WP24 Commutative Algebra

Comprehensive notes by
Lina Weitzenböck

Lectures by
Prof. Fabien Morel
Summer 2026

Course description. This lecture will begin with some complementary topics in Galois theory. Specifically, for a polynomial P with integer coefficients and its splitting field $L|\mathbb{Q}$, we will describe how the decomposition of P modulo a prime number p provides information about the Galois group of L over $\mathbb{Q}$. An application of this will be proving that the Galois group over $\mathbb{Q}$ of a generic polynomial of degree n with integer coefficients is the symmetric group S_n.

We will then introduce the main topic of this course: commutative algebra. Starting with some generalities, we will introduce the spectrum of a commutative ring, which is the set of its prime ideals endowed with the Zariski topology. We will discuss various properties, including localization, Noetherianity, and dimension. We will also introduce the notion of a Dedekind domain and study its primary properties. These rings are crucial in both number theory and algebraic geometry, serving as the algebraic equivalent of "smooth curves".

Next, we will move to the study of the category of R-modules over a commutative ring R, covering exact sequences, tensor products, localization, flatness, and related concepts. After that, we will examine in greater detail the theory of Krull dimension for finite-type algebras over a field k. This will involve the Noether normalization lemma and some of its applications, such as Hilbert's Nullstellensatz and the fact that the dimension of a finite-type algebra A over a field k that is an integral domain equals the transcendence degree of its fraction field over k. We will conclude by introducing and studying the notion of regular rings, including a characterization of regular local rings.

Course page. https://www.mathematik.uni-muenchen.de/~morel/Teaching/AAteaching.html.

Exam date. Thu 16 Jul, 10:00-12:00, B132.

Contents

Date of version: Sun 31 May 2026

Part I
Complementary Topics in Galois Theory

1 Quick Recap of Galois Theory

Let $K \subset L$ be a field extension. We recall the basic definitions and results of finite Galois extensions, which serve as a prerequisite for Kummer theory and the study of solvable extensions.

Definition (Finite Galois Extension). Let $K \subset L$ be a finite field extension with degree $[L : K] = n \in \mathbb{N}$. The extension is said to be *Galois* if it satisfies both of the following properties:

1. *Normal*: For all $x \in L$, its minimal polynomial over K, $P_x \in K[X]$, splits completely in L. That is, $P_x = \prod_a (X - a) \in L[X]$.
2. *Separable*: For all $x \in L$, the minimal polynomial P_x is separable, meaning it has no multiple roots. Equivalently, $\gcd(P_x', P_x) = 1$.

A basic, yet fundamental, example is the splitting field of a polynomial. Let $P \in \mathbb{Q}[X]$ and let L be its splitting field over $\mathbb{Q}$. This can be viewed inside the complex numbers as $\mathbb{Q} \subset \mathbb{Q}(R_P(\mathbb{C})) \subset \mathbb{C}$, where $R_P(\mathbb{C})$ denotes the roots of P in $\mathbb{C}$. This extension is Galois, and we have a natural injective group homomorphism into the symmetric group of the roots:

$$\mathrm{Gal}(L/\mathbb{Q}) \hookrightarrow S_{R_P(\mathbb{C})}.$$

Furthermore, if the polynomial decomposes as $P = \prod_i P_i$ with $P_i \in \mathbb{Q}[X]$, then the Galois group embeds into the direct product of the symmetric groups of the respective roots:

$$\mathrm{Gal}(L/\mathbb{Q}) \hookrightarrow \prod_i S_{R_{P_i}(\mathbb{C})}.$$

Theorem (Fundamental Theorem of Galois Theory). If $K \subset L$ is a finite Galois extension, we denote the Galois group by $G = \mathrm{Gal}(L/K)$, where $|G| = [L : K]$. There exist inclusion-reversing bijections between the set of intermediate fields M and the set of subgroups H of G:

$$\begin{aligned} \{K \subset M \subset L\} &\longleftrightarrow \{H \subset G\} \\ M &\longmapsto \mathrm{Gal}(L/M) \\ L^H &\longmapsfrom H \end{aligned}$$

These mappings are mutual inverses.

$$\begin{array}{ccc} L & & 1 \\ | & & | \\ M & \overset{1:1}{\longleftrightarrow} & H \\ | & & | \\ K & & G \end{array}$$

An immediate corollary of the Galois correspondence is that the fixed field of the full Galois group G is exactly the base field, $K = L^G$. Consequently, if $x \in L \setminus K$, there exists at least one automorphism $\sigma \in G$ such that $\sigma(x) \neq x$. The set of images $\{\sigma(x) \mid \sigma \in G\}$ are called the *conjugates* of x.

Theorem (Artin's Theorem). Let L be a field and let $H \subset \mathrm{Aut}(L)$ be a finite subgroup of automorphisms of L. Then the fixed field $K = L^H \subset L$ is a Galois extension of degree $[L : K] = |H|$, and the Galois group is precisely $\mathrm{Gal}(L/K) = H$.

Another core consequence of the Galois correspondence pertains to normal subgroups and normal extensions. Let $K \subset L$ be a finite Galois extension with Galois group $G = \mathrm{Gal}(L/K)$, and let $K \subset M \subset L$ be an intermediate field. Then the extension $K \subset M$ is *Galois* if and only

if the corresponding subgroup $\mathrm{Gal}(L/M)$ is a *normal subgroup* of G. In this case, we have the isomorphism:

$$\mathrm{Gal}(M/K) \cong \frac{G}{\mathrm{Gal}(L/M)}.$$

2 Solvability and Kummer Theory

We now shift our focus to solvable extensions and their precise correspondence to radical extensions, culminating in Kummer's theorem.

2.1 Radical Extensions and Solvable Groups

Definition (Elementary Radical and Radical Extensions). Let $K \subset L$ be a finite extension of degree n.

1. The extension is called *elementary radical* if the base field K contains all n-th roots of unity, i.e., $|\mu_n(K)| = n$, and there exists an element $\alpha \in L$ such that $L = K[\alpha]$ and $\alpha^n = a \in K$. We often denote this as $\alpha = \sqrt[n]{a}$. The polynomial $X^n - a$ decomposes as $\prod_{\delta \in \mu_n(K)}(X - \delta\alpha)$.
2. The extension $K \subset L$ is called *radical* if there exists a tower of intermediate fields $K \subset L_1 \subset \cdots \subset L_s = L$ such that each step $L_{i-1} \subset L_i$ is an elementary radical extension.

For an elementary radical extension $L = K[\alpha]$ with $\alpha^n = a$, there is a natural injective group homomorphism:

$$\begin{aligned} \mathrm{Gal}(L/K) &\hookrightarrow \mu_n(K) \\ \sigma &\longmapsto \frac{\sigma(\alpha)}{\alpha}. \end{aligned}$$

A standard example of a radical extension is the splitting field of $X^n - 1$ over an arbitrary field K.

Theorem (Solvability and Radical Extensions Theorem). Let $K \subset L$ be a finite Galois extension with Galois group $G = \mathrm{Gal}(L/K)$ and degree $[L : K] = n = |G|$. Assume that the base field contains all n-th roots of unity, i.e., $|\mu_n(K)| = n$. Then:

$$G \text{ is a solvable group} \iff K \subset L \text{ is a radical extension.}$$

If the assumption $|\mu_n(K)| = n$ is not met, but $\gcd(\mathrm{char}(K), n) = 1$, we can adjoin the roots of unity. Let K_n be the splitting field of $X^n - 1$ over K, and L_n be the compositum of L and K_n. Then $\mathrm{Gal}(L/K)$ is solvable if and only if $\mathrm{Gal}(L_n/K_n)$ is solvable.

Proof Sketch (Solvability and Radical Extensions). We sketch the proof, which is structurally straightforward.

Radical $\implies$ Solvable: Assume $K \subset L$ is radical, and $|\mu_n(K)| = n$. There exists a tower $K \subset L_1 \subset \cdots \subset L_s = L$ where each $L_{i-1} \subset L_i$ is elementary radical. By the Galois correspondence, let $H_i = \mathrm{Gal}(L/L_i) \subset G$. Then $H_i \subset H_{i-1}$ is a normal subgroup, and the quotient $H_{i-1}/H_i \cong \mathrm{Gal}(L_i/L_{i-1})$ embeds into $\mu_{n_i}(K)$, meaning it is cyclic. This yields a subnormal series $1 = G_s \subset \cdots \subset G_1 \subset G$ with abelian (cyclic) quotients, proving G is solvable.

Solvable $\implies$ Radical: Assume $\mathrm{Gal}(L/K)$ is solvable and $|\mu_n(K)| = n$. We have a sequence of subgroups $1 = H_s \subset \cdots \subset H_1 = G$ such that $H_i \subset H_{i-1}$ is normal and the quotient H_{i-1}/H_i is cyclic, isomorphic to $\mathbb{Z}/p_i\mathbb{Z}$ for some prime $p_i \mid n$. Setting $L_i = L^{H_i}$ gives a tower $K \subset L_1 \subset \cdots \subset L_s = L$ where each L_i is Galois over L_{i-1} with a cyclic Galois group. By Kummer Theory (proved below), each step is elementary radical, hence $K \subset L$ is radical.

□

2.2 Kummer's Theorem

Theorem (Kummer's Theorem). Let $K \subset L$ be a finite Galois extension of degree $n = [L : K]$. Assume that the Galois group is cyclic, $G \cong \mathbb{Z}/n\mathbb{Z}$, and that K contains the n-th roots of unity, $|\mu_n(K)| = n$. Then there exists an element $\alpha \in L$ such that $L = K[\alpha]$ and $\alpha^n = a \in K$.

(In other words, L is an elementary radical extension of K).

Note that under these hypotheses, the minimal polynomial of α is exactly $X^n - a$, which has n distinct roots of the form $\delta\alpha$ for $\delta \in \mu_n(K)$. The condition $|\mu_n(K)| = n$ ensures that $X^n - a$ is irreducible over K, which is equivalent to the statement that the class $\overline{a}$ has order n in the quotient group $K^\times/(K^\times)^n$.

Proof (Kummer's Theorem). The assumption $|\mu_n(K)| = n$ implies that $n \neq 0$ in K, so $\gcd(\mathrm{char}(K), n) = 1$. Let σ be a generator of the cyclic Galois group $\mathrm{Gal}(L/K) \cong \mathbb{Z}/n\mathbb{Z}$. We can view σ as a K-endomorphism of L. Since the order of σ is n, we have $\sigma^n = \mathrm{id}_L$. Thus, the polynomial $X^n - 1 \in K[X]$ annihilates σ.
Because $\gcd(\mathrm{char}(K), n) = 1$, the polynomial $X^n - 1$ separates into distinct linear factors over K: $X^n - 1 = \prod_{\delta\in\mu_n(K)}(X - \delta)$. The minimal polynomial of the linear operator σ, $\mu_\sigma \in K[X]$, must divide $X^n - 1$, and therefore μ_σ also splits into a product of distinct linear factors. Consequently, σ is diagonalizable over K! We obtain the eigenspace decomposition:

$$L = \bigoplus_{\delta\in\mu_n(K)} L_\delta, \quad \text{where } L_\delta = \ker(\sigma - \delta\,\mathrm{id}_L).$$

Observe how these eigenspaces multiply. If $x \in L_\alpha$ and $y \in L_\beta$ for some $\alpha, \beta \in \mu_n(K)$, then:

$$\sigma(xy) = \sigma(x)\sigma(y) = (\alpha x)(\beta y) = (\alpha\beta)(xy).$$

Thus, $xy \in L_{\alpha\beta}$. Since L is a field and $L \neq \{0\}$, at least one $L_\alpha \neq \{0\}$. Take any non-zero element $x_0 \in L_\alpha \setminus \{0\}$. For any $\beta \in \mu_n(K)$, the multiplication map gives a K-linear isomorphism:

$$\begin{aligned} L_\beta &\xrightarrow{\sim} L_{\alpha\beta} \\ y &\longmapsto yx_0 \end{aligned}$$

with inverse $z \mapsto zx_0^{-1}$ (since $x_0^{-1} \in L_{\alpha^{-1}}$). This implies that all the eigenspaces L_δ are isomorphic to each other as K-vector spaces. Since their direct sum is L, which has dimension n, and there are n such eigenspaces, it must be that $\dim_K L_\delta = 1$ for all $\delta \in \mu_n(K)$.
Now, let $\delta_0 \in \mu_n(K)$ be a primitive n-th root of unity (a generator of the group $\mu_n(K)$). Choose a non-zero eigenvector $x_0 \in L_{\delta_0} \setminus \{0\}$. The elements $1, x_0, x_0^2, \ldots, x_0^{n-1}$ fall into distinct eigenspaces corresponding to $1, \delta_0, \delta_0^2, \ldots, \delta_0^{n-1}$, and thus they form a K-basis for L.
Finally, observe the n-th power: x_0^n belongs to the eigenspace $L_{\delta_0^n} = L_1$. The eigenspace for eigenvalue 1 is precisely the fixed field: $L_1 = \{x \in L \mid \sigma(x) = x\} = K$. Letting $a = x_0^n \in K$, we conclude that $L \cong K[X]/(X^n - a)$, which is exactly the definition of an elementary radical extension.

□

End of Lec 01.

3 Reduction modulo p and the Galois Group over $\mathbb{Q}$

3.1 Polynomial Reduction and Algebraic Integers

Let $\mathbb{Q} \subset K$ be a finite field extension, so K is a number field. By the Primitive Element Theorem, there exists an element $\alpha \in K$ such that $K = \mathbb{Q}[\alpha]$. Let $[K : \mathbb{Q}] = n$.
Let $Q \in \mathbb{Q}[X]$ be the minimal polynomial of α over $\mathbb{Q}$, which takes the form:

$$Q = X^n + a_{n-1}X^{n-1} + \cdots + a_0.$$

In general, Q does not necessarily have integer coefficients. However, there exists a nonzero integer $b \in \mathbb{Z} \setminus \{0\}$ such that evaluating Q at X/b and multiplying by b^n clears all denominators appropriately: $b^nQ(X/b) \in \mathbb{Z}[X]$. Consequently, we can define a monic polynomial $P \in \mathbb{Z}[X]$ as follows:

$$P(X) := b^nQ\left(\frac{X}{b}\right) = X^n + ba_{n-1}X^{n-1} + \cdots + b^na_0.$$

Evaluating this polynomial at $b\alpha$ yields $P(b\alpha) = b^nQ(\alpha) = 0$. Thus, $b\alpha$ is also a generator of the extension $K|\mathbb{Q}$, and its minimal polynomial over $\mathbb{Q}$ is precisely P, which is both monic and has integer coefficients.

3.1.1 The Ring of Integers of $\mathbb{Q}$

Definition (Integral Elements and the Ring of Integers). An element $\alpha \in \overline{\mathbb{Q}}$ is called *integral* over $\mathbb{Z}$ (or an algebraic integer) if it is a root of some monic polynomial in $\mathbb{Z}[X]$.
For a finite field extension $\mathbb{Q} \subset K$, the set of all elements in K that are integral over $\mathbb{Z}$ forms a ring, denoted $\mathcal{O}_K$, called the *ring of integers* of K:

$$\mathcal{O}_K = \{x \in K \mid x \text{ is integral over } \mathbb{Z}\}.$$

As a basic example, let us determine the ring of integers of the base field $\mathbb{Q}$.

Corollary (The Ring of Integers of $\mathbb{Q}$). The algebraic integers within the rational numbers are precisely the standard integers: $\mathcal{O}_{\mathbb{Q}} = \mathbb{Z}$.

Proof ($\mathcal{O}_{\mathbb{Q}} = \mathbb{Z}$). Let $x \in \mathbb{Q}$ and write $x = \frac{a}{b}$ for $(a, b) \in \mathbb{Z}^2$ with $\gcd(a, b) = 1$. Suppose $x \in \mathcal{O}_{\mathbb{Q}}$. By definition, there exists a monic polynomial $P \in \mathbb{Z}[X]$ such that $P(x) = 0$. Let $P = X^d + c_{d-1}X^{d-1} + \cdots + c_0$. Substituting x gives:

$$\left(\frac{a}{b}\right)^d + c_{d-1}\left(\frac{a}{b}\right)^{d-1} + \cdots + c_0 = 0.$$

Multiplying the entire equation by b^d to clear the denominators, we obtain:

$$a^d + c_{d-1}a^{d-1}b + \cdots + c_0 b^d = 0.$$

Rearranging this yields $a^d = -b(c_{d-1}a^{d-1} + \cdots + c_0 b^{d-1})$. Since the term in the parentheses is an integer, this implies that b divides a^d. However, we assumed $\gcd(a, b) = 1$, which forces $b = \pm 1$. Therefore, $x = \pm a \in \mathbb{Z}$.

□

3.1.2 The Subring of Roots and its Reduction Modulo p

Let $P \in \mathbb{Z}[X]$ be a monic, separable polynomial. Let $L \subset \mathbb{C}$ be the splitting field of P over $\mathbb{Q}$. Let $n = [L : \mathbb{Q}]$ and let $G = \mathrm{Gal}(L/\mathbb{Q})$, so $|G| = n$. We denote the set of roots of P in $\mathbb{C}$ by $R = \{\alpha_1, \ldots, \alpha_l\}$, where $l = \deg P = |R|$.
We introduce a subring $A \subset L$ defined as $A = \mathbb{Z}[R] = \mathbb{Z}[\alpha_1, \ldots, \alpha_l]$, the subring generated by the roots of P. Since each root α_i is a root of a monic polynomial in $\mathbb{Z}[X]$, $\alpha_i \in \mathcal{O}_L$. Consequently, $A \subset \mathcal{O}_L$, although in general $A \neq \mathcal{O}_L$.

Lemma (Structure of the Ring A). The ring A is a free abelian group (a free $\mathbb{Z}$-module) of rank $n = [L : \mathbb{Q}]$.

Proof (Structure of A). Consider the natural evaluation homomorphism $\varphi : \mathbb{Z}[X_1, \ldots, X_l] \to A$ given by $X_i \mapsto \alpha_i$. This induces a surjection from the quotient:

$$\frac{\mathbb{Z}[X_1, \ldots, X_l]}{(P(X_1), \ldots, P(X_l))} \twoheadrightarrow A.$$

The domain of this induced map is a free abelian group of rank l^l (which can be seen by induction on l). Let $A' \subset A$ be the subgroup generated by elements of the form $\alpha^{\underline{u}} = \alpha_1^{u_1} \ldots \alpha_l^{u_l}$ where the exponents satisfy $u_i \leq l - 1$ for all i. We claim $A' = A$.
To see this, consider what happens if we multiply an element of A' by α_1. If $u_1 < l - 1$, the degree increases but remains within the bounds defining A'. If $u_1 = l - 1$, we use the fact that $P(\alpha_1) = 0$. Since P is monic of degree l, we can write $\alpha_1^l = -c_{l-1}\alpha_1^{l-1} - \cdots - c_0$, where $c_k \in \mathbb{Z}$. Thus, α_1^l reduces to a $\mathbb{Z}$-linear combination of lower powers. By symmetry, $\alpha_i A' \subset A'$ for all i. Since $1 \in A'$, this ideal property implies $A' = A$. Therefore, A is generated by a finite basis over $\mathbb{Z}$.
Furthermore, for any $x \in L$, there exists $b \in \mathbb{Z} \setminus \{0\}$ such that $bx \in A$. Since $L = \mathbb{Q}[\alpha_1, \ldots, \alpha_l]$, we deduce that $L = \mathrm{Frac}(A)$. Since $A \subset L$, A is torsion-free. A finitely generated torsion-free module over $\mathbb{Z}$ is free, so $A \cong \mathbb{Z}^r$.
Let $(x_1, \ldots, x_r)$ be a $\mathbb{Z}$-basis of A. This basis serves as a generating family for L over $\mathbb{Q}$, and it is also linearly independent over $\mathbb{Q}$. Therefore, the rank must be $r = [L : \mathbb{Q}] = n$.

□

Now, let p be a prime number such that the reduction $\overline{P} \in \mathbb{F}_p[X]$ is separable. We consider the quotient ring $\overline{A} = A/pA$. Since $A \cong \mathbb{Z}^n$, we naturally have $\overline{A} \cong (\mathbb{Z}/p\mathbb{Z})^n = \mathbb{F}_p^n$ as an $\mathbb{F}_p$-vector space.
The Galois group G operates on A (as it permutes the roots R), and this action leaves the ideal pA stable. Thus, G acts on the finite $\mathbb{F}_p$-algebra $\overline{A}$.
Let $\mathcal{M}$ be the set of maximal ideals of $\overline{A}$. Let L' be the splitting field of $\overline{P}$ over $\mathbb{F}_p$, and let $\overline{R} = R_{\overline{P}}(L')$ be the roots of $\overline{P}$ in L'. Since $\overline{P}$ is separable, $|\overline{R}| = l$.

Proposition (Homomorphisms into the Splitting Field).

1. There is a natural bijection $\mathrm{Hom}_{\mathbb{F}_p\text{-alg}}(\overline{A}, L') \cong \mathrm{Hom}_{\mathbb{Z}\text{-alg}}(A, L')$. Any such homomorphism is entirely determined by an *injective mapping* from the generators of $\overline{A}$ into $\overline{R}$.
2. For all $s \in \mathrm{Hom}_{\mathbb{F}_p\text{-alg}}(\overline{A}, L')$, its kernel $\ker(s)$ is a maximal ideal $\mathfrak{m} \in \mathcal{M}$.
3. The Galois group G operates freely on the set $\mathrm{Hom}_{\mathbb{F}_p\text{-alg}}(\overline{A}, L')$. That is, if $g \cdot s = s$ for some $g \in G$, then $g = \mathrm{id}_G$.

Proof (Homomorphisms into the Splitting Field). [1]

1. We can view $\overline{A}$ as a quotient: $\overline{A} = A/pA \cong \mathbb{F}_p[\overline{\alpha}_1, \dots, \overline{\alpha}_l]/\dots$, where $\overline{\alpha}_i = \alpha_i \pmod p$. An $\mathbb{F}_p$-algebra homomorphism s from $\overline{A}$ to L' is completely determined by the images of the generators $\overline{\alpha}_i$. Since $P(\alpha_i) = 0$, we must have $\overline{P}(s(\overline{\alpha}_i)) = 0$, meaning the generators are mapped to roots of $\overline{P}$ in L'.
   The set of all possible mappings from the generators to $\overline{R}$ has a cardinality of l^l. The separability of $\overline{P}$ ensures that its discriminant $\Delta(P)$ is non-zero modulo p. If s were to map two distinct generators $\overline{\alpha}_i$ and $\overline{\alpha}_j$ to the same root in L', their difference would lie in $\ker(s)$, forcing $s(\overline{\Delta(P)}) = 0$. Since $\overline{\Delta(P)} \in \mathbb{F}_p^\times$ and s preserves the identity, this is a contradiction. Consequently, s restricts to an injective mapping from the generators into $\overline{R}$.
2. For any map s, the first isomorphism theorem gives an injection $\overline{A}/\ker(s) \hookrightarrow L'$. Since $\overline{A}$ is a finite ring, the quotient $\overline{A}/\ker(s)$ is a finite integral domain, which is necessarily a field. Hence, $\ker(s)$ must be a maximal ideal.
3. If $g \in G$ stabilizes a homomorphism s, the action of g on the roots commutes with their mapping into $\overline{R}$. This implies $s(\overline{g \cdot \alpha}) = s(\overline{\alpha})$ for all complex roots $\alpha \in R$. As established, s acts as an injective assignment on these reduced roots. Therefore, the relation $s(\overline{g \cdot \alpha}) = s(\overline{\alpha})$ forces $\overline{g \cdot \alpha} = \overline{\alpha}$. Because the discriminant $\Delta(P)$ is non-zero modulo p, the roots remain distinct modulo p. This ensures that any equality of roots in the characteristic p reduction uniquely implies an equality in characteristic 0, meaning $g \cdot \alpha = \alpha$. Because this lift holds for all roots generating L over $\mathbb{Q}$, g must be the identity automorphism id_G.

□

We recall a standard consequence of the Chinese Remainder Theorem:

Lemma (Product of Quotients). Let R be a commutative ring and $\mathfrak{m}_1, \dots, \mathfrak{m}_k$ be maximal ideals such that $\mathfrak{m}_i + \mathfrak{m}_j = R$ for $i \neq j$. The canonical map $R \to \prod_{i=1}^k R/\mathfrak{m}_i$ is surjective.
As a corollary, the natural map $\overline{A} \to \prod_{\mathfrak{m} \in \mathcal{M}} \overline{A}/\mathfrak{m}$ is surjective. By computing dimensions over $\mathbb{F}_p$, we know that $|G| = n = \dim_{\mathbb{F}_p}(\overline{A})$. Setting $\overline{G} = \mathrm{Gal}(L'/\mathbb{F}_p)$, we can deduce a factorization of the size of the homomorphism set:

$$|\mathrm{Hom}_{\mathbb{F}_p\text{-alg}}(\overline{A}, L')| = |\mathcal{M}| \times |\overline{G}|.$$

Further developments regarding these relations and the structure of G will be discussed in the next lecture.

End of Lec 02.

3.1.3 Surjectivity of the Canonical Projection

To complete the discussion from the previous lecture regarding the Chinese Remainder Theorem, we explicitly demonstrate the surjectivity of the canonical projection map.

[1] Note: The live lecture omitted the discriminant argument for injectivity and conflated the domain generators with the target root set.

Proof (Surjectivity of the Canonical Map). Let A be a commutative ring and let $\mathfrak{m}_1, \ldots, \mathfrak{m}_r$ be finitely many maximal ideals in A. Since the ideals are maximal and distinct, they are pairwise coprime, meaning $\mathfrak{m}_i + \mathfrak{m}_j = A$ for all $i \neq j$. Therefore, for any pair $i \neq j$, there exist elements $u_{ij} \in \mathfrak{m}_i$ and $v_{ij} \in \mathfrak{m}_j$ such that

$$u_{ij} + v_{ij} = 1.$$

For each $i \in \{1, \ldots, r\}$, we define the element

$$e_i := \prod_{j \neq i} v_{ij}.$$

Because $v_{ij} = 1 - u_{ij} \equiv 1 \pmod{\mathfrak{m}_i}$, it follows that $e_i \equiv 1 \pmod{\mathfrak{m}_i}$. Conversely, for any $j \neq i$, the product defining e_i contains the factor $v_{ij} \in \mathfrak{m}_j$, which directly implies that $e_i \equiv 0 \pmod{\mathfrak{m}_j}$. Now, consider an arbitrary tuple $(\lambda_1, \ldots, \lambda_r) \in A/\mathfrak{m}_1 \times \cdots \times A/\mathfrak{m}_r$. We can lift each λ_i to an element $x_i \in A$. Construct the element

$$x := \sum_{i=1}^{r} x_i e_i \in A.$$

Reducing x modulo $\mathfrak{m}_i$ yields $x \equiv x_i e_i \equiv x_i \pmod{\mathfrak{m}_i}$, since all other terms $x_j e_j$ vanish modulo $\mathfrak{m}_i$. Thus, x maps precisely to $(\lambda_1, \ldots, \lambda_r)$ under the canonical projection, proving that the map is surjective.

□

3.2 Transitivity and the Decomposition Group

Let us briefly recall the setting from the end of the last lecture: $P \in \mathbb{Z}[X]$ is a monic polynomial whose reduction $\overline{P} \in \mathbb{F}_p[X]$ is separable. We denoted by L the splitting field of P over $\mathbb{Q}$, with Galois group $G = \mathrm{Gal}(L/\mathbb{Q})$. We defined the ring $A = \mathbb{Z}[R_P(L)]$, where $R_P(L)$ are the roots of P in L, and considered its reduction $\overline{A} = A/pA$. The set of maximal ideals of $\overline{A}$ is denoted by $\mathcal{M}$. Furthermore, let L' be the splitting field of $\overline{P}$ over $\mathbb{F}_p$.
We previously established that the Galois group G acts freely on the set of homomorphisms $S := \mathrm{Hom}_{\mathbb{F}_p\text{-alg}}(\overline{A}, L')$. Through dimensionality arguments over $\mathbb{F}_p$, we also found the cardinality:

$$|S| = |\mathcal{M}| \times [L' : \mathbb{F}_p].$$

Proposition (Transitivity of the Galois Action). The action of the Galois group $G = \mathrm{Gal}(L/\mathbb{Q})$ on the set $S = \mathrm{Hom}_{\mathbb{F}_p\text{-alg}}(\overline{A}, L')$ is both free and transitive.

Proof (Transitivity of the Galois Action). We already know from the previous lecture that the action is free. This implies that the cardinality of the group must be less than or equal to the cardinality of the set it acts upon, yielding $|G| \leq |S|$.
Since L is the splitting field of P over $\mathbb{Q}$, the order of the Galois group is $|G| = [L : \mathbb{Q}]$. By the structural properties of A and its reduction, we established that $[L : \mathbb{Q}] = \mathrm{rank}_{\mathbb{Z}}(A) = \dim_{\mathbb{F}_p}(\overline{A})$.
The Chinese Remainder Theorem gives a natural *surjective* map $\overline{A} \twoheadrightarrow \prod_{\mathfrak{m} \in \mathcal{M}} \overline{A}/\mathfrak{m}$. Because each quotient $\overline{A}/\mathfrak{m}$ is a splitting field of $\overline{P}$ over $\mathbb{F}_p$, we have $\overline{A}/\mathfrak{m} \cong L'$. Taking dimensions over $\mathbb{F}_p$, the surjectivity implies:

$$|G| = \dim_{\mathbb{F}_p}(\overline{A}) \geq \dim_{\mathbb{F}_p}\left(\prod_{\mathfrak{m} \in \mathcal{M}} \overline{A}/\mathfrak{m}\right) = |\mathcal{M}| \times [L' : \mathbb{F}_p] = |S|.$$

Combining the inequalities $|G| \leq |S|$ and $|G| \geq |S|$ forces $|G| = |S|$. Because a finite group G acts freely on a set S of the exact same cardinality, the action must necessarily be transitive.
(Note: As a byproduct of this equality, the dimensions are equal, proving that the surjective map from the Chinese Remainder Theorem is in fact an isomorphism $\overline{A} \cong \prod_{\mathfrak{m} \in \mathcal{M}} \overline{A}/\mathfrak{m}$.)

□

Definition (The Decomposition Group). Let $\mathfrak{m}_0 \in \mathcal{M}$ be a fixed maximal ideal of $\overline{A}$. The *decomposition group* at $\mathfrak{m}_0$ is defined as the isotropy subgroup (or stabilizer) of $\mathfrak{m}_0$ under the action of G on $\mathcal{M}$:

$$D_0 = \{g \in G \mid g \cdot \mathfrak{m}_0 = \mathfrak{m}_0\} \subseteq G.$$

It is crucial to note that the condition $g \cdot \mathfrak{m}_0 = \mathfrak{m}_0$ means that the ideal is preserved as a set; it does *not* mean that every individual element $x \in \mathfrak{m}_0$ is fixed by g.
For any element $g \in D_0$, the automorphism g restricts to the subring A and induces an automorphism on $\overline{A}$ that preserves the ideal $\mathfrak{m}_0$. Therefore, it descends to a well-defined $\mathbb{F}_p$-algebra automorphism on the quotient field:

$$\overline{g} : \overline{A}/\mathfrak{m}_0 \xrightarrow{\sim} \overline{A}/\mathfrak{m}_0.$$

This construction naturally defines a group homomorphism from the decomposition group to the Galois group of the residue field extension:

$$\Phi : D_0 \longrightarrow \mathrm{Gal}((\overline{A}/\mathfrak{m}_0)/\mathbb{F}_p).$$

$$\begin{array}{ccc} \overline{A} & \xrightarrow{\quad g \quad} & \overline{A} \\ {\scriptstyle q_0}\downarrow & & \downarrow{\scriptstyle q_0} \\ \overline{A}/\mathfrak{m}_0 & \xrightarrow[\quad \overline{g} \quad]{} & \overline{A}/\mathfrak{m}_0 \end{array}$$

Theorem (Epimorphism of the Decomposition Group). The natural group homomorphism $\Phi : D_0 \to \mathrm{Gal}((\overline{A}/\mathfrak{m}_0)/\mathbb{F}_p)$ is an epimorphism (i.e., it is surjective).

Proof (Epimorphism of the Decomposition Group). Let $\tau \in \mathrm{Gal}((\overline{A}/\mathfrak{m}_0)/\mathbb{F}_p)$ be an arbitrary automorphism of the residue field. Let $\varphi : \overline{A}/\mathfrak{m}_0 \xrightarrow{\sim} L'$ be a fixed $\mathbb{F}_p$-algebra isomorphism. We can compose φ with the canonical projection $q_0 : \overline{A} \twoheadrightarrow \overline{A}/\mathfrak{m}_0$ to obtain an element $s \in S$:

$$s = \varphi \circ q_0 \in \mathrm{Hom}_{\mathbb{F}_p\text{-alg}}(\overline{A}, L').$$

Now consider the composition $s' = \varphi \circ \tau \circ q_0$. Since s' is another homomorphism mapping $\overline{A}$ to L', we have $s' \in S$. Because G acts transitively on S, there exists a unique element $g \in G$ such that $s' = s \circ g$.
We first verify that $g \in D_0$. Since τ and φ are isomorphisms, the kernel of $s' = \varphi \circ \tau \circ q_0$ is precisely the kernel of q_0, which is $\mathfrak{m}_0$. On the other hand, $s' = s \circ g$ implies that $\ker(s') = g^{-1}(\ker(s)) = g^{-1}(\mathfrak{m}_0)$. Equating the two yields $g^{-1}(\mathfrak{m}_0) = \mathfrak{m}_0$, which means $g \cdot \mathfrak{m}_0 = \mathfrak{m}_0$. Thus, $g \in D_0$.
Because $g \in D_0$, it induces the automorphism $\overline{g}$ on $\overline{A}/\mathfrak{m}_0$, satisfying $q_0 \circ g = \overline{g} \circ q_0$. By substituting this into our relation $s' = s \circ g$, we obtain:

$$\varphi \circ \tau \circ q_0 = \varphi \circ q_0 \circ g = \varphi \circ \overline{g} \circ q_0.$$

Since φ is injective and q_0 is surjective, we can cancel them from both sides to conclude that $\overline{g} = \tau$. This exactly means $\Phi(g) = \tau$, proving that Φ is surjective.

□

Remark. In Algebraic Number Theory, the ring $A = \mathbb{Z}[\alpha]$ we are working with is often strictly contained in the full ring of integers $\mathcal{O}_L$ of the number field L ($A \subsetneq \mathcal{O}_L$). However, for primes p that do not divide the index $[\mathcal{O}_L : A]$, the quotient structure $\overline{A}$ behaves identically to $\mathcal{O}_L/p\mathcal{O}_L$. The subgroup $D_0 \subset G$ defined above is a specific realization of the classical decomposition group for prime ideals in $\mathcal{O}_L$.

3.3 Dedekind's Theorem on Cycle Types

To state Dedekind's fundamental result connecting the factorization of polynomials over finite fields to the cycle types of elements in the Galois group over $\mathbb{Q}$, we first provide a quick recollection of permutations and finite fields.

Definition (Type of a Permutation). Let $\sigma \in S_n$ be a permutation in the symmetric group. The cyclic subgroup $\langle \sigma \rangle \cong \mathbb{Z}/o(\sigma)\mathbb{Z}$ acts on the finite set $\{1, \ldots, n\}$. The disjoint union of the orbits of this action, say $\{1, \ldots, n\} = \coprod_{i=1}^{s} \mathcal{O}_i$, corresponds to the disjoint cycle decomposition of σ. The tuple of lengths of these orbits, $(|\mathcal{O}_1|, \ldots, |\mathcal{O}_s|)$, defines the *type* of the permutation σ (typically omitting fixed points of length 1).

Definition (Galois Extensions of $\mathbb{F}_p$). Let $\mathbb{F}_p \subset L'$ be a finite field extension of degree $d = [L' : \mathbb{F}_p]$. We know that L' is exactly the splitting field of the polynomial $X^{p^d} - X$ over $\mathbb{F}_p$. The

Galois group $\mathrm{Gal}(L'/\mathbb{F}_p)$ is a cyclic group of order d, and it is naturally generated by the *Frobenius automorphism*:

$$\begin{aligned} \mathrm{Frob}_p : L' &\longrightarrow L' \\ x &\longmapsto x^p. \end{aligned}$$

We can now apply this to our polynomial $P \in \mathbb{Z}[X]$. Suppose the reduction $\overline{P} \in \mathbb{F}_p[X]$ factors into irreducible monic polynomials over $\mathbb{F}_p$ as follows:

$$\overline{P} = \overline{\pi}_1 \times \cdots \times \overline{\pi}_s,$$

where each $\overline{\pi}_i$ has degree $d_i = \deg(\overline{\pi}_i)$. Since we assumed $\overline{P}$ is separable modulo p, these irreducible factors are distinct.
The set of roots of $\overline{P}$ in its splitting field L' decomposes as $\overline{R} = \coprod_{i=1}^{s} R_{\overline{\pi}_i}(L')$. The Frobenius automorphism generates $\mathrm{Gal}(L'/\mathbb{F}_p)$ and thus operates on the set of roots $\overline{R}$. For each irreducible factor $\overline{\pi}_i$, the Frobenius automorphism acts transitively on its d_i roots, effectively acting as a cycle of length d_i. Therefore, the cycle type of the Frobenius automorphism, when viewed as a permutation of the roots $\overline{R}$, is precisely $(d_1, d_2, \ldots, d_s)$.

Theorem (Dedekind's Theorem). Let $P \in \mathbb{Z}[X]$ be a monic polynomial with splitting field L over $\mathbb{Q}$, and let $G = \mathrm{Gal}(L/\mathbb{Q}) \hookrightarrow S_{R_P(L)}$ be its Galois group. Let p be a prime such that the reduction $\overline{P} \in \mathbb{F}_p[X]$ is separable, and suppose $\overline{P}$ factors into irreducible polynomials of degrees $d_1, \ldots, d_s$.
Then the Galois group G contains an element $\sigma \in G$ whose permutation type on the roots of P is exactly $(d_1, \ldots, d_s)$.

Proof (Dedekind's Theorem). By the epimorphism theorem proved earlier, there is a surjective group homomorphism $\Phi : D_0 \twoheadrightarrow \mathrm{Gal}((\overline{A}/m_0)/\mathbb{F}_p)$.
Since the target group $\mathrm{Gal}((\overline{A}/m_0)/\mathbb{F}_p)$ is generated by the Frobenius automorphism, we can choose an element $g \in D_0 \subset G$ such that $\Phi(g) = \mathrm{Frob}_p$.
Recall that we have a canonical projection map from the subring of roots to the residue field: $A \twoheadrightarrow \overline{A}/m_0 \cong L'$. When we restrict this projection to the set of complex roots $R_P(L) \subset A$, it maps them to the finite field roots $R_{\overline{P}}(L')$. Because $\overline{P}$ is separable, it possesses exactly $l = \deg(P)$ distinct roots in L'. Since this matches the exact number of complex roots in $R_P(L)$, the projection induces a bijection between the complex roots and the finite field roots.
Furthermore, by the very definition of the induced automorphism on the quotient, the action of $g \in D_0$ on A descends compatibly to the action of $\Phi(g)$ on L'. This makes our root bijection *equivariant*: the action of g on a complex root perfectly mirrors the action of $\Phi(g) = \mathrm{Frob}_p$ on the corresponding finite field root.
Because $\Phi(g)$ acts on $R_{\overline{P}}(L')$ with cycle type $(d_1, \ldots, d_s)$, the element $g \in G$ must act on the complex roots $R_P(L)$ with the exact same permutation type $(d_1, \ldots, d_s)$. This completes the proof.

□

End of Lec 03.

3.4 Application: A Polynomial with Galois Group S_n

As a powerful application of Dedekind's Theorem, we can explicitly construct a monic polynomial $P \in \mathbb{Z}[X]$ of degree n whose Galois group over $\mathbb{Q}$ is the full symmetric group S_n. By the insolvability of the quintic, for $n \geq 5$, the roots of such a polynomial are not expressible by nested radicals.

Construction of the Polynomial. Let $n \geq 2$. We choose three monic polynomials $Q_1, Q_2, Q_3 \in \mathbb{Z}[X]$ of degree n satisfying specific irreducibility conditions modulo the primes $2, 3$, and 5:

1. Modulo 2: The reduction $\overline{Q}_1 \in \mathbb{F}_2[X]$ is irreducible of degree n.
2. Modulo 3: The reduction $\overline{Q}_2 \in \mathbb{F}_3[X]$ factors as $\overline{Q}_2 = X \cdot \tilde{Q}_2$, where $\tilde{Q}_2 \in \mathbb{F}_3[X]$ is irreducible of degree $n-1$.
3. Modulo 5: The reduction $\overline{Q}_3 \in \mathbb{F}_5[X]$ factors as $\overline{Q}_3 = \tilde{Q}_3 \cdot \tilde{S}$, where $\tilde{Q}_3 \in \mathbb{F}_5[X]$ is irreducible of degree 2, and $\tilde{S} \in \mathbb{F}_5[X]$ is a product of irreducible polynomials of odd

degree.

We then define the global polynomial as the linear combination:

$$P(X) := -15Q_1(X) + 10Q_2(X) + 6Q_3(X) \in \mathbb{Z}[X].$$

Notice the carefully chosen coefficients. Modulo 2, $P \equiv Q_1$; modulo 3, $P \equiv Q_2$; and modulo 5, $P \equiv Q_3$. Let L be the splitting field of P over $\mathbb{Q}$ and $G = \mathrm{Gal}(L/\mathbb{Q}) \hookrightarrow S_n$ its Galois group acting on the n roots. By applying Dedekind's Theorem to the reductions of P, we deduce the cycle types present in G:

1. Reduction modulo 2 yields an irreducible polynomial of degree n. Thus, G contains an n-cycle, say γ_1.
2. Reduction modulo 3 yields factors of degree 1 and $n-1$. Thus, G contains an $(n-1)$-cycle, say γ_2.
3. Reduction modulo 5 yields a factor of degree 2 and factors of odd degrees. Thus, G contains a permutation of the form $\tau \circ \sigma$, where τ is a transposition (a 2-cycle) and σ is a disjoint product of cycles of odd lengths.

Let k be the order of the permutation σ. Since σ consists entirely of odd cycles, k is an odd integer. Therefore, taking the k-th power gives $(\tau \circ \sigma)^k = \tau^k \circ \sigma^k = \tau \circ \mathrm{id} = \tau$. This proves that G inherently contains a transposition τ.

3.4.1 Explicit Group Theory Calculations

To conclude the argument, we must explicitly demonstrate that the subgroup $G \subseteq S_n$ generated by these specific cycle types is the entire symmetric group S_n.

Lemma (Generation of S_n by Specific Cycles). Let $H \subseteq S_n$ be a subgroup such that it contains an n-cycle γ_1, an $(n-1)$-cycle γ_2, and a transposition τ. Then $H = S_n$.

Proof (Generation of S_n). By renumbering the elements of the set on which H acts, we may assume without loss of generality that the n-cycle is $\gamma_1 = (1, 2, \ldots, n)$.
First, we show that H contains a transposition involving the element 1. Let $\tau = (i, j)$ be the given transposition in H. Since γ_1 acts transitively on $\{1, \ldots, n\}$, there exists an integer k such that $\gamma_1^{-k}(i) = 1$. Conjugating τ by γ_1^{-k} yields a new transposition in H:

$$\gamma_1^{-k} \tau \gamma_1^{k} = (\gamma_1^{-k}(i), \gamma_1^{-k}(j)) = (1, m),$$

where $m = \gamma_1^{-k}(j) \neq 1$. Thus, H contains a transposition of the form $(1, m)$.
Second, we show that H contains an $(n-1)$-cycle that fixes 1. We are given that H contains some $(n-1)$-cycle, say σ, which fixes exactly one element $x \in \{1, \ldots, n\}$. Again, because γ_1 is transitive, there exists an integer p such that $\gamma_1^p(x) = 1$. Conjugating σ by γ_1^p produces another element in H:

$$\gamma_2 = \gamma_1^p \circ \sigma \circ \gamma_1^{-p}.$$

Since σ fixes x, γ_2 fixes $\gamma_1^p(x) = 1$. Furthermore, because conjugation preserves cycle structure, γ_2 is an $(n-1)$-cycle. Since it fixes 1, it must act transitively as a single cycle on the remaining set $\{2, \ldots, n\}$.
Finally, we conjugate our transposition $(1, m)$ by powers of γ_2:

$$\gamma_2^q \circ (1, m) \circ \gamma_2^{-q} = (1, \gamma_2^q(m)).$$

Because γ_2 acts transitively on $\{2, \ldots, n\}$, letting q vary allows $\gamma_2^q(m)$ to take every possible value in $\{2, \ldots, n\}$. This guarantees that H contains the entire set of transpositions $\{(1, 2), (1, 3), \ldots, (1, n)\}$. It is a standard fact in group theory that this specific set of transpositions generates the full symmetric group S_n. Applying this to our Galois group G, which contains the requisite n-cycle, $(n-1)$-cycle, and transposition, yields $G = S_n$.

□

4 Cyclotomic Fields and Cyclotomic Polynomials

4.1 Roots of Unity and the Galois Group

Let $n \geq 2$ be an integer. We consider the roots of the polynomial $X^n - 1 \in \mathbb{Q}[X]$. Over the complex numbers $\mathbb{C}$, the roots of this polynomial are the n-th roots of unity, given by $e^{2i\pi k/n}$ for $k = 0, \ldots, n-1$.

Definition (Cyclotomic Field and Primitive Roots). The *cyclotomic field* of order n is defined as the splitting field of $X^n - 1$ over $\mathbb{Q}$, denoted $K = \mathbb{Q}[\zeta_n]$, where $\zeta_n = e^{2i\pi/n}$. The roots of unity form a cyclic multiplicative group $\mu_n(\mathbb{C}) \subset \mathbb{C}^\times$ of order n, generated by ζ_n.
The generators of this cyclic group are called *primitive n-th roots of unity*. They are of the form ζ_n^k where $\gcd(k, n) = 1$. The set of primitive roots is denoted $\Pi_n \subset \mu_n(\mathbb{C})$, and its cardinality is given by Euler's totient function, $|\Pi_n| = \varphi(n)$.

Since K is the splitting field of the separable polynomial $X^n - 1$, the extension $K|\mathbb{Q}$ is a finite Galois extension. Let $G = \mathrm{Gal}(K/\mathbb{Q})$. Any automorphism $\sigma \in G$ is completely determined by its action on the generator ζ_n. Since σ must map ζ_n to another root of its minimal polynomial, $\sigma(\zeta_n)$ must be another primitive n-th root of unity. Hence, there exists an integer $c(\sigma)$, coprime to n, such that:

$$\sigma(\zeta_n) = \zeta_n^{c(\sigma)}.$$

This mapping $\sigma \mapsto c(\sigma) \pmod n$ defines a canonical, injective group homomorphism:

$$\mathrm{Gal}(K/\mathbb{Q}) \hookrightarrow (\mathbb{Z}/n\mathbb{Z})^\times.$$

As a direct consequence, the Galois group is abelian, and the cyclotomic extensions are abelian extensions of $\mathbb{Q}$.

Theorem (Galois Group of Cyclotomic Fields). The canonical monomorphism $\mathrm{Gal}(\mathbb{Q}[\zeta_n]/\mathbb{Q}) \hookrightarrow (\mathbb{Z}/n\mathbb{Z})^\times$ is an isomorphism. Consequently, $[K : \mathbb{Q}] = \varphi(n)$.

Proof (Galois Group of Cyclotomic Fields). Let $A = \mathbb{Z}[\zeta_n]$. To prove surjectivity, we examine the reduction of the extension modulo prime numbers. Let p be a prime number such that $p \nmid n$, meaning $\gcd(p, n) = 1$. Under this condition, the reduction of $X^n - 1$ modulo p remains separable. We apply our previous results on the decomposition group. The Frobenius automorphism in characteristic p, given by $\mathrm{Frob}_p(x) = x^p$, acts on the roots of unity in the residue field precisely by raising them to the p-th power. By the epimorphism property of the decomposition group, there exists an element $g \in \mathrm{Gal}(K/\mathbb{Q})$ that reduces to the Frobenius automorphism. This element g corresponds exactly to the class of p modulo n in $(\mathbb{Z}/n\mathbb{Z})^\times$.
It is a known fact from elementary number theory that the group of units $(\mathbb{Z}/n\mathbb{Z})^\times$ is generated by the residue classes of prime numbers p that do not divide n. Since each such class lies in the image of our monomorphism, the map is surjective, and therefore an isomorphism.

□

4.2 Cyclotomic Polynomials

Having established the structure of the cyclotomic field, we now define the polynomials whose roots are precisely the primitive roots of unity.

Definition (The Cyclotomic Polynomial). The n-th *cyclotomic polynomial*, denoted $\Phi_n(X)$, is defined as the product over all primitive n-th roots of unity:

$$\Phi_n(X) = \prod_{\zeta \in \Pi_n} (X - \zeta) \in \mathbb{C}[X].$$

By definition, it is a monic polynomial, and its degree is exactly $\deg(\Phi_n) = \varphi(n)$.

Every n-th root of unity is a primitive d-th root of unity for exactly one divisor d of n. This allows us to partition the set of all n-th roots of unity, yielding a fundamental factorization identity:

$$\mu_n(\mathbb{C}) = \coprod_{d|n} \Pi_d \quad \Longrightarrow \quad X^n - 1 = \prod_{d|n} \Phi_d(X).$$

Taking degrees on both sides recovers the classical identity $n = \sum_{d|n} \varphi(d)$.

Fact (Irreducibility over $\mathbb{Z}$). For all $n \geq 1$, the cyclotomic polynomial $\Phi_n(X)$ has integer coefficients, $\Phi_n(X) \in \mathbb{Z}[X]$, and is irreducible over $\mathbb{Q}$.

Proof (Integer Coefficients of Cyclotomic Polynomials). We prove that $\Phi_n(X) \in \mathbb{Z}[X]$ by strong induction on n. For $n = 1$, $\Phi_1(X) = X - 1 \in \mathbb{Z}[X]$. Assume the statement holds for all $d < n$.
We can rewrite the factorization identity as $(X^n - 1) = \Phi_n(X) \cdot \Psi_n(X)$, where $\Psi_n(X) = \prod_{d|n, d<n} \Phi_d(X)$. By the inductive hypothesis, $\Psi_n(X) \in \mathbb{Z}[X]$, and it is monic because it is a product of monic polynomials.
Since $(X^n - 1) \in \mathbb{Z}[X]$ and $\Psi_n(X) \in \mathbb{Z}[X]$ is monic, we can perform polynomial long division within the ring $\mathbb{Z}[X]$. The division algorithm yields $(X^n - 1) = Q(X)\Psi_n(X) + R(X)$ with $Q, R \in \mathbb{Z}[X]$ and $\deg(R) < \deg(\Psi_n)$. However, viewed in $\mathbb{C}[X]$, $X^n - 1$ is perfectly divisible by $\Psi_n(X)$, implying $R(X) = 0$ uniquely. Thus, $Q(X) = \Phi_n(X)$ must reside in $\mathbb{Z}[X]$.
The irreducibility over $\mathbb{Q}$ immediately follows from our previous theorem: since $\Phi_n(X) \in \mathbb{Z}[X]$ is the minimal polynomial of ζ_n over $\mathbb{Q}$, its degree $\varphi(n)$ matches the field extension degree $[K : \mathbb{Q}] = |(\mathbb{Z}/n\mathbb{Z})^\times| = \varphi(n)$. Therefore, it cannot factor further.

□

End of Lec 04.

5 Infinite Galois Theory and Profinite Groups

In this section, we transition from the study of finite field extensions to infinite ones. We consider an algebraic extension $K \subset L$ which is not necessarily finite.

5.1 Projective Limits and the Krull Topology

Definition (Infinite Galois Extension). An algebraic field extension $K \subset L$ is said to be *Galois* if it is both separable and normal.

If $K \subset L$ is a Galois extension, we can express L as the union of its finite Galois subextensions. Specifically, there exists a family of intermediate fields $(L_\alpha)_{\alpha \in I}$ such that for every α, the extension $K \subset L_\alpha$ is a finite Galois extension, and

$$L = \bigcup_{\alpha \in I} L_\alpha.$$

We define the Galois group of the infinite extension exactly as in the finite case: $\mathrm{Gal}(L/K) = \mathrm{Aut}_K(L)$. For any finite Galois subextension L_α, the restriction map provides a natural group homomorphism:

$$\begin{aligned} \mathrm{Gal}(L/K) &\longrightarrow \mathrm{Gal}(L_\alpha/K) \\ \sigma &\longmapsto \sigma|_{L_\alpha}. \end{aligned}$$

Because every element of $\mathrm{Gal}(L_\alpha/K)$ can be extended to an automorphism of L (since $L|K$ is normal), this restriction map is surjective.
Furthermore, if $L_\alpha \subset L_\beta$ is an inclusion of finite Galois subextensions over K, we have canonical surjective transition maps $\mathrm{Gal}(L_\beta/K) \twoheadrightarrow \mathrm{Gal}(L_\alpha/K)$. The Galois group of the infinite extension can then be understood as the projective limit of these finite Galois groups:

$$\mathrm{Gal}(L/K) = \varprojlim_{\alpha} \mathrm{Gal}(L_\alpha/K).$$

Because each $\mathrm{Gal}(L_\alpha/K)$ is a finite group equipped with the discrete topology, their projective limit naturally inherits the structure of a *profinite group*. A profinite group is a topological group that is totally disconnected, compact, and Hausdorff.

Theorem (Finiteness of Galois Groups). Let $K \subset L$ be a Galois extension. The topological Galois group $\mathrm{Gal}(L/K)$ is a finite group if and only if the extension $L|K$ is a finite extension.

As a simple example of a non-trivial finite Galois group structure over an infinite field, if $K = K_s$ is a separably closed field contained in an algebraically closed field L, then the absolute Galois group $\mathrm{Gal}(L/K)$ is either trivial or isomorphic to $\mathbb{Z}/2\mathbb{Z}$.

5.2 The Absolute Galois Group of Finite Fields

We now explore the Galois theory of finite fields by considering the algebraic closure $\overline{\mathbb{F}_p}$ of the finite field $\mathbb{F}_p$. Is the extension $\mathbb{F}_p \subset \overline{\mathbb{F}_p}$ Galois?
For any integer $n \geq 1$, let K_n be the splitting field of the polynomial $X^{p^n} - X$ over $\mathbb{F}_p$. The roots of this polynomial form exactly the field $\mathbb{F}_{p^n}$, and thus $K_n = \mathbb{F}_{p^n}$. The extension $\mathbb{F}_p \subset K_n$ is a finite Galois extension. The Galois group is cyclic and generated by the Frobenius automorphism:

$$\mathrm{Frob}_p : x \longmapsto x^p.$$

Thus, $\mathrm{Gal}(K_n/\mathbb{F}_p) = \langle \mathrm{Frob}_p \rangle \cong \mathbb{Z}/n\mathbb{Z}$.
Observe that $K_n \subset K_m$ if and only if $n \mid m$. The algebraic closure $\overline{\mathbb{F}_p}$ is the union of all these finite fields K_n. Therefore, the absolute Galois group of $\mathbb{F}_p$ is the projective limit of these finite cyclic groups, ordered by divisibility:

$$\mathrm{Gal}(\overline{\mathbb{F}_p}/\mathbb{F}_p) = \varprojlim_n (\mathbb{Z}/n\mathbb{Z}) =: \widehat{\mathbb{Z}}.$$

The group $\widehat{\mathbb{Z}}$ is the profinite completion of the integers.
Alternatively, one can view the finite field K_n as the splitting field of $X^{p^n-1}-1$. In this perspective, K_n is generated by roots of unity over $\mathbb{F}_p$, linking the study of finite fields to cyclotomic extensions.

5.3 Cyclotomic Extensions of $\mathbb{Q}$ and Kronecker-Weber

Recall from the finite case that the cyclotomic extension $\mathbb{Q} \subset K_n = \mathbb{Q}[\zeta_n]$ is Galois, where ζ_n is a primitive n-th root of unity, and its Galois group is isomorphic to $(\mathbb{Z}/n\mathbb{Z})^\times$. Note that $K_n \subset K_m$ whenever $n \mid m$.
We define the maximal cyclotomic extension of $\mathbb{Q}$ as the field generated by all roots of unity in $\mathbb{C}$:

$$K_\infty = \mathbb{Q}[\mu_\infty(\mathbb{C})] = \bigcup_{n=1}^{\infty} K_n \subset \mathbb{C}.$$

By taking the projective limit, we find the Galois group of this infinite extension:

$$\mathrm{Gal}(K_\infty/\mathbb{Q}) = \varprojlim_n (\mathbb{Z}/n\mathbb{Z})^\times = \widehat{\mathbb{Z}}^\times.$$

Because the Galois group of any finite cyclotomic extension is abelian, the Galois group of K_∞ over $\mathbb{Q}$ is an infinite abelian profinite group. This naturally leads to one of the crowning achievements of classical algebraic number theory.

Theorem (Kronecker-Weber Theorem). Let K be a finite Galois extension of $\mathbb{Q}$. If the Galois group $\mathrm{Gal}(K/\mathbb{Q})$ is abelian, then there exists an integer $N \gg 0$ such that $K \subset K_N = \mathbb{Q}[\mu_N(\mathbb{C})]$.

Consequently, the maximal abelian extension of $\mathbb{Q}$ (denoted $\mathbb{Q}^{\mathrm{ab}}$) is exactly the field K_∞ generated by all roots of unity. Thus, $\mathrm{Gal}(\mathbb{Q}^{\mathrm{ab}}/\mathbb{Q}) \cong \widehat{\mathbb{Z}}^\times$.
To conclude our discussion on Galois theory, we state an open conjecture regarding the full absolute Galois group of $\mathbb{Q}$:

Conjecture (Absolute Galois Group of $\mathbb{Q}$). Is the absolute Galois group $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ a free profinite group? The deep structure of this group and its representations remain a central object of study, heavily intertwined with the Langlands Program.

Part II
Introduction to Commutative Algebra

6 Commutative Rings

6.1 Rings, Homomorphisms, and the Multiplicative Group

Definition (Commutative Ring). A *commutative ring* is a triple $(R, +, \times)$ consisting of a set R equipped with two binary operations, addition $(+)$ and multiplication $(\times)$, satisfying the following axioms:

1. $(R, +)$ is an abelian group. In particular, there exists a neutral element for addition denoted $0_R = 0$, and every element $x \in R$ has an additive inverse denoted $-x$.
2. $(R, \times)$ is a commutative monoid. That is, multiplication is associative, commutative, and there exists a multiplicative identity denoted $1_R = 1$.
3. *Compatibility (Distributivity):* Multiplication distributes over addition. For all $a, b, c \in R$, we have $(a + b) \times c = a \times c + b \times c$.

Basic algebraic manipulation directly from the axioms reveals that in any commutative ring R:

$$0 \times x = x \times 0 = 0,$$
$$-(x \times y) = (-x) \times y = x \times (-y).$$

It is a curious observation that $1_R = 0_R$ if and only if $R = \{0\}$. In this case, R is called the *trivial ring*. Unless otherwise stated, we often implicitly assume $R \neq 0$.
Standard examples of commutative rings include the ring of integers $\mathbb{Z}$, any field K (such as $\mathbb{Q}, \mathbb{R}$, or $\mathbb{C}$), and polynomial rings $R[X_1, \ldots, X_n]$ over any commutative ring R. Furthermore, quotients of polynomial rings, such as $R[X_1, \ldots, X_n]/(P_1, \ldots, P_m)$, inherit the structure of a commutative ring.

Definition (Invertible Elements and the Multiplicative Group). Let R be a commutative ring. An element $x \in R$ is called *invertible* (or a *unit*) if there exists an element $y \in R$ such that $x \times y = 1$.
The set of all invertible elements is denoted by $R^\times$:

$$R^\times = \{x \in R \mid \exists y \in R,\ x \times y = 1\}.$$

The set $R^\times$, equipped with the ring's multiplication operation $\times$ and the identity element 1, forms a group, called the *multiplicative group of* R. If R is the trivial ring, $R^\times = \{0\}$.

Definition (Ring Homomorphism). Let A and B be two commutative rings. A map $\varphi : A \to B$ is called a *ring homomorphism* if it respects both the additive and multiplicative structures, as well as the identities:

1. $\varphi(a + b) = \varphi(a) + \varphi(b)$ for all $a, b \in A$.
2. $\varphi(a \times b) = \varphi(a) \times \varphi(b)$ for all $a, b \in A$.
3. $\varphi(1_A) = 1_B$.

An immediate consequence of the additive group homomorphism property is that $\varphi(0_A) = 0_B$ and $\varphi(-x) = -\varphi(x)$.
Examples of ring homomorphisms include the natural inclusions $\mathbb{Z} \subset \mathbb{Q} \subset \mathbb{R} \subset \mathbb{C}$. Moreover, if $K \subset L$ is a field extension and $x \in L$ is a fixed element, the *evaluation map* $K[X] \to L$ defined by $P \mapsto P(x)$ is a ring homomorphism.

Lemma (Homomorphisms Preserve Units). Let $\varphi : A \to B$ be a ring homomorphism. Then φ maps units of A to units of B, i.e., $\varphi(A^\times) \subset B^\times$. Furthermore, the restriction $\varphi|_{A^\times} : A^\times \to B^\times$ is a group homomorphism.

Proof (Homomorphisms Preserve Units). Let $x \in A^\times$. By definition, there exists $y \in A$ such that $xy = 1_A$. Applying the ring homomorphism φ yields:

$$\varphi(x)\varphi(y) = \varphi(xy) = \varphi(1_A) = 1_B.$$

Thus, $\varphi(x)$ possesses a multiplicative inverse in B (namely $\varphi(y)$), meaning $\varphi(x) \in B^\times$. The fact that the restriction is a group homomorphism trivially follows from the multiplicative property of φ.

□

From a categorical perspective, let **comRing** be the category whose objects are commutative rings and whose morphisms are ring homomorphisms. We have two natural functors mapping from **comRing** to the category of abelian groups:

1. The additive functor: $R \mapsto (R, +)$.
2. The multiplicative unit functor: $R \mapsto (R^\times, \times)$.

6.2 Ideals and Quotient Rings

6.2.1 Definition and Properties of Ideals

Definition (Ideal). Let R be a commutative ring. A subset $I \subset R$ is called an *ideal* of R if it satisfies the following two conditions:

1. I is an additive subgroup of $(R, +)$. (Explicitly: $0 \in I$, and for all $x, y \in I$, we have $-x \in I$ and $x + y \in I$).
2. I absorbs multiplication by elements of R. That is, for all $x \in I$ and $r \in R$, the product $r \cdot x \in I$.

An ideal is said to be *proper* if $I \neq R$. We denote the set of all proper ideals of R by $\mathcal{I}(R)$. The set $\mathcal{I}(R)$, ordered by set inclusion ($\subset$), forms a partially ordered set (poset).
When an ideal I is given, we can form the quotient group $(R/I, +)$ from the underlying additive abelian groups.

Fact (Quotient Ring Structure). Let $I \subset R$ be an ideal. The additive quotient group R/I canonically inherits a ring structure. The multiplication is defined via representatives: for any classes $\overline{x}, \overline{y} \in R/I$, we define $\overline{x} \cdot \overline{y} = \overline{x \cdot y}$.
Under this structure, R/I is called the *quotient ring*, and the canonical projection map $\pi : R \twoheadrightarrow R/I$ defined by $x \mapsto \overline{x}$ is a surjective ring homomorphism.
Furthermore, the quotient ring satisfies a universal property: any ring homomorphism $\varphi : R \to B$ whose kernel contains I factors uniquely through R/I.

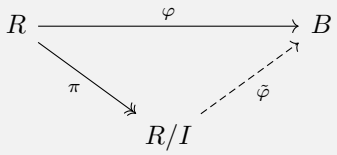

A classical example is the ideal $n\mathbb{Z}$ in the ring $\mathbb{Z}$, which yields the quotient ring $\mathbb{Z}/n\mathbb{Z}$.

6.2.2 Maximal and Prime Ideals

Definition (Maximal Ideal). An ideal $\mathfrak{m} \subset R$ is called a *maximal ideal* if it is a proper ideal ($\mathfrak{m} \neq R$) and it is maximal with respect to set inclusion among all proper ideals. Equivalently, $\mathfrak{m}$ is maximal if for any ideal $I \subset R$ such that $\mathfrak{m} \subset I \subset R$, either $I = \mathfrak{m}$ or $I = R$.

The existence of maximal ideals is guaranteed in any non-trivial ring by an application of Zorn's Lemma.

Lemma (Existence of Maximal Ideals). If $R \neq 0$, then the set of proper ideals $\mathcal{I}(R)$ is non-empty and possesses maximal elements. More precisely, for any proper ideal $I \in \mathcal{I}(R)$, there exists a maximal ideal $\mathfrak{m}$ such that $I \subset \mathfrak{m}$.

Proof (Existence of Maximal Ideals). We apply Zorn's Lemma to the poset $(\mathcal{I}(R), \subset)$. Let $\mathcal{C}$ be a totally ordered chain of proper ideals in $\mathcal{I}(R)$. We define $J = \bigcup_{I \in \mathcal{C}} I$. It is straightforward to verify that J is an ideal. Furthermore, J is proper: if $1 \in J$, then $1 \in I$ for some $I \in \mathcal{C}$, which would imply $I = R$, contradicting that the chain consists of proper ideals. Thus $J \in \mathcal{I}(R)$, and J provides an upper bound for the chain $\mathcal{C}$. By Zorn's Lemma, $\mathcal{I}(R)$ contains a maximal element. □

Maximal ideals have a beautiful characterization in terms of their quotient rings.

Lemma (Characterization of Maximal Ideals). Let $I \subset R$ be a proper ideal. Then I is a maximal ideal if and only if the quotient ring R/I is a field.

Proof (Characterization of Maximal Ideals). There is a one-to-one inclusion-preserving correspondence between the ideals of the quotient ring R/I and the ideals of R that contain I. Specifically, an ideal J such that $I \subset J \subset R$ corresponds perfectly to the ideal $J/I \subset R/I$.
The ideal I is maximal if and only if there are no intermediate proper ideals strictly containing I. This is equivalent to saying that the quotient ring R/I has exactly two ideals: the zero ideal $\{0\}$ and the entire ring R/I. A commutative ring has exactly two ideals if and only if every non-zero element is invertible, which is the definition of a field. □

For example, in a Principal Ideal Domain (P.I.D.) R, any ideal is generated by a single element, say $I = (x)$. The ideal (x) is maximal if and only if x is an irreducible element of R.
We can relax the condition that the quotient be a field to the condition that the quotient merely be an integral domain.

Definition (Prime Ideal). An ideal $\mathfrak{p} \subset R$ is said to be *prime* if it is a proper ideal ($\mathfrak{p} \neq R$) and the quotient ring $R/\mathfrak{p}$ is an integral domain.
Equivalently, $\mathfrak{p}$ is a prime ideal if for all $x, y \in R$, the condition $x \cdot y \in \mathfrak{p}$ implies that either $x \in \mathfrak{p}$ or $y \in \mathfrak{p}$.

Since every field is an integral domain, it immediately follows that *every maximal ideal is a prime ideal*. However, the converse is not true in general.
Consider the polynomial ring $k[X, Y]$ over a field k. We have a chain of prime ideals:

$$0 \subsetneq (X) \subsetneq (X, Y) \subset k[X, Y].$$

The ideal (X, Y) is maximal, as the quotient $k[X, Y]/(X, Y) \cong k$ is a field. The ideal (X) is prime, but not maximal, because the quotient $k[X, Y]/(X) \cong k[Y]$ is an integral domain but not a field. (Outlook: The length of chains of prime ideals will be deeply related to dimension theory later in the course).

7 The Prime Spectrum and Localization

7.1 Noetherian Rings and Finitely Generated Ideals

Definition (Finitely Generated Ideal). An ideal $I \subset R$ is said to be *finitely generated* (or of finite type) if there exists a finite family of elements $(x_1, \ldots, x_r) \in I^r$ such that every element $x \in I$ can be written as an R-linear combination of the x_i:

$$x = \sum_{i=1}^{r} \lambda_i x_i \quad \text{for some } (\lambda_1, \ldots, \lambda_r) \in R^r.$$

In a P.I.D., every ideal is generated by a single element, meaning every ideal is trivially finitely generated. We generalize this property to define a highly important class of rings.

Definition (Noetherian Ring). A commutative ring R is said to be *Noetherian* if every ideal $I \subset R$ is finitely generated.

By definition, any P.I.D. (such as $\mathbb{Z}$ or $K[X]$) is a Noetherian ring. A fundamental result, known as the *Hilbert Basis Theorem* (which we will prove later), states that if R is Noetherian, then the polynomial ring $R[X]$ is also Noetherian.

7.1.1 The Prime Spectrum

Definition (The Spectrum). Let R be a commutative ring. We denote by $\mathrm{Spec}(R)$ the set of all prime ideals of R. We occasionally use the notation $\mathrm{Spm}(R)$ or $\mathrm{Max}(R)$ to denote the subset of $\mathrm{Spec}(R)$ consisting only of maximal ideals.

If $\mathfrak{p} \in \mathrm{Spec}(R)$, we know that the quotient $R/\mathfrak{p}$ is an integral domain. We can therefore form its field of fractions, which we denote by $\kappa(\mathfrak{p}) = \mathrm{Frac}(R/\mathfrak{p})$. This field $\kappa(\mathfrak{p})$ is called the *residue field* at $\mathfrak{p}$.
For example, let $R = \mathbb{Z}$. If $\mathfrak{p} \in \mathrm{Spec}(\mathbb{Z})$ is generated by a prime number $p \neq 0$, the quotient $\mathbb{Z}/(\mathfrak{p}) = \mathbb{F}_p$ is already a field, so $\kappa(\mathfrak{p}) = \mathbb{F}_p$. For the zero ideal $0 \in \mathrm{Spec}(\mathbb{Z})$, the quotient is $\mathbb{Z}/0 \cong \mathbb{Z}$, and its residue field is $\kappa(0) = \mathbb{Q}$.
A ring homomorphism induces a natural map on spectra, flowing in the opposite direction.

Lemma (Induced Map on Spectra). Let $\varphi : A \to B$ be a ring homomorphism. If $\mathfrak{q} \in \mathrm{Spec}(B)$ is a prime ideal, then its preimage $\varphi^{-1}(\mathfrak{q})$ is a prime ideal in A, i.e., $\varphi^{-1}(\mathfrak{q}) \in \mathrm{Spec}(A)$. This naturally defines an induced map $\mathrm{Spec}(\varphi) : \mathrm{Spec}(B) \to \mathrm{Spec}(A)$.

Proof (Induced Map on Spectra). The homomorphism φ induces an injective ring homomorphism from the quotient $A/\varphi^{-1}(\mathfrak{q})$ to $B/\mathfrak{q}$. Since $\mathfrak{q}$ is prime in B, $B/\mathfrak{q}$ is an integral domain. A subring of an integral domain (in this case, the image of $A/\varphi^{-1}(\mathfrak{q})$) must also be an integral domain. Therefore, $A/\varphi^{-1}(\mathfrak{q})$ is an integral domain, which proves that $\varphi^{-1}(\mathfrak{q})$ is a prime ideal. □

7.2 The Zariski Topology

In modern algebraic geometry, an element $f \in R$ is viewed conceptually as a "regular function" defined on the topological space $\mathrm{Spec}(R)$. The "value" of the function f at a point $\mathfrak{p} \in \mathrm{Spec}(R)$ is precisely its image under the composition $R \twoheadrightarrow R/\mathfrak{p} \hookrightarrow \kappa(\mathfrak{p})$.
We endow $\mathrm{Spec}(R)$ with a natural topology.

Definition (The Zariski Topology). Let A be a commutative ring. For any element $f \in A$, we define the set:

$$V(f) = \{\mathfrak{p} \in \mathrm{Spec}(A) \mid f \in \mathfrak{p}\}.$$

We declare the sets $V(f)$ to be the fundamental *closed sets* of $\mathrm{Spec}(A)$. Intersections of these sets form the closed sets of the *Zariski topology.*
Taking the complement, we define the standard open sets:

$$D(f) = \mathrm{Spec}(A) \setminus V(f) = \{\mathfrak{p} \in \mathrm{Spec}(A) \mid f \notin \mathfrak{p}\}.$$

A subset $U \subset \mathrm{Spec}(A)$ is open in the Zariski topology if and only if it can be written as a union of sets of the form $D(f_\alpha)$. In other words, the family $\{D(f)\}_{f \in A}$ forms a *basis* for the Zariski topology on $\mathrm{Spec}(A)$.

End of Lec 05.

7.2.1 Properties, Closures, and Specializations

We recall and expand upon the properties of the Zariski topology introduced previously. Let R be a commutative ring. For any element $f \in R$, we defined the set $V(f) = \{\mathfrak{p} \in \mathrm{Spec}(R) \mid \overline{f} = 0 \text{ in } R/\mathfrak{p}\}$, which constitutes a closed set in $\mathrm{Spec}(R)$. The complement of this set, $D(f) = \mathrm{Spec}(R) \setminus V(f) = \{\mathfrak{p} \in \mathrm{Spec}(R) \mid f \notin \mathfrak{p}\}$, is an open set.
The collection of open sets $\{D(f)\}_{f \in R}$ forms a basis for the Zariski topology. Consequently, any open set $\Omega \subset \mathrm{Spec}(R)$ can be written as a union $\Omega = \bigcup_i D(f_i)$ for some family of elements $f_i \in R$. The entire space and the empty set are given by $\mathrm{Spec}(R) = D(1)$ and $\emptyset = D(0)$, respectively.

Fact (Closed Sets and Ideals). Let $(f_i)_{i \in I}$ be a family of elements in R, and let J be the ideal generated by this family. The intersection of the corresponding closed sets yields a closed set defined by the ideal:

$$\bigcap_{i \in I} V(f_i) = \{\mathfrak{p} \in \mathrm{Spec}(R) \mid \forall i, f_i \in \mathfrak{p}\} = V(J).$$

Therefore, for any ideal $J \subset R$, $V(J)$ is a closed subset of $\mathrm{Spec}(R)$.

As a brief outlook, we note that the natural projection map $R \twoheadrightarrow R/J$ induces an injective continuous map on the spectra, $\mathrm{Spec}(R/J) \hookrightarrow \mathrm{Spec}(R)$, whose image is precisely the closed set $V(J)$. We now examine the closure of points within the Zariski topology, which reveals the non-Hausdorff nature of the space.

Lemma (Closure of a Point). Let $\mathfrak{p} \in \mathrm{Spec}(R)$ be a prime ideal, viewed as a point in the topological space. Its topological closure is given by the closed set it defines:

$$\overline{\{\mathfrak{p}\}} = V(\mathfrak{p}).$$

Proof (Closure of a Point). By definition, $\mathfrak{p} \in V(\mathfrak{p})$, and $V(\mathfrak{p})$ is a closed set. To show it is the closure, we must demonstrate it is the smallest closed set containing $\mathfrak{p}$. Let $V(I)$ be any closed subset containing $\mathfrak{p}$, where I is an ideal of R. Since $\mathfrak{p} \in V(I)$, it follows that $I \subset \mathfrak{p}$.
Now, take any other point $\mathfrak{q} \in V(\mathfrak{p})$. This means $\mathfrak{p} \subset \mathfrak{q}$. By transitivity of inclusion, $I \subset \mathfrak{p} \subset \mathfrak{q}$, which implies $\mathfrak{q} \in V(I)$. Hence, $V(\mathfrak{p}) \subset V(I)$, confirming that $V(\mathfrak{p})$ is indeed the closure of $\{\mathfrak{p}\}$. □

Definition (Specialization and Generic Points). Let $\mathfrak{p}, \mathfrak{q} \in \mathrm{Spec}(R)$ be prime ideals. If $\mathfrak{p} \in \overline{\{\mathfrak{q}\}}$, which is equivalent to $\mathfrak{p} \in V(\mathfrak{q})$ or $\mathfrak{q} \subset \mathfrak{p}$, we say that $\mathfrak{p}$ is a *specialization* of $\mathfrak{q}$.
If R is an integral domain, the zero ideal (0) is a prime ideal, so $(0) \in \mathrm{Spec}(R)$. The closure of this point is $\overline{\{(0)\}} = V((0)) = \mathrm{Spec}(R)$. A point whose closure is the entire space is called a *generic point*.

An immediate consequence of the closure lemma pertains to maximal ideals. If $\mathfrak{m} \subset R$ is a maximal ideal, then for any $\mathfrak{p} \in \mathrm{Spec}(R)$, the condition $\mathfrak{p} \in \overline{\{\mathfrak{m}\}}$ means $\mathfrak{m} \subset \mathfrak{p}$. Because $\mathfrak{m}$ is maximal and $\mathfrak{p}$ is a proper ideal, this forces $\mathfrak{p} = \mathfrak{m}$. Thus, points corresponding to maximal ideals are closed points in the Zariski topology.

Example ($\mathrm{Spec}(\mathbb{Z})$). Let $R = \mathbb{Z}$. The spectrum is $\mathrm{Spec}(\mathbb{Z}) = \{0\} \cup \{(p) \mid p \text{ is prime}\}$. The zero ideal is the generic point. For any non-zero element $f \in \mathbb{Z}$, the closed set $V(f)$ consists of the prime ideals generated by the prime factors of f. Thus, $V(f) = \{(p_1), \ldots, (p_s)\}$, which is a finite set of closed points. Consequently, any open set in $\mathrm{Spec}(\mathbb{Z})$ is the complement of a finite set of non-zero prime ideals (or is the empty set).

7.3 Local Rings and Localization

Definition (Local Ring and Residue Field). A commutative ring R is said to be a *local ring* if it possesses exactly one maximal ideal, $\mathfrak{m} \in \mathrm{Spec}(R)$.
The quotient ring $\kappa = R/\mathfrak{m}$ is necessarily a field and is called the *residue field* of R.

Examples.

1. Any commutative field K is a local ring. Its unique maximal ideal is the zero ideal (0), and its spectrum consists of a single point, $\mathrm{Spec}(K) = \{0\}$.
2. Let p be a prime number. Consider the subring of the rational numbers defined by $\mathbb{Z}_{(p)} = \{\frac{a}{b} \mid a \in \mathbb{Z}, b \in \mathbb{Z} \setminus \{0\}, \gcd(p, b) = 1\}$. This is a local ring. Its unique maximal ideal is generated by p, specifically $\mathfrak{m}_p = p \cdot \mathbb{Z}_{(p)} = \{\frac{a}{b} \mid a \in \mathbb{Z}, p \mid a, b \in \mathbb{Z} \setminus \{0\}, \gcd(p, b) = 1\}$. The only other prime ideal in this ring is the zero ideal.

7.3.1 Characterization of Local Rings

Local rings can be completely characterized by the behavior of their non-invertible elements.

Lemma (Characterization of Local Rings). Let R be a commutative ring and let $\mathfrak{m} \subset R$ be a proper ideal. Then R is a local ring with maximal ideal $\mathfrak{m}$ if and only if $R \setminus R^\times = \mathfrak{m}$, where $R^\times$ denotes the group of invertible elements in R.

Proof (Characterization of Local Rings). Forward Implication ($\implies$): Assume $R \setminus R^\times = \mathfrak{m}$. Equivalently, the units are precisely the elements outside the ideal, $R^\times = R \setminus \mathfrak{m}$. We first verify that $\mathfrak{m}$ is a maximal ideal. Let I be an ideal such that $\mathfrak{m} \subset I \subset R$. If $I \neq \mathfrak{m}$, there exists an element $x \in I \setminus \mathfrak{m}$. Because $x \notin \mathfrak{m}$, it must be that $x \in R^\times$. Since the ideal I contains an invertible element, it must be the entire ring, $I = R$. Thus, $\mathfrak{m}$ is maximal. It is the unique maximal ideal because any other proper ideal must consist entirely of non-invertible elements, and therefore must be contained within $\mathfrak{m}$.

Reverse Implication ($\impliedby$): Assume R is a local ring with a unique maximal ideal $\mathfrak{m}$. Let $f \in R \setminus \mathfrak{m}$. Consider the principal ideal (f) generated by f. If $(f) \neq R$, then it is a proper ideal, and by Zorn's Lemma, it must be contained in some maximal ideal. However, R has only one maximal ideal, $\mathfrak{m}$, which implies $(f) \subset \mathfrak{m}$. This leads to $f \in \mathfrak{m}$, a direct contradiction to our choice of f. Thus, we must have $(f) = R$, which means f is invertible. Therefore, every element outside $\mathfrak{m}$ is a unit, yielding $R \setminus R^\times = \mathfrak{m}$.

□

7.3.2 Construction and Universal Property

We now detail the formal construction of localization, an algebraic procedure that forces a specified set of elements in a ring to become invertible, thereby generalizing the construction of the field of fractions for an integral domain.

Definition (Multiplicative Subset). Let R be a commutative ring. A subset $S \subset R$ is called a *multiplicative subset* if it satisfies two properties:

1. The multiplicative identity is in the set: $1 \in S$.
2. The set is closed under multiplication: for all $s, t \in S$, the product $s \cdot t \in S$.

In other words, S is a submonoid of the multiplicative monoid $(R, \times)$.

For an integral domain R, the set of non-zero elements $R \setminus \{0\}$ naturally forms a multiplicative subset. Another pervasive example arises from prime ideals: if $\mathfrak{p} \in \operatorname{Spec}(R)$, the complement $S_\mathfrak{p} := R \setminus \mathfrak{p}$ is a multiplicative subset precisely because $\mathfrak{p}$ is prime (if $s, t \notin \mathfrak{p}$, then $s \cdot t \notin \mathfrak{p}$).
Let $S \subset R$ be a multiplicative subset. We consider the Cartesian product $S \times R$. We define a relation on this set, declaring $(s, x) \sim (t, y)$ if there exists an element $u \in S$ such that

$$u(tx - sy) = 0 \quad \iff \quad utx = usy.$$

If R is an integral domain and $0 \notin S$, the relation simplifies to the familiar $tx = sy$.

Lemma (Equivalence Relation). The relation $\sim$ is an equivalence relation on $S \times R$.

Proof (Equivalence Relation for Localization). Reflexivity ($(s, x) \sim (s, x)$ via $u = 1$) and symmetry ($(s, x) \sim (t, y) \implies (t, y) \sim (s, x)$) are straightforward. We verify transitivity in detail. Suppose $(s_1, x_1) \sim (s_2, x_2)$ and $(s_2, x_2) \sim (s_3, x_3)$. By definition, there exist $u, v \in S$ such that:

$$us_2x_1 = us_1x_2$$
$$vs_3x_2 = vs_2x_3$$

To relate (s_1, x_1) to (s_3, x_3), we multiply the first equation by vs_3 and the second by us_1:

$$uvs_2s_3x_1 = uvs_1s_3x_2$$
$$uvs_1s_3x_2 = uvs_1s_2x_3$$

Equating the two expressions yields $uvs_2s_3x_1 = uvs_1s_2x_3$, which can be factored as:

$$(uvs_2)s_3x_1 = (uvs_2)s_1x_3$$

Let $w = uvs_2$. Because S is closed under multiplication, $w \in S$. The equation $ws_3x_1 = ws_1x_3$ perfectly satisfies the definition of the relation, proving $(s_1, x_1) \sim (s_3, x_3)$.

□

We denote the equivalence class of (s, x) as a formal fraction $\frac{x}{s}$. The set of all equivalence classes is denoted $S^{-1}R$. This set inherits a unique ring structure where addition and multiplication are defined by the standard arithmetic of fractions:

$$\frac{x}{s} + \frac{y}{u} = \frac{xu + ys}{su}$$
$$\frac{x}{s} \cdot \frac{y}{u} = \frac{xy}{su}$$

The zero element is $\frac{0}{1}$ (or equivalently $\frac{0}{s}$ for any $s \in S$), and the identity element is $\frac{1}{1}$. There is a canonical ring homomorphism $\varphi_{R,S} : R \to S^{-1}R$ given by $x \mapsto \frac{x}{1}$.
The localization $S^{-1}R$ satisfies a definitive universal property, acting as the "most general" ring where elements of S become invertible.

Proposition (Universal Property of Localization). Let $S \subset R$ be a multiplicative subset, and let $\varphi : R \to S^{-1}R$ be the canonical homomorphism. Let A be any commutative ring, and let $\psi : R \to A$ be a ring homomorphism such that $\psi(S) \subset A^\times$ (meaning ψ maps every element of S to a unit in A).
Then there exists a unique ring homomorphism $\Theta : S^{-1}R \to A$ such that $\Theta \circ \varphi = \psi$.

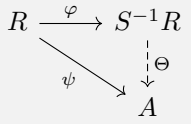

Proof (Universal Property of Localization). Uniqueness: Suppose such a homomorphism Θ exists. For any fraction $\frac{x}{s} \in S^{-1}R$, we can write $\frac{x}{s} = \frac{x}{1} \cdot (\frac{s}{1})^{-1} = \varphi(x) \cdot \varphi(s)^{-1}$. Applying Θ yields:

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

This forces Θ to be completely determined by ψ, ensuring uniqueness.

Existence: We use the forced formula to define Θ. Let $\Theta(s, x) := \psi(s)^{-1} \cdot \psi(x) \in A$. We must verify this is well-defined on equivalence classes. Suppose $(s, x) \sim (t, y)$. Then there exists $u \in S$ such that $usy = utx$. Applying ψ to both sides gives:

$$\psi(u)\psi(s)\psi(y) = \psi(u)\psi(t)\psi(x).$$

Because $u \in S$, the assumption ensures $\psi(u) \in A^\times$. Multiplying by $\psi(u)^{-1}$ cancels it out:

$$\psi(s)\psi(y) = \psi(t)\psi(x).$$

Similarly, since $s, t \in S$, $\psi(s)$ and $\psi(t)$ are invertible. Multiplying by $\psi(s)^{-1}\psi(t)^{-1}$ yields:

$$\psi(t)^{-1}\psi(y) = \psi(s)^{-1}\psi(x).$$

This confirms that $\Theta(t, y) = \Theta(s, x)$, so the map is well-defined. It routinely follows that Θ is a ring homomorphism satisfying the required diagram.

□

Remarks and Examples.

1. If $0 \in S$, then for any $\frac{x}{s} \in S^{-1}R$, the relation $0 \cdot (1 \cdot x - s \cdot 0) = 0$ implies $(s, x) \sim (1, 0)$, so every element is equivalent to zero. Thus, $S^{-1}R$ is the trivial ring 0.

2. If R is an integral domain and $S \subset R \setminus \{0\}$, then the canonical map $\varphi : R \to S^{-1}R$ is injective, and $S^{-1}R$ can be canonically identified with a subring of the fraction field, $S^{-1}R \subset \mathrm{Frac}(R)$. However, in an arbitrary ring, φ is *not* generally injective. For instance, if S contains zero divisors or nilpotents, non-zero elements of R may map to zero in $S^{-1}R$.

3. Let $f \in R$, and let the multiplicative set be the powers of f: $S_f = \{f^n\}_{n \in \mathbb{N}}$. The resulting localization is denoted $R_f = S_f^{-1}R$, representing "R with f inverted". There is a canonical isomorphism $R_f \cong R[X]/(Xf - 1)$. Notice that if f is a nilpotent element, then $0 \in S_f$, leading back to $R_f = 0$.

End of Lec 06.

7.3.3 Further Examples of Local Rings and Localization

We begin by examining the ring of formal power series over a field to provide another crucial example of a local ring.

Example (Formal Power Series Ring). Let k be a commutative field, and let $R = k[[X]]$ be the ring of formal power series with coefficients in k. An element $f \in k[[X]]$ is of the form $f = \sum_{i=0}^{\infty} a_i X^i$, where $a_i \in k$.
The ideal generated by X is given by $(X) = \{f \in k[[X]] \mid f = \sum_{i=1}^{\infty} a_i X^i\}$. Consider the evaluation homomorphism $k[[X]] \to k$ defined by mapping a series to its constant term, $\overline{f} \mapsto a_0$. The kernel of this surjective ring homomorphism is precisely the ideal (X). Therefore, the quotient ring is isomorphic to the field k ($k[[X]]/(X) \cong k$), which implies that (X) is a maximal ideal.
We claim that $k[[X]]$ is a *local ring* with unique maximal ideal (X).

Proof (Formal Power Series is a Local Ring). By our previous characterization of local rings, it suffices to show that the set of non-invertible elements is exactly the ideal (X). That is, we must prove $k[[X]] \setminus (X) = k[[X]]^{\times}$. Equivalently, an element $f = \sum_{i=0}^{\infty} a_i X^i$ is invertible if and only if its constant term satisfies $a_0 \neq 0$. (For instance, the geometric series yields $\frac{1}{1-X} = \sum_{i=0}^{\infty} X^i$).
Let $f = a_0 + a_1 X + a_2 X^2 + \dots$ with $a_0 \neq 0$. We seek an inverse $g = b_0 + b_1 X + b_2 X^2 + \dots$ such that $f \cdot g = 1$. This translates to the formal equation:

$$(a_0 + a_1 X + \dots)(b_0 + b_1 X + \dots) = 1 + 0X + 0X^2 + \dots$$

Equating coefficients degree by degree yields:

- Degree 0: $a_0 b_0 = 1$. Since $a_0 \neq 0$ in the field k, we can uniquely solve for $b_0 = a_0^{-1}$.
- Degree 1: $a_0 b_1 + a_1 b_0 = 0$. We can solve this to find $b_1 = -\frac{a_1 b_0}{a_0}$.

Continuing inductively, the coefficient of X^n in the product provides a linear equation for b_n in terms of the previously determined coefficients $b_0, \dots, b_{n-1}$. Because a_0 is invertible, we can always uniquely solve for b_n. Thus, the inverse g exists in $k[[X]]$.

□

Next, we explicitly describe the localization at a single element.

Lemma (Localization at an Element). Let R be a commutative ring and $f \in R$. Consider the multiplicative subset $S_f = \{1, f, f^2, \dots\} \subset R$. Then there is a canonical ring isomorphism:

$$R_f := S_f^{-1} R \cong \frac{R[X]}{(fX - 1)}.$$

Proof (Localization at an Element). By the universal property of localization, we construct a homomorphism into the quotient ring. Consider the canonical map $R \to R[X]/(fX - 1)$. In this quotient, the class $\overline{f}$ satisfies $\overline{f} \cdot \overline{X} = 1$, meaning $\overline{f}$ is invertible. Because $\overline{f}$ is a unit, every element in the image of S_f becomes a unit. The universal property yields a unique ring homomorphism $\eta : R_f \to R[X]/(fX - 1)$.
Conversely, we define a homomorphism $\theta : R[X] \to R_f$. Since R_f is an R-algebra, we map the indeterminate X to the element $\frac{1}{f} \in R_f$. This map sends the polynomial $fX - 1$ to $f(\frac{1}{f}) - 1 = 0$. Thus, it factors through the quotient, yielding a well-defined homomorphism $\overline{\theta} : R[X]/(fX - 1) \to R_f$.
One can routinely check that $\eta \circ \overline{\theta} = \mathrm{id}_{R[X]/(fX-1)}$ and $\overline{\theta} \circ \eta = \mathrm{id}_{R_f}$, establishing the isomorphism.

□

Remark. This structure generalizes elegantly to arbitrary multiplicative subsets. If $S \subset R$ is any multiplicative subset, let $R[(X_s)_{s \in S}]$ be the polynomial ring with arbitrary variables X_s for each $s \in S$. Then there is an isomorphism:

$$\frac{R[(X_s)_{s \in S}]}{((sX_s - 1)_{s \in S})} \cong S^{-1} R.$$

As a consequence, if S as a monoid is generated by finitely many elements $(f_1, \dots, f_r) \in S^r$, then the localization is isomorphic to the localization at their product: $S^{-1} R \cong R_{f_1 \cdots f_r}$.

Lemma (Vanishing of Localization). Let $f \in R$. The localized ring $R_f = 0$ if and only if f is *nilpotent*.

Proof (Vanishing of Localization). Forward Implication ($\implies$): Assume $R_f = 0$. In this ring, the identity element is equivalent to zero, meaning $\frac{1}{1} = \frac{0}{1}$. By the definition of the equivalence relation for fractions, there exists an element in S_f, which must be of the form f^N for some $N \geq 1$, such that $f^N(1 \cdot 1 - 1 \cdot 0) = 0$. This simplifies to $f^N \cdot 1 = 0$, meaning $f^N = 0$. Thus, f is nilpotent.

Reverse Implication ($\impliedby$): Assume f is nilpotent, so there exists $n \geq 1$ such that $f^n = 0$. Since $f^n \in S_f$, we necessarily have $0 \in S_f$. Any localization with respect to a multiplicative subset containing 0 is the trivial ring, because for any fraction $\frac{x}{s}$, the relation $0 \cdot (1 \cdot x - s \cdot 0) = 0$ forces $(s, x) \sim (1, 0)$. Thus, $R_f = 0$.

□

7.3.4 Localization at a Prime Ideal

Definition (Localization at a Prime). For any prime ideal $\mathfrak{p} \in \operatorname{Spec}(R)$, the complement $S_\mathfrak{p} := R \setminus \mathfrak{p}$ naturally forms a multiplicative subset of R. The localization of R with respect to this subset is denoted $R_\mathfrak{p}$, representing the *localization of R at $\mathfrak{p}$*.

Remark. If we take finitely many prime ideals $\mathfrak{p}_1, \ldots, \mathfrak{p}_r$, the complement $R \setminus \bigcup_{i=1}^{r} \mathfrak{p}_i$ is also a valid multiplicative subset. The resulting localization yields what is known as a *semi-local ring*.

Proposition (Spectra of Localized Rings). Let $S \subset R$ be a multiplicative subset, and let $\varphi : R \to S^{-1}R$ be the canonical homomorphism. The induced map on spectra, $\operatorname{Spec}(\varphi) : \operatorname{Spec}(S^{-1}R) \to \operatorname{Spec}(R)$ defined by $q \mapsto \varphi^{-1}(q)$, is an injection. Its image is precisely the set of prime ideals in R that do not intersect S:

$$\operatorname{Im}(\operatorname{Spec}(\varphi)) = \{\mathfrak{p} \in \operatorname{Spec}(R) \mid \mathfrak{p} \cap S = \emptyset\}.$$

Proof (Spectra of Localized Rings). Let $I \subset R$ be an ideal. We construct its extension in the localized ring, $S^{-1}I = \{\frac{x}{s} \in S^{-1}R \mid x \in I, s \in S\}$, which forms an ideal in $S^{-1}R$. Clearly, we have $I \subset \varphi^{-1}(S^{-1}I)$.
In particular, if $\mathfrak{p} \in \operatorname{Spec}(R)$ satisfies $\mathfrak{p} \cap S = \emptyset$, we claim that $\varphi^{-1}(S^{-1}\mathfrak{p}) = \mathfrak{p}$. Let $x \in \varphi^{-1}(S^{-1}\mathfrak{p})$. Then $\frac{x}{1} \in S^{-1}\mathfrak{p}$, meaning there exists $t \in S$ such that $tx \in \mathfrak{p}$. Since $\mathfrak{p}$ is a prime ideal and $t \in S \implies t \notin \mathfrak{p}$, it must be that $x \in \mathfrak{p}$. Thus, $\varphi^{-1}(S^{-1}\mathfrak{p}) = \mathfrak{p}$.
Next, we establish as a fact that $S^{-1}\mathfrak{p}$ is a prime ideal in $S^{-1}R$ for any prime ideal $\mathfrak{p}$ satisfying $S \cap \mathfrak{p} = \emptyset$. First, it is a proper ideal: if $S^{-1}\mathfrak{p} = S^{-1}R$, then $\frac{1}{1} = \frac{x}{s}$ for some $x \in \mathfrak{p}$ and $s \in S$. This implies there exists $t \in S$ such that $t(1 \cdot s - 1 \cdot x) = 0$, giving $ts = tx \in \mathfrak{p}$. Since S is multiplicatively closed, $t, s \in S \implies ts \in S$, which means $S \cap \mathfrak{p} \neq \emptyset$, contradicting our assumption. Second, if $(\frac{x}{s})(\frac{y}{t}) \in S^{-1}\mathfrak{p}$, then $\frac{xy}{st} = \frac{z}{u}$ for some $z \in \mathfrak{p}$ and $u \in S$. There exists $v \in S$ such that $vuxy = vstz \in \mathfrak{p}$. Since $v, u \in S$, they are not in $\mathfrak{p}$. Because $\mathfrak{p}$ is prime, $xy \in \mathfrak{p}$, implying $x \in \mathfrak{p}$ or $y \in \mathfrak{p}$. Thus, $\frac{x}{s} \in S^{-1}\mathfrak{p}$ or $\frac{y}{t} \in S^{-1}\mathfrak{p}$, confirming primality.
We have established a well-defined mapping from $\{\mathfrak{p} \in \operatorname{Spec}(R) \mid \mathfrak{p} \cap S = \emptyset\}$ to $\operatorname{Spec}(S^{-1}R)$ via $\mathfrak{p} \mapsto S^{-1}\mathfrak{p}$. The composition $\mathfrak{p} \mapsto S^{-1}\mathfrak{p} \mapsto \varphi^{-1}(S^{-1}\mathfrak{p}) = \mathfrak{p}$ is the identity.
Conversely, let $\mathfrak{q} \in \operatorname{Spec}(S^{-1}R)$. We want to show $\mathfrak{q} = S^{-1}\varphi^{-1}(\mathfrak{q})$. Let $y \in \mathfrak{q}$, and write $y = \frac{x}{s} = \frac{x}{1} \cdot \frac{1}{s}$. Since $\frac{1}{s}$ is invertible in $S^{-1}R$, the ideal $\mathfrak{q}$ absorbs it, meaning $\frac{x}{1} \in \mathfrak{q}$. Thus $x \in \varphi^{-1}(\mathfrak{q})$, giving $y = \frac{x}{s} \in S^{-1}\varphi^{-1}(\mathfrak{q})$. This proves that the maps are mutually inverse bijections, rendering $\operatorname{Spec}(\varphi)$ an injection with the stated image.

□

Corollary 1 (The Local Ring $R_\mathfrak{p}$). For any $\mathfrak{p} \in \operatorname{Spec}(R)$, the localized ring $R_\mathfrak{p}$ is a local ring. Its unique maximal ideal is given by $\mathfrak{m}_\mathfrak{p} := S_\mathfrak{p}^{-1}\mathfrak{p} = \{\frac{x}{s} \in R_\mathfrak{p} \mid x \in \mathfrak{p}\}$.

Proof (The Local Ring $R_\mathfrak{p}$). By the preceding proposition, there is an order-preserving bijection between $\operatorname{Spec}(R_\mathfrak{p})$ and the set $\{\mathfrak{p}' \in \operatorname{Spec}(R) \mid \mathfrak{p}' \cap S_\mathfrak{p} = \emptyset\}$. Recall that $S_\mathfrak{p} = R \setminus \mathfrak{p}$, so the condition $\mathfrak{p}' \cap (R \setminus \mathfrak{p}) = \emptyset$ is equivalent to $\mathfrak{p}' \subset \mathfrak{p}$. Under this inclusion ordering, $\mathfrak{p}$ is the unique maximal element in the target set. Consequently, its image under the bijection, $S_\mathfrak{p}^{-1}\mathfrak{p}$, must be the unique maximal ideal in $\operatorname{Spec}(R_\mathfrak{p})$, proving that $R_\mathfrak{p}$ is a local ring.

□

Corollary 2 (Characterization of the Nilradical). Let R be a commutative ring, and let $\mathrm{Nil}(R) = \{f \in R \mid \exists n \geq 1 \text{ s.t. } f^n = 0\}$ denote the *Nilradical* of R. Then the Nilradical is exactly the intersection of all prime ideals of R:

$$\mathrm{Nil}(R) = \bigcap_{\mathfrak{p} \in \mathrm{Spec}(R)} \mathfrak{p}.$$

Remark. Analogously, the intersection of all maximal ideals, $\bigcap_{\mathfrak{m} \in \mathrm{Spm}(R)} \mathfrak{m}$, is called the *Jacobson radical* of R.

Proof (Characterization of the Nilradical). Let $f \in \mathrm{Nil}(R)$. Then there exists $n \geq 1$ such that $f^n = 0$. For any prime ideal $\mathfrak{p} \in \mathrm{Spec}(R)$, since $0 \in \mathfrak{p}$, we have $f^n \in \mathfrak{p}$. By the definition of primality, this forces $f \in \mathfrak{p}$. Since this holds for all $\mathfrak{p}$, we have $\mathrm{Nil}(R) \subset \bigcap_{\mathfrak{p}} \mathfrak{p}$.
Conversely, let $f \notin \mathrm{Nil}(R)$. We must show there exists some prime ideal $\mathfrak{p}$ such that $f \notin \mathfrak{p}$. Since f is not nilpotent, our previous lemma ensures that the localized ring R_f is non-trivial ($R_f \neq 0$). Because R_f is a non-zero ring, it possesses at least one maximal ideal; let us call it $\mathfrak{m} \in \mathrm{Spec}(R_f)$. Let $\varphi : R \to R_f$ be the canonical map, and define $\mathfrak{p} = \varphi^{-1}(\mathfrak{m}) \in \mathrm{Spec}(R)$. The element f maps to $\frac{f}{1}$, which is invertible in R_f. A proper ideal $\mathfrak{m}$ cannot contain an invertible element, so $\frac{f}{1} \notin \mathfrak{m}$. This directly means $f \notin \varphi^{-1}(\mathfrak{m}) = \mathfrak{p}$. We have found a prime ideal not containing f, proving the reverse inclusion.

□

Remark. Let A and B be rings, and let $\psi : A \to B$ be a ring homomorphism such that $\ker(\psi) \subset \mathrm{Nil}(A)$. Under these conditions, the induced map $\mathrm{Spec}(\psi) : \mathrm{Spec}(B) \to \mathrm{Spec}(A)$ is a bijection.

Definition (Residue Field). The quotient of the local ring $R_\mathfrak{p}$ by its maximal ideal $\mathfrak{m}_\mathfrak{p}$ is a field, called the *residue field* of R at $\mathfrak{p}$, denoted $\kappa(\mathfrak{p})$:

$$\kappa(\mathfrak{p}) := \frac{R_\mathfrak{p}}{\mathfrak{m}_\mathfrak{p}}.$$

We have a natural commutative diagram relating these structures. The composition of the canonical maps $R \to R_\mathfrak{p} \twoheadrightarrow R_\mathfrak{p}/\mathfrak{m}_\mathfrak{p}$ has a kernel exactly equal to $\varphi^{-1}(\mathfrak{m}_\mathfrak{p}) = \mathfrak{p}$. Furthermore, the integral domain $R/\mathfrak{p}$ naturally embeds into $\kappa(\mathfrak{p})$, and $\kappa(\mathfrak{p})$ can be canonically identified with the field of fractions of $R/\mathfrak{p}$.

$$\begin{array}{ccc} R & \xrightarrow{\varphi} & R_\mathfrak{p} \\ \downarrow & & \downarrow \\ R/\mathfrak{p} & \hookrightarrow & R_\mathfrak{p}/\mathfrak{m}_\mathfrak{p} \cong \kappa(\mathfrak{p}) \end{array}$$

7.4 Stability Properties of Noetherian Rings

We conclude the lecture by stating two crucial stability properties of Noetherian rings. Recall that a commutative ring R is Noetherian if every ideal is finitely generated. This condition is equivalent to the *Ascending Chain Condition (ACC)*: for any increasing sequence of ideals $I_0 \subset I_1 \subset I_2 \subset \dots$, there exists an integer $N \gg 0$ such that $I_N = I_{N+1} = I_{N+2} = \dots$.

Theorem (Stability of Noetherianity). Let R be a Noetherian ring.

1. For any multiplicative subset $S \subset R$, the localized ring $S^{-1}R$ is Noetherian.
2. *Hilbert's Basis Theorem*: The polynomial ring $R[X]$ is Noetherian.

We will discuss Krull dimension and geometric examples in the next lecture.

End of Lec 07.

8 Dimension Theory

8.1 Krull Dimension and Height

Let R be a commutative ring. We study the topological dimensions of $\mathrm{Spec}(R)$ by looking at the lengths of chains of prime ideals.

Definition (Krull Dimension). An *ascending chain of prime ideals* in R of length $r \geq 0$ is a strictly increasing sequence of prime ideals:

$$\mathfrak{p}_0 \subsetneq \mathfrak{p}_1 \subsetneq \cdots \subsetneq \mathfrak{p}_r \subset R,$$

where $\mathfrak{p}_i \in \operatorname{Spec}(R)$ for all i.
The *Krull dimension* of R, denoted $\dim_{\text{Krull}}(R)$ or simply $\dim(R)$, is the supremum of all $r \geq 0$ for which there exists an ascending chain of prime ideals of length r. Thus, $\dim(R) \in \mathbb{N} \cup \{\infty\}$.

Examples and Remarks on Krull Dimension.

1. Fields. Let $R = K$ be a field. The only proper ideal is the zero ideal, so $\operatorname{Spec}(K) = \{(0)\}$. The maximal length of an ascending chain is 0, hence $\dim(K) = 0$.

2. Principal Ideal Domains (P.I.D.s). Let R be a P.I.D. which is not a field. There exists at least one irreducible element $\pi \in R \setminus \{0\} \cup R^\times$. The ideal (π) is maximal (and thus prime). We have a chain of length 1:

   $$(0) \subsetneq (\pi).$$

   For any other non-zero prime ideal $\mathfrak{p} \in \operatorname{Spec}(R)$, if we have a chain $(0) \subsetneq (\pi_1) \subset (\pi_2)$ where π_1 is irreducible, the primality implies $(\pi_1) = (\pi_2)$. Thus, the length cannot exceed 1. Consequently, $\dim(R) = 1$. For instance, a Discrete Valuation Ring (D.V.R.) has dimension 1. For any maximal ideal $\mathfrak{p} \in \operatorname{Spec}(R)$, its closure is $\overline{\{\mathfrak{p}\}} = V(\mathfrak{p}) = \{\mathfrak{p}\}$. E.g., for the P.I.D. $k[X]$, localizing at the prime ideal (X) yields $k[X]_{(X)}$, which is a D.V.R. and naturally has dimension 1.

3. Polynomial Rings in Two Variables. Let $R = k[X, Y]$. We can easily construct a chain of length 2:

   $$(0) \subsetneq (X) \subsetneq (X, Y).$$

   Hence, $\dim(k[X, Y]) \geq 2$. Geometrically, this corresponds to the regular functions on the affine plane $\mathbb{A}^2_k$, tracking the sequence of subvarieties: the whole plane, a line, and the origin $(0, 0)$.

4. Localization. If $S \subset R$ is a multiplicative subset, the dimension of the localized ring satisfies $\dim(S^{-1}R) \leq \dim(R)$. This follows from the order-preserving correspondence of prime ideals: if $\mathfrak{q}_0 \subsetneq \cdots \subsetneq \mathfrak{q}_r$ is an ascending chain in $\operatorname{Spec}(S^{-1}R)$, its preimage yields a strict ascending chain $\pi^{-1}(\mathfrak{q}_0) \subsetneq \cdots \subsetneq \pi^{-1}(\mathfrak{q}_r)$ in $\operatorname{Spec}(R)$.

Definition (Height). Let $\mathfrak{p} \in \operatorname{Spec}(R)$ be a prime ideal. The *height* of $\mathfrak{p}$, denoted $\operatorname{ht}(\mathfrak{p})$, is the supremum over all lengths r of ascending chains of prime ideals that terminate exactly at $\mathfrak{p}$:

$$\mathfrak{p}_0 \subsetneq \mathfrak{p}_1 \subsetneq \cdots \subsetneq \mathfrak{p}_r = \mathfrak{p}.$$

Fact (Height and Localization). For any prime ideal $\mathfrak{p} \in \operatorname{Spec}(R)$, the height of $\mathfrak{p}$ is exactly the Krull dimension of the localization of R at $\mathfrak{p}$:

$$\operatorname{ht}(\mathfrak{p}) = \dim(R_\mathfrak{p}).$$

This holds because the prime ideals of $R_\mathfrak{p}$ are in inclusion-preserving bijection with the prime ideals of R contained in $\mathfrak{p}$ (i.e., $\{\mathfrak{q} \in \operatorname{Spec}(R) \mid \mathfrak{q} \cap S_\mathfrak{p} = \emptyset\} = \{\mathfrak{q} \in \operatorname{Spec}(R) \mid \mathfrak{q} \subset \mathfrak{p}\}$).

We will encounter two major examples of finite-dimensional rings in this course:

1. *Finite type commutative k-algebras.* If R is an integral domain and a finite type k-algebra, its Krull dimension equals the transcendence degree of its fraction field over k: $\dim(R) = \operatorname{tr.deg}_k(\operatorname{Frac}(R))$. (For example, $\dim(k[X_1, \ldots, X_n]) = n$).

2. *Noetherian local rings.* Any Noetherian local ring R has finite Krull dimension, $\dim(R) < \infty$.

Regarding the first example, for any prime ideal $\mathfrak{p} \in \operatorname{Spec}(R)$, the quotient $R/\mathfrak{p}$ is a finite k-algebra and an integral domain, and we have the fundamental dimension formula: $\operatorname{ht}(\mathfrak{p}) + \dim(R/\mathfrak{p}) = \dim(R)$.

8.2 Dimension of Polynomial Rings

We now analyze how the dimension changes when passing from a ring R to the polynomial ring $R[X]$.

Lemma (Contraction of Prime Ideals from $R[X]$ to R). Let R be a commutative ring. Let $\mathfrak{q} \subsetneq \mathfrak{q}'$ be a strict inclusion of two prime ideals in $R[X]$, so $\mathfrak{q}, \mathfrak{q}' \in \operatorname{Spec}(R[X])$. Then they cannot contract to the same prime ideal in R. That is, we cannot have:

$$\mathfrak{q} \cap R = \mathfrak{q}' \cap R = \mathfrak{p}.$$

As a consequence, if two prime ideals in $R[X]$ contract to $\mathfrak{p}$, the smaller ideal must precisely be $\mathfrak{p}[X]$.

Proof (Contraction of Prime Ideals). Suppose for the sake of contradiction that we have $\mathfrak{q} \subsetneq \mathfrak{q}'$ with $\mathfrak{q} \cap R = \mathfrak{q}' \cap R = \mathfrak{p}$. First, we can pass to the quotient $R \to R/\mathfrak{p}$. Since $\mathfrak{p}[X] \subset \mathfrak{q} \subsetneq \mathfrak{q}'$, these ideals correspond to distinct prime ideals in the quotient $(R/\mathfrak{p})[X] \cong R[X]/\mathfrak{p}[X]$. By doing so, we reduce to the case where R is an integral domain and the contraction is the zero ideal: $\mathfrak{p} = \mathfrak{q} \cap R = \mathfrak{q}' \cap R = (0)$.
Let $S = R \setminus \{0\}$ be the multiplicative subset of non-zero elements of R. Then the localization $S^{-1}(R[X])$ is isomorphic to $(S^{-1}R)[X] = K[X]$, where $K = \operatorname{Frac}(R)$ is the fraction field of R. By our assumption, both $\mathfrak{q}$ and $\mathfrak{q}'$ intersect R trivially, meaning $\mathfrak{q} \cap S = \emptyset$ and $\mathfrak{q}' \cap S = \emptyset$. Thus, they perfectly extend to prime ideals in the localized ring $K[X]$, maintaining the strict inclusion:

$$S^{-1}\mathfrak{q} \subsetneq S^{-1}\mathfrak{q}' \subset K[X].$$

However, $K[X]$ is a Principal Ideal Domain, meaning $\dim(K[X]) = 1$. The longest possible strict chain of prime ideals is of length 1 (from the zero ideal to a maximal ideal). Therefore, the strictly smaller prime ideal $S^{-1}\mathfrak{q}$ must be the zero ideal. Since $S^{-1}\mathfrak{q} = (0)$ and $\mathfrak{q}$ avoids S, the natural map implies $\mathfrak{q} = (0)$. But if $\mathfrak{q} = (0)$, then $\mathfrak{q}$ is just $\mathfrak{p}[X]$ (which is (0) in our reduced case). This proves that the lower ideal is entirely determined by its contraction, preventing two distinct non-trivial prime ideals from sharing the same contraction.

□

Corollary (Dimension of $R[X]$). Let R be a commutative ring of finite Krull dimension. Then:

$$\dim(R) + 1 \leq \dim(R[X]) \leq 2\dim(R) + 1.$$

All intermediate values in this range are realizable. In particular, $\dim(R[X]) < \infty$. Furthermore, if R is a Noetherian ring, the lower bound is achieved: $\dim(R[X]) = \dim(R) + 1$.

Proof (Dimension of $R[X]$).]Dimension of R[X] We first prove the lower bound, $\dim(R[X]) \geq \dim(R) + 1$. Let $r = \dim(R)$, and let $\mathfrak{p}_0 \subsetneq \cdots \subsetneq \mathfrak{p}_r$ be an ascending chain of prime ideals in R. We define $\mathfrak{q}_i = \mathfrak{p}_i[X]$ for $0 \leq i \leq r$. This yields an ascending chain $\mathfrak{q}_0 \subsetneq \cdots \subsetneq \mathfrak{q}_r$ in $R[X]$. Observe that we can extend this chain strictly by adding the ideal generated by $\mathfrak{q}_r$ and the indeterminate X, yielding $\mathfrak{q}_r + (X)$. This ideal is prime because the quotient $R[X]/(\mathfrak{p}_r[X] + (X)) \cong R/\mathfrak{p}_r$ is an integral domain. Thus, we obtain a chain of length $r + 1$:

$$\mathfrak{q}_0 \subsetneq \cdots \subsetneq \mathfrak{q}_r \subsetneq \mathfrak{q}_r + (X).$$

This confirms $\dim(R[X]) \geq \dim(R) + 1$.
Next, we prove the upper bound. Let $\mathfrak{q}_0 \subsetneq \mathfrak{q}_1 \subsetneq \cdots \subsetneq \mathfrak{q}_s$ be an ascending chain of prime ideals in $R[X]$. Consider their contractions to R: $\mathfrak{p}_i = \mathfrak{q}_i \cap R$. This produces a sequence of prime ideals in R:

$$\mathfrak{p}_0 \subseteq \mathfrak{p}_1 \subseteq \cdots \subseteq \mathfrak{p}_s.$$

By the preceding Lemma, we cannot have three consecutive equal contractions (i.e., $\mathfrak{p}_i = \mathfrak{p}_{i+1} = \mathfrak{p}_{i+2}$ is impossible). Therefore, in the sequence of contractions, equality can occur at most in pairs. Extracting the strictly increasing elements yields a chain in R of length at least $s/2$ (or precisely $(s-1)/2$ if s is odd). Since this chain in R can have length at most $\dim(R)$, we deduce $s \leq 2\dim(R) + 1$.

□

9 Module Theory and Algebras

9.1 R-Modules, Submodules, and Quotients

Definition (R-Module). Let R be a commutative ring. An *R-module* M consists of:

1. An abelian group $(M, +)$, possessing a zero element 0, and additive inverses $-x$.
2. A scalar multiplication map $\cdot : R \times M \to M$, mapping $(\lambda, x) \mapsto \lambda \cdot x$, satisfying the following axioms for all $x, y \in M$ and $\lambda, \mu \in R$:
   - Identity: $1 \cdot x = x$.
   - Associativity: $\lambda \cdot (\mu \cdot x) = (\lambda\mu) \cdot x$.
   - Distributivity over module addition: $\lambda \cdot (x + y) = \lambda \cdot x + \lambda \cdot y$.
   - Distributivity over ring addition: $(\lambda + \mu) \cdot x = \lambda \cdot x + \mu \cdot x$.

Examples and Remarks on Modules.

1. By standard deductions from the axioms, $0_R \cdot x = 0_M$ for all $x \in M$.
2. If $R = \mathbb{Z}$, the axioms of a $\mathbb{Z}$-module exactly coincide with those of an abelian group. Thus, every abelian group is naturally a $\mathbb{Z}$-module.
3. If $R = K$ is a field, an R-module is precisely a K-vector space.
4. If R is merely an associative ring with unit (not necessarily commutative), the above definition yields the concept of a *left R-module*. (Right R-modules are defined analogously using right multiplication).

Definition (Morphism of R-Modules). Let M and N be R-modules. A map $f : M \to N$ is called a *morphism of R-modules* (or an R-linear map) if it is a homomorphism of the underlying abelian groups and preserves scalar multiplication:

$$f(x + y) = f(x) + f(y),$$
$$f(\lambda \cdot x) = \lambda \cdot f(x),$$

for all $x, y \in M$ and $\lambda \in R$.

If $R = \mathbb{Z}$, these morphisms are exactly group homomorphisms between abelian groups. If $R = K$, they are linear transformations between vector spaces.

The Category of R-Modules. For R-modules M and N, we denote the set of all morphisms from M to N by $\mathrm{Hom}_{R\text{-mod}}(M, N)$. This set inherently carries the structure of an R-module itself, with point-wise addition and scalar multiplication.
The class of all R-modules together with these Hom sets defines the *category of R-modules*, denoted R-Mod.

Definition (Submodule and Associated Constructions).

1. Submodules: A *sub-R-module* $N \subset M$ of an R-module M is an additive subgroup of M that is stable under scalar multiplication (i.e., $\forall \lambda \in R$ and $\forall x \in N$, $\lambda x \in N$).
2. Kernel and Image: Let $f : M \to N$ be a morphism of R-modules. The kernel, $\ker(f) = \{x \in M \mid f(x) = 0\}$, is a submodule of M. The image, $\mathrm{Im}(f) = \{f(x) \in N \mid x \in M\}$, is a submodule of N.
3. Cokernel: Given $f : M \to N$, we construct the quotient module $\mathrm{Coker}(f) := N/\mathrm{Im}(f)$.

Any morphism $f : M \to N$ canonically factors as the composition of a surjection, an isomorphism, and an injection:

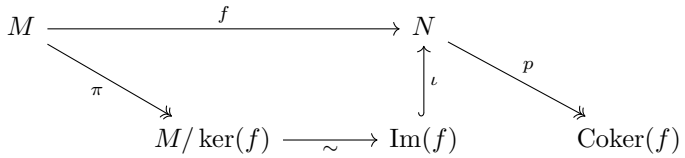

As a final remark, viewing the ring R as a module over itself, an ideal $I \subset R$ is exactly a sub-R-module of R.

End of Lec 08.

9.2 Finitely Generated and Noetherian Algebras

We begin by stating important stability properties of Noetherian rings and modules.

Fact (Stability of Noetherian Modules). Let R be a Noetherian commutative ring. If M is a finitely generated R-module (i.e., a module of finite type), then M is a Noetherian R-module.

Recall Hilbert's Basis Theorem, which asserts that if R is a Noetherian ring, then the polynomial ring $R[X]$ is also Noetherian. This extends directly to finitely generated algebras.

Corollary (Noetherian Algebras). Let R be a Noetherian ring. If A is a finite type R-algebra, then A is a Noetherian ring.

Proof (Noetherian Algebras). By definition, a finite type R-algebra A is generated by a finite set of elements $x_1, \dots, x_n$ over R. This means there exists a surjective ring homomorphism from the polynomial ring $R[X_1, \dots, X_n]$ onto A, given by evaluation at the generators. Thus, by the first isomorphism theorem, $A \cong R[X_1, \dots, X_n]/I$ for some ideal I. By Hilbert's Basis Theorem applied inductively, $R[X_1, \dots, X_n]$ is a Noetherian ring. Since any quotient of a Noetherian ring is Noetherian, A is necessarily a Noetherian ring.

□

10 Homological Algebra

10.1 Exact Sequences and Cochain Complexes

Definition (Exact Sequence). Let R be a commutative ring and let M_1, M_2, M_3 be R-modules. A sequence of R-module morphisms

$$M_1 \xrightarrow{f} M_2 \xrightarrow{g} M_3$$

is said to be *exact* (at M_2) if $g \circ f = 0$ and $\ker(g) = \mathrm{Im}(f)$.

A sequence of arbitrary length $\cdots \to M_i \xrightarrow{\delta_i} M_{i+1} \xrightarrow{\delta_{i+1}} M_{i+2} \to \dots$ is called exact if it is exact at every intermediate module M_{i+1} (i.e., $\ker(\delta_{i+1}) = \mathrm{Im}(\delta_i)$ for all i).

A *Short Exact Sequence* (SES) is an exact sequence of the form:

$$0 \longrightarrow M_1 \xrightarrow{f} M_2 \xrightarrow{g} M_3 \longrightarrow 0.$$

Exactness at each node translates to the following conditions:

1. Exactness at M_1: $\ker(f) = 0$, meaning f is injective.
2. Exactness at M_2: $\ker(g) = \mathrm{Im}(f)$.
3. Exactness at M_3: $\mathrm{Im}(g) = M_3$, meaning g is surjective (equivalently, $\mathrm{Coker}(g) = 0$).

If we relax the exactness condition to only require $g \circ f = 0$, we obtain the notion of a complex.

Definition (Cochain Complex and Cohomology). A *cochain complex* (M_*, δ_*) is a sequence of R-modules and morphisms (differentials):

$$\cdots \longrightarrow M_n \xrightarrow{\delta_n} M_{n+1} \xrightarrow{\delta_{n+1}} M_{n+2} \longrightarrow \dots$$

such that for all $n \in \mathbb{Z}$, the composition is zero: $\delta_{n+1} \circ \delta_n = 0$. This is equivalent to the inclusion $\mathrm{Im}(\delta_n) \subset \ker(\delta_{n+1})$.

The *n-th cohomology group* of the complex is defined as the quotient:

$$H^n(M_*, \delta_*) := \frac{\ker(\delta_{n+1})}{\mathrm{Im}(\delta_n)}.$$

Elements of $\ker(\delta_{n+1})$ are called n-cocycles (often denoted Z^n), and elements of $\mathrm{Im}(\delta_n)$ are called n-coboundaries (often denoted B^n).

A cochain complex is an exact sequence if and only if $H^n(M_*, \delta_*) = 0$ for all $n \in \mathbb{Z}$.

Examples of Complexes.

1. De Rham Complex: Let M be a C^∞ differentiable manifold and let T^*_M denote its cotangent bundle. The spaces of differential k-forms $\Omega^k(M)$ along with the exterior derivative d form a complex:

$$0 \longrightarrow \Omega^0(M) \xrightarrow{d^0} \Omega^1(M) \xrightarrow{d^1} \dots \xrightarrow{d^{n-1}} \Omega^{\dim M}(M) \longrightarrow 0$$

   Here, $\Omega^0(M) = C^\infty(M, \mathbb{R})$. The differential $d^0(f) = df$ maps functions to 1-forms. The resulting cohomology groups $H^n(\Omega^*(M), d)$ are the De Rham cohomology groups of M, denoted $H^n_{dR}(M)$.

2. Kähler Differentials: Let A be an R-algebra. The module of Kähler differentials $\Omega^1_{A/R}$ provides an algebraic analogue to differential forms, defined equipped with a universal derivation $d : A \to \Omega^1_{A/R}$.

10.2 Free Modules and Presentations

Let $(M_\alpha)_{\alpha \in I}$ be a family of R-modules. We distinguish between their direct product and direct sum.

Definition (Direct Product and Direct Sum). The *direct product* $\prod_{\alpha \in I} M_\alpha$ is the Cartesian product of the underlying sets, equipped with component-wise addition and scalar multiplication. The *direct sum* $\bigoplus_{\alpha \in I} M_\alpha$ is the submodule of the direct product consisting of sequences $(m_\alpha)_{\alpha \in I}$ such that $m_\alpha = 0$ for all but finitely many indices $\alpha \in I$.

For any R-module N, the sets of homomorphisms distribute over these constructions as follows:

$$\mathrm{Hom}_R\left(N, \prod_{\alpha \in I} M_\alpha\right) \cong \prod_{\alpha \in I} \mathrm{Hom}_R(N, M_\alpha)$$

$$\mathrm{Hom}_R\left(\bigoplus_{\alpha \in I} M_\alpha, N\right) \cong \prod_{\alpha \in I} \mathrm{Hom}_R(M_\alpha, N)$$

Definition (Free Module). An R-module L is called a *free module* if there exists a subset $S \subset L$ such that the canonical map:

$$\begin{aligned} \bigoplus_{s \in S} R &\longrightarrow L \\ (\lambda_s)_{s \in S} &\longmapsto \sum_{s \in S} \lambda_s s \end{aligned}$$

is an R-module isomorphism. The set S is called a basis of L.

This construction generalizes the concept of a basis from linear algebra over fields to arbitrary rings.

Lemma (Rank of a Free Module). If $R \neq 0$ and L is a free R-module, the cardinality of the basis S depends only on L and not on the choice of basis. This cardinality is called the *rank* of L.

Proof (Well-definedness of Rank). Suppose L admits two bases, S and S', yielding isomorphisms $L \cong \bigoplus_{s \in S} R \cong \bigoplus_{s' \in S'} R$. Because $R \neq 0$, it contains at least one maximal ideal $\mathfrak{m}$.
We can tensor the entire isomorphism with the quotient field $R/\mathfrak{m}$. Algebraically, this corresponds to quotienting the module by $\mathfrak{m}L$. The isomorphism reduces to an isomorphism of vector spaces over the field $R/\mathfrak{m}$:

$$\bigoplus_{s \in S} R/\mathfrak{m} \cong L/\mathfrak{m}L \cong \bigoplus_{s' \in S'} R/\mathfrak{m}.$$

From standard linear algebra, the dimensions of isomorphic vector spaces are equal, dictating that $|S| = |S'|$.

□

If $|S| = n < \infty$, we write $L \cong R^n$, and L is a free module of finite rank.

Definition (Presentation of a Module). A *presentation* of an R-module M is an exact sequence of the form:

$$L_1 \longrightarrow L_0 \longrightarrow M \longrightarrow 0$$

where both L_0 and L_1 are free R-modules. This implies $M \cong L_0/\operatorname{Im}(L_1)$, meaning M is defined by a set of generators (basis of L_0) and a set of relations (image of the basis of L_1).

Every R-module M admits at least one presentation. We can trivially choose $L_0 = \bigoplus_{m \in M} R$ and define the surjection $(\lambda_m) \mapsto \sum \lambda_m m$. The kernel of this map is another R-module, which in turn admits a surjection from a free module L_1, yielding the sequence.

Definition (Finite Type and Finite Presentation). An R-module M is said to be:

1. Of *finite type* (finitely generated) if there exists an exact sequence $L_0 \to M \to 0$ where L_0 is a free R-module of finite rank.
2. Of *finite presentation* if there exists a full presentation $L_1 \to L_0 \to M \to 0$ where both L_0 and L_1 are free R-modules of finite rank.

If R is a Noetherian ring, an R-module M is of finite type if and only if it is of finite presentation. This follows because the kernel of $L_0 \to M$ is a submodule of a Noetherian module, hence is finitely generated and admits a surjection from a finite free module L_1.

10.3 Extensions of Modules

Definition (Module Extension). Let M, N be R-modules. An *extension* of M by N is a short exact sequence of the form:

$$\mathcal{E}: \quad 0 \longrightarrow N \xrightarrow{i} E \xrightarrow{\pi} M \longrightarrow 0.$$

Two extensions $\mathcal{E}$ and $\mathcal{E}'$ (with central modules E and E') are isomorphic if there exists an R-module morphism $\varphi : E \to E'$ such that the following diagram commutes:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & N & \xrightarrow{i} & E & \xrightarrow{\pi} & M & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle \mathrm{Id}_N} & & \downarrow{\scriptstyle \varphi} & & \downarrow{\scriptstyle \mathrm{Id}_M} & & \\ 0 & \longrightarrow & N & \xrightarrow[i']{} & E' & \xrightarrow[\pi']{} & M & \longrightarrow & 0 \end{array}$$

By the Five Lemma, any such morphism φ is automatically an isomorphism. We define $\operatorname{Ext}_R(M, N)$ (more specifically $\operatorname{Ext}^1_R(M, N)$) as the set of isomorphism classes of extensions of M by N.

10.3.1 Pull-back and Push-forward

Extensions can be manipulated via morphisms on the base module M or the sub-module N.

Definition (Pull-back of an Extension). Let $\mathcal{E} : 0 \to N \xrightarrow{i} E \xrightarrow{\pi} M \to 0$ be an extension, and let $f : M' \to M$ be a morphism. The *pull-back* extension $f^*(\mathcal{E})$ is constructed via the fiber product.
We define $E' \subset E \times M'$ as the submodule:

$$E' = \{(e, m') \in E \times M' \mid \pi(e) = f(m')\}.$$

The projection $(e, m') \mapsto m'$ is surjective, and the kernel is naturally isomorphic to N via $n \mapsto (i(n), 0)$. This yields the pulled-back exact sequence:

$$f^*(\mathcal{E}): \quad 0 \longrightarrow N \longrightarrow E' \longrightarrow M' \longrightarrow 0.$$

This construction provides a map $f^* : \operatorname{Ext}_R(M, N) \to \operatorname{Ext}_R(M', N)$.

Definition (Push-forward of an Extension). Let $\mathcal{E} : 0 \to N \xrightarrow{i} E \xrightarrow{\pi} M \to 0$ be an extension, and let $g : N \to N'$ be a morphism. The *push-forward* extension $g_*(\mathcal{E})$ is constructed via the pushout.
We define E' as the quotient of the direct sum $N' \oplus E$:

$$E' = \frac{N' \oplus E}{\{(-g(n), i(n)) \mid n \in N\}}.$$

The natural map $E' \to M$ given by $\overline{(n', e)} \mapsto \pi(e)$ is well-defined and surjective, with kernel isomorphic to N'. This yields the pushed-forward exact sequence:

$$g_*(\mathcal{E}) : \quad 0 \longrightarrow N' \longrightarrow E' \longrightarrow M \longrightarrow 0.$$

This construction provides a map $g_* : \mathrm{Ext}_R(M, N) \to \mathrm{Ext}_R(M, N')$.

10.3.2 The Baer Sum and Monoid Structure

The set $\mathrm{Ext}_R(M, N)$ can be equipped with an algebraic structure using the Baer sum.

Definition (Baer Sum). Let $\mathcal{E}_1$ and $\mathcal{E}_2$ be two extensions of M by N:

$$\mathcal{E}_1 : \quad 0 \longrightarrow N \longrightarrow E_1 \longrightarrow M \longrightarrow 0$$
$$\mathcal{E}_2 : \quad 0 \longrightarrow N \longrightarrow E_2 \longrightarrow M \longrightarrow 0$$

Taking their direct sum produces an extension of $M \oplus M$ by $N \oplus N$:

$$\mathcal{E}_1 \oplus \mathcal{E}_2 : \quad 0 \longrightarrow N \oplus N \longrightarrow E_1 \oplus E_2 \longrightarrow M \oplus M \longrightarrow 0.$$

We define the diagonal map $\Delta_M : M \to M \oplus M$ by $m \mapsto (m, m)$, and the codiagonal map $\nabla_N : N \oplus N \to N$ by $(n_1, n_2) \mapsto n_1 + n_2$.
The *Baer sum* of the two extensions is defined by pulling back along Δ_M and pushing forward along ∇_N:

$$\mathcal{E}_1 + \mathcal{E}_2 := \nabla_{N*}\Delta_M^*(\mathcal{E}_1 \oplus \mathcal{E}_2) \in \mathrm{Ext}_R(M, N).$$

Theorem (Group Structure of Ext). Under the Baer sum, the set $\mathrm{Ext}_R(M, N)$ forms an abelian group.

- The addition is associative and commutative.
- The neutral element $0 \in \mathrm{Ext}_R(M, N)$ is the equivalence class of the split (trivial) extension:

$$0 \longrightarrow N \longrightarrow N \oplus M \longrightarrow M \longrightarrow 0.$$

- The additive inverse of an extension $\mathcal{E}$ is obtained by scaling the inclusion map by -1, i.e., $-\mathcal{E} = (-1_N)_*(\mathcal{E}) = (-1_M)^*(\mathcal{E})$.

As a concrete example, consider $R = \mathbb{Z}$ and $M = N = \mathbb{Z}/p\mathbb{Z}$ for a prime p. The group is $\mathrm{Ext}_{\mathbb{Z}}(\mathbb{Z}/p, \mathbb{Z}/p) \cong \mathbb{Z}/p$. The zero element is the trivial split extension $\mathbb{Z}/p \oplus \mathbb{Z}/p$. A generator of the group corresponds to the non-trivial extension:

$$0 \longrightarrow \mathbb{Z}/p \xrightarrow{\times p} \mathbb{Z}/p^2 \longrightarrow \mathbb{Z}/p \longrightarrow 0.$$

End of Lec 09.

10.4 The Snake Lemma and Long Exact Sequences

Theorem (The Snake Lemma). Let R be a commutative ring. Consider the following commutative diagram of R-modules, where the two horizontal lines are exact:

$$\begin{array}{ccccccc} & & M_1 & \xrightarrow{\varphi} & M_2 & \xrightarrow{\psi} & M_3 \longrightarrow 0 \\ & & \downarrow{\scriptstyle f_1} & & \downarrow{\scriptstyle f_2} & & \downarrow{\scriptstyle f_3} \\ 0 & \longrightarrow & N_1 & \xrightarrow[\varphi']{} & N_2 & \xrightarrow[\psi']{} & N_3 \end{array}$$

Then there exists a canonical exact sequence:

$$\ker(f_1) \longrightarrow \ker(f_2) \longrightarrow \ker(f_3) \xrightarrow{\delta} \operatorname{Coker}(f_1) \longrightarrow \operatorname{Coker}(f_2) \longrightarrow \operatorname{Coker}(f_3).$$

where δ is called the connecting homomorphism.
Furthermore, if φ is an injective morphism, then the induced map $\ker(f_1) \to \ker(f_2)$ is injective.
If ψ' is a surjective morphism, then the induced map $\operatorname{Coker}(f_2) \to \operatorname{Coker}(f_3)$ is surjective.

Proof (Construction of the Connecting Homomorphism). We prove exactness at $\ker(f_3)$ by explicitly constructing the connecting homomorphism δ.
Let $x \in \ker(f_3)$. Because the top row is exact at M_3, the map $\psi : M_2 \to M_3$ is surjective. Therefore, there exists an element $y \in M_2$ such that $\psi(y) = x$. By the commutativity of the right square, we evaluate the image of y in N_3:

$$\psi'(f_2(y)) = f_3(\psi(y)) = f_3(x) = 0.$$

This implies that $f_2(y) \in \ker(\psi')$. By the exactness of the bottom row, $\ker(\psi') = \operatorname{Im}(\varphi')$. Since the map $\varphi' : N_1 \to N_2$ is injective, there exists a unique element $z \in N_1$ such that $\varphi'(z) = f_2(y)$. We define the connecting homomorphism by evaluating the class of this element in the cokernel:

$$\delta(x) := \overline{z} \in \operatorname{Coker}(f_1) = N_1 / \operatorname{Im}(f_1).$$

We must check the independence of δ from the choice of the preimage y. Let $y' \in M_2$ be another element satisfying $\psi(y') = x$. Then $\psi(y - y') = \psi(y) - \psi(y') = 0$, meaning $y - y' \in \ker(\psi)$. By the exactness of the top row, $\ker(\psi) = \operatorname{Im}(\varphi)$, so there exists an element $p \in M_1$ such that $y' = y + \varphi(p)$.
Applying f_2 to this relation yields:

$$f_2(y') = f_2(y) + f_2(\varphi(p)).$$

By the commutativity of the left square, $f_2 \circ \varphi = \varphi' \circ f_1$, so $f_2(y') = f_2(y) + \varphi'(f_1(p))$. Let $z' \in N_1$ be the unique element such that $\varphi'(z') = f_2(y')$. Then:

$$\varphi'(z') = \varphi'(z) + \varphi'(f_1(p)) = \varphi'(z + f_1(p)).$$

Since φ' is injective, we conclude $z' = z + f_1(p)$. When passing to the quotient $\operatorname{Coker}(f_1)$, the term $f_1(p)$ is absorbed, yielding $\overline{z'} = \overline{z}$. Hence, the map δ is well-defined. It routinely follows via diagram chasing that δ is a morphism of R-modules and that the resulting six-term sequence is exact everywhere.

□

Corollary (Generalization to Chain Complexes). Let $0 \to C_* \to D_* \to E_* \to 0$ be a short exact sequence of chain complexes. Specifically, for each degree n, we have a short exact sequence of R-modules:

$$0 \longrightarrow C_n \longrightarrow D_n \longrightarrow E_n \longrightarrow 0$$

compatible with the boundary operators. Then there exists a canonical long exact sequence of homology groups:

$$\cdots \longrightarrow H_n(C_*) \longrightarrow H_n(D_*) \longrightarrow H_n(E_*) \xrightarrow{\delta} H_{n-1}(C_*) \longrightarrow \dots$$

Proof (Long Exact Sequence). The exact sequence is constructed by applying the Snake Lemma iteratively. For a specific degree n, one considers the commutative diagram formed by the boundary operators acting on the modules. The boundary maps induce morphisms between the respective cycles and boundaries, and taking the cokernel of the boundaries alongside the kernel of the differentials produces the homology groups. The connecting homomorphism δ directly

descends from the application of the Snake Lemma on these factored spaces. The full proof of exactness everywhere is completed by a standard diagram chase.

□

10.5 Projective and Injective Modules

Definition (Projective and Injective Modules). Let R be a commutative ring.

1. An R-module P is said to be *projective* if for every surjective R-module morphism $\pi : M \twoheadrightarrow N$ and every morphism $f : P \to N$, there exists a morphism $q : P \to M$ such that $\pi \circ q = f$.
2. Dually, an R-module I is said to be *injective* if for every injective R-module morphism $\iota : N \hookrightarrow M$ and every morphism $f : N \to I$, there exists a morphism $q : M \to I$ such that $q \circ \iota = f$.

Example (Free Modules). Let S be an arbitrary set. The free R-module $R[S] = \bigoplus_{s \in S} R$ is a projective module.

Proof (Projectivity of Free Modules). Recall that by the universal property of free modules, any map from the basis S to a module N extends uniquely to a module homomorphism. Therefore, there is a natural isomorphism $\mathrm{Hom}_R(R[S], N) \cong \mathrm{Map}(S, N) = N^S$.
Let $\pi : M \to N$ be a surjective R-module morphism. Because π is surjective on the underlying sets, the induced mapping between the Cartesian products of sets, $M^S \to N^S$, is also surjective. Consequently, the induced homomorphism $\mathrm{Hom}_R(R[S], M) \to \mathrm{Hom}_R(R[S], N)$ is surjective. Thus, any map $f : R[S] \to N$ lifts to a map $q : R[S] \to M$, satisfying the definition of projectivity.

□

Lemma (Summands of Free Modules). An R-module P is projective if and only if it is a direct summand of a free R-module.

Proof (Summands of Free Modules). For any module P, we can construct a free module generated by the elements of P, denoted $R[P]$. The canonical evaluation map $R[P] \twoheadrightarrow P$ is surjective. If P is projective, the identity map $\mathrm{id}_P : P \to P$ lifts to a morphism $q : P \to R[P]$ such that the composition with the evaluation is the identity. This implies that the exact sequence $0 \to \ker \to R[P] \to P \to 0$ splits, rendering P a direct summand of the free module $R[P]$. The converse holds because a direct summand of a projective module inherits the lifting property, and free modules are projective.

□

Excursion: Grothendieck Group. The properties of projective modules allow for the definition of the Grothendieck group, $K_0(R)$. Let $\mathcal{P}(R)$ be the set of isomorphism classes of finitely generated projective R-modules. The group $K_0(R)$ is defined as the quotient of the free abelian group $\mathbb{Z}[\mathcal{P}(R)]$ by the subgroup generated by elements of the form $[P \oplus Q] - [P] - [Q]$.

Fact (Divisible Groups as Injective Modules). Over the ring of integers $R = \mathbb{Z}$, the module of rational numbers $\mathbb{Q}$ is an injective module. Consequently, the quotient $\mathbb{Q}/\mathbb{Z}$ is also an injective $\mathbb{Z}$-module.

The injectivity of these specific modules corresponds to their property as divisible groups. A module is divisible if for any non-zero scalar $n \in \mathbb{Z} \setminus \{0\}$ and any element x, the equation $ny = x$ admits a solution. For example, in $\mathbb{Q}/\mathbb{Z}$, an element takes the form of a fraction class $x = (a/b)$, and division by n trivially yields $y = (a/(nb))$.
The general proof of injectivity for divisible modules uses Zorn's Lemma. If one has an injection $N \hookrightarrow M$ and a map from N to an injective candidate I, one considers the poset of pairs (C, φ) where $N \subset C \subset M$ and $\varphi : C \to I$ extends the original map. Zorn's Lemma guarantees a maximal element. If this maximal extension does not cover M, one can select an element outside the domain and extend the map explicitly, contradicting maximality.
For arbitrary R-modules, one can construct injective modules using a "coinduced" construction. The module $\mathrm{Hom}_{\mathbb{Z}}(R, \mathbb{Q}/\mathbb{Z})$ serves as an injective cogenerator for the category of R-modules.

11 Multilinear Algebra

11.1 The Tensor Product

11.1.1 Bilinear Maps and Construction

Definition (Bilinear Map). Let M_1, M_2, and N be R-modules. A map $\Phi : M_1 \times M_2 \to N$ is called an *R-bilinear map* if it satisfies the following conditions:

1. For all $\lambda \in R$ and for all $(m_1, m_2) \in M_1 \times M_2$:
$$\Phi(\lambda m_1, m_2) = \lambda\Phi(m_1, m_2) = \Phi(m_1, \lambda m_2).$$
2. The map is additive in each coordinate separately:
$$\Phi(m_1 + m_1', m_2) = \Phi(m_1, m_2) + \Phi(m_1', m_2),$$
$$\Phi(m_1, m_2 + m_2') = \Phi(m_1, m_2) + \Phi(m_1, m_2').$$

The set of all such R-bilinear maps is denoted by $B_R(M_1, M_2; N)$. This set inherently carries the structure of an R-module.

Standard instances of bilinear maps include:

1. For any module M, let $M^* = \mathrm{Hom}_R(M, R)$ be its dual. The evaluation map $M \times M^* \to R$ given by $(m, l) \mapsto l(m)$ is a bilinear form.
2. More generally, the evaluation $M \times \mathrm{Hom}_R(M, N) \to N$ is bilinear.
3. Composition of homomorphisms is bilinear: $\mathrm{Hom}_R(M, N) \times \mathrm{Hom}_R(N, Q) \to \mathrm{Hom}_R(M, Q)$.
4. If k is a field, a non-degenerate symmetric bilinear form $\Phi : V \times V \to k$ over a k-vector space V induces an isomorphism $V \cong \mathrm{Hom}_k(V, k) = V^*$.

The definition of a bilinear map can be interpreted as an instance of a right adjoint functor. If $\Phi : M_1 \times M_2 \to N$ is an R-bilinear map, fixing the first coordinate yields a module homomorphism $\Phi(m_1, -) : M_2 \to N \in \mathrm{Hom}_{R\text{-mod}}(M_2, N)$. This mapping assigns to each $m_1 \in M_1$ a homomorphism in $\mathrm{Hom}_R(M_2, N)$, defining a representable functor that provides the canonical isomorphism:

$$B_R(M_1, M_2; N) \cong \mathrm{Hom}_R(M_1, \mathrm{Hom}_R(M_2, N)).$$

We construct a specific R-module, denoted $M \otimes_R N$, which universalizes the concept of bilinear maps. The objective is to achieve the isomorphism $B_R(M, N; Q) \cong \mathrm{Hom}_R(M \otimes_R N, Q)$.

Definition (The Tensor Product). Let $M, N \in R$-mod. We define $R[M \times N] = \bigoplus_{M \times N} R$ to be the free R-module generated by the elements of the Cartesian product $M \times N$.
Let $R_{M,N} \subset R[M \times N]$ be the sub-R-module generated by the following four types of formal differences for all $m, m_1, m_2 \in M$, $n, n_1, n_2 \in N$, and $\lambda \in R$:

1. $(\lambda m, n) - \lambda(m, n)$
2. $(m, \lambda n) - \lambda(m, n)$
3. $(m_1 + m_2, n) - (m_1, n) - (m_2, n)$
4. $(m, n_1 + n_2) - (m, n_1) - (m, n_2)$

The *tensor product* of M and N over R is defined as the quotient module:

$$M \otimes_R N := \frac{R[M \times N]}{R_{M,N}}.$$

For any pair $(m, n) \in M \times N$, we denote its equivalence class in the quotient module by $m \otimes n$. The canonical projection map mapping the Cartesian product into the tensor product, $M \times N \to M \otimes_R N$ defined by $(m, n) \mapsto m \otimes n$, satisfies all the linearity relations by definition, making it an R-bilinear map.

End of Lec 10.

11.1.2 Universal Property and Structural Isomorphisms

Let M and N be R-modules. We recall the construction of the tensor product $M \otimes_R N = R[M \times N]/R_{M,N}$. For any pair $(m, n) \in M \times N$, we denote its equivalence class in the quotient module by $m \otimes n := \overline{(m, n)}$.
By construction, the canonical map $\Theta_{M,N} : M \times N \to M \otimes_R N$ defined by $(m, n) \mapsto m \otimes n$ is an R-bilinear map. It directly follows from the relations defining $R_{M,N}$ that the elements $m \otimes n$ satisfy the following arithmetic identities for all $m, m_i \in M$, $n, n_i \in N$, and $\lambda \in R$:

1. $\lambda(m \otimes n) = (\lambda m) \otimes n = m \otimes (\lambda n)$
2. $(m_1 + m_2) \otimes n = m_1 \otimes n + m_2 \otimes n$
3. $m \otimes (n_1 + n_2) = m \otimes n_1 + m \otimes n_2$
4. $0 \otimes n = m \otimes 0 = 0$
5. $(-m) \otimes n = -(m \otimes n) = m \otimes (-n)$

Theorem (Universal Property of the Tensor Product). Let M, N, and Q be R-modules. There is a canonical isomorphism of R-modules between the module of bilinear maps and the module of homomorphisms from the tensor product:

$$B_R(M, N; Q) \xrightarrow{\sim} \mathrm{Hom}_{R\text{-mod}}(M \otimes_R N, Q).$$

$$\begin{array}{ccc} M \times N & \xrightarrow{\Phi} & Q \\ {\scriptstyle \Theta_{M,N}}\downarrow & \nearrow_{\tilde{\Phi}} & \\ M \otimes_R N & & \end{array}$$

Proof (Universal Property of the Tensor Product). Let S be an arbitrary set, and $R[S]$ the free module generated by S. For any module Q, a homomorphism $f \in \mathrm{Hom}_R(R[S], Q)$ is uniquely determined by its values on the basis S. If we quotient $R[S]$ by a submodule T generated by elements $t = \sum_s \lambda_{t,s} s$, a homomorphism from the quotient $R[S]/T \to Q$ corresponds precisely to a map $f : S \to Q$ such that $\sum_s \lambda_{t,s} f(s) = 0$ in Q for all $t \in T$.
We apply this to $S = M \times N$ and $T = R_{M,N}$. A homomorphism $\varphi \in \mathrm{Hom}_R(M \otimes_R N, Q)$ corresponds to a map $M \times N \to Q$ that vanishes on the relations $R_{M,N}$. Evaluating the conditions for vanishing on the generators of $R_{M,N}$ yields:

$$\begin{aligned} \varphi(\lambda m, n) - \lambda\varphi(m, n) &= 0, \\ \varphi(m, \lambda n) - \lambda\varphi(m, n) &= 0, \\ \varphi(m_1 + m_2, n) - \varphi(m_1, n) - \varphi(m_2, n) &= 0, \\ \varphi(m, n_1 + n_2) - \varphi(m, n_1) - \varphi(m, n_2) &= 0. \end{aligned}$$

These are exactly the axioms defining an R-bilinear map. Thus, every bilinear map uniquely factors through the tensor product, establishing the isomorphism.

□

Lemma (Generators of the Tensor Product). Let $(m_i)_{i \in I}$ be a generating family for M, and let $(n_j)_{j \in J}$ be a generating family for N. Then the family of pure tensors $(m_i \otimes n_j)_{(i,j) \in I \times J}$ generates the tensor product $M \otimes_R N$.

Proof (Generators of the Tensor Product). Let $m \otimes n \in M \otimes_R N$. Since $(m_i)_{i \in I}$ and $(n_j)_{j \in J}$ are generating families, there exist scalars $\lambda_i, \mu_j \in R$ such that $m = \sum_i \lambda_i m_i$ and $n = \sum_j \mu_j n_j$ (where all but finitely many scalars are zero). By applying the bilinearity of the tensor product, we expand the expression:

$$\begin{aligned} m \otimes n &= \left(\sum_i \lambda_i m_i \right) \otimes \left(\sum_j \mu_j n_j \right) \\ &= \sum_{i,j} (\lambda_i \mu_j)(m_i \otimes n_j). \end{aligned}$$

Since elements of the form $m \otimes n$ generate $M \otimes_R N$ by definition, the family $(m_i \otimes n_j)_{(i,j)\in I\times J}$ generates the entire module.

□

As an immediate corollary, if M and N are R-modules of finite type (finitely generated), then their tensor product $M \otimes_R N$ is also a module of finite type.

Fact (Canonical Isomorphisms). The tensor product satisfies the following canonical isomorphisms for all R-modules M, N, Q:

1. Identity: $R \otimes_R N \cong N$.
2. Commutativity: $M \otimes_R N \cong N \otimes_R M$.
3. Associativity: $(M \otimes_R N) \otimes_R Q \cong M \otimes_R (N \otimes_R Q)$.
4. Distributivity: $(\bigoplus_{i\in I} M_i) \otimes_R N \cong \bigoplus_{i\in I}(M_i \otimes_R N)$.

Proof (Structural Isomorphisms).

1. For the identity, the map $R \times N \to N$ given by $(\lambda, n) \mapsto \lambda n$ is R-bilinear. By the universal property, it induces a homomorphism $\varphi : R \otimes_R N \to N$ satisfying $\varphi(\lambda \otimes n) = \lambda n$. This is an epimorphism since $1 \otimes n \mapsto n$. The inverse map is uniquely given by $n \mapsto 1 \otimes n$.
2. For commutativity, the map $M \times N \to N \otimes_R M$ given by $(m, n) \mapsto n \otimes m$ is bilinear. It induces a homomorphism $M \otimes_R N \to N \otimes_R M$. By symmetry, the reverse map exists and they are mutually inverse.
3. Associativity follows a similar universal property argument, requiring the definition of multilinear maps extending over the three components.
4. For distributivity, the map $(\bigoplus_{i\in I} M_i) \times N \to \bigoplus_{i\in I}(M_i \otimes_R N)$ given by $((m_i)_{i\in I}, n) \mapsto \sum_{i\in I} m_i \otimes n$ is bilinear and induces the requisite isomorphism.

□

Further Examples.

1. Ideals: Let $I, J \subset R$ be ideals. The tensor product of the quotient rings acts as the sum of the ideals:
$$R/I \otimes_R R/J \cong R/(I + J).$$
2. Coprime Integers: Let p, q be distinct prime numbers. Viewing $\mathbb{Z}/p\mathbb{Z}$ and $\mathbb{Z}/q\mathbb{Z}$ as $\mathbb{Z}$-modules, we have $\mathbb{Z}/p\mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Z}/q\mathbb{Z} = 0$. This holds because p and q are coprime; there exist integers a, b such that $ap + bq = 1$. For any $m \otimes n$, we can write $m \otimes n = (ap + bq)(m \otimes n) = a(pm \otimes n) + b(m \otimes qn) = a(0 \otimes n) + b(m \otimes 0) = 0$. Conversely, $\mathbb{Z}/p\mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Z}/p\mathbb{Z} \cong \mathbb{Z}/p\mathbb{Z}$.
3. Free Modules: By the distributivity property, the tensor product of free modules generated by sets S and T yields a free module generated by their Cartesian product: $R[S] \otimes_R R[T] \cong R[S \times T]$. In particular, if V and W are vector spaces over a field k of dimensions $\dim_k(V)$ and $\dim_k(W)$, then $\dim_k(V \otimes_k W) = \dim_k(V) \times \dim_k(W)$.

11.2 Flatness and the Tensor-Hom Adjunction

11.2.1 The Tensor-Hom Adjunction

Theorem (Tensor-Hom Adjunction). Let M, N, Q be R-modules. There exists a canonical natural isomorphism of R-modules:
$$\mathrm{Hom}_R(M \otimes_R N, Q) \cong \mathrm{Hom}_R(M, \mathrm{Hom}_R(N, Q)).$$

$$B_R(M,N;Q) \xleftrightarrow{\sim} \mathrm{Hom}_R(M,\mathrm{Hom}_R(N,Q))$$

$$\begin{array}{ccc} \Phi & \longmapsto & \Psi \\ \updownarrow & & \updownarrow \\ \Phi(m,n) & = & (\Psi(m))(n) \end{array}$$

Proof (Tensor-Hom Adjunction). By the universal property of the tensor product, we have an isomorphism $\mathrm{Hom}_R(M \otimes_R N, Q) \cong B_R(M,N;Q)$. A bilinear map $\Phi \in B_R(M,N;Q)$ is fully determined by its behavior when the first coordinate is fixed. For a fixed $m \in M$, the map $n \mapsto \Phi(m,n)$ is an element of $\mathrm{Hom}_R(N,Q)$. Thus, the map $m \mapsto \Phi(m,-)$ is a linear homomorphism from M into $\mathrm{Hom}_R(N,Q)$. The mapping $\Phi \mapsto (m \mapsto \Phi(m,-))$ is an isomorphism of modules $B_R(M,N;Q) \xrightarrow{\sim} \mathrm{Hom}_R(M,\mathrm{Hom}_R(N,Q))$.

□

Remark. The naturality of this isomorphism implies an adjunction between the functors $- \otimes_R N$ and $\mathrm{Hom}_R(N,-)$, where the tensor product acts as the left adjoint to the Hom functor.

11.2.2 Exactness Properties and Flat Modules

Lemma (Contravariant Hom and Exact Sequences). Let $M_1 \xrightarrow{f} M_2 \xrightarrow{g} M_3 \to 0$ be a sequence of R-modules. This sequence is exact if and only if for every R-module Q, the induced contravariant sequence:

$$0 \longrightarrow \mathrm{Hom}_R(M_3,Q) \xrightarrow{g^*} \mathrm{Hom}_R(M_2,Q) \xrightarrow{f^*} \mathrm{Hom}_R(M_1,Q)$$

is exact.

Proof (Contravariant Hom Exactness). Forward Implication ($\implies$): Assume $M_1 \xrightarrow{f} M_2 \xrightarrow{g} M_3 \to 0$ is exact.

1. *Exactness at* $\mathrm{Hom}_R(M_3,Q)$*:* Let $\varphi \in \mathrm{Hom}_R(M_3,Q)$ be in the kernel of g^*, meaning $g^*(\varphi) = \varphi \circ g = 0$. Since the original sequence is exact at M_3, the map g is surjective. A composition with a surjective map is zero if and only if the outer map is identically zero. Thus, $\varphi = 0$, proving g^* is injective.

2. *Exactness at* $\mathrm{Hom}_R(M_2,Q)$*:* We must show $\mathrm{Im}(g^*) = \ker(f^*)$. First, since $g \circ f = 0$, we have $f^* \circ g^* = (g \circ f)^* = 0$, implying $\mathrm{Im}(g^*) \subset \ker(f^*)$. Conversely, let $\psi \in \ker(f^*)$, so $\psi \circ f = 0$. This implies $\mathrm{Im}(f) \subset \ker(\psi)$. By the exactness of the original sequence, $\mathrm{Im}(f) = \ker(g)$, so $\ker(g) \subset \ker(\psi)$. By the universal property of quotients (since $M_3 \cong M_2/\ker(g)$), the map ψ factors through M_3. Thus, there exists a unique $\varphi \in \mathrm{Hom}_R(M_3,Q)$ such that $\psi = \varphi \circ g = g^*(\varphi)$. Therefore, $\ker(f^*) \subset \mathrm{Im}(g^*)$.

The reverse implication follows by strategically testing specific modules Q (e.g., setting $Q = M_3$ and $Q = \mathrm{Coker}(f)$) to force exactness constraints on the underlying modules.

□

Theorem (Right Exactness of the Tensor Product). Let N be an R-module. If the sequence $M_1 \xrightarrow{f} M_2 \xrightarrow{g} M_3 \longrightarrow 0$ is exact, then the tensored sequence:

$$M_1 \otimes_R N \xrightarrow{f \otimes \mathrm{id}_N} M_2 \otimes_R N \xrightarrow{g \otimes \mathrm{id}_N} M_3 \otimes_R N \longrightarrow 0$$

is also exact.

Proof (Right Exactness of the Tensor Product). By the preceding lemma, the original sequence is exact if and only if for all R-modules Q, the sequence:

$$0 \to \mathrm{Hom}_R(M_3,\mathrm{Hom}_R(N,Q)) \to \mathrm{Hom}_R(M_2,\mathrm{Hom}_R(N,Q)) \to \mathrm{Hom}_R(M_1,\mathrm{Hom}_R(N,Q))$$

is exact. By the Tensor-Hom Adjunction, we can substitute the terms with their canonical isomorphisms:

$$0 \to \mathrm{Hom}_R(M_3 \otimes_R N, Q) \to \mathrm{Hom}_R(M_2 \otimes_R N, Q) \to \mathrm{Hom}_R(M_1 \otimes_R N, Q).$$

Because this Hom sequence is exact for all R-modules Q, we apply the reverse direction of the lemma to conclude that the underlying sequence of tensor products is exact. This establishes that $- \otimes_R N$ is a right exact functor.

□

Definition (Flat Module). An R-module N is called *flat* if the functor $- \otimes_R N$ is exact. Equivalently, since the tensor product is inherently right exact, N is flat if and only if it preserves injections: for every injective R-module morphism $i : M_1 \hookrightarrow M_2$, the induced morphism $i \otimes \mathrm{id}_N : M_1 \otimes_R N \to M_2 \otimes_R N$ remains injective.

Examples and Counterexamples of Flatness.

1. Every free R-module is flat. By extension, any projective module is flat.
2. Over a Principal Ideal Domain (P.I.D.), an R-module is flat if and only if it is torsion-free. (Since the tensor product preserves inclusions locally, flat modules over well-behaved domains correspond geometrically to objects without torsion pathology).
3. The tensor product is not left exact in general. Consider the short exact sequence of $\mathbb{Z}$-modules:

   $$0 \longrightarrow \mathbb{Z}/2\mathbb{Z} \xrightarrow{\times 2} \mathbb{Z}/4\mathbb{Z} \longrightarrow \mathbb{Z}/2\mathbb{Z} \longrightarrow 0.$$

   Tensoring this entire sequence with $\mathbb{Z}/2\mathbb{Z}$ yields:

   $$\mathbb{Z}/2\mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Z}/2\mathbb{Z} \xrightarrow{(\times 2)\otimes \mathrm{id}} \mathbb{Z}/4\mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Z}/2\mathbb{Z} \longrightarrow \mathbb{Z}/2\mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Z}/2\mathbb{Z} \longrightarrow 0.$$

   Using the ideal tensor isomorphism $(R/I \otimes_R R/J \cong R/(I+J))$, this simplifies to:

   $$\mathbb{Z}/2\mathbb{Z} \xrightarrow{\times 2} \mathbb{Z}/2\mathbb{Z} \longrightarrow \mathbb{Z}/2\mathbb{Z} \longrightarrow 0.$$

   However, the map defined by multiplication by 2 on $\mathbb{Z}/2\mathbb{Z}$ is the zero map. It possesses a non-trivial kernel (the entire space $\mathbb{Z}/2\mathbb{Z}$), demonstrating that the induced map is not injective. Consequently, $\mathbb{Z}/2\mathbb{Z}$ is not a flat $\mathbb{Z}$-module.

End of Lec 11.

11.2.3 Induced Morphisms and the Dual Module

Let R be a commutative ring. The tensor product construction naturally extends to morphisms. If $f \in \mathrm{Hom}_R(M, M')$ and $g \in \mathrm{Hom}_R(N, N')$ are R-module homomorphisms, they induce a canonical R-module homomorphism between the respective tensor products:

$$\begin{aligned} f \otimes g : M \otimes_R N &\longrightarrow M' \otimes_R N' \\ m \otimes n &\longmapsto f(m) \otimes g(n). \end{aligned}$$

It is straightforward to verify that this map is well-defined by confirming that the assignment $(m, n) \mapsto f(m) \otimes g(n)$ is R-bilinear, which allows it to factor uniquely through the tensor product $M \otimes_R N$.
Recall that for any R-module M, tensoring with the base ring leaves the module unchanged up to canonical isomorphism, $M \otimes_R R \cong M$.

Definition (Dual Module). Let M be an R-module. We define the *dual module* of M, denoted M^*, as the module of R-linear forms on M:

$$M^* := \mathrm{Hom}_R(M, R).$$

To illustrate the behavior of the dual module, consider a finitely generated R-module M over a Principal Ideal Domain (P.I.D.) R. By the structure theorem, M decomposes into a torsion submodule and a free part, $M \cong M_{\mathrm{tor}} \oplus R^r$. The canonical evaluation map into the double dual, $M \to (M^*)^*$, exhibits the following exact sequence:

$$0 \longrightarrow M_{\mathrm{tor}} M \longrightarrow (M^*)^*,$$

where the image of M in $(M^*)^*$ is exactly the torsion-free quotient $M/M_{\mathrm{tor}} \cong R^r$, since linear forms annihilate all torsion elements.

Fact (The Canonical Map $\Theta_{M,N}$). Let M, N be R-modules. There exists a canonical R-module homomorphism from the tensor product of the dual module and N to the module of homomorphisms from M to N:

$$\Theta_{M,N} : M^* \otimes_R N \longrightarrow \mathrm{Hom}_R(M, N).$$

This morphism is constructed via the universal property of the tensor product applied to the R-bilinear map $M^* \times N \to \mathrm{Hom}_R(M, N)$ defined by:

$$(l, n) \longmapsto (m \mapsto l(m) \cdot n)\,.$$

Lemma (Isomorphism for Finite Free Modules). If M is a free R-module of finite type, then the canonical map $\Theta_{M,N}$ is an isomorphism for any R-module N.

Proof. We proceed by analyzing the case of a single generator and then extending via additivity. First, let $M = R$. The dual module is $R^* = \mathrm{Hom}_R(R, R) \cong R$, generated by the identity map Id_R. The left-hand side simplifies to $R^* \otimes_R N \cong R \otimes_R N \cong N$. The right-hand side simplifies to $\mathrm{Hom}_R(R, N) \cong N$. Under these natural identifications, the map $\Theta_{R,N}$ takes $1_{R^*} \otimes n$ to the homomorphism mapping $1_R \mapsto n$. This corresponds exactly to the identity map on N, meaning $\Theta_{R,N}$ is an isomorphism.
Now, let M be a free R-module of finite rank r, so $M \cong \bigoplus_{i=1}^r R$. Let $(e_1, \ldots, e_r)$ be a basis for M. The dual module is $M^* \cong \bigoplus_{i=1}^r R^*$, with the corresponding dual basis $(e_1^*, \ldots, e_r^*)$ where $e_i^*(e_j) = \delta_{ij}$. Because the tensor product distributes over finite direct sums, we have:

$$M^* \otimes_R N \cong \left(\bigoplus_{i=1}^r R^*\right) \otimes_R N \cong \bigoplus_{i=1}^r (R^* \otimes_R N)\,.$$

Similarly, the Hom functor distributes over finite direct sums in the first coordinate:

$$\mathrm{Hom}_R(M, N) \cong \mathrm{Hom}_R\left(\bigoplus_{i=1}^r R, N\right) \cong \bigoplus_{i=1}^r \mathrm{Hom}_R(R, N).$$

Under these direct sum decompositions, the map $\Theta_{M,N}$ breaks down entirely into the direct sum of the components:

$$\Theta_{M,N} = \bigoplus_{i=1}^r \Theta_{R,N}.$$

Since each component $\Theta_{R,N}$ is an isomorphism, their finite direct sum $\Theta_{M,N}$ is necessarily an isomorphism.

□

The finiteness condition on the rank of the free module is strictly required. If M is a free module of infinite rank, the map $\Theta_{M,N}$ fails to be surjective.

Proof (Failure for Infinite Rank). Let $M = N = \bigoplus_{i \in I} R$ be a free module where the index set I is infinite ($|I| = \infty$). We consider the canonical map into the endomorphism ring:

$$\Theta_{M,M} : M^* \otimes_R M \longrightarrow \mathrm{Hom}_R(M, M).$$

Suppose, for the sake of contradiction, that $\Theta_{M,M}$ is surjective. Then the identity morphism $\mathrm{Id}_M \in \mathrm{Hom}_R(M, M)$ must lie in the image. Therefore, there exists a finite sum of pure tensors $\sum_{k=1}^n l_k \otimes m_k \in M^* \otimes_R M$ such that:

$$\Theta_{M,M}\left(\sum_{k=1}^n l_k \otimes m_k\right) = \mathrm{Id}_M\,.$$

Evaluating this equality on an arbitrary element $x \in M$, we obtain:

$$x = \mathrm{Id}_M(x) = \sum_{k=1}^n l_k(x) \cdot m_k.$$

Because each $l_k(x)$ is a scalar in R, this equation states that every element $x \in M$ can be written as a linear combination of the finite set of vectors $\{m_1, \ldots, m_n\}$. This implies that the module M is finitely generated. However, we assumed M is a free module of infinite rank, yielding a direct contradiction. Thus, Id_M cannot be in the image, and $\Theta_{M,M}$ is not an isomorphism.

□

11.3 Extension of Scalars and Algebras

The tensor product provides a natural mechanism for changing the base ring of a module, allowing us to transport algebraic structures along ring homomorphisms.
Let $\varphi : A \to B$ be a homomorphism of commutative rings. Through this homomorphism, the ring B naturally acquires the structure of an A-module, where the scalar multiplication is defined by $a \cdot b := \varphi(a)b$. Furthermore, the internal ring multiplication of B, denoted $\mu_B : B \times B \to B$, is an A-bilinear map. This induces an A-module morphism out of the tensor product:

$$\mu_B : B \otimes_A B \longrightarrow B.$$

Now, let M be an arbitrary A-module. We can form the tensor product $B \otimes_A M$, which natively carries the structure of an A-module. We can upgrade this to a B-module.

Definition (Extension of Scalars). The tensor product $B \otimes_A M$ can be endowed with a B-module structure by defining the scalar action on pure tensors as:

$$x \cdot (y \otimes m) := (xy) \otimes m, \quad \text{for all } x, y \in B \text{ and } m \in M.$$

To formalize this, consider the map $B \times (B \otimes_A M) \to B \otimes_A M$ mapping $(x, \sum y_i \otimes m_i) \mapsto \sum (xy_i) \otimes m_i$. This map satisfies all the axioms of scalar multiplication. The resulting B-module is called the *extension of scalars* of M from A to B.

Examples of Extension of Scalars.

1. Quotient Rings: If $B = A/I$ where $I \subset A$ is an ideal, the extension of scalars acts as a quotient operation on the module. There is a canonical A/I-module isomorphism:

$$A/I \otimes_A M \cong M/IM.$$

2. Localization: Let $S \subset A$ be a multiplicative subset. The localized ring $S^{-1}A$ is an A-algebra. Extending scalars from A to $S^{-1}A$ is naturally isomorphic to localizing the module directly:

$$S^{-1}A \otimes_A M \cong S^{-1}M.$$

   As a specific instance, if A is an integral domain and $K = \mathrm{Frac}(A)$ is its fraction field, then $K \otimes_A M \cong S^{-1}M$ where $S = A \setminus \{0\}$.

   If we take $A = \mathbb{Z}$, $B = \mathbb{Z}/p\mathbb{Z}$ for a prime p, and $M = \mathbb{Q}$, we observe the annihilation of divisible modules upon tensoring with torsion modules:

$$\mathbb{Z}/p\mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Q} = 0.$$

   This holds because p is invertible in $\mathbb{Q}$. For any $1 \otimes q$, we can write $1 \otimes q = 1 \otimes (p \cdot \frac{q}{p}) = p \otimes \frac{q}{p} = 0 \otimes \frac{q}{p} = 0$.

11.3.1 Multilinear Maps and Adjunction

The concept of bilinear maps extends naturally to an arbitrary number of variables.

Definition (Multilinear Maps). Let $M_1, \ldots, M_n$ and N be R-modules. A mapping $\Phi : M_1 \times \cdots \times M_n \to N$ is called an *n-multilinear map* over R if it is R-linear in each of its n coordinates when all other coordinates are held constant.

By iteratively applying the universal property of the tensor product, we find that the module of n-multilinear maps is canonically isomorphic to the module of linear homomorphisms mapping from the n-fold tensor product:

$$n\text{-lin}_R(M_1 \times \cdots \times M_n; N) \cong \mathrm{Hom}_R(M_1 \otimes_R M_2 \otimes_R \cdots \otimes_R M_n, N).$$

If N is a B-module, we can always retrieve an A-module structure by restricting the scalars along the ring homomorphism $\varphi : A \to B$. We denote this induced A-module by N^{φ}, where the action is defined as $a \cdot n = \varphi(a)n$. This process is known as restriction of scalars. The extension of scalars acts as the left adjoint to this restriction.

Lemma (Adjunction of Extension and Restriction). Let $\varphi : A \to B$ be a commutative ring homomorphism. Let M be an A-module and N be a B-module. There is a canonical isomorphism of A-modules:

$$\mathrm{Hom}_{B\text{-mod}}(B \otimes_A M, N) \cong \mathrm{Hom}_{A\text{-mod}}(M, N^{\varphi}).$$

Proof. We construct the mutually inverse mappings explicitly.
Let $f \in \mathrm{Hom}_{B\text{-mod}}(B \otimes_A M, N)$ be a morphism of B-modules. We define a map $g : M \to N^{\varphi}$ by evaluating f on the canonical inclusion of M:

$$g(m) := f(1_B \otimes m).$$

We must verify that g is an A-module homomorphism. For $a \in A$ and $m \in M$, we have:

$$g(a \cdot m) = f(1_B \otimes (a \cdot m)) = f((1_B \cdot \varphi(a)) \otimes m) = f(\varphi(a)(1_B \otimes m)).$$

Because f is B-linear, we can pull the scalar $\varphi(a) \in B$ outside:

$$f(\varphi(a)(1_B \otimes m)) = \varphi(a) f(1_B \otimes m) = a \cdot g(m),$$

where the final equality uses the definition of the A-action on N^{φ}. Thus, g is A-linear.
Conversely, let $g \in \mathrm{Hom}_{A\text{-mod}}(M, N^{\varphi})$ be an A-module homomorphism. We define a mapping $B \times M \to N$ via $(b, m) \mapsto b \cdot g(m)$. We check that this mapping is A-bilinear. The linearity in B is trivial. For the linearity in M, considering the action of $a \in A$:

$$(b, a \cdot m) \longmapsto b \cdot g(a \cdot m) = b \cdot (\varphi(a) \cdot g(m)) = (b\varphi(a)) \cdot g(m).$$

This matches the evaluation on $(b \cdot \varphi(a), m)$, confirming A-bilinearity. By the universal property of the tensor product, this induces a unique A-linear map $\tilde{g} : B \otimes_A M \to N$ satisfying $\tilde{g}(b \otimes m) = b \cdot g(m)$. Finally, we verify that $\tilde{g}$ is in fact B-linear. Let $b' \in B$:

$$\tilde{g}(b' \cdot (b \otimes m)) = \tilde{g}((b'b) \otimes m) = (b'b) \cdot g(m) = b' \cdot (b \cdot g(m)) = b' \cdot \tilde{g}(b \otimes m).$$

The mappings $f \mapsto g$ and $g \mapsto \tilde{g}$ are directly inverse to each other, establishing the canonical bijection.

□

11.3.2 Tensor Product of Algebras

Let A be a commutative ring. We can take the tensor product of two A-algebras to form a new A-algebra. While the component algebras need not be commutative initially (provided the structure map from A lands within their centers), we will primarily focus on the commutative case.
Let B and C be A-algebras. We consider the A-module $B \otimes_A C$. We endow this module with a ring multiplication by constructing an A-multilinear map:

$$\begin{aligned} B \times C \times B \times C &\longrightarrow B \otimes_A C \\ (b, c, b', c') &\longmapsto (bb') \otimes (cc'). \end{aligned}$$

By the universal property of multilinear maps, this induces a well-defined A-bilinear operation on the tensor product, $(B \otimes_A C) \times (B \otimes_A C) \to B \otimes_A C$, providing the multiplication rule:

$$(b \otimes c) \cdot (b' \otimes c') := (bb') \otimes (cc').$$

Under this multiplication, $B \otimes_A C$ becomes an A-algebra.

Example (Matrix Algebras). Let k be a field. The tensor product of matrix algebras over k is isomorphic to a larger matrix algebra:

$$M_n(k) \otimes_k M_m(k) \cong M_{nm}(k).$$

For the remainder of this discussion, we assume that all rings and algebras are strictly commutative.

Proposition (Pushout in Commutative Algebras). Let A, B, and C be commutative rings. Let $A \to B$ and $A \to C$ be ring homomorphisms, making B and C into commutative A-algebras. Then the tensor product $B \otimes_A C$ is a commutative A-algebra, and it satisfies the universal property of the pushout (or coproduct) in the category of commutative A-algebras. Specifically, for any commutative A-algebra D, there is a canonical bijection:

$$\mathrm{Hom}_{A\text{-alg}}(B \otimes_A C, D) \cong \mathrm{Hom}_{A\text{-alg}}(B, D) \times \mathrm{Hom}_{A\text{-alg}}(C, D).$$

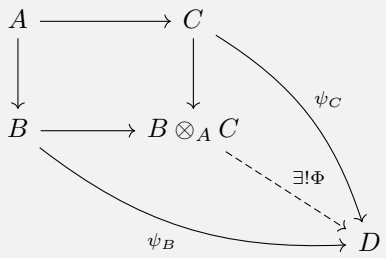

Geometrically, the Spec functor is contravariant and converts colimits in the category of rings into limits in the category of schemes. Therefore, the algebraic tensor product corresponds precisely to the geometric fiber product of the respective spectra:

$$\begin{array}{ccc} \mathrm{Spec}(B \otimes_A C) & \longrightarrow & \mathrm{Spec}(C) \\ \downarrow & \lrcorner & \downarrow \\ \mathrm{Spec}(B) & \longrightarrow & \mathrm{Spec}(A) \end{array}$$

yielding the topological identity:

$$\mathrm{Spec}(B \otimes_A C) \cong \mathrm{Spec}(B) \times_{\mathrm{Spec}(A)} \mathrm{Spec}(C).$$

End of Lec 12.

Appendix

Bibliography

Lecture Recordings

[1] R. E. Borcherds, *Galois Theory*, YouTube Lecture Series, https://youtube.com/playlist?list=PL8yHsr3EFj53Zxu3iRGMYL_89GDMvdkgt&si=tTjS7vNW9FIYM8Nj.

[2] R. E. Borcherds, *Commutative Algebra*, YouTube Lecture Series, https://youtube.com/playlist?list=PL8yHsr3EFj53rSexSz7vsYt-3rpHPR3HB&si=IjQa58LUzA7hKNQr.

[3] R. E. Borcherds, *Homological Algebra*, YouTube Lecture Series, https://youtube.com/playlist?list=PL8yHsr3EFj53tlhAM-ZfopRKNB541ZTi2&si=iyNoA2YuS638rgSJ.

[4] R. E. Borcherds, *Category Theory*, YouTube Lecture Series, https://youtube.com/playlist?list=PL8yHsr3EFj51F9XZ_Ka4bLnQoxTdMxoAL&si=oq_XZ-9IwQYxr2wW.

Modern and Standard Texts

[5] Atiyah, Michael F., and Macdonald, Ian G. *Introduction to Commutative Algebra*. Addison-Wesley Publishing Co., Reading, Mass., 1969.

[6] Eisenbud, David. *Commutative Algebra: with a View Toward Algebraic Geometry*. Graduate Texts in Mathematics, 150. Springer-Verlag, New York, 1995.

[7] Matsumura, Hideyuki. *Commutative Ring Theory*. Translated from the Japanese by M. Reid. Cambridge Studies in Advanced Mathematics, 8. Cambridge University Press, Cambridge, 1986.

[8] Serre, Jean-Pierre. *Local Algebra*. Translated from the French by CheeWhye Chin and revised by the author. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2000.

Classic and Foundational Texts

[9] Bourbaki, Nicolas. *Commutative Algebra. Chapters 1–7*. Elements of Mathematics. Springer-Verlag, Berlin, 1989.

[10] Grothendieck, Alexander, and Dieudonné, Jean. *Éléments de géométrie algébrique* (EGA). Publications Mathématiques de l'IHÉS, Bures-sur-Yvette, 1960–1967.

[11] Nagata, Masayoshi. *Local Rings*. Interscience Tracts in Pure and Applied Mathematics, No. 13. Interscience Publishers, New York-London, 1962.

[12] Zariski, Oscar, and Samuel, Pierre. *Commutative Algebra. Vol. I, II*. Graduate Texts in Mathematics, Vol. 28, 29. Springer-Verlag, New York-Heidelberg, 1975.

Historical Texts

[13] Macaulay, Francis S. *The Algebraic Theory of Modular Systems*. Cambridge Tracts in Mathematics and Mathematical Physics, No. 19. Cambridge University Press, Cambridge, 1916.

Index of Important Concepts

Index of Proofs and Proof Sketches

Index of Regular Concepts