}$.

Let $M \in \mathcal{F}$. We want to show that $(\Psi \circ \Phi)(M) = M$, i.e. $L^{\mathrm{Aut}_M(L)} = M$. Write $H' := \mathrm{Aut}_M(L)$. Every $x \in M$ satisfies $\sigma(x) = x$ for all $\sigma \in \mathrm{Aut}_M(L)$, so $M \subseteq L^{H'}$ by definition of the fixed field.

For the reverse inclusion, observe that because $K \subseteq M \subseteq L$ and $K \to L$ is Galois, $M \to L$ is also Galois. Applying Artin's theorem to $H' = \mathrm{Aut}_M(L)$ acting on L gives $[L : L^{H'}] = |H'| = |\mathrm{Aut}_M(L)|$. On the other hand, applying the first part of this theorem to the Galois extension $M \to L$ gives $[L : M] = |\mathrm{Gal}(L/M)| = |\mathrm{Aut}_M(L)|$. Hence $[L : L^{H'}] = [L : M]$. Since $M \subseteq L^{H'} \subseteq L$,

$$[L : M] = [L : L^{H'}] \cdot [L^{H'} : M],$$

so $[L^{H'} : M] = 1$, which means $L^{H'} = M$. Thus $\Psi \circ \Phi = \mathrm{id}_{\mathcal{F}}$. □

3.6 Radical Extensions and Solvability

Definition 3.6.1. Let $K \subset L$ be a finite extension, and Galois. We say that it is **elementary radical** (of type n) if there exist $a \in K$ and $n \geq 1$ such that L is the splitting field of $X^n - a$ over K.

Remark 3.6.2. Recall that for any field L and $n \geq 1$, $\mu_n(L) = \{x \in L \mid x^n = 1\} \subset L^\times$ is a cyclic group. If $K \subset L$ is elementary radical of type n, let $\alpha \in L$ be a root of $X^n - a$. The roots of $X^n - a$ are precisely $\omega\alpha$ for $\omega \in \mu_n(\bar{K})$; since L is a splitting field, all these roots lie in L. In particular, $\omega = (\omega\alpha)\alpha^{-1} \in L$ for each such ω, so $\mu_n(\bar{K}) \subset L$, and thus $|\mu_n(L)| = n$.

Definition 3.6.3. A finite Galois extension $K \subset L$ is called a **Galois radical extension** if there exists a tower

$$K \subset L_1 \subset L_2 \subset \cdots \subset L_r = L$$

such that each step $L_{i-1} \subset L_i$ is elementary radical of type n_i (where $L_0 := K$).

A polynomial $f \in K[X]$ is **solvable by radicals** if its splitting field is contained in a Galois radical extension of K.

Lemma 3.6.4. Let $K \subset L$ be an elementary radical extension of type n. If $|\mu_n(K)| = n$ (i.e., K contains a primitive n-th root of unity), then $\mathrm{Gal}(L/K)$ is isomorphic to a subgroup of $\mathbb{Z}/n\mathbb{Z}$ and in particular cyclic.

Proof. By assumption, L is the splitting field of $X^n - a$ for some $a \in K$. Let $\alpha \in L$ be a root, so $\alpha^n = a$ and $L = K(\alpha)$. Let $\zeta \in K$ be a primitive n-th root of unity. The roots of $X^n - a$ are $\zeta^k\alpha$ for $k = 0, \ldots, n-1$.

Any $\sigma \in \mathrm{Gal}(L/K)$ must permute these roots, so $\sigma(\alpha) = \zeta^{k_\sigma}\alpha$ for a unique $k_\sigma \in \mathbb{Z}/n\mathbb{Z}$. We claim that $\varphi : \mathrm{Gal}(L/K) \to \mathbb{Z}/n\mathbb{Z}$, $\sigma \mapsto k_\sigma$, is an injective group homomorphism. Indeed, since $\zeta \in K$,

$$(\sigma \circ \tau)(\alpha) = \sigma(\zeta^{k_\tau}\alpha) = \zeta^{k_\tau}\sigma(\alpha) = \zeta^{k_\tau + k_\sigma}\alpha,$$

so $k_{\sigma\tau} = k_\sigma + k_\tau$ and φ is a homomorphism. If $k_\sigma = 0$, then $\sigma(\alpha) = \alpha$ and since $L = K(\alpha)$, $\sigma = \mathrm{id}$, so φ is injective. □

Theorem 3.6.5. Let $K \subset L$ be a Galois radical extension. Then $\mathrm{Gal}(L/K)$ is solvable.

Proof. We may assume $K \subset L \subset \bar{K}$, with $\bar{K}$ a fixed algebraic closure of K. Let

$$K \subset L_1 \subset L_2 \subset \cdots \subset L_r = L$$

be a tower of elementary radical extensions, with $L_{i-1} \subset L_i$ the splitting field of $X^{n_i} - a_i$. Set $m := \prod_{i=1}^r n_i$.

Note that while each step $L_{i-1} \subset L_i$ guarantees $\mu_{n_i}(\bar{K}) \subset L_i$, this does not imply $\mu_m(\bar{K}) \subset L$ in general, (for instance, $\mu_2 \subset L$ does not imply $\mu_4 \subset L$). To fix this, we enlarge L by defining $L' := L[\mu_m(\bar{K})]$. Clearly, the extension $K \subset L'$ is finite and Galois.

We now build a tower inside L' with solvable steps. Define $L'_0 := K[\mu_m(\bar{K})]$ and $L'_i := L_i[\mu_m(\bar{K})]$ for $i = 1, \ldots, r$, so that $L'_r = L'$. This gives the tower

$$K \subset L'_0 \subset L'_1 \subset L'_2 \subset \cdots \subset L'_r = L'.$$

For the initial step, $L'_0 = K[\mu_m(\bar{K})]$ is the splitting field of $X^m - 1$ over K. Any $\sigma \in \mathrm{Gal}(L'_0/K)$ is determined by its action on $\mu_m(\bar{K})$, which must be a group automorphism of the cyclic group $\mu_m(\bar{K})$. This gives an embedding $\mathrm{Gal}(L'_0/K) \hookrightarrow \mathrm{Aut}(\mu_m(\bar{K})) \cong (\mathbb{Z}/m\mathbb{Z})^\times$, so $\mathrm{Gal}(L'_0/K)$ is abelian and thus solvable.

For $i \geq 1$, we observe that $L'_i = L_i[\mu_m(\bar{K})]$ is the splitting field of $X^{n_i} - a_i$ over L'_{i-1}. Since $n_i \mid m$, we have $\mu_{n_i}(\bar{K}) \subseteq \mu_m(\bar{K}) \subset L'_{i-1}$, so $|\mu_{n_i}(L'_{i-1})| = n_i$. By the previous lemma, the Galois group $\mathrm{Gal}(L'_i/L'_{i-1})$ embeds into $\mathbb{Z}/n_i\mathbb{Z}$ and is therefore cyclic.

The tower now yields a subnormal series of $G' := \mathrm{Gal}(L'/K)$ via the Galois correspondence, whose successive quotients are the groups $\mathrm{Gal}(L'_0/K)$ (abelian) and $\mathrm{Gal}(L'_i/L'_{i-1})$ (cyclic) for $i = 1, \ldots, r$. Since extensions of solvable groups by solvable groups are solvable, G' is solvable.

To conclude, we use the fact that L/K is Galois and $L \subseteq L'$, so the restriction map gives a surjective homomorphism $\mathrm{Gal}(L'/K) \twoheadrightarrow \mathrm{Gal}(L/K)$ with kernel $\mathrm{Gal}(L'/L)$. Since quotients of solvable groups are solvable, $\mathrm{Gal}(L/K)$ is solvable. □

To prove the converse, we need Kummer's theorem, which requires the following:

Lemma 3.6.6 (Dedekind)**.** Let L be a field and let $\sigma_1, \ldots, \sigma_n : L \to L$ be pairwise distinct field automorphisms. Then $\sigma_1, \ldots, \sigma_n$ are L-linearly independent: if $c_1, \ldots, c_n \in L$ satisfy $\sum_{i=1}^n c_i\, \sigma_i(x) = 0$ for all $x \in L$, then $c_i = 0$ for all i.

Proof. Suppose for contradiction that a nontrivial L-linear relation exists, and choose one of minimal length $r \geq 1$:

$$c_1\, \sigma_1(x) + \cdots + c_r\, \sigma_r(x) = 0 \quad \text{for all } x \in L,$$

with all $c_i \neq 0$. Since the automorphisms are pairwise distinct, $r \geq 2$. Because $\sigma_1 \neq \sigma_r$, there exists $y \in L$ with $\sigma_1(y) \neq \sigma_r(y)$. Replacing x by yx and using multiplicativity of σ_i, we obtain

$$c_1\, \sigma_1(y)\sigma_1(x) + \cdots + c_r\, \sigma_r(y)\sigma_r(x) = 0 \quad \text{for all } x \in L.$$

Multiplying the original relation by $\sigma_r(y)$ and subtracting gives

$$c_1(\sigma_1(y) - \sigma_r(y))\, \sigma_1(x) + \cdots + c_{r-1}(\sigma_{r-1}(y) - \sigma_r(y))\, \sigma_{r-1}(x) = 0 \quad \text{for all } x \in L.$$

This is a nontrivial relation of length $r-1$, since the first coefficient $c_1(\sigma_1(y) - \sigma_r(y)) \neq 0$. This contradicts the minimality of r. □

Theorem 3.6.7 (Kummer)**.** Let $K \subset L$ be a finite Galois extension of degree n with $|\mu_n(K)| = n$. Then $\mathrm{Gal}(L/K) \cong \mathbb{Z}/n\mathbb{Z}$ if and only if there exists $a \in K$ such that L is the splitting field of $X^n - a$ over K.

Proof. The "if" direction is the preceding lemma: if L is the splitting field of $X^n - a$ and $|\mu_n(K)| = n$, the Galois group embeds into $\mathbb{Z}/n\mathbb{Z}$, and since $[L : K] = |\mathrm{Gal}(L/K)| = n$, the embedding is an isomorphism.

We now prove the "only if" direction. Assume $\mathrm{Gal}(L/K) = \langle \sigma \rangle$ is cyclic of order n and let $\zeta \in K$ be a primitive n-th root of unity.

By Dedekind's lemma, the automorphisms $\mathrm{id}, \sigma, \sigma^2, \ldots, \sigma^{n-1}$ are L-linearly independent. In particular, the L-linear map $\beta \mapsto \sum_{k=0}^{n-1} \zeta^{-k}\, \sigma^k(\beta)$ is not identically zero, so there exists $\beta \in L$ with

$$\alpha := \sum_{k=0}^{n-1} \zeta^{-k}\, \sigma^k(\beta) \neq 0.$$

This element is called the **Lagrange resolvent** of β with respect to σ and ζ.

We now show that $\sigma(\alpha) = \zeta\,\alpha$. Since $\zeta \in K$,

$$\sigma(\alpha) = \sum_{k=0}^{n-1} \zeta^{-k}\,\sigma^{k+1}(\beta) = \sum_{j=1}^{n} \zeta^{-(j-1)}\,\sigma^{j}(\beta) = \zeta \sum_{j=1}^{n} \zeta^{-j}\,\sigma^{j}(\beta).$$

Since $\sigma^n = \mathrm{id}$ and $\zeta^n = 1$, the $j = n$ term $\zeta^{-n}\sigma^n(\beta) = \beta$ coincides with the $j = 0$ term $\zeta^0\,\sigma^0(\beta) = \beta$, so the last sum equals $\sum_{j=0}^{n-1} \zeta^{-j}\,\sigma^j(\beta) = \alpha$. Thus $\sigma(\alpha) = \zeta\,\alpha$.

It follows that $\sigma(\alpha^n) = (\zeta\alpha)^n = \zeta^n\alpha^n = \alpha^n$, and since σ generates $\mathrm{Gal}(L/K)$, every element of the Galois group fixes α^n. Thus $a := \alpha^n \in L^G = K$.

We now verify that $L = K(\alpha)$. Since $\sigma(\alpha) = \zeta\alpha \neq \alpha$ (as ζ is a primitive n-th root and $\alpha \neq 0$), the generator σ does not fix α. More generally, $\sigma^m(\alpha) = \zeta^m\alpha$, which equals α only when $n \mid m$. So the stabiliser of α in G is trivial, giving $|G \cdot \alpha| = n$. This implies $[K(\alpha) : K] \geq n = [L : K]$, so $K(\alpha) = L$.

Finally, $X^n - a$ has α as a root, and its roots in $\bar{K}$ are $\zeta^j\alpha$ for $j = 0, \ldots, n-1$. Since $\zeta \in K$ and $\alpha \in L$, all roots lie in L, and they generate $K(\alpha) = L$. Thus L is the splitting field of $X^n - a$ over K. □

Theorem 3.6.8. Let $\mathrm{char}(K) = 0$. If $K \subset L$ is a finite Galois extension with $\mathrm{Gal}(L/K)$ solvable, then there exist $n \geq 1$ and a Galois radical tower

$$K \subset K[\mu_n(\bar{K})] \subset \cdots \subset L[\mu_n(\bar{K})]$$

such that $L \subset L[\mu_n(\bar{K})]$ and $K \subset L[\mu_n(\bar{K})]$ is Galois radical.

Proof. Set $G := \mathrm{Gal}(L/K)$. Since G is solvable and finite, there is a subnormal series

$$\{e\} = G_r \trianglelefteq G_{r-1} \trianglelefteq \cdots \trianglelefteq G_1 \trianglelefteq G_0 = G$$

with G_{i-1}/G_i cyclic of order m_i for each i. By the Galois correspondence, the fixed fields $F_i := L^{G_i}$ give a tower

$$K = F_0 \subset F_1 \subset F_2 \subset \cdots \subset F_r = L,$$

and each step $F_{i-1} \subset F_i$ is a Galois extension with $\mathrm{Gal}(F_i/F_{i-1}) \cong G_{i-1}/G_i$, which is cyclic of order m_i.

Set $n := \prod_{i=1}^{r} m_i$. Since $\mathrm{char}(K) = 0$, the polynomial $X^n - 1$ is separable, so its splitting field over K contains exactly n roots and $|\mu_n(\bar{K})| = n$.

We adjoin all n-th roots of unity to each level of the tower. Define $L_i' := F_i[\mu_n(\bar{K})]$ for $i = 0, \ldots, r$, so that $L_0' = K[\mu_n(\bar{K})]$ and $L_r' = L[\mu_n(\bar{K})]$. This gives the tower

$$K \subset L_0' \subset L_1' \subset L_2' \subset \cdots \subset L_r'.$$

We show that each step is elementary radical.

Consider the step $L_{i-1}' \subset L_i'$ for some $i \in \{1, \ldots, r\}$. The Galois group of L_i'/L_{i-1}' can be understood via the restriction map: since $L_i' = F_i[\mu_n(\bar{K})]$ and $L_{i-1}' = F_{i-1}[\mu_n(\bar{K})]$, an automorphism $\tau \in \mathrm{Gal}(L_i'/L_{i-1}')$ restricts to an automorphism of F_i that fixes F_{i-1}, because τ fixes $L_{i-1}' \supset F_{i-1}$ pointwise and stabilises F_i (since $F_i \subset L_i'$ and F_i is Galois over F_{i-1}). The restriction map $\mathrm{Gal}(L_i'/L_{i-1}') \to \mathrm{Gal}(F_i/F_{i-1})$ is injective: if τ fixes both F_i and $\mu_n(\bar{K})$, it fixes $L_i' = F_i[\mu_n(\bar{K})]$ pointwise. Thus

$$\mathrm{Gal}(L_i'/L_{i-1}') \hookrightarrow \mathrm{Gal}(F_i/F_{i-1}) \cong \mathbb{Z}/m_i\mathbb{Z},$$

so $\mathrm{Gal}(L_i'/L_{i-1}')$ is cyclic, say of order $d_i \mid m_i$. Since $m_i \mid n$, we have $d_i \mid n$, so $|\mu_{d_i}(L_{i-1}')| = d_i$ (because $\mu_{d_i}(\bar{K}) \subset \mu_n(\bar{K}) \subset L_{i-1}'$). By the Kummer theorem, there exists $a_i \in L_{i-1}'$ such that

L'_i is the splitting field of $X^{d_i} - a_i$ over L'_{i-1}, so the step $L'_{i-1} \subset L'_i$ is elementary radical of type d_i.

It remains to show that the initial step $K \subset L'_0 = K[\mu_n(\bar{K})]$ can also be decomposed into elementary radical steps. Since L'_0 is the splitting field of $X^n - 1$ over K, the Galois group $\mathrm{Gal}(L'_0/K)$ embeds into $\mathrm{Aut}(\mu_n(\bar{K})) \cong (\mathbb{Z}/n\mathbb{Z})^\times$, which is abelian. In particular, $\mathrm{Gal}(L'_0/K)$ is abelian, hence solvable. Being a finite abelian group, it admits a composition series with cyclic quotients of prime order. By the Galois correspondence, this gives a tower of intermediate fields

$$K = E_0 \subset E_1 \subset \cdots \subset E_s = L'_0$$

with each $\mathrm{Gal}(E_j/E_{j-1})$ cyclic of prime order p_j. Since $p_j \mid n$, the field E_{j-1} contains a primitive p_j-th root of unity (because $\mu_{p_j}(\bar{K}) \subset \mu_n(\bar{K}) \subset E_1 \subset \cdots \subset E_{j-1}$ as soon as $j \geq 2$). The first step $E_0 \subset E_1$ requires separate attention: if $p_1 = 2$, any field of characteristic $\neq 2$ contains -1, a primitive second root of unity; if p_1 is odd, we use the fact that $\mathrm{Gal}(L'_0/K)$ is abelian, so we may reorder the composition factors to place a step that already has the required roots of unity first (or note that E_1/K is Galois of prime degree, and it is itself the splitting field of the p_1-th cyclotomic polynomial, which factors into linear terms over E_1).

In either case, for $j \geq 2$, the Kummer theorem applies to $E_{j-1} \subset E_j$ and produces $b_j \in E_{j-1}$ with E_j the splitting field of $X^{p_j} - b_j$. Thus each step is elementary radical.

Concatenating this tower with the one constructed above, we obtain a Galois radical tower from K to $L'_r = L[\mu_n(\bar{K})]$, and $L \subset L[\mu_n(\bar{K})]$. □

Corollary 3.6.9. Let $\mathrm{char}(K) = 0$, let $f \in K[X]$ be non-constant, and let L be the splitting field of f over K. Then

$$f \text{ is solvable by radicals} \iff \mathrm{Gal}(L/K) \text{ is solvable.}$$

Proof. ($\Leftarrow$): This is precisely the preceding theorem: L is contained in the Galois radical extension $L[\mu_n(\bar{K})]$ of K.

($\Rightarrow$): Suppose L is contained in a Galois radical extension M of K. We may assume M/K is Galois.

By the first theorem of this section, $\mathrm{Gal}(M/K)$ is solvable. Since L/K is Galois, the restriction map gives a surjection $\mathrm{Gal}(M/K) \twoheadrightarrow \mathrm{Gal}(L/K)$ with kernel $\mathrm{Gal}(M/L)$, so

$$\mathrm{Gal}(L/K) \cong \mathrm{Gal}(M/K)/\mathrm{Gal}(M/L).$$

Since quotients of solvable groups are solvable, $\mathrm{Gal}(L/K)$ is solvable. □