

Notes on the Mathematical Tripos

Sky Wilshaw

PART III

University of Cambridge
2020–2024

Contents

I.	Category Theory	
	<i>Lectured in Michaelmas 2023 by PROF. P. T. JOHNSTONE</i>	5
II.	Commutative Algebra	
	<i>Lectured in Michaelmas 2023 by DR. O. BECKER</i>	85
III.	Algebraic Geometry	
	<i>Lectured in Michaelmas 2023 by DR. D. RANGANATHAN</i>	163
IV.	Model Theory and Non-Classical Logic	
	<i>Lectured in Michaelmas 2023 by DR. J. SIQUEIRA</i>	217
V.	Group Cohomology	
	<i>Lectured in Lent 2024 by DR. C. J. B. BROOKES</i>	273
VI.	Large Cardinals	
	<i>Lectured in Lent 2024 by PROF. B. LÖWE</i>	295
VII.	Forcing and the Continuum Hypothesis	
	<i>Lectured in Lent 2024 by DR. R. MATTHEWS</i>	333

I. Category Theory

Lectured in Michaelmas 2023 by PROF. P. T. JOHNSTONE

(Course description goes here.)

Contents

1. Definitions and examples	8
1.1. Categories	8
1.2. Functors	10
1.3. Natural transformations	11
1.4. Equivalence of categories	13
1.5. Monomorphisms and epimorphisms	16
2. The Yoneda lemma	18
2.1. Statement and proof	18
2.2. Representable functors	21
2.3. Separating and detecting families	23
2.4. Projectivity	25
3. Adjunctions	27
3.1. Definition and examples	27
3.2. Comma categories	28
3.3. Units and counits	30
3.4. Reflections	33
4. Limits	35
4.1. Cones over diagrams	35
4.2. Limits	36
4.3. Preservation and creation	39
4.4. Interaction with adjunctions	41
4.5. General adjoint functor theorem	42
4.6. Special adjoint functor theorem	43
5. Monads	47
5.1. Definition	47
5.2. Eilenberg–Moore algebras	48
5.3. Kleisli categories	49
5.4. Comparison functors	51
5.5. Monadic adjunctions	53
6. Monoidal and enriched categories	60
6.1. Monoidal categories	60
6.2. The coherence theorem	61
6.3. Monoidal functors	63
6.4. Closed monoidal categories	65
6.5. Enriched categories	66
7. Additive and abelian categories	68
7.1. Additive categories	68

7.2.	Kernels and cokernels	71
7.3.	Abelian categories	71
7.4.	Exact sequences	73
7.5.	The five lemma	75
7.6.	The snake lemma	76
7.7.	Complexes in abelian categories	76
7.8.	Projective resolutions	79
7.9.	Derived functors	81

1. Definitions and examples

1.1. Categories

Definition. A category \mathcal{C} consists of

- (i) a collection of *objects* $\text{ob } \mathcal{C}$, denoted A, B, C, \dots ;
- (ii) a collection of *morphisms* $\text{mor } \mathcal{C}$, denoted f, g, h, \dots ;
- (iii) two operations $\text{dom}, \text{cod} : \text{mor } \mathcal{C} \rightarrow \text{ob } \mathcal{C}$, and we write $f : A \rightarrow B$ or $A \xrightarrow{f} B$ to state that f is a morphism with domain A and codomain B ;
- (iv) an operation $A \mapsto 1_A : A \rightarrow A$;
- (v) a *composition* operation $(f, g) \mapsto fg : \text{dom } g \rightarrow \text{cod } f$, defined exactly when $\text{cod } g = \text{dom } f$; satisfying
- (vi) $f1_A = f$ and $1_Bg = g$ whenever the composites are defined; and
- (vii) $(fg)h = f(gh)$ whenever the composites are defined.

Remark. (i) The collections of objects and morphisms may be sets or classes in some set theory, but our definitions are built to be interpretable in any system supporting first-order logic. If $\text{ob } \mathcal{C}$ and $\text{mor } \mathcal{C}$ are sets, we call \mathcal{C} a *small* category; otherwise we call it *large*.

- (ii) We could formulate a definition of category with no mention of objects, since objects biject with the identity morphisms. We will not take this approach here.
- (iii) Note that we choose fg to mean ‘first g and then f ’; this choice is a convention and the other one may be adopted.

Example. (i) **Set** is the category where the objects are all of the sets, and the morphisms are all of the functions between them, each of which is suitably tagged with an appropriate codomain. This must be done because set-theoretic functions do not ‘remember’ their codomain: $f(x) = x$ as a function $f : \mathbb{R} \rightarrow \mathbb{R}$ or $\mathbb{R} \rightarrow \mathbb{C}$ are equal sets.

- (ii) **Gp** is the category where the objects are all of the groups, and the morphisms are all of the group homomorphisms.
- (iii) **Rng** is the category where the objects are all of the rings, and the morphisms are all of the ring homomorphisms.
- (iv) For a field k , **Vect_k** is the category where the objects are all of the k -vector spaces, and the morphisms are all of the k -linear maps.
- (v) **Top** is the category where the objects are all of the topological spaces, and the morphisms are all of the continuous functions.

1. Definitions and examples

- (vi) **Met** is the category where the objects are all of the metric spaces, and the morphisms are all of the nonexpansive mappings, i.e. functions that do not increase the distance between points. One could choose a different convention, for example by letting morphisms be arbitrary continuous functions.
- (vii) **Mfd** is the category where the objects are all of the smooth manifolds, and the morphisms are C^∞ maps.
- (viii) **TopGp** is the category where the objects are all of the topological groups, and the morphisms are the continuous homomorphisms.
- (ix) **Htpy** is the category where the objects are all of the topological spaces, and the morphisms are equivalence classes of continuous functions under homotopy.
- (x) More generally, if \simeq is an equivalence relation on the morphisms of \mathcal{C} such that $f \simeq g$ implies $\text{dom } f = \text{dom } g$ and $\text{cod } f = \text{cod } g$, and the relation is stable under composition so $f \simeq g$ implies $fh \simeq gh$ and $kf \simeq kg$, we call \simeq a *congruence*. In this case, we can form the *quotient category* \mathcal{C}/\simeq , which has the same objects as \mathcal{C} , but its objects are equivalence classes of morphisms in \mathcal{C} under \simeq .
- (xi) **Rel** is the category where the objects are all of the sets, and the morphisms $A \rightarrow B$ are the relations $R \subseteq A \times B$, where composition is given by

$$S \circ R = \{(a, c) \mid \exists b \in B, (a, b) \in R \wedge (b, c) \in S\}$$

Note that if R and S happen to be functions, \circ is the standard composition operator. Therefore, **Set** is a subcategory of **Rel**.

- (xii) **Part** is the category where the objects are all of the sets, and the morphisms $A \rightarrow B$ are the partial functions $A \rightarrow B$. This is a subcategory of **Rel**, and **Set** is a subcategory of **Part**.
- (xiii) Given a category \mathcal{C} , we can construct its *opposite category* \mathcal{C}^{op} , where the objects and morphisms are the same as in \mathcal{C} , but dom and cod are swapped. We also reverse composition in the opposite category. This gives a duality principle: whenever a statement about categories is proven, a dual statement follows from applying the statement to an opposite category.
- (xiv) A small category with one object \star is a *monoid*, a group without inverses. In particular, every group can be seen as a small category on a single object in which every morphism is an isomorphism, i.e. invertible.
- (xv) A *groupoid* is a category in which every morphism is an isomorphism. For example, we can construct the *fundamental groupoid* of a topological space X . Here, the objects correspond to points x in X , and represent $\pi_1(X, x)$. Morphisms $x \rightarrow y$ are homotopy classes of paths starting at x and ending at y . Composition is path concatenation.
- (xvi) A category with at most one morphism between any pair of objects is a *preorder*. The existence of a morphism $A \rightarrow B$ corresponds to stating $A \leq B$ in the preorder. In partic-

I. Category Theory

ular, a partially ordered set (poset) is a small preorder in which the only isomorphisms are identity morphisms.

- (xvii) For a field k , \mathbf{Mat}_k is the category where the objects are the natural numbers, and the morphisms $n \rightarrow p$ are the $p \times n$ matrices over k . Composition is multiplication of matrices. The identity morphisms are the identity matrices.

1.2. Functors

Definition. Let \mathcal{C}, \mathcal{D} be categories. A *functor* $F : \mathcal{C} \rightarrow \mathcal{D}$ consists of a map $\text{ob } \mathcal{C} \xrightarrow{F} \text{ob } \mathcal{D}$ and a map $\text{mor } \mathcal{C} \xrightarrow{F} \text{mor } \mathcal{D}$, such that

- (i) $F(\text{dom } f) = \text{dom } Ff$;
- (ii) $F(\text{cod } f) = \text{cod } Ff$;
- (iii) $F(1_A) = 1_{FA}$; and
- (iv) $F(fg) = (Ff)(Fg)$ whenever fg is defined.

Example. (i) The *forgetful functors* $\mathbf{Gp} \rightarrow \mathbf{Set}, \mathbf{Rng} \rightarrow \mathbf{Set}, \mathbf{Top} \rightarrow \mathbf{Set}$ and so on forget that the objects are structures and forget the conditions on morphisms. Similarly, there are forgetful functors $\mathbf{Rng} \rightarrow \mathbf{AbGp}, \mathbf{Met} \rightarrow \mathbf{Top}, \mathbf{TopGp} \rightarrow \mathbf{Top}, \mathbf{TopGp} \rightarrow \mathbf{Gp}$.

- (ii) Any mapping $f : A \rightarrow UG$ from a set A to the underlying set of a group G extends uniquely to a homomorphism $FA \rightarrow G$, where FA is the free group on the set A . This can be made into a functor $F : \mathbf{Set} \rightarrow \mathbf{Gp}$: given $f : A \rightarrow B$, the homomorphism Ff is the unique homomorphism extending $A \xrightarrow{f} B \rightarrow FB$. Given $g : B \rightarrow C$, then $F(gf)$ and $(Fg)(Ff)$ both extend the same mapping $A \rightarrow FC$, so by the uniqueness property they are equal.
- (iii) The power-set construction $P : \mathbf{Set} \rightarrow \mathbf{Set}$ is a functor. PA is the set of all subsets of A , and given $f : A \rightarrow B$, Pf is the map sending S to the image of S under f .
- (iv) There is another power-set functor $P^* : \mathbf{Set}^{\text{op}} \rightarrow \mathbf{Set}$ (or $\mathbf{Set} \rightarrow \mathbf{Set}^{\text{op}}$). This has the same object map, but given $f : A \rightarrow B$, P^*f maps $S \subseteq B$ to its inverse image under f . A functor like this that reverses the direction of arrows is sometimes called *contravariant*; functors which do not are called *covariant*.
- (v) The construction of dual spaces in linear algebra gives rise to a functor $(-)^* : \mathbf{Vect}_k^{\text{op}} \rightarrow \mathbf{Vect}_k$. V^* is the space of linear maps $V \rightarrow k$, and a linear map $f : V \rightarrow W$ gives rise to $f^* : W^* \rightarrow V^*$ given by composition.
- (vi) \mathbf{Cat} is the category where the objects are the small categories and the morphisms are functors. This is well-defined as functors have identities and compositions.
- (vii) The assignment $\mathcal{C} \rightarrow \mathcal{C}^{\text{op}}$ defines a (covariant) functor $\mathbf{Cat} \rightarrow \mathbf{Cat}$.

- (viii) A functor between monoids is a monoid homomorphism.
- (ix) A functor between groups is a group homomorphism.
- (x) A functor between posets is an order-preserving map.
- (xi) If G is a group, a functor $F : G \rightarrow \mathbf{Set}$ defines a set $A = F\star$, together with a collection of endomorphisms of A denoted $a \mapsto g \cdot a$ for each $g \in G$. This collection of endomorphisms is compatible with the identity and composition, so is precisely the definition of a group action or permutation representation of G .
- (xii) If G is a group, a functor $F : G \rightarrow \mathbf{Vect}_k$ is a k -linear representation of G .
- (xiii) The fundamental group of a topological space defines a functor $\pi_1 : \mathbf{Top}_\star \rightarrow \mathbf{Gp}$, where \mathbf{Top}_\star is the category of pointed topological spaces.

1.3. Natural transformations

Definition. Let \mathcal{C}, \mathcal{D} be categories, and $F, G : \mathcal{C} \Rightarrow \mathcal{D}$ be functors. A *natural transformation* $\alpha : F \rightarrow G$ is a mapping $\text{ob } \mathcal{C} \rightarrow \text{mor } \mathcal{D}$ denoted $A \mapsto \alpha_A$, such that

- (i) $\alpha_A : FA \rightarrow GA$ for all A ; and
- (ii) for any morphism $f : A \rightarrow B$ in \mathcal{C} , the square

$$\begin{array}{ccc} FA & \xrightarrow{Ff} & FB \\ \alpha_A \downarrow & & \downarrow \alpha_B \\ GA & \xrightarrow{Gf} & GB \end{array}$$

commutes. Such squares are called *naturality* squares.

If we have a natural transformation $\beta : G \rightarrow H$, we can define $\beta\alpha$ by $(\beta\alpha)_A = \beta_A\alpha_A$. We therefore have a category $[\mathcal{C}, \mathcal{D}]$ whose objects are the functors $\mathcal{C} \rightarrow \mathcal{D}$ and whose morphisms are the natural transformations between them.

Example. (i) Given a vector space V , we have a linear map $\alpha_V : V \rightarrow V^{**}$ sending $v \in V$ to the map $f \mapsto f(v)$. This is a natural transformation $\alpha : 1_{\mathbf{Vect}_k} \rightarrow (-)^{**}$. The naturality squares are of the form

$$\begin{array}{ccc} V & \xrightarrow{f} & W \\ \alpha_V \downarrow & & \downarrow \alpha_W \\ V^{**} & \xrightarrow{f^{**}} & W^{**} \end{array}$$

where

$$\alpha_V(v) = f \mapsto f(v); \quad f^{**}(g)(h) = g(f^*h) = g(h \circ f)$$

I. Category Theory

We show the naturality square commutes.

$$\begin{aligned}
 ((g \mapsto h \mapsto g(h \circ f)) \circ \alpha_V)(v) &= (g \mapsto h \mapsto g(h \circ f))(\alpha_V v) \\
 &= (g \mapsto h \mapsto g(h \circ f))(k \mapsto kv) \\
 &= h \mapsto (k \mapsto kv)(h \circ f) \\
 &= h \mapsto (h \circ f)v \\
 &= h \mapsto (h(fv)) \\
 &= \alpha_W(fv) \\
 &= (\alpha_W \circ f)v
 \end{aligned}$$

- (ii) There is an inclusion from any set A to its free group FA . The map sending a set A to the inclusion $A \rightarrow FA$ is a natural transformation $1_{\mathbf{Set}} \rightarrow UF$. Naturality is built into the definition of F on morphisms.

$$\begin{array}{ccc}
 A & \xrightarrow{f} & B \\
 \alpha_A \downarrow & & \downarrow \alpha_B \\
 UFA & \xrightarrow{UF(f)} & UFB
 \end{array}$$

- (iii) There is a mapping $\alpha_A : A \rightarrow PA$ by mapping $a \in A$ to $\{a\} \in PA$. This is a natural transformation $1_{\mathbf{Set}} \rightarrow P$, since $Pf\{a\} = \{fa\}$.

$$\begin{array}{ccc}
 A & \xrightarrow{f} & B \\
 \alpha_A \downarrow & & \downarrow \alpha_B \\
 PA & \xrightarrow{Pf} & PB
 \end{array}$$

- (iv) Let $f, g : P \Rightarrow Q$ be order-preserving maps between posets. Then for $x \leq y$ in P , the naturality square is

$$\begin{array}{ccc}
 fx & \longrightarrow & fy \\
 \alpha_x \downarrow & & \downarrow \alpha_y \\
 gx & \longrightarrow & gy
 \end{array}$$

In particular, the existence of α_x proves that $fx \leq gx$. Thus a natural transformation $f \rightarrow g$ exists if and only if $fx \leq gx$ pointwise for all $x \in P$. Note that every square of morphisms in a poset commutes.

- (v) Let $u, v : G \Rightarrow H$ be group homomorphisms. For $g \in G$, the naturality square is

$$\begin{array}{ccc}
 * & \xrightarrow{ug} & * \\
 \alpha_* \downarrow & & \downarrow \alpha_* \\
 * & \xrightarrow{vg} & *
 \end{array}$$

1. Definitions and examples

A natural transformation $\alpha : u \rightarrow v$ is an element $\alpha_* = h \in H$ such that $hu(g) = v(g)h$ for all g , or equivalently, $v(g) = hu(g)h^{-1}$. Thus a natural transformation exhibits a conjugacy between two homomorphisms. In particular, the natural transformations $u \rightarrow u$ are the elements of the centraliser of $u(G)$.

(vi) Let A, B be permutation representations of G , that is, functors $G \rightarrow \mathbf{Set}$.

$$\begin{array}{ccc} A_* & \xrightarrow{Ag} & A_* \\ f \downarrow & & \downarrow f \\ B_* & \xrightarrow{Bg} & B_* \end{array}$$

A natural transformation $f : A \rightarrow B$ is a mapping of the underlying sets $A_* \rightarrow B_*$ satisfying $g \cdot f(a) = f(g \cdot a)$ for all $a \in A$ and $g \in G$. This is the definition of a G -equivariant map.

(vii) For any (nice) pointed topological space X with base point x , the *Hurewicz homomorphism* is a map $h_{n,x} : \pi_n(X, x) \rightarrow H_n(X)$. This is a natural transformation $\pi_n \rightarrow H_n U$ where U is the forgetful functor $\mathbf{Top}_* \rightarrow \mathbf{Top}$.

1.4. Equivalence of categories

There is a notion of isomorphism of categories, namely, isomorphism in the category \mathbf{Cat} . For example, $\mathbf{Rel} \cong \mathbf{Rel}^{\text{op}}$ via the functor

$$A \mapsto A; \quad R \mapsto R^\circ = \{(b, a) \mid (a, b) \in R\}$$

However, there is a weaker notion that is often more useful in practice, called equivalence. To define this, we need a notion of ‘natural isomorphism’. There are two obvious definitions, which we show are equivalent.

Lemma. Let $\alpha : F \rightarrow G$ be a natural transformation between functors $\mathcal{C} \rightrightarrows \mathcal{D}$. Then α is an isomorphism in the functor category $[\mathcal{C}, \mathcal{D}]$ if and only if each component α_A is an isomorphism in \mathcal{D} .

Proof. The forward direction is clear as composition in $[\mathcal{C}, \mathcal{D}]$ is pointwise; if β is an inverse for α , then β_A is an inverse for α_A . Suppose β_A is an inverse for α_A for each A . We show the β collectively form a natural transformation by verifying the naturality squares. Given $f : A \rightarrow B$ in \mathcal{C} , consider

$$\begin{array}{ccc} FA & \xrightarrow{Ff} & FB \\ \beta_A \uparrow \downarrow \alpha_A & & \alpha_B \downarrow \uparrow \beta_B \\ GA & \xrightarrow{Gf} & GB \end{array}$$

Then

$$(Ff)\beta_B = \beta_B\alpha_B(Ff)\beta_A = \beta_B(Gf)\alpha_A\beta_A = \beta_B(Gf)$$

using naturality of α . Thus β is natural, and an inverse for α . □

I. Category Theory

Definition. Let \mathcal{C}, \mathcal{D} be categories. An *equivalence* between \mathcal{C} and \mathcal{D} is a pair of functors

$$F : \mathcal{C} \rightarrow \mathcal{D}; \quad G : \mathcal{D} \rightarrow \mathcal{C}$$

and a pair of natural isomorphisms

$$\alpha : 1_{\mathcal{C}} \rightarrow GF; \quad \beta : FG \rightarrow 1_{\mathcal{D}}$$

If \mathcal{C} and \mathcal{D} are equivalent, we write $\mathcal{C} \simeq \mathcal{D}$.

The reason the natural isomorphisms point in opposite directions will be clarified later. A property P of categories that is called *categorical* if whenever \mathcal{C} satisfies P and $\mathcal{C} \simeq \mathcal{D}$, then \mathcal{D} satisfies P . For example, the properties of being a preorder or being a groupoid are categorical. Being a partial order or being a group are not categorical. Generally, properties that rely on equality of objects, not isomorphism, will not be categorical.

Example. (i) Let \mathbf{Set}_* be the category of pointed sets and functions preserving the base point. Then $\mathbf{Set}_* \simeq \mathbf{Part}$ by

$$F : \mathbf{Set}_* \rightarrow \mathbf{Part}; \quad F(A, a) = A \setminus \{a\}; \quad F((A, a) \xrightarrow{f} (B, b))(x) = f(x)$$

and

$$G : \mathbf{Part} \rightarrow \mathbf{Set}_*; \quad G(A) = A \cup \{A\}; \quad G(A \xrightarrow{f} B \text{ partial})(x) = \begin{cases} f(x) & \text{if } f \text{ is defined at } x \\ B & \text{otherwise} \end{cases}$$

Note that $FG = 1_{\mathbf{Part}}$, but GF is not equal to $1_{\mathbf{Set}_*}$. It is not possible for these two categories to be isomorphic, because there is an isomorphism class of \mathbf{Part} that has only one member, namely $\{\emptyset\}$, but this cannot occur in \mathbf{Set}_* .

(ii) Let \mathbf{fdVect}_k be the category of finite-dimensional vector spaces over k . This category is equivalent to its opposite category $\mathbf{fdVect}_k^{\text{op}}$ via the dual space functors in both directions. The natural isomorphisms α and β are both as in the double dual example given above.

(iii) We show $\mathbf{fdVect}_k \simeq \mathbf{Mat}_k$. Define

$$F : \mathbf{Mat}_k \rightarrow \mathbf{fdVect}_k; \quad F(n) = k^n$$

and sending a matrix A to the linear map it represents in the standard basis. For each finite-dimensional vector space V , choose a particular basis. Define

$$G : \mathbf{fdVect}_k \rightarrow \mathbf{Mat}_k; \quad G(V) = \dim V$$

and let $G(\theta)$ be the matrix representing θ with respect to the particular bases chosen above. Then $GF = 1_{\mathbf{Mat}_k}$, as long as we chose the bases above in such a way that the k^n have the standard basis. Further, FG is naturally isomorphic to $1_{\mathbf{fdVect}_k}$, since the chosen bases define isomorphisms $k^{\dim V} \rightarrow V$, which are natural in V .

1. Definitions and examples

In line with the idea that we do not want to consider equality of objects but only equality of morphisms, we make the following definitions.

Definition. Let $F : \mathcal{C} \rightarrow \mathcal{D}$ be a functor. We say that F is

- (i) *faithful*, if for each $f, g \in \text{mor } \mathcal{C}$ with equal domain and codomain, $Ff = Fg$ implies $f = g$;
- (ii) *full*, if for each $FA \xrightarrow{g} FB$, there exists a morphism $A \xrightarrow{f} B$ such that $Ff = g$;
- (iii) *essentially surjective*, if every $B \in \text{ob } \mathcal{D}$ is isomorphic to some FA for $A \in \text{ob } \mathcal{C}$.

Note that if F is full and faithful, it is *essentially injective*: if $FA \xrightarrow{g} FB$ is an isomorphism, the unique $A \xrightarrow{f} B$ with $Ff = g$ is an isomorphism, because its inverse is the unique $B \rightarrow A$ mapped to g^{-1} .

Lemma. Let $F : \mathcal{C} \rightarrow \mathcal{D}$ be a functor. Then F is part of an equivalence $\mathcal{C} \simeq \mathcal{D}$ if and only if F is full, faithful, and essentially surjective.

Proof. Suppose G, α, β make F into an equivalence. The existence of β ensures that $B \simeq FGB$ for any $B \in \text{ob } \mathcal{D}$, giving essential surjectivity. For faithfulness, for any $A \xrightarrow{f} B$ in \mathcal{C} , we have $f = \alpha_B^{-1}(GFf)\alpha_A$, allowing us to reproduce f from its domain, codomain, and image under F . For fullness, consider $FA \xrightarrow{g} FB$, and define $f = \alpha_B^{-1}(Gg)\alpha_A : A \rightarrow B$. Then, $GFf = Gg$. As G is faithful by symmetry, $Ff = g$.

For the converse, for each object $B \in \mathcal{D}$, we choose an isomorphism $\beta_B : FA \rightarrow B$ where $A \in \mathcal{C}$, and define the action of G at B to be this A . Then we define G on morphisms by letting $G(B \xrightarrow{g} C)$ be the unique $GB \rightarrow GC$ whose image under F is $\beta_C^{-1} \circ g \circ \beta_B$, thus making the following diagram commute.

$$\begin{array}{ccc} FGB & \xrightarrow{FGg} & FGC \\ \beta_B \downarrow & & \uparrow \beta_C^{-1} \\ B & \xrightarrow{g} & C \end{array}$$

This is functorial: given $h : C \rightarrow D$, we can form $G(hg)$ and $(Gh)(Gg)$ which have the same image under F , so must be equal.

$$\begin{array}{ccccc} & FGB & \xrightarrow{FG(hg)} & FGD & \\ \beta_B \swarrow & & \searrow FGg & & \nwarrow \beta_D^{-1} \\ B & & & & D \\ & \searrow g & & \nearrow FGh & \\ & C & \xrightarrow{\beta_C^{-1}} & FGC & \xrightarrow{\beta_C} & C \\ & & \nearrow 1_C & & \searrow h \end{array}$$

I. Category Theory

By construction, β is a natural isomorphism $FG \rightarrow 1_{\mathcal{D}}$. It suffices to construct the natural isomorphism $\alpha : 1_{\mathcal{C}} \rightarrow GF$. Its component at A is the unique isomorphism whose image under F is

$$FA \xrightarrow{\beta_{FA}^{-1}} FGFA$$

Consider a naturality square for α .

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ \alpha_A \downarrow & & \downarrow \alpha_B \\ GFA & \xrightarrow{GFf} & GFB \end{array}$$

As F is faithful, to show this diagram commutes, it suffices to show that its image under F commutes.

$$\begin{array}{ccc} FA & \xrightarrow{Ff} & FB \\ F\alpha_A = \beta_{FA}^{-1} \downarrow & & \downarrow F\alpha_B = \beta_{FB}^{-1} \\ FGFA & \xrightarrow{FGFf} & FGFB \end{array}$$

This commutes by naturality of β^{-1} . □

We call a subcategory full if its inclusion functor is full.

Definition. A category is called *skeletal* if every isomorphism class has a single member. A *skeleton* of \mathcal{C} is a full subcategory \mathcal{C}' containing exactly one object for each isomorphism class.

Note that an equivalence of skeletal categories is bijective on objects, and hence is an isomorphism of categories.

1.5. Monomorphisms and epimorphisms

Definition. A morphism $f : A \rightarrow B$ is a *monomorphism*, and is called *monic*, if $fg = fh$ implies $g = h$ whenever the compositions are defined. Dually, f is an *epimorphism*, and is called *epic*, if $gf = hf$ implies $g = h$ whenever the compositions are defined.

Monomorphisms are left-cancellable; epimorphisms are right-cancellable. We will often denote a monomorphism with an arrow with a tail $A \rightarrowtail B$, and denote epimorphisms with double-headed arrows $A \twoheadrightarrow B$. Isomorphisms are clearly monic and epic; if all monic and epic morphisms in a category are isomorphisms, we call the category *balanced*.

Example. (i) In **Set**, the monomorphisms are precisely the injective functions, and the epimorphisms are precisely the surjective functions. Thus **Set** is balanced.

(ii) In **Gp**, the monomorphisms are the injective functions, and the epimorphisms are the surjective functions.

1. Definitions and examples

- (iii) In **Rng**, the monomorphisms are again the injective functions, but there are epimorphisms that are not surjective, for example the inclusion $\mathbb{Z} \rightarrow \mathbb{Q}$.
- (iv) In **Top**, the monomorphisms are the injective functions, and the epimorphisms are the surjective functions. However, **Top** is not balanced, because continuous bijections need not have continuous inverses.
- (v) In a preorder, any morphism is monic and epic. The category is balanced if and only if it is an equivalence relation (or equivalently, symmetric).

2. The Yoneda lemma

2.1. Statement and proof

Definition. A category \mathcal{C} is called *locally small* if the collection of morphisms $A \rightarrow B$ are parametrised by a set. In this case, we write $\mathcal{C}(A, B)$ for the set of such morphisms.

Given an object A of a locally small category, we can define a functor

$$\mathcal{C}(A, -) : \mathcal{C} \rightarrow \mathbf{Set}$$

given by

$$B \mapsto \mathcal{C}(A, B); \quad (B \xrightarrow{f} C) \mapsto ((A \xrightarrow{g} B) \mapsto fg)$$

This is functorial by associativity of function composition. We can also define

$$\mathcal{C}(-, A) : \mathcal{C}^{\text{op}} \rightarrow \mathbf{Set}$$

by

$$B \mapsto \mathcal{C}(B, A); \quad (B \xrightarrow{f} C) \mapsto ((C \xrightarrow{g} A) \mapsto gf)$$

Lemma (Yoneda lemma). Let \mathcal{C} be a locally small category. Let $A \in \text{ob } \mathcal{C}$, and let $F : \mathcal{C} \rightarrow \mathbf{Set}$ be a functor. Then,


(i) there is a bijection

$$\{\text{natural transformations } \mathcal{C}(A, -) \rightarrow F\} \leftrightarrow \{\text{elements of } FA\}$$

(ii) and further, this bijection is natural in both A and F .

This shows that we can consider a natural transformation $\mathcal{C}(A, -) \rightarrow F$ as a way to evaluate morphisms at a point $x \in FA$.

Example. Consider the category \mathcal{C} of the form



and the functor $F : \mathcal{C} \rightarrow \mathbf{Set}$ given by

$$F(A) = \{1, 2\}; \quad F(B) = \{3\}; \quad F(C) = \{4, 5, 6\}$$

and

$$F(f)(1) = F(f)(2) = 3; \quad F(g)(1) = 4; \quad F(g)(2) = 5$$

A natural transformation $\alpha : \mathcal{C}(A, -) \rightarrow F$ is given by its components

$$\alpha_A : \{1_A\} \rightarrow \{1, 2\}; \quad \alpha_B : \{f\} \rightarrow \{3\}; \quad \alpha_C : \{g\} \rightarrow \{4, 5, 6\}$$

subject to the naturality square

$$\begin{array}{ccc} \{1_A\} & \xrightarrow{\mathcal{C}(A,g)} & \{g\} \\ \alpha_A \downarrow & & \downarrow \alpha_C \\ \{1, 2\} & \xrightarrow{Fg} & \{4, 5, 6\} \end{array}$$

which enforces that

$$(Fg)(\alpha_A) = \alpha_C(g)$$

This means that such a natural transformation α is defined uniquely by a choice of $(Fg)(\alpha_A)$; that is, a choice of an element of FA .

Example. Let G be a group in the set-theoretic sense. Let us represent G as the category \mathcal{C} ; that is, let

$$\text{ob } \mathcal{C} = \{\star\}; \quad \text{mor } \mathcal{C} = G$$

Consider the functor $F : \mathcal{C} \rightarrow \mathbf{Set}$ given by

$$F(\star) = G; \quad F(g)(h) = gh$$

If $\alpha : \mathcal{C}(\star, -) \rightarrow F$ is a natural transformation, for each $g \in G$, $\alpha_\star(g)$ is a map $G \rightarrow G$. The naturality condition ensures that α respects the group structure. Applying the Yoneda lemma, we find that every map $G \rightarrow G$ that respects the group structure in this way is just the action of multiplication by some element of the group.

We prove part (i) now, and postpone (ii) until some corollaries have been established.

Proof. We want to show that a natural transformation $\alpha : \mathcal{C}(A, -) \rightarrow F$ is a way to evaluate morphisms at a point $x \in FA$. To find a sensible value for x , we evaluate the identity morphism $1_A : A \rightarrow A$.

$$\Phi : (\mathcal{C}(A, -) \rightarrow F) \rightarrow FA; \quad \Phi(\alpha) = \alpha_A(1_A) \in FA$$

Now, given a point $x \in FA$, we want to create a natural transformation that evaluates functions $A \rightarrow B$ and yields a point in FB . We define

$$\Psi : FA \rightarrow (\mathcal{C}(A, -) \rightarrow F); \quad \Psi(x)_B(A \xrightarrow{f} B) = (Ff)x$$

For $h : B \rightarrow C$, the naturality square is as follows.

$$\begin{array}{ccc} \mathcal{C}(A, B) & \xrightarrow{\mathcal{C}(A, h)} & \mathcal{C}(A, C) \\ \Psi(x)_B \downarrow & & \downarrow \Psi(x)_C \\ FB & \xrightarrow{Fh} & FC \end{array}$$

I. Category Theory

Here, $\mathcal{C}(A, h)$ denotes the operation $g \mapsto hg$. For $f : A \rightarrow B$,

$$\Psi(x)_C(\mathcal{C}(A, h)(f)) = \Psi(x)_C(hf) = (F(hf))x$$

and

$$(Fh)(\Psi(x)_B(f)) = (Fh)((Ff)x) = (F(hf))x$$

as required. Hence the ‘evaluate at x ’ map $\Psi(x)$ is a natural transformation. We show that these two constructions are inverses.

$$\Phi\Psi(x) = \Psi(x)_A(1_A) = (F1_A)x = 1_{FA}x = x$$

Let $\alpha : \mathcal{C}(A, -) \rightarrow F$ be a natural transformation, let $B \in \text{ob } \mathcal{C}$, and let $f : A \rightarrow B$. Then $\alpha_B(f)$ and $(\Psi\Phi(\alpha))_B(f)$ are elements of FB ; we show they coincide.

$$(\Psi\Phi(\alpha))_B(f) = (Ff)(\Phi(\alpha)) = (Ff)(\alpha_A(1_A))$$

Naturality of α shows that the following diagram commutes.

$$\begin{array}{ccc} \mathcal{C}(A, A) & \xrightarrow{\mathcal{C}(A, f)} & \mathcal{C}(A, B) \\ \alpha_A \downarrow & & \downarrow \alpha_B \\ FA & \xrightarrow{Ff} & FB \end{array}$$

Thus,

$$(\Psi\Phi(\alpha))_B(f) = \alpha_B(f1_A) = \alpha_B(f)$$

Hence, Φ and Ψ are inverse bijections. □

Corollary. For any locally small category \mathcal{C} , the map

$$A \mapsto \mathcal{C}(A, -)$$

is a full and faithful functor

$$Y : \mathcal{C}^{\text{op}} \rightarrow [\mathcal{C}, \mathbf{Set}]$$

This is called the *Yoneda embedding*.

Proof. Let $F = \mathcal{C}(B, -)$ in the Yoneda lemma. Then there is a bijection

$$\mathcal{C}(B, A) \leftrightarrow \{\text{natural transformations } \mathcal{C}(A, -) \rightarrow \mathcal{C}(B, -)\}$$

This bijection maps $f : B \rightarrow A$ to the natural transformation given by composition with f . This is functorial as composition in \mathcal{C} is associative. □

2. The Yoneda lemma

This says that any locally small category \mathcal{C} is equivalent to a full subcategory of a functor category $[\mathcal{C}^{\text{op}}, \mathbf{Set}]$. The category $[\mathcal{C}^{\text{op}}, \mathbf{Set}]$ is sometimes called the category of *presheaves* on \mathcal{C} , so any category embeds into its category of presheaves.

We now explain and prove part (ii) of the Yoneda lemma. Suppose that \mathcal{C} were small, so $[\mathcal{C}, \mathbf{Set}]$ were locally small. Then we have two functors

$$\mathcal{C} \times [\mathcal{C}, \mathbf{Set}] \rightarrow \mathbf{Set}$$

The first is the evaluation functor

$$(A, F) \mapsto FA$$

The second is the composite

$$\mathcal{C} \times [\mathcal{C}, \mathbf{Set}] \xrightarrow{Y \times 1} [\mathcal{C}, \mathbf{Set}]^{\text{op}} \times [\mathcal{C}, \mathbf{Set}] \xrightarrow{[\mathcal{C}, \mathbf{Set}](-, -)} \mathbf{Set}$$

The naturality condition is that Φ and Ψ are natural transformations between these two functors, and thus are natural isomorphisms.

Proof. Let $f : A \rightarrow A'$, $\alpha : F \rightarrow F'$, and $x \in FA$. If x' is the image of x under the diagonal of the naturality square

$$\begin{array}{ccc} FA & \xrightarrow{Ff} & FA' \\ \alpha_A \downarrow & & \downarrow \alpha_{A'} \\ F'A & \xrightarrow{F'f} & F'A' \end{array}$$

we want to show that $\Psi(x')$ is the composite

$$\mathcal{C}(A', -) \xrightarrow{\mathcal{C}(f, -)} \mathcal{C}(A, -) \xrightarrow{\Psi(x)} F \xrightarrow{\alpha} F'$$

But this can be easily verified, as the composite maps

$$1_{A'} \mapsto f \mapsto (Ff)(x) \mapsto \alpha_{A'}(Ff)(x) = x'$$

as required. □

2.2. Representable functors

Definition. Let \mathcal{C} be a locally small category. A functor $F : \mathcal{C} \rightarrow \mathbf{Set}$ is called *representable* if it is isomorphic to $\mathcal{C}(A, -)$ for some A . A *representation* of F is a pair (A, x) where $A \in \text{ob } \mathcal{C}$, and $x \in FA$ is such that

$$\Psi(x) : \mathcal{C}(A, -) \rightarrow F$$

is a natural isomorphism. In this case, we say that x is a *universal element* of F .

Corollary. Suppose (A, x) and (B, y) are representations of $F : \mathcal{C} \rightarrow \mathbf{Set}$. Then there is a unique isomorphism $f : A \rightarrow B$ such that $Ff(x) = y$.

I. Category Theory

Proof. The Yoneda lemma shows that the elements of FA correspond to natural transformations $\mathcal{C}(A, -) \rightarrow F$, and similarly for the elements of FB . Thus, $Ff(x) = y$ equivalently says that

$$\begin{array}{ccc} \mathcal{C}(B, -) & \xrightarrow{\mathcal{C}(f, -)} & \mathcal{C}(A, -) \\ & \searrow \Psi(y) & \swarrow \Psi(x) \\ & F & \end{array}$$

commutes. But $\Psi(x)$ and $\Psi(y)$ are isomorphisms, so this holds if and only if f is the unique isomorphism sent by the Yoneda embedding to $\Psi(x)^{-1}\Psi(y)$. \square

- (i) Consider the forgetful functor $\mathbf{Gp} \rightarrow \mathbf{Set}$. This is representable by the free group on one generator, \mathbb{Z} . Similarly, the forgetful functor $\mathbf{Rng} \rightarrow \mathbf{Set}$ is represented by the free ring on one generator, $\mathbb{Z}[x]$.
- (ii) The forgetful functor $\mathbf{Top} \rightarrow \mathbf{Set}$ is representable by the one-point space.
- (iii) The contravariant power set functor $P^* : \mathbf{Set}^{\text{op}} \rightarrow \mathbf{Set}$ is representable by the two-element set $2 = \{0, 1\}$ via the bijection mapping $f : A \rightarrow 2$ to $f^{-1}(1)$.
- (iv) The covariant power set functor $P : \mathbf{Set} \rightarrow \mathbf{Set}$ is not representable. $\mathbf{Set}(A, 1) \cong 1$ for any A , but $P1 \cong 2 \not\cong 1$.
- (v) Define $\Omega : \mathbf{Top}^{\text{op}} \rightarrow \mathbf{Set}$ to be the functor mapping a space X to its set of open subsets. If $f : X \rightarrow Y$ is continuous, this induces a map $\Omega f : \Omega Y \rightarrow \Omega X$. This is representable by the *Sierpiński space* Σ with two points $\{0, 1\}$ and open sets

$$\emptyset; \{1\}; \Sigma$$

The continuous maps $f : X \rightarrow \Sigma$ are exactly the characteristic functions of the open subsets of X , because continuity is just that $f^{-1}(\{1\})$ is open.

- (vi) The dual vector space functor $(-)^* : \mathbf{Vect}_k^{\text{op}} \rightarrow \mathbf{Vect}_k$ is not representable because its codomain is not \mathbf{Set} , but composing with the forgetful functor makes it representable by the one-dimensional space k .
- (vii) Let G be a group. The (unique up to isomorphism) representable functor $G \rightarrow \mathbf{Set}$ is the *Cayley representation* of the group; that is, the set G acting on itself by multiplication.
- (viii) Let A, B be objects of a locally small category \mathcal{C} . Then there is a functor $\mathcal{C}^{\text{op}} \rightarrow \mathbf{Set}$ sending C to the Cartesian product

$$\mathcal{C}(C, A) \times \mathcal{C}(C, B)$$

If this is representable, we call the representing object a categorical *product* of A and B , and denote it $A \times B$. The universal element is a pair of morphisms $\pi_1 : A \times B \rightarrow A, \pi_2 : A \times B \rightarrow B$, called *projections*. This has the property that for any pair $(f : C \rightarrow A, g : C \rightarrow B)$ there exists a unique morphism $h = (f, g) : C \rightarrow A \times B$ satisfying $\pi_1 h = f, \pi_2 h = g$.

- (ix) Dually, there is the notion of a *coproduct* $A + B$, which is a representing object of the functor mapping C to

$$\mathcal{C}(A, C) \times \mathcal{C}(B, C)$$

with *coprojections* $\nu_1 : A \rightarrow A + B, \nu_2 : B \rightarrow A + B$.

- (x) Let $f, g : A \rightrightarrows B$ be a parallel pair of morphisms in a locally small category \mathcal{C} . Define a functor $F : \mathcal{C}^{\text{op}} \rightarrow \mathbf{Set}$ by sending C to

$$\{h : C \rightarrow A \mid fh = gh\}$$

If this is representable, we call the representation an *equaliser* of f and g . This consists of a representing object E with a morphism $e : E \rightarrow A$ satisfying $fe = ge$. Moreover, for any morphism h with $fh = gh$, h factors uniquely through e . Hence, e is a monomorphism. Monomorphisms that occur in this way are called *regular*.

- (xi) Dually, there is also a notion of coequaliser, giving rise to an epimorphism. We again call epimorphisms *regular* if they arise in this way.

In **Set**, the categorical product is the Cartesian product, and the categorical coproduct is the disjoint union. The equaliser of $f, g : A \rightrightarrows B$ is the set

$$\{a \in A \mid fa = ga\}$$

The coequaliser of f, g is the quotient

$$B/\sim$$

where \sim is the equivalence relation generated by $fa \sim ga$.

In **Gp**, the product is the direct product, but the coproduct is the *free product* $A * B$. The equaliser of $f, g : A \rightrightarrows B$ is as in **Set**, which is a subgroup of A . The coequaliser of f, g is the quotient by the smallest congruence containing all pairs (fa, ga) . In **Set** and **Gp**, all monomorphisms and epimorphisms are regular.

In **Top**, not all injections or surjections are regular monomorphisms or epimorphisms.

2.3. Separating and detecting families

Definition. Let \mathcal{C} be a locally small category, and \mathcal{G} a class of objects of \mathcal{C} . We say that

- (i) \mathcal{G} is a *separating family* for \mathcal{C} if the functors $\mathcal{C}(G, -)$ for $G \in \mathcal{G}$ are collectively faithful; that is, if $f, g : A \rightrightarrows B$, the equations $fh = gh$ for all $h : G \rightarrow A$ with $G \in \mathcal{G}$ imply $f = g$.

$$G \xrightarrow{h} A \begin{array}{c} \xrightarrow{f} \\ \xrightarrow{g} \end{array} B$$

I. Category Theory

- (ii) \mathcal{G} is a *detecting family* for \mathcal{C} if the functors $\mathcal{C}(G, -)$ for $G \in \mathcal{G}$ collectively *reflect isomorphisms*; that is, if $f : A \rightarrow B$ is such that every $h : G \rightarrow B$ with $G \in \mathcal{G}$ factors uniquely through A , then f is an isomorphism.

$$\begin{array}{ccc} G & \xrightarrow{g} & A \\ & \searrow h & \downarrow f \\ & & B \end{array}$$

If $\mathcal{G} = \{G\}$, we call G a *separator* or *detector* respectively.

Separating and detecting families are both sometimes called *generating families*.

Lemma. (i) If \mathcal{C} has equalisers, then any detecting family is separating.

(ii) If \mathcal{C} is balanced, then any separating family is detecting.

Proof. Part (i). Suppose \mathcal{G} is detecting, and $f, g : A \rightrightarrows B$ such that every morphism $h : G \rightarrow A$ with $G \in \mathcal{G}$ has $fh = gh$. Then every such $h : G \rightarrow A$ with $G \in \mathcal{G}$ factors uniquely through the equaliser of f and g .

$$\begin{array}{ccccc} G & & & & \\ \downarrow & \searrow h & & & \\ E & \xrightarrow{e} & A & \xrightleftharpoons[f]{g} & B \end{array}$$

Thus this equaliser e must be an isomorphism as \mathcal{G} is detecting. Since $ef = eg$, we must have $f = g$, as required.

Part (ii). Suppose \mathcal{G} is separating, and $f : A \rightarrow B$ is such that every $h : G \rightarrow B$ with $G \in \mathcal{G}$ factors uniquely through f . As \mathcal{C} is balanced, it suffices to show that f is both monic and epic.

If $fg = fh$ for some $g, h : C \rightrightarrows A$, then any $k : G \rightarrow C$ with $G \in \mathcal{G}$ satisfies $gk = hk$, since both are factorisations of $fgk = fhk$ through f .

$$\begin{array}{ccccc} G & \xrightarrow{k} & C & \xrightleftharpoons[g]{h} & A \xrightarrow{f} B \end{array}$$

Since \mathcal{G} is separating, $g = h$. As this is true for all pairs g, h , we must have that f is monic.

Similarly, if $\ell, m : B \rightrightarrows D$ satisfy $\ell f = mf$, then any $n : G \rightarrow B$ with $G \in \mathcal{G}$ satisfies $\ell n = mn$, since it factors through f .

$$\begin{array}{ccccc} & & G & & \\ & \swarrow & \downarrow n & \searrow \ell & \\ A & \xrightarrow{f} & B & \xrightleftharpoons[m]{\ell} & D \end{array}$$

So $\ell = m$, giving that f is epic. \square

Example. (i) In **Gp**, the forgetful functor is represented by \mathbb{Z} . This functor is faithful and reflects isomorphisms, so it is a separator and a detector.

(ii) In **Rng**, the forgetful functor is represented by $\mathbb{Z}[x]$, so similarly $\mathbb{Z}[x]$ is a separator and a detector.

(iii) If \mathcal{C} is small, the set $\{\mathcal{C}(A, -) \mid A \in \text{ob } \mathcal{C}\}$ is a separating and detecting set for $[\mathcal{C}, \mathbf{Set}]$ by the Yoneda lemma.

(iv) In **Top**, the one-point space 1 is a separator, but **Top** has no detecting set. If κ is an infinite cardinal, let X_κ be a discrete space of cardinality κ , and let Y_κ be the same set with the co- $< \kappa$ topology:

$$U \text{ open} \iff U = \emptyset \text{ or } |Y_\kappa \setminus U| < \kappa$$

The identity $X_\kappa \rightarrow Y_\kappa$ is continuous but not a homeomorphism. Given any set \mathcal{G} of spaces, if κ is larger than $|G|$ for all $G \in \mathcal{G}$, then \mathcal{G} cannot detect the fact that the map $X_\kappa \rightarrow Y_\kappa$ is not a homeomorphism.

(v) Let \mathcal{C} be the category whose objects are the (von Neumann) ordinals, and in addition to the identity morphisms, there are precisely two morphisms $f, g : \alpha \rightrightarrows \beta$ when $\alpha < \beta$. We define composition in such a way that $ff = fg = gf = gg = f$. Now, 0 is a detector for \mathcal{C} : it detects that $f, g : 0 \rightrightarrows \alpha$ are not isomorphisms, as neither factors through the other, and it detects that $f, g : \alpha \rightrightarrows \beta$ are not isomorphisms for $0 < \alpha < \beta$ since the morphism $g : 0 \rightarrow \beta$ does not factor through either of them. There is no separating set for \mathcal{C} : for any set of ordinals \mathcal{G} , if $\alpha > \gamma$ for all $\gamma \in \mathcal{G}$, \mathcal{G} cannot separate $f, g : \alpha \rightrightarrows \alpha + 1$.

(vi) **Gp** has no *coseparating* or *codetecting* set of objects. Given any set \mathcal{G} of groups, let H be a simple group with cardinality greater than that of each element of \mathcal{G} . Then the only homomorphisms from H to elements of \mathcal{G} are trivial. In particular, \mathcal{G} cannot detect that the map $H \rightarrow 1$ is not an isomorphism.

2.4. Projectivity

The functors $\mathcal{C}(A, -) : \mathcal{C} \rightarrow \mathbf{Set}$ preserve monomorphisms. They do not, in general, preserve epimorphisms.

Definition. We say that an object P of a locally small category \mathcal{C} is *projective* if $\mathcal{C}(P, -)$ preserves epimorphisms. In more elementary terms, given a diagram

$$\begin{array}{ccc} & P & \\ & \downarrow f & \\ Q & \xrightarrow{g} & R \end{array}$$

I. Category Theory

there exists $h : P \rightarrow Q$ such that $gh = f$.

$$\begin{array}{ccc} & P & \\ h \swarrow & & \downarrow f \\ Q & \xrightarrow{g} & R \end{array}$$

If this holds for all g in some class \mathcal{E} of epimorphisms, we say that P is \mathcal{E} -projective. The dual notion is called *injectivity*.

We will consider the class of pointwise epimorphisms in $[\mathcal{C}, \mathbf{Set}]$; that is, those natural transformations α whose components α_A are surjective.

Corollary. Objects of the form $\mathcal{C}(A, -)$ are pointwise projective in $[\mathcal{C}, \mathbf{Set}]$.

Proof. If $P = \mathcal{C}(A, -)$, an f in the above diagram corresponds to some $\Phi(f) \in RA$ by the Yoneda lemma. But g_A is surjective, so there exists $\Phi(h) \in QA$ mapping to $\Phi(f)$. \square

Proposition. If \mathcal{C} is small, then $[\mathcal{C}, \mathbf{Set}]$ has *enough pointwise projectives*; that is, for any object F there exists a pointwise epimorphism $P \rightarrow F$ with P pointwise projective.

Proof. Let $P = \coprod_{(A,x)} \mathcal{C}(A, -)$ where the disjoint union is taken over all pairs (A, x) with $A \in \text{ob } \mathcal{C}$ and $x \in FA$. Then P is pointwise projective, since the $\mathcal{C}(A, -)$ are. There is a natural transformation $\alpha : P \rightarrow F$ where the (A, x) -indexed term is $\Psi(x) : \mathcal{C}(A, -) \rightarrow F$. This is pointwise epic, since any $x \in FA$ is in the image of $\Psi(x)$. \square

3. Adjunctions

3.1. Definition and examples

Definition. Let \mathcal{C}, \mathcal{D} be categories. An *adjunction* between \mathcal{C} and \mathcal{D} is a pair of functors $F : \mathcal{C} \rightarrow \mathcal{D}$ and $G : \mathcal{D} \rightarrow \mathcal{C}$, together with a bijection between morphisms $FA \rightarrow B$ in \mathcal{D} and $A \rightarrow GB$ in \mathcal{C} , which is natural in both variables A, B . We say that F is the *left adjoint* to G , and that G is the *right adjoint* to F , and write $F \dashv G$.

If \mathcal{C}, \mathcal{D} are locally small, then the naturality condition is that

$$\mathcal{D}(F-, -); \quad \mathcal{C}(-, G-)$$

are naturally isomorphic functors $\mathcal{C}^{\text{op}} \times \mathcal{D} \rightarrow \mathbf{Set}$.

Example. (i) The free group functor $F : \mathbf{Set} \rightarrow \mathbf{Gp}$ is left adjoint to the forgetful functor $U : \mathbf{Gp} \rightarrow \mathbf{Set}$.

$$\mathbf{Gp}(FA, G) \leftrightarrow \mathbf{Set}(A, UG)$$

(ii) The forgetful functor $U : \mathbf{Top} \rightarrow \mathbf{Set}$ has a left adjoint $D : \mathbf{Set} \rightarrow \mathbf{Top}$ which equips each set with its discrete topology.

$$\mathbf{Top}(DX, Y) \leftrightarrow \mathbf{Set}(X, UY)$$

It also has a right adjoint $I : \mathbf{Set} \rightarrow \mathbf{Top}$ which equips each set with its indiscrete topology.

$$\mathbf{Set}(UX, Y) \leftrightarrow \mathbf{Top}(X, IY)$$

(iii) Consider the functor $\text{ob} : \mathbf{Cat} \rightarrow \mathbf{Set}$ which maps each category to its set of objects. It has a left adjoint D which turns each set X into a discrete category in which the objects are elements of X , and the only morphisms are identities. It also has a right adjoint I which turns each set X into an indiscrete category in which the objects are elements of X , and there is exactly one morphism between any two elements of X . In addition, $D : \mathbf{Set} \rightarrow \mathbf{Cat}$ has a left adjoint $\pi_0 : \mathbf{Cat} \rightarrow \mathbf{Set}$, where $\pi_0 \mathcal{C}$ is the set of connected components of $\text{ob } \mathcal{C}$ under the graph induced by its morphisms.

$$\mathbf{Set}(\pi_0 \mathcal{C}, X) \leftrightarrow \mathbf{Cat}(\mathcal{C}, DX); \quad \mathbf{Cat}(DX, \mathcal{C}) \leftrightarrow \mathbf{Set}(X, \text{ob } \mathcal{C}); \quad \mathbf{Set}(\text{ob } \mathcal{C}, X) \leftrightarrow \mathbf{Cat}(\mathcal{C}, IX)$$

Thus we have a chain

$$\pi_0 \dashv D \dashv \text{ob} \dashv I$$

(iv) For any set A , we have a functor $(-) \times A : \mathbf{Set} \rightarrow \mathbf{Set}$. This functor has a right adjoint, which is the functor $\mathbf{Set}(A, -) : \mathbf{Set} \rightarrow \mathbf{Set}$.

$$\mathbf{Set}(B \times A, C) \leftrightarrow \mathbf{Set}(B, \mathbf{Set}(A, C))$$

Applying this bijection is sometimes called *currying* or *λ -conversion*. We say that a category \mathcal{C} with binary products is *cartesian closed* if $(-) \times A : \mathcal{C} \rightarrow \mathcal{C}$ has a right adjoint, written $[A, -]$ or $(-)^A$, for each A . For example, \mathbf{Cat} is cartesian closed, where $\mathcal{D}^{\mathcal{C}} = [\mathcal{C}, \mathcal{D}]$ is the functor category that this notation already refers to.

I. Category Theory

- (v) An equivalence $F : \mathcal{C} \rightarrow \mathcal{D}, G : \mathcal{D} \rightarrow \mathcal{C}$ forms adjunctions both ways: $F \dashv G, G \dashv F$.
- (vi) Let **Idem** be the category of pairs (A, e) where A is a set and e is an idempotent endomorphism $A \rightarrow A$. The morphisms in **Idem** are the maps of sets which commute with the idempotents. We have a functor $F : \mathbf{Set} \rightarrow \mathbf{Idem}$ sending A to $(A, 1_A)$. Consider $G : \mathbf{Idem} \rightarrow \mathbf{Set}$ sending (A, e) to the set of fixed points of e . Then $F \dashv G$ since any morphism $FA \rightarrow (B, e)$ takes values in $G(B, e)$. But also $G \dashv F$, since a morphism $(A, e) \rightarrow FB$ is entirely determined by its action on the fixed points in A under e , because $f(a) = f(ea)$. This is not an equivalence of categories, because G is not faithful. So not all pairs of functors that are adjoint in both directions form an equivalence.
- (vii) Let \mathcal{C} be a category. There is a unique functor $G : \mathcal{C} \rightarrow \mathbf{1}$, where $\mathbf{1}$ is the discrete category on a single object. A left adjoint for G , if it exists, sends the object in $\mathbf{1}$ to an *initial object* I of \mathcal{C} , which is an object with a unique morphism to every object in \mathcal{C} . Dually, a right adjoint sends the object in $\mathbf{1}$ to a *terminal object* T , which is an object with a unique morphism from every object in \mathcal{C} . In **Set**, the empty set is initial, and any singleton is terminal. In **Gp**, the trivial group is initial and terminal.
- (viii) Let $f : A \rightarrow B$ be a function of sets, and let $A' \subseteq A, B' \subseteq B$. Then $Pf(A') \subseteq B'$ if and only if $A' \subseteq P^*f(B')$. Thus $Pf \dashv P^*f$ as functors between PA and PB as posets.
- (ix) Let A, B be sets with a relation $R \subseteq A \times B$. We define mappings $(-)^r : PA \rightarrow PB$ by

$$S^r = \{b \in B \mid \forall a \in S, (a, b) \in R\}$$

and $(-)^{\ell} : PB \rightarrow PA$ by

$$T^{\ell} = \{a \in A \mid \forall b \in T, (a, b) \in R\}$$

These are contravariant functors, and

$$S \subseteq T^{\ell} \iff S \times T \subseteq R \iff T \subseteq S^r$$

We say that $(-)^{\ell}$ and $(-)^r$ are *adjoint on the right*. This pair is called a *Galois connection*.

- (x) The contravariant power-set functor P^* is self-adjoint on the right, since functions $A \rightarrow P^*B$ and $B \rightarrow P^*A$ naturally correspond bijectively to subsets of $A \times B$.
- (xi) The dual vector space functor $(-)^* : \mathbf{Vect}_k \rightarrow \mathbf{Vect}_k$ is self-adjoint on the right, as linear maps $V \rightarrow W^*$ and linear maps $W \rightarrow V^*$ both naturally correspond to bilinear forms on $V \times W$.

3.2. Comma categories

Definition. Let $G : \mathcal{D} \rightarrow \mathcal{C}$ be a functor and $A \in \text{ob } \mathcal{C}$. Then, the *comma category* $(A \downarrow G)$ is the category whose objects are pairs (B, f) where $B \in \text{ob } \mathcal{D}$ and $f : A \rightarrow GB$ in \mathcal{C} , and

3. Adjunctions

whose morphisms $(B, f) \rightarrow (B', f')$ are morphisms $g : B \rightarrow B'$ which commute with f, f' :

$$\begin{array}{ccc} A & \xrightarrow{f} & GB \\ & \searrow f' & \downarrow Gg \\ & & GB' \end{array}$$

Theorem. Let $G : \mathcal{D} \rightarrow \mathcal{C}$ be a functor. Then specifying a left adjoint for G is equivalent to specifying an initial object of the comma categories $(A \downarrow G)$ for each A .

Proof. First, note that an object (B, f) is initial in $(A \downarrow G)$ if and only if for every (B', f') , there is a unique morphism $g : B \rightarrow B'$ such that the following triangle commutes.

$$\begin{array}{ccc} A & \xrightarrow{f} & GB \\ & \searrow f' & \downarrow Gg \\ & & GB' \end{array}$$

Suppose $F \dashv G$. Then let $\eta_A : A \rightarrow GFA$ correspond to the identity 1_{FA} under the adjunction. We show that (FA, η_A) is initial in $(A \downarrow G)$. Indeed, given $f : A \rightarrow GB$, then

$$\begin{array}{ccc} A & \xrightarrow{\eta_A} & GFA \\ & \searrow f & \downarrow Gg \\ & & GB \end{array}$$

commutes if and only if g is the morphism corresponding to f under the adjunction. In particular, for any f , there is a unique such g .

Conversely, suppose (FA, η_A) is initial in $(A \downarrow G)$ for each A . Then we define the action of F on objects by mapping A to FA . We make F into a functor by mapping $f : A \rightarrow A'$ to the unique morphism that makes the following square commute; this exists as (FA, η_A) is initial.

$$\begin{array}{ccc} A & \xrightarrow{\eta_A} & GFA \\ f \downarrow & & \downarrow GFf \\ A' & \xrightarrow{\eta_{A'}} & GFA' \end{array}$$

Functoriality of F follows from the uniqueness of Ff . The bijection between morphisms $f : A \rightarrow GB$ and $g : FA \rightarrow B$ sends f to the unique g giving $(Gg)\eta_A = f$. Naturality of the bijection in A was built in to the definition of F as a functor, and naturality in B is easy. \square

Corollary. Let $F, F' : \mathcal{C} \rightarrow \mathcal{D}$ be left adjoints to $G : \mathcal{D} \rightarrow \mathcal{C}$. Then $F \simeq F'$ in $[\mathcal{C}, \mathcal{D}]$.

I. Category Theory

Proof. (FA, η_A) and $(F'A, \eta'_A)$ are both initial objects in $(A \downarrow G)$, and so there is a unique isomorphism $\alpha_A : (FA, \eta_A) \rightarrow (F'A, \eta'_A)$ in this category. The map $A \mapsto \alpha_A$ is natural, because given $f : A \rightarrow A'$, $\alpha_{A'}(Ff)$ and $(F'f)\alpha_A$ are both morphisms $(FA, \eta_A) \Rightarrow (F'A', \eta'_{A'}, f)$ from an initial object in $(A \downarrow G)$, so must be equal. \square

Lemma. Suppose

$$\mathcal{C} \begin{array}{c} \xrightarrow{F} \\ \xleftarrow{G} \end{array} \mathcal{D} \begin{array}{c} \xrightarrow{H} \\ \xleftarrow{K} \end{array} \mathcal{E}$$

where $F \dashv G$ and $H \dashv K$. Then $HF \dashv GK$.

Proof. We have bijections

$$\mathcal{E}(HFA, C) \leftrightarrow \mathcal{D}(FA, KC) \leftrightarrow \mathcal{C}(A, GKC)$$

which are natural in A and C , so their composite is also natural. \square

Corollary. Suppose the square of functors

$$\begin{array}{ccc} \mathcal{C} & \xrightarrow{F} & \mathcal{D} \\ G \downarrow & & \downarrow H \\ \mathcal{E} & \xrightarrow{K} & \mathcal{F} \end{array}$$

commutes, and all of the functors F, G, H, K have left adjoints F', G', H', K' . Then the square of left adjoints

$$\begin{array}{ccc} \mathcal{C} & \xleftarrow{F'} & \mathcal{D} \\ G' \uparrow & & \uparrow H' \\ \mathcal{E} & \xleftarrow{K'} & \mathcal{F} \end{array}$$

commutes up to natural isomorphism.

This result holds for any shape of diagram, not just a square. The hypothesis can be weakened to only require that the first diagram commutes up to natural isomorphism.

Proof. The two composites $F'H'$ and $G'K'$ are left adjoints to $HF = KG$, so must be naturally isomorphic. \square

3.3. Units and counits

Given an adjunction $F \dashv G$, the proof of the previous theorem demonstrated a naturality square between the morphisms $\eta_A : A \rightarrow GFA$ corresponding to 1_{FA} under the adjunction. We call $\eta : 1_{\mathcal{C}} \rightarrow GF$ the *unit* of the adjunction. Dually, the map $\epsilon : FG \rightarrow 1_{\mathcal{D}}$ is called the *counit* of the adjunction; each $\epsilon_B : FGB \rightarrow B$ corresponds to 1_{GB} .

3. Adjunctions

Theorem. Let $F : \mathcal{C} \rightarrow \mathcal{D}$, $G : \mathcal{D} \rightarrow \mathcal{C}$. Specifying an adjunction $F \dashv G$ is equivalent to specifying natural transformations $\eta : 1_{\mathcal{C}} \rightarrow GF$, $\epsilon : FG \rightarrow 1_{\mathcal{D}}$, satisfying the *triangular identities*

$$\begin{array}{ccc} F & \xrightarrow{F\eta} & FGF \\ & \searrow 1_F & \downarrow \epsilon_F \\ & & F \end{array} \qquad \begin{array}{ccc} G & \xrightarrow{\eta_G} & GFG \\ & \searrow 1_G & \downarrow G\epsilon \\ & & G \end{array}$$

Proof. Suppose we have an adjunction $F \dashv G$. We have seen how to define η and ϵ ; it thus suffices to check the triangular identities. Since they are dual to each other, it suffices to check the first. The morphism ϵ_{FA} corresponds under the adjunction to 1_{GFA} , so by naturality, the composite $\epsilon_{FA}(F\eta_A)$ corresponds to $1_{GFA}\eta_A = \eta_A$. But 1_{FA} corresponds to η_A , giving the commutative triangle $\epsilon_{FA}(F\eta_A) = 1_{FA}$.

Conversely, suppose η and ϵ are natural transformations satisfying the triangular identities. We map $f : A \rightarrow GB$ to the composite $\Phi(f)$ given by

$$FA \xrightarrow{Ff} FGB \xrightarrow{\epsilon_B} B$$

and $g : FA \rightarrow B$ to the composite $\Psi(g)$ given by

$$A \xrightarrow{\eta_A} GFA \xrightarrow{Gg} GB$$

These assignments are natural in A and B as η and ϵ are natural transformations. Thus it suffices to show $\Psi\Phi$ and $\Phi\Psi$ are the relevant identity maps; again they are dual so it suffices to show $\Psi\Phi(f) = f$. $\Psi\Phi(f)$ is the composite

$$A \xrightarrow{\eta_A} GFA \xrightarrow{GFf} GFGB \xrightarrow{G\epsilon_B} GB$$

which by naturality of η is equal to

$$A \xrightarrow{f} GB \xrightarrow{\eta_{GB}} GFGB \xrightarrow{G\epsilon_B} GB$$

which is equal to f by the triangular identity. \square

Recall that an equivalence of categories consisted of isomorphisms $\alpha : 1_{\mathcal{C}} \rightarrow GF$ and $\beta : FG \rightarrow 1_{\mathcal{D}}$. These isomorphisms may not satisfy the triangular identities, but we can always choose α and β in such a way that these identities hold.

Proposition. Let (F, G, α, β) be an equivalence of categories. Then there exist natural isomorphisms $\alpha' : 1_{\mathcal{C}} \rightarrow GF$ and $\beta' : FG \rightarrow 1_{\mathcal{D}}$ which satisfy the triangular identities. In particular, $F \dashv G \dashv F$.

Proof. We will set $\alpha' = \alpha$, and construct β' to be the composite

$$FG \xrightarrow{(FG\beta)^{-1}} FGFG \xrightarrow{(F\alpha_G)^{-1}} FG \xrightarrow{\beta} 1_{\mathcal{D}}$$

I. Category Theory

Note that $FG\beta = \beta_{FG}$, since

$$\begin{array}{ccc} FGFG & \xrightarrow{FG\beta} & FG \\ \beta_{FG} \downarrow & & \downarrow \beta \\ FG & \xrightarrow{\beta} & 1_{\mathcal{D}} \end{array}$$

commutes by naturality of β . Note also that β is monic. Dually, note that $GF\alpha = \alpha_{GF}$. For the triangular identities, consider the diagrams

$$\begin{array}{ccccc} F & \xrightarrow{F\alpha} & FGF & \xrightarrow{(\beta_{FGF})^{-1}} & FGF GF \\ & \searrow 1_F & \downarrow (F\alpha)^{-1} & & \downarrow (F\alpha_{GF})^{-1} = (FGF\alpha)^{-1} \\ & & F & \xrightarrow{\beta_F} & FGF \\ & & & \searrow 1_F & \downarrow \beta_F \\ & & & & F \end{array}$$

and

$$\begin{array}{ccccc} G & \xrightarrow{\alpha_G} & GFG & \xrightarrow{(GF\beta)^{-1}} & GFG FG \\ & \searrow 1_G & \downarrow \alpha_G^{-1} & & \downarrow (GF\alpha_G)^{-1} = (\alpha_{GFG})^{-1} \\ & & G & \xrightarrow{(G\beta)^{-1}} & GFG \\ & & & \searrow 1_G & \downarrow G\beta \\ & & & & G \end{array}$$

where the squares commute by naturality of β and α respectively. Thus α', β' are the unit and counit of an adjunction $F \dashv G$ as required. Similarly, $(\beta')^{-1}, (\alpha')^{-1}$ are the unit and counit of an adjunction $G \dashv F$. \square

Lemma. Let $F \dashv G$ be an adjunction with counit $\epsilon : FG \rightarrow 1_{\mathcal{D}}$. Then

- (i) ϵ is pointwise epimorphic if and only if G is faithful;
- (ii) ϵ is a (pointwise) isomorphism if and only if G is full and faithful.

Proof. Part (i). Given $g : B \rightarrow B'$ in \mathcal{D} , the composite $g\epsilon_B$ corresponds under the adjunction to $Gg : GB \rightarrow GB'$. Thus for morphisms g with specified domain and codomain, the map $g \mapsto g\epsilon_B$ is injective if and only if the action of G is injective. This is true for all B and B' if and only if ϵ is pointwise epimorphic, if and only if G is faithful.

Part (ii). Similarly, G is full and faithful if and only if the map $g \mapsto g\epsilon_B$ is a bijection on morphisms with specified domain and codomain. This clearly holds if ϵ_B is an isomorphism for all B . Conversely, if the condition holds, there is a unique map $g : B \rightarrow FGB$ such that $\epsilon_B g = 1_B$. Then $\epsilon_B g\epsilon_B = \epsilon_B$, so $g\epsilon_B$ and 1_{FGB} have the same composite with ϵ_B , so they are equal. \square

3.4. Reflections

Definition. An adjunction $F \dashv G$ is called a *reflection* if the counit is an isomorphism. Dually, it is called a *coreflection* if the unit is an isomorphism. A full subcategory is called *reflective* if the inclusion functor has a left adjoint; in this case the adjunction is a reflection.

Remark. If $F \dashv G$ is a reflection, then $G : \mathcal{D} \rightarrow \mathcal{C}$ induces an equivalence of categories between \mathcal{D} and the full subcategory of \mathcal{C} on the objects in the image of G . This subcategory is reflective.

If $\mathcal{D} \subseteq \mathcal{C}$ is a reflective subcategory, there is intuitively a best possible way to get *into* \mathcal{D} from some object in \mathcal{C} . The left adjoint sends an object in \mathcal{C} to its ‘best approximation’ in \mathcal{D} . If \mathcal{D} is coreflective, there is a best possible way to get *out of* \mathcal{D} to some object in \mathcal{C} .

Example. (i) **AbGp** is reflective in **Gp**; the left adjoint to the inclusion map sends a group G to its abelianisation $G^{\text{ab}} = G/H$, the quotient of G by its commutator subgroup

$$H = \{aba^{-1}b^{-1} \mid a, b \in G\} \trianglelefteq G$$

Note that any homomorphism $G \rightarrow A$ where A is abelian factors uniquely through the quotient map $G \rightarrow G^{\text{ab}}$, giving the adjunction as required.

- (ii) Recall that an abelian group is called *torsion* if all of its elements have finite order, and *torsion-free* if all of its nonzero elements have infinite order. For an abelian group A , its set of torsion elements forms a subgroup A_t , which is a torsion group. Any homomorphism from a torsion group to A must factor through A_t . Thus A_t is the coreflection of A in the category of torsion abelian groups, and A/A_t is the reflection of A in the category of torsion-free abelian groups.
- (iii) The full subcategory **KHaus** of compact Hausdorff spaces is reflective in the category **Top** of topological spaces. The left adjoint to the inclusion map is the *Stone–Čech compactification* functor β . We will construct this functor using the special adjoint functor theorem, which is explored in the next section.
- (iv) Recall that a subset C of a topological space X is called *sequentially closed* if for every sequence $x_n \in C$ converging to a limit $x \in X$, we have $x \in C$. We say that X is a *sequential space* if all sequentially closed subsets are closed. The full subcategory **Seq** of sequential spaces is coreflective in **Top**. Given a space X , let X_s denote the same set, but where the topology is such that all sequentially closed sets are also taken to be closed. The identity map $X_s \rightarrow X$ is continuous, and forms the counit of the adjunction.
- (v) The category **Preord** of preorders is reflective in **Cat**. The left adjoint maps a category \mathcal{C} to the quotient category \mathcal{C}/\sim where \sim identifies all parallel pairs of morphisms.
- (vi) Let X be a topological space. Then the poset ΩX of open sets in X is coreflective in the poset PX , since if U is open and A is an arbitrary subset of X , then $U \subseteq A$ if and only if $U \subseteq A^\circ$. Thus the interior operator $(-)^\circ$ is right adjoint to the inclusion $\Omega X \rightarrow PX$.

I. Category Theory

Dually, the poset of closed sets is reflective in PX ; the closure operator $\overline{(-)}$ is left adjoint to the inclusion.

4. Limits


4.1. Cones over diagrams

To formally define limits and colimits, we first need to define more precisely what is meant by a diagram in a category.

Definition. Let J be a category, which will almost always be small, and often finite. A *diagram* of shape J in a category \mathcal{C} is a functor $D : J \rightarrow \mathcal{C}$.


We call the objects $D(j)$ the *vertices* of the diagram, and the morphisms $D(\alpha)$ the *edges* of the diagram.

Example. Let J be the finite category



A diagram of shape J in \mathcal{C} is exactly a commutative square in \mathcal{C} . The diagonal arrow is required to make J into a category.

Example. Let J be the finite category



Then a diagram of shape J in \mathcal{C} is a square of objects in \mathcal{C} whose morphisms may or may not commute.

Definition. Let D be a diagram of shape J in \mathcal{C} . A *cone* over D consists of an object $A \in \text{ob } \mathcal{C}$ called the *apex* of the cone, together with morphisms $\lambda_j : A \rightarrow D(j)$ called the *legs* of the cone, such that all triangles of the following form commute.

$$\begin{array}{ccc}
 & A & \\
 \lambda_j \swarrow & & \searrow \lambda_{j'} \\
 D(j) & \xrightarrow{D(\alpha)} & D(j')
 \end{array}$$

We can define the notion of a morphism between cones.

Definition. Let $(A, \lambda_j), (B, \mu_j)$ be cones over a diagram D of shape J in \mathcal{C} . Then a *morphism of cones* is a morphism $f : A \rightarrow B$ such that all triangles of the following form commute.

$$\begin{array}{ccc}
 A & \xrightarrow{f} & B \\
 \lambda_j \searrow & & \swarrow \mu_j \\
 & D(j) &
 \end{array}$$

I. Category Theory

This makes the class of cones over a diagram D into a category, which will be denoted $\text{Cone}(D)$.

Remark. A cone over a diagram D with apex A is the same as a natural transformation from the constant diagram ΔA to D , as we can expand the commutative triangles into the following form.

$$\begin{array}{ccc} A & \xrightarrow{1_A} & A \\ \lambda_j \downarrow & & \downarrow \lambda_{j'} \\ D(j) & \xrightarrow{D(\alpha)} & D(j') \end{array}$$

Note that Δ is a functor $\mathcal{C} \rightarrow [J, \mathcal{C}]$, and thus $\text{Cone}(D)$ is exactly the comma category $(\Delta \downarrow D)$.

4.2. Limits

Definition. A *limit* for a diagram D of shape J in \mathcal{C} is a terminal object in the category of cones over D . Dually, a *colimit* for D is an initial object in the category of cones under D .

A cone under a diagram is often called a *cocone*.

Remark. Using the fact that $\text{Cone}(D) = (\Delta \downarrow D)$ where $\Delta : \mathcal{C} \rightarrow [J, \mathcal{C}]$, the category \mathcal{C} has limits for all diagrams of shape J if and only if Δ has a right adjoint.

Example. (i) If J is the empty category, there is a unique diagram D of shape J in any category \mathcal{C} . Thus, a cone over this diagram is just an object in \mathcal{C} , and morphisms of cones are just morphisms in \mathcal{C} . In particular, $\text{Cone}(D) \cong \mathcal{C}$, so a limit for D is a terminal object in \mathcal{C} . Dually, a colimit of the empty diagram is an initial object.

(ii) Let J be the discrete category with two objects. A diagram of shape J in \mathcal{C} is thus a pair of objects. A cone over this diagram is a *span*.

$$\begin{array}{ccc} & C & \\ \swarrow & & \searrow \\ A & & B \end{array}$$

A limit cone is precisely a categorical product $A \times B$.


$$\begin{array}{ccc} & A \times B & \\ \swarrow \pi_1 & & \searrow \pi_2 \\ A & & B \end{array}$$

Similarly, the colimit for a pair of objects is a categorical coproduct $A + B$.

(iii) If J is any discrete category, a diagram of shape J is a family of objects A_j in \mathcal{C} indexed by the objects of J . Limits and colimits over this diagram are products and coproducts of the A_j .


4. Limits

- (iv) If J is the category $\bullet \rightrightarrows \bullet$, a diagram of shape J is a parallel pair of morphisms $f, g : A \rightrightarrows B$. A cone over such a parallel pair is




satisfying $fh = k = gh$. Equivalently, it is a morphism $h : C \rightarrow A$ satisfying $fh = gh$. Thus, a limit is an equaliser, and dually, a colimit is a coequaliser.


- (v) Let J be the category



A diagram of shape J is thus a cospan in \mathcal{C} .



A cone over this diagram is



where $\ell = fh = gk$ is redundant. Thus a cone is a span that completes the commutative square. A limit for the cospan is the universal way to complete this commutative square, which is called a *pullback* of f and g . Dually, colimits of spans are called *pushouts*.

If any category \mathcal{C} has binary products and equalisers, we can construct all pullbacks. First, we construct the product $A \times B$, then we form the equaliser of $f\pi_1, g\pi_2 : A \times B \rightrightarrows C$. This yields the pullback.

- (vi) Let M be the two-element monoid $\{1, e\}$ with $e^2 = e$. A diagram of shape M in a category \mathcal{C} is an object of \mathcal{C} equipped with an idempotent endomorphism. A cone over this diagram is a morphism $f : B \rightarrow A$ such that $ef = f$. A limit (respectively colimit) is the monic (respectively epic) part of a splitting of e . This is because the pair $(e, 1_A)$ has an equaliser if and only if e splits.
- (vii) Let \mathbb{N} be the poset category of the natural numbers. A diagram of shape \mathbb{N} is a *direct sequence* of objects, which consists of objects A_0, A_1, \dots and morphisms $f_i : A_i \rightarrow A_{i+1}$.

I. Category Theory

A colimit for this diagram is a *direct limit*, which consists of an object A_∞ and morphisms $g_i : A_i \rightarrow A_\infty$ which are compatible with the f_i . Dually, an *inverse sequence* is a diagram of shape \mathbb{N}^{op} , and a limit for this diagram is called an *inverse limit*. For example, an infinite-dimensional CW-complex X is the direct limit of its n -dimensional skeletons in **Top**. The ring of p -adic integers is the limit of the inverse sequence defined by $A_n = \mathbb{Z}/p^n\mathbb{Z}$ in **Rng**.

Lemma. Let \mathcal{C} be a category.

- (i) If \mathcal{C} has equalisers and all small products, then \mathcal{C} has all small limits.
- (ii) If \mathcal{C} has equalisers and all finite products, then \mathcal{C} has all finite limits.
- (iii) If \mathcal{C} has pullbacks and a terminal object, then \mathcal{C} has all finite limits.

Note that the empty product is implicitly included in (i) and (ii). A terminal object is a product over no factors.

Proof. Parts (i) and (ii). We prove (i) and (ii) in the same way. We will first construct the product P of the $D(j)$ for each $j \in \text{ob } J$. Then, we will use an equaliser to construct the subobject E of P that simultaneously satisfies all of the equations required for E to be the apex of a cone. The fact that we have used an equaliser will show that this is a limit cone.

Let $D : J \rightarrow \mathcal{C}$ be a diagram. We form the products

$$P = \prod_{j \in \text{ob } J} D(j); \quad Q = \prod_{\alpha \in \text{mor } J} D(\text{cod } \alpha)$$

These are small or finite as required. Using the universal property of the product on Q , we have morphisms $f, g : P \rightrightarrows Q$ defined by

$$\pi_\alpha f = \pi_{\text{cod } \alpha} : P \rightarrow D(\text{cod } \alpha); \quad \pi_\alpha g = D(\alpha)\pi_{\text{dom } \alpha} : P \rightarrow D(\text{cod } \alpha)$$

For $\alpha : j \rightarrow j'$ in D , these morphisms are represented by

$$\begin{array}{ccc} P & \xrightarrow{\quad f \quad} & Q \\ \pi_{j'} \searrow & & \downarrow \pi_\alpha \\ & & D(j') \end{array} \qquad \begin{array}{ccc} P & \xrightarrow{\quad g \quad} & Q \\ \pi_j \downarrow & & \downarrow \pi_\alpha \\ D(j) & \xrightarrow{\quad D(\alpha) \quad} & D(j') \end{array}$$

Let $e : E \rightarrow P$ be an equaliser for f and g , and define $\lambda_j = \pi_j e : E \rightarrow D(j)$. Then for each

$\alpha : j \rightarrow j'$, the following diagram commutes.

$$\begin{array}{ccc}
 & E & \\
 \lambda_j \swarrow & \downarrow e & \searrow \lambda_{j'} \\
 & P & \\
 \pi_j \swarrow & & \searrow \pi_{j'} \\
 D(j) & \xrightarrow{D(\alpha)} & D(j')
 \end{array}$$

Therefore, these morphisms form a cone. Given any cone $(A, (\mu_j)_{j \in \text{ob } J})$ over D , we have a unique $\mu : A \rightarrow P$ with $\pi_j \mu = \mu_j$ for all j . Then,

$$\pi_\alpha f \mu = \mu_{\text{cod } \alpha} = D(\alpha) \mu_{\text{dom } \alpha} = \pi_\alpha g \mu$$

for all α , so μ factors uniquely through e .

Part (iii). We show that the hypotheses of (iii) imply those of (ii). If 1 is the terminal object, we form the pullback of the span

$$\begin{array}{ccc}
 & A & \\
 & \downarrow & \\
 B & \longrightarrow & 1
 \end{array}$$

This has the universal property of the product $A \times B$, so \mathcal{C} has binary products and hence all finite products by induction. To construct the equaliser of $f, g : A \rightrightarrows B$, we consider the pullback of

$$\begin{array}{ccc}
 & A & \\
 & \downarrow (1_A, f) & \\
 A & \xrightarrow{(1_A, g)} & A \times B
 \end{array}$$

Any cone over this diagram has its two legs $C \rightrightarrows A$ equal, so a pullback is an equaliser for f, g . \square

Definition. A category is called *complete* if it has all small limits, and *cocomplete* if it has all small colimits.

Example. The categories **Set**, **Gp**, **Top** are complete and cocomplete.

4.3. Preservation and creation

Definition. Let $G : \mathcal{D} \rightarrow \mathcal{C}$ be a functor. We say that G

- (i) *preserves* limits of shape J if whenever $D : J \rightarrow \mathcal{D}$ is a diagram with limit cone $(L, (\lambda_j)_{j \in \text{ob } J})$, the cone $(GL, (G\lambda_j)_{j \in \text{ob } J})$ is a limit for GD ;

I. Category Theory

- (ii) *reflects* limits of shape J if whenever $D : J \rightarrow \mathcal{D}$ is a diagram and $(L, (\lambda_j)_{j \in \text{ob } J})$ is a cone such that $(GL, (G\lambda_j)_{j \in \text{ob } J})$ is a limit for GD , then $(L, (\lambda_j)_{j \in \text{ob } J})$ is a limit for D ;
- (iii) *creates* limits of shape J if whenever $D : J \rightarrow \mathcal{D}$ is a diagram with limit cone $(M, (\mu_j)_{j \in \text{ob } J})$ for GD in \mathcal{C} , there exists a cone $(L, (\lambda_j)_{j \in \text{ob } J})$ over D such that $(GL, (G\lambda_j)_{j \in \text{ob } J}) \cong (M, (\mu_j)_{j \in \text{ob } J})$ in $\text{Cone}(GD)$, and any such cone is a limit for D .

We typically assume in (i) that \mathcal{D} has all limits of shape J , and we assume in (ii) and (iii) that \mathcal{C} has all limits of shape J . With these assumptions, G creates limits of shape J if and only if G preserves and reflects limits, and \mathcal{D} has all limits of shape J .

Corollary. In any of the statements of the previous lemma, we can replace both instances of ‘ \mathcal{C} has’ by either ‘ \mathcal{D} has and $G : \mathcal{D} \rightarrow \mathcal{C}$ preserves’ or ‘ \mathcal{C} has and $G : \mathcal{D} \rightarrow \mathcal{C}$ creates’.

Example. (i) The forgetful functor $U : \mathbf{Gp} \rightarrow \mathbf{Set}$ creates all small limits. It does not preserve colimits, as in particular it does not preserve coproducts.

(ii) The forgetful functor $U : \mathbf{Top} \rightarrow \mathbf{Set}$ preserves all small limits and colimits, but does not reflect them, as we can retopologise the apex of a limit cone.

(iii) The inclusion $\mathbf{AbGp} \rightarrow \mathbf{Gp}$ reflects coproducts, but does not preserve them. A free product of two groups G, H is always nonabelian, except for the case where either G or H is the trivial group, but the coproduct of the trivial group with H is isomorphic to H in both categories.

Lemma. Suppose \mathcal{D} has limits of shape J . Then, for any \mathcal{C} , the functor category $[\mathcal{C}, \mathcal{D}]$ also has limits of shape J , and the forgetful functor $[\mathcal{C}, \mathcal{D}] \rightarrow \mathcal{D}^{\text{ob } \mathcal{C}}$ creates them.

Proof. Given a diagram $D : J \rightarrow [\mathcal{C}, \mathcal{D}]$, we can regard it as a functor $D : J \times \mathcal{C} \rightarrow \mathcal{D}$, so for a fixed object in \mathcal{C} , we obtain a diagram $D(-, A)$ of shape J in \mathcal{D} , which has a limit $(LA, (\lambda_{j,A})_{j \in \text{ob } J})$. Given any $f : A \rightarrow B$ in \mathcal{C} , the composites

$$LA \xrightarrow{\lambda_{j,A}} D(j, A) \xrightarrow{D(j, f)} D(j, B)$$

form a cone over $D(-, B)$, and so factor uniquely through its limit LB . Thus we obtain $Lf : LA \rightarrow LB$. This is functorial because Lf is unique with this property. This is the unique lifting of $(LA)_{A \in \text{ob } \mathcal{C}}$ to an object of $[\mathcal{C}, \mathcal{D}]$ which makes the $\lambda_{j,-}$ into natural transformations. It is a limit cone in $[\mathcal{C}, \mathcal{D}]$: given any cone in $[\mathcal{C}, \mathcal{D}]$ with apex M and legs $(\mu_{j,-})_{j \in \text{ob } J}$ over D , the $\mu_{j,A}$ form a cone over $D(-, A)$, so we obtain a unique $\nu_A : MA \rightarrow LA$ such that $\lambda_{j,A} \nu_A = \mu_{j,A}$ for all A . The ν_A form a natural transformation $M \rightarrow L$, because for any $f : A \rightarrow B$ in \mathcal{C} , the two paths $\nu_B(Mf), (Lf)\nu_A : MA \rightarrow LB$ are factorisations of the same cone over $D(-, B)$ through its limit, so must be equal. \square

Remark. Note that $f : A \rightarrow B$ is monic if and only if

$$\begin{array}{ccc} A & \xrightarrow{1_A} & A \\ 1_A \downarrow & & \downarrow f \\ A & \xrightarrow{f} & B \end{array}$$

is a pullback square. Thus, if \mathcal{D} has pullbacks, any monomorphism in $[\mathcal{C}, \mathcal{D}]$ is a pointwise monomorphism, because the pullback in $[\mathcal{C}, \mathcal{D}]$ is constructed pointwise by the previous lemma. In particular, the monomorphisms and epimorphisms in $[\mathcal{C}, \mathbf{Set}]$ are precisely the pointwise monomorphisms and pointwise epimorphisms respectively.

4.4. Interaction with adjunctions

Lemma. Let $G : \mathcal{D} \rightarrow \mathcal{C}$ be a functor with a left adjoint. Then G preserves all limits which exist in \mathcal{D} .

Proof 1. In this proof, we will assume that \mathcal{C}, \mathcal{D} both have all limits of shape J . If $F \dashv G$, then the diagram

$$\begin{array}{ccc} \mathcal{C} & \xrightarrow{F} & \mathcal{D} \\ \Delta \downarrow & & \downarrow \Delta \\ [J, \mathcal{C}] & \xrightarrow{[J, F]} & [J, \mathcal{D}] \end{array}$$

commutes. All of the functors in this diagram have right adjoints, so the diagram

$$\begin{array}{ccc} \mathcal{C} & \xleftarrow{G} & \mathcal{D} \\ \lim_J \uparrow & & \uparrow \lim_J \\ [J, \mathcal{C}] & \xleftarrow{[J, G]} & [J, \mathcal{D}] \end{array}$$

commutes up to natural isomorphism, where \lim_J sends a diagram of shape J to the apex of its limit cone. But this is exactly the statement that G preserves limits. \square

Proof 2. In this proof, we will not assume that \mathcal{C} has limits of any kind, and only assume a single diagram $D : J \rightarrow \mathcal{D}$ has a limit cone $(L, (\lambda_j)_{j \in \text{ob } J})$ over it. Given any cone over GD with apex A and legs $\mu_j : A \rightarrow GD(j)$, the legs correspond under the adjunction to morphisms $\bar{\mu}_j : FA \rightarrow D(j)$, which form a cone over D by naturality of the adjunction. We obtain a unique factorisation $\bar{\mu} : FA \rightarrow L$ with $\lambda_j \bar{\mu} = \bar{\mu}_j$ for all j , or equivalently, $(G\lambda_j)\mu = \mu_j$, where $\mu : A \rightarrow GL$ corresponds to $\bar{\mu}$ under the adjunction. \square

Suppose that \mathcal{D} has and $G : \mathcal{D} \rightarrow \mathcal{C}$ preserves all limits. The adjoint functor theorems say that G has a left adjoint, under various assumptions.

I. Category Theory

Lemma. Suppose that \mathcal{D} has and $G : \mathcal{D} \rightarrow \mathcal{C}$ preserves limits of shape J . Then for any $A \in \text{ob } \mathcal{C}$, the category $(A \downarrow G)$ has limits of shape J , and the forgetful functor $U : (A \downarrow G) \rightarrow \mathcal{D}$ creates them.

Proof. Let $D : J \rightarrow (A \downarrow G)$ be a diagram. We write each $D(j)$ as $(UD(j), f_j)$ where $f_j : A \rightarrow GUD(j)$. Let $(L, (\lambda_j)_{j \in \text{ob } J})$ be a limit for UD in \mathcal{D} . By assumption, $(GL, (G\lambda_j)_{j \in \text{ob } J})$ is a limit for GUD in \mathcal{C} . But the edges of D are morphisms in $(A \downarrow G)$, so the f_j form a cone over GUD . Thus, we obtain a unique factorisation $f : A \rightarrow GL$ such that $(G\lambda_j)f = f_j$ for all j . In other words, we have a unique lifting of L to an object (L, f) of $(A \downarrow G)$ which makes the λ_j into a cone over D with apex (L, f) . Any cone over D with apex (M, g) becomes a cone over UD with apex M by forgetting the structure map, so we get a unique $h : M \rightarrow L$, and this becomes a morphism in $(A \downarrow G)$ as both $(Gh)g$ and f are factorisations through L of the same cone over UD . \square

Lemma. Let \mathcal{C} be a category. Specifying an initial object of \mathcal{C} is equivalent to specifying a limit for the identity functor $1_{\mathcal{C}} : \mathcal{C} \rightarrow \mathcal{C}$, considered as a diagram of shape \mathcal{C} in \mathcal{C} .

Proof. First, suppose we have an initial object I in \mathcal{C} . Then the unique morphisms $I \rightarrow A$ form a cone over $1_{\mathcal{C}}$, and it is a limit, because for any other cone $(B, (\lambda_A : B \rightarrow A))$, then λ_I is the unique factorisation as required. Conversely, suppose $(I, (\lambda_A : I \rightarrow A))$ is a limit for $1_{\mathcal{C}}$. Then certainly I is *weakly initial*: it has at least one morphism to any other object, given by λ_A . For any morphism $f : I \rightarrow A$, it is an edge of the diagram, so $f\lambda_I = \lambda_A$, so it suffices to show that λ_I is the identity morphism. Using the same equation with $f = \lambda_A$, we obtain $\lambda_A\lambda_I = \lambda_A$, so λ_I is a factorisation of the limit cone through itself. As this factorisation must be unique, we must have $\lambda_I = 1_I$. \square

Proposition (primitive adjoint functor theorem). If \mathcal{D} has and $G : \mathcal{D} \rightarrow \mathcal{C}$ preserves all limits, then G has a left adjoint.

Proof. The categories $(A \downarrow G)$ have all limits, and in particular they have initial objects, so G has a left adjoint. \square

4.5. General adjoint functor theorem

Theorem (general adjoint functor theorem). Suppose \mathcal{D} is complete and locally small. Then a functor $G : \mathcal{D} \rightarrow \mathcal{C}$ has a left adjoint if and only if G preserves small limits and satisfies the *solution-set condition*: given any $A \in \text{ob } \mathcal{C}$, there is a set $\{f_i : A \rightarrow GB_i\}_{i \in I}$ such that every $f : A \rightarrow GB$ factors as

$$A \xrightarrow{f_i} GB_i \xrightarrow{Gg} GB$$

for some $i \in I$ and $g : B_i \rightarrow B$. This set I is called a solution-set at A .

The solution-set condition can be equivalently phrased as the assertion that the categories $(A \downarrow G)$ all have *weakly initial* sets of objects: every object of $(A \downarrow G)$ admits a morphism from a member of the solution set.

Proof. If $F \dashv G$, then G preserves all limits that exist in its domain, so in particular it preserves small limits, and $\{\eta_A : A \rightarrow GFA\}$ is a solution-set at A for any A . Now suppose $A \in \text{ob } \mathcal{C}$. Then $(A \downarrow G)$ is complete, and is locally small as morphisms $(B, f) \rightarrow (B', f')$ in $(A \downarrow G)$ are a subset of $\mathcal{D}(B, B')$. We must then show that if \mathcal{A} is complete and locally small and has a weakly initial set of objects $\{S_i \mid i \in I\}$, then it has an initial object; then, setting $\mathcal{A} = (A \downarrow G)$ and using the solution-set as the weakly initial set, the result follows.

First, we form the product $P = \prod_{i \in I} S_i$. The set $\{P\}$ is weakly initial since we have morphisms $\pi_i : P \rightarrow S_i$ for all i . Now consider the diagram $P \rightrightarrows P$ whose edges are all endomorphisms of P . By assumption, let $i : I \rightarrow P$ be a limit for this diagram; this is an equaliser over a family of morphisms. Then I is weakly initial. For a parallel pair $f, g : I \rightrightarrows C$, we have an equaliser $e : E \rightarrow I$, and can choose some $h : P \rightarrow E$. Then we have the endomorphisms ieh and 1_P of P . Thus $iehi = 1_P i = i$, but i is monic, so $ehi = 1_I$. Hence e is a split epimorphism, and hence $f = g$. \square

Example. (i) Consider the forgetful functor $U : \mathbf{Gp} \rightarrow \mathbf{Set}$. Note that \mathbf{Gp} is complete and locally small, and U creates small limits so in particular it preserves them. Given a set A , any function $f : A \rightarrow UG$ can be factored as

$$A \longrightarrow UG' \longrightarrow UG$$

where G' is the subgroup generated by $\{f(a) \mid a \in A\}$. Note that the cardinality of G' is at most $\max(\aleph_0, |A|)$, so we can fix a set B of this cardinality and consider all possible subsets of B , all possible group structures on those sets, and all possible functions $A \rightarrow B'$; these form a solution-set at A . Hence, free groups exist. Note that the cardinality bound on G' requires most of the technology needed to explicitly construct free groups.

(ii) Let \mathbf{CLat} be the category of complete lattices. The forgetful functor $U : \mathbf{CLat} \rightarrow \mathbf{Set}$ creates all small limits; this can be seen in the same way as was shown with the forgetful functor $\mathbf{Gp} \rightarrow \mathbf{Set}$. In 1964, A. Hales proved that there are arbitrarily large complete lattices with only three generators. Hence U has no solution set at $A = \{a, b, c\}$. Note that U is representable, or equivalently, $(1 \downarrow U)$ has an initial object. If \mathbf{CLat} had all coproducts, we would be able to form initial objects for $(A \downarrow U)$, as every set is a coproduct of singletons. But \mathbf{CLat} does not have even finite coproducts.

4.6. Special adjoint functor theorem

Definition. Let $A \in \text{ob } \mathcal{C}$. A *subobject* of A is a monomorphism with codomain A ; dually, a *quotient* of A is an epimorphism with domain A . The subobjects of A in \mathcal{C} form a preorder $\text{Sub}_{\mathcal{C}}(A)$ by setting $m \leq m'$ when m factors through m' . \mathcal{C} is *well-powered* if $\text{Sub}_{\mathcal{C}}(A)$ is equivalent to a (small) poset for any A . Dually, we say \mathcal{C} is *well-copowered*.

I. Category Theory

Example. **Set** is well-powered, since every monomorphism is isomorphic to a subset inclusion; the power-set axiom encodes this fact. **Set** is also well-copowered, because quotients correspond to equivalence relations up to isomorphism, there is only a set of equivalence relations on a given object A .

Lemma. Let

$$\begin{array}{ccc} P & \xrightarrow{h} & A \\ k \downarrow & & \downarrow f \\ B & \xrightarrow{g} & C \end{array}$$

be a pullback square where f is monic. Then k is also monic.

Informally, monomorphisms are stable under pullback.

Proof. Let $\ell, m : D \rightrightarrows P$ be such that $k\ell = km$. Then $fhl = gk\ell = gkm = fhm$, but f is a monomorphism, so $hl = hm$.

$$\begin{array}{ccccc} & & D & & \\ & \searrow \ell & & \searrow m & \\ & & P & \xrightarrow{h} & A \\ & & k \downarrow & & \downarrow f \\ & & B & \xrightarrow{g} & C \end{array}$$

So ℓ and m are both factorisations of $(h\ell, k\ell)$ through the pullback, so $\ell = m$. \square


Theorem. Let \mathcal{C}, \mathcal{D} be locally small, and suppose that \mathcal{D} is complete, well-powered, and has a coseparating set. Then a functor $G : \mathcal{D} \rightarrow \mathcal{C}$ preserves all small limits if and only if it has a left adjoint.

Proof. As above, any functor with a left adjoint preserves all limits that exist. For the other direction, fix an object A and consider the category $(A \downarrow G)$, which is complete and locally small. Note that the forgetful functor $(A \downarrow G) \rightarrow \mathcal{D}$ preserves monomorphisms, because it preserves pullbacks. Thus, one can show that $(A \downarrow G)$ is well-powered, because the subobjects of a given object (B, f) are the monomorphisms $m : B' \rightarrow B$ for which f factors through Gm . If $\{S_i\}_{i \in I}$ is a coseparating set for \mathcal{D} , we have a coseparating set for $(A \downarrow G)$ by taking the set of all $f : A \rightarrow GS_i$ with $i \in I$; this is a set by local smallness. This is coseparating, because given $h, k : (B, g) \rightrightarrows (B', g')$ with $h \neq k$, there is a morphism $\ell : B' \rightarrow S_i$ with $\ell h \neq \ell k$, and ℓ is a morphism $(B', g') \rightarrow (S_i, (G\ell)g')$ in $(A \downarrow G)$.

It remains to show that there is an initial object in a category \mathcal{A} if it is complete, locally small, well-powered, and has a coseparating set $\{S_i\}_{i \in I}$. First, we form the product

$$P = \prod_{i \in I} S_i$$

and consider the diagram



whose edges are representative monomorphisms for each isomorphism class of subobjects of P . Let I be the apex of a limit cone for this ‘wide pullback’. The legs of the cone are monomorphisms, using the same argument as was described for pullbacks. In particular, the composite maps $I \rightarrow P$ are monomorphisms, so I is a subobject of P . But by construction, it factors through every subobject of P , so is a minimal subobject of P .

It remains to show that I is initial. Note that if $f, g : I \rightrightarrows A$ were different monomorphisms, their equaliser $e : E \rightarrow I$ would yield a subobject of P contained in $I \rightarrow P$, so it would be an isomorphism, giving $f = g$. For an arbitrary object $A \in \text{ob } \mathcal{A}$, form the product

$$Q = \prod_{(i,f)} S_i; \quad f : A \rightarrow S_i$$

and define $g : A \rightarrow Q$ by

$$\pi_{(i,f)} g = f$$

As the S_i form a coseparating family, g is a monomorphism. Thus A is a subobject of Q by g . There is a map $h : P \rightarrow Q$ defined by

$$\pi_{(i,f)} h = \pi_i$$

Thus we can form the pullback

$$\begin{array}{ccc} B & \longrightarrow & A \\ \downarrow k & & \downarrow g \\ P & \xrightarrow{h} & Q \end{array}$$

where k is a monomorphism as it is the pullback of a monomorphism. Hence B is a subobject of P , and thus factors through I .

$$\begin{array}{ccc} I & \dashrightarrow & B \\ & \searrow & \downarrow k \\ & & P \end{array}$$

Hence, we have a morphism $I \rightarrow A$ by composition. □

Example. Let $I : \mathbf{KHaus} \rightarrow \mathbf{Top}$ be the inclusion functor. \mathbf{KHaus} is closed under small products in \mathbf{Top} by Tychonoff’s theorem, and is closed under equalisers since the equaliser of $f, g : X \rightrightarrows Y$ is a closed subspace of X , and thus is compact and Hausdorff. Hence \mathbf{KHaus} is complete, and the inclusion preserves small limits. It is clearly locally small and well-powered, since the subobjects of X are isomorphic to closed subspaces. It has a single coseparator, namely $[0, 1]$, by Urysohn’s lemma. Hence, by the special adjoint functor theorem, I has a left adjoint β , which is the Stone–Čech compactification functor.

I. Category Theory

Remark. Čech's construction of β is almost identical to the construction of left adjoints given above. Given a space X , one can form

$$P = \prod_{f: X \rightarrow [0,1]} [0,1]; \quad g : X \rightarrow P; \quad \pi_f g = f$$

which is the product of the members of coseparating set for $(X \downarrow I)$. Then, βX can be defined to be the closure of the image of g , that is, the smallest subobject of (P, g) in $(X \downarrow I)$.

The general adjoint functor theorem can also be used to construct β . To obtain a solution-set at a space X , observe that any morphism from X to a compact Hausdorff space IY factors as $X \rightarrow IY' \rightarrow IY$ where Y' is the closure of $X' = \{f(x) \mid x \in X\}$. One can show that if Y' is Hausdorff and X' is dense in Y' , then $|Y'| \leq 2^{|X'|}$.

5. Monads

5.1. Definition

Suppose $F \dashv G$ is an adjunction with $F : \mathcal{C} \rightarrow \mathcal{D}$ and $G : \mathcal{D} \rightarrow \mathcal{C}$, where \mathcal{C} is a well-understood category, but \mathcal{D} is not. We can study \mathcal{D} indirectly inside the context of \mathcal{C} by using the adjunction. We have the composite $T = GF : \mathcal{C} \rightarrow \mathcal{C}$, and we have the unit $\eta : 1_{\mathcal{C}} \rightarrow T$. The counit is not directly accessible from \mathcal{C} , but we have $\mu = G\epsilon_F : T^2 \rightarrow T$. The triangular identities give rise to identities linking η and μ .

$$\begin{array}{ccc} T & \xrightarrow{T\eta} & T^2 \\ & \searrow 1_T & \downarrow \mu \\ & & T \end{array} \quad \begin{array}{ccc} T & \xrightarrow{\eta_T} & T^2 \\ & \searrow 1_T & \downarrow \mu \\ & & T \end{array}$$

In addition, naturality of ϵ gives

$$\begin{array}{ccc} T^3 & \xrightarrow{T\mu} & T^2 \\ \mu_T \downarrow & & \downarrow \mu \\ T^2 & \xrightarrow{\mu} & T \end{array}$$

Definition. A *monad* on a category \mathcal{C} is a triple $\mathbb{T} = (T, \eta, \mu)$ where T is a functor $\mathcal{C} \rightarrow \mathcal{C}$, and $\eta : 1_{\mathcal{C}} \rightarrow T$ and $\mu : T^2 \rightarrow T$ are natural transformations satisfying the following commutative diagrams.

$$\begin{array}{ccc} T & \xrightarrow{T\eta} & T^2 \\ & \searrow 1_T & \downarrow \mu \\ & & T \end{array} \quad \begin{array}{ccc} T & \xrightarrow{\eta_T} & T^2 \\ & \searrow 1_T & \downarrow \mu \\ & & T \end{array} \quad \begin{array}{ccc} T^3 & \xrightarrow{T\mu} & T^2 \\ \mu_T \downarrow & & \downarrow \mu \\ T^2 & \xrightarrow{\mu} & T \end{array}$$

η is the *unit* of the monad, and μ is the *multiplication* of the monad.

The dual notion is called a *comonad*.

Example. (i) Let M be a monoid. The functor $M \times (-) : \mathbf{Set} \rightarrow \mathbf{Set}$ has a monad structure. The unit $\eta_A : A \rightarrow M \times A$ maps each a to $(1, a)$, and the multiplication $\mu_A : M \times M \times A \rightarrow M \times A$ maps (m, m', a) to (mm', a) . These maps are natural. The required commutative diagrams encode precisely the left and right unit laws and the associativity law of a monoid. In fact, monoids correspond precisely to monads on \mathbf{Set} whose underlying functors have right adjoints.

(ii) Let $P : \mathbf{Set} \rightarrow \mathbf{Set}$ be the covariant power-set functor. This can be given a monad structure. The unit $\eta_A : A \rightarrow PA$ maps a to its singleton $\{a\}$, and the multiplication $\mu_A : PPA \rightarrow PA$ is the union operation mapping S to $\bigcup S$. One can check that the required laws are satisfied.

I. Category Theory

These examples both arise as a result of adjunctions. Example (a) arises from the free M -set functor $F : \mathbf{Set} \rightarrow [M, \mathbf{Set}]$ and the forgetful functor $U : [M, \mathbf{Set}] \rightarrow \mathbf{Set}$, where $F \dashv U$. For example (b), there is a forgetful functor $U : \mathbf{CSLat} \rightarrow \mathbf{Set}$ from the category of complete (join-)semilattices. This has a left adjoint $P : \mathbf{Set} \rightarrow \mathbf{CSLat}$, which is the free complete semilattice on A . Indeed, given any $f : A \rightarrow UB$, there is a unique extension of f to a join-preserving map $\bar{f} : PA \rightarrow B$ given by

$$\bar{f}(A') = \bigvee \{f(a') \mid a' \in A'\}$$

Note that an M -set is a set A equipped with a map $\alpha : M \times A \rightarrow A$, and a complete semilattice is a set A equipped with a map $\bigvee : PA \rightarrow A$. So the elements of the other category can be defined in terms of the monad.

This holds in general: every monad arises from an adjunction. We present two constructions.

5.2. Eilenberg–Moore algebras

Definition. Let $\mathbb{T} = (T, \eta, \mu)$ be a monad on \mathcal{C} . An *Eilenberg–Moore algebra* or \mathbb{T} -*algebra* is a pair (A, α) where A is an object in \mathcal{C} , and $\alpha : TA \rightarrow A$ is a morphism satisfying

$$\begin{array}{ccc} A & \xrightarrow{\eta_A} & TA \\ & \searrow 1_A & \downarrow \alpha \\ & & A \end{array} \quad \begin{array}{ccc} T^2A & \xrightarrow{T\alpha} & TA \\ \mu_A \downarrow & & \downarrow \alpha \\ TA & \xrightarrow{\alpha} & A \end{array}$$

A homomorphism of algebras $f : (A, \alpha) \rightarrow (B, \beta)$ is a morphism $f : A \rightarrow B$ such that the following diagram commutes.

$$\begin{array}{ccc} TA & \xrightarrow{Tf} & TB \\ \alpha \downarrow & & \downarrow \beta \\ A & \xrightarrow{f} & B \end{array}$$

This forms a category of \mathbb{T} -algebras, denoted $\mathcal{C}^{\mathbb{T}}$.

Proposition. The forgetful functor $G^{\mathbb{T}} : \mathcal{C}^{\mathbb{T}} \rightarrow \mathcal{C}$ has a left adjoint $F^{\mathbb{T}}$, and the adjunction $F^{\mathbb{T}} \dashv G^{\mathbb{T}}$ induces the monad \mathbb{T} on \mathcal{C} .

Proof. We define the *free algebra* of an object A to be $F^{\mathbb{T}}A = (TA, \mu_A)$. This defines an algebra structure on TA for every A by the monad laws. For $f : A \rightarrow B$, we define $F^{\mathbb{T}}f = Tf$; this is a homomorphism by naturality of μ . This is functorial as T is functorial.

We have $G^{\mathbb{T}}F^{\mathbb{T}} = T$. For the unit of the adjunction, we use the unit of the monad η . For the counit, we define

$$\mu_{(A, \alpha)} = \alpha : F^{\mathbb{T}}A \rightarrow (A, \alpha)$$

This is a homomorphism by the definition of an algebra, and it is a natural transformation by the definition of homomorphisms of algebras. It suffices to verify the triangular identities, which follows from the remaining unused diagrams. One can check that the multiplication induced by this monad is equal to that of \mathbb{T} . \square

5.3. Kleisli categories

If $F \dashv G$ with $F : \mathcal{C} \rightarrow \mathcal{D}$ and $G : \mathcal{D} \rightarrow \mathcal{C}$ is an adjunction inducing \mathbb{T} , then $F' \dashv G'$ with $F' : \mathcal{C} \rightarrow \mathcal{D}'$ and $G' : \mathcal{D}' \rightarrow \mathcal{C}$, where \mathcal{D}' is the full subcategory of \mathcal{D} on objects in the image of F . Thus, when finding a construction for \mathcal{D} , we can assume that F is surjective (or, indeed, bijective) on objects. Then, the morphisms $FA \rightarrow FB$ must correspond to morphisms $A \rightarrow GFB$ under the adjunction, but $GF = T$.

Definition. Let $\mathbb{T} = (T, \mu, \eta)$ be a monad on \mathcal{C} . The *Kleisli category* $\mathcal{C}_{\mathbb{T}}$ is the category where the objects are precisely the objects of \mathcal{C} , and the morphisms from A to B in $\mathcal{C}_{\mathbb{T}}$ are the morphisms $A \rightarrow TB$ in \mathcal{C} . To avoid confusion, we will denote morphisms from A to B in this category by $A \rightrightarrows B$. The identity $A \rightrightarrows A$ is $\eta_A : A \rightarrow TA$. The composite of

$$A \xrightarrow{\quad f \quad} B \xrightarrow{\quad g \quad} C$$

is

$$A \xrightarrow{f} TB \xrightarrow{Tg} T^2C \xrightarrow{\mu_C} TC$$

These satisfy the unit and associativity laws.

$$\begin{array}{ccc} A & \xrightarrow{f} & TB \\ & \searrow 1_{TB} & \downarrow \mu_B \\ & & TB \end{array} \quad \begin{array}{ccc} A & \xrightarrow{\eta_A} & TA \\ f \downarrow & & \downarrow Tf \\ TB & \xrightarrow{\eta_{TB}} & T^2B \\ & \searrow 1_{TB} & \downarrow \mu_B \\ & & TB \end{array}$$

$$\begin{array}{ccccccc} A & \xrightarrow{f} & TB & \xrightarrow{Tg} & T^2C & \xrightarrow{T^2h} & T^3D \\ & & & & \downarrow \mu_C & & \downarrow \mu_{TD} \\ & & & & TC & \xrightarrow{Th} & T^2D \\ & & & & & & \downarrow \mu_D \\ & & & & & & TD \end{array}$$

where in the last diagram, the upper composite is $(hg)f$ and the lower composite is $h(gf)$ in $\mathcal{C}_{\mathbb{T}}$.

Proposition. There is an adjunction $F_{\mathbb{T}} \dashv G_{\mathbb{T}}$ where $F_{\mathbb{T}} : \mathcal{C} \rightarrow \mathcal{C}_{\mathbb{T}}$ and $G_{\mathbb{T}} : \mathcal{C}_{\mathbb{T}} \rightarrow \mathcal{C}$ that induces the monad \mathbb{T} .

I. Category Theory

Proof. We define $F_{\mathbb{T}}A = A$, and for $f : A \rightarrow B$, define $F_{\mathbb{T}}f = \eta_B f$. This preserves identities as $1_{F_{\mathbb{T}}A} = \eta_A$, and preserves composites since

$$\begin{array}{ccccc} A & \xrightarrow{f} & B & \xrightarrow{\eta_B} & TB \\ & & \downarrow g & & \downarrow Tg \\ & & C & \xrightarrow{\eta_C} & TC \\ & & & & \downarrow 1_{TC} \end{array} \quad \begin{array}{ccc} & & T^2C \\ & \nearrow T\eta_C & \downarrow \mu_C \\ & & TC \end{array}$$

commutes. For $G_{\mathbb{T}}$, we define $G_{\mathbb{T}}A = TA$, and for $f : A \rightarrow B$, we define $G_{\mathbb{T}}f$ to be the composite

$$TA \xrightarrow{Tf} T^2B \xrightarrow{\mu_B} TB$$

Note that $G_{\mathbb{T}}$ preserves identities by the unit law and preserves composites as

$$\begin{array}{ccccccc} TA & \xrightarrow{Tf} & T^2B & \xrightarrow{T^2g} & T^3C & \xrightarrow{T\mu_C} & T^2C \\ & & \downarrow \mu_B & & \downarrow \mu_{TC} & & \downarrow \mu_C \\ & & TB & \xrightarrow{Tg} & T^2C & \xrightarrow{\mu_C} & TC \end{array}$$

commutes. Then $G_{\mathbb{T}}$ is a functor, and $G_{\mathbb{T}}F_{\mathbb{T}} = T$. The unit of the adjunction is the unit of the monad η . For the counit $\epsilon_A : TA = F_{\mathbb{T}}G_{\mathbb{T}}A \rightarrow A$, we use the identity 1_{TA} . This is natural, as given $f : A \rightarrow B$, the diagram

$$\begin{array}{ccc} TA & \xrightarrow{F_{\mathbb{T}}G_{\mathbb{T}}f} & TB \\ \epsilon_A \downarrow & & \downarrow \epsilon_B \\ A & \xrightarrow{f} & B \end{array}$$

commutes, as the paths are

$$TA \xrightarrow{Tf} T^2B \xrightarrow{\mu_B} TB \xrightarrow{\eta_{TB}} T^2B \xrightarrow{\mu_B} TB$$

and

$$TA \xrightarrow{Tf} T^2B \xrightarrow{\mu_B} TB$$

which coincide. One can show that both triangular identities reduce to a unit law. It suffices to verify that the multiplication of the induced monad is correct. The multiplication law is $G_{\mathbb{T}}\epsilon_{F_{\mathbb{T}}A}$, which is

$$T^2A \xrightarrow{T1_{TA}} T^2A \xrightarrow{\mu_A} TA$$

which is equal to μ_A , as required. □

5.4. Comparison functors

Definition. Let $\mathbb{T} = (T, \eta, \mu)$ be a monad on \mathcal{C} . Then $\text{Adj}(\mathbb{T})$ is the category of adjunctions $F \dashv G$ which induce \mathbb{T} , where the morphisms $F \dashv G$ to $F' \dashv G'$ are the functors $K : \mathcal{D} \rightarrow \mathcal{D}'$ satisfying $KF = F'$ and $G'K = G$.

$$\begin{array}{ccc}
 & \mathcal{C} & \\
 F \swarrow & & \searrow F' \\
 \mathcal{D} & \xrightarrow{K} & \mathcal{D}' \\
 G \searrow & & \swarrow G' \\
 & \mathcal{C} &
 \end{array}$$

Theorem. The Kleisli adjunction $F_{\mathbb{T}} \dashv G_{\mathbb{T}}$ is initial in $\text{Adj}(\mathbb{T})$, and the Eilenberg–Moore adjunction $F^{\mathbb{T}} \dashv G^{\mathbb{T}}$ is terminal in $\text{Adj}(\mathbb{T})$.

Proof. We will first do the case of the Eilenberg–Moore adjunction. Let $F \dashv G$ be an adjunction inducing \mathbb{T} . We define $K : \mathcal{D} \rightarrow \mathcal{C}^{\mathbb{T}}$ by $KB = (GB, G\epsilon_B)$. This is an algebra by the triangular identities and naturality of ϵ . On morphisms $f : B \rightarrow C$ in \mathcal{D} , we define $Kg = Gg$, which is a homomorphism as ϵ is a natural transformation. Clearly $G^{\mathbb{T}}K = G$, and $KFA = (GFA, G\epsilon_{FA}) = F^{\mathbb{T}}A$, and for $f : A \rightarrow A'$, $KFf = GFf = Tf = F^{\mathbb{T}}f$. So K is a morphism of $\text{Adj}(\mathbb{T})$.

For uniqueness, suppose K' were another such morphism. Then $K'B = (GB, \beta_B)$, and $K'g = Gg$ for $g : B \rightarrow C$. Note that β must be a natural transformation $GFG \rightarrow G$. Also, $\beta_{FA} = G\epsilon_{FA}$ for all A , as $K'F = F^{\mathbb{T}}$. But we have naturality squares

$$\begin{array}{ccc}
 GFGFGB & \xrightarrow{GF G \epsilon_B} & GFGB \\
 \beta_{FGB} \downarrow \downarrow G\epsilon_{FGB} & & \beta_B \downarrow \downarrow G\epsilon_B \\
 GFGB & \xrightarrow{G\epsilon_B} & GB
 \end{array}$$

where the left edges are equal and the top edge is a split epimorphism, so the right edges are equal. Thus K is unique.

Given an adjunction $F \dashv G$ inducing \mathbb{T} , we define $H : \mathcal{C}_{\mathbb{T}} \rightarrow \mathcal{D}$ by $HA = FA$, and for $f : A \rightarrow B$, define Hf to be the composite

$$FA \xrightarrow{Ff} FGFB \xrightarrow{\epsilon_{FB}} FB$$

This is functorial. Indeed, for $f : A \rightarrow B$ and $g : B \rightarrow C$, $H(gf)$ is the upper composite and $(Hg)(Hf)$ is the lower composite in the following diagram.

$$\begin{array}{ccccccc}
 FA & \xrightarrow{Ff} & FGFB & \xrightarrow{FGFg} & FGFGFC & \xrightarrow{FG\epsilon_{FC}} & FGFC \\
 \downarrow \epsilon_{FB} & & \downarrow \epsilon_{FGFC} & & \downarrow \epsilon_{FC} & & \\
 FB & \xrightarrow{Fg} & FGFC & \xrightarrow{\epsilon_{FC}} & FC & &
 \end{array}$$

I. Category Theory

Then $HF_{\mathbb{T}}(f) = \epsilon_{FB}(F\eta_B)(Ff) = Ff$. Moreover, $GHA = GFA = TA = G_{\mathbb{T}}A$, and for $f : A \multimap B$, GFf is the composite

$$GFA \xrightarrow{GFf} GFGFB \xrightarrow{\mu_B} GFB$$

which is the definition of $G_{\mathbb{T}}(f)$. Thus H is a morphism of $\text{Adj}(\mathbb{T})$. If $H' : \mathcal{C}_{\mathbb{T}} \rightarrow \mathcal{D}$ were another such morphism, then since $H'F_{\mathbb{T}} = F$, we must have $H'A = FA$ for all A . Note that for $f : A \multimap B$, Hf is the transpose of $f : A \rightarrow GFB$ across $F \dashv G$. Since H' commutes with G and $G_{\mathbb{T}}$, and $F \dashv G$ and $F_{\mathbb{T}} \dashv G_{\mathbb{T}}$ have the same unit η , H' must send the transpose $f : A \multimap B$ of $f : A \rightarrow GFB$ to its transpose across $F \dashv G$, which is precisely the action of H on morphisms. Hence $H' = H$. \square

Definition. The functor $K : \mathcal{D} \rightarrow \mathcal{C}^{\mathbb{T}}$ is called the *Eilenberg–Moore comparison functor*. Similarly, the functor $H : \mathcal{C}_{\mathbb{T}} \rightarrow \mathcal{D}$ is called the *Kleisli comparison functor*.

Remark. Note that $\mathcal{C}_{\mathbb{T}}$ has coproducts if \mathcal{C} does, since $F_{\mathbb{T}}$ preserves them and is bijective on objects. However, it has few other limits or colimits in general. In contrast, $\mathcal{C}^{\mathbb{T}}$ inherits many limits and colimits from \mathcal{C} .

Proposition. (i) The forgetful functor $G = G^{\mathbb{T}} : \mathcal{C}^{\mathbb{T}} \rightarrow \mathcal{C}$ creates any limits which exist in \mathcal{C} .

(ii) If \mathcal{C} has colimits of shape J , then $G = G^{\mathbb{T}}$ creates colimits of shape J if and only if T preserves them.

Proof. Part (i). Let $D : J \rightarrow \mathcal{C}^{\mathbb{T}}$ be a diagram of shape J . Write $D(j) = (GD(j), \delta_j)$ for $j \in \text{ob } J$. Let $(L, (\lambda_j : L \rightarrow GD(j))_{j \in \text{ob } J})$ be a limit for GD in \mathcal{C} . Then $(TL, (T\lambda_j)_{j \in \text{ob } J})$ is a cone over TGD , so $(TL, (\delta(T\lambda_j))_{j \in \text{ob } J})$ is a cone over TGD , and induces a unique $\theta : TL \rightarrow L$ making squares of the form

$$\begin{array}{ccc} TL & \xrightarrow{T\lambda_j} & TGD(j) \\ \theta \downarrow & & \downarrow \delta_j \\ L & \xrightarrow{\lambda_j} & GD(j) \end{array}$$

commute for each j . Note that θ is an algebra structure on L , since the required diagrams commute by uniqueness of factorisation through limits. It is the unique algebra structure on L which make the λ_j into a cone in $\mathcal{C}^{\mathbb{T}}$, and one can easily show it is a limit cone.

Part (ii). In the forward direction, if G creates colimits of shape J , then it certainly preserves them, as they exist in both categories. But F preserves all colimits, so $T = GF$ preserves them. Given $D : J \rightarrow \mathcal{C}^{\mathbb{T}}$ and a colimit cone $\lambda_j : GD(j) \rightarrow L$ under GD , we know that $T\lambda_j : TGD(j) \rightarrow TL$ is a colimit cone, so there is a unique $\theta : TL \rightarrow L$ satisfying $\theta(T\lambda_j) = \lambda_j \delta_j$ for all j , and θ is an algebra structure since TTL is also a colimit. Hence (L, θ) is a colimit for D in $\mathcal{C}^{\mathbb{T}}$. \square

Remark. One can show that $\mathcal{C}^\mathbb{T}$ has colimits of any shape which exist in \mathcal{C} , provided that it has *reflexive coequalisers*.

5.5. Monadic adjunctions

It can be useful to know, for an arbitrary adjunction, if the Eilenberg–Moore comparison functor $K : \mathcal{D} \rightarrow \mathcal{C}^\mathbb{T}$ is part of an equivalence of categories. Note that the Kleisli comparison functor H is always full and faithful, so is part of an equivalence if and only if it is essentially surjective, and since its action on objects is F , this holds if and only if F is essentially surjective.

Definition. An adjunction $F \dashv G$ is *monadic*, or the right adjoint G is *monadic*, if K is part of an equivalence.

Lemma. Let $F \dashv G$ be an adjunction inducing the monad \mathbb{T} , and suppose that for every \mathbb{T} -algebra (A, α) , the pair

$$FGFA \begin{array}{c} \xrightarrow{F\alpha} \\ \xrightarrow{\epsilon_{FA}} \end{array} FA$$

has a coequaliser in \mathcal{D} . Then the comparison functor $K : \mathcal{D} \rightarrow \mathcal{C}^\mathbb{T}$ has a left adjoint L .

Proof. Let $\lambda_{(A, \alpha)} : FA \rightarrow L(A, \alpha)$ be a coequaliser for $F\alpha, \epsilon_{FA}$. We can make L into a functor $\mathcal{C}^\mathbb{T} \rightarrow \mathcal{D}$. Given $f : (A, \alpha) \rightarrow (B, \beta)$, the composite $\lambda_{(B, \beta)}(Ff)$ coequalises $F\alpha$ and ϵ_{FA} , so it induces a unique map $Lf : L(A, \alpha) \rightarrow L(B, \beta)$. This makes L into a functor by uniqueness.

$$\begin{array}{ccccc} FGFA & \begin{array}{c} \xrightarrow{F\alpha} \\ \xrightarrow{\epsilon_{FA}} \end{array} & FA & \xrightarrow{\lambda_{(A, \alpha)}} & L(A, \alpha) \\ FGFf \downarrow & & \downarrow Ff & & \downarrow Lf \\ FGFB & \begin{array}{c} \xrightarrow{F\beta} \\ \xrightarrow{\epsilon_{FB}} \end{array} & FB & \xrightarrow{\lambda_{(B, \beta)}} & L(B, \beta) \end{array}$$

For any object B of \mathcal{D} , morphisms $L(A, \alpha) \rightarrow B$ correspond to morphisms $f : FA \rightarrow B$ satisfying $f(F\alpha) = f\epsilon_{FA}$. If $\bar{f} : A \rightarrow GB$ is the transpose of f across $F \dashv G$, then by naturality, the transpose of $f(F\alpha)$ is $\bar{f}\alpha$, and the transpose of $f\epsilon_{FA}$ is Gf since ϵ_{FA} transposes to 1_{GFA} . But we have $f = \epsilon_B(G\bar{f})$, so $(G\epsilon_B)(G\bar{f}) = (G\epsilon_B)(T\bar{f})$. Thus $f(F\alpha) = f\epsilon_{FA}$ if and only if $\bar{f}\alpha = (G\epsilon_B)(T\bar{f})$, which is to say that \bar{f} is an algebra homomorphism $(A, \alpha) \rightarrow (GB, G\epsilon_B) = KB$. Naturality of this bijection follows from the fact that the map $f \mapsto \bar{f}$ is natural, so $L \dashv K$ as required. \square

Definition. A parallel pair $f, g : A \rightrightarrows B$ is *reflexive* if there exists $r : B \rightarrow A$ such that $fr = gr = 1_B$.

$$\begin{array}{ccc} A & \begin{array}{c} \xrightarrow{f} \\ \xrightarrow{g} \end{array} & B \\ r \uparrow & & \uparrow 1_B \\ B & & \end{array}$$

I. Category Theory

Note that the parallel pair

$$FGFA \begin{array}{c} \xrightarrow{F\alpha} \\ \xrightarrow{\epsilon_{FA}} \end{array} FA$$

is a reflexive pair, and the common right inverse is $r = F\eta_A$.

Definition. A *split coequaliser diagram* is a diagram

$$\begin{array}{ccccc} A & \xrightarrow{f} & B & \xrightarrow{h} & C \\ & \xrightarrow{g} & & \xleftarrow{s} & \\ & \xleftarrow{t} & & & \end{array}$$

such that $hf = hg$, $hs = 1_C$, $gt = 1_B$, $ft = sh$. That is, h has equal composites with f and g , and the following diagrams commute.

$$\begin{array}{ccc} A & \xrightarrow{g} & B & \xrightarrow{h} & C \\ \uparrow t & \nearrow 1_B & \uparrow s & \nearrow 1_C & \\ B & & C & & \end{array} \quad \begin{array}{ccc} B & \xrightarrow{t} & A \\ h \downarrow & & \downarrow f \\ C & \xrightarrow{s} & B \end{array}$$

The equations $hs = 1_C$, $gt = 1_B$ enforce that s is a section of h , and t is a section of g . The equation $ft = sh$ enforces that the two non-identity paths from B to itself coincide.

Note that this implies that h is a coequaliser of f and g . Indeed, if $k : B \rightarrow D$ satisfies $kf = kg$, then $k = kgt = kft = ksh$, so k factors through h . Moreover, this factorisation is unique as h is split epic. Any functor preserves split coequaliser diagrams.

Definition. Given a functor $G : \mathcal{D} \rightarrow \mathcal{C}$, we say that a parallel pair $f, g : A \rightrightarrows B$ in \mathcal{D} is *G-split* if there is a split coequaliser diagram

$$\begin{array}{ccccc} GA & \xrightarrow{Gf} & GB & \xrightarrow{h} & C \\ & \xrightarrow{Gg} & & \xleftarrow{s} & \\ & \xleftarrow{t} & & & \end{array}$$

in \mathcal{C} .

Note that the pair

$$FGFA \begin{array}{c} \xrightarrow{F\alpha} \\ \xrightarrow{\epsilon_{FA}} \end{array} FA$$

is *G-split*, as

$$\begin{array}{ccccc} GFGFA & \xrightarrow{GF\alpha} & GFA & \xrightarrow{\alpha} & C \\ & \xrightarrow{G\epsilon_{FA}=\mu_A} & & \xleftarrow{\eta_A} & \\ & \xleftarrow{\eta_{GFA}} & & & \end{array}$$

is a split coequaliser diagram.

Theorem (Beck's precise monadicity theorem). A functor $G : \mathcal{D} \rightarrow \mathcal{C}$ is monadic if and only if G has a left adjoint and creates coequalisers of G -split pairs.

Theorem (Beck's crude monadicity theorem). Suppose $G : \mathcal{D} \rightarrow \mathcal{C}$ has a left adjoint, and G reflects isomorphisms. Suppose further that \mathcal{D} has and G preserves reflexive coequalisers. Then G is monadic.

We prove both theorems together.

Proof. First, suppose $G : \mathcal{D} \rightarrow \mathcal{C}$ is monadic. Then G has a left adjoint by definition. It suffices to show that $G^\top : \mathcal{C}^\top \rightarrow \mathcal{C}$ creates coequalisers of G^\top -split pairs. This follows from the argument of a previous lemma: if $f, g : (A, \alpha) \rightrightarrows (B, \beta)$ are algebra homomorphisms, and

$$\begin{array}{ccccc} A & \xrightarrow{f} & B & \xrightarrow{h} & C \\ & \xrightarrow{g} & & \xleftarrow{s} & \\ & \xleftarrow{t} & & & \end{array}$$

is a split coequaliser, then since the coequaliser is preserved by T and T^2 , C acquires a unique algebra structure $\gamma : TC \rightarrow C$ such that h is a coequaliser in \mathcal{C}^\top .

For the converse, either set of assumptions ensures that \mathcal{D} has coequalisers of parallel pairs of the form

$$FGFA \begin{array}{c} \xrightarrow{F\alpha} \\ \xrightarrow{\epsilon_{FA}} \end{array} FA$$

so the comparison functor $K : \mathcal{D} \rightarrow \mathcal{C}^\top$ has a left adjoint L . We must now show that the unit and counit of $L \dashv K$ are isomorphisms. The unit $(A, \alpha) \rightarrow KL(A, \alpha)$ is the unique factorisation of $G\lambda_{(A, \alpha)} : GFA \rightarrow GL(A, \alpha)$ through the $(G^\top$ -split) coequaliser $\alpha : GFA \rightarrow A$ of $GF\alpha, G\epsilon_{FA} : GFGFA \rightrightarrows GFA$ in \mathcal{C}^\top . But either set of hypotheses implies that G preserves the coequaliser of $F\alpha, \epsilon_{FA}$, so the factorisation is an isomorphism. The counit $LKB \rightarrow B$ is the unique factorisation of $\epsilon_B : FGB \rightarrow B$ through $\lambda_{KB} : FGB \rightarrow LKB$. The hypothesis in the precise theorem implies directly that ϵ_B is a coequaliser of $FG\epsilon_B, \epsilon_{GFB}$, because the pair is G -split. From the hypotheses of the crude theorem, we can see that both ϵ_B and λ_{KB} map to coequalisers in \mathcal{C} , so the counit maps to an isomorphism in \mathcal{C} , so it is an isomorphism as G reflects isomorphisms. \square

Remark. (i) Let J be the finite category

$$\begin{array}{ccc} & \overset{s}{\curvearrowright} & \xrightarrow{f} \\ A & \xleftarrow{r} & B \\ & \underset{t}{\curvearrowright} & \xrightarrow{g} \end{array}$$

with $fr = gr = 1_B, rf = s, rg = t$, then a diagram D of this shape is a reflexive pair. A cone under it is determined by $h : DB \rightarrow L$, which must satisfy $h(Df) = h(Dg)$. A colimit for this diagram is a coequaliser for f, g .

I. Category Theory

- (ii) All small (respectively finite) colimits can be constructed from small (respectively finite) coproducts and reflexive coequalisers. The pair $f, g : P \rightrightarrows Q$ in the proof form a coreflexive pair, with common left inverse $r : Q \rightarrow P$ given by $\pi_j r = \pi_{1_j}$ for all j .
- (iii) Given a reflexive pair $f, g : A \rightrightarrows B$, a morphism $h : B \rightarrow C$ is a coequaliser for it if and only if the diagram

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ g \downarrow & & \downarrow h \\ B & \xrightarrow{h} & C \end{array}$$

is a pushout, since any cone under the span given by f and g has its two legs equal. The dual of this statement has already been proven.

- (iv) In any cartesian closed category, reflexive coequalisers commute with finite products: if the following are reflexive coequaliser diagrams,

$$A_1 \begin{array}{c} \xrightarrow{f_1} \\ \xrightarrow{g_1} \end{array} B_1 \xrightarrow{h_1} C_1 \qquad A_2 \begin{array}{c} \xrightarrow{f_2} \\ \xrightarrow{g_2} \end{array} B_2 \xrightarrow{h_2} C_2$$

then the following diagram is also a coequaliser.

$$A_1 \times A_2 \begin{array}{c} \xrightarrow{f_1 \times f_2} \\ \xrightarrow{g_1 \times g_2} \end{array} B_1 \times B_2 \xrightarrow{h_1 \times h_2} C_1 \times C_2$$

Indeed, consider the diagram

$$\begin{array}{ccccc} A_1 \times A_2 & \begin{array}{c} \xrightarrow{\quad} \\ \xrightarrow{\quad} \end{array} & B_1 \times A_2 & \longrightarrow & C_1 \times A_2 \\ \downarrow \downarrow & & \downarrow \downarrow & & \downarrow \downarrow \\ A_1 \times B_2 & \begin{array}{c} \xrightarrow{\quad} \\ \xrightarrow{\quad} \end{array} & B_1 \times B_2 & \longrightarrow & C_1 \times B_2 \\ \downarrow & & \downarrow & & \downarrow \\ A_1 \times C_2 & \begin{array}{c} \xrightarrow{\quad} \\ \xrightarrow{\quad} \end{array} & B_1 \times C_2 & \longrightarrow & C_1 \times C_2 \end{array}$$

All rows and columns are coequalisers, since functors of the form $(-) \times D$ preserve coequalisers. It then follows that the lower right square is a pushout. By reflexivity, if $k : B_1 \times B_2 \rightarrow D$ coequalises

$$f_1 \times f_2, g_1 \times g_2 : A_1 \times A_2 \rightrightarrows B_1 \times B_2$$

then it also coequalises $B_1 \times A_2 \rightrightarrows B_1 \times B_2$ and $A_1 \times B_2 \rightrightarrows B_1 \times B_2$, as they both factor through the diagonal pair. Therefore, it factors through the top and left edges of the lower right square, and hence through its diagonal.

Example. (i) The forgetful functor $U : \mathbf{Gp} \rightarrow \mathbf{Set}$ satisfies the hypotheses of the crude monadicity theorem. Indeed, it has a left adjoint and reflects isomorphisms, and it creates reflexive coequalisers. Given a reflexive pair $f, g : A \rightrightarrows B$ in \mathbf{Gp} , consider its coequaliser $h : UB \rightarrow C$ in \mathbf{Set} . As reflexive coequalisers commute with products in \mathbf{Set} ,

$$UA \times UA \begin{array}{c} \xrightarrow{f} \\ \xrightarrow{g} \end{array} UB \times UB \longrightarrow C \times C$$

is a coequaliser. So we obtain a binary operation $C \times C \rightarrow C$ making h into a homomorphism, C into a group, and h a coequaliser in \mathbf{Gp} . The same procedure applies for many other algebraic structures, such as rings, modules over a given ring, and lattices. For infinitary algebraic categories such as complete semilattices and complete lattices, we can use the precise monadicity theorem whenever a left adjoint exists.

(ii) Any reflection is monadic. If $I : \mathcal{D} \rightarrow \mathcal{C}$ is the inclusion of a reflective subcategory and $f, g : A \rightrightarrows B$ is an I -split pair in \mathcal{D} , then the splitting $t : B \rightarrow A$ belongs to \mathcal{D} , and so its composite $ft = sh$ also lies in \mathcal{D} . But \mathcal{D} is closed under limits that exist in \mathcal{C} , so in particular it is closed under splittings of idempotents.

(iii) Consider the composite adjunction

$$\mathbf{Set} \begin{array}{c} \xrightarrow{F} \\ \xleftarrow{U} \end{array} \mathbf{AbGp} \begin{array}{c} \xrightarrow{L} \\ \xleftarrow{I} \end{array} \mathbf{tfAbGp}$$

Both factors are monadic: we have already shown that $F \dashv U$ is monadic, and $L \dashv I$ is a reflection. However, the composite $LF \dashv UI$ is not monadic. Indeed, free abelian groups are torsion-free, so the monad induced by the composite adjunction coincides with that induced by $F \dashv U$.

(iv) The contravariant power-set functor $P^* : \mathbf{Set}^{\text{op}} \rightarrow \mathbf{Set}$ is monadic as it satisfies the hypotheses of the crude monadicity theorem. Its left adjoint is $P^* : \mathbf{Set} \rightarrow \mathbf{Set}^{\text{op}}$, and it reflects isomorphisms. Let

$$A \xrightarrow{e} B \begin{array}{c} \xrightarrow{f} \\ \xrightarrow{g} \end{array} C$$

be a coreflexive equaliser in \mathbf{Set} . Then the square

$$\begin{array}{ccc} A & \xrightarrow{e} & B \\ e \downarrow & & \downarrow g \\ B & \xrightarrow{f} & C \end{array}$$

is a pullback. Thus, the composite

$$PB \xrightarrow{P^*e} PA \xrightarrow{Pe} PB$$

I. Category Theory

coincides with

$$PB \xrightarrow{Pg} PC \xrightarrow{P^*f} PB$$

Also, $(P^*e)(Pe) = 1_{PA}$ and $(P^*g)(Pg) = 1_{PB}$, so we obtain the following split coequaliser diagram in **Set**.

$$\begin{array}{ccccc} PC & \xrightarrow{P^*f} & PB & \xrightarrow{P^*e} & PA \\ & \xrightarrow{P^*g} & & \xleftarrow{Pe} & \\ & \xleftarrow{Pg} & & & \end{array}$$

- (v) The forgetful functor $U : \mathbf{Top} \rightarrow \mathbf{Set}$ is not monadic. The monad induced by $D \dashv U$ is $1_{\mathbf{Set}}$, and the unit and multiplication are the identity natural transformations. Hence its category of algebras is isomorphic to **Set**. This example demonstrates that reflection of isomorphisms is necessary for the crude theorem.
- (vi) The composite

$$\mathbf{Set} \xrightleftharpoons[U]{D} \mathbf{Top} \xrightleftharpoons[I]{\beta} \mathbf{KHaus}$$

is monadic, where β is the Stone-Ćech compactification functor; we will prove this using the precise monadicity theorem. Consider a UI -split pair $f, g : X \rightrightarrows Y$ in **KHaus**.

$$\begin{array}{ccccc} UX & \xrightarrow{Uf} & UY & \xrightarrow{h} & Z \\ & \xrightarrow{Ug} & & \xleftarrow{s} & \\ & \xleftarrow{t} & & & \end{array}$$


There is a unique topology on Z making h into a coequaliser in **Top**, which is the quotient topology. This is compact as it is a continuous image of the compact space Y . Hence h will be a coequaliser in **KHaus** if and only if this topology is Hausdorff. Note that the quotient topology is the only possible candidate topology on Z that could make h into a morphism in **KHaus**.

It is a general fact that for every compact Hausdorff space Y and equivalence relation $S \subseteq Y \times Y$, the quotient is Hausdorff if and only if S is closed as a subset of $Y \times Y$. Suppose $(y_1, y_2) \in S$, so $h(y_1) = h(y_2)$. Then the elements $x_1 = t(y_1)$ and $x_2 = t(y_2)$ satisfy

$$g(x_1) = y_1; \quad g(x_2) = y_2; \quad f(x_1) = f(x_2)$$

and if x_1, x_2 satisfy these three equations, then $h(y_1) = h(y_2)$. Thus S is the image under $g \times g : X \times X \rightarrow Y \times Y$ of the equivalence relation R on X given by $\{(x_1, x_2) \mid f(x_1) = f(x_2)\}$. But R is closed in $X \times X$, as it is the equaliser of $f\pi_1, f\pi_2 : X \times X \rightrightarrows Y$ into a Hausdorff space, so it is compact. Hence S is compact, and thus closed.

Definition. Let $F \dashv G$ be an adjunction with $F : \mathcal{C} \rightarrow \mathcal{D}, G : \mathcal{D} \rightarrow \mathcal{C}$. Suppose that \mathcal{D} has reflexive coequalisers. The *monadic tower* of $F \dashv G$ is the diagram



where \mathbb{T} is the monad induced by $F \dashv G$, K is the comparison functor, L is the left adjoint to K which exists as \mathcal{D} has reflexive coequalisers, \mathbb{S} is the monad induced by $L \dashv K$, and so on. We say that $F \dashv G$ has *monadic length* n , or that \mathcal{D} has *monadic height* n over \mathcal{C} , if the tower reaches an equivalence after n steps.

If $F \dashv G$ is an equivalence, it has monadic length zero. Monadic length one means that $F \dashv G$ is monadic but not an equivalence, and example (iii) above has monadic length two.

6. Monoidal and enriched categories

6.1. Monoidal categories

There are many examples of categories \mathcal{C} equipped with a functor $\otimes : \mathcal{C} \times \mathcal{C} \rightarrow \mathcal{C}$ and an object $I \in \text{ob } \mathcal{C}$ that turn \mathcal{C} into a monoid up to isomorphism. Such a structure on a category is called a *monoidal structure*, which will be defined precisely at the end of this subsection.

- Example.** (i) Let \mathcal{C} be a category with finite products. Let \otimes be the categorical product \times , and let $I = 1$ be the terminal object. This is known as the *cartesian monoidal structure*. Dually, if \mathcal{C} is a category with finite coproducts, it has a *cocartesian monoidal structure*, given by $\otimes = +$ and $I = 0$.
- (ii) In **Met**, the different metrics on $X \times Y$ yield different monoidal structures on **Met**. Each of these have the one-point space, which is the terminal object, as the unit of the monoid.
- (iii) In **AbGp**, the tensor product gives a monoidal structure, where \mathbb{Z} is the unit. Recall that if A, B, C are abelian groups, then morphisms $A \otimes B \rightarrow C$ (that is, \mathbb{Z} -linear maps) correspond to \mathbb{Z} -bilinear maps $A \times B \rightarrow C$. Similarly, if R is a commutative ring, the tensor product \otimes_R gives a monoidal structure on **Mod** $_R$ with unit R . The R -linear maps $A \otimes B \rightarrow C$ correspond to R -bilinear maps $A \times B \rightarrow C$.
- (iv) For any category \mathcal{C} , its category of endofunctors $[\mathcal{C}, \mathcal{C}]$ has a monoidal structure given by composition. The unit is the identity endofunctor $1_{\mathcal{C}}$.
- (v) For posets with top and bottom elements 1 and 0, we can define the *ordinal sum* $A * B$ to be the poset obtained from their disjoint union, by identifying the top element of A with the bottom element of B . This is a monoidal structure, where the unit is the one-element poset.

Definition. A *monoidal category* is a category \mathcal{C} equipped with a functor $\otimes : \mathcal{C} \times \mathcal{C} \rightarrow \mathcal{C}$ and a distinguished object I , together with three natural isomorphisms

$$\alpha_{A,B,C} : (A \otimes B) \otimes C \rightarrow A \otimes (B \otimes C); \quad \lambda_A : I \otimes A \rightarrow A; \quad \rho_A : A \otimes I \rightarrow A$$

such that the diagrams

$$\begin{array}{ccc} ((A \otimes B) \otimes C) \otimes D & \xrightarrow{\alpha_{ABC} \otimes 1_D} & (A \otimes (B \otimes C)) \otimes D \\ \alpha_{A \otimes B, C, D} \downarrow & & \downarrow \alpha_{A, B \otimes C, D} \\ (A \otimes B) \otimes (C \otimes D) & & A \otimes ((B \otimes C) \otimes D) \\ & \searrow \alpha_{A, B, C \otimes D} \quad \swarrow 1_A \otimes \alpha_{BCD} & \\ & A \otimes (B \otimes (C \otimes D)) & \end{array}$$

$$\begin{array}{ccc}
 (A \otimes I) \otimes B & \xrightarrow{\alpha_{A,I,B}} & A \otimes (I \otimes B) \\
 \searrow \rho_A \otimes 1_B & & \swarrow 1_A \otimes \lambda_B \\
 & A \otimes B &
 \end{array}$$

commute, and $\lambda_I = \rho_I : I \otimes I \rightarrow I$. A monoidal category is *strict* if α, λ, ρ are identities.

α is called the *associator*, and λ and ρ are the *left* and *right unitors*.

These diagrams suffice to prove the commutativity of the following two diagrams.

$$\begin{array}{ccc}
 (I \otimes A) \otimes B & \xrightarrow{\alpha_{I,A,B}} & I \otimes (A \otimes B) \\
 \lambda_A \otimes 1_B \downarrow & \swarrow \lambda_{A \otimes B} & \\
 A \otimes B & &
 \end{array}
 \qquad
 \begin{array}{ccc}
 (A \otimes B) \otimes I & \xrightarrow{\alpha_{A,B,I}} & A \otimes (B \otimes I) \\
 \rho_{A \otimes B} \searrow & & \downarrow 1_A \otimes \rho_B \\
 & A \otimes B &
 \end{array}$$

Note that in the category of abelian groups with the usual tensor product, the obvious choice for $\alpha_{A,B,C}$ is the map sending $(a \otimes b) \otimes c$ to $a \otimes (b \otimes c)$. However, there is also a natural isomorphism sending $(a \otimes b) \otimes c$ to $-a \otimes (b \otimes c)$. But this choice does not satisfy the pentagon equation, as a pentagon has an odd number of sides.

6.2. The coherence theorem

Given a monoidal category $(\mathcal{C}, \otimes, I)$, we define a *word* recursively.

- (i) We have a stack of *variables* A, B, C, \dots , which are all words.
- (ii) The unit I is a word.
- (iii) If u, v are words, then $u \otimes v$ is a word.

A word with n variables defines a functor $\mathcal{C}^n \rightarrow \mathcal{C}$.

Theorem (Mac Lane's coherence theorem). For any two words w, w' with the same sequence of variables in the same order, there is a unique natural isomorphism $w \rightarrow w'$ obtained by composing instances of α, λ, ρ and their inverses.


Proof. We define the *height* of a word w to be $a(w) + i(w)$, where

- (i) $a(w)$ is the *associator height*, which is the number of closing parentheses occurring immediately before \otimes in w ;
- (ii) $i(w)$ is the number of occurrences of I in w .


Applying any instance of α, λ, ρ to a word reduces its height. For example, if $\alpha \dots : w \rightarrow w'$, then $a(w') < a(w)$ and $i(w') = i(w)$, and correspondingly if $\lambda \dots w \rightarrow w'$, then $i(w') = i(w) - 1$ and $a(w') \leq a(w)$. In particular, any string of instances of α, λ, ρ starting from w has length at most $a(w) + i(w)$.

I. Category Theory

We say that a word w is *reduced* if either $a(w) = i(w) = 0$ or $w = I$. If $a(w) > 0$, then w is the domain of an instance of α , and if $i(w) > 0$ and $w \neq I$, then w is the domain of an instance of either λ or ρ . Thus, for any word w , there is a string $w \rightarrow \dots \rightarrow w_0$ where w_0 is the unique reduced word containing the same variables of w in the same order. We must show that any two such strings have the same composite. Given




where φ, ψ are instances of α, λ , or ρ , we need to find a word w''' completing the commutative square



where θ, χ are composites of instances of α, λ , and ρ .

If φ, ψ act on disjoint subwords of w , so $w = u \otimes v$ where $\varphi = \varphi' \otimes 1_v$ and $\psi = 1_u \otimes \psi'$, then we can fill in the square as follows.



Now suppose one acts within the argument of the other, for example, if φ is $\alpha_{t,u,v}$ and $\psi = (1_t \otimes \psi') \otimes 1_v$. Then by naturality of α , we can complete the diagram with $1_t \otimes (\psi' \otimes 1_v)$ and $\alpha_{t,u',v}$.

Now suppose that φ and ψ interfere. If φ and ψ are both instances of α , then the pentagon equation completes the commutative square.

Suppose one is an instance of α and the other is an instance of λ or ρ . Then I must occur as one of the three arguments to α . If it is the middle argument, the two diagrams in the definition of a monoidal category complete the square. If it is the left or right argument, the other two diagrams defined immediately after will complete the square.


Finally, if one is an instance of λ and the other is an instance of ρ , then they must be λ_I and ρ_I , and so must agree. This completes the proof that there is a unique natural isomorphism to a reduced word.

6. Monoidal and enriched categories

Now suppose we have a string

$$w_1 \longrightarrow w_2 \longleftarrow w_3 \longrightarrow w_4 \quad \cdots \quad w_n$$

Then there are unique ‘forwards’ morphisms



to w_0 , which is the reduced word with the same sequence of variables. Each of the triangles must commute by the uniqueness result proven above. Hence the composite of the arrows along the top edge is equal to the composite $w_1 \rightarrow w_0 \leftarrow w_n$. \square

Definition. A *symmetry* on a monoidal category $(\mathcal{C}, \otimes, I)$ is a natural isomorphism $\gamma_{A,B} : A \otimes B \rightarrow B \otimes A$ such that the following diagrams commute.

$$\begin{array}{ccc} (A \otimes B) \otimes C & \xrightarrow{\alpha_{A,B,C}} & A \otimes (B \otimes C) \xrightarrow{1_A \otimes \gamma_{B,C}} A \otimes (C \otimes B) \\ \gamma_{A \otimes B, C} \downarrow & & \downarrow \alpha_{A,C,B}^{-1} \\ C \otimes (A \otimes B) & \xleftarrow{\alpha_{C,A,B}} & (C \otimes A) \otimes B \xleftarrow{\gamma_{A,C} \otimes 1_B} (A \otimes C) \otimes B \end{array}$$

$$\begin{array}{ccc} A \otimes I & \xrightarrow{\gamma_{A,I}} & I \otimes A \\ \rho_A \searrow & & \swarrow \lambda_A \\ & A & \end{array} \quad \begin{array}{ccc} A \otimes B & \xrightarrow{\gamma_{A,B}} & B \otimes A \\ 1_{A \otimes B} \searrow & & \downarrow \gamma_{B,A} \\ & A \otimes B & \end{array}$$

For the weaker notion of a *braiding*, we can omit the last of the three diagrams, but add an additional hexagonal equation, since it can no longer be derived from the first.

There is a coherence theorem for symmetric monoidal categories, which is also due to Mac Lane. The theorem shows that for any two words w, w' involving the same set of variables without repetition, there is a unique natural isomorphism between w and w' obtained from compositions of instances of α, λ, γ and their inverses. Note that ρ is not necessary, as it can be produced from instances of λ and γ . The examples of monoidal categories above are all symmetric, except for (iv) and (v).

6.3. Monoidal functors

Definition. Let $(\mathcal{C}, \otimes, I), (\mathcal{D}, \oplus, J)$ be monoidal categories. A (lax) *monoidal functor* $F : (\mathcal{C}, \otimes, I) \rightarrow (\mathcal{D}, \oplus, J)$ is a functor $F : \mathcal{C} \rightarrow \mathcal{D}$ equipped with a natural transformation

I. Category Theory

$\varphi_{A,B} : FA \oplus FB \rightarrow F(A \otimes B)$ and a morphism $\iota : J \rightarrow FI$, such that the following diagrams commute.

$$\begin{array}{ccc}
 (FA \oplus FB) \oplus FC & \xrightarrow{\varphi_{A,B} \oplus 1_{FC}} & F(A \otimes B) \oplus FC \xrightarrow{\varphi_{A \otimes B, C}} F((A \otimes B) \otimes C) \\
 \alpha_{FA, FB, FC} \downarrow & & \downarrow F\alpha_{A, B, C} \\
 FA \oplus (FB \oplus FC) & \xrightarrow{1_{FA} \oplus \varphi_{B, C}} FA \oplus F(B \otimes C) \xrightarrow{\varphi_{A, B \otimes C}} F(A \otimes (B \otimes C))
 \end{array}$$

$$\begin{array}{ccc}
 J \oplus FA & \xrightarrow{\iota \oplus 1_{FA}} & FI \oplus FA \\
 \lambda_{FA} \downarrow & & \downarrow \varphi_{I, A} \\
 FA & \xleftarrow{F\lambda_A} & F(I \otimes A)
 \end{array}
 \qquad
 \begin{array}{ccc}
 FA \oplus J & \xrightarrow{1_{FA} \oplus \iota} & FA \oplus FI \\
 \rho_{FA} \downarrow & & \downarrow \varphi_{A, I} \\
 FA & \xleftarrow{F\rho_A} & F(A \otimes I)
 \end{array}$$

We say F is *strong monoidal* (respectively *strict monoidal*) if φ and ι are isomorphisms (respectively identities). An *oplax monoidal* functor is the same definition, but where the directions of the maps φ and ι are reversed.

Note that the same letters are used for the associators and unitors in both monoidal categories.

Example. (i) The forgetful functor $U : (\mathbf{AbGp}, \otimes, \mathbb{Z}) \rightarrow (\mathbf{Set}, \times, 1)$ is lax monoidal. We define $\iota : 1 \rightarrow \mathbb{Z}$ to map the element of 1 to the generator $1 \in \mathbb{Z}$, and define $\varphi : UA \times UB \rightarrow U(A \otimes B)$ by $(a, b) \mapsto a \otimes b$. One can easily verify that the required diagrams commute.

(ii) The free functor $F : (\mathbf{Set}, \times, 1) \rightarrow (\mathbf{AbGp}, \otimes, \mathbb{Z})$ is strong monoidal, because $F1 \cong \mathbb{Z}$ and $F(A \times B) \cong FA \otimes FB$.

(iii) Let R be a commutative ring. Then the forgetful functor $\mathbf{Mod}_R \rightarrow \mathbf{AbGp}$ is lax monoidal, where $\iota : \mathbb{Z} \rightarrow R$ is the natural map, and $\varphi : A \otimes_{\mathbb{Z}} B \rightarrow A \otimes_R B$ is the quotient map. Its left adjoint, the free functor $\mathbf{AbGp} \rightarrow \mathbf{Mod}_R$, is strong monoidal.

(iv) If \mathcal{C} and \mathcal{D} have the cartesian monoidal structure, then any functor $F : \mathcal{C} \rightarrow \mathcal{D}$ is oplax monoidal. $\iota : F1 \rightarrow 1$ is the unique morphism to the terminal object of \mathcal{D} , and $\varphi_{A,B} : F(A \times B) \rightarrow FA \times FB$ is given by $(F\pi_1, F\pi_2)$. F is strong monoidal if and only if it preserves finite products.

(v) If X and Y are metric spaces, then $1_{X \times Y}$ is non-expansive as a map $(X \times Y, d_1) \rightarrow (X \times Y, d_\infty)$, making the identity functor $1_{\mathbf{Met}}$ into a monoidal functor $(\mathbf{Met}, \times_\infty, 1) \rightarrow (\mathbf{Met}, \times_1, 1)$. Note that the d_∞ metric on $X \times Y$ defines the categorical product.

Lemma. Let \mathcal{C} and \mathcal{D} be monoidal categories. Let $F \dashv G$, where $F : \mathcal{C} \rightarrow \mathcal{D}$ and $G : \mathcal{D} \rightarrow \mathcal{C}$. Then there is a bijection between lax monoidal structures on G and oplax monoidal structures on F .

Proof sketch. Suppose we have (φ, ι) on G . Then the transpose of $\iota : J \rightarrow GI$ is a morphism

$FJ \rightarrow I$, and we have a natural transformation

$$F(A \otimes B) \xrightarrow{F(\eta_A \times \eta_B)} F(GFA \otimes GFB) \xrightarrow{F\varphi_{FA, FB}} FG(FA \oplus FB) \xrightarrow{\epsilon_{FA \oplus FB}} FA \oplus FB$$

One can check that each of the required diagrams commute, defining an oplax monoidal structure on F . By duality, an oplax monoidal structure on F yields a lax monoidal structure on G , and it can be shown that these constructions are inverse to each other. \square

6.4. Closed monoidal categories

Definition. We say that a monoidal category $(\mathcal{C}, \otimes, I)$ is *(left/right/bi)-closed* if $A \otimes (-)$, $(-) \otimes A$, or both have right adjoints for all A . If \otimes is symmetric, we say in any of these cases that \mathcal{C} is *closed*.

Right adjoints for $(-) \otimes A$ are denoted $[A, -]$ if they exist.

- Example.** (i) A cartesian closed category is a monoidal category with $\otimes = \times$, that is closed as a monoidal category. In particular, **Set** and **Cat** are cartesian closed.
- (ii) The metric d_1 on the set $[X, Y]$ of non-expansive maps $X \rightarrow Y$ yields a closed structure on $(\mathbf{Met}, \times_1, 1)$.
- (iii) **AbGp** and **Mod_R** for any commutative ring R are monoidal closed, where $[A, B]$ is the set of homomorphisms $A \rightarrow B$, turned into an abelian group or R -module by pointwise addition and scalar multiplication. The homomorphisms $C \rightarrow [A, B]$ correspond under λ -conversion to bilinear maps $C \times A \rightarrow B$, and thus to homomorphisms $C \otimes_R A \rightarrow B$.
- (iv) The cartesian monoidal structure on the category of pointed sets **Set_{*}** is not closed, but the monoidal structure given by the *smash product* $(-) \wedge (-)$ is closed, where

$$(A, a_0) \wedge (B, b_0) = A \times B / \sim$$

and \sim identifies all elements where either coordinate is the basepoint. Basepoint-preserving maps $A \wedge B \rightarrow C$ correspond to basepoint-preserving maps from A to the set $[B, C]$ of basepoint-preserving maps $B \rightarrow C$.

- (v) Consider the set $\text{Rel}(A \times A) = P(A \times A)$ of relations on A . This is a poset under inclusion, and is a monoid under relational composition. Composition is order-preserving in each variable, making $\text{Rel}(A \times A)$ into a strict monoidal category. It is not symmetric, but biclosed. For the right adjoint to $(-) \circ R$, we define $R \Rightarrow T$ to be

$$(R \Rightarrow T) = \{(b, c) \in A \times A \mid \forall a \in A, (a, b) \in R \Rightarrow (a, c) \in T\}$$

Then $S \subseteq (R \Rightarrow T)$ if and only if $S \circ R \subseteq T$.

6.5. Enriched categories

Definition. Let $(\mathcal{E}, \otimes, I)$ be a monoidal category. An \mathcal{E} -enriched category consists of

- (i) a collection $\text{ob } \mathcal{C}$ of objects;
- (ii) an object $\mathcal{C}(A, B)$ of \mathcal{E} for each pair of objects $A, B \in \text{ob } \mathcal{C}$;
- (iii) morphisms $\iota_A : I \rightarrow \mathcal{C}(A, A)$ for each A ;
- (iv) morphisms $\kappa_{A,B,C} : \mathcal{C}(B, C) \otimes \mathcal{C}(A, B) \rightarrow \mathcal{C}(A, C)$ for objects A, B, C ,

such that the following diagrams commute.

$$\begin{array}{ccc}
 I \otimes \mathcal{C}(A, B) & \xrightarrow{\iota_B \otimes 1_{\mathcal{C}(A, B)}} & \mathcal{C}(B, B) \otimes \mathcal{C}(A, B) \\
 & \searrow \lambda_{\mathcal{C}(A, B)} & \downarrow \kappa_{A, B, B} \\
 & & \mathcal{C}(A, B)
 \end{array}$$

$$\begin{array}{ccc}
 \mathcal{C}(A, B) \otimes I & \xrightarrow{1_B \otimes \iota_A} & \mathcal{C}(A, B) \otimes \mathcal{C}(A, A) \\
 & \searrow \rho_{\mathcal{C}(A, B)} & \downarrow \kappa_{A, A, B} \\
 & & \mathcal{C}(A, B)
 \end{array}$$

$$\begin{array}{ccc}
 (\mathcal{C}(C, D) \otimes \mathcal{C}(B, C)) \otimes \mathcal{C}(A, B) & \xrightarrow{\kappa \otimes 1} & \mathcal{C}(B, D) \otimes \mathcal{C}(A, B) \\
 \downarrow \alpha & & \searrow \kappa \\
 \mathcal{C}(C, D) \otimes (\mathcal{C}(B, C) \otimes \mathcal{C}(A, B)) & \xrightarrow{1 \otimes \kappa} & \mathcal{C}(C, D) \otimes \mathcal{C}(A, C) \\
 & & \nearrow \kappa \\
 & & \mathcal{C}(A, D)
 \end{array}$$

Definition. Let \mathcal{C}, \mathcal{D} be \mathcal{E} -enriched categories. An \mathcal{E} -enriched functor $\mathcal{C} \rightarrow \mathcal{D}$ consists of a map of objects $F : \text{ob } \mathcal{C} \rightarrow \text{ob } \mathcal{D}$ together with morphisms $F_{A,B} : \mathcal{C}(A, B) \rightarrow \mathcal{D}(FA, FB)$ for each pair of objects $A, B \in \text{ob } \mathcal{C}$, in such a way that is compatible with identities and composition.

Definition. Let $F, G : \mathcal{C} \rightarrow \mathcal{D}$ be \mathcal{E} -enriched functors between \mathcal{E} -enriched categories. An \mathcal{E} -enriched natural transformation $F \rightarrow G$ assigns a morphism $\theta_A : I \rightarrow \mathcal{D}(FA, GA)$ to each $A \in \text{ob } \mathcal{C}$, satisfying the naturality condition

$$\begin{array}{ccccc}
 \mathcal{C}(A, B) & \xrightarrow{F_{A,B}} & \mathcal{D}(FA, FB) & \xrightarrow{\lambda^{-1}} & I \otimes \mathcal{D}(FA, FB) \\
 G_{A,B} \downarrow & & & & \downarrow \theta_B \otimes 1 \\
 \mathcal{D}(GA, GB) & & & & \mathcal{D}(FB, GB) \otimes \mathcal{D}(FA, FB) \\
 \rho^{-1} \downarrow & & & & \downarrow \kappa \\
 \mathcal{D}(GA, GB) \otimes I & \xrightarrow{1 \otimes \theta_A} & \mathcal{D}(GA, GB) \otimes \mathcal{D}(FA, GA) & \xrightarrow{\kappa} & \mathcal{D}(FA, GB)
 \end{array}$$

6. Monoidal and enriched categories

If \mathcal{C} is an \mathcal{E} -enriched category, its *underlying ordinary category* $|\mathcal{C}|$ is the category where the objects are those of \mathcal{C} , the morphisms $A \rightarrow B$ are the morphisms $I \rightarrow \mathcal{C}(A, B)$ in \mathcal{E} , where the identity morphisms are given by ι_A , and the composition of $g : C \rightarrow B$ and $f : A \rightarrow B$ given by

$$I \xrightarrow{\lambda_I^{-1}} I \otimes I \xrightarrow{g \otimes f} \mathcal{C}(B, C) \otimes \mathcal{C}(A, B) \xrightarrow{\kappa} \mathcal{C}(A, C)$$

One can check that this indeed forms a category. An \mathcal{E} -enrichment of an ordinary category \mathcal{C}_0 is an \mathcal{E} -enriched category \mathcal{C} such that $|\mathcal{C}| \cong \mathcal{C}_0$.

Example. (i) A category enriched over $(\mathbf{Set}, \times, 1)$ is a locally small category.

(ii) A category enriched over the poset $2 = \{0, 1\}$ with $0 < 1$ is a preorder.

(iii) A category enriched over $(\mathbf{Cat}, \times, \mathbf{1})$ is a *2-category*. Its morphisms or *1-arrows* $A \rightarrow B$ are the objects of a category $\mathcal{C}(A, B)$. It has *2-arrows* between parallel pairs $f, g : A \rightrightarrows B$, which are the morphisms $f \rightarrow g$ in the category $\mathcal{C}(A, B)$. \mathbf{Cat} is a 2-category, by taking the 2-arrows to be the natural transformations. The category of small \mathcal{E} -enriched categories with \mathcal{E} -enriched functors is a 2-category.

(iv) A category enriched over $(\mathbf{AbGp}, \otimes, \mathbb{Z})$ is an *additive category*.

(v) If \mathcal{E} is a right closed monoidal category, it has a canonical enrichment structure over itself. Take $\mathcal{E}(A, B)$ to be $[A, B]$, where $[A, -]$ is the right adjoint of $(-) \otimes A$. The identity $I \rightarrow [A, A]$ is the transpose $\lambda_A : I \otimes A \rightarrow A$, and the composition κ is the transpose of

$$([B, C] \otimes [A, B]) \otimes A \xrightarrow{\alpha} [B, C] \otimes ([A, B] \otimes A) \xrightarrow{1 \otimes \text{ev}} [B, C] \otimes B \xrightarrow{\text{ev}} C$$

where ev is the evaluation map, which is precisely the counit of the adjunction.

(vi) A one-object \mathcal{E} -enriched category is an (*internal*) *monoid* in \mathcal{E} ; it consists of an object M of \mathcal{E} , equipped with morphisms $e : I \rightarrow M$ and $m : M \otimes M \rightarrow M$ satisfying the left and right unit laws and the associativity law.

(a) An internal monoid in \mathbf{Set} is a monoid.

(b) An internal monoid in \mathbf{AbGp} is a ring.

(c) An internal monoid in \mathbf{Cat} is a strict monoidal category.

(d) An internal monoid in $[\mathcal{C}, \mathcal{C}]$ is a monad on \mathcal{C} .

7. Additive and abelian categories

7.1. Additive categories

In this section, we will study categories enriched over $(\mathbf{AbGp}, \otimes, \mathbb{Z})$; these are called *additive* categories. We will also consider other weaker enrichments: a category enriched over $(\mathbf{Set}_*, \wedge, 2)$ is called *pointed*, and a category enriched over $(\mathbf{CMon}, \otimes, \mathbb{N})$, where \mathbf{CMon} is the category of commutative monoids, is called *semi-additive*.

In a pointed category \mathcal{C} , each $\mathcal{C}(A, B)$ has a distinguished element 0, and all composites with zero morphisms are zero morphisms. In a semi-additive category \mathcal{C} , each $\mathcal{C}(A, B)$ has a binary addition operation which is associative, commutative, and has an identity 0. Composition in a semi-additive category is bilinear, so $(f + g)(h + k) = fh + gh + fk + gk$ whenever the composites are defined. In an additive category, each morphism $f \in \mathcal{C}(A, B)$ has an additive inverse $-f \in \mathcal{C}(A, B)$.

Lemma. (i) For an object A in a pointed category \mathcal{C} , the following are equivalent.

- (a) A is a terminal object of \mathcal{C} .
- (b) A is an initial object of \mathcal{C} .
- (c) $1_A = 0 : A \rightarrow A$.

(ii) For objects A, B, C in a semi-additive category \mathcal{C} , the following are equivalent.

- (a) there exist morphisms $\pi_1 : C \rightarrow A$ and $\pi_2 : C \rightarrow B$ making C into a product of A and B ;
- (b) there exist morphisms $\nu_1 : A \rightarrow C$ and $\nu_2 : B \rightarrow C$ making C into a coproduct of A and B ;
- (c) there exist morphisms $\pi_1 : C \rightarrow A, \pi_2 : C \rightarrow B, \nu_1 : A \rightarrow C, \nu_2 : B \rightarrow C$ satisfying

$$\pi_1 \nu_1 = 1_A; \quad \pi_2 \nu_2 = 1_B; \quad \pi_1 \nu_2 = 0; \quad \pi_2 \nu_1 = 0; \quad \nu_1 \pi_1 + \nu_2 \pi_2 = 1_C$$

Proof. In each part, as (a) and (b) are dual and (c) is self-dual, it suffices to prove the equivalence of (a) and (c).

Part (i). If A is terminal, then it has exactly one morphism $A \rightarrow A$, so this must be the zero morphism. Conversely, if $1_A = 0$, then A is terminal, as for any $f : B \rightarrow A$, we have $f = 1_A f = 0 f = 0$, so the only morphism $B \rightarrow A$ is the zero morphism.

Part (ii). If (a) holds, take ν_1, ν_2 to be defined by the first four equations in (c); it suffices to verify the last equation, $\nu_1 \pi_1 + \nu_2 \pi_2 = 1_C$. Composing with π_1 ,

$$\pi_1 \nu_1 \pi_1 = 1_A \pi_1 + 0 \pi_2 = \pi_1$$

and similarly, composing with π_2 gives π_2 . So by uniqueness of factorisations through limit cones, $\nu_1 \pi_1 + \nu_2 \pi_2$ must be the identity. Conversely, if (c) holds, given a pair $f : D \rightarrow A$ and

7. Additive and abelian categories

$g : D \rightarrow B$, the morphism

$$h = \nu_1 f + \nu_2 g$$

satisfies

$$\pi_1 h = 1_A f + 0g = f; \quad \pi_2 h = 0f + 1_A g = g$$

giving a factorisation, and if h' also satisfies these equations, then

$$h' = (\nu_1 \pi_1 + \nu_2 \pi_2) h' = \nu_1 f + \nu_2 g = h$$

so the factorisation is unique. \square

In any category, an object which is both initial and terminal is called a *zero object*, denoted 0 . An object that is a product and a coproduct of A and B is called a *biproduct*, denoted $A \oplus B$.

Lemma. Let \mathcal{C} be a locally small category.

- (i) If \mathcal{C} has a zero object, then it has a unique pointed structure.
- (ii) Suppose \mathcal{C} has a zero object and has binary products and coproducts. Suppose further that for each pair $A, B \in \text{ob } \mathcal{C}$, the canonical morphism $c : A + B \rightarrow A \times B$ defined by

$$\pi_i c \nu_j = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$

is an isomorphism. Then \mathcal{C} has a unique semi-additive structure.

We adopt the convention that morphisms into a product are denoted with column vectors, and morphisms out of a coproduct are denoted with row vectors.

Proof. Part (i). The unique morphism $0 \rightarrow 0$ is both the identity and a zero morphism. So for any two $A, B : \text{ob } \mathcal{C}$, the unique composite $A \rightarrow 0 \rightarrow B$ must be the zero element of $\mathcal{C}(A, B)$. We can define a pointed structure on \mathcal{C} in this way.

Part (ii). This technique is known as the *Eckmann–Hilton argument*. Given $f, g : A \rightrightarrows B$, we define the *left sum* $f +_\ell g$ to be the composite

$$A \xrightarrow{\begin{pmatrix} f \\ g \end{pmatrix}} B \times B \xrightarrow{c^{-1}} B + B \xrightarrow{\begin{pmatrix} 1 & 1 \end{pmatrix}} B$$

and the *right sum* $f +_r g$ to be

$$A \xrightarrow{\begin{pmatrix} 1 \\ 1 \end{pmatrix}} A \times A \xrightarrow{c^{-1}} B + B \xrightarrow{\begin{pmatrix} f & g \end{pmatrix}} B$$

I. Category Theory

Note that $(f +_\ell g)h = fh +_\ell gh$, since

$$\begin{pmatrix} f \\ g \end{pmatrix} h = \begin{pmatrix} fh \\ gh \end{pmatrix}$$

and similarly,

$$k(f +_r g) = kf +_r kg$$

So if we show that the two sums coincide, we obtain the required distributive laws. First, note that $0 : A \rightarrow B$ is a two-sided identity for both $+_\ell$ and $+_r$. For example, $f +_\ell 0 = f$, since

$$\begin{array}{ccccc} A & \xrightarrow{f} & B & \xrightarrow{1_B} & B \\ & \searrow \begin{pmatrix} f \\ 0 \end{pmatrix} & \swarrow \begin{pmatrix} 1 \\ 0 \end{pmatrix} & \searrow \nu_1 & \swarrow \begin{pmatrix} 1 & 1 \end{pmatrix} \\ & B \times B & \xrightarrow{\begin{pmatrix} 1 \\ 0 \end{pmatrix} c^{-1}} & B + B & \end{array}$$

commutes. Suppose we have morphisms $f, g, h, k : A \rightarrow B$, and consider the composite

$$A \xrightarrow{\begin{pmatrix} 1 \\ 1 \end{pmatrix}} A \times A \xrightarrow{c^{-1}} A + A \xrightarrow{\begin{pmatrix} f & g \\ h & k \end{pmatrix}} B \times B \xrightarrow{c^{-1}} B + B \xrightarrow{\begin{pmatrix} 1 & 1 \end{pmatrix}} B$$

The composite of the first three factors is

$$\begin{pmatrix} f +_r g \\ h +_r k \end{pmatrix}$$

so the whole composite is $(f +_r g) +_\ell (h +_r k)$. Evaluating from other end, we obtain

$$(f +_r g) +_\ell (h +_r k) = (f +_\ell h) +_r (g +_\ell k)$$

This is known as the *interchange law*. Substituting $g = k = 0$, we obtain $f +_\ell k = f +_r k$. Substituting $f = k = 0$ (and dropping the subscripts) we obtain the commutative law $g + h = h + g$. Substituting $h = 0$, we obtain the associativity law $(f + g) + k = f + (g + k)$.

For uniqueness, suppose we have some semi-additive structure $+$ on \mathcal{C} . Then $\nu_1\pi_1 + \nu_2\pi_2$ must be the inverse of $c = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} : A + B \rightarrow A \times B$, since

$$\nu_1\pi_1 c = \nu_1(1 \ 0) = (\nu_1 \ 0); \quad \nu_2\pi_2 c = (0 \ \nu_2)$$

so

$$(\nu_1\pi_1 + \nu_2\pi_2)c = (\nu_1 + 0 \ 0 + \nu_2) = (\nu_1 \ \nu_2) = 1_{A+B}$$

Hence the definitions of $+_\ell$ and $+_r$ both reduce to $+$. □

Note that if \mathcal{C} and \mathcal{D} are semi-additive categories with finite biproducts, then a functor $F : \mathcal{C} \rightarrow \mathcal{D}$ is semi-additive (that is, enriched over **CMon**) if and only if it preserves either finite products or finite coproducts. In particular, if F has either a left or right adjoint, then it is semi-additive, and the adjunction is enriched over **CMon**; the bijection $\mathcal{C}(A, GB) \rightarrow \mathcal{D}(FA, B)$ is an isomorphism of commutative monoids, since the operations $F(-)$ and $(-)\epsilon_B$ both respect addition.

7.2. Kernels and cokernels

Definition. Let $f : A \rightarrow B$ be a morphism in a pointed category \mathcal{C} . The *kernel* of f is the equaliser of the pair $(f, 0)$; dually the *cokernel* is the coequaliser of $(f, 0)$. A monomorphism that occurs as the kernel of a morphism is called *normal*.

In an additive category, the normal monomorphisms are precisely the regular monomorphisms, since the equaliser of (f, g) is the kernel of $f - g$. In \mathbf{Gp} , all inclusions of subgroups are regular, but not all inclusions are normal, since a normal monomorphism corresponds to a normal subgroup. In \mathbf{Set}_* , all surjections are regular epimorphisms, but $(A, a_0) \rightarrow (B, b_0)$ is a normal epimorphism if f is bijective on elements not mapped to b_0 . We say that a morphism $f : A \rightarrow B$ is a *pseudomonomorphism* if its kernel is a zero morphism; that is, $fg = 0$ implies $g = 0$.

Lemma. In a pointed category with kernels and cokernels, $f : A \rightarrow B$ is normal monic if and only if $f \cong \ker \operatorname{coker} f$.

Proof. If $f \cong \ker \operatorname{coker} f$, it is clearly normal. Now suppose $f = \ker g$. Then g factors through the cokernel of f , so $g(\ker \operatorname{coker} f) = 0$. Thus $\ker \operatorname{coker} f \leq f$ in $\operatorname{Sub}(B)$. But $(\operatorname{coker} f)f = 0$, so $f \leq \ker \operatorname{coker} f$, so they are isomorphic as subobjects of B . \square

Corollary. In a pointed category with kernels and cokernels, the operations \ker and coker induce an order-reversing bijection between isomorphism classes of normal subobjects and isomorphism classes of normal quotients of any object.

Remark. For any morphism $f : A \rightarrow B$ in such a category, $\ker \operatorname{coker} f$ is the smallest normal subobject of B through which f factors.

7.3. Abelian categories

Definition. An *abelian category* is an additive category with all finite limits and colimits. Equivalently, an abelian category is a category with a zero object, finite biproducts, kernels, and cokernels, such that all monomorphisms and epimorphisms are normal.

Example. (i) The category \mathbf{AbGp} is abelian; more generally, for any ring R , the category \mathbf{Mod}_R is abelian.

(ii) If \mathcal{A} is abelian and \mathcal{C} is small, then $[\mathcal{C}, \mathcal{A}]$ is abelian, with all structures defined point-wise.

(iii) If \mathcal{A} is abelian and \mathcal{C} is small and additive, then the category of additive functors $\mathcal{C} \rightarrow \mathcal{A}$, denoted $\operatorname{Add}(\mathcal{C}, \mathcal{A})$, is also abelian, as it is closed under all of the structures on $[\mathcal{C}, \mathcal{A}]$. Note that this covers the case of R -modules, as an additive category with a single object is a ring, and the category of modules over such a ring is isomorphic to the category of additive functors from this category to \mathbf{AbGp} .

I. Category Theory

Remark. If $f : A \rightarrow B$ in an abelian category, then $\ker \operatorname{coker} f$ is the smallest subobject $I \rightarrowtail B$ through which f factors. This is called the *image* of f , denoted $\operatorname{im} f = \ker \operatorname{coker} f$. The other part of the factorisation $A \rightarrow I$ is epic, as it cannot factor through the equaliser of any nonequal parallel pair $I \rightrightarrows C$. Thus, it is also the smallest quotient of A through which f factors, so it is the *coimage* of f , given by $\operatorname{coim} f = \operatorname{coker} \ker f$. The composition $A \twoheadrightarrow I \rightarrowtail B$ is the unique epi-mono factorisation of f .

To show that this factorisation is stable under pullback, it suffices to show that the pullback of an epimorphism in an abelian category is epic, as the corresponding statement for monomorphisms has already been shown.

Lemma (flattening lemma). Consider a square

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ g \downarrow & & \downarrow h \\ C & \xrightarrow{k} & D \end{array}$$

in an abelian category \mathcal{A} . Its *flattening* is the sequence

$$A \xrightarrow{\begin{pmatrix} f \\ g \end{pmatrix}} B \oplus C \xrightarrow{\begin{pmatrix} h & -k \end{pmatrix}} D$$

Then

- (i) the square commutes if and only if the composite of the flattening $\begin{pmatrix} h & -k \end{pmatrix} \begin{pmatrix} f \\ g \end{pmatrix}$ is the zero morphism;
- (ii) the square is a pullback if and only if $\begin{pmatrix} f \\ g \end{pmatrix} = \ker \begin{pmatrix} h & -k \end{pmatrix}$;
- (iii) the square is a pushout if and only if $\begin{pmatrix} h & -k \end{pmatrix} = \operatorname{coker} \begin{pmatrix} f \\ g \end{pmatrix}$.

Proof. Part (i). The composite $\begin{pmatrix} h & -k \end{pmatrix} \begin{pmatrix} f \\ g \end{pmatrix}$ is $hf - kg$, so it vanishes if and only if the square commutes.

Part (ii). $\begin{pmatrix} f \\ g \end{pmatrix}$ is the kernel of $\begin{pmatrix} h & -k \end{pmatrix}$ if and only if

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ g \downarrow & & \\ & & C \end{array}$$

is universal among spans completing the cospan

$$\begin{array}{ccc} & B & \\ & \downarrow h & \\ C & \xrightarrow{k} & D \end{array}$$

into a commutative square.

Part (iii). Follows by duality, taking care of the asymmetric negation. \square

Corollary. In an abelian category \mathcal{A} , epimorphisms are stable under pullback.

Proof. Suppose we have a pullback square

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ g \downarrow & & \downarrow h \\ C & \xrightarrow{k} & D \end{array}$$

By part (ii) of the above result, $\begin{pmatrix} f \\ g \end{pmatrix} = \ker(h - k)$. But h is an epimorphism, so $(h - k)$ is also an epimorphism. Thus $(h - k) = \operatorname{coker}\begin{pmatrix} f \\ g \end{pmatrix}$, so the square is also a pushout. We show that g is a pseudoepimorphism; this suffices as \mathcal{A} is abelian. Suppose we have $\ell : C \rightarrow E$ with $\ell g = 0$. Then $\begin{pmatrix} \ell & (B \xrightarrow{0} E) \end{pmatrix}$ factors uniquely through the pushout.

$$\begin{array}{ccccc} A & \xrightarrow{f} & B & & \\ g \downarrow & & \downarrow h & \searrow 0 & \\ C & \xrightarrow{k} & D & \xrightarrow{m} & E \\ & \searrow \ell & & \nearrow & \end{array}$$

But then $mh = 0$ and h is epic, so $m = 0$, giving $\ell = mk = 0$. \square

Thus image factorisations are stable under pullback, and dually, under pushout.

7.4. Exact sequences

Definition. A sequence

$$\cdots \longrightarrow A_{n+1} \xrightarrow{f_{n+1}} A_n \xrightarrow{f_n} A_{n-1} \longrightarrow \cdots$$

in an abelian category \mathcal{A} is *exact* at A_n if $\ker f_n = \operatorname{im} f_{n+1}$. The entire sequence is said to be *exact* if it is exact at every vertex.

I. Category Theory

By duality, the sequence is exact at A_n if and only if $\text{coker } f_{n+1} = \text{coim } f_n$.

Example.

$$0 \longrightarrow A \xrightarrow{f} B \xrightarrow{g} C$$

is exact at A if and only if f is monic, and is exact at A and B if and only if $f = \ker g$.

Definition. A functor between abelian categories $F : \mathcal{A} \rightarrow \mathcal{B}$ is *exact* if it preserves arbitrary exact sequences.

This implies that F preserves kernels and cokernels, and the converse is true as images are defined in terms of kernels and cokernels.

Definition. F is *left exact* if it preserves exact sequences of the form

$$0 \longrightarrow A \xrightarrow{f} B \xrightarrow{g} C$$

Proposition. Let $F : \mathcal{A} \rightarrow \mathcal{B}$ be a functor between abelian categories. Then

- (i) F is left exact if and only if it preserves all finite limits (and hence is additive);
- (ii) F is exact if and only if it is left exact and preserves epimorphisms.

Proof. Part (i). One direction is trivial as kernels are finite limits. Conversely, note that for any A, B , the sequence

$$0 \longrightarrow A \xrightarrow{\begin{pmatrix} 1 \\ 0 \end{pmatrix}} A \oplus B \xrightarrow{\begin{pmatrix} 0 & 1 \end{pmatrix}} B \longrightarrow 0$$

is exact, and conversely, if we have an exact sequence

$$0 \longrightarrow A \xrightarrow{f} C \xrightarrow{g} B \longrightarrow 0$$

and either f is a split monomorphism or g is a split epimorphism, then $C \cong A \oplus B$. Indeed, suppose that f is split, so $rf = 1_A$. Then $g = \text{coker } f = \text{coker } fr$ is the coequaliser of $(1_C - fr, 1_C)$, so it is the epic part of a splitting of the idempotent $1_C - fr$. If $s : B \rightarrow C$ is the monic part of this splitting, then the four morphisms (r, g, f, s) satisfy the equations of a biproduct. So F maps

$$0 \longrightarrow A \xrightarrow{\begin{pmatrix} 1 \\ 0 \end{pmatrix}} A \oplus B \xrightarrow{\begin{pmatrix} 0 & 1 \end{pmatrix}} B \longrightarrow 0$$

to a sequence identifying $F(A \oplus B)$ as $FA \oplus FB$, and thus preserves biproducts. Hence F preserves all finite limits.

Part (ii). If F is left exact and preserves epimorphisms, then it preserves the exactness of sequences of the form

$$0 \longrightarrow A \xrightarrow{f} C \xrightarrow{g} B \longrightarrow 0$$

Thus it preserves kernels and cokernels. \square

7.5. The five lemma

Lemma. Suppose we have a commutative diagram in an abelian category

$$\begin{array}{ccccccccc} A_1 & \xrightarrow{f_1} & A_2 & \xrightarrow{f_2} & A_3 & \xrightarrow{f_3} & A_4 & \xrightarrow{f_4} & A_5 \\ \downarrow u_1 & & \downarrow u_2 & & \downarrow u_3 & & \downarrow u_4 & & \downarrow u_5 \\ B_1 & \xrightarrow{g_1} & B_2 & \xrightarrow{g_2} & B_3 & \xrightarrow{g_3} & B_4 & \xrightarrow{g_4} & B_5 \end{array}$$

where the rows are exact sequences. Then,

- (i) if u_1 is epic and u_2, u_4 are monic, then u_3 is monic;
- (ii) if u_5 is monic and u_2, u_4 are epic, then u_3 is epic.

Thus if u_1, u_2, u_4, u_5 are isomorphisms, u_3 is an isomorphism.

Proof. By duality it suffices to show (i). We show u_3 is a pseudomonomorphism. Suppose we have $x : C \rightarrow A_3$ with $u_3x = 0$. Then $u_4f_3x = g_4u_3x = 0$, so as u_4 is a monomorphism, $f_3x = 0$. Hence x factors through the kernel of f_3 , which is the image of f_2 . Form the pullback of f_2 and x to obtain

$$\begin{array}{ccccccccc} & & D & \xrightarrow{y} & C & & & & \\ & & \downarrow z & & \downarrow x & & & & \\ A_1 & \xrightarrow{f_1} & A_2 & \xrightarrow{f_2} & A_3 & \xrightarrow{f_3} & A_4 & \xrightarrow{f_4} & A_5 \\ \downarrow u_1 & & \downarrow u_2 & & \downarrow u_3 & & \downarrow u_4 & & \downarrow u_5 \\ B_1 & \xrightarrow{g_1} & B_2 & \xrightarrow{g_2} & B_3 & \xrightarrow{g_3} & B_4 & \xrightarrow{g_4} & B_5 \end{array}$$

Then y is also the pullback of this factorisation of x along $\text{coim } f_2$, so y is an epimorphism as epimorphisms are stable under pullback. Then $g_2u_2z = u_3f_2z = u_3xy = 0$. Thus u_2z factors through $\ker g_2 = \text{im } g_1$. Consider the pullback square

$$\begin{array}{ccc} E & \xrightarrow{v} & D \\ w \downarrow & & \downarrow u_2z \\ A_1 & \xrightarrow{g_1u_1} & B_2 \end{array}$$

I. Category Theory

So v is epic, as it is the pullback of $\text{coim}(g_1 u_1)$.

$$\begin{array}{ccccccccc}
 E & \xrightarrow{v} & D & \xrightarrow{y} & C & & & & \\
 \downarrow w & & \downarrow z & & \downarrow x & & & & \\
 A_1 & \xrightarrow{f_1} & A_2 & \xrightarrow{f_2} & A_3 & \xrightarrow{f_3} & A_4 & \xrightarrow{f_4} & A_5 \\
 \downarrow u_1 & & \downarrow u_2 & & \downarrow u_3 & & \downarrow u_4 & & \downarrow u_5 \\
 B_1 & \xrightarrow{g_1} & B_2 & \xrightarrow{g_2} & B_3 & \xrightarrow{g_3} & B_4 & \xrightarrow{g_4} & B_5
 \end{array}$$

Thus $u_2 z v = g_1 u_1 w$, and u_2 is monic, so $z v = f_1 w$. Then $x y v = f_2 z v = f_2 f_1 w = 0$, and $y v$ is epic, hence $x = 0$. \square

7.6. The snake lemma

Lemma. Consider a diagram in an abelian category

$$\begin{array}{ccccccc}
 B_1 & \longrightarrow & B_2 & \longrightarrow & B_3 & \longrightarrow & 0 \\
 \downarrow v_1 & & \downarrow v_2 & & \downarrow v_3 & & \\
 0 & \longrightarrow & C_1 & \longrightarrow & C_2 & \longrightarrow & C_3
 \end{array}$$

where the rows are exact and the squares commute. Then we obtain an exact sequence

$$\begin{array}{ccccccc}
 \text{Ker } v_1 & \longrightarrow & \text{Ker } v_2 & \longrightarrow & \text{Ker } v_3 & & \\
 \downarrow & & \downarrow & & \downarrow & & \\
 B_1 & \longrightarrow & B_2 & \longrightarrow & B_3 & \longrightarrow & 0 \\
 \downarrow v_1 & & \downarrow v_2 & & \downarrow v_3 & & \\
 0 & \longrightarrow & C_1 & \longrightarrow & C_2 & \longrightarrow & C_3 \\
 \downarrow & & \downarrow & & \downarrow & & \\
 \text{Coker } v_1 & \longrightarrow & \text{Coker } v_2 & \longrightarrow & \text{Coker } v_3 & &
 \end{array}$$

\xrightarrow{s}

7.7. Complexes in abelian categories

Definition. Let \mathcal{A} be an abelian category. A *(chain) complex* in \mathcal{A} is an infinite sequence of objects and morphisms

$$\cdots \longrightarrow C_{n+1} \xrightarrow{d_{n+1}} C_n \xrightarrow{d_n} C_{n-1} \longrightarrow \cdots$$

where the composite of any two consecutive morphisms is zero.

7. Additive and abelian categories

Note that a complex may be identified with an additive functor $\mathcal{Z} \rightarrow \mathcal{A}$, where \mathcal{Z} is the additive category with $\text{ob } \mathcal{Z} = \mathbb{Z}$ and

$$\mathcal{Z}(n, m) = \begin{cases} \mathbb{Z} & \text{if } m = n \text{ or } m = n - 1 \\ 0 & \text{otherwise} \end{cases}$$

Thus, complexes on \mathcal{A} are the objects of an abelian category $\text{CA} = \text{Add}(\mathcal{Z}, \mathcal{A})$, where the morphisms are natural transformations.

Definition. Let C_\bullet be a complex. We define

- (i) $Z_n(C_\bullet) \rightarrow C_n$ to be the kernel of d_n ;
- (ii) $I_n(C_\bullet) \rightarrow C_n$ to be the image of d_{n+1} ;
- (iii) $Z_n(C_\bullet) \rightarrow H_n(C_\bullet)$ to be the cokernel of $I_n(C_\bullet) \rightarrow Z_n(C_\bullet)$.

We say that $H_n(C_\bullet)$ is the n th homology object of C_\bullet .

Note that Z_n, I_n, H_n are additive functors $\text{CA} \rightarrow \mathcal{A}$.

Lemma. The construction of $H_n(C_\bullet)$ is self-dual.

Proof. Write $C_n \twoheadrightarrow Q_n(C_\bullet)$ for the cokernel of d_{n+1} . Then we have the diagram

$$\begin{array}{ccccccc} C_{n+1} & \xrightarrow{d_{n+1}} & C_n & \xrightarrow{d_n} & C_{n-1} \\ \downarrow & & \nearrow & & \uparrow \\ I_n & \rightarrow & Z_n & \twoheadrightarrow & H_n & \twoheadrightarrow & Q_n & \twoheadrightarrow & I_{n-1} \end{array}$$

By definition, $I_n \rightarrow C_n$ is $\ker(C_n \rightarrow Q_n)$. As $Z_n \rightarrow C_n$ is a monomorphism, $I_n \rightarrow Z_n$ is $\ker(Z_n \rightarrow C_n \rightarrow Q_n)$. Hence $Z_n \rightarrow H_n$ is $\text{coim}(Z_n \rightarrow Q_n)$, so we obtain

$$\begin{array}{ccccccc} C_{n+1} & \xrightarrow{d_{n+1}} & C_n & \xrightarrow{d_n} & C_{n-1} \\ \downarrow & & \nearrow & & \uparrow \\ I_n & \rightarrow & Z_n & \twoheadrightarrow & H_n & \twoheadrightarrow & Q_n & \twoheadrightarrow & I_{n-1} \end{array}$$

and $Z_n \twoheadrightarrow H_n \twoheadrightarrow Q_n$ is the image factorisation of $Z_n \rightarrow Q_n$. □

Theorem (Mayer-Vietoris sequence). Suppose we have a short exact sequence of complexes in \mathcal{A} .

$$0 \longrightarrow A_\bullet \xrightarrow{f_\bullet} B_\bullet \xrightarrow{g_\bullet} C_\bullet \longrightarrow 0$$

Then there is a long exact sequence of homology objects

$$\cdots \longrightarrow H_n(A_\bullet) \xrightarrow{H_n(f_\bullet)} H_n(B_\bullet) \xrightarrow{H_n(g_\bullet)} H_n(C_\bullet) \longrightarrow H_{n-1}(A_\bullet) \xrightarrow{H_{n-1}(f_\bullet)} H_{n-1}(B_\bullet) \xrightarrow{H_{n-1}(g_\bullet)} H_{n-1}(C_\bullet) \longrightarrow \cdots$$

I. Category Theory

Proof. First, we apply the snake lemma to

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A_{n+1} & \xrightarrow{f_{n+1}} & B_{n+1} & \xrightarrow{g_{n+1}} & C_{n+1} & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & A_n & \xrightarrow{f_n} & B_n & \xrightarrow{g_n} & C_n & \longrightarrow & 0 \end{array}$$

to obtain exact sequences

$$0 \longrightarrow Z_{n+1}(A_\bullet) \longrightarrow Z_{n+1}(B_\bullet) \longrightarrow Z_{n+1}(C_\bullet)$$

and

$$Q_n(A_\bullet) \longrightarrow Q_n(B_\bullet) \longrightarrow Q_n(C_\bullet) \longrightarrow 0$$

Thus Z_n is a left exact functor and Q_n is right exact. We now apply the snake lemma again to the diagram

$$\begin{array}{ccccccccc} Q_{n+1}(A_\bullet) & \longrightarrow & Q_{n+1}(B_\bullet) & \longrightarrow & Q_{n+1}(C_\bullet) & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & Z_n(A_\bullet) & \longrightarrow & Z_n(B_\bullet) & \longrightarrow & Z_n(C_\bullet) \end{array}$$

Here, the cokernel of $Q_{n+1} \rightarrow Z_n$ coincides with that of $I_n \rightarrow Z_n$ as $Q_{n+1} \rightarrow I_n$ is epic. Their kernels coincide with $H_{n+1} \rightarrow Q_{n+1}$ as homology is self-dual. Hence we obtain

$$H_{n+1}(A_\bullet) \longrightarrow H_{n+1}(B_\bullet) \longrightarrow H_{n+1}(C_\bullet) \longrightarrow H_n(A_\bullet) \longrightarrow H_n(B_\bullet) \longrightarrow H_n(C_\bullet)$$

as required. \square

Note that $Z_n : \mathcal{CA} \rightarrow \mathcal{A}$ is the right adjoint to the functor $A \mapsto A[n]$, where $A[n]$ is the complex that has A in dimension n and 0 everywhere else; this gives another proof that Z is left exact. Dually, Q_n is the left adjoint to this functor.

Definition. Let $f_\bullet, g_\bullet : C_\bullet \rightrightarrows D_\bullet$ be two morphisms of \mathcal{CA} . A *homotopy* from f_\bullet to g_\bullet is a sequence of morphisms $h_n : C_n \rightarrow D_{n+1}$ such that

$$g_n - f_n = d_{n+1}h_n + h_{n-1}d_n$$

for all n . We say that f_\bullet, g_\bullet are *homotopic* and write $f_\bullet \simeq g_\bullet$ if there exists such a sequence h_\bullet .

Homotopy is an equivalence relation on morphisms of \mathcal{CA} . It is a congruence, as it is compatible with composition on both sides; indeed, if $k_\bullet : D_\bullet \rightarrow E_\bullet$, and $h_\bullet : f_\bullet \simeq g_\bullet$, then the morphisms $k_{n+1}h_n$ form a homotopy $k_\bullet f_\bullet \rightarrow k_\bullet g_\bullet$, and similarly for the other side. We write \mathcal{HA} for the quotient of \mathcal{CA} by the homotopy congruence. Also, homotopy is compatible with addition, by adding the relevant homotopies, so the quotient category inherits an additive structure, and the quotient $\mathcal{CA} \rightarrow \mathcal{HA}$ is an additive functor. In particular, \mathcal{HA} has finite biproducts, although it is not an abelian category.

Lemma. If $f. \simeq g. : C. \rightrightarrows D.$, then $H_n(f.) = H_n(g.)$ for all n .

Thus, the H_n can be regarded as additive functors $\mathcal{HA} \rightarrow \mathcal{A}$.

Proof. Let $h.$ be a homotopy from $f.$ to $g.$, so $g_n - f_n = d_{n+1}h_n + h_{n-1}d_n$. Then $Z_n(g.) - Z_n(f.)$ is the restriction of $d_{n+1}h_n$ to $Z_n(C.)$, since $h_{n-1}d_n$ is zero on this subobject. Similarly, $H_n(g.) - H_n(f.)$ is zero, as $d_{n+1}h_n$ vanishes when factoring through the quotient. \square

7.8. Projective resolutions

Definition. A category \mathcal{C} has *enough projectives* if for every object A , there exists an epimorphism $P \twoheadrightarrow A$ where P is projective.

Note that this holds in **AbGp** and **Mod_R** for any commutative ring R , because free modules are projective, and every module can be written as a quotient of a free module.

Definition. Let \mathcal{A} be an abelian category and let A be an object of \mathcal{A} . A *projective resolution* of A is a complex $P.$ where the objects P_n are projective, $P_n = 0$ for all $n < 0$, and

$$H_n(P.) = \begin{cases} A & \text{if } n = 0 \\ 0 & \text{otherwise} \end{cases}$$

Equivalently, a projective resolution is an exact sequence

$$\cdots \longrightarrow P_2 \longrightarrow P_1 \longrightarrow P_0 \longrightarrow A \longrightarrow 0$$

where the P_i are projective.

Lemma. Let \mathcal{A} be an abelian category that has enough projectives. Then every object of \mathcal{A} has a projective resolution.

Proof. Given an object A , choose some projective object P_0 with an epimorphism $P_0 \twoheadrightarrow A$. Let $K_0 \hookrightarrow P_0$ be its kernel, and choose P_1 to be a projective object with an epimorphism $P_1 \twoheadrightarrow K_0$, then continue by induction. \square

Lemma. Suppose $P., Q.$ are projective resolutions of objects A, B . Then for any $f : A \rightarrow B$, there is a morphism of complexes $f. : P. \rightarrow Q.$ with $H.(f.) = f$. Moreover, any two such morphisms $P. \rightarrow Q.$ are homotopic.

Proof. Consider the diagram

$$\begin{array}{ccccccccc} P_2 & \longrightarrow & K_1 & \longrightarrow & P_1 & \longrightarrow & K_0 & \longrightarrow & P_0 & \longrightarrow & A \\ & & & & & & & & & & \downarrow f \\ Q_2 & \longrightarrow & L_1 & \longrightarrow & Q_1 & \longrightarrow & L_0 & \longrightarrow & Q_0 & \longrightarrow & B \end{array}$$

I. Category Theory

By projectivity of P_0 , we obtain f_0 completing the right-hand square.

$$\begin{array}{ccccccccc}
 P_2 & \longrightarrow & K_1 & \longrightarrow & P_1 & \longrightarrow & K_0 & \longrightarrow & P_0 & \longrightarrow & A \\
 & & & & & & & & \downarrow f_0 & & \downarrow f \\
 Q_2 & \longrightarrow & L_1 & \longrightarrow & Q_1 & \longrightarrow & L_0 & \longrightarrow & Q_0 & \longrightarrow & B
 \end{array}$$

The morphism $P_1 \rightarrow P_0 \rightarrow A$ is zero by exactness, so $P_1 \rightarrow P_0 \rightarrow Q_0 \rightarrow B$ is also zero. Thus $P_1 \rightarrow Q_0$ factors through the kernel $L_0 \rightarrow Q_0$. We then obtain f_1 by projectivity.

$$\begin{array}{ccccccccc}
 P_2 & \longrightarrow & K_1 & \longrightarrow & P_1 & \longrightarrow & K_0 & \longrightarrow & P_0 & \longrightarrow & A \\
 & & & & \downarrow f_1 & \searrow & & & \downarrow f_0 & & \downarrow f \\
 Q_2 & \longrightarrow & L_1 & \longrightarrow & Q_1 & \longrightarrow & L_0 & \longrightarrow & Q_0 & \longrightarrow & B
 \end{array}$$

Continue by induction.

Now suppose we have another morphism of chains g_* with $H_0(g_*) = f$. Then $g_0 - f_0$ factors through $L_0 \rightarrow Q_0$ as they have the same composite with $Q_0 \rightarrow B$. Thus we obtain

$$\begin{array}{ccccccccc}
 P_2 & \longrightarrow & K_1 & \longrightarrow & P_1 & \longrightarrow & K_0 & \longrightarrow & P_0 & \longrightarrow & A \\
 & & & & \downarrow & \searrow & & & \downarrow & & \downarrow f \\
 Q_2 & \longrightarrow & L_1 & \longrightarrow & Q_1 & \longrightarrow & L_0 & \longrightarrow & Q_0 & \longrightarrow & B
 \end{array}$$

$\swarrow h_0$

where $d'_1 h_0 = g_0 - f_0$. Then

$$d'_1(g_1 - f_1 - h_0 d_1) = d'_1 g_1 - d'_1 f_1 - d'_1 h_0 d_1 = g_0 d_1 - f_0 d_1 - d'_1 h_0 d_1 = 0$$

Hence $g_1 - f_1 - h_0 d_1$ factors through $L_1 \rightarrow Q_1$, so we obtain h_1 as follows.

$$\begin{array}{ccccccccc}
 P_2 & \longrightarrow & K_1 & \longrightarrow & P_1 & \longrightarrow & K_0 & \longrightarrow & P_0 & \longrightarrow & A \\
 & & & & \downarrow & \searrow & & & \downarrow & & \downarrow f \\
 Q_2 & \longrightarrow & L_1 & \longrightarrow & Q_1 & \longrightarrow & L_0 & \longrightarrow & Q_0 & \longrightarrow & B
 \end{array}$$

$\swarrow h_1$

Then $d'_2 h_1 + h_0 d_1 = g_1 - f_1$ as required. Continue similarly by induction to construct all components of the homotopy. \square

Thus construction of projective resolution is a functor. Note that in this proof we never made use of projectivity of Q_* . In particular, this shows that the construction of projective resolutions is left adjoint to $H_0 : \mathcal{C} \rightarrow \mathcal{A}$ where $\mathcal{C} \subseteq \mathbf{HA}$ is the full subcategory on complexes C_* for which $H_n(C_*) = 0$ for all $n > 0$.

7.9. Derived functors

Let $F : \mathcal{A} \rightarrow \mathcal{B}$ be an additive functor between abelian categories. Then F extends to a functor $\mathcal{C}F : \mathcal{C}\mathcal{A} \rightarrow \mathcal{C}\mathcal{B}$ which respects homotopy. Hence F induces a functor $\mathcal{H}F : \mathcal{H}\mathcal{A} \rightarrow \mathcal{H}\mathcal{B}$.

Definition. Let $F : \mathcal{A} \rightarrow \mathcal{B}$ be an additive functor between abelian categories, and suppose \mathcal{A} has enough projectives. Then the *left derived functor* $L^n F$ of F is the composite

$$\mathcal{A} \xrightarrow{\text{PR}} \mathcal{H}\mathcal{A} \xrightarrow{\mathcal{H}F} \mathcal{H}\mathcal{B} \xrightarrow{H_n} \mathcal{B}$$

for any $n \geq 0$, where PR is the projective resolution functor.

Note that if F is exact, we have $L^0 F \cong F$ and $L^n F = 0$ for $n > 0$. More generally, if F is right exact, then it preserves exactness of

$$P_1 \longrightarrow P_0 \longrightarrow A \longrightarrow 0$$

for any projective resolution P_\bullet of A . In particular, $L^0 F \cong F$ in this case.

Lemma. Let

$$0 \longrightarrow A \xrightarrow{f} B \xrightarrow{g} C \longrightarrow 0$$

be a short exact sequence in an abelian category \mathcal{A} with enough projectives. Then we can choose projective resolutions $P_\bullet, Q_\bullet, R_\bullet$ of A, B, C and morphisms f_\bullet, g_\bullet extending f, g making the sequence

$$0 \longrightarrow P_\bullet \xrightarrow{f_\bullet} Q_\bullet \xrightarrow{g_\bullet} R_\bullet \longrightarrow 0$$

exact. Moreover, the exactness of this sequence is preserved by arbitrary additive functors.

Proof. We choose P_\bullet, R_\bullet arbitrarily, and take $Q_n = P_n \oplus R_n$; this is projective as the coproduct of projective objects is projective. Consider the diagram

$$\begin{array}{ccccccc} \cdots & \longrightarrow & P_1 & \longrightarrow & K_0 & \longrightarrow & P_0 \xrightarrow{e_1} A \\ & & & & & & \downarrow \begin{pmatrix} 1 \\ 0 \end{pmatrix} & & \downarrow f \\ & & & & & & P_0 \oplus R_0 & & B \\ & & & & & & \downarrow \begin{pmatrix} 0 & 1 \end{pmatrix} & & \downarrow g \\ \cdots & \longrightarrow & R_1 & \longrightarrow & M_0 & \longrightarrow & R_0 \xrightarrow{e_3} C \end{array}$$

By projectivity of R_0 , we obtain $h : R_0 \rightarrow B$, and so we define $e_2 = (fe_1 \ h)$.

$$\begin{array}{ccccccc} \cdots & \longrightarrow & P_1 & \longrightarrow & K_0 & \longrightarrow & P_0 \xrightarrow{e_1} A \\ & & & & & & \downarrow & & \downarrow f \\ & & & & & & P_0 \oplus R_0 & \xrightarrow{e_2} & B \\ & & & & & & \downarrow & \nearrow h & \downarrow g \\ \cdots & \longrightarrow & R_1 & \longrightarrow & M_0 & \longrightarrow & R_0 & \xrightarrow{e_3} & C \end{array}$$

I. Category Theory

This makes both right-hand squares commute:

$$e_2 \begin{pmatrix} 1 \\ 0 \end{pmatrix} = fe_1; \quad ge_2 = (gfe \quad gh) = (0 \quad e_3)$$

To show e_2 is epic, suppose we have a morphism $k : B \rightarrow D$ such that $ke_2 = 0$.

$$\begin{array}{ccccccc} \cdots & \longrightarrow & P_1 & \longrightarrow & K_0 & \longrightarrow & P_0 \xrightarrow{e_1} A \\ & & & & \downarrow & & \downarrow f \\ & & & & P_0 \oplus R_0 & \xrightarrow{e_2} & B \xrightarrow{k} D \\ & & & & \downarrow & & \downarrow g \\ \cdots & \longrightarrow & R_1 & \longrightarrow & M_0 & \longrightarrow & R_0 \xrightarrow{e_3} C \end{array}$$

Then $kfe_1 = 0$, so k factors as ℓg for some ℓ .

$$\begin{array}{ccccccc} \cdots & \longrightarrow & P_1 & \longrightarrow & K_0 & \longrightarrow & P_0 \xrightarrow{e_1} A \\ & & & & \downarrow & & \downarrow f \\ & & & & P_0 \oplus R_0 & \xrightarrow{e_2} & B \xrightarrow{k} D \\ & & & & \downarrow & & \downarrow g \\ \cdots & \longrightarrow & R_1 & \longrightarrow & M_0 & \longrightarrow & R_0 \xrightarrow{e_3} C \end{array} \quad \begin{array}{c} \nearrow \ell \end{array}$$

Now $\ell e_3 \begin{pmatrix} 0 & 1 \end{pmatrix} = \ell ge_2 = ke_2 = 0$, so $\ell = 0$ as e_3 and $\begin{pmatrix} 0 & 1 \end{pmatrix}$ are pseudoepimorphisms. Thus $k = 0$. Forming the kernel, we obtain

$$\begin{array}{ccccccc} \cdots & \longrightarrow & P_1 & \longrightarrow & K_0 & \longrightarrow & P_0 \xrightarrow{e_1} A \\ & & & & \downarrow & & \downarrow f \\ & & & & L_0 & \longrightarrow & P_0 \oplus R_0 \xrightarrow{e_2} B \\ & & & & \downarrow & & \downarrow g \\ \cdots & \longrightarrow & R_1 & \longrightarrow & M_0 & \longrightarrow & R_0 \xrightarrow{e_3} C \end{array}$$

Applying the snake lemma to the diagram

$$\begin{array}{ccccccc}
 & & & & 0 & & \\
 & & & & \downarrow & & \\
 K_0 & \longrightarrow & P_0 & \xrightarrow{e_1} & A & & \\
 \downarrow & & \downarrow & & \downarrow f & & \\
 L_0 & \longrightarrow & P_0 \oplus R_0 & \xrightarrow{e_2} & B & & \\
 \downarrow & & \downarrow & & \downarrow g & & \\
 M_0 & \longrightarrow & R_0 & \xrightarrow{e_3} & C & & \\
 & & \downarrow & & & & \\
 & & 0 & & & &
 \end{array}$$

the left-hand column extends to a short exact sequence.

$$0 \longrightarrow K_0 \longrightarrow L_0 \longrightarrow M_0 \longrightarrow 0$$

Hence, as before, we can define an epimorphism $P_1 \oplus R_1 \rightarrow L_0$ making the two left-hand squares commute.

$$\begin{array}{ccccccc}
 \dots & \longrightarrow & P_1 & \longrightarrow & K_0 & \longrightarrow & P_0 \xrightarrow{e_1} A \\
 & & \left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix} \right) \downarrow & & \downarrow & & \downarrow \\
 & & P_1 \oplus R_1 & \longrightarrow & L_0 & \longrightarrow & P_0 \oplus R_0 \xrightarrow{e_2} B \\
 & & \left(\begin{smallmatrix} 0 & 1 \end{smallmatrix} \right) \downarrow & & \downarrow & & \downarrow \\
 \dots & \longrightarrow & R_1 & \longrightarrow & M_0 & \longrightarrow & R_0 \xrightarrow{e_3} C
 \end{array}$$

Continue by induction. As the columns

$$0 \longrightarrow P_n \longrightarrow Q_n \longrightarrow R_n \longrightarrow 0$$

are biproduct diagrams, they are preserved by arbitrary additive functors. \square

This proof does not show that $Q_n \cong P_n \oplus R_n$ in \mathcal{CA} . Indeed, if it were, then $d'_n : Q_n \rightarrow Q_{n-1}$ would have matrix

$$\begin{pmatrix} d_n & 0 \\ 0 & d''_n \end{pmatrix}$$

where $d_n : P_n \rightarrow P_{n-1}$ and $d''_n : R_n \rightarrow R_{n-1}$. Our construction above was of the form

$$\begin{pmatrix} d_n & x \\ 0 & d''_n \end{pmatrix}$$

I. Category Theory

Theorem. Let $F : \mathcal{A} \rightarrow \mathcal{B}$ be an additive functor between abelian categories, and suppose \mathcal{A} has enough projectives. Then, for any short exact sequence

$$0 \longrightarrow A \xrightarrow{f} B \xrightarrow{g} C \longrightarrow 0$$

in \mathcal{A} , we obtain an exact sequence

$$\cdots \longrightarrow L^1FA \longrightarrow L^1FB \longrightarrow L^1FC \longrightarrow L^0FA \longrightarrow L^0FB \longrightarrow L^0FC \longrightarrow 0$$

Proof. Choose projective resolutions $P_., Q_., R_.$ for A, B, C as above. Then applying F , we obtain an exact sequence of complexes

$$0 \longrightarrow FP_. \longrightarrow FQ_. \longrightarrow FR_. \longrightarrow 0$$

in \mathcal{B} . Then the result follows from the Mayer–Vietoris sequence. □

In particular, L^0F is always right exact, so $L^0F \cong F$ if and only if F is right exact.

II. Commutative Algebra

Lectured in Michaelmas 2023 by DR. O. BECKER

(Course description goes here.)

Contents

1.	Chain conditions	88
1.1.	Modules	88
1.2.	Noetherian and Artinian modules	88
1.3.	Exact sequences	90
1.4.	Algebras	91
2.	Tensor products	93
2.1.	Introduction	93
2.2.	Definition and universal property	93
2.3.	Zero tensors	95
2.4.	Monoidal structure	97
2.5.	Tensor products of maps	102
2.6.	Tensor products of algebras	103
2.7.	Restriction and extension of scalars	105
2.8.	Extension of scalars on morphisms	108
2.9.	Extension of scalars in algebras	109
2.10.	Exactness properties of the tensor product	109
2.11.	Flat modules	113
3.	Localisation	119
3.1.	Definitions	119
3.2.	Universal property for rings	120
3.3.	Functoriality	121
3.4.	Universal property for modules	123
3.5.	Exactness	124
3.6.	Extension and contraction of ideals	125
3.7.	Local properties	128
3.8.	Localisations as quotients	130
4.	Integrality, finiteness, and finite generation	132
4.1.	Nakayama's lemma	132
4.2.	Integral and finite extensions	133
4.3.	Integral closure	136
4.4.	Noether normalisation	138
4.5.	Hilbert's Nullstellensatz	139
4.6.	Integrality over ideals	142
4.7.	Cohen–Seidenberg theorems	144
5.	Primary decomposition	147
6.	Direct and inverse limits	150
6.1.	Limits and completions	150

6.2.	Graded rings and modules	152
6.3.	Artin–Rees lemma	154
7.	Dimension theory	156
7.1.	???	156
7.2.	Hilbert polynomials	157
7.3.	Dimension theory of local Noetherian rings	159

1. Chain conditions

1.1. Modules

In this course, a *ring* is taken to mean a commutative unital ring R . We do however allow for one noncommutative exception, the endomorphism ring $\text{End}(M)$ of an abelian group M . This is a ring where composition is the multiplication operation.

Definition. An R -module is an abelian group M with a fixed ring homomorphism $\rho : R \rightarrow \text{End}(M)$. If $r \in R$ and $m \in M$, we define $r \cdot m = \rho(r)(m)$.

Remark. Note that as $\rho(r)$ is a group homomorphism,

$$r(m_1 + m_2) = \rho(r)(m_1 + m_2) = \rho(r)(m_1) + \rho(r)(m_2) = r \cdot m_1 + r \cdot m_2$$

Also, as ρ is a ring homomorphism,

$$(r_1 + r_2)m = \rho(r_1 + r_2)(m) = (\rho(r_1) + \rho(r_2))m = r_1 \cdot m + r_2 \cdot m$$

Example. (i) Let k be a field. Then a k -module is a k -vector space.

(ii) Every abelian group M is a \mathbb{Z} -module in a unique way, because the morphism $\mathbb{Z} \rightarrow \text{End } M$ must map 1 to id.

(iii) Every ring R is an R -module, by taking $\rho(r) = r_0 \mapsto r_0 r$.

Definition. The *direct product* of abelian groups $(M_i)_{i \in I}$ is the set of I -tuples $(a_i)_{i \in I}$ where $a_i \in M_i$, with elementwise addition as the group operation.

Definition. The *direct sum* of abelian groups $(M_i)_{i \in I}$ is the set of I -tuples $(a_i)_{i \in I}$ where $a_i \in M_i$ and all but finitely many of the a_i are zero, again with elementwise addition as the group operation.

Direct products are written $\prod_{i \in I} M_i$, and direct sums are written $\bigoplus_{i \in I} M_i$. These constructions coincide if the index set I is finite. Direct products and direct sums of R -modules are also R -modules.

The universal property of the direct sum states that each collection of module homomorphisms $\varphi_i : M_i \rightarrow R$ can be combined into a unique homomorphism $\varphi : \bigoplus_{i \in I} M_i \rightarrow R$. Similarly, the universal property of the direct product states that each collection of module homomorphisms $\varphi_i : R \rightarrow M_i$ can be combined into a unique homomorphism $\varphi : R \rightarrow \prod_{i \in I} M_i$.

1.2. Noetherian and Artinian modules

Definition. An R -module M is *Noetherian* if one of the following conditions holds.

(i) Every ascending chain of submodules $M_0 \subseteq M_1 \subseteq \dots$ inside M stabilises. That is, for some k , every $j \in \mathbb{N}$ has $M_{k+j} = M_k$.

(ii) Every nonempty set Σ of submodules of M has a maximal element.

Lemma. The two conditions above are equivalent.

Proof. (i) implies (ii). Let Σ be a nonempty set of submodules of M . If it has no maximal element, then for each $M' \in \Sigma$ there exists $M'' \in \Sigma$ with $M' \subsetneq M''$. We can then use the axiom of choice to pick a sequence $M_0 \subsetneq M_1 \subsetneq M_2 \subsetneq \dots$ of elements in Σ . This contradicts (i).

(ii) implies (i). Let $M_0 \subseteq M_1 \subseteq \dots$ be an ascending chain of submodules in M . Then let $\Sigma = \{M_0, M_1, \dots\}$. This has a maximal element M_k by (ii). Then for all $j \in \mathbb{N}$, $M_{k+j} = M_k$ as required. \square

Definition. M is *Artinian* if one of the following conditions holds.

(i) Every descending chain of submodules $M_0 \supseteq M_1 \supseteq \dots$ inside M stabilises.

(ii) Every nonempty set Σ of submodules of M has a minimal element.

Again, both conditions are equivalent.

Lemma. An R -module M is Noetherian if and only if every submodule of M is finitely generated.

Proof. Suppose M is Noetherian, and let $N \subseteq M$ be a submodule. Pick $m_1 \in N$, and consider the submodule $M_1 \subseteq N$ generated by m_1 . If $M_1 = N$, then we are done. Otherwise, pick $m_2 \in M_1 \setminus N$, and consider $M_2 \subseteq N$ generated by m_2 . This construction will always terminate, as if it did not, we would have constructed an infinite strictly ascending chain of submodules of M , contradicting that M is Noetherian.

Now suppose every submodule of M is finitely generated, and let $M_0 \subseteq M_1 \subseteq \dots$ be an ascending chain of submodules of M . Let $N = \bigcup_{i=0}^{\infty} M_i$; this is a submodule of M as the M_i form a chain. Then N is finitely generated, say, by generators $m_1, \dots, m_k \in N$. As the M_i form a chain increasing to N , there exists n such that $m_1, \dots, m_k \in M_n$. In particular, $N \subseteq M_n \subseteq N$, so $M_n = N$. Thus the chain stabilises. \square

Note that every Noetherian module is finitely generated. Let $R = \mathbb{Z}[T_1, T_2, \dots]$, and let $M = R$ as an R -module. M is generated by 1_R , so in particular it is finitely generated. But it has a \mathbb{Z} -submodule $\langle T_1, T_2, \dots \rangle$ that is not finitely generated. So in the above lemma we indeed must check every submodule.

Definition. A ring R is Noetherian (respectively Artinian) if R is Noetherian (resp. Artinian) as an R -module.

Example. (i) \mathbb{Z} over itself is a Noetherian module as it is a principal ideal domain, but it is not an Artinian module because we can take the chain $(2) \supsetneq (4) \supsetneq (8) \supsetneq \dots$.

(ii) \mathbb{Z} is similarly a Noetherian ring but not an Artinian ring by unfolding the definition and using (i).

II. Commutative Algebra

- (iii) $\mathbb{Z}\left[\frac{1}{2}\right]/\mathbb{Z}$ is an Artinian \mathbb{Z} -module but not a Noetherian \mathbb{Z} -module. This can be seen from the fact that the only submodules are of the form $\left(\frac{1}{2^k} + \mathbb{Z}\right)$ for $k \in \mathbb{N}$.
- (iv) In fact, a ring R is Artinian if and only if R is Noetherian and R has *Krull dimension* 0.

1.3. Exact sequences

Definition. A sequence

$$\cdots \longrightarrow M_{i-1} \xrightarrow{f_i} M_i \xrightarrow{f_{i+1}} M_{i+1} \longrightarrow \cdots$$

is *exact* if the image of f_i is equal to the kernel of f_{i+1} for each i , where the M_i are modules and the f_i are module homomorphisms.

Definition. A *short exact sequence* is an exact sequence of the form

$$0 \longrightarrow M' \xrightarrow{\text{injective}} M \xrightarrow{\text{surjective}} M'' \longrightarrow 0$$

In this situation, $M'' \simeq M/i(M')$. This is a way to encode M'' as a quotient by a submodule.

Lemma. Let

$$0 \longrightarrow N \xrightarrow{\iota} M \xrightarrow{\varphi} L \longrightarrow 0$$

be a short exact sequence of R -modules. Then M is Noetherian (resp. Artinian) if and only if both N and L are Noetherian (resp. Artinian).

Proof. We show the statement for Noetherian modules.

Suppose M is Noetherian. If $N_0 \subseteq N_1 \subseteq \cdots$ is an ascending chain of submodules inside N , then by taking images,

$$\iota(N_0) \subseteq \iota(N_1) \subseteq \cdots$$

is also naturally an ascending chain of submodules inside M , so it stabilises. As ι is injective, the original sequence also stabilises. Hence N is Noetherian.

If $L_0 \subseteq L_1 \subseteq \cdots$ is an ascending chain of submodules inside L , then by taking preimages,

$$\varphi^{-1}(L_0) \subseteq \varphi^{-1}(L_1) \subseteq \cdots$$

is an ascending chain of submodules inside M , where

$$\varphi^{-1}(L_i) = \{m \in M \mid \varphi(m) \in L_i\}$$

So this chain stabilises at $\varphi^{-1}(L_k)$. But as φ is surjective, $\varphi(\varphi^{-1}(L_i)) = L_i$, so the original sequence must stabilise at L_k .

1. Chain conditions

Now suppose N and L are Noetherian, and let $M_0 \subseteq M_1 \subseteq \dots$ be an ascending chain of submodules in M . Then

$$\iota^{-1}(M_0) \subseteq \iota^{-1}(M_1) \subseteq \dots$$

is an ascending chain of submodules in N , so stabilises at $\iota^{-1}(M_{k_N})$ for some k_N . Similarly,

$$\varphi(M_0) \subseteq \varphi(M_1) \subseteq \dots$$

is an ascending chain of submodules in L , so stabilises at $\varphi^{-1}(M_{k_L})$ for some k_L . Take $k \geq k_N, k_L$, and let $j \geq 0$. We show $M_{k+j} \subseteq M_k$, proving that the sequence stabilises.

Let $m \in M_{k+j}$. As $\varphi(M_{k+j}) = \varphi(M_k)$, there exists $m' \in M_k$ such that $\varphi(m) = \varphi(m')$. Then $\varphi(m - m') = 0$, so by exactness, $m - m'$ is in the image of ι , say, $\iota(x) = m - m'$. Since $m - m' \in M_{k+j}$, we must have $x \in \iota^{-1}(M_{k+j})$. But then $x \in \iota^{-1}(M_k)$, so $\iota(x) = m - m' \in M_k$. Hence $m \in M_k$. \square

Corollary. If M_1, \dots, M_n are Noetherian (resp. Artinian) modules, then so is $M_1 \oplus \dots \oplus M_n$.

Proof. Consider the sequence

$$0 \longrightarrow M_1 \xrightarrow{\iota} M_1 \oplus M_2 \xrightarrow{\pi} M_2 \longrightarrow 0$$

where $\iota(x) = (x, 0)$ and $\pi(x, y) = y$. This is exact, so $M_1 \oplus M_2$ is Noetherian. We then proceed by induction on n . \square

Proposition. For a Noetherian (resp. Artinian) ring R , every finitely generated R -module is Noetherian (resp. Artinian).

Proof. M is finitely generated if and only if there is a surjective module homomorphism $\varphi : R^n \rightarrow M$ for some $n \geq 0$. That is, M is a quotient of R^n . The fact that R^n is Noetherian (or Artinian) passes through to its quotients. \square

1.4. Algebras

Definition. An R -algebra is a ring A together with a fixed ring homomorphism $\rho : R \rightarrow A$.

Example. The map $k \rightarrow k[T_1, \dots, T_n]$ makes the polynomial ring $k[T_1, \dots, T_n]$ a k -algebra.

We will write $ra = \rho(r)a$. Note that $\rho(r) = \rho(r) \cdot 1_A = r \cdot 1_A$, so we can write $r \cdot 1_A$ for $\rho(r)$.

Remark. Every R -algebra is an R -module.

Example. As a k -module, $k[T_1, \dots, T_n]$ is infinite-dimensional. As a k -algebra, $k[T_1, \dots, T_n]$ is generated by the n elements T_1, \dots, T_n .

Definition. $\varphi : A \rightarrow B$ is an R -algebra homomorphism if φ is a ring homomorphism and preserves all elements of R . That is, $\varphi(r \cdot 1_A) = r \cdot 1_B$.

II. Commutative Algebra

An R -algebra A is finitely generated if and only if there is some $n \geq 0$ and a surjective algebra homomorphism $R[T_1, \dots, T_n] \rightarrow A$.

Theorem (Hilbert's basis theorem). Every finitely generated algebra A over a Noetherian ring R is Noetherian.

For example, the polynomial algebra over a field is Noetherian.

Proof. It suffices to prove this for a polynomial ring, as every finitely generated algebra is a quotient of a polynomial ring. It further suffices to prove this for a univariate polynomial ring $A = R[T]$ by induction. Let \mathfrak{a} be an ideal of $R[T]$; we need to show that \mathfrak{a} is finitely generated. For each $i \geq 0$, define

$$\mathfrak{a}(i) = \{c_0 \mid c_0 T^i + \dots + c_i T^0 \in \mathfrak{a}\}$$

Thus $\mathfrak{a}(i)$ is the set of leading coefficients of polynomials of degree i that lie in \mathfrak{a} . Each $\mathfrak{a}(i)$ is an ideal in R , and $\mathfrak{a}(i) \subseteq \mathfrak{a}(i+1)$ by multiplying by T . As R is Noetherian, each $\mathfrak{a}(i)$ is a finitely generated ideal, and this ascending chain stabilises at $\mathfrak{a}(m)$, say. Let

$$\mathfrak{a}(i) = (b_{i,1}, \dots, b_{i,n_i})$$

We can choose $f_{i,j}$ of degree i with leading coefficient $b_{i,j}$. Define the ideal

$$\mathfrak{b} = (f_{i,j})_{i \leq m, j \leq n_i}$$

Note that \mathfrak{b} is finitely generated. Defining $\mathfrak{b}(i)$ in the same way as $\mathfrak{a}(i)$, we have

$$\forall i, \mathfrak{a}(i) = \mathfrak{b}(i)$$

By construction, $\mathfrak{b} \subseteq \mathfrak{a}$; we claim that the reverse inclusion holds, then the proof will be complete. Suppose that $\mathfrak{a} \not\subseteq \mathfrak{b}$, and take $f \in \mathfrak{a} \setminus \mathfrak{b}$ of minimal degree i . As $\mathfrak{a}(i) = \mathfrak{b}(i)$, there is a polynomial g in \mathfrak{b} of degree i that has the same leading coefficient. Then $f - g$ has degree less than i , and lies in \mathfrak{a} . But then by minimality, $f - g \in \mathfrak{b}$, giving $f \in \mathfrak{b}$. \square

Therefore, if $S \subseteq R[T_1, \dots, T_n]/I$ where R is Noetherian, then $(S) = (S_0)$ where $S_0 \subseteq S$ is finite.

2. Tensor products

2.1. Introduction

Let M and N be R -modules. Informally, the tensor product of M and N over R is the set $M \otimes_R N$ of all sums

$$\sum_{i=1}^{\ell} m_i \otimes n_i; \quad m_i \in M, n_i \in N$$

subject to the relations

$$\begin{aligned} (m_1 + m_2) \otimes n &= m_1 \otimes n + m_2 \otimes n \\ m \otimes (n_1 + n_2) &= m \otimes n_1 + m \otimes n_2 \\ (rm) \otimes n &= r(m \otimes n) \\ m \otimes (rn) &= r(m \otimes n) \end{aligned}$$

This is a module that abstracts the notion of bilinearity between two modules.

Example. Consider $\mathbb{Z}/2\mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Z}/3\mathbb{Z}$. In this \mathbb{Z} -module,

$$x \otimes y = (3x) \otimes y = x \otimes (3y) = x \otimes 0 = x \otimes (0 \cdot 0) = 0(x \otimes 0) = 0$$

Hence $\mathbb{Z}/2\mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Z}/3\mathbb{Z} = 0$.

Example. Now consider $\mathbb{R}^n \otimes_{\mathbb{R}} \mathbb{R}^{\ell}$. We will show later that this is isomorphic to $\mathbb{R}^{n\ell}$.

2.2. Definition and universal property

Definition. A map of R -modules $f : M \times N \rightarrow L$ is R -bilinear if for each $m_0 \in M$ and $n_0 \in N$, the maps $n \mapsto f(m_0, n)$ and $m \mapsto f(m, n_0)$ are R -linear (or equivalently, a homomorphism of R -modules).

Definition. Let M, N be R -modules. Let $\mathcal{F} = R^{\oplus(M \times N)}$ be the free R -module with coordinates indexed by $M \times N$. Define $K \subseteq \mathcal{F}$ to be the submodule generated by the following set of relations:

$$\begin{aligned} (m_1 + m_2, n) - (m_1, n) - (m_2, n) \\ (m, n_1 + n_2) - (m, n_1) - (m, n_2) \\ r(m, n) - (rm, n) \\ r(m, n) - (m, rn) \end{aligned}$$

The tensor product $M \otimes_R N$ is \mathcal{F}/K . We further define the R -bilinear map

$$i_{M \otimes N} : M \times N \rightarrow M \otimes N; \quad i_{M \otimes N}(m, n) = e_{(m, n)} = m \otimes n$$

II. Commutative Algebra

Proposition (universal property of the tensor product). The pair $(M \otimes_R N, i_{M \otimes_R N})$ satisfies the following universal property. For every R -module L and every R -bilinear map $f : M \times N \rightarrow L$, there exists a unique homomorphism $h : M \otimes_R N \rightarrow L$ such that the following diagram commutes.

$$\begin{array}{ccc} M \times N & \xrightarrow{i_{M \otimes_R N}} & M \otimes_R N \\ & \searrow f & \downarrow h \\ & & L \end{array}$$

Equivalently, $h \circ i_{M \otimes_R N} = f$.

Proof. The conclusion $h \circ i_{M \otimes_R N} = f$ holds if and only if for all m, n , we have

$$h(m \otimes n) = f(m, n)$$

Note that the elements $\{m \otimes n\}$ generate $M \otimes N$ as an R -module, so there is at most one h . We now show that the definition of h on the pure tensors $m \otimes n$ extends to an R -linear map $M \otimes N \rightarrow L$. The map $R^{\oplus(M \times N)} \rightarrow L$ given by $(m, n) \mapsto f(m, n)$ exists by the universal property of the direct sum. However, this map vanishes on the generators of K , so it factors through the quotient \mathcal{F}/K as required. \square

The universal property given above characterises the tensor product up to isomorphism.

Proposition. Let M, N be R -modules, and (T, j) be an R -module and an R -bilinear map $M \times N \rightarrow T$. Suppose that (T, j) satisfies the same universal property as $M \otimes N$. Then there is a unique isomorphism of R -modules $\varphi : M \otimes N \xrightarrow{\sim} T$ such that $\varphi \circ i_{M \otimes N} = j$.

Proof. By using the universal property of $M \otimes N$ and T , we obtain φ and ψ as follows.

$$\begin{array}{ccc} M \otimes N & \xleftarrow{\psi} & T \\ & \nwarrow i_{M \otimes N} \quad \nearrow j & \\ & M \times N & \end{array}$$

The universal property states that $\varphi \circ i_{M \otimes N} = j$ and $\psi \circ j = i_{M \otimes N}$. Hence, $\psi \circ \varphi \circ i_{M \otimes N} = i_{M \otimes N}$. This means that the following diagram commutes.

$$\begin{array}{ccc} M \times N & \xrightarrow{i_{M \otimes N}} & M \otimes N \\ & \searrow i_{M \otimes N} & \downarrow \text{id} \quad \downarrow \psi \circ \varphi \\ & & M \otimes N \end{array}$$

By the uniqueness condition of the universal property, $\text{id} = \psi \circ \varphi$. Similarly, $\text{id} = \varphi \circ \psi$. Hence, φ is an isomorphism $M \otimes N \rightarrow T$ with $\varphi \circ i_{M \otimes N} = j$. Uniqueness of φ is guaranteed by the universal property: it is the only solution to $\varphi \circ i_{M \otimes N} = j$. \square

In particular, we have

$$\text{Bilin}_R(M \times N, L) \simeq \text{Hom}(M \otimes_R N, L)$$

given by the universal property, and the inverse is given by $h \mapsto h \circ i_{M \otimes N}$.

2.3. Zero tensors

Proposition. Let M, N be R -modules. Then

$$\sum m_i \otimes n_i = 0$$

if and only if for every R -module L and every R -bilinear map $f : M \times N \rightarrow L$, we have

$$\sum f(m_i, n_i) = 0$$

To show an element of $M \otimes N$ is nonzero, it suffices to find a single R -module L and bilinear map $M \times N \rightarrow L$ with mapping the required sum to a nonzero value.

Proof. Assume $\sum m_i \otimes n_i = 0$. f factors through the map $i_{M \otimes N}$, giving

$$\begin{array}{ccc} M \times N & \xrightarrow{i_{M \otimes N}} & M \otimes N \\ & \searrow f & \downarrow h \\ & & L \end{array}$$

So

$$\sum f(m_i, n_i) = \sum h(i_{M \otimes N}(m_i, n_i)) = h\left(\sum i_{M \otimes N}(m_i, n_i)\right) = h(0) = 0$$

In the other direction, suppose $\sum m_i \otimes n_i \neq 0$. Then, taking $f = i_{M \otimes N}$, we obtain $\sum i_{M \otimes N}(m_i, n_i) \neq 0$ as required. \square

Example. Let k be a field, and consider $k^m \otimes k^\ell$. Let k^m have basis $\{e_1, \dots, e_m\}$ and k^ℓ have basis f_1, \dots, f_ℓ . Then

$$k^m \otimes k^\ell = \text{span}_k \{v \otimes w \mid v \in k^m, w \in k^\ell\} = \text{span}_k \{e_i \otimes f_j\}$$

This is in fact a basis. Suppose $\sum_{i,j} \alpha_{i,j} e_i \otimes f_j = 0$. For each $a \leq m, b \leq \ell$, define $T_{a,b} : k^m \times k^\ell \rightarrow k$ by

$$T_{a,b}((v_i)_{i=1}^m, (w_j)_{j=1}^\ell) = v_a w_b$$

By the above proposition,

$$0 = \sum_{i,j} \alpha_{i,j} T_{a,b}(e_i, f_j) = \alpha_{a,b}$$

So $k^m \otimes k^\ell \simeq k^{m\ell}$. Note that this construction only relied on the existence of a free basis, not on k being a field.

II. Commutative Algebra

Example. Consider $\mathbb{R}^2 \otimes_{\mathbb{R}} \mathbb{R}^2$. There are infinitely many pure tensors, but there is a basis consisting of the four pure vectors

$$e_1 \otimes f_1; \quad e_1 \otimes f_2; \quad e_2 \otimes f_1; \quad e_2 \otimes f_2$$

A pure tensor in $\mathbb{R}^2 \otimes_{\mathbb{R}} \mathbb{R}^2$ is of the form

$$(\alpha e_1 + \beta e_2) \otimes (\gamma f_1 + \delta f_2)$$

which expands to

$$(\alpha\gamma)(e_1 \otimes f_1) + (\alpha\delta)(e_1 \otimes f_2) + (\beta\gamma)(e_2 \otimes f_1) + (\beta\delta)(e_2 \otimes f_2)$$

Note that there is a linear dependence relation between the coefficients $\alpha\gamma, \alpha\delta, \beta\gamma, \beta\delta$, so in some sense ‘most’ tensors are not pure. For example,

$$1(e_1 \otimes f_1) + 2(e_1 \otimes f_2) + 3(e_2 \otimes f_1) + 4(e_2 \otimes f_2)$$

is not pure.

Example. Consider $\mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Z}/2\mathbb{Z}$. In this module,

$$2 \otimes (1 + 2\mathbb{Z}) = 1 \otimes (2 + 2\mathbb{Z}) = 1 \otimes 0 = 0$$

Note that \mathbb{Z} has a \mathbb{Z} -submodule $2\mathbb{Z}$. In $2\mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Z}/2\mathbb{Z}$, the element also denoted with $2 \otimes (1 + 2\mathbb{Z})$ is nonzero. For example, we can define a bilinear map to $\mathbb{Z}/2\mathbb{Z}$ given by

$$b(2n, x + 2\mathbb{Z}) = nx + 2\mathbb{Z}$$

Then $b(2, 1 + 2\mathbb{Z}) = 1 \neq 0$. So it is not the case that tensor products of submodules are submodules of tensor products.

However, if $M' \subseteq M$ and $N' \subseteq N$ and $\sum m_i \otimes n_i = 0$ in $M' \otimes N'$, then $\sum m_i \otimes n_i = 0$ in $M \otimes N$.

Proposition. If $\sum m_i \otimes n_i = 0$ in $M \otimes_R N$, then there are finitely generated R -submodules $M' \subseteq M$ and $N' \subseteq N$ such that the expression $\sum m_i \otimes n_i$ also evaluates to zero in $M' \otimes_R N'$.

This is the last proof that will use the direct construction of the tensor product instead of the universal property directly.

Proof. We know that $\sum m_i \otimes n_i = 0$ in $M \otimes_R N = R^{\oplus(M \times N)} / K$, so in particular $\sum e_{(m_i, n_i)} \in K$, where e_x maps $x \in M \times N$ to its basis element in $R^{\oplus(M \times N)}$. So this is a finite sum of $\alpha_i k_i$ with $\alpha_i \in R, k_i \in K$, and so we can take the m'_1, \dots, m'_a that appear on the left-hand sides of the k_i as the generators for M' , and similarly for N' . \square

Corollary. Let A, B be torsion-free abelian groups. Then $A \otimes_{\mathbb{Z}} B$ is torsion-free.

2. Tensor products

Proof. Suppose $n(\sum a_i \otimes b_i) = 0$ with $n \geq 1$. By the previous proposition, there are finitely generated subgroups $A' \leq A$ and $B' \leq B$ such that $n(\sum a_i \otimes b_i) = 0$ in $A' \otimes_{\mathbb{Z}} B'$. But as A' and B' are finitely generated abelian groups, the structure theorem shows that $A' = \mathbb{Z}^m$ and $B' = \mathbb{Z}^\ell$, showing that $A' \otimes_{\mathbb{Z}} B' \simeq \mathbb{Z}^{m\ell}$ is torsion-free. Thus $\sum a_i \otimes b_i = 0$ in $A' \otimes_{\mathbb{Z}} B'$, so also $\sum a_i \otimes b_i = 0$ in $A \otimes_{\mathbb{Z}} B$. \square

Example.

$$\mathbb{C}^2 \otimes_{\mathbb{C}} \mathbb{C}^3 \simeq \mathbb{C}^6 \simeq \mathbb{R}^{12}$$

However,

$$\mathbb{C}^2 \otimes_{\mathbb{R}} \mathbb{C}^3 \simeq \mathbb{R}^4 \otimes_{\mathbb{R}} \mathbb{R}^6 \simeq \mathbb{R}^{24}$$

This is to be expected: tensoring over a larger ring introduces more relations, so the amount of distinguishable elements should shrink.

2.4. Monoidal structure

We will prove a number of elementary propositions in detail to show how tensor products are used in practice.

Proposition (commutativity). There is an isomorphism $M \otimes N \simeq N \otimes M$ mapping a pure tensor $m \otimes n$ to $n \otimes m$.

Proof. Define $f : M \times N \rightarrow N \otimes M$ by $f(m, n) = n \otimes m$; this is bilinear. The universal property yields

$$\begin{array}{ccc} M \times N & \xrightarrow{i_{M \otimes N}} & M \otimes N \\ & \searrow f & \downarrow h \\ & & N \otimes M \end{array}$$

such that $h(m \otimes n) = n \otimes m$. Similarly, we obtain $h' : N \otimes M \rightarrow M \otimes N$ with $h'(n \otimes m) = m \otimes n$. Hence, the following diagram commutes.

$$\begin{array}{ccc} M \times N & \xrightarrow{i_{M \otimes N}} & M \otimes N \\ & \searrow i_{M \otimes N} & \downarrow \text{id} \downarrow h' \circ h \\ & & M \otimes N \end{array}$$

So by the uniqueness condition in the universal property, $h' \circ h$ is the identity. Similarly, $h \circ h'$ is the identity, thus h is an isomorphism. \square

Proposition (associativity). There is an isomorphism $(M \otimes N) \otimes P \simeq M \otimes (N \otimes P)$ mapping $(m \otimes n) \otimes p$ to $m \otimes (n \otimes p)$.

II. Commutative Algebra

Proof. For each $p \in P$, define the bilinear map $f_p : M \times N \rightarrow M \otimes (N \otimes P)$ by

$$f_p(m, n) = m \otimes (n \otimes p)$$

Thus, each f_p factors through $h_p : M \otimes N \rightarrow M \otimes (N \otimes P)$. Then, define the bilinear map $f : (M \otimes N) \times P \rightarrow M \otimes (N \otimes P)$ by

$$f(x, p) = h_p(x)$$

We show this is bilinear in p . Note that

$$\begin{aligned} h_{p_1+p_2}(m \otimes n) &= f_{p_1+p_2}(m, n) \\ &= m \otimes (n \otimes (p_1 + p_2)) \\ &= m \otimes (n \otimes p_1) + m \otimes (n \otimes p_2) \\ &= f_{p_1}(m, n) + f_{p_2}(m, n) \\ &= h_{p_1}(m \otimes n) + h_{p_2}(m \otimes n) \end{aligned}$$

So $h_{p_1+p_2}$ coincides with $h_{p_1} + h_{p_2}$ on the pure tensors, so by the universal property they coincide everywhere. Similarly,

$$\begin{aligned} h_{rp}(m \otimes n) &= f_{rp}(m, n) \\ &= m \otimes (n \otimes rp) \\ &= r(m \otimes (n \otimes p)) \\ &= rf_p(m, n) \\ &= rh_p(m \otimes n) \end{aligned}$$

so $h_{rp} = rh_p$. Then, by the universal property, f factors through $h : (M \otimes N) \otimes P \rightarrow M \otimes (N \otimes P)$, so

$$h((m \otimes n) \otimes p) = m \otimes (n \otimes p)$$

We can similarly construct $h' : M \otimes (N \otimes P) \rightarrow (M \otimes N) \otimes P$ with

$$h'(m \otimes (n \otimes p)) = (m \otimes n) \otimes p$$

Since $h \circ h'$ and $h' \circ h$ are the identity on pure vectors, they are the identity everywhere, and hence are inverse isomorphisms. \square

Proposition (identity). There is an isomorphism $R \otimes M \simeq M$ mapping $r \otimes m$ to rm .

Proof. The map $f : R \times M \rightarrow M$ given by $f(r, m) = rm$ factors through some $h : R \otimes M \rightarrow M$.

$$\begin{array}{ccc} R \times M & \xrightarrow{i_{R \otimes M}} & R \otimes M \\ & \searrow f & \downarrow h \\ & & M \end{array}$$

2. Tensor products

Now define the R -module homomorphism $h' : M \rightarrow R \otimes M$ by $h'(m) = 1 \otimes m = i_{R \otimes M}(1, m)$. Then

$$(h \circ h')(m) = h(i_{R \otimes M}(1, m)) = f(1, m) = m$$

giving $h \circ h' = \text{id}$. Further,

$$(h' \circ h)(r \otimes m) = 1 \otimes h(r \otimes m) = 1 \otimes f(r, m) = 1 \otimes rm = r \otimes m$$

So by the uniqueness condition in the universal property, $h' \circ h$ is the identity, and hence h is an isomorphism. \square

These operations, together with coherence conditions, make the category of R -modules into a *braided monoidal category*, where the monoid operation is \otimes and the unit is R .

Proposition (distributivity). There is an isomorphism $(\bigoplus_i M_i) \otimes P \simeq \bigoplus_i (M_i \otimes P)$ mapping $(m_i)_i \otimes p$ to $(m_i \otimes p)_i$.

Proof. Define f by

$$f((m_i)_i, p) = (m_i \otimes p)_i$$

Then there is a unique h such that the following diagram commutes.

$$\begin{array}{ccc} (\bigoplus_i M_i) \times P & \xrightarrow{i_{(\bigoplus_i M_i) \otimes P}} & (\bigoplus_i M_i) \otimes P \\ & \searrow f & \downarrow h \\ & & \bigoplus_i (M_i \otimes P) \end{array}$$

For each i , define the map $f'_i : M_i \times P \rightarrow (\bigoplus_i M_i) \otimes P$ by

$$f'_i(m_i, p) = m_i \otimes p$$

By the universal property of the tensor product, this factors through a unique h'_i .

$$\begin{array}{ccc} M_i \times P & \xrightarrow{i_{M_i \otimes P}} & M_i \otimes P \\ & \searrow f'_i & \downarrow h'_i \\ & & (\bigoplus_i M_i) \otimes P \end{array}$$

Then, by the universal property of the direct sum, the h'_i can be combined into a single h' , so this diagram commutes for each i .

$$\begin{array}{ccc} M_i \otimes P & \longrightarrow & \bigoplus_i (M_i \otimes P) \\ & \searrow h'_i & \downarrow h' \\ & & (\bigoplus_i M_i) \otimes P \end{array}$$

II. Commutative Algebra

It remains to show that h and h' are inverses. To show $h \circ h' = \text{id}_{\bigoplus_i (M_i \otimes P)}$, it suffices by the universal property of the direct sum to show that $(h \circ h')(x) = x$ for all $x \in M_i \otimes P$, for each i . Then, by the universal property of the tensor product, it further suffices to show this result only for pure tensors.

$$\begin{aligned} (h \circ h')(m_i \otimes p) &= h(h'(m_i \otimes p)) \\ &= h(h'_i(m_i \otimes p)) \\ &= h(f'_i(m_i, p)) \\ &= h(m_i \otimes p) \\ &= f(m_i, p) \\ &= m_i \otimes p \end{aligned}$$

To show $h' \circ h = \text{id}_{(\bigoplus_i M_i) \otimes P}$, it suffices by the universal property of the tensor product to show that $(h' \circ h)((m_i)_i \otimes p) = (m_i)_i \otimes p$. By linearity of h and h' , we can reduce to the case where $(m_i)_i$ has a single non-zero element m_i .

$$\begin{aligned} (h' \circ h)(m_i \otimes p) &= h'(h(m_i \otimes p)) \\ &= h'(f(m_i, p)) \\ &= h'(m_i \otimes p) \\ &= h'_i(m_i \otimes p) \\ &= f'_i(m_i \otimes p) \\ &= f'_i(m_i, p) \\ &= m_i \otimes p \end{aligned}$$

□

Example.

$$R^m \otimes_R R^\ell = \left(\bigoplus_{i=1}^m R \right) \otimes_R \left(\bigoplus_{j=1}^\ell R \right) \simeq \bigoplus_{i=1}^m \bigoplus_{j=1}^\ell (R \otimes R) \simeq \bigoplus_{i=1}^m \bigoplus_{j=1}^\ell R \simeq R^{m\ell}$$

Proposition (quotients). Let $M' \subseteq M$ and $N' \subseteq N$ be R -modules. Then there is an isomorphism

$$M/M' \otimes N/N' \simeq (M \otimes N)/L$$

where L is the submodule of $M \otimes N$ generated by

$$\{m' \otimes n \mid (m', n) \in M' \times N\} \cup \{m \otimes n' \mid (m, n') \in M \times N'\}$$

and mapping

$$(m + M') \otimes (n + N') \mapsto m \otimes n + L$$

2. Tensor products

Proof. Define

$$f : M/M' \times N/N' \rightarrow (M \otimes N)/L$$

by

$$f(m + M', n + N') = m \otimes n + L$$

This is well-defined: if $m \in M'$ or $n \in N'$, then $m \otimes n \in L$. By the universal property of the tensor product, f factors through some h .

$$\begin{array}{ccc} M/M' \times N/N' & \xrightarrow{i_{M/M' \otimes N/N'}} & M/M' \otimes N/N' \\ & \searrow f & \downarrow h \\ & & (M \otimes N)/L \end{array}$$

Now define

$$f' : M \times N \rightarrow M/M' \otimes N/N'$$

by

$$f'(m, n) = (m + M') \otimes (n + N')$$

This is clearly bilinear. Thus, we have

$$\begin{array}{ccc} M \times N & \xrightarrow{i_{M \otimes N}} & M \otimes N \\ & \searrow f' & \downarrow h' \\ & & M/M' \otimes N/N' \end{array}$$

We show that if $x \in L$, then $h'(x) = 0$. By linearity it suffices to show this for the generators.

$$h'(m' \otimes n) = f'(m', n) = 0 \otimes (n + N') = 0; \quad h'(m \otimes n') = f'(m, n') = (m + M') \otimes 0 = 0$$

Thus h' factors through the quotient.

$$\begin{array}{ccc} M \otimes N & \xrightarrow{\pi} & (M \otimes N)/L \\ & \searrow h' & \downarrow h'' \\ & & M/M' \otimes N/N' \end{array}$$

We show h and h'' are inverses. To show $h \circ h'' = \text{id}_{(M \otimes N)/L}$, it suffices by the universal properties of the quotient and the tensor product to consider the images of pure tensors

II. Commutative Algebra

under the quotient map π .

$$\begin{aligned}
 (h \circ h'')(m \otimes n + L) &= h(h''(\pi(m \otimes n))) \\
 &= h(h'(m \otimes n)) \\
 &= h(f'(m, n)) \\
 &= h((m + M') \otimes (n + N')) \\
 &= f(m + M', n + N') \\
 &= m \otimes n + L
 \end{aligned}$$

To show $h'' \circ h = \text{id}_{M/M' \otimes N/N'}$, it suffices to show the result for expressions of the form $(m + M') \otimes (n + N')$.

$$\begin{aligned}
 (h'' \circ h)((m + M') \otimes (n + N')) &= h''(h((m + M') \otimes (n + N'))) \\
 &= h''(f(m + M', n + N')) \\
 &= h''(m \otimes n + L) \\
 &= h'(m \otimes n) \\
 &= f'(m + M', n + N') \\
 &= (m + M') \otimes (n + N')
 \end{aligned}$$

□

2.5. Tensor products of maps

Proposition. Let $f : M \rightarrow M'$ and $g : N \rightarrow N'$ be R -module homomorphisms. There is a unique R -module homomorphism $f \otimes g : M \otimes N \rightarrow M' \otimes N'$ such that

$$(f \otimes g)(m \otimes n) = f(m) \otimes g(n)$$

Proof. We apply the universal property to the map $T : M \times N \rightarrow M' \otimes N'$ given by

$$T(m, n) = f(m) \otimes g(n)$$

which can be checked to be R -bilinear. □

Example. We can show

$$(f \otimes g) \circ (h \otimes i) = (f \circ h) \otimes (g \circ i)$$

For example, if $T : k^a \rightarrow k^b$ and $S : k^c \rightarrow k^d$,

$$T \otimes S : k^a \otimes_k k^c \rightarrow k^b \otimes_k k^d$$

is given by

$$(T \otimes S)(e_i \otimes e_j) = (Te_i) \otimes (Se_j) = \sum_{\ell, t} [T]_{\ell i} [S]_{t j} (f_\ell \otimes f_t)$$

where $[T]$ denotes T in the standard basis. Ordering the basis elements of $k^a \otimes k^c$ as

$$e_1 \otimes e_1, \dots, e_1 \otimes e_c, e_2 \otimes e_1, \dots, e_a \otimes e_c$$

and similarly for $k^b \otimes k^d$,

$$[T \otimes S] = \begin{pmatrix} [T]_{11} \cdot [S] & \cdots & [T]_{1a} \cdot [S] \\ \vdots & \ddots & \vdots \\ [T]_{b1} \cdot [S] & \cdots & [T]_{ba} \cdot [S] \end{pmatrix}$$

This is known as the *Kronecker product* of matrices.

Proposition. Let $f : M \rightarrow M', g : N \rightarrow N'$ be R -module homomorphisms. Then,

- (i) if f, g are isomorphisms, then so is $f \otimes g$;
- (ii) if f, g are surjective, then so is $f \otimes g$.

Proof. Part (i). $f^{-1} \otimes g^{-1}$ is a two-sided inverse for $f \otimes g$, as

$$(f^{-1} \otimes g^{-1}) \circ (f \otimes g) = (f^{-1} \circ f) \otimes (g^{-1} \circ g) = \text{id}$$

and similarly for the other side.

Part (ii). The image of $f \otimes g$ contains all pure tensors of $M' \otimes N'$, so it must be surjective. \square

The analogous result for injectivity does not hold in the general case. Consider $f : \mathbb{Z} \rightarrow \mathbb{Z}$ given by multiplication by p , and $g : \mathbb{Z}/p\mathbb{Z} \rightarrow \mathbb{Z}/p\mathbb{Z}$ given by the identity. Here,

$$(f \otimes g)(a \otimes b) = (pa) \otimes b = a \otimes (pb) = a \otimes 0 = 0$$

So $f \otimes g$ is the zero map, but $\mathbb{Z} \otimes \mathbb{Z}/p\mathbb{Z} \simeq \mathbb{Z}/p\mathbb{Z}$ is not the zero ring.

2.6. Tensor products of algebras

Let B, C be R -algebras. The usual tensor product of modules $B \otimes_R C$ can be made into a ring and then an R -algebra. This allows us to define the tensor product of algebras in a natural way. We want the ring structure to satisfy

$$(b \otimes c)(b' \otimes c') = (bb') \otimes (cc')$$

This extends to a well-defined map on all of $B \otimes C$. Indeed, for a fixed $(b, c) \in B \times C$, there is an R -bilinear map $B \times C \rightarrow B \otimes C$ given by

$$(b', c') \mapsto (bb') \otimes (cc')$$

so we can use the universal property to extend this to a map $B \otimes C \rightarrow B \otimes C$ that acts on pure tensors in the obvious way. One can show that the ring axioms are satisfied. To define the R -algebra structure, we define the ring homomorphism $R \rightarrow B \otimes C$ by

$$r \mapsto (r \cdot 1_B) \otimes 1_C = 1_B \otimes (r \cdot 1_C)$$

II. Commutative Algebra

Example. There is an isomorphism of R -algebras

$$\varphi : R[X_1, \dots, X_n] \otimes_R R[T_1, \dots, T_r] \simeq R[X_1, \dots, X_n, T_1, \dots, T_r]$$

An R -basis for the left-hand side as an R -module is given by elements of the form $a \otimes b$ where a and b are monomials. The right hand side has a basis of elements of the form ab , where $a \in R[X_1, \dots, X_n]$ and $b \in R[T_1, \dots, T_r]$ are monomials as above. Mapping $\varphi(a \otimes b) = ab$, we obtain an R -module isomorphism. To check this is an R -algebra isomorphism, we verify multiplication and its action on scalars.

$$\varphi(r \otimes 1) = r \cdot 1; \quad \varphi(1 \otimes 1)$$

and for monomials p_i, q_i, h_i, g_i ,

$$\begin{aligned} \varphi\left(\left(\sum_i p_i \otimes q_i\right)\left(\sum_j h_j \otimes g_j\right)\right) &= \sum_{i,j} (p_i h_j)(q_i g_j) \\ &= \sum_{i,j} (p_i q_i)(h_j g_j) \\ &= \sum_{i,j} \varphi(p_i \otimes q_i) \varphi(h_j \otimes g_j) \\ &= \left(\sum_i \varphi(p_i \otimes q_i)\right) \left(\sum_j \varphi(h_j \otimes g_j)\right) \\ &= \varphi\left(\sum_i p_i \otimes q_i\right) \varphi\left(\sum_j h_j \otimes g_j\right) \end{aligned}$$

More generally,

$$R[X_1, \dots, X_n]_{/I} \otimes R[T_1, \dots, T_r]_{/J} \simeq R[X_1, \dots, X_n] \otimes R[T_1, \dots, T_r]_{/L} \simeq R[X_1, \dots, X_n, T_1, \dots, T_r]_{/I^e + J^e}$$

where L is constructed as above when quotients were discussed, and I^e is the extension of I in the larger ring $R[X_1, \dots, X_n, T_1, \dots, T_r]$. For example,

$$\mathbb{C}[X, Y, Z]_{/(f,g)} \otimes_{\mathbb{C}} \mathbb{C}[W, U]_{/(h)} \simeq \mathbb{C}[X, Y, Z, W, U]_{/(f,g,h)}$$


Proposition (universal property of tensor product of algebras). Let A, B be R -algebras. For every algebra C and R -algebra homomorphisms $f_1 : A \rightarrow C$ and $f_2 : B \rightarrow C$, there is a unique R -algebra homomorphism $h : A \otimes_R B \rightarrow C$ such that the following diagram commutes:

$$\begin{array}{ccc} A & & B \\ & \searrow i_A \quad \swarrow i_B & \\ & A \otimes B & \\ & \downarrow h & \\ & C & \end{array} \quad \begin{array}{c} f_1 \nearrow \\ f_2 \nearrow \end{array}$$

where $i_A(a) = a \otimes 1$ and $i_B(b) = 1 \otimes b$. Furthermore, this characterises the triple $(A \otimes_R B, i_A, i_B)$ uniquely up to unique isomorphism.

Proof. $A \otimes_R B$ is generated as an R -algebra by $\{a \otimes 1 \mid a \in A\} \cup \{1 \otimes b \mid b \in B\}$. This implies the uniqueness of h . For existence, we can define an R -bilinear map $A \times B \rightarrow C$ by $(a, b) \mapsto f_1(a)f_2(b)$, then apply the universal property of the tensor product of modules. This produces an R -linear map $h : A \otimes B \rightarrow C$. It remains to show that this is a homomorphism of algebras. \square

Example.



An algebra homomorphism from a polynomial ring is defined uniquely by giving its action on its variables, thus

$$R[X_1, \dots, X_n] \otimes R[T_1, \dots, T_r] \simeq R[X_1, \dots, X_n, T_1, \dots, T_r]$$

as was shown above.

Remark. (i) If $f : A \rightarrow A', g : B \rightarrow B'$ are R -algebra homomorphisms, then $f \otimes g : A \otimes B \rightarrow A' \otimes B'$ is not only an R -module homomorphism but is also an R -algebra homomorphism.

(ii) There are R -algebra homomorphisms

- (a) $R/I \otimes R/J \simeq R/I + J$;
- (b) $A \otimes B \simeq B \otimes A$;
- (c) $A \otimes (B \times C) \simeq (A \otimes B) \times (A \otimes C)$;
- (d) $A \otimes B^n \simeq (A \otimes B)^n$;
- (e) $(A \otimes B) \otimes C \simeq A \otimes (B \otimes C)$.

2.7. Restriction and extension of scalars

Let $f : R \rightarrow S$ be a ring homomorphism. Let M be an S -module. Then we can *restrict scalars* to make M into an R -module by

$$r \cdot m = f(r) \cdot m$$

II. Commutative Algebra

The composition $R \rightarrow S \rightarrow \text{End } M$ is a ring homomorphism, so this makes M into an R -module automatically without needing to check axioms.

Example. Let $f : \mathbb{R} \rightarrow \mathbb{C}$ be the inclusion. Then any \mathbb{C} -module is an \mathbb{R} -module.

Now suppose $f : R \rightarrow S$ is a ring homomorphism, M is an S -module, and N is an R -module. We can form the R -module $M \otimes_R N$, as M is an R -module by restriction of scalars. *Extension of scalars* shows that $M \otimes_R N$ is also an S -module. The action of $s \in S$ on pure tensors is

$$s \cdot (m \otimes n) = sm \otimes n$$

We have an R -bilinear map $M \times N \rightarrow M \otimes_R N$ by

$$(m, n) \mapsto sm \otimes n$$

so by the universal property this gives rise to a map $h_s : M \otimes_R N \rightarrow M \otimes_R N$ with the desired action on pure tensors. h_s is R -linear by the universal property. Defining $\varphi : S \rightarrow \text{End}(M \otimes_R N)$ by $\varphi(s) = h_s$, one can check that h_s is a well-defined endomorphism and that φ is a ring homomorphism.

Example. $S \otimes_R R \simeq S$ as R -modules, by $s \otimes r \mapsto s \cdot f(r)$. This is also S -linear, since

$$s'(s \otimes r) = (s's \otimes r) \mapsto s's \cdot f(r) = s'(s \cdot f(r))$$

For example, $\mathbb{C} \otimes_{\mathbb{R}} \mathbb{R} \simeq \mathbb{C}$ as \mathbb{C} -modules.

Example. Let M be an S -module and $(N_i)_{i \in I}$ are R -modules. Then

$$M \otimes \left(\bigoplus_i N_i \right) \simeq \bigoplus_i (M \otimes N_i)$$

as S -modules. So $\mathbb{C} \otimes_{\mathbb{R}} \mathbb{R}^n \simeq \mathbb{C}^n$ as \mathbb{C} -modules.

Example. Restrict the \mathbb{C} -module \mathbb{C}^n to an \mathbb{R} -module to obtain \mathbb{R}^{2n} . Then, extending to \mathbb{C} ,

$$\mathbb{C} \otimes_{\mathbb{R}} \mathbb{R}^{2n} \simeq \mathbb{C}^{2n}$$

Similarly, extending \mathbb{R}^n to \mathbb{C} , we find $\mathbb{C} \otimes_{\mathbb{R}} \mathbb{R}^n \simeq \mathbb{C}^n$ over \mathbb{C} . Restricting to \mathbb{R} , $\mathbb{C}^n \simeq \mathbb{R}^{2n}$. So the operations of restriction and extension of scalars are not inverses in either direction.

Example. Consider \mathbb{Z}^n as a \mathbb{Z} -module. Consider the quotient map $f : \mathbb{Z} \rightarrow \mathbb{Z}/2\mathbb{Z}$. Extending scalars to $\mathbb{Z}/2\mathbb{Z}$,

$$\mathbb{Z}/2\mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Z}^n \simeq (\mathbb{Z}/2\mathbb{Z})^n$$

Example. Consider $\mathbb{C}^n \otimes_{\mathbb{R}} \mathbb{R}^\ell$ as a \mathbb{C} -module. As \mathbb{R} -modules,

$$\mathbb{C}^n \otimes_{\mathbb{R}} \mathbb{R}^\ell \simeq \mathbb{R}^{2n} \otimes_{\mathbb{R}} \mathbb{R}^\ell \simeq \mathbb{R}^{2n\ell} \simeq \mathbb{C}^{n\ell}$$

We would like to make this into an isomorphism of \mathbb{C} -modules. We will show that in fact

$$\mathbb{C}^n \otimes_{\mathbb{R}} \mathbb{R}^\ell \simeq \mathbb{C}^n \otimes_{\mathbb{C}} (\mathbb{C} \otimes_{\mathbb{R}} \mathbb{R}^\ell)$$

where

$$v \otimes u \mapsto v \otimes (1 \otimes u)$$

giving

$$\mathbb{C}^n \otimes_{\mathbb{R}} \mathbb{R}^\ell \simeq \mathbb{C}^n \otimes_{\mathbb{C}} \mathbb{C}^\ell \simeq \mathbb{C}^{n\ell}$$

as \mathbb{C} -modules. The isomorphism

$$\mathbb{C}^n \otimes_{\mathbb{R}} \mathbb{R}^\ell \simeq \mathbb{C}^n \otimes_{\mathbb{C}} \mathbb{C}^\ell$$

maps a pure tensor $v \otimes u$ to $v \otimes u$.

Proposition. Let M be an S -module and N be an R -module. Then

$$M \otimes_R N \simeq M \otimes_S (S \otimes_R N)$$

as S -modules, where

$$m \otimes n \mapsto m \otimes (1 \otimes n); \quad sm \otimes n \mapsto m \otimes (s \otimes n)$$

Proof. The map $(m, n) \mapsto m \otimes (1 \otimes n)$ is R -bilinear, so the map f mapping $m \otimes n$ to $m \otimes (1 \otimes n)$ is well-defined as a map of R -modules. We show it is S -linear on pure tensors.

$$f(s(m \otimes n)) = f(sm \otimes n) = sm \otimes (1 \otimes n) = s(m \otimes (1 \otimes n)) = sf(m \otimes n)$$

For a fixed $m \in M$, the map $s \otimes n \mapsto sm \otimes n$ is well-defined and S -linear. This collection of maps is S -linear in its parameter m , so we obtain an S -bilinear map $(m, s \otimes n) \mapsto sm \otimes n$. Hence, we obtain a map g mapping $m \otimes (s \otimes n)$ to $sm \otimes n$, as desired. One can easily check that f and g are inverses on pure tensors. \square

Proposition. Let M, M' be S -modules and N, N' be R -modules. Then we have S -module isomorphisms

$$\begin{aligned} M \otimes_R N &\simeq N \otimes_R M \\ (M \otimes_R N) \otimes_R N' &\simeq M \otimes_R (N \otimes_R N') \\ (M \otimes_R N) \otimes_S M' &\simeq M \otimes_S (N \otimes_R M') \\ M \otimes_R \left(\bigoplus_i N_i \right) &\simeq \bigoplus_i (M \otimes_R N_i) \end{aligned}$$

Heuristically, the tensor products in the above isomorphisms always operate over the largest possible ring: S if both operands are S -modules, else R . We prove only the third result.

II. Commutative Algebra

Proof. By the previous proposition,

$$\begin{aligned} (M \otimes_R N) \otimes_S M' &\simeq (M \otimes_S (N \otimes_R S)) \otimes_S M' \\ &\simeq M \otimes_S ((N \otimes_R S) \otimes_S M') \\ &\simeq M \otimes_S (N \otimes_R M') \end{aligned}$$

□

Corollary. Let N, N' be R -modules. Then

$$S \otimes_R (N \otimes_R N') \simeq (S \otimes_R N) \otimes_S (S \otimes_R N')$$

as S -modules.

Proof.

$$S \otimes_R (N \otimes_R N') \simeq (S \otimes_R N) \otimes_R N' \simeq (S \otimes_R N) \otimes_S (S \otimes_R N')$$

□

Example.

$$\mathbb{C} \otimes_{\mathbb{R}} (\mathbb{R}^\ell \otimes_{\mathbb{R}} \mathbb{R}^k) \simeq (\mathbb{C} \otimes_{\mathbb{R}} \mathbb{R}^\ell) \otimes_{\mathbb{C}} (\mathbb{C} \otimes_{\mathbb{R}} \mathbb{R}^k) \simeq \mathbb{C}^\ell \otimes_{\mathbb{C}} \mathbb{C}^k \simeq \mathbb{C}^{\ell k}$$

By induction, one can see that

$$S \otimes_R (N_1 \otimes_R \cdots \otimes_R N_\ell) = (S \otimes_R N_1) \otimes_S \cdots \otimes_S (S \otimes_R N_\ell)$$

2.8. Extension of scalars on morphisms

Let $f : N \rightarrow N'$ be an R -linear map, and M be an S -module. Then the map

$$\text{id}_M \otimes f : M \otimes_R N \rightarrow M \otimes_R N'$$

is S -linear. Indeed,

$$(\text{id}_M \otimes f)(s(m \otimes n)) = \text{id}_M sm \otimes f(n) = s(m \otimes f(n)) = s((\text{id}_M \otimes f)(m \otimes n))$$

Example. Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^\ell$ be R -linear, and use bases e_1, \dots, e_n and f_1, \dots, f_ℓ . Then

$$\text{id}_{\mathbb{C}} \otimes T : \mathbb{C} \otimes_{\mathbb{R}} \mathbb{R}^n \rightarrow \mathbb{C} \otimes_{\mathbb{R}} \mathbb{R}^\ell$$

is given by

$$(\text{id}_{\mathbb{C}} \otimes T)(1 \otimes e_i) = 1 \otimes T(e_i) = 1 \otimes \sum_{j=1}^{\ell} [T]_{ji} \cdot f_j = \sum_{j=1}^{\ell} [T]_{ji} (1 \otimes f_j)$$

This shows that the matrix $[\text{id}_{\mathbb{C}} \otimes T]$ has all real elements, and is the same as the matrix $[T]$.

2.9. Extension of scalars in algebras

Let A, B be R -algebras. Then the module $A \otimes_R B$ is also an R -algebra. Furthermore, can see that $A \otimes_R B$ is an A -algebra and a B -algebra by the maps $a \mapsto a \otimes 1$ and $b \mapsto 1 \otimes b$.

Example. Consider $R[X_1, \dots, X_n]$ and $f : R \rightarrow S$. Then

$$\varphi : S \otimes_R R[X_1, \dots, X_n] \xrightarrow{\sim} S[X_1, \dots, X_n]$$

as S -algebras. Indeed, φ already exists as an isomorphism of S -modules given by

$$\varphi(s \otimes p) = sp$$

and one can verify that unity and multiplication are preserved. Further,

$$S \otimes \left(R[X_1, \dots, X_n] / I \right) \simeq S[X_1, \dots, X_n] / I^e$$

Proposition. Let A be an R -algebra and B be an S -algebra. Then

$$A \otimes_R B \simeq (A \otimes_R S) \otimes_S B$$

as S -algebras.

Proposition. Let A, B be R -algebras. Then

$$S \otimes_R (A \otimes_R B) \simeq (S \otimes_R A) \otimes_S (S \otimes_R B)$$

as S -algebras.

The proofs are omitted, but trivial.

2.10. Exactness properties of the tensor product

Let M be an R -module. There is a functor

$$T_M : \mathbf{Mod}_R \rightarrow \mathbf{Mod}_R$$

from the category of R -modules to itself given by

$$T_M(N) = M \otimes_R N; \quad T_M(N \xrightarrow{f} N') = \text{id}_M \otimes f$$

We intend to show that if

$$A \xrightarrow{f} B \xrightarrow{g} C \longrightarrow 0$$

is an exact sequence of R -modules, then

$$M \otimes_R A \xrightarrow{T_M(f)} M \otimes_R B \xrightarrow{T_M(g)} M \otimes_R C \longrightarrow 0$$

is also an exact sequence. This shows that T_M is a *right exact* functor.

II. Commutative Algebra

Definition. Let Q, P be R -modules. Then

$$\mathrm{Hom}_R(Q, P) = \{f : Q \rightarrow P \mid f \text{ is } R\text{-linear}\}$$

This is also an R -module: if $\varphi \in \mathrm{Hom}_R(Q, P)$,

$$(r \cdot \varphi)(q) = r \cdot \varphi(q)$$

Definition. Let Q, P be R -modules. Then

$$\mathrm{Hom}_R(Q, -) : \mathbf{Mod}_R \rightarrow \mathbf{Mod}_R$$

and

$$\mathrm{Hom}_R(-, P) : \mathbf{Mod}_R^{\mathrm{op}} \rightarrow \mathbf{Mod}_R$$

are functors, with action on morphisms $f : N' \rightarrow N$ given by

$$\mathrm{Hom}_R(Q, f)(\varphi) = f \circ \varphi = f_*(\varphi) : \mathrm{Hom}_R(Q, N') \rightarrow \mathrm{Hom}_R(Q, N')$$

and

$$\mathrm{Hom}_R(f, P)(\varphi) = \varphi \circ f = f^*(\varphi) : \mathrm{Hom}_R(N, Q) \rightarrow \mathrm{Hom}_R(N', Q)$$

Proposition. Suppose

$$0 \longrightarrow A \xrightarrow{f} B \xrightarrow{g} C$$

is exact. Then, so is

$$0 \longrightarrow \mathrm{Hom}_R(Q, A) \xrightarrow{f_*} \mathrm{Hom}_R(Q, B) \xrightarrow{g_*} \mathrm{Hom}_R(Q, C)$$

Thus, the covariant hom-functor is *left exact*.

Proof. First, we show f_* is injective. Suppose $f_*(\varphi) = 0$, so $f \circ \varphi = 0$. Then as f is injective, $f(\varphi(x)) = 0$ implies $\varphi(x) = 0$, giving $\varphi = 0$ as required.

Now consider $\varphi : Q \rightarrow A$. Then

$$g_*(f_*(\varphi)) = g \circ (f \circ \varphi) = (g \circ f) \circ \varphi = 0 \circ \varphi = 0$$

so $\mathrm{im} f_* \subseteq \ker g_*$. Now suppose $\varphi : Q \rightarrow B$ has $g_*(\varphi) = g \circ \varphi = 0$. So for all $x \in Q$, $g(\varphi(x)) = 0$. By exactness of the original sequence, $\varphi(x) \in \mathrm{im} f$. As f is injective, $\varphi(x)$ has a unique preimage $\psi(x)$ under f . As f is R -linear, so is $\psi : Q \rightarrow A$. Hence $f_*(\psi) = \varphi$ as required. \square

Proposition. Suppose

$$A \xrightarrow{f} B \xrightarrow{g} C \longrightarrow 0$$

is exact. Then, so is

$$0 \longrightarrow \mathrm{Hom}_R(C, P) \xrightarrow{g^*} \mathrm{Hom}_R(B, P) \xrightarrow{f^*} \mathrm{Hom}_R(A, P)$$

Thus, the contravariant hom-functor is also left-exact.

Proof. First, we show g^* is injective. Suppose $g^*(\varphi) = 0$, so $\varphi \circ g = 0$. As g is surjective, we must have $\varphi = 0$.

Now consider $\varphi : C \rightarrow P$. Then

$$f^*(g^*(\varphi)) = (\varphi \circ g) \circ f = \varphi \circ (g \circ f) = \varphi \circ 0 = 0$$

so $\text{im } g^* \subseteq \ker f^*$. Now suppose $\varphi : B \rightarrow P$ has $f^*(\varphi) = \varphi \circ f = 0$. So for all $x \in A$, $\varphi(f(x)) = 0$. Define $\psi : C \rightarrow P$ by

$$\psi(g(x)) = \varphi(x)$$

We show this is well-defined. If $g(x) = g(y)$, then $g(x - y) = 0$, so $x - y = f(a)$ for some $a \in A$. But then $\varphi(f(a)) = 0$, so $\varphi(x) = \varphi(y)$. As φ and g are R -linear, so is ψ . Hence $g^*(\psi) = \varphi$ as required. \square

Lemma. Consider a sequence of R -modules

$$A \xrightarrow{f} B \xrightarrow{g} C$$

Suppose that for each R -module P ,

$$\text{Hom}_R(C, P) \xrightarrow{g^*} \text{Hom}_R(B, P) \xrightarrow{f^*} \text{Hom}_R(A, P)$$

is exact. Then the original sequence

$$A \xrightarrow{f} B \xrightarrow{g} C$$

is exact.

Proof. First, take $P = C$. By hypothesis, the following sequence is exact.

$$\text{Hom}_R(C, C) \xrightarrow{g^*} \text{Hom}_R(B, C) \xrightarrow{f^*} \text{Hom}_R(A, C)$$

Consider

$$\text{id}_C \mapsto \text{id}_C \circ g \mapsto \text{id}_C \circ g \circ f$$

By exactness, id_C must be mapped to zero under $f^* \circ g^*$, so $g \circ f = 0$. Hence $\text{im } f \subseteq \ker g$.

Now, take $P = B/\text{im } f = \text{coker } f$.

$$\text{Hom}_R(C, B/\text{im } f) \xrightarrow{g^*} \text{Hom}_R(B, B/\text{im } f) \xrightarrow{f^*} \text{Hom}_R(A, B/\text{im } f)$$

Let $h : B \rightarrow B/\text{im } f$ be the quotient map. Then,

$$f^*(h) = h \circ f; \quad h(f(x)) = 0$$

Thus by exactness, h has a preimage $e : C \rightarrow B/\text{im } f$. Then $g^*(e) = e \circ g = h$, so $\ker g \subseteq \ker h = \text{im } f$, giving the reverse inclusion. \square

II. Commutative Algebra

By the universal property of the tensor product,

$$\mathrm{Hom}_R(M \otimes_R N, L) \simeq \mathrm{Bilin}_R(M \times N, L) \simeq \mathrm{Hom}_R(N, \mathrm{Hom}_R(M, L))$$

given by

$$\varphi \mapsto (n \mapsto m \mapsto \varphi(m \otimes n)); \quad (m \otimes n \mapsto \varphi(m)(n)) \mapsto \varphi$$

This bijection is *natural*, in the sense that many commutative diagrams involving them will commute.

Proposition. Let M be an R -module. Then the functor $T_M = M \otimes_R (-)$ is right exact.

Proof. Consider an exact sequence of R -modules

$$A \xrightarrow{f} B \xrightarrow{g} C \longrightarrow 0$$

We must show that

$$M \otimes_R A \xrightarrow{\mathrm{id}_M \otimes f} M \otimes_R B \xrightarrow{\mathrm{id}_M \otimes g} M \otimes_R C \longrightarrow 0$$

is exact. Let P be an R -module, and consider apply the functor $\mathrm{Hom}(-, P)$ to this sequence. As this is left exact, the resulting sequence will be exact.

$$0 \longrightarrow \mathrm{Hom}_R(C, P) \xrightarrow{g^*} \mathrm{Hom}_R(B, P) \xrightarrow{f^*} \mathrm{Hom}_R(A, P)$$

Then, apply the functor $\mathrm{Hom}(M, -)$, which is also left exact.

$$0 \longrightarrow \mathrm{Hom}_R(M, \mathrm{Hom}_R(C, P)) \xrightarrow{(g^*)^*} \mathrm{Hom}_R(M, \mathrm{Hom}_R(B, P)) \xrightarrow{(f^*)^*} \mathrm{Hom}_R(M, \mathrm{Hom}_R(A, P))$$

We thus obtain

$$\begin{array}{ccccccc} 0 & \longrightarrow & \mathrm{Hom}_R(M, \mathrm{Hom}_R(C, P)) & \longrightarrow & \mathrm{Hom}_R(M, \mathrm{Hom}_R(B, P)) & \longrightarrow & \mathrm{Hom}_R(M, \mathrm{Hom}_R(A, P)) \\ \downarrow \simeq & & \downarrow \simeq & & \downarrow \simeq & & \downarrow \simeq \\ 0 & \longrightarrow & \mathrm{Hom}_R(M \otimes_R C, P) & \longrightarrow & \mathrm{Hom}_R(M \otimes_R B, P) & \longrightarrow & \mathrm{Hom}_R(M \otimes_R A, P) \end{array}$$

As this diagram commutes, the bottom sequence is exact. Since this holds for all P , by the previous lemma, we can cancel P to give exact sequences

$$0 \longrightarrow M \otimes_R C \longrightarrow M \otimes_R B \longrightarrow M \otimes_R A$$

which combine into the longer sequence as required. \square

Remark. It is not the case that if

$$A \longrightarrow B \longrightarrow C$$

is exact, then

$$M \otimes_R A \longrightarrow M \otimes_R B \longrightarrow M \otimes_R C$$

is also exact; the fact that the sequence has a zero on the right is important. Consider the exact sequence

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

and tensor with $\mathbb{Z}/2\mathbb{Z}$. We would then obtain

$$\begin{array}{ccccc} 0 & \longrightarrow & \mathbb{Z}/2\mathbb{Z} \otimes \mathbb{Z} & \xrightarrow{\times 2} & \mathbb{Z}/2\mathbb{Z} \otimes \mathbb{Z} \\ \downarrow \simeq & & \downarrow \simeq & & \downarrow \simeq \\ 0 & \longrightarrow & \mathbb{Z}/2\mathbb{Z} & \xrightarrow{\times 2} & \mathbb{Z}/2\mathbb{Z} \end{array}$$

but this sequence is not exact.

2.11. Flat modules

Definition. An R -module M is *flat* if whenever $f : N \rightarrow N'$ is R -linear and injective, the map

$$\text{id}_M \otimes f : M \otimes_R N \rightarrow M \otimes_R N'$$

is injective.

Example. (i) $\mathbb{Z}/2\mathbb{Z}$ is not a flat \mathbb{Z} -module.

(ii) Free modules are flat. Suppose $f : N \rightarrow N'$ is an injective R -linear map. Then

$$\begin{array}{ccc} R^{\oplus I} \otimes_R N & \xrightarrow{\text{id}_{R^{\oplus I}} \otimes f} & R^{\oplus I} \otimes_R N' \\ \simeq \downarrow & & \downarrow \simeq \\ N^{\oplus I} & \xrightarrow{g} & (N')^{\oplus I} \end{array}$$

commutes, where

$$g((n_i)_{i \in I}) = (f(n_i))_{i \in I}$$

But g is injective, so $\text{id}_{R^{\oplus I}} \otimes f$ must also be injective.

(iii) The base ring matters. One can see that $\mathbb{Z}/2\mathbb{Z}$ is not a flat \mathbb{Z} -module but it is a flat $\mathbb{Z}/2\mathbb{Z}$ -module as it is a free $\mathbb{Z}/2\mathbb{Z}$ -module.

Definition. An R -module M is *torsion-free* if $rm \neq 0$ whenever r is not a zero divisor in R and $m \neq 0$.

Proposition. Flat modules are torsion-free.

II. Commutative Algebra

Proof. Suppose M is not torsion-free. Then there is $r_0 \in R$ not a zero divisor and $m_0 \neq 0$, such that $r_0 m_0 = 0$. Consider the R -linear map $f : R \rightarrow R$ given by multiplication by r_0 . Its kernel is zero, as r_0 is not a zero divisor. So f is injective. The following diagram commutes.

$$\begin{array}{ccc} M \otimes_R R & \xrightarrow{\text{id}_M \otimes f} & M \otimes_R R \\ \downarrow \simeq & & \downarrow \simeq \\ M & \xrightarrow{m \mapsto r_0 m} & M \end{array}$$

If M were flat, $\text{id}_M \otimes f$ would be injective, but then the map $m \mapsto r_0 m$ would also be injective, which is a contradiction. \square

Example. Let R be an integral domain, and let I be a nonzero ideal of R . Then R/I is not flat. Indeed, if $I = R$ then $R/I = 0$ is not flat. Instead, suppose $I \subsetneq R$, and let $0 \neq x \in I$. Tensoring with R/I , the map $R/I \rightarrow R/I$ given by multiplication by x is the zero map, but R/I is not the zero module, so R/I is not torsion-free.

Proposition. Let M be an R -module. Then the following are equivalent.

- (i) T_M preserves exactness of all exact sequences;
- (ii) T_M preserves exactness of short exact sequences;
- (iii) M is flat;
- (iv) if $f : N \rightarrow N'$ is R -linear and injective, and N, N' are finitely generated R -modules, then $\text{id}_M \otimes f$ is injective.

Note that a map $f : M \rightarrow N$ is injective exactly when the sequence

$$0 \longrightarrow M \xrightarrow{f} N$$


is exact, so all of these conditions relate exact sequences.

Proof. Note that (i) implies (ii) which implies (iii) which implies (iv).

(ii) *implies (i).* Suppose the sequence

$$A \xrightarrow{f} B \xrightarrow{g} C$$

is exact. Then, the following diagram is exact.



After applying $T = T_M$, the diagram still commutes, and the diagonal lines remain exact.

$$\begin{aligned} \text{im}(TA \rightarrow TB) &= \text{im}(TA \rightarrow T(\text{im } f) \rightarrow TB) \\ &= \text{im}(T(\text{im } f) \rightarrow TB) \\ &= \ker(TB \rightarrow T(\text{im } g)) \\ &= \ker(TB \rightarrow T(\text{im } g) \rightarrow TC) \\ &= \ker(TB \rightarrow TC) \end{aligned}$$

(iii) implies (ii). Suppose the sequence

$$0 \longrightarrow A \xrightarrow{f} B \xrightarrow{g} C \longrightarrow 0$$

is exact. As T_M is right exact, we obtain the exact sequence

$$M \otimes_R A \xrightarrow{\text{id}_M \otimes f} M \otimes_R B \xrightarrow{\text{id}_M \otimes g} M \otimes_R C \longrightarrow 0$$

It suffices to show that $\text{id}_M \otimes f$ is injective, but this is precisely the hypothesis of (iii).

(iv) *implies (iii).* Let $f : N \rightarrow N'$ be R -linear and injective. Let $\sum m_i \otimes n_i \in M \otimes_R N$ be such that

$$0 = (\text{id}_M \otimes f)\left(\sum m_i \otimes n_i\right) \in M \otimes N'$$

Then there are finitely generated submodules L, L' of N, N' such that the n_i are elements of L and

$$0 = (\text{id}_M \otimes f)\left(\sum m_i \otimes n_i\right) \in M \otimes L'$$

By (iv), we obtain

$$0 = \sum m_i \otimes n_i \in M \otimes L$$

But L is a submodule of N , so

$$0 = \sum m_i \otimes n_i \in M \otimes N$$

Hence $\text{id}_M \otimes f : M \otimes_R N \rightarrow M \otimes_R N'$ is injective.

☐

II. Commutative Algebra

Proposition. Let $f : R \rightarrow S$ be a ring homomorphism, and let M be a flat R -module. Then $S \otimes_R M$ is a flat S -module.

Proof. Let $g : N \rightarrow N'$ be an S -linear injective map. Then

$$\begin{array}{ccc} (S \otimes_R M) \otimes_S N & \xrightarrow{\text{id}_{S \otimes_R M} \otimes g} & (S \otimes_R M) \otimes_S N' \\ \downarrow \simeq & & \downarrow \simeq \\ M \otimes_R N & \xrightarrow{\text{id}_M \otimes g} & M \otimes_R N' \end{array}$$

commutes. The map $\text{id}_M \otimes g$ is injective as M is flat, so the map $\text{id}_{S \otimes_R M} \otimes g$ is also injective. Thus $S \otimes_R M$ is a flat S -module. \square

We now explore some further examples of tensor products.

Example. Consider $\mathbb{Q} \otimes_{\mathbb{Z}} \mathbb{Z}/n\mathbb{Z}$. In this ring,

$$x \otimes y = n \cdot \frac{x}{n} \otimes y = \frac{x}{n} \otimes ny = \frac{x}{n} \otimes 0 = 0$$

So this ring is trivial. To prove this, we used the fact that for all $x \in \mathbb{Q}$ and $n \geq 1$, there is an element $y \in \mathbb{Q}$ such that $ny = x$. We say that \mathbb{Q} is a *divisible group*. We also needed the fact that $\mathbb{Z}/n\mathbb{Z}$ is a *torsion group*: all elements are of finite order. Hence the tensor product of a divisible group with a torsion group is zero. In particular, it follows that

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

However, for an R -module $M \neq 0$, if M is finitely generated then $M \otimes_R M \neq 0$.

Example. Let V be a vector space over \mathbb{Q} . Then $\mathbb{Q} \otimes_{\mathbb{Q}} V \simeq V$ as \mathbb{Q} -modules, given by the map $x \otimes v \mapsto xv$. However, $\mathbb{Q} \otimes_{\mathbb{Z}} V$ is also isomorphic to V , given by the same map. First, note that every tensor in $\mathbb{Q} \otimes_{\mathbb{Z}} V$ is pure.

$$\sum \frac{a_i}{b_i} \otimes v_i = \sum \frac{1}{b_i} \otimes a_i v_i = \sum \frac{1}{b_i} \otimes b_i \frac{a_i}{b_i} v_i = \sum 1 \otimes \frac{a_i}{b_i} v_i = 1 \otimes \sum \frac{a_i}{b_i} v_i$$

Surjectivity of the map is clear as $1 \otimes v \rightarrow v$. We check injectivity on pure tensors. If $xv = 0$, then $x = 0$ or $v = 0$, and in any case, $x \otimes v = 0$.

Example. Consider

$$M \otimes_R \left(\bigoplus_{i \in I} N_i \right) \simeq \bigoplus_{i \in I} (M \otimes_R N_i)$$

given by $m \otimes (n_i)_{i \in I} \mapsto (m \otimes n_i)_{i \in I}$. This is not true with the direct product. However, we do have a map

$$M \otimes_R \left(\prod_{i \in I} N_i \right) \rightarrow \prod_{i \in I} (M \otimes_R N_i)$$

given by the same formula, but this is in general not an isomorphism. Consider

$$\mathbb{Q} \otimes_{\mathbb{Z}} \prod_{n=1}^{\infty} \mathbb{Z}/2^n \mathbb{Z} \rightarrow \prod_{n=1}^{\infty} (\mathbb{Q} \otimes_{\mathbb{Z}} \mathbb{Z}/2^n \mathbb{Z})$$

The right-hand side is zero, as each factor is a tensor product of a divisible group by a torsion group. However, the left-hand side is nonzero. Let

$$g = (1, 1, 1, \dots) \in \prod_{n=1}^{\infty} \mathbb{Z}/2^n \mathbb{Z}$$

This is an element of infinite order, so $\langle g \rangle \simeq \mathbb{Z}$ as a subgroup of $\prod_{n=1}^{\infty} \mathbb{Z}/2^n \mathbb{Z}$. Thus

$$\mathbb{Q} \otimes_{\mathbb{Z}} \langle g \rangle \simeq \mathbb{Q}$$

as \mathbb{Z} -modules. But we have an injective inclusion map

$$\langle g \rangle \rightarrow \prod_{n=1}^{\infty} \mathbb{Z}/2^n \mathbb{Z}$$

We will later show that \mathbb{Q} is a flat \mathbb{Z} -module. This justifies the fact that there is an inclusion

$$\mathbb{Q} \otimes_{\mathbb{Z}} \langle g \rangle \hookrightarrow \mathbb{Q} \otimes_{\mathbb{Z}} \prod_{n=1}^{\infty} \mathbb{Z}/2^n \mathbb{Z}$$

showing that in particular the module in question is nonzero.

Example. Consider $\mathbb{C} \otimes_{\mathbb{R}} \mathbb{C}$. We will choose to extend scalars on the left, treating the right-hand copy of \mathbb{C} as an \mathbb{R} -module isomorphic to \mathbb{R}^2 . As a module, $\mathbb{C} \otimes_{\mathbb{R}} \mathbb{C} \simeq \mathbb{C} \otimes_{\mathbb{R}} \mathbb{R}^2$ is isomorphic to \mathbb{C}^2 . The basis for \mathbb{C}^2 is given by $1 \otimes 1, 1 \otimes i$.

As a \mathbb{C} -algebra, we again choose to extend scalars on the left, considering the right-hand copy of \mathbb{C} as an \mathbb{R} -algebra.

$$\begin{aligned} \mathbb{C} \otimes_{\mathbb{R}} \mathbb{C} &\simeq \mathbb{C} \otimes_{\mathbb{R}} \mathbb{R}[T]/(T^2 + 1) \\ &\simeq \mathbb{C}[T]/(T^2 + 1) \\ &\simeq \mathbb{C}[T]/(T - i)(T + i) \\ &\simeq \mathbb{C}[T]/(T - i) \times \mathbb{C}[T]/(T + i) \\ &\simeq \mathbb{C} \times \mathbb{C} \end{aligned}$$

II. Commutative Algebra

using the Chinese remainder theorem, which will be explored later. The action of this isomorphism on a pure tensor is

$$\begin{aligned}x \otimes y &= (a + bi) \otimes (c + di) \mapsto (a + bi) \otimes (c + dT + (T^2 + 1)\mathbb{R}[T]) \\&\mapsto (a + bi)(c + dT) + (T^2 + 1)\mathbb{C}[T] \\&= \underbrace{(ac + bdiT) + (ibc + adT)}_P + (T^2 + 1)\mathbb{C}[T] \\&\mapsto (P + (T - i)\mathbb{C}[T], P + (T + i)\mathbb{C}[T]) \\&\mapsto ((ac - bd) + i(bc + ad), (ac + bd) + i(bc - ad)) = (xy, x\bar{y})\end{aligned}$$

3. Localisation

3.1. Definitions

Definition. A *multiplicative set* or *multiplicatively closed set* $S \subseteq R$ is a subset such that $1 \in S$ and if $a, b \in S$, then $ab \in S$. If $U \subseteq R$ is any set, its *multiplicative closure* \bar{U} is the set

$$\left\{ \prod_{i=1}^n u_i \mid n \geq 0, u_i \in U \right\}$$

which is the smallest multiplicatively closed set containing U .

Example. (i) If R is an integral domain, then $S = R \setminus \{0\}$ is multiplicative.

(ii) More generally, if \mathfrak{p} is a prime ideal in R , then $S = R \setminus \mathfrak{p}$ is multiplicative.

(iii) If $x \in R$, then the set $\{x^n \mid n \geq 0\}$ is multiplicative.

Remark. \mathbb{Q} is obtained from \mathbb{Z} by adding inverses for the elements of the multiplicative subset $\mathbb{Z} \setminus \{0\}$. We have a ring homomorphism $\mathbb{Z} \rightarrow \mathbb{Q}$. We generalise this construction to arbitrary rings and multiplicative sets. In general, injectivity of the ring homomorphism in question may fail.

Definition. Let $S \subseteq R$ be a multiplicative set, and let M be an R -module. Then the *localisation* of M by S is the set $S^{-1}M = M \times S / \sim$ where $(m_1, s_1) \sim (m_2, s_2)$ if and only if there exists $u \in S$ such that $u(s_2m_1 - s_1m_2) = 0$. We write $\frac{m}{s}$ for the equivalence class corresponding to (m, s) . We make $S^{-1}M$ into an R -module by defining

$$\frac{m_1}{s_1} + \frac{m_2}{s_2} = \frac{m_1s_2 + m_2s_1}{s_1s_2}; \quad r \cdot \frac{m}{s} = \frac{rm}{s}$$

We can make $S^{-1}R$ into a ring by defining

$$\frac{r_1}{s_1} \cdot \frac{r_2}{s_2} = \frac{r_1r_2}{s_1s_2}$$

Then $S^{-1}M$ is an $S^{-1}R$ -module by

$$\frac{r}{s} \cdot \frac{m}{t} = \frac{rm}{st}$$

We have the localisation map $R \rightarrow S^{-1}R$ given by $r \mapsto \frac{r}{1}$, which is a ring homomorphism.

We also have the localisation map $M \rightarrow S^{-1}M$ given by $m \mapsto \frac{m}{1}$, which is a homomorphism of R -modules.

We must show that \sim is an equivalence relation. The only nontrivial thing to prove is transitivity. Let

$$u(s_2m_1 - s_1m_2) = 0 = v(s_3m_2 - s_2m_3); \quad u, v \in S$$

Then

$$0 = uv(s_2s_3m_1 - s_1s_3m_2) + uv(s_1s_3m_2 - s_1s_2m_3) = uvs_2(s_3m_1 - s_1m_3); \quad uvs_2 \in S$$

as required. All other operations mentioned are well-defined; the proofs are not enlightening so are omitted.

II. Commutative Algebra

3.2. Universal property for rings

Proposition. Let $U \subseteq R$, and let $S \subseteq R$ be its multiplicative closure. Let $f : R \rightarrow B$ be a ring homomorphism such that $f(u)$ is a unit for all $u \in U$. Then there is a unique ring homomorphism $h : S^{-1}R \rightarrow B$ such that the following diagram commutes.

$$\begin{array}{ccc} R & \xrightarrow{\iota_{S^{-1}R}} & S^{-1}R \\ & \searrow f & \downarrow h \\ & & B \end{array}$$

where $\iota_{S^{-1}R}(r) = \frac{r}{1}$, so in particular, $f(r) = h\left(\frac{r}{1}\right)$.

Thus

$$\text{Hom}_{\text{Ring}}(S^{-1}R, B) \simeq \{\varphi \in \text{Hom}_{\text{Ring}}(R, B) \mid \varphi(U) \subseteq B^\times\}$$

mapping

$$f \mapsto \left(r \mapsto \frac{r}{1}\right); \quad \left(\frac{r}{s} \mapsto \frac{\varphi(r)}{\varphi(s)}\right) \leftarrow \varphi$$

Proof. Let $f : R \rightarrow B$ be a ring homomorphism such that $f(u)$ is a unit for all $u \in U$. Then $f(s)$ is a unit for all $s \in S$. We want to construct a ring homomorphism $h : S^{-1}R \rightarrow B$ such that $f(r) = h\left(\frac{r}{1}\right)$ for all $r \in R$. Such an h must satisfy the following condition.

$$1 = h(1) = h\left(\frac{1}{s} \cdot \frac{s}{1}\right) = h\left(\frac{1}{s}\right)f(s)$$

Thus $h\left(\frac{1}{s}\right) = f(s)^{-1}$. Hence, we must have

$$h\left(\frac{r}{s}\right) = h\left(\frac{1}{s}\right)h\left(\frac{r}{1}\right) = f(s)^{-1}f(r)$$

It thus suffices to show that this h is well-defined; it is then a ring homomorphism satisfying the correct property. If $\frac{r_1}{s_1} = \frac{r_2}{s_2}$, then there is $t \in S$ such that $ts_2r_1 = ts_1r_2$. Applying f ,

$$f(t)f(s_2)f(r_1) = f(t)f(s_1)f(r_2)$$

As $f(t), f(s_1), f(s_2)$ are invertible,

$$\frac{f(r_1)}{f(s_1)} = \frac{f(r_2)}{f(s_2)}$$

so h is well-defined. □

Proposition. Suppose (A, j) has the same universal property of $(S^{-1}R, \iota_{S^{-1}R})$ where $\iota_{S^{-1}R}(r) = \frac{r}{1}$, then there is a unique ring isomorphism $S^{-1}R \rightarrow A$ mapping $\frac{r}{s}$ to $j(s)^{-1}j(r)$.

Remark. (i) Let $\frac{r}{s} \in S^{-1}R$. Then $\frac{r}{s} = \frac{0}{1}$ if and only if there exists $u \in S$ such that $ur = 0$.

- (ii) In particular, $S^{-1}R = 0$ when $\frac{1}{1} = \frac{0}{1}$, which occurs precisely when $0 \in S$.
- (iii) $\ker \iota_{S^{-1}R} = \{r \in R \mid \exists u \in S, ur = 0\}$.
- (iv) $\iota_{S^{-1}R}$ is injective if and only if S contains no zero divisors.
- (v) $\iota_{S^{-1}R}$ is always an epimorphism, but usually not surjective. For example, the map $\iota : \mathbb{Z} \rightarrow \mathbb{Q}$ is epic. Indeed, for $f, g : \mathbb{Q} \rightarrow A$ are such that $f \circ \iota = g \circ \iota$, then

$$f\left(\frac{p}{q}\right) = \frac{f(\iota(p))}{f(\iota(q))} = \frac{g(\iota(p))}{g(\iota(q))} = g\left(\frac{p}{q}\right)$$

Example. (i) Let $f \in R$ and define $S = \{f^n \mid n \geq 0\}$. Define $R_f = S^{-1}R$. Taking for instance $R = \mathbb{Z}$ and $f = 2$,

$$R_f = \left\{ \frac{a}{2^n} \mid a \in \mathbb{Z}, n \geq 0 \right\} = \mathbb{Z}\left[\frac{1}{2}\right]$$

producing the ring of dyadic rational numbers. Since we write $\mathbb{Z}/n\mathbb{Z}$ for the finite quotient ring and \mathbb{Z}_2 for the 2-adic integers, we must use the notation $\mathbb{Z}\left[\frac{1}{2}\right]$ for this particular construction instead. Thus R_f is the zero ring if and only if f is nilpotent.

- (ii) Let $\mathfrak{p} \in \text{Spec } R$, where $\text{Spec } R$ is the set of prime ideals in R . Then $S = R \setminus \mathfrak{p}$ is a multiplicative set. Consider $(R \setminus \mathfrak{p})^{-1}R = R_{\mathfrak{p}}$. For example,

$$\mathbb{Z}_{(3)} = \left\{ \frac{a}{b} \mid a, b \in \mathbb{Z}, 3 \nmid b \right\}$$

3.3. Functoriality

Proposition. Let M be an R -module and $S \subseteq R$ be a multiplicative set. Then there is an isomorphism of $S^{-1}R$ -modules

$$S^{-1}R \otimes_R M \rightarrow S^{-1}M$$

given by $\frac{r}{s} \otimes m \mapsto \frac{rm}{s}$.

Thus the localisation of any module can be reduced to a tensor product with the localisation of a ring.

Proof. Define the map $S^{-1}R \times M \rightarrow S^{-1}M$ mapping $\left(\frac{r}{s}, m\right) \mapsto \frac{rm}{s}$; this is bilinear and thus gives rise to an R -linear map $\varphi : S^{-1}R \otimes_R M \rightarrow S^{-1}M$ with the desired action on pure tensors. One can check that this is in fact $S^{-1}R$ -linear. Clearly φ is surjective by $\frac{1}{s} \otimes m \mapsto \frac{m}{s}$. For injectivity, we first show that every tensor

$$\sum_i \frac{r_i}{s_i} \otimes m_i \in S^{-1}R \otimes_R M$$

II. Commutative Algebra

is pure. We define

$$s = \prod_i s_i; \quad t_j = \prod_{j \neq i} s_j$$

hence

$$\sum_i \frac{r_i}{s_i} \otimes m_i = \sum_i \frac{1}{s_i} \otimes r_i m_i = \sum_i \frac{t_i}{s} \otimes r_i m_i = \sum_i \frac{1}{s} \otimes t_i r_i m_i = \frac{1}{s} \otimes \sum_i t_i r_i m_i$$

as required. Now, it suffices to prove injectivity on pure tensors. If $\varphi\left(\frac{1}{s} \otimes m\right) = \frac{0}{1}$, then there exists $u \in S$ such that

$$u(1m - 0s) = 0 \implies um = 0$$

Thus

$$\frac{1}{s} \otimes m = \frac{u}{us} \otimes m = \frac{1}{us} \otimes um = \frac{1}{us} \otimes 0 = 0$$

as required. \square

The map $S^{-1}R \otimes (-)$ acts on modules and on morphisms. The map $S^{-1}(-)$ acts on modules, and can be extended to act on morphisms in the following way. If $f : N \rightarrow N'$ is R -linear, we produce the commutative diagram

$$\begin{array}{ccc} S^{-1}R \otimes_R N & \xrightarrow{\text{id}_{S^{-1}R} \otimes f} & S^{-1}R \otimes_R N' \\ \sim \downarrow & & \downarrow \sim \\ S^{-1}N & \xrightarrow{S^{-1}(f)} & S^{-1}N' \end{array}$$

with action

$$\begin{array}{ccc} \frac{1}{s} \otimes n & \longmapsto & \frac{1}{s} \otimes f(n) \\ \uparrow & & \downarrow \\ \frac{n}{s} & \longmapsto & \frac{f(n)}{s} \end{array}$$

Then the functor $S^{-1}R \otimes_R (-)$ is naturally isomorphic to the functor $S^{-1}(-)$.

Remark. If A is an R -algebra, then we have an $S^{-1}R$ -linear isomorphism $S^{-1}R \otimes_R A \simeq S^{-1}A$; this is also an isomorphism of $S^{-1}R$ -algebras.

Lemma. Let M be an $S^{-1}R$ -module. Treating M as an R -module, we can define $S^{-1}M$. Then,

$$S^{-1}M \simeq M$$

as $S^{-1}R$ -modules, mapping $\frac{m}{s} \mapsto \frac{1}{s}m$.

Equivalently, $M \simeq S^{-1}R \otimes_R M$ as $S^{-1}R$ -modules, mapping $m \mapsto \frac{1}{1} \otimes m$.

Proof. The localisation map $M \rightarrow S^{-1}M$ maps $m \mapsto \frac{m}{1}$. This is $S^{-1}R$ -linear, and surjective as $\frac{1}{s} \cdot m \mapsto \frac{m}{s}$. To show injectivity, note that $\frac{m}{1} = \frac{0}{1}$ implies there exists $u \in S$ with $um = 0$. Multiplying by $\frac{1}{u}$ as M is an $S^{-1}R$ -module we obtain $m = 0$ as required. \square

3.4. Universal property for modules

Recall that if U has multiplicative closure S ,

$$\text{Hom}_{\text{Ring}}(S^{-1}R, B) \simeq \{\varphi \in \text{Hom}_{\text{Ring}}(R, B) \mid \varphi(U) \subseteq B^\times\}$$

If M is a fixed R -module and L is an $S^{-1}R$ -module, we have

$$\text{Hom}_R(M, L) \simeq \text{Hom}_{S^{-1}R}(S^{-1}M, L)$$

Proposition. Let M be an R -module and L be an $S^{-1}R$ -module. Let $f : M \rightarrow L$ be R -linear. Then there exists a unique $S^{-1}R$ -linear map $h : S^{-1}M \rightarrow L$ such that $f = h \circ i_{S^{-1}M}$.

$$\begin{array}{ccc} M & \xrightarrow{i_{S^{-1}M}} & S^{-1}M \\ & \searrow f & \downarrow h \\ & & L \end{array}$$

As usual with universal properties, this characterises $S^{-1}M$ uniquely up to unique isomorphism.

Proof. We use the natural isomorphism between $S^{-1}(-)$ and $S^{-1}R \otimes_R (-)$. After applying this, we have a map

$$\iota : M \rightarrow S^{-1}R \otimes_R M; \quad m \mapsto \frac{1}{1} \otimes m$$

Let $f : M \rightarrow L$ be R -linear, and define

$$h = \text{id}_{S^{-1}R} \otimes f : S^{-1}R \otimes_R M \rightarrow S^{-1}R \otimes_R L$$

Note that $S^{-1}R \otimes_R L \simeq L$, so we can consider h as mapping to L , with action

$$h\left(\frac{r}{s} \otimes m\right) = \frac{r}{s} f(m)$$

Uniqueness of h follows from the fact that $\{1 \otimes m\}_{m \in M}$ generate $S^{-1}R \otimes_R M$ as an $S^{-1}R$ -module. \square

II. Commutative Algebra

3.5. Exactness

Proposition. The functor $S^{-1}(-)$ is exact. More explicitly, if

$$A \xrightarrow{f} B \xrightarrow{g} C$$

is an exact sequence of R -modules, then

$$S^{-1}A \xrightarrow{S^{-1}f} S^{-1}B \xrightarrow{S^{-1}g} S^{-1}C$$

is an exact sequence of $S^{-1}R$ -modules.

Proof. First,

$$(S^{-1}g) \circ (S^{-1}f) = S^{-1}(g \circ f) = S^{-1}0 = 0$$

so $\text{im } S^{-1}f \subseteq \ker S^{-1}g$. Now suppose $\frac{b}{s} \in \ker S^{-1}g$, so $\frac{g(b)}{s} = \frac{0}{1}$. Hence there exists $u \in S$ such that $ug(b) = 0$. As g is R -linear and $u \in R$, we have $g(ub) = 0$. By exactness, $ub \in \ker g = \text{im } f$. Thus there exists $a \in A$ such that $f(a) = ub$. Hence,

$$\frac{b}{s} = \frac{ub}{us} = \frac{f(a)}{us} = S^{-1}f\left(\frac{a}{us}\right)$$

□

In particular, $S^{-1}R$ is a flat R -module, so for example \mathbb{Q} is a flat \mathbb{Z} -module.

Remark. Suppose $N \subseteq M$ are R -modules, and $\iota : N \rightarrow M$ is the inclusion map. Then applying the localisation, the map $S^{-1}\iota : S^{-1}N \rightarrow S^{-1}M$ given by $\frac{n}{s} \mapsto \frac{n}{s}$ is still injective. Note that the similar result for tensor products fails.

Proposition. Let M be an R -module and N, P be submodules of M . Then,

- (i) $S^{-1}(N + P) = S^{-1}N + S^{-1}P$;
- (ii) $S^{-1}(N \cap P) = S^{-1}N \cap S^{-1}P$;
- (iii) $S^{-1}M / S^{-1}N \simeq S^{-1}(M/N)$ given by $\frac{m}{s} + S^{-1}N \mapsto \frac{m+N}{s}$.

Parts (i) and (ii) rely on a slight abuse of notation, thinking of $S^{-1}N$ as a submodule of $S^{-1}M$. Due to the above remark, this should not cause confusion.

Proof. Part (i). Note that

$$\frac{n+p}{s} = \frac{n}{s} + \frac{p}{s} \in S^{-1}N + S^{-1}P$$

and

$$\frac{n}{s_1} + \frac{p}{s_2} = \frac{s_2n + s_1p}{s_1s_2} \in S^{-1}(N + P)$$

3. Localisation

Part (ii). The forward inclusion is clear. Conversely, suppose $x \in S^{-1}N \cap S^{-1}P$, so $x = \frac{n}{s_1} = \frac{p}{s_2}$. Hence, there exists $u \in S$ such that $us_2n = us_1p = w$. Note $us_2n \in N$ and $us_1p \in P$, so $w \in N \cap P$. Now,

$$x = \frac{n}{s_1} = \frac{us_2n}{us_1s_2} = \frac{w}{us_1s_2} \in S^{-1}(N \cap P)$$

Part (iii). Consider the short exact sequence

$$0 \longrightarrow N \xrightarrow{\iota} M \xrightarrow{\pi} M/N \longrightarrow 0$$

Applying the exact functor $S^{-1}(-)$, we obtain the short exact sequence

$$0 \longrightarrow S^{-1}N \xrightarrow{S^{-1}\iota} S^{-1}M \xrightarrow{S^{-1}\pi} S^{-1}(M/N) \longrightarrow 0$$

Thus

$$(S^{-1}\iota)(S^{-1}N) = S^{-1}N \subseteq S^{-1}M$$

and

$$(S^{-1}\pi)\left(\frac{m}{s}\right) = \frac{m + N}{s}$$

giving the isomorphism as required. \square

Proposition. Let M, N be R -modules. Then

$$S^{-1}M \otimes_{S^{-1}R} S^{-1}N \simeq S^{-1}(M \otimes_R N)$$

Proof. We have already proven that

$$(S^{-1}R \otimes_R M) \otimes_{S^{-1}R} (S^{-1}R \otimes_R N) \simeq S^{-1}R \otimes_R (M \otimes_R N)$$

giving the result as required. \square

Example. Let \mathfrak{p} be a prime ideal in R . Then by setting $S = R \setminus \mathfrak{p}$,

$$M_{\mathfrak{p}} \otimes_{R_{\mathfrak{p}}} N_{\mathfrak{p}} \simeq (M \otimes_R N)_{\mathfrak{p}}$$

3.6. Extension and contraction of ideals

If $f : A \rightarrow B$ is a ring homomorphism and \mathfrak{b} is an ideal in B , the preimage $f^{-1}(\mathfrak{b}) = \mathfrak{b}^c$ is an ideal in A , called its *contraction*. If \mathfrak{a} is an ideal in A , we can generate an ideal $(f(\mathfrak{a})) = \mathfrak{a}^e$ in B , called its *extension*. We show on the first example sheet that for any ring homomorphism $f : A \rightarrow B$, there is a bijection

$$\{\text{contracted ideals of } A\} \leftrightarrow \{\text{extended ideals of } B\}$$

II. Commutative Algebra

noting that the contracted ideals are those ideals with $\mathfrak{a} = \mathfrak{a}^{ec}$, and the extended ideals are those ideals with $\mathfrak{b} = \mathfrak{b}^{ce}$, where the bijection maps $\mathfrak{a} \mapsto \mathfrak{a}^e$ and $\mathfrak{b}^c \mapsto \mathfrak{b}$.

We now study the special case where $f : R \rightarrow S^{-1}R$ is the localisation map of a ring, given by $r \mapsto \frac{r}{1}$. In this case, the extension of an ideal is written $S^{-1}\mathfrak{a} = \mathfrak{a}^e$. We claim that

$$\mathfrak{a}^e = \left\{ \frac{a}{s} \mid a \in \mathfrak{a}, s \in S \right\}$$

Indeed, \mathfrak{a}^e is generated by $\left\{ \frac{a}{1} \mid a \in \mathfrak{a} \right\}$, so \mathfrak{a}^e must contain $\left\{ \frac{a}{s} \mid a \in \mathfrak{a}, s \in S \right\}$, but this is already an ideal. We also claim that

$$\mathfrak{a}^{ec} = \bigcup_{s \in S} (\mathfrak{a} : s); \quad (\mathfrak{a} : s) = \{r \in R \mid rs \in \mathfrak{a}\}$$

Indeed, for $r \in \bigcup_{s \in S} (\mathfrak{a} : s)$, we have $rs = a$ in R for some $s \in S$ and $a \in \mathfrak{a}$, so $\frac{rs}{1} = \frac{a}{1}$, giving $\frac{r}{1} = \frac{a}{s}$, so $r \in \mathfrak{a}^{ec}$ as required. In the other direction, if $r \in \mathfrak{a}^{ec}$, then $\frac{r}{1} = \frac{a}{s}$ for some $s \in S$ and $a \in \mathfrak{a}$, so there exists $u \in S$ such that $rus = ua \in \mathfrak{a}$, so $r \in (\mathfrak{a} : us)$ as required.

Now, let \mathfrak{b} be an ideal of $S^{-1}R$. Then

$$\mathfrak{b}^c = \left\{ r \in R \mid \frac{r}{1} \in \mathfrak{b} \right\}$$

We claim that $\mathfrak{b}^{ce} = \mathfrak{b}$, so all ideals in $S^{-1}R$ are extended. Note that the inclusion $\mathfrak{b}^{ce} \subseteq \mathfrak{b}$ holds for any pair of rings. For the reverse inclusion, consider $\frac{r}{s} \in \mathfrak{b}$, so $\frac{r}{1} \in \mathfrak{b}$. Hence $r \in \mathfrak{b}^c$, so $\frac{r}{1} \in \mathfrak{b}^{ce}$, thus $\frac{r}{s} \in \mathfrak{b}^{ce}$ as \mathfrak{b}^{ce} is an ideal in $S^{-1}R$.

Proposition. Consider the localisation map $R \rightarrow S^{-1}R$ given by $r \mapsto \frac{r}{1}$.

- (i) Every ideal of $S^{-1}R$ is extended.
- (ii) An ideal \mathfrak{a} of R is contracted if and only if the image of S in R/\mathfrak{a} contains no zero divisors.
- (iii) $\mathfrak{a}^e = S^{-1}R$ if and only if $\mathfrak{a} \cap S \neq \emptyset$.
- (iv) There is a bijection

$$\{\mathfrak{p} \in \text{Spec } R \mid \mathfrak{p} \cap S = \emptyset\} \leftrightarrow \text{Spec } S^{-1}R$$

given by $\mathfrak{p} \mapsto \mathfrak{p}^e, \mathfrak{q}^c \mapsto \mathfrak{q}$.

Proof. Part (i). Follows from the fact that $\mathfrak{b}^{ce} = \mathfrak{b}$ for all ideals \mathfrak{b} in $S^{-1}R$.

Part (ii). \mathfrak{a} is contracted if and only if $\mathfrak{a}^{ec} \subseteq \mathfrak{a}$, because the reverse inclusion always holds. This happens if and only if

$$\bigcup_{s \in S} (\mathfrak{a} : s) \subseteq \mathfrak{a}$$

which occurs if and only if

$$\forall r \in R, (Sr \cap \mathfrak{a} \neq \emptyset \implies r \in \mathfrak{a})$$

$$\forall r \in R, (0 + \mathfrak{a} \in S(r + \mathfrak{a}) \implies r + \mathfrak{a} = 0 + \mathfrak{a})$$

which in turn occurs if and only if the image of S in R/\mathfrak{a} contains no zero divisors.

Part (iii). Suppose $\mathfrak{a} \cap S \neq \emptyset$, so let $x \in \mathfrak{a} \cap S$. Then $\frac{x}{x} \in \mathfrak{a}^e$, so $\mathfrak{a}^e = (1) = S^{-1}R$. Conversely, if $\mathfrak{a}^e = S^{-1}R$, then $\frac{1}{1} \in \mathfrak{a}^e$, so $\frac{1}{1} = \frac{a}{s}$ for some $a \in \mathfrak{a}, s \in S$. Therefore there exists $u \in S$ such that $us = ua \in S \cap \mathfrak{a}$.

Part (iv). Consider the contraction map $\text{Spec } S^{-1}R \rightarrow \{\mathfrak{p} \in \text{Spec } R \mid \mathfrak{p} \cap S = \emptyset\}$ given by $\mathfrak{q} \mapsto \mathfrak{q}^e$. We show this is well-defined. In general, a contraction of a prime ideal is always prime. Further, $\mathfrak{p} \in \text{Spec } R$ is contracted if and only if the image of S in R/\mathfrak{p} contains no zero divisors, but R/\mathfrak{p} is an integral domain, so its only zero divisor is zero itself. So this condition is equivalent to the condition $\mathfrak{p} \cap S = \emptyset$. In particular, $\{\mathfrak{p} \in \text{Spec } R \mid \mathfrak{p} \cap S = \emptyset\}$ is precisely the set of contracted prime ideals of R . The map is injective, since if $\mathfrak{q} \in \text{Spec } S^{-1}R$, then $\mathfrak{q}^{ce} = \mathfrak{q}$.

In the other direction, for $\mathfrak{p} \in \text{Spec } R$ such that $\mathfrak{p} \cap S = \emptyset$, it must be contracted, so $\mathfrak{p}^{ec} = \mathfrak{p}$. It therefore remains to show that \mathfrak{p}^e is a prime ideal. We want to show that $S^{-1}R/\mathfrak{p}^e$ is an integral domain. We have that $\mathfrak{p}^e \neq S^{-1}R$ by (iii), so $S^{-1}R/\mathfrak{p}^e$ is not the zero ring, so it suffices to show that this quotient has no zero divisors. To show this, we embed $S^{-1}R/\mathfrak{p}^e$ in the field $FF(R/\mathfrak{p})$.

Consider the composite map

$$R \rightarrow R/\mathfrak{p} \rightarrow FF(R/\mathfrak{p})$$

which is a surjection followed by an injection. This has the property that all elements of S are mapped to units, because $S \cap \mathfrak{p} = \emptyset$. By the universal property of the localisation, we have a map

$$\varphi : S^{-1}R \rightarrow FF(R/\mathfrak{p}); \quad \frac{r}{s} \mapsto \frac{r + \mathfrak{p}}{s + \mathfrak{p}}$$

It suffices to show that $\ker \varphi = \mathfrak{p}^e$, then the result holds by the isomorphism theorem. Let $\frac{r}{s} \in \ker \varphi$, so $\frac{r + \mathfrak{p}}{s + \mathfrak{p}} = \frac{0}{1}$ in $FF(R/\mathfrak{p})$. Observe that $\text{im } \varphi \subseteq \bar{S}^{-1}(R/\mathfrak{p})$, where \bar{S} is the image of S in R/\mathfrak{p} . Restricting the range, we can consider φ as a map from $S^{-1}R$ to $\bar{S}^{-1}(R/\mathfrak{p})$. So $\varphi\left(\frac{r}{s}\right) = \frac{0}{1}$ implies that there exists $u + \mathfrak{p} \in \bar{S}$ such that $(u + \mathfrak{p})(r + \mathfrak{p}) = 0$, so $ur + \mathfrak{p} = 0$. In particular, $u \in S$ and $ur \in \mathfrak{p}$. Hence $\frac{r}{s} = \frac{ur}{us}$ where $ur \in \mathfrak{p}$ and $us \in S$, so $\frac{r}{s} \in \mathfrak{p}^e$.

For the other direction, take $x \in \mathfrak{p}^e$, so $x = \frac{p}{s}$ for $p \in \mathfrak{p}, s \in S$. Then $\varphi(x) = \frac{p + \mathfrak{p}}{s + \mathfrak{p}} = 0$, so $x \in \ker \varphi$. \square

II. Commutative Algebra

It is not true in general that the extensions of prime ideals are prime.

Definition. If I is an ideal in R , the *radical* of I is the ideal

$$\sqrt{I} = \{r \in R \mid \exists n \geq 1, r^n \in I\}$$

Proposition. Let I be an ideal in a ring R . Then

$$\sqrt{I} = \bigcap_{I \subseteq \mathfrak{p} \in \text{Spec } R} \mathfrak{p}$$

Proof. Let $x \in \sqrt{I}$. Then $x^n \in I$ for some $n \geq 1$. For every $\mathfrak{p} \in \text{Spec } R$, if $I \subseteq \mathfrak{p}$, then $x^n \in \mathfrak{p}$, so $x \in \mathfrak{p}$. Conversely, suppose $x^n \notin I$ for all $n \geq 1$. As $I \neq R$, we have $R/I \neq 0$. Let \bar{x} be the image of x in R/I , and consider

$$(R/I)_{\bar{x}} = \{\bar{x}^n \mid n \geq 1\}^{-1}(R/I)$$

This is not the zero ring, because $x^n \notin I$ for all $n \geq 1$. Therefore, $(R/I)_{\bar{x}}$ has a prime ideal, as it contains a maximal ideal. By the bijection described in part (iv) of the previous result, this prime ideal corresponds to a prime ideal of R/I that avoids \bar{x} . This in turn corresponds to a prime ideal $\mathfrak{p} \in \text{Spec } R$ that contains I and avoids x . Hence $x \notin \bigcap_{I \subseteq \mathfrak{p} \in \text{Spec } R} \mathfrak{p}$. \square

3.7. Local properties

Definition. A ring R is *local* if it has exactly one maximal ideal.

We write $\text{mSpec } R$ for the set of maximal ideals of R .

Example. Let $\mathfrak{p} \in \text{Spec } R$. Then there is a bijection between the prime ideals of R contained within \mathfrak{p} to $\text{Spec } R_{\mathfrak{p}}$, mapping $\mathfrak{n} \mapsto \mathfrak{n}R_{\mathfrak{p}}$ and $\mathfrak{q}^c \mapsto \mathfrak{q}$. Hence, all prime ideals of $R_{\mathfrak{p}}$ are contained in $\mathfrak{p}^e = \mathfrak{p}R_{\mathfrak{p}}$. Thus $(R_{\mathfrak{p}}, \mathfrak{p}R_{\mathfrak{p}})$ is a local ring.

Example. Recall that

$$\mathbb{Z}_{(2)} = \left\{ \frac{a}{b} \mid a, b \in \mathbb{Z}, 2 \nmid b \right\}$$

This ring is local, and the unique maximal ideal is

$$(2)\mathbb{Z}_{(2)} = \left\{ \frac{2a}{b} \mid a, b \in \mathbb{Z}, 2 \nmid b \right\}$$

Proposition. Let M be an R -module. The following are equivalent.

- (i) M is the zero module;
- (ii) $M_{\mathfrak{p}}$ is the zero module for all prime ideals $\mathfrak{p} \in \text{Spec } R$;
- (iii) $M_{\mathfrak{m}}$ is the zero module for all maximal ideals $\mathfrak{m} \in \text{mSpec } R$.

Informally, for modules, being zero is a local property.

Proof. First, note that (i) implies (ii) and (ii) implies (iii). We show that (iii) implies (i). Suppose that M is not the zero module, so let $m \in M$ be a nonzero element. Consider $\text{Ann}_R(m) = \{r \in R \mid rm = 0\}$. This is an ideal of R , but is a proper ideal because $1 \notin \text{Ann}_R(m)$. Let \mathfrak{m} be a maximal ideal of R containing $\text{Ann}_R(m)$. Now, $\frac{m}{1} \in M_{\mathfrak{m}} = 0$. Thus, $\frac{m}{1} = \frac{0}{1}$, so $um = 0$ for some $u \in R \setminus \mathfrak{m}$. But then $u \notin \text{Ann}_R(m)$, giving a contradiction. \square

Proposition. Let $f : M \rightarrow N$ be an R -linear map. The following are equivalent.

- (i) f is injective;
- (ii) $f_{\mathfrak{p}} : M_{\mathfrak{p}} \rightarrow N_{\mathfrak{p}}$ is injective for every prime ideal $\mathfrak{p} \in \text{Spec } R$;
- (iii) $f_{\mathfrak{m}} : M_{\mathfrak{m}} \rightarrow N_{\mathfrak{m}}$ is injective for every maximal ideal $\mathfrak{m} \in \text{mSpec } R$.

The same result holds for surjectivity.

Proof. The fact that (i) implies (ii) follows directly from the fact that localisation at \mathfrak{p} is an exact functor. Clearly (ii) implies (iii). Suppose that $f_{\mathfrak{m}}$ is injective for each $\mathfrak{m} \in \text{mSpec } R$. We have the following exact sequence.

$$0 \longrightarrow \ker f \longrightarrow M \xrightarrow{f} N$$

As $(-)_{\mathfrak{p}}$ is exact, the sequence

$$0 \longrightarrow (\ker f)_{\mathfrak{m}} \longrightarrow M_{\mathfrak{m}} \xrightarrow{f_{\mathfrak{m}}} N_{\mathfrak{m}}$$

is exact. But by assumption, $(\ker f)_{\mathfrak{m}} = \ker(f_{\mathfrak{m}}) = 0$. So $(\ker f)_{\mathfrak{m}} = 0$ for all maximal ideals $\mathfrak{m} \in \text{mSpec } R$, so $\ker f = 0$. \square

Proposition. Let M be an R -module. The following are equivalent.

- (i) M is a flat R -module;
- (ii) $M_{\mathfrak{p}}$ is a flat $R_{\mathfrak{p}}$ -module for every prime ideal $\mathfrak{p} \in \text{Spec } R$;
- (iii) $M_{\mathfrak{m}}$ is a flat $R_{\mathfrak{m}}$ -module for every maximal ideal $\mathfrak{m} \in \text{mSpec } R$.

Proof. (i) implies (ii). Note that $M_{\mathfrak{p}} \simeq R_{\mathfrak{p}} \otimes_R M$ as $R_{\mathfrak{p}}$ -modules, by extension of scalars. Since extension of scalars preserves flatness, $M_{\mathfrak{p}}$ is flat.

Clearly (ii) implies (iii).

(iii) implies (i). Let $f : N \rightarrow P$ be an R -linear injective map. Let $\mathfrak{m} \in \text{mSpec } R$. Then $f_{\mathfrak{m}} : N_{\mathfrak{m}} \rightarrow P_{\mathfrak{m}}$ is injective by the previous proposition. Note that the following diagram

II. Commutative Algebra

commutes.

$$\begin{array}{ccc} N_{\mathfrak{m}} \otimes_{R_{\mathfrak{m}}} M_{\mathfrak{m}} & \xrightarrow{f_{\mathfrak{m}} \otimes \text{id}_{M_{\mathfrak{m}}}} & P_{\mathfrak{m}} \otimes_{R_{\mathfrak{m}}} M_{\mathfrak{m}} \\ \sim \downarrow & & \downarrow \sim \\ (N \otimes_R M)_{\mathfrak{m}} & \xrightarrow{(f \otimes \text{id}_M)_{\mathfrak{m}}} & (P \otimes_R M)_{\mathfrak{m}} \end{array}$$

Hence $(f \otimes \text{id}_M)_{\mathfrak{m}}$ is injective. Since this holds for each $\mathfrak{m} \in \text{mSpec } R$, the map $f \otimes \text{id}_M$ must be injective, as required. \square

Example. An R -module M is *locally free* if $M_{\mathfrak{p}}$ is a free $R_{\mathfrak{p}}$ -module for every prime ideal $\mathfrak{p} \in \text{Spec } R$. Consider $R = \mathbb{C} \otimes \mathbb{C}$. Then

$$\text{Spec } R = \{\mathfrak{p} \times \mathbb{C} \mid \mathfrak{p} \in \text{Spec } \mathbb{C}\} \cup \{\mathbb{C} \times \mathfrak{p} \mid \mathfrak{p} \in \text{Spec } \mathbb{C}\} = \{\mathbb{C} \times (0), (0) \times \mathbb{C}\}$$

The map $\mathbb{C} \times \mathbb{C} \rightarrow \mathbb{C}$ given by $(a, b) \mapsto b$ sends $(\mathbb{C} \times \mathbb{C}) \setminus \mathbb{C} \times (0)$ to units. Thus, by the universal property of the localisation, we have a map

$$(\mathbb{C} \times \mathbb{C})_{\mathbb{C} \times (0)} \rightarrow \mathbb{C}; \quad \frac{(a, b)}{(c, d)} \mapsto \frac{b}{d}$$

This is clearly surjective, and one can check that this is also injective. Thus $(\mathbb{C} \times \mathbb{C})_{\mathbb{C} \times (0)} \simeq \mathbb{C}$ is a field. Similarly, $(\mathbb{C} \times \mathbb{C})_{(0) \times \mathbb{C}}$ is a field. So for every $\mathbb{C} \times \mathbb{C}$ -module M and prime ideal $\mathfrak{p} \in \text{Spec}(\mathbb{C} \times \mathbb{C})$, the module $M_{\mathfrak{p}}$ is a \mathbb{C} -vector space, so is free. Thus every module over $\mathbb{C} \times \mathbb{C}$ is locally free, but not every module over $\mathbb{C} \times \mathbb{C}$ is free. For example, take $M = \mathbb{C} \times \{0\}$ as a $\mathbb{C} \times \mathbb{C}$ -module. One can show that M is not the zero module, and not free of rank at least 1, so cannot be free.

3.8. Localisations as quotients

Let $U \subseteq R$, and let $S \subseteq R$ be its multiplicative closure. We can define

$$R_U = R[\{T_u\}_{u \in U}] / I_U; \quad I_U = (\{uT_u - 1\}_{u \in U})$$

We claim that $R_U = S^{-1}R$ as rings, and also as R -algebras. Writing \bar{u} and \bar{T}_u to be the images of these elements in R_U , the isomorphism maps

$$\bar{T}_u \mapsto \frac{1}{u}; \quad rT_{u_1} \dots T_{u_\ell} + I_U \mapsto \frac{r}{u_1 \dots u_\ell}$$

This is because R_U has the universal property of $S^{-1}R$. Indeed, for any $f : R \rightarrow A$ mapping U to units, there is a unique h making the following diagram commute.

$$\begin{array}{ccc} R & \xrightarrow{\quad} & R_U \\ & \searrow f & \downarrow \text{\scriptsize h} \\ & & A \end{array}$$

3. Localisation

Note that A is an R -algebra via f , so the diagram commutes if and only if h is an R -algebra homomorphism. We have

$$\mathrm{Hom}_{R\text{-algebra}}(R_U, A) \simeq \{\varphi : U \rightarrow A \mid f(u)\varphi(u) = 1\}$$

But the the right hand side is a singleton.

Example. Let $x \in R$, and consider $R_x = R_{\{1, x, x^2, \dots\}}$. Here,

$$R_x \simeq R[T]_{(xT - 1)}$$

4. Integrality, finiteness, and finite generation

4.1. Nakayama's lemma

Proposition (Cayley–Hamilton theorem). Let M be a finitely generated R -module, and let $f : M \rightarrow M$ be an R -linear endomorphism. Let \mathfrak{a} be an ideal in R such that $f(M) \subseteq \mathfrak{a}M$. Then, we have an equality in $\text{End}_R M$

$$f^n + a_1 f^{n-1} + \cdots + a_n f^0 = 0; \quad f^r = \underbrace{f \circ \cdots \circ f}_{r \text{ times}}$$

where $a_i \in \mathfrak{a}$.

Proof. Let $M = \text{span}_R \{m_1, \dots, m_n\}$, so $\mathfrak{a}M = \text{span}_{\mathfrak{a}} \{m_1, \dots, m_n\}$. Then

$$\begin{pmatrix} f(m_1) \\ \vdots \\ f(m_n) \end{pmatrix} = P \begin{pmatrix} m_1 \\ \vdots \\ m_n \end{pmatrix}; \quad P \in M_{n \times n}(\mathfrak{a})$$

Let $\rho : R \rightarrow \text{End } M$ be the structure ring homomorphism of M as an R -module. Then we can define $R[T] \rightarrow \text{End } M$ by $T \mapsto f$, making M into an $R[T]$ -module. Hence,

$$T \begin{pmatrix} m_1 \\ \vdots \\ m_n \end{pmatrix} = P \begin{pmatrix} m_1 \\ \vdots \\ m_n \end{pmatrix}$$

Thus

$$Q \begin{pmatrix} m_1 \\ \vdots \\ m_n \end{pmatrix} = 0; \quad Q = TI_n - P$$

Multiplying by the adjugate matrix $\text{adj } Q$ on the left on both sides,

$$(\det Q) \begin{pmatrix} m_1 \\ \vdots \\ m_n \end{pmatrix} = 0$$

In particular, $(\det Q)m = 0$ for all $m \in M$, as the m_i generate M . Hence, $m \mapsto (\det Q)m = (\det Q)|_{T=f}$ is 0 in $\text{End}_R M$. Finally, note that $\det Q$ is a monic polynomial, and all other coefficients lie in \mathfrak{a} . \square

Corollary. Let M be a finitely generated R -module, and let \mathfrak{a} be an ideal in R . If $\mathfrak{a}M = M$, then there exists $a \in \mathfrak{a}$ such that $am = m$ for all $m \in M$.

Proof. Apply the Cayley–Hamilton theorem with $f = \text{id}_M$. We obtain a polynomial

$$(1 + a_1 + \cdots + a_n) \text{id}_M = 0$$

Take $a = -(a_1 + \cdots + a_n)$. \square

4. Integrality, finiteness, and finite generation

Definition. The *Jacobson radical* of a ring R , denoted $J(R)$, is the intersection of all maximal ideals of R .

Example. (i) If (R, \mathfrak{m}) is a local ring, then $J(R) = \mathfrak{m}$.

(ii) $J(\mathbb{Z}) = \{0\}$.

Proposition. Let $x \in R$. Then $x \in J(R)$ if and only if $1 - xy$ is a unit for every $y \in R$.

Proof. First, let $x \in J(R)$, and suppose $y \in R$ is such that $1 - xy$ is not a unit. Then $(1 - xy)$ is a proper ideal, so it is contained in a maximal ideal \mathfrak{m} . But as $x \in J(R)$, we must have $x \in \mathfrak{m}$, giving $1 = 1 - xy + xy \in \mathfrak{m}$, contradicting that \mathfrak{m} is a maximal ideal.

Now suppose $x \notin J(R)$, so there is a maximal ideal \mathfrak{m} such that $x \notin \mathfrak{m}$. Then $\mathfrak{m} + (x) = R$ as \mathfrak{m} is maximal. In particular, there exists $t \in \mathfrak{m}$ and $y \in R$ such that $t + xy = 1$, or equivalently, $1 - xy = t \in \mathfrak{m}$. Note that t cannot be a unit, because it is contained in a proper ideal. \square

Proposition (Nakayama's lemma). Let M be a finitely generated R -module, and let $\mathfrak{a} \subseteq J(R)$ be an ideal of R such that $\mathfrak{a}M = M$. Then $M = 0$.

This lemma is more useful when $J(R)$ is large, so is particularly useful when applied to local rings.

Proof. By the above corollary, there exists $a \in \mathfrak{a}$ such that $am = m$ for all $m \in M$, or equivalently, $(1 - a)m = 0$. By assumption, $a \in J(R)$, so $1 - a$ is a unit in R . Hence $m = 0$. \square

Corollary. Let M be a finitely generated R -module, and let $N \subseteq M$ be a submodule. Let $\mathfrak{a} \subseteq J(R)$ be an ideal in R such that $N + \mathfrak{a}M = M$. Then $N = M$.

This can be applied to find generating sets for M .

Proof. Note that

$$\mathfrak{a}(M/N) = (\mathfrak{a}M + N)/N = M/N$$

so $M/N = 0$ by Nakayama's lemma. \square

4.2. Integral and finite extensions

Definition. Let A be an R -algebra, and let $x \in A$. Then x is *integral* over R if there exists a monic polynomial $f \in R[T]$ such that $f(x) = 0$.

Example. (i) If $R = k$ is a field, then x is integral over k if and only if x is algebraic over k .

(ii) We will prove later that

II. Commutative Algebra

- (a) the \mathbb{Z} -integral elements of \mathbb{Q} are \mathbb{Z} ;
- (b) the \mathbb{Z} -integral elements of $\mathbb{Q}[\sqrt{2}]$ are $\mathbb{Z}[\sqrt{2}]$;
- (c) the \mathbb{Z} -integral elements of $\mathbb{Q}[\sqrt{5}]$ are $\mathbb{Z}\left[\frac{1+\sqrt{5}}{2}\right] \not\subseteq \mathbb{Z}[\sqrt{5}]$.

Definition. Let M be an R -module. We say that M is *faithful* if the structure homomorphism $\rho : R \rightarrow \text{End } M$ is injective. Equivalently, for every nonzero ring element r , there exists $m \in M$ such that $rm \neq 0$.

Example. Let $R \subseteq A$ be rings, and let A be an R -module in the natural way. Then A is a faithful R -module, as if $r \neq 0$, then $r1_A = r \neq 0$.

Proposition. Let $R \subseteq A$ be rings and $x \in A$, and consider A as an $R[x]$ -module. Then x is integral over R if and only if there exists $M \subseteq A$ such that

- (i) M is a faithful $R[x]$ -module; and
- (ii) M is finitely generated as an R -module.

Condition (i) is that M is an R -submodule of A , $xM \subseteq M$, and M is faithful over $R[x]$.

Proof. First, assume conditions (i) and (ii) hold. We have an R -linear map $f : M \rightarrow M$ given by multiplication by x , as $xM \subseteq M$. As M is a finitely generated R -module, we can apply the Cayley–Hamilton theorem to find

$$f^n + r_1 f^{n-1} + \cdots + r_n f^0 = 0; \quad r_i \in R$$

in $\text{End}_R M$. Then, evaluating at $m \in M$,

$$(x^n + r_1 x^{n-1} + \cdots + r_n x^0)m = 0$$

As this holds for all m , and M is a faithful $R[x]$ -module, we must have

$$x^n + r_1 x^{n-1} + \cdots + r_n x^0 = 0$$

Thus x is integral over R .

Now suppose x is integral over R . Then

$$x^n + r_1 x^{n-1} + \cdots + r_n x^0 = 0$$

for some $r_1, \dots, r_n \in R$. We define

$$M = \text{span}_R \{x^0, \dots, x^{n-1}\}$$

This is finitely generated, and satisfies $xM \subseteq M$. M is faithful over $R[x]$ as it contains $x^0 = 1$. □

Definition. Let A be an R -algebra. Then A is

4. Integrality, finiteness, and finite generation

- (i) *integral* over R , if all of its elements are integral over R ;
- (ii) *finite* over R , if A is finitely generated as an R -module.

Proposition. Let A be an R -algebra. Then the following are equivalent.

- (i) A is a finitely generated R -algebra and is integral over R ;
- (ii) A is generated as an R -algebra by a finite set of integral elements;
- (iii) A is finite over R .

Proof. (i) implies (ii). The generators for A are integral.

(ii) implies (iii). Suppose A is generated by $\alpha_1, \dots, \alpha_m$ as an R -algebra, and the α_i are integral over R . As α_i is integral,

$$\alpha_i^{n_i} + r_{i,1}\alpha_i^{n_i-1} + \dots + r_{i,n_i}\alpha_i^0 = 0$$

Hence $\alpha_i^{n_i}$ lies in the R -linear span of $\{\alpha_i^0, \dots, \alpha_i^{n_i-1}\}$. Thus, every element is an R -linear combination of products of the form $\alpha_1^{e_1} \dots \alpha_n^{e_n}$, which in turn lies in the R -linear span of products of the same form where all e_i are less than the corresponding n_i . This is a finite set, so A is finitely generated as an R -module.

(iii) implies (i). As A is finitely generated as an R -module, it must be finitely generated as an R -algebra. Let $\alpha \in A$; we show α is integral over R . Let $\rho : R \rightarrow A$ be the structure homomorphism of A as an R -algebra. Then $\rho(R) \subseteq A$, and consider $(\rho(R))[\alpha] \subseteq A$. Now, A is a $(\rho(R))[\alpha]$ -module, and is faithful because $1_A \in A$. As A is a finitely generated $\rho(R)$ -module, the previous proposition shows that α is $\rho(R)$ -integral. Equivalently, α is R -integral. \square

Proposition. Let A be an R -algebra and let \mathcal{O} be the set of elements of A that are integral over R . Then \mathcal{O} is an R -subalgebra of A .

Proof. Let $x, y \in \mathcal{O}$. Then $\{x, y\}$ is a finite set of R -integral elements, so the set generates an integral R -subalgebra of A . Hence $x+y, xy$ lie in this subalgebra, and so they are integral. \square

Proposition. Let $A \subseteq B \subseteq C$ be rings. Then,

- (i) if C is finite over B and B is finite over A , then C is finite over A ;
- (ii) if C is integral over B and B is integral over A , then C is integral over A .

Proof. Part (i). Suppose that

$$C = \text{span}_B \{\gamma_1, \dots, \gamma_n\}; \quad B = \text{span}_A \{\beta_1, \dots, \beta_\ell\}$$

Then

$$C = \text{span}_A \{\gamma_i \beta_j \mid i \leq n, j \leq \ell\}$$

II. Commutative Algebra

Part (ii). Let $c \in C$, so $f(c) = 0$ for

$$f(T) = T^n + b_1 T^{n-1} + \cdots + b_n T^0 \in B[T]$$

Then $f \in A'[T]$, where $A' = A[b_1, \dots, b_n]$. The inclusion $A \subseteq A'$ is generated as an A -algebra by finitely many integral elements. Similarly, $A' \subseteq A'[c]$ is generated as an A -algebra by c , which is integral over A' as $f \in A'[T]$. By the previous result, both extensions are finite. Then, by part (i), $A \subseteq A'[c]$ is finite, so c is integral over A . \square

4.3. Integral closure

Definition. Let $A \subseteq B$ be rings. The *integral closure* of A in B is the set \bar{A} of elements of B that are integral over A , which is an A -algebra. We say that A is *integrally closed* in B if $\bar{A} = A$.

Definition. Let A be an integral domain. In this case, the *integral closure* of A is the integral closure of A in its field of fractions $FF(A)$. We say that A is *integrally closed* if it is integrally closed in its field of fractions.

Example. (i) $\mathbb{Z}[\sqrt{5}]$ is not integrally closed, because $\alpha = \frac{1+\sqrt{5}}{2} \in FF(\mathbb{Z}[\sqrt{5}]) = \mathbb{Q}[\sqrt{5}]$, and $\alpha^2 - \alpha - 1 = 0$ so it is $\mathbb{Z}[\sqrt{5}]$ -integral.

(ii) \mathbb{Z} is integrally closed.

(iii) If k is a field, $k[T_1, \dots, T_n]$ is integrally closed.

Examples (ii) and (iii) are special cases of the following result.

Proposition. Let A be a unique factorisation domain. Then A is integrally closed.

Proof. Let $x \in FF(A) \setminus A$, and write $x = \frac{a}{b}$ with $a \in A, b \in A \setminus \{0\}$. As A is a unique factorisation domain, we can assume there is a prime p such that $p \mid b$ and $p \nmid a$. If x is integral over A , then

$$\left(\frac{a}{b}\right)^n + a_1 \left(\frac{a}{b}\right)^{n-1} + \cdots + a_n \left(\frac{a}{b}\right)^0 = 0$$

Multiplying by b^n ,

$$a^n = -b(a_1 b_0 a^{n-1} + \cdots + a_n b^{n-1} a^0)$$

But as $p \mid b$, we must have $p \mid a^n$, so $p \mid a$, which is a contradiction. \square

Lemma. Let $A \subseteq B$ be rings, and let \bar{A} be the integral closure of A in B . Then \bar{A} is integrally closed in B .

Taking the integral closure is an idempotent operation.

4. Integrality, finiteness, and finite generation

Proof. Let $x \in B$ be integral over \bar{A} . Then, we have

$$A \subseteq \bar{A} \subseteq \bar{A}[x]$$

The first extension is integral by definition, and the second is integral by the above proposition, as x is integral over \bar{A} . By transitivity of integrality, $\bar{A}[x]$ is integral over A , so in particular, x is integral over A . Thus $x \in \bar{A}$. \square

Proposition. Let $A \subseteq B$ be rings.

- (i) if B is integral over A and \mathfrak{b} is an ideal in B , then B/\mathfrak{b} is integral over $A/\mathfrak{b}c$;
- (ii) if B is integral over A and $S \subseteq A$ is a multiplicative set, then $S^{-1}B$ is integral over $S^{-1}A$;
- (iii) if \bar{A} is the integral closure of A in B and $S \subseteq A$ is a multiplicative set, then $S^{-1}\bar{A}$ is the integral closure of $S^{-1}A$ in $S^{-1}B$, so $\overline{S^{-1}A} = S^{-1}\bar{A}$.

The proofs follow directly from the definitions.

Lemma. Let $A \subseteq B$ be an integral extension of rings. Then

- (i) $A \cap B^\times = A^\times$;
- (ii) if A, B are integral domains, then A is a field if and only if B is a field.

Proof. *Part (i).* One inclusion is clear: $A^\times \subseteq A \cap B^\times$. Suppose $a \in A$ and a is a unit in B with inverse $b \in B$; we show that $b \in A$. As b is integral over A ,

$$b^n + a_1 b^{n-1} + \cdots + a_n b^0 = 0; \quad a_i \in A$$

Multiplying by a^{n-1} ,

$$b + \underbrace{a_1 + a_2 a^1 + \cdots + a_n a^{n-1}}_{\in A} = 0$$

Hence b must lie in A .

Part (ii). Suppose B is a field. Then

$$A^\times = A \cap (B \setminus \{0\}) = A \setminus \{0\}$$

Hence A is a field. Conversely, suppose A is a field. Let $b \in B$ be a nonzero element; we want to show that b is a unit in B . As b is integral over A ,

$$b^n + a_1 b^{n-1} + \cdots + a_n b^0 = 0; \quad a_i \in A$$

Let n be minimal with this property. Then

$$b \underbrace{(b^{n-1} + a_1 b^{n-2} + \cdots + a_{n-1} b^0)}_{\in A} = -a_n$$

II. Commutative Algebra

Note that $b \neq 0$ by assumption, and $\Delta \neq 0$ by minimality. As B is an integral domain, $a_n \neq 0$. Because A is a field, a_n is invertible. Thus

$$b(-a_n^{-1}\Delta) = 1 \implies b \in B^\times$$

□

Corollary. Let $A \subseteq B$ be an integral extension of rings, and let \mathfrak{q} be a prime ideal in B . Then \mathfrak{q} is a maximal ideal of B if and only if $\mathfrak{q}^c = \mathfrak{q} \cap A$ is a maximal ideal in A .

Proof. We have an embedding of rings

$$A/\mathfrak{q} \cap A \hookrightarrow B/\mathfrak{q}$$

which is an integral extension of integral domains. By the previous result, one is a field if and only if the other is, so $\mathfrak{q} \cap A$ is maximal in A if and only if \mathfrak{q} is maximal in B . □

4.4. Noether normalisation

Definition. Let A be a k -algebra, and let $x_1, \dots, x_n \in A$. We say that x_1, \dots, x_n are *k -algebraically independent* if for every nonzero polynomial $p \in k[T_1, \dots, T_n]$, we have $p(x_1, \dots, x_n) \neq 0$. Equivalently, the k -algebra homomorphism $k[T_1, \dots, T_n] \rightarrow A$ given by $T_i \mapsto x_i$ is injective.

Theorem (Noether's normalisation theorem). Let k be a field, and let $A \neq 0$ be a finitely generated k -algebra. Then there exist $x_1, \dots, x_n \in A$ which are k -algebraically independent and A is finite over $A' = k[x_1, \dots, x_n]$.

We first present an example of the method used in the proof.

Example. Let $A = k[T, T^{-1}] \simeq k[X, Y]/(XY - 1)$. We claim that $k[T] \subseteq k[T, T^{-1}]$ is not a finite extension. Indeed, suppose it were finite. Then T^{-1} would be integral over $k[T]$, so

$$(T^{-1})^n \in \text{span}_{k[T]} \{(T^{-1})^0, \dots, (T^{-1})^{n-1}\}$$

Multiplying by T^n , we have

$$1 \in \text{span}_{k[T]}(T^n, \dots, T)$$

which is false. However, if $c \in k$ is a scalar which we will choose later,

$$A = k[T, T^{-1}] = k[T, T^{-1} - cT]$$

We claim that $k[T^{-1} - cT] \subseteq A$ is a finite extension for most values of c , and in particular, for at least one. First, note $T^{-1}T - 1 = 0$, and then change variables to

$$((T^{-1} - cT) + cT)T - 1 = 0 \implies \underbrace{c}_{\in k} T^2 + \underbrace{(T^{-1} - cT)T}_{\in k[T^{-1} - cT]} - \underbrace{1}_{\in k[T^{-1} - cT]} = 0$$

Hence if $c \neq 0$, T is integral over $k[T^{-1} - cT]$.

4. Integrality, finiteness, and finite generation

Proof. In this proof, we will assume k is infinite, although the theorem is true even if k is finite. We will proceed by induction on the minimal number of generators of A as a k -algebra, which we will denote m . For the case $m = 0$, we have $A = k$, so we can take $A' = k$.

Suppose that A is generated as a k -algebra by $x_1, \dots, x_m \in A$. If x_1, \dots, x_m are algebraically independent, then we can take $A' = A$. Otherwise, we claim that there are $c_1, \dots, c_{m-1} \in k$ such that x_m is integral over

$$B = k[x_1 - c_1 x_m, \dots, x_{m-1} - c_{m-1} x_m]$$

Assuming that this holds, we have $A = B[x_m]$, so $B \subseteq A$ is a finite extension. But B is generated by $m - 1$ elements, so by induction B contains $z_1, \dots, z_n \in B$ which are k -algebraically independent, and B is finite over $A' = k[z_1, \dots, z_n]$. Then A is finite over A' by transitivity of finiteness.

We now prove the claim. As x_1, \dots, x_m are not algebraically independent over k , there is a nonzero polynomial $f \in k[T_1, \dots, T_m]$ such that $f(x_1, \dots, x_m) = 0$. We want to show that x_m is integral over B . Write f as the sum of its homogeneous parts, and let F be the part of highest degree $\deg f = r$. For scalars $c_1, \dots, c_{m-1} \in k$ which will be chosen later, we define

$$\begin{aligned} g(T_1, \dots, T_m) &= f(T_1 + c_1 T_m, \dots, T_{m-1} + c_{m-1} T_m, T_m) \\ &= \underbrace{F(c_1, \dots, c_{m-1}, 1)}_{\in k} T_m^r + \text{terms of lower degree in } T_m \text{ with coefficients in } k[T_1, \dots, T_{m-1}] \end{aligned}$$

Note that

$$g(x_1 - c_1 x_m, \dots, x_{m-1} - c_{m-1} x_m, x_m) = f(x_1, \dots, x_m) = 0$$

but as a polynomial in T_m over $k[T_1, \dots, T_{m-1}]$, it has degree at most r , and the coefficient of T_m^r is $F(c_1, \dots, c_{m-1}, 1)$. As $F(T_1, \dots, T_m)$ is a nonzero homogeneous polynomial, $F(T_1, \dots, T_{m-1}, 1)$ is not the zero polynomial. Thus there are c_1, \dots, c_{m-1} such that $F(c_1, \dots, c_{m-1}, 1) \neq 0$ as k is an infinite field. \square

4.5. Hilbert's Nullstellensatz

Proposition (Zariski's lemma). Let $k \subseteq L$ be fields where L is finitely generated as a k -algebra. Then $\dim_k L$ is finite.

Proof. By Noether normalisation, we have

$$k \subseteq k[x_1, \dots, x_n] \subseteq L$$

where x_1, \dots, x_n are algebraically independent over k , and L is finite over $k[x_1, \dots, x_n]$. As this is an integral extension of integral domains and L is a field, $k[x_1, \dots, x_n]$ must be a field. But as $k[x_1, \dots, x_n]$ is a polynomial algebra over k , the x_i cannot be invertible. Hence $n = 0$, so $k \subseteq L$ is finite as required. \square

Definition. Let $k \subseteq \Omega$ be an extension of fields, where Ω is algebraically closed.

II. Commutative Algebra

(i) Let $S \subseteq k[T_1, \dots, T_n]$. We define

$$\mathbb{V}(S) = \{\mathbf{x} \in \Omega^n \mid \forall f \in S, f(\mathbf{x}) = 0\}$$

Sets of this form are called *k-algebraic* subsets of Ω^n .

(ii) Let $X \subseteq \Omega^n$. We define

$$I(X) = \{f \in k[T_1, \dots, T_n] \mid \forall \mathbf{x} \in X, f(\mathbf{x}) = 0\}$$

Note that $\mathbb{V}(S) = \mathbb{V}(I)$, where I is the ideal generated by S . Recall that for every finite field extension $k \subseteq L$, there is a k -algebra embedding $L \rightarrow \Omega$, because Ω is algebraically closed.

Theorem. Let $\mathfrak{a} \subseteq k[T_1, \dots, T_n]$ be an ideal. Then

(i) (weak Nullstellensatz) $\mathbb{V}(\mathfrak{a}) = \emptyset$ if and only if $1 \in \mathfrak{a}$;

(ii) (strong Nullstellensatz) $I(\mathbb{V}(\mathfrak{a})) = \sqrt{\mathfrak{a}}$.

Proof. Weak Nullstellensatz. Clearly if $1 \in \mathfrak{a}$ then $\mathbb{V}(\mathfrak{a}) = \emptyset$, as $1 \neq 0$. Now suppose $1 \notin \mathfrak{a}$. There is a maximal ideal $\mathfrak{m} \in \text{mSpec } k[T_1, \dots, T_n]$ such that $\mathfrak{a} \subseteq \mathfrak{m}$. Then $L = k[T_1, \dots, T_n]_{\mathfrak{m}}$ is a field, which is finitely generated over k as an algebra. By Zariski's lemma, this extension is finitely generated as a module. Hence, there is an injective k -algebra homomorphism $L \rightarrow \Omega$. Composing with the quotient map, we obtain a k -algebra homomorphism $\varphi : k[T_1, \dots, T_n] \rightarrow \Omega$ with kernel \mathfrak{m} . Now, let

$$\mathbf{x} = (\varphi(T_1), \dots, \varphi(T_n)) \in \Omega^n$$

We claim that this is a common solution to all polynomials in \mathfrak{a} . Note that for $f \in k[T_1, \dots, T_n]$, we have $\varphi(f) = f(\mathbf{x})$. Therefore, for all $f \in \mathfrak{a}$, we have $f \in \ker \varphi$ so $f(\mathbf{x}) = \varphi(f) = 0$.

Strong Nullstellensatz. Let $f \in \sqrt{\mathfrak{a}}$. Then $f^\ell \in \mathfrak{a}$ for some $\ell \geq 1$, and therefore, $f^\ell(\mathbf{x}) = 0$ for all $\mathbf{x} \in \mathbb{V}(\mathfrak{a})$. As Ω is an integral domain, $f(\mathbf{x}) = 0$ for all $\mathbf{x} \in \mathbb{V}(\mathfrak{a})$. Hence $f \in I(\mathbb{V}(\mathfrak{a}))$.

Conversely, suppose $f \in I(\mathbb{V}(\mathfrak{a}))$, so for all $\mathbf{x} \in \mathbb{V}(\mathfrak{a})$, we have $f(\mathbf{x}) = 0$. We want to show that $f \in \sqrt{\mathfrak{a}}$. To do this, we show that \bar{f} is nilpotent in $k[T_1, \dots, T_n]_{\mathfrak{a}}$. It suffices to show that

$$\left(k[T_1, \dots, T_n]_{\mathfrak{a}}\right)_{\bar{f}} = 0$$

Note that

$$\left(k[T_1, \dots, T_n]_{\mathfrak{a}}\right)_{\bar{f}} \simeq k[T_1, \dots, T_n, T_{n+1}]_{\mathfrak{b}}; \quad \mathfrak{b} = \mathfrak{a}^e + (T_{n+1}f - 1)$$

We will show that $1 \in \mathfrak{b}$, or equivalently by the weak Nullstellensatz, $\mathbb{V}(\mathfrak{b}) = \emptyset$.

Suppose $\mathbf{x} = (x_1, \dots, x_{n+1}) \in \mathbb{V}(\mathfrak{b}) \subseteq \Omega^{n+1}$. Define $\mathbf{x}_0 = (x_1, \dots, x_n)$, so $\mathbf{x}_0 \in \mathbb{V}(\mathfrak{a})$. In particular, $f(\mathbf{x}_0) = 0$, as $f \in I(\mathbb{V}(\mathfrak{a}))$. Thus $f(\mathbf{x}) = 0$. Now, $(T_{n+1}f - 1)(\mathbf{x}) = -1 \neq 0$, but $(T_{n+1}f - 1) \in \mathfrak{b}$, so \mathbf{x} is not a common solution to all polynomials in \mathfrak{b} , which is a contradiction. \square

4. Integrality, finiteness, and finite generation

One can easily derive the weak Nullstellensatz from the strong Nullstellensatz.

Note that

- (i) $\sqrt{\sqrt{\mathfrak{a}}} = \sqrt{\mathfrak{a}}$.
- (ii) If $X \subseteq Y \subseteq \Omega^n$, then $I(X) \supseteq I(Y)$.
- (iii) If $S \subseteq T \subseteq k[T_1, \dots, T_n]$, then $\mathbb{V}(S) \supseteq \mathbb{V}(T)$.
- (iv) If $S \subseteq k[T_1, \dots, T_n]$, then $S \subseteq I(\mathbb{V}(S))$.
- (v) If $X \subseteq \Omega^n$, then $X \subseteq \mathbb{V}(I(X))$.
- (vi) If $X \subseteq \Omega^n$ is an algebraic set, then $X = \mathbb{V}(I(X))$, as $X = \mathbb{V}(\mathfrak{a})$ gives

$$\mathbb{V}(\mathfrak{a}) \subseteq \mathbb{V}(I(\mathbb{V}(\mathfrak{a}))) \subseteq \mathbb{V}(\mathfrak{a})$$

- (vii) If $X \subseteq \Omega^n$, then $I(X)$ is a radical ideal.

Proposition. Let $k = \Omega$ be an algebraically closed field, and let $n \geq 0$. Then we have an inclusion-reversing bijection

$$\{k\text{-algebraic subsets of } \Omega^n\} \leftrightarrow \{\text{radical ideals of } k[T_1, \dots, T_n]\}$$

given by $X \mapsto I(X)$ and $\mathbb{V}(\mathfrak{a}) \mapsto \mathfrak{a}$.

Proof. We have already shown that $I(X)$ is radical, and $X = \mathbb{V}(I(X))$ if X is an algebraic set. For the converse, let $\mathfrak{a} \subseteq k[T_1, \dots, T_n]$ be a radical ideal. Then $I(\mathbb{V}(\mathfrak{a})) = \sqrt{\mathfrak{a}} = \mathfrak{a}$ by the strong Nullstellensatz. \square

Remark. Every prime ideal \mathfrak{p} is radical, as $x^n \in \mathfrak{p}$ implies $x \in \mathfrak{p}$. In particular, every maximal ideal is radical.

Corollary. Let $k = \Omega$ be an algebraically closed field. Then we have a bijection

$$\Omega^n \leftrightarrow \text{mSpec } k[T_1, \dots, T_n]$$

given by $\mathbf{x} = (x_1, \dots, x_n) \mapsto (T_1 - x_1, \dots, T_n - x_n) = \mathfrak{m}_{\mathbf{x}}$.

Proof. First, note that $\mathfrak{m}_{\mathbf{x}}$ is a maximal ideal for every \mathbf{x} , since it is the kernel of the map $k[T_1, \dots, T_n] \rightarrow \Omega$ given by $T_i \mapsto x_i$. Also, $\mathfrak{m}_{\mathbf{x}} = I(\{\mathbf{x}\})$. Indeed, the inclusion $\mathfrak{m}_{\mathbf{x}} \subseteq I(\{\mathbf{x}\})$ is clear, and $I(\{\mathbf{x}\})$ is a proper ideal of $k[T_1, \dots, T_n]$, so they must be equal by maximality. Note that $\mathbb{V}(\mathfrak{m}_{\mathbf{x}}) = \{\mathbf{x}\}$. Hence the claim follows from the inclusion-reversing bijection, as maximal ideals correspond to minimal nonempty k -algebraic sets. \square

Definition. We say that $X \subseteq \Omega^n$ is *irreducible* if X cannot be expressed as the union of two strictly smaller algebraic subsets.

Proposition. $X \subseteq \Omega^n$ is irreducible if and only if $I(X)$ is prime.

II. Commutative Algebra

4.6. Integrality over ideals

Definition. Let $A \subseteq B$ be an extension of rings, and let $\mathfrak{a} \subseteq A$ be an ideal. We say that $x \in B$ is integral over \mathfrak{a} if

$$x^n + a_1x^{n-1} + \cdots + a_nx^0 = 0$$

for some $a_1, \dots, a_n \in \mathfrak{a}$. The *integral closure* of \mathfrak{a} in B is the set of elements of B that are integral over \mathfrak{a} .

Proposition. Let $A \subseteq B$ be an extension of rings, and let \bar{A} be the integral closure of A in B . Let \mathfrak{a} be an ideal of A . Then the integral closure of \mathfrak{a} in B is $\sqrt{\mathfrak{a}\bar{A}}$, the radical in \bar{A} of the extension of \mathfrak{a} to \bar{A} .

Proof. If $b \in B$ is integral over \mathfrak{a} , then

$$b^n + a_1b^{n-1} + \cdots + a_nb^0 = 0; \quad a_i \in \mathfrak{a}$$

In particular, b lies in \bar{A} , and so all of its powers lie in \bar{A} as \bar{A} is a ring. Using the integrality equation for b , we observe that $b^n \in \mathfrak{a}\bar{A}$, hence $b \in \sqrt{\mathfrak{a}\bar{A}}$.

Now, suppose $b \in \sqrt{\mathfrak{a}\bar{A}}$. Then $b^n \in \mathfrak{a}\bar{A}$ for some n , so

$$b^n = \sum_{i=1}^m a_i x_i; \quad a_i \in \mathfrak{a}, x_i \in \bar{A}$$

Define $M = A[x_1, \dots, x_m]$. The generators lie in \bar{A} , so M is an A -algebra generated by finitely many integral elements over A . Hence M is a finite A -algebra. Note that $b^n M \subseteq \mathfrak{a}M$ by the equation for b^n , thought of as an extension of A -modules.

Now define $f : M \rightarrow M$ by multiplication by b^n . This satisfies $f(M) \subseteq \mathfrak{a}M$, and f is A -linear. Thus by the Cayley–Hamilton theorem,

$$f^\ell + \alpha_1 f^{\ell-1} + \cdots + \alpha_\ell f^0 = 0 \in \text{End}_R M; \quad \alpha_i \in \mathfrak{a}$$

Evaluating this at $1_A \in M$,

$$b^{n\ell} + \alpha_1 b^{n(\ell-1)} + \cdots + \alpha_\ell b^0 = 0 \in B$$

This is an integrality relation for b is \mathfrak{a} -integral. □

Hence, the integral closure of an ideal is closed under sums and products.

Corollary. Let $A \subseteq B$ be an extension of rings, and let \mathfrak{a} be an ideal of A . Then $b \in B$ is \mathfrak{a} -integral if and only if b is $\sqrt{\mathfrak{a}}$ -integral.

4. Integrality, finiteness, and finite generation

Proof. By the previous proposition, it suffices to show that

$$\sqrt{\mathfrak{a}A} = \sqrt{\sqrt{\mathfrak{a}}A}$$

The forwards inclusion is clear. For the other direction, it is a general fact that $\sqrt{I}^e \subseteq \sqrt{I^e}$, so

$$\sqrt{\mathfrak{a}}A \subseteq \sqrt{\mathfrak{a}A}$$

Taking radicals on both sides,

$$\sqrt{\sqrt{\mathfrak{a}}A} \subseteq \sqrt{\sqrt{\mathfrak{a}A}} = \sqrt{\mathfrak{a}A}$$

□

Proposition. Let A be an integrally closed integral domain (in its field of fractions). Let $A \subseteq B$ be an extension of rings, let \mathfrak{a} be an ideal in A , and let $b \in B$. The following are equivalent:

- (i) b is integral over \mathfrak{a} ;
- (ii) b is algebraic over $FF(A)$ with minimal polynomial over $FF(A)$ of the form

$$T^n + a_1 T^{n-1} + \cdots + a_n T^0 = 0; \quad a_i \in \sqrt{\mathfrak{a}}$$

Note that there is an embedding $FF(A) \subseteq FF(B)$.

Proof. Suppose (ii) holds. Then b is integral over $\sqrt{\mathfrak{a}}$ by definition. Thus, by the above corollary, b is integral over \mathfrak{a} .

Now suppose (i) holds. We have an integrality equation

$$b^n + a_1 b^{n-1} + \cdots + a_n b^0 = 0; \quad a_i \in \mathfrak{a}$$

Define

$$h = T^n + a_1 T^{n-1} + \cdots + a_n T^0 \in (FF(A))[T]$$

so $h(b) = 0$, so certainly b is algebraic over $FF(A)$. Let $f \in (FF(A))[T]$ be the minimal polynomial of b over $FF(A)$. Let $FF(A) \subseteq \Omega$ where Ω is an algebraically closed field, so

$$f = \prod_{i=1}^{\ell} (T - \alpha_i); \quad \alpha_1 = b, \alpha_i \in \Omega$$

We want to show that the coefficients of f are in $\sqrt{\mathfrak{a}}$. By the previous proposition, together with the fact that A is integrally closed, the integral closure of \mathfrak{a} in $FF(A)$ is $\sqrt{\mathfrak{a}} \subseteq A$. So it suffices to show that the coefficients of f lie in $FF(A)$ and are integral over \mathfrak{a} . As f is the minimal polynomial over $FF(A)$, the first part holds by definition.

II. Commutative Algebra

Expanding brackets in the equation for f , the coefficients of f are sums of products of the α_i . The proposition above implies that the integral closure of \mathfrak{a} in Ω is closed under sums and products, so it suffices to show that the α_i are all integral over \mathfrak{a} . As the α_i and b have the same minimal polynomial f over $FF(A)$, there is an isomorphism of $FF(A)$ -algebras $\varphi_i : FF(A)[b] \rightarrow FF(A)[\alpha_i]$ that maps b to α_i . Then as $h(b) = 0$ and $h \in (FF(A))[T]$, we must have $h(\alpha_i) = h(\varphi_i(b)) = \varphi_i(h(b)) = \varphi_i(0) = 0$. \square

4.7. Cohen–Seidenberg theorems

If $A \subseteq B$ is an extension of rings, the inclusion $\iota : A \rightarrow B$ gives rise to $\iota^* : \text{Spec } B \rightarrow \text{Spec } A$ given by $\iota^*(\mathfrak{q}) = \mathfrak{q} \cap A$. We will study the fibres of this induced map on spectra.

Proposition (incomparability). Let $A \subseteq B$ be an integral extension, and let $\mathfrak{q}, \mathfrak{q}'$ be prime ideals of B . Suppose that \mathfrak{q} and \mathfrak{q}' contract to the same prime ideal $\mathfrak{p} = \mathfrak{q} \cap A = \mathfrak{q}' \cap A$ of A , and that $\mathfrak{q} \subseteq \mathfrak{q}'$. Then $\mathfrak{q} = \mathfrak{q}'$.

In this section, we will write $B_{\mathfrak{p}}$ to abbreviate $(A \setminus \mathfrak{p})^{-1}B$.

Proof. Define $S = A \setminus \mathfrak{p}$. Then \mathfrak{q} and \mathfrak{q}' are prime ideals of B not intersecting S . Hence $\mathfrak{q} = (S^{-1}\mathfrak{q})^c$, where $S^{-1}\mathfrak{q} = \mathfrak{q}B_{\mathfrak{p}}$ is the extension of \mathfrak{q} to $S^{-1}B$, due to the bijection

$$\{\mathfrak{p} \in \text{Spec } R \mid \mathfrak{p} \cap S = \emptyset\} \leftrightarrow \text{Spec } S^{-1}R$$

It suffices to show that $\mathfrak{q}B_{\mathfrak{p}} = \mathfrak{q}'B_{\mathfrak{p}}$, as then they are the contractions of the same ideal. Note that

$$\mathfrak{q}B_{\mathfrak{p}} \cap A_{\mathfrak{p}} = S^{-1}\mathfrak{q} \cap S^{-1}A = S^{-1}(\mathfrak{q} \cap A) = S^{-1}\mathfrak{p} = \mathfrak{p}A_{\mathfrak{p}}$$

Similarly, $\mathfrak{q}'B_{\mathfrak{p}} \cap A_{\mathfrak{p}} = \mathfrak{p}A_{\mathfrak{p}}$, which is a maximal ideal of $A_{\mathfrak{p}}$. As $A \subseteq B$ is an integral extension, $A_{\mathfrak{p}} \subseteq B_{\mathfrak{p}}$ is also an integral extension. Recall that if the contraction of a prime ideal is maximal in an integral extension, the prime ideal was maximal. Thus, $\mathfrak{q}B_{\mathfrak{p}} \subseteq \mathfrak{q}'B_{\mathfrak{p}}$ are maximal ideals of $B_{\mathfrak{p}}$, so they must coincide. \square

Proposition (lying over). Let $A \subseteq B$ be an integral extension of rings, and let $\mathfrak{p} \in \text{Spec } A$. Then there is a prime ideal $\mathfrak{q} \in \text{Spec } B$ such that $\mathfrak{q} \cap A = \mathfrak{p}$. In other words, $\iota^* : \text{Spec } B \rightarrow \text{Spec } A$ is surjective.

Proof. We have a commutative diagram

$$\begin{array}{ccc} A & \longrightarrow & B \\ \downarrow & & \downarrow \beta \\ A_{\mathfrak{p}} & \longrightarrow & B_{\mathfrak{p}} = (A \setminus \mathfrak{p})^{-1}B \end{array}$$

Let \mathfrak{m} be a maximal ideal of $B_{\mathfrak{p}}$. Then $A_{\mathfrak{p}} \subseteq B_{\mathfrak{p}}$ is an integral extension, so \mathfrak{m} contracts to a maximal ideal $\mathfrak{m} \cap A_{\mathfrak{p}}$ of $A_{\mathfrak{p}}$. But there is exactly one maximal ideal in $A_{\mathfrak{p}}$, namely $\mathfrak{p}A_{\mathfrak{p}}$. Note that $\mathfrak{p}A_{\mathfrak{p}}$ contracts to \mathfrak{p} under the map $A \rightarrow A_{\mathfrak{p}}$.

4. Integrality, finiteness, and finite generation

We have that \mathfrak{m} contracts to \mathfrak{p} under the map $A \rightarrow A_{\mathfrak{p}} \rightarrow B_{\mathfrak{p}}$, but this is the same as the map $A \rightarrow B \rightarrow B_{\mathfrak{p}}$, so $\beta^{-1}(\mathfrak{m}) \cap A = \mathfrak{p}$. Note that $\beta^{-1}(\mathfrak{m})$ is a prime ideal, as required. \square

Theorem (going up). Let $A \subseteq B$ be an integral extension of rings. Let $\mathfrak{p}_1 \subseteq \mathfrak{p}_2$ be prime ideals in A , and let $\mathfrak{q}_1 \in \text{Spec } B$ be a prime ideal such that $\mathfrak{q}_1 \cap A = \mathfrak{p}_1$. Then there is a prime ideal $\mathfrak{q}_2 \in \text{Spec } B$ such that $\mathfrak{q}_1 \subseteq \mathfrak{q}_2$, and $\mathfrak{q}_2 \cap A = \mathfrak{p}_2$.

$$\begin{array}{ccc} \mathfrak{q}_1 & \xrightarrow{\subseteq} & \mathfrak{q}_2 \\ \cap A \downarrow & & \downarrow \cap A \\ \mathfrak{p}_1 & \xrightarrow{\subseteq} & \mathfrak{p}_2 \end{array}$$

Proof. We have an injection $A/\mathfrak{p}_1 \rightarrow B/\mathfrak{q}_1$ given by $a + \mathfrak{p}_1 \mapsto a + \mathfrak{q}_1$. This is an integral extension, so by lying over, there is a prime ideal $\mathfrak{q}_2/\mathfrak{q}_1$ of B/\mathfrak{q}_1 that contracts to $\mathfrak{p}_2/\mathfrak{p}_1$ in A/\mathfrak{p}_1 . We claim that $\mathfrak{q}_2 \cap A = \mathfrak{p}_2$. In the diagram

$$\begin{array}{ccc} A & \longrightarrow & B \\ \downarrow & & \downarrow \\ A/\mathfrak{p}_1 & \longrightarrow & B/\mathfrak{q}_1 \end{array}$$

we obtain contractions of prime ideals

$$\begin{array}{ccc} \mathfrak{p}_2 & & \mathfrak{q}_2 \\ \uparrow & & \uparrow \\ \mathfrak{p}_2/\mathfrak{p}_1 & \longleftarrow & \mathfrak{q}_2/\mathfrak{q}_1 \end{array}$$

hence \mathfrak{q}_2 contracts to \mathfrak{p}_2 , as required. \square

Theorem (going down). Let $A \subseteq B$ be an integral extension of integral domains, and suppose that A is integrally closed (in its field of fractions). Let $\mathfrak{p}_1 \supseteq \mathfrak{p}_2$ be prime ideals in A , and let $\mathfrak{q}_1 \in \text{Spec } B$ be a prime ideal such that $\mathfrak{q}_1 \cap A = \mathfrak{p}_1$. Then there is a prime ideal $\mathfrak{q}_2 \in \text{Spec } B$ such that $\mathfrak{q}_1 \supseteq \mathfrak{q}_2$, and $\mathfrak{q}_2 \cap A = \mathfrak{p}_2$.

$$\begin{array}{ccc} \mathfrak{q}_1 & \xleftarrow{\supseteq} & \mathfrak{q}_2 \\ \cap A \downarrow & & \downarrow \cap A \\ \mathfrak{p}_1 & \xleftarrow{\supseteq} & \mathfrak{p}_2 \end{array}$$

Proof. Consider the map $A \rightarrow B \rightarrow B_{\mathfrak{q}_1}$. These maps are injective as B is an integral domain, so we can think of these as inclusions of rings. We want to prove that there is a prime ideal $\mathfrak{n} \in \text{Spec } B_{\mathfrak{q}_1}$ such that $\mathfrak{n} \cap A = \mathfrak{p}_2$. This suffices, as $(\mathfrak{n} \cap B) \cap A = \mathfrak{p}_2$ is a contraction of a

II. Commutative Algebra

prime ideal $\mathfrak{q}_2 = \mathfrak{n} \cap B$ of B contained in \mathfrak{q}_1 to $\mathfrak{p}_2 \in \text{Spec } A$. In other words, we want to show that \mathfrak{p}_2 is a contracted ideal under the map $A \rightarrow B_{\mathfrak{q}_1}$. As contracted ideals are contracted from their own extension, it suffices to show that $(\mathfrak{p}_2 B_{\mathfrak{q}_1}) \cap A \subseteq \mathfrak{p}_2$, noting that the converse inclusion always holds.

Note that $\mathfrak{p}_2 B_{\mathfrak{q}_1} = (\mathfrak{p}_2 B) B_{\mathfrak{q}_1}$. Let $\frac{y}{s} \in (\mathfrak{p}_2 B) B_{\mathfrak{q}_1} \cap A$, where $y \in \mathfrak{p}_2 B$ and $s \in B \setminus \mathfrak{q}_1$. As $A \subseteq B$ is an integral extension, the integral closure of \mathfrak{p}_2 in B is $\sqrt{\mathfrak{p}_2 B}$. In particular, y is integral over \mathfrak{p}_2 . Since A is integrally closed and y is integral over \mathfrak{p}_2 , the minimal polynomial of $y \in FF(B)$ over $FF(A)$ has the form

$$y^r + u_1 y^{r-1} + \cdots + u_r y^0 = 0; \quad u_i \in \sqrt{\mathfrak{p}_2} = \mathfrak{p}_2$$

We can write $y = \frac{y}{s} \cdot s$, where $y, s \in FF(B)$ and $\frac{y}{s} \in FF(A)$. Hence,

$$\left(\frac{y}{s} \cdot s\right)^r + u_1 \left(\frac{y}{s} \cdot s\right)^{r-1} + \cdots + u_r \left(\frac{y}{s} \cdot s\right)^0 = 0$$

Multiplying by $\left(\frac{s}{y}\right)^r$,

$$s^r + \left(\frac{s}{y}\right)^1 u_1 s^{r-1} + \cdots + \left(\frac{s}{y}\right)^r u_r s^0 = 0; \quad u_i \in \sqrt{\mathfrak{p}_2} = \mathfrak{p}_2$$

This must be the same minimal polynomial of s as an element of $FF(B)$ over $FF(A)$. As $s \in B$, s is integral over A , so the coefficients in this polynomial must lie in A .

$$\left(\frac{s}{y}\right)^1 u_1, \dots, \left(\frac{s}{y}\right)^r u_r \in A$$

Suppose $\frac{y}{s} \notin \mathfrak{p}_2$. Then

$$u_i = \left(\frac{y}{s}\right)^i \cdot \left(\frac{s}{y}\right)^i u_i$$

But

$$u_i \in \mathfrak{p}_2; \quad \left(\frac{y}{s}\right)^i \in A \setminus \mathfrak{p}_2; \quad \left(\frac{s}{y}\right)^i u_i \in A$$

By primality, $\left(\frac{s}{y}\right)^i u_i \in \mathfrak{p}_2$. As this holds for all i , the coefficients in the equation for s lie in \mathfrak{p}_2 , so

$$s^r \in \mathfrak{p}_2 B \subseteq \mathfrak{p}_1 B = (\mathfrak{q}_1 \cap A) B \subseteq \mathfrak{q}_1$$

Hence $s \in \mathfrak{q}_1$ by primality, giving a contradiction. □

5. Primary decomposition

Definition. Let I be an ideal of R . I is

- (i) *prime* if $R/I \neq 0$ and 0 is the only zero divisor of R/I ;
- (ii) *radical* if the only nilpotent element of R/I is zero;
- (iii) *primary* if $R/I \neq 0$ and every zero divisor in R/I is nilpotent.

The prime ideals precisely those ideals that are both radical and primary. R is radical but not prime or primary.

Example. (i) Let $R = \mathbb{Z}$. The ideal (6) is radical but not primary, as $R/(6)$ contains zero divisors 2, 3 which are not nilpotent. The ideal (9) is primary but not radical.

- (ii) More generally, let $R = \mathbb{Z}$ and $x \neq 0$. Then (x) is prime if and only if $x = 0$ or $|x|$ is prime, and (x) is radical if and only if x is squarefree. (x) is primary if and only if $x = p^n$ for some prime p and $n \geq 1$.

Proposition. Let I be a proper ideal in R . Then

- (i) If I is primary, then $\mathfrak{p} = \sqrt{I}$ is prime. We say I is \mathfrak{p} -primary.
- (ii) If \sqrt{I} is maximal, then I is primary.
- (iii) If $\mathfrak{q}_1, \dots, \mathfrak{q}_n$ are \mathfrak{p} -primary, then $\bigcap_{i=1}^n \mathfrak{q}_i$ is also \mathfrak{p} -primary.
- (iv) If I has a *primary decomposition* $I = \bigcap_{i=1}^n \mathfrak{q}_i$ where the \mathfrak{q}_i are primary, then I has a minimal primary decomposition $\bigcap_{j=1}^m \mathfrak{r}_j$ where the $\sqrt{\mathfrak{r}_j}$ are distinct and no \mathfrak{r}_j can be dropped.
- (v) If R is Noetherian, then every proper ideal has a primary decomposition.

In \mathbb{Z} ,

$$(90) = (2) \cap (3^2) \cap (5)$$

Primary decomposition therefore generalises prime factorisation. Note that for a prime ideal \mathfrak{p} , if \mathfrak{p}^n is primary, then \mathfrak{p}^n is \mathfrak{p} -primary, because $\sqrt{\mathfrak{p}^n} = \mathfrak{p}$.

Example. (i) Not every primary ideal is a power of a prime ideal. For instance, consider $R = k[X, Y]$ and $\mathfrak{q} = (X, Y^2)$. We claim that this is primary. For instance, $\sqrt{\mathfrak{q}} = (X, Y)$ is maximal, so \mathfrak{q} is (X, Y) -primary. Alternatively,

$$k[X, Y]_{(X, Y^2)} \simeq k[Y]_{(Y^2)}$$

If $f \in k[Y]$ satisfies $f \in (Y^2)$ so it is a zero divisor, then $Y \mid f$, so $f + (Y^2)$ is nilpotent. Now, if $\mathfrak{q} = \mathfrak{p}^n$, then

$$(X, Y) = \sqrt{\mathfrak{q}} = \sqrt{\mathfrak{p}^n} = \mathfrak{p}$$

II. Commutative Algebra

But

$$(X, Y) \supsetneq (X, Y^2) \supsetneq (X, Y)^2$$

So \mathfrak{q} is not a power of $\mathfrak{p} = (X, Y)$.

(ii) If \mathfrak{p} is prime, \mathfrak{p}^n need not be primary. Let

$$R = k[X, Y, Z]/(XY - Z^2) = k[\bar{X}, \bar{Y}, \bar{Z}]; \quad \mathfrak{p} = (\bar{X}, \bar{Z})$$

where $\bar{X}, \bar{Y}, \bar{Z}$ are the images of X, Y, Z under the quotient map. We claim that \mathfrak{p} is prime, but \mathfrak{p}^2 is not primary. Indeed,

$$R/\mathfrak{p} \simeq k[X, Y, Z]/(X, Z, XY - Z^2) \simeq k[X, Y, Z]/(X, Z) \simeq k[Y]$$

which is an integral domain, so \mathfrak{p} is prime. For the second part,

$$\mathfrak{p}^2 = (\bar{X}^2, \bar{X} \cdot \bar{Z}, \bar{Z}^2)$$

Then $\bar{X} \cdot \bar{Y} = \bar{Z}^2 \in \mathfrak{p}^2$, that is,

$$(\bar{X} + \mathfrak{p}^2)(\bar{Y} + \mathfrak{p}^2) = 0 + \mathfrak{p}^2$$

But $\bar{X} + \mathfrak{p}^2 \neq 0$ and $\bar{Y} + \mathfrak{p}^2 \neq 0$. Hence $\bar{Y} + \mathfrak{p}^2$ is a zero divisor in R/\mathfrak{p}^2 . Note that

$$R/\mathfrak{p}^2 \simeq k[X, Y, Z]/(XY - Z^2, X^2, XZ, Z^2) \simeq k[X, Y, Z]/(XY, X^2, Z^2)$$

so $Y + \mathfrak{p}^2$ is not nilpotent.

Theorem. Let $\bigcap_{i=1}^n \mathfrak{q}_i$ be a minimal primary decomposition for an ideal I of R , and let $\mathfrak{p}_i = \sqrt{\mathfrak{q}_i}$ for each i . Then

- (i) (*associated prime ideals of I*) The prime ideals $\mathfrak{p}_1, \dots, \mathfrak{p}_n$ are determined only by I , even though there may not be a unique minimal primary decomposition.
- (ii) (*isolated prime ideals of I*) The minimal elements of $\{\mathfrak{p}_1, \dots, \mathfrak{p}_n\}$, ordered by inclusion, are exactly the minimal prime ideals of R that contain I . An associated prime ideal that is not isolated is called *embedded*.
- (iii) (*isolated primary components of I*) If $\mathfrak{p}_1, \dots, \mathfrak{p}_t$ are the isolated prime ideals of I for $t \leq n$, then $\mathfrak{q}_1, \dots, \mathfrak{q}_t$ are determined only by I .

Example. Let $R = k[X, Y]$ and $I = (X^2, XY)$. We have primary decompositions

$$I = (X) \cap (X, Y)^2 = (X) \cap (X^2, Y)$$

Note that

$$\sqrt{(X)} = (X); \quad \sqrt{(X, Y)^2} = (X, Y); \quad \sqrt{(X^2, Y)} = (X, Y)$$

The associated primes of I are (X) and (X, Y) . The isolated prime is (X) and the embedded prime is (X, Y) .

5. Primary decomposition

Remark. Let $I = \bigcap_{i=1}^n q_i$ be a minimal primary decomposition with $\sqrt{q_i} = \mathfrak{p}_i$. Suppose $\mathfrak{p}_1, \dots, \mathfrak{p}_t$ are the isolated primes. Then

$$\sqrt{I} = \sqrt{\bigcap_{i=1}^n q_i} = \bigcap_{i=1}^n \sqrt{q_i} = \bigcap_{i=1}^n \mathfrak{p}_i = \bigcap_{i=1}^t \mathfrak{p}_i$$

This is a primary decomposition of \sqrt{I} , and one can check that this is minimal. All associated primes in this decomposition are isolated. Going from I to \sqrt{I} , we only ‘remember’ the isolated primes.

Analogously, let $R = k[T_1, \dots, T_n]$, where $k \subseteq \mathbb{C}$. Then $\mathbb{V}(I) = \mathbb{V}(\sqrt{I})$ and $I(\mathbb{V}(I)) = \sqrt{I}$. Hence, taking the algebraic set of I ‘remembers’ the radical of I and nothing else.

6. Direct and inverse limits

6.1. Limits and completions

Definition. Let \mathcal{C} be a category.

- (i) A *directed set* (I, \leq) is a partially ordered set such that for all $a, b \in I$, there exists $c \in I$ such that $a, b \leq c$.
- (ii) A *direct system* on a directed set (I, \leq) is a pair $((X_i)_{i \in I}, (f_{ij})_{i \leq j})$ where $X_i \in \text{ob } \mathcal{C}$ and $f_{ij} : X_i \rightarrow X_j$, such that $f_{ii} = 1_{X_i}$ and $f_{ik} = f_{jk} \circ f_{ij}$.
- (iii) An *inverse system* on (I, \leq) is a pair $((Y_i)_{i \in I}, (h_{ij})_{i \leq j})$ where $Y_i \in \text{ob } \mathcal{C}$ and $h_{ij} : Y_j \rightarrow Y_i$, such that $h_{ii} = 1_{Y_i}$ and $h_{ik} = h_{ij} \circ h_{jk}$.

Remark. An inverse system in \mathcal{C} is the same as a direct system in \mathcal{C}^{op} .

Example. Let $I = (\mathbb{N}, \leq)$.

- (i) Let p be a prime, and let $X_i = \mathbb{F}_{p^{i!}}$. Recall that if $a \mid b$, then there is an embedding $\varphi : \mathbb{F}_{p^a} \rightarrow \mathbb{F}_{p^b}$. The collection of embeddings $\mathbb{F}_{p^a} \rightarrow \mathbb{F}_{p^b}$ is then given by $x \mapsto (\varphi(x))^{p^c}$ where $0 \leq c < a - 1$. The map $f_{i(i+1)} : \mathbb{F}_{p^{i!}} \rightarrow \mathbb{F}_{p^{(i+1)!}}$ is defined to be one such embedding. A general embedding f_{ij} is given by the composite $f_{(j-1)j} \circ \cdots \circ f_{i(i+1)}$. This creates a direct system on I .
- (ii) Let $Y_i = \mathbb{Z}/p^i\mathbb{Z}$, and let $h_{ij} : \mathbb{Z}/p^j\mathbb{Z} \rightarrow \mathbb{Z}/p^i\mathbb{Z}$ be the natural projection. This is an inverse system on I .

Definition. Let (I, \leq) be a directed set.

- (i) Let $D = ((X_i)_{i \in I}, (f_{ij})_{i \leq j})$ be a direct system on I . Then the *direct limit* of D is

$$\varinjlim X_i = \left(\prod_{i \in I} X_i \right) / \sim$$

where for $x_i \in X_i$ and $x_j \in X_j$,

$$x_i \sim x_j \iff \exists k \geq i, j, f_{ik}(x_i) = f_{jk}(x_j)$$

Equivalently, one can define \sim to be the smallest equivalence relation containing $x_i \sim f_{ij}(x_i)$.

- (ii) Let $E = ((Y_i)_{i \in I}, (h_{ij})_{i \leq j})$ be an inverse system on I . Then the *inverse limit* of E is

$$\varprojlim Y_i = \left\{ \mathbf{y} \in \prod_{X_i} \mid \forall i \leq j, y_i = h_{ij}(y_j) \right\}$$

Example. (i) $\mathbb{F}_p^{\text{alg}} = \varinjlim \mathbb{F}_{p^{i!}}$ is an algebraic closure of \mathbb{F}_p . First, $\mathbb{F}_p^{\text{alg}}$ is algebraic over \mathbb{F}_p . Indeed, for $[x] \in \mathbb{F}_p^{\text{alg}}$, we have $x \in \mathbb{F}_{p^{i!}}$ for some $i \geq 1$. Then $x^{p^{i!}} - x = 0$. Hence

$$[x]^{p^{i!}} - [x] = [x^{p^{i!}} - x] = [0]$$

Further, $\mathbb{F}_p^{\text{alg}}$ is algebraically closed. Any polynomial $h \in \mathbb{F}_p^{\text{alg}}[T]$ has coefficients in $\mathbb{F}_p^{\text{alg}}$, so in particular h arises from an element of $\mathbb{F}_{p^{i!}}[T]$ for some i . This element splits under some $\mathbb{F}_{p^{i!}} \rightarrow \mathbb{F}_{p^{\ell!}}$, so it splits under some $\mathbb{F}_{p^{i!}} \rightarrow \mathbb{F}_{p^{\ell!}}$. Hence it splits under $h_{ij} : \mathbb{F}_{p^{i!}} \rightarrow \mathbb{F}_{p^{j!}}$, so h splits in the direct limit $\mathbb{F}_p^{\text{alg}}$.

- (ii) Define $\mathbb{Z}_p = \varprojlim \mathbb{Z}/p^i\mathbb{Z}$. This is the ring of *p-adic integers*. For example, writing numbers in base $p = 5$,

$$\begin{aligned} 1 &= (1 + 5^1\mathbb{Z}, 1 + 5^2\mathbb{Z}, 1 + 5^3\mathbb{Z}, \dots) \\ -1 &= (4 + 5^1\mathbb{Z}, 44 + 5^2\mathbb{Z}, 444 + 5^3\mathbb{Z}, \dots) \end{aligned}$$

In every position in such an expansion, we ‘expose’ another digit of the p -adic integer to the left.

Definition. Let R be a ring, and let \mathfrak{a} be an ideal of R . Then the *\mathfrak{a} -adic completion* of R is

$$\hat{R} = \varprojlim R/\mathfrak{a}^i$$

where the inverse limit is taken over the directed system (\mathbb{N}, \leq) with morphisms given by the natural projections.

Example. (i) If $R = \mathbb{Z}$ and $\mathfrak{a} = (p)$, then $\hat{R} = \mathbb{Z}_p$.

- (ii) If $R = k[T]$ and $\mathfrak{a} = (T)$, then

$$\hat{R} = \varprojlim k[T]/(T^i) = k[[t]]$$

Definition. Let M be an R -module, and let \mathfrak{a} be an ideal of R . Then the *\mathfrak{a} -adic completion* of M is

$$\hat{M} = \varprojlim M/\mathfrak{a}^i M$$

which is naturally an \hat{R} -module.

We can make the following more general definition.

Definition. Let M be an R -module.

- (i) A *filtration* of M is a sequence $(M_n)_{n \geq 1}$ of submodules of M such that $M_0 = M$ and $M_n \supseteq M_{n+1}$ for each n .
- (ii) The *completion* of M with respect to a filtration $(M_n)_{n \geq 1}$ is $\varprojlim M/M_n$.

Theorem. Let R be a Noetherian ring, and let \mathfrak{a} be an ideal of R . Then,

- (i) the \mathfrak{a} -adic completion \hat{R} is Noetherian;
- (ii) the functor $\hat{R} \otimes_R (-)$ is exact;

II. Commutative Algebra

- (iii) if M is a finitely generated R -module, then the natural map $\hat{R} \otimes_R M \rightarrow \hat{M}$ is an \hat{R} -linear isomorphism.

Thus \mathfrak{a} -adic completion is an exact functor from the category of finitely generated R -modules if R is Noetherian.

6.2. Graded rings and modules

Definition. A *graded ring* is a ring $A = \bigoplus_{n=0}^{\infty} A_n$, where each A_n is an additive subgroup of A , such that $A_m A_n \subseteq A_{m+n}$.

Proposition. A_0 is a subring of A .

Proof. It is clearly a subgroup closed under multiplication, so it suffices to check that it contains the identity element of A . We have

$$1_A = \sum_{i=0}^m y_i; \quad y_i \in A_i$$

For $z_n \in A_n$,

$$z_n = \sum_{i=0}^m y_i z_n$$

z_n is an element of A_n , and each term $y_i z_n$ is an element of A_{n+i} . But since the sum is direct, we must have $z_n = y_0 z_n$, so $z = y_0 z$ for all $z \in A$. Hence $y_0 \in A_0$ is the identity element. \square

Remark. Each A_n is an A_0 -module as $A_0 A_n \subseteq A_n$.

Example. The polynomial ring in finitely many variables has a grading: $k[T_1, \dots, T_m] = \bigoplus_{n=0}^{\infty} A_n$ where A_n is the set of homogeneous polynomials of degree n .

Definition. Let $A = \bigoplus_{n=0}^{\infty} A_n$ be a graded ring. A *graded A -module* is an A -module $M = \bigoplus_{n=0}^{\infty} M_n$ such that $A_m M_n \subseteq M_{m+n}$.

For a graded ring A , we define $A_+ = \bigoplus_{n=1}^{\infty} A_n = \ker(A \rightarrow A_0)$. This is an ideal of A , and $A/A_+ \simeq A_0$.

Proposition. Let $A = \bigoplus_{n=0}^{\infty} A_n$ be a graded ring. Then the following are equivalent:

- (i) A is Noetherian;
- (ii) A_0 is Noetherian and A is finitely generated as an A_0 -algebra.

Proof. Hilbert's basis theorem shows that (ii) implies (i). For the converse, A_0 is Noetherian as it is isomorphic to a quotient of the Noetherian ring A . Note that A_+ is generated by the set of homogeneous elements of positive degree. By (i), A_+ is an ideal in a Noetherian ring

so is generated by a finite set $\{x_1, \dots, x_s\}$, and we can take each x_i to be homogeneous, say, $x_i \in A_{k_i}$ where $k_i > 0$. Let A' be the A_0 -subalgebra of A generated by $\{x_1, \dots, x_s\}$; we want to show $A' = A$. It suffices to show that $A_n \subseteq A'$ for every $n \geq 0$, which we will show by induction. The case $n = 0$ is clear.

Let $n > 0$, and let $y \in A_n$. Note that $y \in A_+$, so

$$y = \sum_{i=1}^s r_i x_i$$

where $r_i \in A$ and $x_i \in A_{k_i}$. Applying the projection to A_n ,

$$y = \sum_{i=1}^s a_i x_i; \quad a_i \in A_{n-k_i}$$

where a_i is the $(n - k_i)$ homogeneous part of r_i . As k_i is positive, the inductive hypothesis implies that each a_i can be written as a polynomial in x_1, \dots, x_s with coefficients in A_0 , giving $y \in A'$ as required. \square

Definition. Let \mathfrak{a} be an ideal of R , and let M be an R -module. Then a filtration $(M_n)_{n \geq 0}$ is an \mathfrak{a} -filtration if $\mathfrak{a}M_n \subseteq M_{n+1}$ for each $n \geq 0$. An \mathfrak{a} -filtration $(M_n)_{n \geq 0}$ is *stable* if there exists $n_0 \geq 0$ such that $\mathfrak{a}M_n = M_{n+1}$ for all $n \geq n_0$.

Example. $(\mathfrak{a}^n M)_{n \geq 0}$ is a stable \mathfrak{a} -filtration of M .

Definition. Let \mathfrak{a} be an ideal in R . The *associated graded ring* is

$$G_{\mathfrak{a}}(R) = \bigoplus_{n \geq 0} \mathfrak{a}^n / \mathfrak{a}^{n+1}; \quad \mathfrak{a}^0 = R$$

This is a ring by defining

$$(x + \mathfrak{a}^{n+1})(y + \mathfrak{a}^{m+1}) = xy + \mathfrak{a}^{n+m+1}; \quad x \in \mathfrak{a}^n, y \in \mathfrak{a}^m$$

Definition. Let M be an R -module, and let \mathfrak{a} be an ideal of R . Let $(M_n)_{n \geq 0}$ be an \mathfrak{a} -filtration of M . The *associated graded module* is

$$G(M) = \bigoplus_{n \geq 0} M_n / M_{n+1}$$

This is a module over $G_{\mathfrak{a}}(R)$ by defining

$$(x + \mathfrak{a}^{n+1})(m + M_{\ell+1}) = xm + M_{n+\ell+1}$$

Proposition. Let R be a Noetherian ring, and let \mathfrak{a} be an ideal of R . Then

- (i) the associated graded ring $G_{\mathfrak{a}}(R)$ is Noetherian; and

II. Commutative Algebra

- (ii) if M is a finitely generated R -module and $(M_n)_{n \geq 0}$ is a stable \mathfrak{a} -filtration of M , then the associated graded module $G(M)$ is a finitely generated $G_{\mathfrak{a}}(R)$ -module.

Proof. Part (i). Let R be Noetherian. Then let $\mathfrak{a} = (x_1, \dots, x_s)$, and write \bar{x}_i for the image of x_i in $\mathfrak{a}/\mathfrak{a}^2$. Note that

$$G_{\mathfrak{a}}(R) = R/\mathfrak{a} \oplus \mathfrak{a}/\mathfrak{a}^2 \oplus \mathfrak{a}^2/\mathfrak{a}^3 \oplus \dots$$

$G_{\mathfrak{a}}(R)$ is generated as an R/\mathfrak{a} -algebra by $\bar{x}_1, \dots, \bar{x}_s$, by taking sums and products. Note that R/\mathfrak{a} is Noetherian, so $G_{\mathfrak{a}}(R)$ is Noetherian by Hilbert's basis theorem.

Part (ii). Let $(M_n)_{n \geq 0}$ be a stable \mathfrak{a} -filtration of M . Then there exists n_0 such that for all $n \geq n_0$, we have $M_{n_0+r} = \mathfrak{a}^r M_{n_0}$. Thus $G(M)$ is generated as a $G_{\mathfrak{a}}(R)$ -module by

$$M_0/M_1 \oplus M_1/M_2 \oplus \dots \oplus M_{n_0}/M_{n_0+1}$$

Each factor M_i/M_{i+1} is a Noetherian R -module, as they are quotients of Noetherian modules, and are annihilated by \mathfrak{a} . In particular, $G(M)$ is a finitely generated $G_{\mathfrak{a}}(R)$ -module, say by x_1, \dots, x_s . \square

Definition. Let M be an R -module. We say that filtrations $(M_n), (M'_n)$ of M are *equivalent* if there exists n_0 such that for all $n \geq 0$, we have $M_{n+n_0} \subseteq M'_n$ and $M'_{n+n_0} \subseteq M_n$.

Lemma. Let \mathfrak{a} be an ideal of R . Let M be an R -module, and let $(M_n)_{n \geq 0}$ be a stable \mathfrak{a} -filtration of M . Then $(M_n)_{n \geq 0}$ is equivalent to $(\mathfrak{a}^n M)_{n \geq 0}$.

Proof. As $(M_n)_{n \geq 0}$ is an \mathfrak{a} -filtration, for all $n \geq 0$, we have

$$M_n \supseteq \mathfrak{a} M_{n-1} \supseteq \mathfrak{a}^2 M_{n-2} \supseteq \dots \supseteq \mathfrak{a}^n M \supseteq \mathfrak{a}^{n+n_0} M$$

For the other direction, as the filtration is stable, there exists n_0 such that for each $n \geq n_0$, we have $\mathfrak{a} M_n = M_{n+1}$. Then $M_{m+n_0} = \mathfrak{a}^n M_{n_0} \subseteq \mathfrak{a}^n M$ as required. \square

6.3. Artin-Rees lemma

Definition. Let \mathfrak{a} be an ideal of R . Let M be an R -module, and let $(M_n)_{n \geq 0}$ be an \mathfrak{a} -filtration of M . Then we define

$$R^* = \bigoplus_{n \geq 0} \mathfrak{a}^n; \quad M^* = \bigoplus_{n \geq 0} M_n$$

Note that R^* is a graded ring, as for $x \in \mathfrak{a}^n, y \in \mathfrak{a}^\ell$, we have $xy \in \mathfrak{a}^{n+\ell}$. As $(M_n)_{n \geq 0}$ is an \mathfrak{a} -filtration, M^* is a graded R^* -module. Indeed, for $x \in \mathfrak{a}^n$ and $m \in M_\ell$, we have $xm \in M_{n+\ell}$ as required.

If R is Noetherian, the ideal \mathfrak{a} is finitely generated, say by x_1, \dots, x_r . Then R^* is generated as an R -algebra by x_1, \dots, x_r by taking sums and products. By Hilbert's basis theorem, R^* is a Noetherian ring.

Lemma. Let R be a Noetherian ring, and let \mathfrak{a} be an ideal of R . Let M be a finitely generated R -module, and let $(M_n)_{n \geq 0}$ be an \mathfrak{a} -filtration of M . Then, the following are equivalent:

- (i) M^* is finitely generated as an R^* -module;
- (ii) the \mathfrak{a} -filtration $(M_n)_{n \geq 0}$ is stable.

Proof. First, note that each M_n is a finitely generated R -module. Indeed, R is a Noetherian ring and M is finitely generated, so M is a Noetherian module, or equivalently, every submodule is finitely generated. Now, consider

$$M_n^* = M_0 \oplus \cdots \oplus M_n \oplus \mathfrak{a}M_n \oplus \mathfrak{a}^2M_n \oplus \cdots$$

This is an R^* -submodule of M^* . Note that $(M_n^*)_{n \geq 0}$ is an ascending chain of R^* -submodules of M^* , and this chain stabilises if and only if the \mathfrak{a} -filtration $(M_n)_{n \geq 0}$ is stable.

(i) *implies* (ii). As R is Noetherian, so is R^* by the discussion above. By assumption, M^* is finitely generated as a module over a Noetherian ring, so it is Noetherian. Hence the ascending chain $(M_n^*)_{n \geq 0}$ stabilises, giving the result.

(ii) *implies* (i). Suppose $(M_n)_{n \geq 0}$ is stable. Then $(M_n^*)_{n \geq 0}$ stabilises at some $n_0 \geq 0$, so

$$M^* = \bigcup_{n \geq 0} M_n^* = M_{n_0}^*$$

Now, note that $M_0 \oplus \cdots \oplus M_{n_0}$ generates $M_{n_0}^*$ as an R^* -module. Each M_n is a finitely generated R -module, so $M_0 \oplus \cdots \oplus M_{n_0}$ is also finitely generated as an R -module. So these generators span $M_{n_0}^* = M^*$ as an R^* -module, as required. \square

Proposition (Artin–Rees). Let R be a Noetherian ring, and let \mathfrak{a} be an ideal of R . Let M be a finitely generated R -module, and let $(M_n)_{n \geq 0}$ be a stable \mathfrak{a} -filtration of M . Then for any submodule $N \leq M$, $(N \cap M_n)_{n \geq 0}$ is a stable \mathfrak{a} -filtration of N .

Thus, stable filtrations pass to submodules.

Proof. First, we show that $(N \cap M_n)_{n \geq 0}$ is indeed an \mathfrak{a} -filtration.

$$\mathfrak{a}(N \cap M_n) \subseteq N \cap \mathfrak{a}M_n \subseteq N \cap M_{n+1}$$

Now, define

$$M^* = \bigoplus_{n \geq 0} M_n; \quad N^* = \bigoplus_{n \geq 0} (N \cap M_n)$$

Note that M^* is an R^* -submodule of N^* . As R is Noetherian, so is R^* . Then as $(M_n)_{n \geq 0}$ is stable, M^* is a finitely generated R^* -module by the previous lemma. Thus M^* is a Noetherian R^* -module. Its submodule N^* is then finitely generated, so $(N \cap M_n)_{n \geq 0}$ is stable. \square

7. Dimension theory

7.1. ???

Definition. Let \mathfrak{p} be a prime ideal of R . The *height* of \mathfrak{p} , denoted $\text{ht}(\mathfrak{p})$, is

$$\text{ht}(\mathfrak{p}) = \sup \{d \mid \mathfrak{p}_0 \subsetneq \mathfrak{p}_1 \subsetneq \cdots \subsetneq \mathfrak{p}_d = \mathfrak{p}; \mathfrak{p}_i \in \text{Spec } R\}$$

The (Krull) *dimension* of R is

$$\dim R = \sup \{\text{ht}(\mathfrak{p}) \mid \mathfrak{p} \in \text{Spec } R\} = \sup \{\text{ht}(\mathfrak{m}) \mid \mathfrak{m} \in \text{mSpec } R\}$$

Remark. The height of a prime ideal \mathfrak{p} is the Krull dimension of the localisation $R_{\mathfrak{p}}$. In particular,

$$\dim R = \sup \{\dim R_{\mathfrak{p}} \mid \mathfrak{p} \in \text{Spec } R\} = \sup \{\dim R_{\mathfrak{m}} \mid \mathfrak{m} \in \text{mSpec } R\}$$

So the problem of computing dimension can be reduced to computing dimension of local rings.

Definition. Let I be a proper ideal of R . Then the *height* of I is

$$\text{ht}(I) = \inf \{\text{ht}(\mathfrak{p}) \mid I \subseteq \mathfrak{p}\}$$

Proposition. Let $A \subseteq B$ be an integral extension of rings. Then,

- (i) $\dim A = \dim B$; and
- (ii) if A, B are integral domains and k -algebras for some field k , they have the same transcendence degree over k .

We prove part (i); the second part is not particularly relevant for this course.

Proof. First, we show that $\dim A \leq \dim B$. Consider a chain of prime ideals $\mathfrak{p}_0 \subsetneq \cdots \subsetneq \mathfrak{p}_d$ in $\text{Spec } A$. By the lying over theorem and the going up theorem, we obtain a chain of prime ideals $\mathfrak{q}_0 \subseteq \cdots \subseteq \mathfrak{q}_d$ in $\text{Spec } B$. As $\mathfrak{p}_i = \mathfrak{q}_i \cap A$ and $\mathfrak{p}_i \neq \mathfrak{p}_{i+1}$, we must have $\mathfrak{q}_i \neq \mathfrak{q}_{i+1}$. So this produces a chain of length d in B , as required.

Now consider a chain $\mathfrak{q}_0 \subsetneq \cdots \subsetneq \mathfrak{q}_d$ in $\text{Spec } B$. Contracting each ideal, we produce a chain $\mathfrak{p}_0 \subseteq \cdots \subseteq \mathfrak{p}_d$ in $\text{Spec } A$. Suppose that \mathfrak{q}_i and \mathfrak{q}_{i+1} contract to the same prime ideal \mathfrak{p}_i in $\text{Spec } A$. Note that $\mathfrak{q}_i \subseteq \mathfrak{q}_{i+1}$, so by incomparability, they must be equal, but this is a contradiction. \square

Remark. If A is a finitely generated k -algebra for some field k , then by Noether normalisation, we obtain a k -algebra embedding $k[T_1, \dots, T_d] \rightarrow A$, and the extension is integral. Thus $\dim A = \dim k[T_1, \dots, T_d]$. One can show that $\dim k[T_1, \dots, T_d] = d$, and hence that the integer d obtained by Noether normalisation is uniquely determined by A and k .

7.2. Hilbert polynomials

Let $A = \bigoplus_{n \geq 0} A_n$ be a Noetherian graded ring, so A_0 is Noetherian and A is finitely generated as an A_0 -algebra. Now let $M = \bigoplus_{n \geq 0} M_n$ be a finitely generated graded A -module. Then each M_n is an A_0 -module.

We claim that M_n is finitely generated as an A_0 -module. Indeed, $M = \text{span}_A \{m_1, \dots, m_t\}$, and the m_i can be taken to be homogeneous, say, $m_i \in M_{r_i}$. Then

$$M_n = \{a_1 m_1 + \dots + a_t m_t \mid a_i \in A_{n-r_i}\}$$

Let x_1, \dots, x_s generate A as an A_0 -algebra, where $x_i \in A_{k_i}, k_i > 0$. Then

$$M_n = \text{span}_{A_0} \left\{ x_1^{e_1} \dots x_t^{e_t} m_i \mid 1 \leq i \leq t, e_i \geq 0, \sum_{i=1}^s k_i e_i = n - r_i \right\}$$

and the right-hand side is a finite set.

We will make the further assumption that A_0 is Artinian. Hence, each M_n is a finitely generated module over a ring that is both Noetherian and Artinian, so each M_n is Noetherian and Artinian as an A_0 -module. Further, each M_n is of finite length $\ell(M_n) < \infty$; it has a *composition series* of finite length. Note that if $A_0 = k$ is a field, then $\ell(M_n) = \dim_k M_n$.

Definition. Let A, M be as above. Then the *Poincaré series* of M is

$$P(M, T) = \sum_{n=0}^{\infty} \ell(M_n) T^n \in \mathbb{Z}[[T]]$$

Theorem (Hilbert–Serre theorem). Let A be generated by x_1, \dots, x_s as an A_0 -module with $x_i \in A_{k_i}$ for $k_i > 0$. The Poincaré series $P(M, T)$ is a rational function of the form

$$\frac{f(T)}{\prod_{i=1}^s (1 - T^{k_i})}; \quad f \in \mathbb{Z}[T]$$

Proof. For the base case $s = 0$, we must have $A = A_0$, so M is a finitely generated A_0 -module, say, $M = \text{span}_{A_0} S$ where S is a finite subset of $M_0 \oplus \dots \oplus M_n$. Thus there exists n_0 such that $M_m = 0$ for all $m > n_0$. In particular, $P(M, T)$ is a polynomial.

For the inductive step, let

$$M = \bigoplus_{n \in \mathbb{Z}} M_n; \quad M_\ell = 0 \text{ if } \ell < 0$$

Let $f : M_n \rightarrow M_{n+k_s}$ be the homomorphism given by multiplication by x_s . We obtain the exact sequence

$$0 \longrightarrow K_n \longrightarrow M_n \xrightarrow{f} M_{n+k_s} \longrightarrow L_{n+k_s} \longrightarrow 0$$

II. Commutative Algebra

where $K_n = \ker f$ and $L_{n+k_s} = \operatorname{coker} f$. Then let $K = \bigoplus_{n \in \mathbb{Z}} K_n$ and $L = \bigoplus_{n \in \mathbb{Z}} L_n$. These are graded A -modules, and K is a submodule of M . Note that K and L are annihilated by x_s . Applying the length function to the exact sequence, we obtain

$$\ell(K_n) - \ell(M_n) + \ell(M_{n+k_s}) - \ell(L_{n+k_s}) = 0$$

Multiplying by T^{n+k_s} ,

$$\ell(M_{n+k_s})T^{n+k_s} - T^{k_s}\ell(M_n)T^n = \ell(L_{n+k_s})T^{n+k_s} - T^{k_s}\ell(K_n)T^n$$

Then, taking the sum over all integers,

$$P(M, T) - T^{k_s}P(M, T) = (1 - T^{k_s})P(M, T) = P(L, T) - T^{k_s}P(K, T)$$

By the inductive hypothesis,

$$(1 - T^{k_s})P(M, T) = \frac{f_1(T)}{\prod_{i=1}^{s-1}(1 - T^{k_i})} + \frac{f_2(T)}{\prod_{i=1}^{s-1}(1 - T^{k_i})}$$

as required. \square

In particular, this rational function is holomorphic almost everywhere, with potentially a pole of some order at 1. Let $d(M)$ be the order of the pole of $P(M, T)$ at $T = 1$. One can show that if $M \neq 0$, then $d(M) \geq 0$.

Example. Let $A = k[T_1, \dots, T_s] = \bigoplus_{n \geq 0} A_n$ where A_n is the set of homogeneous polynomials of degree n . Then A is generated as an $A_0 = k$ -algebra by $\{T_1, \dots, T_s\}$. For this choice of generators, $k_1 = \dots = k_s = 1$. The length of A_n is $\dim_k A_n = \binom{n+s-1}{n}$, which is a polynomial of degree $s-1$ in n over \mathbb{Q} . The Poincaré series of A over itself is

$$P(A, T) = \sum_{n \geq 0} \binom{n+s-1}{n} T^n = \frac{1}{(1-T)^s}$$

Proposition. If $k_1 = \dots = k_s = 1$, then there exists a Hilbert polynomial $HP_M \in \mathbb{Q}[T]$ and $n_0 \geq 0$ such that

$$\ell(M_n) = HP_M(n)$$

for all $n \geq n_0$. In addition, $\deg HP_M = d(M) - 1$ where $d(M)$ is the order of the pole of $P(M, T)$ at $T = 1$.

Proof. Let $d = d(M) \geq 0$. Then,

$$P(M, T) = \sum_{n \geq 0} \ell(M_n)T^n = \frac{f(T)}{(1-T)^d}; \quad f \in \mathbb{Z}[T], f(1) \neq 0$$

Let

$$f = \sum_{k=0}^{\deg f} a_k T^k; \quad a_k \in \mathbb{Z}$$

Note that

$$\frac{1}{(1-T)^d} = \sum_{j=0}^{\infty} \underbrace{\binom{j+d-1}{j}}_{b_j} T^j$$

Thus, for $n \geq \deg f$,

$$\ell(M_n) = \sum_{i=0}^{\deg f} a_i b_{n-i}$$

Note that b_j is a polynomial in j over \mathbb{Q} of degree $d-1$ with leading coefficient $\frac{1}{(d-1)!}$. Then $\ell(M_n)$ is a polynomial p in n over \mathbb{Q} for $n \geq \deg f$. Then $\deg p \leq d-1$, and the coefficient of T^{d-1} in p is

$$\sum_{i=0}^{\deg f} a_i \cdot \frac{1}{(d-1)!} = \frac{f(1)}{(d-1)!} \neq 0$$

so the degree is exactly $d-1$. □

7.3. Dimension theory of local Noetherian rings

Lemma. Let (A, \mathfrak{m}) be a Noetherian local ring. Then

- (i) an ideal \mathfrak{q} of A is \mathfrak{m} -primary if and only if there exists $t \geq 1$ such that $\mathfrak{m}^t \subseteq \mathfrak{q} \subseteq \mathfrak{m}$;
- (ii) if \mathfrak{q} is \mathfrak{m} -primary, then A/\mathfrak{q} is Artinian.

Proof. *Part (i).* Given an ideal \mathfrak{q} between \mathfrak{m}^t and \mathfrak{m} , taking radicals we obtain

$$\sqrt{\mathfrak{m}^t} \subseteq \sqrt{\mathfrak{q}} \subseteq \sqrt{\mathfrak{m}}$$

Hence $\sqrt{\mathfrak{q}} = \mathfrak{m}$ and thus \mathfrak{q} is \mathfrak{m} -primary. Conversely, if \mathfrak{q} is \mathfrak{m} -primary, $(\sqrt{\mathfrak{q}})^t \subseteq \mathfrak{q}$ for some t as A is Noetherian, so $\mathfrak{m}^t \subseteq \mathfrak{q} \subseteq \mathfrak{m}$ as required.

Part (ii). $(A/\mathfrak{q}, \mathfrak{m}/\mathfrak{q})$ is a Noetherian local ring. If $\mathfrak{q} \subseteq \mathfrak{p} \subseteq \mathfrak{m}$, then taking radicals,

$$\mathfrak{m} = \sqrt{\mathfrak{q}} \subseteq \mathfrak{p} \subseteq \mathfrak{m}$$

Hence $\mathfrak{p} = \mathfrak{m}$. In particular, the spectrum of A/\mathfrak{q} is the single ideal $\mathfrak{m}/\mathfrak{q}$. Thus its dimension is zero, and so the quotient is Artinian. □

Theorem (dimension theorem). If A is a Noetherian local ring, then

$$\dim A = \delta(A) = d(G_{\mathfrak{m}}(A))$$

where $\delta(A) = \min \{\delta(\mathfrak{q}) \mid \mathfrak{q} \subseteq A \text{ is } \mathfrak{m}\text{-primary}\}$ and $\delta(\mathfrak{q})$ is the minimal number of generators of \mathfrak{q} , and where the right-hand side is the order of the pole at $T = 1$ of the rational function equal to the Poincaré series

$$\sum_{n \geq 0} \ell(\mathfrak{m}^n / \mathfrak{m}^{n+1}) T^n$$

II. Commutative Algebra

of the associated graded ring.

Proof. We will show that $\delta \geq d \geq \dim \geq \delta$.

Let \mathfrak{q} be an \mathfrak{m} -primary ideal of A , generated by x_1, \dots, x_s where $s = \delta(\mathfrak{q})$. Then

$$G_{\mathfrak{q}}(A) = A/\mathfrak{q} \oplus \mathfrak{q}/\mathfrak{q}^2 \oplus \bigoplus_{n \geq 2} \mathfrak{q}^n/\mathfrak{q}^{n+1}$$

The first factor A/\mathfrak{q} is Artinian, and the images of x_1, \dots, x_s generate $G_{\mathfrak{q}}(A)$ as an A/\mathfrak{q} -algebra, where the x_i are of degree 1. Then $\ell(\mathfrak{q}^n/\mathfrak{q}^{n+1}) < \infty$. From the theorem on Hilbert polynomials, $\ell(\mathfrak{q}^n/\mathfrak{q}^{n+1})$ is a polynomial in n of degree at most $\delta(\mathfrak{q}) - 1$, for sufficiently large n .

Fix some \mathfrak{m} -primary ideal \mathfrak{q}_0 such that $\delta(\mathfrak{q}_0) = \delta(A)$. We consider two special cases: $\mathfrak{q} = \mathfrak{q}_0$ and $\mathfrak{q} = \mathfrak{m}$. For \mathfrak{q}_0 , we have

$$\deg \ell(\mathfrak{q}_0^n/\mathfrak{q}_0^{n+1}) \leq \delta(A) - 1$$

As

$$\ell(A/\mathfrak{q}_0^n) = \sum_{i=0}^{n-1} \ell(\mathfrak{q}_0^i/\mathfrak{q}_0^{i+1})$$

we have

$$\deg \ell(A/\mathfrak{q}_0^n) \leq \delta(A)$$

For \mathfrak{m} ,

$$\deg \ell(\mathfrak{m}^n/\mathfrak{m}^{n+1}) = d(G_{\mathfrak{m}}(A)) - 1$$

and hence

$$\deg \ell(A/\mathfrak{m}^n) = d(G_{\mathfrak{m}}(A))$$

Now, there exists $t \geq 1$ such that $\mathfrak{m}^t \subseteq \mathfrak{q}_0 \subseteq \mathfrak{m}$. Then

$$\ell(A/\mathfrak{m}^n) \leq \ell(A/\mathfrak{q}_0^n) \leq \ell(A/\mathfrak{m}^{tn})$$

But all of these terms are eventually polynomial, and the degrees of the left-hand and right-hand sides are the same, so we must have $\ell(A/\mathfrak{q}_0^n) = \ell(A/\mathfrak{m}^n)$.

Proposition. $\delta(A) \geq d(G_{\mathfrak{m}}(A))$.

Proof.

$$\delta(A) = \delta(\mathfrak{q}_0) \geq \deg \ell(A/\mathfrak{q}_0^n) = \deg \ell(A/\mathfrak{m}^n) = d(G_{\mathfrak{m}}(A))$$

□

Proposition. If $x \in \mathfrak{m}$ is not a zero divisor, then

$$d(G_{(\mathfrak{m}/xA)}(A/xA)) \leq d(G_{\mathfrak{m}}(A)) - 1$$

This proposition allows us to prove results by induction on d .

Proof. We have a local ring $(A/xA, \mathfrak{m}/xA)$. Then

$$d(G_{\mathfrak{m}}(A)) = \deg \ell(A/\mathfrak{m}^n)$$

and

$$d(G_{(\mathfrak{m}/xA)}(A/xA)) = \deg \ell(A/xA/(\mathfrak{m}/xA)^n) = \deg \ell((\mathfrak{m}^n + xA)/xA)$$

We want to show that

$$\deg \ell((\mathfrak{m}^n + xA)/xA) \leq \deg \ell(A/\mathfrak{m}^n) - 1$$

We have the short exact sequence

$$0 \longrightarrow (\mathfrak{m}^n + xA)/\mathfrak{m}^n \longrightarrow A/\mathfrak{m}^n \longrightarrow A/(\mathfrak{m}^n + xA) \longrightarrow 0$$

By the second isomorphism theorem,

$$(\mathfrak{m}^n + xA)/\mathfrak{m}^n \cong xA/(\mathfrak{m}^n \cap xA)$$

Thus, by additivity of length,

$$\ell(A/\mathfrak{m}^n + xA) = \ell(A/\mathfrak{m}^n) - \ell(xA/(\mathfrak{m}^n \cap xA))$$

Note that $(\mathfrak{m}^n)_{n \geq 0}$ is a stable \mathfrak{m} -filtration of A , so $(\mathfrak{m}^n \cap xA)_{n \geq 0}$ is a stable \mathfrak{m} -filtration of the submodule xA by the Artin–Rees lemma. Then $(\mathfrak{m}^n \cap xA)_{n \geq 0}$ is equivalent to the \mathfrak{m} -filtration $(\mathfrak{m}^n xA)_{n \geq 0}$. This equivalence implies that there exists n_0 such that

$$\ell(xA/(\mathfrak{m}^n xA)) \leq \ell(xA/(\mathfrak{m}^{n+n_0} \cap xA)); \quad \ell(xA/(\mathfrak{m}^n \cap xA)) \leq \ell(xA/(\mathfrak{m}^{n+n_0} xA))$$

Hence the polynomials have the same leading term, and so the degree of $\ell(A/\mathfrak{m}^n)$ must decrease. \square

Proposition. $d(G_{\mathfrak{m}}(A)) \geq \dim A$.

Proof. We can prove this by induction using the previous proposition. \square

Proposition. $\dim A \leq \delta(A)$. That is, there exists an \mathfrak{m} -primary ideal \mathfrak{q} that is generated by $d = \dim A$ elements.

II. Commutative Algebra

Proof. As \mathfrak{m} is the unique maximal ideal, we must have $\text{ht}(\mathfrak{m}) = d$. Also, $\text{ht}(\mathfrak{p}) < d$ for any prime $\mathfrak{p} \neq \mathfrak{m}$. We will form an ideal \mathfrak{q} generated by d elements such that $\text{ht}(\mathfrak{q}) \geq d$. This suffices, as then for every minimal prime ideal \mathfrak{p} of \mathfrak{q} , we must have $\text{ht}(\mathfrak{p}) = d$ and thus $\mathfrak{p} = \mathfrak{m}$, giving $\sqrt{\mathfrak{q}} = \mathfrak{m}$ so \mathfrak{p} is \mathfrak{m} -primary as required.

Construct x_1, \dots, x_d inductively such that $\text{ht}(\mathfrak{q}_i) \geq i$ where $\mathfrak{q}_i = (x_1, \dots, x_i)$. For the base case, we take $\mathfrak{q}_0 = (0)$. For the inductive step, we assume that $\mathfrak{q}_{i-1} = (x_1, \dots, x_{i-1})$ has already been constructed, with $i-1 < d$ and $\text{ht}(\mathfrak{q}_{i-1}) \geq i-1$. We claim that there are only finitely many prime ideals $\mathfrak{p}_1, \dots, \mathfrak{p}_t$ that contain \mathfrak{q}_{i-1} and have height exactly $i-1$. Indeed, $\text{ht}(\mathfrak{q}_{i-1}) \geq i-1$, so each \mathfrak{p}_j is a minimal prime ideal of \mathfrak{q}_{i-1} , and in a Noetherian ring, every ideal has only finitely many minimal primes. We know that $i-1 < d = \text{ht}(\mathfrak{m})$, so $\mathfrak{m} \not\subseteq \mathfrak{p}_j$ for all j . Therefore, $\mathfrak{m} \not\subseteq \bigcup_j \mathfrak{p}_j$ by the prime avoidance lemma. Take $x_i \in \mathfrak{m} \setminus \bigcup_j \mathfrak{p}_j$, and define $\mathfrak{q}_i = (x_1, \dots, x_{i-1}, x_i)$. Now, if \mathfrak{p} is a prime ideal that contains \mathfrak{q}_i , as $\mathfrak{p} \notin \{\mathfrak{p}_1, \dots, \mathfrak{p}_t\}$, we must have $\text{ht}(\mathfrak{p}) \geq i$ as required. \square

\square

Corollary (Krull's height theorem). Let A be a Noetherian ring, and let $\mathfrak{a} = (x_1, \dots, x_r)$ be an ideal of A . Let \mathfrak{p} be a minimal prime ideal of \mathfrak{a} . Then $\text{ht}(\mathfrak{p}) \leq r$.

Proof. First, we claim that $\sqrt{\mathfrak{a}A_{\mathfrak{p}}}$ is the unique maximal ideal $\mathfrak{p}A_{\mathfrak{p}}$ of the localisation. Indeed, suppose $\mathfrak{a}A_{\mathfrak{p}} \subseteq \mathfrak{n} \in \text{Spec } A_{\mathfrak{p}}$. Contracting, we obtain $\mathfrak{a} \subseteq (\mathfrak{a}A_{\mathfrak{p}})^c \subseteq \mathfrak{n}^c \subseteq \mathfrak{p}$. But as \mathfrak{p} is a minimal prime ideal of \mathfrak{a} , we must have $\mathfrak{n}^c = \mathfrak{p}$. Extending, $\mathfrak{n}^{ce} = \mathfrak{p}^e = \mathfrak{p}A_{\mathfrak{p}}$, but $\mathfrak{n}^{ce} = \mathfrak{n}$ as required. Hence, $\sqrt{\mathfrak{a}A_{\mathfrak{p}}}$ is the intersection of the primes containing it, which is just $\mathfrak{p}A_{\mathfrak{p}}$.

As the radical is maximal, the ideal $\mathfrak{a}A_{\mathfrak{p}}$ is $\mathfrak{p}A_{\mathfrak{p}}$ -primary. Note that $\mathfrak{a}A_{\mathfrak{p}} = \left(\frac{x_1}{1}, \dots, \frac{x_r}{1}\right)$, so by applying the dimension theorem,

$$\text{ht}(\mathfrak{p}) = \dim A_{\mathfrak{p}} = \delta(A_{\mathfrak{p}}) \leq \delta(\mathfrak{a}A_{\mathfrak{p}}) \leq r$$

\square

III. Algebraic Geometry

Lectured in Michaelmas 2023 by DR. D. RANGANATHAN

(Course description goes here.)

Contents

1. Introduction	166
1.1. Course description	166
1.2. Motivation from moduli theory	166
1.3. Motivation from the Weil conjectures	167
1.4. Summary of classical algebraic geometry	168
1.5. Limitations of classical algebraic geometry	169
1.6. Spectrum of a ring	169
1.7. Distinguished opens and localisation	171
2. Sheaves	173
2.1. Presheaves	173
2.2. Sheaves	174
2.3. Stalks	174
2.4. Sheafification	176
2.5. Kernels and cokernels	176
2.6. Moving between spaces	178
3. Schemes	180
3.1. Localisation	180
3.2. Sheaves on a base	180
3.3. The structure sheaf	181
3.4. Definitions and examples	182
3.5. Gluing sheaves	183
3.6. Gluing schemes	184
3.7. The Proj construction	185
4. Morphisms	188
4.1. Morphisms of ringed spaces	188
4.2. Morphisms of schemes	188
4.3. Immersions	190
4.4. Fibre products	190
4.5. Schemes over a base	192
4.6. Separatedness	192
4.7. Properness	195
4.8. Valuative criteria	197
5. Modules over the structure sheaf	199
5.1. Definitions	199
5.2. Quasi-coherence	200
5.3. Coherent sheaves on projective schemes	202

6.	Divisors	204
6.1.	Height and dimension	204
6.2.	Weil divisors	205
6.3.	Cartier divisors	206
7.	Sheaf cohomology	209
7.1.	Introduction and properties	209
7.2.	Čech cohomology	209
7.3.	Choice of cover	213
7.4.	Further topics in cohomology	214

1. Introduction

1.1. Course description

The course consists of four parts.

- (i) The theory of sheaves on topological spaces.
- (ii) The definitions of schemes and morphisms between them.
- (iii) Properties of schemes, such as the algebraic geometry analogues of compactness and other similar properties.
- (iv) Rapid introduction to the cohomology of sheaves.

1.2. Motivation from moduli theory

In moduli theory, we study families of varieties instead of one at a time. In the extreme, we study all varieties of a given ‘type’ simultaneously. For now, let

$$\mathbb{P}^n = \mathbb{P}_{\mathbb{C}}^n = \mathbb{C}^{n+1} \setminus \{\mathbf{0}\} / \sim$$

where $\mathbf{x} \sim \lambda \mathbf{x}$ for nonzero λ , \mathbf{x} . A variety is the vanishing locus $\mathbb{V}(S)$ of a set S of homogeneous polynomials in $n + 1$ variables. These are subsets of \mathbb{P}^n . We present some examples of moduli.

Example. The set of all lines in \mathbb{P}^2 . A line in \mathbb{P}^2 is given by

$$\{aX_0 + bX_1 + cX_2 = 0\}$$

where not all of a, b, c are zero. The set of all lines in \mathbb{P}^2 are given by triples (a, b, c) . Note that $(\lambda a, \lambda b, \lambda c)$ gives the same line as (a, b, c) , so really lines in \mathbb{P}^2 correspond exactly to points in \mathbb{P}^2 . We call the set of all lines in \mathbb{P}^2 the dual space $\mathbb{P}_{\text{dual}}^2$. This property is known as projective duality.

The same logic applies to the set of degree d hypersurfaces in \mathbb{P}^n ; this space corresponds directly to

$$\mathbb{P}^{\binom{n+d}{d}-1}$$

There is an unfortunate consequence of this method of study. Some polynomials are of the form $f = f_1^2 f_2$ for some non-constant f_1 , but then $\mathbb{V}(f) = \mathbb{V}(f_1 f_2)$. For example, $(X_0 + X_1 + X_2)^2 \subseteq \mathbb{P}^2$ is a line not a conic. In particular, the limit of a sequence of conics may not be a conic. The solution is to take the set

$$U_d \subseteq \mathbb{P}^{\binom{n+d}{d}-1}$$

in which $[f] \in U_d$ has no repeated factors. But then, U_d is ‘not compact’, as some points have been removed.

We will now describe the impact of scheme theory on this situation. Fix some \mathbb{P}^n , and we will produce a ‘space’

$$\mathrm{Var}(\mathbb{P}^n) \subsetneq \mathrm{Hilb}(\mathbb{P}^n)$$

The set $\mathrm{Var}(\mathbb{P}^n)$ bijects onto the set of varieties of \mathbb{P}^n . The set $\mathrm{Hilb}(\mathbb{P}^n)$ bijects onto the set of subschemes of \mathbb{P}^n , and is compact in the Euclidean topology. In particular, limits of varieties need not be varieties, but limits of schemes are always schemes. One consequence is that in scheme theory,

$$\mathbb{V}(X_0 + X_1 + X_2), \quad \mathbb{V}((X_0 + X_1 + X_2)^2)$$

are not isomorphic as schemes in \mathbb{P}^2 .

1.3. Motivation from the Weil conjectures

Fix some homogeneous polynomial $f \in \mathbb{Z}[X_0, \dots, X_{n+1}]$. First, consider

$$X = \mathbb{V}(f) \subseteq \mathbb{P}_{\mathbb{C}}^{n+1}$$

and assume that X is smooth. As X is a compact topological space, we can find its Betti numbers $b_0(X), \dots, b_{2n}(X)$, where

$$b_i(X) = \mathrm{rank} H_i(X; \mathbb{Z})$$

In particular, we can find its Euler characteristic.

$$\chi(X) = \sum (-1)^i b_i(X)$$

Second, fix a prime p and let N_m be the number of solutions of f over \mathbb{F}_{p^m} . Define the Weil zeta function

$$\zeta(X; t) = \exp\left(\sum_m \frac{N_m}{m} \cdot t^m\right)$$

One of the Weil conjectures states the following.

Theorem (Grothendieck). (i) $\zeta(X; t)$ is a rational function in t , so

$$\zeta(X; t) = \frac{P_X(t)}{Q_X(t)}$$

(ii) Further, $\zeta(X; t)$ can be written as a ratio of the form

$$\frac{P_0(t)P_2(t) \dots P_{2n}(t)}{P_1(t)P_3(t) \dots P_{2n-1}(t)}$$

where

$$\deg P_i(t) = b_i(X)$$

The proof relies fundamentally on scheme theory: we need a space \mathcal{X} that interpolates between the algebraic closure $\overline{\mathbb{F}_p}$ and \mathbb{C} .

III. Algebraic Geometry

1.4. Summary of classical algebraic geometry

Let $k = \bar{k}$ be an algebraically closed field. The notation $\mathbb{A}_k^n = \mathbb{A}^n$ denotes affine space of dimension n over the field k . As a set, this is equal to k^n . An *affine variety* is a subset $V \subseteq \mathbb{A}^n$ of the form

$$V = \mathbb{V}(S) = \{x \in \mathbb{A}^n \mid \forall f \in S, f(x) = 0\}$$

where $S \subseteq k[X_1, \dots, X_n]$. Note that $\mathbb{V}(S) = \mathbb{V}(I(S))$, where $I(S)$ is the ideal generated by S . By Hilbert's basis theorem, or equivalently the fact that $k[\mathbf{X}]$ is Noetherian, $\mathbb{V}(S)$ is the vanishing locus of a finite set (even a finite subset of S). In fact, $\mathbb{V}(I) = \mathbb{V}(\sqrt{I})$ where

$$\sqrt{I} = \{f \in k[\mathbf{X}] \mid \exists n > 0, f^n \in I\}$$

Note that \sqrt{I} is an ideal, and is called the *radical ideal* of I . For example, in $k[X]$, if $I = (X^2)$ then $\sqrt{I} = (X)$. Notice that an affine variety is a subset of \mathbb{A}^n for some n , so we have really defined varieties with a chosen n ; we have not defined an abstract variety.

A *morphism* between varieties $V \subseteq \mathbb{A}^n$ and $W \subseteq \mathbb{A}^m$ is a set-theoretic map $\varphi : V \rightarrow W$ such that if $\varphi(f_1, \dots, f_m)$, each f_i is the restriction of a polynomial in $\{X_1, \dots, X_n\}$ to V . Note that the polynomials f_i are not part of the definition; a given set-theoretic map may be represented by multiple polynomials. This indicates that the ambient spaces $\mathbb{A}^n, \mathbb{A}^m$ are not relevant to this definition. Isomorphisms are those morphisms with two-sided inverses.

The basic correspondence of the theory of algebraic varieties is

$$\frac{\{\text{affine varieties over } k\}}{\text{isomorphism}} \leftrightarrow \{\text{finitely generated } k\text{-algebras without nilpotent elements}\}$$

We explain each direction of the correspondence. Given a variety V representing an isomorphism class of affine varieties over k , we can write V as the vanishing locus of some radical ideal $I \subseteq k[X_1, \dots, X_n]$. We can then produce the finitely generated k -algebra given by the quotient

$$k[X_1, \dots, X_n]_I$$

This is nilpotent-free as I is radical. In reverse, if A is a finitely generated nilpotent-free k -algebra, then by definition we can write A as

$$k[Y_1, \dots, Y_m]_J$$

where J is radical, or at least up to isomorphism. Then we can produce the affine variety $V = \mathbb{V}(J)$. One must show that the choices we made in the above explanation do not matter.

Note that, for example, $k[X]_{(X^2)}$ has a nilpotent element X . The theory of schemes explains the relevance of these nilpotent elements, but the theory of varieties 'ignores' nilpotent elements.

The algebra associated to V is classically denoted $k[V]$, and is called the *coordinate ring* of V . There is a bijection between morphisms $V \rightarrow W$ and k -algebra homomorphisms $k[W] \rightarrow k[V]$. In category theoretic terminology, the category whose objects are affine varieties up to isomorphism is equivalent to the category of finitely generated k -algebras up to isomorphism.

Let $V = \mathbb{V}(I) \subseteq \mathbb{A}^n$ be a variety with coordinate ring $k[V]$. The *Zariski topology* on V is defined such that the closed sets are exactly those sets of the form $\mathbb{V}(S)$ where $S \subseteq k[V]$. One can show that this really induces a topology. If $V \cong W$, then V and W are homeomorphic as topological spaces.

Let V be a variety and $k[V]$ be its coordinate ring. For all points $P \in V$, we can produce a homomorphism $\text{ev}_P : k[V] \rightarrow k$ mapping f to $f(P)$; one can check that this is well-defined. Note that ev_P is surjective by considering the constant functions. Thus the kernel of ev_P is a maximal ideal \mathfrak{m}_P . We thus obtain

$$\{\text{points of } V\} \rightarrow \{\text{maximal ideals in } k[V]\}$$

Hilbert's *Nullstellensatz* states, among other things, that this is a bijection.

1.5. Limitations of classical algebraic geometry

The description of varieties given above always retains information about its ambient affine space, so we cannot define an abstract variety. Similarly to manifolds which locally look like vector spaces, we want to consider ‘spaces’ that locally look like affine varieties. For example, projective space does not live inside an affine space.

Let $I = (X^2 + Y^2 + 1) \subseteq \mathbb{R}[X, Y]$. Observe that $\mathbb{V}(I)$ is empty in \mathbb{R}^2 , but I is prime and hence radical. Hence the Nullstellensatz fails in this case. It is then natural to ask on which topological space $\mathbb{R}[X, Y]_{(X^2 + Y^2 + 1)}$ is naturally the set of functions. Similar questions can be asked about \mathbb{Z} or $\mathbb{Z}[X]$, for example.

Consider $C = \mathbb{V}(Y - X^2) \subseteq \mathbb{A}_k^2$ and $D = \mathbb{V}(Y)$. Then $C \cap D = \mathbb{V}(X^2, Y) = \mathbb{V}(X, Y) = \{(0, 0)\}$. If $D_\delta = \mathbb{V}(Y + \delta)$ for $\delta \in k$, $C \cap D_\delta$ is two points unless $\delta = 0$. This breaks a continuity property. Therefore, the intersection of two affine varieties is not naturally an affine variety.

1.6. Spectrum of a ring

Let A be a commutative unital ring.

Definition. The *Zariski spectrum* of A is $\text{Spec } A = \{\mathfrak{p} \trianglelefteq A \text{ prime}\}$.

Remark. Given a ring homomorphism $\varphi : A \rightarrow B$, we have an induced map of sets $\varphi^{-1} : \text{Spec } B \rightarrow \text{Spec } A$ given by $\mathfrak{q} \mapsto \varphi^{-1}(\mathfrak{q})$, as the preimage of a prime ideal is always prime. Note, however, that this property would fail if we only considered maximal ideals, because the preimage of a maximal ideal need not be maximal.

III. Algebraic Geometry

Given $f \in A$ and a point $\mathfrak{p} \in \operatorname{Spec} A$, we have an induced $\bar{f} \in A/\mathfrak{p}$ obtained by taking the quotient. We can think of this operation as ‘evaluating’ an $f \in A$ at a point $\mathfrak{p} \in \operatorname{Spec} A$, with the caveat that the codomain of this evaluation depends on \mathfrak{p} .

Example. (i) Let $A = \mathbb{Z}$. Then $\operatorname{Spec} A = \operatorname{Spec} \mathbb{Z}$ is the set $\{(p) \mid p \text{ prime}\} \cup \{(0)\}$. Consider an element of \mathbb{Z} , say, 132. Given a prime p , we can ‘evaluate it at p ’, giving $132 \bmod p \in \mathbb{Z}/p\mathbb{Z}$. Thus $\operatorname{Spec} \mathbb{Z}$ is a space, 132 is a function on $\operatorname{Spec} \mathbb{Z}$, and $132 \bmod p$ is the value of this function at p .

(ii) Let $A = \mathbb{R}[X]$. Then $\operatorname{Spec} A$ is naturally \mathbb{C} modulo complex conjugation, together with the zero ideal.

(iii) If $A = \mathbb{C}[X]$, then $\operatorname{Spec} A$ is naturally \mathbb{C} , together with the zero ideal.

Definition. Let $f \in A$. Then we define

$$\mathbb{V}(f) = \{\mathfrak{p} \in \operatorname{Spec} A \mid f = 0 \bmod \mathfrak{p}, \text{ or equivalently, } f \in \mathfrak{p}\} \subseteq \operatorname{Spec} A$$

Similarly, for $J \trianglelefteq A$ an ideal,

$$\mathbb{V}(J) = \{\mathfrak{p} \in \operatorname{Spec} A \mid \forall f \in J, f \in \mathfrak{p}\} = \{\mathfrak{p} \in \operatorname{Spec} A \mid J \subseteq \mathfrak{p}\}$$

Proposition. The sets $\mathbb{V}(J) \subseteq \operatorname{Spec} A$ ranging over all ideals $J \trianglelefteq A$ form the closed sets of a topology.

This topology is called the *Zariski topology* on A .

Proof. We have $\emptyset = \mathbb{V}(1)$ and $\operatorname{Spec} A = \mathbb{V}(0)$, so they are closed. Note that

$$\mathbb{V}\left(\sum_{\alpha} I_{\alpha}\right) = \bigcap_{\alpha} \mathbb{V}(I_{\alpha})$$

It remains to show $\mathbb{V}(I_1) \cup \mathbb{V}(I_2) = \mathbb{V}(I_1 \cap I_2)$. The containment $\mathbb{V}(I_1) \cup \mathbb{V}(I_2) \subseteq \mathbb{V}(I_1 \cap I_2)$ is clear. Conversely, note $I_1 I_2 \subseteq I_1 \cap I_2$. If $I_1 \cap I_2 \subseteq \mathfrak{p}$, then by primality of \mathfrak{p} , either $I_1 \subseteq \mathfrak{p}$ or $I_2 \subseteq \mathfrak{p}$. \square

Example. Consider $\operatorname{Spec} \mathbb{C}[x, y]$. The point $(0) \in \operatorname{Spec} \mathbb{C}[x, y]$ is dense in the Zariski topology, so $\overline{\{(0)\}} = \operatorname{Spec} \mathbb{C}[x, y]$. This is because all prime ideals in integral domains contain the zero ideal. (0) is sometimes called the *generic point*.

Consider the prime ideal $(y^2 - x^3)$, and consider a maximal ideal $\mathfrak{m}_{a,b} = (x - a, y - b)$ corresponding to the point (a, b) . Then one can show that

$$\mathfrak{m}_{a,b} \in \overline{\{(y^2 - x^3)\}} \iff b^2 = a^3$$

In general, points are not closed.

1.7. Distinguished opens and localisation

Definition. Let $f \in A$. Define the *distinguished open* corresponding to f to be

$$U_f = \operatorname{Spec} A \setminus \mathbb{V}(f)$$

Example. (i) Let $A = \mathbb{C}[x]$, and recall that $\operatorname{Spec} A$ is $\mathbb{C} \cup \{(0)\}$, where the complex number a represents the maximal ideal $(x - a)$. Let $f = x$ and consider

$$\mathbb{V}(x) = \{\mathfrak{p} \mid x \in \mathfrak{p}\} = \{(x)\}$$

Hence $U_x = \operatorname{Spec} A \setminus \{(x)\}$, which is $\operatorname{Spec} A$ without the complex number 0.

(ii) More generally, suppose we fix $a_1, \dots, a_r \in \mathbb{C}$. Then

$$U = \operatorname{Spec} A \setminus \{(x - a_i)\}_{i=1}^r = U_f; \quad f = \prod_{i=1}^r (x - a_i)$$

Lemma. The distinguished opens U_f , taken over all $f \in A$, form a basis for the Zariski topology on $\operatorname{Spec} A$; that is, every open set in $\operatorname{Spec} A$ is a union of some collection of the U_f .

Proof. Let $U = \operatorname{Spec} A \setminus \mathbb{V}(J)$ be an open set. Then

$$\mathbb{V}(J) = \mathbb{V}\left(\sum_{f \in J} (f)\right) = \bigcap_{f \in J} \mathbb{V}(f)$$

So

$$U = \bigcup_{f \in J} U_f$$

□

Definition. Let $f \in A$. The *localisation* of A at f is

$$A_f = A[x]_{(xf - 1)}$$

Informally, we adjoin $\frac{1}{f}$ to A .

Lemma. The distinguished open $U_f \subseteq \operatorname{Spec} A$ is naturally homeomorphic to $\operatorname{Spec} A_f$ via the ring homomorphism $j : A \rightarrow A_f$.

Proof. We will exhibit a bijection between the prime ideals in A_f and the prime ideals in A that do not contain f , producing a homeomorphism as required. Given $\mathfrak{q} \subseteq A_f$ prime, its contraction $j^{-1}(\mathfrak{q})$ is a prime ideal in A .

III. Algebraic Geometry

Now suppose $\mathfrak{p} \subseteq A$ is a prime ideal, and let $\mathfrak{p}_f = j(\mathfrak{p}) \cdot A_f$. We show that $j(\mathfrak{p}) \cdot A_f$ is a prime ideal if and only if $f \notin \mathfrak{p}$, giving the result. If $f \in \mathfrak{p}$, then the unit f lies in \mathfrak{p}_f . Thus $\mathfrak{p}_f = (1)$, so is not prime. If $f \notin \mathfrak{p}$, observe that

$$A_f/\mathfrak{p}_f \cong (A/\mathfrak{p})_{\bar{f}}; \quad \bar{f} = f + \mathfrak{p}$$

But then,

$$(A/\mathfrak{p})_{\bar{f}} \subseteq FF(A/\mathfrak{p})$$

Since \mathfrak{p} is prime, A/\mathfrak{p} is an integral domain, so its fraction field is well-defined. So \mathfrak{p}_f is a prime ideal. One can then check that our two constructions are inverse to each other, providing a bijection between prime ideals as required. \square

Remark. (i) $U_f \cap U_g = U_{fg}$. Indeed, if $\mathfrak{p} \in U_{fg}$, then $fg \notin \mathfrak{p}$, so clearly neither f nor g can lie in \mathfrak{p} ; conversely, if $\mathfrak{p} \in U_f \cap U_g$, then $f \notin \mathfrak{p}$ and $g \notin \mathfrak{p}$, so by primality, $fg \notin \mathfrak{p}$.

(ii) The distinguished opens U_f do not uniquely define an element $f \in A$. For instance, one can easily show that $U_{f^n} = U_f$ for all $n \geq 1$, using the properties of prime ideals.

(iii) In line with (ii), the localisations A_f and A_{f^n} are homeomorphic in a natural way. If

$$A_f = A[x]/(xf - 1); \quad A_{f^n} = A[y]/(yf^n - 1)$$

then consider the inverse A -algebra homomorphisms given by

$$x \mapsto f^{n-1}y; \quad y \mapsto x^n$$

Informally, we map $\frac{1}{f}$ to $f^{n-1}\frac{1}{f^n}$, and $\frac{1}{f^n}$ to $\left(\frac{1}{f}\right)^n$.

(iv) The containment $U_f \subseteq U_g$ holds if and only if f^n is a multiple of g for some $n \geq 1$. First, if f^n is a multiple of g , then the claim holds by (i). Now suppose $U_f \subseteq U_g$, so $\mathbb{V}(f) \supseteq \mathbb{V}(g)$. Hence, all prime ideals that contain g also contain f . But since

$$\sqrt{I} = \bigcap_{\mathfrak{p} \text{ prime} \supseteq I} \mathfrak{p}$$

we must have

$$\sqrt{(f)} \supseteq \sqrt{(g)}$$

giving the result.

Remark. For a fixed ring A , we have made an assignment

$$\{\text{distinguished opens in } \text{Spec } A\} \rightarrow \mathbf{Rng}$$

given by $U_f \mapsto A_f$, where \mathbf{Rng} denotes the class of rings. This association is functorial: if $U_{f_1} \subseteq U_{f_2}$, there is a natural map $A_{f_2} \rightarrow A_{f_1}$, which should be viewed as the restriction map from functions defined on U_{f_2} to those defined on U_{f_1} . This produces a *sheaf*; we now explore these in more generality.

2. Sheaves

2.1. Presheaves

Definition. Let X be a topological space. Let $\text{Open } X$ be the set of open sets on X , and \mathbf{AbGp} be the class of abelian groups. A *presheaf* \mathcal{F} on X of abelian groups is an association

$$\text{Open } X \rightarrow \mathbf{AbGp}$$

and for open sets $U \subseteq V$, a *restriction map*

$$\text{res}_U^V : \mathcal{F}(V) \rightarrow \mathcal{F}(U)$$

such that

$$\text{res}_U^U = \text{id}; \quad \text{res}_U^V \circ \text{res}_V^W = \text{res}_U^W$$

Example. For any topological space X , the presheaf of real-valued continuous functions on X is defined by

$$\mathcal{F}(U) = \{f : U \rightarrow \mathbb{R} \mid f \text{ continuous}\}$$

and

$$\text{res}_U^V(f) = f|_U$$

One can also define presheaves of rings, sets, or other objects by simply replacing the words ‘abelian groups’ in the definition.

Definition. A *morphism* φ of presheaves \mathcal{F}, \mathcal{G} on X is, for each open set U in X , a homomorphism

$$\varphi(U) : \mathcal{F}(U) \rightarrow \mathcal{G}(U)$$

such that

$$\begin{array}{ccc} \mathcal{F}U & \xrightarrow{\text{res}_U^V} & \mathcal{F}V \\ \varphi(U) \downarrow & & \downarrow \varphi(V) \\ \mathcal{G}U & \xrightarrow{\text{res}_U^V} & \mathcal{G}V \end{array}$$

commutes.

Remark. A presheaf on a topological space X is just a functor $(\text{Open } X)^{\text{op}} \rightarrow \mathbf{AbGp}$, where \mathbf{AbGp} is the category of abelian groups, and $\text{Open } X$ is the category where the objects are the open sets in X , and there is a morphism $U \rightarrow V$ if and only if $U \subseteq V$. A morphism of presheaves is just a natural transformation between two such functors. Replacing \mathbf{AbGp} with an arbitrary category \mathcal{C} , we can define presheaves on X of objects in \mathcal{C} .

Definition. A morphism $\varphi : \mathcal{F} \rightarrow \mathcal{G}$ of presheaves is *injective* (respectively *surjective*) if $\varphi(U) : \mathcal{F}(U) \rightarrow \mathcal{G}(U)$ is injective (respectively surjective) for all open sets U of X .

III. Algebraic Geometry

2.2. Sheaves

Definition. A *sheaf* on X is a presheaf \mathcal{F} on X such that

- (i) if $U \subseteq X$ is open and $\{U_i\}$ is an open cover of U , then for $s \in \mathcal{F}(U)$, if $\text{res}_{U_i}^U s = 0$ for all i , then $s = 0$; and
- (ii) if $U, \{U_i\}$ are as in (i), given $s_i \in \mathcal{F}(U_i)$ such that $\text{res}_{U_i \cap U_j}^{U_i} s_i = \text{res}_{U_i \cap U_j}^{U_j} s_j$ for all i, j , then there exists $s \in \mathcal{F}(U)$ such that $\text{res}_{U_i}^U s = s_i$.

Remark. These two axioms imply that $\mathcal{F}(\emptyset) = 0$.

A morphism of sheaves is a morphism of the underlying presheaves.

Example. (i) Let X be a topological space. Then the presheaf \mathcal{F} given by

$$\mathcal{F}(U) = \{f : U \rightarrow \mathbb{R} \mid f \text{ continuous}\}$$

is a sheaf.

- (ii) Let $X = \mathbb{C}$ with the usual Euclidean topology, and let

$$\mathcal{F}(U) = \{f : U \rightarrow \mathbb{C} \mid f \text{ bounded and holomorphic}\}$$

Then \mathcal{F} is not a sheaf, because the functions id_U on bounded open sets U do not glue together to a bounded holomorphic function on all of \mathbb{C} . This is a failure of locality in our definition of \mathcal{F} ; whether f is bounded is a global condition.

- (iii) Let G be a group and set $\mathcal{F}(U) = G$, giving the constant presheaf. This is not in general a sheaf. For example, if U_1, U_2 are disjoint, then $\mathcal{F}(U_1 \cup U_2) \simeq G \times G$. Instead, we can give G the discrete topology, and define

$$\mathcal{F}(U) = \{f : U \rightarrow G \mid f \text{ continuous}\} = \{f : U \rightarrow G \mid f \text{ locally constant}\}$$

This is now a sheaf, called the constant sheaf.

- (iv) Let V be an irreducible variety over k . Let

$$\mathcal{O}_V(U) = \{f \in k(V) \mid \forall p \in U, f \text{ regular at } p\}$$

where a function f is regular at p precisely if it can be represented as a quotient $\frac{g}{h}$ in a neighbourhood of p on which h is nonzero. This is called the *structure sheaf* of V ; it is a sheaf since regularity is a local condition.

2.3. Stalks

Definition. Let \mathcal{F} be a presheaf. A *section* of \mathcal{F} over U is an element $s \in \mathcal{F}(U)$.

Definition. Let $p \in X$, and \mathcal{F} a presheaf on X . Then the *stalk* of \mathcal{F} at p is

$$\mathcal{F}_p = \{(U, s) \mid s \in \mathcal{F}(U), p \in U\} / \sim$$

where

$$(U, s) \sim (V, s') \iff \exists W \subseteq U \cap V \text{ open with } p \in W \text{ such that } \text{res}_W^U s = \text{res}_W^V s'$$

Elements of \mathcal{F}_p are called *germs*.

Example. Let \mathbb{A}^1 be the affine line, and let $\mathcal{O}_{\mathbb{A}^1}$ be the sheaf of regular functions. Its stalk at 0 is

$$\mathcal{O}_{\mathbb{A}^1, 0} = \left\{ \frac{f(t)}{g(t)} \mid g(0) \neq 0 \right\} = k[t]_{(t)}$$

Proposition. Let $f : \mathcal{F} \rightarrow \mathcal{G}$ be a morphism of sheaves on X . Suppose that for all $p \in X$, the induced map $f_p : \mathcal{F}_p \rightarrow \mathcal{G}_p$ given by

$$f_p((U, s)) = (U, f_U(s))$$

is an isomorphism. Then f is an isomorphism.

Proof. We will show that $f_U : \mathcal{F}(U) \rightarrow \mathcal{G}(U)$ are isomorphisms for each U , then define f^{-1} by $(f^{-1})_U = (f_U)^{-1}$.

To show f_U is injective, consider $s \in \mathcal{F}(U)$ with $f_U(s) = 0$. Since f_p is injective, $(U, s) = 0$ in \mathcal{F}_p for every point $p \in U$. Thus for each $p \in U$, there exists an open neighbourhood $U_p \subseteq U$ such that $\text{res}_{U_p}^U s = 0$. The sets $\{U_p \mid p \in U\}$ cover U , so as \mathcal{F} is a sheaf, $s = 0$.

To show f_U is surjective, let $t \in \mathcal{G}(U)$. For each $p \in U$, there is an element $(U_p, s_p) \in \mathcal{F}_p$ such that $f_p((U_p, s_p)) = (U, t) \in \mathcal{G}_p$. By shrinking U_p if necessary, we can assume $f_{U_p}(s_p) = \text{res}_{U_p}^U t$. For points $p, p' \in U$,

$$f_{U_p \cap U_{p'}}(\text{res}_{U_p \cap U_{p'}}^{U_p} s - \text{res}_{U_p \cap U_{p'}}^{U_{p'}} s') = \text{res}_{U_p \cap U_{p'}}^U t - \text{res}_{U_p \cap U_{p'}}^U t = 0$$

Thus

$$\text{res}_{U_p \cap U_{p'}}^{U_p} s - \text{res}_{U_p \cap U_{p'}}^{U_{p'}} s' = 0$$

by injectivity of $f_{U_p \cap U_{p'}}$. So there exists a section s of \mathcal{F} over U such that $\text{res}_{U_p}^U s = s_p$. We now show $f_U(s) = t$. Consider

$$\text{res}_{U_p}^U f_U(s) = f_{U_p}(\text{res}_{U_p}^U s) = f_{U_p}(s_p) = \text{res}_{U_p}^U t$$

Thus $f_U(s) = t$. □

Remark. (i) Consider the map $\mathcal{F}(U) \rightarrow \prod_{p \in U} \mathcal{F}_p$ given by $s \mapsto ((U, s))_{p \in U}$. This is injective by the first sheaf axiom.

(ii) Given two morphisms of sheaves $\varphi, \psi : \mathcal{F} \rightarrow \mathcal{G}$ with $\varphi_p = \psi_p$ for all $p \in X$, we have $\varphi = \psi$.

III. Algebraic Geometry

2.4. Sheafification

Definition. Let \mathcal{F} be a presheaf on X . Then a morphism $\text{sh} : \mathcal{F} \rightarrow \mathcal{F}^{\text{sh}}$ to a sheaf \mathcal{F}^{sh} is a *sheafification* if for any map $\varphi : \mathcal{F} \rightarrow \mathcal{G}$ where \mathcal{G} is a sheaf, φ factors uniquely through sh .

$$\begin{array}{ccc} \mathcal{F} & \xrightarrow{\text{sh}} & \mathcal{F}^{\text{sh}} \\ & \searrow \varphi & \downarrow \text{hook} \\ & & \mathcal{G} \end{array}$$

Remark. (i) As this is a definition by a universal property, \mathcal{F}^{sh} along with the map $\text{sh} : \mathcal{F} \rightarrow \mathcal{F}^{\text{sh}}$ are unique up to unique isomorphism if they exist.

(ii) A morphism of presheaves $\mathcal{F} \rightarrow \mathcal{G}$ induces a morphism of sheaves $\mathcal{F}^{\text{sh}} \rightarrow \mathcal{G}^{\text{sh}}$.

$$\begin{array}{ccc} \mathcal{F} & \xrightarrow{\text{sh}} & \mathcal{F}^{\text{sh}} \\ & \searrow \varphi & \downarrow \text{hook} \\ & & \mathcal{G} \xrightarrow{\text{sh}} \mathcal{G}^{\text{sh}} \end{array}$$

Proposition. Every presheaf admits a sheafification.

Corollary. The stalks of \mathcal{F} and \mathcal{F}^{sh} coincide.

Proof. Suppose (U, f) is a germ of \mathcal{F}^{sh} at $p \in X$. Then $f(p) \in \mathcal{F}_p$ is a germ of \mathcal{F} at p . If $(U, s) \in \mathcal{F}_p$, we can produce the germ $(U, (U, s)_{p \in U})$ of \mathcal{F}^{sh} at $p \in X$. These are inverse operations, and hence give a bijection of stalks. \square

2.5. Kernels and cokernels

Let $\varphi : \mathcal{F} \rightarrow \mathcal{G}$ be a morphism of presheaves. Then we can define presheaves $\ker \varphi$, $\text{coker } \varphi$, $\text{im } \varphi$ by

$$\begin{aligned} (\ker \varphi)(U) &= \ker \varphi_U \\ (\text{coker } \varphi)(U) &= \text{coker } \varphi_U \\ (\text{im } \varphi)(U) &= \text{im } \varphi_U \end{aligned}$$

One can check that these are indeed presheaves.

Proposition. The presheaf kernel for a morphism of sheaves is a sheaf.

Proof. Let $\varphi : \mathcal{F} \rightarrow \mathcal{G}$ be a morphism of sheaves, let $U \subseteq X$ be open, and let $\{U_i\}_{i \in I}$ be an open cover of U . Let $f \in (\ker \varphi)(U)$ be such that $\text{res}_{U_i}^U f = 0$ for each f . Then as $f \in \mathcal{F}(U)$, we can use the fact that \mathcal{F} is a sheaf to conclude $f = 0$.

Now suppose $f_i \in (\ker \varphi)(U_i)$ agree on their intersections. Then they can be glued as elements of $\mathcal{F}(U_i)$ into $f \in \mathcal{F}(U)$. As $\varphi_{U_i}(f_i) = 0$ for each $i \in I$,

$$0 = \varphi_{U_i}(\text{res}_{U_i}^U f) = \text{res}_{U_i}^U \varphi_U(f)$$

So as \mathcal{G} is a sheaf, $\varphi_U(f) = 0$ in $\mathcal{G}(U)$. \square

However, the presheaf cokernel of a morphism of sheaves is not in general a sheaf.

Example. Consider $X = \mathbb{C}$ with the Euclidean topology, and let \mathcal{O}_X be the sheaf of holomorphic functions on X under addition. Let \mathcal{O}_X^* be the sheaf of nowhere vanishing holomorphic functions under multiplication. We have a morphism of sheaves

$$\exp : \mathcal{O}_X \rightarrow \mathcal{O}_X^*$$

given by

$$f \in \mathcal{O}_X(U) \mapsto \exp(f) \in \mathcal{O}_X^*(U)$$

The kernel of \exp is $2\pi i\mathbb{Z}$, where \mathbb{Z} is the constant sheaf. The cokernel is not a sheaf. To show this, consider the cover

$$U_1 = \mathbb{C} \setminus [0, \infty); \quad U_2 = \mathbb{C} \setminus (-\infty, 0]$$

and take $U = U_1 \cup U_2 = \mathbb{C} \setminus \{0\}$. Let $f(z) = z$, so $f \in \mathcal{O}_X^*(U)$, but f is not in the image of $\exp : \mathcal{O}_X(U) \rightarrow \mathcal{O}_X^*(U)$ as there is no single-valued logarithm on $\mathbb{C} \setminus \{0\}$. Hence f defines a nonzero section of $(\text{coker } \exp)(U)$. However, restricting to U_i , a single-valued branch of logarithm is defined, so f is in the image of $\exp : \mathcal{O}_X(U_i) \rightarrow \mathcal{O}_X^*(U_i)$. Thus $\text{res}_{U_i}^U f = 1$, but $f \neq 1$, violating the first sheaf axiom.

Similarly, the image presheaf may not be a sheaf.

Definition. Let $\varphi : \mathcal{F} \rightarrow \mathcal{G}$ be a morphism of sheaves. We define the *sheaf cokernel* and the *sheaf image* of φ to be the sheafifications of the presheaf cokernel and presheaf image respectively.

Remark. It turns out that the sequence

$$0 \longrightarrow 2\pi i\mathbb{Z} \longrightarrow \mathcal{O}_X \xrightarrow{\exp} \mathcal{O}_X^* \longrightarrow 1$$

is an exact sequence of sheaves. In particular,

$$\ker \exp = 2\pi i\mathbb{Z}; \quad \text{coker } \exp = 1$$

Remark. $\ker \varphi, \text{coker } \varphi$ satisfy the category-theoretic definitions of kernels and cokernels. For kernels, the universal property to be satisfied is

$$\begin{array}{ccccc} & & \mathcal{L} & & \\ & \swarrow \exists! & \downarrow \psi & \searrow 0 & \\ \ker \varphi & \longrightarrow & \mathcal{F} & \xrightarrow{\varphi} & \mathcal{G} \\ & \searrow 0 & & & \end{array}$$

III. Algebraic Geometry

For cokernels, we reverse the arrows.

$$\begin{array}{ccccc}
 & & \mathcal{L} & & \\
 & \nearrow \exists! & \uparrow \psi & \nwarrow 0 & \\
 \text{coker } \varphi & \xleftarrow{\quad} & \mathcal{F} & \xleftarrow{\varphi} & \mathcal{G} \\
 & \nwarrow & \downarrow 0 & \nearrow & \\
 & & & &
 \end{array}$$

Definition. We say that \mathcal{F} is a *subsheaf* of \mathcal{G} , written $\mathcal{F} \subseteq \mathcal{G}$, if there are inclusions $\mathcal{F}(U) \subseteq \mathcal{G}(U)$ compatible with the restriction maps.

Kernels are examples of subsheaves.

2.6. Moving between spaces

Let $f : X \rightarrow Y$ be a continuous map of topological spaces, and let \mathcal{F} and \mathcal{G} be sheaves on X and Y respectively.

Definition. The presheaf *pushforward* or *direct image* $f_*\mathcal{F}$ is the presheaf on Y given by

$$f_*\mathcal{F}(U) = \mathcal{F}(f^{-1}(U))$$

Proposition. The presheaf pushforward of a sheaf is a sheaf.

Proof. Let $\{U_i\}_{i \in I}$ be an open cover of U , and let $s \in f_*\mathcal{F}(U)$ with $\text{res}_{U_i}^U s = 0$ for each U_i . Then $\{f^{-1}(U_i)\}_{i \in I}$ is an open cover of $f^{-1}(U)$ and satisfies $\text{res}_{f^{-1}(U_i)}^{f^{-1}(U)} s = 0$ in $\mathcal{F}(f^{-1}(U_i))$. So $s = 0$ as \mathcal{F} is a sheaf.

Similarly, if $s_i \in f_*\mathcal{F}(U)$ are compatible sections, then they can be glued into an element of $\mathcal{F}(f^{-1}(U))$. But this is precisely an element of $f_*\mathcal{F}(U)$, as required. \square

Definition. The *inverse image presheaf* $(f^{-1}\mathcal{G})^{\text{pre}}$ is the presheaf on X given by

$$(f^{-1}\mathcal{G})^{\text{pre}}(V) = \{(s_U, U) \mid f(V) \subseteq U, s_U \in \mathcal{G}(U)\} / \sim$$

where \sim identifies pairs that agree on a smaller open set containing $f(V)$. The *inverse image sheaf* is $f^{-1}\mathcal{G} = ((f^{-1}\mathcal{G})^{\text{pre}})^{\text{sh}}$.

Example. The inverse image presheaf need not be a sheaf, even when f is an open map. Let Y be a topological space, and let $X = Y \sqcup Y$. Take $\mathcal{G} = \underline{\mathbb{Z}}$ the constant sheaf, and $\mathcal{F} = (f^{-1}\mathcal{G})^{\text{pre}}$. Let $U \subseteq Y$ be open, and let $V = f^{-1}(U)$. Then $\mathcal{F}(V) = \mathcal{G}(U) = \mathbb{Z}$, assuming U is connected. But $V = U \sqcup U$, so $\mathcal{F}^{\text{sh}}(V) = \mathcal{G}(U) \times \mathcal{G}(U) = \mathbb{Z}^2$.

Example. Let \mathcal{F} be a sheaf on X , and let π be the map from X to a point. Then $f_*\mathcal{F}$ is a sheaf on a point, which is just an abelian group, specifically $\mathcal{F}(\pi^{-1}(\{\bullet\})) = \mathcal{F}(X)$.

We will use the notation

$$\mathcal{F}(X) = \Gamma(X, \mathcal{F}) = H^0(X, \mathcal{F})$$

where Γ is called the *global sections*, and H_0 is called the *0th cohomology* with coefficients in \mathcal{F} .

For $p \in X$, $i : \{p\} \rightarrow X$. Let \mathcal{G} be a sheaf on $\{p\}$, which is an abelian group A . Consider the sheaf $i_*\mathcal{G}$ on X , defined by

$$(i_*\mathcal{G})(U) = \begin{cases} 0 & \text{if } p \notin U \\ A & \text{if } p \in U \end{cases}$$

This is called the *skyscraper* at p with value A .

3. Schemes

We will now use the notation $f|_U$ for $\text{res}_U^V f$.

3.1. Localisation

Definition. Let A be a ring and $S \subseteq A$ be a multiplicatively closed set. The *localisation* of A at S is

$$S^{-1}A = \{(a, s) \mid a \in A, s \in S\} / \sim$$

where

$$(a, s) \sim (a', s') \iff \exists s'' \in S, s''(as' - a's) = 0 \in A$$

Examples of multiplicatively closed sets include the set of powers of a fixed element, or the complement of a prime ideal. The pair (a, s) represents $\frac{a}{s}$. The extra s'' term represents a unit in this new ring, which may be needed in rings that are not integral domains.

Remark. The natural map $A \rightarrow S^{-1}A$ need not be injective, for example, if S contains a zero divisor.

3.2. Sheaves on a base

Definition. Let X be a topological space and \mathcal{B} be a basis for the topology. A *sheaf on the base \mathcal{B}* consists of assignments $B_i \mapsto F(B_i)$ of abelian groups, with restriction maps $\text{res}_{B_j}^{B_i} : F(B_i) \rightarrow F(B_j)$ whenever $B_j \subseteq B_i$ such that,

- (i) $\text{res}_{B_i}^{B_i} = \text{id}_{B_i}$;
- (ii) $\text{res}_{B_k}^{B_j} \circ \text{res}_{B_j}^{B_i} = \text{res}_{B_k}^{B_i}$

with the additional axioms that

- (i) if $B = \bigcup B_i$ with $B, B_i \in \mathcal{B}$ and $f, g \in F(B)$ such that $f|_{B_i} = g|_{B_i}$ for all i , then $f = g$;
- (ii) if $B = \bigcup B_i$ as above, with $f_i \in F(B_i)$ such that for all i, j and $B' \subseteq B_i \cap B_j$ with $B' \in \mathcal{B}$, $f_i|_{B'} = f_j|_{B'}$, then there exists $f \in F(B)$ with $f|_{B_i} = f_i$.

This is very similar to the definition of a sheaf, but only specified on the basis.

Proposition. Let F be a sheaf on a base \mathcal{B} of X . This determines a sheaf \mathcal{F} on X such that $\mathcal{F}(B) = F(B)$ for all $B \in \mathcal{B}$, agreeing with restriction maps. Moreover, \mathcal{F} is unique up to unique isomorphism.

Proof. We first define the stalks using F :

$$\mathcal{F}_p = \{(s_B, B) \mid p \in B \in \mathcal{B}, s_B \in F(B)\} / \sim$$

We then use a sheafification idea to define $\mathcal{F}(U)$. The elements are the dependent functions $f \in \prod_{p \in U} \mathcal{F}_p$ such that for each $p \in U$, there exists a basic open set B containing p and a section $s \in F(B)$ such that $s_q = f_q$ in \mathcal{F}_q for all $q \in B$. This is then clearly a sheaf. The natural maps $F(B) \rightarrow \mathcal{F}(B)$ are isomorphisms by the sheaf axioms. \square

3.3. The structure sheaf

Recall that the distinguished opens U_f, U_g coincide if and only if f, g are powers of some $h \in A$. Also, if $U_f = U_g$ then $A_f \cong A_g$. Therefore, the assignment $U_f \mapsto A_f$ is well-defined.

Proposition. The assignment $U_f \mapsto A_f$ defines a sheaf of rings on the base $\{U_f\}$ of the topological space $\text{Spec } A$.

Remark. If $\{U_{f_i}\}_{i \in I}$ covers $\text{Spec } A$, there exists a finite subcover. Indeed, since the U_{f_i} cover $\text{Spec } A$, there is no prime ideal $\mathfrak{p} \subseteq A$ containing all $(f_i)_{i \in I}$. Equivalently, $\sum_{i \in I} (f_i) = (1)$. In particular, $1 = \sum_{i \in J} a_i f_i$ for $J \subseteq I$ finite. So $\sum_{i \in J} (f_i) = (1)$, and thus $\{U_{f_i}\}_{i \in J}$ covers $\text{Spec } A$. We say that $\text{Spec } A$ is *quasi-compact*; traditionally the word ‘compact’ is reserved for Hausdorff spaces in the context of algebraic geometry.

Proof. We will check the axioms for the basic open set $B = \text{Spec } A$; the general case follows by applying this result to a localisation. Suppose $\text{Spec } A = \bigcup_{i=1}^n U_{f_i}$; this union is finite by the previous remark. Let $s \in A$ be such that $s|_{U_i} = 0$ for all i . By the definition of localisation, as the set $\{U_{f_i}\}$ is finite there exists m such that $f_i^m s = 0$ for all i . But note that $(1) = (f_i^m)_{i=1}^n$ for any $m > 0$ because the $\{U_{f_i}\}_{i=1}^n$ cover $\text{Spec } A$. Thus $\{U_{f_i^m}\}_{i=1}^n$ cover $\text{Spec } A$.

$$1 = \sum r_i f_i^m \implies s = \sum r_i f_i^m s = 0$$

Now suppose $\text{Spec } A = \bigcup_{i \in I} U_{f_i}$, and $s_i \in A_{f_i}$ are elements that agree in $A_{f_i f_j}$. We need to build an element in A with these restrictions.

First, suppose I is finite. On U_{f_i} , we have chosen $\frac{a_i}{f_i^{\ell_i}} \in A_{f_i}$; we write $g_i = f_i^{\ell_i}$, noting that $U_{f_i} = U_{g_i}$. On the overlaps, by hypothesis we have

$$(g_i g_j)^{m_{ij}} (a_i g_j - a_j g_i) = 0$$

Rewriting this using the fact that $U_f = U_{f^k}$ for all $k > 0$, and assuming $m = m_{ij}$ by taking the largest, we obtain

$$b_i = a_i g_i^m, \quad h_i = g_i^{m+1}$$

so on each U_{h_i} we have chosen an element $\frac{b_i}{h_i}$. Now, as the $U_{h_i} = U_{f_i}$ cover $\text{Spec } A$, we have $1 = \sum r_i h_i$ for some $r_i \in A$. We can thus construct $r = \sum r_i b_i$ with the r_i as above. This construction then has the correct restrictions to $\frac{b_i}{h_i}$ in U_{h_i} .

III. Algebraic Geometry

When I is infinite, choose $(f_i)_{i=1}^n$ such that the U_{f_i} for $i \in \{1, \dots, n\}$ form a cover, and use the finite case to build $r \in A$. This has the correct restrictions to the U_{f_i} for $i \in \{1, \dots, n\}$. Given $(f_1, \dots, f_n, f_\alpha) = A$, the same construction gives a new $r' \in A$, but then by the first sheaf axiom, $r = r'$. \square

Definition. The *structure sheaf* on $\text{Spec } A$ is the sheaf $\mathcal{O}_{\text{Spec } A}$ associated to the sheaf on the base of distinguished opens mapping U_f to A_f .

Remark. The stalk $\mathcal{O}_{\text{Spec } A, \mathfrak{p}}$ is equal to $A_{\mathfrak{p}}$.

3.4. Definitions and examples

Definition. A *ringed space* (X, \mathcal{O}_X) is a topological space X with a sheaf of rings \mathcal{O}_X . An isomorphism of ringed spaces $(X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ is a homeomorphism $\pi : X \rightarrow Y$ and an isomorphism $\mathcal{O}_Y \rightarrow \pi_* \mathcal{O}_X$ of sheaves on Y .

Note that for $U \subseteq X$ open, U is naturally a ringed space with $\mathcal{O}_U(V) = \mathcal{O}_X(V)$.

Definition. An *affine scheme* is a ringed space (X, \mathcal{O}_X) that is isomorphic to $(\text{Spec } A, \mathcal{O}_{\text{Spec } A})$.

Definition. A *scheme* is a ringed space (X, \mathcal{O}_X) where every point $p \in X$ has a neighbourhood U_p such that the ringed space (U_p, \mathcal{O}_{U_p}) is isomorphic to some affine scheme.

Proposition. Let X be a scheme, $U \subseteq X$ an open set, and $i : U \hookrightarrow X$ be the inclusion map. Then, the ringed space (U, \mathcal{O}_U) is a scheme, where

$$\mathcal{O}_U = \mathcal{O}_X \Big|_U = i^{-1} \mathcal{O}_X$$

For example, take $X = \text{Spec } A$ and $U = U_f$ for some $f \in A$. Then $(U, \mathcal{O}_U) \cong (\text{Spec } A_f, \mathcal{O}_{\text{Spec } A_f})$.

Proof. Let $p \in U \subseteq X$. Since X is a scheme, we can find $(V_p, \mathcal{O}_X|_{V_p})$ inside X with $p \in V_p$, such that V_p is isomorphic to an affine scheme. Then take $V_p \cap U \subseteq U$ with structure sheaf given by the inclusion map. Note that $V_p \cap U$ may not be affine, but $V_p \cong \text{Spec } B$, and the distinguished opens in $\text{Spec } B$ form a basis. This reduces the problem to the example of a distinguished open set above. \square

Definition. *Affine space* of dimension n over k is defined to be

$$\mathbb{A}_k^n = \text{Spec } k[x_1, \dots, x_n]$$

Example. Let

$$U = \mathbb{A}_k^{n^2} \setminus \{\det(x_{ij}) = 0\}$$

which is the open set representing $GL_n(k)$. We will show that the multiplication map $U \times U \rightarrow U$ is a morphism of schemes.

Example. Let $U = \mathbb{A}_k^2 \setminus (x, y)$. This is a scheme representing a plane without an origin. We claim that U is not an affine scheme. Suppose that U were affine; we aim to calculate $\mathcal{O}_U(U)$. Write

$$U_x = \mathbb{V}(x)^c \subseteq \mathbb{A}_k^2; \quad U_y = \mathbb{V}(y)^c \subseteq \mathbb{A}_k^2$$

These two open sets cover U , and

$$U_x \cap U_y = U_{xy} = \mathbb{A}_k^2 \setminus \mathbb{V}(xy)$$

Then,

$$\mathcal{O}_U(U_x) = k[x, x^{-1}, y]; \quad \mathcal{O}_U(U_y) = k[x, y, y^{-1}]; \quad \mathcal{O}_U(U_x \cap U_y) = k[x, x^{-1}, y, y^{-1}]$$

The restriction maps $\mathcal{O}_U(U_x) \rightarrow \mathcal{O}_U(U_{xy})$ and $\mathcal{O}_U(U_y) \rightarrow \mathcal{O}_U(U_{xy})$ are the obvious ones. By the sheaf axioms,

$$\mathcal{O}_U(U) = k[x, x^{-1}, y] \cap k[x, y, y^{-1}] \subseteq k[x, x^{-1}, y, y^{-1}]$$

Thus, $\mathcal{O}_U(U) = k[x, y]$. This is a contradiction: one way to see this is that there exists a maximal ideal (x, y) in the ring of global sections in (U, \mathcal{O}_U) with empty vanishing locus.

In general, if X is a scheme, $f \in \Gamma(X, \mathcal{O}_X) = \mathcal{O}_X(X)$, and $p \in X$, then there is a well-defined stalk $\mathcal{O}_{X,p}$ at p , which is of the form $A_{\mathfrak{p}}$ up to isomorphism, where \mathfrak{p} is a prime ideal. To say this, we are using an isomorphism of an open set $V_{\mathfrak{p}}$ containing p to $\text{Spec } A$. In particular, $A_{\mathfrak{p}}$ has a unique maximal ideal, namely $\mathfrak{p}A_{\mathfrak{p}}$. We say that f vanishes at p if its image in $A_{\mathfrak{p}}/\mathfrak{p}A_{\mathfrak{p}}$, or equivalently, $f \in \mathfrak{p}A_{\mathfrak{p}}$. As a consequence, the vanishing locus $\mathbb{V}(f) \subseteq X$ is well-defined.

3.5. Gluing sheaves

Let X be a topological space with a cover $\{U_{\alpha}\}$. Let $\{\mathcal{F}_{\alpha}\}$ be sheaves on $\{U_{\alpha}\}$, with isomorphisms

$$\varphi_{\alpha\beta} : \mathcal{F}_{\alpha} \Big|_{U_{\alpha} \cap U_{\beta}} \rightarrow \mathcal{F}_{\beta} \Big|_{U_{\alpha} \cap U_{\beta}}$$

such that

$$\varphi_{\alpha\alpha} = \text{id}; \quad \varphi_{\alpha\beta} = \varphi_{\beta\alpha}^{-1}; \quad \varphi_{\beta\gamma} \circ \varphi_{\alpha\beta} = \varphi_{\alpha\gamma}$$

The last equation is called the *cocycle condition*. This combination of conditions resembles the definition of an equivalence relation, with reflexivity, symmetry, and transitivity.

We will construct a sheaf \mathcal{F} on X . Given $V \subseteq X$ open, we define

$$\mathcal{F}(V) = \left\{ (s_{\alpha}) \in \prod_{\alpha} \mathcal{F}_{\alpha}(U_{\alpha} \cap V) \mid \varphi_{\alpha\beta} \left(s_{\alpha} \Big|_{V \cap U_{\alpha} \cap U_{\beta}} \right) = s_{\beta} \Big|_{V \cap U_{\alpha} \cap U_{\beta}} \right\}$$

\mathcal{F} is a presheaf. Indeed, given $(s_{\alpha}) \in \mathcal{F}(V)$ and $W \subseteq V$ open, we take

$$(s_{\alpha}) \Big|_W = \left(\text{res}_{W \cap U_{\alpha}}^{V \cap U_{\alpha}}(s_{\alpha}) \right)_{\alpha}$$

This lies in $\mathcal{F}(W)$ by the sheaf axioms. One can easily check that this is a sheaf.

III. Algebraic Geometry

Proposition. $\mathcal{F}|_{U_\gamma}$ and \mathcal{F}_γ are canonically isomorphic as sheaves on U_γ .

Proof. First, we construct a map $\mathcal{F}_\gamma \rightarrow \mathcal{F}|_{U_\gamma}$. Let $V \subseteq U_\gamma$ and $s \in \mathcal{F}_\gamma(V)$. Define its image in $\mathcal{F}|_{U_\gamma}$ to be

$$\varphi_{\gamma\alpha} \left(s \Big|_{V \cap U_\alpha} \right)_\alpha$$

We must check that this tuple lies in $\mathcal{F}|_{U_\gamma}(V) = \mathcal{F}(V)$.

$$\varphi_{\alpha\beta} \circ \varphi_{\gamma\alpha} \left(s \Big|_{V \cap U_\alpha \cap U_\beta} \right) = \varphi_{\gamma\beta} \left(s \Big|_{V \cap U_\alpha \cap U_\beta} \right)$$

□

3.6. Gluing schemes

Let (X, \mathcal{O}_X) and (Y, \mathcal{O}_Y) be schemes with open sets $U \subseteq X, V \subseteq Y$, and let $\varphi : (U, \mathcal{O}_X|_U) \rightarrow (V, \mathcal{O}_Y|_V)$ be an isomorphism. The topological spaces X, Y can be glued on U, V using φ .

First, take $S = X \sqcup Y / U \sim V$. By definition of the quotient topology, the images of X and Y in S form an open cover, and their intersection is the image of U , or equivalently, the image of V . Now, we can glue the structure sheaves on these open sets as described in the previous subsection. Note that in this case, there is no cocycle condition.

Example (the bug-eyed line; the line with doubled origin). Let k be a field. Let $X = \text{Spec } k[t]$ and $Y = \text{Spec } k[u]$. Let

$$U = \text{Spec } k[t, t^{-1}] = \text{Spec } k[t]_t = U_t \subseteq X; \quad V = \text{Spec } k[u, u^{-1}] = \text{Spec } k[u]_u = U_u \subseteq Y$$

We define the isomorphism $\varphi : U \rightarrow V$ given by $t \mapsto u$. Technically, we define an isomorphism of rings $k[u, u^{-1}] \rightarrow k[t, t^{-1}]$ by $u \mapsto t$ and then apply Spec . At the level of topological spaces, $X = \mathbb{A}_k^1$ and $Y = \mathbb{A}_k^1$, so $U = \mathbb{A}_k^1 \setminus \{(t)\}$ and $V = \mathbb{A}_k^1 \setminus \{(u)\}$. Gluing along this isomorphism, we obtain a scheme S which is a copy of \mathbb{A}_k^1 but with two origins. Note that the generic points in X and Y lie in U and V respectively, and thus are glued into a single generic point in S .

Consider the open sets in S . Open sets entirely contained within X and Y yield open sets in S . We also have open sets of the form $W = S \setminus \{\mathfrak{p}_1, \dots, \mathfrak{p}_r\}$ where \mathfrak{p}_i is contained in U or V . One example is $W = S$; we can calculate $\mathcal{O}_S(S)$ using the sheaf axioms, and one can show that it is isomorphic to $k[t]$. We can conclude that S is not an affine scheme, because there is a maximal ideal in $k[t]$ where the vanishing locus is precisely two points.

Example (the projective line). Let $X = \text{Spec } k[t]$ and $Y = \text{Spec } k[s]$, and define $U = \text{Spec } k[t, t^{-1}], V = \text{Spec } k[s, s^{-1}]$ as above. We glue these schemes using the isomorphism $s \mapsto t^{-1}$, giving the projective line \mathbb{P}_k^1 .

Proposition. $\mathcal{O}_{\mathbb{P}_k^1}(\mathbb{P}_k^1) = k$.

Proof sketch. We use the same idea as in the previous example. The only elements of $k[t, t^{-1}]$ that are both polynomials in t and t^{-1} are the constants. \square

In particular, \mathbb{P}_k^1 is not an affine scheme.

Example. We can similarly build a scheme S which is a copy of \mathbb{A}_k^2 with a doubled origin. This has the interesting property that there exist affine open subschemes $U_1, U_2 \subseteq S$ such that $U_1 \cap U_2$ is not affine; we can take U_1 and U_2 to be S but with one of the origins deleted. Note that \mathbb{A}_k^1 without the origin is affine.

Let $\{X_i\}_{i \in I}$ be schemes, $X_{ij} \subseteq X_i$ be open subschemes, and $f_{ij} : X_{ij} \rightarrow X_{ji}$ be isomorphisms such that

$$f_{ii} = \text{id}_{X_i}; \quad f_{ij} = f_{ji}^{-1}; \quad f_{ik} = f_{jk} \circ f_{ij}$$

where the last equality holds whenever it is defined. Then there is a unique scheme X with an open cover by the X_i , glued along these isomorphisms. This is an elaboration of the above construction, which is discussed on the first example sheet.

Let A be a ring, and let $X_i = \text{Spec } A\left[\frac{x_0}{x_i}, \dots, \frac{x_n}{x_i}\right]$. Let $X_{ij} = \mathbb{V}\left(\frac{x_j}{x_i}\right)^c \subseteq X_i$. We define the isomorphisms $X_{ij} \rightarrow X_{ji}$ by $\frac{x_k}{x_i} \mapsto \frac{x_k}{x_j} \left(\frac{x_i}{x_j}\right)^{-1}$. The resulting glued scheme is called *projective n -space*, denoted \mathbb{P}_A^n .

3.7. The Proj construction

Definition. A \mathbb{Z} -grading on a ring A is a decomposition

$$A = \bigoplus_{i \in \mathbb{Z}} A_i$$

as abelian groups, such that $A_i A_j \subseteq A_{i+j}$.

Example. Let $A = k[x_0, \dots, x_n]$, and let A_d be the set of degree d homogeneous polynomials, together with the zero polynomial.

Example. Let $I \subseteq k[x_0, \dots, x_n]$ be a homogeneous ideal; that is, an ideal generated by homogeneous elements of possibly different degrees. Then, for $A = k[x_0, \dots, x_n]$, the ring A/I is also naturally graded.

Note that by definition, A_0 is a subring of A . For simplicity, we will always assume in this course that the degree 1 elements of a graded ring generate A as an algebra over A_0 . We also typically assume that $A_i = 0$ for $i < 0$. We define

$$A_+ = \bigoplus_{i \geq 1} A_i \subseteq A$$

III. Algebraic Geometry

This forms an ideal in A , called the *irrelevant ideal*. If A is a polynomial ring with the usual grading, the irrelevant ideal corresponds to the point $\mathbf{0}$ in the theory of varieties. This aligns with the definition of projective space in classical algebraic geometry, in which the point $\mathbf{0}$ is deleted.

A *homogeneous element* $f \in A$ is an element contained in some A_d . An ideal I of A is called *homogeneous* if it is generated by homogeneous elements.

Definition. Let A be a graded ring. $\text{Proj } A$ is the set of homogeneous prime ideals in A that do not contain the irrelevant ideal. If $I \subseteq A$ is homogeneous, we define

$$\mathbb{V}(I) = \{\mathfrak{p} \in \text{Proj } A \mid I \subseteq \mathfrak{p}\}$$

The *Zariski topology* on $\text{Proj } A$ is the topology where the closed sets are of the form $\mathbb{V}(I)$ where I is a homogeneous ideal.

The Spec construction allows us to convert rings into schemes; the Proj construction allows us to convert graded rings into schemes. Unlike Spec, the construction of Proj is not functorial.

Let $f \in A_1$ and $U_f = \text{Proj } A \setminus \mathbb{V}(f)$. Observe that the set $\{U_f\}_{f \in A_1}$ covers $\text{Proj } A$, because the f generate the unit ideal. The ring $A\left[\frac{1}{f}\right] = A_f$ is naturally \mathbb{Z} -graded by defining $\deg \frac{1}{f} = -\deg f$. Note that A_f may have negatively graded elements, even though A does not.

Example. Let $A = k[x_0, x_1]$ and $f = x_0$. Then in $A\left[\frac{1}{f}\right] = k[x_0, x_1, x_0^{-1}]$, the degree zero elements include k and elements such as $\frac{x_1}{x_0}, \frac{x_1^2 + x_1 x_0}{x_0^2}$. There are degree one elements such as $\frac{x_1^2}{x_0}$.

Proposition. There is a natural bijection

$$\{\text{homogeneous prime ideals in } A \text{ that miss } f\} \leftrightarrow \{\text{prime ideals in } (A_f)_0\}$$

Note also that the set of homogeneous prime ideals in A that miss f are naturally in bijection with the homogeneous prime ideals in A_f .

Proof. Suppose \mathfrak{q} is a prime ideal in $\left(A\left[\frac{1}{f}\right]\right)_0$. Then let $\psi(\mathfrak{q})$ be the ideal

$$\psi(\mathfrak{q}) = \left(\bigcup_{d \geq 0} \left\{ a \in A_d \mid \frac{a}{f^d} \in \mathfrak{q} \right\} \right) \subseteq A$$

One can check that this is prime. Now suppose \mathfrak{p} is a homogeneous prime ideal missing f . Define $\varphi(\mathfrak{p})$ to be

$$\varphi(\mathfrak{p}) = \left(p \cdot A\left[\frac{1}{f}\right] \cap \left(A\left[\frac{1}{f}\right]\right)_0 \right)$$

This ideal is also prime.

One can easily check that $\varphi \circ \psi$ is the identity. For the other direction, suppose \mathfrak{p} is a homogeneous prime ideal missing f ; we show that $\mathfrak{p} = \psi(\varphi(\mathfrak{p}))$ by antisymmetry. If $a \in \mathfrak{p} \in A_d$, then $\frac{a}{f^d} \in \varphi(\mathfrak{p})$, so $a \in \psi(\varphi(\mathfrak{p}))$ by construction. Conversely, if $a \in \psi(\varphi(\mathfrak{p}))$, then $\frac{a}{f^d} \in \varphi(\mathfrak{p})$ for some d , so there exists $b \in \mathfrak{p}$ such that $\frac{b}{f^e} = \frac{a}{f^d}$ in $A\left[\frac{1}{f}\right]$. Hence for some $k \geq 0$, we have $f^k(f^d b - f^e a) = 0$, and $f^{e+k} \notin \mathfrak{p}$. But by primality, $a \in \mathfrak{p}$, as required. \square

The bijection constructed is compatible with ideal containment, so is a homeomorphism of topological spaces

$$U_f \leftrightarrow \text{Spec}(A_f)_0$$

Thus $\text{Proj } A$ is covered by open sets homeomorphic to an affine scheme. If $f, g \in A_1$, then $U_f \cap U_g$ is naturally homeomorphic to

$$\left(\text{Spec } A\left[\frac{1}{f}\right] \right)_0 \left[\frac{f}{g} \right] = \text{Spec} (A[f^{-1}, g^{-1}])_0$$

Take the open cover $\{U_f\}$ with structure sheaf $\mathcal{O}_{\text{Spec}(A_f)_0}$ on each U_f , and isomorphisms on $U_f \cap U_g$ by the condition above. The cocycle condition follows from the formal properties of the localisation. Therefore, $\text{Proj } A$ is a scheme.

If $A = k[x_0, \dots, x_n]$ with the standard grading, we write \mathbb{P}_k^n for $\text{Proj } A$.

4. Morphisms

4.1. Morphisms of ringed spaces

Let (X, \mathcal{O}_X) be a scheme. The stalks $\mathcal{O}_{X,p}$ are local rings: they have a unique maximal ideal, which is the set of all non-unit elements. Given $f \in \mathcal{O}_X(U)$, we can meaningfully ask whether f vanishes at p ; that is, if the image of f in $\mathcal{O}_{X,p}$ is contained in the maximal ideal.

Definition. A morphism of ringed spaces $f : (X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ consists of a continuous function $f : X \rightarrow Y$ and a morphism $f^\# : \mathcal{O}_Y \rightarrow f_*\mathcal{O}_X$ between sheaves of rings on Y .

$f^\#$ represents function composition with f^{-1} , although the ring \mathcal{O}_X may not be a ring of functions. It is possible to find a morphism $(f, f^\#)$ between schemes (X, \mathcal{O}_X) and (Y, \mathcal{O}_Y) such that there exists $q \in U \subseteq Y$ and $h \in \mathcal{O}_Y(U)$ such that h vanishes at q but $f^\#(h) \in \mathcal{O}_X(f^{-1}(U))$ does not vanish at some $p \in X$ with $f(p) = q$. This motivates the definition of a morphism of schemes.

Let $f : X \rightarrow Y$ be a morphism of ringed spaces. Given any point $p \in X$, there is an induced map $f^\# : \mathcal{O}_{Y,f(p)} \rightarrow \mathcal{O}_{X,p}$. Explicitly, given $s \in \mathcal{O}_{Y,f(p)}$, we can represent it by (s_U, U) where U is open, $f(p) \in U$, and $s_U \in \mathcal{O}_Y(U)$. Now, $f^\#(s_U) \in \mathcal{O}_X(f^{-1}(U))$, so the pair $(f^\#(s_U), f^{-1}(U))$ defines an element of $\mathcal{O}_{X,p}$.

Definition. A ringed space (X, \mathcal{O}_X) is called a *locally ringed space* if for all $p \in X$, the stalk $\mathcal{O}_{X,p}$ is a local ring. A morphism of locally ringed spaces $(f, f^\#) : (X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ is a morphism of ringed spaces such that if \mathfrak{m}_p denotes the maximal ideal in $\mathcal{O}_{X,p}$, then $f^\#(\mathfrak{m}_{f(p)}) \subseteq \mathfrak{m}_p$.

This encapsulates the idea that functions vanishing on the codomain must also vanish on the domain after the inverse image, as the maximal ideal represents functions vanishing at the point.

4.2. Morphisms of schemes

Note that all schemes are locally ringed spaces.

Definition. A *morphism of schemes* $X \rightarrow Y$ is a morphism of locally ringed spaces $X \rightarrow Y$.

Theorem. There is a natural bijection

$$\{\text{morphisms of schemes } \text{Spec } B \rightarrow \text{Spec } A\} \leftrightarrow \{\text{homomorphisms of rings } A \rightarrow B\}$$

Proof. First, recall that a section s of a sheaf \mathcal{F} on U is a coherent collection of elements of the stalks $s(p) \in \mathcal{F}_p$ for all $p \in U$. We will construct a map of schemes $\text{Spec } B \rightarrow \text{Spec } A$ for every ring homomorphism $A \rightarrow B$, and then show that every morphism of schemes arises in this way.

4. Morphisms

Let $\varphi : A \rightarrow B$ be a ring homomorphism. Let $\varphi^{-1} : \text{Spec } B \rightarrow \text{Spec } A$ be the map of topological spaces; this is a continuous function. We now build

$$\varphi^\# : \mathcal{O}_{\text{Spec } A} \rightarrow \varphi_*^{-1} \mathcal{O}_{\text{Spec } B}$$

At the level of stalks, the map $A_{\varphi^{-1}(\mathfrak{p})} \rightarrow B_{\mathfrak{p}}$ is induced by φ by mapping $\frac{a}{s}$ to $\frac{\varphi(a)}{\varphi(s)}$. This is well-defined, as for $s \notin \varphi^{-1}(\mathfrak{p})$, then $\varphi(s) \notin \mathfrak{p}$. Observe that this is automatically a local homomorphism.

We must now show that this choice of maps on stalks extends to a map between sheaves. Given $U \subseteq \text{Spec } A$, we need to define

$$\varphi^\# : \mathcal{O}_{\text{Spec } A}(U) \rightarrow \mathcal{O}_{\text{Spec } B}((\varphi^{-1})^{-1}(U))$$

An element $s \in \mathcal{O}_{\text{Spec } A}(U)$ is a collection of assignments $(\mathfrak{p} \mapsto s(\mathfrak{p}))_{\mathfrak{p} \in U}$ for $\mathfrak{p} \in U$ and $s(\mathfrak{p}) \in A_{\mathfrak{p}}$. We then define $\varphi^\#$ by

$$(\mathfrak{p} \mapsto s(\mathfrak{p}))_{\mathfrak{p} \in U} \mapsto (q \mapsto \varphi_q(s(\varphi^{-1}(q))))_{q \in (\varphi^{-1})^{-1}(U)}$$

One can check that the gluing conditions are satisfied.

Conversely, suppose $(f, f^\#) : \text{Spec } B \rightarrow \text{Spec } A$ is a morphism of schemes. Using the fact that we have a map of global sections $\mathcal{O}_{\text{Spec } A}(\text{Spec } A) \rightarrow \mathcal{O}_{\text{Spec } B}(\text{Spec } B)$, we obtain a ring homomorphism $g : A \rightarrow B$. We must check that $g^{-1} : \text{Spec } B \rightarrow \text{Spec } A$ gives the correct map f on topological spaces, and that the construction above yields the correct map $f^\#$ on sheaves. The maps on stalks are compatible with restriction, so the following diagram commutes for all $\mathfrak{p} \in \text{Spec } B$.

$$\begin{array}{ccc} \Gamma(\text{Spec } A, \mathcal{O}_{\text{Spec } A}) & \longrightarrow & \Gamma(\text{Spec } B, \mathcal{O}_{\text{Spec } B}) \\ \downarrow & & \downarrow \\ \mathcal{O}_{\text{Spec } A, f(\mathfrak{p})} & \longrightarrow & \mathcal{O}_{\text{Spec } B, \mathfrak{p}} \end{array}$$

Equivalently, the following diagram commutes for all $\mathfrak{p} \in \text{Spec } B$.

$$\begin{array}{ccc} A & \longrightarrow & B \\ \downarrow & & \downarrow \\ A_{f(\mathfrak{p})} & \longrightarrow & B_{\mathfrak{p}} \end{array}$$

Since the morphism is local, $(f^\#)^{-1}(\mathfrak{p}B_{\mathfrak{p}}) = f(\mathfrak{p})A_{f(\mathfrak{p})}$. As the above diagram commutes, $g^{-1} = f$ as maps of topological spaces, and the maps of structure sheaves agree at the level of stalks by construction so they must agree everywhere. \square

4.3. Immersions

Definition. Let X, Y be schemes. A morphism of schemes $f : X \rightarrow Y$ is an *open immersion* if f induces an isomorphism of X onto an open subscheme $(U, \mathcal{O}_Y|_U)$ of Y . A morphism $f : X \rightarrow Y$ is a *closed immersion* if f is a homeomorphism onto a closed subset of Y , and $g^\# : \mathcal{O}_Y \rightarrow g_* \mathcal{O}_X$ is surjective.

Example. Let $k[t] \rightarrow k[t]/(t^2)$. The induced map $\text{Spec } k[t]/(t^2) \rightarrow \text{Spec } k[t]$ is a closed immersion. More generally, let A be a ring and I be an ideal in A . Then the induced map $\text{Spec } A/I \rightarrow \text{Spec } A$ is a closed immersion.

Definition. Let Y be a scheme. A *closed subscheme* of Y is an equivalence class of closed immersions $X \rightarrow Y$, where we say $f : X \rightarrow Y$ and $f' : X' \rightarrow Y$ are equivalent if there is a commutative triangle

$$\begin{array}{ccc} X & \xleftarrow{\sim} & X' \\ & \searrow f & \swarrow f' \\ & Y & \end{array}$$

4.4. Fibre products

The notion of fibre product will simultaneously generalise the notions of product, intersections of closed subschemes, and inverse images of subschemes (such as points) along morphisms.

Definition. Consider a diagram

$$\begin{array}{ccc} & X & \\ & \downarrow & \\ Y & \longrightarrow & S \end{array}$$

The *fibre product* is a scheme $X \times_S Y$ making the following diagram commute:

$$\begin{array}{ccc} X \times_S Y & \xrightarrow{p_X} & X \\ p_Y \downarrow & & \downarrow \\ Y & \longrightarrow & S \end{array}$$

such that for any other scheme Z together with morphisms q_X, q_Y completing the square, there is a unique factorisation through $X \times_S Y$, making the following diagram commute.

$$\begin{array}{ccccc} Z & & & & \\ & \searrow q_X & & & \\ & & X \times_S Y & \xrightarrow{p_X} & X \\ & \swarrow q_Y & \downarrow p_Y & & \downarrow \\ & & Y & \longrightarrow & S \end{array}$$

Note that as this is a definition by universal property, if $X \times_S Y$ exists, it is unique up to unique isomorphism. The fibre product in schemes is the category-theoretic *pullback*.

Example. (i) In the category of sets, the fibre product of the diagram

$$\begin{array}{ccc} & X & \\ & \downarrow r_X & \\ Y & \xrightarrow{r_Y} & S \end{array}$$

is the set

$$X \times_S Y = \{(x, y) \in X \times Y \mid r_X(x) = r_Y(y)\}$$

(ii) In the category of topological spaces, the fibre product is defined to be the same set, assigning $X \times_S Y$ the subspace topology as a subset of $X \times Y$.

(iii) Let $r_X : X \rightarrow S$ be a map of sets, and let $Y = \{*\}$ with $r_Y(*) = s \in S$. Then

$$X \times_S Y = r_X^{-1}(s)$$

(iv) Let $r_X : X \rightarrow S$ and $r_Y : Y \rightarrow S$ be inclusions of subsets. Then

$$X \times_S Y = X \cap Y$$

Theorem. Fibre products of schemes exist.

Proof sketch. *Step 1.* Let X, Y, S be affine schemes, with associated rings A, B, R . Then the fibre product $X \times_S Y$ exists, and is isomorphic to $\text{Spec}(A \otimes_R B)$. Note that the tensor product is the category-theoretic pushout in the category of rings. We must now check that the universal property of the fibre product is satisfied. Consider the commutative square

$$\begin{array}{ccc} Z & \longrightarrow & X \\ \downarrow & & \downarrow \\ Y & \longrightarrow & S \end{array}$$

If Z is an affine scheme, the result holds. It is a general fact that a map of schemes $Z \rightarrow \text{Spec}(A \otimes_R B)$ is the same data as a map $A \otimes_R B \rightarrow \Gamma(Z, \mathcal{O}_Z)$.

Step 2. Let X, Y, S be arbitrary schemes. If $X \times_S Y$ exists and $U \subseteq X$ is an open subscheme, then $U \times_S Y$ also exists, by taking the inverse image of U under the projection $X \times_S Y \rightarrow X$ endowed with the structure of an open subscheme.

Step 3. If X is covered by open subschemes $\{X_i\}$, then if $X_i \times_S Y$ exists for all i , then $X \times_S Y$ exists, by gluing each of the $X_i \times_S Y$ together. Note that the ability to glue these schemes together relies on Step 2, and the fact that there is no cocycle condition.

III. Algebraic Geometry

Step 4. If Y and S are affine, then $X \times_S Y$ exists by Step 3, by covering X by affine subschemes. As X and Y are interchangeable, $X \times_S Y$ exists for any X and Y as long as S is affine.

Step 5. Now, cover S by affine subschemes $\{S_i\}$. Let X_i, Y_i be the preimages of S_i in X and Y respectively. Now, $X_i \times_{S_i} Y_i$ exists. Observe by the universal property that $X_i \times_{S_i} Y_i = X_i \times SY_i$. Finally, gluing gives $X \times_S Y$ as required. \square

Example. (i) We have

$$\mathbb{P}_{\mathbb{C}}^n = \mathbb{P}_{\mathbb{Z}}^n \times_{\text{Spec } \mathbb{Z}} \text{Spec } \mathbb{C}$$

where the map $\text{Spec } \mathbb{C} \rightarrow \text{Spec } \mathbb{Z}$ is induced by the ring homomorphism $\mathbb{Z} \rightarrow \mathbb{C}$, and the map $\mathbb{P}_{\mathbb{Z}}^n \rightarrow \text{Spec } \mathbb{Z}$ is induced locally by the inclusion $\mathbb{Z} \rightarrow \mathbb{Z}\left[\frac{x_0}{x_i}, \dots, \frac{x_n}{x_i}\right]$. Note also that

$$\mathbb{Z}[\mathbf{x}] \otimes_{\mathbb{Z}} \mathbb{C} = \mathbb{C}[\mathbf{x}]$$

(ii) Let $C = \text{Spec } \mathbb{C}[x, y]_{(y-x^2)}$ and $L = \text{Spec } \mathbb{C}[x, y]_{(y)}$. We have natural closed immersions $C \rightarrow \mathbb{A}_{\mathbb{C}}^2$ and $L \rightarrow \mathbb{A}_{\mathbb{C}}^2$. One can show that

$$C \times_{\mathbb{A}_{\mathbb{C}}^2} L = \text{Spec } \mathbb{C}[x]_{(x^2)}$$

representing the intersection.

4.5. Schemes over a base

In scheme theory, we often fix a scheme S called the *base scheme*, and consider other schemes with a fixed map to S . These form a category of schemes *over* S , where the morphisms are the morphisms of schemes $f : X \rightarrow Y$ such that the following diagram commutes.

$$\begin{array}{ccc} X & \xrightarrow{f} & Y \\ & \searrow & \swarrow \\ & S & \end{array}$$

This is known as Grothendieck's *relative point of view*. Typically, S is the spectrum of a field or a ring. Note that every scheme has a unique morphism to $\text{Spec } \mathbb{Z}$, so the category of schemes is isomorphic to the category of schemes over $\text{Spec } \mathbb{Z}$. The product of X and Y in the category of schemes over S is the fibre product $X \times_S Y$. Analogously, in commutative algebra, we often consider algebras of a fixed ring, and the category of rings is isomorphic to the category of \mathbb{Z} -algebras.

4.6. Separatedness

Recall that a topological space X is Hausdorff if and only if the diagonal $\Delta_X \subseteq X \times X$ is closed.

4. Morphisms

Definition. Let $X \rightarrow S$ be a morphism of schemes. Then the *diagonal* is the morphism $\Delta_{X/S} : X \rightarrow X \times_S X$ induced using the universal property by the following diagram.

$$\begin{array}{ccccc}
 X & & \xrightarrow{\text{id}_X} & & X \\
 & \searrow \text{dashed} & & \searrow & \\
 & & X \times_S X & \xrightarrow{\quad} & X \\
 & \swarrow \text{id}_X & \downarrow & & \downarrow \\
 & & X & \xrightarrow{\quad} & S
 \end{array}$$

We write Δ for $\Delta_{X/S}$ if X and S are clear from context.

Remark. If U, V are open subschemes of X and $S = \text{Spec } k$ for a field k , then

$$\Delta^{-1}(U \times_S V) = U \cap V$$

Definition. A morphism $X \rightarrow S$ is *separated* if $\Delta_{X/S} : X \rightarrow X \times_S X$ is a closed immersion.

Example. Let $X = \text{Spec } \mathbb{C}[t]$, let $S = \text{Spec } \mathbb{C}$, and induce the map $X \rightarrow S$ by the \mathbb{C} -algebra homomorphism $\mathbb{C} \rightarrow \mathbb{C}[t]$. Then

$$X \times_S X = \text{Spec}(\mathbb{C}[t] \otimes_{\mathbb{C}} \mathbb{C}[t])$$

and the diagonal map Δ is induced by the multiplication map

$$\mathbb{C}[t] \otimes_{\mathbb{C}} \mathbb{C}[t] \rightarrow \mathbb{C}[t]$$

Note that Δ is closed, as the map $\mathbb{C}[t] \otimes_{\mathbb{C}} \mathbb{C}[t] \rightarrow \mathbb{C}[t]$ is surjective.

Proposition. Let $g : X \rightarrow S$ be a morphism of schemes. Then there is a factorisation of $\Delta_{X/S}$ as follows.

$$\begin{array}{ccc}
 & U & \\
 \text{closed immersion} \nearrow & & \searrow \text{open immersion} \\
 X & \xrightarrow{\Delta_{X/S}} & X \times_S X
 \end{array}$$

We say that $g : X \rightarrow S$ is a *locally closed immersion*.

Proof. Let S be covered by open affine subschemes $\{V_i\}$, and suppose X is covered by open affine subschemes $\{U_{ij}\}$, where for some fixed i , the U_{ij} cover $g^{-1}(V_i)$. We have morphisms $U_{ij} \rightarrow V_i$ induced by

$$\begin{array}{ccccc}
 U_{ij} & \longrightarrow & g^{-1}(V_i) & \longrightarrow & V_i \\
 & & \downarrow & & \downarrow \\
 & & X & \longrightarrow & S
 \end{array}$$

III. Algebraic Geometry

where the commutative square is a fibre product. Observe that $U_{ij} \times_{V_i} U_{ij}$ is affine and open in $X \times_S X$, and their union contains the image of the diagonal $\Delta_{X/S}$. Also,

$$\Delta^{-1}(U_{ij} \times_{V_i} U_{ij}) = U_{ij} \subseteq X$$

Let U be the union of the $U_{ij} \times_{V_i} U_{ij}$ over all i, j . Then the second map in the statement is clearly an open immersion. Observe that to check if $f : T \rightarrow T'$ is a closed immersion, it suffices to check locally on the codomain. For each U_{ij} , the diagonal is a map $U_{ij} \rightarrow U_{ij} \times_{V_i} U_{ij}$, which one can show is a closed immersion. \square

Proposition. If $X \rightarrow S$ is a morphism of affine schemes, then $\Delta_{X/S}$ is a closed immersion.

Proof. Let $X = \text{Spec } A, S = \text{Spec } B$, and let the map $X \rightarrow S$ be given by a map $B \rightarrow A$. Then the map $A \otimes_B A \rightarrow A$ is surjective as required. \square

Thus every morphism of affine schemes is separated.

Corollary. Let $X \rightarrow S$ be a morphism of schemes. If the image of $\Delta_{X/S}$ is closed as a topological subspace, then $X \rightarrow S$ is separated.

Proof. A locally closed immersion onto a closed subset is a closed immersion. \square

Example. (i) Recall the bug-eyed line

$$\mathbb{A}_k^1 \sqcup \mathbb{A}_k^1 / \sim$$

where if $U = \mathbb{A}_k^1 \setminus \{0\} \subseteq \mathbb{A}_k^1$ and V is defined similarly, we define the isomorphism $V \rightarrow U$ by the map $u \mapsto t : k[u, u^{-1}] \rightarrow k[t, t^{-1}]$. We claim that the bug-eyed line is not separated over $\text{Spec } k$. We can compute $X \times_S X$ by the gluing construction of the fibre product. This is a plane with doubled axes and four origins. The diagonal only contains two of the four origins, and this is not a closed subset.

(ii) Open and closed immersions are always separated.

(iii) All monomorphisms are separated.

(iv) Compositions of separated morphisms are separated.

(v) Suppose $X \rightarrow S$ is separated and $S' \rightarrow S$ is an embedding. Then the map $X \times_S S' \rightarrow S'$ that comes from

$$\begin{array}{ccc} X \times_S S' & \longrightarrow & X \\ \downarrow & & \downarrow \\ S' & \longrightarrow & S \end{array}$$

is also separated. This is called a *base extension*: the right-hand side of the diagram is the original morphism $X \rightarrow S$, and the left-hand side can be thought of as the same morphism under a base change.

Proposition. Let R be a ring. The morphism $\mathbb{P}_R^n \rightarrow \operatorname{Spec} R$ is separated.

Proposition. We want to show that the map Δ in the following diagram is closed, where the commutative square is a fibre product.

$$\begin{array}{ccccc} \mathbb{P}_R^n & \xrightarrow{\Delta} & \mathbb{P}_R^n \times_R \mathbb{P}_R^n & \longrightarrow & \mathbb{P}_R^n \\ & & \downarrow & & \downarrow \\ & & \mathbb{P}_R^n & \longrightarrow & \operatorname{Spec} R \end{array}$$

It suffices to check this result on an open cover of $\mathbb{P}_R^n \times_R \mathbb{P}_R^n$. Let $A = R[x_0, \dots, x_n]$ with the usual grading, so $\operatorname{Proj} A = \mathbb{P}_R^n$. Then let $U_i = \operatorname{Spec} \left(A \left[\frac{1}{x_i} \right] \right)_0$. These U_i form an open cover of \mathbb{P}_R^n . Now,

$$U_i \times_R U_j = \operatorname{Spec} R \left[\frac{x_0}{x_i}, \dots, \frac{x_n}{x_i}, \frac{y_0}{y_j}, \dots, \frac{y_n}{y_j} \right]$$

Observe that the restriction of Δ to $\Delta^{-1}(U_i \times_R U_j)$ is

$$U_i \cap U_j \rightarrow U_i \times_R U_j$$

given on rings by the map

$$R \left[\frac{x_0}{x_i}, \dots, \frac{x_n}{x_i} \right] \left[\frac{x_i}{x_j} \right] \leftarrow R \left[\frac{x_0}{x_i}, \dots, \frac{x_n}{x_i}, \frac{y_0}{y_j}, \dots, \frac{y_n}{y_j} \right]$$

by changing y_k into x_k . This is surjective, and the $U_i \times_R U_j$ cover $\mathbb{P}_R^n \times_R \mathbb{P}_R^n$, so Δ is closed.

Definition. Let $k = \bar{k}$ be an algebraically closed field. Let $X \rightarrow \operatorname{Spec} k$ be a scheme over $\operatorname{Spec} k$. We say that X is of *finite type* over $\operatorname{Spec} k$ if there is a cover of X by affines $\{U_\alpha\}_\alpha$ such that $\mathcal{O}_X(U_\alpha)$ is finitely generated k -algebra. We say that X is *reduced* if for all open $U \subseteq X$, $\mathcal{O}_X(U)$ has no nilpotent elements.

Definition. A morphism $X \rightarrow \operatorname{Spec} k$ is a *variety* if it is reduced, of finite type, and separated.

4.7. Properness

Definition. Let $f : X \rightarrow S$ be a morphism. Then f is of *finite type* if there exists an affine cover of S by open $\{V_\alpha\}_\alpha$ where $V_\alpha = \operatorname{Spec} A_\alpha$, and covers $\{U_{\alpha\beta}\}_\beta$ of $f^{-1}(V_\alpha)$ by open affine subschemes with $U_{\alpha\beta} = \operatorname{Spec} B_{\alpha\beta}$, such that $B_{\alpha\beta}$ is a finitely generated A_α -algebra, and $\{U_{\alpha\beta}\}_\beta$ can be chosen to be finite.

Definition. A morphism $f : X \rightarrow S$ is *closed* if it is closed as a map of topological spaces. It is *universally closed* if for any $S' \rightarrow S$, the induced map $X \times_S S' \rightarrow S'$ is also closed. f is *proper* if it is separated, of finite type, and universally closed.

Example. (i) Closed immersions are proper.

III. Algebraic Geometry

- (ii) The obvious map $\mathbb{A}_k^1 \rightarrow \operatorname{Spec} k$ is not proper, because it is not universally closed. Indeed, consider the fibre product

$$\begin{array}{ccc} \mathbb{A}_k^2 & \longrightarrow & \mathbb{A}_k^1 \\ \downarrow & & \downarrow \\ \mathbb{A}_k^1 & \longrightarrow & \operatorname{Spec} k \end{array}$$

Consider $Z \subseteq \mathbb{A}_k^2 = \operatorname{Spec} k[x, y]$ given by the vanishing locus of $xy - 1$. Then the projection of Z onto each axis is not Zariski closed.

- (iii) The bug-eyed line is neither separated nor universally closed.

Remark. If $X \rightarrow S$ is universally closed, then any base extension $X \times_S S' \rightarrow S'$ is also universally closed. Similarly, separatedness, properness and being of finite type are stable under base extension.

Proposition. Let R be a commutative ring. Then the morphism $\mathbb{P}_R^n \rightarrow \operatorname{Spec} R$ is proper.

Proof. We have already shown that $\mathbb{P}_R^n \rightarrow \operatorname{Spec} R$ is separated. It is of finite type by construction. It suffices to prove that the morphism is universally closed for $R = \mathbb{Z}$, because $\mathbb{P}_R^n = \mathbb{P}_{\mathbb{Z}}^n \times_{\operatorname{Spec} \mathbb{Z}} \operatorname{Spec} R$. We must show that for any $Y \rightarrow \operatorname{Spec} \mathbb{Z}$, the base extension $\mathbb{P}_{\mathbb{Z}}^n \times_{\operatorname{Spec} \mathbb{Z}} Y \rightarrow Y$ is closed. But Y is covered by affine schemes of the form $\operatorname{Spec} R$, and closedness is local on the codomain, it suffices to show that $\mathbb{P}_R^n \rightarrow \operatorname{Spec} R$ is closed.

Let $Z \subseteq \mathbb{P}_R^n$ be Zariski closed, so Z is the vanishing locus of homogeneous polynomials $\{g_1, g_2, \dots\}$. We want to show that if π is the map $\mathbb{P}_R^n \rightarrow \operatorname{Spec} R$, then $\pi(Z)$ is closed. We need to find equations for $\pi(Z)$, or equivalently, we need to characterise the prime ideals \mathfrak{p} of R such that $\pi^{-1}(\mathfrak{p}) \cap Z$ is nonempty. Let $k(\mathfrak{p}) = FF(R/\mathfrak{p})$. We have a morphism $\operatorname{Spec} k(\mathfrak{p}) \rightarrow \operatorname{Spec} R$. Let $Z_{\mathfrak{p}} = Z \times_{\operatorname{Spec} R} \operatorname{Spec} k(\mathfrak{p})$; we want to know for which \mathfrak{p} this scheme is nonempty. If we take the equations g_1, g_2, \dots and reduce modulo \mathfrak{p} , we obtain equations $\bar{g}_1, \bar{g}_2, \dots$ which are homogeneous polynomials in $k(\mathfrak{p})$. Thus $Z_{\mathfrak{p}}$ is nonempty if and only if $\bar{g}_1, \bar{g}_2, \dots$ cut out more than the origin in $\mathbb{A}_{k(\mathfrak{p})}^{n+1}$. In particular, $Z_{\mathfrak{p}}$ is nonempty if and only if

$$\sqrt{(\bar{g}_1, \bar{g}_2, \dots)} \not\supseteq (x_0, \dots, x_n); \quad \mathbb{P}_R^n = \operatorname{Proj} R[x_0, \dots, x_n]$$

Equivalently, for all positive integers d ,

$$(x_0, \dots, x_n)^d \not\subseteq (\bar{g}_1, \bar{g}_2, \dots)$$

Write $A = R[\mathbf{x}]$ with the usual grading. The non-containment condition above holds if and only if the map

$$\bigoplus_i A_{d-\deg g_i} \rightarrow A_d$$

given by $f_i \mapsto f_i g_i$ in the i th factor is not surjective modulo \mathfrak{p} , or equivalently in $k(\mathfrak{p})$, for all degrees d . This condition is given by the maximal minors of the matrix associated to $\bigoplus_i A_{d-\deg g_i} \rightarrow A_d$, which is a set of infinitely many polynomials, each in the coefficients of the g_i . \square

4.8. Valuative criteria

From here, we will assume that all schemes are Noetherian; that is, it has a finite cover by spectra of Noetherian rings.

Definition. A *discrete valuation ring* is a local principal ideal domain.

Example. (i) $\mathbb{C}[[t]]$ is a discrete valuation ring.

(ii) $\mathcal{O}_{\mathbb{A}^1,0} = \left\{ \frac{f(t)}{g(t)} \mid g(0) \neq 0 \right\}$ is a discrete valuation ring.

(iii) Similarly, $\mathbb{Z}_{(p)}, \mathbb{Z}_p$ are discrete valuation rings, where $\mathbb{Z}_{(p)}$ denotes the localisation of \mathbb{Z} at the prime ideal (p) , and \mathbb{Z}_p denotes the p -adic integers.

We will often drop the word ‘discrete’.

Remark. Let A be a valuation ring. In discrete valuation rings, every nonzero prime ideal is maximal, so $\text{Spec } A$ consists of two points, (0) and the unique maximal ideal \mathfrak{m} . The topology on $\text{Spec } A = \{(0), \mathfrak{m}\}$ has the property that (0) is dense and \mathfrak{m} is closed. This is called the *Sierpiński topology*.

Any generator π for \mathfrak{m} is called a *uniformiser* or a *uniformising parameter*. For example, in $\mathbb{C}[[t]]$, every power series with nonzero constant term is a unit, and t is a uniformiser.

Given a uniformiser, any nonzero element $a \in A$ can be written as $u\pi^k$ where u is a unit and k is a unique natural number called the *valuation* of a . This gives a map $A \setminus \{0\} \rightarrow \mathbb{N}$ mapping a value a to its valuation; this is independent of the choice of uniformiser.

The field of fractions of A is a *valued field* $K = FF(A)$; the valuation extends to a multiplicative function $K \setminus \{0\} \rightarrow \mathbb{Z}$ given by the difference of valuations of the numerator and denominator.

Example. Let $A = k[[t]]$, then $K = k((t))$ is the field of Laurent series in one variable in k . The valuation is the order of vanishing at zero.

One can consider the open immersion $\text{Spec } K \rightarrow \text{Spec } A$ as the inclusion from a disc with a punctured origin to a disc.

Theorem. Let $f : X \rightarrow Y$ be a morphism of schemes. Then f is separated if and only if for any (discrete) valuation ring A with function field K and diagram

$$\begin{array}{ccc} \text{Spec } K & \longrightarrow & X \\ \downarrow & & \downarrow \\ \text{Spec } A & \longrightarrow & Y \end{array}$$

III. Algebraic Geometry

then there exists at most one lift $\text{Spec } A \rightarrow X$ that makes the following diagram commute.

$$\begin{array}{ccc} \text{Spec } K & \longrightarrow & X \\ \downarrow & \nearrow \text{dashed} & \downarrow \\ \text{Spec } A & \longrightarrow & Y \end{array}$$

Similarly, f is universally closed if and only if there exists at least one lift $\text{Spec } A \rightarrow X$ that makes the diagram commute.

In particular, a morphism is proper if and only if there is a unique lift, and the morphism is of finite type. The proof is omitted.

Remark. (i) The map $\mathbb{P}_R^n \rightarrow \text{Spec } R$ is proper.

(ii) The map $\mathbb{A}_R^n \rightarrow \text{Spec } R$ is not proper, but is separated.

(iii) Closed immersions are proper. In particular, if $Z \rightarrow \mathbb{P}_R^n$ is closed, then $Z \rightarrow \text{Spec } R$ is proper.

(iv) Compositions of proper (respectively separated) morphisms are proper (separated).

(v) If $f : X \rightarrow Y$ is proper, then for any $Y' \rightarrow Y$, the base extension $X \times_Y Y' \rightarrow Y'$ is also proper.

Example. We show that $\mathbb{A}_k^1 \rightarrow \text{Spec } k$ is not proper by showing it is not universally closed. Write $\mathbb{A}_k^1 = \text{Spec } k[x]$, and consider $A = k[[t]]$ and $K = k((t))$.

$$\begin{array}{ccc} \text{Spec } k((t)) & \xrightarrow{\varphi} & \mathbb{A}_k^1 \\ \downarrow & & \downarrow \\ \text{Spec } k[[t]] & \longrightarrow & \text{Spec } k \end{array}$$

The map $\text{Spec } k[[t]] \rightarrow \text{Spec } k$ is the obvious morphism. Let φ be induced by the map on rings $k[x] \rightarrow k((t))$ given by $x \mapsto \frac{1}{t}$. Then the map does not factor through $\text{Spec } k[[t]] \rightarrow \text{Spec } k((t))$, as required. However, if we replace \mathbb{A}_k^1 with \mathbb{P}_k^1 , there is always an affine chart in \mathbb{P}^1 such that φ is of the form $x \mapsto t$.

5. Modules over the structure sheaf

5.1. Definitions

Example. Let $\mathbb{C}P^n$ be the variety $\mathbb{C}^{n+1} \setminus \{0\}$ modulo scaling by \mathbb{C} . We have a structure sheaf $\mathcal{O}_{\mathbb{C}P^n}$, where if $U \subseteq \mathbb{C}P^n$ is Zariski open, we define

$$\mathcal{O}_{\mathbb{C}P^n}(U) = \left\{ \frac{P(\mathbf{x})}{Q(\mathbf{x})} \mid P, Q \text{ homogeneous of the same degree, and the ratio is regular at all } p \in U \right\}$$

For any integer d , we can consider a sheaf $\mathcal{O}_{\mathbb{C}P^n}(d)$ given by

$$\mathcal{O}_{\mathbb{C}P^n}(d)(U) = \left\{ \frac{P(\mathbf{x})}{Q(\mathbf{x})} \mid P, Q \text{ homogeneous, } \deg P - \deg Q = d, \text{ and regular at all } p \in U \right\}$$

This is a sheaf of groups, but not a sheaf of rings as it is not closed under multiplication for $d \neq 0$. Note that $\mathcal{O}_{\mathbb{C}P^n}(d)(U)$ is a module over $\mathcal{O}_{\mathbb{C}P^n}(U)$, and the multiplication commutes with restriction.

Example. Let A be a ring, and let M be an A -module. We define the sheaf $\mathcal{F}_M = M^{\text{sh}}$ on $\text{Spec } A$ as follows. If $U \subseteq \text{Spec } A$ is a distinguished open $U = U_f$, then we set

$$\mathcal{F}_M(U) = M_f$$

which is the module M localised at f . This defines a sheaf on a base, and hence extends to a unique sheaf on $\text{Spec } A$.

Definition. Let (X, \mathcal{O}_X) be a ringed space. A *sheaf of \mathcal{O}_X -modules* on X is a sheaf \mathcal{F} of abelian groups together with a multiplication $\mathcal{F}(U) \times \mathcal{O}_X(U) \rightarrow \mathcal{F}(U)$ that makes $\mathcal{F}(U)$ into an $\mathcal{O}_X(U)$ -module, that is compatible with restriction.

$$\begin{array}{ccc} \mathcal{F}(V) \times \mathcal{O}_X(V) & \longrightarrow & \mathcal{F}(V) \\ \downarrow & & \downarrow \\ \mathcal{F}(U) \times \mathcal{O}_X(U) & \longrightarrow & \mathcal{F}(U) \end{array}$$

Similarly, we can define a sheaf of \mathcal{O}_X -algebras. A morphism between sheaves of modules $\varphi : \mathcal{F} \rightarrow \mathcal{G}$ on X is a homomorphism of sheaves of abelian groups that is compatible with multiplication.

Given morphisms of sheaves of modules on X , we can locally take kernels, cokernels, images, direct sums, tensor products, hom functors, and all of these extend to sheaves of modules. In the case of cokernels, images, and tensor products, we require a sheafification step. For example, the presheaf tensor product $\mathcal{F} \otimes_{\mathcal{O}_X} \mathcal{G}$ associated to an open set $U \subseteq X$ is given by $\mathcal{F}(U) \otimes_{\mathcal{O}_X(U)} \mathcal{G}(U)$; the sheaf tensor product is given by sheafification.

Given a morphism of ringed spaces or schemes $f : X \rightarrow Y$, the pushforward of an \mathcal{O}_X -module \mathcal{F} is the sheaf of abelian groups $f_*\mathcal{F}$. As a morphism of ringed spaces, we also

III. Algebraic Geometry

have a map $f^\# : \mathcal{O}_Y \rightarrow f_*\mathcal{O}_X$, giving $f_*\mathcal{F}$ an \mathcal{O}_Y -module structure. Given an open set $U \subseteq Y$, $a \in \mathcal{O}_Y(U)$, and $m \in f_*\mathcal{F}(U) = \mathcal{F}(f^{-1}(U))$, we define $a \cdot m = f^\#(a) \cdot m$, where $f^\#(a) \in \mathcal{O}_X(f^{-1}(U))$.

Conversely, if \mathcal{G} is a sheaf of \mathcal{O}_Y -modules, we define

$$f^*\mathcal{G} = f^{-1}\mathcal{G} \otimes_{f^{-1}\mathcal{O}_Y} \mathcal{O}_X$$

where the $f^{-1}\mathcal{O}_Y$ -module structure on \mathcal{O}_X is defined via the adjoint to $f^\#$.

5.2. Quasi-coherence

Definition. A *quasi-coherent sheaf* \mathcal{F} on a scheme X is a sheaf of \mathcal{O}_X -modules such that there exists a cover of X by affines $\{U_i\}$ such that $\mathcal{F}|_{U_i}$ is the sheaf associated to a module over the ring $\mathcal{O}_X(U_i)$. If these modules can be taken to be finitely generated, we say \mathcal{F} is *coherent*.

Example. (i) On any scheme X , \mathcal{O}_X is quasi-coherent (and, in fact, coherent).

(ii) $\bigoplus_I \mathcal{O}_X$ is quasi-coherent, but not coherent if I is infinite.

(iii) If $i : X \rightarrow Y$ is a closed immersion, then $i_*\mathcal{O}_X$ is a quasi-coherent \mathcal{O}_Y -module. Let $U \subseteq Y$ be an affine open set, so $U = \text{Spec } A$. Then $X \cap U \rightarrow U$ gives an ideal $I \subseteq A$ which is the kernel of the surjection $\mathcal{O}_Y(U) \rightarrow \mathcal{O}_X(X \cap U)$. On U , $i_*\mathcal{O}_X|_U$ is the sheaf associated to the A -module A/I .

Proposition. An \mathcal{O}_X -module \mathcal{F} is quasi-coherent if and only if for any affine open $U \subseteq X$ with $U = \text{Spec } A$, $\mathcal{F}|_U$ is the sheaf associated to a module over A .

We first prove the following key technical lemma.

Lemma. Let $X = \text{Spec } A$, $f \in A$, and \mathcal{F} a quasi-coherent \mathcal{O}_X -module. Let $s \in \Gamma(X, \mathcal{F})$. Then

(i) If s restricts to 0 on U_f , then $f^n s = 0$ for some $n \geq 1$.

(ii) If $t \in \mathcal{F}(U_f)$, then $f^n t$ is the restriction of a global section of \mathcal{F} over X for some $n \geq 1$.

Proof. There exists some cover of X by schemes of the form $\text{Spec } B = V$, such that $\mathcal{F}|_V = M^{\text{sh}}$ for M a B -module. We can cover each such V by distinguished affines of the form U_g for some $g \in A$. Then $\mathcal{F}|_{U_g} = (M \otimes_B A_g)^{\text{sh}}$, as $\mathcal{F}|_V$ is quasi-coherent. But recall that $\text{Spec } A$ is quasi-compact: every open cover has a finite subcover. So finitely many U_{g_i} will suffice to cover X by open sets such that \mathcal{F} restricts to M_i^{sh} on U_{g_i} . Then the lemma follows from formal properties of localisation. \square

We now prove the main proposition.

Proof. Given $U \subseteq X$, observe that $\mathcal{F}|_U$ is also quasi-coherent. We can thus reduce the statement to the case where $X = \operatorname{Spec} A$. Now we take $M = \Gamma(X, \mathcal{F})$, and let M^{sh} be the associated sheaf. We claim that $M^{\text{sh}} \cong \mathcal{F}$. Let $\alpha : M^{\text{sh}} \rightarrow \mathcal{F}$ be the map given by restriction (for example via stalks). Then α is an isomorphism at the level of stalks by the above lemma, so is an isomorphism globally. \square

In particular, the quasi-coherent sheaves of modules over $\operatorname{Spec} A$ are precisely the modules over A . The coherent sheaves of modules over $\operatorname{Spec} A$ are precisely the finitely-generated modules over A .

Proposition. (i) Images, kernels, and cokernels of maps of (quasi-)coherent sheaves remain (quasi-)coherent.

(ii) If $f : X \rightarrow S$ is a morphism of schemes and \mathcal{F} is a (quasi-)coherent sheaf of modules on S , then $f^*\mathcal{F}$ is also (quasi-)coherent.

(iii) If $f : X \rightarrow S$ is a morphism of schemes and \mathcal{G} is a quasi-coherent sheaf on X , then $f_*\mathcal{G}$ is also quasi-coherent.

The proofs are omitted and non-examinable. Note that (iii) need not hold for coherent sheaves: let $f : \mathbb{A}_k^1 \rightarrow \operatorname{Spec} k$ be the obvious map, and consider $f_*\mathcal{O}_{\mathbb{A}_k^1}$. This is a quasi-coherent sheaf on $\operatorname{Spec} k$, so is a k -vector space, which is $k[t]$. As a module, this is not finitely generated. Observe that if $f : \mathbb{P}_k^1 \rightarrow \operatorname{Spec} k$, then $f_*\mathcal{O}_{\mathbb{P}_k^1}$ is the sheaf associated to k . In general, if \mathcal{G} is a coherent sheaf on X and $f : X \rightarrow S$ is proper, then $f_*\mathcal{G}$ is coherent.

Let A be a graded ring, with the usual assumptions on its generators. To build $\operatorname{Proj} A$, we consider the cover by $\operatorname{Spec} \left(A \left[\frac{1}{f} \right]_0 \right)$ for $f \in A_1$. We can produce a similar construction for modules.

Let M be a *graded* A -module, that is,

$$M = \bigoplus_{d \in \mathbb{Z}} M_d$$

where each M_d is an abelian group, M is an A -module, and $A_i M_j \subseteq M_{i+j}$. Consider the sheaf determined by the association

$$\operatorname{Proj} A \supseteq U_f \mapsto \left(M \left[\frac{1}{f} \right]_0 \right)$$

To each $U_f = \mathbb{V}(f)^c$, we associate the degree zero elements of the localisation of M at f . This gives a quasi-coherent sheaf on $\operatorname{Proj} A$ by identical arguments as in the Proj construction.

Definition. Let X be a scheme and \mathcal{F} be a quasi-coherent \mathcal{O}_X -module. We say that \mathcal{F} is

(i) *free*, if $\mathcal{F} \simeq \mathcal{O}_X^{\oplus I}$ for some set I ;

III. Algebraic Geometry

- (ii) an *(algebraic) vector bundle* or *locally free* if there exists an open cover $\{U_i\}$ such that $\mathcal{F}|_{U_i}$ is free;
- (iii) a *line bundle* or an *invertible sheaf* if it is a vector bundle that is locally isomorphic to \mathcal{O}_X .

Note that such sheaves are coherent if and only if the index sets I can be taken to be finite.

5.3. Coherent sheaves on projective schemes

Definition. Let A be a graded ring, and let M be a graded A -module. For $d \in \mathbb{Z}$, we define $M(d)$, called M twisted by d , to be the module such that

$$(M(d))_k = M_{k+d}$$

Definition. Let $X = \text{Proj } A$ where A is a graded ring and let $d \in \mathbb{Z}$. The sheaf $\mathcal{O}_X(d)$ is defined to be the sheaf associated to the graded module $A(d)$. In particular, $\mathcal{O}_X(1)$ is called the *twisting sheaf*.

Remark. $\mathcal{O}_X(d) = \mathcal{O}_X(1)^{\otimes d}$. Note that the tensor product of graded modules is additive in the grading.

Example. Consider $\text{Proj } k[x_0, \dots, x_n] = \mathbb{P}_k^n$. The global sections of $\mathcal{O}_{\mathbb{P}_k^n}(d)$ are homogeneous degree d polynomials in the x_i . In particular, if $d < 0$, then $\Gamma(\mathbb{P}_k^n, \mathcal{O}_{\mathbb{P}_k^n}(d)) = 0$.

Definition. An \mathcal{O}_X -module \mathcal{F} is called *globally generated* or *generated by global sections* if it is a quotient of $\mathcal{O}_X^{\oplus r}$ for some r ; that is, there is a surjective map of coherent sheaves $\mathcal{O}_X^{\oplus r} \rightarrow \mathcal{F}$. Equivalently, there exist elements $s_1, \dots, s_r \in \Gamma(X, \mathcal{F})$ such that $\{s_i\}$ generate the stalks \mathcal{F}_p over $\mathcal{O}_{X,p}$ for all $p \in X$.

Theorem. Let $i : X \rightarrow \mathbb{P}_R^n$ be a closed immersion. Let $\mathcal{O}_X(1)$ be the restriction of $\mathcal{O}_{\mathbb{P}_R^n}(1)$, so $\mathcal{O}_X(1) = i^* \mathcal{O}_{\mathbb{P}_R^n}(1)$. Let \mathcal{F} be a coherent sheaf on X . Then there exists an integer d_0 such that for all $d \geq d_0$, the sheaf

$$\mathcal{F}(d) = \mathcal{F} \otimes_{\mathcal{O}_X} \mathcal{O}_X(d)$$

is globally generated.

Proof. By formal properties, it is equivalent to show the statement for $i_* \mathcal{F}$; that is, $i_* \mathcal{F}(d)$ is globally generated on \mathbb{P}_R^n . Write $\mathbb{P}_R^n = \text{Proj}[x_0, \dots, x_n]$, and cover \mathbb{P}_R^n by $U_i = \text{Spec } B_i$ where $B_i = R\left[\frac{x_0}{x_i}\right]$. We know that $\mathcal{F}|_{U_i} = M_i^{\text{sh}}$, and M_i is a finitely generated B_i -module. Let $\{s_{ij}\}$ be generators for M_i . We claim that the sections $\{x_i^d s_{ij}\}_j$ of $\mathcal{F}(d)|_{U_i}(U_i)$ are restrictions of global sections t_{ij} of $\mathcal{F}(d)$ for sufficiently large d . Such d can be chosen to be independent of i and j . Indeed, if s_{ij} is an element of $M_i = \mathcal{F}(U_i)$ and $x_i \in \mathcal{O}_X(1) = \mathcal{O}_{\mathbb{P}_R^n}(1)$, we can show that $x_i^d s_{ij} \in (F \otimes \mathcal{O}(d))(U_i)$ is a restriction of a global section.

5. Modules over the structure sheaf

Now, on U_i , the s_{ij} generate M_i^{sh} , but we have a morphism of sheaves $\mathcal{F} \rightarrow \mathcal{F}(d)$, mapping s to $x_i^d s := s \otimes x_i^d$. This map is globally defined, but on U_i this restricts to an isomorphism $\mathcal{F}|_{U_i} \rightarrow \mathcal{F}(d)|_{U_i}$ as x_i is invertible on U_i . Since the $\{s_{ij}\}$ generate $\mathcal{F}|_{U_i}$, the $x_i^d s_j$ generate $\mathcal{F}(d)|_{U_i}$. Thus, the t_{ij} globally generate $\mathcal{F}(d)$. \square

Corollary. Let $i : X \hookrightarrow \mathbb{P}_R^n$ be a closed immersion. Let \mathcal{F} be a coherent sheaf on X . Then \mathcal{F} is a quotient of $\mathcal{O}(-d)^{\oplus N}$ for some sufficiently large N and some $d \in \mathbb{Z}$.

6. Divisors

6.1. Height and dimension

Recall that for a prime ideal \mathfrak{p} in R , its *height* is the largest n such that there exists a chain of inclusions of prime ideals

$$\mathfrak{p}_0 \subsetneq \mathfrak{p}_1 \subsetneq \cdots \subsetneq \mathfrak{p}_n = \mathfrak{p}$$

For example, if R is an integral domain, a prime ideal is of height 1 if and only if no nonzero prime ideal is strictly contained within it.

Example. (i) In any integral domain, (0) has height 0.

(ii) In $\mathbb{C}[x, y]$, the ideal (x) has height 1, and the ideal (x, y) has height 2.

It can be shown that in a unique factorisation domain, every prime ideal of height 1 is principal.

We will globalise the notion of height 1 prime ideals, giving *Weil divisors*, and also the notion of principal ideals, giving *Cartier divisors*. In the case of Weil divisors, we will assume that the ambient scheme X is Noetherian, integral, separated, and *regular in codimension 1*.

If X is integral and $U = \text{Spec } A$ is an open affine, then the ideal $(0) \subseteq A$ is called the *generic point* of X . Each open affine is dense as they are irreducible, so they have a nontrivial intersection, including their generic points. The generic points given by each U therefore coincide in X . This point is often denoted by η or η_X .

Definition. Let X be a scheme.

(i) The *dimension* of X is the length n of the longest chain of nonempty closed irreducible subsets

$$Z_0 \subsetneq Z_1 \subsetneq \cdots \subsetneq Z_n$$

(ii) Let $Z \subseteq X$ be closed and irreducible. The *codimension* of X is the length n of the longest chain

$$Z = Z_0 \subsetneq Z_1 \subsetneq \cdots \subsetneq Z_n$$

(iii) If X is a *Noetherian topological space*, so every decreasing sequence of closed subsets stabilises, then every closed $Z \subseteq X$ has a decomposition into finitely many irreducible closed subsets.

(iv) Suppose X is Noetherian, integral, and separated. We say that X is *regular in codimension 1* if for every subspace $Y \subseteq X$ that is closed, irreducible, and of codimension 1, if η_Y denotes the generic point of Y , then \mathcal{O}_{X, η_Y} is a discrete valuation ring, or equivalently a local principal ideal domain.

6.2. Weil divisors

Definition. Let X be Noetherian, integral, separated, and regular in codimension 1. A *prime divisor* on X is an integral closed subscheme of codimension 1. A *Weil divisor* on X is an element of the free abelian group $\text{Div}(X)$ generated by the prime divisors.

We will write $D \in \text{Div}(X)$ as $\sum_i n_{Y_i} [Y_i]$ where the Y_i are prime divisors.

Definition. A Weil divisor $\sum_i n_{Y_i} [Y_i]$ is *effective* if all n_{Y_i} are nonnegative.

If X is integral, for $\text{Spec } A = U \subseteq X$, the local ring $\mathcal{O}_{X,\eta}$ is a field, as it is in particular the fraction field of A . Indeed, because η is contained in every open affine, $\mathcal{O}_{X,\eta}$ permits arbitrary denominators.

Let $f \in \mathcal{O}_{X,\eta_X} = k(X)$ be nonzero. Since for every prime divisor $Y \subseteq X$, the ring \mathcal{O}_{X,η_Y} is a discrete valuation ring, we can calculate the valuation $\nu_Y(f)$ of f in this ring. We thus define the divisor

$$\text{div}(f) = \sum_{Y \subseteq X \text{ prime}} \nu_Y(f) [Y]$$

We claim that this is a Weil divisor; that is, the sum is finite.

Proposition. The sum

$$\sum_{Y \subseteq X \text{ prime}} \nu_Y(f) [Y]$$

is finite.

Proof. Let $f \in k(X)^\times$, and choose A such that $U = \text{Spec } A$ is an affine open, so $FF(A) = k(X)$. We can also require that $f \in A$ by localising at the denominator, so f is *regular* on U . Then $X \setminus U$ is closed and of codimension at least 1, so only finitely many prime Weil divisors Y of X are contained in $X \setminus U$. On U , as f is regular, $\nu_Y(f) \geq 0$ for all Y . But $\nu_Y(f) > 0$ if and only if Y is contained in $\mathbb{V}(f) \subseteq U$, and by the same argument, there are only finitely many such Y . \square

Definition. A Weil divisor of the form $\text{div}(f)$ is called *principal*. In $\text{Div}(X)$, the set of principal divisors form a subgroup $\text{Prin}(X)$, and we define the *Weil divisor class group* of X to be

$$\text{Cl}(X) = \text{Div}(X) / \text{Prin}(X)$$

Remark. (i) Let A be a Noetherian domain. Then A is a unique factorisation domain if and only if A is integrally closed and $\text{Cl}(\text{Spec } A)$ is trivial. This is related to the fact that in unique factorisation domains, all primes of height 1 are principal. In particular, there exist rings with nontrivial class groups of their spectra.

(ii) $\text{Cl}(\mathbb{A}_k^n) = 0$.

(iii) $\text{Cl}(\mathbb{P}_k^n) \cong \mathbb{Z}$; we will prove this shortly.

III. Algebraic Geometry

- (iv) Let $Z \subseteq X$ be closed, and let $U = X \setminus Z$. Then there is a surjective map $\text{Cl}(X) \rightarrow \text{Cl}(U)$, defined by $[Y] \mapsto [Y \cap U]$, but instead mapping $[Y]$ to zero if $Y \cap U = \emptyset$. This is well-defined, as $k(X)$ and $k(U)$ are naturally isomorphic, so principal divisors are mapped to principal divisors. For surjectivity, note that given a prime Weil divisor $D \subseteq U$, its closure \bar{D} in X is a prime Weil divisor that restricts to D under the map.
- (v) If Z has codimension at least 2, then $\text{Cl}(X) \rightarrow \text{Cl}(U)$ is an isomorphism. This is because Z does not enter the definition of $\text{Cl}(X)$.
- (vi) If $Z \subseteq X$ is integral, closed, and of codimension 1, there is an exact sequence

$$\mathbb{Z} \xrightarrow{1 \mapsto [Z]} \text{Cl}(X) \longrightarrow \text{Cl}(U) \longrightarrow 0$$

called the *excision* exact sequence. Indeed, the kernel of $\text{Cl}(X) \rightarrow \text{Cl}(U)$ are exactly the divisors in X contained in Z .

Proposition. Let k be a field. Then, $\text{Cl}(\mathbb{P}_k^n) \cong \mathbb{Z}$.

Proof. Let $D \subseteq \mathbb{P}^n$ be integral, closed, and of codimension 1. Then $D = \mathbb{V}(f)$ where f is homogeneous of some degree d ; we will define $\deg(D) = d$. We extend linearly to obtain a homomorphism $\deg : \text{Div}(\mathbb{P}_k^n) \rightarrow \mathbb{Z}$. We claim that this gives an isomorphism $\text{Cl}(\mathbb{P}_k^n) \rightarrow \mathbb{Z}$. First, this is well defined on classes, since if $f = \frac{g}{h}$ is a rational function, then g and h are homogeneous polynomials of the same degree, so $\deg(\text{div}(f)) = 0$. This is surjective, by taking $H = \mathbb{V}(x_0)$ for x_0 homogeneous linear. For injectivity, suppose $D = \sum n_{Y_i} [Y_i]$ with $\sum n_{Y_i} \deg(Y_i) = 0$. Write $Y_i = \mathbb{V}(g_i)$, and let $f = \prod g_i^{n_{Y_i}}$. Now f is a homogeneous rational function of degree zero. \square

6.3. Cartier divisors

Let X be a scheme. Consider the presheaf on X given by mapping $U = \text{Spec } A$ to $S^{-1}A$ where S is the set of all elements that are not zero divisors. Sheafification yields the sheaf of rings \mathcal{K}_X . Define $\mathcal{K}_X^* \subseteq \mathcal{K}_X$ to be the subsheaf of invertible elements; this is a sheaf of abelian groups under multiplication. If X is integral, then \mathcal{K}_X is the constant sheaf, where the constant field is $\mathcal{O}_{X, \eta_X} = FF(A)$ for any affine open $\text{Spec } A$.

Similarly, let $\mathcal{O}_X^* \subseteq \mathcal{O}_X$ be the subsheaf of invertible elements. Thus, every section of $\mathcal{K}_X^*/\mathcal{O}_X^*$ can be prescribed by $\{(U_i, f_i)\}$ where U_i is a cover of X , f_i is a section of $\mathcal{K}_X^*(U_i)$, and that on $U_i \cap U_j$, the ratio f_i/f_j lies in $\mathcal{O}_X^*(U_i \cap U_j)$.

Definition. A *Cartier divisor* is a global section of the sheaf $\mathcal{K}_X^*/\mathcal{O}_X^*$.

We have a surjective sheaf homomorphism $\mathcal{K}_X^* \rightarrow \mathcal{K}_X^*/\mathcal{O}_X^*$, but a global section of $\mathcal{K}_X^*/\mathcal{O}_X^*$ is not necessarily the image of a global section of \mathcal{K}_X^* .

Definition. The image of $\Gamma(X, \mathcal{K}_X^*)$ in $\Gamma(X, \mathcal{K}_X^*/\mathcal{O}_X^*)$ is the set of *principal* Cartier divisors. The *Cartier class group* is the quotient

$$\Gamma(X, \mathcal{K}_X^*/\mathcal{O}_X^*) / \text{im } \Gamma(X, \mathcal{K}_X^*)$$

A section $\mathcal{D} \in \Gamma(X, \mathcal{K}_X^*/\mathcal{O}_X^*)$ can be specified by $\{(U_i, f_i)\}$ where the $\{U_i\}$ form an open cover and $f_i \in \mathcal{K}_X^*(U_i)$, such that on $U_i \cap U_j$, the quotient $\frac{f_i}{f_j}$ lies in $\mathcal{O}_X^*(U_i \cap U_j)$.

Let X be Noetherian, integral, separated, and regular in codimension 1. Given a Cartier divisor $\mathcal{D} \in \Gamma(X, \mathcal{K}_X^*/\mathcal{O}_X^*)$, we obtain a Weil divisor as follows. If $Y \subseteq X$ is a prime Weil divisor and its generic point is η_Y , we represent \mathcal{D} by $\{(U_i, f_i)\}$ and set n_Y to be $\nu_Y(f_i)$ for some U_i containing η_Y . Then we obtain the Weil divisor

$$\sum_{Y \subseteq X} n_Y [Y]$$

This is well-defined: if η_Y is contained in both U_i and U_j , the valuations of f_i and f_j differ by $\nu_Y\left(\frac{f_i}{f_j}\right)$, but $\frac{f_i}{f_j}$ is a unit, so has valuation zero. Similarly, one can show that this is independent of the choice of representative of \mathcal{D} .

Proposition. Let X be Noetherian, integral, separated, and regular in codimension 1. Suppose that all local rings $\mathcal{O}_{X,x}$ are unique factorisation domains. Then the association of a Weil divisor to each Cartier divisor is a bijection, and furthermore, is a bijection of principal divisors.

Proof sketch. If R is a unique factorisation domain, then all height 1 prime ideals are principal. If $x \in X$, then $\mathcal{O}_{X,x}$ is a unique factorisation domain by hypothesis, so given a Weil divisor D , we can restrict it to $\text{Spec } \mathcal{O}_{X,x} \rightarrow X$. But on $\text{Spec } \mathcal{O}_{X,x}$, D is given by $\mathbb{V}(f_x)$ as $\mathcal{O}_{X,x}$ is a unique factorisation domain. f_x extends to some neighbourhood U_x containing x , then the f_x can be glued to form a Cartier divisor. This can be checked to be bijective. \square

Given a Cartier divisor \mathcal{D} on X with representative $\{(U_i, f_i)\}$, we can define $L(\mathcal{D}) \subseteq \mathcal{K}_X$ to be the sub- \mathcal{O}_X -module generated on U_i by f_i^{-1} . Note that if $X = \text{Spec } A$ where A is integral, and $\mathcal{D} = \{(X, f)\}$ where $f \in A$, then $A_f \subseteq FF(A)$ is an A -module.

Proposition. The sheaf $L(\mathcal{D})$ is a line bundle.

Proposition. On U_i , we have an isomorphism $\mathcal{O}_{U_i} \rightarrow L(\mathcal{D})|_{U_i}$ given by $1 \mapsto f_i^{-1}$.

Consider $X = \mathbb{P}_k^n$, and let D be the Weil divisor $\mathbb{V}(x_0)$. Let \mathcal{D} be the corresponding Cartier divisor. One can show that $\mathcal{O}_{\mathbb{P}_k^n}(1) \cong L(\mathcal{D})$.

Remark. A line bundle L on X has an ‘inverse’ under the tensor product; that is, defining $L^{-1} = \text{Hom}_{\mathcal{O}_X}(L, \mathcal{O}_X)$, we obtain $L \otimes_{\mathcal{O}_X} L^{-1} = \mathcal{O}_X$. Tensor products of line bundles are also line bundles. If all Weil divisors are Cartier, then $L(\mathcal{D} + \mathcal{E}) = L(\mathcal{D}) \otimes L(\mathcal{E})$.

III. Algebraic Geometry

Definition. The *Picard group* of X is the set of line bundles on X up to isomorphism, which forms an abelian group under the tensor product.

Under mild assumptions, for example assuming that X is integral, the map $\mathcal{D} \mapsto L(\mathcal{D})$ is surjective, and the kernel is exactly the set of principal Cartier divisors.

7. Sheaf cohomology

7.1. Introduction and properties

We have previously seen that if $X = \mathbb{A}^2 \setminus \{(0, 0)\}$, then $\mathcal{O}_X(X) \cong \mathcal{O}_{\mathbb{A}^2}(\mathbb{A}^2) \cong k[x, y]$. Given a topological space X and a sheaf \mathcal{F} of abelian groups, there is a series of *cohomology* groups $H^i(X, \mathcal{F})$ for $i \in \mathbb{N}$. The definition will be omitted. These groups have the following features.

- (i) The group $H^0(X, \mathcal{F})$ is precisely $\Gamma(X, \mathcal{F})$.
- (ii) If $f : Y \rightarrow X$ is continuous, there is an induced map $f^* : H^i(X, \mathcal{F}) \rightarrow H^i(Y, f^{-1}\mathcal{F})$.
- (iii) Given a short exact sequence of sheaves

$$0 \longrightarrow \mathcal{F} \longrightarrow \mathcal{F}' \longrightarrow \mathcal{F}'' \longrightarrow 0$$

we obtain a long exact sequence

$$\begin{array}{ccccccc} 0 & \longrightarrow & H^0(X, \mathcal{F}) & \longrightarrow & H^0(X, \mathcal{F}') & \longrightarrow & H^0(X, \mathcal{F}'') \\ & & & & & \swarrow & \\ & & H^1(X, \mathcal{F}) & \longrightarrow & H^1(X, \mathcal{F}') & \longrightarrow & H^1(X, \mathcal{F}'') \\ & & & & & \swarrow & \\ & & H^2(X, \mathcal{F}) & \longrightarrow & \dots & & \end{array}$$

- (iv) If X is an affine scheme and \mathcal{F} is a quasi-coherent sheaf, then $H^i(X, \mathcal{F}) = 0$ for all $i > 0$.
- (v) Cohomology commutes with taking direct sums of sheaves.
- (vi) If X is a Noetherian separated scheme, then $H^i(X, \mathcal{F})$ can be computed from the sections of \mathcal{F} on an open affine cover $\{U_i\}$ and from the data of the restrictions to $\mathcal{F}(U_i \cap U_j)$, $\mathcal{F}(U_i \cap U_j \cap U_k)$ and so on. This can be done by considering *Čech cohomology*.

7.2. Čech cohomology

Let X be a topological space, and let \mathcal{F} be a sheaf on X . Let $\mathcal{U} = \{U_i\}_{i \in I}$ be a fixed open cover of X , indexed by a well-ordered set I . In this course, we will take $I = \{1, \dots, N\}$, and write $U_{i_0 \dots i_p} = U_{i_0} \cap \dots \cap U_{i_p}$. Čech cohomology attaches data to the triple $(X, \mathcal{F}, \mathcal{U})$. The group of Čech p -cochains is

$$C^p(\mathcal{U}, \mathcal{F}) = \prod_{i_0 < \dots < i_p} \mathcal{F}(U_{i_0 \dots i_p})$$

There is a *differential*

$$d : C^p(\mathcal{U}, \mathcal{F}) \rightarrow C^{p+1}(\mathcal{U}, \mathcal{F})$$

III. Algebraic Geometry

where the i_0, \dots, i_{p+1} component of $d\alpha$ is given by

$$(d\alpha)_{i_0 \dots i_{p+1}} = \sum_{k=0}^{p+1} (-1)^k \alpha_{i_0 \dots \hat{i}_k \dots i_{p+1}} \Big|_{U_{i_0 \dots i_{p+1}}}$$

where \hat{i}_k denotes that the element i_k of the sequence is omitted. One can easily show that $d^2 : C^p \rightarrow C^{p+2}$ is the zero map. Thus, $\{C^p(\mathcal{U}, \mathcal{F})\}_p$ has the structure of a *cochain complex*.

Definition. The *ith Čech cohomology* of $(X, \mathcal{F}, \mathcal{U})$ is the *ith* cohomology group of the cochain complex:

$$\check{H}^i(X, \mathcal{F}) = \frac{\ker(C^i(\mathcal{U}, \mathcal{F}) \xrightarrow{d} C^{i+1}(\mathcal{U}, \mathcal{F}))}{\operatorname{im}(C^{i-1}(\mathcal{U}, \mathcal{F}) \xrightarrow{d} C^i(\mathcal{U}, \mathcal{F}))}$$

Example. Let $X = S^1$ be the usual circle. Let \mathcal{F} be the constant sheaf $\underline{\mathbb{Z}}$; on any connected open set this sheaf has value \mathbb{Z} , and for a general open set with n connected components, this sheaf has value \mathbb{Z}^n . Let $\mathcal{U} = \{U, V\}$ where U, V are obtained by deleting disjoint closed intervals from the circle, giving an open cover with $U, V \cong \mathbb{R}$. We have

$$C^0(\mathcal{U}, \underline{\mathbb{Z}}) = \mathbb{Z}^2$$

as there is one copy of \mathbb{Z} for U and one for V . Also,

$$C^1(\mathcal{U}, \underline{\mathbb{Z}}) = \mathbb{Z}^2$$

given by $\underline{\mathbb{Z}}(U \cap V)$. The differential is $(a, b) \mapsto (b - a, b - a)$, so

$$\check{H}^0(\mathcal{U}, \underline{\mathbb{Z}}) \cong \mathbb{Z} = \ker d$$

and

$$\check{H}^1(\mathcal{U}, \underline{\mathbb{Z}}) \cong \mathbb{Z} = \operatorname{coker} d$$

Remark. (i) These Čech cohomology groups are equal to the corresponding singular cohomology groups of S^1 .

(ii) Note that \check{H} is typically only well-behaved when \mathcal{U} is also well-behaved. That is, $\check{H}^i(\mathcal{U}, \mathcal{F})$ depends on \mathcal{U} and not just X . In the example above, we could have chosen $\mathcal{U} = \{S^1\}$, and in this case, $\check{H}^1(\mathcal{U}, \underline{\mathbb{Z}}) = 0$. Also note that $\underline{\mathbb{Z}}$ is not a quasi-coherent sheaf.

(iii) Let $X = \mathbb{P}_k^1$, $U = X \setminus \{0\}$, $V = X \setminus \{\infty\}$, $\mathcal{U} = \{U, V\}$. Then

$$\check{H}^0(\mathcal{U}, \mathcal{O}_X) = k; \quad \check{H}^1(\mathcal{U}, \mathcal{O}_X) = 0$$

7. Sheaf cohomology

- (iv) Let X be Noetherian and separated, and let $\{U_i\}_{i \in I}$ be an affine cover of X , so all $U_{i_0 \dots i_p}$ are affine. Let \mathcal{F} be a quasi-coherent sheaf on X . Then

$$\check{H}^p(\mathcal{U}, \mathcal{F}) \cong H^p(X, \mathcal{F})$$

and the isomorphism is natural. Thus, in this particular case, the cohomology is easy to calculate by going via Čech cohomology.

Theorem. Let $X = \mathbb{P}_k^n$ and $\mathcal{F} = \bigoplus_{d \in \mathbb{Z}} \mathcal{O}_{\mathbb{P}_k^n}(d)$. Then there are isomorphisms of graded k -vector spaces

- (i) $H^0(X, \mathcal{F}) \cong k[x_0, \dots, x_n]$;
- (ii) $H^n(X, \mathcal{F}) \cong \frac{1}{x_0 \dots x_n} k[x_0^{-1}, \dots, x_n^{-1}]$;
- (iii) $H^p(X, \mathcal{F}) = 0$ for $p \neq 0, n$.

In particular, $H^0(\mathbb{P}_k^n, \mathcal{O}(d))$ has dimension $\binom{n+d}{d}$, and $H^n(\mathbb{P}_k^n, \mathcal{O}(d))$ has dimension $\binom{-d-1}{n}$.

Proof. We prove this result using Čech cohomology. Part (i) follows from earlier discussions, as $H^0(X, \mathcal{F}) = \bigoplus_{d \in \mathbb{Z}} \Gamma(\mathbb{P}_k^n, \mathcal{O}(d))$.

Part (ii). Consider the standard cover \mathcal{U} of \mathbb{P}_k^n by affines $U_i = \mathbb{V}(x_i)^c$. Observe that

$$\mathcal{F}(U_{i_0 \dots i_p}) = k[x_0, \dots, x_n]_{x_{i_0} \dots x_{i_p}}$$

This k -module is spanned by monomials $x_0^{k_0} \dots x_n^{k_n}$ where $k_{i_0}, \dots, k_{i_p} \in \mathbb{Z}$ and the other coefficients are nonnegative. In the associated Čech complex, we have

$$\check{C}^{n-1} = \bigoplus_{i=0}^n k[x_0, \dots, x_n]_{x_0 \dots \hat{x}_i \dots x_n}; \quad \check{C}^n = k[x_0, \dots, x_n]_{x_0 \dots x_n}$$

Since \mathcal{U} contains only $n+1$ elements, \check{C}^{n+1} vanishes. Thus,

$$\begin{aligned} H^n(\mathbb{P}_k^n, \mathcal{F}) &= \check{H}^n(\mathcal{U}, \mathcal{F}) \\ &= \frac{\check{C}^n}{\text{im}(\check{C}^{n-1} \rightarrow \check{C}^n)} \\ &= \frac{\text{span}_k \{x_0^{k_0} \dots x_n^{k_n} \mid k_i \in \mathbb{Z}\}}{\text{span}_k \{x_0^{k_0} \dots x_n^{k_n} \mid \text{at least one } k_i \geq 0\}} \end{aligned}$$

as required.

Part (iii). We will use the long exact sequence associated to a short exact sequence of sheaves and use induction on the dimension n . First, observe that \mathbb{P}_k^{n-1} is isomorphic to the closed

III. Algebraic Geometry

subscheme $\mathbb{V}(x_0) \subseteq \mathbb{P}_k^n$. Let $i : \mathbb{P}_k^{n-1} \rightarrow \mathbb{P}_k^n$ be the inclusion. Recall that $\mathcal{O}_{\mathbb{P}_k^n}(-1) = L(-H)$ where $H = \mathbb{V}(x_0)$. By a result on the example sheets, we obtain the *ideal sheaf sequence*

$$0 \longrightarrow \mathcal{O}_{\mathbb{P}_k^n}(-1) \longrightarrow \mathcal{O}_{\mathbb{P}_k^n} \longrightarrow i_* \mathcal{O}_{\mathbb{P}_k^{n-1}} \longrightarrow 0$$

where the map $\mathcal{O}_{\mathbb{P}_k^n}(-1) \rightarrow \mathcal{O}_{\mathbb{P}_k^n}$ is given by multiplication by x_0 . This is analogous to the fact that for an ideal I of a ring A , we have a short exact sequence

$$0 \longrightarrow I \longrightarrow A \longrightarrow A/I \longrightarrow 0$$

We obtain an associated long exact sequence for the homology. Assuming the result for dimension up to $n - 1$, we can break this into three smaller exact sequences.

$$0 \rightarrow H^0(\mathbb{P}_k^n, \mathcal{F}) \xrightarrow{\cdot x_0} H^0(\mathbb{P}_k^n, \mathcal{F}) \rightarrow H^0(\mathbb{P}_k^{n-1}, \mathcal{F}_{\mathbb{P}_k^{n-1}}) \rightarrow H^1(\mathbb{P}_k^n, \mathcal{F}) \xrightarrow{\cdot x_0} H^1(\mathbb{P}_k^n, \mathcal{F}) \rightarrow 0 \quad (a)$$

where $\mathcal{F}_{\mathbb{P}_k^{n-1}} = \bigoplus_{d \in \mathbb{Z}} \mathcal{O}_{\mathbb{P}_k^{n-1}}(d)$;

$$0 \longrightarrow H^p(\mathbb{P}_k^n, \mathcal{F}) \xrightarrow{\cdot x_0} H^p(\mathbb{P}_k^n, \mathcal{F}) \longrightarrow 0 \quad (b)$$

for $1 < p < n - 1$; and

$$0 \rightarrow H^{n-1}(\mathbb{P}_k^n, \mathcal{F}) \xrightarrow{\cdot x_0} H^{n-1}(\mathbb{P}_k^n, \mathcal{F}) \rightarrow H^{n-1}(\mathbb{P}_k^{n-1}, \mathcal{F}_{\mathbb{P}_k^{n-1}}) \rightarrow H^n(\mathbb{P}_k^n, \mathcal{F}) \xrightarrow{\cdot x_0} H^n(\mathbb{P}_k^n, \mathcal{F}) \rightarrow 0 \quad (c)$$

By using (a) and (c), we observe that (b) is also exact for $p = 1$ and $p = n - 1$ by explicit computation in the Čech complex. Now, multiplication by x_0 makes $H^p(\mathbb{P}_k^n, \mathcal{F})$ into a $k[x_0]$ -module. We will calculate the localisation $H^p(\mathbb{P}_k^n, \mathcal{F})_{x_0}$. As localisation is exact, $H^p(\mathbb{P}_k^n, \mathcal{F})_{x_0} = H^p(U_0, \mathcal{F}|_{U_0})$. But the right-hand side vanishes for $p > 0$ as U_0 is affine. Hence, for any $\alpha \in H^p(\mathbb{P}_k^n, \mathcal{F})$, there exists k such that $x_0^k \alpha = 0$. But multiplication by x_0 is an isomorphism on cohomology by (b), so in fact $H^p(\mathbb{P}_k^n, \mathcal{F}) = 0$ for all $1 \leq p \leq n - 1$. \square

Given the exact sequence

$$0 \longrightarrow \mathcal{O}_{\mathbb{P}_k^n}(-1) \longrightarrow \mathcal{O}_{\mathbb{P}_k^n} \longrightarrow i_* \mathcal{O}_{\mathbb{P}_k^{n-1}} \longrightarrow 0$$

taking the tensor product with $\mathcal{O}_{\mathbb{P}_k^n}(d)$, one can show that we obtain an exact sequence

$$0 \longrightarrow \mathcal{O}_{\mathbb{P}_k^n}(d-1) \longrightarrow \mathcal{O}_{\mathbb{P}_k^n}(d) \longrightarrow i_* \mathcal{O}_{\mathbb{P}_k^{n-1}}(d) \longrightarrow 0$$

Note that $\mathcal{O}_{\mathbb{P}_k^n}(d)$ is locally free.

Let X be proper over $\text{Spec } k$ and let \mathcal{F} be a coherent sheaf on X .

Remark. (i) We have observed that $H^0(X, \mathcal{F})$ is a finite-dimensional k -vector space. The same holds for all $H^p(X, \mathcal{F})$.

(ii) If X has dimension n , then $H^p(X, \mathcal{F})$ vanishes for $p > n$. Thus, given (X, \mathcal{F}) , there are finitely many numbers $h^p(X, \mathcal{F}) = \dim_k H^p(X, \mathcal{F})$.

Definition. The *Euler characteristic* of \mathcal{F} is

$$\chi(\mathcal{F}) = \sum_{p=0}^{\infty} (-1)^p h^p(X, \mathcal{F})$$

Suppose that

$$0 \longrightarrow \mathcal{F} \longrightarrow \mathcal{F}' \longrightarrow \mathcal{F}'' \longrightarrow 0$$

is an exact sequence of such sheaves. Then the associated long exact sequence gives

$$\chi(\mathcal{F}') = \chi(\mathcal{F}) + \chi(\mathcal{F}'')$$

7.3. Choice of cover

Given a Noetherian separated scheme X , a quasi-coherent sheaf \mathcal{F} on X , and an open affine cover \mathcal{U} which we typically take to be finite, we can construct the Čech cohomology $\check{H}^i(\mathcal{U}, \mathcal{F})$. In this subsection, we show that the Čech cohomology is independent of the choice of cover in this case.

Theorem. Let X be affine and let \mathcal{F} be quasi-coherent. For any finite cover \mathcal{U} of X by affine opens, the groups $\check{H}^i(\mathcal{U}, \mathcal{F})$ vanish for $i > 0$.

Proof. Define the ‘sheafified’ Čech complex as follows.

$$\mathcal{C}^p(\mathcal{F}) = \prod_{i_0 < \dots < i_p} i_* \mathcal{F} \Big|_{U_{i_0 \dots i_p}}$$

where $i : U_{i_0 \dots i_p} \rightarrow X$ is the inclusion. Then the $\mathcal{C}^p(\mathcal{F})$ are quasi-coherent sheaves. By taking global sections,

$$\Gamma(X, \mathcal{C}^p(\mathcal{F})) = C^p(\mathcal{F})$$

where $C^p(\mathcal{F})$ is the usual group of Čech p -cochains. The same formula used to build the Čech complex gives differentials

$$\mathcal{C}^p(\mathcal{F}) \rightarrow \mathcal{C}^{p+1}(\mathcal{F})$$

as a morphism of sheaves. We intend to show that the usual Čech complex

$$C^0(\mathcal{F}) \longrightarrow C^1(\mathcal{F}) \longrightarrow C^2(\mathcal{F}) \longrightarrow \dots$$

III. Algebraic Geometry

is exact. By a result on the example sheet, on affines, taking local sections preserves exactness. Thus, it suffices to prove that

$$\mathcal{C}^0(\mathcal{F}) \longrightarrow \mathcal{C}^1(\mathcal{F}) \longrightarrow \mathcal{C}^2(\mathcal{F}) \longrightarrow \dots$$

is an exact sequence of sheaves. However, the exactness of this sequence can be checked locally on stalks. Let $q \in X$, and suppose $q \in U_j$. Now define the map on stalks $\kappa : \mathcal{C}_q^p(\mathcal{F}) \rightarrow \mathcal{C}_q^{p-1}(\mathcal{F})$, where for a cochain α , the $(i_0 \dots i_{p-1})$ -component of $\kappa(\alpha)$ is equal to the $(ji_0 \dots i_{p-1})$ -component of α , where by convention if $ji_0 \dots i_{p-1}$ is not in increasing order, but $\sigma \in S_{p+1}$ brings it into increasing order and σ has sign -1 , we instead take the negation of the component. By direct calculation, one can show that $d\kappa + \kappa d = \text{id}$ on \mathcal{C}^p for all p .

We can now verify exactness at each stalk. We know that $\text{im}(\mathcal{C}^{p-1} \rightarrow \mathcal{C}^p) \subseteq \ker(\mathcal{C}^p \rightarrow \mathcal{C}^{p+1})$. Conversely, if $\alpha \in \ker(\mathcal{C}^p \rightarrow \mathcal{C}^{p+1})$, then

$$\alpha = (\kappa d + d\kappa)(\alpha) = d(\kappa\alpha) \in \text{im}(\mathcal{C}^{p-1} \rightarrow \mathcal{C}^p)$$

□

Lemma. Let X be a scheme and let \mathcal{F} be a quasi-coherent sheaf on X . Let $\mathcal{U} = \{U_1, \dots, U_k\}$ and $\tilde{\mathcal{U}} = \{U_0, \dots, U_k\}$. That $\check{H}^i(\mathcal{U}, \mathcal{F})$ and $\check{H}^i(\tilde{\mathcal{U}}, \mathcal{F})$ are naturally isomorphic.

Proof sketch. Let $C^p(\mathcal{F})$ and $\tilde{C}^p(\mathcal{F})$ be the cochain groups for $\mathcal{U}, \tilde{\mathcal{U}}$ respectively. There are maps $\tilde{C}^p(\mathcal{F}) \rightarrow C^p(\mathcal{F})$ given by dropping the U_0 data. To make this precise, observe that $\tilde{\alpha} \in \tilde{C}^p(\mathcal{F})$ can be viewed as a pair (α, α_0) where $\alpha \in C^p(\mathcal{F})$ and $\alpha_0 \in C^{p-1}$ for the sheaf $\mathcal{F}|_{U_0}$ with open cover $\mathcal{U}|_{U_0}$. These maps commute with the differentials, so we have an induced map $\check{H}^i(\tilde{\mathcal{U}}, \mathcal{F}) \rightarrow \check{H}^i(\mathcal{U}, \mathcal{F})$. By reducing to a calculation on the affine U_0 , we can deduce using the previous result that this induced map is surjective and injective. □

Corollary. $\check{H}^i(\mathcal{U}, \mathcal{F})$ is independent of the choice of \mathcal{U} .

Proof. If $\mathcal{U}, \tilde{\mathcal{U}}$ are two finite open covers by affines, we can interpolate between them by using $\mathcal{U} \cup \tilde{\mathcal{U}}$ and use the previous result. □

7.4. Further topics in cohomology

- (i) Let $X_d \subseteq \mathbb{P}_k^3$ be the vanishing locus of a homogeneous polynomial f_d of degree $d \neq 2$. Then X_d is not isomorphic to a product over $\text{Spec } k$ of schemes of dimension 1. Conversely, X_2 can be isomorphic to $\mathbb{P}_k^1 \times_{\text{Spec } k} \mathbb{P}_k^1$, using the Segre embedding. This is a consequence of the sheaf Künneth formula, and in particular, the fact that $h^1(X_d, \mathcal{O}_{X_d}) = 0$.
- (ii) The different X_d are non-isomorphic as schemes. This follows from calculating $\chi(X_d)$.

7. Sheaf cohomology

- (iii) One next direction in cohomology is *duality theory*. Given a closed immersion $i : Z \subseteq X$, the *ideal sheaf* I_Z is the kernel of the map $i^* : \mathcal{O}_X \rightarrow \mathcal{O}_Z$, which is a coherent sheaf on X . The *conormal sheaf* to the closed immersion i , denoted $N_{Z/X}^\vee$, is given by $i^*(I_Z/I_Z^2)$, where I_Z^2 is the sheafification of the presheaf $U \mapsto I_Z(U)^2$. If $X \rightarrow S$ is separated, then the *cotangent sheaf* is

$$\Omega_{X/S} = N_{\Delta_{X/S}}^\vee$$

A scheme X over $\text{Spec } k$ is called *nonsingular* if Ω_X is locally free. The *dualising sheaf* ω_X is the sheafification of $U \mapsto \bigwedge^{\dim X} \Omega_X(U)$.

Theorem (Serre duality). If X is as above and has dimension n , then if \mathcal{F} is a locally free \mathcal{O}_X -module, there is an isomorphism of cohomology groups

$$H^i(X, \mathcal{F}) \rightarrow H^{n-1}(X, \mathcal{F}^\vee \otimes \omega_X)^\vee$$

where

$$\mathcal{F}^\vee = \text{Hom}_{\mathcal{O}_X}(\mathcal{F}, \mathcal{O}_X)$$

IV. Model Theory and Non-Classical Logic

Lectured in Michaelmas 2023 by DR. J. SIQUEIRA

(Course description goes here.)

Contents

1.	Substructures	220
1.1.	Notation	220
1.2.	Homomorphisms and substructures	220
1.3.	Elementary equivalence	222
1.4.	Categorical and complete theories	223
1.5.	Tarski–Vaught test	223
1.6.	Universal theories and the method of diagrams	224
2.	Quantifier elimination	227
2.1.	Skolem functions	227
2.2.	Skolemisation theorem	228
2.3.	Elimination sets	229
2.4.	Amalgamation	231
2.5.	Inductive classes	234
2.6.	Characterisations of quantifier elimination	235
2.7.	Applications	239
3.	Ultraproducts	242
3.1.	Products	242
3.2.	Lattices	242
3.3.	Filters	243
3.4.	Łoś’ theorem	244
4.	Types	247
4.1.	Definitions	247
4.2.	Stone spaces	249
4.3.	Isolated points	249
4.4.	Omitting types	250
5.	Indiscernibles	252
5.1.	Introduction	252
5.2.	Existence of Ehrenfeucht–Mostowski functors	253
6.	Intuitionistic logic and lambda calculi	255
6.1.	The Brouwer–Heyting–Kolmogorov interpretation	255
6.2.	Natural deduction	256
6.3.	The simply typed lambda calculus	258
6.4.	Basic properties	260
6.5.	The normalisation theorems	261
7.	Intuitionistic semantics	264
7.1.	Propositions as types	264
7.2.	Full simply typed lambda calculus	265

7.3.	Heyting semantics	267
7.4.	Kripke semantics	269

1. Substructures

1.1. Notation

The interpretation of a function symbol f in a model \mathcal{M} is denoted by $f^{\mathcal{M}}$, and similarly the interpretation of a relation symbol R in \mathcal{M} is denoted by $R^{\mathcal{M}}$. If \mathcal{M} is an \mathcal{L} -structure, and $A \subseteq \mathcal{M}$ is a subset, we will write \mathcal{L}_A for the language obtained by adding a new constant symbol a to the signature of \mathcal{L} for each element a of A . Then \mathcal{M} is naturally an \mathcal{L}_A -structure by interpreting the constants in the obvious way. We will allow for the empty set to be an \mathcal{L} -structure.

1.2. Homomorphisms and substructures

Definition. Let \mathcal{M} and \mathcal{N} be \mathcal{L} -structures. An \mathcal{L} -homomorphism is a map $\eta : \mathcal{M} \rightarrow \mathcal{N}$ that preserves the interpretations of the symbols in the language: given $\mathbf{a} = (a_1, \dots, a_n) \in \mathcal{M}^n$,

- (i) for all function symbols f of arity n , we have that

$$\eta(f^{\mathcal{M}}(\mathbf{a})) = f^{\mathcal{N}}(\eta(\mathbf{a}))$$

- (ii) for all relation symbols r of arity n , we have that

$$\mathbf{a} \in R^{\mathcal{M}} \iff \eta(\mathbf{a}) \in R^{\mathcal{N}}$$

An injective \mathcal{L} -homomorphism is called an \mathcal{L} -embedding. An invertible \mathcal{L} -homomorphism is called an \mathcal{L} -isomorphism.

Definition. If $\mathcal{M} \subseteq \mathcal{N}$ and the inclusion map is an \mathcal{L} -homomorphism, we say that \mathcal{M} is a *substructure* of \mathcal{N} , and that \mathcal{N} is an *extension* of \mathcal{M} . We will typically use the notation $\mathcal{M} \subseteq \mathcal{N}$ to indicate that \mathcal{M} is a substructure of \mathcal{N} when both are \mathcal{L} -structures, not just that it is a subset.

Example. (i) Let \mathcal{L} be the language of groups. Then $(\mathbb{N}, +, 0)$ is a substructure of $(\mathbb{Z}, +, 0)$, but it is not a subgroup.

- (ii) If \mathcal{M} is an \mathcal{L} -structure, X is the domain of a substructure of \mathcal{M} if and only if it is closed under the interpretations of all function symbols. The forward implication is clear. If f is a function symbol of arity n and X is closed under $f^{\mathcal{M}}$, $f^{\mathcal{M}}|_{X^n}$ is a function $X^n \rightarrow X$ interpreting f on the domain X , as required. In particular, any substructure should also contain all of the constants in the language.

- (iii) The substructure *generated* by a subset $X \subseteq \mathcal{M}$ is given by the smallest set that contains X and is closed under the interpretations of all function symbols in \mathcal{M} . This is denoted $\langle X \rangle_{\mathcal{M}}$, and one can check that for infinite \mathcal{L} (but not necessarily infinite signature),

$$|\langle X \rangle_{\mathcal{M}}| \leq |X| + |\mathcal{L}|$$

We prove this by iteratively closing up X by applying interpretations of function symbols to elements of X , and then taking the union of the resulting sets. At each stage, for each function symbol f of arity n , we add at most $|X|^n \leq |X| \cdot \aleph_0$ new elements. So in a single stage, we add at most $|X| \cdot \aleph_0 \cdot |\mathcal{L}| = |X| \cdot |\mathcal{L}|$ new elements to X . Repeating this ω times, the final set has size at most

$$\begin{aligned} |X| + |X| \cdot |\mathcal{L}| + |X| \cdot |\mathcal{L}|^2 + \dots &= |X|(1 + |\mathcal{L}| + |\mathcal{L}|^2 + \dots) \\ &\leq |X|(|\mathcal{L}| + |\mathcal{L}| + |\mathcal{L}| + \dots) \\ &= |X| \cdot |\mathcal{L}| \cdot \aleph_0 \\ &= |X| \cdot |\mathcal{L}| \end{aligned}$$

We say that \mathcal{M} is *finitely generated* if there exists a finite subset $X \subseteq \mathcal{M}$ such that $\mathcal{M} = \langle X \rangle_{\mathcal{M}}$.

(iv) Consider

$$(\mathbb{R}, \cdot, -1) \models \neg \exists x. (x^2 = -1)$$

But it has an extension $(\mathbb{C}, \cdot, -1)$ that does not model this sentence.

Proposition. Let $\varphi(\mathbf{x})$ be a quantifier-free \mathcal{L} -formula with n free variables. Let \mathcal{M} be an \mathcal{L} -structure, and let \mathbf{a} be an n -tuple in \mathcal{M} . Then for every extension \mathcal{N} of \mathcal{M} ,

$$\mathcal{M} \models \varphi(\mathbf{a}) \iff \mathcal{N} \models \varphi(\mathbf{a})$$

Proof. We proceed by induction on the structure of formulae. First, we show that if $t(\mathbf{x})$ is a term with k free variables, then

$$t^{\mathcal{M}}(\mathbf{b}) = t^{\mathcal{N}}(\mathbf{b})$$

for all $\mathbf{b} \in \mathcal{M}^k$. It is clearly the case if $t = x_i$ is a variable, as both structures interpret $t(\mathbf{b})$ as b_i . Suppose t is a term of the form $t = f(q_1, \dots, q_\ell)$ for f a function symbol of arity ℓ and the q_i are terms. By the inductive hypothesis we have

$$q_i^{\mathcal{M}}(\mathbf{b}) = q_i^{\mathcal{N}}(\mathbf{b})$$

Therefore,

$$\begin{aligned} t^{\mathcal{M}}(\mathbf{b}) &= f^{\mathcal{M}}(q_1^{\mathcal{M}}(\mathbf{b}), \dots, q_\ell^{\mathcal{M}}(\mathbf{b})) \\ &= f^{\mathcal{N}}(q_1^{\mathcal{M}}(\mathbf{b}), \dots, q_\ell^{\mathcal{M}}(\mathbf{b})) \\ &= f^{\mathcal{N}}(q_1^{\mathcal{N}}(\mathbf{b}), \dots, q_\ell^{\mathcal{N}}(\mathbf{b})) \\ &= t^{\mathcal{N}}(\mathbf{b}) \end{aligned}$$

Thus terms are interpreted the same way in both models. For terms t_1, t_2 with the same free variables \mathbf{x} , then for any choice of \mathbf{a} ,

$$\begin{aligned} \mathcal{M} \models (t_1(\mathbf{x}) = t_2(\mathbf{x})) &\iff t_1^{\mathcal{M}}(\mathbf{a}) = t_2^{\mathcal{M}}(\mathbf{a}) \\ &\iff t_1^{\mathcal{N}}(\mathbf{a}) = t_2^{\mathcal{M}}(\mathbf{b}) \\ &\iff \mathcal{N} \models (t_1(\mathbf{x}) = t_2(\mathbf{x})) \end{aligned}$$

IV. Model Theory and Non-Classical Logic

Let R be a relation symbol of arity n , and let t_1, \dots, t_n be terms with the same free variables \mathbf{x} .

$$\begin{aligned}\mathcal{M} \models R(t_1(\mathbf{x}), \dots, t_n(\mathbf{x})) &\iff (t_1^{\mathcal{M}}(\mathbf{a}), \dots, t_n^{\mathcal{M}}(\mathbf{a})) \in R^{\mathcal{M}} \\ &\iff (t_1^{\mathcal{N}}(\mathbf{a}), \dots, t_n^{\mathcal{N}}(\mathbf{a})) \in R^{\mathcal{N}} \\ &\iff (t_1^{\mathcal{N}}(\mathbf{a}), \dots, t_n^{\mathcal{N}}(\mathbf{a})) \in R^{\mathcal{N}} \\ &\iff \mathcal{N} \models R(t_1(\mathbf{x}), \dots, t_n(\mathbf{x}))\end{aligned}$$

So the result holds for all atomic formulae. For connectives, note that

$$\begin{aligned}\mathcal{M} \models \neg\varphi &\iff \mathcal{M} \not\models \varphi \\ &\iff \mathcal{N} \not\models \varphi \\ &\iff \mathcal{N} \models \neg\varphi\end{aligned}$$

and

$$\begin{aligned}\mathcal{M} \models \varphi \wedge \psi &\iff (\mathcal{M} \models \varphi) \wedge (\mathcal{M} \models \psi) \\ &\iff (\mathcal{N} \models \varphi) \wedge (\mathcal{N} \models \psi) \\ &\iff \mathcal{N} \models \varphi \wedge \psi\end{aligned}$$

As quantifier-free formulae can be built out of atomic formulae, negation, and conjunction, we have completed the proof. \square

1.3. Elementary equivalence

Definition. Structures \mathcal{M}, \mathcal{N} are called *elementarily equivalent* if for every \mathcal{L} -sentence,

$$\mathcal{M} \models \varphi \iff \mathcal{N} \models \varphi$$

A map $f : \mathcal{M} \rightarrow \mathcal{N}$ is an *elementary embedding* if it is injective, and for all \mathcal{L} -formulae $\varphi(x_1, \dots, x_n)$ and elements $m_1, \dots, m_n \in \mathcal{M}$, we have

$$\mathcal{M} \models \varphi(m_1, \dots, m_n) \iff \mathcal{N} \models \varphi(f(m_1), \dots, f(m_n))$$

If there is an elementary embedding between two structures, they are elementarily equivalent. If \mathcal{M} and \mathcal{N} are elementarily equivalent, we write $\mathcal{M} \equiv \mathcal{N}$.

Remark. If \mathcal{M} and \mathcal{N} are \mathcal{L} -structures, and $\mathbf{m} \in \mathcal{M}, \mathbf{n} \in \mathcal{N}$ are ordered tuples of the same length k , then by

$$(\mathcal{M}, \mathbf{m}) \equiv (\mathcal{N}, \mathbf{n})$$

we view $(\mathcal{M}, \mathbf{m})$ and $(\mathcal{N}, \mathbf{n})$ as structures over \mathcal{L} with k additional constants, interpreting these new constants as the elements of \mathbf{m} and \mathbf{n} respectively.

Proposition. If $\mathcal{M} \cong \mathcal{N}$, then $\mathcal{M} \equiv \mathcal{N}$.

This can be easily shown by induction. The converse is generally not true, for example if the structures are infinite.

Definition. A substructure $\mathcal{M} \subseteq \mathcal{N}$ is an *elementary substructure* if the inclusion map is an elementary embedding. In this case, we also say that \mathcal{N} is an *elementary extension* of \mathcal{M} . We write $\mathcal{M} \leq \mathcal{N}$.

1.4. Categorical and complete theories

Recall that a theory \mathcal{T} is *complete* if either $\mathcal{T} \vdash \varphi$ or $\mathcal{T} \vdash \neg\varphi$ for all sentences φ . Then any two models of a complete theory are elementarily equivalent, but they may have different cardinalities.

Definition. A theory \mathcal{T} is *model-complete* if every embedding between models of \mathcal{T} is elementary.

Definition. Let κ be an infinite cardinal. A theory \mathcal{T} is κ -*categorical* if all models of \mathcal{T} of cardinality κ are isomorphic.

It turns out that if theory on a countable language is categorical for some uncountable cardinal, then it is categorical for all infinite cardinals.

Proposition (Vaught's test). Let \mathcal{T} be a consistent \mathcal{L} -theory that has no finite models. If \mathcal{T} is κ -categorical for some infinite $\kappa \geq |\mathcal{L}|$, then \mathcal{T} is complete.

Proof. Suppose there is some φ such that $\mathcal{T} \not\vdash \varphi$ and $\mathcal{T} \not\vdash \neg\varphi$. Then $\mathcal{T} \cup \{\varphi\}$ and $\mathcal{T} \cup \{\neg\varphi\}$ are consistent theories, so have models. As \mathcal{T} has no finite models, these two models are infinite. In fact, by the Löwenheim–Skolem theorem, the models can be forced to have size κ . But these models are in particular models of \mathcal{T} , so they must be isomorphic. Since they are isomorphic, they are elementarily equivalent. But the models disagree on the truth value of φ , giving a contradiction. \square

Example. (i) Any two countable dense linear orders are isomorphic, so the theory of dense linear orders without endpoints is \aleph_0 -categorical. Thus, by Vaught's test, the theory DLO of dense linear orders without endpoints is complete.

(ii) Let F be a field. The theory of infinite (not infinite-dimensional) F -vector spaces is κ -categorical for $\kappa > |F|$. Hence, the theory is complete.

1.5. Tarski–Vaught test

Proposition. Let \mathcal{N} be an \mathcal{L} -structure, and let $M \subseteq \mathcal{N}$. Then M is the domain of an elementary substructure if and only if for any formula $\varphi(x, \mathbf{t})$ and tuple $\mathbf{m} \in M$, if there exists a witness $n \in \mathcal{N}$ such that $\mathcal{N} \models \varphi(n, \mathbf{m})$, then there is a witness $\hat{n} \in M$ such that $\mathcal{N} \models \varphi(\hat{n}, \mathbf{m})$.

IV. Model Theory and Non-Classical Logic

Proof. If M is the domain of an elementary substructure \mathcal{M} , then $\mathcal{N} \models \exists x. \varphi(x, \mathbf{m})$ implies that $\mathcal{M} \models \exists x. \varphi(x, \mathbf{m})$. Thus $\mathcal{M} \models \varphi(\hat{m}, \mathbf{m})$ for some $\hat{m} \in M$. But then $\mathcal{N} \models \varphi(\hat{m}, \mathbf{m})$, as required.

For the other implication, if $M \subseteq \mathcal{N}$ has the stated property, we first show that M is closed under the interpretation of function symbols. Consider the formulae $\varphi_f(x, \mathbf{t}) = (x = f(\mathbf{t}))$ for each function symbol f in \mathcal{L} . Then for any $\mathbf{m} \in M$, there exists $n \in \mathcal{N}$ such that $\mathcal{N} \models n = f(\mathbf{m})$, but then by hypothesis, there exists $\hat{m} \in M$ such that $\mathcal{N} \models \hat{m} = f(\mathbf{m})$. Thus $f(\mathbf{m}) = \hat{m} \in M$. Interpreting relation symbols on M in the obvious way, we turn M into an \mathcal{L} -structure \mathcal{M} , which is clearly a substructure of \mathcal{N} .

It now remains to show that the substructure \mathcal{M} of \mathcal{N} is elementary. This follows from induction over the number of quantifiers in formulae, noting that the truth values of quantifier-free formulae are always preserved under any extension. \square

1.6. Universal theories and the method of diagrams

Definition. A formula φ is *universal* if it is of the form $\forall \mathbf{x}. \psi(\mathbf{x}, \mathbf{y})$ where ψ is quantifier-free. A theory is *universal* if all its axioms are universal sentences.

Definition. Let \mathcal{N} be an \mathcal{L} -structure. We define the *diagram* of \mathcal{N} to be the set

$$\text{Diag } \mathcal{N} = \{\varphi(n_1, \dots, n_k) \mid \varphi \text{ is a quantifier-free } \mathcal{L}_{\mathcal{N}}\text{-formula, } \mathcal{N} \models \varphi(n_1, \dots, n_k)\}$$

The *elementary diagram* of \mathcal{N} is

$$\text{Diag}_{\text{el}} \mathcal{N} = \{\varphi(n_1, \dots, n_k) \mid \varphi \text{ is an } \mathcal{L}_{\mathcal{N}}\text{-formula, } \mathcal{N} \models \varphi(n_1, \dots, n_k)\}$$

The diagram of a group is a slight generalisation of its multiplication table. Note that a model of a diagram is the same as an extension, and a model of an elementary diagram is the same as an elementary extension.

Lemma. Let \mathcal{T} be a consistent theory, and let \mathcal{T}_{\forall} be the theory of universal sentences proven by \mathcal{T} . If \mathcal{N} is a model of \mathcal{T}_{\forall} , then $\mathcal{T} \cup \text{Diag } \mathcal{N}$ is consistent.

Proof. Suppose $\mathcal{T} \cup \text{Diag } \mathcal{N}$ is inconsistent. As \mathcal{T} is consistent, by compactness there must be a finite number of sentences in the diagram $\text{Diag } \mathcal{N}$ that are inconsistent with \mathcal{T} . Taking the conjunction, we can reduce to the case where there is a single sentence $\varphi(\mathbf{n})$ that is inconsistent with \mathcal{T} . Then as $\mathcal{T} \cup \{\varphi(\mathbf{n})\}$ is inconsistent, $\mathcal{T} \vdash \neg\varphi(\mathbf{n})$. Since \mathcal{T} has nothing to say about the new constants \mathbf{n} , we must in fact have $\mathcal{T} \vdash \forall \mathbf{x}. \neg\varphi(\mathbf{x})$. This is a universal consequence of \mathcal{T} , so by assumption \mathcal{N} models it, giving a contradiction. \square

Corollary (Tarski, Łoś). An \mathcal{L} -theory \mathcal{T} has a universal axiomatisation if and only if it is preserved under substructures. That is, if $\mathcal{M} \subseteq \mathcal{N}$ are substructures and $\mathcal{M} \models \mathcal{T}$ then $\mathcal{N} \models \mathcal{T}$. Dually, a theory has an existential axiomatisation if and only if it is preserved under extensions.

Proof. One direction is clear. Suppose \mathcal{T} is preserved under taking substructures. If $\mathcal{N} \models \mathcal{T}$, then $\mathcal{N} \models \mathcal{T}_\forall$; we show that the converse also holds. By the previous proposition, $\mathcal{T} \cup \text{Diag } \mathcal{N}$ is consistent. Let \mathcal{N}^* be a model of this theory. So \mathcal{N}^* is an extension of \mathcal{N} , and also models \mathcal{T} . But as \mathcal{T} is preserved under substructures, \mathcal{N} must model \mathcal{T} . \square

We can show much more with the same method.

Theorem (elementary amalgamation theorem). Let \mathcal{M}, \mathcal{N} be \mathcal{L} -structures, and $\mathbf{m} \in \mathcal{M}, \mathbf{n} \in \mathcal{N}$ be tuples of the same size such that $(\mathcal{M}, \mathbf{m}) \equiv (\mathcal{N}, \mathbf{n})$. Then there is an elementary extension \mathcal{K} of \mathcal{M} and an elementary embedding $g : \mathcal{N} \rightarrow \mathcal{K}$ mapping each n_i to m_i .

Proof. Replacing \mathcal{N} with an isomorphic copy if required, we can assume $\mathbf{m} = \mathbf{n}$, and that \mathcal{M} and \mathcal{N} have no other common elements. We show that the theory

$$\mathcal{T} = \text{Diag}_{\text{el}} \mathcal{M} \cup \text{Diag}_{\text{el}} \mathcal{N}$$

is consistent, using compactness. Suppose that Φ is a finite subset of sentences in \mathcal{T} , which of course includes only finitely many sentences in $\text{Diag}_{\text{el}} \mathcal{N}$. Let the conjunction of those sentences be written as $\varphi(\mathbf{m}, \mathbf{k})$, where $\varphi(\mathbf{x}, \mathbf{y})$ is an $\mathcal{L}_{\mathcal{N}}$ -formula, and \mathbf{k} are pairwise distinct elements of $\mathcal{N} \setminus \mathbf{m}$. If Φ is inconsistent, then

$$\text{Diag}_{\text{el}} \mathcal{M} \vdash \neg \varphi(\mathbf{m}, \mathbf{k})$$

Since the elements of \mathbf{k} are distinct and not in \mathcal{M} , we in fact have

$$\text{Diag}_{\text{el}} \mathcal{M} \vdash \forall \mathbf{y}. \neg \varphi(\mathbf{m}, \mathbf{y})$$

In particular,

$$(\mathcal{M}, \mathbf{m}) \models \forall \mathbf{y}. \neg \varphi(\mathbf{m}, \mathbf{y})$$

By hypothesis,

$$(\mathcal{N}, \mathbf{n}) \models \forall \mathbf{y}. \neg \varphi(\mathbf{m}, \mathbf{y})$$

This is a contradiction, as $\varphi(\mathbf{m}, \mathbf{k}) \in \text{Diag}_{\text{el}} \mathcal{N}$. Hence \mathcal{T} is consistent. Take \mathcal{K} to be the \mathcal{L} -reduct of a model of \mathcal{T} . \square

We can also use this technique to constrain the size of a model.

Theorem (Löwenheim–Skolem theorem). Let \mathcal{M} be an infinite \mathcal{L} -structure. Let $\kappa \geq |\mathcal{L}|$ be an infinite cardinal. Then,

- (i) if $\kappa < |\mathcal{M}|$, there is an elementary substructure of \mathcal{M} of size κ ;
- (ii) if $\kappa > |\mathcal{M}|$, there is an elementary extension of \mathcal{M} of size κ .

We postpone the proof of part (i).

IV. Model Theory and Non-Classical Logic

Proof. Expand the language \mathcal{L} by adding constant symbols for each $m \in \mathcal{M}$ and $c \in \kappa$. Let

$$\mathcal{T} = \text{Diag}_{\text{el}} \mathcal{M} \cup \bigcup_{c \neq c' \in \kappa} \{\neg(c = c')\}$$

\mathcal{T} has a model by compactness, and this model must be an elementary extension of \mathcal{M} with size at least κ . We then apply the downward Löwenheim–Skolem theorem if necessary to obtain a model of size exactly κ . \square

For example, if \mathcal{L} is countable, every infinite \mathcal{L} -structure has a countable elementary substructure.

2. Quantifier elimination

2.1. Skolem functions

Definition. Let \mathcal{T} be an \mathcal{L} -theory, and let $\varphi(\mathbf{x}, y)$ be an \mathcal{L} -formula where \mathbf{x} is nonempty. A *Skolem function* for φ is an \mathcal{L} -term t such that

$$\mathcal{T} \vdash \forall \mathbf{x}. (\exists y. \varphi(\mathbf{x}, y) \rightarrow \varphi(\mathbf{x}, t(\mathbf{x})))$$

A *skolemisation* of an \mathcal{L} -theory \mathcal{T} is a language $\mathcal{L}^+ \supseteq \mathcal{L}$ and an \mathcal{L}^+ -theory $\mathcal{T}^+ \supseteq \mathcal{T}$ such that

- (i) every \mathcal{L} -structure that models \mathcal{T} can be expanded to an \mathcal{L}^+ -structure that models \mathcal{T}^+ ;
- (ii) \mathcal{T}^+ has Skolem functions for any \mathcal{L}^+ -formula $\varphi(\mathbf{x}, y)$ where \mathbf{x} is nonempty.

A theory is called a *Skolem theory* if it is a skolemisation of itself.

By ‘expanded’, we mean that \mathcal{T} is given interpretations to the elements of $\mathcal{L}^+ \setminus \mathcal{L}$, but no new objects are added.

Proposition. Let \mathcal{T} be an \mathcal{L} -theory, and let \mathcal{F} be a collection of \mathcal{L} -formulae including all atomic formulae and closed under Boolean operations. Suppose that for every formula $\psi(\mathbf{x}, y) \in \mathcal{F}$, there exists $\varphi(\mathbf{x}) \in \mathcal{F}$ with

$$\mathcal{T} \vdash \forall \mathbf{x}. (\exists y. \psi(\mathbf{x}, y) \leftrightarrow \varphi(\mathbf{x}))$$

Then, every \mathcal{L} -formula is equivalent to one in \mathcal{F} with the same free variables modulo \mathcal{T} (that is, \mathcal{T} proves they are equivalent).

Proof. We proceed by induction on the length of formulae. The case of existential formulae is the only nontrivial inductive step. Consider the formula $\exists y. \psi(\mathbf{x}, y)$. By the inductive hypothesis, $\psi(\mathbf{x}, y)$ is \mathcal{T} -equivalent to $\psi'(\mathbf{x}, y) \in \mathcal{F}$. Then, there is some $\varphi(\mathbf{x}) \in \mathcal{F}$ such that

$$\mathcal{T} \vdash \forall \mathbf{x}. (\exists y. \psi'(\mathbf{x}, y) \leftrightarrow \varphi(\mathbf{x}))$$

Thus the formula $\exists y. \psi(\mathbf{x}, y)$ in question is \mathcal{T} -equivalent to $\varphi(\mathbf{x}) \in \mathcal{F}$. □

Proposition. Let \mathcal{T} be a Skolem theory. Then,

- (i) every \mathcal{L} -formula $\varphi(\mathbf{x})$ where \mathbf{x} is nonempty is equivalent modulo \mathcal{T} to some quantifier-free $\varphi^*(\mathbf{x})$;
- (ii) if $\mathcal{N} \models \mathcal{T}$ and $X \subseteq \mathcal{N}$, then either $\langle X \rangle_{\mathcal{N}} = \emptyset$ or $\langle X \rangle_{\mathcal{N}} \leq \mathcal{N}$.

Remark. When \mathcal{N} is a model of a Skolem theory, $\langle X \rangle_{\mathcal{N}}$ is sometimes called the *Skolem hull* of X .

Proof. Part (i). Clearly, $\varphi(\mathbf{x}, t(\mathbf{x})) \rightarrow \exists y. \varphi(\mathbf{x}, y)$ in any model. So having Skolem functions means that

$$\mathcal{T} \vdash \forall \mathbf{x}. (\exists y. \varphi(\mathbf{x}, y) \leftrightarrow \varphi(\mathbf{x}, t(\mathbf{x})))$$

IV. Model Theory and Non-Classical Logic

completing the proof by the previous proposition.

Part (ii). We proceed by the Tarski–Vaught test. Let $\mathcal{M} = \langle X \rangle_{\mathcal{N}}$, $\mathbf{m} \in \mathcal{M}$, and let $\varphi(\mathbf{x}, y)$ be such that

$$\mathcal{N} \models \exists y. \varphi(\mathbf{m}, y)$$

Then as \mathcal{N} has Skolem functions, there exists an \mathcal{L} -term t such that

$$\mathcal{N} \models \varphi(\mathbf{m}, t(\mathbf{m}))$$

But \mathcal{M} is closed under the interpretation of function symbols as it is a substructure, so $t(\mathbf{m}) \in \mathcal{M}$. Thus

$$\mathcal{M} \models \exists y. \varphi(\mathbf{m}, y)$$

as required. □

2.2. Skolemisation theorem

Theorem. Every first-order language \mathcal{L} can be expanded to some $\mathcal{L}^+ \supseteq \mathcal{L}$ that admits an \mathcal{L}^+ -theory Σ such that

- (i) Σ is a Skolem \mathcal{L}^+ -theory;
- (ii) any \mathcal{L} -structure can be expanded to an \mathcal{L}^+ -structure that models Σ ; and
- (iii) $|\mathcal{L}^+| = |\mathcal{L}|$.

Proof. We will design \mathcal{L}^+ to include Skolem functions for each suitable formula. If $\chi(\mathbf{x}, y)$ is an \mathcal{L} -formula with \mathbf{x} nonempty, we add a function symbol F_χ of arity $|\mathbf{x}|$. Performing this for all \mathcal{L} -formulae of this form, we obtain a language $\mathcal{L}' \supseteq \mathcal{L}$. Next, define $\Sigma(\mathcal{L})$ to be the set of \mathcal{L} -sentences that enforce the correct behaviour of the F_χ :

$$\forall \mathbf{x}. (\exists y. \chi(\mathbf{x}, y) \rightarrow \chi(\mathbf{x}, F_\chi(\mathbf{x})))$$

Note that $\Sigma(\mathcal{L})$ is an \mathcal{L}' -theory, not an \mathcal{L} -theory; there may be existentials in \mathcal{L}' without explicit witnesses. We can overcome this issue by iterating this construction ω times and taking the union. Formally, we recursively define

$$\mathcal{L}_0 = \mathcal{L}; \quad \mathcal{L}_{n+1} = \mathcal{L}'_n; \quad \Sigma_0 = \emptyset; \quad \Sigma_{n+1} = \Sigma_n \cup \Sigma(\mathcal{L}_n)$$

Then we can set

$$\mathcal{L}^+ = \bigcup_{n < \omega} \mathcal{L}_n; \quad \Sigma = \bigcup_{n < \omega} \Sigma_n$$

First, note that Σ is a Skolem theory. This is because each \mathcal{L}^+ -formula is in \mathcal{L}_n for some $n < \omega$, so $\Sigma_{n+1} \subseteq \Sigma$ asserts that it has a Skolem function. It is also clear to see that $|\mathcal{L}^+| = |\mathcal{L}|$ using basic cardinal arithmetic.

To prove property (ii), it suffices to show that each \mathcal{L} -theory can be expanded into an \mathcal{L}' -theory that models $\Sigma(\mathcal{L})$; we can then proceed by induction. Note that this argument will

2. Quantifier elimination

use the axiom of choice. Let \mathcal{M} be an \mathcal{L} -structure. We can assume $\mathcal{M} \neq \emptyset$; if $\mathcal{M} = \emptyset$ then all sentences in Σ would be vacuously true and there would be nothing to prove. We now expand \mathcal{M} into an \mathcal{L}' -structure \mathcal{M} in the following way. Consider $\chi(\mathbf{x}, y)$ where \mathbf{x} is nonempty and $\mathbf{m} \in \mathcal{M}$. If

$$\mathcal{M} \models \exists b. \chi(\mathbf{m}, b)$$

then we can choose such a b and interpret $F_\chi(\mathbf{m})$ as this value. If

$$\mathcal{M} \not\models \exists b. \chi(\mathbf{m}, b)$$

then we interpret $F_\chi(\mathbf{m})$ as an arbitrary model element, say, \mathbf{m}_0 . By construction, \mathcal{M}' models $\Sigma(\mathcal{L})$. \square

Corollary. Any \mathcal{L} -theory \mathcal{T} admits a skolemisation \mathcal{T}^+ in a language \mathcal{L}^+ of the same size as \mathcal{L} .

Proof. Take $\mathcal{T}^+ = \mathcal{T} \cup \Sigma$. Any model of \mathcal{T}^+ models Σ , so \mathcal{T}^+ has Skolem functions. Moreover, any \mathcal{L} -structure that models \mathcal{T} can be extended to one that models Σ , which will therefore model \mathcal{T}^+ . \square

Corollary (downward Löwenheim–Skolem theorem). Let \mathcal{M} be an \mathcal{L} -structure, and let $X \subseteq \mathcal{M}$. Let κ be a cardinal such that

$$|\mathcal{L}| + |X| \leq \kappa \leq |\mathcal{M}|$$

Then \mathcal{M} has an elementary substructure of size κ that contains X .

Proof. Let $X \subseteq Y \subseteq \mathcal{M}$ and $|Y| = \kappa$. Let \mathcal{M}' be an expansion of \mathcal{M} to a Skolem theory, and consider the Skolem hull $\langle Y \rangle_{\mathcal{M}'}$. $\langle Y \rangle_{\mathcal{M}'}$ must be an elementary substructure of \mathcal{M}' as $Y \neq \emptyset$. Let \mathcal{N} be the \mathcal{L} -reduct of $\langle Y \rangle_{\mathcal{M}'}$. Then \mathcal{N} is an elementary substructure of \mathcal{N} , and $X \subseteq \mathcal{N}$. It remains to check $|\mathcal{N}| = \kappa$.

$$|\mathcal{N}| \leq |Y| + |\mathcal{L}^+| = \kappa + |\mathcal{L}| = \kappa = |Y| \leq |\mathcal{N}|$$

So $|\mathcal{N}| = \kappa$. \square

2.3. Elimination sets

Definition. Let \mathcal{T} be an \mathcal{L} -theory. A set F of \mathcal{L} -formulae is an *elimination set* for \mathcal{T} if, for every \mathcal{L} -formula φ , there is a Boolean combination φ^* of formulae in F such that

$$\mathcal{T} \vdash \varphi \leftrightarrow \varphi^*$$

A theory \mathcal{T} has *quantifier elimination* if the family of quantifier-free formulae forms an elimination set for \mathcal{T} .

IV. Model Theory and Non-Classical Logic

Note that a theory having quantifier elimination depends on its underlying language. Every Skolem theory has quantifier elimination.

Example. (i) Let $p \in \mathbb{C}[x]$ be the polynomial $x^3 - 31x^2 + 6$ over \mathbb{C} . The sentence $\exists x. p(x) = 0$ contains a quantifier. But as \mathbb{C} is algebraically closed, it is equivalent to the quantifier-free sentence $1 \neq 0 \vee (-31) \neq 0$.

(ii) A real-valued matrix is invertible if there exists a two-sided inverse. This has a quantifier, but there is a quantifier-free sentence equivalent to it, namely, ‘its determinant is nonzero’.

Remark. (i) We can check if two models of \mathcal{T} are elementarily equivalent by considering just those formulae in an elimination set. In particular, to check if a theory is complete, it suffices to check that all sentences in an elimination set are either deducible from the theory or inconsistent with it.

(ii) Suppose \mathcal{L} is a recursive language, and the map $\varphi \mapsto \varphi^*$ is computable. Then an algorithm to decide whether \mathcal{T} proves any sentence can be produced from one that operates only on the elimination set.

(iii) The elementary embeddings $\mathcal{M} \rightarrow \mathcal{N}$ are precisely those embeddings that preserve φ and $\neg\varphi$ for all φ in F . So a theory with quantifier elimination is model-complete.

(iv) The definable sets of a model are precisely the Boolean combinations of sets definable with only formulae in an elimination set.

In the next result, we use the notation $\neg F$ for the set of negations of formulae in F .

Proposition (syntactic quantifier elimination). Let \mathcal{T} be an \mathcal{L} -theory, and let F be a family of \mathcal{L} -formulae including all atomic formulae. Suppose that, for every \mathcal{L} -formula of the form

$$\theta(\mathbf{x}) = \exists y. \bigwedge_{i < n} \varphi_i(\mathbf{x}, y); \quad \varphi_i \in F \cup \neg F$$

there exists a Boolean combination $\theta^*(\mathbf{x})$ of formulae in F such that

$$\mathcal{T} \vdash \forall \mathbf{x}. (\theta(\mathbf{x}) \leftrightarrow \theta^*(\mathbf{x}))$$

Then F is an elimination set for \mathcal{T} .

The proof is similar to a previous proposition.

Example. Consider the theory \mathcal{T}_∞ of infinite sets in the language with empty signature. The only atomic formulae are equalities, and the only terms in the language are variables. Using the above proposition, it suffices to eliminate the existential quantifier in formulae $\varphi(x_0, \dots, x_{n-1})$ of the form

$$\exists y. \left(\bigwedge_{i \in I} y = x_i \right) \wedge \left(\bigwedge_{i \in J} y \neq x_i \right) \wedge \left(\bigwedge_{i, j \in K} x_i = x_j \right) \wedge \left(\bigwedge_{i, j \in L} x_i \neq x_j \right)$$

2. Quantifier elimination

where $I, J, K, L \subseteq \{0, \dots, n-1\}$. Without loss of generality we can assume I is empty, as we can easily remove the quantifier in this situation. We may also push the quantifier inside the first conjunct.


$$\left(\exists y. \bigwedge_{i \in J} y \neq x_i \right) \wedge \psi(x_0, \dots, x_{n-1}); \quad \psi(x_0, \dots, x_{n-1}) = \left(\bigwedge_{i,j \in K} x_i = x_j \right) \wedge \left(\bigwedge_{i,j \in L} x_i \neq x_j \right)$$

But the theory of infinite sets proves $\exists y. \bigwedge_{i \in J} y \neq x_i$, so we can conclude that φ and ψ are equivalent modulo \mathcal{T} .

2.4. Amalgamation

Definition. Let \mathcal{M} and \mathcal{N} be \mathcal{L} -structures. We write $\mathcal{M} \rightarrow_1 \mathcal{N}$ if every existential sentence modelled by \mathcal{M} is also modelled by \mathcal{N} .

Theorem (existential amalgamation). Let \mathcal{M} and \mathcal{N} be \mathcal{L} -structures, with $S \subseteq \mathcal{M}$. Suppose there is a homomorphism $f : \langle S \rangle_{\mathcal{M}} \rightarrow \mathcal{N}$ such that $(\mathcal{N}, f(S)) \rightarrow_1 (\mathcal{M}, S)$. Then there is an elementary extension \mathcal{K} of \mathcal{N} and an embedding $g : \mathcal{M} \hookrightarrow \mathcal{K}$ making the following diagram commute.



Proof. Let \mathcal{M}, \mathcal{N} be disjoint without loss of generality. Consider the $\mathcal{L}_{\mathcal{M} \sqcup \mathcal{N}}$ -theory

$$\mathcal{T} = \text{Diag}_{\text{el}} \mathcal{M} \cup \text{Diag} \mathcal{N} \cup \bigcup_{s \in S} \{s = f(s)\}$$

We show this is consistent by compactness; then, a model \mathcal{K} will be an elementary extension of \mathcal{M} , and \mathcal{N} embeds into it in such a way that makes the above diagram commute due to the sentences $s = f(s)$. If \mathcal{T} is inconsistent, there is a finite set of formulae in $\text{Diag} \mathcal{N}$ that are inconsistent with

$$\mathcal{T}' = \text{Diag}_{\text{el}} \mathcal{M} \cup \bigcup_{s \in S} \{s = f(s)\}$$

Taking the conjunction, we can suppose it is a single formula $\varphi(\mathbf{n})$, where $\mathbf{n} \in \mathcal{N}$ is a tuple of pairwise distinct elements.

$$\mathcal{T}' \vdash \neg \varphi(\mathbf{n})$$

Then, using the sentences $s = f(s)$ and the fact that $\langle S \rangle_{\mathcal{M}}$ is generated by S , the formula $\varphi(\mathbf{n})$ is equivalent modulo \mathcal{T}' to some quantifier-free formula $\psi(\mathbf{s}, \mathbf{n}')$ where $\mathbf{s} \in S$ and $\mathbf{n}' \in \mathcal{N} \setminus \text{im } f$.

$$\mathcal{T}' \vdash \neg \psi(\mathbf{s}, \mathbf{n}')$$

IV. Model Theory and Non-Classical Logic

Now, note that \mathcal{T}' has nothing to say about \mathbf{n}' , so in fact

$$\mathcal{T}' \vdash \forall \mathbf{x}. \neg \psi(\mathbf{s}, \mathbf{x})$$

As $(\mathcal{N}, f(S)) \rightarrow_1 (\mathcal{M}, S)$, we can convert the universal quantifier above into the negation of an existential quantifier to conclude

$$\mathcal{N} \vdash \neg \exists \mathbf{x}. \psi(\mathbf{s}, \mathbf{x})$$


so

$$\mathcal{N} \vdash \neg \exists \mathbf{x}. \psi(\mathbf{s}, \mathbf{x})$$


But $\varphi(\mathbf{n})$ is in the diagram of \mathcal{N} , so $\mathcal{N} \vdash \exists \mathbf{x}. \psi(\mathbf{s}, \mathbf{x})$, giving a contradiction. \square

We can make the following more general definition.

Definition. A class \mathbb{K} of \mathcal{L} -structures has the *amalgamation property* if, given a diagram of elements of \mathbb{K}



there is a structure \mathcal{D} in \mathbb{K} and embeddings making the following diagram commute.



Definition. Let \mathbb{K} be a class of \mathcal{L} -structures and $\mathcal{M} \in \mathbb{K}$. We say that \mathcal{M} is *existentially closed* in \mathbb{K} if, for every existential formula $\psi(\mathbf{x})$ and tuple $\mathbf{m} \in \mathcal{M}$, the existence of an extension $\mathcal{M} \subseteq \mathcal{N} \in \mathbb{K}$ with $\mathcal{N} \models \psi(\mathbf{m})$ forces $\mathcal{M} \models \psi(\mathbf{m})$.

Note that being existentially closed in \mathbb{K} depends on the choice of \mathbb{K} . For example, an existentially closed ordered field need not be an existentially closed field.

Example. (i) Every field that is existentially closed in the class of fields is algebraically closed. Let A be an existentially closed field, and view a nontrivial polynomial $f(\mathbf{y})$ over A as a statement $p(\mathbf{a}, y)$ where $p(\mathbf{x}, y)$ is a term in the language of rings, and \mathbf{a} is a tuple. For instance, $y^2 + 2y - 3$ can be seen as $p(1, 2, 3, y)$, where $p(x_0, x_1, x_2, y) = x_0 y^2 + x_1 y + (-x_2)$. We can replace f with an irreducible factor and consider the quotient ring $A[y]_{(f)}$, which is an extension of A over which f has a root.

$$A[y]_{(f)} \models \exists y. p(\mathbf{a}, y) = 0$$

2. Quantifier elimination

Since f is irreducible, this is an extension of fields. Thus, as A is existentially closed,

$$A \models \exists y. p(\mathbf{a}, y) = 0$$

so f has a root in A . The converse is true, and is one way that Hilbert's Nullstellensatz can be stated.

- (ii) The existentially closed linear orders are precisely the dense linear orders without endpoints.
- (iii) The existentially closed ordered fields are precisely the *real closed fields*, which are the ordered fields elementarily equivalent to the real numbers. Equivalently, all nonnegative elements are squares, and all odd-degree elements have a root.

Theorem. Let \mathbb{K} be a class of \mathcal{L} -structures that is closed under isomorphism. Suppose that the class of all of the substructures of the structures in \mathbb{K} has the amalgamation property. Then, every existential \mathcal{L} -formula $\varphi(\mathbf{x})$ is equivalent to a quantifier-free \mathcal{L} -formula in all existentially closed structures in \mathbb{K} . In particular, if \mathcal{T} is a theory axiomatising existentially closed structures in \mathbb{K} , then \mathcal{T} has quantifier elimination.

Proof. Let $\varphi(\mathbf{x})$ be an existential formula. We will call a pair $(\mathcal{M}, \mathbf{m})$ a *witnessing pair* if \mathcal{M} is existentially closed in \mathbb{K} and $\mathcal{M} \models \varphi(\mathbf{m})$. For each such pair, let

$$\theta_{(\mathcal{M}, \mathbf{m})}(\mathbf{x}) = \bigwedge \{ \psi(\mathbf{x}) \text{ a literal} \mid \mathcal{M} \models \psi(\mathbf{m}) \}$$

where the *literals* are the atomic formulae and their negations. Let

$$\chi(\mathbf{x}) = \bigvee_{(\mathcal{M}, \mathbf{m})} \theta_{(\mathcal{M}, \mathbf{m})}(\mathbf{x})$$

It suffices to show that if \mathcal{N} is existentially closed in \mathbb{K} then

$$(\mathcal{N} \models \varphi(\mathbf{n})) \iff (\mathcal{N} \models \chi(\mathbf{n}))$$


Then we can use the compactness theorem twice to reduce χ to a first-order finitary formula as required. If $\mathbf{n} \in \mathcal{N}$ is such that $\mathcal{N} \models \varphi(\mathbf{n})$, then $(\mathcal{N}, \mathbf{n})$ is a witnessing pair, and thus $\mathcal{N} \models \chi(\mathbf{n})$ by construction. For the converse, if $\mathcal{N} \models \chi(\mathbf{n})$, there is a witnessing pair $(\mathcal{M}, \mathbf{m})$ such that $\mathcal{N} \models \theta_{(\mathcal{M}, \mathbf{m})}(\mathbf{n})$. Hence, for each literal $\psi(\mathbf{x})$,

$$(\mathcal{M} \models \psi(\mathbf{m})) \implies (\mathcal{N} \models \psi(\mathbf{n}))$$

There is thus an embedding $e : \langle \mathbf{m} \rangle_{\mathcal{M}} \hookrightarrow \mathcal{N}$ mapping \mathbf{m} to \mathbf{n} . Applying the amalgamation

IV. Model Theory and Non-Classical Logic

property, we obtain



where $\mathcal{D} \in \mathbb{K}$, and both \mathcal{M}, \mathcal{N} embed into \mathcal{C} and therefore into \mathcal{D} . Note that $g(\mathbf{m}) = h(\mathbf{n})$. Replacing \mathcal{D} with an isomorphic copy if required, we may assume that h is an inclusion, so $g(\mathbf{m}) = \mathbf{n}$. We know that $(\mathcal{M}, \mathbf{m})$ is a witnessing pair, so $\mathcal{M} \models \varphi(\mathbf{m})$. Then $\mathcal{D} \models \varphi(g(\mathbf{m}))$ as existential formulae are preserved under taking extensions. Since \mathcal{N} is existentially closed in \mathbb{K} , $\mathcal{D} \in \mathbb{K}$, and $\mathcal{N} \subseteq \mathcal{D}$, we conclude that $\mathcal{N} \models \varphi(e(\mathbf{m}))$ so $\mathcal{N} \models \varphi(\mathbf{n})$ as required.

In particular, if \mathcal{T} is a theory axiomatising existentially closed structures in \mathbb{K} , then \mathcal{T} has quantifier elimination by applying the completeness theorem and then using the syntactic criterion for quantifier elimination proven previously. \square

Example. We show that the theory ACF of algebraically closed fields has quantifier elimination. First, recall that ACF axiomatises the existentially closed fields, so it suffices to check that the class of substructures of fields has the amalgamation property. Note that a substructure of a field must satisfy all universal sentences in the theory of fields, so the substructures of fields are precisely the integral domains. General field theory shows that the class of fields has the amalgamation property; we can then prove that the class of integral domains has the amalgamation property by passing to fraction fields.

Example. The theory DLO of dense linear orders without endpoints has quantifier elimination. The class of substructures of dense linear orders has the amalgamation property: indeed, any two linear orders embed into a poset, which can be extended into a linear order by Zorn's lemma, and is thus a substructure of some dense linear order.

2.5. Inductive classes

Definition. A class \mathbb{K} of \mathcal{L} -structures is *inductive* if it is closed under isomorphisms and under unions of chains of embeddings.

Theorem. Let \mathcal{M} be a structure in an inductive class \mathbb{K} . Then $\mathcal{M} \subseteq \mathcal{N}$ for some \mathcal{N} existentially closed in \mathbb{K} .

This is analogous to the theorem that every field has an algebraic closure, and is proven in a similar way.

2. Quantifier elimination

Proof. We show that \mathcal{M} can be extended to some structure $\mathcal{M}^* \in \mathbb{K}$ with the property that for all $\mathbf{m} \in \mathcal{M}$ and $\varphi(\mathbf{x})$ an existential \mathcal{L} -formula, if $\varphi(\mathbf{m})$ holds in some extension of \mathcal{M}^* in \mathbb{K} , then $\varphi(\mathbf{m})$ holds in \mathcal{M}^* .

We now show that this suffices to complete the proof. Indeed, we then recursively define a chain of \mathbb{K} -structures by setting $\mathcal{M}^{(0)} = \mathcal{M}$ and $\mathcal{M}^{(i+1)} = (\mathcal{M}^{(i)})^*$, then taking their union to form \mathcal{N} . Then \mathcal{N} lies in \mathbb{K} as \mathbb{K} is inductive, and moreover it extends \mathcal{M} .

This \mathcal{N} is existentially closed in \mathbb{K} . Suppose $\varphi(\mathbf{x})$ is an existential formula, $\mathbf{n} \in \mathcal{N}$, and \mathcal{D} is a structure in \mathbb{K} such that $\mathcal{D} \models \varphi(\mathbf{n})$. As $\mathbf{n} \in \bigcup_{i < \omega} \mathcal{M}^{(i)}$ and the $\mathcal{M}^{(i)}$ form a chain, there must be $k < \omega$ such that $\mathbf{n} \in \mathcal{M}^{(k)}$. Then $(\mathcal{M}^{(k)})^* = \mathcal{M}^{(k+1)} \models \varphi(\mathbf{n})$, so in particular, $\mathcal{N} \models \varphi(\mathbf{n})$.

We now construct \mathcal{M}^* . Using the axiom of choice, create an ordinal-indexed list of pairs $(\varphi_\beta, \mathbf{m}_\beta)_\beta$ where φ is an existential formula and $\mathbf{m} \in \mathcal{M}$, and β ranges over all ordinals less than some ordinal δ . We then construct a chain of \mathbb{K} -structures by transfinite induction. Let $\mathcal{M}_0 = \mathcal{M}$. At each successor stage, let $\mathcal{M}_{\beta+1}$ be a \mathbb{K} -structure \mathcal{D} that extends \mathcal{M}_β and models $\varphi_\beta(\mathbf{m}_\beta)$, if this exists. If such a model does not exist, define $\mathcal{M}_{\beta+1} = \mathcal{M}_\beta$. At each limit stage, let $\mathcal{M}_\lambda = \bigcup_{\beta < \lambda} \mathcal{M}_\beta$. Finally, set $\mathcal{M}^* = \mathcal{M}_\delta$.

If $\varphi(\mathbf{x})$ is existential, $\mathbf{m} \in \mathcal{M}$, and \mathcal{D} is some \mathbb{K} -structure that extends \mathcal{M}^* and models $\varphi(\mathbf{m})$, then $(\varphi, \mathbf{m}) = (\varphi_\beta, \mathbf{m}_\beta)$ for some $\beta < \delta$. Then $\mathcal{M}_\beta \subseteq \mathcal{M}^* \subseteq \mathcal{D}$, so $\mathcal{M}_{\beta+1}$ models $\varphi_\beta(\mathbf{m}_\beta) = \varphi(\mathbf{m})$ by definition. But as φ is existential and \mathcal{M}^* extends \mathcal{M}_β , we must also have that \mathcal{M}^* models $\varphi(\mathbf{m})$, as required. \square

2.6. Characterisations of quantifier elimination

Theorem. Let \mathcal{T} be an \mathcal{L} -theory. Then the following are equivalent.


- (i) The theory \mathcal{T} is model-complete.
- (ii) Every model of \mathcal{T} is an existentially closed model of \mathcal{T} .
- (iii) Given an embedding $e : \mathcal{A} \hookrightarrow \mathcal{B}$ between models of \mathcal{T} , there is an elementary extension \mathcal{D} of \mathcal{A} and an embedding $g : \mathcal{B} \hookrightarrow \mathcal{D}$ such that $g \circ e = \text{id}_{\mathcal{A}}$.
- (iv) For any quantifier-free \mathcal{L} -formula $\varphi(\mathbf{x}, \mathbf{y})$, the formula $\exists \mathbf{y}. \varphi(\mathbf{x}, \mathbf{y})$ is equivalent to some universal \mathcal{L} -formula $\psi(\mathbf{x})$ modulo \mathcal{T} .
- (v) Every \mathcal{L} -formula is equivalent to some universal \mathcal{L} -formula modulo \mathcal{T} .

Proof. (i) implies (ii). As all embeddings between models are elementary, if a superstructure has a witness to an existential, so does the substructure.

(ii) implies (iii). We use the existential amalgamation theorem. Take S to be the set of all

IV. Model Theory and Non-Classical Logic

elements of \mathcal{A} , then by (ii), $(\mathcal{B}, e(S)) \rightarrow_1 (\mathcal{A}, S)$. We obtain



as required.

(iii) *implies* (iv). By the theorem of Tarski and Łoś characterising theories preserved under substructures, it suffices to show that existential formulas are preserved under substructures. Let $e : \mathcal{A} \rightarrow \mathcal{B}$ be such that $\mathcal{B} \models \varphi(e(\mathbf{a}))$, where φ is an existential formula, and $\mathbf{a} \in \mathcal{A}$. By (ii), there is an elementary extension \mathcal{D} of \mathcal{A} and an embedding $g : \mathcal{B} \rightarrow \mathcal{D}$ such that $g \circ e = \text{id}_{\mathcal{A}}$. Existential formulas are preserved under extensions, so $\mathcal{D} \models \varphi(\mathbf{a})$. As $\mathcal{A} \leq \mathcal{D}$, we must have $\mathcal{A} \models \varphi(\mathbf{a})$, as required.

(iv) *implies* (v). We proceed by induction on the structure of \mathcal{L} -formulae. We can iteratively convert existential quantifiers to universal quantifiers, noting that (iv) allows us to convert a sequence of existentials to a sequence of universals simultaneously.

(v) *implies* (i). Note that universal formulae are preserved under extensions, and every formula and its negation can be represented as a universal formula. This directly gives the result. \square

Let \mathcal{M}, \mathcal{N} be \mathcal{L} -structures. If \mathcal{M}, \mathcal{N} satisfy the same quantifier-free sentences, we write $\mathcal{M} \equiv_0 \mathcal{N}$.

Theorem. Let \mathcal{T} be an \mathcal{L} -theory. Then the following are equivalent.

- (i) \mathcal{T} has quantifier elimination.
- (ii) If $\mathcal{A}, \mathcal{B} \models \mathcal{T}$ and $\mathbf{a} \in \mathcal{A}, \mathbf{b} \in \mathcal{B}$ are tuples of the same length, then $(\mathcal{A}, \mathbf{a}) \equiv_0 (\mathcal{B}, \mathbf{b})$ implies $(\mathcal{A}, \mathbf{a}) \rightarrow_1 (\mathcal{B}, \mathbf{b})$.
- (iii) Whenever $\mathcal{A}, \mathcal{B} \models \mathcal{T}$, $S \subseteq \mathcal{A}$ and $e : \langle S \rangle_{\mathcal{A}} \rightarrow \mathcal{B}$, then there is an elementary extension \mathcal{D} of \mathcal{B} and an embedding $f : \mathcal{A} \rightarrow \mathcal{D}$ extending e .
- (iv) \mathcal{T} is model-complete and \mathcal{T}_{\forall} has the amalgamation property.
- (v) For every quantifier-free \mathcal{L} -formula $\varphi(\mathbf{x}, y)$, the formula $\exists y. \varphi(\mathbf{x}, y)$ is \mathcal{T} -equivalent to a quantifier-free formula $\psi(\mathbf{x})$.


Proof. (i) *implies* (ii) is clear.

(ii) *implies* (iii). It suffices to show that $(\mathcal{B}, e(S)) \rightarrow_1 (\mathcal{A}, S)$ by the existential amalgamation theorem. Since a sentence in \mathcal{L}_S is finite, it can only mention finitely many of the new constants in S , so it is enough to check that $(\mathcal{B}, e(\mathbf{a})) \rightarrow_1 (\mathcal{A}, \mathbf{a})$ for all tuples \mathbf{a} obtainable


2. Quantifier elimination

from S . Now, if \mathbf{a} is such a tuple and $e : \langle S \rangle_{\mathcal{A}} \rightarrow \mathcal{B}$ is an embedding, then $(\mathcal{A}, \mathbf{a}) \equiv_0 (\mathcal{B}, e(\mathbf{a}))$, giving the required result by (ii).

(iii) *implies* (iv). By the previous theorem, to check model-completeness it suffices to check that for each embedding $h : \mathcal{M} \rightarrow \mathcal{N}$ between models of \mathcal{T} , there is an elementary extension \mathcal{D} of \mathcal{M} and an embedding $g : \mathcal{N} \rightarrow \mathcal{D}$ such that $g \circ h = \text{id}_{\mathcal{M}}$. Consider the instance of (iii) where $S = h(\mathcal{M})$ and $e = h^{-1}$ as a map $h(\mathcal{M}) \simeq \mathcal{M}$. Then there is an elementary extension \mathcal{D} of \mathcal{M} and an embedding $g : \mathcal{M} \rightarrow \mathcal{D}$ extending e .



This means that for all $m \in \mathcal{M}$, we have $g(h(m)) = e(h(m)) = m$. To see that \mathcal{T}_{\forall} has the amalgamation property, consider models $\mathcal{A}', \mathcal{B}', \mathcal{C}$ of \mathcal{T}_{\forall} where \mathcal{C} embeds into both \mathcal{A}' and \mathcal{B}' . Models of \mathcal{T}_{\forall} are precisely the substructures of models of \mathcal{T} , so \mathcal{A}' and \mathcal{B}' are substructures of models \mathcal{A} and \mathcal{B} of \mathcal{T} respectively. Consider the instance of (iii) where $S = \mathcal{C} = \langle \mathcal{C} \rangle_{\mathcal{A}}$ and e is the embedding of \mathcal{C} into \mathcal{B} . Then we have an elementary extension \mathcal{D} of \mathcal{B} and an embedding $f : \mathcal{A} \rightarrow \mathcal{D}$ that extends e .



Now, $\mathcal{D} \equiv \mathcal{B} \models \mathcal{T} \vdash \mathcal{T}_{\forall}$, we must have that \mathcal{D} is a model of \mathcal{T}_{\forall} giving the amalgamation property as desired.

(iv) *implies* (v). Model-completeness implies that every model of \mathcal{T} is an existentially closed model of \mathcal{T} . Then, by the theorem characterising theories axiomatising existentially closed structures, this proof is complete, as the models of \mathcal{T}_{\forall} are precisely the substructures of models of \mathcal{T} .

(v) *implies* (i). Immediate from the syntactic criterion for quantifier elimination. \square

Corollary. Let \mathcal{A} be a finite \mathcal{L} -structure. The theory $\text{Th}(\mathcal{A})$ of \mathcal{A} has quantifier elimination if and only if every isomorphism between finitely generated substructures of \mathcal{A} can be extended to an automorphism of \mathcal{A} .

IV. Model Theory and Non-Classical Logic

Proof. For the forward direction, consider case (iii) of the previous theorem applied to $\mathcal{A} = \mathcal{B}$ where e is the composite $\langle \mathbf{a} \rangle_{\mathcal{A}} \simeq \langle \mathbf{b} \rangle_{\mathcal{A}} \rightarrow \mathcal{A}$. We obtain an elementary extension \mathcal{D} of \mathcal{A} . If $|\mathcal{A}| = n < \aleph_0$, then the theory of \mathcal{A} must include a sentence that states this fact. Thus \mathcal{D} models the same sentence, so $|\mathcal{D}| = n = |\mathcal{A}|$. Thus \mathcal{A} and \mathcal{D} are elementarily equivalent finite structures, so the elementary embedding $h : \mathcal{A} \rightarrow \mathcal{D}$ is an isomorphism.

$$\begin{array}{ccc} \mathcal{A} & \xrightarrow{f} & \mathcal{D} \xrightarrow{h^{-1}} \mathcal{A} \\ \uparrow & & \uparrow \\ \langle \mathbf{a} \rangle_{\mathcal{A}} & \xrightarrow{\sim} & \langle \mathbf{b} \rangle_{\mathcal{A}} \end{array}$$

Now, as $|\mathcal{A}| = |\mathcal{D}| = n < \aleph_0$ and f is an embedding, it must also be surjective by the pigeonhole principle, and thus an isomorphism. Hence $h^{-1} \circ f$ is an automorphism of \mathcal{A} extending our isomorphism $\langle \mathbf{a} \rangle_{\mathcal{A}} \simeq \langle \mathbf{b} \rangle_{\mathcal{A}}$, as required.

For the converse, we prove case (ii) in the previous theorem. Let $\mathbf{b} \in \mathcal{B} \models \text{Th}(\mathcal{A})$ and $\mathbf{c} \in \mathcal{C} \models \text{Th}(\mathcal{A})$ be tuples of the same length. As $\text{Th}(\mathcal{A})$ is a complete theory, the models \mathcal{B} and \mathcal{C} are elementarily equivalent to \mathcal{A} , and thus by finiteness they are isomorphic. Thus, without loss of generality, we can set $\mathcal{A} = \mathcal{B} = \mathcal{C}$. By hypothesis, $(\mathcal{A}, \mathbf{b}) \equiv_0 (\mathcal{A}, \mathbf{c})$. Thus we obtain an isomorphism $\langle \mathbf{b} \rangle_{\mathcal{A}} \simeq \langle \mathbf{c} \rangle_{\mathcal{A}}$ mapping \mathbf{b} to \mathbf{c} , which can be extended to an automorphism of \mathcal{A} by assumption. If \mathbf{m} is a witness to

$$(\mathcal{A}, \mathbf{b}) \models \exists y. \varphi(\mathbf{b}, y)$$

then $f(\mathbf{m})$ must witness the truth of

$$(\mathcal{A}, \mathbf{c}) \models \exists y. \varphi(\mathbf{c}, y)$$

Thus, $(\mathcal{A}, \mathbf{b}) \rightarrow_1 (\mathcal{A}, \mathbf{c})$ as required. \square

Example. Let V be a finite vector space. Any isomorphism between subspaces can be extended to an automorphism using the Steinitz exchange lemma, so $\text{Th}(V)$ has quantifier elimination.

Corollary. Let \mathcal{T} be an \mathcal{L} -theory such that

- (i) \mathcal{T} preserves existential formulas under substructures: if $\mathcal{A}, \mathcal{B} \models \mathcal{T}$ with $\mathcal{A} \subseteq \mathcal{B}$, and $\varphi(\mathbf{x}, y)$ is a quantifier-free formula, then for all $\mathbf{a} \in \mathcal{A}$,

$$(\mathcal{B} \models \exists y. \varphi(\mathbf{a}, y)) \implies (\mathcal{A} \models \exists y. \varphi(\mathbf{a}, y))$$

- (ii) For any $\mathcal{C} \subseteq \mathcal{A} \models \mathcal{T}$, there is an *initial intermediate model* $\mathcal{A}' \models \mathcal{T}$: that is, $\mathcal{C} \subseteq \mathcal{A}' \subseteq \mathcal{A}$, and for any other model $\mathcal{C} \subseteq \mathcal{B} \subseteq \mathcal{A}$, there is an embedding $\mathcal{A}' \rightarrow \mathcal{B}$ that fixes \mathcal{C} .

Then \mathcal{T} has quantifier elimination.

2. Quantifier elimination


Proof. We show that condition (ii) of the theorem above holds. Let \mathcal{A}, \mathcal{B} be models of \mathcal{T} , and $\mathbf{a} \in \mathcal{A}, \mathbf{b} \in \mathcal{B}$ be such that $(\mathcal{A}, \mathbf{a}) \equiv_0 (\mathcal{B}, \mathbf{b})$. It suffices to show that $(\mathcal{A}, \mathbf{a}) \rightarrow_1 (\mathcal{B}, \mathbf{b})$. Let $\varphi(\mathbf{x}, y)$ be quantifier-free, and such that $\mathcal{A} \models \exists y. \varphi(\mathbf{a}, y)$. Let $\mathbf{c} = (c_0, \dots, c_{k-1}) \in \mathcal{A}$ be such a witness, so $\mathcal{A} \models \varphi(\mathbf{a}, \mathbf{c})$.

We claim that there is an elementary extension \mathcal{B}_0 of \mathcal{B} and an element $d_0 \in \mathcal{B}_0$ such that $(\mathcal{A}, \mathbf{a}, c_0) \equiv_0 (\mathcal{B}_0, \mathbf{b}, d_0)$. If we can do this, we can iterate the process to obtain a chain of elementary extensions

$$\mathcal{B} \leq \mathcal{B}_0 \leq \mathcal{B}_1 \leq \dots \leq \mathcal{B}_{k-1}$$

and elements $d_i \in \mathcal{B}_i$ such that $(\mathcal{A}, \mathbf{a}, \mathbf{c}) \equiv_0 (\mathcal{B}, \mathbf{b}, \mathbf{d})$. Then $\mathcal{B}_{k-1} \models \varphi(\mathbf{b}, \mathbf{d})$ as φ is quantifier-free, so $\mathcal{B}_{k-1} \models \exists y. \varphi(\mathbf{b}, y)$, giving $\mathcal{B} \models \exists y. \varphi(\mathbf{b}, y)$ as $\mathcal{B}_{k-1} \equiv \mathcal{B}$ as required.

To find \mathcal{B}_0 and d_0 , we use the hypotheses and the compactness theorem. As $(\mathcal{A}, \mathbf{a}) \equiv_0 (\mathcal{B}, \mathbf{b})$, there is an isomorphism $\langle \mathbf{a} \rangle_{\mathcal{A}} \rightarrow \langle \mathbf{b} \rangle_{\mathcal{B}}$ mapping \mathbf{a} to \mathbf{b} . Take $\mathcal{C} = \langle \mathbf{a} \rangle_{\mathcal{A}} \subseteq \mathcal{A}$. By hypothesis (ii), there is an initial intermediate model $\mathcal{C} \subseteq \mathcal{A}' \subseteq \mathcal{A}$ with $\mathcal{A}' \models \mathcal{T}$, and there is an embedding $\mathcal{A}' \hookrightarrow \mathcal{B}$ fixing \mathcal{C} . Without loss of generality, let us assume that this embedding is an inclusion.



Write

$$\Psi = \{\psi(\mathbf{x}, y) \mid \mathcal{A} \models \psi(\mathbf{a}, c_0), \psi \text{ quantifier-free}\}$$

As $\mathbf{a} \in \mathcal{A}'$, we have that $\mathcal{A}' \models \exists y. \psi(\mathbf{a}, y)$ for all $\psi \in \Psi$ by hypothesis (i). Now, $\mathcal{A}' \subseteq \mathcal{B}$, and existential formulae are preserved under extension, so $\mathcal{B} \models \exists y. \psi(\mathbf{b}, y)$ for all $\psi \in \Psi$. We conclude that every finite subset of Ψ is satisfied by some element of \mathcal{B} , as finite conjunctions of quantifier-free formulae are also quantifier-free. Thus, by compactness, there is an elementary extension $\mathcal{B} \leq \mathcal{B}_0$ and $d_0 \in \mathcal{B}_0$ satisfying the formulae in Ψ . In particular, $(\mathcal{A}, \mathbf{a}, c_0) \equiv_0 (\mathcal{B}_0, \mathbf{b}, d_0)$. \square

2.7. Applications

Example. The theory RCF of *real closed fields* is the theory of ordered fields for which every nonnegative element is a square, and that all odd polynomials have a root. Equivalently, it is the theory of ordered fields elementarily equivalent to \mathbb{R} . We show that this theory, with signature $(+, \times, 0, 1, <)$, has quantifier elimination. We will assume that every ordered field has a *real closure*, and that a real closed field satisfies the intermediate value theorem for polynomials.

We show that hypothesis (i) of the corollary above holds. Suppose we have an embedding $\mathcal{A} \subseteq \mathcal{B}$ of real closed fields, $\mathbf{a} \in \mathcal{A}$, and a quantifier-free formula $\varphi(\mathbf{x}, y)$ such that $\mathcal{B} \models$

IV. Model Theory and Non-Classical Logic

$\exists y. \varphi(\mathbf{a}, y)$. By considering the disjunctive normal form, we may assume that φ is a disjunction of a conjunction of literals. Moreover, the formulae $y \neq z$ and $y \not< z$ can be written in terms of $=$ and $<$. Thus, we may assume that $\varphi(\mathbf{a}, y)$ is of the form

$$\left(\bigwedge_{i < r} p_i(y) = 0 \right) \vee \left(\bigwedge_{j < s} 0 < q_j(y) \right)$$

where p_i, q_j are polynomials with coefficients in \mathcal{A} . If φ contains a nontrivial equation $p_i(y) = 0$, then if a witness exists in \mathcal{B} , it must be algebraic over \mathcal{A} . One can show algebraically that this witness must lie in \mathcal{A} . Therefore, let us suppose $r = 0$.

There are only finitely many points $c_0, \dots, c_{n-1} \in \mathcal{A}$ that are roots for the $q_j(y)$. Since the real closed fields satisfy the intermediate value theorem for polynomials, the $q_j(y)$ can only change sign at the c_i . Note that

$$\mathcal{A} \models \forall xy. x < y \rightarrow \exists z. (x < z \wedge z < y)$$

Since the c_i lie in \mathcal{A} , there is an element of \mathcal{A} between any pair of distinct c_i . Suppose b witnesses $\exists y. \varphi(\mathbf{a}, y)$ in \mathcal{B} . If there is a smallest interval (c_i, c_j) containing \mathcal{B} , we can pick $a \in \mathcal{A}$ also inside this interval, giving $\mathcal{A} \models \varphi(\mathbf{a}, a)$ as required. The other cases are similar.

We now show hypothesis (ii). Suppose $\mathcal{C} \subseteq \mathcal{A}$ where \mathcal{A} is a real closed field. Then \mathcal{C} is an ordered integral domain. The field of fractions of \mathcal{C} can be made an ordered field in a canonical way, by saying $\frac{a}{b} > 0$ if $ab > 0$. The embedding \mathcal{C} into \mathcal{A} is an injective homomorphism of ordered rings, into an ordered field. By the universal property of the fraction field, there is a unique homomorphism of ordered fields from $FF(\mathcal{C})$ to \mathcal{A} that extends the inclusion of \mathcal{C} into \mathcal{A} . Let \mathcal{A}' be the real closure of $FF(\mathcal{C})$, so that $\mathcal{C} \subseteq FF(\mathcal{C}) \subseteq \mathcal{A}' \subseteq \mathcal{A}$. If $\mathcal{B} \models \text{RCF}$ and $\mathcal{C} \subseteq \mathcal{B}$, then by the same argument we have a unique ordered ring homomorphism $FF(\mathcal{C}) \rightarrow \mathcal{B}$ extending the embedding $\mathcal{C} \subseteq \mathcal{B}$. Thus $\mathcal{A}' \subseteq \mathcal{B}$ as well, and this embedding fixes \mathcal{C} .

Corollary (Hilbert's Nullstellensatz). Let k be an algebraically closed field, and I be a proper ideal of $k[x_1, \dots, x_n]$. Then there exists $\mathbf{a} \in k^n$ such that $f(\mathbf{a}) = 0$ for all $f \in I$.

Proof. By Zorn's lemma, every proper ideal can be extended to a maximal ideal, so without loss of generality we may assume that I is a maximal ideal. Let L be the residue field $k[x_1, \dots, x_n]_I$, and let \bar{L} be its algebraic closure. By Hilbert's basis theorem, there exists a finite set of generators f_1, \dots, f_r for I . Note that $\mathbf{0}$ is a witness to

$$\bar{L} \models \exists \mathbf{x}. (f_1(\mathbf{x}) = 0 \wedge \dots \wedge f_r(\mathbf{x}) = 0)$$

We have embeddings $k \subseteq L \subseteq \bar{L}$, where both k and \bar{L} are algebraically closed fields. The theory of algebraically closed fields has quantifier elimination, so is model-complete. Thus the embedding $k \subseteq \bar{L}$ is elementary, so

$$k \models \exists \mathbf{x}. (f_1(\mathbf{x}) = 0 \wedge \dots \wedge f_r(\mathbf{x}) = 0)$$

We can then take \mathbf{a} to be a witness to this existential. □

2. Quantifier elimination

Corollary (Chevalley's theorem). Let k be an algebraically closed field. Then the image of a constructible set in k^n under a polynomial map is constructible.

Proof. The quantifier-free-definable subsets of k^n are precisely the finite Boolean combinations of the Zariski closed subsets of k^n , which are by definition the constructible sets. As ACF has quantifier elimination, these are exactly the definable subsets using arbitrary formulae. Now, if $X \subseteq k^n$ is constructible and $p : k^n \rightarrow k^m$ is a polynomial map, then

$$p(X) = \{y \in k^m \mid \exists x. p(x) = y\}$$

This is definable in the same language, so is a constructible set. □

3. Ultraproducts

3.1. Products

We will use the symbol λ to define functions without giving them explicit names. The syntax $\lambda x. y$ represents the function f such that $f(x) = y$.

Let $\{\mathcal{M}_i\}_{i \in I}$ be a set of \mathcal{L} -structures. The *product* $\prod_{i \in I} \mathcal{M}_i$ of this family is the \mathcal{L} -structure with carrier set

$$\prod_{i \in I} \mathcal{M}_i = \left\{ \alpha : I \rightarrow \bigcup M_i \mid \alpha(i) \in \mathcal{M}_i \right\}$$

such that

- an n -ary function symbol f is interpreted as

$$f_{\prod_I \mathcal{M}_i} : \left(\prod_I \mathcal{M}_i \right)^n \rightarrow \prod_I \mathcal{M}_i$$

given by

$$(\alpha_1, \dots, \alpha_n) \mapsto \lambda i. f^{\mathcal{M}_i}(\alpha_1(i), \dots, \alpha_n(i))$$

- an n -ary relation symbol R is interpreted as the subset

$$R_{\prod_I \mathcal{M}_i} \subseteq \left(\prod_I \mathcal{M}_i \right)^n$$

given by

$$R_{\prod_I \mathcal{M}_i} = \left\{ (\alpha_1, \dots, \alpha_n) \in \left(\prod_I \mathcal{M}_i \right)^n \mid \forall i \in I. (\alpha_1(i), \dots, \alpha_n(i)) \in R^{\mathcal{M}_i} \right\}$$

The relation symbols in this kind of product are not particularly useful. We want to construct a different kind of product in such a way that φ holds in the product if the set of \mathcal{M}_i that model φ is ‘large’.

3.2. Lattices

Definition. A *lattice* is a set L equipped with binary operations \wedge and \vee that are associative and commutative, and satisfy the *absorption laws*

$$a \vee (a \wedge b) = a; \quad a \wedge (a \vee b) = a$$

A lattice is called

- *distributive*, if $a \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c)$;

- *bounded*, if there are elements \perp and \top such that $a \vee \perp = a$ and $a \wedge \top = a$;
- *complemented*, if it is bounded and for each $a \in L$ there exists $a^* \in L$ called its *complement* such that $a \wedge a^* = \perp$ and $a \vee a^* = \top$;
- a *Boolean algebra*, if it is distributive, bounded, and complemented.

Remark. (i) Distributive lattices model the fragment of a deduction system with only the conjunction and disjunction operators. Boolean algebras model classical propositional logic.

- (ii) Every lattice has an ordering, defined by $a \leq b$ when $a \wedge b = a$. This ordering models the provability relation between propositions.

Example. (i) Let I be a set. The power set $\mathcal{P}(I)$ can be made into a Boolean algebra by taking $\wedge = \cap$ and $\vee = \cup$.

- (ii) More generally, let X be a topological space. The set of closed and open sets of X form a Boolean algebra; they can also be thought of as the propositions in classical logic. In fact, all Boolean algebras are of this form. This result is known as Stone's representation theorem.

- (iii) For any \mathcal{L} -structure \mathcal{M} and subset $B \subseteq \mathcal{M}$, the set $\{\varphi(\mathcal{M}) \mid \varphi(\mathbf{x}) \in \mathcal{L}_B\}$ of definable subsets with parameters in B is a Boolean algebra.

3.3. Filters

Definition. Let X be a lattice. A *filter* \mathcal{F} on X is a subset of X such that

- (i) $\mathcal{F} \neq \emptyset$;
- (ii) \mathcal{F} is *upward closed*: if $f \leq x$ and $f \in \mathcal{F}$ then $x \in \mathcal{F}$;
- (iii) \mathcal{F} is *downward directed*: if $x, y \in \mathcal{F}$, then $x \wedge y \in \mathcal{F}$.

A filter on X may be thought of as a collection of 'large' subsets of X : subsets that are so large that the intersection of any two large subsets is also large. For property (ii), we might also say that \mathcal{F} is a *terminal segment* of X .

Example. (i) Given an element $j \in I$, the family \mathcal{F}_j of all subsets of I containing j is a filter on $\mathcal{P}(I)$. A filter of this form is called *principal*. A filter that is not principal is called *free*.

- (ii) The family of all cofinite subsets of I forms a filter on $\mathcal{P}(I)$, called the *Fréchet filter*. One can show that any free maximal filter on an infinite set must contain the Fréchet filter.

- (iii) The family of measurable subsets of $[0, 1]$ with Lebesgue measure 1 is a filter.

Definition. A filter \mathcal{F} on a lattice L is *proper* if it is not equal to L . A maximal proper filter is called an *ultrafilter*.

IV. Model Theory and Non-Classical Logic

The ultrafilters on $\mathcal{P}(I)$ are precisely those filters \mathcal{F} where for each $U \subseteq I$, either $U \in \mathcal{F}$ or $I \setminus U \in \mathcal{F}$.

Proposition (the ultrafilter principle). Given a set I , every proper filter on $\mathcal{P}(I)$ can be extended to an ultrafilter.

The ultrafilter principle is a choice principle that is strictly weaker than the axiom of choice.

Proof. Apply Zorn's lemma. □

3.4. Łoś' theorem

For $\alpha \in \prod_{i \in I} \mathcal{M}_i$ and $\varphi(\mathbf{x})$ an \mathcal{L} -formula, we write

$$[\varphi(\alpha)] = \{i \in I \mid \mathcal{M}_i \models \varphi(\alpha(i))\}$$

Let I be a set and \mathcal{F} be a filter on $\mathcal{P}(I)$. Let $\{\mathcal{M}_i\}_{i \in I}$ be a family of \mathcal{L} -structures. The carrier set for the reduced product $\prod \mathcal{M}_i / \mathcal{F}$ is the quotient of the cartesian product $\prod_{i \in I} \mathcal{M}_i$ by the equivalence relation defined by $\alpha \sim \beta$ if and only if $[\alpha = \beta] \in \mathcal{F}$. We write $\langle \alpha \rangle$ for the equivalence class of α in the reduced product. If \mathcal{F} is an ultrafilter, we call the reduced product an *ultraproduct*. If all of the factors \mathcal{M}_i are equal, the ultrapower is called an *ultrapower*.

We turn the reduced product into an \mathcal{L} -structure as follows.

$$\begin{aligned} f^{\prod \mathcal{M}_i / \mathcal{F}}(\langle \alpha_1 \rangle, \dots, \langle \alpha_n \rangle) &= \langle \lambda i. f^{\mathcal{M}_i}(\alpha_1(i), \dots, \alpha_n(i)) \rangle \\ (\langle \alpha_1 \rangle, \dots, \langle \alpha_n \rangle) \in R^{\prod \mathcal{M}_i / \mathcal{F}} &\iff [R(\alpha_1, \dots, \alpha_n)] \in \mathcal{F} \end{aligned}$$

Note that if $\mathcal{F} = \mathcal{F}_j$ is a principal filter, then $\prod \mathcal{M}_i / \mathcal{F} \cong \mathcal{M}_j$.

Theorem. Let $\{\mathcal{M}_i\}_{i \in I}$ be a set of \mathcal{L} -structures, and \mathcal{U} be an ultrafilter on $\mathcal{P}(I)$. Then for all $(\langle \alpha_1 \rangle, \dots, \langle \alpha_n \rangle) \in \left(\prod \mathcal{M}_i / \mathcal{U}\right)^n$ and \mathcal{L} -formulae $\varphi(x_1, \dots, x_n)$,

$$\prod \mathcal{M}_i / \mathcal{U} \models \varphi(\langle \alpha_1 \rangle, \dots, \langle \alpha_n \rangle) \iff [\varphi(\alpha_1, \dots, \alpha_n)] \in \mathcal{U}$$

In particular, if each \mathcal{M}_i is a model for some theory \mathcal{T} , then so is the ultrapower.

Proof. We prove the result by induction on the length of φ . The result holds for atomic formulae by the definition of the interpretations of function and relation symbols. Since all first-order formulae are equivalent to one composed of atomic formulae under negations, conjunctions, and existential quantification, it suffices to check these cases.

If the theorem holds for ψ , and $\varphi = \neg\psi$, we can negate both sides of the induction hypothesis to show that

$$\prod \mathcal{M}_i / \mathcal{U} \models \neg\psi \iff [\psi] \notin \mathcal{U}$$

As \mathcal{U} is an ultrafilter, the right hand side holds if and only if the complement of $[\psi]$ lies in \mathcal{U} . But this complement is precisely $[\neg\psi]$, as required.

If the theorem holds for ψ_1, ψ_2 , then

$$\prod \mathcal{M}_i / \mathcal{U} \models \psi_i \iff [\psi_i] \in \mathcal{U}$$

$$\begin{aligned} \prod \mathcal{M}_i / \mathcal{U} \models \psi_1 \wedge \psi_2 &\iff [\psi_1] \in \mathcal{U} \text{ and } [\psi_2] \in \mathcal{U} \\ &\iff [\psi_1 \wedge \psi_2] \in \mathcal{U} \end{aligned}$$

Indeed, if $[\psi_1 \wedge \psi_2] \in \mathcal{U}$, then both $[\psi_1]$ and $[\psi_2]$ are in \mathcal{U} , since $[\psi_1 \wedge \psi_2] \subseteq [\psi_1], [\psi_2]$. Conversely, if $[\psi_1], [\psi_2] \in \mathcal{U}$, then $[\psi_1] \cap [\psi_2] \subseteq [\psi_1 \wedge \psi_2]$ as they are equal, but $[\psi_1] \cap [\psi_2] \in \mathcal{U}$, so $[\psi_1 \wedge \psi_2] \in \mathcal{U}$.

For the case of existential quantification, we will use the axiom of choice. Let x be free in ψ . We have

$$\prod \mathcal{M}_i / \mathcal{U} \models \exists x. \psi(x) \iff \exists \langle \alpha \rangle. \prod \mathcal{M}_i / \mathcal{U} \models \psi(\langle \alpha \rangle)$$

By the inductive hypothesis, the right hand side holds if and only if $[\psi(\alpha)] \in \mathcal{U}$. Suppose that

$$\prod \mathcal{M}_i / \mathcal{U} \models \psi(\langle \alpha \rangle)$$

Then $[\psi(\alpha)] \subseteq [\exists x. \psi(x)] \in \mathcal{U}$, as \mathcal{U} is a filter.

Conversely, suppose $[\exists x. \psi(x)] \in \mathcal{U}$. Using the axiom of choice, we can choose a witness $\alpha(i)$ to $\mathcal{M}_i \models \exists x. \psi(x)$ for each $i \in [\exists x. \psi(x)]$. For each $i \notin [\exists x. \psi(x)]$, we choose an arbitrary element of \mathcal{M}_i . Hence,

$$\prod \mathcal{M}_i / \mathcal{U} \models \psi(\langle \alpha \rangle)$$

□

Remark. (i) Since \mathcal{U} is an ultrafilter, the complement of $[\exists x. \psi(x)]$ is not in \mathcal{U} . Thus, the set of indices I for which $\alpha(i)$ was chosen arbitrarily does not lie in the ultrafilter, so this choice does not change the equivalence class of α .

(ii) The use of the axiom of choice in the above theorem is essential.

Example. We will show that the class of torsion groups is not first-order axiomatisable in the usual language of abelian groups with signature $(+, 0)$. Let \mathcal{U} be a free ultrafilter on ω , and consider the ultraproduct

$$G = \prod_{i < \omega} C_{i+1} / \mathcal{U}$$

where C_i is the cyclic group of order i , generated by g_i . Consider the element

$$g = \langle \lambda i. g_i \rangle \in G$$

IV. Model Theory and Non-Classical Logic

This has finite order if and only if $[ng = 0] \in \mathcal{U}$ for some $n > 0$. However, for each such n , the set $[ng = 0]$ is finite, so $[ng \neq 0] \in \mathcal{U}$ as \mathcal{U} contains the Fréchet filter, thus $[ng = 0] \notin \mathcal{U}$. But if the class of torsion groups were axiomatisable, this ultraproduct would also model that theory, and thus would be torsion.

Example. Let \mathcal{U} be a free ultrafilter on ω , and consider the ultrapower

$$\mathbb{N}^{\mathcal{U}} = \prod_{i < \omega} \mathbb{N} / \mathcal{U}$$

Its elements are equivalence classes of sequences of natural numbers, where $\langle (a_n) \rangle = \langle (b_n) \rangle$ if and only if $\{n \mid a_n = b_n\} \in \mathcal{U}$. It has elements such as $\langle (n)_{n < \omega} \rangle$, which represent infinitely large numbers. If \mathbb{N} has its usual structure for the language of arithmetic $\mathcal{L}_{\text{arith}}$, then the ultrapower $\mathbb{N}^{\mathcal{U}}$ is a *nonstandard model* of Peano arithmetic by Łoś' theorem, and is an elementary extension of \mathbb{N} .

Example. Let \mathcal{U} be a free ultrafilter on ω , and consider the ultrapower $\mathbb{R}^{\mathcal{U}}$, which is an elementary extension of \mathbb{R} . This includes 'large numbers' bigger than any standard real number, such as $\omega = \langle (n)_{n < \omega} \rangle$, and also includes 'infinitesimal numbers' such as $\frac{1}{\omega}$. This is not zero, but is smaller than any positive standard real.

We can give a semantic proof of the compactness theorem without using completeness, by using Łoś' theorem.

Corollary. Let \mathcal{T} be a first-order theory such that every finite subset of \mathcal{T} has a model. Then \mathcal{T} has a model.

Proof. If \mathcal{T} is finite, the result is trivial, so we may suppose it is infinite. Let I be the set of all finite subtheories of \mathcal{T} , and let

$$D = \{Y \subseteq I \mid \exists \Delta \in I. \forall X \in Y. \Delta \subseteq X\}$$

Then D is a proper filter on I , so by the ultrafilter principle, it can be extended to an ultrafilter \mathcal{U} . Using the axiom of choice, let \mathcal{M}_{Δ} be a model of Δ for each finite subtheory $\Delta \in I$. Then, for any $\varphi \in \mathcal{T}$, we have

$$\{Y \subseteq I \mid \forall X \in Y. \varphi \in X\} \in D \subseteq \mathcal{U}$$

Then by Łoś' theorem, the ultraproduct $\prod_{\Delta \in I} \mathcal{M}_{\Delta} / \mathcal{U}$ models φ . In particular, the ultraproduct models \mathcal{T} . \square

4. Types

4.1. Definitions

Definition. Let $X \subseteq \mathcal{M}^n$ be a subset of an \mathcal{L} -structure \mathcal{M} , and let $P \subseteq \mathcal{M}$. We say that X is *definable* in \mathcal{L} with *parameters* in P if there is a tuple $\mathbf{p} \in P$ and an \mathcal{L}_P -formula $\varphi(\mathbf{x}, \mathbf{y})$ such that

$$X = \varphi(\mathbf{x}, \mathbf{p}) = \{\mathbf{m} \in \mathcal{M}^n \mid \mathcal{M} \models \varphi(\mathbf{m}, \mathbf{p})\}$$

If $P = \mathcal{M}$, we say that X is *definable*.

Example. Consider the usual natural numbers as a structure for the language generated by the signature $(+, \cdot, 0, 1)$. Then there is an \mathcal{L} -formula $T(e, x, s)$ such that $\mathbb{N} \models T(e, x, s)$ if and only if the Turing machine encoded by the number e halts on input x in at most s steps. Thus, the set of halting computations is definable in this language. In particular, this implies that the theory of \mathbb{N} is not decidable.

Definition. Let \mathcal{T} be a theory and $n \in \mathbb{N}$. We obtain an equivalence relation \sim on the set $\mathcal{L}(\mathbf{x})$ of \mathcal{L} -formulae with free variables \mathbf{x} , where \mathbf{x} is a tuple of length n , by setting

$$\varphi(\mathbf{x}) \sim \psi(\mathbf{x}) \iff \mathcal{T} \vdash \forall \mathbf{x}. (\varphi(\mathbf{x}) \leftrightarrow \psi(\mathbf{x}))$$

The quotient $\mathcal{B}_n(\mathcal{T}) = \mathcal{L}(\mathbf{x}) / \sim$ becomes a Boolean algebra by setting $[\varphi] \bowtie [\psi] = [\varphi \bowtie \psi]$ for any logical connective \bowtie , called the *Lindenbaum–Tarski algebra* of \mathcal{T} on variables \mathbf{x} .

Definition. Let \mathcal{M} be an \mathcal{L} -structure and $A \subseteq \mathcal{M}$. Let \mathcal{T} be the \mathcal{L}_A -theory of sentences with parameters in A that hold in \mathcal{M} , denoted $\text{Th}_A(\mathcal{M})$. The proper filters on the Boolean algebra $\mathcal{B}_n(\mathcal{T})$ are called the *n-types* of \mathcal{M} over A .

Remark. If \mathcal{F} is a proper filter on $\mathcal{B}_n(\mathcal{T})$, it cannot include the bottom element $[\perp]$. This motivates the following more convenient definition of an *n-type*.

Definition. Let \mathcal{M} be an \mathcal{L} -structure and $A \subseteq \mathcal{M}$. A set p of \mathcal{L}_A -formulae with n free variables \mathbf{x} is an *n-type* of \mathcal{M} over A if $p \cup \text{Th}_A(\mathcal{M})$ is satisfiable. More generally, if \mathcal{T} is a theory, we say that a set p of \mathcal{L} -formulae with n free variables \mathbf{x} is an *n-type* of \mathcal{T} if

$$\mathcal{T} \cup \{\exists \mathbf{x}. \bigwedge \Psi\}$$

is consistent for all finite subsets Ψ of p . An *n-type* p is called *complete* if it is maximal among the collection of *n-types*, in the sense that for any \mathcal{L} -formula $\varphi(\mathbf{x})$, either $\varphi \in p$ or $\varphi \notin p$. We denote the set of complete *n-types* by $S_n(\mathcal{T})$, or $S_n^{\mathcal{M}}(A)$ if $\mathcal{T} = \text{Th}_A(\mathcal{M})$. An element $\mathbf{m} \in \mathcal{M}^n$ *realises* an *n-type* p in \mathcal{M} if $\mathcal{M} \models \varphi(\mathbf{m})$ holds for all φ in p . If no element realises a type, we say that the type is *omitted* in \mathcal{M} .

Example. (i) Let $\mathcal{M} = (\mathbb{Q}, <)$, and consider the formulae $n < x$ for each natural number n . This collection of formulae is a 1-type, as any finite subset is consistent with $\text{Th}_{\mathbb{N}}(\mathbb{Q})$. This type is omitted in \mathbb{Q} as no rational number x satisfies all of the formulae $n < x$ for $n \in \mathbb{N}$. However, this type is realised in an elementary extension of \mathbb{Q} . The realisers can be thought of as imaginary, infinitely large rationals.

IV. Model Theory and Non-Classical Logic

(ii) Consider \mathbb{R} as a structure for the theory of ordered fields. The set of formulae

$$\left\{0 < x < \frac{1}{n} \mid 0 < n \in \mathbb{N}\right\}$$

form a 1-type of infinitesimal real numbers. This type is omitted in \mathbb{R} , but there is an elementary extension realising this type, such as the ultrapower with respect to a free ultrafilter.

(iii) For any \mathcal{L} -structure \mathcal{M} , subset $A \subseteq \mathcal{M}$, and tuple $\mathbf{m} \in \mathcal{M}$, we can form the n -type of all of the \mathcal{L}_A -formulae that hold in \mathcal{M} of \mathbf{m} .

$$\text{tp}^{\mathcal{M}}(\mathbf{m}/A) = \{\varphi(\mathbf{x}) \in \mathcal{L}_A \mid \mathcal{M} \models \varphi(\mathbf{m})\}$$

This is a complete n -type, called *the type of \mathbf{m} over A* . This is a type corresponding to the principal filter on an equivalence class corresponding to an equality formula.

Proposition. Let \mathcal{M} be an \mathcal{L} -structure with $A \subseteq \mathcal{M}$ and let p be an n -type of \mathcal{M} over A . Then there is an elementary extension \mathcal{N} of \mathcal{M} that realises p .

Proof. We use the method of diagrams, and show that

$$\Gamma = p \cup \text{Diag}_{\text{el}}(\mathcal{M})$$

is satisfiable by compactness. Let Δ be a finite subset of Γ , and let

$$\varphi = \bigwedge_{\varphi' \in \Delta \cap p} \varphi'; \quad \psi = \bigwedge_{\psi' \in \Delta \cap \text{Diag}_{\text{el}}(\mathcal{M})} \psi'$$

Note that Δ is satisfiable if and only if

$$\varphi(\mathbf{x}, \mathbf{a}) \wedge \psi(\mathbf{a}', \mathbf{b})$$

is satisfiable, where $\mathbf{a}, \mathbf{a}' \in A$ and $\mathbf{b} \in \mathcal{M} \setminus A$, and

$$\varphi \in p; \quad \mathcal{M} \models \psi(\mathbf{a}', \mathbf{b})$$

As p is an n -type, there is an \mathcal{L}_A -structure \mathcal{N}_0 that satisfies $p \cup \text{Th}_A(\mathcal{M})$. As $\mathcal{M} \models \psi(\mathbf{a}', \mathbf{b})$, we have $\mathcal{M} \models \exists \mathbf{y}. \psi(\mathbf{a}', \mathbf{y})$. Note that this is an \mathcal{L}_A -formula, so

$$(\exists \mathbf{y}. \psi(\mathbf{a}', \mathbf{y})) \in \text{Th}_A(\mathcal{M})$$

Hence,

$$\mathcal{N}_0 \models \varphi(\mathbf{c}, \mathbf{a}) \exists \mathbf{y}. \psi(\mathbf{a}', \mathbf{y})$$

for some $\mathbf{c} \in \mathcal{N}_0$. Note that \mathcal{N}_0 is an \mathcal{L}_A -structure, not an $\mathcal{L}_{\mathcal{M}}$ -structure. However, by interpreting \mathbf{b} in \mathcal{N}_0 as the witness \mathbf{y} to $\exists \mathbf{y}. \psi(\mathbf{a}', \mathbf{y})$, we make \mathcal{N}_0 into an $\mathcal{L}_{\mathcal{M}}$ -structure; elements of \mathcal{M} not in A or \mathbf{b} are interpreted arbitrarily. In this $\mathcal{L}_{\mathcal{M}}$ -structure, Δ is satisfiable. Thus Γ is satisfiable by compactness.

Now, let \mathcal{N} be an $\mathcal{L}_{\mathcal{M}}$ -structure satisfying Γ , so \mathcal{N} is an elementary extension of \mathcal{M} . As \mathcal{N} satisfies p , there must be a tuple $\mathbf{n} \in \mathcal{N}$ with $\mathcal{N} \models \varphi(\mathbf{n})$ for each $\varphi \in p$. In other words, \mathbf{n} realises p in \mathcal{N} . \square

Corollary. An n -type p of \mathcal{M} over $A \subseteq \mathcal{M}$ is complete if and only if there is an elementary extension \mathcal{N} of \mathcal{M} and some $\mathbf{a} \in \mathcal{N}$ such that $p = \text{tp}^{\mathcal{N}}(\mathbf{a}/A)$.

Proof. If \mathcal{N} is an elementary extension of \mathcal{M} and $\mathbf{a} \in \mathcal{N}$, then

$$\text{tp}^{\mathcal{N}}(\mathbf{a}/A) \in S_n^{\mathcal{N}}(A) = S_n^{\mathcal{M}}(A)$$

as the extension is elementary.

Conversely, if p is a complete n -type, then by the previous result, there is an elementary extension \mathcal{N} of \mathcal{M} with a tuple \mathbf{a} realising the type. As p is complete, every \mathcal{L}_A -formula φ , either $\varphi \in p$ or $\varphi \notin p$, but not both. If $\varphi \in \text{tp}^{\mathcal{N}}(\mathbf{a}/A)$, then $\mathcal{N} \models \varphi(\mathbf{a})$, so we cannot have $\varphi \notin p$, thus $\varphi \in p$. Conversely, if $\varphi \in p$, then $\mathcal{N} \models \varphi(\mathbf{a})$ as \mathbf{a} realises p , so $\varphi \in \text{tp}^{\mathcal{N}}(\mathbf{a}/A)$. Thus $p = \text{tp}^{\mathcal{N}}(\mathbf{a}/A)$ as required. \square

4.2. Stone spaces

Let \mathcal{M} be an \mathcal{L} -structure and let $A \subseteq \mathcal{M}$. For each formula φ on n variables, we consider the set of all complete types that include this formula, denoted

$$\llbracket \varphi \rrbracket = \{p \in S_n^{\mathcal{M}}(A) \mid \varphi \in p\}$$

Note that

$$\llbracket \varphi \vee \psi \rrbracket = \llbracket \varphi \rrbracket \cup \llbracket \psi \rrbracket; \quad \llbracket \varphi \wedge \psi \rrbracket = \llbracket \varphi \rrbracket \cap \llbracket \psi \rrbracket$$

These serve as the basic open sets for a topology on $S_n^{\mathcal{M}}(A)$, so an open set is an arbitrary union of open sets of this form. Moreover, each of these basic open sets $\llbracket \varphi \rrbracket$ is the complement of another basic open set $\llbracket \neg \varphi \rrbracket$, so these open sets are also closed. The $S_n^{\mathcal{M}}(A)$ are called *Stone spaces*, which are compact and totally disconnected topological spaces.

Example. Let F be an algebraically closed field, and let k be a subfield of F . The complete n -types $p \in S_n^F(k)$ are determined by the prime ideals of $k[x_1, \dots, x_n]$. For such a type p , we can define a prime ideal by

$$I_p = \{f \in k[x_1, \dots, x_n] \mid (f(x_1, \dots, x_n) = 0) \in p\}$$

These ideals are prime, and all prime ideals arise in this way. The map $p \mapsto I_p$ is a continuous bijection from the type space $S_n^F(k)$ to the prime spectrum $\text{Spec } k[x_1, \dots, x_n]$ with the Zariski topology. Also, note that $|S_n^F(k)| \leq |k| + \aleph_0$ by Hilbert's basis theorem.

4.3. Isolated points

Recall that a point p in a topological space is *isolated* if $\{p\}$ is an open set. If p is isolated in $S_n^{\mathcal{M}}(A)$, then

$$\{p\} = \bigcup_I \llbracket \varphi_i \rrbracket$$

so as $\{p\}$ is a singleton, there must be a single formula $\varphi = \varphi_i$ such that $\{p\} = \llbracket \varphi \rrbracket$; we say that φ *isolates* the type.

IV. Model Theory and Non-Classical Logic

Definition. Let \mathcal{T} be an \mathcal{L} -theory. We say that a formula $\varphi(x_1, \dots, x_n)$ isolates the n -type p of \mathcal{T} if $\mathcal{T} \cup \{\varphi\}$ is satisfiable, and

$$\mathcal{T} \models \forall \mathbf{x}. (\varphi(\mathbf{x}) \rightarrow \psi(\mathbf{x}))$$

for all $\psi \in p$.

Proposition. If φ isolates p , then p is realised in any model of $\mathcal{T} \cup \{\exists \mathbf{x}. \varphi(\mathbf{x})\}$. In particular, if \mathcal{T} is a complete theory, then all isolated types are realised.

Proof. If \mathcal{M} is a model of \mathcal{T} and there exists \mathbf{a} such that $\mathcal{M} \models \varphi(\mathbf{a})$, then clearly \mathbf{a} realises p in \mathcal{M} . If \mathcal{T} is complete, then either

$$\mathcal{T} \models \exists \mathbf{x}. \varphi(\mathbf{x})$$

or

$$\mathcal{T} \models \forall \mathbf{x}. \neg \varphi(\mathbf{x})$$

If φ isolates \mathcal{T} , then $\mathcal{T} \cup \{\varphi\}$ is satisfiable by definition, so the latter case is impossible. \square

4.4. Omitting types

Theorem (omitting types theorem). Let \mathcal{L} be a countable language and let \mathcal{T} be an \mathcal{L} -theory. Let p be a non-isolated n -type of \mathcal{T} . Then there is a countable model $\mathcal{M} \models \mathcal{T}$ that omits p .

Proof. Let $C = \{c_0, c_1, \dots\}$ be a countable set of new constants. We expand \mathcal{T} to a consistent \mathcal{L}_C -theory \mathcal{T}^* by adding recursively defined sentences $\theta_0, \theta_1, \dots$. We will do this in such a way that $\theta_t \rightarrow \theta_s$ for all $s < t$. To build the θ , we first enumerate the n -tuples $C^n = \{\mathbf{d}_0, \mathbf{d}_1, \dots\}$, and enumerate the \mathcal{L}_C -sentences $\varphi_0, \varphi_1, \dots$.

Start with $\theta_0 = \forall x. x = x$, which is trivially true. Suppose we have already constructed θ_s in such a way that $\mathcal{T} \cup \{\theta_s\}$ is consistent.

First, suppose $s = 2i$. These sentences will be designed to turn C into the domain of an elementary substructure of some model of \mathcal{T}^* . Suppose that $\varphi_i = \exists x. \psi(x)$ is existential, with parameters in C as φ is an \mathcal{L}_C -formula. Suppose also that $\mathcal{T} \models \theta_s \rightarrow \varphi_i$. As only finitely many constants from C have been used so far, we can find some unused $c \in C$. Let

$$\theta_{s+1} = \theta_s \wedge \psi(c)$$

If \mathcal{N} models $\mathcal{T} \cup \{\theta_s\}$, then there is a witness to ψ in \mathcal{N} , so we can interpret c as this witness. Thus, \mathcal{N} models $\mathcal{T} \cup \{\theta_{s+1}\}$, so this theory is consistent. If φ_i is not existential, or $\mathcal{T} \not\models \theta_s \rightarrow \varphi_i$, then define $\theta_{s+1} = \theta_s$.

Now, suppose $s = 2i + 1$. These sentences will be designed to ensure that C omits p . Let $\mathbf{d}_i = (e_1, \dots, e_n)$. Remove every occurrence of the e_j from θ_s by replacing it with the variable x_j , and replace every occurrence of other constants in C with a fresh variable x_c , together

with a quantifier $\exists x_c$ in front of the formula. This yields an \mathcal{L} -formula $\psi(x_1, \dots, x_n)$. For example, if

$$\theta_s = \forall x. \exists y. (rx + e_1 e_2 = y^2 + te_2); \quad r \neq t \in C$$

then

$$\psi(x_1, x_2) = \exists x_r. \exists x_t. \forall x. \exists y. (x_r x + x_1 x_2 = y^2 + x_t x_2)$$

As p is not isolated, there is no \mathcal{L} -formula that isolates it, so there must be some $\varphi(\mathbf{x}) \in p$ that is not implied by $\psi(\mathbf{x})$; otherwise ψ would isolate the type p . We define θ_{s+1} in such a way that \mathbf{d}_i cannot realise p .

$$\theta_{s+1} = \theta_s \wedge \neg\varphi(\mathbf{d}_i)$$

This is consistent, because there must be some $\mathbf{n} \in \mathcal{N} \models \mathcal{T}$ such that

$$\mathcal{N} \models \psi(\mathbf{n}) \wedge \neg\varphi(\mathbf{n})$$

and we can turn \mathcal{N} into an \mathcal{L}_C -structure that models θ_{s+1} by interpreting \mathbf{d}_i as \mathbf{n} , and interpreting the constants in C but not in \mathbf{d} as the respective witnesses to the existential statements $\exists x_c$ within ψ .

Let \mathcal{T}^* be \mathcal{T} together with all of the θ_s . Note that each $\mathcal{T} \cup \{\theta_s\}$ is consistent, and each θ_{s+1} implies θ_s , so by compactness, \mathcal{T}^* must be consistent. Moreover, if \mathcal{M} is a model of \mathcal{T}^* , the construction of θ_{2i+1} ensures that C has a witness to φ_i that holds in \mathcal{M} . Thus, by the Tarski–Vaught test, C is the domain of an elementary substructure of \mathcal{M} . If $\mathbf{c} \in C \models \mathcal{T}^*$, then $\mathbf{c} = \mathbf{d}_i$ for some i . As $C \models \theta_{2i+2}$, we have $\neg\varphi(\mathbf{c})$ for some φ in the type p . Hence \mathbf{c} cannot realise the type p in C . \square

Remark. The proof can be generalised to omit countably many types at the same time.

5. Indiscernibles

5.1. Introduction

Given a linear order η , we will write $[\eta]^k$ for the set of ordered k -tuples in η :

$$[\eta]^k = \{\mathbf{a} \in \eta^k \mid a_0 <^\eta a_1 <^\eta \dots <^\eta a_{k-1}\}$$

Definition. Let \mathcal{M} be an \mathcal{L} -structure, let Φ be a set of \mathcal{L} -formulae, and let η be a strict chain of elements of \mathcal{M} . We say that η is Φ -indiscernible in \mathcal{M} if

$$\mathcal{M} \models \varphi(\mathbf{a}) \leftrightarrow \varphi(\mathbf{b})$$

for all $\mathbf{a}, \mathbf{b} \in [\eta]^k$ of the correct length and $\varphi \in \Phi$. We simply say that η is a sequence of indiscernibles if the above holds where Φ is the set of every \mathcal{L} -formula.

Example. (i) Any linearly ordered basis \mathcal{B} for a vector space provides a sequence of indiscernibles. Indeed, given $\mathbf{a}, \mathbf{b} \in [\mathcal{B}]^k$, there is an automorphism of the vector space that maps \mathbf{a} to \mathbf{b} .

(ii) Any chain of algebraically independent elements in a field $k \models \text{ACF}_0$ is a sequence of indiscernibles.

(iii) If R is a ring, then the variables X_1, \dots, X_n form a set of indiscernibles of $R[X_1, \dots, X_n]$.

Definition. An *Ehrenfeucht–Mostowski functor* is a mapping F that takes each linear order η to an \mathcal{L} -structure $F(\eta)$, and each order embedding $g : \eta \rightarrow \varepsilon$ to an embedding of \mathcal{L} -structures $F(g) : F(\eta) \rightarrow F(\varepsilon)$, in such a way that

- (i) each η generates $F(\eta)$, that is, $\eta \subseteq F(\eta)$ as sets, and every element of $F(\eta)$ is of the form $t^{F(\eta)}(\mathbf{a})$ where $t(\mathbf{x})$ is an \mathcal{L} -term and $\mathbf{a} \in [\eta]^k$;
- (ii) for each order embedding $g : \eta \rightarrow \varepsilon$, the embedding of \mathcal{L} -structures $F(g)$ extends g ;
- (iii) for every linear order η , we have $F(1_\eta) = 1_{F(\eta)}$;
- (iv) for each composable pair of embeddings f, g , we have $F(g \circ f) = F(g)F(f)$.

In particular, every automorphism of a linear order η induces an automorphism of $F(\eta)$.

Proposition (sliding property). Let F be an Ehrenfeucht–Mostowski functor, let η, ε be linear orders, and let $\mathbf{a} \in [\eta]^k, \mathbf{b} \in [\varepsilon]^k$. Then for every quantifier-free formula $\varphi(x_1, \dots, x_k)$, we have

$$F(\eta) \models \varphi(\mathbf{a}) \iff F(\varepsilon) \models \varphi(\mathbf{b})$$

Proof. Embed η and ε into some linear order ρ in which \mathbf{a} and \mathbf{b} are identified. Let $f : \eta \rightarrow \rho$ and $g : \varepsilon \rightarrow \rho$ be the embeddings. Suppose that $F(\eta) \models \varphi(\mathbf{a})$. As embeddings preserve quantifier-free formulae and the map $F(f) : F(\eta) \rightarrow F(\rho)$ extends f , we must have that $F(\rho) \models \varphi(f(\mathbf{a}))$. As $f(\mathbf{a}) = g(\mathbf{b})$, we must have $F(\rho) \models \varphi(g(\mathbf{b}))$, and so for the same reason, $F(\varepsilon) \models \varphi(\mathbf{b})$. \square

We see that the chain $\eta \subseteq F(\eta)$ is indiscernible by quantifier-free formulas.

Definition. Let \mathcal{M} be an \mathcal{L} -structure containing a linear order $\eta \subseteq \mathcal{M}$ as sets. Then, we define the theory of η in \mathcal{M} , denoted $\text{Th}(\mathcal{M}, \eta)$, to be the set of all \mathcal{L} -formulae $\varphi(\mathbf{x})$ that are satisfiable in \mathcal{M} by every ordered tuple $\mathbf{a} = a_0 < \dots < a_{k-1}$ in η . The theory $\text{Th}(F)$ of an Ehrenfeucht–Mostowski functor F is the set of all \mathcal{L} -formulae $\varphi(\mathbf{x})$ such that $F(\eta) \models \varphi(\mathbf{a})$ for every linear order η and ordered tuple \mathbf{a} in η .

Lemma. Let η be an infinite linear order, let F be an Ehrenfeucht–Mostowski functor, and let φ be a universal sentence that is true in $F(\eta)$. Then $\varphi \in \text{Th}(F)$.

Proof. Let $\varphi = \forall \mathbf{x}. \psi(\mathbf{x})$ where ψ is quantifier-free. Let ε be a linear order, and let $\mathbf{a} \in F(\varepsilon)$; we need to show $F(\varepsilon) \models \psi(\mathbf{a})$. As ε generates $F(\varepsilon)$, there is a finite suborder ε_0 such that $\mathbf{a} \in F(\varepsilon_0)$. But η is infinite, so there is an embedding $f : \varepsilon_0 \rightarrow \eta$. By assumption, $F(f)(\mathbf{a})$ satisfies ψ in $F(\eta)$, so $F(\varepsilon_0) \models \psi(\mathbf{a})$, as ψ is quantifier-free so is preserved under substructures. Similarly, $F(\varepsilon) \models \psi(\mathbf{a})$, as required. \square

5.2. Existence of Ehrenfeucht–Mostowski functors

Lemma (stretching property). Let \mathcal{M} be an \mathcal{L} -structure that contains the linear order ω as a generating set. Suppose that ω is indiscernible by quantifier-free formulae. Then there is an Ehrenfeucht–Mostowski functor F such that $\mathcal{M} = F(\omega)$. Moreover, if G is another such functor, then there is an isomorphism $\alpha : F(\eta) \rightarrow G(\eta)$ for each linear order η , and $\alpha|_\eta = 1_\eta$.

F is unique up to natural isomorphism.

Definition. Let F be an Ehrenfeucht–Mostowski functor, and let \mathcal{T} be a theory. The models of \mathcal{T} that are of the form $F(\eta)$ are called *Ehrenfeucht–Mostowski models* of \mathcal{T} .

Theorem (Ramsey). Let X be a countable linear order, and let k, n be positive integers. Then for every function $f : [X]^k \rightarrow n$, there is an infinite subset $Y \subseteq X$ such that f is constant on $[Y]^k$.

We will use Ramsey’s theorem to show that Ehrenfeucht–Mostowski models for Skolem theories with infinite models always exist.

Lemma. Let F be an Ehrenfeucht–Mostowski functor such that $\text{Th}(F(\omega))$ is Skolem. Then $\text{Th}(F)$ includes either $\varphi(\mathbf{x})$ or $\neg\varphi(\mathbf{x})$ for every \mathcal{L} -formula $\varphi(\mathbf{x})$. In particular, all of the $F(\eta)$ are elementarily equivalent, and each linear order η is indiscernible in $F(\eta)$.

Proof. Since $\text{Th}(F(\omega))$ is Skolem, it admits a universal axiomatisation. Moreover, every formula is equivalent to a quantifier-free formula modulo $\text{Th}(F(\omega))$. The result then follows from the sliding property and the lemma on universal sentences. \square

Theorem (Ehrenfeucht–Mostowski theorem). Let \mathcal{M} be an \mathcal{L} -structure, and suppose that $\text{Th}(\mathcal{M})$ is Skolem. If η is infinite linear order that is contained as a set in \mathcal{M} , then there is an Ehrenfeucht–Mostowski functor F in \mathcal{L} whose theory expands $\text{Th}(\mathcal{M}, \eta)$.

IV. Model Theory and Non-Classical Logic

Proof. We want to build a theory expanding $\text{Th}(\mathcal{M}, \eta)$, whose models include an indiscernible copy of ω . First, expand \mathcal{L} to add ω -many constants $C = \{c_i \mid i \in \omega\}$, and we build an \mathcal{L}_C -theory \mathcal{T} with the following axioms:

- (i) $\varphi(\mathbf{a}) \leftrightarrow \varphi(\mathbf{b})$, for each \mathcal{L} -formula $\varphi(\mathbf{x})$ and ordered tuples $\mathbf{a}, \mathbf{b} \in [C]^{|\mathbf{x}|}$;
- (ii) $\varphi(c_0, \dots, c_{k-1})$, for each formula $\varphi(x_0, \dots, x_{k-1})$ in $\text{Th}(\mathcal{M}, \eta)$.

We will show that this theory has a model by compactness. Let \mathcal{U} be a finite subset of \mathcal{T} , and list the formulae in \mathcal{U} as $\varphi_0, \dots, \varphi_{m-1}$. Note that there is some finite k such that the new constants that show up in the formulae in \mathcal{U} are among c_0, \dots, c_{k-1} . By adding redundant variables, we may assume that each of these formulae all have free variables c_0, \dots, c_{k-1} for simplicity.

Define an equivalence relation \sim on $[\eta]^k$ by declaring that $\mathbf{a} \sim \mathbf{b}$ if $\mathcal{M} \models \varphi_j(\mathbf{a})$ if and only if $\mathcal{M} \models \varphi_j(\mathbf{b})$ for each $j < m$. This equivalence relation partitions $[\eta]^k$ into finitely many equivalence classes. Hence, by Ramsey's theorem, there is an infinite sequence $\mathbf{e} = e_0 < e_1 < \dots < e_{2k-1}$ in η such that any two ordered k -tuples extracted from \mathbf{e} are in the same equivalence class. We can interpret each c_j in \mathcal{M} as e_j for each $j < k$, making \mathcal{M} into an \mathcal{L}_C -structure that models \mathcal{U} .

Let \mathcal{N} be a model of \mathcal{T} . The new constants c_i must be interpreted as different elements of \mathcal{N} , as $\text{Th}(\mathcal{M}, \eta)$ includes the sentence $x_0 \neq x_1$. Hence \mathcal{N} contains a copy of ω , by seeing c_i in \mathcal{N} as i . Consider \mathcal{N}^* , which is the \mathcal{L} -reduct of \mathcal{N} , and let $\mathcal{S} = \langle \omega \rangle_{\mathcal{N}^*}$. Note that $\text{Th}(\mathcal{M}, \eta)$ is contained in $\text{Th}(\mathcal{N}^*, \omega)$. This in particular implies that $\text{Th}_{\mathcal{L}}(\mathcal{N}^*)$ is Skolem, as $\text{Th}(\mathcal{M})$ is Skolem and $\text{Th}(\mathcal{M}) \subseteq \text{Th}(\mathcal{M}, \eta)$. It then follows that \mathcal{S} is an elementary substructure of \mathcal{N}^* , and is generated by ω . Then, $\text{Th}(\mathcal{M}, \eta) \subseteq \text{Th}(\mathcal{S}, \omega)$. Finally, sentences in \mathcal{T} ensure that ω is indiscernible in \mathcal{S} by construction, so the stretching lemma gives an Ehrenfeucht–Mostowski functor F with $\mathcal{S} = F(\omega)$, which completes the proof by the previous lemma. \square

6. Intuitionistic logic and lambda calculi

6.1. The Brouwer–Heyting–Kolmogorov interpretation

We will construct a system of logic in which every proof contains evidence of its truth. Our system will have the following properties, known as the Brouwer–Heyting–Kolmogorov interpretation.

- (i) \perp has no proof.
- (ii) To prove $\varphi \wedge \psi$, one must provide a proof of φ together with a proof of ψ .
- (iii) To prove $\varphi \rightarrow \psi$, one must provide a mechanism for translating a proof of φ into a proof of ψ . In particular, to prove $\neg\varphi$, we must provide a way to turn a proof of φ into a contradiction.
- (iv) To prove $\varphi \vee \psi$, we must specify either φ or ψ , and then provide a proof for it. Note that in a classical setting, a proof of $\varphi \vee \psi$ need not specify which of the two disjuncts is true.
- (v) The law of the excluded middle LEM, which states $\varphi \vee \neg\varphi$, is not valid. If this held for some proposition, we could decide whether the proposition was true or its negation is true, because any proof of $\varphi \vee \neg\varphi$ contains this information.
- (vi) To prove $\exists x. \varphi(x)$, one must provide a term t together with a proof of $\varphi(t)$.
- (vii) To prove $\forall x. \varphi(x)$, one must provide a mechanism that converts any term t into a proof of $\varphi(t)$.

This will be called *intuitionistic (propositional) logic* IPC.

Theorem (Diaconescu). In intuitionistic ZF set theory, the law of the excluded middle LEM can be deduced from the axiom of choice AC.

Proof. Let φ be a proposition; we want a proof of $\varphi \vee \neg\varphi$. Using the axiom of separation, we have proofs that the following sets exist.

$$A = \{x \in \{0, 1\} \mid \varphi \vee (x = 0)\}; \quad B = \{x \in \{0, 1\} \mid \varphi \vee (x = 1)\}$$

These sets are *inhabited*: there exists an element in each of them; in particular, $0 \in A$ and $1 \in A$ are intuitionistically valid. Note that being inhabited is strictly stronger than being nonempty in intuitionistic logic. This is because any proof that a set is inhabited contains information about an element in the set. The set $\{A, B\}$ is a family of inhabited sets, so by the axiom of choice, we have a choice function $f : \{A, B\} \rightarrow A \cup B$, and we have a proof that $f(A) \in A$ and $f(B) \in B$. Thus, we have a proof of

$$(\varphi \vee (f(A) = 0)) \wedge (\varphi \vee (f(B) = 1))$$

We also have a proof that $f(A), f(B) \in \{0, 1\}$. In particular, we either have a proof that $f(A) = 0$ or we have a proof that $f(A) = 1$, and the same holds for B . We have the following cases.

IV. Model Theory and Non-Classical Logic

- (i) Suppose we have a proof that $f(A) = 1$. Then we have a proof of $\varphi \vee (1 = 0)$, so we must have a proof of φ .
- (ii) Suppose we have a proof that $f(B) = 0$. Then similarly we have a proof of $\varphi \vee (0 = 1)$, so we must have a proof of φ .
- (iii) Suppose we have proofs that $f(A) = 0$ and $f(B) = 1$. We will prove $\neg\varphi$. Suppose that we have a proof of φ . Then from a proof of $\varphi \vee (x = 0)$ or $\varphi \vee (x = 1)$ we can derive a proof of the other, so by the axiom of extensionality, $A = B$. Then $0 = f(A) = f(B) = 1$ as f is a function, giving a contradiction. Thus, we have constructed a proof of $\neg\varphi$.

We can always specify a proof of φ or a proof of $\neg\varphi$, so we have $\varphi \vee \neg\varphi$. \square

Remark. (i) Intuitionistic mathematics is more general than classical mathematics, because it operates on fewer assumptions.

- (ii) Notions that are classically conflated may be different in intuitionistic logic. For example, there is no classical distinction between inhabited and nonempty sets, but they are not the same in intuitionistic logic. Other examples include finiteness, or disequality and apartness.
- (iii) Intuitionistic proofs have computational content attached to them, but classical proofs may not.
- (iv) Intuitionistic logic is the internal logic of an arbitrary topos.

6.2. Natural deduction

We will use the notation $\Gamma \vdash \varphi$, or $\Gamma \vdash_{\text{IPC}} \varphi$, to denote that the set of *open assumptions* Γ let us conclude φ . Γ is also called the *context*. We will inductively define this provability relation. Some rules, called *introduction rules*, let us construct proofs.

$$\begin{array}{c} \wedge\text{-I} \\ \frac{\Gamma \vdash A \quad \Gamma \vdash B}{\Gamma \vdash A \wedge B} \end{array} \qquad \begin{array}{c} \vee\text{-I} \\ \frac{\Gamma \vdash A}{\Gamma \vdash A \vee B} \end{array} \qquad \begin{array}{c} \vee\text{-I} \\ \frac{\Gamma \vdash B}{\Gamma \vdash A \vee B} \end{array}$$

Dually, some rules, called *elimination rules*, let us extract information from proofs.

$$\begin{array}{c} \wedge\text{-E} \\ \frac{\Gamma \vdash A \wedge B}{\Gamma \vdash A} \end{array} \qquad \begin{array}{c} \wedge\text{-E} \\ \frac{\Gamma \vdash A \wedge B}{\Gamma \vdash B} \end{array} \qquad \begin{array}{c} \vee\text{-E} \\ \frac{\Gamma, A \vdash C \quad \Gamma, B \vdash C \quad \Gamma \vdash A \vee B}{\Gamma \vdash C} \end{array}$$

We now define the principle of explosion, which is an elimination rule for \perp . We do not construct an introduction rule for \perp .

$$\begin{array}{c} \perp\text{-E} \\ \frac{\Gamma \vdash \perp}{\Gamma \vdash A} \end{array}$$

6. Intuitionistic logic and lambda calculi

We now define the introduction and elimination rules for implication. The elimination rule is known as *modus ponens*.

$$\begin{array}{c} \rightarrow\text{-I} \\ \frac{\Gamma, A \vdash B}{\Gamma \vdash A \rightarrow B} \end{array} \qquad \begin{array}{c} \rightarrow\text{-E} \\ \frac{\Gamma \vdash A \rightarrow B \quad \Gamma \vdash A}{\Gamma \vdash B} \end{array}$$

We finally define a rule called the *axiom schema*, that allows us to prove our assumptions.

$$\begin{array}{c} \text{Ax} \\ \hline \Gamma, A \vdash A \end{array}$$

If an inference rule moves an assumption out of the context, we say that the assumption is *discharged* or *closed*. We are allowed to drop assumptions that we do not use; this is called the *weakening* rule. We obtain classical propositional logic CPC by additionally adding one of the following two rules.

$$\begin{array}{c} \text{LEM} \\ \hline \Gamma \vdash A \vee \neg A \end{array} \qquad \begin{array}{c} \neg\neg\text{-E} \\ \frac{\Gamma, \neg A \vdash \perp}{\Gamma \vdash A} \end{array}$$

We will additionally use the informal notation

$$\frac{\begin{array}{c} [A] \\ \vdots \\ X \end{array} \quad \begin{array}{c} [B] \\ \vdots \\ Y \end{array}}{C} (A, B)$$

to mean that if we can prove X assuming A and we can prove Y assuming B , then we can infer C by discharging the open assumptions A and B . For example, we can write an instance of $\rightarrow\text{-I}$ as

$$\frac{\begin{array}{c} \Gamma, [A] \\ \vdots \\ B \end{array}}{\Gamma \vdash A \rightarrow B} (A)$$

To extend this to intuitionistic predicate logic IQC, we need to add rules for quantifiers.

$$\begin{array}{c} \exists\text{-I} \\ \frac{\Gamma \vdash \varphi[x := t]}{\Gamma \vdash \exists x. \varphi(x)} \end{array} \qquad \begin{array}{c} \forall\text{-I} \\ \frac{\Gamma \vdash \varphi \quad x \text{ not free in } \Gamma}{\Gamma \vdash \forall x. \varphi} \end{array}$$

$$\begin{array}{c} \exists\text{-E} \\ \frac{\Gamma \vdash \exists x. \varphi \quad \Gamma, \varphi \vdash \psi \quad x \text{ not free in } \Gamma}{\Gamma \vdash \psi} \end{array} \qquad \begin{array}{c} \forall\text{-E} \\ \frac{\Gamma \vdash \forall x. \varphi}{\Gamma \vdash \varphi[x := t]} \end{array}$$

IV. Model Theory and Non-Classical Logic

Example. We will show that $\vdash_{\text{IPC}} A \wedge B \rightarrow B \wedge A$.

$$\frac{\frac{\frac{[A \wedge B]}{B} \wedge\text{-E} \quad \frac{[A \wedge B]}{A} \wedge\text{-E}}{B \wedge A} \wedge\text{-I}}{A \wedge B \rightarrow B \wedge A} \rightarrow\text{-I}$$

Example. We will show that the logical axioms

$$\varphi \rightarrow (\psi \rightarrow \varphi); \quad (\varphi \rightarrow (\psi \rightarrow \chi)) \rightarrow ((\varphi \rightarrow \psi) \rightarrow (\varphi \rightarrow \chi))$$

are intuitionistically valid.

$$\frac{\frac{\frac{[\varphi]}{\varphi} \text{Ax} \quad [\psi]}{\psi \rightarrow \varphi} (\rightarrow\text{-I}, \psi)}{\varphi \rightarrow (\psi \rightarrow \varphi)} (\rightarrow\text{-I}, \varphi)$$

For the second axiom,

$$\frac{\frac{\frac{[\varphi \rightarrow (\psi \rightarrow \chi)] \quad [\varphi]}{\psi \rightarrow \chi} \rightarrow\text{-E} \quad \frac{\frac{[\varphi \rightarrow \psi] \quad [\varphi]}{\psi} \rightarrow\text{-E}}{\chi} \rightarrow\text{-E}}{\varphi \rightarrow \chi} (\varphi)}{(\varphi \rightarrow \psi) \rightarrow (\varphi \rightarrow \chi)} (\varphi \rightarrow \psi)}{(\varphi \rightarrow (\psi \rightarrow \chi)) \rightarrow ((\varphi \rightarrow \psi) \rightarrow (\varphi \rightarrow \chi))} (\varphi \rightarrow (\psi \rightarrow \chi))$$

Lemma. If $\Gamma \vdash_{\text{IPC}} \varphi$, then $\Gamma, \psi \vdash_{\text{IPC}} \varphi$. Moreover, if p is a primitive proposition and ψ is any proposition, then

$$\Gamma[p := \psi] \vdash_{\text{IPC}} \varphi[p := \psi]$$

Proof. This follows easily by induction over the length of the proof. □

6.3. The simply typed lambda calculus

For now, we will assume we are given a set Π of *simple types*, generated by the grammar

$$\Pi ::= \mathcal{U} \mid \Pi \rightarrow \Pi$$

where \mathcal{U} is a countable set of *primitive types* or *type variables*.

6. Intuitionistic logic and lambda calculi

Let V be an infinite set of variables. The set Λ_Π of *simply typed λ -terms* is defined by the grammar

$$\Lambda_\Pi ::= V \mid \underbrace{\lambda V : \Pi. \Lambda_\Pi}_{\lambda\text{-abstraction}} \mid \underbrace{\Lambda_\Pi \Lambda_\Pi}_{\lambda\text{-application}}$$

A *context* Γ is a set of pairs $\{x_1 : \tau_1, \dots, x_n : \tau_n\}$, where the x_i are distinct variables, and the τ_n are types. We write \mathcal{C} for the set of all contexts. Given a context $\Gamma \in \mathcal{C}$, we also write $\Gamma, x : \tau$ for the context $\Gamma \cup \{x : \tau\}$. The *domain* of Γ is the set $\text{dom } \Gamma$ of variables that appear in Γ ; similarly, the *range* of Γ is the set $|\Gamma|$ of types that appear in Γ .

The *typability relation* $(-) \Vdash (-) : (-)$ is a relation on $\mathcal{C} \times \Lambda_\Pi \times \Pi$, defined recursively using the following rules.

- (i) For every context Γ , variable $x \notin \text{dom } \Gamma$, and type τ , we have $\Gamma, x : \tau \Vdash x : \tau$.
- (ii) Let Γ be a context, $x \notin \text{dom } \Gamma$, let σ, τ be types, and let M be a λ -term. If $\Gamma, x : \sigma \Vdash M : \tau$, then $\Gamma \Vdash (\lambda x : \sigma. M) : \sigma \rightarrow \tau$.
- (iii) Let Γ be a context, σ, τ be types, and let M and N be λ -terms. If $\Gamma \Vdash M : (\sigma \rightarrow \tau)$ and $\Gamma \Vdash N : \sigma$, then $\Gamma \Vdash (MN) : \tau$.

We will refer to the λ -calculus of Λ_Π with this typability relation as $\lambda(\rightarrow)$.

An occurrence of a variable x in a λ -abstraction is called *bound*, otherwise it is called *free*. A term with no free variables is called *closed*. λ -terms that differ only in the names of bound variables are called *α -equivalent*, so for example, $(\lambda x : \sigma. x)$ and $(\lambda y : \sigma. y)$ are α -equivalent. Whenever it is convenient, we will replace terms with α -equivalent terms to avoid reusing variable names.

If M and N are λ -terms and x is a variable, we can define the *substitution* of N for x in M recursively:

- (i) $x[x := N] = N$;
- (ii) $y[x := N] = y$ if $x \neq y$;
- (iii) $(\lambda y : \sigma. M)[x := N] = (\lambda y : \sigma. M[x := N])$ if $x \neq y$ (which can be done without loss of generality by α -equivalence);
- (iv) $(PQ)[x := N] = (P[x := N])(Q[x := N])$.

We define the *β -reduction* relation \rightarrow_β on Λ_Π to be the smallest relation that is closed under the following rules:

- (i) $(\lambda x : \sigma. P)Q \rightarrow_\beta P[x := Q]$;
- (ii) if $P \rightarrow_\beta P'$, then for any $x \in V$ and $\sigma \in \Pi$, we have $(\lambda x : \sigma. P) \rightarrow_\beta (\lambda x : \sigma. P')$;
- (iii) if $P \rightarrow_\beta P'$ and Z is a λ -term, then $PZ \rightarrow_\beta P'Z$ and $ZP \rightarrow_\beta ZP'$.

IV. Model Theory and Non-Classical Logic

We define the β -equivalence relation \equiv_β to be the smallest equivalence relation containing \rightarrow_β . For example, we have

$$(\lambda x : \mathbb{Z}. (\lambda y : \tau. x)) 2 \equiv_\beta (\lambda y : \tau. 2)$$

An expression $(\lambda x : \sigma. P) Q$ to be β -reduced is called a β -redex; the resulting term $P[x := Q]$ is called its β -reduct or β -contractum. If no β -reductions can be carried out on a λ -term, we say that the term is in β -normal form. We write $M \rightarrow_\beta N$ if M reduces to N after potentially multiple applications of β -reduction.

If x is not free in P , the term $(\lambda x : \sigma. (P x))$ is said to η -reduce to P , written $(\lambda x : \sigma. (P x)) \rightarrow_\eta P$, and we say that $(\lambda x : \sigma. (P x))$ and P are η -equivalent.

By convention, we will write

- (i) KLM for $(KL)M$;
- (ii) $\lambda x : \sigma. \lambda y : \tau. M$ for $\lambda x : \sigma. (\lambda y : \tau. M)$;
- (iii) $\lambda x : \sigma. MN$ for $\lambda x : \sigma. (MN)$;
- (iv) $M \lambda x : \sigma. N$ for $M(\lambda x : \sigma. N)$.

6.4. Basic properties

The following technical lemmas can be proven by induction.

Lemma (generation lemma). (i) For every variable x , context Γ , and type σ , if $\Gamma \Vdash x : \sigma$, then $x : \sigma \in \Gamma$.

(ii) If $\Gamma \Vdash (\lambda x : \tau. N) : \sigma$, then there is a type ρ such that $\Gamma, x : \tau \Vdash N : \rho$, and $\sigma = (\tau \rightarrow \rho)$.

(iii) If $\Gamma \Vdash (MN) : \sigma$, then there is a type τ such that $\Gamma \Vdash M : \tau \rightarrow \sigma$ and $\Gamma \Vdash N : \tau$.

Lemma (free variables lemma). Suppose that $\Gamma \Vdash M : \sigma$. Then


- (i) if $\Gamma \subseteq \Delta$, then $\Delta \Vdash M : \sigma$;
- (ii) the free variables of M occur in Γ ;
- (iii) $\Delta \Vdash M : \sigma$ for some $\Delta \subseteq \Gamma$ containing only the free variables of M in its domain.

Lemma (substitution lemma). The typability relation respects substitution.

Lemma (subject reduction). If $\Gamma \Vdash M : \sigma$ and $M \rightarrow_\beta N$, then $\Gamma \Vdash N : \sigma$.

The following theorem establishes the *confluence* property of λ -terms.

Theorem (Church–Rosser theorem for $\lambda(\rightarrow)$). Suppose that $\Gamma \Vdash M : \sigma$. If $M \twoheadrightarrow_{\beta} N_1$ and $M \twoheadrightarrow_{\beta} N_2$, then there exists P such that $N_1 \twoheadrightarrow_{\beta} L$ and $N_2 \twoheadrightarrow_{\beta} L$, and $\Gamma \Vdash L : \sigma$.



Corollary. If a simply typed λ -term admits a β -normal form, then this β -normal form is unique.

Proposition (uniqueness of types). (i) Suppose $\Gamma \Vdash M : \sigma$ and $\Gamma \Vdash M : \tau$. Then $\sigma = \tau$.
(ii) Suppose $\Gamma \Vdash M : \sigma$ and $\Gamma \Vdash N : \tau$, and that $M \equiv_\beta N$. Then $\sigma = \tau$.

Proof. The first part is by induction on M . For the second part, by the Church–Rosser theorem there is a term L to which M and N both eventually reduce, so the result holds by subject reduction. \square


Example. There is no way to assign a type to the expression $\lambda x. x\ x$. Indeed, if x has type τ , then it must also have type $\tau \rightarrow \sigma$ for some σ , but this contradicts uniqueness of types.

6.5. The normalisation theorems

We will measure the complexity of a type by looking at it as a binary tree. For example, for

$$\rho = \mu \rightarrow [((\varphi \rightarrow \psi) \rightarrow \chi) \rightarrow ((\varphi \rightarrow \chi) \rightarrow (\varphi \rightarrow \psi))]$$

the corresponding binary tree is



The height of this tree is the complexity of the type, which in this case is 4. For convenience, we will annotate types of terms with superscripts.

IV. Model Theory and Non-Classical Logic

Definition. The *height* function is the map $h : \Pi \rightarrow \mathbb{N}$ that maps a type variable to 0, and maps a function type $\sigma \rightarrow \tau$ to $1 + \max(h(\sigma), h(\tau))$. We extend the height function to β -redexes: if $(\lambda x : \sigma. P^\tau)^{\sigma \rightarrow \tau} R^\sigma$ is a redex, its height is $h(\sigma \rightarrow \tau)$.

Theorem (weak normalisation theorem). Suppose $\Gamma \Vdash M : \sigma$. Then there is a finite reduction path

$$M = M_0 \rightarrow_\beta M_1 \rightarrow_\beta \dots \rightarrow_\beta M_n$$

where M_n is in β -normal form.

Proof (taming the hydra). First, we define the function $m : \Lambda_\Pi \rightarrow \mathbb{N} \times \mathbb{N}$ by $m(M) = (0, 0)$ if M is in β -normal form, and otherwise, $m(M)$ is the pair $(h(M), \text{redex}(M))$ where $h(M)$ is the maximal height of redexes in M and $\text{redex}(M)$ is the number of redexes in M . We will use induction on the well-founded relation given by the lexicographic order on $\mathbb{N} \times \mathbb{N}$ to show that if M is typeable, it can be reduced to β -normal form.

If $\Gamma \Vdash M : \sigma$ and M is in β -normal form, then the claim is trivial. Otherwise, let Δ be the rightmost redex of maximal height $h = h(M)$. By reducing Δ , we may introduce copies of existing redexes, or create new redexes. Creation of new redexes can occur in one of the following ways.

(i) Suppose Δ is of the form

$$(\lambda x : (\rho \rightarrow \mu). \dots x P^\rho \dots)(\lambda y : \rho. Q^\mu)^{\rho \rightarrow \mu}$$

Then it reduces to

$$\dots (\lambda y : \rho. Q^\mu)^{\rho \rightarrow \mu} P^\rho \dots$$

which is a new redex of height $h(\rho \rightarrow \mu) < h$.

(ii) Suppose Δ is of the form

$$(\lambda x : \tau. \lambda y : \rho. R^\mu) P^\tau$$

occurring in the position $\Delta^{\rho \rightarrow \tau} Q^\rho$. Suppose that Δ reduces to $\lambda y : \rho. R_1^\mu$. Then we have created a new redex $(\lambda y : \rho. R_1^\mu) Q^\rho$ of height $h(\rho \rightarrow \mu) < h(\tau \rightarrow \rho \rightarrow \mu) = h$.

(iii) Suppose Δ is of the form

$$(\lambda x : (\rho \rightarrow \mu). x)(\lambda y : \rho. P^\mu)$$

occurring in the position $\Delta^{\rho \rightarrow \mu} Q^\rho$. Then this reduces to $(\lambda y : \rho. P^\mu) Q^\rho$ of height $h(\rho \rightarrow \mu) < h$.

There is still the possibility that reduction of Δ introduces copies of existing redexes. Suppose Δ is of the form

$$(\lambda x : \rho. P^\rho) Q^\tau$$

and P has more than one free occurrence of x . Then the reduction of Δ will copy all redexes in Q . But as Δ was chosen to be rightmost with maximal height, the height of all redexes in Q have height less than h .

6. Intuitionistic logic and lambda calculi

So if $M \rightarrow_\beta M'$ by reducing Δ , it is always the case that $m(M') < m(M)$ in the lexicographic order. By the inductive hypothesis, M' can be reduced to β -normal form, so the result also holds for M . \square

Theorem (strong normalisation theorem). Let $\Gamma \Vdash M : \sigma$. Then there is no infinite sequence

$$M \rightarrow_\beta M_1 \rightarrow_\beta M_2 \rightarrow_\beta \cdots$$

The proof is omitted.

7. Intuitionistic semantics

7.1. Propositions as types

We will work with the fragment of IPC, denoted $\text{IPC}(\rightarrow)$, where the only connective is \rightarrow , and the deduction rules are \rightarrow -I, \rightarrow -E, AX.

If \mathcal{L} is a propositional language for $\text{IPC}(\rightarrow)$ and P is its set of primitive propositions, we can generate a simply typed λ -calculus $\lambda(\rightarrow)$ by taking the set of primitive types \mathcal{U} to be P . Then the types Π and the propositions \mathcal{L} are generated by the same grammar

$$\mathcal{U} \mid \Pi \rightarrow \Pi$$

A proposition is thus the type of its proofs, and a context is a set of hypotheses.

Proposition (Curry–Howard correspondence for $\text{IPC}(\rightarrow)$). Let Γ be a context for $\lambda(\rightarrow)$, and let φ be a proposition. Then

- (i) If $\Gamma \Vdash M : \varphi$, then

$$|\Gamma| = \{\tau \in \Pi \mid \exists x. (x : \tau) \in \Gamma\} \vdash_{\text{IPC}(\rightarrow)} \varphi$$

- (ii) If $\Gamma \vdash_{\text{IPC}(\rightarrow)} \varphi$, then there is a simply typed λ -term M such that

$$\{(x_\tau : \tau) \mid \tau \in \Gamma\} \Vdash M : \varphi$$

Proof. Part (i). We use induction over the derivation of $\Gamma \Vdash M : \varphi$. If x is a variable not occurring in Γ' , and the derivation is of the form $\Gamma', x : \varphi \Vdash x : \varphi$, then we must prove that $|\Gamma', x : \varphi| \vdash \varphi$, and this holds as $\varphi \vdash \varphi$.

If the derivation has M of the form $\lambda x : \sigma. N$ and $\varphi = \sigma \rightarrow \tau$, then we must have that $\Gamma, x : \sigma \Vdash N : \tau$. By the inductive hypothesis, we have $|\Gamma, x : \sigma| \vdash \tau$, so $|\Gamma|, \sigma \vdash \tau$. Thus we obtain a proof of $\sigma \rightarrow \tau$ from $|\Gamma|$ by \rightarrow -I.

If the derivation is of the form $\Gamma \Vdash (PQ) : \varphi$, then we must have $\Gamma \Vdash P : \sigma \rightarrow \varphi$ and $\Gamma \Vdash Q : \sigma$ for some σ . By the inductive hypothesis, $|\Gamma| \vdash \sigma \rightarrow \varphi$ and $|\Gamma| \vdash \sigma$. Then the result holds by \rightarrow -E.

Part (ii). We use induction over the proof tree of $\Gamma \vdash_{\text{IPC}(\rightarrow)} \varphi$. We write

$$\Delta = \{(x_\tau : \tau) \mid \tau \in \Gamma\}$$

Suppose that we are at a stage of the proof that uses AX, so $\Gamma, \varphi \vdash \varphi$. If $\varphi \in \Gamma$, then clearly $\Delta \Vdash x_\varphi : \varphi$. Otherwise, $\Delta, x_\varphi : \varphi \Vdash x_\varphi : \varphi$ as required.

Suppose that we are at a stage of the proof that uses \rightarrow -E, so

$$\frac{\Gamma \vdash \varphi \rightarrow \psi \quad \Gamma \vdash \varphi}{\Gamma \vdash \psi}$$

By the inductive hypothesis, there are λ -terms M, N such that $\Delta \Vdash M : \varphi \rightarrow \psi$ and $\Delta \Vdash N : \varphi$. Then $\Delta \Vdash (MN) : \psi$ as required.

Finally, suppose we are at a stage of the proof that uses \rightarrow -I, so

$$\frac{\Gamma, \varphi \vdash \psi}{\Gamma \vdash \varphi \rightarrow \psi}$$

If $\varphi \in \Gamma$, then by the inductive hypothesis, there is a λ -term M such that $\Delta \Vdash M : \psi$. By the weakening rule, $\Delta, x : \varphi \Vdash M : \psi$ where x is a variable that does not occur in Δ . Then $\Delta \Vdash (\lambda x : \varphi. M) : \varphi \rightarrow \psi$ as required. Now suppose $\varphi \notin \Gamma$. By the inductive hypothesis we obtain a λ -term M such that $\Delta, x_\varphi : \varphi \Vdash M : \psi$. Then similarly $\Delta \Vdash (\lambda x_\varphi : \varphi. M) : \varphi \rightarrow \psi$. \square

This justifies the Brouwer–Heyting–Kolmogorov interpretation of intuitionistic logic.

Example. Let φ, ψ be primitive propositions, and consider the λ -term

$$\lambda f : (\varphi \rightarrow \psi) \rightarrow \varphi. \lambda g : \varphi \rightarrow \psi. g(fg)$$

This term has type

$$((\varphi \rightarrow \psi) \rightarrow \varphi) \rightarrow ((\varphi \rightarrow \psi) \rightarrow \psi)$$

The term encodes a proof of this proposition in $\vdash_{\text{IPC}(\rightarrow)}$. The corresponding proof tree is

$$\frac{\frac{\frac{g : [\varphi \rightarrow \psi] \quad f : [(\varphi \rightarrow \psi) \rightarrow \varphi]}{fg : \varphi} \rightarrow\text{-E} \quad \frac{g : [\varphi \rightarrow \psi]}{g(fg) : \psi} \rightarrow\text{-E}}{g(fg) : \psi} \rightarrow\text{-I} \quad \frac{\lambda g : \varphi \rightarrow \psi. g(fg) : (\varphi \rightarrow \psi) \rightarrow \psi}{\lambda f : (\varphi \rightarrow \psi) \rightarrow \varphi. \lambda g : \varphi \rightarrow \psi. g(fg) : ((\varphi \rightarrow \psi) \rightarrow \varphi) \rightarrow ((\varphi \rightarrow \psi) \rightarrow \psi)} \rightarrow\text{-I}$$

7.2. Full simply typed lambda calculus

The types of the full simply typed λ -calculus $\text{ST}\lambda\text{C}$ are generated by the following grammar.

$$\Pi ::= \mathcal{U} \mid \Pi \rightarrow \Pi \mid \Pi \times \Pi \mid \Pi + \Pi \mid 1 \mid 0$$

where \mathcal{U} is a set of primitive types or type variables. The terms are of the form

$$\begin{aligned} \Lambda_\Pi ::= & V \mid (\lambda x : \Pi. \Lambda_\Pi) \mid \Lambda_\Pi \Lambda_\Pi \mid \\ & \langle \Lambda_\Pi, \Lambda_\Pi \rangle \mid \pi_1(\Lambda_\Pi) \mid \pi_2(\Lambda_\Pi) \mid \\ & \iota_1(\Lambda_\Pi) \mid \iota_2(\Lambda_\Pi) \mid \text{case}(\Lambda_\Pi; V. \Lambda_\Pi; V. \Lambda_\Pi) \mid \\ & * \mid !_\Pi \Lambda_\Pi \end{aligned}$$

IV. Model Theory and Non-Classical Logic

where V is an infinite set of variables, and \star is a constant. This expanded syntax comes with new typing rules.

$$\begin{array}{c}
\frac{\Gamma \Vdash M : \psi \times \varphi}{\Gamma \Vdash \pi_1(M) : \psi} \quad \frac{\Gamma \Vdash M : \psi \times \varphi}{\Gamma \Vdash \pi_2(M) : \varphi} \quad \frac{\Gamma \Vdash M : \psi \quad \Gamma \Vdash N : \varphi}{\Gamma \Vdash \langle M, N \rangle : \psi \times \varphi} \quad \frac{\Gamma \Vdash M : \psi}{\Gamma \Vdash \iota_1(M) : \psi + \varphi} \\
\\
\frac{\Gamma \Vdash M : \varphi}{\Gamma \Vdash \iota_2(M) : \psi + \varphi} \quad \frac{\Gamma \Vdash L : \psi + \varphi \quad \Gamma, x : \psi \Vdash M : \rho \quad \Gamma, y : \varphi \Vdash N : \rho}{\Gamma \Vdash \text{case}(L; x^\psi.M; y^\varphi.N) : \rho} \\
\\
\frac{}{\Gamma \Vdash \star : 1} \quad \frac{\Gamma \Vdash M : 0}{\Gamma \Vdash !_\varphi M : \varphi}
\end{array}$$

This typing relation captures the Brouwer–Heyting–Kolmogorov interpretation when paired with new reduction rules.

$$\begin{array}{c}
\pi_1(\langle M, N \rangle) \rightarrow_\beta M \quad \pi_2(\langle M, N \rangle) \rightarrow_\beta N \quad \langle \pi_1(M), \pi_2(M) \rangle \rightarrow_\eta M \\
\\
\text{case}(\iota_1(M); x^\psi.K; y^\varphi.L) \rightarrow_\beta K[x := M] \quad \text{case}(\iota_2(M); x^\psi.K; y^\varphi.L) \rightarrow_\beta L[y := M] \\
\\
\text{if } \Gamma \Vdash M : 1 \text{ then } M \rightarrow_\eta \star
\end{array}$$

We can expand propositions-as-types to our new types:

- (i) 0 corresponds to \perp ;
- (ii) 1 corresponds to \top ;
- (iii) product types correspond to conjunctions;
- (iv) coproduct types correspond to disjunctions.

In this way, propositions correspond to types. Redexes are now those expressions consisting of a constructor (pair formation, λ -abstraction, and injections) followed by the corresponding destructor (projections, applications, and case expressions).

Example. Consider the following proof of $(\varphi \wedge \chi) \rightarrow (\psi \rightarrow \varphi)$.

$$\frac{\frac{\frac{[\varphi \wedge \chi]}{\varphi}}{\psi \rightarrow \varphi}}{(\varphi \wedge \chi) \rightarrow (\psi \rightarrow \varphi)} \quad [\psi]$$

Annotating the corresponding λ -terms, we obtain

$$\frac{\frac{\frac{p : [\varphi \wedge \chi]}{\pi_1(p) : \varphi}}{\lambda b^\psi. \pi_1(p) : \psi \rightarrow \varphi}}{\lambda p^{\varphi \times \chi}. \lambda b^\psi. \pi_1(p) : (\varphi \wedge \chi) \rightarrow (\psi \rightarrow \varphi)} \quad b : [\psi]$$

Hence this proof tree corresponds to the λ -term

$$\lambda p^{\varphi \times \chi}. \lambda b^\psi. \pi_1(p) : (\varphi \times \chi) \rightarrow (\psi \rightarrow \varphi)$$

In summary, the Curry–Howard correspondence for the whole of IPC and ST λ C states that

- (i) (primitive) types correspond to (primitive) propositions;
- (ii) variables correspond to hypotheses;
- (iii) λ -terms correspond to proofs;
- (iv) inhabitation of a type corresponds to provability of a proposition;
- (v) term reduction corresponds to proof normalisation.

7.3. Heyting semantics

Boolean algebras represent truth-values of classical propositions. We can generalise this notion to intuitionistic logic.

Definition. A *Heyting algebra* H is a bounded lattice equipped with a binary operation $\Rightarrow : H \times H \rightarrow H$ such that

$$a \wedge b \leq c \iff a \leq b \Rightarrow c$$

A *morphism* of Heyting algebras is a function that preserves all finite meets and joins (including true and false) and \Rightarrow .

In particular, if f is a morphism of Heyting algebras and $a \leq b$, then $f(a) \leq f(b)$.

Example. (i) Every Boolean algebra is a Heyting algebra by defining $a \Rightarrow b$ to be $\neg a \vee b$. Note that $\neg a = a \Rightarrow \perp$.

(ii) Every topology is a Heyting algebra, where $U \Rightarrow V = ((X \setminus U) \cup V)^\circ$.

(iii) Every finite distributive lattice is a Heyting algebra.

(iv) The Lindenbaum–Tarski algebra of a propositional theory \mathcal{T} with respect to IPC is a Heyting algebra.

Definition. Let H be a Heyting algebra and let \mathcal{L} be a propositional language with a set P of primitive propositions. An H -valuation is a function $v : P \rightarrow H$, recursively expanded to \mathcal{L} by the rules

- (i) $v(\perp) = \perp$;
- (ii) $v(A \wedge B) = v(A) \wedge v(B)$;
- (iii) $v(A \vee B) = v(A) \vee v(B)$;
- (iv) $v(A \rightarrow B) = v(A) \Rightarrow v(B)$.

IV. Model Theory and Non-Classical Logic

We say that a proposition A is H -valid if $v(A) = \top$ for all valuations v . A is an H -consequence of a finite set of propositions Γ if $v(\bigwedge \Gamma) \leq v(A)$, and write $\Gamma \models_H A$.

Lemma (soundness). Let H be a Heyting algebra and let $v : \mathcal{L} \rightarrow H$ be an H -valuation. If $\Gamma \vdash_{\text{IPC}} A$, then $\Gamma \models_H A$.

Proof. We proceed by induction over the derivation of $\Gamma \vdash_{\text{IPC}} A$.

$$(i) \text{ (AX) } v((\bigwedge \Gamma) \wedge A) = v(\bigwedge \Gamma) \wedge v(A) \leq v(A).$$

(ii) (\wedge -I) In this case, $A = B \wedge C$ and we have derivations $\Gamma_1 \vdash B$, $\Gamma_2 \vdash C$ with $\Gamma_1, \Gamma_2 \subseteq \Gamma$. By the inductive hypothesis, $v(\Gamma_1) \leq v(B)$ and $v(\Gamma_2) \leq v(C)$, hence

$$v(\bigwedge \Gamma) \leq v(\Gamma_1) \wedge v(\Gamma_2) \leq v(B) \wedge v(C) = v(B \wedge C) = v(A)$$

(iii) (\rightarrow -I) In this case, $A = B \rightarrow C$ and we have $\Gamma \cup \{B\} \vdash C$. By the inductive hypothesis, $v(\bigwedge \Gamma) \wedge v(B) \leq v(C)$. But then $v(\bigwedge \Gamma) \leq v(B) \Rightarrow v(C)$ by definition, so $v(\bigwedge \Gamma) \leq v(B \rightarrow C)$ as required.

(iv) (\vee -I) In this case, $A = B \vee C$, and without loss of generality, we have $\Gamma \vdash B$. By the inductive hypothesis, $v(\bigwedge \Gamma) \leq v(B)$, but $v(B) \leq v(B) \vee v(C) = v(B \vee C)$ as required.

(v) (\wedge -E) By the inductive hypothesis, we have $v(\bigwedge \Gamma) \leq v(A \wedge B) = v(A) \wedge v(B) \leq v(A), v(B)$ as required.

(vi) (\rightarrow -E) We know that $v(A \rightarrow B) = (v(A) \Rightarrow v(B))$. From the inequality $v(A \rightarrow B) \leq (v(A) \Rightarrow v(B))$, we deduce $v(A \rightarrow B) \wedge v(A) \leq v(B)$. Thus, if $v(\bigwedge \Gamma) \leq v(A \rightarrow B)$ and $v(\bigwedge \Gamma) \leq v(A)$, we have $v(\bigwedge \Gamma) \leq v(B)$ as required.

(vii) (\vee -E) By the inductive hypothesis,

$$v(A \wedge \bigwedge \Gamma) \leq v(C); \quad v(B \wedge \bigwedge \Gamma) \leq v(C); \quad v(\bigwedge \Gamma) \leq v(A \vee B) = v(A) \vee v(B)$$

Hence,

$$v(\bigwedge \Gamma) = v(\bigwedge \Gamma) \wedge (v(A) \vee v(B)) = (v(\bigwedge \Gamma) \wedge v(A)) \vee (v(\bigwedge \Gamma) \wedge v(B)) \leq v(C) \vee v(C) = v(C)$$

as every Heyting algebra is a distributive lattice.

(viii) (\perp -E) If $v(\bigwedge \Gamma) \leq v(\perp) = \perp$, then $v(\bigwedge \Gamma) = \perp$. Hence, $v(\bigwedge \Gamma) \leq v(A)$ for any A .

□

Example. The law of the excluded middle LEM is not provable in IPC. Let p be a primitive proposition, and consider the Heyting algebra given by the Sierpiński topology $\{\emptyset, \{1\}, \{1, 2\}\}$ on $X = \{1, 2\}$. We define the valuation given by $v(p) = \{1\}$. Then

$$v(\neg p) = \{1\} \Rightarrow \emptyset = (\{1, 2\} \setminus \{1\})^\circ = \{2\}^\circ = \emptyset$$

Hence,

$$v(p \vee \neg p) = \{1\} \cup \emptyset = \{1\} \neq \{1, 2\} = \top$$

Thus, by soundness, $p \vee \neg p$ is not provable (from the empty context, which has valuation $\top = \{1, 2\}$) in IPC.

Example. Peirce's law $((p \rightarrow q) \rightarrow p) \rightarrow p$ is not intuitionistically valid. Let H be the Heyting algebra given by the usual topology on the plane \mathbb{R}^2 , and let

$$v(p) = \mathbb{R}^2 \setminus \{(0, 0)\}; \quad v(q) = \emptyset$$

Classical completeness can be phrased as

$$\Gamma \vdash_{\text{CPC}} A \iff \Gamma \models_2 A$$

where 2 is the Boolean algebra $\{0, 1\}$. For intuitionistic logic, we cannot replace 2 with a single finite Heyting algebra, so we will instead quantify over all Heyting algebras.

Theorem (completeness). A proposition is provable in IPC if and only if it is H -valid for every Heyting algebra H .

Proof. For the forward direction, if $\vdash_{\text{IPC}} A$, then $\top \leq v(A)$ for every Heyting algebra H and valuation v , by soundness. Then $\top = v(A)$, so A is H -valid.

For the backward direction, suppose A is H -valid for every Heyting algebra H . Note that the Lindenbaum–Tarski algebra \mathcal{L}/\sim for the empty theory, with respect to IPC, is a Heyting algebra. Consider the valuation given by mapping each primitive proposition to its equivalence class in \mathcal{L}/\sim . Then, one can easily show by induction that $v : \mathcal{L} \rightarrow \mathcal{L}/\sim$ is the quotient map by considering the construction of the Lindenbaum–Tarski algebra. Now, A is valid in every Heyting algebra and with respect to every valuation, so in particular, $v(A) = \top$ in \mathcal{L}/\sim . But then $v(A) \in [\top]$, so $\vdash_{\text{IPC}} A \leftrightarrow \top$, so $\vdash_{\text{IPC}} A$ as required. \square

7.4. Kripke semantics

Definition. Let S be a poset. For each $a \in S$, we define its *principal up-set* to be

$$a \uparrow = \{s \in S \mid a \leq s\}$$

Note that $U \subseteq S$ is a terminal segment if and only if it contains $a \uparrow$ for each $a \in U$.

Proposition. Let S be a poset. Then the set $T(S)$ of terminal segments of S has the structure of a Heyting algebra.

Proof. The order is given by inclusion: $U \leq V$ if and only if $U \subseteq V$. We define

$$\begin{aligned} U \wedge V &= U \cap V \\ U \vee V &= U \cup V \\ U \Rightarrow V &= \{s \mid s \uparrow \cap U \subseteq V\} \end{aligned}$$

IV. Model Theory and Non-Classical Logic

One can check that this forms a Heyting algebra as required. \square

Definition. Let P be a set of primitive propositions. A *Kripke model* is a triple (S, \leq, \Vdash) where S is a poset and $(\Vdash) \subseteq S \times P$ is a relation satisfying the *persistence property*: if $p \in P$ is such that $s \Vdash p$ and $s \leq s'$, then $s' \Vdash p$.

S is a set of possible *worlds*, or states of knowledge, ordered by how knowledgeable they are. The relation \Vdash is called the *forcing* relation; we say that a world *forces* a proposition to be true.

Every valuation v on $T(S)$ induces a Kripke model by setting $s \Vdash p \iff s \in v(p)$. The persistence property corresponds to the fact that $T(S)$ contains only terminal segments.

Definition. Let (S, \leq, \Vdash) be a Kripke model. We can extend the forcing relation to a relation $(\Vdash) \subseteq S \times \mathcal{L}$ recursively as follows.

- (i) $s \not\Vdash \perp$;
- (ii) $s \Vdash \varphi \wedge \psi$ if and only if $s \Vdash \varphi$ and $s \Vdash \psi$;
- (iii) $s \Vdash \varphi \vee \psi$ if and only if $s \Vdash \varphi$ or $s \Vdash \psi$;
- (iv) $s \Vdash \varphi \rightarrow \psi$ if and only if for all $s' \geq s$, $s' \Vdash \varphi$ implies $s' \Vdash \psi$.


One can check by induction that persistence holds for arbitrary propositions.

Remark. $s \Vdash \neg\varphi$ if and only if no more knowledgeable world than s forces φ . $s \Vdash \neg\neg\varphi$ is the statement that φ is consistent with every extension of s but need not hold in s itself; that is, for each $s' \geq s$, there exists $s'' \geq s'$ with $s'' \Vdash \varphi$.

We say that $S \Vdash \varphi$ if every world s forces φ . If S has a bottom element s , then $S \Vdash \varphi$ if and only if $s \Vdash \varphi$ by persistence.


Example. Consider the Kripke models

(i)



where $s' \Vdash p$;

(ii)



where $s'' \Vdash p$;

(iii)

$$\begin{array}{c} s' \\ | \\ s \end{array}$$

where $s' \Vdash p$ and $s' \Vdash q$.

Note that in (i), we have $s \nVdash \neg p$, since $s' \geq s$ and $s' \Vdash p$. But also $s \nVdash p$ by assumption, thus $s \nVdash p \vee \neg p$. Note that $s \Vdash \neg\neg p$, but $s \nVdash p$, so we also have $s \nVdash \neg\neg p \rightarrow p$.

In (ii), $s \nVdash \neg\neg p$, since $s' \geq s$ cannot access a world that forces p . We also have $s \nVdash \neg p$, since $s'' \geq s'$ and $s'' \Vdash p$. Thus $s \nVdash \neg\neg p \vee \neg p$.

In (iii), $s \nVdash (p \rightarrow q) \rightarrow (\neg p \vee q)$. Indeed, all worlds force $p \rightarrow q$, and we have $s \nVdash q$, so it suffices to check that $s \nVdash \neg p$, but this holds as $s' \geq s$ and $s' \Vdash p$.

A proper filter \mathcal{F} is called *prime* if whenever $x \vee y \in \mathcal{F}$, either $x \in \mathcal{F}$ or $y \in \mathcal{F}$.

Lemma. Let H be a Heyting algebra and let v be an H -valuation. Then there is a Kripke model (S, \leq, \Vdash) such that for each proposition φ , we have $v \models_H \varphi$ if and only if $S \Vdash \varphi$.

Thus we can convert between Kripke models and valuations on Heyting algebras. This will allow us to prove the completeness theorem for Kripke semantics.

Proof. Let S be the set of prime filters on H ordered by inclusion. We say that $\mathcal{F} \Vdash p$ if and only if $v(p) \in \mathcal{F}$, and prove by induction that this extends to arbitrary propositions. Here, we will prove the case of implications; the other connectives are easy, and primality of the filter is required for the case of disjunction. Let $\mathcal{F} \Vdash (\psi \rightarrow \psi')$ and suppose $v(\psi \rightarrow \psi') = v(\psi) \Rightarrow v(\psi') \notin \mathcal{F}$. Let \mathcal{G}' be the smallest filter containing \mathcal{F} and $v(\psi)$. Then

$$\mathcal{G}' = \{b \mid \exists f \in \mathcal{F}. f \wedge v(\psi) \leq b\}$$

Note that $v(\psi') \notin \mathcal{G}'$, otherwise $f \wedge v(\psi) \leq v(\psi')$ for some $f \in \mathcal{F}$, and then $f \leq v(\psi) \Rightarrow v(\psi') \in \mathcal{F}$, giving a contradiction. In particular, \mathcal{G}' is a proper filter, so by Zorn's lemma there is a prime filter \mathcal{G} containing \mathcal{G}' that does not contain $v(\psi')$.

By the inductive hypothesis, $\mathcal{G} \Vdash \psi$, and since $\mathcal{F} \Vdash (\psi \rightarrow \psi')$ and \mathcal{G}' contains \mathcal{G} which contains \mathcal{F} , we must have $\mathcal{G} \Vdash \psi'$. Then $v(\psi') \in \mathcal{G}$, which is a contradiction. Thus $\mathcal{F} \Vdash \psi \rightarrow \psi'$ implies that $v(\psi \rightarrow \psi') \in \mathcal{F}$.

Conversely, suppose

$$v(\psi \rightarrow \psi') \in \mathcal{F} \subseteq \mathcal{G} \Vdash \psi$$

By the inductive hypothesis, $v(\psi) \in \mathcal{G}$, and so $v(\psi) \Rightarrow v(\psi') \in \mathcal{G}$ as $\mathcal{F} \subseteq \mathcal{G}$. Then $v(\psi') \geq v(\psi) \wedge (v(\psi) \Rightarrow v(\psi')) \in \mathcal{G}$, so again by the inductive hypothesis, $\mathcal{G} \Vdash \psi'$ as required.

It thus suffices to show that $v \models_H \varphi$ if and only if $S \Vdash \varphi$. If $v \models_H \varphi$, then $v(\varphi) = \top$, so $v(\varphi)$ is contained in every filter of H . So $\mathcal{F} \Vdash \varphi$ for every prime filter \mathcal{F} . Conversely, suppose $S \Vdash \varphi$ but $v \not\models_H \varphi$. Then since $v(\varphi) \neq \top$, there must be a proper filter \mathcal{F} that does not contain

IV. Model Theory and Non-Classical Logic

$v(\varphi)$. We extend this as above to a prime filter \mathcal{G} that does not contain $v(\varphi)$. Then $\mathcal{G} \not\vdash \varphi$, contradicting the assumption that $S \Vdash \varphi$. \square

Theorem (completeness). For every proposition φ , we have $\Gamma \vdash_{\text{IPC}} \varphi$ if and only if for all Kripke models (S, \leq, \Vdash) , if $S \Vdash \Gamma$ then $S \Vdash \varphi$.

Proof. Soundness holds by induction. For adequacy, suppose $\Gamma \not\vdash_{\text{IPC}} \varphi$. Then by completeness of Heyting semantics, there is a Heyting algebra H and H -valuation v such that $v \models_H \Gamma$ but $v \not\models_H \varphi$. By the previous lemma, there is a Kripke model (S, \leq, \Vdash) such that $S \Vdash \Gamma$ but $S \not\vdash \varphi$, contradicting the hypothesis. \square

V. Group Cohomology

Lectured in Lent 2024 by DR. C. J. B. BROOKES

(Course description goes here.)

Contents

1.	Definitions and resolutions	275
1.1.	???	275
1.2.	???	278
1.3.	Cohomology	280
1.4.	Independence of cohomology groups	282
2.	Low degree cohomology and group extensions	285
2.1.	Degree 1	285
2.2.	Degree 2	286
2.3.	Central extensions	289
2.4.	Generators and relations	291
2.5.	Homology groups	292

1. Definitions and resolutions

1.1. ???

Let G be a group.

Definition. The *integral group ring* $\mathbb{Z}G$ is the set of formal sums $\sum n_g g$, where $n_g \in \mathbb{Z}$, $g \in G$, and only finitely many of the n_g are nonzero. An addition operation makes this set a free abelian group:

$$\left(\sum m_g g\right) + \left(\sum n_g g\right) = \sum (m_g + n_g)g$$

Multiplication is defined by

$$\left(\sum_{h \in G} m_h h\right) \left(\sum_{k \in G} n_k k\right) = \sum \left(\sum_{hk=g} m_h n_k\right) g$$

The multiplicative identity is $1e$ where e is the identity of G . This produces an associative ring, which underlies the integral representation theory of G .

Definition. A (left) $\mathbb{Z}G$ -module M is an abelian group under addition together with a map $\mathbb{Z}G \times M \rightarrow M$ denoted $(r, m) \mapsto rm$, satisfying

- (i) $r(m_1 + m_2) = rm_1 + rm_2$;
- (ii) $(r_1 + r_2)m = r_1m + r_2m$;
- (iii) $r_1(r_2m) = (r_1r_2)m$;
- (iv) $1m = m$.

A module is *trivial* if $gm = m$ for all $g \in G$ and $m \in M$. We call \mathbb{Z} *the* trivial module, given by the trivial action $gn = n$ for all $n \in \mathbb{Z}$ and $g \in G$.

The *free* $\mathbb{Z}G$ -module on a set X is the module of formal sums $\sum r_x x$ where $r_x \in \mathbb{Z}G$ and $x \in X$, and only finitely many of the r_x are nonzero. This has the obvious G -action. This module will be denoted $\mathbb{Z}G\{X\}$.

We can define submodules, quotient modules, and so on as one would expect.

Definition. A (left) $\mathbb{Z}G$ -map or *morphism* $\alpha : M_1 \rightarrow M_2$ is a map of abelian groups with $\alpha(rm) = r\alpha(m)$ for all $r \in \mathbb{Z}G$ and $m \in M_1$.

Example. The *augmentation map* $\varepsilon : \mathbb{Z}G \rightarrow \mathbb{Z}$ is the $\mathbb{Z}G$ -map between left $\mathbb{Z}G$ -modules given by

$$\sum n_g g \mapsto \sum n_g$$

This is also a right $\mathbb{Z}G$ -map, and also a map of rings.

We will write $\text{Hom}_G(M, N)$ to be the set of $\mathbb{Z}G$ -maps $M \rightarrow N$, which is made into an abelian group under pointwise addition.

V. Group Cohomology

Example. Regarding $\mathbb{Z}G$ as a left $\mathbb{Z}G$ -module, then

$$\text{Hom}_G(\mathbb{Z}G, M) \cong M$$

for any left $\mathbb{Z}G$ -module M . This isomorphism is given by $\varphi \mapsto \varphi(1)$; the $\mathbb{Z}G$ -map is determined by the image of 1.

$$\varphi(r) = \varphi(r \cdot 1) = r\varphi(1)$$

Note that $\text{Hom}_G(\mathbb{Z}G, M)$ can be viewed as a left $\mathbb{Z}G$ -module, given by

$$(s\varphi)(r) = \varphi(rs); \quad s \in \mathbb{Z}G$$

Note that the isomorphism

$$\text{Hom}_G(\mathbb{Z}G, \mathbb{Z}G) \cong \mathbb{Z}G; \quad \varphi \mapsto \varphi(1)$$

satisfies $\varphi(r) = r\varphi(1)$ and so φ corresponds to multiplication on the right by $\varphi(1)$.

Remark. G may not be abelian, and so we must carefully distinguish left and right actions.

Definition. If $f : M_1 \rightarrow M_2$ is a $\mathbb{Z}G$ -map, its *dual maps* f^* are $\mathbb{Z}G$ -maps $\text{Hom}_G(M_2, N) \rightarrow \text{Hom}_G(M_1, N)$ for each $\mathbb{Z}G$ -module N , given by composition on the right with f . If $f : N_1 \rightarrow N_2$, its *induced maps* f_* are $\text{Hom}_G(M, N_1) \rightarrow \text{Hom}_G(M, N_2)$ given by composition on the left with f . These are maps of abelian groups.

We will now present a prototypical example.

Example. Let $G = \langle t \rangle$ be an infinite cyclic group. Consider the graph whose vertices are v_i for $i \in \mathbb{Z}$, where v_i is joined to v_{i+1} and v_{i-1} . Let V be its set of vertices, and E be its set of edges. G acts by translations on this graph, where t maps v_i to v_{i+1} . The formal sums $\mathbb{Z}V$ and $\mathbb{Z}E$ can be regarded as $\mathbb{Z}G$ -modules. They are free: $\mathbb{Z}V = \mathbb{Z}G\{v_0\}$, and $\mathbb{Z}E = \mathbb{Z}G\{e\}$ where e is the edge connecting v_0 and v_1 . The boundary map is a $\mathbb{Z}G$ -map $d : \mathbb{Z}E \rightarrow \mathbb{Z}V$ given by $e \mapsto v_1 - v_0$. There is also a $\mathbb{Z}G$ -map $\mathbb{Z}V \rightarrow \mathbb{Z}$ given by $v_0 \mapsto 1$; this corresponds to the augmentation map.

Definition. A *chain complex* of $\mathbb{Z}G$ -modules is a sequence

$$M_s \xrightarrow{d_s} M_{s-1} \xrightarrow{d_{s-1}} M_{s-2} \longrightarrow \cdots \xrightarrow{d_{t+1}} M_t$$

such that for every $t < n < s$, we have $d_n d_{n+1} = 0$, and so $\text{im } d_{n+1} \subseteq \ker d_n$. We will refer to the entire sequence as $M_\bullet = (M_n, d_n)_{t \leq n \leq s}$.

We say that M_\bullet is *exact* at M_n if $\text{im } d_{n+1} = \ker d_n$, and we say it is *exact* if it is exact at all M_n for $t < n < s$. The *homology* of this chain complex is

$$H_s(M_\bullet) = \ker d_s; \quad H_n(M_\bullet) = \ker d_n / \text{im } d_{n+1}; \quad H_t(M_\bullet) = \text{coker } d_{t-1} = M_t / \text{im } d_{t+1}$$

A *short exact sequence* is an exact chain complex of the form

$$0 \longrightarrow M_1 \xrightarrow{\alpha} M_2 \xrightarrow{\beta} M_3 \longrightarrow 0$$

That is, α is injective, β is surjective, and $\text{im } \alpha = \ker \beta$.

Example. In our example above, we have the short exact sequence

$$0 \longrightarrow \mathbb{Z}E \longrightarrow \mathbb{Z}V \longrightarrow \mathbb{Z} \longrightarrow 0$$

This corresponds to a short exact sequence

$$0 \longrightarrow \mathbb{Z}G \longrightarrow \mathbb{Z}G \longrightarrow \mathbb{Z} \longrightarrow 0$$

where $G = \langle t \rangle$ is an infinite cyclic group, and the map $\mathbb{Z}G \rightarrow \mathbb{Z}G$ is given by multiplication on the right by $t - 1$.

Definition. A $\mathbb{Z}G$ -module P is *projective* if, for every surjective $\mathbb{Z}G$ -map $\alpha : M_1 \rightarrow M_2$ and every $\mathbb{Z}G$ -map $\beta : P \rightarrow M_2$, there is a map $\bar{\beta} : P \rightarrow M_1$ such that $\alpha \circ \bar{\beta} = \beta$.

$$\begin{array}{ccc} & P & \\ \swarrow \bar{\beta} & \downarrow \beta & \\ M_1 & \xrightarrow{\alpha} & M_2 \longrightarrow 0 \end{array}$$

Given any short exact sequence

$$0 \longrightarrow N \xrightarrow{f} M_1 \xrightarrow{\alpha} M_2 \longrightarrow 0$$

we can consider

$$0 \longrightarrow \text{Hom}_G(P, N) \xrightarrow{f_*} \text{Hom}_G(P, M_1) \xrightarrow{\alpha_*} \text{Hom}_G(P, M_2) \longrightarrow 0$$

We could have defined projectivity by saying that this new sequence is exact. Note that this sequence is always a chain complex regardless if P is projective, and we always have exactness except possibly at $\text{Hom}_G(P, M_2)$.

Lemma. Free modules are projective.

Proof. Let $\alpha : M_1 \rightarrow M_2$ be a surjective $\mathbb{Z}G$ -map, and let $\beta : \mathbb{Z}G\{X\} \rightarrow M_2$. Then for each generator $x \in X$, there exists some $m_x \in M_1$ such that $\alpha(m_x) = \beta(x)$. We then define $\bar{\beta} : \mathbb{Z}G\{X\} \rightarrow M_1$ by mapping

$$\sum r_x x \mapsto \sum r_x m_x$$

which satisfies the required equation $\alpha \bar{\beta} = \beta$. □

Definition. A *projective (free) resolution* of the trivial module \mathbb{Z} is an exact sequence

$$\cdots \xrightarrow{d_2} P_1 \xrightarrow{d_1} P_0 \xrightarrow{d_0} \mathbb{Z} \longrightarrow 0$$

where the P_i are projective (respectively free). This is a chain complex.

V. Group Cohomology

Example. Let $G = \langle t \rangle$ be an infinite cyclic group. Then we have a finite free resolution of \mathbb{Z} given by the exact sequence

$$0 \longrightarrow \mathbb{Z}G \xrightarrow{\cdot(t-1)} \mathbb{Z}G \xrightarrow{\varepsilon} \mathbb{Z} \longrightarrow 0$$

where ε is the augmentation map.

Example. Let $G = \langle t \rangle$ be a cyclic group of order n . Then we have a resolution

$$\cdots \longrightarrow \mathbb{Z}G \xrightarrow{\beta} \mathbb{Z}G \xrightarrow{\alpha} \mathbb{Z}G \xrightarrow{\beta} \mathbb{Z}G \xrightarrow{\alpha} \mathbb{Z}G \xrightarrow{\varepsilon} \mathbb{Z} \longrightarrow 0$$

where

$$\alpha(x) = x(t-1); \quad \beta(x) = x(1+t+\cdots+t^{n-1})$$

From algebraic topology, if we have a connected simplicial complex X with fundamental group $\pi_1(X) = G$, such that the universal cover \tilde{X} is contractible, we obtain a free resolution of \mathbb{Z} given by the universal cover. In this way, the simplicial complex X contains a lot of information about its fundamental group; this is what we aim to replicate algebraically.

For calculation purposes, we are interested in ‘small’ resolutions, for instance where the free modules have small rank. However, for theory development, we often want general constructions, and resolutions given by generic theory tend to be large.

Definition. G is of type FP_n if \mathbb{Z} has a projective resolution

$$\cdots \xrightarrow{d_2} P_1 \xrightarrow{d_1} P_0 \xrightarrow{d_0} \mathbb{Z} \longrightarrow 0$$

which may be infinite, but where P_n, P_{n-1}, \dots, P_0 are finitely generated as $\mathbb{Z}G$ -modules.

We say G is of type FP_∞ if \mathbb{Z} has a projective resolution where all of the P_i are finitely generated as $\mathbb{Z}G$ -modules. Finally, G is of type FP if \mathbb{Z} has a projective resolution where all of the P_i are finitely generated as $\mathbb{Z}G$ -modules, and the resolution is of finite length, so $P_s = 0$ for sufficiently large s .

Example. (i) Let $G = \langle t \rangle$ be the infinite cyclic group. Then G is of type FP .

(ii) Let $G = \langle t \rangle$ be a finite cyclic group. Then G is of type FP_∞ ; we will show later that it is not of type FP .

These can be regarded as finiteness conditions on the group G . The topological version of FP_n would be that a simplicial complex X with fundamental group G has a finite n -skeleton.

1.2. ???

Consider a partial projective resolution

$$P_s \longrightarrow P_{s-1} \longrightarrow \cdots \longrightarrow P_1 \longrightarrow P_0 \longrightarrow \mathbb{Z} \longrightarrow 0$$

Then we can set P_{s+1} to be the free module $\mathbb{Z}G\{X_{s+1}\}$ where X_{s+1} is the kernel of d_s . We can then set d_{s+1} to be

$$\sum_{x \in P_{s+1}} r_x x \mapsto \sum_{x \in P_s} r_x x$$

where the left-hand side is a formal sum, and the right-hand sum takes place in P_s . We thus obtain a longer partial projective resolution

$$P_{s+1} \xrightarrow{d_{s+1}} P_s \longrightarrow P_{s-1} \longrightarrow \cdots \longrightarrow P_1 \longrightarrow P_0 \longrightarrow \mathbb{Z} \longrightarrow 0$$

since exactness holds at P_s by construction. We could alternatively take X_{s+1} to be a $\mathbb{Z}G$ -generating set of $\ker d_s$; this would have the effect of reducing the size of P_{s+1} , which is most useful in direct calculation if $\ker d_s$ is finitely generated. Continuing in this way, we obtain a resolution of \mathbb{Z} .

Definition. The *standard* or *bar* resolution of \mathbb{Z} is constructed as follows. Let $G^{(n)}$ be the set of formal symbols

$$G^{(n)} = \{[g_1 | \dots | g_n] \mid g_1, \dots, g_n \in G\}$$

where $G^{(0)}$ is the set containing only the empty symbol $[\]$. Let $F_n = \mathbb{Z}G\{G^{(n)}\}$ be the corresponding free modules. We define the boundary maps $d_n : F_n \rightarrow F_{n-1}$ by

$$\begin{aligned} d_n([g_1 | \dots | g_n]) &= g_1[g_2 | \dots | g_n] \\ &\quad - [g_1 g_2 | g_3 | \dots | g_n] \\ &\quad + [g_1 | g_2 g_3 | \dots | g_n] - \dots \\ &\quad + (-1)^{n-1} [g_1 | \dots | g_{n-1} g_n] \\ &\quad + (-1)^n [g_1 | \dots | g_{n-1}] \end{aligned}$$

One can verify explicitly that these are chain maps as required, giving a free resolution

$$\cdots \longrightarrow F_1 \longrightarrow F_0 \longrightarrow \mathbb{Z}$$

Remark. The bar resolution corresponds to the standard resolution in algebraic topology. Consider the free abelian group $\mathbb{Z}G^{n+1}$ generated by the $(n+1)$ -tuples with elements in G . Then G acts on G^{n+1} diagonally:

$$g(g_0, \dots, g_n) = (gg_0, \dots, gg_n)$$

Thus $\mathbb{Z}G^{n+1}$ is a free $\mathbb{Z}G$ -module on the basis of $(n+1)$ -tuples with first element 1. The symbol $[g_1 | \dots | g_n]$ corresponds to the $(n+1)$ -tuple

$$(1, g_1, g_1 g_2, \dots, g_1 \dots g_n)$$

Removing the first entry gives

$$g_1(1, g_2, g_2 g_3, \dots, g_2 \dots g_n)$$

and removing the second entry gives

$$(1, g_1 g_2, \dots, g_1 \dots g_n)$$

V. Group Cohomology

Lemma. The bar resolution is exact.

Proof. We will just consider the d_n as maps of abelian groups. F_n has basis $G \times G^{(n)}$ as a free abelian group.

$$G \times G^{(n)} = \{g_0[g_1 | \dots | g_n] \mid g_0, \dots, g_n \in G\}$$

We define \mathbb{Z} -maps $s_n : F_n \rightarrow F_{n+1}$ such that

$$\text{id}_{F_n} = d_{n+1}s_n + s_{n-1}d_n$$

by

$$s_n(g_0[g_1 | \dots | g_n]) = [g_0 | g_1 | \dots | g_n]$$

This is not a $\mathbb{Z}G$ -map. One can check that the required equation holds. If $x \in \ker d_n$, then

$$x = \text{id } x = d_{n+1}s_n(x) + s_{n-1}d_n(x) = d_{n+1}s_n(x) \in \text{im } d_{n+1}$$

□

Corollary. Any finite group is of type FP_∞ .

Proof. The bar resolution gives a suitable resolution. □

1.3. Cohomology

Definition. Consider a projective resolution

$$\dots \longrightarrow P_{n+1} \longrightarrow P_n \longrightarrow \dots \longrightarrow P_1 \longrightarrow P_0 \longrightarrow \mathbb{Z} \longrightarrow 0$$

of \mathbb{Z} by $\mathbb{Z}G$ -modules. Let M be a (left) $\mathbb{Z}G$ -module. Applying $\text{Hom}_G(-, M)$, we obtain a sequence

$$\dots \longleftarrow \text{Hom}_G(P_{n+1}, M) \longleftarrow \text{Hom}_G(P_n, M) \longleftarrow \dots \longleftarrow \text{Hom}_G(P_1, M) \xleftarrow{d^1} \text{Hom}_G(P_0, M)$$

where $d^n = d_n^*$. Then the n th cohomology group $H^n(G, M)$ with coefficients in M is

$$H^n(G, M) = \ker d^{n+1} / \text{im } d^n; \quad H^0(G, M) = \ker d^1$$

Remark. We have removed the \mathbb{Z} term in the $\text{Hom}_G(-, M)$ sequence. These cohomology groups are the homology groups of a chain complex $C_n = \text{Hom}_G(P_n, M)$ for $n \leq 0$. We will show that these cohomology groups are independent of the choice of projective resolution.

Example. Let $G = \langle t \rangle$ be an infinite cyclic group. We have a projective resolution

$$0 \longrightarrow \mathbb{Z}G \xrightarrow{\cdot(t-1)} \mathbb{Z}G \longrightarrow \mathbb{Z} \longrightarrow 0$$

For $\varphi \in \text{Hom}_G(\mathbb{Z}G, M)$ and $x \in \mathbb{Z}G$,

$$d^1(\varphi)(x) = \varphi(d_1(x)) = \varphi(x(t-1))$$

Recall that we have an isomorphism $i : \text{Hom}_G(\mathbb{Z}G, M) \cong M$ by $\theta \mapsto \theta(1)$. In particular,

$$d^1(\varphi) \mapsto d^1(\varphi)(1) = \varphi(t-1) = (t-1)\varphi(1) = (t-1)i(\varphi)$$

We thus obtain

$$0 \longleftarrow M \xleftarrow{\alpha} M$$

where α is multiplication on the left by $t-1$. Therefore, the cohomology groups are

$$H^0(G, M) = \{m \in M \mid tm = m\} = M^G; \quad H^1(G, M) = M / (t-1)M = M_G; \quad H^n(G, M) = 0 \text{ for } n \neq 0, 1$$

Note that the group of invariants M^G is the largest submodule with trivial G -action, and the group of coinvariants M_G is the largest quotient module with trivial G -action.

Remark. It is generally true that $H^0(G, M) = M^G$, but in general $H^1(G, M) = M_G$ does not hold. In general, M_G is the 0th homology group, which will be discussed later. Note that for any group of type FP , the cohomology groups vanish for all but finitely many indices n .

Definition. G is of cohomological dimension m over \mathbb{Z} if there exists some $\mathbb{Z}G$ -module M with $H^m(G, M) \neq 0$ but $H^n(G, M_1) = 0$ for all $n > m$ and all $\mathbb{Z}G$ -modules M_1 .

Remark. For all G , we have $H^0(G, \mathbb{Z}) = \mathbb{Z} \neq 0$ so all groups have dimension at least zero.

Example. Infinite cyclic groups have cohomological dimension 1 over \mathbb{Z} . One can show that if G is a free group of finite rank, then it is also of cohomological dimension 1 over \mathbb{Z} . Stallings showed in 1968 that the converse is true: a finitely generated group of cohomological dimension 1 is free. Swan strengthened this in 1969 by removing the assumption of finite generation.

We now consider the bar resolution in our definition of cohomology. Note that

$$\text{Hom}_G(\mathbb{Z}G\{G^{(n)}\}, M) \cong C^n(G, M)$$

where $C^n(G, M)$ is the set of functions $G^{(n)} \rightarrow M$, since a $\mathbb{Z}G$ -map is determined by its action on a basis. Moreover, $C^n(G, M)$ corresponds to the set of functions $G^n \rightarrow M$. For $n = 0$, note that $C^0(G, M)$ is the set of functions $G^0 \rightarrow M$ which bijects with M .

V. Group Cohomology

Definition. The abelian group of n -cochains of G with coefficients in M is $C^n(G, M)$. The n th coboundary map $d^n : C^{n-1}(G, M) \rightarrow C^n(G, M)$ is dual to the d_n from the bar resolution:

$$\begin{aligned} d^n(\varphi)(g_1, \dots, g_n) &= g_1\varphi(g_2, \dots, g_n) \\ &\quad - \varphi(g_1g_2, g_3, \dots, g_n) \\ &\quad + \varphi(g_1, g_2g_3, \dots, g_n) - \dots \\ &\quad + (-1)^{n-1}\varphi(g_1, g_2, \dots, g_{n-1}g_n) \\ &\quad + (-1)^n\varphi(g_1, g_2, \dots, g_{n-1}) \end{aligned}$$

The group of n -cocycles is $Z^n(G, M) = \ker d^{n+1} \leq C^n(G, M)$. The group of n -coboundaries is $B^n(G, M) = \operatorname{im} d^n \leq C^n(G, M)$. Thus the n th cohomology group is

$$H^n(G, M) = Z^n(G, M) / B^n(G, M)$$

Corollary. $H^0(G, M) = M^G$ for all G .

Definition. A derivation of G with coefficients in M is a function $\varphi : G \rightarrow M$ such that $\varphi(gh) = g\varphi(h) + \varphi(g)$.

Note that $Z^1(G, M)$ is exactly the set of derivations of G with coefficients in M , so a derivation is precisely a 1-cocycle.

Definition. An inner derivation of G with coefficients in M is a function $\varphi : G \rightarrow M$ of the form $\varphi(g) = gm - m$ for a fixed $m \in M$.

Such maps are derivations.

Corollary. $H^1(G, M)$ is the group of derivations modulo the inner derivations. In particular, if M is a trivial $\mathbb{Z}G$ -module, then

$$H^1(G, M) = \{\text{group homomorphisms } G \rightarrow M\}$$

treating M as an abelian group under addition.

1.4. Independence of cohomology groups

We now prove that cohomology groups are independent of the choice of resolution.

Definition. Let $(A_n, \alpha_n), (B_n, \beta_n)$ be chain complexes of $\mathbb{Z}G$ -modules. A chain map (f_n) is a sequence of $\mathbb{Z}G$ -maps $f_n : A_n \rightarrow B_n$ such that the following diagram commutes.

$$\begin{array}{ccccccc} \cdots & \longrightarrow & A_n & \xrightarrow{\alpha_n} & A_{n-1} & \xrightarrow{\alpha_{n-1}} & A_{n-2} \longrightarrow \cdots \\ & & \downarrow f_n & & \downarrow f_{n-1} & & \downarrow f_{n-2} \\ \cdots & \longrightarrow & B_n & \xrightarrow{\beta_n} & B_{n-1} & \xrightarrow{\beta_{n-1}} & B_{n-2} \longrightarrow \cdots \end{array}$$

Lemma. A chain map (f_n) as above induces a map on homology groups

$$f_* : H_n(A_*) \rightarrow H_n(B_*)$$

Proof. Let $x \in \ker \alpha_n$, and define $f_*([x]) = [f_n(x)]$, where square brackets denote the quotient maps to the relevant homology classes. Observe that $f_n(x) \in \ker \beta_n$, since $\beta_n f_n(x) = f_{n-1} \alpha_n(x) = 0$. Further, if $x' = x + \alpha_{n+1}(y)$ for some y , we obtain

$$f_n(x') = f_n(x) + f_n \alpha_{n+1}(y) = f_n(x) + \beta_{n+1} f_{n+1}(y) \in f_n(x) + \text{im } b_{n+1}$$

Therefore, this map is well-defined. One can check that this is a map of abelian groups, as required. \square

Theorem. The definition of $H^n(G, M)$ does not depend on the choice of resolution.

Proof. Take projective resolutions (P_n, d_n) and (P'_n, d'_n) of \mathbb{Z} by projective $\mathbb{Z}G$ -modules. We will produce $\mathbb{Z}G$ -maps $f_n : P_n \rightarrow P'_n$ and $g_n : P'_n \rightarrow P_n$ satisfying

$$f_{n-1} d_n = d'_n f_n; \quad g_{n-1} d'_n = d_n g_n$$

as well as maps $s_n : P_n \rightarrow P_{n+1}$ and $s'_n : P'_n \rightarrow P'_{n+1}$ satisfying

$$d_{n+1} s_n + s_{n-1} d_n = g_n f_n - \text{id}; \quad d'_{n+1} s'_n + s'_{n-1} d'_n = f_n g_n - \text{id}$$

Thus, the f_n and g_n form chain maps, and the s_n and s'_n form *chain homotopies*. The chain maps $(f_n), (g_n)$ give rise to chain maps

$$\text{Hom}_G(P', M) \rightarrow \text{Hom}_G(P, M); \quad \text{Hom}_G(P, M) \rightarrow \text{Hom}_G(P', M)$$

giving maps between the respective homology groups by the previous lemma. We now observe that if $\varphi \in \ker d^{n+1} \in \text{Hom}(P, M)$, we have

$$\begin{aligned} f_n^* g_n^* (\varphi)(x) &= \varphi(g_n f_n(x)) \\ &= \varphi(x) + \varphi(d_{n+1} s_n(x)) + \varphi(s_{n-1} d_n(x)) \\ &= \varphi(x) + s_n^* d^{n+1} \varphi(x) + d^n s_{n-1}^* (\varphi)(x) \\ &= \varphi(x) + 0 + d^n s_{n-1}^* (\varphi)(x) \end{aligned}$$

Thus $f_n^* g_n^* (\varphi) = \varphi + d^n s_{n-1}^* (\varphi)$, and so $f_n^* g_n^*$ induces the identity map on $\ker d^{n+1} / \text{im } d^n$. The same holds for $g_n^* f_n^*$, and so f_n^*, g_n^* define isomorphisms of homology groups as desired.

It remains to construct the maps f_n, g_n, s_n, s'_n . At the end of the resolutions, we set $f_{-1} : \mathbb{Z} \rightarrow \mathbb{Z}$ and $f_{-2} : 0 \rightarrow 0$ to be the identity maps. Suppose that we have already defined f_{n-1} and f_n ; we will define f_{n+1} . We have $f_n d_{n+1} : P_{n+1} \rightarrow P'_n$ and $d'_n \circ (f_n d_{n+1}) = f_{n-1} d_n d_{n+1} = 0$. Hence,

V. Group Cohomology

the map $f_n d_{n+1}$ has image inside $\ker d'_n$. We then define f_{n+1} to complete the following diagram, which exists by projectivity.

$$\begin{array}{ccccccc}
 & & P_{n+1} & \xrightarrow{d_{n+1}} & P_n & \xrightarrow{d_n} & P_{n-1} \\
 & \swarrow f_{n+1} & \downarrow f_n d_{n+1} & & \downarrow f_n & & \downarrow f_{n-1} \\
 P'_{n+1} & \xrightarrow[d'_{n+1}]{} & \ker d'_n & \xrightarrow{\quad} & P'_n & \xrightarrow{d'_n} & P'_{n-1}
 \end{array}$$

We can define g_{n+1} in the same way. Now set $h_n = g_n f_n - \text{id} : P_n \rightarrow P_n$; this gives a chain map $P_\bullet \rightarrow P_\bullet$. Set $s_{-1} : \mathbb{Z} \rightarrow P_0$ to be the zero map. Note that $d_0 h_0 = h_{-1} d_0 = 0$, and so $\text{im } h_0 \subseteq \ker d_0$. We now use projectivity to define

$$\begin{array}{ccccccc}
 & & P_0 & \xrightarrow{\quad} & \mathbb{Z} & & \\
 & \swarrow s_0 & \downarrow h_0 & \searrow h_0 & \searrow 0 & & \\
 P_1 & \xrightarrow[d_1]{} & \ker d_0 & \xrightarrow{\quad} & P_0 & \xrightarrow{d_0} & \mathbb{Z}
 \end{array}$$

Suppose that s_{n-1} and s_{n-2} are already defined. Consider $t_n = h_n - s_{n-1} d_n : P_n \rightarrow P_n$. We have

$$d_n t_n = d_n h_n - d_n s_{n-1} d_n = h_{n-1} d_n - (h_{n-1} - s_{n-2} d_{n-1}) d_n = s_{n-2} d_{n-1} d_n = 0$$

Thus $\text{im } t_n \subseteq \ker d_n$.

$$\begin{array}{ccccccc}
 & & P_n & \xrightarrow{d_n} & P_{n-1} & & \\
 & \swarrow s_n & \downarrow t_n & \searrow h_n & \downarrow s_{n-1} & \searrow h_{n-1} & \\
 P_{n+1} & \xrightarrow{\quad} & \ker d_n & \xrightarrow{\quad} & P_n & \xrightarrow{d_n} & P_{n-1}
 \end{array}$$

We define the s'_n similarly. □

Remark. For any left $\mathbb{Z}G$ -module N , we can take a resolution of N by projective or free $\mathbb{Z}G$ -modules.

$$\cdots \longrightarrow P_2 \longrightarrow P_1 \longrightarrow P_0 \longrightarrow N \longrightarrow 0$$

Repeating the constructions outlined in this section, applying $\text{Hom}_G(-, M)$ gives homology groups called $\text{Ext}_{\mathbb{Z}G}^n(N, M)$. Thus

$$H^n(G, M) = \text{Ext}_{\mathbb{Z}G}^n(\mathbb{Z}, M)$$

2. Low degree cohomology and group extensions

2.1. Degree 1

Recall that $H^0(G, M)$, the group M^G of invariants of M under G . A derivation is a 1-cocycle, or equivalently a map $\varphi : G \rightarrow M$ such that $\varphi(g_1 g_2) = g_1 \varphi(g_2) + \varphi(g_1)$, and an inner derivation is a map of the form $\varphi(g) = gm - m$. We present two interpretations of (inner) derivations.

First interpretation. Consider possible $\mathbb{Z}G$ -actions on the abelian group $M \oplus \mathbb{Z}$ of the form $g(m, n) = (gm + n\varphi(g), n)$. Then

$$g_1(g_2(m, n)) = g_1(g_2 m + n\varphi(g_2), n) = (g_1 g_2 m + n g_1 \varphi(g_2) + n\varphi(g_1), n)$$

and

$$(g_1 g_2)(m, n) = (g_1 g_2 m + n\varphi(g_1 g_2), n)$$

For these to coincide, we must require $\varphi(g_1 g_2) = g_1 \varphi(g_2) + \varphi(g_1)$, which is to say that φ is a derivation. In particular, if M is a free \mathbb{Z} -module of finite rank, then we obtain a map

$$g \mapsto \begin{pmatrix} \theta_1(g) & \varphi(g) \\ 0 & 1 \end{pmatrix}$$

where $\theta_1(g)$ is a matrix corresponding to the action of g on M . This is a group homomorphism only if φ is a derivation. One can check that φ is an inner derivation if $(-m, 1)$ generates a $\mathbb{Z}G$ -submodule of M which is the trivial module.

Second interpretation. We first make the following definition.

Definition. Let G be a group and M be a left $\mathbb{Z}G$ -module. We construct the *semidirect product* $M \rtimes G$ by defining a group operation on the set $M \times G$ as follows.

$$(m_1, g_1) * (m_2, g_2) = (m_1 + g_1 m_2, g_1 g_2)$$

Then $M \cong \{(m, 1) \mid m \in M\}$ is a normal subgroup of $M \rtimes G$. Also, $G \cong \{(0, g) \mid g \in G\}$, and conjugation by $\{(0, g) \mid g \in G\}$ corresponds to the G -action on the module M . Further,

$$M \rtimes G / \{(m, 1) \mid m \in M\} \cong G$$

There is a group homomorphism $s : G \rightarrow M \rtimes G$ given by $g \mapsto (0, g)$, such that $\pi_2 \circ s = \text{id}$ where π_2 is the second projection. Such a map s is called a *splitting*. Given another splitting $s_1 : G \rightarrow M \rtimes G$ such that $\pi_2 \circ s_1 = \text{id}$, we define $\psi_{s_1} : G \rightarrow M$ by

$$s_1(g) = (\psi_{s_1}(g), g) \in M \rtimes G$$

Then ψ_{s_1} is a 1-cocycle. Given two splittings s_1, s_2 , the difference $\psi_{s_1} - \psi_{s_2}$ is a coboundary precisely when there exists m such that $(m, 1)s_1(g)(m, 1)^{-1} = s_2(g)$. Conversely, given a 1-cocycle $\varphi \in Z^1(G, M)$, there is a splitting $s_1 : G \rightarrow M \rtimes G$ such that $\varphi = \psi_{s_1}$.

Theorem. $H^1(G, M)$ bijects with the M -conjugacy classes of splittings.

V. Group Cohomology

2.2. Degree 2

Definition. Let G be a group and M be a $\mathbb{Z}G$ -module. An *extension* of G by M is a group E with an exact sequence of group homomorphisms

$$0 \longrightarrow M \longrightarrow E \longrightarrow G \longrightarrow 1$$

M embeds into E , so its image (also called M) is an abelian normal subgroup of E . This is acted on by conjugation by E , and so we obtain an induced action of $E/M \cong G$, which must match the given G -action on M .

Example. The semidirect product $M \rtimes G$ is an extension of G by M .

$$0 \longrightarrow M \longrightarrow M \rtimes G \longrightarrow G \longrightarrow 1$$

In this case, the extension is called a *split extension*, since there is a splitting.

Definition. Two extensions are *equivalent* if there is a commutative diagram of homomorphisms

$$\begin{array}{ccccccc} & & & E & & & \\ & & \nearrow & \vdots & \searrow & & \\ 0 & \longrightarrow & M & & & G & \longrightarrow 1 \\ & & \searrow & \vdots & \nearrow & & \\ & & & E' & & & \end{array}$$

If E, E' are equivalent extensions, then E and E' are isomorphic as groups. The converse is false.

Definition. A *central* extension is an extension where the given $\mathbb{Z}G$ -module is a trivial module (that is, it has trivial G -action).

Proposition. Let E be an extension of G by M . If there is a splitting homomorphism $s_1 : G \rightarrow E$, then the extension is equivalent to

$$0 \longrightarrow M \longrightarrow M \rtimes G \longrightarrow G \longrightarrow 1$$

and thus $E \cong M \rtimes G$.

Theorem. Let G be a group and let M be a $\mathbb{Z}G$ -module. Then there is a bijection from $H^2(G, M)$ to the set of equivalence classes of extensions of G by M .

Proof. Given an extension

$$0 \longrightarrow M \longrightarrow E \longrightarrow G \longrightarrow 1$$

2. Low degree cohomology and group extensions

there is a set-theoretic section $s : G \rightarrow E$ such that

$$\begin{array}{ccc} G & \xrightarrow{s} & E \\ & \searrow \text{id} & \downarrow \pi \\ & & G \end{array}$$

commutes. Note that s need not be a group homomorphism. Without loss of generality, we can suppose $s(1) = 1$. We define a map

$$\varphi(g_1, g_2) = s(g_1)s(g_2)s(g_1g_2)^{-1}$$

which measures the failure of s to be a group homomorphism. Then $\pi(\varphi(g_1, g_2)) = 1$, and so $\varphi(g_1, g_2) \in M$. Thus $\varphi : G^2 \rightarrow M$ is a 2-cochain, and we can show it is a 2-cocycle. We have

$$\begin{aligned} s(g_1)s(g_2)s(g_3) &= \varphi(g_1, g_2)s(g_1g_2)s(g_3) \\ &= \varphi(g_1, g_2)\varphi(g_1g_2, g_3)s(g_1g_2g_3) \end{aligned}$$

and similarly,

$$\begin{aligned} s(g_1)s(g_2)s(g_3) &= s(g_1)\varphi(g_2, g_3)s(g_2g_3) \\ &= s(g_1)\varphi(g_2, g_3)s(g_1)^{-1}s(g_1)s(g_2g_3) \\ &= s(g_1)\varphi(g_2, g_3)s(g_1)^{-1}\varphi(g_1, g_2g_3)s(g_1g_2g_3) \end{aligned}$$

We therefore obtain

$$\begin{aligned} \varphi(g_1, g_2)\varphi(g_1g_2, g_3)s(g_1g_2g_3) &= s(g_1)\varphi(g_2, g_3)s(g_1)^{-1}\varphi(g_1, g_2g_3)s(g_1g_2g_3) \\ \varphi(g_1, g_2)\varphi(g_1g_2, g_3) &= s(g_1)\varphi(g_2, g_3)s(g_1)^{-1}\varphi(g_1, g_2g_3) \end{aligned}$$

Converting into additive notation,

$$\varphi(g_1, g_2) + \varphi(g_1g_2, g_3) = g_1\varphi(g_2, g_3) + \varphi(g_1, g_2g_3)$$

and so

$$(d^3\varphi)(g_1, g_2, g_3) = 0$$

Hence φ is a 2-cocycle as claimed. Note that φ is a *normalised* cocycle: it satisfies $\varphi(1, g) = \varphi(g, 1) = 0$. We have therefore proven that an extension of G by M , with a choice of set-theoretic section $s : G \rightarrow E$, yields a normalised 2-cocycle $\varphi \in Z^2(G, M)$.

Now take another choice of section s' with $s'(1) = 1$. We show that the normalised cocycles φ, φ' differ by a coboundary, and so we have a map defined from equivalence classes of extensions to $H^2(G, M)$. We have $\pi(s(g)s'(g)^{-1}) = 1$, so $s(g)s'(g)^{-1} \in \ker \pi = M$. Let $\psi(g)$ denote $s(g)s'(g)^{-1}$. Thus $\psi : G \rightarrow M$. We have

$$\begin{aligned} s'(g_1)s'(g_2) &= \psi(g_1)s(g_1)\psi(g_2)s(g_2) \\ &= \psi(g_1)s(g_1)\psi(g_2)s(g_1)^{-1}s(g_1)s(g_2) \\ &= \psi(g_1)s(g_1)\psi(g_2)s(g_1)^{-1}\varphi(g_1, g_2)s(g_2) \\ &= \psi(g_1)s(g_1)\psi(g_2)s(g_1)^{-1}\varphi(g_1, g_2)\psi(g_1g_2)^{-1}s'(g_1g_2) \end{aligned}$$

V. Group Cohomology

Switching to additive notation,

$$\begin{aligned}\varphi'(g_1, g_2) &= \psi(g_1) + g_1\psi(g_2) + \varphi(g_1, g_2) - \psi(g_1g_2) \\ &= \varphi(g_1, g_2) + (d^2\psi)(g_1, g_2)\end{aligned}$$

Thus φ and φ' differ by a coboundary, and so we have a well-defined map from extensions of G by M to $H^2(G, M)$.

To complete the proof, we must check that equivalent extensions give rise to the same cohomology class, and that there is an inverse map from cohomology classes to equivalence classes of extensions. To produce the inverse, we use the following lemma.

Lemma. Let $\varphi \in Z^2(G, M)$. Then there is a cochain $\psi \in C^1(G, M)$ such that $\varphi + d^2\psi$ is a normalised cocycle. Hence, every cohomology class can be represented by a normalised cocycle.

Proof. Let $\psi(g) = -\varphi(1, g)$. Then

$$\begin{aligned}(\varphi + d^2\psi)(1, g) &= \varphi(1, g) - (\varphi(1, g) - \varphi(1, g) + \varphi(1, 1)) \\ &= \varphi(1, g) - \varphi(1, 1)\end{aligned}$$

Similarly, we obtain

$$(\varphi + d^2\psi)(g, 1) = \varphi(g, 1) - g\varphi(1, 1)$$

But we know that

$$d^3\varphi(1, 1, g) = 0 = d^3\varphi(g, 1, 1)$$

since φ is a cocycle. Hence, one can check computationally that both equations above are zero. \square

We now take a normalised cocycle φ representing a given cohomology class. We construct an extension

$$0 \longrightarrow M \longrightarrow E_\varphi \longrightarrow G \longrightarrow 1$$

by

$$(m_1, g_1) * (m_2, g_2) = (m_1 + g_1m_2 + \varphi(g_1, g_2), g_1g_2)$$

For this to be a group operation, we use the fact that φ is normalised. This yields an extension

$$0 \longrightarrow M \longrightarrow E_\varphi \xrightarrow{\pi} G \longrightarrow 1$$

where π is the projection onto the second component. Note that if φ' is another normalised 2-cocycle representing the given cohomology class, then $\varphi - \varphi'$ is a coboundary, so we can define a map $E_\varphi \rightarrow E_{\varphi'}$ by

$$(m, g) \mapsto (m + \psi(g), g)$$

One can check that this induces an equivalence of extensions. These constructions are inverses. \square

2.3. Central extensions

Example. Consider central extensions of \mathbb{Z}^2 by \mathbb{Z} . We already know of two such extensions. The first is

$$0 \longrightarrow \mathbb{Z} \longrightarrow \mathbb{Z}^3 \longrightarrow \mathbb{Z}^2 \longrightarrow 0$$

$$m \longmapsto (m, 0, 0)$$

$$(m, r, s) \longmapsto (r, s)$$

Let H denote the *Heisenberg group*

$$H = \left\{ \begin{pmatrix} 1 & r & m \\ 0 & 1 & s \\ 0 & 0 & 1 \end{pmatrix} \middle| r, s, m \in \mathbb{Z} \right\}$$

Then we have the extension

$$0 \longrightarrow \mathbb{Z} \longrightarrow H \longrightarrow \mathbb{Z}^2 \longrightarrow 0$$

$$m \longmapsto \begin{pmatrix} 1 & 0 & m \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

$$\begin{pmatrix} 1 & r & m \\ 0 & 1 & s \\ 0 & 0 & 1 \end{pmatrix} \longmapsto (r, s)$$

Writing multiplicatively, let $T \cong \mathbb{Z}^2$ be generated by a and b . We have the following free resolution of the trivial $\mathbb{Z}T$ -module \mathbb{Z} .

$$0 \longrightarrow \mathbb{Z}T \xrightarrow{\beta} \mathbb{Z}T^2 \xrightarrow{\alpha} \mathbb{Z}T \xrightarrow{\varepsilon} \mathbb{Z} \longrightarrow 0$$

where

$$\begin{aligned} \beta(z) &= (z(1-b), z(a-1)) \\ \alpha(x, y) &= x(a-1) + y(b-1) \end{aligned}$$

and ε is the augmentation map. Apply $\text{Hom}_T(-, \mathbb{Z})$ to obtain the chain complex

$$0 \longleftarrow \text{Hom}_T(\mathbb{Z}T, \mathbb{Z}) \xleftarrow{\beta^*} \text{Hom}_T(\mathbb{Z}T^2, \mathbb{Z}) \xleftarrow{\alpha^*} \text{Hom}_T(\mathbb{Z}T, \mathbb{Z})$$

V. Group Cohomology

We claim that α^* and β^* are both zero maps, and so

$$H^2(T, \mathbb{Z}) = \text{Hom}_T(\mathbb{Z}T, \mathbb{Z}) \cong \mathbb{Z}$$

and the generator is represented by the augmentation map $\varepsilon : \mathbb{Z}T \rightarrow \mathbb{Z}$.

Take a $\mathbb{Z}T$ -map $f : \mathbb{Z}T^2 \rightarrow \mathbb{Z}$. Then

$$\begin{aligned} (\beta^* f)(z) &= f(\beta)(z) \\ &= f(z(1-b), z(a-1)) \\ &= f(z-zb, 0) + f(0, za-z) \\ &= (1-b)f(z, 0) + (a-1)f(0, z) \\ &= 0 \end{aligned}$$

where the last line holds as T acts trivially. Similarly, $\alpha^* = 0$.

Next, we interpret $H^2(T, \mathbb{Z})$ in terms of 2-cocycles arising from the bar resolution. We construct a chain map as follows.

$$\begin{array}{ccccccc} \mathbb{Z}T\{T^{(2)}\} & \xrightarrow{d_2} & \mathbb{Z}T\{T^{(1)}\} & \xrightarrow{d_1} & \mathbb{Z}T\{T^{(0)}\} & \xrightarrow{\varepsilon} & \mathbb{Z} \longrightarrow 0 \\ f_2 \downarrow & & f_1 \downarrow & & \text{id} \downarrow & & \parallel \\ \mathbb{Z}T & \xrightarrow{\beta} & \mathbb{Z}T^2 & \xrightarrow{\alpha} & \mathbb{Z}T & \longrightarrow & \mathbb{Z} \longrightarrow 0 \end{array}$$

To construct f_1 such that $\alpha f_1 = d_1$, we need to give images of the symbols $[a^r b^s]$ with $r, s \in \mathbb{Z}$. We must have

$$[a^r b^s] \mapsto (x_{r,s}, y_{r,s}) \in \mathbb{Z}T^2$$

where

$$\alpha(x_{r,s}, y_{r,s}) = d_1([a^r b^s]) = a^r b^s - 1 = (a^r - 1)b^s + (b^s - 1)$$

We define

$$S(a, r) = \begin{cases} 1 + a + \dots + a^{r-1} & \text{if } r > 0 \\ -a^{-1} - \dots - a^r & \text{if } r \leq 0 \end{cases}$$

Note that

$$S(a, r)(a-1) = a^r - 1$$

for any $r \in \mathbb{Z}$. Then

$$\begin{aligned} \alpha(S(a, r)b^s, S(b, s)) &= S(a, r)b^s(a-1) + S(b, s)(b-1) \\ &= d_1([a^r b^s]) \end{aligned}$$

as required. So we may define

$$f_1([a^r b^s]) = (S(a, r)b^s, S(b, s))$$

2. Low degree cohomology and group extensions

To define f_2 , we need to give images of the symbols $[a^r b^s | a^t b^u]$. For each such symbol, we find $z_{r,s,t,u} \in \mathbb{Z}T$ such that

$$f_1 d_2([a^r b^s | a^t b^u]) = \beta(z_{r,s,t,u})$$

We can explicitly calculate

$$\begin{aligned} f_1 d_2([a^r b^s | a^t b^u]) &= f_1(a^r b^s [a^t b^u] - [a^{r+t} b^{s+u}] - [a^r b^s]) \\ &= (a^r b^s S(a, t) b^u - S(a, r+t) b^{s+u} + S(a, r) b^s, a^r b^s S(b, u) - S(b, s+u) + S(b, s)) \end{aligned}$$

So defining

$$z_{r,s,t,u} = S(a, r) b^s S(b, u)$$

gives the required equation.

$$f_2([a^r b^s | a^t b^u]) = S(a, r) b^s S(b, u)$$

Now we find a cochain $\varphi : T^2 \rightarrow \mathbb{Z}$ representing the cohomology class $p \in \mathbb{Z} = \text{Hom}_T(\mathbb{Z}T, \mathbb{Z}) = H^2(T, \mathbb{Z})$. Such a cochain is given by the composition

$$T^2 \xrightarrow{f_2} \mathbb{Z}T \xrightarrow{p\varepsilon} \mathbb{Z}$$

Since $\varepsilon(S(a, r)) = r$, we find

$$\varphi(a^r b^s, a^t b^u) = p\varepsilon(z_{r,s,t,u}) = pr u$$

The group structure on $\mathbb{Z} \times T$ corresponding to this is

$$(m, a^r b^s) * (n, a^t b^u) = (m + n + pr u, a^{r+t} b^{s+u})$$

This corresponds to the group of matrices

$$\left\{ \begin{pmatrix} 1 & pr & m \\ 0 & 1 & s \\ 0 & 0 & 1 \end{pmatrix} \middle| r, s, m \in \mathbb{Z} \right\}$$

2.4. Generators and relations

Another approach to considering extensions, and in particular central extensions, is the use of partial resolutions arising from generators and relations. Given a group G , for any generating set X there is a canonical map $F \rightarrow G$ where F is the free group on X . Let R be the kernel of this map, and so we have a short exact sequence

$$1 \longrightarrow R \longrightarrow F \longrightarrow G \longrightarrow 1$$

This is a presentation for G , where the subgroup R can be thought of as the set of relations. Since it is a normal subgroup, F acts on it by conjugation. Often we take a set of generators of R as a normal subgroup of F .

V. Group Cohomology

Let $R_{\text{ab}} = R/R'$ be the largest abelian quotient of R . We say that R' is the *derived subgroup* of R , and is given by the commutator subgroup $[R, R]$ of F . It inherits an action of F , but R acts trivially, so we have an induced action by $G = F/R$. Clearly R_{ab} is a \mathbb{Z} -module, and it is a $\mathbb{Z}G$ -module. This is called the *relation module*. We have an extension

$$1 \longrightarrow R_{\text{ab}} \longrightarrow F/R' \longrightarrow G \longrightarrow 1$$

To get a central extension, we instead consider

$$1 \longrightarrow R/[R, F] \longrightarrow F/[R, F] \longrightarrow G \longrightarrow 1$$

where $[R, F]$ is the commutator subgroup. There is not a largest or universal central extension, since we can always form the direct product with an abelian group, but this particular central extension above does have some good properties that we will now explore.

Theorem. Let

$$1 \longrightarrow R \longrightarrow F \longrightarrow G \longrightarrow 1$$

be a presentation of G . Let M be a left $\mathbb{Z}G$ -module. Then there is an exact sequence

$$H^1(F, M) \longrightarrow \text{Hom}_G(R_{\text{ab}}, M) \longrightarrow H^2(G, M) \longrightarrow 0$$

Thus, any equivalence class of extensions of G by M corresponding to a cohomology class in $H^2(G, M)$ arises from a $\mathbb{Z}G$ -map $R_{\text{ab}} \rightarrow M$.

Note that M is a $\mathbb{Z}F$ -module via the map $F \rightarrow G$.

Corollary. In the above situation, if M is a trivial $\mathbb{Z}G$ -module, then we have an exact sequence

$$\text{Hom}(F, M) \longrightarrow \text{Hom}_G(R/[R, F], M) \longrightarrow H^2(G, M) \longrightarrow 0$$

Proof. M is a trivial $\mathbb{Z}F$ -module, so $H^1(F, M) = \text{Hom}(F, M)$, which is a set of group homomorphisms to an abelian group, and any such morphism factors uniquely through the abelianisation so this is equal to $\text{Hom}(F_{\text{ab}}, M)$. Similarly, $\text{Hom}_G(R_{\text{ab}}, M) = \text{Hom}_G(R/[R, F], M)$. \square

2.5. Homology groups

There is also a connection with homology groups. Given a projective resolution of the trivial $\mathbb{Z}G$ -module \mathbb{Z} , we can apply the map $\mathbb{Z} \otimes_{\mathbb{Z}G} -$ and obtain homology groups. The homology groups do not depend on the choice of resolution, and are written $H_n(G, \mathbb{Z})$.

Definition. The *Schur multiplier* $M(G)$ of a group G is the second homology group $H_2(G, \mathbb{Z})$.

2. Low degree cohomology and group extensions

Theorem (universal coefficients theorem). Let G be a group and M be a trivial $\mathbb{Z}G$ -module. Then there is a short exact sequence

$$0 \longrightarrow \text{Ext}^1(G_{\text{ab}}, M) \longrightarrow H^2(G, M) \longrightarrow \text{Hom}(M(G), M) \longrightarrow 0$$

where $\text{Ext}^1(G_{\text{ab}}, M)$ arises from applying $\text{Hom}_{\mathbb{Z}}(-, M)$ to a projective resolution of the abelian group G_{ab} .

Corollary. Suppose that $G = G'$, and so $G_{\text{ab}} = 1$. Then $H^2(G, M) \cong \text{Hom}(M(G), M)$.

In some texts, the Schur multiplier is defined to be $H^2(G, \mathbb{C}^\times)$, where \mathbb{C}^\times is the a trivial module written multiplicatively. This approach can be useful when considering projective representations $G \rightarrow PGL(\mathbb{C})$. Such a map lifts to give a linear representation of central extension of G .

Theorem (Hopf's formula). Given a presentation

$$1 \longrightarrow R \longrightarrow F \longrightarrow G \longrightarrow 1$$

we have

$$M(G) \cong F' \cap R / [R, F]$$

Note that this is not necessarily all of $F/[R, F]$, and this shows that $F' \cap R / [R, F]$ is independent of the choice of presentation.

VI. Large Cardinals

Lectured in Lent 2024 by DR. B. LÖWE

(Course description goes here.)

Contents

1.	Inaccessible cardinals	297
1.1.	Large cardinal properties	297
1.2.	Weakly inaccessible and inaccessible cardinals	298
1.3.	Second order replacement	299
1.4.	Countable transitive models of set theory	301
1.5.	Worldly cardinals	302
1.6.	The consistency strength hierarchy	303
2.	Measurability and compactness	306
2.1.	The measure problem	306
2.2.	Real-valued measurable cardinals	307
2.3.	Measurable cardinals	308
2.4.	Weakly compact cardinals	310
2.5.	Strongly compact cardinals	312
3.	Reflection	314
3.1.	The Keisler extension property	314
3.2.	Ultrapowers of the universe	316
3.3.	Properties above the critical point	319
3.4.	The fundamental theorem on measurable cardinals	322
4.	Towards inconsistency	325
4.1.	Strong cardinals	325
4.2.	Removing the inaccessible	326
4.3.	Supercompact cardinals	327
4.4.	The upper limit	328

1. Inaccessible cardinals

1.1. Large cardinal properties

Modern set theory largely concerns itself with the consequences of the incompleteness phenomenon. Given any ‘reasonable’ set theory T , Gödel’s first incompleteness theorem shows that there is a sentence φ such that $T \not\vdash \varphi$ and $T \not\vdash \neg\varphi$. To be ‘reasonable’, the set of axioms must be computably enumerable, among other similar restrictions. In particular, Gödel’s second incompleteness theorem shows that $T \not\vdash \text{Con}(T)$, where $\text{Con}(T)$ is the statement that T is consistent. Hence,

$$\{\psi \mid T \vdash \psi\} \subsetneq \{\psi \mid T + \varphi \vdash \psi\}$$

We might say

$$T <_{\text{consequence}} T + \varphi$$

so T has strictly fewer consequences than $T + \varphi$. Modern set theory is about understanding the relation $\leq_{\text{consequence}}$ and other similar relations. It turns out that large cardinal axioms are the most natural hierarchy that we can use to measure the strength of set theories.

In this course we will not provide a definition for the notion of ‘large cardinal’, but we will provide an informal description. A *large cardinal property* is a formula Φ such that $\Phi(\kappa)$ implies that κ is a very large cardinal, so large that its existence cannot be proven in ZFC. A *large cardinal axiom* is an axiom of the form $\exists\kappa. \Phi(\kappa)$, which we will abbreviate ΦC . We begin with some non-examples.

- (i) κ is called an *aleph fixed point* if $\kappa = \aleph_\kappa$. Note that, for example, ω , ω_1 , and \aleph_ω are not aleph fixed points. However, we have the following result. We say that $F : \text{Ord} \rightarrow \text{Ord}$ is *normal* if $\alpha < \beta$ implies $F(\alpha) < F(\beta)$, and if λ is a limit, $F(\lambda) = \bigcup_{\alpha < \lambda} F(\alpha)$. One can show that every normal ordinal operation has arbitrarily large fixed points, and in particular that these fixed points may be enumerated by the ordinals. In particular, since the operation $\alpha \mapsto \aleph_\alpha$ is normal, it admits fixed points.

- (ii) Let $\Phi(\kappa)$ be the property

$$\kappa = \aleph_\kappa \wedge \text{Con}(\text{ZFC})$$

Clearly ΦC implies $\text{Con}(\text{ZFC})$, so $\text{ZFC} \not\vdash \Phi\text{C}$. We would like our large cardinal axioms to be unprovable by ZFC because of the size of the cardinal in question, not because of any other arbitrary reasons that we may attach to these axioms.

One source of large cardinal axioms is as follows. Consider the ordinal ω ; it is much larger than any ordinal smaller than it. We can consider properties that encapsulate the notion that ω is much larger than any smaller ordinal, and use these as large cardinal properties.

- (i) If $n < \omega$, then $n^+ < \omega$, where n^+ is the cardinal successor of n . We define

$$\Lambda(\kappa) \iff \forall\alpha. (\alpha < \kappa \rightarrow \alpha^+ < \kappa)$$

VI. Large Cardinals

where α^+ is the least cardinal strictly larger than α . Then, $\Lambda(\kappa)$ holds precisely when κ is a limit cardinal. These need not be very large, for example, \aleph_ω is a limit cardinal, and the existence of this cardinal is proven by ZFC.

- (ii) If $n < \omega$, then $2^n < \omega$, where 2^n is the size of the power set of n .

$$\Sigma(\kappa) \iff \forall \alpha. (\alpha < \kappa \rightarrow 2^\alpha < \kappa)$$

where 2^α is the cardinality of $\mathcal{P}(\alpha)$. Such cardinals are called *strong limit cardinals*. We will show that these exist in all models of ZFC. Similarly to the aleph hierarchy, we can define the *beth* hierarchy, denoted \beth_α . This is given by

$$\beth_0 = \aleph_0; \quad \beth_{\alpha+1} = 2^{\beth_\alpha}; \quad \beth_\lambda = \bigcup_{\alpha < \lambda} \beth_\alpha$$

Cantor's theorem shows that $\aleph_\alpha \leq \beth_\alpha$, and the continuum hypothesis is the assertion that $\aleph_1 = \beth_1$. Note that κ is a strong limit cardinal if and only if $\kappa = \beth_\lambda$ for some limit ordinal λ . In particular, $\text{ZFC} \vdash \Sigma\mathcal{C}$.

- (iii) If $s : n \rightarrow \omega$ for $n < \omega$, then $\sup(s) = \bigcup \text{ran}(s) < \omega$. This gives rise to the following definition.

Definition. Let λ be a limit ordinal. We say that $C \subseteq \lambda$ is *cofinal* or *unbounded* if $\bigcup C = \lambda$. We define the *cofinality* of λ , denoted $\text{cf}(\lambda)$, to be the cardinality of the smallest cofinal subset. If λ is a cardinal, then $\text{cf}(\lambda) \leq \lambda$. If this inequality is strict, the cardinal is called *singular*; if this is an equality, it is called *regular*.

Note that if κ is regular, then if $\lambda < \kappa$, and for each $\alpha < \lambda$ we have a set $X_\alpha \subseteq \kappa$ of size $|X_\alpha| < \kappa$, then $\bigcup X_\alpha \neq \kappa$. It is easy to show that this property is equivalent to regularity.

We have therefore shown that ω is a regular cardinal. Note that \aleph_1 is also regular, since countable unions of countable sets are countable. This argument generalises to all successor cardinals, so all successor cardinals $\aleph_{\alpha+1}$ are regular. The cardinal \aleph_ω is not regular, as it is the union of $\{\aleph_n \mid n \in \mathbb{N}\}$, which is a subset of \aleph_ω of cardinality \aleph_0 , giving $\text{cf}(\aleph_\omega) = \aleph_0$. The cofinality of \aleph_{\aleph_ω} is also \aleph_0 . Limit cardinals are often singular.

1.2. Weakly inaccessible and inaccessible cardinals

Motivated by these examples of properties of ω , we make the following definition.

Definition. A cardinal κ is called *weakly inaccessible* if it is an uncountable regular limit, and (*strongly*) *inaccessible* if it is an uncountable regular strong limit. We write $\text{WI}(\kappa)$ to denote that κ is weakly inaccessible, and $\text{I}(\kappa)$ if κ is inaccessible.

To argue that these are large cardinal properties, we will show that they are very large, and that the existence of such cardinals cannot be proven in ZFC. Note that we cannot actually

prove this statement; if ZFC were inconsistent, it would prove every statement. Whenever we write statements such as $\text{ZFC} \not\vdash \text{IC}$, it should be interpreted to mean ‘if ZFC is consistent, it does not prove IC’.

Many things in the relationship of WI and I are unclear: 2^{\aleph_0} is clearly not inaccessible as it is not a strong limit, but it is not clear that this is not a limit. The *generalised continuum hypothesis* GCH is that for all cardinals α , we have $2^{\aleph_\alpha} = \aleph_{\alpha+1}$, and so $\aleph_\alpha = \beth_\alpha$. Under this assumption, the notions of limit and strong limit coincide, and so the notions of inaccessible cardinals and weakly inaccessible cardinals coincide.

Proposition. Weakly inaccessible cardinals are aleph fixed points.

Proof. Suppose κ is weakly inaccessible but $\kappa < \aleph_\kappa$. Fix α such that $\kappa = \aleph_\alpha$, then $\alpha < \kappa$. As κ is a limit cardinal, α must be a limit ordinal. But then $\aleph_\alpha = \bigcup_{\beta < \alpha} \aleph_\beta$, so in particular, the set $\{\aleph_\beta \mid \beta < \alpha\}$ is cofinal in κ , but this set is of size $|\alpha| < \kappa$. Hence κ is singular, contradicting regularity. \square

1.3. Second order replacement

We will now show that ZFC does not prove IC, and we omit the result for weakly inaccessible cardinals. We could do this via model-theoretic means: we assume $M \models \text{ZFC}$, and construct a model $N \models \text{ZFC} + \neg \text{IC}$. However, there is another approach we will take here. By Gödel’s second incompleteness theorem, under the assumption that ZFC is consistent, we have $\text{ZFC} \not\vdash \text{Con}(\text{ZFC})$, so it suffices to show $\text{IC} \rightarrow \text{Con}(\text{ZFC})$. Gödel’s completeness theorem states that $\text{Con}(T)$ holds if and only if there exists a model M with $M \models T$. Thus, it suffices to show that under the assumption that there is an inaccessible cardinal, we can construct a model of ZFC. Note that the metatheory in which the completeness theorem is proven actually matters; both theories and models are actually sets in the outer theory.

Recall that the *cumulative hierarchy* inside a model of set theory is given by

$$V_0 = \emptyset; \quad V_{\alpha+1} = \mathcal{P}(V_\alpha); \quad V_\lambda = \bigcup_{\alpha < \lambda} V_\alpha$$

- (i) The axiom of foundation is equivalent to the statement that every set is an element of V_α for some α .
- (ii) (V_ω, \in) is a model of all of the axioms of set theory except for the axiom of infinity. This collection of axioms is called *finite set theory* FST.
- (iii) $(V_{\omega+\omega}, \in)$ is a model of all of the axioms of set theory except for the axiom of replacement. This theory is called *Zermelo set theory with choice* ZC. In fact, for any limit ordinal $\lambda > \omega$, ZFC proves that $(V_\lambda, \in) \models \text{ZC}$. That is, ZFC proves the existence of a model of ZC, or equivalently, $\text{ZFC} \vdash \text{Con}(\text{ZC})$. Hence, ZC cannot prove replacement, since Gödel’s second incompleteness theorem applies to ZC. In this way, replacement behaves like a large cardinal axiom for ZC. The same holds for infinity and FST.

VI. Large Cardinals

We briefly discuss why replacement fails in $V_{\omega+\omega}$. Consider the set of ordinals $\omega + n$ for $n < \omega$; this set does not belong to $V_{\omega+\omega}$ as its rank is $\omega + \omega$. However, the class function F given by $n \mapsto \omega + n$ is definable by a simple formula, and applying this to the set $\omega \in V_{\omega+\omega}$ gives a counterexample to replacement. Our counterexample is thus a cofinal subset of $V_{\omega+\omega}$ whose union does not lie in $V_{\omega+\omega}$. In some sense, the fact that $\omega + \omega$ is singular is the reason why $V_{\omega+\omega}$ does not satisfy replacement.

Now, consider $\alpha = \aleph_1$, which is regular. Consider $\mathcal{P}(\omega) \in V_{\omega+2} \subseteq V_{\omega_1}$. There is a definable surjection from $\mathcal{P}(\omega)$ to ω_1 , motivated by the proof of Hartogs' lemma. Indeed, subsets of ω can encode well-orders, and every countable well-order is encoded by a subset of ω , so the map

$$g : A \mapsto \begin{cases} \alpha & \text{if } A \text{ codes a well-order of order type } \alpha \\ 0 & \text{otherwise} \end{cases}$$

is a surjection $\mathcal{P}(\omega) \rightarrow \omega_1$. This class function has cofinal range in ω_1 , and so V_{ω_1} does not satisfy replacement.

We will prove that $I(\kappa)$ implies that V_κ models replacement. A set M is said to satisfy *second-order replacement* SOR if for every function $F : M \rightarrow M$ and every $x \in M$, the set $\{F(y) \mid y \in x\}$ is in M . Any model of V_α that satisfies second-order replacement is a model of ZFC, as the counterexamples to replacement are special cases of violations of second-order replacement.

Theorem (Zermelo). If κ is inaccessible, then V_κ satisfies second-order replacement.

We first prove the following lemmas.

Lemma. If κ is inaccessible and $\lambda < \kappa$, then $|V_\lambda| < \kappa$.

Proof. This follows by induction. Note $|V_0| = 0 < \kappa$. If $|V_\alpha| < \kappa$, then as κ is a strong limit, $|V_{\alpha+1}| = |\mathcal{P}(V_\alpha)| = 2^{|V_\alpha|} < \kappa$. If λ is a limit and $|V_\alpha| < \kappa$ for all $\alpha < \lambda$, then if $|V_\lambda| = \kappa$, we have written κ as a union of less than κ sets of size less than κ , contradicting regularity. \square

Lemma. If κ is inaccessible and $x \in V_\kappa$, then $|x| < \kappa$.

Proof. Suppose $x \in V_\kappa = \bigcup_{\alpha < \kappa} V_\alpha$. Then there exists $\alpha < \kappa$ such that $x \in V_\alpha$. Then $x \subseteq V_\alpha$ as the V_α are transitive, but then $|x| \leq |V_\alpha| < \kappa$. \square

We can now prove Zermelo's theorem.

Proof. Let $F : V_\kappa \rightarrow V_\kappa$, and $x \in V_\kappa$; we must show that $R = \{F(y) \mid y \in x\} \in V_\kappa$. By the second lemma above, $|x| < \kappa$, hence $|R| < \kappa$. For each $y \in x$, define α_y to be the rank of $F(y)$. This is an ordinal less than κ . Consider $A = \{\alpha_y \mid y \in x\}$; its cardinality is bounded by that of x , so $|A| < \kappa$. But as κ is regular, $|A|$ is not cofinal, so there is $\gamma < \kappa$ such that $A \subseteq V_\gamma$. By definition, $R \subseteq V_\gamma$, so $R \in V_{\gamma+1} \subseteq V_\kappa$, as required. \square

The definition of inaccessibility is precisely what is needed for this proof to work. The following converse holds.

Theorem (Shepherdson). If V_κ satisfies second-order replacement, then κ is inaccessible.

Proof. Suppose κ is not inaccessible, so either κ is singular or there is $\lambda < \kappa$ such that $2^\lambda \geq \kappa$. If κ is singular, then $\kappa = \bigcup_{\alpha < \lambda} \kappa_\alpha$ for $\lambda < \kappa$ and $\kappa_\alpha < \kappa$. Consider $C = \{\kappa_\alpha \mid \alpha < \lambda\}$; this set is cofinal in κ , but the cardinality of C is λ . Therefore, $C \notin V_\kappa$. We simply take the function $F : \alpha \mapsto \kappa_\alpha$, then the image of λ under F is $C \notin V_\kappa$, so F witnesses that V_κ violates second-order replacement.

Suppose there is $\lambda < \kappa$ such that $2^\lambda \geq \kappa$. Let $F : \mathcal{P}(\lambda) \rightarrow \kappa$ be a surjection. Since $\lambda < \kappa$, we must have $\mathcal{P}(\lambda) \in V_{\lambda+2} \subseteq V_\kappa$. Then the image of $\mathcal{P}(\lambda)$ under F is $\kappa \notin V_\kappa$ as required. \square

1.4. Countable transitive models of set theory

It is not generally the case that if $V_\kappa \models \text{ZFC}$ then κ is inaccessible. Moreover, the existence of an inaccessible cardinal is strictly stronger than the consistency of ZFC. We will show this second statement first.

Suppose κ is inaccessible, so $V_\kappa \models \text{ZFC}$. A standard model-theoretic argument shows there is a countable elementary substructure $(N, \in) \leq (V_\kappa, \in)$. In particular, $(N, \in) \models \text{ZFC}$. The proof of the downwards Löwenheim–Skolem theorem that we will use is a Skolem hull construction, given by

$$N_0 = \emptyset; \quad N_{k+1} = N_k \cup W(N_k); \quad N = \bigcup_{k \in \mathbb{N}} N_k$$

where $W(N_k)$ is a set of witnesses for all formulas of the form $\exists x. \varphi$ with parameters in N_k . The fact that this is an elementary substructure follows from the Tarski–Vaught test. We will now explore this model in more detail.

If $n \in \omega$, there is a formula φ_n such that $V_\kappa \models \varphi_n(x)$ if and only if $x = n$. Clearly, the formula $\exists x. \varphi_n(x)$ has precisely one witness, so $\omega \subseteq N_1$. Similarly, there are formulas $\varphi_\omega, \varphi_{\omega+\omega}, \varphi_{\omega \cdot 3}$ and so on. There is also a formula φ_{ω_1} such that $x = \omega_1$ if and only if $V_\kappa \models \varphi_{\omega_1}(x)$. As before, because there is a unique witness to this formula in V_κ , we must have $\omega_1 \in N_1$. But since the model N is countable, there must be a gap in the ordinals at some point below ω_1 . By the same argument, the model contains ω_2, ω_3 and so on. Therefore, N is a nontransitive model.

As (N, \in) is well-founded and extensional, by Mostowski's collapsing theorem there is a unique transitive M such that $(M, \in) \cong (N, \in)$. This fills all of the gaps in our model. As this is an isomorphism, we obtain $(M, \in) \leq (N, \in) \leq (V_\kappa, \in)$, so (M, \in) is a countable transitive model of ZFC. In particular, its height $\alpha = \text{Ord} \cap M$ is a countable ordinal. There is an elementary embedding of M into V_κ given by the inverse of the Mostowski collapse. In particular, some $\beta < \alpha$ has the property that $M \models \varphi_{\omega_1}(\beta)$.

VI. Large Cardinals

Therefore, the property ‘ x is a cardinal’ cannot be an *absolute* property between M and V_κ . A property is said to be absolute between M and some larger structure N if it holds in M precisely if it holds in N , where parameters are allowed to take values in the smaller structure M . If the truth of the property in the smaller structure implies the truth in the larger structure, we say the property is *upwards absolute*; conversely, if truth in the larger structure implies truth in the smaller one, we say the property is *downwards absolute*. The theory of absoluteness concerns the following classes of formulas, among others.

- (i) Δ_0 formulas, in which only bounded quantifiers are permitted, for example in ZFC, ‘ x is an ordinal’, ‘ f is a function’, ‘ x is a subset of y ’, ‘ x is ω ’.
- (ii) Σ_1 formulas, which are Δ_0 formulas surrounded by a single existential quantifier.
- (iii) Π_1 formulas, which are Δ_0 formulas surrounded by a single universal quantifier, for example ‘ x is a cardinal’ or ‘ x is the power set of y ’.

One can show that Δ_0 formulas are absolute between transitive models. Further, Σ_1 formulas are upwards absolute and Π_1 formulas are downwards absolute. The example above shows that ‘ x is a cardinal’ cannot be Δ_0 as it is not upwards absolute. Similarly, ‘ x is the power set of y ’ cannot be Δ_0 , because the object p that M believes is the power set of ω must be countable, and so cannot be the real power set in V_κ . As being a subset is absolute, this object p must consist of subsets of ω , but must only contain very few of them.

As being ω is Δ_0 , in fact all arithmetical statements (and therefore, by encoding, all syntactic statements) are Δ_0 .

Theorem. $IC \rightarrow \text{Con}(\text{ZFC})$ but $\text{Con}(\text{ZFC}) \nrightarrow IC$.

Proof. The forward direction has already been proven. Since IC proves the consistency of ZFC, there is a countable transitive model $M \subseteq V_\kappa \subseteq V$ of ZFC. By absoluteness, $M \models \text{Con}(\text{ZFC})$, so $M \models \text{ZFC}^*$ where we define $\text{ZFC}^* = \text{ZFC} + \text{Con}(\text{ZFC})$. We have thus proven that IC implies the consistency of ZFC^* . So, by the second incompleteness theorem, $\text{ZFC}^* \not\models IC$. \square

1.5. Worldly cardinals

We now show that if $V_\kappa \models \text{ZFC}$, it is not necessarily the case that κ is inaccessible.

Observe that $M \neq V_\alpha$ for any α . Clearly $M \neq V_\omega$. But $|V_{\omega+1}| = |\mathcal{P}(\omega)| = 2^{\aleph_0}$, and $|V_\alpha| > 2^{\aleph_0}$ for all $\alpha > \omega + 1$. But M is countable, so it cannot be any of these.

Recall the definition of N by

$$N_0 = \emptyset; \quad N_{k+1} = W(N_k); \quad N = \bigcup_{k \in \mathbb{N}} N_k$$

We wish to create a similar structure that is of the form V_α for some α . We define

$$\alpha_0 = 0; \quad \alpha_{k+1} = \sup \{ \text{rank}(x) \mid x \in W(V_{\alpha_k}) \}; \quad \alpha = \sup \{ \alpha_n \mid n \in \mathbb{N} \}$$

Note that $N \subseteq V_{\alpha_1}$.

Theorem. $V_\alpha \leq V_\kappa$ and $\alpha < \kappa$.

Proof. The first statement follows from the Tarski–Vaught test. To show $\alpha < \kappa$, we first show by induction that $\alpha_k < \kappa$. This is clearly true for $k = 0$. Now, if $\alpha_k < \kappa$, we have $|V_{\alpha_k}| < \kappa$ by a previous lemma. Thus,

$$|W(V_{\alpha_k})| \leq \aleph_0 \cdot |V_{\alpha_k}^{<\omega}| = |V_{\alpha_k}| < \kappa$$

where $X^{<\omega}$ is the set of finite sequences of elements of X . Hence $\{\text{rank}(x) \mid x \in W(V_{\alpha_k})\}$ is a set of less than κ ordinals less than κ , so it must be bounded by regularity. Finally, as α is a countable union of the α_k , regularity again shows $\alpha < \kappa$. \square

Remark. The ordinal α produced in this way has countable cofinality, so cannot be inaccessible. In particular, $V_\alpha \models \text{ZFC}$ but α is not inaccessible.

Definition. We call an ordinal α *worldly* if $V_\alpha \models \text{ZFC}$, and write $\text{Wor}(\alpha)$.

We have shown $\text{I}(\kappa) \rightarrow \text{Wor}(\kappa)$, but not the other way round given that a worldly cardinal exists. In particular,

$$\text{IC} \rightarrow \text{WorC} \rightarrow \text{Con}(\text{ZFC})$$

Theorem. If κ is a worldly ordinal, κ is a cardinal.

Proof. First, observe that κ is a limit ordinal; otherwise, its predecessor would be the largest ordinal in the model, but ZFC proves that there is no largest ordinal. Suppose κ is not a cardinal, so there is $\lambda < \kappa$ such that there is a bijection $\lambda \rightarrow \kappa$. In particular, $\lambda < \kappa < \lambda^+$. By the proof of Hartogs' lemma, there is a relation $R \subseteq \lambda \times \lambda$ such that $(\lambda, R) \cong (\kappa, \in)$. Assuming Kuratowski's definition of ordered pairs, an element of $\lambda \times \lambda$ is an element of V_λ , so $\lambda \times \lambda \in V_{\lambda+1}$ and $R \in V_{\lambda+1}$. The pair (λ, R) is an element of $V_{\lambda+3} \subseteq V_\kappa$. Thus V_κ contains a well-order (λ, R) of order type κ . But ZFC proves that every well-ordering is isomorphic to a unique ordinal, so we must have $\kappa \in V_\kappa$, which is a contradiction. \square

1.6. The consistency strength hierarchy

Let B be a base theory; we will often use ZFC. If T, S are extensions of B , we say that T has lower *consistency strength* than S , written $T \leq_{\text{Con}} S$, if $B \vdash \text{Con}(S) \rightarrow \text{Con}(T)$. We say that T and S is *equiconsistent*, written $T \equiv_{\text{Con}} S$, if $T \leq_{\text{Con}} S$ and $S \leq_{\text{Con}} T$, and write $T <_{\text{Con}} S$ if $T \leq_{\text{Con}} S$ but $S \not\leq_{\text{Con}} T$.

Remark. (i) If I is inconsistent, then $T \leq_{\text{Con}} I$ for all T . All inconsistent theories are equiconsistent. In particular, T is consistent if and only if $T <_{\text{Con}} I$. We typically write \perp for an inconsistent theory.

VI. Large Cardinals

- (ii) $<_{\text{Con}}$ is more than just ‘proving more theorems’. If φ is such that $\text{ZFC} \not\vdash \varphi$ and $\text{ZFC} \not\vdash \neg\varphi$, it is not necessarily the case that $\text{ZFC} <_{\text{Con}} \text{ZFC} + \varphi$ or $\text{ZFC} <_{\text{Con}} \text{ZFC} + \neg\varphi$. For example, $\text{ZFC} + \text{CH}$, $\text{ZFC} + \neg\text{CH}$, and ZFC are all equiconsistent.
- (iii) The second incompleteness theorem shows, for suitably nice theories T , that if $T \neq \perp$ then $T <_{\text{Con}} T + \text{Con}(T)$. Note that it is possible that T is consistent but $T + \text{Con}(T)$ is inconsistent, so the incompleteness theorem does not necessarily give an infinite chain of strict consistency strength inequalities. For example, consider

$$\text{ZFC}^\dagger = \text{ZFC} + \neg \text{Con}(\text{ZFC})$$

Since $\text{ZFC}^\dagger \supseteq \text{ZFC}$, we must have $\text{Con}(\text{ZFC}^\dagger) \rightarrow \text{Con}(\text{ZFC})$, but $\text{ZFC}^\dagger \rightarrow \neg \text{Con}(\text{ZFC})$, so $\text{ZFC}^\dagger + \text{Con}(\text{ZFC}^\dagger)$ is inconsistent.

In conclusion,

$$\text{ZFC} <_{\text{Con}} \text{ZFC} + \text{Con}(\text{ZFC}) <_{\text{Con}} \text{ZFC} + \text{WorC} <_{\text{Con}} \text{ZFC} + \text{IC}$$

where the second inequality uses the same argument as $\text{IC} \rightarrow \text{Con}(\text{ZFC} + \text{Con}(\text{ZFC}))$.

We will see that $\text{ZFC} \equiv_{\text{Con}} \text{ZFC} + \neg \text{IC}$. Many large cardinal axioms have this property that their negations are weak.

If κ is the least inaccessible cardinal, then V_κ is a model of ZFC , but we can show that it cannot satisfy IC . Note that the statement ‘ λ is inaccessible’ is a Π_1 statement, so is downwards absolute. Given a model with two inaccessible cardinals $\kappa_0 < \kappa_1$, we have $V_{\kappa_1} \models \text{ZFC} + \text{I}(\kappa_0)$ so in particular, $V_{\kappa_1} \models \text{ZFC} + \text{IC}$.

Lemma. If α is a limit ordinal, then the formula ‘ λ is inaccessible’ is absolute for V_α and V . In particular, V_κ above does not satisfy IC .

Proof. By downwards absoluteness, it suffices to show that if $V_\alpha \models \text{I}(\lambda)$ then $\text{I}(\lambda)$. Suppose not, so λ is singular or not a strong limit.

Let λ be singular, so there is a cofinal set $C \subseteq \lambda$ with $|C| = \gamma < \lambda$, so there is a bijection $f : \gamma \rightarrow C$. Note that being singular is Σ_1 , witnessed by C, γ, f . We have $C \in V_{\lambda+1}$, $\gamma \in V_\lambda$, and $f \in V_{\lambda+2}$. All of these are subsets of V_α , so these witnesses exist in V_α . Hence V_α believes that C is a cofinal set of cardinality less than λ , so it believes λ is singular, contradicting inaccessibility.

Now let λ not be a strong limit. Let $\gamma < \lambda$, and let $f : \mathcal{P}(\gamma) \rightarrow \lambda$ be a surjection. Then $\mathcal{P}(\gamma) \in V_{\gamma+2} \subseteq V_\lambda \subseteq V_\alpha$, and so this function is an element of $V_{\lambda+2} \subseteq V_\alpha$. The statement that it is a surjection is absolute, so V_α believes f is a surjection from $\mathcal{P}(\gamma)$ to λ , contradicting its belief that λ is a strong limit. \square

Therefore, we have the following.

Theorem. Suppose $\text{ZFC} + \text{IC}$, and let κ be the least inaccessible. Then $V_\kappa \models \text{ZFC} + \neg \text{IC}$.

Proof. Suppose $V_\kappa \models \text{ZFC} + \text{IC}$. Then there is $\lambda < \kappa$ such that $V_\kappa \models \text{I}(\lambda)$, but by the previous lemma this contradicts minimality of κ . \square

Therefore, we have the following.

$$\text{ZFC} + \text{IC} \vdash \text{there is a transitive model of } \text{ZFC} + \neg \text{IC}$$

For any theory T , we write

$$T^* = T + \text{Con}(T)$$

We make the following remarks.

- (i) Observe that if S proves that there is a transitive model of T , then $S \vdash \text{Con}(T^*)$ because consistency statements are downwards absolute between transitive models.
- (ii) Note also that if S proves every axiom of T , then $\text{Con}(S) \rightarrow \text{Con}(T)$.
- (iii) If T is not equiconsistent with \perp , then $\text{Con}(T) \nleftrightarrow \text{Con}(T^*)$.

We can therefore show

$$\text{Con}(\text{ZFC} + \neg \text{IC}) \nleftrightarrow \text{Con}(\text{ZFC} + \text{IC})$$

assuming that $\text{ZFC} + \neg \text{IC}$ is consistent. We have that $\text{ZFC} + \text{IC}$ yields a transitive model of $\text{ZFC} + \neg \text{IC}$. Thus, by (i), $\text{ZFC} + \text{IC}$ implies $\text{Con}((\text{ZFC} + \neg \text{IC})^*)$. Hence $\text{Con}(\text{ZFC} + \text{IC}) \rightarrow \text{Con}((\text{ZFC} + \neg \text{IC})^*)$, so if the given implication were to hold, it would contradict Gödel's second incompleteness theorem. Thus, if $\text{ZFC} + \neg \text{IC}$ is consistent,

$$\text{ZFC} + \neg \text{IC} <_{\text{Con}} \text{ZFC} + \text{IC}$$

Observe that none of the proofs given in this section work for weakly inaccessible cardinals, so it is not clear that weakly inaccessible cardinals qualify as large cardinals. However, under the generalised continuum hypothesis, we have $\aleph_\alpha = \beth_\alpha$ and so the notions of weakly inaccessible cardinal and inaccessible cardinal coincide. In Part III Forcing and the Continuum Hypothesis, we see that if $M \models \text{ZFC}$, there is $L \subseteq M$ such that L is transitive in M , L contains all the ordinals of M , and $L \models \text{ZFC} + \text{GCH}$. Thus, given a model $M \models \text{ZFC} + \text{WIC}$, we obtain $L \models \text{ZFC} + \text{IC}$, and thus the two axioms WIC and IC are equiconsistent.

Note that 2^{\aleph_0} is not a strong limit, but it is consistent that 2^{\aleph_0} is weakly inaccessible (under suitable assumptions), so the notions of weakly inaccessible cardinals and inaccessible cardinals do not coincide.

2. Measurability and compactness

2.1. The measure problem

Let \mathbb{I} denote the unit interval $[0, 1] \subseteq \mathbb{R}$. A function $\mu : \mathcal{P}(\mathbb{I}) \rightarrow \mathbb{I}$ is called a *measure* if

- (i) $\mu(\mathbb{I}) = 1$ and $\mu(\emptyset) = 0$;
- (ii) (translation invariance) if $X \subseteq \mathbb{I}$, $r \in \mathbb{R}$, and $X + r = \{x + r \mid x \in X\} \subseteq \mathbb{I}$, then $\mu(X) = \mu(X + r)$; and
- (iii) (countable additivity) if $(A_n)_{n \in \mathbb{N}}$ is a family of pairwise disjoint subsets of \mathbb{I} , then $\mu\left(\bigcup_{n \in \mathbb{N}} A_n\right) = \sum_{n \in \mathbb{N}} \mu(A_n)$.

The *Lebesgue measure problem* was the question of whether such a measure function exists. Vitali proved that a measure cannot be defined on all of $\mathcal{P}(\mathbb{I})$. This proof requires the axiom of choice nontrivially. In 1970, Solovay proved that if $\text{ZFC} + \text{IC}$ is consistent, then, there is a model of ZF in which all sets are Lebesgue measurable. In 1984, Shelah showed that the inaccessible cardinal was necessary to construct this model.

Now, replace translation invariance with the requirement that for all $x \in \mathbb{I}$, we have $\mu(\{x\}) = 0$, and call such measures *Banach measures*. *Banach's measure problem* was the question of whether a Banach measure exists. Note that every Lebesgue measure is a Banach measure. If $\mu(\{x\}) > 0$ for some x , then by translation invariance, every singleton has the same measure $\mu(\{x\}) > 0$. There is some natural number n such that $n\mu(\{x\}) > 1$, but this contradicts countable additivity using a set of n reals. Observe that for any $\varepsilon > 0$, there can be only finitely many pairwise disjoint sets with measure at least ε .

Banach and Kuratowski proved in 1929 that the continuum hypothesis implies that there are no Banach measures on \mathbb{I} . We can define Banach measures on any set S by also replacing property (i) with the requirement that $\mu(S) = 1$ and $\mu(\emptyset) = 0$. Note that if $|S| = |S'|$, then there is a Banach measure on S if and only if there is one on S' . Thus, having a Banach measure is a property of cardinals.

For larger cardinals, it may not be natural to just consider countable additivity.

Definition. A Banach measure μ is called λ -*additive* if for all $\gamma < \lambda$ and pairwise disjoint families $\{A_\alpha \mid \alpha < \gamma\}$, then

$$\mu\left(\bigcup A_\alpha\right) = \sup \left\{ \sum_{\alpha \in F} \mu(A_\alpha) \mid F \subseteq \gamma \text{ finite} \right\}$$

Theorem. If κ is the smallest cardinal that has a Banach measure, then that measure is κ -additive.

2.2. Real-valued measurable cardinals

Definition. A cardinal κ is *real-valued measurable*, written $\text{RVM}(\kappa)$, if there is a κ -additive Banach measure on κ .

Proposition. Every real-valued measurable cardinal is regular.

Proof. Suppose that κ is a real-valued measurable cardinal, and that $C \subseteq \kappa$ is cofinal with $|C| = \lambda < \kappa$. We can write

$$C = \{\gamma_\alpha \mid \alpha < \lambda\}$$

where γ_α is increasing in α . Consider

$$C_\alpha = \{\xi \mid \gamma_\alpha \leq \xi < \gamma_{\alpha+1}\}$$

Then $\bigcup_{\alpha < \lambda} C_\alpha = \kappa$ as C is cofinal, and the C_α are disjoint. Note that $|C_\alpha| \leq |\gamma_{\alpha+1}| < \kappa$. Writing $C_\alpha = \bigcup_{x \in C_\alpha} \{x\}$, we observe by κ -additivity that $\mu(C_\alpha) = 0$. But again by κ -additivity, $\mu(\kappa) = 0$, contradicting property (i). \square

Proposition (the pigeonhole principle). Let κ be regular, $\lambda < \kappa$, and $f : \kappa \rightarrow \lambda$. Then there is some $\alpha \in \lambda$ such that $|f^{-1}(\alpha)| = \kappa$.

Proof. We have

$$\kappa = \bigcup_{\alpha \in \lambda} f^{-1}(\alpha)$$

giving the result immediately by regularity of κ . \square

Proposition. All successor cardinals are regular.

Proposition. If μ is a Banach measure on S , and C is a family of pairwise disjoint sets of positive μ -measure, then C is countable.

Proof. Consider the collection

$$C_n = \left\{ A \in C \mid \mu(A) > \frac{1}{n} \right\}$$

Observe that each C_n is finite, so $C = \bigcup_{n \in \mathbb{N}} C_n$ must be countable. \square

Lemma (Ulam). For any cardinal λ , there is an *Ulam matrix* A_α^ξ indexed by $\alpha < \lambda^+$ and $\xi < \lambda$ such that

(i) for a given ξ , the set $\{A_\alpha^\xi \mid \alpha < \lambda^+\}$ is a pairwise disjoint family; and

(ii) for a given α , we have

$$\left| \lambda^+ \setminus \bigcup_{\xi < \lambda} A_\alpha^\xi \right| \leq \lambda$$

VI. Large Cardinals

Proof. For each $\gamma < \lambda^+$, fix a surjection $f_\gamma : \lambda \rightarrow \gamma + 1$. Define

$$A_\alpha^\xi = \{\gamma \mid f_\gamma(\xi) = \alpha\}$$

It is clear that property (i) holds. For property (ii), suppose

$$\gamma \in \lambda^+ \setminus \bigcup_{\xi < \lambda} A_\alpha^\xi$$

Then $\gamma < \lambda^+$ and for all ξ , we have $f_\gamma(\xi) \neq \alpha$. Hence

$$\lambda^+ \setminus \bigcup_{\xi < \lambda} A_\alpha^\xi \subseteq \alpha$$

so the size of this set is at most λ . □

Theorem. Every real-valued measurable cardinal is weakly inaccessible.

Remark. If there is a Banach measure on $[0, 1]$, then in particular 2^{\aleph_0} is weakly inaccessible.

Proof. We have already shown regularity. Suppose κ is not a limit cardinal, so $\kappa = \lambda^+$. Let $(A_\alpha^\xi)_{\alpha < \lambda^+; \xi < \lambda}$ be an Ulam matrix for λ . By (ii),

$$|Z_\alpha| \leq \lambda; \quad Z_\alpha = \lambda^+ \setminus \bigcup_{\xi < \lambda} A_\alpha^\xi$$

so by κ -additivity, $\mu(Z) = 0$. Hence

$$\mu\left(\bigcup_{\xi < \lambda} A_\alpha^\xi\right) = 1$$

This is a small union of sets of measure 1, so again by κ -additivity there is some ξ_α such that $\mu(A_\alpha^{\xi_\alpha}) > 0$. Let $f : \lambda^+ \rightarrow \lambda$ be the map $\alpha \mapsto \xi_\alpha$. By the pigeonhole principle, there is some ξ and a set $A \subseteq \lambda^+$ with $|A| = \lambda^+$ such that for all $\alpha \in A$, we have $\xi_\alpha = \xi$. By property (i), the collection $\{A_\alpha^\xi \mid \alpha \in A\}$ is a collection of uncountable size λ^+ of pairwise disjoint sets, all of which have positive measure, but we have already shown that such a collection must be countable. □

2.3. Measurable cardinals

Definition. A Banach measure μ is called *two-valued* if μ takes values in $\{0, 1\}$.

This removes any mention of the real numbers from the definition of a Banach measure.

Remark. Two-valued measures correspond directly to ultrafilters. Recall that F is a *filter* on S if

- (i) $\emptyset \notin F, S \in F$;
- (ii) if $A \subseteq B$ then $A \in F \rightarrow B \in F$;
- (iii) if $A, B \in F$ then $A \cap B \in F$.

We say that F is an *ultrafilter* if $A \in F$ or $S \setminus A \in F$ for all $A \subseteq S$. F is *nonprincipal* if for all $x \in S$, the singleton $\{x\}$ is not in F . An ultrafilter is λ -complete if for all $\gamma < \lambda$ and all families $\{A_\alpha \mid \alpha < \gamma\} \subseteq F$, we have $\bigcap_{\alpha < \gamma} A_\alpha \in F$. In this way, the collection of sets of a two-valued Banach measure μ that are assigned measure 1 form a nonprincipal ultrafilter. This filter is λ -complete if and only if μ is λ -additive.

Definition. An uncountable cardinal κ is *measurable*, written $M(\kappa)$, if there is a κ -complete nonprincipal ultrafilter on κ .

- Remark.* (i) ZFC proves that there is an \aleph_0 -complete nonprincipal ultrafilter on \aleph_0 , because \aleph_0 -completeness is equivalent to closure under finite intersections, which is trivial.
- (ii) A cardinal κ is called *Ulam measurable* if there is an \aleph_1 -complete nonprincipal ultrafilter on κ . With this definition, the least Ulam measurable cardinal is measurable. So the existence of an Ulam measurable cardinal is equivalent to the existence of a measurable cardinal.
- (iii) The theories ZFC + MC and ZFC + RVMC are equiconsistent. This can be shown analogously to inaccessible and weakly inaccessible cardinals, this time using a variant of Gödel's constructible universe.

Theorem. Every measurable cardinal is inaccessible.

Proof. We have already shown regularity in the real-valued measurable cardinal case. Let κ be measurable with ultrafilter U . Suppose it is not a strong limit, so there is $\lambda < \kappa$ such that $2^\lambda \geq \kappa$. Then there is an injection $f : \kappa \rightarrow B_\lambda$, where B_λ is the set of functions $\lambda \rightarrow 2$. Fix some $\alpha < \lambda$, then for each $\gamma < \kappa$, either

$$f(\gamma)(\alpha) = 0 \text{ or } f(\gamma)(\alpha) = 1$$

Let

$$A_0^\alpha = \{\gamma \mid f(\gamma)(\alpha) = 0\}; \quad A_1^\alpha = \{\gamma \mid f(\gamma)(\alpha) = 1\}$$

These two sets are disjoint and have union κ . So there is exactly one number $b \in \{0, 1\}$ such that $A_b^\alpha \in U$. Define $c \in B_\lambda$ by $c(\alpha) = b$. Then

$$X_\alpha = A_{c(\alpha)}^\alpha \in U$$

This is a collection of λ -many sets that are all in U , so by κ -completeness, their intersection $\bigcap_{\alpha < \lambda} X_\alpha$ also lies in U . Suppose $\gamma \in \bigcap_{\alpha < \lambda} X_\alpha$, so for all $\alpha < \lambda$, we have $\gamma \in A_{c(\alpha)}^\alpha$. Equivalently, for all $\alpha < \lambda$, we have $f(\gamma)(\alpha) = c(\alpha)$. So γ lies in this intersection if and only if $f(\gamma)$

VI. Large Cardinals

is precisely the function c . Hence

$$\bigcap_{\alpha < \lambda} X_\alpha \subseteq \{f^{-1}(c)\}$$

So this intersection has either zero or one element, and in particular, it is not in the ultrafilter, giving a contradiction. \square

Nonprincipal ultrafilters on κ are not κ^+ -complete, because κ itself is a union of κ -many singletons. Principal ultrafilters are complete for any cardinal. However, we can emulate completeness for nonprincipal ultrafilters at the cardinal κ^+ using the following method. If $(A_\alpha)_{\alpha \leq \kappa}$ is a sequence of subsets of κ , its *diagonal intersection* is

$$\Delta_{\alpha \leq \kappa} A_\alpha = \left\{ \xi \in \kappa \mid \xi \in \bigcap_{\alpha < \xi} A_\alpha \right\}$$

A filter on κ is called *normal* if it is closed under diagonal intersections.

Theorem. If κ is measurable, then there is a κ -complete normal nonprincipal ultrafilter on κ .

The proof will be given later, and is also on an example sheet.

2.4. Weakly compact cardinals

Let $[X]^n$ be the set of n -element subsets of X . A *2-colouring* of \mathbb{N} is a map $c : [\mathbb{N}]^2 \rightarrow \{\text{red}, \text{blue}\}$. Ramsey's theorem states that for each 2-colouring c , there is an infinite subset $X \subseteq \mathbb{N}$ such that $c|_{[X]^2}$ is *monochromatic* (or *homogeneous*): each 2-element subset is given the same colour under c .

This property is invariant under bijection, so this is really a property of the cardinal \aleph_0 . In *Erdős' arrow notation*, we write

$$\kappa \rightarrow (\lambda)_m^n$$

if for every colouring $c : [\kappa]^n \rightarrow m$, there is a monochromatic subset $X \subseteq \kappa$ of size λ :

$$|c[[X]^n]| = 1$$

In this notation, Ramsey's theorem becomes the statement

$$\aleph_0 \rightarrow (\aleph_0)_2^2$$

We can now make the following definition.

Definition. An uncountable cardinal κ is called *weakly compact*, written $W(\kappa)$, if $\kappa \rightarrow (\kappa)_2^2$.

The name will be explained later.

Theorem (Erdős). Every weakly compact cardinal is inaccessible.

2. Measurability and compactness

Proof. Suppose κ is weakly compact but not regular. Then $\kappa = \bigcup_{\alpha < \lambda} X_\alpha$ for $\alpha < \kappa$ and disjoint sets X_α with $|X_\alpha| < \kappa$. We define a colouring c as follows. A pair $\{\gamma, \delta\}$ is red if γ, δ lie in the same X_α , and blue if they are in different X_α . Let $H \subseteq \kappa$ be a monochromatic subset of size κ for c . If H is red, then one of the X_α is large, which is a contradiction. But if H is blue, then λ must be large, which also gives a contradiction.

Suppose that κ is not a strong limit, so $2^\lambda \geq \kappa$ for $\lambda < \kappa$. Let B_λ be the set of functions $\lambda \rightarrow 2$, and give it the *lexicographic order*: we say that $f < g$ if $f(\alpha) < g(\alpha)$ at the first position α at which f and g disagree. For this proof, we will use the combinatorial fact that this ordered structure $(B_\lambda, \leq_{\text{lex}})$ is a totally ordered set with no increasing or decreasing chains of length $\kappa > \lambda$. The proof is on an example sheet.

If $2^\lambda \geq \kappa$, there is a family of pairwise distinct elements $(f_\alpha)_{\alpha < \kappa}$ of B_λ of length κ . Define a colouring c of κ as follows. A pair α, β is red if the truth value of $\alpha \leq \beta$ is the same as the truth value of $f_\alpha \leq_{\text{lex}} f_\beta$. A pair is blue otherwise. Let H be a monochromatic set for c . If H is red, then f_α forms a \leq_{lex} -increasing sequence of length κ . If H is blue, then f_α forms a \leq_{lex} -decreasing sequence of length κ . Both results contradict the combinatorial result above. \square

Theorem. Every measurable cardinal is weakly compact.

Proof. Let $f : [\kappa]^2 \rightarrow 2$ be a colouring of a measurable cardinal κ . Let

$$X_0^\alpha = \{\beta \mid f(\{\alpha, \beta\}) = 0\}; \quad X_1^\alpha = \{\beta \mid f(\{\alpha, \beta\}) = 1\}$$

For a given α , these are disjoint, and $X_0^\alpha \cup X_1^\alpha = \kappa \setminus \{\alpha\}$, so precisely one of them lies in the ultrafilter U . Define $c : \kappa \rightarrow 2$ be such that $X_{c(\alpha)}^\alpha \in U$. Now, let

$$X_0 = \{\alpha \mid c(\alpha) = 0\}; \quad X_1 = \{\alpha \mid c(\alpha) = 1\}$$

Precisely one of these two sets lies in U .

We claim that if $X_i \in U$, then there is a monochromatic set H for colour i with $|H| = \kappa$. Without loss of generality, we may assume $i = 0$. Define

$$Z_\alpha = \begin{cases} X_0^\alpha & \text{if } c(\alpha) = 0 \\ \kappa & \text{if } c(\alpha) = 1 \end{cases}$$

Each of the Z_α lie in the ultrafilter U . As we may assume U is normal, the diagonal intersection of the Z_α also lies in U . So we can define

$$H = X_0 \cap \bigtriangleup_{\alpha < \kappa} Z_\alpha \in U$$

and $|H| = \kappa$. Let $\gamma < \delta$ with $\gamma, \delta \in H$. Then $\gamma, \delta \in X_0$, so $c(\gamma) = 0 = c(\delta)$. Hence $Z_\gamma = X_0^\gamma$ and $Z_\delta = X_0^\delta$. In particular,

$$\delta \in \bigtriangleup_{\alpha < \kappa} Z_\alpha \subseteq \bigcap_{\xi < \delta} Z_\xi \subseteq Z_\gamma = X_0^\gamma$$

Hence $f(\{\gamma, \delta\}) = 0$. \square

VI. Large Cardinals

The large cardinal axioms discussed so far fall into a linear hierarchy of consistency strength. This is known as the *linearity phenomenon*.

2.5. Strongly compact cardinals

The compactness theorem for first-order logic says that for any first-order language L_S and set of axioms $\Phi \subseteq L_S$,

$$\Phi \text{ is satisfiable} \leftrightarrow (\forall \Phi_0 \subseteq \Phi. |\Phi_0| < \aleph_0 \rightarrow \Phi_0 \text{ is satisfiable})$$

This result cannot work for languages with infinitary conjunctions and disjunctions. Indeed, if we write

$$\varphi_F \equiv \bigvee_{i \in \mathbb{N}} \varphi_{=n}; \quad \varphi_{=n} \equiv \text{there are precisely } n \text{ elements}; \quad \varphi_{\geq n} \equiv \text{there are at least } n \text{ elements}$$

then

$$\{\varphi_{\geq n} \mid n \in \mathbb{N}\} \cup \{\varphi_F\}$$

is finitely satisfiable but not satisfiable.

Definition. An $\mathcal{L}_{\kappa\kappa}$ -language is defined by

- a set of variables;
- a set S of function, relation, and constant symbols of finite arity;
- the logical symbols $\wedge, \vee, \neg, \exists, \forall$; and
- the infinitary logical symbols $\bigwedge_{\alpha < \lambda}, \bigvee_{\alpha < \lambda}, \exists^\lambda, \forall^\lambda$ for $\lambda < \kappa$.

We define the new syntactic rules as follows. If φ_α are L_S -formulas for $\alpha < \lambda$, then so are $\bigwedge_{\alpha < \lambda} \varphi_\alpha$ and $\bigvee_{\alpha < \lambda} \varphi_\alpha$. If \mathbf{v} is a sequence of variables of length λ and φ is an L_S -formula, then $\exists^\lambda \mathbf{v}. \varphi$ and $\forall^\lambda \mathbf{v}. \varphi$ are L_S -formulas.

We say that M is a model of $\bigvee_{\alpha < \lambda} \varphi_\alpha$ if $M \models \varphi_\alpha$ for all $\alpha < \lambda$. Similarly, M models $\exists^\lambda \mathbf{v}. \varphi$ if there is a function $a : \lambda \rightarrow M$ such that

$$M \models \varphi \left[\frac{a(0)a(1) \dots a(\xi) \dots}{v_0 v_1 \dots v_\xi \dots} \right]$$

Definition. An $\mathcal{L}_{\kappa\kappa}$ -language L_S satisfies compactness if for all $\Phi \subseteq L_S$,

$$\Phi \text{ is satisfiable} \leftrightarrow (\forall \Phi_0 \subseteq \Phi. |\Phi_0| < \kappa \rightarrow \Phi_0 \text{ is satisfiable})$$

Note that if $\kappa = \omega$, we recover the standard notion of a first-order language, so all $\mathcal{L}_{\omega\omega}$ -languages satisfy compactness.

Definition. An uncountable cardinal κ is called *strongly compact*, denoted $\text{SC}(\kappa)$, if every $\mathcal{L}_{\kappa\kappa}$ -language satisfies compactness.

2. Measurability and compactness

Theorem (Keisler–Tarski theorem). Suppose κ is a strongly compact cardinal. Then every κ -complete filter on κ can be extended to a κ -complete ultrafilter.

Proof. We define a language L extending the usual language of set theory by creating a constant symbol c_A for each $A \subseteq \kappa$, giving 2^κ -many symbols. Now let L^* be L with an extra constant symbol c . Let

$$M = (\mathcal{P}(\kappa), \in, \{A \mid A \subseteq \kappa\})$$

so c_A is interpreted by A . Let $\Phi = \text{Th}_L(M)$ be the L -theory of M . In particular,

$$M \models \forall x. x \in c_A \rightarrow x \text{ is an ordinal}$$

and

$$M \models \forall x. x \text{ is an ordinal} \rightarrow x \in c_A \vee x \in c_{\kappa \setminus A}$$

Now let

$$\Phi^* = \Phi \cup \{c \in c_A \mid A \in F\}$$

This is a subset of L^* . We show that Φ^* is κ -satisfiable. If $(A_\alpha)_{\alpha < \lambda}$ are subsets of κ such that $c \in c_{A_\alpha}$ occurs in a κ -small subset of Φ^* , then any element $\eta \in \bigcap_{\alpha < \lambda} A_\alpha$ can be chosen as the interpretation of c . As F is κ -complete, this intersection lies in F and so is nonempty as required.

Hence, by strong compactness of κ , the theory Φ^* is satisfiable. Let M be a model of Φ^* . Define

$$U = \{A \mid M \models c \in c_A\}$$

We claim that this is a κ -complete ultrafilter extending F . The fact that U extends F holds by construction of Φ^* . It is an ultrafilter because M believes that $c \in c_A$ or $c \in c_{\kappa \setminus A}$. It is κ -complete because if $\{A_\alpha \mid \alpha < \lambda\} \subseteq U$, let $A = \bigcap_{\alpha < \lambda} A_\alpha$, then

$$M \models \forall x. \left(x \in c_A \leftrightarrow \bigwedge_{\alpha < \lambda} x \in c_{A_\alpha} \right)$$

As this holds in particular for c , we obtain $A \in U$. □

Corollary. Every strongly compact cardinal is measurable.

Proof. Let

$$F = \{A \subseteq \kappa \mid |\kappa \setminus A| < \kappa\}$$

In the case $\kappa = \omega$, this is known as the Fréchet filter. This is a κ -complete filter on κ . If U extends F then U must be nonprincipal, so by the Keisler–Tarski theorem, F can be extended to a κ -complete nonprincipal ultrafilter on κ as required. □

3. Reflection

3.1. The Keisler extension property

Definition. A cardinal κ has the *Keisler extension property*, written $\text{KEP}(\kappa)$, if there is $\kappa \in X \supsetneq V_\kappa$ transitive such that $V_\kappa \leq X$.

Proposition. If κ is inaccessible and satisfies the Keisler extension property, there is an inaccessible cardinal $\lambda < \kappa$.

Proof. Fix X as in the Keisler extension property. As κ is inaccessible, $X \models I(\kappa)$ because $\kappa \in X$ and inaccessibility is downwards absolute for transitive models. Also, $V_\kappa \models \text{ZFC}$, so $X \models \text{ZFC}$ as it is an elementary superstructure. Therefore, $X \models \text{ZFC} + \text{IC}$, so $V_\kappa \models \text{ZFC} + \text{IC}$. So as inaccessibility is absolute between V_κ and V , there is an inaccessible $\lambda < \kappa$. \square

The phenomenon that properties of X occur below κ is called *reflection*. This argument can be improved in the following sense. For a given $\alpha < \kappa$,

$$X \models \exists \lambda > \alpha. I(\lambda)$$

But as $\alpha \in V_\kappa$, elementarity gives

$$V_\kappa \models \exists \lambda > \alpha. I(\lambda)$$

So the set

$$\{\lambda < \kappa \mid I(\lambda)\}$$

is not only nonempty, but cofinal in κ .

Corollary. Let A be the axiom

$$\exists \kappa. I(\kappa) \wedge \text{KEP}(\kappa)$$

Then

$$\text{ZFC} + \text{IC} <_{\text{Con}} \text{ZFC} + A$$

Proof. It suffices to show that $\text{ZFC} + A \models \text{Con}(\text{ZFC} + \text{IC})$. We have seen that $\text{ZFC} + A$ proves the existence of (at least) two inaccessible cardinals below κ , and in particular the larger of the two is a model of $\text{ZFC} + \text{IC}$. \square

Remark. This is the main technique for proving strict inequalities of consistency strength. Given two large cardinal properties Φ, Ψ with the appropriate amount of absoluteness properties, we show that $\text{ZFC} + \Phi(\kappa)$ proves that the set

$$\{\lambda < \kappa \mid \Psi(\lambda)\}$$

is cofinal in κ . Then $\text{ZFC} + \Phi \models \text{Con}(\text{ZFC} + \Psi)$.

Example. Consider the proof that every inaccessible cardinal has a worldly cardinal below it. In the construction, we produce a sequence of ordinals $(\alpha_i)_{i \in \omega}$, and the worldly cardinal is $\sup \alpha_i$. But we can set $\alpha_0 = \lambda + 1$ for a given worldly cardinal $\lambda < \kappa$, so this gives a cofinal sequence of worldly cardinals below every given inaccessible.

Theorem. Every strongly compact cardinal has the Keisler extension property.

Proof. We want to use the method of (elementary) diagrams to produce a model with V_κ as a substructure. However, we have no way to control whether such a model is well-founded using standard first-order model-theoretic techniques. To bypass this issue, we will use infinitary operators.

Let c_x be a constant symbol for each $x \in V_\kappa$, and let L be the language with \in and the c_x . Let

$$\mathcal{V} = (V_\kappa, \in, \{x \mid x \in V_\kappa\})$$

In first-order logic, $\text{Th}(X)$ is the elementary diagram of V_κ , so if $M \models \text{Th}(X)$, then $V_\kappa \subseteq M$. Let L_κ be the $\mathcal{L}_{\kappa\kappa}$ -language with the same symbols. Consider

$$\psi \equiv \forall^\omega \mathbf{v}. \bigvee_{i \in \omega} v_{i+1} \notin v_i$$

This expresses well-foundedness (assuming AC). Writing $\Phi = \text{Th}_{L_\kappa}(\mathcal{V})$ for the L_κ -theory of \mathcal{V} , we must have $\psi \in \Phi$ since V_κ is well-founded. Thus, if $M \models \Phi$, then M is a well-founded model containing V_κ . By taking the Mostowski collapse, we may also assume that any such M is transitive.

Extend L_κ to L_κ^+ with one extra constant c , and let

$$\Phi^+ = \Phi \cup \{c \text{ is an ordinal}\} \cup \{c \neq c_x \mid x \in V_\kappa\}$$

Any model of Φ^+ induces a transitive elementary superstructure of V_κ that contains an ordinal at least κ , so by transitivity, κ is in this model.

We show that Φ^+ is satisfiable by showing that it is κ -satisfiable, using the fact that κ is strongly compact. Let $\Phi^0 \subseteq \Phi^+$ be a subset of size less than κ . Then we can interpret c as some ordinal α greater than all ordinals β occurring in the sentences $c \neq c_\beta$ in Φ^+ . Then \mathcal{V} , together with this interpretation of c , is a model of Φ_0 . \square

Corollary.

$$\text{ZFC} + \text{IC} <_{\text{Con}} \text{ZFC} + \text{SCC}$$

The proof above only used languages with at most κ -many symbols. Let $\text{WC}(\kappa)$ be the axiom that every $\mathcal{L}_{\kappa\kappa}$ -language with at most κ -many symbols satisfies κ -compactness. Then we have shown that $\text{WC}(\kappa)$ implies the Keisler extension property. One can show that

$$\text{W}(\kappa) \leftrightarrow \text{WC}(\kappa)$$

VI. Large Cardinals

So the cardinals κ that satisfy $WC(\kappa)$ are precisely the weakly compact cardinals. In particular,

$$ZFC + IC <_{\text{con}} ZFC + WCC$$

Note that in the proof that strongly compact cardinals are measurable, we used a language with 2^κ -many symbols.

3.2. Ultrapowers of the universe

In order to avoid proper classes, we will consider ultrapowers of particular set universes. Later, we will briefly explain how all of this could have been done in a proper class universe such as V . For convenience, we will assume that $\kappa < \lambda$ where κ is measurable and λ is inaccessible, so $V_\lambda \models ZFC + MC$. We will take the ultrapower of V_λ .

Let U be a κ -complete nonprincipal ultrafilter on κ , and form the ultrapower of V_λ , consisting of equivalence classes of functions $f : \kappa \rightarrow V_\lambda$ where $f \sim g$ when $\{\alpha \mid f(\alpha) = g(\alpha)\} \in U$.

$$V_\lambda^\kappa / U = \{[f] \mid f : \kappa \rightarrow V_\lambda\}$$

The membership relation on the ultrapower is given by

$$[f] E [g] \leftrightarrow \{\alpha \mid f(\alpha) \in g(\alpha)\} \in U$$

We have an embedding ℓ from V_λ into the ultrapower by mapping $x \in V_\lambda$ to the equivalence class of its constant function $c_x : \kappa \rightarrow V_\lambda$. This is an elementary embedding by Łoś' theorem. Hence

$$(V_\lambda, \in) \equiv (V_\lambda^\kappa / U, \in)$$

so they both model $ZFC + MC$, and in particular, $[c_\kappa]$ is a measurable cardinal.

Remark. (i) Suppose $V_\lambda^\kappa / U \models [f]$ is an ordinal. By Łoś' theorem,

$$X = \{\alpha \mid f(\alpha) \text{ is an ordinal}\} \in U$$

We can define

$$f'(\alpha) = \begin{cases} f(\alpha) & \text{if } \alpha \in X \\ 0 & \text{otherwise} \end{cases}$$

Note that $f \sim f'$, so $[f] = [f']$. So without loss of generality, we can assume f is a function into $\text{Ord} \cap \lambda = \lambda$, so $f : \kappa \rightarrow \lambda$. Since λ is inaccessible, f cannot be cofinal, so there is $\gamma < \lambda$ such that $f : \kappa \rightarrow \gamma$. Note also that, for example, we can define $f + 1$ by

$$(f + 1)(\alpha) = f(\alpha) + 1$$

so

$$\{\alpha \mid (f + 1)(\alpha) \text{ is the successor of } f(\alpha)\} = \kappa \in U$$

hence by Łoś' theorem, $[f + 1]$ is the successor of $[f]$.

(ii) If $f : \kappa \rightarrow V_\lambda$ is arbitrary, the set

$$\{\text{rank } f(\alpha) \mid \alpha \in \kappa\}$$

cannot be cofinal in λ , so there is $\gamma < \lambda$ such that $f \in V_\gamma$. However, the union of the equivalence class $[f]$ is unbounded in V_λ .

(iii) Given f , by (ii) we may assume $f \in V_\gamma$ for some $\gamma < \lambda$. If $[g] E [f]$, then

$$X = \{\alpha \mid g(\alpha) \in f(\alpha)\} \in U$$

Now we can define

$$g'(\alpha) = \begin{cases} g(\alpha) & \text{if } \alpha \in X \\ 0 & \text{otherwise} \end{cases}$$

Then $g \sim g'$ so $[g] = [g']$, and $g' \in V_\gamma$. Therefore,

$$|\{[g] \mid [g] E [f]\}| \leq |V_\gamma| < \lambda$$

Lemma. V_λ^κ / U is E -well-founded.

Proof. Suppose not, so let $\{[f_n] \mid n \in \mathbb{N}\}$ be a strictly decreasing sequence, so

$$[f_{n+1}] E [f_n]$$

By definition,

$$X_n = \{\alpha \mid f_{n+1}(\alpha) \in f_n(\alpha)\} \in U$$

But as U is κ -complete,

$$\bigcap_{n \in \mathbb{N}} X_n \in U$$

In particular, there must be an element $\alpha \in \bigcap_{n \in \mathbb{N}} X_n$. Hence, $f_n(\alpha)$ is an \in -decreasing sequence in V_λ , which is a contradiction. \square

Note that we only used \aleph_1 -completeness of U .

We can take the Mostowski collapse to produce a transitive set M such that

$$\pi : (V_\lambda^\kappa / U, E) \cong (M, \in)$$

Combining ℓ and π , we obtain

$$j = \pi \circ \ell : (V_\lambda, \in) \rightarrow (M, \in)$$

given by

$$j(x) = \pi(\ell(x)) = \pi([c_x])$$

For convenience, will write (f) to abbreviate $\pi([f])$, so $j(x) = (c_x)$.

VI. Large Cardinals

Lemma. $M \subseteq V_\lambda$.

Proof. Note that because λ is inaccessible, $V_\lambda = H_\lambda$, where

$$H_\lambda = \{x \mid |\text{tcl}(x)| < \lambda\}$$

Since M is transitive, if $|x| < \lambda$ for each $x \in M$, then $M \subseteq H_\lambda$. But remark (iii) above shows precisely what is required. \square

Lemma. $\text{Ord} \cap M = \lambda$.

Proof. Under the elementary embedding j , ordinals in V_λ are mapped to ordinals in M . So j restricts to an order-preserving embedding from λ into a subset of λ . Thus this embedding is unbounded, and therefore by transitivity, $\text{Ord} \cap M = \lambda$. \square

Lemma. $j|_{V_\kappa} = \text{id}$, so in particular, $V_\kappa \subseteq M$.

Proof. We show this by \in -induction on V_κ . Suppose that $x \in V_\kappa$ is such that for all $y \in x$, $j(y) = y$. For any $y \in x$, by elementarity, $j(y) \in j(x)$, but $j(y) = y$ so $y \in j(x)$ as required. For the converse, suppose $y \in j(x)$. Then define f such that $y = (f)$, so $(f) \in (c_x)$. Hence

$$X = \{\alpha \mid f(\alpha) \in c_x(\alpha)\} = \{\alpha \mid f(\alpha) \in x\} \in U$$

But

$$\{\alpha \mid f(\alpha) \in x\} = \bigcup_{z \in x} \{\alpha \mid f(\alpha) = z\}$$

This is a union of $|x|$ -many sets. By κ -completeness, there must be some $z \in x$ such that

$$\{\alpha \mid f(\alpha) = z\} \in U$$

Hence $f \sim c_z$. Therefore, $(f) = j(z)$, and by the inductive hypothesis, $j(z) = z$. Hence $y \in x$. \square

Lemma. $j \neq \text{id}$, as $j(\kappa) > \kappa$.

Proof. We know that $j(\kappa) = (c_\kappa)$. By the previous lemma, for each $\alpha < \kappa$, $j(\alpha) = (c_\alpha) = \alpha$. Consider the identity map $\text{id}_\kappa : \kappa \rightarrow \kappa$. We have

$$\begin{aligned} (c_\alpha) < (\text{id}_\kappa) &\leftrightarrow \{\gamma \mid c_\alpha(\gamma) < \text{id}_\kappa(\gamma)\} \in U \\ &\leftrightarrow \{\gamma \mid \alpha < \gamma\} \in U \end{aligned}$$

But by a size argument, $\{\gamma \mid \gamma \leq \alpha\} \notin U$ as U is nonprincipal, so we must have $\alpha < (\text{id})$. Also,

$$\begin{aligned} (\text{id}_\kappa) < (c_\kappa) &\leftrightarrow \{\gamma \mid \text{id}_\kappa(\gamma) < c_\kappa(\gamma)\} \in U \\ &\leftrightarrow \{\gamma \mid \gamma < \kappa\} \in U \end{aligned}$$

This is certainly in U . So for all $\alpha < \kappa$,

$$\alpha < (\text{id}_\kappa) < j(\kappa)$$

giving

$$\kappa \leq (\text{id}_\kappa) < j(\kappa)$$

as required. \square

Remark. (i) This implies that $j|_{V_{\kappa+1}} \neq \text{id}$, so the identity result above cannot be strengthened.

(ii) This also shows that many of the elements of M arise from non-constant functions.

(iii) The set

$$\{j(x) \mid x \in V_\lambda\}$$

is isomorphic to V_λ . Therefore, there is a (non-transitive) copy of V_λ that sits strictly inside M .

(iv) Let $f : \kappa \rightarrow \kappa$ be a function such that for all $\gamma < \kappa$, $\text{id}_\kappa(\gamma) < f(\gamma)$. Then $(\text{id}_\kappa) < (f)$. For example, the functions $f_2(\gamma) = \gamma \cdot 2$ and $f_3(\gamma) = \gamma \cdot 3$ satisfy $(\text{id}_\kappa) < (f_2) < (f_3)$.

(v) At the moment, we do not know whether $(\text{id}_\kappa) = \kappa$. Consider

$$f(\gamma) = \begin{cases} \gamma - 1 & \text{if } \gamma \text{ is a successor} \\ \gamma & \text{if } \gamma \text{ is a limit} \end{cases}$$

Then

$$(f) < (\text{id}_\kappa) \leftrightarrow \{\alpha \mid \alpha \text{ is a limit}\} \notin U$$

We will discuss this in more detail later.

3.3. Properties above the critical point

Definition. Let $j : V_\lambda \rightarrow M$ be an elementary embedding such that $M \subseteq V_\lambda$ is transitive. An ordinal μ is called the *critical point* of j , written $\text{crit}(j)$, if $j \neq \text{id}$ and μ is the least ordinal α such that $j(\alpha) > \alpha$.

Note that if $j \neq \text{id}$, it moves the rank of some set, so moves some ordinal. Therefore, if $j \neq \text{id}$, it has a critical point.

In this terminology, the critical point of the embedding j above is κ .

Remark. (i) M is closed under finite intersections: if $A, B \in M$, then $A \cap B \in M$.

(ii) $V_\kappa \in M$. To show this, we claim that the set

$$W = \{y \in M \mid M \models \text{rank } y < \kappa\}$$

is equal to V_κ . Then, since M models ZFC, the set W is V_κ^M , so $W \in M$.

VI. Large Cardinals

If $x \in V_\kappa$, then $\text{rank } x = \alpha < \kappa$, so $j(x) = x$. By elementarity, $\text{rank } x = \text{rank } j(x) = j(\alpha) = \alpha$ as required. Conversely, suppose that $M \models \text{rank } y = \gamma$ for $\gamma < \kappa$. There is f such that $y = (f)$, and without loss of generality we can take $f : \kappa \rightarrow V_{\gamma+1}$. But $|V_{\gamma+1}| < \kappa$, and so by the argument in the lemma proving $j|_{V_\kappa} = \text{id}$, there is some $x \in V_{\gamma+1}$ such that $\{\alpha \mid f(\alpha) = x\} \in U$. Hence $f \sim c_x$, and so $y = j(x) = x$.

Lemma. $V_{\kappa+1} \subseteq M$.

Note that $j|_{V_{\kappa+1}} \neq \text{id}$.

Proof. Let $A \in V_{\kappa+1}$, so $A \subseteq V_\kappa$. We claim that $A = j(A) \cap V_\kappa$. Then, by the two remarks above, this implies $A \in M$.

Suppose $x \in A \subseteq V_\kappa$. By elementarity, $j(x) \in j(A)$, but $x = j(x)$, so $x \in j(A)$. Conversely, suppose $x \in j(A) \cap V_\kappa$. Then $x = j(x)$, so $j(x) \in j(A)$. So by elementarity in the other direction, $x \in A$. \square

Lemma. $V_\lambda \models |j(\kappa)| \leq 2^\kappa$.

Proof. Recall that if $f \in V_\gamma$ then $|(f)| \leq |V_\gamma|$. So if $(f) \in j(\kappa) = (c_\kappa)$, we can assume $f : \kappa \rightarrow \kappa$, and there are only 2^κ -many such functions. \square

In particular, V_λ believes that $j(\kappa)$ is not a strong limit cardinal. Hence,

Lemma. $M \neq V_\lambda$.

Proof. M believes that $j(\kappa)$ is measurable, so in particular it believes $j(\kappa)$ is a strong limit. Hence $M \neq V_\lambda$. \square

There is a strengthening of this result which exhibits a witness to $M \subsetneq V_\lambda$, discussed on the example sheets. Namely, we can show that $U \notin M$. In order to show this, we prove that for arbitrary transitive $N \subseteq V_\lambda$ with $U \in N$, we have $N \models |j(\kappa)| \leq 2^\kappa$. In particular, $V_{\kappa+2} \not\subseteq M$.

Note that M might still believe that κ is measurable, even though $U \notin M$. There could be some other $U' \in V_{\kappa+2}$ which is κ -complete and nonprincipal.

Recall that the Keisler extension property for a transitive model X is the statement that there is $\kappa \in X$ such that $V_\kappa \leq X$. Properties of X reflect down into V_κ : if $\alpha \in \text{Ord}^{V_\kappa}$ and Φ is a property such that $X \models \Phi(\kappa)$, then

$$X \models \exists \mu. \alpha < \mu \wedge \Phi(\mu)$$

so

$$V_\kappa \models \exists \mu. \alpha < \mu \wedge \Phi(\mu)$$

hence

$$C_\Phi = \{\gamma < \kappa \mid \Phi(\gamma)\} \subseteq \kappa$$

is cofinal in κ . Now, if Φ is any property such that $M \models \Phi(\kappa)$, then for any $\alpha < \kappa$,

$$M \models \exists \mu. j(\alpha) < \mu < j(\kappa) \wedge \Phi(\mu)$$

By elementarity,

$$V_\lambda \models \exists \mu. \alpha < \mu < \kappa \wedge \Phi(\mu)$$

Note that $\alpha = j(\alpha)$. So

$$C_\Phi = \{\gamma < \kappa \mid \Phi(\gamma)\}$$

is cofinal in κ .

Example. (i) Let $\Phi(\kappa) = I(\kappa)$ be the statement that κ is inaccessible. By absoluteness, $M \models I(\kappa)$, so

$$C_I = \{\gamma < \kappa \mid I(\gamma)\}$$

is cofinal. So if κ is measurable, it is the κ th inaccessible cardinal.

(ii) Let $\Phi(\kappa) = W(\kappa)$ be the statement that κ is weakly compact. We show that $M \models W(\kappa)$. Let $c : [\kappa]^2 \rightarrow 2$ be a colouring in M ; we find $H \in [\kappa]^\kappa$ in M that is monochromatic for c . By the fact that $V_\lambda \models W(\kappa)$, we obtain H as above in V_λ . But this H is a subset of κ , so is an element of $V_{\kappa+1} \subseteq M$ as required. By the reflection argument,

$$C_W = \{\gamma < \kappa \mid W(\gamma)\}$$

is cofinal in κ . So the least weakly compact cardinal is not measurable.

Definition. A property Φ is called β -stable if for all transitive models M and all κ , if $\Phi(\kappa)$ holds and $V_{\kappa+\beta} \subseteq M$ then $M \models \Phi(\kappa)$.

Remark. (i) Weak compactness is 1-stable, and 1-stable properties of measurable cardinals reflect at a measurable cardinal.

(ii) Measurability is 2-stable, because the property Ξ of being a κ -complete nonprincipal ultrafilter is absolute, but the existence of the ultrafilter requires two power set operations:

$$M(\kappa) \leftrightarrow \exists U \in V_{\kappa+2}. \Xi(U)$$

Example. Suppose that $M \models M(\kappa)$. Then by the same reflection argument, the set C_M is cofinal in κ , so κ is the κ th measurable cardinal, and so is not the least.

Definition. A cardinal κ is called *surviving*, written $\text{Surv}(\kappa)$, if there is $\lambda > \kappa$ inaccessible, a κ -complete nonprincipal ultrafilter on κ , a transitive model M such that $M \cong V_\lambda^\kappa / U$ and j is the elementary embedding derived from U , where $M \models M(\kappa)$.

By the example above, if κ is the first surviving cardinal, it is the κ th measurable. Under sufficient consistency assumptions, we have the following.

Corollary. $\text{MC} <_{\text{Con}} \text{SurvC}$.

VI. Large Cardinals

Proof. Let κ be a surviving cardinal. By the previous results, we can find $\lambda_0 < \lambda_1 < \kappa$ such that λ_0, λ_1 are both measurable. Then λ_1 is inaccessible, so $V_{\lambda_1} \models \text{ZFC} + M(\lambda_0)$ by 2-stability of measurability and the fact that $V_{\lambda_0+2} \subseteq V_{\lambda_1}$. \square

3.4. The fundamental theorem on measurable cardinals

Theorem. Suppose λ is inaccessible and $\kappa < \lambda$. Then the following are equivalent.

- (i) κ is measurable.
- (ii) There is a transitive model M of ZFC with $V_{\kappa+1} \subseteq M$ and an elementary embedding $j : V_\lambda \rightarrow M$ such that $j \neq \text{id}$ and $\kappa = \text{crit}(j)$.

Proof. We have already shown that (i) implies (ii). For the converse, we define an ultrafilter U by

$$U = \{A \subseteq \kappa \mid \kappa \in j(A)\}$$

Note that if $A \subseteq \kappa$, then $j(A) \subseteq j(\kappa)$, so it could in fact be the case that $\kappa \in j(A)$. We show that U is a κ -complete nonprincipal ultrafilter.

- We have $\kappa \in U$ precisely if $\kappa \in j(\kappa)$, but this is true as κ is the critical point of j .
- $\emptyset \in U$ precisely if $\kappa \in j(\emptyset)$, but $j(\emptyset) = \emptyset$ as j is an elementary embedding.
- If $A \in U$ and $B \supseteq A$, then $\kappa \in j(A)$, but $j(B) \supseteq j(A)$ by elementarity, so $\kappa \in j(B)$ giving $B \in U$.
- Suppose $A \notin U$. Then $\kappa \notin j(A)$. We want to show $\kappa \setminus A \in U$, or equivalently, $\kappa \in j(\kappa \setminus A)$. By elementarity, $j(\kappa \setminus A) = j(\kappa) \setminus j(A)$. But $\kappa \in j(\kappa) \setminus j(A)$ as required.
- We show U is nonprincipal. Let $\alpha \in \kappa$. Then $\{\alpha\} \in U$ precisely when $\kappa \in j(\{\alpha\}) = \{j(\alpha)\}$. But $\alpha < \kappa$, so $j(\alpha) = \alpha \neq \kappa$, hence U cannot be principal.
- Finally, we show κ -completeness; this will also show the finite intersection property required for U to be a filter. Let $\gamma < \kappa$, and fix $(A_\alpha)_{\alpha < \gamma}$ such that $A_\alpha \in U$ for each $\alpha < \gamma$. Then $\kappa \in j(A_\alpha)$ for all $\alpha < \gamma$. Then $\bigcap_{\alpha < \gamma} A_\alpha \in U$ if and only if $\kappa \in j\left(\bigcap_{\alpha < \gamma} A_\alpha\right)$. Note that being an element of $\bigcap_{\alpha < \gamma} A_\alpha$ is a formula that says that \mathbf{A} is a sequence of objects A_α , the α th element of this sequence is A_α , and β is an element of each element of the sequence. Therefore $\beta \in j\left(\bigcap_{\alpha < \gamma} A_\alpha\right)$ if and only if β is an element of all elements of the sequence $j(\mathbf{A})$. Clearly, $j(\mathbf{A})$ is a sequence of subsets of $j(\kappa)$ of length $j(\gamma) = \gamma$. Since A_α is the α th element of \mathbf{A} , $j(A_\alpha)$ is the $j(\alpha)$ th element of $j(\mathbf{A})$, but $j(\alpha) = \alpha$. Hence $j(\mathbf{A})$ is the sequence $(j(A_\alpha))_{\alpha < \gamma}$. Then

$$j\left(\bigcap_{\alpha < \gamma} A_\alpha\right) = \bigcap_{\alpha < \gamma} j(A_\alpha)$$

giving κ -completeness as required.

□

Remark. Given a sequence \mathbf{A} of subsets of κ of length γ , then $j(\mathbf{A})$ is a sequence of subsets of $j(\kappa)$ of length $j(\gamma)$. Moreover, if A_α is the α th element of \mathbf{A} , then $j(A_\alpha)$ is the $j(\alpha)$ th element of $j(\mathbf{A})$. In the situation above, $\gamma < \kappa$, so $j(\gamma) = \gamma$ and $j(\alpha) = \alpha$, so $j(\mathbf{A}) = \{j(A_\alpha) \mid \alpha < \gamma\}$. If, for example, $\gamma = \kappa$, then $j(\mathbf{A})$ is a sequence of length $j(\kappa)$, which is strictly longer. Despite this, the first κ -many elements of the sequence are still $j(A_\alpha)$ for $\alpha < \kappa$. Beyond κ , we do not know what the elements of $j(\mathbf{A})$ look like. This remark suffices for the following result.

Proposition. For arbitrary embeddings j with critical point κ , the ultrafilter U_j constructed above is normal.

Proof. Suppose $A_\alpha \in U_j$ for each $\alpha < \kappa$, or equivalently, $\kappa \in j(A_\alpha)$. We must show $\kappa \in j(\Delta_{\alpha < \kappa}(A_\alpha))$. We have

$$\begin{aligned} \xi \in \Delta_{\alpha < \kappa}(A_\alpha) &\leftrightarrow \xi \in \bigcap_{\alpha < \xi} A_\alpha \\ &\leftrightarrow \forall \alpha < \xi. \xi \in A_\alpha \\ \xi \in j\left(\Delta_{\alpha < \kappa}(A_\alpha)\right) &\leftrightarrow \forall \alpha < \xi. \xi \in j(\mathbf{A})_{j(\alpha)} \end{aligned}$$

Substitute κ for ξ and obtain

$$\begin{aligned} \kappa \in j\left(\Delta_{\alpha < \kappa}(A_\alpha)\right) &\leftrightarrow \forall \alpha < \kappa. \kappa \in j(\mathbf{A})_{j(\alpha)} \\ &\leftrightarrow \forall \alpha < \kappa. \kappa \in j(\mathbf{A})_\alpha \\ &\leftrightarrow \forall \alpha < \kappa. \kappa \in j(A_\alpha) \end{aligned}$$

which holds by assumption. □

Remark. (i) This gives an alternative proof of the existence of a normal ultrafilter on a measurable cardinal.

(ii) The operations $U \mapsto j_U$ and $j \mapsto U_j$ are not inverses in general. In particular, if U is not normal, $U_{j_U} \neq U$.

Proposition. Let U be a κ -complete nonprincipal ultrafilter on κ . Then the following are equivalent.

- (i) U is normal;
- (ii) $(\text{id}) = \kappa$.

This proposition provides an alternative view of reflection. Suppose that the ultrafilter U on κ is normal. If $M \models \Phi(\kappa)$, then $M \models \Phi((\text{id}))$. By Łoś' theorem,

$$\{\alpha < \kappa \mid \Phi(\text{id}(\alpha))\} \in U$$

VI. Large Cardinals

So Φ reflects not only on a set of size κ , but on an ultrafilter set. In particular, if $\Phi = M$ and $M \models M(\kappa)$, so if κ is surviving, then the set of α that are measurable is in U . Using this result, we can characterise the surviving cardinals in a more elegant way.

Theorem. κ is surviving if and only if there is a normal ultrafilter on κ such that $\{\alpha < \kappa \mid M(\alpha)\} \in U$.

Proof. We have just shown one direction. For the converse, suppose the set $C = \{\alpha < \kappa \mid M(\alpha)\}$ is in U . Then for each $\alpha \in C$, one can find an α -complete nonprincipal ultrafilter on α called U_α . Define

$$f(\alpha) = \begin{cases} U_\alpha & \text{if } \alpha \in C \\ \emptyset & \text{if } \alpha \notin C \end{cases}$$

Thus the set of α such that $f(\alpha)$ is an α -complete nonprincipal ultrafilter on α is C , so in U . Equivalently, the set of α such that $f(\alpha)$ is an $\text{id}(\alpha)$ -complete nonprincipal ultrafilter on $\text{id}(\alpha)$ is in U . So by Łoś' theorem, M believes that (f) is an (id) -complete nonprincipal ultrafilter on (id) . So (f) witnesses that κ is measurable in M . \square

This shows that whether a cardinal κ is surviving depends only on $V_{\kappa+2}$, and is therefore a 2-stable property.

Definition. If U, U' are normal ultrafilters on κ , we write $U <_M U'$ if

$$C = \{\alpha \mid M(\alpha)\} \in U$$

and there is a sequence of ultrafilters U_α on $\alpha \in C$ such that

$$A \in U' \leftrightarrow \{\alpha \mid A \cap \alpha \in U_\alpha\} \in U$$

This is known as the *Mitchell order*.

Then κ is surviving if and only if there are U, U' on κ such that $U <_M U'$, because of the fact that if $h(\alpha) = A \cap \alpha$ then $(h) = A$. Note that talking about sequences of Mitchell-ordered ultrafilters is also 2-stable.

4. Towards inconsistency

4.1. Strong cardinals

Definition. A large cardinal axiom Φ is called an *embedding axiom* if $\Phi(\kappa)$ holds if and only if there is a transitive model M and elementary embedding $j : V_\lambda \rightarrow M$ with critical point κ with certain additional properties.

$M(\kappa)$ is the simplest embedding axiom. The remaining large cardinal axioms in this course will take the form of embedding axioms.

Definition. An embedding $j : V_\lambda \rightarrow M$ with critical point κ is called β -strong if $V_{\kappa+\beta} \subseteq M$. A cardinal κ is called β -strong if there is a β -strong embedding with critical point κ .

β -stable properties are preserved by β -strong embeddings. In particular, by the reflection argument, if Φ is β -stable and κ is β -strong with $\Phi(\kappa)$, then κ is the κ th cardinal with property Φ .

Note that κ is measurable if and only if κ is 1-strong, and if κ is 2-strong then $\{\alpha < \kappa \mid M(\alpha)\}$ and $\{\alpha < \kappa \mid \text{Surv}(\alpha)\}$ are of size κ . If we write $\beta\text{--}S(\kappa)$ to denote that κ is β -strong, then

$$\text{SurvC} <_{\text{Con}} 2\text{--}S(\kappa)$$

This also gives an example of $j_{U_j} \neq j$, as the ultrapower embedding of any ultrafilter is never 2-strong.

Definition. A large cardinal property Φ is said to have *witness objects* of rank β if there is a formula Ψ that is downwards absolute for transitive models such that

$$\Phi(\kappa) \leftrightarrow \forall x. \exists y \in V_{\kappa+\beta}. \Psi(x, y, \kappa)$$

Any large cardinal property with witness objects of rank β is β -stable.

Example. (i) Weakly compact cardinals have witness objects of rank 1: for all colourings, there exists a homogeneous set in $V_{\kappa+1}$.

(ii) Measurable cardinals have witness objects of rank 2: there is a κ -complete nonprincipal ultrafilter on κ . The initial $\forall x$ quantifier is not needed in this case.

(iii) Surviving cardinals also have witness objects of rank 2, namely, a pair of ultrafilters.

In particular, inaccessibility is 0-stable, weak compactness is 1-stable, and measurability and survivability are 2-stable.

Remark. If β -strong cardinals have witness objects, they cannot be of rank β , because then they would reflect below. Witness objects for strength exist and are called *extenders*, and if μ is the least \beth fixed point larger than $|V_{\kappa+\beta}|$, then the witness object for β -strength has rank at most μ .

Definition. A cardinal κ is called *strong* if it is β -strong for all $\beta < \lambda$.

VI. Large Cardinals

Importantly, the quantifiers are

$$\forall\beta. \exists j. V_{\kappa+\beta} \subseteq M$$

This does not say that there exists an embedding where all of the $V_{\kappa+\beta}$ are subsets of the same M . This notion cannot have a single witness object of a fixed rank, since otherwise, strength would reflect strength.

4.2. Removing the inaccessible

The ultrapower constructions used an inaccessible cardinal above a measurable cardinal, so that we could obtain a set-sized universe containing a measurable cardinal. When trying to do this with the real universe, we encounter several problems.

- (i) The definition of ultrapowers requires a set model.
- (ii) In the fundamental theorem of measurable cardinals, we have a quantification over j and M . If these are proper classes, this quantification cannot be expressed in the usual language of set theory.
- (iii) Also, in the fundamental theorem of measurable cardinals, we use the notion of an elementary embedding, which is only definable for set models.

To solve problem (i), we would like to construct V^κ / \sim_U . Note that V^κ is a well-defined class; it is the class of all functions with domain κ . For such functions, it is easy to define the equivalence relation \sim_U . However, the equivalence classes $[f]_U = \{g \in V^\kappa \mid f \sim_U g\}$ are all proper classes. So V^κ / \sim_U is no longer a standard class; classes containing proper classes are typically not allowed. This can be resolved using *Scott's trick*. If C is a nonempty class, then there is a minimal α such that $C \cap V_\alpha \neq \emptyset$. This is a nonempty set. Define $\text{scott}(C) = C \cap V_\alpha$ for this α . Hence, if $[f]_U \neq [g]_U$, we have $\text{scott}([f]_U) \neq \text{scott}([g]_U)$. We can therefore define

$$V^\kappa / \sim_U = \{\text{scott}([f]_U) \mid \text{dom } f = \kappa\}$$

To obtain our model M , we took the Mostowski collapse of V^κ / \sim_U . Therefore, we need a class version of the Mostowski collapse. Recall that a relation $E \subseteq C \times C$ is *set-like* if for all $x \in C$, the class $\{y \in C \mid y E x\}$ is a set.

Theorem. Let C be a class, and let $E \subseteq C \times C$ be a binary relation on C that is well-founded, extensional, and set-like. Then there is a unique transitive class T such that $(T, \in) \cong (C, E)$.

This may be proven in an almost identical fashion to Mostowski's collapsing theorem for sets.

For problems (ii) and (iii), recall that the fundamental theorem of measurable cardinals was that $M(\kappa)$ is equivalent to the statement that there is an elementary embedding $j : V_\lambda \rightarrow M$ with critical point κ . Measurability is witnessed at $\kappa + 2$, but the elementary embedding is not witnessed anywhere below λ , so we cannot extend this definition to the usual universe. We can solve this by extending our set theory to an appropriate class theory. Standard class

theories include *von Neumann–Bernays–Gödel* or NBG, and *Morse–Kelley* or MK. These theories have very different notions of class. NBG set theory is based upon the idea that definable formulas give the classes. It is a ‘minimal class theory’ where all classes are definable. MK is based on the idea that Ord behaves externally like an inaccessible cardinal. In this theory, there could be undefinable classes, and more classes than sets.

This resolves problem (ii), as we are permitted to work in a language in which we may quantify over proper classes. However, this does not solve problem (iii). Elementarity cannot be expressed as a single formula, but becomes a schema. This causes additional problems as we need the existential over j and M to be part of each formula. This could be solved by extending the language to add symbols for j and M . Another resolution is to observe that Σ_1 -elementarity suffices, as is explored in Kanamori’s book *The Higher Infinite* on page 45. This can be defined using a single formula, therefore solving problem (iii).

4.3. Supercompact cardinals

Definition. M is closed under μ -sequences if $M^\mu \subseteq M$.

Theorem. If κ is measurable and $j : V_\lambda \rightarrow M$ is the ultrapower embedding, then M is closed under κ -sequences but not κ^+ -sequences.

Proof. Let $S = \{(f_\alpha) \mid \alpha < \kappa\} \in M^\kappa$. We must show that $S \in M$. Find h such that $(h) = \kappa$. For $\xi \in \kappa$, define $g(\xi)$ to be a function with domain $h(\xi)$ such that for all $\alpha \in h(\xi)$,

$$g(\xi)(\alpha) = f_\alpha(\xi)$$

Then

$$\{\xi \mid \text{dom } g(\xi) = h(\xi)\} = \kappa \in U$$

By Łoś’ theorem, $\text{dom}(g) = (h) = \kappa$. Further,

$$\{\xi \mid \forall \alpha \in \text{dom } g(\xi). g(\xi)(\alpha) = f_\alpha(\xi)\} = \kappa \in U$$

so again by Łoś’ theorem, if $\alpha \in \text{dom}(g) = \kappa$, then $(g)(\alpha) = (f_\alpha)$. Hence $(g) = S$.

Let

$$T = \{j(\alpha) \mid \alpha < \kappa^+\} \in M^{\kappa^+}$$

We claim that $T \notin M$. To prove this, we first show that T is unbounded in $j(\kappa^+)$, which is equal to $j(\kappa)^+$ by elementarity. Indeed, consider an arbitrary $(f) < j(\kappa^+)$. Then $j(\kappa^+) = (c_{\kappa^+})$, so without loss of generality we can assume $f : \kappa \rightarrow \kappa^+$. As κ^+ is regular, f is bounded by some $\alpha < \kappa^+$, so $f : \kappa \rightarrow \alpha$. Then $(f) < (c_\alpha) = j(\alpha) \in T$.

Now, note that $j(\kappa^+) = j(\kappa)^+$ is a regular cardinal, so cannot have small unbounded subsets. But $|T| = \kappa^+ < j(\kappa)^+$, so $T \notin M$. \square

Definition. An embedding j is called μ -supercompact if $M^\mu \subseteq M$. A cardinal κ is called μ -supercompact if there is a μ -supercompact embedding with critical point κ .

VI. Large Cardinals

Therefore, the theorem above shows that if κ is measurable, then it is κ -supercompact, and the ultrapower embedding is not κ^+ -supercompact, although there could be other embeddings that are.

Definition. A cardinal κ is called *supercompact* if it is μ -supercompact for all $\mu < \lambda$.

As with strong cardinals, the quantifiers are in the order

$$\forall \mu. \exists j. M^\mu \subseteq M$$

If κ is 2^κ -supercompact, then κ is 2-strong. First note that $V_{\kappa+2} = \mathcal{P}(V_{\kappa+1})$ and $|V_{\kappa+1}| = |\mathcal{P}(V_\kappa)| = 2^\kappa$. Every $A \in V_{\kappa+1}$ is a 2^κ -sequence of elements of M , so if every 2^κ -length sequence lies in M , then $A \in M$ as required. In general, if κ is $|V_{\kappa+\beta}| = \beth_{\kappa+\beta}$ -supercompact, then κ is $(\beta + 1)$ -strong. In particular,

Corollary. Every supercompact cardinal is strong.

4.4. The upper limit

We now consider reversing the quantifier order in the definition of a strong cardinal.

Definition. A cardinal κ is called *Reinhardt* if there is an embedding j such that for all β , we have $V_{\kappa+\beta} \subseteq M$, or equivalently, $M = V_\lambda$. In other words, there is an elementary embedding $j : V_\lambda \rightarrow V_\lambda$ with critical point κ .

Theorem (Kunen). ZFC proves that there are no Reinhardt cardinals.

It is an open problem whether ZF without AC proves there are no Reinhardt cardinals.

Proof. Suppose $j : V_\lambda \rightarrow M$ has critical point κ . Find the least j -fixed point above κ , by defining

$$\kappa_0 = \kappa; \quad \kappa_{i+1} = j(\kappa_i); \quad \hat{\kappa} = \bigcup_{i \in \omega} \kappa_i$$

so $j(\hat{\kappa}) = \hat{\kappa}$. We will show that $V_{\hat{\kappa}+1} \not\subseteq M$, which is a result called *Kunen's lemma*. This contradicts the assumption that $M = V_\lambda$.

We need a combinatorial lemma due to Erdős and Hajnal. For a cardinal δ , we say that $f : [\delta]^\omega \rightarrow \delta$ is ω -Jónsson if for every $X \subseteq \delta$ such that $|X| = \delta$, we have $\{f(A) \mid A \in [X]^\omega\} = \delta$. The lemma states that every cardinal δ has an ω -Jónsson function.

Suppose that $V_{\hat{\kappa}+1} \subseteq M$. Let $f : [\hat{\kappa}]^\omega \rightarrow \hat{\kappa}$ be an ω -Jónsson function for $\hat{\kappa}$. Then M believes that $j(f)$ is an ω -Jónsson function for $j(\hat{\kappa}) = \hat{\kappa}$. Define

$$X = \{j(\alpha) \mid \alpha \in \hat{\kappa}\} \in V_{\hat{\kappa}+1}$$

We claim that $X \notin M$, finishing the proof. Suppose $X \in M$. Clearly $|X| = \hat{\kappa}$, so then $M \models |X| = \hat{\kappa}$. We can apply the definition of an ω -Jónsson function in M , which shows that

$$M \models \{j(f)(A) \mid A \in [X]^\omega\} = \hat{\kappa}$$

4. Towards inconsistency

Any $A \in [X]^\omega$ is of the form $\{j(\alpha_i) \mid i \in \omega\}$ for some $a = \{\alpha_i \mid i \in \omega\} \in [\hat{\kappa}]^\omega$. Then $j(a) = \{j(\alpha_i) \mid i \in \omega\} = A$.

In general, if $g(x) = y$, then g is a function, $x \in \text{dom } g$, and $\langle x, y \rangle \in g$. Applying j , we have that $j(g)$ is a function, $j(x) \in \text{dom } j(g)$, and $\langle j(x), j(y) \rangle \in j(g)$. So $j(g)(j(x)) = j(y) = j(g(x))$. Therefore, we obtain $j(f)(A) = j(f)(j(a)) = j(f(a))$. But $f(a) \in \hat{\kappa}$, so $j(f(a)) \in X$. Therefore,

$$M \models \hat{\kappa} = \{j(f)(A) \mid A \in [X]^\omega\} \subseteq X$$

But this cannot be true, for example because $\kappa \notin X$ but $\kappa \in \hat{\kappa}$. □

Remark. (i) The combinatorial lemma was proven using AC, and it is not known whether the proof works without it.

- (ii) To prove Kunen's lemma, we did not need that λ is inaccessible. More explicitly, if $j : V_\alpha \rightarrow M$ is an elementary embedding with critical point κ such that $\hat{\kappa} + 2 \leq \alpha$ (to guarantee that $f \in V_\alpha$), then $V_{\hat{\kappa}+1} \not\subseteq M$.

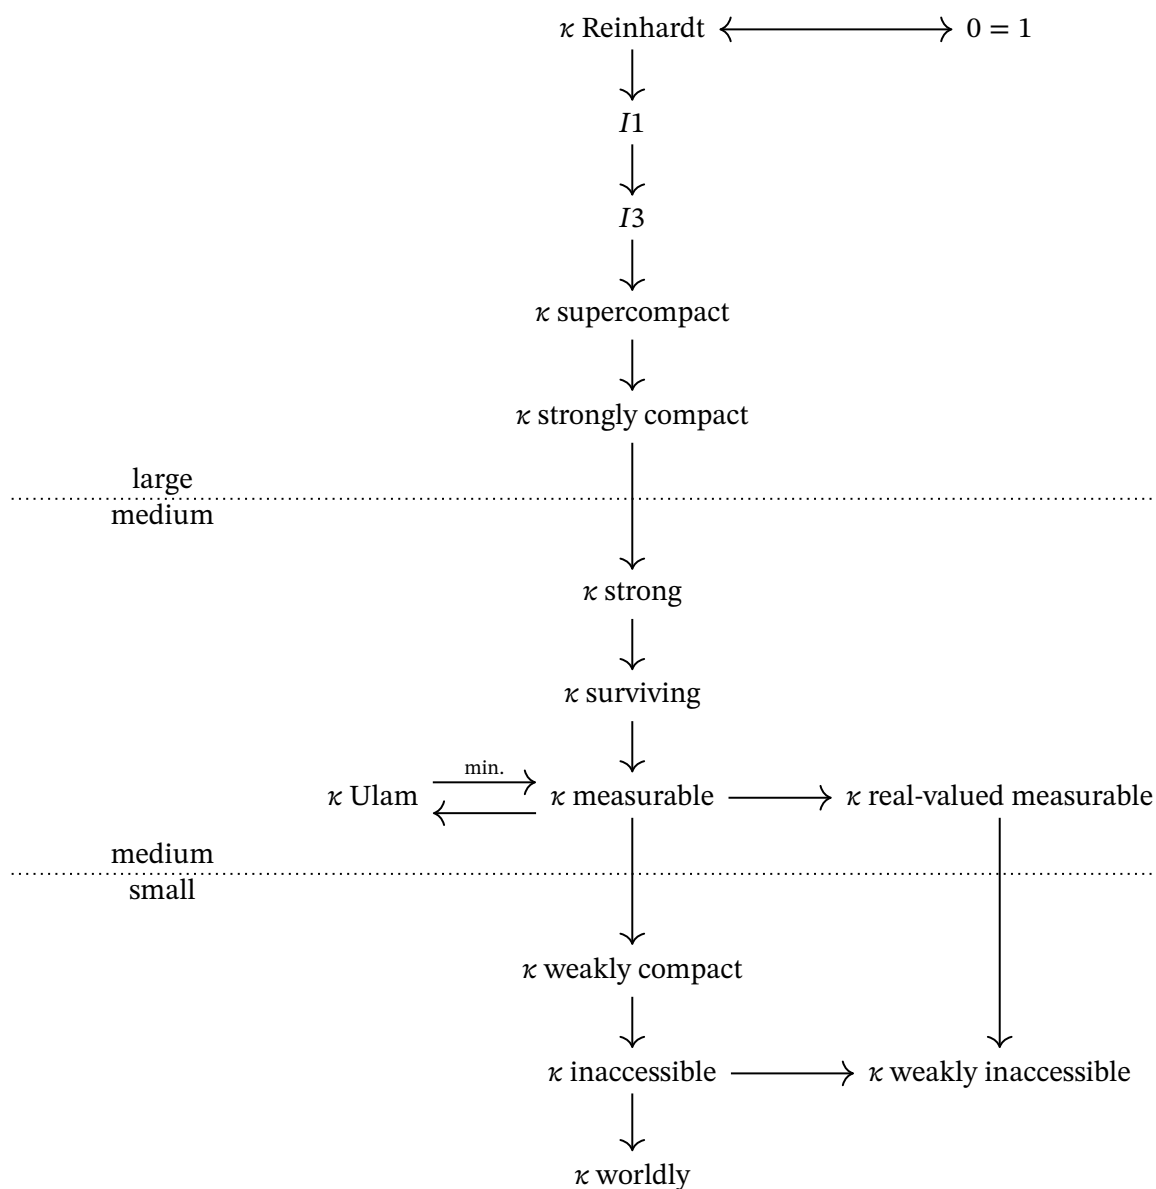
Corollary. For any ordinal δ , there is no elementary embedding $j : V_{\delta+2} \rightarrow V_{\delta+2}$ with critical point less than $\delta + 2$.

Proof. Observe that if $\kappa < \delta + 2$ is the critical point, we cannot have $\kappa = \delta$ or $\kappa = \delta + 1$, because δ and $\delta + 1$ are definable in $V_{\delta+2}$. Then $j(\kappa) < \delta$, so by induction all of the iterated images of κ under j are less than δ , so $\hat{\kappa} \leq \delta$. Thus $\hat{\kappa} + 2 \leq \delta + 2$, so by remark (ii), $V_{\hat{\kappa}+1} \not\subseteq V_{\delta+2}$, giving a contradiction. □

The axiom stating the existence of an analogous $j : V_{\delta+1} \rightarrow V_{\delta+1}$ is called *I1*, and the existence of $j : V_\delta \rightarrow V_\delta$ is called *I3*; there is an axiom *I2* in between. Clearly, if $j : V_{\delta+1} \rightarrow V_{\delta+1}$ is elementary, then so is $j|_{V_\delta} : V_\delta \rightarrow V_\delta$, so *I1* implies *I3*. It has been hypothesised that *I1* and *I3* are inconsistent, but we do not yet have a proof.

Diagram of large cardinal properties

Under suitable consistency assumptions, large cardinal properties that appear in higher positions on this diagram have strictly higher consistency strength than properties appearing lower down the diagram.



The 'small large cardinals' are usually considered those cardinals consistent with $V = L$, and such large cardinal properties are usually downwards absolute. Note that L has no measurable cardinals. Indeed, if $V = L$ and U is a κ -complete nonprincipal ultrafilter on κ , then

4. Towards inconsistency

the ultrapower embedding $j_U : L \rightarrow M$ must map to an inner model strictly smaller than L , but such an inner model cannot exist.

There are certain large cardinals called *Woodin cardinals* which sit between strong and strongly compact cardinals. They represent another boundary between sizes of large cardinal axioms, just like measurable cardinals; smaller large cardinals are sometimes called ‘medium-sized large cardinals’, and the others are called ‘large large cardinals’. Woodin cardinals are crucial for understanding the connection between large cardinals and infinite games. We know very little about large cardinal axioms beyond Woodin cardinals.

VII. Forcing and the Continuum Hypothesis

Lectured in Lent 2024 by DR. R. MATTHEWS

(Course description goes here.)

Contents

1.	Set theoretic preliminaries	336
1.1.	Introduction to independence results	336
1.2.	Systems of set theory	337
1.3.	Adding defined functions	339
1.4.	Absoluteness	340
1.5.	The Lévy hierarchy	341
1.6.	Transfinite recursion	346
1.7.	The reflection theorem	347
1.8.	Cardinal arithmetic	350
2.	Constructibility	354
2.1.	Definable sets	354
2.2.	Defining the constructible universe	354
2.3.	Gödel functions	356
2.4.	The axiom of constructibility	362
2.5.	Well-ordering the universe	363
2.6.	The generalised continuum hypothesis in L	364
2.7.	Combinatorial properties	366
3.	Forcing	368
3.1.	Introduction	368
3.2.	Forcing posets	369
3.3.	Chains and Δ -systems	370
3.4.	Dense sets and genericity	372
3.5.	Names	373
3.6.	Canonical names	374
3.7.	Verifying the axioms: part one	375
3.8.	The forcing relation	378
3.9.	The forcing theorem	381
3.10.	Verifying the axioms: part two	384
4.	Forcing and independence results	387
4.1.	Independence of the constructible universe	387
4.2.	Cohen forcing	388
4.3.	Preservation of cardinals	389
4.4.	The failure of the continuum hypothesis	391
4.5.	Possible sizes of the continuum	392
4.6.	Larger chain conditions	394
4.7.	Closure and distributivity	396
4.8.	The mixing lemma	398

4.9.	Forcing successor cardinals	399
4.10.	Product forcing	401

1. Set theoretic preliminaries

1.1. Introduction to independence results

Independence results are found across mathematical disciplines.

- (i) The *parallel postulate* is independent from the other four postulates of Euclidean geometry. It states that for any given point not on a line, there is a unique line passing through that point that does not intersect the given line. In the 19th century, it was shown that the other four postulates are satisfied by hyperbolic geometry, but this postulate is not satisfied. This shows that the other four axioms are insufficient to prove the parallel postulate.
- (ii) Let φ be the statement in the language of fields describing the existence of a square root of 2. We know that \mathbb{Q} is a field satisfying $\neg\varphi$, and $\mathbb{Q}[\sqrt{2}]$ satisfies φ . The fields \mathbb{Q} and $\mathbb{Q}[\sqrt{2}]$ are models of the theory of fields, one of which satisfies φ , and one of which satisfies $\neg\varphi$. This shows that the theory of fields does not prove φ or $\neg\varphi$. A similar result holds for the statement φ that says that there are no roots of $x^4 = -1$.
- (iii) Gödel's incompleteness theorem implies that there must always be an independence result in a sufficiently powerful consistent set theory.

We will show that there are other independence results in set theory that are not self-referential like the Gödel incompleteness theorems.

Theorem (Cantor). $|\mathbb{N}| < |\mathcal{P}(\mathbb{N})|$.

The continuum hypothesis is that there are no sets of cardinality strictly between $|\mathbb{N}|$ and $|\mathcal{P}(\mathbb{N})| = |\mathbb{R}|$.

Definition. The *continuum hypothesis* CH states that if $X \subseteq \mathcal{P}(\mathbb{N})$ is infinite, then either $|X| = |\mathbb{N}|$ or $|X| = |\mathcal{P}(\mathbb{N})|$, or equivalently,

$$2^{\aleph_0} = \aleph_1$$

Progress was made on the continuum hypothesis in the 19th and 20th centuries.

- (i) In 1883, Cantor showed that any closed subset of \mathbb{R} satisfies CH.
- (ii) In 1916, Alexandrov and Hausdorff showed that any Borel set of \mathbb{R} satisfies CH.
- (iii) In 1930, Suslin strengthened this result to analytic subsets of \mathbb{R} .
- (iv) In 1938, Gödel showed that if ZF is consistent, then so is ZFC + CH.
- (v) However, in 1963, Cohen showed that if ZF is consistent, then so is ZFC + \neg CH.

In this course, we will prove results (iv) and (v), thus establishing the independence of the continuum hypothesis from ZFC.

1.2. Systems of set theory

The language of set theory $\mathcal{L} = \mathcal{L}_\in$ is a first-order predicate logic with equality and membership as primitive relations. We assume the existence of infinitely many variables v_1, v_2, \dots denoting sets. We will only use the logical connectives \vee and \neg as well as the existential quantifier \exists . Conjunction, implication, and universal quantification can be defined in terms of disjunction, negation, and existential quantification.

We say that an occurrence of a variable x is *bound* in a formula φ if it is in a quantifier $\exists x$ or lies in the scope of such a quantifier. An occurrence is called *free* if it is not bound. We write $FV(\varphi)$ for the set of free variables of φ . We will write $\varphi(u_1, \dots, u_n)$ to emphasise the dependence of φ on its free variables u_1, \dots, u_n . By doing so, we will allow ourselves to freely change the names of the free variables, and assume that substituted variables are free. The syntax $\varphi(u_1, \dots, u_n)$ does not imply that u_i occurs freely, or even at all.

Some of the most common axioms of set theory are as follows.

(i) *Axiom of extensionality.*

$$\forall x. \forall y. (\forall z. (z \in x \leftrightarrow z \in y) \rightarrow x = y)$$

(ii) *Axiom of empty set.*

$$\exists x. \forall y \in x. y \neq y$$

(iii) *Axiom of pairing.*

$$\forall x. \forall y. \exists z. (x \in z \wedge y \in z)$$

(iv) *Axiom of union.*

$$\forall a. \exists x. \forall y. (y \in x \leftrightarrow \exists z \in a. y \in z)$$

(v) *Axiom of foundation.*

$$\forall x. (\exists y. y \in x \rightarrow \exists y \in x. \neg \exists z \in x. z \in y)$$

(vi) *Axiom scheme of separation.* For any formula φ ,

$$\forall a. \exists x. \forall y. (y \in x \leftrightarrow (y \in a \wedge \varphi(y)))$$

(vii) *Axiom of infinity.*

$$\exists a. (\exists x. (x \in a) \wedge \forall x \in a. \exists y \in a. x \in y)$$

(viii) *Axiom of power set.*

$$\forall a. \exists x. \forall y. (y \in x \leftrightarrow \forall z. (z \in y \rightarrow z \in a))$$

VII. Forcing and the Continuum Hypothesis

(ix) *Axiom scheme of replacement.* For any formula φ ,

$$\forall a. (\forall x \in a. \exists! y. \varphi(x, y) \rightarrow \exists b. \forall x \in a. \exists y \in b. \varphi(x, y))$$

(ix') *Axiom scheme of collection.* For any formula φ ,

$$\forall a. (\forall x \in a. \exists y. \varphi(x, y) \rightarrow \exists b. \forall x \in a. \exists y \in b. \varphi(x, y))$$

(x) *Axiom of choice.*

$$\forall X. (\emptyset \notin X \rightarrow \exists f : (X \rightarrow \bigcup X). \forall a \in X. f(a) \in a)$$

(x') *Well-ordering principle.*

$$\forall a. \exists R. R \text{ is a well-ordering of } a$$

Some common set theories are as follows.

- *Zermelo set theory Z* consists of axioms (i) to (viii). Axioms (ix) and (ix') are equivalent relative to Z.
- *Zermelo–Fraenkel set theory ZF* consists of axioms (i) to (ix). Axioms (x) and (x') are equivalent relative to ZF.
- *Zermelo–Fraenkel set theory with choice ZFC* consists of axioms (i) to (x).
- *Zermelo–Fraenkel set theory without power set ZF^-* consists of axioms (i) to (vii), with the axiom of collection (ix') instead of replacement (ix); it has been shown that (ix) is weaker than (ix') in the presence of axioms (i) to (vii).
- *Zermelo–Fraenkel set theory with choice and without power set ZFC^-* consists of axioms (i) to (vii), with the axiom of collection (ix') and the well-ordering principle (x').

In this course, our main metatheory will be ZF, and we will be explicit about the use of choice.

We say that a class X is *definable* over M if there exists a formula φ and sets $a_1, \dots, a_n \in M$ such that for all $z \in M$, we have $z \in X$ if and only if $\varphi(z, a_1, \dots, a_n)$. A class is *proper* over M if it is not a set in M .

Under suitable hypotheses, there is a countable transitive model M of ZFC. In this case, $|\mathbb{R} \cap M|$ is countable, so there exists a real v that is not in M . Hence, v is a proper class over M . However, it is not definable, and we cannot ‘talk about it’ in the language of set theory. The only proper classes that affect our theory are the definable ones.

In this course, we will assume that all mentioned classes are definable. We can then use formulas of the form

$$\exists C. (C \text{ is a class} \wedge \forall x \in C. \varphi)$$

by defining it to mean that there is a formula θ giving a class C satisfying $\forall x \in C. \theta$. For example, the universe class $V = \{x \mid x = x\}$, the Russell class $R = \{x \mid x \notin x\}$, and the class of ordinals Ord are all definable. Any set is a definable class. Classes are heavily dependent on the underlying model: if $M = 2$ then $\text{Ord} = 2 = M$, and if $M = 3 \cup \{1\}$ then $\text{Ord} = 3 \neq M$.

Suppose that M is a set model of ZF; that is, M is a set. Let \mathcal{D} be the collection of definable classes over M . Then one can show that \mathcal{D} is a set in our metatheoretic universe V , and (M, \mathcal{D}) is a model of a second-order version of ZF, known as *Gödel–Bernays set theory*.

1.3. Adding defined functions

Often in set theory, we use symbols such as $0, 1, \subseteq, \cap, \wedge, \forall$; they do not exist in our language.

Definition. Suppose that $\mathcal{L} \subseteq \mathcal{L}'$ and T is a set of sentences in \mathcal{L} . We say that P is a *defined n -ary predicate* symbol over T if there is a formula φ in \mathcal{L} such that

$$T \vdash \forall x_1, \dots, x_n. (P(x_1, \dots, x_n) \leftrightarrow \varphi(x_1, \dots, x_n))$$

Similarly, we say that f is a *defined n -ary function* symbol over T if there is a formula φ in \mathcal{L} such that

$$f(x_1, \dots, x_n) = y \text{ if and only if } T \vdash \varphi(x_1, \dots, x_n, y)$$

and

$$T \vdash \forall x_1, \dots, x_n. \exists! y. \varphi(x_1, \dots, x_n, y)$$

We say that a set of sentences T' of \mathcal{L}' is an *extension by definitions* of T over \mathcal{L} when $T' = T \cup S$ and $S = \{\varphi_s \mid s \in \mathcal{L}' \setminus \mathcal{L}\}$ and each φ_s is a definition of s in the language \mathcal{L} over T .

Commonly used symbols such as $0, 1, \subseteq, \cap, \mathcal{P}, \bigcup$ are defined over ZF.

Theorem. Suppose that $\mathcal{L} \subseteq \mathcal{L}'$, and that T is a set of \mathcal{L} -sentences and T' is an extension by definitions of T to \mathcal{L}' . Then

- (i) (conservativity) If φ is a sentence of \mathcal{L} , then $T \vdash \varphi \leftrightarrow T' \vdash \varphi$.
- (ii) (abbreviations) If φ is a formula of \mathcal{L}' , then there exists a formula $\hat{\varphi}$ of \mathcal{L} whose free variables are exactly those of φ , such that $T' \vdash \forall x. (\varphi \leftrightarrow \hat{\varphi})$.

Example. The intersection $a \cap b$ can be defined as the unique set c such that

$$\forall x. (x \in c \iff x \in a \wedge x \in b)$$

This definition makes sense only if there is a unique c satisfying this formula $\varphi(a, b, c)$. If

$$M = \{a, c, d, \{a\}, \{a, b\}, \{a, b, c\}, \{a, b, d\}\}$$

then it is easy to check that both $\{a\}$ and $\{a, b\}$ satisfy $\varphi(\{a, b, c\}, \{a, b, d\}, -)$, so intersection cannot be defined.

VII. Forcing and the Continuum Hypothesis

1.4. Absoluteness

It is often the case that definitions appear to give the same set regardless of which model we are working inside. For example, $\{x \mid x \neq x\}$ is the empty set in any model, and $\{x \mid x = a \vee x = b\}$ gives a pair set. Other definitions need not, for example $\mathcal{P}(\mathbb{N})$, which need not be the true power set in a given transitive model. To quantify this behaviour, we need to define what it means for φ to hold in an arbitrary structure; this concept is called *relativisation*.

Definition. The quantifier $\forall x \in a. \varphi$ is an abbreviation of $\forall x. (x \in a) \rightarrow \varphi$. Similarly, $\exists x \in a. \varphi$ is an abbreviation of $\exists x. (x \in a) \wedge \varphi$. Let W be a class; we define by recursion the *relativisation* φ^W of φ as follows.

$$\begin{aligned} (x \in y)^W &\equiv x \in y \\ (x = y)^W &\equiv x = y \\ (\varphi \vee \psi)^W &\equiv \varphi^W \vee \psi^W \\ (\neg \varphi)^W &\equiv \neg \varphi^W \\ (\exists x. \varphi)^W &\equiv \exists x \in W. \varphi^W \end{aligned}$$

One can easily show that

$$\begin{aligned} (\varphi \wedge \psi)^W &\equiv \varphi^W \wedge \psi^W \\ (\varphi \rightarrow \psi)^W &\equiv \varphi^W \rightarrow \psi^W \\ (\forall x. \varphi)^W &\equiv \forall x \in W. \varphi^W \end{aligned}$$

Proposition. Suppose that $M \subseteq N$ and M is a definable class over N . Then the relation $M \models \varphi$ is first-order expressible in N .

Proof. Suppose M is defined by θ , so

$$\forall z \in N. \theta(z) \leftrightarrow z \in M$$

We claim that $(N, \in) \models \varphi^M$ if and only if $(M, \in) \models \varphi$. We proceed by induction on the length of formulae. For example,

$$N \models (x \in y)^M \text{ iff } N \models x \in y \text{ and } x, y \in M \text{ iff } \theta(x), \theta(y), M \models x \in y$$

The case for equality is similar, and disjunction and negation are simple. Finally,

$$N \models (\exists x. \varphi(x))^M \text{ iff } N \models \exists x. x \in M \wedge \varphi^M(x)$$

which holds precisely when there is some $x \in N$ such that $N \models x \in M$ and $N \models \varphi^M(x)$, but $N \models x \in M$ if and only if $\theta(x)$, giving the result as required. \square

Thus, relativisation is a way to express truth in definable classes.

Definition. Suppose that $M \subseteq N$ are classes and $\varphi(u_1, \dots, u_n)$ is a formula. Then φ is called

(i) *upwards absolute* for M, N if

$$\forall x_1, \dots, x_n \in M. (\varphi^M(x_1, \dots, x_n) \rightarrow \varphi^N(x_1, \dots, x_n))$$

(ii) *downwards absolute* for M, N if

$$\forall x_1, \dots, x_n \in M. (\varphi^N(x_1, \dots, x_n) \rightarrow \varphi^M(x_1, \dots, x_n))$$

(iii) *absolute* for M, N if it is both upwards and downwards absolute, or equivalently,

$$\forall x_1, \dots, x_n \in M. (\varphi^M(x_1, \dots, x_n) \leftrightarrow \varphi^N(x_1, \dots, x_n))$$

If $N = V$, we simply say that φ is (upwards or downwards) absolute for M . If Γ is a set of formulas, we say that Γ is (upwards or downwards) absolute for M, N if and only if φ is (upwards or downwards) absolute for M, N for each $\varphi \in \Gamma$. Suppose T is a set of sentences and f is a defined function by φ . Then for $M \subseteq N$ models of T , we say that f is absolute for M, N precisely when φ is absolute for M, N .

Example. If $M \subseteq N$ both satisfy extensionality, then the empty set is absolute for M, N by the formula $\forall x \in a. (x \neq x)$. The power set of 2 is not absolute between 4 and V , because in 4, it has only two elements.

Example. $\varphi \leftrightarrow \psi$ does not imply $\varphi^M \leftrightarrow \psi^M$. Let $\varphi(v)$ be the statement $\forall x. (x \notin v)$; in ZF this defines \emptyset . Now, the following are two ways to express $0 \in z$.

$$\psi(z) \equiv \exists y. (\varphi(y) \wedge y \in z); \quad \theta(z) \equiv \forall y. (\varphi(y) \rightarrow y \in z)$$

Note that if there exists a unique y such that $\varphi(y)$, then these are equivalent. However, this is often not the case, for example if

$$a = 0; \quad b = \{0\}; \quad c = \{\{0\}\}; \quad M = \{a, b, c\}$$

then $\varphi^M(a)$ holds, so $\psi^M(b)$, but $\varphi^M(c)$ also holds, so $\theta^M(b)$ fails.

The main obstacle to absoluteness for basic statements turns out to be transitivity of the model.

Definition. Given classes $M \subseteq N$, we say that M is *transitive* in N if

$$\forall x, y \in N. (x \in M \wedge y \in x \rightarrow y \in M)$$

1.5. The Lévy hierarchy

Definition. The class of formulas Δ_0 is the smallest class Γ closed under the following conditions.

VII. Forcing and the Continuum Hypothesis

- (i) if φ is atomic, $\varphi \in \Gamma$ (that is, $(v_i \in v_j) \in \Gamma$ and $(v_i = v_j) \in \Gamma$);
- (ii) if $\varphi, \psi \in \Gamma$, then $\varphi \vee \psi \in \Gamma$ and $\neg\varphi \in \Gamma$; and
- (iii) if $\varphi \in \Gamma$, then $(\forall v_i \in v_j. \varphi) \in \Gamma$ and $(\exists v_i \in v_j. \varphi) \in \Gamma$.

That is, Δ_0 is the class of formulas generated from atomic formulas by Boolean operations and bounded quantification.

Definition. We proceed by induction to define Σ_n and Π_n as follows.

- (i) $\Sigma_0 = \Pi_0 = \Delta_0$;
- (ii) if φ is Π_{n-1} then $\exists v_i. \varphi$ is Σ_n ;
- (iii) if φ is Σ_{n-1} then $\forall v_i. \varphi$ is Π_n .

Example. The formula $\forall v_1. \exists v_2. \forall v_3. (v_4 = v_3)$ is Π_3 . But $(\forall v_1. v_1 = v_2) \wedge v_3 = v_4$ is not Π_n or Σ_n for any n .

Definition. Given an \mathcal{L}_\in -theory T , let Σ_n^T be the class of formulas Γ such that for any $\varphi \in \Gamma$, there exists $\psi \in \Sigma_n$ such that $T \vdash \varphi \leftrightarrow \psi$. We define Π_n^T analogously. A formula is in Δ_n^T if there exists $\psi \in \Sigma_n$ and $\theta \in \Pi_n$ such that $T \vdash \varphi \leftrightarrow \psi$ and $T \vdash \varphi \leftrightarrow \theta$.

Note that Δ_n only makes much sense with respect to some theory T for $n > 0$.

Lemma. If φ and ψ are in Σ_n^{ZF} , then so are

$$\exists v. \varphi; \quad \varphi \vee \psi; \quad \varphi \wedge \psi; \quad \exists v_i \in v_j. \varphi; \quad \forall v_i \in v_j. \varphi$$

If φ is in Σ_n^{ZF} , then $\neg\varphi$ is in Π_n^{ZF} . Further, for every φ , there exists n such that φ is in Σ_n^{ZF} , and if φ is in Σ_n^{ZF} , then φ is in Σ_m^{ZF} for all $m \geq n$.

Remark. $\exists x_1. \forall x_2. \exists x_3. \forall y. (y \in v \rightarrow v \neq y)$ is Σ_4 , but is logically equivalent to the statement $\forall y \in v. v \neq y$, which is Σ_0 . The fact that Σ_n^{ZF} is closed under bounded quantification depends on the axiom of collection. In particular, in Zermelo set theory, there is a Σ_1^{Z} formula φ such that $\forall x \in a. \varphi$ is not Σ_1^{Z} . In intuitionistic logic, these classes are very badly behaved; for instance, we could have a Π_1^T formula φ such that $\neg\varphi$ is not Σ_1^T .

We can now show absoluteness for Δ_0 formulas between transitive models.

Theorem. Let M be transitive in N and $M \subseteq N$, and let $\varphi(\mathbf{u})$ be a Δ_0 -formula. Then, for any $\mathbf{a} \in M$,

$$M \models \varphi(\mathbf{a}) \text{ if and only if } N \models \varphi(\mathbf{a})$$

Proof. We prove this by induction on the class Δ_0 . The cases of atomic formulas and propositional connectives are immediate, so it suffices to show the result for $\exists x \in a. \varphi$ where φ is absolute between M and N . Suppose $M \models \exists x \in a. \varphi(x)$, so there exists $b \in M$ such that $M \models b \in a \wedge \varphi(b)$. Then we also have $N \models b \in a \wedge \varphi(b)$ by absoluteness of φ , as required. Conversely, suppose $N \models \exists x \in a. \varphi(x)$, so there exists $b \in N$ such that $N \models b \in a \wedge \varphi(b)$. Since M is transitive in N , we obtain $b \in M$, so $M \models b \in a \wedge \varphi(b)$ by absoluteness of φ . \square

Proposition. The following are Δ_0^{ZF} , and therefore absolute between transitive models.

- (i) $x \subseteq y$;
- (ii) $a = \{x, y\}$ (the unordered pair);
- (iii) $a = \langle x, y \rangle$ (the ordered pair);
- (iv) $a = x \times y$;
- (v) $a = \bigcup b$;
- (vi) a is a transitive set;
- (vii) $x = \emptyset$;
- (viii) r is a relation;
- (ix) r is a function;
- (x) r is a relation with domain a and range b ;
- (xi) x is the pointwise image of r on a , denoted $r''a = \{y \mid \exists x \in a. \langle x, y \rangle \in r\}$;
- (xii) $r|_a$.

Remark. The following are not absolute between transitive models, and thus not Δ_0^{ZF} .

- (i) the cofinality function $\alpha \mapsto \text{cf}(\alpha)$;
- (ii) being a cardinal;
- (iii) ω_1 ;
- (iv) $y = \mathcal{P}(x)$.

Lemma. The statement that a given set a is finite is Δ_1^{ZF} .

Proposition. Let M be transitive in N and $M \subseteq N$. Then Σ_1 formulas are upwards absolute between M and N , and Π_1 formulas are downwards absolute between M and N .

Corollary. Δ_1^{ZF} formulas are absolute between transitive models.

Lemma. (ZF) The statement that α is an ordinal is absolute.

Proof. First, note that α is an ordinal in ZF if and only if it is a transitive set of transitive sets. This can be written as

$$(\forall \beta \in \alpha. \forall \gamma \in \beta. \gamma \in \alpha) \wedge (\forall \beta \in \alpha. \forall \gamma \in \beta. \forall \delta \in \gamma. \delta \in \beta)$$

which is Δ_0 , as required. □

We can give a slightly better rephrasing of this lemma.

Lemma. The statement that r is a strict total ordering of a is Δ_0 .

VII. Forcing and the Continuum Hypothesis

Proof. The statement that r is a transitive relation on a is that

$$\forall xyz \in a. (\langle x, y \rangle \in r \wedge \langle y, z \rangle \in r \rightarrow \langle x, z \rangle \in r)$$

Trichotomy is

$$\forall xy \in a. (\langle x, y \rangle \in r \vee \langle y, x \rangle \in r \vee x = y)$$

Irreflexivity is

$$\forall x \in a. \langle x, x \rangle \notin r$$

□

Corollary. The statement that x is a transitive set totally ordered by \in is Δ_0 , and thus being an ordinal is Δ_0 .

Lemma. (ZF) The statement that r is well-founded on a is Δ_1^{ZF} .

Proof. The Π_1 formula is

$$r \text{ is a relation on } a \wedge [\forall X. (\exists x \in X. (z = z) \wedge X \subseteq a) \rightarrow \exists x \in X. \forall y \in X. \langle y, z \rangle \notin r]$$

For the Σ_1 formula, we first show that a relation is well-founded on a if and only if there exists a function $a \rightarrow \text{Ord}$ such that $\langle y, x \rangle \in r$ implies $f(y) \in f(x)$. Suppose r is well-founded; we then define $f : a \rightarrow \text{Ord}$ by $f(x) = \sup \{f(y) + 1 \mid \langle y, x \rangle \in r\}$, and one can show that this satisfies the required property. For the other direction, let $X \subseteq a$ be a nonempty subset, and consider the pointwise image $f''X$. This has a minimal element α , then for any $z \in X$, if $f(z) = \alpha$ then for all $y \in X$, we have $f(y) \geq \alpha$, so $\langle y, z \rangle \notin r$. We then define well-foundedness with a Σ_1 formula as follows.

$$\exists f. (f \text{ is a function} \wedge \forall u \in \text{ran } f. (u \in \text{Ord}) \wedge \forall xy \in a. (\langle y, x \rangle \in r \rightarrow f(y) \in f(x)))$$

□

Proposition. The following are Δ_0^{ZF} .

- (i) x is a limit ordinal;
- (ii) x is a successor ordinal;
- (iii) x is a finite ordinal;
- (iv) $x = \omega$;
- (v) $x = n$ for any finite ordinal n .

Proposition. The following are Π_1^{ZF} and hence downwards absolute between transitive models.

- (i) κ is a cardinal;
- (ii) κ is regular;

- (iii) κ is a limit cardinal;
- (iv) κ is a strong limit cardinal.

Lemma. (ZF) Let W be a nonempty transitive class. Then the axioms of extensionality, empty set, and foundation all hold in W .

Proof. For extensionality, the relativisation of

$$\forall x. \forall y. (\forall z. (z \in x \leftrightarrow z \in y) \rightarrow x = y)$$

to W is

$$\forall x \in W. \forall y \in W. (\forall z \in W. (z \in x \leftrightarrow z \in y) \rightarrow x = y)$$

Suppose $x \in W, y \in W$, but $x \neq y$. Then by extensionality in the metatheory, without loss of generality we can fix $z \in x$ with $z \notin y$. But since W is transitive, we must have $z \in W$, contradicting $x = y$, as required.

As W is nonempty, we can use foundation to fix $x \in W$ such that $x \cap W = \emptyset$. Since W is transitive, $x \subseteq W$, and therefore $x = \emptyset \in W$. Moreover, the statement that $x = \emptyset$ is Δ_0 and therefore absolute.

□

Lemma. (ZF) Let W be a transitive class. Then

- (i) if for any pair $x, y \in W$, the real pair set $\{x, y\}$ lies in W , then the axiom of pairing holds in W ;
- (ii) if for any set $x \in W$, the union $\bigcup x$ lies in W , then the axiom of union holds in W ;
- (iii) if $\omega \in W$, then the axiom of infinity holds in W ;
- (iv) if, for every formula φ with free variables in $\{x, a, v_1, \dots, v_n\}$, we have

$$\forall a, v_1, \dots, v_n \in W. \{x \in a \mid \varphi^W(x, a, v_1, \dots, v_n)\} \in W$$

then the axiom of separation holds in W ;

- (v) if, for every formula φ with free variables in $\{x, y, a, v_1, \dots, v_n\}$, for all $a, v_1, \dots, v_n \in W$, if

$$\forall x \in a. \exists! y \in W. \varphi^W(x, y, a, v_1, \dots, v_n)$$

then

$$\exists b \in W. \{y \mid \exists x \in a. \varphi^W(x, y, a, v_1, \dots, v_n)\} \subseteq b$$

then the axiom of replacement holds in W ;

- (vi) if, for every $a \in W$, there exists $b \in W$ such that $\mathcal{P}(a) \cap W = b$, then the axiom of power set holds in W .

Corollary. (ZF) If W is a nonempty transitive class satisfying the conditions of the previous lemma, it is a model of ZF.

VII. Forcing and the Continuum Hypothesis

1.6. Transfinite recursion

Definition. A relation R is *set-like* on a class A if for all $x \in A$, the collection of R -predecessors of x is a set.

Example. \in is set-like on V , but \ni is not set-like on V .

Let A be a class, and let φ be such that $A = \{x \mid \varphi(x)\}$. Then $A^W = \{x \mid \varphi^W(x)\}$. We say that A is absolute for W if $A^W = A \cap W$. Viewing a class relation $R \subseteq V \times V$ as a collection of ordered pairs $\{\langle x, y \rangle \mid \psi(x, y)\}$, we have $R^W = \{\langle x, y \rangle \mid \psi^W(x, y)\}$, and say that R is absolute for W if $R^W = R \cap W^2$. Observe that if R is a class function, we can only refer to the *function* R^W if we first check that $(\forall x. \exists! y. \varphi(x, y))^W$. In this case, we have $R^W : W \rightarrow W$, and we say that R is an absolute function for W iff $R^W = R|_W$.

We briefly recall the transfinite recursion theorem.

Theorem. Let R be a relation which is well-founded and set-like on a class A . Let $F : A \times V \rightarrow V$ be a class function. Given $x \in A$, let $\text{pred}(A, x, R) = \{y \in A \mid y R x\}$ be the set of R -predecessors of x in A . Then there is a unique function $G : A \rightarrow V$ such that for all $x \in A$,

$$G(x) = F\left(x, G\Big|_{\text{pred}(A, x, R)}\right)$$

We now prove the absoluteness of transfinite recursion.

Theorem. Let R be a relation which is well-founded and set-like on a class A . Let $F : A \times V \rightarrow V$ be a class function, and let $G : A \rightarrow V$ be the unique function given by applying transfinite recursion to F . Suppose that W is a transitive model of ZF, and suppose that the following hold.

- (i) A and F are absolute for W ;
- (ii) R is absolute for W and $(R \text{ is set-like on } A)^W$;
- (iii) for all $x \in W$, $\text{pred}(A, x, R) \subseteq W$.

Then G is absolute for W .

Proof. By absoluteness, $A^W = A \cap W$ and $R^W = R \cap W^2$. Hence, every nonempty subset of A^W has an R^W -minimal element. In particular, $(R \text{ is well-founded on } A)^W$. We can then apply transfinite recursion in W to define a unique function $G^W : A^W \rightarrow W$ such that for all $x \in A^W$,

$$G^W(x) = F^W\left(x, G^W\Big|_{\text{pred}^W(A^W, x, R^W)}\right)$$

To prove absoluteness for G , it suffices to show that $G^W = G|_{A^W}$. We show this by transfinite induction in W . Suppose that for all $y R x$, we have $G^W(y) = G(y)$. By absoluteness, (iii),

and the inductive hypothesis, we obtain

$$G^W(x) = F^W\left(x, G^W \Big|_{\text{pred}^W(A^W, x, R^W)}\right) = F\left(x, G \Big|_{\text{pred}(A, x, R)}\right) = G(x)$$

□

Corollary. The following are absolute for transitive models of ZFC:

- (i) the rank function;
- (ii) the transitive closure of a set;
- (iii) the addition and multiplication operations of ordinal arithmetic.

1.7. The reflection theorem

In this subsection, we will not use choice.

Recall the Tarski–Vaught test: if \mathcal{M} is a substructure of \mathcal{N} with universes M and N respectively, then the following two statements are equivalent.

- (i) \mathcal{M} is an elementary substructure of \mathcal{N} ;
- (ii) for any formula $\varphi(v, \mathbf{w})$ and $\mathbf{a} \in M$, if there exists $b \in N$ such that $\mathcal{N} \models \varphi(b, \mathbf{a})$, then there exists $c \in M$ such that $\mathcal{M} \models \varphi(c, \mathbf{a})$.

Definition. A finite list of formulas $\boldsymbol{\varphi} = \varphi_1, \dots, \varphi_n$ is said to be *subformula closed* if every subformula of the φ_i is contained on the list.

We can now state a version of the Tarski–Vaught test for classes.

Lemma. Let $\boldsymbol{\varphi}$ be a subformula closed list of formulas, and suppose $W \subseteq Z$ are nonempty classes. Then the following two statements are equivalent.

- (i) each formula in $\boldsymbol{\varphi}$ is absolute for W and Z ;
- (ii) whenever φ_i is of the form $\exists x. \varphi_j(x, \mathbf{y})$ where the free variables of φ_j are equal to x or contained in \mathbf{y} , then

$$\forall \mathbf{y} \in W. (\exists x \in Z. \varphi_j^Z(x, \mathbf{y}) \rightarrow \exists x \in W. \varphi_j^Z(x, \mathbf{y}))$$

Proof. (i) implies (ii). Suppose that each formula in $\boldsymbol{\varphi}$ is absolute. Let φ_i be of the form $\exists x. \varphi_j(x, \mathbf{y})$, and fix $\mathbf{y} \in W$. Then $\varphi_i^Z(\mathbf{y})$ is $\exists x \in Z. \varphi_j^Z(x, \mathbf{y})$. If this holds, by absoluteness $\varphi_i^W(\mathbf{y})$ holds, so there is $x \in W$ such that $\varphi_j^W(x, \mathbf{y})$. Finally, $W \subseteq Z$ and absoluteness of φ_j gives $\exists x \in W. \varphi_j^Z(x, \mathbf{y})$.

(ii) implies (i). We show this by induction on the length of φ_i . The result if φ_i is atomic or of the form $\varphi_j \vee \varphi_k$ or $\neg \varphi_j$ is immediate. Suppose φ_i is of the form $\exists x. \varphi_j(x, \mathbf{y})$, and fix $\mathbf{y} \in W$. Then $\varphi_i^Z(\mathbf{y})$ is equivalent to the statement $\exists x \in Z. \varphi_j^Z(x, \mathbf{y})$. By (ii), this gives

VII. Forcing and the Continuum Hypothesis

$\exists x \in W. \varphi_j^Z(x, \mathbf{y})$. Since $W \subseteq Z$, the reverse implication is trivial. But $\exists x \in W. \varphi_j^Z(x, \mathbf{y})$ is equivalent to the statement that $\varphi_i^W(\mathbf{y})$ holds, as required. \square

Theorem (reflection theorem). Let W be a nonempty class, and suppose that there is a class function F_W such that for any ordinal α , $F_W(\alpha) = W_\alpha \in V$. Suppose that

- (i) if $\alpha < \beta$, then $W_\alpha \subseteq W_\beta$;
- (ii) if λ is a limit ordinal, then $W_\lambda = \bigcup_{\alpha < \lambda} W_\alpha$;
- (iii) $W = \bigcup_{\alpha \in \text{Ord}} W_\alpha$.

Then for any finite list of formulas $\varphi = \varphi_1, \dots, \varphi_n$, ZF proves that for every α there is a limit ordinal $\beta > \alpha$ such that the φ_i are absolute between W_β and W .

One example of such a class function is $W_\alpha = V_\alpha$.

Corollary (Montague–Lévy reflection). For any finite list of formulas $\varphi = \varphi_1, \dots, \varphi_n$, ZF proves that for every α there is a limit ordinal $\beta > \alpha$ such that the φ_i are absolute for V_β .

We now prove the reflection theorem.

Proof. Let $\varphi = \varphi_1, \dots, \varphi_n$ be a finite list of formulas. By extending the list and taking logical equivalences if necessary, we will assume that this list is subformula-closed and that there are no universal quantifiers. For $i \leq n$, we will define a function $G_i : \text{Ord} \rightarrow \text{Ord}$ as follows. If φ_i is of the form $\exists x. \varphi_j(x, \mathbf{y})$ where \mathbf{y} is a tuple of length k_i , we will define a function $F_i : W^{k_i} \rightarrow \text{Ord}$ by setting

$$F_i(\mathbf{y}) = \begin{cases} 0 & \text{if } \neg \exists x \in W. \varphi_j^W(x, \mathbf{y}) \\ \eta & \text{where } \eta \text{ is the least ordinal such that } \exists x \in W_\eta. \varphi_j^W(x, \mathbf{y}) \end{cases}$$

We set

$$G_i(\delta) = \sup \{ F_i(\mathbf{y}) \mid \mathbf{y} \in W_\delta^{k_i} \}$$

If φ_i is not of this form, we set $G_i(\delta) = 0$ for all δ . Finally, we let

$$K(\delta) = \max \{ \delta + 1, G_1(\delta), \dots, G_n(\delta) \}$$

Note that the F_i work in an analogous way to Skolem functions, but does not require choice. The F_i are well-defined, and, using replacement in V , since W_δ is a set, $F_i'' W_\delta^{k_i}$ is also a set in V , so G_i and K are both defined and take values in Ord . Also, G_i is monotone: if $\delta \leq \delta'$ then $G_i(\delta) \leq G_i(\delta')$.

We claim that for every α there is a limit ordinal $\beta > \alpha$ such that for all $\delta < \beta$ and $i \leq n$, we have $G_i(\delta) < \beta$; that is, β is closed under this process of finding witnesses. Set $\lambda_0 = \alpha$ and let $\lambda_{t+1} = K(\lambda_t)$. Then we set $\beta = \sup_{t \in \omega} \lambda_t$, which is a limit ordinal as it is the supremum of a strictly increasing sequence of ordinals. If $\delta < \beta$, then $\delta < \lambda_t$ for some t , so $G_i(\delta) \leq G_i(\lambda_t)$ by monotonicity, but $G_i(\lambda_t) \leq K(\lambda_t) = \lambda_{t+1} < \beta$ as required.

1. Set theoretic preliminaries

To complete the theorem, it suffices to consider φ_i of the form $\exists x. \varphi_j(x, \mathbf{y})$ by the Tarski–Vaught test for classes above. Fix $\mathbf{y} \in W_\beta$, and suppose there exists $x \in W$ such that $\varphi_j^W(x, \mathbf{y})$. Since β is a limit ordinal and \mathbf{y} is a finite sequence in W_β , we must have $\mathbf{y} \in W_\gamma$ for some $\gamma < \beta$. Thus

$$0 < F_i(\mathbf{y}) \leq G_i(\gamma) < \beta$$

so by construction, there exists a witness $x \in W_\beta$ such that $\varphi_j^W(x, \mathbf{y})$. Hence φ is absolute between W_β and W as required. \square

Remark. This is a theorem scheme; for every choice of formulas φ , it is a theorem of ZF that φ are absolute for some V_β . We cannot prove that for every collection of formulas φ , for all ordinals α there exists $\beta > \alpha$ such that φ is absolute for W_β, W . Note that even if φ is absolute for W_β and W , we need not have φ^{W_β} .

If φ is any finite list of axioms of ZF, then there are arbitrarily large β such that φ holds in V_β . If β is a limit ordinal, $V_\beta \models Z(C)$, so we may restrict our attention to instances of replacement.

Corollary. Let T be an extension of ZF in \mathcal{L}_\in , and let $\varphi_1, \dots, \varphi_n$ be a finite list of axioms from T . Then T proves that for every α there exists $\beta > \alpha$ such that $(\bigwedge_{i=1}^n \varphi_i)^{V_\beta}$.

Corollary. (ZFC) Let W be a class and let $\varphi_1, \dots, \varphi_n$ be a finite list of formulas in \mathcal{L}_\in . Then ZFC proves that for every transitive $x \subseteq W$, there exists some transitive $y \supseteq x$ such that the φ_i are absolute between y and W , and $|y| \leq \max\{w, |x|\}$.

Taking $x = \omega$ and $W = V$ gives the following result.

Corollary. Let T be any set of sentences in \mathcal{L}_\in such that $T \vdash \text{ZFC}$. Let $\varphi_1, \dots, \varphi_n \in T$. Then T proves that there is a transitive set y of cardinality \aleph_0 such that $(\bigwedge_{i=1}^n \varphi_i)^y$.

Corollary. Let T be any consistent set of sentences in \mathcal{L}_\in such that $T \vdash \text{ZF}$. Then T is not finitely axiomatisable. That is, for any finite set of sentences Γ in \mathcal{L}_\in such that $T \vdash \Gamma$, there exists a sentence φ such that $T \vdash \varphi$ but $\Gamma \not\vdash \varphi$.

This only holds for first-order theories without classes; for example, Gödel–Bernays set theory is finitely axiomatisable.

Proof. Let $\varphi_1, \dots, \varphi_n$ be a set of sentences such that $T \vdash \bigwedge_{i=1}^n \varphi_i$. Suppose that $\bigwedge_{i=1}^n \varphi_i$ proves every axiom of T . By reflection, T proves that for every α there is $\beta > \alpha$ such that the φ_i hold in V_β if and only if they hold in V . Since they hold in V , they must hold in some V_β . Fix β_0 to be the least ordinal such that $\bigwedge_{i=1}^n \varphi_i^{V_{\beta_0}}$. Then all of the axioms of T hold in V_{β_0} , so $V_{\beta_0} \models T$. Since T extends ZF, our basic absoluteness results hold, so in particular, if $\alpha \in V_{\beta_0}$ then

$$V_\alpha^{V_{\beta_0}} = V_\alpha \cap V_{\beta_0} = V_\alpha$$

VII. Forcing and the Continuum Hypothesis

So V_α is absolute for V_{β_0} . Note that T proves that there exists α such that $\bigwedge_{i=1}^n \varphi_i^V$, but as V_{β_0} satisfies every axiom of T , this must be true in V_{β_0} . That is, there must be $\alpha < \beta_0$ such that $\bigwedge_{i=1}^n \varphi_i^V$. This contradicts minimality of β_0 . \square

1.8. Cardinal arithmetic

In this subsection, we will use the axiom of choice. We recall the following basic definitions and results.

Definition. The *cardinality* of a set x , written $|x|$, is the least ordinal α such that there is a bijection $x \rightarrow \alpha$.

This definition only makes sense given the well-ordering principle.

Definition. The cardinal arithmetic operations are defined as follows. Let κ, λ be cardinals.

- (i) $\kappa + \lambda = |\{0\} \times \kappa \cup \{1\} \times \lambda|$;
- (ii) $\kappa \cdot \lambda = |\kappa \times \lambda|$;
- (iii) $\kappa^\lambda = |\kappa^\lambda|$, the cardinality of the set of functions $\lambda \rightarrow \kappa$;
- (iv) $\kappa^{<\lambda} = \sup \{\kappa^\alpha \mid \alpha < \lambda, \alpha \text{ a cardinal}\}$.

Theorem (Hessenberg). If κ, λ are infinite cardinals, then

$$\kappa + \lambda = \kappa \cdot \lambda = \max\{\kappa, \lambda\}$$

Lemma. If κ, λ, μ are cardinals, then

$$\kappa^{\lambda+\mu} = \kappa^\lambda \cdot \kappa^\mu; \quad (\kappa^\lambda)^\mu = \kappa^{\lambda \cdot \mu}$$

Definition. A map between ordinals $\alpha \rightarrow \beta$ is *cofinal* if $\sup \text{ran } f = \beta$. The *cofinality* of an ordinal γ , written $\text{cf}(\gamma)$, is the least ordinal that admits a cofinal map to γ . A limit ordinal γ is *singular* if $\text{cf}(\gamma) < \gamma$, and *regular* if $\text{cf}(\gamma) = \gamma$.

Remark. (i) Since the identity map is always cofinal, we have $\text{cf}(\gamma) \leq \gamma$.

(ii) $\omega = \text{cf}(\omega) = \text{cf}(\omega + \omega) = \text{cf}(\aleph_\omega)$.

(iii) $\text{cf}(\gamma) \leq |\gamma|$.

Theorem. Let γ be a limit ordinal. Then

- (i) if γ is regular, γ is a cardinal;
- (ii) the cardinal successor γ^+ is a regular cardinal;
- (iii) $\text{cf}(\text{cf}(\gamma)) = \text{cf}(\gamma)$, so $\text{cf}(\gamma)$ is regular;
- (iv) \aleph_α is regular whenever $\alpha = 0$ or a successor;

1. Set theoretic preliminaries

(v) if λ is a limit ordinal, $\text{cf}(\aleph_\lambda) = \text{cf}(\lambda)$.

Theorem. Let κ be a regular cardinal. If \mathcal{F} is a family of sets with $|\mathcal{F}| < \kappa$ and each $|X| < \kappa$ for $X \in \mathcal{F}$, then $|\bigcup \mathcal{F}| < \kappa$.

Proof. We show this by induction on $|\mathcal{F}| = \gamma < \kappa$. Suppose the claim holds for γ , and consider $\mathcal{F} = \{X_\alpha \mid \alpha < \gamma + 1\}$. Then, assuming the sets involved are infinite,

$$|\bigcup \mathcal{F}| = \left| \bigcup_{\alpha < \gamma} X_\alpha \cup X_\gamma \right| = \left| \bigcup_{\alpha < \gamma} X_\alpha \right| + |X_\gamma| = \max \left\{ \left| \bigcup_{\alpha < \gamma} X_\alpha \right|, |X_\gamma| \right\} < \kappa$$

Now suppose γ is a limit, and suppose the claim holds for all $\beta < \gamma$. Let $\mathcal{F} = \{X_\alpha \mid \alpha < \gamma\}$, and define $g : \gamma \rightarrow \kappa$ by

$$g(\beta) = \left| \bigcup_{\alpha < \beta} X_\alpha \right|$$

But κ is regular and $\gamma < \kappa$, so this map is not cofinal. Hence $g''\gamma = |\bigcup \mathcal{F}| < \kappa$. \square

We can generalise the notions of cardinal sum and product as follows.

Definition. Let $(\kappa_i)_{i \in I}$ be an indexed sequence of cardinals, and let $(X_i)_{i \in I}$ be a sequence of pairwise disjoint sets with $|X_i| = \kappa_i$ for all $i \in I$. Then the *cardinal sum* of (κ_i) is

$$\sum_{i \in I} \kappa_i = \left| \bigcup_{i \in I} X_i \right|$$

The *cardinal product* is

$$\prod_{i \in I} \kappa_i = \left| \prod_{i \in I} X_i \right|$$

where $\prod_{i \in I} X_i$ denotes the set of functions $f : I \rightarrow \bigcup_{i \in I} X_i$ such that $f(i) \in X_i$ for each i .

The following theorem generalises Cantor's diagonal argument.

Theorem (König's theorem). Let I be an indexing set, and suppose that $\kappa_i < \lambda_i$ for all $i \in I$. Then

$$\sum_{i \in I} \kappa_i < \prod_{i \in I} \lambda_i$$

Proof. Let $(B_i)_{i \in I}$ be a sequence of disjoint sets with $|B_i| = \lambda_i$, and let $B = \prod_{i \in I} B_i$. It suffices to show that for any sequence $(A_i)_{i \in I}$ of subsets of B such that for all $i \in I$, $|A_i| = \kappa_i$, then

$$\bigcup_{i \in I} A_i \neq B$$

Given such a sequence, we let S_i be the projection of A_i onto its i th coordinate.

$$S_i = \{f(i) \mid f \in A_i\}$$

VII. Forcing and the Continuum Hypothesis

Then by definition, $S_i \subseteq B_i$, and

$$|S_i| \leq |A_i| = \kappa_i < \lambda_i = |B_i|$$

Fix $t_i \in B_i \setminus S_i$. Finally, we define $g \in B$ by $g(i) = t_i$; by construction, we have $g \notin A_i$ for all i , so $g \in B$ but $g \notin \bigcup_{i \in I} A_i$. \square

Corollary. If $\kappa \geq 2$ and λ is infinite, then

$$\kappa^\lambda > \lambda$$

Proof.

$$\lambda = \sum_{\alpha < \lambda} 1 < \prod_{\alpha < \lambda} 2 = 2^\lambda \leq \kappa^\lambda$$

\square

Corollary. $\text{cf}(2^\lambda) > \lambda$.

Proof. Let $f : \lambda \rightarrow 2^\lambda$, we show that $|\bigcup f''\lambda| < 2^\lambda$. Since for all $i \in I$, we have $f(i) < 2^\lambda$, we deduce

$$|\bigcup f''\lambda| \leq \sum_{i < \lambda} |f(i)| < \prod_{i < \lambda} 2^\lambda = (2^\lambda)^\lambda = 2^{\lambda \cdot \lambda} = 2^\lambda$$

\square

Corollary. $2^{\aleph_0} \neq \kappa$ for any κ of cofinality \aleph_0 . In particular, $2^{\aleph_0} \neq \aleph_\omega$.

Corollary. $\kappa^{\text{cf}(\kappa)} > \kappa$ for every infinite cardinal κ .

We can prove very little in general about cardinal exponentiation given ZFC.

Definition. The *generalised continuum hypothesis* is the statement that $2^\kappa = \kappa^+$ for every infinite cardinal κ . Equivalently, $2^{\aleph_\alpha} = \aleph_{\alpha+1}$.

Under this assumption, we can show the following.

Theorem. (ZFC + GCH) Let κ, λ be infinite cardinals.

- (i) if $\kappa < \lambda$, then $\kappa^\lambda = \lambda^+$;
- (ii) if $\text{cf}(\kappa) \leq \lambda < \kappa$, then $\kappa^\lambda = \kappa^+$;
- (iii) if $\lambda < \text{cf}(\kappa)$, then $\kappa^\lambda = \kappa$.

When we construct models with certain properties of cardinal arithmetic, we will often want to start with a model satisfying GCH so that we have full control over cardinal exponentiation. Without this assumption, we know much less. The following theorems are essentially the only restrictions that we have on regular cardinals that are provable in ZFC.

Theorem. Let κ, λ be cardinals. Then

- (i) if $\kappa < \lambda$, then $2^\kappa \leq 2^\lambda$;
- (ii) $\text{cf}(2^\kappa) > \kappa$;
- (iii) if κ is a limit cardinal, then $2^\kappa = (2^{<\kappa})^{\text{cf}(\kappa)}$.

Theorem. Let κ, λ be infinite cardinals. Then

- (i) if $\kappa \leq \lambda$, then $\kappa^\lambda = 2^\lambda$;
- (ii) if $\mu < \kappa$ is such that $\mu^\lambda \geq \kappa$, then $\kappa^\lambda = \mu^\lambda$;
- (iii) if $\kappa > \lambda$ and $\mu^\lambda < \kappa$ for all $\mu < \kappa$, then
 - (a) if $\text{cf}(\kappa) > \lambda$, then $\kappa^\lambda = \kappa$;
 - (b) if $\text{cf}(\kappa) \leq \lambda$, then $\kappa^\lambda = \kappa^{\text{cf}(\kappa)}$.

Theorem (Silver). Suppose that κ is a singular cardinal such that $\text{cf}(\kappa) > \aleph_0$ and $2^\alpha = \alpha^+$ for all $\alpha < \kappa$. Then $2^\kappa = \kappa^+$.

This theorem therefore states that the generalised continuum hypothesis cannot first break at a singular cardinal with cofinality larger than \aleph_0 .

Remark. It is consistent (relative to large cardinals, such as a measurable cardinal) to have $2^{\aleph_n} = \aleph_{n+1}$ for all $n \in \omega$, but $2^{\aleph_\omega} = \aleph_{\omega+2}$.

Theorem (Shelah). Suppose that $2^{\aleph_n} < \aleph_\omega$ for all $n \in \omega$, so \aleph_ω is a strong limit cardinal. Then $2^{\aleph_\omega} < \aleph_{\omega_4}$.

It is not known if this bound can be improved, but it is conjectured that $2^{\aleph_\omega} < \aleph_{\omega_1}$.

2. Constructibility

In this section, we will prove

$$\text{Con}(\text{ZF}) \rightarrow \text{Con}(\text{ZFC} + \text{GCH})$$

2.1. Definable sets

Recall that the V_α hierarchy has the property that $V_{\alpha+1} = \mathcal{P}(V_\alpha)$. We will construct a universe L in which we restrict to the ‘nice’ subsets.

Definition. A set x is said to be *definable* over (M, \in) if there exist $a_1, \dots, a_n \in M$ and a formula φ such that

$$x = \{z \in M \mid (M, \in) \models \varphi(z, a_1, \dots, a_n)\}$$

We write

$$\text{Def}(M) = \{x \subseteq M \mid x \text{ is definable over } M\}$$

Remark. (i) $M \in \text{Def}(M)$.

(ii) $M \subseteq \text{Def}(M) \subseteq \mathcal{P}(M)$.

This definition involves a quantification over infinitely many formulas, so is not yet fully formalised. One method to do this is to code formulas as elements of V_ω , called *Gödel codes*. We can then use Tarski’s *satisfaction relation* to define a formula Sat , and can then prove

$$\text{Sat}(M, \in, \ulcorner \varphi \urcorner, x_1, \dots, x_n) \leftrightarrow (M, \in) \models \varphi(x_1, \dots, x_n)$$

where $\ulcorner \varphi \urcorner \in V_\omega$ is the Gödel code for φ . We will later use a different method to formalise it, but for now we will assume that this is well-defined.

2.2. Defining the constructible universe

We define the L_α hierarchy by transfinite recursion as follows.

$$L_0 = \emptyset; \quad L_{\alpha+1} = \text{Def}(L_\alpha); \quad L_\lambda = \bigcup_{\alpha < \lambda} L_\alpha; \quad L = \bigcup_{\alpha \in \text{Ord}} L_\alpha$$

Lemma. For any ordinals α, β ,

- (i) if $\beta \leq \alpha$ then $L_\beta \subseteq L_\alpha$;
- (ii) if $\beta < \alpha$ then $L_\beta \in L_\alpha$;
- (iii) L_α is transitive;
- (iv) the ordinals of L_α are precisely α ;
- (v) L is transitive and $\text{Ord} \subseteq L$.

Definition. Let T be a set of axioms in \mathcal{L}_\in , and let W be a class. Then W is called an *inner model* of T if

- (i) W is a transitive class;
- (ii) $\text{Ord} \subseteq W$;
- (iii) T^W is true; that is, for every formula φ in T , we have φ^W .

Theorem. L is an inner model of ZF.

This is a theorem scheme; for every axiom of ZF, we can prove its relativisation to L .

Proof. By the previous lemma, it suffices to check that ZF^L holds.

- Since L is transitive, L satisfies extensionality and foundation.
- For the axiom of empty set, we use the fact that $\emptyset^L = \emptyset = L_0 \in L$.
- For pairing, given $a, b \in L$, we must show $\{a, b\} \in L$. Fix α such that $a, b \in L_\alpha$. Then

$$\{a, b\} = \{x \in L_\alpha \mid (L_\alpha, \in) \models x = a \vee x = b\} \in \text{Def}(L_\alpha)$$

- For union, let $a \in L_\alpha$. By transitivity, $\bigcup a \subseteq L_\alpha$. Then

$$\bigcup a = \{x \in L_\alpha \mid (L_\alpha, \in) \models \exists z. (z \in a \wedge x \in z)\} \in \text{Def}(L_\alpha)$$

- For infinity, note that

$$\omega = \{n \in L_\omega \mid (L_\omega, \in) \models n \in \text{Ord}\} \in \text{Def}(L_\omega)$$

- Consider separation. Let φ be a formula, and let $a, \mathbf{u} \in L_\alpha$. We claim that

$$b = \{x \in a \mid \varphi^L(x, \mathbf{u})\} \in L$$

This implicitly uses the fact that L is definable. Using the reflection theorem, there is $\beta > \alpha$ such that

$$\text{ZF} \vdash \forall x \in L_\beta. (\varphi^L(x, \mathbf{u}) \leftrightarrow \varphi^{L_\beta}(x, \mathbf{u}))$$

Moreover, $\varphi^{L_\beta}(x, \mathbf{u})$ holds if and only if $(L_\beta, \in) \models \varphi(x, \mathbf{u})$. We thus obtain

$$\{x \in a \mid \varphi^L(x, \mathbf{u})\} = \{x \in a \mid \varphi^{L_\beta}(x, \mathbf{u})\} = \{x \in L_\beta \mid (L_\beta, \in) \models \varphi(x, \mathbf{u}) \wedge x \in a\} \in \text{Def}(L_\beta)$$

- We now consider replacement. It suffices to show that if $a \in L$ and $f : a \rightarrow L$ is a definable function, then there exists $\gamma \in \text{Ord}$ such that $f''a \subseteq L_\gamma$, since then we can use separation. First, observe that for every $x \in a$, there exists $\beta \in \text{Ord}$ such that $f(x) \in L_\beta$. Using replacement in V , there exists an ordinal γ such that for all $x \in a$, there exists $\beta < \gamma$ such that $f(x) \in L_\beta$. As $L_\beta \subseteq L_\gamma$, we thus obtain for all $x \in a$ that $f(x) \in L_\gamma$.

VII. Forcing and the Continuum Hypothesis

- Finally, consider the axiom of power set. It suffices to prove that if $x \in L$ then $\mathcal{P}(x) \cap L \in L$. Take $x \in L$. Using replacement in V , we can fix an ordinal γ such that $\mathcal{P}(x) \cap L \subseteq L_\gamma$. Then

$$\mathcal{P}(x) \cap L = \{z \in L_\gamma \mid (L, \in) \models z \subseteq x\} \in \text{Def}(L_\gamma)$$

□

2.3. Gödel functions

We will now formally define L . For clarity, we will define the ordered triple $\langle a, b, c \rangle$ to be $\langle a, \langle b, c \rangle \rangle$.

Definition. The *Gödel functions* are the following collection of functions on two variables.

- (i) $\mathcal{F}_1(x, y) = \{x, y\}$;
- (ii) $\mathcal{F}_2(x, y) = \bigcup x$;
- (iii) $\mathcal{F}_3(x, y) = x \setminus y$;
- (iv) $\mathcal{F}_4(x, y) = x \times y$;
- (v) $\mathcal{F}_5(x, y) = \text{dom } x = \{\pi_1(z) \mid z \in x \wedge z \text{ is an ordered pair}\}$;
- (vi) $\mathcal{F}_6(x, y) = \text{ran } x = \{\pi_2(z) \mid z \in x \wedge z \text{ is an ordered pair}\}$;
- (vii) $\mathcal{F}_7(x, y) = \{\langle u, v, w \rangle \mid \langle u, v \rangle \in x, w \in y\}$;
- (viii) $\mathcal{F}_8(x, y) = \{\langle u, w, v \rangle \mid \langle u, v \rangle \in x, w \in y\}$;
- (ix) $\mathcal{F}_9(x, y) = \{\langle v, u \rangle \in y \times x \mid u = v\}$;
- (x) $\mathcal{F}_{10}(x, y) = \{\langle v, u \rangle \in y \times x \mid u \in v\}$.

Proposition. The following can all be written as a finite combination of Gödel functions (i)–(vii).

$$\{x\}; \quad x \cup y; \quad x \cap y; \quad \langle x, y \rangle; \quad \langle x, y, z \rangle$$

Proposition. For every $i \in \{1, \dots, 10\}$, the statement $z = \mathcal{F}_i(x, y)$ can be written using a Δ_0 formula. Hence, these formulas are absolute.

Lemma (Gödel normal form). For every Δ_0 formula $\varphi(x_1, \dots, x_n)$ with free variables contained in $\{x_1, \dots, x_n\}$, there is a term \mathcal{F}_φ built from the symbols $\mathcal{F}_1, \dots, \mathcal{F}_{10}$ such that

$$\text{ZF} \vdash \forall a_1, \dots, a_n. \mathcal{F}_\varphi(a_1, \dots, a_n) = \{\langle x_n, \dots, x_1 \rangle \in a_n \times \dots \times a_1 \mid \varphi(x_1, \dots, x_n)\}$$

Remark. (i) The reversed order of the free variables is done purely for technical reasons.

- (ii) \mathcal{F}_2 will correspond to disjunction for Δ_0 formulas, intersection will correspond to conjunction, \mathcal{F}_3 will give negation, and \mathcal{F}_9 and \mathcal{F}_{10} will give atomic formulas.

- (iii) \mathcal{F}_7 and \mathcal{F}_8 will deal with ordered n -tuples. For example, the triple $\langle x_1, x_2, x_3 \rangle$ is formed using x_1 and $\langle x_2, x_3 \rangle$, but it cannot be formed using $\langle x_1, x_2 \rangle$ and x_3 without \mathcal{F}_7 or \mathcal{F}_8 .

Proof. We show this by induction on the class Δ_0 . We call a formula φ a *termed formula* if the conclusion of the lemma holds for φ ; we aim to show that every Δ_0 -formula is a termed formula. We will only use the logical symbols $\wedge, \vee, \neg, \exists$, and the only occurrence of existential quantification will be in formulas of the form

$$\varphi(x_1, \dots, x_n) \equiv \exists x_{n+1} \in x_j. \psi(x_1, \dots, x_{n+1})$$

where $j \leq m \leq n$. For example, we allow $\exists x_3 \in x_1. (x_1 \in x_2 \wedge x_3 = x_1)$, but we disallow $\exists x_1 \in x_2. \psi$ and $\exists x_3 \in x_1. (x_3 \in x_2 \wedge \exists x_4 \in x_1. \psi)$. Every Δ_0 -formula is equivalent to one of this form. We allow for dummy variables, so $\varphi(x_1, x_2) \equiv x_1 \in x_2$ and $\varphi(x_1, x_2, x_3) \equiv x_1 \in x_2$ are distinct. This proof will take place in four parts: first some logical points, then we consider propositional formulas, then atomic formulas, and finally bounded existentials.

Part (i): logical points. We make the following remarks.

- If $\text{ZF} \vdash \varphi(\mathbf{x}) \leftrightarrow \psi(\mathbf{x})$ and $\varphi(\mathbf{x})$ is a termed formula, then ψ is also a termed formula. This is immediate from the definition, since we can let $\mathcal{F}_\psi = \mathcal{F}_\varphi$.
- For all m, n , if $\varphi(x_1, \dots, x_n) \equiv \psi(x_1, \dots, x_m)$ and ψ is a termed formula, then so is φ . If $n \geq m$, we can show this by induction on n . The base case $n = m$ is trivial. For the inductive step, suppose

$$\varphi(x_1, \dots, x_{n+1}) \equiv \psi(x_1, \dots, x_m)$$

Then, we can write

$$\varphi(x_1, \dots, x_{n+1}) \equiv \theta(x_1, \dots, x_n)$$

where θ is a termed formula. Then

$$\mathcal{F}_\varphi(a_1, \dots, a_n, a_{n+1}) = a_{n+1} \times \mathcal{F}_\theta(a_1, \dots, a_n) = \mathcal{F}_4(a_{n+1}, \mathcal{F}_\theta(a_1, \dots, a_n))$$

giving the result by the inductive hypothesis. This is the reason for reversing the order: because the ordered triple $\langle x, y, z \rangle$ is $\langle x, \langle y, z \rangle \rangle$, the map

$$\{\langle x_1, x_2 \rangle \in a_1 \times a_2 \mid \theta(x_1, x_2)\} \mapsto \{\langle x_1, x_2, x_3 \rangle \in a_1 \times a_2 \times a_3 \mid \theta(x_1, x_2)\}$$

is much more complicated to implement in Gödel functions. We prove the case $n \leq m$ by induction; if

$$\varphi(x_1, \dots, x_{n-1}) \equiv \psi(x_1, \dots, x_m)$$

then

$$\varphi(x_1, \dots, x_{n-1}) \equiv \theta(x_1, \dots, x_n)$$

and

$$\{0\} = \{\mathcal{F}_3(a_1, a_1)\} = \mathcal{F}_1(\mathcal{F}_3(a_1, a_1), \mathcal{F}_3(a_1, a_1))$$

VII. Forcing and the Continuum Hypothesis

Then

$$\begin{aligned}\mathcal{F}_\varphi(a_1, \dots, a_{n-1}) &= \{\langle x_{n-1}, \dots, x_1 \rangle \in a_{n-1} \times \dots \times a_1 \mid \varphi(x_1, \dots, x_{n-1})\} \\ &= \text{ran}(\{\langle 0, x_{n-1}, \dots, x_1 \rangle \in \{0\} \times a_{n-1} \times \dots \times a_1 \mid \theta(x_1, \dots, x_{n-1}, 0)\}) \\ &= \mathcal{F}_6(\mathcal{F}_\theta(a_1, \dots, a_{n-1}, \mathcal{F}_1(\mathcal{F}_3(a_1, a_1), \mathcal{F}_3(a_1, a_1))), a_1)\end{aligned}$$

- If $\psi(x_1, \dots, x_n)$ is a termed formula and

$$\varphi(x_1, \dots, x_{n+1}) = \psi(x_1, \dots, x_{n-1}, x_{n+1}/x_n)$$

then φ is a termed formula. First, if $n = 1$, we have a termed formula $\psi(x_1)$ and consider $\psi(x_2/x_1)$. Then

$$\begin{aligned}\mathcal{F}_\varphi(a_1, a_2) &= \{\langle x_2, x_1 \rangle \in a_2 \times a_1 \mid \psi(x_2)\} \\ &= \{\langle x_2, x_1 \rangle \mid x_1 \in a_1 \wedge x_2 \in \mathcal{F}_\psi(a_2)\} \\ &= \mathcal{F}_\psi(a_2) \times a_1 \\ &= \mathcal{F}_4(\mathcal{F}_\psi(a_2), a_1)\end{aligned}$$

If $n > 1$, we have

$$\begin{aligned}\mathcal{F}_\varphi(a_1, \dots, a_{n+1}) &= \{\langle x_{n+1}, \dots, x_1 \rangle \mid x_n \in a_n \wedge \langle x_{n+1}, x_{n-1}, \dots, x_1 \rangle \in \mathcal{F}_\psi(a_1, \dots, a_{n-1}, a_{n+1})\} \\ &= \mathcal{F}_8(\mathcal{F}_\psi(a_1, \dots, a_{n-1}, a_{n+1}), a_n)\end{aligned}$$

- If $\psi(x_1, x_2)$ is a termed formula, and

$$\varphi(x_1, \dots, x_n) \equiv \psi(x_{n-1}/x_1, x_n/x_2)$$

then φ is a termed formula. This is trivial if $n = 2$, so we assume $n > 2$. Then

$$\begin{aligned}\mathcal{F}_\varphi(a_1, \dots, a_n) &= \{\langle x_n, \dots, x_1 \rangle \in a_n \times \dots \times a_1 \mid \langle x_n, x_{n-1} \rangle \in \mathcal{F}_\psi(a_{n-1}, a_n)\} \\ &= \mathcal{F}_7(\mathcal{F}_\psi(a_{n-1}, a_n), a_{n-2} \times \dots \times a_1)\end{aligned}$$

Part (ii): propositional connectives.

- If φ is a termed formula, then so is $\neg\varphi$.

$$\mathcal{F}_{\neg\varphi}(a_1, \dots, a_n) = (a_n \times \dots \times a_1) \setminus \mathcal{F}_\varphi(a_1, \dots, a_n)$$

- If φ, ψ are termed formulas, then so is $\varphi \vee \psi$.

$$\mathcal{F}_{\varphi \vee \psi}(a_1, \dots, a_n) = \mathcal{F}_\varphi(a_1, \dots, a_n) \cup \mathcal{F}_\psi(a_1, \dots, a_n)$$

It is easy to see that unions can be formed using Gödel functions.

- Conjunctions are similar to disjunctions.

$$\mathcal{F}_{\varphi \wedge \psi}(a_1, \dots, a_n) = \mathcal{F}_{\varphi}(a_1, \dots, a_n) \cap \mathcal{F}_{\psi}(a_1, \dots, a_n)$$

Part (iii): atomic formulas.

- Consider $\varphi(x_1, \dots, x_n) \equiv x_i = x_j$. We show that this is a termed formula for all $i, j \leq n$. Suppose $i = 1$ and $j = 2$. In this case,

$$\mathcal{F}_9(a_1, a_2) = \{\langle x_2, x_1 \rangle \in a_2 \times a_1 \mid x_1 = x_2\}$$

so \mathcal{F}_{φ} is formed using \mathcal{F}_9 and the discussion on dummy variables. Now suppose $j \geq i$. We prove this by induction. First, if $i = j$, then

$$\mathcal{F}_{\varphi} = \{\langle x_n, \dots, x_1 \rangle \in a_n \times \dots \times a_1 \mid x_i = x_i\} = a_n \times \dots \times a_1$$

Now, if $j = i + 1$, we let

$$\theta(x_1, \dots, x_{i+1}) = (x_1 = x_2)[x_i/x_1, x_{i+1}/x_2]$$

This is a termed formula by the result on substitutions. We thus obtain \mathcal{F}_{φ} by adding the required dummy variables. Now suppose we have $\varphi(x_1, \dots, x_n) \equiv x_i = x_{j+1}$. Then we can write

$$\varphi(x_1, \dots, x_{j+1}) = (x_i, x_j)[x_{j+1}, x_j]$$

which is a termed formula by substitution. This concludes the case $i \leq j$ by induction. Finally, suppose $i > j$. As $x_i = x_j$ is logically equivalent to $x_j = x_i$, which is a termed formula, φ is also a termed formula.

- Now consider $\varphi(x_1, \dots, x_n) \equiv x_i \in x_j$. As with equality, we first consider the case $i = 1, j = 2$. In this case, we can form \mathcal{F}_{10} with dummy variables. If $i = j$, the formula is always false, so we have

$$\mathcal{F}_{\varphi}(a_1, \dots, a_n) = \emptyset = a_1 \setminus a_1 = \mathcal{F}_3(a_1, a_1)$$

Now, let

$$\psi(x_1, \dots, x_{n+2}) \equiv (x_i = x_{n+1}) \wedge (x_j = x_{n+2}) \wedge (x_{n+1} \in x_{n+2})$$

We note that $x_{n+1} \in x_{n+2}$ is a termed formula as it is given by the substitution $(x_1 \in x_2)[x_{n+1}/x_1, x_{n+2}/x_2]$. The equalities are termed formulas as above, so ψ is a termed formula. Then

$$\begin{aligned} \mathcal{F}_{\varphi}(a_1, \dots, a_n) &= \text{ran ran}\{\langle x_{n+2}, \dots, x_1 \rangle \times a_j \times a_i \times a_n \times \dots \times a_1 \mid \\ &\quad x_i = x_{n+1} \wedge x_j = x_{n+2} \wedge x_{n+1} \in x_{n+2}\} \\ &= \mathcal{F}_6(\mathcal{F}_6(\mathcal{F}_{\psi}(a_1, \dots, a_n), a_1), a_1) \end{aligned}$$

VII. Forcing and the Continuum Hypothesis

Part (iv): *bounded quantifiers*. We required that the only occurrence of \exists was in the form

$$\varphi(x_1, \dots, x_n) \equiv \exists x_{m+1} \in x_j. \psi(x_1, \dots, x_{m+1})$$

where $j \leq m \leq n$. Due to this restriction, it suffices to show that if $\psi(x_1, \dots, x_{n+1})$ is a termed formula, then so is the formula

$$\varphi(x_1, \dots, x_n) \equiv \exists x_{n+1} \in x_j. \psi(x_1, \dots, x_{n+1})$$

Let $\theta(x_1, \dots, x_{n+1}) \equiv x_{n+1} \in x_j$. Then $\theta \wedge \psi$ is a termed formula. Now

$$\begin{aligned} \mathcal{F}_{\theta \wedge \psi}(a_1, \dots, a_n, \mathcal{F}_2(a_j, a_j)) &= \mathcal{F}_{\theta \wedge \psi}(a_1, \dots, a_n, \bigcup a_j) \\ &= \left\{ \langle x_{n+1}, \dots, x_1 \rangle \in \left(\bigcup a_j \right) \times a_n \times \dots \times a_1 \mid \right. \\ &\quad \left. x_{n+1} \in x_j \wedge \forall k \leq n. x_k \in a_k \wedge \psi(x_1, \dots, x_{n+1}) \right\} \end{aligned}$$

So

$$\begin{aligned} \text{ran}(\mathcal{F}_{\theta \wedge \psi}(a_1, \dots, a_n, \bigcup a_j)) &= \left\{ \langle x_n, \dots, x_1 \rangle \in a_n \times \dots \times a_1 \mid \right. \\ &\quad \left. \exists u. \langle u, x_n, \dots, x_1 \rangle \in \mathcal{F}_{\theta \wedge \psi}(a_1, \dots, a_n, \bigcup a_j) \right\} \\ &= \left\{ \langle x_n, \dots, x_1 \rangle \in a_n \times \dots \times a_1 \mid \right. \\ &\quad \left. \exists x_{n+1} \in x_j. \psi(x_1, \dots, x_{n+1}) \right\} \end{aligned}$$

□

Definition. A class C is *closed under Gödel functions* if whenever $x, y \in C$, we have $\mathcal{F}_i(x, y) \in C$ for $i \in \{1, \dots, 10\}$. Given a set b , we let $\text{cl}(b)$ be the smallest set C containing b as a subset that is closed under Gödel functions.

For example, $\text{cl}(\emptyset) = \emptyset$, $a, b \in \text{cl}(\{a, b\})$, and $\text{cl}(b) = \text{cl}(\text{cl}(b))$.

Definition. Let b be a set. Define $\mathcal{D}^n(b)$ inductively by

$$\mathcal{D}^0(b) = b; \quad \mathcal{D}^{n+1}(b) = \{\mathcal{F}_i(x, y) \mid x, y \in \mathcal{D}^n(b), i \in \{1, \dots, 10\}\}$$

One can easily check that $\text{cl}(b) = \bigcup_{n \in \omega} \mathcal{D}^n(b)$.

Lemma. If M is a transitive class that is closed under Gödel functions, then M satisfies Δ_0 -separation.

Proof. Let $\varphi(x_1, \dots, x_n)$ be a Δ_0 -formula, and let $a, b_1, \dots, b_{i-1}, b_{i+1}, \dots, b_n \in M$. Let

$$Y = \{x_i \in a \mid \varphi(b_1, \dots, b_{i-1}, x_i, b_{i+1}, \dots, b_n)\}$$

We must show $Y \in M$. Let \mathcal{F}_φ be the formula built from Gödel's normal form theorem. Then for any $c_1, \dots, c_n \in M$, we have

$$\mathcal{F}_\varphi(c_1, \dots, c_n) = \{\langle x_n, \dots, x_1 \rangle \in c_n \times \dots \times c_1 \mid \varphi(x_1, \dots, x_n)\} \in M$$

Hence, as $\{b_j\} = \mathcal{F}_1(b_j, b_j) \in M$, we obtain

$$\mathcal{F}_\varphi(\{b_1\}, \dots, \{b_{i-1}\}, a, \{b_{i+1}\}, \dots, \{b_n\}) \in M$$

Then, we can show that $Y \in M$ by taking the range \mathcal{F}_6 a total of $n - i$ times and then taking the domain \mathcal{F}_5 . \square

Theorem. For every transitive set M , the collection of definable subsets is

$$\text{Def}(M) = \text{cl}(M \cup \{M\}) \cap \mathcal{P}(M)$$

Proof. We first prove the forward direction. Let φ be a formula. Then φ^M is Δ_0 , so there is a term \mathcal{G} built from the Gödel functions $\mathcal{F}_1, \dots, \mathcal{F}_{10}$ such that for $a_1, \dots, a_n \in M$, we have

$$\{x \in M \mid (M, \in) \models \varphi(x, a_1, \dots, a_n)\} = \{x \in M \mid \varphi^M(x, a_1, \dots, a_n)\} = \mathcal{G}(M, a_1, \dots, a_n) \in \text{cl}(M \cup \{M\})$$

We now show the converse. We first claim that if \mathcal{G} is built from the Gödel functions, then for any x, a_1, \dots, a_n , the formulas

$$x = \mathcal{G}(a_1, \dots, a_n); \quad x \in \mathcal{G}(a_1, \dots, a_n)$$

are Δ_0 . This can be proven inductively using the iterative construction of $\text{cl}(M \cup \{M\})$. For example, if $X, Y \in \mathcal{D}^k(a_1, \dots, a_n)$, then $x = \mathcal{F}_1(X, Y)$ is equivalent to the statement

$$(\forall z \in x. z = X \vee z = Y) \wedge (\exists w \in x. w = X) \wedge (\exists w \in x. w = Y)$$

so the result holds for \mathcal{F}_1 ; very similar proofs show the result for both equality and membership for all other Gödel functions.

Let $Z \in \text{cl}(M \cup \{M\}) \cap \mathcal{P}(M)$. Since $Z \in \text{cl}(M \cup \{M\})$, we can fix a term \mathcal{G} built from the $\mathcal{F}_1, \dots, \mathcal{F}_{10}$ such that $Z = \mathcal{G}(M, a_1, \dots, a_n)$. Let φ be a Δ_0 formula such that $x \in \mathcal{G}(M, a_1, \dots, a_n)$ if and only if $\varphi(x, M, a_1, \dots, a_n)$. Then $\mathcal{G}(M, a_1, \dots, a_n) = \{x \in M \mid \varphi(x, M, a_1, \dots, a_n)\}$ as $Z \subseteq M$. It therefore remains to prove that there is a formula ψ such that

$$\psi^M(x, a_1, \dots, a_n) \leftrightarrow \varphi(x, M, a_1, \dots, a_n)$$

For example, we can define ψ from φ by the following replacements.

- (i) $\exists v_i \in M \mapsto \exists v_i$;
- (ii) $v_i \in M \mapsto v_i = v_i$;
- (iii) $M = M \mapsto v_0 = v_0$;
- (iv) $M \in M, M \in v_i, M = v_i \mapsto v_0 \neq v_0$.

Finally, we obtain

$$Z = \mathcal{G}(M, a_1, \dots, a_n) = \{x \in M \mid \psi^M(x, a_1, \dots, a_n)\} \in \text{Def}(M)$$

\square

VII. Forcing and the Continuum Hypothesis

2.4. The axiom of constructibility

Definition. The *axiom of constructibility* is the statement $V = L$. Equivalently, $\forall x. \exists \alpha \in \text{Ord}. (x \in L_\alpha)$.

We will show that if ZF is consistent, then so is $ZF + (V = L)$, by demonstrating that L is a model of $ZF + (V = L)$. To do this, we will show that being constructible is absolute.

Lemma. $Z = \text{cl}(M)$ is Δ_1^{ZF} .

Proof. The Π_1 definition is simply being the smallest set closed under Gödel functions. More explicitly,

$$\forall W. \left(M \cup \{M\} \subseteq W \wedge \forall x, y \in W. \bigwedge_{i \leq 10} \mathcal{F}_i(x, y) \in W \right) \rightarrow Z \subseteq W$$

The Σ_1 definition will use the inductive definition of the closure.

$$\begin{aligned} \exists W. W \text{ is a function} \wedge \text{dom } W = \omega \wedge Z = \bigcup \text{ran } W \\ \wedge W(0) = M \wedge W(n) \subseteq W(n+1) \\ \wedge \left(\forall x, y \in W(n). \bigwedge_{i \leq 10} \mathcal{F}_i(x, y) \in W(n+1) \right) \\ \wedge \left(\forall z \in W(n+1). \exists x, y \in W(n). \bigvee_{i \leq 10} z = \mathcal{F}_i(x, y) \right) \end{aligned}$$

□

Lemma. The function mapping $\alpha \mapsto L_\alpha$ is absolute between transitive models of ZF.

Proof. Define $G : \text{Ord} \times V \rightarrow V$ by

$$G(\alpha, x) = \begin{cases} \text{cl}(x(\beta) \cup \{x(\beta)\}) & \text{if } \alpha = \beta + 1 \text{ and } x \text{ is a function with domain } \beta \\ \bigcup_{\beta < \alpha} x(\beta) & \text{if } \alpha \text{ is a limit} \\ \emptyset & \text{otherwise} \end{cases}$$

All of these conditions and constructions are absolute, so G is an absolute function. Therefore, by transfinite recursion, there exists $F : \text{Ord} \rightarrow V$ where $F : \alpha \mapsto G(x, F|_\alpha)$. By absoluteness of transfinite recursion, F is absolute. Finally, $F(\alpha) = L_\alpha$ for all ordinal α . □

Theorem. (i) L satisfies the axiom of constructibility.

(ii) L is the smallest inner model of ZF. That is, if M is an inner model of ZF, then $L \subseteq M$.

Proof. Part (i). We must show

$$(\forall x. \exists \alpha \in \text{Ord}. x \in L_\alpha)^L$$

which is

$$\forall x \in L. \exists \alpha \in \text{Ord}. x \in (L_\alpha)^L$$

Since the L_α hierarchy is absolute, $x \in (L_\alpha)^L$ if and only if $x \in L_\alpha$. As L contains every ordinal, if $x \in L$ then $x \in L_\alpha$ for some α , and thus $x \in (L_\alpha)^L$. Hence $L \models \alpha \in L \wedge x \in L_\alpha$.

Part (ii). Let M be an arbitrary inner model of ZF. We construct L inside M to give L^M . By absoluteness, for every $\alpha \in M \cap \text{Ord}$, we have $L_\alpha = (L_\alpha)^M$. Thus $L_\alpha \subseteq M$ for every $\alpha \in M \cap \text{Ord} = \text{Ord}$. Hence $L \subseteq M$ as required. \square

2.5. Well-ordering the universe

We will show that L satisfies a strong version of the axiom of choice, namely that there is a definable global well-order. We will define well-orderings $<_\alpha$ on L_α such that $<_{\alpha+1}$ *end-extends* $<_\alpha$: if $y \in L_\alpha$ and $x \in L_{\alpha+1} \setminus L_\alpha$, then $y <_{\alpha+1} x$. Then we set $<_L = \bigcup_\alpha <_\alpha$.

Theorem. There is a well-ordering of L .

Proof. For each ordinal α , we will construct a well-order $<_\alpha$ on L_α such that if $\alpha < \beta$, the following hold:

- (i) if $x <_\alpha y$ then $x <_\beta y$; and
- (ii) if $x \in L_\alpha$ and $y \in L_\beta \setminus L_\alpha$, then $x <_\beta y$.

For limit cases, we take unions:

$$<_\gamma = \bigcup_{\alpha < \gamma} <_\alpha$$

We now describe the construction of $<_{\alpha+1}$. To do this, we consider the ordering on L_α , and append the singleton $\{L_\alpha\}$. We then follow that by the elements of $\mathcal{D}(L_\alpha \cup \{L_\alpha\}) \setminus (L_\alpha \cup \{L_\alpha\})$. We then add $\mathcal{D}^2(L_\alpha \cup \{L_\alpha\}) \setminus \mathcal{D}(L_\alpha \cup \{L_\alpha\})$, and so forth. In order to do this, we define $<_{\alpha+1}^n$ for $n \in \omega$ as follows.

- (i) $<_{\alpha+1}^0$ is the well-ordering of $L_\alpha \cup \{L_\alpha\}$ given by making $\{L_\alpha\}$ the maximal element.
- (ii) Suppose that $<_{\alpha+1}^n$ is defined. We end-extend $<_{\alpha+1}^n$ to form $<_{\alpha+1}^{n+1}$ as follows. Suppose $x, y \notin \mathcal{D}^n(L_\alpha \cup \{L_\alpha\})$. We say $x <_{\alpha+1}^{n+1} y$ if either
 - (a) the least $i \leq 10$ such that $\exists u, v \in \mathcal{D}^n(L_\alpha \cup \{L_\alpha\})$ with $x = \mathcal{F}_i(u, v)$ is less than the least $i \leq 10$ such that $\exists u, v \in \mathcal{D}^n(L_\alpha \cup \{L_\alpha\})$ with $y = \mathcal{F}_i(u, v)$; or
 - (b) these indices i are equal, and the $<_{\alpha+1}^n$ -least $u \in \mathcal{D}^n(L_\alpha \cup \{L_\alpha\})$ such that there exists $v \in \mathcal{D}^n(L_\alpha \cup \{L_\alpha\})$ with $x = \mathcal{F}_i(u, v)$ is less than the $<_{\alpha+1}^n$ -least $u \in \mathcal{D}^n(L_\alpha \cup \{L_\alpha\})$ such that there exists $v \in \mathcal{D}^n(L_\alpha \cup \{L_\alpha\})$ with $y = \mathcal{F}_i(u, v)$; or

VII. Forcing and the Continuum Hypothesis

- (c) both of these coincide, and $<_{\alpha+1}^n$ -least $v \in \mathcal{D}^n(L_\alpha \cup \{L_\alpha\})$ with $x = \mathcal{F}_i(u, v)$ is less than the least $v \in \mathcal{D}^n(L_\alpha \cup \{L_\alpha\})$ with $y = \mathcal{F}_i(u, v)$.

□

The restriction of $<_L$ to any set $x \in L$ is a well-ordering of x . Since every set can be well-ordered, the axiom of choice holds.

Lemma. The relation $<_L$ is Σ_1 -definable. Moreover, for every limit ordinal δ and $y \in L_\delta$, we have $x <_L y$ if and only if $x \in L_\delta$ and $(L_\delta, \in) \models x <_L y$.

2.6. The generalised continuum hypothesis in L

Lemma. (ZFC)

- (i) For all $n \in \omega$, we have $L_n = V_n$.
- (ii) If M is infinite, then $|M| = |\text{Def}(M)|$.
- (iii) If α is an infinite ordinal, then $|L_\alpha| = |\alpha|$.

Lemma (Gödel's condensation lemma). For every limit ordinal δ , if $(M, \in) < (L_\delta, \in)$, then there exists some $\beta \leq \delta$ such that $(M, \in) \cong (L_\beta, \in)$.

Proof. Let $\pi : (M, \in) \rightarrow (N, \in)$ be the Mostowski collapse, and set $\beta = N \cap \text{Ord}$. Since N is transitive, $\beta \in \text{Ord}$. We will prove that $\beta \leq \delta$ and $N = L_\beta$.

First, suppose $\delta < \beta$. Then $\delta \in N$, so $\pi^{-1}(\delta) \in M$. Since being an ordinal is absolute between transitive models, $N \models \delta \in \text{Ord}$, so $M \models \pi^{-1}(\delta) \in \text{Ord}$. Note that this does not immediately imply that $\pi^{-1}(\delta)$ is an ordinal in V since M is not necessarily transitive. But as $M < L_\delta$, we obtain $L_\delta \models \pi^{-1}(\delta) \in \text{Ord}$, and since L_δ is transitive, $\pi^{-1}(\delta)$ is an ordinal in V .

Also, $M \models x \in \pi^{-1}(\delta)$ if and only if $N \models \pi(x) \in \delta$. Hence,

$$\pi : (\pi^{-1}(\delta) \cap M) \rightarrow \delta$$

is an isomorphism. Therefore, the order type of $\pi^{-1}(\delta) \cap M$ is δ . Let $f : \delta \rightarrow \pi^{-1}(\delta) \cap M$ be a strictly increasing enumeration. Then, for any $\alpha \in \delta$, we must have $\alpha \leq f(\alpha) < \pi^{-1}(\delta)$. Hence $\delta \leq \pi^{-1}(\delta)$. On the other hand, $\pi^{-1}(\delta) \in M < L_\delta$, so $\pi^{-1}(\delta) < \delta$. This gives a contradiction.

We now show $\beta > 0$. Since

$$L_\delta \models \exists x. \forall y \in x. (y \neq y)$$

the elementary substructure M must also believe this statement, and so N does. In particular, since N believes in the existence of an empty set, we must have $\emptyset \in N \cap \text{Ord} = \beta$ as required.

We show β is a limit. We know that

$$L_\delta \models \forall \alpha \in \text{Ord}. \exists x. x = \alpha + 1$$

So M and hence N believe this statement. Let $\alpha \in \beta = N \cap \text{Ord}$, then by absoluteness, $\alpha + 1 \in N$.

Now we show $L_\beta \subseteq N$.

$$L_\delta \models \forall \alpha \in \text{Ord}. \exists y. y = L_\alpha$$

So N satisfies this sentence. Since the L_α hierarchy is absolute, for all $\alpha \in N \cap \text{Ord} = \beta$, we have $L_\alpha \in N$.

Finally, we show $N \subseteq L_\beta$.

$$L_\delta \models \forall x. \exists y. \exists z. y \in \text{Ord} \wedge z = L_y \wedge x \in z$$

As N satisfies this sentence, for a fixed $a \in N$ there are $\gamma \in N$ and $z \in N$ such that

$$N \models \gamma \in \text{Ord} \wedge z = L_\gamma \wedge a \in z$$

By absoluteness, $a \in L_\gamma \subseteq L_\beta$ as required. \square

Theorem. If $V = L$, then $2^{\aleph_\alpha} = \aleph_{\alpha+1}$ for every ordinal α . In particular, GCH holds.

Proof. We will show that $\mathcal{P}(\omega_\alpha) \subseteq L_{\omega_{\alpha+1}}$. Then, as $|L_{\omega_{\alpha+1}}| = \aleph_{\alpha+1}$, the proof follows. To do this, it suffices to show that if $X \subseteq \omega_\alpha$, then there exists some $\gamma < \omega_{\alpha+1}$ such that $X \in L_\gamma$.

Let $X \subseteq \omega_\alpha$ and let $\delta > \omega_\alpha$ be a limit ordinal such that $X \in L_\delta$. Let M be an elementary submodel of L_δ such that $\omega_\alpha \subseteq M$, $X \in M$, and $|M| = \aleph_\alpha$. This exists by the downward Löwenheim–Skolem theorem. By Gödel’s condensation lemma, if N is the Mostowski collapse of M , then there is a limit ordinal $\gamma \leq \delta$ such that $N = L_\gamma$. As $|N| = |M| = \aleph_\alpha$, we have $|L_\gamma| = \aleph_\alpha$, so $\gamma < \omega_{\alpha+1}$. Finally, as $\omega_\alpha \subseteq M$, the collapsing map is the identity on ω_α . Thus, the map fixes X , and so $X \in L_\gamma$. \square

This gives the following theorem.

Theorem. $\text{Con}(\text{ZF})$ implies $\text{Con}(\text{ZFC} + V = L + \text{GCH})$.

Proof. We have shown that there is a definable class L such that ZF proves

$$(\text{ZFC} + V = L + \text{GCH})^L$$

Suppose that $\text{ZFC} + V = L + \text{GCH}$ were inconsistent. Then fix φ such that

$$\text{ZFC} + V = L + \text{GCH} \vdash \varphi \wedge \neg \varphi$$

Then

$$\text{ZF} \vdash (\varphi \wedge \neg \varphi)^L$$

By relativisation, $\varphi^L \wedge \neg(\varphi^L)$. Hence ZF is inconsistent. \square

VII. Forcing and the Continuum Hypothesis

Lemma (Shepherdson). There is no class W such that

$$\text{ZFC} \vdash W \text{ is an inner model} \wedge (\neg\text{CH})^W$$

Therefore, the technique of inner models does not let us prove the independence of CH from ZFC. In order to do this, we will introduce the notion of *forcing*.

2.7. Combinatorial properties

Definition. Let Ω be either a regular cardinal or the class of all ordinals. A subclass $C \subseteq \Omega$ is said to be a *club*, or *closed and unbounded*, if it is

- (i) *closed*: for all $\gamma \in \Omega$, we have $\sup(C \cap \gamma) \in C$;
- (ii) *unbounded*: for all $\alpha \in \Omega$ there exists $\beta \in C$ with $\beta > \alpha$.

A class $S \subseteq \Omega$ is *stationary* if it intersects every club.

Note that being a stationary class for Ord is not first-order definable.

The property \diamond states that there is a single sequence of length ω_1 which can approximate any subset of ω_1 in a suitable sense.

Definition. We say that the *diamond principle* \diamond holds if there is a sequence $(A_\alpha)_{\alpha < \omega_1}$ such that

- (i) for each $\alpha < \omega_1$, we have $A_\alpha \subseteq \alpha$; and
- (ii) for all $X \subseteq \omega_1$, the set $\{\alpha \mid X \cap \alpha = A_\alpha\}$ is stationary.

Lemma. $\text{ZF} \vdash \diamond \rightarrow \text{CH}$.

Proof. If $(A_\alpha)_{\alpha < \omega_1}$ is a \diamond -sequence, then for all $X \subseteq \omega$, there is $\alpha > \omega$ such that $X = A_\alpha$. Thus $\{A_\alpha \mid \alpha \in \omega_1 \wedge A_\alpha \subseteq \omega\} = \mathcal{P}(\omega)$. \square

Theorem. If $V = L$, then \diamond holds.

Remark. \diamond is used in many inductive constructions in L to build combinatorial objects such as Suslin trees.

Definition. Let κ be an uncountable cardinal. Then the *square principle* \square_κ is the assertion that there exists a sequence (C_α) indexed by the limit ordinals α in κ^+ , such that

- (i) C_α is a club subset of α ;
- (ii) if β is a limit ordinal of C_α then $C_\beta = C_\alpha \cap \beta$; and
- (iii) if $\text{cf}(\alpha) < \kappa$ then $|C_\alpha| < \kappa$.

Theorem (Jensen). If $V = L$, then \square_κ holds for every uncountable cardinal κ .

2. Constructibility

Lemma. If \square_{ω_1} , then there exists a stationary set $S \subseteq \{\beta \in \omega_2 \mid \text{cf}(\beta) = \omega\}$ such that for all $\alpha \in \omega_2$ with $\text{cf}(\alpha) = \omega_1$, $S \cap \alpha$ is not stationary in α .

Remark. If κ is a weakly compact cardinal, then every stationary subset of κ *reflects*: there is $\alpha \in \kappa$ such that $S \cap \alpha$ is stationary in α . In fact, the claim that every stationary subset of $\{\beta \in \omega_2 \mid \text{cf}(\beta) = \omega\}$ reflects at a point of cofinality ω_1 is equiconsistent with ZFC together with the assertion that there is a Mahlo cardinal.

3. Forcing

3.1. Introduction

The idea behind forcing is to widen a given model of ZFC to ‘add lots of reals’. But if we work over V , we already have added all of the sets, so there is nothing left to add. Instead, we will work over countable transitive set models of ZFC. However, this means that we will not immediately get $\text{Con}(\text{ZF}) \rightarrow \text{Con}(\text{ZFC} + \neg\text{CH})$. We will then use the reflection theorem to obtain this result.

If M is such a countable transitive model, we want to add ω_2^M -many reals to M . We will try to do this in a ‘minimal way’; for example, we do not want to add any ordinals. This gives us much more control over the model that we build.

Recall the argument that the sentence $\varphi(x) \equiv \exists x. x^2 = 2$ is independent of the axioms of fields: we began with a field in which the sentence failed, namely \mathbb{Q} , and then extended it in a minimal way to $\mathbb{Q}[\sqrt{2}]$. The model $\mathbb{Q}[\sqrt{2}]$ does not just contain $\mathbb{Q} \cup \{\sqrt{2}\}$, it also contains everything that can be built from \mathbb{Q} and $\sqrt{2}$ using the axioms of fields. The field $\mathbb{Q}[\sqrt{2}]$ is the minimal field extension of \mathbb{Q} satisfying φ .

We may encounter some difficulties when adding arbitrary reals to our model. Suppose that M is of the form L_γ , where γ is a countable ordinal. Then γ can be coded as a subset c of ω , which can be viewed as a real. If we added c to M , we could decode it to form $\gamma = \text{Ord} \cap M$. This would violate the principle of not adding any new ordinals.

Suppose we enumerate all formulas as $\{\varphi_n \mid n \in \omega\}$. Let $r = \{n \mid M \models \varphi_n\}$. If we added r to M , we could then build a truth predicate for M . This would cause self-referential problems discussed by Tarski.

The main issues we must overcome are the following.

- (i) We need a method to choose the ω_2^M -many subsets of M to be added.
- (ii) Given these, we need to ensure that the extension satisfies ZFC.
- (iii) We must ensure that ω_1^M and ω_2^M are still cardinals in the extension.

We will build these reals from within M itself. Note that if r is a real, then each of its finite decimal approximations is already in M . The issue is that from within M , we do not know what the real we want to add is. So we may not know from within M which reals we will add. Instead, we will add a *generic* real. To be generic, we will not specify any particular digits, but its decimal expansion will contain every finite sequence. We will call a specification *dense* if any finite approximation can be extended to one satisfying the specification. For example, ‘beginning with a 7’ is not dense, but ‘containing the subsequence 746’ is dense. We will define that a real is generic precisely when it meets every dense specification.

Note that there are explicit, absolute bijections $f : \mathcal{P}(\omega) \rightarrow \omega^\omega$, $g : \omega^\omega \rightarrow 2^\omega$, $h : 2^\omega \rightarrow \mathbb{R}$ and so on. So if $M \models \text{ZFC}$, knowledge of $\mathcal{P}^M(\omega)$ gives us $(\omega^\omega)^M, (2^\omega)^M, \mathbb{R}^M$. Because of

this, by a ‘real’ we mean either an element of \mathbb{R} , a function $\omega \rightarrow \omega$, a function $\omega \rightarrow 2$, or a subset of ω . In formal arguments, reals will normally be either subsets of ω or functions $\omega \rightarrow 2$.

The axiom of choice is not needed in the basic machinery of forcing, so we will work primarily over ZF and state explicitly where choice is used.

3.2. Forcing posets

Definition. A *preorder* is a pair (\mathbb{P}, \leq) such that

- \mathbb{P} is nonempty;
- \leq is a binary relation on \mathbb{P} ;
- \leq is transitive, so $p \leq q$ and $q \leq r$ implies $p \leq r$;
- \leq is reflexive, so $p \leq p$.

A preorder is called a *partial order* if \leq is antisymmetric, so $p \leq q$ and $q \leq p$ implies $p = q$.

Definition. A *forcing poset* is a triple $(\mathbb{P}, \leq_{\mathbb{P}}, 1_{\mathbb{P}})$, where $(\mathbb{P}, \leq_{\mathbb{P}})$ is a preorder and $1_{\mathbb{P}}$ is a maximal element. Elements of \mathbb{P} are called *conditions*, and we say q is *stronger* than p or an *extension* of p if $q \leq p$. We say that p, q are *compatible*, written $p \parallel_{\mathbb{P}} q$, if there exists r such that $r \leq_{\mathbb{P}} p, q$. Otherwise, we say they are *incompatible*, written $p \perp q$.

Remark. In some texts, the partial order is reversed. This is called *Jerusalem notation*.

The notation $\mathbb{P} \in M$ abbreviates $(\mathbb{P}, \leq_{\mathbb{P}}, 1_{\mathbb{P}}) \in M$. Note that by transitivity if \mathbb{P} is an element of M , then $1_{\mathbb{P}} \in M$, but we do not necessarily have $\leq_{\mathbb{P}} \in M$.

Definition. A preorder is *separative* if whenever $p \neq q$, exactly one of the following two cases holds:

- (i) $q \leq p$ and $p \not\leq q$; or
- (ii) there exists $r \leq q$ such that $r \perp p$.

Proposition. (i) If (\mathbb{P}, \leq) is a separative preorder, it is a partial order.

- (ii) If (\mathbb{P}, \leq) is a poset, then it is separative if and only if whenever $q \not\leq p$, there is $r \leq q$ such that $r \perp p$.

Proposition. Suppose that (\mathbb{P}, \leq) is a preorder. Define $p \sim q$ by

$$p \sim q \leftrightarrow \forall r \in \mathbb{P}. (r \parallel p \leftrightarrow r \parallel q)$$

Then there is a separative preorder on \mathbb{P}/\sim such that

$$[p] \perp [q] \leftrightarrow p \perp q$$

and if \mathbb{P} has a maximal element, so does \mathbb{P}/\sim .

VII. Forcing and the Continuum Hypothesis

Example. For sets I, J , we let $\text{Fn}(I, J)$ denote the set of all finite partial functions from I to J .

$$\text{Fn}(I, J) = \{p \mid |p| < \omega \wedge p \text{ is a function} \wedge \text{dom } p \subseteq I \wedge \text{ran } p \subseteq J\}$$

We let \leq be the reverse inclusion on $\text{Fn}(I, J)$, so $q \leq p$ if and only if $q \supseteq p$. The maximal element $\mathbb{1}$ is the empty set. Then $(\text{Fn}(I, J), \supseteq, \emptyset)$ is a forcing poset, and moreover, the preorder is separative.

Remark. When α is an ordinal, the forcing poset $\text{Fn}(\alpha \times \omega, 2)$ is often written $\text{Add}(\omega, \alpha)$, denoting the idea that we are adding α -many subsets of ω .

3.3. Chains and Δ -systems

Definition. Let \mathbb{P} be a forcing poset.

- (i) A *chain* is a subset $C \subseteq \mathbb{P}$ such that for every $p, q \in C$, either $p \leq q$ or $q \leq p$.
- (ii) An *antichain* is a subset $A \subseteq \mathbb{P}$ such that for every $p, q \in A$, either $p = q$ or $p \perp q$. An antichain is *maximal* if it is not strictly contained in any other antichain.
- (iii) We say that \mathbb{P} has the *countable chain condition* if every antichain is countable.

Example. (i) Consider the tree $\text{Fn}(\omega, 2)$. A chain is a branch through the tree, and an antichain is a collection of points on different branches.

- (ii) The set of functions $\{\langle 0, 0 \rangle, \langle 1, n \rangle \mid n \in \omega\}$ forms an antichain of length ω in $\text{Fn}(I, \omega)$ if $\{0, 1\} \subseteq I$.

Definition. A family of sets \mathcal{A} forms a Δ -system with root R when $X \cap Y = R$ for all $X \neq Y$ in \mathcal{A} .

Example. If $R = \emptyset$, then \mathcal{A} is a family of pairwise disjoint sets.

Definition. Let A be a set and θ a cardinal. Then we write $[A]^\theta$ for the set of subsets of A of size θ .

$$[A]^\theta = \{x \subseteq A \mid |x| = \theta\}$$

We write $[A]^{<\theta}$ for the set of subsets of A of size strictly less than θ .

$$[A]^{<\theta} = \{x \subseteq A \mid |x| < \theta\}$$

Similarly, $[A]^{\leq \theta} = [A]^\theta \cup [A]^{<\theta}$.

Recall that for regular cardinals κ , if \mathcal{F} is a family of sets of size less than κ and each element of \mathcal{F} has size less than κ , then $\bigcup \mathcal{F}$ has size less than κ .

Lemma (Δ -system lemma). (ZFC) Let κ be an uncountable regular cardinal, and let \mathcal{A} be a family of finite sets with $|\mathcal{A}| = \kappa$. Then there exists $\mathcal{B} \in [\mathcal{A}]^\kappa$ that forms a Δ -system.

Proof. To begin, we construct $\mathcal{C} \in [\mathcal{A}]^\kappa$ such that all elements of \mathcal{C} have the same cardinality. By assumption, each element of \mathcal{A} is finite, and so we can define $Y_n = \{X \in \mathcal{A} \mid |X| = n\}$, and suppose each of the Y_n had size less than κ . Then $|\mathcal{A}| = |\bigcup Y_n| < \kappa$, giving a contradiction.

Fix $n \in \omega$ such that $\mathcal{C} = Y_n$ has size κ . We show by induction on n that if $\mathcal{C} = \{X \in \mathcal{A} \mid |X| = n\}$, then there is $\mathcal{B} \subseteq \mathcal{C}$ of size κ that forms a Δ -system. If $n = 1$, we have a collection of pairwise disjoint singletons, so \mathcal{C} is already a Δ -system with root \emptyset as required. Now suppose $n > 1$ and the claim holds for $n - 1$. For each $p \in \bigcup \mathcal{C}$, let $C_p = \{X \in \mathcal{C} \mid p \in X\}$. There are two cases to consider.

Suppose $|C_p| = \kappa$ for some $p \in \bigcup \mathcal{C}$. Then for such a p , we set $\mathcal{D} = \{X \setminus \{p\} \mid X \in C_p\}$. This set has size κ , and each element of \mathcal{D} has size $n - 1$. By the inductive hypothesis, we can find some $\mathcal{E} \in [\mathcal{D}]^\kappa$ such that \mathcal{E} forms a Δ -system with root R . Then $\{Y \cup \{p\} \mid Y \in \mathcal{E}\}$ is a Δ -system with root $R \cup \{p\}$.

Now suppose all of the C_p have size less than κ . Then as κ is regular, for any set S of size less than κ ,

$$\{X \in \mathcal{C} \mid X \cap S \neq \emptyset\} = \bigcup_{p \in S} C_p$$

has size less than κ . Therefore, there exists some $X \in \mathcal{C}$ such that $X \cap S = \emptyset$. We recursively choose $X_\alpha \in \mathcal{C}$ for each $\alpha < \kappa$ such that $X_\alpha \cap \bigcup_{\beta < \alpha} X_\beta = \emptyset$. Then $\{X_\alpha \mid \alpha < \kappa\} \in [\mathcal{C}]^\kappa$ is a Δ -system with empty root. \square

We can show that assumptions in the above lemma were required.

Proposition. Suppose κ is ω or singular. Then there exists a family \mathcal{A} of finite sets with $|\mathcal{A}| = \kappa$ but no $\mathcal{B} \in [\mathcal{A}]^\kappa$ forms a Δ -system.

Lemma. (ZFC) $\text{Fn}(I, J)$ has the countable chain condition if and only if I is empty or J is countable.

Proof. First, we observe that if I or J are empty, then $\text{Fn}(I, J)$ is empty and so trivially has the countable chain condition. Now let us assume that both I and J are nonempty.

Suppose that J is uncountable. Then for any $i \in I$, the set

$$\{\langle i, j \rangle \mid j \in J\}$$

is an uncountable antichain.

Now suppose J is countable, and let $\{p_\alpha \mid \alpha \in \omega_1\}$ be a collection of distinct elements of $\text{Fn}(I, J)$. Let $\mathcal{A} = \{\text{dom } p_\alpha \mid \alpha \in \omega_1\}$, which is a collection of ω_1 -many finite sets. By the Δ -system lemma, we can find an uncountable subset $\mathcal{B} \subseteq \mathcal{A}$ with a root $R \subseteq I$. By definition, $R \subseteq \text{dom}(p_\alpha)$ for all $\text{dom } p_\alpha \in \mathcal{B}$, the root R must be finite. Since J is countable, there are only countably many functions $R \rightarrow J$. Therefore, as \mathcal{B} is uncountable, there are $\alpha \neq \beta$ such that $\text{dom } p_\alpha$ and $\text{dom } p_\beta$ are both in \mathcal{B} and $p_\alpha|_R = p_\beta|_R$. But then since R is a root, $\text{dom } p_\alpha \cap \text{dom } p_\beta = R$, so $p_\alpha \parallel p_\beta$, witnessed by their union $p_\alpha \cup p_\beta$. So the $\{p_\alpha \mid \alpha \in \omega_1\}$ cannot form an antichain. \square

3.4. Dense sets and genericity

Definition. Let \mathbb{P} be a forcing poset.

- (i) $D \subseteq \mathbb{P}$ is *dense* if for all $p \in \mathbb{P}$ there exists $q \in D$ such that $q \leq p$.
- (ii) $D \subseteq \mathbb{P}$ is *open* if for all $p \in D$ and $q \in \mathbb{P}$, if $q \leq p$ then $q \in D$.

A set of conditions is dense if every condition can be extended to one in that set, and a set is open if it is closed under strengthening conditions.

Example. Let I be infinite and J nonempty. Then for all $i \in I$ and $j \in J$, the following are dense.

- (i) $D_i = \{q \in \text{Fn}(I, J) \mid i \in \text{dom } q\}$;
- (ii) $R_j = \{q \in \text{Fn}(I, J) \mid j \in \text{ran } q\}$.

Definition. A subset G of a forcing poset \mathbb{P} is a *filter* if

- (i) $1 \in G$;
- (ii) for all $p, q \in G$ there is $r \in G$ such that $r \leq p$ and $r \leq q$;
- (iii) for all $p, q \in G$, if $q \leq p$ and $q \in G$ then $p \in G$.

A filter G is \mathbb{P} -*generic over* M if $G \cap D$ is nonempty for every \mathbb{P} -dense subset $D \in M$.

Lemma (generic filter existence lemma). Let M be an arbitrary countable set, and let $\mathbb{P} \in M$ be a forcing poset. Then for any condition $p \in \mathbb{P}$, there is a filter $G \subseteq \mathbb{P}$ containing p which is \mathbb{P} -generic over M .

Proof. Let $(D_n)_{n \in \omega}$ enumerate all dense subsets of \mathbb{P} which lie in M . We inductively define $X \subseteq \mathbb{P}$ by $X = \{q_n \mid n \in \omega\}$ as follows. Let $q_0 = p$, and given q_n , we choose $q_{n+1} \in D_n$ such that $q_{n+1} \leq q_n$. Finally, let $G = \{r \in \mathbb{P} \mid \exists n. q_n \leq r\}$. Then G is a filter as the q_n form a chain, and it is clearly generic. \square

Definition. A condition $p \in \mathbb{P}$ is *minimal* if whenever $q \leq p$, we have $q = p$.

Lemma. Let M be a countable transitive model of ZF, and let $\mathbb{P} \in M$ be a separative partial order. Then either \mathbb{P} has a minimal element, or for every filter G which is \mathbb{P} -generic over M , we have $G \notin M$.

Proof. Suppose \mathbb{P} has no minimal element. Let G be a \mathbb{P} -generic filter over M . We show that if $F \subseteq \mathbb{P}$ is a filter in M , then the set $D_F = \mathbb{P} \setminus F \in M$ is a dense set. Then $G \cap D_F$ is nonempty for all filters F , so G cannot be equal to any filter $F \in M$.

Fix $p \in \mathbb{P}$. If $p \notin F$, then $p \in D_F$ as required. Otherwise, suppose $p \in F$. As p is not minimal, we can fix some $q \in F$ with $q < p$. Then $p \not\leq q$, so by separativity, there is $r \leq p$ such that $r \perp q$. But all conditions in F are compatible, so one of r and q is not in F . \square

Proposition. For sets I, J such that $|I| \geq \omega$ and $|J| \geq 2$, the forcing poset $\text{Fn}(I, J)$ is a separative partial order without a minimal element.

Proposition. (ZFC) Let $\mathbb{P} \in M$ be a forcing poset, and let $G \subseteq \mathbb{P}$. Then the following are equivalent.

- (i) G is \mathbb{P} -generic over M , that is, for all dense sets $D \in M$, we have $G \cap D \neq \emptyset$;
- (ii) for all $p \in G$ and $D \in M$, if D is dense below p in \mathbb{P} , then $G \cap D \neq \emptyset$;
- (iii) for all open dense sets $D \in M$, we have $G \cap D \neq \emptyset$;
- (iv) for all $D \in M$ that are maximal antichains in \mathbb{P} , we have $G \cap D \neq \emptyset$.

3.5. Names

Definition. Let \mathbb{P} be a forcing poset. We define the class of \mathbb{P} -names $M^\mathbb{P}$ recursively as follows.

- (i) $M_0^\mathbb{P} = \emptyset$;
- (ii) $M_{\alpha+1}^\mathbb{P} = \mathcal{P}^M(\mathbb{P} \times M_\alpha^\mathbb{P})$;
- (iii) at limit stages λ , $M_\lambda^\mathbb{P} = \bigcup_{\alpha < \lambda} M_\alpha^\mathbb{P}$;
- (iv) $M^\mathbb{P} = \bigcup_{\alpha \in \text{Ord}} M_\alpha^\mathbb{P}$.

Being a \mathbb{P} -name is absolute for transitive models. \mathbb{P} -names are denoted with overdots, such as in \dot{x} .

Definition. The *range* of a \mathbb{P} -name \dot{x} is

$$\text{ran}(\dot{x}) = \{\dot{y} \mid \exists p \in \mathbb{P}. \langle p, \dot{y} \rangle \in \dot{x}\}$$

Remark. Alternatively, by transfinite recursion on rank, we could define the class of \mathbb{P} -names over V in the following way. If $\text{rank } x = \alpha$, then x is a \mathbb{P} -name if and only if it is a relation such that for all $\langle p, \dot{y} \rangle \in x$, we have $p \in \mathbb{P}$ and \dot{y} is a \mathbb{P} -name in V_α . Finally, $M^\mathbb{P} = V^\mathbb{P} \cap M$.

Definition. The \mathbb{P} -rank of a name \dot{x} , written $\text{rank}_\mathbb{P} \dot{x}$, is the least α such that $\dot{x} \subseteq \mathbb{P} \times M_\alpha^\mathbb{P}$.

Definition. Let \dot{x} be a \mathbb{P} -name and G be an arbitrary subset of \mathbb{P} . We define the *interpretation of \dot{x} by G* recursively by

$$\dot{x}^G = \{\dot{y}^G \mid \exists p \in G. \langle p, \dot{y} \rangle \in \dot{x}\}$$

Definition. The *forcing extension of M by G* , written $M[G]$, is

$$M[G] = \{\dot{x}^G \mid \dot{x} \in M^\mathbb{P}\}$$

VII. Forcing and the Continuum Hypothesis

Example. If $\emptyset \in M$, then $\emptyset^G = \emptyset$. Let

$$\dot{x} = \{\langle p, \emptyset \rangle, \langle r, \{\langle q, \emptyset \rangle\} \rangle\}$$

If $p, q, r \in G$, then

$$\begin{aligned}\dot{x}^G &= \{(\langle p, \emptyset \rangle)^G, (\langle r, \{\langle q, \emptyset \rangle\} \rangle)^G\} \\ &= \{\emptyset, \{(\langle q, \emptyset \rangle)^G\}\} \\ &= \{\emptyset, \{\emptyset\}\}\end{aligned}$$

If $p, r \notin G$, then

$$\dot{x}^G = \emptyset$$

If $r \in G$ but $p, q \notin G$, then

$$\dot{x}^G = \{(\langle q, \emptyset \rangle)^G\} = \{\emptyset\}$$

Finally, if $p \in G$ but $r \notin G$, then

$$\dot{x}^G = \{\emptyset\}$$

We aim to show the following major theorem.

Theorem (generic model theorem). Let M be a countable transitive model of ZF, let \mathbb{P} be a forcing poset, and let G be a \mathbb{P} -generic filter. Then

- (i) $M[G]$ is a transitive set;
- (ii) $|M[G]| = \aleph_0$;
- (iii) $M[G] \models \text{ZF}$, and if $M \models \text{AC}$ then $M[G] \models \text{AC}$;
- (iv) $\text{Ord}^M = \text{Ord}^{M[G]}$;
- (v) $M \subseteq M[G]$;
- (vi) $M[G]$ is the smallest countable transitive model of ZF such that $M \subseteq M[G]$ and G is a set in $M[G]$.

Countability is only needed to show the existence of a generic filter, so parts (i) and (iii)–(vi) of this theorem hold without this assumption.

3.6. Canonical names

We can prove some parts of the generic model theorem by introducing the notion of *canonical names*.

Definition. Given a forcing poset $(\mathbb{P}, \leq, \mathbb{1})$ and a set $x \in M$, we define the *canonical name* of x by

$$\check{x} = \{\langle \mathbb{1}, \check{y} \rangle \mid y \in x\}$$

The symbol \check{x} is pronounced *x-check*.

Lemma. If M is a transitive model of ZF, $\mathbb{P} \in M$, and $1 \in G \subseteq \mathbb{P}$, then

- for all $x \in M$, $\check{x} \in M^{\mathbb{P}}$ and $\check{x}^G = x$;
- $M \subseteq M[G]$;
- $M[G]$ is transitive.

Proof. Part (i). We show $\check{x} \in M^{\mathbb{P}}$ by induction, using the definition of \mathbb{P} -names by transfinite recursion. Hence

$$\check{x}^G = \{\check{y}^G \mid y \in x\} = \{y \mid y \in x\} = x$$

Part (ii) follows directly from part (i).

Part (iii). Suppose that $x \in y$ and $y \in M[G]$. By definition, $y = \dot{y}^G$ for some \mathbb{P} -name \dot{y} . By construction, any element of \dot{y} is of the form \dot{z}^G , so in particular, $x = \dot{x}^G$ for some \mathbb{P} -name $\dot{x} \in M^{\mathbb{P}}$. \square

Remark. Even if $G \notin M$, we can still define a name for G in M . From this, it follows that if $G \notin M$, then $M[G] \neq M$.

Proposition. Let

$$\dot{G} = \{\langle p, \check{p} \rangle \mid p \in \mathbb{P}\}$$

Then $\dot{G}^G = G$.

Proof.

$$\dot{G}^G = \{\check{p}^G \mid p \in G\} = \{p \mid p \in G\} = G$$

\square

3.7. Verifying the axioms: part one

We can define unordered and ordered pairs of names, with sensible interpretations.

Definition. Given \mathbb{P} -names \dot{x}, \dot{y} , let

$$\text{up}(\dot{x}, \dot{y}) = \{\langle 1, \dot{x} \rangle, \langle 1, \dot{y} \rangle\}$$

and

$$\text{op}(\dot{x}, \dot{y}) = \text{up}(\text{up}(\dot{x}, \dot{x}), \text{up}(\dot{x}, \dot{y}))$$

Proposition. For $\dot{x}, \dot{y} \in M^{\mathbb{P}}$ and $1 \in G \subseteq \mathbb{P}$,

$$(\text{up}(\dot{x}, \dot{y}))^G = \{\dot{x}^G, \dot{y}^G\}$$

and

$$(\text{op}(\dot{x}, \dot{y}))^G = \langle \dot{x}^G, \dot{y}^G \rangle$$

VII. Forcing and the Continuum Hypothesis

Lemma. Suppose M is a transitive model of ZF and $\mathbb{P} \in M$ is a forcing poset. If $1 \in G \subseteq \mathbb{P}$, then $M[G]$ is a transitive model of extensionality, empty set, foundation, and pairing.

Lemma. Suppose that M is a transitive model of ZF and $\mathbb{P} \in M$ is a forcing poset. Let $G \subseteq \mathbb{P}$ be such that $1 \in G$. Then

- (i) $\text{rank}(\dot{x}^G) \leq \text{rank } \dot{x}$ for all $\dot{x} \in M^{\mathbb{P}}$;
- (ii) $\text{Ord}^M = \text{Ord}^{M[G]}$;
- (iii) $|M[G]| = |M|$.

Proof. *Part (i).* We show this result by induction on x . $\emptyset^G = \emptyset$, and both have rank 0. We have

$$\begin{aligned} \text{rank}(\dot{x}^G) &= \sup \{ \text{rank } u + 1 \mid u \in \dot{x}^G \} \\ &\leq \sup \{ \text{rank}(\dot{y}^G) + 1 \mid \dot{y} \in \text{ran } \dot{x} \} \\ &\leq \sup \{ \text{rank } \dot{y} + 1 \mid \dot{y} \in \text{ran } \dot{x} \} \\ &\leq \sup \{ \text{rank } u + 1 \mid u \in \dot{x} \} \\ &\leq \text{rank } \dot{x} \end{aligned}$$

Part (ii). Since $M \subseteq M[G]$ and being an ordinal is absolute, $\text{Ord}^M \subseteq \text{Ord}^{M[G]}$. For the reverse inclusion, suppose $\alpha \in M[G]$ is an ordinal, and fix a name $\dot{x} \in M^{\mathbb{P}}$ such that $\alpha = \dot{x}^G$. Then α is an ordinal in the universe, so

$$\alpha = \text{rank } \alpha \leq \text{rank } \dot{x}$$

so since M is transitive, $\alpha \in \text{Ord}^M$.

Part (iii). Since any element of $M[G]$ is of the form \dot{x}^G for some $\dot{x} \in M^{\mathbb{P}} \subseteq M \subseteq M[G]$, we must have

$$|M[G]| \leq |M^{\mathbb{P}}| \leq |M| \leq |M[G]|$$

so the inequalities must be equalities. □

Corollary. $M[G]$ satisfies the axiom of infinity.

Proof. $\omega \in \text{Ord}^M$ so $\omega \in \text{Ord}^{M[G]} \subseteq M[G]$. □

Lemma. Suppose M is a transitive model of ZF, $\mathbb{P} \in M$ is a forcing poset, and $G \subseteq \mathbb{P}$ is such that $1 \in G$. Then if N is another transitive model of ZF with $M \subseteq N$ a definable class in N and $G \in N$, then $M[G] \subseteq N$.

Proof. We carry out the construction of $M[G]$ in N . Namely, we will show that for all \mathbb{P} -names \dot{x} , we have $\dot{x}^G \in N$, from which it follows that $M[G] \subseteq N$. We proceed by induction

on x . As the axiom of empty set holds in N and it is a transitive set, $\emptyset^G = \emptyset \in N$. Moreover, since

$$M^{\mathbb{P}} = V^{\mathbb{P}} \cap M \subseteq V^{\mathbb{P}} \cap N = N^{\mathbb{P}}$$

if \dot{x} is a \mathbb{P} -name of M , it must be a \mathbb{P} -name of M . In particular, $x \in N$. Now, suppose that for every $\langle p, \dot{y} \rangle \in \dot{x}$, we have $\dot{y}^G \in N$. Then

$$\begin{aligned} (\dot{x}^G)^N &= \{\dot{y}^G \mid \exists p \in G. \langle p, \dot{y} \rangle \in \dot{x}\}^N \\ &= \{(\dot{y}^G)^N \mid (\exists p \in G. \langle p, \dot{y} \rangle \in \dot{x})^N\} \\ &= \{\dot{y}^G \mid \exists p \in G. \langle p, \dot{y} \rangle \in \dot{x}\} \\ &= \dot{x}^G \end{aligned}$$

Thus $\dot{x}^G \in N$ as required. \square

To prove the generic model theorem, it now suffices to prove the remaining axioms of ZF, which are union, power set, replacement, and separation. We can prove the axiom of union now.

Lemma. Suppose M is a transitive model of ZF, $\mathbb{P} \in M$ is a forcing poset, and $G \subseteq \mathbb{P}$ is such that $1 \in G$. Additionally, suppose that G is a filter. Then $M[G]$ satisfies the axiom of union.

Proof. It suffices to prove that for all $a \in M[G]$, there is some $b \in M[G]$ such that $\bigcup a = b$. Fix $\dot{a} \in M^{\mathbb{P}}$ such that $\dot{a}^G = a$, and let \dot{b} be the following name.

$$\dot{b} = \{\langle p, \dot{z} \rangle \mid \exists \langle q, \dot{y} \rangle \in \dot{a}. \exists r \in \mathbb{P}. \langle r, \dot{z} \rangle \in \dot{y} \wedge p \leq r, q\}$$

Observe that \dot{b} is a \mathbb{P} -name in M : since \dot{a} is a \mathbb{P} -name, any $\dot{y} \in \text{ran } \dot{a}$ is a \mathbb{P} -name, so \dot{b} consists of pairs $\langle p, \dot{z} \rangle$ where $p \in \mathbb{P}$ and $\dot{z} \in \text{ran } \dot{y}$ for some $\dot{y} \in \text{ran } \dot{a}$. Thus \dot{z} is a \mathbb{P} -name in V . Moreover $\dot{b} \in M$ since $\dot{b} \in \mathbb{P} \times \text{tcl}(\dot{a})$.

We claim that $\bigcup a \subseteq \dot{b}^G$. Let $w \in \bigcup a$, so $w \in v$ for some $v \in a$. Since $M[G]$ is transitive, we can fix names \dot{y}, \dot{z} and conditions $q, r \in G$ such that

$$\dot{y}^G = v; \quad \dot{z}^G = w; \quad \langle q, \dot{y} \rangle \in \dot{a}; \quad \langle r, \dot{z} \rangle \in \dot{y}$$

As G is a filter, by directedness there is a condition $p \leq q, r$ in G . Then, by definition, $\langle p, \dot{z} \rangle \in \dot{b}$, and $w = \dot{z}^G \in \dot{b}^G$.

For the converse, we claim that $\dot{b}^G \subseteq \bigcup a$. Let $\langle p, \dot{z} \rangle \in \dot{b}^G$, so $p \in G$ and $\dot{z}^G = c$. By definition, we can fix $\langle q, \dot{y} \rangle \in \dot{a}$ and $r \in \mathbb{P}$ such that $\langle r, \dot{z} \rangle \in \dot{y}$ and $p \leq q, r$. Using the fact that G is a filter, we must have $q, r \in G$. Hence $\dot{z}^G \in \dot{y}^G$ and $\dot{y}^G \in \dot{a}^G$, so $c \in \dot{y}^G$ for some $\dot{y}^G \in a$. \square

Example (motivation for genericity). Note that $\mathbb{P}, G \in M[G]$. If $M[G]$ models any reasonable theory, we should have $\mathbb{P} \setminus G \in M[G]$. We will try to build a name for $\mathbb{P} \setminus G$. A natural name to consider is

$$\dot{c} = \{\langle q, \dot{p} \rangle \mid p, q \in \mathbb{P}, p \perp q\}$$

VII. Forcing and the Continuum Hypothesis

Then

$$\dot{c}^G = \{p \mid \exists q \in G. p \perp q\}$$

If G is a filter, its elements are pairwise compatible, so $G \cap \dot{c}^G = \emptyset$. But we still need to show that $G \cup \dot{c}^G = \mathbb{P}$. For each condition p , set

$$D_p = \{q \in \mathbb{P} \mid p \perp q \vee q \leq p\}$$

It is easy to check that $D_p \in M$ is dense. Now, if G is \mathbb{P} -generic, we could fix some $q \in G \cap D_p$ for any given p . Then if $p \perp q$, by definition $p \in \dot{c}^G$, and if $q \leq p$, then $p \in G$ by upwards closure. From this, it follows that $G \cup \dot{c}^G = \mathbb{P}$.

In fact, we have the following.

Proposition. Let M be a countable transitive model of ZF. Then there exists a forcing poset $\mathbb{P} \in M$ and a (non-generic) filter $G \subseteq \mathbb{P}$ such that $\mathbb{P} \setminus G \notin M[G]$.

3.8. The forcing relation

To show separation, we need to show that if $\varphi(x, y)$ is a formula and \dot{a}, \dot{b} are \mathbb{P} -names, then

$$C = \{\dot{z}^G \in \dot{a}^G \mid (\varphi(\dot{z}^G, \dot{b}^G))^{M[G]}\} \in M[G]$$

This is unclear, even for simple formulas such as $\varphi(x, y) \equiv x \notin y$. We will build a way to formally reason about $M[G]$ from within M , without having to rely on G . To do this, we will define a relation $p \Vdash \varphi$ between conditions $p \in \mathbb{P}$ and names in $V^{\mathbb{P}}$. Its relativisation $(p \Vdash \varphi)^M$ will provide a way to work in M . Our aim is to define \Vdash such that $p \Vdash \varphi(\dot{u})$ if and only if for every generic subset $G \subseteq \mathbb{P}$ with $p \in G$, we have $M[G] \models \varphi(\dot{u}^G)$.

Naively, we might say that if $\langle p, \dot{x} \rangle \in \dot{y}$ then $p \Vdash \dot{x} \in \dot{y}$. The converse cannot be made to hold. Consider $\dot{x} = \{\langle p, \emptyset \rangle\}$ where $p \neq 1$. Then $p \Vdash \emptyset \in \dot{x}$. Suppose $q \perp p$, then we have $q \Vdash \dot{x} = \emptyset$. Therefore, we should have $q \Vdash \dot{x} \in \check{1}$. If we enforce the converse above, we would have $\langle q, \dot{x} \rangle \in \check{1}$, which is incorrect since $\check{1} = \{\langle 1, \emptyset \rangle\}$. Instead, we will define the forcing relation in terms of dense sets, leveraging the fact that generics meet all dense sets.

Definition. Let \mathbb{P} be a forcing poset. The \mathbb{P} -forcing language $\mathcal{FL}_{\mathbb{P}}$ is the class of logical formulas formed using the binary relation \in and constant symbols from $V^{\mathbb{P}}$.

Definition. Let \mathbb{P} be a forcing poset and let $p \in \mathbb{P}$. Let $\dot{x}, \dot{y}, \dot{u}$ be \mathbb{P} -names in V . We define the forcing relation $p \Vdash \varphi(\dot{u})$ recursively as follows.

- (i) $p \Vdash \varphi(\dot{u}) \wedge \psi(\dot{u})$ if and only if $p \Vdash \varphi(\dot{u})$ and $p \Vdash \psi(\dot{u})$;
- (ii) $p \Vdash \neg \varphi(\dot{u})$ if and only if there is no $q \leq p$ such that $q \Vdash \varphi(\dot{u})$;
- (iii) $p \Vdash \exists x. \varphi(x, \dot{u})$ if and only if the set

$$\{q \leq p \mid \exists \dot{x} \in V^{\mathbb{P}}. q \Vdash \varphi(\dot{x}, \dot{u})\}$$

is dense below p ;

(iv) $p \Vdash \dot{x} \in \dot{y}$ if and only if the set

$$\{q \leq p \mid \exists \langle r, \dot{z} \rangle \in \dot{y}. q \leq r \wedge (q \Vdash \dot{x} = \dot{z})\}$$

is dense below p ;

(v) $p \Vdash \dot{x} \subseteq \dot{y}$ if and only if for all $\langle q_1, \dot{z}_1 \rangle \in \dot{x}$, the set

$$\{r \leq p \mid r \leq q_1 \rightarrow \exists \langle q_2, \dot{z}_2 \rangle \in \dot{y}. r \leq q_2 \wedge (r \Vdash \dot{z}_1 = \dot{z}_2)\}$$

is dense below p ; and

(vi) $p \Vdash \dot{x} = \dot{y}$ if and only if $p \Vdash \dot{x} \subseteq \dot{y}$ and $p \Vdash \dot{y} \subseteq \dot{x}$.

Remark. (i) The definitions for \subseteq and $=$ are defined recursively, and thus require transfinite recursion to define formally.

(ii) All of the clauses except for the existential use only absolute notions. In particular, it does not depend on M . When relativising to a model, $(p \Vdash \exists x. \varphi(x))^M$ precisely when the set

$$\{q \leq p \mid \exists \dot{x} \in M^{\mathbb{P}}. q \Vdash \varphi(\dot{x}, \dot{u})\}$$

is dense below p .

Proposition. Let p be a condition, φ be an $\mathcal{FL}_{\mathbb{P}}$ -formula, and $\dot{x}_1, \dots, \dot{x}_n$ be \mathbb{P} -names in V . Then the following are equivalent.

- (i) $p \Vdash \varphi(\dot{x}_1, \dots, \dot{x}_n)$;
- (ii) for all $q \leq p$, $q \Vdash \varphi(\dot{x}_1, \dots, \dot{x}_n)$;
- (iii) there is no $q \leq p$ such that $q \Vdash \neg \varphi(\dot{x}_1, \dots, \dot{x}_n)$;
- (iv) the set $\{r \mid r \Vdash \varphi(\dot{x}_1, \dots, \dot{x}_n)\}$ is dense below p .

Proof. (ii) implies (iii). If (iii) did not hold, there would be some $q \leq p$ such that $q \Vdash \neg \varphi$. Then there is no $r \leq q$ such that $r \Vdash \varphi$. So in particular, $q \nVdash \varphi$, contradicting (ii).

(iii) implies (iv). Suppose that there is no $q \leq p$ such that $q \Vdash \neg \varphi$. Take $q \leq p$. Then by assumption, $q \nVdash \neg \varphi$, so there is $r \leq q$ such that $r \Vdash \varphi$, so the set is dense as required.

(i) implies (ii). We show this by induction on formula complexity.

- For atomic formulas, let \square be either \in or \subseteq . Then $p \Vdash \dot{x} \square \dot{y}$ if and only if some set A is dense below p . Take $q \leq p$, then A is dense below q . Then $q \Vdash \dot{x} \square \dot{y}$ as required.
- If $p \Vdash \neg \varphi$, then there is no $r \leq p$ such that $r \Vdash \varphi$. Then there is no $r \leq q$ such that $r \Vdash \varphi$, so by definition, $q \Vdash \neg \varphi$.
- If $p \Vdash \varphi \wedge \psi$ then $p \Vdash \varphi$ and $p \Vdash \psi$, so by the inductive hypothesis, $q \Vdash \varphi$ and $q \Vdash \psi$, giving $q \Vdash \varphi \wedge \psi$.

VII. Forcing and the Continuum Hypothesis

- If $p \Vdash \exists x. \varphi(x)$, then A is dense below p for some set A , but then A is dense below q , so $q \Vdash \exists x. \varphi(x)$.

(iv) *implies (i)*. Again, we show this by induction.

- For atomic formulas, let \Box be either \in or \subseteq . To prove that $p \Vdash \dot{x} \Box \dot{y}$, we must show that some set A is dense below p . By assumption, the set $\{r \mid r \Vdash \dot{x} \Box \dot{y}\}$ is dense below p . Fix $q \leq p$, then there is $r \leq q$ such that $r \Vdash \dot{x} \Box \dot{y}$. Hence there is some $s \leq r \leq q \leq p$ such that $s \in A$. Therefore $p \Vdash \dot{x} \Box \dot{y}$ as required. The proof for existentials is the same.
- Suppose that $\{r \mid r \Vdash \varphi \wedge \psi\}$ is dense below p . So $\{r \mid r \Vdash \varphi\}$ and $\{r \mid r \Vdash \psi\}$ are also dense below p . By the inductive hypothesis, $p \Vdash \varphi$ and $p \Vdash \psi$. Hence $p \Vdash \varphi \wedge \psi$.
- Suppose that $\{r \mid r \Vdash \neg\varphi\}$ is dense below p . To show $p \Vdash \neg\varphi$, we fix $q \leq p$ and suppose $q \Vdash \varphi$. By the fact that (i) implies (iii), there is no $r \leq q$ such that $r \Vdash \neg\varphi$, contradicting density of the set $\{r \mid r \Vdash \neg\varphi\}$.

□

Proposition. Let \mathbb{P} be a forcing poset, let $p, q \in \mathbb{P}$, and let $\dot{a}, \dot{b} \in V^{\mathbb{P}}$. Then

- (i) $p \Vdash \dot{a} = \dot{a}$;
- (ii) if $\langle q, \dot{b} \rangle \in \dot{a}$ and $p \leq q$, then $p \Vdash \dot{b} \in \dot{a}$;
- (iii) if M is a transitive model of ZF and $\mathbb{P} \in M$, then for any φ, ψ ,

$$\{\langle q, \dot{x} \rangle \mid \langle q, \dot{x} \rangle \in \dot{a} \wedge (q \Vdash \varphi(\dot{x}))^M\} \in M$$

and

$$\{q \in \mathbb{P} \mid (q \Vdash \psi(\dot{a}))^M\} \in M$$

- (iv) $p \Vdash \varphi \vee \psi$ if and only if

$$\{q \leq p \mid q \Vdash \varphi \text{ or } q \Vdash \psi\}$$

is dense below p ;

- (v) $p \Vdash \varphi \rightarrow \psi$ if and only if there is no $q \leq p$ such that $q \Vdash \varphi$ and $q \Vdash \neg\psi$;
- (vi) $p \Vdash \forall x. \varphi(x)$ if and only if for all $\dot{x} \in V^{\mathbb{P}}$, $p \Vdash \varphi(\dot{x})$;
- (vii) for any φ , the set

$$\{p \in \mathbb{P} \mid p \Vdash \varphi \text{ or } p \Vdash \neg\varphi\}$$

is a dense open set;

- (viii) there is no p and formula φ such that

$$p \Vdash \varphi \wedge \neg\varphi$$

3.9. The forcing theorem

Theorem (the forcing theorem). Suppose M be a transitive model of ZF, $\mathbb{P} \in M$ is a forcing poset, $\varphi(u)$ is a formula, and G is \mathbb{P} -generic over M . Then for any $\dot{x} \in M^{\mathbb{P}}$,

- (i) if $p \in G$ and $(p \Vdash \varphi(\dot{x}))^M$, then $M[G] \models \varphi(\dot{x}^G)$; and
- (ii) if $M[G] \models \varphi(\dot{x}^G)$, then there is a condition $p \in G$ such that $(p \Vdash \varphi(\dot{x}))^M$.

Once we have shown this theorem, we will have the following result.

Corollary. Suppose that M is a countable transitive model of ZF, $\mathbb{P} \in M$ is a forcing poset, and $\varphi(u)$ is a formula. Then for any name $\dot{x} \in M^{\mathbb{P}}$,

$$(p \Vdash \varphi(\dot{x}))^M \leftrightarrow \text{for any } \mathbb{P}\text{-generic filter } G \text{ with } p \in G, M[G] \models \varphi(\dot{x}^G)$$

The only reason we need countability is so that every condition is contained in a generic filter.

Proof. The forward direction is part (i) of the forcing theorem. For the backward direction, suppose that $(p \nVdash \varphi(\dot{x}))^M$. Then, by definition, there is some $q \leq p$ such that $(q \Vdash \neg \varphi(\dot{x}))^M$. Let G be a \mathbb{P} -generic filter over M such that $q \in G$. Then, since G is upwards closed, $p \in G$. Hence $M[G] \models \varphi(\dot{x}^G)$ by assumption. But as $q \in G$, by the forcing theorem we obtain $M[G] \models \neg \varphi(\dot{x}^G)$. This contradicts part (viii) of the proposition above by the forcing theorem. \square

Definition. Suppose M is a countable transitive model of ZF, $\mathbb{P} \in M$ is a forcing poset, $\dot{x}_1, \dots, \dot{x}_n \in M^{\mathbb{P}}$, $p \in \mathbb{P}$, and $\varphi(v_1, \dots, v_n)$ is a formula. Then we can define a relation $\Vdash_{\mathbb{P}, M}^*$ by

$$p \Vdash_{\mathbb{P}, M}^* \varphi(\dot{x}_1, \dots, \dot{x}_n)$$

if and only if $M[G] \models \varphi(\dot{x}_1^G, \dots, \dot{x}_n^G)$ for all $G \subseteq \mathbb{P}$ such that $p \in G$ and G is a \mathbb{P} -generic filter.

Corollary. $p \Vdash \varphi \leftrightarrow p \Vdash_{\mathbb{P}, M}^* \varphi$.

We will now prove the forcing theorem.

Proof. We show the result by induction on the complexity of formulas. Note that we need to work with relativised formulas with parameters $(p \Vdash \varphi(\mathbf{v}))^M$, but this only changes the existential case, so for all other cases we will suppress the relativisation and the parameters. We write $\Psi(\varphi)$ for the claim that for any name $\dot{x} \in M^{\mathbb{P}}$, if $p \in G$ and $(p \Vdash \varphi(\dot{x}))^M$, then $M[G] \models \varphi(\dot{x}^G)$, and if $M[G] \models \varphi(\dot{x}^G)$, then there exists $p \in G$ such that $(p \Vdash \varphi(\dot{x}))^M$.

Part (i): negations. Suppose $\Psi(\varphi)$ holds. Let $p \in G$ and $p \Vdash \neg \varphi$. Suppose for a contradiction that $M[G] \models \varphi$, or equivalently, $\varphi^{M[G]}$. Then as $\Psi(\varphi)$ holds, there is $q \in G$ such that $q \Vdash \varphi$. As G is a filter, there is $r \in G$ such that $r \leq p, q$. Then $r \Vdash \varphi$, which contradicts the definition of $p \Vdash \neg \varphi$. Hence $\neg(\varphi^{M[G]})$, so by definition $(\neg \varphi)^{M[G]}$, so $M[G] \models \neg \varphi$.

VII. Forcing and the Continuum Hypothesis

For the converse, suppose $M[G] \models \neg\varphi$. Let

$$D = \{p \in \mathbb{P} \mid p \Vdash \varphi \vee p \Vdash \neg\varphi\}$$

Then D is dense, because if $q \nVdash \varphi$, then there is $p \leq q$ such that $p \Vdash \neg\varphi$, and $p \in D$. So as G is generic, we can fix $p \in G \cap D$. If $p \Vdash \varphi$, then by $\Psi(\varphi)$ we must have $M[G] \models \varphi$, but we assumed $M[G] \models \neg\varphi$. Hence $p \Vdash \neg\varphi$.

Part (ii): conjunctions. Suppose $\Psi(\varphi)$ and $\Psi(\psi)$. Suppose $p \Vdash \varphi \wedge \psi$ for some $p \in G$, so by definition, $p \Vdash \varphi$ and $p \Vdash \psi$. By $\Psi(\varphi)$ and $\Psi(\psi)$, we have $M[G] \models \varphi$ and $M[G] \models \psi$. So $M[G] \models \varphi \wedge \psi$.

For the converse, suppose $M[G] \models \varphi \wedge \psi$. Then $M[G] \models \varphi$ and $M[G] \models \psi$, so there are $p, q \in G$ such that $p \Vdash \varphi$ and $q \Vdash \psi$. But G is a filter, so there is $r \leq p, q$ such that $r \Vdash \varphi$ and $r \Vdash \psi$. Hence $r \Vdash \varphi \wedge \psi$, as required.

Part (iii): existential quantifiers. For this case, we will not suppress relativisation and parameters. Suppose $\Psi(\varphi(\dot{x}))$; we show $\Psi(\exists x. \varphi(x))$. To be more precise, for all names $\dot{x} \in M^{\mathbb{P}}$, we assume the forcing theorem holds for $\varphi(\dot{x})$. Suppose $p \in G$ is such that $(p \Vdash \exists x. \varphi(x))^M$. Let

$$D = (\{q \leq p \mid \exists \dot{x} \in V^{\mathbb{P}}. (q \Vdash \varphi(\dot{x}))\})^M = \{q \leq p \mid \exists \dot{x} \in M^{\mathbb{P}}. (q \Vdash \varphi(\dot{x}))^M\} \in M$$

By definition of forcing existentials, D is a dense set. Since G is generic, there is some $q \in G \cap D$. Then we can fix some \mathbb{P} -name \dot{x} such that $(q \Vdash \varphi(\dot{x}))^M$. Since the forcing theorem holds for $\varphi(\dot{x})$, we have $M[G] \models \varphi(\dot{x}^G)$. Hence $M[G] \models \exists x. \varphi(x)$.

Now suppose $M[G] \models \exists x. \varphi(x)$. We can fix $\dot{x} \in M^{\mathbb{P}}$ such that $M[G] \models \varphi(\dot{x}^G)$. By the fact that $\Psi(\varphi(\dot{x}))$ holds, there is a condition p such that $(p \Vdash \varphi(\dot{x}))^M$. Then

$$\{q \leq p \mid (q \Vdash \varphi(\dot{x}))^M\}$$

is dense. Hence, by definition, $(p \Vdash \exists x. \varphi(x))^M$.

Part (iv): equality. Recall that $p \Vdash \dot{x} = \dot{y}$ if and only if

- (a) for all $\langle q_1, \dot{z}_1 \rangle \in \dot{x}$, $\{r \leq p \mid r \leq q_1 \rightarrow \exists \langle q_2, \dot{z}_2 \rangle \in \dot{y}. r \leq q_2 \wedge (r \Vdash \dot{z}_1 = \dot{z}_2)\}$ is dense below p ; and
- (b) for all $\langle q_2, \dot{z}_2 \rangle \in \dot{y}$, $\{r \leq p \mid r \leq q_2 \rightarrow \exists \langle q_1, \dot{z}_1 \rangle \in \dot{x}. r \leq q_1 \wedge (r \Vdash \dot{z}_1 = \dot{z}_2)\}$ is dense below p .

We show that for any \dot{x}, \dot{y} , we have $\Psi(\dot{x} = \dot{y})$. We will show this by transfinite induction on the pair $\langle \dot{x}, \dot{y} \rangle$ ordered lexicographically.

Suppose that $p \Vdash \dot{x} = \dot{y}$ and $p \in G$. We show $M[G] \models \dot{x}^G \subseteq \dot{y}^G$; the converse holds by symmetry, and then we obtain $M[G] \models \dot{x}^G = \dot{y}^G$ by extensionality. Any element of \dot{x}^G is of the form \dot{z}_1^G where $\langle q_1, \dot{z}_1 \rangle \in \dot{x}$ and $q_1 \in G$. Since G is a filter, we can fix $s \in G$ such that $s \leq p, q_1$. Then, as $s \leq p$, we have $s \Vdash \dot{x} = \dot{y}$, so the set in (a) above is dense below s . Hence there is $r \in G$ such that $r \leq s \leq q_1$ and there exists $\langle q_2, \dot{z}_2 \rangle \in \dot{y}$ such that $r \leq q_2$

and $r \Vdash \dot{z}_1 = \dot{z}_2$. As G is a filter, $q_2 \in G$, so $\dot{z}_2^G \in \dot{y}^G$. By using the inductive hypothesis on $\langle \dot{z}_1, \dot{z}_2 \rangle$, as $r \in G$ we have $M[G] \models \dot{z}_1^G = \dot{z}_2^G$. Hence $\dot{z}_1^G \in \dot{y}^G$, so $\dot{x}^G \subseteq \dot{y}^G$.

For the converse, $M[G] \models \dot{x}^G = \dot{y}^G$. Define D to be the set of $r \in \mathbb{P}$ such that at least one of the following hold.

- (0) $r \Vdash \dot{x} = \dot{y}$;
- (a') there exists $\langle q_1, \dot{z}_1 \rangle \in \dot{x}$ such that $r \leq q_1$ and for all $\langle q_2, \dot{z}_2 \rangle \in \dot{y}$ and $s \in \mathbb{P}$, if $s \leq q_2$ and $s \Vdash \dot{z}_1 = \dot{z}_2$ then $s \perp r$;
- (b') there exists $\langle q_2, \dot{z}_2 \rangle \in \dot{y}$ such that $r \leq q_2$ and for all $\langle q_1, \dot{z}_1 \rangle \in \dot{x}$ and $s \in \mathbb{P}$, if $s \leq q_1$ and $s \Vdash \dot{z}_1 = \dot{z}_2$ then $s \perp r$.

Note that by separation in M and absoluteness, D is a set in M . We claim that D is dense. Fix $p \in \mathbb{P}$, and suppose $p \nVdash \dot{x} = \dot{y}$. Then at least one of (a) and (b) above fails. Suppose that the set in (a) fails; the result for (b) holds by symmetry. Then there is $\langle q_1, \dot{z}_1 \rangle \in \dot{x}$ such that

$$\{r \leq p \mid r \leq q_1 \rightarrow \exists \langle q_2, \dot{z}_2 \rangle \in \dot{y}. r \leq q_2 \wedge (r \Vdash \dot{z}_1 = \dot{z}_2)\}$$

is not dense below p . Then there is $s \leq p$ such that for all $r \leq s$, we have $r \leq q_1$, and for all $\langle q_2, \dot{z}_2 \rangle \in \dot{y}$ such that $\neg((r \Vdash \dot{z}_1 = \dot{z}_2) \wedge r \leq q_2)$. In particular, this gives $s \leq q_1$. Now, if $\langle q_1, \dot{z}_2 \rangle \in \dot{y}$, $r \leq q_2$, and $r \Vdash \dot{z}_1 = \dot{z}_2$, then it must be the case that $s \perp r$, as any common extension of s and r would contradict the fact that the set in (a) was not dense. Thus $s \leq p$ and s satisfies (a'). Hence D is dense.

D is dense below $p \in G$ and G is \mathbb{P} -generic so we can fix $r \in G \cap D$. We will show that r satisfies (0), which finishes the proof. Suppose not, so suppose r satisfies (a') without loss of generality. Then we can fix $\langle q_1, \dot{z}_1 \rangle \in \dot{x}$ such that $r \leq q_1$ and for all $\langle q_2, \dot{z}_2 \rangle \in \dot{y}$ such that for all $s \in \mathbb{P}$ with $s \leq q_2$ and $s \Vdash \dot{z}_1 = \dot{z}_2$, we have $s \perp r$. Since $r \in G$ and $r \leq q_1$, we must have $q_1 \in G$ by upwards closure. Therefore, $M[G] \models \dot{z}_1^G \in \dot{x}^G = \dot{y}^G$. So we can fix $\langle q_2, \dot{z}_2 \rangle \in \dot{y}$ such that $q_2 \in G$ and $M[G] \models \dot{z}_1^G = \dot{z}_2^G$. By the inductive hypothesis, we can fix $p' \in G$ such that $p' \Vdash \dot{z}_1 = \dot{z}_2$. Since G is a filter and both $p', q_2 \in G$, we obtain $s \in G$ with $s \leq p', q_2$. Hence $s \Vdash \dot{z}_1 = \dot{z}_2$. Hence, by (a'), we have $s \perp r$. But $s, r \in G$, so $s \parallel r$, giving a contradiction.

Part (v): membership. Suppose that $p \Vdash \dot{x} \in \dot{y}$ for $p \in G$. Let

$$D = \{q \leq p \mid \exists \langle r, \dot{z} \rangle \in \dot{y}. q \leq r \wedge (q \Vdash \dot{x} = \dot{z})\}$$

By definition, D is dense. We can fix $q \in G \cap D$. Since $q \in D$, we may also fix $\langle r, \dot{z} \rangle \in \dot{y}$ such that $q \leq r$ and $q \Vdash \dot{x} = \dot{z}$. As $q \in G$, by the forcing theorem for equality, $M[G] \models \dot{x}^G = \dot{z}^G$. Since G is a filter and $q \leq r$, then $r \in G$ and so $\dot{z}^G \in \dot{y}^G$. Hence $M[G] \models \dot{x}^G \in \dot{y}^G$.

Now suppose $M[G] \models \dot{x}^G \in \dot{y}^G$. Fix $\langle r, \dot{z} \rangle \in \dot{y}$ such that $r \in G$ and $\dot{z}^G = \dot{x}^G$. Now, by the forcing theorem for equality, there is $q \in G$ such that $q \Vdash \dot{x} = \dot{z}$. Since G is a filter and $q, r \in G$, we can fix $p \in G$ such that $p \leq q, r$. Then $p \Vdash \dot{z} \in \dot{y}$ and $p \Vdash \dot{x} = \dot{z}$. So for all $s \leq p$, we have $s \leq r$ and $s \Vdash \dot{x} = \dot{z}$, so D is dense below p . Hence $p \Vdash \dot{x} \in \dot{y}$, as required. \square

3.10. Verifying the axioms: part two

Lemma. Suppose that M is a countable transitive model of ZF, $\mathbb{P} \in M$ is a forcing poset, and $G \subseteq \mathbb{P}$ is a generic filter. Then $M[G]$ models separation.

Proof. Let $\varphi(x, v)$ be a formula with free variables x, v . It suffices to show that for any $a, v \in M[G]$,

$$b = \{x \in a \mid M[G] \models \varphi(x, v)\} \in M[G]$$

Fix names \dot{a}, \dot{v} such that $\dot{a}^G = a$ and $\dot{v}^G = v$. Any member of \dot{a}^G is of the form \dot{x}^G where $\langle p, \dot{x} \rangle \in \dot{a}$ and $p \in G$. Then

$$b = \{\dot{x}^G \mid \exists p \in G. \langle p, \dot{x} \rangle \in \dot{a} \wedge M[G] \models \varphi(\dot{x}^G, \dot{v}^G)\}$$

We define

$$\dot{b} = \{\langle p, \dot{x} \rangle \mid \langle p, \dot{x} \rangle \in \dot{a} \wedge (p \Vdash \varphi(\dot{x}, \dot{v}))^M\} \in M^{\mathbb{P}}$$

Thus, $\dot{b}^G \in M[G]$, so it suffices to show $\dot{b}^G = b$. We have $x \in \dot{b}^G$ if and only if there is some \mathbb{P} -name \dot{x} in M and $p \in G$ such that $\dot{x}^G = x$, $\langle p, \dot{x} \rangle \in \dot{a}$, and $(p \Vdash \varphi(\dot{x}, \dot{v}))^M$. By the forcing theorem, this is equivalent to the statement $x \in \dot{a}^G$ and $M[G] \models \varphi(x, v)$, which is precisely the statement $x \in b$. \square

The arguments for collection and power set will follow the same pattern.

Lemma. Suppose that M is a countable transitive model of ZF, $\mathbb{P} \in M$ is a forcing poset, and $G \subseteq \mathbb{P}$ is a generic filter. Then $M[G]$ models collection.

Proof. Let $\varphi(x, y, v)$ be a formula with free variables x, y, v . Fix $a, v \in M[G]$ with names \dot{a}, \dot{v} . Suppose $M[G] \models \forall x \in a. \exists y. \varphi(x, y, v)$. We claim that there is $b \in M[G]$ such that $M[G] \models \forall x \in a. \exists y \in b. \varphi(x, y, v)$. Let

$$C = \{\langle p, \dot{x} \rangle \mid p \in \mathbb{P} \wedge \dot{x} \in \text{ran } \dot{a} \wedge \exists \dot{y} \in M^{\mathbb{P}}. (p \Vdash \varphi(\dot{x}, \dot{y}, \dot{v}))^M\}$$

Then for all $\langle p, \dot{x} \rangle \in C$, there is $\dot{y} \in M^{\mathbb{P}}$ such that $(p \Vdash \varphi(\dot{x}, \dot{y}, \dot{v}))^M$. Note that the collection of such \dot{y} might not form a set, for example with the formula $\varphi(x, y) \equiv x \in y$. However, using collection in M , we may form a set $B \in M$ such that $B \subseteq M^{\mathbb{P}}$ and

$$\forall \langle p, \dot{x} \rangle \in C. \exists \dot{y} \in B. (p \Vdash \varphi(\dot{x}, \dot{y}, \dot{v}))^M$$

Finally, set

$$\dot{b} = \{\langle 1, \dot{y} \rangle \mid \dot{y} \in B\} \in M^{\mathbb{P}}$$

We show that $b = \dot{b}^G$ satisfies the required property. Fix some $x \in a$, then by definition there is $\langle q, \dot{x} \rangle \in \dot{a}$ such that $q \in G$ and $\dot{x}^G = x$. By assumption, $M[G] \models \exists y \in b. \varphi(x, y, v)$. So fix \dot{z}^G such that $M[G] \models \varphi(x, \dot{z}^G, v)$. By the forcing theorem, there is $p \in G$ such that $(p \Vdash \varphi(\dot{x}, \dot{z}, \dot{v}))^M$. Hence $\langle p, \dot{x} \rangle \in C$. So we can fix $\dot{y} \in B$ such that $(p \Vdash \varphi(\dot{x}, \dot{y}, \dot{v}))^M$. Therefore, $\langle 1, \dot{y} \rangle \in \dot{b}$. Since $1 \in G$, $\dot{y}^G \in \dot{b}^G$. By the forcing theorem again,

$$M[G] \models \dot{y}^G \in \dot{b}^G \wedge \varphi(\dot{x}^G, \dot{y}^G, v)$$

Hence, collection holds. \square

Note that since power set has not been used in any of the previous proofs, if $M \models \text{ZF}^-$, then $M[G] \models \text{ZF}^-$.

Lemma. Suppose that M is a countable transitive model of ZF , $\mathbb{P} \in M$ is a forcing poset, and $G \subseteq \mathbb{P}$ is a generic filter. Then $M[G]$ models the axiom of power set.

Proof. By separation, it suffices to show that if $a \in M[G]$, then

$$\mathcal{P}(a) \cap M[G] = \{x \in M[G] \mid x \subseteq a\} \subseteq b$$

for some set $b \in M[G]$. Fix $a \in M[G]$ with name $\dot{a} \in M^{\mathbb{P}}$, and define

$$S = \{\dot{x} \in M^{\mathbb{P}} \mid \text{ran } \dot{x} \subseteq \text{ran } \dot{a}\} = \mathcal{P}(\mathbb{P} \times \text{ran } \dot{a})^M$$

and let

$$\dot{b} = \{\langle 1, \dot{x} \rangle \mid \dot{x} \in S\} \in M^{\mathbb{P}}$$

Let $c \in \mathcal{P}(a) \cap M[G]$; we must show that $c \in \dot{b}^G$. Let $\dot{c} \in M^{\mathbb{P}}$ be a name for c , and let

$$\dot{x} = \{\langle p, \dot{z} \rangle \mid \dot{z} \in \text{ran } \dot{a} \wedge (p \Vdash \dot{z} \in \dot{c})^M\} \in S$$

We claim $\dot{x}^G = \dot{c}^G = c$. First, we show $\dot{x}^G \subseteq c$. Fix $\dot{z}^G \in \dot{x}^G$. By definition, we can fix $p \in G$ such that $\langle p, \dot{z} \rangle \in \dot{x}$. From this, it follows that $\dot{z} \in \text{ran } \dot{a}$ and $p \Vdash \dot{z} \in \dot{c}$. Since $p \in G$, by the forcing theorem, $M[G] \models \dot{z}^G \in \dot{c}^G$, as required.

Conversely, since $M[G] \models c \subseteq \dot{a}^G$, every element of c is of the form \dot{z}^G for $\langle q, \dot{z} \rangle \in \dot{a}$ with $q \in G$. Also, if $M[G] \models \dot{z}^G \in c$, then by the forcing theorem, there is p such that $p \Vdash \dot{z} \in \dot{c}$. Then $\langle p, \dot{z} \rangle \in \dot{x}$, so $\dot{z}^G \in \dot{x}^G$. \square

Lemma. Suppose that M is a countable transitive model of ZFC^- , $\mathbb{P} \in M$ is a forcing poset, and $G \subseteq \mathbb{P}$ is a generic filter. Then $M[G]$ models the well-ordering principle, and hence models ZFC^- .

Proof. It suffices to show that any $a \in M[G]$ can be well-ordered in $M[G]$. Fix a name \dot{a} for a . Using the well-ordering principle in M , we can enumerate the elements of $\text{ran } \dot{a}$ as

$$\{\dot{x}_\alpha \mid \alpha < \delta\}$$

Let

$$\dot{f} = \{\langle 1, \text{op}(\check{\alpha}, \dot{x}_\alpha) \rangle \mid \alpha < \delta\} \in M^{\mathbb{P}}$$

So in $M[G]$,

$$\dot{f}^G = \{\langle \alpha, \dot{x}_\alpha^G \rangle \mid \alpha < \delta\}$$

Hence \dot{f}^G is a function with domain δ , and $a \subseteq \text{ran } \dot{f}^G$. We can now define a well-order $<$ on a by defining that $x < y$ if and only if

$$\min\{\alpha < \delta \mid \dot{f}^G(\alpha) = x\} < \min\{\alpha < \delta \mid \dot{f}^G(\alpha) = y\}$$

\square

VII. Forcing and the Continuum Hypothesis

Remark. (i) \dot{f}^G may not be injective, since we could have $\dot{x}_\alpha^G = \dot{x}_\beta^G$ for $\alpha \neq \beta$.

- (ii) $\text{ran } \dot{f}^G$ may not equal a . Elements of \dot{a} are conditions $\langle p, \dot{x}_\alpha \rangle$, and if $p \notin G$, we may not have $\dot{x}_\alpha^G \in a$.
- (iii) For power set, it sufficed to find a set of names which contained enough names to represent all possible subsets of a . However, there are a proper class of names for the empty set, so we could not produce a set of all such names.
- (iv) The statement $M[G] \models \varphi$ should be considered a ternary relation between M , G , and φ . It is possible that G and H are both generic, but $M[G] \models \varphi$ and $M[H] \models \neg\varphi$.
- (v) The relativisation $(p \Vdash \varphi)^M$ will be dropped when clear in subsequent sections.

Lemma. Let M be a countable transitive model of ZFC and let $\mathbb{P} \in M$ be a forcing poset. Let φ, ψ be $\mathcal{FL}_{\mathbb{P}}$ -formulas. Then, for any $p \in \mathbb{P}$ and $\dot{x} \in M^{\mathbb{P}}$,

- (i) if $\text{ZFC} \vdash \forall v. \varphi(v) \rightarrow \psi(v)$ then $(p \Vdash \varphi(\dot{x}))^M \rightarrow (p \Vdash \psi(\dot{x}))^M$; and
- (ii) if $\text{ZFC} \vdash \forall v. \varphi(v) \leftrightarrow \psi(v)$ then $(p \Vdash \varphi(\dot{x}))^M \leftrightarrow (p \Vdash \psi(\dot{x}))^M$.

Informally, forcing is closed under logical equivalence.

Proof. Clearly (ii) follows from (i). Suppose that $\text{ZFC} \vdash \forall v. \varphi(v) \rightarrow \psi(v)$ and $(p \Vdash \varphi(\dot{x}))^M$. Since M is countable, we can let G be a \mathbb{P} -generic filter over M such that $p \in G$. By the forcing theorem, $M[G] \models \varphi(\dot{x}^G)$. Since $M[G] \models \text{ZFC}$, we have $M[G] \models \psi(\dot{x}^G)$. Hence, by the forcing theorem in the reverse direction, as this is true for all generics containing p we have $(p \Vdash \psi(\dot{x}))^M$. \square

4. Forcing and independence results

4.1. Independence of the constructible universe

In this subsection, we show $\text{Con}(\text{ZFC} + V \neq L)$, and thus $V \neq L$ is independent of the axioms of ZFC.

Theorem. Let M be a countable transitive model of ZFC. Then there is a countable transitive model $N \supseteq M$ such that $N \models \text{ZFC} + V \neq L$.

Proof. Let M be a countable transitive model of ZFC, and let $\mathbb{P} \in M$ be any atomless forcing poset (that is, it has no minimal elements), for example $\text{Fn}(\omega, 2)$. Since M is countable, we can let G be a \mathbb{P} -generic filter over M . As \mathbb{P} is atomless, $G \notin M$. Hence $M \subsetneq M[G] \models \text{ZFC}$.

We show that $M[G] \models V \neq L$. We have

$$L_{\text{Ord} \cap M} = L^M \subseteq M \subsetneq M[G]$$

By the generic model theorem, $\text{Ord} \cap M = \text{Ord} \cap M[G]$, so $M[G] \neq L_{\text{Ord} \cap M[G]} = L^{M[G]}$. In particular, we have $(V \neq L)^{M[G]}$. \square

We will now discuss how to remove the assumption that we have a countable transitive model of ZFC.

Theorem. If $\text{Con}(\text{ZFC})$, then $\text{Con}(\text{ZFC} + V \neq L)$. Hence, $\text{ZFC} \not\vdash V = L$.

Proof. Suppose that $\text{ZFC} + V \neq L$ gives rise to a contradiction. Then, from a finite set of axioms $\Gamma \subseteq \text{ZFC} + V \neq L$, we can find ψ such that $\Gamma \vdash \psi \wedge \neg\psi$. By following the previous proofs, there is a finite set of axioms $\Lambda \subseteq \text{ZFC}$ such that ZFC proves that if there is a countable transitive model of Λ , then there is a countable transitive model of Γ . This set Λ should be sufficient to do the following:

- (i) to prove basic properties of forcing and constructibility;
- (ii) to prove the necessary facts about absoluteness, such as absoluteness of finiteness, partial orders and so on;
- (iii) to prove facts about forcing, including the forcing theorem; and
- (iv) if M is a countable transitive model of Λ with $\mathbb{P} \in M$ and G is \mathbb{P} -generic over M , then Λ proves that $M[G] \models \Gamma$.

As Λ is finite and a subset of the axioms of ZFC, then by the reflection theorem there is a countable transitive model of Λ . Hence, there is a countable transitive model N of Γ . But $\Gamma \vdash \psi \wedge \neg\psi$, so $N \models \psi \wedge \neg\psi$. Hence $(\psi \wedge \neg\psi)^N$, so in ZFC we can prove $\psi^N \wedge \neg\psi^N$, so ZFC is inconsistent. \square

Remark. Gunther, Pagano, Sánchez Terraf, and Steinberg recently completed a formalisation of the countable transitive model approach to forcing in the interactive theorem prover

VII. Forcing and the Continuum Hypothesis

Isabelle. To obtain $\text{Con}(\text{ZFC}) \rightarrow \text{Con}(\text{ZFC} + \neg\text{CH})$, they used ZC together with 21 instances of replacement, which are explicitly enumerated in the paper.

4.2. Cohen forcing

Fix a countable transitive model M of ZFC. Recall that for $I, J \in M$,

- (i) $\text{Fn}(I, J) = \{p \mid p \text{ is a finite partial function } I \rightarrow J\}$, together with \supseteq and \emptyset , has the structure of a forcing poset.
- (ii) $\text{Fn}(I, J)$ is always a set in M .
- (iii) $\text{Fn}(I, J)$ has the countable chain condition if and only if I is empty or J is countable.
- (iv) The sets $D_i = \{q \in \text{Fn}(I, J) \mid i \in \text{dom } q\}$ and $R_j = \{q \in \text{Fn}(I, J) \mid i \in \text{ran } q\}$ are dense for all $i \in I$ and $j \in J$.

Now, suppose that $G \subseteq \text{Fn}(I, J)$ is generic over M . Since G is a filter, if $p, q \in G$ then $p \cap q \in G$. Hence, if $p, q \in G$, then p, q agree on the intersection of their domains. Let $f_G = \bigcup G$. Then f_G is a function with domain contained in I and range contained in J . Note that this function has name

$$\dot{f} = \{\langle p, \text{op}(i, j) \rangle \mid p \in \mathbb{P}, \langle i, j \rangle \in p\}$$

Since D_i, R_j are dense, we obtain $G \cap D_i \neq \emptyset$, so we must have $i \in \text{dom } f_G$. Similarly, $j \in \text{ran } f_G$. We therefore obtain the following.

Proposition. Let $G \subseteq \text{Fn}(I, J)$ be a generic filter over M , and suppose I, J are nonempty. Then $M[G] \models f_G : I \rightarrow J$ is a surjection.

Proposition. Suppose that I, J are nonempty sets, at least one of which is infinite. Then

$$|\text{Fn}(I, J)| = \max(|I|, |J|)$$

In particular, $|\text{Fn}(\omega, 2)| = \aleph_0$.

Proof. Each condition $p \in \text{Fn}(I, J)$ is a finite function, so from this it follows that

$$\text{Fn}(I, J) \subseteq (I \times J)^{<\omega}$$

Hence

$$|\text{Fn}(I, J)| \subseteq |(I \times J)^{<\omega}| = |I \times J| = \max(|I|, |J|)$$

For the reverse direction, if we fix $i_0 \in I$ and $j_0 \in J$, then

$$\{\langle i_0, j \rangle \mid j \in J\} \cup \{\langle i, j_0 \rangle \mid i \in I\}$$

is a collection of $|I \cup J|$ -many distinct elements of $\text{Fn}(I, J)$. Thus

$$\max(|I|, |J|) = |I \cup J| \leq |\text{Fn}(I, J)|$$

as required. □

4. Forcing and independence results

We aim to provide a model in which CH fails. To do this, we will consider the forcing poset $\text{Fn}(\omega_2^M \times \omega, 2)$. We may consider $f_G : \omega_2^M \times \omega \rightarrow 2$, and let $g_\alpha : \omega \rightarrow 2$ be the function defined by $g_\alpha(n) = f_G(\alpha, n)$. This provides ω_2^M -many reals in $M[G]$. To show that $M[G] \models \text{ZFC} + \neg\text{CH}$, we must show that all of the g_α are distinct, and that

$$\omega_1^{M[G]} = \omega_1^M; \quad \omega_2^{M[G]} = \omega_2^M$$

It will turn out that the countable chain condition guarantees that all cardinals in M remain cardinals in $M[G]$.

Example. Let κ be an uncountable cardinal in M , and consider $\text{Fn}(\omega, \kappa)$, which does not satisfy the countable chain condition. Then in $M[G]$, the function $f_G : \omega \rightarrow \kappa$ is a surjection. Hence, κ has been collapsed into a countable ordinal in $M[G]$.

4.3. Preservation of cardinals

Definition. Let $\mathbb{P} \in M$ be a forcing poset. We say that \mathbb{P} *preserves cardinals* if and only if for every generic filter $G \subseteq \mathbb{P}$ over M and every $\kappa \in \text{Ord} \cap M$,

$$(\kappa \text{ is a cardinal})^M \leftrightarrow (\kappa \text{ is a cardinal})^{M[G]}$$

Also, \mathbb{P} *preserves cofinalities* if and only if for every generic filter $G \subseteq \mathbb{P}$ over M ,

$$\text{cf}^M(\gamma) = \text{cf}^{M[G]}(\gamma)$$

for all limit ordinals γ .

Recall that being a cardinal is Π_1 -definable so downwards absolute. In particular, cardinals of $M[G]$ are automatically cardinals of M . Also, note that finiteness and being ω are absolute.

Lemma. Let $\mathbb{P} \in M$ be a forcing poset. Then

- (i) \mathbb{P} preserves cofinalities if and only if for every generic filter G , for all limit ordinals β with $\omega < \beta < \text{Ord} \cap M$,

$$(\beta \text{ is regular})^M \rightarrow (\beta \text{ is regular})^{M[G]}$$

and

- (ii) if \mathbb{P} preserves cofinalities, then \mathbb{P} preserves cardinals.

The converse of (ii) is not true. Note that the definition of regularity did not require being a cardinal, but is a consequence.

Proof. Part (i). Suppose \mathbb{P} preserves cofinalities and G is \mathbb{P} -generic. Fix a limit ordinal β such that $\omega < \beta < \text{Ord} \cap M$. Then if β is regular in M , we have

$$\beta = \text{cf}^M(\beta) = \text{cf}^{M[G]}(\beta)$$

VII. Forcing and the Continuum Hypothesis

Hence β is regular in $M[G]$. Conversely, suppose γ is a limit ordinal such that $\omega < \gamma < \text{Ord} \cap M$. Let $\beta = \text{cf}^M(\gamma)$. Then β is a regular cardinal in M . Let $f \in M$ be a strictly increasing cofinal function $\beta \rightarrow \gamma$. If β is uncountable in M , then β is regular in $M[G]$ by assumption. Otherwise, $\beta = \omega$, and then $\beta = \omega^{M[G]}$ by absoluteness, and so again β is regular in $M[G]$. As $f \in M$, also $f \in M[G]$, so there is a strictly increasing cofinal map $\beta \rightarrow \gamma$ in $M[G]$, so

$$\text{cf}^{M[G]}(\gamma) = \text{cf}^{M[G]}(\beta) = \beta = \text{cf}^M(\gamma)$$

Part (ii). Suppose that \mathbb{P} preserves cofinalities. Let κ be a cardinal in M . One of three cases occur.

- (a) If $\kappa \leq \omega$, then $(\kappa \leq \omega)^{M[G]}$, so κ is a cardinal in $M[G]$;
- (b) If κ is regular in M , then κ is regular in $M[G]$ by (i), so it is a cardinal in $M[G]$.
- (c) Suppose κ is singular in M . In this case, one can show that κ is the supremum of a set S of regular cardinals in M . One way to show this is that if κ is the supremum of a set T of cardinals, we can set $S = \{\lambda^+ \mid \lambda \in T\}$. Since \mathbb{P} preserves regular cardinals, every element of S is regular in $M[G]$, and in particular they are cardinals. Hence κ is the supremum of a set of cardinals, and is therefore a cardinal.

□

Lemma (the approximation lemma). Let $A, B, \mathbb{P} \in M$, and suppose that $(\mathbb{P}$ has the countable chain condition) M . Let G be \mathbb{P} -generic over M . Then for any function $f \in M[G]$ with $f : A \rightarrow B$, there is a function $F \in M$ with $F : A \rightarrow \mathcal{P}^M(B)$ such that for all $a \in A$, we have $f(a) \in F(a)$ and $|F(a)| \leq \aleph_0^M$.

This proof requires that M is countable. Note that the relativisation of the countable chain condition to M ensures that the hypothesis is non-vacuous, as any forcing poset in M is externally countable.

Proof. Suppose that $M[G] \models f : A \rightarrow B$. Since $A, B \in M$, we have canonical names $\check{A}, \check{B} \in M^{\mathbb{P}}$. Let \dot{f} be a name for f . By the forcing theorem, there is a condition $p \in G$ such that

$$p \Vdash \dot{f} : \check{A} \rightarrow \check{B} \text{ is a function}$$

Define $F : A \rightarrow \mathcal{P}^M(B)$ by

$$F(a) = \{b \in B \mid \exists q \leq p. q \Vdash \dot{f}(\check{a}) = \check{b}\}$$

Note that $F(a) \in M$ by the definability of the forcing relation, so as $A \in M$, the set

$$F = \{\langle a, F(a) \rangle \mid a \in A\}$$

is a set in M . We now show that this definition has the desired properties. Observe that as F is a function in M , it is also a function in V . We show that $f(a) \in F(a)$. Suppose that

4. Forcing and independence results

$M[G] \models f(a) = b$ for $b \in B$. By the forcing theorem, there is $q \in G$ such that $q \Vdash \dot{f}(\check{a}) = \check{b}$. As G is a filter, there is $r \leq p, q$ with $r \in G$ witnessing $b \in F(a)$ as required.

We now show that $|F(a)| \leq \aleph_0$. Working in M , and in particular using the axiom of choice in M , for each $b \in F(a)$ there is a condition $q_b \leq p$ such that $q_b \Vdash \dot{f}(\check{a}) = \check{b}$. It suffices to show that $q_b \perp q_c$ for $b \neq c$, because then they form an antichain, so by the countable chain condition we may conclude $|F(a)| \leq \aleph_0$. Suppose not, so let $r \leq q_b, q_c$. Then

$$r \Vdash \dot{f} : \check{A} \rightarrow \check{B} \text{ is a function } \wedge \dot{f}(\check{a}) = \check{b} \wedge \dot{f}(\check{a}) = \check{c} \wedge \check{b} \neq \check{c}$$

Let H be a generic filter with $r \in H$; this exists by countability of M . Then $r \leq p$ and

$$M[H] \models f : A \rightarrow B \text{ is a function } \wedge f(a) = b \wedge f(a) = c \wedge b \neq c$$

But $M[H] \models \text{ZFC}$, giving a contradiction. \square

Theorem. If $\mathbb{P} \in M$ is a forcing poset and $(\mathbb{P} \text{ has the countable chain condition})^M$, then \mathbb{P} preserves cofinalities and hence cardinals.

Proof. Using the previous lemma, it suffices to show that \mathbb{P} preserves regular cardinals. That is, if $\omega < \beta < \text{Ord} \cap M$ and β is a limit, then if β is a regular cardinal in M , then β is a regular cardinal in $M[G]$. Suppose this is not the case, so there is such a β that is a regular cardinal in M but singular in $M[G]$. In $M[G]$, we can fix a cofinal map $f : \alpha \rightarrow \beta$ for some ordinal $\alpha < \beta$. As $\alpha, \beta \in M$, we can use the approximation lemma to find a function $F : \alpha \rightarrow \mathcal{P}^M(\beta)$ in M such that for all $\gamma \in \alpha$, we have $f(\gamma) \in F(\gamma)$ and $|F(\gamma)| \leq \aleph_0$. Working in M , let $X = \bigcup_{\gamma < \alpha} F(\gamma)$. This is a union of countable sets indexed by $\alpha < \beta$. So $X \subseteq \beta$ and is a subset of less than β -many countable sets. Hence $X \neq \beta$ as β is a regular cardinal in M . But f was cofinal, so $\beta = \bigcup_{\gamma < \alpha} f(\gamma) \subseteq X$, giving a contradiction. \square

4.4. The failure of the continuum hypothesis

Theorem. Let $\alpha < \text{Ord} \cap M$, and let $\kappa = (\aleph_\alpha)^M$. Let $\mathbb{P} = \text{Fn}(\kappa \times \omega, 2)$, and let G be \mathbb{P} -generic over M . Then $M[G]$ contains a κ -length sequence of distinct elements of 2^ω . Hence, $M[G] \models \text{ZFC} + (\aleph_\alpha = \kappa \leq 2^{\aleph_0})$.

Proof. Let $f = \bigcup G \in M[G]$. Then f is a function $\kappa \times \omega \rightarrow 2$. For $\beta < \kappa$, let $g_\beta : \omega \rightarrow 2$ be the function given by $g_\beta(n) = f(\beta, n)$. We claim that for $\alpha \neq \beta$, we have $g_\alpha \neq g_\beta$. Define a dense set $E_{\alpha, \beta} \in M$ as follows.

$$E_{\alpha, \beta} = \{q \in \mathbb{P} \mid \exists n. \langle \beta, n \rangle, \langle \alpha, n \rangle \in \text{dom } q \wedge q(\langle \beta, n \rangle) \neq q(\langle \alpha, n \rangle)\}$$

To show this is dense, fix $p \in \mathbb{P}$. Since p is finite, there is some m such that $\langle \beta, m \rangle, \langle \alpha, m \rangle \notin \text{dom } p$. Define $q \leq p$ with $q : \text{dom } p \cup \{\langle \beta, m \rangle, \langle \alpha, m \rangle\} \rightarrow 2$ by

$$q(z) = \begin{cases} p(z) & \text{if } z \in \text{dom } p \\ 1 & \text{if } z = \langle \beta, m \rangle \\ 0 & \text{if } z = \langle \alpha, m \rangle \end{cases}$$

VII. Forcing and the Continuum Hypothesis

Since G is \mathbb{P} -generic, we can fix $q' \in G \cap E_{\alpha,\beta}$. Then

$$g_\beta(m) = f(\beta, m) = q(\langle \beta, m \rangle) \neq q(\langle \alpha, m \rangle) = f(\alpha, m) = g_\alpha(m)$$

Hence $g_\alpha \neq g_\beta$. Finally, since \mathbb{P} has the countable chain condition in M , it preserves cardinals, so it preserves the \aleph hierarchy. \square

In particular, if $\alpha = 2$, the model $M[G]$ satisfies $\neg\text{CH}$.

Theorem. If ZFC is consistent, then so is $\text{ZFC} + \neg\text{CH}$.

The proof proceeds in the same way as the independence of $\text{V} = \text{L}$.

Definition. The g_β defined above are called *Cohen reals*. More precisely, we say that $c : \omega \rightarrow 2$ is a Cohen real over M if there exists H which is $\text{Fn}(\omega, 2)$ -generic over M and $c = \bigcup H$.

4.5. Possible sizes of the continuum

We have a way to add Cohen reals into a model M , but in general this process will add many more reals. In this subsection, we determine the possible sizes that the continuum can be. Recall that by König's theorem, $2^{\aleph_0} \neq \kappa$ for any κ with cofinality \aleph_0 . We will show that this is the only restriction on the possible sizes of the continuum. Note that under GCH, for any κ , $\text{cf}(\kappa) \neq \omega$ if and only if $\kappa^\omega = \kappa$.

Recall that in our proof that the axiom of power set holds in $M[G]$, given a name $\dot{a} \in M^{\mathbb{P}}$, the set $\mathcal{P}(\mathbb{P} \times \text{ran } \dot{a})$ is a name for its power set. We will show that there is a better name that gives a tighter bound on the sizes of power sets.

Theorem. Let M be a transitive model of ZFC, and assume $(\kappa = \aleph_\alpha \wedge \kappa^\omega = \kappa)^M$. Let $\mathbb{P} = \text{Fn}(\kappa \times \omega, 2)$, and let G be \mathbb{P} -generic over M . Then $M[G] \models 2^{\aleph_0} = \aleph_\alpha = \kappa$.

Proof. We have already shown that $M[G] \models \text{ZFC}$ and $M[G] \models \kappa = \aleph_\alpha \leq 2^{\aleph_0}$; it therefore remains to show that $2^{\aleph_0} \leq \aleph_\alpha$. Let \dot{x} be a name for a subset of ω . For $n \in \omega$, let

$$E_{\dot{x},n} = \{p \in \mathbb{P} \mid (p \Vdash \check{n} \in \dot{x}) \vee (p \Vdash \check{n} \notin \dot{x})\}$$

This is dense in \mathbb{P} . For each $n \in \omega$, choose a maximal antichain $A_{\dot{x},n} \subseteq E_{\dot{x},n}$. This is shown to be possible on an example sheet using the axiom of choice. Define

$$\dot{z}_{\dot{x}} = \bigcup_{n \in \omega} \{\langle p, \check{n} \rangle \mid p \in A_{\dot{x},n} \wedge p \Vdash \check{n} \in \dot{x}\}$$

Such names are called *nice*. We will show that $\dot{z}_{\dot{x}}$ and \dot{x} are both names for the same subset of ω , and since we can produce a bound on the amount of nice names, we can bound the size of 2^{\aleph_0} .

We claim that $1 \Vdash \dot{x} = \dot{z}_{\dot{x}}$. To do this, it suffices to prove that for all $n \in \omega$,

$$D_{\dot{x},n} = \{q \in E_{\dot{x},n} \mid (q \Vdash \check{n} \in \dot{x}) \leftrightarrow (q \Vdash \check{n} \in \dot{z}_{\dot{x}})\}$$

4. Forcing and independence results

is dense. Fix $n \in \omega$ and $p \in \mathbb{P}$. Since $E_{\dot{x},n}$ is dense, we can fix $p_0 \leq p$ such that $p_0 \in E_{\dot{x},n}$. As $A_{\dot{x},n}$ is a maximal antichain, there is $q_0 \in A_{\dot{x},n}$ such that $p_0 \parallel q_0$. Fix $r \leq p_0, q_0$. We will prove that $r \in D_{\dot{x},n}$. If $r \Vdash \check{n} \in \dot{x}$, then $q_0 \Vdash \check{n} \in \dot{x}$ as $q_0 \in E_{\dot{x},n}$. Hence, $\langle q_0, \check{n} \rangle \in \dot{z}_{\dot{x}}$ by definition, so $r \Vdash \check{n} \in \dot{z}_{\dot{x}}$. For the converse, suppose $r \Vdash \check{n} \in \dot{z}_{\dot{x}}$. By definition,

$$\{s \leq r \mid \exists \langle q_1, \check{m} \rangle \in \dot{z}_{\dot{x}}. s \leq q_1 \wedge (s \Vdash \check{m} = \check{n})\}$$

is dense below r . This can only happen if there is some q_1 with $\langle q_1, \check{n} \rangle \in \dot{z}_{\dot{x}}$ such that $r \parallel q_1$. Therefore, by definition, $q_1 \in A_{\dot{x},n}$. Since $A_{\dot{x},n}$ is an antichain containing q_0 and q_1 which are both compatible with r , we must have $q_0 = q_1$. Hence, $\langle q_0, \check{n} \rangle \in \dot{x}$. Thus $q_0 \Vdash \check{n} \in \dot{x}$ by definition, so since $r \leq q_0$, we have $r \Vdash \check{n} \in \dot{x}$. Therefore $D_{\dot{x},n}$ is dense as required.

The total number of subsets of ω is therefore bounded by the number of nice names. First, note that $|\mathbb{P}| = \kappa$. Furthermore, since \mathbb{P} has the countable chain condition, each $A_{\dot{x},n}$ is countable. Therefore, the amount of nice names is bounded by $(\kappa^\omega)^\omega \times (2^\omega)^\omega = \kappa$. As every subset of ω has a nice name, $M[G] \models 2^{\aleph_0} \leq \kappa$. \square

Corollary. $\text{Con}(\text{ZFC})$ implies $\text{Con}(\text{ZFC} + (2^{\aleph_0} = \aleph_2))$, and (for example) $\text{Con}(\text{ZFC} + (2^{\aleph_0} = \aleph_{\omega_1}))$.

Corollary. The following are equiconsistent.

- (i) $\text{ZFC} +$ there exists a weakly inaccessible cardinal;
- (ii) $\text{ZFC} + \text{GCH} +$ there exists a strongly inaccessible cardinal;
- (iii) $\text{ZFC} + 2^{\aleph_0}$ is weakly inaccessible;
- (iv) $\text{ZFC} +$ there exists a cardinal that is weakly inaccessible but not strongly inaccessible.

Proof. To show (i) implies (ii) we move to L. To show (iii) implies (iv), we note that 2^{\aleph_0} is not strongly inaccessible. It is trivial that (iv) implies (i). It therefore suffices to show that the continuum can be weakly inaccessible given (ii), which follows by considering the forcing $\mathbb{P} = \text{Fn}(\kappa \times \omega, 2)$. \square

Remark. When building models of $\text{ZFC} + (2^{\aleph_0} = \kappa)$, we often assume GCH for convenience. This can normally be done without loss of generality because we are usually only concerned with consistency results.

Example. Consider $\mathbb{P} = \text{Fn}(\aleph_\omega^M \times \omega, 2)$. Let G be a \mathbb{P} -generic filter. Then in $M[G]$, we must have $2^{\aleph_0} \geq \aleph_\omega$. By König's theorem, this inequality must be strict. For convenience, assume GCH holds. Under this assumption, if $\text{cf}(\kappa) = \omega$, then $\kappa^\omega = \kappa^+$, so there must be at most κ^+ -many nice names. Hence $M[G] \models \aleph_\omega < 2^{\aleph_0} \leq \aleph_\omega^+$ which gives $M[G] \models 2^{\aleph_0} = \aleph_{\omega+1}$.

Remark. (i) Note that it is possible that $2^{\aleph_0} < \aleph_\omega$ but $\aleph_\omega^{\aleph_0} = \aleph_{\omega+1}^{\aleph_0} = \aleph_{\omega+2}$ without GCH. This can be proven using large cardinals.

- (ii) If $M \vdash 2^{\aleph_0} = \aleph_\alpha > \aleph_\beta$ and $\mathbb{P} = \text{Fn}(\aleph_\beta^M \times \omega, 2)$, then $M[G] \models 2^{\aleph_0} = \aleph_\alpha$.

VII. Forcing and the Continuum Hypothesis

(iii) The following are equiconsistent.

- (a) ZFC + there exists a measurable cardinal + CH;
- (b) ZFC + there exists a measurable cardinal + \neg CH.

The same holds for other large cardinal axioms such as huge cardinals and I_0 to I_3 . We may also replace CH with GCH and the same holds.

(iv) The *proper forcing axiom*, which is a combinatorial axiom about forcing posets, implies that $2^{\aleph_0} = \aleph_2$ under ZFC.

4.6. Larger chain conditions

We now discuss generalised Cohen forcing. Suppose that we want a model of ZFC + CH + $(2^{\aleph_1} = \aleph_3)$. Naively, we might consider the forcing poset $\text{Fn}(\omega_3 \times \omega_1, 2)$, but we can show that CH fails in this model.

Proposition. Let M be a countable transitive model of ZFC + GCH, and let $(\kappa = \aleph_\alpha \wedge \kappa^\omega = \kappa)^M$. Let $\mathbb{P} = \text{Fn}(\kappa \times \omega, 2)$. Then, for any cardinal λ in M such that $\aleph_0 \leq \lambda < \kappa$, then in $M[G]$ we have

$$2^\lambda = \begin{cases} \kappa & \text{if } \text{cf } \kappa > \lambda \\ \kappa^+ & \text{if } \text{cf } \kappa \leq \lambda \end{cases}$$

There is a natural bijection between $\omega_3 \times \omega$ and $\omega_3 \times \omega_1$, and from this it follows that $2^{\aleph_0} = 2^{\aleph_1} = \aleph_3$.

Definition. Let I, J be sets and let κ be a regular cardinal. Define $\text{Fn}_\kappa(I, J)$ to be the partial functions $I \rightarrow J$ of size less than κ . Its maximal element is \emptyset under the order $q \leq p$ if and only if $p \subseteq q$.

Remark. (i) $\text{Fn}_\omega(I, J) = \text{Fn}(I, J)$.

(ii) The reason that $\text{Fn}(I, J)$ was absolute is that finite objects are absolute. In general, $\text{Fn}_\kappa(I, J)$ is not absolute. Moreover, if M is a countable transitive model, then $\text{Fn}_\kappa(I, J) \notin M$. We instead need to consider the relativisation $(\text{Fn}_\kappa(I, J))^M$.

(iii) If $\kappa > \omega$ and $I, J \neq \emptyset$, $\text{Fn}_\kappa(I, J)$ does not have the countable chain condition.

(iv) If G is $\text{Fn}_\kappa(I, J)$ -generic over M , then $f = \bigcup G$ is a function $I \rightarrow J$.

Let $\mathbb{P} = \text{Fn}_\kappa(\lambda \times \kappa, 2)$ where $\lambda \geq \kappa$ and κ is regular. Suppose also that $\lambda^\kappa = \lambda$. By a similar argument to the ω case, if $f = \bigcup G$ and $h_\alpha : \kappa \rightarrow 2$ is defined by $h_\alpha(\beta) = f(\alpha, \beta)$, then this gives a sequence of λ -many distinct functions $\kappa \rightarrow 2$. Similarly, by the nice names argument, there are precisely λ -many functions $\kappa \rightarrow 2$ because $\lambda^\kappa = \lambda$. We need to explicitly check that we have preserved all cardinals, using a generalisation of the countable chain condition. Once we have shown this, we will obtain $M[G] \models 2^\kappa = \lambda$.

4. Forcing and independence results

Definition. For a cardinal κ , we say that \mathbb{P} has the κ -chain condition if every antichain has cardinality less than κ .

The countable chain condition is equivalent to the \aleph_1 -chain condition. All of the proofs above immediately generalise to the κ -chain condition.

Definition. We say that \mathbb{P} preserves cofinalities above κ if and only if for all \mathbb{P} -generic filters G and limit ordinals $\gamma \in \text{Ord} \cap M$ with $\text{cf}^M(\gamma) \geq \kappa$, we have $\text{cf}^M(\gamma) = \text{cf}^{M[G]}(\gamma)$.

Lemma. Let $\mathbb{P} \in M$ be a forcing poset and $(\kappa \text{ is regular})^M$. Then

- (i) \mathbb{P} preserves cofinalities above κ if and only if for all \mathbb{P} -generic filters G and all limit ordinals β with $\kappa \leq \beta \in \text{Ord} \cap M$, we have $(\beta \text{ is regular})^M \rightarrow (\beta \text{ is regular})^{M[G]}$;
- (ii) If \mathbb{P} preserves cofinalities above κ , then \mathbb{P} preserves cardinals above κ .

Lemma. Let $A, B, \mathbb{P} \in M$, let $(\kappa \text{ is regular})^M$, let $(\mathbb{P} \text{ has the } \kappa\text{-chain condition})^M$, and let G be a \mathbb{P} -generic filter over M . Then for any $f : A \rightarrow B$ in $M[G]$, there is $F : A \rightarrow \mathcal{P}(B)$ in M such that for all $a \in A$, we have $f(a) \in F(a)$ and $(|F(a)| < \kappa)^M$.

Theorem. Let $\mathbb{P} \in M$ be a forcing poset such that $(\kappa \text{ is regular})^M$ and $(\mathbb{P} \text{ has the } \kappa\text{-chain condition})^M$. Then \mathbb{P} preserves cofinalities above κ , and hence cardinals above κ .

On the example sheet, we show that for any infinite cardinal κ , $\text{Fn}_\kappa(I, J)$ has the $(|J|^{<\kappa})^+$ -chain condition. In particular, $\text{Fn}_\kappa(\lambda \times \kappa, 2)$ has the $(2^{<\kappa})^+$ -chain condition. We will show a different version of this theorem.

Lemma. Let κ be a regular cardinal in M , and suppose that $(2^{<\kappa} = \kappa)^M$. Then, if $(1 \leq |J| \leq 2^{<\kappa})^M$, the forcing poset $\mathbb{P} = \text{Fn}_\kappa(I, J)^M$ has the κ^+ -chain condition.

Proof. If I is empty, the result is trivial, so we may assume I is nonempty. Let W be an antichain in \mathbb{P} . To show that $|W| \leq \kappa$, we will construct chains $(A_\alpha)_{\alpha < \kappa}$ in I and $(W_\alpha)_{\alpha \in \kappa}$ such that

- (i) for all $\alpha < \beta < \kappa$, we have $A_\alpha \subseteq A_\beta \subseteq I$ and $W_\alpha \subseteq W_\beta \subseteq W$;
- (ii) for limit ordinals γ , we have $A_\gamma = \bigcup_{\beta < \gamma} A_\beta$ and $W_\gamma = \bigcup_{\beta < \gamma} W_\beta$;
- (iii) $W = \bigcup_{\alpha < \kappa} W_\alpha$;
- (iv) for all $\alpha < \kappa$, $|A_\alpha| \leq \kappa$ and $|W_\alpha| \leq \kappa$.

The result then follows by regularity of κ^+ . Set $A_0 = W_0 = \emptyset$. It remains to define successor cases. Suppose we have constructed A_α, W_α . For each $p \in \mathbb{P}$ with $\text{dom } p \subseteq A_\alpha$, using the axiom of choice we choose $q_p \in W$ such that $p = q_p \upharpoonright_{A_\alpha}$, if it exists. Note that if $\text{dom } p \subseteq A_\beta$ for any $\beta < \alpha$, we will choose q_p to coincide with the q_p chosen at stage β . Then define

$$W_{\alpha+1} = W_\alpha \cup \{q_p \mid \text{dom } p \subseteq A_\alpha\}$$

and

$$A_{\alpha+1} = \bigcup \{\text{dom } q \mid q \in W_{\alpha+1}\}$$

VII. Forcing and the Continuum Hypothesis

Finally, set $A = \bigcup_{\alpha < \kappa} A_\alpha$.

We claim that $W = \bigcup_{\alpha < \kappa} W_\alpha$. By construction, we have $\bigcup_{\alpha < \kappa} W_\alpha \subseteq W$. For any $q \in W$, note that $\text{dom } q \cap A \neq \emptyset$, otherwise take $q_1 \in W_1$, and $\text{dom } q_1 \subseteq A$, so if $\text{dom } q_1 \cap \text{dom } q = \emptyset$, then $q_1 \parallel q$, contradicting $q_1, q \in W$. Since $\text{dom } q \cap A = \emptyset$ and $|\text{dom } q| < \kappa$, we must have $\text{dom } q \cap A = \text{dom } q \cap A_\alpha$ for some $\alpha < \kappa$. Define $p = q|_{A_\alpha}$. By definition, there is some $q' \in W_{\alpha+1}$ such that $q'|_{A_\alpha} = p$. Since $\text{dom } q' \subseteq A$, we have $q \parallel q'$. As W is an antichain, this is only possible if $q = q'$, so $q \in \bigcup_{\alpha < \kappa} W_\alpha$.

We now show that for all $\alpha < \kappa$, the sets W_α and A_α have size at most κ . We show this by induction on α . The result for limit cases follows from regularity. If $|W_{\alpha+1}| \leq \kappa$, then clearly $|A_{\alpha+1}| \leq \kappa$, so it remains to show $|W_{\alpha+1}| \leq \kappa$. Since every condition q that is added to W_α is chosen from some condition p with $\text{dom } p \subseteq A_\alpha$, then

$$|W_{\alpha+1}| \leq |W_\alpha| + |\{p \in \mathbb{P} \mid \text{dom } p \subseteq A_\alpha\}|$$

As $|A_\alpha| \leq \kappa$ and $|\text{dom } p| < \kappa$, then

$$|[A_\alpha]^{<\kappa}| \leq \kappa^{<\kappa} = 2^{<\kappa} = \kappa$$

Hence $|W_{\alpha+1}| \leq \kappa$ as required. \square

Hence, if $\mathbb{P} = \text{Fn}_\kappa(\lambda \times \kappa, 2)$, then $M[G] \models 2^\kappa = \lambda$ and all cardinals at least κ^+ are preserved.

4.7. Closure and distributivity

Definition. A poset \mathbb{P} is $< \kappa$ -closed if for every $\delta < \kappa$, every decreasing sequence of length δ in \mathbb{P} has a lower bound.

Definition. \mathbb{P} is $< \kappa$ -distributive if the intersection of less than κ -many open dense sets is an open dense set.

Lemma. If \mathbb{P} is $< \kappa$ -closed then \mathbb{P} is $< \kappa$ -distributive.

Lemma. If κ is regular in M , then $\text{Fn}_\kappa(I, J)^M$ is $< \kappa$ -closed.

Theorem. Let $A, B, \mathbb{P} \in M$, let κ be a cardinal in M with $(|A| < \kappa)^M$, and suppose \mathbb{P} is $< \kappa$ -distributive in M . Let G be \mathbb{P} -generic. Then if $f \in M[G]$ with $f : A \rightarrow B$, then $f \in M$.

Informally, forcing over a distributive poset cannot add any new small functions.

Proof. It suffices to prove the statement for $A = \delta$ where $\delta < \kappa$. Suppose that $M[G] \models f : \delta \rightarrow B$. By the forcing theorem, there is $p \in G$ such that $p \Vdash \dot{f} : \check{\delta} \rightarrow \check{B}$. For $\alpha < \delta$, let

$$D_\alpha = \{q \leq p \mid \exists x \in B. q \Vdash \dot{f}(\check{\alpha}) = \check{x}\}$$

These sets are clearly open, and they are dense below p because p forces that \dot{f} is a function. Since \mathbb{P} is $< \kappa$ -distributive, their intersection $D = \bigcap_{\alpha < \delta} D_\alpha$ is also (open and) dense below p .

4. Forcing and independence results

Let $q \in D \cap G$. Now, in M , for each $\alpha < \delta$, we can choose $x_\alpha \in B$ such that $q \Vdash \dot{f}(\check{\alpha}) = \check{x}_\alpha$, so we may define $g : \delta \rightarrow B$ by $\alpha \mapsto x_\alpha$. This g lies in M . But for any $\alpha < \delta$, we have $q \Vdash \dot{f}(\check{\alpha}) = \check{x}_\alpha = \check{g}(\check{\alpha})$, so $M[G] \models f = g$. \square

Theorem. Let $I, J, \kappa \in M$. Suppose that κ is a regular cardinal in M , and $(2^{<\kappa} = \kappa \wedge |J| \leq \kappa)^M$. Then $\text{Fn}_\kappa(I, J)^M$ preserves cofinalities and hence cardinals.

Proof. Recall that it suffices to show that for every limit ordinal $\beta \in \text{Ord} \cap M$, if β is regular in M then β is regular in $M[G]$. Let β be regular in M .

Suppose that $\beta > \kappa$. Since $|J| \leq \kappa = 2^{<\kappa}$ in M , the forcing poset $\text{Fn}_\kappa(I, J)^M$ has the κ^+ -chain condition. So it preserves all cofinalities and cardinals at least κ^+ , so in particular, β is regular in $M[G]$.

Now suppose that $\beta \leq \kappa$. Suppose that β is singular in $M[G]$. Fix $\delta < \beta$ and a cofinal map $f : \delta \rightarrow \beta$ in $M[G]$. Note that $\delta \in M$. Since \mathbb{P} is $< \kappa$ -closed, it is $< \kappa$ -distributive, so $f \in M$, contradicting the assumption that β is regular in M . \square

Theorem. Let κ, λ be cardinals in M such that $\aleph_0 \leq \kappa \leq \lambda$. Suppose that κ is regular, $2^{<\kappa} = \kappa$, and $\lambda^\kappa = \lambda$ in M . Let $\mathbb{P} = \text{Fn}_\kappa(\lambda \times \kappa, 2)$, and let G be \mathbb{P} -generic. Then \mathbb{P} preserves cardinals, and $M[G] \models 2^\kappa = \lambda$.

We can use this to fix multiple sizes of power sets at once.

Theorem. Let M be a countable transitive model of ZFC + GCH. Then there is a countable transitive model of ZFC satisfying any of the following statements.

- (i) $\text{CH} + 2^{\aleph_1} = \aleph_3$;
- (ii) $2^{\aleph_0} = 2^{\aleph_1} = \aleph_5$ and $2^{\aleph_2} = \aleph_{\omega+5}$;
- (iii) for a fixed $n \in \omega$, for all $m \leq n$, $2^{\aleph_m} = \aleph_{2m+3}$.

Proof. *Part (i).* Let $\mathbb{P} = \text{Fn}_{\aleph_1}(\omega_3 \times \omega_1, 2)^M$. If G is \mathbb{P} -generic, then $M[G] \models 2^{\aleph_1} = \aleph_3$. As \mathbb{P} is ω_1 -closed, it does not add any new functions $\omega \rightarrow 2$, so CH still holds in $M[G]$.

Part (ii). Let $\mathbb{P}_0 = \text{Fn}_{\aleph_2}(\omega_{\omega+5} \times \omega_2, 2)^M$. Let G_0 be \mathbb{P}_0 -generic. By closure, $2^{<\aleph_1} = \aleph_1$ in $M[G_0]$, and $\aleph_5^{\aleph_1} = \aleph_5$. Then let $\mathbb{P}_1 = \text{Fn}_{\aleph_0}(\omega_5 \times \omega, 2)^{M[G_0]}$. Let G_1 be \mathbb{P}_1 -generic. Then $M[G_1] \models 2^{\aleph_0} = 2^{\aleph_1} = \aleph_5$, where the latter equality is due to the fact that if M is a model of ZFC + GCH and G is $\text{Fn}(\kappa \times \omega, 2)$ -generic, then for any cardinal $\lambda \in M$ with $\aleph_0 \leq \lambda < \kappa$, the value of 2^λ in $M[G]$ is κ if $\text{cf}(\kappa) > \lambda$ and κ^+ if $\text{cf}(\kappa) \leq \lambda$. Also, $M[G_1] \models 2^{\aleph_2} = \aleph_{\omega+5}$ by preservation of cardinals.

Part (iii) is similar; we first make $2^{\aleph_m} = \aleph_{2m+3}$, then make $2^{\aleph_{m-1}} = \aleph_{2(m-1)+3}$, and continue downwards. \square

Remark. (i) It is necessary to start at the largest cardinal and work downwards; this ensures that the cardinal arithmetic in our forcing models remains correct.

VII. Forcing and the Continuum Hypothesis

- (ii) The iterative approach works for any finite number of cardinals. We will see later how we can force $2^{\aleph_n} = \aleph_{2n+3}$ for all $n \in \omega$.

We give an example to show that the order described in (i) is necessary.

Proposition. Let M be a countable transitive model of ZFC with $M \models 2^{\aleph_0} = \aleph_\alpha$. Let $\mathbb{P} = \text{Fn}_{\aleph_1}(\kappa \times \aleph_1, 2)$ for some $\kappa \geq 1$. Then if G is \mathbb{P} -generic, $M[G] \models \text{CH}$, and all cardinals δ of M with $\aleph_1 \leq \delta \leq \aleph_\alpha$ in M are no longer cardinals in $M[G]$. In particular, $\aleph_\alpha^M \neq \aleph_\alpha^{M[G]}$.

This is on the example sheets.

4.8. The mixing lemma

Recall that $p \Vdash \exists x. \varphi(x)$ if and only if

$$\{q \leq p \mid \exists \dot{x} \in V^{\mathbb{P}}. q \Vdash \varphi(\dot{x})\}$$

is dense below p . In most cases, the witness \dot{x} does not depend on G . For example, in $p \Vdash \exists x. (\dot{a} \in x \wedge \dot{b} \in x)$, we can find a name $\dot{x} = \text{op}(\dot{a}, \dot{b})$ without needing to know G . Informally, the mixing lemma says that this is always the case, as long as M has AC.

Theorem (the mixing lemma). (ZFC) Suppose that $(p \Vdash \exists x. \varphi(x))^M$. Then there is a name $\dot{x} \in M^{\mathbb{P}}$ such that $(p \Vdash \varphi(\dot{x}))^M$.

Proof. Since

$$\{q \leq p \mid \exists \dot{x} \in M^{\mathbb{P}}. q \Vdash \varphi(\dot{x})\}$$

is dense below p , it contains a maximal antichain D . Now, for each $q \in D$, choose some \dot{x}_q such that $q \Vdash \varphi(\dot{x}_q)$. Without loss of generality, we may assume that if $\langle r, \dot{y} \rangle \in \dot{x}_q$, then $r \leq q$. This is because

- (i) if $r \perp q$, then $q \Vdash \dot{x}_q = (\dot{x}_q \setminus \langle r, \dot{y} \rangle)$; and

- (ii) if $r \parallel q$, then define

$$\dot{x}'_q = (\dot{x}_q \setminus \langle r, \dot{y} \rangle) \cup \{\langle s, \dot{y} \rangle \mid s \leq r, q\}$$

$$\text{so } q \Vdash \dot{x}_q = \dot{x}'_q.$$

Now, if $q, q' \in D$ are such that $q \neq q'$, we must have $q \perp q'$ as D is an antichain. So $q' \Vdash \dot{x}_q = \emptyset$. We ‘mix’ the \dot{x}_q together to form

$$\dot{x} = \bigcup \{\dot{x}_q \mid q \in D\}$$

Then if $q \in D$, we have $q \Vdash \dot{x} = \dot{x}_q$. By the forcing theorem, $q \Vdash \varphi(\dot{x})$.

It remains to show that $p \Vdash \varphi(\dot{x})$. Suppose otherwise, so there is $r \leq p$ such that $r \Vdash \neg \varphi(\dot{x})$. As D is a maximal antichain of conditions below p , there is a condition $q \in D$ such that $q \parallel r$. Now if $s \leq q, r$, we have $s \Vdash \varphi(\dot{x})$ and $s \Vdash \neg \varphi(\dot{x})$, giving a contradiction. \square

4.9. Forcing successor cardinals

We would now like to find forcing posets that collapse $\kappa < \lambda$ such that $\lambda = \kappa^+$. Observe that this can only happen if λ is regular in M . This is because if $f : \alpha \rightarrow \lambda$ is cofinal with $\alpha < \lambda$ and $f \in M$, then $f \in M[G]$, so

$$\text{cf}^{M[G]}(\lambda) \leq \text{cf}^{M[G]}(\alpha) \leq |\alpha|^{M[G]} < \lambda$$

Assuming GCH in the ground model, this is the only restriction. We will prove this in the case where λ is a successor cardinal, and in the case where λ is strongly inaccessible; given GCH, these are the only options.

Theorem. Let κ be a regular cardinal in M , and let $\delta > \kappa$ be a cardinal in M . Let $\lambda = \delta^+$ in M . Let G be $\text{Fn}_\kappa(\kappa, \delta)$ -generic over M . Then in $M[G]$,

- (i) $|\delta| = \kappa$;
- (ii) every cardinal $\alpha \leq \kappa$ in M remains a cardinal in $M[G]$;
- (iii) if $\delta^{<\kappa} = \delta$ then every cardinal $\alpha > \delta$ in M remains a cardinal in $M[G]$.

In particular, if $\delta^{<\kappa} = \delta$, then $M[G] \models \lambda = \kappa^+$.

Observe that if δ is a cardinal in M and $\delta > |\mathbb{P}|$ in M , then δ remains a cardinal in $M[G]$. This is because \mathbb{P} has the $|\mathbb{P}|^+$ -chain condition.

Proof. Part (i). Note that $\bigcup G : \kappa \rightarrow \delta$ is a surjection, so $|\delta| = |\kappa|$ in $M[G]$. In particular, there are no cardinals between δ and λ .

Part (ii). Since κ is regular, $\text{Fn}_\kappa(\kappa, \delta)$ is $< \kappa$ -closed, so every cardinal $\alpha \leq \kappa$ is preserved.

Part (iii). Finally, if $\delta^{<\kappa} = \delta$, then $|\text{Fn}_\kappa(\kappa, \delta)| = \delta$, so $\text{Fn}_\kappa(\kappa, \delta)$ has the δ^+ -chain condition, so every cardinal $\alpha > \delta$ (in particular, λ) is preserved. \square

We can force inaccessible cardinals λ to become successor cardinals. To do this, we will use a forcing poset called the *Lévy collapse*.

Definition. Let $\lambda > \kappa$ be infinite ordinals. Then $\text{Col}(\kappa, < \lambda)$ consists of all functions p such that

- (i) p is a partial function from $\kappa \times \lambda \rightarrow \lambda$;
- (ii) $|\text{dom } p| < \kappa$;
- (iii) $p(\alpha, \beta) < \beta$ for each $(\alpha, \beta) \in \text{dom } p$.

We make this into a forcing poset by writing $q \leq p$ if and only if q extends p as a function.

Informally, for each $\beta < \lambda$, we add a surjection $\kappa \rightarrow \beta$.

Theorem (Lévy). Let κ be a regular cardinal in M , and suppose $\lambda > \kappa$ is strongly inaccessible in M . Let G be $\text{Col}(\kappa, < \lambda)$ -generic over M . Then in $M[G]$,

VII. Forcing and the Continuum Hypothesis

- (i) every ordinal β with $\kappa \leq \beta < \lambda$ has cardinality κ ; and
- (ii) every cardinal at most κ or at least λ remains a cardinal.

In particular, $M[G] \models \lambda = \kappa^+$.

Proof. If $\beta < \lambda$, we can define $G_\beta : \kappa \rightarrow \beta$ by $G_\beta(\alpha) = (\bigcup G)(\alpha, \beta)$. By density, this is a surjection, so if $\kappa \leq \beta < \lambda$, we have $M[G] \models |\beta| = |\kappa|$.

Note that $\text{Col}(\kappa, < \lambda)$ is $< \kappa$ -closed, so preserves cardinals at most κ . In particular, κ remains a cardinal.

Now, $|\text{Col}(\kappa, < \lambda)| = \lambda$. Therefore, $\text{Col}(\kappa, < \lambda)$ has the λ^+ -chain condition and therefore preserves cardinals at least λ^+ .

Finally, we show that λ is still a cardinal in $M[G]$, which follows from the λ -chain condition. Given $p \in \text{Col}(\kappa, < \lambda)$, define the *support* of p to be

$$\text{sp}(p) = \{\beta \mid \exists \alpha. \langle \alpha, \beta \rangle \in \text{dom } p\}$$

As $|p| < \kappa$, we must have $|\text{sp}(p)| < \kappa$. Let W be an antichain. We will construct chains $(A_\alpha)_{\alpha < \kappa}$ and $(W_\alpha)_{\alpha < \kappa}$ such that

- (i) for $\alpha < \beta < \kappa$, $A_\alpha \subseteq A_\beta \subseteq \lambda$ and $W_\alpha \subseteq W_\beta \subseteq W$;
- (ii) if $\gamma < \kappa$ is a limit, then $A_\gamma = \bigcup_{\alpha < \gamma} A_\alpha$ and $W_\gamma = \bigcup_{\alpha < \gamma} W_\alpha$;
- (iii) $W = \bigcup_{\alpha < \kappa} W_\alpha$;
- (iv) for all $\alpha < \kappa$, $|A_\alpha|, |W_\alpha| < \lambda$.

Assuming this can be done, since λ is regular, we have $|W| = |\bigcup_{\alpha < \kappa} W_\alpha| < \lambda$. To do this, first set $A_0 = W_0 = \emptyset$. To define successor cases, suppose A_α, W_α are defined. Suppose that $p \in \text{Col}(\kappa, < \lambda)$ has $\text{sp}(p) \subseteq A_\alpha$. Using the axiom of choice, choose $q_p \in W$ such that $p = q_p \restriction_{\kappa \times \text{sp}(p)}$ if this exists. Define

$$W_{\alpha+1} = \{q_p \mid \text{sp}(p) \subseteq A_\alpha\}; \quad A_{\alpha+1} = \bigcup \{\text{sp}(q) \mid q \in W_{\alpha+1}\}$$

One can show that $W = \bigcup_{\alpha < \kappa} W_\alpha$ in the same way that we proved this for $\text{Fn}_\kappa(I, J)$. We show by induction that for $\alpha < \kappa$, $|A_\alpha|, |W_\alpha| < \lambda$. Limit cases follow by regularity. If $|W_{\alpha+1}| < \lambda$, then $|A_{\alpha+1}| < \kappa \cdot \lambda = \lambda$. Suppose $|A_\alpha| < \lambda$. Then, since every q added in stage $\alpha + 1$ is chosen from some condition with support contained in A_α , we must have

$$|W_{\alpha+1}| \leq |A_\alpha|^{<\kappa}$$

Then as λ is a strong limit, $|A_\alpha|^{<\kappa} < \lambda$. □

Remark. (i) The requirement that κ was regular allowed us to deduce κ -closure.

4. Forcing and independence results

- (ii) Suppose λ is weakly inaccessible and $2^{\aleph_0} > \lambda$. Then $\text{Col}(\aleph_1, < \lambda)$ has an antichain of length 2^{\aleph_0} , so will not satisfy the λ -chain condition. Indeed, for $A \subseteq \omega$, we define $p_A : \{\omega\} \times [\omega, \omega + \omega) \rightarrow 2$ by

$$p_A(\alpha, \omega + n) = \begin{cases} 0 & \text{if } n \in A \\ 1 & \text{if } n \notin A \end{cases}$$

Then if $A \neq B$, the functions p_A, p_B are incompatible.

- (iii) One can show that λ is weakly compact if and only if it is inaccessible and satisfies the *tree property*. We claim that if G is $\text{Col}(\aleph_0, < \lambda)$ -generic, then in $M[G]$, \aleph_1 has the tree property. In general, we can use forcing to add combinatorial properties from large cardinals to \aleph_1 .
- (iv) This shows that λ being a limit cardinal is not absolute between M and N , even if λ being a cardinal is absolute for M, N .

Corollary. If $\text{ZFC} + \text{IC}$ is consistent, then so is $\text{ZFC} + (\aleph_1^V \text{ is inaccessible in } L)$.

Proof. Start with a model of $V = L$ where λ is inaccessible, and let G be $\text{Col}(\aleph_1, < \lambda)$ -generic. Then $M[G] \models \lambda = \aleph_1$, but also $M[G] \models (\lambda \text{ is inaccessible})^L$. \square

Remark. If $V \models \text{ZFC} + \kappa$ is measurable, then for example, \aleph_1^V is inaccessible in L .

4.10. Product forcing

In this subsection, we will show that is consistent that, for example, each $n \in \omega$ satisfies $2^{\aleph_n} = \aleph_{2n+3}$. We have already shown that for a fixed $N \in \omega$, it is consistent that all $n < N$ have $2^{\aleph_n} = \aleph_{2n+3}$. However, we cannot get this result using the iterated forcing process described in previous sections, and will instead use *product forcing*. This technique will allow us to exactly determine the restrictions on the continuum function $F : \text{Card} \rightarrow \text{Card}$ given by $F(\aleph_\alpha) = 2^{\aleph_\alpha}$.

Definition. Suppose $(\mathbb{P}, \leq_{\mathbb{P}})$ and $(\mathbb{Q}, \leq_{\mathbb{Q}})$ are posets. The *product order* \leq on $\mathbb{P} \times \mathbb{Q}$ is defined by

$$\langle p_1, q_1 \rangle \leq \langle p_0, q_0 \rangle \leftrightarrow p_1 \leq_{\mathbb{P}} p_0 \wedge q_1 \leq_{\mathbb{Q}} q_0$$

Given a $\mathbb{P} \times \mathbb{Q}$ -generic filter G over M , we can produce the *projections*

$$\begin{aligned} G_0 &= \{p \in \mathbb{P} \mid \exists q \in \mathbb{Q}. \langle p, q \rangle \in G\} \\ G_1 &= \{q \in \mathbb{Q} \mid \exists p \in \mathbb{P}. \langle p, q \rangle \in G\} \end{aligned}$$

Lemma. Let M be a transitive model of ZFC with $\mathbb{P}, \mathbb{Q} \in M$. Let $G \subseteq \mathbb{P}$ and $H \subseteq \mathbb{Q}$. Then the following are equivalent.

- (i) $G \times H$ is $\mathbb{P} \times \mathbb{Q}$ -generic over M ;

VII. Forcing and the Continuum Hypothesis

- (ii) G is \mathbb{P} -generic over M and H is \mathbb{Q} -generic over $M[G]$;
- (iii) H is \mathbb{Q} -generic over M and G is \mathbb{P} -generic over $M[H]$.

Moreover, when this is the case, $M[G \times H] = M[G][H] = M[H][G]$.

Proof. The first part is left as an exercise. For the last part, recall that the generic model theorem shows that if N is a transitive model of ZF containing M as a definable class and containing G as a set, then $M[G] \subseteq N$. Since $M \subseteq M[G][H]$, and $G \times H$ is an element of $M[G][H]$, we obtain $M[G \times H] \subseteq M[G][H]$. For the other direction, $G \in M[G \times H]$ and $M \subseteq M[G \times H]$ so $M[G] \subseteq M[G \times H]$, but also $H \in M[G \times H]$ so $M[G][H] \subseteq M[G \times H]$. \square

Recall that we started with a model of ZFC + GCH and forced with

$$G_0 \text{ is } \text{Fn}(\omega_3 \times \omega, 2)^M\text{-generic}; \quad G_1 \text{ is } \text{Fn}(\omega_5 \times \omega_1, 2)^{M[G_0]}\text{-generic}$$

and found that $M[G_0][G_1] \models \text{CH}$. But if instead we used

$$G_0 \text{ is } \mathbb{P}_0 = \text{Fn}(\omega_5 \times \omega_1, 2)^M\text{-generic}; \quad G_1 \text{ is } \mathbb{P}_1 = \text{Fn}(\omega_3 \times \omega, 2)^{M[G_0]}\text{-generic}$$

then we obtain $M[G_0][G_1] \models 2^{\aleph_0} = \aleph_3 + 2^{\aleph_1} = \aleph_5$. However, \mathbb{P}_0 is $< \omega_1$ -closed, so does not add new sequences of length ω . Thus $\mathbb{P}_1 = \text{Fn}(\omega_3 \times \omega, 2)^M$. We can therefore define the forcing poset $\mathbb{P}_0 \times \mathbb{P}_1$ -over M , and $G_0 \times G_1$ is $\mathbb{P}_0 \times \mathbb{P}_1$ -generic over M . To simultaneously force $2^{\aleph_n} = \aleph_{2n+3}$, we use the poset

$$\mathbb{P} = \prod_{n \in \omega} \text{Fn}_{\omega_n}(\omega_{2n+3} \times \omega_n, 2)$$

Easton's theorem shows that this works.

Theorem (Easton's theorem for sets). Let M be a countable transitive model of ZFC + GCH. Let S be a set of regular cardinals in M , and let $F : S \rightarrow \text{Card}^M$ be a function in M such that for all $\kappa \leq \lambda$ in S ,

- (i) $F(\kappa) > \kappa$ (Cantor's theorem);
- (ii) $F(\kappa) \leq F(\lambda)$ (monotonicity);
- (iii) $\text{cf}(F(\kappa)) > \kappa$ (König's theorem).

Then there is a generic extension $M[G]$ of M such that $M, M[G]$ have the same cardinals, and for all $\kappa \in S$, $M[G] \models 2^\kappa = F(\kappa)$.

The proof is non-examinable.

By essentially the same proof, this result can be generalised to proper classes of M , and in particular $S = \text{Reg}^M$. This needs a notion of *class forcing*, as \mathbb{P} is a proper class. The main obstacle with class forcing is that $M[G]$ need not be a model of ZFC. For example, consider $\text{Fn}(\text{Ord} \times \omega, 2)$, which makes 2^{\aleph_0} a proper class. Alternatively, consider $\text{Fn}(\omega, \text{Ord})$, which creates a surjection $\bigcup G : \omega \rightarrow \text{Ord}$. In fact, the forcing relation \Vdash may not even be

4. Forcing and independence results

definable. However, one can show that the particular forcing poset used in Easton's theorem also satisfies all of the required results for the proofs to work. In conclusion, we can say almost nothing about the values of the continuum function.