Math 615

Algebraic Topology I

Fall 2026

David Gepner

Lectures 1–12 and Appendices A–E

Math 615 · Lectures 1–12

Contents

Lecture 1. Additive and Abelian Categories **8**

1.1 Categories, Zero Objects, and Products 8
1.1.1 Categories, Functors, and Natural Transformations 8
1.1.2 Isomorphisms, Monomorphisms, and Epimorphisms 9
1.1.3 Zero Objects, Products, and Coproducts 10

1.2 Additive and Abelian Categories 10
1.2.1 Biproducts and Additive Categories 10
1.2.2 Kernels, Cokernels, and Abelian Categories 12
1.2.3 Normal Morphisms and Pullbacks 12
1.2.4 Coimage, Image, and the Comparison Isomorphism 13

1.3 Exact Sequences and Chain Complexes 13
1.3.1 Exact Sequences, Complexes, and Homology 14
1.3.2 Chain Homotopies 14

Lecture 2. Modules, Derived Categories, and the Tensor Product **15**

2.1 Exact Sequences and Mapping Cones 15
2.1.1 The Snake Lemma 15
2.1.2 The Homology Long Exact Sequence 17
2.1.3 The Cone Sequence 18
2.1.4 Recovering a Short Exact Sequence from a Cone 19

2.2 Homotopy and Derived Categories 19
2.2.1 Passing to Chain Homotopy Classes 19
2.2.2 Inverting Quasi-Isomorphisms 19
2.2.3 Complexes and Essentially Small Categories 20
2.2.4 Local Smallness of Derived Module Categories 21

2.3 Tensor Products and Their Computation 21
2.3.1 The Universal Bilinear Map 21
2.3.2 Tensor Products from Presentations 23

Lecture 3. Resolutions, Tor, and Ext **25**

3.1 Resolutions and Derived Functors 25
3.1.1 Projective and Injective Modules 25
3.1.2 Projective Resolutions and Their Comparison 26
3.1.3 Left and Right Derived Functors 27

3.2 Tor and the Failure of Exactness 29
3.2.1 Definition and Computation of Tor 29
3.2.2 Flatness and Long Exact Sequences 30

3.3 Cochain Complexes and Ext Groups 31
3.3.1 Cochain Complexes and Morphism Functors 31
3.3.2 Ext Groups, Variance, and Computation 33
3.3.3 Extensions as Degree-One Classes 34

Contents

3.4 Computations of Ext 35
3.4.1 An Explicit Noncommutative Ring 35
3.4.2 Kernels and Cokernels for Cyclic Groups 36
3.4.3 Computing Ext with Injective Resolutions 37
3.4.4 Comparing Tor and Ext Computations 38

Lecture 4. Simplicial Sets, Spaces, and Singular Chains 39

4.1 Simplicial Sets and Their Chains 39
4.1.1 Simplicial Sets and Geometric Constructions 39
4.1.2 Free Simplicial Modules and the Boundary 40
4.1.3 Simplicial Chain Complexes in Low Degrees 40

4.2 Presheaves, Simplices, and Simplicial Sets 41
4.2.1 Presheaves and the Yoneda Lemma 41
4.2.2 Topological Simplices and Affine Operators 42
4.2.3 Simplicial Sets 43

4.3 Geometric Realization and Singular Simplices 44
4.3.1 Geometric Realization 44
4.3.2 Singular Simplices 45
4.3.3 The Realization–Singular Adjunction 46

4.4 Singular Chains and Coefficients 47
4.4.1 The Free Simplicial Abelian Group Functor 47
4.4.2 Alternating Faces and Degenerate Simplices 48
4.4.3 The Singular Chain Functor 49
4.4.4 Coefficients and the Real Projective Plane 50

Lecture 5. Generalized Homology and the Eilenberg–Steenrod Axioms 51

5.1 Homotopy Invariance and Homotopy Pushouts 51
5.1.1 Graded Functors 51
5.1.2 Homotopies and Homotopy Invariance 51
5.1.3 Homotopy Pushouts 53

5.2 Homology Axioms and Coefficients 53
5.2.1 The Homology Axioms and Naturality 54
5.2.2 Coefficients and the Dimension Axiom 55

5.3 Relative Homology and Spheres 56
5.3.1 Reduced Homology and Homotopy Cofibers 56
5.3.2 The Long Exact Sequence of a Pair 57
5.3.3 Suspension and the Homology of Spheres 57

Lecture 6. Homotopy Invariance and First Computations 59

6.1 The Prism Operator and Homotopy Invariance 59
6.1.1 Chain Homotopies and the Prism Operator 59
6.1.2 Cutting the Prism Into Simplices 60
6.1.3 Prism Faces and Their Cancellation 61
6.1.4 The Prism Boundary Identity 62

Contents

6.2 Absolute and Relative Homology 63
6.2.1 Components, Points, and Disjoint Unions 63
6.2.2 Contractible Spaces and Deformation Retractions 64
6.2.3 Pairs, Quotients, and the Exact Sequence 65
6.2.4 Relative Homotopies and Reduced Homology 67

Lecture 7. Homotopy Pushouts and Mayer–Vietoris 69

7.1 Homotopy Pushouts and Their Chain Models 69
7.1.1 Double Mapping Cylinders and Their Universal Property 69
7.1.2 Cofibrations and Geometric Gluing 70
7.1.3 Subdivision 72
7.1.4 Algebraic and Geometric Double Mapping Cylinders 77

7.2 Mayer–Vietoris and Homology Calculations 77
7.2.1 Singular chains preserve homotopy pushouts 78
7.2.2 Mayer–Vietoris and Excision 79
7.2.3 The Gluing Axiom and Its Meaning 81
7.2.4 Cell Attachments and Mapping Cones 82

Lecture 8. Computations from the Eilenberg–Steenrod Axioms 83

8.1 Cell Complexes and Degree 83
8.1.1 Cell Attachments and Skeletal Quotients 83
8.1.2 Finite Product Cell Structures 84
8.1.3 Degree of Sphere Maps 84

8.2 Cellular Homology and Products 87
8.2.1 Relative Cells and Cellular Homology 87
8.2.2 Product Cells 90
8.2.3 Even Cells and Products of Spheres 92

8.3 Projective Spaces and Mapping Tori 92
8.3.1 Complex Projective Spaces 92
8.3.2 Real Projective Spaces 93
8.3.3 Mapping Tori 95

Lecture 9. Homotopy Groups and the Fundamental Group Action 96

9.1 Based Homotopy Groups 96
9.1.1 Based Maps and Homotopy Classes 96
9.1.2 Concatenation and the Group Axioms 97
9.1.3 The Eckmann–Hilton Argument 98

9.2 Functoriality and Basepoint Transport 99
9.2.1 Functoriality and Elementary Examples 99
9.2.2 Transport Along a Path 100
9.2.3 Composition, Inverses, and Naturality 101

9.3 The Fundamental Group Action and Homotopy Invariance 102
9.3.1 A Left Action and Free Homotopy Classes 102
9.3.2 Homotopies and Homotopy Equivalences 103

Contents

9.4 The Whitehead and Hurewicz Theorems 104
9.4.1 Weak Equivalences and the Whitehead Theorem 104
9.4.2 The Hurewicz Map and Theorem 105

9.5 Sphere Calculations and the Next Step 106
9.5.1 Connectivity of Spheres Without Cellular Approximation 106
9.5.2 The First Nonzero Group and Degree 107
9.5.3 References and Scope 108

Lecture 10. Fibrations and the Homotopy Exact Sequence 109

10.1 Lifting Homotopies 109
10.1.1 The Homotopy Lifting Property 109
10.1.2 Bundles and Examples 110

10.2 The Homotopy Exact Sequence 111
10.2.1 The Connecting Map 111
10.2.2 Exactness and Components 111

10.3 The Path–Loop Fibration 112
10.3.1 Paths Beginning at a Basepoint 112
10.3.2 Looping Lowers the Homotopy Degree 113

10.4 The Hopf Fibration 113
10.4.1 The Bundle over the Projective Line 113
10.4.2 Computing the Third Homotopy Group 114

Lecture 11. The Fundamental Group and Its Applications 115

11.1 The Fundamental Group of the Circle 115
11.1.1 Local Inverses and Lifting a Path 115
11.1.2 Lifting a Homotopy on a Square 116
11.1.3 Winding Number and Degree 116

11.2 Free Products and Relations 117
11.2.1 Reduced Words and the Universal Property 117
11.2.2 Normal Closures and Amalgamation 118

11.3 The Seifert–van Kampen Theorem 119
11.3.1 The Pushout Statement 119
11.3.2 Subdivided Loop Words and Their Choices 119
11.3.3 The Relations from a Homotopy Square 120

11.4 Graphs, Attaching Maps, and Surface Groups 121
11.4.1 Bouquets of Circles and Maximal Trees 121
11.4.2 Attaching Disks and Higher Cells 122
11.4.3 Tori, Projective Planes, and Closed Surfaces 123

11.5 Polynomial Roots and Antipodal Points 124
11.5.1 The Fundamental Theorem of Algebra 124
11.5.2 Odd Circle Maps Have Odd Degree 125
11.5.3 Borsuk–Ulam in Dimensions One and Two 125

Contents

Lecture 12. Covering Spaces and Subgroups of Free Groups 127

12.1 Lifting Paths, Homotopies, and Maps 127
12.1.1 Covering Maps and Unique Lifting 127
12.1.2 The Fundamental Group and the Lifting Criterion 128

12.2 Constructing and Classifying Connected Covers 129
12.2.1 Local Hypotheses and the Path-Class Construction 129
12.2.2 Based Classification, Conjugacy, and Universal Covers 130

12.3 Deck Transformations and Intermediate Coverings 132
12.3.1 The Deck Group and Normal Subgroups 132
12.3.2 The Galois Correspondence and Circle Covers 133

12.4 Graphs and Subgroups of Free Groups 134
12.4.1 Maximal Trees and Infinite Graphs 134
12.4.2 Covering Graphs, Free Subgroups, and Their Ranks 135

Appendix A. Mapping Complexes and Chain Homotopies 138

A.1 Mapping Complexes and Chain Homotopies 138
A.1.1 The Tensor–Mor Adjunction 138
A.1.2 The Interval and Chain Homotopies 139

A.2 Homotopy Pushouts of Chain Complexes 139
A.2.1 The Comparison Map and Its Signs 139
A.2.2 The Long Exact Homology Sequence 141

Appendix B. Attaching Maps and Euler Characteristic 142

B.1 Attaching Maps and Cellular Boundaries 142
B.1.1 A Disk Attached to a Graph 142
B.1.2 The Degree of an Attaching Map 143

B.2 Surfaces from Polygonal Attaching Maps 144
B.2.1 Orientable Polygonal Surfaces 144
B.2.2 Nonorientable Polygonal Surfaces 144

B.3 Euler Characteristic and Finite Chain Complexes 145
B.3.1 The Alternating-Dimension Identity 145
B.3.2 Attachments, Unions, and Products 146

B.4 Computations and Consequences 147
B.4.1 Graphs, Spheres, and Surfaces 147
B.4.2 Products and a Practical Computation Procedure 147

Appendix C. A Survey of the Kunneth Theorem 149

C.1 The Cross Product and the Product Formula 149
C.1.1 The Canonical Graded Map 149
C.1.2 The Formula over a Field 150

C.2 Torsion in the Künneth Formula 150
C.2.1 The Tor Term and Its Degree 150
C.2.2 Why the Exact Sequence Splits 151

Contents

C.3 Computations for Familiar Products 152
C.3.1 A Circle Factor and Products of Surfaces 152
C.3.2 The Product of Two Real Projective Planes 153
C.4 Maps of Products and Orientation Signs 154
C.4.1 Product Maps on Top Homology 154
C.4.2 Interchanging the Factors 154
Appendix D. Brouwer Fixed Points and Invariance of Dimension 155
D.1 Retractions and Fixed Points 155
D.1.1 The Obstruction to a Retraction 155
D.1.2 The Ray Construction and Brouwer's Theorem 156
D.2 Local Homology and Euclidean Dimension 156
D.2.1 Removing One Point 157
D.2.2 Invariance of Dimension 157
Appendix E. Invariance of Domain 159
E.1 Invariance of Domain and Boundary Degree 159
E.1.1 Degree and Extension over a Ball 159
E.1.2 An Injective Ball Map Contains a Neighborhood 160
E.2 The Degree of an Odd Map 160
E.2.1 Antipodal Hemisphere Chains 161
E.2.2 Comparing the Chains under an Odd Map 161
E.3 Consequences for Embeddings and Manifolds 162
E.3.1 Embeddings into Euclidean Space 162
E.3.2 Open Embeddings of Manifolds 163
References 164

Additive and Abelian Categories

Algebraic Topology I | Fall 2026 | August 31 and September 2

1.1 Categories, Zero Objects, and Products

1.1.1 Categories, Functors, and Natural Transformations

Definition 1.1.1. We assume categories are locally small. A *category* $\mathcal{C}$ consists of the following data:

1. A class $\mathrm{Ob}(\mathcal{C})$ of objects.
2. For every ordered pair of objects (A, B), a set $\mathrm{Mor}_{\mathcal{C}}(A, B)$ of morphisms with source A and target B. A member is written $f\colon A \to B$. Source and target are part of the data of a morphism.
3. For every three objects A, B, C, a composition function
$$\mathrm{Mor}_{\mathcal{C}}(B, C) \times \mathrm{Mor}_{\mathcal{C}}(A, B) \longrightarrow \mathrm{Mor}_{\mathcal{C}}(A, C), \qquad (g, f) \longmapsto g \circ f.$$
4. For every object A, a specified identity morphism $\mathrm{id}_A \in \mathrm{Mor}_{\mathcal{C}}(A, A)$.

The following diagrams are required to commute for every composable triple $A \xrightarrow{f} B \xrightarrow{g} C \xrightarrow{h} D$ (associativity) and every $f\colon A \to B$ (identity laws):

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ {\scriptstyle g\circ f}\downarrow & & \downarrow{\scriptstyle h\circ g} \\ C & \xrightarrow[h]{} & D \end{array} \qquad \begin{array}{ccc} A & \xrightarrow{\mathrm{id}_A} & A \\ & \searrow^{f} & \downarrow{\scriptstyle f} \\ & & B \end{array} \qquad \begin{array}{ccc} A & \xrightarrow{f} & B \\ & \searrow^{f} & \downarrow{\scriptstyle \mathrm{id}_B} \\ & & B. \end{array}$$

A diagram commutes when the composites along paths with the same endpoints agree. A category is *small* if its objects also form a set. Rings are unital and modules are right modules unless specified otherwise. We write Mod_R for right R-modules and ${}_R\mathrm{Mod}$ for left R-modules. For rings A, B, we write ${}_A\mathrm{Mod}_B$ for bimodules with commuting left A- and right B-actions. Examples of categories are sets with functions and Mod_R with linear maps.

Definition 1.1.2. A functor $F\colon \mathcal{C} \to \mathcal{B}$ assigns an object $F(A)$ to each object A and a function $\mathrm{Mor}_{\mathcal{C}}(A, B) \to \mathrm{Mor}_{\mathcal{B}}(F(A), F(B))$, written $f \mapsto F(f)$, to every pair. It preserves identities and composition: the following diagrams commute.

$$\begin{array}{ccc} F(A) & \xrightarrow{F(f)} & F(B) \\ & \searrow^{F(g\circ f)} & \downarrow{\scriptstyle F(g)} \\ & & F(C) \end{array} \qquad \begin{array}{ccc} & \overset{F(\mathrm{id}_A)}{\frown} & \\ F(A) & & F(A). \\ & \underset{\mathrm{id}_{F(A)}}{\smile} & \end{array}$$

It is *faithful* if each of these functions is injective and *full* if each is surjective. A *full subcategory* is obtained by choosing some objects and retaining every morphism between them.

Definition 1.1.3. The opposite category $\mathcal{C}^{\mathrm{op}}$ has the same objects and
$$\mathrm{Mor}_{\mathcal{C}^{\mathrm{op}}}(A, B) = \mathrm{Mor}_{\mathcal{C}}(B, A).$$

Composition is reversed. A contravariant functor on $\mathcal{C}$ means a functor defined on $\mathcal{C}^{\mathrm{op}}$. For example, $\mathrm{Mor}_{\mathcal{C}}(A,-)$ is covariant, whereas $\mathrm{Mor}_{\mathcal{C}}(-,B)$ is contravariant.

Definition 1.1.4. For functors $F, G\colon \mathcal{C} \to \mathcal{B}$, a natural transformation $\eta\colon F \Rightarrow G$ consists of morphisms $\eta_A\colon F(A) \to G(A)$ such that the following square commutes for every $f\colon A \to B$:

$$\begin{array}{ccc} F(A) & \xrightarrow{F(f)} & F(B) \\ {\scriptstyle \eta_A}\downarrow & & \downarrow{\scriptstyle \eta_B} \\ G(A) & \xrightarrow[G(f)]{} & G(B). \end{array}$$

A natural isomorphism is a natural transformation whose components are isomorphisms.

For the origins of the language of categories, functors, and natural transformations, see Eilenberg and Mac Lane [EM45].

1.1.2 Isomorphisms, Monomorphisms, and Epimorphisms

Definition 1.1.5. A morphism $f\colon A \to B$ is:

1. an *isomorphism* if there is a morphism $g\colon B \to A$ with $g \circ f = \mathrm{id}_A$ and $f \circ g = \mathrm{id}_B$;
2. a *monomorphism* (or *mono*) if, for every object X and every $g, h\colon X \to A$, the equality $f \circ g = f \circ h$ implies $g = h$;
3. an *epimorphism* (or *epi*) if, for every object Y and every $g, h\colon B \to Y$, the equality $g \circ f = h \circ f$ implies $g = h$.

An inverse, when it exists, is unique: if g and h are inverses, then $g = g \circ (f \circ h) = (g \circ f) \circ h = h$.

Every isomorphism is both mono and epi, by cancellation with its inverse. The converse depends on the category.

Example. In the category **Top** of topological spaces and continuous maps, let $X = \{0,1\}$ have the discrete topology and let $Y = \{0,1\}$ have the indiscrete topology. The identity function $f\colon X \to Y$ is continuous and bijective. Cancellation of functions shows that it is both mono and epi. Its inverse is not continuous, since a singleton is open in X but not in Y. Thus f is not an isomorphism.

In **Set**, monomorphisms are injections: to test injectivity, use maps from a singleton. Epimorphisms are surjections: if $f\colon A \to B$ misses $b \in B$, the zero function $B \to \{0,1\}$ and the characteristic function of $\{b\}$ agree after composition with f but are distinct. The converse implications follow by cancellation of functions. Hence a map in **Set** that is both mono and epi is bijective and therefore an isomorphism.

The same conclusion holds in Mod_R. For a monomorphism $f\colon M \to N$, compare the inclusion $f^{-1}(0) \to M$ with the zero map to see that f is injective. For an epimorphism f, compare the quotient map $N \to N/f(M)$ with the zero map to see that f is surjective. The set-theoretic inverse of a bijective module homomorphism is additive and R-linear: apply f to each required identity and use injectivity. Thus a morphism that is both mono and epi is an isomorphism. The same argument applies to left modules.

Proposition 1.1.6. A morphism $f\colon A \to B$ is an isomorphism if and only if, for every object X, postcomposition induces a bijection

$$f_*\colon \mathrm{Mor}_{\mathcal{C}}(X,A) \longrightarrow \mathrm{Mor}_{\mathcal{C}}(X,B).$$

Proof. An inverse to f induces an inverse to each f_*. Conversely, surjectivity for $X = B$ gives $g\colon B \to A$ with $f \circ g = \mathrm{id}_B$. Since $f \circ (g \circ f) = f = f \circ \mathrm{id}_A$, injectivity for $X = A$ gives $g \circ f = \mathrm{id}_A$. $\square$

1.1.3 Zero Objects, Products, and Coproducts

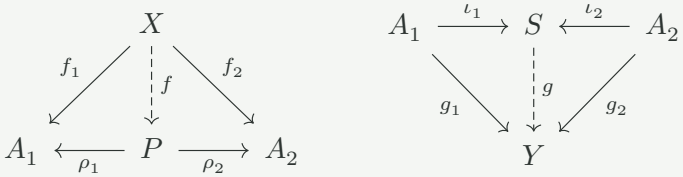
Definition 1.1.7. An object I is *initial* if $\mathrm{Mor}_{\mathcal{C}}(I, X)$ is a singleton for every X. An object T is *terminal* if $\mathrm{Mor}_{\mathcal{C}}(X, T)$ is a singleton for every X. A *zero object* is both initial and terminal and is usually denoted 0.

Definition 1.1.8. In a category with a zero object, the zero morphism $0_{A,B} \colon A \to B$ is the unique composite $A \to 0 \to B$. It does not depend on the choice of zero object. For $f \colon A \to B$, $g \colon X \to A$, and $h \colon B \to Y$,

$$f \circ 0_{X,A} = 0_{X,B}, \qquad 0_{A,B} \circ g = 0_{X,B}, \qquad h \circ 0_{A,B} = 0_{A,Y}.$$

We usually write 0 for the zero morphism when its source and target are clear.

Definition 1.1.9. A product of A_1, A_2 is an object P with projections $\rho_i \colon P \to A_i$ such that for every X and $f_i \colon X \to A_i$, there is a unique $f \colon X \to P$ making the left diagram commute. Dually, a coproduct is an object S with inclusions $\iota_i \colon A_i \to S$ such that for every Y and $g_i \colon A_i \to Y$, there is a unique $g \colon S \to Y$ making the right diagram commute:



Equivalently, the projections induce a bijection $\mathrm{Mor}_{\mathcal{C}}(X, P) \to \mathrm{Mor}_{\mathcal{C}}(X, A_1) \times \mathrm{Mor}_{\mathcal{C}}(X, A_2)$ for every X; the coproduct has the dual property.

Finite products include the empty product, a terminal object; finite coproducts include the empty coproduct, an initial object. Universal properties determine these objects up to unique isomorphism compatible with their structure maps.

1.2 Additive and Abelian Categories

1.2.1 Biproducts and Additive Categories

When the meaning is clear, we may write fg for the composite $f \circ g$.

Definition 1.2.1. In a category with zero morphisms, a biproduct $A_1 \oplus A_2$ is an object equipped with inclusions $\iota_i \colon A_i \to A_1 \oplus A_2$ and projections $\rho_i \colon A_1 \oplus A_2 \to A_i$ such that the projections exhibit a product, the inclusions exhibit a coproduct, and

$$\rho_i \iota_j = \begin{cases} \mathrm{id}_{A_i} & i = j, \\ 0 & i \neq j. \end{cases}$$

A category is *semiadditive* if it has a zero object and all binary biproducts, hence all finite biproducts.

Theorem 1.2.2. In a semiadditive category, each $\mathrm{Mor}_{\mathcal{C}}(A, B)$ has a canonical commutative monoid structure. Its identity is the zero morphism, and composition preserves addition and zero in each variable.

Proof. The diagonal $\Delta_A\colon A \to A \oplus A$ and codiagonal $\nabla_B\colon B \oplus B \to B$ are the unique maps making these diagrams commute for $i = 1, 2$:

$$\begin{array}{ccc} A & \xrightarrow{\Delta_A} & A \oplus A \\ & \searrow^{\mathrm{id}_A} & \downarrow{\scriptstyle \rho_i} \\ & & A \end{array} \qquad \begin{array}{ccc} B & \xrightarrow{\iota_i} & B \oplus B \\ & \searrow^{\mathrm{id}_B} & \downarrow{\scriptstyle \nabla_B} \\ & & B. \end{array}$$

Given $f, g\colon A \to B$, let $\Sigma_{f,g}\colon A \oplus A \to B$ be the unique map whose restrictions to the two summands are f, g. Define

$$f + g = \Sigma_{f,g}\Delta_A = \nabla_B(f \oplus g)\Delta_A,$$

where $f \oplus g$ is uniquely specified by its diagonal components f, g and its two zero off-diagonal components.

If $g = 0$, then $\Sigma_{f,g} = f\rho_1$, so $f + 0 = f$; similarly $0 + f = f$. Interchanging the summands fixes Δ_A and exchanges the restrictions of $\Sigma_{f,g}$, proving commutativity.

Let $\Delta_A^{(3)}\colon A \to A \oplus A \oplus A$ have all three components id_A, and let $\Sigma_{f,g,h}\colon A \oplus A \oplus A \to B$ have restrictions f, g, h. Comparing components and restrictions shows that both $(f + g) + h$ and $f + (g + h)$ equal $\Sigma_{f,g,h}\Delta_A^{(3)}$.

For $k\colon B \to D$, the restrictions of $k\Sigma_{f,g}$ are kf, kg, so $k(f + g) = kf + kg$. For $\alpha\colon X \to A$, the restrictions of $\Sigma_{f,g}(\alpha \oplus \alpha)$ are $f\alpha, g\alpha$, and $\Delta_A\alpha = (\alpha \oplus \alpha)\Delta_X$; hence $(f+g)\alpha = f\alpha + g\alpha$. Composing a zero morphism with any morphism is zero. □

Proposition 1.2.3. In a semiadditive category,

$$\iota_1\rho_1 + \iota_2\rho_2 = \mathrm{id}_{A_1 \oplus A_2}.$$

More generally a morphism $f\colon \bigoplus_j A_j \to \bigoplus_i B_i$ is uniquely determined by its entries $f_{ij} = \rho_i f \iota_j$, and

$$f = \sum_{i,j} \iota_i f_{ij} \rho_j, \qquad (gf)_{kj} = \sum_i g_{ki} f_{ij}.$$

Proof. Composing the first asserted identity with either projection gives the corresponding projection, so it follows from uniqueness for products. Apply this identity to the source and target of f, then use bilinearity and $\rho_i\iota_j = \delta_{ij}$ to obtain both formulas. □

Definition 1.2.4. A category is *preadditive* if every morphism set is equipped with an abelian group structure and composition is bilinear. It is *additive* if it is preadditive and has finite biproducts. Equivalently, it is semiadditive and all its canonical morphism monoids are groups.

The structures agree: the map $\Sigma_{f,g}$ used above is $f\rho_1 + g\rho_2$, so $\Sigma_{f,g}\Delta_A = f + g$ also for the given addition.

Definition 1.2.5. A functor $F\colon \mathcal{A} \to \mathcal{B}$ between preadditive categories is *additive* if every induced map on morphism groups is a group homomorphism. In particular, $F(0) = 0$ on morphisms and $F(f + g) = F(f) + F(g)$.

Between additive categories, an additive functor preserves finite biproducts, because it preserves the biproduct identities. Conversely, a functor preserving finite products or finite coproducts preserves zero objects and the biproduct structure, and hence the formula defining addition in Theorem 1.2.2. Thus it is additive. The functors $\mathrm{Mor}_{\mathcal{A}}(M, -)\colon \mathcal{A} \to \mathbf{Ab}$ and $\mathrm{Mor}_{\mathcal{A}}(-, N)\colon \mathcal{A}^{\mathrm{op}} \to \mathbf{Ab}$ are additive.

1.2.2 Kernels, Cokernels, and Abelian Categories

Definition 1.2.6. In a category with zero morphisms, a kernel of $f\colon A \to B$ is a morphism $\kappa\colon K \to A$ such that $f\kappa = 0$ and, for every $g\colon X \to A$ with $fg = 0$, there is a unique $h\colon X \to K$ with $\kappa h = g$.

A cokernel of f is a morphism $q\colon B \to Q$ such that $qf = 0$ and, for every $g\colon B \to Y$ with $gf = 0$, there is a unique $h\colon Q \to Y$ with $hq = g$.

We also write $\ker f = K$ and $\operatorname{coker} f = Q$ for the corresponding objects. For modules, $K = f^{-1}(0)$ and $Q = B/f(A)$.

Lemma 1.2.7. Kernels are monomorphisms and cokernels are epimorphisms. In an additive category, f is mono if and only if $fu = 0$ implies $u = 0$ for all u with target A. Dually, f is epi if and only if $uf = 0$ implies $u = 0$ for all u with source B. If a kernel of f exists, f is mono precisely when its kernel object is a zero object; dually for epis and cokernels.

Proof. If $kv = kw$ for a kernel k of f, both v, w are factorizations of the same morphism through k, so uniqueness gives $v = w$. The cokernel assertion is dual. In an additive category, $fu = fv$ is equivalent to $f(u - v) = 0$. This proves the cancellation criterion. If the kernel object is a zero object, every morphism whose composite with f is the zero morphism is itself the zero morphism. Conversely, if f is mono, $fk = 0$ forces $k = 0$; since k is mono, $k\,\mathrm{id}_K = k0_{K,K}$ forces $\mathrm{id}_K = 0_{K,K}$, which makes K a zero object. □

Definition 1.2.8. A category $\mathcal{A}$ is *abelian* if:

1. $\mathcal{A}$ is additive;
2. every morphism has a kernel and a cokernel;
3. Every monomorphism is a kernel and every epimorphism is a cokernel.

The definition is self-dual. Examples are **Ab**, Mod_R, and finite-dimensional vector spaces.

For a parallel categorical treatment of kernels, cokernels, and abelian categories, see Mac Lane [ML98, Chapter VIII, §§1–3, pp. 191–201].

1.2.3 Normal Morphisms and Pullbacks

Lemma 1.2.9. In an abelian category, every mono is a kernel of its own cokernel, and every epi is a cokernel of its own kernel.

Proof. Let $m\colon M \to B$ be mono, so m is a kernel of some $g\colon B \to D$. Let $c\colon B \to Q$ be a cokernel of m. Since $gm = 0$, write $g = tc$. If $cu = 0$, then $gu = tcu = 0$, so u factors uniquely through m. This proves that m is a kernel of c. Reverse arrows for the epi statement. □

Proposition 1.2.10. A morphism in an abelian category that is both mono and epi is an isomorphism.

Proof. If $f\colon A \to B$ is mono, it is a kernel of some $g\colon B \to D$. Since $gf = 0$ and f is epi, $g = 0$. The kernel property applied to id_B gives $s\colon B \to A$ with $fs = \mathrm{id}_B$. Monicity and $f(sf) = f$ then give $sf = \mathrm{id}_A$. □

Lemma 1.2.11. If $e\colon E \to B$ is epi and $t\colon X \to B$ is any morphism in an abelian category, their pullback exists and its projection $p\colon P \to X$ is epi.

Proof. Put $\alpha = e\rho_1 - t\rho_2 : E \oplus X \to B$, and let $\kappa : P \to E \oplus X$ be its kernel. Set $\rho_E = \rho_1\kappa$ and $p = \rho_2\kappa$. Then $e\rho_E = tp$. For morphisms $f : Y \to E$ and $g : Y \to X$, let $h : Y \to E \oplus X$ be the unique map with components f, g. The equality $ef = tg$ is equivalent to $\alpha h = 0$, so the kernel property gives a unique $Y \to P$ inducing f, g. This is the pullback universal property.

The map α is epi since $\alpha\iota_1 = e$. By Lemma 1.2.9, it is a cokernel of κ. If $\beta : X \to Y$ satisfies

$\beta p = 0$, then $\beta\rho_2\kappa = 0$, so there is $\gamma : B \to Y$ with $\beta\rho_2 = \gamma\alpha$. Restricting to E gives $\gamma e = 0$, whence $\gamma = 0$. Restricting to X now gives $\beta = 0$. The epi criterion in Lemma 1.2.7 proves that p is epi. □

1.2.4 Coimage, Image, and the Comparison Isomorphism

Definition 1.2.12. For $f\colon A \to B$, choose its kernel $k\colon K \to A$ and cokernel $c\colon B \to Q$. Define the coimage and image together with their structure maps by

$$m\colon \operatorname{im} f = \ker c \longrightarrow B, \qquad q\colon A \longrightarrow \operatorname{coim} f = \operatorname{coker} k.$$

There is a unique morphism $\overline{f}\colon \operatorname{coim} f \to \operatorname{im} f$ such that $f = m\overline{f}q$:

$$\begin{array}{ccc} A & \xrightarrow{\ f\ } & B \\ {\scriptstyle q}\downarrow & & \uparrow{\scriptstyle m} \\ \operatorname{coim} f & \xrightarrow[\ \overline{f}\ ]{} & \operatorname{im} f. \end{array}$$

Indeed, $cf = 0$ first gives a unique $u\colon A \to \operatorname{im} f$ with $mu = f$. Then $muk = fk = 0$ and monicity of m give $uk = 0$, so u factors uniquely through q.

Theorem 1.2.13. In an abelian category the canonical morphism $\overline{f}\colon \operatorname{coim} f \to \operatorname{im} f$ is an isomorphism. In particular, every morphism factors as an epimorphism followed by a monomorphism.

Proof. We first prove that $\overline{f}$ is mono. Let $j\colon L \to \operatorname{coim} f$ be its kernel, and form the pullback

$$\begin{array}{ccc} P & \xrightarrow{\ \alpha\ } & A \\ {\scriptstyle p}\downarrow & & \downarrow{\scriptstyle q} \\ L & \xrightarrow[\ j\ ]{} & \operatorname{coim} f. \end{array}$$

The projection p is epi by Lemma 1.2.11. Since $q\alpha = jp$ and $\overline{f}j = 0$,

$$f\alpha = m\overline{f}q\alpha = m\overline{f}jp = 0.$$

By the kernel property of k, write $\alpha = k\beta$. Then $jp = q\alpha = qk\beta = 0$. Since p is epi, $j = 0$. Because j is mono, its domain L is a zero object. Thus $\overline{f}$ has zero kernel and is mono.

Apply this same argument in $\mathcal{A}^{\mathrm{op}}$. The coimage of the reversed morphism f^{op} is the image of f, its image is the coimage of f, and its comparison is the reversed arrow $\overline{f}$. The argument proves this reversed arrow mono, which says precisely that $\overline{f}$ is epi in $\mathcal{A}$. Proposition 1.2.10 now makes $\overline{f}$ an isomorphism. Hence $f = m(\overline{f}q)$ is an epi followed by a mono. □

For modules, $\operatorname{coim} f = A/\ker f$, $\operatorname{im} f = f(A) \subseteq B$, and $\overline{f}(a + \ker f) = f(a)$.

1.3 Exact Sequences and Chain Complexes

1.3.1 Exact Sequences, Complexes, and Homology

Definition 1.3.1. For $A \xrightarrow{f} B \xrightarrow{g} C$ with $gf = 0$, there is a canonical morphism $\operatorname{im} f \to \ker g$ compatible with their inclusions into B. The pair is *exact at* B if this map is an isomorphism. A sequence is exact if it is exact at each object where two adjacent maps are specified. A short exact sequence is an exact sequence

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

Equivalently, f is a kernel of g and g is a cokernel of f.

Definition 1.3.2. An additive functor is *left exact* if it preserves kernels, *right exact* if it preserves cokernels, and *exact* if it preserves short exact sequences.

Definition 1.3.3. A chain complex $A_\bullet$ in an abelian category consists of objects A_n, $n \in \mathbb{Z}$, and differentials $\partial^A_n \colon A_n \to A_{n-1}$ satisfying $\partial^A_{n-1}\partial^A_n = 0$. A chain map $f \colon A_\bullet \to B_\bullet$ consists of maps $f_n \colon A_n \to B_n$ making the following squares commute:

$$\begin{array}{ccc} A_n & \xrightarrow{\partial^A_n} & A_{n-1} \\ {\scriptstyle f_n}\downarrow & & \downarrow{\scriptstyle f_{n-1}} \\ B_n & \xrightarrow[\partial^B_n]{} & B_{n-1} \end{array} \qquad (n \in \mathbb{Z}).$$

Composition and identities are defined degreewise; the resulting category is denoted $\mathbf{Ch}(\mathcal{A})$.

Definition 1.3.4. The cycle object is $Z_n(A) = \ker \partial^A_n$, and the boundary object is $B_n(A) = \operatorname{im} \partial^A_{n+1}$. The equation $\partial^A_n \partial^A_{n+1} = 0$ gives a canonical mono $B_n(A) \to Z_n(A)$. Define

$$\mathrm{H}_n(A) = \operatorname{coker}(B_n(A) \longrightarrow Z_n(A)).$$

A chain map induces maps on cycles and boundaries by their universal properties and therefore induces $\mathrm{H}_n(f) \colon \mathrm{H}_n(A) \to \mathrm{H}_n(B)$. The category $\mathbf{Ch}(\mathcal{A})$ is abelian, with kernels and cokernels computed degreewise; the induced differentials exist uniquely by the corresponding universal properties. A complex is *acyclic* if every $\mathrm{H}_n(A)$ is zero, equivalently if its underlying sequence is exact in every degree.

1.3.2 Chain Homotopies

Definition 1.3.5. For chain maps $f, g \colon A_\bullet \to B_\bullet$, a chain homotopy from g to f is a family $h_n \colon A_n \to B_{n+1}$ such that

$$f_n - g_n = \partial^B_{n+1} h_n + h_{n-1} \partial^A_n \qquad (n \in \mathbb{Z}).$$

A complex is *contractible* if its identity chain map is chain homotopic to zero.

Proposition 1.3.6. Chain-homotopic maps induce the same maps on homology. In particular every contractible complex is acyclic.

Proof. Let $z \colon Z_n(A) \to A_n$ be the kernel inclusion. Since $\partial^A_n z = 0$, the homotopy identity restricts to $(f_n - g_n)z = \partial^B_{n+1} h_n z$. This factors through $B_n(B)$ and becomes zero in $\mathrm{H}_n(B)$. Thus $\mathrm{H}_n(f) = \mathrm{H}_n(g)$. If $f = \mathrm{id}_A$ and $g = 0$, the identity of $\mathrm{H}_n(A)$ is zero, so $\mathrm{H}_n(A)$ is a zero object. $\square$

The converse fails for modules over a general ring. The exact complex of abelian groups

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

(in degrees $2, 1, 0$) is not contractible: a contracting homotopy in degree 0 would give a section $\mathbb{Z}/2 \to \mathbb{Z}$ of the quotient map, which is impossible. Over a field, every acyclic complex is contractible after choosing splittings of its short exact sequences of cycles.

Modules, Derived Categories, and the Tensor Product

Algebraic Topology I | Fall 2026 | September 9

2.1 Exact Sequences and Mapping Cones

Let $\mathcal{A}$ be a locally small abelian category. We use homological grading.

A concrete example is $\mathcal{A} = \mathrm{Mod}_R$ for a commutative ring R. This is an abelian category: for a module map $f : M \to N$, its kernel is the usual submodule $\ker f \subseteq M$ and its cokernel is $N/\operatorname{im} f$. For an ideal $I \subseteq R$, the inclusion and quotient maps give the short exact sequence $0 \to I \to R \to R/I \to 0$, also viewed as a sequence of complexes concentrated in degree zero.

Recall: Use the chain complexes and chain maps of Lecture 1, Definition 1.3.3, and the cycle objects, boundary objects, homology, and acyclicity of Definition 1.3.4. Write $\mathrm{Ch}(\mathcal{A})$ for their category. A complex is *acyclic* if its homology is zero.

Definition 2.1.1. A chain map is a *quasi-isomorphism* if every induced homology map is an isomorphism.

2.1.1 The Snake Lemma

Lemma 2.1.2. Consider a commutative diagram in an abelian category with exact rows:

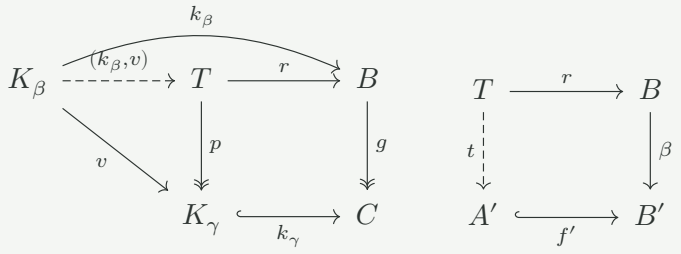
$$\begin{array}{ccccccccc} A & \xrightarrow{f} & B & \xrightarrow{g} & C & \longrightarrow & 0 & & \\ \downarrow{\scriptstyle\alpha} & & \downarrow{\scriptstyle\beta} & & \downarrow{\scriptstyle\gamma} & & & & \\ 0 & \longrightarrow & A' & \xrightarrow[f']{} & B' & \xrightarrow[g']{} & C'. & & \end{array}$$

There is a natural exact sequence

$$\ker\alpha \longrightarrow \ker\beta \longrightarrow \ker\gamma \xrightarrow{\delta} \operatorname{coker}\alpha \longrightarrow \operatorname{coker}\beta \longrightarrow \operatorname{coker}\gamma.$$

If f is monic, the sequence starts with 0; if g' is epic, it ends with 0.

Proof. Write $k_\eta : K_\eta \hookrightarrow \mathrm{dom}(\eta)$ and $q_\eta : \mathrm{cod}(\eta) \twoheadrightarrow Q_\eta$ for the kernel and cokernel of $\eta \in \{\alpha, \beta, \gamma\}$. Denote the induced maps by $u : K_\alpha \to K_\beta$, $v : K_\beta \to K_\gamma$, $\bar{f} : Q_\alpha \to Q_\beta$, and $\bar{g} : Q_\beta \to Q_\gamma$. Form the pullback $T = B \times_C K_\gamma$. The first square makes p epic; since $g'\beta r = \gamma k_\gamma p = 0$, the second defines t uniquely.



The map $j = (f, 0) : A \to T$ has image $\ker p$. Since $tj = \alpha$, the cokernel property of p gives the unique dashed arrow below:

$$\begin{array}{ccccccc} A & \xrightarrow{j} & T & \xrightarrow{p} & K_\gamma & \longrightarrow & 0 \\ {\scriptstyle\alpha}\downarrow & & \downarrow{\scriptstyle t} & & \downarrow{\scriptstyle\delta} & & \\ A' & = & A' & \xrightarrow[q_\alpha]{} & Q_\alpha & \longrightarrow & 0. \end{array}$$

Thus $\delta p = q_\alpha t$. The lift in the first picture satisfies $t(k_\beta, v) = 0$, so $\delta v = 0$. Moreover, $\bar{f}\delta p = q_\beta \beta r = 0$, hence $\bar{f}\delta = 0$ because p is epic. The other composites vu and $\bar{g}\bar{f}$ vanish by exactness of the rows.

To test whether a map $x : X \to Y$ factors through an image, it suffices to lift xe through the original map after an epimorphism $e : E \twoheadrightarrow X$. Such lifts are obtained by pullback along the epimorphism onto that image; epimorphisms detect factorization through subobjects. In each panel below, e may be refined by these pullbacks, and every dashed lift has source E.

Exactness at K_β. For $x : X \to K_\beta$ with $vx = 0$, lift $k_\beta x e$ through f.

$$\begin{array}{ccccccc} & & \overset{\lambda}{\frown} & & & & \\ E & \overset{\tilde{\lambda}}{\dashrightarrow} & K_\alpha & \overset{k_\alpha}{\hookrightarrow} & A & \xrightarrow{\alpha} & A' \\ {\scriptstyle e}\downarrow & & \downarrow{\scriptstyle u} & & \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle f'} \\ X & \xrightarrow[x]{} & K_\beta & \xhookrightarrow[k_\beta]{} & B & \xrightarrow[\beta]{} & B'. \end{array}$$

Indeed, $f'\alpha\lambda = \beta k_\beta x e = 0$ gives $\alpha\lambda = 0$, so the kernel property supplies $\tilde{\lambda}$ and $u\tilde{\lambda} = xe$.

Exactness at K_γ. For $x : X \to K_\gamma$ with $\delta x = 0$, the lifts are

$$\begin{array}{ccccc} E & \overset{s}{\dashrightarrow} & T & \xrightarrow{t} & A' \\ {\scriptstyle e}\downarrow & & \downarrow{\scriptstyle p} & & \downarrow{\scriptstyle q_\alpha} \\ X & \xrightarrow[x]{} & K_\gamma & \xrightarrow[\delta]{} & Q_\alpha \end{array} \qquad \begin{array}{ccc} E & \overset{\lambda}{\dashrightarrow} & A \\ {\scriptstyle ts}\downarrow & & \downarrow{\scriptstyle\alpha} \\ A' & = & A' \end{array} \qquad \begin{array}{ccc} E & \overset{h}{\dashrightarrow} & K_\beta \\ {\scriptstyle e}\downarrow & & \downarrow{\scriptstyle v} \\ X & \xrightarrow[x]{} & K_\gamma. \end{array}$$

Here $q_\alpha ts = 0$ permits λ, and the corrected lift h is defined by

$$k_\beta h = rs - f\lambda, \qquad \beta(rs - f\lambda) = f'(ts - \alpha\lambda) = 0, \qquad g(rs - f\lambda) = k_\gamma x e.$$

Exactness at Q_α. For $x : X \to Q_\alpha$ with $\bar{f}x = 0$, lift first through q_α and then through β:

$$\begin{array}{ccccc} E & \overset{a}{\dashrightarrow} & A' & \xrightarrow{f'} & B' \\ {\scriptstyle e}\downarrow & & \downarrow{\scriptstyle q_\alpha} & & \downarrow{\scriptstyle q_\beta} \\ X & \xrightarrow[x]{} & Q_\alpha & \xrightarrow[\bar{f}]{} & Q_\beta \end{array} \qquad \begin{array}{ccc} E & \overset{b}{\dashrightarrow} & B \\ {\scriptstyle a}\downarrow & & \downarrow{\scriptstyle\beta} \\ A' & \xhookrightarrow[f']{} & B' \end{array} \qquad \begin{array}{ccccc} & & \overset{\nu}{\frown} & & \\ E & \overset{(b,\nu)}{\dashrightarrow} & T & \xrightarrow{p} & K_\gamma \\ {\scriptstyle e}\downarrow & & \downarrow{\scriptstyle q_\alpha t} & & \downarrow{\scriptstyle\delta} \\ X & \xrightarrow[x]{} & Q_\alpha & = & Q_\alpha. \end{array}$$

The kernel property defines ν by $k_\gamma \nu = gb$, since $\gamma g b = g'\beta b = g' f' a = 0$. The pullback lift $(b, \nu) : E \to T$ has $t(b, \nu) = a$, hence $\delta\nu = q_\alpha a = xe$.

Exactness at Q_β. For $x : X \to Q_\beta$ with $\bar{g}x = 0$, lift successively through q_β, γ, and g:

$$\begin{array}{ccccc} E & \overset{b'}{\dashrightarrow} & B' & \xrightarrow{g'} & C' \\ {\scriptstyle e}\downarrow & & \downarrow{\scriptstyle q_\beta} & & \downarrow{\scriptstyle q_\gamma} \\ X & \xrightarrow[x]{} & Q_\beta & \xrightarrow[\bar{g}]{} & Q_\gamma \end{array} \qquad \begin{array}{ccc} E & \overset{b}{\dashrightarrow} & B \\ {\scriptstyle c}\vdots & & \downarrow{\scriptstyle g} \\ C & = & C \end{array} \qquad \begin{array}{ccc} E & \overset{a'}{\dashrightarrow} & A' \\ {\scriptstyle e}\downarrow & & \downarrow{\scriptstyle q_\beta f'} \\ X & \xrightarrow[x]{} & Q_\beta. \end{array}$$

These lifts satisfy $\gamma c = g'b'$ and $gb = c$. The corrected lift a' is the unique map with

$$f'a' = b' - \beta b, \qquad g'(b' - \beta b) = \gamma(c - gb) = 0, \qquad \bar{f} q_\alpha a' = q_\beta(b' - \beta b) = xe.$$

The four panels prove the reverse image–kernel inclusions. If f is monic, then $k_\beta u = f k_\alpha$ makes u monic. If g' is epic, then $\bar{g} q_\beta = q_\gamma g'$ makes $\bar{g}$ epic. Finally, a morphism of the original row diagrams induces a morphism of their pullbacks T; uniqueness of t and δ makes the induced square on connecting maps commute. □

For modules the construction is the familiar rule: lift $c \in \ker \gamma$ to $b \in B$, solve $f'a' = \beta b$, and put $\delta(c) = [a']$. Changing b to $b + fa$ changes a' to $a' + \alpha a$, so the cokernel class does not change.

2.1.2 The Homology Long Exact Sequence

Theorem 2.1.3. A degreewise short exact sequence of complexes

$$0 \longrightarrow A \xrightarrow{\iota} B \xrightarrow{\rho} C \longrightarrow 0 \tag{1}$$

induces natural connecting morphisms $\delta_n : \mathrm{H}_n(C) \to \mathrm{H}_{n-1}(A)$ and a long exact sequence. Five consecutive terms are

$$\mathrm{H}_n(A) \xrightarrow{\mathrm{H}_n(\iota)} \mathrm{H}_n(B) \xrightarrow{\mathrm{H}_n(\rho)} \mathrm{H}_n(C) \xrightarrow{\delta_n} \mathrm{H}_{n-1}(A) \xrightarrow{\mathrm{H}_{n-1}(\iota)} \mathrm{H}_{n-1}(B). \tag{2}$$

For modules, if $\zeta \in C_n$ is a cycle, choose $\beta \in B_n$ with $\rho\beta = \zeta$ and write $\partial^B \beta = \iota\alpha$. Then α is a cycle and δ_n sends the class of ζ to the class of α.

Proof. Put $Q_n(X) = \operatorname{coker}(\partial^X_{n+1}) = X_n / B_n(X)$. The differential induces $\partial_X : Q_n(X) \to Z_{n-1}(X)$. Its kernel is $\mathrm{H}_n(X)$ and its cokernel is $\mathrm{H}_{n-1}(X)$: its image is precisely $B_{n-1}(X)$.

The following commutative diagram has exact rows:

$$\begin{array}{ccccccccc} & & Q_n(A) & \longrightarrow & Q_n(B) & \longrightarrow & Q_n(C) & \longrightarrow & 0 \\ & & {\scriptstyle \partial_A}\downarrow & & \downarrow{\scriptstyle \partial_B} & & \downarrow{\scriptstyle \partial_C} & & \\ 0 & \longrightarrow & Z_{n-1}(A) & \longrightarrow & Z_{n-1}(B) & \longrightarrow & Z_{n-1}(C). & & \end{array} \tag{3}$$

The upper row is exact because cokernels preserve right exact sequences of arrows, applied to the differentials in degrees $n+1, n$. The lower row is exact because kernels preserve left exact sequences of arrows, applied in degrees $n-1, n-2$.

Lemma 2.1.2, in its form with a right exact upper row and a left exact lower row, gives

$$\ker \partial_A \to \ker \partial_B \to \ker \partial_C \xrightarrow{\delta_n} \operatorname{coker} \partial_A \to \operatorname{coker} \partial_B \to \operatorname{coker} \partial_C.$$

Substituting the identifications above proves exactness at all three types of homology terms. The six-term sequences for consecutive n join because their other maps are the maps induced

by ι, ρ. Naturality follows from the snake lemma; see [Stacks, Tag 0117, Long Exact Cohomology Sequence].

For modules, $\rho\partial^B\beta = \partial^C\zeta = 0$, so exactness gives a unique α with $\iota\alpha = \partial^B\beta$. Since $\iota(\partial^A\alpha) = \partial^B\partial^B\beta = 0$ and ι is mono, $\partial^A\alpha = 0$. Changing β to $\beta + \iota\alpha'$ changes α to $\alpha + \partial^A\alpha'$. Changing ζ to $\zeta + \partial^C\zeta'$ and choosing a lift β' of ζ' changes a chosen lift to $\beta + \partial^B\beta'$ and leaves its differential unchanged. Thus the formula is well defined and agrees with the snake map. □

Example 2.1.4. If $A \subseteq B$ is a subcomplex, define the quotient complex B/A degreewise. The theorem gives a long exact sequence with five consecutive terms

$$\mathrm{H}_n(A) \to \mathrm{H}_n(B) \to \mathrm{H}_n(B/A) \xrightarrow{\delta_n} \mathrm{H}_{n-1}(A) \to \mathrm{H}_{n-1}(B).$$

For singular chains of a subspace inclusion $X' \subseteq X$, take $B = \mathrm{C}_*(X; R)$ and $A = \mathrm{C}_*(X'; R)$. A relative cycle is represented by a chain in X whose boundary lies in X'; δ takes the homology class of that boundary in X'.

2.1.3 The Cone Sequence

For cones and termwise split sequences, see [Stacks, Tag 014D]. The Stacks Project uses cohomological grading; these notes use homological grading.

Definition 2.1.5. For $r \in \mathbb{Z}$, set $A[r]_n = A_{n-r}$ and $\partial_n^{A[r]} = (-1)^r\partial_{n-r}^A$. A chain map shifts by $f[r]_n = f_{n-r}$. Thus $\mathrm{H}_n(A[r]) = \mathrm{H}_{n-r}(A)$ canonically. For any biproduct, ι_j and ρ_j denote the canonical inclusions and projections.

Definition 2.1.6. For a chain map $f : A \to B$, define

$$\mathrm{Cone}(f)_n = B_n \oplus A_{n-1}, \qquad \partial_n^{\mathrm{Cone}(f)} = \begin{pmatrix} \partial_n^B & f_{n-1} \\ 0 & -\partial_{n-1}^A \end{pmatrix}.$$

The upper-right entry of its squared differential is $\partial_{n-1}^B f_{n-1} - f_{n-2}\partial_{n-1}^A = 0$; all other entries vanish. There is a degreewise split short exact sequence

$$0 \to B \xrightarrow{\jmath=\iota_1} \mathrm{Cone}(f) \xrightarrow{q=\rho_2} A[1] \to 0. \tag{4}$$

Proposition 2.1.7. The connecting map of (4) is $\mathrm{H}_{n-1}(f)$. Hence

$$\cdots \to \mathrm{H}_n(A) \xrightarrow{\mathrm{H}_n(f)} \mathrm{H}_n(B) \xrightarrow{\mathrm{H}_n(\jmath)} \mathrm{H}_n(\mathrm{Cone}(f)) \xrightarrow{\mathrm{H}_n(q)} \mathrm{H}_{n-1}(A) \xrightarrow{\mathrm{H}_{n-1}(f)} \mathrm{H}_{n-1}(B) \to \cdots \tag{5}$$

is exact. In particular, f is a quasi-isomorphism if and only if $\mathrm{Cone}(f)$ is acyclic.

Proof. A cycle $\zeta \in A_{n-1}$ lifts to $(0, \zeta)$, whose cone differential is $(f(\zeta), 0)$. Thus the connecting map has the asserted positive sign. Categorically, the second-summand section of (4) has the same differential defect f, so this computes the snake map in $\mathcal{A}$. Apply Theorem 2.1.3. If every $\mathrm{H}_n(f)$ is invertible, exactness forces all cone homology to vanish. Conversely, its vanishing makes $\mathrm{H}_n(f)$ both mono and epi, hence an isomorphism in $\mathcal{A}$. □

Example 2.1.8. For $A = B = \mathbb{Z}$ concentrated in degree zero and $f = m$, the cone is $\mathbb{Z} \xrightarrow{m} \mathbb{Z}$ in degrees $1, 0$. Its homology is $\ker(m)$ in degree 1 and $\mathbb{Z}/m\mathbb{Z}$ in degree 0.

More generally, for a commutative ring R and $r \in R$, the cone of multiplication by r on R is $R \xrightarrow{r} R$ in degrees $1, 0$. Its homology is $\{a \in R : ra = 0\}$ in degree 1 and R/rR in degree 0, and is zero in all other degrees. If r is a non-zero-divisor, the quotient map from this cone to R/rR concentrated in degree zero is a quasi-isomorphism.

2.1.4 Recovering a Short Exact Sequence from a Cone

Proposition 2.1.9. For (1), let $\pi : \mathrm{Cone}(\iota) \to C$ have restrictions $\rho : B \to C$ and $0 : A[1] \to C$. Then π is a quasi-isomorphism and the diagram

$$\begin{array}{ccccc} A & \xrightarrow{\iota} & B & \xrightarrow{j} & \mathrm{Cone}(\iota) \\ \| & & \| & & \sim\downarrow\pi \\ A & \xrightarrow[\iota]{} & B & \xrightarrow[\rho]{} & C \end{array}$$

commutes. With $q : \mathrm{Cone}(\iota) \to A[1]$ the second projection, the connecting morphism in Theorem 2.1.3 is

$$\delta_n = -\mathrm{H}_n(q)\,\mathrm{H}_n(\pi)^{-1}. \tag{6}$$

Proof. The map π is a degreewise epimorphism of complexes. Its kernel identifies with $\mathrm{Cone}(\mathrm{id}_A)$: the inclusion sends (α', α) to $(\iota\alpha', \alpha)$. On this kernel, $s_n(\alpha', \alpha) = (0, \alpha')$ is a contracting homotopy, since $\partial s + s\partial = \mathrm{id}$. Theorem 2.1.3 therefore implies that π is a quasi-isomorphism.

For the sign, lift a cycle ζ to β and write $\partial^B\beta = \iota\alpha$. The corresponding cone cycle is $(\beta, -\alpha)$, so $q(\beta, -\alpha) = -\alpha$, while δ_n sends the class of ζ to the class of α. The same computation with the biproduct maps gives the identity in $\mathcal{A}$. □

Thus the cone sequence recovers the long exact sequence of (1); formula (6) identifies the chosen connecting map, including its sign.

2.2 Homotopy and Derived Categories

2.2.1 Passing to Chain Homotopy Classes

Use chain homotopies and contractible complexes as in Lecture 1, Definition 1.3.5. The new construction here is the quotient by chain homotopy.

Definition 2.2.1. Let $N(A, B)$ be the subgroup of nullhomotopic chain maps. The category $K(\mathcal{A})$ has the same objects as $\mathrm{Ch}(\mathcal{A})$ and morphism groups

$$\mathrm{Mor}_{K(\mathcal{A})}(A, B) = \mathrm{Mor}_{\mathrm{Ch}(\mathcal{A})}(A, B)/N(A, B).$$

Composition is induced by composition of chain maps; homotopy is as above.

Proposition 2.2.2. Nullhomotopic maps form an additive ideal. Hence $K(\mathcal{A})$ is additive and every homology functor factors through it. Its isomorphisms are precisely chain homotopy equivalences.

Proof. The sum of homotopies is a homotopy for the sum. If $f = \partial h + h\partial$ and α, β are composable chain maps, then $\beta f \alpha = \partial(\beta h\alpha) + (\beta h\alpha)\partial$. Thus composition descends to the quotient. Lecture 1, Proposition 1.3.6, supplies equality of the induced homology maps. An inverse in the quotient means exactly inverse chain maps up to homotopy. □

Example 2.2.3. The acyclic, noncontractible complex of Lecture 1, following Proposition 1.3.6, is therefore a nonzero object of $K(\mathbf{Ab})$.

2.2.2 Inverting Quasi-Isomorphisms

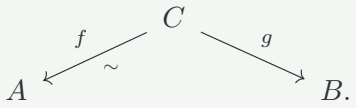
Assume here that $\mathcal{A}$ is essentially small or is a category of modules. The discussion of the size issue begins in Subsection 2.2.3. For a complementary account of the passage from complexes to the homotopy and derived categories, see Weibel [Wei94, §§10.1,10.3].

Definition 2.2.4. Let S be the class of quasi-isomorphisms in $K(\mathcal{A})$. Define

$$D(\mathcal{A}) = K(\mathcal{A})[S^{-1}].$$

It has the same objects as $K(\mathcal{A})$ and a functor $Q : K(\mathcal{A}) \to D(\mathcal{A})$ sending each $s \in S$ to an isomorphism. Its universal property is that precomposition with Q identifies functors out of $D(\mathcal{A})$ with functors out of $K(\mathcal{A})$ which invert S; natural transformations correspond as well.

Morphisms are represented by spans in $K(\mathcal{A})$:



The represented map is $Q(g)Q(f)^{-1}$.

Here $\sim$ marks a quasi-isomorphism: f induces an isomorphism in every homology degree, but need not be an isomorphism of complexes. Two spans are equivalent if they admit a common refinement in $K(\mathcal{A})$ on which both maps agree, with the common map to A a quasi-isomorphism. This is the standard calculus-of-fractions description for quasi-isomorphisms; we use it here without proof.

Homology descends to $D(\mathcal{A})$ by the universal property. A complex becomes zero precisely when it is acyclic: use the quasi-isomorphism $0 \to A$ in one direction, and homology in the other. Similarly, a chain map becomes invertible precisely when it is a quasi-isomorphism.

Definition 2.2.5. The quasi-isomorphism $\pi : \mathrm{Cone}(\iota) \to C$ gives a morphism

$$\delta = Q(-q)Q(\pi)^{-1} : C \longrightarrow A[1]$$

in $D(\mathcal{A})$, with $\mathrm{H}_n(\delta) = \delta_n$. Thus the short exact sequence determines

$$A \xrightarrow{\iota} B \xrightarrow{\rho} C \xrightarrow{\delta} A[1].$$

This is its distinguished triangle. With our signs, cone triangles have the form $X \xrightarrow{f} Y \xrightarrow{\jmath} \mathrm{Cone}(f) \xrightarrow{-q} X[1]$.

2.2.3 Complexes and Essentially Small Categories

Proposition 2.2.6. If $\mathcal{A}$ is locally small, then $\mathrm{Ch}(\mathcal{A})$ and $K(\mathcal{A})$ are locally small, even for unbounded complexes.

Proof. For fixed A, B, every chain map lies in the set

$$\prod_{n \in \mathbb{Z}} \mathrm{Mor}_{\mathcal{A}}(A_n, B_n).$$

The chain-map equations select a subset. The nullhomotopic maps form a subgroup of this set-sized abelian group; quotienting by it gives a set. The countable product here is a product of sets, not a product that must exist in $\mathcal{A}$. □

Proposition 2.2.7. If $\mathcal{A}$ is essentially small, then $D(\mathcal{A})$ is locally small.

Proof. Replace $\mathcal{A}$ by a small skeleton. Its complexes and chain maps form sets. Finite words in arrows and formal inverses also form a set; localization quotients this set by relations. □

2.2.4 Local Smallness of Derived Module Categories

Proposition 2.2.8. For a set-sized ring R, the unbounded category $D(\mathrm{Mod}_R)$ is locally small.

Proof. Fix complexes A, B, and an infinite cardinal κ at least $|R|$, $\aleph_0$, and the cardinality of the disjoint union of the underlying sets A_n. Consider any span $A \xleftarrow{f} C \xrightarrow{g} B$ with f a quasi-isomorphism.

Choose cycles in C representing every homology class of A. They generate a subcomplex $C^{(0)}$ of total cardinality at most κ. From $C^{(m)}$, form $C^{(m+1)}$ by adjoining one bounding element in C for each cycle whose image bounds in A, then closing under the differential and module operations. These bounds exist since $\mathrm{H}_*(f)$ is injective; the cardinality remains at most κ.

The union $C' = \bigcup_{m\geq 0} C^{(m)}$ still has size at most κ, and $f|_{C'}$ is a quasi-isomorphism: chosen cycles give surjectivity, and the successive bounds give injectivity. Thus every span refines to one whose middle complex has size at most κ. There is a set of such complexes up to isomorphism and a set of maps from each to A, B. Quotienting this set of spans gives the required morphism set. □

Remark 2.2.9. For background on size questions beyond the module case, see [CN08]. General treatments using injective replacements appear in [Stacks, Tag 070G and Tag 079P].

2.3 Tensor Products and Their Computation

2.3.1 The Universal Bilinear Map

All rings have an identity and all modules are unital right modules; unless specified, modules and presentations need not be finitely generated. Write Mod_R for the category of right R-modules and $\mathrm{Mor}_{\mathrm{Mod}_R}(A, B)$ for its morphism groups. Fix a commutative ring R. Every right R-module then has the canonical left action $r \cdot n = nr$, which is used on the second factor of a tensor product. An R-bilinear map $\beta : A \times B \to C$ is linear in each variable separately. The tensor product turns such a map into a linear map with a single input.

Definition 2.3.1. The tensor product of A and B is an R-module $A \otimes_R B$ with a bilinear map $\tau : A \times B \to A \otimes_R B$ such that, for every bilinear $\beta : A \times B \to C$, there is a unique linear $\tilde{\beta}$ with $\tilde{\beta}\tau = \beta$. Write $\tau(m, n) = m \otimes n$.

$$\begin{array}{ccc} A \times B & \xrightarrow{\ \beta\ } & C \\ {\scriptstyle\tau}\downarrow & \nearrow_{\tilde{\beta}} & \\ A \otimes_R B & & \end{array}$$

Proposition 2.3.2. Choose arbitrary free resolutions

$$\cdots \longrightarrow P_1 \xrightarrow{\partial_1^P} P_0 \xrightarrow{\epsilon} A \longrightarrow 0, \qquad \cdots \longrightarrow Q_1 \xrightarrow{\partial_1^Q} Q_0 \xrightarrow{\eta} B \longrightarrow 0.$$

Here a free resolution means an exact augmented sequence of free R-modules. Choose bases $P_a = \bigoplus_{\lambda \in \Lambda_a} R$ and $Q_a = \bigoplus_{\gamma \in \Gamma_a} R$ for $a = 0, 1$, and write $\partial_1^P(p_{\lambda'}) = \sum_{\lambda \in \Lambda_0} p_\lambda a_{\lambda\lambda'}$ for $\lambda' \in \Lambda_1$ and $\partial_1^Q(q_{\gamma'}) = \sum_{\gamma \in \Gamma_0} q_\gamma b_{\gamma\gamma'}$ for $\gamma' \in \Gamma_1$. There is a natural coequalizer formula in R-modules

$$A \otimes_R B \cong \mathrm{coeq}\left(\left(\bigoplus_{\Lambda_1 \times \Gamma_0} R \right) \oplus \left(\bigoplus_{\Lambda_0 \times \Gamma_1} R \right) \underset{0}{\overset{D}{\rightrightarrows}} \bigoplus_{\Lambda_0 \times \Gamma_0} R \right),$$

where, writing $[\lambda,\gamma]$ for the target basis,

$$D([\lambda',\gamma],0)=\sum_{\lambda}[\lambda,\gamma]a_{\lambda\lambda'},\qquad D(0,[\lambda,\gamma'])=\sum_{\gamma}[\lambda,\gamma]b_{\gamma\gamma'}.$$

Thus D is the map induced by the two degree-one differentials, $\partial_1^P\otimes\mathrm{id}+\mathrm{id}\otimes\partial_1^Q$, under the standard identification of tensor products of free modules with free modules on pairs of basis elements.

Proof. A linear map from $\bigoplus_{\Lambda_0\times\Gamma_0}R$ to an R-module C is the same as an R-bilinear map $\beta : P_0\times Q_0\to C$: specify its values on basis pairs and extend in both variables. Such a linear map coequalizes D and zero exactly when $\beta(\partial_1^P a,q)=0$ for all $a\in P_1$, $q\in Q_0$, and $\beta(p,\partial_1^Q b)=0$ for all $p\in P_0$, $b\in Q_1$. Equivalently, β vanishes on $\operatorname{im}\partial_1^P\times Q_0$ and on $P_0\times\operatorname{im}\partial_1^Q$. Since $A=P_0/\operatorname{im}\partial_1^P$ and $B=Q_0/\operatorname{im}\partial_1^Q$, these conditions say precisely that β descends uniquely to an R-bilinear map $A\times B\to C$. The coequalizer therefore has the universal property in Definition 2.3.1. In a module category the coequalizer of D and zero is the quotient by $\operatorname{im}D$. The universal property also gives uniqueness up to a unique isomorphism compatible with the universal bilinear map. □

Corollary 2.3.3. In other words, with the free resolutions and bases of Proposition 2.3.2, $A\otimes_R B$ is the free module on symbols $[\lambda,\gamma]$, for $(\lambda,\gamma)\in\Lambda_0\times\Gamma_0$, modulo the relations

$$\sum_{\lambda}[\lambda,\gamma]a_{\lambda\lambda'}=0\quad(\lambda'\in\Lambda_1,\ \gamma\in\Gamma_0),\qquad\sum_{\gamma}[\lambda,\gamma]b_{\gamma\gamma'}=0\quad(\lambda\in\Lambda_0,\ \gamma'\in\Gamma_1).$$

The class of $[\lambda,\gamma]$ represents $\epsilon(p_\lambda)\otimes\eta(q_\gamma)$. All sums are finite because each differential sends a basis element to a finite linear combination. Only degrees zero and one of the chosen resolutions enter this formula.

Proof. The coequalizer of D and zero is $\operatorname{coker}D$. The image of D is generated by its values on the basis elements of its two source summands, which are exactly the displayed relations. □

Every tensor is a finite sum of elementary tensors. Such an expression need not be unique, and a tensor need not be elementary. For example, over a field k, the tensor $e_1\otimes e_1+e_2\otimes e_2$ in $k^2\otimes_k k^2$ corresponds to a rank-two matrix, whereas an elementary tensor corresponds to a matrix of rank at most one.

Linear maps $f : A\to A'$ and $g : B\to B'$ induce $f\otimes g : A\otimes_R B\to A'\otimes_R B'$ by $m\otimes n\mapsto f(m)\otimes g(n)$. The universal property also gives the natural isomorphisms

$$R\otimes_R B\cong B,\qquad A\otimes_R B\cong B\otimes_R A,\qquad\left(\bigoplus_{\lambda\in\Lambda}R\right)\otimes_R B\cong\bigoplus_{\lambda\in\Lambda}B.$$

The first map sends $r\otimes n$ to nr, the second interchanges the factors, and the third sends $e_\lambda\otimes n$ to n in the summand indexed by λ. Their inverse maps are specified on the same generators. Each vector in a direct sum has finite support. For a finite indexing set Λ, we write R^Λ and B^Λ.

Proposition 2.3.4. There is a natural isomorphism

$$\mathrm{Mor}_{\mathrm{Mod}_R}(A\otimes_R B,C)\cong\mathrm{Mor}_{\mathrm{Mod}_R}(A,\mathrm{Mor}_{\mathrm{Mod}_R}(B,C)).$$

For fixed B, tensoring preserves cokernels and direct sums. In particular, if $A'\xrightarrow{u}A\xrightarrow{v}A''\to 0$ is exact, then

$$A'\otimes_R B\xrightarrow{u\otimes\mathrm{id}}A\otimes_R B\xrightarrow{v\otimes\mathrm{id}}A''\otimes_R B\longrightarrow 0$$

is exact.

Proof. Commutativity makes $\mathrm{Mor}_{\mathrm{Mod}_R}(B,C)$ an R-module by $(\phi r)(n)=\phi(n)r$. A bilinear map $\beta : A\times B\to C$ determines the linear map $\tilde{\beta} : A\to \mathrm{Mor}_{\mathrm{Mod}_R}(B,C)$ specified by $\tilde{\beta}(m)(n)=\beta(m,n)$. This correspondence identifies bilinear maps with the right-hand side. The universal property identifies them with the left-hand side. Thus $-\otimes_R B$ is a left adjoint, so it preserves colimits, including cokernels and direct sums. Explicitly, a bilinear map on $A\times B$ vanishing on $u(A')\times B$ factors uniquely through $A''\times B$; this is exactly the asserted cokernel property. □

2.3.2 Tensor Products from Presentations

A presentation is an exact sequence $\bigoplus_{\gamma\in\Gamma} R \xrightarrow{\phi} \bigoplus_{\lambda\in\Lambda} R \to A\to 0$. Its basis elements indexed by Λ give generators m_λ of A, and its columns give relations among those generators. Such a presentation always exists: take a free module surjecting onto A, then a free module surjecting onto the kernel.

Proposition 2.3.5. For an arbitrary presentation of A and an arbitrary module B,

$$A\otimes_R B \cong \mathrm{coker}(\phi_B : \bigoplus_{\gamma\in\Gamma} B \longrightarrow \bigoplus_{\lambda\in\Lambda} B). \tag{7}$$

If $\phi(e_\gamma)=\sum_\lambda e_\lambda a_{\lambda\gamma}$, then

$$\phi_B((n_\gamma)_{\gamma\in\Gamma}) = \left(\sum_{\gamma\in\Gamma} n_\gamma a_{\lambda\gamma}\right)_{\lambda\in\Lambda}.$$

Each column of ϕ has finite support, and each input has finite support, so this formula is defined even for infinite sets.

If also $\bigoplus_{\gamma'\in\Gamma'} R \xrightarrow{\psi} \bigoplus_{\lambda'\in\Lambda'} R \to B\to 0$ is a presentation, then

$$\begin{aligned}&\left(\bigoplus_{(\gamma,\lambda')\in\Gamma\times\Lambda'} R\right)\oplus\left(\bigoplus_{(\lambda,\gamma')\in\Lambda\times\Gamma'} R\right)\\ &\xrightarrow{(\phi\otimes\mathrm{id}\ \ \mathrm{id}\otimes\psi)} \bigoplus_{(\lambda,\lambda')\in\Lambda\times\Lambda'} R \longrightarrow A\otimes_R B\longrightarrow 0.\end{aligned} \tag{8}$$

This sequence is exact. The first arrow applies $\phi\otimes\mathrm{id}$ to the first summand and $\mathrm{id}\otimes\psi$ to the second, then adds their images in the free module on $\Lambda\times\Lambda'$.

Proof. Apply Proposition 2.3.4 to the presentation of A and identify the tensors of free modules with direct sums of B. This proves (7). For (8), the generator indexed by (λ,λ') maps to $m_\lambda\otimes n_{\lambda'}$. The relations from ϕ impose every relation among the m_λ for each fixed $n_{\lambda'}$, and those from ψ impose every relation among the $n_{\lambda'}$ for each fixed m_λ. A linear map out of this cokernel is therefore precisely a bilinear map $A\times B\to C$. Apply Definition 2.3.1. □

Remark 2.3.6. For finite presentations $\phi : R^a\to R^b$ and $\psi : R^c\to R^d$, formula (8) is the cokernel of the block matrix

$$\left[\,\phi\otimes I_d \;\middle|\; I_b\otimes\psi\,\right] : R^{ad+bc}\longrightarrow R^{bd},$$

using the basis order specified by the tensor factors. Over $\mathbb{Z}$ one can reduce finite matrices by Smith normal form. Over an arbitrary ring, the cokernel presentation is still a complete algebraic description, but a finite effective algorithm for simplifying it requires additional input about the ring. Arbitrary modules need not admit finite presentations.

Example 2.3.7. For any ideal $I\subseteq R$,

$$(R/I)\otimes_R B\cong B/IB.$$

The map sends $(r+I)\otimes n$ to $nr+IB$, with inverse induced by the map $B\to (R/I)\otimes_R B$ given by $n\mapsto (1+I)\otimes n$. Here IB consists of finite sums of elements $an=na$ with $a\in I$ and $n\in B$. If B is a commutative R-algebra, IB is the ideal generated by the image of I. Hence

$$(R/I)\otimes_R (R/J)\cong R/(I+J).$$

For positive integers a,b, this gives $\mathbb{Z}/a\mathbb{Z}\otimes_{\mathbb{Z}}\mathbb{Z}/b\mathbb{Z}\cong \mathbb{Z}/\gcd(a,b)\mathbb{Z}$. In particular, nonzero modules can have zero tensor product.

Example 2.3.8. Let $A=\operatorname{coker}\phi$ over $\mathbb{Z}$, where

$$\phi=\begin{pmatrix}2&0\\4&6\end{pmatrix}:\mathbb{Z}^2\longrightarrow\mathbb{Z}^2,\qquad B=\mathbb{Z}/4\mathbb{Z}.$$

Tensoring replaces ϕ by its action on B^2, which is the diagonal matrix with entries $2,2$. Therefore

$$A\otimes_{\mathbb{Z}} B\cong B^2/(2B\oplus 2B)\cong(\mathbb{Z}/2\mathbb{Z})^2.$$

In coordinates, the map on $(\mathbb{Z}/4\mathbb{Z})^2$ is

$$\phi_B(a,b)=(2a,4a+6b)=(2a,2b).$$

Its image consists of $(0,0),(2,0),(0,2),(2,2)$. The map $(\mathbb{Z}/4\mathbb{Z})^2\to(\mathbb{Z}/2\mathbb{Z})^2$ given by $(a,b)\mapsto(a \bmod 2, b \bmod 2)$ is surjective with precisely this kernel, so it gives the displayed quotient isomorphism. For comparison, the integer row operation $\text{row}_2\leftarrow \text{row}_2-2\text{row}_1$ changes ϕ to $\operatorname{diag}(2,6)$. Thus $A\cong\mathbb{Z}/2\mathbb{Z}\oplus\mathbb{Z}/6\mathbb{Z}$, and tensoring the two summands with $\mathbb{Z}/4\mathbb{Z}$ gives the same two copies of $\mathbb{Z}/2\mathbb{Z}$.

Resolutions, Tor, and Ext

Algebraic Topology I | Fall 2026 | September 14 and 16

3.1 Resolutions and Derived Functors

3.1.1 Projective and Injective Modules

The tensor-product construction in Lecture 2, Section 2.3, computes degree zero. To measure the exactness lost by tensoring, we extend a presentation to an exact sequence. The constructions in this section work for right modules over any ring; we continue with commutative R until Section 3.3.

Definition 3.1.1. Let $\mathcal{A}$ be an abelian category. An object $P \in \mathcal{A}$ is *projective* if for every epimorphism $q : E \to L$ in $\mathcal{A}$ and morphism $f : P \to L$, there exists a morphism $\tilde{f} : P \to E$ such that $q \circ \tilde{f} = f$. A right R-module is projective if it is a projective object of the abelian category of right R-modules.

Definition 3.1.2. Let $\mathcal{A}$ be an abelian category. An object $I \in \mathcal{A}$ is *injective* if for every monomorphism $j : A \to B$ in $\mathcal{A}$ and morphism $f : A \to I$, there exists a morphism $\tilde{f} : B \to I$ such that $\tilde{f} \circ j = f$. A right R-module is injective if it is an injective object of the abelian category of right R-modules.

Remark 3.1.3. The projectivity condition is equivalent to exactness of $\mathrm{Mor}_{\mathcal{A}}(P, -)$; injectivity is equivalent to exactness of $\mathrm{Mor}_{\mathcal{A}}(-, I) : \mathcal{A}^{\mathrm{op}} \to \mathrm{Ab}$. In each case the lifting property gives the required surjection, while exactness at the other terms follows from the kernel or cokernel universal property. Free modules are projective: lift the images of their basis elements. A projective module is a direct summand of a free module because a free surjection onto it splits. Conversely, a direct summand of a free module has the lifting property.

Definition 3.1.4. An abelian category $\mathcal{A}$ has *enough projectives* if for every $M \in \mathcal{A}$ there exist a projective object $P \in \mathcal{A}$ and an epimorphism $P \to M$. It has *enough injectives* if for every $M \in \mathcal{A}$ there exist an injective object $I \in \mathcal{A}$ and a monomorphism $M \to I$.

Module categories have both properties. Free modules supply enough projectives; the following construction supplies enough injectives.

Proposition 3.1.5. Every right module over a unital ring embeds in an injective right module.

Proof. First, a divisible abelian group D is injective. Given a map from a subgroup A of an abelian group B, use Zorn's lemma to choose a maximal extension to a subgroup C. If $b \notin C$, the set of integers n with $nb \in C$ is an ideal. If it is zero, choose any image of b. Otherwise it is $d\mathbb{Z}$ for $d > 0$; choose $a \in D$ with $da = f(db)$ and set $f(c + nb) = f(c) + na$. The ideal description proves this is well defined. Both cases extend the map to $C + \mathbb{Z}b$, contradicting maximality. Hence $C = B$.

The group $\mathbb{Q}/\mathbb{Z}$ is divisible, and its characters separate points of any abelian group M: a nonzero element has a cyclic subgroup admitting a character nonzero on that element, which extends by injectivity. Therefore M embeds by evaluation into $D = (\mathbb{Q}/\mathbb{Z})^{\mathrm{Mor}_{\mathrm{Ab}}(M, \mathbb{Q}/\mathbb{Z})}$. This product is injective because each component of a map extends.

For a right R-module M, take such an embedding $j : M \hookrightarrow D$ of underlying abelian groups.

Put $E = \mathrm{Mor}_{\mathrm{Ab}}(R, D)$ and define $(\phi r)(s) = \phi(rs)$. The bijection

$$\mathrm{Mor}_{\mathrm{Mod}_R}(X, E) \cong \mathrm{Mor}_{\mathrm{Ab}}(X, D)$$

sends $f : X \to E$ to the map $X \to D$ whose value at a is $f(a)(1)$. Its inverse sends $g : X \to D$ to the map $X \to E$ whose value at a is the homomorphism $R \to D$ given by $s \mapsto g(as)$. Since forgetting the R-action is exact and D is injective, E is injective. The map $M \to E$ given by the formula $m \mapsto \phi_m$, where $\phi_m : R \to D$ is given by $\phi_m(s) = j(ms)$ is R-linear and injective, as evaluation at 1 recovers j. □

Example 3.1.6. The abelian group $\mathbb{Z}$ is free and therefore projective. It is not injective: the map $2\mathbb{Z} \to \mathbb{Z}$ sending 2 to 1 cannot extend to $\mathbb{Z}$. Both $\mathbb{Q}$ and $\mathbb{Q}/\mathbb{Z}$ are injective by divisibility. The nonzero group $\mathbb{Q}/\mathbb{Z}$ is not projective: projective abelian groups are summands of free abelian groups and hence torsion-free, whereas $\mathbb{Q}/\mathbb{Z}$ is torsion. Thus

$$0 \longrightarrow \mathbb{Z} \longrightarrow \mathbb{Q} \longrightarrow \mathbb{Q}/\mathbb{Z} \longrightarrow 0$$

is an injective resolution of $\mathbb{Z}$, with $\mathbb{Q}$ in cochain degree zero.

3.1.2 Projective Resolutions and Their Comparison

Definition 3.1.7. A projective resolution of M is an exact augmented sequence

$$\cdots \longrightarrow P_2 \xrightarrow{\partial_2} P_1 \xrightarrow{\partial_1} P_0 \xrightarrow{\epsilon} M \longrightarrow 0 \tag{9}$$

with every P_i projective. The associated chain complex $P_\bullet$ has $P_i = 0$ for $i < 0$, differential $\partial_0 = 0$, $\mathrm{H}_0(P_\bullet) \cong M$, and $\mathrm{H}_i(P_\bullet) = 0$ for $i > 0$. The augmentation is not part of the unaugmented complex.

Every module has a free resolution. Choose a free surjection $P_0 \to M$; choose a free surjection $P_1 \to \ker \epsilon$; choose a free surjection $P_2 \to \ker \partial_1$; and continue. The successive kernels are called syzygies. For a completely general choice, use a free module on the underlying set of each kernel. The process may require infinite rank and may continue indefinitely.

Remark 3.1.8. A presentation supplies $\partial_1 : P_1 \to P_0$, but does not assert that ∂_1 is injective. To compute degree-one derived functors one generally also needs ∂_2, which records the relations among the original relations. To compute degree i, one needs the resolution through $P_{i+1} \to P_i \to P_{i-1}$.

Theorem 3.1.9. Let $P_\bullet \to M$ and $Q_\bullet \to L$ be projective resolutions. Every map $f : M \to L$ lifts to a chain map $\tilde{f} : P_\bullet \to Q_\bullet$. Any two lifts of the same f are chain homotopic. In particular, any two projective resolutions of M are chain homotopy equivalent by maps lifting id_M.

Proof. Lift $f\epsilon_P$ through $Q_0 \twoheadrightarrow L$ using projectivity of P_0. Having constructed the lift through degree $n-1$, the map $\tilde{f}_{n-1}\partial_n$ lands in the kernel of the next differential (or of the augmentation when $n = 1$). Exactness identifies this kernel with the image of Q_n. Projectivity of P_n gives $\tilde{f}_n$.

For two lifts α, β, construct $\gamma_n : P_n \to Q_{n+1}$ satisfying

$$\alpha_n - \beta_n = \partial^Q_{n+1}\gamma_n + \gamma_{n-1}\partial^P_n, \qquad \gamma_{-1} = 0.$$

In degree zero, $\alpha_0 - \beta_0$ is killed by the augmentation, so lift it through ∂^Q_1. Inductively, the chain-map identities and the formula already established in degree $n-1$ show that

$$\partial^Q_n(\alpha_n - \beta_n - \gamma_{n-1}\partial^P_n) = 0.$$

Exactness and projectivity give the required γ_n. Finally, lifts of id_M in both directions have composites homotopic to the appropriate identities, by the uniqueness statement. □

Theorem 3.1.10. Let R be a principal ideal domain. Every submodule of a free R-module is free, with no restriction on the rank.

Proof. Well-order a basis $(e_\alpha)_{\alpha<\kappa}$ of F and, for $N \subseteq F$, put $F_\beta = \bigoplus_{\alpha<\beta} Re_\alpha$ and $N_\beta = N \cap F_\beta$. The e_α-coordinate map identifies $N_{\alpha+1}/N_\alpha$ with an ideal $I_\alpha \subseteq R$. If $I_\alpha = (a_\alpha) \neq 0$, choose $x_\alpha \in N_{\alpha+1}$ with coordinate a_α. Subtracting multiples of x_α shows $N_{\alpha+1} = N_\alpha \oplus Rx_\alpha$: the sum is direct because $ra_\alpha = 0$ implies $r = 0$. If $I_\alpha = 0$, then $N_{\alpha+1} = N_\alpha$. At a limit ordinal λ, $N_\lambda = \bigcup_{\beta<\lambda} N_\beta$, since every vector has finite support. Transfinite induction therefore gives

$$N_\beta = \bigoplus_{\substack{\alpha<\beta \\ I_\alpha \neq 0}} Rx_\alpha.$$

Indeed the successor step is the displayed direct-sum decomposition; at a limit, every vector and every finite linear relation already occur at an earlier stage. Taking $\beta = \kappa$ gives the required basis of N. □

Corollary 3.1.11. Every module M over a principal ideal domain admits a free resolution $0 \to P_1 \to P_0 \to M \to 0$, allowing arbitrary ranks.

Proof. Take a free surjection $P_0 \to M$. Its kernel P_1 is free by Theorem 3.1.10. □

3.1.3 Left and Right Derived Functors

Let $\mathcal{A}, \mathcal{B}$ be abelian categories and let all functors in this subsection be additive. The constructions apply to objects of $\mathcal{A}$; complexes of objects require a further construction. The construction by resolutions is developed further in Weibel [Wei94, Chapter 2, pp. 30–65].

Definition 3.1.12. Let $\mathcal{A}$ be an abelian category and $M \in \mathcal{A}$. An *injective resolution* of M is an exact sequence $0 \to M \to I^0 \to I^1 \to \cdots$ in which every I^q is injective. Its associated cochain complex $I^\bullet$ has $I^q = 0$ for $q < 0$; the augmentation $M \to I^0$ is not part of this complex.

Definition 3.1.13. If $\mathcal{A}$ has enough projectives and $F : \mathcal{A} \to \mathcal{B}$ is right exact, choose a projective resolution $P_\bullet \to M$ and set

$$\mathrm{L}_n F(M) = \mathrm{H}_n(F(P_\bullet)), \qquad n \geq 0.$$

If $\mathcal{A}$ has enough injectives and $G : \mathcal{A} \to \mathcal{B}$ is left exact, construct an exact sequence $0 \to M \to I^0 \to I^1 \to \cdots$ with each I^q injective, by successively embedding the cokernels. Set

$$\mathrm{R}^n G(M) = \mathrm{H}^n(G(I^\bullet)), \qquad n \geq 0.$$

For a morphism $f : M \to N$, choose a chain map $\tilde{f}$ between projective resolutions lifting f and define $\mathrm{L}_n F(f) = \mathrm{H}_n(F(\tilde{f}))$. Dually, choose a cochain map $\hat{f}$ between injective resolutions extending f and define $\mathrm{R}^n G(f) = \mathrm{H}^n(G(\hat{f}))$. These assignments define functors $\mathrm{L}_n F, \mathrm{R}^n G : \mathcal{A} \to \mathcal{B}$; their well-definedness is proved in Proposition 3.1.14. For a cochain complex, H^n means the kernel of the differential out of degree n modulo the image of the differential into degree n.

Proposition 3.1.14. These constructions give functors independent of the resolutions up to canonical natural isomorphism. There are natural isomorphisms $\mathrm{L}_0 F \cong F$ and $\mathrm{R}^0 G \cong G$. For $n > 0$, $\mathrm{L}_n F(P) = 0$ for projective P and $\mathrm{R}^n G(I) = 0$ for injective I.

Proof. The comparison theorem and its dual lift maps of objects to maps of resolutions uniquely up to chain or cochain homotopy. Additive functors preserve the homotopy equations,

so their induced homology maps are unique and respect identities and composition. Comparing identity lifts between two resolutions gives mutually inverse isomorphisms; uniqueness shows their compatibility with a third resolution. Right exactness identifies the cokernel of $F(P_1) \to F(P_0)$ with $F(M)$. Left exactness identifies the kernel of $G(I^0) \to G(I^1)$ with $G(M)$. Finally use the resolution concentrated in degree zero when the object is itself projective or injective, respectively. □

Lemma 3.1.15. Let $0 \to M' \to M \to M'' \to 0$ be exact in an abelian category. Given projective resolutions $P'_\bullet \to M'$ and $P''_\bullet \to M''$, there is a projective resolution $P_\bullet \to M$ fitting into a degreewise split exact sequence $0 \to P'_\bullet \to P_\bullet \to P''_\bullet \to 0$. Dually, injective resolutions of M' and M'' extend to a degreewise split exact sequence $0 \to I'^\bullet \to I^\bullet \to I''^\bullet \to 0$ resolving the given sequence.

Proof. Start with $(K'_0, K_0, K''_0) = (M', M, M'')$. Inductively suppose $0 \to K'_n \xrightarrow{i_n} K_n \xrightarrow{\pi_n} K''_n \to 0$ is exact. The chosen end resolutions give epimorphisms $e'_n : P'_n \to K'_n$ and $e''_n : P''_n \to K''_n$ with kernels K'_{n+1} and K''_{n+1}. Projectivity gives a lift $t_n : P''_n \to K_n$ of e''_n. Set $P_n = P'_n \oplus P''_n$ and $e_n = (i_n e'_n, t_n)$. Then

$$\begin{array}{ccccccccc} 0 & \longrightarrow & P'_n & \hookrightarrow & P'_n \oplus P''_n & \twoheadrightarrow & P''_n & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle e'_n} & & \downarrow{\scriptstyle e_n} & & \downarrow{\scriptstyle e''_n} & & \\ 0 & \longrightarrow & K'_n & \xrightarrow[i_n]{} & K_n & \xrightarrow[\pi_n]{} & K''_n & \longrightarrow & 0 \end{array}$$

commutes. The map e_n is epi: if $ue_n = 0$, then $ui_n = 0$, so $u = v\pi_n$; now $ve''_n = ut_n = 0$ forces $v = 0$. Put $K_{n+1} = \ker e_n$. To verify the next exact sequence, write $j''_{n+1} : K''_{n+1} \hookrightarrow P''_n$ and factor $t_n j''_{n+1} = i_n h_n$. The kernel universal property identifies $K_{n+1} \to K''_{n+1}$ with the pullback of e'_n along $-h_n$. It is therefore epi by Lecture 1, Lemma 1.2.11, and its kernel is K'_{n+1}. Thus

$$0 \longrightarrow K'_{n+1} \longrightarrow K_{n+1} \longrightarrow K''_{n+1} \longrightarrow 0$$

is exact, completing the induction. If $j_n : K_n \hookrightarrow P_{n-1}$ is the kernel inclusion, define $d_n = j_n e_n$ for $n \geq 1$, with augmentation e_0. The equality $e_{n-1} j_n = 0$ gives $d^2 = 0$; monicity of j_n and epimorphicity of e_{n+1} give $\ker d_n = K_{n+1} = \operatorname{im} d_{n+1}$ inside P_n. The construction of the kernel maps makes inclusion and projection chain maps, and the objects P_n are projective.

For the dual, apply this proof in the opposite abelian category, where injectives are projectives. Explicitly, at stage n take $0 \to C'^n \xrightarrow{i^n} C^n \xrightarrow{\pi^n} C''^n \to 0$, starting with the given sequence. The end resolutions supply injections $a'^n : C'^n \to I'^n$ and $a''^n : C''^n \to I''^n$. Extend a'^n to $s^n : C^n \to I'^n$, put $I^n = I'^n \oplus I''^n$, and set

$$a^n = (s^n, a''^n \pi^n) : C^n \longrightarrow I'^n \oplus I''^n.$$

This is monic; the dual of the kernel argument proves that its cokernel fits into $0 \to C'^{n+1} \to C^{n+1} \to C''^{n+1} \to 0$. With $q^n : I^n \to C^{n+1}$ the cokernel map, the differentials $d^n = a^{n+1} q^n$ give the asserted exact resolutions and commuting maps. Thus the dual uses precisely the proved induction with arrows reversed.

These constructions also lift a morphism of short exact sequences. At each projective stage first lift the two end maps. The discrepancy on P''_n between the desired middle map and the chosen end lift lands in the target K'_n; lift it through the target $P'_n \twoheadrightarrow K'_n$ and add it as the off-diagonal component. This gives a commuting diagram and hence maps on the next kernels. The dual argument applies to injectives. The comparison theorem makes the induced derived maps independent of these choices, so the resulting long exact sequences are natural. □

Remark 3.1.16. Lemma 3.1.15 resolves a short exact sequence by a degreewise split short exact sequence of projective resolutions. Additivity preserves each splitting, so applying F

and taking homology gives

$$\cdots \to \mathrm{L}_1 F(M) \to \mathrm{L}_1 F(M'') \to F(M') \to F(M) \to F(M'') \to 0.$$

The dual construction with injectives gives

$$0 \to G(M') \to G(M) \to G(M'') \to \mathrm{R}^1 G(M') \to \mathrm{R}^1 G(M) \to \cdots.$$

In particular, an exact sequence $0 \to K \to P \to M \to 0$ with P projective gives $\mathrm{L}_{n+1} F(M) \cong \mathrm{L}_n F(K)$ for $n \geq 1$; an exact sequence $0 \to M \to I \to C \to 0$ with I injective gives $\mathrm{R}^{n+1} G(M) \cong \mathrm{R}^n G(C)$ for $n \geq 1$. These are the dimension-shifting isomorphisms.

3.2 Tor and the Failure of Exactness

3.2.1 Definition and Computation of Tor

Again let R be commutative. Tensoring the injection $2 : \mathbb{Z} \to \mathbb{Z}$ with $\mathbb{Z}/2\mathbb{Z}$ gives the zero map $\mathbb{Z}/2\mathbb{Z} \to \mathbb{Z}/2\mathbb{Z}$, which is not injective. Tensor products preserve cokernels but can fail to preserve kernels.

Definition 3.2.1. Choose a projective resolution $P_\bullet \to M$. Define

$$\mathrm{Tor}_i^R(M, N) = \mathrm{H}_i(P_\bullet \otimes_R N) = \frac{\ker(\partial_i \otimes \mathrm{id}_N)}{\mathrm{im}(\partial_{i+1} \otimes \mathrm{id}_N)} \quad (i \geq 0),$$

where the differential out of degree zero is zero. These are R-modules. Right exactness gives

$$\mathrm{Tor}_0^R(M, N) \cong M \otimes_R N.$$

The comparison theorem makes these groups independent of the resolution up to canonical isomorphism: a chain homotopy remains a chain homotopy after tensoring. Lifting maps of M gives functoriality in the first variable; tensoring maps of N gives functoriality in the second. Thus Tor is covariant in both variables.

For a free resolution $P_i = \bigoplus_{\lambda \in \Lambda_i} R$, the computation is completely explicit:

$$\cdots \longrightarrow \bigoplus_{\lambda \in \Lambda_2} N \xrightarrow{(\partial_2)_N} \bigoplus_{\lambda \in \Lambda_1} N \xrightarrow{(\partial_1)_N} \bigoplus_{\lambda \in \Lambda_0} N \longrightarrow 0.$$

Use the same coefficient matrices as in the free resolution, now acting on direct sums of N; take kernels modulo images. A presentation by itself computes $\mathrm{Tor}_0^R(M, N)$, while $\mathrm{Tor}_1^R(M, N)$ generally requires the next map.

Example 3.2.2. Suppose multiplication by $a \in R$ is injective on R. Then $0 \to R \xrightarrow{a} R \to R/(a) \to 0$ is a free resolution. Writing $N[a] = \{n \in N : na = 0\}$, we obtain

$$\mathrm{Tor}_0^R(R/(a), N) = N/Na, \qquad \mathrm{Tor}_1^R(R/(a), N) = N[a], \qquad \mathrm{Tor}_i^R(R/(a), N) = 0 \quad (i \geq 2).$$

In particular, for $a, b > 0$, $\mathrm{Tor}_1^{\mathbb{Z}}(\mathbb{Z}/a\mathbb{Z}, \mathbb{Z}/b\mathbb{Z}) \cong \mathbb{Z}/\gcd(a, b)\mathbb{Z}$. The kernel description is natural; an identification with a standard cyclic group involves a choice of generator.

Example 3.2.3. For the matrix in Lecture 2, Example 2.3.8, $\det \phi = 12 \neq 0$, so $\phi : \mathbb{Z}^2 \to \mathbb{Z}^2$ is injective. It therefore gives a length-one free resolution of M. Modulo four, $\phi : (\mathbb{Z}/4\mathbb{Z})^2 \to (\mathbb{Z}/4\mathbb{Z})^2$ acts by $(a, b) \mapsto (2a, 2b)$, whose kernel is $\{0, 2\}^2$. Thus

$$\mathrm{Tor}_1^{\mathbb{Z}}(M, \mathbb{Z}/4\mathbb{Z}) \cong (\mathbb{Z}/2\mathbb{Z})^2, \qquad \mathrm{Tor}_i^{\mathbb{Z}}(M, \mathbb{Z}/4\mathbb{Z}) = 0 \quad (i \geq 2).$$

Indeed, $\phi(a,b) = 0$ over $\mathbb{Z}$ implies $2a = 0$, hence $a = 0$, and then $6b = 0$, hence $b = 0$. The tensor complex is $0 \to (\mathbb{Z}/4\mathbb{Z})^2 \xrightarrow{(2a,2b)} (\mathbb{Z}/4\mathbb{Z})^2 \to 0$, in degrees one and zero. A degree-one cycle has each coordinate either zero or two. There is no degree-two term and hence no nonzero degree-one boundary. The map $(\mathbb{Z}/2\mathbb{Z})^2 \to (\mathbb{Z}/4\mathbb{Z})^2$ given by $(\bar{a},\bar{b}) \mapsto (2a,2b)$ identifies $(\mathbb{Z}/2\mathbb{Z})^2$ with this kernel. In degree zero the cokernel was computed in Lecture 2, Example 2.3.8.

Example 3.2.4. Let $R = k[\varepsilon]/(\varepsilon^2)$ for a field k, and let $K = R/(\varepsilon)$. Since the kernel and image of multiplication by ε are both (ε),

$$\cdots \xrightarrow{\varepsilon} R \xrightarrow{\varepsilon} R \xrightarrow{\varepsilon} R \longrightarrow K \longrightarrow 0$$

is a free resolution. To check this directly, write an element of R uniquely as $a + b\varepsilon$. Multiplication by ε sends it to $a\varepsilon$, so its kernel and image are both $k\varepsilon$. The augmentation has the same kernel. Since $K\varepsilon = 0$, tensoring gives one copy of k in each nonnegative degree and zero differential. Every vector is a cycle and the only boundary is zero, so $\operatorname{Tor}_i^R(K,K) \cong k$ for every $i \geq 0$. The non-zero-divisor hypothesis in Example 3.2.2 cannot be dropped.

3.2.2 Flatness and Long Exact Sequences

Definition 3.2.5. A module F is flat if $- \otimes_R F$ is exact. Since tensoring is right exact, this means that it preserves injections. Free modules are flat because tensoring with a free module is a direct sum of copies of the original sequence; direct sums of exact sequences of modules are exact. Direct summands of flat modules are flat, so projective modules are flat.

Theorem 3.2.6. If $P_\bullet \to M$ and $Q_\bullet \to N$ are projective resolutions, there are natural isomorphisms

$$\mathrm{H}_i(P_\bullet \otimes_R N) \cong \mathrm{H}_i(M \otimes_R Q_\bullet).$$

Consequently $\operatorname{Tor}_i^R(M,N) \cong \operatorname{Tor}_i^R(N,M)$ naturally.

Proof. Form the first-quadrant double complex $P_p \otimes_R Q_q$ and its total complex

$$T_n = \bigoplus_{p+q=n} P_p \otimes_R Q_q, \qquad \partial(\alpha \otimes \beta) = \partial_P \alpha \otimes \beta + (-1)^p \alpha \otimes \partial_Q \beta.$$

Write $\partial_h = \partial_P \otimes \mathrm{id}$ and $\partial_v = \mathrm{id} \otimes \partial_Q$, so the total differential on column p is $\partial_h + (-1)^p \partial_v$. The sign makes the two mixed terms in ∂^2 cancel. Since P_p is flat, each augmented column is exact, with degree-zero homology $P_p \otimes_R N$. Similarly, each augmented row is exact, with degree-zero homology $M \otimes_R Q_q$.

Write $\eta : Q_0 \to N$ for the resolution augmentation and $u : T \to P_\bullet \otimes_R N$ for the induced total-complex augmentation. In degree n it is induced by $\mathrm{id}_{P_n} \otimes \eta$ on the summand with $q = 0$ and is zero on the other summands. It is a degreewise surjective chain map. Its kernel is the total complex of the double complex obtained by replacing the bottom term of column p by

$$K_{p,0} = \ker(P_p \otimes_R Q_0 \longrightarrow P_p \otimes_R N)$$

and retaining $K_{p,q} = P_p \otimes_R Q_q$ for $q > 0$. Each resulting column is exact, including at degree zero; the horizontal differential preserves these kernels.

Let ζ be a total cycle of degree n in this kernel, and choose its largest nonzero column index p. Put $q = n - p$. The cycle equation in bidegree $(p, q-1)$ says that the vertical differential of $\zeta_{p,q}$ is zero, since the component in column $p+1$ vanishes. When $q = 0$, the vertical differential out of the bottom term is zero by definition. Column exactness gives $\omega \in K_{p,q+1}$ with

$$\partial_v \omega = (-1)^p \zeta_{p,q}.$$

The total boundary of ω has component $\zeta_{p,q}$ in column p and its only other possible component in column $p-1$. Subtract it from ζ. Repeating strictly decreases the largest nonzero column index until the cycle is zero. Thus every cycle in $\ker u$ bounds. The long exact sequence of $0 \to \ker u \to T \to P_\bullet \otimes_R N \to 0$ proves that u is a quasi-isomorphism.

For $v : T \to M \otimes_R Q_\bullet$, replace the leftmost term of each row by the kernel of its augmentation. These rows are exact. Eliminate a total cycle starting with its largest row index q, choosing a horizontal primitive. Its total boundary has the prescribed component in that row and its only other possible component in row $q-1$, with the sign already specified in the total differential. The same finite induction proves that $\ker v$ is acyclic. Both augmentations commute with comparison maps of resolutions, so their homology isomorphisms are natural. Taking homology gives the first assertion, and interchanging the tensor factors gives symmetry. □

Theorem 3.2.7. For a short exact sequence $0 \to N' \to N \to N'' \to 0$, there is a natural long exact sequence. For $i \geq 1$, five consecutive terms are

$$\operatorname{Tor}_i^R(M,N') \to \operatorname{Tor}_i^R(M,N) \to \operatorname{Tor}_i^R(M,N'') \xrightarrow{\delta} \operatorname{Tor}_{i-1}^R(M,N') \to \operatorname{Tor}_{i-1}^R(M,N).$$

Using the identification $\operatorname{Tor}_0^R(M,N) \cong M \otimes_R N$, its final five terms are

$$\operatorname{Tor}_1^R(M,N'') \xrightarrow{\delta} M \otimes_R N' \to M \otimes_R N \to M \otimes_R N'' \to 0.$$

There is also a long exact sequence in the first variable. A module F is flat if and only if $\operatorname{Tor}_1^R(M,F) = 0$ for every M; in that case all $\operatorname{Tor}_i^R(M,F)$ vanish for $i > 0$.

Proof. Each term of a projective resolution of M is flat, so tensoring the short exact sequence gives a degreewise short exact sequence of chain complexes. Its homology long exact sequence is the displayed one. Theorem 3.2.6 gives the sequence in the first variable.

If F is flat, tensoring the augmented resolution of M preserves exactness, so all positive homology vanishes. Conversely, apply the long exact sequence in the first variable to $0 \to A \to B \to C \to 0$. Vanishing of $\operatorname{Tor}_1^R(C,F)$ implies that $A \otimes_R F \to B \otimes_R F$ is injective. This is flatness. □

Remark 3.2.8. For a particular short exact sequence, the kernel of $M \otimes_R N' \to M \otimes_R N$ is the image of the connecting map from $\operatorname{Tor}_1^R(M,N'')$. That connecting map need not be injective. Thus the kernel is a quotient of this Tor group, not generally the whole group.

3.3 Cochain Complexes and Ext Groups

3.3.1 Cochain Complexes and Morphism Functors

From now on, R is an arbitrary unital ring and all modules are right R-modules. A cochain complex uses an increasing differential. No commutativity is needed for its definition or for the usual exact-sequence arguments in R-modules.

Definition 3.3.1. A cochain complex $A^\bullet$ consists of maps $\partial_A^n : A^n \to A^{n+1}$ with $\partial_A^{n+1}\partial_A^n = 0$. Its cocycles, coboundaries, and cohomology are

$$Z^n(A) = \ker \partial_A^n, \qquad B^n(A) = \operatorname{im} \partial_A^{n-1}, \qquad \mathrm{H}^n(A) = Z^n(A)/B^n(A).$$

A cochain map $f : A \to B$ satisfies $\partial_B^n f^n = f^{n+1}\partial_A^n$. It induces a map on cohomology. It is a quasi-isomorphism when those maps are isomorphisms for all n.

Reindexing a chain complex by $A^n = A_{-n}$ and $\partial^n = \partial_{-n}$ gives a cochain complex, with $\mathrm{H}^n(A^\bullet) = \mathrm{H}_{-n}(A_\bullet)$. By contrast, applying the contravariant morphism functor to a chain complex reverses its arrows and gives cochains in nonnegative degrees; that is the construction used for Ext.

Definition 3.3.2. A cochain homotopy from g to f is a family $h^n : A^n \to B^{n-1}$ with

$$f^n - g^n = \partial_B^{n-1} h^n + h^{n+1} \partial_A^n.$$

On cocycles the difference is a coboundary, so homotopic maps induce the same cohomology maps. The shift is $A[r]^n = A^{n+r}$ with differential $(-1)^r \partial_A^{n+r}$. In particular, $\mathrm{H}^n(A[r]) = \mathrm{H}^{n+r}(A)$.

For a cochain map $f : A \to B$, its cone is

$$\mathrm{Cone}(f)^n = B^n \oplus A^{n+1}, \qquad \partial^n = \begin{pmatrix} \partial_B^n & f^{n+1} \\ 0 & -\partial_A^{n+1} \end{pmatrix}.$$

The chain-map equation makes this differential square to zero. Inclusion of the first summand and projection onto the second give $0 \to B \to \mathrm{Cone}(f) \to A[1] \to 0$.

Proposition 3.3.3. A degreewise short exact sequence $0 \to A \xrightarrow{\iota} B \xrightarrow{\rho} C \to 0$ of cochain complexes gives a long exact sequence. Five consecutive terms are

$$\mathrm{H}^n(A) \to \mathrm{H}^n(B) \to \mathrm{H}^n(C) \xrightarrow{\delta} \mathrm{H}^{n+1}(A) \to \mathrm{H}^{n+1}(B).$$

For a cocycle $\zeta \in C^n$, choose $\beta \in B^n$ with $\rho\beta = \zeta$ and write $\partial_B \beta = \iota\alpha$. Then $\delta[\zeta] = [\alpha]$.

Proof. The equation $\rho \partial_B \beta = \partial_C \zeta = 0$ gives α, and $\iota \partial_A \alpha = \partial_B^2 \beta = 0$ makes it a cocycle. Changing β by $\iota\alpha'$ changes α by $\partial_A \alpha'$; changing ζ by a coboundary and lifting its primitive does not change the resulting cohomology class. The usual kernel–image chase, or the snake lemma on kernels and cokernels of consecutive differentials, gives exactness. Equivalently, reindex all three complexes as chain complexes and apply the homology long exact sequence; its connecting map is exactly this lift-and-differentiate formula. □

Definition 3.3.4. For a chain complex $P_\bullet$ in nonnegative degrees and a module N, put

$$\mathrm{Mor}_{\mathrm{Mod}_R}(P_\bullet, N)^n = \mathrm{Mor}_{\mathrm{Mod}_R}(P_n, N), \qquad \partial^n(\phi) = \phi \circ \partial_{n+1}.$$

Thus the arrows run as follows:

$$0 \to \mathrm{Mor}_{\mathrm{Mod}_R}(P_0, N) \xrightarrow{\partial^0} \mathrm{Mor}_{\mathrm{Mod}_R}(P_1, N) \xrightarrow{\partial^1} \mathrm{Mor}_{\mathrm{Mod}_R}(P_2, N) \to \cdots.$$

The relation $\partial^{n+1}\partial^n = 0$ follows from $\partial_{n+1}\partial_{n+2} = 0$. For general R, this is a cochain complex of abelian groups, canonically of modules over the center $Z(R)$. It is not generally a complex of right R-modules: multiplying the values of an R-linear map on the right by a noncentral element need not give an R-linear map.

Remark 3.3.5. Write $\mathrm{Ch}^\bullet(\mathrm{Mod}_R)$ for cochain complexes of right R-modules. For cochain complexes A, B, define the mapping complex of abelian groups

$$\underline{\mathrm{Mor}}_{\mathrm{Ch}^\bullet(\mathrm{Mod}_R)}(A, B)^n = \prod_p \mathrm{Mor}_{\mathrm{Mod}_R}(A^p, B^{p+n}), \qquad (\partial f)^p = \partial_B^{p+n} f^p - (-1)^n f^{p+1} \partial_A^p.$$

Its degree-zero cocycles are cochain maps, and its degree-zero coboundaries are nullhomotopic maps. If $A^{-n} = P_n$ and $B = N$ is concentrated in degree zero, this convention gives $\partial\phi = (-1)^{n+1} \phi \partial_{n+1}$. Multiplication in degree n by $(-1)^{n(n+1)/2}$ identifies the unsigned complex in Definition 3.3.4 with this signed convention. We use the unsigned complex for the Ext computations below.

3.3.2 Ext Groups, Variance, and Computation

The ring R remains arbitrary. Projective resolutions of right R-modules exist by Section 3.1.

Definition 3.3.6. For right modules M, N, choose a projective resolution $P_\bullet \to M$ and define

$$\operatorname{Ext}_R^n(M,N) = \mathrm{H}^n(\operatorname{Mor}_{\mathrm{Mod}_R}(P_\bullet, N)) = \frac{\ker(\partial^n)}{\operatorname{im}(\partial^{n-1})} \quad (n \geq 0),$$

where the image in degree zero is zero. These are abelian groups with a natural $Z(R)$-module structure; if R is an algebra over a field k, they are k-vector spaces. When R is commutative, they are naturally R-modules.

The comparison theorem again proves independence of resolutions. More explicitly, if $f, g : P_\bullet \to Q_\bullet$ satisfy $f - g = \partial_Q h + h\partial_P$, precomposition gives cochain-homotopic maps: on $\phi : Q_n \to N$ use the homomorphism $\operatorname{Mor}_{\mathrm{Mod}_R}(Q_n, N) \to \operatorname{Mor}_{\mathrm{Mod}_R}(P_{n-1}, N)$ given by $\phi \mapsto \phi h_{n-1}$. Thus Ext is contravariant in M and covariant in N. In degree zero, a map $P_0 \to N$ lies in $\ker \partial^0$ exactly when it kills $\operatorname{im} \partial_1$, so it factors through M and

$$\operatorname{Ext}_R^0(M,N) \cong \operatorname{Mor}_{\mathrm{Mod}_R}(M,N).$$

For an arbitrary free resolution $P_n = \bigoplus_{\lambda \in \Lambda_n} R$, note the essential change from tensor products:

$$\operatorname{Mor}_{\mathrm{Mod}_R}\Big(\bigoplus_{\lambda \in \Lambda_n} R, N\Big) \cong \prod_{\lambda \in \Lambda_n} N.$$

Maps out of a free direct sum may have arbitrary values on its basis, so this is a product, not a direct sum. If $\partial_{n+1}(e_{\lambda'}) = \sum_{\lambda \in \Lambda_n} e_\lambda r_{\lambda\lambda'}$ for $\lambda' \in \Lambda_{n+1}$, with coefficients on the right, then

$$\partial^n((n_\lambda)_{\lambda \in \Lambda_n}) = \left(\sum_{\lambda \in \Lambda_n} n_\lambda r_{\lambda\lambda'}\right)_{\lambda' \in \Lambda_{n+1}}.$$

Every sum is finite because each column of the differential has finite support. Compute Ext by taking the kernel of this map modulo the image of the preceding map. Over a commutative ring with finite free terms, these are the transposed differential matrices acting on finite powers of N. For a noncommutative ring, the displayed order of the coefficients specifies the maps unambiguously.

Proposition 3.3.7. For $0 \to N' \to N \to N'' \to 0$, there is a natural long exact sequence. For $n \geq 0$, five consecutive terms are

$$\operatorname{Ext}_R^n(M,N') \to \operatorname{Ext}_R^n(M,N) \to \operatorname{Ext}_R^n(M,N'') \xrightarrow{\delta} \operatorname{Ext}_R^{n+1}(M,N') \to \operatorname{Ext}_R^{n+1}(M,N).$$

Recall that $\operatorname{Ext}_R^0(M,N) \cong \operatorname{Mor}_{\mathrm{Mod}_R}(M,N)$. The initial map $\operatorname{Mor}_{\mathrm{Mod}_R}(M,N') \to \operatorname{Mor}_{\mathrm{Mod}_R}(M,N)$ is injective. If P is projective, $\operatorname{Ext}_R^n(P,N) = 0$ for $n > 0$. Conversely, if $\operatorname{Ext}_R^1(M,N) = 0$ for every N, then M is projective.

Proof. Apply $\operatorname{Mor}_{\mathrm{Mod}_R}(P_i, -)$ to the short exact sequence. Each resulting sequence is exact by projectivity, so Proposition 3.3.3, applied in abelian groups, gives the displayed sequence. A projective P has a resolution concentrated in degree zero, proving its higher vanishing. For the converse, take a free surjection $F \to M$ with kernel K. The exact sequence makes $\operatorname{Mor}_{\mathrm{Mod}_R}(M,F) \to \operatorname{Mor}_{\mathrm{Mod}_R}(M,M)$ surjective because $\operatorname{Ext}_R^1(M,K) = 0$. A preimage of id_M splits $F \to M$, so M is projective. □

There is also a long exact sequence in the first, contravariant variable. For $0 \to M' \to M \to M'' \to 0$ and $n \geq 0$, five consecutive terms are

$$\operatorname{Ext}^n_R(M'', N) \to \operatorname{Ext}^n_R(M, N) \to \operatorname{Ext}^n_R(M', N) \xrightarrow{\delta} \operatorname{Ext}^{n+1}_R(M'', N) \to \operatorname{Ext}^{n+1}_R(M, N).$$

Under the same degree-zero identification, the initial map $\operatorname{Mor}_{\operatorname{Mod}_R}(M'', N) \to \operatorname{Mor}_{\operatorname{Mod}_R}(M, N)$ is injective. Apply Lemma 3.1.15 to obtain compatible resolutions with $P_i = P'_i \oplus P''_i$. Applying $\operatorname{Mor}_{\operatorname{Mod}_R}(-, N)$ gives a degreewise split short exact sequence of cochain complexes in the reverse order. Proposition 3.3.3 gives the displayed sequence, and the compatible comparison maps in the horseshoe proof give its naturality.

3.3.3 Extensions as Degree-One Classes

Theorem 3.3.8. The group $\operatorname{Ext}^1_R(M, N)$ classifies short exact sequences of right R-modules

$$0 \longrightarrow N \xrightarrow{\iota} E \xrightarrow{\rho} M \longrightarrow 0 \tag{10}$$

up to isomorphisms of sequences that are the identity on N and M. The zero class consists precisely of the split extensions.

Proof. Choose $0 \to K \xrightarrow{\jmath} P_0 \to M \to 0$ with P_0 projective. The definition of Ext from a resolution identifies

$$\operatorname{Ext}^1_R(M, N) \cong \operatorname{coker}(\operatorname{Mor}_{\operatorname{Mod}_R}(P_0, N) \longrightarrow \operatorname{Mor}_{\operatorname{Mod}_R}(K, N)). \tag{11}$$

Indeed, a map $P_1 \to N$ is a cocycle exactly when it kills $\operatorname{im}(\partial_2)$, so it factors through $P_1/\operatorname{im}(\partial_2) \cong K$.

Given $\alpha : K \to N$, form the quotient

$$E_\alpha = (N \oplus P_0)/\operatorname{im}(-\alpha, \jmath).$$

The construction fits into the commutative diagram with exact rows and columns

$$\begin{array}{ccccccccc}
0 & \longrightarrow & 0 & \longrightarrow & K & = & K & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow{\scriptstyle (-\alpha,\jmath)} & & \downarrow{\scriptstyle \jmath} & & \\
0 & \longrightarrow & N & \xrightarrow{\iota} & N \oplus P_0 & \xrightarrow{\mathrm{pr}_2} & P_0 & \longrightarrow & 0 \\
 & & \| & & \downarrow{\scriptstyle q} & & \downarrow{\scriptstyle \epsilon} & & \\
0 & \longrightarrow & N & \xrightarrow{\iota_\alpha} & E_\alpha & \xrightarrow{\pi_\alpha} & M & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
 & & 0 & & 0 & & 0. & &
\end{array}$$

The middle column exhibits $E_\alpha = \operatorname{coker}(K \xrightarrow{(-\alpha,\jmath)} N \oplus P_0)$. Here $\iota : N \to N \oplus P_0$ satisfies $\iota(a) = (a, 0)$; $q(a, p) = [(a, p)]$, $\iota_\alpha(a) = [(a, 0)]$, and $\pi_\alpha([(a, p)]) = \epsilon(p)$. Thus the trivial extension of P_0 by N maps surjectively onto the desired extension of M by N. The map ι_α is injective because $\jmath$ is injective. If $\epsilon(p) = 0$, write $p = \jmath(k)$; then

$$[(a, p)] = [(a + \alpha(k), 0)],$$

which proves exactness at E_α.

Conversely, given an extension $0 \to N \xrightarrow{\iota} E \to M \to 0$, projectivity gives a lift $s : P_0 \to E$ of ϵ. There is a unique $\alpha : K \to N$ with $s\jmath = \iota\alpha$. The surjection

$$N \oplus P_0 \longrightarrow E, \qquad (a, p) \longmapsto \iota(a) + s(p)$$

has kernel $\operatorname{im}(-\alpha, \jmath)$, so it identifies the given extension with the bottom row of the diagram.

Changing the lift replaces α by $\alpha + \beta\jmath$ for some $\beta : P_0 \to N$. Conversely, if $\alpha' = \alpha + \beta\jmath$, the map $N \oplus P_0 \to N \oplus P_0$ given by $(a, p) \mapsto (a - \beta(p), p)$ descends to an equivalence $E_\alpha \cong E_{\alpha'}$. An equivalence of extensions also compares their lifts in exactly this way, so the equivalence classes are precisely the displayed cokernel. Finally, $E_0 \cong N \oplus M$, and a split extension admits a lift with $\alpha = 0$. Hence the zero class corresponds precisely to split extensions. □

The addition transported by this classification is the Baer sum: take the direct sum of two extensions, pull back along the diagonal $M \to M \oplus M$, and push out along addition $N \oplus N \to N$. In the description above, these operations replace the two maps $\alpha, \beta : K \to N$ by $\alpha + \beta$.

3.4 Computations of Ext

3.4.1 An Explicit Noncommutative Ring

Example 3.4.1. Let k be a field and

$$R = \left\{ \begin{pmatrix} a & b \\ 0 & c \end{pmatrix} : a, b, c \in k \right\}, \qquad e_1 = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad e_2 = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, \quad u = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}.$$

This ring is noncommutative: $e_1 u = u$ whereas $u e_1 = 0$. Define right modules $S_1 = S_2 = k$ as vector spaces, with actions

$$\mu \cdot \begin{pmatrix} a & b \\ 0 & c \end{pmatrix} = \mu a \quad \text{on } S_1, \qquad \mu \cdot \begin{pmatrix} a & b \\ 0 & c \end{pmatrix} = \mu c \quad \text{on } S_2.$$

The right ideals $P_1 = e_1 R$ and $P_2 = e_2 R$ are projective because $R = e_1 R \oplus e_2 R$. They have bases e_1, u and e_2, respectively. There is a projective resolution

$$0 \longrightarrow P_2 \xrightarrow{\partial} P_1 \xrightarrow{\rho} S_1 \longrightarrow 0, \qquad \partial(\zeta) = u\zeta, \qquad \rho(a e_1 + b u) = a. \tag{12}$$

Left multiplication by u is a map of right modules. It sends e_2 to u, so it is injective with image $ku = \ker \rho$.

For any right R-module N, evaluation at e_i gives

$$\operatorname{Mor}_{\operatorname{Mod}_R}(e_i R, N) \cong N e_i.$$

Indeed, $f(e_i) = f(e_i)e_i$, and conversely $n \in N e_i$ determines $f_n(e_i r) = nr$. This formula is well defined: if $e_i r = e_i r'$, then $n(r - r') = n e_i (r - r') = 0$. Under these identifications, precomposition with ∂ is

$$N e_1 \longrightarrow N e_2, \qquad n \longmapsto nu,$$

because $\partial(e_2) = u = e_1 u$. Thus the entire Ext computation is the two-term cochain complex

$$N e_1 \xrightarrow{\cdot u} N e_2 \qquad \text{in degrees } 0, 1.$$

Consequently

$$\begin{aligned}\mathrm{Mor}_{\mathrm{Mod}_R}(S_1, N) &\cong \ker(\cdot u : Ne_1 \to Ne_2),\\ \mathrm{Ext}^1_R(S_1, N) &\cong Ne_2/(Ne_1)u,\\ \mathrm{Ext}^n_R(S_1, N) &= 0 \qquad (n \geq 2).\end{aligned}$$

For $N = S_2$, the actions give $S_2e_1 = 0$ and $S_2e_2 = k$. Thus the cochain complex is $0 \to 0 \to k \to 0$, in degrees zero and one. Its degree-zero cohomology is zero, and its degree-one cohomology is $k/0 = k$. Hence

$$\mathrm{Ext}^1_R(S_1, S_2) \cong k.$$

For $N = S_1$, the complex instead is $0 \to k \to 0 \to 0$, so $\mathrm{Mor}_{\mathrm{Mod}_R}(S_1, S_1) = k$ and $\mathrm{Ext}^1_R(S_1, S_1) = 0$. For $N = P_1$, we have $P_1e_1 = ke_1$ and $P_1e_2 = ku$; the differential $P_1e_1 \to P_1e_2$ given by $ae_1 \mapsto au$ is an isomorphism, so both cohomology groups vanish. By contrast, $S_2 \cong P_2$ is projective, so $\mathrm{Ext}^1_R(S_2, S_1) = 0$. Unlike Tor over a commutative ring, Ext is not symmetric in its variables.

The resolution (12) itself, after identifying P_2 with S_2, is a nonsplit extension of S_1 by S_2. If it had a section, the image v of $1 \in S_1$ would satisfy $\rho(v) = 1$, so $v = e_1 + bu$ for some $b \in k$. But u acts trivially on S_1, while $vu = u \neq 0$, contradicting linearity of the section.

More explicitly, the class $\lambda \in k$ corresponds to $E_\lambda = k\xi \oplus k\eta$ with

$$\xi \begin{pmatrix} a & b \\ 0 & c \end{pmatrix} = c\xi, \qquad \eta \begin{pmatrix} a & b \\ 0 & c \end{pmatrix} = a\eta + b\lambda\xi.$$

Define $\iota : S_2 \to E_\lambda$ by $\iota(1) = \xi$ and $\rho : E_\lambda \to S_1$ by $\rho(\xi) = 0$, $\rho(\eta) = 1$. The sequence $0 \to S_2 \xrightarrow{\iota} E_\lambda \xrightarrow{\rho} S_1 \to 0$ has class λ under (11): the lift $P_1 \to E_\lambda$ determined by $e_1 \mapsto \eta$ sends $\partial(e_2) = u$ to $\eta u = \lambda\xi$. It splits exactly when $\lambda = 0$; $E_1 \cong P_1$ via $\xi \mapsto u$ and $\eta \mapsto e_1$. Nonzero values can give isomorphic middle modules while representing different extension classes, because equivalence of extensions must fix both end modules.

3.4.2 Kernels and Cokernels for Cyclic Groups

Example 3.4.2. The free resolution $0 \to \mathbb{Z} \xrightarrow{12} \mathbb{Z} \to \mathbb{Z}/12\mathbb{Z} \to 0$ is exact because multiplication by 12 is injective and its cokernel is $\mathbb{Z}/12\mathbb{Z}$. Set $N = \mathbb{Z}/8\mathbb{Z}$. Tensoring gives

$$0 \longrightarrow N \xrightarrow{4} N \longrightarrow 0 \quad \text{in chain degrees } 1, 0.$$

The image is $\{0, 4\}$, and the kernel is $\{0, 2, 4, 6\}$. The quotient map $N \to \mathbb{Z}/4\mathbb{Z}$ has kernel $\{0, 4\}$; hence $\mathbb{Z}/12\mathbb{Z} \otimes_{\mathbb{Z}} N \cong \mathbb{Z}/4\mathbb{Z}$. The map $\mathbb{Z}/4\mathbb{Z} \to N$, $[t]_4 \mapsto [2t]_8$, is injective with image $\{0, 2, 4, 6\}$. There is no incoming degree-two boundary, so it identifies $\mathrm{Tor}^{\mathbb{Z}}_1(\mathbb{Z}/12\mathbb{Z}, N)$ with $\mathbb{Z}/4\mathbb{Z}$. Higher Tor vanishes.

Applying $\mathrm{Mor}_{\mathrm{Ab}}(-, N)$ reverses the arrows and gives

$$0 \longrightarrow N \xrightarrow{4} N \longrightarrow 0 \quad \text{in cochain degrees } 0, 1.$$

Thus $\mathrm{Mor}_{\mathrm{Ab}}(\mathbb{Z}/12\mathbb{Z}, N) = \{0, 2, 4, 6\}$, while $\mathrm{Ext}^1_{\mathbb{Z}}(\mathbb{Z}/12\mathbb{Z}, N) = N/\{0, 4\} \cong \mathbb{Z}/4\mathbb{Z}$. Higher Ext vanishes. The same map is being used, but homological and cohomological grading put its kernel and cokernel in different degrees.

For general $a, b > 0$, put $g = \gcd(a, b)$, $a = ga'$, and $b = gb'$. The image of multiplication by a on $\mathbb{Z}/b\mathbb{Z}$ equals the subgroup generated by g, by Bézout's identity. Its cokernel is therefore $\mathbb{Z}/g\mathbb{Z}$. For the kernel, $b \mid ac$ is equivalent to $b' \mid c$, since $\gcd(a', b') = 1$. Its elements are the g multiples

of b' modulo b. The map $\mathbb{Z}/g\mathbb{Z} \to \mathbb{Z}/b\mathbb{Z}$ given by $[t]_g \mapsto [b't]_b$ gives its cyclic-group identification. This computes tensor, Tor, morphism groups, and Ext in the preceding manner.

Example 3.4.3. For $M = \operatorname{coker}\left(\begin{smallmatrix} 2 & 0 \\ 4 & 6 \end{smallmatrix}\right)$ and $N = \mathbb{Z}/4\mathbb{Z}$, a homomorphism $\mathbb{Z}^2 \to N$ is specified by a row (a, b). Precomposition with ϕ sends it to $(2a+4b, 6b)$. Equivalently, on column coordinates the cochain differential is ϕ^{T}. Modulo four this is the map $N^2 \to N^2$ given by $(a, b) \mapsto (2a, 2b)$. Its kernel is $\{0, 2\}^2$, and its image is also $\{0, 2\}^2$ in the next copy of N^2. Consequently

$$\operatorname{Mor}_{\mathrm{Ab}}(M, N) \cong (\mathbb{Z}/2\mathbb{Z})^2, \qquad \operatorname{Ext}^1_{\mathbb{Z}}(M, N) \cong (\mathbb{Z}/2\mathbb{Z})^2,$$

and all higher groups vanish. The map defining Ext is the transposed map, even though in this particular example reduction modulo four makes it coincide with the map used for Tor.

3.4.3 Computing Ext with Injective Resolutions

For a fixed right R-module M, the functor $\operatorname{Mor}_{\mathrm{Mod}_R}(M, -)$ is left exact. Its right derived functors are the Ext groups defined above.

Proposition 3.4.4. If $P_\bullet \to M$ is a projective resolution and $0 \to N \to I^\bullet$ is an injective resolution, there are natural isomorphisms

$$\mathrm{H}^n \operatorname{Mor}_{\mathrm{Mod}_R}(P_\bullet, N) \cong \mathrm{H}^n \operatorname{Mor}_{\mathrm{Mod}_R}(M, I^\bullet) = \mathrm{R}^n \operatorname{Mor}_{\mathrm{Mod}_R}(M, -)(N).$$

Proof. Use the first-quadrant double complex $C^{p,q} = \operatorname{Mor}_{\mathrm{Mod}_R}(P_p, I^q)$. Its horizontal differential ∂_h is precomposition with ∂_{p+1}, and its vertical differential ∂_v is postcomposition with the differential of I. These maps commute; on the total complex take $\partial = \partial_h + (-1)^p \partial_v$. Thus $\partial^2 = 0$. The total degree-n group is $\bigoplus_{p+q=n} C^{p,q}$, a finite sum.

For fixed p, exactness of $\operatorname{Mor}_{\mathrm{Mod}_R}(P_p, -)$ makes the augmented column $0 \to \operatorname{Mor}_{\mathrm{Mod}_R}(P_p, N) \to C^{p,0} \to C^{p,1} \to \cdots$ exact. For fixed q, injectivity of I^q makes the augmented row $0 \to \operatorname{Mor}_{\mathrm{Mod}_R}(M, I^q) \to C^{0,q} \to C^{1,q} \to \cdots$ exact. Both augmentations are injective cochain maps into the total complex: the first occupies the row $q = 0$, the second the column $p = 0$. Consider the quotient by the first. It is the total complex of the double complex with

$$D^{p,0} = C^{p,0} / \operatorname{Mor}_{\mathrm{Mod}_R}(P_p, N), \qquad D^{p,q} = C^{p,q} \quad (q > 0).$$

Its columns are exact, including injectivity of the differential out of degree zero. The horizontal differential descends to these quotients.

For a total cocycle ζ of degree n in this quotient, choose its largest nonzero vertical degree q, and put $p = n - q$. Its component in bidegree $(p, q+1)$ under the total differential is $(-1)^p \partial_v \zeta^{p,q}$: the other possible contribution has vertical degree $q + 1$ before differentiation and vanishes. Hence $\partial_v \zeta^{p,q} = 0$. If $q > 0$, choose $\omega \in D^{p,q-1}$ with

$$\partial_v \omega = (-1)^p \zeta^{p,q}.$$

Subtracting $\partial\omega$ removes that component and introduces a possible component only in vertical degree $q - 1$. Repeat. At $q = 0$, injectivity of the bottom vertical differential makes the remaining component zero. Thus the quotient total complex is acyclic. The cohomology long exact sequence proves that the first augmentation is a quasi-isomorphism.

For the second augmentation take the quotient of $C^{0,q}$ by $\operatorname{Mor}_{\mathrm{Mod}_R}(M, I^q)$ instead. Its rows are exact. Starting with the largest horizontal degree p, choose a horizontal primitive in degree $p - 1$; subtracting its total coboundary introduces a component only in horizontal degree $p - 1$. At $p = 0$ the bottom horizontal differential is injective. This proves the other quotient acyclic and the other augmentation a quasi-isomorphism. The two augmentation

diagrams commute with comparison maps of resolutions. Their homology isomorphisms, and therefore the resulting comparison, are natural. □

Example 3.4.5. To compute $\operatorname{Ext}^n_{\mathbb{Z}}(\mathbb{Z}/m\mathbb{Z}, \mathbb{Z})$ for $m \geq 2$, use $0 \to \mathbb{Z} \to \mathbb{Q} \to \mathbb{Q}/\mathbb{Z} \to 0$ from Example 3.1.6. A homomorphism $\mathbb{Z}/m\mathbb{Z} \to \mathbb{Q}$ must send 1 to an element killed by m; the only such rational number is zero. A homomorphism to $\mathbb{Q}/\mathbb{Z}$ is instead determined by one of

$$0, \quad \tfrac{1}{m} + \mathbb{Z}, \quad \tfrac{2}{m} + \mathbb{Z}, \quad \dots, \quad \tfrac{m-1}{m} + \mathbb{Z}.$$

The resulting cochain complex is $0 \to 0 \to (\frac{1}{m}\mathbb{Z})/\mathbb{Z} \to 0$, with its two terms in degrees zero and one. Hence

$$\operatorname{Mor}_{\mathrm{Ab}}(\mathbb{Z}/m\mathbb{Z}, \mathbb{Z}) = 0, \qquad \operatorname{Ext}^1_{\mathbb{Z}}(\mathbb{Z}/m\mathbb{Z}, \mathbb{Z}) \cong \mathbb{Z}/m\mathbb{Z},$$

and higher Ext vanishes. The projective-resolution calculation gives $0 \to \mathbb{Z} \xrightarrow{m} \mathbb{Z} \to 0$, with the same kernel zero and cokernel $\mathbb{Z}/m\mathbb{Z}$. This checks the comparison by two distinct resolutions.

3.4.4 Comparing Tor and Ext Computations

For comparison with Section 3.2, suppose R is commutative and a is a non-zero-divisor. The same resolution used there now gives

$$0 \longrightarrow N \xrightarrow{a} N \longrightarrow 0 \qquad \text{in cochain degrees } 0, 1.$$

Hence $\operatorname{Mor}_{\mathrm{Mod}_R}(R/(a), N) = N[a]$, $\operatorname{Ext}^1_R(R/(a), N) = N/Na$, and higher Ext vanishes. The kernel and cokernel appear in the opposite degrees from the tensor computation. For the dual-number example, applying $\operatorname{Mor}_{\mathrm{Mod}_R}(-, K)$ makes every differential zero, so $\operatorname{Ext}^n_R(K, K) \cong k$ for every $n \geq 0$.

The two procedures can be compared directly:

	Tensor and Tor, R commutative	Morphism Groups and Ext, R arbitrary
Resolution	$P_\bullet \twoheadrightarrow M$	$P_\bullet \twoheadrightarrow M$
Complex	$P_\bullet \otimes_R N$	$\operatorname{Mor}_{\mathrm{Mod}_R}(P_\bullet, N)$
Free term $\bigoplus_{\lambda\in\Lambda} R$	$\bigoplus_{\lambda\in\Lambda} N$	$\prod_\Lambda N$
Degree zero	$M \otimes_R N$	$\operatorname{Mor}_{\mathrm{Mod}_R}(M, N)$
Higher degrees	$\mathrm{H}_i = \operatorname{Tor}^R_i(M, N)$	$\mathrm{H}^i = \operatorname{Ext}^i_R(M, N)$
Variance	covariant in both	contravariant, then covariant

In either computation, the central work is to construct enough of a resolution and then compute a kernel modulo an image. A single presentation suffices for a tensor product; its higher relations control Tor and Ext.

For further reading on tensor products and Tor, see [Stacks, Tensor Products, Tag 00CV; Tor Groups and Flatness, Tag 00LY].

Simplicial Sets, Spaces, and Singular Chains

Algebraic Topology I | Fall 2026 | September 21 and 23

4.1 Simplicial Sets and Their Chains

In this section, R is commutative. Write Mod_R for the category of right R-modules; the usual identification $mr = rm$ identifies left and right modules over a commutative ring. A simplicial set records simplices and the ways they meet. Its simplices are initially elements of sets, so there is no addition and no alternating sum of faces. Passing to free modules introduces the algebra needed for a chain complex.

4.1.1 Simplicial Sets and Geometric Constructions

Definition 4.1.1. The simplex category Δ has objects $[n] = \{0 < 1 < \cdots < n\}$ for $n \geq 0$, and order-preserving maps as morphisms. A simplicial set is a functor $X : \Delta^{\mathrm{op}} \to \mathrm{Set}$. Write $X_n = X([n])$. A map of simplicial sets is a natural transformation, and their category is $\mathcal{P}(\Delta)$.

Definition 4.1.2. The injection $[n-1] \to [n]$ omitting i induces a face map $d_i : X_n \to X_{n-1}$. The surjection $[n+1] \to [n]$ identifying i and $i+1$ induces a degeneracy map $s_i : X_n \to X_{n+1}$. These satisfy the simplicial identities

$$
\begin{aligned}
d_i d_j &= d_{j-1} d_i && (i < j),\\
s_i s_j &= s_{j+1} s_i && (i \leq j),\\
d_i s_j &= \begin{cases} s_{j-1} d_i, & i < j,\\ \mathrm{id}, & i = j \text{ or } i = j+1,\\ s_j d_{i-1}, & i > j+1. \end{cases}
\end{aligned}
$$

They express the equalities among order-preserving maps in Δ. A simplex in the image of a degeneracy is degenerate; otherwise it is nondegenerate.

Example 4.1.3. The standard simplicial n-simplex is the representable functor

$$\Delta^n_k = \mathrm{Mor}_{\Delta}([k], [n]).$$

Its boundary $\partial\Delta^n$ consists of the maps whose images miss at least one vertex. The quotient $\Delta^1/\partial\Delta^1$ is a simplicial circle: it has one nondegenerate vertex v and one nondegenerate edge e, with $d_0 e = d_1 e = v$, and no nondegenerate simplices in higher degrees. It still has degenerate simplices in every degree.

Definition 4.1.4. For a topological space Y, let

$$|\Delta^n| = \{(t_0, \dots, t_n) \in \mathbb{R}^{n+1} : t_i \geq 0,\ \textstyle\sum_i t_i = 1\}.$$

The singular simplicial set $\mathrm{Sing}(Y)$ has continuous maps $|\Delta^n| \to Y$ as its n-simplices. Its structure maps are precomposition with the affine maps between topological simplices induced by maps in Δ.

For a simplicial set X, its geometric realization is

$$|X| = \left(\coprod_{n \geq 0} X_n \times |\Delta^n| \right) / \sim,$$

where $(X(\alpha)(x), t) \sim (x, |\alpha|(t))$ for $\alpha : [m] \to [n]$, $x \in X_n$, and $t \in |\Delta^m|$. Each X_n is given the discrete topology, and the quotient has the quotient topology.

The equivalence relation glues faces and collapses degenerate simplices. The definitions give the natural adjunction

$$\mathrm{Mor}_{\mathrm{Top}}(|X|, Y) \cong \mathrm{Mor}_{\mathcal{P}(\Delta)}(X, \mathrm{Sing}(Y)) :$$

a continuous map out of $|X|$ is exactly a compatible family of maps from its topological simplices. In Example 4.1.3, realization identifies the two endpoints of an interval, giving the usual circle.

4.1.2 Free Simplicial Modules and the Boundary

Definition 4.1.5. A simplicial R-module is a functor $A_\bullet : \Delta^{\mathrm{op}} \to \mathrm{Mod}_R$. Its associated chain complex has $\mathrm{C}_n(A_\bullet) = A_n$ and differential

$$\partial_n = \sum_{i=0}^{n} (-1)^i d_i \quad (n > 0), \qquad \partial_0 = 0.$$

For a simplicial set X, let $R[X_n]$ be the free R-module on X_n. Extending the structure maps linearly gives a simplicial module $R[X]$, and hence a chain complex $\mathrm{C}_\bullet(X; R) = \mathrm{C}_\bullet(R[X])$.

Proposition 4.1.6. For a simplicial module, the alternating face differential satisfies $\partial_{n-1}\partial_n = 0$.

Proof. In the expanded double sum, pair the term $(-1)^{i+j} d_i d_j$ with $i < j$ with the term $(-1)^{i+j-1} d_{j-1} d_i$. The two composites agree by the face identity and have opposite signs. Every term belongs to one such pair. □

In dimension two this is the familiar calculation

$$\partial[0,1,2] = [1,2] - [0,2] + [0,1], \qquad \partial^2[0,1,2] = 0.$$

Taking $X = \mathrm{Sing}(Y)$ gives the usual singular chain complex, since its basis in degree n is precisely the set of singular n-simplices of Y.

4.1.3 Simplicial Chain Complexes in Low Degrees

For a simplicial abelian group $A_\bullet$, the construction is simply

$$\cdots \longrightarrow A_3 \xrightarrow{d_0 - d_1 + d_2 - d_3} A_2 \xrightarrow{d_0 - d_1 + d_2} A_1 \xrightarrow{d_0 - d_1} A_0 \longrightarrow 0.$$

All simplices are retained. The degeneracy maps do not occur in this formula for the differential, though they remain part of the simplicial structure.

It is useful to check the first nontrivial composition in full:

$$\begin{aligned}(d_0 - d_1)(d_0 - d_1 + d_2) &= d_0 d_0 - d_0 d_1 + d_0 d_2 \\ &\quad - d_1 d_0 + d_1 d_1 - d_1 d_2 = 0.\end{aligned}$$

The cancellations use $d_0 d_0 = d_0 d_1$, $d_0 d_2 = d_1 d_0$, and $d_1 d_1 = d_1 d_2$. Proposition 4.1.6 is the same cancellation in every degree.

A map of simplicial abelian groups $f : A_\bullet \to B_\bullet$ is a family of homomorphisms commuting with all face and degeneracy maps. In particular,

$$f_{n-1}\partial_n = \sum_i (-1)^i f_{n-1} d_i = \sum_i (-1)^i d_i f_n = \partial_n f_n.$$

Thus the associated chain complex construction is a functor: simplicial maps become chain maps, and hence induce maps on homology.

Example 4.1.7. Let G be an abelian group. Set $A_n = G$ in every degree and take every face and degeneracy map to be the identity. For $n > 0$ the differential is

$$\partial_n = \left(\sum_{i=0}^{n} (-1)^i\right) \mathrm{id}_G = \begin{cases} \mathrm{id}_G, & n \text{ even}, \\ 0, & n \text{ odd}. \end{cases}$$

The associated chain complex is therefore

$$\cdots \xrightarrow{0} G \xrightarrow{\mathrm{id}} G \xrightarrow{0} G \xrightarrow{\mathrm{id}} G \xrightarrow{0} G \longrightarrow 0,$$

where the rightmost G is in degree zero. It has $\mathrm{H}_0 = G$ and $\mathrm{H}_n = 0$ for every $n > 0$. A constant simplicial group is not sent to a chain complex with G only in degree zero; it is sent to a larger complex with exactly that homology.

Example 4.1.8. For a simplicial set X, take $A_n = \mathbb{Z}[X_n]$. An edge e has boundary $d_0 e - d_1 e$, its terminal vertex minus its initial vertex. Thus

$$\mathrm{H}_0(\mathrm{C}_\bullet(X;\mathbb{Z})) = \mathbb{Z}[X_0]/\langle d_0 e - d_1 e : e \in X_1 \rangle.$$

This is the free abelian group on the equivalence classes of vertices under the equivalence relation generated by the edges. These are the components of the simplicial set.

A triangle t has boundary $d_0 t - d_1 t + d_2 t$. Hence $\mathrm{H}_1(X;\mathbb{Z})$ is the group of integral sums of edges with zero total boundary, modulo boundaries of sums of triangles. In the simplicial circle of Example 4.1.3, the edge e is a cycle because both of its endpoints are v. Degenerate simplices are retained in all of these chain groups as well.

4.2 Presheaves, Simplices, and Simplicial Sets

The simplicial constructions of Section 4.1 are instances of functor categories. We now describe their natural transformations and universal properties.

4.2.1 Presheaves and the Yoneda Lemma

Definition 4.2.1. Let $\mathcal{D}$ be a small category, regarded as a diagram shape, and let $\mathcal{C}$ be a locally small category. The category of $\mathcal{C}$-valued presheaves on $\mathcal{D}$ is

$$\mathcal{P}(\mathcal{D};\mathcal{C}) = \mathrm{Fun}(\mathcal{D}^{\mathrm{op}}, \mathcal{C}).$$

Its objects are functors $F : \mathcal{D}^{\mathrm{op}} \to \mathcal{C}$, equivalently contravariant diagrams on $\mathcal{D}$. Thus F assigns an object $F(d)$ to each $d \in \mathcal{D}$ and a restriction map $f^* = F(f) : F(d') \to F(d)$ to each $f : d \to d'$, with

$$\mathrm{id}_d^* = \mathrm{id}_{F(d)}, \qquad (g \circ f)^* = f^* \circ g^*.$$

Its morphisms are natural transformations: a morphism $\eta : F \to G$ is a family $\eta_d : F(d) \to G(d)$ satisfying

$$G(f) \circ \eta_{d'} = \eta_d \circ F(f) \qquad \text{for every } f : d \to d'.$$

Identities and composition are componentwise: $(\mathrm{id}_F)_d = \mathrm{id}_{F(d)}$ and $(\theta \circ \eta)_d = \theta_d \circ \eta_d$.

The first argument of $\mathcal{P}(\mathcal{D};\mathcal{C})$ specifies the diagram shape; the second specifies the category of values. The opposite category $\mathcal{D}^{\mathrm{op}}$ reverses arrows. Here small means that objects and morphisms form sets; locally small means that morphisms between any two fixed objects form a set.

More generally, the functor category $\mathrm{Fun}(\mathcal{I},\mathcal{C})$ has covariant functors $\mathcal{I} \to \mathcal{C}$ as objects and natural transformations as morphisms. For $f : i \to j$, naturality reads $G(f)\eta_i = \eta_j F(f)$; composition is again componentwise. The definition above is the case $\mathcal{I} = \mathcal{D}^{\mathrm{op}}$.

Definition 4.2.2. Specializing to $\mathcal{C} = \mathrm{Set}$, write

$$\mathcal{P}(\mathcal{D}) = \mathcal{P}(\mathcal{D}; \mathrm{Set}) = \mathrm{Fun}(\mathcal{D}^{\mathrm{op}}, \mathrm{Set}).$$

Thus an omitted second argument means sets. By contrast, $\mathcal{P}(\mathcal{D}; \mathrm{Ab})$ consists of contravariant diagrams of abelian groups and homomorphisms.

Each object $d \in \mathcal{D}$ defines the representable presheaf $\mathrm{Mor}_{\mathcal{D}}(-, d)$, whose value at e is $\mathrm{Mor}_{\mathcal{D}}(e, d)$ and whose restriction maps are precomposition. Simplicial sets are the case $\mathcal{D} = \Delta$.

Proposition 4.2.3. For $F \in \mathcal{P}(\mathcal{D})$ and $d \in \mathcal{D}$, evaluation at id_d gives a natural bijection

$$\mathrm{Mor}_{\mathcal{P}(\mathcal{D})}(\mathrm{Mor}_{\mathcal{D}}(-, d), F) \cong F(d).$$

Proof. An element $x \in F(d)$ defines $\eta_e(f) = f^* x$ for $f : e \to d$. Functoriality makes η natural, and naturality forces this formula from $x = \eta_d(\mathrm{id}_d)$. These constructions are inverse and natural. □

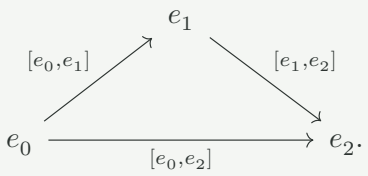
4.2.2 Topological Simplices and Affine Operators

Definition 4.2.4. Use the standard topological simplex of Definition 4.1.4. Its vertices are the coordinate vectors $e_0, \dots, e_n$, and its coordinates t_i are called barycentric coordinates. The realization notation will be justified below.

Thus $|\Delta^0|$ is a point, $|\Delta^1|$ is an interval, and $|\Delta^2|$ is a filled triangle. The ordering of the vertices determines the signs in its oriented boundary:

$$\partial[e_0, e_1, e_2] = [e_1, e_2] - [e_0, e_2] + [e_0, e_1].$$

The triangle can be displayed as



This is an oriented geometric triangle. In a general simplicial set, its edges are not assumed to be morphisms in a category, and the triangle is not an assertion of equality between composites.

Definition 4.2.5. For the simplex category Δ of Definition 4.1.1, a map $\alpha : [m] \to [n]$ determines the affine map

$$\alpha_* : |\Delta^m| \longrightarrow |\Delta^n|, \qquad (\alpha_*(t))_j = \sum_{\alpha(i)=j} t_i.$$

An empty sum is zero. In particular, $\alpha_*(e_i) = e_{\alpha(i)}$.

For $\beta : [\ell] \to [m]$, grouping the coordinates gives $(\alpha\beta)_* = \alpha_* \beta_*$. Thus we obtain a covariant functor $\Delta \to \mathrm{Top}$ that sends $[n]$ to $|\Delta^n|$.

Two kinds of maps generate Δ:

1. The coface $d^i : [n-1] \hookrightarrow [n]$ omits i. Its affine map inserts a zero into barycentric coordinates at position i.
2. The codegeneracy $\sigma^i : [n+1] \twoheadrightarrow [n]$ identifies i and $i+1$. Its affine map replaces the adjacent coordinates t_i, t_{i+1} by their sum.

Every order-preserving map factors uniquely as a surjection followed by an injection. Geometrically, it first collapses repeated vertices and then includes a face.

4.2.3 Simplicial Sets

Definition 4.2.6. The simplicial sets of Definition 4.1.1, form the presheaf category $\mathcal{P}(\Delta) = \mathrm{Fun}(\Delta^{\mathrm{op}}, \mathrm{Set})$. For a simplicial set K, write $K_n = K([n])$ and $\alpha^* = K(\alpha)$ for its restriction maps.

The constructions fit into the following diagram of functors:

$$\mathrm{Top} \underset{|-|}{\overset{\mathrm{Sing}}{\rightleftarrows}} \mathcal{P}(\Delta) \underset{U}{\overset{\mathbb{Z}[-]}{\rightleftarrows}} \mathcal{P}(\Delta;\mathrm{Ab}) \xrightarrow{\mathrm{C}^{\mathrm{alt}}} \mathrm{Ch}_{\geq 0}(\mathrm{Ab}).$$

Here $|-| \dashv \mathrm{Sing}$ and $\mathbb{Z}[-] \dashv U$: realization is left adjoint to the singular functor, and the free simplicial abelian group functor is left adjoint to forgetting addition. The composite from left to right is the singular chain functor.

Elements of K_n are called n-simplices. The following theorem gives an equivalent description using only faces and degeneracies.

Theorem 4.2.7. The data of a simplicial set is equivalently a sequence of sets $K_0, K_1, K_2, \ldots$ equipped with face maps and degeneracy maps

$$\begin{aligned} d_i &: K_n \longrightarrow K_{n-1} \quad (n \geq 1,\ 0 \leq i \leq n), \\ s_i &: K_n \longrightarrow K_{n+1} \quad (n \geq 0,\ 0 \leq i \leq n), \end{aligned}$$

subject to the simplicial identities

$$\begin{aligned} d_i d_j &= d_{j-1} d_i && (i < j), \\ s_i s_j &= s_{j+1} s_i && (i \leq j), \\ d_i s_j &= \begin{cases} s_{j-1} d_i, & i < j, \\ \mathrm{id}, & i = j \text{ or } i = j+1, \\ s_j d_{i-1}, & i > j+1. \end{cases} \end{aligned}$$

The identities are required whenever the indicated composites are defined; composition is read from right to left.

More precisely, every such system extends uniquely to a functor $K : \Delta^{\mathrm{op}} \to \mathrm{Set}$ with $K([n]) = K_n$, $K(d^i) = d_i$, and $K(\sigma^i) = s_i$. Under this correspondence, simplicial maps are exactly families of functions $f_n : K_n \to L_n$ commuting with every face and degeneracy map.

Proof. A functor gives the stated maps by applying it to cofaces and codegeneracies. Their relations in Δ, with composition reversed, give the displayed identities.

Conversely, let $\alpha : [m] \to [n]$ have image of size $r+1$. Its unique factorization is $\alpha = \iota\pi$, where $\pi : [m] \twoheadrightarrow [r]$ records the position of each value in the ordered image and $\iota : [r] \hookrightarrow [n]$ includes that image. If $a_1 < \cdots < a_k$ are the omitted target values and $b_1 < \cdots < b_\ell$ are the indices with $\alpha(b_j) = \alpha(b_j + 1)$, its canonical expression is

$$\alpha = d^{a_k} \cdots d^{a_1} \sigma^{b_1} \cdots \sigma^{b_\ell}, \qquad k = n - r, \quad \ell = m - r.$$

An empty string denotes an identity. Each factor is given the domain required by the next factor to its right. Deleting repeated positions from largest to smallest proves the formula for π; inserting omitted positions from smallest to largest proves the formula for ι.

For completeness, the relations give a terminating reduction to this expression. A mixed adjacent pair consisting of a codegeneracy to the left of a coface is rewritten as a coface followed by a codegeneracy, or cancelled. At fixed word length each such interchange reduces the number of pairs in the wrong order; cancellation reduces length. Thus all cofaces can

be placed to the left. Within the two blocks, their elementary relations order the omitted and repeated indices as displayed. This last ordering can be performed by insertion: insert each new generator into an already ordered block, moving it across successive neighbors and adjusting its index by the relation. The block is finite, so each insertion terminates. An ordered block is uniquely determined by its omitted values or repeated positions. Hence two words for the same map reduce to the same expression. These relations therefore present Δ by generators and relations.

Assigning d_i and s_i to these generators therefore defines the pullback along every map of Δ, independently of its expression. The same relations ensure functoriality. Generation gives uniqueness and also shows that commuting with faces and degeneracies is equivalent to naturality. For the generators-and-relations formulation, see [Stacks, Lemmas 14.2.4 and 14.3.2], in the sections *The Category of Finite Ordered Sets* and *Simplicial Objects*. □

The faces of an edge e are its target $d_0 e$ and source $d_1 e$. The simplex $s_0 v$ is the degenerate edge at v; both of its endpoints are v. Geometrically, the first identity says that deleting vertex j and then vertex $i < j$ has the same effect as deleting vertex i and then the renumbered vertex $j-1$.

A simplex is *degenerate* if it is in the image of some s_i; otherwise it is *nondegenerate*. A degenerate simplex is still an element of K_n. It is not removed from the simplicial set.

Example 4.2.8. Use the representable standard simplex Δ^n of Example 4.1.3. Precomposition gives $\beta^*(\alpha) = \alpha\beta$. Its nondegenerate m-simplices are precisely the injective order-preserving maps $[m] \hookrightarrow [n]$. Thus Δ^1 has two nondegenerate vertices and one nondegenerate edge, while retaining simplices in every degree.

The notations $[n]$, Δ^n, and $|\Delta^n|$ refer respectively to an ordered set, a simplicial set, and a topological space.

Proposition 4.2.9. There is a natural bijection

$$\mathrm{Mor}_{\mathcal{P}(\Delta)}(\Delta^n, K) \cong K_n.$$

Proof. The standard simplex Δ^n is the representable presheaf $\mathrm{Mor}_\Delta(-,[n])$. Apply Yoneda: $x \in K_n$ corresponds to the simplicial map $\Delta^n \to K$ whose degree-m component is $\mathrm{Mor}_\Delta([m],[n]) \to K_m$, $\alpha \mapsto \alpha^* x$. □

4.3 Geometric Realization and Singular Simplices

4.3.1 Geometric Realization

The construction is discussed in [Kerodon, Section 1.2.3], *The Geometric Realization of a Simplicial Set*, especially Proposition 1.2.3.4 and Corollary 1.2.3.5. The realization–singular bijection is proved explicitly in Subsection 4.3.3.

Definition 4.3.1. Use the quotient realization $|K|$ of Definition 4.1.4, with its quotient topology. In the notation of Definition 4.2.5, its relations are $(\alpha^* x, t) \sim (x, \alpha_*(t))$. Write $[x,t]$ for the equivalence class of (x,t).

The face relations attach the boundary of one simplex to the specified lower-dimensional simplices. The degeneracy relations collapse a redundant simplex onto its lower-dimensional source. For example,

$$[s_0 v, (t_0, t_1)] = [v, (1)],$$

so a degenerate edge becomes a point, not an extra loop.

Every $x \in K_n$ has a characteristic map

$$\chi_x : |\Delta^n| \longrightarrow |K|, \qquad t \longmapsto [x, t].$$

These maps obey the commutative diagrams

$$\begin{array}{ccc} |\Delta^m| & \xrightarrow{\alpha_*} & |\Delta^n| \\ {\scriptstyle \chi_{\alpha^* x}}\downarrow & & \downarrow{\scriptstyle \chi_x} \\ |K| & = & |K|. \end{array}$$

Define the category of simplices Δ/K to have objects $([n], x)$ with $x \in K_n$; a morphism $([m], y) \to ([n], x)$ is a map $\alpha : [m] \to [n]$ with $y = \alpha^* x$. Then

$$|K| \cong \operatorname*{colim}_{([n],x) \in \Delta/K} |\Delta^n|.$$

This is a colimit in Top: continuous maps $|K| \to X$ correspond exactly to families of continuous maps $f_x : |\Delta^n| \to X$ satisfying $f_{\alpha^* x} = f_x \alpha_*$. Indeed, such a family defines a continuous map from the disjoint union, and compatibility is precisely what makes it descend through the quotient.

A simplicial map $f : K \to L$ induces

$$|f| : |K| \to |L|, \qquad [x, t] \longmapsto [f_n(x), t].$$

Compatibility with the relation follows from $f_m(\alpha^* x) = \alpha^* f_n(x)$. This defines a functor $|-| : \mathcal{P}(\Delta) \to \mathrm{Top}$.

Example 4.3.2.

1. The realization of the standard simplicial simplex is canonically the topological simplex introduced in Subsection 4.2.2. The map $|\Delta^n| \to \{t \in \mathbb{R}^{n+1} : t_i \geq 0,\ \sum_i t_i = 1\}$ is given by $[\alpha, t] \mapsto \alpha_*(t)$. Its inverse sends u to $[\mathrm{id}_{[n]}, u]$; the defining relation proves that the two composites are identities. We use $|\Delta^n|$ for this identified space.
2. A constant simplicial set on a set S has value S and every structure map is the identity. Its realization is the discrete space S.
3. Write S^1_Δ for the simplicial circle $\Delta^1/\partial\Delta^1$ of Example 4.1.3. Its realization was identified there with the topological circle.

Repeated vertices, loops, and distinct simplices with the same boundary are all allowed in a simplicial set. This flexibility goes beyond an abstract simplicial complex.

4.3.2 Singular Simplices

Definition 4.3.3. Use the singular simplicial set of Definition 4.1.4. We work in Top, without separation assumptions, and write $(\mathrm{Sing}\, X)_n = \mathrm{Mor}_{\mathrm{Top}}(|\Delta^n|, X)$. In the affine notation of Definition 4.2.5, its restriction maps are $\alpha^*(\tau) = \tau\alpha_*$. These morphism sets carry no additional topology here.

Precomposition and $(\alpha\beta)_* = \alpha_*\beta_*$ give the required contravariance. Thus $\mathrm{Sing}\, X$ is a simplicial set, called the *singular simplicial set* of X.

The first few degrees have direct geometric meanings:

1. $(\mathrm{Sing}\, X)_0$ is the underlying set of points of X.

2. $(\operatorname{Sing} X)_1$ is the set of continuous paths, with the interval parametrized by $(1-u, u)$ for $0 \le u \le 1$. Thus $d_1\gamma = \gamma(0)$ and $d_0\gamma = \gamma(1)$.
3. $(\operatorname{Sing} X)_2$ consists of continuous maps of a filled triangle into X. Its three faces are the restrictions to its three ordered edges.

A singular simplex need not be injective or linear, and its image need not look like a geometric simplex. In particular, an arbitrary space need not come with a chosen triangulation for its singular simplicial set to be defined.

For $\tau : |\Delta^n| \to X$, the degeneracy $s_i\tau : |\Delta^{n+1}| \to X$ is

$$(t_0, \dots, t_{n+1}) \longmapsto \tau(t_0, \dots, t_{i-1}, t_i + t_{i+1}, t_{i+2}, \dots, t_{n+1}).$$

For $n = 0$, this is a constant path. A path that goes out and comes back is generally *not* a degenerate edge: degeneracy means the specified factorization through a lower-dimensional standard simplex, not merely that endpoints coincide.

If $f : X \to Y$ is continuous, postcomposition defines

$$(\operatorname{Sing} f)_n : (\operatorname{Sing} X)_n \to (\operatorname{Sing} Y)_n, \qquad \tau \longmapsto f\tau.$$

Since postcomposition commutes with precomposition, this is simplicial and defines the functor $\operatorname{Sing} : \mathrm{Top} \to \mathcal{P}(\Delta)$.

Example 4.3.4. The singular simplicial set of a point is Δ^0: it has one simplex in every degree. If X is discrete, every map $|\Delta^n| \to X$ is constant, so $\operatorname{Sing} X$ is the constant simplicial set on X. For $X = [0, 1]$, the simplicial set $\operatorname{Sing} X$ is much larger than Δ^1: it includes every continuous path and every continuous singular simplex in the interval.

4.3.3 The Realization–Singular Adjunction

An adjunction here is a natural bijection between two sets of maps. The relevant bijection says that mapping a glued space into X is the same as assigning compatible singular simplices in X to its pieces.

Theorem 4.3.5. Use the realization–singular adjunction established after Definition 4.1.4. The proof below expresses its bijection through the colimit universal property, preparing the unit and counit construction.

Proof. By the colimit universal property, a continuous map $|K| \to X$ is a compatible family $f_x : |\Delta^n| \to X$. Since $(\operatorname{Sing} X)_n = \mathrm{Mor}_{\mathrm{Top}}(|\Delta^n|, X)$, this is exactly a family of functions

$$K_n \longrightarrow (\operatorname{Sing} X)_n, \qquad x \longmapsto f_x,$$

commuting with restriction maps: a morphism of presheaves $K \to \operatorname{Sing} X$. Explicitly, $f : |K| \to X$ determines $g : K \to \operatorname{Sing} X$ with $g_n(x) = f\chi_x$. Conversely, $g : K \to \operatorname{Sing} X$ determines $f : |K| \to X$ by $f([x, t]) = g_n(x)(t)$. Both commute with precomposition in K and postcomposition in X, so the bijection is natural. □

The identity of $|K|$ corresponds to the *unit*

$$(\eta_K)_n : K_n \longrightarrow (\operatorname{Sing} |K|)_n, \qquad x \longmapsto \chi_x.$$

Thus each abstract simplex becomes the singular simplex that parametrizes its own image in the realization.

The identity of $\operatorname{Sing} X$ corresponds to the *counit*, or evaluation map,

$$\epsilon_X : |\operatorname{Sing} X| \longrightarrow X, \qquad [\tau, t] \longmapsto \tau(t).$$

The triangle identities are

$$\epsilon_{|K|} \circ |\eta_K| = \mathrm{id}_{|K|}, \qquad \mathrm{Sing}(\epsilon_X) \circ \eta_{\mathrm{Sing}\, X} = \mathrm{id}_{\mathrm{Sing}\, X}.$$

For example, the first composite sends $[x,t]$ to $[\chi_x, t]$ and then to $\chi_x(t) = [x,t]$. The second says that evaluating the characteristic simplex associated to τ recovers τ.

Remark 4.3.6. The map η_{Δ^1} includes the simplices determined by the ordered vertices into all singular simplices of the interval; most continuous paths are not in its image. Likewise, evaluation need not be a homeomorphism. The deeper comparison with homotopy types is a theorem beyond the adjunction itself, and is not needed to construct singular chains. For further context on this adjunction and its relation to cellular homology, see [May99, Chapter 16, §4, pp. 126–127], where the chain complexes are normalized.

4.4 Singular Chains and Coefficients

4.4.1 The Free Simplicial Abelian Group Functor

Definition 4.4.1. For a set S, the free abelian group on S is

$$\mathbb{Z}[S] = \bigoplus_{s \in S} \mathbb{Z} s.$$

Its elements are finite formal sums $\sum_s a_s s$ with $a_s \in \mathbb{Z}$. For a function $f : S \to T$, set

$$\mathbb{Z}[f]\left(\sum_s a_s s\right) = \sum_s a_s f(s).$$

We write s also for the formal generator labelled by $s \in S$; addition in $\mathbb{Z}[S]$ is addition of formal sums. If f identifies two elements, their coefficients add. For example, if $f(s) = f(s') = t$, then $2s - s'$ maps to t.

For every abelian group B, a function $S \to U(B)$ extends uniquely to a homomorphism $\mathbb{Z}[S] \to B$ by linearity. Here $U(B)$ denotes its underlying set. Thus

$$\mathrm{Mor}_{\mathsf{Ab}}(\mathbb{Z}[S], B) \cong \mathrm{Mor}_{\mathsf{Set}}(S, U(B)).$$

Definition 4.4.2. The simplicial modules of Definition 4.1.5, specialize over $\mathbb{Z}$ to simplicial abelian groups. They form $\mathcal{P}(\Delta; \mathsf{Ab})$. Forgetting their group structures degreewise defines $U : \mathcal{P}(\Delta; \mathsf{Ab}) \to \mathcal{P}(\Delta)$.

For a simplicial set K, define

$$\mathbb{Z}[K]_n = \mathbb{Z}[K_n], \qquad \alpha^*\left(\sum_x a_x x\right) = \sum_x a_x \alpha^* x.$$

The simplicial identities hold on generators because they hold in K, and therefore hold on all finite sums. In particular,

$$d_i x = d_i x, \qquad s_i x = s_i x.$$

A simplicial map $f : K \to L$ induces the degreewise homomorphisms $\mathbb{Z}[f_n] : \mathbb{Z}[K_n] \to \mathbb{Z}[L_n]$ given by $x \mapsto f_n(x)$. We obtain the functor

$$\mathbb{Z}[-] : \mathcal{P}(\Delta) \longrightarrow \mathcal{P}(\Delta; \mathsf{Ab}).$$

The group in degree n is free on *all* n-simplices, including degenerate simplices. It is not the free group on the vertices alone.

Proposition 4.4.3. Let $L : \mathcal{D} \rightleftarrows \mathcal{E} : R$ be an adjunction and $\mathcal{I}$ a small category. Postcomposition induces an adjunction

$$L_* : \mathrm{Fun}(\mathcal{I}, \mathcal{D}) \rightleftarrows \mathrm{Fun}(\mathcal{I}, \mathcal{E}) : R_*.$$

Proof. For $F : \mathcal{I} \to \mathcal{D}$ and $G : \mathcal{I} \to \mathcal{E}$, transpose each component $LF(i) \to G(i)$ under $L \dashv R$. Naturality of the adjunction bijection identifies the naturality equations for these components with those for $F(i) \to RG(i)$. Thus componentwise transposition gives a natural bijection between the two sets of natural transformations. □

Proposition 4.4.4. For $K \in \mathcal{P}(\Delta)$ and $A \in \mathcal{P}(\Delta; \mathrm{Ab})$, there is a natural bijection

$$\mathrm{Mor}_{\mathcal{P}(\Delta;\mathrm{Ab})}(\mathbb{Z}[K], A) \cong \mathrm{Mor}_{\mathcal{P}(\Delta)}(K, U(A)).$$

Proof. Apply the pointwise adjunction proposition to $\mathbb{Z}[-] : \mathrm{Set} \rightleftarrows \mathrm{Ab} : U$ with $\mathcal{I} = \Delta^{\mathrm{op}}$. Concretely, $g : K \to U(A)$ extends uniquely by

$$\overline{g}_n \left(\sum_x a_x x \right) = \sum_x a_x g_n(x).$$

Naturality of g makes these extensions commute with every simplicial operator. Restriction to generators is the inverse. □

The unit in degree n is the function $K_n \to \mathbb{Z}[K_n]$ sending a simplex to its formal generator. The counit is the simplicial homomorphism $\mathbb{Z}[U(A)] \to A$ that evaluates each finite formal sum using the addition in A. A generator labelled by a sum in A is distinct from the formal sum of the corresponding generators; the counit sends them to the same element.

Remark 4.4.5. The functor $\mathbb{Z}[-]$ freely introduces sums and additive inverses while retaining all faces and degeneracies. A simplicial abelian group still has $n+1$ face maps in degree n. The next construction combines them into one differential.

4.4.2 Alternating Faces and Degenerate Simplices

Definition 4.4.6. Use the alternating-face complex of Definition 4.1.5, retaining degenerate simplices. Here $\mathrm{C}^{\mathrm{alt}}(A)$ denotes that construction for simplicial abelian groups, with homological grading and zero groups in negative degrees.

The first two differentials are $\partial_1 = d_0 - d_1$ and $\partial_2 = d_0 - d_1 + d_2$. Thus the boundary of an edge is target minus source, and the boundary of a triangle is its three edges with alternating signs.

Proposition 4.4.7. Proposition 4.1.6 and Section 4.1.3 prove functoriality of the alternating-face construction:

$$\mathrm{C}^{\mathrm{alt}} : \mathcal{P}(\Delta; \mathrm{Ab}) \longrightarrow \mathrm{Ch}_{\geq 0}(\mathrm{Ab}).$$

Proof. The face-pairing proof and the compatibility with simplicial maps are given in Proposition 4.1.6 and Section 4.1.3. □

Each codimension-two face occurs twice, with opposite signs. See also [Stacks, Section 14.23, Simplicial Objects and Chain Complexes]; our notation $\mathrm{C}^{\mathrm{alt}}(A)$ denotes the alternating-face complex called $s(A)$ there.

For a simplicial set K, write

$$\mathrm{C}_\bullet(K; \mathbb{Z}) = \mathrm{C}^{\mathrm{alt}}(\mathbb{Z}[K]), \qquad \mathrm{H}_n(K; \mathbb{Z}) = \ker(\partial_n)/\operatorname{im}(\partial_{n+1}).$$

These are the unnormalized simplicial chains and their homology. In positive degrees,

$$\partial x = \sum_{i=0}^{n} (-1)^i d_i x.$$

A degenerate simplex is not the zero element of a free chain group. For instance, the constant edge $\alpha = s_0 v$ has $\partial\alpha = 0$, but α is a nonzero generator of C_1. It becomes a boundary: if $\beta = s_0 s_0 v$, then all three faces of β equal α, so $\partial\beta = \alpha - \alpha + \alpha = \alpha$.

Example 4.4.8. Apply the constant simplicial abelian-group computation of Example 4.1.7, to $G = \mathbb{Z}$. Since $\mathbb{Z}[\Delta^0]$ is constant, this gives $\mathrm{H}_0(\Delta^0; \mathbb{Z}) = \mathbb{Z}$ and zero homology in positive degrees, while retaining degenerate simplices.

4.4.3 The Singular Chain Functor

Definition 4.4.9. The singular chain complex of a topological space X is

$$\mathrm{C}_\bullet^{\mathrm{sing}}(X; \mathbb{Z}) = \mathrm{C}^{\mathrm{alt}}(\mathbb{Z}[\mathrm{Sing}\, X]).$$

Its degree-n group is freely generated by continuous maps $\tau : |\Delta^n| \to X$. For $n \geq 1$, its boundary is

$$\partial\tau = \sum_{i=0}^{n} (-1)^i \tau d^i_*.$$

The singular homology of X is $\mathrm{H}_n^{\mathrm{sing}}(X; \mathbb{Z}) = \mathrm{H}_n(\mathrm{C}_\bullet^{\mathrm{sing}}(X; \mathbb{Z}))$.

The construction separates three steps:

1. Sing records all continuous simplices and their restrictions and repetitions.
2. $\mathbb{Z}[-]$ permits finite integral combinations of these simplices.
3. $\mathrm{C}^{\mathrm{alt}}$ adds the face maps with signs to define a boundary operator.

None of these steps requires choosing a triangulation of X.

For a continuous map $f : X \to Y$, the induced chain map sends τ to $f\tau$. For every $n \geq 1$, we have the commutative square

$$\begin{array}{ccc} \mathrm{C}_n^{\mathrm{sing}}(X; \mathbb{Z}) & \xrightarrow{f_*} & \mathrm{C}_n^{\mathrm{sing}}(Y; \mathbb{Z}) \\ \downarrow{\scriptstyle \partial} & & \downarrow{\scriptstyle \partial} \\ \mathrm{C}_{n-1}^{\mathrm{sing}}(X; \mathbb{Z}) & \xrightarrow[f_*]{} & \mathrm{C}_{n-1}^{\mathrm{sing}}(Y; \mathbb{Z}). \end{array}$$

It follows that singular homology is functorial.

Example 4.4.10. For a path $\gamma : x \to y$, its boundary is $y - x$. Consequently,

$$\mathrm{H}_0^{\mathrm{sing}}(X; \mathbb{Z}) \cong \mathbb{Z}[\pi_0^{\mathrm{path}} X],$$

where $\pi_0^{\mathrm{path}} X$ is the set of path components, not merely connected components.

For a singular triangle τ with ordered edges $\gamma_{01}, \gamma_{02}, \gamma_{12}$,

$$\partial\tau = \gamma_{12} - \gamma_{02} + \gamma_{01}.$$

Thus the cycle given by the two successive edges minus the long edge is a boundary. This is a relation among chains modulo boundaries; the individual nonclosed edges need not represent classes in $\mathrm{H}_1^{\mathrm{sing}}(X; \mathbb{Z})$.

4.4.4 Coefficients and the Real Projective Plane

Definition 4.4.11. Let X be a simplicial set and let G be an abelian group. With $A_\bullet = \mathrm{C}_\bullet(X;\mathbb{Z})$, define

$$\mathrm{H}_n(X;G) = \mathrm{H}_n(A_\bullet \otimes_{\mathbb{Z}} G).$$

For a space Y, apply this construction to $X = \mathrm{Sing}(Y)$ to obtain singular homology with coefficients in G. The chain group is the direct sum of one copy of G for each singular simplex, and its differential is the alternating sum of restrictions to faces.

Example 4.4.12. This is a forward example: the topological input will be proved in Lecture 8, Theorems 8.2.2 and 8.3.4 (cellular comparison and the differential for real projective space). Appendix B, Proposition B.1.2 and the projective-plane calculation immediately following it, also explain the coefficient 2 from the attaching word aa. Using those results, the cellular chain complex of $\mathbb{RP}^2$ over $\mathbb{Z}$ is

$$0 \longrightarrow \mathbb{Z} \xrightarrow{2} \mathbb{Z} \xrightarrow{0} \mathbb{Z} \longrightarrow 0$$

in degrees $2, 1, 0$. Thus $\mathrm{H}_1(\mathbb{RP}^2;\mathbb{Z}) = \mathbb{Z}/2$ and $\mathrm{H}_2(\mathbb{RP}^2;\mathbb{Z}) = 0$. The cellular comparison theorem identifies these groups with singular homology.

For coefficients $\mathbb{Z}/2$, the multiplication-by-two differential vanishes, giving

$$\mathrm{H}_2(\mathbb{RP}^2;\mathbb{Z}/2) \cong \mathbb{Z}/2.$$

This class is invisible in $\mathrm{H}_2(\mathbb{RP}^2;\mathbb{Z}) \otimes \mathbb{Z}/2$, which is zero. Its value agrees with

$$\mathrm{Tor}_1^{\mathbb{Z}}(\mathrm{H}_1(\mathbb{RP}^2;\mathbb{Z}), \mathbb{Z}/2) = \mathrm{Tor}_1^{\mathbb{Z}}(\mathbb{Z}/2, \mathbb{Z}/2) \cong \mathbb{Z}/2.$$

A simplicial chain complex is free in each degree, but it is usually not a resolution of its degree-zero homology: it may have higher homology. To compute Tor of a single module, use an actual resolution. To compute homology with coefficients of a space, retain the full chain complex of that space.

Generalized Homology and the Eilenberg–Steenrod Axioms

Algebraic Topology I | Fall 2026 | September 28 and 30

5.1 Homotopy Invariance and Homotopy Pushouts

The first part of a homology theory is functorial data. The gluing axiom needs a construction that remembers the homotopy along which two maps agree.

5.1.1 Graded Functors

For a space W, write $\mathrm{pr}_W : W \longrightarrow *$ for its unique map to a point.

Write Top for the category of topological spaces and continuous maps, and Ab for the category of abelian groups. Homotopy pushouts are understood up to homotopy equivalence.

Definition 5.1.1. A graded homology functor is a family of covariant functors

$$A_n : \mathrm{Top} \longrightarrow \mathrm{Ab} \qquad (n \in \mathbb{Z}).$$

For a map $f : X \longrightarrow Y$, write $f_* = A_n(f) : A_n(X) \longrightarrow A_n(Y)$. Functoriality means that identities induce identities and $(gf)_* = g_* f_*$. The degree is specified by the source and target when it is omitted from f_*.

5.1.2 Homotopies and Homotopy Invariance

Definition 5.1.2. Write $I = [0,1]$. Given maps

$$f_0, f_1 : X \longrightarrow Y,$$

a *homotopy* from f_0 to f_1 is a continuous map

$$\alpha : X \times I \longrightarrow Y, \qquad \alpha(x,0) = f_0(x), \quad \alpha(x,1) = f_1(x).$$

The endpoint inclusions are

$$\iota_0 : X \longrightarrow X \times I, \qquad x \longmapsto (x,0),$$
$$\iota_1 : X \longrightarrow X \times I, \qquad x \longmapsto (x,1).$$

Thus the endpoint conditions say that the following triangles commute:

$$\begin{array}{ccccc} X & \xrightarrow{\iota_0} & X \times I & \xleftarrow{\iota_1} & X \\ & \searrow^{f_0} & \downarrow{\scriptstyle \alpha} & \swarrow_{f_1} & \\ & & Y & & \end{array}$$

We write $f_0 \simeq f_1$. Constant homotopies, reversal of I, and concatenation show that this is an equivalence relation on the set of maps between two fixed spaces.

Definition 5.1.3. The category hTop has spaces as objects and morphism sets

$$\mathrm{Mor}_{\mathrm{hTop}}(X,Y) = \mathrm{Mor}_{\mathrm{Top}}(X,Y)/\simeq .$$

The class of a map f is written $[f]$. For composable maps

$$f : X \longrightarrow Y, \qquad g : Y \longrightarrow Z,$$

composition and identities are defined by

$$[g] \circ [f] = [g \circ f], \qquad \mathrm{id}_X = [\mathrm{id}_X].$$

To check composition, take homotopies

$$\alpha : X \times I \longrightarrow Y, \qquad \beta : Y \times I \longrightarrow Z.$$

The composite homotopy is the continuous map

$$\gamma : X \times I \longrightarrow Z, \qquad (x,t) \longmapsto \beta(\alpha(x,t),t).$$

Its endpoints are the composites of the endpoint maps of α and β. Hence composition of homotopy classes is well defined. Associativity and the identity laws follow from those in Top.

The quotient functor is

$$q : \mathrm{Top} \longrightarrow \mathrm{hTop}, \qquad q(X) = X, \quad q(f) = [f].$$

Definition 5.1.4. A *homotopy equivalence* is a map admitting a homotopy inverse:

$$f : X \longrightarrow Y, \qquad g : Y \longrightarrow X, \qquad gf \simeq \mathrm{id}_X, \quad fg \simeq \mathrm{id}_Y.$$

Thus f is a homotopy equivalence exactly when $[f]$ is an isomorphism in hTop, with inverse $[g]$.

Definition 5.1.5. Let $\mathcal{C}$ be a category. A functor

$$F : \mathrm{Top} \longrightarrow \mathcal{C}$$

is *homotopy invariant* if homotopic maps have equal images: for

$$f_0, f_1 : X \longrightarrow Y, \qquad f_0 \simeq f_1,$$

one has

$$F(f_0) = F(f_1) : F(X) \longrightarrow F(Y).$$

Equivalently, there is a unique functor

$$\overline{F} : \mathrm{hTop} \longrightarrow \mathcal{C}, \qquad \overline{F}(X) = F(X), \quad \overline{F}([f]) = F(f),$$

making the triangle commute:

$$\begin{array}{ccc} \mathrm{Top} & \xrightarrow{\ q\ } & \mathrm{hTop} \\ & \searrow^{F} & \downarrow{\scriptstyle \overline{F}} \\ & & \mathcal{C}. \end{array}$$

Homotopy invariance is exactly what makes the displayed formula on morphisms well defined. The identity and composition laws descend from F, and the formulas force uniqueness.

Proposition 5.1.6. Every homotopy invariant functor sends homotopy equivalences to isomorphisms.

Proof. For a homotopy equivalence and its homotopy inverse, write

$$f : X \longrightarrow Y, \qquad g : Y \longrightarrow X.$$

Their images have types

$$F(f) : F(X) \longrightarrow F(Y), \qquad F(g) : F(Y) \longrightarrow F(X).$$

Functoriality and homotopy invariance give

$$F(g)F(f) = F(gf) = \mathrm{id}_{F(X)}, \qquad F(f)F(g) = F(fg) = \mathrm{id}_{F(Y)}.$$

Thus $F(g)$ is the inverse of $F(f)$. □

5.1.3 Homotopy Pushouts

The gluing axiom uses the double mapping cylinder to retain the homotopy between the two composites.

Definition 5.1.7. For a span $X \xleftarrow{f} W \xrightarrow{g} Y$, its double mapping cylinder is

$$M(f,g) = (X + (W \times I) + Y)/\sim,$$

where $(w,0) \sim f(w)$ and $(w,1) \sim g(w)$. Here $+$ denotes a finite disjoint union. Let $\iota_X : X \longrightarrow M(f,g)$ and $\iota_Y : Y \longrightarrow M(f,g)$ be the canonical maps. The cylinder supplies a specified homotopy from $\iota_X f$ to $\iota_Y g$.

A map from $M(f,g)$ to a space Z is equivalent to maps $h : X \longrightarrow Z$, $k : Y \longrightarrow Z$, and a homotopy from hf to kg. This follows from the quotient universal property: the homotopy defines the map on $W \times I$, and its endpoint conditions are precisely the identifications in the quotient.

Definition 5.1.8. Consider maps and a specified homotopy $\alpha : hf \simeq kg$ in a square

$$\begin{array}{ccc} & X & \\ {\scriptstyle f}\nearrow & & \searrow{\scriptstyle h} \\ W & \Downarrow \alpha & Z. \\ {\scriptstyle g}\searrow & & \nearrow{\scriptstyle k} \\ & Y & \end{array}$$

The square, with α, is a homotopy pushout if its comparison map

$$\theta_\alpha : M(f,g) \longrightarrow Z$$

is a homotopy equivalence. This comparison restricts to h, k and sends (w,t) to $\alpha(w,t)$.

Remark 5.1.9. A commutative pushout square need not be a homotopy pushout. For the span $* \xleftarrow{\mathrm{pr}_{S^0}} S^0 \xrightarrow{\mathrm{pr}_{S^0}} *$, where S^0 has two points, the ordinary pushout is a point. Its double mapping cylinder consists of two intervals joining the two target points, hence is a circle. Thus the comparison to the ordinary pushout is not a homotopy equivalence. The specified homotopy in Definition 5.1.8 is part of the gluing data.

5.2 Homology Axioms and Coefficients

We now impose homotopy invariance, additivity, and Mayer–Vietoris. The connecting homomorphisms in the third axiom are additional structure; their existence is not a consequence of functoriality alone. The gluing-first formulation uses homotopy pushouts in place of a separate excision axiom. The original pair-based axioms and uniqueness announcement appear in [ES45, §§4–5, pp. 118–119].

5.2.1 The Homology Axioms and Naturality

Homotopic maps induce the same homomorphism, and disjoint unions correspond to direct sums.

Axiom 5.2.1. Homotopy invariance: for every $n \in \mathbb{Z}$ and every pair of homotopic maps $f_0, f_1 : X \longrightarrow Y$,

$$f_{0,*} = f_{1,*} : A_n(X) \longrightarrow A_n(Y).$$

Axiom 5.2.2. Additivity: for every set-indexed family $(X_\lambda)_{\lambda\in\Lambda}$, the inclusions $\iota_\lambda : X_\lambda \longrightarrow \coprod_{\lambda\in\Lambda} X_\lambda$ induce an isomorphism

$$\bigoplus_{\lambda\in\Lambda} A_n(X_\lambda) \xrightarrow{\cong} A_n\left(\coprod_{\lambda\in\Lambda} X_\lambda\right).$$

The empty family is included, so $A_n(\emptyset) = 0$.

The homotopy in a double mapping cylinder makes $\iota_{X,*}f_* = \iota_{Y,*}g_*$. Consequently the composite of the following two maps is zero:

$$A_n(W) \xrightarrow{(f_*,-g_*)} A_n(X) \oplus A_n(Y) \xrightarrow{\iota_{X,*}+\iota_{Y,*}} A_n(M(f,g)).$$

The gluing axiom asserts exactness and supplies the connecting maps that continue this sequence.

Axiom 5.2.3. Mayer–Vietoris: for every span $X \xleftarrow{f} W \xrightarrow{g} Y$, there are homomorphisms

$$\delta_n : A_n(M(f,g)) \longrightarrow A_{n-1}(W)$$

for which the following five consecutive terms are exact in every degree:

$$A_n(W) \xrightarrow{(f_*,-g_*)} A_n(X) \oplus A_n(Y) \xrightarrow{\iota_{X,*}+\iota_{Y,*}} A_n(M(f,g)) \xrightarrow{\delta_n} A_{n-1}(W) \xrightarrow{(f_*,-g_*)} A_{n-1}(X) \oplus A_{n-1}(Y). \quad (13)$$

The homomorphisms δ_n are natural for maps of spans and their induced maps of double mapping cylinders.

A map of spans consists of maps

$$\alpha : W \longrightarrow W', \qquad \beta : X \longrightarrow X', \qquad \gamma : Y \longrightarrow Y'$$

satisfying $\beta f = f'\alpha$ and $\gamma g = g'\alpha$. It induces

$$\mu : M(f,g) \longrightarrow M(f',g'),$$

whose restrictions are β on X, γ on Y, and $(w,t) \mapsto (\alpha(w),t)$ on the cylinder. Naturality of the connecting maps means that

$$\begin{array}{ccc} A_n(M(f,g)) & \xrightarrow{\delta_n} & A_{n-1}(W) \\ \downarrow{\scriptstyle \mu_*} & & \downarrow{\scriptstyle \alpha_*} \\ A_n(M(f',g')) & \xrightarrow[\delta'_n]{} & A_{n-1}(W') \end{array}$$

commutes. Exactness in (13) means that the image of each homomorphism is the kernel of the next one.

Definition 5.2.4. A generalized homology theory in this formulation consists of the functors of Definition 5.1.1 and the natural connecting homomorphisms of Axiom 5.2.3, satisfying Axioms 5.2.1–5.2.3. A morphism of theories is a family of natural transformations commuting with the connecting homomorphisms.

Proposition 5.2.5. For a homotopy pushout square as in Definition 5.1.8, there is a long exact sequence with five consecutive terms

$$A_n(W) \xrightarrow{(f_*,-g_*)} A_n(X) \oplus A_n(Y) \xrightarrow{h_*+k_*} A_n(Z) \xrightarrow{\delta_n^Z} A_{n-1}(W) \xrightarrow{(f_*,-g_*)} A_{n-1}(X) \oplus A_{n-1}(Y).$$

Here $\delta_n^Z = \delta_n(\theta_{\alpha,*})^{-1}$.

Proof. Axiom 5.2.1 sends homotopy equivalences to isomorphisms: the image of a homotopy inverse is an inverse. Thus $\theta_{\alpha,*}$ is invertible. Transport (13) through this isomorphism. The identities $\theta_\alpha \iota_X = h$ and $\theta_\alpha \iota_Y = k$ identify its middle map with $h_* + k_*$. Naturality holds for compatible maps of these squares and their chosen cylinder comparisons. □

Remark 5.2.6. An open cover $X = U \cup V$ gives the familiar square

$$\begin{array}{ccc} & U & \\ \nearrow & & \searrow \\ U \cap V & & X. \\ \searrow & & \nearrow \\ & V & \end{array}$$

For a numerable cover, this is a homotopy pushout. Here numerable means that there is a subordinate partition of unity; for two sets, continuous functions $\lambda_U, \lambda_V : X \longrightarrow [0,1]$ have sum 1 and supports contained in U, V, respectively. Let $p : M \to X$ be the comparison from the double mapping cylinder of $U \leftarrow U \cap V \to V$. Define a section $s : X \to M$ by sending $x \in U \cap V$ to $[(x, \lambda_V(x))]$, $x \in U - V$ to its copy in U, and $x \in V - U$ to its copy in V. These formulas agree at height zero and height one under the cylinder identifications. To check continuity, a point outside V has a neighborhood disjoint from $\operatorname{supp}(\lambda_V)$, where the section is the inclusion into U; the analogous statement holds outside U. On $U \cap V$ the formula is continuous. These neighborhoods cover X, proving continuity, and $ps = \mathrm{id}_X$.

On a cylinder point $[(x,t)]$, define the homotopy at time u by the height $(1-u)t + u\lambda_V(x)$. On a point of the U copy use height $u\lambda_V(x)$ when $x \in U \cap V$, and leave it fixed outside V. On the V copy use height $1 - u + u\lambda_V(x)$ on the overlap, and leave it fixed outside U. The same support-neighborhood argument proves continuity on the two copies, and the formulas agree with the cylinder formula at its ends. The quotient universal property therefore gives a continuous homotopy on $M \times I$. Here taking the product of the defining quotient map with I preserves a quotient map as proved in Lecture 7, Lemma 7.1.7. Its endpoints are id_M and sp. Thus p is a homotopy equivalence, and Proposition 5.2.5 applies. For the sphere covers used below, no general partition-of-unity theorem is needed. Write t for the last coordinate on S^m, and take $U = \{t < 1\}$ and $V = \{t > -1\}$. The functions

$$u = \max(0, \tfrac{1}{2} - t), \qquad v = \max(0, t + \tfrac{1}{2}), \qquad \lambda_U = \frac{u}{u+v}, \qquad \lambda_V = \frac{v}{u+v}$$

have positive denominator, sum to one, and have supports contained in U and V, respectively.

5.2.2 Coefficients and the Dimension Axiom

The values on a point remain part of the information in a generalized theory. They determine its values on discrete and contractible spaces, but need not be concentrated in degree zero.

Definition 5.2.7. The coefficient groups of A are the graded abelian groups $G_n = A_n(*)$. The dimension axiom is the additional condition $G_n = 0$ for $n \neq 0$. For ordinary integral homology one further chooses $G_0 \cong \mathbb{Z}$.

Proposition 5.2.8. For every nonempty contractible space X and every discrete set E,

$$A_n(X) \cong G_n, \qquad A_n(E) \cong \bigoplus_{e \in E} G_n.$$

The first isomorphism is induced by the unique map to a point, and the second by the point inclusions.

Proof. The map $\mathrm{pr}_X : X \longrightarrow *$ is a homotopy equivalence, so Axiom 5.2.1 gives the first isomorphism. A discrete set is a disjoint union of points; Axiom 5.2.2 gives the second, including $E = \emptyset$. □

Example 5.2.9. Once ordinary integral homology satisfies the axioms, its shift $A_n(X) = \mathrm{H}_{n-1}(X; \mathbb{Z})$ also satisfies them: shift every degree in its Mayer–Vietoris sequence and connecting maps. Its only nonzero coefficient group is $G_1 = \mathbb{Z}$. This theory violates the dimension axiom while retaining homotopy invariance, additivity, and Mayer–Vietoris. A more general family is $A_n(X) = \bigoplus_{r \in \mathbb{Z}} \mathrm{H}_{n-r}(X; B_r)$ for a fixed graded family of coefficient groups (B_r). Direct sums of exact sequences are exact, so the same argument applies.

Remark 5.2.10. Generalized theories have groups in every integer degree; their long exact sequences do not stop at degree zero. Ordinary singular homology is obtained with its dimension normalization and zero negative-degree groups. For the broader question of uniqueness among theories satisfying the axioms, see [May99, Chapter 15, §2, pp. 119–120].

5.3 Relative Homology and Spheres

The homotopy cofiber defines relative homology. Its Mayer–Vietoris sequence yields the long exact sequence of a pair.

5.3.1 Reduced Homology and Homotopy Cofibers

Write $*$ for a one-point space. Choosing a point of a space splits off its coefficient summand.

Definition 5.3.1. For a space Y with a chosen point y_0, let $\mathrm{pr}_Y : Y \longrightarrow *$ be the unique map and $s : * \longrightarrow Y$ select y_0. Define

$$\widetilde{A}_n(Y) = \ker((\mathrm{pr}_Y)_* : A_n(Y) \longrightarrow G_n).$$

Since $(\mathrm{pr}_Y)_* s_* = \mathrm{id}$, there is an isomorphism $A_n(Y) \cong \widetilde{A}_n(Y) \oplus G_n$. Explicitly, a is sent to $(a - s_*((\mathrm{pr}_Y)_*(a)), (\mathrm{pr}_Y)_*(a))$.

Definition 5.3.2. For a subspace inclusion $\iota_X : X' \hookrightarrow X$, define its homotopy cofiber $C(\iota_X)$ to be the double mapping cylinder of $X \xleftarrow{\iota_X} X' \xrightarrow{\mathrm{pr}_{X'}} *$. Its distinguished point is the copy of $*$. Define

$$A_n(X, X') = \widetilde{A}_n(C(\iota_X)).$$

A map of pairs is a commutative square

$$\begin{array}{ccc} X' & \xhookrightarrow{\iota_X} & X \\ {\scriptstyle f'}\downarrow & & \downarrow{\scriptstyle f} \\ Y' & \xhookrightarrow[\iota_Y]{} & Y. \end{array}$$

It gives a basepoint-preserving map $C(\iota_X) \longrightarrow C(\iota_Y)$, hence a homomorphism of relative groups. Functoriality follows directly from the cylinder construction.

5.3.2 The Long Exact Sequence of a Pair

Mayer–Vietoris for a homotopy cofiber gives the relative exact sequence.

Proposition 5.3.3. Every pair has a natural long exact sequence with five consecutive terms

$$A_n(X') \xrightarrow{\iota_{X,*}} A_n(X) \xrightarrow{q_*} A_n(X, X') \xrightarrow{\delta_n} A_{n-1}(X') \xrightarrow{\iota_{X,*}} A_{n-1}(X).$$

Proof. Let $\kappa : X \longrightarrow C(\iota_X)$ and $s : * \longrightarrow C(\iota_X)$ be the two canonical maps. Mayer–Vietoris gives

$$\cdots \longrightarrow A_n(X') \xrightarrow{(\iota_{X,*}, -(\mathrm{pr}_{X'})_*)} A_n(X) \oplus G_n \xrightarrow{\kappa_* + s_*} A_n(C(\iota_X)) \xrightarrow{\delta_n} A_{n-1}(X') \longrightarrow \cdots,$$

where $\mathrm{pr}_{X'} : X' \longrightarrow *$ is the unique map. Under the splitting of Definition 5.3.1, set $q_*(b) = \kappa_*(b) - s_*((\mathrm{pr}_X)_*(b))$ and define the connecting map δ_n in the claimed sequence by restricting the displayed connecting map to the reduced summand. Change coordinates on the middle direct sum by $(b, g) \mapsto (b, (\mathrm{pr}_X)_* b + g)$. The first map becomes $a \mapsto (\iota_{X,*} a, 0)$; the next becomes $(b, g) \mapsto (q_* b, g)$. The displayed connecting map vanishes on the coefficient summand by exactness. Thus the displayed sequence splits into the claimed sequence and the identity map $G_n \longrightarrow G_n$. Exactness and naturality pass to the claimed summand. □

Definition 5.3.4. For a space X, let $X_+ = X + *$ be the pointed space obtained by adjoining a disjoint basepoint, with basepoint the added point.

Remark 5.3.5. For a cell pair, or a subspace inclusion satisfying the homotopy extension property, the map from the mapping cone to X/X' is a homotopy equivalence (using X_+ when $X' = \emptyset$). Thus the relative groups are the reduced groups of this quotient. For an arbitrary subspace, the homotopy cofiber is the construction to use; the ordinary quotient need not have its homotopy type. The homotopy extension property says that a homotopy on X' extending a map on X extends to a homotopy on X. See May, Chapter 6 and Chapter 14, Section 2.

5.3.3 Suspension and the Homology of Spheres

The homology of spheres is determined by the graded coefficient groups G_n.

Proposition 5.3.6. For every $m \geq 0$ and every $n \in \mathbb{Z}$,

$$\widetilde{A}_n(\mathrm{S}^m) \cong G_{n-m}, \qquad A_n(\mathrm{S}^m) \cong G_n \oplus G_{n-m}.$$

The reduced group is defined using a chosen point of the sphere.

Proof. For $\mathrm{S}^0 = \{x_0, x_1\}$, additivity identifies $A_n(\mathrm{S}^0)$ with $G_n \oplus G_n$. The map to a point is addition, whose kernel consists of $(a, -a)$ and is isomorphic to G_n.

For $m \geq 1$, cover S^m by the complements U, V of its two poles. Both are contractible, and their intersection deformation retracts onto S^{m-1}. The resulting square is a homotopy pushout by Remark 5.2.6. Under the coefficient identifications, its first map is

$$A_n(\mathrm{S}^{m-1}) \longrightarrow G_n \oplus G_n, \qquad a \longmapsto ((\mathrm{pr}_{\mathrm{S}^{m-1}})_*(a), -(\mathrm{pr}_{\mathrm{S}^{m-1}})_*(a)).$$

Its image is the antidiagonal copy of G_n; its kernel is $\widetilde{A}_n(\mathrm{S}^{m-1})$. Exactness gives

$$0 \longrightarrow G_n \longrightarrow A_n(\mathrm{S}^m) \xrightarrow{\delta_n} \widetilde{A}_{n-1}(\mathrm{S}^{m-1}) \longrightarrow 0.$$

The first homomorphism is the map induced by a chosen point in U, so the map $\mathrm{pr}_{\mathrm{S}^m} : \mathrm{S}^m \longrightarrow *$ splits it. Hence δ_n restricts to an isomorphism from $\widetilde{A}_n(\mathrm{S}^m)$ onto $\widetilde{A}_{n-1}(\mathrm{S}^{m-1})$. Induction proves the formula. □

Example 5.3.7. With $G_0 = \mathbb{Z}$ and $G_n = 0$ for $n \neq 0$, Proposition 5.3.6 gives the usual integral homology of spheres. For the shifted theory of Example 5.2.9, the additional sphere summand occurs in degree $m+1$, and the coefficient summand occurs in degree 1. A connected space can therefore have several nonzero groups unrelated to its dimension.

Homotopy Invariance and First Computations

Algebraic Topology I | Fall 2026 | October 5 and 7

6.1 The Prism Operator and Homotopy Invariance

Singular chains are free abelian on continuous simplices. A triangulation of a simplex times an interval converts a homotopy into a homomorphism of degree one.

6.1.1 Chain Homotopies and the Prism Operator

Use the integer singular chains, their differential, and functoriality from Lecture 4, Definition 4.4.9 and the following paragraphs. We suppress the superscript sing, writing $\mathrm{C}_n(X)$ for the integral singular chains and $\mathrm{H}_n(X;\mathbb{Z})$ for their homology. Write Top for topological spaces and Ab for abelian groups. The affine coface maps of Lecture 4, Definition 4.2.5, are denoted here by $d^i : |\Delta^{n-1}| \to |\Delta^n|$; their vertices remain $e_0, \dots, e_n$. For the classical prism proof of homotopy invariance, see [Hat02, Theorem 2.10 and its proof, pp. 111–113].

Definition 6.1.1. Use chain homotopy from Lecture 1, Definition 1.3.5. For maps $f, g : A_\bullet \to B_\bullet$, a degree-one homotopy K from f to g satisfies $g - f = \partial_B K + K\partial_A$. This endpoint order agrees with the prism construction below.

Proposition 6.1.2. The homology invariance of chain homotopy is Lecture 1, Proposition 1.3.6. In particular, a chain homotopy equivalence induces an isomorphism on homology.

Proof. Apply Lecture 1, Proposition 1.3.6, to the two composites and their homotopies to the identities. □

Write $I = [0,1]$. A homotopy of maps of spaces gives a chain homotopy by triangulating each singular prism.

Theorem 6.1.3. Suppose that

$$f, g : X \longrightarrow Y, \qquad h : X \times I \longrightarrow Y$$

satisfy $h(x,0) = f(x)$ and $h(x,1) = g(x)$. There are homomorphisms

$$P_n : \mathrm{C}_n(X) \longrightarrow \mathrm{C}_{n+1}(Y) \qquad (n \geq 0)$$

with $P_{-1} = 0$ and

$$\partial_{n+1} P_n + P_{n-1}\partial_n = g_{*,n} - f_{*,n} : \mathrm{C}_n(X) \longrightarrow \mathrm{C}_n(Y). \tag{14}$$

Thus singular homology is a homotopy invariant functor in each degree:

$$\mathrm{H}_n(-;\mathbb{Z}) : \mathrm{Top} \longrightarrow \mathrm{Ab}.$$

Proof. The prism operator is constructed in Definition 6.1.5; Proposition 6.1.9 proves (14). For a cycle α, that identity gives

$$g_{*,n}(\alpha) - f_{*,n}(\alpha) = \partial_{n+1} P_n(\alpha).$$

The two images therefore represent the same homology class. □

6.1.2 Cutting the Prism Into Simplices

Fix the base dimension $n \geq 0$. For each prism index $j = 0, \dots, n$, the jth prism simplex has the ordered vertices

$$(e_0, 0), \dots, (e_j, 0), (e_j, 1), \dots, (e_n, 1).$$

Write $s_i = (e_i, 0)$ for the bottom vertices. The affine parametrization is

$$\tau_j : |\Delta^{n+1}| \longrightarrow |\Delta^n| \times I, \qquad \tau_j(e_k) = \begin{cases} (e_k, 0), & 0 \leq k \leq j, \\ (e_{k-1}, 1), & j+1 \leq k \leq n+1. \end{cases}$$

Here k is a vertex position in the prism simplex. The e_k on the left is a standard vertex of the source; those on the right are standard vertices of the base. Affine extension gives the coordinate formula

$$\tau_j : |\Delta^{n+1}| \longrightarrow |\Delta^n| \times I,$$
$$\tau_j(t_0, \dots, t_{n+1}) = \left(\sum_{k=0}^{j} t_k e_k + \sum_{k=j+1}^{n+1} t_k e_{k-1}, \ \sum_{k=j+1}^{n+1} t_k \right).$$

Thus τ_j identifies the two source coordinates in positions $j, j+1$ in the base, and adds the coordinates after position j to obtain the height.

Remark 6.1.4. For a point (x, t) in the prism, now write $x = \sum_{i=0}^{n} t_i e_i$. It belongs to the jth piece exactly when

$$\sum_{i=j+1}^{n} t_i \ \leq t \ \leq \sum_{i=j}^{n} t_i. \tag{15}$$

Indeed, the unique coordinates in the source of τ_j are

$$\left(t_0, \dots, t_{j-1}, \ \sum_{i=j}^{n} t_i - t, \ t - \sum_{i=j+1}^{n} t_i, \ t_{j+1}, \dots, t_n \right).$$

They sum to 1, and they are nonnegative exactly under (15). An empty sum is zero, and an empty string of coordinates is omitted.

For a fixed base point, the endpoints of the intervals in (15) occur in order:

$$0 \leq t_n \leq t_{n-1} + t_n \leq \cdots \leq t_0 + \cdots + t_n = 1.$$

Hence the pieces cover the prism. For $j < j'$, membership in both pieces forces

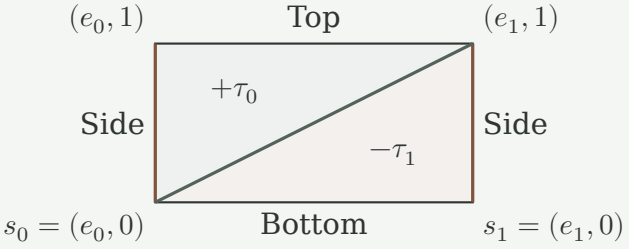
$$t = \sum_{i=j+1}^{n} t_i = \sum_{i=j'}^{n} t_i, \qquad t_{j+1} = \cdots = t_{j'-1} = 0.$$

Their intersection is exactly the common face whose ordered vertices are

$$(e_0, 0), \dots, (e_j, 0), (e_{j'}, 1), \dots, (e_n, 1).$$

Here a face means all linear combinations of the listed vertices with nonnegative coefficients summing to 1. Thus the pieces meet along faces and form a triangulation.

For $n = 1$ the picture is a square divided into two triangles:



The open diagonal is inside the square. The vertical edges lie over the boundary points of the base interval. This distinction persists in every dimension.

Definition 6.1.5. For a singular simplex σ, define σ_j by the composite

$$|\Delta^{n+1}| \xrightarrow{\tau_j} |\Delta^n| \times I \xrightarrow{\sigma \times \mathrm{id}_I} X \times I \xrightarrow{h} Y.$$

Explicitly, the middle map is

$$\sigma \times \mathrm{id}_I : |\Delta^n| \times I \longrightarrow X \times I, \qquad (x,t) \longmapsto (\sigma(x), t).$$

The free abelian universal property extends the following assignment uniquely to a homomorphism:

$$P_n : \mathrm{C}_n(X) \longrightarrow \mathrm{C}_{n+1}(Y), \qquad \sigma \longmapsto \sum_{j=0}^{n} (-1)^j \sigma_j.$$

6.1.3 Prism Faces and Their Cancellation

The boundary of the signed prism consists of its two ends and the prisms on its base faces. This identity gives homotopy invariance and the first absolute computations.

The terms in the boundary are indexed by a Cartesian product:

$$(j,k) \in \{0,\dots,n\} \times \{0,\dots,n+1\}, \qquad \partial_{n+1} P_n(\sigma) = \sum_{j=0}^{n} \sum_{k=0}^{n+1} (-1)^{j+k} \sigma_j d^k.$$

The first index chooses the prism piece; the second deletes its vertex in position k. Each corresponding affine face has type

$$\tau_j d^k : |\Delta^n| \longrightarrow |\Delta^n| \times I.$$

The base vertex e_j appears twice in the piece, in positions j and $j+1$. Every other base vertex appears once.

Remark 6.1.6. For $\epsilon \in \{0,1\}$ the endpoint inclusion of the base is

$$\iota_\epsilon : |\Delta^n| \longrightarrow |\Delta^n| \times I, \qquad x \longmapsto (x, \epsilon).$$

The equalities of affine maps, including their signs, are

$$\tau_0 d^0 = \iota_1 : |\Delta^n| \longrightarrow |\Delta^n| \times I, \qquad (-1)^{0+0} = +1,$$

$$\tau_n d^{n+1} = \iota_0 : |\Delta^n| \longrightarrow |\Delta^n| \times I, \qquad (-1)^{n+n+1} = -1.$$

After composition with the homotopy, these give $g\sigma - f\sigma$.

Remark 6.1.7. For $0 \le j < n$, the two faces

$$\tau_j d^{j+1} = \tau_{j+1} d^{j+1} : |\Delta^n| \longrightarrow |\Delta^n| \times I$$

have the same ordered vertices:

$$(e_0,0), \dots, (e_j,0), (e_{j+1},1), \dots, (e_n,1).$$

They are equal as parametrized simplices, not merely as subsets. The equality is the commutative square

$$\begin{array}{ccc} |\Delta^n| & \xrightarrow{d^{j+1}} & |\Delta^{n+1}| \\ {\scriptstyle d^{j+1}}\downarrow & & \downarrow{\scriptstyle \tau_j} \\ |\Delta^{n+1}| & \xrightarrow[\tau_{j+1}]{} & |\Delta^n| \times I. \end{array}$$

Its relative interior lies inside the prism: all base coordinates are positive and both heights occur. Its boundary may meet the outer boundary.

The coefficients of this same ordered face in the two boundaries are

$$(-1)^j(-1)^{j+1} = -1, \qquad (-1)^{j+1}(-1)^{j+1} = +1.$$

Thus the signed prism simplices induce opposite orientations on the face. Their contributions cancel, including after composition with h.

Remark 6.1.8. Deleting the unique vertex over e_i, with $i \neq j$, gives a side face over $t_i = 0$. Its position is i for $i < j$ and $i+1$ for $i > j$. In the lower-dimensional prism the symbols τ_{j-1} and τ_j use the same affine definition, now with source $|\Delta^n|$ and target $|\Delta^{n-1}| \times I$. The two cases give the squares

$$\begin{array}{ccc} |\Delta^n| & \xrightarrow{d^i} & |\Delta^{n+1}| \\ {\scriptstyle \tau_{j-1}}\downarrow & & \downarrow{\scriptstyle \tau_j} \\ |\Delta^{n-1}| \times I & \xrightarrow[d^i \times \mathrm{id}_I]{} & |\Delta^n| \times I \end{array} \qquad (i < j),$$

$$\begin{array}{ccc} |\Delta^n| & \xrightarrow{d^{i+1}} & |\Delta^{n+1}| \\ {\scriptstyle \tau_j}\downarrow & & \downarrow{\scriptstyle \tau_j} \\ |\Delta^{n-1}| \times I & \xrightarrow[d^i \times \mathrm{id}_I]{} & |\Delta^n| \times I \end{array} \qquad (i > j).$$

Both squares commute by equality of ordered vertices. Omitting e_i changes the repeated vertex index from j to $j-1$ when $i < j$ and leaves it unchanged when $i > j$.

6.1.4 The Prism Boundary Identity

Proposition 6.1.9. For the prism operator of Definition 6.1.5 and every $n \geq 0$, with $P_{-1} = 0$, one has

$$\partial_{n+1}P_n + P_{n-1}\partial_n = g_{*,n} - f_{*,n} : \mathrm{C}_n(X) \longrightarrow \mathrm{C}_n(Y).$$

Proof. Fix a singular n-simplex σ with $n \geq 1$. Expand the two compositions:

$$\partial_{n+1}P_n(\sigma) = \sum_{j=0}^{n}\sum_{k=0}^{n+1}(-1)^{j+k}\sigma_j d^k \qquad \text{Definition 6.1.5 and the boundary formula.}$$

$$P_{n-1}\partial_n(\sigma) = \sum_{i=0}^{n}\sum_{j=0}^{n-1}(-1)^{i+j}(\sigma d^i)_j \qquad i \text{ selects the base face; } j \text{ its prism piece.}$$

Here $(\sigma d^i)_j$ denotes the jth prism simplex of the singular face σd^i. The endpoint terms in the first expansion are

$$+\sigma_0 d^0 = h(\sigma \times \mathrm{id}_I)\iota_1 = g\sigma \qquad (j,k) = (0,0); \text{ Remark 6.1.6.}$$
$$-\sigma_n d^{n+1} = -h(\sigma \times \mathrm{id}_I)\iota_0 = -f\sigma \qquad (j,k) = (n, n+1); \text{ Remark 6.1.6.}$$

The remaining vertices have height 1 and 0, respectively. These are the two full base simplices; the second expansion consists of prisms on base faces.

The table matches every other term. The ordered pairs refer to the prism-piece index and deleted position in ∂P.

Indices in ∂P	Signed term	Opposite signed term
$(j, j+1),\ 0 \le j < n$	$-\sigma_j d^{j+1}$	$+\sigma_{j+1} d^{j+1}$ in ∂P
$(j+1, i),\ 0 \le i \le j < n$	$(-1)^{i+j+1}\sigma_{j+1} d^{i}$	$(-1)^{i+j}(\sigma d^{i})_{j}$ in $P\partial$
$(j, i+1),\ 0 \le j < i \le n$	$(-1)^{i+j+1}\sigma_{j} d^{i+1}$	$(-1)^{i+j}(\sigma d^{i})_{j}$ in $P\partial$

In the first row the paired index is $(j+1, j+1)$, and equality of the two faces is Remark 6.1.7. The last two rows use the commuting squares in Remark 6.1.8; their index ranges partition all (i, j) with $0 \le i \le n$ and $0 \le j < n$. Thus the $(n+1)(n+2)$ terms of ∂P consist of 2 endpoints, $2n$ interior terms, and $n(n+1)$ side terms. The interior terms cancel in pairs; the side terms match all $n(n+1)$ terms of $P\partial$ with opposite signs. Explicitly,

$$
\begin{aligned}
\partial_{n+1} P_n(\sigma) &= g\sigma - f\sigma && \text{endpoint indices } (0,0),(n,n+1).\\
&\quad + \sum_{j=0}^{n-1} \left(-\sigma_j d^{j+1} + \sigma_{j+1} d^{j+1}\right) && \text{interior pairs, first row.}\\
&\quad - \sum_{j=0}^{n-1} \sum_{i=0}^{j} (-1)^{i+j} (\sigma d^{i})_{j} && \text{side indices } (j+1, i), \text{ second row.}\\
&\quad - \sum_{j=0}^{n-1} \sum_{i=j+1}^{n} (-1)^{i+j} (\sigma d^{i})_{j} && \text{side indices } (j, i+1), \text{ third row.}\\
&= g\sigma - f\sigma - P_{n-1}\partial_n(\sigma) && \text{cancel interior pairs and join the side ranges.}
\end{aligned}
$$

For $n = 0$, the prism is one interval and $\partial_1 P_0(\sigma) = g\sigma - f\sigma$, with $P_{-1} = 0$. Adding $P_{n-1}\partial_n$ and extending from generators proves (14) in every degree. □

By the general homotopy-invariance proposition in Lecture 5, homotopy equivalences induce homology isomorphisms.

6.2 Absolute and Relative Homology

6.2.1 Components, Points, and Disjoint Unions

Write $\pi_0^{\mathrm{path}}(X)$ for the set of path components, and $\mathbb{Z}[E]$ for the free abelian group on a set E.

Proposition 6.2.1. There is a natural isomorphism

$$\eta_X : \mathrm{H}_0(X;\mathbb{Z}) \longrightarrow \mathbb{Z}[\pi_0^{\mathrm{path}}(X)], \qquad [x] \longmapsto \text{the generator for the component of } x.$$

Proof. Let B be an abelian group. A homomorphism

$$\alpha : \mathrm{C}_0(X) \longrightarrow B$$

is an assignment of an element of B to each point. The boundary of a path is its endpoint minus its starting point. Thus $\alpha\partial_1 = 0$ exactly when the assignment is constant on path components.

The cokernel property gives the unique factorization

$$\begin{array}{ccccc} \mathrm{C}_1(X) & \xrightarrow{\partial_1} & \mathrm{C}_0(X) & \xrightarrow{q} & \mathrm{H}_0(X;\mathbb{Z}) \\ & & & \searrow^{\alpha} & \downarrow{\scriptstyle \overline{\alpha}} \\ & & & & B. \end{array}$$

This is the free abelian universal property for the set of path components. It proves the isomorphism, and its description on points proves naturality. □

Proposition 6.2.2. Use the homology of a point from Lecture 4, Example 4.4.8: the coefficient group is $\mathbb{Z}$ in degree zero and zero otherwise.

Proof. The singular simplicial set of a point is Δ^0. Apply the constant simplicial group computation of Lecture 4, Example 4.1.7, as specialized in Lecture 4, Example 4.4.8. □

Proposition 6.2.3. For a family $(X_\lambda)_{\lambda\in\Lambda}$, the summand inclusions induce isomorphisms

$$\bigoplus_{\lambda\in\Lambda} \mathrm{C}_\bullet(X_\lambda) \xrightarrow{\cong} \mathrm{C}_\bullet\left(\coprod_{\lambda\in\Lambda} X_\lambda\right), \qquad \bigoplus_{\lambda\in\Lambda} \mathrm{H}_n(X_\lambda;\mathbb{Z}) \xrightarrow{\cong} \mathrm{H}_n\left(\coprod_{\lambda\in\Lambda} X_\lambda;\mathbb{Z}\right).$$

Proof. Each singular simplex has connected nonempty domain, so its image lies in one summand of the disjoint union. The generators therefore split by summand, and every face stays in the same summand. This proves the chain isomorphism. In a direct sum every element has finite support, so both cycles and boundaries are direct sums of the cycles and boundaries in the summands. Taking their quotient proves the homology isomorphism. □

The same argument splits singular chains over path components, even when those components are not open: a simplex has path connected domain and hence lies in one path component, with its subspace topology. In particular, for a discrete space E,

$$\mathrm{H}_0(E;\mathbb{Z}) = \mathbb{Z}[E], \qquad \mathrm{H}_n(E;\mathbb{Z}) = 0 \quad (n>0).$$

For the empty space all chain and homology groups are zero.

6.2.2 Contractible Spaces and Deformation Retractions

Definition 6.2.4. A space is *contractible* if it is homotopy equivalent to the one-point space.

Remark 6.2.5. Equivalently, it is nonempty and its identity is homotopic to a constant map: the map to the point and the inclusion of that constant value are homotopy inverses.

Proposition 6.2.6. For a contractible space X,

$$\mathrm{H}_0(X;\mathbb{Z}) \cong \mathbb{Z}, \qquad \mathrm{H}_n(X;\mathbb{Z}) = 0 \quad (n>0).$$

Proof. Choose a point x_0 to which X contracts. The maps

$$\mathrm{pr}_X : X \longrightarrow *, \qquad g : * \longrightarrow X, \qquad * \longmapsto x_0$$

satisfy $\mathrm{pr}_X\, g = \mathrm{id}_*$ and $g\,\mathrm{pr}_X \simeq \mathrm{id}_X$. Their induced maps

$$(\mathrm{pr}_X)_* : \mathrm{H}_n(X;\mathbb{Z}) \longrightarrow \mathrm{H}_n(*;\mathbb{Z}), \qquad g_* : \mathrm{H}_n(*;\mathbb{Z}) \longrightarrow \mathrm{H}_n(X;\mathbb{Z})$$

are therefore inverse isomorphisms. Apply the computation for a point. □

For a nonempty convex subset $V \subseteq \mathbb{R}^m$ and $x_0 \in V$, the contraction is

$$h : V \times I \longrightarrow V, \qquad (x,t) \longmapsto (1-t)x + tx_0.$$

This applies to Euclidean spaces, balls, intervals, and standard simplices. It also applies when V is star-shaped about x_0, meaning that it contains the line segment joining x_0 to each of its points.

Definition 6.2.7. A *deformation retraction* consists of an inclusion and a retraction

$$\iota : X' \hookrightarrow X, \qquad r : X \longrightarrow X'$$

and a homotopy $h : X \times I \longrightarrow X$ satisfying

$$r\iota = \mathrm{id}_{X'}, \qquad h(x,0) = x, \qquad h(x,1) = \iota r(x), \qquad h(\iota(x'),t) = \iota(x').$$

Then ι and r are homotopy inverse maps, so

$$\iota_* : \mathrm{H}_n(X';\mathbb{Z}) \longrightarrow \mathrm{H}_n(X;\mathbb{Z}), \qquad r_* : \mathrm{H}_n(X;\mathbb{Z}) \longrightarrow \mathrm{H}_n(X';\mathbb{Z})$$

are inverse isomorphisms. For example,

$$r : \mathbb{R} - \{0\} \longrightarrow \{-1,1\}, \qquad x \longmapsto x/|x|,$$

$$h : (\mathbb{R} - \{0\}) \times I \longrightarrow \mathbb{R} - \{0\}, \qquad (x,t) \longmapsto (1-t)x + t\,x/|x|$$

give a deformation retraction. The interpolation remains in the half-line containing x. Thus the homology is $\mathbb{Z}^2$ in degree zero and zero in positive degrees. Removing m distinct points from $\mathbb{R}$ similarly leaves $m+1$ contractible components.

6.2.3 Pairs, Quotients, and the Exact Sequence

Quotient chain complexes give the relative exact sequence. Homotopies preserving the subspaces induce chain homotopies on these quotients.

Definition 6.2.8. A pair is a subspace inclusion

$$\iota : X' \hookrightarrow X.$$

Write (X, X') for this pair. The induced inclusion identifies the singular chains of X' with a subcomplex. Define

$$\mathrm{C}_\bullet(X,X') = \mathrm{C}_\bullet(X)/\mathrm{C}_\bullet(X'), \qquad \mathrm{H}_n(X,X';\mathbb{Z}) = \mathrm{H}_n(\mathrm{C}_\bullet(X,X')).$$

A map of pairs is a commutative square

$$\begin{array}{ccc} X' & \xhookrightarrow{\iota_X} & X \\ {\scriptstyle f'}\downarrow & & \downarrow{\scriptstyle f} \\ Y' & \xhookrightarrow[\iota_Y]{} & Y. \end{array}$$

The relative chain complex is a cokernel. For any chain complex $B_\bullet$ and chain map

$$\alpha : \mathrm{C}_\bullet(X) \longrightarrow B_\bullet, \qquad \alpha\iota_* = 0,$$

there is a unique induced map

$$\overline{\alpha} : \mathrm{C}_\bullet(X,X') \longrightarrow B_\bullet$$

making the right square commute; the left square is a pushout:

$$\begin{array}{ccccc} \mathrm{C}_\bullet(X') & \xrightarrow{\iota_*} & \mathrm{C}_\bullet(X) & \xrightarrow{\alpha} & B_\bullet \\ \downarrow & & \downarrow{\scriptstyle q} & & \| \\ 0 & \longrightarrow & \mathrm{C}_\bullet(X,X') & \xrightarrow[\overline{\alpha}]{} & B_\bullet. \end{array}$$

For a map of pairs this constructs the relative chain map in the square

$$\begin{array}{ccc} \mathrm{C}_\bullet(X) & \xrightarrow{q_X} & \mathrm{C}_\bullet(X,X') \\ {\scriptstyle f_*}\downarrow & & \downarrow{\scriptstyle f_*} \\ \mathrm{C}_\bullet(Y) & \xrightarrow[q_Y]{} & \mathrm{C}_\bullet(Y,Y'). \end{array}$$

Uniqueness gives compatibility with identities and composition, hence functoriality.

A relative cycle is represented by $\alpha \in \mathrm{C}_n(X)$ with $\partial_n \alpha \in \mathrm{C}_{n-1}(X')$. Two such representatives have the same relative homology class exactly when their difference is a boundary in X plus a chain in X'.

Theorem 6.2.9. Every pair has a natural long exact sequence with five consecutive terms

$$\mathrm{H}_n(X';\mathbb{Z}) \xrightarrow{\iota_*} \mathrm{H}_n(X;\mathbb{Z}) \xrightarrow{q_*} \mathrm{H}_n(X,X';\mathbb{Z}) \xrightarrow{\delta_n} \mathrm{H}_{n-1}(X';\mathbb{Z}) \xrightarrow{\iota_*} \mathrm{H}_{n-1}(X;\mathbb{Z})$$

ending in

$$\mathrm{H}_0(X';\mathbb{Z}) \xrightarrow{\iota_*} \mathrm{H}_0(X;\mathbb{Z}) \xrightarrow{q_*} \mathrm{H}_0(X,X';\mathbb{Z}) \longrightarrow 0.$$

The connecting homomorphism is

$$\delta_n : \mathrm{H}_n(X,X';\mathbb{Z}) \longrightarrow \mathrm{H}_{n-1}(X';\mathbb{Z}), \qquad [\alpha] \longmapsto [\partial_n \alpha].$$

Proof. Apply the long exact sequence of chain complexes from Lecture 2 to

$$0 \longrightarrow \mathrm{C}_\bullet(X') \xrightarrow{\iota_*} \mathrm{C}_\bullet(X) \xrightarrow{q} \mathrm{C}_\bullet(X,X') \longrightarrow 0.$$

For a relative cycle, $\partial_n \alpha$ lies in the subcomplex and is a cycle because $\partial_{n-1}\partial_n = 0$. Replacing α by $\alpha + \partial_{n+1}\beta + \alpha'$ with $\alpha' \in \mathrm{C}_n(X')$ changes its boundary by $\partial_n \alpha'$, so the connecting class is well defined. The chain-complex theorem supplies exactness and naturality. □

Example 6.2.10. Take $X = I$ and $X' = \partial I = \{0,1\}$. Contractibility of I and the computation for discrete spaces reduce the exact sequence to

$$0 \longrightarrow \mathrm{H}_1(I,\partial I;\mathbb{Z}) \xrightarrow{\delta_1} \mathbb{Z}^2 \xrightarrow{\varepsilon} \mathbb{Z} \longrightarrow \mathrm{H}_0(I,\partial I;\mathbb{Z}) \longrightarrow 0,$$

where

$$\varepsilon : \mathbb{Z}^2 \longrightarrow \mathbb{Z}, \qquad (a,b) \longmapsto a+b.$$

Both endpoint generators map to the component generator of I. The map is surjective, and its kernel is $\mathbb{Z}(-1,1)$. Thus

$$\mathrm{H}_n(I,\partial I;\mathbb{Z}) = \begin{cases} \mathbb{Z}, & n = 1, \\ 0, & n \neq 1. \end{cases}$$

The singular edge

$$\sigma : |\Delta^1| \longrightarrow I, \qquad (t_0,t_1) \longmapsto t_1$$

has boundary equal to the point generator at 1 minus the point generator at 0. Thus $\delta_1[\sigma] = (-1, 1)$ and $[\sigma]$ is a generator.

Reflection is the map of pairs induced by

$$\rho : I \longrightarrow I, \qquad t \longmapsto 1 - t.$$

Its restriction to the endpoints is

$$\rho' : \partial I \longrightarrow \partial I, \qquad t \longmapsto 1 - t.$$

Naturality gives

$$\begin{array}{ccc} \mathrm{H}_1(I, \partial I; \mathbb{Z}) & \xrightarrow{\delta_1} & \mathrm{H}_0(\partial I; \mathbb{Z}) \\ \downarrow{\scriptstyle \rho_*} & & \downarrow{\scriptstyle \rho'_*} \\ \mathrm{H}_1(I, \partial I; \mathbb{Z}) & \xrightarrow[\delta_1]{} & \mathrm{H}_0(\partial I; \mathbb{Z}). \end{array}$$

The right map exchanges the endpoint generators, sending $(-1, 1)$ to its negative. Injectivity of δ_1 therefore gives $\rho_*([\sigma]) = -[\sigma]$.

6.2.4 Relative Homotopies and Reduced Homology

A homotopy preserving the subspaces descends to the quotient chain complexes.

Proposition 6.2.11. Suppose that

$$f, g : (X, X') \longrightarrow (Y, Y')$$

are joined by a homotopy

$$h : X \times I \longrightarrow Y, \qquad h(X' \times I) \subseteq Y'.$$

Then the induced maps agree:

$$f_* = g_* : \mathrm{H}_n(X, X'; \mathbb{Z}) \longrightarrow \mathrm{H}_n(Y, Y'; \mathbb{Z}).$$

Proof. For a simplex in X', all its prism simplices have image in Y'. Hence the prism homomorphisms restrict and descend:

$$P_n : \mathrm{C}_n(X') \longrightarrow \mathrm{C}_{n+1}(Y'), \qquad \overline{P}_n : \mathrm{C}_n(X, X') \longrightarrow \mathrm{C}_{n+1}(Y, Y').$$

Identity (14) descends to the quotient. It makes the relative chain maps chain homotopic, so they induce equal homology maps. □

The subspace condition is essential. The contraction

$$h : I \times I \longrightarrow I, \qquad (x, t) \longmapsto (1 - t)x$$

moves the endpoint 1 into the interior for $0 < t < 1$. It is therefore not a homotopy of the pair $(I, \partial I)$.

Definition 6.2.12. For a nonempty space X, the *augmentation* is the chain map

$$\varepsilon : \mathrm{C}_\bullet(X) \longrightarrow \mathbb{Z}[0],$$

whose degree-zero component is

$$\varepsilon_0 : \mathrm{C}_0(X) \longrightarrow \mathbb{Z}, \qquad \sum_i a_i x_i \longmapsto \sum_i a_i,$$

and whose positive-degree components are zero. Here $\mathbb{Z}[0]$ is concentrated in degree zero. Every path boundary has coefficient sum zero, so this is a chain map. It is surjective because X is nonempty.

Definition 6.2.13. The *reduced chain complex* of a nonempty space X is the kernel of its augmentation, equivalently the pullback:

$$\begin{array}{ccc} \widetilde{\mathrm{C}}_\bullet(X) & \longrightarrow & \mathrm{C}_\bullet(X) \\ \downarrow & & \downarrow{\scriptstyle \varepsilon} \\ 0 & \longrightarrow & \mathbb{Z}[0]. \end{array}$$

Its homology is the *reduced homology* $\widetilde{\mathrm{H}}_n(X;\mathbb{Z})$.

The kernel exact sequence gives

$$\widetilde{\mathrm{H}}_n(X;\mathbb{Z}) \cong \mathrm{H}_n(X;\mathbb{Z}) \quad (n>0), \qquad \widetilde{\mathrm{H}}_0(X;\mathbb{Z}) = \ker(\varepsilon_* : \mathrm{H}_0(X;\mathbb{Z}) \longrightarrow \mathbb{Z}).$$

The prism homomorphisms restrict to the reduced complexes, so reduced homology is homotopy invariant.

If X is contractible and X' is nonempty, the exact sequence of the pair gives isomorphisms

$$\delta_n : \mathrm{H}_n(X,X';\mathbb{Z}) \xrightarrow{\cong} \widetilde{\mathrm{H}}_{n-1}(X';\mathbb{Z}) \quad (n\geq 1), \qquad \mathrm{H}_0(X,X';\mathbb{Z}) = 0.$$

For $n \geq 2$ both adjacent positive-degree groups of X vanish. For $n = 1$ the image of δ_1 is the kernel of the augmentation on $\mathrm{H}_0(X';\mathbb{Z})$.

For example, let X' be a finite discrete subspace consisting of $m \geq 1$ points in a contractible space X. Then

$$\mathrm{H}_1(X,X';\mathbb{Z}) \cong \ker(\varepsilon : \mathbb{Z}^m \longrightarrow \mathbb{Z}) \cong \mathbb{Z}^{m-1},$$

where ε sums the coordinates; all other relative groups vanish. Choose one marked point and paths from it to the others. Their relative classes form a basis because their boundaries are a basis of the augmentation kernel. In particular, three marked points in an interval or a disk give relative first homology $\mathbb{Z}^2$.

For an open cover, the small-chain subdivision theorem identifies the homology of chains supported in individual cover members with singular homology. The resulting short exact sequence of chain complexes induces the Mayer–Vietoris sequence.

Homotopy Pushouts and Mayer–Vietoris

Algebraic Topology I | Fall 2026 | October 12 and 14

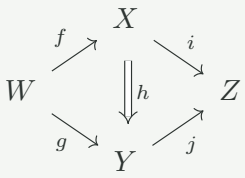
7.1 Homotopy Pushouts and Their Chain Models

7.1.1 Double Mapping Cylinders and Their Universal Property

Write $I = [0,1]$. For a space W, write $\mathrm{pr}_W : W \to *$ for its unique map to a point. We use Top for the category of topological spaces and continuous maps. As in Lecture 5, homotopy pushout means that the comparison is a homotopy equivalence.

Definition 7.1.1. Use the double mapping cylinder $\mathrm{Cyl}(f,g)$ of Lecture 5, Definition 5.1.7. Write ι_X, ι_Y for its canonical maps and $\eta(w,t) = [w,t]$ for its canonical homotopy from $\iota_X f$ to $\iota_Y g$.

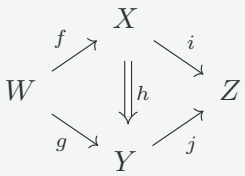
Proposition 7.1.2. Use the quotient universal property following Lecture 5, Definition 5.1.7. In the present notation, a map $\mathrm{Cyl}(f,g) \to Z$ is specified by the diagram



naturally in Z.

Proof. The quotient argument is given after Lecture 5, Definition 5.1.7; postcomposition gives naturality in Z. □

Definition 7.1.3. Use homotopy pushout squares as in Lecture 5, Definition 5.1.8, in the form



The square is equipped with a homotopy h from if to jg. Its comparison map is $\theta_h : \mathrm{Cyl}(f,g) \to Z$.

The specified homotopy is part of this definition. Even when $if = jg$, a nonconstant chosen homotopy can change the comparison map. For a strictly commutative square, the constant homotopy is the usual choice.

A map out of the ordinary pushout $X \underset{W}{+} Y$ requires equality of the two composites. A map out of $\mathrm{Cyl}(f,g)$ requires a chosen homotopy. Passing to homotopy classes therefore does not turn this construction into an ordinary pushout in the homotopy category: the homotopy between the composites remains part of the mapping data.

Example 7.1.4. Take $X = Y = *$, where $*$ is a one-point space. For the two maps pr_W, the double mapping cylinder $\mathrm{Cyl}(\mathrm{pr}_W, \mathrm{pr}_W)$ is the *unreduced suspension*

$$\mathrm{S}(W) = ((W \times I) + \{*_0\} + \{*_1\})/((w,0) \sim *_0,\ (w,1) \sim *_1).$$

Here $*_0$ and $*_1$ are distinct vertices. In particular, $\mathrm{S}(\mathrm{S}^0)$ is a circle: its two intervals have the same two endpoints. The ordinary pushout of $* \xleftarrow{\mathrm{pr}_{\mathrm{S}^0}} \mathrm{S}^0 \xrightarrow{\mathrm{pr}_{\mathrm{S}^0}} *$ is a point. Thus an ordinary pushout need not be a homotopy pushout. More generally, $\mathrm{S}(\mathrm{S}^m) \cong \mathrm{S}^{m+1}$ for $m \geq 0$.

7.1.2 Cofibrations and Geometric Gluing

A Hurewicz cofibration permits a homotopy on a subspace to extend to the whole space. This extension makes ordinary gluing agree with homotopy gluing. The comparison below is the double-mapping-cylinder lemma of May, Chapter 10, §7, p. 80; we give the extension argument explicitly.

Definition 7.1.5. A continuous map $f : X \longrightarrow Y$ is a *Hurewicz cofibration* if it has the homotopy extension property: for every space Z, every map $h : Y \longrightarrow Z$, and every homotopy $\alpha : X \times I \longrightarrow Z$ starting at hf, there is a homotopy $\beta : Y \times I \longrightarrow Z$ starting at h satisfying $\beta(f(x), t) = \alpha(x, t)$ for all $(x, t) \in X \times I$.

For a subspace inclusion, equivalently, the homotopy extension problem is the extension in the diagram

$$\begin{array}{ccc} (Y \times \{0\}) \underset{X \times \{0\}}{\cup} (X \times I) & \longrightarrow & Z \\ \downarrow & \nearrow_{\beta} & \\ Y \times I & & \end{array}$$

for every continuous map on the upper-left subspace. We prove the disk and finite-cell cases used below; see also May, Chapter 6, §1, and Chapter 10, §1.

Regard D^n as the unit disk in $\mathbb{R}^n$ and put $\mathrm{S}^{-1} = \emptyset$. For $(x, t) \in \mathrm{D}^n \times I$, define

$$m(x,t) = \max\{\|x\|, 1 - t/2\},$$
$$\rho(x,t) = \left(\frac{x}{m(x,t)},\ 2 + \frac{t-2}{m(x,t)} \right).$$

The denominator lies in $[1/2, 1]$, the first coordinate has norm at most one, and the second lies in $[0, t]$. According to which term attains the maximum, the image lies on the side or the bottom of the cylinder. On either of those subspaces $m = 1$, so ρ is a continuous retraction onto

$$(\mathrm{D}^n \times \{0\}) \cup (\mathrm{S}^{n-1} \times I).$$

For $n = 0$, it retracts the interval onto its bottom endpoint. A map on the disk and a homotopy on its boundary paste continuously on this union when they agree at time zero. Composing with ρ proves the homotopy extension property for $\mathrm{S}^{n-1} \hookrightarrow \mathrm{D}^n$.

If a disk is attached to Y' through $f : \mathrm{S}^{n-1} \to Y'$, pull the given homotopy on Y' back along f, extend it over the disk by this construction, and glue. Continuity follows because taking the product of the attaching pushout with I preserves the pushout, by Lemma 7.1.7 below. Repetition proves the extension property for any finite sequence of disk attachments. In particular, a subcomplex inclusion $X' \hookrightarrow X$ in a finite cell complex has this property: attach the remaining cells in increasing dimension, so that their boundaries are already present.

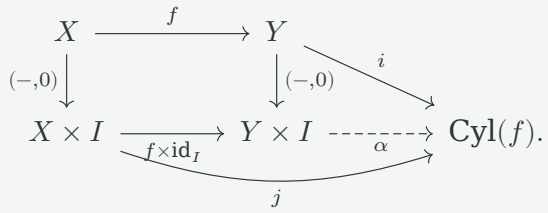
Proposition 7.1.6. Let $f : X \hookrightarrow Y$ be a Hurewicz cofibration and $g : X \longrightarrow Z$ any continuous map. Then the quotient map

$$q : \mathrm{Cyl}(f, g) \longrightarrow Y \underset{X}{+} Z$$

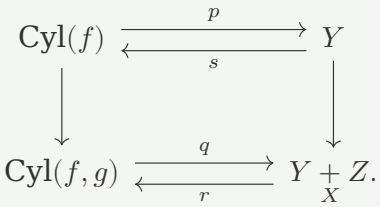
given on the cylinder by $[x, t] \mapsto [f(x)] = [g(x)]$ is a homotopy equivalence. Here $[x, t]$ denotes the image of $(x, t) \in X \times I$ under the cylinder structure map to the double mapping cylin-

der. Consequently, the ordinary pushout square, with its constant homotopy, is a homotopy pushout.

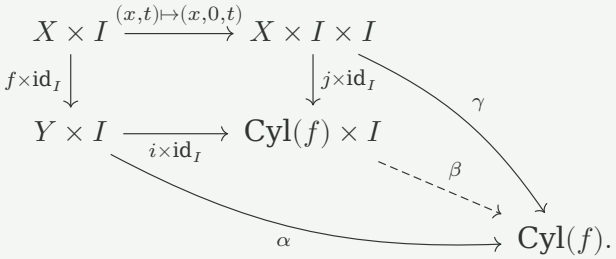
Proof. Let $i : Y \to \mathrm{Cyl}(f)$ and $j : X \times I \to \mathrm{Cyl}(f)$ be the canonical maps, and put $k(x) = j(x, 1)$. The collapse $p : \mathrm{Cyl}(f) \to Y$ satisfies $pi = \mathrm{id}_Y$ and $pj(x, u) = f(x)$. The homotopy extension property supplies the dashed extension:



Set $s(y) = \alpha(y, 1)$. Since $sf = k$ and $pk = f$, the maps s, p glue to the identity on Z, giving the lower maps in the two commuting squares below



The homotopy $p\alpha$ joins id_Y to ps relative to $f(X)$. To construct the corresponding relative homotopy on $\mathrm{Cyl}(f)$, glue α to $\gamma(x, u, t) = j(x, u + t(1 - u))$ through the cylinder pushout:



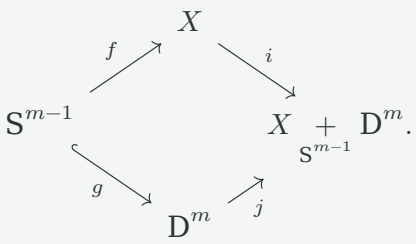
The restrictions agree because $\alpha(f(x), t) = j(x, t) = \gamma(x, 0, t)$. Thus β joins $\mathrm{id}_{\mathrm{Cyl}(f)}$ to sp relative to $k(X)$: at $u = 1$, the map γ is constantly $k(x)$. Gluing $p\alpha$ and β to the constant homotopy on Z gives $\mathrm{id} \simeq qr$ and $\mathrm{id} \simeq rq$.

These gluings are continuous by Lemma 7.1.7. □

Lemma 7.1.7. If $q : X \to Y$ is a quotient map, then $q \times \mathrm{id}_I : X \times I \to Y \times I$ is a quotient map. Consequently, taking the product with I preserves pushouts.

Proof. Suppose $U \subseteq Y \times I$ has open inverse image. For $(y, t) \in U$, its slice U_y is open: choose $x \in q^{-1}(y)$ and restrict that inverse image to $\{x\} \times I$. Choose a compact interval $K \subseteq U_y$ with $t \in \mathrm{int}_I K$, and put $V = \{y' \in Y : \{y'\} \times K \subseteq U\}$. For each $x \in q^{-1}(V)$, compactness of K gives a neighborhood of x whose product with K lies in the inverse image of U. Thus $q^{-1}(V)$ is open, so V is open by the quotient property. The neighborhood $V \times \mathrm{int}_I K$ of (y, t) lies in U, proving the first assertion. A pushout is the quotient of a disjoint union by the gluing identifications; products with I preserve that quotient and distribute over the disjoint union. □

Example 7.1.8. For a map $f : \mathrm{S}^{m-1} \longrightarrow X$, attaching an m-cell is the pushout



The inclusion $g : \mathrm{S}^{m-1} \hookrightarrow \mathrm{D}^m$ is a Hurewicz cofibration. Proposition 7.1.6, with the two legs interchanged, makes this a homotopy pushout. The same conclusion holds for gluing finite cell complexes along a common subcomplex. Taking both legs to be $\mathrm{S}^{m-1} \hookrightarrow \mathrm{D}^m$ gives S^m, obtained by gluing two disks along their boundary.

Example 7.1.9. For $f : W \longrightarrow X$, the space $\mathrm{Cyl}(f, \mathrm{pr}_W)$ is the *mapping cone*

$$\mathrm{Cone}(f) = (X + (W \times I) + \{*\})/((w,0) \sim f(w),\ (w,1) \sim *).$$

Thus the square $W \longrightarrow X$, $* \longrightarrow \mathrm{Cone}(f)$, with its cylinder homotopy, is a homotopy pushout. If W is nonempty and f is a Hurewicz cofibration, Proposition 7.1.6 gives $\mathrm{Cone}(f) \simeq X/W$, where X/W is defined by $w \sim *$ for all $w \in W$. For the identity map of W, this cone is contractible. When $W = \emptyset$, the mapping cone is X_+.

7.1.3 Subdivision

Singular chains preserve homotopy pushouts. We prove this first over $\mathbb{Z}$, using small chains, and then tensor with arbitrary coefficients. Mayer–Vietoris and excision follow from this preservation theorem.

For a space Z, write $\mathrm{C}_\bullet(Z)$ for its singular chain complex with integer coefficients. For an open cover $Z = U \cup V$, put $W = U \cap V$ and let $\mathrm{C}_\bullet^{\{U,V\}}(Z)$ denote the subcomplex of $\mathrm{C}_\bullet(Z)$ generated by singular simplices whose images lie in U or V. Faces of such simplices remain in the same set, so this is a subcomplex.

The small-chains theorem below makes the passage from an open cover to singular homology possible. We prove both the subdivision construction and the chain homotopy equivalence; the same construction will be used for excision. For the classical small-chains argument using subdivision, see [Hat02, Proposition 2.21 and its proof, pp. 119–123].

Definition 7.1.10. A map f between convex subsets of Euclidean spaces is *affine* if it preserves convex combinations: for any points $x_0, \dots, x_m$ in its domain and real numbers $t_i \geq 0$ with $\sum_i t_i = 1$, it satisfies $f(\sum_i t_i x_i) = \sum_i t_i f(x_i)$.

Construction 7.1.11 (Barycentric subdivision). Use the standard topological simplex

$$|\Delta^n| = \{(t_0, \dots, t_n) \in \mathbb{R}^{n+1} : t_i \geq 0,\ \textstyle\sum_i t_i = 1\}.$$

The affine face inclusion $d^i : |\Delta^{n-1}| \to |\Delta^n|$ inserts a zero in the ith coordinate.[1] The affine map $s^i : |\Delta^{n+1}| \to |\Delta^n|$ adds the ith and $(i+1)$st coordinates. These maps make $|\Delta^\bullet|$ cosimplicial.

The simplicial set $\mathrm{Mor}_{\mathrm{Top}}(|\Delta^\bullet|, Z)$ has set of n-simplices $\mathrm{Mor}_{\mathrm{Top}}(|\Delta^n|, Z)$. The singular-chain differential is $\partial\sigma = \sum_i (-1)^i \sigma \circ d^i$.

For a convex subset E of a Euclidean space and an affine simplex $\tau : |\Delta^n| \to E$, denote its vertex images by $e_0, \dots, e_n$. The ordered affine simplex $[e_0, \dots, e_n]$ means the map $(t_0, \dots, t_n) \mapsto$

[1] This is the map induced by the simplicial coface under geometric realization; we use the same notation d^i.

$\sum_i t_i e_i$; the order is part of the map. For a nonempty set $I \subseteq \{0, \dots, n\}$, define its *barycenter* by

$$e_I = \frac{1}{|I|} \sum_{i \in I} e_i.$$

In particular, the full barycenter is $e_{0\cdots n} = \tau(1/(n+1), \dots, 1/(n+1))$. This definition uses the indexed vertices and remains valid when some vertex images coincide. Thus e_{02} is the midpoint of the face with indices $0, 2$, and $e_{123} = (e_1 + e_2 + e_3)/3$ is the barycenter of the face with indices $1, 2, 3$. Points e_I and ordered simplices $[e_{i_0}, \dots, e_{i_k}]$ are different kinds of objects.

For a convex subset E of a Euclidean space, work with the augmented chain complex whose degree-n group is the free abelian group

$$\mathbb{Z}\{\tau \in \mathrm{Mor}(|\Delta^n|, E) \mid \tau \text{ is affine}\} \qquad (n \geq 0),$$

with differential induced by the cofaces. Its group in degree -1 is $\mathbb{Z}$, and the augmentation sends each 0-simplex to 1. For a point $e \in E$, coning is the degree-one homomorphism

$$e \star - : \mathbb{Z}\{\tau : |\Delta^n| \to E \text{ affine}\} \longrightarrow \mathbb{Z}\{\tau : |\Delta^{n+1}| \to E \text{ affine}\} \quad (n \geq 0).$$

It joins an affine simplex τ to e by straight line segments, by the formula

$$(e \star \tau)(t_0, \dots, t_{n+1}) = t_0 e + (1 - t_0)\tau\left(\frac{t_1}{1-t_0}, \dots, \frac{t_{n+1}}{1-t_0}\right),$$

with value e at $t_0 = 1$. In degree -1, the homomorphism $e \star - : \mathbb{Z} \to \mathbb{Z}\{\tau : |\Delta^0| \to E \text{ affine}\}$ sends 1 to the constant 0-simplex at e. Extend the cone formula linearly in every degree.

For $n > 0$, the face of $e \star \tau$ indexed by 0 is τ, and the face indexed by $i+1$ is $e \star (\tau \circ d^i)$. Hence

$$\partial(e \star \tau) = \tau + \sum_{i=0}^{n} (-1)^{i+1} e \star (\tau \circ d^i) = \tau - e \star \partial\tau.$$

For $n = 0$, the boundary is $\tau - (e \star 1)$, and in degree -1 it is $\partial(e \star 1) = 1$. Thus the same formula holds on all augmented chains by linearity.

For every $n \geq 0$, we construct homomorphisms

$$\begin{aligned} \mathrm{Sd} &: \mathbb{Z}\{\tau : |\Delta^n| \to E \text{ affine}\} \longrightarrow \mathbb{Z}\{\tau : |\Delta^n| \to E \text{ affine}\}, \\ T &: \mathbb{Z}\{\tau : |\Delta^n| \to E \text{ affine}\} \longrightarrow \mathbb{Z}\{\tau : |\Delta^{n+1}| \to E \text{ affine}\}, \end{aligned}$$

of degrees zero and one, respectively, with the augmented degree -1 treated below. We define Sd and T inductively. For the base cases, define $\mathrm{Sd} = \mathrm{id}$ and $T = 0$ in degrees -1 and 0. These choices satisfy the required identities: for every chain α of degree 0 or -1,

$$\partial \mathrm{Sd}(\alpha) = \partial\alpha = \mathrm{Sd}(\partial\alpha), \qquad \partial T(\alpha) + T(\partial\alpha) = 0 = \alpha - \mathrm{Sd}(\alpha).$$

Here ∂ in degree zero is the augmentation and in degree -1 is zero.

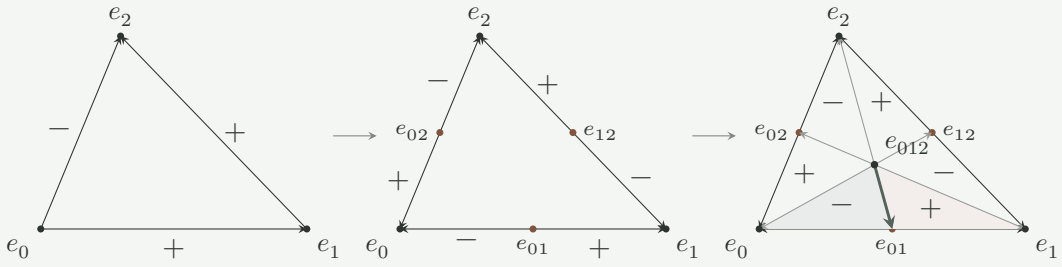
For the inductive step, fix $n > 0$ and assume that Sd and T have been defined in degrees $-1, 0, \dots, n-1$. For every chain α of degree less than n, including degree -1, assume

$$\partial \mathrm{Sd}(\alpha) = \mathrm{Sd}(\partial\alpha), \qquad \partial T(\alpha) + T(\partial\alpha) = \alpha - \mathrm{Sd}(\alpha).$$

For an affine n-simplex τ, put $e = e_{0\cdots n}$, its full barycenter. Define $\mathrm{Sd}(\tau)$ first and then $T(\tau)$ by

$$\mathrm{Sd}(\tau) = e \star \mathrm{Sd}(\partial\tau), \qquad T(\tau) = e \star (\tau - \mathrm{Sd}(\tau) - T(\partial\tau)).$$

For $\tau = [e_0, e_1, e_2]$, the recursion is pictured below, with full barycenter e_{012}. An arrow records the order of the vertices of an affine 1-simplex; its sign is its coefficient in the chain.



For example, the bottom edge contributes $[e_{01}, e_1] - [e_{01}, e_0]$ to $\mathrm{Sd}(\partial\tau)$. Coning gives the two shaded terms $[e_{012}, e_{01}, e_1] - [e_{012}, e_{01}, e_0]$. Their boundaries contain $+[e_{012}, e_{01}]$ and $-[e_{012}, e_{01}]$, so the highlighted interior edge cancels as the *same parametrized simplex* with opposite coefficients. The remaining interior terms cancel in the same way. The six terms are indexed by flags $\{e_i\} \subset [e_i, e_j] \subset \tau$; their vertex order is $[e_{012}, e_{ij}, e_i]$, with the coefficient inherited from $\mathrm{Sd}(\partial\tau)$. Thus the picture expresses the cone identity $\partial(e_{012} \star \mathrm{Sd}(\partial\tau)) = \mathrm{Sd}(\partial\tau) - e_{012} \star \partial\mathrm{Sd}(\partial\tau)$; the next proof verifies that the last term vanishes in every degree.

Proposition 7.1.12. On the augmented chain complex of affine simplices in E from Construction 7.1.11, Sd is a chain map from this complex to itself, and T is a chain homotopy between Sd and the identity. Explicitly,

$$\partial\mathrm{Sd} = \mathrm{Sd}\partial, \qquad \partial T + T\partial = \mathrm{id} - \mathrm{Sd}.$$

Both Sd and T commute with affine maps.

Proof. The identities hold in degrees $-1, 0$ by the computation in the construction. Assume they hold in degrees less than n. Each face $\tau \circ d^i$ has degree $n-1$. The induction hypothesis applies to these faces, without assuming the identity on τ:

$$\begin{aligned}\partial(\mathrm{Sd}(\partial\tau)) &= \sum_{i=0}^{n} (-1)^i \partial(\mathrm{Sd}(\tau \circ d^i)) \\ &= \sum_{i=0}^{n} (-1)^i \mathrm{Sd}(\partial(\tau \circ d^i)) = \mathrm{Sd}(\partial^2\tau) = 0.\end{aligned}$$

The boundary formula for coning gives

$$\partial(e \star \alpha) = \alpha - e \star \partial\alpha, \qquad \partial\alpha = 0 \implies \partial(e \star \alpha) = \alpha.$$

Apply this to $\mathrm{Sd}(\partial\tau)$ and then to $\tau - \mathrm{Sd}(\tau) - T(\partial\tau)$, using the lower-degree homotopy identity:

$$\begin{aligned}\partial\mathrm{Sd}(\tau) &= \mathrm{Sd}(\partial\tau) - e \star \partial\mathrm{Sd}(\partial\tau) = \mathrm{Sd}(\partial\tau), \\ \partial T(\partial\tau) &= (\mathrm{id} - \mathrm{Sd})(\partial\tau) - T(\partial^2\tau) = \partial\tau - \mathrm{Sd}(\partial\tau), \\ \partial(\tau - \mathrm{Sd}(\tau) - T(\partial\tau)) &= \partial\tau - \mathrm{Sd}(\partial\tau) - (\partial\tau - \mathrm{Sd}(\partial\tau)) = 0, \\ \partial T(\tau) &= \partial(e \star (\tau - \mathrm{Sd}(\tau) - T(\partial\tau))) \\ &= (\tau - \mathrm{Sd}(\tau) - T(\partial\tau)) - e \star \partial(\tau - \mathrm{Sd}(\tau) - T(\partial\tau)) \\ &= \tau - \mathrm{Sd}(\tau) - T(\partial\tau).\end{aligned}$$

In the last equality, the term $e \star \partial(\tau - \mathrm{Sd}(\tau) - T(\partial\tau))$ vanishes because the third line shows that the chain in parentheses has zero boundary, and $e \star 0 = 0$ by linearity. That cycle calculation

uses the chain-map identity just proved on τ, the inductive homotopy identity on the lower-degree chain $\partial\tau$, and $\partial^2\tau = 0$. Thus $\partial\mathrm{Sd} = \mathrm{Sd}\partial$ and $\partial T + T\partial = \mathrm{id} - \mathrm{Sd}$ in degree n as well. If $h : E \to E'$ is affine, then $h(e_I) = |I|^{-1}\sum_{i\in I} h(e_i)$ and $h_*(e \star \alpha) = h(e) \star h_*(\alpha)$. These identities, the recursion, and induction on dimension give $h_*\mathrm{Sd}(\tau) = \mathrm{Sd}(h_*\tau)$ and $h_*T(\tau) = T(h_*\tau)$. □

On singular chains, define the degree-zero and degree-one homomorphisms

$$\mathrm{Sd} : \mathrm{C}_n(Z;\mathbb{Z}) \longrightarrow \mathrm{C}_n(Z;\mathbb{Z}), \qquad T : \mathrm{C}_n(Z;\mathbb{Z}) \longrightarrow \mathrm{C}_{n+1}(Z;\mathbb{Z}) \quad (n \geq 0)$$

by the following formulas on each singular n-simplex $\sigma : |\Delta^n| \to Z$: $\mathrm{Sd}(\sigma) = \sigma_*\mathrm{Sd}(\mathrm{id}_{|\Delta^n|})$ and $T(\sigma) = \sigma_*T(\mathrm{id}_{|\Delta^n|})$. Here σ_* means postcomposition of affine simplices in $|\Delta^n|$ with the continuous map σ; σ need not be affine. Put $\iota_n = \mathrm{id}_{|\Delta^n|}$. For $n > 0$, affine naturality for d^i gives

$$\partial\mathrm{Sd}(\sigma) = \sigma_*\mathrm{Sd}(\partial\iota_n) = \sum_{i=0}^{n}(-1)^i(\sigma \circ d^i)_*\mathrm{Sd}(\iota_{n-1}) = \mathrm{Sd}(\partial\sigma),$$
$$\partial T(\sigma) = \sigma_*(\iota_n - \mathrm{Sd}(\iota_n) - T(\partial\iota_n)) = \sigma - \mathrm{Sd}(\sigma) - T(\partial\sigma).$$

For $n = 0$ these equalities hold directly, with the unaugmented singular differential. Every cone used in the recursion stays in the convex hull of its simplex's vertex images. Consequently all terms after postcomposition lie in $\sigma(|\Delta^n|)$, so Sd, T preserve small chains.

For a nonempty subset A of a metric space (K, d_K), its *diameter* is

$$\mathrm{diam}(A) = \sup\{d_K(x,y) : x, y \in A\},$$

with value $+\infty$ when these distances are unbounded.

The vertices of each affine map occurring in $\mathrm{Sd}(\tau)$ are barycenters of nested indexed faces: this holds for a vertex, and the recursion cones each flag in a codimension-one face from the full barycenter, adding the largest face to that flag. For $I \subsetneq J$, put $r = |I|$ and $s = |J|$. Then

$$e_J = \frac{r}{s}e_I + \frac{s-r}{s}e_{J-I}, \qquad \|e_J - e_I\| = \frac{s-r}{s}\|e_{J-I} - e_I\| \leq \frac{n}{n+1}\,\mathrm{diam}(\tau(|\Delta^n|)).$$

Here $1 \leq r < s \leq n+1$, and both barycenters lie in the original convex hull. A convex hull has the same diameter as its vertex set: if $x = \sum_i a_i e_i$ and $y = \sum_j b_j e_j$ are convex combinations, then $\|x - y\| \leq \sum_{i,j} a_i b_j \|e_i - e_j\| \leq \max_{i,j}\|e_i - e_j\|$. Thus each subdivided simplex has diameter at most $n/(n+1)$ times the original diameter. Applying this bound inside each new simplex proves the factor $(n/(n+1))^m$ after m subdivisions for $n > 0$. In dimension zero every simplex is already small.

In the 2-simplex picture, let d be the diameter of τ. Each term has vertices e_{012}, e_{ij}, e_i. The three distances are bounded by

$$\|e_{ij} - e_i\| = \tfrac{1}{2}\|e_j - e_i\| \leq \tfrac{1}{2}d,$$
$$\|e_{012} - e_i\| = \tfrac{2}{3}\|e_{jk} - e_i\| \leq \tfrac{2}{3}d,$$
$$\|e_{012} - e_{ij}\| = \tfrac{1}{3}\|e_k - e_{ij}\| \leq \tfrac{1}{3}d,$$

where $\{i,j,k\} = \{0,1,2\}$. A simplex's diameter is the largest distance between its vertices. Each of the six terms therefore has diameter at most $\frac{2}{3}d$; repeating this estimate inside each term gives the bound $(2/3)^m d$ after m subdivisions. The bound concerns the affine simplices in the domain, before postcomposition with the singular simplex $\sigma : |\Delta^n| \to Z$.

The passage from small diameter to small chains uses the following compactness lemma.

Lemma 7.1.13 (Lebesgue number lemma). Let (K, d_K) be a nonempty compact metric space, and let $\mathcal{U}$ be an open cover of K. There is a number $\delta > 0$ such that every subset of K of diameter less than δ is contained in a member of $\mathcal{U}$.

Proof. For each $x \in K$, choose $\varepsilon_x > 0$ and $U_x \in \mathcal{U}$ with $B(x, 2\varepsilon_x) \subset U_x$. By compactness, finitely many balls $B(x_i, \varepsilon_{x_i})$ cover K. Put $\delta = \min_i \varepsilon_{x_i} > 0$. If a nonempty subset $A \subset K$ has diameter less than δ, choose $a \in A$ and an index i with $a \in B(x_i, \varepsilon_{x_i})$. For every $y \in A$,

$$d_K(y, x_i) \le d_K(y, a) + d_K(a, x_i) < \delta + \varepsilon_{x_i} \le 2\varepsilon_{x_i}.$$

Thus $A \subset B(x_i, 2\varepsilon_{x_i}) \subset U_{x_i}$. The conclusion for $A = \emptyset$ is immediate. □

Apply the Lebesgue number lemma (Lemma 7.1.13) to the cover $\sigma^{-1}(U), \sigma^{-1}(V)$ of the compact metric simplex $|\Delta^n|$. Once the subdivision mesh is less than its Lebesgue number δ, each subdivided domain simplex maps into U or V. Taking sufficiently many subdivisions for the finitely many simplices in a chain α therefore makes $\mathrm{Sd}^m(\alpha)$ small. Define the degree-one homomorphism $T_m : \mathrm{C}_n(Z;\mathbb{Z}) \to \mathrm{C}_{n+1}(Z;\mathbb{Z})$ by $T_m = \sum_{r=0}^{m-1} \mathrm{Sd}^r T$; then

$$\partial T_m + T_m \partial = \mathrm{id} - \mathrm{Sd}^m,$$

and T_m preserves small chains.

We now apply these constructions to small chains.

Theorem 7.1.14 (Small chains). For an open cover $Z = U \cup V$, the inclusion

$$\mathrm{C}_\bullet^{\{U,V\}}(Z) \hookrightarrow \mathrm{C}_\bullet(Z)$$

is a chain homotopy equivalence and hence induces an isomorphism in every homology degree.

Proof. Write $\iota : \mathrm{C}_\bullet^{\{U,V\}}(Z) \hookrightarrow \mathrm{C}_\bullet(Z)$. Construct a degree-zero retraction and a degree-one homomorphism

$$\rho : \mathrm{C}_n(Z) \longrightarrow \mathrm{C}_n^{\{U,V\}}(Z), \qquad D : \mathrm{C}_n(Z) \longrightarrow \mathrm{C}_{n+1}(Z) \quad (n \ge 0)$$

inductively, setting $\rho(\sigma) = \sigma$ and $D(\sigma) = 0$ on small basis simplices. For each remaining basis simplex, put $\alpha = \sigma - D(\partial\sigma)$, choose m making $\mathrm{Sd}^m(\alpha)$ small, and set

$$\iota(\rho(\sigma)) = \mathrm{Sd}^m(\alpha) + T_m(\partial\alpha), \qquad D(\sigma) = T_m(\alpha).$$

All faces of a small simplex are small, so the prescribed values on small simplices satisfy the inductive identities. For a remaining simplex, the lower-dimensional homotopy identity gives

$$\partial\alpha = \partial\sigma - \partial D(\partial\sigma) = \partial\sigma - (\partial\sigma - \iota(\rho(\partial\sigma)) - D(\partial^2\sigma)) = \iota(\rho(\partial\sigma)).$$

Thus both terms defining $\rho(\sigma)$ are small. Moreover,

$$\partial(\mathrm{Sd}^m(\alpha) + T_m(\partial\alpha)) = \mathrm{Sd}^m(\partial\alpha) + (\mathrm{id} - \mathrm{Sd}^m)(\partial\alpha) = \partial\alpha,$$

since $\partial^2\alpha = 0$. The displayed homotopy identity gives

$$\partial\iota(\rho(\sigma)) = \partial\alpha = \iota(\rho(\partial\sigma)), \qquad \partial D(\sigma) = \sigma - D(\partial\sigma) - \iota(\rho(\sigma)).$$

Starting with vertices and extending linearly proves $\rho\iota = \mathrm{id}$, $D\iota = 0$, and $\partial D + D\partial = \mathrm{id} - \iota\rho$. Thus ι is a chain homotopy equivalence. □

Let $f : W \hookrightarrow U$, $g : W \hookrightarrow V$, $i : U \hookrightarrow Z$, and $j : V \hookrightarrow Z$ be the inclusions. There is a short exact sequence of chain complexes

$$0 \longrightarrow \mathrm{C}_\bullet(W) \xrightarrow{(f_*, g_*)} \mathrm{C}_\bullet(U) \oplus \mathrm{C}_\bullet(V) \xrightarrow{j_* - i_*} \mathrm{C}_\bullet^{\{U,V\}}(Z) \longrightarrow 0. \tag{16}$$

The first map is injective because inclusions preserve the basis of singular simplices. In the middle, the equation $\mu - \alpha = 0$ implies that $\alpha = \mu$ is supported on simplices lying in both U and V. This identifies the kernel with the image of the first map. The last map is surjective by definition of small chains.

7.1.4 Algebraic and Geometric Double Mapping Cylinders

The mapping cone gives an explicit homotopy pushout of chain complexes. Its comparison retains the specified homotopy in the original square.

Definition 7.1.15. For chain maps $f_* : \mathrm{C}_\bullet(W) \longrightarrow \mathrm{C}_\bullet(X)$ and $g_* : \mathrm{C}_\bullet(W) \longrightarrow \mathrm{C}_\bullet(Y)$, their *algebraic double mapping cylinder*, or *homotopy pushout*, is the complex $D_\bullet$ with degree n group and differential

$$D_n = \mathrm{C}_n(X) \oplus \mathrm{C}_n(Y) \oplus \mathrm{C}_{n-1}(W),$$
$$\partial_D(\lambda, \mu, \nu) = (\partial\lambda + f_*(\nu),\ \partial\mu + g_*(\nu),\ -\partial\nu).$$

The chain-map identities imply $\partial_{\mathrm{D}}^2 = 0$. The structure maps are $i(\lambda) = (-\lambda, 0, 0)$ and $j(\mu) = (0, \mu, 0)$. The degree-one map $T(\nu) = (0, 0, \nu)$ satisfies

$$\partial_D T + T\partial = jg_* - if_*.$$

For a set S, write $\mathbb{Z}\{S\}$ for the free abelian group on S. For the representable simplicial set Δ^1, use the complex $\mathbb{Z}\{\Delta^1\}^{\mathrm{nd}}_\bullet$: its degree-$n$ group is free abelian on the set of nondegenerate n-simplices of Δ^1. It has generators e_0, e_1 in degree zero and e_{01} in degree one, and differential $\partial e_{01} = e_1 - e_0$. The cylinder chain map

$$h : \mathrm{C}_\bullet(W) \otimes \mathbb{Z}\{\Delta^1\}^{\mathrm{nd}}_\bullet \longrightarrow D_\bullet$$

is given, for $\nu \in \mathrm{C}_n(W)$, by

$$h(\nu \otimes e_0) = if_*(\nu), \qquad h(\nu \otimes e_1) = jg_*(\nu), \qquad h(\nu \otimes e_{01}) = (-1)^n T(\nu).$$

Thus h has endpoints if_* and jg_*, while T is its degree-one component with the displayed sign. With the mapping-cone convention of Lecture 2, this explicitly defined complex is $\mathrm{Cone}(u)$ for $u = (f_*, g_*)$.

Let $f : W \longrightarrow X$ and $g : W \longrightarrow Y$ be continuous maps. Consider the subspaces

$$U = X \underset{W\times\{0\}}{+} (W \times [0, 2/3)), \qquad V = (W \times (1/3, 1]) \underset{W\times\{1\}}{+} Y$$

of $\mathrm{Cyl}(f, g)$, using the identifications in Definition 7.1.1. Their preimages in the defining disjoint union are open and saturated, so U, V are open in $\mathrm{Cyl}(f, g)$. They cover $\mathrm{Cyl}(f, g)$, and $U \cap V$ is $W \times (1/3, 2/3)$.

Let $p_U : U \longrightarrow X$ and $p_V : V \longrightarrow Y$ be the identity on X and Y, respectively, with $p_U([w, t]) = f(w)$ and $p_V([w, t]) = g(w)$ on the cylinder portions. They are deformation retractions, by moving heights linearly to 0 and 1. Projection $p : U \cap V \longrightarrow W$ is a homotopy equivalence. These maps give the commuting diagram of spans

$$\begin{array}{ccccc} U & \hookleftarrow & U \cap V & \hookrightarrow & V \\ {\scriptstyle p_U}\downarrow & & \downarrow{\scriptstyle p} & & \downarrow{\scriptstyle p_V} \\ X & \xleftarrow{f} & W & \xrightarrow{g} & Y \end{array}$$

where the vertical maps on the left and right are $p_U : U \longrightarrow X$ and $p_V : V \longrightarrow Y$, respectively.

7.2 Mayer–Vietoris and Homology Calculations

7.2.1 Singular chains preserve homotopy pushouts

Theorem 7.2.1. For every abelian group A, the singular chains functor

$$\mathrm{C}_\bullet(-;A) : \mathrm{Top} \longrightarrow \mathrm{Ch}(\mathrm{Ab})$$

preserves homotopy pushouts.

Remark 7.2.2. Explicitly, for a homotopy pushout square of spaces as in Definition 7.1.3, its comparison from $D_\bullet \otimes_{\mathbb{Z}} A$ to $\mathrm{C}_\bullet(Z;A)$ is a quasi-isomorphism, where $D_\bullet$ is the complex of Definition 7.1.15. In particular, $\mathrm{C}_\bullet(Z)$ is isomorphic to $D_\bullet$ in $\mathcal{D}(\mathbb{Z})$. The chain square induced by any open cover $Z = U \cup V$, with $W = U \cap V$, is also a homotopy pushout.

Proof. First work over $\mathbb{Z}$. For an open cover $Z = U \cup V$, put $W = U \cap V$ and let $\iota : \mathrm{C}_\bullet(W) \to \mathrm{C}_\bullet(U) \oplus \mathrm{C}_\bullet(V)$ be the pair of inclusions. The comparison for the strictly commuting cover square with zero chain homotopy factors as

$$\begin{array}{ccc} \mathrm{Cone}(\iota) & \xrightarrow{\ q\ } & \mathrm{C}_\bullet^{\{U,V\}}(Z) \\ & \searrow & \downarrow \\ & & \mathrm{C}_\bullet(Z). \end{array}$$

In degree n, the cone is

$$\mathrm{Cone}(\iota)_n = \mathrm{C}_n(U) \oplus \mathrm{C}_n(V) \oplus \mathrm{C}_{n-1}(W).$$

Here the three coordinates λ, μ, and ν lie in $\mathrm{C}_n(U)$, $\mathrm{C}_n(V)$, and $\mathrm{C}_{n-1}(W)$, respectively. Writing $i : U \hookrightarrow Z$ and $j : V \hookrightarrow Z$ for the inclusions, the first map is the row matrix

$$\begin{pmatrix} -i_* & j_* & 0 \end{pmatrix} : \mathrm{Cone}(\iota)_n \longrightarrow \mathrm{C}_n^{\{U,V\}}(Z).$$

Write $S_\bullet = \mathrm{C}_\bullet^{\{U,V\}}(Z)$ and denote this map by q. It is surjective: every generating small simplex lies in U or V, and is the image of $(-\sigma, 0, 0)$ or $(0, \sigma, 0)$, respectively. It is a chain map because the two inclusions of W into Z agree, so the contributions of ν cancel in $q\partial$.

Identify chains on U, V, and W with their images in chains on Z. The common simplex basis gives

$$\mathrm{C}_n(U) \cap \mathrm{C}_n(V) = \mathrm{C}_n(W).$$

Thus, writing $K_\bullet = \ker(q)$,

$$\begin{aligned} K_n &= \{(a,a,b) : a \in \mathrm{C}_n(W),\ b \in \mathrm{C}_{n-1}(W)\}, \\ d_K(a,a,b) &= (\partial a + b, \partial a + b, -\partial b), \\ s_n : K_n &\longrightarrow K_{n+1}, \qquad (a,a,b) \longmapsto (0,0,a). \end{aligned}$$

The contraction is verified in the kernel itself:

$$\begin{aligned} d_K s_n(a,a,b) &= (a,a,-\partial a), \\ s_{n-1} d_K(a,a,b) &= (0,0,\partial a + b), \\ (d_K s_n + s_{n-1} d_K)(a,a,b) &= (a,a,b). \end{aligned}$$

Hence s is a chain homotopy from 0 to $\mathrm{id}_{K_\bullet}$: $K_\bullet$ is contractible. In particular, every cycle (a,a,b) satisfies $b = -\partial a$ and is the boundary $d_K(0,0,a)$. The long exact homology sequence of

$$0 \longrightarrow K_\bullet \longrightarrow \mathrm{Cone}(\iota) \xrightarrow{q} S_\bullet \longrightarrow 0$$

now proves that q is a quasi-isomorphism. The second map is a chain homotopy equivalence by Theorem 7.1.14. This proves the open-cover assertion over $\mathbb{Z}$.

Apply the open-cover assertion to the cover of $\mathrm{Cyl}(f,g)$ constructed above. The retractions p_U, p_V and projection p give a strictly commuting diagram

$$\begin{array}{ccccc} \mathrm{C}_\bullet(U) & \longleftarrow & \mathrm{C}_\bullet(U\cap V) & \longrightarrow & \mathrm{C}_\bullet(V) \\ \downarrow{\scriptstyle p_{U,*}} & & \downarrow{\scriptstyle p_*} & & \downarrow{\scriptstyle p_{V,*}} \\ \mathrm{C}_\bullet(X) & \xleftarrow{f_*} & \mathrm{C}_\bullet(W) & \xrightarrow{g_*} & \mathrm{C}_\bullet(Y). \end{array}$$

The vertical maps are chain homotopy equivalences by homotopy invariance of singular chains, since the corresponding maps of spaces were proved to be homotopy equivalences above. Homotopy invariance of derived pushouts therefore identifies the derived pushouts of these two spans. The upper one is $\mathrm{C}_\bullet(\mathrm{Cyl}(f,g))$ by the open-cover assertion; the lower one is $D_\bullet$ by Definition 7.1.15.

This identification respects the specified cylinder homotopy $\eta(w,t) = [w,t]$. Indeed, on the overlap the contraction in U, read backwards, gives the path from $[w,0]$ to $[w,t]$, and the contraction in V gives the path from $[w,t]$ to $[w,1]$. Their concatenation is the cylinder path with a break at height t. Linear interpolation of its height parametrization with the identity gives a homotopy relative to the endpoints, continuously in (w,t). Thus the transported cover cocone is homotopic to the cocone determined by η, so its derived comparison from $D_\bullet$ to $\mathrm{C}_\bullet(\mathrm{Cyl}(f,g))$ is an isomorphism.

For the given homotopy pushout square, the map $\theta_h : \mathrm{Cyl}(f,g) \to Z$ is a homotopy equivalence and satisfies $\theta_h \eta = h$. Homotopy invariance of singular chains and composition with $\theta_{h,*}$ now identify the comparison for the specified homotopy h with an isomorphism in $\mathcal{D}(\mathbb{Z})$. Hence that comparison is a quasi-isomorphism, proving preservation over $\mathbb{Z}$.

Every integral complex used here is nonnegative and degreewise free. A quasi-isomorphism between such complexes is a chain homotopy equivalence: its acyclic cone is contractible, since its cycle groups are free by Lecture 3, Theorem 3.1.10, and its sequences of cycles and boundaries split. In the splitting $C_n = Z_n \oplus Z_{n-1}$ of an acyclic free complex, the differential is $(z,w) \mapsto (w,0)$ and the contracting homotopy is $(z,w) \mapsto (0,z)$. Tensoring with A preserves chain homotopies and commutes with mapping cones. It therefore preserves these homotopy pushouts and gives the result for $\mathrm{C}_\bullet(-;\mathbb{Z}) \otimes_{\mathbb{Z}} A$, without a flatness hypothesis on A. Equivalently, the free complexes compute derived tensoring. □

7.2.2 Mayer–Vietoris and Excision

The preservation theorem gives long exact sequences and relative comparisons by applying the corresponding constructions in chain complexes.

Corollary 7.2.3. For every abelian group A, a homotopy pushout square

$$\begin{array}{ccc} & X & \\ {\scriptstyle f}\nearrow & \Downarrow h & \searrow{\scriptstyle i} \\ W & & Z \\ {\scriptstyle g}\searrow & & \nearrow{\scriptstyle j} \\ & Y & \end{array}$$

as in Definition 7.1.3 induces a natural long exact sequence. Five consecutive terms are

$$\mathrm{H}_n(W;A) \xrightarrow{(f_*,g_*)} \mathrm{H}_n(X;A)\oplus \mathrm{H}_n(Y;A) \xrightarrow{j_*-i_*} \mathrm{H}_n(Z;A) \xrightarrow{\delta_n^Z} \mathrm{H}_{n-1}(W;A) \xrightarrow{(f_*,g_*)} \mathrm{H}_{n-1}(X;A)\oplus \mathrm{H}_{n-1}(Y;A). \quad (17)$$

The sequence ends with $\mathrm{H}_0(Z;A) \longrightarrow 0$. It is natural for maps preserving the square and its specified homotopy.

Proof. Apply the mapping-cone long exact sequence of Lecture 2 to the homotopy pushout of Theorem 7.2.1. Its cone triangle is

$$\mathrm{C}_\bullet(W;A) \xrightarrow{(f_*,g_*)} \mathrm{C}_\bullet(X;A) \oplus \mathrm{C}_\bullet(Y;A) \longrightarrow D_\bullet \otimes_{\mathbb{Z}} A \longrightarrow \mathrm{C}_\bullet(W;A)[1].$$

The comparison identifies the middle map with $j_* - i_*$. Taking the negative of the shifted projection as connecting map gives the stated convention. The shift has differential $-\partial$, and all constructions are natural in the specified square. □

Corollary 7.2.4 (Mayer–Vietoris). For an open cover $Z = U \cup V$, put $W = U \cap V$. Let f, g be the inclusions of W into U, V, and let i, j be the inclusions of U, V into Z, respectively. For every abelian group A, they induce a natural long exact sequence. Five consecutive terms are

$$\mathrm{H}_n(W;A) \xrightarrow{(f_*,g_*)} \mathrm{H}_n(U;A) \oplus \mathrm{H}_n(V;A) \xrightarrow{j_*-i_*} \mathrm{H}_n(Z;A) \xrightarrow{\delta_n} \mathrm{H}_{n-1}(W;A) \xrightarrow{(f_*,g_*)} \mathrm{H}_{n-1}(U;A) \oplus \mathrm{H}_{n-1}(V;A).$$

The sequence ends with $\mathrm{H}_0(Z;A) \longrightarrow 0$. If a class is represented by a small cycle $\zeta = \mu - \lambda$, with $\lambda \in \mathrm{C}_n(U;A)$ and $\mu \in \mathrm{C}_n(V;A)$, then $\delta_n[\zeta] = [\partial\lambda] = [\partial\mu]$ in $\mathrm{H}_{n-1}(W;A)$.

Proof. Apply the mapping-cone long exact sequence to the open-cover assertion of Theorem 7.2.1. To compute its connecting map, use the degreewise split sequence (16) tensored with A. A lift (λ, μ) of $\zeta = \mu - \lambda$ has boundary $(\partial\lambda, \partial\mu) = (\partial\lambda, \partial\lambda)$, so $\delta_n[\zeta] = [\partial\lambda]$. This agrees with the negative shifted projection in Corollary 7.2.3. Naturality and independence of the lift follow from the long exact sequence construction. □

The second map means $(\lambda, \mu) \mapsto j_*(\mu) - i_*(\lambda)$. Its composite with (f_*, g_*) is zero because $if = jg$.

Corollary 7.2.5 (Excision). For an open cover $Z = U \cup V$, put $W = U \cap V$. For every abelian group A, the inclusions induce isomorphisms

$$\mathrm{H}_n(V, W;A) \xrightarrow{\cong} \mathrm{H}_n(Z, U;A).$$

Proof. Write $S_\bullet = \mathrm{C}_\bullet^{\{U,V\}}(Z;\mathbb{Z})$. The maps ρ, D in Theorem 7.1.14 satisfy $\rho\iota = \mathrm{id}$, $D\iota = 0$, and $\partial D + D\partial = \mathrm{id} - \iota\rho$. In particular, ρ fixes $\mathrm{C}_\bullet(U;\mathbb{Z})$ and D vanishes there. Hence these maps descend to quotients and exhibit

$$S_\bullet/\mathrm{C}_\bullet(U;\mathbb{Z}) \longrightarrow \mathrm{C}_\bullet(Z;\mathbb{Z})/\mathrm{C}_\bullet(U;\mathbb{Z})$$

as a chain homotopy equivalence. The map from $\mathrm{C}_\bullet(V;\mathbb{Z})$ onto the left quotient is surjective because every small simplex lies in U or V. Its kernel consists exactly of chains supported on simplices lying in both U and V, namely $\mathrm{C}_\bullet(W;\mathbb{Z})$. It therefore identifies that quotient with $\mathrm{C}_\bullet(V, W;\mathbb{Z})$, and the resulting comparison is the inclusion-induced relative chain map:

$$\begin{array}{ccc} \mathrm{C}_\bullet(V,W;\mathbb{Z}) & \xrightarrow{\cong} & S_\bullet/\mathrm{C}_\bullet(U;\mathbb{Z}) \\ & \searrow^{\text{inclusion}} & \downarrow \simeq \\ & & \mathrm{C}_\bullet(Z,U;\mathbb{Z}). \end{array}$$

These relative-chain quotients are degreewise free on the remaining singular simplices. Tensoring the maps and homotopy identities with A identifies their quotients with the corresponding relative chains with coefficients in A and retains the chain homotopy equivalence. Taking homology proves the assertion. □

7.2.3 The Gluing Axiom and Its Meaning

The Mayer–Vietoris axiom of Lecture 5 is the abstract form of Corollary 7.2.3. Suspensions, cell attachments, and mapping cones give its basic computations.

Let $A_n : \mathrm{Top} \longrightarrow \mathrm{Ab}$, for $n \in \mathbb{Z}$, be functors equipped with natural Mayer–Vietoris connecting maps for homotopy pushout squares. Assume homotopy invariance, additivity, and exactness of the Mayer–Vietoris sequences. Homotopy invariance sends homotopy equivalences to isomorphisms; additivity identifies $A_n(X+Y)$ with $A_n(X) \oplus A_n(Y)$.

Proposition 7.2.6. For a homotopy pushout square of topological spaces

$$\begin{array}{ccc} & X & \\ {\scriptstyle f}\nearrow & \Downarrow h & \searrow{\scriptstyle i} \\ W & & Z \\ {\scriptstyle g}\searrow & & \nearrow{\scriptstyle j} \\ & Y & \end{array}$$

for every integer n, five consecutive terms of the long exact sequence are

$$A_n(W) \xrightarrow{(f_*,g_*)} A_n(X) \oplus A_n(Y) \xrightarrow{j_*-i_*} A_n(Z) \xrightarrow{\delta_n^Z} A_{n-1}(W) \xrightarrow{(f_*,g_*)} A_{n-1}(X) \oplus A_{n-1}(Y).$$

If $\delta_n^{\mathrm{Cyl}(f,g)}$ is the connecting map for the span (f,g), then $\delta_n^Z = \delta_n^{\mathrm{Cyl}(f,g)}(\theta_{h,*})^{-1}$.

Proof. The Mayer–Vietoris sequence of Lecture 5 is put in the displayed convention by multiplying the $A_n(Y)$ summand and the $A_n(\mathrm{Cyl}(f,g))$ term by -1. Thus the cylinder connecting map here is the negative of the one in Lecture 5. Homotopy invariance makes $\theta_{h,*}$ invertible, so transport that sequence to $A_n(Z)$. The identities on the two restrictions of θ_h identify the adjacent homomorphism with $j_* - i_*$. Naturality follows for compatible maps of the chosen squares, as in Lecture 5. □

Remark 7.2.7. Degreewise homology groups need not form a pushout square in Ab. The ordinary pushout of $A_n(X) \leftarrow A_n(W) \rightarrow A_n(Y)$ is canonically isomorphic to the cokernel of (f_*, g_*), by changing the sign on the $A_n(X)$ summand. Exactness in Proposition 7.2.6 identifies this cokernel with $\ker(\delta_n^Z)$, a subgroup of $A_n(Z)$. It is all of $A_n(Z)$ if and only if $\delta_n^Z = 0$. For the suspension square of S^0 in Example 7.1.4, ordinary integral homology in degree one vanishes at the other three corners but $\mathrm{H}_1(\mathrm{S}^1;\mathbb{Z}) = \mathbb{Z}$. The preservation statement is the long exact sequence, or the derived pushout of chains in Theorem 7.2.1.

Use reduced generalized homology and its basepoint splitting from Lecture 5, Definition 5.3.1. Pointed maps preserve this splitting. For ordinary homology the coefficient summand is $\mathrm{H}_n(*;\mathbb{Z})$, computed in Lecture 6, Proposition 6.2.2.

Proposition 7.2.8. Let W be a nonempty pointed topological space. For its unreduced suspension $\mathrm{S}(W)$, pointed at the top vertex, the connecting map gives isomorphisms

$$\widetilde{A}_n(\mathrm{S}(W)) \xrightarrow{\cong} \widetilde{A}_{n-1}(W) \qquad (n \in \mathbb{Z}).$$

Proof. Apply Proposition 7.2.6 to $* \xleftarrow{\mathrm{pr}_W} W \xrightarrow{\mathrm{pr}_W} *$. In degree $n-1$, the map from $A_{n-1}(W)$ to the two point groups is $((\mathrm{pr}_W)_*, (\mathrm{pr}_W)_*)$, whose kernel is $\widetilde{A}_{n-1}(W)$. Exactness makes this kernel the image of δ_n.

The map from the two point groups into $A_n(\mathrm{S}(W))$ is the second vertex map minus the first. These vertices are joined by the interval over the basepoint of W, so their maps on homology agree. The image of their difference on the direct sum is the coefficient summand

$A_n(*)$ in $A_n(\mathrm{S}(W))$. Exactness therefore identifies $A_n(\mathrm{S}(W))/A_n(*)$ with $\widetilde{A}_{n-1}(W)$. The basepoint splitting identifies this quotient with $\widetilde{A}_n(\mathrm{S}(W))$, proving the assertion. □

The sphere groups and their coefficient dependence are already computed in Lecture 5, Proposition 5.3.6 and Example 5.3.7. Proposition 7.2.8 supplies the suspension interpretation used in the cell calculations below.

7.2.4 Cell Attachments and Mapping Cones

The exact sequence retains both the gluing maps and the connecting map, so it determines which classes survive an attachment and which new classes occur.

Example 7.2.9. Cover S^1 by the complements U, V of two distinct points. Each is an open interval and $W = U \cap V$ has two interval components. The sequence of Corollary 7.2.4 contains

$$0 \longrightarrow \mathrm{H}_1(\mathrm{S}^1;\mathbb{Z}) \xrightarrow{\delta_1} \mathbb{Z}^2 \xrightarrow{(f_*,g_*)} \mathbb{Z}^2 \longrightarrow \mathrm{H}_0(\mathrm{S}^1;\mathbb{Z}) \longrightarrow 0.$$

Each component of W maps to the unique component of each cover member. Thus the middle map sends (u,v) to $(u+v,u+v)$. Its kernel is generated by $(1,-1)$, so $\mathrm{H}_1(\mathrm{S}^1;\mathbb{Z}) \cong \mathbb{Z}$. Its cokernel is $\mathbb{Z}$, in agreement with connectedness. Splitting an oriented circuit into its two arcs gives a boundary equal to the difference of the two overlap points; this is the formula for δ_1 in Corollary 7.2.4.

Proposition 7.2.10. Let $m \geq 2$ and let $f : \mathrm{S}^{m-1} \longrightarrow X$ attach an m-cell, producing $Z = X \underset{\mathrm{S}^{m-1}}{+} \mathrm{D}^m$. There is an exact sequence

$$0 \to \mathrm{H}_m(X;\mathbb{Z}) \to \mathrm{H}_m(Z;\mathbb{Z}) \to \mathbb{Z} \xrightarrow{f_*} \mathrm{H}_{m-1}(X;\mathbb{Z}) \to \mathrm{H}_{m-1}(Z;\mathbb{Z}) \to 0.$$

The homomorphism f_* sends a chosen generator of $\mathrm{H}_{m-1}(\mathrm{S}^{m-1};\mathbb{Z})$ to the homology class of the attaching map. In positive degrees other than $m, m-1$, the inclusion induces $\mathrm{H}_n(X;\mathbb{Z}) \cong \mathrm{H}_n(Z;\mathbb{Z})$; it also induces an isomorphism on $\mathrm{H}_0(-;\mathbb{Z})$.

Proof. Example 7.1.8 gives a homotopy pushout. Apply Corollary 7.2.3, using the sphere homology from Lecture 5, Proposition 5.3.6 and the contractibility of D^m. In the relevant positive degrees, the disk groups vanish and $\mathrm{H}_{m-1}(\mathrm{S}^{m-1};\mathbb{Z}) = \mathbb{Z}$, giving the displayed sequence. The remaining positive-degree terms vanish on the sphere and disk, giving the stated isomorphisms.

Since $m \geq 2$, the attaching sphere is connected. In degree zero the map $\mathrm{H}_0(\mathrm{S}^{m-1};\mathbb{Z}) \longrightarrow \mathrm{H}_0(X;\mathbb{Z}) \oplus \mathrm{H}_0(\mathrm{D}^m;\mathbb{Z})$ identifies the disk component with the component of X containing the image of f. Its cokernel is $\mathrm{H}_0(X;\mathbb{Z})$, proving the final assertion. □

For a pointed map $f : W \longrightarrow X$ between nonempty pointed spaces, Example 7.1.9 and Proposition 7.2.6 also give the reduced cone long exact sequence, with five consecutive terms

$$\widetilde{A}_n(W) \xrightarrow{f_*} \widetilde{A}_n(X) \xrightarrow{-i_*} \widetilde{A}_n(\mathrm{Cone}(f)) \longrightarrow \widetilde{A}_{n-1}(W) \xrightarrow{f_*} \widetilde{A}_{n-1}(X).$$

Here $i : X \longrightarrow \mathrm{Cone}(f)$ is the canonical inclusion. To see the reduction, split off the point groups in the unreduced Mayer–Vietoris sequence. On those summands the maps are the split exact sequence $A_n(*) \xrightarrow{u \mapsto (u,u)} A_n(*)^2 \xrightarrow{(u,v) \mapsto v-u} A_n(*)$. The cone on the basepoint gives the homotopy identifying the two point inclusions, and its connecting map is zero. The remaining summands are the displayed reduced sequence. If f is a Hurewicz cofibration, it computes $\widetilde{A}_n(X/W)$ by Proposition 7.1.6.

Computations from the Eilenberg–Steenrod Axioms

Algebraic Topology I | Fall 2026 | October 2026

8.1 Cell Complexes and Degree

We describe cell complexes by successive attachments. Write D^n for the closed unit disk and $\mathrm{S}^{n-1} = \partial \mathrm{D}^n$; in dimension zero, $\mathrm{D}^0 = *$ and $\mathrm{S}^{-1} = \emptyset$.

8.1.1 Cell Attachments and Skeletal Quotients

Definition 8.1.1. To *attach n-cells* to a space X, choose a set Λ and a continuous map $f : W \to X$, where $W = \coprod_{\lambda \in \Lambda} \mathrm{S}^{n-1}$. Let $Y = \coprod_{\lambda \in \Lambda} \mathrm{D}^n$ and let $g : W \to Y$ be the boundary inclusions. The resulting space is the pushout

$$\begin{array}{ccc} W & \xrightarrow{f} & X \\ {\scriptstyle g}\downarrow & & \downarrow{\scriptstyle i} \\ Y & \xrightarrow{j} & Z. \end{array}$$

Thus $Z = (X + Y)/\!\sim$, with the quotient topology and identifications $f(w) \sim g(w)$ for $w \in W$. The restrictions $f_\lambda : \mathrm{S}^{n-1} \to X$ are the *attaching maps*; the restrictions $j_\lambda : \mathrm{D}^n \to Z$ are the *characteristic maps*. Each j_λ restricts to a homeomorphism from the disk interior onto an *open n-cell*. For $n = 0$, this construction adds one isolated point for each $\lambda \in \Lambda$.

Definition 8.1.2. A *cell complex* is a directed colimit

$$X = \operatorname*{colim}_{n \geq 0} \mathrm{sk}_n X \quad \text{in Top}, \qquad \emptyset = \mathrm{sk}_{-1} X \subseteq \mathrm{sk}_0 X \subseteq \mathrm{sk}_1 X \subseteq \cdots,$$

where $\mathrm{sk}_n X$ is obtained from $\mathrm{sk}_{n-1} X$ by attaching n-cells as in Definition 8.1.1. The colimit supplies the topology on X. Its *n-skeleton* $\mathrm{sk}_n X$ is the union of its cells of dimension at most n, and $\mathrm{sk}_n X = \emptyset$ for $n < 0$. The complex is *finite* if it has finitely many cells in total; then $X = \mathrm{sk}_m X$ for some m. A characteristic map need not be injective on its boundary.

Remark 8.1.3. The homology calculations below concern finite cell complexes. All characteristic maps, including those whose boundary maps are not injective, are interpreted through the attaching pushout.

Proposition 8.1.4. Let Λ_n index the n-cells of a finite cell complex. For $n \geq 0$, the pointed quotient satisfies

$$\frac{(\mathrm{sk}_n X)_+}{(\mathrm{sk}_{n-1} X)_+} \cong \bigvee_{\lambda \in \Lambda_n} \mathrm{S}^n.$$

The empty wedge is $*$. When $\mathrm{sk}_{n-1} X \neq \emptyset$, the left side is the usual quotient $\mathrm{sk}_n X / \mathrm{sk}_{n-1} X$; when $n = 0$, it is $(\mathrm{sk}_0 X)_+$.

Proof. Collapse the previous skeleton and the added point in the attaching quotient. Each n-disk then has its entire boundary collapsed to the common basepoint, and its interior remains unchanged. For $n > 0$ this gives one $\mathrm{D}^n/\mathrm{S}^{n-1} \cong \mathrm{S}^n$ per cell. For $n = 0$ each new point together with the common basepoint is S^0. If there are no cells, only the basepoint remains. Both descriptions impose the same identifications, so the quotient universal property gives the homeomorphism. □

8.1.2 Finite Product Cell Structures

Products of characteristic maps give the cells used in Lecture 8.

Proposition 8.1.5. If X and Y are finite cell complexes, their ordinary product has a finite cell structure whose cells are $e \times e'$, of dimension $p+q$ when e and e' have dimensions p and q. Its skeleta are

$$\mathrm{sk}_n(X \times Y) = \bigcup_{p+q \le n} \mathrm{sk}_p X \times \mathrm{sk}_q Y \qquad (p, q \ge 0).$$

Proof. Take products of characteristic maps. The pair consisting of $\mathrm{D}^p \times \mathrm{D}^q$ and its boundary is homeomorphic to $(\mathrm{D}^{p+q}, \mathrm{S}^{p+q-1})$. Its boundary is

$$(\mathrm{S}^{p-1} \times \mathrm{D}^q) \cup (\mathrm{D}^p \times \mathrm{S}^{q-1}),$$

where the two faces are glued along $\mathrm{S}^{p-1} \times \mathrm{S}^{q-1}$. Both faces map into the preceding skeleton. The resulting finite attachment space maps continuously and bijectively to $X \times Y$. Each finite cell complex is compact Hausdorff: this holds inductively when finitely many compact disks are attached along their closed boundaries. Thus the map is a homeomorphism, which verifies the ordinary product topology. □

8.1.3 Degree of Sphere Maps

By Lecture 5, Proposition 5.3.6, the reduced integral homology of an n-sphere is infinite cyclic in degree n. This gives the following invariant of maps between spheres.

Definition 8.1.6. For $n \ge 1$, a *homological orientation* of S^n is a choice of generator of $\widetilde{\mathrm{H}}_n(\mathrm{S}^n; \mathbb{Z}) \cong \mathbb{Z}$. After choosing such generators for the source and target, the *degree* of a map $f : \mathrm{S}^n \to \mathrm{S}^n$ is the integer $\deg(f)$ by which its induced homomorphism multiplies them.

Throughout the following two propositions, let $n \ge 1$. When source and target are the same oriented sphere, reversing its orientation on both sides does not change this integer. Reversing just one of the two orientations changes its sign.

Proposition 8.1.7. For maps $f, g : \mathrm{S}^n \to \mathrm{S}^n$, one has $\deg(gf) = \deg(g)\deg(f)$. Homotopic maps have the same degree. The identity has degree 1, a constant map has degree 0, and a homotopy equivalence has degree 1 or -1. Every map of nonzero degree is surjective.

Proof. Functoriality of integral homology proves the composition formula and the assertion about the identity. Homotopy invariance proves the second assertion. A constant map factors through a point, whose positive-dimensional integral homology vanishes. If g is a homotopy inverse to f, then $\deg(g)\deg(f) = 1$ in $\mathbb{Z}$, so each integer is 1 or -1.

If a point y is absent from the image, then f factors through $\mathrm{S}^n - \{y\}$. Stereographic projection identifies this complement with the contractible space $\mathbb{R}^n$. The induced map on $\mathrm{H}_n(-; \mathbb{Z})$ is therefore zero. This proves the final assertion by contraposition. □

Proposition 8.1.8. If $f : \mathrm{S}^n \to \mathrm{S}^n$ extends to a continuous map $F : \mathrm{D}^{n+1} \to \mathrm{S}^n$, then $\deg(f) = 0$.

Proof. The extension gives the commutative triangle

$$\begin{array}{ccc} \mathrm{S}^n & \xrightarrow{\ i\ } & \mathrm{D}^{n+1} \\ & \searrow^{f} & \downarrow{\scriptstyle F} \\ & & \mathrm{S}^n. \end{array}$$

After applying $\mathrm{H}_n(-;\mathbb{Z})$, the map f_* factors through the zero group $\mathrm{H}_n(\mathrm{D}^{n+1};\mathbb{Z})$. Hence its degree is zero. □

This is a necessary condition for extension. Its proof uses only homology and does not require a classification of sphere maps up to homotopy. It is the obstruction that will be used in the proof of Brouwer's fixed-point theorem.

Example 8.1.9. Attach an m-cell to S^{m-1} along a map of degree d, with $m \geq 2$. Lecture 7, Proposition 7.2.10 gives

$$\mathrm{H}_m(\mathrm{S}^{m-1} \underset{\mathrm{S}^{m-1}}{+} \mathrm{D}^m;\mathbb{Z}) = \ker(d : \mathbb{Z} \longrightarrow \mathbb{Z}), \qquad \mathrm{H}_{m-1}(\mathrm{S}^{m-1} \underset{\mathrm{S}^{m-1}}{+} \mathrm{D}^m;\mathbb{Z}) = \mathbb{Z}/d\mathbb{Z}.$$

For $d \neq 0$ these groups are 0 and $\mathbb{Z}/d\mathbb{Z}$; for $d = 0$ both are $\mathbb{Z}$. All other reduced homology groups vanish. For $m = 2$ and $d = 2$, this is the disk attachment describing $\mathbb{R}P^2$; its homology was computed in Lecture 4, Example 4.4.12.

The following degree computations use homological orientations throughout. They will determine the attaching coefficients of real projective spaces and provide explicit maps of every integral degree.

Lemma 8.1.10. For $n \geq 1$, suspension preserves the degree of maps $\mathrm{S}^n \to \mathrm{S}^n$. Every coordinate reflection of $\mathrm{S}^n \subseteq \mathbb{R}^{n+1}$ has degree -1, and its antipodal map has degree $(-1)^{n+1}$. For every integer m, the map $p_m : \mathrm{S}^1 \to \mathrm{S}^1$ defined by $p_m(z) = z^m$ has degree m. Its iterated suspensions give pointed degree-m maps in every positive dimension.

Proof. Let $f : \mathrm{S}^n \to \mathrm{S}^n$ with $n \geq 0$. For ordinary reduced integral homology, naturality of the connecting map in Lecture 7, Proposition 7.2.8 gives the commutative square

$$\begin{array}{ccc} \widetilde{\mathrm{H}}_{n+1}(\mathrm{S}(\mathrm{S}^n);\mathbb{Z}) & \xrightarrow[\cong]{\delta} & \widetilde{\mathrm{H}}_n(\mathrm{S}^n;\mathbb{Z}) \\ {\scriptstyle (\mathrm{S}(f))_*}\downarrow & & \downarrow{\scriptstyle f_*} \\ \widetilde{\mathrm{H}}_{n+1}(\mathrm{S}(\mathrm{S}^n);\mathbb{Z}) & \xrightarrow[\cong]{\delta} & \widetilde{\mathrm{H}}_n(\mathrm{S}^n;\mathbb{Z}). \end{array}$$

This also holds for $n = 0$: the reduced group is the kernel of the augmentation, and every map preserves this kernel, whether or not it preserves a chosen basepoint. Choose the suspended generators to correspond under δ. The square proves preservation of degree.

On $\mathrm{S}^0 = \{-1, 1\}$, interchange sends the generator $[1] - [-1]$ to its negative. Under the homeomorphism

$$\mathrm{S}(\mathrm{S}^{n-1}) \longrightarrow \mathrm{S}^n, \qquad [u, t] \longmapsto (\sin(\pi t)u, \cos(\pi t)),$$

reflection of the first coordinate is the suspension of the same reflection in one lower dimension. Induction therefore gives degree -1. Functoriality gives $\deg(gf) = \deg(g)\deg(f)$ and shows that conjugation by a homeomorphism preserves degree. Coordinate permutations conjugate the first reflection to each other coordinate reflection. Their composite is the antipodal map, so its degree is $(-1)^{n+1}$. In particular, this sign has been obtained on homology without an identification of degree with a determinant sign.

Let $\gamma(t) = \exp(2\pi i t)$ for $0 \leq t \leq 1$. This singular 1-simplex is a cycle. Its class generates $\mathrm{H}_1(\mathrm{S}^1;\mathbb{Z})$. Indeed, use the cover of Lecture 7, Example 7.2.9 with U the complement of $-i$ in the circle and V the complement of i. The upper and lower semicircle pieces lie in U and V, respectively. The connecting-map formula sends their sum to $[1] - [-1]$, a generator of the kernel in that example. To justify this splitting at the chain level, for any path $\eta : I \to X$ and $0 < a < 1$, the singular 2-simplex

$$(t_0, t_1, t_2) \longmapsto \eta(at_1 + t_2)$$

has boundary $\eta|_{[a,1]} - \eta + \eta|_{[0,a]}$, where the restricted paths are linearly reparametrized to I. Thus subdivision of a path into finitely many pieces changes its chain by a boundary.

For $m > 0$, subdivide $p_m\gamma$ at $0, 1/m, \dots, 1$. Each reparametrized piece is γ, so $(p_m)_*([\gamma]) = m[\gamma]$. The map p_0 is constant and has degree zero. Complex conjugation p_{-1} is a coordinate reflection of the circle, and $p_{-m} = p_{-1}p_m$ for $m > 0$, proving the formula for negative integers. Iterated suspension preserves the degree and the identities $p_a p_b = p_{ab}$, since suspension respects composition. These identities will give commuting attaching-map squares as well as the required degrees. □

A local contribution will mean an integer defined on relative homology, rather than a sign inferred from a picture. The next lemma proves that these integers add to the homological degree. For a finite subset E of S^n, write S^n_E for its complement. If $E = \{x\}$, write S^n_x.

Lemma 8.1.11. Let $f : \mathrm{S}^n \to \mathrm{S}^n$ with $n \geq 1$, and let y have finite inverse image $F = \{x_1, \dots, x_s\}$. Suppose f is a homeomorphism from a neighborhood of each x_i onto a neighborhood of y. Choose integral homological orientations of the source and target spheres. They induce generators of the local groups

$$\mathrm{H}_n(\mathrm{S}^n, \mathrm{S}^n_{x_i}; \mathbb{Z}), \qquad \mathrm{H}_n(\mathrm{S}^n, \mathrm{S}^n_y; \mathbb{Z}).$$

Let $\varepsilon_i \in \{1, -1\}$ be the coefficient of the local map induced by f in these generators. Then $\deg(f) = \sum_i \varepsilon_i$.

Proof. For any point $x \in \mathrm{S}^n$, stereographic projection identifies S^n_x with $\mathbb{R}^n$, so this complement is contractible. The reduced long exact sequence of its inclusion therefore makes the canonical map

$$j_x : \mathrm{H}_n(\mathrm{S}^n; \mathbb{Z}) \xrightarrow{\cong} \mathrm{H}_n(\mathrm{S}^n, \mathrm{S}^n_x; \mathbb{Z})$$

an isomorphism, including when $n = 1$. The local generator is, by definition, the image of the global generator under j_x.

Choose pairwise disjoint open neighborhoods U_i of the x_i on which f is a homeomorphism onto a neighborhood of y. Put $U = \coprod_i U_i$ and $U_i^\times = U_i \cap \mathrm{S}^n_F$. The sets U and S^n_F form an open cover. Lecture 7, Corollary 7.2.5, followed by additivity, identifies

$$\begin{aligned} \mathrm{H}_n(\mathrm{S}^n, \mathrm{S}^n_F; \mathbb{Z}) &\cong \bigoplus_i \mathrm{H}_n(U_i, U_i^\times; \mathbb{Z}) \\ &\cong \bigoplus_i \mathrm{H}_n(\mathrm{S}^n, \mathrm{S}^n_{x_i}; \mathbb{Z}). \end{aligned}$$

The last isomorphisms are excision for the covers $U_i \cup (\mathrm{S}^n_{x_i})$. The map from the first relative group to the relative group at x_i is the projection onto the ith summand: every other summand is represented inside $\mathrm{S}^n_{x_i}$ and hence maps to zero. Consequently the image of the global generator has coefficient 1 in every summand.

The identity $f^{-1}(y) = F$ gives a map of subspace inclusions from $\mathrm{S}^n_F \hookrightarrow \mathrm{S}^n$ to $\mathrm{S}^n_y \hookrightarrow \mathrm{S}^n$. On the ith excision summand its relative homology map is the local map of f. This map is an isomorphism because the restriction is a homeomorphism and both local identifications are excision isomorphisms; thus its coefficient is $\varepsilon_i \in \{1, -1\}$. Naturality now gives, in the chosen generators, the diagram

$$\begin{array}{ccc} \mathbb{Z} & \xrightarrow{1 \mapsto (1,\dots,1)} & \mathbb{Z}^s \\ {\scriptstyle \times \deg(f)}\downarrow & & \downarrow{\scriptstyle (a_i) \mapsto \sum_i \varepsilon_i a_i} \\ \mathbb{Z} & = & \mathbb{Z}. \end{array}$$

Its commutativity proves the formula. If F is empty, the map factors through the contractible complement of y, so its degree is zero, the empty sum. □

The local-degree formula requires a finite inverse image and local homeomorphisms. Counting inverse images without their signs does not in general compute degree. The local maps are obtained on small neighborhoods by excision, not from maps of the whole sphere punctured at a single inverse image.

For example, consider $p_m : \mathrm{S}^1 \to \mathrm{S}^1$, $p_m(z) = z^m$, for an integer $m > 0$. Every target point has m inverse images. In angular coordinates the local map is $t \mapsto mt$. On a small arc this map extends to an orientation-preserving piecewise linear homeomorphism of the circle. Such a homeomorphism is homotopic to the identity by linearly interpolating an increasing lift to $\mathbb{R}$ with the identity lift. Its global degree, and hence its local degree, is 1. Each contribution is therefore 1, and $\deg(p_m) = m$. Complex conjugation has degree -1, and composition gives $\deg(p_m) = m$ also for negative m. For $m = 0$ the map is constant.

8.2 Cellular Homology and Products

We use ordinary singular homology with coefficients in an abelian group A: in addition to homotopy invariance, excision, exactness, and additivity, the dimension axiom says that $\mathrm{H}_0(*; A) = A$ and $\mathrm{H}_q(*; A) = 0$ for $q \neq 0$. In the computations below, cell complexes are finite, as in Definition 8.1.2. Their n-skeleta, denoted $\mathrm{sk}_n X$, consist of their cells of dimension at most n; put $\mathrm{sk}_r X = \emptyset$ for $r < 0$.

8.2.1 Relative Cells and Cellular Homology

For $n \geq 0$, the sphere calculation of Lecture 5, Proposition 5.3.6, gives

$$\widetilde{\mathrm{H}}_q(\mathrm{S}^n; A) = \begin{cases} A, & q = n, \\ 0, & q \neq n. \end{cases}$$

Here reduced homology has the convention of Lecture 5, Definition 5.3.1, including S^0. For $n > 0$, ordinary homology also has its copy of A in degree zero.

To apply open-cover excision to the characteristic disks, fix $n > 0$ and write $Y = \mathrm{sk}_n X$ and $Y' = \mathrm{sk}_{n-1} X$. Let D be the disjoint union of the characteristic disks, let $E \subset D$ be their boundaries, and let $j : D \to Y$ be the characteristic map. The attaching construction is the pushout

$$\begin{array}{ccc} E & \longrightarrow & Y' \\ \downarrow & & \downarrow \\ D & \xrightarrow[j]{} & Y. \end{array}$$

Let $N \subset D$ be the union of the collars $\|x\| > 1/2$, including the boundaries, and put $U = Y' \cup j(N)$. Let V_D be the union of the disk interiors and $V = j(V_D)$ the union of the open n-cells. The quotient topology makes U, V an open cover of Y: their inverse images in $Y' + D$ are $Y' + N$ and V_D. Likewise, N, V_D cover D by open sets. On the interiors, j gives a homeomorphism of pairs from $(V_D, N \cap V_D)$ to $(V, U \cap V)$, regardless of any boundary identifications.

Radial motion $r \mapsto (1-t)r + t$ retracts N onto E. With Y' fixed, it also retracts U onto Y'. These homotopies are strong deformation retractions: they fix the boundaries, descend through the attaching quotient, and are continuous by the product-quotient lemma of Lecture 7, Lemma 7.1.7, applied to the open saturated inverse image $Y' + N$. The relative long exact sequences

and open-cover excision, Lecture 7, Corollary 7.2.5, now give the commutative diagram

$$\begin{array}{ccc}
\mathrm{H}_q(D,E;A) & \xrightarrow{\;j_*\;} & \mathrm{H}_q(Y,Y';A) \\
\cong\downarrow & & \downarrow\cong \\
\mathrm{H}_q(D,N;A) & \xrightarrow[\;j_*\;]{} & \mathrm{H}_q(Y,U;A) \\
\cong\uparrow & & \uparrow\cong \\
\mathrm{H}_q(V_D,N\cap V_D;A) & \xrightarrow[\cong]{} & \mathrm{H}_q(V,U\cap V;A).
\end{array}$$

The lower vertical maps are excision, and the bottom map is the homeomorphism of pairs. Hence the top map is an isomorphism. Additivity for the disjoint disks identifies the characteristic-map homomorphism

$$\bigoplus_{n\text{-cells of } X} \mathrm{H}_q(\mathrm{D}^n,\mathrm{S}^{n-1};A) \xrightarrow{\;\cong\;} \mathrm{H}_q(\mathrm{sk}_n X,\mathrm{sk}_{n-1} X;A).$$

For $n=0$, this is the direct-sum decomposition for a finite set of points, using $(\mathrm{D}^0,\emptyset)$.

Choose an orientation of each cell. The disk-pair long exact sequence and the sphere calculation show that these groups vanish unless $q=n$, when each cell contributes one copy of A. A degree-d map of an oriented sphere acts by multiplication by d with coefficients in A: the coefficient maps $\mathbb{Z}\to A$, sending 1 to each element of A, carry the integral degree formula to this formula by naturality.

Definition 8.2.1. The *cellular chain complex* of X with coefficients in A has chain groups

$$\mathrm{C}_n^{\mathrm{cell}}(X;A)=\mathrm{H}_n(\mathrm{sk}_n X,\mathrm{sk}_{n-1} X;A).$$

Its differential is the composite in the diagram

$$\mathrm{H}_n(\mathrm{sk}_n X,\mathrm{sk}_{n-1} X;A) \xrightarrow{\;\delta\;} \mathrm{H}_{n-1}(\mathrm{sk}_{n-1} X;A) \longrightarrow \mathrm{H}_{n-1}(\mathrm{sk}_{n-1} X,\mathrm{sk}_{n-2} X;A).$$

Write $j_r:\mathrm{H}_r(\mathrm{sk}_r X;A)\to\mathrm{H}_r(\mathrm{sk}_r X,\mathrm{sk}_{r-1} X;A)$ for the relative comparison. Then $\partial_n=j_{n-1}\delta_n$, and $\delta_{n-1}j_{n-1}=0$ by exactness for the pair $(\mathrm{sk}_{n-1} X,\mathrm{sk}_{n-2} X)$. Thus $\partial_{n-1}\partial_n=0$; put $\partial_0=0$.

Here is the attaching-coefficient calculation for $n\geq 2$. Fix an oriented n-cell with characteristic map $\chi:(\mathrm{D}^n,\mathrm{S}^{n-1})\to(\mathrm{sk}_n X,\mathrm{sk}_{n-1} X)$ and attaching map $a=\chi|_{\mathrm{S}^{n-1}}$. Its relative generator ω has boundary the chosen sphere generator ζ. Naturality gives $\delta_n(\chi_*(\omega))=a_*(\zeta)$. For an $(n-1)$-cell e, collapse the complement of its open cell in $\mathrm{sk}_{n-1} X$ to a point. The resulting quotient is its characteristic disk with boundary collapsed, hence an oriented S^{n-1}; denote the collapse by q_e. It gives a map of pairs

$$q_e:(\mathrm{sk}_{n-1} X,\mathrm{sk}_{n-2} X)\longrightarrow(\mathrm{S}^{n-1},*).$$

On the characteristic relative classes this map kills every other cell and sends the class of e to the chosen generator: on the latter disk it is exactly the disk-boundary collapse. This last assertion follows also by applying the earlier collar-excision argument to the disk and its collapsed boundary. Thus $(q_e)_*$ is projection onto the e summand. The map $\mathrm{H}_{n-1}(\mathrm{S}^{n-1};\mathbb{Z})\to\mathrm{H}_{n-1}(\mathrm{S}^{n-1},*;\mathbb{Z})$ is an isomorphism, including $n=2$, by the reduced relative exact sequence. Applying $(q_e)_*j_{n-1}$ to $a_*(\zeta)$ shows that the coefficient is precisely $\deg(q_e\circ a)$. For $n=1$, use the interval boundary directly: its terminal vertex minus its initial vertex, with cancellation if

they coincide. No degree of a map of 0-spheres is needed. Naturality with respect to $\mathbb{Z} \to A$, $1 \mapsto a$ for each $a \in A$, now gives the same integer incidence coefficients with coefficients in A. Therefore

$$\mathrm{C}_\bullet^{\mathrm{cell}}(X; A) = \mathrm{C}_\bullet^{\mathrm{cell}}(X; \mathbb{Z}) \otimes_{\mathbb{Z}} A.$$

Theorem 8.2.2. For every finite cell complex X and abelian group A, there are isomorphisms $\mathrm{H}_n(\mathrm{C}_\bullet^{\mathrm{cell}}(X; A)) \cong \mathrm{H}_n(X; A)$.

Proof. For brevity write $X_r = \mathrm{sk}_r X$ and $C_r = \mathrm{C}_r^{\mathrm{cell}}(X; A)$. The relative-cell calculation gives $\mathrm{H}_q(X_r, X_{r-1}; A) = 0$ for $q \neq r$. Starting with X_0, a finite set of points, induction in the exact sequence

$$\mathrm{H}_q(X_{r-1}; A) \longrightarrow \mathrm{H}_q(X_r; A) \longrightarrow \mathrm{H}_q(X_r, X_{r-1}; A)$$

shows $\mathrm{H}_q(X_r; A) = 0$ for $q > r$.

Fix $n \geq 1$. The following exact segments display all maps needed for the comparison:

$$\begin{aligned} 0 = \mathrm{H}_n(X_{n-1}; A) \longrightarrow \mathrm{H}_n(X_n; A) &\xrightarrow{j_n} C_n \xrightarrow{\delta_n} \mathrm{H}_{n-1}(X_{n-1}; A), \\ 0 = \mathrm{H}_{n-1}(X_{n-2}; A) \longrightarrow \mathrm{H}_{n-1}(X_{n-1}; A) &\xrightarrow{j_{n-1}} C_{n-1}, \\ C_{n+1} \xrightarrow{\delta_{n+1}} \mathrm{H}_n(X_n; A) \longrightarrow \mathrm{H}_n(X_{n+1}; A) &\longrightarrow \mathrm{H}_n(X_{n+1}, X_n; A) = 0. \end{aligned}$$

The first two lines make j_n, j_{n-1} injective and give $\ker \partial_n = \ker(j_{n-1}\delta_n) = \ker \delta_n = \operatorname{im} j_n$. Write $p : \mathrm{H}_n(X_n; A) \to \mathrm{H}_n(X_{n+1}; A)$ for the inclusion-induced map. The relevant diagram has exact rows:

$$\begin{array}{ccccccc} C_{n+1} & \xrightarrow{\delta_{n+1}} & \mathrm{H}_n(X_n; A) & \xrightarrow{p} & \mathrm{H}_n(X_{n+1}; A) & \longrightarrow & 0 \\ \| & & \downarrow{\scriptstyle j_n} & & \uparrow{\scriptstyle \theta} & & \\ C_{n+1} & \xrightarrow[\partial_{n+1}]{} & \ker \partial_n & \xrightarrow[\pi]{} & \ker \partial_n / \operatorname{im} \partial_{n+1} & \longrightarrow & 0. \end{array}$$

Here j_n is regarded as a map onto its image $\ker \partial_n$, and the left square commutes because $j_n \delta_{n+1} = \partial_{n+1}$. Define θ by the following chase.

- For $c \in \ker \partial_n$, there is a unique $a \in \mathrm{H}_n(X_n; A)$ with $j_n(a) = c$. Set $\theta([c]) = p(a)$.
- If c is replaced by $c + \partial_{n+1} b$, its unique lift is $a + \delta_{n+1} b$. Since $p\delta_{n+1} = 0$, the value of θ is unchanged. Thus θ is a well-defined homomorphism.
- Given $z \in \mathrm{H}_n(X_{n+1}; A)$, exactness gives a with $p(a) = z$; then $\theta([j_n(a)]) = z$. Thus θ is surjective.
- If $\theta([c]) = 0$, lift $c = j_n(a)$. Exactness gives $a = \delta_{n+1} b$, hence $c = j_n \delta_{n+1} b = \partial_{n+1} b$ and $[c] = 0$. Thus θ is injective.

Consequently

$$\ker \partial_n / \operatorname{im} \partial_{n+1} \cong \mathrm{H}_n(X_{n+1}; A).$$

For $n = 0$, $C_0 = \mathrm{H}_0(X_0; A)$ and $\partial_0 = 0$. The third exact segment, with $n = 0$, gives the same quotient conclusion directly.

Finally, for $r \geq n + 2$, both outside terms in

$$\mathrm{H}_{n+1}(X_r, X_{r-1}; A) \longrightarrow \mathrm{H}_n(X_{r-1}; A) \longrightarrow \mathrm{H}_n(X_r; A) \longrightarrow \mathrm{H}_n(X_r, X_{r-1}; A)$$

vanish. The middle map is therefore an isomorphism. There are finitely many skeleta before reaching X, so $\mathrm{H}_n(X_{n+1}; A) \cong \mathrm{H}_n(X; A)$. This proves the theorem in every nonnegative degree; both sides vanish in negative degrees. □

The dimension axiom is used in the vanishing of the relative-cell groups outside one degree. The same argument does not give this cellular complex for a generalized homology theory with nonzero coefficient groups in other degrees.

8.2.2 Product Cells

For ordinary homology, products of cells identify the chain complex of a product with a tensor product. We first construct the canonical homology cross product and then give hypotheses under which it is an isomorphism. A geometric treatment of product cells and their tensor-product signs is given in [May99, Chapter 13, §4, pp. 101–103].

A cell of $X \times Y$ is the product of a cell of X and a cell of Y, oriented in that order. Put $\mathrm{S}^{-1} = \emptyset$ when a factor is a 0-cell. The oriented boundary of $\mathrm{D}^p \times \mathrm{D}^q$ is

$$\partial(\mathrm{D}^p \times \mathrm{D}^q) = (\mathrm{S}^{p-1} \times \mathrm{D}^q) \cup (\mathrm{D}^p \times \mathrm{S}^{q-1}).$$

The two faces are glued along $\mathrm{S}^{p-1} \times \mathrm{S}^{q-1}$, and the product characteristic map sends their union into $\mathrm{sk}_{p+q-1}(X \times Y)$. Give the product disk the orientation specified by the signed shuffle triangulation below. The boundary calculation there gives sign $+1$ on the first kind of face and $(-1)^p$ on the second kind. Consequently, identifying oriented product cells with tensors gives an isomorphism

$$\mathrm{C}_\bullet^{\mathrm{cell}}(X \times Y; \mathbb{Z}) \cong \mathrm{C}_\bullet^{\mathrm{cell}}(X; \mathbb{Z}) \otimes_{\mathbb{Z}} \mathrm{C}_\bullet^{\mathrm{cell}}(Y; \mathbb{Z}), \qquad \partial(\lambda \otimes \mu) = \partial\lambda \otimes \mu + (-1)^p \lambda \otimes \partial\mu \quad (\lambda \in \mathrm{C}_p^{\mathrm{cell}}(X; \mathbb{Z})).$$

Indeed, restrict each product characteristic map to these two boundary pieces and use naturality of the relative connecting homomorphism. The relative product boundary identity proved below gives the original incidence coefficients on the first piece and $(-1)^p$ times those on the second. Their common boundary maps into the $(p+q-2)$-skeleton, so contributes zero to the relative group defining the differential. These account for every $(p+q-1)$-cell. The product disk is a disk of dimension $p+q$: for positive-dimensional factors, radial rescaling from the Euclidean unit ball by the continuous gauge $(x, y) \mapsto \max(\|x\|, \|y\|)$ gives a boundary-preserving homeomorphism. A zero-dimensional factor changes nothing.

Let R be a commutative ring. Triangulate $|\Delta^p| \times |\Delta^q|$ by monotone lattice paths from $(0, 0)$ to (p, q). A path with successive lattice vertices (i_k, j_k), $0 \le k \le p + q$, determines the ordered affine simplex with vertices (e_{i_k}, e_{j_k}). Give it coefficient $(-1)^N$, where N counts pairs consisting of a vertical step followed later by a horizontal step. These simplices form a triangulation: the cumulative coordinates of a point in the two factors can be interleaved in nondecreasing order, and the resulting order determines its path simplex; ties give their common faces. Thus their interiors are disjoint and their union is the product. For singular simplices σ and τ, composing this signed triangulation with $\sigma \times \tau$ and extending R-bilinearly defines the natural chain map

$$\boxtimes : \mathrm{C}_\bullet(X; R) \otimes_R \mathrm{C}_\bullet(Y; R) \longrightarrow \mathrm{C}_\bullet(X \times Y; R).$$

Here is the boundary calculation, including the signs. Deleting the kth vertex of an ordered path simplex contributes $(-1)^k$. If the two steps adjacent to that vertex have different directions, swapping those steps gives the same ordered deleted-vertex face in a second path. The two values of N differ by one, so these interior faces cancel as identical parametrized simplices.

The remaining faces lie on a face of one factor. For two horizontal steps adjacent to a vertex (i, j), deleting that vertex gives the ith face of the first factor. Since $k = i + j$, and removing the redundant horizontal step changes N by j, its coefficient relative to the smaller shuffle is $(-1)^{i+j}(-1)^j = (-1)^i$. For two vertical steps, the analogous change in N is $p - i$, and the sign is $(-1)^{i+j}(-1)^{p-i} = (-1)^{p+j}$, the jth face sign of the second factor with the extra $(-1)^p$. Deleting

the initial or terminal vertex satisfies the same formulas, according to the direction of the initial or terminal step. These cases exhaust the faces, proving

$$\partial(\sigma \boxtimes \tau) = \partial\sigma \boxtimes \tau + (-1)^p \sigma \boxtimes \partial\tau.$$

For a zero-dimensional factor this identity follows directly from the single path and the zero differential in degree zero.

The same formula passes to relative chains: chains in either specified subspace map into its product with the other factor. On oriented disks the signed triangulation defines the product relative generator. One can check it is a generator by restricting to a point in the interior of a top-dimensional triangulation simplex: all other simplices disappear in local relative homology, and its coefficient is $+1$ or -1. Excision identifies the local group with that of the interior simplex; radial retraction of the punctured disk onto its boundary and the relative exact sequences identify it with the disk's relative top homology. The displayed boundary calculation therefore proves the relative product boundary identity, and hence the product-cell signs used above. Thus products of cycles define a canonical homomorphism of graded R-modules

$$\boxtimes : \mathrm{H}_*(X;R) \otimes_R \mathrm{H}_*(Y;R) \longrightarrow \mathrm{H}_*(X \times Y;R), \qquad (M_* \otimes_R N_*)_n = \bigoplus_{p+q=n} M_p \otimes_R N_q.$$

For the interchange map $t : X \times Y \to Y \times X$, exchanging horizontal and vertical steps replaces N by $pq - N$. Thus every corresponding simplex changes sign by $(-1)^{pq}$, proving $t_*(\alpha \boxtimes \beta) = (-1)^{pq}\beta \boxtimes \alpha$ in degrees p, q. On oriented relative cells, this construction gives the product orientation used above.

Theorem 8.2.3. Let R be a principal ideal domain and let X, Y be finite cell complexes whose homology groups over R are free. The canonical cross product is an isomorphism

$$\mathrm{H}_*(X;R) \otimes_R \mathrm{H}_*(Y;R) \xrightarrow{\cong} \mathrm{H}_*(X \times Y;R).$$

In particular, this holds over any field. For finite cell complexes $X_1, \dots, X_r$ with free integral homology, the graded tensor product over $\mathbb{Z}$ gives

$$\mathrm{H}_*\left(\prod_{i=1}^{r} X_i;\mathbb{Z}\right) \cong \bigotimes_{i=1}^{r} \mathrm{H}_*(X_i;\mathbb{Z}).$$

For coefficients in an abelian group A, tensor this graded group over $\mathbb{Z}$ with A. Its degree-k part is, with all tensor products over $\mathbb{Z}$,

$$\mathrm{H}_k\left(\prod_{i=1}^{r} X_i;A\right) \cong \bigoplus_{k_1+\cdots+k_r=k} \left(\bigotimes_{i=1}^{r} \mathrm{H}_{k_i}(X_i;\mathbb{Z})\right) \otimes_{\mathbb{Z}} A.$$

Proof. Over a principal ideal domain, a finite free chain complex with free homology splits into its homology, with zero differential, and contractible two-term complexes. To see this, first split cycles onto homology; then split each chain group onto the boundaries in the next degree, which are free by Lecture 3, Theorem 3.1.10. Tensoring preserves these contracting homotopies. The relative cross product sends oriented cell generators to the oriented product-cell generator. A cellular cycle of degree p corresponds to a class in $\mathrm{H}_p(\mathrm{sk}_p X;R)$. Naturality for the skeletal pairs therefore identifies products of cellular cycles with the singular cross product under Theorem 8.2.2. Thus the induced homology map is the displayed canonical map. Apply the product-cell isomorphism and the splittings to obtain the isomorphism. For the integral statement, iterate for finitely many factors and then tensor with A. No flatness assumption on A is needed. □

8.2.3 Even Cells and Products of Spheres

Definition 8.2.4. A finite cell complex is *even-cell* if every cell has even dimension.

Its cellular differential is zero. Products of even-cell complexes are again even-cell, because dimensions add. Thus if c_{2r} is the number of $2r$-cells, then

$$\mathrm{H}_{2r}(X;A) \cong A^{\oplus c_{2r}}, \qquad \mathrm{H}_{2r+1}(X;A) = 0.$$

This is a condition on cells. If instead one only knows that integral homology is even and free, Theorem 8.2.3 gives the corresponding product calculation.

For $n_1, \dots, n_r > 0$, give each sphere S^{n_i} its one 0-cell and one n_i-cell. Each differential is zero, so product cells give

$$\mathrm{H}_k\left(\prod_{i=1}^{r} \mathrm{S}^{n_i}; A\right) \cong \bigoplus_{\substack{J\subseteq\{1,\dots,r\} \\ \sum_{j\in J} n_j = k}} A.$$

The special case of products of circles gives the torus calculation.

Theorem 8.2.5. For $r \geq 0$ and every abelian group A, the torus $T^r = (\mathrm{S}^1)^r$ has homology

$$\mathrm{H}_k(T^r;A) \cong A^{\oplus \binom{r}{k}} \qquad (0 \leq k \leq r),$$

and zero homology in other degrees. Over $\mathbb{Z}$, the isomorphism is the iterated graded cross product of the circle homology groups.

Proof. The integral homology of a circle is $\mathbb{Z}$ in degrees zero and one and vanishes otherwise. By Theorem 8.2.3, the degree-k part of its r-fold graded tensor product has one generator for every choice of k factors contributing degree one. There are $\binom{r}{k}$ such choices. Tensoring with A gives the claim. Equivalently, the two endpoints of the 1-cell of a circle attach to the same vertex with opposite signs; the product boundary formula gives zero differential in every torus cell. □

8.3 Projective Spaces and Mapping Tori

8.3.1 Complex Projective Spaces

Definition 8.3.1. *Complex projective space* CP^n is the space of complex lines in $\mathbb{C}^{n+1}$.

The standard inclusions give a filtration $\mathrm{CP}^0 \subseteq \mathrm{CP}^1 \subseteq \cdots \subseteq \mathrm{CP}^n$.

Theorem 8.3.2. For $n \geq 0$ and an abelian group A,

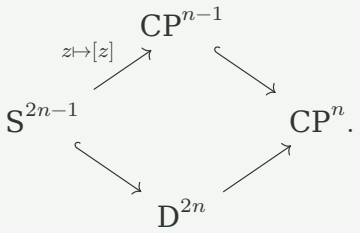
$$\mathrm{H}_k(\mathrm{CP}^n;A) \cong \begin{cases} A, & k = 0, 2, \dots, 2n, \\ 0, & \text{otherwise.} \end{cases}$$

For $n \geq 1$, the standard inclusion $\mathrm{CP}^{n-1} \to \mathrm{CP}^n$ induces isomorphisms in degrees below $2n$.

Proof. For $n > 0$, the lines with last coordinate nonzero form an affine chart $\mathbb{C}^n$. A characteristic map for this cell is

$$\mathrm{D}^{2n} \longrightarrow \mathrm{CP}^n, \qquad z \longmapsto [z : \sqrt{1 - \|z\|^2}],$$

where $z \in \mathbb{C}^n$ and $\|z\| \le 1$. On the boundary it is $z \mapsto [z : 0]$, so it gives the pushout



On the boundary $\|z\| = 1$, this map is $z \mapsto [z : 0] \in \mathrm{CP}^{n-1}$. On the interior it identifies with the homeomorphism $z \mapsto z/\sqrt{1 - \|z\|^2}$ onto the affine chart $\mathbb{C}^n$; its inverse is $w \mapsto w/\sqrt{1 + \|w\|^2}$. The induced map from the displayed pushout to CP^n is therefore bijective: its interior points give the chart, and its boundary identifications give exactly CP^{n-1}. The pushout is compact, and CP^n is Hausdorff, so this continuous bijection is a homeomorphism. This proves the pushout assertion. Since the sphere inclusion into the disk is a cofibration, the square is also a homotopy pushout. Starting with $\mathrm{CP}^0 = *$ gives one cell in each dimension $0, 2, \dots, 2n$. Its cellular differential is zero. The cellular theorem proves the formula and the assertion about inclusions. □

For example, $\mathrm{CP}^m \times \mathrm{CP}^n$ has one copy of A in degree $2k$ for each pair (i, j) with $0 \le i \le m$, $0 \le j \le n$, and $i + j = k$; its odd homology vanishes. For a chosen finite cell structure on X, its *cell-count polynomial* is $\sum_k c_k t^k$, where c_k is the number of k-cells. Product-cell dimensions add, so multiplication of these polynomials counts product cells: the coefficient of t^k is $\sum_{i+j=k} c_i(X) c_j(Y)$. Thus the cell-count polynomial of a finite product of complex projective spaces is

$$\prod_{i=1}^{r} (1 + t^2 + \cdots + t^{2n_i}) \quad \text{for} \quad \prod_{i=1}^{r} \mathrm{CP}^{n_i}.$$

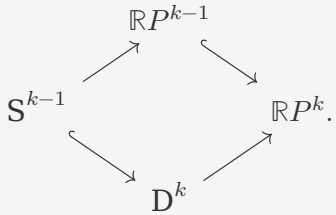
Each factor records the one cell in every even dimension of the corresponding projective space. In general a cell-count polynomial depends on the chosen cell structure and is not the homology-rank polynomial. Here every cellular differential is zero, so its coefficients also give the integral Betti numbers. The coefficient of t^k counts the copies of A in degree k.

8.3.2 Real Projective Spaces

Definition 8.3.3. *Real projective space* $\mathbb{R}P^n$ is the space of real lines in $\mathbb{R}^{n+1}$.

Its standard filtration has one cell in every dimension from 0 through n; unlike the complex case, adjacent cell dimensions occur.

The upper hemisphere of S^k gives a characteristic disk for the k-cell. On its boundary, antipodal points are identified, giving the pushout



For $k \ge 2$, let $q_{k-1} : \mathrm{S}^{k-1} \to \mathbb{R}P^{k-1}$ send a unit vector to its line, and let $c : \mathbb{R}P^{k-1} \to \mathbb{R}P^{k-1}/\mathbb{R}P^{k-2}$ be the collapse map. The characteristic disk identifies this quotient with $\mathrm{D}^{k-1}/\mathrm{S}^{k-2} \cong \mathrm{S}^{k-1}$. Write $f_k = c q_{k-1}$ for the map whose degree is the cellular coefficient.

The two open hemispheres $U_+ = \{u_k > 0\}$ and $U_- = \{u_k < 0\}$ in $\mathrm{S}^{k-1} \subseteq \mathbb{R}^k$ map homeomorphically under f_k onto the complement V of the collapsed point. Indeed, a line outside $\mathbb{R}P^{k-2}$ has a unique unit representative with last coordinate positive, and a unique one with last coordinate negative; these representatives depend continuously on the line. Choose y in that open cell, and write its inverse image as $\{x, -x\}$ with $x \in U_+$. The local-degree formula of Lemma 8.1.11, gives

$$\deg(f_k) = \varepsilon_+ + \varepsilon_-.$$

Let $U_+^\times$, $U_-^\times$, and $V^\times$ be U_+, U_-, and V with x, $-x$, and y, respectively, removed. Let $a(u) = -u$. Since $q_{k-1}a = q_{k-1}$, the two local maps of f_k satisfy the commutative square

$$\begin{array}{ccc} \mathrm{H}_{k-1}(U_+, U_+^\times; \mathbb{Z}) & \xrightarrow[\cong]{a_*} & \mathrm{H}_{k-1}(U_-, U_-^\times; \mathbb{Z}) \\ \downarrow{\scriptstyle (f_k|_{U_+})_*} & & \downarrow{\scriptstyle (f_k|_{U_-})_*} \\ \mathrm{H}_{k-1}(V, V^\times; \mathbb{Z}) & = & \mathrm{H}_{k-1}(V, V^\times; \mathbb{Z}). \end{array}$$

The local generators are the images of the global generator, as in the proof of that lemma. Naturality of these images shows that a_* multiplies them by the global degree of a, which is $(-1)^k$ by Lemma 8.1.10. The square therefore gives $\varepsilon_+ = (-1)^k \varepsilon_-$. Choose the orientation of the k-cell so that $\varepsilon_+ = 1$. It follows that $\varepsilon_- = (-1)^k$ and hence $\deg(f_k) = 1 + (-1)^k$.

For $k = 1$, the two endpoints of the characteristic interval attach to the same vertex. The relative boundary of the interval is its terminal endpoint minus its initial endpoint, so their images cancel. The coefficient formula holds over every abelian group A by naturality in coefficients. Therefore

$$\partial_k = 1 + (-1)^k = \begin{cases} 2, & k \text{ even}, \\ 0, & k \text{ odd}, \end{cases} \qquad (1 \le k \le n),$$

$$\mathrm{C}_k^{\mathrm{cell}}(\mathbb{R}P^n; A) = A \quad (0 \le k \le n).$$

Set $\partial_0 = 0$. The orientation of each even-dimensional cell may be chosen to give $+2$; this does not change neighboring nonzero differentials, since the odd differentials vanish.

Theorem 8.3.4. The cellular differential of real projective space alternates between zero and multiplication by two. Write $A[2] = \{a \in A : 2a = 0\}$. For $n \ge 1$, the resulting homology is

$$\mathrm{H}_k(\mathbb{R}P^n; A) \cong \begin{cases} A, & k = 0, \\ A/2A, & 0 < k < n \text{ and } k \text{ odd}, \\ A[2], & 0 < k < n \text{ and } k \text{ even}, \\ A, & k = n \text{ and } n \text{ odd}, \\ A[2], & k = n \text{ and } n \text{ even}, \\ 0, & \text{otherwise}. \end{cases}$$

For $n = 0$, projective space is a point.

Proof. The attaching-degree calculation gives the cellular complex. In an odd degree below the top, its homology is $A/2A$; in a positive even degree it is $\ker(2 : A \to A)$. At the top there is no incoming differential. The cellular theorem gives all the stated groups. □

With integral coefficients, the positive groups below the top are $\mathbb{Z}/2\mathbb{Z}$ in odd degrees and zero in even degrees. The top group is $\mathbb{Z}$ for odd n and zero for positive even n. With coefficients in $\mathbb{Z}/2\mathbb{Z}$, every differential vanishes, so every degree from 0 to n contains one copy of $\mathbb{Z}/2\mathbb{Z}$.

For instance, $\mathbb{R}P^2$ has integral groups $\mathbb{Z}$, $\mathbb{Z}/2\mathbb{Z}$, and 0 in degrees $0,1,2$, while its mod-two homology is nonzero in all three degrees. These are computations of groups; a cohomology ring calculation additionally requires a cup product.

8.3.3 Mapping Tori

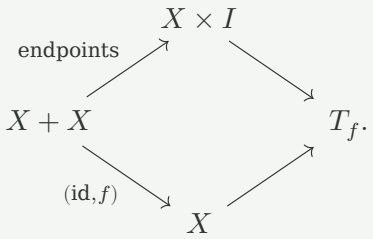
A final gluing calculation distinguishes the torus from the Klein bottle.

Definition 8.3.5. For a continuous map $f : X \to X$, let $i_0, i_1 : X \to X \times I$ be the endpoint inclusions. Its *mapping torus* is the coequalizer in Top

$$X \underset{i_0 \circ f}{\overset{i_1}{\rightrightarrows}} X \times I \xrightarrow{\;q\;} T_f.$$

Thus $qi_1 = qi_0 f$, and every continuous map $g : X \times I \to Y$ satisfying $gi_1 = gi_0 f$ factors uniquely as $g = \bar{g}q$ with $\bar{g} : T_f \to Y$ continuous. Concretely, q identifies $(x, 1)$ with $(f(x), 0)$, and T_f carries the quotient topology.

It is the pushout



The endpoint inclusion is a cofibration, so Lecture 7, Proposition 7.1.6, makes this a homotopy pushout. After retracting $X \times I$ onto X, its Mayer–Vietoris map has matrix

$$\begin{pmatrix} 1 & 1 \\ 1 & f_* \end{pmatrix}.$$

Subtract the first row from the second and then the first column from the second. This separates an identity summand from $f_* - 1$. Cancelling the identity summand in the exact sequence, and changing the sign of the remaining summand, gives the *Wang sequence*

$$\cdots \longrightarrow \mathrm{H}_k(X; A) \xrightarrow{1 - f_*} \mathrm{H}_k(X; A) \longrightarrow \mathrm{H}_k(T_f; A) \longrightarrow \mathrm{H}_{k-1}(X; A) \xrightarrow{1 - f_*} \mathrm{H}_{k-1}(X; A) \longrightarrow \cdots.$$

If $X = \mathrm{S}^1$ and f has degree d, this gives

$$\mathrm{H}_1(T_f; A) \cong A \oplus A/(d-1)A, \qquad \mathrm{H}_2(T_f; A) \cong \{a \in A : (d-1)a = 0\},$$

with $\mathrm{H}_0(T_f; A) = A$ and higher groups zero. Naturality identifies the quotient map in degree one with the map induced by $T_f \to \mathrm{S}^1$. Choose $x \in \mathrm{S}^1$ and a path γ from $f(x)$ to x. The map $t \mapsto [\gamma(t), t]$ defines a section, since its endpoints are identified in T_f. It therefore splits the sequence in degree one.

For the identity map, T_f is the torus. For reflection of the circle, it is the Klein bottle, and $d = -1$ gives

$$\mathrm{H}_1(T_f; A) \cong A \oplus A/2A, \qquad \mathrm{H}_2(T_f; A) \cong A[2].$$

In particular, the Klein bottle has integral groups $\mathrm{H}_1(T_f; \mathbb{Z}) = \mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z}$ and $\mathrm{H}_2(T_f; \mathbb{Z}) = 0$, but its mod-two top homology is $\mathbb{Z}/2\mathbb{Z}$.

Homotopy Groups and the Fundamental Group Action

Algebraic Topology I | Fall 2026 | October 2026

9.1 Based Homotopy Groups

The homology groups studied in Lectures 5–8 measure cycles modulo boundaries. Homotopy groups start with a different geometric object: a map from a sphere. Two such maps represent the same element when one can be continuously deformed into the other while a chosen point remains fixed. Keeping track of that point is essential. It makes multiplication possible, and the ambiguity in moving it is itself measured by the fundamental group.

This lecture begins Part III of the syllabus. We construct homotopy groups, prove their basic properties, and explain their dependence on basepoints. We then state the Whitehead and Hurewicz theorems, whose proofs belong to Math 616. The sphere calculations at the end use the statement of Hurewicz and the integral homology of spheres already computed in Lecture 7. Throughout, $I = [0, 1]$; all maps and homotopies are continuous. Every occurrence of space homology includes its coefficient group explicitly.

9.1.1 Based Maps and Homotopy Classes

A basepoint is part of the data of a space in this discussion. It must be preserved at every stage of a based homotopy, not just at its two endpoints.

Definition 9.1.1. A *based space* (X, x) is a space with a distinguished point $x \in X$. A based map $f : (X, x) \to (Y, y)$ is a map satisfying $f(x) = y$. A based homotopy from f to g is a map $H : X \times I \to Y$ satisfying

$$H(z, 0) = f(z), \qquad H(z, 1) = g(z), \qquad H(x, s) = y.$$

We write $[X, Y]_*$ for the set of based homotopy classes when the basepoints have been specified. An ordinary homotopy, without the last condition, is also called a *free homotopy* in this context.

Based homotopy is an equivalence relation. The constant homotopy proves reflexivity; replacing s by $1 - s$ proves symmetry; concatenating two homotopies in the parameter s proves transitivity. The pasting lemma proves continuity of concatenation because the two formulas agree at the joining time. These same elementary operations will occur repeatedly below.

Definition 9.1.2. The set $\pi_0(X, x)$ is the set of path components of X, with the component containing x distinguished. A based map induces a function of pointed sets by sending the component of z to the component of $f(z)$. In general, $\pi_0(X, x)$ has no natural group operation.

Equivalently, give the two-point space S^0 one distinguished point. A based map $\mathrm{S}^0 \to X$ is specified by the image of its other point, and a based homotopy is a path between those images. Thus $\pi_0(X, x) = [\mathrm{S}^0, X]_*$ as pointed sets. Path components, rather than connected components, are the appropriate invariant: a homotopy supplies actual paths.

For $n \geq 1$, write I^n for the n-cube and ∂I^n for the points having at least one coordinate equal to 0 or 1. A map of pairs $f : (I^n, \partial I^n) \to (X, x)$ means simply a map $f : I^n \to X$ with $f(\partial I^n) = \{x\}$.

Definition 9.1.3. For $n \geq 1$, define $\pi_n(X, x)$ to be the set of maps $f : (I^n, \partial I^n) \to (X, x)$ modulo homotopies relative to ∂I^n. Such a homotopy is a map $H : I^n \times I \to X$ with the prescribed endpoint maps and $H(u, s) = x$ whenever $u \in \partial I^n$. The constant map c_x gives the distinguished class.

The quotient $I^n/\partial I^n$ is a sphere S^n, with the collapsed boundary as basepoint. One way to see this is to identify the cube with a disk, radially from their centers, and then identify the disk with a closed hemisphere whose boundary is collapsed to the opposite pole. Equivalently, the quotient is the one-point compactification of the cube interior, which is homeomorphic to $\mathbb{R}^n$. The quotient universal property identifies cube maps with based sphere maps. It also identifies their homotopies: the quotient map times I is a quotient map because it is a continuous surjection between compact Hausdorff spaces. Consequently

$$\pi_n(X, x) = [\mathrm{S}^n, X]_*.$$

A representative can therefore be pictured either as a sphere in X or as a cube whose entire boundary maps to x.

For $n = 1$ this is a loop at x. The group we are about to construct is the *fundamental group*. The interval model has two endpoints, but both are sent to the same point; this condition makes one loop follow another.

9.1.2 Concatenation and the Group Axioms

Our convention is that paths and loops are multiplied in *travel order*: $a * b$ traverses a first and b second. Composition of maps retains the usual convention, so $g \circ f$ applies f first. Distinguishing these conventions prevents a reversal later in the basepoint action.

Construction 9.1.4. Write $u = (t, v)$ with $t \in I$ and $v \in I^{n-1}$. For based cube maps $f, g : I^n \to X$, set

$$(f * g)(t, v) = \begin{cases} f(2t, v), & 0 \leq t \leq \tfrac{1}{2}, \\ g(2t - 1, v), & \tfrac{1}{2} \leq t \leq 1. \end{cases}$$

The two formulas agree on $t = \frac{1}{2}$ because each has value x. They define a continuous based cube map. Define $[f][g] = [f * g]$.

Theorem 9.1.5. For every $n \geq 1$, concatenation makes $\pi_n(X, x)$ a group. Its identity is $[c_x]$, and the inverse of $[f]$ is represented by

$$\bar{f}(t, v) = f(1 - t, v).$$

Proof. First suppose F and G are based homotopies of f and g. Applying the concatenation formula to $F(-, s)$ and $G(-, s)$ for every s gives a based homotopy of the products. Thus the operation is well-defined on classes.

For associativity, allocate intervals of lengths $a, b, c > 0$, with $a + b + c = 1$, to the three maps f, g, h. On these intervals use respectively

$$f(t/a, v), \qquad g((t - a)/b, v), \qquad h((t - a - b)/c, v).$$

These formulas agree at the joining faces. The lengths for $(f * g) * h$ are $(1/4, 1/4, 1/2)$, while those for $f * (g * h)$ are $(1/2, 1/4, 1/4)$. Interpolate linearly between these triples. All lengths stay positive, so the piecewise formulas vary continuously and define a based homotopy. This proves associativity of classes; it does not assert equality of the original parametrized maps.

For the right identity, let $a_s = (1 + s)/2$ and put

$$R_s(t, v) = \begin{cases} f(t/a_s, v), & t \leq a_s, \\ x, & t \geq a_s. \end{cases}$$

At $s = 0$ this is $f * c_x$, and at $s = 1$ it is f. The closed-piece pasting lemma gives continuity also at $t = a_s$. The analogous left-identity homotopy keeps a constant interval $0 \le t \le (1-s)/2$ and rescales the remaining interval to run through f.

Finally, $f * \bar{f}$ runs the first coordinate from 0 to 1 and back to 0. A contraction of that excursion is

$$K(t, v, s) = \begin{cases} f(2t(1-s), v), & t \le \frac{1}{2}, \\ f(2(1-t)(1-s), v), & t \ge \frac{1}{2}. \end{cases}$$

At $s = 0$ it is $f * \bar{f}$; at $s = 1$ it is x, since $f(0, v) = x$. All boundary points stay at x. Replacing f by $\bar{f}$ proves that $\bar{f} * f$ also contracts. This proves both inverse identities and completes all group axioms. □

The same argument applies to paths with different endpoints whenever the concatenations are defined. In particular, a path followed by its reverse is homotopic relative to its endpoints to the constant path. The identity and associativity statements concern endpoint-preserving homotopy classes, exactly as for loops.

9.1.3 The Eckmann–Hilton Argument

When $n \ge 2$, there is a second coordinate in which to concatenate. The resulting interchange law is the geometric reason higher homotopy groups are abelian.

Lemma 9.1.6. Suppose a set has two operations $\star$ and $\circ$ with the same two-sided identity e, and suppose

$$(a \star b) \circ (c \star d) = (a \circ c) \star (b \circ d)$$

for every four elements. Then the operations coincide and are commutative.

Proof. Inserting identity elements in the interchange law gives

$$a \circ b = (a \star e) \circ (e \star b) = (a \circ e) \star (e \circ b) = a \star b.$$

Inserting them in the other two positions gives

$$a \circ b = (e \star a) \circ (b \star e) = (e \circ b) \star (a \circ e) = b \star a.$$

Together these equalities prove the assertions. □

Theorem 9.1.7. For $n \ge 2$, the group $\pi_n(X, x)$ is abelian, and concatenation in any coordinate induces the same operation.

Proof. Let $\star$ concatenate in the first coordinate and let $\circ$ concatenate in the second. The group-law proof applies to either coordinate, so both operations have identity $[c_x]$. Their interchange law holds already for representative maps. Subdivide the first two coordinates at $1/2$. Both sides send the four resulting subcubes through f, g, h, k, respectively, in the same positions and with the same linear rescalings. All joining faces map to x, so no extra homotopy is required. The preceding lemma proves that the two operations agree and commute. Repeating the argument with the first and the jth coordinates proves independence of the coordinate. □

We therefore use additive notation for $\pi_n(X, x)$ when $n \ge 2$: the constant class is 0, reflection represents the negative, and concatenation represents addition. We retain multiplicative notation for $\pi_1(X, x)$. A one-dimensional interval has no second independent coordinate, so the preceding argument says nothing about commutativity of the fundamental group.

9.2 Functoriality and Basepoint Transport

9.2.1 Functoriality and Elementary Examples

Postcomposition carries a sphere in one space to a sphere in another. The definitions make this operation compatible with both homotopies and concatenation.

Proposition 9.2.1. A based map $f : (X, x) \to (Y, y)$ induces

$$f_* : \pi_n(X, x) \longrightarrow \pi_n(Y, y), \qquad [a] \longmapsto [f \circ a].$$

For $n = 0$ this is a function of pointed sets; for $n \geq 1$ it is a homomorphism, and for $n \geq 2$ a homomorphism of abelian groups. Identity maps induce identities, and $(g \circ f)_* = g_* \circ f_*$.

Proof. Postcomposing a based homotopy of cube maps with f gives a based homotopy in Y, so the formula is independent of representatives. For positive n, the defining formulas give the equality of maps

$$f \circ (a * b) = (f \circ a) * (f \circ b),$$

which proves preservation of multiplication. Constant maps remain constant. Identity and composition laws follow from the corresponding equalities for ordinary maps. For $n = 0$, a path in X maps to a path in Y, and the distinguished component maps to the distinguished component. □

Proposition 9.2.2. If $f, g : (X, x) \to (Y, y)$ are based homotopic, then $f_* = g_*$ on $\pi_n(X, x)$ for all $n \geq 0$. A based homotopy equivalence induces isomorphisms in positive degrees and a bijection in degree zero.

Proof. Let H be a based homotopy from f to g. For a based cube map a, the formula $(u, s) \mapsto H(a(u), s)$ is a homotopy relative to the boundary between $f \circ a$ and $g \circ a$. For $n = 0$, the path $s \mapsto H(z, s)$ connects $f(z)$ to $g(z)$. If g is a based homotopy inverse of f, functoriality now gives $g_* f_* = (gf)_* = \mathrm{id}$ and $f_* g_* = \mathrm{id}$. □

Thus π_n is a functor on the based homotopy category, with target pointed sets, groups, or abelian groups according to the degree. This is the based counterpart of the homotopy-invariant functors considered in Lecture 5. An ordinary homotopy need not be based; its precise effect is the subject of the next section.

Before making difficult calculations, it is useful to record what the construction does to elementary geometric operations. These arguments work directly with representatives.

Proposition 9.2.3. For based spaces (X, x) and (Y, y), the projections induce a natural isomorphism

$$\pi_n(X \times Y, (x, y)) \cong \pi_n(X, x) \times \pi_n(Y, y)$$

for $n \geq 1$, and a bijection of pointed sets for $n = 0$.

Proof. A map into $X \times Y$ is exactly a pair of maps into the factors. The same is true of homotopies, and the boundary condition holds exactly when it holds in both factors. The correspondence on classes is consequently bijective. Concatenation is performed coordinate-wise in the target product, so it is a group isomorphism in positive degrees. For $n = 0$, two pairs of points are joined by a path precisely when their coordinates can be joined by paths; the two coordinate paths use the same parameter interval. □

Proposition 9.2.4. If $i : (A, a) \to (X, a)$ and $r : (X, a) \to (A, a)$ satisfy $r \circ i = \mathrm{id}_A$, then i_* is injective and r_* is surjective in every degree. If A is a deformation retract of X by a homotopy fixed on A, then these maps are inverse isomorphisms in positive degrees and inverse bijections in degree zero.

Proof. Functoriality gives $r_*i_* = \mathrm{id}$, proving the first assertions, including for pointed sets. In the deformation-retract case, ir is based homotopic to id_X, since $a \in A$ stays fixed. Hence $i_*r_* = \mathrm{id}$ as well. □

Example 9.2.5. A point has trivial homotopy groups in positive degrees and a one-element π_0. If X is discrete, every map from a connected cube is constant, so $\pi_n(X,x) = 0$ for $n \geq 1$, whereas $\pi_0(X,x)$ is the underlying pointed set of X. A convex subset of a real vector space contracts to any of its points by straight lines, fixing that point; its positive homotopy groups therefore vanish.

For arbitrary contractible spaces, a given contraction need not fix a prescribed basepoint. We will prove the same vanishing conclusion after accounting for moving basepoints. That small distinction is a useful test of whether a homotopy-invariance argument has actually handled its basepoints.

9.2.2 Transport Along a Path

Let $\alpha : I \to X$ be a path with $\alpha(0) = x$ and $\alpha(1) = y$. We want to carry a based sphere at x to one at y. Geometrically, shrink its cube and surround it by a collar on which the path runs from x on the inner edge to y on the outer edge.

For explicit formulas, let $c = (1/2, \dots, 1/2)$ and

$$r(u) = \max_{1 \leq i \leq n} |2u_i - 1|.$$

The equation $r(u) = 1$ describes the boundary; $r(u) \leq 1/2$ describes the concentric cube with half the side length. Scalar multiplication of $u - c$ below is coordinatewise.

Construction 9.2.6. For $f : (I^n, \partial I^n) \to (X, x)$, define $T_\alpha f : I^n \to X$ by

$$(T_\alpha f)(u) = \begin{cases} f(c + 2(u-c)), & r(u) \leq \frac{1}{2}, \\ \alpha(2r(u) - 1), & r(u) \geq \frac{1}{2}. \end{cases}$$

At the common boundary both expressions equal x, and on ∂I^n the value is y. Thus $T_\alpha f$ is a based cube at y. Set

$$\tau_\alpha : \pi_n(X,x) \longrightarrow \pi_n(X,y), \qquad \tau_\alpha([f]) = [T_\alpha f].$$

Applying the same formula to a based homotopy of f, or to an endpoint-preserving homotopy of α, proves that τ_α depends only on the two relevant homotopy classes. It remains to prove the algebraic properties. The following lemma is a convenient way to do all of them without relying on a picture of a moving collar.

Lemma 9.2.7. There is a homotopy from f to $T_\alpha f$ whose boundary has value $\alpha(s)$ at time s. More generally, if $H : I^n \times I \to X$ satisfies

$$H(u,0) = f(u), \quad H(u,1) = g(u), \quad H(u,s) = \alpha(s) \text{ for } u \in \partial I^n,$$

then $[g] = \tau_\alpha([f])$ in $\pi_n(X,y)$.

Proof. For the first assertion, put $a_s = 1 - s/2$ and define

$$M(u,s) = \begin{cases} f(c + (u-c)/a_s), & r(u) \leq a_s, \\ \alpha(2r(u) - 2 + s), & r(u) \geq a_s. \end{cases}$$

On the joining surface, the first argument lies on ∂I^n and the second argument is 0, so both values are x. The formula is continuous on its two closed regions, including at $s = 0$, since $a_s \geq 1/2$. At time 0 it gives f, at time 1 it gives $T_\alpha f$, and on the boundary it gives $\alpha(s)$.

For the general assertion, convert H into a homotopy with boundary fixed at y. Define

$$K(u,s) = \begin{cases} H(c + 2(u-c), s), & r(u) \le \frac{1}{2}, \\ \alpha(s + (1-s)(2r(u)-1)), & r(u) \ge \frac{1}{2}. \end{cases}$$

The joining values are $\alpha(s)$, and the outer boundary value is $\alpha(1) = y$. Hence K is a based homotopy at y. It starts at $T_\alpha f$ and ends at the map obtained by shrinking g into the central cube and filling the collar constantly with y. That last map is based homotopic to g: use the first formula of this proof for the constant path at y, and run the resulting homotopy backward. Therefore $[T_\alpha f] = [g]$. □

9.2.3 Composition, Inverses, and Naturality

Transport should compose in the same order that paths are traversed. The moving-boundary lemma proves this and also shows that transport respects the group operation.

Theorem 9.2.8. For every $n \ge 1$, transport is a group isomorphism. If α runs from x to y, β runs from y to z, and c_x is the constant path, then

$$\tau_{c_x} = \mathrm{id}, \qquad \tau_{\alpha * \beta} = \tau_\beta \circ \tau_\alpha, \qquad \tau_{\bar\alpha} = \tau_\alpha^{-1},$$

where $\bar\alpha(t) = \alpha(1-t)$. For any map $q : X \to Y$, transport is natural:

$$q_* \circ \tau_\alpha = \tau_{q\circ\alpha} \circ q_*.$$

Here the induced maps have the basepoints supplied by the displayed path endpoints.

Proof. For the constant path, the first homotopy in the moving-boundary lemma is based, proving $\tau_{c_x} = \mathrm{id}$. Apply that lemma first along α to f and then along β to $T_\alpha f$. Concatenating the two homotopies gives a moving-boundary homotopy from f to $T_\beta T_\alpha f$, with boundary path $\alpha * \beta$. The second part of the lemma gives

$$[T_\beta T_\alpha f] = \tau_{\alpha*\beta}([f]),$$

which is the composition formula. The paths $\alpha * \bar\alpha$ and $\bar\alpha * \alpha$ contract relative to their endpoints by the excursion contraction used in Theorem 9.1.5. Since transport depends only on the endpoint-preserving path class, the composition and identity formulas prove the inverse formula.

To prove preservation of multiplication, choose the moving-boundary homotopies M_f and M_g for two cubes at x. At each time, concatenate them in the first cube coordinate. Along the common face their values are both $\alpha(s)$, so the formulas paste continuously. The resulting homotopy runs from $f * g$ to $(T_\alpha f) * (T_\alpha g)$ and has boundary path α. The moving-boundary lemma gives

$$\tau_\alpha([f][g]) = \tau_\alpha([f])\tau_\alpha([g]).$$

Together with the inverse formula, this proves that transport is a group isomorphism. Finally, the collar formula gives the equality of representative maps $q \circ T_\alpha f = T_{q\circ\alpha}(q \circ f)$, proving naturality. □

For $n = 1$, reading the collar from the left endpoint to the right gives the especially useful formula

$$\tau_\alpha([f]) = [\bar\alpha * f * \alpha].$$

The three pieces have lengths $1/4, 1/2, 1/4$ in the collar model. Changing these positive lengths does not change their homotopy class, by the reparametrization argument in the group-law

proof. The order has a geometric meaning: a loop at y must first travel backward along α to reach the old basepoint x.

If X is path connected, all its groups $\pi_n(X,x)$ with fixed $n \geq 1$ are therefore isomorphic. A particular isomorphism still requires a path. The next section identifies exactly the ambiguity in that choice.

9.3 The Fundamental Group Action and Homotopy Invariance

9.3.1 A Left Action and Free Homotopy Classes

Returning a basepoint to its starting point produces an automorphism rather than a new group. Reversing the transporting loop converts this construction into a left action with the usual conjugation formula in degree one.

Definition 9.3.1. For $n \geq 1$, define a left action of $\pi_1(X,x)$ on $\pi_n(X,x)$ by

$$[a] \cdot u = \tau_{\bar{a}}(u).$$

For $n \geq 2$ these are automorphisms of an abelian group. In degree 1 the same formula is the action by inner automorphisms.

Proposition 9.3.2. The preceding formula is a well-defined left action by group automorphisms. In degree 1 it is

$$[a] \cdot [f] = [a][f][a]^{-1}.$$

For a based map $q : (X,x) \to (Y,y)$, induced maps respect the actions:

$$q_*([a] \cdot u) = q_*([a]) \cdot q_*(u).$$

Proof. Path-homotopy independence and the automorphism property follow from Theorem 9.2.8. The constant loop acts as the identity. Reversal of a concatenation gives $\overline{a * b} = \bar{b} * \bar{a}$, so

$$\tau_{\overline{a*b}} = \tau_{\bar{b}*\bar{a}} = \tau_{\bar{a}} \circ \tau_{\bar{b}}.$$

This says exactly $([a][b]) \cdot u = [a] \cdot ([b] \cdot u)$, the left-action law. In degree 1, apply the explicit transport formula to the path $\bar{a}$ to obtain $[a * f * \bar{a}]$. Naturality of transport, together with $q \circ \bar{a} = \overline{q \circ a}$, proves equivariance. □

Using τ_a itself would instead define a right action. Both conventions occur in the literature. Our convention is fixed by two requirements: paths multiply in travel order, and the left action on π_1 is conjugation $u \mapsto aua^{-1}$.

Proposition 9.3.3. Suppose α and β both run from x to y. Then

$$\tau_{\beta}^{-1} \circ \tau_{\alpha} = \tau_{\alpha * \bar{\beta}} \quad \text{on } \pi_n(X,x).$$

Consequently, if X is simply connected, transport between any two points is independent of the chosen path.

Proof. The inverse and composition formulas give $\tau_{\beta}^{-1}\tau_{\alpha} = \tau_{\bar{\beta}}\tau_{\alpha} = \tau_{\alpha*\bar{\beta}}$. By definition, simply connected means path connected with trivial fundamental group. In that case every loop at x, including $\alpha * \bar{\beta}$, represents the constant class, so the displayed automorphism is the identity.□

A path-connected space is called *simple* if the fundamental-group action on every positive homotopy group is trivial. In degree 1 this requires the fundamental group to be abelian.

Abelianity alone does not assert that the actions in higher degrees vanish. A simply connected space is automatically simple. The point of this terminology is that even abelian higher homotopy groups can carry additional information through their fundamental-group actions.

Write $[\mathrm{S}^n, X]$ for free homotopy classes of maps, with no basepoint condition on maps or homotopies. Forgetting the basepoint gives a function from $\pi_n(X, x)$ to this set.

Proposition 9.3.4. If X is path connected and $n \geq 1$, forgetting the basepoint induces a bijection

$$\pi_n(X, x)/\pi_1(X, x) \longrightarrow [\mathrm{S}^n, X],$$

where the left side is the set of orbits for the action just defined. In particular, free homotopy classes of loops correspond to conjugacy classes in $\pi_1(X, x)$.

Proof. First consider an arbitrary sphere map f. Its chosen sphere basepoint has some image $y \in X$. Choose a path from y to x. In the cube model, the moving-boundary homotopy transports f to a map based at x and is an ordinary sphere homotopy throughout. Thus forgetting the basepoint is surjective.

Now let f and g both be based at x, and suppose H is a free homotopy between them. The trace $a(s) = H(s_0, s)$ of the sphere basepoint s_0 is a loop at x. Pulling back to the cube gives a moving-boundary homotopy with boundary path a. Hence $[g] = \tau_a([f]) = [\bar{a}] \cdot [f]$, so the two based classes belong to the same orbit.

Conversely, if $[g] = [a] \cdot [f] = \tau_{\bar{a}}([f])$, the collar construction supplies a free homotopy from f to a representative of $[g]$. Follow it by a based homotopy to g. Thus the fibers of the forgetful map are exactly the orbits. In degree 1, Proposition 9.3.2 identifies these with conjugacy classes. □

For a simply connected target, the action is trivial, so every free homotopy class of sphere maps has a unique based homotopy class after a basepoint is specified. For a general target, forgetting the basepoint can genuinely identify different elements; the orbit set need not itself carry a useful group structure.

9.3.2 Homotopies and Homotopy Equivalences

A free homotopy carries its own basepoint comparison: follow the chosen point through the homotopy. This path supplies precisely the correction needed for the induced homomorphisms.

Theorem 9.3.5. Let $H : X \times I \to Y$ be a homotopy from f to g, and choose $x \in X$. The path $\alpha(s) = H(x, s)$ runs from $f(x)$ to $g(x)$. For $n \geq 1$,

$$g_* = \tau_\alpha \circ f_* : \pi_n(X, x) \longrightarrow \pi_n(Y, g(x)).$$

For $n = 0$, the two maps induce the same function on the underlying sets of path components.

Proof. For a based cube $a : I^n \to X$, the composite $(u, s) \mapsto H(a(u), s)$ has endpoint maps $f \circ a$ and $g \circ a$. On the cube boundary it equals $H(x, s) = \alpha(s)$. The moving-boundary lemma gives the formula. In degree zero, $s \mapsto H(z, s)$ is a path from $f(z)$ to $g(z)$ for every z, so both images lie in the same component. □

If f and g preserve a basepoint but the homotopy does not, α is a loop. Their induced homomorphisms need not be equal: the formula relates them by the action of its reversed class. This is why an unqualified assertion that “homotopic maps induce the same map on homotopy groups” is incomplete. Equality follows when the homotopy is based, or when the relevant transport action is trivial.

Even if neither a homotopy inverse nor its inverse homotopies preserve the selected basepoints, the trace formula still gives isomorphisms. The following argument keeps the intermediate basepoints explicit.

Corollary 9.3.6. If $f : X \to Y$ is an ordinary homotopy equivalence, then for every $x \in X$ the map

$$f_* : \pi_n(X, x) \longrightarrow \pi_n(Y, f(x))$$

is an isomorphism for $n \geq 1$, and f induces a bijection on path components.

Proof. Choose a homotopy inverse $g : Y \to X$. For each x, a homotopy $gf \simeq \mathrm{id}_X$ and the trace formula show that

$$g_* f_* : \pi_n(X, x) \longrightarrow \pi_n(X, gf(x))$$

is an isomorphism: composing it with the transport along the trace gives the identity. Thus f_* is injective at every basepoint. Applying the same argument to $fg \simeq \mathrm{id}_Y$ shows that g_* is injective at every basepoint of Y.

To see that f_* is surjective at the chosen x, let $v \in \pi_n(Y, f(x))$. Since $g_* f_*$ is surjective, choose $u \in \pi_n(X, x)$ with $g_* f_*(u) = g_*(v)$. Injectivity of g_* at $f(x)$ then gives $f_*(u) = v$. This proves the isomorphism without requiring either homotopy to fix a basepoint. The homotopies induce identities on component sets, so f and g give inverse bijections on π_0 as well. □

Example 9.3.7. If X is nonempty and contractible, the map $X \to *$ is a homotopy equivalence. Therefore $\pi_0(X, x)$ has one element and $\pi_n(X, x) = 0$ for every $n \geq 1$ and every choice of x. No assumption that a chosen contraction fixes x is needed.

9.4 The Whitehead and Hurewicz Theorems

9.4.1 Weak Equivalences and the Whitehead Theorem

A homotopy equivalence preserves all the invariants constructed so far. The converse needs geometric hypotheses on the spaces and a single map realizing the isomorphisms.

Definition 9.4.1. A map $f : X \to Y$ is a *weak homotopy equivalence* if it induces a bijection on path components and, for every $x \in X$ and every $n \geq 1$, an isomorphism

$$f_* : \pi_n(X, x) \longrightarrow \pi_n(Y, f(x)).$$

By naturality of transport, it is enough to test one basepoint in each path component of X.

The last assertion follows by choosing a path from the tested point to another point in its component. The two transport maps give a commutative square for f_* with isomorphisms on the vertical sides. Thus the isomorphism property is unchanged by the choice of basepoint within a component. The condition on π_0 must be stated separately; positive homotopy groups at one point cannot see other components.

Theorem 9.4.2. A weak homotopy equivalence $f : X \to Y$ between CW complexes is a homotopy equivalence. In particular, for connected CW complexes it is enough that f_* be an isomorphism on every positive homotopy group at one basepoint.

The proof of the Whitehead theorem is deferred to Math 616. The CW hypotheses are part of the theorem: for arbitrary spaces, a weak homotopy equivalence need not be a homotopy equivalence. The conclusion also requires an actual map inducing the isomorphisms; abstract isomorphisms between lists of homotopy groups are not the stated hypothesis. Recall from

Lecture 8 that CW complexes are built by attaching disks in successive dimensions, with their weak topology. That geometry is what makes the theorem possible.

As an immediate application of the statement, a nonempty CW complex with one path component and all positive homotopy groups zero is contractible. Indeed, its map to a point is a weak homotopy equivalence, and Whitehead makes it a homotopy equivalence. This is stronger than the already-proved fact that a contractible space has vanishing homotopy groups.

9.4.2 The Hurewicz Map and Theorem

Fix an orientation class $[\mathrm{S}^n] \in \mathrm{H}_n(\mathrm{S}^n;\mathbb{Z})$ for each $n \geq 1$. We use the cube orientation and the associated quotient sphere orientation. The natural way to turn a sphere map into a homology class is to push forward this class.

Construction 9.4.3. For $n \geq 1$, the *Hurewicz map* is

$$h_n : \pi_n(X,x) \longrightarrow \mathrm{H}_n(X;\mathbb{Z}), \qquad h_n([f]) = f_*([\mathrm{S}^n]).$$

It is well-defined because singular homology is homotopy invariant. Its target is ordinary integral singular homology, without a basepoint choice.

Proposition 9.4.4. The Hurewicz map is a natural homomorphism. It is unchanged by basepoint transport: for every path α from x to y,

$$h_n \circ \tau_\alpha = h_n.$$

In particular, it is constant on the orbits of the fundamental-group action.

Proof. The sphere model for concatenation is the pinch map

$$p : \mathrm{S}^n \longrightarrow \mathrm{S}^n \vee \mathrm{S}^n,$$

which collapses the middle face of the cube, together with its boundary. Orient the two quotient cubes by their inherited coordinates. The induced map on integral homology sends the sphere class to the sum of the two oriented sphere classes. To check the signs and coefficients, triangulate the cube so that its middle face is a subcomplex. Its oriented relative fundamental chain is the sum of the oriented chains of its two halves; the shared face cancels with opposite boundary orientations. After pinching, the two summands are exactly the respective sphere fundamental classes. This also works when $n = 1$, with the two positively oriented half-intervals.

The map representing the product is $(f \vee g) \circ p$, so pushing the last identity into X gives

$$h_n([f][g]) = h_n([f]) + h_n([g])$$

in multiplicative notation for the source. This proves the homomorphism assertion in all positive degrees. Naturality follows from $(qf)_* = q_* f_*$ on homology. Finally, f and $T_\alpha f$ are freely homotopic by Lemma 9.2.7; homotopy invariance of singular homology gives the transport identity. Replacing α by a loop or its reverse gives the orbit assertion. □

The range in which homotopy agrees with homology is controlled by the vanishing of lower homotopy groups. We first fix terminology so that the theorem's hypotheses are unambiguous.

Definition 9.4.5. A space is 0-*connected* if it is nonempty and path connected. For $r \geq 1$, it is r-*connected* if it is path connected and $\pi_i(X,x) = 0$ for $1 \leq i \leq r$. Transport makes this condition independent of the chosen basepoint. Thus 1-connected means simply connected.

Theorem 9.4.6. If $n \geq 2$ and X is $(n-1)$-connected, then $\widetilde{\mathrm{H}}_i(X;\mathbb{Z}) = 0$ for $0 \leq i < n$, and

$$h_n : \pi_n(X,x) \longrightarrow \mathrm{H}_n(X;\mathbb{Z})$$

is an isomorphism. If X is path connected, the degree-one map is surjective with kernel the commutator subgroup, and hence induces an isomorphism

$$\pi_1(X,x)^{\mathrm{ab}} \cong \mathrm{H}_1(X;\mathbb{Z}).$$

Both assertions of the Hurewicz theorem are stated here without proof; their proofs are deferred to Math 616. The degree-one assertion is also commonly called the first Hurewicz theorem. More explicitly, if $G = \pi_1(X,x)$ and $[G,G]$ is its commutator subgroup, it gives the exact sequence of groups

$$1 \longrightarrow [G,G] \longrightarrow G \xrightarrow{h_1} \mathrm{H}_1(X;\mathbb{Z}) \longrightarrow 0.$$

The final group is written additively. The kernel is exactly $[G,G]$, not merely a subgroup containing all commutators.

For $n = 2$, the hypothesis is simple connectivity and the conclusion is $\pi_2(X,x) \cong \mathrm{H}_2(X;\mathbb{Z})$. For larger n, all lower positive homotopy groups must vanish. The theorem compares the first possible nonzero positive homotopy group of a simply connected space with integral homology; it does not identify every homotopy group with homology in the same degree. The Hopf fibration in Lecture 10 will exhibit a nonzero third homotopy group of S^2, although $\mathrm{H}_3(\mathrm{S}^2;\mathbb{Z}) = 0$.

9.5 Sphere Calculations and the Next Step

9.5.1 Connectivity of Spheres Without Cellular Approximation

To apply the Hurewicz statement to spheres, we need a separate connectivity argument. The following proof uses only compactness, fine triangulations of a cube, and elementary Euclidean geometry. In particular, it does not presume the cellular approximation theorem.

Lemma 9.5.1. For $m \geq 1$, the sphere S^m is path connected, and

$$\pi_k(\mathrm{S}^m, x) = 0 \qquad (1 \leq k < m).$$

Thus S^m is $(m-1)$-connected.

Proof. The unit sphere in $\mathbb{R}^{m+1}$ is path connected for $m \geq 1$: nonantipodal points can be joined by normalizing their straight segment; antipodal points can be joined using any third point not on their common line. Fix $1 \leq k < m$ and a based cube map $f : I^k \to \mathrm{S}^m$ with boundary value x.

By uniform continuity, choose a finite triangulation of I^k, with its boundary a subcomplex, sufficiently fine that the values of f at any two points of one simplex are less than $1/2$ apart in Euclidean distance. Such triangulations are obtained by subdividing a cubical grid into simplices. Define $F : I^k \to \mathbb{R}^{m+1}$ by taking the value $f(v)$ at each vertex v and extending affinely over each simplex. The affine definitions agree on common faces. Since every point of a simplex is a convex combination of its vertices,

$$\|F(u) - f(u)\| < \tfrac{1}{2} \quad \text{for every } u \in I^k.$$

Also $F = x$ on the boundary, because all boundary vertices have value x. Every vector $(1-s)f(u)+sF(u)$ is within $1/2$ of the unit vector $f(u)$ and is therefore nonzero. Normalizing gives a based homotopy from f to

$$g(u) = \frac{F(u)}{\|F(u)\|}.$$

On any simplex, F takes values in the linear span of at most $k+1$ vertex vectors, and g takes values in the intersection of that span with the sphere. Since $k+1 \le m < m+1$, each such span is a proper linear subspace of $\mathbb{R}^{m+1}$. Finitely many proper linear subspaces cannot cover the sphere: each is annihilated by a nonzero linear functional, and the product of these functionals is a nonzero homogeneous polynomial. If it vanished on the unit sphere, homogeneity would make it vanish at every nonzero vector and hence everywhere, which is impossible. Choose $p \in \mathrm{S}^m$ outside all these spans, also different from x. The image of g avoids p.

Let $U_p = \{z \in \mathrm{S}^m : z \neq p\}$ be the sphere punctured at p. Stereographic projection gives a homeomorphism $\sigma : U_p \to \mathbb{R}^m$. The formula

$$(u,s) \longmapsto \sigma^{-1}\big((1-s)\sigma(g(u)) + s\sigma(x)\big)$$

is a homotopy from g to the constant map at x, fixed on the boundary. Combining the two homotopies proves $[f] = 0$. □

9.5.2 The First Nonzero Group and Degree

The connectivity lemma verifies the hypothesis of the stated Hurewicz theorem for each positive-dimensional sphere. In dimensions at least two, the previously computed integral homology now determines the first nonzero homotopy group.

Corollary 9.5.2. For $n \ge 2$, the Hurewicz map identifies

$$\pi_n(\mathrm{S}^n, x) \cong \mathrm{H}_n(\mathrm{S}^n; \mathbb{Z}) \cong \mathbb{Z},$$

and the identity sphere map represents the positive generator when source and target orientations agree. In particular, $\pi_2(\mathrm{S}^2) \cong \mathbb{Z}$ and $\pi_3(\mathrm{S}^3) \cong \mathbb{Z}$.

Proof. Lemma 9.5.1 verifies the $(n-1)$-connectedness hypothesis of the stated Hurewicz theorem. Lecture 7 computed the integral homology group in dimension n as $\mathbb{Z}$, generated by the chosen sphere orientation. By its definition, the Hurewicz map sends the identity class to that generator. □

This proof uses Hurewicz as a stated theorem, rather than supplying its deferred proof. The case $n = 1$ will be established in Lecture 11 by lifting loops along $\mathbb{R} \to \mathrm{S}^1$; the higher-dimensional argument above does not replace that lifting calculation.

Corollary 9.5.3. For $n \ge 2$, two based maps $f, g : \mathrm{S}^n \to \mathrm{S}^n$ are based homotopic if and only if they have the same degree. The same classification by an integer holds for free homotopy classes of sphere self-maps.

Proof. Degree was defined in Lecture 8, Definition 8.1.6 by $f_*([\mathrm{S}^n]) = \deg(f)[\mathrm{S}^n]$ in $\mathrm{H}_n(\mathrm{S}^n;\mathbb{Z})$. Thus the preceding Hurewicz isomorphism sends $[f]$ to $\deg(f)$, giving the based classification and realizing every integer. For $n \ge 2$, the target sphere is simply connected by Lemma 9.5.1. Proposition 9.3.4 therefore identifies based and free homotopy classes, proving the second assertion. □

These results prepare the exact-sequence methods of Lecture 10. A fibration will compare the homotopy groups of its fiber, total space, and base, with basepoints and fundamental-group actions controlling the low-degree end. In the Hopf fibration, the group $\pi_3(\mathrm{S}^3) \cong \mathbb{Z}$ just

established becomes the source of a computation of $\pi_3(S^2)$, a group that cannot be read from $H_3(S^2;\mathbb{Z})$.

9.5.3 References and Scope

The core constructions and proofs above are self-contained, using homotopy invariance of singular homology and sphere homology from earlier lectures where indicated. The two major comparison theorems are deliberately statements only. No proof of Whitehead or Hurewicz is included, and no relative Hurewicz theorem, cellular approximation theorem, or covering-space classification is assumed.

For further reading, consult the following author-hosted texts. The page numbers are the printed page numbers; the links open the corresponding PDF locations.

1. J. P. May, *A Concise Course in Algebraic Topology*: Chapter 9, Sections 1 and 5, pp. 65 and 68–69, for homotopy groups and basepoint change; Chapter 10, Section 3, pp. 75–76, for Whitehead; and Chapter 15, Section 1, pp. 117–118, for the Hurewicz map and both forms of the theorem. The collar proofs here do not require the fibration machinery used in May's basepoint discussion.
2. Allen Hatcher, *Algebraic Topology*, provides supplementary statements in Theorem 4.5, p. 346, for Whitehead, Theorem 4.32, pp. 366–367, for Hurewicz, and Section 2.A, pp. 166–168, for abelianization in degree one. The conventions and proofs in this lecture follow the course's own map-based development.

Fibrations and the Homotopy Exact Sequence

Algebraic Topology I | Fall 2026 | November 2026

10.1 Lifting Homotopies

A map $p : Y \to X$ describes a space parametrized over X: the space lying over $x \in X$ is the fiber $Y_x = p^{-1}(x)$. A fibration is a parametrized space with a homotopy lifting property. This property governs how the fibers vary along paths in X, by allowing a deformation in the base to be followed upstairs from a prescribed initial lift. The lifts need not be unique, so this does not specify a canonical transport map between the fibers.

For a chosen point $y_0 \in Y$ over $x_0 \in X$, put $Z = Y_{x_0}$. The resulting fiber sequence has the order $Z \to Y \to X$, formally dual to the order $X \to Y \to Z$ in a cofiber sequence. Its homotopy exact sequence relates the fiber Z, the total space Y, and the base X.

We use the based homotopy groups of Lecture 9. Write $I = [0, 1]$, and give path spaces the compact-open topology. Since I is compact Hausdorff, a map from any parameter space into the compact-open path space is continuous exactly when its evaluation map is continuous. We use this exponential law when constructing families of paths.

10.1.1 The Homotopy Lifting Property

The lifting condition specifies both the deformation in the base and its initial value upstairs.

Definition 10.1.1. A map $p : Y \to X$ has the *homotopy lifting property* for a space K if every commutative square of solid arrows admits a dotted lift:

$$\begin{array}{ccc} K \times \{0\} & \xrightarrow{u} & Y \\ \downarrow & \overset{\widetilde{H}}{\nearrow} & \downarrow p \\ K \times I & \xrightarrow[H]{} & X. \end{array}$$

A *Hurewicz fibration* has this property for every space K. A *Serre fibration* has it for every disk D^m, including $m = 0$. Thus every Hurewicz fibration is a Serre fibration. In this lecture, the unqualified word fibration will mean a Serre fibration.

The condition for $m = 0$ is existence of path lifts with prescribed starting point. It does not assert uniqueness. For example, a path in X has many lifts under the product projection $X \times Z \to X$ when the fiber Z has nonconstant paths.

Lemma 10.1.2. Let $p : Y \to X$ be a Serre fibration and (K, A) a finite CW pair. Given a homotopy $H : K \times I \to X$ and a lift on

$$(K \times \{0\}) \cup (A \times I),$$

there is a lift on $K \times I$ extending the given lift.

Proof. Attach the cells of K relative to A, one at a time. For an m-cell, the required extension problem has prescribed lift on the bottom and sides of $\mathrm{D}^m \times I$. This subset is a closed m-disk in the boundary of the $(m+1)$-disk $\mathrm{D}^m \times I$. The pair consisting of the prism and its bottom and sides is homeomorphic to the pair consisting of the same prism and its bottom. For $m > 0$, rescale the cylinder to $\mathrm{D}^m \times [-1, 1]$ and radially identify its boundary with S^m. The bottom

and the bottom together with the sides become latitude caps with boundary heights $-1/\sqrt{2}$ and $1/\sqrt{2}$. An increasing piecewise linear change of the latitude coordinate taking the first height to the second, with the longitudinal direction unchanged, gives the required boundary homeomorphism. Extend it radially over the disk. For $m = 0$ no change is needed. Applying the disk lifting property after this change of coordinates gives the desired extension. The attaching pushout then combines it with the lift already constructed. Finitely many such extensions finish the proof. □

This relative version is important: based homotopies must keep the basepoint fixed, and comparisons of lifts must preserve prescribed faces of a cube.

10.1.2 Bundles and Examples

Local product coordinates provide a concrete way to construct lifts, even when no global product description exists.

Definition 10.1.3. A *fiber bundle* with fiber Z is a map $p : Y \to X$ such that every point of X has an open neighborhood U and a homeomorphism

$$\phi : p^{-1}(U) \longrightarrow U \times Z, \qquad \mathrm{pr}_1 \, \phi = p.$$

Such a homeomorphism is a local trivialization. The fibers $p^{-1}(x)$ are therefore homeomorphic to Z.

Theorem 10.1.4. Every fiber bundle is a Serre fibration.

Proof. Consider $H : I^m \times I \to X$ and a lift at time zero. The inverse images under H of trivializing neighborhoods cover the compact cube. Choose a finite cubical subdivision, with a common sufficiently small mesh in all coordinates, so that every closed small cube is carried into one trivializing neighborhood. This is possible by the Lebesgue number lemma.

Proceed through the time intervals of this subdivision in order. On one time interval, proceed through the spatial cells in increasing dimension. Suppose the lift has been constructed on the bottom and sides of a spatial cell times this interval. Choose a trivialization containing its entire image under H. In this chart the prescribed lift is a map into Z on the bottom and sides. A prism retracts onto its bottom and sides. For the unit prism $\mathrm{D}^d \times I$, an explicit retraction is $r(x, t) = (\lambda x, 2 + \lambda(t - 2))$, where $\lambda = 2/\max\{2\|x\|, 2 - t\}$. It fixes the bottom and sides and lands in them; rescale the formula for the small prism. For a cubical cell use the maximum norm. Extend the fiber-coordinate map by composing with this retraction. Pair it with H and apply the inverse trivialization. This extends the lift and agrees with every previously constructed face. The finite pasting lemma supplies a continuous lift on each slab and then on the whole cube. Cubes and disks are homeomorphic, so this proves the Serre condition. □

Example 10.1.5. The exponential map $e : \mathbb{R} \to \mathrm{S}^1$, defined by $e(t) = \exp(2\pi i t)$, is a fiber bundle with discrete fiber $\mathbb{Z}$. On an open arc shorter than a full circle, choose a continuous argument. Its translates by integers give a trivialization. Lecture 11 will use the uniqueness of these lifts to compute $\pi_1(\mathrm{S}^1, 1)$.

Lemma 10.1.6. If $p : Y \to X$ is a Serre fibration and $f : T \to X$ is continuous, then the pullback projection $T \times_X Y \to T$ is a Serre fibration.

Proof. Compose a lifting problem over T with f and lift it to Y. The resulting lift and the original map to T have the same composite to X. The universal property of the pullback gives the required lift to $T \times_X Y$. □

10.2 The Homotopy Exact Sequence

Fix $y_0 \in Y$, put $x_0 = p(y_0)$, and let $Z = p^{-1}(x_0)$ with basepoint y_0. Write $i : Z \to Y$ for inclusion. The groups in the sequence depend on these choices, even when we suppress repeated basepoints in explanatory prose.

10.2.1 The Connecting Map

The connecting map records the failure of a lifted sphere to close up upstairs: the remaining face lies in the fiber.

Construction 10.2.1. For $n \geq 1$, represent a class in $\pi_n(X, x_0)$ by a map

$$u : (I^{n-1} \times I, \partial(I^{n-1} \times I)) \longrightarrow (X, x_0).$$

Lift u starting with the constant lift y_0 on the bottom and keeping the lift equal to y_0 on the sides, using Lemma 10.1.2. The top face of the lift is a based map $I^{n-1} \to Z$. Its class is

$$\partial : \pi_n(X, x_0) \longrightarrow \pi_{n-1}(Z, y_0).$$

For $n = 1$, this means the path component of the endpoint of a lifted loop.

Lemma 10.2.2. The connecting map is independent of the lift and of the representative. It is a homomorphism for $n \geq 2$ and is natural for maps of based fibrations.

Proof. Let u_s be a based homotopy between representatives, and choose lifts at both parameter endpoints. Apply relative lifting in the last cube coordinate t, with parameter pair

$$K = I^{n-1} \times I_s, \qquad A = (\partial I^{n-1} \times I_s) \cup (I^{n-1} \times \{0, 1\}).$$

Prescribe the constant lift on $K \times \{0\}$ and on the x-side faces, and the two chosen lifts on the s-end faces. These data agree on intersections. Lemma 10.1.2 extends them, and restriction to $t = 1$ gives a based homotopy in Z between the two top maps. Taking u_s constant also proves independence of the chosen lift.

For $n \geq 2$, concatenate representatives in one of the first $n-1$ coordinates. Their chosen lifts concatenate in that coordinate as well, and their top faces represent the product in $\pi_{n-1}(Z, y_0)$. Thus ∂ is a homomorphism. Finally, composing a chosen lift with a map of total spaces gives a lift of the corresponding map downstairs. Independence of choices proves naturality. □

10.2.2 Exactness and Components

A class vanishes after one map precisely when the preceding map can supply it. The cubical description makes each of these assertions a lifting or filling problem.

Theorem 10.2.3. For a based Serre fibration, the following sequence is exact:

$$\pi_n(Z, y_0) \xrightarrow{i_*} \pi_n(Y, y_0) \xrightarrow{p_*} \pi_n(X, x_0) \xrightarrow{\partial} \pi_{n-1}(Z, y_0) \xrightarrow{i_*} \pi_{n-1}(Y, y_0) \qquad (n \geq 2).$$

It continues as a sequence of pointed sets

$$\pi_1(Y, y_0) \xrightarrow{p_*} \pi_1(X, x_0) \xrightarrow{\partial} \pi_0(Z) \xrightarrow{i_*} \pi_0(Y) \xrightarrow{p_*} \pi_0(X).$$

Exactness at a pointed set means that the image of the preceding map is the inverse image of the distinguished point under the next map.

Proof. At $\pi_n(Y, y_0)$, a class coming from Z projects to the constant class. Conversely, if a based map $v : I^n \to Y$ projects to a nullhomotopic map, lift a based nullhomotopy of pv starting from v and fixed on the boundary. Its final map lies in Z and is based homotopic in Y to v.

At $\pi_n(X, x_0)$, a class lifted by a based map into Y has constant endpoint map in Z, so its boundary is zero. Conversely, if the top face of a lift of u is nullhomotopic in Z, append that nullhomotopy to the lift in the last coordinate. The result is a based cube in Y. Its projection is u followed by a constant slab, hence represents the original class.

At $\pi_{n-1}(Z, y_0)$, the top face of a lifted cube bounds that cube in Y, with all other faces at y_0. Thus its image in Y is zero. Conversely, a based nullhomotopy in Y of a map into Z, read from the constant map to that map, projects to a cube in X whose entire boundary is x_0. The original nullhomotopy is a lift, and its top face is the given map into Z.

The same arguments in dimension one use paths rather than higher cubes. In particular, a lifted loop ends in the component of y_0 in Z exactly when one can append a path in Z to obtain a loop in Y. At $\pi_0(Z)$, a point of Z is connected to y_0 in Y exactly when it is the endpoint of a path whose projection is a loop at x_0. At $\pi_0(Y)$, a point whose image is connected to x_0 can be joined, by lifting such a path, to a point of Z. These prove the remaining assertions. □

The component terms are generally not groups. Exactness there does not justify subtracting components or applying an abelian-group argument. If Z is path connected, however, the sequence immediately shows that $p_* : \pi_1(Y, y_0) \to \pi_1(X, x_0)$ is surjective.

10.3 The Path–Loop Fibration

10.3.1 Paths Beginning at a Basepoint

The most useful universal example comes from allowing all paths that start at one fixed point.

Construction 10.3.1. For a based space (X, x_0), define

$$PX = \{\gamma : I \to X \mid \gamma(0) = x_0\}, \qquad \Omega X = \{\gamma \in PX \mid \gamma(1) = x_0\}.$$

The endpoint map $\mathrm{ev}_1 : PX \to X$ has fiber ΩX over x_0. Both path spaces have the constant path as basepoint.

Proposition 10.3.2. The endpoint map $\mathrm{ev}_1 : PX \to X$ is a Hurewicz fibration, and PX is contractible through a contraction fixing the constant path.

Proof. Suppose $H : K \times I \to X$ is a homotopy and $u : K \to PX$ satisfies $u(k)(1) = H(k, 0)$. Define a path $\widetilde{H}(k, t)$ by first following $u(k)$ and then following $H(k, -)$ up to time t. To make its value at $t = 0$ exactly $u(k)$, use the formula

$$\widetilde{H}(k,t)(s) = \begin{cases} u(k)((1+t)s), & 0 \le s \le (1+t)^{-1}, \\ H(k, (1+t)s - 1), & (1+t)^{-1} \le s \le 1. \end{cases}$$

The two formulas agree at their common endpoint, and their values at $s = 1$ equal $H(k, t)$. Continuity follows from the pasting lemma for the evaluation map and the compact-open exponential law, since I is compact Hausdorff. This gives the required lift for every K.

The formula $C(\gamma, t)(s) = \gamma((1 - t)s)$ contracts PX to the constant path and fixes that path throughout. □

10.3.2 Looping Lowers the Homotopy Degree

A contractible total space leaves only the relation between the base and the fiber in the exact sequence.

Corollary 10.3.3. For $n \geq 2$, the path–loop fibration gives a natural isomorphism

$$\partial : \pi_n(X, x_0) \xrightarrow{\cong} \pi_{n-1}(\Omega X, \mathrm{const}_{x_0}).$$

For $n = 1$, it gives the natural bijection between $\pi_1(X, x_0)$ and the set $\pi_0(\Omega X)$.

Proof. Every positive homotopy group of PX vanishes because of its based contraction. The exact sequence therefore makes ∂ injective and surjective in positive degrees. At the component end, PX is path connected, and the same lifting argument identifies components of ΩX exactly with based homotopy classes of loops in X. □

A map $I^{n-1} \to \Omega X$ carrying ∂I^{n-1} to the constant loop is equivalently a map $I^{n-1} \times I \to X$ that is constant on the full boundary. This adjunction is the direct geometric explanation for the degree shift. In particular, looping does not simplify a space pointwise; it packages higher-dimensional based homotopies as lower-dimensional maps into a new space.

Corollary 10.3.4. For every $n \geq 2$, one has $\pi_n(\mathrm{S}^1, 1) = 0$.

Proof. Apply the fibration sequence to $\mathbb{Z} \to \mathbb{R} \xrightarrow{e} \mathrm{S}^1$. The based contraction of $\mathbb{R}$ makes its positive homotopy groups zero. Every positive homotopy group of the discrete space $\mathbb{Z}$ is zero, since the image of a connected cube is a point. Thus the exact segment

$$\pi_n(\mathbb{R}, 0) \longrightarrow \pi_n(\mathrm{S}^1, 1) \longrightarrow \pi_{n-1}(\mathbb{Z}, 0)$$

has zero outer terms for $n \geq 2$. □

10.4 The Hopf Fibration

10.4.1 The Bundle over the Projective Line

Regard S^3 as the unit sphere in $\mathbb{C}^2$. The complex projective line $\mathbb{CP}^1$ is the space of complex lines in $\mathbb{C}^2$, with the quotient topology. The quotient map below sends a unit vector to its line:

$$h : \mathrm{S}^3 \longrightarrow \mathbb{CP}^1, \qquad (z_0, z_1) \longmapsto [z_0 : z_1].$$

Its fibers are precisely the orbits of the scalar action of S^1 on S^3.

Proposition 10.4.1. The map h is a fiber bundle with fiber S^1, and $\mathbb{CP}^1$ is homeomorphic to S^2.

Proof. On the open chart $U_0 = \{[z_0 : z_1] \mid z_0 \neq 0\}$, use the coordinate $w = z_1/z_0$. A section is

$$s_0(w) = \frac{(1, w)}{\sqrt{1 + |w|^2}}.$$

The map $U_0 \times \mathrm{S}^1 \to h^{-1}(U_0)$ given by $(w, \lambda) \mapsto \lambda s_0(w)$ is a homeomorphism over U_0. Its inverse takes a unit vector to $(z_1/z_0, z_0/|z_0|)$. On $U_1 = \{[z_0 : z_1] \mid z_1 \neq 0\}$, the analogous section is $(v, 1)/\sqrt{1 + |v|^2}$, where $v = z_0/z_1$. These two trivializations prove the bundle assertion.

The map

$$(z_0, z_1) \longmapsto (2\,\mathrm{Re}(\overline{z}_0 z_1),\ 2\,\mathrm{Im}(\overline{z}_0 z_1),\ |z_0|^2 - |z_1|^2)$$

takes values in the unit sphere in $\mathbb{R}^3$ and is constant on scalar orbits. It therefore induces a continuous map $\mathbb{CP}^1 \to \mathrm{S}^2$. For a point $(x, y, z) \in \mathrm{S}^2$ with $z \neq -1$, the unique preimage has chart coordinate $w = (x + iy)/(1 + z)$. The remaining point $(0, 0, -1)$ has preimage $[0 : 1]$. Thus the induced map is bijective. The domain is compact as a quotient of S^3, and the target is Hausdorff. A continuous bijection from a compact space to a Hausdorff space is a homeomorphism. □

With this identification, $h : \mathrm{S}^3 \to \mathrm{S}^2$ is the *Hopf map*. Choose $y_0 = (1, 0)$ and its image x_0 as basepoints, and identify the fiber with S^1 by $\lambda \mapsto (\lambda, 0)$.

10.4.2 Computing the Third Homotopy Group

The circle fiber has no higher homotopy groups, so the Hopf bundle transfers the third homotopy group of its total space to its base.

Theorem 10.4.2. The Hopf map induces an isomorphism

$$h_* : \pi_3(\mathrm{S}^3, y_0) \xrightarrow{\cong} \pi_3(\mathrm{S}^2, x_0).$$

Consequently $\pi_3(\mathrm{S}^2, x_0) \cong \mathbb{Z}$, with generator represented by h after an orientation of S^3 is chosen.

Proof. The relevant five terms of the Hopf fibration sequence are

$$\pi_3(\mathrm{S}^1, 1) \longrightarrow \pi_3(\mathrm{S}^3, y_0) \xrightarrow{h_*} \pi_3(\mathrm{S}^2, x_0) \longrightarrow \pi_2(\mathrm{S}^1, 1) \longrightarrow \pi_2(\mathrm{S}^3, y_0).$$

The first and fourth terms vanish by Corollary 10.3.4. Exactness proves that h_* is an isomorphism.

Lemma 9.5.1 proves that S^3 is 2-connected. Corollary 9.5.2 applies the stated Hurewicz theorem to give

$$\pi_3(\mathrm{S}^3, y_0) \xrightarrow{\cong} \mathrm{H}_3(\mathrm{S}^3; \mathbb{Z}) = \mathbb{Z}.$$

The identity map represents the oriented generator under this isomorphism. Its image under h_* is the homotopy class of h. This last identification uses the Hurewicz theorem whose proof, as specified in the syllabus, is deferred to Math 616; the fibration argument above is proved in full here. □

This is a first indication that homotopy groups differ sharply from homology groups. Although $\mathrm{H}_3(\mathrm{S}^2; \mathbb{Z}) = 0$, the group $\pi_3(\mathrm{S}^2, x_0)$ is infinite cyclic. A map from a higher-dimensional sphere can carry essential homotopy information that its map on homology does not detect.

For further reading, see May, *A Concise Course in Algebraic Topology*, Chapter 7 for the lifting viewpoint and Chapter 9, Sections 3–4, pp. 66–67, for the fibration sequence and Hopf example. May works with Hurewicz fibrations; the disk-lifting arguments above isolate the weaker Serre hypothesis needed for homotopy groups. A supplementary comparison is Hatcher, *Algebraic Topology*, Theorem 4.41, pp. 376–377, and Proposition 4.48, pp. 379–380, for the Serre sequence and the local bundle argument.

The Fundamental Group and Its Applications

Algebraic Topology I | Fall 2026 | November 9 and 11, 2026

11.1 The Fundamental Group of the Circle

The fundamental group records the order in which a loop passes through different parts of a space. Our first calculation makes this order visible by lifting a loop in the circle to a path in the real line. We then prove a gluing theorem for fundamental groups and use it to compute groups of graphs and surfaces. The final applications concern roots of complex polynomials and antipodal points on spheres.

We retain the conventions of Lecture 9. The interval is $I = [0,1]$, and the product $\alpha * \beta$ traverses α first and β second. The reverse path is $\overline{\alpha}(t) = \alpha(1-t)$. Path homotopies fix both endpoints unless stated otherwise. For a based map f, the induced homomorphism is denoted by f_*. Write $\mathrm{S}^1 = \{z \in \mathbb{C} : |z| = 1\}$, with basepoint 1, and define

$$p : \mathbb{R} \longrightarrow \mathrm{S}^1, \qquad p(t) = e^{2\pi i t}.$$

A lift of a map $g : Y \to \mathrm{S}^1$ is a map $\tilde{g} : Y \to \mathbb{R}$ satisfying $p \circ \tilde{g} = g$. We establish the required lifting properties directly, before the general theory of covering spaces in Lecture 12.

11.1.1 Local Inverses and Lifting a Path

An angle can be chosen continuously on any sufficiently short open arc of the circle. The ambiguity between two choices is an integer. These two elementary facts provide both existence and uniqueness of lifts.

For $a \in \mathbb{R}$, let $J_a = (a - \frac{1}{4}, a + \frac{1}{4})$ and let $A_a = p(J_a)$. The restriction $p|_{J_a} : J_a \to A_a$ is a homeomorphism. Its inverse is a continuous choice of angle. Every point of $p^{-1}(A_a)$ belongs to a unique interval $J_a + k$, with $k \in \mathbb{Z}$, and restriction to each such interval is again a homeomorphism onto A_a. These assertions follow from the elementary identities

$$p(s) = p(t) \quad \Longleftrightarrow \quad s - t \in \mathbb{Z}$$

and the usual angle parametrization of an open semicircle.

Lemma 11.1.1. For every path $\gamma : I \to \mathrm{S}^1$ and every $a \in \mathbb{R}$ with $p(a) = \gamma(0)$, there is a unique lift $\tilde{\gamma} : I \to \mathbb{R}$ satisfying $\tilde{\gamma}(0) = a$.

Proof. The inverse images $\gamma^{-1}(A_b)$ form an open cover of the compact interval. The Lebesgue number lemma gives a subdivision

$$0 = t_0 < t_1 < \cdots < t_m = 1$$

such that the image of each closed subinterval lies in one arc A_b. On the first subinterval, choose the inverse of p whose value at $\gamma(0)$ is a. Once the lift has been constructed through t_{j-1}, there is a unique integer translate of an inverse on the next arc that agrees with this value. Use that inverse on $[t_{j-1}, t_j]$. The finitely many continuous pieces agree at their endpoints, so the pasting lemma gives a continuous lift on I.

For uniqueness, two lifts u, v have $u(t) - v(t) \in \mathbb{Z}$ for every t. Their difference is continuous and I is connected, so it is constant. Agreement at 0 makes that constant zero. □

The same uniqueness argument applies on any connected domain: two real-valued lifts that agree at one point agree everywhere. Existence, by contrast, needs an argument adapted to the domain.

11.1.2 Lifting a Homotopy on a Square

To distinguish loop classes, we must know that the endpoint of a lifted loop does not change during a based homotopy. Continuity in the homotopy parameter is the essential point; lifting each path separately would not yet prove it.

Lemma 11.1.2. Let $H : I \times I \to \mathrm{S}^1$ be continuous, with coordinates (s, t). Given a lift $u : I \to \mathbb{R}$ of $H(s, 0)$, there is a unique continuous lift $\widetilde{H} : I \times I \to \mathbb{R}$ satisfying $\widetilde{H}(s, 0) = u(s)$.

Proof. Apply the Lebesgue number lemma to the cover of the square by the sets $H^{-1}(A_a)$. Choose a rectangular grid fine enough that each closed rectangle Q is mapped into one such arc. Process the bottom row from left to right, then the next row from left to right, and continue upward.

On the first rectangle, choose an inverse branch over its arc and translate it by an integer so that its composition with H agrees with u at the lower-left corner. Along the bottom edge, both functions lift the same map. Their difference is integer-valued and continuous, so agreement at the corner gives agreement on the entire edge.

For each subsequent rectangle, its bottom edge is already lifted, either by u or by the preceding row. Its left edge is also already lifted unless it is the first rectangle in its row. Choose the integer translate of a local inverse that agrees at the lower-left corner. Uniqueness on a connected interval makes it agree along the entire bottom edge and, when present, the entire left edge. These edges contain its intersection with the previously processed rectangles. Thus each new formula agrees with the existing lift wherever both are defined. Finite closed pasting now gives a continuous lift on the whole square.

Two lifts of H differ by a continuous integer-valued function on the connected square. Their agreement on its bottom edge proves uniqueness. □

Corollary 11.1.3. For a based homotopy of loops at 1, the lifts that start at 0 all have the same endpoint.

Proof. Lift the homotopy by Lemma 11.1.2, starting with the lift of its bottom loop. Each vertical boundary edge has image in $p^{-1}(1) = \mathbb{Z}$. Its lift is continuous, hence constant. The left edge is therefore constantly 0, and the right edge is the common endpoint of all the lifted loops. □

11.1.3 Winding Number and Degree

A loop lifts to a path whose endpoints differ by an integer. This integer counts its signed turns around the circle. The calculation below identifies the group law as addition of these displacements.

Theorem 11.1.4. For a loop γ at 1, let $\tilde{\gamma}$ be its lift starting at 0. The assignment

$$w : \pi_1(\mathrm{S}^1, 1) \longrightarrow \mathbb{Z}, \qquad w([\gamma]) = \tilde{\gamma}(1)$$

is an isomorphism. The loop $\omega_n(t) = e^{2\pi i n t}$ represents the integer n.

Proof. The endpoint is an integer because $p(\tilde{\gamma}(1)) = 1$. Corollary 11.1.3 makes the assignment well-defined on based homotopy classes. If $\tilde{\gamma}(1) = m$ and $\tilde{\eta}(1) = n$, the lift of $\gamma * \eta$ first follows $\tilde{\gamma}$ and then follows $m + \tilde{\eta}$. Its endpoint is $m + n$, so w is a homomorphism.

The lift of ω_n starting at 0 is $t \mapsto nt$, which proves surjectivity. If the lift of γ ends at n, the straight-line homotopy

$$(t, r) \longmapsto (1 - r)\tilde{\gamma}(t) + rnt$$

fixes its endpoints 0 and n. Composing with p gives a based homotopy from γ to ω_n. Thus a loop is determined up to based homotopy by its endpoint displacement, proving injectivity. □

For a loop γ whose basepoint is any point of the circle, define its winding number by $\tilde{\gamma}(1)-\tilde{\gamma}(0)$. A different lift adds the same integer at both endpoints, so the definition is independent of the lift. If H is a homotopy through loops whose basepoints may move, square lifting still shows that

$$t \longmapsto \widetilde{H}(1,t) - \widetilde{H}(0,t)$$

is a continuous integer-valued function. Consequently winding number is invariant under these homotopies as well.

Recall that the degree of a continuous map $f : \mathrm{S}^1 \to \mathrm{S}^1$ is the integer by which it acts on $\mathrm{H}_1(\mathrm{S}^1;\mathbb{Z})$. The winding number of $t \mapsto f(e^{2\pi i t})$ computes this degree. Indeed, the degree-one Hurewicz theorem stated in Lecture 9 identifies $\pi_1(\mathrm{S}^1,1)$ with $\mathrm{H}_1(\mathrm{S}^1;\mathbb{Z})$, sending the positive loop to the positive homology generator. Naturality identifies the two multipliers. For a map not preserving 1, a rotation makes it based without changing either multiplier. The following proof obtains the properties needed here directly from winding number.

Proposition 11.1.5. Degree is invariant under homotopy. For continuous maps $f, g : \mathrm{S}^1 \to \mathrm{S}^1$, it satisfies

$$\deg(f \circ g) = \deg(f)\deg(g), \qquad \deg(z \mapsto z^n) = n \quad (n \in \mathbb{Z}).$$

If f extends to a continuous map $\mathrm{D}^2 \to \mathrm{S}^1$, then $\deg(f) = 0$.

Proof. Homotopy invariance follows from the preceding winding-number argument. The displayed power map has lift $t \mapsto nt$, giving its degree.

Lift $t \mapsto f(e^{2\pi i t})$ on $[0,1]$ and denote this lift by F. Write $d = \deg(f)$. Extend the lift to $\mathbb{R}$ by the rule

$$F(t+k) = F(t) + kd \qquad (0 \le t \le 1,\ k \in \mathbb{Z}).$$

The endpoint identity $F(1)-F(0) = d$ makes the formulas agree at integers, and hence defines a continuous function. In particular, $F(u+k) = F(u)+kd$ for every real u and integer k. Construct a lift G for g in the same manner, with displacement $e = \deg(g)$. The function $F \circ G$ lifts the composite, and its displacement is

$$F(G(1)) - F(G(0)) = F(G(0)+e) - F(G(0)) = de.$$

Finally, if $\overline{f} : \mathrm{D}^2 \to \mathrm{S}^1$ extends f, the maps $z \mapsto \overline{f}((1-r)z)$ give a homotopy from f to a constant map. Such a map has degree zero. □

11.2 Free Products and Relations

When a loop passes between two open sets, it breaks into a word of loops drawn from those sets. The relevant algebra must preserve the order of the letters. Free products provide the universal group in which words from the two factors can be multiplied without imposing extra commutativity relations.

11.2.1 Reduced Words and the Universal Property

For groups A and B, regard their nonidentity elements as distinct letters, tagged by their factor even when the underlying sets overlap. A word is reduced if consecutive letters come from different factors. The empty word is allowed.

Proposition 11.2.1. There is a group $A*B$, with homomorphisms $j_A : A \to A*B$ and $j_B : B \to A*B$, such that for every group K and every pair of homomorphisms $a : A \to K$ and $b : B \to K$, there is a unique homomorphism $h : A*B \to K$ satisfying $hj_A = a$ and $hj_B = b$. Every element of $A*B$ has a unique reduced-word expression. Both factor homomorphisms are injective.

Proof. Let $\mathcal{W}$ be the set of reduced words. For $g \in A$, define a permutation L_g of $\mathcal{W}$ by inserting g on the left and reducing at that end. More explicitly, L_1 is the identity. For $g \neq 1$, if the word is empty or starts with a B-letter, prepend g. If its first letter is $a_1 \in A$, replace that letter by ga_1 when this product is nonidentity, and delete it when the product is the identity. Define L_g for $g \in B$ in the same way.

For g, h in the same factor, these operations satisfy $L_g L_h = L_{gh}$. To verify this, a word starting in the other factor, or an empty word, acquires the single product gh, with that letter omitted if it is 1. If the first letter is in the same factor, the result is obtained by multiplying g, h and that first letter there. If h cancels the first letter, the next letter belongs to the other factor; prepending g gives the same result as first multiplying gh by the original first letter. These cases cover every word. It follows in particular that $L_g^{-1} = L_{g^{-1}}$.

Define $A * B$ to be the subgroup of the permutation group of $\mathcal{W}$ generated by these permutations. The preceding identities give the two factor homomorphisms. Every product of generators reduces by multiplying adjacent letters from the same factor and omitting identities. A reduced product $L_{g_1} \cdots L_{g_r}$ sends the empty word to $(g_1, \dots, g_r)$. Thus different reduced words represent different permutations. This proves existence and uniqueness of the normal form, including the injectivity of the factors.

Given a and b, send a reduced word to the ordered product of the images of its letters in K. Combining adjacent letters from the same factor or omitting an identity does not change that product. Hence the assignment respects multiplication and defines h. Since the two factors generate $A * B$, this homomorphism is unique. □

The construction works for any set of groups by allowing letters tagged by an arbitrary factor and requiring consecutive factors to differ. The same left-insertion proof gives unique reduced words. In particular, for a set S, the free group $F(S)$ is the free product of infinite cyclic groups with generators indexed by S. Specifying a homomorphism $F(S) \to K$ is equivalent to choosing an arbitrary element of K for each member of S.

Example 11.2.2. For two generators a, b, the word $aba^{-1}b^{-1}$ is a nonempty reduced word in $F(\{a, b\}) = \mathbb{Z} * \mathbb{Z}$. It is therefore nonidentity. Thus a free product of abelian groups need not be abelian; this is the algebraic distinction between a bouquet of two circles and a torus.

11.2.2 Normal Closures and Amalgamation

Attaching a disk makes its boundary loop nullhomotopic. A group quotient recording this fact must also make every conjugate of that loop trivial. This explains why normal closures, rather than ordinary generated subgroups, occur in topological presentations.

For a subset R of a group G, write $\langle\langle R \rangle\rangle$ for its normal closure, the smallest normal subgroup containing R. Concretely, its elements are finite products of conjugates $gr^{\varepsilon}g^{-1}$, with $g \in G$, $r \in R$, and $\varepsilon \in \{1, -1\}$. These products form a normal subgroup containing R, and every normal subgroup containing R contains them. A presentation $\langle S \mid R \rangle$ denotes $F(S)/\langle\langle R \rangle\rangle$, where the elements of R are words in the generators and their inverses.

Definition 11.2.3. For homomorphisms $u : C \to A$ and $v : C \to B$, the amalgamated free product is the quotient

$$A *_C B = (A * B) \big/ \langle\langle j_A(u(c)) j_B(v(c))^{-1} : c \in C \rangle\rangle.$$

The maps u and v need not be injective.

The universal property of the free product and the quotient says that homomorphisms $A *_C B \to K$ correspond to homomorphisms $A \to K$ and $B \to K$ whose restrictions along u and v agree. Equivalently, this quotient is the pushout of the two homomorphisms from C in the

category of groups. Even when the original maps into the free product are injective, their composites into this quotient need not be.

11.3 The Seifert–van Kampen Theorem

An open cover lets us subdivide paths and homotopies into pieces lying in individual members of the cover. Path subdivision gives generators. Subdivision of a homotopy square proves that the overlap supplies all the relations. We keep both arguments explicit because the second is what makes the theorem a computation of the group, rather than only a list of generators.

11.3.1 The Pushout Statement

Let $X = U \cup V$, where U, V are open and path-connected, and suppose that $W = U \cap V$ is nonempty and path-connected. Fix $x_0 \in W$. Denote the inclusions from W by i_U, i_V, and those into X by j_U, j_V.

Theorem 11.3.1. Under these hypotheses, the square of induced homomorphisms

$$\begin{array}{ccc} \pi_1(W, x_0) & \xrightarrow{(i_U)_*} & \pi_1(U, x_0) \\ {\scriptstyle (i_V)_*}\downarrow & & \downarrow{\scriptstyle (j_U)_*} \\ \pi_1(V, x_0) & \xrightarrow[(j_V)_*]{} & \pi_1(X, x_0) \end{array}$$

is a pushout of groups. In particular, the natural homomorphism

$$\pi_1(U, x_0) *_{\pi_1(W, x_0)} \pi_1(V, x_0) \longrightarrow \pi_1(X, x_0)$$

is an isomorphism.

Write P for the displayed amalgamated free product. Inclusion of loops gives a homomorphism $\theta : P \to \pi_1(X, x_0)$, since the two images of every overlap loop agree in X. We will construct its inverse by assigning to each loop a subdivided word in P.

11.3.2 Subdivided Loop Words and Their Choices

Take a based loop $\gamma : I \to X$. The Lebesgue number lemma gives a subdivision $0 = t_0 < \cdots < t_m = 1$ such that every restricted path γ_j lies in one chosen set $C_j \in \{U, V\}$. Here and below, restrictions to subintervals are linearly reparametrized to I.

Choose paths λ_j from x_0 to $\gamma(t_j)$ as follows. At $j = 0, m$, use the constant path. At an interior division point where the neighboring colors differ, choose λ_j in W. This is possible because the point lies in both sets and W is path-connected. Where the colors agree, choose λ_j in their common set. The loop

$$\lambda_{j-1} * \gamma_j * \overline{\lambda_j}$$

then lies in C_j. Map its class into P, and define $\Psi(\gamma)$ to be the ordered product of these m classes.

Lemma 11.3.2. The element $\Psi(\gamma) \in P$ is independent of the subdivision, the colors C_j, and the paths λ_j.

Proof. First change one interior path from λ_j to μ_j. The loop $\lambda_j * \overline{\mu_j}$ lies in the common color when the two colors agree, and in W when they differ. Write d for its image in P. The jth factor changes from a_j to $a_j d$, while the next changes from a_{j+1} to $d^{-1} a_{j+1}$. In the different-color case,

the defining overlap relation in P identifies the two images of d. Thus the complete product is unchanged. Repeating this proves independence of all the chosen paths.

Next insert a division point in a segment of color C. Choose any connecting path to it in C, keeping the two new segments of color C. Multiplication within $\pi_1(C, x_0)$ cancels this connecting path and its reverse, so the two new factors have the same product as the old one. This proves invariance under subdivision and its inverse operation.

Finally, suppose that a segment can be assigned either color. Its entire image then lies in W. Before changing its color, choose the connecting paths at its endpoints in W, using the already proved path independence. Those paths are admissible for either color and for every neighboring segment. The factor associated to the segment is now represented by a loop in W, so the overlap relation identifies its two possible images in P.

Given any two initial choices, pass to the common refinement of their subdivisions. Each refined segment lies in both of its inherited colors if those colors disagree. Change these colors one at a time by the preceding argument, then change the connecting paths. All choices therefore give the same element. □

The product of the loops defining $\Psi(\gamma)$, viewed in X, is homotopic to γ: adjacent connecting paths occur as $\overline{\lambda_j} * \lambda_j$ and contract relative to their endpoints. Consequently $\theta(\Psi(\gamma)) = [\gamma]$. In particular, θ is surjective.

11.3.3 The Relations from a Homotopy Square

It remains to show that the subdivided word depends only on the based homotopy class. For this we label all oriented edges of a sufficiently fine grid and use the relation supplied by each small rectangle.

Proof. Let $H : I^2 \to X$ be a based homotopy from γ_0 to γ_1, with path coordinate horizontal and homotopy coordinate vertical. Subdivide the square by a rectangular grid such that each closed cell has image in one of U, V, and give the cell such a color. This is possible by applying the Lebesgue number lemma to $H^{-1}(U), H^{-1}(V)$.

At each grid vertex v, choose a path λ_v from x_0 to $H(v)$. If incident cells have both colors, choose it in W; if they all have one color, choose it in that set. On the two vertical boundary edges every vertex has image x_0, and we choose the constant paths there. Each choice is possible by the assumed path-connectedness. Moreover, the chosen path at a vertex lies in the color of every cell incident to that vertex.

For an oriented grid edge $e : v \to w$, let H_e be the restriction of H along it and assign to it the element

$$a_e = [\lambda_v * H_e * \overline{\lambda_w}] \in P,$$

computed in either incident cell's group. If the two incident colors differ, the edge image and its two connecting paths lie in W, so the two answers agree in P. Reversing the edge replaces a_e by a_e^{-1}.

For one small rectangle, denote the bottom and top edge labels, oriented to the right, by a and d, and the left and right labels, oriented upward, by c and b. All four connecting paths and the whole rectangle image lie in its chosen color. Within the parameter rectangle, the bottom-then-right and left-then-top paths are homotopic relative endpoints by straight-line interpolation. Applying H gives such a homotopy in its color. Cancellation of the connecting paths at the intermediate vertices therefore gives the relation

$$ab = cd \quad \text{in } P.$$

Thus the elementary change of a grid path across this rectangle preserves its product of edge labels.

Every monotone grid path from the lower-left corner to the upper-right corner consists of rightward and upward steps. By successively exchanging an adjacent rightward-then-upward pair for an upward-then-rightward pair, any such path changes to the path that follows the left and then the top boundary. Each exchange occurs across one grid rectangle and uses the relation just proved. It follows that the bottom-then-right boundary product equals the left-then-top boundary product. The vertical boundary labels are identities, since their paths and connecting paths are constant. By Lemma 11.3.2, the remaining bottom and top products are $\Psi(\gamma_0)$ and $\Psi(\gamma_1)$. They are therefore equal.

We have proved that Ψ descends to a map $\pi_1(X, x_0) \to P$. Concatenating subdivisions for two loops shows that this map is a homomorphism: the connecting path at their shared basepoint is constant, and the resulting word is the product of their words. We already know that $\theta\Psi$ is the identity. For a loop lying entirely in U or entirely in V, choose the single-segment subdivision and constant endpoint paths. Its image under $\Psi\theta$ is its original class in P. These factor images generate P, so $\Psi\theta$ is the identity as well. Hence θ is an isomorphism, and the universal property of the amalgamated free product proves the pushout statement. □

If the overlap is simply connected, the theorem gives the ordinary free product. The overlap hypothesis has geometric content: two contractible open arcs can cover the circle, but their intersection has two components. Applying the two-set statement without path-connectedness would incorrectly give a trivial group. A groupoid version retains several basepoints and handles this situation; it is not needed for the applications below.

11.4 Graphs, Attaching Maps, and Surface Groups

The gluing theorem turns decompositions into presentations. Graphs contribute generators, and attached disks contribute relations. We begin with finite graphs so that every collapse and every open cover can be described directly.

11.4.1 Bouquets of Circles and Maximal Trees

A finite graph is a space obtained from finitely many vertices and finitely many copies of I by identifying the two endpoints of each interval with vertices. Loops and multiple edges are allowed. A tree is a connected graph with no embedded circle. A spanning tree is a subgraph containing every vertex and forming a tree.

For the bouquet $B_r = \bigvee_{j=1}^{r} \mathrm{S}^1$, choose the common vertex as basepoint. When $r \geq 2$, let U contain the first $r-1$ circles and a small open arc around the vertex in the last circle. Let V contain the last circle and small open arcs around the vertex in the others. These are open sets. They deformation retract onto B_{r-1} and S^1, while their intersection is a finite tree of arcs and contracts to the vertex. Van Kampen and induction therefore give

$$\pi_1(B_r) \cong F(\{a_1, \dots, a_r\}),$$

where a_j traverses the jth circle positively. The cases $r = 0, 1$ are a point and the circle calculation, respectively.

Theorem 11.4.1. Let Γ be a finite connected graph with v vertices and e edges. Choose a vertex x_0 and a spanning tree T. Then $\pi_1(\Gamma, x_0)$ is free of rank $e - v + 1$. A free basis consists of the loops that travel from x_0 along T to the initial vertex of an oriented edge outside T, traverse that edge, and return to x_0 along T.

Proof. A spanning tree can be constructed starting at x_0 and adjoining an edge to a new vertex whenever not all vertices have been reached. Connectedness supplies such an edge,

and adjoining a new vertex creates no cycle. The resulting tree has $v-1$ edges, one for each vertex added after x_0.

Collapse T to a point, obtaining a quotient map $q : \Gamma \to \Gamma/T$. The quotient is a bouquet with one circle for each edge outside T, hence with $e-v+1$ circles. We verify that q is a based homotopy equivalence rather than merely assuming that this collapse preserves the fundamental group.

Give each tree edge length one. Between x_0 and any point of T there is a unique simple path: two distinct simple paths would contain a cycle. Contract T by moving a point at distance d from x_0 to the point at distance $(1-t)d$ on this path. This contraction C_t is continuous, as is seen on each of the finitely many edges, and fixes x_0.

For an oriented edge a outside T, with endpoints v_a, w_a, extend this contraction over that edge at time t as follows. Traverse the tree path from $C_t(v_a)$ to v_a, then the original edge a, then the tree path from w_a to $C_t(w_a)$. Allocate parameter lengths $t/(1+2t)$, $1/(1+2t)$, and $t/(1+2t)$ to these three pieces. At $t=0$, the first and last pieces have length zero and this is the original edge map. Their images shrink to the endpoints as t tends to zero, so the family is continuous there as well. The endpoint values agree with C_t, so the edge maps and the tree contraction paste to a based homotopy on Γ.

At time one, the map is constant on T and therefore factors as sq, where $s : \Gamma/T \to \Gamma$ sends each circle to the corresponding tree-edge-tree loop. We have proved $sq \simeq \mathrm{id}_\Gamma$. On each circle of Γ/T, the map qs inserts constant initial and final segments around one traversal of that circle. A linear change of the traversal interval gives a based homotopy to the identity. These homotopies agree at the common vertex, so $qs \simeq \mathrm{id}_{\Gamma/T}$. The asserted basis and rank now follow from the bouquet computation. □

Example 11.4.2. A theta graph has two vertices joined by three edges a, b, c, all oriented from the first vertex to the second. Take the first vertex as basepoint and c as the spanning tree. The group is free on the loops $a * \bar{c}$ and $b * \bar{c}$, since $3-2+1=2$. For instance, the loop $a * \bar{b}$ represents the product $(a\bar{c})(b\bar{c})^{-1}$.

11.4.2 Attaching Disks and Higher Cells

A disk attachment identifies a boundary circle with a specified loop. The van Kampen calculation depends on an open annular neighborhood of that boundary; applying the theorem directly to the closed disk and the old space would not satisfy its hypotheses.

Proposition 11.4.3. Let X be path-connected, let $x_0 \in X$, and let $f : \mathrm{S}^1 \to X$ satisfy $f(1) = x_0$. Form the attached space by the pushout

$$\begin{array}{ccc} \mathrm{S}^1 & \xrightarrow{\ f\ } & X \\ \downarrow & & \downarrow \\ \mathrm{D}^2 & \longrightarrow & Y = X +_{\mathrm{S}^1} \mathrm{D}^2. \end{array}$$

Then inclusion induces an isomorphism

$$\pi_1(Y, x_0) \cong \pi_1(X, x_0) / \langle\langle [f] \rangle\rangle.$$

Here $[f]$ means the class of the loop $t \mapsto f(e^{2\pi i t})$.

Proof. Write $q : X + \mathrm{D}^2 \to Y$ for the quotient map. Let U be the image of X together with the disk points satisfying $\|z\| > 1/3$, and let V be the image of the disk points satisfying $\|z\| < 2/3$. Their inverse images under q are open saturated subsets, so U, V are open. They cover Y.

Radial motion toward the disk boundary gives a deformation retraction of U onto X, and V is an open disk. Their intersection is the annulus $1/3 < \|z\| < 2/3$.

Choose $b = q(1/2)$ in the intersection. Moving radially from b to $q(1) = x_0$ supplies the basepoint comparison. Under the deformation retraction of U onto X, the positive circle generating the annulus group becomes the loop f. Its image in $\pi_1(V, b)$ is trivial. Van Kampen therefore identifies $\pi_1(Y, b)$ with the quotient of $\pi_1(X, x_0)$ by the normal closure of $[f]$. Transporting the basepoint to x_0 gives the stated isomorphism, with the homomorphism induced by inclusion. □

For an attaching circle whose chosen point maps elsewhere, choose a path from x_0 to that image point. The relation is represented by the attaching loop conjugated along this path. Changing the path changes that relation by conjugacy, which leaves its normal closure unchanged. Repeated application handles any finite collection of attached disks.

Lecture 9 already established the simple connectedness of higher spheres. The following alternate proof by van Kampen also supplies the local calculation for higher-dimensional cell attachments.

Proposition 11.4.4. For every $n \geq 2$, the sphere S^n is simply connected. Attaching a cell of dimension $n \geq 3$ to a path-connected space does not change its fundamental group.

Proof. Let N and S be the two poles of S^n, and let U and V be their respective complements. Stereographic projection identifies each with $\mathbb{R}^n$, so both are contractible. Their intersection is homeomorphic to $\mathrm{S}^{n-1} \times \mathbb{R}$ and is path-connected for $n \geq 2$. Van Kampen is the pushout of two trivial groups, and hence gives $\pi_1(\mathrm{S}^n) = 1$.

For an attached n-cell, use the same two open sets as in Proposition 11.4.3, with D^n in place of D^2. Their intersection deformation retracts onto S^{n-1}. When $n \geq 3$, that intersection and the inner open ball have trivial fundamental group. The pushout therefore identifies the fundamental group of the attached space with that of the original space. □

11.4.3 Tori, Projective Planes, and Closed Surfaces

Polygonal descriptions of surfaces now become presentations: oriented edges give generators and the boundary word of the polygon gives a relation. The word must be read in order, including the orientation of each side.

Example 11.4.5. The torus is obtained from a bouquet with generators a, b by attaching a disk along $aba^{-1}b^{-1}$. Thus

$$\pi_1(\mathrm{S}^1 \times \mathrm{S}^1) \cong \langle a, b \mid aba^{-1}b^{-1} \rangle \cong \mathbb{Z}^2.$$

For the last isomorphism, the relation says that a and b commute, so every word is $a^m b^n$. Sending a to $(1, 0)$ and b to $(0, 1)$ defines a homomorphism to $\mathbb{Z}^2$, and $(m, n) \mapsto a^m b^n$ is its inverse. This also proves uniqueness of this expression.

Example 11.4.6. The real projective plane is the quotient of a disk by identifying antipodal boundary points. Its boundary quotient is $\mathbb{RP}^1 \cong \mathrm{S}^1$, and the attaching map traverses this quotient circle twice. Consequently

$$\pi_1(\mathbb{RP}^2) \cong \langle a \mid a^2 \rangle \cong \mathbb{Z}/2\mathbb{Z}.$$

Indeed, under the identification $\mathbb{RP}^1 \to \mathrm{S}^1$ given by $[e^{i\theta}] \mapsto e^{2i\theta}$, the attaching map is $z \mapsto z^2$ and has winding number 2.

For $g \geq 1$, define the orientable surface Σ_g by identifying the paired sides of a $4g$-gon according to the boundary word

$$a_1 b_1 a_1^{-1} b_1^{-1} \cdots a_g b_g a_g^{-1} b_g^{-1}.$$

Its vertices become one vertex, its sides form a bouquet of $2g$ circles, and its interior is one attached disk. If $[a,b]$ denotes $aba^{-1}b^{-1}$, the attachment calculation gives

$$\pi_1(\Sigma_g) \cong \langle a_1, b_1, \dots, a_g, b_g \mid [a_1,b_1]\cdots[a_g,b_g]\rangle.$$

For genus two, this is one relation involving four generators, rather than two separate commutativity relations. In fact, for every $g \geq 2$ there is a surjection onto the free group on two generators: send a_1, a_2 to those generators and all remaining generators to the identity. The displayed relation maps to the identity. Since a quotient of an abelian group is abelian and this free group is not, these surface groups are nonabelian.

The nonorientable surface N_h, defined by the polygon with boundary word $a_1^2 \cdots a_h^2$, similarly has

$$\pi_1(N_h) \cong \langle a_1, \dots, a_h \mid a_1^2 \cdots a_h^2 \rangle.$$

The case $h = 1$ recovers the projective plane. A usual square description of the Klein bottle instead gives $\langle a, b \mid aba^{-1}b\rangle$, or equivalently $aba^{-1} = b^{-1}$. This relation records that transport once around the a-direction reverses the b-direction.

11.5 Polynomial Roots and Antipodal Points

The winding number detects an obstruction to extending a circle map over a disk. In the fundamental theorem of algebra, a polynomial without roots would provide such an extension for a map of positive degree. In the Borsuk–Ulam theorem, an assumed separation of antipodal images would provide an extension of an odd circle map, whose degree must be odd.

11.5.1 The Fundamental Theorem of Algebra

The only estimates needed are on a sufficiently large circle, where the leading term of a polynomial dominates the sum of the lower terms. This estimate gives an explicit homotopy through nonzero complex numbers.

Theorem 11.5.1. Every nonconstant polynomial with complex coefficients has a complex root. Consequently every polynomial of degree $n \geq 1$ is a product of its leading coefficient and n linear factors, with roots repeated according to multiplicity.

Proof. Write

$$P(z) = a_n z^n + \sum_{k=0}^{n-1} a_k z^k, \qquad n \geq 1, \quad a_n \neq 0.$$

Suppose that P has no root. Choose $R \geq 1$ so large that

$$\sum_{k=0}^{n-1} |a_k| R^k < |a_n| R^n.$$

For example, any $R > \max\{1, \sum_{k<n} |a_k|/|a_n|\}$ has this property. The map $z \mapsto P(Rz)/|P(Rz)|$ is then defined on the entire closed unit disk. Its boundary restriction $f_R : \mathrm{S}^1 \to \mathrm{S}^1$ has degree zero by Proposition 11.1.5.

For $z \in \mathrm{S}^1$ and $t \in I$, set

$$P_t(z) = a_n R^n z^n + (1-t)\sum_{k=0}^{n-1} a_k R^k z^k.$$

The reverse triangle inequality gives

$$|P_t(z)| \geq |a_n| R^n - \sum_{k=0}^{n-1} |a_k| R^k > 0.$$

Thus $P_t(z)/|P_t(z)|$ is a homotopy from f_R to $z \mapsto (a_n/|a_n|)z^n$. Multiplication by the fixed unit complex number $a_n/|a_n|$ is homotopic to the identity rotation: choose an angle for that one number and rotate through that angle. Hence $\deg(f_R) = n$, contradicting $\deg(f_R) = 0$.

A root r gives a factorization $P(z) = (z - r)Q(z)$ by polynomial division. Apply the root-existence statement repeatedly to Q and the subsequent factors of decreasing degree. After n steps, the remaining factor is the leading coefficient. □

11.5.2 Odd Circle Maps Have Odd Degree

A map between spheres is odd if it commutes with the antipodal maps. On the circle, oddness determines the change in a lifted angle over a half-turn, which determines the parity of its degree.

Proposition 11.5.2. For every continuous map $g : \mathrm{S}^1 \to \mathrm{S}^1$ satisfying $g(-z) = -g(z)$, the integer $\deg(g)$ is odd.

Proof. As in the proof of Proposition 11.1.5, choose a lift $G : \mathbb{R} \to \mathbb{R}$ of $t \mapsto g(e^{2\pi i t})$ satisfying

$$G(t+1) - G(t) = \deg(g).$$

Oddness gives

$$e^{2\pi i G(t+1/2)} = -e^{2\pi i G(t)},$$

so the continuous function $G(t + 1/2) - G(t) - 1/2$ is integer-valued. Since $\mathbb{R}$ is connected, it is a constant $m \in \mathbb{Z}$. Applying this half-turn identity twice yields

$$\deg(g) = G(t+1) - G(t) = 2m + 1.$$

Thus the degree is odd. □

In particular, an odd circle map cannot extend over a disk. The obstruction concerns all continuous odd maps, with no differentiability or injectivity assumption.

11.5.3 Borsuk–Ulam in Dimensions One and Two

For a map to Euclidean space, the difference between the values at antipodal points is automatically odd. If that difference never vanishes, normalization produces a map to a smaller sphere. In dimension two, its equatorial restriction contradicts the preceding degree calculation.

Theorem 11.5.3. For $n = 1$ and $n = 2$, every continuous map $f : \mathrm{S}^n \to \mathbb{R}^n$ has a point x such that $f(x) = f(-x)$.

Proof. For $n = 1$, consider the continuous real-valued function $h(x) = f(x) - f(-x)$. If h vanishes anywhere, we are done. Otherwise choose any x and a path in the circle from x to $-x$. The values at the two ends are nonzero and opposite, because $h(-x) = -h(x)$. The intermediate value theorem gives a zero on that path, a contradiction.

For $n = 2$, suppose that $f(x) \neq f(-x)$ for every $x \in \mathrm{S}^2$. The formula

$$g(x) = \frac{f(x) - f(-x)}{\|f(x) - f(-x)\|}$$

defines a continuous map $g : \mathrm{S}^2 \to \mathrm{S}^1$ satisfying $g(-x) = -g(x)$. Identify the equator with S^1 using $(x_1, x_2, 0) \leftrightarrow x_1 + i x_2$. The equatorial restriction is odd, so Proposition 11.5.2 gives it odd degree.

The same restriction extends over the closed northern hemisphere through g. That hemisphere is a disk: the homeomorphism from D^2 sends (u, v) to $(u, v, \sqrt{1-u^2-v^2})$ and identifies its boundary with the equator. Proposition 11.1.5 therefore gives degree zero, a contradiction.□

The proof in dimension two uses only circle lifting and an extension over a hemisphere. In higher dimensions, the analogous argument requires the higher-dimensional statement that an odd self-map of a sphere has odd degree. The low-dimensional theorem proved here is sufficient for the syllabus at this point.

Further reading. J. P. May, *A Concise Course in Algebraic Topology*, Chapter 1, Sections 5 and 7, pp. 8–11, treats circle lifting and polynomial roots. His Chapter 2, Sections 7–8, pp. 17–20, develops van Kampen as a categorical colimit, and Chapter 4, Sections 2–4, pp. 35–37, treats trees and graph groups. Further comparisons are Allen Hatcher, *Algebraic Topology*, Theorems 1.7, 1.8, and 1.10, pp. 29–33, for circle lifting and the two final applications; Theorem 1.20, pp. 43–46, for van Kampen; and Proposition 1.26, pp. 50–51, for cell attachments.

Covering Spaces and Subgroups of Free Groups

Algebraic Topology I | Fall 2026 | October 19, 2026

12.1 Lifting Paths, Homotopies, and Maps

The exponential map used in Lecture 11 describes the circle by a simply connected space together with a group of translations. Covering space theory extends this description to spaces whose sufficiently small neighborhoods have no essential loops. Its classification theorem identifies connected covering spaces with subgroups of the fundamental group. The lifting results themselves require fewer hypotheses, and we prove them first.

12.1.1 Covering Maps and Unique Lifting

Write $I = [0,1]$. For composable paths α and β, the concatenation $\alpha * \beta$ traverses α first and β second. Write $\overline{\alpha}(t) = \alpha(1-t)$ for the reversed path. All homotopies of paths preserve their endpoints unless otherwise stated.

Definition 12.1.1. A *covering map* is a continuous surjection $p : Y \to X$ such that every point of X has an open neighborhood U for which $p^{-1}(U)$ is a disjoint union of open subsets of Y, each mapped homeomorphically onto U by p. Such a neighborhood is *evenly covered*, and these open subsets are its *sheets*. A *lift* of $f : T \to X$ is a continuous map $\tilde{f} : T \to Y$ satisfying $p\tilde{f} = f$.

Each fiber $p^{-1}(x)$ is discrete. Every covering map is a local homeomorphism, but the simultaneous decomposition over one neighborhood is part of the definition of a covering. For example, the restriction $\exp : (0,2) \to \mathrm{S}^1$, defined by $t \mapsto e^{2\pi i t}$, is a surjective local homeomorphism that is not a covering: the fiber over 1 has one point, whereas nearby fibers have two.

Lemma 12.1.2. Let $p : Y \to X$ be a covering map, let T be connected, and let $f : T \to X$ be continuous. If two lifts of f agree at one point, then they agree everywhere.

Proof. Let A be the set of points where the lifts agree. Near any $u \in T$, choose an evenly covered neighborhood of $f(u)$ and a neighborhood of u on which each lift lies in one sheet. If the sheets are equal, the two lifts agree throughout this neighborhood, because they are obtained by composing f with the same inverse of p. If the sheets differ, the lifts disagree throughout this neighborhood. Thus A and its complement are open. Connectedness and the assumed agreement imply $A = T$. □

Proposition 12.1.3. For every path $\alpha : I \to X$ and every $y \in p^{-1}(\alpha(0))$, there is a unique lift $\tilde{\alpha} : I \to Y$ satisfying $\tilde{\alpha}(0) = y$.

Proof. Pull back evenly covered neighborhoods along α. The Lebesgue number lemma supplies a subdivision $0 = t_0 < \cdots < t_m = 1$ such that the image of each closed interval $[t_{j-1}, t_j]$ lies in one evenly covered neighborhood U_j. On the first interval, compose α with the inverse of p on the sheet containing y. Its endpoint selects a sheet over U_2, and this determines the lift on the next interval. Continue through the finite subdivision. The resulting maps agree at their endpoints, so the pasting lemma proves continuity. Uniqueness follows from Lemma 12.1.2. □

The same local inverses lift a continuously varying family of paths. The parameter space need not be locally connected or compact.

Theorem 12.1.4. Let $p : Y \to X$ be a covering map, let T be any space, and let $F : T \times I \to X$ be a homotopy. Given a lift $g : T \to Y$ of $F|_{T\times\{0\}}$, there is a unique continuous map $\widetilde{F} : T \times I \to Y$ with $p\widetilde{F} = F$ and $\widetilde{F}(u, 0) = g(u)$.

Proof. For each u, lift the path $t \mapsto F(u, t)$ starting at $g(u)$. This gives a uniquely determined function $\widetilde{F}$. We prove that this function is continuous near $\{u_0\} \times I$ for an arbitrary $u_0 \in T$.

Choose a finite subdivision such that $F(\{u_0\} \times [t_{j-1}, t_j])$ lies in an evenly covered open set U_j. By compactness of each closed interval, there is a neighborhood N of u_0 such that

$$F(N \times [t_{j-1}, t_j]) \subseteq U_j \qquad (1 \le j \le m).$$

Indeed, the open set $F^{-1}(U_j)$ contains the compact set $\{u_0\} \times [t_{j-1}, t_j]$; finitely many product neighborhoods give the required neighborhood of u_0. Intersect these neighborhoods over the finitely many values of j.

Shrink N so that $g(N)$ lies in the sheet over U_1 containing $g(u_0)$. On $N \times [t_0, t_1]$, the desired lift is the continuous map obtained from the inverse on that sheet. Its endpoint is a continuous function of u. Shrink N again so that these endpoints lie in the sheet over U_2 containing $\widetilde{F}(u_0, t_1)$. This gives a continuous lift on the next interval. After finitely many such restrictions, the local formulas agree at every junction and define a continuous map on $N \times I$. Path-lifting uniqueness identifies this map with $\widetilde{F}$. Such neighborhoods exist for every u_0, proving continuity everywhere. Uniqueness also follows by restricting to each vertical path. □

In particular, every covering map is a Hurewicz fibration in the terminology of Lecture 10. If a homotopy fixes a point of T, its lifted track lies in a discrete fiber and is constant. Thus a homotopy of paths lifts to a homotopy with the same lifted endpoints.

12.1.2 The Fundamental Group and the Lifting Criterion

A choice of $y_0 \in Y$ determines the basepoint $x_0 = p(y_0)$ and a subgroup of $\pi_1(X, x_0)$. This subgroup records which loops return to the chosen sheet.

Proposition 12.1.5. For every covering map $p : (Y, y_0) \to (X, x_0)$, the homomorphism

$$p_* : \pi_1(Y, y_0) \longrightarrow \pi_1(X, x_0)$$

is injective. Its image consists of the classes represented by loops whose lifts starting at y_0 end at y_0.

Proof. If a loop λ in Y projects to a nullhomotopic based loop, lift a based nullhomotopy of $p\lambda$ with initial lift λ. The endpoints remain y_0. The final path is the lift of a constant path, so it is constant. This is a based nullhomotopy of λ, proving injectivity.

A loop that lifts to a loop represents an element of the image of p_*. Conversely, if α represents an element of that image, it is based homotopic to the projection of a loop in Y at y_0. Lifting this homotopy shows that its lift also ends at y_0. □

Theorem 12.1.6. Let $p : (Y, y_0) \to (X, x_0)$ be a covering map. Let T be path connected and locally path connected, with basepoint u_0, and let $f : (T, u_0) \to (X, x_0)$ be continuous. A based lift $\tilde{f} : (T, u_0) \to (Y, y_0)$ exists if and only if

$$f_*(\pi_1(T, u_0)) \subseteq p_*(\pi_1(Y, y_0)).$$

When it exists, the based lift is unique.

Proof. A based factorization $f = p\tilde{f}$ implies the subgroup inclusion by functoriality. Conversely, assume the inclusion. For $u \in T$, choose a path γ from u_0 to u and define $\tilde{f}(u)$ to be the endpoint of the lift of $f\gamma$ starting at y_0.

If γ' is another choice, the loop $f\gamma * \overline{f\gamma'}$ represents an element of $p_*(\pi_1(Y, y_0))$ and therefore lifts to a loop at y_0. Its second part, reversed, is the lift of $f\gamma'$ starting at y_0. Hence the lifts of $f\gamma$ and $f\gamma'$ have the same endpoint. The definition is independent of the path.

To prove continuity at u, choose an evenly covered neighborhood U of $f(u)$ and a path-connected open neighborhood V of u contained in $f^{-1}(U)$. For every $z \in V$, a path from u to z inside V can be appended to γ. The lifted appended path stays in the sheet W over U containing $\tilde{f}(u)$. Consequently

$$\tilde{f}|_V = (p|_W)^{-1} \circ f|_V.$$

This proves continuity. The constant path gives $\tilde{f}(u_0) = y_0$, and Lemma 12.1.2 proves uniqueness. □

The local path-connectedness hypothesis enters in choosing V. No local hypothesis on X was used in this theorem. In particular, every based map from a simply connected, locally path-connected space has a unique based lift.

12.2 Constructing and Classifying Connected Covers

The lifting criterion compares covering spaces that already exist. To realize every subgroup, we must construct a topology on suitable classes of paths. The required local condition says that small loops become trivial in the whole space, rather than necessarily in the neighborhood itself.

12.2.1 Local Hypotheses and the Path-Class Construction

The local hypotheses distinguish paths inside small neighborhoods from nullhomotopies that may leave those neighborhoods.

Definition 12.2.1. A space is *locally path connected* if its path-connected open subsets form a basis. It is *semilocally simply connected* if each point x has an open neighborhood U such that every loop in U based at x is nullhomotopic in X.

From now through Section 12.3, assume that X is connected, locally path connected, and semilocally simply connected. Connectedness and local path-connectedness imply path-connectedness: path components are open, and their complements are unions of open path components.

Call a path-connected open set $U \subseteq X$ *admissible* if every loop in U is nullhomotopic in X. Admissible sets form a basis. To see this, intersect a neighborhood from the semilocal condition with any prescribed open neighborhood and choose a path-connected open neighborhood inside the intersection. Changing the basepoint along a path in this smaller neighborhood shows that its loops at every basepoint become trivial in X.

Fix $x_0 \in X$, put $G = \pi_1(X, x_0)$, and let $H \leq G$ be a subgroup. For paths α, β starting at x_0, define

$$\alpha \sim_H \beta \quad \Longleftrightarrow \quad \alpha(1) = \beta(1) \text{ and } [\alpha * \overline{\beta}] \in H.$$

This is an equivalence relation. Reflexivity and symmetry use the identity and inverses in H; transitivity follows by cancelling the intervening path followed by its reverse. Let Y_H be the set of equivalence classes, write $[\alpha]_H$ for a class, and define $p_H : Y_H \to X$ by $p_H([\alpha]_H) = \alpha(1)$. Let y_H be the class of the constant path at x_0.

For an admissible neighborhood U of $\alpha(1)$, define

$$B([\alpha]_H, U) = \{[\alpha * \delta]_H \mid \delta \text{ is a path in } U \text{ starting at } \alpha(1)\}.$$

This set depends only on the class of α: appending the same path preserves the defining equivalence relation. We give Y_H the topology having these sets as a basis. This is a specified basis topology; no quotient topology on a function space is being assumed.

Proposition 12.2.2. The sets $B([\alpha]_H, U)$ form a basis, and the endpoint map $p_H : Y_H \to X$ is a path-connected covering space. Its associated subgroup is

$$(p_H)_*(\pi_1(Y_H, y_H)) = H.$$

Proof. First consider one admissible set U. Any two paths in U with the same endpoints differ by a loop nullhomotopic in X. After appending them to α, the resulting paths are equivalent. It follows that the restriction

$$p_H : B([\alpha]_H, U) \longrightarrow U$$

is a bijection. Surjectivity here uses path-connectedness of U.

If $[\beta]_H$ belongs to $B([\alpha]_H, U)$ and W is an admissible neighborhood of $\beta(1)$ contained in U, then

$$B([\beta]_H, W) \subseteq B([\alpha]_H, U).$$

Indeed, represent $[\beta]_H$ by $\alpha * \delta$ with δ in U, and append a path in W. Given a point in the intersection of two proposed basis sets, choose W inside the intersection of their endpoint neighborhoods. The displayed inclusion proves the basis intersection axiom.

The endpoint map is continuous: for an open set $O \subseteq X$, every class ending in O has a basic neighborhood whose endpoint neighborhood lies in O. Each displayed bijection is a homeomorphism. To check continuity of its inverse, take a basic neighborhood of a point in the sheet and shrink its endpoint neighborhood to an admissible $W \subseteq U$. The inverse then sends W into that basic neighborhood.

Fix an admissible U and a point $u \in U$. The sets $B([\alpha]_H, U)$ with $\alpha(1) = u$ cover $p_H^{-1}(U)$. In fact, for a path β ending at $x \in U$, choose a path δ in U from u to x and use the path $\beta * \overline{\delta}$ ending at u. Two of these sheets that meet are equal: append the reverse of a path in U from u to the common endpoint to see that their classes over u agree. Thus U is evenly covered. Path-connectedness of X makes p_H surjective.

For a path α starting at x_0, set $\alpha_t(s) = \alpha(ts)$. The function $t \mapsto [\alpha_t]_H$ is a continuous path from y_H to $[\alpha]_H$. Near any t_0, choose a small interval of parameters whose image under α lies in one admissible neighborhood. The classes $[\alpha_t]_H$ are then given by the local inverse on the sheet containing $[\alpha_{t_0}]_H$; reversing the short segment handles $t < t_0$. This proves continuity and path-connectedness of Y_H.

The same path is the unique lift of α from y_H. For a loop α at x_0, it closes if and only if $[\alpha]_H = y_H$, which means $[\alpha] \in H$. Proposition 12.1.5 now identifies the image subgroup with H. □

12.2.2 Based Classification, Conjugacy, and Universal Covers

An isomorphism of covers is a homeomorphism over X. For based covers it must also preserve the chosen points above x_0. The classification is more precise with those points retained.

Theorem 12.2.3. For a connected, locally path-connected, semilocally simply connected space X, the assignment

$$(p : (Y, y_0) \to (X, x_0)) \longmapsto p_*(\pi_1(Y, y_0))$$

gives a bijection from based isomorphism classes of connected covering spaces to subgroups of $G = \pi_1(X, x_0)$. If H and K are the associated subgroups, there is a based map of covers $Y_H \to Y_K$ if and only if $H \subseteq K$, and this map is unique.

Proof. A covering of a locally path-connected space is locally path connected, because its sheets are homeomorphic to open subsets of the base. Thus a connected covering space is path connected. The lifting criterion applied to the projection of one based covering space into another gives the stated existence and uniqueness of maps over X.

If two covers have the same subgroup, there are based maps over X in both directions. Each composite and the appropriate identity map are lifts of the same projection and agree at the basepoint. Unique lifting makes the composites identities, so the maps are inverse homeomorphisms. Proposition 12.2.2 realizes every subgroup. These statements prove the bijection and the assertion about maps. □

The classification also counts sheets and explains the conjugacy ambiguity. For a covering $p : Y \to X$, lift a loop representing $g \in G$ from a point $y \in p^{-1}(x_0)$ and denote its endpoint by $y \cdot g$. Homotopy lifting makes this well defined. Lifting concatenations gives a right action:

$$(y \cdot g) \cdot h = y \cdot (gh).$$

For a connected cover, the action is transitive, since a path between two points of the fiber projects to a loop. Its stabilizer at y_0 is $H = p_*(\pi_1(Y, y_0))$ by Proposition 12.1.5.

Corollary 12.2.4. For a connected based cover with subgroup $H \leq G$, the map

$$\{Hg \mid g \in G\} \longrightarrow p^{-1}(x_0), \qquad Hg \longmapsto y_0 \cdot g$$

is a bijection of right G-sets. Consequently the number of sheets is the index $[G : H]$, allowing an infinite cardinal. Unbased isomorphism classes of connected covers correspond to conjugacy classes of subgroups of G.

Proof. Transitivity proves surjectivity. The endpoints $y_0 \cdot g$ and $y_0 \cdot h$ agree if and only if the loop with class gh^{-1} lifts to a loop at y_0. This is equivalent to $gh^{-1} \in H$, and hence to $Hg = Hh$. The formula respects the right action.

If the chosen point is changed from y_0 to $y_0 \cdot g$, its stabilizer is $g^{-1}Hg$. Thus an unbased isomorphism identifies conjugate subgroups. Conversely, if the subgroup of a second based cover is $g^{-1}Hg$, change the chosen point in the first cover to $y_0 \cdot g$ and apply the based classification. This gives an unbased isomorphism. □

Here each coset Hg consists of the products hg with $h \in H$. The right action on these cosets is induced by multiplication on the right in G.

Corollary 12.2.5. The cover $p_{\{1\}} : Y_{\{1\}} \to X$ corresponding to the trivial subgroup is simply connected. It admits a unique based map over X to every connected based covering space, and this map is itself a covering map. Every simply connected connected cover is based isomorphic to $Y_{\{1\}}$. Such a cover is called a *universal cover*.

Proof. The homomorphism $(p_{\{1\}})_*$ is injective and has trivial image, so $\pi_1(Y_{\{1\}}, y_{\{1\}})$ is trivial. The classification gives the unique map to any based connected cover because the trivial subgroup is contained in every subgroup. We verify below, in Proposition 12.3.2, that every such based map between connected covers is a covering map. Conversely, a simply connected cover has trivial image subgroup, so the uniqueness statement in the classification identifies it with $Y_{\{1\}}$. □

The semilocal hypothesis is also necessary for the existence of a simply connected covering space over a locally path-connected base. Indeed, choose a path-connected evenly covered neighborhood of a base point. Every loop there lifts to a loop in one sheet; if the total space is simply connected, its projection is nullhomotopic in X.

For example, the Hawaiian earring is the union in $\mathbb{R}^2$ of the circles of radius $1/n$ centered at $(1/n, 0)$, for $n \geq 1$. Every neighborhood of their common point contains whole circles of

sufficiently small radius. Each such circle is essential: the map that fixes it and collapses every other circle to the common point is a continuous retraction, so it injects on fundamental groups. Therefore the Hawaiian earring is not semilocally simply connected and has no simply connected covering space. This distinguishes it from the CW wedge of countably many circles used below.

12.3 Deck Transformations and Intermediate Coverings

The automorphisms of a covering are constrained by path lifting: the image of one point determines the entire automorphism. The monodromy action on the fiber translates this constraint into a normalizer computation. It also identifies the quotient constructions that occur in the Galois correspondence.

12.3.1 The Deck Group and Normal Subgroups

A *deck transformation* of $p : Y \to X$ is a homeomorphism $d : Y \to Y$ satisfying $pd = p$. Write $\mathrm{Deck}(p)$ for their group, with multiplication given by composition. A connected cover is *regular* if its deck group acts transitively on the fiber over x_0. Path lifting then gives transitivity on every fiber.

Let $H \leq G$. Its *normalizer* is the subgroup

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

The subgroup H is normal in $N_G(H)$, so $N_G(H)/H$ is a group.

Theorem 12.3.1. For the connected based cover corresponding to $H \leq G$, there is a group isomorphism

$$\mathrm{Deck}(p_H) \cong N_G(H)/H.$$

This cover is regular if and only if H is normal in G. In that case its deck group is G/H. In particular, the universal cover has deck group G.

Proof. A deck transformation commutes with the monodromy action: applying it to a lifted path gives the lift of the same path starting at the transformed point. Its permutation of the fiber is therefore a G-equivariant bijection of $\{Hg \mid g \in G\}$. If it sends H to Hn, equivariance requires

$$Hg \longmapsto Hng.$$

The stabilizer of H is H, whereas the stabilizer of Hn is $n^{-1}Hn$. An equivariant bijection preserves stabilizers, so n belongs to $N_G(H)$. Unique lifting shows that the deck transformation is determined by this value at the chosen point.

Conversely, represent $n \in N_G(H)$ by a loop ν at x_0 and define

$$D_n : Y_H \longrightarrow Y_H, \qquad D_n([\alpha]_H) = [\nu * \alpha]_H.$$

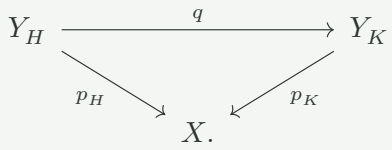
If $[\alpha * \overline{\beta}] \in H$, the difference between the new paths represents $n[\alpha * \overline{\beta}]n^{-1} \in H$. Thus the formula is well defined, and changing the representative of n does not change it. It carries every basic sheet $B([\alpha]_H, U)$ homeomorphically onto $B([\nu * \alpha]_H, U)$ and has inverse $D_{n^{-1}}$. Therefore it is a deck transformation.

The formulas give $D_n \circ D_m = D_{nm}$ and $D_n = \mathrm{id}$ if and only if $n \in H$. All deck transformations have already been accounted for, so they give the asserted isomorphism. The possible images of the chosen point are the cosets Hn with $n \in N_G(H)$. These constitute the entire fiber if and only if $N_G(H) = G$, which is the normality condition. □

No nonidentity deck transformation of a connected cover fixes a point, by unique lifting. Thus a regular finite-sheeted cover has as many deck transformations as sheets. A nonregular cover can have a much smaller deck group; an example appears in Section 12.4.2.

12.3.2 The Galois Correspondence and Circle Covers

Fix a based connected cover $p_H : Y_H \to X$. A based intermediate covering is a factorization through a connected based cover, with maps preserving all projections and basepoints:



Isomorphisms of factorizations are based isomorphisms of the middle covers commuting with both maps.

Proposition 12.3.2. The based intermediate coverings of p_H are classified by the subgroups K satisfying $H \le K \le G$. The map $q : Y_H \to Y_K$ is a covering map of degree $[K : H]$. It is regular if and only if H is normal in K, and its deck group is $N_K(H)/H$.

Proof. For $H \le K$, the unique based map is given in the path-class models by

$$q([\alpha]_H) = [\alpha]_K.$$

It is well defined because the equivalence relation for H implies that for K. Over an admissible open set $U \subseteq X$, the map carries each H-sheet homeomorphically onto the corresponding K-sheet: both sheets are identified with U by their projections. The inverse image of any K-sheet is the disjoint union of the H-sheets mapping to it. Every point of the K-sheet has a representative path and therefore a preimage. This proves that q is a covering map.

Conversely, an intermediate factorization gives

$$(p_H)_*(\pi_1(Y_H, y_H)) \subseteq (p_K)_*(\pi_1(Y_K, y_K)) \subseteq G.$$

The classification and uniqueness of based lifts identify the factorization with the one just constructed. The fiber of q over y_K consists of the cosets Hk for $k \in K$, so its degree is $[K : H]$.

The injection $(p_K)_*$ identifies $\pi_1(Y_K, y_K)$ with K and identifies the subgroup defined by q with H. The space Y_K is locally path connected by its sheet neighborhoods. A loop in a sheet over an admissible open set projects to a nullhomotopic loop in X; injectivity of $(p_K)_*$ makes the original loop nullhomotopic in Y_K. Thus Y_K is also semilocally simply connected. Applying Theorem 12.3.1 to q proves the regularity criterion and the deck-group formula. □

The direction of the correspondence is important: $K_1 \subseteq K_2$ gives a map $Y_{K_1} \to Y_{K_2}$ over X. A smaller subgroup gives a covering space lying above a larger subgroup. If the relevant indices are finite, counting fibers gives

$$[G : H] = [G : K]\,[K : H].$$

Corollary 12.3.3. Suppose H is normal in G, and identify $D = \mathrm{Deck}(p_H)$ with G/H. For every subgroup $L \le D$, let $K \le G$ be its inverse image. The orbit space Y_H/L is naturally isomorphic over X to Y_K. This gives the Galois correspondence between subgroups of D and based intermediate covers of p_H. The cover $Y_K \to X$ is regular if and only if L is normal in D, and then its deck group is D/L.

Proof. The group $L = K/H$ acts through the deck transformations D_k. Two classes $[\alpha]_H$ and $[\beta]_H$ have the same image in Y_K if and only if they have the same endpoint and $[\alpha * \overline{\beta}] \in K$. In

that case, put $k = [\alpha * \overline{\beta}]$; cancelling $\overline{\beta} * \beta$ shows that $D_k([\beta]_H) = [\alpha]_H$. Thus the fibers of $Y_H \to Y_K$ are the L-orbits.

The covering map $Y_H \to Y_K$ is open and surjective, hence is a quotient map. Its fiber description therefore identifies the quotient topology on Y_H/L with the topology of Y_K. Finally, K is normal in G if and only if $K/H = L$ is normal in $G/H = D$. The normalizer theorem and the group isomorphism $(G/H)/(K/H) \cong G/K$ prove the final assertion. □

In particular, every connected cover of X is an orbit space of the universal cover by a subgroup of its deck group. The orbit space is formed using that specified deck action, so its projection is a covering map by the preceding proof.

Example 12.3.4. For each integer $n \geq 1$, the map $p_n : \mathrm{S}^1 \to \mathrm{S}^1$ defined by $p_n(z) = z^n$ is an n-sheeted covering. Local inverse branches on short arcs give its sheets. The winding-number calculation in Lecture 11 shows that $(p_n)_*$ is multiplication by n on $\mathbb{Z}$, so the associated subgroup is $n\mathbb{Z}$. Its deck group consists of the rotations $z \mapsto \zeta z$ with $\zeta^n = 1$, and is isomorphic to $\mathbb{Z}/n\mathbb{Z}$.

The universal cover is $p : \mathbb{R} \to \mathrm{S}^1$, given by $p(t) = e^{2\pi i t}$, with deck transformations $t \mapsto t + k$ for $k \in \mathbb{Z}$. Since every subgroup of $\mathbb{Z}$ is $n\mathbb{Z}$ for a unique $n \geq 0$, these are all connected covers: $n = 0$ gives the universal cover, and $n \geq 1$ gives p_n.

For positive integers m, n, the inclusion $n\mathbb{Z} \subseteq m\mathbb{Z}$ holds if and only if m divides n. Writing $n = mk$, the corresponding intermediate map is $q : \mathrm{S}^1 \to \mathrm{S}^1$, given by $q(z) = z^k$, and $p_m q = p_n$. Its degree is k, illustrating the direction of the subgroup correspondence.

12.4 Graphs and Subgroups of Free Groups

Covering spaces turn subgroup questions into questions about spaces. For free groups, the spaces are graphs, and a maximal tree gives an explicit basis for the fundamental group. We include the infinite-graph argument because a subgroup of a finitely generated free group need not be finitely generated.

12.4.1 Maximal Trees and Infinite Graphs

A *graph* here is a one-dimensional CW complex; loops, multiple edges, and infinitely many edges are allowed. Each edge is a copy of an interval with its endpoints attached to vertices, and the topology is the CW weak topology. A *tree* is a connected graph containing no finite edge cycle. A *spanning tree* of a connected graph is a subtree containing every vertex.

Every connected graph has a spanning tree. Start with all vertices and use Zorn's lemma on the spanning subgraphs containing no finite cycle, ordered by inclusion. The union of a chain has no cycle, because a finite cycle in the union would lie in one member of the chain. A maximal such subgraph is connected: if two of its components could be joined by an edge, that edge could be added without creating a cycle. The ambient graph is connected in the edge-path sense, since its edge-path components are open and closed by the weak topology. Therefore an edge joining distinct components would exist unless the maximal subgraph were connected. Between two vertices of a tree there is a unique reduced finite edge path; two distinct such paths would give a finite cycle.

We first justify the finiteness reduction needed for loops and homotopies.

Lemma 12.4.1. Every compact subset of a graph is contained in a finite subgraph.

Proof. Suppose a compact set met infinitely many open cells, and choose points $z_1, z_2, \ldots$ in distinct such cells. Let A be their set. Every subset of A is closed in the graph: its intersection with each closed edge consists of at most the two endpoints and one point in the interior, so

its inverse image under every characteristic map is closed. The weak topology gives the assertion.

In particular, A is a closed subset of the compact set, and is therefore compact. But every subset of A is closed, so A is an infinite discrete space, which is not compact. This contradiction shows that the compact set meets only finitely many open cells. The union of their closures is a finite subgraph containing it. □

Choose a base vertex v_0 in a connected graph Γ and a spanning tree T. For each vertex v, let τ_v be the unique reduced edge path in T from v_0 to v. Choose an orientation of each edge outside T. If such an edge e runs from u to v, define the based loop

$$\lambda_e = \tau_u * e * \overline{\tau_v}.$$

Theorem 12.4.2. The fundamental group $\pi_1(\Gamma, v_0)$ is free on the classes $[\lambda_e]$, with one generator for each oriented edge outside the spanning tree T.

Proof. For a finite connected graph and a spanning tree, collapse the tree. This is a homotopy equivalence onto a finite wedge of circles, one for each remaining edge, as proved in Lecture 11. The van Kampen theorem identifies the fundamental group with the free group on these circles. Under this identification, the loops λ_e give the stated basis.

Now allow Γ to be infinite. The image of a based loop is contained in a finite subgraph by Lemma 12.4.1. Enlarge this subgraph by adjoining the finitely many tree paths from v_0 to its vertices. The result is a finite connected subgraph Γ_0 for which $T \cap \Gamma_0$ is a spanning tree. The finite case expresses the given loop class as a word in the classes $[\lambda_e]$ for edges of Γ_0 outside T. Hence these classes generate $\pi_1(\Gamma, v_0)$.

Suppose a reduced nonempty word in finitely many of these generators were nullhomotopic in Γ. The representing loops and the image of a nullhomotopy are contained in a finite subgraph, again by the compactness lemma. Adjoin the tree paths from v_0 to its vertices to obtain a finite connected subgraph Γ_1 with spanning tree $T \cap \Gamma_1$. The alleged nullhomotopy now takes place in Γ_1. In its finite-graph free basis the word is still reduced and nonempty, contradicting freeness. Thus there are no relations among the indicated generators. □

12.4.2 Covering Graphs, Free Subgroups, and Their Ranks

Graphs satisfy the local hypotheses of covering space classification. At an interior point of an edge, use a small interval. At a vertex, use an open star consisting of that vertex and a sufficiently short half-interval in each incident edge; for a loop edge, use disjoint half-intervals at its two ends. These stars form a neighborhood basis, even at a vertex of infinite valence. Radially contracting every half-interval gives a contraction of the star. Continuity at the vertex follows because a smaller star remains inside any prescribed star throughout the contraction. Thus graphs are locally path connected and semilocally simply connected.

Lemma 12.4.3. The total space of a covering of a graph has a graph structure in which vertices and edges project to vertices and edges of the base.

Proof. Take the inverse images of the vertices as the new vertices. Orient each edge of the base and lift its characteristic path $I \to \Gamma$ once from every point above its initial vertex. Each lift supplies a closed edge with endpoints among the new vertices. Over the interior of the base edge, path lifting or the lifting criterion for the interval identifies each component of its inverse image homeomorphically with the interval. Thus the lifted open edges are disjoint and cover the inverse images of all open edges.

These attachments define a CW graph with the same underlying set as the covering space. Their topologies agree. On an evenly covered interval this follows from the inverse branches

of the covering. Near a lifted vertex, choose an evenly covered open star downstairs. Its sheet is homeomorphic to that star and meets each incident lifted edge in the corresponding half-interval. The weak graph topology on such a star is the same topology, since openness can be tested separately on its incident half-intervals. These neighborhoods cover the total space, proving the claim. □

Theorem 12.4.4. Every subgroup of a free group is free.

Proof. Let F be the free group on a set S, and let $H \leq F$. Form the CW graph R_S with one vertex and one oriented loop for each element of S. Theorem 12.4.2, with the single vertex as spanning tree, identifies $\pi_1(R_S)$ with F.

Covering space classification gives a connected based cover $p : Y_H \to R_S$ whose image subgroup is H. The covering-space injection identifies $\pi_1(Y_H)$ with H. By Lemma 12.4.3, the space Y_H is a graph, and Theorem 12.4.2 makes its fundamental group free. This proves the assertion for arbitrary sets of generators and arbitrary subgroups. □

Corollary 12.4.5. Let F_r be a free group on $r \geq 1$ generators, and let $H \leq F_r$ have finite index d. Then H is free of rank

$$1 + d(r-1).$$

Proof. Use the d-sheeted connected covering of the rose with r loops. There are d vertices, one for each point over the base vertex. Each of the r oriented base edges has one lift starting at each lifted vertex, giving dr edges. A spanning tree has $d-1$ edges, so Theorem 12.4.2 gives $dr-(d-1)$ free generators. □

Example 12.4.6. Let R be the rose with loops a, b. Construct a graph with vertices $1, 2, 3$, an oriented a-edge from each vertex i to $(12)(i)$, and an oriented b-edge from each i to $(23)(i)$. Map each labeled edge to the corresponding loop of R. At every vertex there is one incoming and one outgoing edge of each label, so small vertex stars map homeomorphically to the base star. The edge interiors do as well. This is a connected three-sheeted covering.

With base vertex 1, its subgroup $H \leq F(a,b)$ consists of words whose successive edge lifts return to 1. The word b belongs to H, whereas aba^{-1} carries 1 to 3 and does not. Hence H is not normal. A deck transformation permutes the vertices and commutes with both monodromy permutations. Since (23) fixes only vertex 1, the image of vertex 1 under a deck transformation must be 1. Unique lifting then makes that transformation the identity. Thus this three-sheeted cover has trivial deck group.

There are six edges and three vertices, so H has rank four. Choose the spanning tree consisting of the a-edge from 1 to 2 and the b-edge from 2 to 3. The four edges outside this tree give the explicit free basis

$$b, \qquad a^2, \qquad abab^{-1}a^{-1}, \qquad ab^2a^{-1}.$$

The labels in these words record the projected loops $\tau_u * e * \overline{\tau_v}$ from Theorem 12.4.2.

Example 12.4.7. Consider the homomorphism $F(a,b) \to \mathbb{Z}$ sending a to 1 and b to 0. Its kernel corresponds to the graph whose vertices are the integers, with an a-edge from n to $n+1$ and a b-loop at each n. Projection to the two-loop rose is a covering, and lifting a word changes the vertex by its total exponent of a. Its associated subgroup is therefore the stated kernel.

The line of a-edges is a spanning tree. The b-loops give the free basis

$$\{a^n b a^{-n} \mid n \in \mathbb{Z}\}.$$

Thus a subgroup of a free group on two generators can be free on countably infinitely many generators. Translation of the integer vertices gives the deck group $\mathbb{Z}$, in agreement with the normal-subgroup calculation.

For covering spaces and their classification in categorical form, see J. P. May, *A Concise Course in Algebraic Topology*, Chapter 3, especially Sections 7–8, pp. 28–33. Chapter 4, pp. 35–38, treats graphs and their applications to groups. The author's text is available at https://www.math.uchicago.edu/~may/CONCISE/ConciseRevised.pdf.

For supplementary comparisons, see Allen Hatcher, *Algebraic Topology*, Section 1.3: Propositions 1.30–1.34, pp. 60–62, for lifting; the construction and classification on pp. 63–68, including Theorem 1.38; and Proposition 1.39, pp. 71–72, for deck transformations. Section 1.A, pp. 83–85, proves the graph results and the theorem that subgroups of free groups are free. The author's chapter is available at https://pi.math.cornell.edu/~hatcher/AT/ATch1.pdf.

Mapping Complexes and Chain Homotopies

Algebraic Topology I | Fall 2026

A.1 Mapping Complexes and Chain Homotopies

The tensor product and mapping complex express chain homotopies through an interval complex. We use the homological grading of Lecture 2. Write Ab for the category of abelian groups and Ch(Ab) for its category of chain complexes. For a complex $A_\bullet$, write ∂_A for its differential.

A.1.1 The Tensor–Mor Adjunction

The mapping complex is the homological form of Lecture 3, Remark 3.3.5, obtained by reindexing cochains. It retains graded maps as well as chain maps. Its differential records which graded maps commute with the differentials of the two complexes.

Definition A.1.1. For complexes $A_\bullet, B_\bullet$ of abelian groups, define the tensor product by

$$(A \otimes B)_n = \bigoplus_{p+q=n} A_p \otimes_{\mathbb{Z}} B_q, \qquad \partial(\lambda \otimes \mu) = \partial_A \lambda \otimes \mu + (-1)^p \lambda \otimes \partial_B \mu \quad (\lambda \in A_p).$$

Define the *mapping complex* by

$$\underline{\mathrm{Mor}}_{\mathrm{Ch(Ab)}}(A, B)_r = \prod_{n \in \mathbb{Z}} \mathrm{Mor}_{\mathrm{Ab}}(A_n, B_{n+r}), \qquad \partial\phi = \partial_B \phi - (-1)^r \phi \partial_A.$$

Thus ∂ lowers degree by one. Its degree-zero cycles are chain maps, and a degree-one element s with $\partial s = f_1 - f_0$ is a chain homotopy from f_0 to f_1.

For three complexes $A_\bullet, B_\bullet, C_\bullet$, composition of graded maps gives a chain map

$$\underline{\mathrm{Mor}}_{\mathrm{Ch(Ab)}}(B, C) \otimes \underline{\mathrm{Mor}}_{\mathrm{Ch(Ab)}}(A, B) \longrightarrow \underline{\mathrm{Mor}}_{\mathrm{Ch(Ab)}}(A, C).$$

Indeed, for ψ of degree s, $\partial(\psi\phi) = (\partial\psi)\phi + (-1)^s \psi(\partial\phi)$. Composition is associative, and the identity maps are degree-zero cycles. Hence the morphisms between complexes themselves form complexes, compatibly with composition.

Proposition A.1.2. For complexes of abelian groups, there is a natural isomorphism of chain complexes

$$\underline{\mathrm{Mor}}_{\mathrm{Ch(Ab)}}(A \otimes B, C) \cong \underline{\mathrm{Mor}}_{\mathrm{Ch(Ab)}}(A, \underline{\mathrm{Mor}}_{\mathrm{Ch(Ab)}}(B, C)).$$

It sends a homogeneous map F to the map $\widetilde{F}$ defined by $\widetilde{F}(\lambda)(\mu) = F(\lambda \otimes \mu)$.

Proof. Degreewise, the isomorphism is the tensor–Mor adjunction for abelian groups, with the direct sums in the tensor product becoming products of morphism groups. For F of degree r and $\lambda \in A_p$, the two differentials agree since

$$\begin{aligned} \widetilde{\partial F}(\lambda)(\mu) &= \partial_C F(\lambda \otimes \mu) - (-1)^r F(\partial_A \lambda \otimes \mu) \\ &\qquad - (-1)^{r+p} F(\lambda \otimes \partial_B \mu) \\ &= (\partial \widetilde{F})(\lambda)(\mu). \end{aligned}$$

□

For fixed B, this says concretely that forming $A \otimes B$ is left adjoint to forming $\underline{\mathrm{Mor}}_{\mathrm{Ch(Ab)}}(B, C)$, including at the level of mapping complexes.

A.1.2 The Interval and Chain Homotopies

Use the standard topological simplices $|\Delta^n|$ of Lecture 7, Construction 7.1.11. The affine face inclusion d^i inserts a zero in the ith coordinate, and s^i adds adjacent coordinates. In particular, $|\Delta^1| = I$.

For a set S, write $\mathbb{Z}\{S\}$ for the free abelian group on S. Use the complex $\mathbb{Z}\{\Delta^1\}^{\mathrm{nd}}_\bullet$ from Lecture 7, Definition 7.1.15, whose degree-n group is free abelian on the set of nondegenerate n-simplices of Δ^1. Thus it has generators e_0, e_1 in degree zero and e_{01} in degree one, with $\partial e_{01} = e_1 - e_0$, and is zero in other degrees.

The tensor differential explains the sign in the degree-one component of its cylinder homotopy. For chain maps $f_0, f_1 : A_\bullet \to B_\bullet$, a chain homotopy from f_0 to f_1 is equivalently a chain map

$$h : A_\bullet \otimes \mathbb{Z}\{\Delta^1\}^{\mathrm{nd}}_\bullet \longrightarrow B_\bullet$$

whose restrictions to the endpoints are f_0, f_1. Indeed, for $\lambda \in A_n$, a degree-one map $s : A_\bullet \to B_{\bullet+1}$ determines the graded map h by

$$h(\lambda \otimes e_0) = f_0(\lambda), \qquad h(\lambda \otimes e_1) = f_1(\lambda), \qquad h(\lambda \otimes e_{01}) = (-1)^n s(\lambda).$$

With the tensor differential of Definition A.1.1, the chain-map identity on $\lambda \otimes e_{01}$ is

$$(-1)^n \partial_B s(\lambda) = (-1)^{n-1} s(\partial_A \lambda) + (-1)^n (f_1(\lambda) - f_0(\lambda)),$$

or $\partial_B s + s\partial_A = f_1 - f_0$. The identities on the two vertices say that f_0, f_1 are chain maps.

Remark A.1.3. For an additive category $\mathcal{A}$, write $\underline{\mathrm{Mor}}_{\mathrm{Ch}(\mathcal{A})}$ for the mapping complex, defined in the same way using $\mathrm{Mor}_{\mathcal{A}}$, with the products taken in abelian groups. Its zeroth homology gives the morphisms in the homotopy category:

$$\mathrm{Mor}_{K(\mathcal{A})}(A_\bullet, B_\bullet) = \mathrm{H}_0 \underline{\mathrm{Mor}}_{\mathrm{Ch}(\mathcal{A})}(A_\bullet, B_\bullet).$$

Indeed, cycles are chain maps and boundaries are differences of chain-homotopic maps.

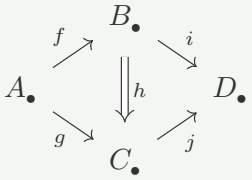
A.2 Homotopy Pushouts of Chain Complexes

A.2.1 The Comparison Map and Its Signs

A square with a chosen chain homotopy induces a comparison from the cone construction of Lecture 7. The degree-one component of the homotopy supplies the third term of that comparison.

The following square is the algebraic setting for the comparison. The endpoint maps and the chosen homotopy are retained as part of its data.

Definition A.2.1. A homotopy-commutative square of complexes consists of chain maps



and a specified chain map

$$h : A_\bullet \otimes \mathbb{Z}\{\Delta^1\}^{\mathrm{nd}}_\bullet \longrightarrow D_\bullet, \qquad h(\nu \otimes e_0) = if(\nu), \quad h(\nu \otimes e_1) = jg(\nu).$$

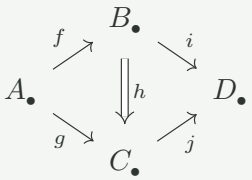
For $\nu \in A_n$, define its degree-one component by $s(\nu) = (-1)^n h(\nu \otimes e_{01})$. The chain-map identity for h is equivalent to

$$\partial_D s + s\partial_A = jg - if. \tag{18}$$

Thus h is a homotopy from if to jg, represented by the degree-one map s.

For the span $B_\bullet \xleftarrow{f} A_\bullet \xrightarrow{g} C_\bullet$, use the cone construction of Lecture 7, Definition 7.1.15, with $A_\bullet, B_\bullet, C_\bullet$ in place of the three singular chain complexes. Denote the resulting complex by $Z_\bullet$ and its structure maps by i_Z, j_Z, h_Z. Write s_Z for the degree-one component of h_Z. A degree-n element of $Z_\bullet$ is written (λ, μ, ν), where $\lambda \in B_n$, $\mu \in C_n$, and $\nu \in A_{n-1}$. The construction applies to arbitrary chain maps because its differential and homotopy identities use only the chain-map identities for f and g.

Proposition A.2.2. For a homotopy-commutative square



as in Definition A.2.1, there is a chain map

$$\Phi_h : Z_\bullet \to D_\bullet, \qquad \Phi_h(\lambda, \mu, \nu) = -i(\lambda) + j(\mu) + s(\nu). \tag{19}$$

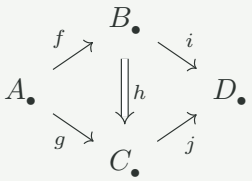
Its composites with the two structure maps i_Z, j_Z are i and j.

Proof. The chain-map condition is the homotopy identity (18). For $(\lambda, \mu, \nu) \in Z_n$, expand the two sides:

$$\begin{aligned}
\Phi_h \partial_Z(\lambda, \mu, \nu) &= -i\partial_B\lambda + j\partial_C\mu + (jg - if)(\nu) - s\partial_A\nu \\
&= -i\partial_B\lambda + j\partial_C\mu + \partial_D s(\nu) && \text{by (18)} \\
&= \partial_D(-i(\lambda) + j(\mu) + s(\nu)) && \text{since } i, j \text{ are chain maps} \\
&= \partial_D \Phi_h(\lambda, \mu, \nu).
\end{aligned}$$

Finally, $\Phi_h(-\lambda, 0, 0) = i(\lambda)$ and $\Phi_h(0, \mu, 0) = j(\mu)$, giving the required composites. □

Definition A.2.3. A homotopy-commutative square of complexes



is a *homotopy pushout square* if its comparison map Φ_h in Proposition A.2.2 is a quasi-isomorphism.

This is a condition in the homotopy theory of complexes with quasi-isomorphisms as equivalences. With the structure maps i_Z, j_Z, h_Z, the canonical complex $Z_\bullet$ is a homotopy pushout: its comparison map to itself is the identity. The square with rightmost complex $D_\bullet$ is a homotopy pushout when Φ_h is a quasi-isomorphism. The sign in its first structure map makes the comparison map restrict to the given maps i and j.

A.2.2 The Long Exact Homology Sequence

The shifted common complex occurs as a quotient of the double mapping cylinder. This quotient produces the connecting map, including its sign.

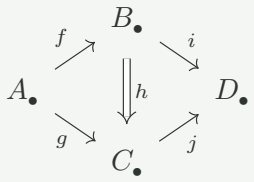
Apply the mapping-cone long exact sequence of Lecture 2 to $(f,g) : A_\bullet \to B_\bullet \oplus C_\bullet$. Let $k : B_\bullet \oplus C_\bullet \to Z_\bullet$ be the summand inclusion and let $p : Z_\bullet \to A_\bullet[1]$ be the projection onto the shifted summand.

For a cycle $\nu \in A_{n-1}$, the lift $(0,0,\nu) \in Z_n$ has boundary $(f(\nu), g(\nu), 0)$. Hence the boundary homomorphism out of $\mathrm{H}_n(A[1])$ is (f_*, g_*). The projection sends $[(\lambda,\mu,\nu)]$ to $[\nu]$ under the shift identification. We take the negative of this projection on homology, so that the convention agrees with the Mayer–Vietoris sequence in Lecture 7:

$$\delta_Z[(\lambda,\mu,\nu)] = -[\nu]. \tag{20}$$

Negating a homomorphism leaves its kernel and image unchanged, so this convention preserves exactness. The resulting class is independent of the chosen cycle representative because p is a chain map.

Theorem A.2.4. For a homotopy pushout square of complexes



there is a long exact sequence whose five consecutive terms are

$$\mathrm{H}_n(A) \xrightarrow{(f_*,g_*)} \mathrm{H}_n(B) \oplus \mathrm{H}_n(C) \xrightarrow{j_*-i_*} \mathrm{H}_n(D) \xrightarrow{\delta_D} \mathrm{H}_{n-1}(A) \xrightarrow{(f_*,g_*)} \mathrm{H}_{n-1}(B) \oplus \mathrm{H}_{n-1}(C). \tag{21}$$

Its connecting map is $\delta_D = \delta_Z(\Phi_{h,*})^{-1}$.

Proof. By Definition A.2.3, the homomorphism $\Phi_{h,*} : \mathrm{H}_n(Z) \to \mathrm{H}_n(D)$ is an isomorphism in every degree. Transport the mapping-cone exact sequence along these isomorphisms, using the convention (20). Formula (19) gives

$$\Phi_h k(\lambda,\mu) = j(\mu) - i(\lambda),$$

so the map into $\mathrm{H}_n(D)$ is $j_* - i_*$. The transported connecting map is the one stated, proving (21). □

A cycle $\zeta \in D_n$ can therefore be handled as follows. Choose a cycle $(\lambda,\mu,\nu) \in Z_n$ whose image under Φ_h is homologous to ζ. The cycle conditions and the image are

$$\partial_A \nu = 0, \qquad \partial_B \lambda = -f(\nu), \qquad \partial_C \mu = -g(\nu),$$
$$\zeta \sim -i(\lambda) + j(\mu) + s(\nu).$$

Then $\delta_D[\zeta] = -[\nu]$. The common chain ν records the boundary that has to be matched between the two pieces. The preceding projection argument proves that this class is independent of the representative in Z_n.

For the mapping-cone specialization $C_\bullet = 0$, the map from $\mathrm{H}_n(B)$ is induced by the negative of the usual summand inclusion. This is the same sign convention as in Lecture 7; the mapping-cone and suspension calculations themselves are given there and in Lecture 2.

Attaching Maps and Euler Characteristic

Algebraic Topology I | Fall 2026 | October 2026

B.1 Attaching Maps and Cellular Boundaries

A torus and a projective plane can each be obtained by attaching a disk to a graph. Their homology differs because the boundary circle traverses the edges with different signed multiplicities. The cellular theorem of Lecture 8 turns those multiplicities into a differential. This lecture develops that interpretation and then explains why alternating cell counts give an invariant even when the individual cell counts do not.

Throughout, homology means ordinary singular homology. We write its coefficient group after a semicolon, including when that group is $\mathbb{Z}$. All cell complexes used for cell counts are finite. An orientation of a positive-dimensional sphere means a chosen generator of its top integral homology. Orientations of disks are chosen so that the boundary homomorphism gives the specified orientation of their boundary spheres.

B.1.1 A Disk Attached to a Graph

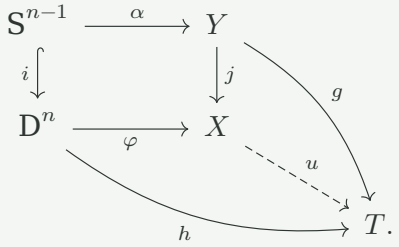
Begin with a graph G having one vertex v and two oriented loop edges a, b. Thus G is the wedge $\mathrm{S}^1 \vee \mathrm{S}^1$. A path denoted by a^{-1} traverses the edge a with the opposite orientation. A word such as $aba^{-1}b^{-1}$ specifies a continuous map from a subdivided circle to G, with each successive interval traversing the indicated edge.

Construction B.1.1. For a space Y, an integer $n \geq 1$, and a continuous map $\alpha : \mathrm{S}^{n-1} \to Y$, define the attached space X by the pushout

$$\begin{array}{ccc} \mathrm{S}^{n-1} & \xrightarrow{\ \alpha\ } & Y \\ {\scriptstyle i}\downarrow & & \downarrow{\scriptstyle j} \\ \mathrm{D}^n & \xrightarrow[\ \varphi\ ]{} & X = Y +_{\mathrm{S}^{n-1}} \mathrm{D}^n. \end{array}$$

Concretely, X is the quotient of $Y + \mathrm{D}^n$ by the identifications $u \sim \alpha(u)$ for $u \in \mathrm{S}^{n-1}$. The map α is the *attaching map*, and φ is the *characteristic map*. Its restriction to the disk interior is a homeomorphism onto the new open cell.

The universal property specifies every continuous map out of this space. Given maps $g : Y \to T$ and $h : \mathrm{D}^n \to T$ satisfying $g\alpha = hi$, there is a unique continuous map $u : X \to T$ with $uj = g$ and $u\varphi = h$. The factorization is displayed by



Continuity follows from the quotient topology. Conversely, the commutativity of the pushout square shows that any map u gives a compatible pair g, h.

For the word $aba^{-1}b^{-1}$, the resulting quotient of a square identifies each pair of opposite edges in the same direction. This is the usual torus $\mathrm{S}^1 \times \mathrm{S}^1$. Its cellular groups with coefficients in an abelian group A are A, A^2, and A in dimensions $0, 1, 2$. Each loop edge has identical initial and terminal vertices, so its boundary is zero. The boundary of the square traverses each edge once in each direction, so its cellular boundary is also zero. Consequently

$$\mathrm{H}_0(\mathrm{S}^1 \times \mathrm{S}^1; A) = A, \qquad \mathrm{H}_1(\mathrm{S}^1 \times \mathrm{S}^1; A) = A^2, \qquad \mathrm{H}_2(\mathrm{S}^1 \times \mathrm{S}^1; A) = A.$$

The next subsection justifies the signed traversal calculation and applies it beyond this example.

B.1.2 The Degree of an Attaching Map

An attaching map usually takes values in a whole skeleton, rather than in a sphere. To obtain an integer for a particular target cell, first collapse all other cells in the relevant dimension.

Write $X^{(r)} = \mathrm{sk}_r X$ for the r-skeleton, and set $X^{(-1)} = \emptyset$. For an oriented n-cell e with $n \geq 2$, let $\alpha_e : \mathrm{S}^{n-1} \to X^{(n-1)}$ be its attaching map. The skeletal quotient is a wedge of spheres. For an oriented $(n-1)$-cell b, the composite

$$\mathrm{S}^{n-1} \xrightarrow{\alpha_e} X^{(n-1)} \xrightarrow{q} X^{(n-1)}/X^{(n-2)} \xrightarrow{p_b} \mathrm{S}^{n-1}$$

collapses every sphere of that wedge except the one corresponding to b. Choose its identification with the target sphere to respect the cell orientation.

The cellular attaching-degree formula following Lecture 8, Definition 8.2.1, says that the coefficient of b in ∂e is $\deg(p_b q \alpha_e)$. Thus

$$\partial e = \sum_b \deg(p_b q \alpha_e)\, b \quad \text{in } \mathrm{C}^{\mathrm{cell}}_{n-1}(X; \mathbb{Z}).$$

This formula follows from the definition of the cellular differential as the relative boundary followed by the quotient map. On each disk, naturality identifies that boundary with the oriented boundary sphere. Projecting the target relative group onto the summand for b gives the displayed sphere map. We use Lecture 8, Theorem 8.2.2, to pass from this chain complex to $\mathrm{H}_*(X; A)$; its proof is not needed again here. With coefficients in A, the same integer acts by repeated addition on A.

For $n = 1$, the formula has the familiar form $\partial e = v_+ - v_-$, where v_- and v_+ are the initial and terminal vertices. This treats distinct vertices as well as loops, without introducing degree for a map between disconnected zero-spheres.

Proposition B.1.2. Let G be a graph with one vertex and oriented edges $a_1, \dots, a_r$. Attach a 2-cell e along a finite word in the letters a_i, a_i^{-1}. If m_i is the number of occurrences of a_i minus the number of occurrences of a_i^{-1}, then

$$\partial e = \sum_{i=1}^{r} m_i a_i.$$

Proof. Collapse all edges except a_i to the vertex. Choose a point in the interior of the surviving edge. Its inverse images lie in the interiors of the intervals of the attaching circle labeled a_i or a_i^{-1}. On each such interval the map is a homeomorphism near the inverse image. A positively traversed interval contributes 1 to the local degree; a negatively traversed interval contributes -1. These signs follow by comparison with the identity and a reflection of the oriented circle. The local-degree sum and the reflection calculation were proved in

Lecture 8, Lemmas 8.1.10–8.1.11. Their sum is m_i, and the attaching-degree formula proves the assertion. □

For a projective plane, use one loop edge a and attach the disk along aa. The boundary circle covers the edge circle twice in the same direction. The cellular complex is therefore

$$0 \longrightarrow A \xrightarrow{2} A \xrightarrow{0} A \longrightarrow 0.$$

It follows that $\mathrm{H}_0(\mathbb{R}P^2; A) = A$, $\mathrm{H}_1(\mathbb{R}P^2; A) = A/2A$, and $\mathrm{H}_2(\mathbb{R}P^2; A) = \ker(2 \colon A \to A)$. In particular, the top integral group vanishes, whereas the top group with coefficients in $\mathbb{Z}/2\mathbb{Z}$ is nonzero. Both conclusions arise from the same attaching map.

B.2 Surfaces from Polygonal Attaching Maps

A polygonal presentation exhibits a surface as a graph with one attached 2-cell. The resulting differential records the total signed number of traversals of each edge. It is often much simpler than the attaching word itself.

B.2.1 Orientable Polygonal Surfaces

For $g \geq 1$, let Σ_g be the polygonal surface with one vertex, edges $a_1, b_1, \dots, a_g, b_g$, and one face attached along

$$a_1 b_1 a_1^{-1} b_1^{-1} \cdots a_g b_g a_g^{-1} b_g^{-1}.$$

This is the usual orientable surface of genus g; its polygon is the definition used here, so no classification theorem for surfaces is needed. Every letter has exponent sum zero. Proposition B.1.2 therefore gives zero cellular differentials. For every abelian group A,

$$\mathrm{H}_k(\Sigma_g; A) \cong \begin{cases} A, & k = 0, 2, \\ A^{2g}, & k = 1, \\ 0, & k > 2. \end{cases}$$

For $g = 0$, take $\Sigma_0 = \mathrm{S}^2$, with one vertex and one 2-cell. The same formula holds with $A^0 = 0$.

The face supplies a cycle because the two occurrences of each edge cancel. That explanation concerns the cellular differential; it does not assert that the attaching map is homotopic to a constant. The torus example already illustrates how much simpler a homological invariant can be than the map defining it.

B.2.2 Nonorientable Polygonal Surfaces

For $h \geq 1$, let N_h be the polygonal surface with one vertex, edges $a_1, \dots, a_h$, and a face attached along $a_1 a_1 \cdots a_h a_h$. Thus $N_1 = \mathbb{R}P^2$. Its cellular boundary sends the face generator to $2a_1 + \cdots + 2a_h$.

To compute with an arbitrary abelian group A, apply the invertible coordinate change

$$A^h \longrightarrow A^h, \qquad (x_1, x_2, \dots, x_h) \longmapsto (x_1, x_2 - x_1, \dots, x_h - x_1).$$

It sends the boundary vector $(2x, \dots, 2x)$ to $(2x, 0, \dots, 0)$. Thus

$$\begin{aligned} \mathrm{H}_0(N_h; A) &\cong A, \\ \mathrm{H}_1(N_h; A) &\cong (A/2A) \oplus A^{h-1}, \\ \mathrm{H}_2(N_h; A) &\cong \ker(2 \colon A \to A). \end{aligned}$$

All higher groups vanish. With integral coefficients, the last group is zero and the first homology group is $\mathbb{Z}^{h-1} \oplus \mathbb{Z}/2\mathbb{Z}$. With coefficients in $\mathbb{Z}/2\mathbb{Z}$, the face is a cycle and the first homology group has dimension h. The change in coefficients changes both homology groups in a compatible way; the Euler characteristic below remains the same.

B.3 Euler Characteristic and Finite Chain Complexes

The sphere has cell structures with different numbers of cells. For example, a 2-sphere has a structure with one vertex and one face, and a tetrahedral structure with four vertices, six edges, and four faces. Both give alternating sum 2. The explanation is an algebraic cancellation of boundaries in consecutive dimensions.

B.3.1 The Alternating-Dimension Identity

For a finite cell complex X, let $c_k(X)$ be the number of its k-cells. Define its *Euler characteristic* by

$$\chi(X) = \sum_{k \geq 0} (-1)^k c_k(X).$$

The sum is finite. A priori this definition uses a cell structure. The following identity proves its independence of that structure.

Proposition B.3.1. Let $C_\bullet$ be a chain complex of finite-dimensional vector spaces over a field K, with $C_k = 0$ outside finitely many nonnegative degrees. Then

$$\sum_k (-1)^k \dim_K C_k = \sum_k (-1)^k \dim_K \mathrm{H}_k(C_\bullet).$$

Proof. Write $Z_k = \ker(\partial_k : C_k \to C_{k-1})$ and $B_k = \mathrm{im}(\partial_{k+1} : C_{k+1} \to C_k)$, and set $B_{-1} = 0$. The two exact sequences

$$0 \longrightarrow Z_k \longrightarrow C_k \xrightarrow{\partial_k} B_{k-1} \longrightarrow 0,$$

$$0 \longrightarrow B_k \longrightarrow Z_k \longrightarrow \mathrm{H}_k(C_\bullet) \longrightarrow 0$$

give $\dim C_k = \dim B_{k-1} + \dim B_k + \dim \mathrm{H}_k(C_\bullet)$. Multiply by $(-1)^k$ and sum. Every boundary dimension occurs once with each sign, so the boundary terms cancel. All sums are finite, including the zero boundary groups beyond the complex's endpoints. □

Theorem B.3.2. For every finite cell complex X and every field K,

$$\chi(X) = \sum_{k \geq 0} (-1)^k \dim_K \mathrm{H}_k(X; K).$$

In particular, Euler characteristic is independent of the cell structure and is invariant under homotopy equivalence between finite cell complexes.

Proof. The cellular complex $\mathrm{C}_\bullet^{\mathrm{cell}}(X; K)$ has dimension $c_k(X)$ in degree k. Apply Proposition B.3.1, then the cellular homology theorem. The resulting expression depends only on the homology groups with coefficients in K. Homotopy invariance of these groups proves the last assertion. □

With $K = \mathbb{Q}$, these dimensions are the rational Betti numbers. The integral cellular complex gives the equivalent formula

$$\chi(X) = \sum_k (-1)^k \operatorname{rank}_{\mathbb{Z}} \mathrm{H}_k(X; \mathbb{Z}).$$

Indeed, tensoring an exact sequence of abelian groups with $\mathbb{Q}$ preserves exactness: if an element becomes zero after localization, an integer multiple of it was already zero. Consequently tensoring commutes with the kernels and images of the integral cellular differential. The rational homology dimensions are therefore the ranks of the integral homology groups. Finite torsion summands have rank zero.

For $\mathbb{R}P^2$, rational homology has dimensions $1, 0, 0$, while mod-two homology has dimensions $1, 1, 1$. Both alternating sums equal 1. Individual Betti numbers can depend on the field even though their alternating sum does not.

B.3.2 Attachments, Unions, and Products

The cell-count definition is often the most efficient way to compute Euler characteristic once its invariance is known.

Proposition B.3.3. For finite cell complexes, the following formulas hold.

1. For $n \geq 1$ and an attaching map $\alpha : \mathrm{S}^{n-1} \to Y^{(n-1)}$, attaching one n-cell to Y gives $\chi(Y +_{\mathrm{S}^{n-1}} \mathrm{D}^n) = \chi(Y) + (-1)^n$.
2. If $X = U \cup V$ for subcomplexes U, V, then $\chi(X) = \chi(U) + \chi(V) - \chi(U \cap V)$.
3. For finite cell complexes X, Y, one has $\chi(X \times Y) = \chi(X)\chi(Y)$.

Proof. For the first assertion, exactly one cell has been added. For the second, every cell of $U \cap V$ is counted twice in $c_k(U) + c_k(V)$ and every other cell of X is counted once. Subtract the intersection count in each dimension and take the alternating sum.

For the third assertion, the product-cell construction of Lecture 8, Proposition 8.1.5, gives one $(p+q)$-cell for every pair consisting of a p-cell of X and a q-cell of Y. Therefore

$$\begin{aligned} \chi(X \times Y) &= \sum_{p,q} (-1)^{p+q} c_p(X) c_q(Y) \\ &= \left(\sum_{p} (-1)^p c_p(X) \right) \left(\sum_{q} (-1)^q c_q(Y) \right). \end{aligned}$$

The sums are finite. This argument uses product cells and does not require a Künneth theorem. □

The subcomplex hypothesis in the union formula supplies a common cell structure. For arbitrary subsets, the displayed cell-count argument is not available. A separate Mayer–Vietoris argument can prove an analogous formula under appropriate finiteness and excision hypotheses.

If X is connected and has a chosen vertex, then identifying that vertex with a chosen vertex of another connected finite cell complex Y gives

$$\chi(X \vee Y) = \chi(X) + \chi(Y) - 1.$$

There is one shared vertex and no other shared cell. For a nonempty subcomplex $Y \subseteq X$, collapsing Y to one vertex gives

$$\chi(X/Y) = \chi(X) - \chi(Y) + 1.$$

The quotient has the cells outside Y and one additional vertex. Both formulas include the contribution of degree-zero homology; omitting that contribution would give the wrong constant term.

B.4 Computations and Consequences

Euler characteristic retains less information than the full cellular complex, but it can rule out proposed homotopy equivalences immediately. The attaching map determines the differential, whereas the Euler characteristic depends only on how many cells were attached in each dimension.

B.4.1 Graphs, Spheres, and Surfaces

A finite connected graph with v vertices and e edges has $\chi = v - e$. To see its first Betti number directly, the image of the cellular boundary consists of vectors in $\mathbb{Z}^v$ whose coordinates sum to zero. Connectivity implies that this entire subgroup occurs: an edge path from a fixed vertex to any other vertex has boundary equal to the difference of those vertices. The boundary therefore has rank $v - 1$. Its kernel is free of rank $e - v + 1$, and there are no 2-cells. Thus

$$\mathrm{H}_1(G;\mathbb{Z}) \cong \mathbb{Z}^{e-v+1}, \qquad \chi(G) = 1 - (e - v + 1).$$

This expresses the alternating count as one minus the number of independent homological cycles.

For $n \geq 1$, the sphere's structure with one vertex and one n-cell gives $\chi(\mathrm{S}^n) = 1 + (-1)^n$. A nonempty contractible finite cell complex has Euler characteristic 1, since its rational homology is that of a point. Therefore no positive-dimensional sphere is contractible.

The polygonal structures above give

$$\chi(\Sigma_g) = 2 - 2g, \qquad \chi(N_h) = 2 - h.$$

These agree with the rational homology computations. In particular, orientable surfaces of distinct genera cannot be homotopy equivalent. Euler characteristic alone does not distinguish Σ_g from N_{2g} when $g \geq 1$, since the displayed values coincide. Their integral second homology does distinguish them: it is $\mathbb{Z}$ for the former and zero for the latter.

B.4.2 Products and a Practical Computation Procedure

Every positive-dimensional torus has Euler characteristic zero, since $\chi(\mathrm{S}^1) = 0$ and the product formula applies. More generally, $\chi(X \times \mathrm{S}^1) = 0$ for every finite cell complex X. The cell structure of $\mathbb{C}P^m$ from Lecture 8 has one cell in each even dimension from 0 through $2m$, so

$$\chi(\mathbb{C}P^m) = m + 1, \qquad \chi(\mathbb{C}P^m \times \mathbb{C}P^n) = (m + 1)(n + 1).$$

These are consequences of the cell structure even before one writes the product's homology groups.

For a finite cell complex presented by attaching maps, the computations in this lecture suggest the following procedure.

1. Choose orientations and list the cellular groups $\mathrm{C}_k^{\mathrm{cell}}(X;\mathbb{Z})$.
2. For each adjacent pair of dimensions, collapse to one target-cell sphere and compute the resulting attaching degree. For polygonal words, use signed exponent sums.
3. Write the resulting integer matrices. For coefficients in an abelian group A, let the same matrices act on direct sums of copies of A.
4. Compute kernels modulo images to obtain $\mathrm{H}_k(X;A)$. Independently compute the alternating cell count and compare it with the alternating homology dimensions over a field.

The last comparison can detect an erroneous rank, but it cannot detect every erroneous differential. The torus and the nonorientable surface N_2 both have Euler characteristic zero; their integral homology groups differ. Degree and Euler characteristic are most effective when used together: degree determines the cellular maps, and Euler characteristic records the alternating total of the resulting homology dimensions.

For further reading, see Allen Hatcher, *Algebraic Topology*, Section 2.2: the local-degree formula is Proposition 2.30, the cellular boundary formula is on pp. 140–141, and the Euler characteristic identity is Theorem 2.44 on pp. 146–147.

A Survey of the Künneth Theorem

Algebraic Topology I | Fall 2026 | October 2026

C.1 The Cross Product and the Product Formula

A product cell has one direction from each factor. The homology of a product is therefore related to the graded tensor product of the homology of its factors. Over a field, this gives the whole answer. Over the integers, torsion can contribute additional classes one degree higher. We study both cases through the product chain complex established in Lecture 8.

Throughout this lecture, all spaces under consideration are finite cell complexes. Coefficients are always displayed. When R is a principal ideal domain, tensor products of R-modules are taken over R. The notation $\mathrm{H}_n(C_\bullet)$ denotes the algebraic homology of a specified chain complex of R-modules; it is distinct from the space-homology notation $\mathrm{H}_n(X;R)$.

C.1.1 The Canonical Graded Map

For a commutative coefficient ring R, write $\mathrm{H}_*(X;R)$ for the graded R-module whose degree-n part is $\mathrm{H}_n(X;R)$. Its graded tensor product with the homology of Y has degree-n part

$$\bigl(\mathrm{H}_*(X;R)\otimes_R \mathrm{H}_*(Y;R)\bigr)_n = \bigoplus_{p+q=n} \mathrm{H}_p(X;R)\otimes_R \mathrm{H}_q(Y;R).$$

The cross product constructed in Lecture 8 is the canonical graded map

$$\boxtimes : \mathrm{H}_*(X;R)\otimes_R \mathrm{H}_*(Y;R) \longrightarrow \mathrm{H}_*(X\times Y;R).$$

It comes from triangulating products of singular simplices. On relative cells it sends the two oriented cell classes to the product-cell class, with the directions of X placed before those of Y.

For continuous maps $f : X \to X'$ and $g : Y \to Y'$, naturality is the commutativity of the square

$$\begin{array}{ccc} \mathrm{H}_*(X;R)\otimes_R \mathrm{H}_*(Y;R) & \xrightarrow{\ \boxtimes\ } & \mathrm{H}_*(X\times Y;R) \\ \downarrow{\scriptstyle f_*\otimes g_*} & & \downarrow{\scriptstyle (f\times g)_*} \\ \mathrm{H}_*(X';R)\otimes_R \mathrm{H}_*(Y';R) & \xrightarrow[\ \boxtimes\ ]{} & \mathrm{H}_*(X'\times Y';R). \end{array}$$

Thus a product computation also computes maps, once the maps on the factors are known.

Construction C.1.1. Put $C_\bullet = \mathrm{C}^{\mathrm{cell}}_\bullet(X;R)$ and $D_\bullet = \mathrm{C}^{\mathrm{cell}}_\bullet(Y;R)$. The product-cell construction of Lecture 8 gives

$$\mathrm{C}^{\mathrm{cell}}_\bullet(X\times Y;R) \cong C_\bullet \otimes_R D_\bullet, \qquad (C_\bullet \otimes_R D_\bullet)_n = \bigoplus_{p+q=n} C_p \otimes_R D_q.$$

On the summand $C_p \otimes_R D_q$, the differential is

$$\partial_C \otimes 1 + (-1)^p 1 \otimes \partial_D.$$

The two mixed terms in its square cancel. The cellular comparison theorem, Theorem 8.2.2, identifies the resulting homology map from the graded tensor product with the singular cross product.

C.1.2 The Formula over a Field

The field case of Theorem 8.2.3 is the most direct form of the Künneth theorem. We recall its algebraic reason, which also identifies what changes over the integers.

Theorem C.1.2. For finite cell complexes X, Y and a field k, the cross product is an isomorphism of graded vector spaces

$$\mathrm{H}_*(X;k) \otimes_k \mathrm{H}_*(Y;k) \xrightarrow{\cong} \mathrm{H}_*(X \times Y;k).$$

The same conclusion holds over a principal ideal domain R if both factors have free homology over R in every degree.

Proof. Let $C_\bullet$ be a bounded complex of free R-modules with free homology. Write $Z_n(C)$ and $B_n(C)$ for its cycles and boundaries. The two sequences

$$0 \longrightarrow B_n(C) \longrightarrow Z_n(C) \longrightarrow \mathrm{H}_n(C_\bullet) \longrightarrow 0, \qquad 0 \longrightarrow Z_n(C) \longrightarrow C_n \longrightarrow B_{n-1}(C) \longrightarrow 0$$

split: the first quotient is free by hypothesis, and the second is free by Lecture 3, Theorem 3.1.10. Choosing sections decomposes $C_\bullet$ into its homology with zero differential and contractible two-term complexes. For a contracting homotopy h on a summand of the first factor, the map $H(\alpha \otimes \beta) = h(\alpha) \otimes \beta$ contracts its tensor product with the second factor. If k instead contracts a summand of the second factor, use

$$K(\alpha \otimes \beta) = (-1)^p \alpha \otimes k(\beta) \qquad (\alpha \in C_p).$$

In the first case the terms involving the second differential cancel in $\partial H + H\partial$, leaving $(\partial h + h\partial)(\alpha)\otimes\beta = \alpha\otimes\beta$. In the second case the terms involving $\partial\alpha$ cancel, leaving $\alpha\otimes(\partial k + k\partial)(\beta) = \alpha\otimes\beta$. Apply this decomposition to the cellular complexes of both factors. Only the tensor product of the two homology complexes contributes homology. The isomorphism obtained on homology is the canonical cycle-product map, even though the decompositions used to prove that it is an isomorphism were chosen. Construction C.1.1 identifies this map with the cross product. Over a field every vector space is free. □

C.2 Torsion in the Künneth Formula

The free-homology hypothesis in Theorem C.1.2 is useful, but products such as $\mathbb{R}P^2 \times \mathbb{R}P^2$ require a correction. The correction is the first Tor group from Lecture 3. We recall its computation before deriving the product formula.

C.2.1 The Tor Term and Its Degree

Let M, N be modules over a principal ideal domain R. A free surjection onto M has a free kernel, so there is a free resolution of length one,

$$0 \longrightarrow P_1 \xrightarrow{u} P_0 \longrightarrow M \longrightarrow 0.$$

By the definition of Tor in Lecture 3,

$$\mathrm{Tor}_1^R(M,N) = \ker(u \otimes 1 : P_1 \otimes_R N \longrightarrow P_0 \otimes_R N),$$

with the resolution-independent interpretation supplied by the comparison theorem. Right exactness gives the cokernel $M \otimes_R N$ of the same map. In particular,

$$\mathrm{Tor}_1^{\mathbb{Z}}(\mathbb{Z}/m\mathbb{Z}, N) = N[m], \qquad (\mathbb{Z}/m\mathbb{Z}) \otimes_{\mathbb{Z}} N = N/mN,$$

where $m \geq 1$ and $N[m]$ is the kernel of multiplication by m on N. These formulas follow from the resolution with differential $m : \mathbb{Z} \to \mathbb{Z}$. Tor vanishes if either argument is free. For positive integers m, n, both the tensor product and the first Tor group of $\mathbb{Z}/m\mathbb{Z}$ and $\mathbb{Z}/n\mathbb{Z}$ are cyclic of order $\gcd(m, n)$.

Theorem C.2.1. For finite cell complexes X, Y and a principal ideal domain R, the cross product is the first map in a split short exact sequence

$$0 \longrightarrow \bigoplus_{p+q=n} \mathrm{H}_p(X;R) \otimes_R \mathrm{H}_q(Y;R) \xrightarrow{\boxtimes} \mathrm{H}_n(X \times Y;R) \longrightarrow \bigoplus_{p+q=n-1} \operatorname{Tor}_1^R(\mathrm{H}_p(X;R), \mathrm{H}_q(Y;R)) \longrightarrow 0.$$

The sequence is natural for cellular maps of the factors. A splitting need not be natural. If the homology of one factor is free in every degree, its Tor terms vanish and the cross product is an isomorphism.

The tensor terms in degree n come from degrees summing to n. A Tor term from degrees p, q appears instead in degree $p + q + 1$. The degree shift is part of the theorem and is essential in computations.

C.2.2 Why the Exact Sequence Splits

We prove the algebraic assertion for bounded free complexes. This is sufficient for the finite cell complexes in the theorem and keeps the comparison with cells explicit.

Proof. Let $C_\bullet, D_\bullet$ be the cellular complexes over R. Let $Z(C)$ have $Z_p(C) = \ker \partial_p$ in degree p and zero differential, and let $B(C)[1]$ have $B_{p-1}(C) = \operatorname{im} \partial_p$ in degree p and zero differential. The sequence

$$0 \longrightarrow Z(C) \longrightarrow C_\bullet \xrightarrow{\partial} B(C)[1] \longrightarrow 0$$

is exact and split in each degree, because every boundary module is free. Tensoring with $D_\bullet$ therefore preserves this short exact sequence of complexes.

The long exact homology sequence can be computed term by term. Since cycles and boundaries of $C_\bullet$ are free, its connecting homomorphism is the direct sum of the inclusions

$$B_p(C) \otimes_R \mathrm{H}_q(D_\bullet) \longrightarrow Z_p(C) \otimes_R \mathrm{H}_q(D_\bullet).$$

To identify the connecting map, lift a boundary to C_{p+1} and tensor it with a cycle in D_q. The total differential is its inclusion into $Z_p(C)$ tensored with that cycle. The cokernel of the displayed map is $\mathrm{H}_p(C_\bullet) \otimes_R \mathrm{H}_q(D_\bullet)$, and its kernel is $\operatorname{Tor}_1^R(\mathrm{H}_p(C_\bullet), \mathrm{H}_q(D_\bullet))$. The long exact sequence consequently gives the asserted short exact sequence, with the kernel shifted up one total degree. Its first map is the map induced by tensoring cycles, hence the cross product after cellular comparison. These constructions commute with chain maps, proving the stated naturality.

For splitting, choose sections of $C_n \to B_{n-1}(C)$. They give a chain-complex isomorphism from $C_\bullet$ to a direct sum of two-term complexes

$$P^{(p)} : \quad B_p(C) \hookrightarrow Z_p(C),$$

placed in degrees $p + 1, p$. Make the same choices for $D_\bullet$, giving complexes $Q^{(q)}$. Each $P^{(p)}$ is a free resolution of $\mathrm{H}_p(C_\bullet)$, placed in those degrees. The tensor product decomposes into the pairs $P^{(p)} \otimes_R Q^{(q)}$.

After shifting a pair down to degree zero, write its resolutions as $P_1 \to P_0 \to M$ and $Q_1 \to Q_0 \to N$. The augmentation $Q_\bullet \to N$, with N concentrated in degree zero, gives a chain map

$$\varepsilon : P_\bullet \otimes_R Q_\bullet \longrightarrow P_\bullet \otimes_R N.$$

To prove that this is a quasi-isomorphism, use the map of short exact sequences

$$\begin{array}{ccccccccc} 0 & \longrightarrow & P_0 \otimes_R Q_\bullet & \longrightarrow & P_\bullet \otimes_R Q_\bullet & \longrightarrow & (P_1 \otimes_R Q_\bullet)[1] & \longrightarrow & 0 \\ & & \downarrow & & \downarrow{\scriptstyle \varepsilon} & & \downarrow & & \\ 0 & \longrightarrow & P_0 \otimes_R N & \longrightarrow & P_\bullet \otimes_R N & \longrightarrow & (P_1 \otimes_R N)[1] & \longrightarrow & 0. \end{array}$$

The leftmost and rightmost vertical maps are quasi-isomorphisms: tensoring the resolution of N with either free module P_0 or P_1 preserves its exactness. The long exact homology sequences and the five lemma show that the middle vertical map is a quasi-isomorphism as well. The target has differential $P_1 \otimes_R N \to P_0 \otimes_R N$, whose cokernel is $M \otimes_R N$ and whose kernel is $\operatorname{Tor}_1^R(M, N)$. This computes the homology of the pair in degrees zero and one and proves its vanishing in other degrees. The shift signs are removed by changing a chosen component by its negative when necessary.

Restoring degrees, these contributions lie in degrees $p + q$ and $p + q + 1$. On each pair, the degree-zero contribution is the image of the canonical tensor map, while the degree-one contribution maps isomorphically to its Tor quotient. These identifications follow from the exact sequence already proved, applied to that pair. Taking the inverses of the latter isomorphisms and summing over the pairs gives a section of the Tor quotient in each total degree. Thus the chosen direct-sum decomposition splits the asserted exact sequence. The sections of chain modules were not canonical. The theorem gives a natural injection and quotient map without specifying sections of those quotient maps. □

For a supplementary treatment of this algebraic proof, see May [May, n.d., Section 2.2, Theorem 2.4, pp. 6–7]. We will use the natural cross product when computing induced maps and the existence of a splitting when determining abstract groups.

C.3 Computations for Familiar Products

The product formulas become particularly transparent when one factor is a circle or when the factors have only one torsion group. These two cases exhibit both the canonical isomorphism and the additional Tor contribution.

C.3.1 A Circle Factor and Products of Surfaces

Choose an orientation of the circle, hence a generator of $\mathrm{H}_1(\mathrm{S}^1; \mathbb{Z})$. Since its integral homology is free in degrees zero and one, Theorem C.2.1 gives, for every finite cell complex X,

$$\mathrm{H}_n(X \times \mathrm{S}^1; \mathbb{Z}) \cong \mathrm{H}_n(X; \mathbb{Z}) \oplus \mathrm{H}_{n-1}(X; \mathbb{Z}).$$

The first inclusion is induced by $X \to X \times \mathrm{S}^1$ at a chosen point, and the second is cross product with the oriented circle class. These descriptions specify the isomorphism, including its dependence on the circle orientation.

Let Σ_g denote the closed oriented surface of genus $g \geq 0$. Its cell structure from Appendix B gives integral homology $\mathbb{Z}$, $\mathbb{Z}^{2g}$, $\mathbb{Z}$ in degrees zero, one, two. Therefore

$$\mathrm{H}_n(\Sigma_g \times \mathrm{S}^1; \mathbb{Z}) \cong \begin{cases} \mathbb{Z}, & n = 0, 3, \\ \mathbb{Z}^{2g+1}, & n = 1, 2, \\ 0, & \text{otherwise.} \end{cases}$$

All these products have Euler characteristic zero, but their first homology ranks distinguish different genera. Euler characteristic and the individual homology groups answer different questions.

For a field k, define the homology polynomial of a finite cell complex by

$$P_X(t;k) = \sum_{n \geq 0} \dim_k \mathrm{H}_n(X;k)\, t^n.$$

The field Künneth theorem says that $P_{X\times Y}(t;k) = P_X(t;k)P_Y(t;k)$, since the coefficient in degree n is the sum over $p+q=n$ of the products of dimensions. In particular,

$$\begin{aligned} P_{\Sigma_g\times\Sigma_h}(t;k) &= (1+2gt+t^2)(1+2ht+t^2) \\ &= 1+2(g+h)t+(4gh+2)t^2+2(g+h)t^3+t^4. \end{aligned}$$

Integral homology is free here and has the same ranks. Evaluating the polynomial at $t=-1$ recovers the product formula for Euler characteristic from Appendix B.

C.3.2 The Product of Two Real Projective Planes

The calculation in Lecture 8 gives

$$\mathrm{H}_n(\mathbb{R}P^2;\mathbb{Z}) \cong \begin{cases} \mathbb{Z}, & n=0, \\ \mathbb{Z}/2\mathbb{Z}, & n=1, \\ 0, & \text{otherwise.} \end{cases}$$

The integral Künneth sequence now gives

$$\mathrm{H}_n(\mathbb{R}P^2\times\mathbb{R}P^2;\mathbb{Z}) \cong \begin{cases} \mathbb{Z}, & n=0, \\ (\mathbb{Z}/2\mathbb{Z})^2, & n=1, \\ \mathbb{Z}/2\mathbb{Z}, & n=2,3, \\ 0, & \text{otherwise.} \end{cases}$$

The degree-one terms come from the two factors. The degree-two term is their tensor product. The degree-three term is their first Tor group, which cannot be obtained by tensoring the integral homology groups and adding degrees alone. In each positive degree of this example at most one of the tensor and Tor terms is nonzero, so no splitting choice is required to identify the answer.

There is a direct cellular interpretation of the degree-three class. Let a,u be the oriented one- and two-cells of the first factor, with $\partial u = 2a$, and let b,v be those of the second, with $\partial v = 2b$. The cross-product boundary formula gives

$$\partial(u\otimes b) = 2(a\otimes b), \qquad \partial(a\otimes v) = -2(a\otimes b).$$

Consequently $u\otimes b + a\otimes v$ is a cycle. The boundary of the product four-cell is twice this cycle. The integral kernel in degree three is generated by it, since the boundary row is $(2\ -2)$. Its homology class therefore has order two. This identifies geometrically how the Tor class comes from cancellation of two boundaries.

With coefficients in $\mathbb{F}_2$, the cellular differential of $\mathbb{R}P^2$ is zero in every degree. Thus

$$P_{\mathbb{R}P^2\times\mathbb{R}P^2}(t;\mathbb{F}_2) = (1+t+t^2)^2 = 1+2t+3t^2+2t^3+t^4.$$

In particular, mod-two homology is nonzero in degree four, although integral homology in degree four vanishes. This conclusion follows directly from cellular chains and field Künneth; it does not require a separate change-of-coefficients theorem.

C.4 Maps of Products and Orientation Signs

A canonical product map contains more information than a list of groups. Its naturality computes the action of a product map, and the interchange of two factors records their dimensions through a sign.

C.4.1 Product Maps on Top Homology

For $r, s \geq 1$, choose oriented generators of the top integral homology of S^r and S^s. Their cross product is the resulting generator of $\mathrm{H}_{r+s}(\mathrm{S}^r \times \mathrm{S}^s; \mathbb{Z}) \cong \mathbb{Z}$. If $f : \mathrm{S}^r \to \mathrm{S}^r$ and $g : \mathrm{S}^s \to \mathrm{S}^s$ have degrees a, b, naturality gives the commutative square

$$\begin{array}{ccc} \mathrm{H}_r(\mathrm{S}^r;\mathbb{Z}) \otimes \mathrm{H}_s(\mathrm{S}^s;\mathbb{Z}) & \xrightarrow[\cong]{\boxtimes} & \mathrm{H}_{r+s}(\mathrm{S}^r \times \mathrm{S}^s;\mathbb{Z}) \\ \downarrow{\scriptstyle a\otimes b} & & \downarrow{\scriptstyle (f\times g)_*} \\ \mathrm{H}_r(\mathrm{S}^r;\mathbb{Z}) \otimes \mathrm{H}_s(\mathrm{S}^s;\mathbb{Z}) & \xrightarrow[\cong]{\boxtimes} & \mathrm{H}_{r+s}(\mathrm{S}^r \times \mathrm{S}^s;\mathbb{Z}). \end{array}$$

The left vertical map is multiplication by ab, so the right vertical map is also multiplication by ab. The top-dimensional degree of the product map is the product of the degrees. For example, the product of a degree-two map of S^2 and a degree-three map of S^3 acts by multiplication by six on degree-five integral homology. It is not a homotopy equivalence, because that induced map is not an isomorphism.

C.4.2 Interchanging the Factors

Let $\tau : X \times Y \to Y \times X$ interchange the factors. The graded symmetry map

$$T : \mathrm{H}_*(X;R) \otimes_R \mathrm{H}_*(Y;R) \longrightarrow \mathrm{H}_*(Y;R) \otimes_R \mathrm{H}_*(X;R)$$

is the tensor-factor interchange multiplied by $(-1)^{pq}$ on bidegree (p, q). The square

$$\begin{array}{ccc} \mathrm{H}_*(X;R) \otimes_R \mathrm{H}_*(Y;R) & \xrightarrow{\boxtimes} & \mathrm{H}_*(X \times Y;R) \\ \downarrow{\scriptstyle T} & & \downarrow{\scriptstyle \tau_*} \\ \mathrm{H}_*(Y;R) \otimes_R \mathrm{H}_*(X;R) & \xrightarrow[\boxtimes]{} & \mathrm{H}_*(Y \times X;R) \end{array}$$

commutes. Indeed, interchanging an oriented p-cell and an oriented q-cell requires moving each of p directions past each of q directions, giving pq sign changes. The same sign is obtained by interchanging the two families of steps in the singular shuffle triangulation.

For $X = Y = \mathrm{S}^r$, the interchange acts on top integral homology by $(-1)^{r^2} = (-1)^r$. It reverses the product orientation when r is odd and preserves it when r is even. In particular, the coordinate interchange on $\mathrm{S}^1 \times \mathrm{S}^1$ reverses orientation, whereas the interchange on $\mathrm{S}^2 \times \mathrm{S}^2$ preserves it. These signs are part of the graded tensor product convention, not an additional choice made during the computation.

Brouwer Fixed Points and Invariance of Dimension

Algebraic Topology I | Fall 2026 | October 19, 2026

D.1 Retractions and Fixed Points

The homology of a sphere gives an obstruction to extending its identity map over a ball. Brouwer's fixed point theorem follows by turning a fixed-point-free map into just such an extension. A second application uses relative homology to recover the dimension of a Euclidean neighborhood from its topology.

We use ordinary singular homology, with integer coefficients throughout this lecture. For $n \geq 1$, let $\mathrm{D}^n = \{x \in \mathbb{R}^n : \|x\| \leq 1\}$ and $\mathrm{S}^{n-1} = \partial \mathrm{D}^n$. We write $B(p, r)$ for the open Euclidean ball of center p and radius $r > 0$. In dimension zero, $\mathrm{D}^0 = \mathbb{R}^0$ is a point. The sphere computation from the earlier lectures is

$$\widetilde{\mathrm{H}}_q(\mathrm{S}^d; \mathbb{Z}) \cong \begin{cases} \mathbb{Z}, & q = d, \\ 0, & q \neq d, \end{cases} \qquad d \geq 0.$$

Reduced homology is important here: the case S^0 belongs to the same argument as all the higher-dimensional spheres.

D.1.1 The Obstruction to a Retraction

A retraction would split the inclusion of the boundary at the level of spaces, and hence also after applying homology. The resulting splitting is incompatible with the homology of a ball.

Definition D.1.1. Recall that, for an inclusion $i : A \hookrightarrow X$, a retraction is a continuous map $r : X \to A$ such that $r \circ i = \mathrm{id}_A$. In this situation, A is called a retract of X.

A retraction is not required to come with a homotopy from $i \circ r$ to id_X. Thus proving that no retraction exists is stronger than proving that a specified deformation retraction does not exist.

Proposition D.1.2. For every $n \geq 1$, there is no retraction $\mathrm{D}^n \to \mathrm{S}^{n-1}$.

Proof. Suppose that r were a retraction of the boundary inclusion i. Applying reduced singular homology would give the factorization

$$\widetilde{\mathrm{H}}_{n-1}(\mathrm{S}^{n-1}; \mathbb{Z}) \xrightarrow{i_*} \widetilde{\mathrm{H}}_{n-1}(\mathrm{D}^n; \mathbb{Z}) \xrightarrow{r_*} \widetilde{\mathrm{H}}_{n-1}(\mathrm{S}^{n-1}; \mathbb{Z}), \qquad \mathrm{id}.$$

The middle group is zero because D^n is contractible, whereas both outer groups are infinite cyclic. Their identity cannot factor through the zero group. This includes $n = 1$, since $\widetilde{\mathrm{H}}_0(\mathrm{S}^0; \mathbb{Z}) \cong \mathbb{Z}$. □

The same factorization explains the extension obstruction of Proposition 8.1.8 of Lecture 8: a sphere map extending over a ball induces zero on its top reduced homology with integer coefficients. This formulation also covers S^0. On $\widetilde{\mathrm{H}}_0(\mathrm{S}^0; \mathbb{Z})$, the identity has degree 1, interchange of the two points has degree -1, and either constant map has degree 0.

D.1.2 The Ray Construction and Brouwer's Theorem

Suppose that a self-map of the ball has no fixed point. From its value at x, draw the ray through x and continue to the boundary. To obtain a retraction, one must verify that the endpoint varies continuously, including when x already lies on the boundary.

Theorem D.1.3. For every $n \geq 0$, every continuous map $f : \mathrm{D}^n \to \mathrm{D}^n$ has a fixed point.

Proof. For $n = 0$, the domain is a point. Assume $n \geq 1$, and suppose that $f(x) \neq x$ for every $x \in \mathrm{D}^n$. Set $v(x) = x - f(x)$, a continuous nowhere-zero vector. The nonnegative solution of

$$\|x + sv(x)\|^2 = 1$$

is

$$s(x) = \frac{-\langle x, v(x)\rangle + \sqrt{\langle x, v(x)\rangle^2 + \|v(x)\|^2(1 - \|x\|^2)}}{\|v(x)\|^2}.$$

The expression under the square root is nonnegative, and the denominator is nonzero. Consequently s is a continuous function on the entire closed ball. Define $r : \mathrm{D}^n \to \mathrm{S}^{n-1}$ by $r(x) = x + s(x)v(x)$. The quadratic equation shows that $r(x)$ belongs to the sphere.

If $\|x\| = 1$, then

$$\langle x, v(x)\rangle = 1 - \langle x, f(x)\rangle > 0.$$

Indeed, equality in $\langle x, f(x)\rangle \leq \|f(x)\| \leq 1$ would imply $f(x) = x$. Thus the numerator defining $s(x)$ is zero on the boundary, and $r(x) = x$ there. We have constructed a continuous retraction, contradicting Proposition D.1.2. □

The proof uses a global formula only to verify the continuity of a geometric construction. It neither assumes differentiability nor chooses a boundary point separately at each interior point. For $n = 1$, it recovers the fixed point theorem for a closed interval.

Proposition D.1.4. If a space K is homeomorphic to D^n, every continuous self-map of K has a fixed point. In particular, this holds for the cube $[-1, 1]^n$.

Proof. For a homeomorphism $h : \mathrm{D}^n \to K$ and a continuous map $g : K \to K$, the conjugate $h^{-1} \circ g \circ h$ has a fixed point. Its image under h is a fixed point of g.

For the cube, an explicit homeomorphism sends 0 to 0 and sends each nonzero $x \in \mathrm{D}^n$ to

$$h(x) = \frac{\|x\|_2}{\|x\|_\infty} x.$$

Its inverse sends a nonzero $y \in [-1, 1]^n$ to $(\|y\|_\infty / \|y\|_2)y$. These formulas are inverse on every ray. The inequalities between the two norms prove continuity at the origin, and continuity elsewhere follows directly from the formulas. The case $n = 0$ is again immediate. □

The closedness and boundedness of the ball matter to this statement. Translation on $\mathbb{R}$ has no fixed point, and the self-map $x \mapsto (x + 1)/2$ of the open interval $(-1, 1)$ has no fixed point. The theorem also asserts existence rather than uniqueness: the identity fixes every point.

D.2 Local Homology and Euclidean Dimension

Every Euclidean space is contractible, so its absolute homology cannot record its dimension. Relative homology of the complement of a point retains the missing information. Excision makes this invariant depend only on a neighborhood of the point.

D.2.1 Removing One Point

We first formulate locality without choosing coordinates. The subspace in the relative group is always written explicitly.

Definition D.2.1. For a space X and a point $p \in X$, its local singular homology at p, with integer coefficients, is the family of groups $\mathrm{H}_q(X, X - \{p\}; \mathbb{Z})$, indexed by $q \geq 0$.

Lemma D.2.2. Let X be a Hausdorff space, let $p \in X$, and let $V \subseteq X$ be an open neighborhood of p. Inclusion induces an isomorphism

$$\mathrm{H}_q(V, V - \{p\}; \mathbb{Z}) \xrightarrow{\cong} \mathrm{H}_q(X, X - \{p\}; \mathbb{Z})$$

for every $q \geq 0$.

Proof. Apply excision to the subspace $X - \{p\} \hookrightarrow X$, removing $Z = X - V$. Since V is open, Z is closed. Since $p \notin Z$ and $\{p\}$ is closed in the Hausdorff space X, its closure is contained in the open subspace $X - \{p\}$. The excision hypothesis therefore holds. The remaining pair is the inclusion $V - \{p\} \hookrightarrow V$. □

Proposition D.2.3. For $n \geq 0$, every open set $U \subseteq \mathbb{R}^n$ and every $p \in U$ satisfy

$$\mathrm{H}_q(U, U - \{p\}; \mathbb{Z}) \cong \begin{cases} \mathbb{Z}, & q = n, \\ 0, & q \neq n. \end{cases}$$

Proof. By Lemma D.2.2, inclusion identifies these groups with the groups for $\mathbb{R}^n - \{p\} \hookrightarrow \mathbb{R}^n$. Translation reduces the calculation to $p = 0$.

Assume first that $n \geq 1$. The punctured space deformation retracts onto S^{n-1}: on a nonzero vector x, the deformation at time t is $((1-t) + t/\|x\|)x$. Since $\mathbb{R}^n$ is contractible, the reduced long exact sequence of the pair gives, for $q \geq 1$, the isomorphism

$$\mathrm{H}_q(\mathbb{R}^n, \mathbb{R}^n - \{0\}; \mathbb{Z}) \xrightarrow[\cong]{\partial} \widetilde{\mathrm{H}}_{q-1}(\mathbb{R}^n - \{0\}; \mathbb{Z}) \xrightarrow{\cong} \widetilde{\mathrm{H}}_{q-1}(\mathrm{S}^{n-1}; \mathbb{Z}).$$

The sphere computation gives the asserted answer in positive degrees. In degree zero, $\mathrm{H}_0(\mathbb{R}^n - \{0\}; \mathbb{Z}) \to \mathrm{H}_0(\mathbb{R}^n; \mathbb{Z})$ is surjective because the punctured space is nonempty. Exactness gives zero relative homology in degree zero. This argument allows the punctured line to be disconnected.

When $n = 0$, the nonempty open set U is a point and $U - \{p\}$ is empty. The relative groups are the absolute homology groups of a point, so they are $\mathbb{Z}$ in degree zero and zero otherwise. □

For a point on a line, the nonzero relative class records the two local sides of the point. In higher dimensions, the boundary homomorphism records a small surrounding sphere. These are the same relative construction in different degrees.

D.2.2 Invariance of Dimension

A homeomorphism induces a homeomorphism of the complements of corresponding points. It consequently preserves the degree in which local homology is nonzero.

Theorem D.2.4. Let $m, n \geq 0$. If nonempty open subsets $U \subseteq \mathbb{R}^m$ and $V \subseteq \mathbb{R}^n$ are homeomorphic, then $m = n$.

Proof. Choose a homeomorphism $h : U \to V$ and a point $p \in U$. Its restrictions form a commutative square of inclusions

$$\begin{array}{ccc} U - \{p\} & \hookrightarrow & U \\ {\scriptstyle h|_{U-\{p\}}}\downarrow{\scriptstyle \cong} & & {\scriptstyle \cong}\downarrow{\scriptstyle h} \\ V - \{h(p)\} & \hookrightarrow & V. \end{array}$$

The induced homomorphisms on relative singular homology with coefficients in $\mathbb{Z}$ are isomorphisms. By Proposition D.2.3, the source is nonzero only in degree m, while the target is nonzero only in degree n. Hence these degrees agree. □

In particular, $\mathbb{R}^m$ and $\mathbb{R}^n$ are not homeomorphic when $m \neq n$, despite having identical absolute homology. The nonemptiness assumption is necessary because the empty sets in different Euclidean spaces are homeomorphic.

Definition D.2.5. A topological n-manifold without boundary is a Hausdorff, second-countable space in which each point has an open neighborhood homeomorphic to an open subset of $\mathbb{R}^n$. Such a homeomorphism is called a coordinate chart.

Corollary D.2.6. If M is a topological n-manifold without boundary, then at every point $p \in M$ its integer local homology is $\mathbb{Z}$ in degree n and zero in other degrees. A nonempty space cannot be both an m-manifold and an n-manifold without boundary for distinct m and n.

Proof. Choose a chart neighborhood of p. Lemma D.2.2 identifies the local homology in M with the local homology in this neighborhood. The chart and Proposition D.2.3 compute it. If both dimensions were possible, the same family of local groups would have two different unique nonzero degrees. □

This conclusion concerns homeomorphisms between open sets. To decide whether a continuous injection from one n-dimensional open set into another has open image, an additional argument is needed. Appendix E supplies it using degree with coefficients in $\mathbb{F}_2$.

For the no-retraction and fixed point arguments, see Allen Hatcher, *Algebraic Topology*, Corollary 2.15, pp. 114–115. For local homology and invariance of dimension, see Theorem 2.26, p. 126. The author's Chapter 2 is available at https://pi.math.cornell.edu/~hatcher/AT/ATch2.pdf.

Invariance of Domain

Algebraic Topology I | Fall 2026 | October 19, 2026

E.1 Invariance of Domain and Boundary Degree

Invariance of dimension compares spaces already known to be homeomorphic. Invariance of domain starts with less information: a continuous injection between open subsets of the same Euclidean space. The theorem asserts that its image is open and that its inverse on the image is continuous.

Theorem E.1.1. Let $n \geq 0$, let $U \subseteq \mathbb{R}^n$ be open, and let $f : U \to \mathbb{R}^n$ be continuous and injective. Then $f(U)$ is open, and $f : U \to f(U)$ is a homeomorphism.

We will prove that the image of a small closed ball contains a neighborhood of the image of its center. The boundary of that ball determines a sphere map of nonzero degree. The main homological step is the assertion that an odd sphere map has nonzero degree with coefficients in $\mathbb{F}_2 = \mathbb{Z}/2\mathbb{Z}$. We first give the deduction of the theorem from that assertion, and then prove the assertion directly with singular chains.

We write $\|x\|$ for the Euclidean norm and $B(y, r)$ for the open Euclidean ball of center y and radius $r > 0$. Throughout this lecture, all singular chains and homology groups have coefficients in $\mathbb{F}_2$, indicated explicitly. The use of these coefficients makes orientation signs unnecessary. For $d \geq 0$, a continuous map $g : \mathrm{S}^d \to \mathrm{S}^d$ is called *odd* if $g(-x) = -g(x)$ for every $x \in \mathrm{S}^d$.

E.1.1 Degree and Extension over a Ball

The sphere has a unique nonzero reduced homology class over $\mathbb{F}_2$ in its top degree. We use the action on this class to measure whether a map can extend across a ball.

Definition E.1.2. For $d \geq 0$ and a continuous map $g : \mathrm{S}^d \to \mathrm{S}^d$, its degree with coefficients in $\mathbb{F}_2$, denoted $\deg_2(g)$, is the scalar by which g_* acts on the one-dimensional vector space $\widetilde{\mathrm{H}}_d(\mathrm{S}^d; \mathbb{F}_2)$.

Homotopic maps have the same degree. Composition multiplies degrees, and the identity has degree 1. These statements follow from homotopy invariance and functoriality of ordinary singular homology with coefficients in $\mathbb{F}_2$. They apply to S^0 because reduced homology is used.

Lemma E.1.3. If $g : \mathrm{S}^d \to \mathrm{S}^d$ extends to a continuous map $G : \mathrm{D}^{d+1} \to \mathrm{S}^d$, then $\deg_2(g) = 0$.

Proof. The factorization through D^{d+1} gives the commutative diagram

$$\begin{array}{ccc} \widetilde{\mathrm{H}}_d(\mathrm{S}^d; \mathbb{F}_2) & \longrightarrow & \widetilde{\mathrm{H}}_d(\mathrm{D}^{d+1}; \mathbb{F}_2) \\ & \searrow^{g_*} & \downarrow{\scriptstyle G_*} \\ & & \widetilde{\mathrm{H}}_d(\mathrm{S}^d; \mathbb{F}_2). \end{array}$$

The upper-right group is zero by contractibility. Thus g_* is zero. □

Lemma E.1.4. For every $d \geq 0$, every continuous odd map $g : \mathrm{S}^d \to \mathrm{S}^d$ satisfies $\deg_2(g) = 1$.

The proof of Lemma E.1.4 occupies the next section. The following argument shows its geometric role before introducing the chains needed for its proof.

E.1.2 An Injective Ball Map Contains a Neighborhood

For a continuous injection on a closed ball, distinct source points give nonzero difference vectors in the target. Normalizing such vectors produces sphere maps without any assumption about differentiability or local linearity.

Proposition E.1.5. Let $n \geq 1$ and let $F : \mathrm{D}^n \to \mathbb{R}^n$ be continuous and injective. Then $F(\mathrm{D}^n)$ contains an open neighborhood of $F(0)$.

Proof. For $x \in \mathrm{S}^{n-1}$ and $t \in [0,1]$, define

$$K(x,t) = \frac{F(x) - F(-tx)}{\|F(x) - F(-tx)\|}.$$

The source points x and $-tx$ are distinct, because $\|x\| = 1$ and $t \geq 0$. Injectivity therefore makes the denominator nonzero everywhere. Hence $K : \mathrm{S}^{n-1} \times [0,1] \to \mathrm{S}^{n-1}$ is a continuous homotopy. Its initial map is

$$\beta_{F(0)}(x) = \frac{F(x) - F(0)}{\|F(x) - F(0)\|},$$

and its final map is odd. Lemma E.1.4 and homotopy invariance give $\deg_2(\beta_{F(0)}) = 1$.

The function $x \mapsto \|F(x) - F(0)\|$ is positive and continuous on the compact sphere. Thus its minimum δ is positive. For each $y \in B(F(0), \delta)$, define

$$\beta_y : \mathrm{S}^{n-1} \longrightarrow \mathrm{S}^{n-1}, \qquad \beta_y(x) = \frac{F(x) - y}{\|F(x) - y\|}.$$

For every $s \in [0,1]$, the point $y_s = F(0) + s(y - F(0))$ satisfies

$$\|F(x) - y_s\| \geq \delta - \|y - F(0)\| > 0 \qquad (x \in \mathrm{S}^{n-1}).$$

Consequently the maps β_{y_s} give a homotopy from $\beta_{F(0)}$ to β_y, and $\deg_2(\beta_y) = 1$.

If y did not belong to $F(\mathrm{D}^n)$, the same formula would define a continuous map $\mathrm{D}^n \to \mathrm{S}^{n-1}$ extending β_y. Lemma E.1.3 would give degree zero, a contradiction. We conclude that $B(F(0), \delta) \subseteq F(\mathrm{D}^n)$. □

Compactness of the boundary of a chosen closed ball supplies a uniform positive distance from the image of the center.

Proof. Assume first that $n \geq 1$. For each $p \in U$, openness supplies $r > 0$ such that the closed ball $\overline{B}(p,r)$ is contained in U. The map $F : \mathrm{D}^n \to \mathbb{R}^n$ defined by $F(x) = f(p + rx)$ is continuous and injective. Proposition E.1.5 shows that $f(U)$ contains a neighborhood of $f(p)$. Since this holds for every p, its image is open.

If $W \subseteq U$ is open, it is also open in $\mathbb{R}^n$. Applying the same argument to $f|_W$ shows that $f(W)$ is open. Thus f is an open continuous bijection onto its image, and its inverse is continuous. For $n = 0$, an open subset of the one-point space $\mathbb{R}^0$ is either empty or a point, and the assertion follows directly. □

E.2 The Degree of an Odd Map

We now prove Lemma E.1.4. The proof uses explicit chains for successive coordinate spheres. A chain for one hemisphere has boundary equal to the sum of an equatorial hemisphere chain and its antipodal translate. Vanishing of the intermediate homology groups of the ambient sphere then compares the images of these chains under an odd map.

E.2.1 Antipodal Hemisphere Chains

Fix $d \geq 1$ and use coordinates in $\mathbb{R}^{d+1}$ with unit vectors $v_0, \dots, v_d$. For $0 \leq i \leq d$, let S^i be the unit sphere in the span of $v_0, \dots, v_i$. These are nested subspaces of S^d. Write $a : \mathrm{S}^d \to \mathrm{S}^d$ for the antipodal map, which preserves every S^i.

On singular chains with coefficients in $\mathbb{F}_2$, define the operator $T = 1 + a_*$. It commutes with the boundary, and $T^2 = 0$, since $a_*^2 = 1$ and $2 = 0$ in $\mathbb{F}_2$. These are chain identities, rather than statements only about homology.

Lemma E.2.1. For each $0 \leq i \leq d$, there are chains $\alpha_i \in \mathrm{C}_i(\mathrm{S}^i; \mathbb{F}_2)$ such that α_0 is the point v_0, and, for $i \geq 1$,

$$\partial \alpha_i = T(\alpha_{i-1}).$$

For every i, the chain $\zeta_i = T(\alpha_i)$ represents the nonzero class in $\widetilde{\mathrm{H}}_i(\mathrm{S}^i; \mathbb{F}_2)$.

Proof. Set $\alpha_0 = v_0$ and $\zeta_0 = v_0 + (-v_0)$. This is an augmented zero-cycle representing the generator of $\widetilde{\mathrm{H}}_0(\mathrm{S}^0; \mathbb{F}_2)$.

Suppose ζ_{i-1} has been constructed. Cone each singular simplex in this chain to the pole v_i in the upper hemisphere of S^i, and let α_i be the sum of the resulting cones. More explicitly, for a singular $(i-1)$-simplex $\sigma : |\Delta^{i-1}| \to \mathrm{S}^{i-1}$, its cone sends the barycentric coordinates $(t_0, \dots, t_i)$ to the normalization of

$$t_0 v_i + (1 - t_0)\sigma\left(\frac{t_1}{1-t_0}, \dots, \frac{t_i}{1-t_0}\right)$$

when $t_0 < 1$, and sends the vertex $t_0 = 1$ to v_i. The two summands are orthogonal, so their sum never vanishes. As t_0 tends to 1, the formula tends uniformly to v_i, proving continuity at the apex.

Taking faces gives the cone boundary identity: the face opposite the apex is the original simplex, and the other faces are cones on its boundary faces. All signs are 1 over $\mathbb{F}_2$. Since ζ_{i-1} is a cycle, the latter faces cancel, yielding $\partial \alpha_i = \zeta_{i-1}$. When $i = 1$, the extra apex terms cancel because the augmentation of ζ_0 is zero. Define $\zeta_i = \alpha_i + a_*(\alpha_i)$. Its boundary is $\zeta_{i-1} + a_*(\zeta_{i-1}) = 0$, because $\zeta_{i-1} = T(\alpha_{i-1})$ is antipodally invariant as a chain.

To verify that ζ_i is the sphere generator, use the open cover

$$U_i = \mathrm{S}^i - \{-v_i\}, \qquad V_i = \mathrm{S}^i - \{v_i\}.$$

Both open sets are contractible, and their intersection deformation retracts onto S^{i-1}. The upper chain α_i lies in U_i, and its antipodal translate lies in V_i. The chain formula for the reduced Mayer–Vietoris connecting homomorphism therefore gives

$$\widetilde{\mathrm{H}}_i(\mathrm{S}^i; \mathbb{F}_2) \xrightarrow[\cong]{\partial_{\mathrm{MV}}} \widetilde{\mathrm{H}}_{i-1}(U_i \cap V_i; \mathbb{F}_2) \xrightarrow{\cong} \widetilde{\mathrm{H}}_{i-1}(\mathrm{S}^{i-1}; \mathbb{F}_2),$$

$$[\zeta_i] \longmapsto [\zeta_{i-1}] \longmapsto [\zeta_{i-1}].$$

By induction, the final class is nonzero. The connecting homomorphism is an isomorphism, including for $i = 1$ with reduced degree-zero homology, so $[\zeta_i]$ is nonzero as well. □

The equation $\partial \alpha_i = T(\alpha_{i-1})$ holds before taking homology.

E.2.2 Comparing the Chains under an Odd Map

The map g need not preserve any coordinate equator. We regard every hemisphere chain as a chain in the ambient sphere S^d, where all reduced homology below degree d vanishes.

Proof. For $d = 0$, an odd map permutes the two points of S^0. It therefore fixes the nonzero class in $\widetilde{\mathrm{H}}_0(\mathrm{S}^0; \mathbb{F}_2)$, so its degree is 1.

Assume $d \geq 1$ and use the chains from Lemma E.2.1. Oddness means $g \circ a = a \circ g$, so g_* commutes with T. Since S^d is path connected, there is a singular 1-chain $\beta_1 \in \mathrm{C}_1(\mathrm{S}^d; \mathbb{F}_2)$ such that

$$\partial\beta_1 = g_*(\alpha_0) + \alpha_0.$$

For example, take a path joining the points represented by the two zero-simplices.

Inductively, suppose that β_i has been constructed, and define the i-chain

$$\omega_i = g_*(\alpha_i) + \alpha_i + T(\beta_i).$$

For $i = 1$, its boundary is zero because

$$\partial\omega_1 = T(g_*(\alpha_0) + \alpha_0) + T(\partial\beta_1) = 0.$$

For $i > 1$, the induction will arrange $\partial\beta_i = \omega_{i-1}$. Using $T^2 = 0$ then gives

$$\partial\omega_i = T(g_*(\alpha_{i-1}) + \alpha_{i-1}) + T(g_*(\alpha_{i-1}) + \alpha_{i-1} + T(\beta_{i-1})) = 0.$$

Whenever $i < d$, the vanishing $\mathrm{H}_i(\mathrm{S}^d; \mathbb{F}_2) = 0$ permits us to choose β_{i+1} with $\partial\beta_{i+1} = \omega_i$. This constructs the chains through ω_d, which is a cycle. If $d = 1$, there are no intermediate fillings; the first boundary computation already proves that ω_d is a cycle.

Every homeomorphism of S^d induces an automorphism of its one-dimensional top homology over $\mathbb{F}_2$. The only such automorphism is the identity. In particular, a_* is the identity, and the operator $T_* = 1 + a_*$ on $\mathrm{H}_d(\mathrm{S}^d; \mathbb{F}_2)$ is zero. On the other hand, the chain identity $T^2 = 0$ gives

$$T(\omega_d) = g_*(\zeta_d) + \zeta_d.$$

Taking homology yields $g_*([\zeta_d]) + [\zeta_d] = 0$. Since $[\zeta_d]$ is the nonzero sphere class, g_* is the identity and $\deg_2(g) = 1$. □

This finishes the deferred proof used in Proposition E.1.5 and hence the proof of invariance of domain. The argument required only the sphere homology calculation, the Mayer–Vietoris connecting map, and functoriality on singular chains.

E.3 Consequences for Embeddings and Manifolds

The openness conclusion has consequences beyond subsets already known to be open. It also explains why the dimension equality in the theorem is essential.

E.3.1 Embeddings into Euclidean Space

A topological embedding is a continuous map that is a homeomorphism onto its image with the subspace topology. Invariance of domain says, in particular, that a continuous injection from a Euclidean open set to a Euclidean space of the same dimension is an embedding.

Corollary E.3.1. Let $m > n \geq 0$, and let $U \subseteq \mathbb{R}^m$ be nonempty and open. There is no continuous injection $U \to \mathbb{R}^n$.

Proof. If such a map existed, compose it with the coordinate inclusion $\mathbb{R}^n \hookrightarrow \mathbb{R}^m$. The resulting continuous injection would have open image by Theorem E.1.1. Its image lies in $\mathbb{R}^n \times \{0\}^{m-n}$, which has empty interior in $\mathbb{R}^m$: every ball about a point of this subspace contains a point with a nonzero last coordinate. A nonempty open subset cannot lie in this subspace. □

In the opposite direction, the coordinate inclusion $\mathbb{R}^n \hookrightarrow \mathbb{R}^{n+1}$ is an embedding with nonopen image. Thus the domain theorem requires equal dimensions, even for very simple maps.

E.3.2 Open Embeddings of Manifolds

The Euclidean result passes to manifolds by restricting to coordinate charts in both source and target. Here manifolds are Hausdorff and second countable, as defined in Appendix D.

Corollary E.3.2. Let M and N be topological n-manifolds without boundary. Every continuous injective map $f : M \to N$ is an open embedding. If, moreover, M is nonempty and compact and N is connected, then f is a homeomorphism onto N.

Proof. For $p \in M$, choose a target chart neighborhood V of $f(p)$. By continuity, a source chart can be restricted to an open neighborhood $W \subseteq f^{-1}(V)$ of p. In the two charts, $f|_W$ is a continuous injection between Euclidean open sets of dimension n. Theorem E.1.1 shows that $f(W)$ is open in V, hence in N.

The same construction can be made inside any prescribed open subset of M. Its image is therefore a union of open sets $f(W)$, so f is open. Since it is also continuous and injective, it is a homeomorphism onto its image.

If M is compact, then $f(M)$ is compact and hence closed in the Hausdorff space N. It is also open and nonempty. Connectedness gives $f(M) = N$. □

For example, a continuous injection from a compact connected manifold without boundary to itself is surjective. For $n \geq 1$, a nonempty compact n-manifold without boundary cannot admit a continuous injection into $\mathbb{R}^n$: the corollary would make it homeomorphic to the noncompact connected space $\mathbb{R}^n$.

The absence of boundary is necessary. The inclusion $[0, 1] \hookrightarrow \mathbb{R}$ is a continuous injection whose image is not open. At an endpoint, a neighborhood is a half-interval, not an open interval, so the source-chart step in the proof does not apply.

See Allen Hatcher, *Algebraic Topology*, Theorem 2B.3, pp. 172–173, for invariance of domain, and Proposition 2B.6, pp. 174–176, for the odd-degree theorem. The singular-chain proof above supplies the mod-2 statement needed here without using cohomology or a separation theorem. The author's Chapter 2 is available at https://pi.math.cornell.edu/~hatcher/AT/ATch2.pdf.

References

[EM45] S. Eilenberg and S. Mac Lane, *General theory of natural equivalences*, Transactions of the American Mathematical Society **58** (1945), 231–294.

[ML98] Saunders Mac Lane, *Categories for the Working Mathematician*, second edition, Graduate Texts in Mathematics, vol. 5, Springer-Verlag, New York, 1998.

[Stacks] The Stacks Project Authors, *The Stacks Project*, 2026. https://stacks.math.columbia.edu.

[CN08] Carles Casacuberta and Amnon Neeman, *Brown Representability Does Not Come for Free*, arXiv:0807.1872v1, 2008. arXiv. See Section 1 for the failure of small morphism sets.

[Wei94] Charles A. Weibel, *An Introduction to Homological Algebra*, Cambridge Studies in Advanced Mathematics, vol. 38, Cambridge University Press, 1994.

[May, n.d.] J. Peter May, *Notes on Tor and Ext*. University of Chicago lecture notes.

[Kerodon] *Kerodon*. https://kerodon.net.

[May99] J. P. May, *A Concise Course in Algebraic Topology*, University of Chicago Press, 1999. Further background for Lectures 5 and 7: Chapters 6, 10, and 13–15; Chapter 10, §7, p. 80 treats the double mapping cylinder.

[ES45] S. Eilenberg and N. E. Steenrod, Axiomatic approach to homology theory, *Proceedings of the National Academy of Sciences of the United States of America* **31** (1945), no. 4, 117–120. DOI: 10.1073/pnas.31.4.117.

[Hat02] A. Hatcher, *Algebraic Topology*, Cambridge University Press, 2002. See also §2.2, pp. 149–150 for mapping tori.