Algebra I and II

Kai-Uwe Bux

May 11, 2007

Contents

Chapter 1

Groups and Actions

1.1 Basic Notions

1.1.1 Monoids

Definition 1.1.1.1 (monoid). A monoid is a set $\mathcal{M}$ together with

1. a distinguished element $1 \in \mathcal{M}$, called the identity element of $\mathcal{M}$, and
2. a binary operation in $\mathcal{M}$, called multiplication,

$$\begin{array}{rcl} \mathcal{M} \times \mathcal{M} & \longrightarrow & \mathcal{M} \\ (\mu, \nu) & \mapsto & \mu\nu, \end{array}$$

such that the following axioms are satisfied:

1. The element 1 is a two-sided identity, i.e., $1\mu = \mu = \mu 1$ for all $\mu \in \mathcal{M}$.
2. The law of associativity holds, i.e.: $(\mu_0\mu_1)\mu_2 = \mu_0(\mu_1\mu_2)$ for all $\mu_0, \mu_1, \mu_2 \in \mathcal{M}$.

Observation 1.1.1.2. *In a monoid, the identity element is uniquely determined by the binary operation: suppose 1_1 and 1_2 would both qualify, then*

$$1_1 = 1_1 1_2 = 1_2.$$

Thus, to specify a monoid, we only have to declare the multiplication operation and confirm the existence of an identity element.

Example 1.1.1.3. The natural numbers including 0 form a monoid $\mathbb{N}_+$ with respect to addition where 0 serves as the identity element.

Example 1.1.1.4. The natural numbers including also form a monoid $\mathbb{N}_\times$ with respect to multiplication where 1 serves as the identity element.

Definition 1.1.1.5. Let $\mathcal{M}$ and $\mathcal{N}$ be monoids. A map $\varphi : \mathcal{M} \to \mathcal{N}$ is called a homomorphism if:

1. The map φ preserves the identity element, i.e.,

$$\varphi(1_{\mathcal{M}}) = 1_{\mathcal{N}}.$$

2. The map φ is compatible with multiplication, i.e.,

$$\varphi(\mu\nu) = \varphi(\mu)\,\varphi(\nu)$$

for all $\mu, \nu \in \mathcal{M}$.

A surjective homomorphism is called an epimorphism, an injective (1-1) homomorphism is called a monomorphism, and an invertible homomorphism (i.e., it has an inverse that is also a homomorphism) is called an isomorphism. A homomorphism from $\mathcal{M}$ (in)to $\mathcal{M}$ is called an endomorphism and an invertible endomorphism is called an automorphism.

Exercise 1.1.1.6. Show that a monoid homomorphism is an isomorphism if and only if it is 1-1 and onto.

Exercise 1.1.1.7. The composition of two monoid homomorphisms is a monoid homomorphism.

Example 1.1.1.8. The map

$$\begin{array}{ccc} \mathbb{N}_+ & \longrightarrow & \mathbb{N}_\times \\ n & \mapsto & 2^n \end{array}$$

is a monoid-monomorphism.

Example 1.1.1.9 (free monoid). Let X be a set. The free monoid with basis X is the monoid $X^* := \{(x_1, x_2, \ldots, x_u) \mid u \geq 0,\ x_i \in X\}$ consisting of all finite sequences (ordered tuples of all possible lengths) with entries from X. We consider X as an alphabet and such tuple as a word. The binary operation is concatenation of words. The identity element is the empty word of length 0.

To simplify notation, we shall represent the finite sequence $(x_1, x_2, \ldots, x_u)$ simply as $x_1 x_2 \cdots x_u$. Note that this notation fits perfectly: one can regard X as a subset of X^* by regarding every letter of the alphabet X as a one-letter word in X^*. Then the expression $x_1 x_2 \cdots x_u$ can be interpreted two ways: we can regard it as a shorthand for a certain u-letter word, but we can also regard it as a product of u one-letter words. In fact, any grouping allows us to read a product structure into this expression. Since concatenation is associative, all those readings refer to the same element of X^*.

Theorem 1.1.1.10 (universal property of free monoids). *Let X be a set. For any monoid $\mathcal{M}$ and any map $f : X \to \mathcal{M}$, there exists a unique monoid homomorphism $\varphi : X^* \to \mathcal{M}$ such that*

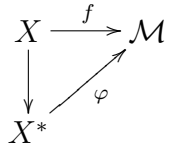

commutes, where the vertical arrow is the canonical inclusion $X \hookrightarrow X^*$.

Proof. Note that we really have no choice but to define

$$\begin{array}{rcl} \varphi : X^* & \longrightarrow & \mathcal{M} \\ x_1 x_2 \cdots x_u & \mapsto & f(x_1)\, f(x_2) \cdots f(x_u) \end{array}$$

Now, we just need to observe that this is a monoid homomorphism. **q.e.d.**

Example 1.1.1.11. Let X be a set. The set $\mathrm{End}_{\mathrm{set}}(X) := \mathrm{Maps}_{\mathrm{set}}(X;X) := \{f : X \to X\}$ of all maps from X to X is a monoid with respect to composition of maps. The sets $\mathrm{Mono}_{\mathrm{set}}(X;X) := \{f : X \to X \mid f \text{ is injective}\}$ and $\mathrm{Epi}_{\mathrm{set}}(X;X) := \{f : X \to X \mid f \text{ is surjective}\}$ are two important submonoids. Their intersection is the monoid $\mathrm{Perm}(X) := \{f : X \to X \mid f \text{ is bijective}\}$.

Example 1.1.1.12. Similarly, in any category, the set of endo-morphisms of an an object is a monoid; it contains as submonoids the set of endo-monomorphism, the set of endo-epimorophisms, and the set of automorphisms.

For instance, for any monoid $\mathcal{M}$, the set

$$\mathrm{End}_{\mathrm{mon}}(\mathcal{M}) = \{\varphi : \mathcal{M} \to \mathcal{M} \mid \varphi \text{ is a monoid homomorphism}\}$$

of all homomorphisms is a monid with respect to composition. It containing the submonoids $\mathrm{Mono}_{\mathrm{mon}}(\mathcal{M};\mathcal{M})$ and $\mathrm{Epi}_{\mathrm{mon}}(\mathcal{M};\mathcal{M})$ and their intersection $\mathrm{Aut}_{\mathrm{mon}}(\mathcal{M}) = \{\varphi : \mathcal{M} \to \mathcal{M} \mid \varphi \text{ is a monoid isomorphism}\}$.

Definition 1.1.1.13 (action). Let $\mathcal{M}$ be a group and let X be a set. A left-action of $\mathcal{M}$ on X is a map

$$\begin{array}{rcl} \alpha : \mathcal{M} \times X & \longrightarrow & X \\ (\mu, x) & \mapsto & \mu \ltimes_\alpha x \end{array}$$

satisfying the following axioms:

1. For all $x \in X$,
$$1 \ltimes_\alpha x = x.$$

2. For all $\mu, \nu \in \mathcal{M}$ and all $x \in X$,

$$\mu \ltimes_\alpha (\nu \ltimes_\alpha x) = (\mu\nu) \ltimes_\alpha x.$$

As for notation, we will drop the subscript α whenever there is no reasonable doubt about which action is meant.

Example 1.1.1.14 (trivial action). Let $\mathcal{M}$ be a monoid. Any set X can be endowed with the trivial $\mathcal{M}$-action:

$$\begin{array}{ccc} \mathcal{M} \times X & \longrightarrow & X \\ (\mu, x) & \mapsto & x \end{array}$$

Example 1.1.1.15 (left multiplication). The multiplication map

$$\begin{array}{ccc} \mathcal{M} \times \mathcal{M} & \longrightarrow & \mathcal{M} \\ (\mu, \nu) & \mapsto & \mu\nu \end{array}$$

is an action of the monoid $\mathcal{M}$ on the set underlying $\mathcal{M}$. Checking the axioms is straight forward. This action is called left multiplication.

Example 1.1.1.16 (tautological action). Let X be a set. Then the monoid $\mathrm{End}_{\mathrm{set}}(X) := \mathrm{Maps}_{\mathrm{set}}(X; X)$ acts on X via

$$\begin{array}{ccc} \mathrm{End}_{\mathrm{set}}(X) \times X & \longrightarrow & X \\ (f, x) & \mapsto & f(x) \end{array}$$

Definition 1.1.1.17. Let $\mathcal{M}$ be a monoid. An $\mathcal{M}$-set is a set X together with a specified $\mathcal{M}$-action.

Let X and Y be two $\mathcal{M}$-sets. A map

$$h : X \to Y$$

is a $\mathcal{M}$-map if for every $\mu \in \mathcal{M}$ and every $x \in X$,

$$h(\mu \ltimes x) = \mu \ltimes h(x)\,.$$

Let $\mathcal{N}$ be another monoid and let $\varphi : \mathcal{M} \to \mathcal{N}$ be a homomorphism. Now, let X be an $\mathcal{M}$-set and let Y be an $\mathcal{N}$-set. A map

$$h : X \to Y$$

is φ-equivariant if for every $\mu \in \mathcal{M}$ and every $x \in X$,

$$h(\mu \ltimes x) = \varphi(\mu) \ltimes h(x)\,.$$

Observation 1.1.1.18 (universal property of the tautological action). *Let X be a set. For any monoid $\mathcal{M}$ and any action $\alpha : \mathcal{M} \times X \to X$ of $\mathcal{M}$ on X, there exits a unique monoid homomorphism*

$$\begin{array}{rcl} \varphi : \mathcal{M} & \longrightarrow & \mathrm{End}_{\mathrm{set}}(X) \\ \mu & \mapsto & f_\mu \end{array}$$

such that for every $\mu \in \mathcal{M}$ and every $x \in X$:

$$\mu \ltimes x = f_\mu(x)$$

Note that this means the identity map $\mathrm{id}_X : X \to X$ is φ-equivariant. **q.e.d.**

Exercise 1.1.1.19. Let $\varphi : \mathcal{M} \to \mathcal{N}$ be a monoid homomorphism and let

$$\begin{array}{rcl} \alpha : \mathcal{N} \times X & \longrightarrow & X \\ (\nu, x) & \mapsto & \nu \ltimes x \end{array}$$

be a left-action of $\mathcal{N}$ on the set X. Show that

$$\begin{array}{rcl} \alpha_\varphi : \mathcal{M} \times X & \longrightarrow & X \\ (\mu, x) & \mapsto & \varphi(\mu) \ltimes x \end{array}$$

is a left-action of $\mathcal{M}$ on X.

Exercise 1.1.1.20. Let $\mathcal{M}$ be a monoid and let X be a set. Show that there is a 1-1 correspondence

$$\{\text{left-actions of } \mathcal{M} \text{ on } X\} \longleftrightarrow \{\text{monoid homomorphisms } \mathcal{M} \to \mathrm{End}_{\mathrm{set}}(X)\}\,.$$

Exercise 1.1.1.21. Let $\mathcal{M}$ act on a set X, and observe that the relation $\sim$ defined on X by

$x \sim y$ if and only if there exist $z \in X$, $\mu, \nu \in \mathcal{M}$ with $x = \mu z$ and $y = \nu z$

is reflexive and symmetric.

1. Let $\equiv$ be the transitive closure of $\sim$, let ${}_{\mathcal{M}}\backslash^{X}$ be the set of $\equiv$-equivalence classes in X, and let $\pi : X \to {}_{\mathcal{M}}\backslash^{X}$ be the canonical projection. Show that π is an $\mathcal{M}$-map when we regard ${}_{\mathcal{M}}\backslash^{X}$ as a $\mathcal{M}$-set with the trivial $\mathcal{M}$-action.

2. Let Y be a set and consider it as a $\mathcal{M}$-set with the trivial $\mathcal{M}$-action. Show that for any $\mathcal{M}$-map $f : X \to Y$, there exists a unique map $f_* : {}_{\mathcal{M}}\backslash^{X} \to Y$ such that

$$\begin{array}{ccc} X & \xrightarrow{f} & Y \\ {\scriptstyle\pi}\downarrow & \nearrow_{f_*} & \\ {}_{\mathcal{M}}\backslash^{X} & & \end{array}$$

   commutes. (This is a universal property.)

3. Let $\mathcal{N}$ be another monoid and let $\varphi : \mathcal{M} \twoheadrightarrow \mathcal{N}$ be a monoid epimorphism. Show that there exists a pair $(X_\varphi, \pi : X \to X_\varphi)$ where X_φ is an $\mathcal{N}$-set and π is a φ-equivariant map so that the following universal property holds:

   > For any $\mathcal{N}$-set Y and any φ-equivariant map $f : X \to Y$ there exists a unique $\mathcal{N}$-map $f_\varphi : X_\varphi \to Y$ such that
   >
   > $$\begin{array}{ccc} X & \xrightarrow{f} & Y \\ {\scriptstyle\pi}\downarrow & \nearrow_{f_\varphi} & \\ X_\varphi & & \end{array}$$
   >
   > commutes.

4. Suppose $(X_1, \pi_1 : X \to X_1)$ and $(X_2, \pi_2 : X \to X_2)$ are two pairs satisfying the universal property above. Show that there is a unique $\mathcal{N}$-isomorphism

$$\psi : X_1 \longrightarrow X_2$$

such that

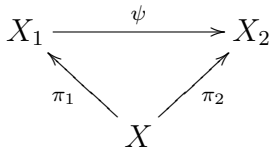

commutes.

Exercise 1.1.1.22. A monoid $\mathcal{M}$ is called left-cancellative if for any $\mu, \nu_1, \nu_2 \in \mathcal{M}$ the implication

$$\text{if } \mu\nu_1 = \mu\nu_2 \text{ then } \nu_1 = \nu_2$$

holds.

1. Show that the free monoid X^* over any set X is left-cancellative.
2. Show that $\mathrm{Mono}_{\mathrm{set}}(X)$ is left-cancellative for any set X.
3. Show that $\mathrm{Epi}_{\mathrm{set}}(X)$ is not left-cancellative for any infinite set X.
4. Show that $\mathrm{End}_{\mathrm{set}}(X)$ is not left-cancellative for any set X that contains at least two elements.
5. Show that a submonoid of a left-cancellative monoid is left-cancellative.
6. Show that $\mathcal{M}$ is left-cancellative if and only if the map

   $$\mathcal{M} \to \mathrm{End}_{\mathrm{set}}(\mathcal{M})$$

   induced by the left-multiplication action has image within $\mathrm{Mono}_{\mathrm{set}}(\mathcal{M})$.

1.1.2 Group Actions

Definition 1.1.2.1 (group). A (left-)group is a set G together with

1. a distinguished element $1 \in G$, called the identity element of G, and

2. a binary operation in G, called multiplication,

$$\begin{aligned} G \times G &\longrightarrow G \\ (g,h) &\mapsto gh, \end{aligned}$$

such that the following axioms are satisfied:

1. The element 1 is a left-identity, i.e., $1g = g$ for all $g \in G$.
2. For each $g \in G$ there is a left-inverse, i.e., an element $g^{\mathrm{l}} \in G$ such that $g^{\mathrm{l}}g = 1$.
3. The law of associativity holds, i.e.: $(g_0g_1)g_2 = g_0(g_1g_2)$ for all $g_0, g_1, g_2 \in G$.

Remark 1.1.2.2. It is not obvious, yet, that groups are monoids: we have not required the identity to be two-sided. However, we shall see later that it actually follows.

Example 1.1.2.3 (permutations). Let X be a set. The symmetric group over X is the set

$$\mathrm{Perm}(X) := \{\sigma : X \to X \,|\, \sigma \text{ is a bijection}\}$$

together with the identity element

$$\begin{aligned} \mathrm{id}_X : X &\longrightarrow X \\ x &\mapsto x \end{aligned}$$

and the multiplication

$$\begin{aligned} \mathrm{Perm}(X) \times \mathrm{Perm}(X) &\longrightarrow \mathrm{Perm}(X) \\ (\sigma, \tau) &\mapsto \sigma \circ \tau \end{aligned}$$

where the composition $\sigma \circ \tau$ is defined by

$$\begin{aligned} \sigma \circ \tau : X &\longrightarrow X \\ x &\mapsto \sigma(\tau(x)). \end{aligned}$$

With these definitions, $\mathrm{Perm}(X)$ is a group.

The group $\mathrm{Perm}(\{1, 2, \ldots, r\})$ is commonly denoted by $\mathbf{S}_r$.

Definition 1.1.2.4 (action). Let G be a group and let X be a set. A left-action of G on X is a map

$$\begin{aligned} \alpha : G \times X &\longrightarrow X \\ (g, x) &\mapsto g \ltimes_\alpha x \end{aligned}$$

satisfying the following axioms:

1. For all $x \in X$,
$$1 \ltimes_\alpha x = x.$$
2. For all $g, h \in G$ and all $x \in X$,
$$g \ltimes_\alpha (h \ltimes_\alpha x) = (gh) \ltimes_\alpha x.$$

As for notation, we will drop the subscript α whenever there is no reasonable doubt about which action is meant.

Remark 1.1.2.5. Once we see that groups are monoids, we can say that a group action is the same as an action of the monoid that the group happens to be.

Example 1.1.2.6 (trivial action). Let G be a group and let X be a set. The trivial action of G on X is given by

$$\begin{aligned} G \times X &\longrightarrow X \\ (g, x) &\mapsto x \end{aligned}$$

In this case, we also say that G acts trivially on X.

Remark 1.1.2.7. Groups are meant to act: almost everything we can prove in group theory is proved by looking at a particular action.

Example 1.1.2.8 (tautological action). Let X be a set. Then the symmetric group $\mathrm{Perm}(X)$ acts on X by evaluation

$$\begin{aligned} \mathrm{ev} : \mathrm{Perm}(X) \times X &\longrightarrow X \\ (\sigma, x) &\mapsto \sigma \ltimes x := \sigma(x) \end{aligned}$$

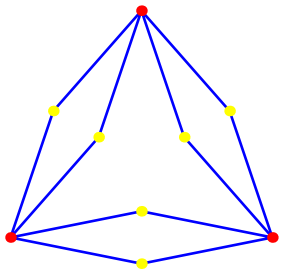

Table 1.1: A highly symmetric graph

Observation and Definition 1.1.2.9 (orbit). *Let α be a left-action of G on X. Then the relation*

$$x \sim_\alpha y :\Longleftrightarrow x = g \ltimes_\alpha y \text{ for some } g \in G$$

is an equivalence relation. The equivalence classes are called orbits of the given action. The orbit of x is denoted by $\mathrm{Orb}(x)$ if the group action is understood. Otherwise, we use a more explicit notation like $\mathrm{Orb}_G(x)$ if we want to stress the group, or $\mathrm{Orb}_\alpha(x)$ if we want to stress the action.

Definition 1.1.2.10. An action with exactly one orbit is called transitive.

Example 1.1.2.11. The tautological action of $\mathrm{Perm}(X)$ on X is transitive.

Example 1.1.2.12. Groups often arise as symmetry groups. For instance, consider the group G of symmetries of the graph[1] Γ given in table 1.1. Note that G naturally acts on several sets:

[1] A graph is a simple combinatorial structure: it is a set (whose elements are called vertices) together with a collection of 2-element subsets (those are called edges). A morphism of graphs is just a map sending vertices of one graph to vertices of another graph that does not tear apart edges. That's it, no strings attached. Note, that the way a graph is drawn is not part of the structure; so in the example, a symmetry (i.e., an automorphism of the graph) is allowed to swap inner yellow vertices and outer yellow vertices.

1. the set of vertices Γ
2. the set of edges in the graph (symmetries do not tear apart edges!)
3. the set of degree 4 vertices in Γ (symmetries preserve valencies!)
4. the set of degree 2 vertices in Γ

Exercise 1.1.2.13. How many elements has the full symmetry group of the graph Γ, and which of the four actions above are transitive?

1.1.3 Homomorphisms

Definition 1.1.3.1. Let G and H be groups. A map $\varphi : G \to H$ is called a homomorphism if:

1. The map φ preserves the identity element, i.e.,
$$\varphi(1_G) = 1_H.$$
2. The map φ is compatible with multiplication, i.e.,
$$\varphi(gh) = \varphi(g)\,\varphi(h)$$
for all $g, h \in G$.

A surjective homomorphism is called an epimorphism, an injective (1-1) homomorphism is called a monomorphism, and an invertible homomorphism (i.e., it has an inverse that is also a homomorphism) is called an isomorphism. A homomorphism from G (in)to G is called an endomorphism and an invertible endomorphism is called an automorphism.

Remark 1.1.3.2. Note that we did not require that homomorphisms preserve inverses. Thus, a map $\varphi : G \to H$ from a group G to a group H is a group homomorphism if and only if it is a monoid homomorphism.

Observation 1.1.3.3. *The composition of two group homomorphisms is a group homomorphism.* **q.e.d.**

Proposition 1.1.3.4. *Let α be a left-action of the group G on the set X. Then, for each $g \in G$, the map*

$$\begin{array}{rcl} \alpha_g : X & \longrightarrow & X \\ x & \mapsto & g \ltimes_\alpha x \end{array}$$

is a bijection. Moreover, the map

$$\begin{array}{rcl} \tilde{\alpha} : G & \longrightarrow & \mathrm{Perm}(X) \\ g & \mapsto & \alpha_g \end{array}$$

is a homomorphism.

Conversely, any homomorphism $\varphi : G \to \mathrm{Perm}(X)$ induces an action

$$\begin{array}{rcl} \alpha_\varphi : G \times X & \longrightarrow & X \\ (g, x) & \mapsto & g \ltimes x := \varphi(g)(x) \end{array}$$

The two constructions are inverses of one another and yield a 1-1 *correspondence:*

$$\{\textit{left-actions of } G \textit{ on } X\} \longleftrightarrow \{\textit{homomorphisms } G \to \mathrm{Perm}(X)\}\,.$$

Proof. First observe that for any $g \in G$ we have

$$\alpha_{g^{-1}} \circ \alpha_g = \mathrm{id}_X$$

whence α_g is 1-1: the first map in a composition has to be 1-1 if the composite map is 1-1. Since this applies to any group element, $\alpha_{g^{-1}}$ is 1-1 as well.

Now, we see that α_g is onto: Suppose $x \in X$ is not in the image of α_g. Since $\alpha_{g^{-1}}$ is 1-1, the element $\alpha_{g^{-1}}(x)$ cannot be in the image of $\alpha_{g^{-1}} \circ \alpha_g$ which is absurd since this composition is the identity on X.

The other statements follow exactly as in the monoid case. **q.e.d.**

Definition 1.1.3.5. An action α of G on X is faithful if 1 is the only element G that acts trivially on X, i.e., if $\alpha_g = \mathrm{id}_X$ then $g = 1$.

Exercise 1.1.3.6. Show that an action α is faithful if and only if the associated homomorphism $\tilde{\alpha}$ is a monomorphism.

Exercise 1.1.3.7. Which of the actions from Example 1.1.2.12 are faithful?

Proposition 1.1.3.8. *Let α be an action of H on X, and let $\varphi : G \to H$ be a homomorphism. Then G acts on X by*

$$\begin{array}{rcl} G \times X & \longrightarrow & X \\ (g, x) & \mapsto & \varphi(g) \ltimes_\alpha x \end{array}$$

Proof. Easy. **q.e.d.**

1.1.4 Example: Left Multiplication

Let G be a group. Then

$$\begin{array}{rcl} \lambda : G \times G & \longrightarrow & G \\ (g, h) & \mapsto & g \ltimes_\lambda h := gh \end{array}$$

is a left-action of G on itself: the group axioms turn directly into the axioms for a left action.

Definition 1.1.4.1. The action λ defined thus is called the left-multiplication action of G on itself.

Consider the corresponding homomorphism $\lambda : G \to \mathrm{Perm}(G)$. It follows that for each $g \in G$, we have $\lambda_g \in \mathrm{Perm}(G)$. In particular, for any $h \in G$, we have:

$$1 = gh \Longleftrightarrow 1 = \lambda_g(h)$$

However, since λ_g is a bijection, there is one and only one $h \in G$ satisfying $1 = gh$. This proves existence and uniqueness of right-inverses.

Let us denote the right-inverse of g by g^{r}. Let g^{l} be any left-inverse. Then we have:

$$\lambda_g \circ \lambda_{g^{\mathrm{r}}} = \mathrm{id} = \lambda_{g^{\mathrm{l}}} \circ \lambda_g$$

whence

$$\lambda_{g^{\mathrm{r}}} = \lambda_g^{-1} = \lambda_{g^{\mathrm{l}}}.$$

In particular

$$g^{\mathrm{r}} g = \lambda_{g^{\mathrm{r}}}(g) = \lambda_{g^{\mathrm{l}}}(g) = 1.$$

I.e., the right-inverse will also be a left-inverse.

Furthermore, this implies:

$$g1 = gg^{\mathrm{r}}g = 1g = g.$$

I.e., the left-identity 1 also serves as a right-identity.

Remark 1.1.4.2. We have now established the right-handed versions of the axioms. Thus, running through the same arguments with left and right reversed, we see that left-inverses also must be unique.

Proposition 1.1.4.3 (Cayley). *The left-multiplication of G in itself is faithful.*

Proof. Easy: 1 is a right-identity. **q.e.d.**

1.1.5 Right-actions

Definition 1.1.5.1 (right-action). Let G be a group and let X be a set. A right-action of G on X is a map

$$\begin{array}{rcl} \alpha : X \times G & \longrightarrow & X \\ (x, g) & \mapsto & x \ltimes_\alpha g \end{array}$$

satisfying the following axioms:

1. For all $x \in X$,

$$x \ltimes_\alpha 1 = x.$$

2. For all $g, h \in G$ and all $x \in X$,

$$x \ltimes_\alpha (gy) = (x \ltimes_\alpha g) \ltimes_\alpha h.$$

As for notation, we will drop the subscript α whenever there is no reasonable doubt about which action is meant.

Remark 1.1.5.2. Since we have established right-handed versions of the group axioms, we see that all considerations for left-actions apply mutatis mutandis to right-actions. In particular, every right on a set X induces a decomposition of X into disjoint orbits.

Example 1.1.5.3 (right-multiplication action). Multiplication

$$\begin{aligned} \rho : G \times G &\longrightarrow G \\ (g, h) &\mapsto g \rtimes_\rho h := gh \end{aligned}$$

also defines a right-action of G on itself. Moreover, for each $g \in G$, the map

$$\begin{aligned} \rho_g : G &\longrightarrow G \\ h &\mapsto hg \end{aligned}$$

is a bijections. However, the map

$$\begin{aligned} \rho : G &\longrightarrow \mathrm{Perm}(G) \\ g &\mapsto \rho_g \end{aligned}$$

is not a homomorphism. It is an anti-homomorphism.

Definition 1.1.5.4. A map $\varphi : G \to H$ is called a anti-homomorphism if:

1. The map φ preserves the identity element, i.e.,

$$\varphi(1_G) = 1_H.$$

2. The map φ swaps the order of multiplication, i.e.,

$$\varphi(gh) = \varphi(h)\,\varphi(g)$$

for all $g, h \in G$.

Proposition 1.1.5.5. *If α is a right-action of G on X, then the map*

$$\begin{array}{rcl} \tilde{\alpha} : G & \longrightarrow & \mathrm{Perm}(X) \\ g & \mapsto & \alpha_g \end{array}$$

is an anti-homomorphism where:

$$\begin{array}{rcl} \alpha_g : X & \longrightarrow & X \\ x & \mapsto & x \rtimes_\alpha g. \end{array}$$

Proof. Easy. **q.e.d.**

Exercise 1.1.5.6. Show that inversion $g \mapsto \bar{g}$ is an anti-automorphism of G.

Exercise 1.1.5.7. Show that the composition of two anti-homomorphisms is a homomorphism.

Corollary and Definition 1.1.5.8. *Let $\alpha : G \times X \to X$ be a left action. Then*

$$\begin{array}{rcl} \bar{\alpha} : X \times X & \longrightarrow & X \\ (x, g) & \mapsto & x \rtimes_{\bar{\alpha}} g := \bar{g} \ltimes_\alpha x \end{array}$$

is a right-action. It is called the right-action <u>induced</u> by α.

Proof. Easy: inverting is an anti-automorphism. **q.e.d.**

1.1.6 Stabilizers, Subgroups and Cosets

Definition 1.1.6.1. A subset U of a group G is a subgroup if it contains the identity element and is closed with respect to taking inverses and products, i.e.:

1. The identity element belongs to U.
2. The inverse of any element in U lies in U.
3. For any two elements $g, h \in U$, we have $gh \in U$.

Observation 1.1.6.2. *Subgroups are groups.*

Example 1.1.6.3. The image of any homomorphism is a subgroup in its range.

Example 1.1.6.4. The kernel of a homomorphism $\varphi : G \to H$ is defined as

$$\ker(\varphi) := \{g \in G \,|\, \varphi(g) = 1_H\}\,.$$

The kernel is a subgroup of G.

Example 1.1.6.5. The intersection of a family of subgroups is a subgroup.

Definition 1.1.6.6. Let $A, B \subseteq G$ be subsets. We put:

$$\begin{aligned} AB &:= \{ab \,|\, a \in A,\ b \in B\} \\ \bar{A} &:= \{\bar{a} \,|\, a \in A\} \end{aligned}$$

Exercise 1.1.6.7. Let A be a non-empty subset of G. Show that A is a subgroup if and only if $\bar{A}A \subseteq A$.

Exercise 1.1.6.8. Let U and V be subgroups. Show that UV is a subgroup if and only if $UV = VU$.

Proposition and Definition 1.1.6.9. *Let α be an action of G on X. For each element $x \in X$, the stabilizer*

$$\mathrm{Stab}(x) := \{g \in G \mid g \ltimes_\alpha x \in A\}$$

is a subgroup of G. We also use the notation $\mathrm{Stab}_\alpha(x)$ and $\mathrm{Stab}_G(x)$ if there is a need to be more specific.

Proof. Easy. **q.e.d.**

Example 1.1.6.10. Consider the tautological action of $\mathrm{Perm}(X)$ on X. The stabilizer of $x \in X$ is

$$\mathrm{Stab}(x) = \{\sigma \in \mathrm{Perm}(X) \mid \sigma(x) = x\}.$$

Note that such bijections σ restrict to bijections on $X - \{x\}$. Thus, $\mathrm{Stab}(x)$ acts on $X - \{x\}$.

Exercise 1.1.6.11. Show that restriction to $X - \{x\}$ induces an isomorphism of groups:

$$\mathrm{Stab}_{\mathrm{Perm}(X)}(x) \longrightarrow \mathrm{Perm}(X - \{x\}).$$

Discussion 1.1.6.12. Let G be a group and $U \leq G$ be a subgroup.

1. Multiplication induces a left- and a right-action of U on G:

$$\begin{array}{rcl} \lambda^U : U \times G & \longrightarrow & G \\ (u, g) & \mapsto & ug \end{array}$$

and:

$$\begin{array}{rcl} \rho^U : G \times U & \longrightarrow & G \\ (g, u) & \mapsto & gu \end{array}$$

The orbits of the left-action are called right-cosets and the orbit of the right-action are called left-cosets.

2. Now, fix an element $g \in G$. The right-coset of g is $Ug = \{ug \mid u \in U\} = \{\rho_g(u) \mid u \in U\} = \rho_g(U)$ and its left-coset is $gU = \lambda_g(U)$. Since ρ_g and λ_g are permutations of G, we find that they induce bijections

$$gU \xleftarrow{\lambda_g} U \xrightarrow{\rho_g} Ug.$$

In particular, all cosets (left and right) have the same cardinality as U.

3. Consider the inversion anti-automorphism in G. Since subgroups are closed with respect to taking inverses, it leaves the subset U invariant. Since anti-automorphisms interchange left and right, inversion takes right-cosets of U to left-cosets and vice versa. Thus, inversion induces a 1-1 correspondence

$$\left.{}_U\backslash^G := \{\text{right-cosets of } U\} \longleftrightarrow \{\text{left-cosets of } U\} =: {}^G/_U.\right.$$

In particular, both sets have the same cardinality.

4. The quotients ${}_U\backslash^G$ and ${}^G/_U$ are not structureless sets. In fact, the left-multiplication action of G on itself is compatible with the partition of G into left-cosets of U. Thus, we have an induced left-action of G on ${}^G/_U$. With respect to this action, we have:

$$\mathrm{Stab}_G(U) = U.$$

This is tricky to read: on the left hand, U represents an element of ${}^G/_U$, i.e., a left-coset of U; recall that every subgroup is a left-coset of itself.

Analogously, the orbit space ${}_U\backslash^G$ carries a right-action induced by the right-multiplication action.

Definition 1.1.6.13. The number of (left-)cosets of U in G is called the index of U. It is denoted by $[G : U]$.

Thus:

Lagrange's Theorem 1.1.6.14. *If G is a finite group, then*

$$\operatorname{card}(G) = [G : U] \operatorname{card}(U) .$$

In particular, the order and index of any subgroup divide the order of G.

Example 1.1.6.15 (cyclic and dihedral groups). Let X_5 be a set of five elements. Fix a cyclic order on X_5 as represented by the following directed graph:

$\Delta :=$ 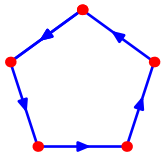

The cyclic group $\mathbf{C}_5$ can be realized as a subgroup of $\operatorname{Perm} X_5$ as follows:

$$\mathbf{C}_5 = \operatorname{Aut}(\Delta) = \{\sigma \in \operatorname{Perm}(X_5) \mid \sigma \text{ preserves the graph } \Delta\} .$$

Now ditch the orientations of edges:

$\Gamma :=$ 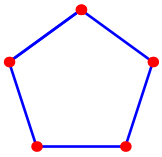

The dihedral group of order 10 can be realized as a subgroup of $\operatorname{Perm}(X_5)$ in the same way:

$$\mathbf{D}_{10} = \operatorname{Aut}(\Gamma) = \{\sigma \in \operatorname{Perm}(X_5) \mid \sigma \text{ preserves the graph } \Gamma\} .$$

The group $\mathbf{D}_{10}$ has index 12 in $\operatorname{Perm}(X_5)$. The 12 cosets have a geometric meaning: they represent the different possible ways of arranging five elements into an undirected cycle:

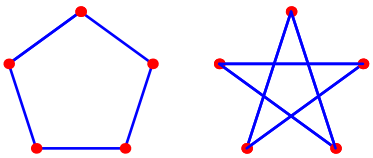

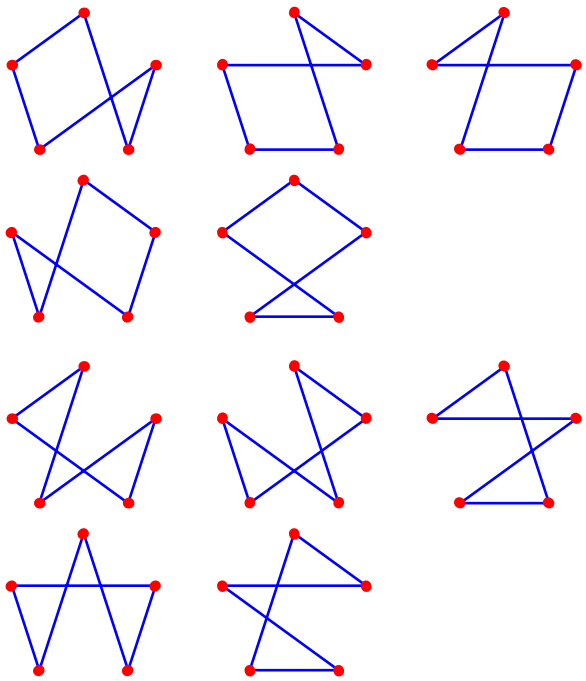

Exercise 1.1.6.16. Explain why and how these pictures represent the cosets of $\mathbf{D}_{10}$ in $\mathrm{Perm}(X_5)$. In particular, decide whether these pictures represent right-cosets or left-cosets.

1.1.7 Equivariant Maps and Invariant Partitions

Definition 1.1.7.1. Let G acting on X by means of an action α and on Y by an action β. A map $f : X \to Y$ is called G-equivariant (should be (α, β)-equivariant, but this is never used since actions are always suppressed anyway) if

$$f(g \ltimes_\alpha x) = g \ltimes_\beta f(x) \text{ for all } g \in G \text{ and all } x \in X.$$

The name G-map is also used because a set for which an action of G is specified is often called a G-set.

Observation 1.1.7.2. *Compositions of G-maps are G-maps, as a straight forward computation shows.*

Remark 1.1.7.3. One might wonder why not just saying "equivariant". If you are determined to suppress the actions involved, why not go

all the way? Well, if U is a subgroup of G it will act also on X and Y. Clearly, every G-map is an U-map. However the converse is not true. Thus you will often want to be specific about the group.

Definition 1.1.7.4. Let α be an action of G on X. An equivalence relation $\sim$ on X is called α-invariant if for all $x, y \in X$ and all $g \in G$, we have

$$x \sim y \iff g \ltimes_\alpha x \sim g \ltimes_\alpha y.$$

A partition of X into equivalence classes of an invariant equivalence relation is called an invariant partition.

Example 1.1.7.5 (graph automorphisms). Let Γ be a graph. An orientation of Γ is an assignment of a direction to each edge in Γ. The distance of two orientations is the number of edges on which they are opposite. Two orientations are called equivalent if they have an even distance.

Let us see that this notion actually defines an equivalence relation. Only transitivity needs proof: Let o_1, o_2 and o_3 be three orientations on Γ. We compute distances:

$$\begin{aligned} \mathrm{d}(o_1, o_3) &= \mathrm{card}(\{e \text{ edge in } \Gamma \mid o_1(e) \neq o_3(e)\}) \\ &= \mathrm{card}(\{e \text{ edge in } \Gamma \mid o_2(e) \neq o_1(e) \text{ or } o_2(e) \neq o_3(e)\}) \\ &\quad - \mathrm{card}(\{e \text{ edge in } \Gamma \mid o_2(e) \neq o_1(e) \text{ and } o_2(e) \neq o_3(e)\}) \end{aligned}$$

and

$$\begin{aligned} \mathrm{d}(o_1, o_2) + \mathrm{d}(o_2, o_3) &= \mathrm{card}(\{e \text{ edge in } \Gamma \mid o_2(e) \neq o_1(e) \text{ or } o_2(e) \neq o_3(e)\}) \\ &\quad + \mathrm{card}(\{e \text{ edge in } \Gamma \mid o_2(e) \neq o_1(e) \text{ and } o_2(e) \neq o_3(e)\}) \end{aligned}$$

Thus,

$$\mathrm{d}(o_1, o_3) \equiv \mathrm{d}(o_1, o_2) + \mathrm{d}(o_2, o_3) \mod 2$$

and transitivity follows.

Equivalence classes of orientations are called orientation types. We denote the set of all orientation types of Γ

by $\mathcal{O}(\Gamma)$. Note that a finite graph has exactly two orientation types.

Note that the action of $\mathrm{Aut}(\Gamma)$ on Γ induces an action of $\mathrm{Aut}(G)$ on the set of all orientations on Γ. For instance, consider the following graph:

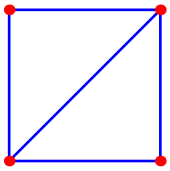

Then a rotation by 180 degrees will map orientations as follows:

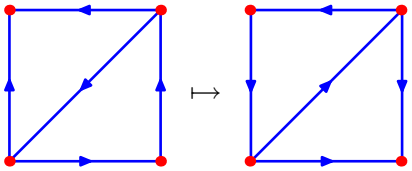

whereas a flip around the diagonal will act like so:

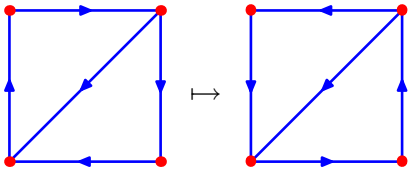

This action preserves distances of orientations: two orientations agree/disagree on an edge e if and only if their translates agree/disagree on the translate of e. Since the group action preserves distances, it preserves parity of distances. Thus equivalence of orientations is a $\mathrm{Aut}(\Gamma)$-invariant equivalence relation.

Example 1.1.7.6 (left-multiplication). The partitions of G invariant with respect to left-multiplication are in 1-1 correspondence with the subgroups of G. The correspondence is given by the following two mutually inverse constructions:

1. Let U be a subgroup of G. We already noted that the decomposition of G into left-cosets of U is an invariant

partition arising from the invariant equivalence relation

$$\begin{aligned} h \equiv_U g \quad &:\Longleftrightarrow \quad h = gu \text{ for some } u \in U \\ &\Longleftrightarrow \quad \bar{g}h \in U. \end{aligned}$$

2. Conversely, suppose $\equiv$ is an equivalence relation on G invariant under left-multiplication. Then

$$U := \{g \in G \mid g \equiv 1\}$$

is a subgroup of G and the $\equiv$-equivalence classes are exactly the left-cosets of U.

To see this, consider $g, h \equiv 1$. Since $\equiv$ is invariant, we find that

$$gh \equiv g1 = g \equiv 1.$$

Also,

$$1 = \bar{g}g \equiv \bar{g}1 = \bar{g}.$$

It follows that U is a subgroup.

Also, note that

$$\begin{aligned} h \equiv g \quad &\Longleftrightarrow \quad \bar{g}h \equiv 1 \\ &\Longleftrightarrow \quad \bar{g}h \in U \\ &\Longleftrightarrow \quad h \equiv_U g \end{aligned}$$

Proposition 1.1.7.7. *Let $\sim$ be an α-invariant equivalence relation on the G-set X. Let ${}^{X}\!/_{\sim}$ be the set of $\sim$-equivalence classes. Then there is a unique G-action ${}^{\alpha}\!/_{\sim}$ on ${}^{X}\!/_{\sim}$ such that the natural projection $X \twoheadrightarrow {}^{X}\!/_{\sim}$ becomes G-equivariant.*

Proof. Let $[x]$ denote the $\sim$-equivalence class of x. If we want the projection $x \mapsto [x]$ to be G-equivariant, we have no choice but to define:

$$g \ltimes [x] := [g \ltimes x]\,.$$

This settles uniqueness. To finish the proof, we need to observe (a) that $g \ltimes [x]$ is well-defined, i.e., independent of the choice of the representative x, and (b) that we have actually defined an action on ${}^{X}\!/_{\sim}$.

Now, (a) is just restating α-invariance of $\sim$, and (b) follows from the fact that α is an action. **q.e.d.**

Example 1.1.7.8 (the sign of a graph automorphism). As a direct consequence of Proposition 1.1.7.7, we see that a group $\mathrm{Aut}(\Gamma)$ of automorphisms of a graph Γ acts on the set of orientation types $\mathcal{O}(\Gamma)$, i.e., we have a canonical homomorphism

$$\mathrm{sign} : \mathrm{Aut}(\Gamma) \longrightarrow \mathrm{Perm}(\mathcal{O}(\Gamma))\,.$$

The kernel of this homomorphism consists of <u>orientation preserving</u> or <u>even</u> automorphisms.

A special case deserves mentioning: The symmetric group $\mathrm{Perm}(X)$ can be regarded as the automorphism group of the complete graph $\mathrm{K}(X)$ over X wherein all pairs of elements are joined by an edge. Thus there is an induced homomorphism

$$\mathrm{sign} : \mathrm{Perm}(X) \longrightarrow \mathrm{Perm}(\mathcal{O}(\mathrm{K}(X)))\,.$$

Since finite graphs have exactly two orientation types, we find a homomorphism

$$\mathrm{sign} : \mathbf{S}_r \longrightarrow \mathbf{S}_2 \cong \{-1, 1\}$$

that sends all transpositions (swaps of exactly two numbers) to -1. [BTW: yes, this is the one used to define determinants.]

Definition 1.1.7.9. Let X be a finite set. The kernel

$$\mathrm{Alt}(X) := \ker(\mathrm{sign} : \mathrm{Perm}(X) \longrightarrow \mathbf{S}_2)$$

is called the <u>alternating group</u> over X. We put

$$\mathbf{A}_r := \mathrm{Alt}(\{1, 2, \ldots, r\})\,.$$

Definition 1.1.7.10. If $f : X \to Y$ is a G-map of G-sets, then

$$x_0 \sim_f x_1 \ :\Longleftrightarrow \ f(x_0) = f(x_1)$$

defines a G-invariant equivalence relation on X, which we call f-equivalence.

Observation 1.1.7.11. *Let $f : X \to Y$ be a G-map, and let $\sim$ denote f-equivalence on X. Then, f factors through the natural projection $X \twoheadrightarrow {}^{X}\!/_{\sim}$ as a product of G-maps*

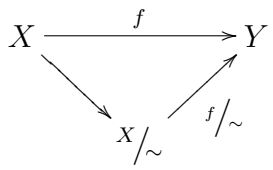

where the induced map ${}^{f}\!/_{\sim}$ is 1-1. In particular, if f is onto, then ${}^{f}\!/_{\sim}$ is a G-bijection.

Orbit-Stabilizer Theorem 1.1.7.12. *Let α be a transitive action of G on X – or, equivalently, let X be an orbit of a G-action. For any $x \in X$, the map*

$$\begin{array}{rcl} {}^{G}\!/_{\mathrm{Stab}(x)} & \longrightarrow & X \\ [g] & \mapsto & g \ltimes_\alpha x \end{array}$$

is a G-equivariant bijection.

In particular, all points in a given orbit have stabilizers of equal index, and that index equals the size of the orbit.

Remark 1.1.7.13. The statement says that we may think of any G-orbit as some quotient ${}^{G}\!/_{U}$.

Proof. We just observe that

$$\begin{array}{rcl} \alpha_x : G & \longrightarrow & X \\ g & \mapsto & g \ltimes_\alpha x \end{array}$$

is a G-map. The stabilizer $\mathrm{Stab}_G(x)$ consists exactly of those elements α_x-equivalent to the identity element. Thus, by (1.1.7.6), the α_x-equivalence classes are the cosets of $\mathrm{Stab}_G(x)$, and the map α_x factors through ${}^{G}\!/_{\mathrm{Stab}_G(x)}$. Clearly the induced map ${}^{G}\!/_{\mathrm{Stab}_G(x)} \to X$ is onto and 1-1. **q.e.d.**

Definition 1.1.7.14. Let $f : X \to Y$ be a map. For any equivalence relation $\approx$ defined on Y, the pull-back is the equivalence relation $\approx^f$ defined on X by

$$x_1 \approx^f x_1 \ :\Longrightarrow\ f(x_0) \approx f(x_1)\,.$$

Remark 1.1.7.15. Note that the pull-back of the identity-relation on Y is just the relation of f-equivalence introduced in (1.1.7.10).

Also note that X and Y are G-sets and f is G-equivariant, then every G-invariant equivalence relation on Y pulls back to a G-invariant equivalence on X.

Definition 1.1.7.16. let $\sim$ and $\approx$ be two equivalence relations on the set X. Then $\approx$ is called coarser than $\sim$ if

$$x_0 \sim x_1 \implies x_0 \approx x_1 \quad \text{for all } x_0, x_1 \in X.$$

This means that $\approx$-equivalence classes are unions of $\sim$-equivalence classes.

Correspondence Theorem for G-sets 1.1.7.17. *Let $f : X \twoheadrightarrow Y$ be surjective. Then pull-back with respect to f induces a 1-1 correspondence*

$$\{\textit{coarsenings of } f\textit{-equivalence}\} \longleftrightarrow \{\textit{equivalences on } Y\}\,.$$

Moreover, if X and Y each carry a G-action and f is G-equivariant, then the above correspondence induces a 1-1 correspondence

$$\{G\textit{-inv. coarsenings of } f\textit{-equivalence}\} \longleftrightarrow \{G\textit{-inv. equivalences on } Y\}\,.$$

Proof. Pull-back preserves the partial order "is coarser than". In particular, since every equivalence relation is coarser than the identity, which pulls back to f-equivalence on X, all pull-backs are coarser than f-equivalence.

Since f is onto, any two difference equivalence relations will pull-back to different relations on X. On the other hand any coarsening $\approx$ of f-equivalence can be "pushed forward" to Y by defining:

$$y_0 \approx_f y_1 \ :\Longleftrightarrow\ x_0 \approx x_1 \text{ for some } x_0 \in f^{-1}(y_0) \text{ and } x_1 \in f^{-1}(y_1)\,.$$

A straight forward computation (using that $\approx$ is coarser than f-equivalence) show that $\approx_f$ is well-defined. Another straight forward computation shows that pushing forward and pulling back are mutually inverses operations.

In the presence of a G-action, all operations are compatible with the action as long as the equivalence relations involved are G-invariant; this is, again, a straight forward computation. [In fact more is true: G acts (from the right!) on the set of all equivalence relations on X and Y, and pull-back and push-forward are G-maps. See Exercise 1.1.7.20.] **q.e.d.**

As a corollary, we obtain the useful:

Correspondence Theorem for Subgroups 1.1.7.18. *Let $U_0 \leq G$ be a subgroup. We have a* 1-1 *correspondence*

$$\{U \mid U_0 \leq U \leq G\} \longleftrightarrow \left\{G\text{-invariant partitions of } {}^{G}\!/_{U_0}\right\}.$$

Proof. Just observe that G-invariant coarsenings of the congruence $\equiv_{U_0}$ are exactly the congruences induced by subgroups between U_0 and G. Now apply the Correspondence Theorem for G-sets. **q.e.d.**

Corollary 1.1.7.19. *Let X be a set and consider the tautological action of* $\mathrm{Perm}(X)$ *on X. For each $x \in X$, the subgroup* $\mathrm{Stab}(x) \leq \mathrm{Perm}(X)$ *is a maximal subgroup.*

The same statement holds for $\mathrm{Alt}(X)$*: one-point stabilizers are maximal subgroups.*

Proof. By the Correspondence Theorem for Subgroups, subgroups containing $\mathrm{Stab}(x)$ are in 1-1 correspondence with $\mathrm{Perm}(X)$-invariant equivalence relations on the set ${}^{\mathrm{Perm}(X)}\!/_{\mathrm{Stab}(x)}$. This set carries a natural $\mathrm{Perm}(X)$-action induced by left-multiplication, and with respect to this action it is $\mathrm{Perm}(X)$-equivariantly bijective to the set X endowed with the tautological $\mathrm{Perm}(X)$-action. Thus, it suffices to show that there are exactly two $\mathrm{Perm}(X)$-invariant equivalence relations on X, namely the identity relation where no two different elements are equivalent (corresponding to the subgroup $\mathrm{Stab}(x)$) and the trivial relation where any two elements are equivalent (this corresponds to the subgroup $\mathrm{Perm}(X)$).

For contradiction, suppose that $\sim$ is an invariant equivalence relation on X such that there are elements $x_0, x_1, x_2 \in X$ with $x_0 \sim x_1$ but $x_0 \not\sim x_2$. Then the transposition interchanging x_1 and x_2 clearly does not stabilize $\sim$.

To deal with $\mathrm{Alt}(X)$, we run the same argument. However, at the very end, we replace the transposition, which is not orientation preserving, by a three-cycle on the elements x_0, x_1, x_2. **q.e.d.**

Exercise 1.1.7.20. Let α be an action of G on X. For any equivalence relation $\sim$ on X (not necessarily invariant), and any g, define the relation $\sim_g$ by

$$x \sim_g y :\Longleftrightarrow g \ltimes x \sim g \ltimes y.$$

1. Let $\mathcal{E}(X)$ be the set of equivalence relations on X. Show that

$$\begin{array}{ccc} \mathcal{E}(X) \times G & \longrightarrow & \mathcal{E}(X) \\ (\sim, g) & \mapsto & \sim_g \end{array}$$

defines a right-action of G on the set of all equivalence relations on X.

2. Show that the fixed points of this right-action are exactly the α-invariant equivalence relations. (If $Y \times G \to Y$ is a right-action, then an element $y \in Y$ is a fixed point of the action, if $yg = y$ for each $g \in G$.)

3. Also consider the natural left-action of G induced on the set $\mathcal{P}(X)$ of partitions of X. Since there is a 1-1 correspondence of equivalence relations on X and partitions of X, one would hope that these two actions are related. Figure out the relationship. Pay particular attention to the fact that one action is from the right and the other is from the left!

Exercise 1.1.7.21. Let G be the automorphism group (group of symmetries) of the following graph Γ:

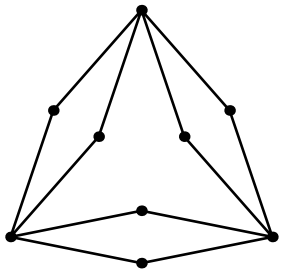

Let H be the group of all symmetries of a solid cube (including orientation reversing symmetries, e.g., reflections).

1. Show that G and H both have 48 elements.

2. Determine whether G and H are isomorphic.

1.1.8 Generating Sets and Cayley Graphs

Definition 1.1.8.1. Let G be a group and let $\mathrm{X} \subseteq G$ be a subset. The Cayley graph of G with respect to X is the graph $\Gamma_{\mathrm{X}}(G)$ whose vertices are the elements of G and any two vertices $g, h \in G$ (group elements!) are connected by an edge if and only $g = h\chi$ for some $\chi \in \mathrm{X}$. (Notice that the element χ goes to the r i g h t!)

Example 1.1.8.2. Here are some Cayley graphs for the infinite cyclic group $\mathbb{Z}$:

- First just the canonical one-element generating set $\{1\}$:

- Now, let us look at the Cayley graph with respect to the set $\{2\}$:

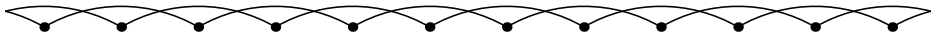

  Note that this graph has two connected components, one for each coset of the subgroup of even integers.

- Low, and behold: this pattern also obtains for the set $\{3\}$:

  We have three connected components, just comprising the cosets of the subgroup of multiples of 3.

- Now, $\mathbb{Z}$ is generated by 2 and 3. The corresponding Cayley graph is connected:

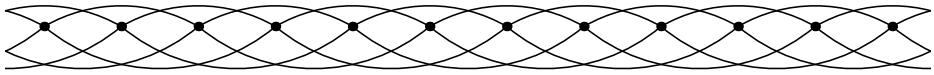

Observation 1.1.8.3. *Left-multiplication induces an action of group G on $\Gamma_{\mathrm{X}}(G)$ by graph automorphisms:*

$$\begin{aligned} & \textit{there is an edge from } h_0 \textit{ to } h_1 \\ \iff\ & h_0 = h_1\chi \textit{ for some } \chi \in \mathrm{X} \\ \iff\ & gh_0 = gh_1\chi \textit{ for some } \chi \in \mathrm{X} \\ \iff\ & \textit{there is an edge from } gh_0 \textit{ to } gh_1 \end{aligned}$$

Definition 1.1.8.4. Define a relation $\sim_{\mathrm{X}}$ on G by

$$g \sim_{\mathrm{X}} h \ :\iff\ g \text{ and } h \text{ are connected by a path in } \Gamma_{\mathrm{X}}(G)\,.$$

Note that this is an equivalence relation and the fact that G acts by automorphisms on the Cayley graph implies that this equivalence relation is invariant. Thus the connected components of the Cayley graph $\Gamma_{\mathrm{X}}(G)$ are exactly the left-cosets of some subgroup $\langle \mathrm{X} \rangle$, called the subgroup generated by X, which of course is the connected component of the identity element. The set X is called a generating set for $\langle \mathrm{X} \rangle$ and its elements are called generators.

Observation 1.1.8.5. *Clearly, $G = \langle \mathrm{X} \rangle$ if and only if $\Gamma_{\mathrm{X}}(G)$ is connected.*

Discussion 1.1.8.6. Let us look at $\langle \mathrm{X} \rangle$ a little closer: it contains exactly those elements $g \in G$ that are connected to the identity element by a path in $\Gamma_{\mathrm{X}}(G)$. That means you can get from 1 to g by repeatedly multiplying (from the right) by an element of X or an inverse thereof. Thus, G can be written as a product of generators and inverses of generators. Such a representation of g is usually called a word.

This describes the subgroup $\langle \mathrm{X} \rangle$ from the inside. Now, we give a description from the outside:

Observation 1.1.8.7. *The subgroup $\langle \mathrm{X} \rangle$ is the smallest subgroup of G containing X, i.e., every subgroup $U \leq G$ that contains X also contains $\langle \mathrm{X} \rangle$.*

Example 1.1.8.8. Let $X = \{1, 2, \ldots, r\}$. A neighbor transposition is a permutation

$$\begin{array}{rccl} \tau_{i,i+1} : X & \longrightarrow & X & \\ & & & \\ & x & \mapsto & \begin{cases} i & \text{if } x = i+1 \\ i+1 & \text{if } x = i \\ x & \text{otherwise.} \end{cases} \end{array}$$

The neighbor transpositions form a generating set for the symmetric group $\mathrm{Perm}(X)$.

Proof. What I will show is that every arrangement of the numbers $1, 2, \ldots, r$ can be sorted into ascending order by a finite sequence of neighbor transformations. That shows that the subgroup generated by neighbor transformations acts transitively on the set of all orders. But then, it must be the whole symmetric group $\mathrm{Perm}(X_r)$ just because of size.

So, suppose you are given an ordering of $\{1, 2, \ldots, r\}$. To sort, look at the position of 1. If it is put at the bottom, fine. If not, use a neighbor transposition to move 1 closer to the bottom. Iterate until 1 is the bottom element. Now use the same trick to put 2 in the second slot. Go on until everything is ordered. **q.e.d.**

Exercise 1.1.8.9. Show that the cyclic rotations of three consecutive elements, i.e., permutations of the type

$$i \ \mapsto \ i+1 \ \mapsto \ i+2 \ \mapsto \ i$$

form a generating set for $\mathbf{A}_r$. Draw the corresponding Cayley graph for $\mathbf{A}_4$. Be sure to make it look cool.

Corollary 1.1.8.10. *The group $\mathbf{A}_r$ is the only subgroup of index 2 in $\mathbf{S}_r$.*

Proof. Observe that any subgroup of index 2 is normal. Thus, we only need to show that the kernel of any homomorphism $\mathbf{S}_r \to \mathbf{C}_2$ contains $\mathbf{A}_r$. Since the cyclic rotations of length 3 generate $\mathbf{A}_r$, we are reduced to proving that any homomorphism $\mathbf{S}_r \to \mathbf{C}_2$ vanishes on three cycles. Note that three cycles are the squares of their inverses. Since homomorphisms take squares to squares and all sqares in $\mathbf{C}_2$ are trivial, the claim follows. **q.e.d.**

Example 1.1.8.11. Table 1.2 shows the Cayley graph of the symmetric group $\mathbf{S}_4 := \mathrm{Perm}(\{1, 2, 3, 4\})$ with respect to the generating set consisting of the three neighbor transpositions. Note:

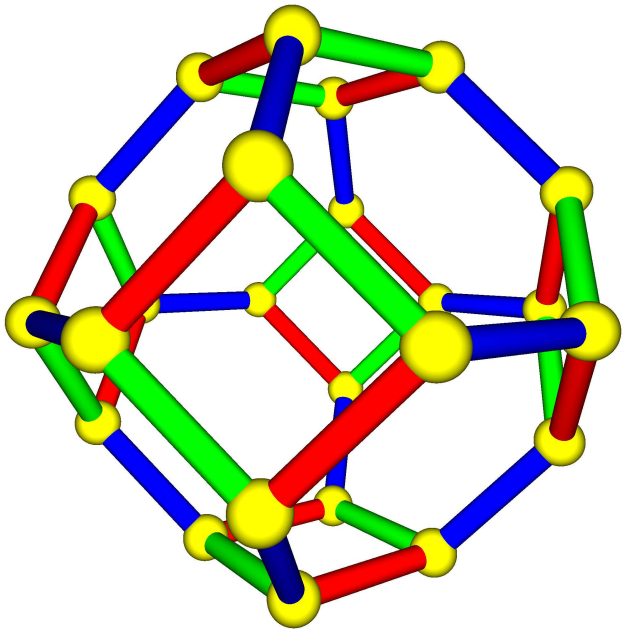

Table 1.2: The Cayley graph for S_4 relative to the generating set of neighbor transpositions

1. The graph has 24 vertices, one for each group element.

2. The edges are left-cosets of the cyclic subgroups $\langle \tau_{1,2} \rangle$, $\langle \tau_{2,3} \rangle$, $\langle \tau_{3,4} \rangle$.

3. The squares are the left-cosets of the subgroup $\langle \tau_{1,2}, \tau_{3,4} \rangle$. Similarly the hexagons correspond to the left-cosets of $\langle \tau_{1,2}, \tau_{2,3} \rangle$ and $\langle \tau_{2,3}, \tau_{3,4} \rangle$.

Exercise 1.1.8.12. Let $\mathrm{GL}_n(\mathbb{Z})$ be the set of all $n \times n$ matrices with integer coefficients and determinant ± 1. Use the identity matrix as identity element and matrix multiplication as multiplication.

1. Show that every matrix in $\mathrm{GL}_n(\mathbb{Z})$ can be transformed into the identity matrix by applying a finite sequence of "very elementary" row operations: (i) swapping two rows, (ii) multiply a row by -1, and (iii) add a row to another row. [Hint: use the ideas from the previous problem.]

2. Show that $\mathrm{GL}_n(\mathbb{Z})$ is a group.

3. Construct a finite generating set for $\mathrm{GL}_n(\mathbb{Z})$.

[Remark: you will receive partial credit for the case $n = 2$.]

Definition 1.1.8.13. A group is cyclic if it is generated by one element.

Exercise 1.1.8.14. Show that every subgroup of a cyclic group is cyclic and that any two cyclic groups are isomorphic if and only if they have the same cardinality.

1.1.9 Fixed Points and Fix Groups

Definition 1.1.9.1. A fixed point of an action α of a group G on a set X is an element $x \in X$ that is not moved by any element in G, i.e., $g \ltimes_\alpha x = x$ for each $g \in G$. The set of all fixed points is denoted by $\mathrm{Fix}(\alpha)$ or, suppressing the action, by X^G.

More generally, for any subgroup $U \leq G$, we put:

$$\mathrm{Fix}_\alpha(U) = X^U := \{x \in X \mid u \ltimes_\alpha x = x \text{ for all } u \in U\}.$$

For any element $g \in G$, we put:

$$\mathrm{Fix}_\alpha(g) := \{x \in X \mid g \ltimes_\alpha x = x\}.$$

Burnside's Lemma 1.1.9.2 (Gaschütz). *For any action of a finite group G on a finite set X,*

$$\mathrm{ord}(G)\,\mathrm{card}\left({}_G\backslash^X\right) = \sum_g \mathrm{card}(\mathrm{Fix}(g)).$$

In other words: the number of orbits equals the average number of fixed points for group elements.

Proof. This is just counting two different ways: Consider the set

$$\mathcal{X} := \{(g, x) \in G \times X \mid g \ltimes x = x\}$$

Grouping by first coordinates yields:

$$\operatorname{card}(\mathcal{X}) = \sum_{g \in G} \operatorname{card}(\operatorname{Fix}(g)) \,.$$

Grouping by second coordinates and then by orbits yields:

$$\begin{aligned} \operatorname{card}(\mathcal{X}) &= \sum_{x \in X} \operatorname{ord}(\operatorname{Stab}(x)) = \sum_{x \in X} \frac{\operatorname{ord}(G)}{\operatorname{card}(\operatorname{Orb}(x))} = \sum_{\mathcal{O} \in G \backslash X} \sum_{x \in \mathcal{O}} \frac{\operatorname{ord}(G)}{\operatorname{card}(\mathcal{O})} \\ &= \sum_{\mathcal{O} \in G \backslash X} \operatorname{ord}(G) = \operatorname{card}\left(G \backslash X\right) \operatorname{ord}(G) \end{aligned}$$

Duh! **q.e.d.**

Corollary 1.1.9.3. *Consider the tautological action of $\mathbf{S}_r$ on a set of r elements. Since the action is transitive, we have*

$$\operatorname{ord}(\mathbf{S}_r) = \sum_{\sigma \in \mathbf{S}_r} \operatorname{card}(\operatorname{Fix}(\sigma))$$

which implies that permutations in $\mathbf{S}_r$ have on average exactly one fixed point.

Exercise 1.1.9.4. For $r > 1$, compute the standard deviation for the number of fixed points of permutations in $\mathbf{S}_r$.

Definition 1.1.9.5. Let G act on X. Note that

$$\begin{aligned} G \times \operatorname{Pow}(X) &\longrightarrow \operatorname{Pow}(X) \\ (g, A) &\mapsto g \ltimes A := \{g \ltimes x \mid x \in A\} \end{aligned}$$

defines an action of G on $\operatorname{Pow}(X)$. The stabilizer of a subset with respect to this action is denoted (unsurprisingly) as follows:

$$\operatorname{Stab}_G(A) := \{g \in G \mid g \ltimes A = A\} \,.$$

Note, that $\operatorname{Stab}(A)$ acts on A. The kernel of this action is the fix group of A:

$$\operatorname{Fix}_G(A) := \{g \in G \mid g \ltimes x = x \text{ for each } x \in A\} = \bigcap_{x \in A} \operatorname{Stab}_G(x)$$

1.1.10 Example: Conjugation

Lemma and Definition 1.1.10.1. *Let G be a group. Then*

$$\begin{aligned} \mathrm{ad} : G \times G &\longrightarrow G \\ (g,h) &\mapsto \mathrm{ad}_g(h) := ghg^{-1} \end{aligned}$$

defines an action of G on itself.

This action is an action by automorphisms, i.e., for each $g \in G$, the map $\mathrm{ad}_g : G \to G$ is an invertible homomorphism called the inner automorphism associated to g.

This action is called conjugation. Its orbits are called conjugacy classes.

Proof. Straight forward computations. **q.e.d.**

Example 1.1.10.2. The problem of deciding whether two given elements are conjugate is called the conjugacy problem. Of course, for finite groups, this problem can be solves by exhaustive search for a conjugating element. However, sometimes more elegant solutions can be obtained. Here, we shall solve the conjugacy problem in $\mathrm{Perm}(X)$ for finite X.

Let $\sigma : X \to X$ be a permutation of X. A σ-cycle (of length u) is a tuple

$$(x_1, x_2, \ldots, x_u)$$

such that

$$\sigma(x_i) = \begin{cases} x_{i+1} & \text{for } i < u \\ x_1 & \text{for } i = u \end{cases}$$

The following are easily checked:

1. If $(x_1, x_2, \ldots, x_u)$ is a σ-cycle, then so is $(x_2, x_3, \ldots, x_u, x_1)$.
2. Two σ-cycles either share all elements or no elements. In the former case, we consider the cycles equivalent (or actually equal); and in the later case, we call them disjoint.

3. If $(x_1, x_2, \ldots, x_u)$ is a σ-cycle, and τ is another permutation of X, then $(\tau(x_1), \tau(x_2), \ldots, \tau(x_u))$ is a $\tau\sigma\bar{\tau}$-cycle.

The cycle-type of σ is the sorted list of all lengths of all (equivalence classes of) σ-cycles. The last observation implies that conjugate permutations have the same cycle-type.

On the other hand, suppose two permutations σ_0 and σ_1 have the same cycle-type. Then, we can "align their cycles". For instance

$$\begin{aligned} \sigma_0 &: (125)(86)(43)(7) \\ \sigma_1 &: (238)(71)(56)(4) \end{aligned}$$

Now, the permutation τ defined by

$$\begin{aligned} \tau : 1 &\mapsto 2 \\ \tau : 2 &\mapsto 3 \\ \tau : 5 &\mapsto 8 \\ \tau : 8 &\mapsto 7 \\ \tau : 6 &\mapsto 1 \\ \tau : 4 &\mapsto 5 \\ \tau : 3 &\mapsto 6 \\ \tau : 7 &\mapsto 4 \end{aligned}$$

conjugates σ_0 to σ_1. This also follows from the last observation above.

Thus, we conclude:

> Two elements of $\mathrm{Perm}(X)$ are conjugate if and only if they have the same cycle-type.

Definition 1.1.10.3. The conjugation action of G on itself extends naturally to an action of G on the set of all subsets of G. Since the conjugacy action is an action by automorphisms of G, the induced

action takes subgroups of G to subgroups. Two subgroups in the same orbit are called conjugate subgroups.

A subgroup $U \leq G$ is normal if it is stabilized by the conjugation action. It is central if conjugation fixes the subgroup pointwise. The stabilizer of U with respect to the conjugation action is called the normalizer $\mathrm{N}_G(U)$. The fix-group of U with respect to conjugacy is called the centralizer $\mathrm{C}_G(U)$.

Exercise 1.1.10.4. Show that two subgroups U_0 and U_1 in G are conjugate if and only if there is a G-equivariant bijection between the coset-sets ${}^{G}\!/_{U_0}$ and ${}^{G}\!/_{U_1}$.

Exercise 1.1.10.5. Let U be a subgroup of G. Show that the set

$$\{H \leq G \,|\, U \text{ is a normal subgroup in } H\}$$

has a unique maximal element, namely the normalizer $\mathrm{N}_G(U)$.

Characterization of Normal Subgroups 1.1.10.6. *Let $N \leq G$ be a subgroup. Then the following are equivalent:*

1. *N is normal in G.*
2. *Every left-coset of N is a right-coset of N (and vice versa).*
3. *The left-multiplication action of G on the coset set ${}^{G}\!/_{N}$ restricts to a trivial action of N on ${}^{G}\!/_{N}$.*
4. *The coset set ${}^{G}\!/_{N}$ carries a (unique) group structure such that the canonical projection $G \to {}^{G}\!/_{N}$ is a homomorphism. (Note that N then is the kernel of this homomorphism!)*
5. *N is the kernel of some homomorphism defined on G.*

Proof.

(1) $\Longrightarrow$ **(2)** Exercise. (You know this already!)

(2) $\Longrightarrow$ **(3)** Pick $n \in N$. Then $ng \in Ng = gN$ whence $ngN = gN$.

(3) $\Longrightarrow$ **(4)** The requirement that $G \twoheadrightarrow {}^{G}\!/_{N}$ be a homomorphism leaves no choice but to define the multiplication on ${}^{G}\!/_{N}$ via $(gN)(hN) := (gh)N$. This is well defined since $(gn_0hn_1)N = g(n_0hN) = g(hN) = (gh)N$. The group axioms are inherited from G.

(4) $\Longrightarrow$ **(5)** Obvious.

(5) $\Longrightarrow$ **(1)** Straightforward computation. **q.e.d.**

Definition 1.1.10.7. A subgroup of G is called characteristic if it is stabilized by all automorphisms of G.

Observation 1.1.10.8. *Every characteristic subgroup is normal.*

Example and Definition 1.1.10.9. The set of all fixed points of the conjugacy action of G is called the center of G, its elements are called central. It is denoted by $\mathrm{Z}(G)$ and is a characteristic subgroup.

Definition 1.1.10.10. Let $g, h \in G$. We say that g and h commute if $gh = hg$. If all elements of G commute pairwise, we call G commutative or Abelian.

Failure to commute is measured by the commutator, defined as:

$$[g,h] := gh\bar{g}\bar{h} = \mathrm{ad}_g(h)\,\bar{h} = g\,\mathrm{ad}_h(g)\,.$$

Definition 1.1.10.11. Let $A, B \subseteq G$ be subsets. We put:

$$[A,B] := \langle\{[a,b] \mid a \in A,\ \ b \in B\}\rangle$$

The subgroup $[G,G]$ is called the commutator subgroup.

Remark 1.1.10.12. Note that AB does not need to be a subgroup even if A and B happen to be subgroups.

Exercise 1.1.10.13. Show that a subgroup $N \leq G$ is normal if and only if $[N,G] \leq N$.

Exercise 1.1.10.14. The commutator subgroup is characteristic (hence normal). The quotient $^{G}/_{[G,G]}$ is Abelian. Every normal subgroup $N \leq G$ with Abelian quotient $^{G}/_{N}$ contains the commutator subgroup.

Exercise 1.1.10.15. Show that for any element $g \in G$, the following are equivalent:

1. The element g is central, i.e., it is fixed by the conjugacy action.
2. The element g commutes with all elements in G, i.e., $gh = hg$ for any $h \in G$.
3. The automorphism $\mathrm{ad}_g : G \to G$ is the identity.

Infer that $\mathrm{Z}(G) = \mathrm{C}_G(G)$.

Observation 1.1.10.16. *Let N_0 and N_1 be two subgroups in G with trivial intersection, i.e., $N_0 \cap N_1 = \{1\}$. If N_0 and N_1 mutually normalize one another (e.g., if both are normal), then they already centralize one another. In other words: any two elements $g \in N_0$ and $h \in N_1$ commute: $gh = hg$.*

Proof. The commutator $[g, h]$ lies in $N_0 \cap N_1$. **q.e.d.**

Exercise 1.1.10.17. Let G act on X. Prove that for every subset $A \subseteq X$, the fix group $\mathrm{Fix}_G(A)$ is normal in the stabilizer $\mathrm{Stab}_G(A)$.

Exercise 1.1.10.18. Let G be a group and put

$$\mathrm{Aut}(G) := \{\varphi : G \to G \,|\, \varphi \text{ is an automorphism}\}.$$

1. Show that $\mathrm{Aut}(G)$ is a subgroup of $\mathrm{Perm}(G)$.
2. Show that the homomorphism

$$\begin{array}{rcl} \mathrm{ad} : G & \longrightarrow & \mathrm{Perm}(G) \\ g & \mapsto & \mathrm{ad}_g \end{array}$$

   takes values in $\mathrm{Aut}(G)$. Let $\mathrm{Inn}(G)$ denote the image of ad.
3. Show that $\mathrm{Inn}(G)$ is a normal subgroup of $\mathrm{Aut}(G)$.
4. Show that G is Abelian if $\mathrm{Inn}(G)$ is cyclic.

1.1.11 Extensions and (Short) Exact Sequences

Definition 1.1.11.1. A sequence

$$\cdots \xrightarrow{\varphi_{i-3}} G_{i-2} \xrightarrow{\varphi_{i-2}} G_{i-1} \xrightarrow{\varphi_{i-1}} G_i \xrightarrow{\varphi_i} G_{i+1} \xrightarrow{\varphi_{i+1}} G_{i+2} \xrightarrow{\varphi_{i+2}} \cdots$$

is called exact at G_i if $\ker(\varphi_i) = \mathrm{im}(\varphi_{i-1})$. A short exact sequence is a sequence

$$1 \longrightarrow N \longrightarrow G \longrightarrow Q \longrightarrow 1$$

that is exact at N, G, and Q. Note that exactness at N just requires $N \to G$ to be 1-1 and exactness at Q just means that $G \to Q$ is onto. In this case, the group G is called an extension of N by Q.

Isomorphism Theorem 1.1.11.2. *Suppose we have a commutative diagram with short exact rows:*

$$\begin{array}{ccccc} N & \rightarrowtail & G & \xrightarrow{\pi_1} & Q_1 \\ \| & & \| & & \\ N & \rightarrowtail & G & \xrightarrow[\pi_2]{} & Q_2 \end{array}$$

Then there is a unique homomorphism $\varphi : Q_1 \to Q_2$ such that

$$\begin{array}{ccccc} N & \rightarrowtail & G & \xrightarrow{\pi_1} & Q_1 \\ \| & & \| & & \downarrow \varphi \\ N & \rightarrowtail & G & \xrightarrow[\pi_2]{} & Q_2 \end{array}$$

commutes. Moreover, φ is an isomorphism.

Proof. Since π_1 and π_2 are both onto, φ is uniquely determined already on the level of sets: we have to satisfy

$$\varphi(\pi_1(g)) = \pi_2(g)\,,$$

which determines φ on all of Q_1. For the above equation to yield a well-defined map φ, we need independence of representatives, i.e:

$$\pi_1(g) = \pi_1(h) \implies \pi_2(g) = \pi_2(h) \qquad \text{for all } g, h \in G.$$

Now, this condition follows from exactness: the epimorphisms π_1 and π_2 have the same kernel. In fact, at this point, we only need that $\ker(\pi_1) \subseteq \ker(\pi_2)$. The map φ is a homomorphism since

$$\varphi(\pi_1(g)\,\pi_1(h)) = \varphi(\pi_1(gh)) = \pi_2(gh) = \pi_2(g)\,\pi_2(h) = \varphi(\pi_1(g))\,\varphi(\pi_1(h))$$

That φ is an isomorphism follows actually from uniqueness by means of a neat trick worth knowing: Consider the diagram

$$\begin{array}{ccccc}
N & \rightarrowtail & G & \xrightarrow{\pi_1} & Q_1 \\
\| & & \| & & \downarrow{\scriptstyle \varphi} \\
N & \rightarrowtail & G & \xrightarrow[\pi_2]{} & Q_2 \\
\| & & \| & & \downarrow{\scriptstyle \psi} \\
N & \rightarrowtail & G & \xrightarrow{\pi_1} & Q_1
\end{array}$$

Then the top and bottom row are a problem of the same kind, clearly solved by $\psi \circ \varphi$. Since the identity map on Q_1 also solves the problem, uniqueness implies $\psi \circ \varphi = \mathrm{id}_{Q_1}$. By the same trick: $\varphi \circ \psi = \mathrm{id}_{Q_2}$. Thus φ and ψ are inverse isomorphisms. **q.e.d.**

Exercise 1.1.11.3. Suppose the following commutative diagram has exact rows:

$$\begin{array}{ccccccccc}
G_{-2} & \longrightarrow & G_{-1} & \longrightarrow & G_0 & \longrightarrow & G_1 & \longrightarrow & G_2 \\
\downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\
H_{-2} & \longrightarrow & H_{-1} & \longrightarrow & H_0 & \longrightarrow & H_1 & \longrightarrow & H_2
\end{array}$$

Assume that the vertical arrows at G_{-2}, G_{-1}, G_1, and G_2 are isomorphisms. Show that the arrow $G_0 \to H_0$ is an isomorphism, too. What hypotheses would you need to prove just injectivity? – just surjectivity?

Exercise 1.1.11.4. Suppose that the following commutative diagram of groups has short exact columns.

$$\begin{array}{ccccc}
N_1 & \longrightarrow & G_1 & \longrightarrow & Q_1 \\
\downarrow & & \downarrow & & \downarrow \\
N_2 & \longrightarrow & G_2 & \longrightarrow & Q_2 \\
\downarrow & & \downarrow & & \downarrow \\
N_3 & \longrightarrow & G_3 & \longrightarrow & Q_3
\end{array}$$

1. Show: if the two top rows are short exact sequences, then so is the bottom row.

2. Disprove: if the top and the bottom row are short exact sequences, then so is the middle row.

3. Show: if the two bottom rows are short exact sequences, then so is the top row.

The technique of proving these claims is called diagram chase. As a model argument, we show that if the two bottom rows are exact, then the right arrow in the top row is onto. When reading the argument, follow the path of the elements as they are running through the diagram:

Let q_1 be in Q_1.

Let q_2 be the image of q_1 in Q_2.

By exactness of the middle row, q_2 is the image of some $g_2 \in G_2$.

Let g_3 be the image of g_2 in G_3.

By exactness, the image of q_2 in Q_3 is 1.

By commutativity, the image of g_3 in Q_3 is 1.

By exactness of the bottom row, q_3 is the image of some $n_3 \in N_3$.

By exactness of the left column, n_3 is the image of some $n_2 \in N_2$.

Let h_2 be the image of n_2 in G_2.

By commutativity, g_2 and h_2 both have image g_3 in G_3.

Put $f_2 := h_2^{-1} g_2 \in G_2$. Note that f_2 goes to 1 in G_3.

By exactness of the middle column, f_2 has a preimage f_1 in G_1.

Claim: f_1 goes to q_1 in Q_1.

By exactness of the middle row, the image of h_2 in Q_2 is 1.

Thus, f_2 and g_2 both go to q_2 in Q_2.

By commutativity, f_1 goes to a preimage of q_2 in Q_1.

By exactness of the right column, q_2 has only one preimage in Q_1.

That preimage is q_1.

Example 1.1.11.5 (yet another isomorphism theorem). Assume N_1 is normal in G and N_0 is normal in both, N_1 and G. Then, we have the diagram:

$$\begin{array}{ccccc} N_0 & \longrightarrow & N_1 & \longrightarrow & {}^{N_1}\!/_{N_0} \\ \downarrow & & \downarrow & & \downarrow \\ N_0 & \longrightarrow & G & \longrightarrow & {}^{G}\!/_{N_0} \\ \downarrow & & \downarrow & & \downarrow \\ 1 & \longrightarrow & {}^{G}\!/_{N_1} & \longrightarrow & \left({}^{G}\!/_{N_0}\right)\Big/\left({}^{N_1}\!/_{N_0}\right) \end{array}$$

which yields an isomorphism that you know already from your undergraduate algebra.

Correspondence Theorem 1.1.11.6. *Let*

$$1 \longrightarrow N \longrightarrow G \xrightarrow{\pi} Q \longrightarrow 1$$

be a short exact sequence of groups. Then π induces 1-1 *correspondences:*

$$\{U \mid N \leq U \leq G\} \longleftrightarrow \{V \mid V \leq Q\}$$

and:

$$\{U \mid N \leq U \leq G,\ U \text{ is normal in } G\} \longleftrightarrow \{V \mid V \leq Q\ V \text{ is normal in } Q\}$$

Proof. To see the first correspondence, we first apply the Correspondence Theorem for subgroups to see that the subgroups of G containing N actually correspond to G-invariant partitions of the quotient ${}^{G}\!/_{N}$. But this quotient is canonically isomorphic to Q. Thus, the G-invariant partitions of Q represent the subgroups in G

above N. Finally, note that the G-action on Q factors through the left-multiplication action of Q on itself. Thus G-invariant partitions are exactly the Q-invariant partitions. These, however, correspond to the subgroups of Q.

Chasing through and unravelling the identifications, you will find that the correspondence is actually realized by sending a subgroup $V \leq Q$ to its preimage $\pi^{-1}(V)$.

This correspondence restricts to a bijection on the level of normal subgroups by the Characterization of Normal Subgroups, e.g., that a subgroup is normal if and only if the set of cosets carries a group structure so that the natural projection is a homomorphism. (elaborate, or do it differently.) **q.e.d.**

A special case of extensions are direct products

Definition 1.1.11.7. Let G and H be groups. The direct product $G \times H$ is the group with set $G \times H = \{(g,h) \mid g \in G \text{ and } h \in H\}$ and multiplication

$$(g_0, h_0)\,(g_1, h_1) := (g_0 g_1, h_0 h_1)\,.$$

Exercise 1.1.11.8. Show that the above actually defines a group. Also verify that the identity element is $(1_G, 1_H)$.

Exercise 1.1.11.9. Verify the following claims by straight forward calculations:

1. The projection onto the first coordinate

$$\begin{aligned} \pi_1 : G \times H &\longrightarrow G \\ (g,h) &\mapsto g \end{aligned}$$

   is a homomorphism.

2. Similarly, the projection onto the second coordinate

$$\begin{aligned} \pi_2 : G \times H &\longrightarrow G \\ (g,h) &\mapsto h \end{aligned}$$

   is a homomorphism.

3. The inclusion of the first direct factor

$$\begin{aligned} \iota_1 : G &\longrightarrow G \times H \\ g &\mapsto (g, 1_H) \end{aligned}$$

is a homomorphism.

4. The inclusion of the second direct factor

$$\begin{aligned} \iota_2 : H &\longrightarrow G \times H \\ g &\mapsto (1_G, h) \end{aligned}$$

is a homomorphism.

5. The sequences

$$G \xrightarrow{\iota_1} G \times H \xrightarrow{\pi_2} H$$

and

$$H \xrightarrow{\iota_2} G \times H \xrightarrow{\pi_1} G$$

are both short exact sequences.

Exercise 1.1.11.10. Let U_0 and U_1 both be subgroups of G. Suppose:

1. U_0 and U_1 are both normal in G.
2. $U_0 \cap U_1 = \{1\}$.
3. The map

$$\begin{aligned} U_0 \times U_1 &\longrightarrow G \\ (h_0, h_1) &\mapsto h_0 h_1 \end{aligned}$$

is onto.

Show that

$$\begin{aligned} U_0 \times U_1 &\longrightarrow G \\ (h_0, h_1) &\mapsto h_0 h_1 \end{aligned}$$

is an isomomorphism of groups.

1.1.12 Solvable and Nilpotent Groups

Definition 1.1.12.1. Let foo, bar, and blah be properties. A group is called foo-by-bar if it is an extension of a foo-group by a bar-group. The class of poly-blah groups is the smallest class that contains all blah-groups and is closed with respect to extensions.

Lemma 1.1.12.2. *A group G is poly-blah if and only if it contains a chain of subgroups*

$$1 = G_0 \leq G_1 \leq G_2 \leq \cdots \leq G_{u-1} \leq G_u = G$$

such that each G_i is normal in G_{i+1} and the quotient ${}^{G_{i+1}}\!/_{G_i}$ is a blah-group.

Proof. Let $\mathcal{C}_\alpha$ be the class of groups that have chains of the given form with length $u \leq \alpha$. We have to show that the class $\bigcup_\alpha \mathcal{C}_\alpha$ is the class of poly-blah groups.

So, let $N \in \mathcal{C}_\alpha$ with subgroup chain

$$1 = N_0 \leq N_1 \leq N_2 \leq \cdots \leq N_{u-1} \leq N_u = N$$

and let $Q \in \mathcal{C}_\beta$ with subgroup chain

$$1 = Q_0 \leq Q_1 \leq Q_2 \leq \cdots \leq Q_{v-1} \leq Q_v = Q$$

Now, we consider a short exact sequence

$$N \hookrightarrow G \twoheadrightarrow Q.$$

Our first claim is that $G \in \mathcal{C}_{\alpha+\beta}$ with subgroup chain

$$G_i := \begin{cases} N_i & \text{for } i \leq u \\ \text{preimage of } Q_{i-u} & \text{for } i \geq u \end{cases}$$

The claim follows immediately from the correspondence theorem and the isomorphism theorem: the quotients ${}^{Q_{i+1}}\!/_{Q_i}$ and ${}^{G_{i+1}}\!/_{G_i}$ are isomorphic. This proves that all poly-blah groups are in $\bigcup_\alpha \mathcal{C}_\alpha$.

Conversely, we show by induction that $\bigcup_\alpha \mathcal{C}_\alpha$ consists only of poly-blah groups. To see that, we note that groups in $\mathcal{C}_1$ are exactly the blah groups. Now, groups in $\mathcal{C}_{\alpha+1}$ are extensions of groups in $\mathcal{C}_\alpha$, which are poly-blah by induction hypothesis, by blah groups. The claim follows. **q.e.d.**

An important special case is:

Definition 1.1.12.3. A group is solvable if it is poly-Abelian.

A stronger condition is also quite frequently used:

Definition 1.1.12.4. A group G, not necessarily finite, is nilpotent if it contains a chain of subgroups

$$1 = G_0 \leq G_1 \leq G_2 \leq \cdots \leq G_u = G$$

satisfying $[P, P_{i+1}] \leq P_i$ for all $i < u$.

Observation 1.1.12.5. *Note that in such a chain all G_i are normal in G.* **q.e.d.**

Proposition 1.1.12.6. *The direct product of two nilpotent groups G and H is nilpotent.*

Proof. Let

$$1 = G_0 \leq G_1 \leq G_2 \leq \cdots \leq G_u = G$$

be a subgroup chain for G and let

$$1 = H_0 \leq H_1 \leq H_2 \leq \cdots \leq H_v = H$$

be a subgroup chain for H. Note that we can assume without loss of generality assume that $u = v$ since we can always extend subgroup chains by repeating terms.

Then

$$1 = G_0 \times H_0 \leq G_1 \times H_1 \leq G_2 \times H_2 \leq \cdots \leq G_u \times H_u = G \times H$$

is a subgroup chain proving $G \times H$ nilpotent. **q.e.d.**

Exercise 1.1.12.7. Show that the class of nilpotent groups is the smallest class containing all Abelian groups that is closed with respect to taking central extensions. [Hint: this is to say that (a) central extensions by nilpotent groups are nilpotent and (b) every nilpotent group can be obtained starting with an Abelian group by passing to a central extension finitely many times. Thus you might want to introduce the notion of being "n step nilpotent" and use induction.]

Corollary 1.1.12.8. *Nilpotent groups are solvable.*

Exercise 1.1.12.9. Give a direct proof of the corollary.

1.1.13 Simple Groups

Definition 1.1.13.1. A group is simple if it does not contain any normal strict non-trivial subgroups. I.e., a group is simple if is cannot be written as an extension in a non-trivial way.

Exercise 1.1.13.2. Show that all finite groups are poly-simple.

Exercise 1.1.13.3. Show that non-Abelian finite simple groups have even order.

Exercise 1.1.13.4. Show that Abelian finite simple groups are cyclic of prime order.

The Hölder Program 1.1.13.5. Since all finite groups can be obtained from finite simple groups by forming a finite number of extensions, ???. Hölder formulated the following research project:

1. Classify all finite simple groups.
2. Classify all extensions that can be obtained from two finite groups.

A solution to the classification of finite simple groups has been announced in 1980. People are still busy writing up a proof (the first "proof" had a gap). Researchers in finite groups can be a little touchy about this. So let us pretend that the classification is correct.

The second part of the Hölder program inspired a lot of research, too. Despite considerable progress, it can safely be considered hopeless.

1.2 Finite Groups

Let G be a finite group. We already observed that all subgroups of G have order and index dividing the order of G. Thus, the length of any G-orbit in an action is also a divisor of $\mathrm{ord}(G)$. It is this small set of observations that fuel the elementary theory of finite groups: we will just apply these insights to various actions. (Uhm, well, ok: since finite groups get smaller when you pass to quotients and subgroups, we might run into an induction occasionally.)

1.2.1 Warm-up Example

Example 1.2.1.1. Let p be the smallest prime divisor of $\mathrm{ord}(G)$. Then any action of G on a set of size p is trivial if it has a fixed point because all other orbits must have length 1 or length at least p.

In particular, any subgroup $U \leq G$ of index p is normal: The quotient ${}^{G}\!/_{U}$ has size p. Restrict the left-multiplication action of G on ${}^{G}\!/_{U}$ to U. Note that U fixes the left-coset $1U$. Since non-trivial U-orbits also have length at least p, there cannot be any. Thus, U acts trivially on ${}^{G}\!/_{U}$, i.e., U is normal in G.

1.2.2 Detecting Subgroups

Definition 1.2.2.1. Let p be a prime number. A finite p-group is a group that has p^m elements for some integer m.

Observation 1.2.2.2. *Every quotient and every subgroup of a finite p-group is a finite p-group.* **q.e.d.**

The importance of p-groups stems from the following:

Observation 1.2.2.3. *Let G act on the finite set X such that all non-trivial orbits have length divisible by n, then*

$$\mathrm{card}(X) \equiv \mathrm{card}(\mathrm{Fix}_G(X)) \mod n.$$

In particular, if G is a finite p-group, all orbits have p-power length. Thus:

$$\operatorname{card}(X) \equiv \operatorname{card}(\operatorname{Fix}_G(X)) \mod p.$$

Example 1.2.2.4 (counting block partitions). Let X be a finite set whose size is a multiple of k. Let $\mathcal{P}_k(X)$ be the set of all decompositions of X into pairwise disjoint subsets of size k. We call such subsets the blocks of a partition and k is the block size.

Let us count the block partitions of block size k in a set of size kd. We obtain:

$$\operatorname{card}(\mathcal{P}_k(X)) = \frac{(kd)!}{(d!)(k!)^d}.$$

For instance, for $k = 4$ and a set X_{24} of 24 elements, we have

$$\operatorname{card}(\mathcal{P}_4(X_{24})) = \frac{24\cdot 23\cdot 22\cdot 21}{4\cdot 3\cdot 2\cdot 1\cdot 6}\frac{20\cdot 19\cdot 18\cdot 17}{4\cdot 3\cdot 2\cdot 1\cdot 5}\frac{16\cdot 15\cdot 14\cdot 13}{4\cdot 3\cdot 2\cdot 1\cdot 4}\cdots\frac{4\cdot 3\cdot 2\cdot 1}{4\cdot 3\cdot 2\cdot 1\cdot 1}$$

Miraculously, the result ends up being odd since all 2-prime factors cancel. This works for other prime numbers, as well: For any finite set X and any prime power p^m, dividing the size of X, we have the congruence:

$$\operatorname{card}(\mathcal{P}_{p^m}(X)) \equiv 1 \mod p.$$

Proof. Let $A \subset X$ be a subset of X of exactly p^m elements. It clearly suffices to show:

$$\operatorname{card}(\mathcal{P}_{p^m}(X)) \equiv \operatorname{card}(\mathcal{P}_{p^m}(X - A)) \mod p.$$

To argue this congruence, let the cyclic group of order p^m act as a cyclic permutation on the elements in A and fixes all elements of $X - A$. This action induces an action on $\mathcal{P}_{p^m}(X)$ whose fixed points are exactly those partitions that feature A as a block. Thus the number of fixed points of this action is given by $\operatorname{card}(\mathcal{P}_{p^m}(X - A))$. Now the congruence follows from (1.2.2.3). **q.e.d.**

Now recall from (1.1.7.6) that there is a 1-1 correspondence of subgroups in G and left-invariant partitions of G. This restricts to a 1-1-correspondence

$$\{\text{subgroups in } G \text{ of size } p^m\} \longleftrightarrow \{\text{invariant partitions in } \mathcal{P}_{p^m}(G)\}$$

We shall use this correspondence and Observation 1.2.2.3 to prove the first substantial result. It says, the world is teeming with p-groups:

Strong Cauchy Theorem 1.2.2.5. *Let G be a finite group and let p^m be a prime power dividing the order of G. Then G has a subgroup of order p^m.*

Proof. We consider the action of G on $\mathcal{P}_{p^m}(G)$ induced by left-multiplication. Since $\operatorname{card}(\mathcal{P}_{p^m}(G)) \equiv 1 \mod p$, there is an orbits whose length is not divisible by p. If this orbit is a fixed point, we are done: we found a subgroup of size p^m.

So assume that the orbits orbit is non-trivial Then the stabilizer of any point in such an orbit is a strict subgroup of G whose order is a multiple of p^m. Now we only have to find a subgroup in there. But this is a problem of smaller size whence the result follows by induction. **q.e.d.**

Definition 1.2.2.6. A p-group is a group all elements of which have p-power order.

Remark 1.2.2.7. Above we have defined the term "finite p-group". Now, this phrase can be parsed two ways. This, however, is harmless: finite p-groups are finite p-groups, i.e., a p-group that happens to be finite has p-power order.

Proof. Let P be a p-group of finite order. If the order had a prime factor other than p, the group P would contain an element of that order. **q.e.d.**

1.2.3 The Structure of Finite p-Groups

For p-groups, we can vastly improve upon the Strong Cauchy Theorem:

Proposition 1.2.3.1. *Let P be a finite p-group. Fix a subgroup $U_0 \leq P$ and let p^m be a p-power dividing the order of P but divisible by the order of U_0. Then there exists a subgroup U of order p^m satisfying $U_0 \leq U \leq P$. In fact, the number of such subgroups is $\equiv 1 \mod p$.*

Moreover, if U_0 is normal in P, then U can be chosen to be normal, as well.

Proof. The Correspondence Theorem for Subgroups (1.1.7.18) implies that subgroups in P containing U_0 correspond exactly to the P-invariant partitions of ${}^{P}\!/_{U_0}$. If we require the order of U to be p^m, then the corresponding partition of ${}^{P}\!/_{U_0}$ has block size $k := \frac{p^m}{\text{ord}(U_0)}$. Thus, we study the action of P on $\mathcal{P}_k\left({}^{P}\!/_{U_0}\right)$ induced by left-multiplication. By (1.2.2.4), the set $\mathcal{P}_k\left({}^{P}\!/_{U_0}\right)$ has size $\equiv 1 \mod p$. Thus, (1.2.2.3) implies that the number of fixed points also is $\equiv 1 \mod p$. This proves the first claim.

Now assume U_0 is normal. Normal subgroups are exactly those stable under the conjugacy action. Thus, we consider the action of P

$$X := \{U \mid U_0 \leq U \leq P,\ \text{ord}(U) = p^m\}$$

induced by conjugation (conjugation takes subgroups to subgroups of the same order!). We just proved that $\text{card}(X) \equiv 1 \mod p$. By (1.2.2.3), any action of a p-group on a set of such size has a fixed point. This is the normal subgroup U for which we were searching. **q.e.d.**

Subgroup Structure of Finite p-Groups 1.2.3.2. *Let P be a group of order p^m.*

1. *Every subgroup U occurs in a subgroup chain*

$$1 = U_0 \leq U_1 \leq U_2 \leq \cdots \leq U_{m-1} \leq U_m = P$$

wherein $\mathrm{ord}(U_i) = p^i$. *In particular, all maximal subgroups of* P *have order* p^{m-1}.

2. *Every normal subgroup* N *occurs in a subgroup chain*

$$1 = N_0 \leq N_1 \leq N_2 \leq \cdots \leq N_{m-1} \leq N_m = P$$

wherein $\mathrm{ord}(N_i) = p^i$ *and all* N_i *are normal in* P. *Any such chain satisfies* $[P, N_{i+1}] \leq N_i$, *for all* $i < m$.

Moreover all maximal subgroups of P *are normal, and all normal subgroups of order* p *are central.*

In particular, P *is nilpotent and has non-trivial center.*

Proof. To argue the first claim, we begin by constructing the upper part of the subgroup chain: By Proposition 1.2.3.1, every subgroup U has an index p-supergroup in P. Iterating this argument yields the desired sequence. For the part of the subgroup chain below U, note that the Strong Cauchy Theorem (Proposition 1.2.3.1) implies that every non-trivial p-group has an index p subgroup. Now iterate again.

The construction of a normal subgroup chain surrounding N is similar: for the upper part apply again Proposition 1.2.3.1 to embed N as an index p subgroup into some normal subgroup in P; then iterate.

The lower part of the chain is a little mor subtle. Consider the set of maximal subgroups in N:

$$X := \{U' \leq N \,|\, U' \text{ has index } p \text{ in } N\}.$$

By Proposition 1.2.3.1, $\mathrm{card}(X) \equiv 1 \mod p$. Since N is normal, P acts on X by conjugation. Now (1.2.2.3) implies that this action has a fixed point, i.e., a subgroup of index p in N that is normal in P.

Maximal subgroups of P are normal by (1.2.1.1). Even more was shown there: any action of a p-group on a set of size p is

trivial as soon as it has at least one fixed point. This also proves that normal subgroups of order p are central. Just look at the conjugation action of P. It fixes the identity element and therefore the other $p-1$ elements are also fixed. But a group that is fixed elementwise by under conjugation is central.

Finally, we argue the inclusion $[P, N_{i+1}] \leq N_i$. Consider the projection $P \twoheadrightarrow P/N_i$. The image of N_{i+1} in the quotient is a normal subgroup of order p. Such a subgroup is central. The inclusion $[P, N_{i+1}] \leq N_i$ is merely restating this fact in terms of coset-representatives upstairs in P. **q.e.d.**

Proposition 1.2.3.3. *Let P be a finite, Abelian p-group. Let x be an element of maximum order in P. Then the cyclic subgroup C generated by x is a direct factor of P.*

Proof. Let V be a subgroup of P maximal among those intersecting C trivially. Let U be the subgroup generated by C and V. Then $U = C \times V$ since everything takes place in an Abelian group. We have to show that $U = P$.

We argue by contradiction and assume that the quotient P/U is non-trivial. Since it is a p-group it contains a element of order p by Cauchys theorem. Let g be a preimage in P of such an element. It follows that $g^p = x^k v \in C \times V$. Let p_{i+1} be the order of g, and let p_j be the order of x. Then $1 = g^{p_{i+1}} = x^{p_i} V^{p_{i+1}} \in C \times V$ whence $x^{kp_i} = 1$. We infer that p_j divides kp_i And since the order of x is maximal, we have $j > i$. Thus, p divides into k.

Put: $h := g\bar{x}^{k/p}$, and note that $gU = hU \in {}^{P}\!/_{U}$. Also, $h^p = g^p \bar{x}_k = v \in V$ and it follows that C intersects $\langle h, V \rangle$ trivially: Any word in h and elements from V can be sorted, moving all powers of h to the left. Hence:

$$\langle h, V \rangle = \left\{ h^k v \,\middle|\, k \in \mathbb{Z} v \in V \right\}.$$

On the other hand, any element $h^k v$ in C represents $1U$ in ${}^{P}\!/_{U}$. It follows that k is a multiple of p, in which case $h^k v \in V \cap C = 1$.

Since h, like g, is non-trivial in P/U, we have a contradiction to maximality of V. **q.e.d.**

Induction yields immediately the following:

Corollary 1.2.3.4. *Every finite Abelian p-group is a direct product of cyclic subgroups.*

1.2.4 Sylow Subgroups

Definition 1.2.4.1. Let p be a prime dividing the order of the finite group G. A p-Sylow subgroup of G is a maximal p-subgroup of G.

Sylow Theorem 1.2.4.2. *Let G be a finite group of order $p^m d$ where p does not divide d. Then G contains a p-Sylow subgroup S of order p^m. Moreover:*

1. *Every p-subgroup of G is contained in some conjugate of S. In particular all p-Sylow subgroups are conjugate and of size p^m.*
2. *Any p-subgroup P normalizes S if and only if $P \leq S$. In particular, a p-Sylow subgroup normalizes itself but no other p-Sylow subgroup.*
3. *The number of p-Sylow subgroups of G divides d and is $\equiv 1 \mod p$.*

Proof. Since p^m divides the order of G, the existence of a subgroup S of that size follows from Theorem 1.2.2.5. This subgroup is clearly a maximal p-subgroup since higher powers of p do not divide evenly into $\mathrm{ord}(G)$. As for the moreover parts, we argue as follows:

1. Let P be any p-subgroup of G. Consider the left-multiplication action of P on the left-cosets of S. Since p does not divide the number of left-cosets, (1.2.2.3) implies that the action must have a fixed point. Thus:

$$PgS = gS$$

for some $g \in G$. It follows that $P \leq gS\bar{g}$.

2. The quotient $\mathrm{N}_G(S)/S$ has order not divisible by p. Thus, P has trivial image therein.

3. Let X be the set of all p-Sylow subgroups in G. First we consider the action of G on X induced by conjugation. As argued above, this action is transitive. Then $\mathrm{Stab}_G(S)$ has index $\mathrm{card}(X)$ in G. However, $S \leq \mathrm{Stab}_G(S)$. It follows that $\mathrm{card}(X)$ divides the complementary factor d.

   Now restrict the conjugation action of G on X to S. Every S-orbit of the conjugation action has length dividing $\mathrm{ord}(S) = p^m$. Since, in the preceding item, we argued that S normalizes exactly one p-Sylow subgroup, namely itself, it follows that the conjugation action of S on X has a unique fixed point. All other orbits have length divisible by p. The congruence follows. **q.e.d.**

Observation 1.2.4.3. *All p-Sylow subgroups are conjugate and every conjugate of a p-Sylow subgroup is a p-Sylow subgroup. Hence a p-Sylow subgroup is normal if and only if it is the only p-Sylow subgroup.* **q.e.d.**

Proposition 1.2.4.4. *Let G be a finite group, let S be a p-Sylow subgroup of G and suppose a subgroup U satisfies $\mathrm{N}_G(S) \leq U \leq G$. Then $\mathrm{N}_G(U) = U$.*

Proof. Let $g \in G$ normalize U. We have to show $g \in U$.

Note that S is a p-Sylow subgroup of U Consider the conjugate $gS\bar{g} \leq U$, which clearly is also a p-Sylow subgroup in U. Therefore, these two groups are already conjugate in U, i.e., there is an element $u \in U$ such that

$$gS\bar{g} = uS\bar{u}.$$

Hence $\bar{u}g$ normalizes S. Thus $\bar{u}g \in \mathrm{N}_G(S) \leq U$ whence $g \in U$. **q.e.d.**

Theorem 1.2.4.5. *Let G be a finite group. Then, the following are equivalent:*

1. *G is nilpotent.*
2. *$U < \mathrm{N}_G(U)$ for each strict subgroup $U < G$.*
3. *Every maximal subgroup in G is normal.*
4. *Every Sylow subgroup in G is normal.*
5. *G is a the direct product of its Sylow subgroups for the various primes.*

Proof.

(1) $\Longrightarrow$ **(2)** Let G be nilpotent with a subgroup chain

$$1 = G_0 \leq G_1 \leq G_2 \leq \cdots \leq G_r = G$$

satisfying $[G, G_{i+1}] \leq G_i$ for each i. Let j be maximal with $G_j \leq U$.

Note

$$[U, G_{j+1}] \leq [G, G_{j+1}] \leq G_j \leq U$$

It follows that G_{j+1} normalizes U. Thus, $G_{j+1} \leq \mathrm{N}_G(U)$ and therefore (by maximality of j), we find $U \neq \mathrm{N}_G(U)$.

(2) $\Longrightarrow$ **(3)** Let $M \leq G$ be maximal. Since $M < \mathrm{N}_G(M)$ we have $\mathrm{N}_G(M) = G$.

(3) $\Longrightarrow$ **(4)** Suppose S was a Sylow subgroup not normal in G, i.e., $\mathrm{N}_G(S) \neq G$. Then there is a maximal subgroup M containing $\mathrm{N}_G(S)$. By Proposition 1.2.4.4, we find $\mathrm{N}_G(M) = M$ contradicting normality of maximal subgroups.

(4) $\Longrightarrow$ **(5)** We have to show three items:

1. For two different primes p and q, any p-Sylow subgroup S_p intersects trivially with any q-Sylow subgroup: Note that the intersection $S_p \cap S_q'$ is a p-group as well as a q-group. Thus, it must be trivial (not much choice for a possible order).

2. For two different primes p and q, any p-Sylow subgroup S_p commutes with any q-Sylow subgroup: By hypothesis, all Sylow subgroups are normal. From S_p being normal, we infer $[S_p, S_q] \leq S_p$. Similarly, $[S_p, S_q] \leq S_q$ whence, by the previous item: $[S_p, S_q] \leq S_p \cap S_q = 1$.

3. G is generated by its Sylow subgroups for the various primes: Let U be the subgroup generated by all Sylow subgroups of G. Note that the order of every Sylow subgroup divides into the order $\mathrm{ord}(U)$. Since the orders of the Sylow subgroups in G form the prime factor decomposition of $\mathrm{ord}(G)$ it follows that the order of G divides the order of the subgroup U. Thus $G = U$.

(5) $\Longrightarrow$ **(1)** Note that the Sylow subgroups of G are finite p-groups. Thus, they are all nilpotent. Direct products of nilpotent groups are nilpotent by Proposition 1.1.12.6. **q.e.d.**

Structure Theorem for Finite Abelian Groups 1.2.4.6. *Every finite Abelian group is the direct product of cyclic subgroups of prime power order.*

Proof. Follows from Theorem 1.2.4.5 and Corollary 1.2.3.4. **q.e.d.**

Exercise 1.2.4.7. Let S be a Sylow subgroup of G. Show that $\mathrm{N}_G(S) = \mathrm{N}_G(\mathrm{N}_G(S))$.

1.2.5 Applications: Groups of "Small" Orders

Of course, "small" orders really are orders with simple prime factor decompositions.

Observation 1.2.5.1. *Any group of prime order is cyclic.*

Exercise 1.2.5.2. Show that a group of order p^2 is Abelian (p is, of course, prime); and infer that it is either $\mathbf{C}_p \times \mathbf{C}_p$ or $\mathbf{C}_{p^2}$. [Hint: show that the group of inner automorphism is cyclic.]

Observation 1.2.5.3. *Let G be a group of order pq^m wherein $p < q$ are prime numbers. Then G is solvable: The number of q-Sylow subgroups in G divides p and is $\equiv 1 \mod q$. Thus, G has a unique, hence normal, q-Sylow subgroup S_q and there is a short exact sequence*

$$S_q \hookrightarrow G \twoheadrightarrow \mathbf{C}_p.$$

Since S_q is nilpotent (hence solvable) and cyclic groups are Abelian (hence solvable), G is solvable-by-solvable and hence solvable. **q.e.d.**

This is a little more tricky, but in the same spirit:

Exercise 1.2.5.4. Let G be a group of order $p^m q$ wherein $p < q$ are prime numbers. Show that G is solvable.

It is worth noting that the methods of elementary finite group theory are not sufficient to prove even the most naive generalization of this type of results:

Fact 1.2.5.5 (Burnside). *Show a group of order $p^m q^n$ is not simple.*

So don't be fooled: advanced finite group theory looks much different from what is presented here.

Exercise 1.2.5.6. Determine the subgroup lattices of $\mathbf{S}_4$ and $\mathbf{D}_8$.

Proposition 1.2.5.7. *Let G be a simple group of order 60. Then $G \cong \mathbf{A}_5$.*

Proof. First, we construct a suitable action.

Claim A. There is a transitive action of G on a set of size 5.

Proof. Let $\mathcal{S}_2$ be the set of all 2-Sylow subgroups in G. Since G is simple, it cannot have a unique 2-Sylow subgroup as such a group would be normal. Also, $\mathcal{S}_2$ cannot consist of three elements: Conjugacy induces a transitive action of G on $\mathcal{S}_2$ which gives rise to a non-trivial homomorphism

$$G \to \mathrm{Perm}(\mathcal{S}_2)\,.$$

Since the target group has order six and the homomorphism is non-trivial, we would find a proper non-trivial kernel. That does not happen since G is simple.

That leaves two possibilities: G could have five or fifteen 2-Sylow subgroups. If the size of $\mathcal{S}_2$ is five, we are done.

So, assume that G has fifteen 2-Sylow subgroups.

Any two distinct 2-Sylow subgroups intersect trivially: In this case, we find a total of 45 non-identity elements in all fifteen 2-Sylow subgroups. All these elements have order 2 or 4. Note that G cannot contain a unique 5-Sylow subgroup as that had to be normal. Thus, we have six 5-Sylow subgroups. These are cyclic of prime order whence they pairwise intersect trivially. Thus, we find another batch of elements, namely 24 non-identity elements of order 5. This exceeds the quota of 59 non-identity elements in G – and we did not even consider elements of order 3.

There is a non-trivial g common in two distinct 2-Sylow subgroups: Let S_1 and S_2 be two distinct 2-Sylow subgroups whose intersection contains a non-trivial element g. Then $\mathrm{C}_G(g)$ is a subgroup of G. Since groups of order 4 are Abelian, we find $S_1 \leq \mathrm{C}_G(g)$. Since $\mathrm{C}_G(g)$ also contains S_2, these inclusion is strict and we infer that the order of $\mathrm{C}_G(g)$ is a strict multiple of 4. That leaves three possibilities: 12, 20, and 60.

Note that $\mathrm{C}_G(g)$ cannot be all of G since in that case, g had to be central. Then G has a non-trivial center. This cannot

happen in simple groups, unless they are cyclic.

Also note that G cannot contain a subgroup U of index 3: The left-multiplication would induce a transitive action on the three cosets in ${}^{G}\!/_{U}$. As above, we obtain a non-trivial homomorphism to a symmetric group of order 6. This rules out order 20.

Thus, $\mathrm{C}_G(g)$ has order 12, i.e., index 5. Then, the left-multiplication action on the five cosets is the transitive action we have been looking for.

We have finished the proof of the main analysis and can now finish the proof. From a transitive action of G on a set of five elements, we obtain a non-trivial homomorphism

$$G \hookrightarrow \mathbf{S}_5$$

which is injective since G cannot host a non-trivial kernel. Since G has order 60, we can regard G as a subgroup of index 2 in $\mathbf{S}_5$. Now the claim follows from Corollary 1.1.8.10. **q.e.d.**

Proposition 1.2.5.8. $\mathbf{A}_5$ *is simple.*

Proof (Zassenhaus). We know the conjugacy classes in $\mathbf{S}_5$. They are given by cycle types. Since $\mathbf{A}_5$ is normal in $\mathbf{S}_5$, a conjugacy class of $\mathbf{S}_5$ is either contained in or disjoint from $\mathbf{A}_5$. The $\mathbf{S}_5$-conjugacy classes inside of $\mathbf{A}_5$ are:

$$\begin{aligned} \{1\} &\quad \text{size: } 1 \\ \{(\cdot\cdot)(\cdot\cdot)(\cdot)\} &\quad \text{size: } 15 \\ \{(\cdot\cdot\cdot)(\cdot)(\cdot)\} &\quad \text{size: } 20 \\ \{(\cdot\cdot\cdot\cdot\cdot)\} &\quad \text{size: } 24 \end{aligned}$$

The first goal of our analysis is to determine the conjugacy classes in $\mathbf{A}_5$. So let $C \subseteq \mathbf{A}_5$ be an $\mathbf{A}_5$-conjugacy class. Fix an odd permutation $\sigma \in \mathbf{S}_5$. Since $\mathbf{S}_5 = \mathbf{A}_5 \cup \sigma\mathbf{A}_5$, the set $C \cup \sigma C \bar{\sigma}$ is an

$\mathbf{S}_5$-conjugacy class. Also, C is either equal to or disjoint from $\sigma C \bar{\sigma}$. It follows that conjugacy classes in $\mathbf{A}_5$ either are $\mathbf{S}_5$-conjugacy classes or have a disjoint parter with whom they form an $\mathbf{S}_5$-conjugacy class. It follows that the conjugacy classes of $\mathbf{A}_5$ have sizes

$$1, 15, \text{ twice } 10 \text{ or once } 20, \text{ twice } 12 \text{ or once } 24.$$

A subgroup in $\mathbf{A}_5$ has order dividing 60. A normal subgroup is a disjoint union of conjugacy classes. Thus, we can rule out the existence of normal subgroups by looking at which divisors of 60 can be written as a sum of terms involving only the numbers listed above. As you can check. there simply is no way to do this (keep in mind that you have to include 1as every subgroup contains the identity.) **q.e.d.**

Theorem 1.2.5.9. *For $r \neq 4$, the alternating group $\mathbf{A}_r$ is simple.*

Proof. We use induction on r. The case $r = 3$ is easy, and the case $r = 5$ was dealt with above. So assume $r \geq 6$ and that $\mathbf{A}_{r-1}$ is simple.

We consider the tautological action of $\mathbf{A}_r$ on $\{1, 2, \ldots, r\}$. Each stabilizer $\mathrm{Stab}(i)$ is isomorphic to $\mathbf{A}_{r-1}$ and hence simple by induction. Let $N \leq \mathbf{A}_r$ be a normal subgroup.

$\underline{N \cap \mathrm{Stab}(1) \neq \{1\}}$: Note that $N \cap \mathrm{Stab}(1)$ is normal in $\mathrm{Stab}(1)$. Since $\mathrm{Stab}(1)$ is simple, we infer $\mathrm{Stab}(1) \leq N$.

Note that the action of $\mathbf{A}_r$ on $\{1, 2, \ldots, r\}$ is transitive. Hence all stabilizers $\mathrm{Stab}(i)$ are conjugate. Since N is normal, we deduce for each i that $\mathrm{Stab}(i) \leq N$.

Since $\mathbf{A}_r$ is generated by the union of the $\mathrm{Stab}(i)$, it follows that $N = \mathbf{A}_r$.

$\underline{N \cap \mathrm{Stab}(1) = \{1\}}$: As above, we observe that all $\mathrm{Stab}(i)$ are conjugate to $\mathrm{Stab}(1)$. This time we infer that $N \cap \mathrm{Stab}(i) = \{1\}$ for all i. Thus, any element of N that fixes but one number is already trivial.

Suppose $n \in N$ is non-trivial.

n is a product of disjoint transpositions: W.l.o.g., we assume that $n = (12)(34)\cdots$ and we put $n' = (356)n(365) = (12)(54)\cdots \in N$. Then the element $\bar{n'}n = (1)(2)\cdots \in N$ fixes a point but is non-trivial. This is a contradiction.

n has a cycle decomposition with at least one cycle of length ≥ 3: W.l.o.g., we assume that $n = (123\ldots)\cdots$ and we put $n' = (356)n(365) = (125\ldots)\cdots \in N$. Then the element $\bar{n'}n = (1)\cdots \in N$ fixes a point but is non-trivial. This is a contradiction. **q.e.d.**

1.3 Infinite Groups

1.3.1 Free Groups

Definition 1.3.1.1. Let X be a set. A
free group with free generating set X is a group F together with a map $\iota : X \to \mathrm{F}$ satisfying the following universal property:

> For every group G and every map $f : X \to G$ there exists a unique group homomorphis $\varphi_f : \mathrm{F} \to G$ that makes the following diagram commute:
>
> 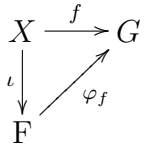
>

In this section, we shall prove the following:

Theorem 1.3.1.2. *For every set X there is free group F_X with free generating set X and this free group is unique up to unique isomorphism.*

As usual, we start with uniqueness:

Lemma 1.3.1.3. *Suppose $\iota_0 : X \to \mathrm{F}_0$ and $\iota_1 : X \to \mathrm{F}_1$ are two free groups with free generating set X. Then there is a unique isomorphism*

$$\varphi : \mathrm{F}_0 \to \mathrm{F}_1$$

that makes the following diagram commute:

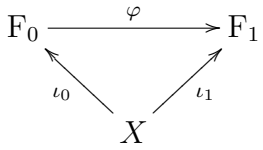

Proof. Note that the uniquenes part of the universal property implies that any homomorphism $\mathrm{F}_i \to \mathrm{F}_i$ that makes

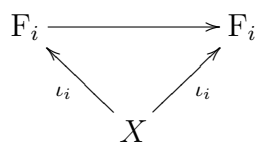

commutative has to be the identity on F_i.

The universal property implies that there is a unique homomorphism

$$\varphi : \mathrm{F}_0 \to \mathrm{F}_1$$

with the desired property. On the other hand, we also have that there is a unique homomorphism in the other direction:

$$\psi : \mathrm{F}_1 \to \mathrm{F}_0$$

Finally, we observe that $\varphi \circ \psi$ is the identity on F_1 and $\psi \circ \varphi$ is the identity on F_0 by our introductory observation. **q.e.d.**

As usual, the main problem is existence of the free object. There are several ways to construct free groups, and I shall present two: a one based upon pictures and one based upon words.

Definition 1.3.1.4. Let X^{-1} be a set of formal inverses of the elements of X disjoint from X. A word over X is an element of the free monoid over $X \cup X^{-1}$. A word is called reduced if it does not contain a subword of the form xx^{-1} or $x^{-1}x$. Two words are called neighbors if you can obtain one from the other by inserting or deleting a subword of either of those two forms. Equivalence of words is defined as the transitive closure of neighborhood.

Observation 1.3.1.5. *Equivalence of words is compatible with concatenation, i.e., if the words w_0 and w_1 are equivalent and if v_0 and v_1 are equivalent, too; then w_0v_1 is equivalent to w_1v_1. Consequently, concatenation of equivalence classes is well-defined.*

Proposition 1.3.1.6. *The set F of equivalence classes of words with the binary operation induced by concatenation is a free group over X where the map $\iota : X \to \mathrm{F}$ sends every generator to the parallelism class of its corresponding one-letter word.*

Proof. We clearly have a monoid before us. Inverses are obtained by formally inverting a word: reverse the order and flip the exponents. Any product ww^{-1} visibly reduces to the empty word.

It remains to argue the universal property. Let G be a group and suppose we have a map $f : X \to G$. Uniqueness of the extending homomorphism is clear – we have to proceed as follows: First, extrend this map to $X \cup X^{-1}$ by sending x^{-1} to $f(x)^{-1}$. Thus, we obtain a monoid homomorphism

$$\left(X \cup X^{-1}\right)^* \longrightarrow G$$

and we want to see that this induces a (group) homomorphism from F to G. Note that F is an epimorphic image of $(X \cup X^{-1})^*$ and that all we need to argue is that different lifts of elements in F to the free monoid have equal images in G. Thus, all we need to observe is that equivalent words map to identical group elements in G. This is obvious since it visibly holds for neighboring words. **q.e.d.**

Definition 1.3.1.7. A diagram over the set X is a finite directed tree D (i.e., a finite graph without cycles each of whose edges is given orientation) drawn in the plane whose edges are labeled by elements from the set X.

Let p be a path in D. Note that the given orientations and edge labels allow us to read a word as we travel along p as follows: we construct the word from left to right; whenever we travel along an edge with label $x \in X$, we write x if we travel along the orientation of the edge, and when we travel against the edge, we write x^{-1}.

A path in a diagram is called a left-boundard path if at any given vertex you take the left-most possible turn. It is a right-boundary path if you always take the right-most possible turn.

For any two distinct vertices in D there is a unique left-boundary and a unique right-boundary path between them. We call the two corresponding boundary words complementary. In general, we call two words parallel if they can be realized as complementary boundary words in some diagram.

Proposition 1.3.1.8. *Parallelsm is an equivalence relation on the set of words that is compatible with concatenation, i.e., if the*

words w_0 and w_1 are parallel and if v_0 and v_1 are parallel, too; then w_0v_1 is parallel to w_1v_1. Consequently, concatenation of equivalence classes is well-defined.

Proof. Clearly parallelism is symmetric. It is reflexive as shown by a straight line path reading a given word. Given two diagrams D_0 and D_1 with a left-boundary path p_0 in D_0 and a right-boundary path p_1 in D_1 that read the same boundary word, we can glue the diagrams together by identifying p_0 with p_1. The result is a diagram that we denote by: $D_0 \cup_{p_0=p_1} D_1$ This construction proves transitivity.

Compatibility with concatenation is obvious from glueing diagrams along terminal vertices. **q.e.d.**

PICTURES

Proposition 1.3.1.9. *The set F of parallelism classes of words with the binary operation induced by concatenation is a free group over X where the map $\iota : X \to \mathrm{F}$ sends every generator to the parallelism class of its corresponding one-letter word.*

Proof. Clearly, the set F is a monoid.

Note that a straight line reading a word w has a boundary path reading ww^{-1} that goes all the way around. Thus, ww^{-1} is parallel to the empty word. Therefore, F has inverses and is a group.

We need to verify the universal property. So let $f : X \to G$ be a map. AS in the previous case (using equivalence of words instead of parallelism), the residual problem that remains to be solved is to show that parallel words evaluate to identical elements in G. This amounts to say that the complete boundary word around a tree evaluates trivially. This can be seen easily using induction on the edges of the tree: adding an edge just extends the boundary word by a cancelling pair. **q.e.d.**

Exercise 1.3.1.10. Show that parellism of words and equivalence of words are the same relations.

Exercise 1.3.1.11. Show that every equivalence class of words contains a unique reduced word.

Remark 1.3.1.12. It follows that one may think of elements in F_X as reduced words over $X \cup X^{-1}$ where multiplication is defined as concatenation followed up by reduction.

Observation 1.3.1.13. *One letter words with different letters are not equivalent. Consequently, the canonical inclusion $\iota : X \to \mathrm{F}_X$ is injective. Thus, we identify X with its image in F_X. Note that F_X is generated by X.*

Exercise 1.3.1.14. Show that the Cayley graph of F_X with respect to the generating set X is a regular tree of degree $2\,\mathrm{card}(X)$, i.e., each vertex has degree $2\,\mathrm{card}(X)$.

1.3.2 Presentations of Groups

Definition 1.3.2.1 (Presentation). Let $\mathcal{X}$ be a set and let $\mathcal{R}$ be a set of elements in $\mathrm{F} := \mathrm{F}_{\mathcal{X}}$. Let $N \trianglelefteq \mathrm{F}$ be the intersection of all normal subgroups in F that contain $\mathcal{R}$, Observe that N is normal in F, i.e., N is the smallest normal subgroup of F containing $\mathcal{R}$. The pair $\langle \mathcal{X} \,|\, \mathcal{R} \rangle$ is called a presentation for the quotient ${}^{\mathrm{F}}\!/_{N}$; and the group ${}^{\mathrm{F}}\!/_{N}$ is said to be defined by the presentation $\langle \mathcal{X} \,|\, \mathcal{R} \rangle$.

Chapter 2

Rings

2.1 Basic Notions

2.1.1 Rings and Modules

Definition 2.1.1.1. A ring is a set R together with two distinct distinguished elements $0 \in R$ (called zero) and $1 \in R - \{0\}$ (called one or identity element) and with two binary operations

$$\begin{aligned} \cdot : R \times R &\longrightarrow R \\ (a, b) &\mapsto ab \end{aligned}$$

(called multiplication) and

$$\begin{aligned} + : R \times R &\longrightarrow R \\ (a, b) &\mapsto a + b \end{aligned}$$

(called addition) such that $(R, +, 0)$ is an Abelian group and such that the following additional axioms hold:

1. $1a = a1 = a$ for each $a \in R$.
2. $(ab)c = a(bc)$ for all $a, b, c \in R$.
3. $(a + b)c = ac + bc$ for all $a, b, c \in R$.
4. $a(b + c) = ab + ac$ for all $a, b, c \in R$.

(The first two axioms say that $(R, \cdot, 1)$ is a monoid.)

R is commutative if, in addition, $ab = ba$ for all $a, b \in R$.

A ring is said to have no zero divisors if $R - \{0\}$ is multiplicatively closed. If, in addition, $(R - \{0\}, \cdot, 1)$ is a group, R is called a division ring or skew field. A commutative ring without zero divisors is called a domain and a commutative division ring is called a field.

Remark 2.1.1.2. A ring without multiplicative identity (i.e., an algebraic structure satisfying the ring axioms except possibly those that require the existence of 1) is a rng. We will not have much use for rngs. Examples include real-valued functions with compact support on $\mathbb{R}$.

Remark 2.1.1.3. In a ring (as opposed to a rng), commutativity of addition is forced by distributivity and the existence of a multiplicative identity:

$$a+b+a+b = 1(a+b)+1(a+b) = (1+1)(a+b) = (1+1)a+(1+1)b = a+a+b+b$$

whence

$$a + b = b + a$$

Remark 2.1.1.4. The additive inverses of a in $(R, +)$ is denoted by $-a$. Its multiplicative inverse, should it have one, is denoted by a^{-1}.

Observation 2.1.1.5.

1. *We have* $ab = a(b+0) = ab + a0$ *whence* $a0 = 0$ *for any* $a \in R$. *Similarly,* $0a = 0$ *for each* $a \in R$.

2. *Also,* $ab + ((-a)b) = (a + (-a))b = 0b = 0$ *and consequently:*

$$-(ab) = (-a)b = a(-b).$$

3. *Above, we defined a ring as a set with two distinguished elements. We already know from group theory, that the additive*

structure uniquely determines the additive identity element 0. *Also, the multiplicative structure determines the multiplicative identity element since*

$$1 = 11' = 1'$$

if 1 *and* $1'$ *are two multiplicative identities.*

4. $R^* := \{a \in R \mid ab = 1 = ba$ *for some* $b \in R\}$ *is a group with respect to multiplication and called the* group of units *of* R. *This also settles the question as to what the elements of* R^* *are called.*

Definition 2.1.1.6. An element $a \in R$ is a left-divisor of $b \in R$ if there is a non-zero element $c \in R - \{0\}$ with $b = ac$. The notion of a right-divisor is defined symmetrically.

A non-zero element $a \in R - \{0\}$ is called a left zero divisor if there is an element $c \in R - \{0\}$ with $ac = 0$. Note that a ring has left zero divisors if and only if it has right zero divisors.

As a rule of thumb, units are good and zero divisors are bad.

Example 2.1.1.7. $\mathbb{Z}$ is a domain. Its group of units is $\{1, -1\}$.

Example 2.1.1.8. For any field K, the polynomial ring $K[x]$ is a domain. Its group of units consits precisely of the non-zero constant polynomials.

Example 2.1.1.9. Let K be a field and let $n \geq 2$. Then

$$\mathbb{M}_n(K) := \{n \times n\text{-matrices over } K\}$$

is a non-commutative ring with zero divisors:

$$\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}.$$

Its group of units is $\mathrm{GL}_n(K)$.

Example 2.1.1.10. $\mathbb{Q}$, $\mathbb{R}$, $\mathbb{C}$ are fields. So is the field of rational functions $K(x)$ over a field K.

Example 2.1.1.11. $\mathbb{Z}\left[\frac{1}{m}\right] := \left\{\frac{n}{m^k} \,\middle|\, n,k \in \mathbb{Z}\right\}$ is the smallest subring of $\mathbb{Q}$ that contains $\frac{1}{m}$. It is a domain.

Remark 2.1.1.12. Any subring of a field is clearly a domain. The converse is true, as well. We shall prove that in (**??**).

Example 2.1.1.13. $\mathbb{Q}[i] := \{a+ib \,|\, a,b \in \mathbb{Q}\}$ is a subfield of $\mathbb{C}$.

Example 2.1.1.14. $\mathbb{Q}\left[\sqrt{2}\right] := \left\{a+\sqrt{2}b \,\middle|\, a,b \in \mathbb{Q}\right\}$ is a subfield of $\mathbb{R}$.

Example 2.1.1.15. Let R be a ring and let $\mathcal{M}$ be a monoid. The set

$$R[\mathcal{M}] := \{f : \mathcal{M} \to R \,|\, f(\mu) = 0 \text{ for all but finitely many } \mu \in \mathcal{M}\}$$

is a ring with respect to pointwise addition

$$(f+g)(\mu) := f(\mu) + g(\mu)$$

and the convolution product

$$(f \cdot g)(\mu) := \sum_{\mu=\nu_1\nu_2} f(\nu_1)\, g(\nu_2)$$

where the sum is really finite since all but finitely many terms vanish. Verification of the ring axioms is a straight forward check.

Note that the monoid $\mathcal{M}$ embeds into the multiplicative monoid of $R[\mathcal{M}]$ via sending every element to its characteristic function:

$$\begin{array}{ccc} \mathcal{M} & \longrightarrow & R[\mathcal{M}] \\ \mu & \mapsto & \hat{\mu} := \chi_\mu : \nu \mapsto \begin{cases} 1 & \text{if } \mu = \nu \\ 0 & \text{if } \mu \neq \nu \end{cases} \end{array}$$

If G is a group, the ring $R[G]$ is usually called the group ring of G.

Example 2.1.1.16. Let $\mathcal{M}$ be a monoid and suppose that the following condition holds:

($\star$) Every $\mu \in \mathcal{M}$ can be written as a product $\mu = \nu_1\nu_2$ in at most finitely many ways.

Let R be a ring. Then, the set

$$R[[\mathcal{M}]] := \{f : \mathcal{M} \to R\}$$

is a ring with respect to pointwise addition and the convolution product as in example 2.1.1.15. In this example, condition $(\star)$ ensures that the sum is finite. Again, verification of the ring axioms is a straight forward check.

Note that $R[\mathcal{M}]$ is a subring of $R[[\mathcal{M}]]$.

Also note that both $(\mathbb{N}_0, +)$ and $(\mathbb{N}_1, \cdot)$ both satisfy $(\star)$.

Remark 2.1.1.17. For the free monoid in one variable, we recover the polynomial ring and the ring of formal power series.

Definition 2.1.1.18. Let R be a ring. A <u>left-R-module</u> is an Abelian group M together with a multiplication $R \times M \to M$ such that the following hold:

- $1m = m$ for all $m \in M$.
- $(a + b)m = am + bm$ for all $a, b \in R$ and all $m \in M$.
- $a(m + n) = am + an$ for all $a \in R$ and all $m, n \in M$.
- $(ab)m = a(bm)$ for all $a, b \in R$ and all $m \in M$.

A <u>right-R-module</u> is defined analogously.

An Abelian group together with a left-R-module and a right-S-module structure is an <u>R-S-bimodule</u> if $(am)b = a(mb)$ for all $a \in R$, $m \in M$, and $b \in S$.

A <u>submodules</u> of a (left-, right-, bi-) module is an additive subgroup that is closed with respect to multiplication by ring elements (so that it inherits a module structure).

Remark 2.1.1.19. Modules are for rings, what actions are for groups. Simple modules are somewhat like orbits.

Example 2.1.1.20. Let A be an Abelian group (written additively). The set of endomorphisms $\mathrm{End}(A) := \{\varphi : A \to A \,|\, \varphi \text{ is a group hom.}\}$ is a ring where

1. The sum of two homomorphisms φ and ψ is given by pointwise addition:
$$(\varphi + \psi)(g) := \varphi(g) + \psi(g)$$
2. Multiplication is composition of homomorphisms.
3. The zero-element is the trivial homomorphism sending every element to $0 \in A$.
4. The one-element is the identity homomorphism sending every element to itself.

The group of units in $\mathrm{End}(A)$ is the automorphism group $\mathrm{Aut}(A)$.

Moreover, the group A is a left $\mathrm{End}(A)$-module via

$$\begin{array}{ccc} \mathrm{End}(A) \times A & \longrightarrow & A \\ (\varphi, g) & \mapsto & \varphi(g) \end{array}$$

Exercise 2.1.2.4 shows that $\mathrm{End}(A)$ can be considered an analogue of the symmetric group and the module structure on A corresponds to the tautological action.

Exercise 2.1.1.21. Any ring R is an R-R-bimodule.

Definition 2.1.1.22. Let M be an R-module and fix a subset $\mathrm{X} \subseteq R$. We say that M is <u>generated by</u> X if M does not contain a proper submodule that contains X.

Note that for each subset $\mathrm{X} \subseteq M$, there is a unique submodule $S \leq M$ that is generated by X, namely the intersection of all submodules that contain X. We write $\langle \mathrm{X} \rangle$ for the submodule generated by X.

Exercise 2.1.1.23. Let M be a left-R-module. Show that, for any subset $\mathrm{X} \subseteq M$,

$$\langle \mathrm{X} \rangle = \left\{ \sum_{\chi \in \mathrm{X}} a_\chi \chi \;\middle|\; a_\chi \in R \text{ and all but finitely many } a_\chi \text{ vanish} \right\}$$

Exercise 2.1.1.24. Let $\mathcal{M}$ be a monoid, let R be a ring and let M be a left-R-module. Define

$$M[\mathcal{M}] := \{ g : \mathcal{M} \to M \mid g(\mu) = 0 \text{ for all but finitely many } \mu \in \mathcal{M} \}$$

Show that

$$\begin{array}{rcl} R[\mathcal{M}] \times M[\mathcal{M}] & \longrightarrow & M[\mathcal{M}] \\ (f, g) & \mapsto & \left(fg : \mu \mapsto \displaystyle\sum_{\mu = \nu_1 \nu_2} f(\nu_1)\, f(\nu_2) \right) \end{array}$$

defines a left-$R[\mathcal{M}]$-structure on $M[\mathcal{M}]$.

Similarly, suppose that $\mathcal{M}$ allows for the definition of $R[[\mathcal{M}]]$. Then

$$M[[\mathcal{M}]] := \{ g : \mathcal{M} \to M \}$$

is a left-$R[[\mathcal{M}]]$-module via

$$\begin{array}{rcl} R[[\mathcal{M}]] \times M[[\mathcal{M}]] & \longrightarrow & M[[\mathcal{M}]] \\ (f, g) & \mapsto & \left(fg : \mu \mapsto \displaystyle\sum_{\mu = \nu_1 \nu_2} f(\nu_1)\, g(\nu_2) \right) \end{array}$$

2.1.2 Homomorphisms and Ideals

Definition 2.1.2.1. Let R and S be rings. A homomorphism from R to S is a map

$$\varphi : R \longrightarrow S$$

that is compatible with with addition and multiplication, i.e.,

$$\varphi(a + b) = \varphi(a) + \varphi(b)$$

and

$$\varphi(ab) = \varphi(a)\, \varphi(b)$$

for all $a, b \in R$. Note that $\varphi(0_R) = 0_S$. If, in addition, $\varphi(1_R) = 1_S$ we say that φ is a unital homomorphism.

Example 2.1.2.2. Embedding $n \times n$ matrices as North-West corners into $(n+1) \times (n+1)$ matrices (with zeros in all other places) defines a non-unital ring homomorphism.

Exercise 2.1.2.3. Show that every surjective ring homomorphism is unital.

Exercise 2.1.2.4. Show that a left-R-module structure on an Abelian group A is given by a ring unital homomorphism $R \to \mathrm{End}(A)$ and that conversely, any such unital ring homomorphism defines a left-R-module structure on A.

Definition 2.1.2.5. Let R be ring. The ring R^{op} is has the same underlying additive group, but the law for multiplication is reversed: $a \times_{\mathrm{op}} b := ba$. An anti-homomorphism from R to S is a map that is a ring homomorphism from R to S^{op}.

Exercise 2.1.2.6. Show that right-R-module structures on A are in 1-1 correspondence to unital ring anti-homomorphisms $R \to \mathrm{End}(A)$.

Definition 2.1.2.7. An additive subgroup $I \subseteq R$ is a left-ideal if $ai \in I$ for any $a \in R$ and any $i \in I$. An additive subgroup I is a right-ideal if $ia \in I$ for any $a \in R$ and any $i \in I$. A subset that is simultaneously a left-ideal and a right-ideal is called a two-sided ideal.

Remark 2.1.2.8. When talking about modules and ideals, I will often forget the prefix. For modules, that means left-modules. For ideals, the statement is actually an abbreviation for all three interpretations, and all of them are supposed to hold (provided you interpret all occurrences of "ideal" consistently).

Observation 2.1.2.9. *A subset of R is a left-ideal if and only if it is a left-R-submodule of R.*

A subset of R is a right-ideal if and only if it is a right-R-submodule of R

A subset of R is a two-sided ideal if and only if it is a R-R-sub-bimodule of R.

Example 2.1.2.10. In any ring, the subsets $\{0\}$ and R are a two-sided ideals.

Example 2.1.2.11. Let $\varphi : R \to S$ be a homomorphism, and let $J \subseteq S$ be an ideal in S, then $\varphi^{-1}(J)$ is an ideal in R. In particular, the kernel

$$\ker(\varphi) := \{a \in R \mid \varphi(a) = 0_S\}$$

of φ is a two-sided ideal in R.

Observation and Definition. Since any ideal (left, right, or two-sided) I is an additive subgroup and since R is an Abelian group with respect to addition, the subgroup I is normal and we can form the quotient group ${}^{R}\!/_{I} := \{a + I \mid a \in R\}$. The following are easily checked by straight forward computations:

1. If I is a left-ideal in R, then the Abelian group R/I is a left-R-module via

$$\begin{array}{rcl} R \times R/I & \longrightarrow & R/I \\ (a, b+I) & \mapsto & ab+I \end{array}$$

2. If I is a right-ideal in R, then the Abelian group R/I is a right-R-module via

$$\begin{array}{rcl} R/I \times R & \longrightarrow & R/I \\ (a+I, b) & \mapsto & ab+I \end{array}$$

3. If I is a two-sided ideal in R, then the Abelian group R/I is a ring with multiplication

$$(a+I)(b+I) := (ab+I)$$

zero $0+I$ and one $1+I$. This ring is called the quotient ring. The natural projection

$$\begin{aligned} \pi_I : R &\longrightarrow {}^{R}\!/_{I} \\ a &\mapsto a+I \end{aligned}$$

is a unital homomorphism (and clearly onto). Its kernel is I.

Theorem 2.1.2.12 (First Isomorphism Theorem). *Let $\varphi : R \to S$ be a ring homomorphism with kernel $I := \ker(\varphi)$. Then, φ factors as*

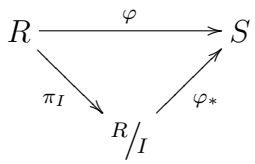

where φ_* is 1-1 and uniquely determined by φ. Also, φ_* is onto if and only if φ is onto; and φ_* is unital if and only if φ is unital.

Proof. Exercise. **q.e.d.**

Definition 2.1.2.13. An ideal that is generated (as a submodule) by a single element, is called a principal ideal.

Definition 2.1.2.14. Let R be commutative. An proper ideal I is called a prime ideal if $ab \in I$ implies that $a \in I$ or $b \in I$; i.e., whenever the ideal contains a product, it contains at least one of the factors.

The set of prime ideals in R is called the spectrum of R and denoted by $\mathrm{Spec}(R)$.

Definition 2.1.2.15. A proper ideal M in the ring R is maximal if there are no proper ideals in R of which M is a proper subset.

Observation 2.1.2.16. *The union along an ascending chain of proper ideals is a proper ideal. Thus, using Zorn's lemma, we see that every proper ideal is contained in a maximal ideal.* **q.e.d.**

We describe a source of prime ideals:

Proposition 2.1.2.17. *Let R be a commutative ring and let $L \subseteq R - \{0\}$ be non-empty and multiplicative (i.e., for any $a, b \in L$, we have $ab \in L$). Then any maximal element of*

$$\{I \trianglelefteq R \,|\, I \cap L = \emptyset\}$$

is a prime ideal. (Note again that Zorn's lemma implies that any element of this set is contained in a maximal element of this set.)

Proof. Let $M \in \mathcal{I} := \{I \trianglelefteq R \,|\, I \cap L = \emptyset\}$ be maximal. Assume $ab \in M$. If $a \notin M$, then the ideal generated by M and a is not in $\mathcal{I}$, i.e., it intersects L. In other words: there are elements $c_a \in R$, $m \in M$ such that $c_a a + m \in L$. Similarly, if $b \notin M$, then there are elements $c_b \in R$, $n \in M$ such that $c_b b + n \in L$. As L is multiplicative, we find

$$L \ni (c_a a + m)(c_b b + n) = c_a a c_b b + c_a a n + m c_b b + mn \in M$$

This is a contradiction. **q.e.d.**

Definition 2.1.2.18. Let I and J be two (left) ideals. Then, the additive subgroup generated by all products ij (where $i \in I$ and $j \in J$) is a (left) ideal, called the product ideal IJ.

Theorem 2.1.2.19 (Correspondence and Second Isomorphism Theorem). *Let I be a two-sided ideal in R. Then, there is a 1-1 correspondence*

$$\begin{array}{ccc} \left\{\text{ideals in } {}^{R}\!/_{I}\right\} & \longrightarrow & \{\text{ideals in } R \text{ containing } I\} \\ \tilde{J} & \mapsto & \pi_I^{-1}\left(\tilde{J}\right) \end{array}$$

In other words, any ideal in ${}^{R}\!/_{I}$ is of the form ${}^{J}\!/_{I}$ for some ideal J containing I.

Moreover, for any ideal J containing I, we have an isomorphism

$$R\Big/_J \cong {}^{R}\Big/_{I}\Big/_{J/I}.$$

The correspondence preserves prime ideals and maximal ideals in both directions. Principal ideals in R above I are send to principal ideals in R/I, and the same is true for finitely generated ideals.

Proof. Exercise. **q.e.d.**

Corollary 2.1.2.20. *Let R be commutative and let $I \trianglelefteq R$ be an ideal. Then I is maximal if and only if R/I is a field. The ideal I is prime if and only if R/I is a domain. In particular, every maximal ideal is prime. (We knew that already.)*

Proof. Exercise. **q.e.d.**

Theorem 2.1.2.21 (Yet Another Isomorphism Theorem). *Let R be a ring and $S \leq R$ be a subring. Let $I \trianglelefteq R$ be a two-sided ideal in R. Then*

1. *$S \cap I$ is a two-sided ideal in S.*
2. *$S + I$ is a subring in R.*
3. *${}^{S}\big/_{S\cap I} = {}^{S+I}\big/_{I}$.*

Proof. Exercise. **q.e.d.**

2.1.3 Module Homomorphisms

Definition 2.1.3.1. Let R be a ring and let M and N be two left-R-modules. A module homomorphism is a homomorphism of Abelian groups

$$\varphi : M \longrightarrow N$$

that is compatible with multiplication:

$$\varphi(am) = a\varphi(m) \qquad \text{for all } a \in R \text{ and } m \in M.$$

Homomorphisms of right-R-modules are defined analogously.

Let R and S be rings and let M and N be two R-S-bimodules. A bimodule homomorphism is a homomorphism of Abelian groups

$$\varphi : M \longrightarrow N$$

that is compatible with multiplication:

$$\varphi(amb) = a\varphi(m)\, b \qquad \text{for all } a \in R \text{ and } m \in M \text{ and } b \in S.$$

The kernel

$$\ker(\varphi) := \{m \in M \mid \varphi(m) = 0\}$$

is a (left-, right-, bi-) submodule of M, and the image

$$\mathrm{im}(\varphi) := \{\varphi(m) \mid m \in M\}$$

is a (left-, right-, bi-) submodule of N.

Correspondence Theorem 2.1.3.2. *Let S be a submodule of the (left-, right-, bi-) module M. Then M/S is a (left-, right-, bi-) module over the same ring(s), and*

$$\pi : M \longrightarrow M/S$$

is a module epimorphism.

Moreover, π induces a 1-1*-corresponcence*

$$\{S' \mid S \leq S' \leq M\} \longleftrightarrow \left\{\tilde{S} \,\middle|\, \tilde{S} \leq M/S\right\}.$$

Isomorphism Theorems 2.1.3.3. *All modules are consistently left-, right-, or bi-modules over fixed rings.*

1. *Suppose we have a commutative diagram with short exact rows:*

$$\begin{array}{ccccc} S & \hookrightarrow & M & \xrightarrow{\pi_1} & Q_1 \\ \| & & \| & & \\ S & \hookrightarrow & M & \xrightarrow[\pi_2]{} & Q_2 \end{array}$$

Then there is a unique homomorphism $\varphi : Q_1 \to Q_2$ *such that*

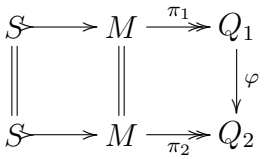

commutes. Moreover, φ *is an isomorphism.*

2. *Let* $S \leq M \leq N$ *be a chain of module inclusions. Then*

$$N/S\Big/_{M/S} \cong N/M$$

3. *Let* S_0 *and* S_1 *be submodules of* M*. Then*

$$S_0 + S_1 := \{m_0 + m_1 \mid m_0 \in S_0 \text{ and } m_1 \in S_1\}$$

is a submodule of M*, and*

$$S_0 + S_1\Big/_{S_0} \cong S_1\Big/_{S_0 \cap S_1}$$

All isomorphisms are naturally induced by looking at representatives.

Proofs are straight forward.

Exercise 2.1.3.4. Let M be a left-R-module. Show that for any element $m \in M$, the annihilator

$$\mathrm{Ann}(m) := \{a \in R \mid am = 0\}$$

is a left-ideal in R.

Show that the intersection

$$\mathrm{Ann}(M) := \{a \in R \mid am = 0 \text{ for all } m \in M\} = \bigcap_{m \in M} \mathrm{Ann}(m)$$

is a two-sided ideal in R.

Let $S \leq M$ be a submodule. Show that M/S is an $R/\mathrm{Ann}(S)$-module.

2.2 Non-Commutative Rings

2.2.1 Noetherian Rings and Modules

Proposition and Definition 2.2.1.1. Let R be a ring. A left-R-module M is noetherian if it satisfies the following equivalent conditions:

1. Every non-empty collection $\mathcal{C}$ of submodules in M contains maximal elements.
2. Every submodule S of M is finitely generated.
3. M does not admit an infinite strictly ascending chain

$$S_0 < S_1 < S_2 < \cdots$$

of submodules.

Proof. (1) $\implies$ (2) Put

$$\mathcal{C} := \{T \leq S \mid T \text{ is finitely generated}\}$$

and let $T_{\max}$ be a maximal element of $\mathcal{C}$, i.e, a maximal finitely generated submodule of S. For any element $m \in S$, the module generated by $T_{\max}$ and m is finitely generated, is contained in S and does contain $T_{\max}$. By maximality of $T_{\max}$, we have $m \in T$ and $T_{\max} = S$.

(2) $\implies$ (3) Suppose there was an infinite strictly increasing chain

$$S_0 < S_1 < S_2 < \cdots$$

and assume the union

$$S := \bigcup_i S_i = \langle m_1, \ldots, m_u \rangle$$

is finitely generated. Then we obtain a contradiction since there is an index j such that $S_j = S$ because at some finite stage all generators made it into the union.

(3) $\implies$ (1) Suppose $\mathcal{C}$ does not contain a maximal element. Since $\mathcal{C}$ is non-empty, we can choose $S_0 \in \mathcal{C}$ and since S_0 is not maximal, the sub-collection $\mathcal{C}_1 := \{S \in \mathcal{C} \mid S_0 < S\}$ is non-empty and has no maximal elements. Inductively, we could then construct an infinite strictly ascending chain. **q.e.d.**

Exercise 2.2.1.2. Let R be a ring, and let

$$S \hookrightarrow M \twoheadrightarrow Q$$

be a short exact sequence of left-R-module. Show that M is noetherian if and only if both S and Q are noetherian.

Definition 2.2.1.3. A ring is called left-noetherian if it is noetherian as a left-R-module. It is called right-notherian if it is noetherian as a right-R-module. It is called noetherian if it is simultaneously left- and right-noetherian

Exercise 2.2.1.4. Proof or disprove: a ring R is noetherian if and only if it is noetherian as an R-R-bimodule.

Theorem 2.2.1.5 (Hilberts Basis Theorem). *Let $\mathbb{N}_0$ denote the monoid of non-negative integers with addition as binary operation. Let R be a ring and let M be a noetherian left-R-module. Then $M[\mathbb{N}_0]$ is a noetherian $R[\mathbb{N}_0]$-module.*

Proof. For any non-zero element f of $M[\mathbb{N}_0]$ or $R[\mathbb{N}_0]$, we define its degree to be the maximum element $n \in \mathbb{N}_0$ for which $f(n) \neq 0$. We call the value $f(n)$ at the degree, the leading coefficient.

Let $S \leq M[\mathbb{N}_0]$ be an infinitely generated submodule. Let f_0 be a non-zero element of minimum degree. Since S is not finitely generated, $S - \langle f_0 \rangle \neq \emptyset$. Let f_1 be an element of $S - \langle f_0 \rangle$ of minimum degree. Since S is not finitely generated, $S - \langle f_0, f_1 \rangle \neq \emptyset$. Continue and define $f_2, f_3, \ldots$

I claim that the leading coefficients, $m_i := f_i(\deg(f_i))$ generate a submodule of M which is not finitely generated. Otherwise, there is an index j such that

$$m_j \in \langle m_i \,|\, i < j \rangle$$

i.e., there are ring elements $a_0, a_1, \ldots$ such that

$$m_j = \sum_{i<j} a_i m_i.$$

Define $g_j := a_j \chi_{\deg(f_j)-\deg(f_i)}$, i.e.:

$$\begin{array}{rcl} g_i : \mathbb{N}_0 & \longrightarrow & R \\ n & \mapsto & \begin{cases} a_i & \text{if } \deg(f_j) = n + \deg(f_i) \\ 0 & \text{otherwise} \end{cases} \end{array}$$

Then, it is easy to check that

$$f_j - \sum_{i<j} g_i f_i \in S - \langle f_i \,|\, i < j \rangle$$

has smaller degree than f_j contrary to our construction principle. **q.e.d.**

Corollary 2.2.1.6. *If R is (left-) noetherian, then so is $R[\mathbb{N}_0]$.* **q.e.d.**

Exercise 2.2.1.7. Suppose that $\mathcal{M}$ is a monoid. We say that μ is a right-divisor of ν there is a complementary divisor $\mu_{\text{compl}} \in \mathcal{M}$ with $\mu_{\text{compl}}\mu = \nu$. Suppose that right-divisibilty defines a total order on $\mathcal{M}$, i.e, for any two elements, one is a right-divisor of the other, and if this holds either way, both elements are equal.

Prove or disprove the folowing generalization of Hilbert's basis theorem: if M is a noetherian left-R-module, then $M[\mathcal{M}]$ is a noetherian left-$R[\mathcal{M}]$-module.

Exercise 2.2.1.8. Let R be a ring and let M be a noetherian left-R-module. Then $M[[\mathbb{N}_0]]$ is a noetherian $R[[\mathbb{N}_0]]$-module.

2.2.2 Artinian Rings and Modules

Proposition and Definition 2.2.2.1. A (left-, right-, bi-) module M is artinian if it satisfies the following equivalent conditions:

1. Every non-empty collection $\mathcal{C}$ of submodules in M contains minimal elements.
2. M does not admit an infinite strictly descending chain
$$S_0 > S_1 > S_2 > \cdots$$
of submodules.

Proof of equivalence. Exercise. **q.e.d.**

Example 2.2.2.2. Let K be a field. Any finite dimensional K-module is an artinian K-module: since all submodules are vector spaces, the dimension along a strictly descending chain has to go down at each step. Thus, the chain cannot be infinite.

Exercise 2.2.2.3. Let
$$S \hookrightarrow M \twoheadrightarrow Q$$
be a short exact sequence of left-R-modules. Show that M is artinian if and only if S and Q are both artinian.

Exercise 2.2.2.4. Let D be a division ring. Show that D^m is an artinian left-D-module.

Definition 2.2.2.5. A ring R is left-artinian if R is an artinian left-R-module. It is right-artinian if R is an artinian right-R-module.

2.2.3 Simple Rings and Modules

Definition 2.2.3.1. A (left-, right-, bi-) module is simple if $\{0\}$ is a maximal submodule, i.e., the module is non-trivial and does not contain proper non-trivial submodules.

A ring is simple if $\{0\}$ is a maximal two-sided ideal, i.e., the ring does not contain proper non-trivial two-sided ideals.

Example 2.2.3.2. Let K be a field. Then K^m is a simple left-$\mathbb{M}_{m\times m}(K)$-module.

Exercise 2.2.3.3. Let D be a division ring. Show that D^m is a simple left-$\mathbb{M}_{m\times m}(D)$-module.

Observation 2.2.3.4. *Let U be a simple left-R-module. Note that the image and the kernel of any homomorphism are submodules of the target and the domain, respectively. Hence every non-trivial homomorphism into U is onto; and every non-trivial homomorphism defined on U is* 1-1.

In particularl, any non-trivial endomorphism of a simple module is 1-1 *and onto, i.e., invertible. We infer that the set of endomorphisms* $\mathrm{End}_R(U)$ *is a division ring: (1) it is a ring in the way illustrate in (??) and (2) every non-zero endomorphism is invertible.* **q.e.d.**

Observation 2.2.3.5. *Let $I \trianglelefteq R$ be a left-ideal and let M be a left-R-module. For any element $m \in M$, the map*

$$\begin{array}{rcl} \rho_m : I & \longrightarrow & M \\ i & \mapsto & im \end{array}$$

is an R-module homomorphism. **q.e.d.**

Lemma 2.2.3.6. *Let U be a simple left-R-module. Let $\mathbf{b}_1, \mathbf{b}_2, \ldots$ be a sequence of non-trivial elements of U. Put*

$$I_k := \bigcap_{i<k} \mathrm{Ann}(\mathbf{b}_i)$$

where $I_0 = R$. Then for any $m \in U$ with $I_k \subseteq m$, there are endomorphisms $\varphi_i : U \to U$ for any $i < k$ such that

$$m = \sum_{i<k} \varphi_i(\mathbf{b}_i)\,.$$

Proof. We use induction. The start is purely formal. So assume $I_k \cap \mathrm{Ann}(\mathbf{b}_k) = I_{k+1} \subseteq \mathrm{Ann}(m)$. Consider the module homomorphisms

$$\rho_{\mathbf{b}_k} : I_k \to U$$

and

$$\rho_m : I_k \to U$$

If $\rho_{\mathbf{b}_k}$ is trivial, we have $I_k \subseteq \mathrm{Ann}(\mathbf{b}_k)$ in which case $I_k = I_{k+1} \subseteq \mathrm{Ann}(m)$, whence the claim follows by induction. Thus, we may assume that $\rho_{\mathbf{b}_k}$ is onto (U is simple). By hypothesis, $\ker(\rho_{\mathbf{b}_k}) \subseteq \ker(\rho_m)$. Hence we have an induced map

$$\varphi_k : U \longrightarrow U$$

such that

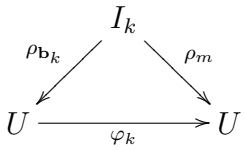

commutes. Note that for each $i \in I_k$,

$$i(m - \varphi_k(\mathbf{b}_k)) = im - \varphi_k(i\mathbf{b}_k) = im - \varphi_k(\rho_{\mathbf{b}_k}(i)) = im - \rho_m(i) = im - im = 0$$

Thus $I_k \subseteq \mathrm{Ann}(m - \varphi_k(\mathbf{b}_k))$ and by induction hypotheses, there are endomorphisms φ_i for $i < k$ with

$$m - \varphi_k(\mathbf{b}_k) = \sum_{i<k} \varphi_i(\mathbf{b}_i)$$

The claim follows. **q.e.d.**

Lemma 2.2.3.7. *Let R be a left-artinian ring, and let U be a simple left-R-module. Then U is a vector space over $D := \mathrm{End}_R(U)$ of finite dimension.*

Proof. Consider the collection of left-ideals

$$\mathcal{C} := \left\{ \bigcap_{\mathbf{b}\in B} \mathrm{Ann}(\mathbf{b}) \;\middle|\; B \subseteq U \text{ is finite} \right\}.$$

Since R is left-artinian, the collection has a minimal element realized by some finite set B. Then, by minimality,

$$\bigcap_{\mathbf{b}\in B} \mathrm{Ann}(\mathbf{b}) \subseteq \mathrm{Ann}(m)$$

for each element $m \in U$. Thus

$$m = \sum_{\mathbf{b}\in B} \varphi_{\mathbf{b}}(\mathbf{b})$$

for suitably chosen $\varphi_{\mathbf{b}} \in D$. **q.e.d.**

Corollary 2.2.3.8 (Jacobsons Density Theorem). *Let R, U and D be as above, and let B be a finite D-linearly independent subset of U. Then for each D-linear map*

$$\varphi : U \longrightarrow U$$

there exists an $a \in R$ such that $\varphi(\mathbf{b}) = a\mathbf{b}$ for all $\mathbf{b} \in B$.

Proof. We use induction on the size of B. The case of $B = \emptyset$ is formal. So assume that $B = B' \cup \{\mathbf{b}\}$ with $\mathbf{b}$ linearly independent from B'; and assume and that we already found a ring element a with $\varphi(\mathbf{b}') = a\mathbf{b}'$ for all $\mathbf{b}' \in B'$. Note that we are free to choose a within $a + \mathrm{Ann}(B')$ where

$$\mathrm{Ann}(B') := \bigcap_{\mathbf{b}'\in B'} \mathrm{Ann}(\mathbf{b}') = \mathrm{Ann}\left(\sum_{\mathbf{b}'} D\mathbf{b}'\right)$$

Thus, all we need to find is an element $i \in \mathrm{Ann}(B')$ with

$$(a + i)\mathbf{b} = \varphi(\mathbf{b})\,.$$

Since $\mathbf{b}$ is linearly independent from B', we have $\mathbf{b}$ does not lie in the D-span of B'. Thus, $\mathrm{Ann}(B') \not\subseteq \mathrm{Ann}(\mathbf{b})$ and we have $U = \mathrm{Ann}(B')\,\mathbf{b}$. Hence there is $i \in \mathrm{Ann}(B')$ with $i = \varphi(\mathbf{b}) - a\mathbf{b}$. **q.e.d.**

Proposition 2.2.3.9. *Let R be a simple artinian ring. Then R is (isomorphic to) the endomorphism ring $\mathrm{End}_D(V)$ of a finite dimensional vector space V over a division ring D.*

Proof. Let U be a simple left-R-module (obtain that from a maximal left-ideal). Put $D := \mathrm{End}_R(U)$. Then U is a D-vector space of finite dimension. Moreover, the map

$$\begin{array}{rcl} R & \longrightarrow & \mathrm{End}_D(U) \\ a & \mapsto & \lambda_a : m \mapsto am \end{array}$$

is a unitial ring homomorphism. Since R is simple, it is 1-1. By Jacobsons Density Theorem (2.2.3.8), it is onto. **q.e.d.**

Exercise 2.2.3.10. Let D be a division ring, and let $V = D^m$ be the left-D-vector space of dimension m. Show:

$$\mathrm{End}_D(V) \cong \mathbb{M}_m(D^{\mathrm{op}})$$

Exercise 2.2.3.11. Let D be a division ring. Show that $\mathbb{M}_m(D)$ is a simple left-artinian ring.

Exercise 2.2.3.12. Let R be a ring. Show that any two minimal left-ideals in R are isomorphic as left-R-modules.

Observation 2.2.3.13. *Let M be an R-S-bimodule. Then we have a ring homomorphism*

$$\begin{array}{rcl} R & \longrightarrow & \mathrm{End}_S(M) \\ a & \mapsto & (\lambda_a : m \mapsto am) \end{array}$$

induced by left-multiplication. We also have a ring antihomomorphism

$$\begin{array}{rcl} S & \longrightarrow & \mathrm{End}_R(M) \\ b & \mapsto & (\rho_a : m \mapsto mb) \end{array}$$

induced by right-multiplication. (This is an anti-homomorphism because, we always let endomorphisms act from the left, thus there is a change of direction.) **q.e.d.**

Example 2.2.3.14. Let D be a division ring. Consider D^m as a $\mathbb{M}_m(D)$-D-bimodule. Thus, we have a ring antihomomorphism

$$D \longrightarrow \mathrm{End}_{\mathbb{M}_m(D)}(D^m)$$

which is non-trivial and hence 1-1.

To see that this homomorphism is onto, as well, let $B = \{\mathbf{e}_1, \ldots, \mathbf{e}_m\}$ be the standard basis of D^m, and let $\varphi : D^m \to D^m$ be any $\mathbb{M}_m(D)$-endomorphism of D^m. Then, using the right-D-vector space structure on D^m, we cann write

$$\varphi(\mathbf{e}_1) = \sum_{i=1}^{m} \mathbf{e}_i a_i$$

for some elements $a_i \in D$. Then

$$\varphi(\mathbf{e}_j) = \varphi\big(\mathbf{E}^{j1}\mathbf{e}_1\big) = \mathbf{E}^{j1}\varphi(\mathbf{e}_1) = \mathbf{E}^{j1}\sum_{i=1}^{m} \mathbf{e}_i a_i = \mathbf{E}^{j1}\mathbf{e}_1 a_1 = \mathbf{e}_j a_1$$

In other words: $\varphi = \rho_{a_1}$.

Thus,

$$D \cong \mathrm{End}_{\mathbb{M}_m(D)}(D^m)^{\mathrm{op}}.$$

Exercise 2.2.3.15. Let D be a division ring. Show that the scheme

$$\begin{pmatrix} * & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ * & 0 & \cdots & 0 \end{pmatrix}$$

defines a minimal left-ideal in $\mathbb{M}_m(D)$.

Theorem 2.2.3.16 (Wedderburn). *Every simple artinian ring R is isomorphic to a ring $\mathbb{M}_m(D)$ where D is a division ring. Moreover, m is uniquely determined by R and D is uniquely determined up to unital isomorphism.*

Proof. Existence of D and m follows from (2.2.3.9) and (2.2.3.10). As for uniqueness, let V be a minimal left-ideal in R. Those exists since R is artinian and are unique up to R-module

isomorphisms by (2.2.3.12). Moreover, the set of all $m \times m$ matrices over D all of whose columns vanish with the possible exception of the first is a minimal left-ideal of $\mathbb{M}_m(D)$. Thus, any minimal left-ideal in R is isomorphic to D^m. Thus, (2.2.3.14) allows us to recover D from R:

$$D \cong \mathrm{End}_R(V)^{\mathrm{op}}$$

We now can regard V as a D-vector and recover m as its dimension. **q.e.d.**

The following is immediate from correspondence theorem and the defintions:

Observation 2.2.3.17. *If M is a maximal left-ideal in R, then R/M is a simple module and $M = \mathrm{Ann}(1 + M)$.* **q.e.d.**

Here is a converse to (2.2.3.17)

Lemma 2.2.3.18. *Let U be a simple left-R-module. For any non-zero $m \in U$, the annihilator $\mathrm{Ann}(m)$ is a maximal left-ideal in R. Moreover, the map*

$$\begin{array}{ccc} R/\mathrm{Ann}(m) & \longrightarrow & U \\ a + \mathrm{Ann}(m) & \mapsto & am \end{array}$$

defines an isomorphism of left-R-modules.

Proof. Consider the map

$$\begin{array}{ccc} R & \longrightarrow & U \\ a & \mapsto & am \end{array}$$

as a homomorphism of left-R-modules. Since $m \neq 0$, its image is a non-trivial submodule of the simple module U. Thus, the map is onto. The kernel is clearly the annihilator $\mathrm{Ann}(m)$. This already proves the second statement. The first statement follows by correspondence theorem from the fact that $R/\mathrm{Ann}(m)$ is simple. **q.e.d.**

Definition 2.2.3.19. Let R be a ring. Call an ideal I a simple annihilator if there exists a simple left-R-module U with $I = \mathrm{Ann}(U)$. Recall that annihilators of modules (as opposed to annihilators of individual elements) are two-sided ideals. The Jacobson radical is the intersection of all simple annihilator ideals:

$$\mathrm{J}(R) := \bigcap_{I \text{ is simple annihilator}} I.$$

It is a two-sided ideal.

Proposition 2.2.3.20. *The Jacobson radical of R is the intersection of all maximal left-ideal of R.*

Proof. By (2.2.3.17) and (2.2.3.18), maximal left-ideal arise precisely as annihilators of non-zero elements in simple modules. The annihilator of a simple module, in turn is the intersection of all annihilators of its non-zero elements. **q.e.d.** **q.e.d.**

Definition 2.2.3.21. An element $a \in R$ is called left-quasiregular if $(1-a)$ has a left-inverse, i.e.,

$$1 = b(1-a)$$

for some $b \in R$. Equivalently, a is left-quasiregular, if the left-ideal generated by $(1-a)$ is all of R.

Lemma 2.2.3.22. *Every element of the Jacobson radical $\mathrm{J}(R)$ is left-quasiregular.*

Proof. Assume that $a \in \mathrm{J}(R)$. Then, a is contained in each maximal left-ideal of R. Thus, $1-a$ does not belong to any maximal left-ideal. Thus, the left-ideal generated by $1-a$ is all of R. **q.e.d.**

Lemma 2.2.3.23. *If a left-ideal $I \trianglelefteq R$ consists entirely of left-quasiregular elements, then $I \subseteq \mathrm{J}(R)$.*

Proof. ??? **q.e.d.**

2.3 Commutative Rings

2.4 Constructions

2.4.1 Polynomials and Power Series

Recall the examples (2.1.1.15) and (2.1.1.16).

Definition 2.4.1.1. Let $\mathbb{N}_0$ be the monoid of non-negative integers with addition as binary operation, let R be a ring, and let M be a left-R-module. The power series ring over R is $R[[\mathbb{N}_0]]$. This ring is usually denoted by $R[[x]]$ (where x is an arbitrary letter denoting a variable) and its elements are written as power series

$$\sum_{i=0}^{\infty} a_i x^i$$

where this series corresponds to the map $i \mapsto a_i$ in $R[[\mathbb{N}_0]]$. This is reasonable since under this correspondence, the Cauchy product of power series corresponds to the convolution product in $R[[\mathbb{N}_0]]$.

The subring $R[\mathbb{N}_0] \leq R[[\mathbb{N}_0]]$ is called the polynomial ring over R. This ring is customarily denoted by $R[x]$ and its elements are written as polynomials in x. Recall that polynomials in x can be regarded as power series that have only finitely many non-zero coefficients and ordinary multiplication of polynomials is nothing but the Cauchy product.

We also put $M[[x]] := M[[\mathbb{N}_0]]$ and $M[x] := M[[\mathbb{N}_0]]$.

The degree of a non-zero element in $R[x]$ or $M[x]$ is the highest degree with a non-vanishing coefficient. We define the degree of 0 to be $-\infty$. Then $\deg(pm) \leq \deg(p) + \deg(m)$ for any $p \in R[x]$ and any $m \in M[x]$. If $\mathrm{Ann}(M) = \{0\}$, equality holds. (Note that $\mathrm{Ann}(R) = \{0\}$ if R does not contain zero divisors.)

The degree of a non-zero element in $R[[x]]$ or $M[[x]]$ is the lowest degree with a non-vanishing coefficient. Now, we define the degree of 0 to be ∞. We have $\deg(pm) \geq \deg(p) + \deg(m)$ for any $p \in R[[x]]$ and any $m \in M[[x]]$. Equality always holds if $\mathrm{Ann}(M) = \{0\}$.

Exercise 2.4.1.2. Compute, if possible, $\mathrm{Ann}_{R[x]}(M[x])$ and $\mathrm{Ann}_{R[[x]]}(M[[x]])$ in terms of $\mathrm{Ann}_R(M)$.

Observation 2.4.1.3. *The group of units in $R[[x]]$ consists of exactly those power series whose constant term is a unit in R.*

Observation 2.4.1.4. *Suppose R does not contain zero divisors. The group of units in $R[x]$ consists precisely of the degree 0 polynomials whose constant term is a unit.*

Exercise 2.4.1.5. Show that R is and integral domain if and only if $R[[x]]$ is an integral domain.

Exercise 2.4.1.6. Show that R is and integral domain if and only if $R[x]$ is an integral domain.

Exercise 2.4.1.7. Let $\mathcal{M}$ be the monoid of strictly positive integers with multiplication as its binary operation. Show that the monoid ring $R[\mathcal{M}]$ is a domain if and only if R is a domain.

Exercise 2.4.1.8. Let G be a finite, non-trivial group. Show that the group ring $R[G]$ is never a domain.

Exercise 2.4.1.9. Let K be a field. Show that $K[[x]]$ has a unique maximal ideal.

Exercise 2.4.1.10. Let R be commutative. Show that for each $a \in R$ there is a unique unital ring homomorphism (called evaluation at a)

$$\varphi_a : R[x] \longrightarrow R$$

satisfying $\varphi_a(b) = b$ for any $b \in R \subset R[x]$ and $\varphi_a(x) = a$.

2.4.2 Localization

In this section on localization, all rings are **commutative**.

Definition 2.4.2.1. Let $L \subseteq R$ be a subset of the commutative ring R. We say that L is multiplicative if

1. $1 \in L$,
2. $0 \notin L$, and
3. L is closed with respect to multiplication, i.e., for any $r, s \in L$, we have $rs \in L$.

Example 2.4.2.2. Let P be a prime ideal in R. Then $R - P$ is a multiplicative set. In particular, $R - \{0\}$ is multiplicative.

Construction 2.4.2.3. Let $L \subset R$ be a multiplicative subset of the commutative ring R. We define an equivalence relation on $R \times L$ via

$$(a, r) \equiv (b, s) \text{ if and only if there is } t \in L \text{ with } trb = tsa.$$

It is a straight forward check that $\equiv$ is an equivalence relation. The equivalence classes are called fractions. The $\equiv$-equivalence class of (a, r) is denoted by $\frac{a}{r}$.

The set of $\equiv$-equivalence classes is denoted by

$$L^{-1}R := \left\{ \frac{a}{r} \,\middle|\, a \in R, r \in L \right\}.$$

We define addition and multiplication for fractions as follows:

$$\begin{aligned} \frac{a}{r} + \frac{b}{s} &:= \frac{as + rb}{rs} \\ \frac{a}{r}\frac{b}{s} &:= \frac{ab}{rs} \end{aligned}$$

As L is multiplicative, the results are actually fractions. A straight forward computation shows (a) that addition and multiplication are well-defined and (b) that $L^{-1}R$ is a commutative ring with these arithmetic operations. The zero element is $\frac{0}{1}$ and the one is $\frac{1}{1}$.

Definition 2.4.2.4. The commutative ring $L^{-1}R$ is called the localization of R with respect to L.

When doing the following exercises pay close attention as to where you need to use that R is commutative.

Exercise 2.4.2.5. Check that $\equiv$ is an equivalence relation.

Exercise 2.4.2.6. Show that addition and multiplication of fractions is well-defined, i.e., independent of the choice of representatives.

Exercise 2.4.2.7. Carry out the verification that $L^{-1}R$ is a commutative ring.

Construction 2.4.2.8. Let M be a left-R-module. The localization $L^{-1}M$ is constructed as follows:

1. Define an equivalence relation on $M \times L$ as follows

   $$(m_1, r_1) \equiv (m_2, r_2)$$

   if there is an element $s \in L$ such that $sr_2m_1 = sr_1m_2$. We denote the equivalence class of the pair (a, r) as a fraction $\frac{a}{r}$.

2. As a set, $L^{-1}M$ is the set of fractions, i.e., the set of equivalence classes in $M \times L$. One checks that

   $$\frac{m}{r} + \frac{n}{s} := \frac{sm + rn}{rs}$$

   and

   $$\frac{a}{r} \times \frac{m}{s} := \frac{am}{rs}$$

   define the structure of a left-$L^{-1}R$-module structure on $L^{-1}M$.

Remark 2.4.2.9. Particularly useful is localization to obtain a vector space over the field of fractions.

Proposition 2.4.2.10. *Let $\varphi : M \to N$ be a module homomorphism. Then, there is an induced module homomorphism*

$$\begin{array}{rcl} \varphi_L : L^{-1}M & \longrightarrow & L^{-1}N \\ \dfrac{m}{r} & \mapsto & \dfrac{\varphi(m)}{r} \end{array}$$

Proof. Straight forward check. **q.e.d.**

Proposition 2.4.2.11. *Let R be a commutative ring and let $L \subset R$ be multiplicative. Then*

$$\begin{array}{rcl} \iota_L : R & \longrightarrow & L^{-1}R \\ a & \mapsto & \dfrac{a}{1} \end{array}$$

is a unital ring homomorphism with

$$\ker(\iota_L) = \bigcup_{r \in L} \mathrm{Ann}_R(r) = \{a \in R \,|\, ra = 0 \text{ for some } r \in L\}\,.$$

If R is a domain, then the localization map ι_L is 1-1.

Proof. To verify that ι_L is a ring homomorphism, we observe

$$\iota_L(a+b) = \frac{a+b}{1} = \frac{1a+b1}{1} = \frac{a}{1} + \frac{b}{1} = \iota_L(a) + \iota_L(b)$$

and

$$\iota_L(ab) = \frac{ab}{1} = \frac{ab}{11} = \frac{a}{1}\frac{b}{1} = \iota_L(a)\,\iota_L(b)\,.$$

We compute the kernel:

$$\begin{array}{rcl} \ker(\iota_L) & = & \left\{ a \in R \,\middle|\, \iota_L(a) = \dfrac{0}{1} \right\} \\ & = & \left\{ a \in R \,\middle|\, \dfrac{a}{1} = \dfrac{0}{1} \right\} \\ & = & \{a \in R \,|\, ra = 0 \text{ for some } r \in L\} \end{array}$$

If L does not contain any zero divisors, it follows that $\ker(\iota_L) = \{0\}$. **q.e.d.**

Theorem 2.4.2.12 (Universal Property of Localizations). *Let L be a multiplicative subset of the commutative ring R. For any commutative ring S and any ring homomorphism $\varphi : R \to S$ with $\varphi(L) \subseteq S^*$, there is a unique ring homomorphism $\varphi_L : L^{-1}R \to S$ such that the diagram*

$$\begin{array}{ccc} R & \xrightarrow{\varphi} & S \\ {\scriptstyle \iota_L}\downarrow & \nearrow_{\varphi_L} & \\ L^{-1}R & & \end{array}$$

commutes.

Proof. As usual, uniqueness is easy:

$$\varphi_L\Big(\frac{a}{r}\Big) = \varphi_L\Big(\frac{a}{1}\frac{1}{r}\Big) = \varphi(a)\,\varphi(r)^{-1}$$

implies that φ determines φ_L.

Also, now we know how to prove existence: we just have to verify that

$$\varphi_L\Big(\frac{a}{r}\Big) := \varphi(a)\,\varphi(r)^{-1}$$

defines a ring homomorphism. Thus, one needs to check that this map is well-defined, i.e., independent of the choice of representatives for the fraction; and one needs to check that this map is a ring homomorphism, i.e., compatible with addition and multiplication. All three checks are simple straight forward calculations and left as an exercise. **q.e.d.**

Exercise 2.4.2.13. Fill in the computational checks in the preceding proof.

In the remainder of this section, we shall discuss the ideals in the localization $L^{-1}R$.

Proposition 2.4.2.14. *Let L be a multiplicative subset of the commutative ring R.*

1. *For any ideal $I \trianglelefteq R$, the set $L^{-1}I := \left\{\frac{i}{r} \,\middle|\, i \in I, r \in L\right\}$ is an ideal in $L^{-1}R$.*

2. *The ideal $L^{-1}I$ is a proper ideal of $L^{-1}R$ if and only if $I \cap L = \emptyset$.*

3. *If I is a principal ideal in R with generator i, then, $L^{-1}I$ is a principal ideal in $L^{-1}R$ with generator $\frac{i}{1}$.*

4. *For any ideal $\tilde{I} \trianglelefteq L^{-1}R$, the preimage $\tilde{I}_R := \iota_L^{-1}\left(\tilde{I}\right)$ is an ideal in R satisfying $\tilde{I} = L^{-1}\tilde{I}_R$. Note that the if R is a domain, the localization map is injective so that we can regard R as a subring of $L^{-1}R$. In this case, $\tilde{I}_R = \tilde{I} \cap R$.*

5. *There is a* 1-1 *correspondence*

$$\begin{aligned} \{P \in \operatorname{Spec}(R) \mid P \cap L = \emptyset\} &\longrightarrow \operatorname{Spec}\left(L^{-1}R\right) \\ P &\mapsto L^{-1}P \end{aligned}$$

This correspondence preserves inclusion relations among ideals.

Proof. Consider $\frac{i}{r}, \frac{j}{s} \in L^{-1}I$. We have

$$\frac{i}{r} + \frac{j}{s} = \frac{is + rj}{ij} \in L^{-1}I$$

since $is + ij \in I$ as I is a two-sided ideal. Closure with respect to multiplication by elements from $L^{-1}R$ is similar. This proves (1).

Note that if I contains an element of L, then the ideal $L^{-1}I$ contains $\frac{1}{1}$ and we have $L^{-1}I = L^{-1}R$. Conversely, if $L^{-1}I = L^{-1}R$. we have that $\frac{i}{r} = \frac{1}{1}$ for some $i \in I$ and some $r \in L$. Then for some $s \in L$, we find $si = r$ which implies that $r \in I$. This proves (2).

Claim (3) is clear.

Now, fix an ideal $\tilde{I} \trianglelefteq L^{-1}R$. Since preimages of ideals are ideals, we know that $\tilde{I}_R = \iota_L^{-1}\left(\tilde{I}\right)$ is an ideal in R. Thus, we only have to show that $\tilde{I} = L^{-1}\tilde{I}_R$. Since $\frac{i}{r} \in \tilde{I}$ implies $\frac{i}{1} = \frac{r}{1}\frac{i}{r} \in \tilde{I}$ whence $i \in \tilde{I}_R$, we find that $\tilde{I} \subseteq L^{-1}\tilde{I}_R$. The inclusion $L^{-1}\tilde{I}_R \subseteq \tilde{I}$ follows similarly as $i \in \tilde{I}_R$ implies that $\frac{i}{1} \in \tilde{I}$ whence $\frac{i}{r} = \frac{1}{r}\frac{i}{1} \in \tilde{I}$. This proves (4).

Finally, we prove (5). First, let p be a prime ideal in R not intersecting L. Then, $L^{-1}P$ is a proper ideal of $L^{-1}R$ by (2).

To see that $L^{-1}P$ is prime, we assume $\frac{a}{r}\frac{b}{s} = \frac{ab}{rs} \in L^{-1}P$. Then, $\frac{ab}{1} \in L^{-1}P$ whence $abt \in P$ for some $t \in L$. Since $t \notin P$, we infer $ab \in P$ and, since P is prime, it follows that $a \in P$ or $b \in P$. Thus $\frac{a}{r} \in L^{-1}P$ or $\frac{b}{s} \in L^{-1}P$. Thus, $L^{-1}P$ is a prime ideal.

Now assume that P and Q are two prime ideals in R both disjoint from L. Also, suppose that $L^{-1}P = L^{-1}Q$. We want to show that $P = Q$. So suppose that $i \in P$. Then $\frac{i}{1} \in L^{-1}P = L^{-1}Q$ which implies $\frac{i}{1} = \frac{j}{s}$ for some $j \in Q$ and some $s \in L$. Thus, for some $t \in L$, we have $tsi = tj \in Q$, which forces $i \in Q$ as $ts \in L$ and $L \cap Q = \emptyset$. Thus, we proved $P \subseteq Q$. The reversed inclusion is shown symmetrically. This shows that the map $P \mapsto L^{-1}P$ is 1-1.

This map is onto since preimages of prime ideals are prime ideals by (**??**). It clearly preserves inclusion relations of ideals. **q.e.d.**

Corollary 2.4.2.15. *If R is a PID, then so is $L^{-1}R$.*

Corollary 2.4.2.16. *If R is noetherian, then so is $L^{-1}R$.*

Theorem 2.4.2.17. *If R is a UFD, then so is $L^{-1}R$.*

Proof. It follows that every element $\frac{a}{r} \in L^{-1}R$ has a prime factor decomposition.

Recall that prime ideals in $L^{-1}R$ are in 1-1 correspondence with prime ideals in R that are disjoint from L. This correspondence restricts to principal prime ideals. Thus, the primes element in $L^{-1}R$ are (up to units) exactly the elements $\frac{p}{1}$ where $p \in R$ does not divide any element from L. Call those primes green. We call the other primes red and observe that they turn into units in $L^{-1}R$.

It follows that every element $\frac{a}{r} \in L^{-1}R$ has a prime factor decomposition.

Now suppose that

$$\alpha \frac{p_1}{1}\frac{p_2}{1}\cdots\frac{p_u}{1} = \frac{q_1}{1}\frac{q_2}{1}\cdots\frac{q_v}{1}$$

is a counter example to uniqueness of prime factor decomposition in $L^{-1}R$. The primes p_i and q_j are all green. The primes that go into the denominator and numerator of α are all red. Thus, we obtain an idenity in R

$$(\text{some red primes})p_1p_2\cdots p_u = (\text{some more red primes})q_1q_2\cdots q_v.$$

Since R is a UFD, we deduce that the green primes must correspond nicely. **q.e.d.**

2.5 Important Classes of Rings

2.5.1 Euclidean Domains

2.5.2 Principal Ideal Domains

Exercise 2.5.2.1. Let R be an integral domain. Let $p(x) \in R[x]$ be a non-zero polynomial whose leading coefficient is a unit in R. Show that for any polynomial $q(x) \in R[x]$ there are unique $f(x), r(x) \in R[x]$ satisfying

1. $q(x) = f(x)\,p(x) + r(x)$
2. $\deg(r(x)) < \deg(p(x))$

Definition 2.5.2.2. An integral domain R is a principal ideal domain (PID) if every ideal in R is principal.

Example 2.5.2.3. Let K be a field, then the polynomial ring $K[x]$ is a PID.

Proof. Let $I \trianglelefteq K[x]$ be an ideal. If $I = \{0\}$, it is clearly principal. So we assume $I \neq \{0\}$. Choose $p \in I - \{0\}$ with minimum degree. I claim that p generated I.

Let $q \in I$. By exercise 2.5.2.1 there exists $f, r \in R$ with

$$q = fp + r$$

and

$$\deg(r) < \deg(p)\,.$$

The first condition implies $r \in I$. Since p was chosen with minimum degree, we conclude $r = 0$. Thus, q is a multiple of p. **q.e.d.**

2.5.3 Noetherian Rings

2.5.4 Unique Factorization Domains

In this section, all rings are commutative. Also, R is a ring.

Definition 2.5.4.1. For $a, b \in R$, we say that a divides b, in shorthand $a|b$, if there is an $c \in R$ with $ac = b$.

Observation 2.5.4.2.

$$\begin{aligned} & a|b \\ \iff \quad & b \in \langle a \rangle \\ \iff \quad & \langle b \rangle \subseteq \langle a \rangle \end{aligned}$$

Observation 2.5.4.3. *For any $a \in R$,*

$$a \in R^* \text{ if and only if } a|1.$$

Also, for any unit α,

$$\alpha|a \text{ and } \alpha a|a$$

These divisors are called trivial divisors of a.

Note: in an integral domain, if $a|b$ and $b|a$, then a and b are trivial divisors of one another.

Definition 2.5.4.4. An element of R is called irreducible if it is neither 0 nor a unit and only has trivial divisors.

An element $p \in R$ is called prime if the principal ideal generated by p is prime.

Observation 2.5.4.5. *An element a is irreducible if and only if the principal ideal $\langle a \rangle$ is a maximal element in the set of all principal ideals of R.*

Observation 2.5.4.6. *For any element $a \in R$, we have*

$$\langle a \rangle = \{b \mid a|b\}.$$

Thus, p is prime if and only if for all $a, b \in R$,

$$p|ab \implies p|a \text{ or } p|b,$$

i.e., in order to be prime: if you divide a product, you have to divide a factor, too.

Observation 2.5.4.7. *In any integral domain, prime elements are irreducible.*

Proof. Let p be prime and suppos $a|p$. We have to show that a is a trivial divisor of p. Since we can write $p = ab$ and p is prime, we find

$$p|a \text{ or } p|b.$$

In the first case, we are done since a and p are mutual divisors, they are mutual trivial divisors.

In the second case, a is a unit. **q.e.d.**

Definition 2.5.4.8. A factorization of a ring element a is a sequence $(\alpha; b_1, b_2, \ldots, b_u)$ with

$$a = \alpha b_1 b_2 \cdots b_u$$

where α is a unit and all b_i are irreducible.

Such a factorization is considered unique if for any other factorization $(\beta; c_1, c_2, \ldots, c_v)$ of a,

1. $u = v$ and

2. there is a permutation $\sigma \in \mathbf{S}_u$ such that for all i there is a unit γ_i satisfying $b_i = \gamma_i c_{\sigma(i)}$.

I.e.: up to reordering, the irreducible factors differ only by units.

An integral domain is called a unique factorization domain (UFD) if every non-zero element has a unique factorization.

Koenig's Lemma 2.5.4.9. *A graph is locally finite if every vertex has only finitely many edges attached to it.*

Every infinite locally finite connected graph contains a repetition-free path of infinite length. In fact, at each vertex you can find such a path issuing from there.

Proof. In a connected graph, we define the distance of two vertices to be the minimum length of an edge path connecting them.

Let v be a vertex. Note that v has only finitely many neighbors (at distance 1). Each of these has only finitely many neighbors, too. Thus, there are only finitely many vertices at distance 2 from v. Inductively, we see that for any specified distance, there are only finitely many vertices that distance away from v. Since the graph is infnite, it follows that there are vertices of arbitrary large distance from v. Since each of these vertices can be connected to v, we infer that there is an infinite set $\mathcal{P}_0$ of paths starting at v such that the lengths of the paths in $\mathcal{P}_0$ are not bounded. **q.e.d.**

Proposition 2.5.4.10 (existence of factorizations). *Suppose that every ascending chain of principal ideals*

$$\langle a_1 \rangle \subseteq \langle a_2 \rangle \subseteq \langle a_3 \rangle \subseteq \cdots$$

stablizes, i.e., there is an index n such that for all $i \geq n$, we have $\langle a_i \rangle = \langle a_{i+1} \rangle$. Then, every element $a \in R$ admits a factorization into irreducible elements.

Note that the second condition holds automatically in any noetherian ring.

Proof. Given an element a **q.e.d.**

Proposition 2.5.4.11 (uniqueness of factorizations). *Suppose R is an integral domain wherein every irreducible element in R is prime. Then, every factorization is unique.*

Proof. Let

$$\alpha a_1 a_2 \cdots a_u = \beta b_1 b_2 \cdots b_v$$

be two factorizations of the same element into irreducible factors. We may assume that this example has minimal total length $u+v$.

Since a_1 is irreducible, it is prime; and since it divides the right hand side, it divides one of the factors there. This factor cannot be the unit. Thus $a_1 | b_i$ for some i. By reordering the right hand side, we may assume $i = 1$. Since b_1 is also irreducible, a_1 is a trivial divisor. Since a_1 is not a unit, we find that $b_1 = \gamma a_1$ for some unit γ. Thus:

$$\alpha a_1 a_2 \cdots a_u = \beta\gamma a_1 b_2 \cdots b_v.$$

Now, we use that R is an integral domain – we may cancel:

$$\alpha a_2 \cdots a_u = \beta\gamma b_2 \cdots b_v.$$

This yields a shorter counter example for a non-unique factorization. **q.e.d.**

Proposition 2.5.4.12. *In a PID, every irreducible is prime.*

Proof. Let a be irreducible. Then $\langle a \rangle$ is a maximal element in the set of all principal ideals. In a PID, this is the set of all ideals. Thus, $\langle a \rangle$ is a maximal ideal. But all maximal ideal are prime. **q.e.d.**

Corollary 2.5.4.13. *Every PID is a UFD.* **q.e.d.**

Proposition 2.5.4.14. *In a UFD, every irreducible element is prime.*

Proof. Let p be irreducible and assume $p|ab$, i.e.,

$$cp = ab.$$

We have to show that $p|a$ or $p|b$.

Fix factorizations into irreducible elements:

$$\begin{aligned} a &= \alpha a_1 \cdots a_{u_a} \\ b &= \beta b_1 \cdots b_{u_b} \\ c &= \gamma a_1 \cdots c_{u_c} \end{aligned}$$

Since p is irreducible, we find that

$$\gamma a_1 \cdots c_{u_c} p = (\alpha\beta) a_1 \cdots a_{u_a} b_1 \cdots b_{u_b}$$

are two factorizations of the same ring element. Since factorizations are unique up to reordering, we infer that p must occur (maybe modified by a unit element) among the irreducible factors on the right hand side. If it is one of the a_i, we have $p|a$; if it is one of the b_i, we have $p|b$. **q.e.d.**

Definition 2.5.4.15. Let R be an integral domain. An element t is called the greates common divisor of the elements a and b if the following two conditions hold:

1. The element t divides both, a and b.
2. Any $s \in R$ dividing a and b also divides t.

Let $A \subseteq R - \{0\}$ be a non-empty subset of non-zero elements. An element t is called a greatest common divisor of A if

1. $A \subseteq \langle t \rangle$
2. For any $s \in R$ with $A \subseteq \langle s \rangle$, we have, $t \in \langle s \rangle$.

Observation 2.5.4.16. *If $\langle t \rangle = \langle a, b \rangle$, then t is a greatest common divisor of a and b. Similarly, if $\langle t \rangle = \langle A \rangle$, then t is the greatest common divisor of A. Warning: the converse implications fail miserably.* **q.e.d.**

Proposition 2.5.4.17. *In a UFD, any non-empty set of non-zero elements has a greatest common divisor. It is to be found as follows: fix a set $\mathcal{P} \subseteq R$ of prime elements that contains exactly one representative from each R^*-orbit of the action of R^* on the set of all prime elements. (Recall that prime elements are essentially unique up to multiplication by units, thus $\mathcal{P}$ contains one prime element from each equivalence class of prime elements.) Then, every $a \in A$ has a unique factorization*

$$a = \alpha_a \prod_{p \in \mathcal{P}} p^{a_p}.$$

For each $p \in \mathcal{P}$, put

$$t_p := \min_{a \in A} a_p.$$

Then

$$t := \prod_{p \in \mathcal{P}} p^{t_p}$$

is a greatest common divisor for A. Note that this product is actually finite since all but finitely many exponents vanish: this is already true for the exponents a_p for any given $a \in A$, and these exponents dominate the exponents $t_p \leq a_p$.

Moreover, A contains a finite subset B for which t is a greatest common divisor.

Proof. It is clear that $t|a$ for each $a \in A$. Thus $A \subset \langle t \rangle$.

Now, let $s = \alpha_s \prod_{p \in \mathcal{P}} p^{s_p}$ be the factorization of another common divisor of A. Then, for each $p \in \mathcal{P}$, we have $s_p \leq t_p$ whence $s|t$.

To argue the moreover part, fix an $a \in A$. Put $\mathcal{Q} := \{q \in \mathcal{P} \mid a_q \neq 0\}$ and note that this is a finite set. For each $q \in \mathcal{Q}$, fix another element $b(q) = \alpha_q \prod_{p \in \mathcal{P}} p^{b_p(q)} \in A$ so that $t_q = b_q(q)$ for each $q \in \mathcal{Q}$. Then, t is the greatest common divisor of the finite set $\{b_q \mid q \in \mathcal{Q}\} \cup \{a\}$. **q.e.d.**

Observation 2.5.4.18. *In any integral domain R, the following are equivalent:*

1. *Any two elements $a, b \in R$ have a common divisor t that is an R-linear combination $t = c_a a + c_b b$.*

2. *Every ideal that is generated by two elements is a principal ideal.*

3. *Every finitely generated ideal is a principal ideal.*

Theorem 2.5.4.19. *Let R be a UFD. Then the following are equivalent:*

1. *Every principal ideal generated by an irreducible element is maximal.*

2. *Every finitely generated ideal is a principal ideal.*

3. *R is a PID.*

Proof. Clearly, (3) implies (2): if every ideal is a principal ideal, then this holds in particular for finitely generated ideals.

Conversely, assume (2). Let I be any non-trivial ideal in R. By (**??**), the set $I - \{0\}$ has a gcd t which is already the gcd of a finite subset $B \subseteq I - \{0\}$. However, then $\langle B \rangle$ is finitely generated and therefore a principal ideal. It follows that

$$\langle t \rangle = \langle B \rangle \subseteq I \subseteq \langle t \rangle$$

which implies that all terms are equal. Thus, I is a principal ideal.

Now, we show that (2) implies (1). Let p be irreducible. In a UFD, this implies that p is prime. We have to show that every coset $a + \langle p \rangle \neq \langle p \rangle$ is a unit in ${}^{R}\!/_{\langle p \rangle}$, i.e., we have to there is a $b \in R$ such that $ab + \langle p \rangle = 1 + \langle p \rangle$. In other words, we have to show that 1 is an R-linear combination of a and p. However, if a is not a multiple of the prime p, the only common divisors of these two elements are units. Since (2) imlpies that we can combine a common divisor, the claim follows. **q.e.d.**

Exercise 2.5.4.20. Complete the proof by supplying a missing implication. Hint: You might find the following lemma useful.

Lemma 2.5.4.21. *Let R be a UFD such that every principal prime ideal is maximal. Then, for any two elements $a, b \in R$, the following are equivalent:*

1. *The elements a and b are relatively prime, i.e., there is no prime element dividing both.*
2. *The unit $1 \in R$ is an R-linear combination of a and b, i.e., there are c_a and c_b in R such that*

$$1 = c_a a + c_b b.$$

3. *The element a projects to a unit in the ring $R/\langle b \rangle$.*
4. *The element b projects to a unit in the ring $R/\langle a \rangle$.*

Proof. First note that (2) is clearly equivalent to (3). Similarly, it is equivalent to (4). In particular, (4)and (3) are equivalent.

Also, since any prime dividing a and b divides any linear combination, we find that (2) implies (1).

Finally, we prove that (1) implies (3). Now, let $p \in R$ be a prime factor of a. Then p does not divide b, i.e., $b \notin \langle p \rangle$. Since $\langle p \rangle$ is maximal, $R/\langle p \rangle$ is a field and b projects to a non-zero element, which is therefore a unit in $R/\langle p \rangle$. Since (4)and (3) are equivalent, we conclude that p projects to a unit in $R/\langle b \rangle$. This applies to all prime factors of a. Since a is the product of its prime factors (with multiplicities) and since the product of units is a unit, a projects to a unit in $R/\langle b \rangle$, as the projection is a ring homomorphism. **q.e.d.**

As opposed to the previous result, the following is somewhat silly:

Theorem 2.5.4.22. *Let R be a noetherian domain. Then the following are equivalent:*

1. *Every principal ideal generated by an irreducible element is maximal.*
2. *Every finitely generated ideal is a principal ideal.*
3. *R is a PID.*

Proof. Note that (1) implies that all irreducibles are prime. Since R is noetherian, R is a UFD and the other two statements follow from (2.5.4.19).

Similarly, (3) implies that R is a UFD. Again the other two statements follow.

Finally, in a noetherian domain, all ideals are finitely generated, whence (2) and (3) are clearly equivalent. **q.e.d.**

Our next goal is to show:

Theorem 2.5.4.23. *If R is a UFD, then so is $R[x]$.*

The main idea is to reduce the question to $K[x]$ where $K \supseteq R$ is the field of fractions of R. Since K is a field, $K[x]$ is a PID and thus a UFD. Thus, we need some lemmas relating irreducibility and factorizations in $K[x]$ and $R[x]$.

Definition 2.5.4.24. A polynomial $p = a_0 + a_1x + a_2x^2 + \cdots + a_mx^m$ is called primitive if units are the only divisors common to all coefficients (or equivalently in a UFD: no prime element divides all the a_i).

Observation 2.5.4.25. *If R is a UFD, any irreducible polynomial in $R[x]$ is primitive.*

Proof. If a polynomial is not primitive, there is a common prime factor for all coefficients. Splitting of that factor is a non-trivial product decomposition. **q.e.d.**

Lemma 2.5.4.26. *If R is a UFD, then any prime $p \in R$ is prime in $R[x]$.*

Proof. Let $p = a_0 + a_1 x + a_2 x^2 + \cdots + a_m x^m$ and $q = b_0 + b_1 x + b_2 x^2 + \cdots + b_n x^n$ be two polynomials. We argue by contradiction and assume that p divides (all coefficients of) pq, but does neither divide p nor q.

Let i be minimal such that p does not divide a_i and let j be minimal such that p does not divide b_j. Since $p | a_0 b_{i+j} + a_1 b_{i+j-1} + \cdots a_i b_j + \cdots + a_{i+j} b_0$ it follows by minimality of i and j that $p | a_i b_j$. Since p is prime, this is a contradiction. **q.e.d.**

Lemma 2.5.4.27 (Gauss). *Let R be a UFD with field of fractions K. A non-constant polynomial $p \in R[x]$ is irreducible in $R[x]$ if and only if it is primitive and irreducible as an element of $K[x]$.*

Proof. First, let us assume that p is irreducible in $R[x]$. We already have seen that p must be primitive. So let us assume that we had a non-trivial factorization

$$p = q' r'$$

in $K[x]$. Note that since all constant polynomials are units in $K[x]$, the polynomials on the right hand side have degrees strictly smaller than p. Multiplying by all denominators of coefficients on the right hand side, we see that we find an element $a \in R$ and polynomials $q, r \in R[x]$ with

$$ap = qr.$$

Let us choose the element a and the polynomials q and r so that the number of prime factors in the unique prime factor decomposition of a is as small as possible. We claim that in this case, a is a unit: This clearly concludes the proof.

So assume that some prime element p divides a. Then p also divides qr. Since p is prime in $R[x]$ we find that p divides on of the polynomials q or r. In this case, we could cancel p on both sides and thus reduce the number of prime factors in a.

Now assume that p is primitive and irreducible in $K[x]$. Let us assume, we had a factorization

$$p = qr$$

in $R[x]$. Since p is irreducible in $K[x]$ one of the polynomials on the right hand side is constant (i.e., a unit in $K[x]$). Since p is also primitive, this constant polynomial must be a unit in R. **q.e.d.**

Proof of Theorem 2.5.4.23. First we argue that every polynomial $p \in R[x]$ has a decomposition into irreducible factors: We decompose p as an element of $K[x]$ where K is the field of fractions for R. After multiplying with a common multiple for all denominators, we find

$$p = \xi q_1 \cdots q_u$$

where $\xi \in K$ and all $q_i \in R[x]$. Moving further prime factors into the field element, we may assume that all q_i are primitive. Then, by the previous lemma, they are irreducible. Moving the denominator of ξ to the left hand side, we obtain:

$$ap = bq_1 \cdots q_u$$

Note that since all q_i are primitive, no prime divisor of a divides into any of the polynomial factors on the right. Thus, all those prime factors are actually in b. It follows that we can cancel a completely. This yields a factorization in $R[x]$.

Now we argue uniqueness. Let

$$p = p_1 \cdots p_u p_1 \cdots p_{u'}$$

and

$$p = q_1 \cdots q_v q_1 \cdots q_{v'}$$

be two decompositions into irreducible elements of $R[x]$ where we set off the constant polynomial factors at the front: those are units

in $K[x]$. The non-constant factors are primitive and irreducible in $K[x]$. Since $K[x]$ is a UFD (it is a PID), we find that $u' = v'$ and that the primitive polynomials coincide up to permutation: note that two primitive polynomials cannot differ by a unit in K since the argument above implies that for

$$ar = a'r'$$

with two primitive polynomials r and r' the ring elements a and a' have the same prime factors (up to units in R).

Since the non-constant polynomials in the decomposition of p coincide, we find that

$$p_1 \cdots p_u$$

and

$$q_1 \cdots q_v$$

differ by a unit in R. Now uniqueness of factorizations in R applies. **q.e.d.**

Theorem 2.5.4.28 (Eisenstein Criterion). *Let R be a UFD and let $p = a_0 + a_1 x + \cdots a_m x^m \in R[x]$ be a primitive polynomial. Assume that there is a prime element $p \in R$ satisfying*

1. *The prime p does not divide the leading coefficient term a_m.*
2. *The prime p divides all other coefficients a_i with $i < m$.*
3. *However, p^2 does not divide the constant term a_0.*

Then p is irreducible in $R[x]$. Note: since p is primitive, it will also be irreducible over the field of fractions of R.

Proof. Let us assume that $p = qr = (b_0 + b_1 x + \cdots b_n x^n)(c_0 + c_1 x + \cdots c_{n'} x^{n'})$. Since $p | a_0 = b_0 c_0$, we may assume that $p | b_0$. Since p^2 does not divide a_0, we infer that p does not divide c_0.

Let i be minimal such that p does not divide b_i. Assume $i < m$, then

$$p | b_0 c_i + \cdots b_i c_0$$

which, by minimality of i, implies $p|b_i c_0$ and therefore $p|b_i$. Hence q and p have the same degree. Since p is primitive, r has to be a unit. Thus p is irreducible. **q.e.d.**

Example 2.5.4.29. The polynomial ring $\mathbb{Z}[x]$ is a noetherian UFD that is not a PID: the elements x and 2 are relatively prime, but it is impossible to write 1 as a linear combination of these two elements.

Example 2.5.4.30. We already saw that the ring of Gaussian integers $\mathbb{Z}[i] = \{a + ib \,|\, a, b \in \mathbb{Z}\}$ is a Euclidean domain. In particular, this ring is a UFD.

Example 2.5.4.31. The ring $\mathbb{Z}\big[\sqrt{-5}\big] = \big\{a + \sqrt{-5}b \,\big|\, a, b \in \mathbb{Z}\big\}$ is not a UFD but every element has a factorization into irreducible elements.

Existence of factorizations follows from considering the norm of elements: there are only finitely many elements with norm below a given bound. Thus, only finitely many ways of splitting off factors arise. Hence we get trapped with a decomposition into irreducible factors.

Non-uniqueness follows from the example

$$6 = 2 \times 3 = \big(1 + \sqrt{-5}\big) \times \big(1 - \sqrt{-5}\big).$$

Chapter 3

Modules and Bi-Modules

3.1 Functors

Let R and S be two fixed rings. Unless otherwise specified, a bimodule is always an R-S-bimodule. Note that a left-R-module is the same as an R-$\mathbb{Z}$-bimodule. Thus, the following notions apply to left and right modules just as well.

3.1.1 Direct Product and Sum

Definition 3.1.1.1. Let I be a set (whose elements we shall be using as indices). Let $(M_i)_{i\in I}$ be a family of bimodules. The direct product is the bimodule

$$\prod_{i\in I} M_i := \left\{(m_i)_{i\in I} \mid m_i \in M_i\right\}$$

where addition and multiplication are defined "slot-wise".

The direct sum is the submodule

$$\bigoplus_{i\in I} M_i := \left\{ (m_i)_{i\in I} \in \prod_{i\in I} M_i \;\middle|\; \text{only finitely many } m_i \neq 0 \right\}$$

Note that for a finite index set I, there is no difference between the direct product and the direct sum.

We have canonical projections

$$\pi_i : \prod_{i\in I} M_i \longrightarrow M_i$$

(these restrict to the direct sum) and canonical injections

$$\iota_i : M_i \longrightarrow \bigoplus_{i\in I} M_i \leq \prod_{i\in I} M_i$$

Functoriality 3.1.1.2. *Let $(\varphi_i : M_i \to N_i)_{i\in I}$ be a family of bimodule homomorphisms. Then there is a unique bimodule homomorphism $\varphi : \prod_{i\in I} M_i \to \prod_{i\in I} M_i$ such that the following diagram commutes for each i:*

$$\begin{array}{ccccc} M_i & \xrightarrow{\iota_i} & \prod_{i\in I} M_i & \xrightarrow{\pi_i} & M_i \\ \downarrow{\scriptstyle \varphi_i} & & \downarrow{\scriptstyle \varphi} & & \downarrow{\scriptstyle \varphi_i} \\ N_j & \xrightarrow{\iota_i} & \prod_{i\in I} N_i & \xrightarrow{\pi_i} & N_i \end{array}$$

Moreover, φ restricts to a bimodule homomorphism from $\bigoplus_{i\in I} M_i$ to $\bigoplus_{i\in I} N_i$. **q.e.d.**

Universal Property of Direct Products 3.1.1.3. *Let M be a bimodule and let $(\varphi_i : M \to M_i)_{i\in I}$ be a family of bimodule homomorphisms. Then there is a unique bimodule homomorphism $\varphi : M \to \prod_{i\in I} M_i$ such that the following diagram commutes for each i:*

$$\begin{array}{ccc} M & \xrightarrow{\varphi} & \prod_{i\in I} M_i \\ & {\scriptstyle \varphi_i}\searrow & \downarrow{\scriptstyle \pi_i} \\ & & M_i \end{array}$$

Proof. Exercise: do uniqueness first. **q.e.d.**

Corollary 3.1.1.4. *Direct products are characterized by the universal property as follows: Define a <u>product</u> of the familiy (M_i) to be a bimodule P together with a familiy of morphisms $(\pi_i : P \to M_i)$. Given two products $(\pi_i : P \to M_i)$ and $(\xi_i : Q \to M_i)$ we say that a*

bimodule homomorphism $\Phi : P \to Q$ is a morphism of products if for each j the diagram

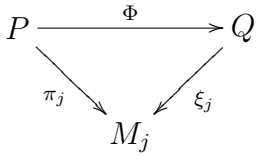

commutes. A product $(\pi_i : P \to M_i)$ is called a direct product if it satisfies the following universal property: whenever $(\pi_i : Q \to M_i)$ is another product, there is a unique morphism of products $P \to Q$.

Direct products are unique up to unique product isomorphism and the direct product $\prod_i M_i$ defined above is a direct product.

Proof. All but the last paragraph is just definitions. The very last sentence just restates the universal property of the direct product. As for the uniqueness claim, we use the usual trick: Let $(\pi_i : P \to M_i)$ and $(\xi_i : Q \to M_i)$ be two direct products. By the universal property, there are unique product morphisms $\Phi : P \to Q$ and $Q : Q \to P$. Note that the following diagrams commute:

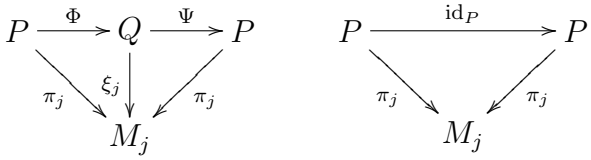

From uniqueness, it follows that $\Psi \circ \Phi = \mathrm{id}_P$. By the same reasoning, $\Phi \circ \Psi = \mathrm{id}_Q$. Hence Φ and Ψ are inverse isomorphisms. **q.e.d.**

Universal Property of Direct Sums 3.1.1.5. *Let M be a bimodule and let $(\varphi_i : M_i \to M)_{i \in I}$ be a family of bimodule homomorphisms. Then there is a unique bimodule homomorphism $\varphi : \bigoplus_{i \in I} M_i \to M$ such that the following diagram commutes for each i:*

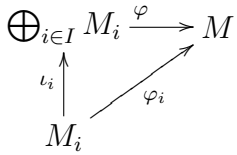

Proof. Exercise: do uniqueness first. **q.e.d.**

Corollary 3.1.1.6. *Direct sums are characterized by the universal property as follows: Define a sum of the familiy (M_i) to be a bimodule P together with a familiy of morphisms $(\iota_i : M_i \to P)$. Given two sums $(\iota_i : M_i \to P)$ and $(\kappa_i : M_i \to Q)$ we say that a bimodule homomorphism $\Phi : P \to Q$ is a morphism of sums if for each j the diagram*

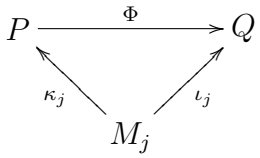

commutes. A sum $(\iota_i : M_i \to P)$ is called a direct sum if it satisfies the following universal property: whenever $(\kappa_i : M_i \to Q)$ is another sum, there is a unique morphism of sums $P \to Q$.

Direct sums are unique up to unique product isomorphism and the direct sum $\bigoplus_i M_i$ defined above is a direct sum.

Proof. Same as above. Details are left to you. **q.e.d.**

Exercise 3.1.1.7. Prove or disprove: Let $(M_i \to N_i \to M_i^*)_{i\in I}$ be a family of sequences exact in the middle. Then the induced sequence

$$\prod_{i\in I} M_i \to \prod_{i\in I} N_i \to \prod_{i\in I} M_i^*$$

is exact in the middle.

Exercise 3.1.1.8. Prove or disprove: Let $(M_i \to N_i \to M_i^*)_{i\in I}$ be a family of sequences in the middle. Then the induced sequence

$$\bigoplus_{i\in I} M_i \to \bigoplus_{i\in I} N_i \to \bigoplus_{i\in I} M_i^*$$

is exact in the middle.

Exercise 3.1.1.9. Let $I = I_1 \cup I_2$ be a decomposition of the index set into two disjoint subsets. Prove or disprove:

1. The sequence

$$\prod_{i\in I_1} M_i \to \prod_{i\in I} M_i \to \prod_{i\in I_2} M_i$$

is a short exact sequence. If so, decide whether there is an obvious splitting.

2. The sequence

$$\bigoplus_{i\in I_1} M_i \to \bigoplus_{i\in I} M_i \to \bigoplus_{i\in I_2} M_i$$

is a short exact sequence. If so, decide whether there is an obvious splitting.

(The maps used in these statements should be obvious.)

3.1.2 Tensor Products

Definition 3.1.2.1. Let X be a set. The free abelian group over X, is the additive group

$$\mathbb{Z}^X := \bigoplus_{i\in X} \mathbb{Z} = \left\{ (m_i)_{i\in X} \,\middle|\, \text{ only finitely many } m_i \neq 0 \right\}.$$

We regard X as a subset of $\mathbb{Z}^X$ by way of characteristic functions: the element $x \in X$ corresponds to the family that is everywhere 0 except at x, where the entry is 1.

Universal Property of Free Abelian Groups. *Let X be a set and let A be an abelian group. For any map $f : X \to A$ there exists a unique homomorphism $\varphi_f : \mathbb{Z}^X \to A$ such that the diagram*

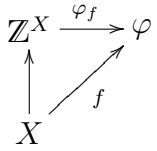

commutes. **q.e.d.**

Proposition 3.1.2.2 (Functoriality). *Let X and Y be sets, and let $f : X \to Y$ be a map. Then, there is a unique homomorphism $f_* : \mathbb{Z}^X \to \mathbb{Z}^Y$ so that the diagram*

$$\begin{array}{ccc} \mathbb{Z}^X & \xrightarrow{f_*} & \mathbb{Z}^Y \\ \uparrow & & \uparrow \\ X & \xrightarrow[f]{} & Y \end{array}$$

commutes. Moreover, the star operation is compatible with composition:

$$(f \circ g)_* = f_* \circ g_*$$

Proof. Straight forward computation: Let the uniqueness claim guide you. **q.e.d.**

Exercise 3.1.2.3. Prove or disprove: $\mathbb{Q}$ is a free abelian group.

Exercise 3.1.2.4. Again the universal property characterizes free abelian groups: Given a set X, let us define an abelian envelope of X to be an abelian group A together with a map $f : X \to A$. Given two abelian envelopes $f : X \to A$ and $g : X \to \psi$, a homomorphism $\alpha : A \to B$ is an X-morphism if the diagram

$$\begin{array}{ccc} A & \xrightarrow{\alpha} & B \\ & {\scriptstyle f}\nwarrow \quad \nearrow{\scriptstyle g} & \\ & X & \end{array}$$

commutes.

An abelian envolpe $f : X \to A$ is called free if it satisfies the following universal property: For any abelian envelope $g : X \to g$, there is a unique X-morphism $\alpha : A \to B$.

Show that that free abelian envelope over X is unique up to unique X-isomorphism and that the free abelian group over X is a free abelian envelope. Consequently, we will drop the terminology "envelope" and we shall just call them free abelian groups.

Exercise 3.1.2.5. Let F_X be the free group over X. Show that its abelianization $\mathrm{F}_X / [\mathrm{F}_X, \mathrm{F}_X]$ is a free abelian group over X.

Definition 3.1.2.6. Let M be an R-S-bimodule, and let N be an S-T-bimodule. Let Q be an R-T-bimodule. A map

$$\beta : M \times N \longrightarrow Q$$

is called R-S-T-bilnear if the following axioms hold:

1. β is left-R-linear in the first argument:

$$\beta\left(\sum_i a_i m_i, n\right) = \sum_i a_i \beta(m_i, n)\,.$$

2. β is right-T-linear in the second argument:

$$\beta\left(m, \sum_j n_j c_j\right) = \sum_j \beta(m, n_j)\, c_j.$$

3. β is balanced:

$$\beta(mb, n) = \beta(m, bn)$$

Remark 3.1.2.7. Note that the first two conditions can be combined into:

$$\beta\left(\sum_i a_i m_i, \sum_j n_j c_j\right) = \sum_{i,j} a_i \beta(m_i, n_j)\, c_j.$$

Definition 3.1.2.8. Let M be an R-S-bimodule, let N be an S-T-bimodule. An R-T-bimodule P together with a R-S-T-bilinear map

$$\mu : M \times N \longrightarrow P$$

is called an S-product. Given two S-products $\mu : M \times N \longrightarrow P$ and $\nu : M \times N \longrightarrow Q$, an R-T-bimodule homomorphism $\varphi : P \to Q$ is an S-product morphism if the diagram

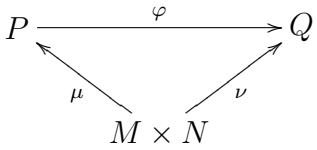

commutes.

An S-product $\mu : M \times N \longrightarrow P$ is called a tensor product of M and N over S if it is universal: for every S-product $\nu : M \times N \longrightarrow Q$, there exists a unique S-product morphism $\Phi : P \to Q$.

Theorem 3.1.2.9. *Tensor products exist and are unique up to unique S-product isomorphism.*

Remark 3.1.2.10. The tensor product is denoted by $M \otimes_S N$ and the structure map is given by

$$\begin{array}{ccc} M \times N & \longrightarrow & M \otimes_S N \\ (m,n) & \mapsto & m \otimes n \end{array}$$

The elements $m \otimes n$ are called elementary tensors.

Proof. Uniqueness is the usual trick and left as an excercise. Existence is what requires proof.

Put:

$$\begin{array}{rcl} A & := & \mathbb{Z}^{M \times N} \\ U & := & \left\langle \left\{ \begin{array}{l|l} (mb,n) - (m,bn) & b \in S \\ (m,n) + (m^*,n) - (m+m^*,n) & m, m^* \in M \\ (m,n) + (m,n^*) - (m,n+n^*) & n, n^* \in N \end{array} \right\} \right\rangle \\ M \otimes_S N & := & A/U \end{array}$$

Let $m \otimes n$ denote the image of the generator (m,n) in $M \otimes_S N$.

For any $a \in R$, let $\lambda_a : A \to A$ be the endomorphism induced by the map $(m,n) \mapsto (am,n)$. We check that $\lambda_a(N) \subseteq N$, which is easy since we just have to see what λ_a does on the generators of U. We find:

$$\begin{array}{rcl} \lambda_a((mb,n) - (m,bn)) & = & (amb,n) - (am,bn) \in U \\ \lambda_a((m,n) + (m^*,n) - (m+m^*,n)) & = & (am,n) + (am^*,n) - (a(m+m^*),n) \\ & = & (am,n) + (am^*,n) - (am+am^*,n) \in U \\ \lambda_a((m,n) + (m,n^*) - (m,n+n^*)) & = & (am,n) + (am,n^*) - (am,n+n^*) \in U \end{array}$$

It follows that λ_a descends to an endormorphism $\lambda_a : M \otimes_S N \to M \otimes_S N$. Analogously, we define $\rho_c : M \otimes_S N \to M \otimes_S N$ using the right multiplication of T on N.

This way, we have equipped $M \otimes_S N$ with an R-T-bimodule structure: the associative laws all can be proved using (3.1.2.2), and the distributive laws follow from the definition of U by computations similar to those above. E.g.:

$$\begin{aligned}(a+b)(m \otimes n) &= ((a+b)m) \otimes n \\ &= (am + bm) \otimes n \\ &= (am) \otimes n + (bm) \otimes n \\ &= a(m \otimes n) + b(m \otimes n)\end{aligned}$$

The next item on the agenda is to check that the canonical map

$$\begin{aligned}M \times N &\longrightarrow M \otimes_S N \\ (m, n) &\mapsto m \otimes n\end{aligned}$$

R-S-T-bilinear. This, again, follows from the defininiton of U by easy computations on elementary tensors.

It remains to check the universal property. Let P be an R-T-bimodule and let $\beta : M \times N \to P$ be bilinear. Since $M \otimes_S N$ is generated (as an abelian group) by elementary tensors, there is at most one homomorphism of abelian groups $\Phi : M \otimes_S N \to P$ that makes the diagram

$$\begin{array}{ccc} M \otimes_S N & \xrightarrow{\Phi} & P \\ \uparrow & \nearrow_{\beta} & \\ M \times N & & \end{array}$$

commute. Thus, it remains to show that the homomorphism

$$\tilde{\Phi} : A \to P$$

that makes

$$\begin{array}{ccc} A & \xrightarrow{\Phi} & P \\ \uparrow & \nearrow_{\beta} & \\ M \times N & & \end{array}$$

commute, descends to an R-T-homomorphism $\Phi : M \otimes_S N \to P$. The proof of this requires two straight forward checks: first, you (not I) have to see that the generators ofg U are in the kernel of $\tilde{\mu}$; finally you (again, not I) have to verify that $\tilde{\Phi}$ is compatible with the R-T-multiplication, which can also be checked on generators. **q.e.d.**

Remark 3.1.2.11. It is aparent from the construction (and has been used in the above proof) that the tensor product $M \otimes_S N$ is generated as an abelian group by the set $\{m \otimes n \,|\, m \in M, n \in N\}$. However, we can also easily deduce this just from the universal property: Put

$$P := \langle \{m \otimes n \,|\, m \in M, n \in N\} \rangle \leq M \otimes_S N$$

and let $\beta : M \times N \to P$ be a bilinear map. Then there exists a homomorphism $M \otimes_S N \to P$ which restricts to a homomorphism $\Phi : P \to P$ so that

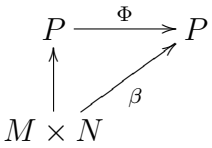

commutes. On the other hand, the homomorphism Φ is determined already by its values on the elementary tensors as they generate P. It follows that P satisfies the universal property of the tensor product. Consequently, the inclusion

$$P \hookrightarrow M \otimes_S N$$

is an isomorphism of tensor products. In particular, it is onto, whence $P = M \otimes_S N$.

Properties

Proposition 3.1.2.12 (Functoriality). *Let $\varphi : M_0 \to M_1$ be an R-S-homomorphism and let $\psi : M_0 \to N_1$ be an S-T-homomorphism. Then there is a unique homomorphism*

$$\varphi \otimes \psi : M_0 \otimes_S N_0 \longrightarrow M_1 \otimes_S N_1$$

that makes the diagram

$$\begin{array}{ccc} M_0 \otimes_S N_0 & \xrightarrow{\varphi\otimes\psi} & M_1 \otimes_S N_1 \\ \uparrow & & \uparrow \\ M_0 \times N_0 & \xrightarrow[\varphi\times\psi]{} & M_0 \times N_0 \end{array}$$

commute.

Morover, tensoring homomorphisms is compatible with composition: Given sequences

$$M_0 \xrightarrow{\varphi_0} M_1 \xrightarrow{\varphi_1} M_2$$

and

$$N_0 \xrightarrow{\psi_0} N_1 \xrightarrow{\psi_1} N_2$$

the diagram

$$\begin{array}{ccc} M_0 \otimes_S N_0 & \xrightarrow{(\varphi_1\circ\varphi_0)\otimes(\psi_1\circ\psi_0)} & M_2 \otimes_S N_2 \\ & {\scriptstyle\varphi_0\otimes\psi_0}\searrow \quad \nearrow{\scriptstyle\varphi_1\otimes\psi_1} & \\ & M_1 \otimes_S N_1 & \end{array}$$

commutes.

Proof. Straight forward consequence of the universal property: the composition

$$M_0 \times N_0 \xrightarrow{\varphi\times\psi} M_1 \times N_1 \longrightarrow M_1 \otimes_S N_1$$

is bilinear. The first claim follows.

The second statement follows from uniqueness and the commutativity of

$$\begin{array}{ccc} M_0 \times N_0 & \xrightarrow{(\varphi_1\circ\varphi_0)\times(\psi_1\circ\psi_0)} & M_2 \times N_2 \\ & {\scriptstyle\varphi_0\times\psi_0}\searrow \quad \nearrow{\scriptstyle\varphi_1\times\psi_1} & \\ & M_1 \times N_1 & \end{array}$$

q.e.d.

Proposition 3.1.2.13 (Distributivity). *Let (M_i) be a family of R-S-bimodules and let N be an S-T-bimodule. Then*

$$\begin{array}{rccc} \Phi : & \left(\bigoplus_i M_i\right) \otimes_S N & \longrightarrow & \bigoplus_i (M_i \otimes_S N) \\ & (m_i)_i \otimes n & \mapsto & (m_i \otimes n)_i \end{array}$$

is a natural isomorphism. That means: given morphisms $\varphi_i : M_i \to M_i^*$ *and* $\psi : N \to N^*$, *the diagram*

$$\begin{array}{ccc} (\bigoplus_i M_i) \otimes_S N & \xrightarrow{(\bigoplus_i \varphi_i)\otimes\psi} & (\bigoplus_i M_i^*) \otimes_S N^* \\ \downarrow & & \downarrow \\ \bigoplus_i (M_i \otimes_S N) & \xrightarrow[\bigoplus_i (\varphi_i \otimes \psi)]{} & \bigoplus_i (M_i^* \otimes_S N^*) \end{array}$$

commutes.

Proof. We shall first specify an inverse homomorphism. Let $\iota_j : M_j \to \bigoplus_i M_i$ be the canonical inclusion. We define

$$\begin{array}{rcl} \Psi : \bigoplus_i (M_i \otimes_S N) & \longrightarrow & \left(\bigoplus_i M_i\right) \otimes_S N \\ (m_i \otimes n_i)_i & \mapsto & \sum_i \iota_i(m_i) \otimes n_i \end{array}$$

First note that the sum on the right hand side is finite since $\iota_i(m_i) \otimes n_i$ vanishes whenever $m_i \otimes n_i$ vanishes. Using the universal property of the direct sum, we can describe Ψ alternatively as the homomorphism that makes

$$\begin{array}{ccc} \bigoplus_i (M_i \otimes_S N) & \xrightarrow{\Psi} & (\bigoplus_i M_i) \otimes_S N \\ \uparrow & \nearrow_{\iota_j \otimes \mathrm{id}_N} & \\ M_j \otimes_S N & & \end{array}$$

commute. Thus, the map before us is a well-defined homomorphism.

It is easy to check that $\Phi \circ \Psi$ takes generators back to themselves, and so does $\Psi \circ \Phi$. Naturality can also be checked on generators and is straight forward.

Alternatively, you could stare at the diagram

$$\begin{array}{ccccc} \bigoplus_i (M_i \otimes_S N) & \xrightarrow{\Psi} & (\bigoplus_i M_i) \otimes_S N & \xrightarrow{\Phi} & \bigoplus_i (M_i \otimes_S N) \\ & \nwarrow & \uparrow{\scriptstyle \iota_j \otimes \mathrm{id}_N} & \nearrow & \\ & & M_j \otimes_S N & & \end{array}$$

and deduce from universal properties (uniqueness strikes again) that $\Phi \circ \Psi$ is the identity.

Similarly, we get that $\Psi \circ \Phi$ is the identity from the diagram

$$\begin{array}{ccccc} (\bigoplus_i M_i) \otimes_S N & \xrightarrow{\Phi} & \bigoplus_i (M_i \otimes_S N) & \xrightarrow{\Psi} & \bigoplus_i (M_i \otimes_S N) \\ & \nwarrow \quad \nearrow & & & \\ & (\bigoplus_i M_i) \times N & & & \end{array}$$

q.e.d.

Exercise 3.1.2.14. Disprove that tensor products distribute over direct products (as opposed to direct sums).

Exercise 3.1.2.15 (Associativity). Let R_i be rings $(i \in \{1,2,3,4\})$ and let M_i be R_i-R_{i+1}-bimodules $(i \in \{1,2,3\})$. Show that there is a unique R_1-R_4-isomorphism

$$\Phi : (M_1 \otimes_{R_2} M_2) \otimes_{R_3} M_3 \longrightarrow M_1 \otimes_{R_2} (M_2 \otimes_{R_3} M_3)$$

that makes

$$\begin{array}{ccc} (M_1 \otimes_{R_2} M_2) \otimes_{R_3} M_3 & \xrightarrow{\Phi} & M_1 \otimes_{R_2} (M_2 \otimes_{R_3} M_3) \\ \uparrow & & \uparrow \\ (M_1 \times M_2) \times M_3 & \xrightarrow[\text{obvious map}]{} & M_1 \times (M_2 \times M_3) \end{array}$$

commute.

Moreover show that this morphism is natural, i.e., given bimodule homomorphisms $\varphi_i : M_i \to N_i$ then the diagram

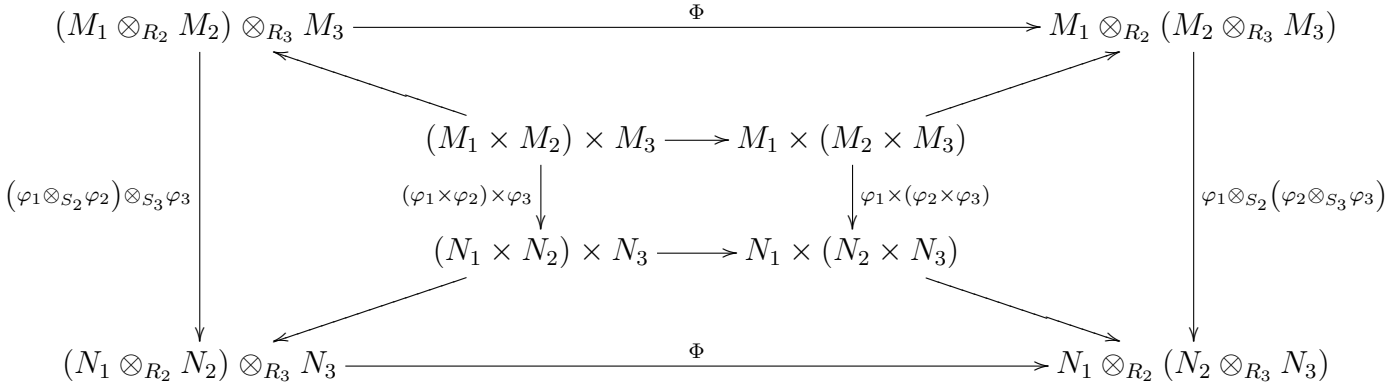

commutes.

Motivated by the above, define a triple tensor product by means of a universal property (involving tri-linear maps) and show that $(M_1 \otimes_{R_2} M_2) \otimes_{R_3} M_3$ and $M_1 \otimes_{R_2} (M_2 \otimes_{R_3} M_3)$ qualify as realizations of the triple tensor product. (Hint: one way to go elegantly about this problem is to do the last part first.)

Proposition 3.1.2.16 (Right Exactness). *The functor $- \otimes_S N$ is right-exact: Let*

$$0 \to M' \xrightarrow{\iota} M \xrightarrow{\pi} M^* \to 0$$

be a short exact sequence of R-S-bimodules, and let N be an S-T-bimodule. Then the induced sequence

$$M' \otimes_S N \xrightarrow{\iota \otimes \mathrm{id}_N} M \otimes_S N \xrightarrow{\pi \otimes \mathrm{id}_N} M^* \otimes_S N \to 0$$

is exact.

In the same way, the functor $M \otimes_S -$ is right-exact. (This statement is left as an exercise.)

Proof. First note that $\pi \otimes \mathrm{id}_N$ is onto, since elementary tensors generate $M^* \otimes_S N$. Moreover, it follows from naturality that

$$(\pi \otimes \mathrm{id}_N) \circ (\iota \otimes \mathrm{id}_N) = (\pi \circ \iota) \otimes \mathrm{id}_N = 0 \otimes \mathrm{id}_N = 0$$

whence $\mathrm{im}(\iota \otimes \mathrm{id}_N) \subseteq \ker(\pi \otimes \mathrm{id}_N)$. It remains to show that $\ker(\pi \otimes \mathrm{id}_N) \subseteq \mathrm{im}(\iota \otimes \mathrm{id}_N)$.

Put

$$\mathcal{I} := \mathrm{im}(\iota \otimes \mathrm{id}_N) \leq M \otimes_S N$$

and

$$\mathcal{CK} := {}^{(M \otimes_S N)}\Big/_{\mathcal{I}} \ .$$

Since $\mathrm{im}(\iota \otimes \mathrm{id}_N) \subseteq \ker(\pi \otimes \mathrm{id}_N)$, the homomorphism $\pi \otimes \mathrm{id}_N$ descends to a homomorphism $\pi_* : \mathcal{CK} \to M^* \otimes_S N$. We finish the proof by constructing a section $\sigma_* : M^* \otimes_S N \to \mathcal{CK}$, which demonstrates that π_* is injective and hence an isomorphism.

For starters define

$$\begin{aligned} \sigma : M^* \times N &\longrightarrow \mathcal{CK} \\ (\pi(m), n) &\mapsto m \otimes n + \mathcal{I} \end{aligned}$$

and note that this definition is well-put: if $\pi(m_1) \otimes n = \pi(m_2) \otimes n$ then $m_1 \otimes n - m_2 \otimes n \in \mathcal{I}$. It is easy to check that σ is bilinear. Hence it induces a homomorphism $\sigma_* : M^* \otimes_S N \to \mathcal{CK}$ so that

$$\begin{array}{ccc} M^* \otimes_S N & \xrightarrow{\sigma_*} & \mathcal{CK} \\ \uparrow & \nearrow & \\ M^* \times N & & \end{array}$$

commutes. By construction, $\sigma_* \circ \pi_* = \mathrm{id}_{\mathcal{CK}}$. You can check this easily on generators. **q.e.d.**

Free Modules

Definition 3.1.2.17. Let R and S be rings and let X be a set. A free R-S-bimodule with basis X is an R-S-bimodule F together with a map $\iota : X \to F$ satisfying the following universal propert:

> For any bimodule M and any map $f : X \to M$, there exists a unique R-S-homomorphism $\varphi : F \to M$ so that the diagram
>
> $$\begin{array}{ccc} F & \xrightarrow{\varphi} & M \\ {\scriptstyle \iota}\uparrow & \nearrow_{f} & \\ X & & \end{array}$$
>
> commutes.

The cardinality of X is called the rank of the free module.

Warning 3.1.2.18. *The rank of a free module is in general not a well-defined notion, i.e., it is possible that free modules over sets of different cardinality are isomorphic as R-S-bimodules. We shall see, however, that in important cases the structure of a free module determines the cardinality of its basis.*

Example 3.1.2.19 (Free Bimodules of Rank 1). Let R and S be two rings. Then R is naturally an R-$\mathbb{Z}$-bimodule and S is naturally a $\mathbb{Z}$-S-bimodule. We shall write out the universal property for the R-S-bimodule

$$F := R \otimes_{\mathbb{Z}} S.$$

We get: for any R-S-bimodule M and any bilinear map $\beta : R \times S \to M$ there exists a unique R-S-homomorphism $\varphi : F \to M$ so that

$$\begin{array}{ccc} F & \xrightarrow{\varphi} & M \\ \uparrow & \nearrow_{\beta} & \\ R \times S & & \end{array}$$

commutes. Note, however, that the bilinear map is uniquely determined by the value $\beta(1,1)$. Conversely, any value maybe specified for this pair and we can always extend it to a bilinear map. Thus, we may equivalently say: Let X be a one point set and let $\iota : X \to F$ send its element to $1 \otimes 1$. Then, for any map $f : X \to M$ there exists a unique homomorphism $\varphi : F \to M$ so that

$$\begin{array}{ccc} F & \xrightarrow{\varphi} & M \\ {\scriptstyle \iota}\uparrow & \nearrow_{f} & \\ X & & \end{array}$$

commutes.

Exercise 3.1.2.20 (Free Modules of rank 1). Let R be a ring.

1. Prove or disprove: R is a free R-R-bimodule of rank 1.

2. Prove or disprove: R is a free R-$\mathbb{Z}$-bimodule of rank 1.

Exercise 3.1.2.21. Let $X = X_1 \cup X_2$ be a disjoint union of sets. Let $X_1 \to F_1$ and $X_2 \to F_2$ be free R-S-bimodules over X_1 and X_2, respectively. Show that

$$X \to F_1 \oplus F_2$$

is a free R-S-bimodule over X.

Corollary 3.1.2.22. *In particular, we note that* $R^X := \bigoplus_{x \in X} (R \otimes_{\mathbb{Z}} S)$ *is a free* R*-*S*-bimodule over* X*.*

Definition 3.1.2.23. Let R and S be rings and let X be a set. A free left-R-module with basis X is an left-R-module F together with a map $\iota : X \to F$ satisfying the following universal propert:

> For any left-R-module M and any map $f : X \to M$, there exists a unique R-homomorphism $\varphi : F \to M$ so that the diagram
>
> 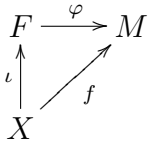
>
>
> commutes.

The cardinality of X is called the rank of the free module.

Exercise 3.1.2.24. Show that every free left-R-module is a free R-$\mathbb{Z}$-bimodule and that every free R-$\mathbb{Z}$-bimodule is a free left-R-module.

Modules over Commutative Rings

Let us assume that R is commutative. We can define an R-R-bimodule structure on any left-R-module M by $ma := am$. This works as $a(mb) = a(bm) = b(am) = (am)b$. Let us call a R-R-bimodule M a commutative R-module if $am = ma$ for all $a \in R$ and all $m \in M$. Thus, we have seen that we can define a commutative R-module structure on any left-R-module.

Conversely, given a commutative R-module M, we can forget the right multiplication and regard it as a left-R-module. However, note that no information is lost: we can reconstruct the right-multiplication from the left-multiplication. (This is why you thaught your Linear Algebra teacher was overly picky when he took points off for multplying from the wrong side, insisting that there is no rule for how to do that.) In short:

Observation 3.1.2.25. *The construction*

$$\{\textit{left-}R\textit{-modules}\} \longrightarrow \{\textit{commutative } R\textit{-modules}\}$$

and

$$\{\textit{commutative } R\textit{-modules}\} \longrightarrow \{\textit{left-}R\textit{-modules}\}$$

described above are mutually inverse. They preseve commutative diagrams. **q.e.d.**

Exercise 3.1.2.26. Give an example of a R-R-bimodule that is not commutative.

Exercise 3.1.2.27. Define a free commutative R-module by means of a universal property. Show that the correspondence from (3.1.2.25) identifies free left-R-modules and free commutative R-modules.

Observation 3.1.2.28. *It is much easier to construct free commutative R-modules than it is to construct free bimodules: for any set X, the module $R^X := \bigoplus_{x\in X} R$ with the obvious map $x \mapsto \chi_x$ is a free commutative R-module.* **q.e.d.**

Exercise 3.1.2.29. Let K be a field. Show that every commutative K-module is a free commutative K-module: Let F be a commutative K-module. Then there exists a set X and a map $\iota : X \to F$ such that for any commutative K-module M and any map $f : X \to M$, there is a unique homomorphism $\varphi : F \to M$ so that

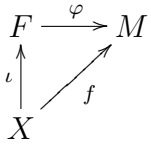

commutes.

Observation 3.1.2.30. *Tensor products of free commutative R-modules*

are easy:

$$\begin{aligned} R^X \otimes_R R^Y &= \left(\bigoplus_{x\in X} R\right) \otimes_R \left(\bigoplus_{y\in Y} R\right) \\ &= \bigoplus_{x\in X} \left(R \otimes_R \left(\bigoplus_{y\in Y} R\right)\right) \\ &= \bigoplus_{x\in X} \left(\bigoplus_{y\in Y} R \otimes_R R\right) \\ &= R^{X\times Y} \end{aligned}$$

In particular, the tensor product is again a free commutative R-module.

Corollary 3.1.2.31 (tensor product of vector spaces). *Dimensions multiply when you tensor vector spaces over the same field.* **q.e.d.**

Example 3.1.2.32 (vector spaces of finite dimension). Note that the universal property gives a 1-1 correspondence

$$\mathrm{Bil}_K(V \times W; K) \longleftrightarrow \mathrm{Lin}_K(V \otimes_K W; K)$$

This correspondence is easily seen to be K-linear and natural. Thus, there is a natural isomorphism of the dual of $V \otimes_K W$ and the K-values bilinear forms on $V \times W$:

$$(V \otimes_K W)^* = \mathrm{Bil}_K(V \times W; K)$$

For vector spaces of finite dimension, we can do better because finite dimensional vector spaces are canonically isomorphic to their double-dual spaces. We obtain the interesting natural isomorphsm

$$V \otimes_K W = (V \otimes_K W)^{**} = \mathrm{Bil}_K(V \times W; K)^*$$

for vector spaces of finite dimension.

Proposition 3.1.2.33. *Let M and N be two commutative R-modules. Then $M \otimes_R N$ is a commutative R-module. (Note that M and N are both R-R-bimodules. Thus, $M \otimes_R N$ is a R-R-bimodule.)*

Proof. For elementary tensors, we have

$$a(m \otimes n) = (am) \otimes n = (ma) \otimes n = m \otimes (an) = m \otimes (na) = (m \otimes n)a$$

and we can extend additively to all of $M \otimes_R N$. **q.e.d.**

Proposition 3.1.2.34 (Commutativity of Tensor Products). *Let M and N be two commutative R-modules. Then*

$$\begin{array}{ccc} M \otimes_R N & \longrightarrow & N \otimes_R M \\ m \otimes n & \mapsto & n \otimes m \end{array}$$

defines a natural isomorphism.

Proof. Note that

$$\begin{array}{ccc} M \times N & \longrightarrow & N \otimes_R M \\ (m, n) & \mapsto & n \otimes m \end{array}$$

is R-R-R-bilinear. Thus, there is a unique R-R-homomorphism

$$M \otimes_R N \longrightarrow N \otimes_R M$$

that makes

$$\begin{array}{ccc} M \otimes_R N & \longrightarrow & N \otimes_R M \\ \uparrow & \nearrow & \uparrow \\ M \times N & \longrightarrow & N \times M \end{array}$$

commute.

From uniqueness, we also get that the obvious homomorphism in the other direction has to be the inverse.

Naturallity is easily checked on elementary tensors. Let $\varphi : M_1 \to M_2$ and $\psi : N_1 \to N_2$ be two homomorphisms. Then, for $m \in M_1$ and $n \in N_1$, we have

$$\begin{array}{ccc} m \otimes n & \longrightarrow & n \otimes m \\ {\scriptstyle \varphi \otimes \psi} \downarrow & & \downarrow {\scriptstyle \psi \otimes \varphi} \\ \varphi(m) \otimes \psi(n) & \longrightarrow & \psi(n) \otimes \varphi(m) \end{array}$$

q.e.d.

Extensions

Example 3.1.2.35 (Scalar Extension). Let R and T be rings, and let M be an R-T-bimodule.

An R-algebra is ring S together with a ring homomorphism $\sigma : R \to S$. Note that S is an S-R-bimodule in a natural way. Then $S \otimes_R M$ is an S-T-bimodule, called the scalar extension of M over S.

The universal property specializes to: For any S-T-bimodule N and any R-$\mathbb{Z}$-T-bilinear map $\beta : S \times M \to N$ there exists a unique S-R-T-homomorphism $\varphi_* : S \otimes_R M \to N$ so that

$$\begin{array}{ccc} S \otimes_R M & \xrightarrow{\varphi_*} & N \\ \uparrow & \nearrow_{\beta} & \\ S \times M & & \end{array}$$

commutes. Now observe that S-R-T-bilinear maps from $S \times M$ to N determined by what they do on $1 \times M$. Moreover, there is a 1-1-correspondence from those S-R-T-bilinear maps to R-T-homomorphisms from M to N regarded as an R-Q via the structure homomorphism $\sigma : R \to S$.

Thus, we obtain the following characterization of the scalar extension by a universal property:

> For any S-T-bimodule N and any R-T-homomorphism $\varphi : M \to N$ there exists a unique S-T-homomorphism $\varphi_* : S \otimes_R M \to N$ so that
>
> $$\begin{array}{ccc} S \otimes_R M & \xrightarrow{\varphi_*} & N \\ {\scriptstyle m \mapsto 1 \otimes m}\uparrow & \nearrow_{\varphi} & \\ M & & \end{array}$$
>
> commutes.

Exercise 3.1.2.36. Let K be a field and let R be an integral domain. Recall that $R[x]$ denotes the ring of polynomials in the variable x whereas $R[[x]]$ denotes the ring of power series in the

variable x. Both are integral domains whose fields of fractions are denoted by $R(x)$ and $R((x))$, respectively. We also fix the multiplicative set $L := \{1, x, x^2, \ldots\}$. Prove or disprove:

1. (a) $K((x)) \cong K(x) \otimes_{K[x]} K[[x]]$
   (b) $K((x)) \cong K(x) \otimes_K K[[x]]$
   (c) $K[x] \otimes_K K[y] = K[x, y]$
2. (a) $R((x)) \cong R(x) \otimes_{R[x]} R[[x]]$
   (b) $L^{-1}R[[x]] \cong R(x) \otimes_{R[x]} R[[x]]$
   (c) $L^{-1}R[[x]] \cong L^{-1}R[x] \otimes_{R[x]} R[[x]]$
   (d) $R((x)) \cong R(x) \otimes_R R[[x]]$
   (e) $L^{-1}R[[x]] \cong R(x) \otimes_R R[[x]]$
   (f) $L^{-1}R[[x]] \cong L^{-1}R[x] \otimes_R R[[x]]$
   (g) $R[x] \otimes_R R[y] = R[x, y]$

Exercise 3.1.2.37 (Localization). Let R be a commutative ring and $L \subset R$ be a multiplicative set. Let M be a left R-module. Show that

$$\begin{array}{ccc} L^{-1}R \otimes_R M & \longrightarrow & L^{-1}M \\ \frac{a}{r} \otimes m & \mapsto & \frac{am}{r} \end{array}$$

is an isomorphism of $L^{-1}R$-left modules. Also show that this isomorphism is natural in the module M.

Corollary 3.1.2.38. *In particular, localization distributes over direct sums and is right-exact.* **q.e.d.**

Corollary 3.1.2.39. *Let X be a set, let R be an integral domain and let K be its field of fractions. Then we have a natural isomorphism $K^X \cong K \otimes_R R^X$. Since K^X is a vector space, dimension theory from linear algebra tells us that the cardinality of X is determined by the vector space. Thus, free modules over integral domains have well-defined ranks.*

Exercise 3.1.2.40. Show that

$$\mathbb{C} \otimes_{\mathbb{R}} \mathbb{M}_{n,m}(\mathbb{R}) \cong \mathbb{M}_{n,m}(\mathbb{C})$$

Exercise 3.1.2.41. Let R and S be rings. Prove or disprove the following statements:

1. Every R-S-bimodule can be regarded as an $R \times S^{\mathrm{op}}$-left module.
2. Every $R \times S^{\mathrm{op}}$-left module can be regarded naturally as an R-S-bimodule.
3. Every R-S-bimodule can be regarded as an $R \otimes_{\mathbb{Z}} S^{\mathrm{op}}$-left module.
4. Every $R \otimes_{\mathbb{Z}} S^{\mathrm{op}}$-left module can be regarded naturally as an R-S-bimodule.

3.1.3 Algebras

Definition 3.1.3.1. An R-algebra is ring S together with a ring homomorphism $\sigma : R \to S$. The homomorphism σ is called the structure homomorphism.

An R-algebra $\sigma : R \to S$ is called central if the image of σ is contained in the center of S. Note that in this case, S is a commutative R-R-bimodule.

Proposition 3.1.3.2 (Tensor Products of Algebras). *Let R be a ring and let S and T be central R-algebras. Show that the S-T-bimodule $S \otimes_R T$ carries a unique ring structure satisfying*

$$(b_1 \otimes c_1)(b_2 \otimes c_2) = (b_1 b_2) \otimes (c_1 c_2)$$

Moreover, show that

$$a \mapsto a \otimes 1 = 1 \otimes a$$

defines a ring homomorophism turning $S \otimes_R T$ into a central R-algebra.

Proof. First of all let us see why the hypothesis of centrality is forced upon us. We observe

$$(b \otimes 1)(1 \otimes c) = b \otimes c = (1 \otimes c)(b \otimes 1)$$

which implies that S and T commute inside $S \otimes_R T$. Since the image of σ sits inside S it commutes with all of T and since it sits inside T it commutes with all of S.

The key problem in this statement is to see that the definition of the multiplication is well-put. If it is, then it is obvious that we defined a ring structure. So, how could the definition fail? Well, for $a_i \in R$, $b_i \in S$, and $c_i \in T$, we find right away: $b_i a_i \otimes c_i = b_i \otimes a_i b_i$ whence, for instance,
$b_1 a_1 b_2 a_2 \otimes c_1 c_2 = (b_1 a_1 \otimes c_1)(b_2 a_2 \otimes c_2) = (b_1 \otimes a_1 c_1)(b_2 \otimes a_2 c_2) = b_1 b_2 \otimes a_1 c_1 a_2 c_2$
which just so happens to be fine, since we assumed centrality. We could check all relations from the construction of the tensor product and verify that the product rule from above extends to a well defined multiplication.

A more structural approach is this. First observe that multiplication $S \times S \to S$ is bilinear and therefore induces a homomorphism $S \otimes_R S \to S$ that extends multiplication from elementary tensors. Similarly, we get for T a homomorphism $T \otimes_R T \to T$. By naturality, we can put these together and get:

$$S \otimes_R S \otimes_R T \otimes_R T \longrightarrow S \otimes_R T.$$

Here, we use associativity and commutativity (the algebras are central!) of the tensor product, and obtain:

$$(S \otimes_R T) \otimes_R (S \otimes_R T) \longrightarrow S \otimes_R T$$

Finally, we write out the commutative diagram

$$\begin{array}{ccc} (S \otimes_R T) \otimes_R (S \otimes_R T) & \longrightarrow & S \otimes_R T \\ \uparrow & \nearrow & \\ (S \otimes_R T) \times (S \otimes_R T) & & \end{array}$$

and observe that the diagonal arrow is exactly the map we wanted to extend. **q.e.d.**

Observation 3.1.3.3. *Let R be a ring and let S be an R-R-bimodule together with*

- *an R-R-homomorphism $\sigma : R \to S$ and*
- *a R-R-homomorphism map $\mu : S \otimes_R S \to S$*

Then, μ defines a multiplication low on S via

$$S \times S \to S \otimes_R S \to S$$

so that $\sigma : R \to S$ is an R-algebra provided the following diagrams commute:

associativity

$$\begin{array}{ccc} S \otimes_R S \otimes_R S & \xrightarrow[\mu\otimes \mathrm{id}]{} & S \otimes_R S \\ \downarrow{\scriptstyle \mathrm{id}\otimes\mu} & & \downarrow{\scriptstyle \mu} \\ S \otimes_R S & \xrightarrow[\mu]{} & S \end{array}$$

structure homomorphism

$$\begin{array}{ccc} R \otimes_R S & \xrightarrow{\cong} & S \\ {\scriptstyle \sigma\otimes\mathrm{id}}\downarrow & \nearrow_{\mu} & \\ S \otimes_R S & & \end{array} \qquad \begin{array}{ccc} S \otimes_R R & \xrightarrow{\cong} & S \\ {\scriptstyle \mathrm{id}\otimes\sigma}\downarrow & \nearrow_{\mu} & \\ S \otimes_R S & & \end{array}$$

This says that $\sigma(1)$ is a multiplicative identity element in S.

$$\begin{array}{ccc} R \otimes_R R & \xrightarrow{\cong} & R \\ \downarrow{\scriptstyle \sigma\otimes\sigma} & & \downarrow{\scriptstyle \sigma} \\ S \otimes_R S & \xrightarrow[\mu]{} & S \end{array}$$

This says that σ is multiplicative.

One nice thing is that distributivity just flows from bilinearity of the multiplication law and does not need a diagram.

Definition 3.1.3.4. Let $\sigma_1 : R \to S_1$ and $\sigma_2 : R \to S_2$ be two R-algebras. A ring homomorphism $\varphi : \sigma_1 \to \sigma_2$ is an R-algebra homomorphism if

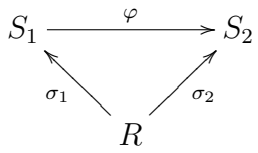

commutes.

The Tensor Algebra

Definition 3.1.3.5 (The Tensor Algebra). Let R be a ring and let M be an R-R-bimodule. We construct the tensor algebra of M over R as follows: Put

$$M^{\otimes m} := \underbrace{M \otimes_R M \otimes_R \cdots \otimes_R M}_{m \text{ factors}}$$

and define a ring structure on

$$\mathcal{T} := \mathcal{T}_R(M) := R \oplus M^{\otimes 1} \oplus M^{\otimes 2} \oplus M^{\otimes 3} \oplus \cdots$$

induced by

$$\begin{array}{ccc} M^{\otimes m} \times M^{\otimes n} & \longrightarrow & M^{\otimes (m+n)} \\ (s,t) & \mapsto & s \otimes t \end{array}$$

The inclusion of R as the degree 0 summand in $\mathcal{T}$ endows this ring with the structure of an R-algebra.

Note that $\mathcal{T}_R(M)$ comes with a direct sum decomposition $\mathcal{T}_R(M) = \bigoplus_{i=0}^{\infty} \mathcal{T}_R^i(M)$. The elements concentrated in one summand are called homogeneous and the index of their summand is called the degree of the element. Note that degrees add up when homogeneous elements are multiplied. This makes $\mathcal{T}_R(M)$ a graded algebra (with an $\mathbb{N}$-grading; other monoids could occur as well, the next important being $\mathbb{Z}_2$: there are many $\mathbb{Z}_2$-graded algebras in physics).

Observation 3.1.3.6. *If M is a commutative R-module over a commutative ring R, the tensor algebra is a central R-algebra.* **q.e.d.**

Exercise 3.1.3.7. Show that polynomial rings arise as tensor algebras: $\mathcal{T}_R(R) \cong R[x]$.

Example 3.1.3.8. Let K be a field, and let V be a K-vector space with basis $X \subset V$. Then $K[X^*] \cong \mathcal{T}_R(V)$, where X^* is the free monoid over X. (Moreover, this morphism is natural in the category of vector spaces with distinguished bases.)

The isomorphism is induced by

$$x_1 x_2 \cdots x_u \mapsto x_1 \otimes x_2 \otimes \cdots \otimes x_u.$$

Proposition 3.1.3.9 (Universal Property of the Tensor Algebra). *Let R be a ring and let M be an R-R-bimodule. For any R-algebra S and any R-R-bimodule homomorphism $\varphi : M \to S$ there exists a unique R-algebra homomorphism φ_* that makes*

$$\begin{array}{ccc} \mathcal{T}_R(M) & \xrightarrow{\varphi_*} & S \\ \uparrow & \nearrow_{\varphi} & \\ M & & \end{array}$$

commute.

Proof. First, we observe that φ_* is nailed on elementary tensors. We have to put: $\varphi_*(m_1 \otimes \cdots \otimes m_u) = \varphi(m_1) \cdots \varphi(m_u)$. Moreover, if there is an R-R-homomorphism satisfying this rule, it will clearly be multiplicative by the way multiplication in $\mathcal{T}_R(M)$ is defined.

As for existence of such an R-R-homomorphism, note that

$$\begin{array}{ccc} M \times \cdots \times M & \longrightarrow & S \\ (m_1, \ldots, m_u) & \mapsto & \varphi(m_1) \cdots \varphi(m_u) \end{array}$$

is multilinear and therefore induces a bimodule homomorphism

$$M^{\otimes u} \to S$$

The direct sum of all these homomorpihsms defines an R-R-homomorphism $\varphi_* : \mathcal{T}_R(M) \to S$ which clearly does what we need on elementary tensors. **q.e.d.**

The Symmetric Algebra

Definition 3.1.3.10. Let R be a ring and let M be a R-R-bimodule. The symmetric algebra over M is the quotient

$$S_R(M) := \mathcal{T}_R(M)\,/I$$

where I is the two-sided ideal generated by all elements of the form $m \otimes n - n \otimes m \in M^{\otimes 2} \leq \mathcal{T}_R(M)$. We denote the multiplication in the symmetric algebra by $\odot$.

Observation 3.1.3.11. *The elementary elements of the form $m_1 \odot \cdots \odot m_u$ with the empty product interpreted as 1 generate the symmetric algebra as a module over R.* **q.e.d.**

Observation 3.1.3.12. *The Ideal I is generated by homogeneous elements. It follows that the symmetric algebra inherits a grading via*

$$S_R^m(M) := \mathrm{im}\big(M^{\otimes m}\big)\,.$$

These R-R-submodules are called symmetric powers of M. The claim here is that $S_R^m(M)$ and $S_R^n(M)$ intersect trivially for $m \neq n$ so that

$$S_R(M) = \bigoplus_m S_R^m(M)$$

is an $\mathbb{N}$-graded algebra.

Proof. Just observe that $I = \bigoplus_m I^m$ where $I^m = I \cap M^{\otimes m}$. This is to say that the ideal I is closed with respect to taking homomgeneous components. **q.e.d.**

Remark 3.1.3.13. Symmetric powers are functorial in M.

Exercise 3.1.3.14. Prove or disprove: $S_R(M)$ is commutative if and only if R is commutative.

Exercise 3.1.3.15. Let K be a field and let V be a vector space of dimension m. Show that $S_K(V) \cong K[x_1, \ldots, x_m]$ but that the above isomorphism is neither unique nor natural (in fact, it depends on the choice of a basis for V).

The Exterior Algebra

Definition 3.1.3.16. Let R be a ring and let M be a R-R-bimodule. The exterior algebra over M is the quotient

$$\Lambda_R(M) := \mathcal{T}_R(M)\,/I$$

where I is the two-sided ideal generated by all elements of the form $m \otimes n + n \otimes m \in M^{\otimes 2} \leq \mathcal{T}_R(M)$. We denote the exterior product by $\wedge$.

Observation 3.1.3.17. *The elementary elements of the form $m_1 \wedge \cdots \wedge m_u$ with the empty product interpreted as 1 generate the exterior algebra as a module over R.* **q.e.d.**

Observation 3.1.3.18. *The Ideal I is generated by homogeneous elements. It follows that the exterior algebra inherits a grading via*

$$\Lambda_R^m(M) := \mathrm{im}\big(M^{\otimes m}\big)\,.$$

These R-R-submodules are called exterior powers of M. The claim here is that $\Lambda_R^m(M)$ and $\Lambda_R^n(M)$ intersect trivially for $m \neq n$. **q.e.d.**

Remark 3.1.3.19. Exterior powers are functorial in M.

Proposition 3.1.3.20. *Let K be a field of characteristic $\neq 2$ and left V be a vector space of dimension $m \leq \infty$ over K.*

Then for any sequence $\mathbf{v}_1, \ldots, \mathbf{v}_u \in V$, we have the equivalence:

$$\mathbf{v}_1 \wedge \cdots \wedge \mathbf{v}_u = 0 \qquad \textit{iff} \qquad \{\mathbf{v}_1, \cdots, \mathbf{v}_u\} \ \textit{is linearly dependent}$$

Fix a basis $\mathbf{e}_1, \ldots, \mathbf{e}_m$ and put $I := \{1, \ldots, m\}$. For each subset $A = \{i_1 < i_2 < \cdots < i_u\} \subseteq I$, put

$$\mathbf{e}_A := \mathbf{e}_{i_1} \wedge \cdots \wedge \mathbf{e}_{i_u}.$$

The set of all $\mathbf{e}_A$ forms a basis of $\Lambda_K(V)$.

Proof. First we note that swapping the order of two factors within an elementary element $\mathbf{v}_1 \wedge \cdots \wedge \mathbf{v}_u$ amounts to flipping the sign. Thus even permutations of the factors leave the sign unchanged whereas odd permutations of the factors introduce a minus sign. It follows that if the same factor is repeated, the product evaluates to 0: interchanging the two equal factors leaves the produc unchanged and introduces a sign change.

Suppose now, the vectors $\mathbf{v}_i$ are linearly dependent, say

$$\mathbf{v}_1 = a_2\mathbf{v}_2 + \cdots + a_u\mathbf{v}_u$$

Then:

$$\mathbf{v}_1 \wedge \cdots \wedge \mathbf{v}_u = (a_2\mathbf{v}_2 + \cdots + a_u\mathbf{v}_u) \wedge \mathbf{v}_2 \wedge \cdots \wedge \mathbf{v}_u = \sum_{i=2}^{u} a_i\mathbf{v}_i \wedge \mathbf{v}_2 \wedge \cdots \wedge \mathbf{v}_u = 0$$

Since any vector is a linear combination of the basis vectors $\mathbf{e}_i$, any elementary element of the exterior algebra is a linear combination of elementary elements whose factors are basis vectors. The above shows that repeated basis vectors may be dropped and that (up to signs) we can insist on putting those factors in ascending order of indices. That shows that the proclaimed basis vectors $\mathbf{e}_A$ span $\Lambda_K(V)$.

To see linear independence, one defines an algebra structure on the abstract vector space with basis $\{\mathbf{e}_A \mid A \subseteq \{1, \ldots, m\}\}$ and shows that the obvious map is an isomorphism of algebras (I am too lazy to think about the details right now). **q.e.d.**

Exercise 3.1.3.21. Finish the prove of (3.1.3.20):

1. Show that the basis vectors $\mathbf{e}_A$ are linearly independent.
2. Show: if $\mathbf{v}_1, \ldots, \mathbf{v}_u$ are linearly independent, then $\mathbf{v}_1 \wedge \cdots \wedge \mathbf{v}_u \neq 0$.

Remark 3.1.3.22. It follows that $\Lambda^m_K(V)$ is a K-vector of dimension 1. Since exterior powers are functorial, this implies that any K-endomorphism $\eta : V \to V$ induces a K-endomorphism $\eta_* : \Lambda^m_K(V) \to \Lambda^m_K(V)$

which of course, is just multiplication by a certain scalar. Surprise: this scalar is the determinant of η. (Just in case you ever wondered what the real reason is that determinants are multiplicative.)

3.1.4 Appendix: Categories and Functors

Definition 3.1.4.1. A category is a class $\mathcal{C}$ (whose elements are called objects) together with a family of sets $(\mathrm{Mor}(A;B))_{A,B\in\mathcal{C}}$ (whose elements are called morphisms from A to B, i.e., A is the source object and B is the target object of these morphisms) and a rule of composition

$$\begin{aligned}\mathrm{Mor}(A;B)\times\mathrm{Mor}(B;C) &\longrightarrow \mathrm{Mor}(A;C)\\ (f,g) &\mapsto g\circ f\end{aligned}$$

so that composition is associative. Also, we require that $\mathrm{Mor}(A;A)$ always contains a multiplicatively neutral element id_A.

Definition 3.1.4.2. Let $\mathcal{C}$ be a category. A diagram over $\mathcal{C}$ is a directed graph whose vertices are labelled with objects from $\mathcal{C}$ and whose arrows are labeled with morphisms from $\mathcal{C}$ (if A is the label of the initial vertex of an arrow and B is the label of its terminal vertex, then we want the arrow to be labelled by a morphism in $\mathrm{Mor}(A;B)$, of course).

A diagram commutes if for any pair of vertices (source and sink) the composition of the arrows along a directed edge path from the source to the sink does not depend on the particular path chosen.

Definition 3.1.4.3. Let $\mathcal{C}$ and $\mathcal{D}$ be categories. A functor $\mathbf{F}:\mathcal{C}\to\mathcal{D}$ assigns to each object A and morphism f in $\mathcal{C}$ and object $\mathbf{F}(A)$ and a morphism $\mathbf{F}(f)$ in $\mathcal{D}$ so that commutative diagrams over $\mathcal{C}$ are taken to commutative diagrams over $\mathcal{D}$.

Exercise 3.1.4.4. Let X be a topological space. For any two points $x, y \in X$ let $\mathrm{Paths}(x, y)$ be the set of homotopy classes of paths from x to y. Let

$$\mathrm{Paths}(x, y) \times \mathrm{Paths}(y, z) \longrightarrow \mathrm{Paths}(x, z)$$

be the obvious concatenation operation. Show that these data define a category. Show that any continuous function $X \to Y$ induces a functor.

Exercise 3.1.4.5. Let G be a group. Let $\mathcal{C}_G$ be a category with one object A and one morphism from A to A for each group element. Composition of morphisms shall obey the group law in G.

1. Show that any group homomorphism $\varphi : G \to H$ induces a functor from $\mathcal{C}_G$ to $\mathcal{C}_H$.

2. Deduce that

$$G \mapsto \mathcal{C}_G$$

defines a functor from the category of groups to the category of categories (where categories are objects and where functors are morphisms).

Limits and Colimits

Definition 3.1.4.6 (Cones and Limits). Let $\mathcal{C}$ be a category and let D be a commutative diagram in $\mathcal{C}$. A <u>cone</u> over D is an object C together with a family of morphisms $(g_v : C \to A_v)_{v \in D}$ where v runs through the vertices of D. A morphism between cones $C_1 \to C_2$ is a $\mathcal{C}$-morphism that makes the resulting diagram with base D commute.

A cone L over D is a <u>limit</u> if it is universal among all cones, i.e., for every cone C there exists a unique cone-morphism $C \to L$.

Remark 3.1.4.7. As usual, limits are unique up to unique cone-isomorphism. However, existence of limits is usually not clear; and there are categories where limits do not necessarily exist.

Example 3.1.4.8. In the category of R-S-bimodules, let D be a just a set of vertices without any arrows. Then the limit over D is just the direct product of its labels.

The dual notion is obtained by reversing arrows:

Definition 3.1.4.9 (Cocones and Colimits). Let $\mathcal{C}$ be a category and let D be a commutative diagram in $\mathcal{C}$. A cocone over D is an object C together with a family of morphisms $(g_v : A_v \to C)_{v \in D}$ where v runs through the vertices of D. A morphism between cocones $C_1 \to C_2$ is a $\mathcal{C}$-morphism that makes the resulting diagram with base D commute.

A cocone L over D is a colimit if it is universal among all cocones, i.e., for every cone C there exists a unique cocone-morphism $L \to C$.

Remark 3.1.4.10. As usual, colimits are unique up to unique cocone-isomorphism. However, existence of colimits is usually not clear; and there are categories where colimits do not necessarily exist.

Example 3.1.4.11. In the category of R-S-bimodules, let D be a just a set of vertices without any arrows. Then the colimit over D is just the direct sum of its labels.

Definition 3.1.4.12. A functor that commutes with limits is called continuous. It is called co-continuous if it commutes with colimits.

Exercise 3.1.4.13. Prove or disprove: $- \otimes_S N$ is a co-continuous functor.

Definition 3.1.4.14 (Diagram Categories). Let $\mathcal{C}$ be a category and let Γ be a directed graph. Define a category $\mathcal{C}_\Gamma$ of diagrams in $\mathcal{C}$ over Γ whose objects are commutative diagrams over $\mathcal{C}$ whose underlying directed graph is Γ. A morphism between two commutative diagrams is a family $(f : A_v \to B_v)_{v \in \Gamma}$ (here v runs through all vertices of Γ) so that the resulting prism is commutative.

Exercise 3.1.4.15. Limits are functorial: Given a diagram morphism $\delta : D_1 \to D_2$ and given two limits $L_1 := \varinjlim D_1$ and $L_2 := \varinjlim D_2$, there exists a unique $\mathcal{C}$-morphism $f : L_1 \to L_1$ that makes the diagram

$$\begin{array}{ccc} L_1 & \xrightarrow{f} & L_2 \\ \uparrow & & \uparrow \\ D_1 & \xrightarrow[\delta]{} & D_2 \end{array}$$

commute.

3.1.5 Appendix: Homotopy

This is a little bit of topology.

Definition 3.1.5.1. Let X and Y be topological spaces. Two continuous maps $f_0, f_1 : X \to Y$ are called homotopic if there is a continuous map $H : X \times [0,1] \to Y$ so that $f_0 = H(-,0)$ and $f_1 = H(-,1)$. One puts $f_t := H(-,t)$ and thinks of this as a continuous family of functions interpolating from f_0 to f_1.

Let $A \subset X$ be a subset and assume that f_0 and f_1 agree on A. We say that f_0 and f_1 are homotopic relative to A if there is a homotopy from f_0 to f_1 such that all intermediate f_t agree with f_0 and f_1 on the subset A.

Definition 3.1.5.2. A path in X is a continuous map from the unit interval $[0,1]$ into X. Two paths are homotopic relative endpoints it they are homotopic as maps relative to the subset $\{0,1\} \subset [0,1]$.

Exercise 3.1.5.3. Show that paths that only differ in parametrization are homotopic: Let $p : [0,1] \to X$ be a path and let $g : [0,1] \to [0,1]$ be a continuous map with $g(0) = 0$ and $g(1) = 1$. Show that p is homotopic to $p \circ g$ relative to $\{0,1\}$.

Exercise 3.1.5.4. Show that homotopy of maps from X to Y relative to a fixed subset of X is an equivalence relation.

Definition 3.1.5.5. For any path $p : [0,1] \to X$ the reversed path $-p$ is given by $-p(t) := p(1-t)$. For two paths $p, q : [0,1] \to X$, with $q(1) = p(0)$ the concatenation is defined as

$$pq : t \mapsto \begin{cases} q(2t) & t \leq \frac{1}{2} \\ p(2t-1) & \frac{1}{2} \leq t \end{cases}$$

Exercise 3.1.5.6. Show that concatenating a path and its inverse yields a path that is homotopic to a constant path.

Exercise 3.1.5.7. Show that homotopy and concatenation interact nicely: if p and p' are homotopic relative endpoints and q and q' are homotopic relative endpoints, then so are the concatenations pq and $p'q'$.

3.1.6 Appendix: Dual Vector Spaces

In this seciton, K is a field.

Definition 3.1.6.1. Let V be a K-vector space. The dual of V is the set of all K-valued linear forms on V:

$$V^* := \mathrm{Lin}_K(V; K)$$

This set is a K-vector space in the obvious way.

Proposition 3.1.6.2 (Functoriality). *The dual construction is a cofunctor, i.e., for any linear map $\varphi : V \to W$ we define:*

$$\begin{array}{rccc} \varphi^* : W^* & \longrightarrow & V^* \\ \lambda & \mapsto & \lambda \circ \varphi \end{array}$$

This construction takes commutative diagrams to commutative diagrams (reversing arrows), i.e.:

$$(\varphi \circ \psi)^* = \psi^* \circ \varphi^*$$

Proof. Straight forward computation. **q.e.d.**

Proposition 3.1.6.3. *The map*

$$\begin{aligned} \Phi_V : V &\longrightarrow V^{**} \\ \mathbf{v} &\mapsto \mathrm{ev}_{\mathbf{v}} := (\lambda \mapsto \lambda(\mathbf{v})) \end{aligned}$$

is a natural monomorphism. In particular, it is a natural isomorphism for finite dimensional vector spaces.

Proof. First you (not I) check that $\mathrm{ev}_{\mathbf{v}} : V^* \to K$ is linear. Then, another straight forward computation shows Φ_V to be a linear map. Moreover, it is easy to see that 0 is the only vector on which all linear forms vanish. Injectivity follows.

As for naturality, let $\varphi : V \to W$ be linear. We have to show that

$$\begin{array}{ccc} V & \xrightarrow{\Phi_V} & V^{**} \\ {\scriptstyle\varphi}\downarrow & & \downarrow{\scriptstyle\varphi^{**}} \\ W & \xrightarrow{\Phi_W} & W^{**} \end{array}$$

commutes. Thus, we have to show that $\Phi_W(\varphi(\mathbf{v})) = \mathrm{ev}_{\varphi(\mathbf{v})}$ is the same linear map on W^* as $\varphi^{**}(\Phi_V(\mathbf{v})) = \varphi^{**}(\mathrm{ev}_{\mathbf{v}}) = \mathrm{ev}_{\mathbf{v}} \circ \varphi^*$. To that these function are the same, we evaluate on a random linear form $\lambda \in W^*$. We get:

$$\mathrm{ev}_{\varphi(\mathbf{v})}(\lambda) = \lambda(\varphi(\mathbf{v}))$$

and

$$\begin{aligned} (\mathrm{ev}_{\mathbf{v}} \circ \varphi^*)(\lambda) &= \mathrm{ev}_{\mathbf{v}}(\varphi^*(\lambda)) \\ &= \mathrm{ev}_{\mathbf{v}}(\lambda \circ \varphi) \\ &= \lambda(\varphi(\mathbf{v})) \end{aligned}$$

This counts as success. **q.e.d.**

Exercise 3.1.6.4. Make sense of the following statement and prove it: The only natural homomorphism from V to its dual V^* is the trivial 0-homomorphism. (Beware of contravariance!)

3.2 Modules over Group Rings (aka Representation Theory)

3.2.1 Representations as Modules

Let G be a group, let K be a field, and let V be a K-vector space.

Definition 3.2.1.1. A representation of G in V is a group homomorphism

$$\begin{array}{ccc} \rho : G & \longrightarrow & \mathrm{Aut}_K(V) \\ g & \mapsto & \rho_g \end{array}$$

The dimension of V is called the degree of the representation.

A subspace $U \leq V$ is ρ-invariant if for every $g \in G$, we have $\rho_g(U) \subseteq U$. It follows that restriction induces a representation of G on U.

A representation without invariant subspaces is called irreducible.

Let $\rho : G \to \mathrm{Aut}_K(V)$ and $\sigma : G \to \mathrm{Aut}_K(W)$ be two representations. A linear map $\varphi : V \to W$ is equivariant with respect to ρ and σ if for any $g \in G$, the diagram

$$\begin{array}{ccc} V & \xrightarrow{\varphi} & W \\ {\scriptstyle \rho_g}\uparrow & & \uparrow{\scriptstyle \sigma_g} \\ V & \xrightarrow[\varphi]{} & W \end{array}$$

commutes.

Proposition 3.2.1.2. *Let $\rho : G \to V$ be a representation. Then*

$$\begin{array}{ccc} K[G] \times V & \longrightarrow & V \\ \left(\displaystyle\sum_{g \in G} \xi_g g, \mathbf{v} \right) & \mapsto & \displaystyle\sum_{g \in G} \xi_g \rho_g(\mathbf{v}) \end{array}$$

endows V with the structure of a left-$K[G]$-module.

Under this construction, irreducible representations correspond to simple left-$K[G]$-modules; and equivariant linear maps turn into $K[G]$-homomorphisms.

Conversely, any left-$K[G]$-module M is already a K-vector space and left-multiplication $g \mapsto \lambda_g$ is a representation of G on M.

The two constructions are inverse functors and realize an equivalence of categories:

$$\{\textit{representations of } G \textit{ over } K\textit{-vector spaces}\} \longleftrightarrow \{\textit{left-}K[G]\textit{-modules}\}$$

Proof. ... **q.e.d.**

Example 3.2.1.3. Representations of the infinite cyclic group $\mathbb{Z}$ are exactly the modules over the Laruent polynomial ring $K[x, x^{-1}]$. Recall that $K[x, x^{-1}]$ is a PID. We will study finitely generated modules over PIDs extensively in the next chapter.

Exercise 3.2.1.4. Let G be a finite group acting on $\mathbb{R}^m$ by linear automorphisms. Show that

$$\begin{aligned} \langle -, - \rangle : \mathbb{R}^m \times \mathbb{R}^m &\longrightarrow \mathbb{R} \\ (\mathbf{v}, \mathbf{w}) &\mapsto \sum_{g \in G} (g\mathbf{v}) \cdot (g\mathbf{w}) \end{aligned}$$

defines an inner product on $\mathbb{R}^m$.

Also show that this inner product is G-invariant, i.e., for any $g \in G$, we have

$$\langle g\mathbf{v}, g\mathbf{w} \rangle = \langle \mathbf{v}, \mathbf{w} \rangle.$$

Deduce that every G-invariant subspace of $\mathbb{R}^m$ has a G-invariant complementary direct summand.

3.2.2 Constructions

Functoriality of various constructions gives a way to create new group representations from old ones.

Direct Sum

Observation and Definition 3.2.2.1. *Let* $\rho_i : G \to \mathrm{Aut}_K(V_i)$ *be a family* $(i \in I)$ *of representations. Then*

$$\begin{array}{rcl} \rho : G & \longrightarrow & \mathrm{Aut}_K\left(\bigoplus_{i\in I} V_i\right) \\ g & \mapsto & \bigoplus_{i\in I} \rho_i(g) \end{array}$$

is a representation, called the direct sum of the ρ_i *and denoted by* $\bigoplus_{i\in I} \rho_i$.

Definition 3.2.2.2. A representation is indecomposable if it does not split non-trivially as a direct sum of two subrepresentations.

Observation 3.2.2.3. *Irreducible representations are indecomposable.* **q.e.d.**

Remark 3.2.2.4. It is generally not true that indecomposable representations are irreducible.

Exercise 3.2.2.5. Construct an indecomposable, non-irreducible representation of the infinite cyclic group.

Proposition 3.2.2.6. *Let* V *be a complex vector space of finite dimension and let* $\rho : G \to \mathrm{Aut}_{\mathbb{C}}(V)$ *be a a representation of the finite group* G. *Then* ρ *decomposes as a direct sum of irreducibles. In particular, any finite dimensional indecomposable complex representation of a finite group is irreducible.*

Proof. It suffices to show that any proper subrepresentation has a complementary summand. This, however, follows from (3.2.1.4) **q.e.d.**

Definition 3.2.2.7. A representation that decomposes as a sum of irreducibles is called completely reducible. The corresponding notions for modules is semi-simplicity: A module is called semi-simple if it is a direct sum of simple modules.

Exercise 3.2.2.8. Let M be a semi-simple left-R-module. Let $S \leq M$ be a submodule.

1. Show S is a direct summand, i.e., show that there exists a complementary direct summand inside M.

2. Show that S and M/S are semi-simple.

Exercise 3.2.2.9. Show that a left-R-module M is semi-simple if and only if it is the sum (not direct!) of its minimal submodules.

Tensor Product

Observation and Definition 3.2.2.10. *Let $\rho : G \to \mathrm{Aut}_K(V)$ and $\sigma : G \to \mathrm{Aut}_K(W)$ be two representations. The tensor product of ρ and σ is the representations*

$$\begin{array}{rcl} \rho \otimes \sigma : G & \longrightarrow & \mathrm{Aut}_K(V \otimes_K W) \\ g & \mapsto & \rho(g) \otimes \sigma(g) \end{array}$$

Exercise 3.2.2.11. Let $\rho_i : G \to \mathrm{Aut}_K(V_i)$ with $i = 1, 2$ and $\sigma : G \to \mathrm{Aut}_K(W)$ be representations. Show that $(\rho_1 \oplus \rho_2) \otimes \sigma$ and $(\rho_1 \otimes \sigma) \oplus (\rho_2 \otimes \sigma)$ are equivalent representations.

Exercise 3.2.2.12. Let

$$\rho : G \longrightarrow \mathrm{GL}_r(K)$$

and

$$\sigma : G \longrightarrow \mathrm{GL}_s(K)$$

be two representations. Describe their tensor product and their direct sum as a homomorphisms

$$G \to \mathrm{GL}_{rs}(K) \qquad \text{and} \qquad G \to \mathrm{GL}_{r+s}(K)$$

The Dual Representation

Observation and Definition 3.2.2.13. *Let $\rho : G \to \mathrm{Aut}_K(V)$ be a representation. The dual representation is given via:*

$$\begin{aligned} \rho^* : G &\longrightarrow \mathrm{Aut}_K(V^*) \\ g &\mapsto \rho\big(g^{-1}\big)^* \end{aligned}$$

Remark 3.2.2.14. Taking inverses in the above formula is braught upon us since taking dual is a contravariant functor. Naturally it turns a left-action into a right-action, and we need to compensate for that by composing with an anti-automorphism.

3.2.3 Example: The Regular Representation

Definition 3.2.3.1. Let G be a group and let K be a field. The left-multiplication action of G on itself induces (by functoriality of free constructions) an action of G on the K-vector space with basis G by linear maps, i.e., we obtain a representation $\rho : G \to \mathrm{Aut}_K\big(K^G\big)$. This representation is called the regular representation of G over K.

Remark 3.2.3.2. From the module point of view, this representation is particularly simple: it is just the group ring $K[G]$ regarded as a left-$K[G]$-module.

We shall work out what this means for finite groups. In this case, $K_G = \mathrm{Maps}\, G; K$ and the action is given by

$$\begin{aligned} G \times \mathrm{Maps}\, G; K &\longrightarrow \mathrm{Maps}\, G; K \\ (g, f) &\mapsto f \circ \lambda_{g^{-1}} \end{aligned}$$

To prove this, check on characteristic functions for group elements: they form the canonical basis for $\mathrm{Maps}\, G; K$.

Proposition 3.2.3.3. *Let G be a finite group. Any irreducible representation of G arises as a subrepresentation of the regular*

representation. (I.e., every simple $K[G]$-module injects into $K[G]$). More precisely, let $G : G \to \mathrm{Aut}_K(V)$ be a representation. For any linear form $\lambda : V \to K$ the map

$$\begin{array}{rcl} \varphi : V & \longrightarrow & \mathrm{Maps}\, G; K \\ \mathbf{v} & \mapsto & \left(f_{\mathbf{v}} : g \mapsto \lambda\left(g^{-1}\mathbf{v}\right)\right) \end{array}$$

is a G-equivariant linear map, i.e., a homomorphism of representations.

Note that the kernel of φ consists of exactly those $\mathbf{v} \in V$ for which λ vanishes on the orbit $G\mathbf{v}$. In particular, φ is injective provided that ρ is irreducible and λ is non-trivial.

Proof. First, check that φ is K-linear. This is easy.

Then check that φ is G-equivariant. This is easy, too:

$$\begin{array}{rcl} (hf_{\mathbf{v}})(g) & = & (f_{\mathbf{v}} \circ \lambda_{h^{-1}})(g) \\ & = & \lambda\Big(\left(h^{-1}g\right)^{-1}\mathbf{v}\Big) \\ & = & \lambda\left(g^{-1}h\mathbf{v}\right) \\ & = & f_{h\mathbf{v}}(g) \end{array}$$

The other statements are clear. **q.e.d.**

Corollary 3.2.3.4. *Let G be a finite group. Let*

$$\mathbb{C}^G = W_1 \oplus \cdots \oplus W_u$$

*be a decomposition of the regular representation into irreducible representations (**??**). Then, any irreducible representation $\rho : G \to \mathrm{Aut}_{\mathbb{C}}(V)$ is equivalent to one of the W_i.*

Proof. This is most easily seen from the modules point of view. In (**??**), we constructed an injective homomorphism

$$\varphi : V \to \mathbb{C}^G.$$

Let $\pi_i \mathbb{C}^G \to W_i$ be the projection to the ith coordinate. Then the composition $\pi_i \circ \varphi$ is a morphism from a simple module to a simple

module which, therefore, is either trivial or an isomorphism. Since not all of the compositions can be trivial (this would imply that φ is trivial), we find that one of those compositions is an isomorphism. **q.e.d.**

Example 3.2.3.5. $\mathbb{Z}_2$ has exactly two irreducible representatitions: the trivial representation and the flip representations.

Question: Let $M = \bigoplus_i M_i$ be a finite decomposition of a $\mathbb{C}[\mathbb{Z}_2]$-module (of finite complex dimension) into irreducibles. Are the number of trivial and flip summands independent of the chosen decomposition?

Answer: yes, you can use the traces of the two group elements to work out the numbers.

3.2.4 Characters

In this section, we restrict ourselves to finite dimensional representations over the field $\mathbb{C}$ of complex numbers. I.e., a representation of G is given to us as a group homomorphism

$$\rho : G \longrightarrow \mathrm{GL}_r(\mathbb{C})$$

and r is the degree of ρ. The <u>character</u> associated to ρ is the map

$$\begin{array}{rccc} \chi_\rho : G & \longrightarrow & \mathbb{C} \\ g & \mapsto & \mathrm{tr}(\rho_g) \end{array}$$

A function arising this way from a representation is called a <u>character</u>.

Observation 3.2.4.1. *The sum of two characters is a character (arising from the direct sum of their underlying representations).*

Definition 3.2.4.2. A character is <u>irreducible</u> if it cannot be written as the sum of two characters.

Exercise 3.2.4.3. The product of two characters is a character (arising from the tensor product of their underlying representations).

Corollary 3.2.4.4. *The set of all characters on G forms a semi-ring.*

Observation 3.2.4.5. *The following follow from properties of the trace. Let $\rho : G \to \mathrm{GL}_r(\mathbb{C})$ be a representation with associated character χ. Then the following hold:*

1. *Characters are constant on conjugacy classes, i.e.:*

$$\chi(g) = \mathrm{tr}(\rho_g) = \mathrm{tr}(\rho_h \rho_g \rho_{h^{-1}}) = \chi\big(hgh^{-1}\big)$$

2. *We have:*

$$\chi(g) = \overline{\chi(g^{-1})}$$

   To see this, note that g has finite order whence the eigenvalues of ρ_g are roots of unity. Thus complex conjugation is taking reciprocals. This way we get the eigenvalues for $\rho_{g^{-1}}$.

Observation 3.2.4.6. *We can recover from a character the degree of the underlying representation: Let $\rho : G \longrightarrow \mathrm{GL}_r(\mathbb{C})$ be a representation. Then $\chi_\rho(1) = m$. In particular $\chi(1)$ is always a positive integer.*

Corollary 3.2.4.7. *Every character is a sum of irreducible characters.* **q.e.d.**

Schur's Lemma 3.2.4.8. *Let V and W be two simple left-$\mathbb{C}[G]$-modules and let*

$$\varphi : V \longrightarrow W$$

be a module homomorphism.

1. *If V and W are not isomorphic, then φ is trivial*

2. *If V and W are isomorphic, then φ is a homothety, i.e., there are $\mathbb{C}$-bases for V and W relative to which φ is described as a multiple of the identity matrix.*

In particular, $\mathrm{End}_{\mathbb{C}[G]}(V)$ *is isomorphic, as a ring, to* $\mathbb{C}$.

Proof. The first statement is obvious: homomorphisms between simple modules are either isomorphisms or trivial.

As for the second claim, we may assume without loss of generality that $V = W$ and that φ is an endomorphism. Since $\mathbb{C}$ is algebraically closed, φ has an eigenvalue λ. Then

$$\lambda \,\mathrm{id} - \varphi \in \mathrm{End}_{\mathbb{C}[G]}(V)$$

is an endomorphism with non-trivial kernel. Hence, it is 0. **q.e.d.**

Observation 3.2.4.9 (Averaging Linear Maps). *Let* $\rho : G \to \mathrm{Aut}_{\mathbb{C}}(V)$ *and* $\sigma : G \to \mathrm{Aut}_{\mathbb{C}}(W)$ *be representations. Let* $\varphi : V \to W$ *be complex linear. Then*

$$\mu_\varphi := \frac{1}{\mathrm{card}(G)} \sum_{g \in G} \sigma_{g^{-1}} \circ \varphi \circ \rho_g$$

is a G-equivariant complex linear map, i.e., a morphism of representations. Hence Schur's Lemma applies and tells us:

1. *If ρ and σ are inequivalent, then μ_φ vanishes.*
2. *If $V = W$ and $\rho = \sigma$, then μ_φ is a homothety of ratio $\frac{\mathrm{tr}(\varphi)}{\dim_{\mathbb{C}}(V)}$.* **q.e.d.**

Corollary 3.2.4.10. *Let* $R : G \to \mathrm{GL}_r(\mathbb{C})$ *and* $S : G \to \mathrm{GL}_s(\mathbb{C})$ *be two inequivalent irrerducible representations. Then for any four indices, i, i', j, j', we have:*

$$\sum_{g \in G} R_{i,i'}\big(g^{-1}\big)\, S_{j,j'}(g) = 0$$

and

$$\sum_{g \in G} \mathrm{tr}\big(R\big(g^{-1}\big)\big)\, \mathrm{tr}(S(g)) = 0$$

Proof. For any $r \times s$-matrix A, we have

$$\sum_g R\big(g^{-1}\big)\, AS(g) = 0$$

for which we can work out the (i, j')-entry:

$$0 = \sum_{g \in G} \sum_{k',k} R_{i,k'}\big(g^{-1}\big)\, A_{k',k} S_{k,j'}(g)$$

For the matrix A that is everywhere 0 except for a 1 in the (i', j)-entry, we have:

$$0 = \sum_{g \in G} R_{i,i'}\big(g^{-1}\big)\, S_{j,j'}(g)$$

This proves the first claim. The second follows easily:

$$\begin{aligned}
\sum_{g \in G} \mathrm{tr}\big(R\big(g^{-1}\big)\big)\, \mathrm{tr}(S(g)) &= \sum_{g \in G} \left(\sum_i R_{i,i}\big(g^{-1}\big) \right) \left(\sum_j S_{j,j}(g) \right) \\
&= \sum_{i,j} \sum_{g \in G} R_{i,i}\big(g^{-1}\big)\, S_{j,j}(g) \\
&= 0
\end{aligned}$$

q.e.d.

Corollary 3.2.4.11. *Let $R : G \to \mathrm{GL}_r(\mathbb{C})$ be an irreducible representation. Then, for any four indices $i, i' j, j'$, we have*

Proof. For any square matrix A we have

$$\sum_{k,k'} \sum_{g \in G} R_{i,k'}\big(g^{-1}\big)\, A_{k',k} R_{k,j'}(g) = \frac{\mathrm{tr}(A)}{r} \delta_i^{j'}$$

Thus:

$$\sum_{g \in G} R_{i,i'}\big(g^{-1}\big)\, R_{j,j'}(g) = \frac{1}{r} \delta_i^{j'}\, \delta_{i'}^{j}$$

Moreover:

$$\begin{aligned}
\sum_{g \in G} \mathrm{tr}\big(R\big(g^{-1}\big)\big)\, \mathrm{tr}(R(g)) &= \sum_{g \in G} \sum_{i,j} R_{i,i}\big(g^{-1}\big)\, R_{j,j}(g) \\
&= \sum_{g \in G} \sum_i \frac{1}{r} \\
&= \mathrm{card}(G)
\end{aligned}$$

q.e.d.

Let us turn to equivalent representations:

Corollary 3.2.4.12. *Let* $R : G \to \mathrm{GL}_r(\mathbb{C})$ *and* $S : G \to \mathrm{GL}_r(\mathbb{C})$ *be two equivalent irrerducible representations. Then for any four indices,* i, i', j, j'*, we have:*

$$\sum_{g \in G} R_{i,i'}\big(g^{-1}\big)\, S_{j,j'}(g) = 0$$

and

$$\sum_{g \in G} \mathrm{tr}\big(R\big(g^{-1}\big)\big)\, \mathrm{tr}(S(g)) = 0$$

For the remaining discussion, we fix the following hermitean inner product:

$$\begin{array}{rcl} \langle -, - \rangle : \mathrm{Maps}(G; \mathbb{C}) \times \mathrm{Maps}(G; \mathbb{C}) & \longrightarrow & \mathbb{C} \\ (f, g) & \mapsto & \dfrac{1}{\mathrm{card}(G)} \displaystyle\sum_{g \in G} f(g)\, \overline{g(g)} \end{array}$$

Using this inner product, we can rephrase part of the previous corollary as follows:

Corollary 3.2.4.13. *Characters associated to inequivalent irreducible representations are orthogonal: Let* $\rho : G \to \mathrm{Aut}_{\mathbb{C}}(V)$ *and* $\sigma : G \to \mathrm{Aut}_{\mathbb{C}}(W)$ *be two irreducible representations. Then*

$$\langle \chi_\rho, \chi_\sigma \rangle = \begin{cases} 0 & \rho \text{ and } \sigma \text{ are inequivalent} \\ 1 & \rho \text{ and } \sigma \text{ are equivalent} \end{cases}$$

Proof. This is just a reformulation of the previous results. Just recall that $\mathrm{tr}(\rho_{g^{-1}}) = \overline{\mathrm{tr}(\rho_g)}$. **q.e.d.**

Observation 3.2.4.14. *Let*

$$\rho = \bigoplus_i \sigma_i^{k_i}$$

be a direct sum decomposition of a representation $\rho : G \to \mathrm{Aut}_{\mathbb{C}}(V)$ *into multiples of pairwise inequivalent irreducible representations* σ_i*. Then, we have:*

$$\chi_\rho = \sum_i k_i \chi_{\sigma_i}$$

whence $\langle \chi_\rho, \chi_{\sigma_i} \rangle = k_i$*. It follows that:*

1. *The multiplicities of irreducible summands in ρ are well-defined, i.e., decompositions of ρ into irrecucible summands are unique up to reordering of summands.*

2. *Two representations are equivalent if and only if their characters agree.* **q.e.d.**

Theorem 3.2.4.15. *The characters associated to the irreducible representations for an orthonormal basis for the $\mathbb{C}$-vector space of class functions.*

Proof. We have already seen that characters are class functions (**??**) and that the characters associated to irreducible representations form an orthonormal set. It remains to show that every class function is a linear combination of irreducible characters.

Let L be the $\mathbb{C}$-span of all complex conjugates of irreducible characters inside the space of class functions. We shall show that L has trivial orthogonal complement. Thus, assume that c is a class function perpendicular to each complex conjugate of an irreducible character. Let $\rho : G \to \mathrm{Aut}_{\mathbb{C}}(V)$ be an irreducible representation. Check that

$$\begin{array}{rcl} \varphi : V & \longrightarrow & V \\ \mathbf{v} & \mapsto & \displaystyle\sum_{g \in G} c(g)\, \rho_g(\mathbf{v}) \end{array}$$

is a G-equivariant linear endomorphism. By Schur's Lemma, it is a homothety of so far unknown ratio α. Computing the trace yields:

$$\dim(V)\, \alpha = \mathrm{tr}(\varphi) = \sum_{g \in G} c(g)\, \mathrm{tr}(\rho_g(\mathbf{v})) = \sum_{g \in G} c(g)\, \chi_{\rho_g}(\mathbf{v}) = \mathrm{card}(G)\, \langle c, \overline{\chi_\rho} \rangle$$

It follows that $\alpha = 0$.

As the regular representation decomposes as a direct sum of irreducible representations, it follows that

$$\begin{array}{rcl} \mathrm{Maps}(G, \mathbb{C}) & \longrightarrow & \mathrm{Maps}(G, \mathbb{C}) \\ f & \mapsto & \displaystyle\sum_{g \in G} c(g)\, f \circ \lambda_{g^{-1}} \end{array}$$

is the trivial map. Evaluation at characteristic functions shows that c must vanish. **q.e.d.**

Corollary 3.2.4.16. *There are as many irreducible representations as conjugacy classes.* **q.e.d.**

Remark 3.2.4.17. Note that we did not discover any natural 1-1-correspondence of irreducible representations and conjugacy classes.

Corollary 3.2.4.18. *Let $\rho = \bigoplus \tau^{k_\tau}$ be a direct sum decomposition of the representation ρ into irreducibles. Then*

$$\langle \chi_\rho, \chi_\rho \rangle = \sum k_\tau^2$$

In particular, ρ is irreducible if and only if

$$\langle \chi_\rho, \chi_\rho \rangle = 1$$ **q.e.d.**

Corollary 3.2.4.19. *An irreducible representation τ occurs in the regular representation with multiplicity* $\dim(\tau)$.

Consequently:

$$\mathrm{card}(G) = \sum_{\chi \ irred.} \langle \chi, \chi \rangle = \sum_{[M] \ simple} \dim_{\mathbb{C}}(M)^2$$

Proof. Let $\rho = \bigoplus_{\tau \in \mathrm{Irr}} \tau^{k_\tau}$ be the regular representation written as a direct sum of irreducibles. Note:

$$\chi_\rho(g) = \begin{cases} \mathrm{card}(G) & g = 1 \\ 0 & g \neq 1 \end{cases}$$

Hence

$$k_\tau = \langle \chi_\rho, \chi_\tau \rangle = \frac{1}{\mathrm{card}(G)} \sum_{g \in G} \chi_\rho(g)\, \chi_\tau\big(g^{-1}\big) = \chi_\tau(1) = \dim(\tau)$$

Now:

$$\mathrm{card}(G) = \sum_{\tau \in \mathrm{Irr}} k_\tau \dim(\tau) = \sum_{\tau \in \mathrm{Irr}} \dim(\tau)^2$$

This completes the proof. **q.e.d.**

So far, we have made little use of the full strength of our result (i.e., we have not used that the irreducible characters span the class functions). That shall change:

Corollary 3.2.4.20. *For any $g \in G$, let g^G denote the conjugacy class of g. For any two group element $g, h \in G$,*

$$\sum_{\chi \text{ irred.}} \chi(g)\,\overline{\chi(h)} = \begin{cases} \frac{\operatorname{card}(G)}{\operatorname{card}(g^G)} & g, h \text{ are conjugate} \\ 0 & \text{otherwise} \end{cases}$$

Proof. Let f_g be the characteristic function of g^G. This is a class function. Thus, it is a linear combination

$$f_g = \sum_{\chi \text{ irred.}} a_\chi \chi \qquad \text{with} \qquad a_\chi = \langle f_g, \chi \rangle = \frac{\operatorname{card}\left(g^G\right)}{\operatorname{card}(G)} \overline{\chi(g)}$$

Thus,

$$\sum_{\chi \text{ irred.}} \frac{\operatorname{card}\left(g^G\right)}{\operatorname{card}(G)} \overline{\chi(g)} \chi(h) = f_g(h) = \begin{cases} 1 & h \in g^G \\ 0 & h \notin g^G \end{cases}$$

The claim follows. **q.e.d.**

Exercise 3.2.4.21. Determine all irreducible characters for the groups $\mathbf{S}_3$ and $\mathbb{Z}_4$.

Exercise 3.2.4.22. The symmetric group $\mathbf{S}_4$ acts tautlogically on the set of four letters. As with the regular representation, this action induces a linear representation (in this case, of degree 4). Write this representation as a direct sum of irreducible representations.

3.3 Modules over Principal Ideal Domains

In this section, R is always a principal ideal domain.

3.3.1 The Smith Normal Form

Let A and B be two $m \times n$-matrix with coefficients in R. We say that A and B are equivalent if there are invertible matrices $M \in \mathrm{GL}_m(R)$ and $N \in \mathrm{GL}_n(R)$ such that:

$$A = MBN$$

Theorem 3.3.1.1. *Every $m \times n$-matrix A over R is equivalent to a matrix in Smith Normal Form, i.e., a block matrix of the form*

$$\begin{pmatrix} B_{1,1} & B_{1,2} \\ B_{2,1} & B_{2,2} \end{pmatrix}$$

where all but the upper left block vanish and we have

$$B_{1,1} = \begin{pmatrix} a_1 & 0 & \cdots & 0 \\ 0 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & a_r \end{pmatrix}$$

with $a_1 | a_2 | \cdots | a_r$.

Moreover, the entries ... are uniquely determined up to units by A. In fact, $a_1 a_2 \cdots a_k$ is the greatest common divisor of all k-minors of A.

Proof. We break this up into a couple of claims. Hidden within this proof is actually an algorithm that one can make effective.

Before we start, we note that elementary matrices are invertible. Thus, we can swap rows and columns freely and we can add multiples of a a row to another row (or multiples of a column to another column) without changing the equivalence class of a matrix. We can, however, not divide. To overcome this, we need one more type of operation: secondary row/column operations. Those are described as left- or right-multiplication by a block-matrix of the following form:

$$\begin{pmatrix} a & b & \\ c & d & 0 \\ & 0 & 1 \end{pmatrix}$$

where the determinant $ad - bc = 1$. Thus, secondary matrices are invertible.

Claim A. A is equivalent to a matrix B whose first row and column vanish everywhere except for the upper left corner, wihch also divides all other entries.

PROOF. Recall that any PID is a UFD. Define the rank of any element to be the exponent-sum of in its prime-factor decomposition. For units, let the length be 0; and we declare the length of 0 to be ∞.

Let B' be a matrix equivalent to A so that its upper left entry has minimum length. We claim that in this case, the upper left entry $a_{1,1}$ divides all other entries in the first row and first column.

First assume that the first column contains another entry not a multiple of $a_{1,1}$. Since we can permute rows, we may assume that $a_{1,1}$ does not divide $a_{2,1}$. Let d be a greatest common divisor of $a_{1,1}$ and $a_{2,1}$ and write it as a combinator

$$d = a_{1,1}a_{1,1} - a_{1,2}a_{2,1}$$

Moreover, write $a_{1,1} = da_{2,1}$ and $a_{2,1} = da_{2,2}$. Then

$$1 = \begin{pmatrix} a_{1,1} & a_{1,2} \\ a_{2,1} & a_{2,2} \end{pmatrix}$$

Thus, we found a secondary operation that will transform B into a matrix that has upper left entry d. This contradicts minimality of the length in that entry.

The same construction shows that $a_{1,1}$ divides all other entries in the first row.

Now we use elementary column and elementary row operations to kill off all entries in the first row and in the first column

apart from the upper left corner. We chose the resulting matrix as B.

We claim that the upper left entry $a_{1,1}$ divides all other entries. Suppose not, then we could use an elementary row or column operation to copy that entry into the first row or column. It follows from the preceeeding considerations that $a_{1,1}$ divides evenly into the copy. □

Claim B. A is equivalent to a matrix in Smith Normal Form.

PROOF. Use the previous claim to kill off the first row and column and make all remaining entries multiples of the upper left corner. Now, recurse into the submatrix obtained by choping off the first row and the first column. By induction, we can put that part into Smith Normal Form. □

Now we shall turn to uniqueness. For any matrix A over R let $\det_k(A)$ denote the greatest common divisor of all k-minors, (recall that a k-minor is the determinant of a $k \times k$-submatrix of A).

Claim C. $\det_k(A)$ divides $\det_k(MAN)$.

PROOF. Determinants are linear forms in the columns. The columns of AN are linear combinations of columns in A. It follows that the determinant of any $k \times k$-submatrix in AN is a linear combination of r-minors in A. Thus, $\det_k(A)$ divides $\det_k(AN)$.

Running the same argument for rows, we can take care of the left-multiplication by M. □

Consequently, $\det_k(A)$ is an invariant (up to units) of the equivalence class of A. To finish the proof, we just need to observe that $\det_k(A) = a_1 \cdots a_k$ when A has Smith Normal Form. **q.e.d.**

Exercise 3.3.1.2. Let $A := \begin{pmatrix} 2 & -1 & 2 \\ 5 & 0 & -3 \end{pmatrix}$. Find $M \in \mathbb{M}_{2,2}(\mathbb{Z})$ and $N \in \mathbb{M}_{3,3}(\mathbb{Z})$ so that MAN has Smith-Normal-Form.

Exercise 3.3.1.3. Let $A := \begin{pmatrix} x & x-1 & x^2-1 \\ 1 & 0 & x+1 \end{pmatrix}$. Find $M \in \mathbb{M}_{2,2}(\mathbb{Q}[x])$ and $N \in \mathbb{M}_{3,3}(\mathbb{Q}[x])$ so that MAN has Smith-Normal-Form.

Exercise 3.3.1.4. Let $A := \begin{pmatrix} x & x+1 & x^3+1 \\ 1 & 0 & x+1 \end{pmatrix}$. Find $M \in \mathbb{M}_{2,2}(\mathbb{F}_2[x])$ and $N \in \mathbb{M}_{3,3}(\mathbb{F}_2[x])$ so that MAN has Smith-Normal-Form.

3.3.2 Presentations of Finitely Generated Modules

Definition 3.3.2.1. Let M be a left-R-module. A presentation of M is an exact sequence

$$F_2 \to F_1 \to M \to 0$$

where F_1 and M_2 are both free left-R-modules. M is called finitely presented if it admits a presentation where F_1 and F_2 are of finite rank.

Proposition 3.3.2.2. *Any submodule S of a free left-R-module R^r of finite rank is free of rank at most r.*

Proof. We induct on r. Note that the staments is trivial for $r = 0$ and amounts to the defintion of PID for $r = 1$.

For $r > 1$, consider the obvious short exact sequence

$$R^{r-1} \hookrightarrow R^r \twoheadrightarrow R$$

Intersecting with S induces the short exact sequence

$$R^{r-1} \cap S \hookrightarrow S \twoheadrightarrow {}^{S}\Big/_{R^{r-1} \cap S}$$

where

$${}^{S}\Big/_{R^{r-1} \cap S} = {}^{S + R^{r-1}}\Big/_{R^{r-1}} \leq {}^{R^r}\Big/_{R^{r-1}} \cong R$$

is free of rank at most one. If follows that

$$R^{r-1} \cap S \hookrightarrow S \twoheadrightarrow {}^{S}\Big/_{R^{r-1} \cap S}$$

splits, i.e.:

$$S = \left(R^{r-1} \cap S\right) \oplus \left({}^{S}\!/_{R^{r-1} \cap S}\right)$$

Moreover, $R^{r-1} \cap S$ is a free module of rank at most $r-1$ by induction. Thus, M is free of rank at most r. **q.e.d.**

Exercise 3.3.2.3. Read up on transfinite induction and decide whether the above argument can be twisted to show the analogous statement for free modules of possibly infinite rank.

Corollary 3.3.2.4. *Every finitely generated R-module M is finitely presented. In fact, one can find a short exact sequence*

$$F_2 \hookrightarrow F_1 \twoheadrightarrow M$$

where F_1 and F_2 are free left-R-modules of finite rank.

Proof. Since M is finitely generated, it is the epimorphic image $F_1 \twoheadrightarrow M$. By the preceeding proposition, the kernel is a free left-R-module F_2. **q.e.d.**

Compatible Basis Theorem 3.3.2.5. *Let R be a PID and let $F' \leq F$ be an inclusion of finitely generated free R-modules. Then there are bases*

$$\begin{aligned} F' &= \langle f'_1, f'_2, \ldots, f'_u \rangle \\ F &= \langle f_1, f_2, \ldots, f_u, f_{u+1}, \ldots, f_v \rangle \end{aligned}$$

and ring elements $a_1, a_2, \ldots, a_u$ satisfying

1. *the identities $f'_i = a_i f_i$, and*
2. *the divisibilty condition $a_1 | a_2 | \cdots | a_u$.*

Moreover, the invariant factors $a_1, a_2, \ldots, a_u$ are uniquely determined up to units.

Proof. Choose bases for F and S. Write the inclusion map as a matrix. Put this matrix into Smith-Normal-Form. It follows that there are bases for F and S for which the inclusion is represented by a Smith-Normal-Form. Those bases are the ones called for.

As for uniqueness, note that given bases as in the theorem, the matrix describing the inclusion has Smith-Normal-Form. Since coefficients of the Smith-Normal-Form are unique, any other choice of bases would give rise to the same invariant factors. **q.e.d.**

Corollary 3.3.2.6. *Every finitely generated module over a PID is a finte direct sum of cyclic modules.*

More precisely: Let M be a finitely generated left-R-module over a PID. Then there exist a number r and elements $a_1, \ldots, a_t \in R$ satisfying

$$a_1 | a_2 | \cdots | a_t$$

so that:

$$M \cong R^r \oplus \bigoplus_{i=1}^{t} R/\langle a_i \rangle$$

Proof. Let

$$R^s \hookrightarrow R^{s'} \twoheadrightarrow M$$

be a finite presentation of M and assume that left arrow is multiplication by a matrix in Smith-Normal-Form. **q.e.d.**

Remark 3.3.2.7. We shall see later that the decomposition given here is essentially unique.

Exercise 3.3.2.8. Consider the free $\mathbb{Z}$-module $F := \mathbb{Z}^3$ of rank 3 and let S be the submodule defined by

$$S := \left\{ \begin{pmatrix} a \\ b \\ c \end{pmatrix} \;\middle|\; 2a + 3b - c = 0 \right\}$$

Find compatible bases for F and S, i.e., bases such as given in the Compatible Bases Theorem.

3.3.3 Torsion and Annihilation

Definition 3.3.3.1. Let M be an R-module. An element $m \in M$ is called a torsion element if there is a non-zero scalar $a \in R - \{0\}$ with $am = 0$.

The ring M is called torsion free if 0 is the only torsion element in M.

The module M is called a torsion module if all elements in M are torsion elements.

Observation 3.3.3.2. *Every quotient and every submodule of a torsion module is a torsion module.*

Exercise 3.3.3.3. Show that all, over integral domains, all free modules are torsion free.

Observation 3.3.3.4. *Let M be a module over an integral domain. The torsion elements in M form a sub-module $T(M)$, called the torsion part of M.*

Proof. Let m_1, m_2 be torsion elements with annihilating scalars $a_1, a_2 \neq 0$. Then $a_1 a_2 \neq 0$ is an annihilating scalar for $m_1 + m_2$. It is even easier to see that multiples of torsion elements are torsion. **q.e.d.**

Corollary 3.3.3.5. *Every finitely generated module M over a PID R decomposes as*

$$M = R^r \oplus T$$

where T is the torsion part. Moreover, the free rank r of M is uniquely determined by M.

Proof. We have

$$M = R^r \oplus \bigoplus_i {}^{R}\!/_{\langle a_i \rangle}$$

and $T = \bigoplus_i {}^{R}\!/_{\langle a_i \rangle}$. It follows that ${}^{M}\!/_{T}$ is free whence r is unique. **q.e.d.**

Corollary 3.3.3.6. *Every finitely generated torsion free module over a PID is free.*

Proof. The torsion part vanishes. **q.e.d.**

Chinese Remainder Theorem 3.3.3.7. *Let a and b be elements in a PID R. The left-R-module homomorphism*

$$\begin{aligned} \varphi : R \oplus R &\longrightarrow R \\ (x, y) &\mapsto xb + ya \end{aligned}$$

is onto if and only if a and b are relatively prime. In this case, we have:

$$\varphi^{-1}(\langle ab \rangle) = \langle a \rangle \oplus \langle b \rangle$$

Consequently, the induced homomorphism

$$\begin{aligned} {}^{R}\!/_{\langle a \rangle} \oplus {}^{R}\!/_{\langle b \rangle} &\longrightarrow {}^{R}\!/_{\langle ab \rangle} \\ (x + \langle a \rangle, y + \langle b \rangle) &\mapsto xb + ya + \langle ab \rangle \end{aligned}$$

is an isomorphism of left-R-modules if and only if a and b are relatively prime.

Proof. The elements a and b are relatively prime if and only if we can combine 1 linearly from them. This is visibly equivalent to φ being onto.

Now let a and b be relatively prime and assume that the linear combination $xb + ya$ is a multiple of ab. Note that xb is always a multiple of b. It follows that ya is a multiple of b, too. Thus, y is a multiple of b since b and a are relatively prime. Similarly, $x \in \langle a \rangle$. Hence, $\varphi^{-1}(\langle ab \rangle) = \langle a \rangle \oplus \langle b \rangle$ as claimed. **q.e.d.**

Corollary 3.3.3.8. *Every finitely generated torsion module M over a PID R is a finite direct sum*

$$M = \bigoplus_i {}^{R}\!/_{\langle p_i^{k_i} \rangle}$$

Proof. We have

$$M = T(M) = \bigoplus_j {}^{R}\!\Big/_{\langle a_j \rangle}$$

Now apply the Chinese Remainder Theorem to all the cyclic torsion modules: each a_j has a prime factor decomposition. **q.e.d.**

Exercise 3.3.3.9. Let $I \trianglelefteq R$ be a left ideal and let M be an R-module. Show that

$$IM := \left\{ \sum_{i=1}^{u} i_i m_i \;\middle|\; u \in \mathbb{N},\ i_i \in I,\ m_i \in M \right\},$$

i.e., the set of all finite combinations of module elements with coefficients from I, is a submodule of M. Moreover, show that ${}^{M}\!/_{IM}$ is an ${}^{R}\!/_{I}$-module.

Use the above to show that any R-module M can be regarded as a module over ${}^{R}\!/_{\mathrm{Ann}(M)}$ via

$$(a + \mathrm{Ann}(M))m := am.$$

3.3.4 The Classification of Finitely Generated Modules

Lemma 3.3.4.1. *Let R be a PID, let $p \in R$ be a prime. For any left-R-module M and any $k \in \mathbb{N}$, put*

$$M_p^k := \left\{ m \in M \,\middle|\, p^k m = 0 \right\}.$$

Then M_p^{k+1}/M_p^k is a vector space over the field $R/\langle p \rangle$.

Proof. Changing M, we may assume $k = 0$. Thus, it suffices to show that $M_p^1 = \{m \in M \mid pm = 0\}$ is a $R/\langle p \rangle$-vector space. I.e., we have to observe that p (and the ideal it generates) acts as 0 multiplicatively. That, however, holds by definition. **q.e.d.**

Now, we compute:

$$\left({}^{R}\!/_{\langle q^l \rangle}\right)_p^k = \left\{ a + \langle q^l \rangle \,\middle|\, q^l | p^k a \right\}$$

If $\langle q \rangle \neq \langle p \rangle$, we have

$$q^l | p^k a \quad \Longleftrightarrow \quad q^l | a$$

whence

$$\left({}^{R}\!/_{\langle q^l\rangle}\right)_p^k = \{0\} \qquad \text{for } \langle q\rangle \neq \langle p\rangle.$$

If $\langle q\rangle = \langle p\rangle$, we have

$$q^l | p^k a \quad \Longleftrightarrow \quad l \leq k \text{ or } p^{l-k} | a$$

whence, in this case,

$$\left({}^{R}\!/_{\langle q^l\rangle}\right)_p^k = \begin{cases} R/\langle p^l\rangle & l \leq k \\ \left\{a + \langle p^l\rangle \mid p^{l-k} | a\right\} = \langle p^{l-k}\rangle/\langle p^l\rangle & k \leq l \end{cases}$$

We summarize:

Lemma 3.3.4.2.

$$\left({}^{R}\!/_{\langle q^l\rangle}\right)_p^k = \begin{cases} \{0\} & \langle q\rangle \neq \langle p\rangle \\ R/\langle p^l\rangle & \langle q\rangle = \langle p\rangle \text{ and } l \leq k \\ \langle p^{l-k}\rangle/\langle p^l\rangle & \langle q\rangle = \langle p\rangle \text{ and } k \leq l \end{cases}$$

Put

$$d_p^k(M) := \dim_{R/\langle p\rangle}\left(M_p^{k+1}/M_p^k\right)$$

Lemma 3.3.4.3.

$$d_p^k\left({}^{R}\!/_{\langle q^l\rangle}\right) = \begin{cases} 0 & \langle q\rangle \neq \langle p\rangle \\ 0 & \langle q\rangle = \langle p\rangle \text{ and } l \leq k \\ 1 & \langle q\rangle = \langle p\rangle \text{ and } k < l \end{cases}$$

Observation 3.3.4.4. *Since*

$$(M \oplus N)_p^k = M_p^k \oplus N_p^l$$

we have

$$d_p^k(M \oplus N) = d_p^k(M) + d_p^k(N)$$

Structure Theorem 3.3.4.5. *Let M be a finitely generated left-R-module over a PID. Then there exist a number r and elements $a_1, \ldots, a_t \in R$ satisfying*

$$a_1 | a_2 | \cdots | a_t$$

so that:

$$M \cong R^r \oplus \bigoplus_{i=1}^{t} R/\langle a_i \rangle$$

Moreover, $r, t,$ *and the* *invariant factors* *are uniquely determined (up to units, that is).*

Equivalently, there exist prime elements p_j *and exponents* k_j *so that*

$$M \cong R^r \oplus \bigoplus_{j} {}^{R}\Big/_{\langle p_j^{k_j} \rangle}$$

where the *elementary divisors* $p_j^{k_j}$ *are unique up to order and units.*

Proof. The free rank r is uniquely determined by (**??**).

The two versions are equivalent by the Chinese Remainder Theorem and uniqueness of prime factor decompositions in PIDs.

It remains to prove uniqueness of elementary divisors. That, however, follows since the numbers $d_p^k(T(M))$ determine how many terms of the form ${}^{R}\big/_{\langle q^l \rangle}$ occur in the torsion part. **q.e.d.**

3.3.5 Advanced Linear Algebra

The crucial observation of advanced linear algebra is the following

Example 3.3.5.1. Let K be a field and let V be a finite dimensional vector space over K. Fix an endomorphism $\varphi : V \to V$. Then, V can be given the structure of a $K[x]$-module as follows via

$$p\mathbf{v} := p(\varphi)(\mathbf{v}),$$

where $p(\varphi) : V \to V \in \mathrm{End}_K(V)$ is the endomorphism obtained by evaluating $p(x)$ at $x = \varphi$. We denote the $K[x]$-module defined this way by V_φ.

Definition 3.3.5.2. Let $\varphi : V \to V$ and $\psi : W \to W$ be two endomorphisms of K-vector spaces V and W. We call φ and ψ similar if there is a

K-isomorphism $\mu : V \to W$ such that

$$\begin{array}{ccc} V & \xrightarrow{\varphi} & V \\ {\scriptstyle\mu}\downarrow & & \downarrow{\scriptstyle\mu} \\ W & \xrightarrow[\psi]{} & W \end{array}$$

commutes.

Exercise 3.3.5.3. Fix a pair of bases for V and W. Show that two endomorphisms $\varphi : V \to V$ and $\psi : W \to W$ are similar if and only they are represented by similar matrices (with respect to the fixed bases for V and W, respectively).

Exercise 3.3.5.4. Show that two endomorphisms $\varphi : V \to V$ and $\psi : W \to W$ are similar if and only if the associated $K[x]$-modules V_φ and W_ψ are $K[x]$-ismorphic.

Observation 3.3.5.5. *The module V_φ is a torsion module.* **q.e.d.**

Observation 3.3.5.6. *Since V has finite dimension over K and since any K-basis of K is a generating set of V_φ, it follows that V_φ is finitely generated torsion module. Thus,*

$$V_\varphi = \bigoplus_i K[x] \Big/ \langle q_i(x) \rangle$$

where the polynomials $q_i(x)$ form a divisor chain $q_1|q_2|\cdots|q_u$. Since this condition determines the q_i up to units, we can eliminate all remaining freedom of choice by requirering that each q_i be monic (i.e., has leading coefficient 1). **q.e.d.**

Definition 3.3.5.7. The polynomial q_u is called the minimal polynomial of φ. Usually, it is denoted by $\mu_\varphi(x)$.

Exercise 3.3.5.8. Show that the minimal polynomial generates the annihilator ideal of V_φ, i.e.:

$$\langle \mu_\varphi(x) \rangle = \{ p \in K[x] \mid p\mathbf{v} = 0 \text{ for all } \mathbf{v} \in V_\varphi \}$$

Example 3.3.5.9 (Cyclic Endomorphisms). Consider the cyclic torsion module $M := {}^{K[x]}\!/_{\langle q(x)\rangle}$ where

$$q(x) = x^m + a_{m-1}x^{m-1} + \cdots a_1 x + a_0$$

is a monic polynomial of degree m. Note that M is is is a K-vector space with ordered basis $B := (1, x, x^2, \ldots, x^{m-1})$. Left-multiplication by x is a K-linear map. It is easy to work out the matrix representing this map relative to the basis B. We find:

$$\begin{pmatrix} 0 & \cdots & \cdots & 0 & -a_0 \\ 1 & 0 & \cdots & 0 & -a_1 \\ 0 & \ddots & \ddots & \vdots & \vdots \\ \vdots & \ddots & \ddots & 0 & -a_{m-2} \\ 0 & \cdots & 0 & 1 & -a_{m-1} \end{pmatrix}$$

We call a matrix of this form an RCF-block.

Corollary 3.3.5.10. *Every square matrix is similar to one and only one matrix in rational canonical form, i.e., in block diagonal form where the blocks along the diagonal are RCF-blocks and the corresponding minimal polynomials form a divisor chain.*

Consequently, two matrices are similar if and only if they have the same rational canonical form.

Corollary 3.3.5.11. *Let K be a subfield of the field M. Two matrices $A, B \in \mathbb{M}_r(K)$ are similar over K if and only they are similar over M. In particular, two matrices in $\mathrm{GL}_r(K)$ are conjugate in $\mathrm{GL}_r(M)$ if they are conjugate in $\mathrm{GL}_r(M)$.*

Proof. Note that the rational canonical form $\mathrm{RCF}_K(A)$ also qualifies as a rational canonical form over M. Uniqueness of rational canonical forms implies $\mathrm{RCF}_K(A) = \mathrm{RCF}_M(A)$. Thus, we have the

following chain of equivalences:

$$\begin{aligned} & A \text{ and } B \text{ are similar orver } K \\ \iff & \mathrm{RCF}_K(A) = \mathrm{RCF}_K(B) \\ \iff & \mathrm{RCF}_M(A) = \mathrm{RCF}_M(B) \\ \iff & A \text{ and } B \text{ are similar orver } M \end{aligned}$$

q.e.d.

Cayley-Hamilton Theorem 3.3.5.12. *For any endomorphism $\varphi : V \to V$ of a finite dimensional K-vector space, the characteristic polynomial $\chi_\varphi(x) = \det(x\,\mathrm{id} - \varphi)$ is the product of the invariant factors. I.e., let*

$$V_\varphi = \bigoplus_i K[x] \Big/ \langle q_i(x) \rangle$$

be the canonical decomposition of V_φ with invariant factors $q_i(x)$. Then

$$\chi_\varphi(x) = \prod_i q_i(x)\,.$$

In particular, the minimal polynomial divides the characteristic polynomial: $\mu_\varphi | \chi_\varphi$. Consequently, $\chi_\varphi(\varphi) = 0 \in \mathrm{End}_K(V)$.

Proof. We just compute the characteristic polynomial using a basis of V relative to which φ has rational canonical form.

Relative to such basis, $x\,\mathrm{id} - \varphi$ has block diagonal form. Thus, it suffices to show that for an individual RFC-block, we have

$$x^m + a_{m-1}x^{m-1} + \cdots a_1 x + a_0 = \det \begin{pmatrix} x & 0 & \cdots & 0 & a_0 \\ -1 & x & \ddots & \vdots & a_1 \\ 0 & \ddots & \ddots & 0 & \vdots \\ \vdots & \ddots & \ddots & x & a_{m-2} \\ 0 & \cdots & 0 & -1 & x + a_{m-1} \end{pmatrix}$$

This is an easy induction. The induction hypothesis yields

$$x^{m-1} + a_{m-1}x^{m-1} + \cdots a_2 x + a_1 = \det \begin{pmatrix} x & 0 & \cdots & 0 & a_1 \\ -1 & x & \ddots & \vdots & a_2 \\ 0 & \ddots & \ddots & 0 & \vdots \\ \vdots & \ddots & \ddots & x & a_{m-2} \\ 0 & \cdots & 0 & -1 & x + a_{m-1} \end{pmatrix}$$

Then, developing along the first row yields:

$$\det \begin{pmatrix} x & 0 & \cdots & 0 & a_0 \\ -1 & x & \ddots & \vdots & a_1 \\ 0 & \ddots & \ddots & 0 & \vdots \\ \vdots & \ddots & \ddots & x & a_{m-2} \\ 0 & \cdots & 0 & -1 & x + a_{m-1} \end{pmatrix} = x \det \begin{pmatrix} x & 0 & \cdots & 0 & a_1 \\ -1 & x & \ddots & \vdots & a_2 \\ 0 & \ddots & \ddots & 0 & \vdots \\ \vdots & \ddots & \ddots & x & a_{m-2} \\ 0 & \cdots & 0 & -1 & x + a_{m-1} \end{pmatrix} \pm a_0 \det \begin{pmatrix} -1 & x & 0 \\ 0 & \ddots & \ddots \\ \vdots & \ddots & \ddots \\ \vdots & \ddots & \ddots \\ 0 & \cdots & \cdots \end{pmatrix}$$

This proves the claim. **q.e.d.**

Remark 3.3.5.13. Often, knowing the minimal polynomial and the characteristic is enough to deduce the invariant factors.

Example 3.3.5.14. We classify the matrices of order 3 in $\mathrm{GL}_5(\mathbb{Q})$ up to similarity. Then, all invariant factors divide the polynomial $x^3 - 1$ which decomposes into irreducibles (over $\mathbb{Q}[x]$) as follows:

$$x^3 - 1 = (x-1)\big(x^2 + x + 1\big)$$

Since the degrees of the invariant factors add up to 5, we have only the following choices for the invariant divisors:

$$x - 1 | x - 1 | x^3 - 1$$

and

$$x^2 + x + 1 | x^3 - 1$$

Note: $x - 1$ cannot be the minimal polynomial since in this case, the order would be 1.

Exercise 3.3.5.15. A matrix A is nilpotent if $A^k = 0$ for some exponent k. Classify, up to similarity, all nilpotent matrices in $\mathbb{M}_{5,5}(\mathbb{Q})$. Does the classification change when you replace $\mathbb{Q}$ by the field $\mathbb{F}_3$ of three elements? What happens for $\mathbb{F}_2$?

Exercise 3.3.5.16. This is a fall back method for computing the rational canonical form. Let A be an $n \times n$-matrix over K. Show that the non-unit entries in the Smith Normal Form of $x\,\mathrm{id}_n - A \in \mathbb{M}_{n,n}(K[x])$ are the invariant factors for A. Hint: The evaluation $K[x]$-homomorphism can be extended to a finite presentation of the $K[x]$-module K^n_A. It is given as

$$\begin{array}{rcl} K[x]^n & \longrightarrow & K^n_A \\ \sum_i p_i(x)\,\mathbf{e}_i & \mapsto & \sum_i p_i(A)\,\mathbf{e}_i \end{array}$$

where $\mathbf{e}_i$ denotes the ith standard basis vector (in both modules!).

Observation 3.3.5.17. *If $\mu_\eta(x)$ splits into linear factors, it is very convenient to use elementary divisors instead of invariant factors. We obtain the decompostion*

$$V_\eta \cong \bigoplus_i {K[x]}\Big/{\langle (x-\lambda_i)^{k_i}\rangle}$$

where $\lambda_i \in K$. The decomposition is unique up to order of summands.

Example 3.3.5.18. We can describe a single summand of the decomposition. Let $M := {K[x]}/{\langle (x-\lambda)^k\rangle}$. We describe multiplication by x relative to the ordered basis

$$\left((x-\lambda)^0, (x-\lambda)^1, \ldots, (x-\lambda)^{k-1}\right)$$

The corresponding matrix is

$$\begin{pmatrix} \lambda & 0 & 0 & \cdots & 0 \\ 1 & \lambda & 0 & \ddots & \vdots \\ 0 & 1 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \lambda & 0 \\ 0 & \cdots & 0 & 1 & \lambda \end{pmatrix}$$

Such a matrix is called a Jordan block.

Corollary 3.3.5.19. *If $\mu_\varphi(x)$ splits into linear factors over K, then φ is representable in Jordan canonical form, i.e., block diagonal form where all block are Jordan blocks. The Jordan canonical form, when it exists, is unique up to order of Jordan blocks.*

Remark 3.3.5.20. Note that $\mu_\varphi(x)$ splits over any algebraically closed field. Consequently, if K is algebraically closed, every K-endomorphism has a Jordan canonical form.

Here is a sample application of Jordan canonical forms:

Proposition 3.3.5.21. *Every matrix is similar to its transpose.*

Proof. First note that by (3.3.5.11), we can assume that we work over an algebraically closed field: every field is a subfield of an algebraically closed field. Thus, it suffices to show that Jordan blocks are similar to their transpose. This is seen through a simple matrix multiplcication:

$$\begin{pmatrix} & & 1 \\ & \diagup & \\ 1 & & \end{pmatrix} \begin{pmatrix} \lambda & 0 & 0 & \cdots & 0 \\ 1 & \lambda & 0 & \ddots & \vdots \\ 0 & 1 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \lambda & 0 \\ 0 & \cdots & 0 & 1 & \lambda \end{pmatrix} \begin{pmatrix} & & 1 \\ & \diagup & \\ 1 & & \end{pmatrix} = \begin{pmatrix} \lambda & 1 & 0 & \cdots & 0 \\ 0 & \lambda & \ddots & \ddots & \vdots \\ 0 & 0 & \ddots & 1 & \vdots \\ \vdots & \ddots & \ddots & \lambda & 1 \\ 0 & \cdots & 0 & 0 & \lambda \end{pmatrix}$$

Note that the matrix $\begin{pmatrix} & & 1 \\ & \diagup & \\ 1 & & \end{pmatrix}$ has order two so that the above equation truly proves similarity. **q.e.d.**

Chapter 4

Fields

4.1 Field Extensions

4.1.1 Basic Definitions

Definition 4.1.1.1. A field is a commutative division ring.

Observation 4.1.1.2. *A domain is a field if and only if it has exactly two ideals.* **q.e.d.**

Corollary 4.1.1.3. *Any ring homomorphism from a field to any ring is* 1-1 *since its kernel is proper ideal (the kernel cannot contain* 1*).*

Since morphisms between fields are always injective, most of the time, we shall just pretend that one field is actually a subfield of another.

Definition 4.1.1.4. A extension of a field K is a field M together with a field morphism $K \hookrightarrow M$. We shall regularly pretend that K is contained as a subset in M and that the inclusion is the specified field morphism.

Definition 4.1.1.5. If M/K is a field extension, then M is a K-vector space. We define the degree of the field extension to be the dimension $[M : K] := \dim_K(M)$. An extension is called finite if it has finite degree.

Exercise 4.1.1.6. Let $M/K/F$ be a tower of fields; and let B_F^K be a basis of K as an F-vector space and let B_K^M be a basis of M as an K-vector space. Then the set

$$C_M^F := \left\{\zeta\xi \,\middle|\, \xi \in B_F^K, \zeta \in B_K^M\right\}$$

is a basis of M as an F-vector space. In particular:

$$[M : F] = [M : K]\,[K : F]\,.$$

Definition 4.1.1.7. Let M/K and N/K be two field extensions of the same base field. A K-morphism from M to N is a field morphism

$$\varphi : M \longrightarrow N$$

that fixes K element-wise, i.e., the following diagram commutes:

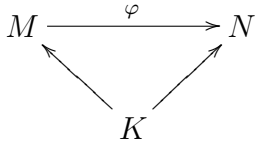

4.1.2 Algebraic and Transcendent Elements

Observation and Definition 4.1.2.1. Let M/K be a field extension and fix an element $\zeta \in M$. Evaluation $x \mapsto \zeta$ induces a ring homomorphism

$$\begin{array}{rccc} \mathrm{ev}_\zeta : K[x] & \longrightarrow & M \\ p(x) & \mapsto & p(\zeta) \end{array}$$

We denote the image of the evaluation homomorphism by $K[\zeta]$. This is the smallest intermediate ring between K and M that contains ζ. Its field of fractions canonically embeds into M; and this way, we obtain $K(\zeta)$, which is the smallest intermediate field between K and M that contains ζ.

Let

$$\mathcal{K}_\zeta := \{p(x) \in K[x] \mid p(\zeta) = 0\}$$

be the kernel of the evaluation homomorphism. Then

$$K[\zeta] \cong {}^{K[x]}\Big/_{\mathcal{K}_\zeta} \ .$$

Note that $\mathcal{K}_\zeta$ is an ideal, and as $K[x]$ is a PID, the ideal $\mathcal{K}_\zeta$ is generated by one element. There are two cases:

$\underline{\mathcal{K}_\zeta = \{0\}}$: In this case, the evaluation homomorphism is injective and the ring $K[\zeta]$ is isomorphic to the polynomial ring. Thus, it has infinite dimension over K (which implies that this cannot happen if M/K is finite). Also, $K(\zeta)$ is isomorphic to the field of rational functions in one variable.

We say that ζ is $\underline{\text{transcendent over } K}$.

$\underline{\mathcal{K}_\zeta \neq \{0\}}$: In this case, the $\mathcal{K}_\zeta$ is generated by a single non-zero polynomial. This polynomial is unique up to units. Thus, we can make it unique by normalizing the leading coefficient. The unique monic generator of $\mathcal{K}_\zeta = \langle \mu_\zeta(x) \rangle$ is called the $\underline{\text{minimal polynomial}}$ of ζ. Its degree equals the (finite!) dimension $\dim_K(K[\zeta])$.

As the image $K[\zeta] \leq M$ has no zero-divisors, the polynomial $\mu_\zeta(x)$ is irreducible, hence the image $K[\zeta] \cong {}^{K[x]}\Big/_{\langle \mu_\zeta(x) \rangle}$ is, in fact, a field and we have

$$K[\zeta] = K(\zeta) \, .$$

We say that ζ is $\underline{\text{algebraic over } K}$.

Note that multiplication by ζ induces a K-linear vector space endomorphism

$$\lambda_\zeta : M \longrightarrow M$$

and it turns out that μ_ζ is also the minimal polynomial of the endomorphism λ_ζ.

Definition 4.1.2.2. An extension M/K is called $\underline{\text{algebraic}}$ if each element of M is algebraic over K.

An extension M/K is called $\underline{\text{simple}}$ if there is an element $\zeta \in M$ such that $M = K(\zeta)$.

Corollary 4.1.2.3. *A simple extension is finite if and only if it is algebraic.* **q.e.d.**

Exercise 4.1.2.4. Let $M = K(\xi)/K$ be a finite simple field extension. Let $\varphi : M \to M$ be left-multiplication by ξ and regard φ as a K-linear endomorphism of the K-vector space M. Show that the minimal polynomial of ξ and the minimal polynomial of φ coincide.

Example 4.1.2.5. Any finite extension of a finite field is simple.

Proof. Finite subgroups of multiplicative groups in fields are cyclic. Thus, any finite field is generated by a single element. **q.e.d.**

Exercise 4.1.2.6. Let $M/K/F$ be a tower of fields and fix $\alpha \in M$. Show that the minimal polynomial $\mu_{\alpha,K}$ divides the minimal polynomial $\mu_{\alpha,F}$.

Exercise 4.1.2.7. Let $M/K/F$ be a tower of fields with M/F finite and fix $\alpha \in M$. Show that there is no proper subfield of K containing F that contains all coefficients of $\mu_{\alpha,K}$.

Exercise 4.1.2.8. Show that a simple extension M/F has only finitely many intermediate fields.

Observation 4.1.2.9. *Let $\xi, \zeta, \zeta, \xi \in F$ with $\zeta \neq \xi$ and assume $F(\xi + \zeta\zeta) = F(\xi + \xi\zeta)$. Then $F(\xi, \zeta) = F(\xi + \zeta\zeta)$.*

Proof. First note that $\zeta = \frac{(\xi+\zeta\zeta)-(\xi+\xi\zeta)}{\zeta-\xi} \in F(\xi + \zeta\zeta)$. It follows that $\xi \in F(\xi + \zeta\zeta)$. **q.e.d.**

Exercise 4.1.2.10. Show that a finite extension M/F is simple if it has only finitely many intermediate fields.

Observation 4.1.2.11. *An extension M/K is algebraic if and only if M is the union of intermediate fields of finite degree over K. In particular, every finite extension is algebraic.* **q.e.d.**

Proposition 4.1.2.12. *Let M/K be a field extension and fix finitely many elements $\zeta_1, \ldots, \zeta_u \in M$. Then the following are equivalent:*

1. *Each ζ_i is algebraic over K.*
2. *The extension $K(\zeta_1, \ldots, \zeta_u)/K$ is finite.*
3. *The extension $K(\zeta_1, \ldots, \zeta_u)/K$ is algebraic.*

Proof. We show that the first condition implies the second. The other implications are obvious.

A simple algebraic extension is finite. Since $K(\zeta_1, \ldots, \zeta_i)/K(\zeta_1, \ldots, \zeta_{i-1})$ is a simple algebraic extension, a simple induction finishes the proof. **q.e.d.**

Corollary and Definition 4.1.2.13. *Let M/K be an extension. The set*

$$K_{\text{alg}} := \{\zeta \in M \mid \zeta \text{ is algebraic over } K\}$$

is a subfield of M containing K. It is called the relative algebraic closure of K in M.

Proof. We have to show that K_{alg} is closed with respect to arithmetic operations.

If two elements in M are algebraic over K, they generate a finite, hence algebraic, extension of K. This proves that their sum, difference, product, and quotient are again algebraic over K. **q.e.d.**

Corollary 4.1.2.14. *If M/K and K/F are both algebraic, then so is M/F.*

Proof. Any element ζ is algebraic over K, i.e., has a minimal polynomial with coefficients $\xi_0, \ldots, \xi_u \in K$. Hence, ζ is already algebraic over $F(\xi_0, \ldots, \xi_u)$ which is a finite extension of F. Hence $F(\xi_0, \ldots, \xi_u, \zeta)$ is a finite extension of F. **q.e.d.**

4.1.3 Splitting Fields

Definition 4.1.3.1. Let $p(x) \in K[x]$ a polynomial with coefficients in the field K. A splitting field for p is a field extension K_p/K that satisfies:

1. The polynomial splits into linear factors as an element of $K_p[x]$.
2. The field K_p is minimal among those satisfying (1), i.e., the polynomial $p(x)$ does not split into linear factors over any proper subfield of K_p.

We shall show that every polynomial has a splitting field and that such splitting field is essentially unique. We will tackle uniqueness first.

Observation 4.1.3.2. *Let $p(x) \in K[x]$ a polynomial with coefficients in the field K. Then K_p is generated as a field extension by the finitely many roots of p. Since all these roots are algebraic, K_p/K is a finite field extension. Now it follows that K_p is generated by the roots of p as an K-algebra.* **q.e.d.**

Proposition 4.1.3.3. *For $i \in \{0,1\}$, let M_i/K_i be a field extension and let $\zeta_i \in M_i$ be an algebraic element with minimal polynomial $\mu_i(x) \in K_i[x]$. Let $\varphi : K_0 \to K_1$ be an isomorphism of fields. Then φ induces an isomorphism of polynomial rings, which we also denote by φ. If $\varphi(\mu_0(x)) = \mu_1(x)$ then there exists a unique field isomorphism $K_0[\zeta_0] \to K_1[\zeta_1]$ extending φ and sending ζ_0 to ζ_1.*

Proof. Since $\varphi : K_0[x] \to K_1[x]$ is an isomorphism sending $\mu_0(x)$ to $\mu_1(x)$, we have induced isomorphisms

$$K_0[\zeta_0] = K_0[x]/\langle \mu_0(x) \rangle \cong K_1[x]/\langle \mu_1(x) \rangle = K_1[\zeta_1].$$

Uniqueness follows as $K_0[\zeta_0]$ is generated by ζ_0 as an K_0-algebra. **q.e.d.**

Proposition 4.1.3.4. *Let $p(x) \in K[x]$ a polynomial with coefficients in the field K, and let K_p/K be a splitting field for p. Let M/K be a field extension such that p splits into linear factors over M. Then there is an K-homomorphism $K_p \to M$.*

Proof. We induct on $[K_p : K]$. The start $[K_p : K] = 1$ is trivial as $K_p = K$.

So assume $K_p \neq K$. Then p has a root $\alpha \in K_p - K$. Since the minimal polynomial μ_α divides p, there are linear factors $x - \zeta_i \in M[x]$ such that $\mu_\alpha = (x - \zeta_1) \cdots (x - \zeta_u)$ and μ_α is the minimal polynomial of each of the ζ_i over K. Thus, there is an K-isomorphism from $K(\alpha) \leq K_p$ to some subfield of M. Since $[K_p : K(\alpha)] < [K_p : K]$ we can apply induction replacing K with $K(\alpha)$ and regarding M as an extension of $K(\alpha)$ via the constructed K-homomorphism. **q.e.d.**

Corollary 4.1.3.5 (Uniqueness of Splitting Fields). *Let $p(x) \in K[x]$ a polynomial with coefficients in the field K. Any two splitting fields K_0 and K_1 for p are K-isomorphic.* **q.e.d.**

Proposition 4.1.3.6 (Vieta, Existence of Splitting Fields). *Let K be a field and fix a monic polynomial $p(x) = a_0 + a_1 x + \cdots + a_{m-1}x^{m-1} + x^m \in K[x]$. Consider the K-algebra*

$$R_p := K[y_1, \ldots, y_m] / I$$

where I is the ideal modelling the relations

$$\begin{aligned}
\sigma_1(y_1, \ldots, y_m) = y_1 + \cdots + y_m &= -a_{m-1} \\
\sigma_2(y_1, \ldots, y_m) = \sum_{i<j} y_i y_j &= a_{m-2} \\
\sigma_3(y_1, \ldots, y_m) = \sum_{i<j<k} y_i y_j y_k &= -a_{m-3} \\
&\vdots \\
\sigma_m(y_1, \ldots, y_m) = y_i y_j \cdots y_m &= (-1)^m a_0
\end{aligned}$$

Then, $p(x)$ splits in R as

$$p(x) = (x - y_1)(x - y_2) \cdots (x - y_m).$$

Moreover, if M is a maximal ideal containing I, the field $K_p := K[y_1, \ldots, y_m]/M$ is a splitting field for $p(x)$.

Proof. Straight forward computation. **q.e.d.**

Corollary 4.1.3.7. *Every polynomial has a splitting field and such field is unique up to isomorphism. It has finite degree over the base field.*

4.1.4 The Algebraic Closure

Proposition and Definition 4.1.4.1. *A field K is called algebraically closed if it satisfies the following equivalent conditions:*

1. *The field K does not admit proper finite extensions.*
2. *The field K does not admit any proper algebraic extension.*
3. *All irreducible polynomials in $K[x]$ have degree 1.*
4. *Every polynomial in $K[x]$ splits into linear factors.*

Proof. We show that adjacent conditions are equivalent.

Suppose there is a proper finite extension. Then there is a proper algebraic extension. Suppose there is a proper algebraic extension. Since an algebraic extensions is the union of the intermediate finite extensions, there is a proper finite extension.

Suppose there is a proper algebraic extension M/K. Then there is an algebraic element $\zeta \in M$ not in K. The minimal polynomial of ζ is not linear and irreducible. Conversely suppose there is a non-linear irreducible polynomial. Then this polynomial generates a maximal ideal in $K[x]$. Quotienting out the maximal ideal defines a proper algebraic extension of K.

Suppose all irreducible polynomials are linear. Since every polynomial splits into irreducibles, every polynomial splits into linear factors. Suppose every polynomial splits into linear factors. It is immediate that irreducibles must be linear. **q.e.d.**

Lemma 4.1.4.2. *Let K/F be an algebraic field extension. Then the following are equivalent:*

1. *The field K is algebraically closed.*
2. *All (irreducible) polynomials in $F[x]$ split into linear factors in $K[x]$.*

Proof. Suppose there is a proper algebraic extension M/K. Then there is an algebraic element $\zeta \in M$ not in K. The minimal polynomial $\mu_F\zeta$ is irreducible over F and does not split into linear factors over K as one of its irreducible factors over K is the minimial polynomial $\mu_K\zeta$.

Conversely, suppose there is an irreducible polynomial over F that does not split into linear factors over K. Then, we can construct an algebraic element over K that has this minimal polynomial over F. **q.e.d.**

Theorem and Definition 4.1.4.3. *Every field F admits an algebraically closed algebraic extension. Such extension is unique up to F-isomorphism. It is called an algebraic closure of F.*

Proof. Existence: well order the polynomials in $F[x]$ and form the transfinite limit of the associate sequence of splitting fields.

Let M_1/F and M_2/F be two algebraic closures. Consider the partially ordered set of pairs

$$\{(K, \varphi : K \to M_2,) |\ M_1/K/F \text{ is a tower and } \varphi \text{ is a } F\text{-hom}\}$$

Note that it satisfies the hypotheses of Zorn's lemma. Note that (4.1.3.3) implies that any maximal element has first coordinate M_1. Thus, there is a F-homomorphism $\varphi : M_1 \to M_2$. It has to be onto since its image will be an algebraic closure for F inside M_2. **q.e.d.**

4.2 Galois Theory

4.2.1 The Galois Group

Definition 4.2.1.1. Let M/F be a field extension. The set of F-automorphisms of M is a group with respect to composition. It is called the Galois group of the extension and denoted by $\mathrm{Aut}_F(M)$. Let $M/K/F$ be a tower of fields. The relative Galois group is defined as $\mathrm{Aut}_F(M/K) := \{\varphi \in \mathrm{Aut}_F(M) \mid \varphi(K) = K\}$.

Observation 4.2.1.2. *Restriction defines a group homomorphism*

$$\mathrm{Aut}_F(M/K) \longrightarrow \mathrm{Aut}_F(K)$$

with kernel $\mathrm{Aut}_K(M)$.

Exercise 4.2.1.3. Prove or disprove: Let $M/K/F$ be a tower of fields with M/F algebraic. Then

$$\mathrm{Aut}_K(M) \hookrightarrow \mathrm{Aut}_F(M/K) \twoheadrightarrow \mathrm{Aut}_F(K)$$

is a short exact sequence of groups.

Observation 4.2.1.4. *Let M/F be a field extension. For any F-automorphism $\varphi \in \mathrm{Aut}_F(M)$ the fix point set $\mathrm{Fix}(\varphi) := \{\zeta \in M \mid \varphi(\zeta) = \zeta\}$ is closed with respect to arithmetic operation and hence a subfield of M containing F. Consequently, for any subgroup $G \leq \mathrm{Aut}_F(M)$, the fix point set $\mathrm{Fix}(G) := \{\zeta \in M \mid \varphi(\zeta) = \zeta$ for all $\varphi \in G\}$ is an intermediate field.* **q.e.d.**

Observation 4.2.1.5. *Consider a simple algebraic extension $K(\alpha)/K$ with minimal polynomial $\mu_\alpha(x)$. Since any K-automorphism of $K(\alpha)$ is determined by where it sends α and since it has to send α to some root of $\mu_\alpha(x)$, we find that* $\mathrm{card}(\mathrm{Aut}_K(K(\alpha))) \leq \deg(\mu_\alpha) = [K(\alpha) : K]$.

We can strengthen this result by recalling (4.1.3.3). It follows that $\mathrm{Aut}_K(K(\alpha))$ acts simply transitively on the set of roots of $\mu_\alpha(x)$ in $K(\alpha)$. **q.e.d.**

Theorem 4.2.1.6. *Let M/F be a field extension and fix a subgroup $G \leq \mathrm{Aut}_F(M)$ and put $K := \mathrm{Fix}(G)$. If M/K is finite, then $\mathrm{card}(G) \leq [M : K]$.*

In particular, for any finite field extension M/F, we have $\mathrm{card}(\mathrm{Aut}_F(M)) \leq [M : F]$.

Proof. We induct on $[M : K]$. If $M = K$, there is nothing to prove. Otherwise, consider $\alpha \in M$, the simple extension $K(\alpha)/K$, and the subgroup $H := G \cap \mathrm{Aut}_{K(\alpha)}(M) \leq G$. By induction, $\mathrm{card}(H) \leq [M : K(\alpha)]$.

On the other hand, cosets in G/H correspond to restrictions $\varphi|_{K(\alpha)}$. Since such a restriction is determined by the image of α and field homomorphisms have to send α to another root of its minimal polynomial, we find $\mathrm{card}(G/H) \leq [K(\alpha) : K]$. Now the claim follows from (4.1.1.6). **q.e.d.**

Definition 4.2.1.7. An algebraic extension M/F is called a Galois extension if $F = \mathrm{Fix}(\mathrm{Aut}_F(M))$.

Example 4.2.1.8. Let M/F be a field extension. Put $K := \mathrm{Fix}(\mathrm{Aut}_F(M))$. Then M/K is a Galois extension.

Proof. We have to show that $K = \mathrm{Fix}(\mathrm{Aut}_K(M))$. This follows from $\mathrm{Aut}_K(M) = \mathrm{Aut}_F(M)$, which we just observe: every K-automorphism of M clearly fixes $F \leq K$, and any F-automorphism of M is a K-automorphism since it fixes $K = \mathrm{Fix}(\mathrm{Aut}_F(M))$. **q.e.d.**

Fundamental Observation 4.2.1.9. *Let M/F be an algebraic extension and fix $G \leq \mathrm{Aut}_F(M)$ with $F = \mathrm{Fix}(G)$. (In particular, M/F is Galois.)*

For any $\alpha \in M$, the minimal polynomial μ_α splits into pairwise different linear factors, and G acts transitively on the roots of μ_α. (In particular, so does $\mathrm{Aut}_F(M)$.)

Proof. Let $\mathcal{O}$ be the G-orbit of α. Then $p(x) := \prod_{\alpha' \in \mathcal{O}} x - \alpha' \in M[x]$ is a polynomial fixed by G. Since $F = \mathrm{Fix}(G)$ its coefficients lie in F. Thus μ_α divides p. On the other hand, each root of p also is a root of μ_α since the action of G fixes μ_α. Thus $\mu_\alpha = p$. **q.e.d.**

4.2.2 Normal Field Extensions

Definition 4.2.2.1. An algebraic extension M/K is normal if every irreducible polynomial $p \in K[x]$ that has a root in M splits completely into linear factors over M. Equivalently, any minimal polynomial $\mu_\alpha(x) \in K[x]$ for $\alpha \in M$ splits into linear factors over M.

Observation 4.2.2.2. *It follows immediately from (4.2.1.9) that every Galois extension is normal.* **q.e.d.**

Lemma 4.2.2.3. *Let $M/K/F$ be a tower of field so that M/F is normal. Then, any F-homomorphism $\varphi : K \hookrightarrow M$ extends to an F-automorphism of M.*

Proof. Consider the set of extensions

$$\{\psi : K \hookrightarrow M \mid F \leq K \leq M,\ \psi \text{ extends } \varphi\}$$

ordered by restriction. By Zorn's lemma, there is a maximal element $\psi : K_{\max} \hookrightarrow M$. We claim that in this case, $K_{\max} = M$ and that ψ is onto.

Assume that $K_{\max}$ is a proper subfield of M. Fix $\xi \in M - K_{\max}$ and let $\mu_{K_{\max}}$ be its minimal polynomial over $K_{\max}$ and let μ_F be its minimal polynomial over F. Since $\mu_{K_{\max}}$ divides μ_F, we infer that $\psi(\mu_{K_{\max}})$ divides $\psi(\mu_F) = \mu_F$. Hence $\psi(\mu_{K_{\max}})$ has a root $\zeta \in M$ of which it is the minimal polynomial over $\psi(K_{\max})$. By (4.1.3.3), the field-isomorphism $\psi : K_{\max} \to \psi(K_{\max})$ extends to an isomorphism $K(\xi) \to \psi(K_{\max})(\zeta)$. This contradicts maximality.

Assume that ψ is not onto. Since ψ is an F-homomorphism, every irreducible polynomial that splits over M already splits over $\psi(M)$. This implies that elements in $M - \psi(M)$ cannot be algebraic over F. **q.e.d.**

Corollary 4.2.2.4. *Let M/F be a normal extension. For each irreducible polynomial $p \in F[x]$, the Galois group $\mathrm{Aut}_F(M)$ acts transitively on the set $\{\alpha \in M \mid p(\alpha) = 0\}$ of M-roots of p.*

Proof. If p has no roots in M, there is nothing to prove.

Let $p \in F[x]$ be an irreducible polynomial that has one and hence all its roots in M. Let α and α' be two roots of p. By (4.1.3.3), there exists an F-isomorphism $\varphi : F(\alpha) \to F(\alpha')$. By (4.2.2.3), the homomorphism φ extends to an F-automorphism of M. **q.e.d.**

Theorem 4.2.2.5 (Characterization of Normal Extensions). *An algebraic extension K/F is normal if and only if, for each algebraic extension M/K, every F-automorphism $\varphi : M \to M$ stabilizes K, i.e., satisfies $\varphi(K) = K$. In fact, if K/F is not normal, for M/K for which M/F is normal admits an automorphism that moves K off itself.*

Proof. First assume that K/F is normal. Consider $\alpha \in K$ with minimal polynomial μ_α. Since φ fixed the μ_α, it has to send α to a root of μ_α, when $\varphi(\alpha) \in K$. The inclusion $\varphi(K) \subseteq K$ follows. The reverse inclusion follows since the same argument can be applied to the inverse automorphism. φ^{-1}.

Conversely, assume that K/F is not normal. Let α be an element of K whose minimal polynomial μ_α does not split into linear factors in K. Let M/K be any extension normal over F, e.g., the algebraic closure will do. By (4.2.2.4), the Galois group $\mathrm{Aut}_F(M)$ acts transitively on the roots of μ_α, there is an F-automorphism of M moving α out of K. **q.e.d.**

Corollary 4.2.2.6. *Let $M/K/F$ be a tower of fields so that M/F is normal. Then, M/K is normal.* **q.e.d.**

Remark 4.2.2.7. Here is a direct proof of the corollary: An irreducible polynomial $p \in K[x]$ that has a root α in M is the minimal polynomial over K of this root. Consider the minimal polynomial μ of α over F. We know that p divides μ. Since M/F is normal, μ splits into linear factors over M and hence so does the

divisor p. Note, how this argument already appeared in the proof of (4.2.2.4).

Corollary 4.2.2.8 (Characterization of Finite Normal Extensions). *A finite field extension K/F is normal if and only if K is the splitting field for some polynomial $p \in K[x]$.*

Proof. Splitting fields are stabilized by all ambient F-automorphisms since they fix the underlying polynomial and can only permute its root. Thus, splitting fields are normal.

Conversely, a finite normal extension is finitely generated and the splitting field of the product of the minimal polynomials of its generators. **q.e.d.**

4.2.3 Separable Field Extensions

Definition 4.2.3.1. An irreducible polynomial $p(x) \in K[x]$ is separable if it does not have multiple roots in any field extension of K.

An algebraic extension M/K is separable if all elements of M have separable minimal polynomials over K. (Recall that minimal polynomials are always irreducible.)

Observation 4.2.3.2. *It follows immediately from (4.2.1.9) that every Galois extension is separable.* **q.e.d.**

Observation 4.2.3.3. *Let $M/K/F$ be a tower of fields with M/F separable. Then M/K is separable.*

Proof. Minimal polynomials over K divide minimal polynomials over F. **q.e.d.**

4.2.4 Characterizations of Galois Extensions

Theorem 4.2.4.1. *A field extension M/F is a Galois extension, i.e., $\mathrm{Fix}(\mathrm{Aut}_F(M)) = F$, if and only if it is normal and separable.*

Proof. We have already seen that Galois extensions are normal (4.2.2.2) and separable (4.2.3.2).

Put $K := \mathrm{Fix}(\mathrm{Aut}_F(M))$ and assume that M/F is normal and separable. Also assume, for contradiction, that $K \neq F$. Then there is an element $\alpha \in K - F$. Since the minimal polynomial μ_α splits into linear factors over M, we can find a splitting field F_μ for μ_α inside M and within this splitting field, we can move α to another root (the roots are distinct because of separability). Thus, there is an element in $\mathrm{Aut}_F(F_\mu)$ not fixing α. This automorphism can be extended to all of M. Contradiction. **q.e.d.**

Corollary 4.2.4.2. *Let $M/K/F$ be a tower of fields with M/F Galois. Then M/K is Galois, i.e.,*

$$\mathrm{Fix}(\mathrm{Aut}_K(M)) = K.$$

Proof. The preceeding characterization together with (4.2.2.6) and (4.2.3.3) does the trick. **q.e.d.**

Proposition 4.2.4.3. *Let M/F be a Galois extension. Then the following are equivalent:*

1. *M/F is finite.*
2. *$\mathrm{Aut}_F(M)$ is finite.*
3. *M/F is simple.*

In particular, every finite Galois extension is simple.

Proof. We have argued (1) $\Longrightarrow$ (2) in (4.2.1.6). The implication (3) $\Longrightarrow$ (1) follows since a finitely generated algebraic extension is finite.

Now we argue (2) $\Longrightarrow$ (3) By (4.1.2.10), it suffices to show that M/F admits only finitely many intermediate fields. By (4.2.4.2), such intermediate fields K are uniquely determined by the subgroup $\mathrm{Aut}_K(M) \leq \mathrm{Aut}_F(M)$. Since the finite group $\mathrm{Aut}_F(M)$ has only finitely many subgroups, the claim follows. **q.e.d.**

Corollary 4.2.4.4. *Let M/F be a finite Galois extension. Then M is the splitting field of a separable polynomial in $F[x]$ and* $\mathrm{card}(\mathrm{Aut}_F(M)) = [M : F]$.

Proof. We have argued $\mathrm{card}(\mathrm{Aut}_F(M)) \leq [M : F]$ in (4.2.1.5).

By (4.2.4.3), $M = F(\alpha)$ for some $\alpha \in M$. Since M/F is normal, and α is a root of μ_α, we find that μ_α splits into linear factors, which are pairwise distinct since α is separable over F. By (4.2.1.9), the Galois group acts transitively on the set of roots of μ_α. This implies $\mathrm{card}(\mathrm{Aut}_F(M)) \geq [M : F]$. Also, M is the splitting field for μ_α since no proper subfield does contain α. **q.e.d.**

Corollary 4.2.4.5. *Let M/F be a finite field extension and put $K := \mathrm{Fix}(\mathrm{Aut}_F(M))$. Then $[M : F] = \mathrm{card}(\mathrm{Aut}_F(M))\,[K : F]$. In particular, if M/F is not Galois, then* $\mathrm{card}(\mathrm{Aut}_F(M)) \neq [M : F]$.

Proof. M/K is Galois and $\mathrm{Aut}_F(M) = \mathrm{Aut}_K(M)$. **q.e.d.**

Corollary 4.2.4.6. *Let M/F be a finite Galois extension and let G be a subgroup of the Galois group $\mathrm{Aut}_F(M)$. Put $K := \mathrm{Fix}(G)$. Then* $G = \mathrm{Aut}_K(M)$.

Proof. Since M/K is Galois, it is simple. Let α be a generator. By (4.2.1.9), the minimal polynomial splits into different linear factors and G acts transitively on these. Since a proper subgroup of $\mathrm{Aut}_K(M)$ has not enough elements to do this, the claim follows. **q.e.d.**

Theorem 4.2.4.7. *For a finite extension M/F, the following are equivalent:*

1. *The extension is a Galois extension.*
2. $[M : F] = \mathrm{card}(\mathrm{Aut}_F(M))$.
3. *M is the splitting field for a separable polynomial in $F[x]$.*

4. *M is the splitting field for a polynomial in $F[x]$ whose irreducible factors are separable.*

Proof. It suffices to show that the last condition implies the first. So let $p(x)$ be a polynomial with separable irreducible factors. If all roots of $p(x)$ are in F, there is nothing to prove. Otherwise, let α be a root in $M - F$. Let $K := F(\alpha)$. Then, the irreducible factors of $p(x)$ in $K[x]$ divide the factors in $F[x]$. We conclude the no multiple roots crop up. Thus, $p(x)$ has separable irreducible factors in $K[x]$.

Now, we are in a position to use induction since M/K has smaller degree than M/F. Thus, M/K is Galois.

Let $\mu(x)$ be the minimal polynomial of α. Observe that $\mu(x)$ divides $p(x)$. Hence, $\mu(x)$ splits over M. Let $Z := \{\alpha' \in M \mid \mu(\alpha') = 0\}$. Since $\mu(x)$ is an irreducible factor of $p(x)$, it is separable and $\mathrm{card}(Z) = \deg(\mu)$. Since M is normal by (4.2.2.8), the Galois group $\mathrm{Aut}_F(M)$ acts transitively on Z by (4.2.2.4).

Put $L := \mathrm{Fix}(\mathrm{Aut}_F(M))$, and let $q(x)$ be the minimal polynomial of α over L. By (4.2.1.9), the $\mathrm{Aut}_L(M)$-orbit of α consists of precisely of the roots of q in M. Since $\mathrm{Aut}_L(M) = \mathrm{Aut}_F(M)$ we have that this orbit it Z. In particular, $\deg(q) \geq \mathrm{card}(Z) = \deg(\mu)$. Hence $[L(\alpha) : L] \geq [F(\alpha) : F]$. Since M is Galois over $F(\alpha)$, we have $F \leq L \leq F(\alpha)$ whence $F(\alpha) = L(\alpha)$. Hence, our degree inequality implies $L = F$, i.e., M/F is Galois. **q.e.d.**

4.2.5 Galois Correspondence

Main Theorem of Galois Theory 4.2.5.1. *Let M/F be a finite Galois extension. The map*

$$\begin{array}{ccc} \{K \mid M/K/F\} & \longrightarrow & \{G \mid G \leq \mathrm{Aut}_F(M)\} \\ K & \mapsto & \mathrm{Aut}_K(M) \end{array}$$

is an inclusion-reversing bijection whose inverse is given by

$$G \mapsto \mathrm{Fix}(G)\,.$$

Moreover, K/F is normal if and only if $\mathrm{Aut}_K(M)$ is normal in $\mathrm{Aut}_F(M)$.

Proof. That $K \mapsto \mathrm{Aut}_K(M)$ and $G \mapsto \mathrm{Fix}(G)$ are inverse operations is the contents of (4.2.4.2) and (4.2.4.4). It is easy to see that they reverse inclusions.

It remains to show that the correspondence preserves normality. This follows from (4.2.2.5). **q.e.d.**

4.2.6 Finite Fields

Theorem 4.2.6.1. *There exists up to isomorphism a unique finite field $\mathbb{F}_{p^k}$ of size p^k for each prime number p and each exponent $k \geq 1$. Such finite field is the splitting field of the polynomial $x^{p^k} - x$ whose roots are exactly the elements of $\mathbb{F}_{p^k}$.*

Proof. Uniqueness: let $K/\mathbb{F}_p$ be a field extension of degree k. The multiplicative group of K has order $p^k - 1$. Hence the elements of K are exactly the roots of $x^{p^k} - x$ and K is the splitting field of this polynomial.

Existence: it suffices to show that $\mathbb{F}_p[x]$ has irreducible polynomials of every degree. This is left as an exercise. **q.e.d.**

Exercise and Definition 4.2.6.2. Let K be a field of characteristic $p > 0$. Then, the Frobenius homomorphism

$$\begin{array}{ccc} K & \longrightarrow & K \\ \xi & \mapsto & \xi^p \end{array}$$

is a field homomorphism.

Theorem 4.2.6.3. *Let $q := p^k$. The Frobenius automorphism $\varphi : \mathbb{F}_q \to \mathbb{F}_q$ is an element of $\mathrm{Aut}_{\mathbb{F}_p}(\mathbb{F}_q)$ of order k. Consequently, $\mathrm{Aut}_{\mathbb{F}_p}(\mathbb{F}_q)$ is cyclic of order k, and $\mathbb{F}_q/\mathbb{F}_p$ is a Galois extension.*

Proof. Let l be the order of the Frobenius automorphism. Then $x^{p^l} - x$ vanishes identically on $\mathbb{F}_q$. Since it cannot have more than p^l roots, we find $l \geq k$. **q.e.d.**

Corollary 4.2.6.4. *The field $\mathbb{F}_{p^k}$ contains exactly one subfield of size p^l for every l dividing k. Such subfield consists precisely of the roots of the polynomial $x^{p^l} - x$. In particular, the subfields of any finite field are totally ordered with respect to inclusion.* **q.e.d.**

4.3 Separability

4.3.1 Perfect Fields

Definition 4.3.1.1. A field K is perfect if every algebraic extension of K is separable. Equivalently, K is perfect if every irreducible polynomial in $K[x]$ is separable.

Exercise 4.3.1.2. Show that for any irreducible polynomial $p(x) \in K[x]$, the following are equivalent:

1. The polynomial $p(x)$ is separable.
2. The polynomial $p(x)$ does not share roots with its derivative $p^{'}(x)$.
3. The polynomial $p(x)$ and its derivative $p^{'}(x)$ are relatively prime, i.e.,
$$\gcd\Big(p(x)\,, p^{'}(x)\Big) = 1.$$

Corollary 4.3.1.3. *In characteristic 0, every irreducible polynomial is separable. Consequently all fields of characteristic 0 are perfect.* **q.e.d.**

Corollary 4.3.1.4. *In characteristic p, an irreducible polynomial $p(x)$ is not separable if and only if it can be regarded as a polynomial in x^p, i.e., if there exists a polynomial $q(x)$ with $p(x) = q(x^p)$.* **q.e.d.**

Exercise 4.3.1.5. Let K be a field of prime characteristic p. If $\xi \in K$ is not a pth power, then the polynomial $x^p - \xi$ is irreducible. In particular, if the Frobenius endomorphism for K is not onto, then K is not perfect.

Exercise 4.3.1.6. Let K be a field of prime characteristic p. Every polynomial of the form $p(x^p)$ is a pth power of another polynomial. In particular, no polynomial of the form $p(x^p)$ is irreducible. It follows that a K is perfect.

From (4.3.1.5) and (4.3.1.6) we infer:

Corollary 4.3.1.7. *A field of positive characteristic is perfect, if and only if the Frobenius homomorphism is onto.* **q.e.d.**

Corollary 4.3.1.8. *Finite fields are perfect.* **q.e.d.**

Exercise 4.3.1.9. Construct a non-perfect field.

4.3.2 xx

Theorem 4.3.2.1. *Let K/F be a field extension.*

1. *If $\xi_1, \xi_2, \ldots, \xi_u \in K$ are all separable over F, then $F(\xi_1, \xi_2, \ldots, \xi_u)$ is a separable extension of F. This says: rational expressions in separable arguments are separable.*

2. *The set*
$$F_{\text{sep}} := \{\xi \in K \mid \xi \text{ is separable over } F\}$$
is a subfield of K. It is called the separable closure of F in K.

Proof. Clearly, the first statement implies the second. To see that $F(\xi_1, \xi_2, \ldots, \xi_u)/F$ is a separable extension, consider the minimal polynomials μ_{ξ_i}. All of these are separable. Hence, the splitting field M of their product is a Galois extension of F. Clearly, $F(\xi_1, \xi_2, \ldots, \xi_u) \subset K \cap M$ whence all elements of $F(\xi_1, \xi_2, \ldots, \xi_u)$ are separable over F. **q.e.d.**

4.3.3 Primitive Elements

Theorem 4.3.3.1 (The Primitive Element). *Let K/F be an algebraic field extension and let $\xi, \xi_1, \xi_2, \ldots, \xi_u \in K$ be chosen so that ξ_i is separable over F for all i. Then there is an element $\beta \in K$ (called a primitive element) such that*

$$F(\xi, \xi_1, \xi_2, \ldots, \xi_u) = F(\beta)$$

In particular, every finite separable extension is simple.

Proof. Induction easily reduces us to the case $u = 1$. Also, we may assume that F is infinite as finite fields always have primitive elements. ... **q.e.d.**

4.4 Determinants

4.4.1 Norms

Let us fix a finite field extension K/F of degree m. Then, K is a vector space of finite dimension over F. For every element $\xi \in K$, left-multiplication induces a F-vector space endomorphism $\lambda_\xi : K \to K$. If we fix a basis B of K over F, we can represent each element ξ by the matrix representing λ_ξ. This way, we can realize the field K as a subfield in the matrix ring $\mathbb{M}_m(F)$.

Observation 4.4.1.1. *Note that $\mathrm{norm}(\xi) = 0$ if and only if $K = 0$. Also,*

$$\mathrm{norm}(\xi\xi) = \mathrm{norm}(\xi)\,\mathrm{norm}(\xi)$$ **q.e.d.**

Observation 4.4.1.2. *The norm map is $\mathrm{Aut}_F(K)$-invariant.*

Proof. Let B be a basis of K as a F vector space and let $\varphi : K \to K$ be an F-automorphism of K. Then, $\varphi(B)$ is another basis for K. Let $\xi \in K$ and consider the matrix representing λ_ξ relative to B. It is the same matrix that represents $\lambda_{\varphi(\zeta)}$ relative to $\varphi(B)$. Consequently, $\mathrm{norm}(\xi) = \mathrm{norm}(\varphi(K))$. **q.e.d.**

Corollary 4.4.1.3. *The set* $\{\xi \in K \mid \mathrm{norm}(\xi) = 1\}$ *is a subgroup of* K^* *on which* $\mathrm{Aut}_F(K)$ *acts by group automorphisms.* **q.e.d.**

4.4.2 Normal Bases

Exercise 4.4.2.1. Let K/F be a field extension. Let G be a subgroup of $\mathrm{Aut}_F(K)$ and put $F := \mathrm{Fix}(G)$. Show that there is an element $\xi \in K$ such that the familiy $(\varphi(\xi))_{\varphi \in G}$ is linearly independent over F.

4.5 Examples

4.5.1 Symmetric Functions

Let F be a field. Let $M := K(\zeta_1, \zeta_2, \ldots, \zeta_u)$ denote the rational function field in u variables over K. Observe that the symmetric group $\mathbf{S}_u$ acts on M by F-automorphisms defined by permuting the variables.

Definition 4.5.1.1. A function $f \in K(\zeta_1, \zeta_2, \ldots, \zeta_u)$ is called symmetric if it is fixed under all permutations of the variables.

Observation 4.5.1.2. *The elementary symmetric functions*

$$\begin{aligned}
\sigma_1 &:= \zeta_1 + \cdots + \zeta_u \\
\sigma_2 &:= \sum_{i<j} \zeta_i \zeta_j \\
\sigma_3 &:= \sum_{i<j<k} \zeta_i \zeta_j \zeta_k \\
&\vdots \\
\sigma_u &:= \zeta_i \zeta_j \cdots \zeta_u
\end{aligned}$$

are symmetric functions. **q.e.d.**

Let $K := F(\sigma_1, \sigma_2, \ldots, \sigma_u)$ be the subfield of M generated by the elementary symmetric functions. We consider the polynomial

$$p(x) := (x - \zeta_1)(x - \zeta_2) \cdots (x - \zeta_u) = x^u - \sigma_1 x^{u-1} + \sigma_2 x^{u-2} - \cdots \pm \sigma_u \in K[x]$$

whose separability is obvious from the factorization.

Observation 4.5.1.3. *M is the splitting field for p over K. In particular, M/K is Galois. Also, since any K-automorphism is uniquely determined by its induced permutation on the roots of p, we see that $\mathrm{Aut}_K(M)$ embeds into the symmetric group $\mathbf{S}_u$.* **q.e.d.**

Observation 4.5.1.4. *The symmetric group $\mathbf{S}_u$ embeds into the Galois group $\mathrm{Aut}_K(M)$ as the group of those automorphisms defined by permuting the variables ζ_i. (Note that such F-automorphisms leave all elementary symmetric functions fixed and therefore fix K elementwise.)* **q.e.d.**

Thus, we have shown:

Proposition 4.5.1.5. *The extension M/K is Galois with Galois group $\mathbf{S}_u$ permuting the variables ζ_i.*

In particular, $K = \mathrm{Fix}(\mathbf{S}_u)$, i.e., any symmetric function is a rational expression of elementary symmetric functions.

4.5.2 The General Polynomial

Let F be a field and let $K := F(a_1, a_2, \ldots, a_u)$ be the rational function field. We consider the polynomial:

$$p(x) := x^u - a_1 x^{u-1} + a_2 x^{u-2} - \cdots \pm a_u \in K[x]\,.$$

This is the general polynomial of degree u over F. Note that it is not a polynomial in $F[x]$. We aim to show:

Proposition 4.5.2.1. *Let M be the splitting field of the general polynomial $p(x)$ as defined above. Then M/K is a Galois extension with Galois group $\mathrm{Aut}_K(M) = \mathbf{S}_u$.*

Proof. We will compare the situation to the previous example. Let $M' := F(\zeta_1, \zeta_2, \ldots, \zeta_u)$ be a rational function field on a different set of indeterminates, and put $K' := F(\sigma_1, \sigma_2, \ldots, \sigma_u)$ where the σ_i are the

elementary symmetric functions in the new indeterminates ζ_j. Put $q(x) := (x-\zeta_1)(x-\zeta_2)\cdots(x-\zeta_u)$.

Consider the evaluation homomorphism

$$\begin{array}{rcl} \varphi : F[a_1, a_2, \ldots, a_u] & \longrightarrow & F[\sigma_1, \sigma_2, \ldots, \sigma_u] \\ \zeta_i & \mapsto & \sigma_i \end{array}$$

We claim that φ is an isomorphism. Note that φ is onto since the σ_i generate the right hand. To see that φ is 1-1, assume that $r(a_1, a_2, \ldots, a_u) \in \ker(\varphi)$, i.e., $r(\sigma_1, \sigma_2, \ldots, \sigma_u) = 0$. Let $\alpha_1, \ldots, \alpha_u \in M$ be the roots of $p(x)$. Recall that the σ_i really are polynomials in the ζ_j. Thus, further evaluating at $\zeta_i = \alpha_i$, we find:

$$r(\alpha_1 + \cdots + \alpha_u, \ldots, \alpha_1 \cdots \alpha_u) = 0$$

Remebering that $a_1 = \alpha_1 + \cdots + \alpha_u, \ldots, \ a_u = \alpha_1 \cdots \alpha_u$. Consequently,

$$r(a_1, a_2, \ldots, a_u) = 0,$$

which proves injectivity.

Since K is the field of fractions for $F[\zeta_1, \zeta_2, \ldots, \zeta_u]$ and K' is the field of fractions for $F[\sigma_1, \sigma_2, \ldots, \sigma_u]$, the homomorphism φ extends to a field isomorphism $\varphi : K \to K'$. Observe that $q = \varphi(p)$ whence (by uniqueness of splitting fields) we have $M \cong M'$ and $\mathrm{Aut}_K(M) \cong \mathrm{Aut}_{K'}(M') = \mathbf{S}_u$. **q.e.d.**

4.5.3 Roots of Unity

Let $\mathbb{Q}_r$ be the splitting field of $x^r - 1$ over $\mathbb{Q}$. Note that $\mathbb{Q}$ is perfect. Hence $\mathbb{Q}_r/\mathbb{Q}$ is Galois.

Note that the roots of $x^r - 1$ form a finite group under multiplication. Thus, this is a cyclic group. Any generator is called a primitive rth root of unity.

Definition 4.5.3.1. The polynomial $\Phi_m(x) := \prod_\omega x - \omega$ where ω ranges over all primitive roots of order m is called the mth cyclotomic polynomial.

Observation 4.5.3.2. *The mth cyclotomic polynomial has degree* $\deg(\Phi_m(x)) = \phi(m)$. **q.e.d.**

Observation 4.5.3.3. *For any r, we have* $x^r - 1 = \prod_{m|r} \Phi_m(x)$. **q.e.d.**

Corollary 4.5.3.4. *The cyclotomic polynomials are primitive integer polynomials.*

Proof. We use induction. Assume that for any $m < r$, the mth cyclotomic polynomial has integer coefficients.

Since cyclotomic polynomials are monic, we can perform division with remainder in $\mathbb{Z}[x]$. Thus, there exist unique polynomials $p(x), q(x) \in \mathbb{Z}[x]$ with:

$$x^r - 1 = p(x) \prod_{m|r, m<r} \Phi_m(x) + q(x) \qquad \text{and} \qquad \deg(q(x)) < \deg\left(\prod_{m|r, m<r} \Phi_m(x) \right)$$

Since the same constraints define the same polynomials in $\mathbb{C}[x]$, we find that $p(x) = \Phi_r(x)$ and $q(x) = 0$.

Since $x^r - 1$ is primitive, so must be all its integer factors. **q.e.d.**

Lemma 4.5.3.5 (Gauss). *The cyclotomic polynomials are irreducible over $\mathbb{Q}$.*

Proof. Since cyclotomic polynomials are primitive integer polynomials, it suffices to show that they don't factor over the integers.

Suppose $\Phi_m(x) = p(x)\, q(x)$ in $\mathbb{Z}[x]$, where we assume that $p(x)$ is monic and irreducible. Let α be a root of p and let p be a prime not dividing m. Note that α^p is another primitive root of unity of degree m. Hence $\Phi_m(\alpha^p) = 0$.

Assume $p(\alpha^p) \neq 0$. Then $q(\alpha^p) = 0$. Note that $p(x)$ is the minimal polynomial of α over $\mathbb{Q}$. It follows that $p(x) \,|q(x^p)$ over $\mathbb{Q}$. However, since $p(x)$ is monic, the same division with remainder trick as above yields $p(x) \,|q(x^p)$ in $\mathbb{Z}[x]$. Now, we reduce this mod p, i.e.,

we regard this as a statement in $\mathbb{Z}_p[x]$. Since $q(x^p) = q(x)^p$, we find that any root of $p(x)$ is also a root of $q(x)$. This, however, is a contradiction since $x^m - 1$ has no multiple roots over $\mathbb{Z}_p$ unless $p|m$. It follows that $p(\alpha^p) = 0$.

From this, we deduce that all primitive roots of degree m are roots of $p(x)$. Thus, by degree, $\Phi_m(x) = p(x)$. **q.e.d.**

Corollary 4.5.3.6. *The extension $\mathbb{Q}_r/\mathbb{Q}$ is Galois with Galois group $\mathrm{Aut}_{\mathbb{Q}}(\mathbb{Q}_r) = \mathrm{Aut}(\mathbf{C}_r)$.*

Proof. Any element has to permute the primitive roots. Also, $\mathbb{Q}_r = \mathbb{Q}(\omega)$, whence we have a monomorphism

$$\mathrm{Aut}_{\mathbb{Q}}(\mathbb{Q}_r) \hookrightarrow \mathrm{Aut}(\mathbf{C}_r)$$

which is onto by observing that both groups have $\phi(r)$ elements. **q.e.d.**

Corollary 4.5.3.7. $\mathrm{GL}_n(\mathbb{Q})$ *has only finitely many conjugacy classes of torsion elements.*

Proof. Let A be a matrix of order k. Then, all invariant factors divide $x^k - 1$ and are therefore products of cyclotomic polynomials $\Phi_m(x)$.

Note that $\phi(m) = \deg(\Phi_m)$ exceeds the number of primes less than r. It follows that only finitely many different cyclotomic polynomials can occur since the degrees of the invariant factors are bounded from above by n. It follows that only finitely many different polynomials can occur as invariant factors and the claim follows. **q.e.d.**

Chapter 5

Appendix: Sets

5.1 Zorn's Lemma and the Well-Ordering Theorem

5.1.1 Ordered Sets

Definition 5.1.1.1. A partially ordered set is a set X together with a relation $\preceq$ that satsifies:

$$x_1 \preceq x_2 \preceq x_3 \implies x_1 \preceq x_3$$
$$x_1 \preceq x_2 \preceq x_1 \iff x_1 = x_2$$

An element x is called minimal if $y \preceq x$ only holds for $y = x$. Similarly, x is called maximal if $x \preceq y$ only holds for $y = x$.

A subset A of a partially ordered set X is closed if it satisfies:

$$x \in A \quad \text{and} \quad y \preceq x \qquad \implies \qquad y \in A$$

If X and Y are partially ordered sets, we say that Y is an initial segment of X if Y is a closed subset of X and the partial order of Y is induced by the partial order on X. In this case, we write $Y \leq X$.

A partially ordered set is totally ordered if we always have:

$$x_1 \preceq x_2 \quad \text{or} \quad x_1 \preceq x_2$$

A totally ordered set is well-ordered if any non-empty subset has a minimal element.

Observation 5.1.1.2. *For any set of partially ordered sets, the relation "is an initial segment of" defines a partial order.* **q.e.d.**

Observation 5.1.1.3. *Let $\mathcal{F}$ be a $\leq$-nested collection of well-ordered sets, i.e., for any two $X, Y \in \mathcal{F}$, we have $X \leq Y$ or $Y \leq X$. Then there is a unique partial order on $\bigcup_{X \in \mathcal{F}} X$ that induces the order relations on each $X \in \mathcal{F}$, and this order is a well-ordering.* **q.e.d.**

5.1.2 The Theorems

Well-Ordering Theorem 5.1.2.1. *Any set X can be well-ordered.*

Zorn's Lemma 5.1.2.2. *A partially ordered set wherein each well-ordered subset has an upper bound has a maximal element.*

Lemma 5.1.2.3. *Let X be a set and let $\mathcal{P}_-$ be the collection of proper subsets of X. Let $\mathrm{ch} : \mathcal{P}_- \to X$ be a function satisfying $\mathrm{ch}(A) \in X - A$ for each proper subset $A \subset X$. Note that such a function always exists by the axiom of choice.*

Call a subset $A \subseteq X$ with a total order $\preceq$ a ch-set if

$$x = \mathrm{ch}(\{y \in A \mid y \preceq x \text{ and } y \neq x\})$$

for all $x \in A$.

Then, for any ch-set A properly contained in X, the set $A \cup \{\mathrm{ch}(A)\}$ is a ch-set containing A as a proper initial segment.

Moreover, the collection of all ch-sets is $\leq$-nested, i.e., for any two ch-sets A and B, we have $A \leq B$ or $B \leq A$.

Proof. Let $\mathcal{F}$ be the collection of those partially ordered sets that are initial segments as well of A as of B. Note that their union C is the unique maximal common initial segment of A and B.

We claim $C = A$ or $C = B$. Suppose neither equality obtains. Then $\mathrm{ch}(C) \in A \cap B$ and $C \cup \mathrm{ch}(C)$ would be a larger common initial segment for A and B, contradicting maximality. **q.e.d.**

Proof of the Well-Order Theorem. Let ch be a function as in Lemma 5.1.2.3, and let A be the union of all ch-sets. Since the ch-sets form a nested family, we know that A is a well-ordered. Observe that A is a ch-set. Indeed, it is maximal among all ch-sets.

Suppose A is a proper subset of X. Then, $A \cup \{\mathrm{ch}(A)\}$ is a larger ch-set. This contradiction implies that $A = X$. **q.e.d.**

Proof of Zorn's Lemma. Suppose that X does not have maximal elements. Then we can choose ch in Lemma 5.1.2.3 so that it assigns to each well-ordered $A \subset X$ an upper bound outside A.

The union Y, of all well-ordered ch-sets is a well-ordered ch-set. Indeed, it is the maximal one. Considering $Y \cup \{\mathrm{ch}(Y)\}$, we deduce that Y cannot be a proper subset of X, whence X is well-ordered and hence has an upper bound which must be a maximal element. **q.e.d.**

Appendix: Solutions to Selected Exercises

(1.1.6.16) The group $\mathrm{Perm}(X_5)$ acts naturally from the left on the set of all graphs with vertex set X_5: just move edges according to where their endpoints move. This action preserves the isomorphism type of a graph. In particular, the group $\mathrm{Perm}(X_5)$ acts on the cycles (graphs isomorphic to a pentagon-graph) with vertex set X_5. With respect to this action, the dihedral group $\mathbf{D}_{10}$ is the stabilizer of the pentagon-graph. Moreover, two elements of $\mathrm{Perm}(X_5)$ take the pentagon-graph to the same picture if and only if they belong to the same left-coset of $\mathbf{D}_{10}$. [You should argue this point a little. It isn't hard. Draw a picture. I allow myself to be brief here!] Thus, the pictures represent left-cosets.

In terms of Section 1.1.7 there is another way of seeing this: there is a natural $\mathrm{Perm}(X_5)$-equivariant bijection:

$$\{\text{cycles with vertex set } X_5\} \longleftrightarrow {}^{\mathrm{Perm}(X_5)}\Big/_{\mathbf{D}_{10}}$$

Both sets carry a left-action of $\mathrm{Perm}(X_5)$.

(1.1.7.21) To count the elements in G and H, we use the Orbit-Stabilizer Theorem. For G, note that this group acts transitively on the set of all twelve edges of the graph Γ. The stabilizer of a given edge preserves its end-points as one of them has degree 4 (degree=number of edges containing a given vertex) and the other

has degree 2. Since symmetries of Γ preserve degrees, an element stabilizing an edge cannot swap its end-points. It follows that the stabilizer of an edge has order 4: it consists of the identity, two independent swaps, and the simultaneous swapping. Thus the order of the group is $4 \times 12 = 48$.

For H consider the transitive action on the eight vertices of the cube. The stabilizer of any such vertex is a symmetric group on the three edges adjacent to the vertex. Thus, the order of H is $8 \times 6 = 48$.

It turns out that the two groups are isomorphic. We will interpret the graph Γ geometrically in the cube. The three degree 4 vertices of Γ corresponds to the three coordinate planes, and the six degree 2 vertices of Γ correspond to the six face midpoints of the cube. An edge in Γ is drawn if and only if the face midpoint lies within the coordinate plane. Here is a picture:

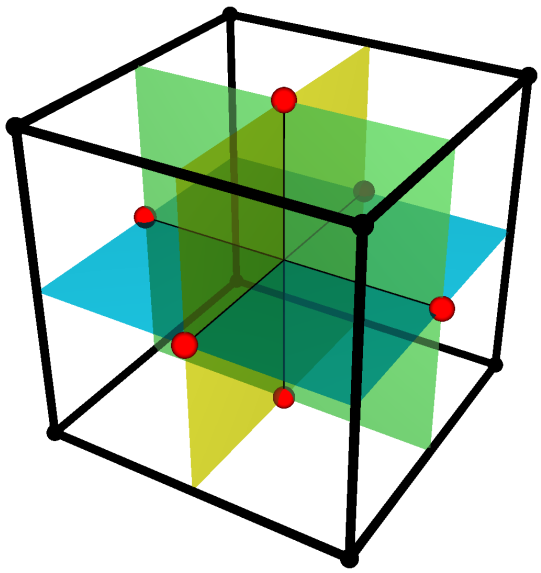

And here is the corresponding reading of Γ:

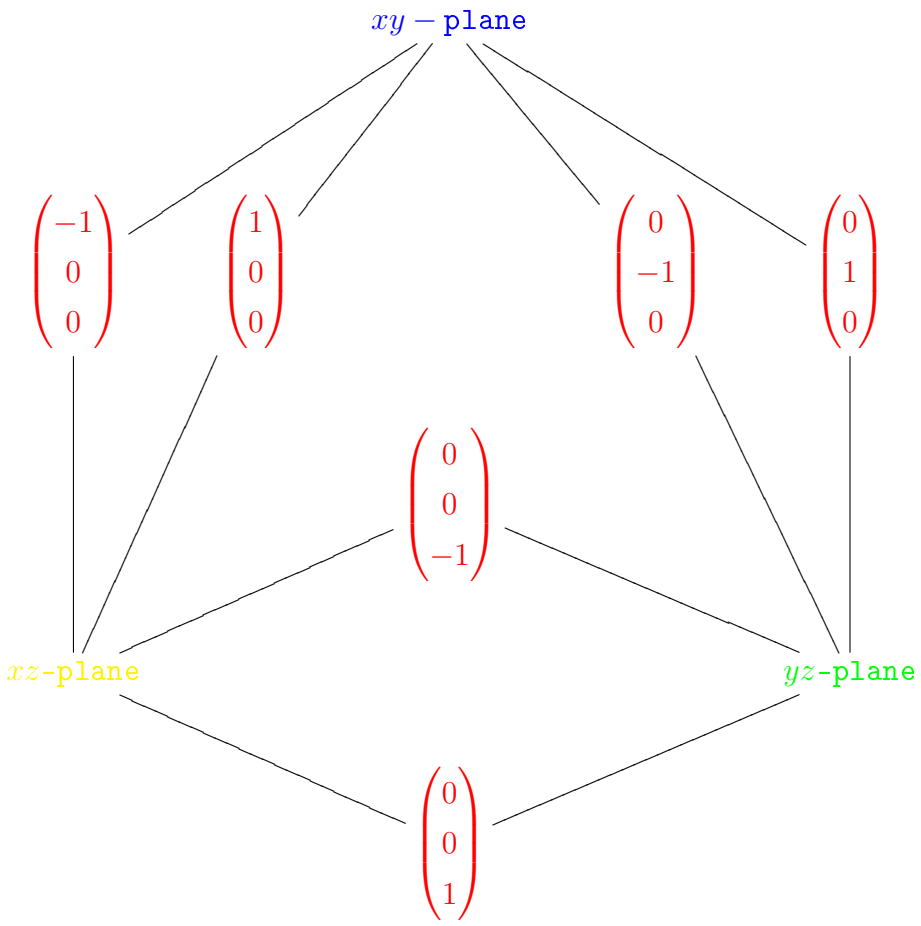

Every symmetry of the cube sends coordinate planes to coordinate planes and face midpoints to face midpoints. Also, if a point lies within a plane before applying the symmetry then its image will lie in the image of the plane afterward. Thus, the group H acts on Γ by graph automorphisms. Therefore, we obtain a group homomorphism

$$H \longrightarrow G$$

and all that remains to argue is injectivity: surjectivity will then follow since we establishes already that H and G have the same number of elements.

However, injectivity amounts to nothing more than the statement that the only cube symmetry that acts as the identity on Γ is,

indeed, the identity of the cube. Now, we see this by looking at the six face midpoints: any non-trivial symmetry of the cube has to move at least some of them. This movement, however, will move some of the degree 2 vertices in Γ.

(1.1.8.12) We use induction on n. The claim is obvious for $n = 1$. So let $A = (a_{ij})$ be a matrix in $\mathrm{GL}_n(\mathbb{Z})$. We devise the following strategy: form:

1. Multiply all rows with negative first-column entry by -1.
2. If there is just one row whose first-column entry is non-zero, then this entry is 1 as seen by developing $\det(A)$ along the first column. Swap, if necessary, to put this 1 in the top row. Now focus on the lower-right $(n-1) \times (n-1)$-submatrix. Its determinant has to be ± 1 by the same argument. Thus, we can apply the induction hypothesis. Since the necessary row-operations never involve the top-row, the leading zeros in the first column will be preserved. Thus, we can obtain a matrix of the form

$$\begin{pmatrix} 1 & * & * & \cdots & * \\ 0 & 1 & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & 1 & 0 \\ 0 & \cdots & 0 & 0 & 1 \end{pmatrix}$$

   And it is clear that we can kill the entries in the top-row.
3. If there is more than one non-zero entry in the first column, swap so that the smallest non-zero entry is in the first row. Now subtract the first row from all other rows as many times as possible without making the first-column entry negative. At the end of this step, all rows below the first will have non-negative first-column entries strictly smaller than the top-left entry.

   Now go back to the previous step.

The apparent loop in this algorithm is traversed only finitely many times since the top-left entry strictly decreases but stays positive in each iteration. This settles the first part.

Recall that elementary row operations can be realized as left-multiplication by elementary matrices. Let $\mathcal{E}$ be the set of those elementary $n \times n$-matrices that we need to realize the "very elementary row-operations" from the first part of the problem. Note that all such matrices lie in $\mathrm{GL}_n(\mathbb{Z})$ and that $\mathcal{E}$ is a finite set.

Now, we argue that $\mathrm{GL}_n(\mathbb{Z})$ is a group. Clearly, this set contains the identity matrix. Also, $\mathrm{GL}_n(\mathbb{Z})$ is closed with respect to matrix multiplication since determinants multiply. As for inverses, we note that the reduction algorithm above constructs the inverse of any matrix $A \in \mathrm{GL}_n(\mathbb{Z})$ as a product of matrices from $\mathcal{E}$. Thus multiplicative closure implies closure with respect to inverses.

The same observation shows that $\mathcal{E}$ generates $\mathrm{GL}_n(\mathbb{Z})$.

(??) The first claim is obvious: automorphisms are invertible, hence every automorphism is a bijection; now the statement follows since composition of maps defines the group law in both sets.

The second claim just says, in a wordy way, that the conjugation

$$\begin{aligned} \mathrm{ad}_g : G &\longrightarrow G \\ h &\mapsto ghg^{-1} \end{aligned}$$

is a homomorphism, which follows from straight forward computations. We just verify that it is multiplicative:

$$\mathrm{ad}_g(h_1h_2) = gh_1h_2g^{-1} = gh_1g^{-1}gh_2g^{-1} = \mathrm{ad}_g(h_1)\,\mathrm{ad}_g(h_2)\,.$$

To show that $\mathrm{Inn}(G)$ is normal in $\mathrm{Aut}(G)$, let $\xi : G \to G$ be an automorphism. We claim that the conjugate $\xi \circ \mathrm{ad}_g \circ \bar{\xi} = \mathrm{ad}_{\xi(g)}$, which

again is a straight forward computation:

$$\begin{aligned}\xi \circ \mathrm{ad}_g \circ \bar{\xi}(h) &= \xi\big(\mathrm{ad}_g\big(\bar{\xi}(h)\big)\big)\\ &= \xi\big(g\bar{\xi}(h)\,\bar{g}\big)\\ &= \xi(g)\,\xi\big(\bar{\xi}(h)\big)\,\xi(\bar{g})\\ &= \xi(g)\,h\inf\xi(g)\\ &= \mathrm{ad}_{\xi(g)}(h)\end{aligned}$$

It follows that the conjugate of a conjugation automorphism by an automorphism is again given by a conjugation. Thus conjugation automorphisms for a subgroup closed under conjugation.

The last claim is the most difficult, because we have to apply the same trick twice. Fix $h \in G$ so that ad_h generates $\mathrm{Inn}(G)$. That means, for every $g \in G$, we have $\mathrm{ad}_g = \mathrm{ad}_{h^k}$ for some exponent k. Now, we compute an arbitrary commutator:

$$\begin{aligned}[g_1, g_2] &= \mathrm{ad}_{g_1}(g_2)\,\bar{g}_2\\ &= \mathrm{ad}_{h^i}(g_2)\,\bar{g}_2\\ &= \big[h^i, g_2\big]\\ &= h^i\,\mathrm{ad}_{g_2}\big(h^{-i}\big)\\ &= h^i\,\mathrm{ad}_{h^j}\big(h^{-i}\big)\\ &= \big[h^i, h^j\big]\\ &= 1\end{aligned}$$

(1.1.12.7) First, we show that nilpotent groups are closed with respect to central extensions. So let

$$N \hookrightarrow G \twoheadrightarrow Q$$

be a short exact sequence, where N and Q are both nilpotent and where N is a central subgroup in G. Then N is automatically Abelian. Let

$$1 = Q_0 \leq Q_1 \leq \cdots \leq Q_u = Q$$

be a subgroup chain in Q satisfying $[Q, Q_{i+1}] \leq Q_i$ for each i. Let G_{i+1} be the preimage of Q_i in G. Put $G_0 := 1$. Note that by exactness, $G_1 = N$. Also note that since G_1 is central, we have $[G, G_1] \leq G_0$. For higher indices i, we find that the condition $[G, G_{i+1}] \leq G_i$ is equivalent to $[Q, Q_i] \leq Q_{i-1}$ by direct computation with coset representatives.

Conversely let G be nilpotent, i.e., suppose there is a subgroup chain

$$1 = G_0 \leq G_1 \leq G_2 \leq \cdots \leq G_u = G$$

satisfying $[G, G_{i+1}] \leq G_i$. We want to show that G can be obtained from Abelian groups within finitely many steps of taking central extensions. Note that $[G, G_1] \leq G_0 = 1$ implies that G_1 is central (and hence normal) in G. Thus, G is a central extension:

$$G_1 \hookrightarrow G \twoheadrightarrow {}^{G}\!/_{G_1}$$

If we can show that ${}^{G}\!/_{G_1}$ has a subgroup chain testifying to its nilpotency that is shorter than the one for G, we can iterate this procedure consistently reducing the length of the subgroup chain whence the process will terminate.

However, the chain

$$1 = {}^{G_1}\!/_{G_1} \leq {}^{G_2}\!/_{G_1} \leq {}^{G_3}\!/_{G_1} \leq \cdots \leq {}^{G_u}\!/_{G_1} = {}^{G}\!/_{G_1}$$

fits the bill.

(1.2.5.6) For $\mathbf{D}_8$, we first classify the seven non-trivial elements: there are four reflections about axes and three rotations. Of these, five have order 2 namely the rotation by π and the four reflections. This immediately yields a complete list of the five subgroups of order 2.

The two rotations by $\frac{\pi}{2}$ both generate the same cyclic group of order 4. Any other group of order 4 cannot contain either of

the two order 4 elements. Thus the other subgroups of order 4 consist of involutions only. Moreover, order four subgroups in $\mathbf{D}_8$ are normal. Thus, we will find them by looking at the orbits of order two elements under conjugation: it turns out that the two reflections about diagonals are conjugate and the two reflections about lines through opposite edge midpoints are another pair of conjugate reflections. The order two rotation is central. Thus, to form a union of conjugacy classes with three non-trivial order two elements, there are exactly two possibilities. Both of these actually yield subgroups: one generated by the first pair of conjugate reflections (also containing the rotation as their product) and the other one generated by the other pair of conjugate reflections.

Note, how in the classification, we also established the inclusion relations. Here is the resulting diagram:

PICTURE

Now for $\mathbf{S}_4$. We have seen in class that $\mathbf{S}_r$ has exactly one index 2 subgroup, namely $\mathbf{A}_r$. This leaves the following orders: 2, 3, 4, 6, and 8. Subgroups of order 8 are 2-Sylow subgroups, of which there are either 1 or 3. Note that any way of assigning the numbers 1 through 4 to the four corners of a square yields an embedding of $\mathbf{D}_8$ into $\mathbf{S}_4$. We find that up to conjugation in $\mathbf{D}_8$ there are three possible arrangements: the corner labeled 1 can be moved anywhere and then we have three choices to label the opposite corner. The remaining choice can be undone by a reflection and thus does not yield a different subgroup. These three 2-Sylows also contain all subgroups of order 4 and 2. Thus, these considerations determined the 2-power subgroups since we know the subgroup lattice of $\mathbf{D}_8$.

We can see four subgroups of order 6 immediately: the stabilizer of any number 1, 2, 3, or 4 is a symmetric group

permuting the remaining numbers. These are all isomorphic to $\mathbf{S}_3$. We shall show, that these are indeed all subgroups of order 6: note that by Cauchys theorem any subgroup of order 6 has an element of order 3. In $\mathbf{S}_4$ such an element must be a 3-cycle. This cycle leave a certain number fixed. Note that any other element of the subgroup must also fix that number for otherwise it would conjugate the given 3-cycle to a 3-cycle moving the fourth number. Then the two cycles would generate $\mathbf{A}_4$. This consideration rules out that our subgroup contains a product of two disjoint transpositions. Hence, its order two elements are transpositions (of which there is at least one in the subgroup). Now, a transposition together with a 3-cycle generates $\mathbf{S}_3$.

Inside our four subgroups of order 6, we also see four subgroups of order 3. By Sylow theorems, there cannot be more than four 3-Sylow subgroups in $\mathbf{S}_4$, which completes our list.

(1.3.1.11) Existence of a reduced word is clear: just delete offending subwords as long as necessary.

For uniqueness, suppose r and s are two equivalent reduced words. Let

$$r = w_0 \sim w_1 \sim \cdots \sim w_u = s$$

be a chain witnessing for the equivalence. Chose the chain so that the sum of lengths $\sum_{i=0}^{u} |w_i|$ is minimal. We claim that in this case, $u = 0$.

Suppose otherwise and let i be an index where $|w_i|$ is maximal. Then the two neighors w_{i-1} and w_{i+1} are obtained from w_i by deleting offending subwords. If those subwords do not overlap, we can replace w_i by the word where both subwords are deleted and replace the deletions by insertions. The total lengthof the chain decreases by 4. The cases where the subwords overlap are even easier since under these circumstances $w_{i-1} = w_{i+1}$.

(1.3.1.14) It is clear that every vertex v in the Cayley graph Γ has two neighbors for each generator $x \in X$, namely the vertices vx and vx^{-1} which are different since a one letter word is not equivalent to its inverse (any two reduced words are inequivalent if they are different). This proves the claim about the degree.

To see that Γ does not contain non-trvial cycles, assume that

$$v \rightarrowtail vx_1^{\varepsilon_1} \rightarrowtail vx_1^{\varepsilon_1}x_2^{\varepsilon_2} \rightarrowtail vx_1^{\varepsilon_1}x_2^{\varepsilon_2}x_3^{\varepsilon_3} \rightarrowtail \cdots \rightarrowtail vx_1^{\varepsilon_1}x_2^{\varepsilon_2}x_3^{\varepsilon_3}\cdots x_u^{\varepsilon_u} \rightarrowtail v$$

is such a non-trival cycle of minimum length. Then, clearly, no vertex can occur twice (otherwise, we could decompose the cycle into two shorter cycles one of which has to be non-trival). In particular, all the words $x_1^{\varepsilon_1}x_2^{\varepsilon_2}x_3^{\varepsilon_3}\cdots x_i^{\varepsilon_i}$ are pairwise different.

(??) ince any basis can be taken to any other basis, any non-zero vector can be taken to any vector by a matrix. Thus, every non-zero vector generated D^m as an left-$\mathbb{M}_{m\times m}(D)$-module.

(3.1.3.14) The statement is true. By distributivity, it suffices to show that elementary elements commute. Note that

$$m_1 \odot \cdots \odot m_i \odot n \odot n' \odot m_{i+1} \odot \cdots \odot m_u = m_1 \odot \cdots \odot m_i \odot n' \odot n \odot m_{i+1} \odot \cdots \odot m_u$$

which implies that we can reorder factors in an elementary element at will (neighbor transpositions generate the symmetric group). In particular, we have

$$(m_1 \odot \cdots \odot m_i) \odot (n_1 \odot \cdots \odot n_j) = (n_1 \odot \cdots \odot n_j) \odot (m_1 \odot \cdots \odot m_i)$$

which proves the claim.

(4.1.1.6) First, we show that B_F^M is a spanning set. Fix $\xi \in M$. We can write ξ as a finite linear combination

$$\xi = \zeta_1\zeta_1 + \cdots + \zeta_u\zeta_u$$

and then, we can write each $\zeta_i \in K$ as a linear combination of finitely many elements of B_F^K with F-coefficients. Substitution yields a finite combination for ξ of elements in B_F^M with F-coefficients.

Now we argue that B_F^M is a linearly independent set. Any linear combination of 0 with F-coefficients of elements in B_F^M can be rearranged as a combination of elements of B_K^M with K-coefficients (which are linear combinations of B_F^K-elements with F-coefficients). Linear independence of B_K^M implies that all K-coefficients have to be 0. Now linear independence of B_F^K implies that all F-coefficients are 0.

(4.1.2.7) $K_* \leq K$ be the field generated over F by the coefficients of $\mu_{\xi/K}$. Since $\mu_{\xi/K}(\xi) = 0$ and $\mu_{\xi/K} \in K_*[x]$, the minmal polynomial μ_{ξ/K_*} divides $\mu_{\xi/K}$. This implies

$$[M/K] = \deg\big(\mu_{\xi/K}\big) \geq \deg\big(\mu_{\xi/K_*}\big) = [M/K_*].$$

On the other hand $K_* \leq K$. It follows that $K = K_*$.

(4.1.2.8) First suppose that $M = F(\xi)$ is a simple extension. For any intermediate field K, the minimal polynomial $\mu_{\xi,K}$ divides $\mu_{\xi,F}$. Moreover, by 4.1.2.7, the field K is generated by the coefficients of $\mu_{\xi,K}$ Since there are only finitely many monic divisors of $\mu_{\xi,F}$ in $M[x]$, there can be only finitely many intermediate fields.

(4.1.2.10) If F is finite, then so is M. By (4.1.2.5), M/F is simple.

Thus, we assume that F is infinite. By induction, it suffices to argue that $F(\xi, \zeta)$ is a simple extension. However, since there are only finitely many intermediate fields, by pigeon hole principle, there are $\zeta \neq \xi$ with $F(\xi + \zeta\zeta) = F(\xi + \xi\zeta)$ whence (4.1.2.9) applies.

(??) IXME: [totatlly bogus] Only surjectivity requires proof. Zorn's lemma and (4.1.3.3).

(4.3.1.2) ...

(4.3.1.5) Let α be a solution of $x^p - \xi$ (in some extension of K). Then

$$x^p - \xi = (x - \alpha)^p$$

and the factors of $x^p - \xi$ are all of the form $(x - \alpha)^k$. Since $k\alpha \notin K$ is a coefficient of $(x - \alpha)^k$ we find that there is no non-trivial divisor of $x^p - \xi$ in $K[x]$.

(4.3.1.6) Just take the pth roots of all coefficients.